NOISE-DRIVEN CLUSTER-LEVEL PERFORMANCE MODELLING AND ENGINEERING

Ayesha Afzal <ayesha.afzal@fau.de >, Georg Hager (advisor), Gerhard Wellein (advisor)

Department of Computer Science | Erlangen National High Performance Computing Center | Friedrich-Alexander University Erlangen-Nürnberg, Germany

Previous related work
- ✓ characterization and sources of noise
- ✓ noise impact on code performance
- ✓ explicit techniques of noise mitigation

Model For Runtime Of Distributed-Memory Parallel Applications = **Execution Runtime Model** + **Communication Runtime Model** ± **Noise-Driven Runtime Model**

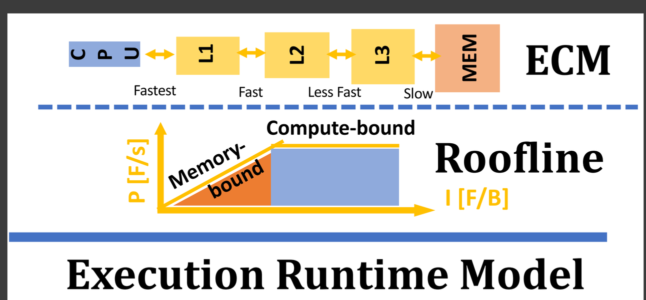

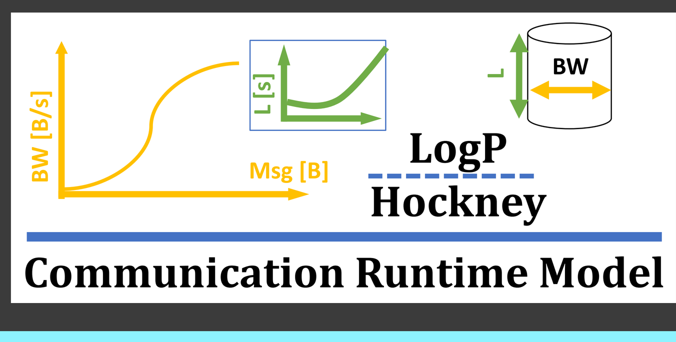

Problem: noise can cause acceleration, speed down, or neither

Our vision: white-box first-principle performance modelling of extreme-scale applications incorporating the influence of noise

1. COMMUNICATION-DRIVEN WAVE PROPAGATION (NO FINE-GRAINED NOISE)

Afzal et al.

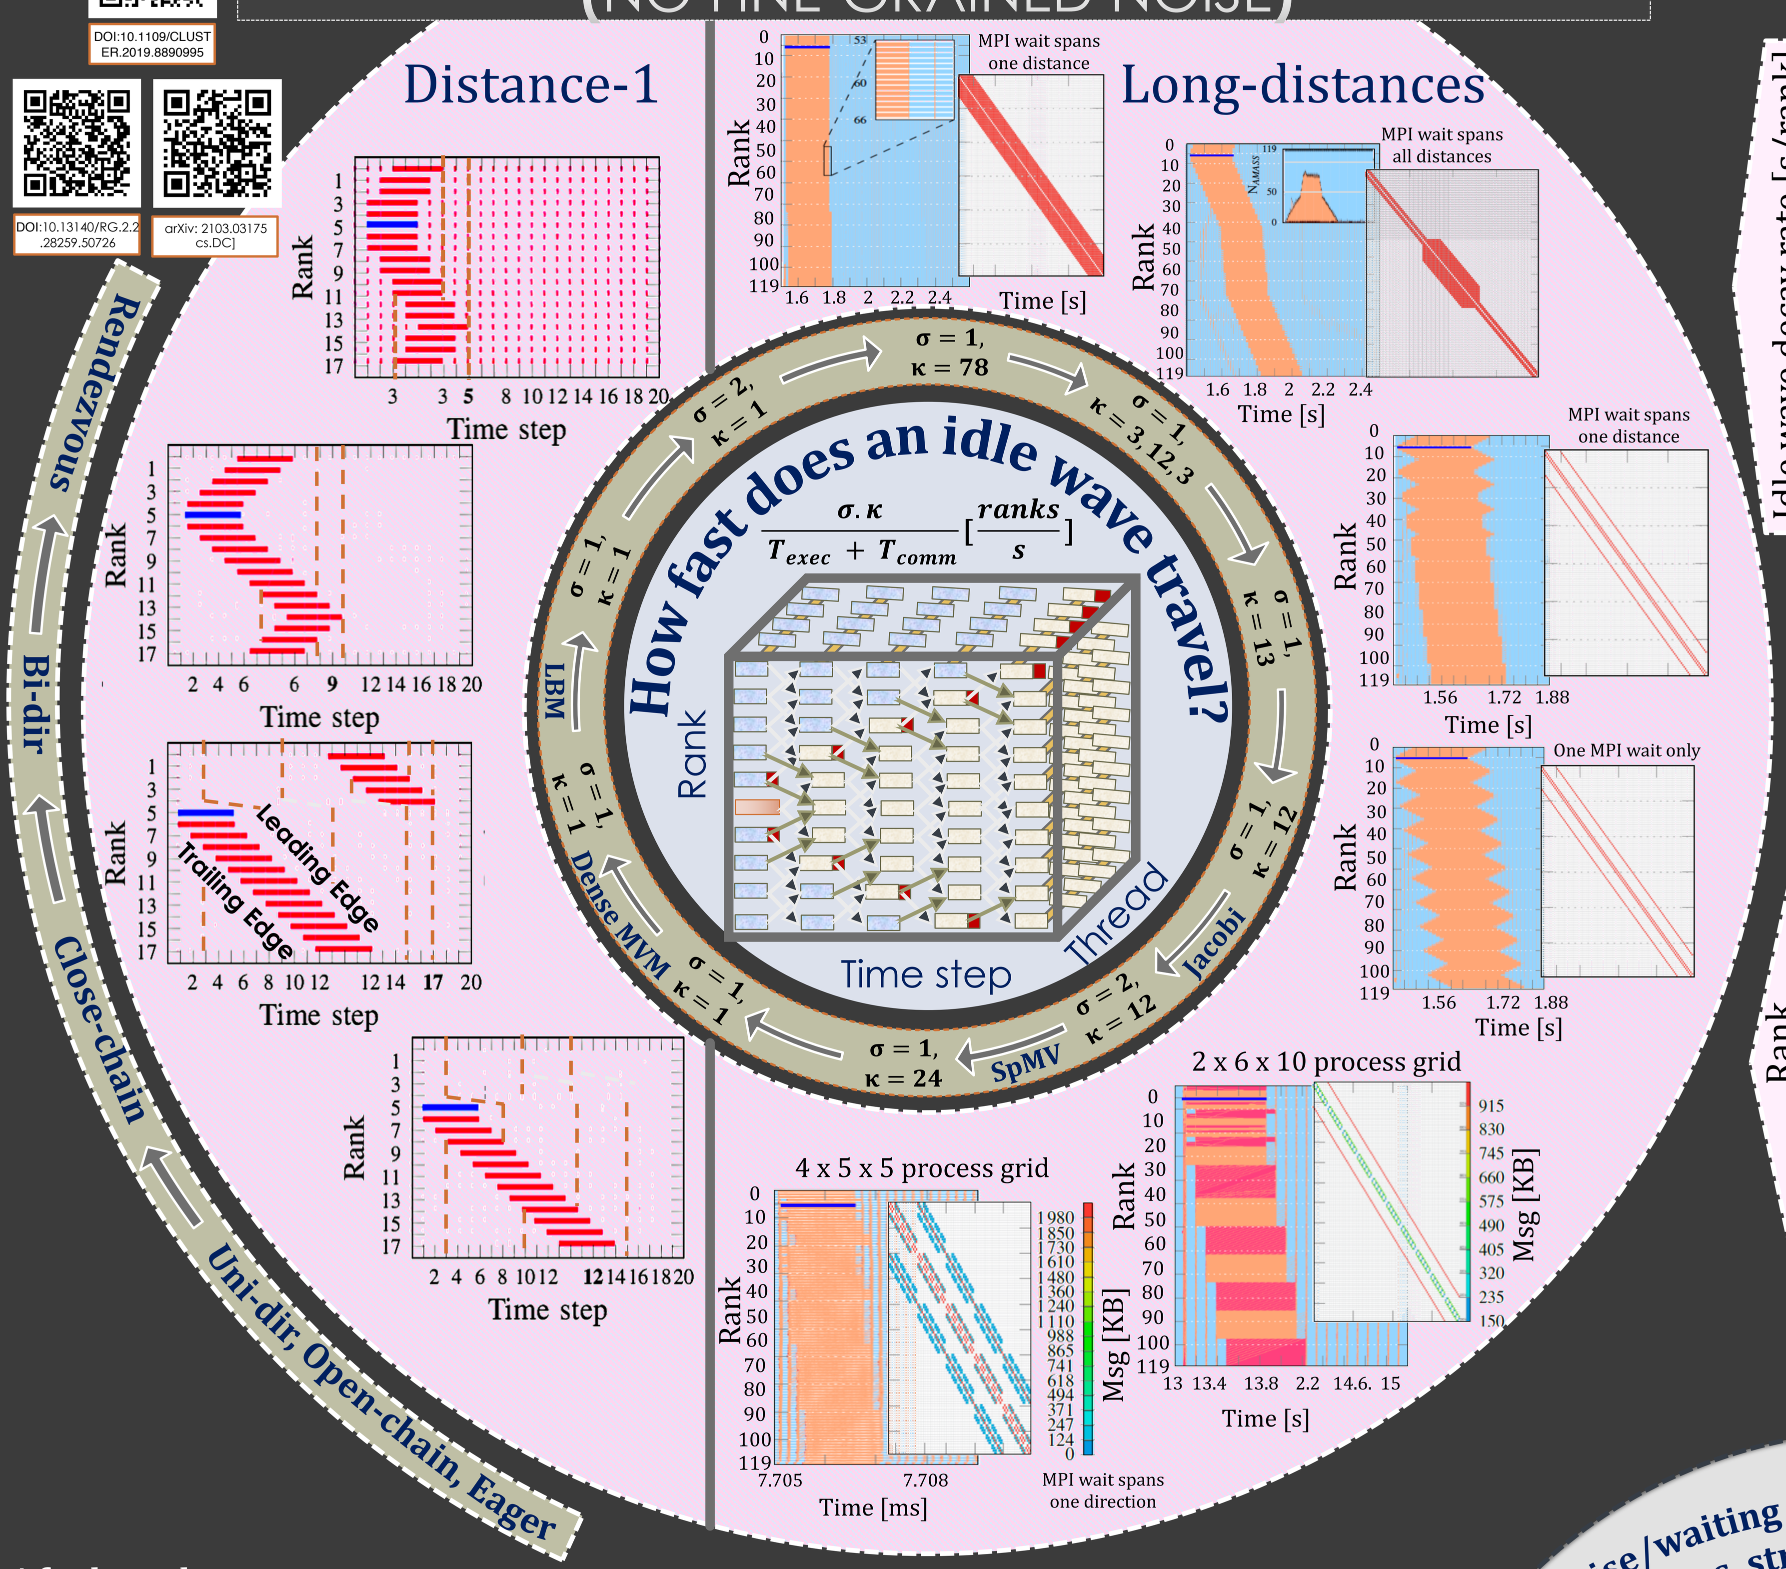

2. INTERACTION-DRIVEN WAVE DECAY (SCALABLE PROCESSES)

Afzal et al.

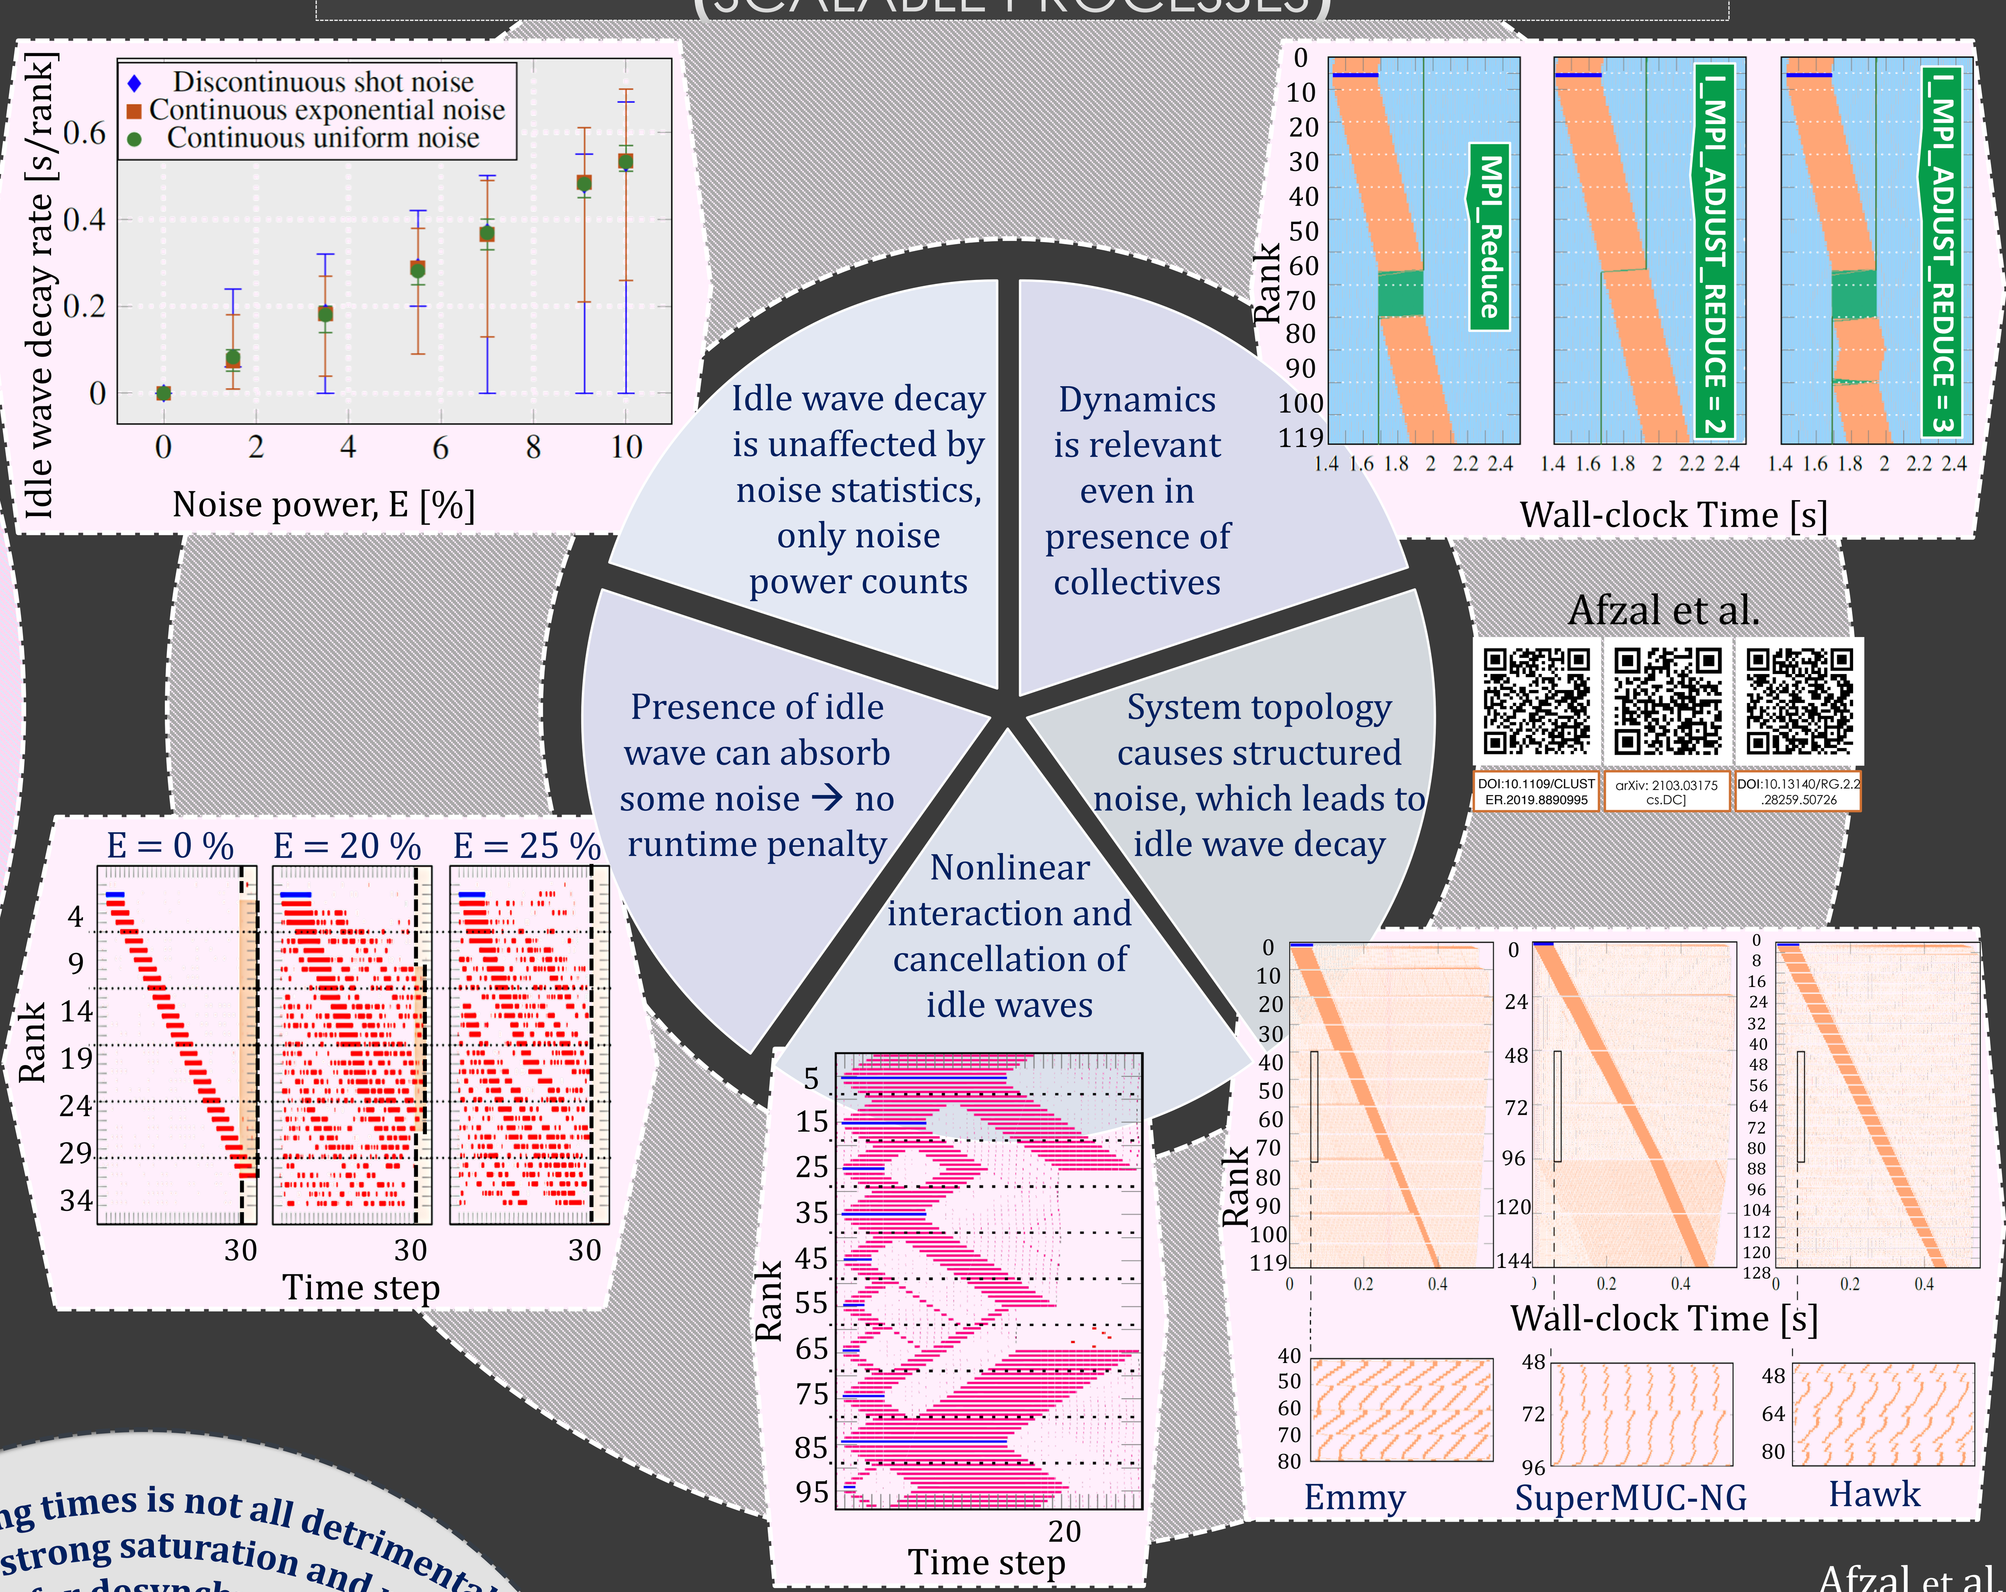

Key-takeaways

Noise/waiting times is not all detrimental slower waves, strong saturation and noise can enable high potential for desynchronization of processes, automatic communication overlap and contention evasion

Optimization techniques

↓

depending on the context, if slower idle waves are desirable, change communication, e.g., by making matrix slimmer in spMV with some reordering techniques

4. BANDWIDTH SHARING MODEL

Afzal et al.

← Model predicts bandwidth shares among different kernels

Extended scope of de-sync multi-phase applications (e.g., HPCG), spatial multitasking in GPUs and task-parallel programs

Characteristics of back-to-back kernels impact whether de-sync is amplified or damped

HPCG Benchmark on Broadwell (BDW-2)

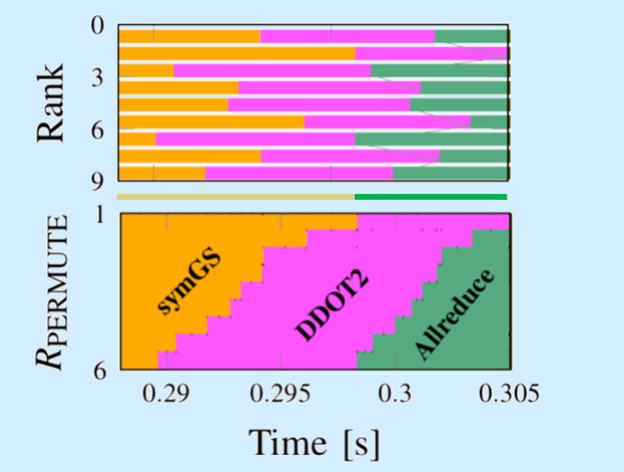

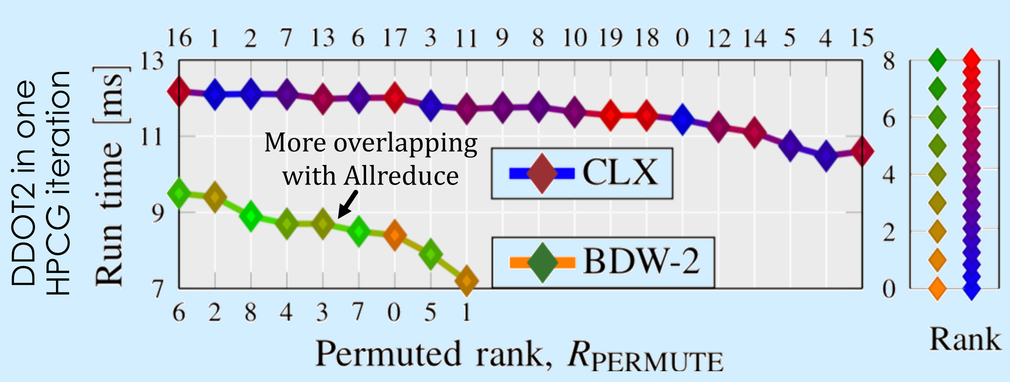

3. DE-SYNC WAVEFRONT FORMATION (CONTENDED PROCESSES)

Afzal et al.

Emergence of de-sync wavefronts requires the presence of a bottleneck across MPI processes

SuperMUC-NG @2.3 GHz, non-temporal stores, bi-dir, 1024 B, close chain, distance-1 communication

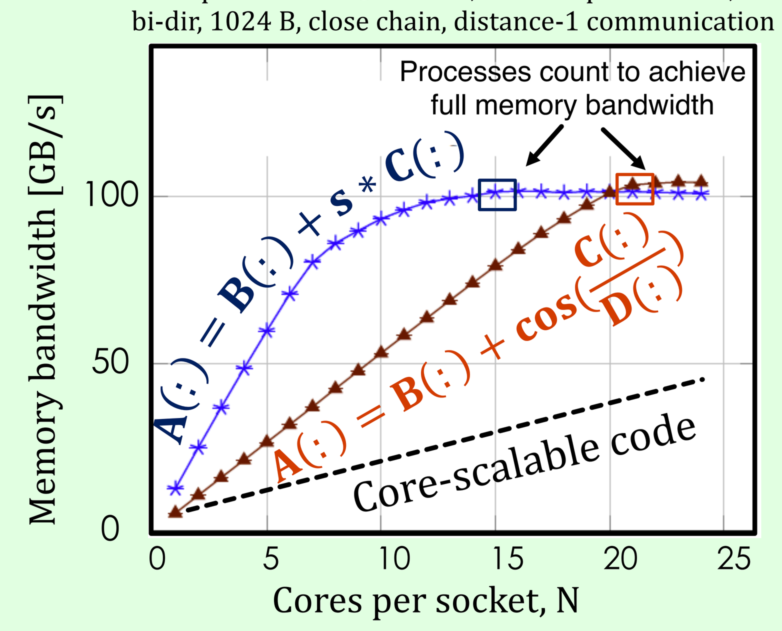

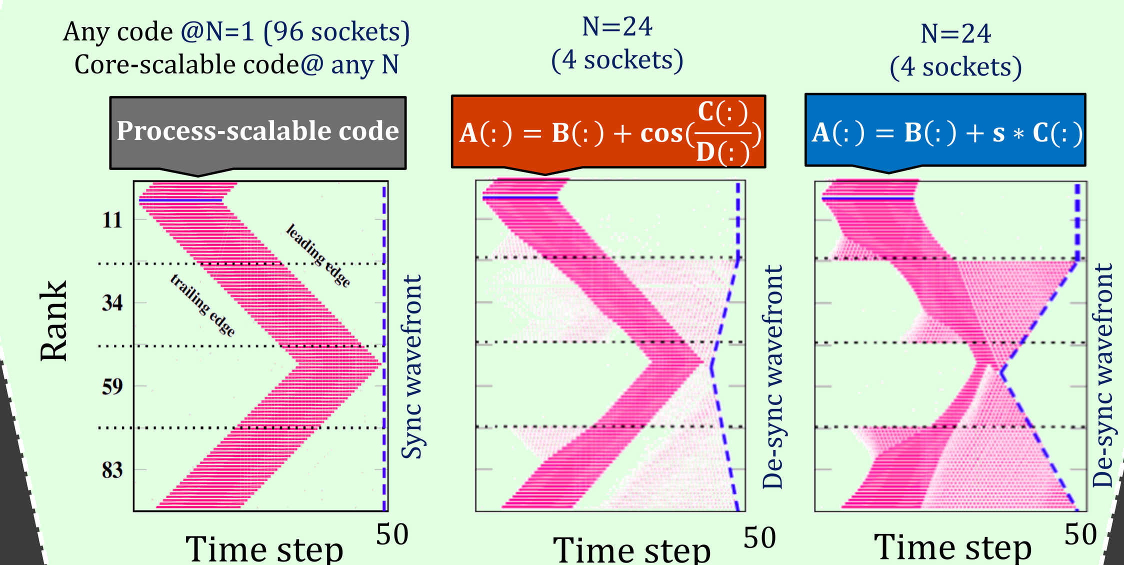

5. USEABILITY & GENERIZATION

Afzal et al.

MPI - parallelized Lattice Boltzmann Fluid Solver

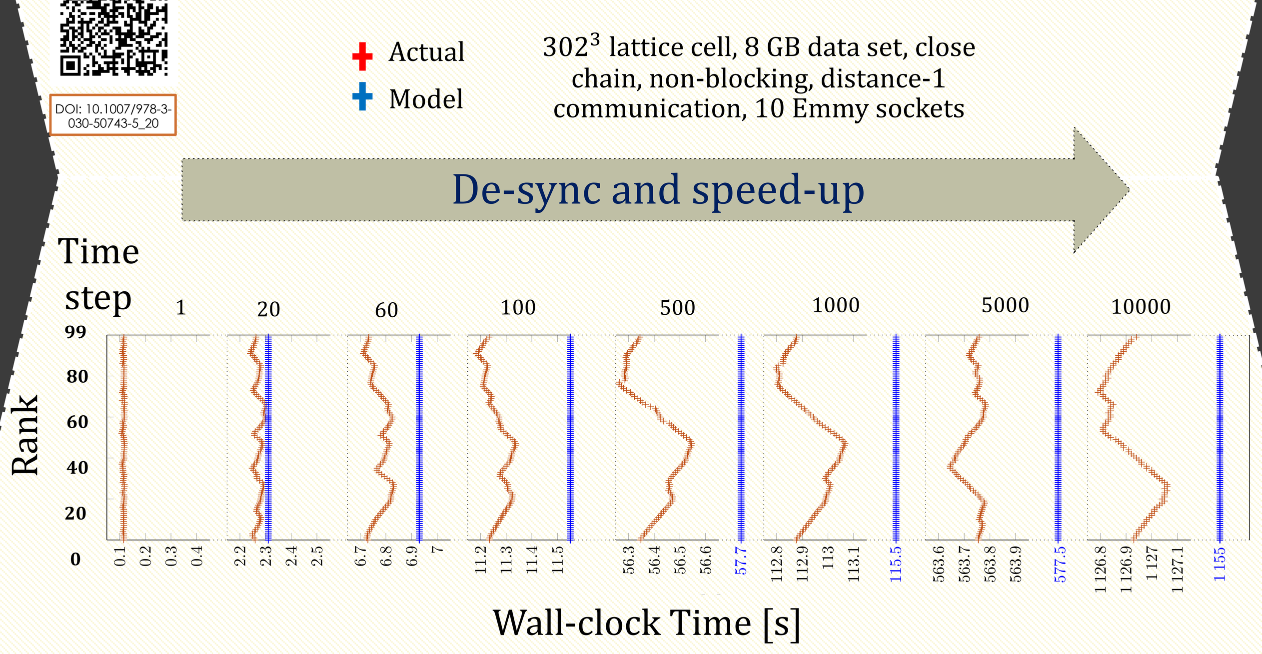

Chebychev Filter Diagonalization

128x64x64 size, 6.7 GB data set, close chain , non-blocking, 16 Emmy sockets

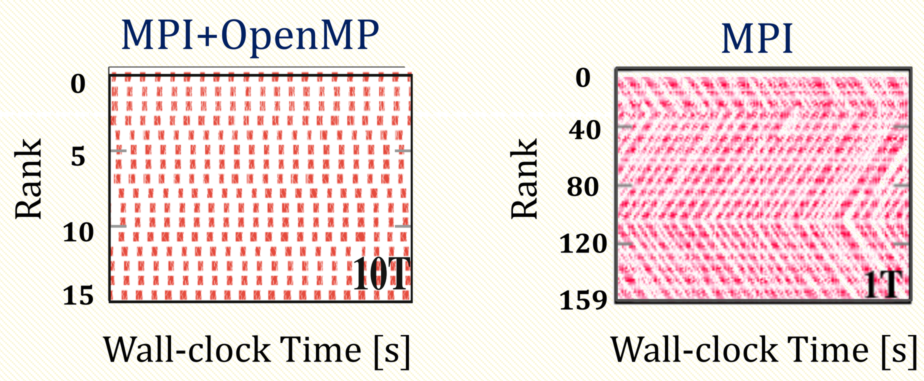

6. END-USER SIMULATION TOOL (IN-DEVELOPMENT)

A complete *analytical-model based* simulation tool will enable the investigation of the dynamics of massively parallel programs

- ➤ uses accurate analytic node-level model (ECM)
- ➤ considers memory and network bandwidth contention
- ➤ enables low-cost architectural exploration

7. PHYSICAL OSSILATOR MODEL (IN-PROGRESS)

Afzal et al.

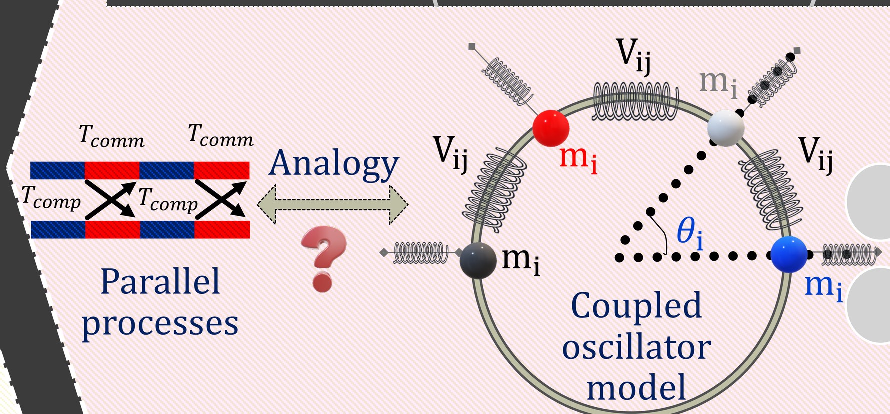

- Cost effective
- Unknown analogy
- Cluster and code characterization tool
- End-user tool

KONWIHR OMI4papps | NHR FAU | FAU Friedrich-Alexander Universität Erlangen-Nürnberg Technische Fakultät

POSTER ARTIFACTS

https://tiny.cc/NDCLPMAE-ISC21