



HPC Performance Almanac

Christian Plessl

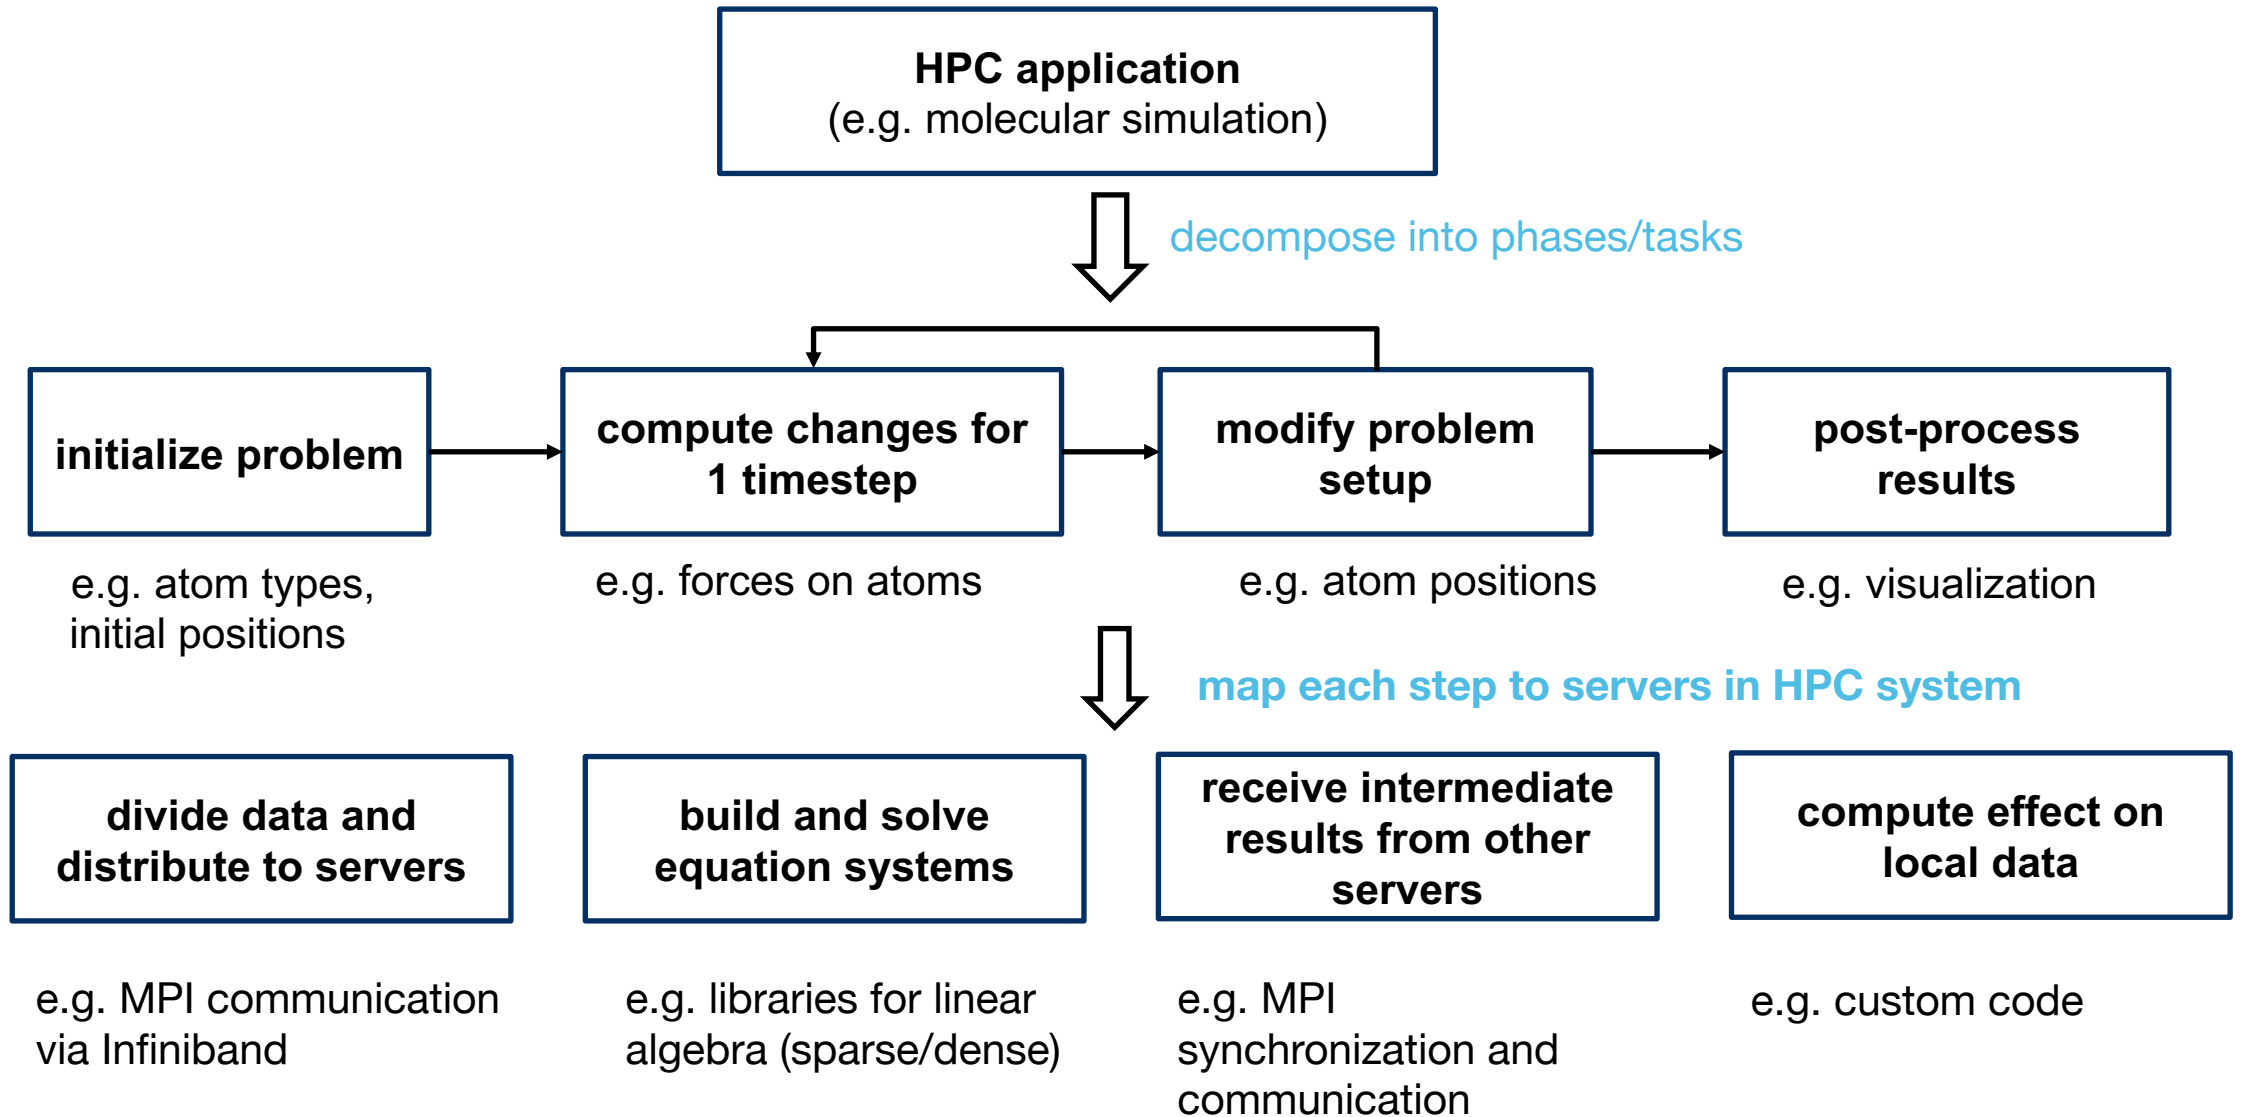
High-Performance IT Systems group

Paderborn University, Germany

Paderborn Center for Parallel Computing



Performance Modeling for HPC Applications



runtime: estimate runtime of each phase, multiply with number of executions

Getting Good Estimations on Runtime is not Easy

- What time does it take, to multiply two matrices, with
 - 1024x1024 elements
 - single precision arithmetic
 - on an Intel Xeon Gold 6148 CPU
- How can we answer this question?

Analytical Performance
Model

time \approx number of operations /
operations per second

Empirical Measurement
(Execution)


compile, execute, measure

Database

rely on measured data by
others

Analytical Performance Models: Challenges

- Performance is determined by overall bottleneck of system
 - instruction processing rate
 - memory throughput
 - I/O throughput
- Concrete algorithm determines number of operations
- Dynamic system behavior




time \approx number of operations /
operations per second

Mode	Base	Turbo Frequency/Active Cores					
		1	2	4	8	16	20
Normal	2400	3700	3700	3500	3400	3300	3100
AVX2	1900	3600	3600	3400	3300	2800	2600
AVX512	1600	3500	3500	3300	3100	2300	2200

Empirical Measurements: Challenges

- Requires that we already have an implementation available
- It is easy to make mistakes in measurements because performance is sensitive to many parameters
 - problem size
 - data types / layout
 - number of threads
 - compiler
 - operating system




which more setting to use for optimal performance?

Database: Challenges

- No comprehensive performance database available
- Suffers from complexity of empirical measurements
 - so many CPU / system variants
 - hard to find data for exactly same system
- Even if data is published, the quality is frequently uncertain

Intel® oneAPI Math Kernel Library (oneMKL) 2021.1 SGEMM & BF16 GEMM on Intel® Xeon® Gold 6348H Processor at 2.30 GHz



Testing Date: Performance results are based on testing by Intel as of November 18, 2020 and may not reflect all publicly available security updates.

Configuration Details and Workload Setup: Intel® Xeon® Gold 6348H Processor @ 2.30 GHz, 384GB, Intel® oneAPI Math Kernel Library 2021.1 Gold (Intel® oneMKL); SGEMM & DGEMM performance for square matrix dimensions between 800 and 30000.

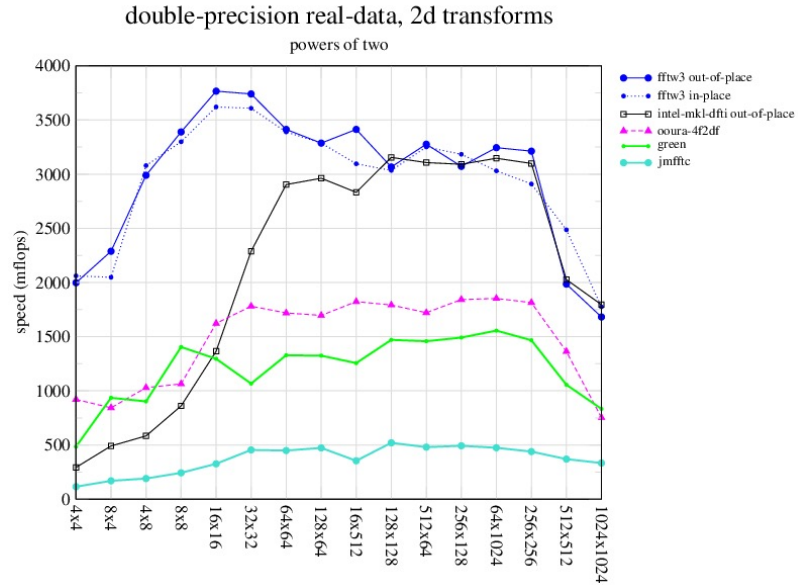
Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details. No product or component can be absolutely secure.

Performance varies by use, configuration, and other factors. Learn more at www.intel.com/PerformanceIndex. Your costs and results may vary.

The HPC Performance Almanac Project Group

- **Goals & Activities**
 - develop a **trusted database of performance data** for common HPC kernels
 - automate the benchmarking process with **continuous integration pipeline**
 - make data available through **web application**
 - performance data
 - system hardware specification
 - compilation and runtime environment
 - release code and data as open source
- **Project requires broad range of **complementary skills****
 - high-performance computing
 - code optimization
 - software engineering and automation
 - usability

Mockup: Fast Fourier Transform Performance



Check the wonderful
<https://ourworldindata.org>
website for inspiration how such a
webapp could look like

1D / 2D / 3D
single / double precision
range
algorithm / package selection (baseline)

compiler / version
optimization flags
used libraries

reference to benchmarking code in Git

download figure / csv / json

kernel / data options

build and runtime
environment

reference code

data download

- **Admission**

- register your interest via questionnaire (motivation statement and background)
- preference will be given to students who attended HPC lecture and passed exam

- **Rules**

- regular meetings with group members and supervisor (weekly)
- intermediate presentations
- mandatory (bi)weekly meetings with supervisor

- **Further information**

- <https://en.cs.uni-paderborn.de/hit/teaching/courses/ws-2021-22/hpc-pa/>
- email to Christian Plessl