




Contents lists available at ScienceDirect

Extreme Mechanics Letters

journal homepage: www.elsevier.com/locate/eml

Computational homogenization at extreme scales



Matthew Mosby, Karel Matouš*

Department of Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, IN, 46556, USA

ARTICLE INFO

Article history:

Received 7 October 2015

Received in revised form 21 December 2015

Accepted 22 December 2015

Available online 29 December 2015

Keywords:

Computational homogenization

High-performance computing

Extreme scale computing

Heterogeneous interfaces

Multi-scale interfacial damage modeling

ABSTRACT

Multi-scale simulations at extreme scales in terms of both physical length scales and computational resources are presented. In this letter, we introduce a hierarchically parallel computational homogenization solver that employs hundreds of thousands of computing cores and resolves $\mathcal{O}(10^5)$ in material length scales (from $\mathcal{O}(\text{cm})$ to $\mathcal{O}(100 \text{ nm})$). Simulations of this kind are important in understanding the multi-scale essence of many natural and synthetically made materials. Thus, we present a simulation consisting of 53.8 Billion finite elements with 28.1 Billion nonlinear equations that is solved on 393,216 computing cores (786,432 threads). The excellent parallel performance of the computational homogenization solver is demonstrated by a strong scaling test from 4,096 to 262,144 cores. A fully coupled multi-scale damage simulation shows a complex crack profile at the micro-scale and the macroscopic crack tunneling phenomenon. Such large and predictive simulations are an important step towards *Virtual Materials Testing* and can aid in development of new material formulations with extreme properties. Furthermore, the high computational efficiency of our computational homogenization solver holds great promise for utilizing the next generation of exascale parallel computing platforms that are expected to accelerate computations through orders of magnitude increase in parallelism rather than speed of each processor.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

The societal and economic pressure for improved performance of engineered systems has placed great emphasis on development of materials with extreme properties and their application in extreme environments. For example, development of advanced high strength multi-phase steels [1] for improved automotive crash worthiness, structured materials [2] and meta-materials [3], fiber/particle reinforced polymeric composites [4,5] with wide application from aerospace to consumer sports equipment, and multi-functional systems such as self-healing [6] or electrically conductive adhesives [7] has become reality. In many cases, these advanced materials are

multi-scale in nature, and accurately predicting their response is essential for improved design and safety assessment.

Of particular interest to this letter are predictive multi-scale simulations of these complex heterogeneous materials in typical mechanical systems. In such engineering analysis and optimal design, phenomenological constitutive models of heterogeneous materials may prove insufficient. Thus, detailed simulations that include microstructural effects and relevant micro-scale physics are required. Direct numerical modeling (DNM), which captures all of the relevant physics and length scales in a single simulation, is an accurate method for predicting the *in situ* multi-scale behavior of heterogeneous materials. However, even for small structures, these simulations can become extremely large as the required numerical resolution leads to a large number of degrees of freedom (DOFs). Such large simulations remain intractable, even for today's

* Corresponding author.

E-mail address: kmatous@nd.edu (K. Matouš).




Fig. 1. Schematic of the multi-scale kinematics for heterogeneous interfaces with a comparison of length-scales.

supercomputers [8]. Moreover, the parallel algorithmic complexity of such computations mitigates the potential improvements gained on high-performance computing (HPC) systems [9–11]. On the other hand, computational homogenization (CH) [12–14] provides an alternative highly accurate modeling strategy with reduced computational requirements. However, until recently [12], CH has frequently been regarded as impractical or limited to small theoretical examples.

In this letter, we present extreme scale simulations (in terms of both physical scales and computing resources) using a hierarchically parallel CH solver [12] that enables the efficient computation of large realistic engineering problems. In particular, we focus on failure of heterogeneous interfaces such as adhesive layers. For the first time, fully-coupled multi-scale simulations are used to predict important fracture properties, such as toughness and crack speed, from the material behavior of the individual micro-scale constituents in the 3D finite strain regime.

In addition, we show the solver's ability to efficiently compute the fully coupled nonlinear multi-scale response of structures with resolution from $\mathcal{O}(\text{cm})$ to $\mathcal{O}(100 \text{ nm})$, containing 53.8 Billion finite elements and 28.1 Billion nonlinear equations. Furthermore, we demonstrate ideal computational strong scaling performance of the hierarchically parallel solver using up to 262,144 computing cores. The ability to compute such large problems is an important step towards predictive simulations and the *Virtual Materials Testing* paradigm. In addition, future exascale HPC resources are expected to accelerate computations through orders of magnitude increase in parallelism rather than increasing the speed of each processor [15–18]. Therefore, the high scalability of CH makes it a promising approach to efficiently use future exascale HPC resources for scientific investigation and discovery.

2. Computational homogenization for interfaces

Before we proceed to extreme scale computations, we review the CH theory and its implementation [19,20,12] for completeness of the presentation. The CH of heterogeneous interfaces is shown schematically in Fig. 1, where two bodies (adherends) denoted as Ω^\pm are separated by a heterogeneous layer with thickness l_c . The layer is collapsed to an interface, Γ , and a representative unit cell (RUC, Θ) is locally attached to each material point on the interface. The RUC contains all of the micro-scale complexity in terms of both structure and constitutive behavior. Under applied load, the deformation of the macro-scale

adherends, Ω^\pm , is described by the deformation gradient, $\mathbf{F} = \mathbf{I} + \bar{\nabla}_{\bar{\mathbf{x}}}\bar{\mathbf{u}}$, where $\bar{\mathbf{u}}$ are the macro-scale displacements. The deformation of a macroscopic point on the interface, Γ , is described by the deformation gradient, $\mathbf{F}_M = \mathbf{I} + \frac{1}{l_c} [[\bar{\mathbf{u}}]] \otimes \bar{\mathbf{N}}$, where $\bar{\mathbf{N}}$ is the normal to the interface (see Fig. 1) and $[[\bar{\mathbf{u}}]] = \bar{\mathbf{u}}^+ - \bar{\mathbf{u}}^-$ is the opening displacement of the interface. The deformation within the micro-structure is a function of both macro- and micro-variables, with $\mathbf{F} = \mathbf{F}_M + \bar{\nabla}_{\bar{\mathbf{y}}}\bar{\mathbf{w}}$, where $\bar{\mathbf{w}}$ are the micro-scale displacement fluctuations.

The weak form of macro-scale equilibrium neglecting inertia and body forces is given by

$$\int_{\Omega^\pm} \mathbf{S}_M : [\mathbf{F}^T \bar{\nabla}_{\bar{\mathbf{x}}}\delta\bar{\mathbf{u}}]^{\text{sym}} d\Omega^\pm + \int_{\Gamma} \bar{\mathbf{t}}_M \cdot [[\delta\bar{\mathbf{u}}]] dA - \int_{\partial\Omega^t} \bar{\mathbf{t}}^p \cdot \delta\bar{\mathbf{u}} dA = 0. \quad (1)$$

In Eq. (1), \mathbf{S}_M is the second Piola–Kirchhoff stress given by a known constitutive model for the adherends, while $\bar{\mathbf{t}}_M$ is the macro-scale traction vector across the interface that is computationally derived from the RUC as described in the sequel.

The Hill–Mandel condition for interfaces is given by [19,12,20]

$$\inf_{[[\bar{\mathbf{u}}]]} \psi([[\bar{\mathbf{u}}]]) = \inf_{[[\bar{\mathbf{u}}]]} \inf_{\bar{\mathbf{w}} \mid |\Theta|} \frac{l_c}{|\Theta|} \int_{\Theta} W_m(\mathbf{F}_M + \bar{\nabla}_{\bar{\mathbf{y}}}\bar{\mathbf{w}}) d\Theta, \quad (2)$$

which relates the unknown macro-scale traction–separation potential, ψ , to the average known micro-scale strain energy density, W_m . Taking variations of Eq. (2) with respect to the macro-scale and micro-scale variables leads to

$$\frac{l_c}{|\Theta|} \int_{\Theta} \mathbf{S}_m : [\mathbf{F}^T \bar{\nabla}_{\bar{\mathbf{y}}}\delta\bar{\mathbf{w}}]^{\text{sym}} d\Theta = 0, \quad (3)$$

$$\bar{\mathbf{t}}_M = \frac{1}{|\Theta|} \left[\int_{\Theta} \mathbf{F} \mathbf{S}_m d\Theta \right] \cdot \bar{\mathbf{N}}. \quad (4)$$

Eq. (3) is the weak form of micro-scale equilibrium, and Eq. (4) is the closure equation for the macro-scale traction vector across the interface employed in Eq. (1). In Eqs. (3)–(4), $\mathbf{S}_m = 2 \partial W_m / \partial \mathbf{C}$ where $\mathbf{C} = \mathbf{F}^T \mathbf{F}$. Note that in this work, we make use of semi-periodic micro-scale boundary conditions ($\bar{\mathbf{w}} = \mathbf{0} \forall \bar{\mathbf{y}} \in \Gamma^\pm$ and $\bar{\mathbf{w}}^+ = \bar{\mathbf{w}}^- \forall \bar{\mathbf{y}} \in \partial\Theta^\pm$) and our RUCs are periodic in the $\bar{Y}_{1,2}$ -directions (see Fig. 1). Other admissible boundary conditions are discussed in [20].

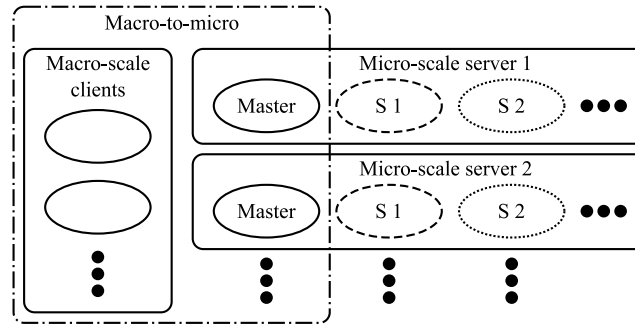


Fig. 2. Schematic of the hierarchically parallel communication framework. The communication between scales is performed only by the “master” processors of each server and the macro-scale client processors. Note that communication between servers is restricted to the equivalent task id (e.g., S1 on server 1 can communicate with S1 on server 2, but not S2 on server 2).

We solve the nonlinear multi-scale equilibrium (Eqs. (1) and (3)) via the implicit finite element method, which requires the iterative solution to the linearized discrete system of equations

$$\begin{bmatrix} K_{MM}^{\Omega\pm} + K_{MM}^{\Gamma} & K_{Mm} \\ K_{mM} & K_{mm} \end{bmatrix} \begin{Bmatrix} \Delta \bar{u} \\ \Delta \bar{w} \end{Bmatrix} = - \begin{Bmatrix} \bar{R}_M \\ \bar{R}_m \end{Bmatrix}$$

by Newton’s method. For realistic engineering problems, the multi-scale system of equations can be extremely large. Therefore, we use a staggered approach described in Algorithm 1. For more details on the staggered solution procedure, please see [12].

Algorithm 1 Multi-scale staggered solution method.

1. Hold \bar{u} constant, compute \bar{w} from Equation (3)
 2. Compute \bar{t}_M from Equation (4)
 3. Assemble $\mathcal{K} = K_{MM}^{\Gamma} - K_{Mm}K_{mm}^{-1}K_{mM}$
 4. Solve $(K_{MM}^{\Omega\pm} + \mathcal{K})\Delta \bar{u} = -\bar{R}_M$
 5. Iterate to convergence ($\bar{R}_M = \bar{R}_m \leq \text{tolerance}$)
-

3. High-performance computational implementation

The CH formulation allows the response of each RUC to be computed independently, which enables high computational efficiency. In many cases the computational domains for both the macro-scale and micro-scale are large, and thus a parallel solution strategy is required at each scale. Therefore, we employ our in-house hierarchically parallel CH solver that uses an efficient parallel finite element code, *PGFem3D* [21–23,20], at each scale in a client–server framework [12]. Verification of the CH solver for hyper-elastic materials was presented in our prior work [12].

Fig. 2 displays a schematic of the client–server communication structure. The macro-scale equilibrium is computed in parallel on the “client” processors, and the contributions from the individual RUCs are computed in parallel on the “servers”. Furthermore, all communications are performed using dynamic point-to-point non-blocking messages that are overlaid with computations at both scales.

Under complex macroscopic loading conditions, a work imbalance among the micro-scale servers usually develops due to some RUCs requiring more computational

effort than others. Furthermore, evolving the micro-scale constitutive equations, e.g., damage or plasticity, can also promote this imbalance between servers. Minimizing this workload imbalance is paramount to maintain high resource utilization and reduce computational time. We use a “largest first” [24] load balancing scheme to reassign RUCs to servers based on how long the previous computations took. The new assignment is only used if it will reduce the overall computational time including the RUC data migration. We overlay the data migration with computation of RUCs that are not reassigned.

4. Extreme scale simulations

In the previous sections, we gave an overview of the CH formulation and its efficient HPC implementation. In the following subsections, we present three distinct numerical examples that demonstrate both the excellent computational efficiency of the CH solver, as well as its ability to compute extremely large simulations with fine numerical resolution and highly nonlinear physical behavior. Such detailed multi-scale simulations can lead to understanding of mechanics by directly linking the underlying morphology and phenomena to the overall performance of engineering structures. Furthermore, with co-designed experiments, such simulations can accelerate the development of new material formulations. Moreover, these examples show the excellent potential of deploying CH on future exascale computing platforms.

4.1. Scaling performance

First, we perform a strong scaling study to demonstrate the computational efficiency of the CH solver using a patch test (see inset of Fig. 3). The macro-scale domain consists of two cubic steel blocks with side lengths of 10 mm separated by an interface. In both this and the next section, each RUC attached to material points on the interface is $210 \times 210 \times 210 \mu\text{m}^3$ and is made of an epoxy resin containing 98 randomly distributed voids with $d = 30 \mu\text{m}$. All materials are modeled as geometrically nonlinear hyper-elastic solids. See [12,20] for more details on the constitutive equations used in this letter.

Usually, computational speedup is defined as $s = t_n/t_N$, where t_n is the time required to compute the simulation

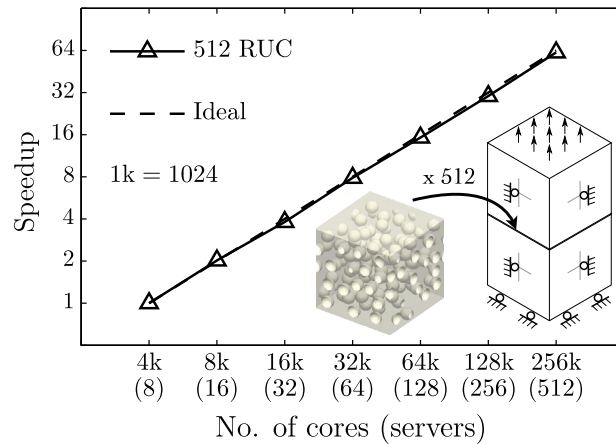


Fig. 3. Computational speedup of the hierarchically parallel multi-scale solver. The macro-scale is discretized with 17.5k finite elements and 9.8k nonlinear DOFs, and is computed using 32 cores. Each of the 512 RUCs is discretized with 1.46M finite elements and 773k nonlinear DOFs, and is computed using 512 cores. The complete multi-scale simulation contains 747M finite elements and 396M DOFs.

using n computational resources (e.g., micro-scale servers) and t_N is the time to compute the same simulation using N resources. In the ideal case, $s = N/n$ for all N , but speedup is typically limited by algorithmic complexity and communication overhead. However, using the hierarchically parallel multi-scale solver, both macro- and micro-domains can be computed using the optimal number of computing cores and further speedup can be obtained by increasing the number of servers used for computing the micro-scale responses. Due to the high parallelism inherent to CH and the efficient implementation described in Section 3, the CH solver maintains ideal speedup from 4096 to 262,144 computing cores on the Vulcan machine at the Lawrence Livermore National Laboratory (LLNL). In this example, the speedup is evaluated from the time to compute one nonlinear fully-coupled multi-scale load increment.

4.2. Hyper-elastic billet compression

The aspects of computational efficiency become important for larger multi-scale simulations spanning several orders of magnitude in length scales. Fig. 4 shows the hyper-elastic multi-scale compression response of a heterogeneous layer computed with numerical resolution spanning $\mathcal{O}(10^5)$ length scales (from $\mathcal{O}(10$ mm) at the macro-scale to $\mathcal{O}(100$ nm) at the micro-scale). The macro-scale domain consists of two steel disks with $t = 10$ mm and $d = 20$ mm separated by an epoxy layer. The macro-scale discretization contains 325k finite elements and 182k nonlinear DOFs, and is computed using 512 cores. Each RUC corresponding to the 5296 macro-scale cohesive elements on the interface is finely discretized, each containing 10.2M finite elements ($h_{\min} = 191$ nm, $h_{\text{mean}} = 1.8$ μm , and $h_{\max} = 2.8$ μm) and 5.3M DOFs. The micro-scale response is computed using 767 servers consisting of 512 cores each. The total implicit multi-scale simulation contains 53.8B finite elements, 28.1B nonlinear equations, and was computed using 393,216 computing cores (786,432 threads) on the Vulcan machine at LLNL. Four nonlinear fully-coupled multi-scale load increments were computed in under 12 h. Fig. 4 shows the multi-scale response at the end of the load

history. As shown, the non-uniform response at both scales is well resolved.

4.3. Progressive failure of a dual cantilever beam

In the previous examples, only the nonlinear hyper-elastic response was considered. Here, we demonstrate the ability of the CH solver to predict the multi-scale failure response of a dual cantilever beam (DCB) using an isotropic viscous damage model at the micro-scale (see [25,19,20]). In addition, we measure the effect of using the load balancing algorithm described in Section 3.

A schematic of the multi-scale DCB and mode-I loading conditions are shown in Fig. 5 (inset). The DCB is loaded under displacement control at 0.375 mm/min to ensure quasi-static isothermal conditions. The DCB adherends are each 42 mm long, 10 mm wide, 5 mm thick, and are discretized with a total of 10k finite elements. The macro-scale interface is 40 mm long (2 mm pre-crack) and is discretized by 322 cohesive elements, each with a corresponding RUC. In this example, the RUC is $250 \times 250 \times 125$ μm^3 and contains 40 randomly distributed voids with $d = 40$ μm . The RUCs are each discretized with 249k finite elements, and the total multi-scale simulation contains 80M finite elements and 42.5M nonlinear DOFs. The multi-scale response is computed using up to 128k cores on the Mira supercomputer at Argonne National Laboratory.

Fig. 5 shows the macro-scale mode-I response of the DCB compared to the analytical solution provided by the Linear Fracture Mechanics (LFM) theory (broken lines). Note that we are using a rate dependent failure model at the micro-scale, and thus each material point on the interface evolves differently with varying fracture properties as shown in Fig. 6. The CH solver predicts that the fracture toughness of material points on the interface is $G_c \in [141, 175]$ J/m^2 due to varying local strain rates. The mean toughness of the interface is $G_c^{\text{mean}} = 162$ J/m^2 . The critical interface traction is predicted as $t_c = \max \|\vec{t}_M\| \in [61, 68.2]$ MPa, with $t_c^{\text{mean}} = 65.6$ MPa.

Returning to Fig. 5, the multi-scale response is bounded by the LFM theory in the failure regime, but the LFM




Fig. 4. Large multi-scale compression simulation of a heterogeneous hyperelastic layer. Clockwise from top left: Schematic of the macro-scale domain and loading conditions, response of the macro-scale adherends, macroscopic response of the interface, micro-scale effective strain of the heterogeneous interfaces at the marked points.




Fig. 5. Macro-scale model-I DCB loading conditions (inset) and computed macroscopic force–displacement curve. Broken lines denote the LFM response for $G_c^{\min} = 141 \text{ J/m}^2$ (dashed), $G_c^{\max} = 175 \text{ J/m}^2$ (dash-dot), and $G_c^{\text{mean}} = 162 \text{ J/m}^2$ (dotted).




Fig. 6. Effective traction–separation response of fully failed material points on the macro-scale interface.




Fig. 7. Macro-scale mode-I response of the DCB in the vicinity of the propagating crack tip. From left: equivalent Mises stress in the deformed adherends (deformations magnified 100×) and magnitude of the cohesive traction on the interface. (The pre-crack is not shown on the macro-scale interface.) Note the crack tunneling effect.




Fig. 8. Extent of micro-scale damage, $\omega \in [0, 1]$, at the marked points in Fig. 7(right). Material points with $\omega = 0$ and $\omega = 1$ are undamaged and fully failed, respectively.

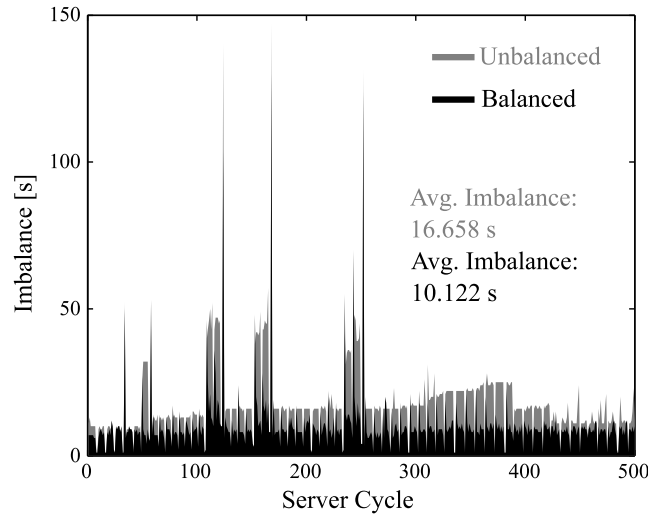


Fig. 9. Evolution of the maximum server imbalance up to the limit point with and without the load balancing scheme. For the rebalancing study, the macro-scale is computed using 16 cores, while 35 micro-servers using 32 cores each are utilized to compute the micro-scale response (1136 computing cores total). Each server is allowed to compute up to 15 micro-scale solutions to allow for non-uniform rebalancing.

theory over-predicts the maximum load by more than 20% ($P_{\max} = 759$ N from the CH). We emphasize the predictive nature of this simulation, and the CH modeling strategy in general, since the fracture behavior of the heterogeneous interface is entirely governed by the bulk behavior of the individual micro-scale constituents.

The macro-scale response of the DCB is shown in Fig. 7. The non-planar progressive failure of the interface governed by the micro-scale failure is shown in Fig. 7 (right). Note that the initial crack tip (point 1) has completely failed and the process zone (point 2) has propagated along the macro-scale interface. Additionally, one can see the crack tunneling in Fig. 7 (right), captured by the three-dimensional analysis. The crack tunneling is

expected due to the ductile failure response of the interface (see Fig. 6) and the transition from a plane-strain condition near the center of the DCB to a stress-free condition at the boundaries [26,27]. We also predict the crack propagation speed, $v_c = 1.9$ mm/s ($\sim 304\times$ the loading rate). Fig. 8 shows the extent of damage in the micro-scale at the marked points along the macro-scale interface (point 1 and point 2). Results in this section show the impact of extreme scale simulations in predicting and better understanding mechanics across vast spatial scales during the complex rate-dependent damage process.

As mentioned in Section 3, the introduction of more complex macro-scale loading conditions and micro-scale constitutive equations containing material nonlinearity

result in computational load imbalance among the micro-scale servers. The effect of the load balancing algorithm was evaluated up to the limit point in the macro-scale load history (see Fig. 9). The average imbalance is reduced by nearly 40% and the simulation up to the limit point is computed 12% faster overall when using the load balancing algorithm.

5. Conclusions

We have presented extreme scale computations, in terms of both physical scales and computing resources, containing $\sim 54\text{B}$ finite elements, over 28B nonlinear DOFs and executed on ~ 400 thousand computing cores. Such large and detailed simulations are necessary for better understanding of complex (i.e., rate-dependent) multi-scale material behavior under nontrivial loading conditions. Moreover, with co-designed experiments and properly validated constitutive models, such large predictive simulations can be the basis of *Virtual Materials Testing* standards, and aid in development of new material formulations with extreme properties. These large simulations were enabled by a hierarchically parallel CH solver with excellent computational scaling behavior. The high scalability of the CH method makes it a promising approach for efficiently utilizing the future massively parallel exascale platforms.

Acknowledgments

This work was supported by the Department of Energy, National Nuclear Security Administration, under the award number DE-NA0002377. We would also like to acknowledge computational resources from the 2015 ASCR Leadership Computing Challenge (ALCC) under project number CSC188.

References

- [1] R. Kuziak, R. Kawalla, S. Waengler, Advanced high strength steels for automotive industry, *Arch. Civ. Mech. Eng.* 8 (2) (2008) 103–117.
- [2] X.W. Gu, J.R. Greer, Ultra-strong architected Cu meso-lattices, *Extreme Mech. Lett.* 2 (2015) 7–14. <http://dx.doi.org/10.1016/j.eml.2015.01.006>.
- [3] S. Krödel, N. Thomè, C. Daraio, Wide band-gap seismic metastructures, *Extreme Mech. Lett.* 4 (2015) 111–117. <http://dx.doi.org/10.1016/j.eml.2015.05.004>.
- [4] P.K. Mallick, *Fiber-Reinforced Composites: Materials, Manufacturing, and Design*, CRC press, 2007.
- [5] S.-Y. Fu, X.-Q. Feng, B. Lauke, Y.-W. Mai, Effects of particle size, particle/matrix interface adhesion and particle loading on mechanical properties of particulate–polymer composites, *Composites B* 39 (6) (2008) 933–961.
- [6] S.R. White, N. Sottos, P. Geubelle, J. Moore, M. Kessler, S. Sriram, E. Brown, S. Viswanathan, Autonomic healing of polymer composites, *Nature* 409 (6822) (2001) 794–797.
- [7] H.H. Lee, K.S. Chou, Z.W. Shih, Effect of nano-sized silver particles on the resistivity of polymeric conductive adhesives, *Int. J. Adhes. Adhes.* 25 (5) (2005) 437–441.
- [8] T.M. Evans, G.G. Davidson, R.N. Slaybaugh, K. Clarno, Three-dimensional full core power calculations for pressurized water reactors, in: *SciDAC, J. Phys. Conf. Ser.* 68 (2010) 367–379.
- [9] P. Ghysels, T.J. Ashby, K. Meerbergen, W. Vanroose, Hiding global communication latency in the GMRES algorithm on massively parallel machines, *SIAM J. Sci. Comput.* 35 (1) (2013) C48–C71.
- [10] X.S. Li, J.W. Demmel, SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Trans. Math. Software (TOMS)* 29 (2) (2003) 110–140.
- [11] A.H. Baker, T. Gamblin, M. Schulz, U.M. Yang, Challenges of scaling algebraic multigrid across modern multicore architectures, in: *Parallel & Distributed Processing Symposium (IPDPS)*, 2011 IEEE International, IEEE, 2011, pp. 275–286.
- [12] M. Mosby, K. Matouš, Hierarchically parallel coupled finite strain multiscale solver for modeling heterogeneous layers, *Internat. J. Numer. Methods Engrg.* 102 (3–4) (2015) 748–765.
- [13] M. Kulkarni, K. Matouš, P. Geubelle, Coupled multi-scale cohesive modeling of failure in heterogeneous adhesives, *Internat. J. Numer. Methods Engrg.* 84 (8) (2010) 916–946.
- [14] M. Geers, V. Kouznetsova, W. Brekelmans, Multi-scale computational homogenization: Trends and challenges, *J. Comput. Appl. Math.* 234 (7) (2010) 2175–2182.
- [15] H. Simon, T. Zacharia, R. Stevens, et al. Modeling and simulation at the exascale for energy and the environment, Department of Energy Technical Report, 2007.
- [16] K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, K. Hill, J. Hiller, et al. Exascale computing study: Technology challenges in achieving exascale systems, Defense Advanced Research Projects Agency Information Processing Techniques Office (DARPA IPTO), Tech. Rep 15, 2008.
- [17] J. Shalf, S. Dosanjhi, J. Morrison, Exascale computing technology challenges, in: *High Performance Computing for Computational Science–VECPAR 2010*, Springer, 2011, pp. 1–25.
- [18] US Department of Energy, Aurora, 2015. URL: <http://www.aurora.alcf.anl.gov/>.
- [19] K. Matouš, M.G. Kulkarni, P.H. Geubelle, Multiscale cohesive failure modeling of heterogeneous adhesives, *J. Mech. Phys. Solids* 56 (4) (2008) 1511–1533.
- [20] M. Mosby, K. Matouš, On mechanics and material length scales of failure in heterogeneous interfaces using a finite strain high performance solver, *Modelling Simul. Mater. Sci. Eng.* 23 (8) (2015) 085014.
- [21] K. Matouš, A.M. Maniatty, Finite element formulation for modelling large deformations in elasto-viscoplastic polycrystals, *Internat. J. Numer. Methods Engrg.* 60 (14) (2004) 2313–2333.
- [22] K. Matouš, A.M. Maniatty, Multiscale modeling of elasto-viscoplastic polycrystals subjected to finite deformations, *Interact. Multiscale Mech.* 2 (4) (2009) 375–396.
- [23] K. Matouš, P.H. Geubelle, Multiscale modelling of particle debonding in reinforced elastomers subjected to finite deformations, *Internat. J. Numer. Methods Engrg.* 65 (2) (2006) 190–223.
- [24] F.T. Leighton, A graph coloring algorithm for large scheduling problems, *J. Res. Natl. Bur. Stand.* 84 (6) (1979) 489–506.
- [25] J. Simo, J. Ju, On continuum damage-elastoplasticity at finite strains, *Comput. Mech.* 5 (5) (1989) 375–400.
- [26] M. James, J. Newman, The effect of crack tunneling on crack growth: experiments and CTOA analyses, *Eng. Fract. Mech.* 70 (3) (2003) 457–468.
- [27] W. Lan, X. Deng, M.A. Sutton, Investigation of crack tunneling in ductile materials, *Eng. Fract. Mech.* 77 (14) (2010) 2800–2812.