Technical Report Series
Center for Data and Simulation Science

Axel Klawonn, Stephan Köhler, Martin Lanser, Oliver Rheinbach

Computational Homogenization with Million-way Parallelism using Domain Decomposition Methods

Technical Report ID: CDS-2019-8
Available at http://kups.ub.uni-koeln.de/id/eprint/9480
Submitted on March 29, 2019

Computational Homogenization with Million-way Parallelism using Domain Decomposition Methods

Axel Klawonn · Stephan Köhler · Martin Lanser · Oliver Rheinbach

Abstract Parallel computational homogenization using the well-known FE^2 approach is described and combined with fast domain decomposition and algebraic multigrid solvers. It is the purpose of this paper to show that and how the FE^2 method can take advantage of the largest supercomputers available and those of the upcoming exascale era for virtual material testing of micro-heterogeneous materials such as advanced steel. The FE^2 method is a computational micro-macro homogenization approach which incorporates micromechanical finite element simulations into macroscopic finite element simulations. In this approach, at each Gauß integration point of the macroscopic finite element problem a microscopic finite element problem, defined on a representative volume element (RVE), is attached. Note that the FE^2 method is not embarassingly parallel since the RVE problems are coupled through the macroscopic problem. Numerical results are presented considering different grids on both, the macroscopic and microscopic level. Unstructured as well as structured grids with different irregular domain decompositions are considered on the microscale. Finally, weak scaling results from a few nodes up to a million parallel processes are presented.

This work was supported in part by Deutsche Forschungsgemeinschaft (DFG) through the Priority Programme 1648 "Software for Exascale Computing" (SPPEXA) under grants KL 2094/4-1, KL 2094/4-2, RH 122/2-1, and RH 122/3-2.

Axel Klawonn · Martin Lanser
Department for Mathematics and Computer Science, University of Cologne, Weyertal 86-90, 50931 Köln, Germany, E-mail: axel.klawonn@uni-koeln.de, martin.lanser@uni-koeln.de, url: `http://www.numerik.uni-koeln.de` and Center for Data and Simulation Science, University of Cologne, url: `http://www.cds.uni-koeln.de`

Stephan Köhler · Oliver Rheinbach
Institut für Numerische Mathematik und Optimierung, Fakultät für Mathematik und Informatik, Technische Universität Bergakademie Freiberg, Akademiestr. 6, 09596 Freiberg, E-mail: oliver.rheinbach@math.tu-freiberg.de, stephan.koehler@math.tu-freiberg.de, url: `http://www.mathe.tu-freiberg.de/nmo/mitarbeiter/oliver-rheinbach`

1 Introduction

In this paper, we present algorithms and software for the direct computational homogenization of micro-heterogeneous media in nonlinear structural mechanics. Our approach is based on the combination of an MPI-parallel implementation of the well-known FE^2 computational homogenization method with efficient MPI-parallel iterative solvers, e.g., from domain decomposition (DD) and algebraic multigrid, for the problems on the micro- as well as on the macroscale. An earlier version of our software has already been demonstrated to be applicable to problems in the simulation of dual-phase steel [36,34]. It is the purpose of this paper to show how the FE^2 method can take advantage of the largest supercomputers available for the computational homogenization of micro-heterogeneous materials using million-way concurrency and beyond.

The FE^2 approach is well-established in engineering and has been validated to experiments in numerous publications, e.g., by our collaborators in the EXASTEEL project [12, Fig. 10]; see below for a short description of EXASTEEL. It is our goal, to pave the way for predictive simulations in virtual material testing, with a focus on modern multiphase steels as a show-case, through robust and scalable computational algorithms, optimized software, and advanced modeling; see [53] for a discussion of mathematics-based advanced computing as a means of discovery and innovation. Our project EXASTEEL is part of the German ex-

ascale initiative SPPEXA[1] (Software for Exascale Computing) to develop algorithms and software for the next generation of supercomputers of the exascale era with parallelism beyond 10^7 parallel processes or threads.

The FE^2 method is a computational micro-macro homogenization approach which incorporates micromechanical finite element simulations into macroscopic finite element simulations. It is well established in the engineering community for many years; see, e.g., [61, 23, 48, 59, 45, 24, 25, 60]. In [49], a comparison of the FE^2 method with Direct Numerical Modeling (DNM), i.e., full simulation without homogenization, was given, for a problem of a heterogenous hyperelastic layer. Later, in [50], computational homogenization for an interface problem based on the same implementation, but using a staggered approach different from FE^2, was scaled to 393 216 computing cores.

In this approach, at each Gauß integration point of the macroscopic finite element problem a microscopic finite element problem, defined on a representative volume element (RVE), is attached. The meshes for the finite element simulations on the RVEs are chosen such that they are able to resolve the micro-heterogeneities. The computational micro-macro approach replaces a phenomenological material law on the macroscale; see [60] for an introduction to the FE^2 method. Computational homogenization approaches such as the FE^2 method rely on the assumption of scale separation, i.e., the features of the microstructure are assumed to be magnitudes smaller than the diameter of the macroscopic finite elements. However, detailed micro-macro simulations using the FE^2 method were still out of reach until recently when supercomputers with million-way concurrency became available.

For certain materials the discretization of the macroscopic mechanical part down to its microstructure may become feasible as the computational power grows towards the exascale era. Such approaches are then also sometimes referred to as direct numerical simulations in solid mechanics [9, 10]. Note that in [9, 10] a voxel approximation (using hexahedral finite elements) of the grain structure is used in order to avoid difficulties in the construction of reliable meshes for realistic grain structures.

In the present paper, to avoid this problem we make use of statistically similar RVEs (SSRVEs); see [12, 7, 6], where different SSRVEs were also validated, first against virtual experiments and then against experiments. This approach allows us to replace RVEs with complicated microstructure geometries (either from EBSD – Electron Backscatter Diffraction – measurements or synthetically constructed, e.g., by a tessalation algorithm) by RVEs using simple geometric bodies such as embedded ellipsoids; see, e.g., Fig. 1. The SSRVE approach also helps to reduce the problem size; see Fig. 4.

The resulting algorithm is still computationally expensive but highly parallelizable since the microscopic problems on the RVEs are only coupled through the macroscopic finite element problem. To solve the nonlinear implicit structural mechanics problems on the RVEs efficient implicit solvers are still needed. In our case, we apply parallel domain decomposition solvers, adding an additional layer of concurrency. Recently, in [47], a closely related software framework was presented, combining the FE^2 method with FETI domain decomposition solvers. The framework includes dynamic scheduling and was tested for three-dimensional problems with unstructured meshes on a compute server with two Xeon E5-2650 processors with 8 cores each.

2 Computational homogenization

The FE^2 computational homogenization method, see, e.g., [61, 23, 48, 59, 45, 24, 25, 60], is well known and widely used. In this method, the numerical simulation of a micro-heterogeneous medium is separated into two scales: a macroscopic finite element problem where the microstructure is not resolved and many microscopic finite element problems, each attached to a Gauß point of the macroscopic finite element problem. The microscopic boundary value problems in the FE^2 method are based on the definition of a representative volume element (RVE). While the boundary conditions for the microscopic problems are imposed by the deformation gradients on the macroscale, the upscaling is performed by averaging the stresses over the microscopic problems; see also Fig. 1.

If the microscopic problems are small, the tangent problems on the RVEs can be solved using a sparse direct solver for each RVE problem. For larger RVEs, efficient parallel nonlinear finite element solvers, which should also be robust for heterogeneous problems, have to be incorporated, e.g., a Newton-Krylov method with an appropriate preconditioner such as (algebraic) multigrid or domain decomposition. Recent nonlinear domain decomposition approaches can also be used as, e.g., ASPIN [13, 43, 31, 14, 30, 28, 27, 26] or Nonlinear-FETI-DP [35, 33, 40] and the related Nonlinear FETI-1 or Neumann-Neumann methods [52, 11]. In our implementation of the FE^2 method, the FE2TI package, we focus on Newton-Krylov-FETI-DP and Nonlinear-FETI-DP approaches. We can, however, currently use also PARDISO [54], UMFPACK [15], MUMPS [2], or the algebraic multigrid implementation BoomerAMG [29, 4] from the hypre [20] package. Fast Fourier transform

[1] http://www.sppexa.de

(FFT) based solvers would also be possible for structured meshes but are out of the scope of this paper; see the discussion in Section 3.

Domain decomposition methods of the FETI-DP type are fast solvers for linear and nonlinear problems in solid and structural mechanics and highly scalable to half a million cores [35] and beyond [36]. BoomerAMG is also highly scalable [3], also for linear elasticity [4] if improved interpolations are used.

Using parallel solvers for the microscopic boundary value problems (BVPs) results in several levels of parallelism. On the first level, we have the parallel RVEs, coupled through the macroscopic problem. Each RVE is attached to a Gauß integration point of the macroscopic problem. On the second level, a parallel approach for the solution of the nonlinear PDEs on the RVEs or SSRVEs is used. Third, we will also consider parallelization of the macroscopic problem. Note that the FE^2 method is not embarassingly parallel since the RVE problems are coupled through the macroscopic problem.

2.1 Description of the FE^2 approach

Our description of the computational homogenization FE^2 method follows [60], and more details can be found in [23,60]; also see our earlier works on computational homogenization in the EXASTEEL project [34,36,8]. A similar description can also be found in the dissertation [46].

We first assume that we have a macroscopic boundary value problem or, more precisely, a problem from the field of solid mechanics with typical length scale $\mathfrak{L}$. The characteristic length scale of the microstructure of the material is defined by $\mathfrak{l}$. Therefore, we assume that the microscopic and heterogeneous structure of the material can be resolved in the scale l and that $\mathfrak{L}$ is larger by orders of magnitudes, i.e., $\mathfrak{L} \gg \mathfrak{l}$. Additionally, we assume that we have a representative volume element (RVE), which can effectively describe the microscopic and heterogeneous material properties. The detailed discussion of the notion of a representative volume element [60] is out of the scope of this paper.

Based on the underlying assumptions, it is sufficient to discretize the macroscopic BVP on a domain $\overline{\mathcal{B}}$ with finite elements in the scale of $\mathfrak{L}$ without considering the microscopic structure. Then, at each Gauß point of the macroscopic finite elements, a microscopic BVP, representing the microstructure, is attached. To define an appropriate microscopic BVP, a finite element discretization of the RVE in the scale of the microstructure l is necessary. The appropriate boundary conditions are induced from the macroscopic deformation gradient at the corresponding Gauß point. A related method for multiscale problems, more established in the mathematical community, is the Heterogeneous Multiscale Method (HMM) [17,1]. A recent interesting paper [18] highlights the close relations to the FE^2 method.

Throughout this paper, we will mark macroscopic quantities with bars, as, e.g., the deformation gradient $\overline{F}$ and the first Piola-Kirchhoff stress tensor $\overline{P}$. On the microscale, we use P for the stress and F for the deformation gradient. We do not have a phenomenological material law on the macroscale, instead, the macroscopic quantities are assumed to be volumetric averages of quantities on the microscale. In Fig. 1, a schematic illustration of the homogenization approach FE^2 is given.

Let us remark that in recent years an approach to reduce the computational effort of an FE^2 simulation was to use RVEs with simplified geometries, e.g., consisting of ellipsoids; see [56,55]. These statistically similar RVEs (SSRVEs) can be discretized with a coarser mesh. In Fig. 1, we depict a typical RVE (top right) as well as a typical SSRVE (bottom right).

2.1.1 Microscopic boundary value problems

We first introduce the microscopic boundary value problem defined on a reference configuration $\mathcal{B}_0$, which represents our microscopic RVE, using the corresponding reference variables X. Let us assume that in a deformed state $\mathcal{B}$, the deformation of our reference configuration can be described by $\varphi : \mathcal{B}_0 \to \mathcal{B}$ and the deformation gradient is defined by $F := \nabla\varphi$. Then, without any further external forces, the balance of momentum in a weak formulation with a test function δx writes

$$-\int_{\mathcal{B}_0} \delta x \cdot (\mathrm{Div}_X\ P(F))\ dV = 0. \tag{1}$$

The tensor $P(F)$ is the first Piola-Kirchhoff stress tensor and in contrast to the macroscale, the relation between P and F has to be described by a phenomenological material law, which describes the considered material sufficiently.

In our numerical experiments, we consider dual-phase steels and use an implementation [42] of a J2-elastoplasticity model in FEAP as a material law on the microscale. Work on crystal plasticity is in progress. The parameter choices and different yield stresses for the martensitic and ferritic phases can be found in [12, Fig. 10]. The nonlinear equation (1) is then solved by Newton's method. More precisely, we typically apply a Newton-Krylov-FETI-DP or Nonlinear-FETI-DP approach to (1); see Section 3 for details.

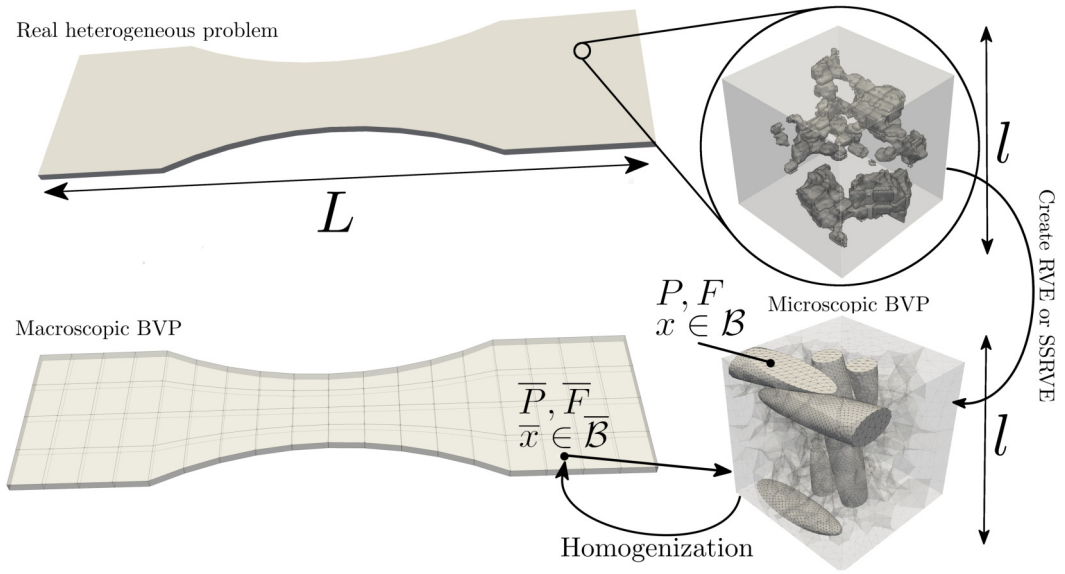

Fig. 1 Illustration of the FE^2 homogenization approach. **Top left:** Realistic and heterogeneous macroscopic boundary value problem of length scale $\mathfrak{L}$. **Top right:** Zoom into the microstructure of the material of length scale l; we have $\mathfrak{L} \gg \mathfrak{l}$. From the microstructure, an RVE or SSRVE can be constructed (see bottom right for SSRVE). **Bottom Left:** Macroscopic and simplified boundary value problem on the domain $\overline{\mathcal{B}}$ with quantities $\overline{P}$ and $\overline{F}$. The deformation gradient $\overline{F}$ induces the boundary conditions of the microscopic BVP on the right. **Bottom right:** Microscopic boundary value problem on domain $\mathcal{B}$, which can be an RVE or SSRVE.

As mentioned before, the boundary conditions on the microscale are induced from the macroscopic deformation gradient $\overline{F}$ at the corresponding Gauß integration point. In the case of Dirichlet conditions, we simply enforce $x := \overline{F}X$ for each boundary node $X \in \partial\mathcal{B}_0$ on the microscale. Here, X are the variables in the reference configuration $\mathcal{B}_0$ and x is the microscopic solution, i.e., x describes the deformed configuration $\mathcal{B}$. As many homogenization approaches do, the FE^2 method assumes that the material can be characterized sufficiently by a repetitive periodic RVE and therefore the use of periodic boundary conditions is more reasonable. We first split the boundary $\partial\mathcal{B}$ into two parts

$$\partial\mathcal{B} = \partial\mathcal{B}^+ \cup \partial\mathcal{B}^-.$$

For each node $X^+ \in \partial\mathcal{B}^+$ exists an associated $X^- \in \partial\mathcal{B}^-$ and both have opposing outer normal vectors. With fluctuation fields $\tilde{w}^+ = x - \overline{F}X^+$ and $\tilde{w}^- = x - \overline{F}X^-$ we enforce the periodic boundary condition

$$\tilde{w}^+ = \tilde{w}^- \quad \forall \text{ pairs } X^+ \in \partial\mathcal{B}^+ \text{ and } X^- \in \partial\mathcal{B}^-.$$

To obtain regular systems, we additionally enforce $\tilde{w}^+ = \tilde{w}^- = 0$ in the eight corners of each RVE.

2.1.2 Homogenization and macroscopic boundary value problem

Similarly to the microscopic boundary value problems, we can formulate the macroscopic problem in a given reference configuration $\overline{\mathcal{B}}_0$ and reference variables $\overline{X}$. Again, the balance of momentum in the weak formulation with a test function $\delta\overline{x}$ writes

$$\int_{\overline{\mathcal{B}}_0} \delta\overline{x} \cdot (\mathrm{Div}_{\overline{X}}\, \overline{P}(\overline{F}) - \overline{f})\, dV = 0, \tag{2}$$

where $\overline{f}$ is a volume force. We here disregard surface forces for simplicity.

On the macroscale, we do not have a material law to deliver an explicit relation between the deformation gradient $\overline{F}$ and the stress tensor $\overline{P}$. Instead, the first Piola-Kirchhoff stress tensor $\overline{P}$ at a macroscopic Gauß point is obtained by homogenization, i.e., as a volumetric average over the Piola-Kirchhoff stresses P of the corresponding RVE. Therefore, we have

$$\overline{P} := \frac{1}{V} \int_{\mathcal{B}_0} P(F)\, dV, \tag{3}$$

where $V := |\mathcal{B}_0|$ is the volume of the corresponding RVE. With (3) and the fact that we consider nonlinear material laws on the microscale, also (2) is nonlinear in the desired solution $\overline{F}$. Thus, also on the macroscale, a Newton approach is used to solve (2). Standard techniques for the globalization of Newton's method can be applied. In our computations, load stepping (pseudo time stepping) is utilized as a homotopy method; see below.

Let us briefly derive the Newton iteration formulated in the macroscopic displacement $\overline{u}$, such that $\overline{F} =$

$\nabla\overline{\varphi} = I + \nabla\overline{u}$, where I is the identity and $\overline{\varphi}$ the macroscopic deformation. We assume that we have an initial value $\overline{u}^{(0)}$. Given finite element shape functions $\overline{\mathbb{N}}_{\overline{T}}$ and their derivatives $\overline{\mathbb{B}}_{\overline{T}}$, we can derive a nonlinear residual in the k-th iteration of Newton's method on a certain finite element from (2):

$$\overline{r}_{\overline{T}}\left(\overline{u}^{(k)}\right) := \int_{\overline{T}} \overline{\mathbb{B}}_{\overline{T}}^{T}\, \overline{P}^{h}\left(I + \overline{u}^{(k)}\right) - \overline{\mathbb{N}}_{\overline{T}}^{T}\, \overline{f}\; dV.$$

Here, $\overline{P}^{h}$ marks the discretized stress tensor on element $\overline{T}$ available in all Gauß points of the element. For the macroscopic tangent, we then have

$$\overline{dk}_{\overline{T}}\left(\overline{u}^{(k)}\right) = \int_{\overline{T}} \overline{\mathbb{B}}_{\overline{T}}^{T}\, \overline{\mathbb{A}}^{h}\left(I + \overline{u}^{(k)}\right) \overline{\mathbb{B}}_{\overline{T}}\, dV, \tag{4}$$

where $\overline{\mathbb{A}}^{h}$ is the discretized macroscopic tangent modulus, which again writes

$$\overline{\mathbb{A}} := \frac{\partial \overline{P}}{\partial \overline{F}} = \frac{\partial}{\partial \overline{F}}\left(\frac{1}{V}\int_{\mathcal{B}_0} P(F) dV\right) \tag{5}$$

and is thus defined in each Gauß point. Using standard finite element routines to assemble the tangent $\overline{DK}\left(\overline{u}^{(k)}\right)$ from $\overline{dk}_{\overline{T}}\left(\overline{u}^{(k)}\right)$ and $\overline{R}\left(\overline{u}^{(k)}\right)$ from $\overline{r}_{\overline{T}}\left(\overline{u}^{(k)}\right)$, we obtain the Newton iteration

$$\overline{u}^{(k+1)} = \overline{u}^{(k)} - \overline{DK}\left(\overline{u}^{(k)}\right)^{-1}\, \overline{R}\left(\overline{u}^{(k)}\right). \tag{6}$$

We now provide a brief description how to compute $\overline{\mathbb{A}}$. Let us therefore remark that $\overline{F}$ is constant on a single RVE. We assume a decomposition of the microscopic deformation gradient $F =: \overline{F} + \widetilde{F}$ on a single RVE into $\overline{F}$ and a fluctuating part $\widetilde{F}$. Using the chain rule and (3), this immediately leads to

$$\overline{\mathbb{A}} := \frac{\partial \overline{P}}{\partial \overline{F}} = \frac{\partial}{\partial \overline{F}}\left(\frac{1}{V}\int_{\mathcal{B}_0} P(F)\; dV\right) \tag{7}$$

$$= \frac{1}{V}\int_{\mathcal{B}_0} \frac{\partial P(F)}{\partial F} : \frac{\partial \overline{F} + \widetilde{F}}{\partial \overline{F}}\; dV \tag{8}$$

$$= \frac{1}{V}\int_{\mathcal{B}_0} \mathbb{A}\; dV + \frac{1}{V}\int_{\mathcal{B}_0} \mathbb{A} : \frac{\partial \widetilde{F}}{\partial \overline{F}}\; dV; \tag{9}$$

see also [60, equation (89)]. The first term in the sum (9) is a volumetric average over the tangent modulus $\mathbb{A}$ of the microscopic problem.

We only have to compute $\overline{P}$ and $\overline{\mathbb{A}}$ after Newton's method has converged on the microscale. Therefore, we can assume an equilibrium state of the weak formulation in (1). Exploiting this fact and using the finite element basis of the RVE, a reformulation of the discretized tangent modulus $\overline{\mathbb{A}}^{h}$ was provided in [60, Section 3.2]:

$$\overline{\mathbb{A}}^{h} := \frac{1}{V}\left(\sum_{T\in\tau}\int_{T} \mathbb{A}^{h}\, dV\right) - \frac{1}{V} L^{T}\,(DK)^{-1}\, L. \tag{10}$$

Here, τ is the finite element discretization of $\mathcal{B}_0$ into finite elements and $\mathbb{A}^{h}$ the discrete microscopic tangent modulus defined in the Gauß points of the finite elements. Then,

$$\frac{1}{V}\left(\sum_{T\in\tau}\int_{T} \mathbb{A}^{h}\, dV\right)$$

is simply the discrete representation of $\frac{1}{V}\int_{\mathcal{B}_0} \mathbb{A}\; dV$ in (9). The second term in (10)

$$\frac{1}{V} L^{T}\,(DK)^{-1}\, L \tag{11}$$

is the discrete form of $\frac{1}{V}\int_{\mathcal{B}_0} \mathbb{A} : \frac{\partial \widetilde{F}}{\partial \overline{F}}\; dV$ exploiting the balance of momentum on the microscale; see [60] for the derivation. Here, the tangential matrix DK of the microscopic BVP is obtained from assembly of finite element matrices

$$k_T := \int_T \mathbb{B}_T^{T} \mathbb{A}^{h} \mathbb{B}_T\, dV,$$

where $T \in \tau$ are finite elements and $\mathbb{B}_T$ are the derivatives of the shape functions. The matrix L has to be assembled also from the element contributions

$$l_T := \int_T \mathbb{A}^{h} \mathbb{B}_T\, dV.$$

Here, L has the dimension $n \times s$, where n is the number of degrees of freedom in the RVE and $s = 4$ in two dimensions and, respectively, $s = 9$ in three spatial dimensions.

Let us remark that DK is simply the tangent in the final Newton step on the RVE and thus the application $(DK)^{-1}L$ can be interpreted as solving a linear system with s right hand sides. If a direct solver is used, a mode for multiple right hand sides can be used or s additional forward backward substitutions have to be performed. However, if using an iterative solver such as a Nonlinear- or Newton-Krylov-FETI-DP method, the iterative solver has to be called s times; of course, information reuse and recycling techniques can be applied to reduce computational cost for the later iteration phases; see Section 4 for details.

We finally provide an algorithmic description of a single load step in Fig. 2, where a Nonlinear- or Newton-Krylov-FETI-DP type method is used to solve the microscopic RVEs. Let us remark that usually the desired

macroscopic deformation cannot be applied in a single step but has to be applied in several consecutive load steps. This often is denoted as pseudo time stepping. The step shown in Fig. 2 has to be repeated for each load step with increasing load. It is possible to use the solution $\overline{u}_n$ from the n-th load step plus additional Dirichlet boundary conditions as an initial value $\overline{u}_{n+1}^{(0)}$ for the Newton iteration in the $(n+1)$-th load step. Alternatively, an extrapolation approach using several former solutions can be beneficial; see also Section 4. Let us remark that we do not consider volume forces $\overline{f}$ in this paper but prescribe a fixed deformation as Dirichlet boundary conditions on parts of the boundary $\partial\overline{\mathcal{B}}_{0,D} \subset \partial\overline{\mathcal{B}}_0$ in incremental load steps.

2.2 Implementation remarks

Our C/C++ implementation, the FE2TI software, uses PETSc 3.5.2 [5] and MPI.

2.2.1 Solving the macro problem

In most of our numerical examples, the macroscopic problem is solved by a sparse direct solver, i.e., UMFPACK or MUMPS, where this is still feasible. In this case, we solve the macroscopic problem redundantly on all MPI ranks. For large FE^2 simulations, the macroscopic problem can also be solved in parallel. We can use a Krylov method combined with, e.g., an algebraic multigrid (BoomerAMG) approach from the hypre library [29] as a preconditioner. Here, the Krylov subspace method and the AMG preconditioner will run on small MPI communicators obtained by an *MPI_Comm_split*, the macro problem is thus solved redundantly on the subcommunicators. A parallel domain decomposition method could also be applied on the macroscale but this approach has not been utilized in this paper. We often reuse the communicators created for the microscopic solvers (see below for details), but for small RVEs an additional split into communicators of efficient size is also possible. The complete macroscopic solution is finally collected on all MPI ranks of the subcommunicator and thus all MPI ranks. Let us also remark that the assembly process of the macroscopic problem is also parallelized and, as should be expected, scales perfectly.

2.2.2 Solving the micro problems

For each Gauß point of the macroscopic problem and thus for each microscopic problem, we introduce a separate MPI communicator. In our implementation, we use *MPI_Comm_split* to create subcommunicators of equal size. Inter-communicator communication is not necessary during the microscopic solves and the averaging of the different microscopic quantities; see points 1 to 4 in Fig. 2. To solve the microscopic BVPs, we use a Nonlinear- or Newton-Krylov-FETI-DP method; see Section 3 for details.

In order to compute $L^T (DK)^{-1} L$ (see (10)), we solve s linear systems with DK as left hand side, i.e., we have nine right hand sides in 3D or four right hand sides in 2D. If an iterative solver is used, this can be an expensive step and, in extreme cases, in our numerical experiments, the computation of the consistent tangent moduli can take up more than 50% of the total time to solution. Approximating $L^T (DK)^{-1} L$ or even discarding the matrix $L^T (DK)^{-1} L$ completely can be feasible alternatives even though superlinear convergence may be lost for the macroscopic problem.

A discussion of different strategies is provided in Section 4. Finally, we use collective communication to provide $\overline{\mathbb{A}}^h$ and $\overline{P}^h$ on all MPI-ranks in order to assemble the linearized macroscopic problem (6). This is efficiently performed by first collecting and summing up all averaged values on the first rank of each microscopic subcommunicator and a consecutive collection step using only those first ranks. Thereby, we avoid global communication involving all MPI ranks but only communicate on independent subsets of MPI ranks at the same time.

3 Parallel domain decomposition solvers

We use iterative multilevel methods to solve our nonlinear solid mechanics problems on the RVEs, i.e., mostly parallel domain decomposition methods of the linear or nonlinear FETI-DP type or algebraic multigrid. Our implementations of nonlinear FETI-DP domain decomposition methods, developed within the SPPEXA EXASTEEL project, have scaled to the full 786 432 cores [36] of the Mira BG/Q Supercomputer (Argonne National Laboratory, USA) for a heterogeneous nonlinear hyperelasticity problem with 60 billion unknowns; also see [35], for details on the method and implementation. In [35], also results for nonlinear domain decomposition on the SuperMUC supercomputer (LRZ, Munich, Germany), the Vulcan supercomputer (Lawrence Livermore National Laboratory, USA), and the JUQUEEN supercomputer (JSC, Jülich, Germany) are included. For robustness, especially for heterogeneous problems, we usually apply sparse direct solvers, i.e., PARDISO [54], UMFPACK [15], or MUMPS [2] as local subdomain solvers. Here, to speed up the Krylov iteration phase, accelerating the forward-backward substitution in the sparse direct solver is of interest [65]. In our implemen-

Choose initial macroscopic deformation $\overline{F}$ which fulfils the boundary conditions
Repeat until convergence (Newton iteration):

1. **Apply boundary conditions to RVE** based on macroscopic deformation gradient; e.g. enforce $x = \overline{F}X$ on the boundary of the microscopic problem $\partial\mathcal{B}$ in the case of Dirichlet constraints.
2. **Solve** microscopic nonlinear boundary value problem using Nonlinear-FETI-DP, Newton-Krylov-FETI-DP, or related methods.
3. **Compute** and return macroscopic stresses as volumetric average over microscopic stresses P^h:

$$\overline{P}^h = \frac{1}{V} \sum_{T \in \tau} \int_T P^h dV.$$

4. **Compute** and return macroscopic tangent moduli as average over microscopic tangent moduli $\mathbb{A}^h$:

$$\overline{\mathbb{A}}^h = \frac{1}{V} \left(\sum_{T \in \tau} \int_T \mathbb{A}^h \, dV \right) - \frac{1}{V} L^T \left(DK\right)^{-1} L$$

5. **Assemble** tangent matrix and right hand side of the linearized macroscopic boundary value problem using $\overline{P}^h$ and $\overline{\mathbb{A}}^h$.
6. **Solve** linearized macroscopic boundary value problem.
7. **Update** macroscopic deformation gradient $\overline{F}$.

Fig. 2 Algorithmic description. Overlined letters denote macroscopic quantities. Blue parts are computations on the microscale and thus the MPI subcommunicators obtained by MPI_Comm_split. Red parts are macroscopic computations performed on all MPI ranks. Algorithmic description taken from [36].

tation it is also possible to use efficient preconditioners such as multigrid for the local problems in nonlinear domain decomposition; see [37,38]. Shared-memory parallelization on the node is an option in our software if a shared-memory parallel subdomain solver (such as PARDISO or BoomerAMG) are used; see [39]. There, parallelization of the finite element assembly, combining PETSc and OpenMP, and the solution of the subdomain problems, using PARDISO, was performed. Good scalability using up to four OpenMP threads for each MPI rank on an Intel Ivy Bridge architecture was observed and incremental improvements were obtained using up to ten threads.

For voxel meshes, see, e.g., Fig. 3 and Fig. 4, geometric multigrid methods or Fast Fourier (FFT) transform based solvers would be efficient alternatives, since they can profit from the tensor structure of such meshes. In [19], the authors compare an FFT-based solution method with standard $Q1$ finite elements for a crystal plasticity problem defined on an RVE. As the FFT-based method exploits the tensor structure of the problem, the authors use structured meshes (voxel meshes) for, both, the FFT-based approach as well as for the $Q1$ finite element meshes. For the FEM problems, a commercial code was applied. The authors concluded that with respect to computational cost the FFT-based approach can be computationally cheaper by one to two orders of magnitude. Indeed, FFT-based computational homogenization approaches have successfully been used by different authors in recent years; see, e.g., [62,32,58, 44,57,16] and references therein. The FFT approaches are valued for the high efficiency when applied to voxel-based RVEs with millions of degrees of freedom, and because voxel meshes from EBSD measurements, as in Fig. 3, can directly be used without further processing. In FFT homogenization, e.g., based on [51], the number of iterations will depend on the contrast; see, e.g., [32].

Note that the approach presented in this paper differs from other approaches, including those using FFT solvers on the microscale. In order to obtain better approximations on the microscale, especially for plasticity, we use second order (P2) finite elements instead of simple linear tets (P1) or trilinear hexahedral elements (Q1). We then use an approach (SSRVEs) to construct RVEs using simple geometries, e.g., ellipsoids. However, the advantages of SSRVEs can only be exploited fully when using unstructured meshes, and we therefore use unstructured tetrahedral meshes, well adapted to the geometry. Our parallel domain decomposition methods can cope with these unstructured meshes, and the number of Krylov iterations will be only slightly higher. We also use direct sparse solvers on the subdomains for high robustness of the numerical methods and usually obtain quadratic convergence of the overall Newton scheme; cf. Fig. 18. We will see in our numerical results that, for our SSRVEs, unstructured meshes are clearly favorable; see Section 4.2. We will also see that the unstructured meshes can also be allowed to be significantly, i.e., more than an order of magnitude, smaller than the structured meshes. Using adaptive mesh refinement could further expand the advantage but this is out of the scope of this paper.

Recent versions of PETSc include an efficient implementation of iterative substructuring methods [66].

However, we apply our own parallel implementation [35, 41] which is earlier and also includes our most recent versions of nonlinear domain decomposition methods [40].

Let us briefly describe the linear and nonlinear FETI-DP domain decomposition methods, which we will utilize to solve the RVE problems on the microscale. In our computational homogenization approach, the micro problems are usually more challenging than the macro problem as a result of the heterogeneities, which have to be homogenized.

3.1 Parallel linear and nonlinear domain decomposition of FETI-DP type

Given a nonoverlapping domain decomposition $\Omega_1, \dots, \Omega_N$ of the computational domain $\Omega \subset \mathbf{R}^3$, we introduce the (generally) nonlinear operators $K_1, \dots, K_N$, defined on the subdomains, and the corresponding right hand sides $f_1, \dots, f_N$. In nonlinear domain decomposition, we usually obtain the operators $K_1, \dots, K_N$ from the minimization of a nonlinear (e.g., hyperelastic) energy on the subdomains; see [33, equation (2.4)]. In Newton-Krylov-DD methods, the operators $K_1, \dots, K_N$ are linear as the domain decomposition is performed after Newton linearization.

In our FE^2 simulations, the domain Ω will be an RVE $\mathcal{B}_0$ and $K_1, \dots, K_N$ as well as $f_1, \dots, f_N$ are determined by (1). We define the block vectors $u = (u_1, \dots, u_N)$, $K(u)^T = [K_1(u_1)^T \dots K_N(u_N)^T]^T$, $\quad f^T = [f_1^T, \dots f_N^T]^T$.

Using finite element assembly in primal variables, we obtain the partially assembled nonlinear operator $\widetilde{K}$,

$$\widetilde{K}(\tilde{u}) = R_\Pi^T K(R_\Pi \tilde{u})$$

and the corresponding right hand side $\tilde{f} = R_\Pi^T f$. As in linear FETI-DP methods, the global coupling by partial finite element assembly is introduced to make the local problems invertible and to obtain numerical scalability. As a result, the Jacobian $D\widetilde{K}(\tilde{u})$ can be assumed to be invertible, while it still maintains a favorable, partially decoupled, structure, i.e.,

$$D\widetilde{K}(\tilde{u}) = \begin{pmatrix} DK_{BB}^{(1)}(\tilde{u}) & & D\widetilde{K}_{B\Pi}^{(1)}(\tilde{u}) \\ & \ddots & \vdots \\ & DK_{BB}^{(N)}(\tilde{u}) & D\widetilde{K}_{B\Pi}^{(N)}(\tilde{u}) \\ D\widetilde{K}_{\Pi B}^{(1)}(\tilde{u}) \dots & D\widetilde{K}_{\Pi B}^{(N)}(\tilde{u}) & D\widetilde{K}_{\Pi\Pi}(\tilde{u}) \end{pmatrix}.$$

We now define nonlinear FETI-DP methods as iterative methods for the solution of the nonlinear system (see [33, 40])

$$A(\tilde{u}, \lambda) := \begin{bmatrix} \widetilde{K}(\tilde{u}) + B^T\lambda - \tilde{f} \\ B\tilde{u} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \tag{12}$$

Since the nonlinear problem $\widetilde{K}(\tilde{u}) = \tilde{f} - B^T\lambda$ can very efficiently be parallelized, such methods can be highly scalable and nonlinear problems with billions of degrees of freedom [35] can be solved in a few minutes. As in standard, linear FETI-DP methods [64, 22, 21, 41] the linear constraint $B\tilde{u} = 0$ enforces the continuity across the subdomain boundaries, using Lagrange multipliers λ. However, as opposed to linear FETI-DP methods, the continuity is only achieved at convergence of the Newton iteration; see [33] for details.

The (generally) nonlinear system (12) can be solved using different strategies; see [40]. These methods are all equivalent to the classical FETI-DP method if K is a linear operator.

For our problems on the RVEs, we can apply Newton's method directly to (12). This was denoted Nonlinear-FETI-DP-1 in [33]. We will either use a classical Newton-Krylov FETI-DP approach, i.e., the microscopic BVP is first linearized and the FETI-DP method is applied to the linearized system, or Nonlinear-FETI-DP-1 is used, i.e., Newton linearization is applied to (12), and the Newton correction $\left[\delta\tilde{u}^{(k)}; \delta\lambda^{(k)}\right]^T$ can be computed from

$$\begin{bmatrix} D\widetilde{K}(\tilde{u}^{(k)}) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \delta\tilde{u}^{(k)} \\ \delta\lambda^{(k)} \end{bmatrix} = \begin{bmatrix} \widetilde{K}(\tilde{u}^{(k)}) + B^T\lambda^{(k)} - \tilde{f} \\ B\tilde{u}^{(k)} \end{bmatrix}. \tag{13}$$

A reduction to the Lagrange multipliers leads to a linear system of the form

$$\mathcal{F}\, \delta\lambda^{(k)} = d, \tag{14}$$

which is solved iteratively by a Krylov method, e.g., GMRES, throughout this paper, preconditioned by the standard FETI-DP Dirichlet preconditioner [64].

4 Numerical results

If not marked otherwise, we use a Newton-Krylov-FETI-DP approach for the solution of the microscopic RVE problems, i.e., we have used the most conservative choice among our FETI-DP methods. We also present results using Nonlinear-FETI-DP-1 on the Theta supercomputer later on. This is the most conservative choice among the recent Nonlinear-FETI-DP and Nonlinear-BDDC methods [35, 33, 40]. In our numerical experiments, if not denoted otherwise, the macroscopic problem is discretized using piecewise linear triangular elements (P1) in 2D and piecewise trilinear or quadratic brick elements (Q1 or Q2) in 3D. In all our experiments we stop the macroscopic Newton iteration if the norm of the update is smaller than $1e-6$. We use a tolerance

#subdomains	unstructured meshes	
	#elements	#degrees of freedom
30	23 537	103 266
64	47 296	199 125
110	222 489	921 648
#subdomains	structured meshes	
	#elements	#degrees of freedom
30	20 480	95 523
64	69 120	311 475
110	214 375	945 108

Table 1 Meshes for the SSRVEs. We use P2 finite elements.

of $1e-8$ for the microscopic Newton iteration and use a relative stopping tolerance of $1e-9$ for all Krylov subspace methods.

We use computing time on the JUQUEEN supercomputer (Jülich Supercomputing Centre), a large BG/Q (PowerBQC 16C 1.6 Ghz) installation with a total of 458 752 cores and 1 GB of memory for each core. JUQUEEN was Europe's fastest supercomputer in 2015 and is still ranked 22nd in the TOP500 list of November 2017.

4.1 Simulation using a voxel-based RVE from EBSD measurements

In this section, we consider FE^2 simulations using a structured mesh based on voxel data obtained from EBSD (Electron Backscatter Diffraction) measurements as an RVE. We use P2 finite elements on the microscale, where each voxel is decomposed into six ten-noded tetrahedra. The RVEs consisting of 663 552 finite elements and 2 738 019 degrees of freedom. Each RVE is decomposed into 512 subdomains. We consider a plate with a hole discretized with 64 finite elements and thus using 512 RVEs. The sum of the number of degrees of freedom on the microscale is thus approximately 1.4 billion. We simulate four different deformations. We twist the plate and apply different pulling and pressing forces; see Fig. 3 (top).

Each of the simulations has been performed using 262 144 MPI ranks using 131 072 cores of the JUQUEEN supercomputer. The numerical scheme remained stable and reliable results have been obtained; see Fig. 3. The I/O-time to write all data for all RVEs to the file system always stayed below 2% of the total runtime.

These results show the scalability of our implementation and demonstrate its ability to run efficiently on large supercomputers, consistent with the earlier results from [34,36,8]. Therefore, in the following sections, we fokus on further optimizing and tailoring the RVE solver. We also investigate if we can reduce the size of the RVEs.

4.2 Considering different discretizations of the microscale

In this section, we study the influence of the structure and resolution of the discretization of a chosen RVE or SSRVE on the homogenized solution on the macroscale. The construction of the SSRVEs using a fixed number of ellipsoidal inclusions based on EBSD measurements is discussed in [55]; also see [12]. Nonetheless, the resolution of the discretization of the SSRVEs is comparably coarse in [55]. Here, we present a comparison of macroscopic and microscopic results using meshes of different refinement levels; see Fig. 4 and Fig. 5 for different SSRVE meshes. We consider structured grids (Fig. 4), which cannot resolve the ellipsoidal inclusions very well, as well as unstructured grids (Fig. 5). These investigations also provide us with a grid convergence study. As a macroscopic problem, we extend a symmetric plate discretized with 72 finite elements in x-direction; see Fig. 6 (left) for the undeformed geometry. We first apply 11 load steps with a deformation of 0.025 percent of the current deformed state in each load step; see Fig. 6 (right) for a visualization of the load.

For our meshes, see Table 1. Our largest unstructured mesh, discretizing the SSRVE using two ellipsoids [55], consists of 222 489 finite elements (see Fig. 5 right) and is chosen as our reference. When we to refer to the error (see Fig. 8), we mean the difference to the solution on the reference mesh.

In Fig. 7, we present the von Mises stresses in the macroscopic problem after 11 load steps using the reference discretization. To compare the effects of using different SSRVE meshes on the macroscopic results, in Fig. 8, we depict the relative difference in the von Mises stresses between the reference solution from Fig. 7 and solutions obtained using coarser unstructured and structured grids.

For the unstructured grids – see Fig. 8 (rows four to five) – we see only small differences for both grids and also observe convergence to the reference solution from the coarser grid (row four) to the finer one (row five). On the other hand, we obtain significant differences using structured grids, even though the finest structured SSRVE is of the same size as the unstructured reference SSRVE; see Fig. 8 (rows one to three). Therefore, we can conclude that the choice of the discretization of the SSRVE has a significant impact. Our unstructured grids result in a significant better approximation of the reference solution than all structured grids, even if the number of finite elements is an order of magnitude smaller.

All in all, in our experiments, the von Mises stresses have been slightly higher when structured grids are

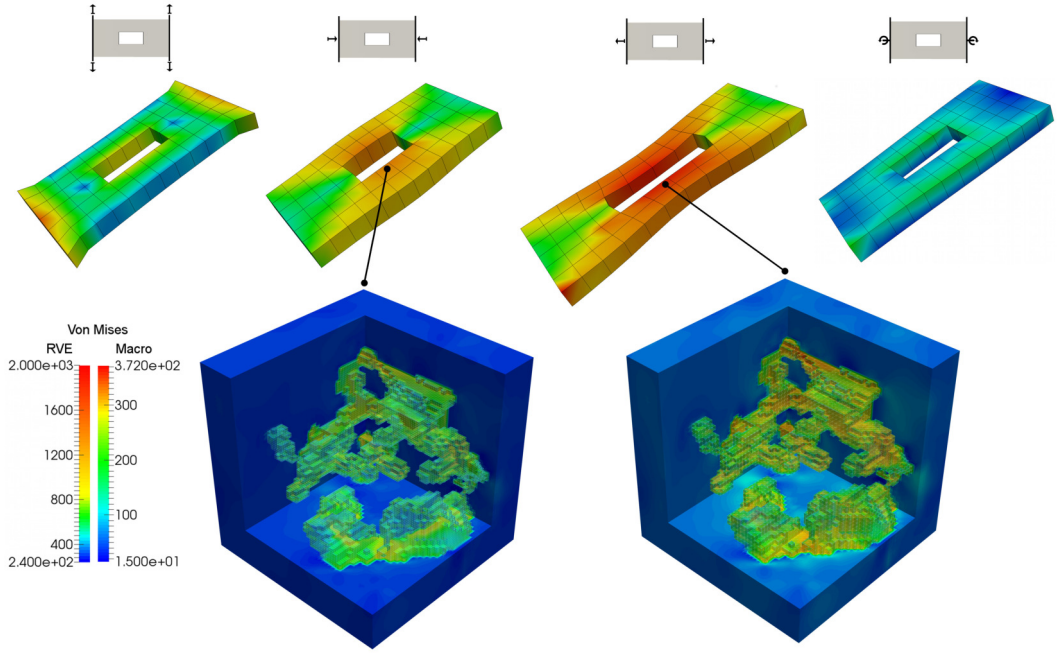

Fig. 3 Four different types of macroscopic deformations. Simulations performed with RVEs with 663 552 finite elements and 2 738 019 degrees of freedom. The geometry of the RVEs corresponds to a small cubic part of a larger structure obtained from EBSD measurements of a dual-phase steel. The computations have been performed with 262 144 MPI ranks on 131 072 cores of JUQUEEN. Based on data from [12].

used. First, this is a result of high stress peaks caused by the unsmooth resolution of the surfaces of the ellipsoids. Second, the minimal stresses in the SSRVEs appear to be slightly higher in the case of structured meshes. The first effect can be observed in the stress distributions in Fig. 9, where we depict the unstructured SSRVE alongside the structured ones. In Fig. 9, four different Gauß points A,B,C, and D are considered as well as two different grid resolutions.

We additionally provide data on the stress peaks and the average stresses in Gauß points A and C in Fig. 10 (left and middle). The second effect is depicted in Fig. 10 (right) for the same two Gauß points A and C.

4.3 Improvements of the numerical scheme

In this section, we study techniques to improve the efficiency of our simulation. We also discuss the use of an approximate tangent in the solution of the macroscopic problem.

Our macroscopic test problem in this section consists of only 16 macroscopic finite elements, resulting in 128 microscopic RVE problems. Here, we always use the unstructured SSRVE with 47 296 finite elements decomposed into 64 FETI-DP subdomains; see Table 1. In this section, a 1% deformation in x-direction using 21 load steps was applied on the macroscale; see also Fig. 11 for the macroscopic solution. We keep the stopping criterion of the macroscopic and microscopic Newton iterations fixed at $1e-6$ and, respectively, $1e-8$. We keep the relative stopping criterion of the FETI-DP solver at $1e-9$ since the high accuracy is necessary for fast Newton convergence on the microscale.

We consider three improvements of our numerical scheme.

1. **Lambda-recycling in FETI-DP:** We reuse the solutions λ of earlier Newton steps as initial values for the FETI-DP solvers; see (14).
2. **Approximate tangent modulus:** We vary the stopping tolerance for the linear FETI-DP solves necessary for the computation of the consistent macroscopic tangent moduli; see (11). Stopping early leads to an approximate tangent on the macroscale.
3. **Extrapolation on the macroscale:** We use an extrapolation approach to obtain a proper initial value for the Newton iteration on the macroscale.

Lambda-recycling In general, the idea of lambda-recycling in Newton-Krylov-FETI-DP is to use the solution $\lambda^{(k)}$ of the linear system from the k-th Newton iteration as an initial value for FETI-DP in the $(k+1)$-th Newton iteration, e.g., to compute $\lambda^{(k+1)}$.

We apply Newton-Krylov-FETI-DP on each RVE in each macroscopic step and use lambda-recycling on each RVE individually. Additionally, we use the solu-

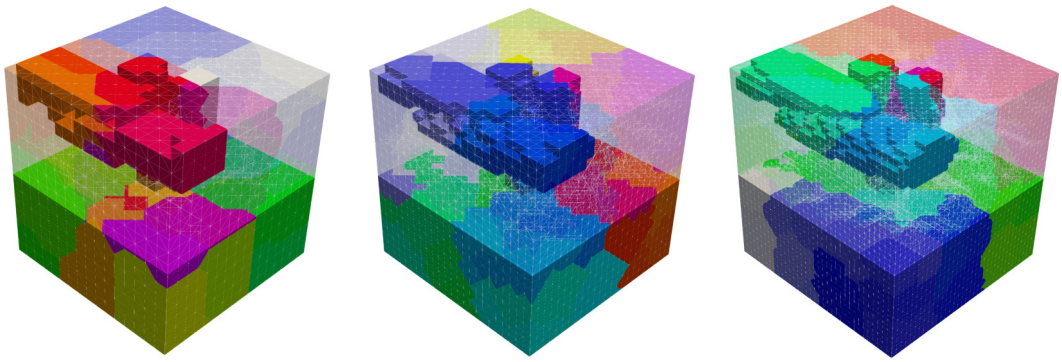

Fig. 4 Structured discretization with similar numbers of finite elements as the unstructured in Fig. 5 and domain decomposition of an SSRVE with two ellipsoidal inclusions. The colors show the decomposition of the lower half of the cube and the ellipsoids. **Left:** 20 480 finite elements and 30 subdomains. **Middle:** 69 120 finite elements and 64 subdomains. **Right:** 214 375 finite elements and 110 subdomains.

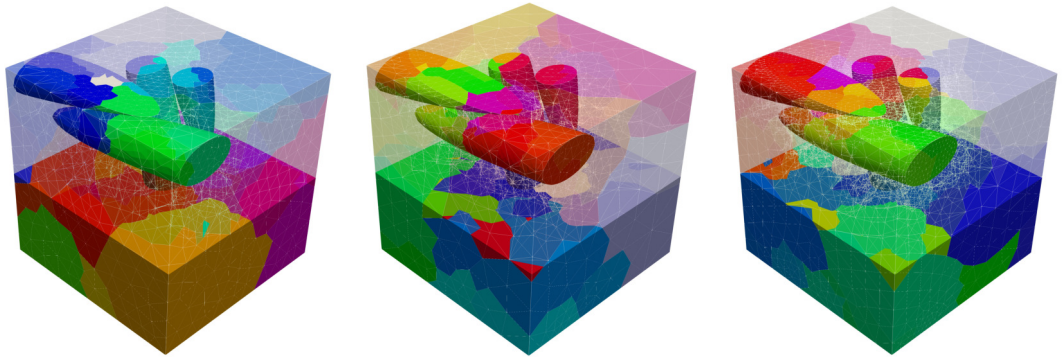

Fig. 5 Unstructured discretization and domain decomposition of an SSRVE with two ellipsoidal inclusions. The colors show the decomposition of the lower half of the cube and the ellipsoids. **Left:** 23 537 finite elements and 30 subdomains. **Middle:** 47 296 finite elements and 64 subdomains. **Right:** 222 489 finite elements and 110 subdomains.

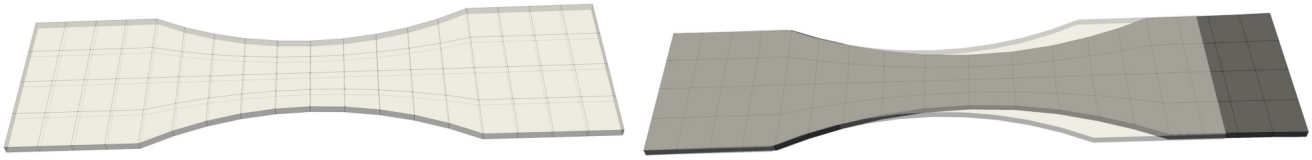

Fig. 6 Macroscopic Problem. **Left:** undeformed Geometry; **Right:** scaled deformation (factor 50) after 10 steps.

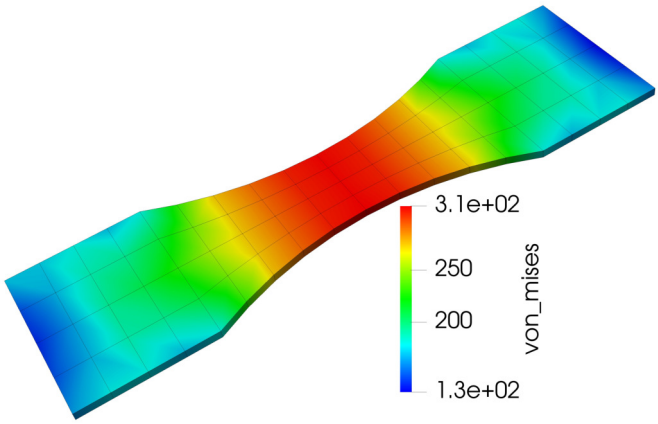

Fig. 7 Von Mises stresses after 10 load steps using the largest of the three SSRVE.

tion from the last Newton-Krylov-FETI-DP iteration of the previous macroscopic step as an initial value for the first Newton-Krylov-FETI-DP iteration of the current macroscopic step. Let us remark that nonlinear FETI-DP methods do not need any lambda-recycling strategies since the reuse of information automatically arises from the nonlinear scheme.

It is interesting to note that when solving for the nine right hand sides in (14), these right hand sides are rather different and the lambda-recycling approach between those nine solves does not seem to be beneficial. It turns out that it is more efficient to use lambda-recyling from the nine individual solutions from the previous macroscopic iteration. This strategy is thus standard in our FE2TI software.

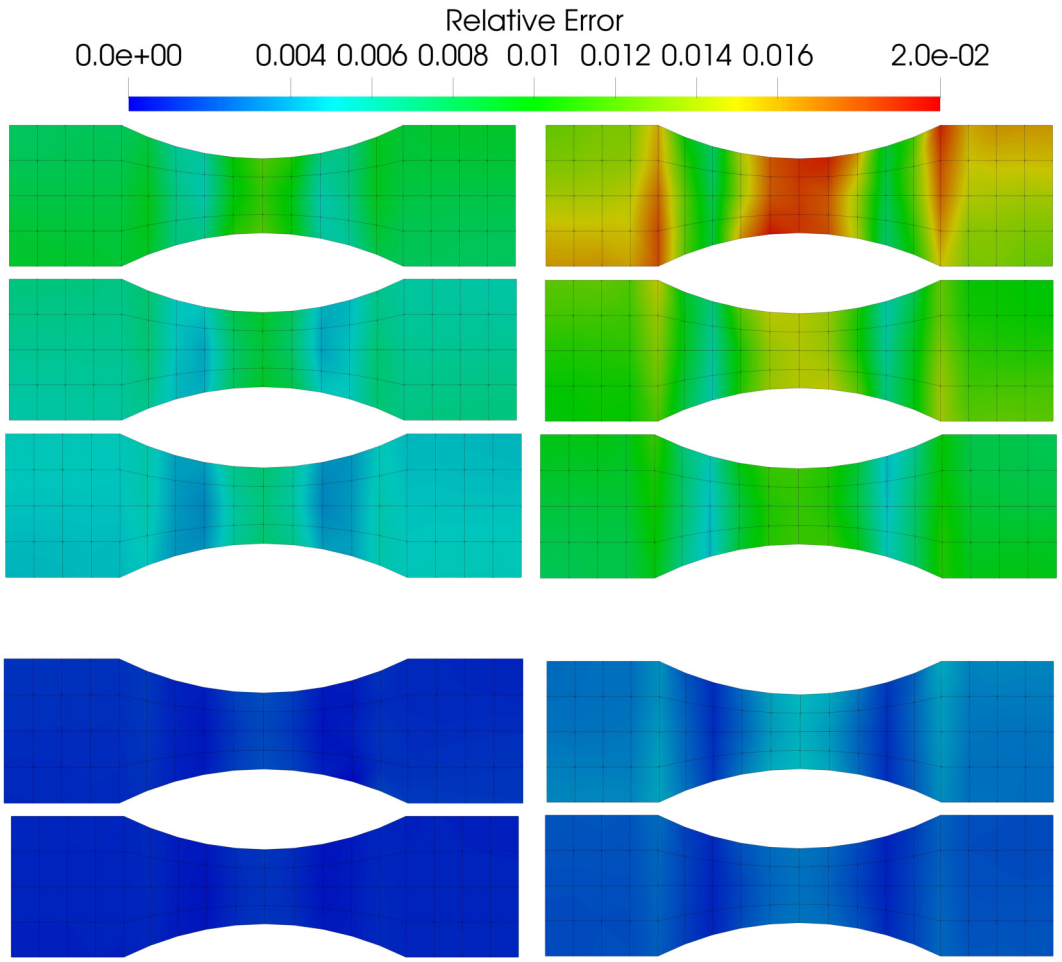

Fig. 8 Pointwise relative error in von Mises stresses on the macroscale using different SSRVEs after 7 (left column) and 11 load steps (right column). The reference solution is obtained using an SSRVE with a large unstructured mesh (see bottom right in Fig. 5). For the results in the first three rows, we use SSRVEs with structured meshes and for the last two rows SSRVEs with unstructured meshes - **First row:** SSRVE with structured mesh decomposed into 30 subdomains and 20 480 finite elements. **Second row:** SSRVE with structured mesh decomposed into 64 subdomains and 69 120 finite elements. **Third row:** SSRVE with structured mesh decomposed into 110 subdomains and 214 375 finite elements. **Fourth row:** SSRVE with unstructured mesh decomposed into 30 subdomains and 23 537 finite elements. **Fifth row:** SSRVE with unstructured mesh decomposed into 64 subdomains and 47 296 finite elements.

Approximate tangent modulus The nine additional linear FETI-DP solves to obtain a consistent tangent modulus and thus a consistent tangent on the macroscale (see (14)) can exhibit a significant portion of the runtime, especially when using iterative solvers. Nevertheless, it is necessary in order to obtain quadratic convergence of Newton's method on the macroscale. We suggest to reduce the accuracy of the nine linear solves and thus use an approximate tangent on the macroscale. We therefore tested different stopping tolerances for FETI-DP for the computation of the consistent tangent modulus.

Extrapolation of macroscopic solutions This approach is rather simple. If an extrapolation of the macroscopic iterates is activated in the FE2TI package, the initial value for the $(k+2)$-th macroscopic load step $u^{(0)}_{macro,k+2}$ is chosen as a linear extrapolation from the solutions of the k-th and $(k+1)$-th load steps $u^{*}_{macro,k}$ and, respectively, $u^{*}_{macro,k+1}$, i.e.,

$$u^{(0)}_{macro,k+2} = 2u^{*}_{macro,k+1} - u^{*}_{macro,k}\,.$$

4.4 Effects of the improvements

Without any of the improvements introduced in the previous section, the FE^2 simulation takes 10 445.5s to perform 21 macroscopic load steps using a total of 82 macroscopic Newton steps. Using lambda-recycling alone, the runtime can be reduced to 7 693.0s, which is 1.36 times faster.

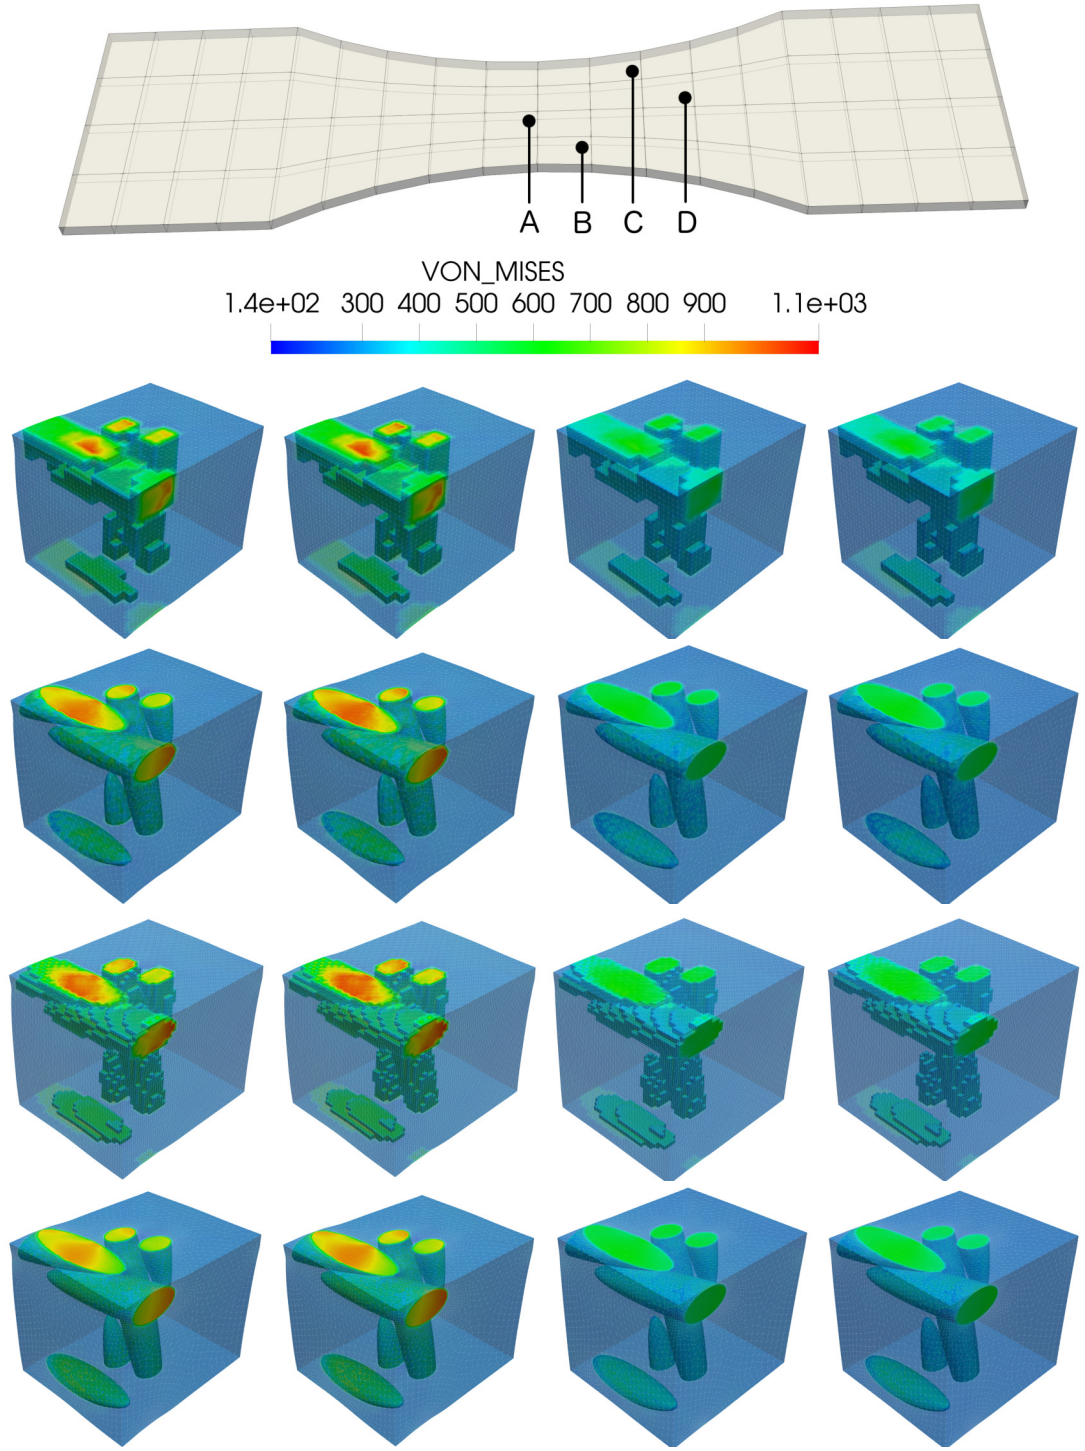

Fig. 9 From left to right: von Mises stresses in four Gauß points A,B,C, and D after 10 load steps - the locations of the points A to D is depicted above. The deformation of the SSRVEs is scaled by a factor of 20. **Rows one and two:** coarse structured and unstructured discretizations with 20 480 and, respectively, 23 537 finite elements, both decomposed into 30 irregular subdomains. **Rows three and four:** fine structured and unstructured discretizations with 214 375 and, respectively, 222 489 finite elements, both decomposed into 110 irregular subdomains.

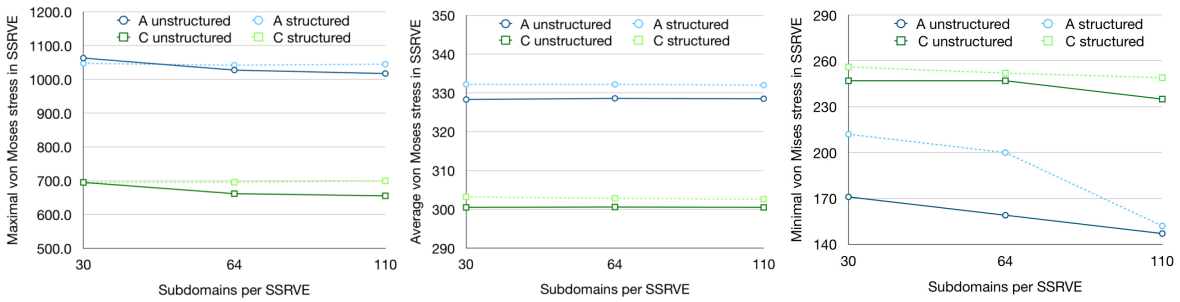

Fig. 10 Maximal stress **(left)**, average stress **(middle)**, and minimal stress **(right)** after 10 load steps in structured and unstructured SSRVEs attached to Gauß points A and C - see Fig. 9 for location of Gauß points; see Table 1 for the number of finite elements.

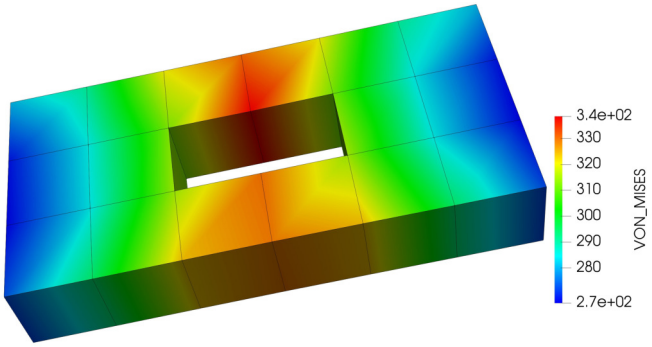

Fig. 11 Macroscopic solution of a small model problem after 17 load steps.

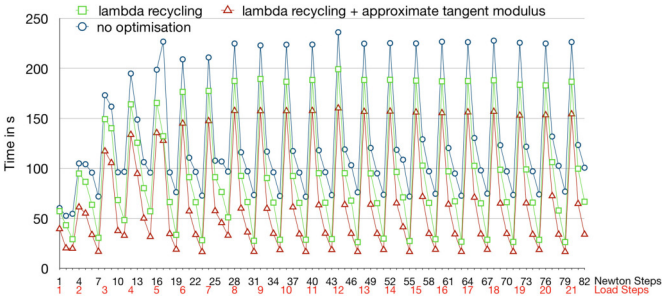

Fig. 12 Runtime of the macroscopic Newton iterations. Comparison between our software FE2TI without optimization, with lambda-recycling, and with lambda-recycling plus using an approximate tangent modulus with tolerance $1e-3$. Additionally, the macroscopic load steps are marked in red on the x-axis.

Choosing different accuracies for the computation of the consistent tangent modulus ($1e-9, 1e-6, 1e-3, 1e-1$), quadratic convergence of the macroscopic Newton iterations was lost only for $1e-1$, where the total number of Newton iterations in 21 load steps increases slightly from 82 to 88; see also Table 2.

The runtime is reduced from 7 693.0s to 5 573.5s for the tolerance $1e-3$, which is our best choice and corresponds to a factor of 1.38 of reduction in runtime.

In Fig. 12, we provide a comparison of three variants, namely without lambda-recycling and with lambda-recycling and different tolerances $1e-9$ and $1e-3$ for the consistent tangent.

Finally, adding the macroscopic extrapolation of solutions, the number of macroscopic Newton iterations is further reduced to 66, and thus also the runtime is reduced to 4 721.0s; see Table 2. Combining all three improvements, we can accelerate the computations by a factor of 2.2.

4.5 Large run using the improvements

Combining our different strategies, i.e., λ-recycling, approximation of the consistent tangent moduli, and extrapolation of macroscopic solutions, we are able to simulate 81 load steps with a 2.1% total deformation in 4 hours and 39 minutes using 18 432 JUQUEEN cores. Here, we use the unstructured SSRVE with 47 296 finite elements, since the approximation of the reference solution showed to be accurate enough. We depict the final state of the solution in Fig. 13.

4.6 Improving scalability by exploiting parallelism on the macroscale

As mentioned above, the capability to solve the macroscopic problem in parallel with an iterative Krylov subspace method (e.g., CG or GMRES) and, e.g., the BoomerAMG preconditioner [29] was included. When using BoomerAMG, we always use *HMIS* coarsening, long range *ext+i* interpolation, and a nodal coarsening approach; see [4] for a discussion of efficient parameters for elasticity.

The macroscopic problem is assembled and solved in parallel on a subset of ranks of the MPI communicators which are assigned to the RVEs. In our tests, these communicators have a reasonable size, e.g., 64 MPI ranks if the RVE is decomposed into 64 subdomains. In the following tests, we always use between 2 and 64 MPI

load steps	lambda recycling	stopping tol. consistent tangent	extrapolation	Newton It. Macroscale	Total Runtime
21	no	1e-9	no	82	10 445.5s
21	yes	1e-9	no	82	7 693.0s
21	yes	1e-6	no	82	6 484.2s
21	yes	1e-3	no	82	5 573.5s
21	yes	1e-1	no	88	5 688.5s
21	yes	1e-3	yes	66	4 721.0s

Table 2 Improvements using lambda-recycling, inexact macroscopic Newton's method, and extrapolation of the macroscopic problem.

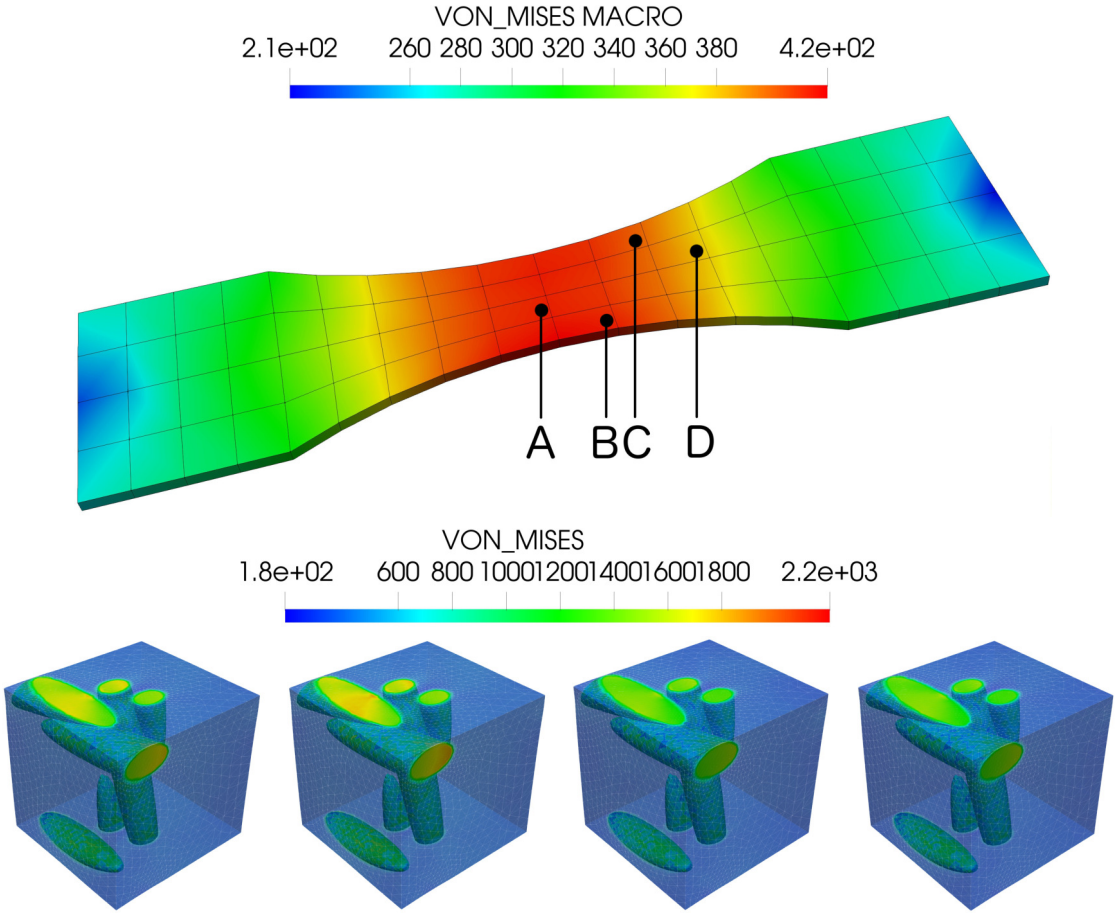

Fig. 13 Von Mises stresses after 81 load steps and a total deformation of approximately 2.1% of the macroscopic reference configuration. Four SSRVEs A,B,C, and D (from left to right) are depicted.

ranks and thus AMG is only applied on a small number of ranks. The result of the macroscopic problem is distributed to the complete communicator and then available on all MPI ranks. We perform two different weak scaling tests using the unstructured SSRVE with 64 subdomains and the best parameter settings found before, i.e., λ-recycling, approximation of the consistent tangent moduli, and extrapolation of macroscopic solutions are used.

We perform 13 load steps for an uniaxial tension test of a macroscopic cube. In the first weak scaling test, we use four MPI ranks for each JUQUEEN core and use up to 1.04 million MPI ranks; see Fig. 14 (top left and top right) and Fig. 15 for the results. The results in Fig. 15 show significant improvements using parallelization of the macro problem.

In the second test, we use two MPI ranks per core and scale up to the complete machine; see Fig. 14 (bottom left and bottom right). Considering the average time per macroscopic Newton iteration, we obtain an excellent efficiency of 74% on 1.04 million MPI ranks using AMG for the macroscopic problem and, respectively, for the setup using two MPI ranks per compute core, 78% parallel efficiency on the whole machine. In both cases the scalability is increased drastically in

comparison to the approach using a direct solver on the macroscale.

4.7 Production runs on Theta using Nonlinear-FETI-DP-1

In this section, we present some results for large production runs using Nonlinear-FETI-DP-1 to solve the microscopic problems.

Theta is a large Intel Xeon Phi (Xeon Phi 7230, 64 Cores, 1.3 Ghz, "Knights Landing") x86 installation at Argonne Leadership Computing Facility (ALCF) with a total of 231 424 cores and ranked 18th in the TOP500 list of November 2017. Theta is a very capable supercomputer. Unfortunately, Intel has discontinued its Xeon Phi line and Theta's successor named Aurora, planned to be the first exascale system of the USA in 2021, will use a different architecture.

Analogously to the last subsection, we use the extrapolation approach and an approximate consistent tangent. With Nonlinear-FETIDP-1, λ-recycling is not necessary. Indeed, the most conservative nonlinear FETI-DP method – Nonlinear-FETI-DP-1 – shows to be as robust and efficient as Newton-Krylov-FETI-DP with λ-recycling.

4.8 First run on Theta

We are able to simulate 171 load steps with a total deformation of 4.3% in exactly 4 hours walltime using 36 864 cores of Theta (Argonne National Laboratory, USA). Here, we use the unstructured SSRVE with 47 296 finite elements (see Table 1) since the approximation of the reference solution showed to be accurate enough. The problem has 302 million degrees of freedom in the microscale and 570 degrees of freedom on the macroscale

We depict the final state of the solution in Fig. 16 and the development of the von Mises stresses over time in Fig. 17 for six different vertices.

We also show the number of necessary Newton and FETI-DP iterations on the microscale. In Fig. 18 the effect of plastification of certain RVEs can be observed nicely by observing the number of Nonlinear-FETI-DP-1 and GMRES iterations in different load steps. Clearly, starting at load step number four, RVEs A and B need more Newton steps to converge than C and D, which marks the point in time when these two RVEs start to show a plastic behavior.

This introduces a certain load imbalance and can also be observed in the corresponding runtimes; see Fig. 19. This load imbalance shows to be of minor interest and does not affect the performance severly, since, after approximately 20 load steps, all RVEs in the critical region show a plastic behavior and the load imbalance vanishes. To illustrate this effect, we also present the Newton iterations and Nonlinear-FETI-DP-1 runtimes of load steps 70 to 79; see Fig. 18 (right) and Fig. 19 (right). We also present the number of GMRES iterations per Nonlinear-FETI-DP-1 solve in Fig. 20. The average number of GMRES iterations per linear solve is approximately 40 for each of the four RVEs, which is a satisfactory result.

4.9 Second run on Theta

We finally present results for a refined macroscopic mesh with 1 792 SSRVEs computed on Theta using Nonlinear-FETI-DP-1. We simulate 319 load steps using 114 688 cores of Theta and obtain a deformation of 7.7% in x-direction. The problem has 693 million degrees of freedom in the microscale and 1 305 degrees of freedom on the macroscale. Some results for different RVEs and different load steps are presented in Fig. 21.

4.10 Conclusion

We have presented a framework for parallel computational homogenization using the FE^2 approach. The usage of parallel domain decomposition solvers on the microscale and parallel algebraic multigrid solvers on the macroscale allows large multiscale simulations for micro-heterogeneous media. We have shown FE^2 simulations with million-way parallelism and billions of degrees of freedom. Larger simulations will be possible once exascale supercomputers will become available.

Acknowledgments The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de) and **JUQUEEN** [63] at Jülich Supercomputing Centre (JSC, www.fz-juelich.de/ias/jsc). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK) and Nordrhein-Westfalen (MKW). This research used resources (**Theta**) of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. The authors acknowledge the use of data from [12] provided through a collaboration in the DFG SPPEXA project **EXASTEEL**.
The authors would also like to thank **Jörg Schröder**, **Dominik Brands**, and **Lisa Scheunemann** (University of Duisburg-Essen) for providing the SSRVEs, the J2 plasticity model, and many fruitful discussions.

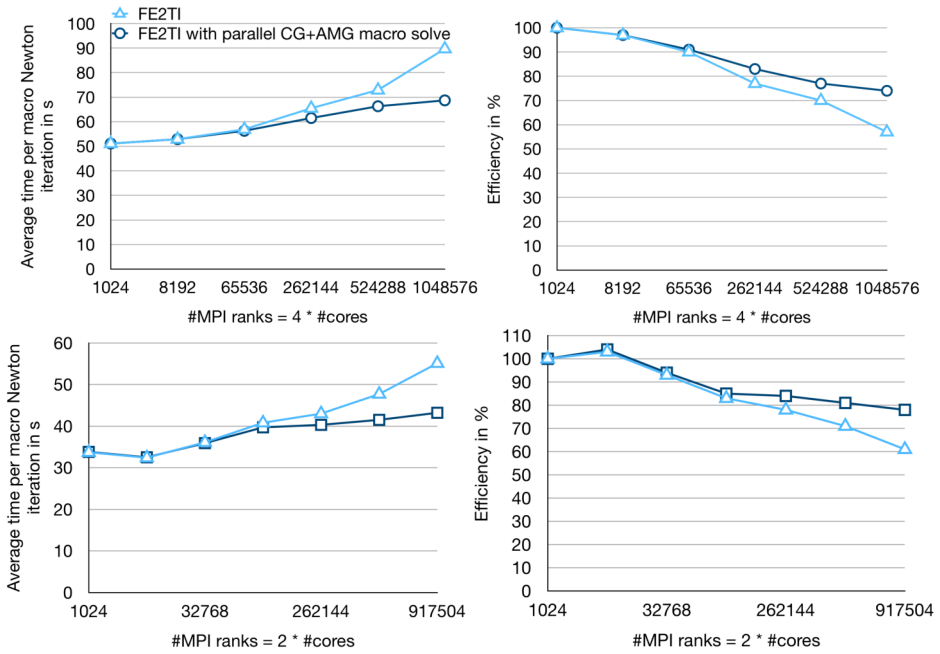

Fig. 14 Top Left: Average Time for each macroscopic Newton iteration scaling to 1.02 million MPI ranks on JUQUEEN using CG plus AMG preconditioner or a direct solver, respectively. **Top Right:** Corresponding efficiency to runtimes presented in top left figure. **Bottom Left:** Average Time for each macroscopic Newton iteration scaling up to the complete JUQUEEN. **Bottom Right:** Corresponding efficiency to runtimes presented in bottom left figure.

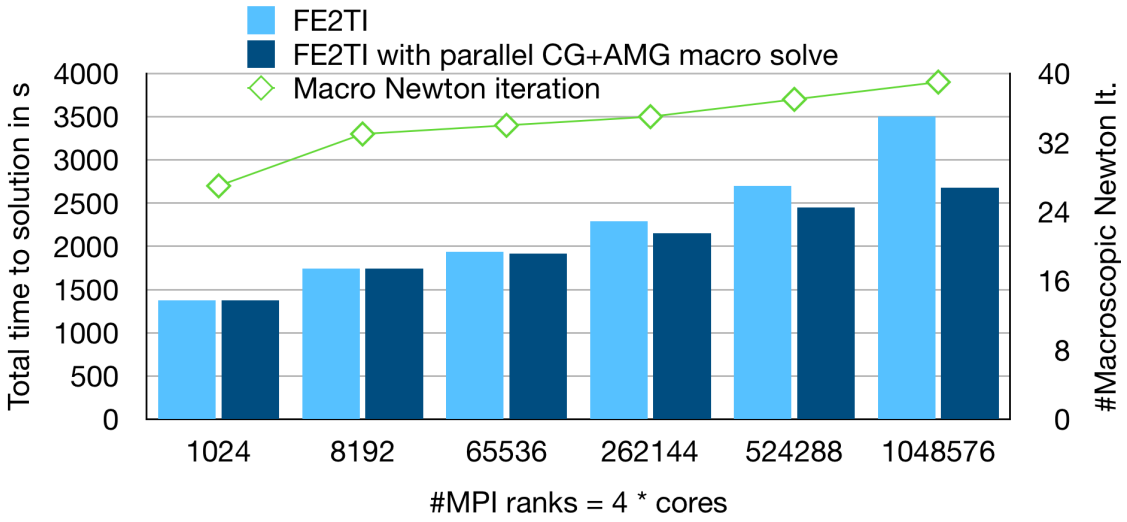

Fig. 15 Weak parallel scalability on JUQUEEN. Total time to solution of our FE^2 implementation FE2TI for 13 load steps using CG with an AMG preconditioner or, respectively, a direct solver on the macroscale. The largest problem has 3.3 billion degrees of freedom on the microscale and 7.8 thousand degrees of freedom on the macroscale.

References

1. Assyr Abdulle, Weinan E, Björn Engquist, and Eric Vanden-Eijnden. The heterogeneous multiscale method. *Acta Numer.*, 21:1–87, 2012.

2. Patrick R. Amestoy, Iain S. Duff, Jean-Yves L'Excellent, and Jacko Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41, 2001.

3. Allison H. Baker, Robert D. Falgout, Tzanio V. Kolev, and Ulrike Meier Yang. Scaling hypre's multigrid solvers to 100,000 cores. In Michael W. Berry, Kyle A. Gallivan, Efstratios Gallopoulos, Ananth Grama, Bernard Philippe, Yousef Saad, and Faisal Saied, editors, *High-Performance Scientific Computing: Algorithms and Applications*, pages 261–279. Springer London, London, 2012.

4. Allison H. Baker, Axel Klawonn, Tzanio Kolev, Martin Lanser, Oliver Rheinbach, and Ulrike Meier Yang. Scalability of classical algebraic multigrid for elastic-

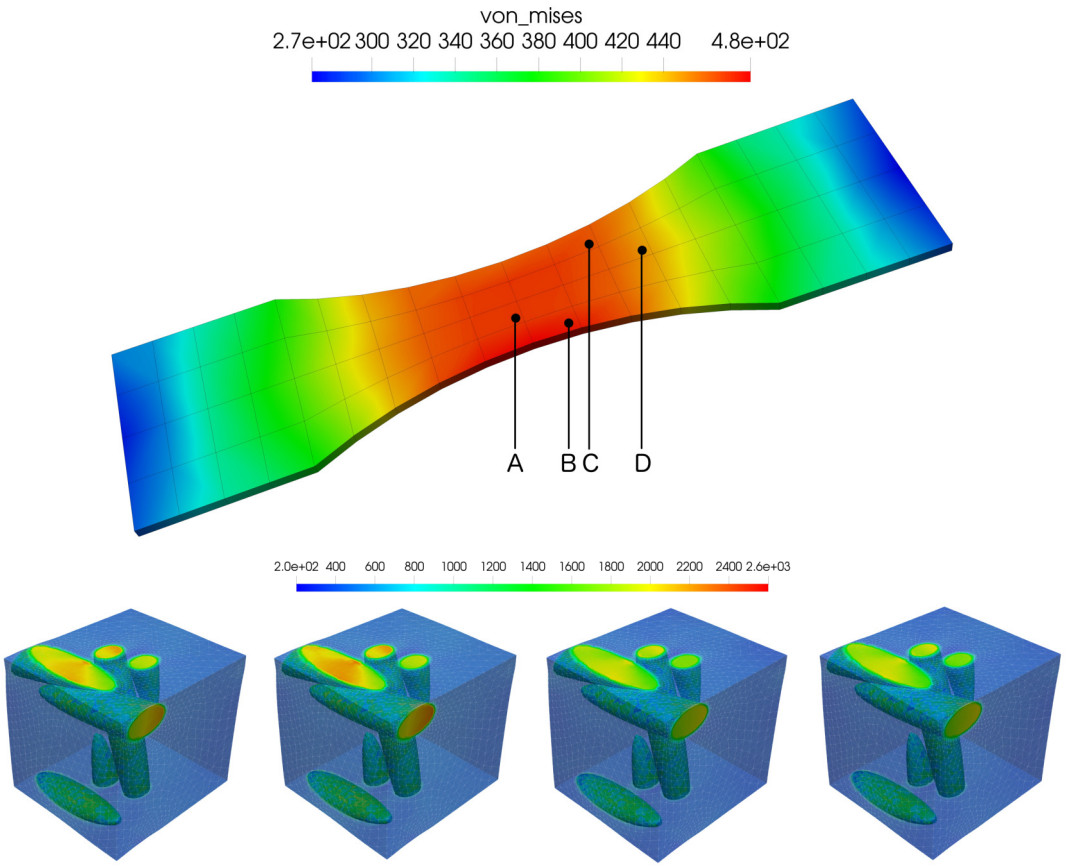

Fig. 16 Von Mises stresses after 171 load steps and a total deformation of approximately 4.3% of the macroscopic reference configuration. Four exemplary SSRVEs A,B,C, and D (from left to right) are depicted.

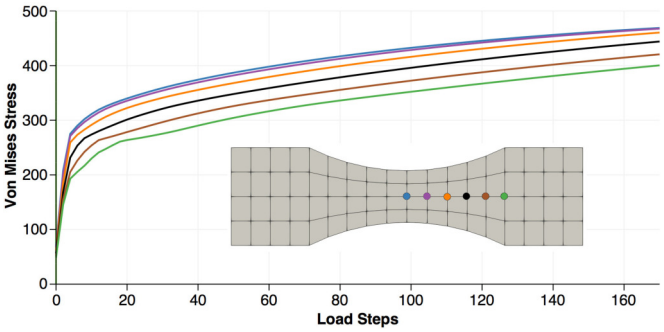

Fig. 17 Von Mises stress in six vertices of the macroscopic problem over pseudo time.

ity to half a million parallel tasks. In Hans-Joachim Bungartz, Philipp Neumann, and E. Wolfgang Nagel, editors, *Software for Exascale Computing - SPPEXA 2013-2015*, pages 113–140, 2016. vol. 113 of Springer Lect. Notes Comput. Sci. and Engrg. Also TUBAF Preprint 2015-14 at `http://tu-freiberg.de/fakult1/forschung/preprints`.

5. Satish Balay, Shrirang Abhyankar, Mark F. Adams, Jed Brown, Peter Brune, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Karl Rupp, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.5, Argonne National Laboratory, 2014.
6. D. Balzani, D. Brands, and J. Schröder. Construction of statistically similar representative volume elements. In J. Schröder and K. Hackl, editors, *Plasticity and beyond: Microstructures, Crystal-Plasticity and Phase Transitions*. CISM Lecture Notes No. 550, 2013.
7. D. Balzani, L. Scheunemann, D. Brands, and J. Schröder. Construction of two- and three-dimensional statistically similar RVEs for coupled micro-macro simulations. *Computational Mechanics*, 54:1269–1284, 2014.
8. Daniel Balzani, Ashutosh Gandhi, Axel Klawonn, Martin Lanser, Oliver Rheinbach, and Jörg Schröder. *One-Way and Fully-Coupled FE^2 Methods for Heterogeneous Elasticity and Plasticity Problems: Parallel Scalability and an Application to Thermo-Elastoplasticity of Dual-Phase Steels*, pages 91–112. Springer International Publishing, Cham, 2016. Also TUBAF Preprint: 2015-13, `http://tu-freiberg.de/fakult1/forschung/preprints`.
9. J.E. Bishop, J.M. Emery, R.V. Field, C.R. Weinberger, and D.J. Littlewood. Direct numerical simulations in solid mechanics for understanding the macroscale effects of microscale material variability. *Comput. Meth. Appl. Mech. Engrg.*, 287:262–289, 2015.
10. Joseph E. Bishop, John M. Emery, Corbett C. Battaile, David J. Littlewood, and Andrew J. Baines. Direct numerical simulations in solid mechanics for quantifying the macroscale effects of microstructure and material model-form error. *JOM*, 68(5):1427–1445, May 2016.

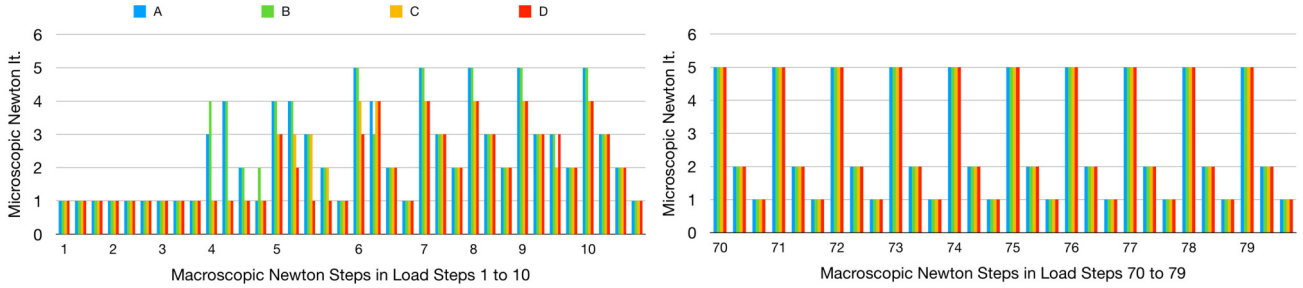

Fig. 18 Newton iterations on the microscale for the four RVEs A, B, C, and D (see Fig. 13 for position of RVEs); each block of four columns represents a single macroscopic Newton step in a certain macroscopic load step (marked on the x-axis). **Left:** Load steps 1 to 10. **Right:** Load steps 70 to 79.

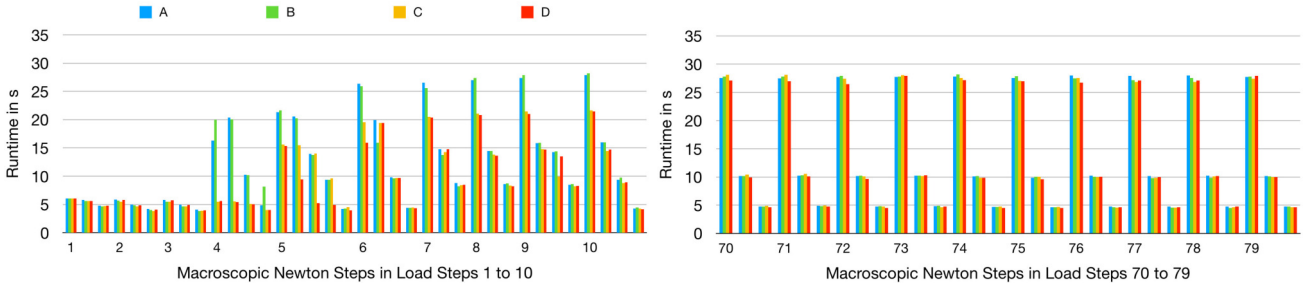

Fig. 19 Runtime of Nonlinear-FETI-DP-1 on the microscale for the four RVEs A, B, C, and D (see Fig. 13 for position of RVEs); each block of four columns represents a single macroscopic Newton step in a certain macroscopic load step (marked on the x-axis). **Left:** Load steps 1 to 10. **Right:** Load steps 70 to 79.

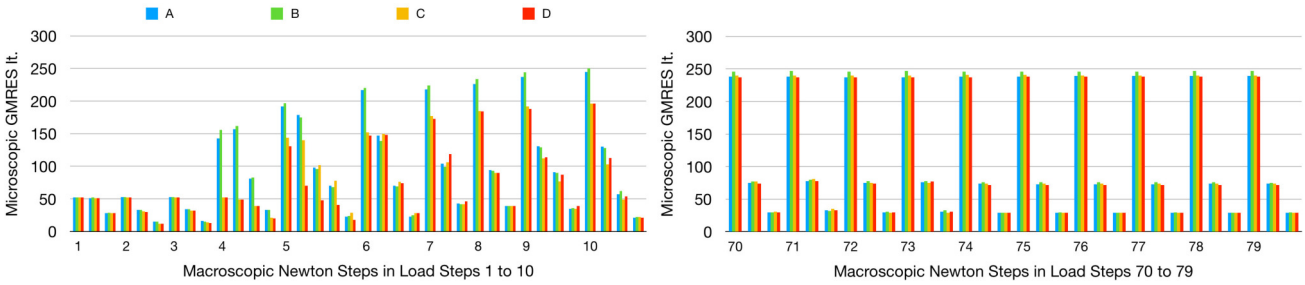

Fig. 20 GMRES iterations per Nonlinear-FETI-DP-1 solve on the microscale for the four RVEs A, B, C, and D (see Fig. 13 for position of RVEs); each block of four columns represents a single macroscopic Newton step in a certain macroscopic load step (marked on the x-axis). **Left:** Load steps 1 to 10. **Right:** Load steps 70 to 79.

11. Felipe Bordeu, Pierre-Alain Boucard, and Pierre Gosselet. Balancing domain decomposition with nonlinear relocalization: Parallel implementation for laminates. In P. Ivnyi B.H.V. Topping, editor, *Proc. 1st Int. Conf. on Parallel, Distributed and Grid Computing for Engineering*, Stirlingshire, UK, 2009. Civil-Comp Press.
12. D. Brands, D. Balzani, L. Scheunemann, J. Schröder, H. Richter, and D. Raabe. Computational modeling of dual-phase steels based on representative three-dimensional microstructures obtained from EBSD data. *Arch. Appl. Mech.*, 86(3):575–598, 2016.
13. Xiao-Chuan Cai and David E. Keyes. Nonlinearly preconditioned inexact Newton algorithms. *SIAM J. Sci. Comput.*, 24(1):183–200 (electronic), 2002.
14. Xiao-Chuan Cai, David E. Keyes, and Leszek Marcinkowski. Non-linear additive Schwarz preconditioners and application in computational fluid dynamics. *Internat. J. Numer. Methods Fluids*, 40(12):1463–1470, 2002.
15. Timothy A. Davis. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, 30(2):165–195, June 2004.
16. T.W.J. de Geus, J. Vondejc, J. Zeman, R.H.J. Peerlings, and M.G.D. Geers. Finite strain FFT-based non-linear solvers made simple. *Computer Methods in Applied Mechanics and Engineering*, 318:412 – 430, 2017.
17. Weinan E and Bjorn Engquist. The heterogeneous multiscale methods. *Commun. Math. Sci.*, 1(1):87–132, 2003.
18. Bernhard Eidel and Andreas Fischer. The heterogeneous multiscale finite element method for the homogenization of linear elastic solids and a comparison with the fe2 method. *Computer Methods in Applied Mechanics and Engineering*, 329(Supplement C):332 – 368, 2018.
19. P. Eisenlohr, M. Diehl, R.A. Lebensohn, and F. Roters. A spectral method solution to crystal elasto-viscoplasticity at finite strains. *International Journal of Plasticity*, 46:37–53, 2013.

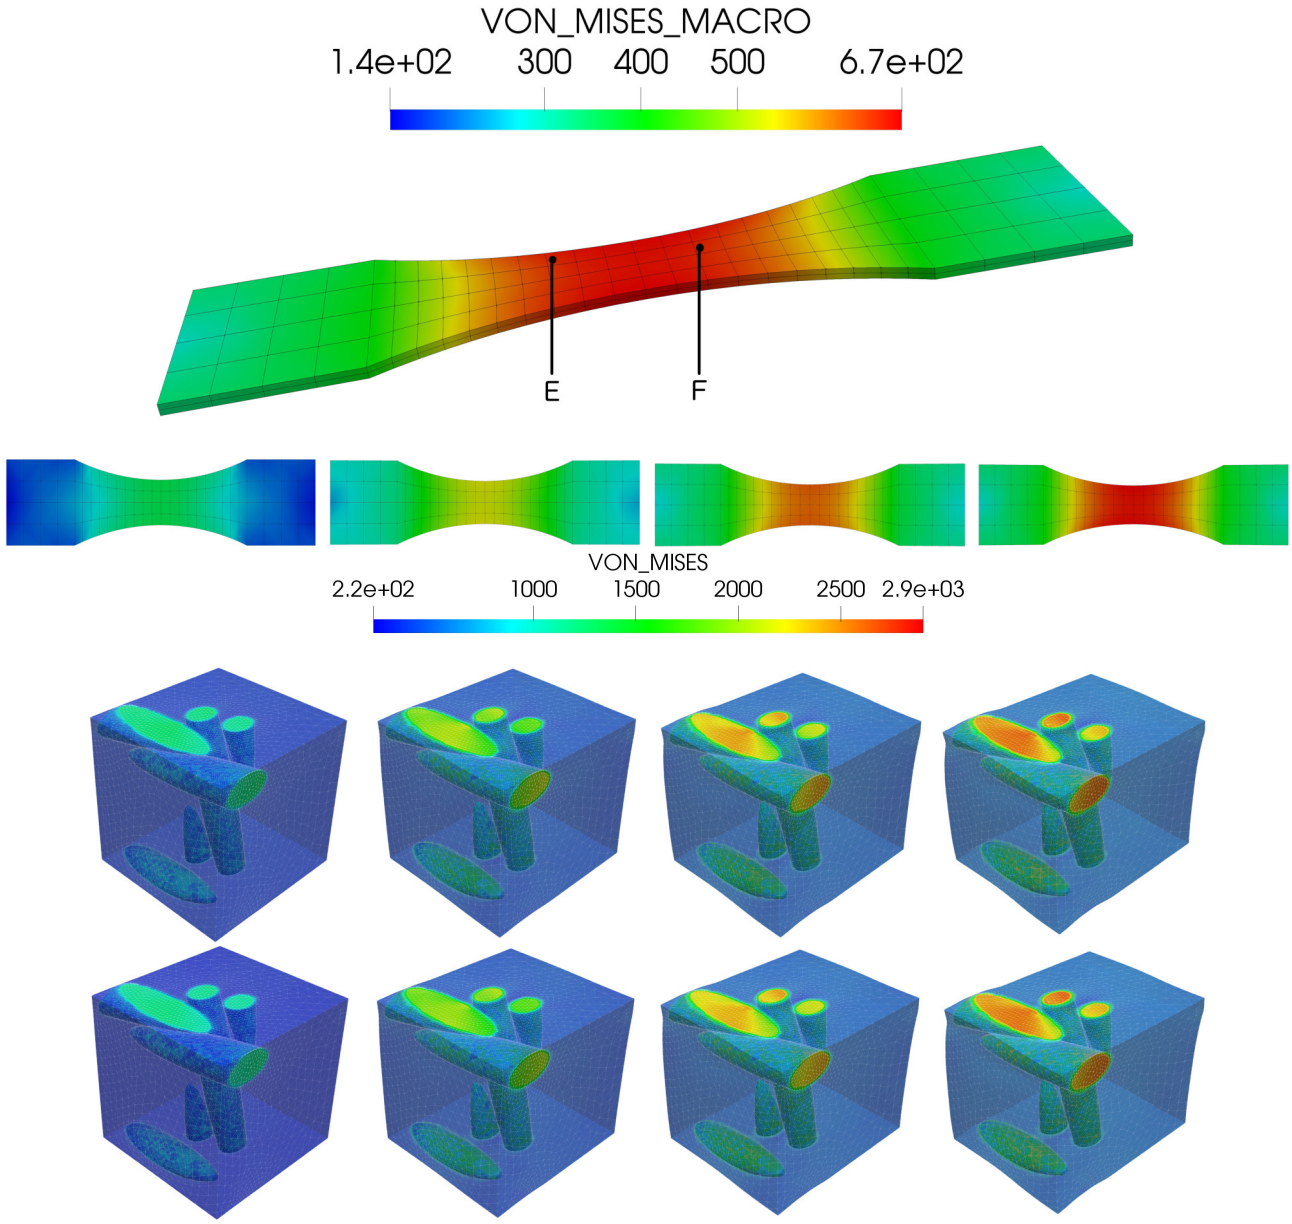

Fig. 21 Top: Macroscopic Von Mises stresses after 319 load steps and a total deformation of approximately 7.7% of the macroscopic reference configuration. **Second row (from left to right):** Macroscopic Von Mises stresses after 20, 120, 220, and 319 load steps. **Last two rows (from left to right):** Von Mises stresses in RVEs E and F after 20, 120, 220, and 319 load steps.

20. Robert D. Falgout, Jim E. Jones, and Ulrike Meier Yang. The design and implementation of hypre, a library of parallel high performance preconditioners. *Chapter in Numerical solution of Partial Differential Equations on Parallel Computers, A.M. Bruaset, P. Bjorstad, and A. Tveito, editors. Springer-Verlag.*, Dec 2005. Also available as LLNL Technical Report UCRL-JRNL-205459, 2004.

21. Charbel Farhat, Michel Lesoinne, Patrick LeTallec, Kendall Pierson, and Daniel Rixen. FETI-DP: A dual-primal unified FETI method - part i: A faster alternative to the two-level FETI method. *Internat. J. Numer. Methods Engrg.*, 50:1523–1544, 2001.

22. Charbel Farhat, Michel Lesoinne, and Kendall Pierson. A scalable dual-primal domain decomposition method. *Numer. Lin. Alg. Appl.*, 7:687–714, 2000.

23. Frédéric Feyel. Multiscale FE^2 elastoviscoplastic analysis of composite structures. *Computational Materials Science*, 16(14):344 – 354, 1999.

24. Frederic Feyel. A multilevel finite element method (FE^2) to describe the response of highly non-linear structures using generalized continua. *Computer Methods in Applied Mechanics and Engineering*, 192(28):3233 – 3244, 2003. Multiscale Computational Mechanics for Materials and Structures.

25. Frederic Feyel and Jean-Louis Chaboche. FE^2 multiscale approach for modelling the elastoviscoplastic behaviour

of long fibre sic/ti composite materials. *Computer Methods in Applied Mechanics and Engineering*, 183(3):309 – 330, 2000.
26. Christian Groß. *A unifying theory for nonlinear additively and multiplicatively preconditioned globalization strategies: Convergence Results and Examples From the field of Nonlinear Elastostatics and Elastodynamics*. PhD thesis, 2009. Deutsche Nationalbibliothek `https://www.deutsche-digitale-bibliothek.de/item/PCLVYPVW5OCPUOTIKRKTMSHMFSNWEFPL`.
27. Christian Groß and Rolf Krause. A generalized recursive trust-region approach - nonlinear multiplicatively preconditioned trust-region methods and applications. Technical Report 2010-09, Institute of Computational Science, Universita della Svizzera italiana, 03 2010.
28. Christian Groß and Rolf Krause. On the globalization of ASPIN employing trust-region control strategies - convergence analysis and numerical examples. Technical Report 2011-03, Inst. Comp. Sci., Universita della Svizzera italiana, 01 2011.
29. Van E. Henson and Ulrike Meier Yang. BoomerAMG: A parallel algebraic multigrid solver and preconditioner. *Appl. Numer. Math.*, 41:155–177, 2002.
30. Feng-Nan Hwang and Xiao-Chuan Cai. Improving robustness and parallel scalability of Newton method through nonlinear preconditioning. In *Domain decomposition methods in science and engineering*, volume 40 of *Lect. Notes Comput. Sci. Eng.*, pages 201–208. Springer, Berlin, 2005.
31. Feng-Nan Hwang and Xiao-Chuan Cai. A class of parallel two-level nonlinear Schwarz preconditioned inexact Newton algorithms. *Comput. Methods Appl. Mech. Engrg.*, 196(8):1603–1611, 2007.
32. Matthias Kabel, Dennis Merkert, and Matti Schneider. Use of composite voxels in FFT-based homogenization. *Comput. Meth. Appl. Mech. Engrg.*, 294:168 – 188, 2015.
33. A. Klawonn, M. Lanser, and O. Rheinbach. Nonlinear FETI-DP and BDDC methods. *SIAM J. Sci. Comput.*, 36(2):A737–A765, 2014.
34. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. FE2TI (ex_nl/fe^2) EXASTEEL - Bridging scales for multiphase steels, 2015.
35. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. Toward extremely scalable nonlinear domain decomposition methods for elliptic partial differential equations. *SIAM Journal on Scientific Computing*, 37(6):C667–C696, 2015.
36. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. FE^{2}TI: Computational scale bridging for dual-phase steels. In *IOS Series Advances in Parallel Computing, Volume 27, Parallel Computing: On the Road to Exascale; Proceedings of ParCo2015*, pages 797–806, 2016. Also TUBAF Preprint: 2015-12, `http://tu-freiberg.de/fakult1/forschung/preprints`.
37. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. A highly scalable implementation of inexact nonlinear feti-dp without sparse direct solvers. In Bülent Karasözen, Murat Manguoğlu, Münevver Tezer-Sezgin, Serdar Göktepe, and Ömür Uğur, editors, *Numerical Mathematics and Advanced Applications ENUMATH 2015*, pages 255–264, Cham, 2016. Springer International Publishing.
38. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. Nonlinear BDDC methods with inexact solvers. 2017. Submitted.
39. Axel Klawonn, Martin Lanser, Oliver Rheinbach, Holger Stengel, and Gerhard Wellein. *Hybrid MPI/OpenMP Parallelization in FETI-DP Methods*, pages 67–84. Springer International Publishing, Cham, 2015.
40. Axel Klawonn, Martin Lanser, Oliver Rheinbach, and Matthias Uran. Nonlinear FETI-DP and BDDC methods: A unified framework and parallel results. *SIAM Journal on Scientific Computing*, 39(6):C417–C451, 2017.
41. Axel Klawonn and Oliver Rheinbach. Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM Z. Angew. Math. Mech.*, 90(1):5–32, 2010.
42. S.O. Klinkel. *Theorie und Numerik eines Volumen-Schalen-Elementes bei finiten elastischen und plastischen Verzerrungen*. Berichte des Instituts für Baustatik, Karlsruher Institut für Technologie. Inst. für Baustatik, 2000.
43. Dana A. Knoll and David E. Keyes. Jacobian-free Newton-Krylov methods: a survey of approaches and applications. *J. Comput. Phys.*, 193(2):357–397, 2004.
44. Julian Kochmann, Stephan Wulfinghoff, Stefanie Reese, Jaber Rezaei Mianroodi, and Bob Svendsen. Two-scale FE-FFT- and phase-field-based computational modeling of bulk microstructural evolution and macroscopic material behavior. *Computer Methods in Applied Mechanics and Engineering*, 305:89 – 110, 2016.
45. V. Kouznetsova, W.A.M. Brekelmans, and F.P.T. Baaijens. An approach to micro-macro modeling of heterogeneous materials. 27:37–48, 2001.
46. Martin Lanser. *Nonlinear FETI-DP and BDDC Methods*. PhD thesis, Universität zu Köln, 2015.
47. I. A. Rodrigues Lopes, F. M. Andrade Pires, and F. J. P. Reis. A mixed parallel strategy for the solution of coupled multi-scale problems at finite strains. *Computational Mechanics*, Sep 2017.
48. C. Miehe, J. Schröder, and J. Schotte. Computational homogenization analysis in finite plasticity. Simulation of texture development in polycrystalline materials. *Computer Methods in Applied Mechanics and Engineering*, 171:387–418, 1999.
49. Matthew Mosby and Karel Matou. Hierarchically parallel coupled finite strain multiscale solver for modeling heterogeneous layers. *International Journal for Numerical Methods in Engineering*, 102(3-4):748–765.
50. Matthew Mosby and Karel Matou. Computational homogenization at extreme scales. *Extreme Mechanics Letters*, 6:68 – 74, 2016.
51. Herv Moulinec and Pierre Suquet. Fast numerical method for computing the linear and nonlinear properties of composites. 318, 01 1994.
52. Julien Pebrel, Christian Rey, and Pierre Gosselet. A nonlinear dual-domain decomposition method: Application to structural problems with damage. *Inter. J. Multiscal Comp. Eng.*, 6(3):251–262, 2008.
53. Ulrich Rüde, Karen Willcox, Lois Curfman McInnes, Hans De Sterck, George Biros, Hans-Joachim Bungartz, James Corones, Evin Cramer, James Crowley, Omar Ghattas, Max Gunzburger, Michael Hanke, Robert J. Harrison, Michael A. Heroux, Jan Hesthaven, Peter K. Jimack, Chris Johnson, Kirk E. Jordan, David E. Keyes, Rolf H. Krause, Vipin Kumar, Stefan Mayer, Juan Meza, Knut Martin Mørken, J. Tinsley Oden, Linda R. Petzold, Padma Raghavan, Suzanne M. Shontz, Anne E. Trefethen, Peter R. Turner, Vladimir V. Voevodin, Barbara I. Wohlmuth, and Carol S. Woodward. Research and education in computational science and engineering. *CoRR*, 2016. Submitted to SIAM Rev.
54. O. Schenk and K. Gärtner. PARDISO. In David A. Padua, editor, *Encyclopedia of Parallel Computing*, pages 1458–1464. Springer, 2011.

55. L. Scheunemann, D. Balzani, D. Brands, and J. Schröder. Design of 3D statistically similar representative volume elements based on Minkowski functionals. *Mechanics of Materials*, 90(Supplement C):185 – 201, 2015. Proceedings of the IUTAM Symposium on Micromechanics of Defects in Solids.
56. Lisa Scheunemann, Daniel Balzani, Dominik Brands, and Jörg Schröder. Construction of statistically similar RVEs. In *Analysis and computation of microstructure in finite plasticity*, volume 78 of *Lect. Notes Appl. Comput. Mech.*, pages 219–256. Springer, Cham, 2015.
57. Matti Schneider, Dennis Merkert, and Matthias Kabel. FFT-based homogenization for microstructures discretized by linear hexahedral elements. *International Journal for Numerical Methods in Engineering*, 109(10):1461–1489, 2017.
58. Matti Schneider, Felix Ospald, and Matthias Kabel. Computational homogenization of elasticity on a staggered grid. *Int. J. Numer. Meth. Engrg.*, 105(9):693–720, 2016.
59. J. Schröder. *Homogenisierungsmethoden der nichtlinearen Kontinuumsmechanik unter Beachtung von Stabilitätsproblemen.* PhD thesis, Bericht aus der Forschungsreihe des Institut für Mechanik (Bauwesen), Lehrstuhl I, Universität Stuttgart, 2000. Habilitationsschrift.
60. J. Schröder. A numerical two-scale homogenization scheme: the FE^2-method. In *Plasticity and Beyond - Microstructures, Crystal-Plasticity and Phase Transitions (CISM Lecture Notes 550, Eds. J. Schröder, K. Hackl)*. Springer, 2013.
61. R.J.M. Smit, W.A.M. Brekelmans, and H.E.H. Meijer. Prediction of the mechanical behavior of nonlinear heterogeneous systems by multi-level finite element modeling. *Computer Methods in Applied Mechanics and Engineering*, 155:181–192, 1998.
62. J. Spahn, H. Andr, M. Kabel, and R. Mller. A multiscale approach for modeling progressive damage of composite materials using fast fourier transforms. *Computer Methods in Applied Mechanics and Engineering*, 268:871 – 883, 2014.
63. Michael Stephan and Jutta Docter. JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre. *Journal of large-scale research facilities*, 1:A1, 2015.
64. Andrea Toselli and Olof Widlund. *Domain decomposition methods—algorithms and theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 2005.
65. Markus Wittmann, Georg Hager, Radim Janalik, Martin Lanser, Axel Klawonn, Oliver Rheinbach, Olaf Schenk, and Gerhard Wellein. Multicore performance engineering of sparse triangular solves using a modified roofline model. 2018. In preparation.
66. Stefano Zampini. PCBDDC: a class of robust dual-primal methods in PETSc. *SIAM J. Sci. Comput.*, 38(5):S282–S306, 2016.