

Controlling Factors in Evaluating Path-sensitive Error Detection Techniques

Matthew B. Dwyer, Suzette Person, Sebastian Elbaum

Department of Computer Science and Engineering
University of Nebraska - Lincoln
Lincoln, Nebraska
{dwyer,sperson,elbaum}@cse.unl.edu

Abstract

Recent advances in static program analysis have made it possible to detect errors in applications that have been thoroughly tested and are in wide-spread use. The ability to find errors that have eluded traditional validation methods is due to the development and combination of sophisticated algorithmic techniques that are embedded in the implementations of analysis tools. Evaluating new analysis techniques is typically performed by running an analysis tool on a collection of subject programs, perhaps enabling and disabling a given technique in different runs. While seemingly sensible, this approach runs the risk of attributing improvements in the cost-effectiveness of the analysis to the technique under consideration, when those improvements may actually be due to details of analysis tool implementations that are uncontrolled during evaluation.

In this paper, we focus on the specific class of path-sensitive error detection techniques and identify several factors that can significantly influence the cost of analysis. We show, through careful empirical studies, that the influence of these factors is sufficiently large that, if left uncontrolled, they may lead researchers to improperly attribute improvements in analysis cost and effectiveness. We make several recommendations as to how the influence of these factors can be mitigated when evaluating techniques.

1. INTRODUCTION

Static program analyses calculate information about the executable behavior of a program without running the program. Traditionally, static analyses have been formulated to provide guarantees about program behavior to support, for example, semantics-preserving code transformations to improve performance. Such analyses must necessarily account for all possible program behaviors. In practice, this requirement forces analysis developers to formulate relatively imprecise analyses to achieve scalability to real programs.

It is also possible to formulate static analyses explicitly to detect errors and issue diagnostic information to users. The intuition behind such approaches is that a static analysis can be engineered to more efficiently cover a broader space of program behavior than can be achieved through testing and, consequently, such analyses have the potential to detect *hard to find* errors. Because they need not account for all possible behaviors, analysis developers have exploited this relaxed requirement to customize very precise path-

sensitive analysis frameworks, for example model-checking tools like Bogor [23], Java Path Finder (JPF) [27], Mur ϕ [3], and Spin [17], to make them cost effective for error detection.

Several recent efforts along these lines involve adaptations of the CMC [20] model checker to make it more effective for finding errors in certain kinds of applications. In [19], the authors adapt CMC for error detection in protocol implementations and have used it to find four errors in the Linux TCP/IP implementation. More recently they have developed FiSC [28], a version of CMC that has been adapted for and used to reason about file system implementations; several significant errors in three widely-used file system implementations have been detected using FiSC. These adaptations have been carefully tuned to use specific heuristics for selectively storing only part of a program's data state during analysis and for prioritizing the order in which statements are analyzed.

Results such as these provide an important *proof of concept* that cost-effective and precise path-sensitive analyses for error detection can be built. They demonstrate that there exists a combination of specific techniques that can provide cost-effective analysis for a specific class of programs. They do not, however, provide information about the relative cost-effectiveness of the individual analysis techniques that they are comprised of nor about the range of programs over which they are effective. For example, two very different heuristics for prioritizing exploration of transitions are described in [19], preferring exploration of new behaviors relative to protocol states and preferring infrequent state changes, but information about the breadth or relative effectiveness of the heuristics is not provided. To be fair, this was not the goal of the authors, but it is important to gain this kind of information so that techniques such as these can be selected and combined for maximum benefit in path-sensitive error detection tools.

Obtaining such results requires careful empirical evaluation of techniques used to achieve cost-effective analysis across a range of programs. This kind of evaluation can be difficult to perform especially when there is a lack of knowledge about the factors that can influence the performance of an analysis tool. In this paper, we take a first step towards enabling controlled empirical studies of path-sensitive static analyses by presenting data on two factors that can *significantly* influence the performance of path-sensitive error detection analyses, to the extent that, if uncontrolled their influence may obscure differences in performance that are attributable to analysis techniques.

The first factor is related to the implementation of path-sensitive analysis algorithms. Most analysis algorithms allow the execution of a program to be under-specified in some way. When analyzing a multi-threaded Java program, for example, the analysis may reach a point where bytecodes in two different threads are enabled for execution. In a JVM, the thread scheduling algorithm will choose

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE Controlling Factors in Evaluating Path-sensitive Error Detection Techniques				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Nebraska - Lincoln, Department of Computer Science and Engineering, Lincoln, NE, 68588				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

one of those bytecodes to execute first, but path-sensitive analysis techniques generally abstract from thread scheduling algorithm details and simply require that each of the schedulings is analyzed. Path-sensitive analysis tools, such as Bogor, Mur ϕ , SPIN, and JPF, implement a specific *default search order* for exploring simultaneously enabled execution steps; in fact, these four tools each implement different default orders. Given that these tools were built to exhaustively analyze all possible program paths the specific default order used was not a concern to their developers. When targeting or customizing such tools to detect errors, however, *we show that variation in search order can give rise to very large variations in path-sensitive analysis cost and fault detection effectiveness across a range of programs*. In Section 3, we support this conclusion with a *retrospective study* that looks back at previously published results and relates them to results from empirical studies we performed.

The second factor is related to the subject programs used to evaluate the cost-effectiveness of path-sensitive analysis techniques. The literature contains many papers that introduce analysis techniques and *illustrate* the performance of those techniques on a few small selected examples, for example, dining philosophers and bounded buffer examples [1, 21, 6]. Recent efforts to establish benchmarks to support the evaluation of testing and analysis techniques for multi-threaded Java programs are focused on making a broader collection of examples available to the community [11, 10]. One thing lacking from the literature and emerging benchmarks is a meaningful characterization of the programs and the faults they contain, and the criteria for their inclusion in the benchmark. Such characterization and criteria would help researchers determine whether the benchmarks are appropriate to assess their particular techniques and, if they are appropriate, it would help them put their findings into perspective leading to claims that are substantiated in the collected data set. More specifically, for evaluating path-sensitive error detection tools we are interested in understanding whether programs contain *hard to find* errors. To address this, we characterize programs in terms of *path error density* - the percentage of program paths that contain an error. Surprisingly, many of the examples currently used in evaluations have very high path error densities which makes them poor subjects for evaluating the merits of path-sensitive analysis techniques. Furthermore, *we show that path error density is a key factor in exposing path-sensitive analysis cost tradeoffs*. In Section 4, we describe a case controlled study that supports this conclusion.

To enable quantitative exploration of these issues, we set our work in the context of path-sensitive analysis of multi-threaded Java programs for detecting safety property violations. We use the JPF Java model checker as the basis for our evaluation of the influence of the factors described above.

We believe that our findings provide important information that can guide the evaluation of path-sensitive analysis techniques. The contributions of the paper lie in: (i) the identification of default search order as a factor that can impact the performance of path-sensitive error detection techniques; (ii) the identification of path error density as a program factor that can impact the performance of path-sensitive error detection techniques; (iii) the results of empirical studies that indicate the frequency and magnitude of performance variation with these factors (Sections 3 and 4); and (iv) recommendations for how to control for the effects of these factors in experimental studies (Section 5).

2. BACKGROUND

This Section gives an overview of the class of path-sensitive analyses realized by multi-threaded Java model checkers, explains the sources of search order variation in such analyses, and charac-

```

1  seen := {s0}
2  push(stack, s0)
3  DFS(s0)

DFS(s)
4  workSet := enabled(s)
5  for each  $\alpha \in workSet$  do
6    s' :=  $\alpha(s)$ 
7    if error(s') then
8      counterexample := stack
9    exit
10   if s'  $\notin$  seen then
11     seen := seen  $\cup$  {s'}
12     push(stack, s')
13     DFS(s')
14   pop(stack)
end DFS()

```

Figure 1: Depth-first search for first error state

terizes the availability of subject multi-threaded Java programs for experiments with path-sensitive error detection tools.

2.1 Path-Sensitive State-Space Search

Many path-sensitive analysis techniques treat programs as guarded transition systems and analyze program behavior via depth-first traversal of transition sequences rooted at the initial state¹. A *guarded transition system* consists of a set of variables, which for our purposes are coalesced into a single composite *state variable* s , and a set of guarded transitions which atomically test, with predicate ϕ , the current state and update the state by executing a transition α , i.e., if $\phi(s)$ then $s = \alpha(s)$. The initial values of program variables are used to define an initial state, s_0 . Figure 1 presents the DFS analysis algorithm. On line 4, $enabled(s)$ returns the set of transitions, α , whose guard, ϕ , is true in the given state. Lines 7-9 test if an error state has been reached, and if so, records the current DFS stack, which encodes the path under analysis, as a counterexample and exits. Even though this analysis does not generate all program paths, it is path-sensitive since it reasons about paths and prefixes of paths independently; a DFS can be thought of as analyzing all acyclic program paths.

Line 5 of this algorithm imposes no order on iterating through the set of enabled transitions in a state. This is an issue if non-singleton sets are produced at line 4 which is actually very common in analyzing realistic programs. This may seem odd since program execution is usually thought of as deterministic for a given execution environment, i.e., sequence of inputs and scheduling decisions. In practice, path-sensitive analyses must perform significant abstraction to gain tractability and to produce results that generalize across multiple specific execution environments.

For example, the past decade has seen a significant amount of work on predicate abstraction [13] which replaces reasoning about specific variable values with sets of values encoded symbolically. These abstractions encode approximations using non-determinism. For example, if a variable is abstracted by predicates $x < 0$, $x == 0$, $x > 0$ then the result of executing a statement $x = x - 2$ in a state in which $x > 0$ could result in any of the three predicates being true and is therefore modeled with an enabled transition for each resultant predicate.

In reasoning about a multi-threaded Java program, a path-sensitive analysis must reflect the possible scheduler decisions that could be

¹We note that all of the search order issues we discuss for DFS are also relevant for other forms of search, such as breadth-first and variants of breadth-first.

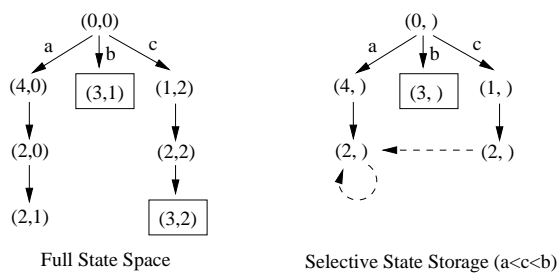


Figure 2: Search Order Example

executed by a JVM. The JVM specification provides only a weak specification of scheduler behavior, i.e., higher priority threads should be scheduled first. It says nothing about the order in which threads at a given priority level should be executed; in fact, it does not even require that threads be executed fairly. All JVMs implement a scheduling policy and that policy calculates a thread ordering. To provide a degree of JVM independence, path-sensitive analyses for multi-threaded Java programs use non-determinism to over-approximate the set of all legal JVM scheduler decisions. Transitions in multiple threads will typically be enabled in a given state and returned on line 4 in the algorithm above. Non-determinism has other uses in tools such as JPF as well; for example, to generate a range of input values [14] or to encode complex specifications [2].

These sources of non-determinism combine to produce a large space of possible paths through the transition system. When performing an exhaustive traversal of the transition system states, the order in which transitions are executed at line 4 is of no consequence; the algorithm is guaranteed to visit all reachable system states regardless of order. Clearly, when an error exists order may matter and intuitively the presence of non-determinism is the key factor in determining the number of paths to be traversed; a deterministic system has a single path. Figure 2 illustrates a sample state space where search order on transitions from the initial state $(0, 0)$ can vary the cost of finding an error state, measured in the number of states explored, from 2 (under the order $b < a < c$) to 7 (under the order $a < c < b$).

2.2 Heuristic Search of Program State-Space

Recent years have seen a growing interest in the incorporation of heuristics into path-sensitive analyses. Heuristics are typically designed to calculate a search order that will reach an error state quickly, e.g., [14], or to reach a particular kind of goal state, e.g., one with a short counter-example [8, 26]. We note also that heuristics can choose to completely drop transitions from consideration [19], which clearly changes the search order.

Heuristics can function in several ways. In traditional heuristic search, one applies a *cost* function to map each enabled transition, α , to a value and then implements line 5 of Figure 1 to iterate over the transitions in cost-order. Cost functions usually calculate a score based on the current state in the exploration or based on the path explored up to the state, for example, Groce and Visser’s [14] “demonic” scheduling heuristic scores thread transitions based on the frequency of thread execution along the path.

It is important to observe that using this type of heuristic does not completely eliminate the issue of transition order. If two transitions evaluate to the same cost value, a common situation when using discrete cost functions [14], then the order in which those transitions execute is left to the stability of the transition sorting algorithm used to implement line 5. A stable sort will honor the underlying order implemented in the model checker, whereas an

unstable sort may modify it. Thus, two different heuristic path-sensitive analysis implementations that use the same cost function may actually explore the paths in the system in a different order.

Another heuristic approach is to be selective in storing the program state. This is realized by modifying line 10 of Figure 1 so that the membership test is not performed on the complete state, s , but rather on a projection of the state $\pi(s)$, and projected seen values, $\{\pi(s') \mid \exists s' \in \text{seen}\}$. For example, Musuvathi and Engler [19] drop a variable from the state if it has been assigned a large number of distinct values on the path explored. This has the effect of forcing backtracking in the DFS earlier than would happen without this modification. In doing this, the search order may change since the continuation of the current path in the original DFS is either eliminated from the search or deferred until later in the search.

Figure 2 illustrates a selective state storage strategy where the second state component is dropped. Under the order $a < c < b$ one can see that the traversal of the state space is curtailed prematurely, potentially reducing the analysis cost, but in this case it forces the error state along the path beginning with transition c to be missed due to matching on the partial state $(2,)$ - the net result is it still requires exploration of 7 states to find an error.

Heuristics are viewed by many as a promising mechanism for mitigating the combinatorial explosion in the cost of path-sensitive state-space analyses. They have the effect of focusing the search on a portion of the state space. Unlike property-preserving state-space reductions, heuristics are oriented towards error detection and the only valid means of evaluating them is through broad experimentation across a variety of programs and properties. Consequently, *we believe that evaluation of heuristic state-space search techniques is especially vulnerable to a lack of control on experimental factors.*

2.3 Java State-Space Analysis Tools

In our studies, we evaluated the performance of path-sensitive analysis tools for multi-threaded Java over a range of subject programs. The primary tool we considered was JPF; Bandera, using Bogor as its model checking engine, will be considered in our replicated studies. We chose these tools since they are the most mature and sophisticated path-sensitive multi-threaded Java analysis tools that we are aware of, both can be applied to a range of Java programs, both provide flexible implementations that make it easy to modify the transition order, and, finally, we are familiar with both tools and are in close contact with their developers so that we can ensure any modifications we make to them are correct.

JPF is built as a virtual machine that stores the set of states visited along a path and uses that state set to force backtracking along other paths as in the general DFS algorithm above. JPF processes JVM bytecode programs directly and as a result it relies on a Java compiler to translate source programs. JPF has a flexible architecture that makes heavy use of Java interfaces and abstract classes to consolidate common functionality in the model checker and enable different algorithms and data structures to be used.

The `Search` interface defines the generic API used by the main analysis module. Multiple search modules are included with JPF including: `DFSSearch` - depth-first search, implementations of all of the heuristics described in [14], and `RandomSearch` - a stateless search that explores a single path in the program. The `Scheduler` abstract class is the base class used to define the strategy for calculating the order in which enabled threads are analyzed; this corresponds to line 5 in Figure 1. JPF includes two `Scheduler` subtypes: `DefaultScheduler` - which implements a fixed strategy for selecting the *next* thread based on the order in which `Thread` (or `Runnable`) objects are allocated along the path being analyzed, and `RandomOrderScheduler` - which randomizes the

order of thread selection based a pseudo-random sequence determined by the current time or a user-specified seed value.

An additional degree of *non-determinism* can be specified in JPF programs, for example, a call to `Verify.random(3)` returns one of $\{0, 1, 2, 3\}$; in a complete stateful search all such values are guaranteed to be returned. The order in which these values are produced is up to the implementation of this method; the current implementation produces them in their value order. Calls such as this create multiple enabled transitions *internal* to a single thread.

Version 3.1.2 of JPF does not allow for randomization of the order in which such internal transitions are explored; a forthcoming version of JPF will have this feature. Given this, the results we report for analyses using JPF with `RandomOrderScheduler` should be interpreted as exploring a subset of all possible random orders. The Java programs we consider in our studies have very limited internal non-determinism, so while we believe that additional variation in search order may be possible when running JPF on those example we regard it as a minor effect. Even if that were not the case, the results we report on the variation in performance of JPF due to search order can be regarded as a lower-bound.

The only modifications we made to JPF were to allow for time-bounded analysis, a command-line parameter specified the maximum number of seconds an analysis may execute, and the reporting of statistics on partial searches that are terminated either when the time bound is reached or memory is exhausted. Neither change affects the path-sensitive analysis implementation.

As a final note, Version 3.1.2 of JPF has not, to the best of our knowledge, been used in any published study of error detection techniques. Previous papers reporting JPF results, such as [21, 14], used older versions that were missing important advances in mitigating the cost of path-sensitive analysis. For example, canonical heap symmetry reductions, which represent all execution states of a Java program that differ only in the physical addresses of objects or in the unreclaimed garbage using a single representative state, e.g., [24] are implemented. In addition sophisticated partial order reductions that are customized for multi-threaded Java programs [7], have also been adapted to JPF and implemented. The cost-effectiveness of these reductions is sufficient to regard the use of these features as the default configuration of JPF. All of our runs use these features. As a consequence, even on the same Java program the performance measures we report may differ significantly from those reported in previous studies.

2.4 Default Transition Search Order

Default transition order is an *implementation detail* that is typically realized in a way that is convenient given the internal path and state representations maintained by an analysis tool. It is no surprise then that the default transition order varies between tools. In fact, each of Spin, Mur ϕ , JPF and Bogor use a different default order to select which enabled thread will be analyzed next. Spin selects the next thread based on the order in which `active` proctypes, i.e., Spin's notion of thread, appear in the source file and then in the order in which dynamically started threads are created. Mur ϕ orders threads in the reverse order of their appearance in the source file. JPF and Bogor are very similar; JPF uses the strategy described above for the `DefaultScheduler`, whereas Bogor uses the order in which the thread `start()` method is invoked. Despite the close similarity in implementations, certain program structures can give rise to dramatically different default orders. For the following code:

```
Thread[] threads[3];
for (int i=0; i<3; i++) threads[i] = new Thread();
for (int i=0; i<3; i++) threads[3-i].start;
```

JPF and Bogor will explore the threads in opposite orders.

2.5 Common Multi-threaded Java Subjects

The past decade has seen significant advances in the development of path-sensitive program analyses. There have, however, been relatively few broad evaluations of the cost-effectiveness of those techniques across a range of systems; a notable exception is Corbett's study of deadlock detection in 30 different multi-tasking Ada programs [1].

For analyzing multi-threaded Java programs, researchers have had two basic choices: (1) adapt existing examples from the concurrency literature, or (2) target real multi-threaded Java programs. Most have chosen option (1), since it can be difficult to obtain faulty versions of real multi-threaded Java programs, even from open source projects, and because using a small set of well-understood examples would seem to provide a means for comparing performance across tools and techniques.

We obtained the suites of programs compiled by the Bandera [6] and JPF projects. These programs cover nearly all of the multi-threaded Java programs that are used in evaluating path-sensitive Java analyses in the literature; some papers have also used Java standard library implementations as analysis subjects. We selected programs that exhibit some type of concurrency error. These examples can be divided into two kinds: concurrency error *kernels* and *realistic* program structures. Concurrency error kernels are very simple programs that distill the essence of a particular concurrency error. Examples include adapted versions from the concurrency literature, such as dining philosophers, as well as programs that exhibit Java-specific errors; we had a student independently implement kernels for the Java *concurrent bug patterns* (CBP) described in [12]. These kernels typically include the control and data structures required to exhibit the error and nothing else. Real programs are small to medium size programs that perform a computation over rich data structures. They tend to be much larger than the concurrency kernels, have a higher degree of multi-threading, often accept input data that parameterizes the computation, and include significant control and data structures that are unrelated to the error.

We were also granted access to a collection of multi-threaded Java programs being developed at IBM to support testing and analysis research [11]. This set of programs overlaps with the Bandera and JPF programs to some extent, but it also includes a number of programs that were developed to encode common Java concurrency bug patterns; we refer to those programs as the IBM *benchmarks*. Many of the IBM benchmarks were written following standard forms for parameterization, e.g., the degree of multi-threading in an example was indicated by a string `little`, `average`, or `lot`, for error reporting, e.g., error messages were printed to a log file, and for perturbing the schedules so as to make errors more difficult to detect by testing, e.g., by inserting random `sleep()` calls. We transformed these examples to further parameterize them, e.g., programs accept an integer that indicates the degree of multi-threading, to indicate errors through assertion violations, and to remove `sleep()` calls; this last step is easily automated. The resultant versions of the IBM benchmark programs we considered are behaviorally equivalent to the original versions.

Table 1 lists each subject program and describes its source, kind (i.e., kernel, realistic, benchmark), parameters, the errors it contains - including CBP designations where appropriate, application class counts, non-comment source lines of code as calculated by JavaNCSS [18], and path error density measures which are discussed in Section 4. We note that the analysis of these programs considers all of the library code used by the applications which can significantly increase analyzed program size.

Subject	Source	Reference	Kind	Parameters	Error	Classes	SLOC	Density
Account	IBM	[10]	benchmark	none	Deadlock, Race	3	66	11.4-66.3%
Airline	IBM	[10]	benchmark	#ticketsIssued, cushion	Race	2	31	65.7-83.6%
AlarmClock	Bandera	[6]	real	none	NullPtrExcpn	6	125	23.2%
AllocateVector	IBM	[10]	benchmark	blockSize, vectorSize, #runs	No Lock	3	85	3.8-75.0%
BoundedBuffer	Bandera	[21, 6]	real	bufferSize, #producers #consumers, modCount	Deadlock	5	65	0-99.6%
Clean	CBP	[12]	kernel	#firstTasks, #secondTasks, #iterations	Deadlock	4	51	0.9-100%
Daisy	other	[22]	real	none	Assert	21	744	0.07%
Deadlock	Bandera		kernel	none	Deadlock	4	24	63.8-75.4%
DEOS	JPF	[14]	real	none	Assert	24	838	0-41.5%
DiningPhil	Bandera	[14]	kernel	#forks/philosophers	Deadlock	3	25	100%
Elevator	other	[9]	real	none	ArrayIdxOOBExcpn	12	934	0%
LinkedList	IBM	[10]	benchmark	#builders, maxSize	Atomicity	5	117	100%
LoseNotify	CBP	[12]	kernel	#waitThreads, #notifyThreads, #iterations	Deadlock	4	41	100%
NestedMonitor	Bandera		kernel	none	Deadlock	6	53	100%
Piper	IBM	[10]	benchmark	#seats/passengeRequests, #passengers, queueCapacity	Lose Notify	2	71	7.3-33.6%
ProducerConsumer	Bandera		kernel	#producers, #consumers #itemsProduced	Race	8	87	14.6-33.8%
Reorder	CBP	[12]	kernel	#setThreads, #checkThreads	Atomicity	4	44	0-0.02%
ReplicatedWorker	Bandera	[21, 6]	real	#workers, #items, min, max, epsilon	Deadlock	14	304	26.1-70.3%
RaxExtended	JPF	[21, 14, 6]	real	gc, wc	Race (Assert)	11	127	76.1-79.4%
RW	Bandera	[21, 6]	real	#readers, #writers, bound	Race (Assert)	6	103	43.3-49.5%
SleepingBarber	Bandera	[6]	kernel	none	Deadlock	4	66	100%
TwoStage	CBP	[12]	kernel	#twoStageThreads, #readThreads	Two-stage (Assert)	5	52	1.2-1.9%

Table 1: Subject Program Descriptions

3. DOES ORDER REALLY MATTER?

It seems obvious that search order *can* matter, but the real questions are: (a) *Can search order cause performance to vary enough to affect the conclusions of carefully performed evaluations?* and (b) *Does order matter across a range of programs?* (as opposed to toy examples like the one shown in the previous section). In this section, we provide anecdotal evidence regarding question (a) and then followup with a broader study to assess question (b).

3.1 An Anecdote

Even if every path-sensitive analysis of every program were susceptible to the influence of varying order, we would not be concerned if that influence was small. Path-sensitive analyses can be quite expensive and the community expects that techniques of practical importance will yield improvements of practical significance; we need not be too concerned with small scale effects.

To assess question (a), we considered the results of an existing carefully performed study that compared four path-sensitive analysis tools on a number of versions of models of the GNU implementation of the UUCP i-Protocol [5]. We downloaded the i3 and i4 Mur ϕ models from the author’s web-site and configured them to be 2fn models as described in their study. The purpose of their study was to understand the performance improvements that could be achieved by applying five different abstractions to the model. The authors were able to order the abstractions based on analysis performance in finding errors in versions of the system; they also considered performance in showing the absence of errors in versions that were free of errors. They determined that i4 was uniformly faster than i3.

We modified the implementation of Mur ϕ version 2.70L to randomly choose the order in which enabled transitions are explored during the search of system paths for violations of safety properties; the random order was seeded by system time when the analysis was initiated. We executed Mur ϕ on the models using the default search

order of the tool and 20 different randomized orders. Since we ran on a much faster platform, the execution times are not comparable to the results published in [5], so we compared the default runs we performed to the randomized runs using the state count measure used in original the study. We found that the default run for i4 explored 218 states and that there was an order of search for i3 that only explored 87 states. Similarly, the default order for i3 explored 397 states and there was an order for i4 that explored 703 states. The point is that the variation in performance due to search order for Mur ϕ on these problems is sufficiently large so that, in some cases, the conclusions about the cost-effectiveness of abstractions i3 and i4 that were originally drawn by considering default orderings would be inverted.

Based on this limited experience, we conjecture that variations in search order may invalidate the results of otherwise carefully performed evaluations of path-sensitive error detection techniques.

3.2 A Retrospective Study

To evaluate this conjecture more broadly, and thereby address questions (a) and (b), we divide this study into two parts. First, we quantify the variation in performance we observed when running JPF to find an error using randomly chosen search orders on a set of selected subjects. In the second part of our study, we relate the variation we observed on subjects utilized in previous studies back to the results reported in the those studies.

3.2.1 Dependent Variable

In this study, we measure the dependent variable in terms of the number of *new program states* explored during the analysis. This measure is commonly used in model checker evaluations. It is also system independent, making it possible to compare analysis performance across platforms which for our study is crucial since we do not have access to the execution platforms used in previously reported studies.

3.2.2 Independent Variable

Our study manipulated one independent variable: the search order. Given the differences we encountered in default search order across existing path-sensitive analysis tools, we believe it would be difficult to characterize the space of all *implementable* orders that an analysis tool developer might choose. Consequently, we chose to randomize the search order. For each subject, we executed JPF configured with depth-first search, using the `DFSearCh` component, and we selected either the `DefaultScheduler` or the `RandomOrderScheduler` as described in Section 2.3.

3.2.3 Study Design and Setup

For both parts of this study we used the following subjects from Table 1: `AlarmClock`, `DEOS`, `DEOSAbstracted`, `DiningPhil`, `ReplicatedWorker`, `RaxExtended`, `RW`, and `SleepingBarber`. Each of these programs appeared in one or more of [21, 14, 6].

In the first part of our study, we also included all non-kernel subjects from Table 1 in order to assess the variation in analysis cost on the more *realistic* subjects. Since these examples were not the subject of previous studies we do not include them in the retrospective part of our analysis.

Given the large number of subject programs to choose from, we also considered a set of secondary factors when selecting our subject population. These factors include the size of the program (in terms of lines of code and thread counts), the type of concurrency fault, and coverage of the main, non-trivial codes bases that we are aware of and have access to. In all cases, our goal was to choose a variety of programs to increase the diversity of our subject population.

In the second part of this study, we relate the variation in observed performance of JPF back to the results reported in the three selected studies by calculating the ratio of the analysis cost for the technique considered in the study to the default analysis cost cited in the study. For example, [6] reports data on the performance of an analysis when program slicing is applied to a program. Specifically, the performance of an analysis with slicing as a pre-processing step (*slice*) is compared to the default call-graph-based reachability pruning pre-processing step. In this case, we would calculate the ratio of analysis cost using the *slice* technique to the cost of using the default technique. For our study, we calculated the ratios for all of the techniques and subject programs used in the three previous studies.² We then selected the largest ratio for each subject program and applied it to *scale* the default performance of JPF that we observed in the first part of the study; we term this the *scaled default* performance. Finally, we compared the scaled default performance to the variations in performance we observed over the range of orders considered in our study. Orders that result in lower analysis cost than the scaled default performance indicate that using a different default search order in the original study could have possibly led the authors to draw a different conclusion about the benefit of the technique they studied. In such cases, we say that the variation in performance due to search order is of *practical significance*.

To perform this study and the follow-on study discussed in Section 4, we compiled all subjects using Java v1.4.2.07 and then model checked each subject using JPF v3.1.2 with partial order reductions enabled. The study was performed on a cluster of dual-Opteron 250's running at 2.4 GHz with 4 GByte of memory and running Fedora Core 3 Linux. Each subject was model checked one time using JPF's `DFSearCher` and then model checked 500 times using JPF's `RandomOrderSearCher` using system time

²All such ratios are less than or equal to one in the studies we considered since the techniques all represent *improvements* over the default.

as the seed. JPF's execution time was limited to one hour of wall time for all runs and all subjects; this limit exceeds most of the time-bounds used in previous studies (except for [6]). Statistics for each run of JPF were gathered from the output sent to standard output and standard error by JPF. For the two examples (`Account` and `RW`) containing multiple errors, we used JPF's ability to distinguish between certain types of property violations and enabled JPF to look for each type of property violation on separate model checks of that subject.

3.2.4 Results and Analysis

We start by providing a comparative summary of the performance of search orders in terms of new states visited for the 16 subjects we considered in this study. This summary is presented in Table 2 which includes the values for the default search order, and the min, max, average and 95% confidence interval for the random order search based on the 500 observations collected for each subject. We also include information for both the default and random runs indicating if the error is found. *OM* indicates the default search order ran out of memory, *TO* indicates the run timed-out, and *Num. Random* represents the number of random runs (out of 500) that found the error. Note that where a default search did not find the error, the number of states listed can be considered a lower bound on the number of states that would be traversed if the search had been allowed to run to completion. The final two columns in Table 2 show the technique and analysis cost ratio calculated based on the results in the previously published studies for those subjects. Note that the names used in this table, and in subsequent text, are the program name from Table 1 with a '.' separated list of parameter values for the program following a '-'.

When observing the reported number of new states traversed by the different search orders, we first note the great variability across subjects; the number of new states reported ranges from the tens to the millions. This is probably not surprising given that our subjects vary significantly in the number of lines of code, number of threads and the general complexity of their control and data structures. We can also see that the default search order visits more states than the average random search in 10 of the 16 subjects. Interestingly, none of the default runs reported new states within the 95% confidence interval computed based on the 500 random runs.

Variability in results between the default search order and the random order for a given subject is also of interest. For example, the default search order for `AllocateVector-2.100.1` explores over 20 million states without finding the error while all of the random runs find the error in an average of 133,908 states (in fact, all of the 500 random order searches find the error in under .5 million states and in as few as 43 states). On the other hand, the default search order for `RW-2.2.100NoDeadLckCk` finds the error in 3.1 million states versus an average of 13.6 million states for the random runs, of which only 214 of 500 find the error. Perhaps, even more important than the number of states traversed is whether a search actually succeeds in locating the error when resources are bounded. Interestingly we note that in the six subject where the default search order does not find the error, at least 8% of the corresponding random searches for each program finds the error, and in half of these programs, over 80% of the random searches find the error.

Clearly, researchers utilizing programs reporting a smaller number of new states such as `AlarmClock` are less likely to be able to discriminate or expose the potential of their error detection technique. However, targeting programs with a larger state space is also challenging because of the variability we observe with the analysis of such programs. For example, for `RW-2.2.100NoDeadLckCk`, the 95% confidence interval around the mean covers a range of over

Subject	New States					Found		Retrospective	
	Default	Average Random	95% CI Random	Minimum Random	Maximum Random	Default	Num. Random	Technique	Ratio
Account-NoExcepCk	34484	1097737	±168942	92	8393836	OM	86		
Account-NoDeadLckCk	1034	695726	±128737	257	8041901	✓	495		
Airline-20.2	6942672	298274	±101032	148	4842487	TO	469		
AlarmClock	290	264	±8	83	438	✓	500	slice[6]	0.78
AllocateVector-2.100.1	20352435	133908	±11448	43	484132	TO	500		
BoundedBuffer-2.4.4.1	802542	27272	±2167	147	132771	✓	500		
Daisy-4.3	101816	70356	±2750	35256	99774	✓	500		
DEOS	260039	506816	±22712	103065	1213707	✓	500	A*[14]	0.32
DEOSAbstracted	54	783	±67	7	2877	✓	500	choose-free[21]	0.65
DiningPhil-8	1652	38	±1	17	131	✓	500	most-blocked[14]	0.034
LinkedList-4.100	60964	7972	±55	5866	9448	✓	500		
Piper-2.16.8	6377527	3739112	±95781	402	4905845	TO	43		
ReplicatedWorker-5.2.0	6231840	4670826	±186362	485	6485120	TO	85	choose-free[21]	1
RaxExtended-4.3	3470398	195492	±33106	46	3108548	OM	497	interleave[14]	0.41
RW-2.2.100NoDeadLckCk	3147356	13619036	±1036042	49	24225772	✓	214	choose-free[21]	0.0026
SleepingBarber	36	34	±1	28	39	✓	500	slice[6]	0.87

Table 2: Comparative summary of random versus default search strategies.

one million new states.

These observations attest to the degree of variability observed when program paths are traversed in different orders, and they also emphasize the importance of properly qualifying findings when evaluating a path-sensitive error detection technique relative to the single default order implemented in an analysis tool, since, as mentioned Section 2, default order varies across tool implementations.

Table 2 provides details on the technique selected for the retrospective study and the cost ratio for that technique relative to the default. Figure 3 shows histograms for six of eight subject programs relating the scaled default analysis performance to the performance of 500 randomized analysis runs. Two subjects are not shown: DEOS, discussed below, and DEOSAbstracted since it is very similar to the plot for SleepingBarber. The x-axis ranges from the minimum to the maximum number of new states across all analyses and is partitioned into regions. The y-axis shows the percentage of the 500 random runs whose performance lies in each of the regions. The dashed line is the scaled default performance for the selected technique.

For one subject, DEOS, the techniques previously reported improved analysis enough relative to the default that they overwhelmed any variation in cost due to randomization observed in our study. For three of the subjects, DiningPhil-8, RaxExtended, and ReplicatedWorker, more than 86% of the 500 random order searches are classified as having practically significant variation from the scaled default value. For the remaining four subjects, the percentage of practically significant orders varies from 7% to 40%.

We believe that these findings tell a strong cautionary tale. Even for carefully planned and conducted studies of path-sensitive error detection techniques, failing to account for the influence of default search order exposes researchers to the possibility that the reported benefits of techniques are attributable to default search order rather than the technique itself.

4. WHAT PROGRAMFACTORS INFLUENCE ANALYSIS COST?

In this Section, we address the question: *What characteristics of programs cause significant increase in the cost of path-sensitive error detection techniques?*

4.1 Two Candidate Factors

The model checking community, in general, believes that there is a strong correlation between the number of threads in a concurrent program and the number of reachable program states. Holzmann, the author of the SPIN model checker, notes that: “In the worst

case, the global reachability graph has the size of the Cartesian product of all component systems. ... Although, in practice, the size of the global reachability graph never approaches the worst case size, the reachable portion of the Cartesian product can also easily become prohibitively expensive to construct exhaustively.” [17]; a “component system” in SPIN is analogous to a thread in Java. While careful studies of the cost of model checking and related path-sensitive analyses are rare, the few that exist, such as Corbett’s [1], suggest that in practice, analysis cost grows exponentially with the number of threads.

Our second factor attempts to capture the intuitive notion of *hard to find* bug that is often used to characterize the sweet spot for path-sensitive error detection techniques, e.g., [20, 28]. Testing researchers have studied a number of measures for characterizing the ease with which a fault in a program can be revealed. For example, Hamlet and Voas [15] defined the notion of program testability as “... the probability that if P contains fault(s), P will fail under test.” In programs with high testability, faults are likely to be revealed as failures, while programs with low testability are unlikely to expose their faults. Unfortunately, the existing body of work on sensitivity analysis and testability has not explicitly considered concurrent or multi-threaded programs. In such programs there is an additional *input* that can lead to faults being exposed or hidden - the thread schedule. Rather than extend existing testability notions to account for scheduling decisions, we fix the program inputs to isolate thread scheduler decisions as the only varying *input* to the program under analysis. By sampling the possible program paths and checking for errors, we generate an estimate of the percentage of schedules that exhibit failures thus, producing a surrogate-measure that we call *path error density*.

Given this context, *we conjecture that the number of threads used during program execution and the path error density of a program are important factors in determining the cost of path-sensitive error detection techniques.*

4.2 A Case Controlled Study

To evaluate this conjecture, we performed a case controlled study. In this type of study, a researcher identifies groups of subjects with different characteristics that are of interest, and then analyzes the relationship between their characteristics and one or more dependent variables. In our study, we focus on the potential effect of the number of threads and path error density on the dependent variables: number of new states, error depth, and whether or not an error is found.

4.2.1 Characterization Variables

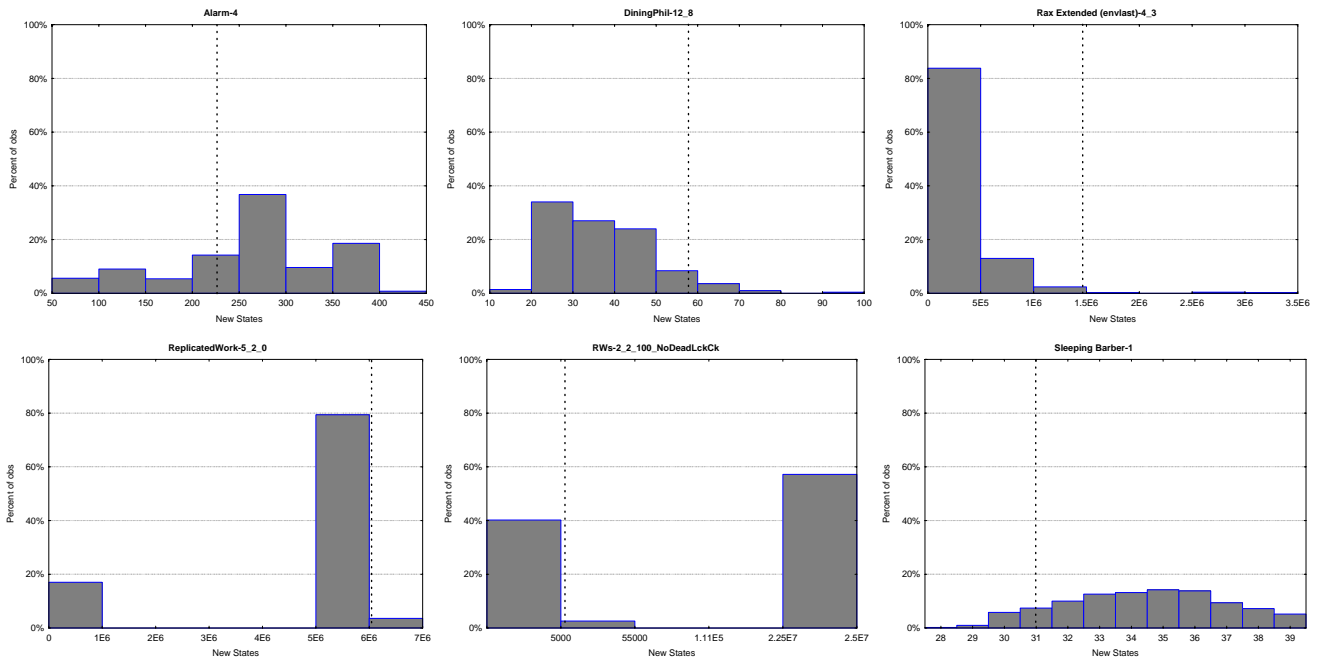


Figure 3: Scaled Default vs. Random Search Order

The characterization variables in our study are (1) the total number of threads created during the execution of the subject, and (2) the path error density. The first is simply a count of the main thread plus any child threads created during execution. The second measures the difficulty of finding a schedule that exhibits the error. We calculate path error density using JPF configured with the `RandomSearcher` and `RandomOrderScheduler`. This has the effect of simulating a single run of the subject program making a randomized sequence of scheduler decisions. We ran between 1000 and 10000 such runs for each subject and report the percentage of runs on which an error is found as the path error density. Subjects with path error densities below 10% after 1000 runs had an additional 9000 runs performed. Note that some of the kernel subjects contain infinite loops; for those we ran the simulation for a bounded number of steps where the bound was four times the shallowest error depth encountered for that subject.

4.2.2 Study Design and Setup

To conduct this study, we generated a population of subject programs that vary in the characterization variables. We began with the set of programs in Table 1 with their default parameter values as our *convenience sample*. We then assessed the number of examples with differing thread count and path error density values. Our goal was to produce three categories of values for each characterization variable with each category containing a sufficient number of subjects.

Many of the subject programs accept parameters to manipulate the thread count, but manipulating path error density was more problematic. We considered randomly generating parameterized versions of the examples and then calculating path error density as described above. This was ineffective in generating versions with diverse density measures. Subsequently, we spent time studying the subject programs in detail to understand the nature of the fault and its relationship to program parameters. This helped us choose parameter values to generate diverse density measures.

We explored several different categorizations of the measures. Based on “rules of thumb” in the model checking community we

chose thread count categories of $\{< 5, [5, 9], \geq 10\}$. The initial categories for density were quite broad, but we found that narrowing the definitions of low and high density to the extremal 10% regions gave us good explanatory power. We settled on density categories of $\{< 10, [10, 90], > 90\}$ (in %). The process of subject parameterization and density measurement was repeated until we ended up with approximately 5 subjects in each category; given the non-uniform nature of the density categories there were more subjects in the middle.

We used the same execution platform and environment configuration as was used in the study described in the previous section.

4.2.3 Results and Analysis

We organize the results into three parts. First, we focus on the effect of path error density, the most novel of the factors of our conjecture, on the dependent variables. Then, we analyze the interactions between the number of threads and the error density as measured by their impact on the dependent variables. Last, we briefly explore a regression model that can be helpful in predicting the dependent variables based on the path error density value.

Figure 4 shows the frequency distribution of new states and errors found with varying path error density over the more than 65 thousand randomized analysis samples; the size of the circle at a point indicates the number of subjects with that pair of values. The left plot shows a strong tendency towards small numbers of new states for high-density subjects - note the large circle in the lower right corner and the lack of any in the upper right portions of the plot. This suggests that analysis quickly finds errors for programs with high path error density. Even though a general pattern is hard to discern, it is clear that once density drops below 80% the cost measure can grow quite large. The right plot shows that for high-density programs it is generally the case that the error can be found - note the large circle in the upper right corner of the plot. Because the nature of a path-sensitive error detection tool like JPF is to continue the analysis until the first error is found, and we know that our population contains only programs with errors, failure to detect a fault can only be due to exhaustion of time or space. Thus, the pair of plots together can be interpreted as showing that low-density




Figure 4: Measure Variation with Density. (Circle’s size is proportional to observations’ frequency.)

programs yield expensive analyses and some of those analyses exhaust resource bounds thereby leading to failures in fault detection.

The variation in density by itself does not fully explain the variations observed in the dependent variables. Following our previous conjecture, we investigated whether observations that consider jointly the number of threads and error density can better explain the variation in the number of new states, error depth, and errors found. Figure 5 depicts two graphs for the pair of dependent variables, where the x-axis describes the three categories of thread count and the lines represents the three categories of error density. The left graph shows that the programs with the high density tend to cover a relatively small number of new states independently of the number of program’s threads. At lower densities, the picture is less clear. Closer inspection of the data revealed that analysis of a single program (RW) in the medium category has an enormous number of new states and raises the average dramatically. This is also confirmed by the large variation in the number of states covered by programs with a medium number of threads as evidence by the large variation observed around the means. The middle graph on error depth shows a more consistent tendency. Only programs with low error density seem to provide larger error depths. Interestingly, on average, programs with a larger number of threads did not end up with more error depth. The right graph presents a similar story. Programs with a higher error density tend to have easy to find faults, independent of the number of threads, while programs with lower error densities have harder to find faults. However, the number of threads does seem to affect whether the error is found or not, and it does compound with error density.

Based on the findings from Figure 5, we decided to further explore whether the number of threads and the error density could be good predictors of each of the dependent variables. Since the observations corresponding to the middle category of threads and error density appeared to change in response to variables other than the ones we are interested in, we decided to focus our model on the observations corresponding to the low and high groups. We then created three multiple regression models, one for each dependent variable, utilizing the standard support for regression analysis with backward stepwise refinement provided by most modern statistical packages (threats to the validity of the application of these models on this data set and our attempt to address those issues are discussed in the next section). The regression models, summarized in Table 3, show that a program’s error density has a significant inverse effect on the number of new states traversed and the error depth (low density programs tend to lead to a large number of new states and deeper errors), and a direct significant effect on whether or not the

Total observations: 8500			
Degrees of freedom: 2, 8497			
Factors	New States	Error Depth	Error Found
Thread Group		0.49	-0.47
Error Density Group	-0.27	-0.42	0.41

Table 3: Multiple regression models. Values in cells represent the model coefficient when the effect was significant ($p < 0.05$).

error is found (programs that have high error density reveal their errors more easily). The results for the number of threads confirms what we expected, the number of threads has a significant effect on error depth and the ease with which the error is found. However, the number of threads did not significantly impact the number of new states explored during analysis.

Part of the point of our study is that the research community does not currently have an adequate understanding of the size and structure of concurrent program state spaces and the effect that has on analysis cost. Consequently, caution must be exercised in generalizing from this model to broader populations of programs.

5. RESULTS IN CONTEXT AND RECOMMENDATIONS

5.1 Threats to Validity

Empirical studies are subject to threats to validity; these threats must be considered in order to determine the soundness and significance of the results. We detail the threats on our study and the steps we took to mitigate their impact on our findings.

Internal Validity. We placed an upper-bound on the execution time and memory that could be used by JPF during any analysis run. The bounds we chose were large, one hour and 4 GBytes respectively, and consistent with settings used in other studies. Changing those bounds may impact the findings on error detection, for instance, unlimited space and time would allow all errors to be found. During the course of our studies several defects were discovered and reported to the JPF team; fixes were provided and all studies were repeated on the updated version of JPF. We know of no defects remaining in JPF that would affect the results of our study.

External Validity. Our studies consider version 3.1.2 of JPF only; different versions of JPF and different path-sensitive analysis frameworks may yield different results. To address this, we are replicating the studies in this paper using the Bandera/Bogor model checking framework. Preliminary findings from our repli-




Figure 5: Interaction between error density and the number of threads.

cated studies confirm the results of the JPF study, however, we defer judgment until we have finalized those studies. The subjects chosen for this study were selected from a variety of sources. While the main goal of selecting subjects was to create a diverse set of multi-threaded Java programs containing safety property violations, we had two additional criteria: (1) selecting subjects that had been used in at least one previously published study in support of the studies in Section 3, and (2) selecting subjects that are either in widespread use, or are proposed as benchmarks, for evaluating path-sensitive analysis techniques in support of the studies in Section 4. Although we do not know if these programs are truly representative of multi-threaded Java programs in general, we believe our selection of programs from this initial population provides meaningful information with regard to our studies.

Construct Validity. The measures we considered in our studies are not the only possible measures of the variation on search order and the influence of program factors on analysis cost. System-independent measures such as the number of seen (visited) states, the number of end states, the number of transitions, etc. could also have been used; however, the number of new states and the error depth are widely used in evaluating state-oriented analysis techniques. System-dependent measures such as memory usage and CPU time were not considered to be valid measures because they could skew the results for a given subject even within an isolated environment due to factors such as execution time spent on garbage collection. Furthermore, execution time for path-sensitive analyses is strongly dependent on the number of new states; an early version of our studies that used execution time bore this out, but did not allow for retrospective comparisons.

Conclusion Validity. One of the major lessons we learned through our studies is that the values of the dependent variables are not only large, but also extremely variable. As a result, the 500 runs of JPF on each subject may have not been enough to appropriately characterize the error density of some programs. Our choice on the number of runs was incremental, increasing the number of runs until the observed standard variation seem to stabilized, and it was also limited by data collection costs. This is one of the reasons we adopted primarily an exploratory rather than a formal analysis, where we try to characterize relationships rather than claim any type of causality between the independent and dependent variables. Along similar lines, we apply similar caveats to the resulting regression models that, although checked through a residual analysis, did not fully meet all the traditional data distribution requirements.

5.2 Exposing and Controlling Search Order

Researchers wishing to evaluate the benefits of techniques that reduce the cost of path-sensitive error detection analyses should control for search order. We believe that it is not cost-effective for researchers to perform studies of the scale reported in this paper; our studies took multiple person-months and CPU-weeks to setup and conduct. We propose instead that developers of path-sensitive

analysis tools *provide the ability to configure the default search order*. Tool frameworks such as JPF and Bogor make this relatively easy, but for other tools, such as SPIN, this would require significant effort. With this ability, cross-tool studies would be able to ensure that tools use the same default order and intra-tool studies of new techniques will be able to evaluate the extent to which a technique’s benefit is independent of default order. Even without the ability to control default search order, tool developers should, at a minimum, clearly describe the default search order implemented in their tool so that researchers using the tool will be able to properly qualify research findings based on the use of the tool.

5.3 Building Better Benchmarks

Our studies clearly demonstrate that programs with high path error density are poor subjects for evaluating error detection techniques. Fortunately, the community is beginning to move away from using concurrency kernels, which almost uniformly have high density, in evaluations.

Researchers have proposed a number of criteria for constructing benchmarks of programs for evaluating multi-threaded program validation tools. For example, simple criteria such as program size or the presence of language constructs (e.g., wait and notify) are sensible ways to build diversity in a benchmark. More insightful selection criteria involve varying thread counts and the types of errors in the benchmark (e.g., the IBM benchmark). The results of our study suggest that additional factors can significantly effect the difficulty of finding an error in a program. Path error density is one such factor, but the variation in Figure 5 suggests that there are others.

The community needs good benchmarks and to build good benchmarks we need to understand the variations in programs to which different validation techniques are sensitive in terms of cost and effectiveness. Clearly, more work is needed to achieve this, but we plan to continue the work on benchmark development [16, 11] by sharing all of the subjects in this study through the Subject Infrastructure Repository [4].

5.4 Future Directions

We believe that the studies reported in this paper provide a wealth of data that can be leveraged for future work. For example, we plan to explore the influence of path error density on other techniques for validating and testing multi-threaded programs to understand their sensitivity to that factor. We also believe path error density is just one measure that can be used to characterize an analysis problem. It has value because it is efficient to calculate, i.e., a small number of randomized simulations can indicate a program’s path error density, and it appears to be useful in predicting when path-sensitive techniques will have an advantage over simpler techniques, such as randomized testing [25] for a given program. There may well be other factors that share these advantages.

In conducting our studies, we performed a very large number of

randomized depth-first search runs using JPF. Data from these runs suggest that a cost-effective strategy for finding errors in systems with low path error density is to run multiple parallel randomized analyses, and then terminate all analysis runs when one of them finds an error. Randomized analysis runs are completely independent which enables very large numbers of them to be executed in parallel. We are performing follow-up studies on parallel randomized path-sensitive error detection using JPF and Bandera/Bogor on large systems and plan to report on our findings in the near future.

Acknowledgments

This work was supported in part by the National Science Foundation through awards 0429149 and 0444167 and through CAREER award 0347518 and by the Army Research Office through DURIP award W911NF-04-1-0104. We would like to thank Radu Iosif, Willem Visser, Peter Melhitz, Robby, and Matt Hoosier for discussions about search order in Spin, JPF and Bogor. We thank Todd Wallentine, John Hatcliff, and Venkatesh Prasad Ranganath for general discussions about transition order in Bogor. Finally, we thank Shmuel Ur and Yaniv Eytani for granting access and supporting our use of their benchmarks.

6. REFERENCES

- [1] J. C. Corbett. Evaluating Deadlock Detection Methods for Concurrent Software. *IEEE Transactions on Software Engineering*, 22(3), Mar. 1996.
- [2] J. C. Corbett, M. B. Dwyer, J. Hatcliff, and Robby. Expressing checkable properties of dynamic systems: The Bandera Specification Language. *International Journal on Software Tools for Technology Transfer*, 2002.
- [3] D. L. Dill, A. J. Drexler, A. J. Hu, and C. H. Yang. Protocol verification as a hardware design aid. In *IEEE International Conference on Computer Design*, October 1992.
- [4] H. Do, S. G. Elbaum, and G. Rothermel. Subject infrastructure repository. <http://esquared.unl.edu/sir>.
- [5] Y. Dong, X. Du, G. J. Holzmann, and S. A. Smolka. Fighting livelock in the gnu i-protocol: a case study in explicit-state model checking. *Int'l. Journal on Software Tools for Tech. Transfer*, 4(4):505–528, 2003.
- [6] M. B. Dwyer, J. Hatcliff, M. Hoosier, V. Ranganath, Robby, and T. Wallentine. Evaluating the effectiveness of program slicing for model reduction of concurrent object-oriented programs. In *Proc. of the Twelfth Int'l. Conf. Tools and Algorithms for the Construction and Analysis of Systems*, 2006. LNCS 3920.
- [7] M. B. Dwyer, J. Hatcliff, V. R. Prasad, and Robby. Exploiting object escape and locking information in partial order reductions for concurrent object-oriented programs. *Formal Methods in System Designs*, 25(2–3):199–240, September–November 2004.
- [8] S. Edelkamp, S. Leue, and A. Lluch-Lafuente. Partial-order reduction and trail improvement in directed model checking. *International Journal on Software Tools for Technology Transfer*, 6(4), 2004.
- [9] <http://home.att.net/~ddavies/NewSmulator.html>.
- [10] Y. Eytani, K. Havelund, S. D. Stoller, and S. Ur. Toward a framework and benchmark for testing tools for multi-threaded programs. *Concurrency and Computation: Practice and Experience*, to appear.
- [11] Y. Eytani and S. Ur. Compiling a benchmark of documented multi-threaded bugs. In *Proc. of the Workshop on Parallel and Distributed Systems: Testing and Debugging*, 2004.
- [12] E. Farchi, Y. Nir, and S. Ur. Concurrent bug patterns and how to test them. In *Proc. of the 17th Int'l. Symp. on Parallel and Distributed Processing*, 2003.
- [13] S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *Proc. 9th International Conference on Computer Aided Verification*, June 1997.
- [14] A. Groce and W. Visser. Heuristics for model checking java programs. *Int'l. Journal on Software Tools for Tech. Transfer*, 6(4):260–276, 2004.
- [15] D. Hamlet and J. Voas. Faults on its sleeve: amplifying software reliability testing. In *Proc. of the 1993 ACM SIGSOFT international symposium on Software testing and analysis*, pages 89–98, 1993.
- [16] K. Havelund, S. D. Stoller, and S. Ur. Benchmark and framework for encouraging research on multi-threaded testing. In *Proc. of the Workshop on Parallel and Distributed Systems: Testing and Debugging*, 2003.
- [17] G. J. Holzmann. The model checker SPIN. *IEEE Transactions on Software Engineering*, 23(5):279–294, May 1997.
- [18] <http://www.kclee.de/clemens/java/javancss>.
- [19] M. Musuvathi and D. R. Engler. Model Checking Large Network Protocol Implementations. In *Proc. of the First Symp. on Networked Systems Design and Implementation*, Mar. 2004.
- [20] M. Musuvathi, D. Park, A. Chou, D. R. Engler, and D. L. Dill. CMC: A Pragmatic Approach to Model Checking Real Code. In *Proc. of the Fifth Symp. on Operating Systems Design and Implementation*, Dec. 2002.
- [21] C. Păsăreanu, M. B. Dwyer, and W. Visser. Finding feasible abstract counter-examples. *Int'l. Journal on Software Tools for Tech. Transfer*, 5(1):34–48, 2003.
- [22] <http://research.microsoft.com/qadeer/cav-issta.htm>.
- [23] Robby, M. B. Dwyer, and J. Hatcliff. Bogor: An extensible and highly-modular model checking framework. In *Proceedings of the 9th European Software Engineering Conference held jointly with the 11th ACM SIGSOFT Symposium on the Foundations of Software Engineering*, 2003.
- [24] Robby, M. B. Dwyer, J. Hatcliff, and R. Iosif. Space-reduction strategies for model checking dynamic systems. In *Proceedings of the 2003 Workshop on Software Model Checking*, July 2003.
- [25] S. D. Stoller. Testing concurrent java programs using randomized scheduling. In *Proc. Workshop on Runtime Verification*, 2002.
- [26] J. Tan, G. S. Avrunin, L. A. Clarke, S. Zilberstein, and S. Leue. Heuristic-guided counterexample search in flavors. In *Proc. of the 12th ACM SIGSOFT Twelfth Int'l. Symp. on Foundations of Software Engineering*, pages 201–210, 2004.
- [27] W. Visser, K. Havelund, G. Brat, and S. Park. Model Checking Programs. In *Proceedings of the 15th IEEE Conference on Automated Software Engineering*, Sept. 2000.
- [28] J. Yang, P. Twohey, D. R. Engler, and M. Musuvathi. Using Model Checking to Find Serious File System Errors. In *Proc. of the Seventh Symp. on Operating Systems Design and Implementation*, Dec. 2004.