DJXPerf: Identifying Memory Inefficiencies via Object-centric Profiling for Java

Bolun Li
North Carolina State University
bli35@ncsu.edu

Pengfei Su
University of California, Merced
psu9@ucmerced.edu

Milind Chabbi
Scalable Machines Research
milind@scalablemachines.org

Shuyin Jiao
North Carolina State University
sjiao2@ncsu.edu

Xu Liu
North Carolina State University
xliu88@ncsu.edu

arXiv:2104.03388v1 [cs.PF] 7 Apr 2021

Abstract

Java is the "go-to" programming language choice for developing scalable enterprise cloud applications. In such systems, even a few percent CPU time savings can offer a significant competitive advantage and cost saving. Although performance tools abound in Java, those that focus on the data locality in the memory hierarchy are rare.

In this paper, we present DJXPerf, a lightweight, object-centric memory profiler for Java, which associates memory-hierarchy performance metrics (e.g., cache/TLB misses) with Java objects. DJXPerf uses statistical sampling of hardware performance monitoring counters to attribute metrics to not only source code locations but also Java objects. DJXPerf presents Java object allocation contexts combined with their usage contexts and presents them ordered by the poor locality behaviors. DJXPerf's performance measurement, object attribution, and presentation techniques guide optimizing object allocation, layout, and access patterns. DJXPerf incurs only ~8% runtime overhead and ~5% memory overhead on average, requiring no modifications to hardware, OS, Java virtual machine, or application source code, which makes it attractive to use in production. Guided by DJXPerf, we study and optimize a number of Java and Scala programs, including well-known benchmarks and real-world applications, and demonstrate significant speedups.

CCS Concepts

• **Software and its engineering** → **Compilers**; **General programming languages**.

Keywords

compiler techniques and optimizations, performance, dynamic analysis, managed languages and runtimes

ACM Reference Format:
Bolun Li, Pengfei Su, Milind Chabbi, Shuyin Jiao, and Xu Liu. 2018. DJXPerf: Identifying Memory Inefficiencies via Object-centric Profiling for Java. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/1122445.1122456

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Woodstock '18, June 03–05, 2018, Woodstock, NY
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

1 Introduction

Java is the "go-to" programming language choice for developing scalable enterprise cloud applications [1, 11, 12, 50, 64, 80, 81]. Performance is critical in such Java programs running on a distributed system comprised of thousands of hosts; on such systems, saving CPU time even by a few percentages offers both competitive advantages (lower latency) and cost savings. Tools abound in Java for "hotspot" performance analysis that can bubble-up code contexts where the execution spends the most time and are also used by developers for tuning performance. However, once such "low hanging targets" are optimized, identifying further optimization opportunities is not easy. Java, like other managed languages, employs various abstractions, runtime support, just-in-time compilation, and garbage collection, which hide important execution details from the plain source code.

In modern computer systems, compute is "free" but memory accesses cost dearly. Long-latency memory accesses are a major cause of execution stalls in modern general-purpose CPUs. CPU's memory hierarchy (different levels of caches) offers a means to reduce average memory access latency by staging data into caches and repeatedly accessing before evicting. Access patterns that reuse previously fetched data are said to exhibit good data locality. There are traditionally two types of data locality: *spatial* and *temporal*. An access pattern exhibits spatial locality when it accesses a memory location and then accesses nearby locations soon afterward. An access pattern exhibits temporal locality when it accesses the same memory multiple times. Programs that do not exploit these features are said to lack spatial or temporal locality.

Maintaining the locality of references in a CPU's memory hierarchy is well-known and mastered to achieve high performance in natively compiled code such as C and C++. Besides the traditional locality problems, garbage collected languages such as Java expose another unique locality issue — memory bloat [95]. Memory bloat occurs by allocating (and initializing) many objects whose lifetimes do not overlap. For example, allocating objects in a loop where the lifetime of the object is only the scope of the loop body. Since the garbage collection happens sometime later in the future, the memory consumption spikes, which results in a higher memory footprint and suboptimal cache utilization. Memory bloat can be seen as a case of lack of both spatial and temporal locality because

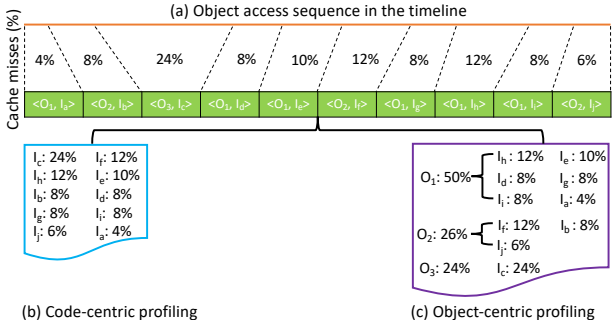

Figure 1: Code-centric vs. object-centric profiling. O denotes an object and I denotes a memory access instruction; tuple $\langle O_m, I_n \rangle$ denotes that instruction I_n accesses object O_m.

accessing a large number of independent objects results in accessing disparate cache lines and little or no reusing of a previously accessed cache line(s).

Exploring data locality in native languages has been under investigation for decades. There exist a number of tools [13–15, 58, 59, 61–63, 101] to measure the data locality using various metrics. Software metrics such as reuse distances [28, 79, 86, 100, 101], memory footprints [10, 54], and cache miss ratio curves [15, 19], which are derived from memory access traces, quantify the locality independent of architectures. In contrast, hardware metrics, such as cache/TLB misses, collected with hardware performance monitoring units (PMU) during program execution, quantify data locality on a given architecture. Attributing PMU metrics to source code is a straightforward way to demonstrate code regions that incur high access latencies; we call it as *code-centric profiling*.

In object-oriented systems, delivering profiles centered around objects is highly desirable. An object may have accesses to it scattered in many places, and each code location may contribute a small fraction to the overall data-access problem. Bridging the memory-hierarchy latency metrics with runtime data objects requires more involved measurement and attribution techniques, which is the focus of our work; we call it as *object-centric profiling*. Note that *object-centric profiling* not only shows the objects subject to high aggregate access latencies but also pinpoints code locations ordered by their contribution to the overall latency to the object under question.

Figure 1 illustrates the difference between code-centric profiling and object-centric profiling. The code-centric profiling associates the cache miss metric with memory accesses, showing that access I_c accounts for most cache misses during program execution. In contrast, the object-centric profiling aggregates the cache miss metric from different accesses that touch the same object to present a unified view. Guided by the object-centric profiling, we find that object O_1 accounts for most cache misses. However, its accesses are scattered across multiple instructions, and each individual access is less significant than the access to object O_3. Thus, instead of checking individual accesses, one can apply various optimization to the allocation of the O_1 object or its data layout.

Collecting object-centric profiles for Java has unique challenges posed by managed runtimes. First, just-in-time compilation and interpretation used in JVM disjoin the program source code and its execution behaviors. Second, automatic runtime memory management (i.e., garbage collection) further impedes understanding memory performance of Java and other languages based on JVM, such as Scala [8] and Clojure [48].

Since developers in native languages have explicit knowledge and understanding of objects and their lifetimes, they pay more attention to objects and their locality; Java developers, on the other hand, lack the precise knowledge of object lifetimes and their influence on locality. Surprisingly, Java lacks any object-centric performance attribution tool, which can pinpoint objects subject to serious latency problems. Tools such as Linux perf [57] and Intel VTune [52] exploit hardware performance monitoring and attribute cache miss metrics to code, such as Java methods, loops, or source locations, but they do not attribute metrics to Java objects. As shown in Figure 1, these tools only attribute metrics to problematic code lines; without object-level information, they cannot tell which objects are problematic and deserve optimization. A few tools such as [69, 94, 97] instrument Java byte code to identify problematic objects. However, these tools suffer from high overhead and lack real execution performance metrics available from the hardware. Often times, the optimization lacking the quantification from the underlying hardware can yield trivial or negative speedups [59].

Most of today's shared memory multiprocessor systems support non-uniform memory access (NUMA), where the loss of data locality not only pervasively exists within a single CPU processor (aka a NUMA socket), but also across CPU processors. NUMA characteristics introduce the explicit NUMA property of the shared memory systems, where differences in local and remote memory access latencies can be up to two orders of magnitude. Shared memory applications with transparent data distributions across all nodes often incur high overheads due to excessive remote memory accesses. Minimizing the number of remote accesses is crucial for high performance, and this, in turn, usually requires a suitable, application-specific data distribution. In general, choosing an appropriate data distribution remains a challenge. We can attribute metrics to individual objects with the object-centric idea, then identifying the objects that suffer from severe remote accesses.

In this paper, we describe DJXPerf, a lightweight object-centric profiler for Java programs. DJXPerf complements existing Java profilers by collecting memory-related performance metrics from hardware PMUs and attributing them with Java objects. DJXPerf provides unique insights into Java's memory locality issues. In the rest of this section, we show two motivating examples, the contribution of this paper and the paper organization.

1.1 Motivating Examples

In this section, we motivate the importance of combining object-level information and PMU metrics for locality optimization. Listing 1 and 2 show two problematic code snippets suffering from memory bloat, which are respectively from `batik` and `lusearch`, both from `Dacapo-9.12` [65]. We run them with the default large inputs using 48 threads.

In Listing 1, the object allocation site at line 5 creates an array of `float` objects in the method `makeRoom`, which is part of the class `ExtendedGeneralPath`. This allocation site is repeatedly invoked 2478 times, resulting in memory bloat. For optimization, one can

```
1 private void makeRoom(int numValues) {
2   ...
3   if ( newSize > values.length) {
4     ...
5 ▶   float [] nvals = new float[nlen];
6     System.arraycopy(values, 0, nvals, 0, numVals);
7     ... }
8   ... }
```

Listing 1: The code snippet from `Dacapo 9.12 batik`. Optimizing memory bloat by moving the object allocation site at line 5 outside of the loop yields a nontrivial speedup ($(1.15 \pm 0.03)\times$) to the entire program.

```
1 public TopDocs search(Weight weight, Filter filter, final int
      nDocs) {
2   ...
3 ▶TopDocCollector collector = new TopDocCollector(nDocs);
4   search(weight, filter, collector);
5   ... }
```

Listing 2: The code snippet from `Dacapo 9.12 lusearch`. Optimizing memory bloat by moving the object allocation site at line 3 outside of the loop does not bring any speedup to the entire program.

move the array allocation outside of the loop that encloses the method `makeRoom` and replace it with a static object array, aka the singleton pattern. This optimization addresses the memory bloat and yields a $(1.15 \pm 0.03)\times$ speedup.

In Listing 2, the memory bloat occurs at line 3, which repeatedly allocates the object `collector` 15179 times. This object is passed as an input parameter to the method `search` and used in many places in that method. One can also apply the singleton pattern by hoisting the allocation site outside of the loop enclosing the method, and declaring it as a static object. While this optimization addresses the memory bloat, it does not bring any noticeable speedup.

The study of these two example code snippets reveals that basing the optimization only on allocation frequency (or the metrics derived from the allocation frequency [95]) does not necessarily yield performance benefits. This motivates the need for the extra locality metrics associated with the object allocation site, which we call as *object-centric profiling*. To be concrete, DJXPerf measures L1 cache misses[1] with PMU on individual memory access instances and aggregates the measurement of memory accesses to the object's allocation site. For example, DJXPerf reports that accessing the `nvals` array object shown in Listing 1 accounts for 21% of total cache misses, while accessing the `collector` objects in Listing 2 accounts for less than 1% of total cache misses only, which explains the different speedups obtained from the locality optimization. **The strength of DJXPerf is its ability to aggregate myriad accesses to the same object, scattered all over the program, back to the same object.** We emphasize that object-centric analysis does not do away with code-centric aspect; underneath each object allocation site C, DJXPerf provides the ability to disaggregate the code contexts contributing towards C's overall locality loss. Thus, object-centric analysis with the locality metrics associated with the object allocation sites is desired to determine whether locality optimization can yield significant speedups.

[1] we can measure myriad other events, for example, L3 cache misses, TLB misses, etc.

1.2 Paper Contributions

In this paper, we propose DJXPerf, an object-centric profiler that guides data locality optimization in Java programs. DJXPerf makes the following contributions.

- DJXPerf develops a novel object-centric profiling technique. It provides rich information to guide locality optimization in Java programs, which yields nontrivial speedups.
- DJXPerf combines hardware performance monitoring units with minimal Java byte code instrumentation, which typically incurs 8% runtime and 5% memory overhead.
- DJXPerf applies to unmodified Java (and languages based on JVM, e.g., Scala) applications, the off-the-shelf Java virtual machine and Linux operating system, running on commodity CPU processors, which can be directly deployed in the production environment.
- DJXPerf provides intuitive optimization guidance for developers. We evaluate DJXPerf with popular Java benchmarks (Dacapo [65], NPB [29], Grande [18], SPECjvm2008 [24], and the most recent Renaissance [76]) and more than 20 real-world applications. Guided by DJXPerf, we are able to obtain significant speedups by improving data locality in various Java programs.

1.3 Paper Organization

The paper is organized as follows. Section 2 describes the related work and distinguishes DJXPerf. Section 3 introduces some background knowledge. Section 4 depicts DJXPerf's methodology. Section 5 describes the implementation details of DJXPerf. Section 6 evaluates DJXPerf's accuracy and overhead. Section 7 shows some case studies. Finally, Section 8 presents some conclusions.

2 Related Work

There are numerous Java performance tools assisting developers in understanding their program behaviors, such as profiling for execution hotspots [7, 22, 35, 44, 57, 67, 75] in CPU cycles or heap usage, and pinpointing redundant computation [26, 27, 73, 87]. These tools target orthogonal problems to DJXPerf, which particularly focuses on data locality. Furthermore, there are many tools [17, 55, 58, 60, 61, 63, 66, 82, 83] pinpointing poor locality issues in native code via OS timers or PMU-based sampling techniques, or code instrumentation. Unlike DJXPerf, these tools do not work for Java applications. In this section, we only review Java profiling techniques that are related to data locality and PMUs.

2.1 Data Locality Analysis in Java

Most existing Java profilers focus on memory bloat, which is one of the locality issues (aka high memory footprint) in Java. Mitchell *et al.* [70] design a mechanism to track data structures that suffer from excessive amounts of memory. Their follow-up work [68, 69] summarizes memory usage to uncover the costs of design decisions, which provides more intuitive guidance for code improvement.

Xu *et al.* [96] develop copy profiling that detects data copies and suggests removal of allocation and propagation of useless objects. Their follow-up work [98] presents a technique that combines static and dynamic analyses to identify underutilized and overpopulated

containers. They also develop a dynamic technique [95] to highlight data structures that can be reused to avoid frequent object allocations.

Nguyen and Xu [53] develop Cachetor, a value profiler for Java. Cachetor identifies operations that keep generating identical data values and suggests memoizing the invariant values for the future usage. Yan *et al.* [99] track object propagation by monitoring object allocation, copy, and reference operations; by constructing a propagation graph, one can identify never-used or rarely-used object allocations. Dufour *et al.* [32, 33] analyze the use and shape of temporary data structures based on a blended escape analysis to find excessive memory usage. JOLT [85] uses dynamic analysis to identify object churn and performs function inlining.

There are few studies in measuring traditional data locality in Java programs. Gu *et al.* develop ViRDA [46], which is perhaps the most related to DJXPerf. They collect memory access trace and compute reuse distance to quantify the temporal and spatial data locality in Java programs.

While these existing efforts can effectively identify some locality issues in Java, they mostly suffer from two limitations. First, they employ fine-grained byte code instrumentation, which incurs high overhead. For example, the work [53, 96] can incur 30-200× runtime overhead. Second, they do not collect performance data from the real execution provided by the hardware; instead, they employ cache simulators. Without the information of the underlying hardware, optimization efforts may be misguided as shown in Section 1.1.

DJXPerf addresses these limitations by introducing object-centric profiling technique, which is based on lightweight data collection from hardware PMUs. DJXPerf is not a replacement for existing tools; it provides complementary information to save non-fruitful optimization efforts.

2.2 Java Profilers Based on PMUs

Sweeney *et al.* [36] develop a system to help interpret results obtained from PMUs when analyzing a Java application's behavior. Cuthbertson *et al.* [25] map the instruction IP address based hardware event information to the JIT server components. Goldshtein [45] monitors CPU bottlenecks, system I/O load and GC with perf in production JVM environments. Hauswirth *et al.* [47] introduce vertical profiling, adding software performance monitors (SPMs) to observe the behavior in the layers (VM, OS, and libraries) above the hardware. Georges *et al.* [41] measure the execution time for each method invocation with PMUs and study method-level phase behaviors in Java applications. Lau *et al.* [56] guide inline decisions in a dynamic compiler with the direct measure of CPU cycles. Eizenberg *et al.* [34] utilize PMUs to identify false sharing in Java programs.

Unlike these existing approaches, DJXPerf leverages PMUs to identify data locality in Java programs. Its usage of lightweight PMU measurement for object-centric analysis is unique among all existing Java profilers.

2.3 NUMA Analysis in Java

Gidra *et al.* [42] study the scalability of throughput-oriented GCs and propose to map pages and balance GC threads across NUMA nodes. Gidra *et al.* [43] propose local mode for NUMA machines to forbid GC threads to steal references from remote NUMA nodes so as to avoid costly cross-node memory accesses. Maria *et al.* [20] show how to optimize three main GCs in OpenJDK, i.e., ParallelOld, ConcurrentMarkSweep, and G1 in multicore NUMA machines. Tikir *et al.* [71] propose NUMA-aware heap configurations for Java server applications to improve the memory performance during the GC phase. Raghavendra *et al.* [78] propose a dynamic compiler scheme for splitting the Java code buffer on a CC-NUMA machine.

While these works can identify some memory access latency issues, they cannot pinpoint the object-level remote memory access issue. DJXPerf is able to identify NUMA locality and does not depend on the GC.

3 Background

DJXPerf leverages facilities available in commodity Java virtual machines (JVM) and CPU processors, which we introduce in this section.

ASM Framework ASM [16] is a Java byte code manipulation and analysis framework. ASM can modify existing classes or dynamically generate classes, supporting custom complex transformations and code analysis tools. ASM focuses on performance, with an emphasis on the low overhead, which makes it suitable for dynamic analysis. ASM can instrument object allocation (e.g., `new`) and capture the object information, such as allocation size and context.

Java Virtual Machine Tool Interface (JVMTI) JVMTI [21] is a native programming interface of the JVM, which supports developing debuggers/profilers (aka JVMTI agents) in C/C++ based native languages to inspect JVM internals. JVMTI provides a number of event callbacks to capture JVM start and end, thread creation and destruction, method loading and unloading, garbage collection epochs, to name a few. User-defined functions are subscribed to these callbacks and invoked when the associated events happen. Also, JVMTI maintains a variety of JVM internal states, such as the map from the machine code of each JITted method to byte code and source code, and the call path for any given point during the execution. Tools based on JVMTI can query these states at any time. JVMTI is available in off-the-shelf Oracle HotSpot JVM [23].

Hardware Performance Monitoring Unit (PMU) Modern CPUs expose programmable registers (aka PMU) that count various hardware events such as memory loads, stores, and CPU cycles. These registers can be configured in sampling mode: when a threshold number of hardware events elapse, PMUs trigger an overflow interrupt. A profiler can capture the interrupt as taking a sample and attribute the metrics collected along with the sample to the execution context. PMUs are per CPU core and virtualized by the operating system (OS) for each thread.

Intel offers Precise Event-Based Sampling (PEBS) [51] in Sandy-Bridge and following generations. PEBS provides the effective address (EA) at the time of the sample when the sample is for a memory load or store instruction. PEBS also reports memory-related metrics of the sampled loads/stores such as cache misses, TLB misses, memory access latency. This facility is often referred to as

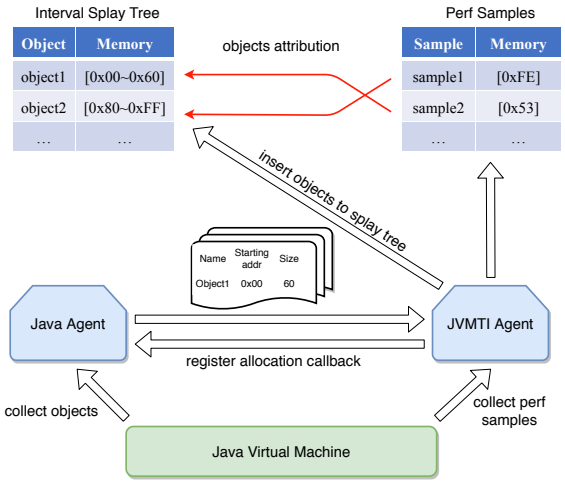

Figure 2: Overview of DJXPerf's object-centric analysis.

address sampling — a building block of DJXPerf. AMD Instruction-Based Sampling [31] and PowerPC Marked Events [88] offer similar capabilities.

Linux perf_event. Linux offers a standard interface to program and sample PMUs via the `perf_event_open` system call [92] as well as the associated `ioctl` calls. The Linux kernel can deliver a signal to the specific thread whose PMU event overflows. The user code can extract PMU performance data and execution contexts at the signal handler.

4 Methodology

Figure 2 overviews DJXPerf's object-centric profiling. DJXPerf includes two agents: a Java agent and a JVMTI agent. The Java agent adds lightweight byte code instrumentation to capture object allocation information during execution, such as the allocation context and address range of objects. The JVMTI agent subscribes to Java thread creation callbacks to enable PMU to sample memory accesses. When PMU interrupts a thread with a sampled address, DJXPerf associates the address seen in the sample with the Java object enclosing that address, as shown in Figure 2. In the rest of this section, we elaborate on each agent and discuss their interactions for the object-centric analysis.

4.1 Java and JVMTI Agents

Capturing Object Addresses via A Java Agent DJXPerf leverages a Java agent to capture object allocation. The Java agent is based on the ASM framework [16]. The Java agent scans Java byte code and instruments four object allocation routines — `new`, `newarray`, `anewarray`, and `multianewarray`. The Java agent inserts pre- and post-allocation hooks to intercept each object allocation and returns the object information (e.g., object pointer, type, and size) via user-defined callbacks. Upon each allocation callback, we follow an existing technique [4] to obtain the memory range allocated for each Java object.

Generating Memory Access Samples via JVMTI Agent DJXPerf leverages a JVMTI agent to sample and collect memory accesses. With the help of JVMTI, DJXPerf can intercept Java thread start, where DJXPerf configures PMUs to sample precise events for cache misses (e.g., `MEM_LOAD_UOPS_RETIRED:L1_MISS`), TLB misses (e.g., `DTLB_LOAD_MISSES`), or memory access latency (e.g., `MEM_TRANS_RETIRED:LOAD_LATENCY`). DJXPerf also installs a signal handler to process PMU samples. On thread termination, DJXPerf stops PMUs and produces a profile for each thread. Besides controlling PMUs, DJXPerf also utilizes the JVMTI agent to capture the calling contexts for both PMU samples and object allocations, which is described in Section 4.4.

4.2 Object-centric Attribution

Identifying Objects Java objects are allocated on the heap. How to represent an object to a developer is a challenging question. We adopt a simple and perhaps most intuitive approach that developers can identify with the allocation call path leading to the object allocation. We represent a call path where an object O was allocated with $\mathcal{P}(O)$. An application may create multiple objects via a single allocation site, for example, in a loop. In our approach, all such objects will be represented by a single call path and become indistinguishable from one another; we accept this trade-off since objects allocated at the same call path are likely to exhibit similar behavior. DJXPerf's Java agent captures each allocation instance and invokes the JVMTI agent to obtain the allocation call path. As these allocation instances share the same call path, DJXPerf associates the PMU metrics for any of those objects with the same call path.

Attributing PMU Samples to Objects DJXPerf utilizes an efficient interval splay tree [30] to maintain the memory ranges allocated for all the monitored objects. On each PMU sample, DJXPerf uses the effective address M presented by the PMU to look up into the splay tree. The lookup for M returns the object O whose memory range encloses the sampled address. DJXPerf then attributes any associated PMU metric related with the sample to $\mathcal{P}(O)$ — the object's allocation call path.

4.3 Object NUMA locality detection

On each PMU sample, DJXPerf uses the effective address M to identify a memory page in libnuma library function `numa_move_pages`, which simply uses the `move_pages` system call. Not only can `move_pages` move a specified page to a specified NUMA node, but also return the NUMA node where the page is currently residing. Then DJXPerf is able to identify which NUMA node ($Node_1$) the current object is allocated. Still on each PMU sample, DJXPerf uses `PERF_SAMPLE_CPU` identifier, which contains the CPU number, to show which NUMA node ($Node_2$) the current object is accessing. If $Node_1$ and $Node_2$ are distinct nodes, then DJXPerf reports a remote memory access for this object.

4.4 Calling Context Determination

We associate an object allocation with the full calling context (aka call path) leading to its allocation. The full call path helps distinguish allocations by the same routine called from different code contexts. The alternative, a flat profile, would be unable to distinguish, for example, an allocation in a common library routine called from two different user code locations.

Oracle Hotspot JVM supports two JVMTI APIs to obtain calling contexts: `GetStackTrace` and `AsyncGetCallTrace`. `GetStackTrace` requires the program to reach a safe point to collect calling context, which produces biased results [49, 72]. Instead, DJXPerf employs `AsyncGetCallTrace` to obtain calling contexts at anytime [74]. `AsyncGetCallTrace` accepts `u_context` obtained at PMU interrupts or object allocations, and returns the byte code index (BCI) and method ID for each frame in the calling context. Usually, a single Java source line may translate to several byte code instructions, and the BCI can tell which byte code instruction was executed. Since an individual method may be JITted multiple times, the method ID helps distinguish different JITted instances of the same method. With the method ID, DJXPerf obtains the corresponding class name and method name by querying JVM. To obtain the line number, DJXPerf maintains a "`BCI→line number`" map for each method instance via JVMTI API `GetLineNumberTable` and queries the line number on demand.

4.5 Interfering with Garbage Collection

The garbage collector (GC) complicates object-centric attribution because GC implicitly reclaims memory of unused objects and moves objects for compact memory layouts. The trigger of the GC thread is determined by JVM, which is transparent to Java applications. Ignoring GC, DJXPerf may yield incorrect object attribution. We assume GC reclaims or moves an object O_1, whose memory can be reused by another allocation, say object O_2. In this case, a tool may incorrectly attribute PMU samples that touch O_2 to O_1. Moreover, if O_1 is moved to a new memory location, any subsequent PMU samples of the new address will be unavailable for us to map via the original mapping maintained in the splay tree.

Handling GC is necessary in DJXPerf. Unfortunately, JVM exposes limited information about GC via JVMTI: JVMTI only provides hooks to register callbacks on GC start and end, with no insight about the individual object behavior (i.e., reclamation and movement). We offer a solution to handle all kinds of GC in the off-the-shelf JVM.

Solution for Object Movement by GC Our solution is based on an important observation from the source code of OpenJDK: *GC moves objects using the* `memmove` *function.* Thus, DJXPerf overloads `memmove` to obtain the source and destination of every moved object and update the memory ranges associated with this object in the splay tree described in Section 4.2. However, updating the splay tree upon each `memmove` invocation is costly. Instead, DJXPerf creates a relocation map for each thread to record the moved objects (e.g., source as the key, destination memory addresses, and size as the value). DJXPerf updates the objects in the map in a batch at the end of each GC invocation.

To capture every GC invocation, DJXPerf utilizes JVM management interface, i.e., `MXBean`. DJXPerf, in the Java agent, registers GC invocation callbacks via `GARBAGE_COLLECTION_NOTIFICATION` event. Upon each GC completion, an `MXBean` instance (i.e., `GarbageCollectorMXBean`) emits a callback; DJXPerf captures this callback and updates all the newly moved objects in the relocation map to the splay tree. The relocation map is reset after the update.

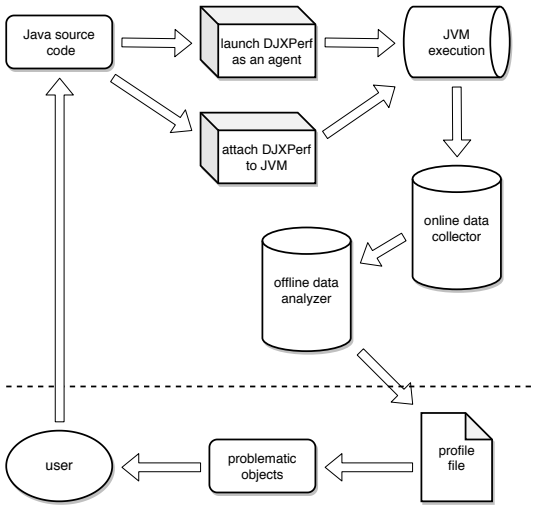

Figure 3: The workflow of DJXPerf.

It is worth noting that DJXPerf may not always capture all the object allocation because its attach mode may omit some allocations (see Section 5.1). If this is the case, DJXPerf directly inserts the new memory intervals for the moved objects.

Solution for Object Reclamation by GC DJXPerf handles object reclamation by overloading the `finalize` method. GC always calls the `finalize` method before reclaiming memory for any object, which cleans up resources allocated to the object. DJXPerf intercepts the `finalize` method, obtains the memory interval reclaimed, and removes it from the splay tree.

5 Implementation

DJXPerf is a user-space tool with no need for any privileged system permissions. DJXPerf requires no modification to hardware, OS, JVM, or monitored applications, which makes DJXPerf applicable to the production environment. Figure 3 shows the workflow of DJXPerf, which consists of an online data collector and an offline data analyzer. There are two ways to enable the collector. If we need to profile Java source code as soon as the JVM starts up, we can launch DJXPerf as an agent by passing JVM options. If the JVM is already started, we can attach DJXPerf to this running JVM. The collector gleans the measurement via the Java and JVMTI agents and generates a profile file per thread. The analyzer then aggregates the files from different threads, sorts the metrics, and highlights the problematic objects for investigation.

The implementation challenges include maintaining a low measurement overhead and scaling the analysis to many threads. In the rest of this section, we discuss how DJXPerf addresses these challenges.

5.1 Online Collector

DJXPerf supports two modes to monitor a Java program. On the one hand, DJXPerf can monitor the end-to-end execution of a program by launching the tool together with the program. On the other hand, DJXPerf can attach and detach to any running Java

program to collect the object-centric profile for a while. This is particularly useful to monitor long-running programs such as web servers or microservices.

DJXPerf accepts any memory-related PMU precise events. In our implementation, DJXPerf presets the event as L1 cache misses (MEM_LOAD_UOPS_RETIRED:L1_MISS). We empirically choose a sampling period to ensure DJXPerf is able to collect 20-200 samples per second per thread, which has a good trade-off between runtime overhead and statistical accuracy [90].

To minimize the thread synchronization, DJXPerf has each thread collect PMU samples independently and maintains the calling contexts of PMU samples in a compact calling context tree (CCT) [9], which merges all the common prefixes of given calling contexts. The only shared data structure between threads is the splay tree for objects because an object allocated by a thread can be accessed by other threads. DJXPerf uses a spin lock to ensure the integrity of the splay tree across threads.

Another source of overhead is monitoring objects, which depends on a given Java application. DJXPerf can either monitor every object allocation, or filter out objects whose sizes are smaller than a configurable value $\mathcal{S}$ from monitoring to trade off the overhead. DJXPerf by default sets $\mathcal{S}$ as 1KB to pinpoint large objects that are more vulnerable to locality issues. We evaluate the impact of $\mathcal{S}$ in Section 6.

5.2 Offline Analyzer

To generate compact profiles, which is important for scalability, DJXPerf's offline analyzer merges profiles from different threads. Object-centric profiles, organized as a CCT per thread, are amenable to coalescing. The analyzer merges CCTs in a top-down way — from call path root to leaf. Call paths for individual objects and their memory accesses can be merged recursively across threads. If object allocation call paths are the same, they have coalesced even if they are from different threads. All memory accesses with their call paths to the same objects are merged as well. Metrics are also summed up when merging CCT nodes. Typically, DJXPerf's analyzer requires less than one minute to merge all profiles in our experiments. DJXPerf provides a Python-based GUI to visualize the profiles for intuitive analysis.

5.3 Discussions

DJXPerf is a sampling-based dynamic profiling tool. The sampling strategy only identifies statistically significant performance issues and ignores some insignificant ones. Also, the sampling rate should be appropriately chosen. A high sampling rate brings high overhead, and a low sampling rate obtains insufficient samples, which can result in over- or under-estimation. Nevertheless, there is sufficient evidence to show that random sampling via PMUs is superior to biased sampling [72]. Like other dynamic profilers DJXPerf's optimization guidance is input dependent. We recommend to use typical program inputs for representative profiles. Additionally, we ensure the optimization is correct across different inputs. Finally, DJXPerf pinpoints objects potentially for optimization, but users need to determine and apply the optimization.

6 Evaluation

We evaluate DJXPerf on a 24-core Intel Xeon E5-2650 v4 (Broadwell) CPU clocked at with 2.2GHz running Linux 4.18. The memory hierarchy consists of a private 32KB L1 cache, a private 256KB L2 cache, a shared 30MB L3 cache, and 256GB main memory. DJXPerf works for any Oracle JDK version with the JVMTI support; in this experiment, we run all applications in Oracle Hotspot JVM with JDK 1.8.0_161.

Accuracy Analysis We show that DJXPerf can accurately pinpoint Java objects with the poor locality. We evaluate DJXPerf's ability to find performance bugs on five benchmarks that have locality issues reported by prior work [95]. These five benchmarks are luindex, bloat, lusearch, and xalan, all from Dacapo 2006 [65], as well as SPECJbb2000 [2]. DJXPerf successfully identified all the locality issues as reported by existing tools [95]. We elaborate the analysis of these benchmarks in the appendix since they have already been discussed elsewhere.

Overhead Analysis The runtime overhead (memory overhead) is the ratio of the runtime (peak memory usage) under monitoring with DJXPerf to the runtime (peak memory usage) of the corresponding native execution. To quantify the overhead of DJXPerf in both runtime and memory, we apply DJXPerf to a number of well-known Java benchmark suites, such as the most recent JVM parallel benchmarks suite Renaissance [76], Dacapo 9.12 [6], and SPECjvm2008 [24]. We run all benchmarks with four threads if applicable. We run every benchmark 30 times and compute the average and error bar. Figure 4 shows the overhead when DJXPerf is enabled at a sampling period of 5M for Renaissance benchmark suite [76], Dacapo 9.12 [6], and SPECjvm2008 [24]. Some Renaissance and Dacapo benchmarks have higher time overhead (larger than 30%) because they usually invoke too many allocation site callbacks (e.g., more than 400 million times for mnemonics, parmnemonics, scrabble, akka-uct, db-shootout, dec-tree, and neo4j-analytics). From Figure 4 we can see that DJXPerf typically incurs 8% runtime and 5% memory overhead.

Further Discussions To trade off the overhead, DJXPerf allows users to set up a configurable value $\mathcal{S}$ to capture memory allocations with the size greater than $\mathcal{S}$. We also test an extreme case: setting $\mathcal{S}$ to 0, which means DJXPerf captures the allocation for each object. With the evaluation on the Renaissance benchmark suite [76], DJXPerf incurs a runtime overhead ranging from 1.8× to 3.6× to monitor every object allocation. With further investigation, we seldom find opportunities for optimizing small-sized objects. Thus, we believe setting $\mathcal{S}$ to 1KB is a good trade-off between the overhead and obtained insights. One can expect DJXPerf to be used in attach and detach mode on production services, where developers collect profiles from multiple instances of their services; hence, the overhead, if any, is introduced only for the short duration of measurement and the samples from multiple instances will offer a good coverage.

7 Case Studies

DJXPerf's low overhead allows us to collect object-centric profiles from a variety of Java and Scala applications, such as the Renaissance benchmark suite [76], ObjectLayout [91], Findbugs [77],

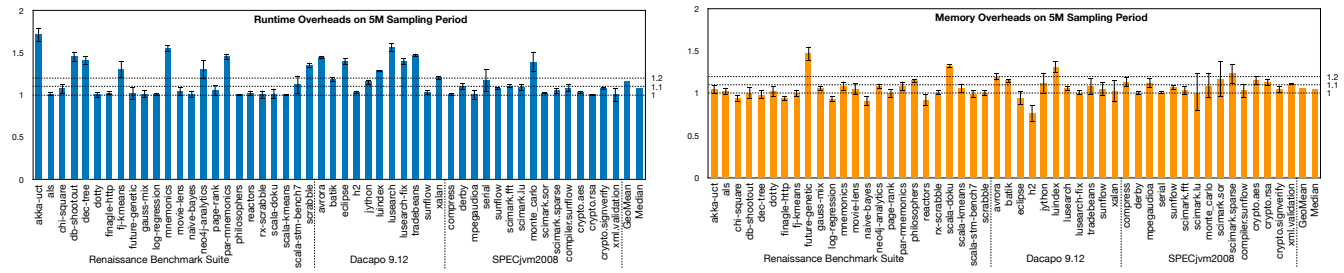

(a) Runtime overheads.

(b) Memory overheads.

Figure 4: DJXPerf's runtime and memory overheads in the unit of times (×) on various benchmarks.

Real-world Applications	Inefficiency		Optimization	
	Problematic Code	Type	Patch	WS (×)
ObjectLayout 1.0.5 [88]	AbstractStructuredArrayBase.java (292) SAHashMapBench.java (85, 120,135)	Excessive memory usage in nested loops	Move problematic allocation sites out of loop and reset them upon request	1.45±0.07
FindBugs 3.0.1 [74]	LoadOfKnownNullValue.java (120) ClassParserUsingASM.java (643)			1.11±0.01
Ranklib 2.3 [81]	CoorAscent.java.java (218) MergeSorter.java (137, 138)			1.25±0.05
Cache2k 1.2.0 [90]	Hash2.java (313)			1.09±0.02
Apache SAMOA 0.5.0 [36]	ArffLoader.java (165)			1.17±0.04
Apache Commons Collections 4.2 [37]	AbstractHashedMap.java (151)			1.08±0.01
Renaissance 0.10: scala-stm-bench7 [73]	AccessHistory.scala (619)		Enlarge the initial allocation size to avoid frequent allocation	1.12±0.04
SPECjvm2008: Scimark.fft.large [3]	FFT.java (171, 172, 174, 175)	Problematic data with high L1 cache misses	Loop interchange	2.37±0.07
JGFMonteCarloBench 2.0 [18]	RatePath.java (205)		Loop tiling for high L1 cache miss data	1.07±0.03
JGFMolDynBench 2.0 [18]	md.java (348, 349, 350)			1.24±0.13
Eclipse Collection [39]	Interval.java (758)	NUMA remote access	Allocate problematic allocation sites interleaved on all NUMA nodes	1.13±0.04
NPB SP [28]	SPBase.java (155)			1.1±0.03
Apache Druid [38]	WrappedImmutableBitSetBitmap.java (37)			1.75±0.05

WS: Whole-program speedup after optimization

Table 1: Overview of performance optimization guided by DJXPerf.

Ranklib [84], cache2k [93], Apache SAMOA [37], Apache Commons Collections [38], Java Grande 2.0 [18], to name a few. We run these applications with the default inputs released with these applications or the inputs that we can find to our best knowledge. To fully utilize CPU resources, we run each parallel application to saturate all CPU cores if not specified. DJXPerf can pinpoint many previously not reported data locality issues and guide optimization choices. To guarantee correctness, we ensure the optimized code passes the application validation tests. To avoid system noises, we run each application 30 times and use a 95% confidence interval for the geometric mean speedup to report the performance improvement, according to the prior approach [89]. Table 1 overviews the performance improvements on several real-world Java applications guided by DJXPerf. In the remaining section, we elaborate on the analysis and optimization of several applications under the guidance of DJXPerf in Section 7.1- 7.6. Due to the page limit, we omit the description of other applications in Table 1. Section 7.7 shows some studies on optimizing insignificant objects, which illustrate the unique usefulness of DJXPerf over other tools.

We also collected code-centric profiles using the Linux perf utility for each case study to compare with DJXPerf's profiles. In several cases, we found it arduous to tie data accesses segregated over many code location back to the object allocation site in code-centric profiles. DJXPerf eased developer's task by showing top data objects subject to locality problems along with an ordered, hierarchical view of code locations contributing to the locality problem.

To optimize the NUMA locality issue, we developed a Java library to access the libnuma library interfaces. As the libnuma library can be used only in native languages, we leverage Java Native Interface (JNI) to enable our Java library to access the native NUMA API. After DJXPerf detected the problematic object that suffers from remote memory access, we modify the Java application allocation code for this object by calling the libnuma `numa_alloc_interleaved` API. Then, we can allocate this problematic object interleaved on all NUMA nodes to reduce the remote accesses.

7.1 ObjectLayout 1.0.5

ObjectLayout [91] is a layout-optimized Java data structure package, which provides efficient data structure implementation. We run ObjectLayout with the SAHashMap input released with the package. DJXPerf reports four problematic objects, which account for 84% of cache misses in the entire program. Figure 5 shows the snapshot of DJXPerf's GUI for intuitive analysis. The top pane of the GUI

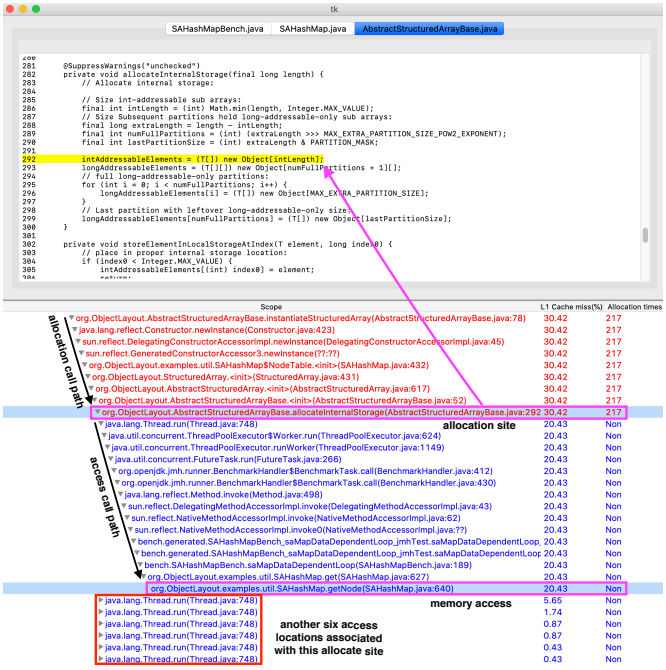

Figure 5: The top-down object-centric GUI view of ObjectLayout shows a problematic object's allocation site in source code, and its allocation call path with its all access call paths.

shows the Java source code; the bottom left pane shows the object allocation (in red) and accesses (in blue) in their full call paths; and the bottom right pane shows the metrics (e.g., L1 cache misses, object allocation times in this example).

The GUI in Figure 5 shows one problematic object, which is at line 292 of method `allocateInternalStorage` in class `AbstractStructuredArrayBase`. The `allocateInternalStorage` method is repeatedly invoked in a loop when a new instance is created by the `newInstance` method. The bottom left pane of the GUI shows one problematic object allocation in the full call path, which is allocated 217 times in a loop. There are multiple sampled accesses associated with this allocation site and they account for 30.4% L1 cache misses of the entire program. To save the space, we only show one access in its call path in blue (method `getNode` in class `SAHashMap`) that accounts for most of the cache misses. For the other accesses, we only show their call paths rooted at `java.lang.Thread.run`; with no object-centric profiling, these accesses are separately presented with no aggregate view.

We further investigate the source code and find that the problematic object is array `intAddressableElements`. The life cycles of different instances do not overlap, which means using singleton pattern of this object (i.e., allocating a single object instance and reusing it without creating more instances) is safe and avoids memory bloat. To apply the singleton pattern, we hoist `intAddressableElements` allocation out of `allocateInternalStorage` method, which is thread safe. We optimize three other problematic objects with similar methods. Our optimization reduces total cache misses by 76% and yield a $(1.45 \pm 0.07)\times$ increase in throughput.

```
634 public void setAppClassList(List<ClassDescriptor>
        appClassCollection) {
635   for (ClassDescriptor appClass : allClassDescriptors) {
636     ...
637     ▶XClass xclass = currentXFactory().getXClass(appClass);
638     //getXClass method will call to parse method below
639   }}
640 public void parse(ClassInfo.Builder builder) {
641   ...
642   ▶char[] buf = new char[1024];
643   ...}
```

Listing 3: The problematic source code highlighted by DJXPerf in FindBugs.

```
111 private void analyzeApplication() throws InterruptedException {
112   for (Detector2 detector : detectorList) {
113     ...
114     ▶detector.visitClass(classDescriptor);
115     //visitClass method will call to analyzeMethod method below
116   }}
117 private void analyzeMethod(ClassContext classContext, Method
        method) {
118   ...
119   ▶IdentityHashMap<InstructionHandle, Object> sometimesGood = new
          IdentityHashMap<InstructionHandle, Object>();
120   ... }
```

Listing 4: The source code highlighted by DJXPerf shows the problematic object `sometimesGood` (allocated at line 120) suffering from poor locality.

```
615 private def grow() {
616   _wCapacity *= 2
617   if (_wCapacity > _wDispatch.length) {
618     ...
619     ▶_wDispatch = new Array[Int](_wCapacity)
620   } ... }
```

Listing 5: DJXPerf pinpoints `_wDispatch` object suffering from high cache misses in scala-stm-bench7.

7.2 FindBugs 3.0.1

FindBugs is a program to find bugs in Java programs. It looks for instances of "bug pattern" — code instances that are likely to be errors [77]. We run FindBugs on Java chart library 1.0.19 version as input. DJXPerf reports two objects that account for 32% of cache misses in the entire program as shown in Listings 3 and 4. The two problematic objects `buf` and `IdentityHashMap` are both repeatedly allocated in loops with no overlap in lifecycles across different instances. We apply singleton pattern by hoisting the two object allocations out of the loops to avoid memory bloat. These optimizations reduce peak memory usage from 1.8GB to 0.9GB, yielding a $(1.11 \pm 0.01)\times$ speedup to the entire program.

7.3 Renaissance 0.10: scala-stm-bench7

scala-stm-bench7 is a Renaissance [76] benchmark, which runs stmbench7 code using ScalaSTM [5] for parallelism. It is written in Scala. We run scala-stm-bench7 using the default 60 repetitions. DJXPerf pinpoints a problematic object `_wDispatch` as shown in Listing 5. This object accounts for 25% of total cache misses. With further investigation, we find that the method `grow` is called frequently to adjust the `_wDispatch` array capacity and create a new `_wDispatch` array. Such frequent invocation of `grow` is because the initial size of `_wDispatch` array is only 8. For optimization, we increase the initial size of `_wDispatch` array to be 512, which reduces array creation and copy by 79%. This optimization yields a $(1.12 \pm 0.04)\times$ speedup to the entire program.

Applications	Problematic Code	Allocation times	L1 Cache Misses (%)	WS (×)
NPB 3.0 SP	SP.java (2086)	400	<1%	0.5%
Dacapo 2006 chart	Datasets.java (397, 408)	3760	<1%	1%
Dacapo 2006 antlr	Preprocessor.java (564)	2840	0%	1%
Dacapo 2006 luindex	DocumentWriter.java (206)	3055	0%	0%
Dacapo 9.12 lusearch	IndexSearcher.java (98)	15179	<1%	0%
Dacapo 9.12 lusearch-fix	FastCharStream.java (54)	225060	<1%	0.5%
Dacapo 9.12 batik	ExtendedGeneralPath.java (743)	2470	0%	0%
SPECjbb2000	StockLevelTransaction.java (173)	116376	<1%	1%
JGFMonteCarloBench 2.0	RatePath.java (296)	60000	0%	0%

WS: Whole-program speedup after optimization

Table 2: Optimizing insignificant objects yields little speedups.

```
165 protected void transform_internal (double data[], int direction) {
166   for (int bit = 0, dual = 1; bit < logn; bit++, dual *= 2) {
167     for (int a = 1; a < dual; a++) {
168       for (int b = 0; b < n; b += 2 * dual) {
169         int i = 2*(b + a);
170         int j = 2*(b + a + dual);
171         ►double z1_real = data[j];
172         ►double z1_imag = data[j+1];
173         ...
174         ►data[j]   = data[i]   - wd_real;
175         ►data[j+1] = data[i+1] - wd_imag;
176         ...
177 }}}}
```

Listing 6: DJXPerf identifies the `data` array with poor locality in SPECjvm2008: Scimark.fft.large.

```
235 //Interval.java
236 public Integer[] toArray() {
237   ►Integer[] result = new Integer[this.size()];
238   ...
239   return result;
240 }
241 //InternalArrayIterate.java
242 private static <T> void batchFastListCollect(T[] array, ...) {
243   ...
244   for (int i = start; i < end; i++)
245     ►castProcedure.value(array[i]);
246   ...
247 }
```

Listing 7: DJXPerf pinpoints the `array` array suffering from NUMA remote accesses in Eclipse Collections.

7.4 SPECjvm2008: Scimark.fft.large

Scimark [3] is a composite Java benchmark measuring the performance of numerical codes occurring in scientific applications. Scimark.fft refers to a fast Fourier transform (FFT) implementation. We run Scimark.fft with its large input released with the benchmark. DJXPerf reports the object `data`, which is an array, suffering most from cache miss (accounting for 75.5% of total cache misses). The most problematic accesses are at lines 171, 172, 174, and 175 in method `transform_internal` of class `FFT`, as shown in Listing 6. From the code listing we can see that array `data` is accessed in a 3-level loop nest. The innermost loop index `b` increases by $2 * dual$ every iteration, and *dual* is also doubled every iteration of the outer-most loop. Thus, the accessing stride for `data` array is large, resulting in the poor spatial locality. For optimization, we interchange loops `a` and `b` to reduce the stride. This optimization reduces cache misses of the entire program by 70%, yielding a $(2.37 \pm 0.07)\times$ speedup.

7.5 Eclipse Collections

Eclipse Collections is a comprehensive Java collections library, which enables productivity and performance by delivering an expressive and efficient set of APIs and types [40]. We run eclipse collections using CollectTest as input. After profiling eclipse collections, DJXPerf reports an object that suffers from NUMA remote accesses, the Integer array `result`, which is allocated at line 758 in method `toArray` of class `Interval` and accessed at line 245 in method `batchFastListCollect` of class `InternalArrayIterate` with a high percentage of NUMA remote accesses (73.4%). By investigating the source code as shown in Listing 7, the program first calls to `toArray` method to allocate and initialize an Integer array `result`. And then, the program passes the `result` to `batchFastListCollect` method and accesses it at line 245. DJXPerf detects that there's a mismatch between the allocation of `result` and initialization by the master thread in one NUMA domain, and accesses by the worker threads executing in other NUMA domains. The workers all compete for memory bandwidth to access the data in the master's NUMA domain. To avoid contending for data allocated in a single NUMA domain, we optimized the program as allocating and initializing the object `result` in every NUMA domain. The optimization reduces remote accesses by 41% and increases the throughput (operations per second) by $(1.13 \pm 0.04)\times$.

7.6 Apache Druid

Apache Druid is a high performance real-time analytics database designed for workflows where fast queries and ingest matter [39]. We run Apache Druid using BitmapIterationBenchmark as input. After profiling SP, DJXPerf reports an array an BitSet object, `bitmap`, which is initialized at line 37 in constructor method `WrappedImmutableBitSetBitmap` of class `WrappedImmutableBitSetBitmap` and accessed at line 120 in method `next` of same class. Listing 8 shows the problematic object is the `bitmap` (accessed at line 120), which more than half of total memory accesses are remote accesses. With further investigation, we find that the problematic object `bitmap` is initialized in constructor method `WrappedImmutableBitSetBitmap` executed in one NUMA domain, but accessed by many threads in other NUMA domains. To address this problem, we parallelize the allocation and initialization for this object `bitmap` to ensure that each thread first touches its own data. With this optimization, we reduce remote accesses by 47% and increases the throughput by $(1.75 \pm 0.05)\times$.

```
112 public class WrappedImmutableBitSetBitmap {
113   protected final BitSet bitmap;
114   public WrappedImmutableBitSetBitmap(BitSet bitmap) {
115     ►this.bitmap = bitmap;
116   }
117   ...
118   public int next() {
119     int pos = nextPos;
120     ►nextPos = bitmap.nextSetBit(pos + 1);
121     return pos;
122   }
123 }
```

Listing 8: DJXPerf pinpoints the BitSet object `bitmap` suffering from NUMA remote accesses in Apache Druid.

7.7 Attempts to Optimization for Insignificant Objects

To demonstrate the importance of PMU metrics (i.e., cache misses in our experiments) associated with the objects, we show a number of studies on attempting to optimize insignificant objects in Table 2. All these code bases have the memory bloat problem: repeatedly allocate objects many times, and different instances have no overlap in their life intervals. In the table, we show the location of problematic object allocations, the number of object instances, the associated cache miss metrics, and the speedups after optimization. Our studies show that these optimizations yield negligible speedups, which emphasize the fact that these objects account for very few cache misses. Thus, DJXPerf's object-centric analysis is useful to filter out insignificant objects for non-fruitful optimization.

8 Conclusions

In this paper, we present DJXPerf, the very first lightweight Java profiler that performs object-centric analysis to identify data locality issues in Java applications. DJXPerf leverages the lightweight Java byte code instrumentation and the hardware PMU available in commodity CPU processors. DJXPerf works for off-the-shelf Linux OS and Oracle Hotspot JVM, as well as unmodified Java applications. DJXPerf incurs low overhead, typically 8% in runtime and 5% in memory. These features make DJXPerf applicable to the production environment. DJXPerf is able to identify a number of locality issues in real-world Java applications. Such locality issues arise due to traditional spatial/temporal data locality and also due to memory bloat. Guided by DJXPerf, we are able to perform optimization, which yields nontrivial speedups. DJXPerf is open-sourced at an anonymous URL.

References

[1] [n.d.]. Java on Google Cloud. https://cloud.google.com/java.
[2] 2000. SPEC JBB2000. https://www.spec.org/jbb2000/.
[3] 2004. Scimark. https://www.spec.org/jvm2008/docs/benchmarks/scimark.html.
[4] 2013. Obtain object address. https://jrebel.com/rebellabs/dangerous-code-how-to-be-unsafe-with-java-classes\-objects-in-memory/4/.
[5] 2017. ScalaSTM. https://nbronson.github.io/scala-stm/.
[6] 2018. DaCapo Benchmark Suite 9.12. https://sourceforge.net/projects/dacapobench/files/9.12-bach-MR1/.
[7] 2018. OProfile. http://oprofile.sourceforge.net.
[8] 2019. The Scala Programming Language. https://www.scala-lang.org.
[9] Matthew Arnold and Peter F. Sweeney. 1999. *Approximating the Calling Context Tree via Sampling.* Technical Report 21789. IBM.
[10] Cédric Bastoul and Paul Feautrier. 2003. Improving Data Locality by Chunking. International Conference on Compiler Construction, 320–334.
[11] Ed Baum, David abd Baum. [n.d.]. Netflix powers through 2 billion content requests per day with Java-driven architecture. https://go.java/netflix.html.
[12] Ed Baum, David abd Baum. [n.d.]. Twitter migrates core infrastructure to Java Virtual Machine and supports more than 400 million tweets a day. https://go.java/twitter.html.
[13] Erik Berg and Erik Hagersten. 2005. Fast data-locality profiling of native execution. *SIGMETRICS Perform. Eval. Rev.* 33, 1 (2005), 169–180.
[14] Kristof Beyls and Erik D'Hollander. 2006. Discovery of Locality-Improving Refactorings by Reuse Path Analysis. *Proc. of the 2nd Intl. Conf. on High Performance Computing and Communications (HPCC)* (2006).
[15] Kristof Beyls and Erik H. D'Hollander. 2009. Refactoring for Data Locality. *Computer* 42, 2 (Feb. 2009), 62 –71.
[16] Eric Bruneton, Romain Lenglet, and Thierry Coupaye. 2019. ASM Framework. https://asm.ow2.io.
[17] Bryan R. Buck and Jeffrey K. Hollingsworth. 2004. Data Centric Cache Measurement on the Intel Itanium 2 Processor. In *SC*.
[18] Mark Bull, Lorna Smith, Martin Westhead, David Henty, and Robert Davey. 2001. Java Grande benchmark suite. https://www.epcc.ed.ac.uk/research/computing/performance-characterisation-and-benchmarking/java-grande-benchmark-suite.
[19] Calin Caşcaval and David A. Padua. 2003. Estimating cache misses and locality using stack distances. In *Proc. of the 17th annual Intl. conference on Supercomputing* (San Francisco, CA, USA) *(ICS '03)*. ACM.
[20] Maria Carpen-Amarie, Patrick Marlier, Pascal Felber, and Gaël Thomas. 2015. A performance study of java garbage collectors on multicore architectures. Proceedings of the 20th International Conference on Architectural Support for Programming Languages and Operating Systems, 20–29.
[21] Oracle Corporation. 2007. JVMTI Tool Interface Online Document. https://docs.oracle.com/javase/8/docs/platform/jvmti/jvmti.html.
[22] Oracle Corporation. 2018. All-in-One Java Troubleshooting Tool. https://visualvm.github.io.
[23] Oracle Corporation. 2019. Oracle HotSpot JVM. https://openjdk.java.net/groups/hotspot/.
[24] Standard Performance Evaluation Corporation. 2008. SPECjvm2008 benchmark suite. https://www.spec.org/jvm2008.
[25] John Cuthbertson, Sandhya Viswanathan, Konstantin Bobrovsky, Alexander Astapchuk, and Eric Kaczmarek. Uma Srinivasan. 2009. A Practical Approach to Hardware Performance Monitoring Based Dynamic Optimizations in a Production JVM. Proceedings of the 7th annual IEEE/ACM International Symposium on Code Generation and Optimization (CGO'09).
[26] Luca Della Toffola, Michael Pradel, and Thomas R. Gross. 2015. Performance Problems You Can Fix: A Dynamic Analysis of Memoization Opportunities. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Pittsburgh, PA, USA) *(OOPSLA 2015)*. ACM, New York, NY, USA, 607–622.
[27] Monika Dhok and Murali Krishna Ramanathan. 2016. Directed Test Generation to Detect Loop Inefficiencies. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering* (Seattle, WA, USA) *(FSE 2016)*. ACM, New York, NY, USA, 895–907.
[28] Chen Ding and Yutao Zhong. 2003. Predicting whole-program locality through reuse distance analysis. In *Proc. of the ACM SIGPLAN 2003 conference on Programming Language Design and Implementation* (San Diego, California, USA) *(PLDI '03)*. ACM, New York, NY, USA, 245–257.
[29] NASA Advanced Supercomputing Division. 2018. NAS Parallel Benchmarks. https://www.nas.nasa.gov/publications/npb.html.
[30] Daniel Dominic Sleator and Robert Endre Tarjan. 1985. Self-Adjusting Binary Search Trees. Journal of the Association for Computing Machinery, 652–686.
[31] Paul J. Drongowski. 2010. Instruction-Based Sampling: A New Performance Analysis Technique for AMD Family 10h Processors.
[32] Bruno Dufour, Barbara G. Ryder, and Gary Sevitsky. 2007. Blended analysis for performance understanding of framework-based applications. In ACM SIGSOFT Int. Symp. Software Testing and Analysis, 118–128.
[33] Bruno Dufour, Barbara G. Ryder, and Gary Sevitsky. 2008. A scalable technique for characterizing the usage of temporaries in framework-intensive Java applications. ACM SIGSOFT Interna- tional Symposium on the Foundations of Software Engineer- ing (FSE), 59–70.
[34] Ariel Eizenberg, Shiliang Hu, Gilles Pokam, and Joseph Devietti. 2016. Remix: Online Detection and Repair of Cache Contention for the JVM. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, CA, USA) *(PLDI '16)*. ACM, New York, NY, USA, 251–265. https://doi.org/10.1145/2908080.2908090
[35] ej-technologies GmbH. 2018. THE AWARD-WINNING ALL-IN-ONE JAVA PROFILER. https://www.ej-technologies.com/products/jprofiler/overview.html.
[36] Peter F. Sweeney, Matthias Hauswirth, Brendon Cahoon, Perry Cheng, Amer Diwan, David Grove, and Michael Hind. 2004. Using hardware performance monitors to understand the behavior of Java applications. In Proceedings of the 3rd Virtual Machine Research and Technology Symposium (VM'04).
[37] Apache Software Foundation. 2017. Apache SAMOA: Scalable Advanced Massive Online Analysis. https://samoa.incubator.apache.org.

[38] Apache Software Foundation. 2019. Apache Commons Collections. https://github.com/apache/commons-collections.

[39] Apache Software Foundation. 2020. Druid. https://github.com/apache/druid.

[40] Eclipse Foundation. 2019. Eclipse Collections. https://github.com/eclipse/eclipse-collections.

[41] Andy Georges, Dries Buytaert, Lieven Eeckhout, and Koen De Bosschere. 2004. Method-Level Phase Behavior in Java Workloads. Proc. of the ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), 270–287.

[42] Lokesh Gidra, Gaël Thomas, Julien Sopena, and Marc Shapiro. 2013. A study of the scalability of stop-the-world garbage collectors on multicores. Proceedings of the eighteenth international conference on Architectural support for programming languages and operating systems, 229–240.

[43] Lokesh Gidra, Gaël Thomas, Julien Sopena, Marc Shapiro, and Nhan Nguyen. 2015. NumaGiC: a Garbage Collector for Big Data on Big NUMA Machines. Proceedings of the 20th International Conference on Architectural Support for Programming Languages and Operating Systems, 661–673.

[44] YourKit GmbH. 2018. The Industry Leader in .NET & Java Profiling. https://www.yourkit.com.

[45] Sasha Goldshtein. 2018. Profiling JVM Applications in Production. USENIX Association.

[46] Xiaoming Gu, Xiao-Feng Li, Buqi Cheng, and Eric Huang. 2009. Virtual Reuse Distance Analysis of SPECjvm2008 Data Locality. PPPJ '09 Proceedings of the 7th International Conference on Principles and Practice of Programming in Java, Calgary, Alberta, Canada, 182–191.

[47] Matthias Hauswirth, Peter F. Sweeney, Amer Diwan, and Michael Hind. 2004. Vertical Profiling: Understanding the Behavior of Object-Oriented Applications. Proc. of Conf. on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), 251–269.

[48] Rich Hickey. 2019. The Clojure Programming Language. https://clojure.org.

[49] Peter Hofer and Hanspeter Mössenböck. 2014. Fast Java Profiling with Scheduling-aware Stack Fragment Sampling and Asynchronous Analysis. In *Proceedings of the 2014 International Conference on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools* (Cracow, Poland) *(PPPJ '14)*. ACM, New York, NY, USA, 145–156.

[50] Charles Humble. [n.d.]. Twitter Shifting More Code to JVM, Citing Performance and Encapsulation As Primary Drivers. https://www.infoq.com/articles/twitter-java-use/.

[51] Intel Corporation. 2010. Intel 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide, Part 2, Number 253669-032.

[52] Intel Corporation. 2019. Intel VTune Performance Analyzer. https://software.intel.com/en-us/vtune.

[53] Nguyen Khanh and Guoqing Xu. 2013. Cachetor: Detecting Cacheable Data to Remove Bloat. ESEC/FSE 2013 Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, 268–278.

[54] Sameer Kumar and Chris Wilkerson. 1998. Exploiting spatial locality in data caches using spatial footprints. Proceedings. 25th Annual International Symposium on Computer Architecture, 357–368.

[55] Renaud Lachaize, Baptiste Lepers, and Vivien Quéma. 2012. MemProf: A memory profiler for NUMA multicore systems. In *USENIX ATC*.

[56] Jeremy Lau, Matthew Arnold, Michael Hind, and Brad Calder. 2006. Online performance auditing: Using hot optimizations without getting burned. In Proc. Conf. on Programming Language Design and Implementation (PLDI 2006), New York, USA, 239–251.

[57] Linux. 2015. Linux Perf Tool. http://www.brendangregg.com/perf.html.

[58] Xu Liu and John Mellor-Crummey. 2011. Pinpointing data locality problems using data-centric analysis. In *Proc. of the 9th IEEE/ACM Intl. Symp. on Code Generation and Optimization*. Washington, DC, 171–180.

[59] Xu Liu and John Mellor-Crummey. 2013. Pinpointing Data Locality Bottlenecks with Low Overheads. In *Proc. of the 2013 IEEE Intl. Symp. on Performance Analysis of Systems and Software*.

[60] Xu Liu and John Mellor-Crummey. 2014. A Tool to Analyze the Performance of Multithreaded Programs on NUMA Architectures. In *PPoPP*.

[61] Xu Liu and John M. Mellor-Crummey. 2013. A Data-centric Profiler for Parallel Programs. In *SC*.

[62] Xu Liu, Kamal Sharma, and John Mellor-Crummey. 2014. ArrayTool: a lightweight profiler to guide array regrouping. In *PACT*.

[63] Xu Liu and Bo Wu. 2015. ScaAnalyzer: A Tool to Identify Memory Scalability Bottlenecks in Parallel Programs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Austin, Texas) *(SC '15)*. ACM, New York, NY, USA, Article 47, 12 pages. https://doi.org/10.1145/2807591.2807648

[64] Lucie Lozinski. [n.d.]. The Uber Engineering Tech Stack, Part I: The Foundation,. https://eng.uber.com/tech-stack-part-one-foundation/.

[65] Stephen M. Blackburn, Robin Garner, Chris Hoffmann, Asjad M. Khang, Kathryn S. McKinley, Rotem Bentzur, Amer Diwan, Daniel Feinberg, Daniel Frampton, Samuel Z. Guyer, Martin Hirzel, Antony Hosking, Maria Jump, Han Lee, J. Eliot B. Moss, Aashish Phansalkar, Darko Stefanović, Thomas VanDrunen, Daniel von Dincklage, and Ben Wiedermann. 2006. The DaCapo benchmarks: java benchmarking development and analysis. OOPSLA '06 Proceedings of the 21st annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications, Portland, Oregon, USA, 169–190.

[66] Collin McCurdy and Jeffrey S. Vetter. 2010. Memphis: Finding and fixing NUMA-related performance problems on multi-core platforms. In *Proc. of 2010 IEEE Intl. Symp. on Performance Analysis of Systems Software*.

[67] Milena Milenkovic, Scott T. Jones, Frank Levine, and Enio Pineda. 2010. Performance Inspector Tools with Instruction Tracing and Per-Thread / Function Profiling. Linux Symposium, Volume Two.

[68] Nike Mitchell, Edith Schonberg, and Gary Sevitsky. 2009. Making Sense of Large Heaps. European Conference on Object-Oriented Programming (ECOOP), 77–97.

[69] Nick Mitchell, Edith Schonberg, and Gary Sevitsky. 2010. Four trends leading to Java runtime bloat. IEEE Software, 56–63.

[70] Nick Mitchell and Gary Sevitsky. 2007. The causes of bloat, the limits of health. ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), 245–260.

[71] Mustafa M.Tikir and Jeffrey K.Hollingsworth. 2005. NUMA-Aware Java Heaps for Server Applications. Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium.

[72] Todd Mytkowicz, Amer Diwan, Matthias Hauswirth, and Peter F. Sweeney. 2010. Evaluating the Accuracy of Java Profilers. In *Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation* (Toronto, Ontario, Canada) *(PLDI '10)*. ACM, New York, NY, USA, 187–197.

[73] Adrian Nistor, Linhai Song, Darko Marinov, and Shan Lu. 2013. Toddler: Detecting Performance Problems via Similar Memory-access Patterns. In *Proceedings of the 2013 International Conference on Software Engineering* (San Francisco, CA, USA) *(ICSE '13)*. IEEE Press, Piscataway, NJ, USA, 562–571.

[74] Nitsan Wakart. 2016. The Pros and Cons of AsyncGetCallTrace Profilers. http://psy-lob-saw.blogspot.com/2016/06/the-pros-and-cons-of-agct.html.

[75] Andrei Pangin. 2018. Async-profiler. https://github.com/jvm-profiling-tools/async-profiler.

[76] Aleksandar Prokopec, Andrea Rosà, David Leopoldseder, Gilles Duboscq, Petr Tůma, Martin Studener, Lubomír Bulej, Yudi Zheng, Alex Villazón, Doug Simon, Thomas Würthinger, and Walter Binder. 2019. Renaissance: Benchmarking Suite for Parallel Applications on the JVM. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) *(PLDI 2019)*. ACM, New York, NY, USA, 31–47. https://doi.org/10.1145/3314221.3314637

[77] Bill Pugh, Andrey Loskutov, and Keith Lea. 2015. FindBugs. http://findbugs.sourceforge.net.

[78] Prakash Raghavendra and Sandya Mannarsamy. 2006. Optimal Placement of Java Objects and Generated Code in a Cell-based CC-NUMA Machine. The 13th annual IEEE International Conference on High Performance Computing.

[79] Ashay Rane and James Browne. 2012. Enhancing performance optimization of multicore chips and multichip nodes with data structure metrics. In *PACT*.

[80] Rhuan Rocha. [n.d.]. Why Java is so hot right now. https://developers.redhat.com/blog/2019/09/05/why-java-is-so-hot-right-now/.

[81] Rhuan Rocha. 2019. Top 5 Programming Languages in Cloud Enterprise Development. https://www.scnsoft.com/blog/cloud-development-languages.

[82] Probir Roy, Xu Liu, and Shuaiwen Leon Song. 2016. SMT-Aware Instantaneous Footprint Optimization. In *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing* (Kyoto, Japan) *(HPDC '16)*. ACM, New York, NY, USA, 267–279. https://doi.org/10.1145/2907294.2907308

[83] Probir Roy, Shuaiwen Leon Song, Sriram Krishnamoorthy, and Xu Liu. 2018. Lightweight Detection of Cache Conflicts. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization* (Vienna, Austria) *(CGO 2018)*. ACM, New York, NY, USA, 200–213. https://doi.org/10.1145/3168819

[84] Shyam Saladi. 2013. ranklib correlation. https://github.com/smsaladi/ranklib-correlation.

[85] Ajeet Shankar, Mattew Arnold, and Rastislav Bodik. 2008. JOLT: Lightweight Dynamic Analysis and Removal of Object Churn. ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), 127–142.

[86] Xipeng Shen, Jonathan Shaw, Brian Meeker, and Chen Ding. 2007. Locality approximation using time. In *Proc. of the 34th annual ACM SIGPLAN-SIGACT Symposium on Principles Of Programming Languages* (Nice, France) *(POPL '07)*. ACM, New York, NY, USA, 55–61.

[87] Linhai Song and Shan Lu. 2017. Performance Diagnosis for Inefficient Loops. In *Proceedings of the 39th International Conference on Software Engineering* (Buenos Aires, Argentina) *(ICSE '17)*. IEEE Press, Piscataway, NJ, USA, 370–380.

[88] M. Srinivas et al. 2011. IBM POWER7 performance modeling, verification, and evaluation. *IBM JRD* 55, 3 (May/June 2011), 4:1–19.

[89] Pengfei Su, Qingsen Wang, Chabbi Milind, and Xu Liu. 2019. Pinpointing Performance Inefficiencies in Java. The 27th ACM Joint European Software

Engineering Conference and Symposium on the Foundations of Software Engineering, Tallinn, Estonia.
[90] Nathan R. Tallent. 2010. *Performance analysis for parallel programs from multicore to petascale.* Ph.D. thesis. Department of Computer Science, Rice University.
[91] Gil Tene and Martin Thompson. 2018. ObjectLayout. https://github.com/ObjectLayout/ObjectLayout.
[92] Vince Weaver. 2013. The Unofficial Linux Perf Events Web-Page. http://web.eece.maine.edu/~vweaver/projects/perf_events. Last accessed: Dec. 12, 2013.
[93] Jens Wike, Ayman Abdel Ghany, Filipe Manuel, Tatsiana Shpikat, and Andreas Worm. 2013. cache2k. https://cache2k.org.
[94] Guoqing Xu. 2011. Analyzing Large-Scale Object-Oriented Software to Find and Remove Runtime Bloat. PhD thesis.
[95] Guoqing Xu. 2012. Finding reusable data structures. ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), 1017–1034.
[96] Guoqing Xu, Matthew Arnold, and Nick Mitchell. 2009. Go with the Flow: Profiling Copies To Find Runtime Bloat. Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '09), 419–430.
[97] Guoqing Xu, Nick Mitchell, Matthew Arnold, Atanas Rountev, and Gary Sevitsky. 2010. Software bloat analysis: Finding, removing, and preventing performance problems in modern large-scale object-oriented applications. FSE/SDP Working Conference on the Future of Software Engineering Research (FoSER), 421–426.
[98] Guoqing Xu and Atanas Rountev. 2010. Detecting Inefficiently-Used Containers to Avoid Bloat. Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '10)., 160–173.
[99] Dacong Yan, Guoqing Xu, and Atanas Rountev. 2012. Uncovering performance problems in Java applications with reference propagation profiling. International Conference on Software Engineering (ICSE), 134–144.
[100] Yutao Zhong and Wentao Chang. 2008. Sampling-based program locality approximation. In *Proc. of the 7th Intl. Symposium on Memory Management* (Tucson, AZ, USA) *(ISMM '08)*. ACM, New York, NY, USA, 91–100.
[101] Yutao Zhong, Xipeng Shen, and Chen Ding. 2009. Program locality analysis using reuse distance. *ACM Trans. Program. Lang. Syst.* 31, 6 (Aug. 2009), 20:1–20:39.