

Paper Review

Philip Gust

4 Feb 2021

Citation

Kishore Pusukuri, Rob Gardner, and Jared Smolens. "An Implementation of Fast `memset()` Using Hardware Accelerators," 8th International Workshop on Runtime and Operating Systems for Supercomputers (ROSS'18), 2018. [ACM Digital Library](#).

Summary

This paper reports on an experiment to add support to an operating system for using a hardware accelerator to clear large blocks of memory through new and enhanced system calls.

Key Concepts

Clearing large blocks of memory has a significant impact on the performance of big-data operations. It also impacts the performance of data-intensive systems like data stores and garbage collectors that need to clear memory during garbage collection.

The experiment described in this paper modified the Linux operating system by implementing systems calls for 8Mb, 2GB, and 16GB *huge pages*, a special resource supported by the operating system to facilitate big-data operations. The paper described the techniques developed by the authors that use the hardware Data Analytics Accelerator (DAX) on the SPARC M7 processor to fill huge pages with zeros.

After describing the SPARC and DAX hardware, the paper explains the modifications to the Linux operating system, the new and enhanced system calls, and the techniques to utilize the capability. The paper reports on measurements taken after modifications to the Linux kernel on operations using huge pages, as well as on database startup and garbage collection in the Java Runtime Environment (JRE).

The paper reports the following contributions made by the experiment:

- Speed up the creation of 2GB huge pages by 11x compared to the kernel `clear_huge_page()` system call. This achieves the peak achievable write-back memory bandwidth available on SPARC M7 hardware.
- Enhance the `memset()` system call to works on non-contiguous physical memory.
- Show that these techniques improve Oracle database startup times by up to 6.0x and Java (JVM) Garbage Collector (GC) latencies by up to 4.0x.

Assessment

The paper is well-written and clearly explains the rationale and steps taken, the measurements made, and the benefits of the results. It also discusses earlier work on using other hardware acceleration techniques. This is a very practical demonstration of the value of using hardware accelerators to improve the performance of operations at the operating system level, and how to make the functionality available through well-chosen system call enhancements.