Must Fault Localization for Program Repair

Bat-Chen Rothenberg and Orna Grumberg^(✉)

Technion - Israel Institute of Technology, Haifa, Israel
`{batg,orna}@cs.technion.ac.il`

Abstract. This work is concerned with fault localization for automated program repair.

We define a novel concept of a *must* location set. Intuitively, such a set includes at least one program location from every repair for a bug. Thus, it is impossible to fix the bug without changing at least one location from this set. A fault localization technique is considered a *must* algorithm if it returns a must location set for every buggy program and every bug in the program. We show that some traditional fault localization techniques are not must.

We observe that the notion of must fault localization depends on the chosen repair scheme, which identifies the changes that can be applied to program statements as part of a repair. We develop a new algorithm for fault localization and prove that it is *must* with respect to commonly used schemes in automated program repair.

We incorporate the new fault localization technique into an existing mutation-based program repair algorithm. We exploit it in order to prune the search space when a buggy mutated program has been generated. Our experiments show that must fault localization is able to significantly speed-up the repair process, without losing any of the potential repairs.

1 Introduction

Fault localization and automated program repair have long been combined. Traditionally, given a buggy program, fault localization suggests locations in the program that might be the cause of the bug. Repair then attempts to change those suspicious locations in order to eliminate the bug.

Bad fault localization may cause a miss of potential repairs, if it is too restrictive, or cause an extra work, if it is too permissive. Studies have shown that for test-based repair imprecise fault localizations happen very often in practice [27]. This identifies the need for fault localization that can narrow down the space of candidates while still promising not to lose potential causes for a bug.

In this work, we define the concept of a *must* location set. Intuitively, such a set includes at least one location from every repair for the bug. Thus, it *must* be

This research was partially supported by the Technion Hiroshi Fujiwara cyber security research center and the Israel cyber bureau and partially by the Israel Science Foundation.

used for repair. In other words, **it is impossible to fix the bug using only locations outside this set**. A fault localization technique is considered a *must* algorithm if it returns a must location set for every buggy program and every bug in the program.

To demonstrate the importance of the *must* notion, consider the program in Fig. 1 for computing the absolute value of a variable x . The program is buggy since the assertion in location 4 is violated when initially $x = -1$. Intuitively, a good repair would replace the condition ($x < -1$) in location 2 with condition $x \leq -1$. Our must fault localization, defined formally in the paper, will include location 2 in the must location set. In contrast, the fault localization techniques defined for instance in [14, 21] do not include 2 in their location sets: They are not must and may miss optional repairs.

Our first observation regarding must notions is that their definition should take into account the *repair scheme* under consideration. A repair scheme identifies the changes that can be applied to program statements as part of a repair. A scheme can allow, for instance, certain syntactic changes in a condition (e.g. replacing $<$ with $>$) or in the right-hand-side expression of an assignment (e.g. replacing $+$ by $-$). A particular location set can be a *must* set using one scheme, but non-*must* using another. We further discuss this observation when presenting our formal definition of a must fault localization.

The setting of our work is as follows. Our approach is formula-based rather than test-based. We handle simple C-programs, with specification given as assertions in the code. Similarly to bounded model checking tools (e.g. [8]), the program and the negated specification are translated to a set of constraints, whose conjunction forms the *program formula*. This formula is satisfiable if and only if the program violates an assertion, in which case a satisfying assignment (also called a *model*) is returned.

We focus on a simple repair scheme of syntactic changes, as described above. We assume that the user prefers repairs that are as close to the original program as possible and will want to get several repair suggestions. Thus, we return *all minimal repairs* (minimal in the number of changes applied to the program code).

Once the notion of must fault localization is defined, we develop a new algorithm for fault localization and prove that it is *must* with respect to syntactic mutation schemes. The input to the algorithm is a program formula φ and a model μ for φ , representing a buggy execution of the program. Our approach is based on a dynamic-slicing-like algorithm that computes dependencies.

For a variable v in φ , its slice F is computed based on dynamic dependencies among variables in φ , whose values influence the value of v in μ . Informally, F is a must location set that contains all assignment to the variables that v depends upon. Some assignment from F thus must be changed in order to eliminate the bug associated with μ .

We incorporated the new fault localization technique into an existing mutation-based program repair algorithm [38]. In [38], the repair scheme is based on a predefined set of mutations. Given a buggy program P , the goal of the algorithm is to return all minimal repairs for P . The algorithm goes through

iterations of generate-validate, where the generate part produces a mutated program of P and the validate part checks whether it is bounded-correct. The bottleneck of the algorithm is the size of the search space, consisting of all possible mutated programs of P . In [38], the search space has been pruned when the generated mutated program has been successfully validated. No pruning has been applied otherwise.

In this work, we exploit our novel *must* fault localization in order to prune the search space when a buggy mutated program P' has been generated (i.e. validation failed). In this case, we compute the *must* location set F of P' . We can now prune from the search space any mutated program whose F locations are identical to those of P' . This is because, by the property of *must* location set, it is guaranteed that the bug cannot be repaired without changing a location in F . Thus, a large set of buggy mutated programs is pruned, without the need for additional validation and without losing any minimally repaired program. It should be noted that the smaller F is, the larger the pruned set is. Our experimental results confirm the effectiveness of this pruning by showing significant speedups.

To summarize, the contributions of this work are:

1. We define a novel notion of *must* fault localization with respect to a repair scheme. We show that many of the formula-based techniques are not *must*.
2. We present a novel fault localization technique and prove that it is *must* for the scheme of syntactic mutations. Our technique also has other advantages, such as low-complexity and incrementality.
3. We show how our new fault localization technique can be incorporated into an existing mutation-based program repair algorithm for pruning its search space. The technique is applied iteratively, whenever a generated mutated program is found to be incorrect.
4. We implemented the algorithm of repair with fault localization as part of the open source tool AllRepair. Our experimental results show that fault-localization is able to significantly speed-up the repair process, without losing any of the potential repairs.

2 Motivating Example

Figure 1 presents a simple program for computing the absolute value of a variable x . The result is computed in the variable `abs`, and the specification states, using an assertion on line 4, that in the end `abs` should always be non-negative. Unfortunately, the program has a bug. The true branch of the if is intended to flip the sign of x whenever x is negative, but it accidentally misses the case where x is -1 . As a result, if x is -1 , the wrong branch of the if is taken, and the assertion is reached with `abs` = -1 , which causes a violation.

```

procedure absValue(x)
1: abs := x
2: if x < -1 then
3:   abs := -x
4: assert (abs >= 0)

```

Fig. 1. A buggy program

Clearly, it is desirable that line number 2 be returned when running fault localization on this bug, as a human written repair is likely to change the condition on this line from $x < -1$ to $x \leq -1$ or $x < 0$. But, as we will show next, some of the existing formula-based fault localization techniques do not include this line in their result.

The error trace representing the bug for input $I = \{x \leftarrow -1\}$ is $\pi = \langle 1, 2, 4 \rangle$ (this is the sequence of program locations visited when executing the program on I). The MAX-SAT-based fault localization technique of [21] and the error-invariant-based technique of [14] use a formula called the *extended trace formula* in order to find faulty statements along the error trace. The extended trace formula for the bug in question is

$$\underbrace{(x = -1)}_{\text{Input}} \wedge \underbrace{(abs = x) \wedge (x \geq -1)}_{\text{Computation}} \wedge \underbrace{(abs \geq 0)}_{\text{Assertion}}$$

This formula encodes three things: a) that the input remains I , b) that the computation is as the trace dictates, and, c) that the assertion holds at the end. Therefore, the formula is unsatisfiable. Both [21] and [14] intuitively look for explanations of its unsatisfiability, and therefore decide that the statement $(x \geq -1)$ on line 2 is irrelevant; The formula remains unsatisfiable even if the constraint $(x \geq -1)$ is removed.

Even the method of [6], which suggests a flow-sensitive encoding of the extended trace formula, with the goal of including all statements affecting control-flow decisions that are relevant to the bug, classifies the statement on line 2 as irrelevant. This is because the error trace does not include any location from the body of the branch that was taken (in our case it is the `else` branch, which is empty), in which case the flow-sensitive formula remains identical to the traditional formula.

The dynamic slicing method of [2, 23] also fails to include line 2 in its result. This method computes the set of statements influencing the evaluation of the assertion along the trace, using data and control dependency relations. A statement st_1 is data dependent on st_2 iff st_1 uses a variable x , and st_2 is the last to assign a value to x along the trace. In our example, the assertion on line 4 is data dependent only on the statement in line 1, which in itself is not data dependent on any other statement. A statement st_1 is control dependent on a conditional statement st_2 iff st_1 is inside the body of either branch of st_2 . None of the statements along our error trace is control dependent on another statement. The slice, which is the set of lines returned, is computed using the transitive closure of these relations. Thus, for our example, only line 1 is part of the slice.

In this example, we have seen how many different fault localization techniques fail to include a statement that is relevant, i.e., where a modification could be made for the bug to be fixed. In contrast, the set of locations returned by our technique for this example is $\{1, 2\}$. The fact that our technique includes line 2 is not a coincidence: We show that, intuitively, whenever a repair can be made by making changes to a single line, this line *must* be included in the result.

<pre> proc. foo(x, w) 1: t := 0 2: y := x - 3 3: z := x + 3 4: if (w > 3) then 5: t := z + w 6: assert (t < x) 7: y := y + 10 8: assert (y > z) </pre>	<pre> proc. simFoo(x, w) t := 0 y := x - 3 z := x + 3 g := w > 3 if (g) then t := z + w assert (t < x) y := y + 10 assert (y > z) </pre>	<pre> proc. SSAFoo(x, w) t0 := 0 y0 := x0 - 3 z0 := x0 + 3 g0 := w0 > 3 t1 := z0 + w0 assert (g0 → t1 < x0) y1 := y0 + 10 t2 := g0 ? t1 : t0 y2 := g0 ? y1 : y0 assert (y2 > z0) </pre>	$\varphi_{foo} = \{$ $t_0 = 0,$ $y_0 = x_0 - 3,$ $z_0 = x_0 + 3,$ $g_0 = w_0 > 3,$ $t_1 = z_0 + w_0,$ $y_1 = y_0 + 10,$ $t_2 = ite(g_0, t_1, t_0),$ $y_2 = ite(g_0, y_1, y_0),$ $\neg(y_2 > z_0) \vee \neg(g_0 \rightarrow t_1 < x_0)$ $\}$
--	--	---	---

Fig. 2. Example of the translation process of a simple program

In general, whenever a repair can be made by making changes to a set of lines, at least one of them must be included in the result.

3 Preliminaries

3.1 Programs and Error Traces

For our purposes, a *program* is a sequential program composed of standard statements: assignments, conditionals, loops and function calls, all with their standard semantics. Each statement is located at a certain *location* (or *line*) l_i , and all statements are defined over the set of program variables X .

In addition to the standard statements, a program may also contain *assume* statements of the form **assume**(*bexpr*), and *assert* statements of the form **assert**(*bexpr*). In both cases *bexpr* is a boolean expression over X . If an assume or an assert statement is located in l_i , execution of the program stops whenever location l_i is reached in a state where *bexpr* is evaluated to false. In the case of an assertion, this early termination has the special name *assertion violation*, and it is an indication that an error has occurred.

A program P has a *bug on input* I if an assertion violation occurs during the execution of P on I . Otherwise, the program is *correct for* I .¹ Whenever P has a bug on I , this bug is associated with an *error trace*, which is the sequence of statements visited during the execution of P on I .

3.2 From Programs to Program Formulas

In this section we explain how a program is translated into a set of constraints, whose conjunction constitutes the program formula. In addition to constraints representing assignments and conditionals, such a formula includes constraints representing assumptions and a constraint representing the negated conjunction of all assertions. Thus, a satisfying assignment (a *model*) of the program formula

¹ Alternatively, one could assume to know the desired output of the program for I and define a bug on I as a case where the program outputs the wrong value for I .

represents an execution of the program that satisfies all assumption but violates at least one assertion. Such an execution is a *counterexample*.

The translation, following [8], goes through four stages. We refer to the example in Fig. 2 to demonstrate certain steps.

1. Simplification: Complex constructs of the language are replaced with equivalent simpler ones. Also, branch conditions are replaced with fresh boolean variables. In the example, the `if` condition (`w > 3`) is assigned to a fresh boolean variable `g`. Branching is then done based on the value of `g`, instead of (`w > 3`).
2. Unwinding: The body of each loop and each function is inlined *wb* times. The set of executions of the new program is called the *wb*-executions of *P*.
3. Conversion to SSA: The program is converted to static single assignment (SSA) form, which means that each variable in the new program is assigned at most once. This is done by replacing all variables with indexed variables, and increasing the index of a variable whenever it appears on the left-hand-side of an assignment. In the example, the first assignment to `t` is replaced by an assignment to `t0` and the second, by an assignment to `t1`. Since `t` is assigned inside a conditional statement and is used after the statement, the if-then-else assignment `t2 := g0?t1:t0` is inserted in order to determine which copy of `t` should be used after the conditional statement. These special if-then-else assignments are called *Φ*-assignments. In the example, there is also a *Φ*-assignment for `y` (`y2=g0?y1:y0`).

Note that, assertions are also expressed by means of indexed variables. The specific indices in the assertion indicate the location in the execution in which the assertion is checked. In addition, if an assumption or an assertion is located within an `if` statement with branch condition *g*, then it is implied by *g* if it is within the `then` part of the `if` and is implied by $\neg g$, if it is within the `else` part. In the example, `assert (t < x)` is encoded by $(g_0 \rightarrow t_1 < x_0)$.

4. Conversion to SMT constraints: Once the program is in SSA form, conversion to SMT is straightforward: An assignment `x:=e` is converted to the constraint $x = e$; A *Φ*-assignment `x:= b?x1:x2` is converted to the constraint $(x = ite(b, x_1, x_2))$, which is an abbreviation of $((b \wedge x = x_1) \vee (\neg b \wedge x = x_2))$; An assume statement `assume(bexpr)` is converted to the constraint `bexpr`, and an assert statement `assert(bexpr)` is converted to the constraint $\neg \text{bexpr}$ (since a model of the SMT formula should correspond to an assertion violation).

If the program includes several assertions, then they are converted to one constraint, representing the negation of their conjunction. In the example, the two assertions are converted to the following constraint:

$$\neg(y2 > z0) \vee \neg(g0 \rightarrow t1 < x0).$$

We say that a constraint *encodes* the statement it came from and we partition constraints into three sets, S_{assign} , S_{phi} and S_{demand} , based on what they encode. S_{assign} contains constraints encoding assignments, including those originated from assigning a fresh boolean variable with a branching condition;

S_{phi} - encoding Φ -assignments; and S_{demand} - encoding demands from assert and assume statements. In particular, it encodes the negated conjunction of all assertions.

The triple $(S_{assign}, S_{phi}, S_{demand})$ is called a *program constraint set*. The program constraint set we get from a program P when using wb as an unwinding bound is denoted CS_P^{wb} . The *program formula* φ_P^{wb} , is the conjunction of all constraints in all three sets of CS_P^{wb} :

$$\varphi_P^{wb} = \left(\bigwedge_{s \in S_{assign}} s \right) \wedge \left(\bigwedge_{s \in S_{phi}} s \right) \wedge \left(\bigwedge_{s \in S_{demand}} s \right).$$

Theorem 1 ([9]). *A program P is wb -violation free iff the formula φ_P^{wb} is unsatisfiable.*

For simplicity of notation, in the rest of the paper we omit the superscript wb .

Since the program formula is the result of translating an SSA program, the formula is defined over indexed variables. Further, each constraint in S_{assign} corresponds to the single variable, which is assigned in the statement encoded by the constraint.

4 Must Fault Localization

In this section, we precisely define when a location should be considered relevant for a bug. This definition is motivated by a repair perspective, taking into account which changes can be made to statements in order to repair a bug.

In order to define the changes allowed, we use repair schemes. A *repair scheme* $\mathcal{S}$ is a function from statements to sets of statements. An $\mathcal{S}$ -*patch* for a program P is a set of pairs of location and statement $\{(l_1, st_1^r), \dots, (l_k, st_k^r)\}$, for which the following holds: for all $1 \leq i \leq k$, let st_i be the statement in location l_i in P , then $st_i^r \in \mathcal{S}(st_i)$. The patch is said to be *defined over* the set of locations $\{l_1, \dots, l_k\}$. Applying an $\mathcal{S}$ -*patch* τ to a program P means replacing for every location l_i in τ , the statement st_i with st_i^r . This results in an $\mathcal{S}$ -*patched* program of P . The set of all $\mathcal{S}$ -*patched* programs created from a program P is the $\mathcal{S}$ -*search space* of P .

Let P be a program with a bug on input I , and $\mathcal{S}$ be a repair scheme. An $\mathcal{S}$ -*repair* for I is an $\mathcal{S}$ -patched program that is correct for I . An $\mathcal{S}$ -*repairable set* is a set of locations F such that there exists an $\mathcal{S}$ -repair defined over F . An $\mathcal{S}$ -repairable set is *minimal* if removing any location from it makes it no longer an $\mathcal{S}$ -repairable set. A location is $\mathcal{S}$ -*relevant* if it is a part of a minimal $\mathcal{S}$ -repairable set.²

In this paper, we focus on two repair schemes that are frequently used for automated program repair: the arbitrary scheme ($\mathcal{S}_{arb}$) and the mutation scheme ($\mathcal{S}_{mut}$). Both schemes only manipulate program expressions, but the

² We sometimes omit $\mathcal{S}$ from notations where $\mathcal{S}$ is clear from context.

mutation scheme is more restrictive than the arbitrary scheme: $\mathcal{S}_{arb}(st)$ is the set of all options to replace the expression of st ³ with an arbitrary expression, while $\mathcal{S}_{mut}(st)$ only contains statements where the expression in st is mutated according to a set of simple syntactic rules. The rules we consider are replacing a $+$ operator with a $-$ operator, and vice versa, replacing a $<$ operator with a $>$ operator, and vice versa, and increasing or decreasing a numerical constant by 1.⁴

Example 1. In this example we demonstrate how different repair schemes define different sets of relevant locations. Consider again the foo program from Fig. 2. This program has a bug on input $I = x \leftarrow 0, w \leftarrow 0$. The error trace associated with the bug is $\langle 1, 2, 3, 4, 8 \rangle$ (the assertion on line 8 is violated).

The location set $\{3, 4\}$ is a minimal $\mathcal{S}_{mut}$ -repairable set: It is an $\mathcal{S}_{mut}$ -repairable set because applying the $\mathcal{S}_{mut}$ -patch $\{(3, z := x - 3), (4, w < 3)\}$, results in an $\mathcal{S}_{mut}$ -patched program that is correct for I . This set is also minimal, because none of the $\mathcal{S}_{mut}$ -patches defined over $\{3\}$ or $\{4\}$ alone is an $\mathcal{S}_{mut}$ -repair for I : Each one of the $\mathcal{S}_{mut}$ -patches $\{(3, z := x - 3)\}$, $\{(3, z := x + 4)\}$, $\{(3, z := x + 2)\}$, $\{(4, w < 3)\}$, $\{(4, w > 4)\}$, $\{(4, w > 2)\}$ results in an assertion violation for I .

On the other hand, $\{3, 4\}$ is *not* a minimal $\mathcal{S}_{arb}$ -repairable set: For example, the $\mathcal{S}_{arb}$ -patch $\{(3, z := -6)\}$ is an $\mathcal{S}_{arb}$ -repair for I . Note that, the $\mathcal{S}_{arb}$ -patch only needs to repair the bug, and not the program. That is, it is sufficient that there is no assertion violation on the specific input I , even though an assertion could be violated in the $\mathcal{S}_{arb}$ -patched program on another input.

The set of all minimal $\mathcal{S}_{arb}$ -repairable sets is $\{\{2\}, \{3\}, \{4, 5\}\}$. Therefore, the set of $\mathcal{S}_{arb}$ -relevant statements is $\{2, 3, 4, 5\}$. The set of all minimal $\mathcal{S}_{mut}$ -repairable sets is $\{\{2, 3\}, \{3, 4\}\}$. Therefore, the set of $\mathcal{S}_{mut}$ -relevant statements is $\{2, 3, 4\}$.

Fault localization should focus the programmer's attention on locations that are relevant for the bug. But, returning the exact set of $\mathcal{S}$ -relevant locations, as defined above, can be computationally hard. In practice, what many fault localization algorithms return is a set of locations that *may* be relevant: The returned locations have a higher chance of being $\mathcal{S}$ -relevant than those who are not, but there is no guarantee that all returned locations are $\mathcal{S}$ -relevant, nor that all $\mathcal{S}$ -relevant locations are returned. We call such an algorithm *may fault localization*. In contrast, we define *must fault localization*, as follows:

Definition 1 ($\mathcal{S}$ -must location set). An $\mathcal{S}$ -must location set is a set of locations that contains at least one location from each minimal $\mathcal{S}$ -repairable set.⁵

³ If st is an assignment, its expression is its right-hand-side. If st is a conditional statement, its expression is its condition.

⁴ This simple definition of the mutation scheme is used only for simplicity of presentation. Our implementation supports a much richer set of mutation rules, as explained in Sect. 7.

⁵ This is, in fact, a hitting set of the set of all minimal $\mathcal{S}$ -repairable sets.

Definition 2 ($\mathcal{S}$ -must fault localization). *An $\mathcal{S}$ -must fault localization algorithm is an algorithm that for every program P and every buggy input I , returns an $\mathcal{S}$ -must location set.*

Note that, an $\mathcal{S}$ -must location set is not required to contain all $\mathcal{S}$ -relevant locations, but only one location from each minimal $\mathcal{S}$ -repairable set. Still, this is a powerful notion since it guarantees that no repair is possible without including at least one element from the set.

Also note, that the set of all locations visited by P during its execution on I is always an $\mathcal{S}$ -must location set. This is because any $\mathcal{S}$ -patch where none of these locations is included is definitely **not** an $\mathcal{S}$ -repair, since the same assertion will be violated along the same path. However, this set of locations may not be minimal. In the sequel, we aim at finding small $\mathcal{S}$ -must location sets.

Example 2. Continuing the previous example, the set $\{2, 3, 4\}$ is an $\mathcal{S}_{arb}$ -must location set, and also an $\mathcal{S}_{mut}$ -must location set. In contrast, the set $\{2, 3\}$ is only an $\mathcal{S}_{mut}$ -must location set, but not an $\mathcal{S}_{arb}$ -must location set, since it does not contain any location from the $\mathcal{S}_{arb}$ -minimal repairable set $\{4, 5\}$. The set $\{2\}$ is neither an $\mathcal{S}_{arb}$ -must location set nor an $\mathcal{S}_{mut}$ -must location set.

Example 3. Consider again the `absValue` procedure of Fig. 1. The set $\{2\}$ is an $\mathcal{S}_{mut}$ -minimal repairable set and an $\mathcal{S}_{arb}$ -minimal repairable set for the bug in question. Therefore, we can say that all algorithms that were shown in Sect. 2 not to include the location 2 in their result [2, 6, 14, 21, 23], are neither $\mathcal{S}_{arb}$ -must nor $\mathcal{S}_{mut}$ -must fault localization algorithms.

5 Fault Localization Using Program Formula Slicing

In this section we formally define the notion of slicing. Based on this, we present an algorithm for computing must fault localization for $\mathcal{S}_{arb}$ and $\mathcal{S}_{mut}$.

5.1 Program Formula Slicing

A central building block in our fault localization technique is *slicing*. But, we do not define slicing in terms of the program directly, but in terms of the program formula representing it, instead. The input to the slicing algorithm is a program formula φ , a model μ of it, and a variable v . Recall that φ is a conjunction of constraints from $\mathcal{S}_{assign}$, $\mathcal{S}_{phi}$ and $\mathcal{S}_{demand}$ (see Sect. 3.2). The goal of the slicing algorithm is to compute the *slice* of the variable v with respect to φ and μ . Intuitively, this slice includes the set of all constraints that influence the value v gets in μ .

Similar to traditional slicing, it is easy to define the slice as the reflexive-transitive closure of a dependency relation. But, unlike traditional slicing, which defines dependencies between statements, our dependency relation is between variables of the formula. These variables are indexed. Each originates from a variable of the underlying SSA program, where it was assigned at most once.

Fig. 3. Illustration of the static and dynamic dependency relations of the `foo` procedure

We refer to variables never assigned as *input variables*, and denote the set containing them by $InputVars$. A variable v that was assigned once is called a *computed variable*, and the (unique) constraint encoding the assignment to it is denoted $Assign(v)$. The set of all computed variables is denoted $ComputedVars$. We also denote by $vars(e)$ the set of variables that appear in a formula or expression e .

Definition 3 (Static Dependency). *The static dependency relation of a program formula φ is $SD_{\varphi} \subseteq vars(\varphi) \times vars(\varphi)$ s.t.*

$$SD_{\varphi} = \{(v_1, v_2) \mid \exists e \text{ s.t. } (v_1 = e) \in S_{assign}, v_2 \in vars(e)\} \cup \{(v, b), (v, v_1), (v, v_2) \mid (v = ite(b, v_1, v_2)) \in S_{phi}\}$$

The left-hand-side of Fig. 3 presents the graph for the static dependency relation of the `foo` procedure of Fig. 2. The nodes in the graph are (indexed) variables and there is an arrow from v_1 to v_2 iff $(v_1, v_2) \in SD_{\varphi}$.

Definition 4 (Dynamic Dependency). *The dynamic dependency relation of a program formula φ and a model μ of φ is $DD_{\varphi, \mu} \subseteq vars(\varphi) \times vars(\varphi)$ s.t.*

$$DD_{\varphi, \mu} = \{(v, v_1) \mid \exists b, v_2 \text{ s.t. } (v = ite(b, v_1, v_2)) \in S_{phi}, \mu[b] = true\} \cup \{(v, v_2) \mid \exists b, v_1 \text{ s.t. } (v = ite(b, v_1, v_2)) \in S_{phi}, \mu[b] = false\} \cup \{(v, b) \mid \exists v_1, v_2 \text{ s.t. } (v = ite(b, v_1, v_2)) \in S_{phi}\} \cup \{(v, v_1) \mid \exists e \text{ s.t. } (v = e) \in S_{assign}, v_1 \in vars(e)\}$$

Note that, dynamic dependency includes only dependencies that coincide with the specific model μ , which determines whether the `then` or the `else` direction

of the `if` is executed. Static dependency, on the other hand, takes both options into account. Thus, $DD_{\varphi,\mu} \subseteq SD_{\varphi}$ for every model μ .

The bold arrows on the right-hand-side of Fig. 3 represent the relation $DD_{\varphi,\mu}$ of the `foo` procedure, for any μ where $\mu[g_0] = \text{false}$.

Definition 5 (Influencing Variables). *Given a program formula φ , a model μ of it, and a computed variable v , the set of influencing variables of v with respect to φ and μ is:*

$$InfluenceVars_{\varphi,\mu}(v) = \{v' \mid (v, v') \in (DD_{\varphi,\mu})^*\}$$

The circled nodes on the right-hand-side of Fig. 3 represents the variables that belong to $InfluenceVars_{\varphi,\mu}(y_2)$.

Definition 6 (Program Formula Slice). *Given a program formula φ , a model μ of it, and a computed variable v , the program formula slice of v with respect to φ and μ is:*

$$Slice_{\varphi,\mu}(v) = \{Assign(v') \mid v' \in (InfluenceVars_{\varphi,\mu}(v) \cap ComputedVars)\}$$

Thus, intuitively, $Slice_{\varphi,\mu}(v)$ includes all constraints (in SSA form) encoding assignments that influence the value of v in μ . More precisely, when considering the conjunction of only the constraints of $Slice_{\varphi,\mu}(v)$, as long as the value of all input variables remains the same as in μ , the value of v will remain the same as well. This is formalized in the following theorem, whose proof can be found in the full version [39].

Theorem 2. *For every φ, μ and v , the following holds:*

$$\left[\bigwedge_{c \in Slice_{\varphi,\mu}(v)} c \wedge \bigwedge_{v_i \in InputVars} (v_i = \mu[v_i]) \right] \implies (v = \mu[v])$$

Continuing with our example of `foo` procedure,

$$Slice_{\varphi,\mu}(y_2) = \{ y_2 = ite(g_0, y_1, y_0), y_0 = x_0 - 3, g_0 = w_0 > 3 \}.$$

5.2 Computing the Program Formula Slice

The computation of the program formula slice is composed of two steps. In the first step, we build a graph based on the static dependency relation, SD_{φ} . In the second step, we compute the slice $Slice_{\varphi,\mu}(v)$ by computing the set of nodes reachable from v in this graph, using a customized reachability algorithm, which makes use of the model μ .

The graph built during the first step is called the *Static Dependency Graph* (*SDG*) of φ . Nodes of this graph are variables of φ and edges are the static dependencies of SD_{φ} . Edges are annotated using the function ψ , mapping every static dependency (v, v') to a boolean formula such that $(v, v') \in DD_{\varphi,\mu}$ iff

$\mu \models \psi[(v, v')]$. Specifically, for every constraint of the form $(v = \text{ite}(b, v_1, v_2))$ in S_{phi} , the edge (v, v_1) is annotated with b and the edge (v, v_2) is annotated with $\neg b$. All other edges of the graph are annotated with *true*. See the left-hand-side of Fig. 3. For simplicity all *true* annotations are omitted.

The algorithm for the second step is presented in Algorithm 1. This algorithm gets a program formula φ , its SDG, a model μ of φ , and a variable v , and computes $\text{Slice}_{\varphi, \mu}(v)$. First, the set $\text{InfluenceVars}_{\varphi, \mu}(v)$ is computed as the set of nodes reachable from v in SDG, except that the reachability algorithm traverses an edge (v, v') only if $\mu \models \psi[(v, v')]$. Thus, an edge (v, v') is traversed iff $(v, v') \in DD_{\varphi, \mu}$, which means that the set of reachable nodes computed this way is in fact $\text{InfluenceVars}_{\varphi, \mu}(v)$. Finally, the slice $\text{Slice}_{\varphi, \mu}(v)$ is the set of constraints encoding assignments to variables in $\text{InfluenceVars}_{\varphi, \mu}(v)$.

Algorithm 1. Compute The Program Formula Slice

Input: a program formula φ , its SDG, a model μ of φ and a variable v .
Output: $\text{Slice}_{\varphi, \mu}(v)$.

Procedure

ComputeSlice($\varphi, \text{SDG}, \mu, v$)

```

1:  $V := \emptyset$ 
2:  $\text{ModelBasedDFS}(\text{SDG}, v, \mu, V)$ 
3:  $\text{Slice} := \{\text{Assign}(v') \mid v' \in V\}$ 
4: return  $\text{Slice}$ 
```

Procedure

ModelBasedDFS(SDG, v, μ, V)

```

1:  $V := V \cup \{v\}$ 
2: for  $(v, w) \in E$  s.t.  $\mu \models \psi[(v, w)]$  do
3:   if  $w \notin V$  then
4:      $\text{ModelBasedDFS}(\text{SDG}, w, \mu, V)$ 
```

Algorithm 2. FOrmula-Slicing-Fault-Localization (FOSFL)

Input: A program formula φ of a program P , and a model μ of φ .
Output: A set of statements F of P .

Procedure *FOSFL*(φ, μ)

```

1:  $\text{SDG} := \text{ComputeDependencyGraph}(\varphi)$ 
2:  $\text{demandFormula} := \bigwedge_{c \in S_{\text{demand}}} c$ 
3:  $V := \text{ImportantVars}(\text{demandFormula}, \mu)$ 
4:  $S := \emptyset$ 
5: for  $v \in V$  do
6:    $S := S \cup \text{ComputeSlice}(\varphi, \text{SDG}, \mu, v)$ 
7:  $F := \emptyset$ 
8: for  $c \in S \cap S_{\text{assign}}$  do
9:    $F := F \cup \{\text{Origin}(c)\}$ 
10: return  $F$ 
```

5.3 The Fault Localization Algorithm

Our fault localization algorithm is presented in Algorithm 2. The input to this algorithm is a program formula φ of a program P , and a model μ of φ . The model μ represents a buggy execution of P on an input I , and the algorithm returns a set of locations, F , that is an $\mathcal{S}_{mut}$ -must location set.

As before, we assume to know the origin of constraints in φ , and use the sets S_{assign} , S_{phi} and S_{demand} . Furthermore, here we also assume that for every constraint $c \in S_{\text{assign}}$, we know exactly which program statement it came from. We call this statement the *origin* of c , and denote it by $\text{Origin}(c)$.

As a first step, the algorithm computes a set of variables V by calling the procedure *ImportantVars*. This procedure receives an SMT formula φ and a model μ of φ , and reduces μ to a partial model of φ . A *partial model* of φ w.r.t. μ is a partial mapping from variables of the formula to values, which is

consistent with μ and is sufficient to satisfy the formula. For example, for the formula $\varphi = (a = 0 \vee b = 0)$ and the model $\mu = \{a \mapsto 0, b \mapsto 1\}$, the valuation $\{a \mapsto 0\}$ is a partial model of φ . Procedure *ImportantVars* will return the set of variables that appear in the partial model ($\{a\}$ in our example). Details of this procedure are presented in the full version [39].

The formula passed to *ImportantVars* in our case is the conjunction of all demands in S_{demand} . Recall that the set S_{demand} contains constraints encoding all conditions that need to be met for an assertion violation to happen: Conditions from assumptions appear as is, while conditions from assertions are negated and disjuncted (See Fig. 2. The last constraint on the right-hand-side represents the disjunction of the negated assertions). Therefore, the set of variables V , returned by *ImportantVars*, is such that as long as their values in μ remain the same, this conjunction will still be satisfied, which means that an assertion violation will still occur.

To make sure that their values do *not* remain the same, we use slicing: The algorithm proceeds by computing the program formula slice for each of the variables in V using Algorithm 1. All slices are united into the combined set S . This set represents all constraints that if remain the same, then *all* the variables in V maintain their value. Thus, at least one element from S must be included in any repair.

Note that, by first applying *ImportantVars*, we reduce the number of variables whose value should be preserved in order to maintain the bug. The smaller this number, the smaller F is. We will explain the usefulness of a small F in Sect. 6.

Finally, we need to translate the constraints in S back to statements of P . Because of how the slicing algorithm works, constraints in S may belong to either S_{assign} or S_{phi} . If they belong to S_{phi} , we ignore them, because they encode the control-flow structure of the program, rather than a particular statement. Otherwise, we add the origin of the constraint, which is a statement of the program, to the set of returned locations, F . Note that, several different constraints may have the same origin, for example due to loop unwinding. In such a case, it is sufficient for one constraint encoding the statement st to be included in S , for st to be included in F . A proof for the following theorem can be found in the full version [39].

Theorem 3. *Algorithm FOSFL is an $\mathcal{S}_{arb}$ -must and also an $\mathcal{S}_{mut}$ -must fault localization algorithm.*

5.4 Incremental Fault Localization

It is often necessary to apply fault localization to several bugs in the same program, or even to several programs with different bugs. Therefore, it is desired that the fault localization algorithm be *incremental*, which means that the computation effort of each fault localization attempt should be proportional to the changes made from the previous attempt. In other words, we should avoid re-computation whenever possible, taking advantage of the fact that the program remains the same, or at least remains similar.

Algorithm FOSFL can be easily made incremental for the case of different bugs of the same program. In this case, several successive calls are made to the algorithm using the same program formula φ , but with different models of it. Since the static dependency relation SD_φ depends solely on the program formula, and not on the model, we can avoid re-computing the SDG for each call. Instead, we can compute the SDG once, upfront, and whenever FOSFL is called, simply skip the first line. We call the incremental version of FOSFL Incremental-Formula-Slicing-Fault-Localization (I-FOSFL).

Note that I-FOSFL is useful not only for fault localization of different bugs of the same program, but also whenever the SDG remains the same during successive fault localization calls. This is the case when considering different mutated programs P' of the same program P , since every change to P' replaces an expression e with an expression e' over the same variables. Thus, the SDG remains the same, since the static dependency relation, in fact, only depends on $\text{vars}(e)$, and not on e itself⁶.

6 Program Repair with Iterative Fault Localization

In [38], a mutation-based algorithm for program repair, named ALLREPAIR, was presented. This algorithm uses the mutation scheme in order to repair programs with respect to assertions in the code. Unlike fault localization, where the motivation is repairing a bug for a specific input, program repair aims at repairing the program for *all* inputs. To avoid confusion, we refer to a repair for all inputs as a *full repair*. In [38], the notion of a *full repair* is bounded: loops are unwound *wb* times, and a program is considered *fully repaired* if no assertion is violated along executions with at most *wb* unwindings. A program that is not fully repaired is said to be *buggy*. For the rest of this section, we refer to an $\mathcal{S}_{mut}$ -patch as a patch, and to an $\mathcal{S}_{mut}$ -patched program as a mutated program.

As its name implies, the goal of ALLREPAIR is to obtain all *minimal* fully repaired mutated programs, where minimality refers to the patch used in the program. It goes through an iterative generate-validate process. The generate phase chooses a mutated program from the search space, and the validate phase checks whether this program is fully repaired, by solving its program formula. The mutated program is fully repaired iff the formula is unsatisfiable.

The generate-validate process is realized using an interplay between a SAT solver and an SMT solver. The SAT solver is used for the generate stage. For every mutation M and line l , there is a boolean variable $B_M(l)$, which is true if and only if mutation M is applied to line l . A boolean formula is constructed and sent to the SAT solver, where each satisfying assignment corresponds to a program in the search space. The SMT solver is used for the validate stage. The program formula of the mutated program is solved to check if it is buggy

⁶ This is true for $\mathcal{S}_{mut}$ but not for $\mathcal{S}_{arb}$, since the latter allows to replace an expression e with an expression e' over different variables.

Fig. 4. Algorithm FL-ALLREPAIR: Mutation-based program repair with iterative fault localization. The notation $P^M \equiv_F P'$ means that P^M and P' agree on the content of all locations in F . The notation $P^M \subseteq P'$ means that the patch used for creating P' is a superset of the patch used for creating P^M .

or not. To achieve minimality, when a mutated program created using a patch τ is fully repaired, every mutated program created using a patch τ' , with $\tau \subseteq \tau'$, is blocked.

Example 4. Let P^M be a fully repaired mutated program obtained by applying the patch τ , consisting of mutating line l_1 using mutation M_1 and mutating line l_2 using mutation M_2 . Then blocking any superset of τ will be done by adding to the boolean formula representing the search space, the blocking clause $\neg(B_{M_1}(l_1) \wedge B_{M_2}(l_2))$, which means “either do not apply M_1 to l_1 or do not apply M_2 to l_2 ”. This clause blocks any mutated program with $\tau \subseteq \tau'$.

Blocking such programs prunes the search space, but only in a limited way. No pruning occurs when the mutated program is buggy.

In this paper, we extend the algorithm of [38] with a fault localization component. The goal of the new component is to prune the search space by identifying sets of mutated programs that are buggy, without inspecting each of the individual programs in the set.

Figure 4 shows the program repair algorithm with the addition of fault localization. In the new algorithm, called FL-ALLREPAIR, whenever a mutated program is found to be buggy during the validation step, its program formula is passed to the fault localization component along with the model obtained when solving the formula. The fault localization component returns a set of locations F , following the I-FOSFL algorithm. Since this set is guaranteed to be an $\mathcal{S}_{mut}$ -must location set, at least one of the locations in it should be changed for the bug to be fixed. Consequently, all mutated programs in which all locations from F remain unchanged are blocked from being explored in the future. As before, blocking is done by adding a blocking clause that disallows such programs.

Example 5. Let P^M be a buggy mutated program for which F consists of $\{l_1, l_2, l_3\}$, where l_1 was mutated with M_1 , l_2 was not mutated, and l_3 was mutated with M_3 . The blocking clause $\neg B_{M_1}(l_1) \vee \neg B_{Original}(l_2) \vee \neg B_{M_3}(l_3)$ will be added

to the boolean formula representing the search space of mutated programs. It restricts the search space to those mutated programs that either do not apply mutation M_1 to l_1 , or do mutate l_2 or do not apply M_3 to l_3 . This will prune from the search space all mutated programs which are identical to P^M on the locations in F . Note that smaller F will result in a larger set of pruned programs.

Proposition 1. *Algorithm FL-ALLREPAIR is sound and complete.*

7 Experimental Results

We have implemented our fault localization technique and its integration with mutated-based program repair in the tool ALLREPAIR, available at <https://github.com/batchenRothenberg/AllRepair>. In this section, we present experiments evaluating the contribution of the new fault localization component to the program repair algorithm. We refer to the algorithm of [38], without fault localization, as AllRepair, and to the algorithm presented in this paper as FL-AllRepair. Both algorithms search for minimal wb -violation free programs, and both are sound and complete. Thus, for every buggy program and every bound wb , both algorithms will eventually produce the same list of repairs.

The difference between the algorithms lies in the repair loop. In case a mutated program is found to be buggy, the AllRepair algorithm will only block the one program, while the FL-AllRepair algorithm might block a set of programs. Therefore, the number of repair iterations required to cover the search space can only decrease using the FL-AllRepair algorithm. On the other hand, the cost of each iteration with fault localization is strictly higher than without it. Our goal in this evaluation is to check if the use of fault localization pays off. That is, to check if repairs are produced faster using FL-AllRepair than using AllRepair.

Benchmarks. For our evaluation, we have used programs from two benchmarks: TCAS and Codeflaws. The TCAS benchmark is part of the Siemens suite [12], and is frequently used for program repair evaluation [5, 34, 38]. The TCAS program implements a traffic collision avoidance system for aircrafts, and consists of approximately 180 lines of code. We have used all 41 faulty versions of the benchmark in our experiments.

The Codeflaws benchmark [41] is also a well-known and widely used benchmark for program repair. Programs in this benchmark are taken from buggy user submissions to the programming contest site Codeforces⁷. In each program, a user tries to solve a programming problem published as part of a contest on the site. The programming problems are varied, and also the users have a diverse level of expertise. The benchmark also provides correct versions for all buggy versions, which are used to classify bug types by computing the syntactic difference. For our experiments we randomly chose 13 buggy versions classified with bug types that can be fixed using mutations. The size of the chosen programs ranges from 17 to 44 lines of code.

⁷ <http://codeforces.com/>.

Mutations. The mutations used in ALLREPAIR (and accordingly in FL-AllRepair) is a subset of the mutations used in [37]. We define two *mutation levels*, where level 1 contains only a subset of the mutations available in level 2. Thus, level 1 involves easier computation but may fail more often in finding repairs.

Table 1 shows the list of mutations used in each mutation level. For example, for the category of arithmetic operator replacement, in mutation level 1, the table specifies two sets: $\{+, -\}$ and $\{/, \%\}$. This means that a $+$ can be replaced by a $-$, and vice versa, and that the operators $/$, $\%$ can be replaced with each other. Constant manipulation mutations apply to a numeric constant and include increasing its value by 1 ($C \rightarrow C + 1$), decreasing it by 1 ($C \rightarrow C - 1$), setting it to 0 ($C \rightarrow 0$) and changing its sign ($C \rightarrow -C$).

Level 1	Level 2
$\{+, -\}, \{/, \%\}$	$\{+, -, *, \{/, \%\}$
$\{>, >=\},$ $\{<, <=\}$	$\{>, >=, <, <=\}, \{==, !=\}$
$\{ , \& \& \}$	
$\{>>, << \}, \{ \&, , ^ \}$	
	$C \rightarrow C + 1, C \rightarrow C - 1, C \rightarrow -C,$ $C \rightarrow 0$

Table 1. Partition of mutations to levels

Setting. All of our experiments were run on a Linux 64-bit Ubuntu 16.0.4 virtual machine with 1 CPU, 4 GB of RAM and 40 GB of storage, provided using the VMWARE vRA service⁸. For each of the buggy versions in our benchmarks we have experimented with both mutation levels 1 and 2. For the Codeflaws benchmarks we additionally experimented with different unwinding bounds: 2 (entering the loop once), 5, 8 and 10. This experiment is irrelevant to the TCAS benchmarks since the TCAS program does not contain loops or recursive calls. Overall we had 186 combinations of buggy programs, mutation levels and unwinding bounds. We refer to each such combination as an *input*. For each input, we run both the AllRepair and the FL-AllRepair algorithms with a timeout of 10 minutes and a mutation size limit of 2 (i.e., at most two mutations could be applied at once).

7.1 Results

In total, 131 different repairs were found during our experiments, for 60 different inputs (for several inputs there was more than one possible repair). In this count, we treat repairs fixing the same program in the same way as different, if they were produced using different mutation levels or unwinding bounds. This is because our evaluation is concerned with the time to find these repairs, and both the mutation level and the unwinding bound greatly influence this time.

Because the time to produce a repair sometimes varied in several orders of magnitude depending on the input, we have chosen to split repairs into three categories: fast, intermediate, and slow, and examine the time difference separately for each category. Splitting repairs to categories was done according to the time it took to find them using the AllRepair algorithm. If that time was

⁸ <https://www.vmware.com/il/products/vrealize-automation.html>.

Fig. 5. Time to find each repair using AllRepair (AR) and FL-AllRepair (FLAR). Each x value represents a single repair, and the corresponding y values represent the time, in seconds, it took to find that repair using both algorithms. Note that the graphs differ in the y axis scale.

under 5 seconds, the repair was considered fast. If it was over 4 minutes, it was considered slow, and otherwise it was considered intermediate.

Figure 5 shows a comparison of the time, in seconds, it took to find repairs in both algorithms. There are three graphs, according to our three categories. In all graphs, each x value represents a single repair, where the corresponding blue dot in the y axis represents the time it took to find that repair using AllRepair, and the red square represents the time using FL-AllRepair. So, whenever the blue dot is above the red square, FL-AllRepair was faster in finding that repair, and the y difference represents the time saved.

For the fast category (Fig. 5a), there is no clear advantage to FL-AllRepair. The majority of the repairs in this category are produced in less than a second using both algorithms. For the remaining repairs, there appears to be as many cases where FL-AllRepair is faster as when it is slower. But, in all cases where there is a time difference, in either direction, it is only of a few seconds.

For the intermediate category (Fig. 5b), the advantage of FL-AllRepair is starting to become clear. There are now only 4 repairs (out of 20) for which FL-AllRepair is slower. Also, on average, it is slower by 4 seconds, but faster by 10 seconds. Finally, for the slow category (Fig. 5c), there is an obvious advantage to FL-AllRepair. First, it is able to find 6 repairs *exclusively*, while AllRepair reaches a time-out. Also, for the remaining 27 repairs, FL-AllRepair is faster in all cases but one. The time difference is now also very significant: FL-AllRepair is faster by 1512 seconds (around 25 minutes) on average.

To sum up, the results show that in many cases our algorithm FL-AllRepair is able to save time in finding repairs. The savings are especially significant in

cases where it takes a long time to produce the repair using the original AllRepair algorithm, and these are the cases where time savings are most needed.

7.2 Comparison with Other Repair Methods

The TCAS benchmark was recently used also in [34], where ALLREPAIR's performance was compared to that of four other automated repair tools: ANGELIX [29], GENPROG [26], FORENSIC [5] and MAPLE [34]. ALLREPAIR was found to be faster by an order of magnitude than all of the compared tools, taking only 16.9 seconds to find a repair on average, where the other tools take 1540.7, 325.4, 360.1, and 155.3 seconds, respectively. Since in our experiments on TCAS FL-ALLREPAIR was faster than ALLREPAIR on average (and even when it was slower it was only by a few seconds), we conclude that FL-ALLREPAIR also compares favorably to these other tools.

In terms of reparability, the repair scheme used by ALLREPAIR (and FL-ALLREPAIR) is limited compared to the other tools: ALLREPAIR only uses mutations on expressions while ANGELIX, FORENSIC and MAPLE allow replacing an expression with a template (e.g., a linear combination of variables), which is then filled out to create a repair. GENPROG allows modifying a statement as well as deleting it or adding a statement after it. Therefore, the other tools are inherently capable of producing repairs in more cases than ALLREPAIR.

In the case of TCAS, the study showed that ALLREPAIR is able to find repairs for 18 versions (a result that we confirm in our experiments as well), while ANGELIX, GENPROG, FORENSIC and MAPLE found 32, 11, 23 and 26, respectively. But, what the study also showed, is that in repair methods that are based on tests, in many cases the repair found only adhered to the test-suite, but was not correct when inspected manually. When counting only correct repairs, ALLREPAIR finds repairs for 18 versions (all of ALLREPAIR's repairs are correct), while ANGELIX, GENPROG, FORENSIC and MAPLE find 9, 0, 15 and 26, respectively. Since FL-ALLREPAIR is able to find all repairs found by ALLREPAIR, the same results also apply to FL-ALLREPAIR.

8 Related Work

Dynamic slicing has been widely used for fault localization in the past [16, 36, 43, 45–47]. But, as we have seen, traditional notations of dynamic slicing [2, 23] are not must (with respect to neither of the presented schemes), and thus, the above techniques may fail to include relevant locations in their results.

Other approaches for fault localization include spectrum-based (SBFL) [1, 13, 20, 31, 44], mutation-based (MBFL) [15, 18, 30, 35] and formula-based (FBFL) [7, 14, 17, 21, 40]. Both SBFL and MBFL techniques compute the suspiciousness of a statement using coverage information from failing and passing test executions. MBFL uses, in addition, information on how test results change after applying different mutations to the program. Both SBFL and MBFL techniques can be seen as may fault localization techniques, in nature: they return locations that

are *likely* to be relevant to the failing execution, based on all executions. We see may fault localization techniques as orthogonal to ours (and to must fault localization techniques in general), since in the trade-off between returning a small set of locations, and returning one that is guaranteed to contain all relevant statements, may techniques prefer the first, while must techniques prefer the second. In the context of repair, there are interesting applications for both.

FBFL techniques represent an error trace using an SMT formula and analyze it to find suspicious locations. These techniques include using error invariants [6, 14, 17, 40], maximum satisfiability [21, 24, 25], and weakest preconditions [7]. What we were able to show in this paper, is that the methods of [6, 14, 21] are not must. In contrast, we believe (though we do not prove it) that the methods of [7, 24, 25] are must. But, what [7, 24, 25] have in common is that they use the semantics of the error trace or the program. Though semantic information can help to further minimize the number of suspicious locations, retrieving it involves using expensive solving-based procedures. Our approach, on the other hand, uses only syntactic information, which makes the fault localization computation relatively cheap; No SMT solving is needed. Thus, these approaches can be seen as complementary to ours.

In the literature there is also a wide range of techniques for automated program repair using formal methods [4, 10, 19, 22, 29, 32, 33, 42]. Both [11] and [37] also use fault localization followed by applying mutations for repair. But, unlike this work, fault localization is applied only for the original program. Also, neither the Tarantula fault localization used in [11] nor the dynamic slicing used in [37] carries the guarantee of being a must fault localization. The tool MUT-APR [3] fixes binary operator faults in C programs, but only targets faults that require one line modification. The tools FORENSiC [5] and MAPLE [34] repair C programs with respect to a formal specification, but they do so by replacing expressions with templates, which are then patched and analysed. SEMGRAFT [28] conducts repair with respect to a reference implementation, but relies on tests for SBFL fault localization of the original program.

9 Conclusion

In this work we define a novel notion of *must* fault localization, that carefully identifies program locations that are relevant for a bug, so that the set is sufficiently small but is guaranteed not to miss desired repairs. We also show that the notion of *must* fault localization should be defined with respect to the repair scheme in use. We show that our notion of must fault localization is particularly useful in pruning the search space of a specific mutation-based repair algorithm.

To the best of our knowledge, we are the first to investigate the widely-used notion of fault localization and to suggest criteria for evaluating its different implementation.

References

1. Abreu, R., Zoetewij, P., Van Gemund, A.J.C.: An evaluation of similarity coefficients for software fault localization. In: Proceedings of the 12th Pacific Rim International Symposium on Dependable Computing, PRDC 2006, pp. 39–46 (2006)
2. Agrawal, H., Horgan, J.R.: Dynamic Program Slicing. In: PLDI, pp. 246–256 (1990)
3. Assiri, F.Y., Bieman, J.M.: MUT-APR: MUTation-based automated program repair research tool. In: Arai, K., Kapoor, S., Bhatia, R. (eds.) FICC 2018. AISC, vol. 887, pp. 256–270. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-03405-4_17
4. Attie, P.C., Dak, K., Bab, A.L., Sakr, M.: Model and program repair via SAT solving. *ACM Trans. Embed. Comput. Syst.* **17**(2), 1–25 (2017)
5. Bloem, R., Drechsler, R., Fey, G., Finder, A., Hofferek, G., Könighofer, R., Raik, J., Repinski, U., Süllow, A.: FoREnSiC— an automatic debugging environment for C programs. In: Biere, A., Nahir, A., Vos, T. (eds.) HVC 2012. LNCS, vol. 7857, pp. 260–265. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39611-3_24
6. Christ, J., Ermis, E., Schäfer, M., Wies, T.: Flow-sensitive fault localization. In: Giacobazzi, R., Berdine, J., Mastroeni, I. (eds.) VMCAI 2013. LNCS, vol. 7737, pp. 189–208. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35873-9_13
7. Christakis, M., Heizmann, M., Mansur, M.N., Schilling, C., Wüstholtz, V.: Semantic fault localization and suspiciousness ranking. In: Vojnar, T., Zhang, L. (eds.) TACAS 2019. LNCS, vol. 11427, pp. 226–243. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17462-0_13
8. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 168–176. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24730-2_15
9. Clarke, E., Kroening, D., Yorav, K.: Behavioral consistency of C and Verilog programs using bounded model checking. In: Proceedings of the Design Automation Conference, 2003, pp. 368–371. IEEE (2003)
10. D’Antoni, L., Samanta, R., Singh, R.: QLOSE: program repair with quantitative objectives. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 383–401. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41540-6_21
11. Debroy, V., Wong, W.E.: Using mutation to automatically suggest fixes for faulty programs. In: 2010 Third International Conference on Software Testing, Verification and Validation (ICST), pp. 65–74. IEEE (2010)
12. Do, H., Elbaum, S., Rothermel, G.: Supporting controlled experimentation with testing techniques: an infrastructure and its potential impact. *Empir. Softw. Eng.* **10**(4), 405–435 (2005)
13. Eric Wong, W., Debroy, V., Choi, B.: A family of code coverage-based heuristics for effective fault localization. *J. Syst. Softw.* **83**(2), 188–208 (2010)
14. Ermis, E., Schäfer, M., Wies, T.: Error invariants. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012. LNCS, vol. 7436, pp. 187–201. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32759-9_17
15. Gong, P., Zhao, R., Li, Z.: Faster mutation-based fault localization with a novel mutation execution strategy. In: Proceedings of the 2015 IEEE 8th International Conference on Software Testing, Verification and Validation Workshops, ICSTW 2015, pp. 1–10. IEEE (2015)

16. Hofer, B., Wotawa, F.: Spectrum enhanced dynamic slicing for better fault localization. *ECAI* **242**, 420–425 (2012)
17. Holzer, A., Schwartz-Narbonne, D., Tabaei Befrouei, M., Weissenbacher, G., Wies, T.: Error invariants for concurrent traces. In: Fitzgerald, J., Heitmeyer, C., Gnesi, S., Philippou, A. (eds.) *FM 2016*. LNCS, vol. 9995, pp. 370–387. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48989-6_23
18. Hong, S., Lee, B., Kwak, T., Jeon, Y., Ko, B., Kim, Y., Kim, M.: Mutation-based fault localization for real-world multilingual programs. In: *ASE*, pp. 464–475 (2015)
19. Jobstmann, B., Griesmayer, A., Bloem, R.: Program repair as a game. In: Etesami, K., Rajamani, S.K. (eds.) *CAV 2005*. LNCS, vol. 3576, pp. 226–238. Springer, Heidelberg (2005). https://doi.org/10.1007/11513988_23
20. Jones, J., Harrold, M., Stasko, J.: Visualization for fault localization. In: *Proceedings of ICSE 2001 Workshop on Software Visualization*, pp. 71–75 (2001)
21. Jose, M., Majumdar, R.: Cause clue clauses: error localization using maximum satisfiability. In: *PLDI*, pp. 437–446 (2011)
22. Kneuss, E., Koukoutos, M., Kuncak, V.: Deductive program repair. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015*. LNCS, vol. 9207, pp. 217–233. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_13
23. Korel, B., Laski, J.: Dynamic program slicing. *Inf. Process. Lett.* **29**, 155–163 (1988)
24. Lamraoui, S.-M., Nakajima, S.: A formula-based approach for automatic fault localization of multi-fault programs. *J. Inf. Process.* **24**, 88–98 (2016)
25. Lamraoui, S.-M., Nakajima, S., Hosobe, H.: Hardened flow-sensitive trace formula for fault localization. In: *ICECCS* (2015)
26. Le Goues, C., Nguyen, T., Forrest, S., Weimer, W.: GenProg: a generic method for automatic software repair. *IEEE Trans. Softw. Eng.* **38**(1), 54–72 (2012)
27. Liu, K., Koyuncu, A., Bissyande, T.F., Kim, D., Klein, J., Le Traon, Y.: You cannot fix what you cannot find! An investigation of fault localization bias in benchmarking automated program repair systems. In: *ICST*, pp. 102–113 (2019)
28. Mechtaev, S., Nguyen, M.-D., Noller, Y., Grunske, L., Roychoudhury, A.: Semantic program repair using a reference implementation. In: *ICSE* (2018)
29. Mechtaev, S., Yi, J., Roychoudhury, A.: Angelix: scalable multiline program patch synthesis via symbolic analysis. In: *ICSE* (2016)
30. Moon, S., Kim, Y., Kim, M., Yoo, S.: Ask the mutants: mutating faulty programs for fault localization. In: *ICST* (2014)
31. Naish, L., Lee, H.J., Ramamohanarao, K.: A model for spectra-based software diagnosis. *ACM Trans. Softw. Eng. Methodol.* **20**(3), 1–32 (2011)
32. Nguyen, H.D.T., Qi, D., Roychoudhury, A., Chandra, S.: SemFix: program repair via semantic analysis. In: *Proceedings of the 2013 International Conference on Software Engineering*, pp. 772–781. IEEE Press (2013)
33. Nguyen, T.V., Weimer, W., Kapur, D., Forrest, S.: Connecting program synthesis and reachability: automatic program repair using test-input generation. In: Legay, A., Margaria, T. (eds.) *TACAS 2017*. LNCS, vol. 10205, pp. 301–318. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_17
34. Nguyen, T.-T., Ta, Q.-T., Chin, W.-N.: Automatic program repair using formal verification and expression templates. In: Enea, C., Piskac, R. (eds.) *VMCAI 2019*. LNCS, vol. 11388, pp. 70–91. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11245-5_4
35. Papadakis, M., Traon, Y.L.: Metallaxis-FL: mutation-based fault localization. *Softw. Test. Verif. Reliab.* **21**(3), 195–214 (2015)

36. Qian, J., Xu, B.: Scenario oriented program slicing. In: Proceedings of the ACM Symposium on Applied Computing, pp. 748–7752 (2008)
37. Repinski, U., Hantson, H., Jenihhin, M., Raik, J., Ubar, R., Guglielmo, G.D., Pravadelli, G., Fummi, F.: Combining dynamic slicing and mutation operators for ESL correction. In: 2012 17th IEEE European Test Symposium (ETS), pp. 1–6. IEEE (2012)
38. Rothenberg, B.-C., Grumberg, O.: Sound and complete mutation-based program repair. In: Fitzgerald, J., Heitmeyer, C., Gnesi, S., Philippou, A. (eds.) FM 2016. LNCS, vol. 9995, pp. 593–611. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48989-6_36
39. Rothenberg, B.-C., Grumberg, O.: Must fault localization for program repair. <https://batg.cswp.cs.technion.ac.il/wp-content/uploads/sites/78/2020/05/MustFaultLocalizationForProgramRepairCav2020.pdf>, May 2020. A full version of the CAV 2020 paper of the same title
40. Schäf, M., Schwartz-Narbonne, D., Wies, T.: Explaining inconsistent code. In: ESEC/FSE, pp. 521–531 (2013)
41. Tan, S.H., Yi, J., Yulis, Mechtaev, S., Roychoudhury, A.: Codeflaws: A programming competition benchmark for evaluating automated program repair tools. In: Proceedings of the 2017 IEEE/ACM 39th International Conference on Software Engineering Companion, ICSE-C 2017, pp. 180–182 (2017)
42. von Essen, C., Jobstmann, B.: Program repair without regret. *Form. Methods Syst. Des.* **47**(1), 26–50 (2015). <https://doi.org/10.1007/s10703-015-0223-6>
43. Wang, Y., Patil, H., Pereira, C., Lueck, G., Gupta, R., Neamtiu, I.: DrDebug: deterministic replay based cyclic debugging with dynamic slicing. In: Proceedings of the 12th ACM/IEEE International Symposium on Code Generation and Optimization, CGO 2014, pp. 98–108 (2014)
44. Wong, W.E., Debroy, V., Gao, R., Li, Y.: The DStar method for effective software fault localization. *IEEE Trans. Reliab.* **63**(1), 290–308 (2014)
45. Wotawa, F.: Fault localization based on dynamic slicing and hitting-set computation. In: Proceedings of the International Conference on Quality Software, pp. 161–170 (2010)
46. Zhang, X., Gupta, N., Gupta, R.: A study of effectiveness of dynamic slicing in locating real faults. *Empir. Softw. Eng.* **12**(2), 143–160 (2007)
47. Zhang, X., Gupta, N., Gupta, R.: Locating faulty code by multiple points slicing. *Softw. Pract. Exp.* **39**(7), 661–699 (2007)

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.