Correcting the Dynamic Call Graph Using Control-Flow Constraints

Byeongcheol Lee, Kevin Resnick, Michael D. Bond, and Kathryn S. McKinley

The University of Texas at Austin

Abstract. To reason about whole-program behavior, dynamic optimizers and analysis tools collect a dynamic call graph using sampling. Previous approaches have not achieved high accuracy with low runtime overhead, and this problem is likely to become more challenging as object-oriented programmers increasingly compose complex programs.
This paper demonstrates how to use static and dynamic *control-flow graph* (CFG) constraints to improve the accuracy of the *dynamic call graph* (DCG). We introduce the *frequency dominator* (FDOM) which is a novel CFG relation that extends the dominator relation to expose relative execution frequencies of basic blocks. We combine conservation of flow and dynamic CFG basic block profiles to further improve the accuracy of the DCG. Together these approaches add minimal overhead (1%) and achieve 85% accuracy compared to a perfect call graph for SPEC JVM98 and DaCapo benchmarks. Compared to sampling alone, accuracy improves by 12 to 36%. These results demonstrate that static and dynamic control-flow information offer accurate information for efficiently improving the DCG.

1 Introduction

Well designed object-oriented programs use language features such as encapsulation, inheritance, and polymorphism to achieve reusability, reliability, and maintainability. As a result, these programs decompose functionality into many small methods, and virtual dispatch often obscures call targets at compile time. The dynamic call graph (DCG) records execution frequencies of call site-callee pairs, and is the key data structure that dynamic optimizers use to analyze and optimize whole-program behavior [2–5, 11, 23].

Prior approaches sample the DCG, trading accuracy for low overhead. Software sampling periodically examines the call stack to construct the DCG [4, 13, 18, 22, 29]. Hardware sampling lowers overhead by examining hardware performance counters instead of the call stack, but gives up portability. All DCG sampling approaches suffer from sampling error, and timer-based sampling suffers from timing bias. Arnold and Grove first measured and noted that the DCG is not very accurate [4]. They introduce *counter-based sampling* (CBS) to improve DCG accuracy by taking multiple samples and skipping some samples, adding overhead. We show this approach leaves room for improvement.

Figure 8(a) (page 15) shows DCG accuracy for the SPEC JVM98 benchmark *raytrace* using Jikes RVM default sampling. Each bar represents the *true* relative

frequency of a DCG edge (call site and callee) from a fully instrumented execution. Each dot is the frequency according to sampling. Edges are grouped by caller and are separated by dashed lines. Notice in particular that many methods make calls with the same frequency (i.e., the bars are the same magnitude within a method), but sampling tells a different story (i.e., the dots are not aligned). Sampling provides poor accuracy for many edges because of timer bias.

We present new *DCG correction* algorithms to improve DCG accuracy with extremely low overhead (1% on average). Our key insight is that a program's static and dynamic *control-flow graph* (CFG) constrains possible DCG frequency values. For example, two calls must execute the same number of times if their basic blocks execute the same number of times. To leverage this insight, we introduce the static *frequency dominator* (FDOM) relation, which extends the dominator and strong region relations on CFGs as follows: given statements x and y, x FDOM y if and only if x executes at least as many times as y.

We also exploit dynamic *basic block profiles* to improve DCG accuracy. Most dynamic optimizers collect accurate control-flow profiles such as basic block and edge profiles to make better optimization decisions [1, 3, 13, 17, 18]. We show how to combine these constraints to further improve the accuracy of the DCG. Our intraprocedural and interprocedural correction algorithms require a single pass over the basic block profile, which we perform periodically.

We evaluate DCG correction in Jikes RVM [3] on the SPEC JVM98 and DaCapo [8] benchmarks. We compare our approach to the default sampling configuration in Jikes RVM and the CBS sampling configuration recommended by Arnold and Grove [4]. Compared to a perfect call graph, default sampling attains 52% accuracy and our DCG correction algorithms boost accuracy to 71%; CBS by itself attains 76% accuracy and our DCG correction boosts its accuracy to 85%, improvements of 36% and 12% respectively. We show that each of FDOM, dynamic intraprocedural control-flow information, and interprocedural control-flow information improve accuracy while adding just 1% overhead.

Clients of the DCG include interprocedural analysis, such as alias and escape analysis, and optimizations such as the selection of which methods to recompile and inline. We evaluate the effect of more accurate DCGs on inlining, one of its clients. The adaptive hotspot compiler in Jikes RVM periodically recompiles and inlines hot methods. We add DCG correction immediately before the system recompiles. We measure the potential effect on inlining using a perfect call graph, which provides only a modest 2% average improvement in application time, significantly improving two programs by 13% and 12%. DCG correction matches these results on one of two: 18% and 2% respectively.

2 Background and Related Work

This section includes background material and compares dynamic call graph (DCG) correction to previous work. We first discuss how dynamic optimizers use sampling to collect a DCG with low overhead. We then compare the new frequency dominator relation to previous work. Finally, we compare DCG correction to previous static compiler analyses that construct a call graph using control-flow information.

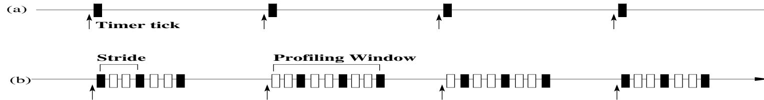

Fig. 1. Sampling. Filled boxes are taken samples and unfilled boxes are skipped samples. (a) One sample per timer tick. (b) CBS takes multiple samples per timer tick and strides between samples.

2.1 Collecting Dynamic Call Graphs

Dynamic optimizers could collect a *perfect* DCG by profiling every call, but the overhead is too high [4]. Some optimizers profile calls fully for some period of time and then turn off profiling to reduce overhead. For example, HotSpot adds call graph instrumentation only in unoptimized code [18]. Suganama et al. use *code patching* to insert call instrumentation, collect call samples for a period of time, and then remove the instrumentation [22]. These *one-time profiling* approaches keep overhead down but lose accuracy when behavior changes.

Many dynamic optimizers use software sampling to profile calls and identify hot methods [4, 6, 13]. Software-based approaches examine the call stack periodically and update the DCG with the call(s) on the top of the stack. For example, Jikes RVM and J9 use a periodic timer that sets a flag that triggers the system to examine the call stack at the next *yield point* and update the DCG [6, 13]. These systems insert yield points on method entry and exit, and on back edges.

Figure 1(a) illustrates timer-based sampling. Arnold and Grove show that this approach suffers from insufficient samples and *timing bias*: some yield points are more likely to be sampled than others, which skews DCG accuracy. They present *counter-based sampling* (CBS), which takes multiple samples per timer tick and *strides* to skip some samples in the profiling window, thus reducing timing bias. Figure 1(b) shows CBS configured to take three samples for each timer tick and to stride by three. By widening the profiling windows, CBS improves DCG accuracy, but increases profiling overhead. They report a few percent overhead to attain an average accuracy of 69%, but to attain 85% accuracy, they hit some pathological case with 1000% overhead. With our benchmarks, their recommended configuration attains 76% accuracy compared to a perfect call graph, whereas our approach improves the accuracy to 86% with an overhead of 1%.

Other dynamic optimizers periodically examine hardware performance counters such as those in Itanium to update the DCG. All sampling approaches suffer from sampling error, and timer-based sampling approaches suffer from timing bias as well. DCG correction can improve the accuracy of any DCG collected by sampling and we demonstrate two in Section 6.

Zhuang et al. [29] present a method for efficiently collecting the calling context tree (CCT), which represents the calling context of edges in call graph profile. Their work is orthogonal to ours since they add another dimension to the DCG (context sensitivity), while we improve DCG accuracy. We believe that our correction approach would improve CCT accuracy as well.

2.2 Constructing the DCG using Control-Flow Information

Static compilers have traditionally used control-flow information to construct a call graph [15, 28]. Hashemi et al. use static heuristics to construct an estimated call frequency profile [15]. Wu and Larus construct an estimated edge profile, which they use to construct an estimated call frequency profile [28]. These approaches rely solely on control-flow information to estimate call frequencies, whereas DCG correction starts with an inaccurate DCG and applies control-flow constraints to improve its accuracy. Hashemi et al. and Wu and Larus report high accuracy but the accuracy metric only considers the relative rank of call sites, whereas our overlap accuracy metric uses call edge frequencies. They construct profiles for C programs, while we target Java, which has richer DCGs and multiple call targets for a call site because of virtual dispatch [14]. These approaches use static *heuristics* to estimate frequencies, while DCG correction uses static *constraints* and combines them with dynamic profile information.

2.3 The Dominator Relation and Strong Regions

This paper introduces the *frequency dominator* (FDOM) relation, which extends *dominators* and *strong regions* [7]. Prosser first introduced dominators, which have a rich history [10, 25]. The set of dominators and post-dominators of x is the set of y that will execute at least once if x does. The set which frequency dominates x, on the other hand, is the subset which executes at least as many times as x. While strong regions find vertices x and y that must execute the same number of times, FDOM goes further and also finds vertices x and y where y must execute at least as many times as x.

3 Call Graph Correction Algorithms

This section describes DCG correction algorithms. We first present formal definitions for a control flow graph (CFG) and the dynamic call graph (DCG). We introduce the *frequency dominator* (FDOM) and show how to apply these static constraints to improve the accuracy of the DCG, and how to combine them with dynamic CFG frequencies to further improve the DCG.

3.1 Terminology

A *control-flow graph* represents *static* intraprocedural control flow in a method, and consists of basic blocks (V) and edges (E). Figure 2 shows an example control-flow graph CFG_p that consists of basic blocks *ENTRY*, a, b, c, d, e, and *EXIT*, as well as edges between them. A *basic block profile* gives the *dynamic* execution frequency of each basic block from some execution.

A *call edge* represents a method call, and consists of a *call site* and a *callee*. An example call edge in Figure 2 is e_5, the call from cs_c to CFG_t. The DCG of a program includes the *dynamic* frequency of each call edge, from some execution. For a call site cs, $OutEdges(cs)$ is the set of call edges that start at call site cs. $OutEdges(cs_a) = \{e_3, e_4\}$ in Figure 2. For a method m, $InEdges(m)$ is the set of call edges that end at m. $InEdges(CFG_t) = \{e_4, e_5\}$ in Figure 2.

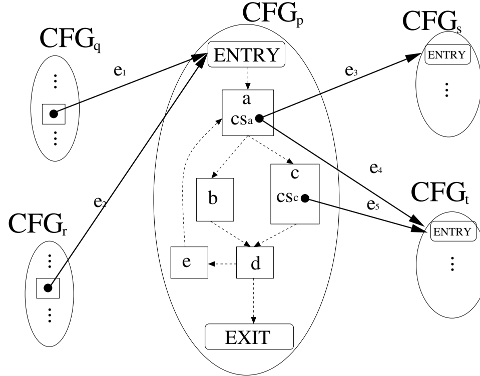

Fig. 2. Example dynamic call graph (DCG) and control flow graphs (CFGs).

DEFINITION 1 *The* INFLOW *of a method* m *is the total* flow *(execution frequency) coming into* m*:*

$$\text{INFLOW}(m) \equiv \sum_{e \in \text{InEdges}(m)} f(e)$$

where $f(e)$ *is the execution frequency of call edge* e*.* $\text{INFLOW}(m)$ *in an accurate DCG is the number of times* m *executes.*

DEFINITION 2 *The* OUTFLOW *of a call site* cs *is:*

$$\text{OUTFLOW}(\text{cs}) \equiv \sum_{e \in \text{OutEdges}(\text{cs})} f(e)$$

$\text{OUTFLOW}(cs)$ *in an accurate DCG is the number of times* cs *executes.*

Because a sampled DCG has timing bias and sampling errors, the DCG yields *inaccurate OUTFLOW* and *INFLOW* values. DCG correction corrects *OUTFLOW* using constraints provided by static and dynamic control-flow information (doing so indirectly corrects method *INFLOW* as well).

DCG correction maintains the relative frequencies between edges coming out of the same call site (which occur because of virtual dispatch), and does not correct their relative execution frequencies. For example, DCG correction maintains the ratio between $f(e_3)$ and $f(e_4)$ in Figure 2.

3.2 The Frequency Dominator (FDOM) Relation

This section introduces the *frequency dominator* (FDOM) relation, a static property of CFGs that represents constraints on program statements' relative execution frequencies. We show two constraints (theorems) it provides from the CFG on the DCG, and we prove these relations in the appendix.

DEFINITION 3 *Frequency Dominator (FDOM). Given statements* x *and* y *in the same method,* x *FDOM* y *if and only if for every possible path through the method,* x *must execute at least as many times as* y*. We also define FDOM(y)* $\equiv \{x \mid x \text{ FDOM } y\}$.

Like the dominator relation, FDOM is reflexive and transitive.

3.3 Static FDOM Constraints

We first propogate the FDOM constraint to DCG frequencies.

THEOREM 1 ***FDOM OUTFLOW Constraint****: Given method m and two call sites cs_1 and cs_2 in m, if cs_1 FDOM cs_2,*

$$\mathrm{OUTFLOW}(\mathrm{cs}_1) \geq \mathrm{OUTFLOW}(\mathrm{cs}_2)$$

Intuitively, the *OUTFLOW* constraint tells us that flow on two call edges is related if they are related by FDOM. For example, in Figure 2, cs_a *FDOM* cs_c and thus $OUTFLOW(cs_a) \geq OUTFLOW(cs_c)$.

THEOREM 2 ***FDOM INFLOW Constraint****: Given method m, if cs FDOM ENTRY (m's entry basic block),*

$$\mathrm{INFLOW}(m) \leq \mathrm{OUTFLOW}(\mathrm{cs})$$

Intuitively, the *INFLOW* constraint specifies that a call site must execute at least as many times as a method that always executes the call site.

3.4 Static FDOM Correction

Figure 3 applies the *FDOM OUTFLOW* constraint to a sampled DCG. The algorithm *FDOMOutflowCorrection* identifies missing edges from sampling based on the *FDOM OUTFLOW* constraint and adds these call edges by predicting their targets and frequencies. Then, the algorithm compares the sampled *OUTFLOW* of pairs of call sites that satisfy the FDOM relation. If their *OUTFLOW*s violate the *FDOM OUTFLOW* constraint, *FDOMOutflowConstraint* sets both *OUTFLOW*s to the maximum of their two *OUTFLOW*s. After processing a method, *FDOMOutflowConstraint* scales the *OUTFLOW*s of all the method's call sites to preserve the sum of the frequencies out of the method.

We also implemented correction algorithms using the *INFLOW* constraint, but they degrade DCG accuracy in some cases. This class of correction algorithms requires high accuracy in the initial *INFLOW* for a method to subsequently correct its *OUTFLOW*. In practice, we found that errors in *INFLOW* information propagated to the *OUTFLOW*s, degrading accuracy.

3.5 Dynamic Basic Block Profile Constraints

This section describes constraints on DCG frequencies provided by basic block profiles, and the following section shows how to correct the DCG with them. The *Dynamic OUTFLOW* constraint says that the ratio between the execution frequencies of two call sites specified by the basic block profile can be applied to the *OUTFLOW* of these two call sites.

THEOREM 3 ***Dynamic OUTFLOW Constraint*** *Given two call sites cs_1 and cs_2, and execution frequencies $f_{\mathrm{bprof}}(\mathrm{cs}_1)$ and $f_{\mathrm{bprof}}(\mathrm{cs}_2)$ provided by a basic block profile,*

$$\frac{\mathrm{OUTFLOW}(\mathrm{cs}_1)}{\mathrm{OUTFLOW}(\mathrm{cs}_2)} = \frac{f_{\mathrm{bprof}}(\mathrm{cs}_1)}{f_{\mathrm{bprof}}(\mathrm{cs}_2)}$$

procedure FDOMOutflowCorrection
input:
caller: a method whose OUTFLOWs are to be corrected
$FDOM(cs)$: a set of call sites that frequency-dominate cs
$f_{sample}(e)$: a function that returns the frequency of call edge e from sampling
$f_{sample}(cs)$: a function that returns the frequency sum of call edges in $OutEdge(cs)$ from sampling
output:
$f_{corrected}(e)$: a function that returns the corrected frequency for the call edge e

1: {**STEP1**: Insert call edges if missing.}
2: $CallSites \leftarrow$ getCallSitesFromCallerInDCG($caller$)
3: $InsertedCallSites \leftarrow \varnothing$
4: **for all** $cs_y \in CallSites$ **do**
5:   **for all** $cs_x \in FDOM(cs_y)$ **do**
6:     **if** $cs_x \notin CallSites$ **then**
7:       $InsertedCallSites \leftarrow InsertedCallSites \cup \{cs_x\}$
8:       {Predicts the virtual target using the class hierachy analysis.}
9:       $callee \leftarrow$ PredictTarget (cs_x)
10:       AddCallEdgeToDCGIfNotExists($caller$, cs_x, $callee$)
11:       $f_{sample}(cs_x) \leftarrow \max(f_{sample}(cs_x),\ f_{sample}(cs_y))$
12:     **end if**
13:   **end for**
14: **end for**
15: $CallSites \leftarrow CallSites \cup InsertedCallSites$

16: {**STEP2**: Initialize outflows.}
17: $sum_{old} \leftarrow 0$
18: $sum_{new} \leftarrow 0$
19: **for all** $cs \in CallSites$ **do**
20:   $Outflow(cs) \leftarrow f_{sample}(cs)$
21:   $sum_{new} \leftarrow sum_{new} + Outflow(cs)$
22:   **if** $cs \notin InsertedCallSites$ **then**
23:     {preserve per-method OUTFLOW.}
24:     $sum_{old} \leftarrow sum_{old} + Outflow(cs)$
25:   **end if**
26: **end for**

27: {**STEP3**: satisfy FDOM Outflow constraint.}
28: **for all** $cs_y \in CallSites$ **do**
29:   **for all** $cs_x \in FDOM(cs_y)$ **do**
30:     {constraint: $OUTFLOW(cs_x) \geq OUTFLOW(cs_y)$}
31:     $outflow_{old} \leftarrow Outflow(cs_x)$
32:     $Outflow(cs_x) \leftarrow \max(Outflow(cs_x), Outflow(cs_y))$
33:     $diff \leftarrow Outflow(cs_x) - outflow_{old}$
34:     $sum_{new} \leftarrow sum_{new} + diff$
35:   **end for**
36: **end for**

37: {**STEP4**: Use new outflow to derive the corrected frequency.}
38: $scale \leftarrow sum_{old}/sum_{new}$
39: **for all** $cs \in CallSites$ **do**
40:   **for all** $e \in OutEdges(cs)$ **do**
41:     $fraction \leftarrow f_{sample}(e)/f_{sample}(cs)$
42:     {Preserve the call target fraction and the frequency sum.}
43:     $f_{corrected}(e) \leftarrow Outflow(cs) \times scale \times fraction$
44:   **end for**
45: **end for**

Fig. 3. DCG Correction with FDOM OUTFLOW Constraints

The *Dynamic OUTFLOW* constraint can be applied to two call sites in different methods if basic block frequencies from different methods are accurate relative to each other (i.e., if the basic block profiles have *interprocedural accuracy*). In our implementation, basic block profiles do *not* have interprocedural accuracy. We experiment with using low-overhead method invocation counters to give basic block profiles interprocedural accuracy, and in this case we do apply *Dynamic OUTFLOW* to call sites in different methods (Section 4).

The *Dynamic INFLOW* constraint says that the call edge flow (frequency) coming into a method with a single basic block constrains the flow leaving any call site in the method.

THEOREM 4 ***Dynamic INFLOW Constraint****: Given a method m with a single basic block and a call site cs in m,*

$$\mathrm{INFLOW}(m) = \mathrm{OUTFLOW}(cs)$$

The *Dynamic INFLOW* constraint is useful for methods with a single basic block because the *Dynamic OUTFLOW* constraint cannot constrain the *OUTFLOW* of call sites in the single basic block (when basic block profiles do not have interprocedural accuracy). The *Dynamic INFLOW* constraint uses the total flow (frequency) coming into the method to constrain call sites' *OUTFLOW*.

3.6 Dynamic Basic Block Profile Correction

Figure 4 presents the algorithm for applying the *Dynamic OUTFLOW* constraint. *DynamicOutflowCorrection* sets the *OUTFLOW* of each call site cs to $f_{bprof}(cs)$, its frequency from the basic block profile. The algorithm then scales all the *OUTFLOW* values so that the method's total *OUTFLOW* is the same as before. This scaling helps to maintain the frequencies due to sampling across disparate parts of the DCG. Like *FDOMOutflowCorrection* in Figure 3, the *DynamicOutflowCorrection* can insert a call edge if its basic block profile frequency is higher than some threshold. For simplicity, we do not include this scheme in Figure 4.

Figure 5 presents the algorithm for applying the *Dynamic INFLOW* constraint to the DCG. For each method with a single basic block, *DynamicInflowCorrection* sets the *OUTFLOW* of each call site in the method to the *INFLOW* of the method. As in the case of the FDOM *INFLOW* constraint, we do not use the *Dynamic INFLOW* constraint together with an intraprocedural edge profile. However, with an interprocedural edge profile, *INFLOW* is accurate enough to improve overall DCG accuracy.

4 Implementing DCG Correction

Dynamic compilation systems perform profiling while they execute and optimize the application, and therefore DCG correction needs to be done at the same time with minimal overhead.

We minimize DCG correction overhead by limiting its frequency and scope. We limit correction's frequency by delaying it until the optimizing compiler requests DCG information. The correction overhead is thus proportional to the

procedure DynamicOutflowCorrection
input:
CallSites: a set of call sites whose OUTFLOWs are to be corrected
$f_{bprof}(cs)$: a function that returns the frequency of the call site cs from basic block profiles
$f_{sample}(e)$: a function that returns the frequency of call edge e from sampling
$f_{sample}(cs)$: a function that returns the frequency sum of call edges in $OutEdge(cs)$ from sampling
output:
$f_{corrected}(e)$: a function that returns the corrected frequency for the call edge e

1: {**STEP1**: Iterate call sites to find scale factor from basic block profile count to the sample count.}
2: $sum_{sample} \leftarrow 0$
3: $sum_{bprof} \leftarrow 0$
4: **for all** $cs \in CallSites$ **do**
5: $\quad sum_{sample} \leftarrow sum_{sample} + f_{sample}(cs)$
6: $\quad sum_{bprof} \leftarrow sum_{bprof} + f_{bprof}(cs)$
7: **end for**
8: $scale \leftarrow sum_{sample}/sum_{bprof}$
9: {**STEP2**: assign corrected call edge frequency.}
10: **for all** $cs \in CallSites$ **do**
11: $\quad$ **for all** $e \in OutEdges(cs)$ **do**
12: $\quad\quad fraction = f_{sample}(e)\ / f_{sample}(cs)$
13: $\quad\quad$ {constraint: $OUTFLOW(cs)\ / f_{bprof}(cs)$ is constant }
14: $\quad\quad$ {Preserve the call target fraction and the frequency sum}
15: $\quad\quad f_{corrected}(e) \leftarrow f_{bprof}(cs)\ \times scale \times fraction$
16: $\quad$ **end for**
17: **end for**

Fig. 4. DCG Correction with Dynamic OUTFLOW Constraints

Correction algorithm	**Correction unit**	**Algorithms**
Static FDOM CF Correction	Call sites within a method to be optimized	*FDOMOutflowCorrection*
Dynamic Intraprocedural CF Correction	Call sites within a method to be optimized	*DynamicOutflowCorrection*
Dynamic Interprocedural CF Correction	All call sites in the DCG	*DynamicOutflowCorrection* & *DynamicInflowCorrection*

Table 1. Call Graph Correction Implementations

number of times the compiler selects optimization candidates during an execution. Correction overhead is thus naturally minimized when the dynamic optimizer is selective about how often and which methods to recompile.

We limit the scope of DCG correction by localizing the range of correction. When the compiler optimizes a method m, it does not require the entire DCG,

```
procedure DynamicInflowCorrection
input:
  p: a single basic block procedure whose OUTFLOWs are
     to be corrected
  f_sample(e): a function that returns the frequency of call edge e
     from sampling
  f_sample(cs): a function that returns the frequency sum of call
     edges in OutEdge(cs) from sampling
output:
  f_corrected(e): a function that returns the corrected frequency for
     the call edge e
 1: CallSites ← getCallSitessInsideProcedure(p)
 2: {STEP1: Compute maxflow for the procedure p.}
 3: inflow ← 0
 4: for all e ∈ InEdges(p) do
 5:    inflow ← inflow + f_sample(e)
 6: end for
 7: maxoutflow ← max_{cs∈CallSites} f_sample(cs)
 8: maxflow ← max(inflow, maxoutflow)
 9: {STEP2: assign corrected frequency.}
10: for all cs ∈ CallSites do
11:    for all e ∈ OutEdges(cs) do
12:       fraction ← f_sample(e) / f_sample(cs)
13:       {constraint: INFLOW(p) = OUTFLOW(cs)}
14:       f_corrected(e) ← maxflow × fraction
15:    end for
16: end for
```

Fig. 5. DCG Correction with Dynamic INFLOW Constraints

but instead considers a localized portion of the DCG relative to m. Because we preserve the call edge frequency sum in the $OUTFLOW$ correction algorithm, we can correct m and all the methods it invokes without compromising the correctness of the other portions of the DCG. Because we preserve the DCG frequency sum, the normalized frequency of a call site in a method remains the same, independent of whether call edge frequencies in other methods are corrected or not.

For better interaction with method inlining, one of the DCG clients, we limit correction to *nontrivial* call edges in the DCG. Trivial call edges by definition are inlined regardless of their measured frequencies because they are so small that inlining them always reduces the code size. To exclude trivial call edges from correction, the inliner informs DCG correction of the trivial edges.

Table 1 summarizes the correction algorithms and their scope. They take as input the set of call sites to be corrected. Clearly, for FDOM correction, the basic unit of correction is the call sites within a procedure boundary. For dynamic basic block profile correction, there are two options. The first one limits the call site set to be within a procedural boundary, and the second one corrects all

the reachable methods. Since many dynamic compilation systems support only high precision intraprocedural basic profiles, the first configuration indicates how much DCG correction would benefit these systems.

Because our system does not collect interprocedural basic block profiles, we implement interprocedural correction by adding method counters. DCG correction multiplies the counter value by the normalized intraprocedural basic block frequency. We find this mechanism is a good approximation to interprocedural basic block profiles.

5 Methodology

This section describes our benchmarks, platform, implementation, and VM compiler configurations. We describe our methodologies for accuracy measurements against the perfect dynamic call graph (DCG), overhead measurements, and performance measurements.

We implement and evaluate DCG correction algorithms in Jikes RVM 2.4.5, a Java-in-Java VM, in its production configuration [3]. This configuration precompiles the VM methods (e.g., compiler and garbage collector) and any libraries it calls into a boot image. Jikes RVM contains two compilers: the *baseline compiler* and *optimizing compiler* with three optimization levels. (There is no interpreter in this system.) When a method is first executed, the baseline compiler generates assembly code (x86 in our experiments). A call-stack sampling mechanism identifies frequently executed (*hot*) methods. Based on these method sample counts, the *adaptive compilation system* then recompiles methods at progressively higher levels of optimization. Because it is sample based, the adaptive compiler is non-deterministic.

Jikes RVM runs by default using *adaptive* methodology, which dynamically identifies frequently executed methods and recompiles them at higher optimization levels. Because it uses timer-based sampling to detect hot methods, the adaptive compiler is non-deterministic. To measure performance, we use *replay compilation* methodology, which is deterministic. Replay compilation forces Jikes RVM to compile the same methods in the same order at the same point in execution on different executions and thus avoids high variability due to the compiler.

Replay compilation uses *advice files* produced by a previous well-performing adaptive run (best of twenty five). The advice files specify (1) the optimization level for compiling each method, (2) the dynamic call graph profile, and (3) the edge profile. Fixing these inputs, we execute two consecutive iterations of the application. During the first iteration, Jikes RVM optimizes code using the advice files. The second iteration executes only the application at a realistic mix of optimization levels.

We use the SPEC JVM98 [20] benchmarks, the DaCapo benchmarks (beta-2006-08) [8], and *ipsixql* [9]. We omit the DaCapo benchmarks *lusearch*, *pmd* and *xalan* because we could not get them to run correctly. We also include *pseudojbb* (labeled as *jbb*), a fixed-workload version of JBB2000 [21].

We perform our experiments on a 3.2 GHz Pentium 4 with hyper-threading enabled. It has a 64-byte L1 and L2 cache line size, an 8KB 4-way set associative

L1 data cache, a 12Kμops L1 instruction trace cache, a 512KB unified 8-way set associative L2 on-chip cache, and 2GB main memory, and runs Linux 2.6.0.

Accuracy Methodology. To measure the accuracy of our technique against the perfect DCG for each application, we first generate a perfect DCG by modifying Jikes RVM call graph sampling to sample every method call (instead of skipping). We also turn off the adaptive optimizing system to eliminate non-determinism due to sampling and since call graph accuracy is not influenced by code quality. We modify the system to optimize (at level 1) every method and to inline only trivial calls. Trivial inlining in Jikes RVM inlines a callee if its size is smaller than the calling sequence. The inliner therefore never needs the frequency information for these call sites. When program execution ends, the sampler has collected the perfect call graph. We restrict the call graph to the application methods by excluding all call edges with both the source and target in the boot image, and calls from the boot image to the application. We include calls edges into the boot image, since these represent calls to libraries that the compiler may want to inline into the application.

To measure and compare call graph accuracy, we compare the final perfect DCG to the final corrected DCG generated by our approach. Because DCG clients use incomplete graphs to make optimization decisions, we could have compared the accuracy of the instantaneous perfect and corrected DCGs as a function of time. However, we follow prior work in comparing the final graphs [4] rather than a time series, and believe these results are representative of the instantaneous DCGs.

Overhead Methodology. To measure the overhead of DCG correction without including its influence on optimization decisions, we configure the call graph correction algorithms to do correction but report old frequency information. We report the first iteration time because the call graph correction is triggered only during the compilation time. We report the execution time as the median of 25 trials to obtain a representative result not swayed by outliers.

Performance Methodology. We use the following configuration to measure the performance of using corrected DCGs to drive inlining. We correct the DCG as the VM optimizes the application, providing a realistic measure of DCG correction's ability to affect inlining decisions. We measure application-only performance by using the second iteration time. We report the median of 25 trials.

6 Results

This section evaluates the accuracy, overhead and performance effects of the DCG correction algorithms.

We use the notation *CBS(SAMPLES, STRIDE)* to refer to an Arnold-Grove sampling configuration [4]. To compare the effect of the sampling configuration on call graph correction, we use two sampling configurations: *CBS(1,1)* and *CBS(16,3)*. The default sampling configuration is *CBS(1,1)* in Jikes RVM. Arnold and Grove recommend *CBS(16,3)*, which takes more samples to increase accuracy, but keeps average overhead down to 1-2%.

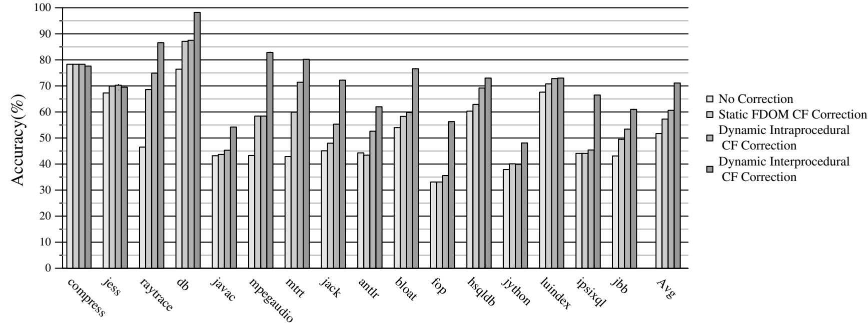

Fig. 6. Accuracy of DCG correction over the *CBS(1,1)* configuration.

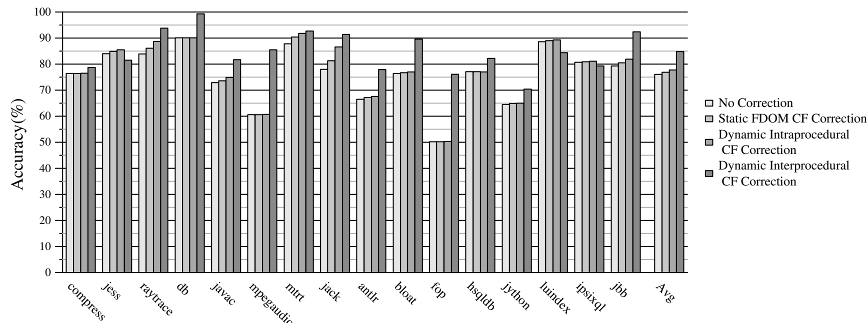

Fig. 7. Accuracy of DCG correction over the *CBS(16,3)* configuration

6.1 Accuracy

We use the overlap accuracy metric from prior work to compare the accuracy of DCGs [4].

$$overlap(DCG_1,\ DCG_2) = \sum_{e \in CallEdges} min(weight(e, DCG_1),\ weight(e,\ DCG_2))$$

where *CallEdges* is the intersection of the two call edge sets in DCG_1 and DCG_2 respectively, and $weight(e,\ DCG_i)$ is the normalized frequency for a call edge e in DCG_i. We use this function to compare the perfect DCG to other DCGs.

Figures 6 and 7 show how DCG correction boosts accuracy over *CBS(1,1)* and *CBS(16,3)* sampling configurations. The perfect DCG is 100% (not shown). The graphs compare the perfect DCG to the base system (*No Correction*), *Static FDOM CF Correction*, *Dynamic Intraprocedural CF Correction* and *Dynamic Interprocedural CF Correction*. Arnold and Grove report an average accuracy of 50% on their benchmarks for *CBS(1,1)*, and 69% for *CBS(16,3)* for 1 to 2% overhead [4]. We show better base results here with an average accuracy of 52% for *CBS(1,1)*, and 76% for *CBS(16,3)*.

These results show that our correction algorithms improve over both of the sampled configurations, and that each of the algorithm components contributes to the increase in accuracy (for example, *raytrace* in Figure 6 and *jack* in Figure 7), but their importance varies with the program. FDOM and intraprocedural correction are most effective when the base graph is less accurate as in

CBS(1,1) because they improve relative frequencies within a method. Interprocedural correction is relatively more effective using a more accurate base graph such as *CBS(16,3)*. This result is intuitive; a global scheme for improving accuracy works best when its constituent components are accurate.

Figure 8 shows how the correction algorithms change the shape of the DCG for *raytrace* for *CBS(1,1)*, our best result. The vertical bar presents normalized frequencies of the 150 most frequently executed call edges from the perfect DCG. The call edges on the x-axis are grouped by their callers, and the vertical dashed lines show the group boundaries. The dots show the frequency from the sampled or corrected DCG. In the base case, call edges have different frequencies due to timing bias and sampling error. *Static FDOM CF Correction* eliminates many of these errors and improves the shape of the DCG; Figure 8(b) shows that FDOM eliminates frequency variations in call edges in the same routine. Since FDOM takes the maximum of edge weights, it raises some frequencies above their true values. *Dynamic Intraprocedural CF Correction* further improves the DCG because it uses fractional frequency between two call sites, while FDOM gives only relative frequency. We can see in Figure 8(c) several frequencies are now closer to their perfect values. Finally, *Interprocedural CF Correction* further improves the accuracy by eliminating interprocedural sampling bias. The most frequently executed method calls, on the left of Figure 8(d), show particular improvement.

6.2 Overhead

Figure 9 presents the execution overhead of DCG correction, which occurs each time the optimizing compiler recompiles a method. Correction could occur on every sample, but this approach aggregates the work and eliminates repeatedly correcting the same edges. We take the median out of 10 trials (shown as dots). *Static FDOM Correction* and *Dynamic Intraprocedural CF Correction* add no detectable overhead. The overhead of the interprocedural correction is on average 1% and at most 3% (on jython). This overhead stems from method counter instrumentation (Section 4).

6.3 Performance

We evaluate the costs and benefits of using DCG correction to drive one client, inlining. We use the default inlining policy with *CBS(1,1)*. Figure 10 shows application-only performance (median of 10 trials) with several DCG correction configurations. The graphs are normalized to the execution time without correction. We first evaluate feeding a perfect DCG to the inliner at the beginning of execution (*Perfect DCG*). The perfect DCG improves performance by a modest 2.3% on average, showing that the Jikes RVM's inliner does not currently benefit significantly from high-accuracy DCGs. *Static FDOM CF Correction* shows the improvement from static FDOM correction, which is 1.1% on average. *Dynamic Intraprocedural CF Correction* improves performance by 1.7% on average. *Dynamic Interprocedural CF Correction* shows 1.3% average improvement. However, a perfect call graph does improve two programs significantly: *raytrace*

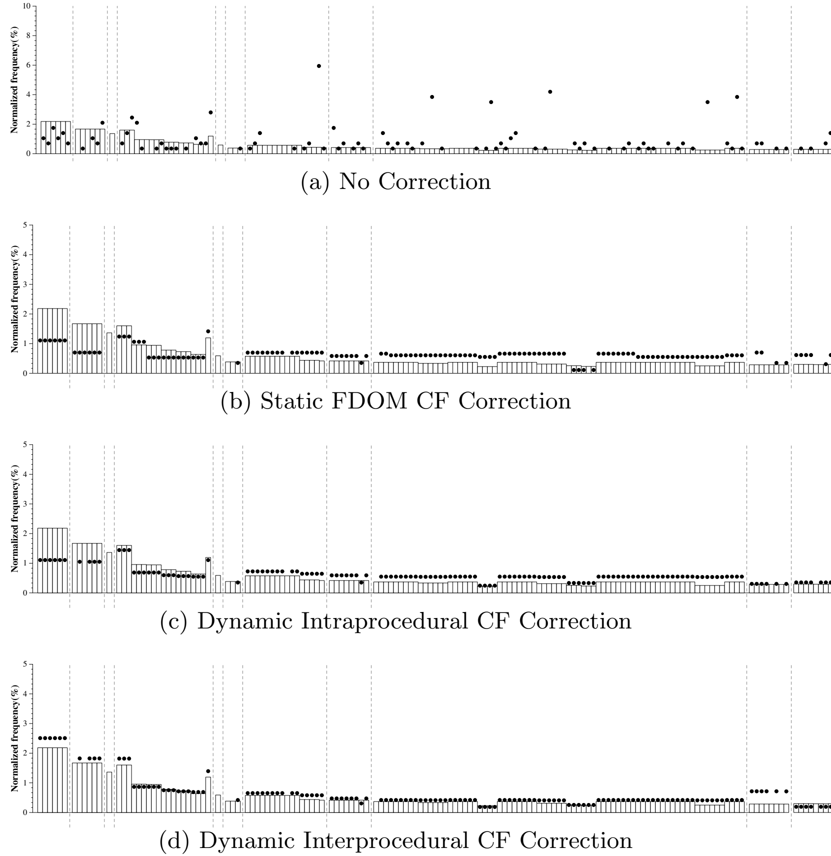

(a) No Correction

(b) Static FDOM CF Correction

(c) Dynamic Intraprocedural CF Correction

(d) Dynamic Interprocedural CF Correction

Fig. 8. Call graph frequencies for *raytrace* in *CBS(1,1)* configuration

ipsixql by 13% and 12% respectively, and DCG correction gains some of these improvements: 18% and 2% respectively.

7 Conclusion

This paper introduces *dynamic call graph (DCG) correction*, a novel approach for increasing DCG accuracy with existing static and dynamic control-flow information. We introduce the *frequency dominator* (FDOM) relation to constrain and correct DCG frequencies, and also use intraprocedural and interprocedural basic block profiles to correct the DCG. By adding just 1% overhead on average, we show that DCG correction increases average DCG accuracy over sampled graphs by 12% to 36% depending on the accuracy of the original. We believe DCG correction will be increasingly useful in the future as object-oriented programs become more complex and more modular.

Acknowledgments

We thank Xianglong Huang, Robin Garner, Steve Blackburn, David Grove, and Matthew Arnold for help with Jikes RVM and the benchmarks. We thank Calvin

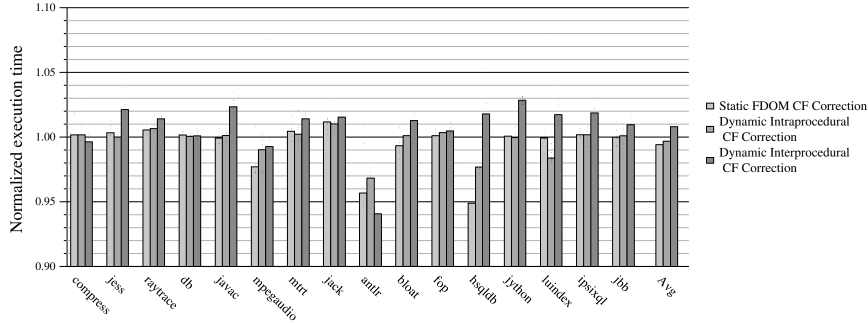

Fig. 9. The runtime overhead of call graph correction in *CBS(1,1)* configuration.

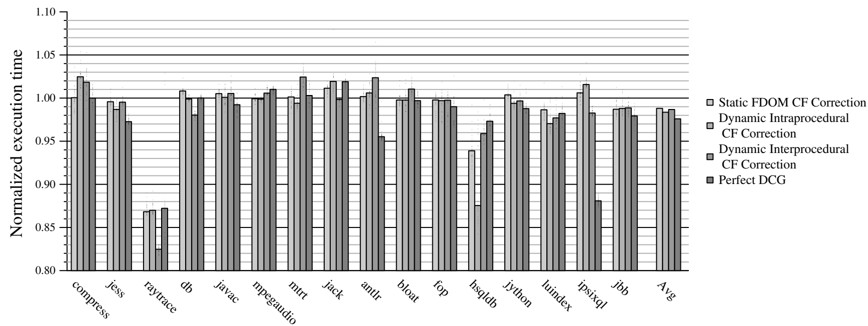

Fig. 10. The performance of correcting inlining decisions in *CBS(1,1)* configuration.

Lin, Curt Reese, Jennifer Sartor and Emmett Witchel for their helpful suggestions for improving the paper.

A Computing FDOM

The next two subsections present our algorithm for computing the *frequency dominator* (FDOM) relation. We first show a correlation between *simple cycles* and FDOM that applies to both *irreducible* and *reducible* graphs. We then a present our algorithm for computing FDOM, which applies to reducible graphs only. Both algorithms assume the existence of a back edge from *EXIT* to *ENTRY*. This edge simplifies the analysis but does not affect the FDOM relation, since following this edge is equivalent to re-calling the method.

General Control Flow Graphs

In this section, we show that the FDOM relation relates to *simple cycles* for general CFGs (both irreducible and reducible). This relation could be used to compute FDOM for irreducible CFGs. However, we only present FDOM computation for reducible CFGs. Lemma 1 shows that FDOM can be computed by considering *cycles* in the CFG.

LEMMA 1 *Given vertices* $x, y \in \text{CFG}$*,* x *FDOM* y*, if and only if* x *is contained in every cycle containing* y*.*

Proof. Show both forward and backward directions.

(⇒) *(by contradiction) Suppose there is a cycle that contains y but not x. Let this cycle be $c_y = \langle a, \dots, y, \dots, a \rangle$. From the definition of CFG, there is one path from* ENTRY *to a and another path from a to* EXIT*. By concatenating these three paths, we can build a path from* ENTRY *to* EXIT*, $c'_y = \langle \text{ENTRY}, \dots, a, \dots, y, \dots, a, \dots, \text{EXIT} \rangle$. If x is not in c'_y then y is already executed more times than x, a contradiction. So let x belong to some execution path that includes c'_y and executes n times, then we can repeat $c_y n+1$ times, resulting in y being executed more times than x, a contradiction.*

(⇐) *(by contrapositive) Suppose there is a path p from* ENTRY *to* EXIT *where the number of executions of y exceeds that of x. For the number of y to exceed x, p must contain at least one y. If the number of executions of y in the path is n, then path p should be of the form, $p = \langle \text{ENTRY}, \dots, y_1, \dots, y_2, \dots, y_n, \dots, \text{EXIT} \rangle$, where each occurrence of y is subscripted. The path p can be divided into $(n+1)$ subpaths: $\langle \text{ENTRY}, \dots, y_1 \rangle$, $\langle y_1, \dots, y_2 \rangle$, ..., $\langle y_n, \dots, \text{EXIT} \rangle$. Because there are fewer executions of x than y in p, x appears in all of the subpaths less than or equal to $(n-1)$ times. From the pigeonhole principle, there exist at least two x-free subpaths. If one of the two x-free subpaths is $\langle y_i, \dots, y_{i+1} \rangle$, this x-free subpath contradicts our assumption that there exists no cycle containing y and not x. If two subpaths are $\langle \text{ENTRY}, \dots, y_1 \rangle$ and $\langle y_n, \dots, \text{EXIT} \rangle$, another x-free cycle $\langle y_n, \dots, \text{EXIT}, \text{ENTRY}, \dots, y_1 \rangle$ can be constructed since we have a back edge from* EXIT *to* ENTRY*, again a contradiction.*

Lemma 1 states that FDOM can be computed by considering every cycle in the CFG. However, it is impractical to check if x is contained in every cycle that contains y because the number of cycles may be unbounded. The number of simple cycles in a method is bounded, and Lemma 2 shows that *simple cycles* are sufficient.

LEMMA 2 *Given vertices x, $y \in$ CFG, x is in every cycle containing y if and only if x is in every simple cycle containing y.*

Proof. Forward direction is trivial because every simple cycle is a cycle. To show the backward direction:

(⇐)(by contradiction) Suppose that x is in every simple cycle containing y, but there exists an x-free cycle c that contains y, $c = \langle a, \dots, y, \dots, a \rangle$ or $c = \langle y, \dots, y \rangle$. By assumption, c cannot be simple, so there must exist some element in c other than the beginning or end. There exists a cycle, therefore, in c that is from $\langle w, \dots, w \rangle$ that is simple, which we will call c'. Since by assumption, x is in every simple cycle containing y, if there exists a y in c', then there also exists an x, and since c' is simple, there can exist at most one of each. Therefore if c is a valid cycle including c', then another valid cycle is c with c' replaced with the single element w. Further, this new c still has the property that there exists an x-free cycle that contains y. Since c is of finite length, this process can be continued until there are no simple cycles left in c. Note that c still contains an x-free path that contains y, but now c is simple too, a contradiction. This proof holds for the entire program because there exists a back edge, by assumption, from EXIT *to* ENTRY*, and hence it is a cycle.*

Theorem 5 provides a method for computing FDOM by enumerating simple cycles [24, 26, 27], and is easily proved from the lemmas.

THEOREM 5 *Given vertices x, $y \in$ CFG, x FDOM y, if only if x is in every simple cycle containing y.*

Proof. From Lemma 1 and Lemma 2.

Reducible Control Flow Graphs

$$\begin{aligned} loophead(y) &= \text{loop entry of innermost loop that contains } y \\ backsrcs(y) &= backsrcs^*(loophead(y)) \\ natloop(y) &= natloop^*(loophead(y)) \\ exits(y) &= exits^*(loophead(y)) \end{aligned}$$

Ball defines "w pd v with respect to a set of vertices V" to capture post dominance within a natural loop [7]. For our algorithm, we only use two sets of vertices, back edges and loop exits for a given natural loop. Thus we use *PDBE* y to mean *backsrcs*(y) and *PDLE* y to mean *exits*(y). We combine these terms and make a new term, *PDL* $y \equiv$ *PDBE* y and *PDLE* y.

DEFINITION 4 *Post dominator with respect to loop (PDL). Given two vertices x, $y \in$ CFG, x PDL y if and only if every possible path from y to any innermost loop back source or exit ($backsrcs(y) \cup exit(y)$) must include x.*

Theorem 6 defines FDOM in terms of other CFG properties. Ball [7] and Johnson et al. [16] use sufficient and necessary conditions for FDOM in Lemma 1 to characterize *control regions* [12]. The algorithm in this section is motivated by Ball's paper [7].

THEOREM 6 *Given two vertices x, $y \in$ CFG, x FDOM y, if and only if $x \in \text{natloop}(y)$ and (x DOM y or x PDL y).*

Proof. (⇒) If x is not in natloop(y), then there exists an x-free cycle that contains y, a contradiction due to Lemma 1. So suppose that x DOM y or x PDL y is not true. If x DOM y is not true, then there exists a path from $\text{loophead}(y)$ to y that does not include x. If x PDL y is not true, then there exists a path from y to either a back edge or an exit of $\text{natloop}(y)$. Hence there exists either a path $\langle \text{loophead}(y), \ldots, y, \ldots, exits(y) \rangle$, which means there exists a path from ENTRY to EXIT that includes y and not x, or there exists a cycle $\langle \text{loophead}(y), \ldots, y, \ldots, b(\in \text{backsrcs}(y)), \text{loophead}(y) \rangle$, which is a cycle that includes y and not x, also a contradiction. Hence x DOM y or x PDL y must also be true.

(⇐) Show that every cycle that starts and ends at y, contains x. Let h_y be $\text{loophead}(y)$, and every path from y to y is one of this form, $c_y = \langle y, \ldots, z, \ldots, h_y, \ldots, y \rangle$, where z is in $backsrcs(y) \cup exits(y)$. Suppose that $x \in natloop(y)$, and x DOM y, then every path from h_y to y must contain x. If x PDL y, then path $\langle y, \ldots, z \rangle \subset c_y$ always contains x. Therefore c_y contains x if either of the two conditions is satisfied.

$FDOM(y)$ is the set of all x s.t. x $FDOM$ y. From Theorem 6, the following equations hold:

$$\begin{aligned} FDOM(y) &= natloop(y) \cap (DOM(y) \cup PDL(y)) \\ &= FDOM_D(y) \cup FDOM_P(y) \\ FDOM_D(y) &= natloop(y) \cap DOM(y) \\ FDOM_P(y) &= natloop(y) \cap PDL(y) \end{aligned}$$

We create an algorithm that can compute $FDOM_D(y)$ and $FDOM_P(y)$ in near-linear time.

Pingali and Bilardi give the paradox of linear computation time for super-linear-sized relations such as *control dependence* and post-dominance [19]. This paradox is resolved by the fact that these relations are transitive and can be factored into the transitive reduction form. The preprocessing time is used to construct some data structure that describes this reduction form, and a query is performed on this data structure. For instance, the post-dominance relation is transitive, and its transitive reduction form is a *post-dominator tree*, which can be computed in $O(E)$ time.[1] We describe an algorithm that computes an FDOM data structure in near-linear time, but retrieving FDOM for all nodes potentially takes super-linear time.

procedure setupFDOM (G:CFG)
1: Compute LoopStructureTree(G)
2: Compute DominatorTree(G)
3: Transform(G)
4: Compute PostdominatorTree(G)

procedure Transform(G:CFG)
1: **for all** loop header h in G **do**
2: Create $tempnode_h$
3: {Redirect each exit edge to go through the temp node.}
4: **for all** exit edges $b \rightarrow k$ for $natloop(h)$ s.t. b and k $\neq$ $tempnode_h$ **do**
5: remove $b \rightarrow k$
6: add $b \rightarrow tempnode_h$
7: add $tempnode_h \rightarrow k$
8: **end for**
9: {Redirect each back edge to go through the temp node.}
10: **for all** back edges $b \rightarrow h$ s.t. b $\neq$ $tempnode_h$ **do**
11: remove $b \rightarrow h$
12: add $b \rightarrow tempnode_h$
13: add $tempnode_h \rightarrow h$
14: **end for**
15: **end for**

Fig. 11. FDOM setup algorithm

[1] Cooper et al. describe the history of algorithms for dominance relation in detail [10].

procedure retrieveFDOM (y:Vertex):Set

```
 1: s ←{y}
 2: {Compute FDOM_D(y)}
 3: current ← dominatorParent(y)
 4: while current ≠ null and current ≠ loophead(y) do
 5:    s ← s ∪ {current}
 6:    current ← dominatorParent(current)
 7: end while
 8: {Compute FDOM_P(y)}
 9: current ← postdominatorParent(y)
10: while current ≠ null and current ≠ tempnode(y) do
11:    s ← s ∪ {current}
12:    current ← postdominatorParent(current)
13: end while
14: return s
```

Fig. 12. FDOM retrieval algorithm

The following theorem shows that we can compute *PDL*(y) inside a loop by applying the post-dominator algorithm on a CFG transformed by *Transform*(G) in Figure 11.

THEOREM 7 *Given two vertices $x, y \in$ CFG G, $x \in L = natloop(y)$ and the transformed CFG,* Transform(G) *as defined in Figure 11, then x* PDL *y in G, if and only if x* PDOM *y in* Transform(G)*.*

Proof. For both directions, let tempnode_L *be a newly added vertex for the natural loop, L, during the transformation from G to* Transform(G)*.*

(⇒) Assume x PDL *y in G is true. In G, there is no x-free path from y to any z in $backsrcs(y) \cup exit(y)$. In* Transform(G)*, there is no x-free path from y to $tempnode_L$. In* Transform(G)*, since $tempnode_L$ postdominates every vertex in L, there is no x-free path from y to* EXIT*. Therefore, x postdominates y in* Transform(G)*.*

(⇐) (by contrapositive) Suppose that x PDL *y in G is false. In G, there is an x-free path from y to some z in $backsrcs(y) \cup exits(y)$. In* Transform($G$)*, there is an x-free path from y to* tempnode_L*. Since x is inside the loop L and* tempnode_L *is the exit of the loop L in* Transform(G)*, there is an x-free path from y to* EXIT*. Therefore, x* PDOM *y does not hold in in* Transform(G)*.*

Figure 11 shows an algorithm that computes *FDOM*(y) in near-linear time using Theorems 6 and 7. The algorithm first computes the dominator tree and the loop structure tree for the CFG, then applies the transformation *Transform*(G) to the CFG and computes the post-dominator tree on the new graph. The transformation creates a single vertex per loop that the loop exits and back edges all go through. The computation is near-linear time since it is dominated by the time complexity of constructing the loop structure tree, or $O(V + E\alpha(E, V))$ [7]. If

the loop structure tree has already been computed, as in many compilers, then the algorithm adds $O(E)$ time.

Figure 12 shows an efficient algorithm for retrieving FDOM. To retrieve $FDOM(y)$ for a given vertex y, we simply iterate up the dominator tree from y until we hit y's loop header. We then iterate up the post-dominator tree on the transformed graph until we hit y's *tempnode*. We are guaranteed that from y to $loophead(y)$ in the dominator tree and from y to $tempnode(y)$ in the post-dominator tree, we will not leave $natloop(y)$. Further, once we go beyond these nodes, we will no longer be in $natloop(y)$. Hence we are only considering the set x s.t. $x \in natloop(y)$. Retrieving *FDOM* for all nodes potentially takes $O(V^2)$ time since there is potentially $O(V^2)$ information.

References

1. J. M. Anderson, L. M. Berc, J. Dean, S. Ghemawat, M. R. Henzinger, S.-T. A. Leung, R. L. Sites, M. T. Vandevoorde, C. A. Waldspurger, and W. E. Weihl. Continuous Profiling: Where Have All the Cycles Gone? In *Symposium on Operating Systems Principles*, pages 1–14, 1997.
2. M. Arnold, S. Fink, V. Sarkar, and P. F. Sweeney. A comparative study of static and profile-based heuristics for inlining. pages 52–64, Boston, MA, July 2000.
3. M. Arnold, S. J. Fink, D. Grove, M. Hind, and P. Sweeney. Adaptive optimization in the Jalapeño JVM. In *ACM Conference on Object Oriented Programming, Systems, Languages, and Applications*, pages 47–65, Minneapolis, MN, October 2000.
4. M. Arnold and D. Grove. Collecting and exploiting high-accuracy call graph profiles in virtual machines. In *Symposium on Code Generation and Optimization*, pages 51–62, Mar. 2005.
5. M. Arnold and B. G. Ryder. A framework for reducing the cost of instrumented code. In *PLDI '01: Proceedings of the ACM SIGPLAN 2001 conference on Programming language design and implementation*, pages 168–179, New York, NY, USA, 2001. ACM Press.
6. M. Arnold and P. F. Sweeney. Approximating the calling context tree via sampling. Technical Report RC 21789, IBM T.J. Watson Research Center, July 2000.
7. T. Ball. What's in a region?: or computing control dependence regions in near-linear time for reducible control flow. *ACM Letters on Programming Languages and Systems*, 2(1-4):1–16, 1993.
8. S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann. The DaCapo Benchmarks: Java benchmarking development and analysis. In *ACM Conference on Object-oriented programing, systems, languages, and applications*, Portland, OR, USA, Oct. 2006. http://www.dacapobench.org.
9. Colorado Bench. http://www-plan.cs.colorado.edu/henkel/projects/-colorado_bench.
10. K. D. Cooper, T. J. Harvey, and K. Kennedy. A simple, fast dominance algorithm. *Software Practice & Experience*, 4:1–10, 2001.
11. J. Dean, C. Chambers, and D. Grove. Selective specialization for object-oriented languages. In *PLDI '95: Proceedings of the ACM SIGPLAN 1995 conference on Programming language design and implementation*, pages 93–102, New York, NY, USA, 1995. ACM Press.

12. J. F., K. J. Ottenstein, and J. D. Warren. The program dependence graph and its use in optimization. *ACM Transactions on Programming Languages and Systems*, 9(3):319–349, 1987.
13. N. Grcevski, A. Kielstra, K. Stoodley, M. G. Stoodley, and V. Sundaresan. Java just-in-time compiler and virtual machine improvements for server and middleware applications. In *Virtual Machine Research and Technology Symposium*, pages 151–162, 2004.
14. D. Grove, G. DeFouw, J. Dean, and C. Chambers. Call graph construction in object-oriented languages. In *Conference on Object-Oriented Programming Languages and Systems*, pages 108–124, Oct. 1997.
15. A. Hashemi, D. Kaeli, and B. Calder. Procedure mapping using static call graph estimation. In *Workshop on Interaction between Compiler and Computer Architecture*, San Antonio, TX, 1997.
16. R. Johnson, D. Pearson, and K. Pingali. The program structure tree: computing control regions in linear time. In *ACM Conference on Programming Language Design and Implementation*, pages 171–185, 1994.
17. M. C. Merten, A. R. Trick, E. M. Nystrom, R. D. Barnes, and W. mei W. Hmu. A Hardware Mechanism for Dynamic Extraction and Relayout of Program Hot Spots. In *International Symposium on Computer Architecture*, pages 59–70, 2000.
18. M. Paleczny, C. Vick, and C. Click. The java hotspot server compiler. In *Usenix Java Virtual Machine Research and Technology Symposium (JVM'01)*, pages 1–12, April 2001.
19. K. Pingali and G. Bilardi. Optimal control dependence computation and the roman chariots problem. *ACM Transactions on Programming Languages and Systems*, 19(3):462–491, 1997.
20. Standard Performance Evaluation Corporation. *SPECjvm98 Documentation*, release 1.03 edition, March 1999.
21. Standard Performance Evaluation Corporation. *SPECjbb2000 (Java Business Benchmark) Documentation*, release 1.01 edition, 2001.
22. T. Suganuma, T. Yasue, M. Kawahito, H. Komatsu, and T. Nakatani. A Dynamic Optimization Framework for a Java Just-in-Time Compiler. In *ACM Conference on Object Oriented Programming, Systems, Languages, and Applications*, pages 180–195, 2001.
23. T. Suganuma, T. Yasue, and T. Nakatani. An empirical study of method inlining for a java just-in-time compiler. July 2002.
24. R. E. Tarjan. Enumeration of the elementary circuits of a directed graph. Technical report, Ithaca, NY, USA, 1972.
25. R. E. Tarjan. Finding dominators in directed graphs. *SIAM Journal of Computing*, 3(1):62–89, 1974.
26. J. C. Tiernan. An efficient search algorithm to find the elementary circuits of a graph. *Commun. ACM*, 13(12):722–726, 1970.
27. H. Weinblatt. A new search algorithm for finding the simple cycles of a finite directed graph. *J. ACM*, 19(1):43–56, 1972.
28. Y. Wu and J. R. Larus. Static branch frequency and program profile analysis. In *ACM/IEEE International Symposium on Microarchitecture*, pages 1–11, 1994.
29. X. Zhuang, M. J. Serrano, H. W. Cain, and J.-D. Choi. Accurate, efficient, and adaptive calling context profiling. In *ACM Conference on Programming Language Design and Implementation*, pages 263–271, 2006.