




# Grid on QPACE 4

Nils Meyer, Peter Georg, Stefan Solbrig, Tilo Wettig  
Department of Physics, University of Regensburg, Germany



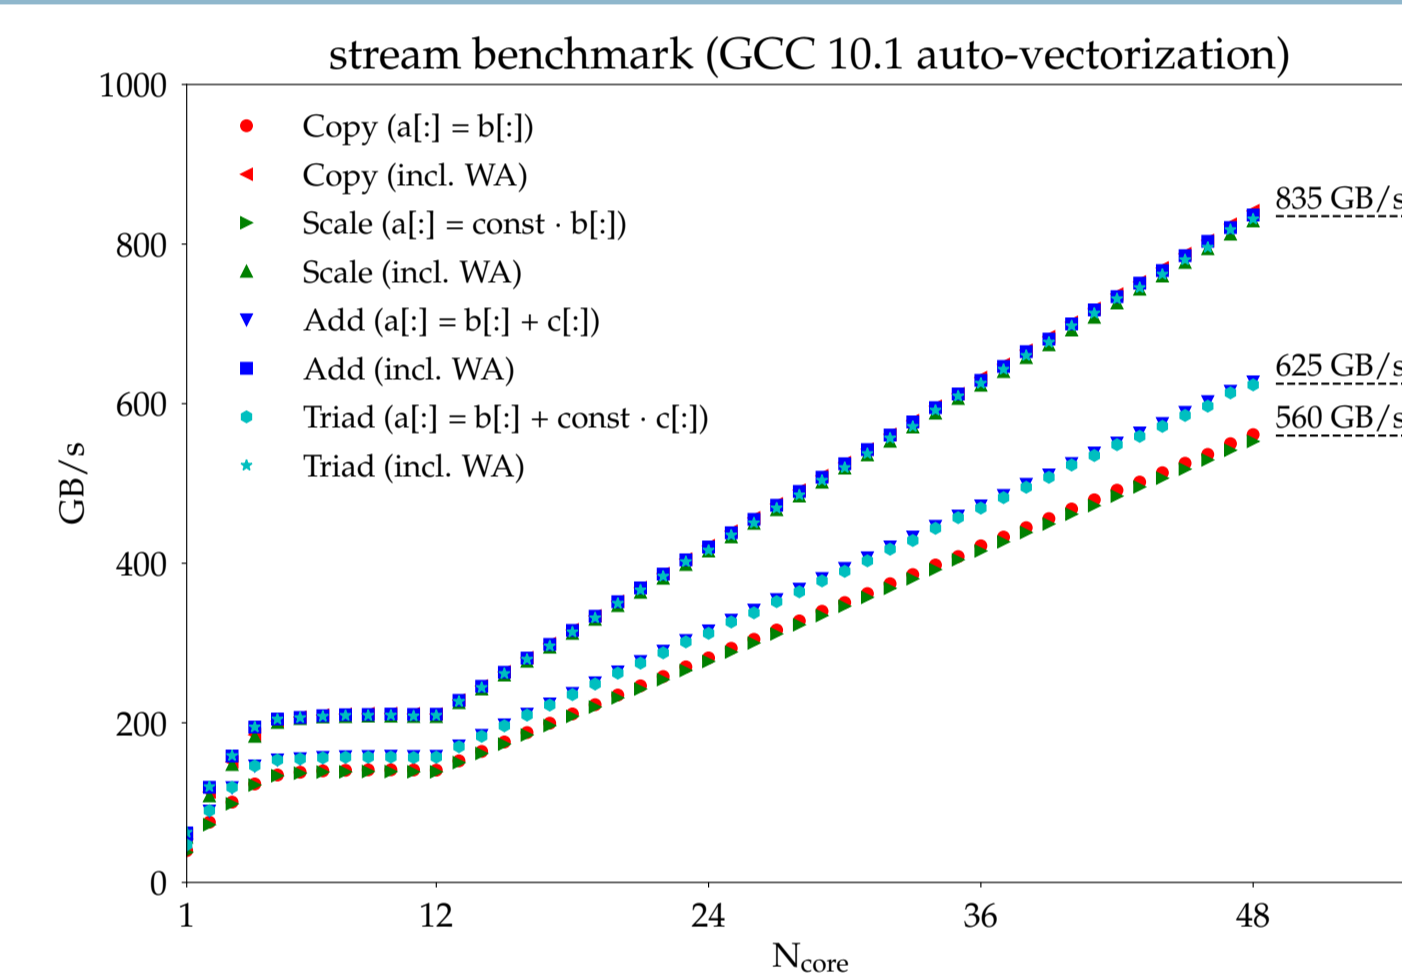
## QPACE 4

- ▶ Latest member of QCD PArallel Compute Engine (QPACE) series
  - ▷ Fujitsu PRIMEHPC FX700 series
  - ▷ Deployed June 2020 at Regensburg University
- ▶ 64 Fujitsu A64FX CPUs (48 cores each, 1.8 GHz)
  - ▷ 512-bit Arm Scalable Vector Extension (SVE)
  - ▷ 177/354 TFlop/s peak in double/single precision (DP/SP)
  - ▷ 2048 GB HBM2 memory total
  - ▷ InfiniBand EDR interconnect (100 Gbit/s)
- ▶ Open-source software stack
  - ▷ CentOS Stream 8, GCC 10.1, OpenMPI 4.0
  - ▷ GlusterFS parallel filesystem
  - ▷ Grid Lattice QCD framework [1]
  - ▷ Grid Python Toolkit (GPT) [2]



## Stream benchmark

- ▶ Standard benchmark for memory throughput evaluation
- ▶  $N_{\text{core}} > 12$ : data throughput scales with number of cores in use
- ▶ Benchmark data throughput: up to 625 GB/s
- ▶ Caches implement write-back policy
  - ▷ Write Allocation (WA): cache block load from memory on write miss
  - ▷ WA causes extra traffic and thus reduces effective memory throughput
  - ▷ Including WA traffic, data throughput is up to 835 GB/s




## Grid port to A64FX (512-bit SVE)

- ▶ Arm C Language Extensions (ACLE) provide access to SVE vector types and SVE instructions in C/C++
- ▶ We use ACLE to implement Grid's lower-level functions
  - ▷ Alternating re/im parts layout (RIRI)
  - ▷ Hardware support for processing complex numbers
- ▶ Hand-optimized Wilson Dslash and Domain Wall kernels using ACLE and RIRI layout
- ▶ Available in upstream Grid develop branch (`configure --enable-simd=A64FX`)

## Grid benchmark: SU(3) matrix multiplication

- ▶ Independent SU(3) matrix multiplication  $z = x \times y$  on each lattice site
- ▶ Well suited for compiler testing
- ▶ GCC 10.1: up to  $\sim 300$  GFlop/s DP (800 GB/s)
  - ▷ GCC 10.2 and 11.1 achieve same performance
- ▶ clang/LLVM compilers underperform (incl. clang/LLVM 12.0)


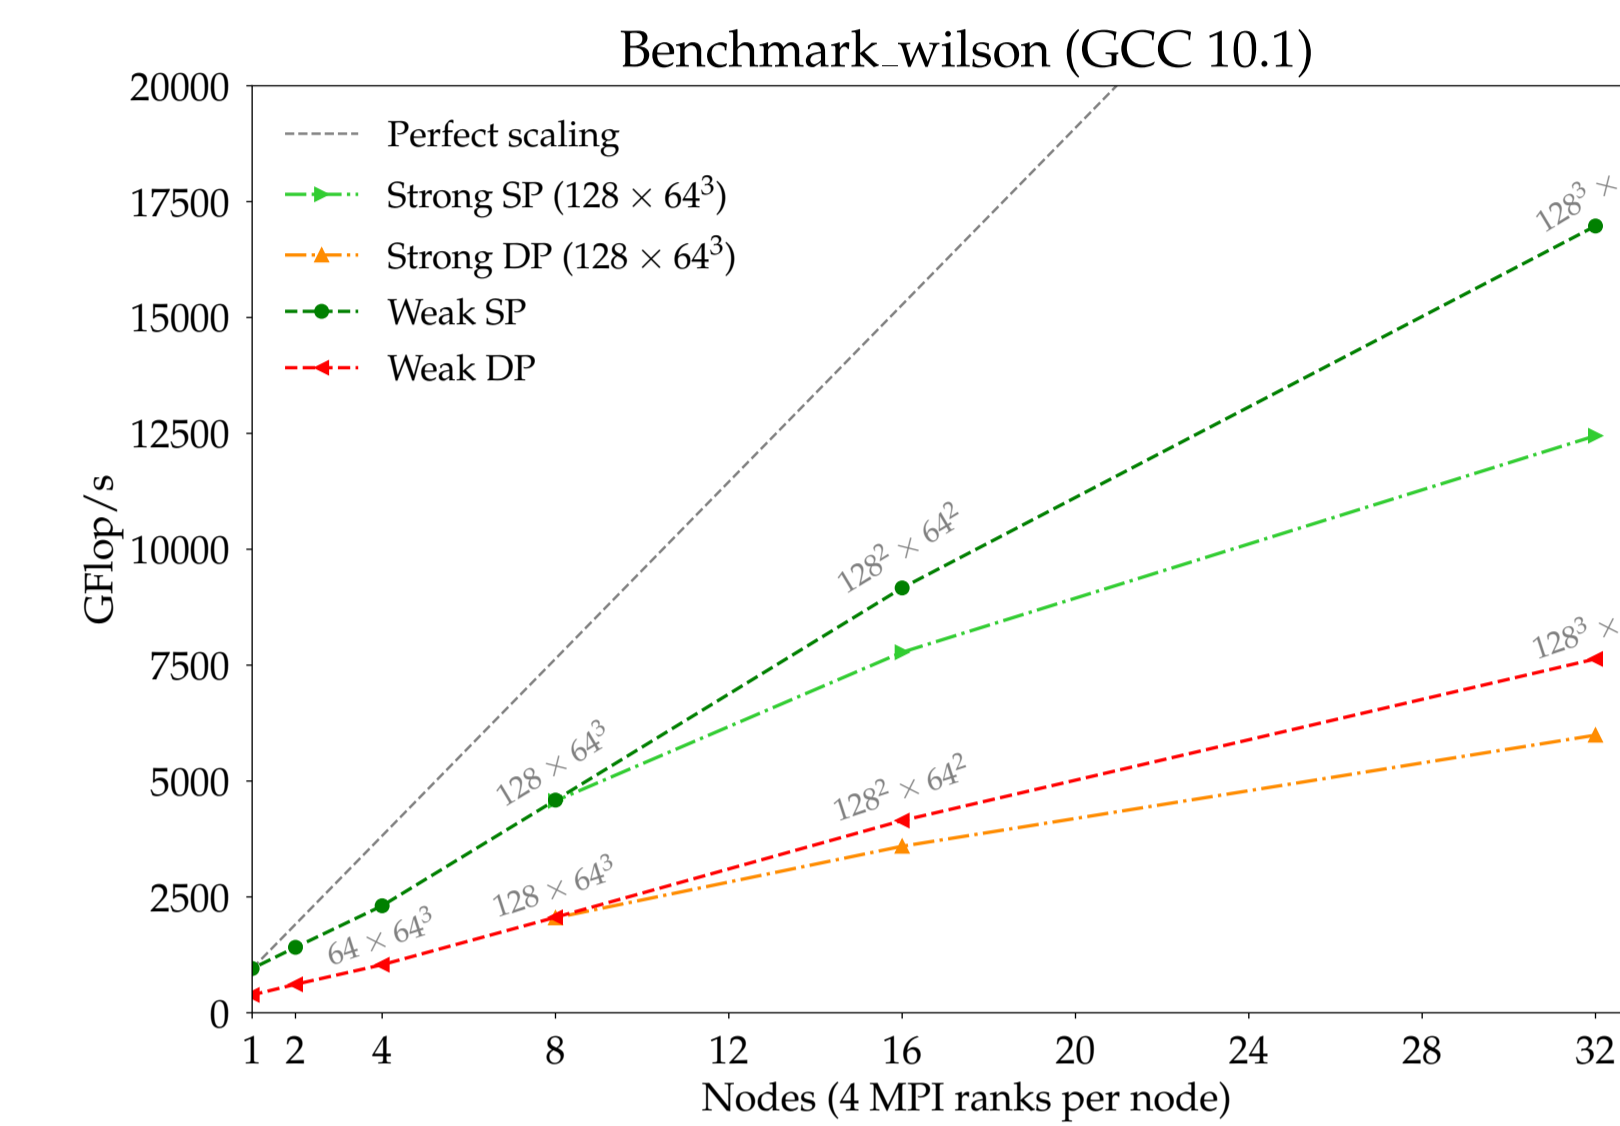


## Wilson Dslash and Domain Wall kernels

- ▶ Performance-relevant part of Domain Wall Dirac operator, same source code as Wilson Dslash ( $L_s = 1$ )


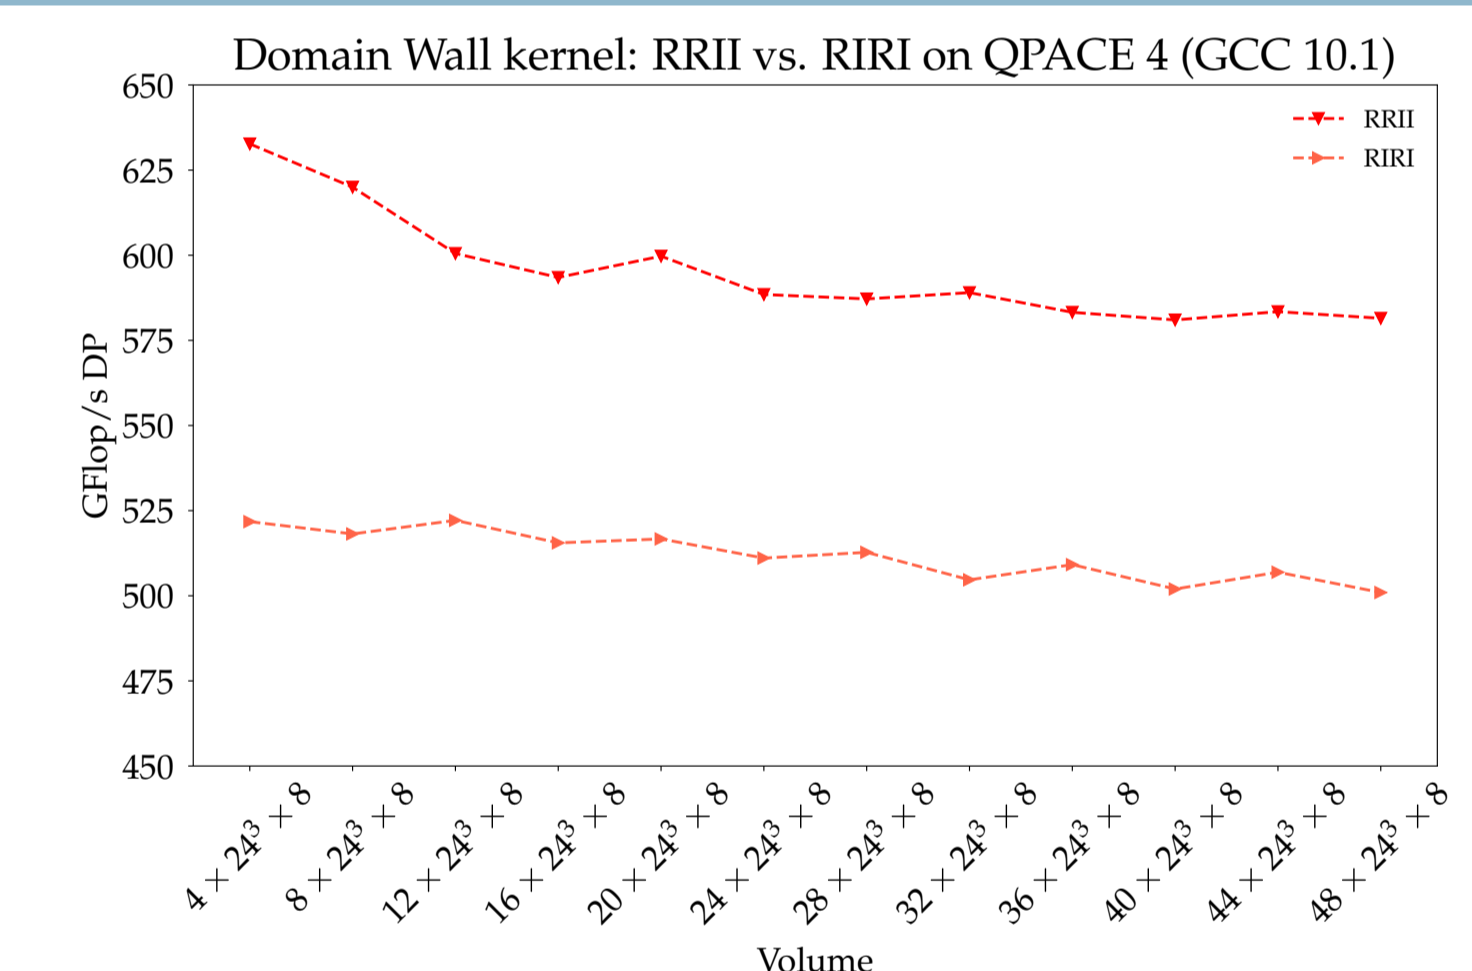
$$\psi'(n, s)_{\alpha\alpha} = (D\psi)(n, s)_{\alpha\alpha} = \sum_{\mu=1}^4 \sum_{\beta=1}^4 \sum_{b=1}^3 \{U_{\mu}(n)_{ab}(1 + \gamma_{\mu})_{\alpha\beta}\psi(n + \hat{\mu}, s)_{\beta b} + U_{\mu}^{\dagger}(n - \hat{\mu})_{ab}(1 - \gamma_{\mu})_{\alpha\beta}\psi(n - \hat{\mu}, s)_{\beta b}\}$$

- ▶ Specialization for A64FX (`--dslash-asm`): manual instruction scheduling and software prefetching using ACLE
- ▶ GCC 10.1 and 10.2 achieve best performance, outperforming other compilers (incl. clang/LLVM 12.0 and GCC 11.1)



## Alternative data layout of complex numbers

- ▶ Collaboration with Erlangen University (FAU), Germany
  - ▷ We extended GridBench [3] to support the A64FX [4]
  - ▷ Performance study of Domain Wall kernel on Fugaku (single node): alternative data layout separating re/im parts (RRII) outperforms RIRI
  - ▷ Details will be published in [5], preprint available
- ▶ QPACE 4 single node: RRII outperforms RIRI up to  $\sim 20\%$
- ▶ RRII currently not supported by Grid (future work)



## Summary and outlook

- ▶ Fujitsu A64FX achieves outstanding performance for Lattice QCD applications
- ▶ GCC 10.1 and 10.2 achieve best overall performance
- ▶ Other compilers underperform (incl. clang/LLVM 12.0 and GCC 11.1)
- ▶ Alternative layout of complex numbers is beneficial on A64FX, but yet to be integrated into Grid (future work)

## References

- [1] P. Boyle et al., *Proceedings of LATTICE 15* (2016) 023 [arxiv:1512.03487]
- [2] C. Lehner et al., "Grid Python Toolkit (GPT)." <https://github.com/lehner/gpt>
- [3] P. Boyle et al., "GridBench – Single CPU benchmarks cutting down Grid." <https://github.com/paboyle/GridBench>
- [4] N. Meyer et al., "GridBench – AVX512 and A64FX extensions." <https://github.com/nmeyer-ur/GridBench/tree/intrinsics>
- [5] C. Alappat et al., *Concurrency and Computation: Practice and Experience (PMBS special edition)* (2021) [arxiv:2103.03013]