

Search strategies for Java bottleneck location by dynamic instrumentation

Douglas J Brear, Thibaut Weise, Tim Wiffen, Kwok Cheung Yeung,
Sarah A M Bennett and Paul H J Kelly

Department of Computing, Imperial College London SW7 2AZ
s.bennett,p.kelly@imperial.ac.uk

Abstract. We have developed a prototype tool that supports instrumentation of distributed Java applications by on-the-fly deployment of interposition code at user-selectable program points. This paper explores the idea, originated in the Paradyn Performance Consultant, of systematically searching for performance bottlenecks by progressive refinement. We present the callgraph search algorithm in detail, and discuss a number of shortcomings with the approach, some of which can be addressed by improving the search strategy. We support our conclusions with two application examples. This is a report of work in progress, aimed at stimulating further investigation of this interesting approach.

1 Introduction

The idea of dynamically patching instrumentation code into a program while it is running immediately leads to the idea of deploying instrumentation in response to earlier measurements. For diagnosing performance problems, the objective in doing this is to direct the programmer towards opportunities for improving performance. The idea was pioneered in the Paradyn Performance Consultant [9], which uses a callgraph-based bottleneck search strategy [4].

This paper presents our progress in exploring this approach in the Java context. We address the following issues:

- **Implementing dynamic instrumentation within a Java Virtual Machine.** Several options are available; our implementation used our Veneer Virtual JVM, which offers the prospect of automatically optimising the bottleneck code once it has been identified.
- **The callgraph search strategy.** We give a detailed exposition of the callgraph search strategy, which clarifies some of the difficulties of the approach.
- **Search strategy enhancements.** To address some of these difficulties, “deep start” enhancements to the callgraph strategy have been proposed [12]. Random call-stack sampling has been shown to be effective but for Java it is difficult to implement; we study a call count alternative.

We demonstrate the usefulness of our prototype tools, JUDI (Java utility for dynamic instrumentation) and JBOLT (Java bottleneck locator toolkit) using two non-trivial applications.

2 Background

The Performance Consultant (PC), in Paradyn [9], structures the search for performance “bottlenecks” in terms of experiments; each experiment tests a

hypothesis. A hypothesis is an assertion that the application spends a substantial amount of time behaving in some pathological way which might be addressed by the performance engineer. Each experiment is targeted on a “focus”. For each experiment, the PC uses dynamic instrumentation to collect data to evaluate whether the experiment’s hypothesis holds true for this focus.

If an experiment results in support for the hypothesis at the specified focus, a further experiment can be formulated to identify the problem more precisely. The experiment can be refined in two ways - by refining the hypothesis (for example by distinguishing different kinds of synchronisation problem), or by refining the focus. A natural focus refinement is to test whether the problem is within a function, or within one of its callees. Other focus refinements might be to distinguish particular threads or transaction types.

An alternative to using dynamic instrumentation might be to instrument the binary to generate data on all possible hypotheses at all foci. The PC aims to get essentially the same result, with much lower overhead. The approach relies on sampling, and works on long-running applications provided behaviour is statistically fairly stable (we return to this issue shortly). The PC monitors the interference caused by the instrumentation it inserts, and throttles the number of ongoing experiments in order to keep the interference within specified limits. Refining the focus by tracing the application’s callgraph confines the search to code which is actually executed (explored in [4]).

For finding CPU bottlenecks, this approach is not clearly better than conventional approaches, such as sampling the program counter at random intervals. The real potential for the idea lies in bottlenecks which are harder to characterise — where instrumenting for all hypotheses would lead to excessive interference.

DynInst [3] is an example of a portable library for dynamic instrumentation, derived from the dynamic instrumentation technology [6] developed as part of Paradyn. In DynInst, a “point” is a location in a program where instrumentation can be inserted. A “snippet” is an abstract syntax tree representing some executable code which is to be inserted into a program at a point. Snippets can include conditionals, functions calls, and loops. Snippets are translated to binary instructions then copied into an array in the application’s address space. The application then has to be modified to branch into the snippet code, using code they term a “trampoline”.

Dynamic instrumentation for Java cannot be implemented this way, without exposing low-level implementation details of the JVM (for example, profile-directed re-optimisation). There are several alternative approaches:

- **Re-define the class** using the Java Debug Interface (JDI) call `VirtualMachine.redefineClasses()`, introduced in Sun’s JDK 1.4. This approach is used in ProbeMeister [10]. The overhead to do this is reported to be around 20 milliseconds for a small example, but increases with large classes since methods cannot be redefined individually, and JIT optimisation must be re-done. To reduce the overheads, Dmitriev [5] advocates refining the JDI with a call to re-define methods individually.
- **Run the JVM in debugging mode**, and set breakpoints to insert instrumentation. This is the approach taken by Popovici et al [11]. Historically,

JVMs have run substantially slower than normal in debugging mode, whether breakpoints are present or not; modern JVMs appear to make this approach more competitive.

- **Run the Java application in a virtual JVM.** This is the approach used in our JUDI tool [16]. We use the native JVM to execute application bytecode as much as possible, but have to intercept execution in order to retain control. The scheme suffers some overhead (see Section 5) on execution of all the application’s code (apart from system libraries), but runs with JIT optimisation.

2.1 The Veneer Virtual Java Virtual Machine (vJVM)

When extending the Java platform, it is often desirable to implement new features directly into the Java Virtual Machine. However, such modifications are tied to a specific JVM, and the complexity and size of many JVM implementations can make this a difficult and time-consuming task. One way to overcome the portability problem might be to write a JVM in Java, that runs on top of an underlying JVM. This has been done before in projects such as JavaInJava [14]. However, such implementations have a tendency to be extremely slow, since they attempt to emulate all aspects of a JVM.

Our approach, similar to the Dynamo/Rio projects [1,2], is to build a Java Virtual Machine using Java that uses the underlying JVM to directly execute as much of the program code as possible, only seizing control of the system when we wish the behaviour to deviate from that of the underlying JVM [17]. It allows us to run most of the application code directly (*i.e.* jumps to the corresponding bytecode), but the vJVM maintains control over execution by intercepting control flow. Veneer is much more powerful than is needed for dynamic instrumentation alone – our long-term goal is to use it to diagnose performance improvement opportunities automatically, then optimise dynamically.

The control flow is intercepted by “fragmenting” each method. There are a number of different fragmentation policies: by basic block, at method level (used by the JBolt extension to JUDI, see Section 3), and at RMI invocations (used for our work on RMI optimisation [17]). The method body is split into blocks, and the method entry is replaced by an “executor loop” that walks the control flow graph, invoking each block in turn. A method’s control flow graph can be updated “on-the-fly” (*i.e.* as the application is running), allowing us to use this as a framework for dynamic instrumentation.

The fragmentation process (which is based on the SOOT framework [15]) includes use/def and liveness analyses. Each fragment carries this as dependence metadata, which can be used in a run-time optimiser.

2.2 Java Utility for Dynamic Instrumentation (JUDI)

JUDI is a prototype dynamic instrumentation tool for Java [16]. It has a client graphical user interface (GUI) which connects to a set of remote vJVM’s running fragmented code. The GUI allows the user to browse the remote systems’ methods, and to upload “instruments” to the remote systems, where they are patched into the running code. The instruments are simple Java objects that can

be compiled and loaded on-the-fly. The tool consists of two components, which can run on separate hosts to avoid interference:

- JUDI-StartApp is responsible for starting the application and registering it with the RMI registry.
- JUDI-GUI allows the user to insert instruments dynamically into the application.

The instrumentation strategy used for the CPU Time bottleneck search (see Section 3) is designed to produce an inclusive timing of instrumented methods. Instruments are placed at the method entry point and at every return statement in the method. The instruments at the return statements are twinned with the single instrument at the entry, and when executed obtain the elapsed time on the `HPTimer`¹ of the entry point instrument.

JUDI’s unit of instrumentation deployment is an “Instrumentation Strategy Component” (ISC). This consists of:

- A set of **Instruments** - subclasses of a generic Instrument plan block. Instruments typically start, stop and log timers, or generate a log entry recording control flow, or data values.
- An **Instrumentation strategy**. This is usually just whether the instrument is to be executed “before”, “after”, or “around” its target block, and whether it applies to the whole method, or every basic block in the method.
- **Instrumentation targets**: the set of program objects (methods, classes) to which the instrumentation strategy should be applied. If not the entire program, this is selected explicitly through the GUI.
- **Instrumentation data class**: instruments generate data, usually either a log or some kind of histogram.
- **Instrumentation analyser**: this is a GUI component for viewing the results from the experiment.

3 The basic bottleneck search algorithm

Automatic bottleneck search is implemented as a JUDI ISC called JBolt, the **J**ava **b**ottleneck **l**ocator **t**oolkit. Fig. 4 shows, in outline form, the automatic bottleneck search algorithm (described informally in [4]). The algorithm starts with an application ready to run:

- It installs instrumentation at the root of the call graph (line 10), then allows execution to proceed.
- When an application thread executes a timer instrument, the application thread blocks and the active instrument object is passed to the search algorithm (line 19).
- The algorithm maintains a profile database which records instrumentation data accumulated so far. At line 24 the algorithm determines whether the new measurement allows us to classify this candidate program point as a bottleneck.

¹ The timer class used to obtain inclusive time measurements, `HPTimer`, was developed by Kwok Yeung [17]. It offers nanosecond resolution.

- When a bottleneck method is identified (line 25), we add it to the output set (line 30), then the instrumentation is refined to determine which, if any, of the method’s callees is responsible (line 33).
- We may instead conclude that we have enough data to decide this method is not a bottleneck (line 36). The instrument is therefore removed.
- Finally, we may decide to leave the instrument in place to accumulate further profile data.

The output consists of a list of bottleneck methods, prioritised by severity and specificity. This is conveniently presented to the user by colouring the nodes of the call graph, as shown in Fig. 2.

The algorithm in Fig. 4 is incomplete in two important ways:

- The algorithm fails to find some bottlenecks – where a bottleneck method is called from many different points, but none of its callers is itself a bottleneck (*e.g.* see Fig. 1). We return to this issue in Section 4.
- The algorithm assumes that we know each method’s callees.

3.1 Finding the callee set

The algorithm maintains two key data structures:

- **Focus:** This is the set of methods which have been identified as bottlenecks, and whose callees are being instrumented to determine whether a refined bottleneck hypothesis holds – *i.e.* whether the problem lies in one of the callees. Focus methods are not instrumented.
- **Frontier:** This is the set of methods currently being instrumented. A method is in the frontier if its caller is in the focus (we ignore recursion for simplicity of presentation).

The problem is that this definition is retrospective: by the time we find out that a method is a bottleneck, it is too late to instrument its callees. This is a fundamental problem, but there are several measures which help:

1. We could analyse each method to find call sites where the target method is statically known. This technique is used in Paradyn, but does not handle Java’s prevalent virtual methods well. Callee methods identified this way can, in principle, be added to the frontier as soon as the caller reaches the bottleneck threshold. A method can be included in the callee set but not actually called, at the cost of some redundant instrumentation.
2. We could instrument each bottleneck method, with code to record its callees and add them to the frontier. This is how Paradyn deals with virtual and indirect method calls. This does lead to some interference: the callee logging overheads are included in the caller method measurement.
3. We add a simple instrument to *all* methods, which checks whether the caller is in the focus. If so, the callee is added to the frontier. Our current prototype uses this technique, but we plan to investigate alternatives.
4. Some bottleneck methods are executed only once in a run of the application. To instrument their callees, we have to re-run the application. The need for this is alleviated to some extent by augmenting the callgraph search as described in Section 4.

5. We could instrument the application to construct its call graph, then re-run the application and use this call graph to guide bottleneck search. The first run would be slowed down due to heavy instrumentation, but the second run would suffer minimum interference. This scheme relies on the call graph being very similar on both runs. This is an attractive alternative strategy for our prototype.

3.2 Bottleneck identification criterion

The objective is to find and prioritise bottlenecks. The criterion for including an candidate bottleneck would ideally be the proportion of the application’s total run-time attributable to that candidate. However, to direct the bottleneck search, we need to classify candidates before the program has finished. For the experiments reported here we used a simple threshold of 10% of the execution time so far. We found that this strategy lead to some problems:

- Object constructors often evaluate as bottlenecks when the application is starting up. When a constructor is called from the ‘main’ method of the application, it is instrumented since the ‘main’ method is initially in the search focus. Later on, when seen in the context of the entire program run, it will probably represent only a small fraction of the total CPU time.
- Small methods that are called very frequently often do not initially appear to be bottlenecks. When instrumented and a time obtained, the method’s relatively short execution time does not make it appear to be a bottleneck. However later on in the program, once that method has been called many times, the combination of high frequency calls and short but un-negligible execution time may mean it is a bottleneck.

In either case, the root cause of the initial misinterpretation of the metric can be attributed to a lack of context; the method can really only be properly evaluated as a bottleneck in the context of the whole program. For this reason, JBolt periodically re-evaluates (currently at the end of each run through the application) all instrument data, in order to get a balanced view for each method.

4 Searching upwards through the call graph

Fig. 1 shows how the callgraph-based bottleneck search algorithm fails to find some bottlenecks. The callees of a method are included in the search only if the caller’s execution time indicates a bottleneck is present. If a method is called by several non-bottleneck methods, it could still account for a large proportion of the run-time.

The idea proposed by Roth and Miller [12] is to augment the search using additional information, and use this to target the search on “deep starters”. Roth and Miller choose deep starters, using call stack sampling, from information which their implementation already collects. In our vJVM implementation, it is possible to capture stack samples but rather expensive. Instead we use a simple instrument to count method executions (the same instrument builds the call graph in order to provide callee information). In our implementation, deep starters are methods whose execution frequency exceeds our chosen threshold (10%).




Fig. 1. Bottleneck hidden from the callgraph search

We use a deep-starter to target the callgraph search, by finding all the paths through the call graph that connect a focus node to the deep starter. All methods on these paths can be added to the frontier, and thus be instrumented. When executed, these instruments generate method timings (collected at line 19 in Fig. 4). If a method exceeds the bottleneck threshold, it is added to the frontier.

The timing instrument is an “around advice” [7]: the timer is started on entry, and logged on exit. However, the deep-starter scheme above adds methods from the call paths to the instrumentation frontier before those methods have returned. If we use only “around” instruments, we will not get any measurements until these methods are re-entered. Methods which are called just once will have to wait until the application is restarted. To improve this situation we experimented with “late instruments” – if a method is already on the call stack, we add an “after” instrument. This is used to measure the time between adding the instrument and method exit. The actual method execution time is sure to be more, but if this lower bound exceeds the threshold we can add the method to the focus immediately. We found that late instruments speed up the search substantially, at the expense of less reliable quantitative results (see Section 5.3).

5 Experimental evaluation

This section presents two examples of using JBolt to detect performance bottlenecks, and validates the results against Sun’s *hprof* profiler. RouteFinder is a railway route planning tool based on Dijkstra’s shortest path algorithm. It is single-threaded and consists 3823 lines of code (55 classes, 74 methods). SpecJVM98 *_208.db* (data management) is taken from the SpecJVM98 benchmark suite [13]. The program performs a variety of database operations on a memory-resident database of name, address and phone number records. It is also single-threaded, and consists of 8541 lines of code (24 classes, 40 methods).

We used the Sun Java 2 platform, standard edition version 1.4.02, running on SuSE Linux 7.2. Most of the experiments were carried out on a system with a single 1400MHz AMD Athlon processor, with separate 64KB L1 data & instruction caches, unified 256KB L2 cache, and 512MB memory. *Hprof* [8] samples at a constant rate, so the results are more accurate on a slower system. To compare measured method timings we used a slower machine, with a 450 MHz Pentium III processor, with separate L1 data & instruction caches of 16Kbytes, L2 unified cache 512Kbytes, and 256Mbytes memory.

Table 1 shows the impact of JBolt and *hprof* on the two applications’ execution time. For *_209.db*, the slowdown is fairly small, and JBolt does better than

Table 1. Benchmark Overheads

	RouteFinder		SpecJVM98 _209_db	
	Time(secs)	Slowdown factor	Time(secs)	Slowdown factor
Unfragmented Application	3.73	1.00	22.02	1.00
Fragmented Application	28.13	7.54	24.30	1.10
Profiled with JBolt	105.14	28.19	26.96	1.22
Profiled with <i>hprof</i>	71.74	19.23	51.30	2.33

hprof. However, for RouteFinder, the slowdown is very severe with both profilers, with JBolt somewhat worse. We believe the reason is that RouteFinder’s bottleneck method is executed many times (7 million), while _209_db spends most of its time in a method which is called a small number of times. In both profilers, method entry is the main source of overhead, but *hprof* incurs performance overheads on primitive Java classes, which Veneer runs at full speed.

5.1 Routefinder results

Fig. 2 shows the view displayed at the end of JBolt’s search for bottlenecks in RouteFinder. There is only one search strand in the application, branching at *Algorithm.findNextNode*, and joining again at *Edge.isMe*. Both *hprof* and JBolt agree on the bottleneck. As a result, the method *Edge.isMe* was modified. Running the new version of the program (with a larger rail network) gave a time of 12.07 seconds, as opposed to 25.13 seconds before optimisation (each averaged over five runs), *i.e.* a speedup of just over two.

Fig. 3 illustrates the efficiency of the hybrid search (top-down + deep-start strategies) in comparison with the callgraph (just top-down) search. The hybrid search locates all the bottlenecks in approximately 55% of the time taken by the callgraph search on its own.

5.2 SpecJVM98 _209_db results

Again, *hprof* and JBolt agree on the main bottleneck for this application; the method *DataBase.shell_sort*. JBolt overheads are low because *DataBase.shell_sort* is called a small number of times.

For the *DataBase.read_db* method, JBolt only attributes 4.43% of total CPU time on average to this method, compared to *hprof*’s 9.8%. Although JBolt and *hprof* agree on where the bottlenecks are in both applications, their attribution of time to methods varies substantially. In general, compared to *hprof*, JBolt underestimates time spent in short-running methods, and over-estimates time spent in longer-running methods.

5.3 Discussion

Our experiments show that JBolt is able to identify the same CPU bottlenecks as a conventional profiler. However, applications can run very slowly under JBolt. This is largely the overhead of fragmentation. We are implementing a number of improvements to Veneer’s basic mechanisms. Another strategy which could help would be to switch between unfragmented and fragmented method variants as instrumentation is added and removed.

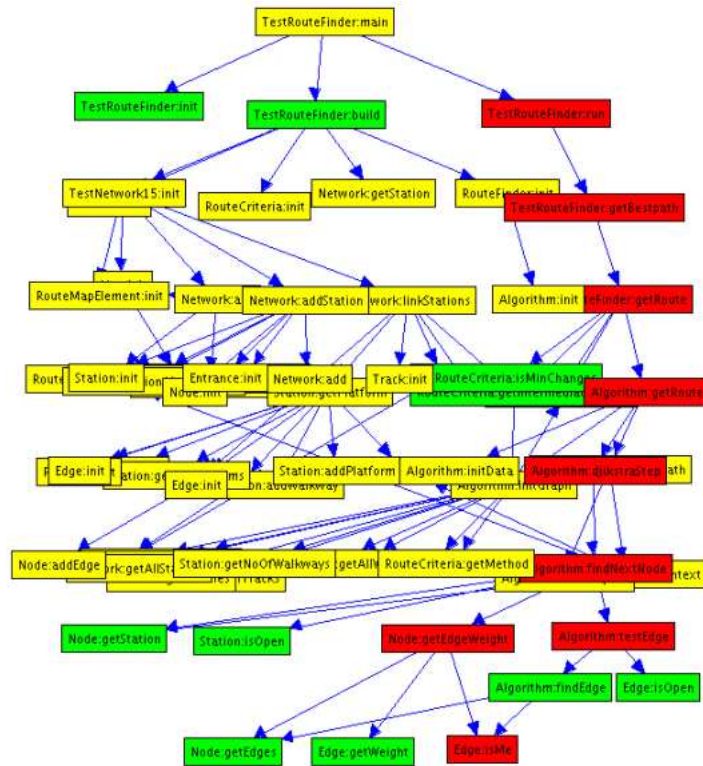


Fig. 2. RouteFinder - bottlenecks identified by JBoIt. The path down the right-hand side (displayed in red on the JUDI GUI) identifies the method Edge.isMe as the main bottleneck in this application.

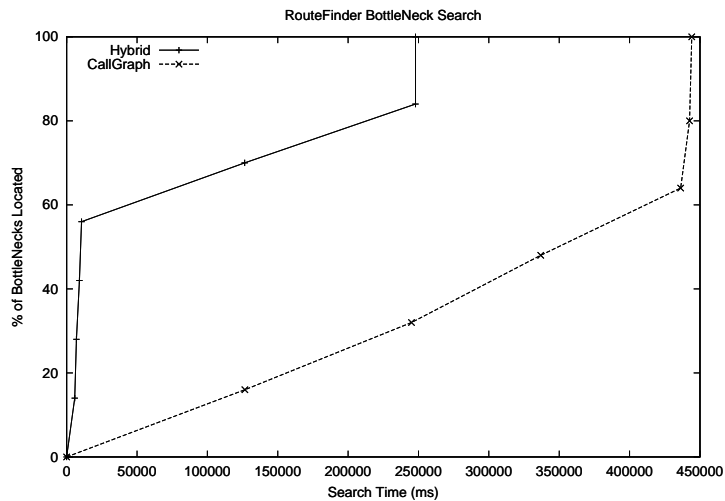


Fig. 3. RouteFinder - the hybrid search finds all bottlenecks in approximately 55% of the time taken by the callgraph-only search.

JBolt profiles the fragmented version of the application, not the real application. Comparing the method timings given by JBolt and *hprof* suggests that there are differences, so it is unclear how much of the variation in method timings and bottleneck ranking are due to differences in the performance characteristics of fragmented code, and how much is due to inaccuracies in the data sampling and analysis. Fragmentation appears to distort results for short lived methods.

Short-running applications like the benchmarks chosen here have to be re-executed several times (up to four), for JBolt to finish its search for bottlenecks. The main reason for this is to account for methods which are called only once per run.

The callgraph strategy has the desirable property that instrumentation overheads are never included in measurements. However, the upward search schemes introduce “late instruments” which are added while measurements higher in the callgraph are in place. We heuristically alleviated this problem by allowing only one late instrument to be in place at a time; we maintain a late instrument queue, and always select the deepest late instrument in the call graph. Relaxing this constraint leads to serious measurement errors, but enormously faster convergence: most applications saw all the bottlenecks identified within the first run. A more sophisticated instrument placement strategy should help here.

6 Conclusions and Further Work

We have reported on our exploration of search strategies for using dynamic instrumentation to locate and characterise performance bottlenecks. This work forms part of our longer-term objective to explore automatic profile-driven optimisation, and it is constructed on top of the Veneer framework which we built for this purpose.

Our results are not entirely positive. The main purpose of dynamic instrumentation is to avoid the performance impact of static approaches. We expect to reduce the performance impact of Veneer dramatically with further development. There are also serious concerns about the statistical significance of JBolt’s results. This seems inherent in the approach: sampling is driven by earlier measurements, so is not very random. For very long-running applications, or applications with a known repetitive structure, this can, perhaps, be overcome.

Perhaps the most promising prospect lies in searching for more subtle performance bottlenecks. Miller et al [9] observe that the search for a bottleneck can involve refinement in three dimensions; they call this the W^3 model:

- **When:** Is the performance problem confined to a particular phase of the computation? A particular time of day?
- **Where:** at what class, method, module, server, component or line of code does the problem occur? The callgraph (for example as shown in Fig. 2) shows a natural example. Others are possible: which threads? Which users? Which transaction types?
- **Why:** What is the reason for the performance problem?

In each dimension (when, where, why), the hierarchy provides a way to structure the search, leading to a successively more refined characterisation of the problem.

This should allow us to target expensive instrumentation on just the parts of the code and the phases of the computation where subtle performance problems are likely to occur.

Acknowledgements

This work was funded by the U.K. Engineering and Physical Sciences Research Council through a PhD studentship and the DESORMI project (GR/R15566).

References

1. Vasanth Bala, Evelyn Duesterwald, and Sanjeev Banerjia. Dynamo: a transparent dynamic optimization system. *ACM SIGPLAN Notices*, 35(5):1–12, 2000.
2. Derek Bruening, Evelyn Duesterwald, and Saman Amarasinghe. Design and implementation of a dynamic optimization framework for windows. In *4th ACM Workshop on Feedback-Directed and Dynamic Optimization (FDDO-4)*, Dec 2001.
3. Bryan Buck and Jeffrey K. Hollingsworth. An API for runtime code patching. *The International Journal of High Performance Computing Applications*, 14(4):317–329, Winter 2000.
4. H. Cain, B. Wylie, and B. P. Miller. A callgraph based search strategy for automated performance diagnosis. In Arndt Bode et al., editors, *Euro-Par 2000 - Munich*, volume 1900 of *Lecture Notes in Computer Science*, pages 108–122, 2000.
5. M. Dmitriev. Application of the HotSwap technology to advanced profiling. In *First International Workshop on Unanticipated Software Evolution (USE2002)*, Malaga, Spain, June 2002. <http://www.joint.org/use2002/sub/dmitriev-hotswapprof.pdf>.
6. Jeffrey K. Hollingsworth, Barton P. Miller, and Jon Cargille. Dynamic program instrumentation for scalable performance tools. In *Scalable High-Performance Computing Conference*, Knoxville, Tennessee, May 1994. ftp://ftp.cs.wisc.edu/paradyn/technical_papers/dyninst.ps.Z.
7. Gregor Kiczales, Erik Hilsdale, Jim Hugunin, et al. An overview of AspectJ. *Lecture Notes in Computer Science*, 2072:327–355, 2001.
8. Sheng Liang and Deepa Viswanathan. Comprehensive profiling support in the Java Virtual Machine. In *Proceedings of the 5th USENIX Conference on Object-Oriented Technologies and Systems (COOTS-99)*, pages 229–240, May 1999.
9. Barton P. Miller, Mark D. Callaghan, Jonathan M. Cargille, et al. The Paradyn parallel performance measurement tool. *IEEE Computer*, 28(11):37–46, 1995.
10. Paul Pazandak and David Wells. ProbeMeister: Distributed runtime software instrumentation. In *USE2002*, Malaga, Spain, June 2002. <http://www.joint.org/use2002/sub/pazandak-ProbeMeister.pdf>.
11. A. Popovici, T. Gross, and G. Alonso. Dynamic weaving for aspect oriented programming. In *AOSD 2002, April 22-26*, Enschede, The Netherlands, 2002. <http://ikplab11.inf.ethz.ch:9000/prose/webthings/aosd02.ps>.
12. Philip C. Roth and Barton P. Miller. Deep Start: A hybrid strategy for automated performance problem searches. In Burkhard Monien and Rainer Feldmann, editors, *Euro-Par 2002 - Paderborn, Germany*, volume 2400 of *Lecture Notes in Computer Science*, pages 86–96. Springer-Verlag, August 2002.
13. Standard Performance Evaluation Corporation (SPEC) JVM98 Suite, 1998. Available from <http://www.spec.org>.
14. Antero Taivalsaari. Implementing a Java Virtual Machine in the Java programming language. Technical Report TR-98-64, Sun Microsystems, 1998. <http://research.sun.com/techrep/1998/abstract-64.html>.

15. Raja Vallée-Rai, Etienne Gagnon, Laurie J. Hendren, et al. Optimizing Java bytecode using the Soot framework: Is it feasible? In David A. Watt, editor, *Compiler Construction (CC2000)*, Berlin, Germany, volume 1781 of *Lecture Notes in Computer Science*, pages 18–34. Springer-Verlag, 2000.
16. Kwok Yeung, Paul H J Kelly, and Sarah Bennett. Dynamic instrumentation for Java using a virtual JVM. In Getov et al., editors, *Performance Analysis and Grid Computing*. Kluwer, 2003.
17. Kwok Cheung Yeung and Paul H. J. Kelly. Optimising Java RMI programs by communication restructuring. In D. Schmidt and M. Endler, editors, *Middleware 2003 - Rio de Janeiro, Brazil*, Lecture Notes in Computer Science, June 2003.

```

1 // Initially no bottlenecks, no frontier, & focus is just application's Main method
2 Set<MethodId> bottleneckSet, focus; focus.add(mainId);
3 Set<Instrument> frontier;
4
5 // Frontier tells us which methods have been instrumented and maintains a list
6 // of each method's instrument objects. Initialise frontier to contain just
7 // the methods called by Main, and add instrumentation to this initial frontier
8
9 Set<MethodId> mainCallees = mainId.calleeSet();
10 frontier.add(TimerInstrumentFactory(mainCallees));
11
12 // Now run application. As it runs it will encounter instruments we put in place.
13 app.start();
14
15 while (!app.finished()) {
16     // wait for application to execute an instrument; resulting callback
17     // enqueues instrument object that has been activated
18
19     Instrument m = app.getNextActivatedInstrument();
20
21     // Consult profile database to determine whether we have enough
22     // information to conclude that this method is a bottleneck
23
24     switch (profileDatabase.isBottleneck(m)) {
25     case YES:
26         // This method turns out to be a bottleneck. Add this method to list of
27         // known bottlenecks, remove it from search frontier, remove its instrument-
28         // ation, & instead add its callees to the search frontier and instrument them
29
30         bottleneckSet.add(m.methodId, m.measurement);
31         frontier.remove(m);
32         focus.add(m.methodId);
33         frontier.add(TimerInstrumentFactory(m.methodId.calleeSet()));
34         break;
35
36     case NO:
37         // This method turns out not to be a bottleneck. Remove it from
38         // search frontier, remove its instrumentation.
39
40         frontier.remove(m);
41         break;
42
43     case MAYBE:
44         break; // leave instrumentation as it is for a while
45     }
46     // Update profile database for future reference
47     profileDatabase.update(m.methodId, m.measurement);
48
49     // Remove parent method from focus if none of its callees remain in the frontier
50     if (frontier.isIn(methodId.parent().calleeSet())) focus.remove(parent);
51 }

```

Fig. 4. Pseudocode outline of the callgraph-based bottleneck search algorithm.