Reproducibility and Performance on modern hardware architectures

M. Wiesenberger[a], L. Einkemmer[b;c], A. Gutierrez-Milla[d], Xavier Sáez[d], M. Held[e], and R. Iakymchuk[f]

a Department of Physics, Technical University of Denmark (DTU), Denmark,
b Department of Mathematics, University of Tübingen, Germany,
c Department of Mathematics, Universität Innsbruck, Austria
d Barcelona Supercomputing Center, Barcelona, Spain
e Institute for Ion Physics and Applied Physics, Universität Innsbruck, Austria
f Department of Computational Science and Technology, Royal Institue of Technology (KTH), Sweden

Feltor library

https://feltor-dev.github.io

FELTOR (Full-F ELectromagnetic code in TORoidal geometry)

6	diagnostics	Physical projects / User zone
5	applications	
4	Advanced numerical schemes	dg library (discontinuous galerkin)/ Developer zone
3	Topology and Geometry	
2	Basic numerical algorithms	
1	Vector and Matrix operations	

- both a **numerical library** and a scientific software package.
- two- and three-dimensional drift- and gyrofluid **simulations**
- **discontinuous Galerkin** methods on structured grids.
- **platform independent code** from laptop CPUs to hybrid CPU+GPU distributed memory systems.

Reproducibility

$$\partial_t n + \{\phi, n\} = \alpha\left(\widetilde{\phi} - \widetilde{\ln(n)}\right),$$
$$\partial_t N + \left\{\phi - (\nabla\phi)^2/2, N\right\} = 0,$$
$$\nabla\cdot(N\nabla_\perp\phi) = n - N,$$

- same exact input
- parallel hardware

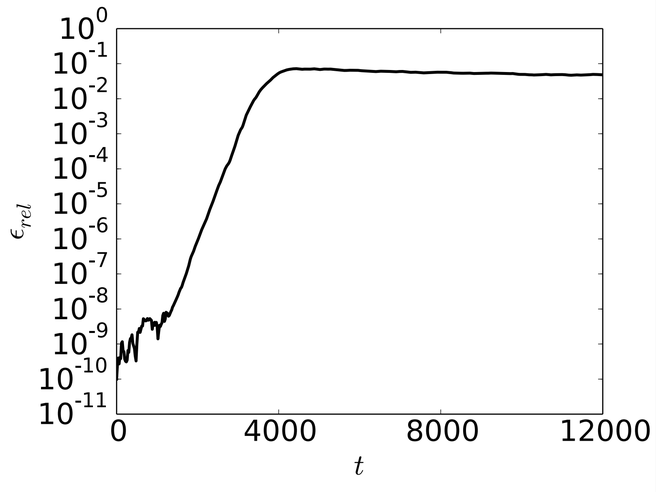

$$(-1 \oplus 1) \oplus 2^{-53} \neq -1 \oplus (1 \oplus 2^{-53})$$

with / without superaccumulators [1]
(a very long floating point number)

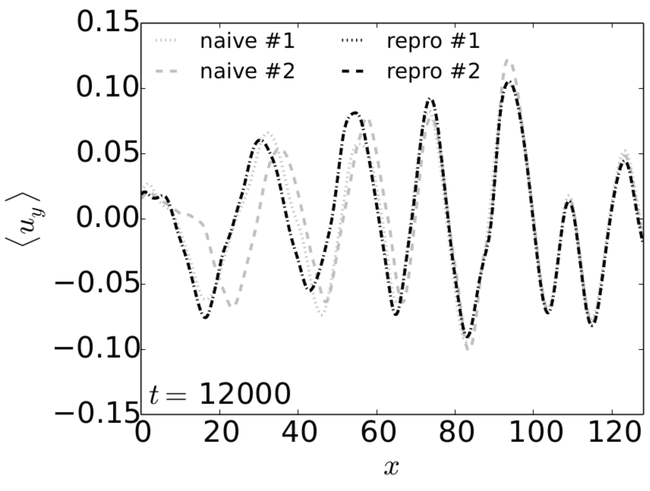

Condition?
Accuracy?
Convergence?

INVARIANTS ARE IMPORTANT

Modern hardware architectures

a*x+b*y

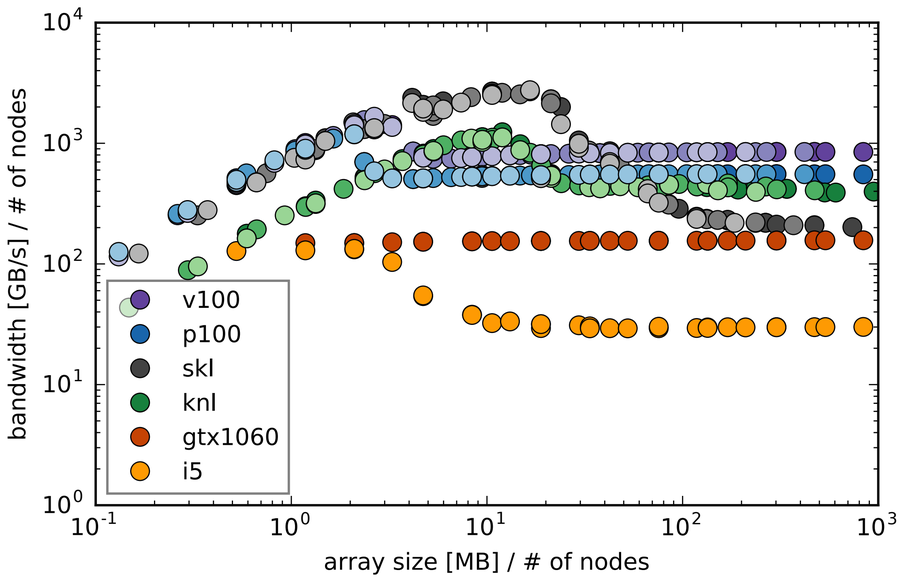

reproducible scalar product

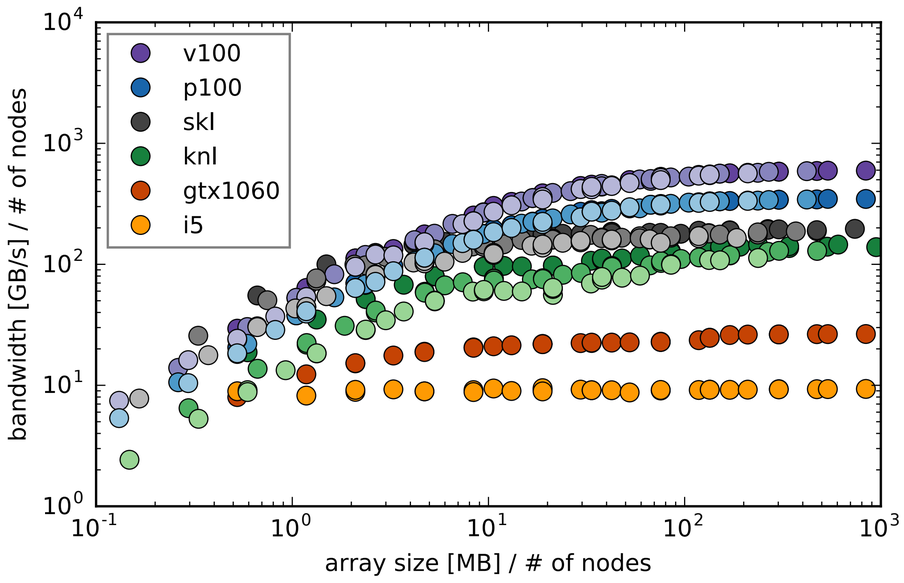

Performance prediction

$$t = T_{\text{lat}}(n) + \frac{mS}{nB},$$

Conjugate gradient iteration

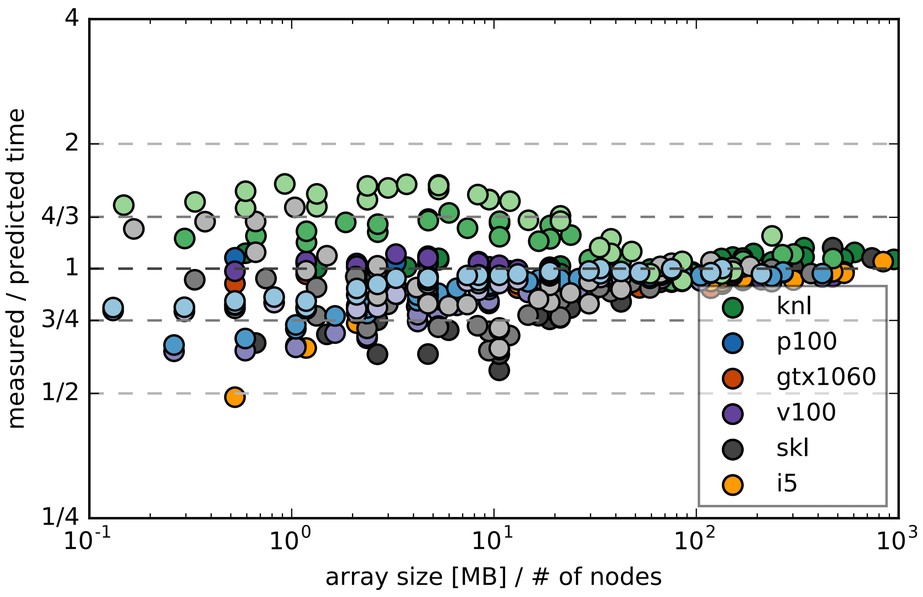

measured average bandwidth and latencies

	B(P=2) [GB/s]	B(P=3) [GB/s]	B(P=4) [GB/s]	B(P=5) [GB/s]	$T_{lat}(1)$ [μs]	$T_{lat}(4)$ [μs]
i5	26 ± 02	27 ± 02	26 ± 01	23 ± 02	01 ± 01	n/a
gtx1060	116 ± 01	108 ± 01	94 ± 09	85 ± 12	09 ± 01	n/a
skl	194 ± 20	183 ± 09	153 ± 15	147 ± 07	14 ± 02	19 ± 02
knl	281 ± 13	232 ± 24	188 ± 20	160 ± 18	13 ± 01	42 ± 02
p100	377 ± 02	333 ± 04	297 ± 02	259 ± 17	06 ± 01	39 ± 01
v100	808 ± 09	763 ± 11	727 ± 10	653 ± 35	06 ± 01	39 ± 01

https://github.com/mwiesenberger/performance

[1] S. Collange, D. Defour, S. Graillat, and R. Iakymchuk.
Parallel Computing, 49:83–97, 2015.