Multi-level μ-Finite Element Analysis for Human Bone Structures

Peter Arbenz[1], Harry van Lenthe[2], Uche Mennel[1], Ralph Müller[2], and Marzio Sala[1]

[1] Institute of Computational Science, ETH Zürich, CH-8092 Zürich
[2] Institute for Biomedical Engineering, ETH Zürich, CH-8092 Zürich

Abstract. Osteoporotic fractures are a major cause of severe long-term pain and physical disability, and have an enormous impact on the individual, society and health care social systems. Osteoporosis is second only to cardiovascular disease as a leading health care problem. Since global parameters like bone density do not admit to predict the fracture risk, patients have to be treated in a more individual way. Today's approach consists of combining 3D high-resolution CT scans of individual bones with a micro-finite element (μFE) analysis. In this paper we investigate the efficient solution of the resulting large systems of linear equations by the preconditioned conjugate gradient algorithm. Emphasis is given on multilevel AMG preconditioners.

1 Introduction

Osteoporosis is a disease characterized by low bone mass and deterioration of bone microarchitecture. It leads to increased bone fragility and risk of fracture, particularly of the hip, spine and wrist. Worldwide, lifetime risk for osteoporotic fractures in women is estimated close to 40%; in men risk is 13% [4]; osteoporosis is second only to cardiovascular disease as a leading health care problem (World Health Organization). Osteoporotic fractures are a major cause of severe long-term pain and physical disability, and have an enormous impact on the individual, society and health care social systems.

With the advent of fast and powerful computers, simulation techniques are becoming popular for investigating the mechanical properties of bone. Using microstructural finite element (μFE) models generated directly from computer reconstructions of trabecular bone it is now possible to perform a 'virtual experiment', i.e. to simulate a mechanical test in great detail and with high precision.

Ideally, the development of a system with microstructural resolution better than 50 μm would allow in-vivo measurement of patients at different instances in time and at different anatomical sites. Unfortunately, such systems are not yet available, but the resolution at peripheral sites is approaching a level that allows elucidation of individual microstructural bone elements. Faster and more precise peripheral quantitative CT systems (pQCT) are now reaching the market, allowing for in-vivo patient measurements with an isotropic resolution better than 100 μm.

The resulting FE models are computationally demanding and require special solution schemes. The preconditioning conjugate gradient method is the obvious solution method. However, finding a proper preconditioners is not easy. Element-by-element approaches have been proposed about 10 years ago and are now common [5]. Their quality however deteriorates as the problems become large. For example, a convergent linear analysis of one human vertebral body requires over 130 million elements with a resolution of 30 microns.

The memory requirements for solving such a problem with the element-by-element preconditioned conjugate gradient (EBE-PCG) algorithm (that is presently employed by researchers at the Institute for Biomedical Engineering) is relatively small. However the convergence rate becomes excessively poor as the problem size increases. Adams et al. [1] suggested a smoothed aggregation multigrid preconditioner. The largest (i.e. most accurate) computations of this kind we know of are those reported by Adams et al. [1]. The number of degrees of freedom (unknowns) of the problems solved there in parallel on 4000 processors exceeded half a billion.

2 Mathematical models

The basic mathematical model of the problem is given by the Lamé equations of elasticity. The weak formulation of the 3D problem is: find the unknown displacements $\mathbf{u}$ such that

$$\int_\Omega [2\mu\, \varepsilon(\mathbf{u}) : \varepsilon(\mathbf{v}) + \lambda\, \nabla\cdot\mathbf{u}\, \nabla\cdot\mathbf{v}]\, d\Omega = \int_\Omega \mathbf{f}^T\mathbf{v}\, d\Omega + \int_{\Gamma_N} \mathbf{g}_S^T\mathbf{v}\, d\Gamma, \quad \text{for all } \mathbf{v}, \tag{1}$$

with the positive Lamé constants λ and μ, the symmetric strains

$$\varepsilon(\mathbf{u}) := 0.5\,(\nabla\mathbf{u} + (\nabla\mathbf{u})^T),$$

the volume forces $\mathbf{f}$, and the boundary tractions $\mathbf{g}$. (We omit details about function spaces and boundary conditions, see [7] for details.) For the problem under consideration, the computational domain Ω is extremely complicated consisting of the union of connected micro cubes. We discretize problem (1) by finite elements. The h-version of the finite element method is applied, which means that the related shape functions are lowest order piecewise polynomials. The FE discretization leads to a the linear algebraic system

$$K\mathbf{u} = \mathbf{f}, \tag{2}$$

where K is the related symmetric and positive definite stiffness matrix.

Since K is symmetric positive definite the preconditioned conjugate gradient method is applied for solving (2). The crucial question is now how to choose the preconditioner. We employ a AMG multilevel preconditioner based on smoothed aggregation. Such preconditioners require relatively little memory in addition to that required to store the linear system matrix K. We usually observed operator

complexities below 1.5. Second, these preconditioners are optimal, i.e. the number of iteration steps does not depend on the fineness of the discretization of the problem at hand and have been proven to work well on dirtributed memory parallel machines.

3 Numerical experiments

We report on numerical experiments conducted on two parallel machines.

- A 160 dual-processor high performance cluster system at ETH Zürich. Each node is configured with two AMD Opteron 250 processors with 2.4 GHz, 8 GB of main memory (shared by the two processors) and a 120 GB disk. The compute nodes are inter-connectedvia a two-layer Quadrics QsNet II network with a sustained bandwidth of 900MB/s and a latency below 2μsec.
- A Cray XT3 computer system consisting of 1100 compute processors. The CPUs are 2.6 GHz AMD Opteron processors with 2GB of main memory each.The processors are connected by the Cray SeaStar high-bandwidth (4 GB/s sustained), low-latency interconnect.

Our parallel C++ code is implemented with the Trilinos frame work [3]. The distribution of the mesh data, the matrix assembly, and the solution with the preconditioned conjugate gradient algorithm are fully parallel. At present, the disk I/O is still sequential.

Some results have been obtained on the Cray XT3 at CSCS in Manno. We have investigated algebraic multigrid preconditioners in structural analysis recently in [2]. We intend to incorporate the smoothed aggregation multigrid preconditioner ML [6] in the EBE-PCG solver of the Institute for Biomedical Engineering.

We performed a scalability test with an artificial problem, a solid reactangular beam the length of which is adjusted to the number p of processors used. For a problem consisting of $40 \times 40 \times 20p$ cubes the execution times of the Jacobi preconditioned conjugate gradient algorithm with element-by-element matrix-vector product [8] are given in Table 1. The last column of this table indicates

p	total time	solver time	iteration count	time/iteration
1	18.03	11.56	318	0.0364
2	25.50	16.69	419	0.0398
4	37.24	24.98	581	0.0430
8	66.30	45.39	925	0.0491
16	98.86	64.77	1616	0.0401
32	195.11	127.42	3085	0.0413
64	451.78	302.91	6001	0.0505

Table 1. Scalable beam test results. Execution times (in seconds) for Jacobi preconditioned cg algorithm to reduce the residual by a factor of 10^8.

that a *single* step of pcg with Jacobi preconditioning scales very well. However, as the iteration count grows with the problem size the solution process as a whole does not scale. The advantage of this simple solver is the minimal storage requirement. The element-by-element matrix-vector multiplication requires negligible memory space. The diagonal of A for the preconditioner and a few vectors for the conjugate gradient algorithm are all that is needed memory-wise.

To lower the iteration count, we employ the smoothed-agregation AMG preconditioner of Trilinos [6]. We show that, although the memory consumption increases significantly (to 1GB/200'000dofs), the execution time is reduced tremendously.

References

1. Adams, M.F., Bayraktar, H.H., Keaveny, T.M., Papadopoulos, P.: Ultrascalable implicit finite element analyses in solid mechanics with over a half a billion degrees of freedom. In: ACM/IEEE Proceedings of SC2004: High Performance Networking and Computing. (2004)
2. Arbenz, P., Hetmaniuk, U.L., Lehoucq, R.B., Tuminaro, R.: A comparison of eigensolvers for large-scale 3D modal analysis using AMG-preconditioned iterative methods. Internat. J. Numer. Methods Engrg. **64** (2005) 204–236
3. Heroux, M.A., Bartlett, R.A., Howle, V.E., Hoekstra, R.J., Hu, J.J., Kolda, T.G., Lehoucq, R.B., Long, K.R., Pawlowski, R.P., Phipps, E.T., Salinger, A.G., Thornquist, H.K., Tuminaro, R.S., Willenbring, J.M., Williams, A., Stanley, K.S.: An overview of the Trilinos project. ACM Trans. Math. Softw. **31** (2005) 397–423
4. Melton, III, L.J., Chrischilles, E.A., Cooper, C., Lane, A.W., Riggs, B.L.: Perspective. How many women have osteoporosis? J. Bone Miner. Res. **7** (1992) 1005–1010
5. van Rietbergen, B., Weinans, H., Huiskes, R., Polman, B.J.W.: Computational strategies for iterative solutions of large FEM applications employing voxel data. Internat. J. Numer. Methods Engrg. **39** (1996) 2743–2767
6. Sala, M., Gee, M., Hu, J., Tuminaro, R.S.: ML 4.0 Smoothed Aggregation User's Guide. Tech. Report SAND2004-4819, Sandia National Laboratories (2005)
7. Schwab, C.: The p and hp Finite Element Methods. Oxford University Press, Oxford (1998)
8. Smith, I., Griffiths, D.: Programming the Finite Element Method. 4th edn. Wiley, Chichester (2004)