Memomania: From Huge to Huge-Huge Pages

Mohammad Agbarya Idan Yaniv Dan Tsafrir
Technion – Israel Institute of Technology
Haifa, Israel
{a.mohammad,idanyani,dan}@cs.technion.ac.il

ABSTRACT

Linux Transparent Huge Pages (THP) are not able to minimize the virtual memory overhead because they are limited to 2MB pages. We introduce Memomania, a flexible memory allocator that is able to transparently back the application with a mixture of 2MB and 1GB pages. We use Memomania to characterize a subset of SPEC CPU2006 benchmarks and find that (1) using 1GB pages provides a notable runtime improvement compared to THP, and that (2) in some cases, allocating more 2MB pages to an application *degrades* its performance. We suggest that Memomania can be used in (1) placement algorithms, and (2) micro-architectural studies.

1 MOTIVATION

The Linux THP feature is able to trasparently back applications memory with huge pages to improve their performance [1]. Besides suffering from several fairness and latency problems [2], the current THP mechanism is limited to use only 2MB-sized huge pages. Users who wish to squeeze every drop of performance must thus resort to explicit allocation of 1GB-sized huge pages.

We developed Memomania, a flexible memory allocator based on Doug Lea's memory allocator [3] (dlmalloc), which hooks the mmap(), munmap(), and brk() system calls and forwards them to memory pools backed by user-defined mixtures of page sizes. Memomania is implemented as a user-space library to be preloaded at runtime and therefore does not require modifying application source code or installing kernel patches.

SYSTOR '18, June 4–7, 2018, HAIFA, Israel

ACM ISBN 978-1-4503-5849-1/18/06...$15.00
https://doi.org/10.1145/3211890.3211918

2 FINDINGS

benchmark	thp speedup (2MB pages)	memowana speedup (1GB pages)
429.mcf	24%	31%
471.omnetpp	0%	9%
473.astar	2%	7%

Table 1: Backing memory with 1GB pages improves performance relative to the Linux THP.

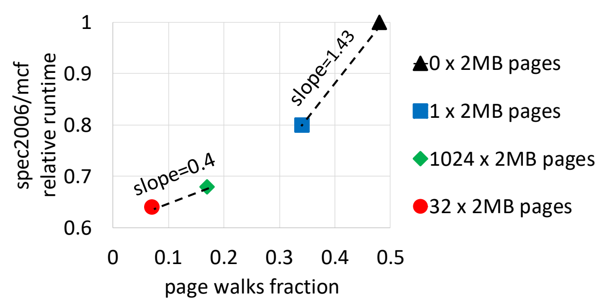

Figure 1: Backing all memory with huge pages is sometimes not the optimal configuration.

3 APPLICATIONS

Besides controlling the memory allocation of applications, Memomania can be used for building runtime models that describe the performance of different huge page configurations. Placement algorithms may use these models to decide which node in an homogeneous cluster will show the best performance for a given application. Additionally, virtual memory studies can use these models to predict the bottom-line runtime from TLB simulations.

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 688386 (OPERA).

REFERENCES

[1] Linux documentation page. 2017. Transparent Hugepage Support. https://www.kernel.org/doc/Documentation/vm/transhuge.txt. (2017). (Accessed: May 2017).
[2] Youngjin Kwon, Hangchen Yu, Simon Peter, Christopher J. Rossbach, and Emmett Witchel. 2016. Coordinated and Efficient Huge Page Management with Ingens. In *OSDI*.
[3] Doug Lea. 2000. A Memory Allocator. http://g.oswego.edu/dl/html/malloc.html. (2000). (Accessed: March 2018).