

Quantitative Interprocedural Analysis

Abstract

We consider the quantitative analysis problem for interprocedural control-flow graphs (*ICFGs*). The input consists of an *ICFG*, a positive weight function that assigns every transition a positive integer-valued number, and a labelling of the transitions (events) as good, bad, and neutral events. The weight function assigns to each transition a numerical value that represents a measure of how good or bad an event is. The quantitative analysis problem asks whether there is a run of the *ICFG* where the ratio of the sum of the numerical weights of good events versus the sum of weights of the bad events in the long-run is at least a given threshold (or equivalently, to compute the maximal ratio among all valid paths in the *ICFG*). The quantitative analysis problem for *ICFGs* can be solved in polynomial time, and we present an efficient and practical algorithm for the problem. We show that several problems relevant for static program analysis, such as estimating the worst-case execution time of a program or the average energy consumption of a mobile application, can be modeled in our framework. We have implemented our algorithm as a tool in the Java Soot framework. We demonstrate the effectiveness of our approach with two case studies. First, we show that our framework provides a sound approach (no false positives) for the analysis of inefficiently-used containers. Second, we show that our approach can also be used for static profiling of programs which reasons about methods that are frequently invoked. Our experimental results show that our tool scales to relatively large benchmarks, and discovers relevant and useful information that can be used to optimize performance of the programs.

Categories and Subject Descriptors D.3.4 [Programming Languages]: Processors—Optimization; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Program Analysis

General Terms Algorithms, Languages, Performance

Keywords Interprocedural analysis, Quantitative objectives, Mean-payoff and ratio objectives, Memory bloat, Static profiling.

1. Introduction

Static and interprocedural analysis. Static analysis techniques provide ways to obtain information about programs without actually running them on specific inputs. Static analysis explores the program behavior for *all* possible inputs and *all* possible executions. For non-trivial programs, it is impossible to explore all the possibilities, and hence static analysis uses approximations (abstract

interpretations) to account for all the possibilities [16]. Static analysis algorithms generally operate on the interprocedural control-flow graphs (for brevity *ICFGs*). An *ICFG* consists of a collection of control-flow graphs (CFGs), one for each procedure of the program. The CFG of each procedure has a *unique entry* node and a *unique exit* node, and other nodes represent statements of the program and conditions (in other words, basic blocks of the program). In addition, there are *call* and *return* nodes for each procedure which represent invoking of procedures and return from procedures. Call-transitions connect call nodes to entry nodes; and return-transitions connect exit nodes to return nodes. Algorithmic analysis of *ICFGs* provides the mathematical framework for static analysis of programs. Interprocedural analysis with objectives such as reachability, set-based information, etc have been deeply studied in the literature [10, 24, 29, 31–33].

Quantitative objectives. A *qualitative* (or Boolean) objective assigns to every run of a program a Boolean value (accept or reject). A *quantitative* objective assigns to every run of a program a real value that represents a quality measure of the run. The analysis of programs with quantitative objectives is gaining huge prominence due to embedded systems with requirements on resource consumption, promptness of responses, performance analysis etc. Quantitative objectives have been proposed in several applications such as for worst-case execution time (see [41] for survey), power consumption [37], prediction of cache behavior for timing analysis [20], performance measures [5, 13, 21], to name a few. Another important feature of quantitative objectives is that they are very well-suited for *anytime algorithms* [8] (where anytime algorithms generate imprecise answers quickly, and proceed to construct progressively better approximate solutions with refinements) (for a more elaborate discussion see [12]).

Mean-payoff and ratio objectives. One of the most well-studied and mathematically elegant quantitative objectives is the *mean-payoff* objective, where a rational-valued weight is associated with every transition and the goal is to ensure that the long-run average of the weights along a run is at least a given threshold [3, 19, 46]. For example, consider a weight function that assigns to every transition the resource (such as power) consumption, then the mean-payoff objective measures the average resource consumption along a run. Along with *ICFGs* with mean-payoff objectives, we also consider *ratio* objectives. For ratio objectives, the transitions (events) of the *ICFGs* are labelled as good, bad, or neutral events, and a positive weight function assigns a positive integer-valued weight to every transition, and the weight function represents how good or bad an event is. The quantitative analysis problem asks if there is a run of the program such that the ratio of the sum of the weights of the good events versus the weights of the bad events in the long-run is at least a given threshold. For example, consider a weight function that assigns weight 1 to each transition, and a labelling of events as follows: whenever a request is made is a bad event, whenever a request is pending is a good event, and whenever no request is pending is a neutral event. The ratio objective assigns the long-run average time between requests and the corresponding grant per request for a run, and measures timeliness of responses to requests. Finite-state systems (or intraprocedural finite-state programs) with mean-payoff objectives have been studied in the lit-

erature in [19, 21, 26, 46] for performance modeling, and more recently applied in synthesis of reactive systems with quality guarantee [5] and robustness [6], reliability requirements and resource bounds of reactive systems [9, 13, 17].

Interprocedural quantitative analysis. Quantitative objectives like mean-payoff and ratio objectives provide the appropriate framework to express several important system properties such as resource consumption and timeliness. While finite-state systems with mean-payoff objectives have been studied in the literature, the static analysis of *ICFGs* with mean-payoff and ratio objectives has largely been ignored. An interprocedural analysis is *precise* if it provides the meet-over-all-*valid-paths* solution (a path is valid if it respects the fact that when a procedure finishes it returns to the site of the most recent call). In the quantitative setting, the problem corresponds to finding the maximal value over all valid paths and to produce a witness (symbolic) path for that value. In this work, we consider precise interprocedural quantitative analysis for *ICFGs* with mean-payoff and ratio objectives.

Our contributions. In this work we present a flexible and general modelling framework for quantitative analysis and show how it can be used to reason about quantitative properties of programs and about potential optimizations in the program. We present an efficient polynomial-time algorithm for precise interprocedural quantitative analysis, which is implemented as a tool. We demonstrate the efficiency of the algorithm with two case studies, and show that our approach scales to programs with thousands of methods.

1. (*Theoretical modeling*). We show that *ICFGs* with mean-payoff and ratio objectives provide a robust framework that naturally captures a wide variety of static program analysis optimization and reasoning problems.
 - (a) (*Detecting container usage*). An exceedingly important problem for performance analysis is detection of runtime bloat that significantly degrades the performance and scalability of programs [27, 45]. A common source of bloat is inefficient use of containers [45]. We show that the problem of detecting usage of containers can be modeled as *ICFGs* with ratio objectives. A good use of a container corresponds to a good event and no use of the container is a bad event, and a misuse is represented as a low ratio of good vs bad events. Hence the container usage problem is naturally modeled as ratio analysis of *ICFGs*. While the problem of detecting container usage was already considered in [45], our different approach has the following benefits (see Section 5.2 for a comparison). First, our approach can handle recursion ([45] does not handle recursion). Second, our approach is sound, and does not yield false positives. Third, the approach of [45] ignored DELETE operations and we are able to take into consideration both ADD and DELETE operations (thus provide a more refined analysis). Moreover, our algorithmic approach for analysis of *ICFGs* is polynomial, whereas the algorithmic approach of [45] in the worst case can be exponential.
 - (b) (*Static profiling of programs*). We use our framework to model a conceptually new way for static profiling of programs for performance analysis. A line in the code (or a code segment) is referred as *hot* if there exists a run of the program where the line of code is frequently executed. For example, a function is referred as *hot* if there exists a run of the program where the function is frequently invoked, i.e., the frequency of calls to the function among all function calls is at least a given threshold. Again this problem is naturally modeled as ratio problem for *ICFGs*, and our approach statically detects methods that are more frequently invoked. Optimization of frequently executed code would naturally lead to performance improvements and reasoning about hot spots in the code can assist the compiler to apply

optimization such as function inlining and loop unrolling (see Subsection 3.2 for more details).

- (c) (*Other applications*). We show the generality of our framework by demonstrating that it provides an appropriate framework for theoretical modeling of diverse applications such as interprocedural worst-case execution time analysis, evaluating speedup in parallel computation, and interprocedural average energy consumption analysis.
 2. (*Algorithmic analysis*). The quantitative analysis of *ICFGs* with mean-payoff objectives can be achieved in polynomial time by a reduction to pushdown systems with mean-payoff objectives (which can be solved in polynomial time [14]). However, the resulting algorithm in the worst case has time complexity that is a polynomial of degree thirteen and space complexity that is a polynomial of degree six (which is prohibitive in practice). We exploit the special theoretical properties of *ICFGs* in order to improve the theoretical upper bound and get an algorithm that in the worst case runs in cubic time and with quadratic space complexity. In addition, we exploit the properties of *real-world* programs and introduce optimizations that give a practical algorithm that is much faster than the theoretical upper bound when the relevant parameters (the total number of entry, exit, call, and returns nodes) are small, which is typical in most applications. Finally we present a linear reduction of the quantitative analysis problem with ratio objectives to mean-payoff objectives.
 3. (*Tool and experimental results*). We have implemented our algorithm and developed a tool in the Java Soot framework [38]. We show through two case studies that our approach scales to relatively large programs from well-known benchmarks. The details of the case studies are as follows:
 - (a) (*Detecting container usage*). Our experimental results show that our tool scales to relatively large benchmarks (DaCapo 2009 [4]), and discover relevant and useful information that can be used to optimize performance of the programs. Our tool could analyze all containers in several benchmarks (like muffin) whereas [45] could analyze them partially (in muffin only half of the containers were analyzed in [45] — before the predefined time bound was exceeded). Our sound approach allows us to avoid false reports (that were reported by [45]) and our simple mathematical modelling even allows us to report misuses that were not reported by [45].
 - (b) (*Static profiling of programs*). We run an analysis to detect hot methods for various thresholds. Our experimental results on the benchmarks report only a small fraction of the functions as hot for high threshold values, and thus give useful information about potential functions to be optimized for performance gain. In addition we perform a dynamic profiling and mark the top 5% of the most frequently invoked functions as hot. Our experiments show a significant correlation between the results of the static and dynamic analysis. In addition, we show that the sensitivity and specificity of the static classification can be controlled by considering different thresholds, where lower thresholds increase the sensitivity and higher thresholds increase the specificity. We investigate the trade-off curve (ROC curve) and demonstrate the prediction power of our approach.
- Thus we show that several conceptually different problems related to program optimizations are naturally modeled in our framework, and demonstrate that we present a flexible and generic framework for quantitative analysis of programs. Moreover, our case studies show that our tool scales to benchmarks with real-world programs.

2. Definitions

In this section we present formal definitions of interprocedural control-flow graphs, and the quantitative analysis problems. We

will use an example program, described in Figure 1 and Figure 2, to demonstrate each definition.

Interprocedural control-flow graphs (ICFGs). A program P with m methods is modeled by an interprocedural control-flow graph (ICFG) \mathcal{A} which consists of a tuple $\langle A_1, \dots, A_m \rangle$ of m modules, where each module $A_i = \langle N_i, En_i, Ex_i, Calls_i, Retns_i, \delta_i \rangle$ represents a method (or the control-flow graph of a method) in the program. A module A_i contains the following components:

- A finite set of nodes N_i .
- An entry node En_i , which represents the first node of the method.
- An exit node Ex_i , which represents termination of the method.
- A finite set $Calls_i$ that denotes the set of calls of the method, and a finite set $Retns_i$ that denotes the set of returns.
- A transition relation δ_i , defined as follows: A transition in A_i is either (i) between two nodes in the same module (internal transitions) or between a return node and a node in the same module (i.e., u, v such that $u \in N_i \cup Retns_i$ and $v \in N_i \cup Calls_i$); or (ii) between a call node of a module A_i and the entry node of a module A_j (which models the invocation of method j from method i); or (iii) between the exit node of module A_i and a return node of a module A_j (which models the case that method i terminated and the run of the program continues in method j which invoked method i).

We denote by $N = \bigcup_{1 \leq i \leq m} N_i$ and similarly, $En = \bigcup_{1 \leq i \leq m} \{En_i\}$; $Ex = \bigcup_{1 \leq i \leq m} \{Ex_i\}$; $Calls = \bigcup_{1 \leq i \leq m} Calls_i$; $Retns = \bigcup_{1 \leq i \leq m} Retns_i$; and $\delta = \bigcup_{1 \leq i \leq m} \delta_i$. In the sequel we use N (resp. N_i) to denote all nodes in \mathcal{A} (resp. A_i), and refer to the nodes of $N \setminus (Calls \cup Retns \cup En \cup Ex)$ as *internal nodes*.

Quantitative ICFGs. A *quantitative ICFG (QICFG)* for brevity consists of an ICFG \mathcal{A} and a weight function w that assigns a rational-valued weight $w(e) \in \mathbb{Q}$ to every transition e , where \mathbb{Q} is the set of all rationals.

Example 1. Consider an example program shown in Figure 1 and Figure 2. In this example, (i) Modules: $A_1 = main$, $A_2 = foo$; (ii) Nodes: $N_2 = \{f:Entry, if(x>1), f:call foo, f:foo ret, x++, f:Exit\}$; (iii) Entry and exit: $En_2 = f:Entry$ and $Ex_2 = f:Exit$; (iv) Calls and returns: $Calls_2 = \{f:call foo\}$ and $Retns_2 = \{f:foo ret\}$; and (v) Transition: for example, $(x++, f:Exit) \in \delta_2$, $(f:Exit, f:foo ret) \in \delta_2$ and $(f:Exit, m:foo ret) \in \delta_2$.

Configurations and paths. A *configuration* consists of a sequence $c = (r_1, \dots, r_j, u)$, where each r_i is a return node (i.e., $r_i \in Retns$) and $u \in N$ is a node in one of the modules. Intuitively, when the module that u belongs to terminates, the program will continue in r_j . A sequence of configurations is *valid* if it does not violate the transition relation, and a *path* π is a valid sequence of configurations. We note that a path can be equivalently represented by the first configuration and a sequence of transitions. For a path $\pi = (c_1, c_2, \dots, c_\ell)$ we denote by (i) n_i : the node of configuration c_i (i.e., $c_i = (r_1, \dots, r_j, n_i)$), and (ii) α_i : the stack string of c_i (i.e., $\alpha_i = r_1, \dots, r_j$). For a path π let $\pi[1, i]$ denote the prefix of length i of π . A run of the program is modeled by a path in \mathcal{A} .

Example 2. Consider the program and the corresponding ICFG shown in Figure 1 and Figure 2. An example of a run in the program modeled as a path in the ICFG is as follows: $\langle (\epsilon, m:Entry), (\epsilon, while(x)), (\epsilon, if(y > 0)), (\epsilon, m:call foo), (m:foo ret, f:Entry), (m:foo ret, if(x > 1)), (m:foo ret, f:call foo), (m:foo ret, f:foo ret), (f:Entry, \dots) \rangle$, where ϵ denotes the empty stack.

Ratio analysis problem. In this work we consider the *ratio analysis* problem, where every transition of a QICFG has a label from the set $\{good, bad, neutral\}$. Intuitively, desirable events are la-

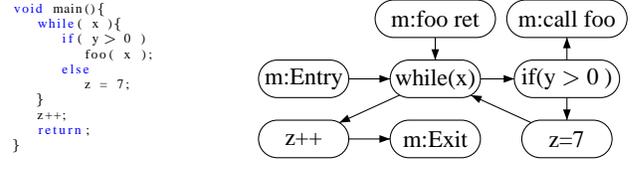


Figure 1. main

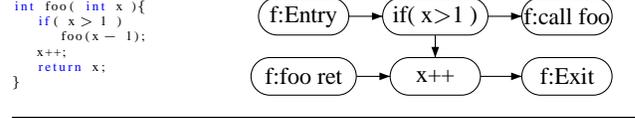


Figure 2. foo

belled as *good*, undesirable events are labelled as *bad*, and other events are labelled as *neutral*. The ratio analysis problem, given a QICFG, a labeling of the events, and a threshold $\lambda > 0$, asks to determine whether there is a run where the ratio of sum of weights of *good* events vs the sum of weights of the *bad* events that is greater than the threshold λ . Formally, we consider a positive integer-weight function w , that assigns a positive integer-valued weight to every transition, and for good and bad events the weight denotes how good or how bad the respective event is. For a finite path π we denote by $good(w(\pi))$ (resp., $bad(w(\pi))$) the sum of weights of the good (resp., bad) events in π . In particular, for the weight function w that assigns weight 1 to every transition, we have that $good(w(\pi))$ (resp. $bad(w(\pi))$) represents the number of good (resp. bad) events. We denote $Rat(w(\pi)) = \frac{good(w(\pi))}{\max\{1, bad(w(\pi))\}}$ the ratio of the sum of weights of good and bad events in π (note that in denominator we have $\max\{1, bad(w(\pi))\}$ to remove the pathological case of division by zero). For an infinite path π we denote

$$LimRat(w(\pi)) = \begin{cases} \liminf_{i \rightarrow \infty} Rat(w(\pi[1, i])) & \text{if } \pi \text{ has infinitely many} \\ & \text{good or bad events;} \\ 0 & \text{otherwise;} \end{cases}$$

informally this represents the ratio as the number of relevant (good/bad) events goes to infinity. Our analysis focusses on paths with unbounded number of relevant events, and infinitely many events provide an elegant abstraction for unboundedness. Hence, we investigate the following problem:

Given a QICFG with labeling of good, bad, and neutral events, a positive integer weight function w , and a threshold $\lambda \in \mathbb{Q}$ such that $\lambda > 0$, determine whether there exists an infinite path π such that $LimRat(w(\pi)) > \lambda$.

Remark 1. Our approach can be extended to reason about finite runs by a adding an auxiliary transition, labeled as a neutral event, from the final state of the program to its initial state (also see Section 3.3).

Mean-payoff analysis problem. In the mean-payoff analysis problem we consider a QICFG with a rational-valued weight function w . For a finite path π in a QICFG we denote by $w(\pi)$ the total weight of the path (i.e., the sum of the weights of the transitions in π), and by $Avg(w(\pi)) = \frac{w(\pi)}{|\pi|}$ the average of the weights, where $|\pi|$ denotes the length of π . For an infinite path π , we denote $LimAvg(w(\pi)) = \liminf_{i \rightarrow \infty} Avg(w(\pi[1, i]))$. The mean-payoff analysis problem asks whether there exists an infinite path π such that $LimAvg(w(\pi)) > 0$.

Assigning context-dependent and path-dependent weights. In our model the numerical weights are assigned to every transition of a QICFG. First, note that since we consider weight functions as an

input and allow all weight functions, the weights could be assigned in a dependent way. Second, in general, we can have an *ICFG*, and a finite-state deterministic automaton (such as a deterministic mean-payoff automaton [13]) that assigns weights. The deterministic automaton can assign weights depending on different contexts (or call strings) of invocations, or even independent of the context but dependent on the past few transitions (i.e., path-dependent), i.e., the automaton has the stack alphabet and transition of the *ICFG* as input alphabet and assigns weights depending on the current state of the automaton and an input letter. We call such a weight function *regular weight function*. Given a regular weight function specified by an automaton A and an *ICFG* we can obtain a *QICFG* (that represents the path-dependent weights) by taking their synchronous product, and hence we will focus on *QICFGs* for algorithmic analysis. The regular weight function can also be an *abstraction* of the real weight function, e.g., the regular weight function is an *over-approximation* if the weights that it assigns to the good (resp. bad) events are higher (resp. lower) than the real weights. If the original weight function is bounded, then an over-approximation with a regular weight function can be obtained (which can be refined to be more precise by allowing more states in the automaton of the regular weight function). Note that the *QICFG* we obtain from an *ICFG* and automaton A has a blowup in the number of states of A , and thus there is a tradeoff between the precision of A and the size of the *QICFG* constructed.

3. Applications: Theoretical Modeling

In this section we show that many diverse problems for static analysis can be reduced to ratio analysis of *QICFGs*. We will present experimental results (in Section 5) for the problems described in Section 3.1 and Section 3.2.

3.1 Container analysis

The inefficient use of containers is the cause of many performance issues in Java. An excellent exposition of the problem with several practical motivations is presented in [45]. The importance of accurate identification of misuse of containers that minimizes (and ideally eliminates) the number of false warnings was emphasized in [45] and much effort was spent to avoid false warnings for real-world programs. We show that the ratio analysis for *QICFGs* provides a mathematically sound approach for the identification of inefficient use of containers.

Two misuses. We aim to capture two common misuses of containers following the definitions in [45]. The first inefficient use is an *underutilized container* that always holds very few number of elements. The cause of inefficiency is two-fold: (i) a container is typically created with a default number of slots, and much more memory is allocated than needed; and (ii) the functionality that is associated with the container is typically not specialized to the case that it has only very few elements. The second inefficiency is caused by *overpopulated containers* that are looked up rarely, though potentially they can have many elements. This causes a memory waste and performance penalty for every lookup. Thus we consider the following two cases of misuse:

1. A container is *underutilized* if there exists a constant bound on the number of elements that it holds for all runs of the program.
 2. For a threshold λ , a container is *overpopulated* if for all runs of the program the ratio of GET vs ADD operations is less than λ .
- We note that our approach is demand-driven (where users can specify to check the misuse of a specific container).

Modeling. The modeling of programs as *ICFGs* is standard. We describe how the weight function and the ratio analysis problem can model the problem of detecting misuses. We abstract the different container operations into GET, ADD, and DELETE operations. For this purpose we require the user to annotate the relevant class methods by GET, ADD, or DELETE; and by a weight function that corresponds to the number of GET, ADD, or DELETE operations

```

12 void foo( int x ){
13     if ( x % 2 ){
14         Hashtable h1 = bar(x);
15         qux(h1, x);
16     }
17     else
18     {
19         for( int y = 0 ; y < x ; y++ ){
20             Hashtable h2 = new Hashtable(y);
21             h2.put(y,x);
22             for( int z = 0 ; z < y ; z++ )
23             {
24                 h2.put(z,y);
25                 ..
26                 h2.remove(z);
27             }
28         }
29     }
30 }

```

Figure 3. An example for underutilized container analysis.

that the method do (typically this number is 1). For example, in the class **HashSet**, the **add** method is annotated by ADD, the **contains** method is annotated by GET and the **remove** operation is annotated by DELETE. The **clear** operation which removes all elements from the set is annotated by DELETE but with a large weight (if **clear** appears in a loop, it dominates the add operations of the loop). We note that the annotation can be automated with the approach that is described in [45].

1. When detecting underutilized containers we define ADD operations as good events and DELETE operations as bad events, and check for threshold 1. Note that the relevant threshold is 1: if the (long-run) ratio of ADD vs DELETE is not greater than 1, then the total number of elements in the container is bounded by a constant.
2. When detecting overpopulated containers we define GET operations as good events and ADD operations as bad events, and check for the given threshold λ .

In addition, since we wish to analyze heap objects, the allocation of the container is a bad event with a large weight (i.e., similar effect as of **clear**); see Example 3. The container is misused iff the answer to the ratio analysis problem is NO (note that in the problem description for container analysis we have quantification over all paths and for ratio analysis of *QICFGs* the quantification is existential). The detection is demand-driven and done for an allocated container c .

Details of modeling. Intuitively, a transition in the call graph is good if it invokes a functionality that is annotated by a good operation (i.e., ADD operation for the underutilized analysis and GET operation for the overpopulated analysis) and the object that invokes the operation points to container c , and it is bad if the invoked operation is annotated as bad. Formally, for a given allocated container c : If at a certain line a variable t that *may point to* c invokes a good functionality, then we denote the transition as good. If t *must point to* c and invokes a bad functionality, then we denote it as a bad events. All other transitions are neutral. Note that our modeling is conservative. The misuse is detected for the container c if all runs of the program have a ratio of good vs bad events that is below the threshold (in other words, the container is not misused if there exists a path where the ratio of good vs bad events is above the threshold, and this exactly corresponds to the ratio analysis of *QICFGs*).

Example 3. We illustrate some important aspects of the container analysis problem with an example. Consider the program shown in Figure 3. We consider the containers that are allocated in line 9 and in line 20 and analyze them for underutilization. There exist runs that go through line 14 and properly use the container that is allocated in line 9, since *qux* method can add unbounded number of elements to the hash table (due to its recursive call). However, the container in line 20 is underutilized, since in every such run the number of elements is bounded by 2. However, note that if the DELETE operation is not handled, then the container is reported as properly used. We note that since we assign large weights to

the allocation of the container, this prevents the analysis from reporting that h2 properly uses the container that is allocated in line 20. In summary, the example illustrates the following important features: (1) the proper usage of the container should be tested also outside of its allocation site (as opposed to the approach of [45]); (2) sometimes the proper usage of a container is due to recursion; and (3) handling DELETE operations appropriately increases the precision of analysis. While these important features are illustrated with the toy example, such behaviors were also manifested in the programs of the benchmarks (see Subsection 5.2 for details).

Soundness. Our ratio analysis approach for *QICFGs* is both sound and complete (given the weighted abstracted *ICFG*). Since we use a conservative approach for assigning bad and good events, the *QICFG* we obtain for the misuse analysis of containers is sound and we get the following result.

Theorem 1. (SOUNDNESS). *The underutilized and overpopulated container analysis through the ratio analysis problem on QICFGs is sound (do not report false positives).*

Remark 2. *We remark about the significance of the soundness of our approach.*

- *The soundness criteria is a very important and desirable feature for container analysis (for details see [43, 45]), because a reported misused container needs to be analyzed manually and incurs a substantial effort for optimization. Hence as argued in [43, 45], spurious warnings (false positive) of misuse must be minimized. In our approach, a misuse is reported iff in every run a misuse is detected, and with a sound (over approximation) annotation of the weights our approach is sound.*
- *The soundness of our approach is with respect to a sound (over approximation) annotation of the ADD, DELETE and GET operations. In addition, for a given QICFG, our ratio analysis is precise (i.e., both sound and complete), hence, our container analysis is sound.*

3.2 Static profiling

Finding the most frequently executed lines in the code can help the programmer to identify the critical parts of the program and focus on the optimization of these parts. It can also assist the compiler (e.g., a C compiler) to decide whether it should apply certain optimizations such as *function inlining* (replacing a function call by the body of the called function) and *loop unrolling* (re-write the loop as a repeated sequence of similar independent statements). These optimizations can reduce the running time of the program, but on the other hand, they increase the size of the (binary) code. Hence, knowing whether the function or loop is *hot* (frequently invoked) is important when considering the time vs. code size tradeoff. In this subsection we present the model for profiling the frequency of function calls (which allows finding hot functions), and we note that our profiling technique is generic and can be scaled to detect other hot spots in the code (e.g., hot loops).

Problem description. Given a program with several functions, a function f is called λ -hot, if there exists a run (of unbounded length) of the program where the frequency of calls to f (among all function calls in the run) is at least λ . Formally, for a run, given a prefix of length i , let $\#f(i)$ denote the number of calls to f and $\#c(i)$ denote the number of function calls in the prefix of length i . The function is λ -hot if there exists a run such that $\liminf_{i \rightarrow \infty} \frac{\#f(i)}{\#c(i)} > \lambda$.

Modeling. The modeling of programs as *ICFGs* is straightforward. We describe the labeling of events and weight function in *ICFGs* to determine if a function f is λ -hot. First we label call-transitions to f as good events and assign weight 1; then we label all other call-transitions as bad events and assign them weight 1. To ensure that the number of calls to f also appear in the denominator (in the total number of calls) we label transitions from the entry

node of f as bad events with weight 1. The function f is λ -hot iff the answer to the ratio analysis problem with threshold λ is YES.

3.3 Estimating worst-case execution time

The approach of [12] for estimating worst-case execution time (WCET) is also naturally captured by ratio analysis. While the intraprocedural problem was considered in [12], our approach allows the more general interprocedural analysis. In this approach, we consider (as in [12]) that each program statement is assigned a cost that corresponds to its running time (e.g., number of CPU cycles).

Modeling. The modelling of WCET analysis of the program is as follows: We add to the *ICFG* of the program a transition from every terminal node to the initial node, and every such transition is a bad event with weight 1. All the other transitions are good events and their weight is their cost (running time). The WCET of the program is at most N cycles if and only if the answer to the ratio analysis problem with threshold $\frac{1}{N}$ is NO.

3.4 Evaluating the speedup in a parallel computation

The speed of a parallel computing is limited by the time needed for the sequential fraction of the program. For example, if a program runs for 10 minutes on a single processor core, and a certain part of the program that takes 2 minutes to execute cannot be parallelized, then the minimum execution time cannot be less than two minutes (regardless of how many processors are devoted to a parallelized execution of this program). Hence, the speedup is at most 5. Amdahl's law [1] states that the theoretical speedup that can be obtained by executing a given algorithm on a system capable of executing n threads of execution is at most $\frac{1}{B + \frac{1}{n}(1-B)}$, where B is the fraction of the algorithm that is strictly serial. Our ratio analysis technique can be used to (conservatively) estimate the value of B and thus to evaluate the outcome of Amdahl's law.

Modeling. As in Section 3.3, we consider that the cost of every program statement is given, and we add to the *ICFG* of the program a transition from every terminal node to the initial node, this time as a neutral event with weight 0. All the transitions of the code that cannot be parallelized are defined as good events, and the other transitions are defined as bad events. We denote by P the fraction of the code that can be parallelized and by S the fraction of the code that is strictly serial. The value of $\frac{S}{P}$ is at most λ if and only if the answer to the ratio analysis problem with threshold λ is NO. Hence B is bounded by $\frac{1}{1+\lambda}$ for which the answer to the ratio analysis problem with threshold λ is NO.

3.5 Average energy consumption

In the case of many consumer electronics devices, especially mobile phones, battery capacity is severely restricted due to constraints on size and weight of the device. This implies that managing energy well is paramount in such devices. Since most mobile applications are non-terminating (e.g., a web browser), the most important metric for measuring energy consumption is the average memory consumption per time unit [11], e.g., watts per second.

Modeling. We consider that the running time and energy consumption of each statement in the application code is given (or is approximated). In our modeling we split each transition in the *ICFG* into two consecutive transitions, the first is a good event and the next is a bad event. The good event is assigned with a weight that corresponds to the energy consumption of the program statement and the bad event is assigned with a weight that corresponds to the running time of the statement. The average energy consumption of the application is at most λ if and only if the answer to the ratio analysis problem is NO.

4. Algorithm for Quantitative Analysis of QICFGs

In this section we present three results. The mean-payoff analysis problem for QICFGs can be solved in polynomial time, this can be derived from [14]. First, we present an algorithm that significantly improves the current theoretical bound for the problem for QICFGs. Second, we present an efficient algorithm that in most practical cases is much faster as compared to the theoretical upper bound. Finally, we present a linear reduction of the ratio analysis problem to the mean-payoff analysis problem for QICFGs.

4.1 Improved algorithm for mean-payoff analysis

In this section we first discuss the basic polynomial-time algorithm for mean-payoff analysis of QICFGs that can be obtained from the results on pushdown systems shown in [14]. Due to space constraints the technical proofs are relegated to the supplementary material.

Results of [14] and reduction. The results of [14] show that pushdown systems with mean-payoff objectives can be solved in polynomial time. Given a pushdown system with state space Q and stack alphabet Γ , the polynomial-time algorithm of [14] can be described as follows. The algorithm is iterative, and in each iteration it constructs a finite graph of size $O(|Q| \cdot |\Gamma|^2)$ and runs a Bellman-Ford style algorithm on the finite graph from each vertex. The Bellman-Ford algorithm on the finite graph from all vertices in each iteration requires $O(|Q|^3 \cdot |\Gamma|^6)$ time and $O(|Q|^2 \cdot |\Gamma|^4)$ space. The number of iterations required is $O(|Q|^2 \cdot |\Gamma|^2)$. Thus the time and space requirement of the algorithm are $O(|Q|^5 \cdot |\Gamma|^8)$ and $O(|Q|^2 \cdot |\Gamma|^4)$, respectively. A QICFG can be interpreted as a pushdown system where N corresponds to Q and $Retns$ corresponds to Γ .

Theorem 2. (BASIC ALGORITHM [14]). *The mean-payoff analysis problem for QICFGs can be solved in $O(|N|^5 \cdot |Retns|^8)$ time and $O(|N|^2 \cdot |Retns|^4)$ space, respectively.*

Improved algorithm. We will present an improved polynomial-time algorithm for the mean-payoff analysis of QICFGs. The improved algorithm relies on the following properties of QICFGs:

1. The transitions of a module are independent of the stack of a configuration, while in pushdown systems the transitions can depend on the top symbol of the stack. This enables to reduce the size of the finite graphs to be considered in every iteration.
2. Every call node has only one corresponding return node. Therefore, if a module A_1 invokes a module A_2 , then the behavior of A_1 after the termination of A_2 is independent of A_2 . This enables us to reduce the number of iterations to $O(|Calls|)$.

To present the improved algorithm and its correctness formally, we need a refined analysis and extensions of the results of [14]. We first describe a key aspect and present an overview of the solution.

Remark 3. (Infinite-height lattice). *Our algorithm will be an iterative algorithm till some fixpoint is reached. However, for interprocedural analysis with finite-height lattices, fixpoints are guaranteed to exist. Unfortunately in our case for mean-payoff objectives, it is an infinite-height lattice. Thus a fixpoint is not guaranteed. For this reason the analysis for mean-payoff objectives is more involved, and this is even in the case of finite graphs. For example, for reachability objectives in finite graphs linear-time algorithms exist, whereas for finite graphs with mean-payoff objectives the best-known algorithms (for over three decades) are quadratic [23].*

Solution overview. In finite graphs the solution for the mean-payoff analysis is to check whether the graph has a cycle C such that the sum of weights of C is positive. If such a cycle exists, then a *lasso* path that leads to the cycle and then follows the cyclic path forever has positive mean-payoff value. For QICFGs we show that it is enough to find either a loop in the program such that the sum of

weights of the loop is positive or a sequence of calls and returns with positive total weight such that the last invoked module is the same as the first invoked module. For this purpose we compute a *summary function* that finds the maximum weight (according to the sum of weights) path between every two statements of a method (i.e., between every two nodes of a module). The computation is an extension of the Bellman-Ford algorithm to QICFGs. We show that it is enough to compute a summary function for QICFGs with a *stack height* that is bounded by some constant, and then all that is left is to mark pairs of nodes such that the weight of a maximal weight path between them is unbounded. In finite graphs the maximum weight between two vertices is unbounded only if the graph has a cycle with positive sum of weights (i.e., a path with positive total weight that can be pumped). For QICFGs it is also possible to *pump* special types of acyclic paths. We first characterize these pumpable paths (up to Lemma 2). We then show how to compute a bounded summary function (Lemma 3 and the paragraph that follows it). Finally we show how to use the summary function to solve the mean-payoff analysis problem. We start with the basic notions related to stack heights and pumpable paths, and their properties which are crucial for the algorithm.

Stack heights. The *configuration stack height* of $c = (r_1, \dots, r_j, u)$, denoted as $SH(c)$, is j . For a finite path $\pi = \langle (\alpha_1, n_1), \dots, (\alpha_\ell, n_\ell) \rangle$, the stack height of the path (denoted by $SH(\pi)$) is the maximal stack height of all the configurations in the path. Formally $SH(\pi) = \max\{|\alpha_1|, \dots, |\alpha_\ell|\}$. The *additional stack height* of π is the additional height of the stack in the segment of the path, i.e., the additional stack height $ASH(\pi)$ is $SH(\pi) - \max(|\alpha_1|, |\alpha_\ell|)$.

Pumpable pair of paths. Let $\pi = \langle c_1 t_1 t_2 \dots \rangle$ be a finite or infinite path (where each t_i is a transition in the QICFG). A *pumpable pair of paths* for π is a pair of non-empty sequences of transitions: $(p_1, p_2) = (t_{i_1} t_{i_1+1} \dots t_{i_1+\ell_1}, t_{i_2} t_{i_2+1} \dots t_{i_2+\ell_2})$, for $\ell_1, \ell_2 \geq 0$, $i_1 \geq 0$ and $i_2 > i_1 + \ell_1$ such that for every $j \geq 0$ the path $\pi_{(p_1, p_2)}^j$ obtained by pumping the pair p_1 and p_2 of paths j times each is a valid path, i.e., for every $j \geq 0$ we have

$$\pi_{(p_1, p_2)}^j = \langle c_1 t_1 t_2 \dots t_{i_1-1} (p_1)^j t_{i_1+\ell_1+1} \dots t_{i_2-1} (p_2)^j t_{i_2+\ell_2+1} \dots \rangle$$

is a valid path. We illustrate the above definitions with the next example.

Example 4. *Consider the program from Figure 1 and Figure 2 and the corresponding ICFG. A possible path in the program is*

```
m:Entry → while(x) → if(y>0) → m:call foo →
f:Entry → if(x>1) → f:call foo → f:Entry → if(x>1) →
x++ → f:Exit → f:foo ret → x++ → f:Exit →
m:foo ret → while(x)
```

and we denote this path with π . Then $ASH(\pi) = 2$, and the pair of paths $f:Entry \rightarrow if(x>1) \rightarrow f:call foo$ and $f:foo ret \rightarrow x++ \rightarrow f:Exit$ is a pumpable pair of paths.

In the next lemmas we first show that every path with large additional stack has a pumpable pair of paths, and then establish the connection of additional stack height and the existence of pumpable pair of paths with positive weights in Lemma 2. The key intuition for the proof of the next lemma is that a path with $ASH(\pi) > |Calls| + 1$ must contain a recursive call that can be pumped.

Lemma 1. *Let π be a finite path with $ASH(\pi) = d > |Calls| + 1$. Then π has a pumpable pair of paths.*

Lemma 2. *Let c_1, c_2 be two configurations and $j \in \mathbb{Z}$. Let $d \in \mathbb{N}$ be the minimal additional stack height of all paths between c_1 and c_2 with total weight at least j . If $d > |Calls| + 1$, then there exists a path π^* from c_1 to c_2 with additional stack height d that has a pumpable pair (p_1, p_2) with $w(p_1) + w(p_2) > 0$.*

Our algorithm for the mean-payoff analysis problem is based on detecting the existence of certain *non-decreasing* paths with posi-

tive weight. The maximal weights of such non-decreasing paths between node pairs are captured with the notion of a *summary function* and *bounded summary functions* (with bounded additional stack height). We now define them, and establish the lemma related to the number of bounded summary functions to be computed.

Local minimum and non-decreasing paths. A configuration c_i in a path $\pi = \langle c_1, \dots, c_\ell \rangle$ is a *local minimum* if the stack height of c_i is minimal in π , i.e., $|\alpha_i| = \min(|\alpha_1|, \dots, |\alpha_\ell|)$. A path from configuration (α, n_1) to $(\alpha\beta, n_2)$ is a *non-decreasing α -path* if (α, n_1) is a local minimum. Note that if a sequence of transitions is a non-decreasing α -path for some $\alpha \in Retns^*$, then the same sequence of transitions is a non-decreasing γ -path for every $\gamma \in Retns^*$. Hence, we say that π is a non-decreasing path if there exists $\alpha \in Retns^*$ such that π is a non-decreasing α -path.

Summary function. Given the *QICFG* \mathcal{A} and $\alpha \in Retns^*$, we define a summary function $s_\alpha : \bigcup_{1 \leq \ell \leq m} (N_\ell \times N_\ell) \rightarrow \{-\infty\} \cup \mathbb{Z} \cup \{\omega\}$ as:

1. $s_\alpha(n_1, n_2) = z \in \mathbb{Z}$ iff the weight of the maximum weight non-decreasing path from configuration (α, n_1) to configuration (α, n_2) is z .
2. $s_\alpha(n_1, n_2) = \omega$ iff for all $j \in \mathbb{N}$ there exists a non-decreasing path from (α, n_1) to (α, n_2) with weight at least j .
3. $s_\alpha(n_1, n_2) = -\infty$ iff there is no non-decreasing path from (α, n_1) to (α, n_2) .

We note that for every $\alpha, \beta \in Retns^*$ it holds that $s_\alpha \equiv s_\beta$. Hence, we consider only $s \equiv s_\epsilon$ (where ϵ is the empty string and corresponds to empty stack). The computation of the summary function is done by considering stack height bounded summary functions defined below.

Stack height bounded summary function. For every $d \in \mathbb{N}$, the *stack height bounded summary function* $s_d : \bigcup_{1 \leq \ell \leq m} (N_\ell \times N_\ell) \rightarrow \{-\infty\} \cup \mathbb{Z} \cup \{\omega\}$ is defined as follows: (i) $s_d(n_1, n_2) = z \in \mathbb{Z}$ iff the weight of the maximum weight non-decreasing path from (ϵ, n_1) to (ϵ, n_2) with additional stack height at most d is z ; (ii) $s_d(n_1, n_2) = \omega$ iff for all $j \in \mathbb{N}$ there exists a non-decreasing path from (ϵ, n_1) to (ϵ, n_2) with weight at least j and additional stack height at most d ; and (iii) $s_d(n_1, n_2) = -\infty$ iff there is no non-decreasing path with additional stack height at most d from (ϵ, n_1) to (ϵ, n_2) .

Facts of summary functions. We have the following facts: (i) for every $d \in \mathbb{N}$, we have $s_{d+1} \geq s_d$ (monotonicity); and (ii) s_{d+1} is computable from s_d and the *QICFG*. By the above facts we get that if $s_d \equiv s_{d+1}$, i.e., if a fix point is reached, then $s \equiv s_d$. For interprocedural analysis with finite-height lattices, fix points are guaranteed to exist. Unfortunately in our case, the image of s_i is infinite and moreover, it is an infinite-height lattice. Thus a fix point is not guaranteed. The next lemma shows that we can compute all the non- ω values of s with the bounded summary function.

Lemma 3. *Let $d = |Calls| + 1$. For all $n_1, n_2 \in N$, if $s(n_1, n_2) \in \mathbb{Z} \cup \{-\infty\}$, then $s(n_1, n_2) = s_d(n_1, n_2)$.*

By Lemma 3 we get that if $s_{d+1}(n_1, n_2) > s_d(n_1, n_2)$ (for $d = |Calls| + 1$), then $s(n_1, n_2) = \omega$. Hence, the summary function s is obtained by the fix point of the following computation: (i) Compute s_{i+1} from s_i up to s_d for $d = |Calls| + 1$; (ii) for $i \geq |Calls| + 1$, if $s_{i+1}(n_1, n_2) > s_i(n_1, n_2)$, then set $s_{i+1}(n_1, n_2) = \omega$; (iii) a fix point is reached after at most $O(|Calls|)$ iterations (say j iterations), and then we set $s \equiv s_j$. This establishes that we require only $O(|Calls|)$ iterations as compared to $O(|N|^2 \cdot |Retns|^2)$ iterations. The number of returns and calls are the same and thus we significantly improve the number of iterations required from the quartic worst-case bound to linear bound. We now describe the computation of every iteration to obtain s_{i+1} from s_i .

Computation of s_{i+1} from s_i . We first compute a partial function, namely, $s'_{i+1} : En \times Ex \rightarrow \{-\infty, \omega\} \cup \mathbb{Z}$ that satisfies

$s'_{i+1}(n_1, n_2) = s_{i+1}(n_1, n_2)$ for every $n_1 \in En$ and $n_2 \in Ex$. We initialize $s'_0(n_1, n_2) = s_0(n_1, n_2)$. For every module A_ℓ we construct a finite graph G_ℓ^i by taking all the nodes and transitions of A_ℓ and by adding a transition between every call node and its corresponding return node. For every transition between a pair of nodes $n_1, n_2 \in N_\ell \setminus (Calls_\ell \cup Retns_\ell)$ we assign the weight according to the original weight in \mathcal{A} . For every transition between a call node that invokes module A_p and a corresponding return node we assign the weight $s'_i(En_p, Ex_p)$. To compute s'_{i+1} for module A_ℓ we run one Bellman-Ford iteration over G_ℓ^i for source node En_ℓ and target node Ex_ℓ . We observe the next two key properties of s'_i :

- For every iteration i , a module A_ℓ , and pair of nodes $n_1, n_2 \in N_\ell$ we have that the weight of the maximum weight path between n_1 and n_2 in G_ℓ^i is exactly $s_{i+1}(n_1, n_2)$ (the proof is by a simple induction over i).
- If $s'_{i+1} \equiv s'_i$, then $s_{i+1} \equiv s_i$ (follows from the first key property).

Hence, to compute s we compute s'_{i+1} from s'_i until we get $s'_{i+1} \equiv s'_i$, and then we compute all pairs maximum weight paths (e.g., by the Floyd-Warshall algorithm) over every G_ℓ^i and get s_{i+1} (and $s_{i+1} \equiv s$). The Floyd-Warshall algorithm has a cubic time complexity and quadratic space complexity [15]. Therefore, the time complexity for computing every iteration of s_i is $O(\sum |N_\ell|^2)$ and the complexity of the last step is $O(\sum |N_\ell|^3)$. The space complexity of the last step is $O(\max\{|N_1|, \dots, |N_m|\}^2)$, but to store s_i we require $O(\sum |N_\ell|^2)$ space.

Summary graph. Given *QICFG* \mathcal{A} with a summary function s , we construct the *summary graph* $Gr(\mathcal{A}) = (V, E)$ of \mathcal{A} with a weight function $w : E \rightarrow \mathbb{Z} \cup \{\omega\}$ as follows: (i) $V = N \setminus (Ex \cup Retns)$; and (ii) $E = E_{internal} \cup E_{calls}$ where $E_{internal} = \{(n_1, n_2) \mid n_1, n_2 \in N_\ell \text{ for some } \ell, \text{ and } s(n_1, n_2) > -\infty\}$ contains the transitions in the same module and $E_{calls} = \{(n_1, n_2) \mid n_1 \in Calls \text{ and } n_2 \in En \text{ and } n_1 \text{ is a call to a module with entry node } n_2\}$ contains the call transitions. The weights of $E_{internal}$ are according to the summary function s and the weights of E_{calls} are according to the weights of these transitions in \mathcal{A} (i.e., according to w). A simple cycle in $Gr(\mathcal{A})$ is a *positive simple cycle* iff one of the following conditions hold: (i) the cycle contains an ω edge; or (ii) the sum of the weights of the cycles according to the weights of the summary graph is positive. Lemma 4 shows the equivalence of the mean-payoff analysis problem and positive cycles in the summary graph.

Lemma 4. *A QICFG \mathcal{A} has a path π with $LimAvg(w(\pi)) > 0$ iff the summary graph $Gr(\mathcal{A})$ has a (reachable) positive cycle.*

Algorithm and analysis. Algorithm 1 solves the mean-payoff analysis problem for *QICFGs*. The computation of the summary function requires $O(|Calls|)$ computations of the partial summary function s'_i , which requires m runs of Bellman-Ford algorithm, each run over a graph of size $|N_\ell|$ (hence, each run takes $O(|N_\ell|^2)$ time). In addition the computation requires m runs of all pairs maximum weight path (Floyd-Warshall) algorithm. Each run is over a graph of size $O(|N_\ell|)$ (hence, each run takes $O(|N_\ell|^3)$ time and $O(|N_\ell|^2)$ space). Finally we detect positive cycles by running Bellman-Ford algorithm once over the summary graph, which takes $O(|N|^2)$ time and $O(|N|)$ space. Thus we obtain the following result.

Theorem 3. (IMPROVED ALGORITHM). *Algorithm 1 solves the mean-payoff analysis problem for QICFGs in $O((|Calls| \cdot (\sum |N_\ell|^2)) + (\sum |N_\ell|^3) + |N|^2)$ time and $O(\sum |N_\ell|^2)$ space.*

Remark 4. *Note that in the worst case the running time of Algorithm 1 is cubic and the space requirement is quadratic.*

Algorithm 1 Mean-payoff *QICFG* Analysis

```
for  $\ell \leftarrow 1$  to  $m$  do
   $s'_0(En_\ell, Ex_\ell) \leftarrow$  BELLMAN-FORD( $A_\ell$ ) {Compute  $s'_0$  by
  running Bellman-Ford algorithm on  $A_\ell$ }
end for
 $i \leftarrow 1$ 
loop
  for  $\ell \leftarrow 1$  to  $m$  do
    Construct  $G_\ell^{i-1}$  according to  $s'_{i-1}$ 
     $s'_i(En_\ell, Ex_\ell) \leftarrow$  BELLMAN-FORD( $G_\ell^{i-1}$ ) {Compute  $s'_i$ 
    by running Bellman-Ford algorithm over  $G_\ell^{i-1}$ }
  end for
  if  $s'_i \equiv s'_{i-1}$  then
    break
  end if
  if  $i > |Calls| + 1$  then
    for  $\ell \leftarrow 1$  to  $m$  do
      if  $s'_i(En_\ell, Ex_\ell) > s'_{i-1}(En_\ell, Ex_\ell)$  then
         $s'_i(En_\ell, Ex_\ell) = \omega$ 
      end if
    end for
  end if
   $i \leftarrow i + 1$ 
end loop
 $s \leftarrow$  FLOYD-WARSHALL( $s'_i$ )
Construct  $Gr(\mathcal{A})$  from  $s$ 
BELLMAN-FORD( $Gr(\mathcal{A})$ ) {Run Bellman-Ford over  $Gr(\mathcal{A})$ }
if  $Gr(\mathcal{A})$  has a positive cycle then
  return YES
else
  return NO
end if
```

4.2 Efficient algorithm for mean-payoff analysis

In this section we further improve the algorithm for the mean-payoff analysis problem for *QICFGs*, and the improvement depends on the fact that typically the number of entry, exit, call, and returns nodes is much smaller than the size of the *QICFGs*. Formally, in most typical cases we have $|Ex \cup Retns \cup Calls \cup En| \ll |N|$. Let $X_\ell = \{Ex_\ell, En_\ell\} \cup Retns_\ell \cup Calls_\ell$ and $X = \bigcup_\ell X_\ell$. We present an improvement that enables us to construct the summary function over graphs of size $O(|X_\ell|)$ (instead of graphs of size $O(|N_\ell|)$ of Section 4.1), and with at most $O(|Calls|)$ iterations. Hence, the algorithm in most typical cases will be much faster and require much smaller space.

Compact representation. The key idea for the improvement is to represent the modules in *compact form*. The compact form of a module A_ℓ , denoted by $Comp(A_\ell)$, consists of the entry, exit, call, and returns node of A_ℓ . There is transition between every node in $Comp(A_\ell)$, and the weight of each transition is the maximum weight path between the nodes with additional stack height 0 (and if there is no such path, then the weight is $-\infty$). Formally, $Comp(A_\ell) = (V, E)$; where $V = X_\ell$; $E = V \times V$, and $w(v_1, v_2) = s_0(v_1, v_2)$ (where s_0 is the bounded height summary function of height 0). If in $Comp(A_\ell)$ there is a cycle with positive weight that is reachable from the entry node, then we say that A_ℓ is a *positive mean-payoff witness*. The computation of the compact form for a module A_ℓ requires $O(|X_\ell| \cdot |N_\ell|^2)$ time and $O(|N_\ell|)$ space (running Bellman-Ford on each A_ℓ), and thus the compact form for all modules can be computed in $O(\sum |X_\ell| \cdot |N_\ell|^2)$ time and $O(\max |N_\ell|)$ space (note that the space can be reused).

Witness in summary graph of compact forms. After constructing the compact forms, we compute a summary function for $Comp(A_1), \dots, Comp(A_m)$, and a corresponding summary graph. We say that there is a path with positive mean-payoff iff

there exists a positive cycle in the summary graph or there exists a path to the entry node of a positive mean-payoff witness. The correctness of the algorithm relies on the next lemma.

Lemma 5. Let $\mathcal{A} = \langle A_1, \dots, A_m \rangle$ be a *QICFG*, let $Gr(\mathcal{A})$ be its summary graph and let $Comp(Gr(\mathcal{A}))$ be the summary graph that is formed by $Comp(A_1), \dots, Comp(A_m)$. The following assertions are equivalent:

1. $Gr(\mathcal{A})$ has a (reachable) positive cycle.
2. $Comp(Gr(\mathcal{A}))$ has a (reachable) positive cycle or a positive mean-payoff witness.

The above lemma establishes the correctness of the computation on compact form graphs, and gives us the following result. The following result is obtained from Theorem 3 by replacing $|N_\ell|$ with $|X_\ell|$ and $|N|$ by $|X|$, and the additional $\sum |X_\ell| \cdot |N_\ell|^2$ time and $\max |N_\ell|$ space are required for the compact form computation.

Theorem 4. (EFFICIENT ALGORITHM). *The mean-payoff analysis problem for QICFGs can be solved in $O(|Calls| \cdot (\sum |X_\ell|^2) + (\sum |X_\ell|^3) + |X|^2 + \sum |X_\ell| \cdot |N_\ell|^2)$ time and $O(\sum |X_\ell|^2 + \max |N_\ell|)$ space, where $X_\ell = \{Ex_\ell, En_\ell\} \cup Retns_\ell \cup Calls_\ell$ and $X = \bigcup_\ell X_\ell$.*

4.3 Reduction: Ratio analysis to mean-payoff analysis

We now establish a linear reduction of the ratio analysis problem to the mean-payoff analysis problem. Given a *QICFG* \mathcal{A} with labeling of good, bad, and neutral events, a positive integer weight function w , and rational threshold $\lambda > 0$, the reduction of the ratio analysis problem to the mean-payoff analysis problem is as follows. We consider a weight function w_λ for the mean-payoff objective defined as follows: for a transition e we have

$$w_\lambda(e) = \begin{cases} w(e) & \text{if } e \text{ is labelled with good} \\ -\lambda \cdot w(e) & \text{if } e \text{ is labelled with bad} \\ 0 & \text{otherwise (if } e \text{ is labelled with neutral)} \end{cases}$$

The next lemma establishes the correctness of the reduction.

Lemma 6. *Given a QICFG \mathcal{A} with a positive integer weight function w , and a rational threshold $\lambda > 0$:*

There exists a path π in \mathcal{A} with $LimRat(w(\pi)) > \lambda$ iff there exists a path π with $LimAvg(w_\lambda(\pi)) > 0$.

Remark 5. *Note that in our reduction from ratio analysis to mean-payoff analysis we do not change the QICFG, but only change the weight function. Thus our algorithms from Theorem 3 and Theorem 4 can also solve the ratio analysis problem for QICFGs. Moreover, our proof of lemma 6 shows that for all paths π , if we have $LimAvg(w_\lambda(\pi)) > 0$, then we also have $LimRat(w(\pi)) > \lambda$, i.e., any witness for the mean-payoff analysis is also a witness for the ratio analysis.*

5. Experimental Results: Two Case Studies

In this section we present our experimental results on two case studies described in Section 3.1 and Section 3.2. We run our case studies on several benchmarks in Java, including DaCapo 2009 benchmarks [4], and we use [7, 25] to assist Soot for the construction of the control-flow graphs. First we present some optimizations that proved useful for speed-up in the benchmarks.

5.1 Optimization for case studies

We present four optimizations for the case studies: the first two are general, and the last two are specific to our case studies.

Faster computation of bounded height summary function. We note that if module A_ℓ invokes only modules A_{j_1}, \dots, A_{j_k} , and $s'_i(En_{j_h}, Ex_{j_h}) = s'_{i-1}(En_{j_h}, Ex_{j_h})$ for all $h \in \{1, \dots, k\}$, then $s'_i(En_\ell, Ex_\ell) = s'_{i-1}(En_\ell, Ex_\ell)$. Hence, when computing s'_i , we

maintain a set $\mathcal{L}^i = \{\ell \mid s'_i(En_\ell, Ex_\ell) > s'_{i-1}(En_\ell, Ex_\ell)\}$, and in the next iteration we run Bellman-Ford algorithm only for the modules that invoke modules from \mathcal{L}^i .

Reducing the number of iterations for fix point. We now present an optimization that allows us to reduce the number of bounded height summary functions from $O(|Calls|)$ to a practically constant number. We note that the $O(|Calls|)$ theoretical bound is tight. However, only pathological cases can reach even a fraction of this bound. We note that in typical programs the average nesting of function calls is practically constant (say 10). So if we do not get a fix point after 10 iterations (i.e., $s'_{11} > s'_{10}$), then it is probably because there is a recursive call with positive weight. If this is the case, then if we build the summary graph according to s_{11} , we will get a positive cycle in the summary graph, that is, we will get a witness for a path with a positive mean-payoff, and we can stop the computation (since by definition $s \geq s_{11}$, we get that this witness is valid). Hence, our optimized algorithm is to compute the bounded height summary function s'_i and if $s'_i > s'_{i-1}$ and $i = 10, 20, 30, \dots$, then we construct the summary graph and look for a witness path. If a path is found, then we are done. Otherwise we continue and compute s'_{i+1} .

Removing redundant modules. Consider an *ICFG* $\mathcal{A} = \langle A_1, \dots, A_m \rangle$ in which every node is reachable from the program entry (the entry node of the main method). We say that module A_i is *non-redundant* if (i) the module has non-zero weight transitions (good or bad events); or (ii) it invokes a non-redundant module, and is called *redundant* otherwise. Let A_i be a redundant module. For every path π that contains a transition to En_i (an invocation of A_i), the segment of π between that transition and the first transition to Ex_i contains only neutral transitions. Because all nodes of \mathcal{A} are reachable, we can safely replace each call node that invokes A_i by an internal node that leads to the corresponding return node, and label it as a neutral event. Our optimization then consists of removing redundant modules, as follows:

1. First, we perform a single-source interprocedural reachability from the program entry, which requires linear time ([31]), and discard all non-reachable nodes in all modules.
2. Then, we perform a backwards reachability computation on the call graph of \mathcal{A} , starting from the set of all modules that contain non-zero weight transitions. All detected redundant modules are discarded, and calls to them are replaced according to the above description.

Hence, when computing the bounded height summary function, the size of the graph is smaller and the Bellman-Ford algorithm takes less time. Additionally, the number of calls $|Calls|$ decreases, which reduces the number of iterations required in the main loop of Algorithm 1. In the first case study, typically more than half of the methods are eliminated in this process.

Incremental computation of summary functions. We present the final optimization which is relevant for our second case study. Let \mathcal{A}^1 be a *QICFG* and let \mathcal{A}^2 be a *QICFG* that is obtained from \mathcal{A}^1 only by increasing some of the transitions weights. Let s^1 be the summary function of \mathcal{A}^1 . Then we can compute the summary function of \mathcal{A}^2 by setting $s_0^2 \equiv s^1$ and by computing s_i^2 from s_{i+1}^2 in the usual way. The correctness is almost trivial. Since the weights of \mathcal{A}^2 are at least as the weights of \mathcal{A}^1 , we get that if we conceptually add a transition (n_1, n_2) with weight $s^1(n_1, n_2)$ for every two nodes (in the same module) in \mathcal{A}^2 , then the weights of the paths with the maximal weight in \mathcal{A}^2 remain the same. By assigning $s_0^2 = s^1$ we only add such conceptual transitions. Hence, the correctness follows. We now describe how this optimization speed up the analysis of the second case study. In the static profiling for function frequencies, we need to build a summary graph for every function f , and then run the mean-payoff analysis for every such graph. Given this optimization, we can first compute (only once) a summary graph for the case that all method

invocations are bad events. We denote this *QICFG* by \mathcal{A}^* and the corresponding summary function by s^* . Note that in \mathcal{A}^* all weights are negative, and the mean-payoff analysis answer is trivially NO. But still the summary function computation, which computes the quantitative information about the maximum weight context-free paths, provides useful information and saves recomputation. To analysis the frequency of f we assign weights to \mathcal{A} and get \mathcal{A}^f . We note that the difference between \mathcal{A}^f and \mathcal{A}^* is only in the weight that is assigned to the invocation of f . We then compute the summary function s^f for \mathcal{A}^f by first assigning $s_0^f = s^*$. In practical cases, programs can have thousands of methods, but only small portion of them will have a path to f . So along with the previous optimizations we get that only few Bellman-Ford runs are required to compute s^f . Overall, the computation of s^* is expensive, and may take several minutes for a large program, but it is done only once, and for every method the computation of s^f is much faster.

5.2 Container analysis

Technical details about experimental results. We discuss a few relevant details about our experiments and results.

- We use the points-to analysis tool of [36]. This tool provides interprocedural on demand analysis for a may-alias relationship of two variables. We say that a variable may point to an allocated container if it may-alias the container, and a variable must point to an allocated container if it may-alias only one allocated container.
- For the underutilized containers the threshold is 1, and for the analysis of overpopulated containers we set a threshold of 0.1 for our experimental results. That is, if the ratio between the number of added elements to the number of lookup operations is more than 10, then the container is overpopulated.

Experimental results. Our experimental results on the benchmarks are reported in Table 1. In the table, # **M** and # **CO** represent the number of methods and containers that are reachable from the main entry of the program, respectively; # **OP** and # **UC** represent the number of overpopulated and underutilized containers discovered by our tool, respectively; and **TA(s)** and **TQ(s)** represent the time required for alias analysis and the time required for the quantitative analysis of *QICFGs* (in seconds), respectively; and the entries of the respective columns represent the time for overpopulated/underutilized container analysis. We now highlight some interesting aspects of our experimental results. First, our approach for container analysis discovers containers that are overpopulated or underutilized, while maintaining soundness. Second, the cases that we identify reveal useful information for optimization, for example, in the first (batik-rasterizer) and the second (batik-svgpp) benchmarks we identify containers that always have a small bounded number of elements.

Benchmark	# M	# CO	# OP	# UC	TA(s)	TQ(s)
batik-rasterizer	21433	9	1	2	124/125	144/143
batik-svgpp	7859	3	0	3	20/20	14/13
mrt	9798	10	1	0	70/13	41/59
java_cup	8173	10	0	0	19/19	25/22
xalan	8729	6	3	2	5/5	41/43
polyglot	8068	8	2	2	0/0	17/17
antlr	8607	15	5	2	11/12	25/24
jflex	21852	43	3	6	2473/2614	178/210
avrorra	13331	75	9	9	145/141	111/113
muffin	22503	50	3	5	2500/157	352/173
bloat06	10675	211	32	14	399/250	2241/2165
eclipse06	9335	74	8	4	37/22	222/164
jython06	12210	66	9	5	154/68	13593/8376

Table 1. Experimental results for container usage analysis

Benchmark	#M	#I	T
antlr	768	326	1.2
bloat	2576	676	30.8
eclipse	1056	215	2.3
fop	429	47	0.4
luindex	567	239	0.7
lusearch	842	237	2.5
pmd	2547	589	11.5

Table 2. Experimental results for frequency of functions.

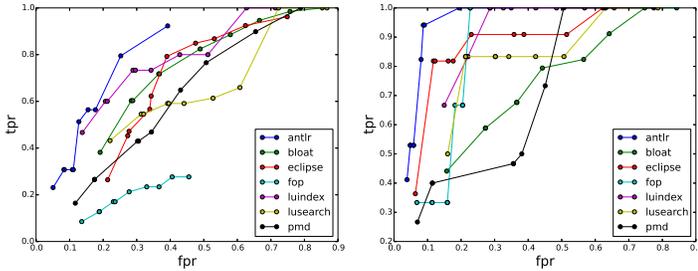


Figure 4. The ROC curves. The left plot shows the results when all methods are analyzed. The right plot shows the results when only the active methods are analyzed.

Comparison with the work of [45]. While the notion of underutilized and overpopulated containers is the same in our work and in [45], the concrete mathematical definition is different.

- (*Conceptual difference in modeling*). The formal definitions in [45] relies on properties of a specifically constructed *inequality graph* that in practice gives a good approximation on whether a container is properly used (for details, see Definition 8 in [45]). Our formal definition is conceptually simpler and relies on a very general and flexible mathematical framework, but on the other hand it crucially relies on the accuracy of the constructed *QICFG*. For example, our technique may not report a container as underutilized even if the witness path for proper utilization is very complex, e.g., has 20 nested function calls, and therefore it is unlikely to be realizable on practice.
- (*Advantages and disadvantages*). Our approach has several advantages: (1) First, our approach can handle recursion, whereas [45] does not handle recursion. (2) Second, we present a sound and complete approach for ration analysis of *QICFGs*, and with a conservative modeling we have a sound analysis approach for detecting containers misuse. (3) Third, our algorithm is polynomial time (once the points-to relation is computed), whereas the algorithmic approach of [45] in the worst case is exponential. (4) Finally, our approach also allows us to handle DELETE operations: in [45] loops with only ADD operations (in ADD vs DELETE) were identified as proper usage, whereas we identify loops where the difference of the ADD and DELETE operations is positive as proper usage (this subsumes the loops of [45] and for example, also loops with two ADD operations and one DELETE operation). Example 3 illustrates the advantages of our approach. One drawback of our approach is that it is conservative: while for under-utilized containers analysis (ADD vs DELETE) our approach captures all cases of [45] (as explained above), our approach for over-populated containers analysis is more conservative (to obtain soundness).
- (*Comparison of experimental outcomes*). With our approach we were able to fully analyze all the containers in all benchmarks, whereas in [45] the analysis for few benchmarks (e.g., muffin) was not done for all containers since a timeout was reached. Below we present example snippets of code from the benchmarks where our analysis gives different results from the

```

1 public void cleared() {
2     ...
3     if (list != null){
4         ...
5     }
6     else{
7         Object o = elementsById.remove(id);
8         if (o != this) // oops not us!
9             elementsById.put(id, o);
10    }
11 }
12 public void run() {
13     while(true) {
14         ...
15         if (...) {
16             ...
17             re.cleared();
18         }
19         ...
20     }
21 }

```

Figure 5. An example from benchmark batik. The method `run` invokes `cleared` in a loop, and in every invocation, one element of `elementsById` is removed and one element is added. Thus in this loop the total number of elements in `elementsById` is bounded.

```

1 void addCNAME(CNAMERecord cname) {
2     if (backtrace == null)
3         backtrace = new Vector();
4     backtrace.insertElementAt(cname, 0);
5 }
6 public SetResponse findRecords(Name name, short type) {
7     ...
8     if (type != Type.CNAME && type != Type.ANY && rrset.getType() == Type.CNAME)
9     {
10        zr = findRecords(cname.getTarget(), type);
11        zr.addCNAME(cname);
12        ...
13    }
14    ...
15    return zr;
16 }
17 }

```

Figure 6. An example from benchmark muffin. The method `findRecords` has a recursive call, and method `addCNAME` adds an element to vector backtrace. A path with recursion depth n adds n elements to backtrace. Hence, backtrace may have unbounded number of elements and it is not underutilized.

analysis of [45]. The example in Figure 5 shows that handling DELETE operations leads to more refined analysis: in the example, if DELETE operations are not handled, then the misuse is not detected. The example in Figure 6 shows that the proper utilization of containers might depend on the recursive calls. Finally, the example in Figure 7 illustrates that the proper use of containers can be outside its allocation site.

5.3 Static profiling: frequency of function calls

Experimental results. We examined ten thresholds, namely $1/30, 2/30, 3/30, \dots, 10/30$, and for each threshold λ we say that a method is *statically hot* if it is λ -hot (according to the definition in Subsection 3.2). We compared the results to dynamic profiling from the DaCapo benchmarks [4]. In the dynamic profiling we define the top 5% of the most frequently invoked functions as *dynamically hot*. For example, if a program has 1000 functions, and in the benchmark 500 functions were invoked at least once, then the 25 most frequently invoked functions are dynamically hot. We note that theoretically speaking, the definitions of dynamic and static hotness are incomparable (basically for any λ), but our experimental results show a good correlation between the two notions. To illustrate the correlation we treat our static analysis as a *classifier* of hot methods, and the *specificity* and *sensitivity* of the classifier are controlled by the threshold λ . The sensitivity of a classifier is measured by the *true positive rate* (tpr), which is

$$\frac{\text{\#dynamically hot methods that are reported as statically hot}}{\text{\#dynamically hot methods}}$$

The specificity is uniquely characterized by the *false positive rate* (fpr), namely,

$$\frac{\text{\#non-dynamically hot methods that are reported as statically hot}}{\text{\#methods}}$$

For high values of λ , the classifier is expected to capture only dynamically hot methods (but it will miss most of the dynamically

```

1 Hashtable cgi(Request request)
2 {
3     Hashtable attrs = new Hashtable(13);
4     ...
5     if (query != null)
6     {
7         StringTokenizer st = new StringTokenizer (decode(query), "&");
8         while (st.hasMoreTokens ())
9         {
10            ...
11            attrs.put(key, value);
12        }
13    }
14    ...
15    return attrs;
16 }
17
18 public Reply recvReply (Request request)
19 {
20     ...
21     else if (request.getPath().equals("/admin/set"))
22     {
23         Hashtable attrs = cgi(request);
24         ...
25         for (int i = 0; i < enabled.size(); i++)
26         {
27             ...
28             prefs.put(key, (String) attrs.get(key));
29         }
30         ...
31     }
32     ...
33     return reply;
34 }

```

Figure 7. An example from benchmark muffin. The method `cgi` allocates the container `attrs` and potentially adds it many elements. The method `recvReply` performs a `get` operation over `attrs` in a loop. Since we analyze not only the operations that are nested in the allocation site, we detect that `attrs` is not overpopulated (the analysis of [45] reports it as overpopulated).

hot methods), and thus it will have very high fpr but very low tpr. For very low values of λ the classifier will report most of the methods as hot, so most of the hot methods will be reported as hot, and we will have very high tpr but very low fpr. A fundamental metric for classifier evaluation is a *receiver operating characteristic (ROC) graph*. A ROC graph is a plot with the false positive rate on the X axis and the true positive rate on the Y axis. The point (0,1) is the perfect classifier, and the area beneath an ROC curve can be used as a measure of accuracy of the classifier.

In our experimental evaluation we only considered application functions (and not the library functions), and the results are presented in Table 2. In the table, $\#M$ represents the number of application methods (that are reachable from the main entry of the program), $\#I$ represents the number of application methods that were actually invoked in the benchmark, T represents the average running time for the static analysis of a single method (i.e., to check whether a single method is λ -hot for a fixed λ) (in seconds). For each λ we present the tpr and fpr values of the classifications. We present an evaluation for two cases. In the first case we statically analyze all the methods and calculate the tpr and fpr accordingly. In the second case we consider only the *active methods*, namely, the methods that were invoked at least once in the program and we remove all the other methods from the program control flow graph. This analysis simulates a typical case where the programmer has a prior knowledge on which methods are definitely not hot and can (manually) instruct the static analysis to ignore them. The ROC curves are presented in Figure 4, where the most left points on the graph are for $\lambda = 10/30$ and the fpr and tpr increases as λ decreases (until it finally reaches $1/30$). In general, for most of the programs the static analysis gives useful and quite accurate information. Specifically, the threshold $\lambda = 7/30$ managed to capture more than half of the hot methods for most benchmarks (namely, other than fop and antlr) and a false positive rate less than 0.3 which means that if a method was statically reported as not hot, then with probability 0.7 it is really not hot. We note that the analysis over fop gives quite poor results because only 10% of the methods were active. However, when we analyzed only the active methods we get better results for fop, see the right hand graph in Figure 4. When we

only consider the active methods, the threshold $\lambda = 9/30$ managed to capture most of the dynamically hot methods while maintaining a false positive rate less than 0.1 (for most programs).

Remarks. We run the experiments over a single thread Intel Pentium 3.80GHz. For Table 1 results, for some benchmarks (such as fop, pmd) the alias analysis did not complete. For Table 2 results, we only show results for benchmarks for which we managed to obtain the dynamic profiles. Also for a few benchmarks (such as jython) the quantitative analysis took too long for the entire benchmark, and in such cases, our tool could be used to focus on specific functions.

6. Related Work

Interprocedural analysis. Algorithms that operate on the interprocedural control-flow graphs provide the framework for static analysis of programs, and have numerous applications. Precise interprocedural analysis is crucial for dataflow analysis and has been studied in several works [29, 31]. The study of interprocedural analysis has also been extended to weighted pushdown systems, where the weight domain is a bounded idempotent semiring [10, 33]. Analysis of such weighted pushdown systems has been used in many applications of program analysis [24, 32, 33]. Our work is different because the objectives (mean-payoff and ratio analysis) we consider are very different from reachability and bounded domains. The mean-payoff objective is a function that assigns a real-valued number to every path. In contrast to bounded domain functions, the range of a mean-payoff function is potentially uncountable. We develop novel techniques to extend the summary graph approach for finite-height lattices to solve mean-payoff analysis of *QICFGs* (which requires computing fix points for infinite-height lattices).

Mean-payoff analysis. Mean-payoff objectives are quantitative metrics for performance modeling in many applications and very well-studied in the context of finite-state graphs and games. Finite-state graphs and games with mean-payoff objectives have been studied in [19, 21, 23, 46] for performance modeling, and robust synthesis of reactive systems [5, 6]. Quantitative abstraction-refinement frameworks for finite-state systems with mean-payoff objectives have also been studied in [12]. While the mean-payoff objectives have been considered in depth for finite-state systems, they have not been considered in depth for interprocedural analysis. Pushdown systems with mean-payoff objectives were considered in [14]. We significantly improve the complexity of the polynomial-time algorithm for interprocedural mean-payoff analysis that can be obtained by a reduction to the results of [14].

Detecting inefficiently-used containers. Bloat detection and detecting inefficiently used containers have been identified in many previous works as a major reason for program inefficiency. Dynamic approaches for the problem were studied in many works such as [18, 27, 28, 30, 34, 35, 44]. A static approach to analyze the problem was first considered in [45], which is the most closely related work to our case study. The work of [45] provides an excellent exposition of the problem with several practical motivations. It also describes the clear advantages of the static analysis tools, and identifies that soundness in detecting inefficiently used containers (with no or low false positive rates) is a very important feature. Our approach for the problem is significantly different from the approach of [45]. A big part of the contribution of [45] is an automated annotation for the functionality of the containers operation. The main algorithmic approach of [45] is to use CFL-reachability (context-free reachability) to identify nesting loop depths and then use this information for detecting misuse of containers. Our algorithmic approach is very different: we use a quantitative analysis approach, i.e., ratio analysis of *QICFGs* to model the problem.

Static profiling of programs. Static and dynamic profiling of programs is in the heart of program optimization. Static profiling are typically used in branch predictions where the goal is to assign probabilities to branches, and typically require some prior knowl-

edge on the probability of inputs. Static profiling of programs for branch predictions has been considered in [2, 22, 39, 42]. Dynamic profiling has also been used in many applications related to performance optimizations, see [40] for a collection of dynamic profiling tools. Two main drawbacks of dynamic profiling are that they require inputs, and they cannot be used for compiler optimizations. We use static profiling to determine if a function is invoked frequently along some run of the program, and do not require any prior knowledge on inputs. The techniques used in [2, 39, 42] involves solving linear equations with sparse matrix solvers, whereas our solution method is different (by quantitative analysis of *QICFGs*).

7. Conclusion

In this work we considered the quantitative (ratio and mean-payoff) analysis for interprocedural programs. We demonstrated how interprocedural quantitative analysis can aid to automatically reason about properties of programs and potential program optimizations. We significantly improved the theoretical known upper-bound for the polynomial-time solution, and presented several practical optimizations that proved to be useful in real programs. We have implemented the algorithm in Java, and showed that it scales to DaCapo benchmarks of real-world programs. This shows that interprocedural quantitative analysis is feasible and useful. Some possible directions of future works are as follows: (1) extend our framework with multiple quantitative objectives and study their applications; and (2) extend [12] to have an abstraction-refinement framework for quantitative interprocedural analysis.

References

- [1] G. M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference*, pages 483–485. ACM, 1967.
- [2] T. Ball and J. R. Larus. Branch prediction for free. In *PLDI*, pages 300–313, 1993.
- [3] H. Bjorklund, S. Sandberg, and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. In *MFCSS'04*, pages 673–685, 2004.
- [4] S. M. Blackburn, R. Garner, C. Hoffmann, A. M. Khang, 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. *SIGPLAN Not.*, 41(10):169–190, 2006.
- [5] R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In *CAV*, pages 140–156, 2009.
- [6] R. Bloem, K. Greimel, T. A. Henzinger, and B. Jobstmann. Synthesizing robust systems. In *FMCAD*, pages 85–92, 2009.
- [7] E. Bodden, A. Sewe, J. Sinschek, M. Mezini, and H. Oueslati. Taming reflection: Aiding static analysis in the presence of reflection and custom class loaders. In *ICSE '11*, pages 241–250. ACM, 2011.
- [8] M. S. Boddy. Anytime problem solving using dynamic programming. In *AAAI*, pages 738–743, 1991.
- [9] U. Boker, K. Chatterjee, T. A. Henzinger, and O. Kupferman. Temporal specifications with accumulative values. In *LICS*, 2011.
- [10] A. Bouajjani, J. Esparza, and T. Touili. A generic approach to the static analysis of concurrent programs with procedures. In *POPL*, 2003.
- [11] A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *USENIX*, 2010.
- [12] P. Cerný, T. A. Henzinger, and A. Radhakrishna. Quantitative abstraction refinement. In *POPL*, pages 115–128, 2013.
- [13] K. Chatterjee, L. Doyen, and T. A. Henzinger. Quantitative languages. *ACM Trans. Comput. Log.*, 11(4), 2010.
- [14] K. Chatterjee and Y. Velner. Mean-payoff pushdown games. In *LICS*, pages 195–204, 2012.
- [15] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.
- [16] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.
- [17] M. Droste and I. Meinecke. Describing average- and longtime-behavior by weighted MSO logics. In *MFCSS*, pages 537–548, 2010.
- [18] B. Dufour, B. G. Ryder, and G. Sevitsky. A scalable technique for characterizing the usage of temporaries in framework-intensive java applications. In *SIGSOFT FSE*, pages 59–70, 2008.
- [19] A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *Int. Journal of Game Theory*, 8(2):109–113, 1979.
- [20] C. Ferdinand, F. Martin, R. Wilhelm, and M. Alt. Cache behavior prediction by abstract interpretation. *Sci. Comput. Program.*, 99.
- [21] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
- [22] J. L. Hennessy and D. A. Patterson. *Computer Architecture, Fourth Edition: A Quantitative Approach*. Morgan Kaufmann, 2006.
- [23] R. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Mathematics*, 23:309–311, 1978.
- [24] A. Lal, T. W. Reps, and G. Balakrishnan. Extended weighted pushdown systems. In *CAV*, pages 434–448, 2005.
- [25] O. Lhoták and L. J. Hendren. Scaling java points-to analysis using spark. In *CC*, pages 153–169, 2003.
- [26] T. A. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *Siam Review*, 11:604–607, 1969.
- [27] N. Mitchell and G. Sevitsky. The causes of bloat, the limits of health. In *OOPSLA*, pages 245–260, 2007.
- [28] N. Mitchell, G. Sevitsky, and H. Srinivasan. Modeling runtime behavior in framework-based applications. In *ECOOP*, 2006.
- [29] M. Müller-Olm and H. Seidl. Precise interprocedural analysis through linear algebra. In *POPL*, pages 330–341, 2004.
- [30] G. Novark, E. D. Berger, and B. G. Zorn. Efficiently and precisely locating memory leaks and bloat. In *PLDI*, pages 397–407, 2009.
- [31] T. W. Reps, S. Horwitz, and S. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL*, pages 49–61, 1995.
- [32] T. W. Reps, A. Lal, and N. Kidd. Program analysis using weighted pushdown systems. In *FSTTCS*, pages 23–51, 2007.
- [33] T. W. Reps, S. Schwoon, S. Jha, and D. Melski. Weighted pushdown systems and their application to interprocedural dataflow analysis. *Sci. Comput. Program.*, 58(1-2):206–263, 2005.
- [34] O. Shacham, M. T. Vechev, and E. Yahav. Chameleon: adaptive selection of collections. In *PLDI*, pages 408–418, 2009.
- [35] A. Shankar, M. Arnold, and R. Bodík. Jolt: lightweight dynamic analysis and removal of object churn. In *OOPSLA*, 2008.
- [36] M. Sridharan and R. Bodík. Refinement-based context-sensitive points-to analysis for java. In *PLDI*, pages 387–400, 2006.
- [37] V. Tiwari, S. Malik, and A. Wolfe. Power analysis of embedded software: a first step towards software power minimization. *IEEE Trans. VLSI Syst.*, 2(4):437–445, 1994.
- [38] R. Vallée-Rai, P. Co, E. Gagnon, L. Hendren, P. Lam, and V. Sundaresan. Soot - a java bytecode optimization framework. In *CASCON'99*.
- [39] T. A. Wagner, V. Maverick, S. L. Graham, and M. A. Harrison. Accurate static estimators for program optimization. In *PLDI*, 1994.
- [40] Wikipedia. List of performance analysis tools.
- [41] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. B. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. P. Puschner, J. Staschulat, and P. Stenström. The worst-case execution-time problem - overview of methods and survey of tools. *ACM Trans. Embedded Comput. Syst.*, 7(3), 2008.
- [42] Y. Wu and J. R. Larus. Static branch frequency and program profile analysis. In *MICRO 27*, pages 1–11. ACM, 1994.
- [43] Y. Xie, M. Naik, B. Hackett, and A. Aiken. Soundness and its role in bug detection systems. In *Proc. of the Workshop on the Evaluation of Software Defect Detection Tools*, 2005.
- [44] G. H. Xu, N. Mitchell, M. Arnold, A. Rountev, E. Schonberg, and G. Sevitsky. Finding low-utility data structures. In *PLDI*, 2010.
- [45] G. H. Xu and A. Rountev. Detecting inefficiently-used containers to avoid bloat. In *PLDI*, pages 160–173, 2010.
- [46] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.