

Experiences with OpenMP on ScaleMP @ RWTH Aachen University

Christian Terboven, Dirk Schmidl, Dieter an Mey

{terboven, schmidl, anmey}@rz.rwth-aachen.de

Center for Computing and Communication
RWTH Aachen University, Germany

Agenda


- FIRE: Image Retrieval Engine
- SHEMAT-Suite: Geothermal Simulation Package
- Closing Remarks

FIRE: Flexible Image Retrieval Engine



- Image Retrieval: A set of query images is compared with all images in a (huge) database and the k most similar images are returned
 - Performance comparison of common features on different databases
 - Analysis of correlation of different features
 - Nested Parallelization – Outer level: Queries, Inner level: Comparison

- **Data Mining is well-suited for Shared-Memory parallelization, but hard with MPI**

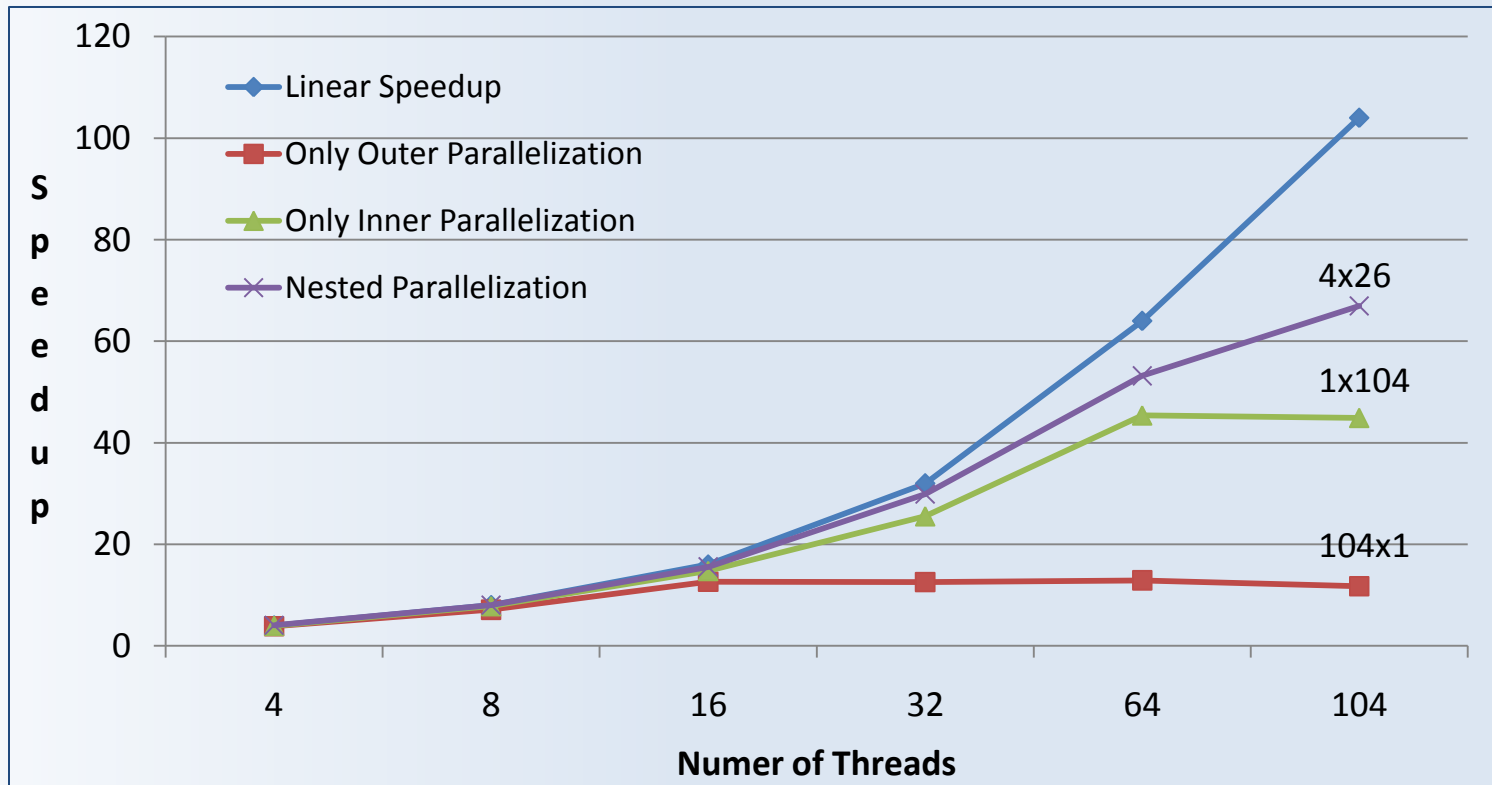


*Thomas Deselaers and Daniel Keysers,
RWTH I6: Chair for Human Language
Technology and Pattern Recognition*

3

FIRE: Scalability on a ScaleMP System

- Nested Parallelization improves speedup by reducing the total overhead. Best effort: Speedup of 66.96 on 104 cores.



4

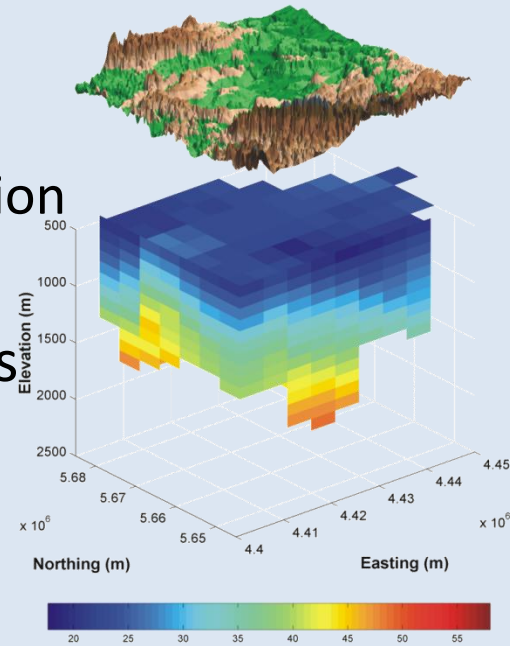
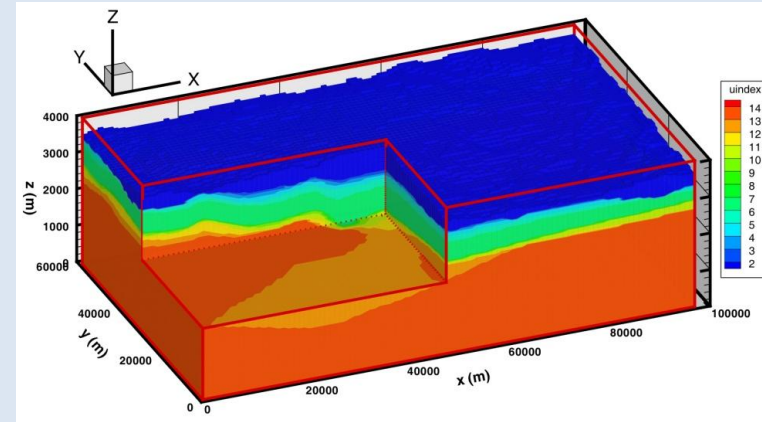
- Our system: 13 boards à 2x Intel Xeon E5420 (2.50 GHz)
- Explicit Thread Binding: `export KMP_AFFINITY=scatter`

Agenda

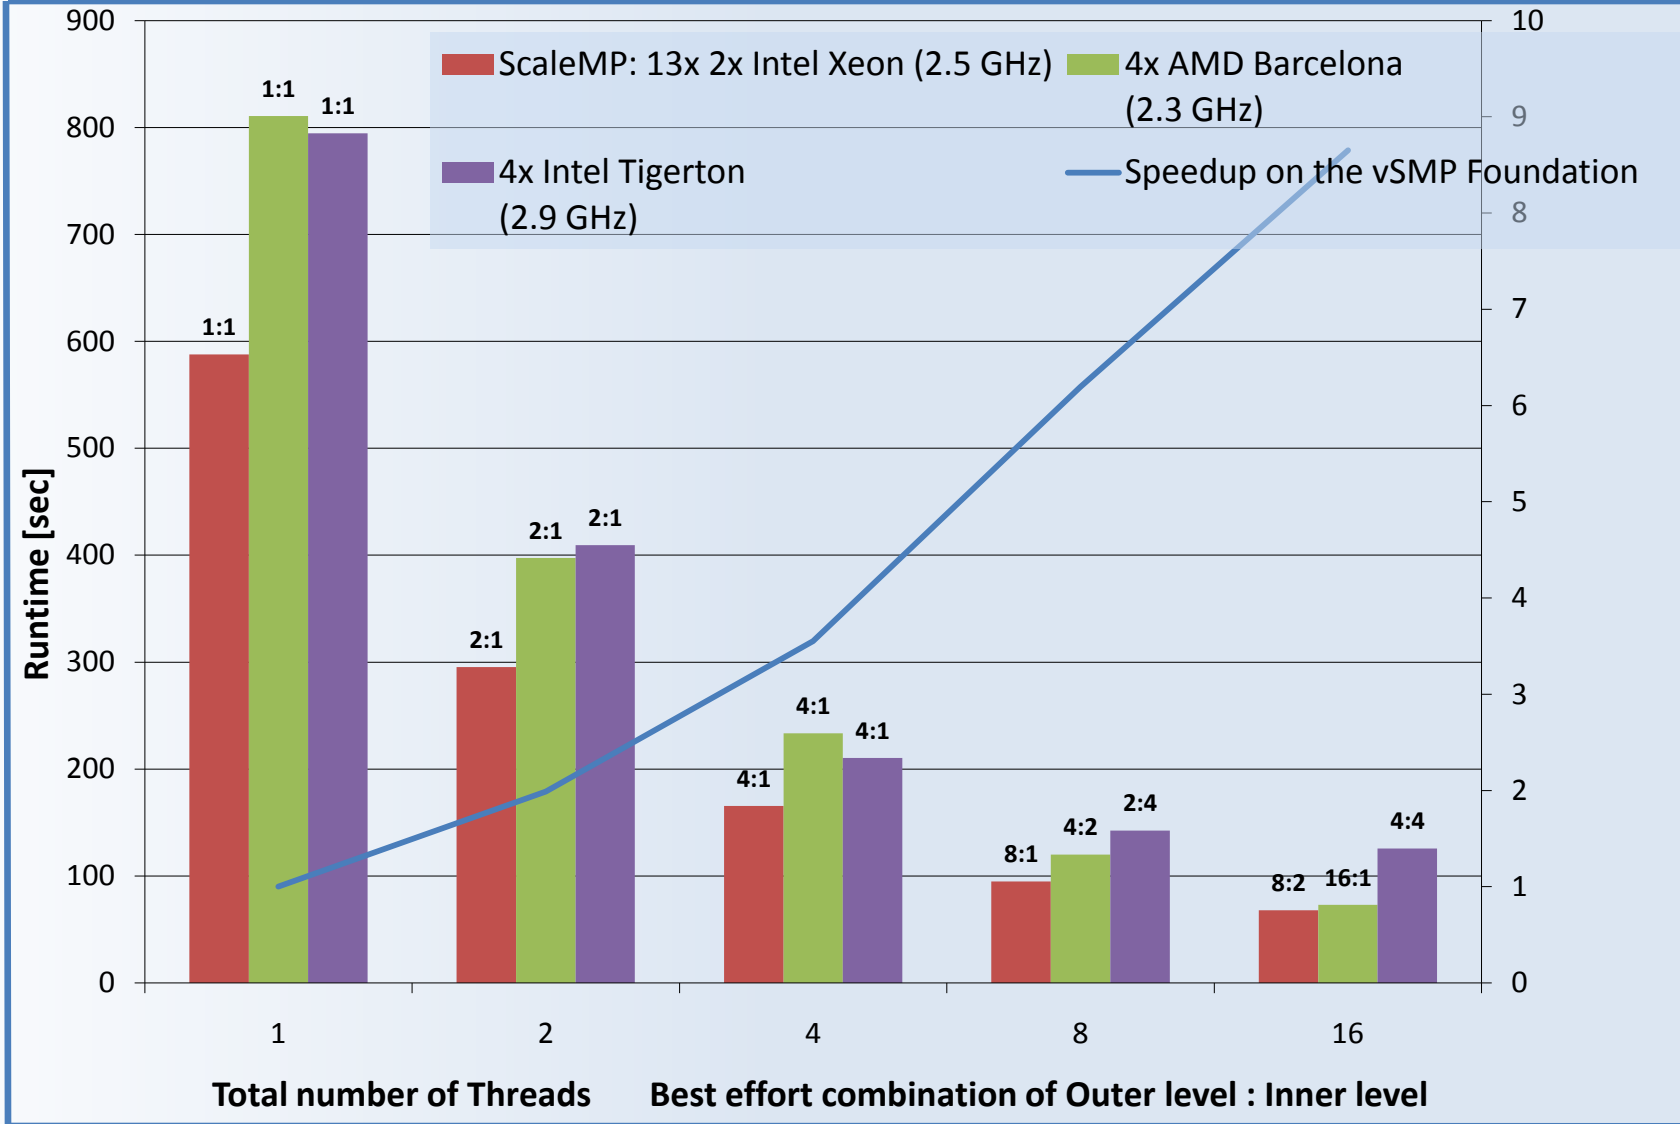
- FIRE: Image Retrieval Engine
- SHEMAT-Suite: Geothermal Simulation Package
- Closing Remarks

SHEMAT-Suite

- Geothermal Simulation Code
 - Simulating Groundwater flow, heat transfer and transport of reactive solutes
 - 3D Finite Difference Solver
 - Parameter estimation of rock properties using AD (Automatic Differentiation) Techniques → Shared-Memory Parallelization
- Outer Parallelization: Independent computations of the directional derivatives (Jacobian columns)
- Inner Parallelization: Solver based on BiCGStab with ILU(0)




SHEMAT-Suite: Machine Comparison (up to 16 threads)



7

SHEMAT-Suite: Speedup of Nested OpenMP on ScaleMP



∞

Center for

Computing and

Communication

○ Nested OpenMP improves scalability on many cores

Agenda

- FIRE: Image Retrieval Engine
- SHEMAT-Suite: Geothermal Simulation Package
- Closing Remarks

Remarks

- Limited application experience so far, still work in progress.
- What's still missing and subject to further research:
 - Experience with Performance Analysis Tools
 - Analysis of the data exchange between boards (over InfiniBand)
 - Analysis of and influence on the thread migration mechanism
 - Experience with Tuning Methodologies on ScaleMP
 - Does manual data „prefetching“ help?
 - Sometime the optimal thread binding is „surprising“: lack of binding support for nested OpenMP in the Intel compiler.

The End

Thank you for
your attention!

11