High-Performance Computing by and for Patient Specific Mechanical Properties

PASC23 Davos Switzerland 26–28 June 2023

HLRS

Johannes Gebert
Ralf Schneider
Michael Resch
High-Performance Computing Center Stuttgart, University of Stuttgart

HLRS High-Performance Computing Center Stuttgart

Mission of HLRS and its People

- Science in Performance, Networking and other HPC related topics.
- Supporting Research (~90%), SMEs and Industry (~10%) with HPC...
  - Know-How
  - Compute capacity
- Systems at the High-Performance Computing Center Stuttgart:
  - HLRS Hawk
    - 26 Petaflops x86
    - 5632 nodes with 2x AMD EPYC ROME 7742, 256 GB memory
    - 192 NVIDIA A100 GPUs
- NEC Cluster (Vulcan)
  - Heterogeneous Cluster

Research for and with supercomputing supports society.

Fig. 1: HLRS Hawk. A 26 Petaflops on-prem High-Performance Computer.

An Example Use Case

- Basic research on human trabecular bone.
- Computing effective or apparent stiffness tensors.
  - Supporting continuum mechanical simulations based on fine tuned stiffness information.

Fig. 2: The X-Ray scan of an original and an artificial hip joint.

Input data

- Computed Tomography (CT) datasets of various resolutions, e.g.:
  - Human femoral heads
  - Human proximal tibia
- Pre processed for:
  - Binary segmentation
  - Direct discretization
- Datasets are available via **DaRus**
  - The **data Repository** of the **University of Stuttgart.**

Biomechanical Challenge

- Trabecular bone consists of many small struts and plates.
  - Highly anisotropic
  - Living organ
  - Patient specific
  - Adapts to stress/strain from external loading.
- Microfocus Computed Tomography for imaging bones:
  - 0.005 - 0.015 mm/Voxel
  - Sufficient for computing effective stiffness tensors.
  - Radioactive doses prohibit scanning in situ.
- Clinical Computed Tomography for imaging in situ:
  - ~0.175-0.6 mm/Voxel
  - Resolution too low for computing stiffness tensors.

Fig. 3: A human femoral head in its pathology vessel.

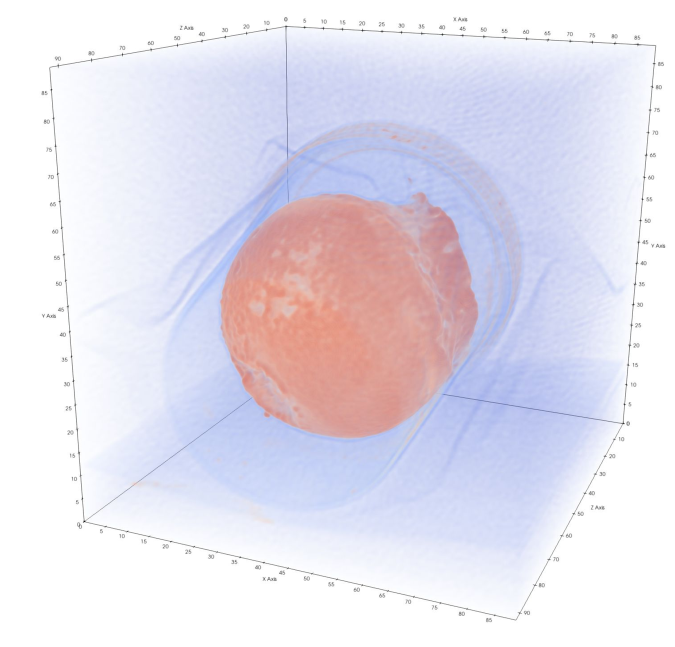

Fig. 4: The corresponding clinical Computed Tomography scan.

Software Stack

- Fortran
  - Pre processing
  - Direct Tensor Computation (DTC)
  - Initial post processing
  - Tensor optimizations
- Parallelization with MPI
- Solver by PETSc
- Mesh partitioning by Metis
- Bash/Python
  - Process steering
  - Analyzes

GitHub

Fig. 5: The microfocus computed tomography scan of the human femoral head with a resolution of 0.01495 mm/Voxel

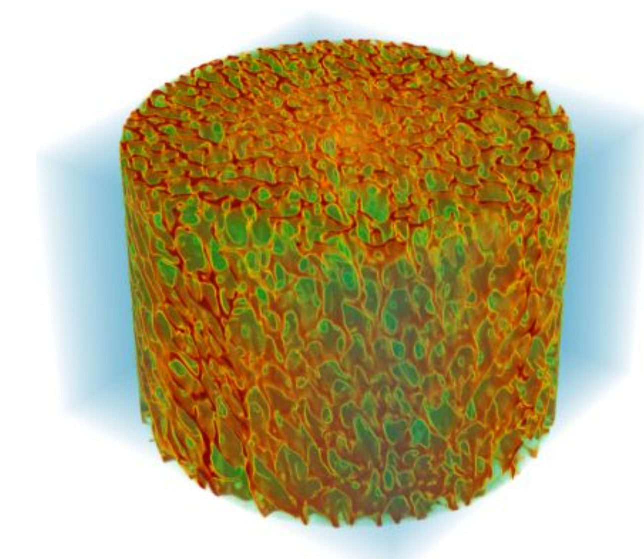

Fig. 6: Scans at the resolution of ~0.005 mm/Voxel show the struts and plates of trabecular bone.

Approach

- Decompositing the bone into cuboid subvolumes:
  - 0.6 mm, 1.2 mm, 2.4 mm, 4.8 mm or 9.6 mm.
- Direct discretization of the voxels with hexahedral finite *micro*-elements.
- Subvolume is considered and treated like a finite *macro*-element.
- Macro-element-stiffness matrix replicates effective anisotropy.

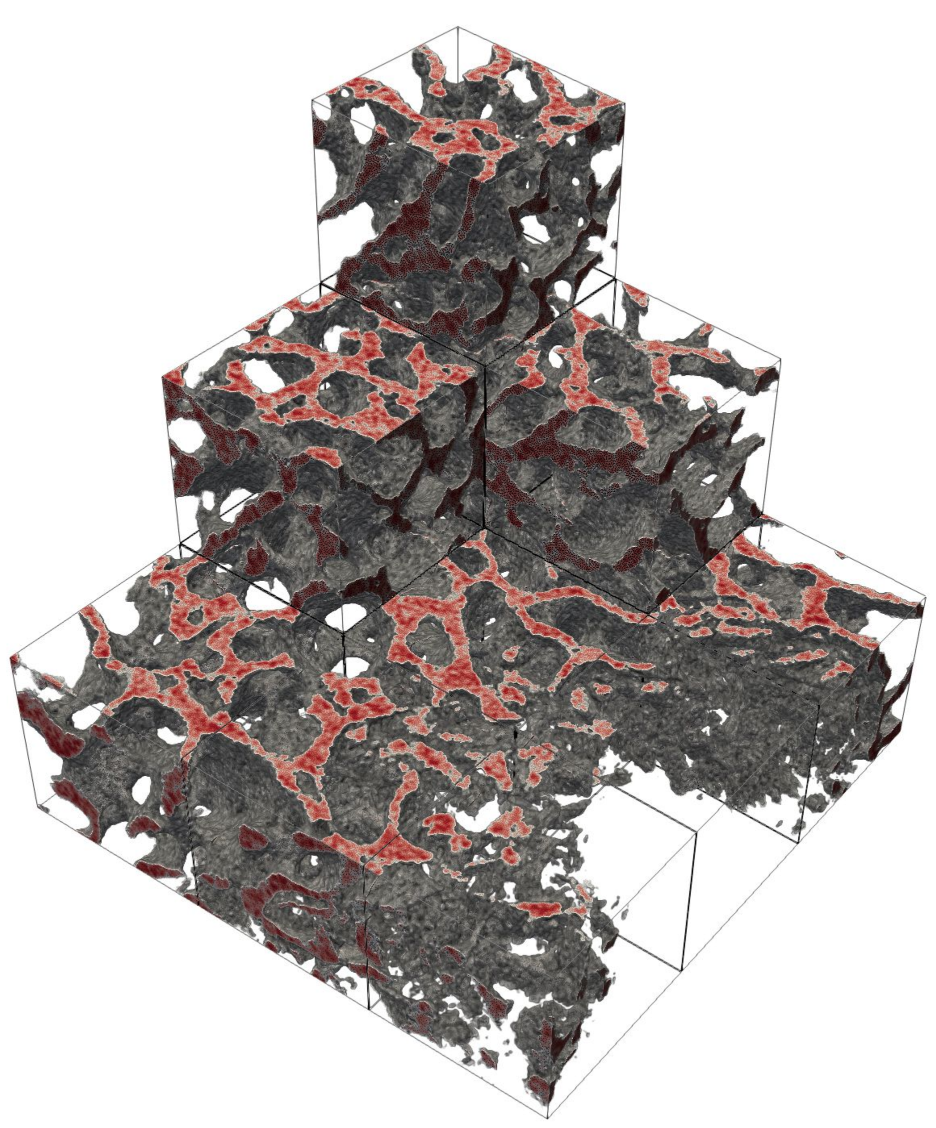

Fig. 7: Decomposition and direct discretization of the image into cuboid subvolumes.

Effective stiffness

- Stiffness → Resistance against deformation.
- A simple compression test of a structure results in the *apparent* stiffness.
- The direct tensor computation computes the effective macro elements stiffness.
  - Linear hexahedral: 20 dof, 20 load cases
  - Quadratic hexahedral: 60 dof, 60 load cases

The resulting element stiffness matrices can be mapped to continuum mechanical FEAs.

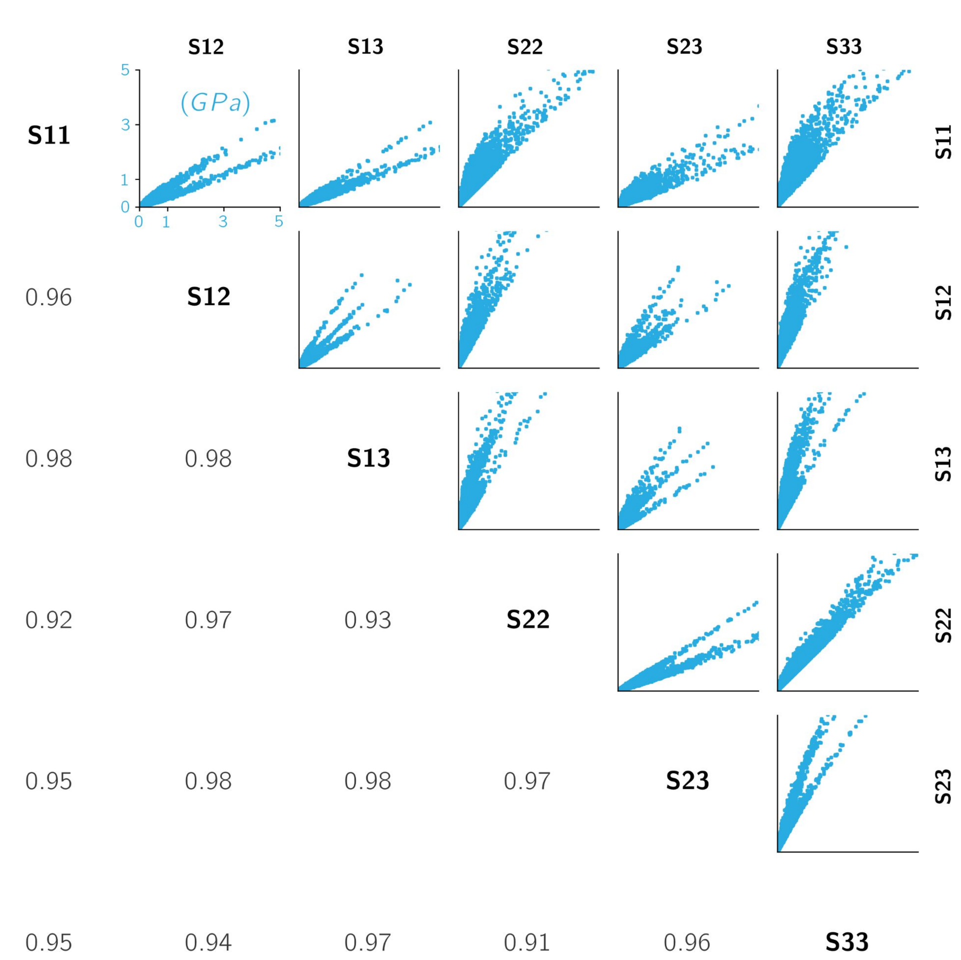

Fig. 8: The correlations of the first 6 orthotropic stiffness parameters of the decomposition size of 1.2 mm.

Profiling and Performance by Morphometry

- Direct discretization leads to:
  - Precise results
  - High computational effort
- Size of cuboid subvolumes fixed.
- Bone Volume/Total Volume is a morphometric measure.
  - Determination is computationally cheap.
- Profiling of the image for better load balancing is feasible.

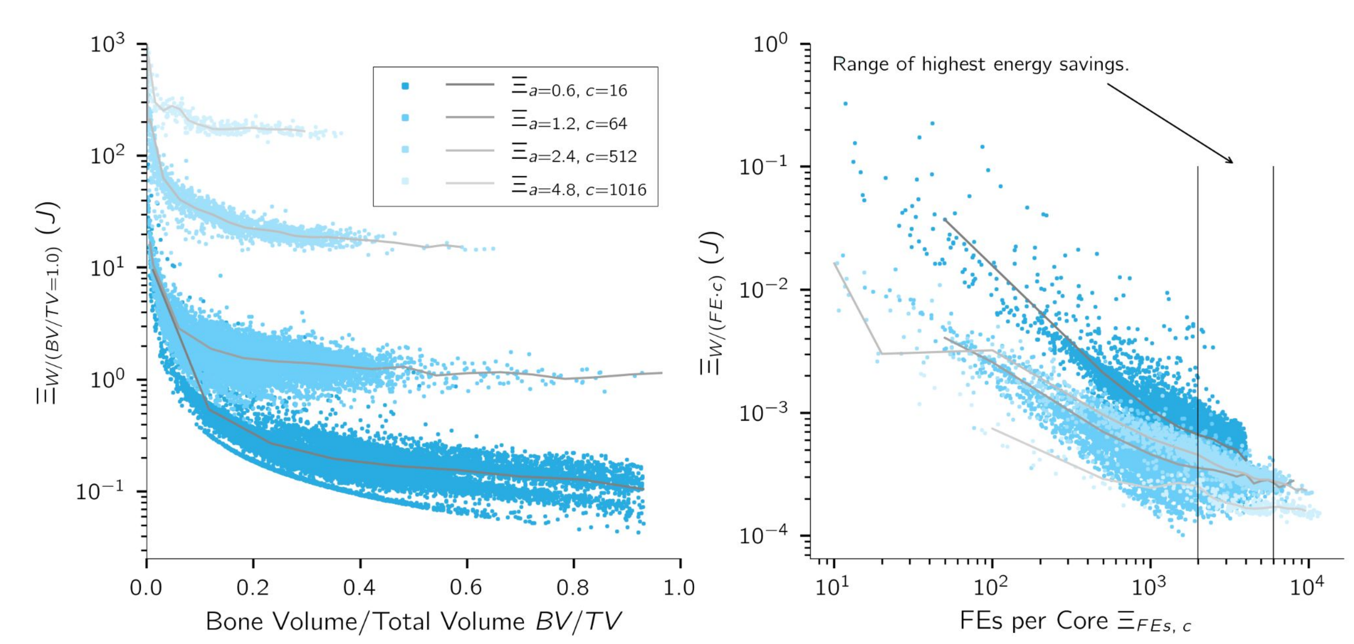

Fig. 9: Different serial and parallel parts of computing a single subvolume lead to a characteristic relationship between the Bone Volume/Total Volume (BV/TV) ratio and the energy usage.

Approach

- Computationally cheap analysis of the bones morphometry for ideal:
  - Load balancing
  - Scheduling
  - Energy efficiency

Implementation

- Step 1
  - Read the CT image.
  - Decompose into domains.
  - Calculate the morphometric quantity.
- Step 2
  - Compute the ideal distribution of domains to MPI communicators of ideal size.
  - Prepare the input data for DTC.
- Step 3
  - Run a modified version of DTC with non-uniform communicators.

Energy savings are expected around 5-25%, depending on the size of the domain.

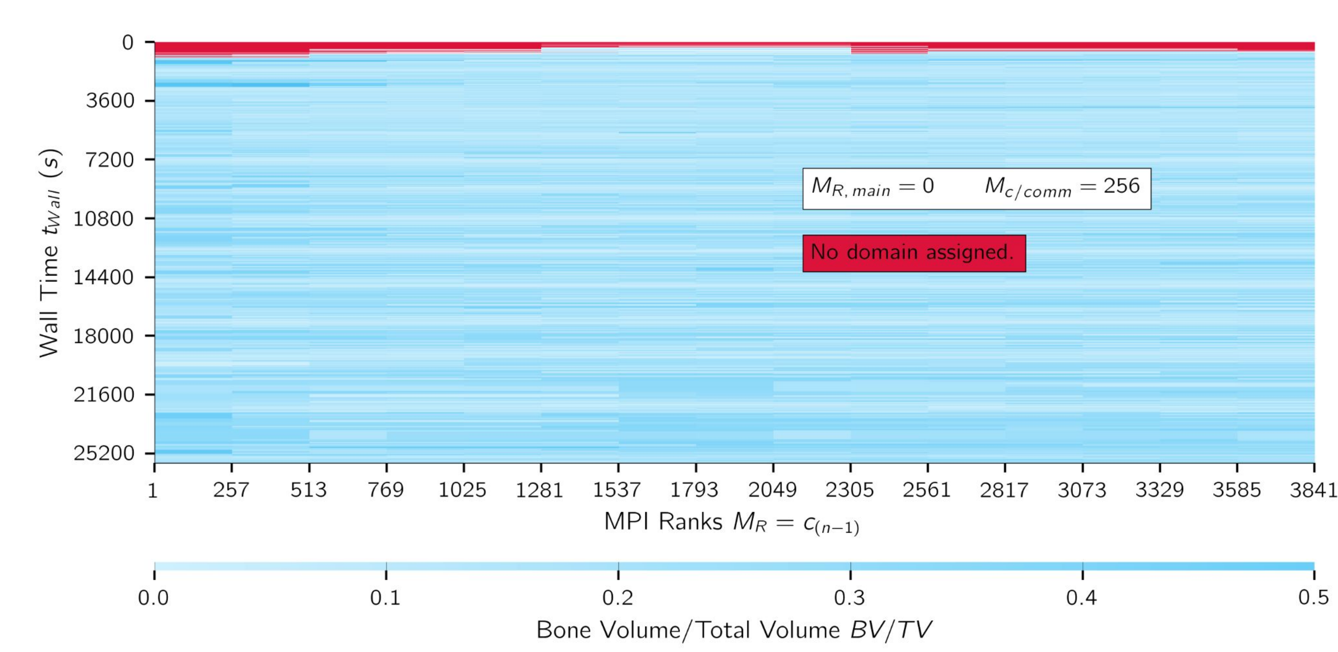

Fig. 10: The Direct Tensor Computation is started with a few to more than 100 000 cores to compute many domains in parallel. The main rank distributes the work packages to the worker communicators.