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

Bachelor's Thesis
Utilizing Memory-mapped Files to Reduce Memory Overhead in AntTracks

Student: Felix Batusic
SKZ/Matr.Nr.: 1455064
Email: felix@batusic.net
Advisor: Dipl.-Ing. Markus Weninger, BSc.
Start date: Feb 2018

AntTracks comprises a modified Java VM based on the Hotspot VM, i.e., AntTracks VM, and an offline post-processing analysis tool.

The VM's aim is to allow tracking of an application's entire life cycle by writing information about certain events to a trace file. This events include object allocations, object movements by the garbage collector, pointers between the objects and so on.
Such an event trace can then be analyzed in the offline post-processing tool. Based on the information parsed from the trace file the tool is able to reconstruct the heap for any garbage collection point.

A potential risk of memory analysis tools, especially when analyzing heap states of applications that ran out of memory on the same machine, is to also run out of memory during analysis. For each heap object, information such as type, allocation site, references, etc. has to be reconstructed and stored in memory. AntTracks currently receives extended support for object graph analysis, which requires us to store additional information about the references between objects (e.g., dominator tree relation).

The goal of this thesis is to refactor AntTracks's current data structures that represent the heap state and the heap state classification (i.e., the heap state analysis result). Currently, both data structures store the data completely in memory, mostly using arrays of Java basic data types (e.g., *int[]*).
Instead of in-memory arrays, the data should be stored in memory-mapped files. This allows the OS to automatically load and flush page-sized data chunks on-demand, reducing the applications memory overhead.
A potential risk when using this technique is performance degradation, since data may have to be transferred between the RAM and the secondary data storage, and vice-versa. Therefore, the thesis should also include an evaluation on how the usage of memory-mapped files influences AntTracks's memory utilization and heap state analysis performance.

The final version of the written thesis must be submitted not later than 20.08.2018.

JOHANNES KEPLER
UNIVERSITY LINZ
Altenberger Str. 69
4040 Linz, Austria
www.jku.at