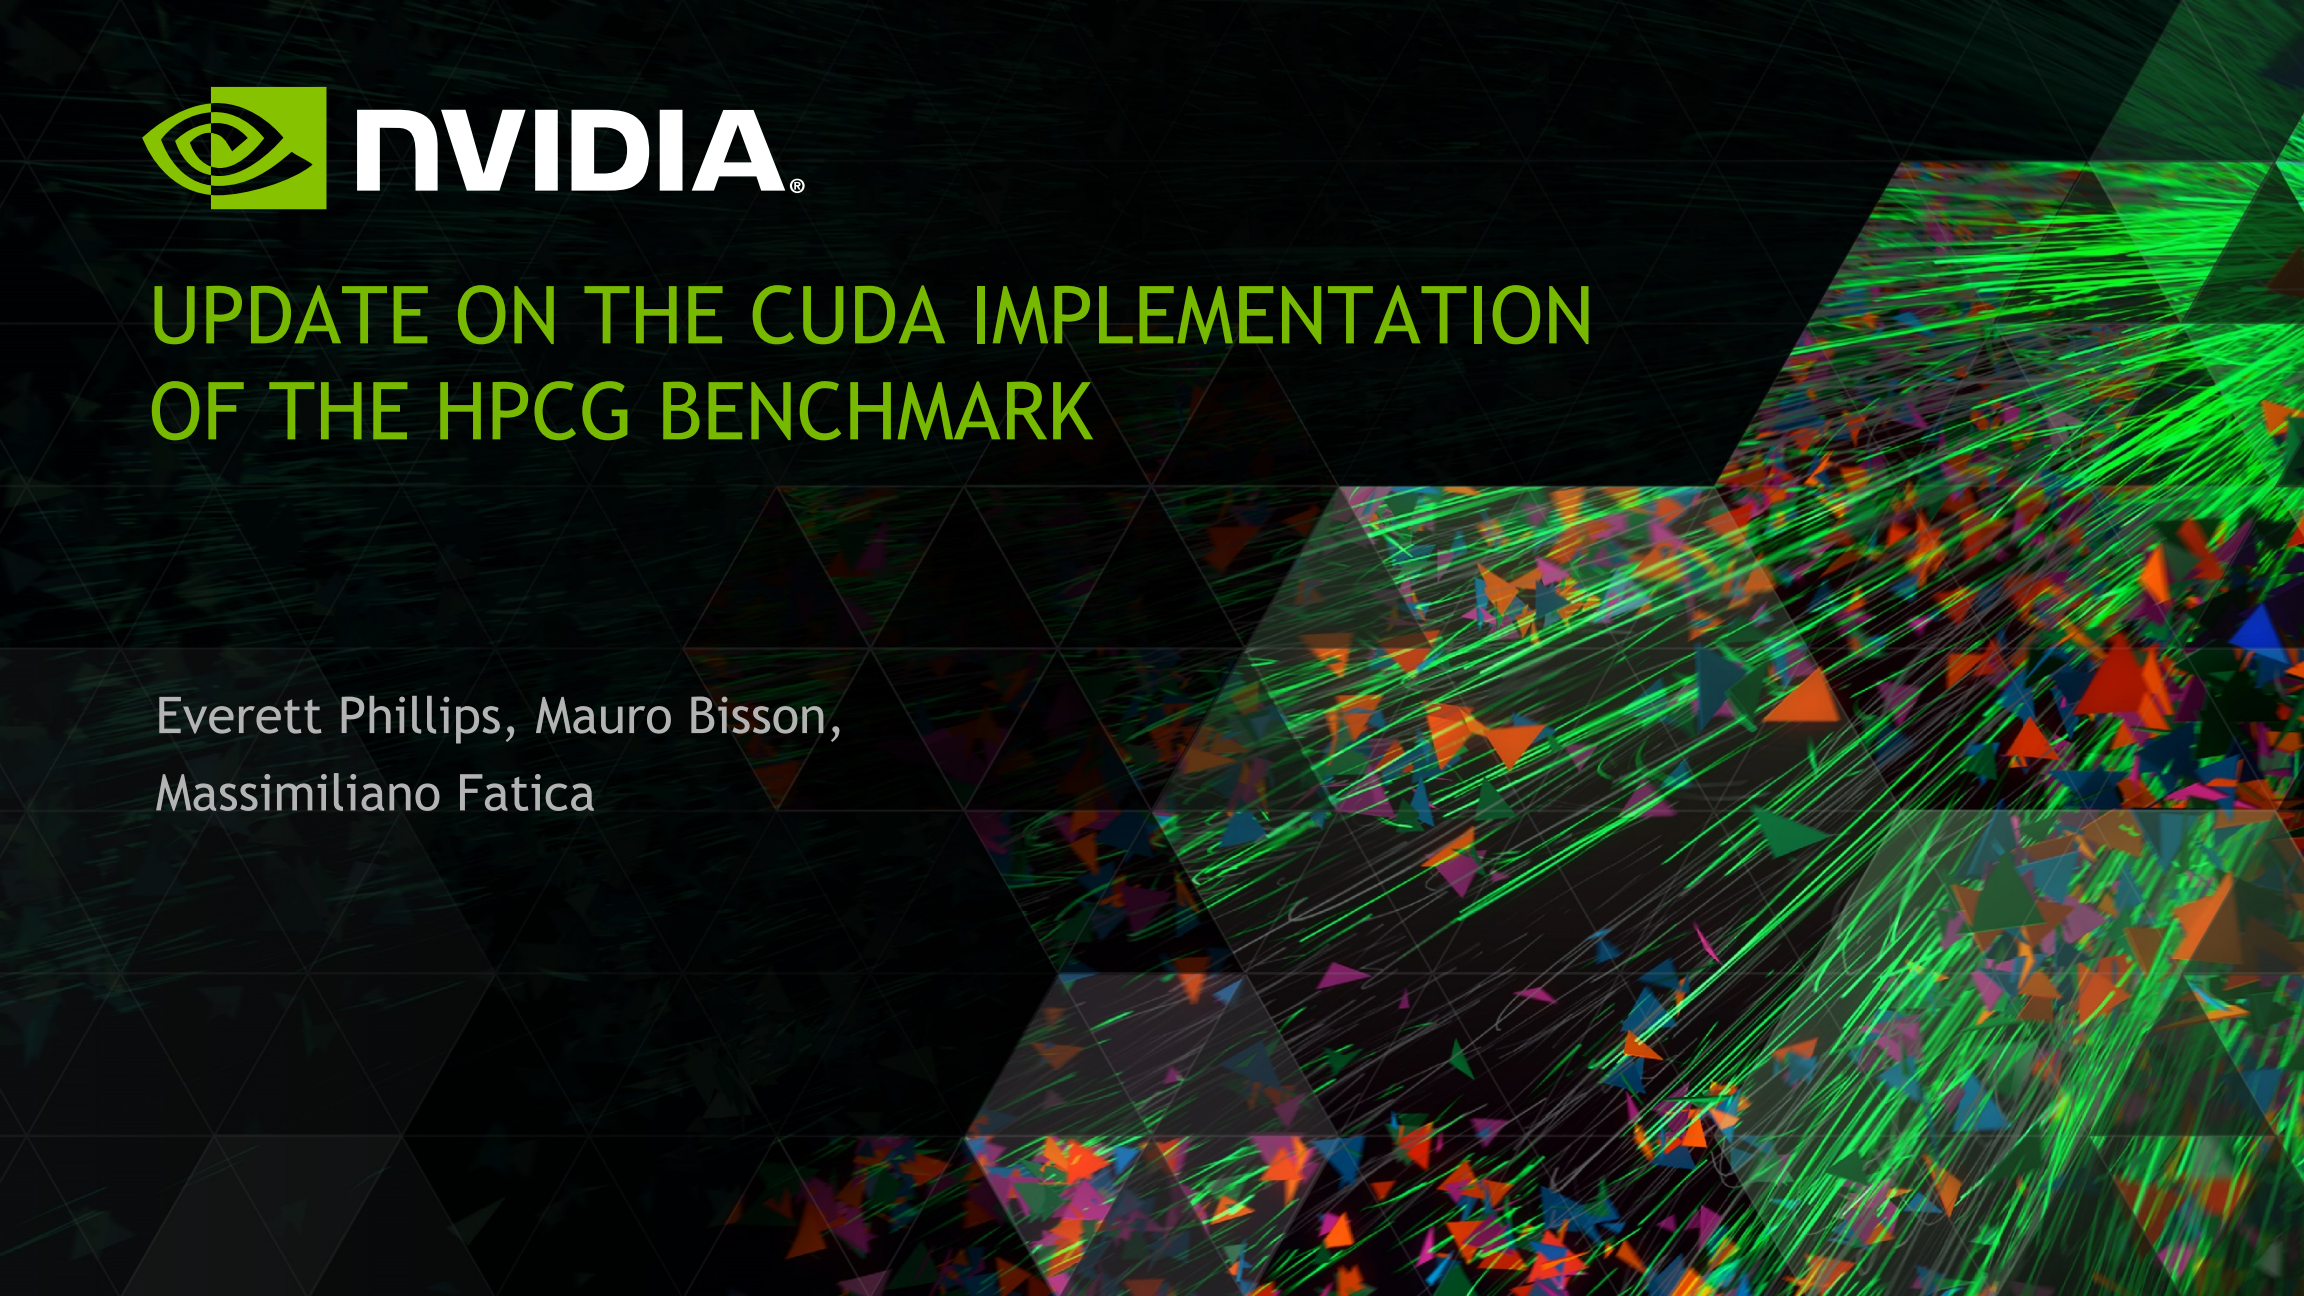

UPDATE ON THE CUDA IMPLEMENTATION OF THE HPCG BENCHMARK

Everett Phillips, Mauro Bisson,
Massimiliano Fatica

HYBRID VERSION

Problem Optimization

- Matrices are analyzed, reordered and split between CPU and GPU

  -Split ratio between 0.0 (full GPU) and to 1.0 (full CPU)

- Data structures are rearranged to best fit the GPU/CPU architecture

  -CPU matrices in sliced-ELLPACK format (column-major)

  -GPU matrices in ELLPACK format (column-major)

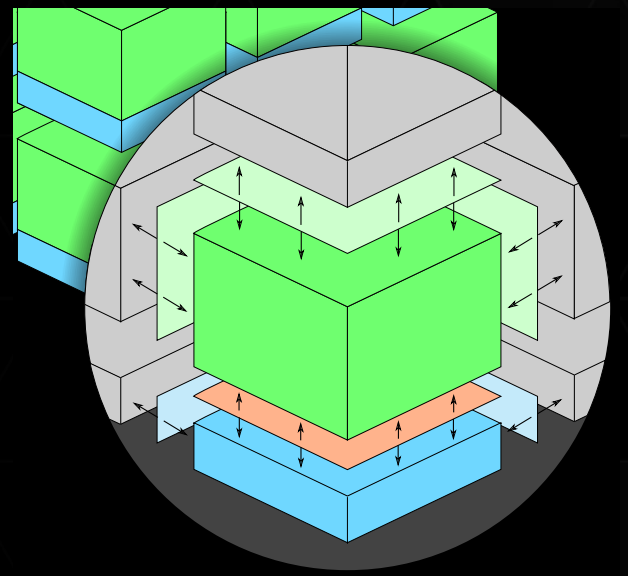

Optimized Run

- SpMV, SymGS, DotProd performed simultaneously on CPU and GPU

  -OpenMP on CPU

  -CUDA on GPU

- External AND internal halos exchnages

  -Inter-node via MPI
  -Intra-node via cudaMemcpy[Async]

HYBRID VERSION

- Gather GPU boundary to send_buf
- Copy send_buf to CPU
- Exchange CPU/GPU frontiers
- Launch GPU SPMV bulk kernel
- Execute CPU SPMV bulk kernel
- MPI_Send / MPI_Recv
- Copy recv_buf to GPU
- Launch GPU SPMV boundary kernel
- Execute CPU SPMV boundary kernel

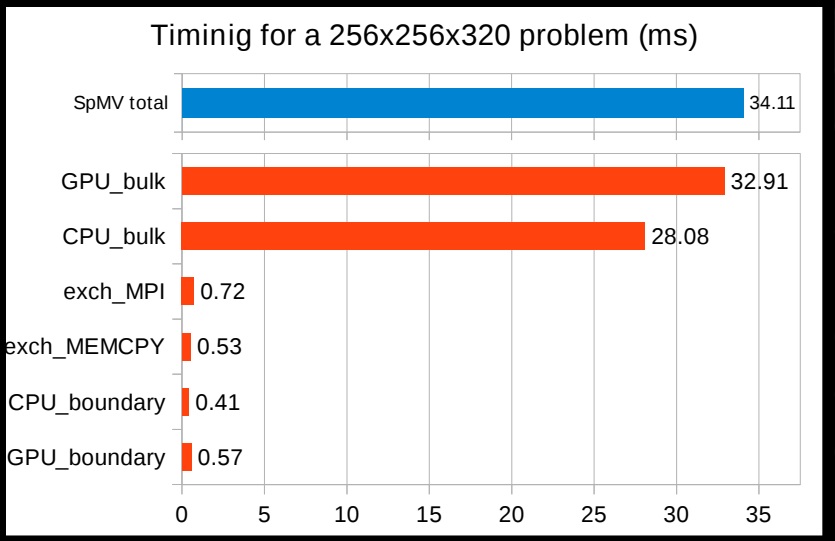

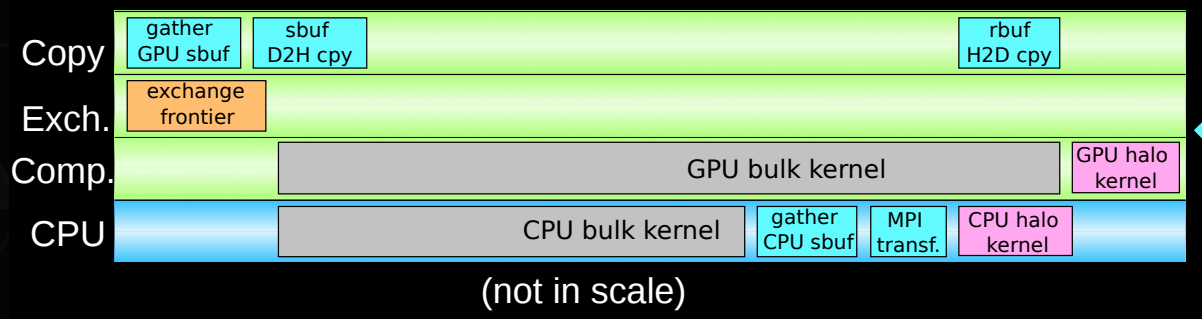

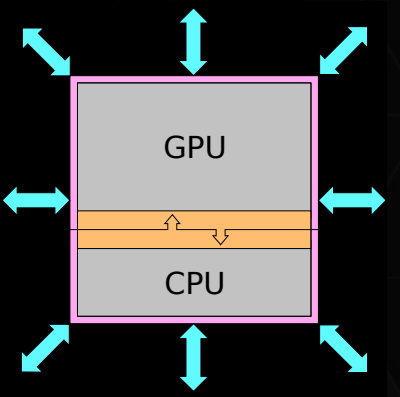

SINGLE NODE HYBRID VERSION

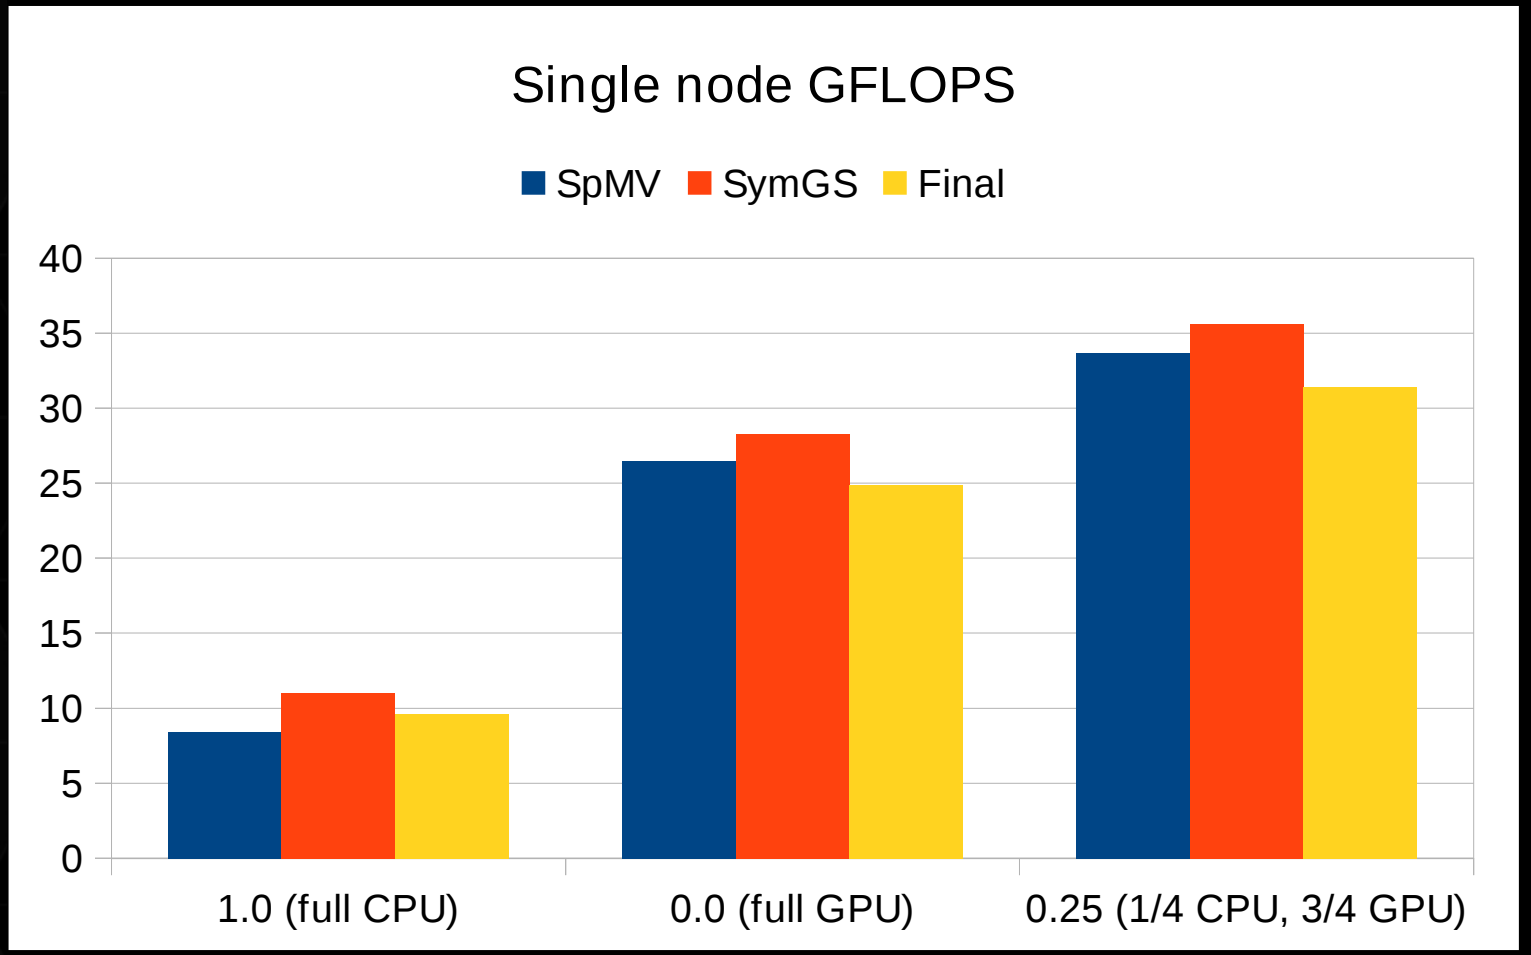

RESULTS ON PIZ DAINT (CSCS)

- GPU only version
  - 97 TFLOPS (5265 K20X)
- Hybrid version
  - 124.5 TFLOPS (5148 K20x)
- We have a newer version with better OpenMP performance waiting for a full machine run.

HPCG 3.0

- New optimized version available from the hpcg web site
- Additional optimizations

Total Time: 6.107714e+01 sec

Setup Overhead: 3.08%

Optimization Overhead: 0.46%

Convergence Overhead: 7.41%

2x2x2 process grid

256x256x256 local domain

SpMV = 195.7 GF (1232.4 GB/s) 24.5 GF_per (154.1 GB/s_per)

SymGS = 231.3 GF (1785.0 GB/s) 28.9 GF_per (223.1 GB/s_per)

total = 221.7 GF (1681.2 GB/s) 27.7 GF_per (210.2 GB/s_per)

final = 198.1 GF (1502.0 GB/s) 24.8 GF_per (187.7 GB/s_per)

HPCG 3.0