

Computational Bone Mechanics: From the Cloud to an Orthopedists Mobile Device

Z. Yosibash, K. Myers, and Y. Levi

Abstract Verified and validated simulations of the mechanical response of femurs, based on CT scans, have been recently presented. These simulations, based on high-order finite element methods (p -FEMs), may be used to diagnose the risk of fracture when used in clinical orthopedic practice. The first part of this chapter describes the methods used to create p -FEM models of patient-specific femurs and the in-vitro experiments used to assess the validity of the simulation results.

Having demonstrated that p -FEMs using CT-scans allow personalized assessment of the risk of fracture of a given femur, we envision a simulation center that may serve a large community of orthopedic doctors, each submitting a CT-scan to be analysed. In such a situation, multiple p -FE models must be simultaneously generated, solved and the verified FE results must be returned to the interested orthopedists. Because such analyses are sought in a short time-scale, and the interaction with the orthopedists will be by mobile devices in the future, we present the implementation of a multi-threaded p -FE solver in the second half of the chapter that will perform multiple solutions simultaneously.

1 Introduction

Quantitative computer tomography (qCT) and advanced computational mechanics algorithms (high order FEMs) have emerged in the past two decades that allow personalized finite element analyses of patient femurs. Such FE analyses require:

- (a) an accurate description of the femur's geometry,
- (b) an accurate description of the material inhomogeneous properties,
- (c) accurate physiologic boundary conditions (loads and constraints),

Z. Yosibash (✉) • Y. Levi

Department of Mechanical Engineering, Ben-Gurion University, Beer-Sheva, Israel

e-mail: zohary@bgu.ac.il

K. Myers

PerSimiO Inc., Beer-Sheva, Israel

e-mail: kent.myers@persimio.com

© Springer International Publishing Switzerland 2015

M. Mehl et al. (eds.), *Recent Trends in Computational Engineering - CE2014*,

Lecture Notes in Computational Science and Engineering 105,

DOI 10.1007/978-3-319-22997-3_14

- (d) a verification process assuring that the numerical errors are bounded by a specified tolerance,
- (e) a validation process in which the FE results are compared against all measurable data (strains and displacements) recorded during simplified in-vitro tests on a variety of femurs (with a wide spread in age, gender and weight).

An accurate description of a patient-specific femur's geometry, and its material inhomogeneous properties can be obtained from the qCT scans as detailed in [5, 12, 13, 16]. Automatic procedures can thereafter be applied to generate a p -FE mesh [10]. The femur may be well approximated by the linear elastic theory with inhomogeneous properties, so a linear FEA can be easily performed once the FE mesh, material properties and boundary conditions are prescribed. Personalized boundary conditions are far more difficult to assess. Such boundary conditions can be either computed by reverse engineering after a gait analysis is performed on the specific patient (this is however impossible to be obtained if the patient has a fracture already in a femur and needs an analysis of the optimal implant), or by a population stochastic analysis as recently presented in [20].

We use p -FEMs because these have many advantages over conventional h -FEMs when applied to the analysis of bones: p -FEMs accurately represent the bone's surfaces by using blending-function techniques, the mesh is kept unchanged and only the polynomial degree p of the shape functions is increased systematically to achieve convergence and p -FEMs allow natural functional variation of the material properties within each element. In addition, p -elements are much larger, may be far more distorted, have larger aspect ratio, and produce considerable faster convergence rates compared to their h -FEM counterparts.

Once the p -FE analyses are finalized and the results are verified so the numerical error is assured to be controlled, one must validate the results by a set of in-vitro experiments on a large cohort of healthy and diseased human fresh-frozen femurs. The validation process is successful only after one compares the experimental and the predicted results—both strains and displacements. To this end, the FE results were validated by in-vitro tests on the largest (so far reported) cohort of 17 fresh-frozen femurs [15] and 14 pathological femurs with metastatic tumors [19]. The verification and validation process and the different experiments are detailed in Sect. 2.

After the simulation has been verified and validated, efficient and automatic algorithms must be available for daily orthopedic usage. A semi-automatic algorithm for the generation of a p -FE femur model from clinical qCT scans was implemented and an efficient FE solver that may solve many models simultaneously is being developed. The p -FE code is written using the Java object-oriented programming language [8] because:

- It is platform independent so it will run on any Java Desktop/Android environment.
- It has a very mature class library of sophisticated data structures, and a broad selection of open source code available that address frequently encountered programming requirements.

- It has automatic garbage collection so there is no need to be concerned with freeing memory.
- It is relatively easy to integrate native libraries written in C++, FORTRAN, etc. for individual computational platforms (Windows, Linux) using Java Native Interface (JNI) wrappers.
- Development and prototyping is more rapid than using explicitly compiled and linked languages like C++. One can easily change code on the fly and continue development without having to stop the execution to recompile and link.
- Developing multi-threaded applications is pretty simple.
- Performance rivals more traditional languages.

We highlight features of the FE code where algorithm efficiency can be compared with a similar p -FE commercial product.

2 Personalized p -FE Models of Femurs from qCT-Scans

For any specific patient that undergoes a qCT scan of a femur, the data is semi-automatically evaluated by in-house Matlab programs. Exterior, interface and interior boundaries are detected representing different boundaries of a given slice. These slices are combined to form a surface representation of the femur and subsequently a solid model. Large surfaces are generated, which are essential for the p -mesh generator. The resulting 3D solid is thereafter auto-meshed by tetrahedral high-order elements. The entire algorithm (qCT to FE) is schematically illustrated in Fig. 1.

The femur's response under physiological loading is well described by the linear theory of elasticity. Although the bone on the macroscopic level is orthotropic, under stance position excellent results are obtained using isotropic inhomogeneous relations (see [13] and references therein). The inhomogeneous properties depend on the local density measure at each and every point within the bone and are determined by the qCT scan. K_2HPO_4 phantoms are placed near the femur of interest and are used to correlate the known mineral density and HU.

$$\rho_{K_2HPO_4} [gr/cm^3] = 10^{-3} \times (0.8072 \times HU - 1.6) \quad (1)$$

The bone ash density ρ_{ash} is thereafter determined based on empirical connections [9], using the connection between hydroxyapatite and K_2HPO_4 phantoms [2]:

$$\rho_{ash} [gr/cm^3] = 0.877 \times 1.21 \times \rho_{K_2HPO_4} + 0.08 \quad (2)$$

The relation reported in [4] includes specimens with a wide density range ($0.092 < \rho_{ash} < 1.22 [g/cm^3]$) while the relation reported in [6] was obtained using ash densities $< 0.3 [g/cm^3]$. The ρ_{ash} threshold between cortical and trabecular tissues is unclear, however all the pixels having HU number larger than 475 are considered

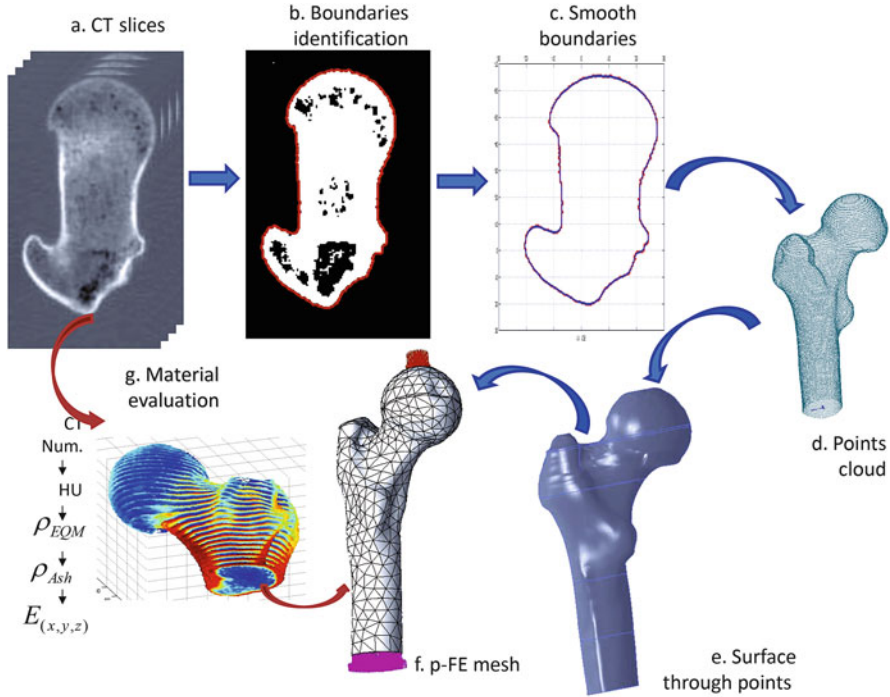


Fig. 1 Schematic flowchart describing the generation of the p -FE model from qCT scans. (a) Typical CT-slice, (b) Contour identification, (c) Smoothing boundary points, (d) Point cloud representing the bone surface, (e) Bone surface, (f) p -FE mesh and (g) Material evaluation from CT data. Figure from [14]

cortical bone. $HU = 475$ leads to $\rho_{ash} = 0.486$ using (1) and (2) based on previous publications [3, 17, 18]. Thus, the following relations were used to determine Young's modulus from ρ_{ash} :

$$E_{cort} = 10200 \times \rho_{ash}^{2.01} [MPa], \quad \rho_{ash} \geq 0.486 \quad (3)$$

$$E_{trab} = 2398 [MPa], \quad 0.3 < \rho_{ash} < 0.486 \quad (4)$$

$$E_{trab} = 33900 \times \rho_{ash}^{2.2} [MPa], \quad \rho_{ash} \leq 0.3 \quad (5)$$

Young's modulus at the transition area between cortical and trabecular bone tissue ($0.3 < \rho_{ash} < 0.486$) was set to $E = 2398 [MPa]$, based on the data reported in the literature. Poisson ratio was set to $\nu = 0.3$.

Having available the femur's geometry and material properties, boundary conditions have to be specified on the FE model according to the activity of interest. Many past studies have been performed to measure physiological loads (also in-vivo) during well monitored activities. Based on these works, models that are intended to predict the hip contact load force on femur's head for example are available (see

e.g. [20] and the references therein). However, to first verify and validate the p -FE models, we consider deterministic forces applied to femurs in in-vitro experiments to simulate a stance position loading, and wish to demonstrate that these FE models predict the experimental observations well.

The p -FE results are being verified to ensure that the numerical error is under a specific tolerance. Convergence is realized by keeping a fixed mesh and increasing the polynomial degree of the approximated solution p until the relative error in energy norm is small, and the strains at the points of interest converge to a given value—see for example Fig. 2. Each femur FE model consists of between 3500 and 4500 elements ($\sim 150,000$ degrees of freedom (DOFs) at $p = 4$ and $\sim 300,000$ DOFs at $p = 5$).

The validation of the p -FE simulations is performed by comparing them to in-vitro biomechanical experiments. The validation process was demonstrated by experiments on 17 healthy fresh-frozen human cadaver femurs and 14 fresh frozen femurs with anticipated metastatic tumors. In these experiments a simple stance position configuration was considered in which the femurs were loaded through their head at different inclination angles (0, 7, 15 and 20 degrees). Vertical and horizontal displacements of the femur head and strains along the femur surface were measured (see for example the test setting and the corresponding FEA in Fig. 3).

A total of 102 displacements and 161 strains on the 17 femurs were used to assess the validity of the p -FE simulations. In Fig. 4 the computed FE strains and displacements are compared to the experimental observations.

Remark 1 Note that for 12 of the 17 femurs a blind comparison was performed, i.e. the group that performed the experiments did not know the FE results, the experimental results were not known by the group that performed the analysis.

One may notice the match between the predicted and measured data for femurs under stance position loading: the slope and R^2 of the linear regression are very close to 1, and the average error in the Bland-Altman is zero.




Fig. 2 Convergence in energy norm and ϵ_{zz} at a representative point of interest in a femur




Fig. 3 (a-left) Sketch of the frontal plane of an instrumented left femur. (b-right) Experimental setup with the optical markers on an instrumented left femur and its corresponding deformed (magnified) FE model. Figure from [11]




Fig. 4 Comparison of the computed strains (*plus* and *asterisk* symbols) and displacements (*open circles*) to the experimental observations normalized to 1000 N load. (Left) Linear regression, (Right) Bland-Altman plot. Figure from [15]

Having demonstrated that verified and validated *p*-FE analyses are obtained for healthy femurs, one is naturally interested in pathological bones. For example, how well will the risk of fracture be predicted in femurs with metastatic tumors? To address this question, radiographs of fresh frozen proximal femur specimens from donors that died of cancer were examined, and seven pairs diagnosed by an orthopedic surgeon to contain metastatic tumors were identified. These femurs were CT-scanned, instrumented by strain-gauges and loaded in stance position

at three inclination angles. Finally the femurs were loaded until fracture that usually occurred at the neck. Histopathology was performed to determine whether metastatic tumors are present at fractured surfaces. Figure 5 presents all femurs and the fracture patterns.

Following each experiment, p -FE models were created from the CT-scans in an effort to reproduce the mechanical experiments by simulation. The predicted displacements, strains and yield loads were compared to experimental observations. The principal strain E_3 at 1000 N computed for the 14 femurs at 0° is shown for example in Fig. 6. The simulated strains and displacements correlated well with the experimental observations, with a linear regression slope of 0.95 and a coefficient of regression $R^2 = 0.967$ (see [19] for details).

A simplified yield strain criterion, shown to predict the yield of healthy fresh frozen femurs reasonably well in [18], was used to estimate the yield of the cancer affected femurs. This criterion estimates the yield initiation to occur at the location where the largest principal strain (by a linear elastic analysis) on the bone's surface reaches a critical value of 7300 μ strains in tension or $-10,400$ μ strains in compression (reported in [1]). The principal strains on the femur's surface were computed for an applied 1000 N. The ratio between the critical strain in tension (respectively compression) to the maximum (respectively minimum) computed principal strain times 1000 N was determined as the predicted yield force. Because pointwise values of FE strains may contain large numerical errors, we used instead an averaged value along a part of an element edge adjacent to the maximum strain location.

Figure 7 shows the linear regression of the yield load predicted by the p -FEA and estimated in the experiments, and the associated Bland-Altman graph. In these graphs the FFM1 femurs are excluded because they fractured at a very high load in a creep-like mode.

One may notice the relatively good correlation obtained between the predicted yield load and the experimental observed yield with a linear regression slope of 0.80 and a coefficient of regression $R^2 = 0.78$. This is considered a good correlation since only a simplified and naive fracture prediction was considered.

We may summarize by concluding that personalized FE models of femurs were demonstrated to predict the mechanical response very well. A simplified yield criterion based on the computation of principal strains was demonstrated to predict the yield force in most of the cases, especially for femurs that failed at small loads. In view of the limited capabilities of present tools to predict risk of fracture in femurs with metastatic tumors, the p -FE methodology may be very valuable in making clinical decisions.

3 A Java Based p -FE Code

Imagine an orthopedic surgeon having for example to decide on the best implant configuration for a given patient that needs a total hip replacement. We envision



Fig. 5 Fracture patterns in femurs with metastatic tumors. Figure from [19]

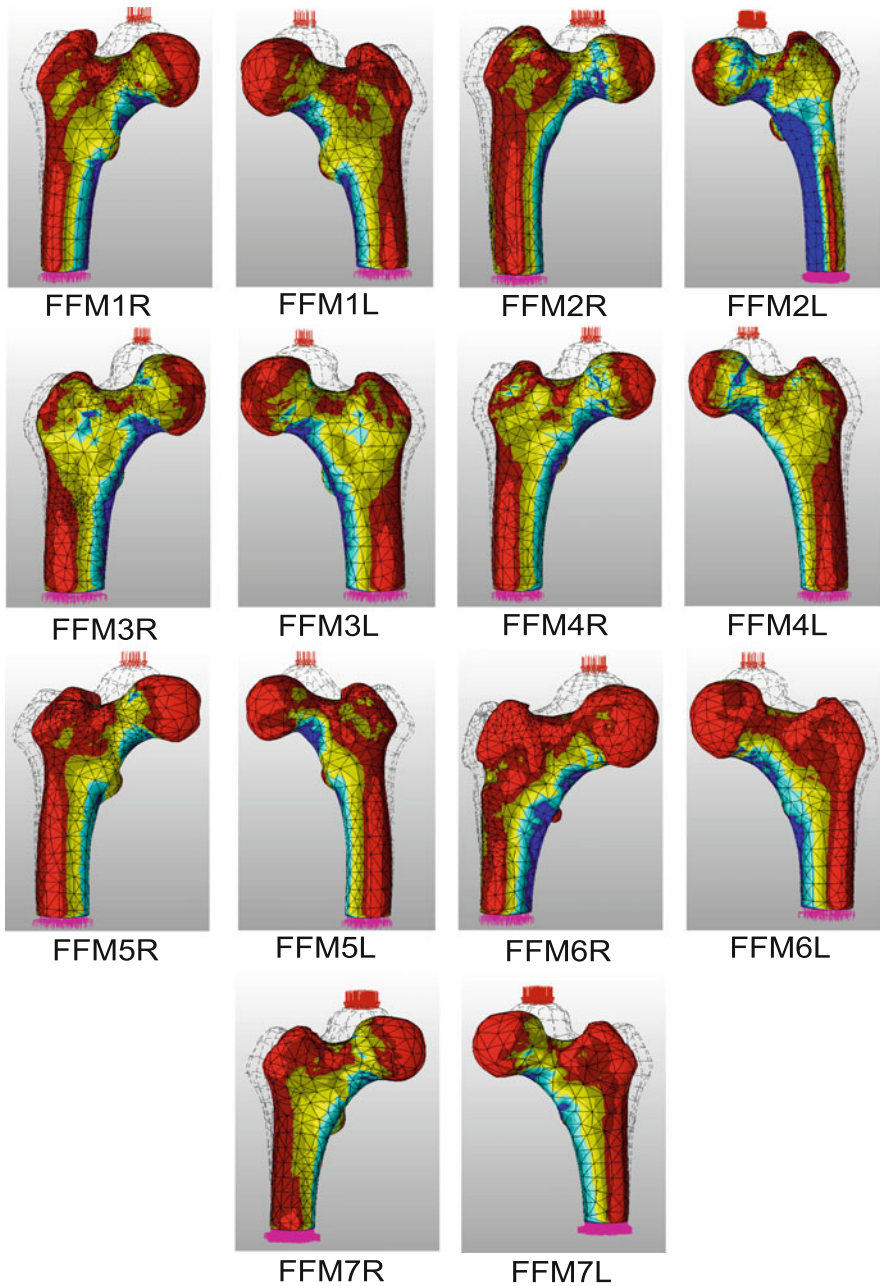


Fig. 6 Principal strain E_3 at 1000 N load computed by p -FE analyses for the 14 femurs at 0° inclination (colors not at same scales for all models). Figure from [19]




Fig. 7 *Left*—Linear correlation for yield load excluding FFM1. *Right*—Bland-Altman error plot for yield load excluding FFM1. Figure from [19]

a system in which a surgeon utilizes a mobile application to select the name of a patient, ensure that CT scans are available for the patient, and subsequently select the best available implant based on the physical characteristics of the selected patient. The surgeon's selections and the patient's CT scans are transferred to a cloud application that will automatically construct the p -FE model of the patient's bone. The bone model will be automatically revised to include each available implant, resulting in several different models. Verified FE analyses are performed, so that the polynomial degree is increased over all elements until a converged solution is obtained to determine the implant that results in the minimum strain-shielding. The cloud application sends the results of the simulation to the orthopedic surgeon's mobile device, with a ranking of relative suitability of each implant. In Fig. 8 this procedure is illustrated. The Java programming language was selected for implementation of the cloud application due to its inherent cross-platform compatibility, mature object oriented class library, and ease of developing multi-threaded applications. Fig. 9 illustrates the system's parallel configuration.

Because p -FE models produce large elemental stiffness matrices an efficient wavefront numbering system was required in order to minimize the memory footprint of the global system of equations. Fig. 10 illustrates by colors the increase in the element numbers along the length of a femur, resulting from the wavefront algorithm. The size of each element stiffness matrix stored was reduced by saving only the non-zero terms of the condensed matrices. As an example, we show in Fig. 11 the sparsity percentage (number of zeroes in the upper triangular stiffness matrix compared to all upper triangular matrix) for a 1000 ($10 \times 10 \times 10$) FE straight-sided hexahedral mesh for increasing p value (the FE mesh is shown in Fig. 12).

One may notice that by retaining only the non-zero terms of the global stiffness matrix, less than 1% of the terms need be retained in memory during the back-substitution phase.

Another characteristic of the p -FE methodology is the lengthy computation time for elemental stiffness matrices (at $p = 8$ for example, and using the trunk space, there are 576 shape functions per element). We show in Fig. 13 the CPU time required to compute a hexahedral elemental stiffness matrix as a function of the




Fig. 8 An illustration on the use of the *p*-FE system by orthopedic surgeons

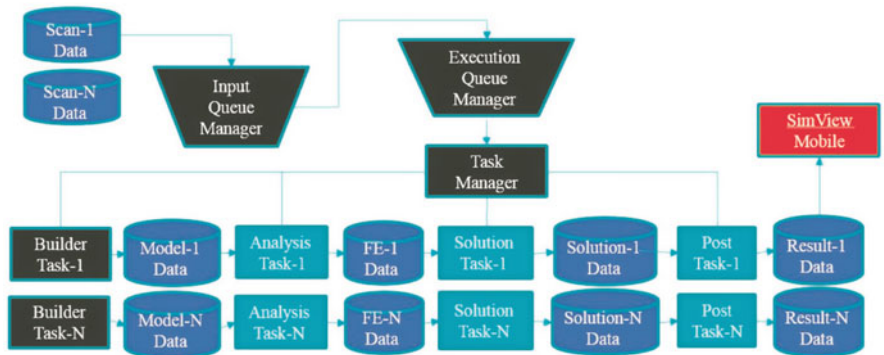


Fig. 9 Multiple analyses performed in parallel—the flowchart

p-level (all computations are performed on a DELL T3500 with a single Intel Xeon E5-1620 3.6 Ghz processor with 4 cores (8 hyperthreads) and 20 GB RAM). One notices the rapid increase in the CPU time as the polynomial degree increases beyond $p = 6$. Because a significant portion of the stiffness matrix computation for each element is the same when material properties are identical, we pre-compute auxiliary arrays once at the start of the simulation, and reuse these precomputed arrays during the computation of the individual element matrices. We also compute the individual element stiffness matrices and load vectors in parallel. For the solution of the system of equations, the parallel C++ solver Pardiso [7] was used by providing a Java Native Interface (JNI) wrapper.




Fig. 10 The numbering sequence of a proximal femur—*different colors* represent the element numbering intervals




Fig. 11 Sparsity percentage as a function of the p -level (DOFs) for a 1000 FE mesh

At this stage of development, the FE system does not yet perform matrix assembly in parallel, and does not take full advantage of the capabilities of the Pardiso solver to perform computations on distributed memory parallel systems. However, the p -FE Java code is faster when compared to the commercial p -FE code




Fig. 12 A 1000 straight sided hexahedral p -FE mesh having 228,723 DOFs

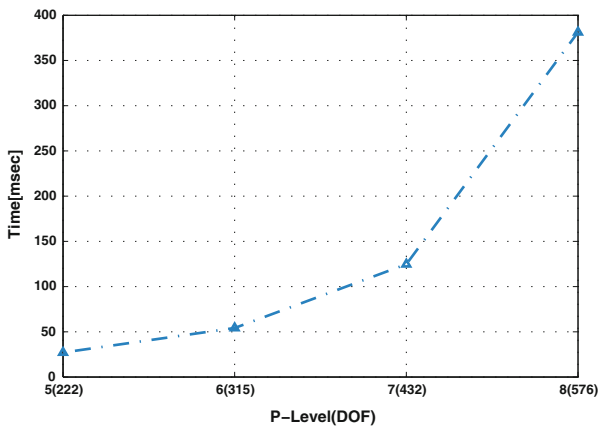


Fig. 13 CPU time for the computation of the elemental stiffness matrix of a hexahedral straight sided element vs. the p -level

StressCheck.¹ We compared performance for the 1000 element model depicted in Fig. 12, where traction boundary conditions were applied on the upper surface and

¹StressCheck is Trademark of ESRD, St. Louis, MO, USA. StressCheck exhibits some multi-threading as evidenced by multiple cores being active during a solution.

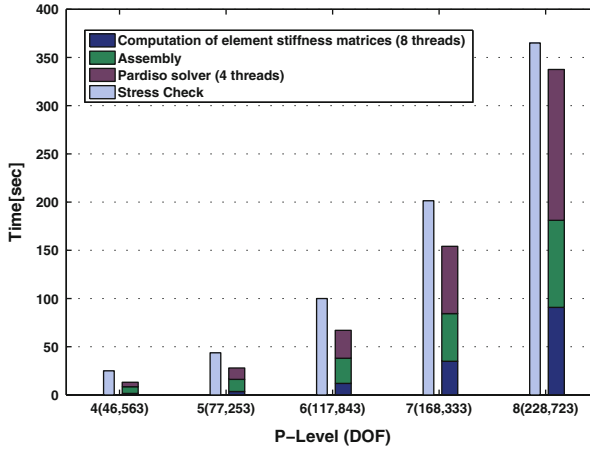


Fig. 14 CPU time of the Java code compared to the commercial code StressCheck for the 1000 FE mesh as the p -level is increased. Computations performed on a DELL T3500 with a single Intel Xeon E5-1620 3.6 Ghz processor (4 cores, 8 threads) and 20 GB RAM

a clamped boundary condition was applied to the lower surface. The model was solved for $p = 1$ to 8 and the computation times are given in Fig. 14.

Although the Java p -FE code is in an early stage of development, our experience indicates that it is as efficient as optimized commercial codes and provides a platform that is convenient for performing parallel computations on a variety of desktop, cloud, and mobile devices.

4 Summary and Conclusions

Personalized, verified, and validated simulations of human bones have been presented, that are suitable for application in daily orthopedic practice. We are developing a p -FE code using Java and parallel algorithms capable of delivering simultaneously computed patient specific simulations from the cloud to the surgeon's mobile device. It is hoped that when such computational tools are available to the orthopedic community, in a form that can be easily interpreted by medical doctors, that FE methods will find their way into daily medical practice, resulting in more effective treatment at a personalized level.

References

1. Bayraktar, H.H., Morgan, E.F., Niebur, G.L., Morris, G.E., Wong, E.K., Keaveny, M.: Comparison of the elastic and yield properties of human femoral trabecular and cortical bone tissue. *J. Biomech.* **37**, 27–35 (2004)

2. Goodsitt, M.M.: Conversion relations for quantitative ct bone mineral densities measured with solid and liquid calibration standards. *Bone Miner.* **19**, 145–158 (1992)
3. Katz, A.: The mechanical response of femurs fixed by metal devices. MSc thesis, Ben-Gurion University of the Negev (2011)
4. Keller, T.S.: Predicting the compressive mechanical behavior of bone. *J. Biomech.* **27**, 1159–1168 (1994)
5. Keyak, J.H., Meagher, J.M., Skinner, H.B., Mote Jr., C.D.: Automated three-dimensional finite element modelling of bone: a new method. *ASME J. Biomech. Eng.* **12**, 389–397 (1990)
6. Keyak, J.H., Fourkas, M.G., Meagher, J.M., Skinner, H.B.: Validation of automated method of three-dimensional finite element modelling of bone. *ASME J. Biomech. Eng.* **15**, 505–509 (1993)
7. Kuzmin, A., Luisier, M., Schenk, O.: Fast methods for computing selected elements of the greens function in massively parallel nanoelectronic device simulations. In: Wolf, F., Mohr, B., Mey, D. (eds.) *Euro-Par 2013 Parallel Processing*, pp. 533–544. Springer, Berlin/Heidelberg (2013)
8. Nikishkov, G.: *Programming Finite Elements in Java*. Springer, London (2010)
9. Schileo, E., DallAra, E., Taddei, F., Malandrino, A., Schotkamp, T., Baleani, M., Viceconti, M.: An accurate estimation of bone density improves the accuracy of subject-specific finite element models. *J. Biomech.* **41**, 2483–2491 (2008)
10. Szabó, B.A., Babuška, I.: *Finite Element Analysis*. Wiley, New York (1991)
11. Trabelsi, N., Yosibash, Z.: Patient-specific FE analyses of the proximal femur with orthotropic material properties validated by experiments. *ASME J. Biomech. Eng.* **155**, 061001-1–061001-11 (2011)
12. Trabelsi, N., Yosibash, Z., Milgrom, C.: Validation of subject-specific automated p-FE analysis of the proximal femur. *J. Biomech.* **42**, 234–241 (2009)
13. Trabelsi, N., Yosibash, Z., Wutte, C., Augat, R., Eberle, S.: Patient-specific finite element analysis of the human femur - a double-blinded biomechanical validation. *J. Biomech.* **44**, 1666–1672 (2011)
14. Yosibash, Z.: p-FEMs in biomechanics: bones and arteries. *Comput. Methods Appl. Mech. Eng.* **249–252**, 169–184 (2012)
15. Yosibash, Z., Trabelsi, N.: Reliable patient-specific simulations of the femur. In: Gefen, A. (ed.) *Patient-Specific Modeling in Tomorrow's Medicine*, pp. 3–26. Springer, Berlin/Heidelberg (2012)
16. Yosibash, Z., Padan, R., Joscowicz, L., Milgrom, C.: A CT-based high-order finite element analysis of the human proximal femur compared to in-vitro experiments. *ASME J. Biomech. Eng.* **129**(3), 297–309 (2007)
17. Yosibash, Z., Trabelsi, N., Milgrom, C.: Reliable simulations of the human proximal femur by high-order finite element analysis validated by experimental observations. *J. Biomech.* **40**, 3688–3699 (2007)
18. Yosibash, Z., Tal, D., Trabelsi, N.: Predicting the yield of the proximal femur using high order finite element analysis with inhomogeneous orthotropic material properties. *Philos. Trans. R. Soc. A* **368**, 2707–2723 (2010)
19. Yosibash, Z., Plitman Mayo, R., Dahan, G., Trabelsi, N., Amir, G., Milgrom, C.: Predicting the stiffness and strength of human femurs with realistic metastatic tumors. *Bone* **69**, 180–190 (2014)
20. Yosibash, Z., Wille, H., Rank, E.: Stochastic description of the peak hip contact force during walking free and going upstairs. *J. Biomech.* **48**, 1015–1022 (2015)