

OPTIMIZING AN HPC LBM APPLICATION USING CUDA GRAPHS

Jayesh Badwaik, Milena Veneva, Andreas Herten

Forschungszentrum Jülich



Introduction

With increasing focus on scalability and performance of high performance computing applications, it has become important for the simulation softwares to be able to utilize the underlying hardware as comprehensively to its maximum performance.

waLberla [2] is a multiphysics software framework that has achieved high scalability and performance. It achieves this excellent performance due to architecture specific code generation algorithms [1] combined with efficient communication and parallel data structures like BlockForest.

In this work, we attempt to improve the GPU utilization of an Lattice-Boltzmann Method (LBM) software.

Objectives and Motivation

Objective:

Exploit CUDA Graphs to improve GPU utilization in Walberla for simulations with small block sizes.




Fig. 1: Domain Partitioning in Walberla. The whole domain is divided into blocks of possibly different sizes. The blocks are then distributed among the available processes.

Motivation

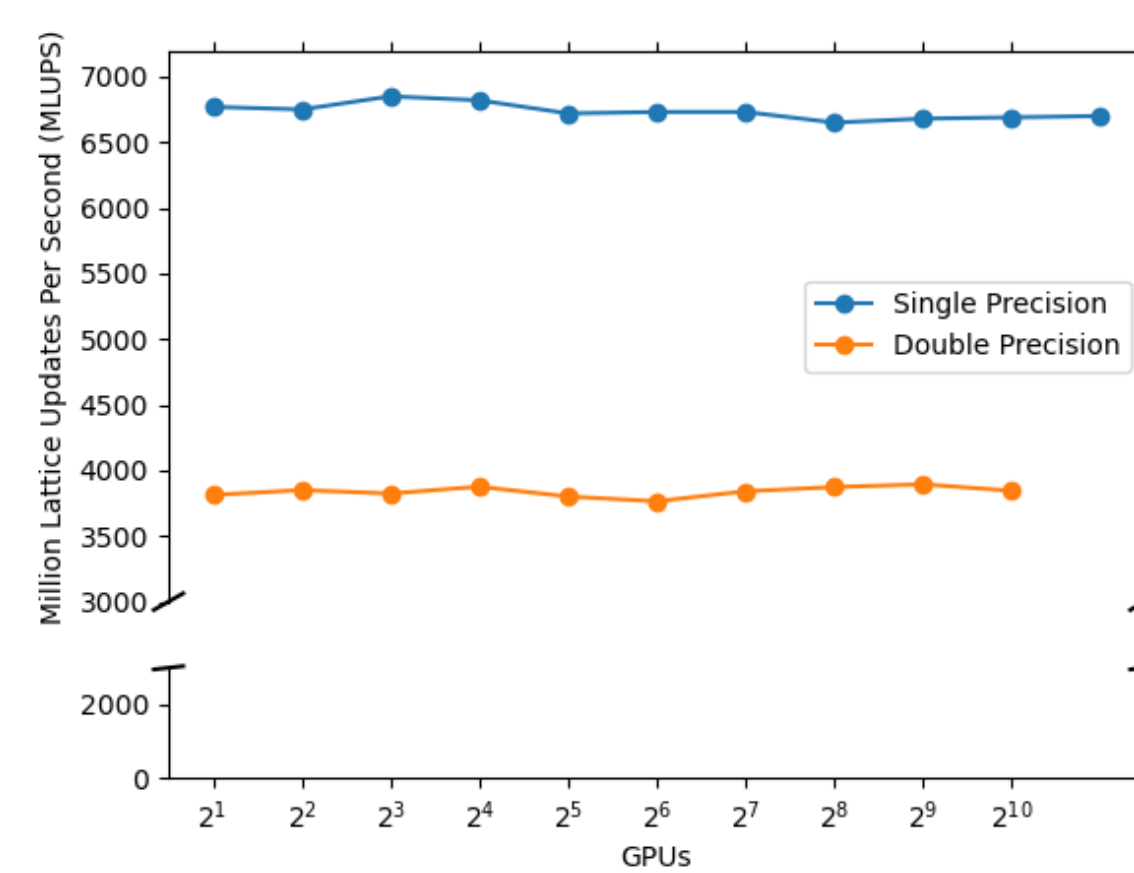


Fig. 2: Weak Scaling for Walberla on Juwels Booster

| NVIDIA GPU | Tesla A100 |
|----------------------|------------|
| Cores | 6912 |
| SMs | 108 |
| Memory | 40 GB HBM2 |
| FP32 [TFLOPS] | 19.5 |
| FP64 [TFLOPS] | 9.7 |
| Capability | 8.0 |
| FREQ [MHz] | 1410 |
| Mem bandwidth [GB/s] | 1555 |

Fig. 3: NVIDIA Tesla A100 specifications..

Walberla achieves high scalability in general. However, for regions with AMR as shown in Figure 1, small block sizes lead to suboptimal utilization of GPUs

- ▶ Small block sizes lead to smaller kernel run times
- ▶ Kernel launch times become comparable to kernel run times
- ▶ Leading to suboptimal utilization of GPUs
- ▶ CUDA Graphs

- Define work as graph rather than single operation
- Requires dependency graphs between kernels

- ▶ Exploit the fixed dependency graphs in WALBERLA

Method

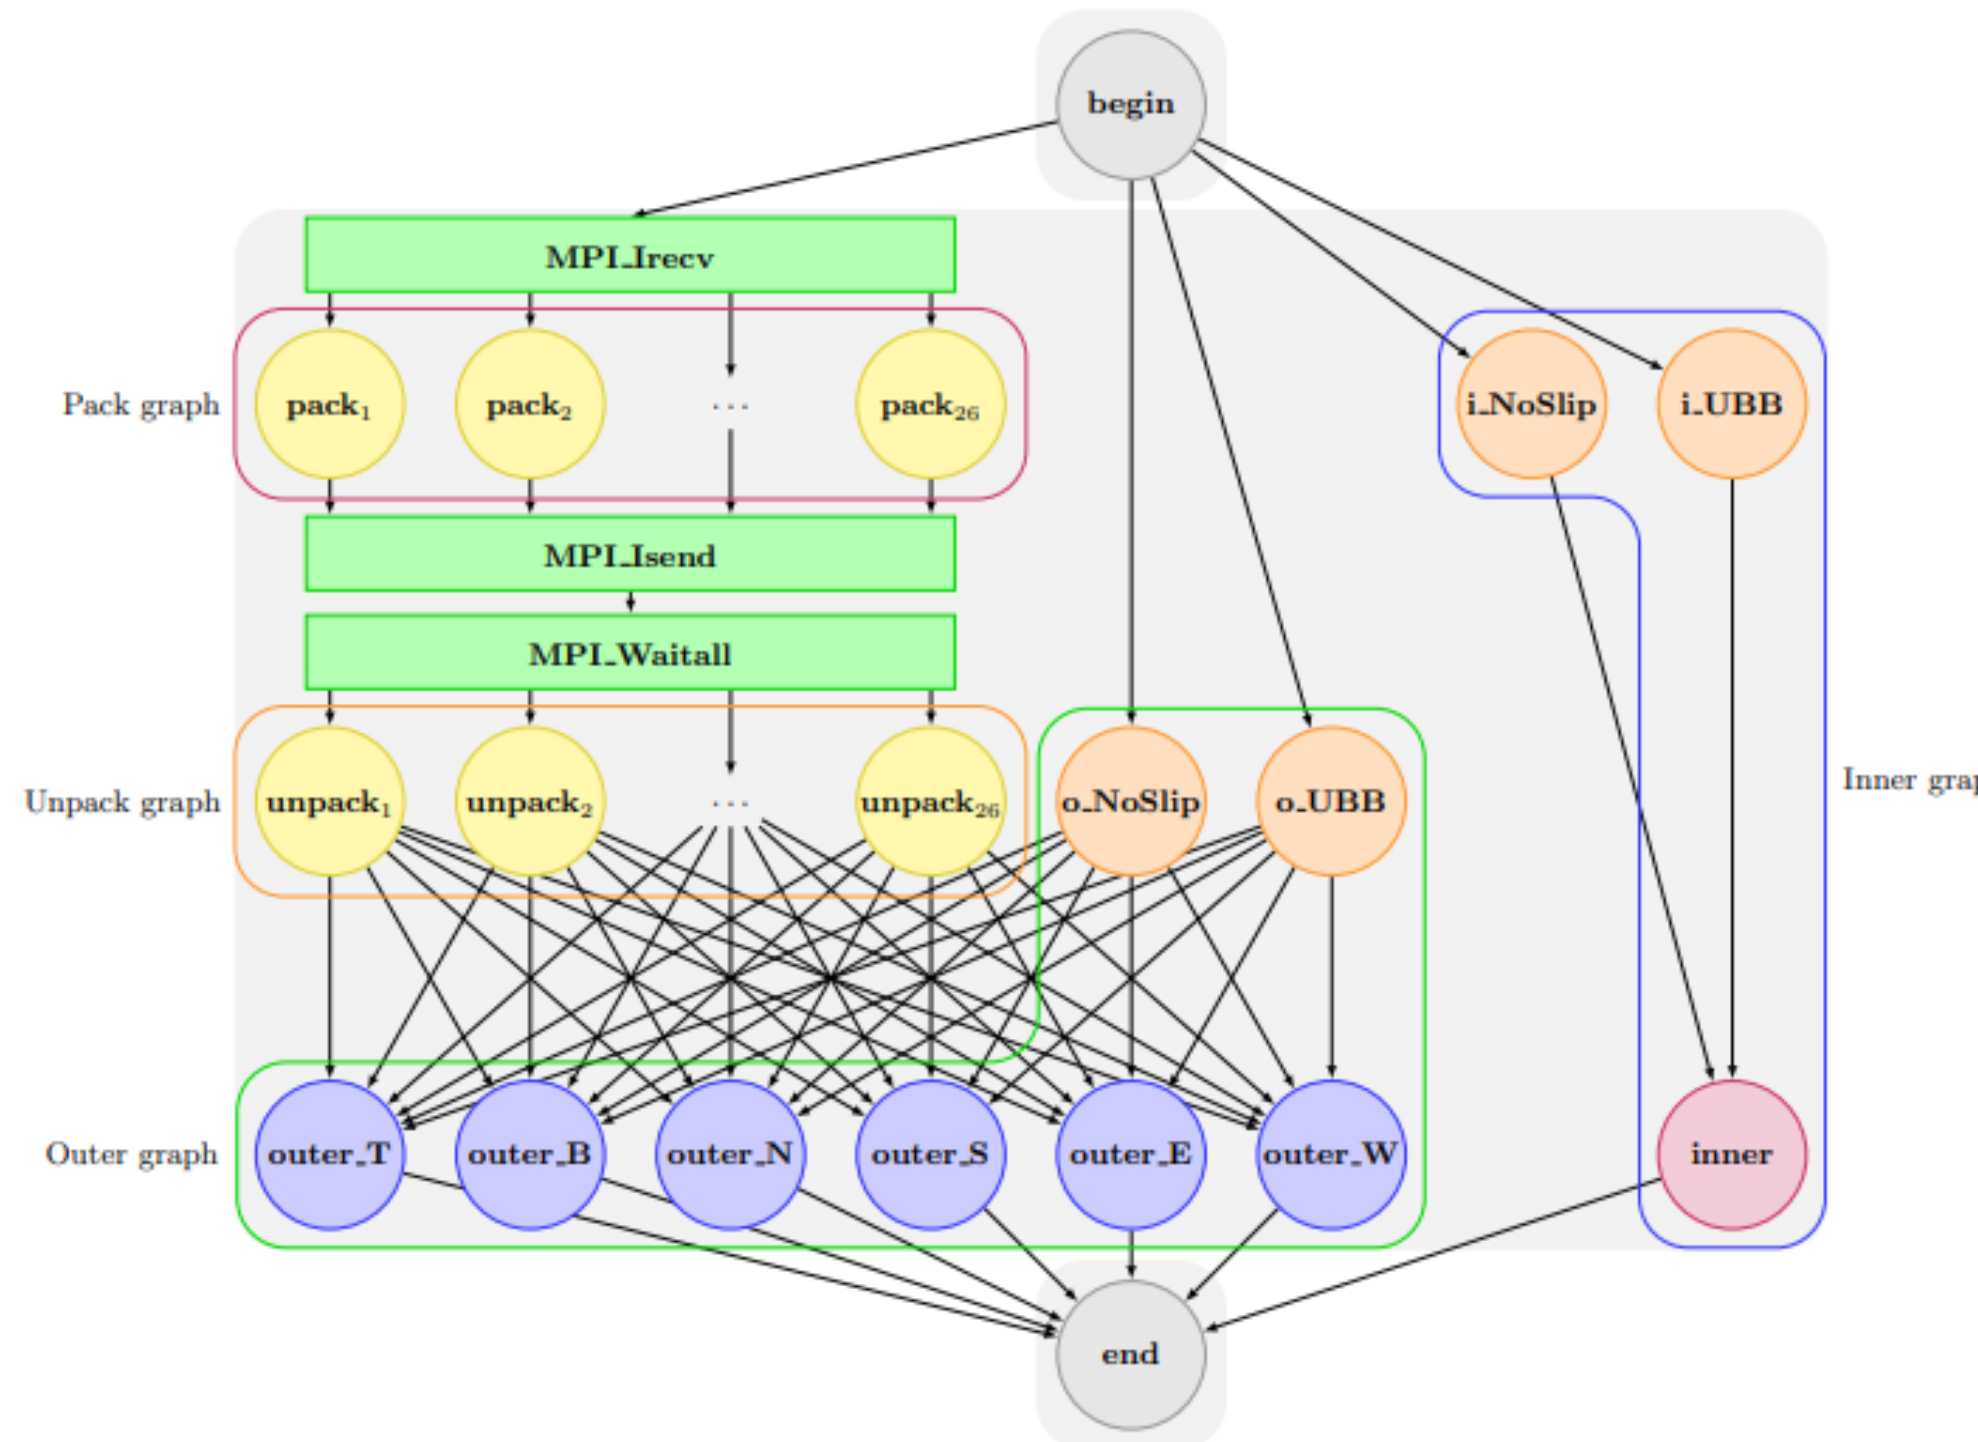


Fig. 4: Dependency Between Kernels in Walberla

(Yellow : Pack/Unpack Kernels, Green: MPI API Calls, Gray: begin/end, Orange: Outer Domain, Others: Inner Domain)

- ▶ Port UniformGridGPU Benchmark to CUDA Graphs using 4 Graphs
 - All pack kernels
 - All unpack kernels
 - Inner Domain Kernels
 - Outer Domain Kernels
- ▶ Exclude CUDA-Aware MPI calls to avoid complexity of the structure of software.
- ▶ Use `cudaStreamSynchronize` to synchronize graph and non-graph components
- ▶ Empty nodes added at beginning and end to help the graph scheduler.

Results

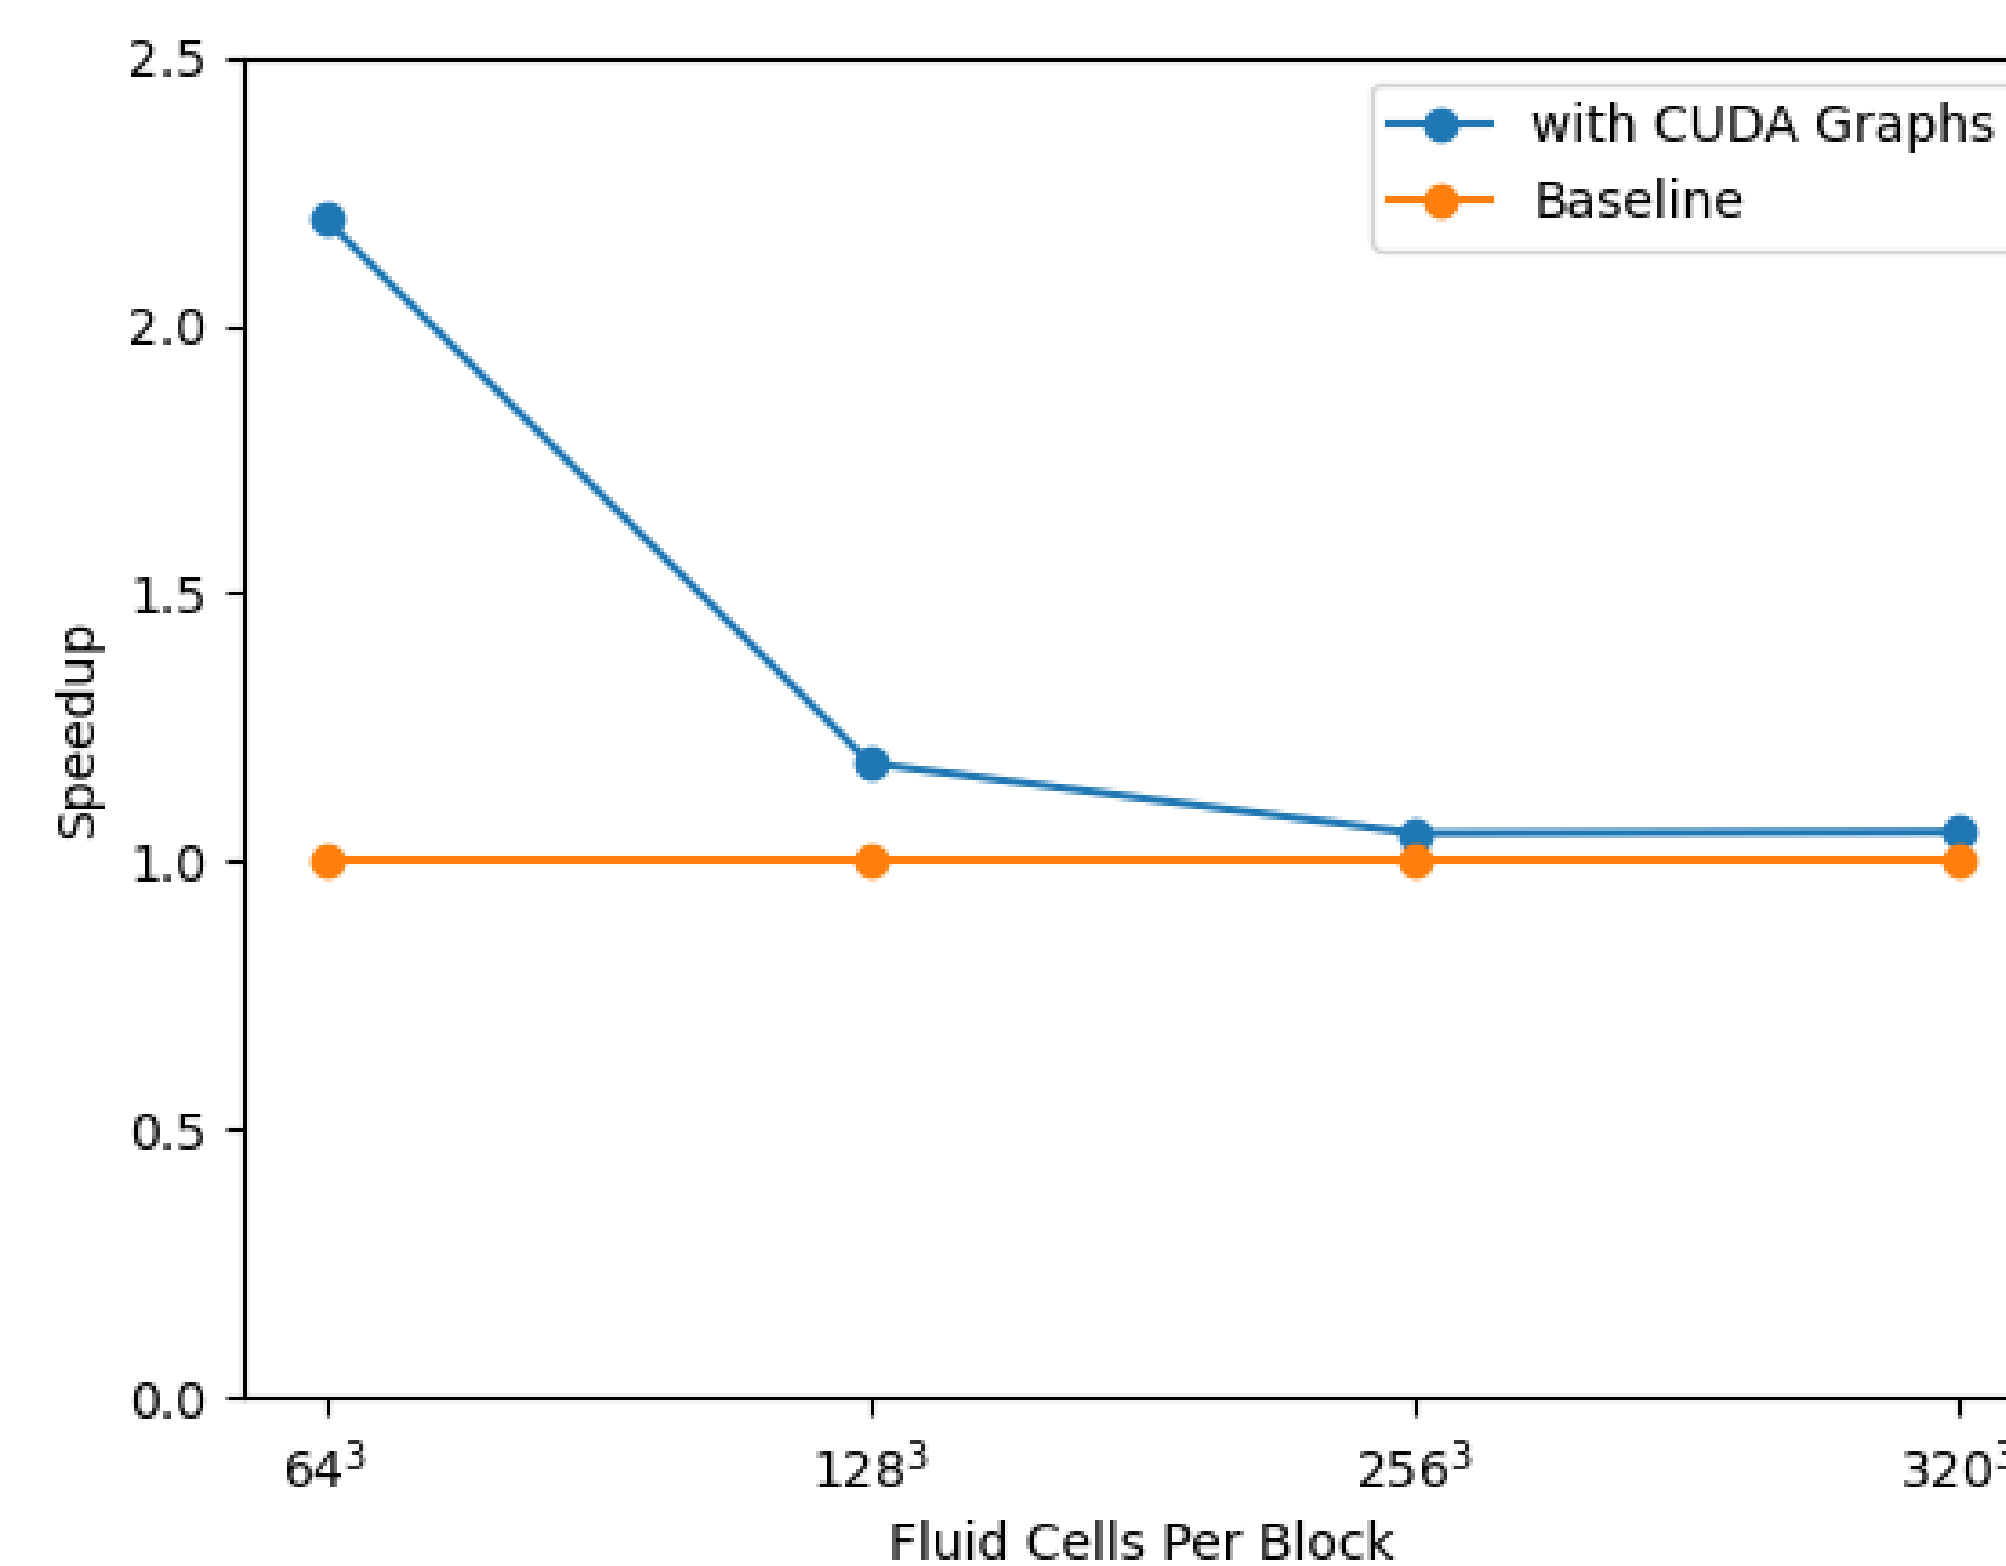


Fig. 5: Results on a Single A100 GPU

An Example Graph Node

```

cudaKernelNodeParams kernelNodeParams = { 0 };
kernelNodeParams.func = (void *)compute_kernel;
kernelNodeParams.gridDim = _grid;
kernelNodeParams.blockDim = _block;
kernelNodeParams.sharedMemBytes = 0;
kernelNodeParams.kernelParams = (void**)kernel_args;
kernelNodeParams.extra = NULL;

```

```

cudaGraphAddKernelNode(&kernelNode, graph, nodeDependencies.data(),
nodeDependencies.size(), &kernelNodeParams)

```

Conclusions

- ▶ More than 100% improvement in performance for smaller block sizes.
- ▶ Between 6% and 23% improvement for large block sizes.
- ▶ Traditional MPI Calls not integrable into CUDA Task Graphs
- ▶ Possibility of further performance benefits by using communication frameworks which can be integrated into graphs (NCCL instead of MPI).
- ▶ Promise of performance improvement while using using Task Graph approach to GPU kernels in applications with smaller kernels.

Next Steps

- ▶ Port More Complex Use Cases to CUDA Task Graphs
- ▶ Port MPI Communications to use NCCL

References

[1] Martin Bauer, Harald Köstler, and Ulrich Rüde. "lbmpy: Automatic code generation for efficient parallel lattice Boltzmann methods". In: *Journal of Computational Science* 49, 101269 (Feb. 2021). DOI: 10.1016/j.jocs.2020.101269.

[2] Martin Bauer et al. "waLberla: A block-structured high-performance framework for multiphysics simulations". In: *Computers and Mathematics with Applications* 81 (2021). Development and Application of Open-source Software for Problems with Numerical PDEs, pp. 478–501. ISSN: 0898-1221. DOI: <https://doi.org/10.1016/j.camwa.2020.01.007>. URL: <https://www.sciencedirect.com/science/article/pii/S0898122120300146>.

Acknowledgements

- ▶ This project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 956000
- ▶ The JU receives support from the European Union's Horizon 2020 research and innovation programme, and the French National Research Agency, the Federal Ministry of Education and Research of Germany, the Ministry of Education, Youth and Sports of the Czech Republic.
- ▶ The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

