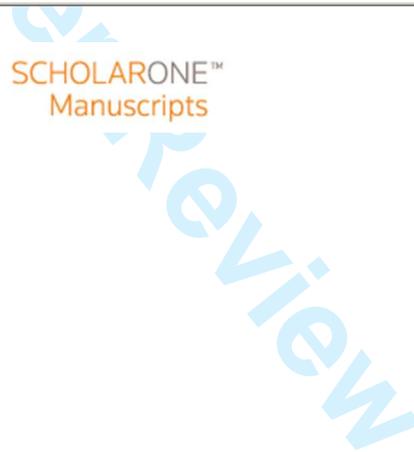


Simulink Fault Localization : an Iterative Statistical Debugging Approach

Journal:	<i>Software Testing, Verification and Reliability</i>
Manuscript ID	STVR-15-0070.R1
Wiley - Manuscript type:	Research Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Liu, Bing; University of Luxembourg, SnT Centre Lucia, Lucia; University of Luxembourg, SnT centre Nejati, Shiva; University of Luxembourg, SnT Centre Briand, Lionel; University of Luxembourg, SnT Centre Bruckmann, Thomas; Delphi Automotive Systems,
Keywords:	Fault localization, Simulink model, Test oracle



SOFTWARE TESTING, VERIFICATION AND RELIABILITY

Softw. Test. Verif. Reliab. 2015; **00**:1–33

Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/stvr

Simulink Fault Localization: an Iterative Statistical Debugging Approach

Bing Liu^{1*} Lucia^{1*} Shiva Nejati^{1*} Lionel Briand¹, and Thomas Bruckmann²

¹*SnT Centre - University of Luxembourg, 4 rue Alphonse Weicker, L-2721, Luxembourg, Luxembourg*

²*Delphi Automotive Systems, Avenue de Luxembourg, L-4940, Bascharage, Luxembourg*

SUMMARY

Debugging Simulink models presents a significant challenge in the embedded industry. In this work, we propose *SimFL*, a fault localization approach for Simulink models by combining statistical debugging and dynamic model slicing. Simulink models, being visual and hierarchical, have multiple outputs at different hierarchy levels. Given a set of outputs to observe for localizing faults, we generate test execution slices, for each test case and output, of the Simulink model. In order to further improve fault localization accuracy, we propose *iSimFL*, an iterative fault localization algorithm. At each iteration, *iSimFL* increases the set of observable outputs by including outputs at lower hierarchy levels, thus increasing the test oracle cost but offsetting it with significantly more precise fault localization. We utilize a heuristic stopping criterion to avoid unnecessary test oracle extension. We evaluate our work on three industrial Simulink models from Delphi Automotive. Our results show that, on average, *SimFL* ranks faulty blocks in the top 8.9% in the list of suspicious blocks. Further, we show that *iSimFL* significantly improves this percentage down to 4.4% by requiring engineers to observe only an average of five additional outputs at lower hierarchy levels on top of high-level model outputs.

KEY WORDS: Fault localization; Simulink model; test oracle

1. INTRODUCTION

The embedded software industry increasingly relies on Simulink to develop software [44, 48, 49, 51]. Simulink is a data-flow-based block diagram language for the modeling, simulation, and development of embedded software. The Simulink language, being supported by advanced automated code generators, has become a prevalent language for implementing embedded software. Automotive software modules are typically first developed as Simulink models from which C code is later generated automatically [35, 57]. These Simulink models are subject to extensive testing and debugging before code generation takes place. Testing Simulink models is the primary testing phase focused on verification of the logic and behavior of automotive software modules. Further, Simulink model testing is more likely to help with fault

*Correspondence to: SnT Centre - University of Luxembourg, 4 rue Alphonse Weicker, L-2721, Luxembourg, Luxembourg. Email : bing.liu@uni.lu, lucia.lucia@uni.lu, shiva.nejati@uni.lu, lionel.briand@uni.lu

finding compared to testing code as Simulink models are more abstract and more informative for engineers. Given the importance of testing and debugging Simulink models, an automated technique to support localization of faults in Simulink models is crucial.

Fault localization in source code is an active research area that focuses on automating various code debugging activities [2, 6, 12, 13, 16, 17, 25, 30, 31, 40, 41, 45, 46, 53, 54, 58, 59]. A well-known approach in this area is *statistical debugging* [2, 25, 30, 31, 45, 46, 53, 54]. Statistical debugging is a light-weight approach to fault localization and has been extensively studied for code (e.g., C programs [2, 25, 45, 60]). This approach utilizes an abstraction of program behavior, also known as *spectra*, (e.g., sequences of executed statements) obtained from testing. The spectra and the testing results, in terms of failed or passed test cases, are used to derive a statistical fault *ranking*, specifying an ordered list of program elements (e.g., statements) likely to be faulty. Developers can consider such ranking to identify faults in their code. These fault localization techniques, however, have never been studied for Simulink models.

In this paper, we propose SimFL, our approach to localize faults in Simulink models based on statistical debugging. Statistical debugging is most effective when it is provided with a large number of observation points (i.e., the spectra size). Existing approaches, where each test case produces one spectrum, require a large test suite to generate a large number of spectra. For Simulink models, however, test suites are typically small. This is mostly because test oracles for embedded software are costly, and further, test suites are required to be eventually applied at the Hardware-in-the-Loop stage where test execution is time consuming and expensive. Hence, we may not obtain a sufficiently large number of spectra if we simply generate one spectrum per each test case as is the case in most existing work [2, 25, 30, 31, 45, 46, 53, 54].

Simulink models, being visual, data-flow based and hierarchical, have multiple observable outputs at different hierarchy levels, each of which can be tested and evaluated independently. For Simulink models, engineers not only identify whether a test case passes or fails, but they also routinely and explicitly determine which specific outputs are correct and which ones are incorrect for each given test case. Relying on this observation, in our work, we use a dynamic slicing technique in conjunction with statistical debugging to generate one spectrum per each output and each test case. Hence, we obtain a set of spectra that is significantly larger than the size of the test suite. We then use this set of spectra to rank model blocks using statistical ranking formulas. In this work, we consider three well-known statistical formulas used for fault localization for source code (i.e., *Tarantula* [25], *Ochiai* [2], and D^* [54]), and study their accuracy in localizing faults in Simulink models. Nevertheless, our approach is not specific to any particular statistical formula and other ones could also be used.

Our approach relies on accounting for as many outputs as possible to increase the number of spectra used to compute rankings. Such expansion, however, should be driven by needs to avoid unreasonable overhead resulting from checking more test oracles on test case executions. In our work, we propose *iSimFL* where we apply *SimFL* iteratively starting with a small test oracle size to obtain some initial ranking results. We provide a heuristic to guide engineers based on the quality of the rankings obtained at each iteration to decide whether they need to extend the oracle or not. We then apply *SimFL* iteratively until the ranking results are satisfactory and do not require any further oracle expansion. In this paper, we make the following contributions:

(1) We propose *SimFL*, a combination of statistical debugging and dynamic slicing to localize faults in Simulink models. Our work is the first to extend statistical debugging to Simulink models. We use dynamic slicing to generate Simulink model spectra such that each spectrum is related to one output and one test case. We then apply statistical ranking formulas to the resulting spectra to compute suspiciousness scores for each Simulink model block.

(2) We propose *iSimFL*, an iterative fault localization approach to refine rankings by increasing the number of observed outputs at each iteration. Our approach utilizes a heuristic stopping criterion to avoid unnecessary expansion of test oracles.

(3) We conduct, for the first time, an empirical study to evaluate statistical debugging for Simulink models using three *industrial* subjects. Our experiments show that: (i) On average, using *SimFL* with *Tarantula*, engineers need to inspect at least 2.1% and at most 8.9% of Simulink models to find faults. Further, we found that the accuracy of *Tarantula*, *Ochiai*, and *D** in localizing faults in Simulink models are considerably close. (ii) We show that increasing the size of test suites does not make any significant improvement in *SimFL*'s accuracy. (iii) We show that in most (but not all) cases, extending the test oracle to include outputs at lower hierarchy levels significantly improves the fault localization capability of *SimFL*. Specifically, *iSimFL* improves *SimFL*'s accuracy by requiring engineers to inspect, on average, at least 1.3% and at most 4.4% of the model blocks while extending test oracles with only five outputs on average. (iv) We investigate the predictability of *iSimFL*'s performance with respect to changes made in to its input heuristic parameters. We, further, provide suggestions, based on the results of our experiments, on how to set these parameters to obtain optimal results.

The paper is organized as follows: In Section 2, we provide background on Simulink models and motivate our approach. In Section 3, we present *SimFL*, our approach to fault localization in Simulink models. In Section 4, we present an iterative fault localization approach, namely *iSimFL* that can further improve the accuracy of *SimFL* in localizing faults. In Section 5, we evaluate our work and discuss our experiment results. We then compare our work with the related work in Section 6. We conclude the paper and discuss our future work in Section 7.

2. BACKGROUND AND MOTIVATION

In this section, we briefly introduce a Simulink model example, describe how Simulink models are tested/simulated in practice, and provide a brief overview of our approach.

Simulink model example. Figure 1 presents a small sanitized snippet of a real-world Simulink model from Delphi. The specific fragment in Figure 1 calculates temperature and pressure of the compressed air. The model is essentially composed of numerical and combinatorial operations as well as some constant blocks, e.g., P_{max} . The model has five inputs, e.g., the position of the Clutch, and two outputs: the output pressure p_{out} and the output temperature T_{out} . The model inputs are specified by input ports (dashed rounded boxes), and the model outputs are shown using output ports (grey rounded boxes). Further, Simulink models are hierarchical and allow the encapsulation of blocks into subsystems. Each subsystem has its own input and output ports. Each input and output port of a subsystem is connected to an atomic block within the subsystem. For example, the model

4

B. LIU ET AL.

in Figure 1 has one hierarchical level i.e., depth 1. At depth 1, the model has two subsystems i.e., Subsystem1 and Subsystem2.

Subsystem1 has two input ports and one output port. Subsystem2 has three input ports and one output port. That is, at depth 1, the model has two outputs. In our work, we denote by depth 0 the root hierarchy level. In our example, at depth 0, the model has two outputs i.e., pOut and TOut.

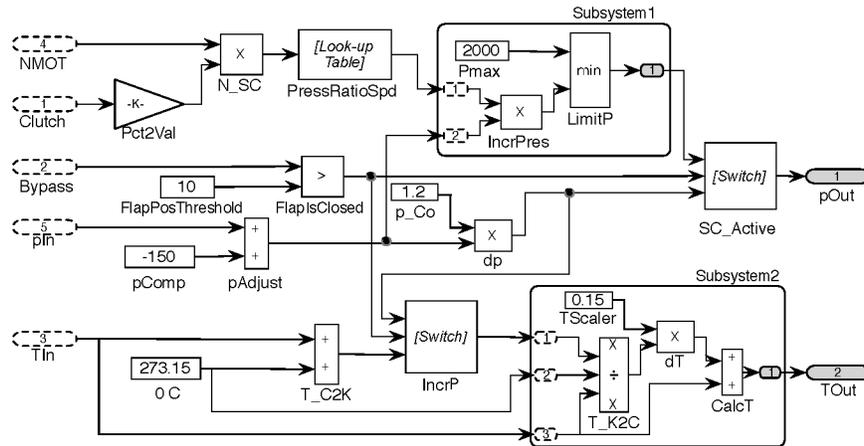


Figure 1. A snippet of a real-world Simulink model.

Input data. Engineers simulate (execute) the model in Figure 1 by providing five input signals, i.e., functions over time. In theory, the input signals can be complex continuous functions. In practice, however, engineers mostly test Simulink models using constant input signals over a fixed time interval. This enables engineers to reproduce the simulation results on different platforms (e.g., when the environment is composed of real hardware or is a real-time simulator). Further, developing test oracles for non-constant input signals is very complex and time consuming. Figure 2(a) shows an input signal example applied to the input pIn. The input signal time interval indicates the simulation length and is chosen to be large enough to let the output signals stabilize.

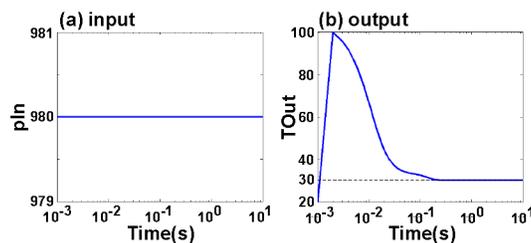


Figure 2. Example of an input (a) and an output (b) signal.

Test output. Similar to the input, the Simulink model output is a signal. Each test case execution (simulation) of a Simulink model results in an individual output signal for each output port of that model. Engineers evaluate each output signal independently. To determine whether an output passes or fails a test case, engineers evaluate various aspects of the output signal, particularly the value at which the output signal stabilizes (if it stabilizes) and the dynamic characteristics of the signal, such as the signal fluctuations (over/undershoot), the response time, and if the signal reaches a steady state. For example, Figure 2(b) shows an example output signal of TOut. As shown in the figure, the

output signal stabilizes after 1 sec of simulation. The output values are the final (stabilized) values of each output signal collected at the end of simulation (e.g., 30 for the signal shown in Figure 2(b)).

For Simulink models from the automotive domain, test oracles are typically developed not just for final model outputs (e.g., pOut and TOut in Figure 1), but also for subsystem outputs (e.g., the output ports of Subsystem1 and Subsystem2 in Figure 1). This is because in this domain Simulink models often capture physical devices (e.g., a supercharger), and as a result, the structure of these models is closely related to the hardware architecture where Simulink subsystems match physical subcomponents. Therefore, a Simulink subsystem often conceptually relates to a real stand-alone entity, and its output is meaningful and can be evaluated independently from the final model outputs. Test oracles in this domain are often costly because they require some degree of engineer's judgement and involvement to determine whether a test case passes or fails. Hence, the size of test suites cannot be very large. In addition, although, in principle, it is possible to derive test oracles for all the outputs at different hierarchy levels, in practice, engineers often develop test oracles for a subset of the most important output ports, e.g., the final model outputs and a few selected subsystem outputs.

Overview of Our Approach. Simulink models consist of blocks and lines. Blocks represent individual functions, whereas lines represent data and control flow relations between blocks. Given a Simulink model with multiple outputs (e.g., Figure 1), for a single test case execution, it often happens that some outputs reveal failures, while others are correct. If any individual output value deviates from its oracle, the engineers typically follow the links connected to that output in a backward manner to identify faulty block(s). This helps engineers focus only on blocks that can reach the erroneous outputs via data and control dependency links. For example, in Figure 1, if pOut shows an error for a test case, we know that the failure at pOut cannot be due to a fault at T_C2K because T_C2K is not in static backward slice of pOut.

Inspired by how Simulink models are debugged in practice, in our work, we define the notion of spectrum per test case and per output. Specifically, we propose in Section 3.2 a dynamic slicing approach to identify, for a given Simulink model, *test execution slices* of that model for each test case and each output. Each test execution slice (spectrum) represents the set of blocks that are executed by a single test case to produce a specific output. As a result, the number of spectra we obtain is larger than the size of the test suite, and hence, we have more observation points to perform fault localization. We propose (in Section 4) to apply our fault localization approach iteratively starting with a small set of outputs. Further, we provide a heuristic which, based on the ranking results from the previous iteration, determines whether extending test oracle is needed or not. If so, engineers extend the output set typically by including subsystem outputs, and perform another fault localization iteration.

Our iterative fault localization approach requires engineers to develop test oracles for several model outputs as well as subsystem outputs. We note that for Simulink models from the automotive domain, test oracles are typically developed not just for final model outputs (e.g., pOut and TOut in Figure 1), but also for subsystem outputs (e.g., the outputs of Subsystem1 and Subsystem2 in Figure 1). This is because in this domain Simulink models often capture physical devices (e.g., a supercharger). Specifically, a Simulink subsystem often conceptually relates to a real stand-alone entity, and its outputs are meaningful and can be evaluated independently from the final model

6

B. LIU ET AL.

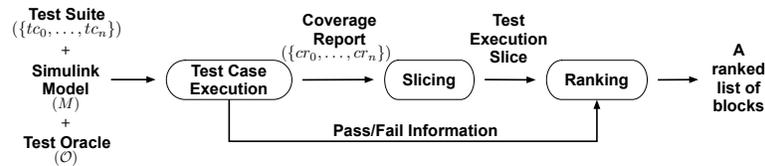


Figure 3. Overview of our fault localization approach for Simulink (*SimFL*).

outputs. Finally, in our approach, we provide a heuristic to guide engineers on when test oracle expansion is unlikely to be worthwhile, thus avoiding unnecessary overhead.

3. FAULT LOCALIZATION FOR SIMULINK

We present *SimFL*, our fault localization approach for Simulink models. Figure 3 shows an overview of *SimFL*. The inputs to our approach is a (faulty) Simulink model (M), a test suite ($TS = \{tc_0, \dots, tc_n\}$), and a test oracle (\mathcal{O}) to determine whether the test cases in TS pass or fail.

Given a Simulink model M , we denote the set of input ports of M by I . For the model in Figure 1, the set I is $\{\text{NMOT}, \text{Clutch}, \text{Bypass}, \text{pIn}, \text{TIn}\}$, and each test case in TS provides a value (i.e., a constant signal) for each element in I . We denote the set of all outputs of M by O , including the model outputs (at depth zero) as well as all the subsystem outputs. For each test case $tc \in TS$, the test oracle \mathcal{O} determines whether each output $o \in O$ passes or fails tc .

The output of the approach in Figure 3 is a ranked list of Simulink (atomic) blocks where the top ranked blocks are more likely to be faulty. This ranked list is generated based on the three main steps of *SimFL*, i.e., Test Case Execution, Slicing, and Ranking, that we discuss in Sections 3.1 to 3.3, respectively.

3.1. Test Case Execution

This step takes as input a test suite TS , a test oracle \mathcal{O} , and a (faulty) Simulink model M . In this step, we execute M for each test case in TS to generate the following information: (1) The PASS/FAIL information corresponding to each output o of M and each test case in TS , and (2) A list $\{cr_0, \dots, cr_n\}$ of coverage reports corresponding to the test cases $\{tc_0, \dots, tc_n\}$.

In Section 2, we discussed how Simulink output signals are typically evaluated to obtain the PASS/FAIL information. In this section, we focus on coverage reports. Given a test case tc_l , Simulink generates a coverage report cr_l after simulating M using tc_l . A coverage report shows the list of atomic blocks that were covered during execution of tc_l .

Using a coverage report describing a list of atomic blocks covered by a test case, we identify which inputs of those blocks were covered by that test case as well. Simulink atomic blocks have two kinds of inputs: data inputs and control inputs. Every (non trivial) atomic block has some data inputs[†]. But they may or may not have control inputs. For a block that has only data inputs, e.g., a multiplication, we know that all its inputs are covered if that block is covered, i.e., appears in the

[†]Some trivial Simulink blocks (e.g., clock) do not have any input.

coverage report. For a block that has control inputs as well as data inputs, e.g., a switch block, the coverage report provides some *block details information* describing which data inputs were covered and which ones were not covered. For example, Figure 4 shows the block details information for a MultiPortSwitch block, which consists of five inputs: one control input (i.e., input 0) and four data inputs (i.e., inputs 1–4). The table in Figure 4 reports which data inputs were actually covered during simulation. Specifically, inputs 1–3 (highlighted in red by Simulink) were not covered, whereas input 4 was covered. That is, the control input 0 selected the data input 4 for the output. Note that the coverage report does not explicitly include the control input 0. However, we know that all the control inputs of a covered block are covered as well. To summarize, from the coverage report in Figure 4, we conclude that among inputs 0–4, only control input 0 and data input 4 were covered during simulation.

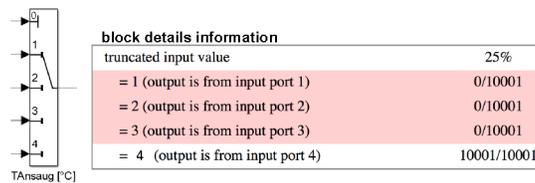


Figure 4. A coverage report snippet generated by Simulink

Note that coverage reports combine the list of covered blocks for all the Simulink outputs. That is, they do not determine which blocks were covered for which specific output. We use slicing (Section 3.2) to determine which blocks were covered for which output.

3.2. Slicing

The second step of *SimFL* is slicing of the input model. This step takes as input the Simulink model M , and the set of coverage reports $\{cr_0, \dots, cr_n\}$ from the first step. The output of this step is a set of test execution slices of M indicating the set of (atomic) blocks that were executed by each test case to generate each output in O .

To generate test execution slices, we first generate a *backward static slice*, denoted by $static_slice_o$, for each output $o \in O$. That is, we set the *slicing criterion*, which also indicates the starting point of the slice, to an output port. Important considerations in slicing are data and control dependencies [44]. As Simulink is a data-flow oriented language, data dependencies are specified by the links between blocks. Starting from an output port, we follow the data dependencies of blocks backwards through the model. If a block is data dependent, then we add this block to our slice. Note that in Simulink, the data dependency links are disconnected at subsystem input/output ports and at Goto/From blocks. In our backward graph traversal, for each subsystem, we ensure to connect its input (respectively, output) ports to the corresponding incoming (respectively, outgoing) links of that subsystem. Similarly, we connect the Goto ports to their matching From ports. The backward traversal stops once we reach the model input ports or constant blocks (shown as orange blocks in Figure 1). Finally, we note that Simulink models may contain feedback loops enabling to use the output of a block for a subsequent calculation [44]. The graphical structure of a Simulink model with a feedback loop is cyclic. To generate (backward) static slices, we detect these cycles and do not go through them more than once.

For example, the model in Figure 5 is an example of a static slice for a MultiPortSwitch block. Suppose that the slicing criterion is the output port Taus. To compute the static slice, we follow the (backward) data dependency to TAnsaug, i.e., the link between TAnsaug and Taus. Hence, we add TAnsaug to the slice. Next, we identify the data dependencies to the blocks TAnsaugdaempfer, Ground, and pEin, and thus, we add these three blocks to the slice. Note that blocks TAnsaugdaempfer and Ground are constant, and hence, do not induce any further data dependencies, while for the block pEin, we may have further data dependent elements that are not shown in the figure.

The backward static slices discussed above are generated for each individual output, but they may contain blocks that do not always affect that output at runtime. For example, the static slice in Figure 5 includes all the blocks connected to the MultiPortSwitch TAnsaug, however, for a given test case, only *some* of the blocks connected to this switch may be executed.

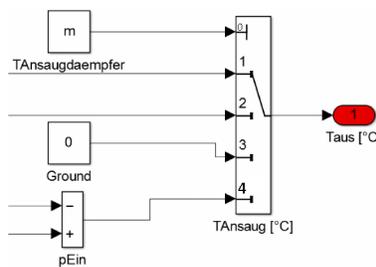


Figure 5. A (partial) static slice for a MultiPortSwitch block.

Having generated backward static slices for each output, we create test execution slices by identifying the subset of these static slices that are executed by each test case. Given an output $o \in O$ and a test case $tc_l \in TS$, in this step, we compute a test execution slice, denoted by $slice_{o,l}$, containing the blocks that are executed by tc_l to generate o . Let $\{cr_0, \dots, cr_n\}$ be the set of coverage reports, and let $\{static_slice_o \mid o \in O\}$ be the set of static backward slices. We define $slice_{o,l}$ as a set of atomics b as follows:

$$slice_{o,l} = \{b \mid b \in static_slice_o \wedge b \in cr_l\}$$

We compute a test execution slice $slice_{o,l}$ by traversing the blocks in the backward static slice of o (i.e., $static_slice_o$) and including in $slice_{o,l}$ those blocks that appear in the coverage report of tc_l (i.e., cr_l). We use cr_l to determine the behavior of Simulink blocks for each test case and for each output. For each block in the static slice, cr_l helps identify which data inputs of that block are selected by its control inputs. For example, as discussed in Section 3.1, the coverage report in Figure 4 indicates that among inputs 0 – 4, only the inputs 0 and 4 are covered during running a test case. Combining this coverage report with the static slice in Figure 5, we obtain a test execution slice which includes TAnsaugdaempfer and pEin blocks. That is, the block Ground is not included in the test execution slice.

A test execution slice $slice_{o,l}$ is passing (respectively failing) if the result of tc_l for output o matches (respectively deviates from) the test oracle for o . Table I shows eight test execution slices corresponding to four test cases (TC1 to TC4) and two (final) outputs for the Simulink model example in Figure 1. In this table, we report for each executed test case and for each output, which blocks were covered (i.e., \checkmark) during the execution. We use the test oracle to determine which execution slices are passing and which ones are failing. The example in Table I consists

of five passing and three failing execution slices. For example, the execution slice for $pOut$ and $TC1$ includes SC_Active , $LimitP$, $IncrPres$, $PressRatioSpd$, etc, because the coverage report for $TC1$ indicated that control block SC_Active selects (for $TC1$) the input coming from control block $LimitP$, and $LimitP$ selects the input coming from $IncrPres$, and so on.

Table I. Test execution slices and suspiciousness scores of model blocks using *Tarantula* for the example model of Figure 1.

Block Name	Test Execution Slices								Scores			Overall Ranking (Min-Max)
	TC1		TC2		TC3		TC4		pOut	TOut	Overall	
	pOut	TOut	pOut	TOut	pOut	TOut	pOut	TOut				
SC_Active	✓		✓		✓		✓		0.5	NaN	0.5	5-13
LimitP	✓		✓						0	NaN	0	14-20
Pmax			✓						0	NaN	0	14-20
IncrPres	✓								0	NaN	0	14-20
PressRatioSpd	✓								0	NaN	0	14-20
N_SC	✓								0	NaN	0	14-20
Pct2Val	✓								0	NaN	0	14-20
FlapIsClosed	✓	✓	✓	✓	✓	✓	✓	✓	0.5	0.5	0.5	5-13
FlapPosThreshold	✓	✓	✓	✓	✓	✓	✓	✓	0.5	0.5	0.5	5-13
dp		✓		✓	✓		✓		0.75	1	0.875	1-2
p_Co		✓		✓	✓		✓		0.75	1	0.875	1-2
pComp	✓	✓	✓	✓			✓		0.6	1	0.8	3-4
pAdjust	✓	✓	✓	✓			✓		0.6	1	0.8	3-4
CalcT		✓		✓		✓		✓	NaN	0.5	0.5	5-13
dT		✓		✓		✓		✓	NaN	0.5	0.5	5-13
TScaler		✓		✓		✓		✓	NaN	0.5	0.5	5-13
T_K2C		✓		✓		✓		✓	NaN	0.5	0.5	5-13
IncrP		✓		✓		✓		✓	NaN	0.5	0.5	5-13
T_C2K						✓		✓	NaN	0	0	14-20
0 C		✓		✓		✓		✓	NaN	0.5	0.5	5-13
Passed/Failed	Passed	Failed	Passed	Failed	Failed	Passed	Passed	Passed				

3.3. Ranking

The third step of our approach is ranking of Simulink blocks. This step takes as input test execution slices from the Slicing step, and the PASS/FAIL information for each test case and for each output from the Test Case Execution step. The output of this step is a ranked list of Simulink (atomic) blocks where each block is ranked with a suspiciousness score. The higher the suspiciousness score of a block, the higher the probability that the block has caused a failure.

To compute the suspiciousness score for a Simulink block, we use three well-known statistical formulas proposed for source code fault localization, namely, *Tarantula* [25], *Ochiai* [2], and D^* [54]. *Tarantula* and *Ochiai* have been the subject of many experiments, and are supported by more substantial empirical evidence than other formulas [7, 21, 23, 24, 32, 33, 39, 46, 55]. Recently, D^* has been shown to outperform 38 statistical formulas in localizing faults for programs [54]. Hence, we decided to focus on these three formulas as a representative set of the many existing statistical ranking formulas. Finally, we note that these formulas are intuitive and easy to explain. This is important as we require involvement of engineers in our experiments. Note that our technique is not tied to any particular ranking formula and can be extended to work with other statistical formulas.

Let s be a statement, and let $passed(s)$ and $failed(s)$ respectively be the number of passed and failed test cases that execute s . Let $totalpassed$ and $totalfailed$ represent the total number of passed and failed test cases, respectively. The suspiciousness score of s according to *Tarantula*, *Ochiai*, D^* denoted by $Score^{Ta}(s)$, $Score^{Oc}(s)$, and $Score^{D^*}(s)$, respectively, are calculated as:

$$Score^{Ta}(s) = \frac{\frac{failed(s)}{totalfailed}}{\frac{passed(s)}{totalpassed} + \frac{failed(s)}{totalfailed}}$$

$$Score^{Oc}(s) = \frac{failed(s)}{\sqrt{totalfailed \times (failed(s) + passed(s))}}$$

$$Score^{D^*}(s) = \frac{failed(s)^*}{(totalfailed - failed(s)) + passed(s)}$$

Similar to [28], we set $*$ to 2 in D^* . Wong et al., [54] show that D^2 is the lowest power D^* variant that still outperforms several existing ranking formulas.

In our work, we compute the suspiciousness score of a Simulink block with respect to each individual output $o \in O$ and denote it by $Score_o$. To compute $Score_o$, we define the functions, $totalpassed_o$, $totalfailed_o$, $passed_o$, and $failed_o$ for every output $o \in O$. Based on the Test Case Execution step, we obtain the set of test execution slices and the pass/fail information for each slice. We define $totalpassed_o$, $totalfailed_o$, $passed_o(b)$ and $failed_o(b)$ for each output $o \in O$ and atomic block b as follows:

$$totalpassed_o = |\{slice_{o,l} \mid \forall tc_l \in TS \wedge slice_{o,l} \text{ is passing}\}|$$

$$totalfailed_o = |\{slice_{o,l} \mid \forall tc_l \in TS \wedge slice_{o,l} \text{ is failing}\}|$$

$$passed_o(b) = |\{slice_{o,l} \mid \forall tc_l \in TS \wedge b \in slice_{o,l} \wedge slice_{o,l} \text{ is passing}\}|$$

$$failed_o(b) = |\{slice_{o,l} \mid \forall tc_l \in TS \wedge b \in slice_{o,l} \wedge slice_{o,l} \text{ is failing}\}|$$

That is, $totalpassed_o$ and $totalfailed_o$ represent the total passing and failing test execution slices, respectively, for output o . The sets $passed_o(b)$, and $failed_o(b)$ represent the numbers of test execution slices that pass and fail, respectively, for output o , and include b . For each output $o \in O$, we define the suspiciousness score of block b for *Tarantula*, $Score_o^{Ta}(b)$, for *Ochiai*, $Score_o^{Oc}(b)$, and for D^* , $Score_o^{D^*}(b)$ as follows:

$$Score_o^{Ta}(b) = \frac{\frac{failed_o(b)}{totalfailed_o}}{\frac{passed_o(b)}{totalpassed_o} + \frac{failed_o(b)}{totalfailed_o}}$$

$$Score_o^{Oc}(b) = \frac{failed_o(b)}{\sqrt{totalfailed_o \times (failed_o(b) + passed_o(b))}}$$

$$Score_o^{D^*}(b) = \frac{failed_o(b)^*}{(totalfailed_o - failed_o(b)) + passed_o(b)}$$

Note that for a block b , $Score_o(b)$ is undefined (*NaN*) if both $passed_o(b)$ and $failed_o(b)$ are zero. This means that b has not appeared in any of the execution slices related to o .

In practice, engineers may either choose to use the scores for each output separately or combine the scores for all outputs. In particular, when there is some indication that failures in different outputs are caused by different faults, e.g., when the test execution slices of different outputs are disjoint, it is preferable to study scores separately. Otherwise, combining scores may improve the accuracy of fault localization, as in typical Simulink models a single faulty block may produce several failures in different outputs.

In our experiment in Section 5, we decided to combine the scores, since we want to assess the overall accuracy for all faults and outputs. We considered and experimented with several alternative

ways of combining the score functions $Score_o$, and based on our experiments computing the average of the scores (see below) yielded the best experiment results. Hence, we use this method to combine the scores of the individual outputs in Section 5.

$$Score(b) = \frac{\sum_{o \in O \wedge Score_o(b) \neq NaN} Score_o(b)}{|\{o \in O \mid Score_o(b) \neq NaN\}|}$$

Having computed the scores, we now rank the blocks based on these scores. The ranking is done by putting the blocks with the same suspiciousness score in the same *rank group*. Given blocks in the same rank group, we do not know in which order the blocks are inspected by engineers to find faults. Hence, we assign a min and a max rank number to each rank group. The min rank for each rank group indicates the least number of blocks that would need to be inspected if the faulty block happens to be in this group and happens to be the first to be inspected. Similarly, the max rank indicates the greatest number blocks that would be inspected if the faulty block happens to be the last to be inspected in that group.

For example, Table I reports the *Tarantula* suspiciousness score for each block and for each of the p0ut and T0ut outputs as well as the mean of these two scores for the example in Figure 1. Note that undefined scores are shown as *NaN* cells and are not used for mean score computation. Table I also shows the block rankings obtained based on the mean scores. According to the overall ranking, the blocks dp and p_Co have the highest ranking (min rank: 1 and max rank: 2). In this example, the block p_Co is faulty causing both p0ut and T0ut to fail for different test cases. Note that if we use the scores for the p0ut and T0ut outputs (without averaging), four blocks dp, p_Co, pComp, and pAdjust appear in the highest rank, whereas the average ranking, which ranks two of these blocks as the highest, is more refined.

4. ITERATIVE FAULT LOCALIZATION

In this section, we describe how the approach in Figure 3 can be applied iteratively, allowing engineers to start with a small test oracle \mathcal{O} , and extend the oracle only when it is necessary. The purpose of iterative fault localization is to enable engineers to select a trade-off between the accuracy of fault localization and the cost of test oracles. The core of our iterative fault localization is a heuristic that guides engineers based on the quality of the ranking obtained at each iteration to determine whether it is worthwhile to continue fault localization with an extended test oracle or not.

Figure 6 shows our iterative fault localization algorithm referred to as *iSimFL*. Similar to *SimFL* (Figure 3), *iSimFL* takes as input a Simulink model M and a test suite TS . Since in *iSimFL*, the test oracle \mathcal{O} is built incrementally, \mathcal{O} is not part of its input. In addition, *iSimFL* receives two input parameters: (1) N which is the number of top most suspicious blocks that engineers typically inspect during fault localization, and (2) g which is a *coarseness threshold*. The coarseness threshold is used to determine whether a given group is too coarse or not. A rank group is too coarse if its size is larger than the maximum number of blocks that engineers can conceivably inspect (i.e., larger than g). These parameters are used in our heuristic and are domain specific. In practice, the values of these parameters are determined by archival analysis of historical fault localization data.

As discussed in Section 2, Simulink models are composed of subsystems blocks that can be hierarchical. Each subsystem at each hierarchical level can have multiple outputs. We denote the

Algorithm. *iSimFL*

Input: - M : Simulink model
 - TS : Test suite
 - N : Number of most suspicious blocks that engineers typically inspect
 - g : Coarseness threshold

Output: - L : A ranked list of blocks

```

1. Let  $h$  be the hierarchy depth of  $M$ , let  $itr = 0$ , and let  $\mathcal{O} = \emptyset$ .
2. do
3.   Let  $\mathcal{O}'$  be the test oracle for outputs at depth  $itr$ , and let  $\mathcal{O} = \mathcal{O} \cup \mathcal{O}'$ 
4.   Let  $L = \{g_0, \dots, g_m\}$  be the ranking list obtained by calling  $SimFL(M, TS, \mathcal{O})$ 
5.   Let  $L' = \{g_0, \dots, g_k\}$  such that  $|\bigcup_{0 \leq i \leq k} g_i| \geq N$  and  $|\bigcup_{0 \leq i \leq k-1} g_i| < N$ 
6.   Let  $g^*$  be the largest group in  $L'$ 
7.   if ( $|g^*| < g$ ) do
8.     break;
9.    $itr++$ 
10. while ( $itr \leq h$ )
11. return  $L$ 

```

Figure 6. Iterative fault localization with *iSimFL*

hierarchy depth of M by h , i.e., the maximum subsystem nesting level. Model M has outputs at hierarchy depths 0 to h . For example, for the model in Figure 1, we have $h = 1$. The outputs p0ut and T0ut are at depth 0, and the outputs of Subsystem1 and Subsystem2 are at depth 1.

In *iSimFL*, we start at hierarchy depth zero ($itr = 0$), and iteratively build test oracle \mathcal{O} such that \mathcal{O} always includes the test oracle data for all the outputs from depth 0 up to depth itr . At each iteration, we call original *SimFL* with test oracle \mathcal{O} (line 4) to obtain a ranked list $L = \{g_0, g_1, \dots, g_m\}$ containing rank groups. Given a ranked list $L = \{g_0, g_1, \dots, g_m\}$, we apply our heuristic to determine whether another iteration of *iSimFL* is worthwhile or not.

Briefly, the intuition behind our heuristic is that engineers cannot effectively localize faults when the ranked list L is *coarse*, particularly within the top blocks in the list. We say a ranked list L is coarse for the top blocks, if, among the rank groups covering the top N blocks, there is a rank group whose size is larger than g (coarseness threshold). Lines 5 to 8 in Figure 6 implement our heuristic. If L happens to be coarse for the N top most blocks, we proceed to the next iteration where we increase itr , extend \mathcal{O} to include test outputs at depth itr , and call *SimFL* with the extended test oracle. Otherwise, we terminate *iSimFL* either when L does not pass our heuristic, i.e., is not coarse (line 8), or when we reach the outputs at depth h of M (line 10).

5. EMPIRICAL EVALUATION

In this section, we present our research questions, describe our industrial subjects, and experimental setup, followed by the analysis of the results.

5.1. Research Questions

RQ1. [*SimFL*'s accuracy] Can *SimFL* help localize faults by ranking the faulty blocks in the top most suspicious blocks? and what is the accuracy of *SimFL* for different statistical formulas? We start by investigating whether *SimFL* can help engineers locate faulty blocks by inspecting a small

subset of the model blocks. Specifically, we report the minimum and maximum number of blocks that engineers have to inspect to identify faulty blocks when they are provided with a ranked list of blocks generated by *SimFL* using *Tarantula*, *Ochiai*, and D^2 . We then compare the accuracy of *SimFL* in localizing faults for these three ranking formulas.

RQ2. [Increasing test suite size] *Does increasing test suite size improve SimFL's accuracy in localizing faults?* In order to increase the spectra size, one can either increase the size of test suites or increase test oracles to include more outputs. Both require effort and have to be investigated. Here, we focus on the former to determine if increasing the size of test suites can improve the accuracy of *SimFL* in localizing faults.

RQ3. [Extending test oracle] *Does extending the set of outputs and correspondingly the test oracle to include more subsystem outputs improve the accuracy of SimFL in localizing faults?* For Simulink models, engineers often try to manually localize faults by inspecting intermediary outputs (i.e., the subsystem outputs at different hierarchy levels) in addition to final model outputs. We investigate the impact of increasing the number of outputs, by including subsystem outputs, on the accuracy of *SimFL* in localizing faults.

RQ4. [*iSimFL* vs *SimFL*] *How do the accuracy results of *iSimFL* compare with those of *SimFL*, and further, does *iSimFL* help limit the size of test oracles while improving accuracy?* Given that developing test oracles for subsystem outputs, though common and feasible, is costly, it is important to evaluate the heuristic we use in *iSimFL* to determine if it can predict when test oracle expansion is worthwhile. That is, when extending test oracles results in significant improvement in fault localization accuracy justifying the expansion overhead.

RQ5. [Impact of *iSimFL*'s parameters] *Does the performance of *iSimFL* change in a predictable way when we vary its input parameters g and N ?* In **RQ4**, we compare the performance of *iSimFL* with that of *SimFL* by giving fixed values to the g and N parameters used in *iSimFL*. It is important to investigate if and how the performance of *iSimFL* is impacted when these parameters change. Specifically, for this question, we report the test oracle size required by *iSimFL* and the accuracy of *iSimFL* for different values of g and N . This data allows us (1) to determine whether the changes to the oracle size and accuracy are monotonic, and hence predictable; and (2) to identify optimal values for g and N . The optimal values of g and N are determined by comparing the results of *SimFL* and *iSimFL* and are those values that lead to a larger oracle size reduction with a negligible accuracy loss.

5.2. Our Industrial Subject

We use three Simulink models developed by Delphi in our experiments. These models simulate physical processes that occur inside the powertrain systems, more specifically, the combustion engine and gearbox behavior. We refer to these three models as *MS*, *MC*, and *MG*. All these three models contain different types of Simulink blocks such as switches, lookup tables, conditional blocks, integrator blocks, From/Gotos, and feedback loops. Table II shows key information about our industrial subjects. For example, Model *MS* contains 37 subsystems, 646 atomic blocks, and 596 links. The hierarchy depth is five, and the model has 12 inputs, 8 outputs at hierarchy depth zero, 8 outputs at depth one, and 7 outputs at depth two. That is, the number of outputs at depths zero and

one (\mathcal{O}_1) is 16, and the number of outputs at depths zero, one, and two (\mathcal{O}_2) is 23. The outputs at depths three to five are redundant because they match those at depths one and two (e.g., in Figure 1, the Subsystem2 output matches T0ut).

Table II. Key information about industrial subjects.

Model Name	# of sub-system	# of atomic blocks	# of links	# of inputs	Hierarchy Depth	# of model outputs		
						(depth 0) \mathcal{O}_0	(depths 0 to 1) \mathcal{O}_1	(depths 0 to 2) \mathcal{O}_2
MS	37	646	596	12	5	8	16	23
MC	64	819	798	13	7	7	11	14
MG	15	295	261	5	4	6	13	17

We asked a Delphi engineer to seed 40 realistic faults in each one of *MS* and *MC*, and 15 realistic faults in *MG*. In total, we generated 95 faulty versions (one fault per each faulty version). The faults were seeded before our experiment took place. The engineer seeded faults based on his past experience in Simulink development and, to achieve diversity in terms of the location and types of faults, we required faults of different types to be seeded in different parts of the models. We categorize the seeded faults into the following three groups: (1) *Wrong Function* which indicates a mistake in the block function type such as choosing $>$ instead of $>=$. (2) *Wrong Connection* which indicates a wrong link between two blocks. For example, engineers may connect the signal A to input 2 instead of input 1 of a block. Note that if the data type of signal A and input 2 does not match, Simulink reports a syntax error. Hence, this fault refers to cases where the types match, but the connection is still wrong. (3) *Wrong Value*, indicating a wrong value in a constant Simulink block or a wrong threshold value in a Simulink control block.

The above classification of faults does not include Stateflow [37], which is the state machine notation of Simulink. This is because (1) our industrial subjects do not include any Stateflows, and (2) we would need to adapt slicing to Stateflow, which is out of the scope of this paper. Further, Simulink models may fail due to the wrong configuration of the simulator, e.g., a wrong step size. In our work, we focus on handling failures caused by faults applied to the model and not those that are due to the wrong configuration of the simulator.

Finally, we note that our industrial subjects are representative in terms of size and complexity among Simulink models developed at Delphi. Our industrial subject models include about ten times more blocks than the publicly available Simulink models from the Mathworks model repository [36]. In addition, most publicly available Simulink models are small exemplars created for the purpose of training for which realistic faults are not available. Hence, we chose to perform our experiments exclusively on industrial subjects for which *realistic* faults could be obtained from an experienced engineer.

5.3. Experiment Settings

In addition to a Simulink model, which is discussed in Section 5.2, *SimFL* requires as input a test suite and a test oracle which are discussed below, along with the experiment design and evaluation metrics.

Test Suite. In this paper, we generated test suites using Adaptive Random Testing [10]. In our experiment, we were provided with the valid ranges of input signals of our industrial subjects.

Adaptive random testing is a black box and lightweight test generation strategy that distributes test cases evenly within the input space (i.e., the valid ranges), and therefore, helps ensure diversity among test cases.

Test Oracle. In practice, the development of test oracles is largely manual and out of scope of this paper. In our experiment, we chose to use a fault-free version of our industrial subject model for the oracle information to automate our large-scale and time-consuming experiments. Note that the Simulink models used in our experiment, when provided with constant input signals, are expected to stabilize and eventually converge to a constant output signal (see Figure 2(b)). If the output signal does not stabilize within a sufficiently large simulation time interval, we mark that as a failure. In this case study, we followed the suggestion from Delphi engineers and set the simulation time (and thus, the moment at which we measure the output signal) to 10 seconds. For each output, to determine if a test case passes or fails, we compared the values of that output from the faulty Simulink model with the fault-free Simulink model at the end of a 10-sec simulation. If they matched, we marked the output as PASS, and otherwise, as FAIL.

Experiments. We perform five separate experiments, referred to as *EXP1*, *EXP2*, *EXP3*, *EXP4*, and *EXP5* to answer our research questions. In our experiments, we consider three different sets of outputs and their corresponding test oracles: (1) test oracle \mathcal{O}_0 for the model outputs at depth zero, (2) test oracle \mathcal{O}_1 for the outputs at depth zero and one, and (3) test oracle \mathcal{O}_2 for the outputs at depths zero, one, and two. Table II shows the sizes of test oracles \mathcal{O}_0 to \mathcal{O}_2 for each industrial subject.

To answer **RQ1**, we perform experiment *EXP1* where we apply *SimFL* (Figure 3) using *Tarantula*, *Ochiai*, and D^2 to our 95 faulty models with a test suite size of 200 and with the smallest test oracle (\mathcal{O}_0). Note that the size selected for test suites was based on typical practice at Delphi given test budget constraints and the cost of oracles. For **RQ2**, we perform experiment *EXP2* where we apply *SimFL* using *Tarantula*, *Ochiai*, and D^2 to our 95 faulty models with the test oracle \mathcal{O}_0 and with nine different test suites of varying size: 200, 300, 400, 500, 600, 700, 800, 900, and 1000. We start from the test suite with 200 test cases and augment the test suites by adding (100) more test cases generated using Adaptive Random Testing. For **RQ3**, we perform experiment *EXP3* where we apply *SimFL* using *Tarantula*, *Ochiai*, and D^2 to our 95 faulty models with a test suite size of 200 and with test oracles \mathcal{O}_1 and \mathcal{O}_2 . Based on the results of *EXP1*, *EXP2*, and *EXP3*, we select one statistical formula to be used by *iSimFL* in *EXP4* to *EXP5*. For **RQ4**, we perform experiment *EXP4* where we apply *iSimFL* (Figure 6) to our 95 faulty models with a test suite size of 200 and rely on the heuristic used in *iSimFL* to determine how many iterations are required for each faulty model. For the parameters N and g used in *iSimFL*, we set their values based on our experience and discussions with domain experts. Specifically, we set $N = 15$ because engineers, when provided with a ranked list of blocks, are able to typically and routinely inspect the top 15 blocks. Further, for each faulty model, we set g to 6% of the size of the model. Finally, for **RQ5**, we perform experiment *EXP5* where we apply *iSimFL* with different values of N and g to our 95 faulty models with a test suite size of 200. *EXP5* consists of two parts i.e., *EXP5a* and *EXP5b*. For *EXP5a*, we apply *iSimFL*

where we fix N to 15 and vary the value of g to 1%, 2%, ..., 10% of the model size. For *EXP5b*, we apply *iSimFL* where we fix g to 6% of the model size and vary the value of N to 5, 10, ..., 25.

Evaluation Metrics. Assuming that engineers inspect block rankings generated by *SimFL* or *iSimFL* to find faults, we evaluate the accuracy of *SimFL* and *iSimFL* using the following metrics from the fault localization literature [13, 24, 31, 33, 41, 45]: The *percentage of blocks inspected* to find faults, the *absolute number of blocks inspected* to find faults, and the *proportion of faults localized* when engineers inspect fixed numbers of the top most suspicious blocks.

For the *absolute number of blocks inspected* to find faults, we consider the min and the max ranks of the rank group that contains the faulty block. For the *percentage of blocks inspected* to find faults, we divide the absolute number of blocks inspected (both for the min and the max ranks) by the total number of blocks. The *proportion of faults localized* is the proportion of localized faults over the total number of faults when engineers inspect a fixed number of the top most suspicious blocks from a ranked list.

5.4. Experiment Results

In this section, we address our research questions based on our experiment results.

RQ1. [*SimFL*'s accuracy] To answer this question, we performed *EXP1* described in Section 5.3. We evaluate *SimFL*'s accuracy in localizing faults in terms of the percentage and the absolute number of blocks inspected, and the proportion of faults localized, as follows:

Percentage and absolute number of blocks inspected. In Table III, we show the percentages and absolute numbers of blocks that engineers need to inspect when they use *SimFL* with the smallest test oracle (i.e. \mathcal{O}_0) for three formulas i.e., *Tarantula*, *Ochiai*, and D^2 . For all 95 models, when using *SimFL* with \mathcal{O}_0 and *Tarantula* as the statistical formula, engineers need to inspect, on average, at least 14 and at most 63 blocks (i.e., 2.1% - 8.9%). Similarly, when using *SimFL* with *Ochiai* as the statistical formula, engineers need to inspect, on average, at least 23 and at most 62 blocks (i.e., 3.1% - 8.8%), and when using *SimFL* with D^2 , engineers need to inspect, on average, at least 16 and at most 56 (i.e., 2.3% - 7.9%).

Table III. Average percentage and absolute number of blocks inspected using *SimFL* with \mathcal{O}_0 for *Tarantula*, *Ochiai*, and D^2 .

Model name	min.#(%) - max.#(%) for <i>SimFL</i> with \mathcal{O}_0		
	<i>Tarantula</i>	<i>Ochiai</i>	D^2
MS	13 (2.1%) - 46 (7.1%)	14(2.2%) - 43(6.7%)	11(1.6%) - 40(6.1%)
MC	19 (2.4%) - 96 (11.7%)	39(4.7%) - 98(12%)	26(3.2%) - 86(10.5%)
MG	4 (1.4%) - 18(6%)	5(1.6%) - 18(6%)	4(1.4%) - 18(6.0%)
All models	14 (2.1%) - 63 (8.9%)	23 (3.1%) - 62 (8.8%)	16 (2.3%) - 56 (7.9%)

Proportion of faults localized. In Figures 7, 8, and 9, we present the proportion of faults localized when engineers inspect a fixed number of the most suspicious blocks in the rank lists generated by *SimFL* with three statistical formulas (i.e., *Tarantula*, *Ochiai*, and D^2) for *MS*, *MC*, and *MG*, respectively. In each figure, the solid line shows the maximum proportion of faults localized, and the dashed line shows the minimum proportion of faults localized.

For 40 faulty versions of *MS* (see Figure 7), by inspecting the top 10% of most suspicious blocks (i.e., 65 blocks), engineers can locate at most 95% and at least 78% of the faults when *SimFL* with

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Tarantula is used; at most 93% and at least 78% of the faults when *SimFL* with *Ochiai* is used; and at most 95% and at least 83% of the faults when *SimFL* with D^2 is used. For 40 faulty versions of *MC* (see Figure 8), by inspecting the top 10% of most suspicious blocks (i.e., 82 blocks), engineers can locate at most 85% and at least 33% of the faults when *SimFL* with *Tarantula* is used; at most 80% and at least 33% of the faults when *SimFL* with *Ochiai* is used; and at most 85% and at least 35% of the faults when *SimFL* with D^2 is used. For 15 faulty versions of *MG* (see Figure 9), by inspecting the top 10% of most suspicious blocks (i.e., 30 blocks), engineers can locate at most 100% and at least 93% of the faults when *SimFL* with *Tarantula* and D^2 are used; and at most 100% and at least 87% of the faults when *SimFL* with *Ochiai* is used.

Note that, for all of the three formulas (i.e., *Tarantula*, *Ochiai*, and D^2), the results of using *SimFL* for *MC* is not as good as the results for the other two models. This is because, compared to *MS* and *MG*, *MC* includes a larger number of lookup tables, integrator blocks, unit convertors, and trigonometry and logarithmic functions that may potentially reduce or mask data discrepancies, and hence, impact the number of observed failures for outputs at depth zero. As a result, the fault localization results for *MC* when we focus on the outputs in \mathcal{O}_0 are less accurate compared to the results for *MS* and *MG*.

Comparing with the state-of-the-art. Since no studies on fault localization for Simulink models are reported, we briefly report the results obtained from applying statistical debugging approaches (with various ranking formulas) to source code implemented in C or Java. We note that, like our work, the approaches discussed here assume that the code under analysis has a single fault.

Comparing the percentage of blocks inspected, on average, according to [13, 31, 33, 45], developers need to inspect at most around 20% of their code (i.e., program blocks) to localize faults, while *SimFL*, on average, requires at most around 8% (i.e., 8.9% for *SimFL* with *Tarantula*, 8.8% for *SimFL* with *Ochiai*, and 7.9% for *SimFL* with D^2) of the model blocks to be inspected to find faults. Comparing the proportion of faults localized, assuming that developers only inspect the top 10% of the most suspicious code elements, on average, the minimum percentage of faults localized is less than 55% [13, 31, 33, 45]. When engineers inspect 10% of the top most suspicious blocks returned by *SimFL*, on average, the minimum percentage of faults localized is around 60% (i.e., 58/95 for *SimFL* with *Tarantula* 57/95 for *SimFL* with *Ochiai*, and 61/95 for *SimFL* with D^2).

While source code debugging and Simulink model debugging have major differences, the above comparison shows that our results are promising and statistical debugging for Simulink models is potentially useful. We note that while inspecting 10% of software code may indeed require developers to review tens or hundreds of KLOC, 10% of a typical Simulink model is often less than 100 blocks. Moreover, engineers are often able to conceptually trace Simulink blocks to abstract functions and concepts, making it easier for them to determine whether an individual block is faulty or not.

Comparing Tarantula, Ochiai, and D^2 . The above results show that the accuracy of ranking results obtained by these three formulas are considerably close. Based on Table III, the percentage and absolute numbers of blocks inspected using *SimFL* with \mathcal{O}_0 for *Tarantula*, *Ochiai*, and D^2 are considerably close. The average maximum percentages of blocks inspected corresponding to the three formulas range from 7.9% to 8.9%. Figure 10 shows the comparison of the proportions of faults localized for all 95 faulty versions when using *SimFL* with *Tarantula*, *Ochiai*, and D^2 . Based on Figure 10, the minimum proportions of faults localized when using the three formulas are close,

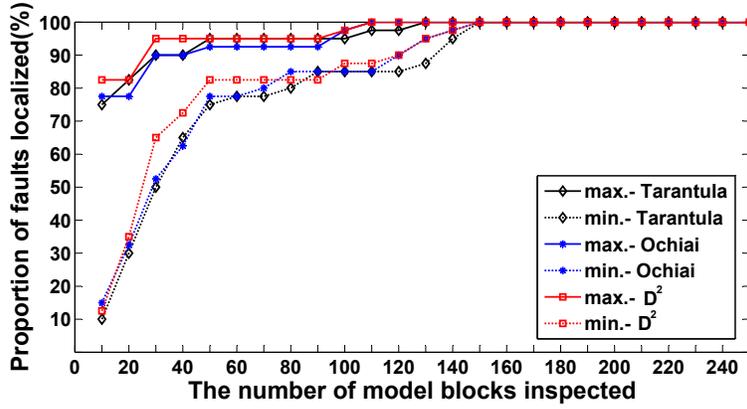


Figure 7. Proportion of faults localized for *MS* using *SimFL* with \mathcal{O}_0 .

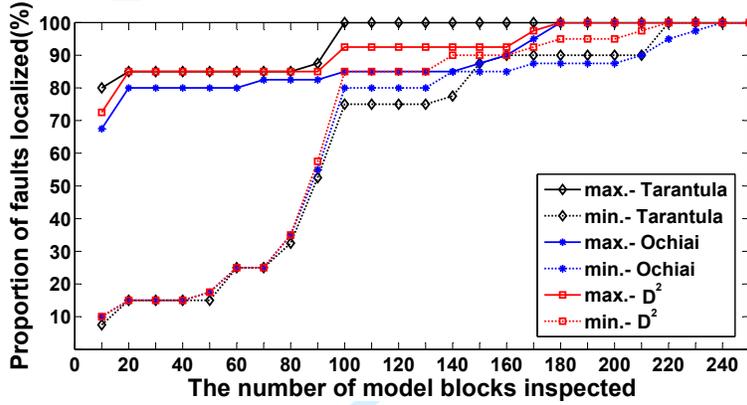


Figure 8. Proportion of faults localized for *MC* using *SimFL* with \mathcal{O}_0 .

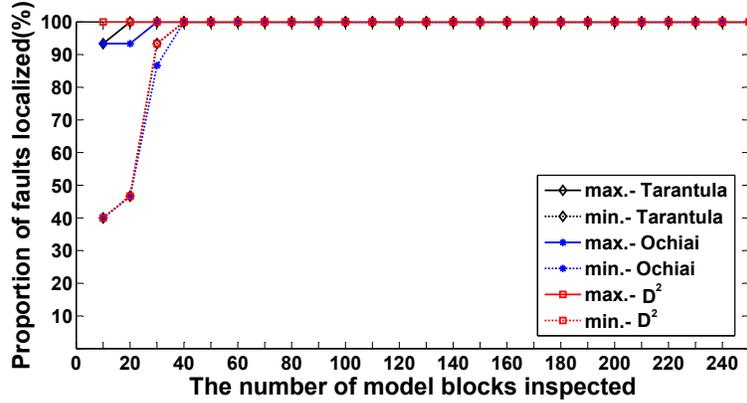


Figure 9. Proportion of faults localized for *MG* using *SimFL* with \mathcal{O}_0 .

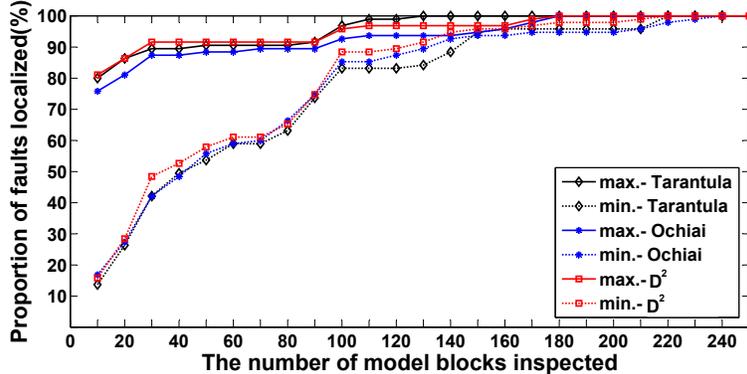


Figure 10. Proportion of faults localized for *all* models using *SimFL* with \mathcal{O}_0 .

1
2
3 in particular when engineers inspect the top 20 blocks. Further, when engineers inspect more than
4 20 blocks, the variations in the minimum proportions of faults localized across the three formulas
5 are less than 7%, and hence, not substantial.
6

7 *In summary*, the answer to **RQ1** is that, on average, *SimFL* is able to rank the faulty blocks as the
8 most suspicious blocks that should be inspected by engineers. Further, the accuracies of *SimFL* using
9 *Tarantula*, *Ochiai*, and D^2 in localizing faults in Simulink models are not substantially different.
10

11 **RQ2.[Increasing test suite size]** To answer this question, we performed *EXP2*. We observed that
12 the maximum number of blocks that engineers need to inspect to find the fault in each faulty model
13 remains almost constant as we apply *SimFL* with *Tarantula*, *SimFL* with *Ochiai*, and *SimFL* with D^2
14 using test suite sizes of 200, 300, . . . , 1000. For all the 95 faulty models and for these three formulas,
15 changes in the rankings of the faulty blocks are negligible as we apply *SimFL* with different test
16 suite sizes. Specifically, for the three formulas, the differences on the maximum percentages of
17 blocks inspected are less than 1.8% as we increase the test suite size. Note that we start with a test
18 suite with size 200 because this size is realistic and comparable to test suite sizes used for Simulink
19 models in Delphi.
20

21 To explain why the rankings of the faulty blocks remain almost constant, we introduce the notion
22 of *Coincidentally Correct Test cases (CCT)* [52]. CCTs are test execution slices that execute faulty
23 blocks but do not result in failure. CCTs are likely to occur in Simulink models because these models
24 often contain various mathematical function blocks that may reduce or mask data discrepancies,
25 resulting in passing test execution slices that exercise faulty blocks. We note that based on the
26 *Tarantula*, *Ochiai*, and D^2 formulas, and given that our faulty models include a single faulty block,
27 the following hold: (1) The *Tarantula* scores of faulty blocks depend on the proportion of CCTs over
28 the total number of passing test execution slices [52], (2) the *Ochiai* scores of faulty blocks depend
29 on the proportion of total failing test execution slices over the number of test execution slices that
30 cover the faulty blocks, and (3) the D^2 scores of faulty blocks depend on the proportion of total
31 failing test execution slices over the number of CCTs.
32

33 In our experiments, we observed that as we increase the test suite size: (1) the proportion of CCT
34 over all passing test execution slices remains almost constant (i.e., changes in this proportion are
35 less than 1%), (2) the proportion of total failing test execution slices over the total number of test
36 execution slices that cover the faulty blocks also remains almost constant with an average difference
37 of 0.004, and (3) the changes in the proportion of total failing test execution slices over the number
38 of CCTs are negligible (i.e., the average difference is 0.22). Hence, increasing the test suite size has
39 no notable impact on the rank of faulty blocks for none of these three formulas.
40

41 *In summary*, the answer to **RQ2** is that increasing the size of test suites, above what can be
42 considered a typical size in our application context, does not make any significant changes in
43 *SimFL*'s accuracy.
44

45 **RQ3.[Extending test oracles]** To answer this question, we performed *EXP3*. In Figures 11(a)
46 to 11(f), we show the maximum percentages of blocks required to be inspected for 58 out of 95
47 faulty models when *SimFL* with *Tarantula* is applied with test oracles \mathcal{O}_0 , \mathcal{O}_1 , and \mathcal{O}_2 . The results
48 for the other 37 models are not shown, as *SimFL* with \mathcal{O}_0 already performs well, i.e., on average, the
49 maximum percentage of blocks inspected is 4.4% and by extending test oracles of those 37 models,
50 the accuracy only slightly improves or remains the same.
51

Table IV. Number of failing execution slices based on test oracles \mathcal{O}_0 , \mathcal{O}_1 , and \mathcal{O}_2 .

Model Name	Test suite size	# of failing slices (min. ~ max.)		
		Test oracle \mathcal{O}_0	Test oracle \mathcal{O}_1	Test oracle \mathcal{O}_2
MS	200	6 ~ 1009	34 ~ 1796	47 ~ 2461
MC	200	8 ~ 1010	8 ~ 1390	8 ~ 1390
MG	200	34 ~ 249	92 ~ 467	102 ~ 467

Among the 58 models under consideration, *SimFL* with \mathcal{O}_0 performs reasonably well for 20 models (Figures 11(a) and 11(b)) as the maximum percentage of blocks inspected for these models is less than 10%. We still chose to show the results for these 20 models because these results are used to answer **RQ4** as well. For the other 38 models (i.e., models in Figures 11(c) to 11(f)), *SimFL* with \mathcal{O}_0 requires engineers to inspect more than 10% of the blocks in order to locate faults (between 10.2% and 26.6%).

For faulty models shown in Figure 11(a), extending test oracles from \mathcal{O}_0 to \mathcal{O}_2 improves *SimFL*'s accuracy slightly (i.e., up to 3%) for 13 models, while for the other three models (i.e., *MS39*, *MG3*, and *MG9*), *SimFL*'s accuracy remains the same. For 29 faulty models (Figures 11(b) to 11(d)), extending test oracles from \mathcal{O}_0 to \mathcal{O}_1 notably improves *SimFL*'s accuracy. Specifically, on average, the maximum percentage of blocks inspected reduces to 1.8%, 3%, and 10% for the models in Figures 11(b), 11(c), and 11(d), respectively. However, for these models, the accuracy improves slightly or remains the same when we extend the oracle to \mathcal{O}_2 . On the contrary, for the 10 models as shown in Figure 11(e), extending the oracle to \mathcal{O}_1 improves *SimFL*'s accuracy slightly, but extending the oracle to \mathcal{O}_2 notably improves the accuracy to, on average, 4%. Note that extending the test oracle could potentially increase the number of failing execution slices that are useful for localizing faults. In Table IV, we show the minimum and maximum numbers of failing execution slices for all the faulty versions of *MS*, *MC*, and *MG*, as we extend the test oracle from \mathcal{O}_0 to \mathcal{O}_2 . For the large difference between the minimum and maximum, we can see that certain faults are much easier to detect than others and hence they result in many more failing execution slices. Based on Table IV, the minimum and maximum numbers of failing execution slices increase or remain the same as we extend the test oracle from \mathcal{O}_0 to \mathcal{O}_2 .

In contrast to the above models, where *SimFL*'s accuracy either improves or stays the same as we expand the oracle, for *MS32*, *MS33*, and *MS14* (Figure 11(f)), *SimFL* may fare worse as we extend the oracle. For *MS32* and *MS33*, the maximum percentages of blocks that engineers need to inspect decrease to below 10% (i.e., 5.8%) when going from \mathcal{O}_0 to \mathcal{O}_1 , but these percentages increase to above 10% (i.e., 15.7% and 10.9%) again when \mathcal{O}_2 is used. As for *MS14*, *SimFL* fares worse when we extend the oracle from \mathcal{O}_0 to \mathcal{O}_1 . But after extending the oracle to \mathcal{O}_2 , we observe a high improvement (i.e., 8.3%).

To explain why test oracle expansion does not always improve accuracy, we note that as we extend the size of test oracles, either the number of CCTs increases or stays the same. In the latter case, *SimFL*'s accuracy either improves or remains the same because none of the new *passing* test execution slices exercise faults, and hence, the block rankings either stay the same or become more accurate. In the former case, however, *SimFL*'s accuracy is unpredictable and may even decrease. Our experiment data confirms this intuition. For the cases where *SimFL*'s accuracy declines as we increase the spectra size, i.e., for *MS32* and *MS33* (from \mathcal{O}_1 to \mathcal{O}_2) and for *MS14* (from \mathcal{O}_0 to \mathcal{O}_1), the size of CCT increases.

Nevertheless, we note that for all but three faulty models, *SimFL*'s accuracy improves or remains the same as we extend the test oracles to include more subsystem outputs. When using *SimFL* with \mathcal{O}_2 , on average, engineers need to inspect, for *MS*, at least 1.7% and at most 4.0% of model blocks (i.e., 11 to 26 blocks); for *MC*, at least 1.1% and at most 4.1% of model blocks (i.e., 9 to 34 blocks); and for *MG*, at least 1.4% and at most 3.4% of model blocks (i.e., 4 to 10 blocks). On average, for all 95 faulty models and using *SimFL* with \mathcal{O}_2 , engineers need to inspect at least 1.4% and at most 4% of model blocks (i.e., 9 to 27 blocks), which is less than the results for *SimFL* with \mathcal{O}_0 (i.e., on average, at least 2.1% and at most 8.9% of model blocks). Furthermore, using *SimFL* with \mathcal{O}_2 , by inspecting only the top 10% of most suspicious blocks, engineers are able to find at least 91 out of 95 faults, which is 33 more faults compared to using *SimFL* with \mathcal{O}_0 .

When we extend the test oracle to \mathcal{O}_2 , the accuracy of *SimFL* with *Ochiai* and *SimFL* with D^2 is similar to the accuracy of *SimFL* with *Tarantula*. Specifically, for all the 95 faulty models and for *SimFL* with *Ochiai*, the number of blocks inspected decreases to, on average, at least 1.8% and at most 4% of the model blocks (i.e., 13 to 27 blocks). Similarly, for *SimFL* with D^2 , the number of blocks inspected decreases to, on average, at least 1.5% and at most 3.6% of the model blocks (i.e., 10 to 24 blocks).

In summary, the answer to **RQ3** is that extending test oracles by including more outputs at lower hierarchy levels may or may not improve *SimFL*'s accuracy in localizing faults on a specific model. But overall, oracle extension leads to the detection of significantly more faults.

Since, as shown in **RQ1** to **RQ3**, there is no significant differences in the accuracies of *Tarantula*, *Ochiai*, and D^2 in localizing faults in Simulink models, we answer **RQ4** and **RQ5** based on *iSimFL* results with one of these formulas. In particular, we report the results for *iSimFL* with *Tarantula* as its results are representative for the other two formulas.

RQ4.[iSimFL vs SimFL] To answer this question, we performed *EXP4*. Our experiment shows that for 37 out of 95 models (not shown in Figure 11), *iSimFL* only performed one iteration before it terminates. That is, the loop in Figure 6 was executed only once and with oracle \mathcal{O}_0 for these 37 models. The maximum percentages of blocks inspected for these 37 models with \mathcal{O}_0 are reasonably low (4.4% on average) and hence, as *iSimFL* correctly predicted, oracle expansion is not necessary. The results of *EXP4* for the other 58 models are shown in Figures 11(g) to 11(l).

For 16 faulty models as shown in Figure 11(g), *iSimFL* extends test oracles although the maximum percentages of blocks inspected using \mathcal{O}_0 are already good (4.2% on average) and oracle expansion does not lead to a substantial improvement. For these models, the *iSimFL* heuristic still extended \mathcal{O}_0 because there were some coarse groups (with size larger than g) below the faulty block but within the top N blocks. Specifically, for eight of these 16 models, *iSimFL* extends test oracle to \mathcal{O}_1 , and for the other eight models (i.e., *MS3*, *MS12*, *MS34*, *MS36*, *MC2*, *MC29*, *MC31*, *MC33*), *iSimFL* extends test oracle to \mathcal{O}_2 .

For four faulty models as shown in Figure 11(h), the maximum percentages of blocks inspected using \mathcal{O}_0 are within an acceptable range (8.8% on average). Nevertheless, extending test oracles to \mathcal{O}_1 is still beneficial. For these models, *iSimFL* correctly extends test oracles to \mathcal{O}_1 . By doing so, on average, the maximum percentage of blocks inspected notably decreases from 8.8% to 1.8% of the model blocks. However, for *MC14* and *MG15*, *iSimFL* unnecessarily extends the oracles to \mathcal{O}_2 while the maximum percentage of blocks inspected remains the same.

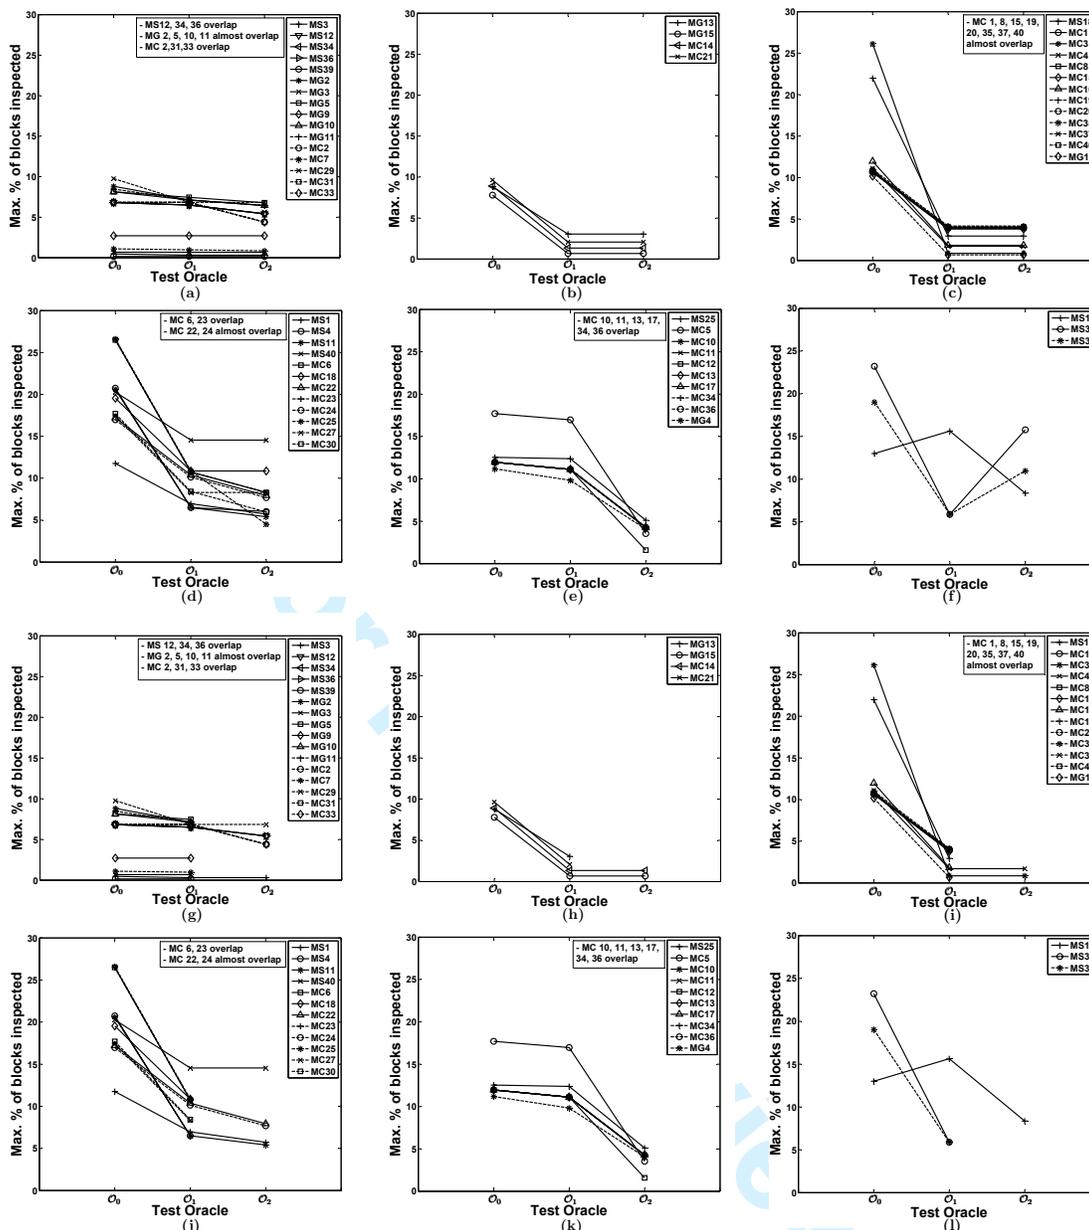


Figure 11. Maximum percentage of blocks that need to be inspected to find faults for *SimFL* with *Tarantula* and test oracles \mathcal{O}_0 , \mathcal{O}_1 , and \mathcal{O}_2 and for *iSimFL* with *Tarantula*: (a) *SimFL*'s accuracy improves slightly or remains the same as we extend the oracle, (b-e) *SimFL*'s accuracy improves notably as we extend the oracle, (f) *SimFL*'s accuracy is unpredictable as we extend the oracle, and (g-l) *iSimFL*'s accuracy for those models where, according to the *iSimFL*'s heuristic, oracle expansion is required.

For the other 38 models (Figures 11(i) to 11(l)), the maximum percentages of blocks inspected using \mathcal{O}_0 are considerably high (15.5% on average). For 34 of the 38 models, *iSimFL* correctly extends oracles which substantially decreases the maximum percentage of blocks inspected. Specifically, *iSimFL* correctly performs two iterations (with \mathcal{O}_0 and \mathcal{O}_1) for 20 models and three iterations (with \mathcal{O}_0 , \mathcal{O}_1 , and \mathcal{O}_2) for 14 models. For these 34 models, *iSimFL* continues extending the oracle either until its accuracy improves and fall below 10%, or until no further extension is possible. Note that only one model (i.e., *MS40*) falls in the latter group. Further, for four models (i.e., *MS1*, *MS11*, *MC3*, and *MC4* (Figure 11(i))), *iSimFL* correctly predicts that extending oracles

to \mathcal{O}_1 is beneficial, though *iSimFL* additionally and unnecessarily extends the oracles to \mathcal{O}_2 while the accuracy remains the same or does not substantially improve.

In summary, oracle extension is not necessary for 53 out of 95 models. For the other 42 models where it is necessary (i.e., leads to considerable improvement in accuracy), 28 models need to extend the test oracle up to depth one (i.e., \mathcal{O}_1), and 14 models require to extend the test oracle up to depth two (i.e., both \mathcal{O}_1 and \mathcal{O}_2).

The *iSimFL* heuristic was able to correctly identify 37 out of 53 models that do not need oracle extension and correctly identify all models (i.e., 42) that require oracle extension. Among these 42 models, the *iSimFL* correctly predict the oracle extension depth for 36 models. For the other six models (i.e., *MS1*, *MS11*, *MC3*, *MC4*, *MC14*, and *MG15*), *iSimFL* correctly extends test oracles to \mathcal{O}_1 , but *iSimFL* unnecessarily extends test oracles further to \mathcal{O}_2 . Further, using *iSimFL*, the average oracle size for each model is about 12 and therefore lower compared to the size of \mathcal{O}_2 (23 for *MS*, 14 for *MC*, and 17 for *MG*). Finally, *iSimFL* was able to properly handle the three cases discussed in **RQ3** where oracle extension caused the accuracy to decline (Figure 11(f)). Specifically, for *MS32* and *MS33*, *iSimFL* stops after applying \mathcal{O}_1 , whereas for *MS14*, it goes all the way to \mathcal{O}_2 .

Table V shows the minimum and maximum numbers (and percentages) of blocks inspected for each industrial subject, comparing *SimFL* with \mathcal{O}_0 , *SimFL* with \mathcal{O}_2 (i.e., extending all oracles), and *iSimFL*. Specifically, after applying *iSimFL* to our 95 faulty models, we obtained the following values for our evaluation metrics:

Percentage and absolute number of blocks inspected. For all models, using *iSimFL*, engineers need to inspect, on average, at least 1.3% and at most 4.4% of model blocks. As shown in Table V, these results are comparable to those obtained by *SimFL* with \mathcal{O}_2 and are better than those obtained by *SimFL* with \mathcal{O}_0 .

Table V. Average of minimum and maximum numbers of blocks inspected and test oracle sizes when using *SimFL-Tarantula* with \mathcal{O}_0 , *SimFL-Tarantula* with \mathcal{O}_2 , and *iSimFL-Tarantula*.

Model name	<i>SimFL</i> with \mathcal{O}_0	<i>SimFL</i> with \mathcal{O}_2	<i>iSimFL</i>
	min. #(%) - max. #(%) ($ \mathcal{O}_0 $)	min. #(%) - max. #(%) ($ \mathcal{O}_2 $)	min. #(%) - max. #(%) (Avg. $ \mathcal{O} $)
MS	13 (2.1%) - 46 (7.1%) (8 outputs)	11 (1.7%) - 26 (4.0%) (23 outputs)	9 (1.5%) - 29 (4.5%) (12 outputs)
MC	19 (2.4%) - 96 (11.7%) (7 outputs)	9 (1.1%) - 34 (4.1%) (14 outputs)	10 (1.2%) - 37 (4.5%) (12 outputs)
MG	4(1.4%) - 18(6.0%) (6 outputs)	4 (1.4%) - 10 (3.4%) (17 outputs)	4 (1.4%) - 11 (3.7%) (11 outputs)

Proportion of faults localized. Using *iSimFL*, engineers can find at least 90 out of 95 faults (i.e., 95%) when only the top 10% of most suspicious blocks are inspected. *iSimFL* is able to locate a similar number of faults compared to *SimFL* with \mathcal{O}_2 (i.e., 90 vs. 91).

In summary, the answer to **RQ4** is that the accuracy of *iSimFL* is similar to the accuracy of *SimFL* with \mathcal{O}_2 , while the average test oracle size for *iSimFL* is 12 compared to a larger size for \mathcal{O}_2 (12 vs. 23 for *MS*, 12 vs. 14 for *MC*, and 11 vs. 17 for *MG*). That is, *iSimFL* achieves the same accuracy as *SimFL* with \mathcal{O}_2 using smaller test oracles. Further, *iSimFL*, with an average oracle size of 12, yields

a significant improvement in accuracy over *SimFL* with \mathcal{O}_0 , which has an average oracle size of 7. That is, *iSimFL* extends only by five outputs the oracle \mathcal{O}_0 .

RQ5.[Impact of *iSimFL*'s parameters] To answer this question, we performed *EXP5a* and *EXP5b* as described in Section 5.3. We evaluated the impact of changes in the values of N and g parameters of *iSimFL* on the average accuracy and the average oracle size extension of *iSimFL*. The reference for comparison is *SimFL* with the maximum oracle (\mathcal{O}_2). Specifically, we want to know, when changing N and g , how the average accuracy and the average oracle size of *iSimFL* fare compared to the accuracy and the test oracle size of *SimFL* with \mathcal{O}_2 .

Figures 12 and 13 show the results of these experiments: In Figure 12, we show the results of *EXP5a* where N is fixed at 15 and we vary the value of g from 1% to 10% of model blocks. Specifically, Figure 12(a) shows the average reduction in the oracle size required by *iSimFL* compared to the size of \mathcal{O}_2 for *MS*, *MC*, and *MG*, and Figure 12(b) shows the average loss in the accuracy of *iSimFL*, which tries to use smaller oracles than \mathcal{O}_2 , compared to the accuracy of *SimFL* with \mathcal{O}_2 for *MS*, *MC*, and *MG*. For example, based on the results in these figures, by applying *iSimFL* to *MS* and when g is set to 3% of the size of *MS*, the average accuracy of the rankings generated by *iSimFL* is around 2 (blocks) less than the average accuracy of rankings obtained by *SimFL* with \mathcal{O}_2 (see Figure 12(b)). But *iSimFL* obtains these rankings with an oracle that contains on average seven less outputs compared to \mathcal{O}_2 (see Figure 12(a)). In Figure 13, we show the results of *EXP5b* where g is set to 6% of the size of the underlying models and N is set to 5, 10, 15, 20, and 25. Similar to Figure 12, Figure 13(a) shows the average reduction in the oracle size required by *iSimFL* compared to the size of \mathcal{O}_2 for *MS*, *MC*, and *MG*, and Figure 13(b) shows the average loss in the accuracy of *iSimFL* compared to the accuracy of *SimFL* with \mathcal{O}_2 for *MS*, *MC*, and *MG*.

As shown in Figure 12, for $N = 15$, as the value of g increases, *iSimFL* extends test oracles less (i.e., the difference between the oracle size required by *iSimFL* and size of \mathcal{O}_2 increases), while the accuracy of ranking results mostly decreases (i.e., engineers on average have to inspect more blocks to find the fault compared to the number of blocks that they need to inspect when *SimFL* with \mathcal{O}_2 is used). This is because for larger g , the probability of finding rank groups with size larger than g decreases and *iSimFL*'s heuristic tends to extend test oracles less often (line 7 in Figure 6). Note that for *MS* and for two points $g = 1\%$ and $g = 5\%$, the accuracy slightly decreases when we increase g . This is because as we observed in **RQ3**, in some few cases by extending test oracles, accuracy may decrease. So although the relationship between g and oracle reduction is monotonic and fully predictable, i.e., oracle size decreases with increasing g , the relationship between g and accuracy loss is not always monotonic. However, as shown in Figure 12(b), in most cases by increasing g , accuracy loss either increases or stays the same, and only in two cases we may slightly gain accuracy by increasing g .

Similarly, when we fix g to 6% of the size of model (Figure 13) and increase N , *iSimFL* extends test oracles more (i.e., the difference between *iSimFL* required oracle size and size of \mathcal{O}_2 decreases), while the accuracy of ranking results increases (i.e., engineers on average have to inspect less blocks to find the fault compared to the number of blocks that they need to inspect when *SimFL* with \mathcal{O}_2 is used). Note that, when the value of N increases, *iSimFL* checks a larger number of most suspicious blocks for deciding whether the suspiciousness ranking is coarse or not. When the set of most suspicious blocks checked by *iSimFL* is larger, *iSimFL* is more likely to find a rank group with

SIMULINK FAULT LOCALIZATION

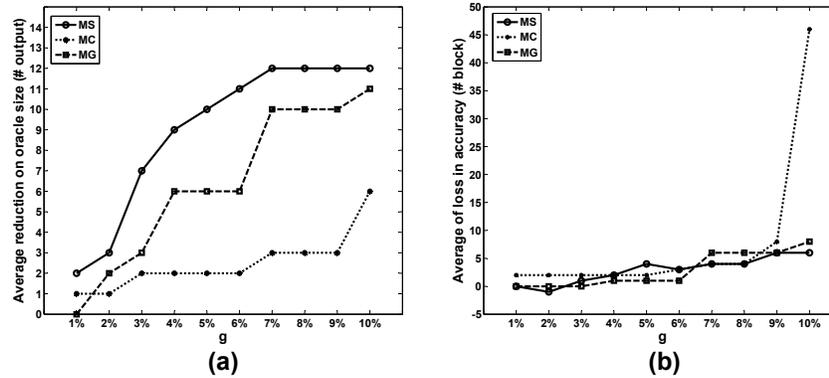


Figure 12. The impact of varying the value of g on the average reduction of oracle size (a) and the average loss in fault localization accuracy (b).

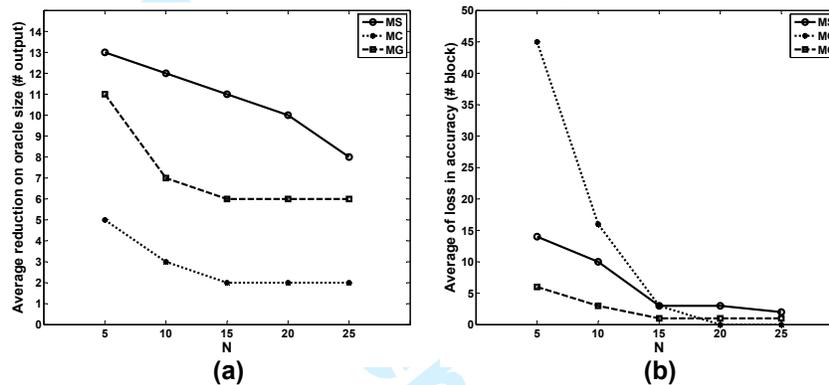


Figure 13. The impact of varying the value of N on the average reduction of oracle size (a) and the average loss in fault localization accuracy (b).

size $> g$ (i.e., coarse ranking results), and hence, is more likely to decide that oracle extension is necessary. As a result, the average reduction on oracle size decreases. On the other hand, as shown in Figure 13(b), as the value of N increases, the accuracy of *iSimFL* gets closer to the accuracy of *SimFL* with \mathcal{O}_2 , i.e., accuracy loss decreases. Note that the trend in Figure 13(b) happens to be monotonic, but as we discussed earlier, the changes in accuracy that are caused by changes in the oracle size are in general unpredictable.

In summary, the answer to **RQ5** is that changing the value of the parameters (i.e., N and g) used in *iSimFL* has a predictable impact on the oracle size required by *iSimFL*. By increasing g , oracle size decreases, and by increasing N , oracle size increases when compared to the size of the maximum oracle (\mathcal{O}_2). The accuracy loss is not always predictable when we change N and g . In a majority of cases, however, by increasing g , the accuracy loss increases, and by increasing N , the accuracy loss decreases when compared with the results obtained by *SimFL* with \mathcal{O}_2 . Finally, based on Figure 12, we observe that when the value of g is between 4% to 6% of the model blocks, the average loss in fault localization accuracy is low (i.e., less than 5 blocks) for all the three models, while reduction in the test oracle size is relatively large (around 8 outputs on average for the three models). Based on Figure 13, we observe that the loss in accuracy is high when N is less than 15, suggesting that checking less than 15 most suspicious blocks may not be enough to assess the coarseness of ranking results and could lead to missing necessary oracle extensions, hence degrading *iSimFL*'s accuracy.

MC is the largest model but also has the smallest variation in oracle size from \mathcal{O}_0 to \mathcal{O}_2 , i.e., there is less room for improvement compared to MS and MG . With the highest value of g and the smallest value of N , the heuristic leads to extending the oracle by only two more outputs, resulting in a larger loss of accuracy compared to MS and MG .

Based on the above results for three distinct models of different sizes, for the experiment whose results are reported in Figure 11, we picked optimal values for g and N , that is 6% and 15, respectively. When setting these parameters in practice, it does not make much sense for g to go higher than 10%, which is already a quite large rank group size. Further, g should not be below 4% of the model size since our results suggest that the reduction in oracle size will be limited. The parameter N is limited by how much time engineers have to inspect blocks. Our results suggest that, for our three subject studies, by setting N to be at least 15, i.e., less than 5% of the model size for our smallest model and less than 2% for our largest model, we are able to provide a reasonable prediction as to when the test oracle extensions are beneficial. Covering such small percentages of blocks is feasible in most practical contexts and situations.

5.5. Threats to Validity

Threats to the external validity relate to the generalizability of our findings. In this work, we evaluated the accuracy of our approach in localizing 95 faulty versions of three industrial Simulink models from the automotive domain. The industrial Simulink models that we analyzed are representative in terms of size and complexity among Simulink models developed at Delphi, and the seeded faults were realistic and were obtained from Delphi engineers. However, it is yet to be seen if our findings are generalizable to Simulink models from other domains.

Threats to the internal validity relate to the assumptions we made in our experiments. In particular, we evaluated our approach on faulty Simulink models where each faulty model contained one fault only. In practice, models may have multiple faults, and these faults may impact one another in unknown ways. However, a large bulk of existing research on applying statistical debugging to code is exclusively evaluated on programs seeded with single faults [2, 5, 9, 11, 12, 13, 24, 25, 30, 31, 32, 41, 45, 46, 56]. Our approach is the first to apply statistical debugging to Simulink models, and no prior empirical results on Simulink fault localization exist. In our work, in order to be able to compare our findings with those reported in the literature, we decided to be consistent with the existing experiment settings and evaluate our approach on models seeded with single faults. Our work is a necessary basis before we can move forward to more complex evaluations involving models seeded with multiple faults. Further, our work opens up opportunities for more research on applying statistical debugging to Simulink models.

6. RELATED WORK

In this section we present the related work to our fault localization approach. First, we discuss fault localization techniques applied to source code that are closely related to our work. Next, we present the existing work on the analysis of Simulink models.

6.1. Software Fault Localization

Many fault localization techniques have been proposed to localize faults in programs [1, 2, 3, 4, 5, 6, 9, 11, 12, 13, 16, 17, 24, 25, 26, 29, 30, 31, 32, 38, 40, 41, 42, 45, 46, 50, 53, 54, 56, 58, 59, 60]. Statistical debugging is one family of fault localization approaches that has been extensively studied to localize faults in programs [1, 2, 3, 4, 5, 9, 11, 24, 25, 26, 29, 30, 31, 32, 38, 42, 45, 46, 50, 54, 56, 60]. Nevertheless, statistical debugging has not been studied to localize faults in Simulink models. In this work, we propose a statistical debugging technique that takes into account the characteristic of Simulink in order to localize faults in Simulink models.

To identify faults in programs, statistical debugging techniques analyze program spectra and use a statistical formula to measure the likelihood of program elements to be faulty. Different types of program spectra have been analyzed to localize faults, e.g. sequences of statements [24, 25, 45, 56, 54], program blocks [1, 2, 32, 45, 60], predicates [30, 31, 42], combination of spectra [46], program path [11]. A number of statistical formulas to measure suspiciousness of program elements have also been proposed e.g., *Tarantula* [25], *Ochiai* [1, 2], formulas from data mining [32], Naish [39], D^* [54], formulas generated using genetic programming [56], SOBER [31], CBI [30]. In this work, we analyze sequences of (atomic) blocks in Simulink and use existing statistical formulas (i.e., *Tarantula*, *Ochiai*, and D^*) to measure the suspiciousness of Simulink (atomic) blocks to be faulty.

Tang et al. [50] build a hierarchy of program predicates using hierarchical clustering and uses it to compute the suspiciousness of each predicate. Parsa et al. [42] focus on comparing the effectiveness of ranking results for different code abstractions. Our work uses intermediary model outputs obtained from different subsystems at different hierarchical levels, and focuses on extending test oracles based on these outputs.

The above techniques [1, 2, 9, 24, 25, 30, 31, 32, 45, 46, 56, 60] localize faults by performing statistical debugging technique only once. Other debugging techniques [5, 11, 42, 50, 61] iteratively apply a statistical debugging technique until developers find the root cause of failures. The techniques proposed in [5, 11] first instrument selected program elements and apply a statistical debugging technique to obtain the most suspicious program element. Developers then check whether the most suspicious program element is faulty or not. If the suspicious element is not faulty, these techniques extend their instrumentation to other program elements, and a statistical debugging technique is applied again to locate faults. Chilimbi et al. [11] search the location of faults by extending their instrumentation to include program elements (i.e., functions) that are highly dependent on the non-suspicious program elements (e.g., functions, branches). A program element is not suspicious if their suspiciousness score is less than a threshold. Nainar and Liblit [5] extend their instrumentation to include program elements (i.e., predicates) that are nearby to the most suspicious program element (i.e., predicates). Their intuition is that predicates that are nearby to the most suspicious predicate are also suspicious. Instead of extending the instrumentation to include other program elements, Zuo et al. [61] search the location of faults using hierarchical instrumentation. They first instrument functions in a program and use a statistical debugging technique to rank functions. They then instrument predicates of the functions that appear in the top rank and run the statistical debugging technique to locate the faulty predicates. The existing iterative debugging techniques [5, 11, 61] focus on extending and improving program instrumentation to reduce memory and time required by the instrumentation. In our work, however, we focus on

extending test oracles to improve the accuracy of statistical debugging in localizing faults. Further, our approach does not require engineers to inspect the ranked list first in order to decide whether or not another iteration is needed, since our heuristic automatically predicts whether another iteration of fault localization is needed or not.

Gong et al. [14] refine suspiciousness rankings returned by a statistical debugging techniques by using developer feedback (i.e., whether a program element is faulty or not) to adjust the suspiciousness scores of program elements and rerank the program elements. Our approach refines the ranked lists by asking engineers whether some selected intermediary outputs are correct or not, and use this information to narrow down the potential faulty Simulink blocks.

Program slicing has been used to refine ranking results produced by statistical debugging techniques [3, 4, 20, 29, 34, 38]. Alves et al. [4] first obtain one spectrum for each test case and apply statistical debugging on these spectra. They then prune the ranking results by removing the statements that are not in the dynamic slicing of incorrect outputs. Our work is different as instead of using dynamic slicing to prune the ranking results, we first obtain one spectrum (test execution slice) per output and per test case, and then apply statistical debugging on all the test execution slices. When multiple outputs are incorrect in a test execution, the techniques proposed in [29, 34] rely on the spectra related only to the first failing output and use only those spectra to compute rankings. Hofer et al. [20] use the spectra from incorrect outputs to generate sets of program elements that can explain failures. In our work, we use spectra from all outputs instead of only one output [29, 34] or only incorrect outputs [20]. Further, we compute block scores based on the combined spectra related to all the outputs. Note that in our work, we notice that some output may fail for all test executions. That is, considering the spectra only from one output is not sufficient to produce meaningful rankings. Finally, we also provide a heuristic to improve the ranking results by guiding engineers on how to extend test oracles.

Statistical debugging has been previously extended to spreadsheets [18, 19, 22] and to multi-agent systems [43] as well. Similar to our work, in both of these approaches, the notion of spectrum is defined per test case and per output. Specifically, in the spreadsheet fault localization approach, only one test case is used for each spreadsheet, and each execution of that test case for each output amounts to one spectrum [18, 19, 22]. In the multi-agent system fault localization, the behavior of each agent at each time step is considered to be a spectrum [43]. Our work bears some similarities with these two approaches with respect to the notion of spectrum and application of ranking formulas. However, our approach is applied to Simulink models that are drastically different from spreadsheets and multi-agent systems in terms of structure and usage, and further, the computation of Simulink model spectra (test execution slices) is significantly different in our work compared to these approaches.

Gonzales et al. [15] and Campos et al. [8], respectively, prioritize and generate test cases with the goal of reducing the size and number of ambiguity groups in ranked lists. In our work, we use the maximum size of the top-most ranked ambiguity groups in a ranked list as a heuristic to choose whether to continue expanding test oracles or not. Specifically, our coarseness measure differs from those in [8, 15] in that we focus on the size of ambiguity groups that ranked in the top of the list as opposed to the size of ambiguity groups in the entire list.

Statistical debugging assumes developers can find faults by inspecting statements in isolation, while in reality they often need context information to decide if a statement is faulty or not [41].

1
2
3 Like existing work, we generate block rankings without including context information. However,
4 Simulink blocks often contain some implicit context information since engineers often label them
5 with terms coming from requirements or architecture. For example, the multiplication block with
6 label `IncrPres` in Figure 1 refers to an operation for increasing the pressure of supercharger. This
7 observation suggests that block rankings could be useful to find faults in Simulink.
8
9

10 11 6.2. Analysis of Simulink Models

12
13 In our work, we relied on model simulations to identify control dependencies between Simulink
14 blocks. Reicherdt and Glesner [44] proposed a slicing method for Simulink models where control
15 dependencies are obtained via Simulink Conditional Execution Contexts (CECs) and are used
16 to create static slices based on a set of blocks. In our work, we chose to use model execution
17 information to identify control dependencies and compute slices since the static slicing of [44, 49]
18 based on CECs may provide over approximations that may not be sufficiently precise to determine
19 control dependencies.
20
21

22 Our work relates to the recent work of Schneider [47] that proposes a technique for tracking
23 the root causes of defects in Simulink. In that technique, engineers identify failures, typically run-
24 time failures, at the level of code generated from Simulink models. The program statement that
25 exhibits the failure is then mapped to a Simulink block, and all the paths leading to that block are
26 collected and assigned weights based on some heuristic. The path with the highest weight is then
27 reported to the engineer as the root cause of the defects. This work focuses on runtime failures
28 (e.g., division by zero), while in our work, we consider a wider range of fault types for Simulink
29 models (see Section 5.2). Further, in [47], the author does not provide any realistic evaluation of the
30 proposed approach. In particular, the number of blocks that engineers need to eventually inspect is
31 not reported. Finally, the scalability of the approach to large models is not discussed as the number
32 of paths leading to a specific block can be very large for real-world Simulink models.
33
34
35
36
37

38 7. CONCLUSION AND FUTURE WORK

39
40 We presented *SimFL*, a new fault localization approach for Simulink models by combining statistical
41 debugging and dynamic model slicing. In our work, we generate finer grained spectra (i.e., one
42 spectrum for each test case and each output) compared to the existing techniques where one test
43 case yields a single spectrum. This allows us to apply statistical debugging to Simulink models
44 where test suites are typically small due to the practical limits of embedded system development.
45 We use backward static slicing and coverage reports to generate test execution slices. We then
46 compute suspiciousness scores per block and per output using three different, well-known statistical
47 ranking formulas and take the average of suspiciousness scores of each block over all outputs to
48 obtain the final scores used for ranking. Our approach considers as many outputs as possible and
49 necessary, potentially increasing test oracle cost. Hence, we propose an iterative fault localization
50 algorithm (*iSimFL*) to help engineers determine when oracle extension is likely to increase accuracy.
51 We applied *SimFL* to 95 faulty models generated based on three different Simulink models from
52 the automotive industry. Our results show that *SimFL*'s accuracy in localizing faults in Simulink
53 models is promising: on average, for example, using *SimFL* with *Tarantula*, the percentage of blocks
54
55
56
57
58
59
60

inspected is at least 2.1% and at most 8.9% of the total model blocks. In contrast to fault localization for source code, we found that the accuracy of *Tarantula*, *Ochiai*, and D^2 in localizing fault in Simulink models are very similar. Further, we show that increasing the size of test suites, above what is common practice in embedded systems, does not significantly change *SimFL*'s accuracy. Hence, to improve accuracy, we extend test oracles using *iSimFL*, a method to iteratively refine them and augment their failure detection capability. We show that *iSimFL* significantly improves *SimFL*'s accuracy (i.e. on average, at least 1.3% and at most 4.4% of the total model blocks need to be inspected) by extending test oracles with only five outputs on average.

The performance of *iSimFL* depends on a stopping criterion heuristic, which is tunable via parameters N (the number of top most suspicious blocks inspected) and g (coarseness threshold). Our analysis shows that changing the value of N and g has a predictable impact on the test oracle size required by *iSimFL*. Further, for the majority of cases, the impact on the accuracy of *iSimFL* is also predictable. This is expected to facilitate the setting of such parameters. In this work, we relied on our experience and discussions with domain experts to determine the value of *iSimFL*'s parameters. Practical guidelines for choosing values for N and g require further studies and are left for future work.

Our results suggest that the three well-known statistical formulas (i.e., *Tarantula*, *Ochiai*, and D^2) yield similar accuracy in localizing faults for Simulink models. In future, we will investigate why different statistical formulas have similar impacts on fault localization accuracy. Moreover, we plan to extend *SimFL* to localize faults in Stateflow (state machine) models. In addition, we intend to perform user studies with engineers to better understand their information needs while debugging, so as to provide additional insights along with the block rankings.

ACKNOWLEDGEMENT

We thank Mehrdad Sabetzadeh and anonymous reviewers for their useful comments on an earlier draft of this paper. This work is supported by the Fonds National de la Recherche, Luxembourg (FNR/P10/03 and FNR8003491), and Delphi Automotive Systems, Luxembourg.

REFERENCES

1. Rui Abreu, Peter Zoetewij, Rob Golsteijn, and Arjan J.C. Van Gemund. A practical evaluation of spectrum-based fault localization. *Journal of Systems and Software*, 82(11):1780–1792, 2009.
2. Rui Abreu, Peter Zoetewij, and Arjan J.C. Van Gemund. On the accuracy of spectrum-based fault localization. In *Testing: Academic and Industrial Conference Practice and Research Techniques-MUTATION, 2007. TAICPART-MUTATION 2007*, pages 89–98. IEEE, 2007.
3. Rui Abreu, Peter Zoetewij, and Arjan J.C. Van Gemund. Spectrum-based multiple fault localization. In *Proceedings of the 24th IEEE/ACM International Conference on Automated Software Engineering, ASE '09*, pages 88–99. IEEE, 2009.
4. Elton Alves, Milos Gligoric, Vilas Jagannath, and Marcelo d'Amorim. Fault-localization using dynamic slicing and change impact analysis. In *Proceedings of the 2011 26th IEEE/ACM International Conference on Automated Software Engineering*, pages 520–523. IEEE Computer Society, 2011.
5. Piramanayagam Arumuga Nainar and Ben Liblit. Adaptive bug isolation. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 1*, pages 255–264. ACM, 2010.

6. Thomas Ball, Mayur Naik, and Sriram K. Rajamani. From symptom to cause: localizing errors in counterexample traces. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '03*, pages 97–105, 2003.
7. Benoit Baudry, Franck Fleurey, and Yves Le Traon. Improving test suites for efficient fault localization. In *Proceedings of the 28th international conference on Software engineering*, pages 82–91. ACM, 2006.
8. Juan Campos, Rui Abreu, Gordon Fraser, and Marcelo d'Amorim. Entropy-based test generation for improved fault localization. In *Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on*, pages 257–267. IEEE, 2013.
9. Mike Y Chen, Emre Kiciman, Eugene Fratkin, Armando Fox, and Eric Brewer. Pinpoint: Problem determination in large, dynamic internet services. In *Proceedings. International Conference on Dependable Systems and Networks, 2002. DSN 2002.*, pages 595–604. IEEE, 2002.
10. Tsong Yueh Chen, Hing Leung, and IK Mak. Adaptive random testing. In *Advances in Computer Science-ASIAN 2004. Higher-Level Decision Making*, pages 320–329. Springer, 2005.
11. Trishul M Chilimbi, Ben Liblit, Krishna Mehra, Aditya V Nori, and Kapil Vaswani. Holmes: Effective statistical debugging via efficient path profiling. In *Software Engineering, 2009. ICSE 2009. IEEE 31st International Conference on*, pages 34–44. IEEE, 2009.
12. Holger Cleve and Andreas Zeller. Finding failure causes through automated testing. In *Proceedings of the Fourth International Workshop on Automated Debugging, AADEBUG 00*, 2000.
13. Holger Cleve and Andreas Zeller. Locating causes of program failures. In *Proceedings of the 27th International Conference on Software Engineering, ICSE '05*, pages 342–351. ACM, 2005.
14. Liang Gong, Daniel Lo, Lingxiao Jiang, and Hongyu Zhang. Interactive fault localization leveraging simple user feedback. In *Software Maintenance (ICSM), 2012 28th IEEE International Conference on*, pages 67–76. IEEE, 2012.
15. Alberto Gonzalez-Sanchez, Rui Abreu, Hans-Gerhard Gross, and Arjan JC Van Gemund. Prioritizing tests for fault localization through ambiguity group reduction. In *Automated Software Engineering (ASE), 2011 26th IEEE/ACM International Conference on*, pages 83–92. IEEE, 2011.
16. Alex Groce, Daniel Kroening, and Flavio Lerda. Understanding counterexamples with explain. In *Computer Aided Verification*, pages 453–456. Springer, 2004.
17. Ralf Hildebrandt and Andreas Zeller. Simplifying failure-inducing input. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA '00*, pages 135–145, 2000.
18. Birgit Hofer, Alexandre Perez, Rui Abreu, and Franz Wotawa. On the empirical evaluation of similarity coefficients for spreadsheets fault localization. *Automated Software Engineering*, 22(1):47–74, 2014.
19. Birgit Hofer, André Riboira, Franz Wotawa, Rui Abreu, and Elisabeth Getzner. On the empirical evaluation of fault localization techniques for spreadsheets. In *Fundamental Approaches to Software Engineering*, pages 68–82. Springer, 2013.
20. Birgit Hofer and Franz Wotawa. Spectrum enhanced dynamic slicing for better fault localization. In *ECAI*, pages 420–425, 2012.
21. Hwa-You Hsu, James A Jones, and Alessandro Orso. Rapid: Identifying bug signatures to support debugging activities. In *Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering*, pages 439–442. IEEE Computer Society, 2008.
22. Dietmar Jannach, Thomas Schmitz, Birgit Hofer, and Franz Wotawa. Avoiding, finding and fixing spreadsheet errors—a survey of automated approaches for spreadsheet qa. *Journal of Systems and Software*, 94:129–150, 2014.
23. James A Jones, James F Bowring, and Mary Jean Harrold. Debugging in parallel. In *Proceedings of the 2007 international symposium on Software testing and analysis*, pages 16–26. ACM, 2007.
24. James A. Jones and Mary Jean Harrold. Empirical evaluation of the tarantula automatic fault-localization technique. In *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering, ASE '05*, pages 273–282, 2005.
25. James A. Jones, Mary Jean Harrold, and John Stasko. Visualization of test information to assist fault localization. In *Proceedings of the 24th International Conference on Software Engineering, ICSE '02*, pages 467–477, 2002.
26. Jeongho Kim, Jonghee Park, and Eunseok Lee. A new hybrid algorithm for software fault localization. In *Proceedings of the 9th International Conference on Ubiquitous Information Management and Communication*, page 50. ACM, 2015.
27. Tien-Duy B Le and Daniel Lo. Will fault localization work for these failures? an automated approach to predict effectiveness of fault localization tools. In *Software Maintenance (ICSM), 2013 29th IEEE International Conference on*, pages 310–319. IEEE, 2013.
28. Tien-Duy B Le, David Lo, and Ferdian Thung. Should i follow this fault localization tool's output? *Empirical Software Engineering*, pages 1–38, 2014.

29. Yan Lei, Xiaoguang Mao, Ziyang Dai, and Chengsong Wang. Effective statistical fault localization using program slices. In *Computer Software and Applications Conference (COMPSAC), 2012 IEEE 36th Annual*, pages 1–10. IEEE, 2012.
30. Ben Liblit, Mayur Naik, Alice X. Zheng, Alex Aiken, and Michael I. Jordan. Scalable statistical bug isolation. *ACM SIGPLAN Notices*, 40(6):15–26, 2005.
31. Chao Liu, Xifeng Yan, Long Fei, Jiawei Han, and Samuel P Midkiff. Sober: statistical model-based bug localization. *ACM SIGSOFT Software Engineering Notes*, 30(5):286–295, 2005.
32. David Lo, Lingxiao Jiang, Aditya Budi, et al. Comprehensive evaluation of association measures for fault localization. In *Software Maintenance (ICSM), 2010 IEEE International Conference on*, pages 1–10. IEEE, 2010.
33. Lucia, David Lo, and Xin Xia. Fusion fault localizers. In *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering, ASE '14*, pages 127–138, 2014.
34. Xiaoguang Mao, Yan Lei, Ziyang Dai, Yuhua Qi, and Chengsong Wang. Slice-based statistical fault localization. *Journal of Systems and Software*, 89:51–62, 2014.
35. MathWorks. Simulink. <http://www.mathworks.nl/solutions/model-based-design/>.
36. MathWorks. Simulink Examples. <http://nl.mathworks.com/help/simulink/examples.html>.
37. MathWorks. Stateflow. <http://www.mathworks.nl/products/stateflow/>.
38. Wolfgang Mayer, Rui Abreu, Markus Stumptner, AJ van Gemund, et al. Prioritising model-based debugging diagnostic reports. In *Proceedings of the International Workshop on Principles of Diagnosis (DX)*, pages 127–134, 2009.
39. Lee Naish, Hua Jie Lee, and Kotagiri Ramamohanarao. A model for spectra-based software diagnosis. *ACM Transactions on software engineering and methodology (TOSEM)*, 20(3):11, 2011.
40. Alessandro Orso, James A. Jones, Mary Jean Harrold, and John T. Stasko. Gammatella: Visualization of program-execution data for deployed software. In *Proceedings of the 26th International Conference on Software Engineering, ICSE '04*, pages 699–700, 2004.
41. Chris Parnin and Alessandro Orso. Are automated debugging techniques actually helping programmers? In *Proceedings of the 20th International Symposium on Software Testing and Analysis, ISSTA'11*, pages 199–209, 2011.
42. Saeed Parsa, Mojtaba Vahidi-Asl, and Maryam Asadi-Aghbolaghi. Hierarchy-debug: a scalable statistical technique for fault localization. *Software Quality Journal*, 22(3):427–466, 2014.
43. Lúcio S Passos, Rui Abreu, and Rosaldo JF Rossetti. Spectrum-based fault localisation for multi-agent systems. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 1134–1140. AAAI Press, 2015.
44. R. Reicherdt and S. Glesner. Slicing matlab simulink models. In *Proceedings of the 34th International Conference on Software Engineering, ICSE'12*, pages 551–561, 2012.
45. Manos Renieris and Steven P. Reiss. Fault localization with nearest neighbor queries. In *Proceedings of the 18th IEEE International Conference on Automated Software Engineering, ASE'03*, pages 30–39, 2003.
46. Raul Santelices, James A Jones, Yanbing Yu, and Mary Jean Harrold. Lightweight fault-localization using multiple coverage types. In *Software Engineering, 2009. ICSE 2009. IEEE 31st International Conference on*, pages 56–66. IEEE, 2009.
47. Johanna Schneider. Tracking down root causes of defects in simulink models. In *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering, ASE '14*, pages 599–604, 2014.
48. P Skruch, M Panek, and B Kowalczyk. Model-based testing in embedded automotive systems. *Model-Based Testing for Embedded Systems*, pages 293–308, 2011.
49. Adepu Sridhar and D. Srinivasulu. Slicing matlab simulink/stateflow models. In *Intelligent Computing, Networking, and Informatics*, pages 737–743. Springer, 2014.
50. Chung Man Tang, WK Chan, and YT Yu. Extending the theoretical fault localization effectiveness hierarchy with empirical results at different code abstraction levels. In *Computer Software and Applications Conference (COMPSAC), 2014 IEEE 38th Annual*, pages 161–170. IEEE, 2014.
51. Andreas Thums and Jochen Quante. Reengineering embedded automotive software. In *Proceedings of the 28th IEEE International Conference on Software Maintenance, ICSM '12*, pages 493–502, 2012.
52. Xinming Wang, Shing-Chi Cheung, Wing Kwong Chan, and Zhenyu Zhang. Taming coincidental correctness: Coverage refinement with context patterns to improve fault localization. In *Proceedings of the 31st International Conference on Software Engineering, ICSE '09*, pages 45–55. IEEE Computer Society, 2009.
53. Eric Wong, Tingting Wei, Yu Qi, and Lei Zhao. A crosstab-based statistical method for effective fault localization. In *Proceedings of the 2008 International Conference on Software Testing, Verification, and Validation, ICST '08*, pages 42–51, 2008.
54. W Eric Wong, Vidroha Debroy, Ruizhi Gao, and Yihao Li. The dstar method for effective software fault localization. *Reliability, IEEE Transactions on*, 63(1):290–308, 2014.

55. Xiaoyuan Xie, Tsong Yueh Chen, Fei-Ching Kuo, and Baowen Xu. A theoretical analysis of the risk evaluation formulas for spectrum-based fault localization. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 22(4):31, 2013.
56. Xiaoyuan Xie, Fei-Ching Kuo, Tsong Yueh Chen, Shin Yoo, and Mark Harman. Provably optimal and human-competitive results in sbse for spectrum based fault localisation. In *Search Based Software Engineering*, pages 224–238. Springer, 2013.
57. Justyna Zander, Ina Schieferdecker, and Pieter J Mosterman. *Model-based testing for embedded systems*. CRC press, 2011.
58. Xiangyu Zhang, Neelam Gupta, and Rajiv Gupta. Pruning dynamic slices with confidence. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '06*, pages 169–180, 2006.
59. Xiangyu Zhang, Rajiv Gupta, and Youtao Zhang. Precise dynamic slicing algorithms. In *Proceedings of the 25th International Conference on Software Engineering, ICSE '03*, pages 319–329, 2003.
60. Peter Zoetewij, Rui Abreu, Rob Golsteijn, and Arjan JC Van Gemund. Diagnosis of embedded software using program spectra. In *14th Annual IEEE International Conference and Workshops on the Engineering of Computer-Based Systems, 2007. ECBS'07.*, pages 213–220. IEEE, 2007.
61. Zhiqiang Zuo, Siau-Cheng Khoo, and Chengnian Sun. Efficient statistical debugging via hierarchical instrumentation. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, pages 457–460. ACM, 2014.