8th. World Congress on Computational Mechanics (WCCM8)
5th European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS 2008)
June 30 –July 5, 2008
Venice, Italy

Extreme Scalability Challenges in Analyses of Human Bone Structures

P. Arbenz[2], C. Flaig[2], G. H. van Lenthe[3], R. Müller[3] A. J. Wirth[3]
***C. Bekas[1], A. Curioni[1]**

[1] IBM, Research Division
Zurich Research Laboratory
{bek,cur}@zurich.ibm.com

[2] Department of Computer Science, ETH Zurich
arbenz@inf.ethz.ch
cflaig@ethz.ch

[3] Institute for Biomechanics
ETH Zurich
{ram,vanlenthe}@ethz.ch
awirth@ethz.ch

Key Words: *Micro Finite Elements, Osteoporosis, Blue Gene/L, Graph Partitioning*

ABSTRACT

High-resolution in vivo peripheral quantitative computed tomography (pQCT) provides detailed information on bone structure, see image below. The underlying voxel model admits to estimate the local bone density. The analysis of bone density (using other, more commonly available technology) is today's approach of predicting bone strength and fracture risk in diseases like osteoporosis that is, according to the WHO, second only to cardiovascular disease as a leading health care problem.

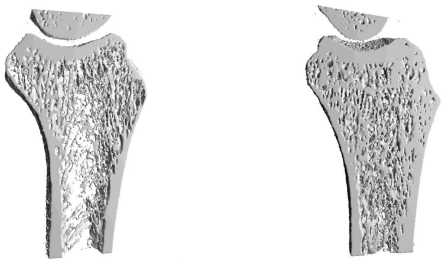

Figure 1: Bone specimens from human subjects. Left: Low density (osteoporotic) specimen from 72 year old male subject. Right: High density (normal) specimen from a 78 year old male subject.

Such a quantitative analysis of bone density does not take into account the microarchitectural structure of the bone. Coupling recent imaging capabilities with microstructural finite element (μFE) analysis offers a powerful means to determine bone stiffness and strength. It shows high potential to improve individual fracture risk prediction, a tool much needed in the diagnosis and treatment of osteoporosis. μFE models are created from CT scans by a direct voxel-to element conversion. The intricate microarchitectural structure of bone entails that these microFE models possess a very large number of elements and, by consequence, degrees of freedom. The computational model is based on linear elasticity and the linear solver is a matrix-free conjugate gradient algorithm preconditioned by aggregation-based AMG. The method is implemented in a software package called ParFE which is parallelized using MPI, and

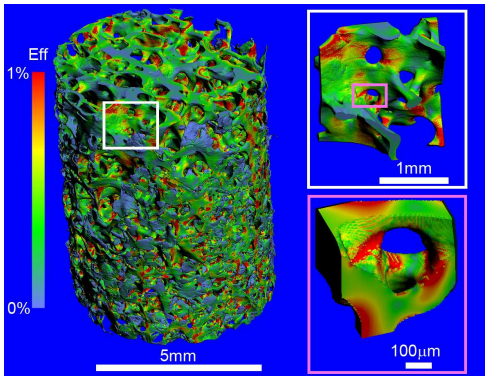

Figure 2: Effective strain on a human bone (vertebra) specimen.

is based on the public-domain software Trilinos, ParMETIS and HDF5. We targeted the IBM BG/L Supercomputer and improved both algorithms and implementation techniques in order to exploit its excellent scale-out potential. We conducted a study of human bone specimens resulting in very large sparse systems of up to about 1.5 billions unknowns. These runs always required less than half an hour, using up to 8 racks (8192 nodes) of the BG/L system at the T.J. Watson Research Center. Pre- and postprocessing of data took place at the Swiss National Supercomputing Centre, (CSCS) in Manno, Switzerland.

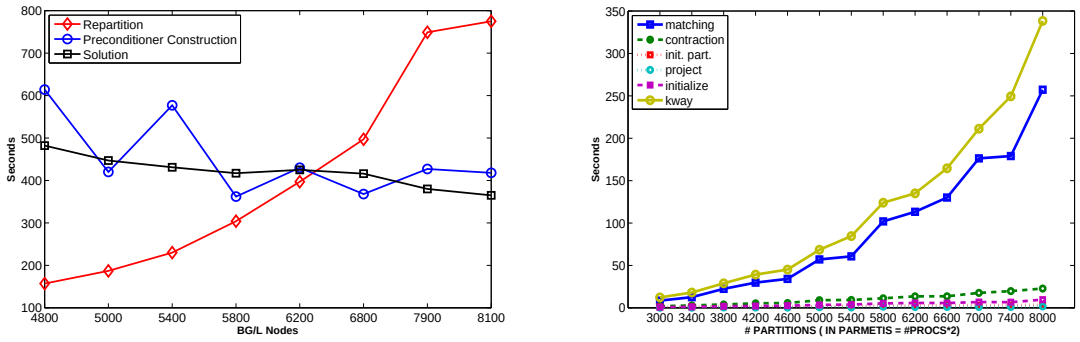

Figure 3: Left: Strong scalability test on the large bone specimen. Timings for the repartition (ParMETIS), construction of preconditioner and solution phase. Right: Timing analysis for repartitioning on a smaller bone sample.

We conducted the largest simulation of its kind so far (1.5 billions dof) and calculated the effective strain of a vertebral bone specimen (see Figures 2-3). This enables a highly detailed analysis of bone deformation under load, and calculation of bone stiffness and strength. Typical problem sizes (c.a. 200 millions dof) are solved in a few minutes. However, the major bottleneck in extreme scale-out of this application (and μFE applications in general) is in graph repartitioning: a) The intricate structure of bone causes significantly imbalanced partitions that have a strong negative effect when thousands of processors are used. b) The scalability of parallel graph partitioning tools (such as ParMETIS, Zoltan and others) on tens of thousands of processors appears to be a formidable task (see right view-graph of Figure 3). Clearly, in anticipation of the petaflop machines, efficient mapping of applications on millions of processing elements will require next generation algorithms and efficient mapping models.