

Multi-Location Program Repair Strategies Learned from Successful Experience

Shangwen Wang^{1,3}, Xiaoguang Mao^{1,3}, Nan Niu², Xin Yi^{1,3}, and Anbang Guo^{1,3}

¹National University of Defense Technology, Changsha, China

²University of Cincinnati, Cincinnati, USA

³Hunan Key Laboratory of Software Engineering for Complex Systems, Changsha, China
{wangshangwen13, xgmao, yixin09, guoanbang12}@nudt.edu.cn, niunn@ucmail.uc.edu

Abstract—Automated program repair (APR) has great potential to reduce the effort and time-consumption in software maintenance and becomes a hot topic in software engineering recently with many approaches being proposed. Multi-location program repair has always been a challenge in this field since its complexity in logic and structure. While some approaches do not claim to have the features for solving multi-location bugs, they generate correct patches for these defects in practice. In this paper, we first make an observation on multi-location bugs in Defects4J and divide them into two categories (i.e., *similar* and *relevant multi-location bugs*) based on the repair actions in their patches. We then summarize the situation of multi-location bugs in Defects4J fixed by current tools. We analyze the twenty-two patches generated by current tools and propose two feasible strategies for fixing multi-location bugs, illustrating them through two detailed case studies. At last, preliminary results prove the feasibility of our methods with the repair of two bugs that have never been fixed before. By learning from successful experience in the past, this paper points out possible ways ahead for multi-location program repair.

Keywords—automated program repair; multi-location bugs; case studies

I. INTRODUCTION

Over the years, researchers develop various *Automated Program Repair* (APR) techniques aiming at reducing the onerous burden of fixing bugs. Generally, these automated repair tools can be classified into two categories, i.e., search-based methodology (e.g., GenProg [1] and RSRepair [2]) and semantics-based methodology (e.g., S3 [3] and Angelix [4]). Search-based repair method (also known as generate-and-validate methodology) generate patch candidates by searching within a predefined fault space determined by Fault Location (FL) techniques and then validate these candidates against the provided test-suite. Semantics-based repair methodology, on the contrary, utilizes semantic information generated by symbolic execution and constraint solving to synthesize patches. These state-of-the-art APR tools make great achievements on single-edit program repair.

Multi-location program repair, which refers to fixing multi-location bugs whose human-written patches contain multiple chunks [4, 5], has become a challenge since the rise of APR due to its complexity. Some recent empirical studies have shown the importance of multi-location repair: The study by Sobreira et al. [5] shows that more than 60% of the bugs in Defects4J [6], a well-known dataset containing 395 real bugs collected from six open-source Java projects, need fixing in multiple points and Zhong and Su [7] draw the conclusion that

programmers make at least two repair actions in total to fix more than 70% of bugs. So far, only Angelix and S3 have been reported to possess special features designed for multi-location bugs by capturing the dependence among multiple program locations. However, as the authors stated in [3], semantic-based repair exclusively modifies expressions in conditions or on the right-hand side of assignments, leading to its not so satisfactory performance. Thus, it is an emerging trend for solving multi-location bugs.

While some tools do not claim that they have the abilities to fix multi-location bugs, they generate correct patches for these bugs when being evaluated. For example, the patch in Listing 1 modifies the types of two variables in different locations and it is a typical multi-location defect. Recently, a tool named ACS [8] has reported to fix this bug successfully. This phenomenon motivates our study. In this paper, we analyze why these patches are generated and how they fix the bugs, aiming to provide practical guidance for future research by learning the experience. We first conduct an empirical study on the multi-location bugs in Defects4J and classify them into two categories according to repair actions in their patches. We then investigate the statistics of multi-location bugs from Defects4J that are successfully fixed by the current tools. We analyze the twenty-two patches for multi-location bugs generated by current tools and propose two suggestions for solving this kind of bugs, illustrating them through two detailed case studies. Preliminary results show the practicalities of our suggestions with the repair of two out of eight multi-location bugs in Defects4J that have never been fixed before. Our methods are successful by micro-adjustment of our suggestions with existing tools.

II. EMPIRICAL OBSERVATION

Two patches for multi-location bugs in Defects4J, Lang#35 and Chart#5, are shown in Listings 1 and 2, respectively. We divide multi-location bugs in this dataset into two categories by analyzing the repair actions in the two instances.

In Listing 1, developer changes two statements at different places into a same oracle-throwing statement. Modifications at each edit point share similar actions in syntax and thus this bug is classified into syntax similar multi-location bugs category (*similar multi-location bugs* for short). In this category, the similar modification we talk about may spread over both a single line of statement (Time#3) and a chunk of codes (Chart#14). Such cases are abundant in Defects4J, such as Time#3 where several similar *if* conditional statements are added, Math#49 where an object instantiation is modified in many functions from a same class, Closure#4 where a conditio-

```

- type = Object.class;
3295 + throw new IllegalArgumentException();
.....
- return (T[]) new Object[] { null };
3574 + throw new IllegalArgumentException();

```

Listing 1. The patch of Lang#35

```

541 if (x == null) {
542     throw new IllegalArgumentException("Null 'x' argument.");
543 }
544 + if (this.allowDuplicateXValues) {
545     add(x);
546 + return null;
547 }
548
549 // if we get to here, we know that duplicate X values are not permitted
550 XYDataItem overwritten = null;
551 int index = indexOf(x);
552 - if (index == 0 && !this.allowDuplicateXValues) {
553 + if (index >= 0) {
554     XYDataItem existing = (XYDataItem) this.data.get(index);
555     if (
        overwritten = (XYDataItem) existing.clone();

```

Listing 2. The patch of Chart#5

nal expression is modified similarly at two places, etc. Among the 244 multi-location patches in our dataset, this category has 70 instances, occupying 28.69% of the total amount. A small part of these cases (23/70) share exactly the same operations at each edit point, such as Chart#14, adding the same conditional block in four edit points.

The fixing shown in Listing 2 is different. It involves the addition of an *if* conditional statement and corresponding operations from lines 544-547 and the modification of the content of an *if* conditional statement in line 552. These modifications are compact and have great logical correlation in the program structure and we name this type semantic relevant multi-location bugs (*relevant multi-location bugs* for short). The criterion is from Abstract Syntax Tree (AST) level: *if the node under one modification appears in other modified places in this patch or the modified places are sub-nodes of a common node*, then the bug belongs to this type. Another example is the patch of Mockito#2 where developer first uses *Method Definition Addition* repair action from lines 30-34¹ and then operates *Method Call Addition* in line 11. This type is more popular in our dataset, holding 67.62% (165/244) of the total amount. This result is consistent with our perception that modifications performed at multiple locations aiming at solving a bug should have logical correlations, in most cases.

A small proportion of the dataset which contains 9 patches like Time#2 and Mockito#11 shows differences from the above two conditions: modifications at different places in these patches are neither similar nor logical related. We have not proposed any method for this kind of situation in this paper due to its peculiarity.

Note that the classification is based on the features of modifications at different places and it can provide guidance for repairing. *Similar* type bugs have no logical correlation at each location and thus we may fix these places one by one, however, *relevant* type bugs possess logical correlations at each location and it may affect other places when operating in one place. That is why *relevant* type bugs are more difficult to repair and it is proved through the results which we will show in the next section: more *similar multi-location bugs* have been fixed than *relevant multi-location bugs*. Also note that when counting the number of each category, we use a *relevant first* strategy which means if a patch contains both similar and relevant edits, it belongs to the latter. For example, in Math#74, two similar loop chunks are added but there is another modification about the loops, making this patch belong to relevant bug. The reason for this strategy is that if both kinds of

¹ Due to space limitation, some code snippets are not shown. Please check them in <https://github.com/program-repair/defects4j-dissection>.

TABLE I. STATISTICAL RESULTS

Bug ID	T	PM	S	jGP	JK	N	A	ssF	J	HDR	SF
C5	R					✓					
C14	S						✓				
C19	S						✓				
CL115	R							✓			
L10	R									✓	
L27	R		✓								
L35	S						✓				
L41	R		✓								
L50	S		✓								
L60	S		✓								
M4	S						✓				
M22	S									✓	
M35	S		✓				✓				
M61	R						✓				
M71	S		✓								
M79	R		✓					✓			
M90	R						✓				
M93	R						✓				
M98	S		✓								
M99	S						✓				
Similar	11	-	5	-	-	-	6	-	-	1	-
Relevant	9	-	3	-	-	1	3	2	-	1	-
Total	20	-	8	-	-	1	9	2	-	2	-

Column “T” means the type of this bug and R refers to *relevant* type while S refers to *similar* type. “✓” denotes this bug is successfully fixed by the tool. “Similar” and “Relevant” denote the numbers of different types of bugs fixed by each tool and “Total” denotes the total number of bugs fixed by each tool. It is marked with “-” if the tool cannot fix any bug.

operations are needed for fixing a bug, then the difficulty degree is near to repairing a *relevant* bug.

III. SITUATION STATISTICS

In this section, we investigate how the contemporary APR tools would handle the multi-location bugs. We select ten tools which have been evaluated on this dataset: ProbabilisticModel (PM) [9], SimFix (S) [10], jGenProg (jGP) [11], jKali (JK) [11], Nopol (N) [12], ACS (A), ssFix (ssF) [13], JAID (J) [14], HDRRepair (HDR) [15], and SketchFix (SF) [16]. Note that we adopt the experimental results for jGenProg, jKali, and Nopol reported by Martinez et al. [11] and the results of other approaches come from the corresponding research papers. The results are illustrated in Table I where each tool is represented by its acronym.

Generally speaking, 22 valid patches are generated and 20 multi-location bugs are successfully fixed including 11 *similar* type and 9 *relevant* type, among which M35 and M79 are fixed by two tools. There are only five tools being able to fix these bugs (i.e., SimFix, Nopol, ACS, ssFix, and HDRRepair) among which SimFix and ACS repair the most bugs with 8 and 9, respectively. Nopol and ssFix repair 1 and 2 bugs respectively and they can only fix *relevant multi-location bugs* at this moment. HDRRepair fixes one bug for each type.

IV. LESSONS LEARNED AND SUGGESTIONS

In this section, we propose two suggestions learned from successful experience and illustrate them through two detailed case studies.

A. Case Study 1: Patch of Lang#35 Generated by ACS

We list the patch of Lang#35 generated by ACS in Fig. 1. Another modification chunk performed at line 3578 is the same as the code in the figure. ACS is especially designed for synthesizing conditional expressions containing two steps: variable selection and predicate selection. It uses a method named *Oracle-Throwing* to avoid the crash, thus, it can generate patch as shown. However, why it is able to generate two modification chunks still needs further explanation. Note that Lang#35 is a *similar multi-location bug* which means the two modification points have no correlation in program structure. Thus, there may be multiple test cases aiming at testing different places in the program and they all fail. We find that the test suite for this project contains two failing test cases and when executing, ACS uses a *fitness function* which enables it to continue fixing if the repair actions that have been performed reduce the number of failing test cases (ACS does not introduce this feature in its paper, we get this information after connecting with the authors). Previous studies such as GenProg and HDRRepair use fitness functions to guide the selection process of candidate patches while ACS exploits the deduction of failing test cases for solving buggy points one by one, bringing a new idea for *similar multi-location bugs*. The main challenge for applying this strategy is the precondition: the test suite must have enough failing test cases to expose the defects and thus we must strengthen the test suite. Test case purification [17], which means recovering the execution of omitted assertions, has the ability to generate more practical test cases and enhance the performance of test suite. EvoSuite is a commonly used tool for automated test suite generation and empirically, it can increase code coverage up to 63% [18]. Thus, if we first purify the test suite and add test cases, leading to an enhanced test suite, and then use this fitness function to repair, we may be able to solve more multi-location bugs. This strategy is suitable for our study subject, Defects4J, since all the projects in this benchmark are open source projects and the original test suites are manually created which means they may not cover all the entities in the code.

Suggestion 1: For *similar multi-location bugs*, use a suitable fitness function for guiding the repair process combined with strengthened test suite.

B. Case Study 2: Patch of Chart#5 Generated by Nopol

The patch of Chart#5 generated by Nopol is shown in Fig. 2. Unlike the human-written patch shown in Listing 2 modifying two code chunks, this patch only modifies a conditional statement to repair this bug. The modification point is at line 563, just under the buggy point. Nopol is a semantic-based program repair tool utilizing angelic values and a Satisfiability Modulo Theory (SMT) solver for synthesizing conditional expressions. The conditional expression it generates really avoids the error. The variable *overwritten* is defined with *null* in line 546 and its value can only be modified if the condition in line 548 is met. When the condition in line 548 is not met, *overwritten* keeps the value *null* and the program goes to the conditional branch in line 563 where the condition is not satisfied, either, after being modified. Then the

```
3300 //ACS's patch begin
3301 + if (element == null){throw new IllegalArgumentException();}
3302 //ACS's patch end
```

Fig. 1. Patch of Lang#35 generated by ACS

```
- // append the value to the list...
+ if (this.autoSort) {
+   if (overwritten!=null) {
+     this.data.add(-index - 1, new XYDataItem(x, y));
```

Fig. 2. Patch of Chart#5 generated by Nopol

program skips this conditional branch and goes to line 566 directly and thus the wrong expression in line 548 does not cause the error in line 564 which means the error is avoided. This strategy is to some extent like **fault tolerance** technique [19] and it is the same principle with ACS not producing human-written patch like Listing 1. We further study the reason for modifying line 563.

Recently, **Error Propagation Chain (EPC)**, which refers to a sequence of statements between program defect and program failure statement, is proposed by Guo et al. [20] to improve the efficiency of fault localization. We check their experiment results and find that line 563 is in this chain. That indicates a new direction for fixing *relevant multi-location bugs*: since modifications at each edit point possess correlation in logic and it is hard for current technologies to fix at each point, we can find out the closest intersection to the buggy points in the EPCs and utilize SMT solver to find a patch which avoids the error. It is possible to generate a patch as long as an intersection can be found no matter how many buggy points the program possesses. All we need to do is selecting out the top-k suspicious statements, calculating their EPCs, and searching for patches at the intersections preferentially. Note that two situations of *relevant multi-location bugs* have been introduced in Section II. If two AST nodes have dependency relations, one statement will appear in another's EPC; if two nodes are both under a common node, then there will be an intersection in their EPCs.

Suggestion 2: For *relevant multi-location bugs*, find out the intersection of several EPCs and search for modifications at that point.

V. EMERGING RESULTS

In this section, we introduce our preliminary experimental results. We first randomly selected four *similar multi-location bugs* whose test cases are not yet capable for exposing all the defects and used SimFix, the latest tool with the same fitness function as ACS contains, to perform the fixing using the strategy of Suggestion 1. We then randomly selected four *relevant multi-location bugs*, each of whose EPC is less than ten lines (due to the time limitation), added EPCs information into their bug locations as we introduced in Suggestion 2, and utilized Nopol for synthesizing patches. All the selected bugs have never been repaired and the results are shown in Table II.

We manually examine the generated patches and consider a patch correct if it is the same or semantically equivalent to human-written one. The results show that our strategies repair two bugs by applying our suggestions with two current tools.

TABLE II. PRELIMINARY RESULTS

Bug ID	SimFix + S1				Nopol + S2			
	M46	M49	L62	T3	M95	CL8	CL50	L22
Fixed?		✓				✓		

Test case purification focuses on recovering the execution of omitted assertions and thus is sometimes useless for strengthening the test suite. For example, in Math#49, the test case *OpenMapRealVector* fails for its first function invocation which leads to an *InvocationTargetException* and thus the following function invocations cannot be executed, being the reason for SimFix not fixing this bug. In our experiment, we added the test cases generated by Evosuite into test suite, successfully exposed the two defects, and at last fixed this bug. The generated patch is the same as the standard one provided by Defects4J. We performed the same operation to the other three *similar* type bugs but SimFix failed to generate patches for them for mainly two reasons. For Lang#62 and Time#3, the reason is SimFix finds for fix ingredients in the original projects but there is no similar code in the source files, indicating that we may combine source files with existing open source projects to enlarge the space for searching for fix ingredients in the future. While for Math#46, the reason is the code snippet is so large (10 lines) that it considers the donor with a return statement the same as human-written patch as not similar. However, after we adjusted the code snippet size to a finer-grained value (2 lines), it still neglected the snippet which contains fix ingredient, which indicates that the similarities for identifying donor code snippets need to be improved.

In Closure#8, the edit point in human-written patch is line 202 in class *CollapseVariableDeclarations*. We calculated the intersections of the EPCs of top-100 ranked suspicious statements and perform synthesis on these points. Finally, Nopol generated a patch which adds an if-conditional statement to avoid the error under the another class. We performed the same operations to the other three *relevant* type bugs but they failed because of not finding angelic value at the interpoints (Lang#22) and not synthesizing a patch (Math#95 and Closure#50), corresponding to two of five limitations (*No angelic value found* and *Timeout in SMT*) the authors discussed in their paper, which means the repair ability of Nopol needs to be improved.

By making some micro-adjustments of our suggestions with current tools, we fixed two multi-location bugs. The failed cases are due to the weaknesses of current tools according to our analysis, indicating the potential of our suggestions to fix more bugs when combined with more powerful tools.

VI. CONCLUSION

In this paper, we divided multi-location bugs in Defects4J into two categories according to the repair actions in their patches, summarized the situation of these bugs fixed by current tools, and learned the successful experience as well as put forward two suggestions for future research (one for each type). Guided by our suggestions, we successfully fixed two multi-location bugs in Defects4J which have never been repaired before. To our best knowledge, we are the first to propose strategies by analyzing patches generated by current tools, bringing new idea for APR techniques as well as pointing out possible ways for multi-location program repair.

ACKNOWLEDGEMENT

This work was supported by the National Natural Science Foundation of China under Grant 61672529.

REFERENCES

- [1] Weimer W, Nguyen T V, Le Goues C, et al. Automatically finding patches using genetic programming[C]//Proceedings of the 31st International Conference on Software Engineering. IEEE Computer Society, 2009: 364-374.
- [2] Qi Y, Mao X, Lei Y, et al. The strength of random search on automated program repair[C]//Proceedings of the 36th International Conference on Software Engineering. ACM, 2014: 254-265.
- [3] Le X B D, Chu D H, Lo D, et al. S3: syntax-and semantic-guided repair synthesis via programming by examples[C]//Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering. ACM, 2017: 593-604.
- [4] Mechtaev S, Yi J, Roychoudhury A. Angelix: Scalable multiline program patch synthesis via symbolic analysis[C]//Proceedings of the 38th international conference on software engineering. ACM, 2016.
- [5] Sobreira V, Durieux T, Madeiral F, et al. Dissection of a bug dataset: Anatomy of 395 patches from Defects4J[C]//2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER). IEEE, 2018: 130-140.
- [6] Just R, Jalali D, Ernst M D. Defects4J: A database of existing faults to enable controlled testing studies for Java programs[C]//Proceedings of the 2014 International Symposium on Software Testing and Analysis. ACM, 2014: 437-440.
- [7] Zhong H, Su Z. An empirical study on real bug fixes[C]//Proceedings of the 37th International Conference on Software Engineering-Volume 1. IEEE Press, 2015: 913-923.
- [8] Xiong Y, Wang J, Yan R, et al. Precise condition synthesis for program repair[C]//Proceedings of the 39th International Conference on Software Engineering. IEEE Press, 2017: 416-426.
- [9] Soto M, Le Goues C. Using a probabilistic model to predict bug fixes[C]//2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER). IEEE, 2018: 221-231.
- [10] Jiang J, Xiong Y, Zhang H, et al. Shaping Program Repair Space with Existing Patches and Similar Code[C]// The International Symposium on Software Testing and Analysis. 2018.
- [11] Martinez M, Durieux T, Sommerard R, et al. Automatic repair of real bugs in java: A large-scale experiment on the defects4j dataset[J]. Empirical Software Engineering, 2017, 22(4): 1936-1964.
- [12] DeMarco F, Xuan J, Le Berre D, et al. Automatic repair of buggy if conditions and missing preconditions with SMT[C]//Proceedings of the 6th International Workshop on Constraints in Software Testing, Verification, and Analysis. ACM, 2014: 30-39.
- [13] Xin Q, Reiss S P. Leveraging syntax-related code for automated program repair[C]//Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering. IEEE, 2017: 660-670.
- [14] Chen L, Pei Y, Furia C A. Contract-based program repair without the contracts[C]//Automated Software Engineering (ASE), 2017 32nd IEEE/ACM International Conference on. IEEE, 2017: 637-647.
- [15] Le X B D, Lo D, Goues C L. History Driven Program Repair[C]// IEEE, International Conference on Software Analysis, Evolution, and Reengineering. IEEE, 2016:213-224.
- [16] Hua J, Zhang M, Wang K, et al. Towards practical program repair with on-demand candidate generation[C]//Proceedings of the 40th International Conference on Software Engineering. ACM, 2018: 12-23.
- [17] Xuan J, Monperrus M. Test case purification for improving fault localization[C]//Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering. ACM, 2014:52-63.
- [18] Galeotti J P, Fraser G, Arcuri A. Improving search-based test suite generation with dynamic symbolic execution[C]//Software Reliability Engineering (ISSRE), 2013 IEEE 24th International Symposium on. IEEE, 2013: 360-369.
- [19] Castro M, Liskov B. Practical Byzantine fault tolerance[C]// Symposium on Operating Systems Design & Implementation. ACM, 1999:173-186.
- [20] Guo A, Mao X, et al. An Empirical Study on the Effect of Dynamic Slicing on Automated Program Repair Efficiency[C]// In: Proceedings of IEEE International Conference on Software Maintenance and Evolution. IEEE, 2018: 580-584.