EUROfusion

GENE-X

Carlos Romero Madrid

HPC ACHs | 15-16 November 2023

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

About GENE-X

GENE-X is a gyrokinetic code to simulate plasma turbulence in X-point geometries in tokamaks and stellarators.

About GENE-X

GENE-X is a gyrokinetic code to simulate plasma turbulence in X-point geometries in tokamaks and stellarators.

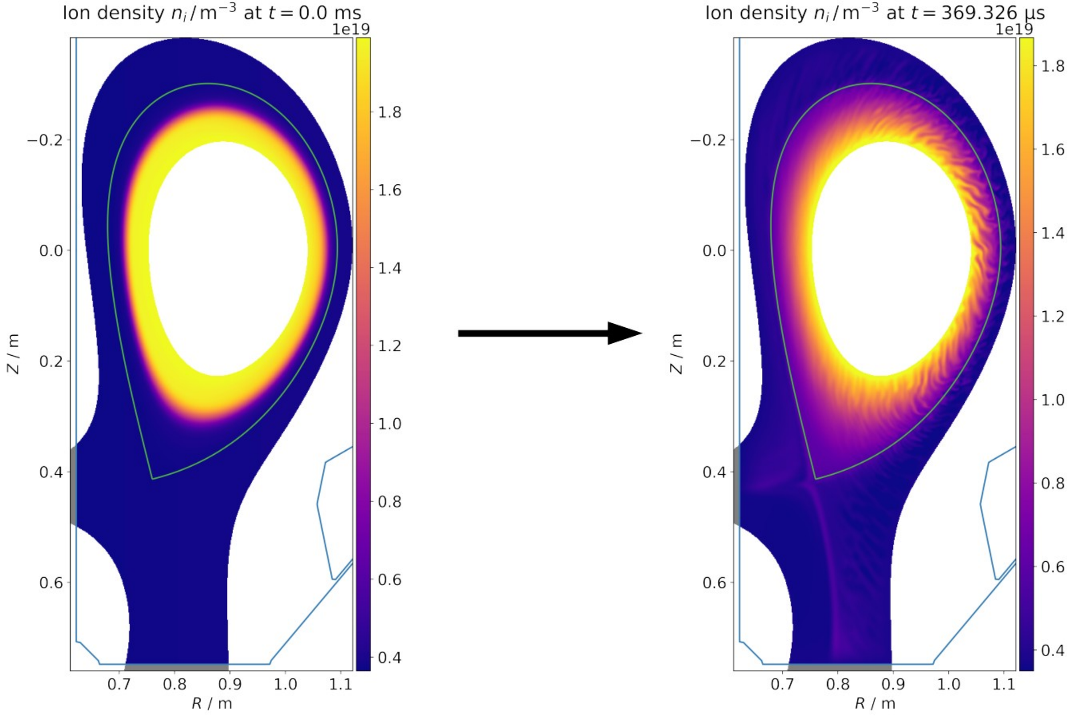

About GENE-X

GENE-X is a gyrokinetic code to simulate plasma turbulence in X-point geometries in tokamaks and stellarators.

- Parallelization of a typical case:
  - Hybrid MPI+OMP. MPI in φ, v|| , μ, α. OMP in RZ.
  - Example:
    - (RZ, φ, v|| , μ, α) = (1, 16, 2, 10, 2).
    - 320 nodes - 72 cores. 640 MPI procs, 36 cores for OMP

Work required into ACH

Assessment of a reordering algorithm based on the multigrid approach aimed at reducing L3 cache misses and reducing simulation time.

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- Scalability/Profiler
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- Conclusions so far
- Steps to follow

Introduction

- About GENE-X
- Work required into ACH
- **Reordering**
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- Scalability/Profiler
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- Conclusions so far
- Steps to follow

Reordering

- Z-Morton

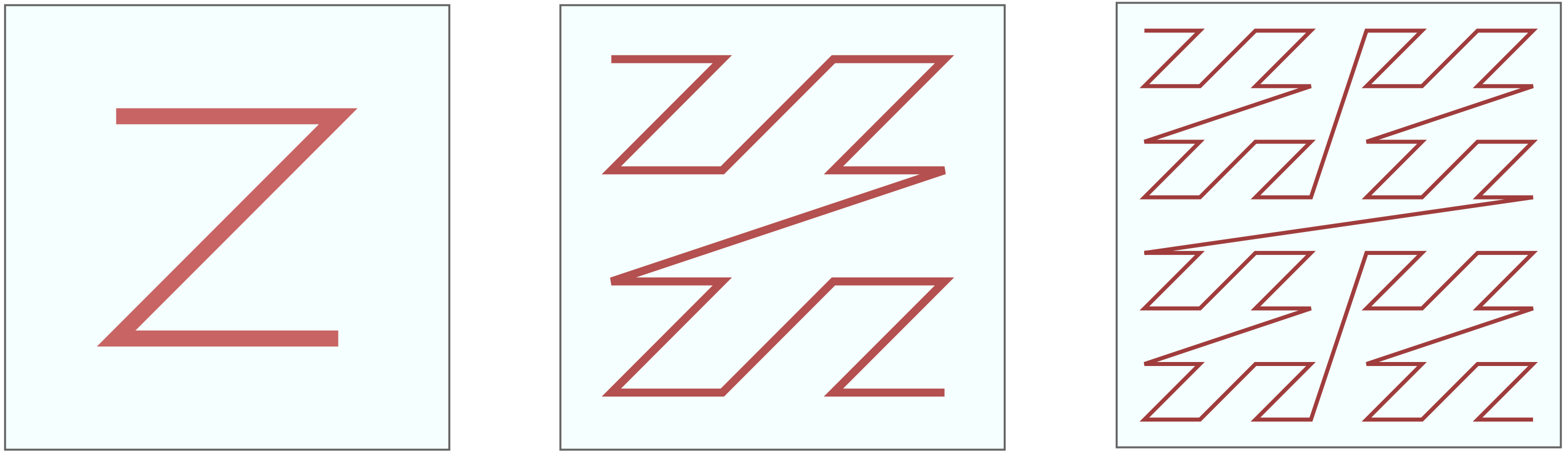

Reordering

- Z-Morton

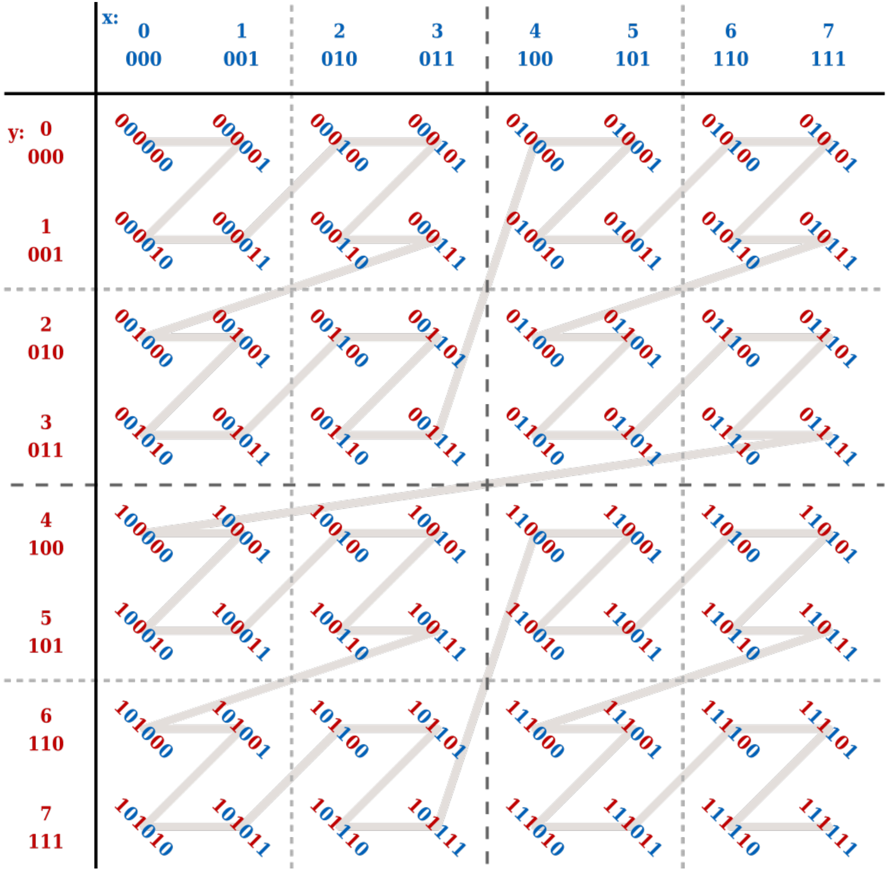

Reordering

The **multigrid** method is an algorithm for solving differential equations using a hierarchy of discretizations, very useful in problems exhibiting multiple scales of behavior.

Reordering

The **multigrid** method is an algorithm for solving differential equations using a hierarchy of discretizations, very useful in problems exhibiting multiple scales of behavior.

- In GENE-X:
  - 3D structure of **inner**, **boundary** and **ghost** layers.
  - Independent reordering on each layer.
  - The variable **multigrid_reorder_size** is an integer m-array.

Reordering

The **multigrid** method is an algorithm for solving differential equations using a hierarchy of discretizations, very useful in problems exhibiting multiple scales of behavior.

- In GENE-X:
  - 3D structure of **inner**, **boundary** and **ghost** layers.
  - Independent reordering on each layer.
  - The variable **multigrid_reorder_size** is an integer m-array.

- Test choices
  [0 0 0] - [0 0 2] - [0 2 0] - [2 0 0] - [0 2 2] - [2 2 2]

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- **Intel Advisor**
- Scalability/Profiler
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- Conclusions so far
- Steps to follow

Intel Advisor

Intel Advisor is used to see how well your application uses hardware resources, what is the main limiting factor, and what are the optimization candidates.

Intel Advisor

Intel Advisor is used to see how well your application uses hardware resources, what is the main limiting factor, and what are the optimization candidates.

- We identify the most time consuming modules: **the field solvers**.
  - Vlasov operator
  - Ampère’s law
  - Ohm’s law
  - Maxwell’s equation

Intel Advisor

Intel Advisor is used to see how well your application uses hardware resources, what is the main limiting factor, and what are the optimization candidates.

- We identify the most time consuming modules: **the field solvers**.
  - Vlasov operator
  - Ampère’s law
  - Ohm’s law
  - Maxwell’s equation

- Let’s have a look at the **CPU/Memory Roofline Model**:

Intel Advisor – CPU/Memory Roofline Model

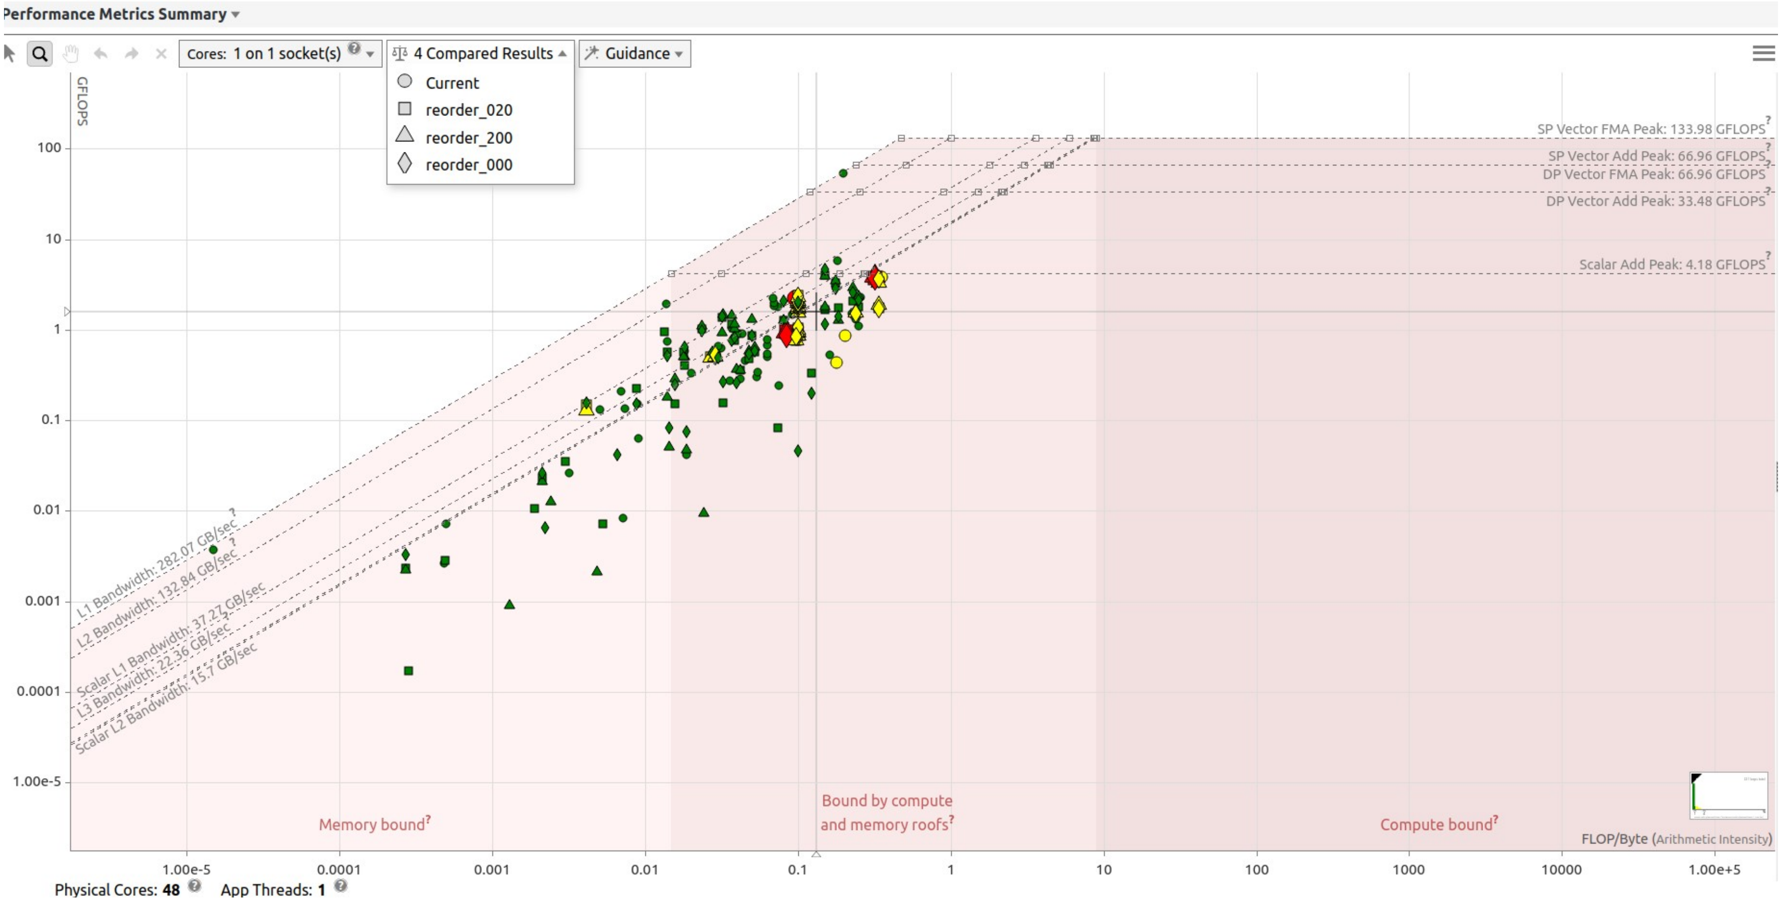

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- **Scalability/Profiler**
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- Conclusions so far
- Steps to follow

Scalability/Profiler

- We have conducted 2 reordering studies:
  - In the local machine → 30.000 grid points
  - In the MN4 cluster → 700.000 grid points

Scalability/Profiler

- We have conducted 2 reordering studies:
  - In the local machine → 30.000 grid points
  - In the MN4 cluster → 700.000 grid points

- Let’s have a look a couple of the computationally heavy modules:
  - Vlasov
  - Ampère
  - Total

Scalability/Profiler – LOCAL - Vlasov

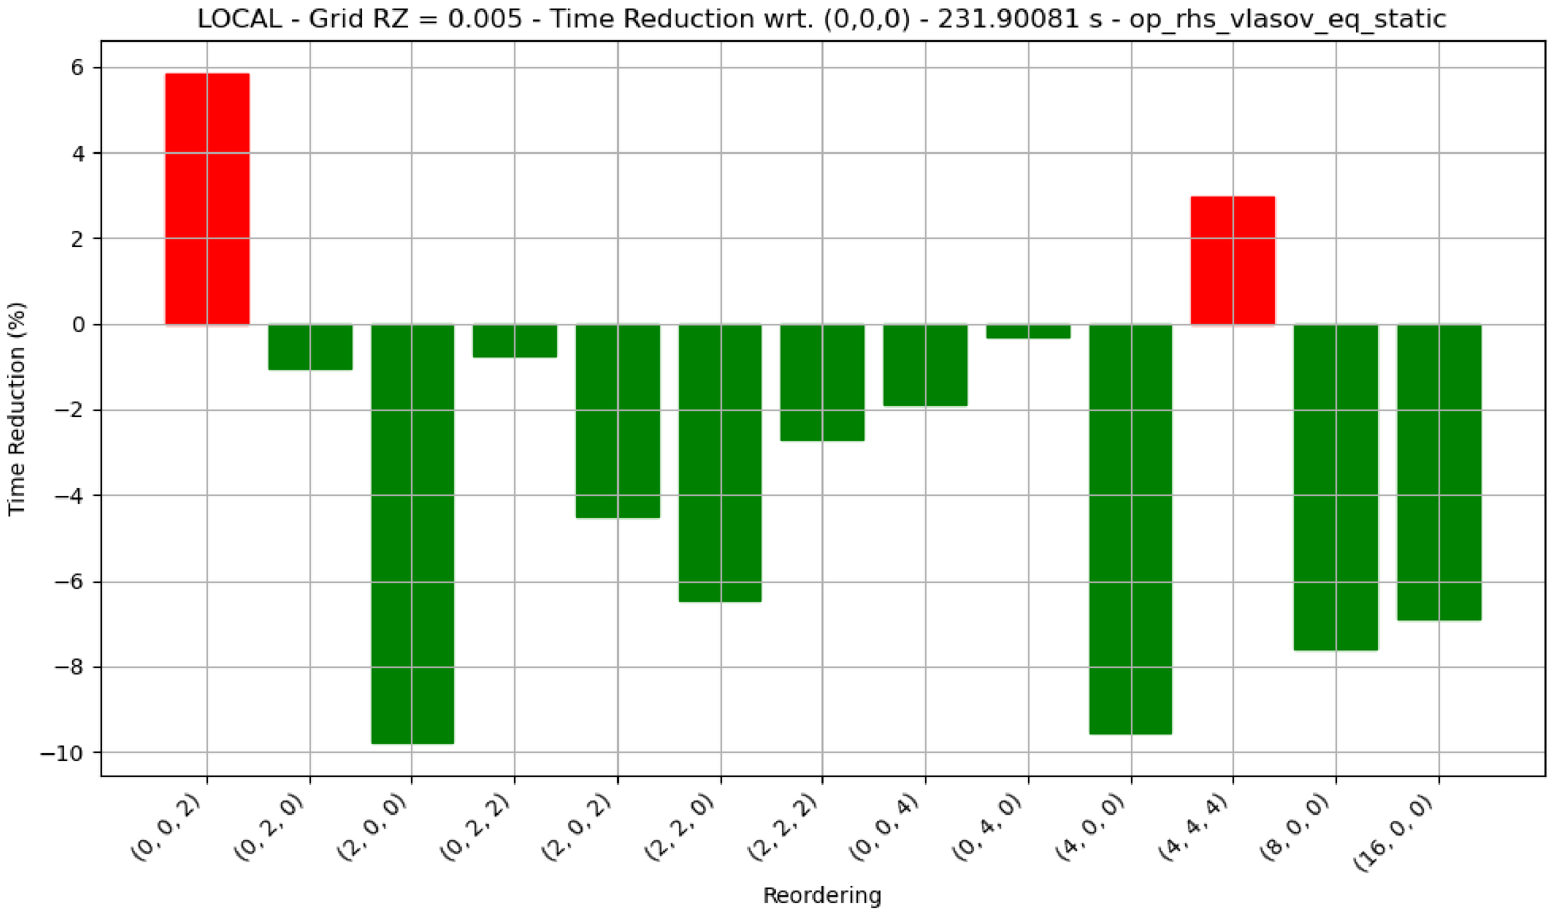

Scalability/Profiler – LOCAL - Ampère

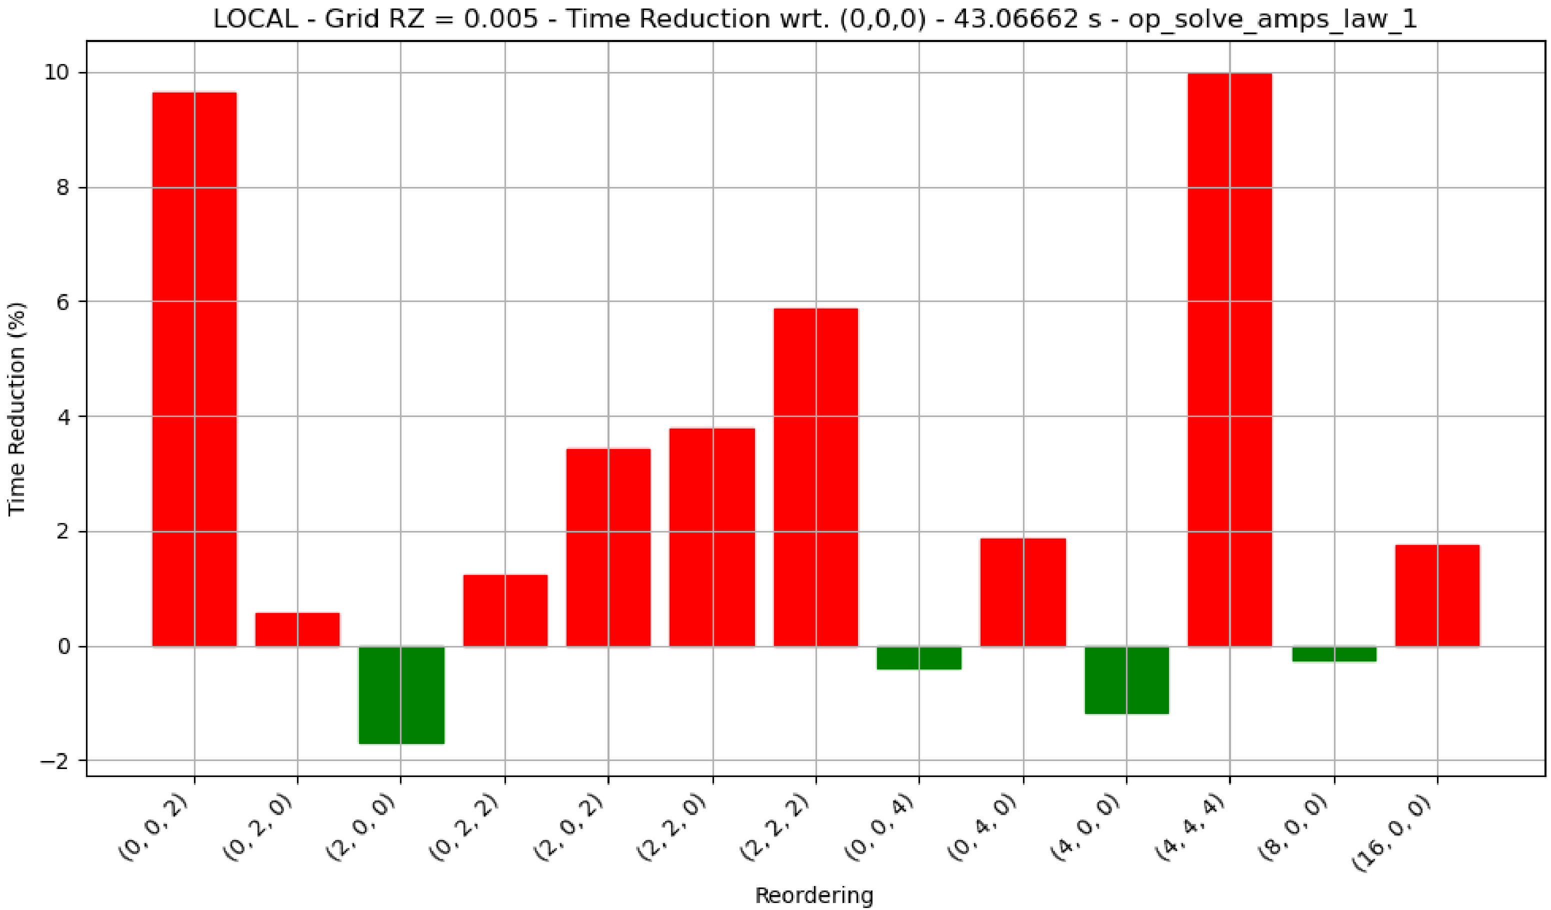

Scalability/Profiler – LOCAL - Total

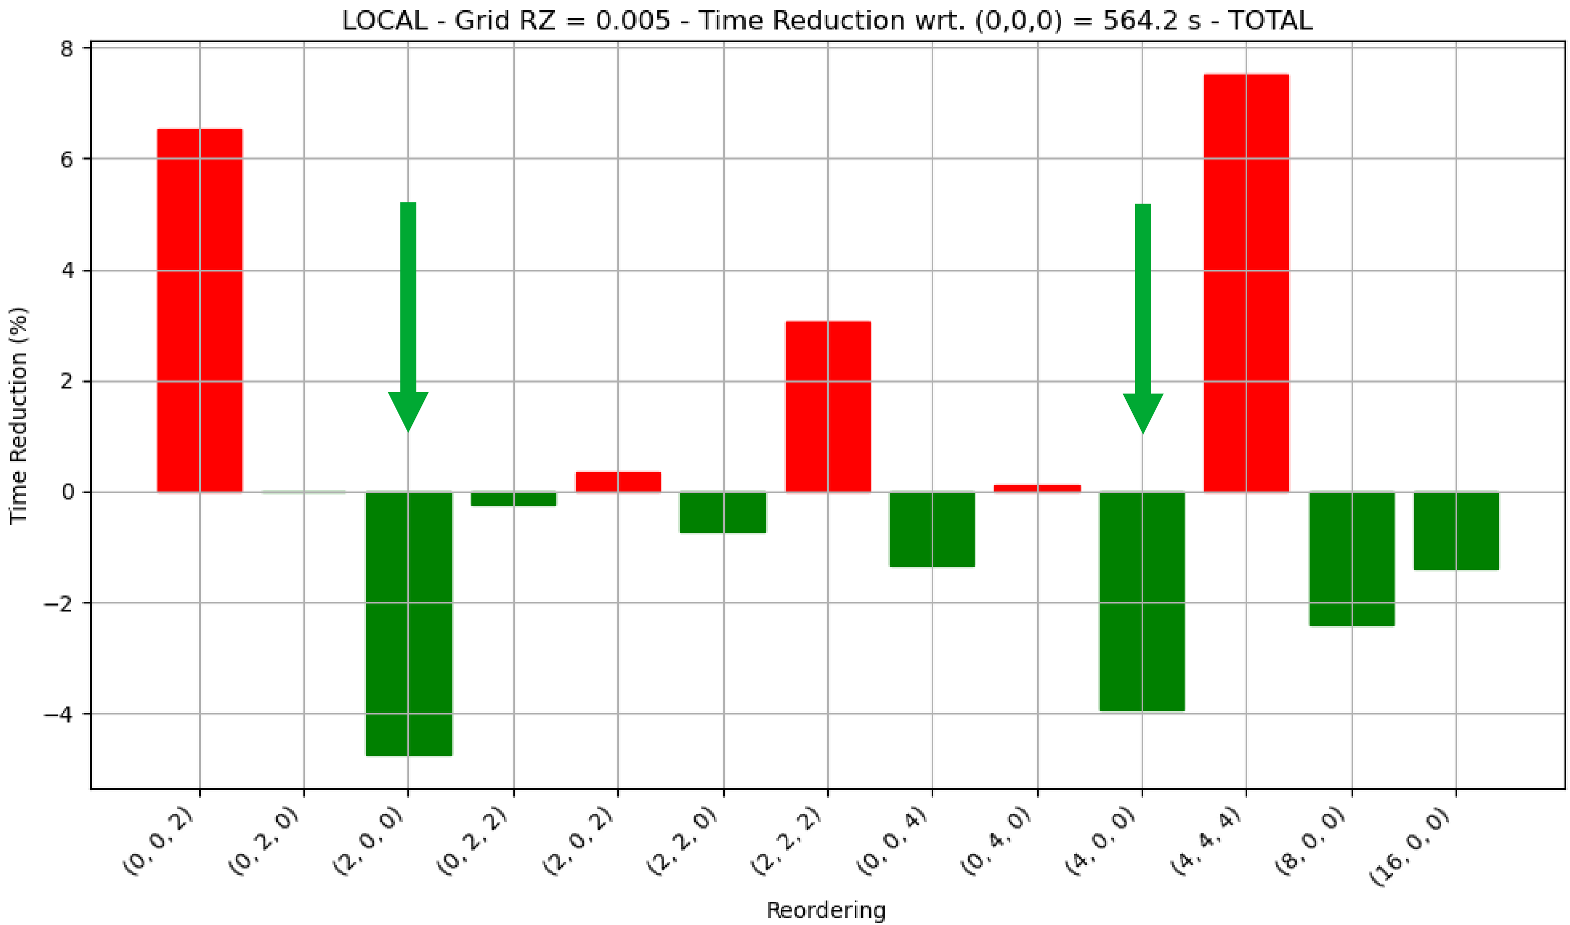

Scalability/Profiler – MN4 - Total

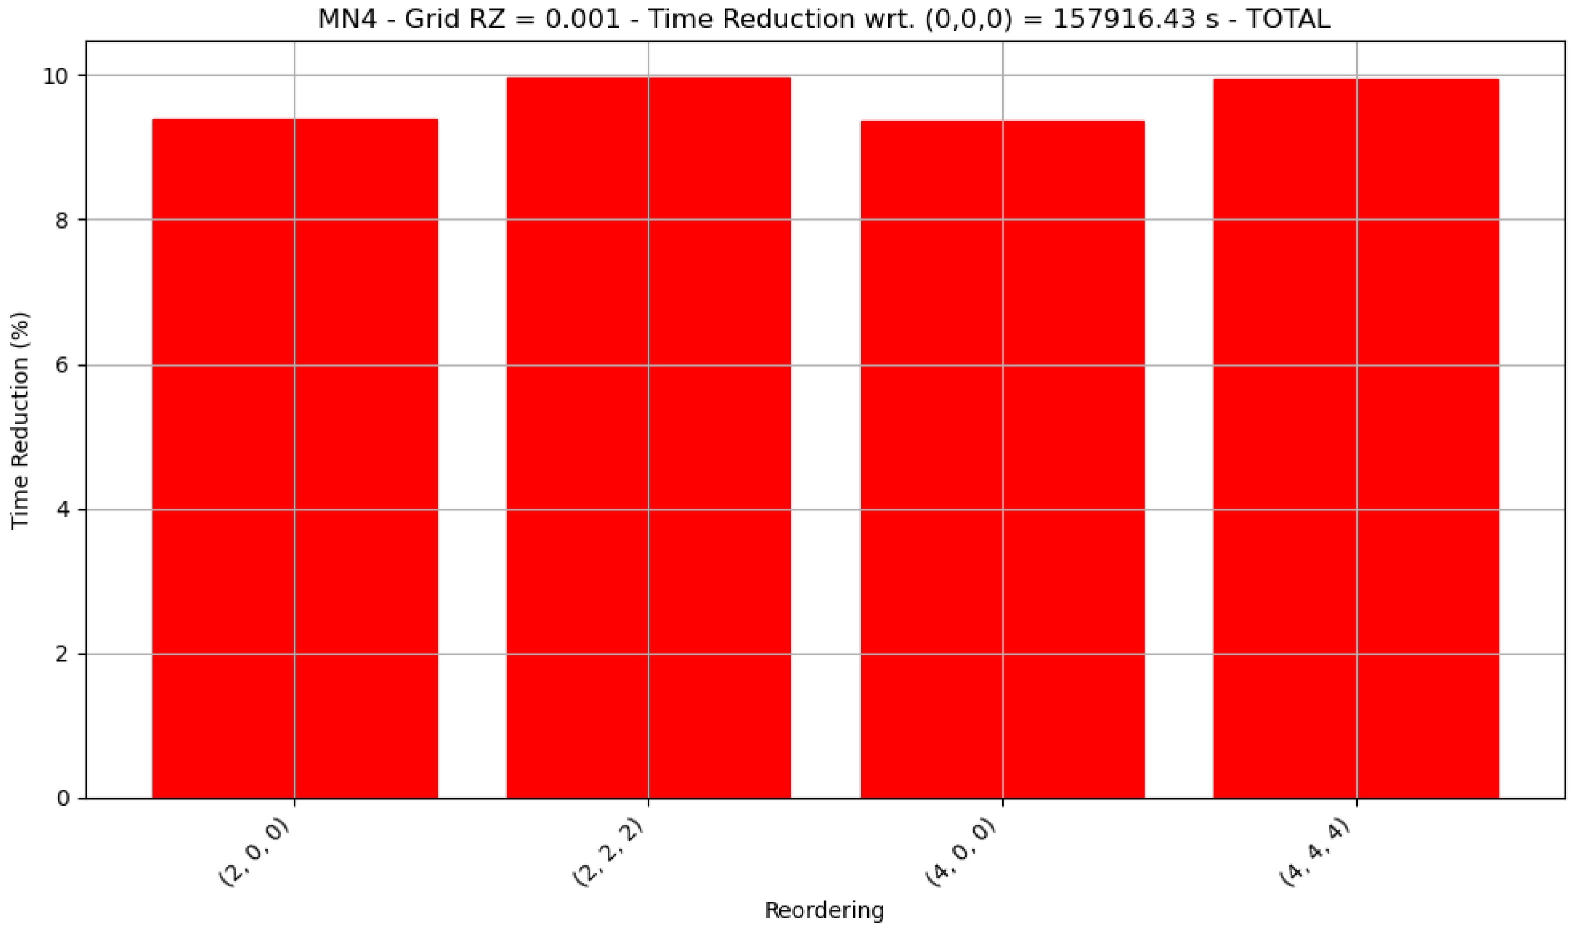

Scalability/Profiler – MN4 - Total

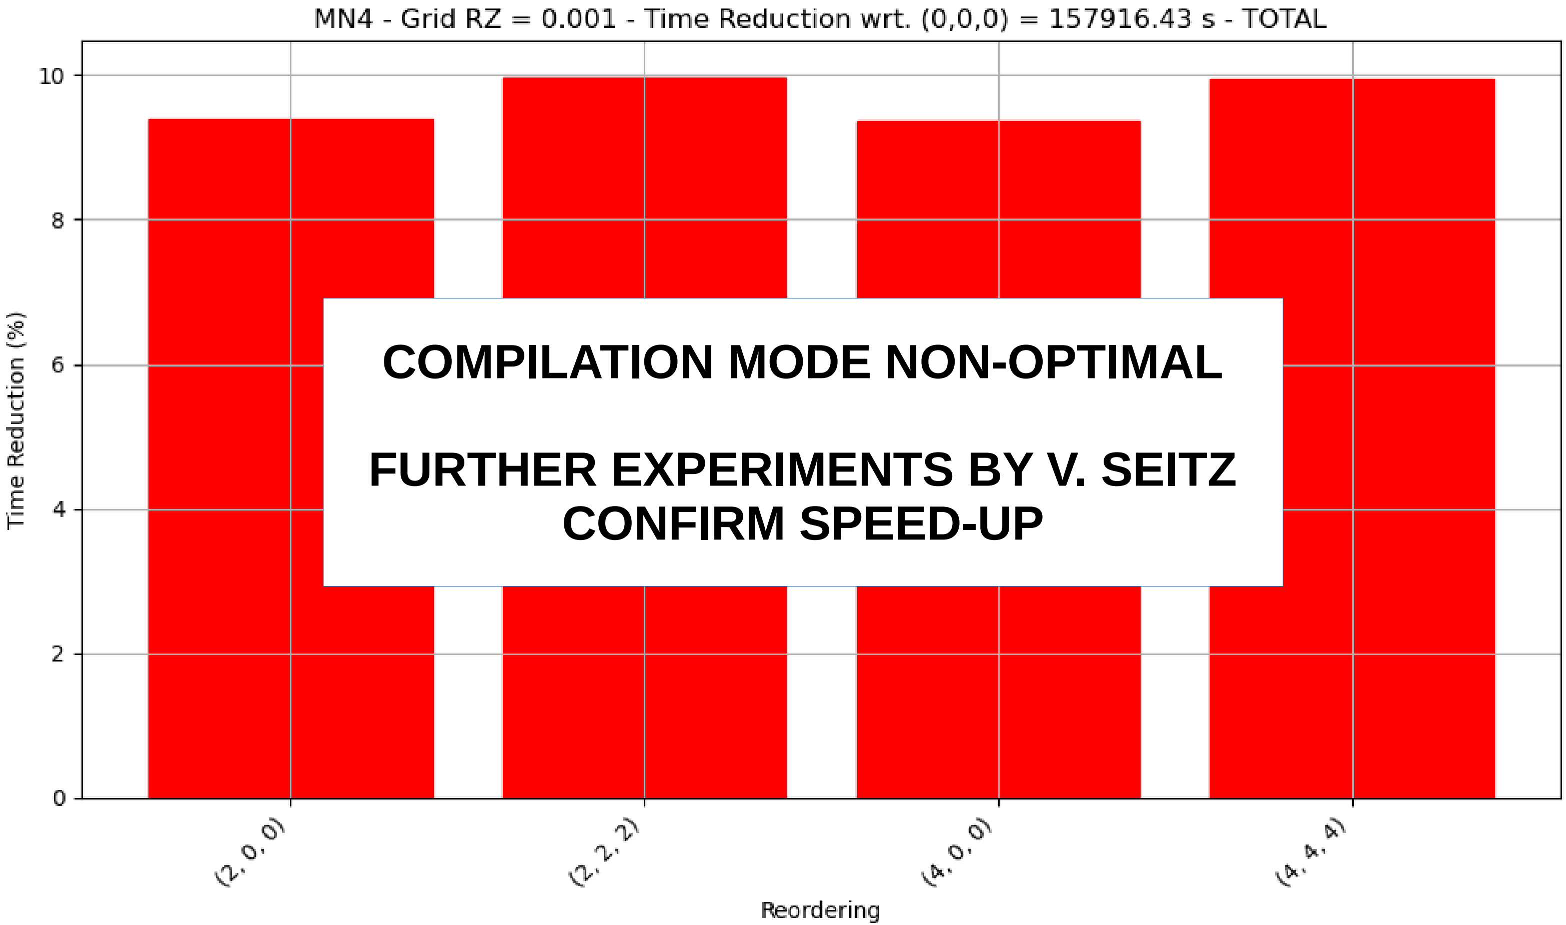

Scalability/Profiler - OMP

- GENE-X implements OpenMP for RZ-grid parallelization.

Scalability/Profiler - OMP

- GENE-X implements OpenMP for RZ-grid parallelization.

- Let's study the scalability with the same 2 studies:
  - In the local machine → 30.000 grid points
  - In the MN4 cluster → 700.000 grid points

Scalability/Profiler - OMP

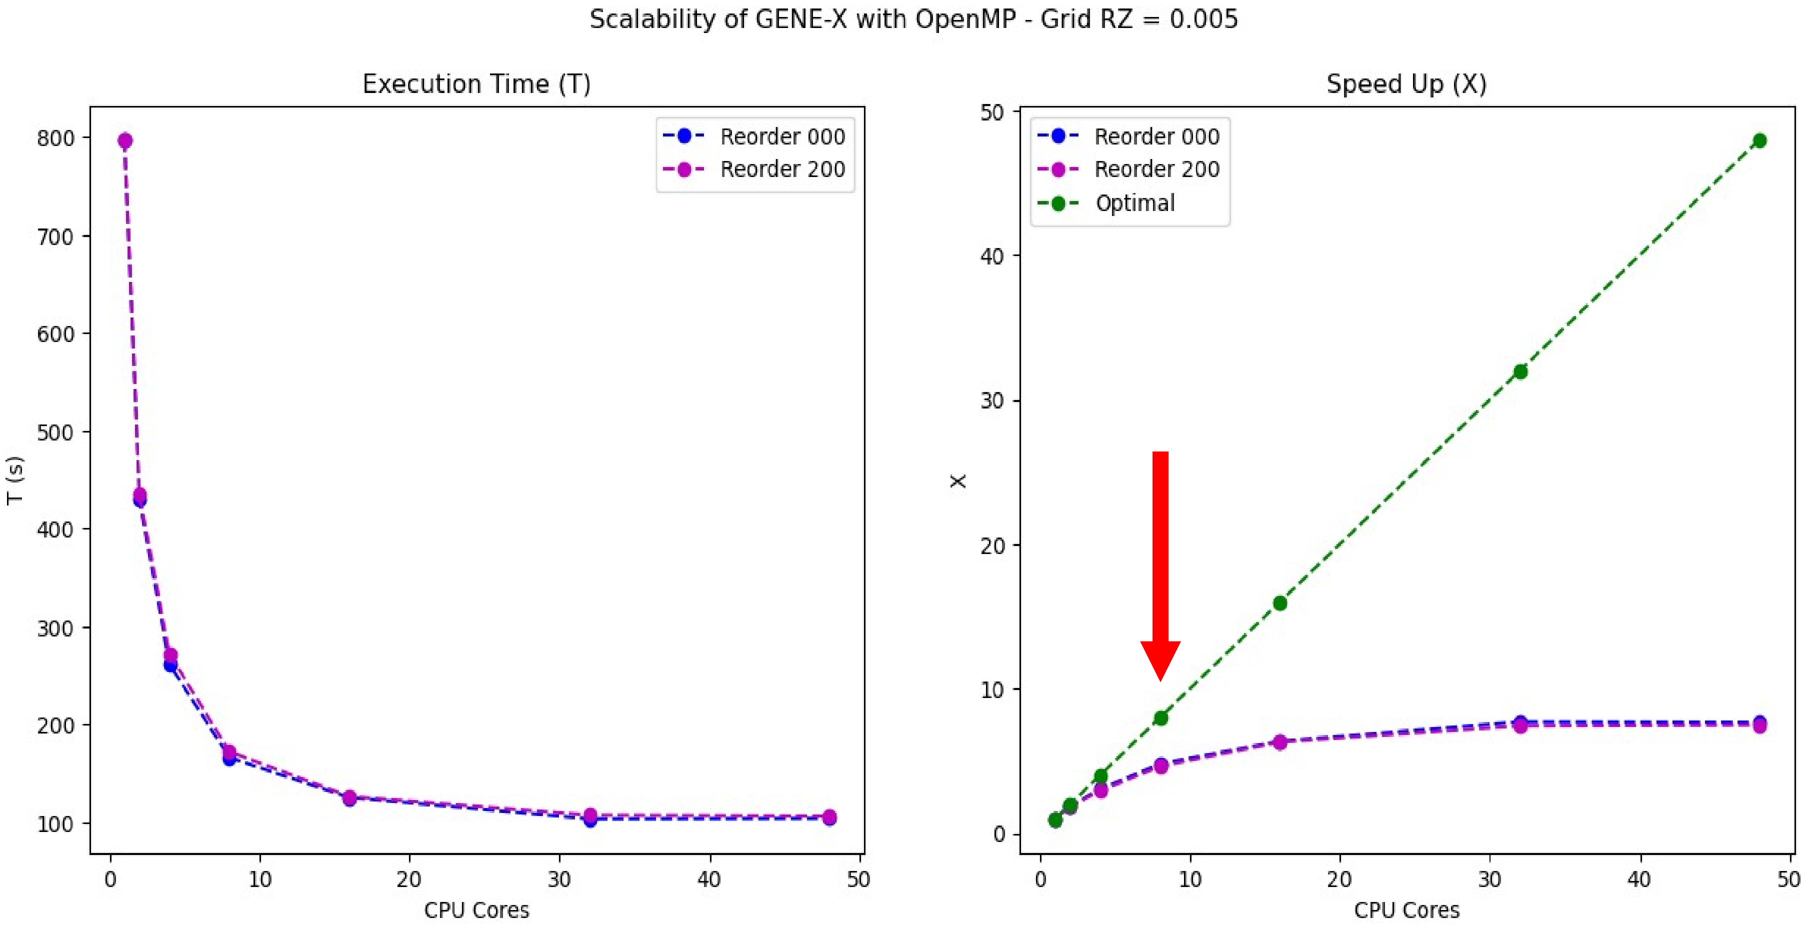

Scalability/Profiler - OMP

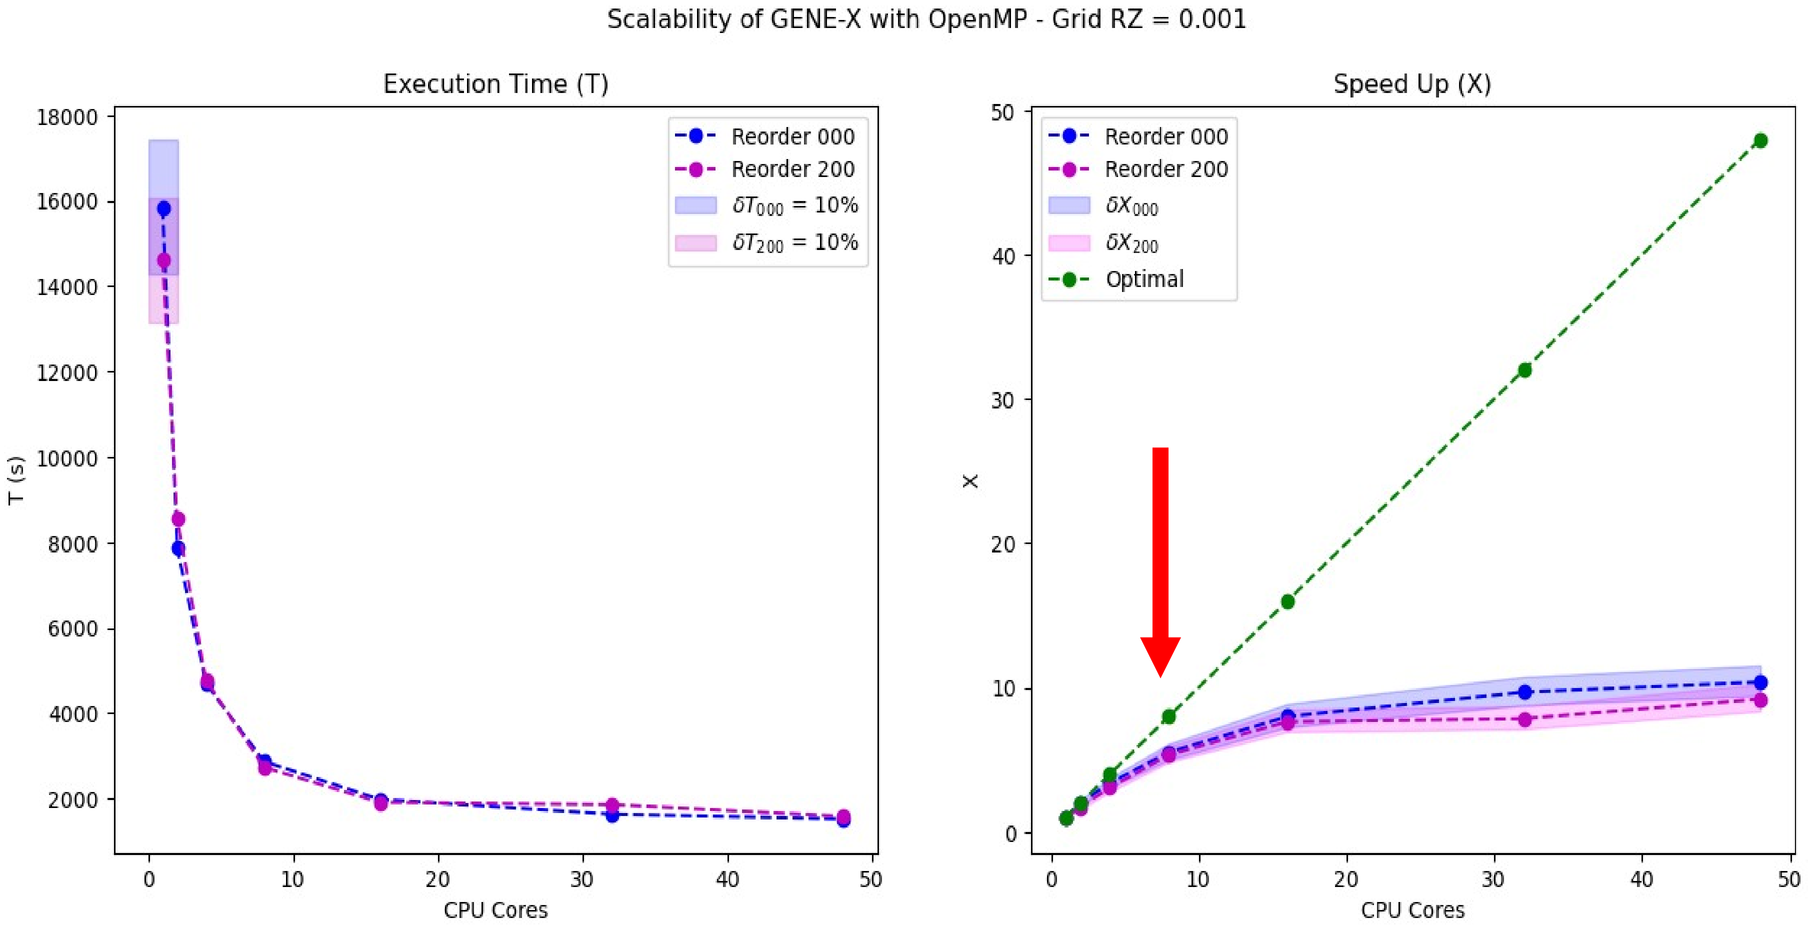

Scalability/Profiler

- LOCAL
  - Saturation achieved at around 8 OMP cores.
  - 5% improvement from case 000 to 200.

Scalability/Profiler

- LOCAL
  - Saturation achieved at around 8 OMP cores.
  - 5% improvement from case 000 to 200.

- MN4
  - Saturation achieved at around 8 OMP cores.
  - No improvement from case 000 to 200.

Scalability/Profiler

- LOCAL
  - Saturation achieved at around 8 OMP cores.
  - 5% improvement from case 000 to 200.

- MN4
  - Saturation achieved at around 8 OMP cores.
  - No improvement from case 000 to 200.

COMPILATION MODE NON-OPTIMAL

FURTHER EXPERIMENTS BY V. SEITZ CONFIRM SPEED-UP

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- Scalability/Profiler
  - Local
  - MN4
- **EXTRAE**
  - L3 cache misses
- Conclusions so far
- Steps to follow

EXTRAE

EXTRAE is a dynamic instrumentation package to trace programs compiled and run with the shared memory model (like OpenMP and pthreads), MPI or both programming models.

EXTRAE

EXTRAE is a dynamic instrumentation package to trace programs compiled and run with the shared memory model (like OpenMP and pthreads), MPI or both programming models.

- Raven at IPP
- Newest master of GENE-X
- 1 MPI and 72 OMP cores for both executions
- ITER geometry
    - GENE-X still crashes due to small L3 caché
    - Used benchmark-operators, a very mesh dependent routine

EXTRAE

- "Performance" of **op_rhs_vlasov_eq_static_t**
  - **went up by a factor of ~1.25x with [2 0 0 0] configuration**
  - also tested [4 0 0 0] and [2 2 0 0] without any big impacts

EXTRAE

- "Performance" of **op_rhs_vlasov_eq_static_t**
  - **went up by a factor of ~1.25x with [2 0 0 0] configuration**
  - also tested [4 0 0 0] and [2 2 0 0] without any big impacts

- Why are we faster? Can we see anything in the hardware counters?
  - **Yes**, we see a cut of L3 cache misses by nearly 1/1.5 which might be responsible for the better IPC we see in that region.
  - **IPC from 2.3 → 2.8**

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- Scalability/Profiler
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- **Conclusions so far**
- Steps to follow

Conclusions so far

- Speed-up is achieved at the coarser levels of the multigrid.
  - The coarser, the higher the reordering factor.

Conclusions so far

- Speed-up is achieved at the coarser levels of the multigrid.
  - The coarser, the higher the reordering factor.

- OMP parallelization achieves saturation around 8 cores.
  - Need to test behaviour with heavy production cases, might improve.

Conclusions so far

- Speed-up is achieved at the coarser levels of the multigrid.
  - The coarser, the higher the reordering factor.
- OMP parallelization achieves saturation around 8 cores.
  - Need to test behaviour with heavy production cases, might improve.
- IPC clearly improves with the reordering of the multigrid.
  - 1.25x speed-up achieved in V. Seitz’s run
  - Need to re-run MN4 cases with proper compilation.

Introduction

- About GENE-X
- Work required into ACH
- Reordering
  - Z-Morton
  - Multigrid
  - Test choices
- Intel Advisor
- Scalability/Profiler
  - Local
  - MN4
- EXTRAE
  - L3 cache misses
- Conclusions so far
- **Steps to follow**

Steps to follow

- **Short term**
  - Fix utests errors
  - Suggestion: move reordering parameter to input
  - Get numbers also with field solvers (benchmark_operators → genex)
    - Idea: test meshes of increasing sizes

Steps to follow

- **Short term**
  - Fix utests errors
  - Suggestion: move reordering parameter to input
  - Get numbers also with field solvers (benchmark_operators → genex)
    - Idea: test meshes of increasing sizes

- **Long term**
  - Heuristic for:
    - number of multigrid levels (mesh size dependent)
    - reordering vector

EUROfusion

GENE-X

Carlos Romero Madrid

HPC ACHs | 15-16 November 2023

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.