JOHANNES KEPLER
UNIVERSITY LINZ

Dipl.-Ing. Markus Weninger
Institute for System Software

P +43 732 2468 4361
F +43 732 2468 4345
markus.weninger@jku.at

Secretary:
Birgit Kranzl
Ext. 4341
birgit.kranzl@jku.at

Master's Thesis

Allocation Profiling Analysis

Student:	Simon Eilmsteiner
Advisor:	Dipl.-Ing. Markus Weninger
Dynatrace Advisor:	Dr. Philipp Lengauer
Start date:	01.04.2019

Linz, March 4, 2019

Dynatrace is performance monitoring tool supporting, among many others, Java. Many Java applications suffer from high GC times, GC spikes, or high object churn. These problems are easy to identify, but very hard to diagnose because the only tool at the develops disposal are heap dumps. Heap dumps can tell what types of objects are currently in the heap, but cannot tell where an object has been allocated. In some cases, the type of an object is hint enough to find the responsible code location, but many types, e.g., strings, are allocated at many different locations. To make matters even worse, the immediate allocation site, i.e., the method allocating the object, does not provide enough additional information because these types are always allocated in some low-level factory methods, e.g., *String::substring* or *String::format*.

Starting with Java 11, the Virtual Machines provide an interface to sample object allocations and also provide information about the allocation site and its stack. The *Dynatrace OneAgent* uses this interface to build a call tree, annotate its nodes with allocation information, and send it periodically to the Dynatrace Cluster.

The goal of this thesis is to implement an analysis based on the data provided by the *OneAgent* to determine the allocation pattern of an application, determine the GC impact of individual allocation sites, and provide an estimate of how much more load the application can support without violating its availability constraints. In addition to the allocation profiling data, the analysis can also rely on heap dump statistics as well as garbage collection statistics to improve its estimations. Collecting and sending the data in the *OneAgent* is not part of this thesis.

The progress of the project should be discussed at least every three weeks with the advisor. A time schedule and a milestone plan must be set up within the first 3 weeks. It should be continuously refined and monitored to make sure that the thesis will be completed in time.

The written thesis must contain both a quantitative and qualitative evaluation based on several different real-world applications and benchmarks. The final version of the thesis should be submitted not later than 31.03.2020.

JOHANNES KEPLER
UNIVERSITY LINZ
Altenberger Straße 69
4040 Linz, Austria
jku.at
DVR 0093696