A Fast Implicit Solver with Low Memory Footprint and High Scalability for Comprehensive Earthquake Simulation System

Kohei Fujita[1], Tsuyoshi Ichimura[2], Kentaro Koyama[3], Masashi Horikoshi[4], Hikaru Inoue[3], Larry Meadows[5],
Seizo Tanaka[6], Muneo Hori[2,1], Maddegedara Lalith[2] and Takane Hori[7]
[1]Advanced Institute for Computational Science, RIKEN
[2]Earthquake Research Institute & Department of Civil Engineering, The University of Tokyo
[3]Frontier Computing Center, Fujitsu Limited
[4]Software and Solutions Group, Intel K.K.
[5]Data Center Group, Intel Corporation
[6]Faculty of Engineering, Information and System, University of Tsukuba
[7]R&D Center for Earthquake and Tsunami, Japan Agency for Marine-Earth Science and Technology

To achieve a reduction in the damage caused by large earthquakes by improvement in reliability of disaster estimations, we have developed an HPC-based comprehensive earthquake simulation system of wave propagation, structural and social responses. Although challenging to attain performance on current computer architectures, we greatly accelerated unstructured implicit finite-element simulations needed to solve nonlinear wave-propagation problems in complex-shaped domains, by both computational and computer science considerations (e.g., Ichimura et al., Gordon Bell Prize Finalists in SC14 [1] and SC15 [2]). By applying this fast and scalable general-purpose finite-element solver to wave propagation, soil amplification, and crust-deformation analyses, we can simulate all phases of earthquake disasters on a larger scale and with higher resolution than existing HPC methods. Furthermore, these physics-based simulation results can be input into social simulations of an earthquake disaster. This system is expected to be a next-generation earthquake disaster estimation method, and joint projects are currently underway with national and local governments as well as with nuclear power and lifeline industries. This poster presents the algorithmic design, implementation, performance, and portability of a fast low-memory footprint implicit solver developed to forecast earthquakes by data assimilation – the next ambitious target of a comprehensive earthquake simulation system.

With physics-based crust-deformation simulation in earthquake forecasting, a resolution of 10^{0-1} m is required to resolve the physics of faulting on a heterogeneous fault plane [3], [4], while a domain of 10^{6-7} m $\times 10^{6-7}$ m $\times 10^{5-6}$ m with many plates having complex three-dimensional geometries needs to be modeled to connect the fault plane to the observation points. This leads to at least 10^{0-1} Tera-Degrees-of-Freedom (T-DOF) unstructured finite-element simulations, even when using mathematical methods such as multiscale analyses based on singular perturbations. This problem is 10^{2-3} fold larger than the 0.01 T-DOF problem solved in the current state-of-the-art simulation using one-tenth of the K computer [5], [6], and thus, a fast implicit solver with a low-memory footprint and high scalability is needed for such computations on present-day 10^{1-2} PFLOP machines with 10^0 petabyte memory. Thus we designed a solver using multi-grid preconditioning in an inexact conjugate gradient solver [7] together with mixed-precision arithmetic. A communication avoiding inexact LU decomposition [8], which has very good convergence characteristics for crust-deformation problems, is used for preconditioning the coarse grid. By using a sophisticated preconditioner on a problem reduced by the multi-grid, we can reduce the iteration counts while attaining a low-memory footprint. Furthermore, memory access, communication, and computation cost per iteration are reduced by using single-precision arithmetic in the preconditioner. For matrix-vector multiplication in these solvers, we use matrix-free methods (i.e., element-by-element method, EBE [9]) to further reduce the memory footprint. As EBE becomes an on-cache computation kernel, we can also circumvent load-imbalance from the differences in memory access patterns in the unstructured computations. Indeed, 96.6% size-up efficiency was attained up to full K computer with 663,552 CPU cores for the whole solver with a simple preconditioner and 88.8% size-up efficiency was attained up to full K computer even when using the targeted ILU preconditioner on a highly heterogeneous problem. Although it is not a straightforward process to achieve high levels of performance on EBE kernels on CPUs with SIMD units and multi-cores due to data recurrence, we could accelerate these hot-spot EBE kernels using novel SIMD buffering and coloring methods. These developments led to high performance levels for an implicit unstructured finite-element method; 27.6% of double-precision peak was attained for the EBE kernels and 11.5% of peak (1.21 PFLOPS) was attained for the whole solver when using the full K computer system. The method is also highly portable; the EBE kernel attained 27.0% of double-precision peak FLOPS on the newest Intel Broadwell CPU [10], and although only at a preliminary stage, we could also confirm high levels of portability using GPU architecture. As expected, the solver is extremely effective for practical problems, and enabled solving a 2.05 T-DOF problem of Eastern Japan, which corresponds to 205 times more DOF compared with the state-of-the-art, in just 3199 s by using the whole K computer. This is 30 times faster compared with a memory-efficient solver developed in SC14 using 3 × 3 block Jacobi preconditioning with EBE scalable up to the whole K computer.

This scalable and portable unstructured computation algorithm capable of utilizing SIMD width and multiple core counts is expected to benefit from large advances in hardware capability in the next generation of many-core architectures (e.g., Post-K supercomputer [11] and Knights Landing [12]). This will enable multiple practical crust-deformation simulations, and thus, will strongly support advances in physics-based earthquake forecasting based on data assimilation. This unstructured finite-element solver is also expected to be useful in a wide range of applications where appropriate geometrical modeling is essential for ensuring appropriate degrees of numerical accuracy.

REFERENCES

[1] T. Ichimura, K. Fujita, S. Tanaka, M. Hori, M. Lalith, Y. Shizawa, and H. Kobayashi. "Physics-based urban earthquake simulation enhanced by 10.7 BlnDOF x 30 K time-step unstructured FE non-linear seismic wave simulation," Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, (SC'14), pp 15–26, 2014.

[2] T. Ichimura, K. Fujita, P.E.B. Quinay, L. Maddegedara, M. Hori, S. Tanaka, Y. Shizawa, H. Kobayashi, and K. Minami, "Implicit nonlinear wave simulation with 1.08T DOF and 0.270T unstructured finite elements to enhance comprehensive earthquake simulation," Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, (SC'15), 2015.

[3] H. Noda, M. Nakatani, and T. Hori, "Large nucleation before large earthquakes is sometimes skipped due to cascade-up–Implications from a rate and state simulation of faults with hierarchical asperities," J. Geophys. Res., 118, 6, 2924-2952, 2013.

[4] M. Hyodo, T. Hori, and Y. Kaneda, "A possible scenario for earlier occurrence of the next Nankai earthquake due to triggering by an earthquake at Hyuga-nada, off southwest Japan," Earth, Planets and Space, 68, 6, 2016.

[5] T. Ichimura, R. Agata, T. Hori, K. Hirahara, C. Hashimoto, M. Hori, and Y. Fukahata, "An elastic/viscoelastic finite element analysis method for crustal deformation using a 3D island-scale high-fidelity model," Geophysical Journal International, doi:10.1093/gji/ggw123 (in advance access, first published online April 4, 2016).

[6] H. Miyazaki, Y. Kusano, N. Shinjou, F. Shoji, M. Yokokawa, and T. Watanabe. "Overview of the K computer system," FUJITSU Sci. Tech. J., 48, 3, 302-309, 2012.

[7] G. H. Golub and Q. Ye Inexact conjugate gradient method with inner-outer iteration. SIAM., Journal on Scientific Computing 21(4), 1305-1320, 1997.

[8] Y. Saad, Iterative methods for sparse linear systems (2nd ed.), *SIAM*, 2003.

[9] J. M. Winget and T. J. R. Hughes, "Solution algorithms for nonlinear transient heat conduction analysis employing element-by-element iterative strategies," Computer Methods in Applied Mechanics and Engineering, 52, 711-815, 1985.

[10] Intel Xeon E5 v4 series product web page [Online]. http://www.intel.com/content/www/us/en/processors/xeon/xeon-e5-brief.html

[11] Y. Ishikawa, "System software in post K supercomputer," SC15: International Conference for High Performance Computing, Networking, Storage and Analysis (Invited Talk), 2015.

[12] HotChips 2015, Knights Landing: 2nd Generation Intel Xeon Phi Processor, http://www.hotchips.org/wp-content/uploads/hc_archives/hc27/HC27.25-Tuesday-Epub/HC27.25.70-Processors-Epub/HC27.25.710-Knights-Landing-Sodani-Intel.pdf