



Available online at [www.sciencedirect.com](http://www.sciencedirect.com)



Journal of Systems and Software 00 (2019) 1–42

# Detection of intermittent faults in software programs through identification of suspicious shared variable access patterns

Panagiotis Sotiropoulos, Costas Vassilakis

*Department of Informatics and Telecommunications, University of the Peloponnese, Tripolis, Greece*

*Tel.: +30-2710-372203*

*Fax: +30-2710-372160*

*panossot@uop.gr costas@uop.gr*

---

## Abstract

Intermittent faults are a very common problem in the software world, while difficult to be debugged. Most of the existing approaches though assume that suitable instrumentation has been provided in the program, typically in the form of assertions that dictate which program states are considered to be erroneous. In this paper we propose a method that can be used to detect probable sources of intermittent faults within a program. Our method proposes certain points in the code, whose data interdependencies combined with their execution interweaving indicate that they could be the cause of intermittent faults. It is the responsibility of the user to accept or reject these proposals. An advantage of this method is that it removes the need for having predefined assertion points in the code, being able to detect potential sources of intermittent faults in the whole bulk of the code, with no instrumentation requirements on the side of the programmer. The proposed approach exploits information from the dynamic behavior of the program. In comparison with parser-based approaches which analyze only the program structure, our approach is immutable to language term changes and in general is not depending on any user-provided assertions or configuration.

*Keywords:* Intermittent faults, Fault detection, Shared variables, Model-based checking

---

## 1. Introduction

2 An intermittent fault in computer software is a malfunction of a soft-  
3 ware program that occurs at intervals, usually irregular, while the software

4 functions normally at other times. Avizienis et al. [1] defines intermittent  
5 faults as the union of (a) *elusive permanent faults*, i.e. faults that mani-  
6 fest themselves conditionally, with their activation conditions depending on  
7 complex combinations of internal state and external requests, that occur  
8 rarely and can be very difficult to reproduce and (b) *transient faults*, which  
9 includes *physical faults* (i.e. faults associated with the hardware) as well as  
10 *interaction faults*, stemming from reciprocal actions with external systems.  
11 The root causes of intermittent faults can be traced to (a) particular hard-  
12 ware conditions, e.g. radiation-induced transient faults are caused by alpha  
13 particles found in chip packages and atmospheric neutrons [2], (b) limit con-  
14 ditions (e.g. out of memory or disk storage, lost interrupts, not initialized  
15 memory, unexpected data from external sources including interactions with  
16 other systems) and (c) concurrency errors, including race conditions and  
17 scheduling decisions [3][1].

18 Software-rooted intermittent faults are referred to as MandelBugs [4][5].  
19 A Mandelbug is a bug residing some location in the code, however apply-  
20 ing test cases on the code even under seemingly exact conditions does not  
21 always lead to a failure. The reason for this non-deterministic behavior  
22 is twofold: firstly, the execution of the buggy code leads to an erroneous  
23 internal condition (e.g. a wrong variable value) which does not necessar-  
24 ily manifest itself as a failure immediately, but rather it may necessitate  
25 a chain between errors (*error propagation*) until the system uses elements  
26 (e.g. variable values) involved in the erroneous internal conditions in a way  
27 that influences a perceivable system behavior. And secondly, other elements  
28 of the software system, including other applications, the operating system  
29 or the hardware, may affect the behavior of a fault in a specific application.  
30 For instance, if a multi-threaded application lacks adequate synchronization  
31 mechanisms, race conditions may occur, depending on the choices made by  
32 the operating system scheduler regarding the exact time points that threads  
33 are dispatched on the CPU for execution, or preempted. Gray [3] uses the  
34 term *Heisenbugs*, to refer to software bugs that either do not appear or  
35 change their behavior when attempts are made to discover them. Typically  
36 this is owing to the fact that when programs are debugged the execution  
37 environment and conditions change [6]: optimization features are turned off;  
38 debugger programs may initialize memory contents to zero or modify the  
39 memory layout during execution; stepwise execution alters timings; state-  
40 ments inserted to print out variable values differentiate register values and  
41 so forth. Grottko et al. [7] identifies aging-related bugs as an interesting

42 a sub-type of Mandelbugs: aging-related bugs are faults capable of causing  
43 degraded performance or increased failure rate because they either accu-  
44 mulate internal error states and/or the activation and/or error propagation  
45 of the fault is influenced by the total time the system has been running.  
46 MandelBugs and Heisenbugs are contrasted to *Bohrbugs* [3], which refers to  
47 the class of bugs that always produce a failure on retying the operation that  
48 involves the bug; in this respect, a Bohrbug is a solid and easily detectable  
49 bug, that can be isolated by standard debugging techniques.

50 In the analysis presented in [7] Mandelbugs correspond to the 36.5%  
51 of the total number of the faults discovered in the on-board software for  
52 18 JPL/NASA space missions. Carrozza et al. [5] studied an industrial  
53 mission-critical software system, in which Mandelbugs accounted for the  
54 14.56% of the total number of faults. Cotroneo et al. [8] examine four  
55 major open source projects, and report that the percentage of Mandelbugs  
56 ranges from 7.5% (for the AXIS project) to 50.2% (for the Linux project);  
57 they also assert that in their sample, the fault densities for Bohrbugs and  
58 Mandelbugs are similar for large software projects, while for smaller projects  
59 the fault density for Bohrbugs tends to be higher and that Mandelbugs take  
60 more time to fix than Bohrbugs. Chillarege [9] concludes that Mandelbugs  
61 predominantly affect non-functional aspects, such as reliability, availability  
62 and serviceability, while they rarely affect software functionality.

63 A common cause of software-rooted intermittent faults in applications, is  
64 the erroneous order of accessing shared variables in multi-threaded applica-  
65 tions, e.g. when a write-after-write (WAW) hazard occurs, a shared variable  
66 is written by a thread while it should have first been written by another  
67 one, and so forth. The more complex the software program, the greater  
68 the likelihood of an intermittent fault to occur and the harder to locate its  
69 root cause. Many research efforts have targeted the issue of intermittent  
70 faults, and in this context a number of concurrency anti-patterns, (i.e. con-  
71 currency control mechanisms that have been proven to be ineffective and  
72 error-prone) and possible solutions have been identified, e.g. [10, 11]).

73 Intermittent faults can be detected both using static and dynamic de-  
74 bugging techniques [12]. For the detection of logical faults, in particular, in  
75 the context of dynamic approaches, the programmer typically needs to add  
76 appropriate *assert* statements expressing program-specific invariants (i.e.  
77 conditions that must always hold), which is evaluated at runtime. When  
78 the invariant is found not to hold, then a fault is flagged and the developer  
79 may use a dynamic debugger to examine the program state, trying to trace

80 back the root cause of the error.

81 The method proposed in this paper, intends to help programmers dis-  
82 cover locations in the code that could cause intermittent faults that are  
83 owing to improper order of accessing shared variables. On top of an exist-  
84 ing debugging and verification tool, we add mechanisms that create traces  
85 of shared variable access sequences and rules that are able to identify such  
86 improper access patterns within these traces; these patterns may be mani-  
87 festations of intermittent fault presence. Then, the system is able to suggest  
88 to the developer code locations that may be the root cause of these inter-  
89 mittent faults. In this paper, we have chosen Java Path Finder (JPF [13];  
90 a brief overview of JPF is given in section 2.3) as the base debugger tool,  
91 on top of which the proposed method is built; we exploit the capabilities of  
92 JPF to extract runtime information from the executing program. Our ap-  
93 proach is immutable to any user configuration (e.g. parser configurations),  
94 as it exploits information from the dynamic behavior of the program, which  
95 is sourced through the mechanisms provided by JPF. More specifically, JPF  
96 functionalities are used to gather all the information about possible inter-  
97 leavings of the accesses of the shared variables from the different threads  
98 in a tree structure, and after the tree structure is shaped, it is searched for  
99 the presence of shared variable access patterns that indicate the presence of  
100 an intermittent fault. Code locations that are involved in the suspectable  
101 shared variable accesses are then identified, and these locations are pro-  
102 posed to the user (i.e. the developer) for check (e.g. code review to verify  
103 whether synchronization mechanisms are used appropriately). The devel-  
104 oper is the one who makes the final decision on whether a suggestion made  
105 by the tool should be accepted or not. Contrary to other algorithms in  
106 the literature, the proposed approach needs no instrumentation (e.g. in-  
107 sertation of appropriate assertions in selected code locations) to work. In  
108 this way, the whole extent of the executed code is always checked, and no  
109 additional effort on the side of the developer is required. The proposed  
110 technique can be used in conjunction with other intermittent fault detec-  
111 tion techniques, both at hardware and software level (e.g. [14][15][16][17]);  
112 combined application can be achieved either by the simultaneous use of in-  
113 dividual techniques (this is directly applicable for other techniques that are  
114 hardware-based, e.g. [14][16]; for software-based techniques, an integration  
115 step will be required), or through a more loosely coupled approach where  
116 the proposed algorithm is run in parallel with other techniques and their  
117 results are combined.

118 In addition, in this paper, we examine the complexity of the proposed  
119 intermittent fault detection algorithm, by experimentally quantifying the  
120 effect that partial order reduction techniques [18] have on to the limitation  
121 of this number of paths.

122 The rest of the paper is structured as follows: section 2 overviews related  
123 work, including static and dynamic verification tools and elaborating on  
124 JPF, which is used in our approach. Section 3 introduces the proposed  
125 algorithm, while section 4 discusses the complexity of the algorithm. Section  
126 5 explores methods for speeding up the execution of the proposed algorithm  
127 by (a) exploiting parallelism and (b) pruning the possible execution paths  
128 tree, with the latter techniques being able to also tackle the state explosion  
129 issue, which is inherent in state space-based approaches. Section 6 presents  
130 an experimental evaluation of the algorithm, and finally section 7 concludes  
131 the paper and outlines future work.

## 132 2. Related work

133 Since reliability is a key objective in software development, numerous  
134 techniques have been proposed and employed to aid developers to localize,  
135 identify and remove faults. Some techniques examine the source code stati-  
136 cally to identify *code smells*, i.e. characteristics that may indicate a deeper  
137 problem. Towards this direction, code smell detectors have been employed  
138 [19][20]. Similarly, software fault prediction aims to identify fault-prone soft-  
139 ware modules by using some underlying properties of the software project  
140 before the actual testing process begins [21].

141 Considering the dynamic behavior of the software, using test cases for  
142 unit-level [22] or integration testing was one of the first tools to verify soft-  
143 ware correctness [22]. Considering the size and complexity of modern soft-  
144 ware, methods for automatically generating comprehensive test case suites  
145 have been developed [23][24][25]. Since test case-based fault detection may  
146 miss certain faults, even under high code coverage, approaches to identifying  
147 faults that evade test-case based detection processes have also been proposed  
148 [26]. Additionally, taking into account that execution of test cases consumes  
149 time and resources, their minimization and management have been explored  
150 [27][28]. With security aspects gaining increasing attention in the past few  
151 years, specialized methods for analyzing and detecting software vulnerabil-  
152 ities have been developed [29].

153 When faults do manifest in software, either in the context of testing or  
154 while execution in production environments, testers and developers need

155 to pinpoint the actual fault location and root cause: to this end, a num-  
156 ber of relevant algorithms and techniques have been proposed. Besides  
157 “traditional” fault localization techniques, which include logging, asser-  
158 tions, breakpoints and profiling, a number of advanced fault localization  
159 techniques have been proposed, which are classified as (a) slice-based, (b)  
160 program spectrum-based, (c) statistics-based, (d) program state-based, (e)  
161 machine learning-based, (f) data mining-based and (g) model-based tech-  
162 niques. Wong et al. [30] provides a survey on fault localization techniques.

163 Intermittent faults however, due to their nature, may evade detection  
164 from typical fault discovery tools [3], therefore specialized methods have  
165 been developed to assist developers in identifying and removing intermittent  
166 faults. In the rest of section we overview related work for intermittent fault  
167 detection. We initially survey work in the domain of static debuggers, and  
168 subsequently we examine approaches using dynamic debuggers. Finally, we  
169 give a brief introduction to JPF, the dynamic debugging tool used for the  
170 instrumentation of the proposed intermittent fault detection approach.

### 171 *2.1. Static debuggers*

172 Static debuggers analyze the software code without running it. Because  
173 these debuggers do not rely on tests, they can be extremely thorough. Theo-  
174 retically, static debuggers can test even code paths which are rarely executed  
175 in practice [31]. Because they are based on static analysis and satisfying  
176 predefined constraints, they could fail to detect some errors. Moreover,  
177 while static debuggers can be used in unsafe languages<sup>1</sup> to reveal potential  
178 bugs, they cannot guarantee that the data in memory is coherent according  
179 to any high-level criteria [32].

180 It is very common that a static analyzer tool is used to analyze the  
181 software code and then symbolic execution with SMT (Satisfiability Modulo  
182 Theory) [33] formulas of defined constraints is used for the verification of  
183 the code [34].

184 Symbiosis is an example of a static debugger [35]. Symbiosis necessitates  
185 the existence of a failing scheduling, which is then analyzed to determine  
186 the root cause of the fault.

---

<sup>1</sup>An unsafe language does not ensure that primitive program operations are applied to arguments of the proper form, e.g. does not ensure that array subscripts are within the allowable range.

187 2.2. *Dynamic debuggers*

188 Dynamic debuggers examine the software code while it is running. The  
189 code is instrumented and all the possible paths are executed in order to  
190 detect candidate errors; the Partial Order Reduction technique ([18]) can be  
191 used to reduce the number of paths tested, by avoiding to re-examine some  
192 path that has been already examined while exploring some other branch.  
193 However, depending on the actual values assigned to input variables of the  
194 code, it is possible that some paths are not executed and thus the tools  
195 may miss certain code defects. In addition, because dynamic debuggers use  
196 information available at run time, which is harder to extract statically from  
197 the source code, dynamic debuggers can detect errors that are harder to  
198 discover when using static analysis tools [31].

199 The CHES tool [36] is an example of a dynamic debugger. CHES  
200 creates multiple versions of the debugged program, each one suitably in-  
201 strumented to control the scheduling of threads. The instrumentation step  
202 generates  $O(2^n)$  versions for a function with  $n$  components, however [36] re-  
203 ports that the execution of  $O(n)$  versions (context switch at one of the com-  
204 ponents each time) is usually enough to activate a concurrency fault; this  
205 however may lead to missing Heisenbugs with complex activation patterns.  
206 Furthermore, [37] reports that CHES necessitates additional scaffolding  
207 and test code, on top of the test code that would be normally needed for  
208 unit or integration testing.

209 The SCURF tool [37] also follows the instrumentation approach to create  
210 particular combinations of thread interleaving. Then, each of these versions  
211 is run against a number of test cases -coupled with test oracles- for checking  
212 system functionality that need to be available, and a spectrum-based fault  
213 localization [38] is utilized to correlate detected errors with concurrently  
214 executing code blocks.

215 CTrigger [39] focuses on atomicity violation bugs; the fault identification  
216 process begins by profiling the software and identifying potential unserial-  
217 izable interleavings, while subsequently infeasible interleavings are pruned  
218 and low-probability interleavings are ranked. Afterwards, the unserializ-  
219 able interleaving space is explored. CTrigger also requires testing inputs  
220 and oracles.

221 Java Path Finder (JPF) [13], is another dynamic debugger which is used  
222 by the proposed algorithm. In the following subsection, we provide a brief  
223 introduction to JPF, to present the core functionalities exploited by the  
224 proposed algorithm. JPF can locate a number of concurrency bugs, such as

225 deadlocks and missed signals, as well as Java-related faults e.g. unhandled  
226 exceptions and improper heap usage; in order to identify faults related to  
227 application semantics (e.g. erroneous variable values), relevant assertions  
228 should be given within the application code.

229 Table 1 summarizes the existing tools and methods, their features and  
230 capabilities and compares them to those of the proposed algorithm.



Table 1. Comparison of existing fault identification tools and methods

| Tool-method        | Scope  | Capabilities   | Limitations   |
|--------------------|--|--|---|
| Test cases         | Unit & integration testing   | Mostly detects Bohrbugs.   | MandelBugs and Heisenbugs typically evade detection; coverage alone cannot guarantee a comprehensive fault detection. Manual creation of test cases is laborious and tedious, however test case generators are available.           |
| Static debuggers   | Static checking of code properties; symbolic execution can be also performed       | Identification of resource leaks, security issues and code smells. Can be used in conjunction with SMT models for increased detection capabilities. With failing schedules available, faults can be localized. | Cannot capture dynamic behavior and may miss some errors; cannot guarantee data coherence in memory according to any high-level criteria. Use with SMT models necessitates definition of constraints ( <i>program invariants</i> ). |
| Chess [36]         | Concurrency faults   | Detects faults owing to thread interleaving  | May miss Heisenbugs with complex activation patterns; necessitates additional scaffolding and test code.  |
| SCURF [37]         | Concurrency faults   | Detects faults owing to thread interleaving and inadequate atomicity guarantees.   | Necessitates pre-crafted test cases and test oracles; errors not foreseen in these cases may be missed.   |
| CTrigger [39]      | Atomicity violation bugs   | Locates faults owing to thread interleaving, catering for efficiency.  | Necessitates pre-crafted test cases and test oracles; errors not foreseen in these cases may be missed.   |
| JPF [13]           | Concurrency faults, including deadlocks and atomicity violations; generic faults.  | Powerful detection engine, extendable via the listener mechanism.  | Needs programmer-provider assertions to detect errors related to high-level data coherence.   |
| Proposed algorithm | Enhances JPF with detection of erroneous/-suspect shared variable access patterns. | Captures all errors detected by JPF and errors related to high-level data coherence, without the need to pre-define assertions, test cases or oracles.   | May flag false positives.   |

## 231 2.3. JPF - A brief overview


232 Java Path Finder (JPF) is an open-source software verification system,  
233 initially developed by NASA, that performs model checking for Java pro-  
234 grams. While test case-based software checks only some of the potential  
235 program executions and may thus miss errors, model checking automati-  
236 cally combines the behavior of state machines with a specification, which  
237 corresponds to the properties that the system should satisfy [13, 40, 41]. In  
238 more detail, the model checker accepts as input the state machine (FSM) of  
239 the program and the specification, and exhaustively explores all executions  
240 in a systematic way, flagging executions where the specification is found not  
241 to hold [42, 43]. The JPF code is available at [44].

242 While the systematic generation of all potential execution paths covers  
243 the whole search space of program states, handling millions of combinations  
244 which are hard to be modeled by manually crafted test cases [42], and  
245 thus expose all errors, this approach entails excessive computation cost,  
246 which renders it infeasible [42]. Two techniques can be used here to reduce  
247 computation costs, namely backtracking and state matching.

248 • *Backtracking* is a technique that allows the restoration of previous  
249 execution states, to examine if there are unexplored choices left. For  
250 instance, if JPF reaches a program end state, it can walk backwards  
251 to find different possible scheduling sequences that have not been ex-  
252 plored yet. While this theoretically can be achieved by re-executing  
253 the program from the beginning, and arranging that a different schedul-  
254 ing sequence is adopted in each execution, backtracking is a much more  
255 efficient mechanism if state storage is optimized.

256  
257 • *State Matching* is another key mechanism to avoid unnecessary work.  
258 The *execution state* of a program mainly consists of the heap and  
259 thread-stack snapshots. While JPF executes, it checks for every new  
260 state, whether an identical one has already been explored (c.f. Fig.  
261 1); in this case, there is no use to explore again from that state on-  
262 wards. When state matching occurs, JPF backtracks to the nearest  
263 non-explored non-deterministic choice.

264 Since concurrent actions can be executed in any arbitrary order, con-  
265 sidering all possible interleavings of concurrent actions can lead to a very  
266 large state space. It can be shown that the number of states increases ex-  
267 ponentially with the number of threads [45]. JPF uses a technique called



The final result is the same no matter which path is followed.

Figure 1. State matching: both execution paths lead to the same state (state 3).

268 *Partial Order Reduction* (POR [46]), which basically identifies statements  
 269 whose order of interleaving does not affect in any way the overall program  
 270 execution, and groups them into a single state transition, reducing thus  
 271 drastically the number of states that must be maintained. For instance, the  
 272 instructions of threads *T1* and *T2* in Fig. 2 can be interleaved in any of the  
 273 six ways listed in the same figure, however to verify program correctness it  
 274 suffices to explore the two paths highlighted in Fig. 3 [47, 46].

275 JPF uses a customizable Virtual Machine that supports various features  
 276 related to model checking, including state storage and state matching. Ac-  
 277 tually, JPF is a virtual machine (VM) running on top of the Java Virtual  
 278 Machine (JVM) and controlling its operation. The core JPF model sup-  
 279 ports checks for generic properties, such as absence of unhandled exceptions,  
 280 deadlocks, and race conditions.

281 Listeners are perhaps the most important extension mechanism of JPF.  
 282 They provide a way to observe, interact with and extend JPF execution  
 283 through code provided in the form of custom classes. Listeners are dynam-  
 284 ically configured at runtime, and therefore they do not require any modifi-  
 285 cation to the JPF core. Listeners are executed at the same authorization  
 286 level as JPF, so no limitations are imposed to their functionality (c.f. Fig.  
 287 4). In our approach, we use one listener that observes shared variable access  
 288 by threads, and logs these accesses for further analysis.

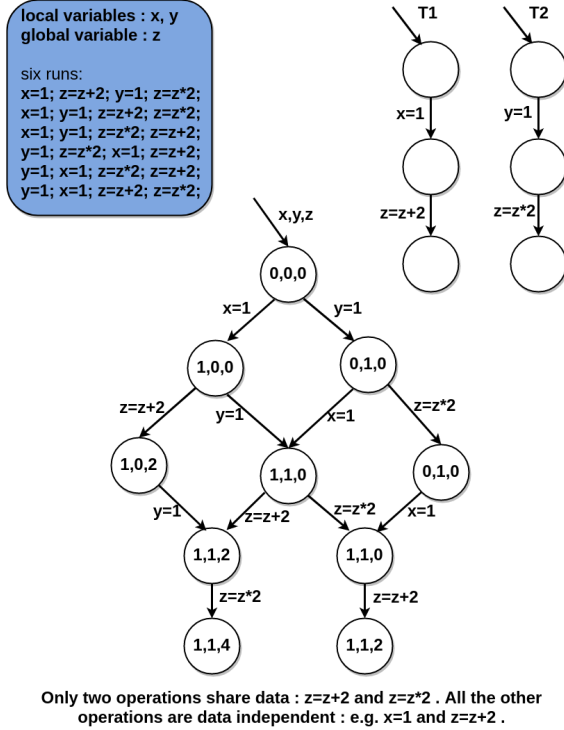


Figure 2. POR - Partial Order Reduction Example

289 **3. The Proposed Intermittent Fault Detection Algorithm**

290 The method proposed in this paper comprises three parts: The first  
 291 one encompasses the development of a rule base for the detection of shared  
 292 variable access patterns that may indicate sources of intermittent faults.  
 293 The second one is about the generation of complete execution traces for  
 294 the target program, to record all possible shared variable access patterns.  
 295 In this part, JPF, augmented with additional logging listeners, is used to  
 296 implement the generation of the program traces. The third one comprises  
 297 the application of the rule base developed in part 1 on the traces generated  
 298 during step 2, in order to detect possible sources of intermittent faults within  
 299 the program. These points are proposed to the user for review and, if  
 300 appropriate, application of the necessary corrections.

301 The first phase (rule base development) need not be performed for each  
 302 program. Instead, a generic rule base can be developed once and be subse-  
 303 quently applied to all target programs. It is possible that derivatives of the  
 304 generic rule base are created to match the requirements of specific program

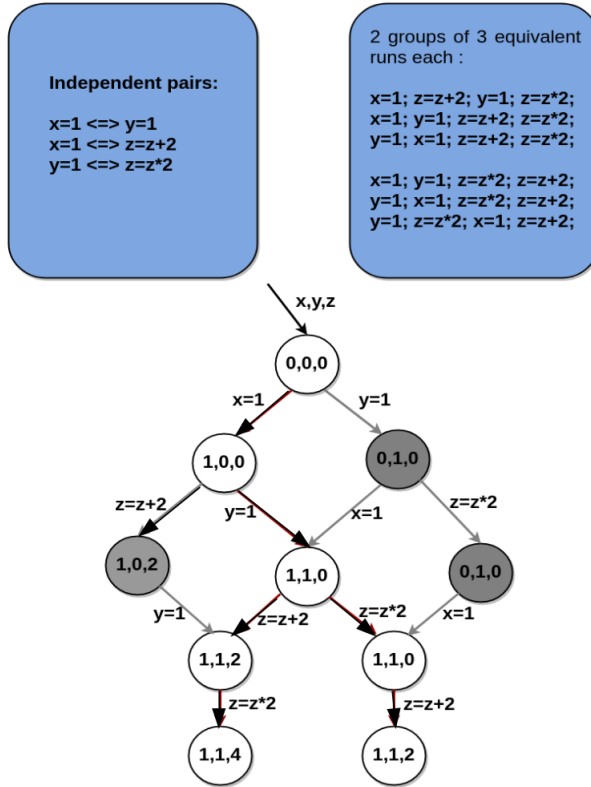


Figure 3. POR - Partial Order Reduction Example

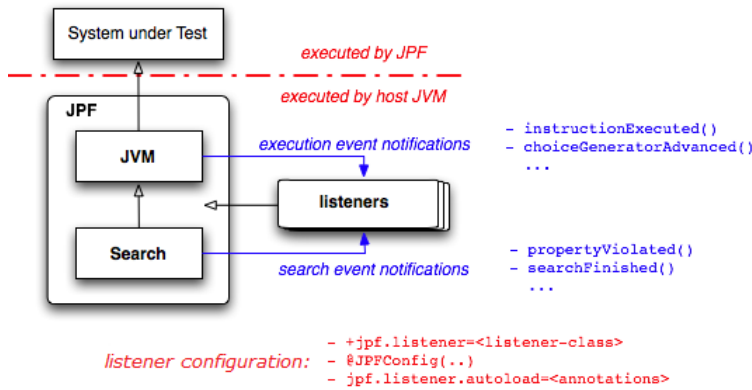


Figure 4. JPF listeners

305 classes, e.g. programs with different isolation levels. The basic rules that  
306 we use in this paper, are :

- 307 1. sequences of operations of the form  $read_{T_1}(X)$ ,  $write_{T_2}(X)$ ,  $write_{T_1}(X)$ ,  
308 corresponding to *write-after-write* hazards (the notation  $read_T(X)$  de-  
309 notes that thread T reads variable X; and similarly for  $write_T(X)$ )
- 310 2. sequences of operation of the form  $read_{T_1}(X)$ ,  $write_{T_2}(X)$ ,  $read_{T_1}(X)$ ,  
311 corresponding to the *read-after-write* hazard

312 When either of these patterns is detected in the program traces, then  
313 the corresponding code *may* be the cause of intermittent faults. In case  
314 of the write-after-write hazard rule, we can observe that this order is not  
315 equivalent to any serializable order: if T1 were scheduled before T2, then  
316 the value finally stored in X would be the one written by T2 instead of the  
317 value written by T1, which is finally stored by the schedule above (this is  
318 known as the lost update problem [48]); and if T2 were scheduled before  
319 T1, T1 would have read the value of X stored by T2 instead of the value  
320 previously stored for X (recall that T1 may have used the value read for X  
321 to compute a new value for X, so reading a different value leads to erroneous  
322 computation). In case of the *read-after-write* hazard rule, before we read  
323 our shared variable X for the second time on thread T1, it is modified on T2  
324 with a write operation, which makes the value of X on thread T1 to have a  
325 new value (without having saved the product of the previous X read value  
326 or simply using different values of X in different parts of the computation  
327 in T1; the latter is also known as the non-repeatable read problem [49]),  
328 which may be a potential cause of some intermittent fault.

329 In the second phase, we create a tree structure, where we store the data of  
330 the states of the JPF run. Each state can be a combination of data from the  
331 different threads that are accessed concurrently. Typical data that we store  
332 during this state are: the variable name, the class name, whether the access  
333 is a read or write one, the thread id, the method name, synchronization  
334 info, the package name, the value of the variable, if there is a monitor enter  
335 or exit operation, lock info, the line of the code that is executed and the  
336 source line, etc. When the JPF run is done, the detailed trace about the  
337 accesses of the state variables is available for analysis. Since each state may  
338 have one or more previous states (recall that the state matching mechanism  
339 specifically searches for cases where multiple execution paths have led to the  
340 same state) and may lead to multiple subsequent states, the trace effectively  
341 forms a directed acyclic graph (DAG).

342 In the third phase (processing phase), we apply our rule base on this  
 343 structure in order to detect the points in the code that match our rules.  
 344 Multiple rule bases or even ad-hoc rules could be applied during this phase.  
 345 The JPF should be executed only once while the tree structure can be stored  
 346 and reused for the application of all user rules.

347 As an example of applying the proposed algorithm, consider the case  
 348 that the rule base contains the two rules listed above, and that the code  
 349 illustrated in Listing 1 is checked for the presence of intermittent faults.  
 350 Before the execution of line (3) of this code, the value of the *filled* variable  
 351 should always be smaller than *MAXNUM*, a condition that is checked by the  
 352 condition at line (1) and the code associated with it. However, in the context  
 353 of concurrent executions of the *put* method, it is possible that two distinct  
 354 threads detect that the value of the *filled* variable is equal to *MAXNUM-1*,  
 355 and subsequently each one increases the value of the variable by 1, therefore  
 356 violating the invariant  $filled \leq MAXNUM$ ; this is owing to the premature  
 357 lock release, occurring at line (2).

Listing 1. A simple code example, which produces faults intermittently

---

```

358 private final static Lock l = new ReentrantLock();
359 private static int filled = 0;
360 private static ArrayList queue = new ArrayList();
361 private static final int MAXNUM = 2;
362
363
364 public void put(Object elem) {
365
366     l.lock();
367 (1)  if (filled < MAXNUM) {
368         //other code
369 (2)  l.unlock();
370     } else {
371         l.unlock();
372         return;
373     }
374     l.lock();
375     // assert (filled < MAXNUM);
376 (3)  filled++;
377     queue.add(elem);
378     l.unlock();
379     return ;
380

```

```

381     }
382
383     public Object get() {
384         Object elem = null;
385         l.lock();
386
387         if (filled > 0) {
388 (4)         filled--;
389             elem = queue.remove(0);
390         }
391         l.unlock();
392
393         return elem;
394     }
395

```

Fig. 5 demonstrates the different access interleavings that may occur when the code of the *put* method is executed concurrently by two threads, *T1* and *T2*, assuming that the condition at line (1) evaluates to *true* for both threads. We can notice that six distinct interleavings are possible, out of which four (the 1st, 2nd, 4th and 5th branches of the tree) entail the appearance of the non-repeatable read problem. For instance, in the second branch of the tree, the following shared variable accesses will be performed:  $read_{T_1}(filled)$ ,  $read_{T_2}(filled)$ ,  $read_{T_1}(filled)$ ,  $write_{T_1}(filled)$ ,  $read_{T_2}(filled)$ ,  $write_{T_2}(filled)$ , with the first two reads corresponding to the checking of condition at line (1), and subsequently each read/write pair corresponding to the variable increment at line (3). In this sequence, we can observe that a *read-after-write* hazard occurs, since a write operation on variable *filled* is performed by thread *T1* between the two read operations performed on the same variable by thread *T2*, therefore the read performed by *T2* is non-repeatable.

Fig. 6 shows the respective states of the *filled* variable, again assuming that the condition at line (1) evaluates to *true* for both threads. Notably, when the *filled* variable is less than *MAXNUM-1* all interleavings lead to a correct state, increasing the *filled* variable by two. However, if  $filled == MAXNUM-1$ , the execution of the branches entailing the *read-after-write* hazard leads to an incorrect state, where the *filled* variable is set to *MAXNUM+1*. The proposed algorithm can thus identify code that is bound to cause intermittent faults, without any knowledge about the correctness of the states.



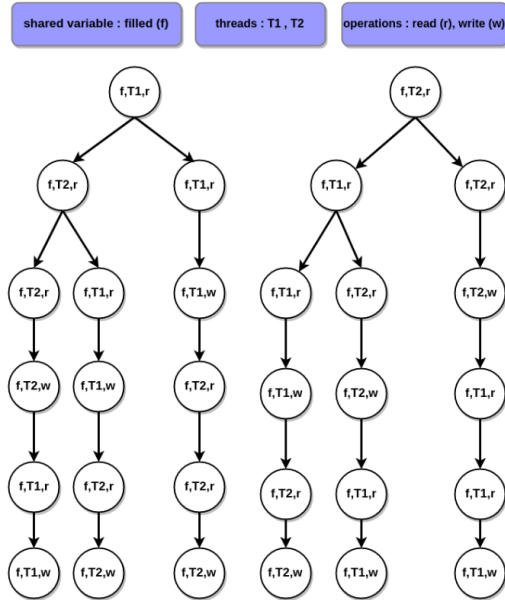


Figure 5. The possible access interleavings of the shared variable "filled" when two put methods of two different threads are executed concurrently.




Figure 6. The states of the shared variable "filled" when two put methods of two different threads are executed concurrently.

420 Using JPF to generate our state tree structure, has the advantage that  
 421 the user does not need to give any a-priori information about the code  
 422 (shared variables, atomic blocks, etc.), while JPF is not sensitive in possible  
 423 code structure changes as a static analyzer could potentially be.

424 An implementation of the algorithm is available in open source at [https:](https://)

425 [//github.com/pansot2/JPF](https://github.com/pansot2/JPF).

#### 426 4. Complexity analysis

427 When a multithreaded program with  $k$  threads executes, at each time  
 428 point the scheduler may pick any of the threads that are not in a suspended  
 429 state to execute, thus having a maximum of  $k$  alternative choices. Since  
 430 we focus on operations that access shared variables (because inappropriate  
 431 shared variable access patterns are a major cause of intermittent faults), if  
 432 we consider that each thread  $t_i$  performs  $n_i$  shared variable accesses, then  
 433 the corresponding states of a multithreaded program can be arranged in a  
 434 tree whose rank is equal to the number of threads  $k$  and its depth is equal  
 435 to:

$$depth = \sum_{i=1}^k n_i \tag{1}$$

436 The root of the tree corresponds to the initial state of the program, while  
 437 an edge denotes a transition from a state to a subsequent one, through  
 438 the execution of an instruction that accesses a shared variable, with the  
 439 instruction belonging to some thread  $t_i$  (c.f. Fig. 5). Sibling tree states  
 440 correspond to different scheduling decisions. The total number of paths in  
 441 the tree is equal to the number of ways to interleave  $k$  ordered sequences.

442 In order to compute the number of paths, we consider that the instruc-  
 443 tions in each thread  $t_i$  ( $1 \leq i \leq k$ ) are essentially ordered lists, and we want  
 444 to interleave the elements of these lists while preserving the order of the  
 445 elements in each ordered list.

446 According to equation (1), there will be  $n_1 + n_2 + \dots + n_k$  places that  
 447 we must fill (one place for each level of the tree). We can proceed by first  
 448 assigning the elements of the first list, corresponding to the instructions of  
 449 the first thread, to places. Therefore, we select  $n_1$  out of the available  $n_1 +$   
 450  $n_2 + \dots + n_k$  places, and we assign to the selected  $n_1$  places the instructions of  
 451 the first thread, preserving their order. The number of possible alternatives  
 452 is  $\binom{n_1+n_2+\dots+n_k}{n_1}$ .

453 Next, we choose  $n_2$  of the remaining places that will accommodate mem-  
 454 bers of the second list, which correspond to the instructions of the second  
 455 thread. Out of the total number of  $n_1 + n_2 + \dots + n_k$  places in the list,  
 456  $n_1$  are now occupied by elements of the first list, therefore the number of  
 457 available places in the list is  $n_2 + \dots + n_k$ . Consequently, the number of

458 possible alternatives is  $\binom{n_2+\dots+n_k}{n_2}$ . Working in the same way with the re-  
 459 remaining ordered lists, when placing the elements of the  $k^{th}$  list there are  $n_k$   
 460 elements to be placed in  $n_k$  positions, therefore there exist  $\binom{n_k}{n_k}$  alternatives.

461 Combining all the above, the mathematical formula that calculates the  
 462 number of ways to interleave  $k$  ordered sequences is:

$$\prod_{i=1}^k \binom{\sum_{j=i}^k n_j}{n_i} \quad (2)$$

463 In each state of the execution tree, we store information about the cur-  
 464 rent accesses of the shared variables, synchronization info, etc. for all active  
 465 threads. If a thread progresses, by accessing a shared variable, recording  
 466 changes in the synchronization info, etc., then a new tree node is created as  
 467 a child of the previous state (c.f. Fig. 7).

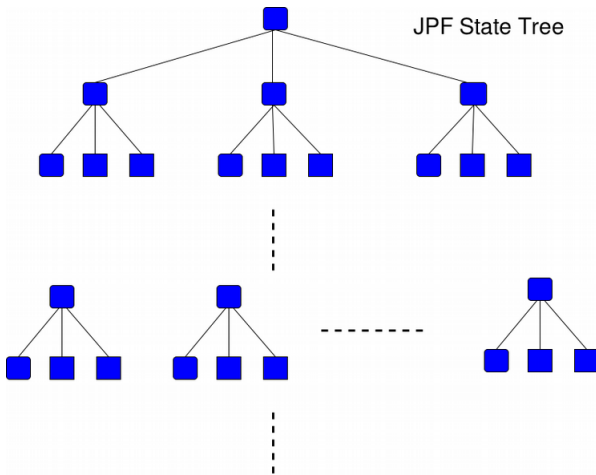


Figure 7. JPF State Tree - Rank of the tree: The maximum number of threads that can run in parallel. Depth of the tree: The sum of the accesses of all shared memory variables for all the threads.

468 In our work, execution path traversal is instrumented via the JPF model  
 469 checker, which is a so-called explicit-state model checker, since it enumer-  
 470 ates all visited states, and therefore suffers from the state-explosion problem  
 471 inherent in analyzing large programs [50], while the number of paths to be  
 472 examined also increases rapidly, with the number of threads and instructions  
 473 per thread (c.f. equation 2. However, JPF employs a number of techniques  
 474 including POR (Partial Order Reduction), state matching, and branch cov-  
 475 erage [50] to reduce the number of states and the number of paths that

476 will be examined. Using these techniques, JPF can scale up to analyzing  
477 programs up to 100,000 lines of code [51].

478 Insofar, there has not been any theoretical analysis of the effect that  
479 POR and state matching have on the complexity of the algorithms that  
480 explore the search space of possible execution paths. This is due to the  
481 fact that the final effect is highly dependent on the specific instruction  
482 placement for each program (which affects the number of cases that POR  
483 can be applied), existence and location of lock/unlock instructions (which  
484 may limits the actual choices available to the scheduler at each step), as  
485 well as volatility of external inputs, which is a determinant factor for the  
486 number of cases that states will be actually matched. Further theoretical  
487 analysis of this aspect is part of our future work.

488 The proposed algorithm dictates that shared variable accesses that are  
489 performed along an execution path are recorded, and access sequences are  
490 scanned for occurrences of the two concurrency hazards, i.e.:

- 491 1. access patterns of the form  $read_{T_1}(X)$ ,  $write_{T_2}(X)$ ,  $write_{T_1}(X)$ , corre-  
492 sponding to *write-after-write* hazards
- 493 2. access patterns of the form  $read_{T_1}(X)$ ,  $write_{T_2}(X)$ ,  $read_{T_1}(X)$ , corre-  
494 sponding to the *read-after-write* hazard

495 The complexity of recording shared variable access sequences within an  
496 execution path is  $O(n)$ , where  $n$  is the number of shared variable access  
497 sequences occurring within the execution path. Once the access sequence is  
498 recorded, the next task is to determine whether this sequence contains any  
499 of the access patterns (1) and (2) above. In the following, we discuss the  
500 matching procedure, considering initially the first form of access pattern  
501 for a specific thread, and subsequently we generalize for the second form of  
502 access patterns and all threads of a program.

503 When searching for access patterns of the form (1) above, it is not neces-  
504 sary that the instructions are found in strict sequence. The following types  
505 of instructions may intervene between the first instruction of the pattern  
506 ( $read_{T_1}(X)$ ) and the last one ( $write_{T_3}(X)$ ):

- 507 • accesses to other shared variables, either by the same or by other  
508 threads, e.g.  $read_{T_1}(Y)$  and  $write_{T_2}(Y)$ ;
- 509 • accesses to the same shared variable by threads other than T1, e.g.  
510  $read_{T_3}(X)$  and  $write_{T_2}(X)$ ;

511 If such instructions occur, then still the hazard can be flagged, since  
 512 they have no effect on the semantics of the pattern. Additionally, we can  
 513 note the following:

- 514 • if a  $read_{T1}(X)$  instruction occurs between the first and the second  
 515 instruction of the pattern (i.e. we have an access sequence  $read_{T1}(X)$ ,  
 516  **$read_{T1}(X)$** ,  $write_{T2}(X)$ ,  $write_{T1}(X)$ ), then the hazard still exists, and  
 517 it actually maps to the instructions 2-4 of the extended access pattern  
 518 (i.e. the first  $read_{T1}(X)$  is not a part of the hazard; the hazard occurs  
 519 later on).
  
- 520 • similarly, if a  $write_{T1}(X)$  instruction occurs between the second and  
 521 the third instruction of the pattern (i.e. we have an access sequence  
 522  $read_{T1}(X)$ ,  $write_{T2}(X)$ ,  **$write_{T1}(X)$** ,  $write_{T1}(X)$ ), then the hazard still  
 523 exists, and it actually maps to the instructions 1-3 of the extended  
 524 access patterns (i.e. the last  $write_{T1}(X)$  is not a part of the hazard;  
 525 the hazard has already occurred upon the execution of the third in-  
 526 struction of the extended access pattern).
  
- 527 • if a  $read_{T1}(X)$  occurs between the second and the third instruction of  
 528 the pattern (i.e. we have an access sequence  $read_{T1}(X)$ ,  $write_{T2}(X)$ ,  
 529  **$read_{T1}(X)$** ,  $write_{T1}(X)$ ), then the hazard *does not occur*, since the  
 530 computation of the value written by the fourth instruction has been  
 531 performed based on a “fresh” copy of variable  $X$  (i.e. a copy obtained  
 532 after thread  $T2$  has written a new value [48]).
  
- 533 • finally, if a  $write_{T1}(X)$  instruction occurs between the first and the  
 534 second instruction of the pattern (i.e. we have an access sequence  
 535  $read_{T1}(X)$ ,  **$write_{T1}(X)$** ,  $write_{T2}(X)$ ,  $write_{T1}(X)$ ) then the hazard *may*  
 536 *occur*, since the value stored by the fourth instruction *may* be depen-  
 537 dent on the value read by thread  $T1$  during the execution of the first  
 538 instruction of the extended access sequence.

539 Considering all the above, the target access pattern may be formulated  
 540 as a regular expression of the form:

541  $read_{T1}(X) (\text{all except } read_{T1}(X))^* write_{T2}(X) (\text{all except } read_{T1}(X))^* write_{T1}(X)$   
 542 where the notation *all except*  $read_{T1}(X)$  means any either a read or write on  
 543 any shared variable, by any thread except for a read access by thread  $T1$ ;  
 544 note here that since both the number of shared variables and the number of  
 545 threads are finite, the notation *all except*  $read_{T1}(X)$  corresponds to a finite

546 set of elements, whose cardinality is  $(\#threads * \#shared\ variables - 1)$ .  
 547 Furthermore, the star operator denotes “zero or more occurrences of the  
 548 preceding element”.

549 In order to match a regular expression against an element sequence, a  
 550 deterministic finite state automaton can be used [52]; Fig. 8 depicts the  
 551 deterministic finite state automaton which matches the regular expression  
 552 described above. The finite state automaton performs the match in linear  
 553 time, performing one state transition for each input symbol. Therefore,  
 554 matching a single instance of a rule, pertaining to a specific thread and a  
 555 specific shared variable, can be performed in linear time.




Figure 8. Deterministic finite state automaton for matching the access pattern

556 In the detection phase multiple instances of rules must be matched  
 557 against the shared variable access traces of each execution path; effectively,  
 558 each of the two rules corresponding to the *write-after-write* and the *read-*  
 559 *after-write* hazard must be specialized for each thread and each shared  
 560 variable. Therefore a maximum of  $(2 * \#threads * \#shared\ variables)$  rules  
 561 must be matched; the number may be lower if some thread  $T$  does not read  
 562 or write some shared variable  $V$ , in which case the respective rule instances  
 563 specialized for thread  $T$  and shared variable  $V$  need not be considered.

Matching of all rule instances can be performed by a single reading of the shared variable access trace, by combining the deterministic finite state automata into a single deterministic finite state automaton capable of recognizing all suspect shared variable access patterns. The procedure for building the automaton is described in [52] and summarized in the following. Let  $M_i = (K_i, \Sigma, s_i, F_i, \delta_i)$  be the automaton that realizes the match of rule instance  $R_i$ , where  $K_i$  is the set of states of the automaton,  $\Sigma$  is the alphabet (read and write operations on shared variables by threads),  $s_i$  is the start state of the automaton,  $F_i$  is the set of final states and  $\delta_i$  the transition function for  $M_i$ . A new non-deterministic automaton  $M_{nd} =$

$(K_{nd}, \Sigma, s_{nd}, F_{nd}, \delta_{nd})$  is constructed where:

$$K_{nd} = S_{nd} \cup \left( \bigcup_i K_i \right) \tag{3}$$

$$s_{nd} = S_{nd} \tag{4}$$

$$F_{nd} = \bigcup_i F_i \tag{5}$$

$$\delta_{nd} = \left( \bigcup_i \delta_i \right) \cup \left( \bigcup_i \{S_{nd} \xrightarrow{\epsilon}\} \right) \tag{6}$$

564 Effectively, a new start state  $S_{nd}$  is introduced which is non-deterministically  
 565 linked to all start states of the individual automata under an  $\epsilon$ -transition  
 566 (i.e. a transition that occurs with no input), and all final states of the  
 567 individual automata are considered as final states in the merged automa-  
 568 ton. The non-deterministic automaton is depicted in Fig. 9. Finally, the  
 569 non-deterministic automaton is converted to a deterministic one, using the  
 570 algorithm described in [52].




Figure 9. The non-deterministic automaton.

571 Since an execution path typically includes other instructions besides ac-  
 572 cesses to shared variables (e.g. accesses to local variables, computations,  
 573 etc.), the overall complexity of shared variable access recording and match-  
 574 ing is inferior to that of the execution of the path. Additionally, note here  
 575 that in the context of the execution performed by JPF as part of the explicit  
 576 state-model checking, some operations such as state matching are expensive  
 577 ones, needing to examine a number of data elements (e.g. values of state  
 578 variables), as contrasted to the shared variable access recording and rule  
 579 matching operations introduced by the algorithm, where each shared vari-  
 580 able access is recorded or processed in the merged deterministic finite state

581 automaton in a  $O(1)$  operation. Therefore, the overall complexity of the  
 582 suspicious shared variable access pattern detection procedure is dominated  
 583 by the complexity of the execution of the different execution paths, which  
 584 is instrumented by JPF which -as noted above- can satisfactorily handle  
 585 programs of the magnitude of 100,000 lines of code.

## 586 5. Optimizing Intermittent Fault Identification

587 While the presented algorithm leverages the intermittent error detection  
 588 potential, it introduces additional overheads. In this section, we examine  
 589 methods for limiting overheads and increasing the efficiency of intermittent  
 590 fault detection.

### 591 5.1. Separate analysis of independent thread partitions

592 Our method targets the identification of access patterns on shared vari-  
 593 ables which may lead to errors; to test whether such patterns may appear,  
 594 all possible execution paths are examined. However, when two threads do  
 595 not access any variable in common, it is not necessary to test all possi-  
 596 ble interleavings of these threads' execution, since obviously no "suspect"  
 597 variable access patterns may be identified among these threads.

598 Generalizing, we can partition the threads in the program in subsets  
 599  $TS_1, TS_2, \dots, TS_n$  where:

- 600 •  $TS_i \cap TS_j = \emptyset, \forall i, j: i \neq j$
- 601 •  $\bigcup_i TS_i = T$ , where T is the set of all program threads
- 602 •  $SVA(TS_i) \cap SVA(TS_j) = \emptyset, \forall i, j: i \neq j$ , where  $SVA(TS_j)$  is the set  
 603 of all shared variables accessed by any thread in  $TS_j$

604 In order to exploit this aspect towards the optimization of the intermit-  
 605 tent fault identification process, we override the default choice generation  
 606 and backtracking behavior of JPF, to allow the user to specify the threads  
 607 whose execution will be monitored in a particular execution. At implemen-  
 608 tation level, this is realized by overriding the *stateAdvanced* method of  
 609 the listener, which controls what happens when a new state is generated.  
 610 In more detail, the user defines in the configuration file the threads that  
 611 contain related shared variables, which form a thread subset, and the rest  
 612 of the threads being partitioned into trivial, single-thread subsets which



613 are ignored in order not to produce any alternative choices; only a single  
 614 choice is considered for these threads. Effectively we have a model involv-  
 615 ing some “interacting threads” and some “independent threads”. This is  
 616 accomplished using the following configuration parameter:

617 *vm.watched.threads* = the threads that should trigger alternative choice gen-  
 618 erations in JPF

619 The code in the listener that handles this functionality is illustrated in  
 620 the following Algorithm (Listing 2):

Listing 2. Process followed by the listener used for choice generation only for a subset of the threads of the software program

- 
- 621
- 622 1. Get information about the watched threads defined
  - 623 with the configuration parameter *vm.watched.threads*
  - 624 2. Get information about the threads that the watched
  - 625 threads depend on; this is defined via the
  - 626 configuration parameter *vm.watched.threads.seqdeps*
  - 627 3. In the stateAdvanced overridden method, ignore the
  - 628 states that are not caused by executing
  - 629 instructions of the watched threads or the threads
  - 630 they depend on. This is accomplished by invoking
  - 631 the `search.getVM().ignoreState()` method.
  - 632
- 

633 As it can be noticed in the process above, the *ignoreState()* method  
 634 that JPF provides is used in order to ignore the states related to a thread  
 635 change that are not included at the *vm.watched.threads* list in the con-  
 636 figuration. There is no need to make alternative choices for the scheduling  
 637 of threads that are not included in the *vm.watched.threads* variable and  
 638 have no watched thread depending on them, as these, in general, do not  
 639 influence the subset of threads for which the intermittent fault analysis is  
 640 conducted.

641 In order to comprehensively analyze the program for existence of in-  
 642 termittent faults, the intermittent fault detection procedure should be run  
 643 for each thread subset  $TS_i$ . If the number of states examined when an-  
 644 alyzing thread subset  $TS_i$  is equal to  $numStates(TS_i)$ , then the analysis  
 645 of all subsets entails the examination of  $\sum_i numStates(TS_i)$  states. Con-  
 646 trary, if a combined analysis of all threads is employed (i.e. the thread  
 647 “independence property” is not exploited), then the analysis will entail the  
 648 examination of  $\prod_i numStates(TS_i)$  states, under the assumption that no  
 649 dependent threads exist for each thread subset. It is clear that the thread  
 650 partitioning scheme introduces significant performance gains.

651 The formulation of independent thread partitions that are separately  
652 examined for suspect shared variable access patterns contributes to the  
653 reduction of the state space that need to be examined, alleviating thus the  
654 issue of state explosion. The gains regarding the aspect of state space size  
655 reduction are quantified through the experiments presented in subsection  
656 6.3.

657 At this stage of development, we have delegated the responsibility of  
658 partitioning threads to thread subsets to the user; in our future work, we  
659 will consider automatic or semi-automated ways to determine independent  
660 thread subsets.

### 661 5.2. Pruning state subtrees of specific nodes

662 In this section we explore the potential to optimize the intermittent fault  
663 analysis time, by reducing the nodes of the JPF tree for different ranks and  
664 depths. This method can be used for software programs employing a Boss-  
665 workers model [53] (or the dispatcher-worker model, as listed in [54]), where  
666 the different tasks are distributed to workers which use code/libraries that  
667 are independent (in terms of shared variables) from the rest of the program.  
668 A typical case of this example is the web server process service loop, where  
669 requests are accepted by the main thread and then their execution is del-  
670 egated to worker threads, with worker threads being totally independent  
671 and not accessing any shared variables. In this case, we could avoid the  
672 expansion of alternatives for worker thread states, since -by virtue of their  
673 independence- are not bound to be the source of intermittent faults (cf. Fig.  
674 10).

675 Pruning state subtrees can be configured by specifying the depth of the  
676 state tree at which pruning will occur and the order of the child nodes  
677 at this level to be allowed: at the present state of development, prun-  
678 ing is regulated via the properties listed in table 2. At runtime, when  
679 the listener detects that a subtree should be pruned it executes the state-  
680 ment *ti.breakTransition(true)*; which breaks the current transition  
681 and forces an end state.

682 Pruning the state subtrees of specific nodes contributes to the alleviation  
683 of the state explosion problem, since the state space that is explored within  
684 the program execution is reduced. The gains regarding the aspect of state  
685 space size reduction are quantified through the experiments presented in  
686 subsection 6.3.

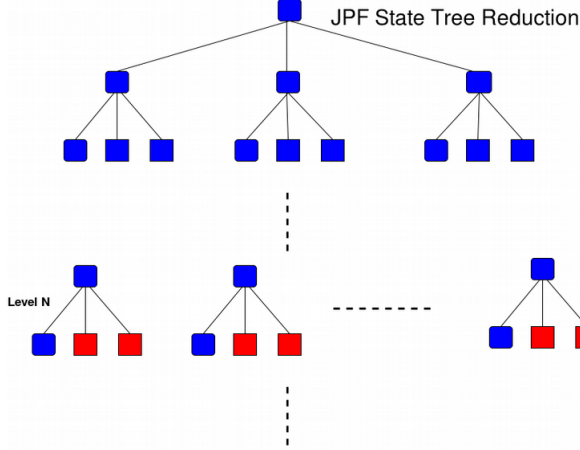


Figure 10. JPF State Tree Reduction : Pruning subtrees of nodes at Level N by allowing execution of the first child of each node only.

| Property name           | Description  |
|-------------------------|--|
| vm.parallel.allow.depth | the level of the nodes where the reduction will be applied (single value or a range)   |
| vm.parallel.allow.child | the order of the child node(s) at the specified level that will be allowed to continue (e.g. the value “2” designates that the 2nd child will be allowed to continue, while execution of other children will be inhibited) |

Table 2. Parameters regulating the pruning of state subtrees

687 *5.3. Exploiting processing power in share-nothing architectures*

688 The proposed method is orchestrated on top of JPF, and can thus ben-  
 689 efit from JPF’s potential to run efficiently on shared memory architectures,  
 690 exploiting multiple execution cores for accelerating the state space search  
 691 procedure [55]. To further scale parallelism potential and take advantage  
 692 of share-nothing architectures, the user could designate specific JPF paths  
 693 whose exploration would be assigned to a different machine. More specifi-  
 694 cally:

- 695 • the listener examines all transitions
- 696 • when a path that is designated to be transferred to another machine is

697 reached, the state space is serialized and transferred over the network  
698 to the destination machine

699 • at the local machine further exploration is inhibited by means of the  
700 *ti.breakTransition(true)*; statement, which breaks the current  
701 transition and forces an end state.

702 • On the remote machine, the state is deserialized, and execution re-  
703 sumes from the point that it was suspended; in this case, the listener  
704 does not issue the *ti.breakTransition(true)*; statement, allowing  
705 the exploration of the path.

706 The state serialization, transfer and execution resumption mechanisms  
707 are currently under implementation.

## 708 6. Experimental evaluation

709 In this section, we present the experiments conducted to:

- 710 1. validate the proposed algorithm in terms of its fault detection poten-  
711 tial,
- 712 2. experimentally assess the complexity of the algorithm and quantify  
713 the overhead introduced over the “plain JPF” software validation and
- 714 3. assess and quantify the gains reaped from applying the optimization  
715 methods presented in section 5.

### 716 6.1. Algorithm validation

#### 717 6.1.1. Small-scale validation

718 Initially, when we ran the non-extended version of JPF against the code  
719 illustrated in listing 1 using three parallel threads. The validation succeeds  
720 without identifying any potential error sources, exhibiting thus a false neg-  
721 ative.

722 Then we run the proposed algorithm to generate data access traces, ap-  
723 plying the rules *read(s, T1)-write(s, T2)-write(s, T1)* and *read(s, T1)-write(s, T2)-*  
724 *read(s, T1)*, which can identify data access patterns that are potentially er-  
725 roneous. After the processing of the tree, which has been generated via  
726 our listener, the following instruction interleavings are identified as possible  
727 intermittent fault causes:

- 728 1.  $t_1(1) - t_2(3) - t_1(3)$

- 729 2.  $t_1(1) - t_2(4) - t_1(3)$   
730 3.  $t_1(1) - t_2(4) - t_1(4)$

731 where  $t_i(j)$  denotes the execution of instruction  $j$  (cf. code example in  
732 Section 3) by thread  $i$ .

733 The user is invited to review the relevant code and accept or reject  
734 those proposals. In the above case, the first case flagged by the algorithm  
735 is a source of intermittent faults, since it may lead to the violation of the  
736 *filled* < *MAXNUM* program invariant, as we would like in every case  
737 before executing line (3) to have a shared variable which is smaller than  
738 *MAXNUM* (*filled* < *MAXNUM*). The two last sequences do not gener-  
739 ate an intermittent fault in our case, however it is worth noting that when  
740 the instructions `l.unlock()`; (line (2) in listing 1) as well as the corre-  
741 sponding `l.lock()` instruction immediately preceding line (3) in the same  
742 listing, effectively thus removing the atomicity violation which is the root  
743 cause of the error flagged in the first case, the second error flagging are  
744 removed and only the third one is reported.

745 Our approach is able to identify previously missed intermittent faults.  
746 On the other hand, it introduces some false positives. An annotation-based  
747 approach could be used to inhibit the reporting of specific patterns that  
748 have been validated not to cause intermittent faults.

#### 749 6.1.2. Validation in real-world scale

750 To validate our approach in a real-world scale, we conducted experiments  
751 using the multithreaded Java webserver available in [56], which extensively  
752 uses multithreading (using a thread pool of configurable size) and shared  
753 variables. We initially ran the non-extended version of JPF against the  
754 simulation code given in GitHub [57], and no errors were flagged. The code  
755 was also checked by the proposed algorithm, and no errors were flagged  
756 either.

757 Subsequently, we followed a fault injection approach [8], to inject faults  
758 within the code and test whether these faults are detected by (a) the non-  
759 extended versions of the JPF and (b) the proposed algorithm (i.e. JPF  
760 extended with our listener and the potentially erroneous access pattern  
761 detection). The results of the tests are summarized in table 3.

762 Effectively, the proposed algorithm was able to detect all faults detected  
763 by JPF (which underpins the proposed algorithm), plus errors related to  
764 erroneous access patterns, which were missed by JPF. Therefore, the pro-  
765 posed algorithm offers more comprehensive error detection, at the expense

| Fault type  | JPF | Proposed algorithm  |
|---|-----|---|
| Deadlock  | Yes | Yes   |
| Unhandled exception                                       | Yes | Yes   |
| Race conditions   | Yes | Yes   |
| Application-specific assertions                           | Yes | Yes   |
| Erroneous shared variable access patterns as in listing 1 | No  | Yes; for the injected erroneous access pattern faults, one related false positive was also raised |

Table 3. Detection of injected faults by JPF and the proposed algorithm

766 of flagging a limited number of false positives and a performance overhead,  
767 which has been quantified to be up to 10.7%, as discussed in subsection  
768 6.2. Recall here that the potential of the proposed algorithm to detect erro-  
769 neous shared variable access patterns is advantageous over the detection of  
770 errors based on application-specific assertions, in that (a) the former does  
771 not necessitate any instrumentation by the programmer (i.e. insertion of  
772 *assert* statements), while the latter does and (b) assertion-based error de-  
773 tection is limited to detecting errors at the locations that assertions have  
774 been inserted and related to the conditions within the assertions, whereas  
775 erroneous shared variable access pattern detection can identify errors at any  
776 location and under any condition.

777 An instance of the code of [56] with injected faults is available at [58]<sup>2</sup>

## 778 6.2. Complexity Assessment Experiments

779 In this subsection, we report on the experiments conducted to gain in-  
780 sight regarding the size of the state space and the execution time needed  
781 under different thread mixtures, and present the obtained metrics. To pro-  
782 mote example clarity, we initially present a complexity analysis on the code  
783 depicted in listing 3, while subsequently we present our complexity analysis  
784 findings on our real world-scale example of the multithreaded Java webserver  
785 [56]. All the experiments reported in this subsection, as well as in the fol-  
786 lowing one, have been performed on a PowerEdge M910 blade server, with

<sup>2</sup>At different phases of the test, different faults were injected; [58] contains a specific set of faults.

787 256 GBytes of physical memory and four 8-core E7-4830 Intel Xeon proces-  
 788 sors. The Java environment had been configured to use up to 40 GBytes of  
 789 memory.

790 The example in listing 3 entails instructions belonging to three threads,  
 791 namely  $T1$ ,  $T2$  and  $T3$ . Threads  $T1$  and  $T2$  access two shared variables  $A$   
 792 and  $B$ , therefore the interleaving of their instructions can be the root cause  
 793 of intermittent fault occurrence. On the contrary, thread  $T3$  accesses only  
 794 local variables, and consequently no intermittent faults can occur due to  
 795 the instructions of this thread.

Listing 3. Thread code

---

```

796 shared int A, B;
797
798
799 T1:
800     B = 2;
801     A = B + 1;
802
803 T2:
804     B = 0;
805     B = B + 2;
806     A = B + 1;
807
808 T3:
809     local int c, d;
810     d = 0;
811     d = d + 2;
812     c = d + 1;
813
    
```

---

814 In terms of shared variable read and write operations, threads  $T1$  and  $T2$   
 815 can be written as:

Listing 4. Thread read and write operations

---

```

816 T1: w(B), r(B), w(A)
817
818 T2: w(B), r(B), w(B), r(B), w(A)
819
820 T3: -
    
```

---

821 Some possible execution schedules of threads  $T1$  and  $T2$  are presented  
 822 in the following list. For conciseness, we have only included those execution  
 823 schedules which end with the last instruction of thread  $T1$ ; inclusion of  
 824 cases that end with the last instruction of  $T2$  is done in an identical fashion.

825 The instructions of  $T1$  (i.e. instructions of thread  $T1$  for which at least  
 826 one instruction of  $T2$  intervenes between them and the last instruction of  
 827  $T1$ ) are denoted using boldface, to promote readability.

828 For the creation of the execution schedules presented in the following  
 829 list, we assume a simple processor addressing mode, where each instruction  
 830 can fetch access only one memory location (corresponding to a variable),  
 831 fetching its contents to a register or storing register contents to it. This is  
 832 in-line with the Java model, where instructions operate on operands individ-  
 833 ually stored on the operand stack, and results are then copied to variables<sup>3</sup>.  
 834 In this sense, read and write operations executed in the context of the same  
 835 instruction (e.g.  $r(B)T1$ ,  $w(A)T1$  corresponding to the instruction  $A = B$   
 836  $+ 1$ ; are separable, in the sense that thread switching can occur between  
 837 these two accesses. However, in some processors it is possible to execute  
 838 multiple variable accesses in a single instruction: for instance in the Pen-  
 839 tium it is possible to map the program instruction  $B = B + 2$ ; to a single  
 840 machine-language instruction *ADD WORD PTR [ESI], 0x2* (assuming that  
 841 the ESI register points to the memory location of variable  $B$ ) [59]. Notably,  
 842 this can happen when Java bytecode is compiled into optimized machine  
 843 instructions e.g. through the Java HotSpot VM<sup>4</sup>. In these cases, execution  
 844 schedules in which the involved read and write operations appear separated  
 845 cannot occur and therefore should not be considered; henceforth, the follow-  
 846 ing list contains only execution schedules where the operations  $r(B)T2$  and  
 847  $w(B)T2$  realizing the  $B = B + 2$ ; instruction of thread  $T2$  are adjacent.

- 848 1.  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $w(B)T1$ ,  $r(B)T1$ ,  $w(A)T1$
- 849 2.  **$w(B)T1$** ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $r(B)T1$ ,  $w(A)T1$
- 850 3.  $w(B)T2$ ,  **$w(B)T1$** ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $r(B)T1$ ,  $w(A)T1$
- 851 4.  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  **$w(B)T1$** ,  $r(B)T2$ ,  $w(A)T2$ ,  $r(B)T1$ ,  $w(A)T1$
- 852 5.  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  **$w(B)T1$** ,  $w(A)T2$ ,  $r(B)T1$ ,  $w(A)T1$
- 853 6.  **$w(B)T1$** ,  **$r(B)T1$** ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $w(A)T1$
- 854 7.  **$w(B)T1$** ,  $w(B)T2$ ,  **$r(B)T1$** ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $w(A)T1$
- 855 8.  **$w(B)T1$** ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  **$r(B)T1$** ,  $r(B)T2$ ,  $w(A)T2$ ,  $w(A)T1$
- 856 9.  **$w(B)T1$** ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  **$r(B)T1$** ,  $w(A)T2$ ,  $w(A)T1$
- 857 10.  $w(B)T2$ ,  **$w(B)T1$** ,  **$r(B)T1$** ,  $r(B)T2$ ,  $w(B)T2$ ,  $r(B)T2$ ,  $w(A)T2$ ,  $w(A)T1$

<sup>3</sup><https://docs.oracle.com/javase/specs/jvms/se8/html/jvms-6.html>

<sup>4</sup><http://www.oracle.com/technetwork/java/javase/tech/index-jsp-136373.html>



- 858 11. w(B)T2, **w(B)T1**, r(B)T2, w(B)T2, **r(B)T1**, r(B)T2, w(A)T2, w(A)T1  
859 12. w(B)T2, **w(B)T1**, r(B)T2, w(B)T2, r(B)T2, **r(B)T1**, w(A)T2, w(A)T1  
860 13. w(B)T2, r(B)T2, w(B)T2, **w(B)T1**, **r(B)T1**, r(B)T2, w(A)T2, w(A)T1  
861 14. w(B)T2, r(B)T2, w(B)T2, **w(B)T1**, r(B)T2, **r(B)T1**, w(A)T2, w(A)T1  
862 15. w(B)T2, r(B)T2, w(B)T2, r(B)T2, **w(B)T1**, **r(B)T1**, w(A)T2, w(A)T1

863 Table 4 depicts the experimental results obtained from running the code  
864 in listing 3 with a varying number of instances of threads  $T1$ ,  $T2$  and  $T3$ .  
865 The experimental results are contrasted with the theoretical maximum of  
866 possible paths, which is equal to the number of possible distinct execution  
867 schedules (c.f. Section 4).

| Complexity Comparison Table                            |                                 |  |  |                                 |  |  |  |  |
|--|---------------------------------|--|--|---------------------------------|--|--|--|--|
|  | 2<br>Threads<br>(1xT1,<br>1xT2) | 3<br>Threads<br>(1xT1,<br>1xT2,<br>1xT3) | 4<br>Threads<br>(1xT1,<br>1xT2,<br>2xT3) | 4<br>Threads<br>(2xT1,<br>2xT2) | 5<br>Threads<br>(1xT1,<br>1xT2,<br>3xT3) | 5<br>Threads<br>(2xT1,<br>2xT2,<br>1xT3) | 6<br>Threads<br>(1xT1,<br>1xT2,<br>4xT3) | 4<br>Threads<br>(2xT1,<br>2xT2,<br>2xT3) |
| Theoretical number of possible paths                   | 56                              | 56                                       | 56                                       | 40360320                        | 56                                       | 40360320                                 | 56                                       | 40360320                                 |
| Experimentally determined number of paths (JPF)        | 13                              | 18                                       | 63                                       | 49766                           | 63                                       | 90627                                    | 63                                       | 154081                                   |
| Practical number of final states (JPF)                 | 2                               | 2  | 2  | 3                               | 2  | 3  | 2  | 3  |
| Experimentally determined number of final states (JPF) | 4                               | 4  | 4  | 11                              | 4  | 11                                       | 4  | 11                                       |
| JPF end states   | 29                              | 33                                       | 37                                       | 331                             | 41                                       | 342                                      | 45                                       | 353                                      |

Table 4. Complexity comparison for varying number of threads

868 At this point, we note the following :

- 869 1. For the number of possible paths, we assume all possible execution  
870 schedules regarding instructions accessing shared variables, as calcu-  
871 lated by equation 2.  
872 2. The metric *Practical number of final states* corresponds to the number  
873 of the different program results, in terms of shared variable values.

874     3. The difference between the *Practical number of final states* and the  
 875         *JPF end states* exists because of the additional information regarding  
 876         shared data used in JPF (e.g thread shared information).

877     Regarding the complexity analysis experiments conducted using our real  
 878     world-scale example of the multithreaded Java webserver [56], table 5 depicts  
 879     the execution statistics of the multithreaded Java webserver [56] regarding  
 880     the number of states, and table 6 depicts the runtimes measured, while  
 881     varying the parameter of the execution tree depth search limit (c.f. subsection  
 882     5.2). In all experiments, all injected faults were flagged, while in the  
 883     setting where the JPF search depth was set to 350 it was observed that the  
 884     limit was not reached within the fault detection process execution, indicat-  
 885     ing that further increments to that parameter would not affect the time and  
 886     resources needed. We can observe that the overhead introduced by the pro-  
 887     posed algorithm over the non-extended version of JPF is up to 10.7%, which  
 888     is deemed acceptable, considering the increased fault detection potential of  
 889     the proposed algorithm.

|                    |              |
|--------------------|--------------|
| new states         | 35,185,856   |
| visited states     | 24,779,490   |
| backtracked states | 59,965,346   |
| end states         | 0            |
| instructions       | 499,329,9768 |
| max memory         | 30,7 GB      |

Table 5. Execution statistics for the fault detection process of the multithreaded Java webserver

890 where:

- 891     • *new states* is the number of unique states visited during the run;
- 892     • *visited states* is the number of states that are examined and have been  
 893         revisited during the same execution;
- 894     • *backtracked states* refers to the states from which the search back-  
 895         tracked, so as to examine different paths;
- 896     • *end states* refers to the concluding states of the program execution,  
 897         from which there are no forward transitions to try.

| JPF execution time |                         |                         |                         |
|--------------------|-------------------------|-------------------------|-------------------------|
|                    | JPF search<br>depth=120 | JPF search<br>depth=240 | JPF search<br>depth=350 |
| JPF (not extended) | 01:27:19                | 04:01:59                | 04:05:09                |
| Proposed algorithm | 01:34:25                | 04:26:15                | 04:31:26                |

Table 6. Fault detection process execution time for Java web server Simulation

898 *6.3. Optimization Experiments*

899 In this subsection we report on the experiments conducted to assess  
900 the gains introduced by the optimization methods presented in section 5,  
901 and present our findings. As noted above, all the experiments reported in  
902 this subsection have been performed on a PowerEdge M910 blade server,  
903 with 256 GBytes of physical memory and four 8-core E7-4830 Intel Xeon  
904 processors. The Java environment had been configured to use up to 40  
905 GBytes of memory.

906 Table 7 illustrates the performance gains obtained by isolating threads  
907 that are not bound to be involved in the occurrence of intermittent faults,  
908 regarding the code depicted in listing 3. In more detail, the first data column  
909 in Table 7 corresponds to the measurements obtained from the execution  
910 of a program whose main thread creates one instance of threads  $T1$  and  
911  $T2$ , as well as four instances of  $T3$  (1xT1, 1xT2, 4xT3); in this execution,  
912 JPF monitors all threads. The second data column in Table 7 corresponds  
913 to the execution of the same program, with JPF being however instructed  
914 to monitor only the main thread and the instances of threads  $T1$  and  $T2$ ,  
915 since no instance of thread  $T3$  is bound to be involved in the generation of  
916 intermittent faults. As described in Section 5.1, this is realized through the  
917 setting `vm.watched.threads=main,1,2`.

918 Table 8 depicts how performance benefits can be obtained from applying  
919 the Children Node Reduction technique described in Section 5.2, regarding  
920 the code depicted in listing 3. The figures in this table refer to the execution  
921 of a java program with 6 threads (1xT1, 1xT2, 4xT3), varying the order  
922 of the child that is allowed to continue. Since in our example the first and  
923 second children correspond to executions of instructions by threads  $T1$  and  
924  $T2$ , which include accesses to shared variables, these choices entail more  
925 states to be examined. Given that only these two choices may actually lead  
926 to intermittent faults, it suffices to examine only these two cases to fully  
927 uncover all intermittent fault root causes.

|                    |             |            |
|--------------------|-------------|------------|
| watched threads    | All Threads | main,T1,T2 |
| elapsed time       | 00:00:10    | 00:00:03   |
| new states         | 5147        | 2192       |
| visited states     | 14365       | 1264       |
| backtracked states | 19512       | 3456       |
| end states         | 45          | -          |
| instructions       | 418886      | 88756      |
| max memory         | 303MB       | 169MB      |

Table 7. Examining all possible threads vs. limiting the set of threads examined by JVM.

| Children Node Reduction |   |   |   |               |
|-------------------------|---|---|---|---------------|
|                         | 1st child node<br>for depths be-<br>tween 10-30 | 2nd child node<br>for depths be-<br>tween 10-30 | 3rd child node<br>for depths be-<br>tween 10-30 | No cut<br>off |
| Time                    | 2 sec   | 1 sec   | 1 sec   | 10 sec        |
| New states              | 1637  | 523   | 302   | 5147          |
| Visited states          | 1118  | 344   | 253   | 14365         |
| Backtracked<br>states   | 2755  | 867   | 555   | 19512         |

Table 8. Children Node Reduction effect applied at different thread orders

928 Table 9 focuses on the scalability of the proposed algorithm under the  
 929 optimization techniques presented in section 5, depicting the time needed to  
 930 execute the proposed algorithm to detect faults injected to the open-source  
 931 multithreaded Java web sever [56] when the thread partitioning and the  
 932 state subtree pruning of specific nodes techniques (cf. subsections 5.1 and  
 933 5.2, respectively) are applied. The configuration used in this experiment is:

934

935 `vm.parallel.allowed.depth=40-350`

936 `vm.watched.threads=main,Thread-1,Thread-2,Thread-3,Thread-4`

937 `vm.parallel.allowed.child=[1] search.depth_limit = 350`

938

939 This configuration effectively scans the full state tree up to the depth  
 940 of 40, and beyond that point limits the detection to the first child only,

| Number of threads in the web server request executor thread pool |          |          |          |          |
|--|----------|----------|----------|----------|
|  | 5        | 10       | 50       | 100      |
| Elapsed time   | 00:32:41 | 00:47:04 | 00:47:41 | 00:47:54 |

Table 9. JPF execution time for the Java web server simulation

941 since the Java web server [56] employs the worker thread model discussed  
 942 in section 5.2 [54], and moreover worker threads are totally independent and  
 943 are thus not bound to generate any intermittent errors.

944 We can observe that the time needed to run the fault detection algorithm  
 945 increases very slowly with the overall number of threads, while additionally  
 946 significant time savings against the non-optimized version (c.f. table 6) are  
 947 introduced; these savings are quantified to 82%.

## 948 7. Conclusions and Future work

949 In this paper we presented a methodology for intermittent fault detec-  
 950 tion that is based on the identification of suspicious shared variable access  
 951 patterns in the code execution traces. Execution traces are generated using  
 952 the JPF tool, which has been enhanced by a customized listener, while the  
 953 suspicious access patterns that are searched for correspond to well-known  
 954 parallel programming hazards. Our method has been shown to be capable  
 955 of detecting intermittent faults that evade detection when other methods  
 956 are used, while on the other hand introducing some false positives. In this  
 957 sense, the programmer is asked to review the potential intermittent fault  
 958 root causes and accept or reject them. In order to leverage the efficiency  
 959 of the proposed method, we have introduced optimization methods which  
 960 can exploit structural properties of the code, such as thread independence  
 961 and thread subtree isolation, as well as parallel hardware capabilities. Our  
 962 experiments on optimizations exploiting structural properties of the code  
 963 have demonstrated that significant performance gains can be reaped.

964 In the context of our future work we plan to examine the following  
 965 dimensions:

- 966 1. fully implement and evaluate the optimization method for exploiting  
 967 parallel hardware capabilities.
- 968 2. take into account dependencies between global and local variables,  
 969 which are established via assignment statements.
- 970 3. Identify and evaluate additional dependency rules.

971 4. study the computational complexity of the proposed algorithm, com-  
972 puting a theoretical upper bound for the number of possible execution  
973 paths that need to be explored.

## 974 References

- 975 [1] A. Avizienis, J.-C. Laprie, B. Randell, C. Landwehr, Basic concepts and taxonomy of  
976 dependable and secure computing, *IEEE Transactions on Dependable and Secure*  
977 *Computing* 1 (1) (2004) 11–33. doi:10.1109/TDSC.2004.2.  
978 URL <http://dx.doi.org/10.1109/TDSC.2004.2>
- 979 [2] S. Mukherjee, *Architecture Design for Soft Errors*, Morgan Kaufmann Publishers  
980 Inc., San Francisco, CA, USA, 2008.
- 981 [3] J. Gray, Why do computers stop and what can be done about it?, Technical Report  
982 TR 85.7, Tandem Computers, [Online; accessed 13-July-2019] (1985).
- 983 [4] M. Grottke, K. S. Trivedi, Fighting bugs: remove, retry, replicate, and rejuvenate,  
984 *Computer* 40 (2) (2007) 107–109. doi:10.1109/MC.2007.55.
- 985 [5] G. Carrozza, D. Cotroneo, R. Natella, R. Pietrantuono, S. Russo, Analysis and  
986 prediction of mandelbugs in an industrial software system, in: 2013 IEEE Sixth  
987 International Conference on Software Testing, Verification and Validation, 2013,  
988 pp. 262–271. doi:10.1109/ICST.2013.21.
- 989 [6] M. Winslett, Bruce lindsay speaks out: On system r, benchmarking, life as an ibm  
990 fellow, the power of dbas in the old days, why performance still matters, heisenbugs,  
991 why he still writes code, singing pigs, and more, *SIGMOD Rec.* 34 (2) (2005) 71–79.  
992 doi:10.1145/1083784.1083803.  
993 URL <http://doi.acm.org/10.1145/1083784.1083803>
- 994 [7] M. Grottke, A. P. Nikora, K. S. Trivedi, An empirical investigation of fault types  
995 in space mission system software, in: 2010 IEEE/IFIP International Conference on  
996 Dependable Systems Networks (DSN), 2010, pp. 447–456. doi:10.1109/DSN.2010.  
997 5544284.
- 998 [8] D. Cotroneo, M. Grottke, R. Natella, R. Pietrantuono, K. S. Trivedi, Fault triggers  
999 in open-source software: An experience report, in: 2013 IEEE 24th International  
1000 Symposium on Software Reliability Engineering (ISSRE), 2013, pp. 178–187. doi:  
1001 10.1109/ISSRE.2013.6698917.
- 1002 [9] R. Chillarege, Comparing four case studies on bohr-mandel characteristics using  
1003 odc, in: 2013 IEEE International Symposium on Software Reliability Engineering  
1004 Workshops (ISSREW), 2013, pp. 285–289. doi:10.1109/ISSREW.2013.6688908.
- 1005 [10] J. S. Bradbury, K. Jalbert, Defining a Catalog of Programming Anti-Patterns for  
1006 Concurrent Java, in: Proceedings of the 3rd International Workshop on Software  
1007 Patterns and Quality, SPAQu’09, 2009, pp. 6–11.
- 1008 [11] J. Duffy, Solving 11 Likely Problems In Your Multithreaded Code, *MSDN Magazine*  
1009 (October 2008).
- 1010 [12] G. Gopalakrishnan, J. Sawaya, Achieving formal parallel program debugging by  
1011 incentivizing cs/hpc collaborative tool development, in: Proceedings of the 1st  
1012 Workshop on The Science of Cyberinfrastructure: Research, Experience, Appli-  
1013 cations and Models, SCREAM ’15, ACM, New York, NY, USA, 2015, pp. 11–18.

- 1014       doi:10.1145/2753524.2753531.  
1015       URL <http://doi.acm.org/10.1145/2753524.2753531>
- 1016 [13] P. C. Mehltz, W. Visser, J. Penix, The JPF Runtime Verification System, [http://](http://www.doc.gold.ac.uk/%7Eemas01sd/classes/jpf_release/doc/JPF.pdf)  
1017 [www.doc.gold.ac.uk/%7Eemas01sd/classes/jpf\\_release/doc/JPF.pdf](http://www.doc.gold.ac.uk/%7Eemas01sd/classes/jpf_release/doc/JPF.pdf) (2005).
- 1018 [14] A. Mahmood, E. J. McCluskey, Concurrent error detection using watchdog  
1019 processors-a survey, *IEEE Transactions on Computers* 37 (2) (1988) 160–174.  
1020 doi:10.1109/12.2145.
- 1021 [15] O. Goloubeva, M. Rebaudengo, M. Sonza Reorda, M. Violante, Soft-error detection  
1022 using control flow assertions, in: *Proceedings 18th IEEE Symposium on Defect and*  
1023 *Fault Tolerance in VLSI Systems, 2003*, pp. 581–588. doi:10.1109/DFTVS.2003.  
1024 1250158.
- 1025 [16] A. Benso, S. Di Carlo, G. Di Natale, P. Prinetto, A watchdog processor to detect  
1026 data and control flow errors, in: *9th IEEE On-Line Testing Symposium, 2003.*  
1027 *IOLTS 2003.*, 2003, pp. 144–148. doi:10.1109/OLT.2003.1214381.
- 1028 [17] A. Li, B. Hong, Software implemented transient fault detection in space  
1029 computer, *Aerospace Science and Technology* 11 (2) (2007) 245 – 252.  
1030 doi:<https://doi.org/10.1016/j.ast.2006.06.006>.  
1031 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S1270963806000800)  
1032 [S1270963806000800](http://www.sciencedirect.com/science/article/pii/S1270963806000800)
- 1033 [18] C. Flanagan, P. Godefroid, Dynamic Partial-Order Reduction for Model Check-  
1034 ing Software, *ACM SIGPLAN Notices - Proceedings of the 32nd ACM SIGPLAN-*  
1035 *SIGACT symposium on Principles of programming languages 1 (40) (2005) 110 –*  
1036 *121.*
- 1037 [19] T. Sharma, D. Spinellis, A survey on software smells, *Journal of Systems and Soft-*  
1038 *ware* 138 (2018) 158 – 173. doi:<https://doi.org/10.1016/j.jss.2017.12.034>.  
1039 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0164121217303114)  
1040 [S0164121217303114](http://www.sciencedirect.com/science/article/pii/S0164121217303114)
- 1041 [20] M. S. Haque, J. Carver, T. Atkison, Causes, impacts, and detection approaches of  
1042 code smell: A survey, in: *Proceedings of the ACMSE 2018 Conference, ACMSE '18,*  
1043 *ACM, New York, NY, USA, 2018*, pp. 25:1–25:8. doi:10.1145/3190645.3190697.  
1044 URL <http://doi.acm.org/10.1145/3190645.3190697>
- 1045 [21] S. S. Rathore, S. Kumar, A study on software fault prediction techniques, *Artificial*  
1046 *Intelligence Review* 51 (2) (2019) 255–327. doi:10.1007/s10462-017-9563-5.  
1047 URL <https://doi.org/10.1007/s10462-017-9563-5>
- 1048 [22] P. Runeson, A survey of unit testing practices, *IEEE Software* 23 (4) (2006) 22–29.  
1049 doi:10.1109/MS.2006.91.
- 1050 [23] V. V. Kuliainin, A. A. Petukhov, A survey of methods for constructing covering  
1051 arrays, *Programming and Computer Software* 37 (3) (2011) 121. doi:10.1134/  
1052 S0361768811030029.  
1053 URL <https://doi.org/10.1134/S0361768811030029>
- 1054 [24] E. P. Enoiu, A. Čaušević, T. J. Ostrand, E. J. Weyuker, D. Sundmark, P. Pet-  
1055 tersson, Automated test generation using model checking: an industrial evaluation,  
1056 *International Journal on Software Tools for Technology Transfer* 18 (3) (2016) 335–  
1057 353. doi:10.1007/s10009-014-0355-9.  
1058 URL <https://doi.org/10.1007/s10009-014-0355-9>
- 1059 [25] A. Salahirad, H. Almulla, G. Gay, Choosing the fitness function for the job: Au-

- 1060 tomated generation of test suites that detect real faults, *Software Testing, Ver-*  
1061 *ification and Reliability* 29 (4-5) (2019) e1701, e1701 stvr.1701. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/stvr.1701>, doi:10.1002/stvr.  
1062 1701.  
1063 URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/stvr.1701>  
1064 [26] A. Schwartz, D. Puckett, Y. Meng, G. Gay, Investigating faults missed by test  
1065 suites achieving high code coverage, *Journal of Systems and Software* 144 (2018)  
1066 106 – 120. doi:<https://doi.org/10.1016/j.jss.2018.06.024>.  
1067 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0164121218301201)  
1068 [S0164121218301201](http://www.sciencedirect.com/science/article/pii/S0164121218301201)  
1069 [27] B. S. Ahmed, Test case minimization approach using fault detection and com-  
1070 binatorial optimization techniques for configuration-aware structural testing,  
1071 *Engineering Science and Technology, an International Journal* 19 (2) (2016) 737 –  
1072 753. doi:<https://doi.org/10.1016/j.jestch.2015.11.006>.  
1073 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S2215098615001706)  
1074 [S2215098615001706](http://www.sciencedirect.com/science/article/pii/S2215098615001706)  
1075 [28] M. Khatibsyarbini, M. A. Isa, D. N. Jawawi, R. Tumeng, Test case pri-  
1076 oritization approaches in regression testing: A systematic literature re-  
1077 view, *Information and Software Technology* 93 (2018) 74 – 93. doi:<https://doi.org/10.1016/j.infsof.2017.08.014>.  
1078 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/S0950584916304888)  
1079 [S0950584916304888](http://www.sciencedirect.com/science/article/pii/S0950584916304888)  
1080 [29] S. M. Ghaffarian, H. R. Shahriari, Software vulnerability analysis and discovery  
1081 using machine-learning and data-mining techniques: A survey, *ACM Comput. Surv.*  
1082 50 (4) (2017) 56:1–56:36. doi:10.1145/3092566.  
1083 URL <http://doi.acm.org/10.1145/3092566>  
1084 [30] W. E. Wong, R. Gao, Y. Li, R. Abreu, F. Wotawa, A survey on software fault  
1085 localization, *IEEE Transactions on Software Engineering* 42 (8) (2016) 707–740.  
1086 doi:10.1109/TSE.2016.2521368.  
1087 [31] N. Bazan, Static and Dynamic Verification Tools, [https://docs.microsoft.com/en-us/windows-hardware/drivers/devtest/static-and-dynamic-  
1088 verification-tools/](https://docs.microsoft.com/en-us/windows-hardware/drivers/devtest/static-and-dynamic-verification-tools/), [Online; accessed 24-March-2019] (2017).  
1089 [32] M. Felleisen, R. Cartwright, Safety as a metric, in: *Proceedings 12th Conference on  
1090 Software Engineering Education and Training (Cat. No.PR00131)*, 1999, pp. 129–  
1091 131. doi:10.1109/CSEE.1999.755192.  
1092 URL <https://doi.org/10.1109/CSEE.1999.755192>  
1093 [33] Wikipedia, Satisfiability modulo theories, [https://en.wikipedia.org/wiki/  
1094 Satisfiability\\_modulo\\_theories](https://en.wikipedia.org/wiki/Satisfiability_modulo_theories), [Online; accessed 24-March-2019] (2019).  
1095 [34] N. Machado, B. Lucia, L. Rodrigues, Production-guided Concurrency Debugging,  
1096 *ACM SIGPLAN Notices - PPOPP '16* 51 (8) (August 2016).  
1097 [35] N. Machado, B. Lucia, L. Rodrigues, Concurrency Debugging with Differential  
1098 Schedule Projections, *ACM SIGPLAN Notices - PLDI 15* 50 (6) (2015) 586–595.  
1099 [36] M. Musuvathi, S. Qadeer, T. Ball, G. Basler, P. A. Nainar, I. Neamtiu, Finding and  
1100 reproducing heisenbugs in concurrent programs, in: *Proceedings of the 8th USENIX  
1101 Conference on Operating Systems Design and Implementation, OSDI'08, USENIX  
1102 Association, Berkeley, CA, USA, 2008*, pp. 267–280.  
1103  
1104  
1105



- 1106 URL <http://dl.acm.org/citation.cfm?id=1855741.1855760>
- 1107 [37] F. Koca, H. Sözer, R. Abreu, Spectrum-based fault localization for diagnosing concurrency faults, in: H. Yenigün, C. Yilmaz, A. Ulrich (Eds.), *Testing Software and Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 239–254.
- 1108
- 1109
- 1110 [38] R. Abreu, P. Zoetewij, A. J. C. v. Gemund, Spectrum-based multiple fault localization, in: *Proceedings of the 2009 IEEE/ACM International Conference on Automated Software Engineering, ASE 09*, IEEE Computer Society, Washington, DC, USA, 2009, pp. 88–99. doi:10.1109/ASE.2009.25.
- 1111
- 1112
- 1113
- 1114 URL <https://doi.org/10.1109/ASE.2009.25>
- 1115 [39] S. Park, S. Lu, Y. Zhou, Ctrigger: Exposing atomicity violation bugs from their hiding places, *SIGPLAN Not.* 44 (3) (2009) 25–36. doi:10.1145/1508284.1508249.
- 1116
- 1117 URL <http://doi.acm.org/10.1145/1508284.1508249>
- 1118 [40] C. S. Păsăreanu, W. Visser, Symbolic execution and model checking for testing, in: K. Yorav (Ed.), *Hardware and Software: Verification and Testing*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 17–18.
- 1119
- 1120
- 1121 [41] E. Clarke, O. Grumberg, S. Jha, Y. Lu, H. Veith, Counterexample-guided abstraction refinement for symbolic model checking, *J. ACM* 50 (5) (2003) 752–794. doi:10.1145/876638.876643.
- 1122
- 1123
- 1124 URL <http://doi.acm.org/10.1145/876638.876643>
- 1125 [42] C. B. Bergersen, Java Path Finder, <https://www.uio.no/studier/emner/matnat/ifi/INF5140/v15/slides/jpf.pdf>, jPF (2015).
- 1126
- 1127 [43] K. Y. Rozier, Survey: Linear temporal logic symbolic model checking, *Comput. Sci. Rev.* 5 (2) (2011) 163–203. doi:10.1016/j.cosrev.2010.06.002.
- 1128
- 1129 URL <http://dx.doi.org/10.1016/j.cosrev.2010.06.002>
- 1130 [44] NASA, Java pathfinder, <https://github.com/javapathfinder/jpf-core>, [Online; accessed 13-July-2019] (2017).
- 1131
- 1132 [45] A. Malkis, A. Podelski, A. Rybalchenko, Precise thread-modular verification, *SAS’07 Proceedings of the 14th international conference on Static Analysis (2007)* 218–232.
- 1133
- 1134
- 1135 [46] NASA, On-the-fly Partial Order Reduction, [http://javapathfinder.sourceforge.net/On-the-fly\\_Partial\\_Order\\_Reduction.html](http://javapathfinder.sourceforge.net/On-the-fly_Partial_Order_Reduction.html), [Online; accessed 24-March-2019] (2009).
- 1136
- 1137
- 1138 [47] R. Alur, R. K. Brayton, T. A. Henzinger, S. Qadeer, S. K. Rajamani, Partial-order reduction in symbolic state space exploration, in: O. Grumberg (Ed.), *Computer Aided Verification*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1997, pp. 340–351.
- 1139
- 1140
- 1141
- 1142 [48] R. Büschkes, M. Borning, D. Kesdogan, Transaction-based anomaly detection, *ID’99 Proceedings of the 1st conference on Workshop on Intrusion Detection and Network Monitoring 1 (April 1999)*.
- 1143
- 1144
- 1145 [49] T. Shpeisman, V. Menon, A.-R. Adl-Tabatabai, S. Balensiefer, D. Grossman, R. L. Hudson, K. F. Moore, B. Saha, Enforcing isolation and ordering in STM, *ACM SIGPLAN Notices - Proceedings of the 2007 PLDI conference* 42 (6) (2007) 78–88.
- 1146
- 1147
- 1148 [50] M. Mansouri-Samani, P. Mehlitz, C. Pasareanu, J. Penix, G. Brat, L. Markosian, O. O’Malley, T. Pressburger, W. Visser, *Program Model Checking: A Practitioner’s Guide*, [https://ti.arc.nasa.gov/m/pub-archive/1439h/1439%20\(Mansouri-Samani\).pdf](https://ti.arc.nasa.gov/m/pub-archive/1439h/1439%20(Mansouri-Samani).pdf), [Online; accessed 11-July-2019] (2012).
- 1149
- 1150
- 1151

- 1152 [51] NASA, Java pathfinder: A model checker for java programs, [https://ti.arc.](https://ti.arc.nasa.gov/tech/rse/vandv/jpf/)  
1153 [nasa.gov/tech/rse/vandv/jpf/](https://ti.arc.nasa.gov/tech/rse/vandv/jpf/), [Online; accessed 11-July-2019] (2012).
- 1154 [52] H. R. Lewis, C. H. Papadimitriou, Elements of the Theory of Computation, 2nd  
1155 Edition, Prentice Hall PTR, Upper Saddle River, NJ, USA, 1997.
- 1156 [53] I. U. Plale, Thread Design Patterns, [https://www.cs.indiana.edu/classes/](https://www.cs.indiana.edu/classes/b534-plal/ClassNotes/thread-design-patterns4.pdf)  
1157 [b534-plal/ClassNotes/thread-design-patterns4.pdf](https://www.cs.indiana.edu/classes/b534-plal/ClassNotes/thread-design-patterns4.pdf), [Online; accessed 26-  
1158 March-2019] (2001).
- 1159 [54] A. S. Tanenbaum, M. van Steen, Distributed systems - principles and paradigms,  
1160 2nd Edition, Pearson Education, 2007.
- 1161 [55] P. Parizek, T. Kalibera, Verifying nested lock priority inheritance in rtems with java  
1162 pathfinder, in: International Conference on Formal Engineering Methods, ICFEM  
1163 2016, Springer, Cham, 2016, pp. 417–432. doi:[https://doi.org/10.1007/978-](https://doi.org/10.1007/978-3-319-47846-3_26)  
1164 [3-319-47846-3\\_26](https://doi.org/10.1007/978-3-319-47846-3_26).
- 1165 URL [https://link.springer.com/chapter/10.1007/978-3-319-47846-3\\_26](https://link.springer.com/chapter/10.1007/978-3-319-47846-3_26)
- 1166 [56] “djessup” GitHub user, Java web server, [https://github.com/djessup/java-](https://github.com/djessup/java-webserver)  
1167 [webserver](https://github.com/djessup/java-webserver), [Online; accessed 29-July-2019] (2016).
- 1168 [57] P. Sotiropoulos, Java web server simulation without injected faults, [https://](https://github.com/pansot2/java-webserver/tree/simulation)  
1169 [github.com/pansot2/java-webserver/tree/simulation](https://github.com/pansot2/java-webserver/tree/simulation), [Online; accessed 29-  
1170 August-2019] (2019).
- 1171 [58] P. Sotiropoulos, Java web server with injected faults, [https://github.com/](https://github.com/pansot2/java-webserver/tree/jpf-simulation)  
1172 [pansot2/java-webserver/tree/jpf-simulation](https://github.com/pansot2/java-webserver/tree/jpf-simulation), [Online; accessed 29-August-  
1173 2019] (2019).
- 1174 [59] S. Dandamudi, Addressing modes, in: T. editor (Ed.), Introduction to Assembly  
1175 Language Programming, Undergraduate Texts in Computer Science, Springer, New  
1176 York, NY, 1998, Ch. 5, pp. 173–206.