

Achieving High Efficiency in Reduced Order Modeling for Large Scale Polycrystal Plasticity Simulations

Aslan Nasirov¹, Xiaoyu Zhang¹, David Wagner²,
Saikumar R. Yeratapally³, Caglar Oskay^{1*}

¹ Department of Civil Engineering, Vanderbilt University, Nashville, TN, 37235

² NASA Langley Research Center, Hampton, VA, 23681

³ Science and Technology Corporation, Hampton, VA, 23666

Abstract

Reduced order models for the nonlinear response of heterogeneous microstructures typically require a construction (or training) stage to build the reduced order basis. In this manuscript, an efficient model construction strategy for the eigenstrain homogenization method (EHM) is presented. The proposed strategy relies on a parallel, element-by-element, conjugate gradient solver. Near linear scaling has been achieved with respect to the number of degrees of freedom used to resolve the microstructure. Linear scaling with respect to the number of pre-analyses required to construct the reduced order model (ROM) follows from the EHM formulation. Furthermore, a parallel implementation for fast evaluation of the constructed ROM has been developed using shared memory parallelization. It has been shown that for large microstructures with $\approx 10,000$ grains, the total computational cost of evaluating the nonlinear response of a polycrystal could be reduced by approximately an order of magnitude using 32 cores with respect to serial ROM simulation. The present methodology has been verified using an additively manufactured polycrystalline microstructure of a nickel-based superalloy, Inconel 625. The capability of the developed framework to construct a ROM for such large microstructures, as well as the ability of the ROM to predict average and local quantities of interest has been demonstrated.

1 Introduction

Polycrystal simulations performed using the crystal plasticity finite element (CPFE) method are widely used for predicting the mechanical response of metals and alloys at fine scales to characterize a wide range of phenomena associated with fatigue initiation [1], manufacturing processes [2, 3], impact response [4] and many others [5, 6, 7, 8, 9]. The high computational cost of simulations limits application of CPFE to relatively small microstructural volumes or requires simplification of subgrain features that may otherwise require very fine mesh resolution. The issue of computational cost is particularly challenging when parametric sensitivity of response fields and quantification of model/parameter uncertainty are of interest, or if the fine scale response is to be used in conjunction with multiscale modeling.

Reduced order models (ROMs), surrogate models, parallel solution algorithms, and their combinations have been previously proposed to address the high computational cost issue (e.g., [10, 11, 12]). Taylor [13] and Sachs [14] bounds, constructed by neglecting grain-to-grain interactions and assuming constant stress (Sachs) or constant strain (Taylor) within each grain, provide first order bounds to the polycrystalline plasticity problem. Using the Taylor model, Knezevic and Savage have implemented a graphics processing units (GPU)-parallelized spectral database solver [15] for crystal plasticity simulations [2]. Authors report a speedup of approximately 200–1,000 over the conventional CPFEM implementation for a 15,000 grain microstructure [15]. In Ref. [16], the same authors presented GPU parallelization of iterative (Newton-Raphson) and direct (spectral database) solvers for evaluating the nonlinear constitutive law. It was shown that a parallelized direct solver is faster by three orders of magnitude compared to a serial iterative solver. Viscoplastic self-consistent method has been introduced by Lebensohn and Tome for modeling rigid-viscoplastic response of polycrystals [17]. This method improves upon first-order bounds by embedding grains into homogenized medium with effective tangent moduli, leading to more accurate predictions of stress-strain response.

ROMs aim to approximate the response of heterogeneous microstructures with a coarse basis that results in a much smaller system of equations. For general heterogeneous materials, various ROMs have been proposed, including the transformation field analysis (TFA) [18, 19], nonuniform TFA (NTFA) [20, 21, 22, 23] and its extensions [24, 25, 26], numerical potentials [27], self-consistent clustering analysis [28], eigendeformation and eigenstrain-based reduced order homogenization [29, 30, 31, 32, 33, 34], multiscale discrete damage theory [35, 36], proper orthogonal decomposition (POD) [37, 38, 39], and proper generalized decomposition [40]. The general structure of the construction and evaluation of reduced-order and surrogate models is shown in Figure 1. Implementation of these ROMs consists of construction (offline) and the evaluation (analysis/online) stages. During the construction stage, a series of problems defined over the domain of the microstructure volume (i.e. a representative volume element (RVE) or a statistical volume element (SVE)) is solved to build the reduced basis. Once the model is constructed, a reduced nonlinear system of equations coupled with crystal plasticity constitutive equations [41] are evaluated (i.e. the evaluation stage) to simulate the response of the microstructure volume under the provided loading condition. The construction stage is typically computationally costly and the total time spent on construction could scale exponentially with increasing mesh size. In the context of the eigenstrain homogenization method (EHM), the ROM construction requires the evaluation of a series of linear elastic analyses to compute the so called influence functions and the coefficient tensors of the model. The NTFA method uses nonuniform reduced order basis functions (modes), which can be determined by series nonlinear analyses under a set of prescribed loads. Compared to EHM, the NTFA method has been shown to provide more accurate effective response of nonlinear composites. This increase in accuracy is achieved at the expense of costly nonlinear analyses during the construction stage, and dependence of the accuracy on the loading conditions used in training. Based on the NTFA ideas, Fritzen et al. presented a massively parallel GPU implementation of the reduced order method (named FE^{2R}) for composites with a viscoplastic matrix [42, 43]. In this study, a speedup is reported in the order of 20–35 for a GPU-parallelized ROM over a central processing unit (CPU) parallelized version of the same code and 104 speedup over full-field computational homogenization (also known as FE^2 [44]) on a single core. POD-based ROMs also use a series of nonlinear pre-analyses in conjunction with an eigenvalue problem. Typically, the lowest frequency eigenmodes are chosen to construct the reduced basis. POD procedure leads to microscale nonlinear finite element problems with a stiffness




Figure 1: General framework for construction and evaluation of ROMs or surrogate models for nonlinear homogenization.

matrix size corresponding to the number of chosen modes and an array of unknowns corresponding to amplitudes of the modes. A speedup of larger than an order of magnitude compared to direct numerical simulations was reported by Yvonnet and He [38].

The spectral method proposed by Moulinec and Suquet [45, 46, 47] solves the microscale boundary value problem using fast Fourier transforms, thus avoiding the formation (and storage) of a global stiffness matrix. The spectral method has been extensively used for modeling polycrystals and implementation of various crystal plasticity constitutive models has been discussed in a comprehensive review by Lebensohn and Rollett [48]. More recently, Liu et al. have developed a model order reduction technique for the spectral method, named self-consistent clustering analysis [28]. This technique requires evaluation of a series of convolutions in frequency space to construct the reduced order model and delivers good agreement with respect to reference simulations, provided that the interaction tensors are recomputed and a sufficient number of clusters/partitions are used to discretize the microstructure.

Several data-driven models have also been recently proposed based on neural networks [49, 50], geometric deep learning [51], and Gaussian processes [52] among others. Many studies develop surrogate models for free energy density function using neural networks [53] and Gaussian processes [52] for nonlinear elasticity problems. In the context of polycrystals undergoing viscoplastic deformations, Liu et al. [50] developed a neural-network-based approach for constructing a map between the input strain space and the volume-averaged stress, and used it in concurrent multiscale modeling. The aforementioned studies focus on macroscopic stress and homogenized tangent updates using the surrogate models, and the microscopic information is not retained after the training process.

Localization operators for computing the microscale fields (e.g. fatigue indicator parameters [54, 1] or other microscale quantities of interest) are typically not available. Furthermore, loading conditions used to train a model must span the six-dimensional strain space and varying proportional and non-proportional loading paths, which may require a very large number of nonlinear simulations.

This study presents a new parallelized implementation of EHM construction and execution to model the mechanical response of large-sized polycrystalline microstructures. The novelty of the proposed implementation lies particularly in acceleration of the construction stage by leveraging the scalable nature of the pre-analyses used to compute the model. For the ROM construction stage, we demonstrate that near linear scaling can be achieved as a function of degrees of freedom used in the discretization of the microstructure volume using a parallel element-by-element preconditioned conjugate gradient solver. Linear scaling as a function of the number of parts in the ROM follows from the independence of the pre-analyses in EHM formulation. With the proposed implementation procedure, we demonstrate the ability to simulate the response of a microstructure that is an order of magnitude larger than previously reported in the literature [55]. The capabilities of the model were demonstrated by a verification study on a microstructure characterized within the Air Force Research Laboratory (AFRL) additive manufacturing (AM) microstructure challenge series [56, 57, 58]. The challenge microstructure contains approximately 10,000 grains. The global as well as local (i.e. grain scale) quantities of interest predicted using the EHM model were compared with the CPFEE simulations. Furthermore, performance of parallel implementation of the ROM evaluation was analyzed and compared to a serial implementation.

The following notation is used unless otherwise stated. Scalars are denoted by lightface letters, α , Cartesian tensor fields are denoted as italic lightface letters, A_{ijkl} (or σ_{ij}), in indicial notation with italic indices, $ijkl$ (or ij), and referred to as boldface italic, \mathbf{A} (or $\boldsymbol{\sigma}$), in text/tensor notation. A comma in the indices indicates partial spatial derivative. Where appropriate, the top right superscript is used as a descriptor, top left superscript is the iteration counter, and bottom left subscript is the increment counter. Overbar indicates macroscopic quantity averaged over the microstructural domain while overhead dot indicates time derivative. Boldface \mathbf{A} and $\boldsymbol{\sigma}$ are used to represent (6×6) Voigt matrices and (6×1) Voigt vectors in the numerical implementation, respectively.

The rest of the manuscript is organized as follows: An overview of the EHM model used in this study is provided in Section 2. The parallel implementation frameworks for the ROM construction and evaluation are presented in Section 3 and Section 4, respectively. Section 5 is an assessment of the scalability and performance of the parallel EHM model construction framework. Accuracy and performance verification in terms of the macroscale and microscale quantities are discussed. A set of numerical experiments conducted to study the parallel performance of the ROM evaluation is also presented in this section. Conclusions are discussed in Section 6.

2 Reduced order eigenstrain homogenization method

An overview of the governing equations for EHM is provided herein (see [32, 55] for detailed derivation). Consider a macroscopic domain, $\Omega \subset \mathbb{R}^3$, that is made of a periodic polycrystalline microstructural volume, $\Theta \subset \mathbb{R}^3$. The domain Θ consists of N grains, where the domain of each grain α is defined as $\Theta^{(\alpha)} \subset \Theta$. A two-scale asymptotic analysis and the EHM procedure lead to two tightly coupled boundary value problems defined over the macroscale and the microscale domains.

The macroscale boundary value problem in the absence of inertial forces is given by

$$\bar{\sigma}_{ij,j}(\mathbf{x}, t) + \bar{b}_i(\mathbf{x}, t) = 0 \quad (1)$$

$$\bar{\epsilon}_{ij}(\mathbf{x}, t) = \frac{1}{2}(\bar{u}_{i,j}(\mathbf{x}, t) + \bar{u}_{j,i}(\mathbf{x}, t)) \quad (2)$$

$$\bar{u}_i(\mathbf{x}, t) = \bar{u}_i^*(\mathbf{x}, t) \quad \mathbf{x} \in \Gamma_u \quad (3)$$

$$\bar{\sigma}_{ij}(\mathbf{x}, t)\hat{n}_j = \bar{t}_i^*(\mathbf{x}, t) \quad \mathbf{x} \in \Gamma_t \quad (4)$$

where $\bar{\boldsymbol{\sigma}}$ denotes the macroscale stress field, $\bar{\mathbf{b}}$ denotes the macroscale body forces, $\bar{\boldsymbol{\epsilon}}$ is the macroscopic strain, $\bar{\mathbf{u}}$ denotes the macroscale deformation field, $\bar{\mathbf{u}}^*$ is the prescribed displacement over Dirichlet boundaries, Γ_u , $\bar{\mathbf{t}}^*$ is the prescribed traction over Neumann boundaries, Γ_t , $\hat{\mathbf{n}}$ is the normal vector, and \mathbf{x} denotes the macroscale spatial position vector. Using the EHM procedure [32], the microscale homogenized equilibrium equation and the stress-strain relationship lead to the governing equation at the microscale

$$M_{ijkl}^{(\beta)} \dot{\sigma}_{kl}^{(\beta)}(t) - \sum_{\alpha=1}^N \left[P_{ijkl}^{(\alpha\beta)} - \delta^{(\alpha\beta)} I_{ijkl} \right] \dot{\mu}_{kl}^{(\alpha)}(t) = A_{ijkl}^{(\beta)} \dot{\epsilon}_{kl}(t) \quad \beta = 1, \dots, N \quad (5)$$

where α and β represent grain indices, \mathbf{I} is the fourth-order identity tensor, $\mathbf{M}^{(\beta)}$ represents the compliance tensor for grain β , $\mathbf{P}^{(\alpha\beta)}$ represents the interaction tensor between grains α and β , $\mathbf{A}^{(\beta)}$ represents the localization tensor, $\boldsymbol{\mu}^{(\beta)}$ is grain-averaged viscoplastic strain, and $\boldsymbol{\sigma}^{(\beta)}$ is grain-averaged stress. The coefficient tensors ($\mathbf{A}^{(\beta)}$, $\mathbf{M}^{(\beta)}$, and $\mathbf{P}^{(\alpha\beta)}$) retain information about microstructural quantities such as the elastic properties of grains, grain orientation, grain morphology, and grain-to-grain interactions. In a more general setting, the grain indices, α , correspond to the number of parts in the reduced order model, i.e., each grain can be further split into multiple parts. Sub-partitioning grains further would relax the uniform field assumption and may improve the accuracy of the model at the expense of some computational efficiency. In the present study, a part-per-grain model is used. We consider the viscoplastic slip evolution in a given grain to be governed by dislocation glide over the slip systems. The flow rule is then expressed using the Schmid law as:

$$\dot{\mu}_{ij}^{(\alpha)} = \sum_{s=1}^{N_{ss}} \dot{\gamma}^{s(\alpha)} \mathbf{Z}_{ij}^{s(\alpha)} \quad (6)$$

where $\dot{\gamma}^{s(\alpha)}$ is a viscoplastic slip rate on the s^{th} slip system of grain α , N_{ss} is the number of slip systems, and $\mathbf{Z}^{s(\alpha)}$ is the Schmid tensor associated with a slip system, s , of grain α . By Rice and Pierce [59, 60, 61], the slip rate for face centered cubic (FCC) crystals (considered in the present study) is expressed using the flow rule

$$\dot{\gamma}^{s(\alpha)} = \dot{\gamma}_0 \left[\frac{|\tau^{s(\alpha)}|}{g^{s(\alpha)}} \right]^{1/m} \text{sign}(\tau^{s(\alpha)}) \quad (7)$$

where the reference self-shearing rate, $\dot{\gamma}_0$, and rate sensitivity, m , are material parameters, $g^{s(\alpha)}$ is the s^{th} slip system strength of grain α , and $\tau^{s(\alpha)} = \boldsymbol{\sigma}^{(\alpha)} : \mathbf{Z}^{s(\alpha)}$ is the grain-averaged resolved shear stress over the s^{th} slip system of grain α . A Voce-type isotropic hardening model [62, 63] is considered:

$$\dot{g}^{s(\alpha)} = h_0 \left[\frac{g_{sa} - g^{s(\alpha)}}{g_{sa} - g_0} \right] \sum_{r=1}^{N_{ss}} |\dot{\gamma}^{r(\alpha)}| \quad (8)$$

where h_0 is reference self-hardening, g_0 is the initial strength of a slip system, and g_{sa} is saturation strength.

In EHM, the construction phase consists of computing the coefficient tensors, whereas the evaluation phase refers to the solution of Eq. 5 along with the constitutive equations describing the viscoplastic flow. The compliance tensor, $\mathbf{M}^{(\beta)}$, is defined by the lattice orientation and the anisotropic tensor of elastic moduli of the material. The localization tensor, $\mathbf{A}^{(\beta)}$, and the interaction tensors, $\mathbf{P}^{(\alpha\beta)}$, are related to each other through the following kinematic consistency formula [55, 18]:

$$\sum_{\alpha=1}^N P_{ijkl}^{(\alpha\beta)} = I_{ijkl} - A_{ijkl}^{(\beta)} \quad (9)$$

Therefore, computation of localization tensors from given interaction tensors is straightforward. Computation of $\mathbf{P}^{(\alpha\beta)}$ requires the evaluation of N phase influence function problems (IFP) governed by

$$\left[L_{ijmn}(\mathbf{y}) \left(h_{(m,n)kl}^{(\alpha)}(\mathbf{y}) - I_{mnkl} 1^{(\alpha)}(\mathbf{y}) \right) \right]_{,j} = 0 \quad \mathbf{y} \in \Theta \quad (10)$$

$$h_{mkl}^{(\alpha)}(\mathbf{y}) = 0 \quad \mathbf{y} \in \partial\Theta^v \quad (11)$$

$$h_{mkl}^{(\alpha)}(\mathbf{y}) = h_{mkl}^{(\alpha)}(\mathbf{y} + \hat{\mathbf{y}}) \quad \mathbf{y} \in \partial\Theta^{fe} \quad (12)$$

where \mathbf{y} is the microscale position vector, $\mathbf{h}^{(\alpha)}$ is the phase influence function, $\mathbf{L} = \mathbf{M}^{-1}$ is the tensor of elastic moduli, $\partial\Theta^v$ is the set of vertices (i.e., corners) of the microstructural volume, $\hat{\mathbf{y}}$ is the period of the microstructure, and $\partial\Theta^{fe}$ is a set defined over the faces and edges of the microstructural volume. Parentheses in the subscript of the phase influence function stand for the symmetric part of the gradient with respect to the indices in parentheses. Indicator function, $1^{(\alpha)}$, is defined as $1^{(\alpha)}(\mathbf{y}) = 1$ if $\mathbf{y} \in \Theta^{(\alpha)}$ (where $\Theta^{(\alpha)}$ is the domain of grain α) and $1^{(\alpha)}(\mathbf{y}) = 0$ elsewhere. The phase influence function is fixed at the vertices of the microstructural volume to eliminate rigid body motion, and periodic boundary conditions are enforced along the faces and edges of the microstructural volume. The interaction tensors are expressed as a function of the influence functions as:

$$P_{ijkl}^{(\alpha\beta)} = \frac{1}{C^{(\beta)}} \int_{\Theta} 1^{(\beta)}(\mathbf{y}) h_{(i,j)kl}^{(\alpha)}(\mathbf{y}) d\mathbf{y} \quad (13)$$

where $C^{(\beta)}$ is the volume fraction of grain β . The general workflow of the EHM consists of two main steps:

1. Construction stage: In this stage, the IFPs (given in Eqs. 12) are solved for each grain α in the microstructure, and the localization, $\mathbf{A}^{(\alpha)}$, and interaction tensors, $\mathbf{P}^{(\alpha\beta)}$, are precomputed. In the present study, the solution of the IFPs is carried out in parallel using Algorithm 1 as discussed in Section 3.
2. ROM evaluation: In this stage, the microscale governing equation (given in Eq. 5) is solved at each macroscale integration point for a given macroscopic strain increment. Since the microscale governing equation is nonlinear (due to viscoplastic strain) a staggering scheme is used to solve the system, as shown in Algorithm 2. Parallel implementation of residual and Jacobian assembly routines of Newton-Rhapson algorithm is further demonstrated in Algorithms 3 and 4.

3 Efficient construction of reduced order model

As shown by Zhang and Oskay [32], by using standard Galerkin-Bubnov finite element discretization and using Voigt notation, phase IFPs are reduced to a system of linear equations

$$\mathbf{K}\mathbf{D} = \mathbf{B} \quad (14)$$

where $\mathbf{K} \in \mathbb{R}^{N_{\text{dof}} \times N_{\text{dof}}}$ is the stiffness matrix, $\mathbf{B} \in \mathbb{R}^{N_{\text{dof}} \times 6N}$ is the external force matrix, $\mathbf{D} \in \mathbb{R}^{N_{\text{dof}} \times 6N}$ is the matrix of nodal coefficients of the influence function with $6N$ columns for three-dimensional problems, and N_{dof} is the number of degrees of freedom in a finite element mesh of the microstructural domain. The \mathbf{B} and \mathbf{D} matrices are given as

$$\mathbf{B} = \{\hat{\mathbf{b}}^{(1)}, \hat{\mathbf{b}}^{(2)}, \dots, \hat{\mathbf{b}}^{(\alpha)}, \dots, \hat{\mathbf{b}}^{(N)}\} \quad (15)$$

$$\mathbf{D} = \{\hat{\mathbf{d}}^{(1)}, \hat{\mathbf{d}}^{(2)}, \dots, \hat{\mathbf{d}}^{(\alpha)}, \dots, \hat{\mathbf{d}}^{(N)}\} \quad (16)$$

where each $\hat{\mathbf{b}}^{(\alpha)}$ is a matrix of external forces and each $\hat{\mathbf{d}}^{(\alpha)}$ is a matrix of displacements with N_{dof} number of rows and six columns.

For microstructures with complex features or many grains, the number of grains (or reduced order parts resulting from subpartitioning), N , may be large, which results in a large number of right hand sides in Eq. 16. These types of microstructures would also require a fine finite element mesh (possibly with higher-order elements) so that the solution of each system for $\hat{\mathbf{d}}^{(\alpha)}$ becomes expensive and storing the global stiffness matrix becomes problematic. Therefore, the cost of ROM construction for such microstructures scales with both N and N_{dof} , necessitating the introduction of parallel computing algorithms.

To this end, the present study proposes an efficient ROM construction methodology based on an element-by-element (EBE) preconditioned conjugate gradient (PCG [64]) solver. The EBE methodology for the solution of linear systems was first introduced in Ref. [65] and then applied to the PCG method in Ref. [66]. The main advantage of the EBE method is that the assembly and storage of the stiffness matrix is not required in order to solve the linear system. The PCG algorithm combined with EBE and parallel computing techniques can converge quickly when an efficient preconditioner is used.

Since each IFP (Eq. 16) is independent of the others, the IFPs are split into batches, and the batches are submitted to separate compute nodes to achieve parallelism. Then, each IFP in a given batch is solved one-by-one using the parallel EBE PCG algorithm. Within each compute node, the EBE PCG method is implemented using shared memory parallelization (OpenMP library [67]).

The general procedure for the parallel EBE PCG method for evaluating the phase IFPs is shown in Algorithm 1. All PCG operations are performed in shared memory parallelism by splitting the mesh of the polycrystal into partitions as indicated by superscript (*partition*). This domain decomposition performed for parallel computation is distinct from the partitioning of the domain for model order reduction introduced in Section 2. These partitionings could be made independent of each other. However in this study, we choose them to be the same, hence each parallel computing partition (and reduced order part) is the domain of a single grain.

The six linear solves (associated with each column of the force vector, $\hat{\mathbf{b}}^{(\alpha)}$) are evaluated simultaneously by vectorizing the displacement and external force matrices such that

$$\mathbf{d}^{(\alpha)} = \{\hat{\mathbf{d}}_1^{(\alpha)T}, \dots, \hat{\mathbf{d}}_6^{(\alpha)T}\}^T \quad (17)$$

$$\mathbf{b}^{(\alpha)} = \{\hat{\mathbf{b}}_1^{(\alpha)T}, \dots, \hat{\mathbf{b}}_6^{(\alpha)T}\}^T \quad (18)$$

This vectorization implies that the number of degrees of freedom at each node is set to 18 (rather than 3). In the algorithm, \mathbf{d} indicates the solution vector, \mathbf{b} is the external force vector, $\tilde{\mathbf{D}}$ is the preconditioner, α is the line search constant, \mathbf{r} indicates the residual vector, β is the Gram-Schmidt orthogonalization constant, and \mathbf{p} stands for orthogonal search directions. In steps 1–6, the solver is initialized using an initial guess for the solution, ${}^0\mathbf{d}$ (IFP index, α , is omitted for brevity). In step 1, an initial guess for the initial force vector, ${}^0\mathbf{f}^{(partition)}$, is assembled using the standard finite element procedures (such that $\mathbf{A}_{el=1}^{N^{(partition)}}$ represents the finite element assembly over elements in a given partition), and $N_{el}^{(partition)}$ stands for the number of elements in a given partition. The internal force vector is computed without explicitly forming and storing the elemental stiffness matrix, \mathbf{K}^{el} , as described in Ref. [68]. The Jacobi preconditioner is assembled similarly (step 2). In step 3, the internal forces, ${}^0\mathbf{f}^{(partition)}$, the external forces, $\mathbf{b}^{(partition)}$, and the preconditioner, $\tilde{\mathbf{D}}^{(partition)}$, computed for individual partitions are synchronized. Nodes shared between the partitions are designated as halo (or ghost) nodes, and synchronization is performed by looping through all halo nodes and summing the quantity of interest over all partitions containing a given halo node. This synchronization process is illustrated in Figure 2 for two partitions. Yellow (see electronic version of this article for colors) circles indicate halo nodes i.e. those nodes that are shared between partition 1 and 2, and blue arrows indicate nodal force vectors. For each halo node, a vector sum of force vectors from two partitions (shown in Figure 2a) is performed, resulting in the force vectors in Figure 2b. After synchronization is completed, the nodal forces for each halo node are equal in partition 1 and partition 2.

Subsequent iterations are described by the rest of the Algorithm 1. In steps 4–6, the residual is initialized and preconditioned, and the search direction is initialized. Steps 8–16 describe the update on the residual and solution at an arbitrary iteration, m . Steps 11, 12, 13, and 15 follow standard PCG procedures with the additional consideration of partitioning and parallel computation. The scalars α and β are updated in steps 10 and 14 such that numerators and denominators are computed by performing summation over all partitions. Iterations are performed until the following convergence criterion is reached

$$\sum_{partition=1}^{N_{partitions}} m_{\mathbf{r}^{(partition)}}^T m_{\mathbf{p}^{(partition)}} < \delta^2 \sum_{partition=1}^{N_{partitions}} {}^0\mathbf{r}^{(partition)T} {}^0\mathbf{p}^{(partition)} \quad (19)$$

for a relative tolerance, δ .

Periodic boundary conditions are enforced over the faces and edges of the microstructure by replacing replica node number with primary node number. Each node on the primary faces and edges is associated with a node on the replica faces and edges of the microstructure as depicted in Figure 3. In this study, the microstructure mesh is generated from a voxelized geometry where each voxel is split into six linear tetrahedra. Therefore, there is a unique replica node for each primary node on all three face pairs. As for the edges, each primary node on a primary edge has three corresponding replica nodes on three replica edges. For a primary node on a given microstructure face or edge, replica nodes are identified using an efficient algorithm based on the node numbering for regular/voxelized meshes. Finally, all degrees of freedom at the vertices of the microstructure are fixed to ensure well-posedness of the phase IFPs.

In the present study, the parallel EBE PCG method was implemented based on the Femera solver [68]. The initial guess for the solution vector, ${}^0\mathbf{d}$, was set to zero. A diagonal Jacobi

Algorithm 1 Parallel EBE PCG method

Solve single phase IFP $\mathbf{Kd} = \mathbf{b}$ using parallel EBE PCG method.

- 1: Initialize internal forces : ${}^0\mathbf{f}^{(partition)} \leftarrow A_{el=1}^{N_{el}^{(partition)}} \mathbf{K}^{el} \left({}^0\mathbf{d}^{el(partition)} \right)$
 - 2: Compute preconditioner : $\tilde{\mathbf{D}}^{(partition)} \leftarrow A_{el=1}^{N_{el}^{(partition)}} \text{diag}(\mathbf{K}^{el})$
 - 3: Synchronize over the partitions : ${}^0\mathbf{f}^{(partition)}, \mathbf{b}^{(partition)}, \tilde{\mathbf{D}}^{(partition)}$
 - 4: Initialize residual : ${}^0\mathbf{r}^{(partition)} \leftarrow \mathbf{b}^{(partition)} - {}^0\mathbf{f}^{(partition)}$
 - 5: Update preconditioned residual : ${}^0\mathbf{z}^{(partition)} \leftarrow \tilde{\mathbf{D}}^{-1(partition)} {}^0\mathbf{r}^{(partition)}$
 - 6: Initialize orthogonal search directions : ${}^0\mathbf{p}^{(partition)} \leftarrow {}^0\mathbf{z}^{(partition)}$
 - 7: **Loop until convergence**
 - 8: Compute internal forces : ${}^m\mathbf{f}^{(partition)} \leftarrow A_{el=1}^{N_{el}^{(partition)}} \mathbf{K}^{el} \left({}^m\mathbf{p}^{el(partition)} \right)$
 - 9: Synchronize internal forces : ${}^m\mathbf{f}^{(partition)}$
 - 10: Update line search constant : ${}^m\alpha \leftarrow \frac{\sum_{partition=1}^{N_{partitions}} {}^m\mathbf{r}^{(partition)T} {}^m\mathbf{z}^{(partition)}}{\sum_{partition=1}^{N_{partitions}} {}^m\mathbf{p}^{(partition)T} {}^m\mathbf{f}^{(partition)}}$
 - 11: Update solution vector : ${}^{m+1}\mathbf{d}^{(partition)} \leftarrow {}^m\mathbf{d}^{(partition)} + {}^m\alpha {}^m\mathbf{p}^{(partition)}$
 - 12: Update residual : ${}^{m+1}\mathbf{r}^{(partition)} \leftarrow {}^m\mathbf{r}^{(partition)} - {}^m\alpha {}^m\mathbf{f}^{(partition)}$
 - 13: Update preconditioned residual : ${}^{m+1}\mathbf{z}^{(partition)} \leftarrow \tilde{\mathbf{D}}^{-1(partition)} {}^{m+1}\mathbf{r}^{(partition)}$
 - 14: Update Gram-Schmidt constant : ${}^m\beta \leftarrow \frac{\sum_{partition=1}^{N_{partitions}} {}^{m+1}\mathbf{r}^{(partition)T} {}^{m+1}\mathbf{z}^{(partition)}}{\sum_{partition=1}^{N_{partitions}} {}^m\mathbf{r}^{(partition)T} {}^m\mathbf{z}^{(partition)}}$
 - 15: Update orthogonal search directions : ${}^{m+1}\mathbf{p}^{(partition)} \leftarrow {}^{m+1}\mathbf{z}^{(partition)} + {}^m\beta {}^m\mathbf{p}^{(partition)}$
 - 16: Update iteration counter : $m \leftarrow m + 1$
 - 17: **End loop**
-

preconditioner, $\tilde{\mathbf{D}}^{(partition)}$, was selected (as shown in Algorithm 1) for its relatively straightforward parallel implementation and low memory requirements. Microstructure vertices were fixed by setting the inverse of the preconditioner, $\tilde{\mathbf{D}}^{-1}$, to zero for all global degrees of freedom corresponding to the microstructure vertex nodes.

4 Efficient evaluation of reduced order model

This section introduces the microscale stress update procedure followed by the parallel implementation of the ROM evaluation. Equation 5 is solved for two sets of unknowns: grain-averaged stresses, $\boldsymbol{\sigma}^{(\alpha)}$, and slip systems strengths, $\mathbf{g}^{(\alpha)}$, for all slip systems within each grain α . For notational convenience, grain-averaged stresses and slip systems strengths are collected into single vectors as follows

$$\boldsymbol{\Sigma} = \{ \boldsymbol{\sigma}^{(1)T}, \boldsymbol{\sigma}^{(2)T}, \dots, \boldsymbol{\sigma}^{(N)T} \}^T \quad (20)$$

$$\mathbf{g} = \{ \mathbf{g}^{(1)T}, \mathbf{g}^{(2)T}, \dots, \mathbf{g}^{(N)T} \}^T \quad (21)$$

where $\boldsymbol{\Sigma}$ denotes a vector that contains all microscale stresses, \mathbf{g} denotes a vector that contains all microscale slip system strengths, and each $\mathbf{g}^{(\alpha)}$ contains slip system strengths across all slip systems.




Figure 2: Synchronization of forces for parallel solver by summing force vectors for each halo node: (a) before synchronization and (b) after synchronization.




Figure 3: Application of periodic boundary conditions to the faces and edges of the microstructure.

The residual and Jacobian of Eq. 5 are given by

$$\boldsymbol{\phi}^{(\beta)} = \sum_{\alpha=1}^N (\delta^{(\alpha\beta)} \mathbf{I} - \mathbf{P}^{(\alpha\beta)})_{l+1} \dot{\boldsymbol{\mu}}^{(\alpha)} + \mathbf{M}^{(\beta)} \frac{\boldsymbol{\sigma}^{(\beta)} - {}_{l-1} \boldsymbol{\sigma}^{(\beta)}}{\Delta t} - \mathbf{A}^{(\beta)} \frac{{}_{l+1} \bar{\boldsymbol{\epsilon}} - {}_{l-1} \bar{\boldsymbol{\epsilon}}}{\Delta t} \quad (22)$$

$$\frac{\partial \boldsymbol{\phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} = (\delta^{(\eta\beta)} \mathbf{I} - \mathbf{P}^{(\eta\beta)}) \frac{\dot{\gamma}_0}{m} \sum_{s=1}^{N_{ss}} \frac{1}{g^{s(\eta)}} \left[\frac{|\tau^{s(\eta)}|}{g^{s(\eta)}} \right]^{\frac{1-m}{m}} \tilde{\mathbf{Z}}^{s(\eta)} + \delta^{(\beta\eta)} \frac{\mathbf{M}^{(\beta)}}{\Delta t} \quad (23)$$

where $\boldsymbol{\phi}$ is the residual of the constitutive law, $\partial \boldsymbol{\phi} / \partial \boldsymbol{\sigma}$ is the Jacobian of the constitutive law, \mathbf{I} is the fourth order identity matrix, indices l indicate values of the fields at time $l t$, $\mathbf{A}^{(\beta)}$ are the localization matrices, $\tilde{\mathbf{Z}}^{s(\eta)} = \mathbf{Z}^{s(\eta)} \mathbf{Z}^{s(\eta)T}$, and $\mathbf{P}^{(\alpha\beta)}$ are the interaction matrices. A two-level staggering scheme [69] is used to update grain-averaged stresses and strengths as shown in Algorithm 2. The macroscopic strain increment and state variables at the previous increment are passed as input to the algorithm. The stresses are updated through Newton-Raphson iterations (steps 4–8) under the condition that the residual is minimized within an acceptable tolerance. The slip system strengths are updated explicitly based on Eq. 8 at step 10. This process is repeated until convergence is

Algorithm 2 Microscale Stress Update Algorithm

 Inputs: Macroscale strain increment $\Delta \bar{\boldsymbol{\epsilon}}$

 Outputs: Macroscale stress $\bar{\boldsymbol{\sigma}}$, Macroscale tangent $\partial \bar{\boldsymbol{\sigma}} / \partial \bar{\boldsymbol{\epsilon}}$

- 1: Microscale iteration counter : $k \leftarrow 0$
 - 2: Initial guess : ${}_{l+1}^k \boldsymbol{\Sigma} \leftarrow {}_l \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g} \leftarrow {}_l \mathbf{g}$
 - 3: **State iterations until convergence**
 - 4: **N-R iterations until convergence**
 - 5: $\Phi \Big|_{{}_{l+1}^k \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}} \leftarrow \text{ComputeResidual}(\dots)$
 - 6: $\left[\frac{\partial \Phi}{\partial {}_{l+1}^k \boldsymbol{\Sigma}} \right] \Big|_{{}_{l+1}^k \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}} \leftarrow \text{ComputeJacobian}(\dots)$
 - 7: Solve linear system : $\Delta \boldsymbol{\Sigma} \leftarrow - \left[\frac{\partial \Phi}{\partial {}_{l+1}^k \boldsymbol{\Sigma}} \Big|_{{}_{l+1}^k \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}} \right]^{-1} \Phi \Big|_{{}_{l+1}^k \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}}$
 - 8: Update total stress vector : ${}_{l+1}^{k+1} \boldsymbol{\Sigma} \leftarrow {}_{l+1}^k \boldsymbol{\Sigma} + \Delta \boldsymbol{\Sigma}$
 - 9: **End N-R iterations**
 - 10: Update hardening : ${}_{l+1}^{k+1} \mathbf{g} \leftarrow {}_l \mathbf{g} + \dot{\mathbf{g}} \Big|_{{}_{l+1}^{k+1} \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}} \Delta t$
 - 11: Update iteration counter : $k \leftarrow k + 1$
 - 12: **End state iterations**
 - 13: Update Tangent Modulus : $\frac{\partial {}_{l+1} \bar{\boldsymbol{\sigma}}}{\partial {}_{l+1} \bar{\boldsymbol{\epsilon}}} \leftarrow \sum_{\alpha=1}^N C^{(\alpha)} \left[\left[\frac{\partial \Phi}{\partial {}_{l+1}^k \boldsymbol{\Sigma}} \Big|_{{}_{l+1}^k \boldsymbol{\Sigma}, {}_{l+1}^k \mathbf{g}} \right]^{-1} \frac{\mathbf{A}}{\Delta t} \right]^{(\alpha)}$
 - 14: Update Macroscopic Stress : ${}_{l+1} \bar{\boldsymbol{\sigma}} \leftarrow \sum_{\alpha=1}^N C^{(\alpha)} \boldsymbol{\sigma}^{(\alpha)}$
-

reached in terms of grain-averaged stresses and slip system strengths within acceptable tolerances. The converged grain-averaged stresses and strengths are then used to update macroscale stress, ${}_{l+1} \bar{\boldsymbol{\sigma}}$, and tangent modulus, $\partial {}_{l+1} \bar{\boldsymbol{\sigma}} / \partial {}_{l+1} \bar{\boldsymbol{\epsilon}}$ (steps 13–14).

In order to speed up convergence of the solver, the matrix decomposition computed for the stress update prior to convergence of the microscale problem is reused for tangent computation. Since the difference between ${}_{l+1}^{k+1} \boldsymbol{\Sigma}$ and ${}_{l+1}^k \boldsymbol{\Sigma}$ is small (within the tolerance used) at the last increment prior to convergence, accuracy of the tangent computation isn't significantly affected by using the decomposition computed from ${}_{l+1}^k \boldsymbol{\Sigma}$ instead of ${}_{l+1}^{k+1} \boldsymbol{\Sigma}$. Furthermore, the initial guess for microscale stresses used in Algorithm 2 was replaced with explicit solutions for ${}_{l+1}^0 \boldsymbol{\Sigma}$ and ${}_{l+1}^0 \mathbf{g}$ computed using the equations

$${}_{l+1}^0 \boldsymbol{\sigma}^{(\beta)} = {}_l \boldsymbol{\sigma}^{(\beta)} + \mathbf{M}^{(\beta)-1} \left[\mathbf{A}'^{(\beta)} \frac{{}_{l+1} \bar{\boldsymbol{\epsilon}} - {}_l \bar{\boldsymbol{\epsilon}}}{\Delta t} - \sum_{\alpha=1}^N (\delta^{(\alpha\beta)} \mathbf{I} - \mathbf{P}'^{(\alpha\beta)}) {}_l \dot{\boldsymbol{\mu}}^{(\alpha)} \right] \Delta t \quad (24)$$

$${}_{l+1}^0 \mathbf{g} = {}_l \mathbf{g} + \dot{\mathbf{g}} \Big|_{{}_l \boldsymbol{\Sigma}, {}_l \mathbf{g}} \Delta t. \quad (25)$$

Explicit estimates speed up convergence of the solver significantly, particularly in the elastic and

Algorithm 3 ComputeResidual

```

1: function COMPUTE_RESIDUAL(...)
2:   Initialize viscoplastic strain and residual :  $\dot{\boldsymbol{\mu}}^{(\beta)} \leftarrow \mathbf{0}, \boldsymbol{\Phi}^{(\beta)} \leftarrow \mathbf{0}$ 

3:   !$OMP PARALLEL DO PRIVATE( $\beta, s$ )
4:   for  $\beta \leftarrow 1$  to  $N$  do
5:     for  $s \leftarrow 1$  to  $N_{ss}$  do
6:       Compute slip rate :  $\dot{\gamma}^{s(\beta)}$ 
7:       Compute inelastic strain rate :  $\dot{\boldsymbol{\mu}}^{(\beta)} \leftarrow \dot{\boldsymbol{\mu}}^{(\beta)} + \dot{\gamma}^{s(\beta)} \mathbf{Z}^{s(\beta)}$ 
8:     end for
9:   end for
10:  !$OMP END PARALLEL DO

11:  !$OMP PARALLEL DO PRIVATE( $\beta, \alpha, \delta^{(\alpha\beta)}$ )
12:  for  $\beta \leftarrow 1$  to  $N$  do
13:    for  $\alpha \leftarrow 1$  to  $N$  do
14:      Add viscoplastic contribution :  $\boldsymbol{\Phi}^{(\beta)} \leftarrow \boldsymbol{\Phi}^{(\beta)} + (\delta^{(\alpha\beta)} \mathbf{I} - \mathbf{P}^{(\alpha\beta)}) \dot{\boldsymbol{\mu}}^{(\alpha)}$ 
15:    end for
16:    Add elastic contribution :  $\boldsymbol{\Phi}^{(\beta)} \leftarrow \boldsymbol{\Phi}^{(\beta)} + \mathbf{M}^{(\beta)} \frac{l_{+1} \boldsymbol{\sigma}^{(\beta)} - l \boldsymbol{\sigma}^{(\beta)}}{\Delta t}$ 
17:    Add total contribution :  $\boldsymbol{\Phi}^{(\beta)} \leftarrow \boldsymbol{\Phi}^{(\beta)} - \mathbf{A}^{(\beta)} \frac{l_{+1} \bar{\boldsymbol{\epsilon}} - l \bar{\boldsymbol{\epsilon}}}{\Delta t}$ 
18:  end for
19:  !$OMP END PARALLEL DO
20:  Assemble  $\boldsymbol{\Phi}$  from  $\boldsymbol{\Phi}^{(\beta)}$ 

21: end function

```

viscoplastic regions. However, several state iterations are still required inside the elasto-viscoplastic transition region due to nonlinearities.

4.1 Parallel computation of the residual and Jacobian

The significant majority of the computational cost of executing Algorithm 2 is due to (1) residual computation (step 5), (2) Jacobian assembly (step 6), and (3) the solution of the system of equations (step 7). A pseudo-algorithm for parallel implementation of the residual computation is shown in Algorithm 3. First, viscoplastic strain rates for all grains are computed in one for-loop, which is parallelized with respect to grain index β . $\dot{\gamma}^{s(\alpha)}$ is a function of stress and hardening of the α^{th} grain only, and is computed using Eq. 7. The residual for each grain, $\boldsymbol{\Phi}^{(\alpha)}$, is computed in the second for-loop, which is parallelized with respect to grain index β . The process of assembly of the Jacobian matrix (Eq. 23) is shown in Algorithm 4. Jacobian assembly consists of computing $\partial \boldsymbol{\mu} / \partial \boldsymbol{\sigma}$, and adding elastic and viscoplastic contributions to the matrix. In both steps of the assembly process, the loops are parallelized with respect to the grain index β . Since computing the outer product of the Schmid tensor, $\tilde{\mathbf{Z}}^{s(\eta)}$, at each iteration is computationally intensive, it is precomputed at the initialization stage. Precomputing this value is possible provided that the texture evolution is considered to be small under the applied loading conditions. Assembly of the Jacobian matrix for the computation of tangent moduli (step 13 in Algorithm 2) follows the same pattern, except that there are six right hand sides (due to the minor symmetry of the localization tensor) instead of one.

4.2 Parallel evaluation of linear system

The evaluation of the stress increment $\Delta \boldsymbol{\Sigma}$ in step 7 of Algorithm 2 is also parallelized. In the classical EHM formulations [32], the resulting Jacobian is dense, hence stress update scales poorly

Algorithm 4 ComputeJacobian

```

1: function COMPUTEJACOBIAN(...)
2:   Initialize  $\frac{\partial \dot{\boldsymbol{\mu}}^{(\eta)}}{\partial \boldsymbol{\sigma}^{(\eta)}}$  and the Jacobian :  $\frac{\partial \dot{\boldsymbol{\mu}}^{(\eta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} \Leftarrow 0, \frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} \Leftarrow 0$ 
3:   !$OMP PARALLEL DO PRIVATE( $\beta, s$ )
4:   for  $\beta \leftarrow 1$  to  $N$  do
5:     for  $s \leftarrow 1$  to  $N_{ss}$  do
6:       
$$\frac{\partial \dot{\boldsymbol{\mu}}^{(\eta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} \Leftarrow \frac{\partial \dot{\boldsymbol{\mu}}^{(\eta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} + \frac{\dot{\gamma}_0}{m} \frac{1}{g^{s(\eta)}} \left[ \frac{|\tau^{s(\eta)}|}{g^{s(\eta)}} \right]^{\frac{1-m}{m}} \dot{\mathbf{Z}}^{s(\eta)}$$

7:     end for
8:   end for
9:   !$OMP END PARALLEL DO

10:  !$OMP PARALLEL DO PRIVATE( $\beta, \eta, \delta^{(\eta\beta)}$ )
11:  for  $\beta \leftarrow 1$  to  $N$  do
12:    for  $\eta \leftarrow 1$  to  $N$  do
13:      Add viscoplastic contribution :  $\frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} \Leftarrow \frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} + (\delta^{(\eta\beta)} \mathbf{I} - \mathbf{P}^{(\eta\beta)}) \frac{\partial \dot{\boldsymbol{\mu}}^{(\eta)}}{\partial \boldsymbol{\sigma}^{(\eta)}}$ 
14:      Add elastic contribution :  $\frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} \Leftarrow \frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}} + \delta^{(\eta\beta)} \frac{\mathbf{M}^{(\beta)}}{\Delta t}$ 
15:    end for
16:  end for
17:  !$OMP END PARALLEL DO
18:  Assemble  $\frac{\partial \boldsymbol{\Phi}}{\partial \boldsymbol{\Sigma}}$  from  $\frac{\partial \boldsymbol{\Phi}^{(\beta)}}{\partial \boldsymbol{\sigma}^{(\eta)}}$ 
19: end function

```

with number of grains. A formulation that employs a sparse Jacobian with better scalability has been proposed in Ref. [55]. The sparse Jacobian is achieved by adjusting the interaction matrices, $\mathbf{P}^{(\alpha\beta)}$ without changing the overall structure of the ROM. This formulation (corresponding to a 1-layer ϵ -sparse EHM model in Ref. [55]) is adopted in this study such that only the interaction between nearest neighbors are considered. Solution of the sparse system of equations is handled using a direct sparse solver - PARDISO [70]. The PARDISO solver is parallelized internally and parallelization is activated by passing the number of cores/threads as an environment variable. Further, additional time is saved during the multiscale simulation by performing symbolic decomposition of the Jacobian at initialization using PARDISO. Due to sparse nature of the interaction tensors, for-loops at step 13 (Algorithm 3) and step 12 (Algorithm 4) have been replaced with loops over all first nearest neighbors (of grain β) since $\mathbf{P}^{(\alpha\beta)} = \mathbf{0}$ for non-neighboring grains α and β . The number of neighbors and neighbor indices are fetched from precomputed maps and denoted as thread-private variables to preserve thread independence. Sparse formulation leads to a sparse Jacobian matrix which is stored in a compressed row storage (CRS) format.

5 Verification : AFRL challenge problem

In this section, the performance of the proposed ROM construction and evaluation framework is assessed using the microstructure of an additively manufactured Inconel 625 alloy that was thoroughly characterized within the AFRL AM Modeling Challenge Series [56, 57, 58].

The grain morphology of the microstructure obtained from Ref. [56] after digital "cleanup" is shown in Figure 4a. The original voxelized microstructure contained approximately 30,000 grains, a small amount of porosity (0.001% by volume) and experimental markers. The microstructure

Table 1: Elastic and viscoplastic properties of Inconel 625.

C_{11}	C_{12}	C_{44}	g_0	g_{sa}	h_0	m	$\dot{\gamma}_0$
243.3 GPa	156.7 GPa	117.8 GPa	138 MPa	170 MPa	400 MPa	0.048	0.001

was processed by deleting the experimental markers and merging pores and border artifacts into neighboring grains using DREAM.3D software [71]. The number of grains was reduced from approximately 30,000 to 9,294 by merging small grains (with diameters $< 10.7 \mu\text{m}$) into neighboring large grains. The voxel resolution was set to $4 \mu\text{m}$, which does not significantly alter the morphology of the grains. The final size of the processed microstructure is $496 \times 700 \times 496 \mu\text{m}$. Since the implementation of EBE PCG algorithm was optimized for linear tetrahedral elements we split each voxel into six tetrahedral elements before solving IFPs. The preprocessed microstructure, shown in Figure 4a, contained 9,294 grains and was discretized using 16.2 million linear tetrahedral elements and 2.75 million nodes. The morphological characteristics (e.g. average grain size, orientation distribution, etc.) of the microstructure are provided in Ref. [72].

A set of 28 grains named "challenge grains" that were experimentally characterized in particular detail [56] is shown in Figure 4b. It is worth noting that challenge grains and their nearest neighbors were not merged to any other grains in preprocessing to preserve their morphology for accurate grain-scale results. The material was subjected to tensile loading (with occasional partial unloads) and the response state was measured at six load states (denoted as S1–S6). For the load state S1, the specimen was loaded to 100 MPa, for S2 it was further loaded to 200 MPa. S3 was achieved by further loading to 300 MPa and then unloading by 50 MPa. Load states S4, S5, and S6 were achieved by continuing to load the specimen to 0.35 %, 0.5 %, and 1 % strain levels, respectively, and unloading by 50 MPa after the target strain level is reached.

The reference CPFE simulation was carried out by Yeratapally et al. [72] and was completed in 44.5 hours on 640 Intel Xeon E5-2670 processing cores on the NASA Langley Research Center (LaRC) K-cluster. Boundary conditions used for the CPFE study are given in Figure 4c where the domain of the cuboid represents the challenge microstructure. Laplacian smoothing was used for the CPFE microstructure to smooth the voxelized grain boundaries. The challenge microstructure was then meshed with quadratic tetrahedral elements. Cubic elastic constants (C_{11} , C_{12} , and C_{44}) and the viscoplastic properties used in the present study were calibrated for Inconel 625 superalloy in Ref. [72] and are given in Table 1. The pole figure (generated using MTEX [73]) depicted in Figure 4d reveals scattered pattern and absence of texture in the microstructure.

In order to construct the ROM for the challenge problem, 9,294 IFPs (Eq. 12) were split into twelve batches with around 773 IFPs in each batch. The batches were submitted to twelve separate compute nodes for parallel evaluation (using Algorithm 1). The number of batches could be varied and, in general, depends on the availability of computational resources. IFPs were solved to a relative tolerance of 10^{-3} using 32 physical cores. Solution of each IFP consumed approximately three minutes on an Intel Xeon Processor E5-2683 v4 processor on average.

For the ROM, a single reduced integration linear hexagonal element was used for the macroscale finite element model. The macroscale boundary conditions were set the same as the reference CPFE simulation and are shown in Figure 4c. Loading is applied along the Y axis, while axes X and Z are transverse to the loading direction. The loading on the top surface along the Y axis was alternated between traction boundary conditions and displacement boundary conditions corresponding to the

description of the load states from the experimental procedure. The target strain rate used in the simulation was $10^{-2}/s$ corresponding to the strain rate used for CPFE simulations. However, since mixed traction/displacement boundary conditions were used to simulate the loading, the exact strain rates slightly vary from the target strain rate. The macroscale simulation was performed using the Abaqus/Standard¹ finite element software package [74]. ROM evaluation (i.e. solving Algorithm 2 for a given macroscopic strain increment) is performed by a user subroutine at the single integration point of macroscopic finite element. In total, around 100 macroscale increments of the ROM consumed less than 16 minutes on an Intel Xeon Gold 6130 CPU @ 2.10 GHz processor employing only one core. It is worth mentioning that the interaction tensors from the sparse formulation [55] were used to speed up the ROM simulations, whereas the localization tensor from the full EHM formulation [32] were employed to prevent the introduction of additional errors in the elastic region. Using localization tensors from the full EHM does not add to the computational cost of the ROM simulation.

The preprocessed microstructure and mesh are not identical to those used for the reference simulations since preprocessing was performed separately and by a different group. It is worthy of note that the ROM results could have been compared to experimental data directly in the form of a validation study. In that case, the experimental macroscale stress-strain curve would be used to calibrate viscoplastic properties used in the ROM, similar to that performed for CPFE in Ref. [72]. Direct comparison of the ROM results with CPFE ensures that the same set of viscoplastic properties (calibrated using CPFE) is used for both simulations. The rest of this section discusses a comparison between ROM and CPFE results, scalability of the ROM construction framework, and performance of the parallel implementation of ROM evaluation.

5.1 Comparison of CPFE and ROM predictions

5.1.1 Macroscopic stress-strain curves

The macroscopic stress-strain curves obtained from CPFE (orange line) and ROM simulations under monotonic tensile loading conditions (black line) are shown in Figure 5a. The gray line shows the response of the ROM including the unloading steps used to extract the microscale data at steps S1 through S6. In the elastic region, a reasonable agreement is observed between the stress-strain curves as shown in Figure 5b. The homogenized elastic moduli in the Y direction are given as ~ 213 GPa and ~ 198 GPa when employing ROM and CPFE, respectively. In the viscoplastic region, the ROM exhibits higher stress levels compared to the CPFE simulation, consistent with results available in the literature for reduced order models [21, 75]. In a verification study over a smaller microstructure, Ref. [32] reports an error of approximately 10% at peak stress levels, whereas in the present study the error is approximately 17%. This discrepancy could be partly attributed to differences in the way the microstructure was preprocessed and meshed for the ROM and CPFE simulations. Differences in preprocessing of the microstructure influences the grain morphologies and crystal orientations and, by extension, the accuracy of localization and interaction tensors. Influence of the relative tolerance used for the PCG algorithm in ROM construction showed negligible effects on the ROM results. The larger discrepancies observed in the present study are also due to the use

¹Specific vendor and manufacturer names are explicitly mentioned only to accurately describe the analytical tools used. The use of vendor and manufacturer names does not imply an endorsement by the authors nor does it imply that the specified equipment is the best available.




Figure 4: (a) 28 challenge grains inside the microstructure, (b) challenge microstructure, (c) macroscale boundary conditions and d) pole figure showing orientation distribution.

of quadratic tetrahedral elements in CPFE simulations (linear tetrahedral elements were used in Ref. [32]). Knezevic and coworkers [76] have demonstrated that overly stiff behavior for microscale stresses is exhibited by linear tetrahedral elements in CPFE simulations which aligns better with the stiffer response predicted by the ROM. Accuracy in the viscoplastic region could be improved by increasing the number of parts in the reduced order model or using techniques for recomputing localization and interaction tensors [75, 77] at the expense of additional computational cost.

5.1.2 Grain-averaged elastic strains

Grain-averaged elastic strain components extracted at load states S1–S6 for the 28 challenge grains are shown in Figure 6. Results are compared to the CPFE elastic strains as scatter plots where each point represents a single challenge grain. Each color represents a strain component and each marker represents a load state. During load states S1–S3 the specimen is still in the elastic regime with relatively small viscoplastic deformation. At S4–S6, the specimen has undergone relatively significant viscoplastic strain. A point that lies along the solid black line represents ideal agreement between the elastic component of the strain computed by the ROM and CPFE. The dashed lines are two standard deviations away from the solid black line. Overall, there is a reasonable agreement between the ROM and CPFE strains. For longitudinal grain-wise elastic strain (ϵ_{yy}), very good agreement is observed for load states S1–S3 and a slight overestimation of the elastic strains by the ROM is observed for load states S4–S6. Overestimation of elastic strains is related to the higher stress levels predicted by the ROM in the viscoplastic region. For lateral grain-averaged elastic strains (ϵ_{xx} and ϵ_{zz}), at load states S4–S6, the results are more spread out, indicating higher relative errors in lateral strains compared with the longitudinal strains. The magnitude of absolute errors are similar for all strain components. R^2 values for each load state and strain component are plotted in Figure 6b. R^2 values were computed by taking the square of the correlation coefficient between the ROM and reference CPFE data. R^2 values for all strain components are greater than 0.8 for elastic load states S1–S3. For load states S4–S6 with viscoplastic strains, R^2 values for longitudinal strains are also higher than 0.8. R^2 values for lateral strains drop to around 0.6 for viscoplastic load states S4–S6. It is worth noting that the R^2 values for shear strain XZ (perpendicular to the tension axis) are lower than those for all other strain components. One possible explanation for this is that the XZ shear magnitudes are small ($\sim 10^{-5}$) in CPFE simulation, while they are overestimated ($\sim 10^{-4}$) in ROM simulation, thus exacerbating the errors. Difference in the shear strain magnitudes may originate from periodicity assumptions used in the solution of IFPs.

5.1.3 Grain-averaged microscale stress distributions

Histograms of grain-averaged microscale stress components for all grains are compared in Figure 7 at 1% macroscopic strain. For tensile stress components, stresses computed with CPFE are more concentrated compared to the stresses from the ROM simulation. For lateral stresses (σ_{xx} and σ_{zz}), both ROM and CPFE distributions are centered at zero stress. Individual grains can experience non-zero transverse stresses due to anisotropy induced by orientation of crystals. The CPFE simulation indicates a smaller deviation in the stress components from the mean compared with the ROM results. However, the deviation is less pronounced for shear strains and better agreement is observed. Consistent with the homogenized stresses, grain-averaged stresses from the ROM simulation show stiffer response compared to the CPFE results.

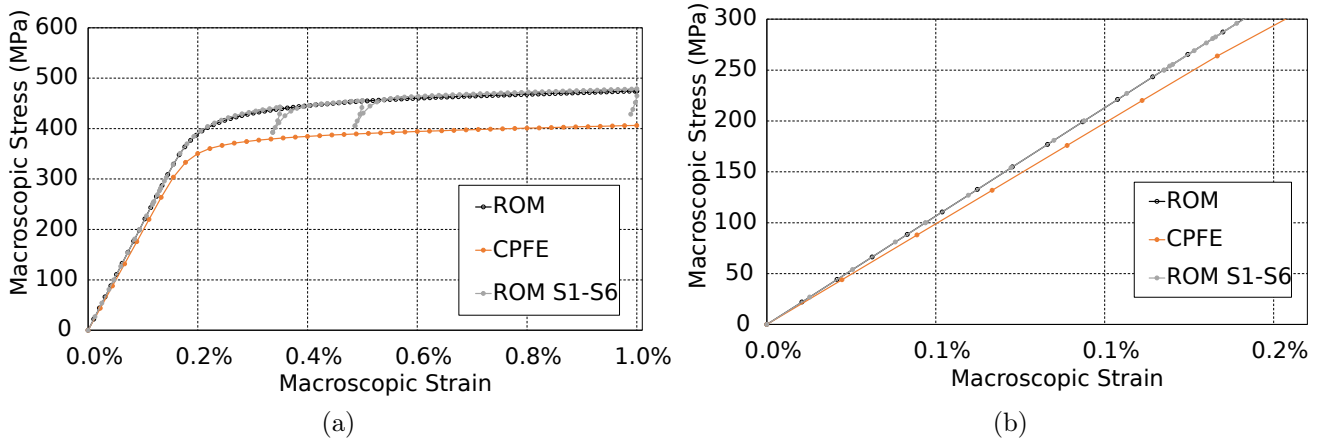


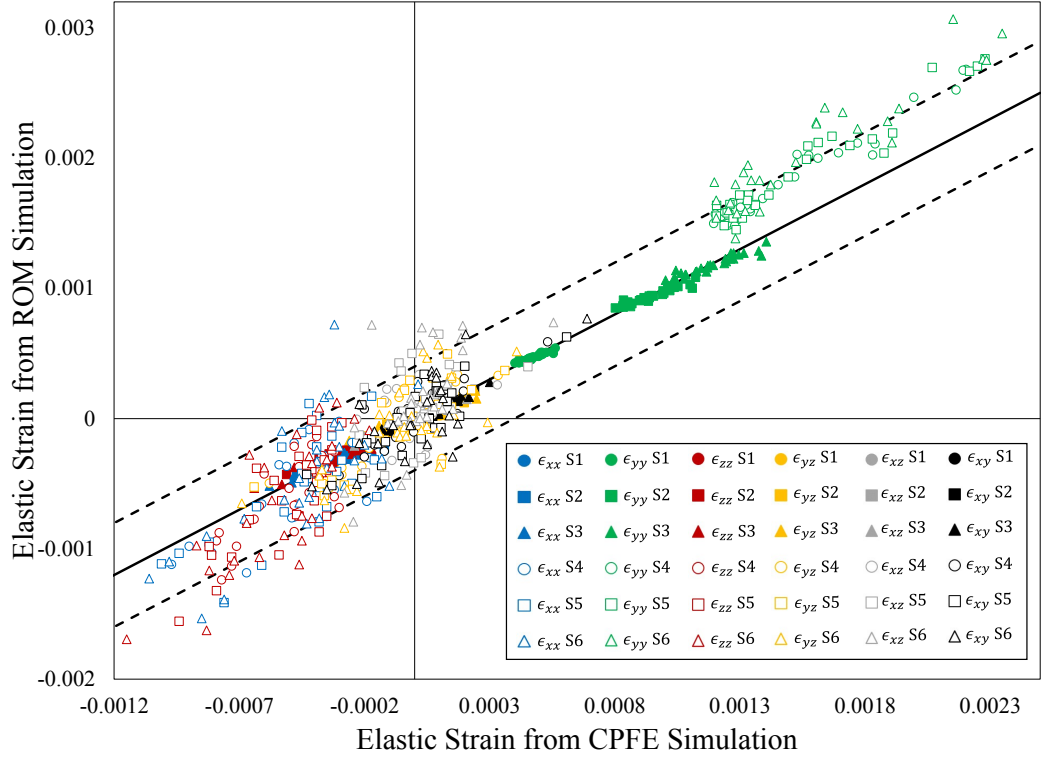
Figure 5: (a) Comparison of stress-strain graphs from CPFE and ROM simulations and (b) stress-strain response in the elastic regime.

5.1.4 Grain-averaged Fatemie-Socie parameter


The ability of the ROM in predicting quantities of interest that are complex functions of stress and strain measures is further assessed. Fatemi-Socie (FS) parameter is a measure of stress, strain, and internal state variables, that is used to indicate initiation of fatigue failure [78]. Accurate characterization of such complex functions of response fields with the ROM offers the potential for accelerated assessment of failure initiation [1, 54]. In this study, we employ the crystal-scale version of the FS parameter [79] for verification purposes. We note that the current study does not attempt to characterize fatigue behavior of the microstructure, but rather assess accuracy of the ROM in capturing quantities of interest. The FS parameter for a slip plane p is given by the following formula [80]

$$FS^p = \int_0^t \sum_{\alpha=1}^{N_d} |\dot{\gamma}^{\alpha,p}| \left(1 + k \frac{\langle \sigma_n^p \rangle}{g_0}\right) dt \quad (26)$$

where N_d is the number of slip directions on a given slip plane p , $k = 0.5$ is a parameter dictating the weight of tensile stress, and $\langle \sigma_n^p \rangle$ is tensile stress projected on slip plane p . Macaulay brackets $\langle \cdot \rangle$ are defined as $\langle \sigma_n^p \rangle = 0$ if $\sigma_n^p \leq 0$ and $\langle \sigma_n^p \rangle = \sigma_n^p$ if $\sigma_n^p \geq 0$. FS parameter for each slip plane extracted from the ROM and CPFE simulations is compared in Figure 8, where each color corresponds to a slip plane for FCC crystals (111), $(\bar{1}11)$, $(1\bar{1}1)$, and $(\bar{1}\bar{1}1)$. For each slip plane there are 28 data points representing grain-averaged FS parameter of the challenge grains at load state S6. The R^2 values were calculated as 0.651, 0.830, 0.919, and 0.935 for slip planes (111), $(\bar{1}11)$, $(1\bar{1}1)$, and $(\bar{1}\bar{1}1)$, respectively. As indicated by the R^2 values, FS parameter from the ROM captures the general trend of the FS parameter from the CPFE simulation. The discrepancies between the reference and ROM simulations are attributed to higher stress estimates in viscoplasticity observed in Figure 5 for ROM, since stresses directly affect the FS parameter as shown in Eq. 26.



(a)



(b)

Figure 6: Grain-averaged elastic strain components are computed for the 28 challenge grains over load states S1–S6 as shown in a). R^2 values computed for each load state for each strain component are provided in b).

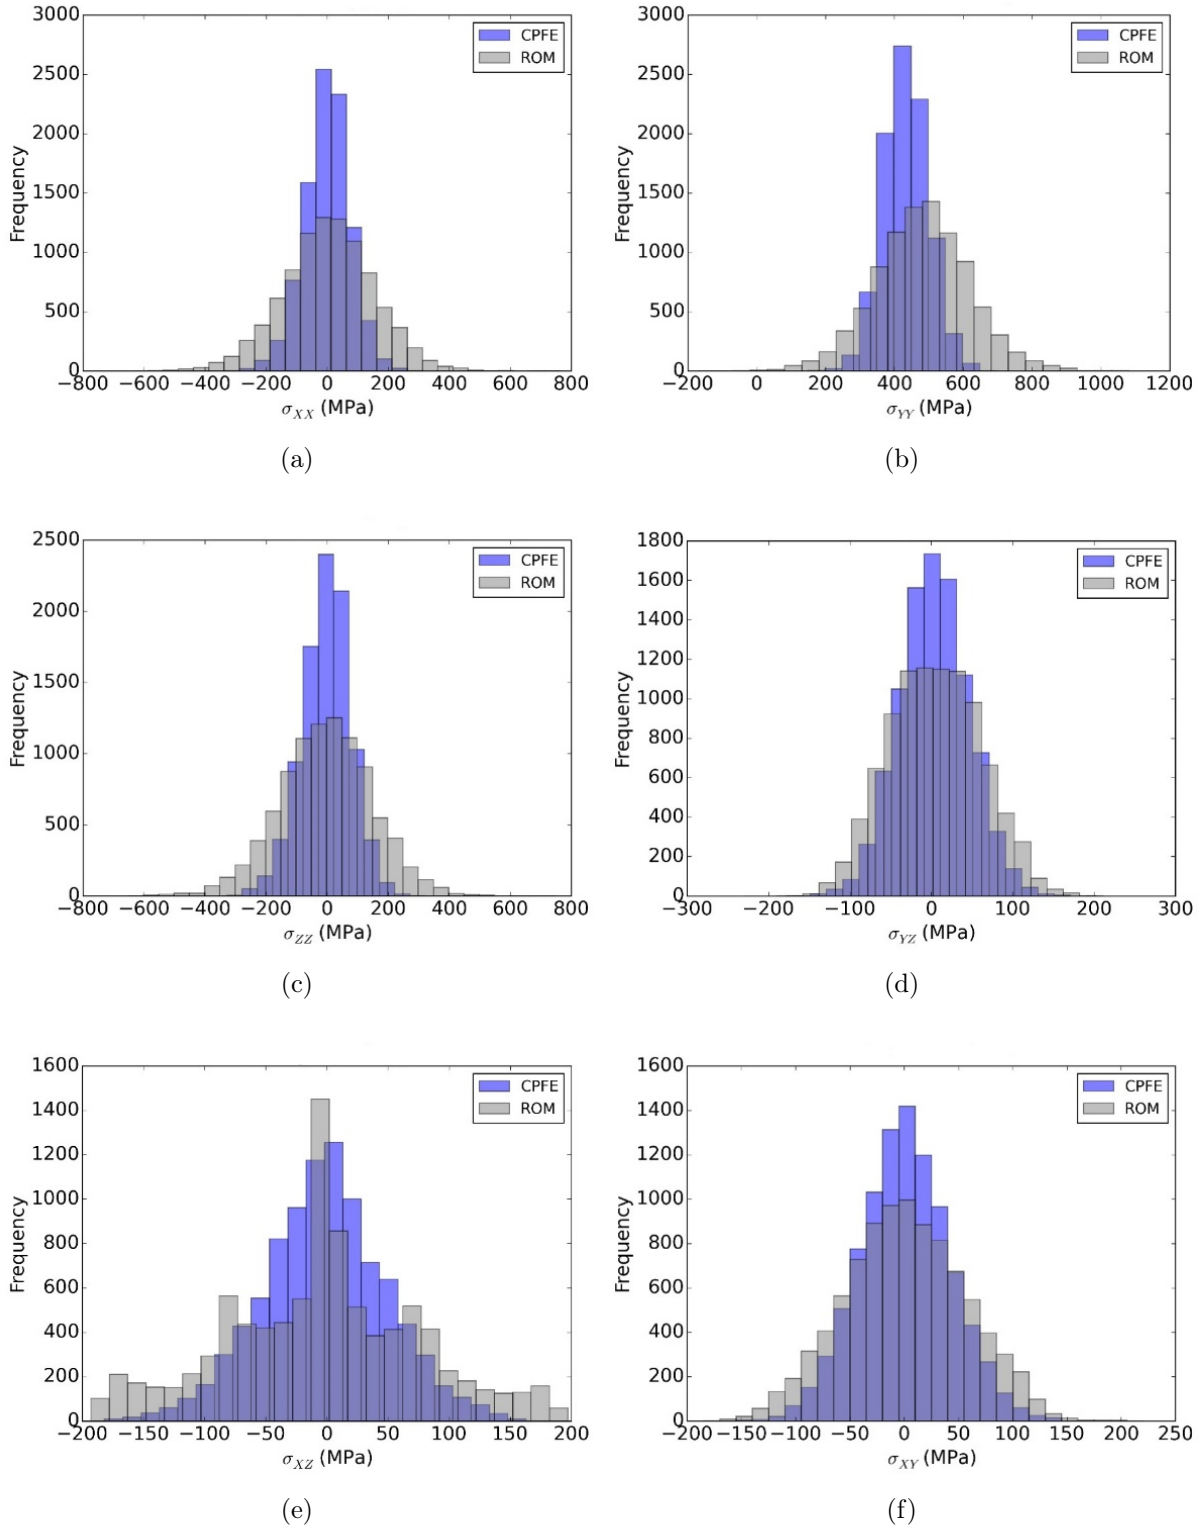


Figure 7: Histogram of grain-averaged stress components a) XX, b) YY, c) ZZ, d) YZ, e) XZ, and f) XY for all grains at 1% macroscopic strain.

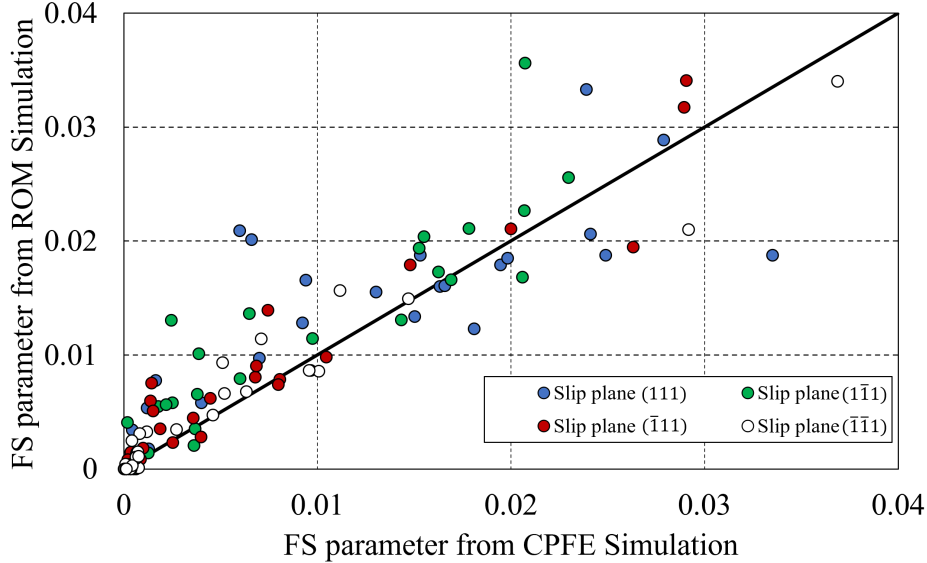


Figure 8: Scatter plot comparing FS parameter for 28 challenge grains at load state S6 for slip planes (111), ($\bar{1}\bar{1}\bar{1}$), ($1\bar{1}\bar{1}$), and ($\bar{1}\bar{1}1$).

5.1.5 Grain-averaged von Mises stress-equivalent strain plots

The grain-averaged von Mises stress vs. grain-averaged equivalent strain plots for two challenge grains (grain IDs 5876 and 21698) are shown in Figure 9. The two grains were selected from the set of challenge grains as they exhibit the highest and lowest stresses according to the reference CPFE simulations. Behavior of these two grains are therefore the upper and lower bounds of the grain-scale microscale stress-strain responses of all challenge grains within the microstructure. The elastic responses of the ROM and CPFE curves are in good agreement. The stress magnitudes in viscoplasticity from the ROM are slightly higher than their CPFE counterparts. Overall, the ROM is able to accurately pick out the extreme values of grain-averaged stresses and strains within the microstructure.

5.2 Scalability of the ROM construction framework

In order to test the performance of the ROM construction, a set of numerical experiments was performed using the challenge problem microstructure. The resolution of the mesh used in ROM construction was varied from 2 μm to 5 μm with an increment of one, producing four meshes with 392.8, 116.6, 49.5 (used for the ROM simulation in previous section), and 25.4 million degrees of freedom, respectively. Practically, the number of degrees of freedom is slightly lower since reported numbers include degrees of freedom associated with the replica and vertex nodes. To obtain the average computational cost of solving an IFP problem (see Algorithm 1), ten IFPs were solved at any given resolution, and the simulation times were then averaged. Computational costs as a function of problem size for relative tolerances of 10^{-1} , 10^{-2} , 10^{-3} , and 10^{-4} are reported in Figure 10. An Intel Xeon CPU E5-2683 v4 @ 2.10 GHz processor with 32 physical cores was used for the simulations. Solving a system with 100 million degrees of freedom to 10^{-3} relative tolerance requires approximately 10 minutes. The slopes of the four lines are given by 1.0127, 1.1644, 1.2163, and 1.2846




Figure 9: Comparison of the history of microscale von Mises stress against equivalent strain for challenge grains 5876 and 21698.

for relative tolerances (δ in Algorithm 1) of 10^{-1} , 10^{-2} , 10^{-3} , and 10^{-4} , respectively. It follows from the slopes that relative tolerance has significant effect on the total simulation time and simulation time increases with increasing relative tolerance. Near linear scaling is observed with respect to number of degrees of freedom for tolerance of 10^{-1} . However, deviation from linearity increases with increasing relative tolerance as indicated by increasing slopes. Due to the independence of the phase IFPs, the scaling of the code with respect to number of IFPs is linear.

5.3 Scalability of the proposed ROM evaluation algorithm

In order to test the scalability of the parallel implementation of the ROM evaluation algorithm, a set of numerical experiments was performed using the coefficient tensors computed for the challenge microstructure. The implementation of Algorithm 2 (together with Algorithms 3 and 4) was compiled with the Intel ifort 2021.3.0 compiler using O3 flag on Intel Xeon CPU E5-2683 v4 @ 2.10 GHz processor with 32 physical cores. To simulate a tensile test with constant strain rate, the following macroscopic strain history was applied : $\Delta\bar{\epsilon}_{xx} = \Delta t \dot{\bar{\epsilon}}$, $\Delta\bar{\epsilon}_{yy} = -\nu\Delta\bar{\epsilon}_{xx}$, and $\Delta\bar{\epsilon}_{zz} = -\nu\Delta\bar{\epsilon}_{xx}$ where $\Delta\bar{\epsilon}$ is the strain increment, Δt is the time increment, $\dot{\bar{\epsilon}}$ is the strain rate, ν is a parameter controlling lateral strains. ν was set to 0.5 to represent volume-preserving tension. The time increment (Δt) was set to 1 second and the strain rate of $10^{-4}/s$ was used.

To test the parallel implementation of ROM evaluation algorithm, simulations were performed for 100 and 1,000 increments with $\bar{\epsilon}_{xx}$ reaching 1% strain and 10% strain, respectively. Varying number of cores, time spent on the linear solver, total simulation time, and ratio of time spent on the linear solver to total simulation time are shown in Table 2. For 100 increment simulations, total simulation time decreases rapidly as the number of cores is increased, reaching 55 seconds for 16 cores and 45 seconds for 32 cores. The rightmost column presents the ratio of the linear solver computations to total time. For 1,000 increment simulations, it takes almost a half an hour to run the simulation in the serial setting. However, as the number of cores was increased to 32 the total simulation time drops to 3 minutes. For both 100 and 1000 increment simulations, as the number of cores increases from 1 to 32, the share of time spent on the linear solver is within 70%–80% and

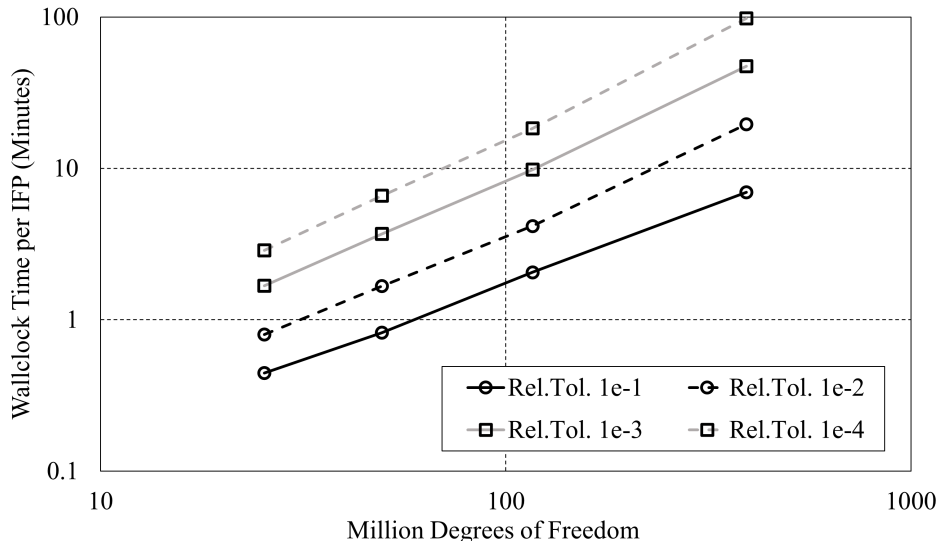


Figure 10: Cost of solving one IFP vs. number of degrees of freedom with varying relative tolerance of the PCG solver.

only 20%–30% of the total time is spent on non-solver related computation.

In both 100 and 1,000 increment tests, the total time was reduced by almost ten times when the number of physical cores reached 32. The total simulation time is plotted against number of cores in Figure 11. Linear scaling with respect to the number of cores is observed up to 8 cores, after which the performance flattens out slightly. The lack of linear scaling can be attributed to saturation of memory bandwidth because of memory intensive numerical decomposition performed by the linear solver. However, since the share of linear solver related computations remains consistently high with varying the number of physical cores, further improvements to other subroutines will yield diminishing returns. Therefore, scaling improvements would require either changes to the solver strategy or reduction of total solver calls. One potential approach to minimize solver time would be to use a fully-coupled Newton-Raphson scheme to solve for both stresses and strengths at the same time [81]. This scheme could reduce the number of microscale iterations needed to converge which would reduce number of linear solver calls. Alternatively, sparse iterative methods could be experimented with to avoid the large cost of numerical decomposition of the Jacobian matrix.

6 Conclusion

In this manuscript, parallel implementation algorithms were presented for construction and evaluation of ROMs developed, based on the EHM approach. The scalability of the implementations was evaluated for both the construction and evaluation stages. For ROM construction, near linear scaling has been achieved with respect to the number of degrees of freedom in the mesh and linear scaling with respect to the number of phase IFPs follows from EHM formulation. Additionally, it was shown that increasing the number of cores for ROM evaluation led to significant improvements in simulation time, with nearly an order of magnitude improvement observed as the number of cores was increased to 32. These algorithms open the possibility of efficiently evaluating the mechanical response of

Table 2: Scaling results for 100 and 1,000 increment parallel ROM simulations.

# of increments	# of cores	Time spent on linear solver (s)	Total time (s)	Ratio of linear solve/total time
100	1	276.8	355.8	78%
	2	154.1	202.0	76%
	4	90.4	120.0	75%
	8	56.5	76.8	74%
	16	38.6	53.9	72%
	32	32.8	45.5	72%
1000	1	1,110.1	1,493.3	74%
	2	626.7	848.6	74%
	4	386.3	515.5	75%
	8	236.7	318.9	74%
	16	171.4	228.5	75%
	32	148.8	191.2	78%

large microstructural volumes or microstructural volumes that contain a large number of complex subgrain features. The proposed implementations were verified against CPFEM simulations using a large microstructure previously investigated in an AFRL study. The comparison between the ROM and CPFEM simulations showed reasonable agreement in terms of macroscale stress-strain behavior and microscale quantities of interest. However, it was observed that the ROM exhibited a stiffer response compared to the reference CPFEM simulations. The discrepancies could be attributed to two main factors: the constrained kinematics due to the assumption of uniform fields (strain, stress, and internal state variables) within the subdomains associated with reduced order parts (grains in this particular study), and the possible locking phenomenon observed in the ROM due to the effects of incompressibility in the viscoplastic flow regime. The occurrence of locking phenomenon is well-known in CPFEM simulations when first order finite elements are used, but this phenomenon has not yet been investigated thoroughly in the context of reduced order modeling. Further improvements will be investigated in the future to enhance the accuracy of the ROM in the viscoplastic regime. Another potential direction for future research is to extend the ROM evaluation to a multiscale solver by leveraging the hierarchical nature of multiscale methods. This multiscale approach could involve coupling parallel domain decomposition methods for the macroscale solver with a parallel ROM evaluation algorithm, enabling scalable multiscale simulations of complex systems beyond the scope of academic benchmarks.

7 Acknowledgements

The authors gratefully acknowledge the financial support from the National Aeronautics and Space Administration (NASA), Space Technology Research Grants (Early State Innovation Grant Nos.: 80NSSC20K0294). Saikumar Yeratapally was sponsored through a NASA cooperative agreement NNL09AA00A with the National Institute of Aerospace (NIA) and through the Research, Engineering,




Figure 11: Scaling of the parallel ROM evaluation algorithm.

and Science Services (RSES) contract at NASA Langley Research Center. The authors would also like to acknowledge the Transformational Tools and Technologies (TTT) Project in the Aeronautics Research Mission Directorate (ARMD) for supporting this work.

References

- [1] X. Zhang, Y. Liu, and C. Oskay. “Uncertainty quantification for microstructure-sensitive fatigue nucleation and application to Titanium alloy, Ti6242”. In: *Frontiers in Materials* 9, 897998 (May 2022), p. 897998. DOI: 10.3389/fmats.2022.897998.
- [2] M. Knezevic and D. J. Savage. “A high-performance computational framework for fast crystal plasticity simulations”. In: *Computational Materials Science* 83 (2014), pp. 101–106. ISSN: 0927-0256. DOI: <https://doi.org/10.1016/j.commatsci.2013.11.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0927025613006848>.
- [3] F. Roters et al. “Overview of constitutive laws, kinematics, homogenization and multiscale methods in crystal plasticity finite-element modeling: Theory, experiments, applications”. In: *Acta materialia* 58.4 (2010), pp. 1152–1211.
- [4] X. Zhang and C. Oskay. “Plastic dissipation sensitivity to mechanical properties in polycrystalline β -HMX subjected to impact loading”. In: *Mech. Mater.* 138 (2019), p. 103079.
- [5] X. Zhang and C. Oskay. “Modeling and numerical investigation of mechanical twinning in β -HMX crystals subjected to shock loading”. In: *Modelling and Simulation in Materials Science and Engineering* 29.7 (2021), p. 075009. DOI: 10.1088/1361-651X/ac21a6. URL: <https://dx.doi.org/10.1088/1361-651X/ac21a6>.

- [6] Rahul and S. De. “Multiscale modeling of irradiated polycrystalline FCC metals”. In: *International Journal of Solids and Structures* 51.23 (2014), pp. 3919–3930. ISSN: 0020-7683. DOI: <https://doi.org/10.1016/j.ijsolstr.2014.07.015>. URL: <https://www.sciencedirect.com/science/article/pii/S0020768314002893>.
- [7] V. -T. Phan et al. “Microscale modeling of creep deformation and rupture in Nickel-based superalloy IN 617 at high temperature”. In: *Mech. Mater.* 114 (2017), pp. 215–227.
- [8] X. Zhang and C. Oskay. “Polycrystal plasticity modeling of nickel-based superalloy IN 617 subjected to cyclic loading at high temperature”. In: *Modelling Simul. Mater. Sci. Eng.* 24.055009 (2016).
- [9] Y. Liu et al. “Experimental and computational study of microstructural effect on ductile fracture of hot-forming materials”. In: *Mater. Sci. Eng. A* 724 (2018), pp. 298–323.
- [10] M. Mosby and K. Matouš. “Computational homogenization at extreme scales”. In: *Extreme Mechanics Letters* 6 (Mar. 2016). ISSN: 2352-4316. DOI: 10.1016/j.eml.2015.12.009. URL: <https://www.osti.gov/biblio/1565430>.
- [11] M. Mosby and K. Matouš. “Hierarchically parallel coupled finite strain multiscale solver for modeling heterogeneous layers”. In: *International Journal for Numerical Methods in Engineering* 102.3-4 (2015), pp. 748–765. DOI: <https://doi.org/10.1002/nme.4755>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.4755>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4755>.
- [12] M. G. D. Geers et al. “Homogenization methods and multiscale modeling: Nonlinear problems”. In: *Encyclopedia of Computational Mechanics Second Edition*. John Wiley & Sons, Ltd, 2017, pp. 1–34. ISBN: 9781119176817. DOI: <https://doi.org/10.1002/9781119176817.ecm2107>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119176817.ecm2107>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781119176817.ecm2107>.
- [13] G.I. Taylor. “Plastic strain in metals”. In: *Journal of Institute of Metals* 62 (1938), pp. 307–24.
- [14] G. Sachs. “Zur ableitung einer fließbedingung”. In: *Z VDI* 72 (1928), pp. 734–6.
- [15] M. Knezevic, H. F. Al-Harbi, and S. R. Kalidindi. “Crystal plasticity simulations using discrete Fourier transforms”. In: *Acta Materialia* 57.6 (Apr. 2009), pp. 1777–1784. DOI: 10.1016/j.actamat.2008.12.017.
- [16] D. J. Savage and M. Knezevic. “Computer implementations of iterative and non-iterative crystal plasticity solvers on high performance graphics hardware”. In: *Computational Mechanics* 56.4 (2015), pp. 677–690.
- [17] R.A. Lebensohn and C.N. Tomé. “A self-consistent anisotropic approach for the simulation of plastic deformation and texture development of polycrystals: Application to zirconium alloys”. In: *Acta Metallurgica et Materialia* 41.9 (1993), pp. 2611–2624. ISSN: 0956-7151. DOI: [https://doi.org/10.1016/0956-7151\(93\)90130-K](https://doi.org/10.1016/0956-7151(93)90130-K). URL: <https://www.sciencedirect.com/science/article/pii/095671519390130K>.
- [18] G. J. Dvorak and Y. Benveniste. “On transformation strains and uniform fields in multiphase elastic media”. In: *Proceedings: Mathematical and Physical Sciences* 437.1900 (1992), pp. 291–310. ISSN: 09628444. URL: <http://www.jstor.org/stable/52199> (visited on 08/22/2022).

- [19] G. J. Dvorak. “Transformation field analysis of inelastic composite materials”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 437.1900 (1992), pp. 311–327. DOI: [10.1098/rspa.1992.0063](https://doi.org/10.1098/rspa.1992.0063). eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1992.0063>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1992.0063>.
- [20] J.C. Michel and P. Suquet. “Computational analysis of nonlinear composite structures using the nonuniform transformation field analysis”. In: *Computer Methods in Applied Mechanics and Engineering* 193.48 (2004). Advances in Computational Plasticity, pp. 5477–5502. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2003.12.071>. URL: <https://www.sciencedirect.com/science/article/pii/S004578250400283X>.
- [21] J.C. Michel and P. Suquet. “Nonuniform transformation field analysis”. In: *International Journal of Solids and Structures* 40.25 (2003). Special issue in Honor of George J. Dvorak, pp. 6937–6955. ISSN: 0020-7683. DOI: [https://doi.org/10.1016/S0020-7683\(03\)00346-9](https://doi.org/10.1016/S0020-7683(03)00346-9). URL: <https://www.sciencedirect.com/science/article/pii/S0020768303003469>.
- [22] F. Covezzi et al. “Comparison of reduced order homogenization techniques: pRBMOR, NUTFA and MxTFA”. In: *Meccanica* 53 (2018), pp. 1291–1312.
- [23] F. Fritzen, S. Marfia, and V. Sepe. “Reduced order modeling in nonlinear homogenization: A comparative study”. In: *Computers & Structures* 157 (2015), pp. 114–131. ISSN: 0045-7949. DOI: <https://doi.org/10.1016/j.compstruc.2015.05.012>. URL: <https://www.sciencedirect.com/science/article/pii/S0045794915001492>.
- [24] Felix Fritzen and Matthias Leuschner. “Reduced basis hybrid computational homogenization based on a mixed incremental formulation”. In: *Computer Methods in Applied Mechanics and Engineering* 260 (2013), pp. 143–154. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2013.03.007>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782513000583>.
- [25] V. Sepe, S. Marfia, and E. Sacco. “A nonuniform TFA homogenization technique based on piecewise interpolation functions of the inelastic field”. In: *International Journal of Solids and Structures* 50.5 (2013), pp. 725–742. ISSN: 0020-7683. DOI: <https://doi.org/10.1016/j.ijsolstr.2012.11.005>. URL: <https://www.sciencedirect.com/science/article/pii/S0020768312004684>.
- [26] F. Covezzi et al. “Homogenization of elastic–viscoplastic composites by the Mixed TFA”. In: *Computer Methods in Applied Mechanics and Engineering* 318 (2017), pp. 701–723. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2017.02.009>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782516308635>.
- [27] J. Yvonnet, D. Gonzalez, and Q.-C. He. “Numerically explicit potentials for the homogenization of nonlinear elastic heterogeneous materials”. In: *Computer Methods in Applied Mechanics and Engineering* 198.33 (2009), pp. 2723–2737. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2009.03.017>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782509001492>.

- [28] Z. Liu, M.A. Bessa, and W. K. Liu. “Self-consistent clustering analysis: An efficient multi-scale scheme for inelastic heterogeneous materials”. In: *Computer Methods in Applied Mechanics and Engineering* 306 (2016), pp. 319–341. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2016.04.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782516301499>.
- [29] C. Oskay and J. Fish. “Eigendeformation-based reduced order homogenization for failure analysis of heterogeneous materials”. In: *Computer Methods in Applied Mechanics and Engineering* 196.7 (2007), pp. 1216–1243. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2006.08.015>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782506002933>.
- [30] C. Oskay and G. Pal. “A Multiscale Failure Model for Analysis of Thin Heterogeneous Plates”. In: *Int. J. Damage Mechanics* 19 (2010), pp. 575–611.
- [31] J. Fish, V. Filonova, and Z. Yuan. “Reduced order computational continua”. In: *Computer Methods in Applied Mechanics and Engineering* 221-222 (2012), pp. 104–116. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2012.02.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782512000527>.
- [32] X. Zhang and C. Oskay. “Eigenstrain based reduced order homogenization for polycrystalline materials”. In: *Computer Methods in Applied Mechanics and Engineering* 297 (2015), pp. 408–436. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2015.09.006>. URL: <https://www.sciencedirect.com/science/article/pii/S004578251500300X>.
- [33] D. Xia, X. Zhang, and C. Oskay. “Large-deformation reduced order homogenization of polycrystalline materials”. In: *Computer Methods in Applied Mechanics and Engineering* 387 (2021), p. 114119. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2021.114119>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782521004503>.
- [34] D. Xia and C. Oskay. “Reduced order mathematical homogenization method for polycrystalline microstructure with microstructurally small cracks”. In: *Int. J. Numer. Methods Eng.* 124 (2023), pp. 3166–3190.
- [35] C. Oskay, Z. Su, and B. Kapusuzoglu. “Discrete eigenseparation-based reduced order homogenization method for failure modeling of composite materials”. In: *Comput. Meth. Appl. Mech. Engrg.* 359 (2020), p. 112656.
- [36] Z. Su and C. Oskay. “Mesh size objective fatigue damage propagation in laminated composites using the multiscale discrete damage theory”. In: *Comput. Mech.* 67 (2021), pp. 969–987.
- [37] P. Krysl, S. Lall, and J. E. Marsden. “Dimensional model reduction in non-linear finite element dynamics of solids and structures”. In: *International Journal for Numerical Methods in Engineering* 51.4 (2001), pp. 479–504. DOI: <https://doi.org/10.1002/nme.167>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.167>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.167>.
- [38] J. Yvonnet and Q.-C. He. “The reduced model multiscale method (R3M) for the non-linear homogenization of hyperelastic media at finite strains”. In: *Journal of Computational Physics* 223.1 (2007), pp. 341–368. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2006.09.019>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999106004402>.

- [39] R. Hatano et al. “FE^{2r} method with surrogate localization model for hyperelastic composite materials”. In: *Advanced Modeling and Simulation in Engineering Sciences* 7.1 (2020), p. 39. ISSN: 2213-7467. DOI: 10.1186/s40323-020-00175-0. URL: <https://doi.org/10.1186/s40323-020-00175-0>.
- [40] F. Chinesta and P. Ladevèze. “3 Proper generalized decomposition”. In: *Volume 2: Snapshot-Based Methods and Algorithms*. Ed. by Peter Benner et al. Berlin, Boston: De Gruyter, 2021, pp. 97–138. ISBN: 9783110671490. DOI: doi:10.1515/9783110671490-003. URL: <https://doi.org/10.1515/9783110671490-003>.
- [41] Y. Liu et al. “Dislocation density informed eigenstrain based reduced order homogenization modeling: verification and application on a titanium alloy structure subjected to cyclic loading”. In: *Modelling and Simulation in Materials Science and Engineering* 28.2 (2020), p. 025004.
- [42] F. Fritzen, M. Hodapp, and M. Leuschner. “GPU accelerated computational homogenization based on a variational approach in a reduced basis framework”. In: *Computer Methods in Applied Mechanics and Engineering* 278 (2014), pp. 186–217. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2014.05.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782514001510>.
- [43] F. Fritzen and M. Hodapp. “The finite element square reduced (FE2R) method with GPU acceleration: towards three-dimensional two-scale simulations”. In: *International Journal for Numerical Methods in Engineering* 107.10 (2016), pp. 853–881. DOI: <https://doi.org/10.1002/nme.5188>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.5188>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.5188>.
- [44] F. Feyel and J. L. Chaboche. “FE2 multiscale approach for modelling the elastoviscoplastic behaviour of long fibre SiC/Ti composite materials”. In: *Computer Methods in Applied Mechanics and Engineering* 183 (2000), pp. 309–330.
- [45] H. Moulinec and P. Suquet. “A numerical method for computing the overall response of nonlinear composites with complex microstructure”. In: *Computer Methods in Applied Mechanics and Engineering* 157.1 (1998), pp. 69–94. ISSN: 0045-7825. DOI: [https://doi.org/10.1016/S0045-7825\(97\)00218-1](https://doi.org/10.1016/S0045-7825(97)00218-1). URL: <https://www.sciencedirect.com/science/article/pii/S0045782597002181>.
- [46] J.C. Michel, H. Moulinec, and P. Suquet. “Effective properties of composite materials with periodic microstructure: a computational approach”. In: *Computer Methods in Applied Mechanics and Engineering* 172.1 (1999), pp. 109–143. ISSN: 0045-7825. DOI: [https://doi.org/10.1016/S0045-7825\(98\)00227-8](https://doi.org/10.1016/S0045-7825(98)00227-8). URL: <https://www.sciencedirect.com/science/article/pii/S0045782598002278>.
- [47] J. Michel, H. Moulinec, and P. Suquet. “A computational scheme for linear and non-linear composites with arbitrary phase contrast”. In: *International Journal for Numerical Methods in Engineering* 52 (Sept. 2001), pp. 139–160. DOI: 10.1002/nme.275.
- [48] R. A. Lebensohn and A. D. Rollett. “Spectral methods for full-field micromechanical modelling of polycrystalline materials”. In: *Computational Materials Science* 173 (2020), p. 109336. ISSN: 0927-0256. DOI: <https://doi.org/10.1016/j.commatsci.2019.109336>. URL: <https://www.sciencedirect.com/science/article/pii/S0927025619306354>.

- [49] B. A. Le, J. Yvonnet, and Q.-C. He. “Computational homogenization of nonlinear elastic materials using neural networks”. In: *International Journal for Numerical Methods in Engineering* 104.12 (2015), pp. 1061–1084. DOI: <https://doi.org/10.1002/nme.4953>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.4953>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.4953>.
- [50] B. Liu et al. “A learning-based multiscale method and its application to inelastic impact problems”. In: *Journal of the Mechanics and Physics of Solids* 158 (2022), p. 104668. ISSN: 0022-5096. DOI: <https://doi.org/10.1016/j.jmps.2021.104668>. URL: <https://www.sciencedirect.com/science/article/pii/S0022509621002982>.
- [51] N. N. Vlassis, R. Ma, and W. Sun. “Geometric deep learning for computational mechanics Part I: anisotropic hyperelasticity”. In: *Computer Methods in Applied Mechanics and Engineering* 371 (2020), p. 113299. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2020.113299>. URL: <https://www.sciencedirect.com/science/article/pii/S0045782520304849>.
- [52] J. N. Fuhg, M. Marino, and N. Bouklas. “Local approximate Gaussian process regression for data-driven constitutive models: development and comparison with neural networks”. In: *Computer Methods in Applied Mechanics and Engineering* 388 (2022), p. 114217. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2021.114217>. URL: <https://www.sciencedirect.com/science/article/pii/S004578252100548X>.
- [53] V. Minh Nguyen-Thanh et al. “A surrogate model for computational homogenization of elastostatics at finite strain using high-dimensional model representation-based neural network”. In: *International Journal for Numerical Methods in Engineering* 121.21 (2020), pp. 4811–4842. DOI: <https://doi.org/10.1002/nme.6493>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.6493>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.6493>.
- [54] Y. Liu, X. Zhang, and C. Oskay. “A comparative study on fatigue indicator parameters for near- α titanium alloys”. In: *Fatigue & Fracture of Engineering Materials & Structures* 46.1 (2023), pp. 271–294.
- [55] X. Zhang and C. Oskay. “Sparse and scalable eigenstrain-based reduced order homogenization models for polycrystal plasticity”. In: *Computer Methods in Applied Mechanics and Engineering* 326 (2017), pp. 241–269. ISSN: 0045-7825. DOI: <https://doi.org/10.1016/j.cma.2017.07.027>. URL: <https://www.sciencedirect.com/science/article/pii/S004578251730347X>.
- [56] D. B. Menasche et al. “AFRL additive manufacturing modeling series: Challenge 4, in situ mechanical test of an IN625 sample with concurrent high-energy diffraction microscopy characterization”. In: *Integrating Materials and Manufacturing Innovation* 10.3 (June 2021). DOI: 10.1007/s40192-021-00218-3.
- [57] *Air Force Research Laboratory (AFRL) additive manufacturing (AM) modeling challenge series*. <https://materials-data-facility.github.io/MID3AS-AM-Challenge/>. Accessed: 2022-09-10.

- [58] M. G. Chapman et al. “AFRL additive manufacturing modeling series: Challenge 4, 3D reconstruction of an IN625 high-energy diffraction microscopy sample using multi-modal serial sectioning”. In: *Integrating Materials and Manufacturing Innovation* 10.2 (2021), pp. 129–141. ISSN: 2193-9772. DOI: [10.1007/s40192-021-00212-9](https://doi.org/10.1007/s40192-021-00212-9). URL: <https://doi.org/10.1007/s40192-021-00212-9>.
- [59] J.R. Rice. “Inelastic constitutive relations for solids: An internal-variable theory and its application to metal plasticity”. In: *Journal of the Mechanics and Physics of Solids* 19.6 (1971), pp. 433–455. ISSN: 0022-5096. DOI: [https://doi.org/10.1016/0022-5096\(71\)90010-X](https://doi.org/10.1016/0022-5096(71)90010-X). URL: <https://www.sciencedirect.com/science/article/pii/002250967190010X>.
- [60] D. Pierce, R.J.Asaro, and A. Needleman. “An analysis of nonuniform and localized deformation in ductile single crystals”. In: *Acta Metallurgica* 30.6 (1982), pp. 1087–1119. ISSN: 0001-6160. DOI: [https://doi.org/10.1016/0001-6160\(82\)90005-0](https://doi.org/10.1016/0001-6160(82)90005-0). URL: <https://www.sciencedirect.com/science/article/pii/0001616082900050>.
- [61] D. Peirce, R.J. Asaro, and A. Needleman. “Material rate dependence and localized deformation in crystalline solids”. In: *Acta Metallurgica* 31.12 (1983), pp. 1951–1976. ISSN: 0001-6160. DOI: [https://doi.org/10.1016/0001-6160\(83\)90014-7](https://doi.org/10.1016/0001-6160(83)90014-7). URL: <https://www.sciencedirect.com/science/article/pii/0001616083900147>.
- [62] L. Anand. “Constitutive equations for the rate-dependent deformation of metals at elevated temperatures”. In: *Journal of Engineering Materials and Technology* 104.1 (Jan. 1982), pp. 12–17. ISSN: 0094-4289. DOI: 10.1115/1.3225028. eprint: https://asmedigitalcollection.asme.org/materialstechnology/article-pdf/104/1/12/5513557/12_1.pdf. URL: <https://doi.org/10.1115/1.3225028>.
- [63] E. Voce. “The relationship between stress and strain for homogeneous deformations”. In: 1948.
- [64] M. R Hestenes and E. Stiefel. “Methods of conjugate gradients for solving”. In: *Journal of research of the National Bureau of Standards* 49.6 (1952), p. 409.
- [65] T. J.R. Hughes, I. Levit, and J. Winget. “An element-by-element solution algorithm for problems of structural and solid mechanics”. In: *Computer Methods in Applied Mechanics and Engineering* 36.2 (1983), pp. 241–254. ISSN: 0045-7825. DOI: [https://doi.org/10.1016/0045-7825\(83\)90115-9](https://doi.org/10.1016/0045-7825(83)90115-9). URL: <https://www.sciencedirect.com/science/article/pii/0045782583901159>.
- [66] T. J.R. Hughes, R. M. Ferencz, and J. O. Hallquist. “Large-scale vectorized implicit calculations in solid mechanics on a Cray X-MP/48 utilizing EBE preconditioned conjugate gradients”. In: *Computer Methods in Applied Mechanics and Engineering* 61.2 (1987), pp. 215–248. ISSN: 0045-7825. DOI: [https://doi.org/10.1016/0045-7825\(87\)90005-3](https://doi.org/10.1016/0045-7825(87)90005-3). URL: <https://www.sciencedirect.com/science/article/pii/0045782587900053>.
- [67] OpenMP Architecture Review Board. *OpenMP Application Program Interface Version 3.0*. May 2008. URL: <http://www.openmp.org/mp-documents/spec30.pdf>.
- [68] NASA. *Femera mini-app*. URL: <https://github.com/waggyz/femera-mini-app>.

- [69] E.B. Marin and P.R. Dawson. “On modelling the elasto-viscoplastic response of metals using polycrystal plasticity”. In: *Computer Methods in Applied Mechanics and Engineering* 165.1 (1998), pp. 1–21. ISSN: 0045-7825. DOI: [https://doi.org/10.1016/S0045-7825\(98\)00034-6](https://doi.org/10.1016/S0045-7825(98)00034-6). URL: <https://www.sciencedirect.com/science/article/pii/S0045782598000346>.
- [70] *Intel Math Kernel Library. Reference Manual*. Intel Corporation, 2009.
- [71] M. Groeber and M. Jackson. “DREAM.3D: a digital representation environment for the analysis of microstructure in 3D”. In: *Integrating Materials and Manufacturing Innovation* 3 (Feb. 2014), p. 5. DOI: 10.1186/2193-9772-3-5.
- [72] S. R. Yeratapally, A. R. Cerrone, and E. H. Glaessgen. “Discrepancy between crystal plasticity simulations and far-field high-energy X-ray diffraction microscopy measurements”. In: *Integrating Materials and Manufacturing Innovation* 10.2 (2021), pp. 196–217.
- [73] F. Bachmann, Ralf Hielscher, and Helmut Schaeben. “Texture Analysis with MTEX – Free and Open Source Software Toolbox”. In: *Texture and Anisotropy of Polycrystals III*. Vol. 160. Solid State Phenomena. Trans Tech Publications Ltd, Mar. 2010, pp. 63–68. DOI: 10.4028/www.scientific.net/SSP.160.63.
- [74] M. Smith. *ABAQUS/Standard User’s Manual, Version 6.9*. English. United States: Dassault Systèmes Simulia Corp, 2009.
- [75] J.L. Chaboche et al. “Towards a micromechanics based inelastic and damage modeling of composites”. In: *International Journal of Plasticity* 17.4 (2001), pp. 411–439. ISSN: 0749-6419. DOI: [https://doi.org/10.1016/S0749-6419\(00\)00056-5](https://doi.org/10.1016/S0749-6419(00)00056-5). URL: <https://www.sciencedirect.com/science/article/pii/S0749641900000565>.
- [76] W. G. Feather, H. Lim, and M. Knezevic. “A Numerical Study into Element Type and Mesh Resolution for Crystal Plasticity Finite Element Modeling of Explicit Grain Structures”. In: *Comput. Mech.* 67.1 (2021), pp. 33–55. ISSN: 0178-7675. DOI: 10.1007/s00466-020-01918-x. URL: <https://doi.org/10.1007/s00466-020-01918-x>.
- [77] J.L. Chaboche, P. Kanoute, and A. Roos. “On the capabilities of mean-field approaches for the description of plasticity in metal matrix composites”. In: *International Journal of Plasticity* 21.7 (2005), pp. 1409–1434. ISSN: 0749-6419. DOI: <https://doi.org/10.1016/j.ijplas.2004.07.001>.
- [78] A. Fatemi and D. F. Socie. “A critical plane approach to multiaxial fatigue damage including out-of-phase loading”. In: *Fatigue & Fracture of Engineering Materials & Structures* 11.3 (1988), pp. 149–165. DOI: <https://doi.org/10.1111/j.1460-2695.1988.tb01169.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1460-2695.1988.tb01169.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-2695.1988.tb01169.x>.
- [79] Y. Liu, X. Zhang, and C. Oskay. “A comparative study on fatigue indicator parameters for near-titanium alloys”. In: *Fatigue Fract Eng Mater Struct.* 46 (2023), pp. 271–294.
- [80] J. Hochhalter. “Finite element simulations of fatigue crack stages in AA 7075-T651 microstructure”. PhD thesis. Cornell University, 2010.

- [81] F.T. Meissonnier, E.P. Busso, and N.P. O’Dowd. “Finite element implementation of a generalised non-local rate-dependent crystallographic formulation for finite strains”. In: *International Journal of Plasticity* 17.4 (2001), pp. 601–640. ISSN: 0749-6419. DOI: [https://doi.org/10.1016/S0749-6419\(00\)00064-4](https://doi.org/10.1016/S0749-6419(00)00064-4). URL: <https://www.sciencedirect.com/science/article/pii/S0749641900000644>.