

EXASTEEL: Towards a Virtual Laboratory for the Multiscale Simulation of Dual-Phase Steel Using High-Performance Computing



Axel Klawonn, Martin Lanser, Matthias Uran, Oliver Rheinbach, Stephan Köhler, Jörg Schröder, Lisa Scheunemann, Dominik Brands, Daniel Balzani, Ashutosh Gandhi, Gerhard Wellein, Markus Wittmann, Olaf Schenk, and Radim Janalík

Abstract We present a numerical two-scale simulation approach of the Nakajima test for dual-phase steel using the software package FE2TI, a highly scalable implementation of the well known homogenization method FE². We consider the incorporation of contact constraints using the penalty method as well as the sample sheet geometries and adequate boundary conditions. Additional software features such as a simple load step strategy and prediction of an initial value by linear extrapolation are introduced.

The macroscopic material behavior of dual-phase steel strongly depends on its microstructure and has to be incorporated for an accurate solution. For a reasonable computational effort, the concept of statistically similar representative volume

A. Klawonn (✉) · M. Lanser · M. Uran

University of Cologne, Köln, Germany

e-mail: axel.klawonn@uni-koeln.de; martin.lanser@uni-koeln.de; m.uran@uni-koeln.de

O. Rheinbach · S. Köhler

Technische Universität Bergakademie Freiberg, Freiberg, Germany

e-mail: oliver.rheinbach@math.tu-freiberg.de; Stephan.Koehler@math.tu-freiberg.de

J. Schröder · L. Scheunemann · D. Brands

Universität Duisburg-Essen, Duisburg, Germany

e-mail: j.schroeder@uni-due.de; lisa.scheunemann@uni-due.de; dominik.brands@uni-due.de

D. Balzani · A. Gandhi

Ruhr-Universität Bochum, Bochum, Germany

e-mail: daniel.balzani@rub.de; ashutosh.gandhi@rub.de

G. Wellein · M. Wittmann

Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

e-mail: gerhard.wellein@rre.uni-erlangen.de

O. Schenk · R. Janalík

Universita della Svizzera italiana, Lugano, Switzerland

e-mail: olaf.schenk@usi.ch; radim.janalik@usi.ch

© The Author(s) 2020

H.-J. Bungartz et al. (eds.), *Software for Exascale Computing - SPPEXA*

2016–2019, Lecture Notes in Computational Science and Engineering 136,

https://doi.org/10.1007/978-3-030-47956-5_13

elements (SSRVEs) is presented. Furthermore, the highly scalable nonlinear domain decomposition methods NL-FETI-DP and nonlinear BDDC are introduced and weak scaling results are shown. These methods can be used, e.g., for the solution of the microscopic problems. Additionally, some remarks on sparse direct solvers are given, especially to PARDISO. Finally, we come up with a computationally derived Forming Limit Curve (FLC).

1 Introduction

In the EXASTEEL project, we are solving challenging **nonlinear multiscale problems from computational material science** showing parallel scalability **beyond a million parallel processes**. Our software package FE2TI solves large scale **contact problems** in sheet metal forming of microheterogeneous materials and **scales to some of the largest supercomputers available today**. Although an exascale computer is not yet available, FE2TI is **exascale ready**: For our current production simulations, we have not pushed the combined parallelism of the FE^2 multiscale computational homogenization method and of our nonlinear solvers to the limit. Both, i.e., the FE^2 method by itself, as well as our nonlinear solvers are scalable to the largest supercomputers currently in production in the leading international computing facilities.¹

In particular, as a problem, we consider the finite element simulation of sheet metal forming processes of **dual-phase (DP) steels**, whose macroscopic material behavior strongly **depends on its microscopic material properties**. A brute force discretization with respect to the microscopic structure would lead to extremely large systems of equations, which are not feasible, even on the upcoming exascale supercomputers. To give an example, a reasonable finite element discretization down to the microscopic scale would require 10^3 – 10^9 finite elements for a three dimensional cube with a volume of $1 \mu\text{m}^3$. Extrapolating this to a sheet with an area of 1 m^2 and a thickness of 1 mm would lead to 10^{18} – 10^{24} finite elements. A brute force simulation would also require knowledge of the complete microstructure of the steel sheet which is not available. Therefore, an efficient multiscale or homogenization approach is indispensable to save 3 to 6 orders of magnitude of

¹In 2011, the overall scientific goal of the German DFG priority program SPP 1648 “Software for Exascale Computing” (SPPEXA) was stated as “*to master the various challenges related to [...] [the] paradigm shift from sequential or just moderately parallel to massively parallel processing*” and thereby to “*advance the frontier of parallel computing*” [4]. From the beginning, SPPEXA aimed at a true co-design, i.e., closely connecting “*computer science with the needs of Computational Science and Engineering (CSE) and HPC*” [4]. The project EXASTEEL addresses three of the main SPPEXA research areas, namely computational algorithms, application software, and programming, i.e., we have, e.g., introduced new nonlinear solver algorithms, implemented our multiscale application software FE2TI, and applied hybrid programming and performance engineering to our codes. This work was only possible in close collaboration of mathematics, computer science, and engineering.

unknowns. Our choice of a computational homogenization approach is the **FE² method** which is well established in engineering; see Sect. 3 for a short introduction and further references. In the FE² method, the microscopic and macroscopic level are discretized independently of each other. No material laws are needed for the macroscopic level, all information needed is obtained from microscopic computations based on material laws and data on the microscopic level. Let us note that the microscopic problems can be solved in parallel once the solution of the macroscopic problem is available as input. The solution of the macroscopic problem, however, requires the information of all microscopic solutions. Thus, the FE² method is not trivially parallelizable but requires a sequential solution of the microscopic and the macroscopic problems; this is similar to the coarse level of a hybrid two-level domain decomposition method with multiplicative coarse level and additive subdomain solvers. The nonlinear problems on both levels, the macroscopic and the microscopic one, can be solved (after linearization) using highly parallel scalable and robust implicit solvers such as parallel algebraic multigrid methods (AMG) or parallel domain decomposition preconditioners such as FETI-DP (Finite Element Tearing and Interconnecting-Dual-Primal) [27, 28, 47–50] or BDDC (Balancing Domain Decomposition by Constraints) [20, 24, 71–73] methods. These preconditioners are usually applied as part of a Newton-Krylov approach, where the tangent problem in each Newton iteration is solved using preconditioned Krylov iteration methods. A more recent approach to nonlinear implicit problems, developed extensively within EXASTEEL, is given by **non-linear parallel domain decomposition methods**, which are applied directly to the nonlinear problem, i.e., before linearization. In such methods, the **nonlinear problem is first decomposed into concurrent nonlinear problems**, which are **then solved by (decoupled) Newton’s methods in parallel**. In this project, **nonlinear FETI-DP and nonlinear BDDC domain decomposition methods** (see also Sect. 6) have been introduced and have successfully **scaled to the largest supercomputers available**—independently of the multiscale context given by the FE² homogenization methods, which adds an additional level of parallelism. It was found that the nonlinear domain decomposition methods can reduce communication and synchronization and thus time to solution. They can, however, also reduce the energy to solution; see Sect. 6.1.1 and [63]. These methods can be applied within our highly scalable software package FE2TI but can also be used for all problems where implicit nonlinear solvers are needed on extreme scale computers. For scaling results of the FE² method to more than one million MPI ranks, see Fig. 3 in Sect. 3.2 and [64]. Note that these scaling results can be obtained only using additional parallelization on the macroscopic level. Note also that our new nonlinear implicit solvers based on nonlinear FETI-DP have scaled to the complete Mira supercomputer, i.e., 7,58,000 MPI ranks (see Fig. 15 and [57]); on the JUQUEEN supercomputer [44] (see [60]) our solver based on nonlinear BDDC has scaled to 2,62,000 MPI ranks for a 3D structural mechanics problem as well as 5,24,000 MPI ranks for a 2D problem. In the present article, the software package is used to derive a virtual forming limit diagram (FLD) by simulating the Nakajima test,

a standard procedure for the derivation of FLDs. An FLD contains a Cartesian coordinate system with major and minor strain values and a regression function of these values, which is called forming limit curve (FLC). An FLC gives the extent to which the material can be deformed by stretching, drawing or any combination of stretching and drawing without failing [77, p. v].

The software and algorithms developed here have participated in scaling workshops at the Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany (see the reports [53, 58]), as well as at the Argonne Leadership Computing Facility (ALCF), Argonne National Laboratory, USA. They have scaled on the following world-leading supercomputers in Europe, the United States, and Asia (TOP500 rank given for the time of use):

- JUQUEEN at the Jülich Supercomputing Centre, Germany; European Tier 0; TOP500 rank 9 in the year 2015 (458,752 cores; 5.8 petaflops); FE2TI and FETI-DP have scaled to the complete machine [53, 56–58, 64]; since 2015 FE2TI is a member of the High-Q Club of the highest scaling codes on JUQUEEN [53].
- JUWELS at Jülich Supercomputing Centre, Germany; European Tier 0; TOP500 rank 23 in the year 2018 (114,480 cores; 9.8 petaflops); main source of compute time for the computation of an FLD; see Sect. 5
- Mira at Argonne Leadership Computing Facility (ALCF), Argonne National Laboratory (ANL), USA; TOP500 rank 5 in the year 2015 (786,432 cores; 10.1 petaflops); FE2TI and nonlinear FETI-DP have scaled to the complete machine [54, 57]
- Theta at ALCF, USA; TOP500 rank 18 in the year 2017 (280,320 cores; 9.6 petaflops); is ANL's bridge to the upcoming first US exascale machine AURORA (or AURORA21) scheduled for 2021; BDDC domain decomposition solver has scaled to 193,600 cores [60] and recently to 262,144 cores
- Oakforest-PACS at Joint Center for Advanced High Performance Computing, Japan; TOP500 rank 6 in the year 2016 (556,104 cores; 24.9 petaflops); first deep drawing computations using FE2TI

The remainder of the article is organized as follows. In Sect. 2, we introduce the experimental test setup of the Nakajima test and the evaluation strategy of major and minor strain values described in DIN EN ISO 12004-2:2008 [77]. In Sect. 3, we briefly describe the ingredients of our highly scalable software package FE2TI, including the computational homogenization method FE^2 and the contact formulation which is integrated into the software package FE2TI since the sheet metal deformation in the Nakajima test is caused by contact. We also present some strategies to reduce computing time. In Sect. 4, we describe the numerical realization of the Nakajima test. Then, in Sect. 5, we present numerical results of several *in silico* Nakajima simulations with different specimens resulting in a virtual FLC. In Sect. 6, we give an overview over our highly scalable linear and nonlinear implicit solvers, including nonlinear FETI-DP and nonlinear BDDC. These methods

can be used to solve all nonlinear problems occurring in FE2TI, as shown, e.g., in [64]. In Sect. 7, performance engineering aspects regarding the sparse direct solver package PARDISO [81] are discussed. The PARDISO sparse direct solver is a central building block in our implicit solvers and in FE2TI. In Sect. 8, we introduce different improvements on the microscopic material model to even better match some experimental results.

2 Nakajima Test

Stricter CO₂ emission regulations in combination with higher passenger safety norms in the automotive industry requires steel grades with higher toughness and less weight. The class of DP steels belongs to the advanced high-strength steels and combines strength and ductility. Its favorable macroscopic properties result from the microscopic heterogeneous structure; see the beginning of Sect. 8 for further remarks.

To demonstrate the macroscopic material behavior of a specific steel grade, different material parameters and forming behaviors have to be proven. A prominent member of material characterization is the forming limit diagram (FLD). It contains major and minor strain values at failure initiation in a Cartesian coordinate system and represents the forming limits of a steel grade for one specific material thickness. In this context, material failure is already associated with the beginning of local necking in the direction of thickness and not only with crack formation [77, p. v]. The major and minor strain values vary from uniaxial to equi-biaxial tension.

The Nakajima test is a standard procedure in material testing. In the Nakajima test, a specimen is clamped between a blank holder and a die and a hemispherical punch is driven into the specimen until a crack can be observed; see Fig. 1 (left). Friction between the forming tool and the specimen has to be avoided as much as possible. Therefore, different lubrication systems can be applied; see [77, Ch. 4.3.3.3]. To get different pairs of major and minor strains, one has to use at least five different shapes of sample geometries and for each shape, one has to carry out three different valid tests [77]. The recommended shapes of the sample sheet geometries are described in [77, Ch. 4.1.2], see also Sect. 4.1 and Fig. 1 (right) for an example of a permissible sample sheet. In experiments, the surface of a specimen is equipped with a regular grid or a stochastic pattern and is recorded by one or more cameras during the deformation process.

There are at least two different strategies to get the pair of major and minor strains for the FLC, namely the cross section method [77] and a method based on thinning rates proposed by W. Volk and P. Hora [97]. Since the FLC gives information about material deformation without failing, we are interested in major and minor strains just before localized necking occurs.




Fig. 1 **Left:** Cross section of the initial test setup of the Nakajima test. **Right:** Dimensions of a specimen used for the simulation of the Nakajima test with a shaft length of 25 mm and a parallel shaft width of 90 mm. The inner (red) circle represents the inner wall of the die and the outer (green) circle represents the beginning of the clamped part between die and blank holder. Material outside the outer (green) circle is only considered for a width of the parallel shaft of 90 mm or more (dark grey)

In the method based on thinning rates, the last recorded image before localized necking occurs is explicitly determined. This specific image is used to derive major and minor strains for the FLC.

The **cross section method** is standardized in DIN EN ISO 12004-2:2008 [77]. It uses knowledge about the position of the crack and evaluates major and minor strain values in the last recorded image before crack along cross sections perpendicular to the crack. Then, from these values, the state immediately before material failure is interpolated. Cross sections have a length of at least 20 mm at both sides of the crack. One cross section cuts exactly through the center of the crack and one or two cross sections are positioned above and below with a distance of about 2 mm. For each cross section, we want to compute a pair of major and minor strains $\bar{\epsilon}_1^{\text{FLC}}$ and $\bar{\epsilon}_2^{\text{FLC}}$, which represent the major and minor strains just before the beginning of plastic instability.² Therefore, we have to fit an inverse second-order polynomial using a least squares fit; see Figs. 2 and 8 (bottom). Instead of fitting inverse second-order polynomials to the values along the cross sections we fit second order polynomials to the inverse of the values. For the least squares fit we have to determine optimal fit windows for both sides of the crack separately; see Figs. 2 and 8 (bottom). For a detailed description of the procedure we refer to [77]. Let us note that $\bar{\epsilon}_1^{\text{FLC}}$ and $\bar{\epsilon}_2^{\text{FLC}}$ in general never exist during the deformation process. Hence, these numbers do not have a physical meaning [97].

²Note that here and in the following, all macroscopic variables and objects are denoted with an overline to distinguish them from microscopic variables and objects.




Fig. 2 Fitted inverse second-order polynomials to the major strain values along the first cross section just before material failure. See also the description of the cross section method in Sect. 2. Optimal fit windows are computed as described in [77, Ch. 5.2.3, 5.2.4]. **Left:** Specimen with a width of the parallel shaft of 70 mm. **Right:** Full circular specimen

3 FE2TI: A Highly Scalable Implementation of the FE^2 Algorithm

For the finite element simulation of the Nakajima test, we use our FE2TI software package [9, 52, 57, 64], which is a C/C++ implementation of the FE^2 computational homogenization approach [29–31, 33, 70, 75, 86, 87, 91]. It is based on PETSc [6] and MPI. The multiscale simulations based on FE2TI and using FETI-DP and BoomerAMG as solvers are a “CSE Success Story” in SIAM Review [80, p. 736].

3.1 FE^2 Algorithm

For DP steel, the overall macroscopic material behavior strongly depends on its microscopic properties. Assuming that the macroscopic length scale \bar{L} is much larger than the length scale l representing the microscopic heterogeneities, i.e., $\bar{L} \gg l$, the scale separation assumption is satisfied and a scale bridging or homogenization approach such as the FE^2 method can be applied.

The idea of the FE^2 approach is to discretize the micro- and macroscopic scale separately using finite elements. The macroscopic domain is discretized without any consideration of the microscopic material properties, i.e., the material is assumed to be homogeneous from a macroscopic point of view. Additionally, a microscopic boundary value problem is defined on a representative volume element (RVE) which is assumed to represent the microscopic heterogeneities sufficiently. One microscopic finite element problem is assigned to each macroscopic Gauß point and the phenomenological law on the macroscopic level is replaced by volumetric averages of stresses and associated consistent tangent moduli of the microscopic solution. Note that the boundary values of the microscopic level are induced through the macroscopic deformation gradient at the integration point the microscopic problem is attached to.

To derive an RVE representing a realistic microstructure, electron backscatter diffraction is used; see [14]. Note that for DP steel the martensitic inclusions in the ferrite are quite small and widely spread, which enforces a fine discretization to incorporate the heterogeneities sufficiently. To overcome this problem, we make use of so called statistically similar RVEs (SSRVEs) [8, 83], which are constructed in an optimization process with only inclusions of simple geometry such as ellipsoids, but describe the mechanical behavior in an approximate way. Note that the constructed ellipsoids are simpler than the realistic microstructure and hence, the SSRVE can be discretized with a coarser grid.

For further details such as the derivation of consistent tangent moduli we refer to the literature [33, 87] and to earlier works on computational homogenization in the EXASTEEL project [9, 52, 57, 64].

3.2 *FE2TI Software Package*

The FE2TI software package was developed within the EXASTEEL project and has been successfully used for the simulation of tension tests of DP steel [9, 52, 57, 64]. It belongs to the highest scaling codes on the European Tier-0-supercomputer JUQUEEN since 2015.³

For comparably small sizes of microscopic problems, we can solve the resulting tangent problems with a sparse direct solver such as PARDISO [81], UMFPACK [22], or MUMPS [2]. For larger sizes of microscopic problems, we have to use efficient parallel solvers which are also robust for heterogeneous problems. Such methods are Newton-Krylov methods with appropriate preconditioners such as domain decomposition or (algebraic) multigrid or nonlinear domain decomposition methods, possibly, combined with algebraic multigrid.

In our software package, Newton-Krylov-FETI-DP and the more recent highly scalable nonlinear FETI-DP methods, which were developed in this project (see [51, 54, 59] and Sect. 6.1), are integrated. Other nonlinear domain decomposition approaches are the related Nonlinear FETI-1 or Neumann-Neumann approaches [13, 78] or ASPIN [17, 18, 35–37, 40, 41, 66]. Furthermore, FE2TI can also use the highly scalable algebraic multigrid implementation BoomerAMG [5, 38] from the hypr package [25] for micro- as well as macroscopic problems. The scalability of BoomerAMG was recently improved for problems in elasticity, and scalability of BoomerAMG to half a million ranks was then achieved within the EXASTEEL project [5] in close collaboration with the authors of BoomerAMG.

For the contact simulations presented here, we consider problem sizes for which we can use the direct solver package PARDISO to solve the resulting tangent problems on the microscopic as well as on the macroscopic level. This limits the

³https://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/FE2TI/_node.html and also <https://juser.fz-juelich.de/record/188191>.

size of our problems but is suitable for mid-sized supercomputers. In our opinion, this is a relevant criterion for the applicability in industry. Using our parallel nonlinear solvers, the FE2TI package scales up to the largest machines even without making use of the full scaling potential of the solvers (see Fig. 3); and for the combination of large macroscopic problems with large RVEs an exascale computer will be necessary in the future. While Fig. 3 (left) represents a weak scaling study with large RVEs and comparably small macroscopic problems, in Fig. 3 (right) the macroscopic problems are larger. Therefore, in the latter case, a parallel macroscopic solve using GMRES with an AMG preconditioner is beneficial. The scalability in Fig. 3 (right) somewhat suffers from an increase in the numbers of Newton iterations. Let us remark that the setup in Fig. 3 (right) is the setup of a typical production run. The strong scaling potential of FE2TI is also presented in Fig. 4; see [9] for details. For more scalability results on different architectures also see [64].

Even if the macroscopic problem is solved with a direct solver, the assembly process is parallelized. For the incorporation of a material law on the microscopic level the software is equipped with an interface to FEAP, and we use an implementation of a J_2 elasto-plasticity model [65]. Material parameters are chosen as in Brands et al. [14, Fig. 10].

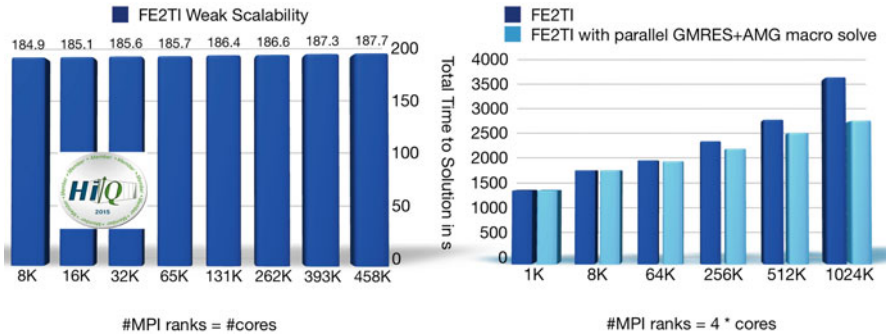


Fig. 3 Weak scalability of the FE2TI software on the JUQUEEN supercomputer [44]. **Left:** Time to solution of a single load step solving a three-dimensional heterogeneous hyperelastic model problem; uses Q1 finite elements (macro) and P2 finite elements (micro); 1.6M d.o.f. on each RVE; 512 FETI-DP subdomains for each RVE; the macroscopic problem size grows proportionally to the number of MPI ranks while the microscopic problem size is fixed; corresponding data in [57, Tab. 2]; High-Q club computation in 2015. **Right:** Total time to solution for 13 load steps solving 3D heterogeneous plasticity; uses Q1 finite elements (macro) and P2 finite elements (micro); 200K d.o.f. on each RVE; 64 FETI-DP subdomains for each RVE; the macroscopic problem is increased proportionally to the number of MPI ranks; for the larger problems using parallel AMG for the problem on the macroscale, instead of a sparse direct solver, is beneficial; see also [64, Fig. 15]




Fig. 4 Strong scalability of the FE2TI software for a nonlinear elasticity problem. Macroscopic problem with 32 finite elements; each RVE with 107K degrees of freedom is solved using 512 FETI-DP subdomains. Simulation of one macroscopic load step. **Left:** Total time to solution; **Right:** Speedup. Figures from [9]

3.3 Contact Kinematics and Incorporation of Contact Constraints for Frictionless Contact in FE2TI

For the simulation of the Nakajima test, we have to consider contact between the deformable specimen $\bar{\mathcal{B}}$ and different rigid tools $\bar{\mathcal{T}}_i$, $i = 1, 2, 3$, namely the hemispherical punch, blank holder, and die; see Fig. 1 (left). Therefore, we implemented a contact algorithm on the macroscopic scale in FE2TI. To simplify the notation, we consider an arbitrary rigid tool $\bar{\mathcal{T}}$ in the following.

A general convention in contact formulations is to consider one contact partner as the master body and one contact partner as the slave body [68, 99].

Only points of the contact surface of the slave body are allowed to penetrate into the master body. Following [68], one can choose the rigid surface as slave surface as well as a master surface, and in [99, Rem. 4.2] it is recommended to use the rigid surface as master surface; we have applied the latter in our simulations. Nevertheless, the contact contributions to the stiffness matrix and the right-hand-side are computed in the coordinate system of the deformable body.

In every iteration, we have to check for all finite element nodes $\bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_{\bar{\mathcal{B}}}$ of the contact surface of $\bar{\mathcal{B}}$ whether it penetrates into $\bar{\mathcal{T}}$ or not; see Fig. 5 for a simplified illustration. For each $\bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_{\bar{\mathcal{B}}}$ we have to determine the related minimum distance point $\bar{x}_{\bar{\mathcal{T}}}^{\min} := \min_{\bar{x}_{\bar{\mathcal{T}}} \in \bar{\Gamma}_{\bar{\mathcal{T}}}} \|\bar{x}_{\bar{\mathcal{B}}} - \bar{x}_{\bar{\mathcal{T}}}\|$ of the contact surface of $\bar{\mathcal{T}}$. Now, we can formulate a non-penetration condition

$$\bar{g}_{NP} = (\bar{x}_{\bar{\mathcal{B}}} - \bar{x}_{\bar{\mathcal{T}}}^{\min}) \cdot \bar{n}_{\bar{\mathcal{T}}}^{\min} \geq 0, \quad \bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_{\bar{\mathcal{B}}}. \quad (1)$$

Alternatively, for all points $\bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_c := \{\bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_{\bar{\mathcal{B}}} \mid \bar{g}_{NP} < 0\}$ which penetrate into the master body, the amount of penetration can be computed by

$$\bar{g}_N = (\bar{x}_{\bar{\mathcal{B}}} - \bar{x}_{\bar{\mathcal{T}}}^{\min}) \cdot \bar{n}_{\bar{\mathcal{T}}}^{\min}, \quad \bar{x}_{\bar{\mathcal{B}}} \in \bar{\Gamma}_c, \quad (2)$$




Fig. 5 Illustration for the determination of active contact nodes and the amount of penetration

and is set to zero for all other points. Here, $\bar{n}_{\bar{T}}^{\min}$ is the outward normal of the rigid tool at $\bar{x}_{\bar{T}}^{\min}$; see Fig. 5.

Since the contact partners of the sheet metal are assumed to be rigid, the tools are not discretized by finite elements but the contact surfaces are characterized by analytical functions. This also simplifies the computation of the related minimum distance point and, hence, the computation of the outward normal direction as well as of the amount of penetration. For a detailed description of contact between two deformable bodies we refer to [99, Ch. 4.1].

As in standard finite element simulations of continuum mechanics problems, we are interested in the minimization of an energy functional $\tilde{\Pi}$, but under additional consideration of an inequality constraint due to the non-penetration condition (Eq. (1)). Constrained optimization can be performed, e.g., using the (quadratic) penalty method, where an additional penalty term is added to the objective function if the constraint is violated; see [76, Ch. 17]. Let us note that the incorporation of contact constraints into a finite element formulation does not change the equations describing the behavior of the bodies coming into contact [99].

There are several different strategies to incorporate contact constraints into finite element formulations, where the penalty method and the Lagrange multiplier method are the most prominent members; see [99, Ch. 6.3]. The penalty method is the most widely used strategy to incorporate contact constraints into finite element simulation software [99, Ch. 10.3.3] because the number of unknowns does not increase. In comparison to the Lagrange multiplier method [99], the contact constraints are only resolved approximately and a small penetration depending on the choice of the penalty parameter $\bar{\varepsilon}_N > 0$ is allowed. For $\bar{\varepsilon}_N \rightarrow \infty$, the exact solution of the Lagrange multiplier method is reached, but for higher penalty parameters $\bar{\varepsilon}_N$, the resulting system of equations becomes ill-conditioned [99]. For a suggestion of a choice of the penalty parameter $\bar{\varepsilon}_N$, see [99, Remark 10.2].

Using the penalty method, we have to add an additional term

$$\tilde{\Pi}_P = \int_{\overline{\Gamma}_c} \frac{1}{2} \cdot \overline{\varepsilon}_N \cdot \overline{g}_N^2 \, dA$$

to the energy functional $\tilde{\Pi}$ for all active contact nodes $\overline{x}_{\mathcal{B}} \in \overline{\Gamma}_c$ [99, Ch. 6.3]. The penalty parameter $\overline{\varepsilon}_N$ can be interpreted as the stiffness of a spring in the contact interface; see [99, Ch. 2.1.3]. Let us note that the definition of an active set is different from standard textbooks as [76, Def. 12.1], where points belong to the active set if they fulfill equality of the inequality constraint. Other authors like Konyukhov and Schweizerhof considered the Heaviside function to follow the common definition of an active set; see, e.g., [67, 68]. Since the energy functional is changed due to the contact constraints, also the resulting stiffness matrix and right-hand-side are affected.

3.4 Algorithmic Improvements in FE2TI

In simulations making use of load stepping (or pseudo time stepping) as a globalization strategy, as is the case in FE2TI (see Sect. 4), the time to solution strongly depends on the number of load steps as well as on the number of macroscopic Newton iterations per load step. The required time of a single macroscopic Newton step again depends on the time to solution of the microscopic problems.

While a large load step may seem desirable, it can lead to slow convergence or even divergence; convergence can be forced by reducing the load step size thus increasing the total number of load steps; this can be observed in Table 1. To adapt the size of the load steps, we use a simple dynamic load step strategy; see Sect. 3.4.1.

Keeping the number of macroscopic Newton iterations as small as possible is directly related to a good choice of the initial value of a single load step. For a better prediction of the initial value, we use linear extrapolation; see Sect. 3.4.2. This is especially relevant for our contact problems where quadratic convergence of Newton’s method can be lost.

Table 1 First 2 mm displacement of the forming tool with constant load increments \bar{l} of $3.125\text{e}−3$ (first row) and $1.0\text{e}−1$ (second row) for the specimen with a parallel shaft width of 50 mm; two MPI ranks per core; computed on JUWELS [45]

	Cov. dist.	Load	Newton	∅ Newton its.		∅ Time per	∅ Time per
	punch	steps	its.	per load step	Runtime	load step	Newton it.
Load increment 3.125e−3	2 mm	640	966	1.51	7595.30 s	11.87 s	7.86 s
Load increment 1.0e−1	2 mm	20	130	6.50	1407.03 s	70.45 s	10.82 s

3.4.1 Dynamic Loadstepping

Our load step strategy depends on macroscopic as well as microscopic information. The macro load increment \bar{l} is reduced when microscopic stagnation is observed or when a maximum number of macroscopic Newton iterations per load step is reached. Stagnation on the RVE level is detected whenever the norm of the current microscopic Newton iteration compared to the previous one does not reduce after more than six microscopic Newton iterations or if the number of microscopic Newton iterations is larger than 20. The load step is increased based on the number of macroscopic Newton iterations per load step. Note that \bar{l} is not allowed to exceed a predefined maximum load increment \bar{l}^{\max} . For details, see Fig. 6.

Whenever stagnation in a microscopic problem occurs, the microscopic solver gives this information to the macroscopic level and the load step is repeated with a reduced load increment. Otherwise, stagnation of a microscopic problem would lead to a simulation break down due to a missing tangent modulus and stresses in the macro Gauß point where the micro problem is attached.

3.4.2 (Linear) Extrapolation

For Newton-type methods, it is important to choose a good initial value. If the initial value is close to the solution, only a few Newton iterations are necessary. As in [64],

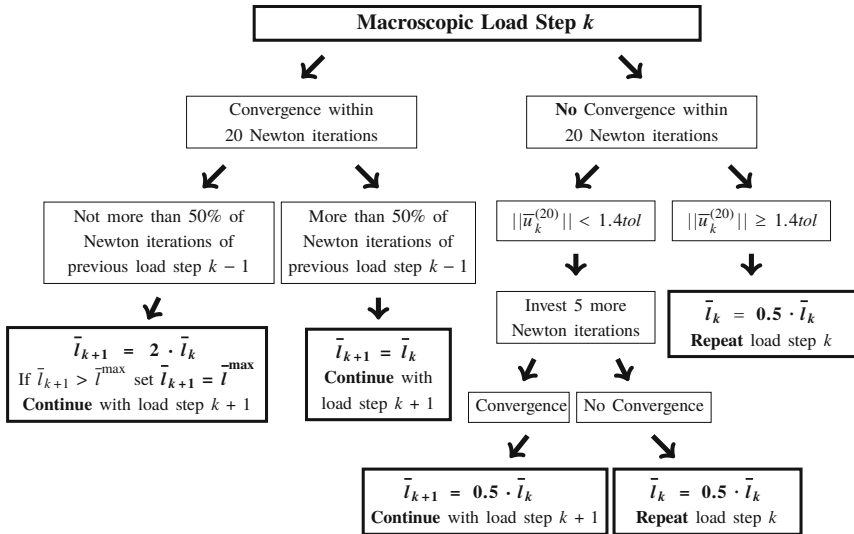


Fig. 6 Schematic procedure of reduction and increase of the load increment \bar{l} depending on macroscopic events. Here, \bar{l}^{\max} is a predefined maximum load, $\|\bar{u}_k^{(20)}\|$ is the norm of the solution of the 20th Newton iteration of the current load step k , and tol is the Newton tolerance

we use linear extrapolation based on the converged solutions \bar{u}_k and \bar{u}_{k-1} of the former two load steps to guess a good initial value $\bar{u}_{k+1}^{(0)}$,

$$\bar{u}_{k+1}^{(0)} = \bar{u}_{k-1} + \frac{\bar{l}_{k+1} - \bar{l}_{k-1}}{\bar{l}_k - \bar{l}_{k-1}} \cdot (\bar{u}_k - \bar{u}_{k-1}).$$

In [64], this technique was already successfully used in the FE² simulations using the FE2TI software package without contact constraints. For results using second order extrapolation, we refer to [95, Ch. 4.2.2].

3.4.3 Checkpoint/Restart

To perform the simulation of the Nakajima test until material failure of a sample sheet, i.e., until a failure zone occurs in the metal sheet, often significant computation times are needed, even if the full supercomputer is available. To reduce the consequences of hardware failures and also to overcome specific wall time limits on HPC systems, we equipped our FE2TI package with synchronous Checkpoint/Restart (CR). We integrated the CRAFT library (Checkpoint/Restart and Automatic Fault Tolerance) [88], which was developed in the second phase of the SPPEXA project ESSEX. Let us note that we use synchronous application level checkpointing with different checkpoints for the macroscopic level and the microscopic level.

In [21], different checkpoint intervals are introduced based on the expected run time of the simulation and the mean time of hardware failure of the HPC system the software is performed on, but in all our simulations presented here, we choose a fixed checkpoint interval of 75 load steps. Here, we do not take into account that the load increment may change during the simulation and that small load increments are usually faster. As an improvement, the checkpointing could take into account the displacement of the forming tool or a fixed wall time interval could be used which also could depend on the mean time of hardware failure.

4 Numerical Simulation of the Nakajima Test Using FE2TI

For the simulation of the Nakajima test, we use our highly scalable parallel software package FE2TI; see Sect. 3. For the solution of the boundary value problems on both scales, we here use the sparse direct solver package PARDISO [81]. Since we are considering a DP600 grade of steel, we use a fitted J₂ elasto-plasticity model on the microscopic level; see [14, Fig. 10]. Throughout this article, we use structured Q2 finite element discretizations for the sample sheet and an unstructured P2 finite element discretization for the RVEs. Both, the macroscopic as well as the

microscopic meshes, are generated using the open source software package GMSH [34]. We use the load stepping and extrapolation described in Sects. 3.4.1 and 3.4.2.

In the Nakajima test, the macroscopic deformation is driven by the rigid punch. Hence, load stepping is applied to the movement of the forming tool, where the hemispherical punch moves a small step in upward direction in each load step.

Since in reality a tribological system is set up to avoid friction between the hemispherical punch and the sheet metal [77], we consider frictionless contact. Hence, we have to deal exclusively with contact conditions in normal direction of the rigid tools. Considering frictionless contact, we neglect friction between the specimen and the blank holder or die.

4.1 Description of Specimen Geometry

In our simulations, we consider specimens with a central parallel shaft and also a completely circular specimen. The length of the parallel shaft is 25 mm and the fillet radius is 30 mm, which both fit to the normed range in [77]; also see Fig. 1 (right).

For all specimens, the material is assumed to be completely clamped by the bead, which has a radius of 86.5 mm. We therefore only consider material points $\bar{p} = (\bar{p}_x, \bar{p}_y, \bar{p}_z)$ which have a distance of at most 86.5 mm to the center of the sample sheet; see Fig. 1 (right) for an example of a sample sheet with a parallel shaft width of 90 mm. In our simulations, the sample sheet is always 1 mm thick, and we consider specimens with a parallel shaft width of 30, 50, 70, 90, 100, 110, 125, and 129 mm as well as the completely circular sample sheet. Note that the center $\bar{c}^b = (\bar{c}_x, \bar{c}_y, \bar{c}_z)$ of the bottom surface of all sample sheets is placed in the origin of the coordinate system.

The specifications of the rigid tools are also within the range given in [77]. The radius of the hemispherical punch is 50 mm. The blank holder is a square plate of 173 mm \times 173 mm with a circular hole in the middle with a radius of 55 mm; see the inner circle in Fig. 1 (right). The die is placed within a distance of 5 mm to the rigid punch, i.e., the inner wall of the die also has a radius of 55 mm; see, again, the inner circle in Fig. 1 (right). The die radius (see Fig. 1 (left)) is 10 mm, i.e., all material points \bar{p} with $\sqrt{\bar{p}_x^2 + \bar{p}_y^2} \geq 65$ are potentially clamped; see the outer circle in Fig. 1 (right).

For all sample sheets with a parallel shaft width less than 90 mm as well as for the completely circular specimen, we only consider material points \bar{p} with $\sqrt{\bar{p}_x^2 + \bar{p}_y^2} \leq 65$ and choose all points with $\sqrt{\bar{p}_x^2 + \bar{p}_y^2} = 65$ as Dirichlet boundary nodes. For specimens with wider parallel shaft widths, we choose boundary conditions analogously to [43]; see also [95].




Fig. 7 **Left:** Symmetric computational domain from the full sample sheet and additional Dirichlet boundary conditions along the symmetric boundary surfaces. **Right:** Orientation of the RVEs for the different quadrants of the full geometry of the sample sheet after mirroring the first quadrant due to the symmetric assumption

4.2 Exploiting Symmetry

The setup of the Nakajima test is symmetric. Assuming that the failure zone evolves symmetrically, i.e., along the vertical center line, it is sufficient to only simulate a quarter of the full geometry and to rebuild the full solution by mirroring; see Fig. 7 (left). This is only exact, if the RVEs are also symmetric, since mirroring of the macroscopic solution also implies mirroring of the RVEs; see Fig. 7 (right). Hence, for an asymmetric RVE, we violate the assumption of a periodic unit cell because mirroring leads to a change in orientation for all four quadrants. In this case, the solution generated by the symmetric assumption is only an approximation to the solution of the simulation using the full geometry of the sample sheet, even for the first quadrant of the full geometry. Nevertheless, we use the symmetric assumption throughout this article, even when the RVEs are asymmetric, to reduce the number of MPI ranks by 75%; see Sect. 3.4. As a sanity check, we have also performed simulations for the full geometry; these will, however, be presented elsewhere due to space limitations.

For the simulation of one quarter, we have to add further boundary conditions along the symmetric boundaries of the quarter. Then, we have to fix all y -coordinates of macroscopic material points \bar{p} with $\bar{p}_y = \bar{c}_y$. Analogously, we have to fix all x -coordinates for macroscopic points with $\bar{p}_x = \bar{c}_x$; see Fig. 7 (left) and Sect. 4.1.

4.3 Failure Criterion

For the detection of macroscopic crack initialization, we have to formulate a failure criterion and a maximum critical value, which will indicate the initialization of failure. Note that, for the computation of the forming limit, we do not need to simulate cracks or crack propagation. Instead, we are only interested to compute when structural failure occurs.

We use the Cockcroft and Latham criterion [19],

$$\overline{W}_k = \overline{W}(\overline{\alpha}_k) = \int_0^{\overline{\alpha}_k} \max(\overline{\sigma}_I(\overline{\alpha}), 0) d\overline{\alpha}, \quad (3)$$

which was used by Tarigopula et al. [93] for analyzing large deformation in DP steels. It depends on the maximum principal stress component $\overline{\sigma}_I$ and the equivalent plastic strain $\overline{\alpha}_k := \overline{\alpha}(t_k)$ at load step k (pseudo time t_k) and is integrated over $\overline{\alpha}$. Since $\overline{\alpha}$ depends on the load step, this also holds for the failure criterion \overline{W} and the stress $\overline{\sigma}$. In general, in FE², $\overline{\alpha}_k$ is not known explicitly but can be approximated by the volumetric average⁴ $\tilde{\alpha}_k := \langle \alpha(t_k) \rangle$ from the microscopic level at load step k .

The value of \overline{W} is computed at each Gauß point and is accumulated throughout the loading process until at least one Gauß point exceeds a critical value \overline{W}_c at which failure initializes, i.e., $\overline{W} \geq \overline{W}_c$ is fulfilled. Tarigopula et al. provide the value $\overline{W}_c = 590 - 610$ MPa for DP800 grade of steel; see [93]. Since we consider DP600 grade of steel, the critical value should be smaller in our case.

Equation (3) is approximated by numerical integration and using $\tilde{\alpha}$

$$\begin{aligned} \overline{W}_k &\approx \overline{W}(\tilde{\alpha}_k) = \int_0^{\tilde{\alpha}_k} \max(\overline{\sigma}_I(\tilde{\alpha}), 0) d\tilde{\alpha} \approx \sum_{i=1}^k \max(\overline{\sigma}_I(\tilde{\alpha}_i), 0) \cdot (\tilde{\alpha}_i - \tilde{\alpha}_{i-1}) \\ &= \overline{W}_{k-1} + \max(\overline{\sigma}_I(\tilde{\alpha}_k), 0) \cdot (\tilde{\alpha}_k - \tilde{\alpha}_{k-1}). \end{aligned} \quad (4)$$

where $(\tilde{\alpha}_i - \tilde{\alpha}_{i-1})$ is nothing else than the increment of the equivalent plastic strain from load step $i - 1$ to load step i and $\overline{W}_0 = \tilde{\alpha}_0 = 0$. Hence, we can sum up the failure criterion \overline{W} over all load steps and summation is performed whenever a load step reached convergence. See Fig. 9 for an example of the evolution of the failure criterion \overline{W} during a Nakajima simulation.

Let us note that the failure criterion is formulated such that ductile failure takes place due to tensile stresses and shear stresses, where the effect of tensile stresses is considered by usage of the maximum positive principal value of Cauchy stress and the effect of shear stresses and work-hardening is considered through the equivalent plastic strain increment.

4.4 Numerical Realization of the Experimental Cross Section Method

In the experiment, major and minor strains $\overline{\varepsilon}_1$ and $\overline{\varepsilon}_2$ are evaluated at the top surface of the sample sheets along cross sections. However, the simulation only provides

⁴Note that here and in the following, all volumetric averages obtained from microscopic level are in brackets $\langle \cdot \rangle$.

exact macroscopic values in the integration points, which generally do not coincide with the finite element nodes. Therefore, the simulations do not provide any exact macroscopic values on the sample sheets surface, and we decided to consider cross sections along those Gauß points closest to the upper surface.

For the computation of major and minor strains $\bar{\epsilon}_1$ and $\bar{\epsilon}_2$ we transform our resulting 3D strain tensor to the 2D plane strain tensor and then follow the strategy as described in [97]. For further details, we refer to [95].

In this article, we show results for computations using the symmetric test setup of the experiment; see Sect. 4.2. This automatically implies that we assume that the failure zone evolves along the vertical center line and the center of the failure zone is identical to the center of the upper surface of the sample. For the computations using symmetry, all discretizations have finite element nodes at the center of the probe. Keeping in mind that we choose cross sections along Gauß points, no cross section cuts through the center of the failure zone. To keep the distance between the first cross section and the center of the failure zone as small as possible, we consider integration points with the smallest distance to the horizontal center line as first cross section. For simplicity, the other cross sections are along the remaining Gauß points of the same finite elements which were used for the first cross section. Hence, the distance between the cross sections depends on the diameter of those finite elements and is smaller than 2 mm in our computations.

Due to the symmetric computations, we only have one side of the cross sections at hand but the other side can be simply generated by mirroring; see Fig. 2 and the upper pictures in Fig. 8.

Note that the cross section method can only be used for specimens with a single failure zone. Unfortunately, in our simulations the failure zone does not always evolve along the vertical center line but parallel to it for sample sheets with a parallel shaft width of 100 mm or more. Hence, mirroring leads to the occurrence of a second failure zone; see Figs. 11 (left) and 12. For these specimens, first simulations using the upper half or the complete domain of the sample sheet also lead to two failure zones parallel to the vertical center line. For simulations obtaining two failure zones, we assume that the maximum major strain along the cross section defines the position of the failure zone. Neglecting all values between the vertical center line and the maximum major strain and shifting the failure zone to the vertical center line, we can proceed as before; see Fig. 8.

So far it is not clear, whether the symmetry assumptions hold for specimens with failure zones parallel to the vertical center line. In future work, we have to perform a comparison with the full geometry.

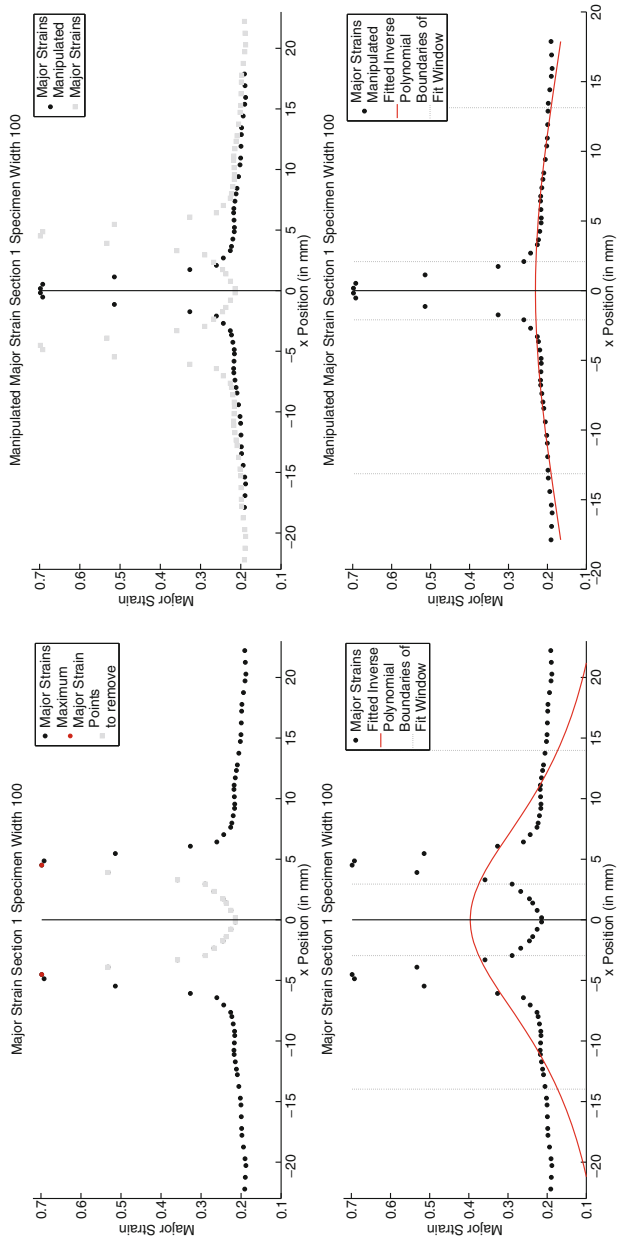


Fig. 8 Major strains along the first cross section for the specimen with a width of the parallel shaft of 100 mm. See also the description of the cross section method in Sect. 2. **Top:** Original values (left) and major strain values after shifting maximum values to the center as described at the end of Sect. 4.4 (right). **Bottom:** Fitted inverse second-order polynomials of original (left) and shifted (right) major strains. Computation of optimal fit windows is described in [77, Ch. 5.2.3, 5.2.4]

5 The Virtual Forming Limit Diagram Computed with FE2TI

Since we use the symmetry assumptions for the numerical simulation of an FLC, we perform our simulations on a quarter of the full specimen and rebuild the solution by mirroring; see also Sect. 4.2. The computational domain is discretized by structured Q2 finite elements, where the number of finite elements depends on the width of the parallel shaft as well as on the chosen boundary conditions; see Sect. 4.1 as well as Table 2. Note that we use (only) two Q2 finite elements in the direction of thickness for all specimens.

For the microscopic problems, we consider an SSRVE with two ellipsoidal inclusions discretized by 1470 P2 finite elements and 7152 d.o.f. in an unstructured manner. As mentioned before, the resulting tangent problems on the microscopic level are solved by using the direct solver package PARDISO [81] and each problem is solved on an individual MPI rank. Hence, the number of macroscopic finite elements automatically determines the number of required MPI ranks, which is 27 times the number of finite elements. We have performed all our simulations on JUWELS [45] using 2 MPI ranks per core and a penalty parameter of 500. For the specifications of the rigid tools, we refer to Sect. 4.1.

As an initial load increment, we choose 0.1 mm and define $\bar{l}^{\max} = 0.2$ mm as maximum allowed load increment. Our stopping criterion is either based on reaching a predefined covered distance of the forming tool of 40 mm or on the load increment and not on the failure criterion, since we have only little experience how to choose the critical value \bar{W}_c to detect failure. Simulation stops if the load increment of 10 successive load steps is smaller than a predefined allowed minimum, which is the initial load increment multiplied by 10^{-4} , or, if the load increment has to be reduced 7 times within a single load step. This is motivated by the fact that small load increments indicate hard numerical problems which are expected in case of material failure.

We have summarized some data on our Nakajima simulations including the number of restarts in Table 2. Note that most restarts are caused by reaching the requested wall time and only in few cases, if any, by hardware errors.

For all specimens with a parallel shaft, we obtain comparable results, which are characterized in the following. After a certain covered distance of the tool, a local increase in the failure value \bar{W} , the major strains $\bar{\epsilon}_1$, the equivalent plastic strain $\tilde{\alpha}$, the thinning rate, and the von Mises stress can be detected almost simultaneously along an area parallel to the vertical center line. Later, the values continue to rise, especially in this area, so that the degree of localization increases; see Fig. 9. Finally, some microscopic problems within the aforementioned localized area cause the simulation to end. At this point, however, a pronounced change in thickness can be observed within the localized area, which can be associated with material failure; see cross sections in Figs. 12 and 10 (top right) as well as the upper right picture in Fig. 11 (left).

Table 2 Nakajima simulations with different specimens

	Width 30	Width 40	Width 50	Width 70	Width 90	Width 100	Width 110	Width 115	Width 129	Full circular
Macro finite elements (Q_2)	424	392	460	482	558	558	580	574	574	542
MPI ranks	11,448	10,584	12,420	13,014	15,066	15,066	15,660	15,498	15,498	14,634
Macro d.o.f.	13,725	12,705	14,804	15,465	17,835	17,835	18,495	18,195	18,195	17,145
Covered dist. punch (mm)	27.156	28.325	29.242	29.896	30.734	31.654	32.593	35.502	36.566	40,000
Load steps	736	751	806	901	985	898	780	1358	1651	569
Newton its.	8148	8510	9272	9823	10,976	8604	8540	9556	10,012	8064
Runtime (h)	17.60	16.00	19.35	22.00	27.00	24.00	21.00	25.25	24.17	19.25
Restarts	1	0	1	1	2	1	1	2	1	1
Overhead load steps	74	–	6	8	68	16	24	104	3	12
Overhead runtime (h)	1.70	–	0.23	0.23	1.71	0.40	0.58	2.46	0.25	0.41

Microscopic problems: SSRVE with two ellipsoidal inclusions; 1470 unstructured P2 finite elements and 7152 d.o.f. Two MPI ranks per core; computed on JUWELS [45]; one microscopic problem computed per MPI rank. Overall problem size is obtained by multiplying number of MPI ranks by d.o.f. of microscopic SSRVE and adding number of macroscopic d.o.f. resulting in 80–112 million d.o.f.

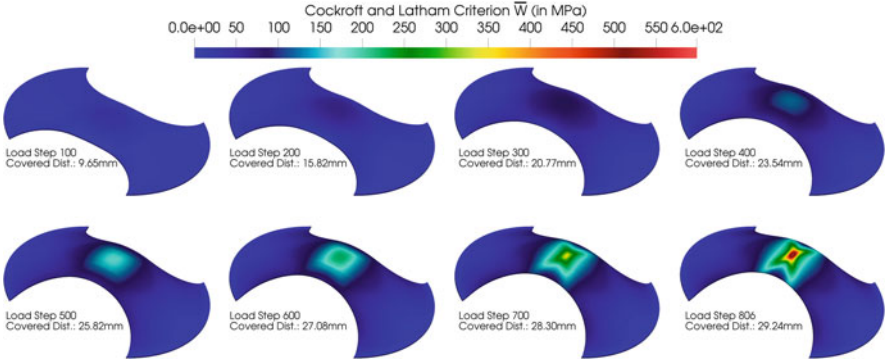


Fig. 9 Evolution of failure criterion \overline{W} during the simulation for the specimen with a width of the parallel shaft of 50 mm

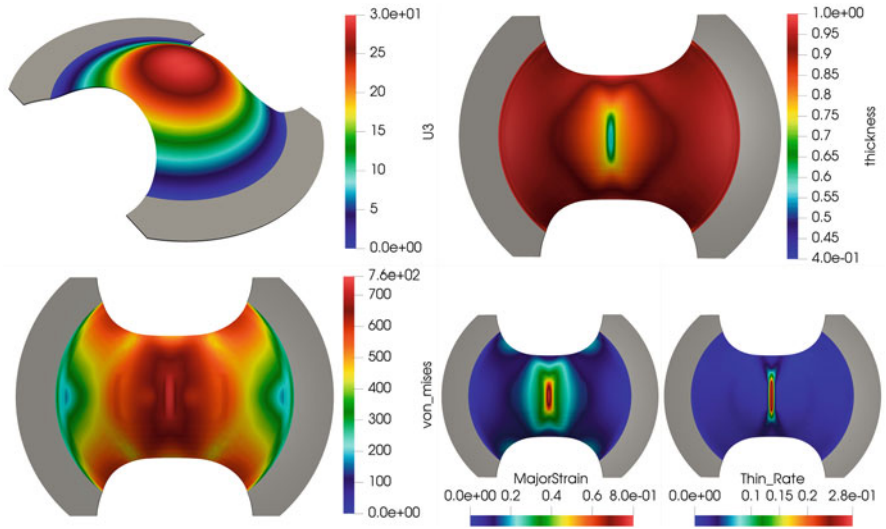


Fig. 10 Final solution of the simulation with a specimen with a width of the parallel shaft of 70 mm; z-direction (top left), thickness (top right), von Mises stress (bottom left), major strain (bottom center), and thinning rate (bottom right). The grey area is not simulated since material is totally clamped between die and blank holder

For the completely circular specimen we do not observe any localized effects along an area parallel to the vertical center line, even if we reach the predefined distance of 40 mm. Instead, we obtain a similar behavior over nearly the complete contact area; see Fig. 11 (right).

The results for specimens with a parallel shaft can be divided into two groups. The first group contains all samples with a parallel shaft width of at most 90 mm. For these specimens material failure occurs along the vertical center line; see Fig. 10.




Fig. 11 Final solution of the simulation with a specimen with a width of the parallel shaft of 100 mm (**left**) and the completely circular specimen (**right**); variables and color bars as in Fig. 10. **Left:** Material between blank holder and die is simulated since material movement is allowed. **Right:** Material between blank holder and die is assumed to be clamped (dark grey) and hence is not simulated

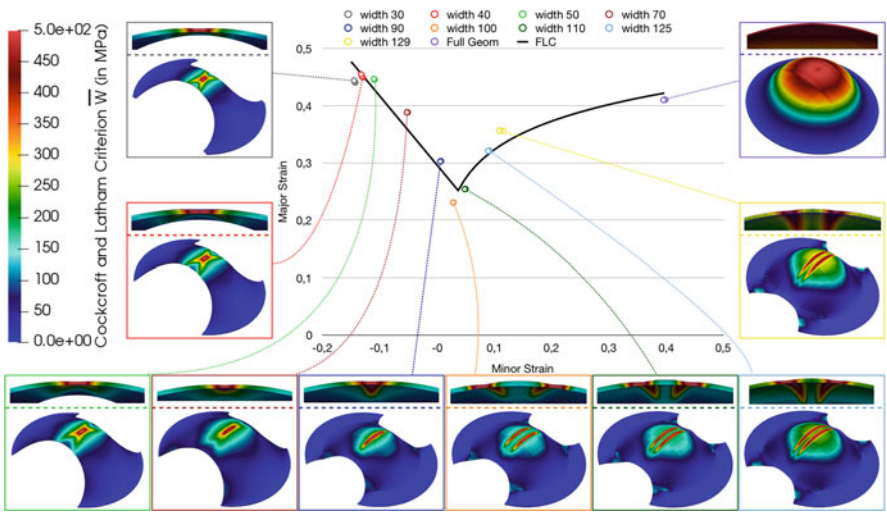


Fig. 12 Final results of distribution of Cockcroft and Latham failure value \overline{W} for all Nakajima simulations and associated Forming Limit Diagram (FLD) with FLC (black curve) for $\overline{W}_c = 450$. In the cross section, one can identify local necking in thickness for all but the completely circular specimen

All specimens with a wider parallel shaft belong to the second group, which can be characterized by material failure parallel to the vertical center line. For these samples, mirroring of the computed solution leads to the occurrence of a second failure zone; see Figs. 11 (left) and 12. As mentioned before, first numerical tests indicate that this also holds when simulating the complete specimen or the upper half of it. Hence, we decided to use the results to determine an FLC. The adaptations to evaluate major and minor strains for failure zones parallel to the vertical center line are mentioned in Sect. 4.4; see also Fig. 8. Without manipulating the values along the cross sections, we would obtain unphysically high values in the FLD.

Based on all available results we have subsequently defined a critical value of $\overline{W}_c = 450$ for the failure criterion. Hence, we have to find for all simulations the corresponding load step for which the failure value in at least one point exceeds the critical value $\overline{W}_c = 450$ for the first time; see Table 3.

When we have found the corresponding load steps, we have to compute the major and minor strain values along the cross sections perpendicular to the failure zone and generate for each cross section one point in the FLD. So, we come up with 30 different pairs of major and minor strains wherein each three belong to one specimen; see Fig. 12. The final FLC is derived by regression; see black curve in Fig. 12 and [95] for details.

6 Linear and Nonlinear Parallel Solvers

For large scale computations using FE2TI, scalable parallel implicit solver algorithms are needed for the problems on the microscale as well as the macroscale [64]. Another focus of the EXASTEEL project therefore was on solver algorithms, i.e., efficient and highly parallel scalable implicit solvers for linear and nonlinear problems arising from a finite element discretization of linear and nonlinear partial differential equations; see, e.g., [5, 54–57, 59–61, 63].

Here, nonlinear domain decomposition (DD) approaches, namely nonlinear FETI-DP (Finite Element Tearing and Interconnecting-Dual Primal) and nonlinear BDDC (Balancing Domain Decomposition by Constraints) methods, have been introduced in the first phase of EXASTEEL and improved during the second phase, where also new variants were introduced.

In [64], our new nonlinear FETI-DP solver algorithms were then applied within large FE2TI simulations for the first time: We have used Nonlinear-FETI-DP-1 as a parallel implicit solver for the microscopic problems using 114,688 KNL cores of the Theta many-core supercomputer (Argonne National Laboratory) [64, section 4.9]. However, the nonlinear DD methods developed in the project have a broad range of applications, e.g., in hyperelasticity, elasto-plasticity, or nonlinear diffusion problems.

To further improve the performance of the nonlinear solvers, also the efficiency and parallel scalability of all building blocks was in the focus of the EXASTEEL project [5, 54–56, 60, 61].

In this section, we describe very briefly our nonlinear domain decomposition approaches and sum up the achievements and highlights obtained within the past 6 years.

Table 3 Identification of first load step with failure values higher than the critical value $\overline{W}_c = 450$; two MPI ranks per core; computed on JUWELS[45]

	Width 30	Width 40	Width 50	Width 70	Width 90	Width 100	Width 110	Width 115	Width 129	Full circular
Covered dist. punch (mm)	26.725	27.813	28.747	29.689	30.361	30.986	31.972	34.976	36.332	38.488
Load step	684	710	743	873	896	785	701	768	800	541
Newton its.	7708	8099	8737	9695	10.635	8318	8253	8931	9142	7676
Maximum \overline{W}	454.49	451.15	451.33	451.24	450.24	450.08	451.92	451.27	450.14	450.91
Runtime (h)	16.55	15.16	17.96	21.39	26.14	23.09	19.80	23.09	22.36	18.15

6.1 Nonlinear FETI-DP Framework

Classical domain decomposition methods are robust and highly scalable divide-and-conquer algorithms for the iterative solution of discretized linear partial differential equations. In the case of FETI-DP methods [27, 28, 47–50], the computational domain is decomposed into nonoverlapping subdomains which are distributed to the available compute cores. FETI-DP methods are well established in structural mechanics and have been awarded a Gordon Bell price [10].

The robustness and scalability originates from the sparse direct solvers applied on the subdomains, combined with a carefully designed coarse problem. The coarse problem, though necessary for robustness, is a potential scaling bottleneck, since its size grows with the number of subdomains, i.e., with the number of parallel cores. In order to retain scalability, the coarse solution can be approximated, e.g., by algebraic multigrid methods; see [46, 48]. Finally, to solve nonlinear problems, a linearization with Newton’s method is usually applied first, and the linearized tangential systems are then solved, e.g., by FETI-DP. We refer to the latter approach as Newton-Krylov-FETI-DP.

In contrast, in nonlinear FETI-DP or BDDC methods, first introduced in [51], the discretized nonlinear partial differential equation is decomposed into small and independent nonlinear problems before linearization. In the case of nonlinear FETI-DP, a nonlinear coarse problem is added by strongly coupling the local nonlinear problems in few variables on the interface of the domain decomposition, as, e.g., vertices or averages over edges or faces. This leads to a nonlinear FETI-DP saddle point system; see, e.g., [51, eq. (3.2)], [54, eq. (3.4)], or [59, eq. (1)]. Also, a partially nonlinear elimination process of sets of variables from the nonlinear FETI-DP saddle point system is possible before linearization. The nonlinear elimination process can be interpreted as a nonlinear right-preconditioner, which we described in [59, Section 2.5] in detail. There we also introduced a unified framework to describe different nonlinear FETI-DP and BDDC methods. The selection of the elimination set finally defines the nonlinear FETI-DP method precisely. We discussed four canonical choices in [59], but other choices are feasible and possibly beneficial. Let us briefly repeat the four variants from [59]. In FETI-DP, all degrees of freedom or variables are divided into three classes. First, all variables belonging to the interior of the subdomains are grouped into the set I (marked as circles in Fig. 13), second, all variables of the global coarse problem are grouped into the set Π of so-called primal variables (marked as squares in Fig. 13), and third, all local degrees of freedom on the interface are grouped into the set Δ of so-called dual variables (marked as dots in Fig. 13). Let us remark that continuity of the solution in the dual degrees of freedom is enforced iteratively by enforcing zero jump constraints with a Lagrangian approach; see [27] for details. Finally, Nonlinear-FETI-DP-1 is defined by eliminating nothing, Nonlinear-FETI-DP-2 by eliminating everything, Nonlinear-FETI-DP-3 by eliminating the interior and dual variables, and finally Nonlinear-FETI-DP-4 by eliminating the interior variables; see Fig. 13 for a visualization of the different variants.




Fig. 13 The coarse problem is marked with squares, the interior degrees of freedom by circles, and the dual degrees of freedom by dots. The colored degrees of freedom are eliminated nonlinearly before linearization in the different variants. **From left to right:** Nonlinear-FETI-DP-1, Nonlinear-FETI-DP-4, Nonlinear-FETI-DP-3, Nonlinear-FETI-DP-2




Fig. 14 Comparison of classical Newton-Krylov-FETI-DP (NK), Nonlinear-FETI-DP-1 (NL1), Nonlinear-FETI-DP-2 (NL2), Nonlinear-FETI-DP-3 (NL3), and Nonlinear-FETI-DP-4 (NL4) using the JUQUEEN supercomputer. **Left:** Nonlinear p -Laplace inclusions not touching the interface; 40k d.o.f. per subdomain; **Right:** Nonlinear p -Laplace channels crossing the interface; 160k d.o.f. per subdomain. For a detailed discussion of the results see also [59, Fig. 10 and 12] and corresponding descriptions. If an appropriate nonlinear elimination set is chosen for the problem (as is the case here in NL3) then the nonlinear method outperforms the classical Newton-Krylov approach significantly

If the elimination set is chosen appropriately, nonlinear FETI-DP methods can outperform classical methods, i.e., Newton-Krylov-FETI-DP. In [42], a heuristic approach is suggested to eliminate the areas with strong nonlinear effects. For illustration let us consider combinations of the nonlinear p -Laplace equation and the linear Laplace equation, where the nonlinearities either touch the subdomain interface or are at a distance. In the latter case, we choose nonlinear inclusions of p -Laplace enclosed in the subdomains and in the first case nonlinear channels of p -Laplace cutting certain edges. For a detailed description of the chosen model problem, see [59, Section 5.1, Fig. 7]. Considering the inclusions, all nonlinear FETI-DP methods work well. Considering the example with channels, it is necessary to at least eliminate the dual interface variables; see Fig. 14 or [59]. Let us note that Nonlinear-FETI-DP-1 performs well in both cases, which is a result of a careful choice of the initial value; see [51, Section 3.3] for details.




Fig. 15 Weak scalability of Nonlinear-FETI-DP-1 for a p -Laplace problem computed on the Mira supercomputer; largest problem has 62.9 billion d.o.f.; results from [57]; see also [57, Fig. 4 (left)]

Another benefit of nonlinear domain decomposition approaches is the localization of work, which increases the scalability of these methods. This can be verified either in Fig. 14 or in our larger weak scaling experiments on Mira published in [57] and presented in Fig. 15, where an algebraic multigrid preconditioner from the BoomerAMG package [38] is used to approximate the coarse solve to obtain scalability without losing robustness.

We have also considered approaches to improve the convergence of nonlinear FETI-DP methods. We have introduced heuristics to determine if a nonlinear elimination is useful in a certain Newton step or not. Additionally, the elimination process is approximated up to a necessary tolerance to save computational cost. This approach is called NL-ane (approximate nonlinear elimination) and is also discussed in [59, 62]. We recently also considered a globalization strategy using the SQP approach; see Sect. 6.1.2.

Finally, we successfully investigated a hybrid parallelization of FETI-DP using PARDISO in [55], and also considered nonlinear FETI-DP and BDDC methods where the sparse direct solvers are replaced by preconditioners; see [56, 60, 61].

6.1.1 Improving Energy Efficiency

In classical Newton-Krylov methods, global synchronization occurs in every Krylov iteration. Global synchronization can be significantly more coarse-grained in our nonlinear domain decomposition methods since the Krylov iteration can be asyn-

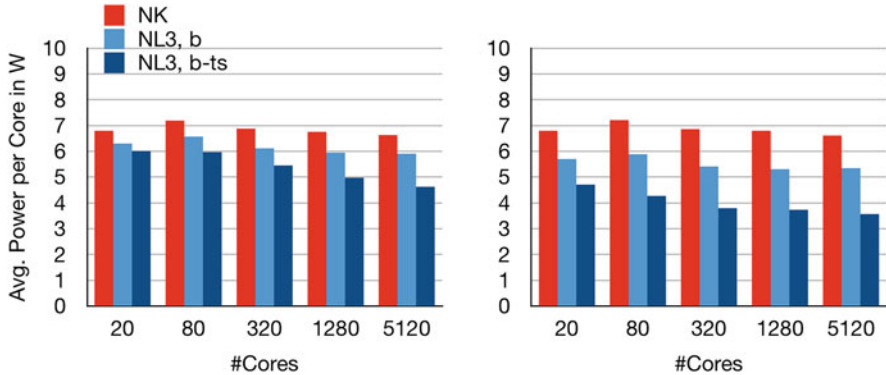


Fig. 16 Power per core for Newton-Krylov-FETI-DP (NK) and Nonlinear-FETI-DP-3 (NL3) with normal barrier (b) and test-sleep approach (b-ts) on Meggie cluster. Each subdomain problem has 160k d.o.f. **Left:** Nonlinearity in each of the subdomains; **Right:** Nonlinearity in a single subdomain. See also [63, Fig. 10] and corresponding descriptions for details

chronous between subdomains. In this section, we describe how this can be exploited to save energy when load imbalances occur.

If the nonlinear elimination set in nonlinear FETI-DP is completely local to the subdomains as, e.g., in Nonlinear-FETI-DP-3, the nonlinear subdomain problems can be solved in parallel and asynchronously. This solution process using Newton's method can introduce a load imbalance, even for perfectly balanced meshes, if the number of Newton iterations is different between subdomains; see [63, Fig. 7]. Note that even for these cases, Nonlinear-FETI-DP-3 typically has a shorter time to solution compared with classical approaches; see Fig. 14.

In [63], we have suggested to use a nonblocking barrier in combination with a sleep statement to set idling cores in deep sleep states, to reduce energy consumption. This is feasible in nonlinear parallel domain decomposition since the synchronization between the cores is coarse grained (typically larger than 1 s). During these phases sleeping cores wake up every 10 ms. The wake-up latency itself for current Intel processors is significantly below 1 ms. Therefore, the overhead of the sleep and wake up approach is insignificant compared to the time spawn of global synchronization and does not impact overall performance or scalability. We call this method test-sleep approach. To investigate the energy saving potential, we measured the energy consumption of Nonlinear-FETI-DP-3 in [63] reading out the RAPL hardware counters with LIKWID [94] on the Meggie⁵ cluster at Erlangen Regional Computing Center (RRZE). In Fig. 16, we present the power consumption per core for two different scenarios, i.e., a single nonlinear inclusion in a single subdomain or a single nonlinear inclusion in each subdomain. The total energy

⁵Standard Cluster with Intel Omnipath Interconnect and two Intel Xeon E5-2630 v4 chips per node.

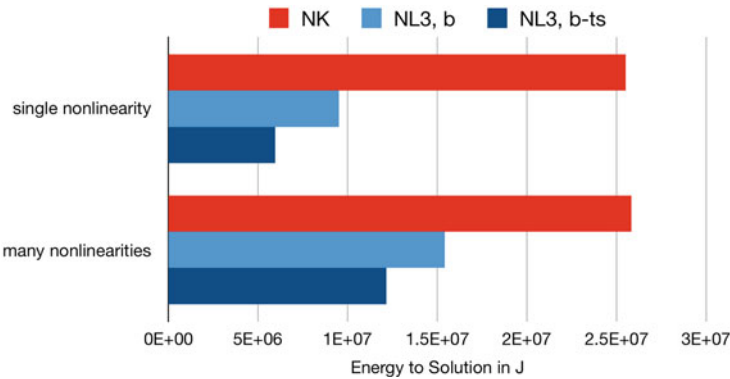


Fig. 17 Total energy to solution for Newton-Krylov-FETI-DP (NK) and Nonlinear-FETI-DP-3 (NL3) with normal barrier (b) and test-sleep approach (b-ts). Computation on 5120 Meggie cores. See also [63, Fig. 9] for a complete weak scaling study

consumed by the nodes during the solution procedure on 5120 cores is shown in Fig. 17. The test-sleep approach also works for alternative nonlinear domain decomposition methods, as, e.g., ASPIN [17]; see [63] for a brief discussion using the ASPIN implementation described in [15] which is provided in PETSc.

6.1.2 Globalization

We consider different globalization strategies for our nonlinear domain decomposition methods. For the different nonlinear FETI-DP methods, we consider trust region methods and also an approach based on the SQP (sequential quadratic programming) method using the exact penalty function $P_\beta^1(\tilde{u}) = J(\tilde{u}) + \beta ||B\tilde{u}||_1$, where J denotes the mechanical energy and $B\tilde{u}$ are the FETI-DP equality constraints; see Table 4 for some results.

Table 4 Number of global iterations for a snap through buckling problem for compressible Neo-Hookean energy with material parameters $E = 210$, $\nu = 0.3$ in 2D; 100 subdomains, 20,402 degrees of freedom; – : failed

Volume force	SQP		Trust region		Without globalization	
	NL-1	NL-4	NL-1	NL-4	NL-1	NL-4
$\begin{bmatrix} 0, -20 \end{bmatrix}^T$	6	6	321	343	–	–
$\begin{bmatrix} 0, -40 \end{bmatrix}^T$	6	6	298	469	–	–
$\begin{bmatrix} 0, -80 \end{bmatrix}^T$	7	7	258	662	–	–




Fig. 18 Left: Weak scaling experiment on the JUQUEEN supercomputer; deformation of a hyperelastic cube with 0.7 billion d.o.f. for the largest computation (heterogeneous Neo-Hooke material model); using Newton-Krylov-BDDC or Nonlinear-BDDC; see [60, Section 4.4, Fig. 15] for details. For 262,144 subdomains Newton-Krylov-BDDC did not converge within 20 Newton iterations. **Right:** Weak scaling of linear BDDC solver with approximate coarse solve (using AMG) on JUQUEEN (BG/Q) and Theta (KNL) supercomputers for a heterogeneous linear elastic model problem in two dimensions with 14k d.o.f. per subdomain; see [60, Fig. 7] for details

6.2 Nonlinear BDDC Methods

Using the same elimination set as in Nonlinear-FETI-DP-4, the nonlinear BDDC method [51] can be derived, which is based on its linear version; see [20, 24, 71–73]. We presented an efficient and scalable implementation of linear and nonlinear BDDC avoiding the computation of Schur-complements in [60]. This approach proved to be faster, more scalable, and more robust for nonlinear hyperelasticity problems (see Fig. 18 (left)) as well as for elasto-plasticity problems using realistic RVE microstructures obtained from EBSD measurements; see [60, Table 4.7]. We also studied the scalability of the embedded linear BDDC solver on different architectures; see Fig. 18 (right).

7 Multicore Performance Engineering of Sparse Triangular Solves in PARDISO Using the Roofline Performance Model

The PARDISO [12, 23, 69, 96] parallel sparse direct solver is a building block in FE2TI. As long as the macroscopic problem is small enough, it can be solved directly by PARDISO; if the microscopic problems are of reasonable size, it is efficient to use PARDISO concurrently on the microscale problems. For large micro and macro problems, the direct solver has to be replaced by linear or nonlinear FETI-DP or BDDC domain decomposition solvers. Here, PARDISO is typically applied as the subdomain and coarse solver.

The PARDISO solver has two phases: factorization and forward/backward substitution, with factorization being more time consuming than substitution. However, the former is only performed once in a FETI-DP or BDDC domain decomposition iterative process, whereas the latter is repeated many times. We are in particular interested in the forward and backward solution process of sparse direct solvers since they build the computational kernel, e.g., in FETI-DP or BDDC methods. FETI-DP and BDDC are known to be highly parallelizable, but most implementations are using sparse direct solvers as building blocks. More precisely, most domain decomposition implementations use sparse direct solves for the local subdomain problems to obtain the necessary robustness. Additionally, the coarse problem is usually solved directly up to a certain size, but for larger problems the coarse solve is often approximated by, e.g., AMG or recursive applications of the domain decomposition approach itself. Therefore, we investigate and analyze the performance of the forward/backward solution process of PARDISO for the local subdomain solves in FETI-DP and BDDC and present not only numerical results, but also a detailed performance analysis for a representative sparse solver kernel based on the roofline model. The goal is to create an analysis of this part of the algorithm and to establish a roofline performance model [98], which considers performance bounds given by the memory bandwidth and the processor peak performance. We modeled both the serial and parallel execution phases. Despite the fact that the roofline model prediction can be inaccurate in the serial case, when the in-core execution or in-cache transfers become dominant, it still provides an easily obtainable upper bound. The simple roofline model brings together the computational capabilities of the processor and the available memory bandwidth with the requirements of an algorithm. In our case the relevant quantities are the number of FLOPs performed and the data transferred between processor and memory, which we determined by an analysis of the forward/backward substitution.

The performance of the forward and backward substitution process is analyzed and benchmarked for a representative set of sparse matrices on modern x86-type multicore architectures. The characteristic quantities of these systems are shown in Table 5 along with the required machine specific input parameters (lower part of Table 5) for the roofline model. The measurement approach, its validation, as well as limitations are discussed in [98]. Our modeling approach covered both the serial and parallel execution phases, allowing for in-socket performance predictions. The performance modeling results for a discretized Laplacian matrix ('lapl2') and a local subdomain matrix ('BDDC') from the EXASTEEL project are shown in Fig. 19; see also Table 6 for dimensions and numbers of nonzeros for the considered matrices. The latter matrix is used in FETI-DP as well as BDDC methods to compute the discrete harmonic extension from the domain decomposition interface to the interior of a certain subdomain. The specific problem represents a typical RVE using the J_2 elasto-plasticity model including the material parameters also used in the computations of the FLD. We verified that the considered subdomain already showed a plastic behavior.

As opposed to the original roofline model, the modified roofline model covers also the performance impact of limited scalability imposed by the algorithm, i.e.,

Table 5 Details of the Intel and AMD hardware systems evaluated

Name	IVB	BDW	SKX	ZEN-S
Processor name	Intel Xeon	Intel Xeon	Intel Xeon	AMD EPYC
	E5-2660 v2	E5-2630 v4	Gold 6148	745
Micro architecture	Ivy Bridge	Broadwell	Skylake	Zen
Frequency [GHz]	2.2	2.2	2.4	2.3
Number of cores	10	10	20	24
Vector instruction set	AVX	AVX2	AVX-512	AVX2
NUMA LDs	1	1	1	4
<i>Scalar read bandwidth</i>				
1 core [GB/s]	9.5	11.5	14.5	19.3
NUMA LD [GB/s]	44.4	56.3	108.0	37.6
<i>Scalar ADD+MUL/FMA</i>				
1 core [F/cy]	2	4	4	4
NUMA LD [F/cy]	20	40	80	24
<i>Scalar machine balance B_m</i>				
1 core [B/F]	2.2	1.3	1.5	2.1
NUMA LD [B/F]	1.0	0.6	0.6	0.7

The NUMA Locality Domain (LD) refers to a number of cores sharing a physical or logical memory controller

Fig. 19 Performance on IVB, BDW, SKX, and ZEN-S (left to right) for two matrices (upper and lower panel) from EXASTEEL. (a–d) lapl2. (e–h) BDDC. Copyright 2019, IEEE, [98]

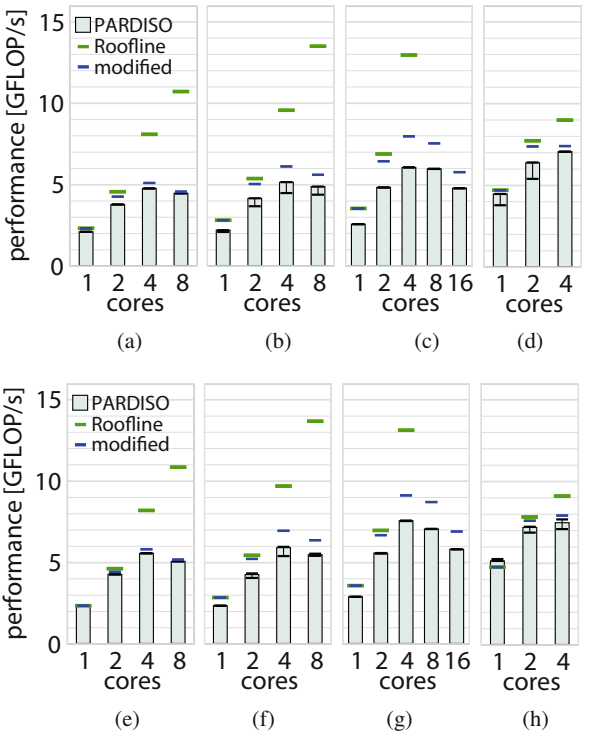


Table 6 Dimension (n) and number of nonzeros (nnz) for A and L for benchmark matrices

Matrix	n	nnz(A)	nnz(L)
lap12	343×10^3	1×10^6	166×10^6
BDDC	750×10^3	31×10^6	1590×10^6

both serial and parallel execution phases of the forward and backward substitution are considered in the model; see [98] for details. It captures the behavior of the measured performance quite well compared to the original roofline model.

For Intel Ivy Bridge systems the modified roofline model error is only up to 5%. Further details are given in [98]. During the second phase of EXASTEEL the close collaboration with ESSEX-II in the context of performance engineering was also extended to iterative solvers, leading to a new promising recursive algebraic coloring scheme [1]. The benefits of the recursive algebraic coloring were demonstrated by applying it to the kernel operation of a symmetric sparse matrix-vector multiplication (SymmSpMV) on various multicore architectures. The performance was compared against standard multicoloring and various other algebraic block multicoloring methods. The coloring implementation resulted in an average and maximum speedup of 1.4 and 2, respectively. Our entire experimental and performance analysis process was also backed by the roofline performance model, corroborating the optimality of the approach in terms of resource utilization of the SymmSpMV kernel on modern hardware; see the ESSEX-II report in this book for details.

8 Improvement of the Mechanical Model for Forming Simulations

In this section, we describe improvements to the modeling developed in the project. Not all of the techniques described here are currently used in our FE2TI production simulations, mainly to reduce computational cost.

As mentioned earlier, the favorable macroscopic properties are to a large extent due to the heterogeneous microstructural features of the DP steels. A sophisticated heat treatment process is used to produce a heterogeneous microstructure with ferritic matrix and martensitic inclusions. This process is also accompanied by several effects which then in conjunction with the difference in mechanical properties of ferrite (soft phase) and martensite (hard phase) generate advantageous macroscopic properties. In this project, an initial volumetric strains approach, c.f. [14], was developed to mimic the micromechanical features resulting from the phase transformations during the production process.

The high yield and work-hardening properties of DP steels, on the other hand, pose a problem in terms of forming complex geometries and designing the metal working tools. One of the major issues associated with the forming of DP steels, is the large springback observed at the end of the forming process, which results in

undesirable geometries of the formed parts. Here, simulating the forming process with an accurate material behavior can help to predict springback precisely and further save valuable resources while optimizing the tooling parameters for the process. The springback behavior is found to be very closely associated to the Bauschinger factor of the material. Therefore, within this project a multiscale modeling strategy to effectively model the DP steel response under cyclic loading was developed. In this context, an efficient neural network based algorithm is employed to identify the associated microstructural material parameters, leading to a reduction in the required computational resources. In order to obtain further understanding of the correlation of the model parameters on the macroscopic behavior of DP steels during cyclic loading, uncertainty quantification studies have been carried out using the developed mechanical models.

Due to their higher accuracy and physical interpretability, crystal plasticity formulations may be used at the RVE level to directly describe plasticity in a polycrystal such as multiphase steels. Due to the fact that such formulations are computationally highly expensive, they may be primarily applied to computationally identify macroscopic yield surfaces. FE^2 simulations of metal forming processes based on such formulations at the small scale will however hardly be feasible. Therefore, as part of this project one goal was to improve the quality of micromechanical models to be used efficiently in the context of FE^2 . The associated micromechanical simulations are mainly making use of a classical finite J_2 elasto-plasticity material model, c.f. [79, 87, 89, 90], which is used to model the micro-constituents (ferrite and martensite) by defining the hardening law

$$\beta^{\text{iso}} = y_{\infty}^{\text{iso}} + (y_0^{\text{iso}} - y_{\infty}^{\text{iso}})\exp(-\eta^{\text{iso}}\alpha) + h^{\text{iso}}\alpha. \quad (5)$$

Herein, y_0^{iso} is the initial yield stress, y_{∞}^{iso} is the saturation yield stress, η^{iso} is the exponent, h^{iso} is the linear hardening at saturation yield stress and α is the equivalent plastic strain variable. The material parameters of the models are calibrated based on uniaxial tests performed on the pure individual constituents martensite and ferrite. As representative microstructure a so-called statistically similar RVE was identified; see [14]. More information on SSRVEs can be obtained in [7, 8], and [83]. Although the individual phases can be represented accurately and the microstructure is reflected with high accuracy, still the experimental stress-strain response cannot be obtained. The main reason is that distributed properties in the ferritic matrix phase, which however result from the production process, were not yet taken into account. In addition to that, when focusing on cyclic loading protocols, the macroscopic kinematic hardening of the real sheet metal cannot be represented. As a suitable quantitative measure for the kinematic hardening, the so-called Bauschinger factor can be computed as $\bar{f}_B = (|\bar{P}_I| - \bar{P}_{II})/|\bar{P}_I|$. Herein, $\bar{P}_I = \bar{P}_x(\Delta\bar{l}_x/\bar{l}_{x,0} = -0.05)$ and $\bar{P}_{II} = \bar{P}_x(\Delta\bar{l}_x/\bar{l}_{x,0}(\bar{P}_x = 0) + 0.002)$, where \bar{P} is the 1st Piola-Kirchhoff stress tensor. Although the Bauschinger factor of the FE^2 simulation with $\bar{f}_B^{\text{comp}} = 0.47$ is interestingly high considering that for the individual phases no kinematic hardening is taken into account up to here, it does not agree well with the experimental value

$\overline{f}_B^{\text{exp}} = 0.66$. Therefore, in this project three major improvements were developed to enhance the model quality for FE² simulations of sheet metal forming problems: (i) A mixed isotropic-kinematic hardening model was implemented at the microscale for the ferrite phase, (ii) an initial volumetric strains approach was developed to model the locally distributed plastic properties in the ferrite phase, and (iii) an implicit fit procedure was constructed based on a neural network to identify the kinematic hardening parameters.

A mixed hardening model was implemented for the ferritic phase, which consists of an exponential isotropic hardening law, see Eq. (5) and a linear kinematic hardening law, c.f. [89]. The yield criterion and the evolution of the back stress ξ are then given by

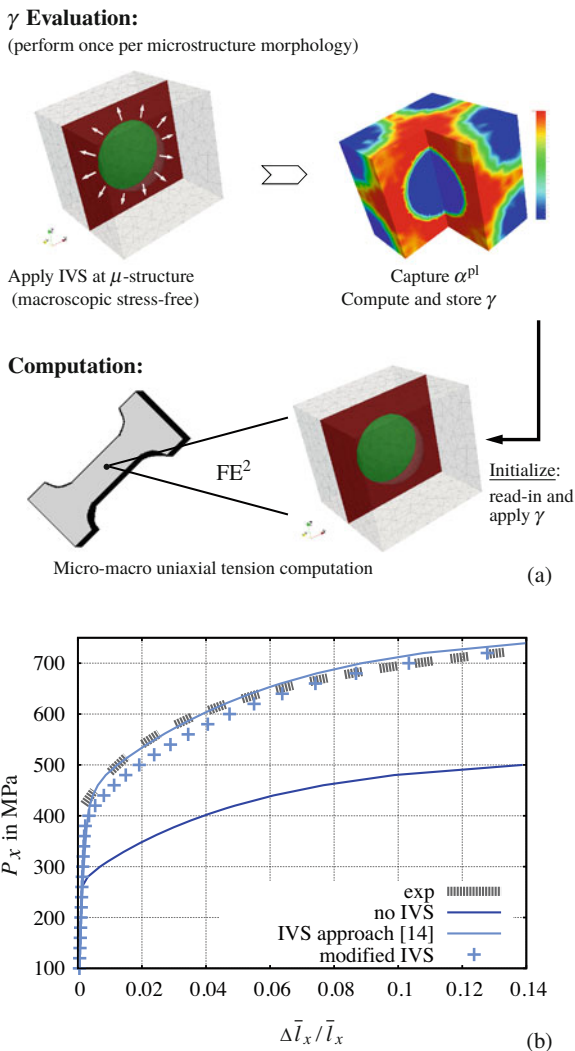
$$\phi^{\text{mix}} = ||\text{dev}\sigma - \xi|| - \sqrt{\frac{2}{3}}\beta^{\text{iso}} \quad \text{and} \quad \dot{\xi} = \frac{2}{3}\dot{\lambda}H^{\text{kin}} \frac{\text{dev}\sigma - \xi}{||\text{dev}\sigma - \xi||}. \quad (6)$$

Here, H^{kin} is an additional material parameter, signifying the linear kinematic hardening. Thus, the material parameters associated with the ferritic phase should be newly identified for the mixed hardening material model. An appropriate multiscale approach has been developed which is described in Sect. 8.2.

8.1 Initial Volumetric Strains (IVS) Approach

The IVS approach proposed in [14] allows the modeling of heterogeneous yield stress distribution in ferrite and results in a good agreement of the predicted stress values with the experiments. Here, the ferritic yield curve is locally modified using modification factor $\gamma(X) \in [1, 1.6]$, quantified based on physical and experimental observations. As a result of this continuous procedure, where the microstructure is subjected to first IVS, followed by subsequent mechanical loading, for, e.g., uniaxial tension, not only the distributed properties are obtained but also the eigenstresses related with the volume expansion of the inclusion phase can be modeled. However, in the context of FE² simulations this procedure is rather expensive since the application of the volumetric strains has to be simulated at each point before the actual loading can be applied. Due to the fact that the above-mentioned eigenstresses are not significant to the macroscopic stresses under loading, here a separated approach is proposed: the first step of applying IVS is considered only to generate the local ferritic yield modification factors which are saved independently of any potential subsequent loading. Then, in the second step of mechanical loading these modification factors are applied to the undeformed microstructure. The main benefit is the reduction of computing time since the IVS has to be performed only once to one microstructure. On the other hand, the eigenstresses obtained from the volume expansion are not included anymore. Note that these eigenstresses are usually removed from the DP steel sheet by a special heat treatment procedure which is why

Fig. 20 (a) Illustrations of the steps involved in the (modified) initial volumetric strains approach, (b) comparison of macroscopic stress-strain curves for FE² uniaxial tension calculations for simplified microstructure with spherical inclusion: (i) modified IVS approach, (ii) IVS approach [14], (iii) no IVS approach and (iv) experiment



the absence of these eigenstresses in the numerical simulation may even be more realistic. The scheme is illustrated in Fig. 20a. Furthermore, the macroscopic stress strain curves obtained under uniaxial tension for various IVS considerations are compared against the experiment in Fig. 20b. Here, it can be seen that the proposed modified (separated) IVS scheme performs equivalently to the continuous IVS as given in [14]. However, as seen in Fig. 20b, not using the IVS approach yields a poor accuracy in representing the experimental curve.

Additionally, it is observed that the choice and extent of ferritic hardening has no effect on the resulting factors γ and that the same set can be applied in mechanical loading computations as long as (i) the microstructure, (ii) the amount of martensitic

volume jump considered (i.e. 4%) and (3) the initial yield stress of ferrite (y_0^{iso}), remain unchanged.

The IVS approach has been implemented in FE2TI but has not been applied in the production runs in Sect. 5.

8.2 Parameter Identification Approach for Ferritic Mixed Hardening

The incorporation of the mixed hardening in ferrite along with the initial volumetric strains approach necessitates the identification of a new set of material parameters for ferrite, i.e. y_∞^{iso} , η^{iso} and H^{kin} . Here, y_0^{iso} and h^{iso} are assumed to be constant. Since no cyclic stress-strain data is available for the pure ferrite, the ferrite properties need to be adjusted such that the macroscopic response matches well the experiment. Due to the micro-macro nature of the computations required for the resulting inverse problem, this parameter identification problem becomes highly time and computation intensive. Therefore, to accelerate the process a neural network based algorithm is proposed.

As illustrated in Fig. 21, a sufficiently trained neural network takes in eight input values from DP steel experiments and outputs the values for the three parameters to be identified. These input values are as illustrated in Fig. 21b, i.e. seven macroscopic stress values and the Bauschinger factor. The neural network consists of one hidden

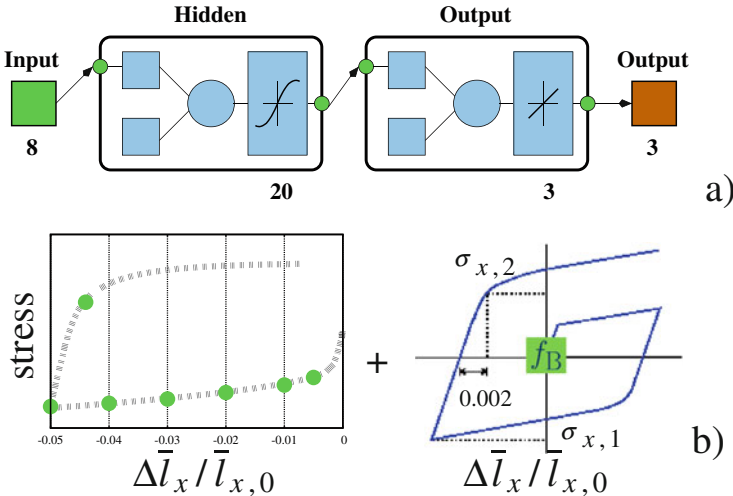


Fig. 21 (a) Schematic representation of the neural network with the respective activation functions at various layers, (b) the input values for the neural network—7 macroscopic stress-strain values and 1 macroscopic Bauschinger factor

Table 7 Neural network training data range for the identification of ferritic material parameters

	H^{kin}	η^{iso}	y_{∞}^{iso}
Range	0.9–2.1 GPa	15.0–25.0	0.35–0.80 GPa
Nos.	5	4	4

Table 8 (a) Target material parameters for the ferritic phase identified with the trained neural network and (b) Bauschinger factor computed with mixed hardening in ferrite (Sim-mix)

(a)	H^{kin}	η^{iso}	y_{∞}^{iso}	(b)		Exp	Sim-mix
	1.71 GPa	25.6	0.375 GPa		f_B	0.66	0.62




Fig. 22 Stress response comparison between simulation model with mixed hardening for ferrite and the experiment

layer with *tanh* type activation functions and one output layer with linear activation functions. Results from 80 simulations with the choice of target parameters in the ranges mentioned in Table 7 using the simplified microstructure (spherical inclusion in a cuboidal matrix) are used as training data. These simulations with different target parameter combinations can all be executed at once on many core machines to accelerate the process of gathering training data.

Additionally, a good choice of training range helps to ensure a robust prediction of the target parameters. The parameters identified by evaluating this algorithm are given in Table 8a. The macroscopic stresses achieved during compression and the overall Bauschinger factor obtained with these parameters are in Fig. 22 and Table 8b, respectively. These indicate a good match with the experimental observations. Additionally, it is found that the identified material parameters predict a higher pure ferritic yield curve than observed in experiments on the lab-synthesized pure ferrite. It is emphasized that it is in principle challenging to synthesize a pure ferrite which corresponds to the ferrite in the DP steel with respect to similarity in chemical composition and grain size distribution. Therefore, the experimental data regarding the pure ferritic behavior should generally be only considered carefully.

8.3 Quantification of Uncertain Material Properties

The material parameters for the micro-constituents of the DP steel are usually obtained based on experiments on limited numbers of samples. Since the material

behavior of the constituents depends on the production process parameters, which may be non-uniform due to the nature of the process over large batch sizes, the measurements might not accurately represent the complete reality. This holds in particular for specialized laboratory productions of samples only consisting of the pure ferrite phase, which matches the microstructure and chemical composition as closely as possible compared to the ferrite in the DP steel. Due to the fact that these ferrite properties are however believed to strongly influence the properties of the overall DP steel behavior an associated uncertainty quantification analysis was conducted as part of the project. Based on such analysis the sensitivity of the macroscopic stress-strain response of DP steel for modified ferritic properties can be investigated. For the analysis employed here, known probability distributions are assumed for selected ferritic parameters which are (i) y_0^{iso} and y_∞^{iso} and (ii) H^{kin} . It should be noted that varying y_0^{iso} and y_∞^{iso} together for ferrite leads to a change in the height of the ferritic yield curve. For each of the cases, 15,000 samples are randomly constructed to generate Gaussian distributions as input uncertainty distributions for the ferrite parameters y_0^{iso} and H^{kin} ; see Fig. 23.

Based on these assumed input distributions of the ferrite parameters the resulting distributions of macroscopic properties are to be computed. Trained neural networks are used here again to evaluate each of these samples and to compute the macroscopic DP steel responses for (i) the yield stress in compression ($\bar{R}_{p0.25}$), (ii) the Bauschinger factor (\bar{f}_B) and (iii) the hardening modulus around 5% compression (\bar{H}_{end}).

The output uncertainties in the above mentioned macroscopic measures are then plotted in the sense of their co-relation with the ferritic yield curve and the prescribed kinematic hardening parameter in Figs. 24 and 25 respectively. The correlation between the output macroscopic initial yield stress ($\bar{R}_{p0.25}$) and the ferritic yield curve as seen in Fig. 24a for the prescribed input, turns out to be a linear relationship. As evident in Fig. 24b the macroscopic Bauschinger factor, changes non-linearly with the ferritic yield curve. However, the overall small variations in the values of the Bauschinger factor suggest that the height of the ferritic yield curve




Fig. 23 Input uncertainty distributions for (a) variation in y_0^{iso} and (b) H^{kin} ferritic material parameters




Fig. 24 (a) Output uncertainty of (a) $\bar{R}_{p0.25}$ stress, (b) \bar{f}_B and (c) \bar{H}_{end} based on the variation of y_0^{iso}




Fig. 25 (a) Output uncertainty of (a) $\bar{R}_{p0.25}$ stress, (b) \bar{f}_B and (c) \bar{H}_{end} based on the variation of H^{kin}

only negligibly influences the macroscopic Bauschinger factor. Thus, it appears that it is the overall large difference in yield stress between the ferrite and the martensite rather than moderate changes within the ferrite itself, which is responsible for the relatively large kinematic hardening of DP steel. Fig. 24c indicates a linear relationship of \bar{H}_{end} with the ferritic yield curve height. Again, the values of macroscopic moduli change only insignificantly, which indicates a small sensitivity of macroscopic response due to modifications of the ferrite yield stress.

Now, the influence of the kinematic hardening is investigated. The variation in the above mentioned macroscopic quantities has been considered for a prescribed uncertainty in the linear kinematic hardening modulus H^{kin} of the pure ferrite phase. The results for the macroscopic initial yield stress $\bar{R}_{p0.25}$ are plotted in Fig. 25a, where a linearly reducing correlation is observed with increasing ferritic H^{kin} . As before, the values indicate a negligible change in $\bar{R}_{p0.25}$ with modifications of H^{kin} . This changes significantly for the macroscopic Bauschinger factor, see the results in Fig. 25b, which indicate a strong influence of the macroscopic Bauschinger factor due to modifications of the ferritic kinematic hardening modulus. The relationship is a linearly increasing one. Likewise, an also rather significant, linearly increasing correlation is observed between \bar{H}_{end} and H^{kin} , which has been plotted in Fig. 25c. The results indicate that while the variation in the height of the ferritic yield curve results in a considerable effect on the macroscopic initial yield stress and hardening modulus, it only negligibly influences the Bauschinger factor. Whereas, the ferritic kinematic hardening has a strong influence on the macroscopic response of the DP

steel, especially the macroscopic Bauschinger factor. This further highlights the necessity of employing a mixed hardening based ferritic material model for the micro-macro simulation of relevant DP steel forming problems where effects such as spring-back are of major importance.

Especially for uncertainty quantification problems, where the variation in the microstructure's morphology is considered as source of uncertainty, a high number of different statistically similar volume elements (SSVEs) needs to be simulated. For this purpose, an optimal decomposition approach in the context of a finite cell integration scheme was developed in this project, see [26]. This approach allows for an automated calculation without the necessity to construct a new mesh for each SSVE while keeping the overall computing time even lower than for a conforming (standard) mesh.

8.4 Crystal Plasticity

A better description of certain phenomena, e.g., localization, in crystalline materials can be achieved by explicitly modeling the polycrystalline structure of the material.

Such materials consist of various single crystals with different orientations which interact through the granular interfaces. By directly modeling the plastic behavior of these single crystals, anisotropic yield and complex flow behavior can be captured directly. As pointed out in Sect. 8, this would lead to computationally highly expensive simulations, which can be overcome using approximations of the response of the underlying polycrystal. However, to illustrate the procedure and complexity of incorporating polycrystalline microstructures directly into multiscale simulations, a single crystal plasticity model for face-centered cubic (fcc) crystals at small strains has been implemented, considering an additive decomposition of the small strain tensor into elastic and plastic part $\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}^e + \boldsymbol{\varepsilon}^p$ where $\dot{\boldsymbol{\varepsilon}}^p = \sum_{\delta} \dot{\gamma}^{\delta} \boldsymbol{P}^{\delta}$ directly connects the inelastic behavior in the individual grains to the inherent crystallographic structure through the dependency of the rate of plastic strain on the projected rate of plastic slip $\dot{\gamma}^{\delta}$ summed over all systems δ . Therein, the projection tensor $\boldsymbol{P}^{\delta} = \text{sym}(\boldsymbol{s}^{\delta} \otimes \boldsymbol{n}^{\delta})$ is defined based on the orthonormal vectors $\boldsymbol{s}^{\delta} \perp \boldsymbol{n}^{\delta}$, describing the slip system δ . Single crystal plasticity models can be distinguished into rate-independent and rate-dependent models. Algorithms of the former type are typically governed by the issue of non-uniqueness of choice of active slip systems among all possible ones, [3, 16], which adds to the complexity of the material model. Different approaches exist to handle this intrinsic problem by e.g. simple perturbation techniques [74], augmented Lagrangian methods [85] or penalty approaches. Recently, an alternative approach to handle the issue of non-uniqueness among the activity of slip systems has been proposed in [84], which uses Infeasible Primal-Dual Interior Point methods for solving the constraint optimization problem. This method uses barrier functions combined with the given constraints of the problem in order to penalize the approach of the boundary of the feasible domain. In contrast to that, rate-dependent algorithms consider all slip systems to be active and

link the rate of slip $\dot{\gamma}^\delta$ on each system δ directly to the Schmid stress $\tau^\delta = \boldsymbol{\sigma} : \mathbf{P}^\delta$. The kinetic law reads

$$\dot{\gamma}^\delta = \dot{\gamma}_0 \left| \frac{\tau^\delta}{g^\delta} \right|^{p-1} \left(\frac{\tau^\delta}{g^\delta} \right) \quad \text{with} \quad \dot{g}^\delta = \sum_{\beta} h^{\delta\beta} |\dot{\gamma}^\beta|, \quad (7)$$

as, e.g., proposed by [39], where the hardening moduli $h^{\delta\beta}$ depends on the strain-like internal variable A with $\dot{A} = \sum_{\delta} |\dot{\gamma}^\delta|$.

8.5 Macroscopic Yield Surface Based on Polycrystalline RVEs

In this section, we use two-scale simulations using crystal plasticity to compute macroscopic yield surfaces. These yield surfaces can then be used in FE2TI simulations without directly incorporating crystal plasticity.

The influence of the microscopic polycrystalline material can be considered to compute the resulting macroscopic anisotropic yield surfaces, as mentioned in Sect. 8 and included in a hierarchical multiscale approach, see [32]. In the following, a microstructure consisting of a polycrystal with multiple grains is considered to model its macroscopic yield behavior. Here, for the computation of macroscopic yield surfaces based on the microscopic behavior of polycrystalline unit cells, the software Neper is used to generate a periodic unit cell with 15 grains. The geometry is meshed using 10-noded tetrahedral finite elements. In order to account for an isotropic orientation distribution of the polycrystalline unit cell, each grain is assigned to a specific orientation following from a geodesic dome. For details, we refer to [82]. With these unit cells, macroscopic yield curves based on macroscopic biaxial loading paths, i.e. $\bar{\sigma}_1 : \bar{\sigma}_2, \bar{\sigma}_3 = 0$, are computed in an FE² scheme. The stress-driven simulation requires small time steps, which is amplified by the small time step size required for the rate-dependent formulation of single crystal plasticity. Figure 26 shows the initial yield surface at $\langle \alpha \rangle = 3.3 \cdot 10^{-8}$ as well as the distribution of α inside the unit cell. Since the rate-dependent formulation does not have a distinct yield point and the rate-independent behavior is here modeled with $p = 200$, see Eq. (7), this value of equivalent plastic strains has been arbitrarily chosen by the authors. The evolved macroscopic yield surface based on a polycrystalline unit cell at $\langle \alpha \rangle = 4.7 \cdot 10^{-4}$ and a respective distribution of α is shown in Fig. 26. As pointed out in [11], the initial yield surface forms the shape of a Tresca-type yield criterion, whereas the further evolved yield surface is of typical elliptical Mises-type.




Fig. 26 Polycrystal with 15 grains. Distribution of equivalent plastic strains α for loading path $\bar{\sigma}_1 : \bar{\sigma}_2 = -1 : 0$ at $\alpha_1^* = \langle \alpha \rangle = 3.3 \cdot 10^{-8}$ (left) and for loading path $\bar{\sigma}_1 : \bar{\sigma}_2 = 0 : 1$ at $\alpha_2^* = \langle \alpha \rangle = 4.7 \cdot 10^{-4}$ (center). Right: Associated successive yield surfaces; Tresca-type for α_1^* and von Mises type for α_2^*

8.6 One-Way Coupled Simulation of Deep-Drawing Using Polycrystalline Unit Cells

Finally, in this section, we demonstrate a two-scale simulation directly incorporating crystal plasticity on the micro scale. Such simulations are computationally expensive. Only a one-way coupling is used here, and J_2 elasto-plasticity is applied on the macroscale.

In the following, a sheet metal forming process of deep-drawing of a hat-profile, adopted from [7], using an Al-Cu alloy is simulated under consideration of the polycrystalline microstructure in a one-way coupled FE scheme. In Fig. 27, the finite element mesh (165 linear quadrilateral elements) is shown. The interaction between the sheet and the tools is realized with a frictionless penalty contact formulation. The macroscale simulation is carried out using a finite J_2 elasto-plasticity model with isotropic von Mises yield behavior based on an algorithmic setting by [65]. The material parameters, cf. Eq. (5), were fitted to a macroscopic uniaxial tension test with the polycrystalline unit cell used on the microscale leading to $\bar{\kappa}^{\text{Al-Cu}} = 50,754 \text{ N/mm}^2$, $\bar{\mu}^{\text{Al-Cu}} = 23,425 \text{ N/mm}^2$, $\bar{y}_0^{\text{Al-Cu}} = 125 \text{ N/mm}^2$, $\bar{y}_\infty^{\text{Al-Cu}} = 160 \text{ N/mm}^2$, $\bar{\eta}^{\text{Al-Cu}} = 750$ and $\bar{h}^{\text{Al-Cu}} = 1 \text{ N/mm}^2$. The final state of the sheet forming simulation is depicted in Fig. 28 and the distribution of equivalent plastic strain is shown. Throughout the simulation, the deformation gradient $\bar{\mathbf{F}}$ is captured at three different positions, marked by \square , \triangle and \circ therein, at the top, center and bottom of the sheet, respectively, leading to nine evaluation points in total.

The recorded deformation is applied to a polycrystalline unit cell in a one-way FE coupling. The single crystal plasticity computation is performed at small strains, as described in Sect. 8.4. Thereby, the applied material parameters are taken from [92] with Lamé constant $\lambda = 35,104.88 \text{ N/mm}^2$, shear modulus $\mu = 23,427.25 \text{ N/mm}^2$,




Fig. 27 Discretization of the macroscopic BVP of the deep-drawing of a hat profile under plane strain conditions. 1: drawing die, 2: blank holder, 3: sheet metal (discretized with 5×33 elements), 4: punch, 5: punch radius of 7 mm, and 6: die radius of 6 mm are used. The contact definitions between punch, drawing die, blank holder, and sheet metal are realized using a frictionless penalty formulation. The analyzed RVEs in the macroscopic BVP are located near the punch radius, ⑦, in the vertical section, ⑧, and near the die radius, ⑨, according to the final deformation, cf. 28 (left). The drawing depth of the hat profile is 45.7 mm with a sheet half width of 100 mm and a thickness of 1.4 mm




Fig. 28 Distribution of $\bar{\alpha}$ in sheet metal at final deformed state ($t = 50$) and strain-like internal variable A in polycrystalline unit cells at different points

initial slip resistance $\tau_0 = 60.84 \text{ N/mm}^2$, saturation stress $\tau_\infty = 109.51 \text{ N/mm}^2$, initial hardening modulus $h_0 = 541.48 \text{ N/mm}^2$, material rate sensitivity parameter $p = 200$, and reference slip rate $\dot{\gamma}_0 = 1 \cdot 10^{-3}$. The small strain tensor $\bar{\epsilon}$ is used to transfer the deformation state from the macroscale to the microscale, however, no coupling back from micro- to macroscale is considered.

In Fig. 28, the distribution of equivalent plastic strain $\bar{\alpha}$ in the hat-profile and the distribution of the strain-like internal variable A as a result of the evaluation of the one-way coupled polycrystalline unit cells is shown. Differences between the positions of the polycrystals in the sheet are obvious as well as the nonhomogeneous distribution of A .

9 Conclusion

The vision of the EXASTEEL project is to develop a virtual HPC laboratory allowing for predictive virtual material testing of modern steels. On this path, we have moved forward in several directions: Since the properties of modern dual-phase steels largely stem from their microstructure, homogenization is indispensable to achieve our goals. We therefore have developed and implemented the FE2TI library, a highly scalable software for computational homogenization based on the FE² approach (Sects. 2–4). This approach was then used, for the first time, to compute a forming limit diagram for DP600 steel using the JUWELS supercomputer (Sect. 5). Let us remark that the computation of an FLD is already a step beyond the achievements envisaged in the original EXASTEEL-2 proposal. We have also shown scalability of the FE2TI package up to the largest supercomputers currently available, e.g., using more than one million MPI ranks for nonlinear production problems, i.e., using unstructured meshes, elasto-plasticity, and full parallel I/O [64] (Sect. 3). These latter simulations use parallel FETI-DP solvers for the RVE problems and made use of the full JUQUEEN supercomputer.

To move towards full use of the future exascale supercomputers, we have worked on extending the parallel scalability of implicit nonlinear FETI-DP and BDDC domain decomposition solvers (Sect. 6). Scalability to 800,000 parallel tasks was achieved for our nonlinear solvers [54], outside of our parallel FE² multiscale context; see Fig. 15. These simulations used the full Mira supercomputer. We have also considered techniques to improve the energy efficiency of our nonlinear domain decomposition solvers (Sect. 6.1.1). Careful performance analysis and engineering was applied to the FE2TI software building blocks, e.g., for the performance engineering of the sparse triangular solves of the PARDISO sparse direct solver (Sect. 7).

For the modeling, considering initial volumetric strains, resulting from the complex steel production process, has shown to be of interest; therefore, an efficient algorithmic approach to IVS was proposed (Sects. 8.1 and 8.2). This IVS approach has been implemented in FE2TI. Further improvements in our modeling

may be achieved by incorporating effects from crystal plasticity (Sect. 8.4). An approach to fit macroscopic yield surfaces to crystal plasticity simulations was presented (Sect. 8.5). The resulting yield surfaces can be used in FE2TI without using an explicit coupling with crystal plasticity simulations. However, we have also demonstrated a two-scale simulation using a one-way coupling with crystal plasticity (Sect. 8.6).

For the quantitatively predictive simulations envisaged in this project, several improvements are planned for the future. First, realistic material models reproducing the physics on the microscale are important. Different advanced approaches beyond the ones considered so far may be of interest, e.g., based on the techniques described in Sect. 8. Second, for the computation of the FLD, the exploitation of the symmetry of the Nakajima specimen has to be reviewed and, especially for strongly anisotropic microstructures, simulations using the full geometry have to be performed for all specimen. Third, a validation with experiments for steels other than DP600 will be necessary. Finally, once exascale supercomputers will be available, predictive virtual steel simulations at the exascale will leverage the combined parallelism of the FE² algorithm and of the parallel nonlinear domain decomposition solvers.

Acknowledgments This work was supported by the German Research Foundation (DFG) through the Priority Programme 1648 “Software for Exascale Computing” (SPPEXA), DFG project 230723766 and by the Swiss National Science Foundation (SNSF), project 200021E-162296. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer **JUWELS** and its predecessor **JUQUEEN** at Jülich Supercomputing Centre (JSC). Computational resource of Fujitsu PRIMERGY CX600M1/CX1640M1 (**Oakforest-PACS**) was awarded by “Large-scale HPC Challenge” project, Joint Center for Advanced High Performance Computing (JCAHPC). The authors also gratefully acknowledge the computing time granted by the Center for Computational Sciences and Simulation (CCSS) of the Universität of Duisburg-Essen and provided on the supercomputer **magnitUDE** (DFG grants INST 20876/209-1 FUGG, INST 20876/243-1 FUGG) at the Zentrum für Informations- und Mediendienste (ZIM) as well as the compute resources on **Meggie** and support provided by the Erlangen Regional Computing Center (RRZE). This research used the resources **Mira** and **Theta** of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. Furthermore, the access to the computing cluster **Taurus** provided by the ZIH, Technical University of Dresden and the computing cluster **SD Cluster** provided by the Faculty of Civil and Environmental Engineering, Ruhr-University Bochum are also gratefully acknowledged.

References

1. Alappat, C., Hager, G., Schenk, O., Thies, J., Basermann, A., Bishop, A., Fehske, H., Wellein, G.: A recursive algebraic coloring technique for hardware-efficient symmetric sparse matrix-vector multiplication. *ACM Trans. Parallel Comput.* (2020, accepted). arXiv e-prints, ArXiv:1907.06487
2. Amestoy, P., Duff, I., L’Excellent, J.Y., Koster, J.: A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.* **23**(1), 15–41 (2002)

3. Anand, L., Kothari, M.: A computational procedure for rate-independent crystal plasticity. *J. Mech. Phys. Solids* **44**, 525–558 (1996)
4. Announcement of the DFG Priority Programme ‘Software for Exascale Computing’ (SPP 1648) (2011). https://www.dfg.de/foerderung/info_wissenschaft/2011/info_wissenschaft_11_59/index.html
5. Baker, A., Klawonn, A., Kolev, T., Lanser, M., Rheinbach, O., Yang, U.: Scalability of classical algebraic multigrid for elasticity to half a million parallel tasks. *Lect. Notes Comput. Sci. Eng.* **113**, 113–140 (2016)
6. Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W., Karpeyev, D., Kaushik, D., Knepley, M., May, D., McInnes, L.C., Mills, R.T., Munson, T., Rupp, K., Sanan, P., Smith, B., Zampini, S., Zhang, H.: PETSc Users Manual. Technical Report ANL-95/11 - Revision 3.11, Argonne National Laboratory (2019). <https://www.mcs.anl.gov/petsc>
7. Balzani, D., Brands, D., Schröder, J.: Construction of statistically similar representative volume elements. In: Schröder, J., Hackl, K. (eds.) *Plasticity and Beyond - Microstructures, Crystal-Plasticity and Phase Transitions (CISM Lecture Notes 550)*, pp. 355–412. Springer, Berlin (2014)
8. Balzani, D., Scheunemann, L., Brands, D., Schröder, J.: Construction of two- and three-dimensional statistically similar RVEs for coupled micro-macro simulations. *Comput. Mech.* **54**, 1269–1284 (2014)
9. Balzani, D., Gandhi, A., Klawonn, A., Lanser, M., Rheinbach, O., Schröder, J.: One-way and fully-coupled FE2 methods for heterogeneous elasticity and plasticity problems: parallel scalability and an application to thermo-elastoplasticity of dual-phase steels. *Lect. Notes Comput. Sci. Eng.* **113**, 91–112 (2016). https://doi.org/10.1007/978-3-319-40528-5_5. https://www.scopus.com/inward/record.uri?eid=2-s2.0-84989948238&doi=10.1007%2f978-3-319-40528-5_5&partnerID=40&md5=5c4efedccb3dab06ef11fcc5b2a61b2e
10. Bhardwaj, M., Pierson, K., Reese, G., Walsh, T., Day, D., Alvin, K., Peery, J., Farhat, C., Lesoinne, M.: Salinas: a scalable software for high-performance structural and solid mechanics simulations. In: *Proceedings of the 2002 ACM/IEEE Conference on Supercomputing, SC '02*, pp. 1–19. IEEE Computer Society Press, Los Alamitos (2002). <http://dl.acm.org/citation.cfm?id=762761.762790>
11. Böhlke, T., Kraska, M., Bertram, A.: Simulation der einfachen Scherung einer polykristallinen Aluminiumprobe. *Tech. Mech., Sonderheft* 47–54 (1997)
12. Bollhöfer, M., Schenk, O., Janalik, R., Hamm, S., Gullapalli, K.: State-of-The-Art Sparse Direct Solvers. *Parallel Algorithms in Computational Science & Engineering - Parallelism as Enabling Technology in CSE Applications*, Birkhauser (2019). ArXiv: 1907.05309
13. Bordeu, F., Boucard, P., Gosselet, P.: Balancing domain decomposition with non-linear relocation: parallel implementation for laminates. *Civil-Comp Proc.* **90** (2009). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84894120361&partnerID=40&md5=fa707dc23f7475e5878d87c48103e443>
14. Brands, D., Balzani, D., Scheunemann, L., Schröder, J., Richter, H., Raabe, D.: Computational modeling of dual-phase steels based on representative three-dimensional microstructures obtained from EBSD data. *Arch. Appl. Mech.* 1–24 (2016). <https://doi.org/10.1007/s00419-015-1044-1>
15. Brune, P., Knepley, M., Smith, B., Tu, X.: Composing scalable nonlinear algebraic solvers. *SIAM Rev.* **57**(4), 535–565 (2015). <https://doi.org/10.1137/130936725>
16. Busso, E., Cailletaud, G.: On the selection of active slip systems in crystal plasticity. *Int. J. Plast.* **21**, 2212–2231 (2005)
17. Cai, X.C., Keyes, D.: Nonlinearly preconditioned inexact Newton algorithms. *SIAM J. Sci. Comput.* **24**(1), 183–200 (2003). <https://doi.org/10.1137/S106482750037620X>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0037248934&doi=10.1137%2fS106482750037620X&partnerID=40&md5=b222326c33fbbec11255e3de162d54d4>

18. Cai, X.C., Keyes, D., Marcinkowski, L.: Non-linear additive Schwarz preconditioners and application in computational fluid dynamics. *Int. J. Numer. Methods Fluids* **40**(12), 1463–1470 (2002). <https://doi.org/10.1002/fld.404>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0037203227&doi=10.1002%2ffld.404&partnerID=40&md5=bdbadff518af94aae714206b017180d6>
19. Cockcroft, M., Latham, D.: Ductility and the workability of metals. *J. Inst. Met.* **48**, 33–39 (1968)
20. Cros, J.M.: A preconditioner for the Schur complement domain decomposition method. In: Herrera, O.W.I., Keyes, D., Yates, R. (eds.) *Domain Decomposition Methods in Science and Engineering*, pp. 373–380. National Autonomous University of Mexico (UNAM), Mexico City (2003). ISBN 970-32-0859-2. *Proceedings of the 14th International Conference on Domain Decomposition Methods in Science and Engineering*. <http://www.ddm.org/DD14>
21. Daly, J.: A model for predicting the optimum checkpoint interval for restart dumps. In: Sloot, P.M.A., Abramson, D., Bogdanov, A.V., Gorbachev, Y.E., Dongarra, J.J., Zomaya, A.Y. (eds) *Computational Science—ICCS 2003 (ICCS 2003)*. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 2660, pp. 3–12. Springer, Berlin (2003). https://doi.org/10.1007/3-540-44864-0_1
22. Davis, T.: A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Trans. Math. Softw.* **30**(2), 165–195 (2004)
23. De Coninck, A., De Baets, B., Kourounis, D., Verbosio, F., Schenk, O., Maenhout, S., Fostier, J.: Needles: toward large-scale genomic prediction with marker-by-environment interaction. *Genetics* **203**(1), 543–555 (2016). <https://doi.org/10.1534/genetics.115.179887>. <http://www.genetics.org/content/203/1/543>
24. Dohrmann, C.: A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comput.* **25**(1), 246–258 (2003). <https://doi.org/10.1137/S1064827502412887>
25. Falgout, R., Jones, J., Yang, U.: The design and implementation of hypre, a library of parallel high performance preconditioners. *Lect. Notes Comput. Sci. Eng.* **51**, 267–294 (2006)
26. Fangye, Y., Miska, N., Balzani, D.: Automated simulation of voxel-based microstructures based on enhanced finite cell approach. *Arch. Appl. Mech.* 2020, accepted
27. Farhat, C., Lesoinne, M., Pierson, K.: A scalable dual-primal domain decomposition method. *Numer. Linear Algebra Appl.* **7**(7-8), 687–714 (2000). [https://doi.org/10.1002/1099-1506\(200010/12\)7:7/8<687::AID-NLA219>3.0.CO;2-S](https://doi.org/10.1002/1099-1506(200010/12)7:7/8<687::AID-NLA219>3.0.CO;2-S). *Preconditioning techniques for large sparse matrix problems in industrial applications* (Minneapolis, MN, 1999)
28. Farhat, C., Lesoinne, M., LeTallec, P., Pierson, K., Rixen, D.: FETI-DP: a dual-primal unified FETI method. I. A faster alternative to the two-level FETI method. *Int. J. Numer. Methods Eng.* **50**(7), 1523–1544 (2001). <https://doi.org/10.1002/nme.76>
29. Feyel, F.: Multiscale FE^2 elastoviscoplastic analysis of composite structures. *Comput. Mater. Sci.* **16**(1-4), 344–354 (1999). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0042142743&partnerID=40&md5=c5b0208bdde06570bc8a5898347b88f1>
30. Feyel, F.: A multilevel finite element method (FE^2) to describe the response of highly non-linear structures using generalized continua. *Comput. Methods Appl. Mech. Eng.* **192**(28-30), 3233–3244 (2003). [https://doi.org/10.1016/S0045-7825\(03\)00348-7](https://doi.org/10.1016/S0045-7825(03)00348-7). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0043127126&doi=10.1016%2fS0045-7825%2803%2900348-7&partnerID=40&md5=75de4641820478ccd41e5aa8e8d4d7f2>
31. Feyel, F., Chaboche, J.L.: FE^2 multiscale approach for modelling the elastoviscoplastic behaviour of long fibre SiC/Ti composite materials. *Comput. Methods Appl. Mech. Eng.* **183**(3-4), 309–330 (2000). [https://doi.org/10.1016/S0045-7825\(99\)00224-8](https://doi.org/10.1016/S0045-7825(99)00224-8). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0034677856&doi=10.1016%2fS0045-7825%2899%2900224-8&partnerID=40&md5=77470c678941b0c2ade865a2e31afeb>
32. Gawad, J., Van Bael, A., Eyckens, P., Samaey, G., Van Houtte, P., Roose, D.: Hierarchical multi-scale modeling of texture induced plastic anisotropy in sheet forming. *Comput. Mater. Sci.* **66**, 65–83 (2013)

33. Geers, M., Kouznetsova, V., Matouš, K., Yvonnet, J.: Homogenization Methods and Multiscale Modeling: Nonlinear Problems, pp. 1–34. American Cancer Society (2017). <https://doi.org/10.1002/9781119176817.ecm2107>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119176817.ecm2107>
34. Geuzaine, C., Remacle, J.F.: Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Methods Eng.* **79**(11), 1309–1331 (2009)
35. Groß, C.: A unifying theory for nonlinear additively and multiplicatively preconditioned globalization strategies: convergence results and examples from the field of nonlinear elastostatics and elastodynamics. Ph.D. thesis, Rheinische Friedrich-Wilhelm Universität Bonn (2009). Deutsche Nationalbibliothek <https://www.deutsche-digitale-bibliothek.de/item/PCLVYPVW5OCPUOTIKRKTMSHMFSNWEFPL>
36. Groß, C., Krause, R.: A generalized recursive trust-region approach - nonlinear multiplicatively preconditioned trust-region methods and applications. Technical report 2010-09, Institute of Computational Science, Università della Svizzera italiana (2010)
37. Groß, C., Krause, R.: On the globalization of ASPIN employing trust-region control strategies - convergence analysis and numerical examples. Technical report 2011-03, Institute of Computational Science, Università della Svizzera italiana (2011)
38. Henson, V., Yang, U.: BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Appl. Numer. Math.* **41**(1), 155–177 (2002). [https://doi.org/10.1016/S0168-9274\(01\)00115-5](https://doi.org/10.1016/S0168-9274(01)00115-5). Developments and trends in iterative methods for large systems of equations—in memoriam Rüdiger Weiss (Lausanne, 2000)
39. Hutchinson, J.: Bounds and self-consistent estimates for creep of polycrystalline materials. *Proc. R. Soc. Lond. A* **348**, 101–127 (1976)
40. Hwang, F.N., Cai, X.C.: Improving robustness and parallel scalability of Newton method through nonlinear preconditioning. *Lect. Notes Comput. Sci. Eng.* **40**, 201–208 (2005). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-33751093392&partnerID=40&md5=214fdd3b1e445839e0232863a17b3076>
41. Hwang, F.N., Cai, X.C.: A class of parallel two-level nonlinear Schwarz preconditioned inexact Newton algorithms. *Comput. Methods Appl. Mech. Eng.* **196**(8), 1603–1611 (2007). <https://doi.org/10.1016/j.cma.2006.03.019>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-33751073924&doi=10.1016%2fj.cma.2006.03.019&partnerID=40&md5=4c60a5f767eb41542ab0a553a1069bf4>
42. Hwang, F.N., Su, Y.C., Cai, X.C.: A parallel adaptive nonlinear elimination preconditioned inexact Newton method for transonic full potential equation. *Comput. Fluids* **110**, 96–107 (2015). <https://doi.org/10.1016/j.compfluid.2014.04.005>
43. Jocham, D.: Bestimmung der lokalen Einschnürung nach linearer und nichtlinearer Umformhistorie sowie Ermittlung dehnungs- und geschwindigkeitsabhängiger Materialkennwerte. Ph.D. thesis, Technische Universität München (2018)
44. Jülich Supercomputing Centre: JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre. *J. Large-Scale Res. Facil.* **1**(A1) (2015). <http://dx.doi.org/10.17815/jlsrf-1-18>
45. Jülich Supercomputing Centre: JUWELS: Modular Tier-0/1 Supercomputer at the Jülich Supercomputing Centre. *J. Large-Scale Res. Facil.* **5**(A135) (2019). <http://dx.doi.org/10.17815/jlsrf-5-171>
46. Klawonn, A., Rheinbach, O.: Inexact FETI-DP methods. *Int. J. Numer. Methods Eng.* **69**(2), 284–307 (2007). <https://doi.org/10.1002/nme.1758>
47. Klawonn, A., Rheinbach, O.: Robust FETI-DP methods for heterogeneous three dimensional elasticity problems. *Comput. Methods Appl. Mech. Eng.* **196**(8), 1400–1414 (2007). <https://doi.org/10.1016/j.cma.2006.03.023>
48. Klawonn, A., Rheinbach, O.: Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM Z. Angew. Math. Mech.* **90**(1), 5–32 (2010). <https://doi.org/10.1002/zamm.200900329>
49. Klawonn, A., Widlund, O.: Dual-primal FETI methods for linear elasticity. *Commun. Pure Appl. Math.* **59**(11), 1523–1572 (2006). <https://doi.org/10.1002/cpa.20156>

50. Klawonn, A., Widlund, O., Dryja, M.: Dual-primal FETI methods for three-dimensional elliptic problems with heterogeneous coefficients. *SIAM J. Numer. Anal.* **40**, 159–179 (2002). <https://doi.org/10.1137/S0036142901388081>
51. Klawonn, A., Lanser, M., Rheinbach, O.: Nonlinear FETI-DP and BDDC methods. *SIAM J. Sci. Comput.* **36**(2), A737–A765 (2014). <https://doi.org/10.1137/130920563>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84899626259&doi=10.1137%2f130920563&partnerID=40&md5=d92241d0098f77167704c1b2cb5aae85>
52. Klawonn, A., Lanser, M., Rheinbach, O.: FE2TI (ex_nl/fe2) EXASTEEL - Bridging scales for multiphase steels (2015). <https://www.swmath.org/software/13908>; see also the report on the JUQUEEN Extreme Scaling Workshop 2015: <http://hdl.handle.net/2128/8435>
53. Klawonn, A., Lanser, M., Rheinbach, O.: Juqueen Extreme Scaling Workshop 2015. Technical Report FZJ-JSC-IB-2015-01 (2015). Brömmel, D., Frings, W., Wylie, B.J.N. (eds.), <http://hdl.handle.net/2128/8435>
54. Klawonn, A., Lanser, M., Rheinbach, O.: Toward extremely scalable nonlinear domain decomposition methods for elliptic partial differential equations. *SIAM J. Sci. Comput.* **37**(6), C667–C696 (2015). <https://doi.org/10.1137/140997907>
55. Klawonn, A., Lanser, M., Rheinbach, O., Stengel, H., Wellein, G.: Hybrid MPI/OpenMP parallelization in FETI-DP methods. In: Mehl, M., Bischoff, M., Schäfer, M. (eds.) *Recent Trends in Computational Engineering - CE2014: Optimization, Uncertainty, Parallel Algorithms, Coupled and Complex Problems*, pp. 67–84. Springer International Publishing, Cham (2015). https://doi.org/10.1007/978-3-319-22997-3_4
56. Klawonn, A., Lanser, M., Rheinbach, O.: A highly scalable implementation of inexact nonlinear FETI-DP without sparse direct solvers. In: Karasözen, B., Manguoğlu, M., Tezer-Sezgin, M., Göktepe, S., Uğur, Ö. (eds.) *Numerical Mathematics and Advanced Applications ENUMATH 2015*, pp. 255–264. Springer International Publishing, Cham (2016)
57. Klawonn, A., Lanser, M., Rheinbach, O.: FE2TI: computational scale bridging for dual-phase steels. *Adv. Parallel Comput.* **27**, 797–806 (2016). <https://doi.org/10.3233/978-1-61499-621-7-797>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84969913450&doi=10.3233%2f978-1-61499-621-7-797&partnerID=40&md5=ae60853a9101d7bad22126f1481cced4>
58. Klawonn, A., Lanser, M., Rheinbach, O.: Juqueen Extreme Scaling Workshop 2016. Technical Report FZJ-JSC-IB-2016-01 (2016). Brömmel, D., Frings, W., Wylie, B.J.N. (eds.), <http://hdl.handle.net/2128/9990>
59. Klawonn, A., Lanser, M., Rheinbach, O., Uran, M.: Nonlinear FETI-DP and BDDC methods: a unified framework and parallel results. *SIAM J. Sci. Comput.* **39**(6), C417–C451 (2017). <https://doi.org/10.1137/16M1102495>
60. Klawonn, A., Lanser, M., Rheinbach, O.: Nonlinear BDDC methods with approximate solvers. *Electron. Trans. Numer. Anal.* **49**, 244–273 (2018)
61. Klawonn, A., Lanser, M., Rheinbach, O.: Using algebraic multigrid in inexact BDDC domain decomposition methods. In: Bjørstad, P., Brenner, S., Halpern, L., Kim, H., Kornhuber, R., Rahman, T., Widlund, O. (eds.) *Domain Decomposition Methods in Science and Engineering XXIV*, pp. 425–433. Springer, Cham (2018)
62. Klawonn, A., Lanser, M., Rheinbach, O., Uran, M.: On the accuracy of the inner Newton iteration in nonlinear domain decomposition. In: Bjørstad, P., Brenner, S., Halpern, L., Kim, H., Kornhuber, R., Rahman, T., Widlund, O. (eds.) *Domain Decomposition Methods in Science and Engineering XXIV*, pp. 435–443. Springer International Publishing, Cham (2018)
63. Klawonn, A., Lanser, M., Rheinbach, O., Wellein, G., Wittmann, M.: Energy efficiency of nonlinear domain decomposition methods. Technical Report, Universität zu Köln (2018). <https://kups.ub.uni-koeln.de/8654/>
64. Klawonn, A., Köhler, S., Lanser, M., Rheinbach, O.: Computational homogenization with million-way parallelism using domain decomposition methods. *Comput. Mech.* **65**(1), 1–22 (2020). <http://dx.doi.org/10.1007/s00466-019-01749-5>
65. Klinkel, S.: Theorie und Numerik eines Volume-Schalen-Elementes bei finiten elastischen und plastischen Verzerrungen. Ph.D. thesis, Universität zu Karlsruhe (2000)

66. Knoll, D., Keyes, D.: Jacobian-free Newton-Krylov methods: a survey of approaches and applications. *J. Comput. Phys.* **193**(2), 357–397 (2004). <https://doi.org/10.1016/j.jcp.2003.08.010>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0348198275&doi=10.1016%2fj.jcp.2003.08.010&partnerID=40&md5=3555dff5a4109f37815938eaf876d64c>
67. Konyukhov, A., Schweizerhof, K.: Contact formulation via a velocity description allowing efficiency improvements in frictionless contact analysis. *Comp. Mech.* **33**(3), 165–173 (2004)
68. Konyukhov, A., Schweizerhof, K.: On some aspects for contact with rigid surfaces: surface-to-rigid surface and curves-to-rigid surface algorithms. *Comput. Methods Appl. Mech. Eng.* **283**, 74–105 (2014)
69. Kourounis, D., Fuchs, A., Schenk, O.: Towards the next generation of multiperiod optimal power flow solvers. *IEEE Trans. Power Syst.* **PP**(99), 1–10 (2018). <https://doi.org/10.1109/TPWRS.2017.2789187>
70. Kouznetsova, V., Brekelmans, W., Baaijens, F.: Approach to micro-macro modeling of heterogeneous materials. *Comput. Mech.* **27**(1), 37–48 (2001). <https://doi.org/10.1007/s004660000212>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0035088146&doi=10.1007%2fs004660000212&partnerID=40&md5=544bd0552f2392765175af3e69542b91>
71. Li, J., Widlund, O.: FETI-DP, BDDC, and Block Cholesky methods. *Int. J. Numer. Methods Eng.* **66**(2), 250–271 (2006)
72. Mandel, J., Dohrmann, C.: Convergence of a balancing domain decomposition by constraints and energy minimization. *Numer. Linear Algebra Appl.* **10**, 639–659 (2003)
73. Mandel, J., Dohrmann, C., Tezaur, R.: An algebraic theory for primal and dual substructuring methods by constraints. *Appl. Numer. Math.* **54**(2), 167–193 (2005). <https://doi.org/10.1016/j.apnum.2004.09.022>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-19044366698&doi=10.1016%2fj.apnum.2004.09.022&partnerID=40&md5=fd842620bb07336ba64bc9dc2168d184>
74. Miehe, C., Schröder, J.: A comparative study of stress update algorithms for rate independent and rate dependent crystal plasticity. *Int. J. Numer. Methods Eng.* **50**, 273–298 (2001)
75. Miehe, C., Schröder, J., Schotte, J.: Computational homogenization analysis in finite plasticity Simulation of texture development in polycrystalline materials. *Comput. Methods Appl. Mech. Eng.* **171**(3), 387–418 (1999)
76. Nocedal, J., Wright, S.: Numerical Optimization. Springer Series in Operations Research and Financial Engineering, 2nd edn. Springer, New York (2006)
77. Norm DIN EN ISO 12004-2:2008: Metallic materials – sheet and strip – determination of forming-limit curves – part 2: determination of forming-limit curves in the laboratory (2008)
78. Pebrel, J., Rey, C., Gosselet, P.: A nonlinear dual-domain decomposition method: application to structural problems with damage. *Int. J. Multiscale Comput. Eng.* **6**(3), 251–262 (2008). <https://doi.org/10.1615/IntJMultCompEng.v6.i3.50>. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-54949155511&doi=10.1615%2fIntJMultCompEng.v6.i3.50&partnerID=40&md5=b1dd3dc619dd458bc8b7b4070690731d>
79. Prüger, S., Gandhi, A., Balzani, D.: Influence of microstructure morphology on multi-scale modeling of low-alloyed TRIP-steels. *Comput. Eng.* **35**(2), 499–528 (2017)
80. Rüde, U., Willcox, K., McInnes, L., Sterck, H.: Research and education in computational science and engineering. *SIAM Rev.* **60**(3), 707–754 (2018). <https://doi.org/10.1137/16M1096840>
81. Schenk, O., Gärtner, K.: Two-level dynamic scheduling in PARDISO: improved scalability on shared memory multiprocessing systems. *Parallel Comput.* **28**(2), 187–197 (2002)
82. Scheunemann, L.: Scale-bridging of elasto-plastic microstructures using statistically similar representative volume elements. Ph.D. thesis, Department Civil Engineering, Faculty of Engineering, University Duisburg-Essen (2017)
83. Scheunemann, L., Balzani, D., Brands, D., Schröder, J.: Design of 3D statistically similar representative volume elements based on Minkowski functionals. *Mech. Mater.* **90**, 185–201 (2015)

84. Scheunemann, L., Nigro, P., Schröder, J., Pimenta, P.: A novel algorithm for rate independent small strain crystal plasticity based on the infeasible primal-dual interior point method. *Int. J. Plast.* **124**, 1–19 (2020). <https://doi.org/10.1016/j.ijplas.2019.07.020>
85. Schmidt-Baldassari, M.: Numerical concepts for rate-independent single crystal Plasticity. *Comput. Methods Appl. Mech. Eng.* **192**, 1261–1280 (2003)
86. Schröder, J.: Homogenisierungsmethoden der nichtlinearen Kontinuumsmechanik unter Beachtung von Stabilitätsproblemen. Habilitationsschrift, Universität Stuttgart (2000). Bericht aus der Forschungsreihe des Instituts für Mechanik (Bauwesen), Lehrstuhl I
87. Schröder, J.: A numerical two-scale homogenization scheme: the FE^2 -method. In: Schröder, K.H.J. (ed.) *Plasticity and Beyond - Microstructure, Crystal Plasticity and Phase Transitions*, CISM International Centre for Mechanical Sciences, vol. 550, pp. 1–64. Springer, Berlin (2014)
88. Shahzad, F., Thies, J., Kreutzer, M., Zeiser, T., Hager, G., Wellein, G.: CRAFT: a library for easier application-level checkpoint/restart and automatic fault tolerance. *IEEE Trans. Parallel Distrib. Syst.* (2018). <https://doi.org/10.1109/tpds.2018.2866794>
89. Simo, J.: Algorithms for static and dynamic multiplicative plasticity that preserve the classical return mapping schemes of the infinitesimal theory. *Comput. Methods Appl. Mech. Eng.* **99**, 61–112 (1992)
90. Simo, J., Hughes, J.: *Computational Inelasticity*. Interdisciplinary Applied Mechanics - Mechanics and Materials, vol. 7. Springer, New York (1998)
91. Smit, R., Brekelmans, W., Meijer, H.: Prediction of the mechanical behavior of nonlinear heterogeneous systems by multi-level finite element modeling. *Comput. Methods Appl. Mech. Eng.* **155**(1-2), 181–192 (1998). [https://doi.org/10.1016/S0045-7825\(97\)00139-4](https://doi.org/10.1016/S0045-7825(97)00139-4). <https://www.scopus.com/inward/record.uri?eid=2-s2.0-0032026085&doi=10.1016%2fS0045-7825%2897%2900139-4&partnerID=40&md5=8b9c27800c66a206d2c5352ace9bd78f>
92. Steinmann, P., Stein, E.: On the numerical treatment and analysis of finite deformation ductile single crystal plasticity. *Comput. Methods Appl. Mech. Eng.* **129**, 235–254 (1996)
93. Tarigopula, V., Hopperstad, O., Langseth, M., Clausen, A., Hild, F., Lademo, O.G., Eriksson, M.: A study of large plastic deformations in dual phase steel using digital image correlation and FE analysis. *Exp. Mech.* **48**(2), 181–196 (2008). <https://doi.org/10.1007/s11340-007-9066-4>
94. Treibig, J., Hager, G., Wellein, G.: LIKWID: a lightweight performance-oriented tool suite for x86 multicore environments. In: *Proceedings of the 2010 39th International Conference on Parallel Processing Workshops, ICPPW '10*, pp. 207–216. IEEE Computer Society, Washington (2010). <http://dx.doi.org/10.1109/ICPPW.2010.38>
95. Uran, M.: High-performance computing two-scale finite element simulations of a contact problem using computational homogenization. Phd thesis, Universität zu Köln (2020)
96. Verbosio, F., Coninck, A.D., Kourounis, D., Schenk, O.: Enhancing the scalability of selected inversion factorization algorithms in genomic prediction. *J. Comput. Sci.* **22**(Supplement C), 99–108 (2017). <https://doi.org/10.1016/j.jocs.2017.08.013>
97. Volk, W., Hora, P.: New algorithm for a robust user-independent evaluation of beginning instability for the experimental FLC determination. *Int. J. Mater. Form.* **4**(3), 339–346 (2011)
98. Wittmann, M., Hager, G., Janalik, R., Lanser, M., Klawonn, A., Rheinbach, O., Schenk, O., Wellein, G.: Multicore performance engineering of sparse triangular solves using a modified roofline model. In: *2018 30th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pp. 233–241 (2018). <https://doi.org/10.1109/CAHPC.2018.8645938>
99. Wriggers, P.: *Computational Contact Mechanics*. Wiley, Chichester (2002)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

