

# Simultaneous Debugging of Software Faults\*

Rui Abreu<sup>†1</sup>      Peter Zoetewej<sup>2</sup>  
Arjan J.C. van Gemund<sup>3</sup>

<sup>1</sup> Department of Informatics Engineering  
Faculty of Engineering  
University of Porto  
Portugal  
rui@computer.org

<sup>2</sup> IntelliMagic B.V.  
The Netherlands  
p.zoetewej@gmail.com

<sup>3</sup> Embedded Software Department  
Faculty of Electronics, Math, and CS  
Delft University of Technology  
The Netherlands  
a.j.c.vangemund@tudelft.nl

## Abstract

(Semi-)automated diagnosis of software faults can drastically increase debugging efficiency, improving reliability and time-to-market. Current automatic diagnosis techniques are predominantly of a statistical nature and, despite typical defect densities, do not explicitly consider multiple faults, as also demonstrated by the popularity of the single-fault benchmark set of programs. We present a reasoning approach, called Zoltar-M(multiple fault), that yields multiple-fault diagnoses, ranked in order of their probability. Although application of Zoltar-M to programs with many faults requires heuristics (trading-off completeness) to reduce the inherent computational complexity, theory as well as experiments on synthetic program models and multiple-fault program versions available from the software infrastructure repository (SIR) show that for multiple-fault programs this approach can outperform statistical techniques, notably spectrum-based fault localization (SFL). As a side-effect of this research, we present a new SFL variant, called Zoltar-S(single fault), that is optimal for single-fault programs, outperforming all other variants known to date.

---

\*This work has been carried out as part of the TRADER project under the responsibility of the Embedded Systems Institute. This project is partially supported by the Netherlands Ministry of Economic Affairs under the BSIK03021 program.

<sup>†</sup>Corresponding author.

**Keywords:** Software fault diagnosis, program spectra, statistical and reasoning approaches.

## 1 Introduction

Automatic software fault localization (also known as fault diagnosis) techniques aid developers to pinpoint the root cause of detected failures, thereby reducing the debugging effort. Two approaches can be distinguished:

- (1) the *spectrum-based fault localization* (SFL) approach, which correlates software component activity with program failures (a *statistical* approach) [3, 15, 20, 21, 26, 30], and
- (2) the *model-based diagnosis* or *debugging* (MBD) approach, which deduces component failure through *logic reasoning* [11, 13, 23, 28].

Because of its low computational complexity, SFL has gained large popularity. Although inherently not restricted to single faults, in most cases these statistical techniques are applied and evaluated in a single-fault context, as demonstrated by the benchmark set of programs widely used by researchers<sup>1</sup>, which is seeded with only 1 fault per program (version). In practice, however, the defect density of even small programs typically amounts to multiple faults. Although the root cause of a particular program failure need not constitute multiple faults that are acting *simultaneously*, many failures will be caused by *different* faults. Hence, the problem of *multiple-fault localization* (diagnosis) deserves detailed study.

Unlike SFL, MBD traditionally deals with multiple faults. However, apart from much higher computational complexity, the logic models that are used in the diagnostic inference are typically based on static program analysis. Consequently, they do not exploit execution behavior, which, in contrast, is the essence of the SFL approach. Combining the dynamic approach of SFL with the multiple-fault logic reasoning approach of MBD, in this paper, we present a multiple-fault reasoning approach that is based on the dynamic, spectrum-based observations of SFL. Additional reasons to study the merits of this approach are the following.

- Diagnoses are returned in terms of multiple faults, whereas statistical techniques return a one-dimensional list of single fault locations only. The information on fault multiplicity is attractive from parallel debugging point of view [19].
- Unlike statistical approaches, multiple-fault diagnoses only include valid candidates, and are asymptotically optimal with increasing test information [4].

---

<sup>1</sup><http://sir.unl.edu/>

- The ranking of the diagnoses is based on probability instead of similarity. This implies that the quality of a diagnosis can be expressed in terms of information entropy or any other metric that is based on probability theory [24].
- The reasoning approach naturally accommodates additional (model) information about component behavior, increasing diagnostic performance when more information about component behavior is available.

To illustrate the difference between multiple-fault and the statistical approach, consider a triple-fault (sub)program with faulty components  $c_1$ ,  $c_2$ , and  $c_3$ . Whereas under ideal testing circumstances a traditional SFL approach would produce multiple single-fault diagnoses (in terms of the component indices) like  $\{\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \dots\}$  (ordered in terms of statistical similarity), a multiple-fault approach would simply produce one single multiple-fault diagnosis  $\{\{1, 2, 3\}\}$ . Although the statistical similarity of the first three items in the former diagnosis would be highest, the latter, single diagnosis unambiguously reveals the actual triple fault.

Despite the above advantages, a reasoning approach is more costly than statistical approaches because an exponential number of multiple-fault candidates need to be processed instead of just the ( $M$ , being  $M$  the number of components in the system under analysis) single-fault candidates. In this paper, we compare our reasoning approach to several statistical approaches. Our study is based on random synthetic spectra, as well as on several benchmark programs, extended by us to accommodate multiple faults. More specifically, this paper makes the following 5 contributions.

- We introduce a multiple-fault diagnosis approach that originates from the model-based diagnosis area, but which is specifically adapted to the interaction dynamics of software. The approach is coined Zoltar-M (Zoltar for the name of our debugging tool set [18]<sup>2</sup>, M for multiple-fault).
- We show how our reasoning approach applies to single-fault programs, yielding a provably optimal SFL variant, called Zoltar-S (S for single-fault), as of yet unknown in literature.
- We introduce a general, multiple-fault, probabilistic program (spectrum) model, parameterized in terms of size, testing code coverage, and testing fault coverage, to theoretically study Zoltar-M, compared to statistical techniques such as Tarantula and Zoltar-S.
- We extend the traditional, single-fault benchmark set of programs (referred to as SIR-S) with a multiple-fault version (SIR-M), by combining the existing single-fault versions, to empirically evaluate debugging performance under realistic, multiple-fault conditions.

---

<sup>2</sup><http://www.fdir.org/zoltar>

- We investigate the ability of all techniques to deduce program fault multiplicity, which is aimed at providing a good estimate to guide parallel debugging, using an approach that substantially differs from [19].

To the best of our knowledge, this is the first paper to specifically address software multiple-fault localization using a spectrum-based, logic reasoning approach, yielding two new localization techniques Zoltar-S and Zoltar-M, implemented within our Zoltar SFL framework. Our experiments confirm that Zoltar-S is superior to all known similarity coefficients for the Siemens-S benchmark. More importantly however, our experiments for multiple-fault programs show that although for synthetic spectra Zoltar-M is outperformed by Zoltar-S, for our SIR-M experiments Zoltar-M outperforms all similarity coefficients known to date.

The paper is organized as follows. In the next section, we present the concepts and terminology used throughout the paper. In Section 3, our multiple-fault localization approach is described, as well as a derivation of the optimal similarity coefficient for single-fault programs. In Section 4, the approaches are theoretically evaluated, and in Section 5, real programs are used to assess the capabilities of the studied techniques for fault localization. Related work is discussed in Section 6. Preliminary results of Section 4 and Section 5 appeared in [2]. We conclude and discuss future work in Section 7.

## 2 Preliminaries

In this section, we introduce basic definitions as well as the traditional SFL approach. As defined in [6], in the remainder of this paper, we use the following terminology.

- A *failure* is an event that occurs when delivered service deviates from correct service.
- An *error* is the part of the total state of the system that may cause a failure.
- A *fault* is the cause of an error in the system.

To illustrate these concepts, consider the C function in Figure 1. It is meant to sort, using the bubble sort algorithm, a sequence of  $n$  rational numbers whose numerators and denominators are passed via parameters `num` and `den`, respectively. There is a fault (bug) in the swapping code of block 4: only the numerators of the rational numbers are swapped. The denominators are left in their original order.

A failure occurs when applying `RationalSort` yields anything other than a sorted version of its input. An error occurs after the code inside the conditional statement is executed, while `den[j]  $\neq$  den[j+1]`. Such errors can be temporary: if we apply `RationalSort` to the sequence  $\langle \frac{4}{1}, \frac{2}{2}, \frac{0}{1} \rangle$ , an error occurs after the first two numerators are swapped. However, this error is “canceled” by later

```

void RationalSort(int n, int *num, int *den)
{ /* block 1 */
  int i,j,temp;

  for ( i=n-1; i>=0; i-- ) {
    /* block 2 */
    for ( j=0; j<i; j++ ) {
      /* block 3 */
      if (RationalGT(num[j], den[j],
                    num[j+1], den[j+1])) {
        /* block 4 */
        temp = num[j];
        num[j] = num[j+1];
        num[j+1] = temp; } } }
}

```

Figure 1: A faulty C function for sorting rational numbers

swapping actions, and the sequence ends up being sorted correctly. Faults do not automatically lead to errors either: no error will occur if the input is already sorted, or if all denominators are equal.

The purpose of *diagnosis* is to locate the faults that are the root cause of detected errors. As such, error detection is a prerequisite for diagnosis. As a rudimentary form of error detection, failure detection can be used, but in software more powerful mechanisms are available, such as pointer checking, array bounds checking, deadlock detection, etc.

In a software context, faults are often called *bugs*, and diagnosis is part of *debugging*. Computer-aided techniques as the one we consider in this paper are known as *automated debugging*.

## 2.1 Basic Definitions

A program that is being diagnosed comprises a set of  $M$  components (statements in the context of this paper), which is executed using  $N$  test cases that either pass or fail. Program (component) activity is recorded in terms of program spectra [16]. This data is collected at run-time, and typically consists of a number of counters or flags for the different components of a program. In the context of this paper, we use the so-called hit spectra, which indicate whether a component was involved in a (test) run or not.

Both spectra and program pass/fail (test) information is input to SFL, as well as to our reasoning technique. The program spectra are expressed in terms of the  $N \times M$  *activity matrix*  $A$ . An element  $a_{ij}$  is equal to 1 if component  $j$  was observed to be *involved* in the execution of run  $i$ , and 0 otherwise. For  $j \leq M$ , the row  $A_{i*}$  indicates whether a component was executed in run  $i$ , whereas the column  $A_{*j}$  indicates in which runs component  $j$  was involved. The pass/fail information is stored in a vector  $e$ , the *error vector*, where  $e_i$  signifies whether

run  $i$  has *passed* ( $e_i = 0$ ) or *failed* ( $e_i = 1$ ). Note that the pair  $(A, e)$  is the only input to the techniques studied in this paper. From  $(A, e)$ , we can derive the probability  $r$  that a component is actually executed in a run (testing code coverage), and the probability  $g$  that a faulty component is actually exhibiting good behavior (testing fault coverage, also known as the “goodness” parameter  $g$  from MBD [9]).

Programs can have multiple faults, the number being denoted  $C$  (fault cardinality). A *diagnosis candidate* is expressed as the set of indices of those components whose *combined* faulty behavior is logically consistent with the observations  $A$  and therefore must be considered as a collective candidate. A *diagnosis* is the ordered set of diagnostic candidates  $D = \{d_1, \dots, d_k\}$ , all of which are an explanation consistent with observed program behavior ( $A$ ), ordered in probability of being the program’s actual multiple fault condition. An example multiple-fault diagnosis is the diagnosis  $\{d_1\} = \{\{1, 2, 3\}\}$  given in the Introduction. For brevity, we will often refer to diagnostic candidates as diagnoses as well, as it is clear from the context whether we refer to a single diagnosis candidate or to the entire diagnosis.

## 2.2 Traditional SFL

In SFL, one measures the similarity between the error vector  $e$  and the activity profile vector  $A_{*j}$  for each component  $j$ . This similarity is quantified by a *similarity coefficient*, expressed in terms of four counters  $n_{pq}(j)$  that count the number of positions in which  $A_{*j}$  and  $e$  contain respective values  $p$  and  $q$ , i.e., for  $p, q \in \{0, 1\}$ , we define

$$n_{pq}(j) = |\{i \mid n_{ij} = p \wedge e_i = q\}|$$

Two examples of well-known coefficients are

$$s_T = \frac{\frac{n_{11}(j)}{n_{11}(j)+n_{01}(j)}}{\frac{n_{11}(j)}{n_{11}(j)+n_{01}(j)} + \frac{n_{10}(j)}{n_{10}(j)+n_{00}(j)}}$$

as used by the Tarantula tool [20], and the Ochiai coefficient

$$s_O = \frac{n_{11}(j)}{\sqrt{(n_{11}(j) + n_{01}(j)) * (n_{11}(j) + n_{10}(j))}} \quad (1)$$

known from molecular biology, introduced in SFL in [3].

Under the assumption that a high similarity to the error vector indicates a high probability that the corresponding parts of the software cause the detected errors, the calculated similarity coefficients rank the parts of the program with respect to their likelihood of containing the faults. Algorithm 1 concisely describes the SFL approach to fault localization.

As an example, suppose we have a program with  $M = 7$  components, of which  $c_1$ ,  $c_2$ , and  $c_3$  are faulty, with  $A$  as given in Table 1. The table also includes the  $n_{pq}$  counts as well as the resulting similarity based on the Tarantula

---

**Algorithm 1** SFL Algorithm

---

**Require:** Activity matrix  $A$ , error vector  $e$ , number of runs  $N$ , number of components  $M$ , and similarity coefficient  $s$

**Ensure:** Diagnostic report  $D$

```
1  $D \leftarrow \emptyset$ 
2 for  $j = 0$  to  $M$  do
3    $n_{11}(j) \leftarrow 0$ 
4    $n_{10}(j) \leftarrow 0$ 
5    $n_{01}(j) \leftarrow 0$ 
6    $n_{00}(j) \leftarrow 0$ 
7    $S[j] \leftarrow 0$  ▷ Similarity  $s$  of component  $j$ 
8 end for
9 for  $i = 0$  to  $N$  do
10  for  $j = 0$  to  $M$  do
11    if  $a[i, j] = 1 \wedge e[i] = 1$  then
12       $n_{11}(j) \leftarrow n_{11}(j) + 1$ 
13    else if  $a[i, j] = 0 \wedge e[i] = 1$  then
14       $n_{01}(j) \leftarrow n_{01}(j) + 1$ 
15    else if  $a[i, j] = 1 \wedge e[i] = 0$  then
16       $n_{10}(j) \leftarrow n_{10}(j) + 1$ 
17    else if  $a[i, j] = 0 \wedge e[i] = 0$  then
18       $n_{00}(j) \leftarrow n_{00}(j) + 1$ 
19    end if
20  end for
21 end for
22 for  $j = 0$  to  $M$  do
23    $S[j] \leftarrow s(n_{11}(j), n_{10}(j), n_{01}(j), n_{00}(j))$ 
24 end for
25  $D \leftarrow \text{SORT}(S)$ 
26 return  $D$ 
```

---

and Ochiai coefficients. Assuming that a developer would follow the ranking produced by the techniques, Tarantula requires him/her to inspect more components in order to find a faulty one. The first faulty component ranked by Tarantula is at the 3<sup>rd</sup> place of the list, whereas with Ochiai it is already at the 2<sup>nd</sup> place. The results shows the sensitivity of Tarantula to components that are not involved in passed runs ( $n_{00}$ ), considering them likely to be the faulty one and not taking into account their involvement in failed runs (e.g.,  $c_4$  and  $c_7$ ). Ochiai, however, exonerates components based on their involvement in passed runs ( $n_{10}$ ), and absence in failed runs ( $n_{01}$ , for detailed comparison, see [3]).

As can be seen, both Tarantula and Ochiai fail to consider  $c_3$  as one of the most suspicious components. Besides,  $c_2$  and  $c_3$  can be considered as multiple fault because all failed runs can be explained either by  $c_2$  or  $c_3$  (but the two by themselves are not a valid explanation for all failures). In the next section, we present our technique that exploits this info and contains multiple-fault

	$c_1$	$c_2$	$c_3$	$c_4$	$c_5$	$c_6$	$c_7$	$e$
	1	0	1	0	1	0	0	0
	0	0	0	0	1	0	0	0
	0	0	1	0	0	1	0	0
	1	1	0	0	1	0	0	0
	1	0	1	0	0	0	0	0
	1	0	0	0	1	0	0	0
	1	0	1	0	0	1	0	0
	1	1	0	0	1	1	0	1
	1	1	0	1	0	1	1	1
	1	0	1	0	0	0	0	1
	1	1	0	1	1	1	1	1
	1	1	0	0	0	1	0	1
	1	0	1	0	0	1	1	1
$n_{11}(j)$	6	4	2	2	2	5	3	
$n_{10}(j)$	5	1	4	0	4	2	0	
$n_{01}(j)$	0	2	4	4	4	1	3	
$n_{00}(j)$	2	6	3	7	3	5	7	
$s_T$	0.58	0.82	0.37	1	0.37	0.74	1	
$s_O$	0.74	0.73	0.33	0.58	0.33	0.77	0.71	

Table 1: Observation Matrix Example  $A$

explanations in its ranking.

### 3 Multiple-Fault Localization

In this section, we present our multiple-fault localization approach Zoltar-M, which is based on reasoning as performed in model-based diagnosis, combined with (Bayesian) probability theory to compute the ranking of the candidates. The major difference with the statistical approach in Section 2.2 is

- that only a *subset* of components is considered (the so-called hitting set) in contrast to all components,
- all computed candidates logically explain the observed failures, and
- that the ranking is based on probability, rather than statistical similarity.

In the remainder of this section, specific details on the two main phases of our Zoltar-M approach are given (see Figure 2): (1) candidate generation and (2) candidate fault probability computation (ranking). In addition, as a by-product of Zoltar-M, we propose an optimal SFL variant for single faults.






Figure 2: Zoltar-M’s main phases

### 3.1 Hitting Set Computation

In model-based diagnosis, one derives a model of the program that, together with the observations of input-output behavior, determines a set of constraints from which diagnostic solutions consistent with this behavior are logically deduced. Unlike the MBD approaches such as presented in [22, 23], in our Zoltar-M approach we refrain from modeling the program in detail, but use  $A$  as the only, dynamic source of information, from which we derive the model and the input-output observations.

Each component  $c_j$  is modeled by the logic proposition

$$h_j \Rightarrow in.ok_j \Rightarrow out.ok_j \tag{2}$$

where  $h_j$  models the health state of  $c_j$  (*true* = healthy, *false* = defect). This so-called *weak* model specifies that a component produces correct output values (*out.ok* true) if (1) healthy ( $h$  is true), and (2) when provided with correct input (*in.ok* true). Note that this model still allows a component to produce correct data (the probability of which is measured by  $g$ ) even though  $h = false$ . Also note that a component can accept erroneous input data and still produce correct output. Finally, this approach allows the inclusion of additional component information as the number of propositions per component is not limited to the above, default model of nominal behavior.

Due to dynamic (data-dependent) control flow each run may involve different components. Consequently, rather than modeling the program by a static composition of component propositions (Eq. 2), we consider a dynamic model that is defined per program run (i.e.,  $A_{i*}$ ). In [4], it is shown that each failed run yields a conjunction of logical constraints (i.e., a sub-model) in terms of the components involved, which is known in MBD as a *conflict* [10]. For instance, a

failed run involving  $c_1$  and  $c_2$  generates the conflict  $\neg h_1 \vee \neg h_2$ , indicating that  $c_1$  and  $c_2$  cannot both be healthy.

The multiple-fault approach is based on compiling each failed run (row  $A_{i*}$ ) to a conflict, after which the diagnosis for  $A$  is derived by computing the hitting set [25] from all conflicts [4] (the hitting set algorithm essentially transforms logic products-of-sums into sums-of-products). For instance, the example observation matrix  $A$  in Table 1 generates the following 6 conflicts

$$\begin{aligned}
&(\neg h_1 \vee \neg h_2 \vee \neg h_5 \vee \neg h_6) \wedge \\
&(\neg h_1 \vee \neg h_2 \vee \neg h_4 \vee \neg h_6 \vee \neg h_7) \wedge \\
&(\neg h_1 \vee \neg h_3) \wedge \\
&(\neg h_1 \vee \neg h_2 \vee \neg h_4 \vee \neg h_5 \vee \neg h_6 \vee \neg h_7) \wedge \\
&(\neg h_1 \vee \neg h_2 \vee \neg h_6) \wedge \\
&(\neg h_1 \vee \neg h_3 \vee \neg h_6 \vee \neg h_7)
\end{aligned}$$

The (minimal) hitting set comprises one single-fault candidate  $\{1\}$ , and two double-fault candidates  $\{2, 3\}$  and  $\{3, 6\}$ . Note that the triple-fault candidate  $\{1, 2, 3\}$ , which equals the actual fault state, is subsumed by both  $\{1\}$  and  $\{2, 3\}$  and therefore does not appear in  $D$  (which is the *minimal* hitting set). The reason why, e.g.,  $\{1\}$  subsumes  $\{1, j\}$ ,  $j = 2, 3, \dots, M$  is that the weak component model (2) allows any faulty component  $j$  to exhibit correct behavior. Hence  $\{1, j\}$  is also a valid explanation. The hitting set can be directly observed from  $A$  by (multi-)column “chains” of ‘1’s from top to bottom formed by all failing rows of  $A$ . Note that this procedure only considers failed runs. Passed runs are considered later on when computing the probability of each diagnostic candidate.

The above example illustrates that the true diagnosis (candidate) can be preceded (or subsumed) by other, more probable candidates. However, for  $N \rightarrow \infty$  Zoltar-M produces the optimal result, where diagnoses such as  $\{1\}$  and  $\{2, 3\}$  eventually disappear, exposing the only correct diagnosis  $\{1, 2, 3\}$ . This can be seen through the following argument. Consider a  $C$ -fault program. While for small  $N$  the minimal hitting set will still contain many members (components) other than the  $C$  faulty components, by increasing  $N$  the probability that a non-faulty component will still be included steadily decreases [4]. Let  $f$  denote the probability of a run failing. As the probability that a run passes equals the probability that none of the  $C$  components cause a failure, which equals  $(1 - r \cdot (1 - g))^C$ , it follows that  $f = 1 - (1 - r \cdot (1 - g))^C$  [4]. For the hitting set analysis, only failing runs matter. For  $f \cdot N$  failing runs, the  $C$ -fault candidate is by definition within the set of candidates that “survive” those  $f \cdot N$  runs (whose chain is still unbroken). However, the probability that other components can be involved in a candidate is less than unity, which forms the basis of those candidates’ eventual elimination.

For example, the probability of a  $B$ -cardinality diagnosis ( $B < C$ ) competing with our  $C$ -cardinality solution equals the probability that at least one out of the  $B$  components is hit every time in a failing run. The latter probability equals  $b =$

$1 - (1 - r)^B$  (derivation similar to  $f$ ). Hence, for  $N$  runs, the probability of this competing diagnosis surviving in the final hitting set is of the order  $b^{f \cdot N}$ , which as  $b < 1$  negative-exponentially decreases to zero for large  $N$ . Note that the above analysis does not consider a particular probability computation regarding the ranking. It simply proves that, for large  $N$ , the diagnosis can only consist of the single surviving  $C$ -cardinality candidate ( $\{1, 2, 3\}$  in the earlier example). Our experiments confirm this optimality. There is one exception to the above argument. Components that are *always* executed (e.g., initialization code) will always appear as single-fault candidate, which is typically ranked higher than a genuine multiple-fault (see next section). As this applies to techniques based on spectral information this problem also occurs with statistical techniques.

### 3.2 Probability Computation

For each multiple-fault candidate, the probability of being the actual diagnosis depends on the extent to which that candidate explains all observations (pass or fail per run). Let  $\Pr(\{j\})$  denote the *a priori* probability that a component  $c_j$  is at fault. Although this value is typically dependent on code complexity, design, etc., we will simply assume  $\Pr(\{j\}) = p$  (we arbitrarily set  $p = 0.01$  in the context of this paper). Assuming components fail independently, and in absence of any observation, the prior probability a particular diagnosis  $d_k$  is correct is given by  $\Pr(d_k) = p^{|d_k|} \cdot (1 - p)^{M - |d_k|}$ . Similar to the incremental compilation of conflicts per run we compute the posterior probability for each candidate based on the pass/fail observation  $obs$  for each sequential run using Bayes' update rule according to

$$\Pr(d_k|obs) = \frac{\Pr(obs|d_k)}{\Pr(obs)} \cdot \Pr(d_k)$$

The denominator  $\Pr(obs)$  is a normalizing term that is identical for all  $d_k$  and thus needs not to be computed directly.  $\Pr(obs|d_k)$  is defined as

$$\Pr(obs|d_k) = \begin{cases} 0 & \text{if } d_k \text{ and } obs \text{ are inconsistent} \\ 1 & \text{if } d_k \text{ logically follows from } obs \\ \epsilon & \text{if neither holds} \end{cases}$$

In the context of model-based diagnosis, many policies exist for  $\epsilon$  (see [9]). In this paper, we define  $\epsilon$  as follows

$$\epsilon = \begin{cases} g(d_k)^\eta & \text{if run passed} \\ 1 - g(d_k)^\eta & \text{if run failed} \end{cases}$$

In this equation,  $\eta$  is the number of faulty components involved in the run (the rationale being that the more faulty components are involved, the more likely it is that the run will fail [4]), and  $g$  is estimated by

$$g(d_k) = \frac{n_{10}(d_k)}{n_{10}(d_k) + n_{11}(d_k)}$$

Technique	$D = \{d_1(s pr), \dots, d_k(s pr)\}$
Tarantula	$\{\{4\}(1), \{7\}(1), \{2\}(0.82), \{6\}(0.74), \{1\}(0.58), \{3\}(0.37), \{5\}(0.37)\}$
Ochiai	$\{\{6\}(0.77), \{1\}(0.74), \{2\}(0.73), \{7\}(0.71), \{4\}(0.58), \{3\}(0.33), \{5\}(0.33)\}$
Zoltar-M	$\{\{1\}(0.98), \{2, 3\}(0.99e^{-2}), \{3, 6\}(0.52e^{-2})\}$

Table 2: Diagnoses for example  $A$

where  $n_{1q}(d_k) = \sum_{i=1..N} [(\bigvee_{j \in d_k} a_{ij} = 1) \wedge e_i = q]$  is a generalization of the definition in Section 2.2 to support multiple fault explanations,  $q \in \{0, 1\}$ , and  $[\cdot]$  denotes Iverson’s operator [17] ( $[\mathbf{true}] = 1$ ,  $[\mathbf{false}] = 0$ ).

To illustrate the differences between the probabilistic approach as presented in this section and the statistical SFL approach (as explained in Section 2.2), again consider the example  $A$  in Table 1. The diagnostic report  $D$  for the different approaches are listed in Table 2. As can be seen, the top ranked candidate for both Tarantula and Ochiai is not one of the three faulty locations, whereas for Zoltar-M one of the faults, namely  $c_1$  would be immediately found. Furthermore, in contrast to Zoltar-M, which contains multiple faults explanations such as  $\{2, 3\}$ , Tarantula and Ochiai only rank single-fault explanations. To conclude, note that Zoltar-M *only* lists candidates that actually *explain* all observed failures.

While the inherent multiple-fault approach used in Zoltar-M is asymptotically optimal, the complexity of the underlying hitting set algorithm and subsequently having to manage a possibly exponential number of multiple-fault candidates (e.g., update their probability) is prohibitive for large  $C$  (and  $N, M$ ). Nevertheless, preliminary experiments with a statistically directed search technique (i.e., using statistical similarity to guide the search) indicates that the complexity of our current hitting set computation can be reduced by several orders of magnitude. In addition, hitting set completeness can be traded-off to further reduce time complexity (see [13] for a greedy stochastic search addressing this issue).

### 3.3 Single-fault Case

In this section, we show how our above reasoning approach can be used to derive an optimal similarity coefficient for *single-fault* programs.

In the single-fault case (such as the SIR benchmark set of programs), we know that all failures relate to only one fault, which, by definition, is included in the minimal hitting set. Hence, any coefficient approach should consider the minimal hitting set only (i.e., only those  $c_j$  which consistently occur in failing runs). This implies that the optimal approach is to select only the failing runs

and compute the similarity coefficient. Since for these components by definition  $n_{01} = 0$ , one only needs to consider  $n_{11}$  and  $n_{10}$ . This, in turn, implies that the ranking is only determined by the exonerating term  $n_{10}$ . In summary, once we only consider the components included in the hitting set, any of the coefficients that includes  $n_{10}$  in the denominator will produce the same, optimal ranking. Experiments using this “hitting set filter” combined with a simple similarity coefficient such as Tarantula indeed confirm that this approach leads to the best performance [27]. For instance,

$$\text{filter} = \begin{cases} s_T & \text{if } n_{01} \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the above filter is only optimal for programs that have only 1 fault as applying this filter to any multiple-fault program would be overly restrictive. It would fail to detect faults that are not always involved in failed runs. For example, the diagnosis for  $A$  in Table 1 when using the filtering approach would yield  $D = \{\{1\}\}$ , entirely ignoring two of the three faults. Hence, instead of considering a single-fault hitting set filter, we modify this approach in order to also allow application to multiple-fault programs. Taking the Ochiai coefficient as (best) starting point (for  $\kappa = 1$ , Eq. 3 follows from Eq. 1 by squaring, and factoring out  $n_{11}(j)$ , none of which changes the ranking) and applying the above filtering approach, we derive the following similarity coefficient [4], coined Zoltar-S, according to

$$s_{Z-S} = \frac{n_{11}(j)}{n_{11}(j) + n_{10}(j) + n_{01}(j) + \kappa \cdot \frac{n_{01}(j) \cdot n_{10}(j)}{n_{11}(j)}} \quad (3)$$

where  $\kappa > 0$  is a constant factor that exonerates a component  $c_j$  that was either seldom executed in failed runs or often in passed runs. We empirically verified that the higher the value of  $\kappa$ , the more identical the diagnosis becomes with the one obtained by the hitting set filter [27]. In the context of this paper, we limit  $\kappa$  to 10,000 to avoid round-off errors.

In [5] it has been proven that, for single-fault programs, given the available input data  $(A, e)$ , the diagnostic ranking produced by our reasoning technique is theoretically optimal. As such theoretical optimality applies to  $C = 1$  (i.e., for single-fault programs), we have adapted the result in [5] to a similarity coefficient in terms of  $\kappa$ . The following theorem proves that for  $\kappa \rightarrow \infty$  the Zoltar-S similarity coefficient has exactly the same behavior.

**Theorem** *For single-fault programs, given the available input data  $(A, e)$ , the diagnostic ranking produced by ZOLTAR-S is theoretically optimal.*

**Proof** Let  $c_f$  be the faulty component and  $c_p$  a (representative) non-faulty component. For single faults,  $n_{01}(f) = 0$ ,  $n_{11}(f) = N_F$  (where  $N_F$  is the number of failed runs), and  $n_{01}(p) \geq 0$ . Therefore,  $\frac{n_{01}(f) \cdot n_{10}(f)}{n_{11}(f)} = 0$  and  $\frac{n_{01}(p) \cdot n_{10}(p)}{n_{11}(p)} \geq 0$ . Consequently, for non-faulty components, the following holds

Program	Faulty Versions	#components ( $M$ )	#runs ( $N$ )	Description
<code>print_tokens</code>	7	539	4,130	Lexical Analyzer
<code>print_tokens2</code>	10	489	4,115	Lexical Analyzer
<code>replace</code>	32	507	5,542	Pattern Recognition
<code>schedule</code>	9	397	2,650	Priority Scheduler
<code>schedule2</code>	10	299	2,710	Priority Scheduler
<code>tcas</code>	41	174	1,608	Altitude Separation
<code>tot_info</code>	23	398	1,052	Information Measure

Table 3: The Siemens benchmark set

$$\lim_{\kappa \rightarrow \infty} s_{Z-S}(p) \simeq 0 \quad (4)$$

$$\lim_{\kappa \rightarrow \infty} s_{Z-S}(f) = \frac{n_{11}(f)}{n_{11}(f) + n_{10}(f)} = \frac{N_F}{N_F + n_{10}(f)} > 0 \quad (5)$$

Hence, the higher  $\kappa$  is, the more is the exoneration factor for non-faulty, and as such, the faulty one will rank high in the diagnostic ranking.  $\square$

To evaluate the diagnostic capabilities of Zoltar-S in comparison with other techniques, the Siemens benchmark set is used. This well-known benchmark is composed of 7 programs (see Table 3; for detailed info, visit <http://sir.unl.edu>). In total, the Siemens benchmark set of programs provides 132 faulty programs. However, as no failures are observed in two of these programs, namely version 9 of `schedule2` and version 32 of `replace`, they are discarded. Besides, we also discard versions 4 and 6 of `print_tokens` because the faults are not in the program itself but in a header file. In summary, we discarded 4 versions out of 132 provided by the suite, using 128 versions in our experiments. To collect the program spectra, the `Zoltar` toolset [18] was used. For compatibility with previous work in (single-) fault localization, we use the effort/score metric [3, 26], which is the percentage of statements that need to be inspected to find the fault - in other words, the rank position of the faulty statement divided by the total number of statements. Note that some techniques such as in [21, 26] do not rank all statements in the code, and their rankings are therefore based on the program dependence graph (PDG) of the program.

Figure 3 plots the percentage of located faults in terms of debugging effort [3]. Apart from the coefficients studied for SFL, the following techniques are also plotted: Intersection and Union [26], Delta Debugging (DD) [30], Nearest Neighbor (NN) [26], and Sober [21], which are among the best SFL techniques (detailed discussion in Section 6). As Sober is publicly available, we run it in our own environment. The values for the other techniques are, however, directly cited from their respective papers.

From Figure 3, we conclude that Zoltar-S and the filter version are consistently the best performing techniques (note that in the single-fault context Zoltar-M simply reduces to Zoltar-S), finding 60% of the faults by examining

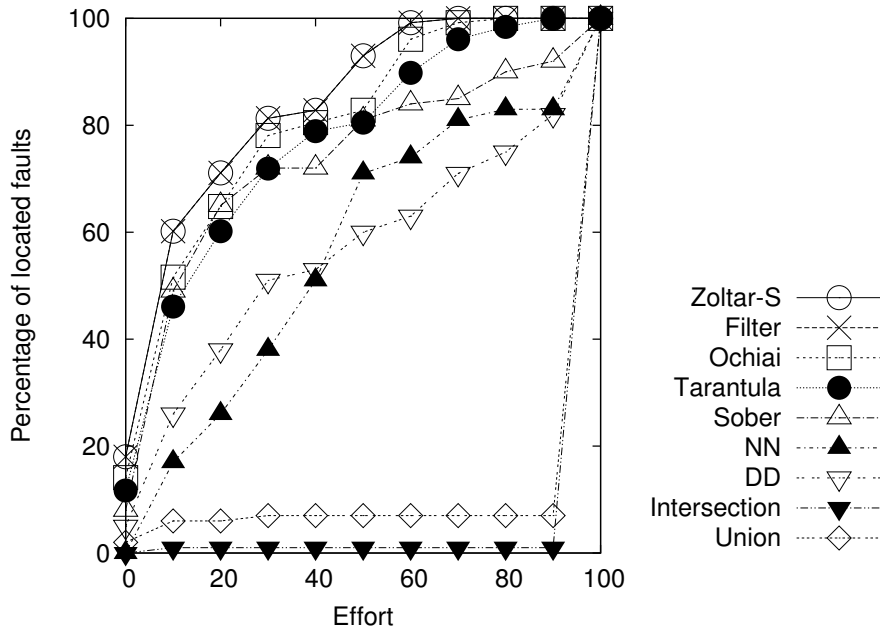


Figure 3: Effectiveness Comparison ( $C = 1$ )

less than 10% of the source code. For the same effort, using Ochiai would lead a developer to find 52% of the faulty versions and with Tarantula only 46% would be found. The Zoltar-S approach is followed by Ochiai, which outperforms Sober and Tarantula, which as concluded in [21], yield similar performance. Finally, the other techniques plotted are clearly outperformed by the spectrum-based techniques.

## 4 Theoretical Evaluation

In order to gain understanding of the effects of the various parameters on the diagnostic performance of the different approaches, we use a simple, probabilistic model of program behavior that is directly based on  $C$ ,  $N$ ,  $M$ ,  $r$ , and  $g$ . Without loss of generality we model the first  $C$  of the  $M$  components to be at fault. For each run, every component has probability  $r$  to be involved in that run. If a selected component is faulty, the probability of exhibiting nominal (“good”) behavior equals  $g$ . When either of the  $C$  components fails, the run will fail. We study the performance of Zoltar-M in comparison to Tarantula, Ochiai, and Zoltar-S for observation matrices that are randomly generated according to the above model.

$P$	1	2	3	4	5	6	...
Tarantula $W/I$	14/0	29/0	29/1	43/1	43/2	43/3	...
Zoltar-M $W/I$	0/1	0/2	0/3	14/3	–	–	...

Table 4: Wasted effort for different developers  $P$

## 4.1 Performance Metrics

Before evaluating the results, we first present our performance metric. As one of the motivations of our multiple-fault approach is the exposure of fault multiplicity (parallel debugging) we refrain from reusing established metrics such as the diagnostic quality [3] or score [26] but evaluate the amount of *wasted debugging effort*  $W$  as a function of the number of parallel debuggers [19], denoted by  $P$ , which more clearly indicates practical debugging parallelism. The wasted debugging effort is computed as follows. From the diagnosis (obtained with either a statistical or reasoning approach), the first  $P$  candidates are examined (debugged) in parallel [19]. Actual faults are assumed to be properly debugged, after which the program is retested. Based on the retest a new diagnosis is obtained (excluding the repaired components, but including the still uncovered faults that may have considerably moved up in the ranking). This  $P$ -parallel process continues until in the last iteration the program retests ok (i.e., all faults have been found).  $W$  measures the percentage of non-faulty components that were debugged in the above process. For  $P = 1$ , the above procedure reduces to a standard sequential debugging process. For instance, consider the diagnostic reports yielded by Tarantula and Zoltar-M (as in Table 2) for Example  $A$  in Table 1. Table 4 shows the performance profile for these two techniques ( $I$  stands for the number of bugs found in the first debugging iteration). As can be seen, Tarantula would need more developers in order to get a bug-free program in one iteration (6 developers against 3 for Zoltar-M). Furthermore, for this example, the wasted effort is consistently higher for Tarantula: with Zoltar-M, 3 developers would eliminate all bugs from the program at the cost of 0% wasted effort, whereas with Tarantula 6 developers would be needed at a cost of 43%. Note that there is no point in putting more than 4 developers to work as the Zoltar-M diagnosis contains only 4 different components.

Another reason not to adopt the aforementioned score metric [26] is that in our synthetic model we do not have program dependence graph information. Furthermore, the choice to exclude the actual faults from the debugging effort (i.e., instead of counting them as effort) is to make our performance metric independent of the number of faults  $C$ .

## 4.2 Experimental Results

In our first experiment, we focus on the effect of  $C$ ,  $N$ ,  $M$ ,  $r$ , and  $g$  on  $W$ . Consequently, we choose  $P = 1$ . We have varied  $M$  between 10 and 30 and after verifying that this does not change our conclusions [4], we choose  $M = 20$



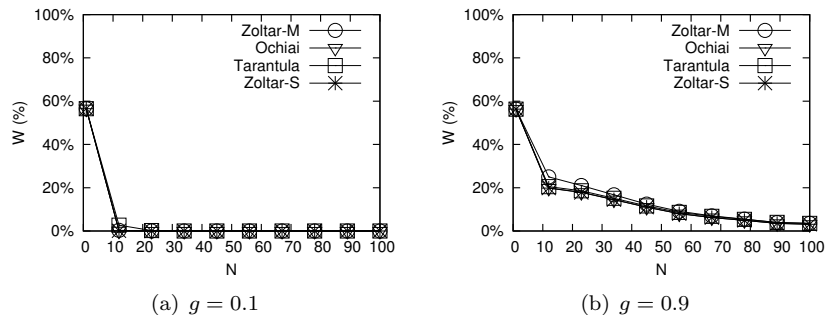


Figure 4: Wasted effort  $W$  for  $C = 1$

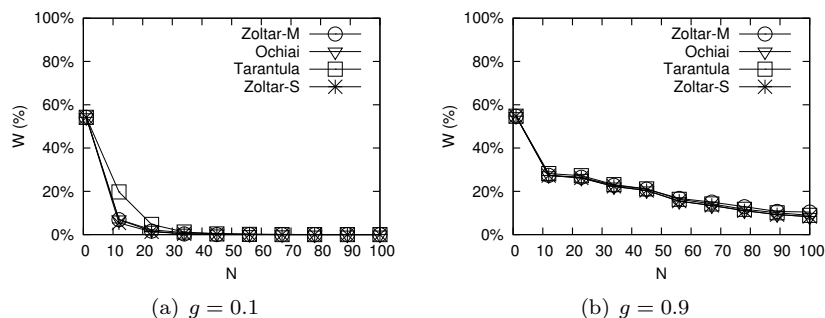


Figure 5: Wasted effort  $W$  for  $C = 2$

for the plots in the paper. Similarly, we also varied  $r$  between  $r = 0.4$  and  $r = 0.6$ , and as there are no significant differences we only include the plots for  $r = 0.6$ , which is roughly the same as the values measured for the Siemens set.

Figures 4, 5, and 6 plot  $W$  versus  $N$  for  $C = 1$ ,  $C = 2$ , and  $C = 5$ , respectively. We have also applied the technique for matrices with  $C = 8$  and the conclusions are essentially the same. Each measurement represents an average over 1,000 sample matrices. The plots show that for small  $N$  all techniques start with equal  $W$  (for  $N = 1$  it follows that  $W = (M - C) \cdot r / M$  [4]), while for sufficiently large  $N$  all techniques produce an optimal diagnosis. The plots clearly show that all techniques yield an optimal diagnosis for sufficiently large  $N$ . This happens earlier for small  $C$  and  $g$ . In the single-fault case, there is hardly any difference in the various techniques. For a small value of  $g$ , almost each run that involves the faulty component yields a failure, already producing near-perfect diagnoses for only small  $N$ . For a large value of  $g$  (which is more realistic, the Siemens set exhibits  $g$  values ranging from 79% (`tot_info`) to 99% (`tcas`)), the fraction of failing runs dramatically decreases (cf.  $f$  in Section 3.1). Consequently, a much larger number of runs is required to obtain a good diagnosis. For  $C = 5$ , we see the same trend, albeit that convergence to good diagnosis is much slower, especially for high  $g$ . This is due to the

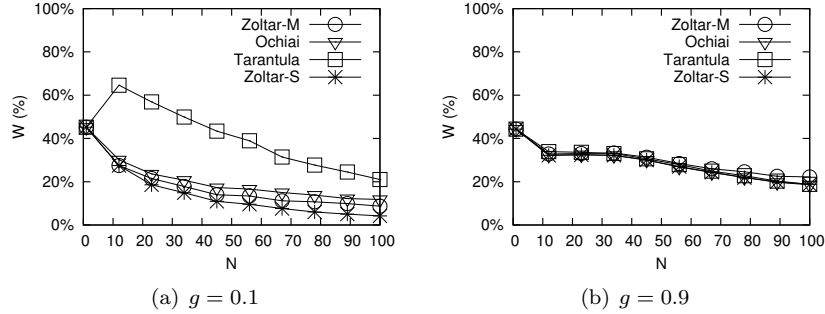


Figure 6: Wasted effort  $W$  for  $C = 5$

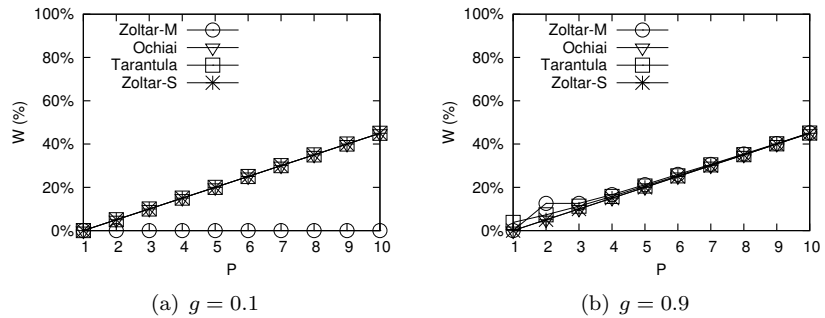


Figure 7:  $W$  vs.  $P$  for  $C = 1$

(combinatorial) fact that the number of “competitor” candidates of cardinality  $B \leq C$  (see Section 3.1) greatly increases with  $C$  (and  $M$  [4]). The main conclusion is that all techniques are very similar for the synthetic matrices and no technique clearly outperforms the others. For  $C = 5$ , we conclude that for small  $g$  Zoltar-S outperforms all other techniques, Zoltar-M is the second-best technique, and Tarantula is the worst performing technique. Besides, for small  $g$  and  $N$ , the Tarantula technique has very poor performance because some of the non-faulty components are only touched in failed runs, hence sharing the first position of the ranking and degrading the diagnosis. This is also the reason why the wasted effort first increases and only then starts to decrease (more passed runs are needed for non-faulty components to be exonerated). The fact that Zoltar-S outperforms Zoltar-M comes from the fact that for the synthetic matrices there are not that many non-faulty components involved in all failed runs, and therefore Zoltar-S manages to rank the faulty components on top. For more realistic cases ( $g = 0.9$ ), all techniques perform equally (poor), and much higher  $N$  is required to produce high diagnostic quality.

In Figures 7, 8, and 9, we measure  $W$  for all approaches as function of  $P$  to study inherent debugging parallelism for  $C = \{1, 2, 5\}$ . For these plots, we set  $N = 500$  to ensure that each technique has reached acceptable diagnostic

quality. For  $g = 0.1$ ,  $W$  starts a linear increase after  $P = C$  (the “knee”), which indicates that all  $C$  faults are indeed at or near the top of the ranking (the bump at  $P = 4$  is due to integer division effects). Except for Ochiai, both Zoltar-S and Tarantula yield similar performance as Zoltar-M. For  $C = 1$  and  $g = 0.1$ , Zoltar-M has zero wasted effort throughout. This occurs because, for  $N = 500$ , the diagnosis only contains the faulty statement (perfect diagnosis), revealing that there is no point in having more than 1 developer debugging the program.

While the above results show to what extent debugging can be efficiently parallelized, in practice information on  $C$  is, of course, not available. In the following, we evaluate the added value of multiple-fault diagnosis in estimating the number of debuggers  $P$  that can be efficiently deployed in parallel. The plots in Figure 10 show the distribution of the probability (Zoltar-M) or similarity (Zoltar-S, Ochiai, Tarantula) versus the ranking position. For multiple-fault diagnoses, each member index is counted as separate position. For cases where the diagnoses are near-perfect ( $g = 0.1$ ), the Zoltar-M distribution clearly exhibits the added information on the program’s fault cardinality  $C$  (corresponding to the “knee” in the previous plots), whereas the statistical techniques fail to produce any information on  $C$  whatsoever (although the Zoltar-S ranking distribution has more dynamics). For a high value of  $g$ , this relative advantage becomes less as diagnostic quality degrades. Note that this can be remedied by further increasing  $N$ .

## 5 Empirical Evaluation

Whereas the synthetic observation matrices used in the previous section are populated using a uniform distribution, this is not the case with observation matrices for the behavior of actual programs (different *spectral distribution*). Therefore, in this section we will evaluate the same diagnosis techniques on the SIR-S set, which provides the programs introduced in Section 3.3 extended with the real-world, large programs `space`, `gzip`, and `sed` (see Table 5). In addition, we also evaluated our techniques in the extended the benchmark set of programs




Figure 8:  $W$  vs.  $P$  for  $C = 2$

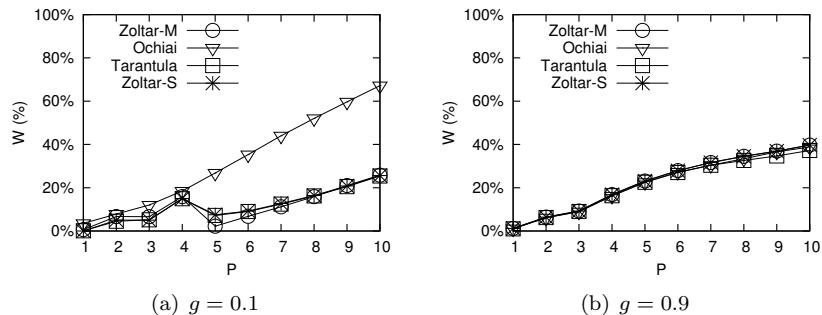


Figure 9:  $W$  vs.  $P$  for  $C = 5$

Program	Faulty Versions	$M$	$N$	Description
space	38	9,564	150	ADL Interpreter
gzip-1.3	7	5,680	210	Data compression
sed-4.1.5	6	14,427	370	Textual manipulator

Table 5: SIR’s real-world programs

to accommodate multiple-fault (SIR-M).

## 5.1 Experimental Setup

The SIR-M set extends the Siemens-S set with program versions that combine several faults from the latter set. The faults can be selectively activated via conditional compilation. The selections of faults that are available in the SIR-M set are limited by

- (1) their nature (e.g., a fault in non-executable code, which is not handled by our techniques would effectively reduce a  $C$ -fault diagnosis to a  $(C - 1)$ -fault diagnosis),
- (2) the number of failed runs (we only consider faults that yield at least one failed test case),
- (3) their locations (several faults in SIR-S have the same statement location), and
- (4) the number of lines of code involved (we only consider faults that can be attributed to a single line).

As mentioned in Section 3.3, the `Zoltar` toolset [18] is used to obtain code coverage information for each of the test cases supplied with the programs in the benchmark set of programs. The error vector in the last column is constructed

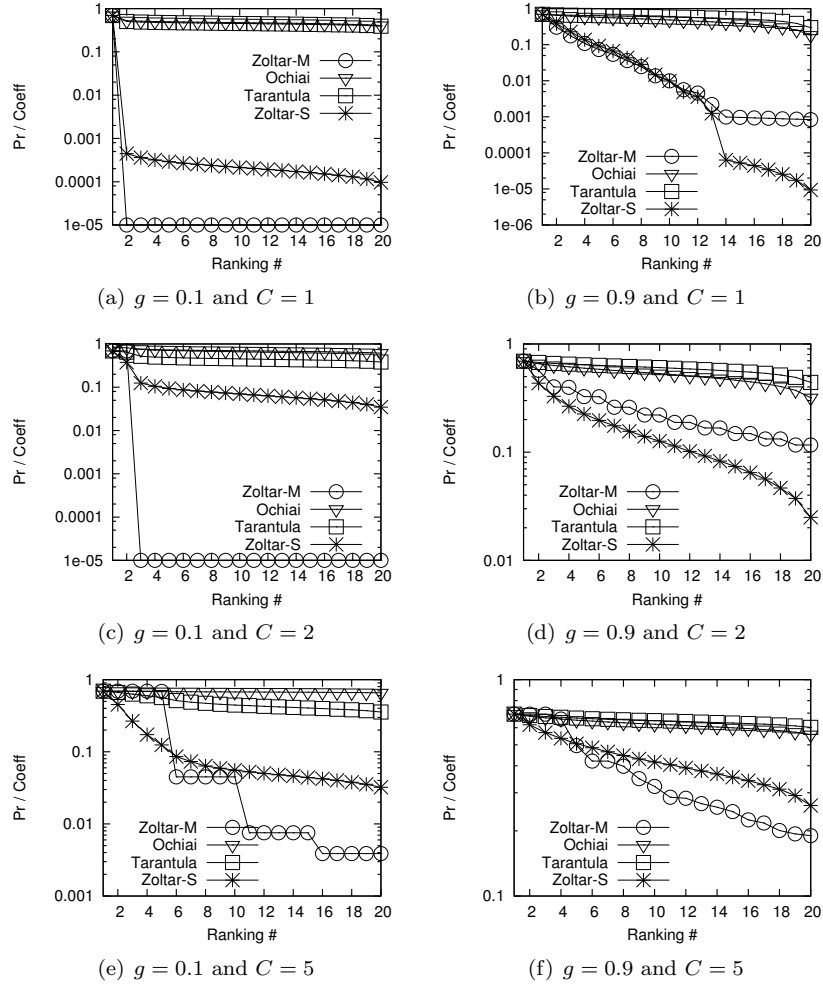


Figure 10: Probability/Similarity distribution

by comparing the output of a faulty version of a program with that of the correct version of the program, on a given test case.

For the resulting set of program spectra, Zoltar supports various diagnosis techniques, including Zoltar-M, and the Tarantula, Ochiai, and Zoltar-S coefficients. In the case of Zoltar-M, the presence of duplicate columns, following from the block structure of a program, is exploited in the hitting-set calculation by grouping all identical columns, while maintaining the set of components (lines of code) that they correspond to. This way, larger numbers of components can be handled than in the case of synthetic observation matrices.

## 5.2 Experimental Results

In Figure 11, we show  $W$  versus  $P$  for `tcas` and `replace`, two representative programs, when seeded with  $C = 1$ ,  $C = 2$ , and  $C = 3$  faults, respectively. Although the minimal hitting set computation is known to be rather expensive, we have used a low-cost, approximate technique dubbed STACCATO [1], which makes Zoltar scale to large, real-world programs [5]. We have also repeated the experiment up to  $C = 5$  (up to  $C = 10$  for `tcas`), but the graphs are similar to those for  $C = 2$  and  $C = 3$  of the representative programs, with the performance of Zoltar-S approaching that of the other methods as  $C$  increases.


The plotted  $W$  values are averaged over several different program versions: in case of the plots for  $C = 1$ , these are all faulty versions in the SIR-S set that can be attributed to a single line of executable code (30 for `tcas`, and 25 for `replace`). In the case of the  $C = 2$  and  $C = 3$  plots, these are 40 – 100 randomly selected combinations of faults. The plateau reached by Zoltar-M for `tcas` at high  $P$  values is caused by the limited size (ambiguity) of Zoltar-M’s diagnosis, removing the need to have them inspected by additional developers.

## 5.3 Evaluation


Figure 11(a) and 11(d) confirm the observation of Section 4 that for single faults, Zoltar-S is optimal. Although at the end of Section 3.3 we noted that in a single-fault context, Zoltar-M reduces to Zoltar-S, here Zoltar-M runs in multiple-fault mode, and the presence of cardinality  $C' = 3$  diagnoses in the hitting set, as explained above, is the cause for the small differences between these two techniques.

Contrary to what we observed in Section 4.2, where Zoltar-S is among the best performing methods for multiple-faults, in Figures 11(b), 11(c), 11(e), and 11(f), Zoltar-S performs worst. This is caused by many non-faulty components that are active in all failed runs. Having  $n_{01}(j) = 0$  in Eq. (3), such components fail to be exonerated via the term  $\kappa \cdot \frac{n_{01}(j) \cdot n_{10}(j)}{n_{11}(j)}$ , and will therefore rank high, leading to a lower quality diagnosis. While in the synthetic observation matrices it is unlikely that a component is active in all failed runs, this is quite common in software (e.g., statements that are always executed).


As shown in Figures 11(b), 11(c), 11(e), and 11(f), for  $C = 2$  and  $C = 3$ , Zoltar-M generally outperforms the statistical techniques, but for `tcas`, its performance is quite close to that of the SFL approaches using the Ochiai and Tarantula coefficients. This can be attributed to the following two related effects. First, the `tcas` faults that are available for making multiple-fault versions have a higher goodness factor ( $g = 0.95$ ) than those available for `replace` ( $g = 0.86$ ), making the diagnosis problems for the multiple-fault `tcas` versions inherently more difficult. The rationale is that for faults whose observation matrix inherently does not permit a good diagnosis (e.g., because the activity of a non-faulty component accidentally coincides with the occurrence of failures), all appropriate techniques will yield an equally bad diagnosis on average. Referring back to the discussion at the end of Section 4.2, the high values for  $g$  (common




(a) `tcas`,  $C = 1$




(b) `tcas`,  $C = 2$




(c) `tcas`,  $C = 3$



(d) `replace`,  $C = 1$



(e) `replace`,  $C = 2$



(f) `replace`,  $C = 3$

Figure 11: Wasted effort for 1 - 64 developers on representative programs of the SIR-M benchmark set

to all programs in the used benchmark set of programs) also explain why on average, no technique achieves optimal diagnostic quality on the Siemens-S and Siemens-M faults, and the consequent absence of a “knee” in the  $P - W$  graph of Zoltar-M.

The second effect that contributes to the difference in the plots for `tcas` and `replace` is that the variations in control flow in the former program are extremely limited, while essentially, this is what the diagnosis methods are based on. As an illustration, for the correct version of `tcas`, the observation matrix that follows from the 1,608 test cases that accompany the program contains many duplicate rows and columns: the number of unique rows (spectra) and columns (component behavior profiles) are 8 and 14, respectively. In comparison, the 5,542 test cases of `replace` lead to 2,023 different spectra, and 91 different behavior profiles, providing much more information to base the diagnosis on.

The latter observation confirms our expectation that the effectiveness of automated diagnosis techniques generally improves with program size. As an illustration, near-zero wasted effort is implied by the experiments with SFL on a 0.5 MLOC industrial software product reported in [33]. In summary, `tcas` is too simple, and Figures 11(e) and 11(f) can be expected to be the more representative of multiple-fault debugging in a realistic development environment. From this, we conclude that Zoltar-M can be expected to yield a significant improvement of debugging efficiency over the statistical methods in the multiple-fault case.

## 5.4 Time/Space Complexity

In this section, we report on the time/space complexity of Zoltar-M, compared to other fault localization techniques. We measure the time efficiency by conducting our experiments on a 2.3 GHz Intel Pentium-6 PC with 4 GB of memory. As most fault localization techniques have been evaluated in the context of single faults, in order to allow us to compare our fault localization approach to related work we limit ourselves to the original, single-fault Siemens benchmark set. We obtained timings for probabilistic dependence graph (PPDG) and Delta Debugging (DD) from published results [7, 30].

Table 6 summarizes the results of the study. The columns show the programs, the average CPU time (in seconds) of Zoltar-M, traditional SFL (Tarantula/Ochiai), PPDG, and DD, respectively. As expected, the less expensive techniques are the statistics-based techniques Tarantula and Ochiai. At the other extreme are PPDG and DD. Zoltar-M costs much less than PPDG and DD.

With respect to space complexity, statistical techniques need to store the counters  $(n_{11}, n_{10}, n_{01}, n_{00})$  for the similarity computation for all  $M$  components. Hence, the space complexity is  $O(M)$ . Zoltar-M also stores similar counters but per diagnosis candidate. Assuming that  $|D|$  scales with  $M$ , these approaches have  $O(M)$  space complexity.



Program	Zoltar-M	Tarantula/Ochiai	PPDG	DD
print_tokens	4.2	0.37	846.7	2590.1
print_tokens2	4.7	0.38	243.7	6556.5
replace	6.2	0.51	335.4	3588.9
schedule	2.5	0.24	77.3	1909.3
schedule2	2.5	0.25	199.5	7741.2
tcas	1.4	0.09	1.7	184.8
tot_info	1.2	0.08	97.7	521.4
space	7.4	0.15	N/A	N/A
gzip	6.2	0.19	N/A	N/A
sed	9.7	0.36	N/A	N/A

Table 6: Diagnosis cost for the single-fault subject programs (time in seconds)

## 5.5 Threats to Validity

Although the empirical study presented in this section provides evidence of the potential usefulness of the simultaneous bug fixing technique, there are threats to the validity of the empirical results that should be taken into account when interpreting the results.

Using only small to medium-sized C programs is a threat to external validity. Although, we believe the results will be identical, we cannot claim that the results generalize (large-sized programs, other programming languages). Yet another threat to external validity is the way multiple fault versions are built. When combining faults we assume an or-model (cf. the  $\epsilon$ -policy). So, we ignore interference between faults (faults can mask other faults).

Quantifying the debugging effort is extremely difficult because developers can recognize some components do not need to be inspected. Besides, we assume developers will inspect components following the ranking given by the techniques and that may not be entirely true (a developer could try to follow a “smell” following the control-data relationship). Finally, we also assume that a developer is able to identify the faulty component once it inspects it.

Deploying several developers to fix the multiple bugs in a system may be more error-prone than a single developer fixing all bugs. The experiments reported do not address this problem. Further studies are needed to investigate this issue.

## 6 Related Work

As mentioned in the introduction, automated debugging techniques can be distinguished into statistical and logic reasoning approaches that use program models.

In model-based reasoning to automatic software debugging (MBSD), the program model is typically generated from the source code using static analy-

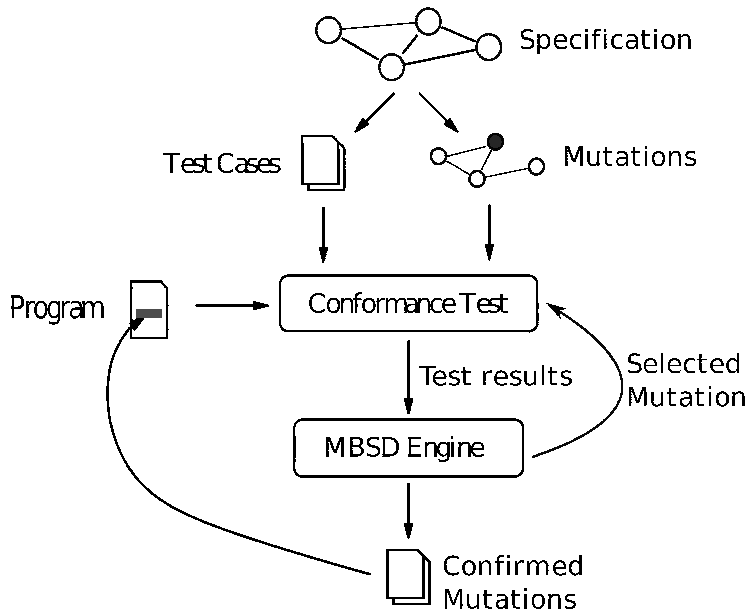


Figure 12: Using Conceptual Models to Enhance Model-based Reasoning

sis, as opposed to the traditional application of model based diagnosis where the model is obtained from a formal specification of the (physical) system [25]. An overview of different techniques to generate program models is given in [23]. The authors conclude that models generated by means of abstract interpretation [22] are the most accurate for debugging, while not suffering from the computational complexity inherent to more precise analysis techniques [23]. Recently, model-based techniques have also been proposed to isolate specific faults stemming from incorrect implementation of high-level conceptual models [29], where mutations are applied to state machine models to detect conceptual errors (see Figure 12), such as incorrect control flow and missing or additional features found in the implementation. Model-based approaches also include the work of Wotawa, Stumptner, and Mayer [28]. Other approaches that fit into this category include `explain` [14] and  $\Delta$ -slicing [14], which are based on comparing execution traces of correct and failed runs using model checkers. Model-based test generation [12] from abstract specifications of systems employs a similar idea where possible faults manifested as differences in abstract state machines are analyzed to generate tests. Although model-based diagnosis inherently considers multiple-faults, thus far the above software debugging approaches only consider single faults. Apart from this, our approach differs in the fact that we use program spectra as dynamic information on component activity, which allows us to exploit execution behavior, unlike static approaches. Furthermore, our approach does not rely on the approximations required by static techniques (i.e., incompleteness).

Statistical approaches are very attractive from complexity-point of view. Well-known examples are the Tarantula tool [20], the Nearest Neighbor technique [26], the Sober tool [21], CP [31], Holmes [8], and the Ochiai coefficient [3]. Although differing in the way they derive the statistical fault ranking, all techniques are based on measuring program spectra. Examples of other techniques that do not require extra knowledge of the program under analysis are the Delta Debugging technique [30] and the dynamic program slicing technique [15].

Essentially all of the above work has mainly been studied in the context of single faults, except for recent work by Jones, Bowring, and Harrold [19] and Zheng, Jordan, Liblit, Naik, and Aiken [32] which are motivated by the obvious advantages of parallel debugging with respect to development time reduction (particularly the work in [19]). They use clustering techniques to identify traces (rows in  $A$ ) that refer to the same fault, after which a single-fault technique is applied to each cluster of rows. While our work has the same motivation, our approach is based on logic reasoning instead of clustering. Although both introduce an increase of computational complexity, compared to the aforementioned statistical approaches, our hitting set analysis approach is asymptotically optimal, while in the clustering approach there is a possibility that multiple developers will still be effectively fixing the same bug. As their parallel debugging approach has only been evaluated in a restricted empirical context, our results, e.g., for the Siemens programs, cannot yet be compared.

## 7 Conclusions and Future Work

In this paper, we have presented a multiple-fault localization technique, Zoltar-M, which is based on the dynamic, spectrum-based measurement approach from statistical fault localization methods, combined with a logic (and probabilistic) reasoning approach from model-based diagnosis, inspired by previous work in both separate disciplines [3, 13]. We have compared the performance of Zoltar-M with Tarantula and Ochiai, which are amongst the best known statistical SFL approaches, as well as a new statistical SFL technique, coined Zoltar-S, derived by us as a by-product of our reasoning approach, and shown to be optimal for single-fault programs ( $C = 1$ ).

Our synthetic experiments show that both the reasoning and statistical approaches have the same general properties with respect to the influence of the parameters we introduced, viz, number of components  $M$ , number of test cases  $N$ , testing code coverage  $r$ , testing fault coverage  $g$ , and fault cardinality  $C$ . For a low value of  $g$ , both approaches yield near-perfect quality for relatively small  $N$ , while for high  $g$  (typical for many components in practice) a much larger  $N$  is required for good diagnosis. In most cases, it is Zoltar-S that outperforms Zoltar-M, which for  $C > 1$  is due to the fact that all components are involved in different runs with the same probability, making it easy for Zoltar-S to pinpoint the faulty ones. Despite these small differences, Zoltar-M's ranking probability distribution clearly provides information on the program's potential debugging parallelism while statistical techniques fail to provide any information.

Our results on two multiple-fault programs of our newly created SIR-M benchmark suggest that for programs with small spectral distribution variability (and high  $g$  value) both approaches do not significantly differ. For the larger program, much more test information is available ( $N$ ), the  $g$  parameter is somewhat lower, and the spectral distribution is highly non-uniform. In this case (for  $C > 1$ ), Zoltar-M clearly outperforms all statistical approaches. The disparity with the synthetic results is due to the particular spectral distribution properties of real programs (such as components being executed in all failed runs). Aimed at providing a first-order understanding of the impact of some of the main parameters on diagnostic performance, our simple, probabilistic program model is still far from being able to accurately account for real program behavior.

Although both the reasoning and statistical approach are based on the same (spectral) information, our reasoning approach generally produces improved diagnostic information, in terms of debugging effort and/or (most notably) potential debugging parallelism. Nevertheless our results also indicate that even in the multiple-fault case statistical approaches are by no means outclassed by our reasoning approach, a result that was not initially anticipated. Given the higher complexity of the reasoning approach, there may be situations where application of a statistical technique such as Ochiai or Zoltar-S may be preferred over Zoltar-M. In this respect, we believe the result may be relevant in the context of the multiple-fault / parallel debugging work by Jones, Bowring, and Harrold [19]. Provided their clustering approach produces spectral partitions that apply to a single fault (or a very low fault multiplicity), our results would suggest the use of Zoltar-S, rather than Tarantula.

## Acknowledgments

We extend our gratitude to Johan de Kleer for discussions which have influenced our multiple-fault reasoning approach. Also thanks to Rafi Vayani for conducting initial experiments on the effect of the hitting set filter in the single-fault case. Finally, we acknowledge the feedback from the discussions with our TRADER project partners.

## References

- [1] R. Abreu and A. J. C. van Gemund. A low-cost approximate minimal hitting set algorithm and its application to model-based diagnosis. In V. Bulitko and J. C. Beck, editors, *Proceedings of the 8th Symposium on Abstraction, Reformulation and Approximation (SARA'09)*, Lake Arrowhead, California, USA, 8 – 10 July 2009. AAAI Press.
- [2] R. Abreu, P. Zoetewij, and A. van Gemund. Localizing software faults simultaneously. In B. Choi, editor, *9th International Conference on Quality of Software (QSIC'09)*. IEEE Computer Society, August 2009.

- [3] R. Abreu, P. Zoetewij, and A. J. C. van Gemund. On the accuracy of spectrum-based fault localization. In P. McMinn, editor, *Proceedings of the Testing: Academia and Industry Conference - Practice And Research Techniques (TAIC PART'07)*, pages 89–98, Windsor, United Kingdom, September 2007. IEEE Computer Society.
- [4] R. Abreu, P. Zoetewij, and A. J. C. van Gemund. An observation-based model for fault localization. In B. Liblit and A. Rountev, editors, *Proceedings of the 6th Workshop on Dynamic Analysis (WODA'08)*, pages 64–70. ACM Press, July 2008.
- [5] R. Abreu, P. Zoetewij, and A. J. C. van Gemund. Spectrum-based multiple fault localization. In G. Taentzer and M. Heimdahl, editors, *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering (ASE'09)*, Auckland, New Zealand, 16 – 20 November 2009. IEEE Computer Society.
- [6] A. Avižienis, J.-C. Laprie, B. Randell, and C. E. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Sec. Comput.*, 1(1):11–33, 2004.
- [7] G. K. Baah, A. Podgurski, and M. J. Harrold. The probabilistic program dependence graph and its application to fault diagnosis. In *Proceedings of International Symposium on Software Testing and Analysis (ISSTA'08)*.
- [8] T. M. Chilimbi, B. Liblit, K. K. Mehra, A. V. Nori, and K. Vaswani. Holmes: Effective statistical debugging via efficient path profiling. In *Proceedings of the 31st International Conference on Software Engineering (ICSE'09)*, pages 34–44, Vancouver, Canada, 16 – 24 May 2009. IEEE CS.
- [9] J. de Kleer. Diagnosing intermittent faults. In G. Biswas, X. Koutsoukos, and S. Abdelwahed, editors, *Proceedings of the 18th International Workshop on Principles of Diagnosis (DX'07)*, pages 45 – 51, Nashville, Tennessee, USA, 29 – 31 May 2007.
- [10] J. de Kleer, A. K. Mackworth, and R. Reiter. Characterizing diagnoses and systems. *Artificial Intelligence*, 56:197–222, 1992.
- [11] J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.
- [12] M. Esser and P. Struss. Automated test generation from models based on functional software specifications. In M. M. Veloso, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 2255–2268, Hyderabad, India, 6 – 12 January 2007. AAAI Press.
- [13] A. Feldman, G. Provan, and A. J. C. van Gemund. Computing minimal diagnoses by greedy stochastic search. In D. Fox and C. P. Gomes, editors, *Proceedings of the 23rd National Conference on Artificial Intelligence*

- (*AAAI'08*), pages 919–924, Chicago, Illinois, USA, 13 – 17 July 2008. AAAI Press.
- [14] A. Groce, S. Chaki, D. Kroening, and O. Strichman. Error explanation with distance metrics. *International Journal on Software Tools for Technology Transfer (STTT)*, 8(3):229–247, 2006.
  - [15] N. Gupta, H. He, X. Zhang, and R. Gupta. Locating faulty code using failure-inducing chops. In D. F. Redmiles, T. Ellman, and A. Zisman, editors, *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering (ASE'05)*, pages 263 – 272, Long Beach, California, USA, 7 – 11 November 2005. IEEE Computer Society.
  - [16] M. Harrold, G. Rothermel, R. Wu, and L. Yi. An empirical investigation of program spectra. volume 33. ACM Press, 1998.
  - [17] K. E. Iverson. *A programming language*. John Wiley & Sons, New York, NY, USA, 1962.
  - [18] T. Janssen, R. Abreu, and A. J. C. van Gemund. Zoltar: A toolset for automatic fault localization. In A. van der Hoek and T. Menzies, editors, *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering (ASE'09) - Tools Track*, Auckland, New Zealand, 16 – 20 November 2009. IEEE Computer Society.
  - [19] J. A. Jones, M. J. Harrold, and J. F. Bowring. Debugging in parallel. In D. S. Rosenblum and S. G. Elbaum, editors, *Proceedings of the ACM/SIGSOFT International Symposium on Software Testing and Analysis (ISSTA'07)*, pages 16–26, London, UK, 9 – 12 July 2007. ACM Press.
  - [20] J. A. Jones, M. J. Harrold, and J. T. Stasko. Visualization of test information to assist fault localization. In M. Young and J. Magee, editors, *Proceedings of the 24th International Conference on Software Engineering (ICSE'02)*, pages 467–477, Orlando, Florida, USA, 19 – 25 May 2002. ACM Press.
  - [21] C. Liu, X. Yan, L. Fei, J. Han, and S. P. Midkiff. SOBER: statistical model-based bug localization. In M. Wermelinger and H. Gall, editors, *Proceedings of the 10th European Software Engineering Conference held jointly with 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/SIGSOFT FSE)*, pages 286–295, Lisbon, Portugal, 5 – 9 September 2005. ACM Press.
  - [22] W. Mayer and M. Stumptner. Abstract interpretation of programs for model-based debugging. In M. M. Veloso, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'07)*, Hyderabad, India, 6 – 12 January 2007. AAAI Press.

- [23] W. Mayer and M. Stumptner. Evaluating models for model-based debugging. In A. Ireland and W. Visser, editors, *Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE'08)*, pages 128–137, L'Aquila, Italy, 15 – 19 September 2008. ACM Press.
- [24] J. Pietersma and A. J. C. van Gemund. Temporal versus spatial observability in model-based diagnosis. In C.-T. Lin, editor, *Proceedings of 2006 IEEE International Conference on Systems, Man, and Cybernetics (SMC'06)*, pages 5325–5331, Taipei, Taiwan, 8 – 11 October 2006. IEEE Computer Society.
- [25] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, April 1987.
- [26] M. Renieris and S. P. Reiss. Fault localization with nearest neighbor queries. In J. Grundy and J. Penix, editors, *Proceedings of the 18th IEEE International Conference on Automated Software Engineering (ASE'03)*, pages 30–39, Montreal, Canada, 6 – 10 October 2003. IEEE Computer Society.
- [27] R. Vayani. Improving automatic software fault localization, July 2007. Master's thesis, Delft University of Technology.
- [28] F. Wotawa, M. Stumptner, and W. Mayer. Model-based debugging or how to diagnose programs automatically. In T. Hendtlass and M. Ali, editors, *Proceedings of IEA/AIE 2002*, volume 2358 of *LNCS*, pages 746–757, Cairns, Australia, 17 – 20 June 2002. Springer-Verlag.
- [29] C. Yilmaz and C. Williams. An automated model-based debugging approach. In R. E. K. Stirewalt, A. Egyed, and B. Fischer, editors, *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering (ASE'07)*, pages 174–183, Atlanta, Georgia, USA, 5 – 9 November 2007. ACM Press.
- [30] A. Zeller. Isolating cause-effect chains from computer programs. In *Proceedings of the 10th ACM SIGSOFT Symposium on Foundations of Software Engineering (FSE'02)*, pages 1 – 10, Charleston, South Carolina, USA, 10 – 12 November 2002. ACM Press.
- [31] Z. Zhang, W. K. Chan, T. H. Tse, B. Jiang, and X. Wang. Capturing propagation of infected program states. In *Proceedings of the the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering (ESEC/FSE'09)*, pages 43–52, Amsterdam, The Netherlands, 2009. ACM.
- [32] A. X. Zheng, M. I. Jordan, B. Liblit, M. Naik, and A. Aiken. Statistical debugging: Simultaneous identification of multiple bugs. In *Proceedings of International Conference on Machine Learning (ICML'06)*, Pittsburgh, Pennsylvania, USA, 25 – 18 June 2006. ACM Press.

- [33] P. Zoetewij, R. Abreu, R. Golsteijn, and A. J. C. van Gemund. Diagnosis of embedded software using program spectra. In J. Leaney, J. Rozenblit, and J. Peng, editors, *Proceedings 14th International Conference on the Engineering of Computer Based Systems (ECBS'07)*, pages 213–218. IEEE Computer Society, 2007.