

Automatic parallelisation from high-level abstractions for mesh-based simulations

Gábor Dániel Balogh, István Zoltán Reguly

In this talk, we will introduce the OPS and OP2 abstractions, targeting block-structured and unstructured mesh applications respectively. The goal of these libraries is to provide a high-level programming abstraction embedded in C/C++/Fortran that enables domain scientists to write their scientific codes without having to understand the intricacies of parallel programming. These libraries then use a combination of code generation and back-end libraries to automatically parallelize this code, enabling near-optimal execution on multi- and many-core CPUs, GPUs, in combination with distributed memory message passing, utilizing large supercomputers. These libraries also provide a number of automated features improving productivity, such as parallel I/O, automated checkpointing and recovery. For better performance when using GPUs these libraries also deploy optimizations automatically: such as colouring to avoid race conflicts, staging to improve memory locality, and data layout transformations for better memory accesses, and they are vehicles for further computer science research, such as lazy execution, cache-blocking tiling and others. With this separation of concerns, scientists gain productivity, performance and portability. Results are demonstrated through a variety of applications, including OpenS-BLI, a large-scale Navier-Stokes CFD solver, the TeaLeaf and CloverLeaf proxy applications, the Volna tsunami simulation code, and Hydra, the production turbomachinery simulation code of Rolls-Royce.