

Cloud-Scale Java Profiling at Alibaba

Fangxi Yin, Denghui Dong, Chuansheng Lu, Tongbao Zhang, Sanhong Li, Jianmei Guo and Kingsum Chow

Alibaba Group, Hangzhou, China

{fangxi.yfx,denghui.ddh,chuansheng.lcs,tongbao.ztb,sanhong.lsh,jianmei.gjm,kingsum.kc}@alibaba-inc.com

ABSTRACT

On the 2017 Double 11 Global Shopping Festival, Alibaba's cloud platform achieved total sales of more than 25 billion dollars and supported peak volumes of 325,000 transactions and 256,000 payments per second. Most of the cloud-based e-commerce transactions were processed by hundreds of thousands of Java applications with above a billion lines of code. It is challenging to achieve comprehensive and efficient performance profiling for large-scale, cloud-based Java applications in production. We developed ZPROFILER, a fine-grained, low-overhead Java performance profiler. ZPROFILER allows developers to load a profiling agent on the fly without restarting Java virtual machines, and its profiling information also facilitates code warmup. ZPROFILER is developed based on Alibaba JDK (AJDK), a customized version of OpenJDK, and it has been rolled out to Alibaba's cloud platform to support large-scale performance tuning for online critical business.

CCS CONCEPTS

• **Software and its engineering** → **Cloud computing**; **Software performance**;

KEYWORDS

Java performance, cloud, profiling, overhead, code warmup

ACM Reference Format:

Fangxi Yin, Denghui Dong, Chuansheng Lu, Tongbao Zhang, Sanhong Li, Jianmei Guo and Kingsum Chow. 2018. Cloud-Scale Java Profiling at Alibaba. In *ICPE '18: ACM/SPEC International Conference on Performance Engineering Companion*, April 9–13, 2018, Berlin, Germany. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3185768.3186295>

1 INTRODUCTION

Performance profiling helps developers to diagnose performance issues and identify optimization opportunities. At Alibaba, the requirements for Java performance optimization come mainly from three aspects: First, how to execute Java code efficiently (e.g., using less time to execute a certain method)? Second, how to tune the behavior of *garbage collector* (GC) to minimize the time of *Stop-The-World* (STW)? Third, how to ensure the peak performance of Java applications for the bursty traffic over a short time period.

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.

ICPE '18, April 9–13, 2018, Berlin, Germany

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5629-9/18/04...\$15.00

<https://doi.org/10.1145/3185768.3186295>

It is challenging to achieve comprehensive and efficient performance profiling for Java applications in production. Firstly, many existing profiling methods and tools [3], such as VisualVM [4], Dynatrace[2], and Java Interactive Profiler [6], rely on *bytecode instrumentation* to trace the execution time of each method. They usually suffer from non-negligible overheads caused by the measurement of additional code. Also, changing the size of the code changes the optimizing decisions made by the JIT. Secondly, it is usually hard for developers to diagnose the issues relevant to GC only through the straightforward analysis of GC logs [5]. For example, when a Java application is experiencing a long GC pause, it is not reliable to depend on GC logs to diagnose issues. Thirdly, a peak of 325,000 transactions per second is a huge challenge for the runtime performance of any Java program. The bursty traffic requires that the application must remain a peak performance at that time. To meet this requirement, a well-established code warmup is often expected [1].

Further, there are more particular challenges for Java profiling at cloud-scale environment. First, the profiling activities take place frequently because of the massive quantity of JVM instances. So the profiling tools' troubleshooting efficiency is critical and an intuitive profiling tool is always favored over. Second, the overhead requirement of profiling is stricter. At cloud-scale environment, even one percent overhead would increase cost significantly. To compensate the overhead, a common way is to purchase more machines.

To address the above challenges, we developed a fine-grained, low-overhead Java performance profiler, called ZPROFILER, based on Alibaba JDK (AJDK). AJDK supports all the Java applications developed at Alibaba and affiliated companies (e.g., Alipay, CaiNiao). Java developers can easily load the profiling agent *on the fly* without restarting the JVM.

In summary, we make the following contributions. Firstly, different from typical Java profilers working at the levels of bytecode or JVM Tool Interface (JVMTI), ZPROFILER traces the execution time and memory usage at the level of each compiled method. Secondly, ZPROFILER provides fine-grained information of GC behavior with much lower overheads, which used to minimize the STW cost. Thirdly, ZPROFILER has the built-in code cache analysis ability to help users to warmup code by ahead-of-time compilation, which makes application ready for bursty traffic handling. Finally, ZPROFILER has been rolled out to Alibaba's cloud platform to support large-scale performance tuning for online critical business.

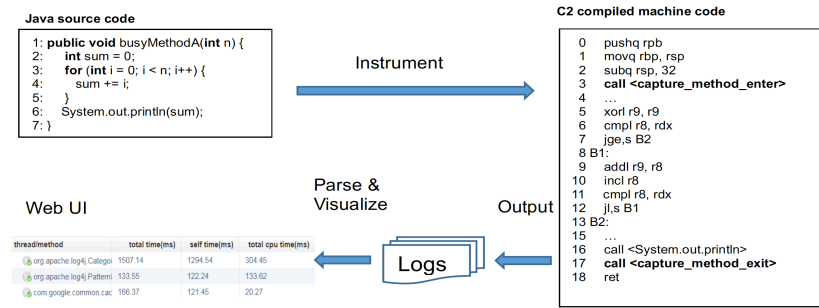


Figure 1: Method Tracing of ZPROFILER in C2 compiler

2 APPROACH AND IMPLEMENTATION

2.1 Method Tracing

Instead of instrumenting at the level of Java bytecode, ZPROFILER modifies the Hotspot VM, including interpreter and C1/C2 compiler, to trace entry/exit of each method. Figure 1 illustrates how ZPROFILER enables method tracing in C2 compiler.

When C2 JIT compiler compiles the application's method, ZPROFILER adds a few machine code for every method, that is, adding capture instructions when compiler encounters method entry or exit. For every method enter or exit, the captured runtime functions record the onsite information, including the current CPU/wall-clock time, method information, and the allocated memory of each thread. When the compiled method is executed, the onsite information would be recorded in the profiling files. When profiling is switched off, ZPROFILER uploads the profiling files, parses the profiling data, and visualizes them in a Web UI.

ZPROFILER's method tracing does not impact JIT inline decisions because its instrumentation is implemented at VM level instead of bytecode level. Moreover, we added only a few extra instructions for non-inline methods, so that the profiling overheads are reduced as far as possible. Most of APM tools [2] often instrument applications at the byte code level. The byte code instrumentation often impacts JIT inline decision thereby affects the performance. A common case is, the getter/setter functions which are widely used can not be inlined and the performance loss is non-negligible.

2.2 GC Tuning

ZPROFILER modifies the implementation of GC in the Hotspot VM and records the information of time consumption at each of GC root processing phases. Developers can use the recorded information to determine which phase gives rise to a longer GC pause. For example, if one acquires the information that the StringTable root processing takes a long time during a GC pause, then the developer can check StringIntern use in the implementation and get clues to fix this problem. Moreover, ZPROFILER provides mxbeans for monitoring the fine-grained information.

2.3 Code Warmup

ZPROFILER extended jcmd utility of OpenJDK to support code cache dump, which helps them to understand if the Java application has been fully warmed up and ready for processing bursty traffic. The

code cache dump provides the summary information of code cache usage, the memory usage of compiled method by class loader, and the distribution of compiled methods at each level (from level 0 to level 4) of code cache. AJDK provides public APIs to trigger compilation explicitly at *ahead-of-time*. For example, if one finds many methods get compiled at Level 2 in code cache, the developer can choose to compile some of them in advance according to business logic.

3 CONCLUSION

This paper presents our Java performance profiler ZPROFILER that supports hundreds of thousands of Java applications deployed in Alibaba's cloud platform. Compared to existing Java profilers, ZPROFILER instruments each method with a few machine code and traces the resource usage of each compiled method with low overheads. Simultaneously, ZPROFILER provides a fine-grained view on GC behavior and code cache, which facilitates comprehensive performance diagnosis and efficient fixing of performance bugs. Furthermore, ZPROFILER allows Java developers to easily load the profiling agent on the fly without restarting the JVM or code de-optimization, which is critical to guarantee the reliability and serviceability of online Java applications in production.

REFERENCES

- [1] Edd Barrett, Carl Friedrich Bolz-Tereick, Rebecca Killick, Sarah Mount, and Laurence Tratt. 2017. Virtual machine warmup blows hot and cold. *Proceedings of the ACM on Programming Languages* 1 (OOPSLA) (2017), 52:1–52:27.
- [2] Dynatrace. 2018. Deliver unrivaled digital experiences. (2018). <https://www.dynatrace.com/>
- [3] Todd Mytkowicz, Amer Diwan, Matthias Hauswirth, and Peter F. Sweeney. 2010. Evaluating the accuracy of Java profilers. In *Proceedings of PLDI*. 187–197.
- [4] Oracle. 2014. Java VisualVM. (2014). <https://docs.oracle.com/javase/8/docs/technotes/guides/visualvm/profiler.html>
- [5] Sun and IBM. 2002. Tool Report: GCViewer. (2002). <http://www.javaperformancetuning.com/tools/gcviewer>
- [6] Andrew Wilcox. 2006. JIP - The Java Interactive Profiler. (2006). <http://jiprof.sourceforge.net/>