

ABSTRACTS COLOC

OPTIMIZING MEMORY BANDWIDTH EFFICIENCY WITH USER-PREFERRED KERNEL MERGE

Nabeeh Jumah and Julian Kunkel

Earth system modeling computations use stencils extensively while running many kernels. Optimal coding of the stencils is essential to efficiently use memory bandwidth of an underlying hardware. This is important as stencil computations are memory bound. Even when the code within one kernel is written to optimally use the memory bandwidth, there are still opportunities for further optimization at the inter-kernel level. Stencils naturally exhibit data locality, and executing a sequence of stencils within separate kernels could waste caching capabilities. Interprocedural optimizations such as merging of kernels bears the potential to improve the use of the caches. However, due to semantic restrictions, it is difficult to achieve on general purpose languages.

Some tools were developed to automatically fuse loops instead of the manual optimization. However, scientists still implement fusion in different levels of loop nests manually to find optimal performance. To allow scientists to still apply loop fusions equal to manual loop fusion, we develop a technique to automatically analyze the code and allow scientists to select their preferred fusions by providing automatic dependency analysis and code transformation; this also bears the potential for automatic tools that make smart choices on behalf of the user. Our work is done using GGDML language extensions which enables performance portability over different architectures using a single source code.

OPPORTUNITIES FOR PARTITIONING NON-VOLATILE MEMORY DIMMS BETWEEN CO-SCHEDULED JOBS ON HPC NODES

Brice Goglin and Andrés Rubio Proaño

The emergence of non-volatile memory DIMMs such as Intel Optane DCPMM blurs the gap between usual volatile memory and persistent storage by enabling byte-accessible persistent memory with reasonable performance. This new hardware supports many possible use cases for high-performance applications, from high performance storage to very-high-capacity volatile memory (terabytes). However the numerous ways to configure the memory subsystem raises the question of how to configure nodes to satisfy applications' needs (memory, storage, fault tolerance, etc.).

We focus on the issue of partitioning HPC nodes with NVDIMMs in the context of co-scheduling multiple jobs. We show that the basic NVDIMM configuration modes would require node reboots and expensive hardware configuration. Moreover it does not allow the co-scheduling of all kinds of jobs, and it does not always allow locality to be taken into account during resource allocation.

Then we show that using 1-Level-Memory and the Device DAX mode by default is a good compromise. It may be easily used and partitioned for storage and memory-bound applications with locality awareness.



Submission deadline for the Euro-Pas PhD Symposium has been extended to 27 May, 2023. Click here for more information - <https://t.co/wwxisICJSC>

17.05.2023 - 11:51

The Euro-Par PhD Symposium is a welcoming and supportive forum for PhD students to present their work. Click here for more information: <https://t.co/wwxisICJSC>

04.04.2023 - 09:25

Submit your paper for EURO-PAR 2023 Workshops and Minisymposia! Click here for more information. <https://t.co/UEseXWb3Dz>

07.03.2023 - 08:18

Abstract submission is due tomorrow 24 Feb, 2023 Click here for more information: <https://t.co/eH2C9CRZA3>

23.02.2023 - 08:18

CONTACT US

HOSTS



SPONSORS



