

Accelerating crack growth simulations through adaptive model order reduction

Rocha, Iuri B.C.M.; van der Meer, Frans P.; Taumaturgo Mororo, Luiz Antonio; Sluys, Lambertus J.

DOI

[10.1002/nme.6303](https://doi.org/10.1002/nme.6303)

Publication date

2019

Document Version

Final published version

Published in

International Journal for Numerical Methods in Engineering

Citation (APA)

Rocha, I. B. C. M., van der Meer, F. P., Taumaturgo Mororo, L. A., & Sluys, L. J. (2019). Accelerating crack growth simulations through adaptive model order reduction. *International Journal for Numerical Methods in Engineering*, 121 (2020)(10), 2147-2173. <https://doi.org/10.1002/nme.6303>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

RESEARCH ARTICLE

WILEY

Accelerating crack growth simulations through adaptive model order reduction

Iuri. B. C. M. Rocha¹ | Frans. P. van der Meer² | Luiz. A. T. Mororó³ | Lambertus. J. Sluys⁴

Faculty of Civil Engineering and
Geosciences, Delft University of
Technology, Delft, The Netherlands

Correspondence

I. B. C. M. Rocha, Delft University of
Technology, Faculty of Civil Engineering
and Geosciences, P.O. Box 5048, 2600GA
Delft, The Netherlands.

Email:

i.barceloscarneiromrocha@tudelft.nl

Summary

Accurate numerical modeling of fracture in solids is a challenging undertaking that often involves the use of computationally demanding modeling frameworks. Model order reduction techniques can be used to alleviate the computational effort associated with these models. However, the traditional *offline-online* reduction approach is unsuitable for complex fracture phenomena due to their excessively large parameter spaces. In this work, we present a reduction framework for fracture simulations that leaves out the *offline* training phase and instead adaptively constructs reduced solutions spaces *online*. We apply the framework to the thick level set (TLS) method, a novel approach for modeling fracture able to model crack initiation, propagation, branching, and merging. The analysis starts with a fully-solved load step, after which two consecutive reduction operations—the proper orthogonal decomposition and the empirical cubature method—are performed. Numerical features specific to the TLS method are used to define an adaptive domain decomposition scheme that allows for three levels of model reduction coexisting on the same finite element mesh. Special solutions are proposed that allow the framework to deal with enriched nodes and a dynamic number of integration points. We demonstrate and assess the performance of the framework with a number of numerical examples.

KEYWORDS

adaptivity, fracture, hyper-reduction, model order reduction, thick level set

1 | INTRODUCTION

The search for numerical methods capable of accurately capturing the complex mechanisms involved in the fracture of solids is still one of the most active research fields in computational mechanics, despite its long history and vast body of literature. Although fracture mechanics can readily predict how a single existing crack grows,¹ modeling of phenomena such as crack initiation, branching, and merging is still a challenge. At the other end of the spectrum, damage models are well suited for predicting initiation and interaction between cracked regions but tend to suffer from spurious strain localization² and mesh bias. Two popular solutions for this mesh size dependency are to either use the crack band method³ to introduce a length scale in the damage formulation or to use a rate-dependent damage evolution in order to retard

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2019 The Authors. *International Journal for Numerical Methods in Engineering* published by John Wiley & Sons, Ltd.

strain localization.⁴ These, however, do not remove the bias introduced by the orientation of the mesh. A number of more sophisticated damage models such as the nonlocal integral damage model⁵ and gradient damage models^{6–8} offer alternative ways to couple damage evolution with an intrinsic material length scale. More recently, two new approaches to couple damage and fracture mechanics in a single regularized framework based on superposed solution fields have been introduced, namely the phase field method⁹ and the thick level set (TLS) method.¹⁰ The latter is the method of choice for modeling fracture in this work.

The TLS method, first proposed by Moës et al.,¹⁰ attempts to bridge the gap between damage and fracture mechanics. In this method, cracks are represented by a level set front that can naturally represent interaction phenomena such as branching and merging. Fracture mechanics is incorporated by making the rate of front growth dependent on the energy release rate of the material. Furthermore, damage is incorporated by introducing a thick band of material behind the front where the damage variable gradually increases until a zero-stiffness, completely cracked region emerges. The introduction of a length scale represented by the thickness of this damaging band effectively introduces nonlocality in the damage formulation, eliminating spurious strain localization. Furthermore, the use of a continuous level set field independent from the underlying finite element discretization alleviates the issue of crack orientation bias. Since its inception, the method has been expanded upon by multiple authors in order to deal with asymmetric constitutive behavior,¹¹ couple the damage formulation with plasticity¹² and cohesive zone laws,¹³ improve the representation of traction-free sliding in shear,¹⁴ and treat fatigue loading.¹⁵

Although being a versatile and robust method, TLS is a computationally demanding analysis approach. Accurate modeling of front growth requires element sizes significantly smaller than the width of the damage band.¹¹ Depending on the geometry of the problem being modeled, this width is constrained to be smaller than the length scale of geometric features in the domain—for example, a crack propagating between two inclusions very close to each other. Moreover, since the level set is usually updated in a staggered fashion, numerical stability of the scheme requires that the front does not advance for more than one element length per time step.^{11,14} These constraints render TLS computationally prohibitive for a number of relevant problems. It is worth mentioning that opting for other regularized damage approaches (eg, gradient-enhanced damage⁶ or phase field models⁹) does not alleviate the issue, as those methods also require dense finite element meshes and involve additional degrees of freedom (eg, nonlocal strains) in the complete domain while the TLS solves a damage update problem that involves additional unknowns only within damaged regions.

Model order reduction (MOR) techniques may be used to reduce the computational effort associated with the TLS method. These techniques essentially consist in finding suitable lower-dimensional solution spaces that can be employed as surrogates for full-order models. The goal is to construct surrogate models with significantly lower computational effort while minimizing the associated loss of solution accuracy. For the equilibrium problem that constitutes the most demanding phase of a TLS analysis,^{10,11} the number of degrees of freedom can be drastically reduced through the proper orthogonal decomposition (POD) technique¹⁶ while the cost of computing the global internal forces and stiffness matrix can be reduced through *hyper-reduction* techniques such as the empirical cubature method (ECM).¹⁷

However, using MOR for highly nonlinear fracture problems such as the ones treated with TLS is a complex endeavor for which no definitive solution is currently available. MOR methods rely on an *offline* training phase during which representative loading cases are executed and the resulting solution snapshots are used to compute the reduced solution spaces. For fracture problems, a difficulty arises when choosing loading cases for training: subtle changes in boundary conditions can lead to considerable differences in crack initiation and propagation behavior. In other words, the parameter space of possible solutions that the surrogate model should be able to approximate is prohibitively large, leading to inefficient training and a lower acceleration level through an increase in the size of the reduced-order solution space.

In recent years, a number of strategies have been proposed in order to alleviate this issue and allow for projection-based MOR techniques to be applied to fracture problems. In Reference 18, the authors employ a mesh morphing technique that effectively reduces the size of the reduced space necessary to accurately approximate the solution, while authors in Reference 19 combine POD with domain decomposition in order to allow for locally-refined bases to be used in zones of strain localization. Although promising, these approaches still rely on the construction of a reduced basis *offline*. For TLS problems, aside from the difficult task of selecting training cases, the adaptive nodal enrichment and numerical integration schemes necessary for the accuracy and robustness of the method¹¹ lead to stress and displacement snapshots of variable sizes and orderings, further complicating the task of a *a posteriori* construction of reduced bases. An alternative class of strategies which seems more suitable to the problem at hand consists in adaptively refining the reduced basis during the *online* analysis either through the use of Krylov subspaces^{16,20,21} or by locally solving critical mesh regions in the full-order space.²²

It seems clear that resorting to pre-trained models would be a cumbersome and inefficient approach to reducing the computational complexity of the TLS method. This paper is instead focused on adaptively reduced models that preclude the need for an *offline* training phase. In a recent work,²³ we propose an adaptive reduction framework that starts with a fully-solved load step and combines a number of state-of-the-art MOR techniques^{17,19,22,24} in order to progressively build a hyper-reduced model that allows for different levels of reduction coexisting on different subdomains of a single finite element mesh. Here, the framework is applied and tailored to the specific needs of a TLS analysis, although it can easily be used to accelerate other types of fracture models. We start by looking at the full-order model of a compact tension test^{14,25} and reviewing the basic components of the TLS method. MOR ingredients are then gradually added to the model while exploiting the full-field information that is a by-product of the TLS front update—the level set field, front velocities, and parametrized front length—in order to achieve a fully-adaptive domain decomposition strategy to accommodate the different levels (full, reduced, hyper-reduced) of MOR. The performance of the modified framework is assessed and additional numerical examples are shown.

1.1 | Mathematical notation

Throughout this work, vectors and matrices are written as boldfaced lower-case and upper-case symbols, respectively. Sets of indices are written in upper-case calligraphic script (eg, $\mathcal{S}$). Entities in the full-order solution space are indicated by the subscript h , while reduced-order entities are left unmarked for compactness. A subscript with a set between parenthesis (eg, $\mathbf{M}_{(\mathcal{S}, \mathcal{T})}$) indicates a selection operation based on one or two sets of indices that extracts submatrices or subvectors.

2 | THE THICK LEVEL SET METHOD

In this section, we review the main analysis steps that comprise the TLS method. Let Ω be a volume in equilibrium subjected to Dirichlet constraints at boundary Γ_u and Neumann boundary conditions at Γ_f such as $\Gamma_u \cap \Gamma_f = \emptyset$. The TLS method consists in using a level set function $\phi = \phi(x, y)$ in 2D or $\phi = \phi(x, y, z)$ in 3D to define a damage front Γ_ϕ that divides Ω into a damaged volume Ω_d and an intact volume Ω_i , as shown in Figure 1.

The level set function ϕ is a signed distance function to the front, satisfying the eikonal equation for unit velocity, written as:

$$|\nabla \phi| = 1, \quad (1)$$

subjected to $\phi = 0$ on Γ_ϕ . The presence of a front introduces in Ω the curvilinear coordinate system $\phi - s$ shown in Figure 1. At distances larger than the length scale parameter l_c in the direction normal to the front (along the ϕ -axis for every value of s), the material is considered to have lost its complete stiffness ($d = 1$). With these definitions, the damage variable can be written as:

$$d(\phi) = \begin{cases} 0, & \phi \leq 0 \\ f(\phi), & 0 < \phi \leq l_c \\ 1, & \phi > l_c \end{cases} \quad (2)$$




FIGURE 1 The Thick Level Set method: the domain Ω is partitioned into damaged (Ω_d) and intact (Ω_i) volumes by a level set ϕ [Colour figure can be viewed at wileyonlinelibrary.com]

where $f(\phi)$ is a damage profile function satisfying $\partial f / \partial \phi \geq 0$ for $0 \leq \phi \leq l_c$. In this paper, the arctangent profile proposed by Bernard et al¹¹ is used for all numerical examples:

$$f(\phi) = c_2 \operatorname{atan} \left(c_1 \left(\frac{\phi}{l_c} - c_3 \right) \right) + c_4, \quad (3)$$

with $c_1 = 10$, $c_3 = 0.5$, and the other coefficients given by:

$$c_2 = \frac{1}{\operatorname{atan}(c_1(1 - c_3)) - \operatorname{atan}(-c_1 c_3)}, \quad c_4 = -c_2 \operatorname{atan}(-c_1 c_3). \quad (4)$$

With the signed distance function for ϕ , the closest distance to the damage front is known at every point in the domain. This information will be useful for the adaptive scheme in Section 3.

The TLS implementation considered in this work is based on a staggered scheme: from an existing level set configuration, an equilibrium problem with unit load* is solved for the displacement field $\mathbf{u}$. The displacements are then used to solve a second system of equations on Ω_d that computes the configurational forces at the damage front Γ_ϕ . By comparing the configurational forces with a material resistance parameter, the evolution of the level set front is computed, along with the load factor λ that is used to scale the unit load equilibrium solution. Finally, the level set field is updated and new damage fronts are allowed to nucleate if the material strength is surpassed at any point on Ω_i . In the following sections, further details on each of these analysis phases (global equilibrium, front evolution, and level set update) will be provided.

2.1 | Equilibrium problem

For the present study, we opt for a simple isotropic damage formulation for which the free energy can be written as:

$$\psi = \frac{1}{2} (1 - d) \boldsymbol{\epsilon}^T \mathbf{D} \boldsymbol{\epsilon}, \quad (5)$$

where $\mathbf{D}$ is the elastic stiffness tensor. The stresses are obtained by differentiating the energy with respect to strain:

$$\boldsymbol{\sigma} = \frac{\partial \psi}{\partial \boldsymbol{\epsilon}} = (1 - d) \mathbf{D} \boldsymbol{\epsilon}. \quad (6)$$

It is important to note that the isotropic nature of this simple damage model leads to unrealistic predictions of damage nucleation and propagation under compression. Although asymmetric damage models have already been proposed for use with TLS,^{10,11} those are also not applicable to general loading directions, leading to spurious stiffness recovery under shear. In Reference 14, van der Meer and Sluys propose an interphase material that eliminates this issue and allows the TLS to capture traction-free sliding, but it can only be employed on real material interfaces. In order to keep the discussion focused on accelerating TLS using MOR techniques, the simple and numerically robust isotropic formulation is adopted.

With the stress tensor, the full-order internal force at integration point i can be written as:

$$\mathbf{f}_h(\mathbf{x}_i) = \mathbf{B}^T(\mathbf{x}_i) \boldsymbol{\sigma}(\boldsymbol{\epsilon}_i), \quad (7)$$

where $\mathbf{B}$ is a matrix that relates nodal displacements to local strains and $\mathbf{f}_h \in \mathbb{R}^N$ is the contribution of point i to the full-order (as indicated by the subscript h) global internal force vector, with N being the total number of degrees of freedom in the model.

In the TLS method, the region with $\phi > l_c$ represents a stress-free crack. Due to the discrete nature of the finite element method, this region must be at least one element wide in order to achieve a stress-free state, which would otherwise introduce mesh dependency to the problem.¹⁴ In order to circumvent this issue, the adaptive enrichment scheme proposed by Bernard et al¹¹ is employed and extra degrees of freedom are added to nodes inside the cracked region whose element support is cut at least twice by the iso- l_c . In practice, this leads to a dynamic number of degrees of freedom N , with

*This is only allowed because a linear material model is employed in the staggered solution scheme. For an alternative approach suitable for nonlinear material models, the reader is referred to Reference 12.

nodes being enriched or unenriched as the front moves. Furthermore, enriched elements will have a different $\mathbf{B}$ matrix than regular elements. These two consequences of the enrichment scheme will be important when building the reduced models in Section 3.

The global internal force vector $\mathbf{f}_h^\Omega$ is numerically integrated by combining the contribution of all M integration points in the mesh:

$$\mathbf{f}_h^\Omega = \int_{\Omega} \mathbf{f}_h d\Omega \approx \sum_{i=1}^M \mathbf{f}_h(\mathbf{x}_i) w_i, \quad (8)$$

where w_i is the integration weight of the point located at $\mathbf{x}_i$. In order to improve the accuracy of the TLS solution, extra integration points are added to the elements cut by the iso-zero as well as to the elements cut by the iso- l_c . This means that M is not constant during the analysis, a fact that must be accounted for when building a hyper-reduced model for $\mathbf{f}^2$ (Section 3). Finally, the displacement vector $\mathbf{u}_h \in \mathbb{R}^N$ that leads to global equilibrium can be obtained by solving:

$$\mathbf{r}_h = \mathbf{f}_h^\Gamma - \mathbf{f}_h^\Omega(\mathbf{u}_h) = \mathbf{0}, \quad (9)$$

where $\mathbf{r}_h \in \mathbb{R}^N$ is a residual vector and $\mathbf{f}_h^\Gamma \in \mathbb{R}^N$ is the global external force vector.

2.2 | Front evolution

As damage is directly coupled to the distance to the level set front, a movement of this front leads to changes in damage distribution and, consequently, to energy dissipation. This dissipation is related to the configurational force Y , obtained from Equation (5) as:

$$Y = -\frac{\partial \psi}{\partial d} = \frac{1}{2} \boldsymbol{\epsilon}^T \mathbf{D} \boldsymbol{\epsilon}. \quad (10)$$

In the TLS method, the classic local dissipation measure $\Xi = Y \dot{d}$ is substituted by a measure of dissipation over the complete damaged band when it moves by a distance $\delta \phi$:

$$\Xi = \int_{\Gamma_\phi} \int_0^l Y(\phi, s) \frac{\partial d}{\partial \phi} \left(1 - \frac{\phi}{\rho(s)} \right) d\phi \delta \phi(s) ds = \int_{\Gamma_\phi} g(s) \delta \phi(s) ds, \quad (11)$$

where $g(s)$ is the band configurational force at coordinate s :

$$g(s) = \int_0^l Y(\phi, s) \frac{\partial d}{\partial \phi} \left(1 - \frac{\phi}{\rho(s)} \right) d\phi, \quad (12)$$

with ρ being the curvature of the iso-zero at coordinate s and $0 \leq l \leq l_c$ is the width of the damaged band. Since $g(s) \rightarrow 0$ when $l \rightarrow 0$, an average configurational force $\bar{Y}$ which is only a function of s is computed in such a way as to enforce:¹¹

$$\int_0^l \bar{Y}(s) \frac{\partial d}{\partial \phi} \left(1 - \frac{\phi}{\rho(s)} \right) d\phi = \int_0^l Y(\phi, s) \frac{\partial d}{\partial \phi} \left(1 - \frac{\phi}{\rho(s)} \right) d\phi. \quad (13)$$

Numerically, this average configurational force along the front is computed by solving Equation (13) on the domain Ω_d with $\bar{Y}$ as nodal values while enforcing the constraint $\nabla \bar{Y} \cdot \nabla \phi = 0$ with Lagrange multipliers.¹¹ The front propagation criterion at a given point i at the front simply becomes $\bar{Y}_i = \bar{Y}_{ci}$, with $\bar{Y}_c$ being the homogenized value of a resistance parameter $Y_c(\phi, s)$ computed by solving a system of equations similar to Equation (13) but substituting $\bar{Y}$ by $\bar{Y}_c$.

Although a constant propagation resistance Y_c could be used throughout the analysis,^{10,11} we follow the approach proposed in Reference 14 instead and make Y_c a function of the size of the damaged zone in order to take into account the

difference in the stress levels necessary for damage initiation and propagation. With values for Y_c related to the strength (Y_c^f) and fracture energy ($\bar{Y}_c^G$), the following interpolation is adopted:

$$\log Y_c = \log Y_c^f + \frac{\bar{\phi}}{\bar{\phi}_{\max}} \left(\log Y_c^G - \log Y_c^f \right), \quad (14)$$

where $\bar{\phi}$ is the average value of ϕ inside Ω_d , $l_c/3 \leq \bar{\phi}_{\max} \leq l_c/2$ is the band size after which $\bar{Y}_c$ ceases to increase, and the two parameters related to initiation and propagation are given by:

$$Y_c^f = \frac{1}{2} \frac{f_t^2}{E}, \quad Y_c^G = \frac{G_c}{2Al_c} \quad \text{with} \quad A = \frac{1}{l_c} \int_0^{l_c} f(\phi) d\phi, \quad (15)$$

where $f(\phi)$ is given in Equation (3).

With values of $\bar{Y}$ and $\bar{Y}_c$ at every node i on $\mathcal{N}_\phi$, defined as the set of nodes on elements cut by the front Γ_ϕ , the load factor λ necessary to promote front growth can be computed by solving:

$$\lambda^2 \max_{i \in \mathcal{N}_\phi} \left\{ \frac{\bar{Y}_i}{\bar{Y}_{ci}} \right\} = 1. \quad (16)$$

Finally, the front velocity v in the direction perpendicular to the level set (in the negative direction along the ϕ -axis of Figure 1) for every node in $\mathcal{N}_\phi$ can be obtained through:¹⁴

$$v_i = k \left\langle \frac{c\lambda^2 \bar{Y}_i}{\bar{Y}_{ci}} - 1 \right\rangle_+ \quad k = \frac{v_{\max}}{c-1}, \quad (17)$$

where the operator $\langle \cdot \rangle_+$ returns the positive part of its operand and $v_{\max}$ is the maximum growth the front can undergo in a single load step. In order to guarantee the stability of the staggered analysis scheme, a value $v_{\max} = h_{\min}/2$ is adopted here, with $h_{\min}$ being the size of the smallest finite element present in the mesh. The parameter c controls the amount of velocity spread to nodes with lower values for the ratio $\bar{Y}/\bar{Y}_c$. For $c \rightarrow 1$, only the point used to compute λ in Equation (16) will have nonzero velocity.

2.3 | Level set update and damage nucleation

With the nodal velocities at the front, the last analysis phase consists in updating the level set field ϕ (and consequently the damage distribution). This in turn involves extending the front velocities to every node not belonging to $\mathcal{N}_\phi$. In order to guarantee that the updated level set will satisfy Equation (1), the velocities are extended along the ϕ -axis, satisfying the constraint:

$$\nabla \phi \cdot \nabla a = 0, \quad (18)$$


which is solved with a fast marching algorithm.²⁶ The level set field at every node is then updated as:

$$\phi_i = \phi_{io} + v_i, \quad (19)$$

with $i \in \mathcal{N}$ (the complete set of nodes in the mesh) and the subscript o indicating values from the previous time step. Even though this extension strategy is designed to guarantee the signed distance property of ϕ , the discrete nature of the level set update leads to increasingly larger violations of Equation (1) as the analysis progresses. A reinitialization procedure in which Equation (1) is solved with a fast marching algorithm is therefore periodically performed.²⁶

During this update phase, a nucleation check is performed in order to allow for new damage fronts to emerge. The size of a newly created damage nucleus is considered to be infinitesimal, therefore reducing the interpolated Y_c value of

FIGURE 2 Analysis flow of a single time step in the TLS method, showing the couplings between the three analysis phases. TLS, thick level set [Colour figure can be viewed at wileyonlinelibrary.com]



Equation (14) to simply Y_c^f and making the initiation criterion local. The check is performed at every integration point and a new nucleus is created at coordinates:

$$\mathbf{x}_{\text{nucl}} = \arg \max_{\bar{\mathbf{x}}} \left\{ \lambda^2 \frac{Y(\bar{\mathbf{x}})}{Y_c^f(\bar{\mathbf{x}})} \right\}, \quad (20)$$

where Y_c^f is made a function of $\mathbf{x}$ to take into account the general case of multiple materials being used in the same mesh. After nucleation, the value of ϕ is checked at every node and updated if the distance between the node and the new nucleus is smaller than ϕ .

With the updated level set field, the damaged band size $\bar{\phi}$ at Γ_ϕ is computed in order to be used in Equation (14) during the following time step. This is done in a similar way as in the computation of $\bar{Y}$ and $\bar{Y}_c$, by solving the following weak form:¹⁴

$$\int_0^l \bar{\phi}(s) \left(1 - \frac{\phi}{\rho(s)} \right) d\phi = \int_0^l \phi \left(1 - \frac{\phi}{\rho(s)} \right) d\phi, \quad (21)$$

while using Lagrange multipliers to enforce the constraint $\nabla \bar{\phi} \cdot \nabla \phi = 0$. The analysis flow of a single TLS step is shown in Figure 2, including the information exchanged between analysis phases.

2.4 | Example: Compact tension test

Before attempting to accelerate TLS simulations with MOR techniques, we show a full-order numerical example to illustrate the method. In Section 3, the same example will be used to gradually introduce each reduction technique.

The example involves the compact tension test used by Li et al²⁵ to investigate mode I crack growth in fiber-reinforced polymers and later modeled with TLS by van der Meer and Sluys.¹⁴ The specimen is a square (100mm × 100mm) plate modeled in plane stress with a notch with circular tip ($r = 1\text{mm}$) pulled vertically from the circular load application points seen in Figure 3. The material parameters are $E = 7\text{GPa}$, $\nu = 0.3$, $G_c = 40\text{ N/mm}^2$, $f_t = 79\text{MPa}$, $l_c = 2\text{ mm}$, $c = 2$, and $\bar{\phi}_{\text{max}} = 0.45l_c$. The plate is discretized with a total of 76 784 constant-strain triangles with one integration point each, resulting in $N = 77\,234$ degrees of freedom (DOFs) and $M = 76\,784$ integration points. A relatively dense discretization is employed along the expected crack path (Figure 4), with a minimum characteristic element size $h_{\text{min}} = 0.15\text{ mm}$.

No damage fronts are present during the first time step. The unit load problem is solved and a very high λ is set in order to guarantee that a first nucleation will occur during the level set update phase. Since the structure behaves elastically if no damage fronts exist, this effectively promotes a jump to the moment at which the first nucleation occurs, at a load of approximately 100 N. From the second step onwards, λ is computed according to Equation (16). The global behavior starts with a hardening branch during which a fracture process zone of radius l_c develops and transitions to a softening branch as the crack grows towards the opposite edge of the plate. The load-displacement curve can be seen in Figure 3, as well as the extended velocity field from a time step during the softening branch.




FIGURE 3 Compact tension test: load-displacement behavior and velocity field at the step indicated in red




FIGURE 4 Compact tension test: enriched nodes and elements with modified integration scheme [Colour figure can be viewed at wileyonlinelibrary.com]

As the damage front propagates, the size of the equilibrium problem (N) changes as nodes from elements cut by the $\text{iso-}l_c$ are enriched. Furthermore, the number of integration points of elements cut by both the iso-zero and the $\text{iso-}l_c$ changes from 1 to 9, increasing the total number of integration points M . Figure 4 shows the damage front after the last analysis step, as well as a zoomed-in view of the damage band indicating which nodes and elements are modified.

The analysis runs in 1670 s on a workstation equipped with a Xeon W-2123 processor running Ubuntu 18.04.2. From this total execution time, 1494 s (92%) is dedicated to solving the equilibrium problem, with the remaining 5% and 3% being spent on computing the front advance and updating the level set field, respectively. Since Equations (13) and (21) are only solved inside Ω_d , the equilibrium problem is the main complexity bottleneck of the TLS method.^{11,14}

3 | ADAPTIVE MODEL ORDER REDUCTION

In this section, the mechanical equilibrium problem associated with the TLS method (the first analysis phase in Figure 2) is accelerated with a number of the MOR components presented in Reference 23. Two main reduction avenues can be identified: reducing the number of DOFs N and the number of integration points M necessary for integrating the internal force vector $\mathbf{f}^\Omega$ of Equation (8). As the parameter space—defined here as the space of possible load histories to which the model can be subjected—associated with fracture problems is prohibitively large, the reduced spaces for N and M are constructed *online* in an adaptive fashion.

3.1 | Reduction by projection

Instead of solving for $\mathbf{u}_h \in \mathbb{R}^N$, the solution to the equilibrium problem of Equation (9) can instead be sought for on a lower-dimensional manifold defined by a set of n orthonormal basis vectors $\boldsymbol{\psi}_i$. After solving for the $n \ll N$ mode

contributions α , the full-order displacement field can be recovered by:

$$\mathbf{u}_h = \Psi \alpha \quad \Psi = [\psi_1 \ \psi_2 \ \cdots \ \psi_n]. \quad (22)$$

A popular approach for constructing the reduced solution space spanned by Ψ is the POD method: a set of P full-order displacement snapshots $\mathbf{X}_u$ is decomposed into orthonormal contributions through a singular value decomposition (SVD) operation:

$$\mathbf{X}_u = \Psi \mathbf{S} \mathbf{V}^T, \quad (23)$$

where the basis Ψ is the left-singular matrix of $\mathbf{X}_u$, $\mathbf{V}$ is a matrix with right-singular vectors, and $\mathbf{S}$ is a diagonal matrix with singular values. When the number of snapshots P is large, the SVD is usually truncated to the first $n \ll P$ modes, associated with the highest singular values. The reduced version of the equilibrium problem is obtained by imposing the Galerkin projection constraint $\Psi^T \mathbf{r}_h = \mathbf{0}$, which yields:

$$\mathbf{f}^\Omega = \Psi^T \mathbf{f}_h^\Omega, \quad \mathbf{f}^\Gamma = \Psi^T \mathbf{f}_h^\Gamma, \quad \mathbf{K} = \Psi^T \mathbf{K}_h \Psi. \quad (24)$$

For pre-trained models, obtaining the snapshot matrix $\mathbf{X}_u$ is the goal of the *offline* training phase. Here, $\mathbf{X}_u$ is updated whenever a load step is solved in the full-order space. Since the analysis starts without an initial basis Ψ , the very first load step is solved fully and subsequent steps are reduced. As the analysis progresses and changes in damage distribution occur, the reduced space Ψ gradually loses the ability to describe the global behavior of the structure being modeled. At this point, even though the reduced problem is in equilibrium ($\mathbf{f}^\Omega = \mathbf{f}^\Gamma$), the full-order equilibrium of Equation (9) is not exactly satisfied. This deviation from equilibrium can be used to trigger a fully-solved retraining step if a certain tolerance threshold ϵ_{force} is crossed:²³

$$\left\| \mathbf{f}_h^\Gamma - \mathbf{f}_h^\Omega \right\| > \epsilon_{\text{force}}, \quad (25)$$

where the load scaling used in Reference 23 is dropped because the mechanical equilibrium problem is based on a unit-load solution. After the retraining step is solved, the field $\mathbf{u}_h$ and its change from the previous time step $\Delta \mathbf{u}_h$ are added to $\mathbf{X}_u$ and a new basis Ψ is computed with Equation (23). In order to limit the computational overhead associated with these SVD operations, $\mathbf{X}_u$ is only allowed to have a number n_f of snapshots, with older solutions being gradually discarded as new snapshots are added. Furthermore, the SVD operation is truncated once singular values drop below a threshold ϵ_{SV} in order to guarantee that only the most relevant basis vectors are included in the basis.

3.1.1 | Compact tension example

We now return to the compact tension example of Section 2.4 and attempt to use the present adaptive POD approach to solve it in the reduced space. With the current approach, we immediately run into a problem when using the nodal enrichment scheme shown in Figure 4: any new enrichments occurring between two consecutive retraining steps cannot be included in the basis, since Ψ can only be updated after a fully-solved step. Delaying these new enrichments would in turn lead to an enlargement of the damage front until the region with $\phi > l_c$ is at least one element wide. Since the level set advancement is irreversible, these thicker cracked regions would persist even after a full step is triggered. The enrichment scheme is therefore disabled until a solution for this issue is proposed in Section 3.2.

For the additional parameters related to the POD reduction, we adopt $n_f = 6$ and $\epsilon_{\text{SV}} = 10^{-6}$. This means that the number of reduced DOFs n is at most 6, a significant reduction when compared to the original $N = 77\,234$. The tolerance ϵ_{force} is adjusted such that the crack is able to correctly propagate from the notch. For higher values of ϵ_{force} , the deficient reduced basis blocks the front advance and causes the cracked band to widen. This is an unphysical effect that does not occur if a low enough tolerance is adopted. Figure 5 shows the comparison between damage front shapes obtained with two different tolerances.

With $\epsilon_{\text{force}} = 0.025$, the load-displacement curve shown in Figure 6 is obtained. Although both the shape of the damage front and the quasi-static equilibrium path are correctly reproduced by the reduced solution, it produces a curve with




FIGURE 5 Front behavior comparison between the full solution and POD solutions with different error tolerances. POD, proper orthogonal decomposition [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 6 Load-displacement behavior obtained with the adaptive POD approach. The detail highlights the oscillatory reduced behavior caused by a high retraining frequency. POD, proper orthogonal decomposition





FIGURE 7 Accumulated number of reduced steps as the analysis progresses. A large deviation from linearity implies a large number of retraining steps

higher numerical noise. These equilibrium path jumps occur whenever a retraining step occurs and the reduced curve snaps back to the full one (cf²³).

The noisy load-displacement response hints at a very high number of retraining steps driven by the highly non-linear structural behavior caused by the propagating crack. Indeed, plotting the cumulative number of reduced steps versus the total number of load steps shows that approximately half of the steps are fully solved (Figure 7). The reduction scheme is particularly inefficient towards the end of the analysis, when the error control condition of Equation (25) is triggered as soon as Ψ is updated and no more reduced steps are allowed. Nevertheless, the reduced mechanical equilibrium model runs approximately 1.4 times faster than the full one. However, it is important to note that even though the quantity of interest in this case (the load shown in Figure 6) is reproduced with reasonable accuracy by the reduced model, the net benefit of reducing the problem might become negative depending on the application: it is difficult to justify even very minor losses in accuracy when the gains in execution time are so limited.

Finally, it is interesting to investigate the spatial distribution of the deviation from equilibrium $\mathbf{r}_h$ that triggers retraining steps. In Figure 8, this error is plotted for a reduced load step immediately prior to a retraining step. As expected, the reduced space obtained with information from the latest retraining step struggles to represent what happens immediately

FIGURE 8 Deviation from full-order equilibrium from a reduced step immediately prior to retraining. The reduced basis Ψ does not adequately represent the behavior of the region immediately ahead of the crack tip



ahead of the crack tip, the area of Ω_d where most of the changes in damage profile occur—that is, the region with the highest slope $\partial d / \partial \phi$ of the arctangent damage expression of Equation (3).

3.2 | Equilibrium system partitioning

As illustrated in Section 3.1, two issues arise when attempting to accelerate the TLS equilibrium problem with an adaptive POD approach. First, the enrichment scheme proposed in Reference 11 is incompatible with the fact that the basis Ψ is only updated during fully-solved steps. Second, the inability of the reduced basis to correctly capture changes in mechanical behavior immediately ahead of the crack tip leads to a high number of retraining steps, negating most of the acceleration associated with solving an equilibrium problem of size $n \ll N$. The partitioning strategy first proposed by Kerfriden et al²² and later adopted as part of the framework in Reference 23 offers a potential solution to both of these issues, as it allows for a group of DOFs to be detached from the reduced solution space and directly solved for in the full-order space.

We divide the complete set of DOFs $\mathcal{U}$ into a set $\mathcal{R}$ of reduced DOFs and a set $\mathcal{F}$ of fully-solved DOFs ($\mathcal{R} \cap \mathcal{F} = \emptyset$). The DOFs in $\mathcal{F}$ are solved for in the full-order space while the reduced basis Ψ is used to approximate the solution in $\mathcal{R}$. With this partitioning, the reconstruction of $\mathbf{u}_h$ becomes:

$$\mathbf{u}_{h(\mathcal{R})} = \Psi_{(\mathcal{R})} \alpha, \quad \mathbf{u}_{h(\mathcal{F})} = \mathbf{u}_f, \quad (26)$$

where $\mathbf{u}_f$ is a vector with displacements associated to DOFs in $\mathcal{F}$ and $\Psi_{(\mathcal{R})}$ is a submatrix obtained by selecting the rows of Ψ associated with the DOF indices in $\mathcal{R}$. After partitioning, the reduced solution vector becomes:

$$\mathbf{u} = \begin{pmatrix} \alpha \\ \mathbf{u}_f \end{pmatrix}, \quad (27)$$

By detaching $\mathcal{F}$ from the Galerkin projection constraint $\Psi^T \mathbf{r}_h = \mathbf{0}$, we obtain the following partially-reduced versions of the force vectors and stiffness matrix:^{22,23}

$$\mathbf{f}^\Omega = \begin{pmatrix} \Psi_{(\mathcal{R})}^T \mathbf{f}_{h(\mathcal{R})}^\Omega \\ \mathbf{f}_{h(\mathcal{F})}^\Omega \end{pmatrix}, \quad \mathbf{f}^\Gamma = \begin{pmatrix} \Psi_{(\mathcal{R})}^T \mathbf{f}_{h(\mathcal{R})}^\Gamma \\ \mathbf{f}_{h(\mathcal{F})}^\Gamma \end{pmatrix}, \quad (28)$$

$$\mathbf{K} = \begin{bmatrix} \Psi_{(\mathcal{R})}^T \mathbf{K}_{h(\mathcal{R}, \mathcal{R})} \Psi_{(\mathcal{R})} & \Psi_{(\mathcal{R})}^T \mathbf{K}_{h(\mathcal{R}, \mathcal{F})} \\ \mathbf{K}_{h(\mathcal{F}, \mathcal{R})} \Psi_{(\mathcal{R})} & \mathbf{K}_{h(\mathcal{F}, \mathcal{F})} \end{bmatrix}. \quad (29)$$

The question now becomes how to define $\mathcal{F}$ in order to minimize the number of retraining steps. In Reference 23, we have proposed a number of strategies to populate $\mathcal{F}$ when the framework is applied to problems with plasticity, including two criteria that search for regions with energy dissipation concentrations caused by plastic strain development. Here, we devise a similar strategy while taking advantage of full-field information already computed by the TLS method.

The region ahead of the crack tip shown in Figure 8, which concentrates most of the deviation from equilibrium that triggers retraining steps, is an intuitive choice for $\mathcal{F}$. As this deviation is caused by the front movement




FIGURE 9 Definition of the fully-solved DOF set $\mathcal{F}$ from the intersection of the level set curve ϕ_f with a viscous cumulative velocity field. DOF, degree of freedom [Colour figure can be viewed at wileyonlinelibrary.com]

as the crack propagates, the extended velocity field computed in order to update the level set field becomes a useful tool to identify regions undergoing changes in damage distribution. A straightforward approach would be to include the DOFs of every node with nonzero velocity in $\mathcal{F}$. However, two additional points must be taken into account:

- The velocity field is extended to the whole domain, which means that points far away from the front also have positive velocities (see Figure 3). This would lead to an overly large $\mathcal{F}$ set;
- Only points directly ahead of the crack tip have nonzero velocities. As the tip moves forward between two consecutive retraining steps, the nodes behind it will immediately drop out of $\mathcal{F}$, but the basis Ψ would still be unable to correctly capture their behavior.

We address the first point by also requiring points in $\mathcal{F}$ to fulfill the constraint $\phi \geq \phi_f$ —that is, $\mathcal{F}$ is the intersection between the region delimited by a level set ϕ_f and the nonzero velocity field. As for the second point, we substitute the instantaneous extended velocity field v by a decaying cumulative velocity measure v^* in order to retard the removal of points behind the crack tip from $\mathcal{F}$. At the beginning of every load step, the velocity field v_o computed during the previous step is used to update v^* as follows:

$$v^* = \exp\left(-\frac{1}{\tau}\right) v_o^* + \tau \left[1 - \exp\left(-\frac{1}{\tau}\right)\right] v_o, \quad (30)$$

where τ is a decay parameter. A node $i \in \mathcal{N}$ is added to $\mathcal{F}$ if the following conditions are satisfied:

$$\phi_i \geq \phi_f \quad \text{and} \quad v_i^* \geq \beta v_{\max}^*, \quad (31)$$

where β is a cutoff velocity parameter that controls the shape of the domain formed by nodes in $\mathcal{F}$ and $v_{\max}^*$ is the value of v^* at the node with the highest accumulated velocity. Figure 9 illustrates how each of the parameters τ , β , and ϕ_f influence the shape of $\mathcal{F}$. At this point, it is important to mention that even though the deviation from equilibrium of Figure 8 is limited to the $\phi > 0$ region, extending ϕ_f beyond the front may prove beneficial in reducing the number of retraining steps. The region ahead of the front is where a gradual increase in configurational force takes place up to the point where Y_c is reached. The movement of the front may therefore be hindered if the reduced basis does not allow for this stress concentration to emerge.

The partitioning strategy also allows for the reintroduction of the enrichment around $\phi = l_c$: any DOFs created by the enrichment scheme (Figure 4) that are not already present in Ψ are added to $\mathcal{F}$. This allows for each snapshot $\mathbf{x}_u$ on $\mathbf{X}_u$ to have a different size, depending on the value of N at the time when it is stored. The basis computed from this set of snapshots has, therefore, a number of rows $N_c = \min_i(\text{size}(\mathbf{x}_{ui}))$. As older snapshots are overwritten, N_c increases and the enriched DOFs now included in the basis are removed from $\mathcal{F}$. With this strategy, the enrichment scheme is allowed to operate between two consecutive retraining instances. It is worth noting in passing that a similar scheme could also be used to combine a POD-reduced model with a remeshing strategy, other methods based on nodal enrichment such as XFEM²⁷ or overlapping domain techniques such as CutFEM.²⁸

3.2.1 | Compact tension example

We now apply this partitioned POD strategy to the compact tension example of the previous sections. Figure 10 shows snapshots at two load steps of three different model executions with $\beta = 0.001$ and different combinations of τ and ϕ_f , resulting in different topologies for the region formed by the nodes included in $\mathcal{F}$. For the first model, the decay parameter is high enough to guarantee that the size of $\mathcal{F}$ increases monotonically and covers the whole damaged band. In contrast, the decay is disabled for the second model by setting τ to a very low value, which means that only points with a positive instantaneous velocity are included in $\mathcal{F}$. Finally, the third model seeks a balance between including nodes behind the crack tip and those ahead of the damage front while keeping $\mathcal{F}$ relatively small throughout the analysis. It is interesting to note that the third model, although having a smaller number of nodes in $\mathcal{F}$, needs significantly less retraining steps and is therefore more efficient in terms of acceleration than the other two. In any case, the fact that $\mathcal{F}$ is updated at the beginning of every load step leads to an efficient and fully-adaptive partitioning scheme that naturally follows the crack as it moves through the domain.




FIGURE 10 Compact tension test with the partitioned adaptive POD scheme. Different topologies of $\mathcal{F}$ lead to different retraining frequencies. The models are executed for 550 load steps [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 11 Total number of retraining steps needed to maintain accuracy with multiple combinations of τ and ϕ_f ($\beta = 0.001$, 550 load steps) [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 12 Maximum size of $\mathcal{F}$ throughout the analysis for multiple combinations of τ and ϕ_f ($\beta = 0.001$, $N = 77\,234$) [Colour figure can be viewed at wileyonlinelibrary.com]

The influence of the shape and size of $\mathcal{F}$ on the efficiency of the reduced solution is further investigated by running the model with different combinations of $0.01 \leq \tau \leq 100$ and $-10 \leq \phi_f \leq 3$, with $\beta = 0.001$. The resultant number of retraining steps for each combination is plotted in Figure 11, while the maximum size attained by $\mathcal{F}$ during the analysis is shown in Figure 12. For $\phi_f = 3$ mm, $\mathcal{F}$ is only used to accommodate DOFs from newly enriched nodes (since $l_c = 2$ mm) and the same number of retraining steps (317) is obtained for all values of τ . As ϕ_f decreases, an increasing number of nodes ahead of the crack tip are added to $\mathcal{F}$.

At first there is a sharp decrease in the number of retraining steps, as the critical region shown in Figure 8 is included in $\mathcal{F}$, with still significant decreases for $\phi_f < 0$, albeit at a gradually decreasing rate. This confirms that fully solving a finite region immediately ahead of the damage front is beneficial in reducing retraining frequency, depending on the problem being modeled and the choice of l_c (see Section 4.1 for an extended discussion on this point). For even lower values of ϕ_f , the number of retraining steps tends to stabilize. This is consistent with the fact that the stress concentrations that drive crack propagation gradually vanish as we move away from the damage front towards regions where deformation can be correctly reproduced in the reduced solution space.

The decay parameter τ also has significant influence on the number of retraining steps. In general, increasing τ leads to a decrease in retraining frequency, although no improvement is observed for $0 \leq \tau \leq 0.1$ due to the exponential nature of Equation (30). However, increasing τ to such extent that $\mathcal{F}$ covers the whole damaged band (eg, the first model in Figure 10) seems to be unnecessary. Because the basis Ψ is constantly being refreshed, regions of Ω_d far away from the crack tip can be well approximated by the POD solution.

It is important, therefore, to seek a balance between populating $\mathcal{F}$ with enough nodes both ahead and behind of the crack tip while minimizing its size. As can be seen in Figure 12, unfavorable choices for τ and ϕ_f can lead to highly inefficient reduced models with sizes of $\mathcal{F}$ that quickly approach the total number of DOFs N with only negligible reductions to the number of retraining steps.

Opting for the combination of parameters used on the third model of Figure 10 ($\phi_f = -2$ mm, $\tau = 10$, $\beta = 0.001$), we obtain the load-displacement curve of Figure 13. In contrast to the curve obtained with $\mathcal{F} = \emptyset$ (cf Figure 6), the reduced

FIGURE 13 Load-displacement behavior obtained with the partitioned POD approach. POD, proper orthogonal decomposition [Colour figure can be viewed at wileyonlinelibrary.com]





FIGURE 14 Accumulated number of reduced steps obtained with the partitioned POD approach with and without populating $\mathcal{F}$ with DOFs around the crack tip. DOF, degree of freedom



response here is virtually indistinguishable from the fully-solved one. Furthermore, the efficient partitioning successfully limits the number of retraining steps (Figure 14) while maintaining a relatively low number of DOFs in $\mathcal{F}$ (7.7%).

With this combination of parameters, the reduced equilibrium problem runs in 406.8s, accounting for 76% of the total execution time if all TLS phases are included. Of this total time, only 0.54s is dedicated to the SVD operations used to update Ψ . Although the resulting speed-up of 3.7 times with respect to the full-order equilibrium problem is already more than twice as high as the one obtained with $\mathcal{F} = \emptyset$, in the following section we attempt to further increase the level of acceleration by reducing the time spent on assembling $\mathbf{K}_h$ and $\mathbf{f}_h^\Omega$.

3.3 | Domain-based hyper-reduction

We now seek a way to compute $\mathbf{f}^\Omega$ —and consequently $\mathbf{K}$ as consistent linearization of $\mathbf{f}^\Omega(\mathbf{u})$ —in a faster way and with minimum loss of accuracy. Because the basis Ψ reduces the size of the equilibrium problem to $n \ll N$, the minimum amount of information needed to assemble $\mathbf{f}^\Omega$ is also greatly reduced. However, in the POD model $\mathbf{f}^\Omega$ is obtained by compressing the data gathered in a loop over all M integration points. In the final example of Section 3.2, 42% of the total time spent on the equilibrium problem was dedicated to this assembly of $\mathbf{f}_h^\Omega$ and $\mathbf{K}_h$. Strategies for accelerating this assembly operation are referred to as hyper-reduction methods.

The problem becomes to choose from the complete set of M integration points $\mathcal{I}$ a $m \ll M$ subset $\mathcal{Z} \subset \mathcal{I}$ with associated modified integration weights ϖ in such a way as to minimize the error between the M -integrated $\mathbf{f}^\Omega$ and its m -integrated counterpart:

$$(\mathcal{Z}, \varpi) = \arg \min_{\substack{\mathcal{Z} \subset \mathcal{I}, \varpi \geq 0}} \left\| \Psi^T \left(\sum_{i \in \mathcal{I}} \mathbf{f}_h(\mathbf{x}_i) w_i \right) - \Psi^T \left(\sum_{j \in \mathcal{Z}} \mathbf{f}_h(\mathbf{x}_j) \overline{\varpi}_j \right) \right\|^2. \quad (32)$$

In order to solve Equation (32), we use the so-called ECM originally proposed by Hernández et al¹⁷ and later adopted in the framework of Reference 23. At the end of every fully-solved step, the reduced internal force contribution $\mathbf{f}(\mathbf{x}_i) \in \mathbb{R}^n$ of each integration point is used to construct a single force snapshot $\mathbf{X}_f \in \mathbb{R}^{M \times n}$ given by:

$$\mathbf{X}_f = \begin{bmatrix} \sqrt{w_1} \left(\mathbf{f}(\mathbf{x}_1) - \frac{1}{\Omega} \mathbf{f}^\Omega \right) \\ \sqrt{w_2} \left(\mathbf{f}(\mathbf{x}_2) - \frac{1}{\Omega} \mathbf{f}^\Omega \right) \\ \vdots \\ \sqrt{w_M} \left(\mathbf{f}(\mathbf{x}_M) - \frac{1}{\Omega} \mathbf{f}^\Omega \right) \end{bmatrix}, \quad (33)$$

where reductions are performed with the updated Ψ basis. Since only a single snapshot is used to perform the reduction, a slightly different approach than the one used in Reference 23 is adopted: instead of storing stress snapshots and using SVD to obtain a basis matrix for stresses, here the SVD is used to decompose $\mathbf{X}_f$:

$$\mathbf{X}_f = \Lambda \mathbf{S} \mathbf{V}^T, \quad (34)$$

where $\Lambda \in \mathbb{R}^{M \times p}$ is a basis matrix truncated to the first p singular vectors. With this basis for internal forces, the minimization problem of Equation (32) is rewritten as the nonnegative least-squares minimization given by:

$$(\boldsymbol{\beta}, \mathcal{Z}) = \arg \min_{\substack{\bar{\mathcal{Z}} \subset \mathcal{I}, \\ \bar{\boldsymbol{\beta}} \geq \mathbf{0}}} \left\| \mathbf{J}(\bar{\mathcal{Z}}) \bar{\boldsymbol{\beta}} - \mathbf{b} \right\|^2, \quad (35)$$

where $\mathbf{J}$ and $\mathbf{b}$ are:

$$\mathbf{J} = [\Lambda \quad \sqrt{\mathbf{w}}]^T, \quad \mathbf{b} = [\mathbf{0} \quad \Omega]^T, \quad (36)$$

and the modified integration weights are obtained through element-wise multiplication of $\boldsymbol{\beta}$ by the original integration weights $\mathbf{w}_{(\mathcal{Z})}$. For more details on the formulation and implementation of the method, the interested reader is referred to References 17, 23.

The preceding development implies that the integration of the whole volume Ω is going to be reduced (see Equations (33) and (36)). For the current framework, this would not be desirable because:


- With ECM, the full-order force vector $\mathbf{f}_h^\Omega$ is only defined at a small set of nodes and the modified weights ϖ deprive it of its physical meaning. Reducing the whole domain Ω would therefore make computing the deviation from equilibrium used to trigger retraining steps (Equation (25)) impossible;
- Solving for DOFs belonging to the full-order set $\mathcal{F}$ requires the internal forces at those nodes to be fully integrated. Reduction of the complete domain would therefore make it impossible to use the equilibrium partitioning scheme of Section 3.2.

It is clear that at least part of Ω should still be fully integrated even after hyper-reduction. We therefore divide Ω into an arbitrary number[†] of hyper-reduced domains $\mathcal{H}$ and one fully-integrated domain $\mathcal{P}$, as illustrated in Figure 15. With this domain decomposition strategy, the reduced internal force vector is computed as:

$$\mathbf{f}^\Omega = \Psi^T \left(\sum_{i \in \mathcal{P}} \mathbf{f}_h(\mathbf{x}_i) w_i + \sum_{j \in \mathcal{Z}_{\mathcal{H}_1}} \mathbf{f}_h(\mathbf{x}_j) \varpi_j + \dots + \sum_{k \in \mathcal{Z}_{\mathcal{H}_n}} \mathbf{f}_h(\mathbf{x}_k) \varpi_k \right), \quad (37)$$

[†]The least-squares problem of Equation (35) cannot produce a set $\mathcal{Z}$ with more than $(p+1)$ ECM points, at which point $\mathbf{J}_{(\mathcal{Z})}$ becomes square (cf¹⁷). This limited set of points might not be enough to accurately represent the whole volume Ω . Defining multiple $\mathcal{H}$ domains allows for this integration error to be controlled.

FIGURE 15 Domain-based hyper-reduction scheme for TLS. A POD domain $\mathcal{P}$ follows the crack tip and the remaining volume is subdivided into ECM domains $\mathcal{H}$. A domain $\mathcal{F}$ with full DOFs is allowed to exist within $\mathcal{P}$. Domain topology is updated after every retraining step. DOF, degree of freedom; ECM, empirical cubature method; POD, proper orthogonal decomposition; TLS, thick level set [Colour figure can be viewed at wileyonlinelibrary.com]



and the minimization problem of Equation (35) is solved separately for each domain $\mathcal{H}_i$. Upon retraining, the domain configuration is updated in order to take into account changes in crack topology and new sets of points and weights are computed for each $\mathcal{H}_i$.

From the preceding discussion, the minimum requirement for $\mathcal{P}$ is that it should at least contain the elements in $\mathcal{F}$, including those with newly-enriched nodes. For simplicity, we adopt here the same strategy used to define $\mathcal{F}$ but using a value $\phi_p < \phi_f$ for the minimum level set value that defines the contour of the domain. This results in a $\mathcal{P}$ domain that is always larger than $\mathcal{F}$ in order to account for the movement of the front between two consecutive retraining steps. The motivation behind this choice stems from the fact that $\mathcal{F}$ cannot cross over the boundary of $\mathcal{P}$ into the region where internal forces are not defined. The difference between ϕ_p and ϕ_f should therefore be large enough as to allow the front to grow until the next retraining step is triggered without allowing $\mathcal{F}$ to touch the $\mathcal{P} - \mathcal{H}$ border.

The cracked region (where $\phi > l_c$) is also included in $\mathcal{P}$ in order to correctly represent the stress discontinuity introduced by the enrichment scheme. Since this is a small region composed of a single line of elements, the acceleration loss caused by extending $\mathcal{P}$ to cover its whole length is negligible. Since now $\mathbf{f}_h^\Omega$ is only correctly defined in $\mathcal{P}$, the retraining check of Equation (25) becomes:

$$\left\| \mathbf{f}_p^\Gamma - \mathbf{f}_p^\Omega \right\| > \epsilon_{\text{force}}, \quad (38)$$

where the subscript p refers to quantities from nodes in $\mathcal{P}$.

The remaining volume $\Omega \cap \mathcal{F}$ is subdivided into ECM domains $\mathcal{H}$. In Reference 23 we have proposed to use a k -means clustering algorithm to divide the domain into clusters with similar strain and use these clusters to define $\mathcal{H}$. Here we exploit the introduction of the $\phi - s$ coordinate system of Figure 1 to propose a simpler and more efficient approach: $\mathcal{H}$ can either be defined as TLS bands (regions between two specific values of ϕ) or as bands defined by a range of s values after using a fast marching algorithm to extend the nodal s values at the front to the rest of the domain. Alternatively, domains may be defined by combining both strategies, forming regions delimited by iso- ϕ and iso- s curves.

It is important to recall that as the front Γ_ϕ moves and new cracked regions are created, the integration scheme of elements cut by iso-zero and iso- l_c changes (Figure 4). It follows that a scenario might arise in which the integration of an element containing an ECM point changes between two consecutive retraining steps. Although this situation is mostly avoided by including cracked regions and regions with positive velocity in $\mathcal{P}$, it might still happen under certain circumstances (eg, if $\phi_p > 0$). As an additional safeguard, we store the coordinates of the points in $\mathcal{Z}$ at the time of training and use them when computing strains at ECM points instead of relying on the updated[‡] integration point locations.

3.3.1 | Compact tension example

We now revisit the compact tension example one last time. The adaptive partitioned POD strategy of Section 3.2 is complemented with the present domain-based hyper-reduction approach. The parameters used are $n_f = 6$, $\epsilon_{SV} = 1 \times 10^{-6}$, $\phi_f = -2$ mm, $\epsilon_{\text{force}} = 0.025$, $\phi_p = -4$ mm, $\tau = 10$, and $\beta = 0.001$. A precision $\epsilon_{\text{greedy}} = 1 \times 10^{-10}$ is used to truncate the SVD

[‡]We still allow the integration scheme in $\mathcal{H}$ to be updated because the two remaining analysis phases (Figure 2) are still fully integrated and rely on a higher integration precision along the damage front.




FIGURE 16 Compact tension test with the domain-based ECM scheme. Resulting domain topology for partitioning strategies based on ϕ (top), s (middle), and on both ϕ and s (bottom). ECM, empirical cubature method

of Equation (34) and as stopping criterion for the Greedy algorithm used to solve Equation (35). Figure 16 shows snapshots of three model executions with different partitioning strategies: using only the level set field ϕ ($n_\phi = 4$), using only the parametric coordinate s ($n_s = 4$), and combining both ϕ and s ($n_\phi = n_s = 4$). The adaptive nature of the present domain decomposition approach is evident: as new retraining steps are triggered, the domain topology changes and automatically follows the damage front as it propagates to the right edge of the specimen.

The benefit of increasing the number of ECM domains can be seen in the load-displacement curves of Figure 17. For $n_\phi = n_s = 1$, the global reduced behavior tends to deviate from the full-order curve in a similar way as observed in Figure 6. Furthermore, a higher number of retraining steps are triggered, with 94 steps for $n_\phi = n_s = 1$ and only 48 steps for $n_\phi = n_s = 10$. Since the number of POD modes is at most $n = 6$, the maximum number of ECM points on a single domain is $p + 1 = 7$. These 7 points are not enough to accurately integrate the whole domain $\mathcal{H} = \Omega \cap \mathcal{P}$, resulting in a poor approximation for $\mathbf{f}^\Omega$ which in turn leads to a higher retraining frequency. For $n_\phi = n_s = 10$, the presence of 100 domains allows for a higher number of integration points to be used (up to 700 points in this case).

Regarding the numerical stability of the hyper-reduced model, ECM guarantees that the reduced-order stiffness matrix retains its positive definiteness and thus its numerical stability.¹⁷ Furthermore, solving Equation (35) leads to the best possible approximation (in a least-squares sense) of the reduced force vector and is found to be robust for $\mathcal{H}$ domains of any size: in the lower-bound situation in which a domain $\mathcal{H}$ comprises only a single element, the error of the greedy procedure drops to a value close to zero after the first point is added and the procedure is stopped.

With $n_\phi = n_s = 10$, the obtained response agrees very well with the full-order one while still achieving a compression ratio of $\frac{M}{m} \approx 110$. Even with 100 ECM domains, assembly remains a very cheap operation, meaning that the increased

FIGURE 17 Load-displacement curves obtained with one ECM domain ($n_s = n_\phi = 1$) and 100 ECM domains ($n_s = n_\phi = 10$). The detail highlights the oscillatory behavior observed for the curve with one ECM domain. ECM, empirical cubature method





FIGURE 18 Accumulated number of reduced steps obtained for $n_\phi = n_s = 1$ domain and $n_\phi = n_s = 10$ domains, with the fully-integrated POD response for comparison. POD, proper orthogonal decomposition [Colour figure can be viewed at wileyonlinelibrary.com]



assembly time when compared to the model with one domain is easily outweighed by the consequent reduction in the number of retraining steps.

Another way to look at differences in efficiency between the models with one and 100 ECM domains is by plotting the cumulative number of reduced steps versus the total number of load steps (Figure 18). The model with 100 domains behaves similar to the fully-integrated adaptive POD model of Figure 14, while the higher number of retraining steps associated with the model with a single ECM domain makes the curve shift further away from the linear upper bound associated with a nonadaptive pre-trained model.

The model is executed for different combinations of $1 \leq n_\phi \leq 10$ and $1 \leq n_s \leq 10$ and the resulting number of retraining steps is shown in Figure 19. Increasing either n_ϕ or n_s in isolation leads to lower retraining frequencies, although increasing the number of domains in the s direction seems to be slightly more beneficial in this case: for $n_s = 5$ and $n_\phi = 1$, the number of retraining steps of the hyper-reduced model is already close to that of the fully-integrated model. Although the retraining frequency tends to converge to the one obtained with only POD, no strict lower bound seems to exist for the number of retraining steps after ECM, being at times lower than the one from the fully-integrated model.

Aside from affecting retraining frequency, the number of ECM domains also dictates the total number of constitutive model computations and matrix assembly operations. Both of these aspects affect the speed-up of the hyper-reduced model. With $n_s = 1$ and $1 \leq n_\phi \leq 100$, the acceleration levels obtained for the equilibrium problem with respect to the full-order solution are shown in Figure 20. For $1 \leq n_\phi \leq 10$, the acceleration sharply increases due to the decrease in retraining frequency. After this point, the frequency stabilizes and further increases in the number of domains do not affect the acceleration level considerably. For $n_\phi = 10$, the total time dedicated to solving the equilibrium problem is 279s, of which 4.1s are spent on updating Ψ , $\mathcal{Z}$, and ϖ . Comparing the time spent by the model from Section 3.2 on updating Ψ (0.54s), we see that hyper-reduction significantly increases the overhead associated with updating the reduced bases, although it still represents a very small fraction of the total execution time. Interestingly, this overhead does not increase as we add more ECM domains, as the complexity of the minimization problem of Equation (35) scales linearly with the total number of integration points of each domain.




FIGURE 19 Number of retraining steps necessary to maintain precision for models with different combinations of n_ϕ and n_s . The number of retraining steps associated with the fully-integrated POD solution is shown for comparison. POD, proper orthogonal decomposition [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 20 Speed-ups with respect to the full-order solution for models with different numbers of ECM domains ($n_s = 1$). Only the time spent in the equilibrium phase of the TLS is compared. ECM, empirical cubature method; TLS, thick level set [Colour figure can be viewed at wileyonlinelibrary.com]

Although the acceleration level of approximately 5 times is higher than the one obtained with the fully-integrated partitioned POD strategy (Section 3.2), the benefit of using hyper-reduction in combination with the linear material employed here is limited: the computational effort associated with computing stresses at the integration points is already very low. It follows that the speed-ups obtained here may be seen as the lower bound of what can be expected from the present hyper-reduction approach. Furthermore, acceleration can be expected if nonlinear material models are employed (cf^{23,29}). However, it is worth mentioning that for nonlinear material models with history, an additional history recovery component must be added to the framework since history must be known at every integration point during retraining steps.

4 | ADDITIONAL EXAMPLES

In this section, two additional numerical examples are shown in order to further assess the performance of the reduction framework. The first is a beam loaded in three-point bending with a single straight crack. The second is a doubly-notched plate loaded in tension with two curved cracks that eventually merge. Apart from demonstrating that the framework can be used on models with various loading conditions and crack topologies, we investigate the effect of changing two parameters which have been kept fixed until now, namely the thickness l_c of the damaged zone and the force tolerance ϵ_{force} .

4.1 | Three-point bending

The first example concerns the three-point bending test shown in Figure 21. The geometry is taken from Bernard et al¹¹ but here we model the beam in plane strain. The material properties are the same ones used for the compact tension example (Section 2.4). The beam is discretized with 73 498 constant-strain triangles with one integration point each ($N = 74\,046$,

FIGURE 21 Three-point bending: finite element discretization, boundary conditions, and crack topology with $l_c = 5$ mm (dimensions in mm) [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 22 Three-point bending: number of retraining steps for multiple values of ϕ_f and two different values for l_c [Colour figure can be viewed at wileyonlinelibrary.com]





FIGURE 23 Three-point bending: number of retraining steps for multiple sizes of the $\mathcal{F}$ domain [Colour figure can be viewed at wileyonlinelibrary.com]



$M = 73\,498$). For the hyper-reduced model, $n_f = 6$, $\epsilon_{\text{force}} = 0.025$, $\epsilon_{\text{SV}} = 1 \times 10^{-6}$, $\tau = 10$, $\beta = 0.001$, $\epsilon_{\text{greedy}} = 1 \times 10^{-10}$, and $n_\phi = n_s = 10$.

Since the introduction of Y_c values related to strength and fracture energy decouples l_c from G_c (see Equation (14) and Reference 14), we are in principle free to choose the most convenient value for l_c and still obtain the same crack propagation behavior. As the definitions of $\mathcal{F}$ and $\mathcal{P}$ are directly linked to the size of the damaged band, it is interesting to investigate how the performance of the framework is affected by l_c . We therefore solve the same problem for both $l_c = 5$ mm and $l_c = 15$ mm.

We start by investigating the influence of ϕ_f on the retraining frequency of the POD-reduced model (fully integrated). The number of fully-solved steps for $-40 \leq \phi_f \leq 15$ is plotted in Figure 22 for both values of l_c . Although the value of ϕ_f at which the retraining frequency stabilizes changes with l_c ($\phi_f = -20$ mm for $l_c = 5$ mm and $\phi_f = -10$ mm for $l_c = 15$ mm), the total width of the material band that comprises $\mathcal{F}$ ($w_F = 2l_c + 2\phi_f$) at stabilization seems to be the same for both models.

Most interesting is the fact that the model with $l_c = 5$ mm needs almost three times more retraining steps than the model with $l_c = 15$ mm when both are executed with an empty[§] $\mathcal{F}$ domain. This can be more clearly visualized in Figure 23, where we plot the size of $\mathcal{F}$ versus the number of retraining steps. Increasing l_c leads to a lower stress at the crack tip, which in turn leads to a slower buildup of the norm of the deviation from equilibrium used to trigger retraining

[§]With the exception of DOFs of enriched nodes not currently covered by the basis Ψ (see Section 3.2).




FIGURE 24 Three-point bending: load-displacement curves of the full and hyper-reduced models ($l_c = 5$ mm)




FIGURE 25 Three-point bending: domain topologies at two different load steps for $l_c = 5$ mm (top) and $l_c = 15$ mm (bottom). For both models, $n_\phi = n_s = 10$




FIGURE 26 Doubly-notched plate: finite element discretization and crack topology (dimensions in mm) [Colour figure can be viewed at wileyonlinelibrary.com]

(Equation (25) and Figure 8). This suggests the possibility of increasing l_c with the specific purpose of improving the efficiency of the reduced model. However, the rate of decrease in retraining frequency is steeper for the model with smaller l_c , and any acceleration gained by increasing l_c is quickly compensated by the sharp reduction in retraining steps as DOFs are added to $\mathcal{P}$.

For this specific example, the ϕ_f values that lead to the highest speed-ups are $\phi_f = -10$ mm for $l_c = 5$ mm and $\phi_f = 0$ mm for $l_c = 15$ mm. This corresponds to an $\mathcal{P}$ domain with $w_F = 2l_c + 2\phi_f = 30$ mm for both models and leads to a similar number of retraining steps (Figure 22).

Adopting $\phi_p = \phi_f - 10$, the time spent on the hyper-reduced equilibrium problem is 4.9 times shorter and 5.2 times shorter than the full-order one for $l_c = 5$ mm and $l_c = 15$ mm, respectively. Although this measure of speed-up only includes the time spent solving the equilibrium phase of the TLS, it is interesting to note that the model with $l_c = 5$ mm spends only 25% of the total analysis time on the two remaining TLS phases, while the model with $l_c = 15$ mm spends 53% of the total time in those phases due to its larger Ω_d domain.

A comparison in terms of load-displacement response between the full and reduced models is shown in Figure 24, while domain topologies at two different load steps are shown in Figure 25.

4.2 | Doubly-notched plate

The final example is the doubly-notched square plate of Figure 26, inspired by the crack merging example by Moës et al.¹⁰ Material properties are kept the same as the ones from the previous examples and a value $l_c = 0.6$ mm is adopted. The full-order model is discretized with 86 068 constant-strain triangles and has $N = 86\,696$ DOFs and $M = 86\,068$ integration points at the beginning of the analysis. Two damage fronts nucleate at the notches and grow towards the center of the plate. The interaction between the fronts cause their trajectory to become curved and eventually merge, leaving a region of intact material at the center of the plate.

We initially adopt the same tolerance $\epsilon_{\text{force}} = 0.025$ used in the previous examples. By running the POD-reduced model for different sizes of $\mathcal{F}$, we find that the combination $\phi_f = -1.2$ mm, $\tau = 7$, $\beta = 0.001$ gives the best balance between the number of DOFs in $\mathcal{F}$ and the number of retraining steps. The optimum value for τ we find here is different from the one used in the previous examples ($\tau = 10$), suggesting that optimum values for the viscosity parameters used to define $\mathcal{F}$ depend to some extent on the problem being solved. With $n_\phi = 100$, $n_s = 1$, and $\phi_p = -2.4$ mm, the integration domain topology at four different load steps is shown in Figure 27, together with the associated displacement u_{top} of the top edge of the plate.

The adaptive reduction approach is capable of automatically following the two damage fronts simultaneously while taking into account changes in propagation direction. We note that even though some reduction domains are spatially




FIGURE 27 Doubly-notched plate: ECM domain topology and crack shapes at four different load steps ($n_\phi = 100$, $n_s = 1$). ECM, empirical cubature method




FIGURE 28 Doubly-notched plate: ECM domain topology and crack shapes at four different load steps ($n_\phi = n_s = 10$). ECM, empirical cubature method





FIGURE 29 Doubly-notched plate: load-displacement curve for the full model and hyper-reduced models with different ϵ_{force} tolerances

disconnected—for example, the $\mathcal{P}$ domain or the domain $\mathcal{H}_{100}$ at the edges of the models—they are treated as single domains for training and numerical integration purposes. As the damage fronts start to merge, the $\mathcal{P}$ domain becomes distorted as the velocity ahead of the crack tip and towards the center of the plate start to vanish. Nevertheless, the model is able to maintain an efficient retraining frequency, with 34 retraining steps out of a total of 440 steps. For $n_\phi = n_s = 10$, the more irregular domain topology shown in Figure 28 is obtained due to the fact that a single s coordinate system is used for both damage fronts (when switching between fronts, the value of s at the starting point of the second front is made to be the same as the value at the final point along the first front). In any case, the performance in terms of acceleration and retraining frequency is similar for both models.

The load versus the displacement at the top of the plate is plotted in Figure 29. Although the curve with $\epsilon_{\text{force}} = 0.025$ closely follows the one from the full-order model, the precision is lower close to the onset of global softening. Although the correct crack topology is still obtained and the post-peak behavior is correctly captured, this loss of precision might

FIGURE 30 Doubly-notched plate: speed-ups with respect to the full-order solution for models with different ϵ_{force} tolerances. The marked points correspond to the two models shown in Figure 29



be undesirable. Running the same model with $\epsilon_{\text{force}} = 0.01$ eliminates the oscillations and produces a response almost indistinguishable from the full-order one (Figure 29).

Naturally, the increased precision obtained with a lower ϵ_{force} comes at the cost of additional retraining steps and reduced speed-ups. We execute the model for multiple values of ϵ_{force} and plot the obtained acceleration levels in Figure 30. The speed-up appears to decrease linearly as ϵ_{force} is decreased up until a value around $\epsilon_{\text{force}} = 0.005$, after which point a sharp loss in efficiency is observed. In any case, acceleration levels similar to the ones from the previous examples are obtained even though the framework needs to cope with two cracks instead of one.

5 | CONCLUSIONS

MOR techniques have been used in an attempt to accelerate the mechanical equilibrium problem that constitutes the main computational bottleneck associated with the TLS method. The need for an *offline* training phase was circumvented by resorting to adaptive reduction techniques that refine their reduced solution spaces through *online* retraining. The framework is based on a domain decomposition strategy that uses the curvilinear coordinate system associated with the level set front to adaptively define a region with fully-solved DOFs, a POD-reduced region with full integration and a user-defined number of hyper-reduced regions.

The framework was gradually constructed starting from a POD-reduced model without partitioning. Although the reduced model was able to capture the correct crack propagation behavior, a large number of retraining steps was necessary because the reduced solution space could not represent changes in damage distribution ahead of the crack tip. Furthermore, the unpartitioned POD solution has been found to be incompatible with the enrichment scheme used in the TLS.

The equilibrium system was therefore partitioned in order to allow for part of the mesh to be solved in the full-order space. The level set and velocity fields were used to define a region with full DOFs that includes enriched nodes and regions both behind and ahead of the crack tip. The partitioning strategy allowed for the enrichment scheme to be used and was found to drastically decrease the number of retraining steps necessary to maintain precision, leading to higher levels of acceleration.

Domain-based hyper-reduction based on the ECM method was used to reduce the effort associated with computing the full-order internal force vector and stiffness matrix. In order to still be able to solve DOFs in the full-order space and estimate the error associated with model reduction, full integration was maintained on regions close to the crack tip. Using a single ECM domain was found to yield a high integration error and increase retraining frequency. This error vanishes as the number of domains is increased, while maintaining a reduction of several orders of magnitude in the number of integration points.

The reduction framework was demonstrated using three different numerical examples. The method was found to be suitable for different loading scenarios and crack topologies and the adaptive domain decomposition strategy was able to follow multiple cracks with complex propagation paths simultaneously. However, a number of parameters were found to be problem-dependent and had to be individually tuned for each example. The hyper-reduced equilibrium problems ran approximately 5 times faster than their full-order counterparts, with even higher speed-ups being expected if more

complex constitutive models are used. Although the framework has been applied to the specific case of the TLS method, it can also be employed to accelerate other failure models as long as suitable strategies for defining the domains $\mathcal{F}$, $\mathcal{P}$, and $\mathcal{H}$ are developed.

ACKNOWLEDGEMENTS

FM acknowledges financial support from the Netherlands Organization for Scientific Research (NWO) under Vidi grant nr. 16464.

ORCID

Iuri. B. C. M. Rocha  <https://orcid.org/0000-0001-8410-3741>

Frans. P. van der Meer  <https://orcid.org/0000-0002-6691-1259>

REFERENCES

1. Krueger R. Virtual crack closure technique: history, approach, and applications. *Appl Mech Rev.* 2004;57:109-143.
2. de Borst R, Sluys LJ, Muhlhaus HB, Pamin J. Fundamental issues in finite element analyses of localisation of deformation. *Eng Comput.* 1993;10:99-121.
3. Bažant ZP, Oh B. Crack band theory for fracture of concrete. *Mater Constr.* 1983;16:155-177.
4. Li Piani T, Weerheijm J, Koene L, Sluys LJ. The adobe delta damage model: a locally regularized rate-dependent model for the static assessment of soil masonry bricks and mortar. *Eng Fract Mech.* 2019;206:114-130.
5. Pijaudier-Cabot G, Bažant ZP. Nonlocal damage theory. *ASCE J Eng Mech.* 1987;113:1512-1533.
6. Peerlings RHJ, de Borst R, Brekelmans WAM, de Vree HP. Gradient enhanced damage for quasi-brittle materials. *Int J Numer Methods Eng.* 1996;39:3391-3403.
7. Geers MGD, de Borst R, Brekelmans WAM, Peerlings RHJ. Strain-based transient-gradient damage model for failure analyses. *Comput Method Appl M.* 1998;160:133-153.
8. Simone A, Wells GN, Sluys LJ. From continuous to discontinuous failure in a gradient-enhanced continuum damage model. *Comput Method Appl M.* 2003;192:4581-4607.
9. Miehe C, Welschinger F, Hofacker M. Thermodynamically consistent phase-field models of fracture: variational principles and multi-field fe implementations. *Int J Numer Methods Eng.* 2010;83:1273-1311.
10. Moës N, Stolz C, Bernard PE, Chevaugeon N. A level set based model for damage growth: the thick level set approach. *Int J Numer Methods Eng.* 2011;86:358-380.
11. Bernard PE, Moës N, Chevaugeon N. Damage growth modeling using the Thick Level Set (TLS) approach: efficient discretization for quasi-static loadings. *Comput Method Appl M.* 2012;233-236:11-27.
12. Mororó LAT, van der Meer FP. Combining the thick level set method with plasticity. *Eur J Mech A-Solid.* 2020;79:14.
13. Lé B, Moës N, Legrain G. Coupling damage and cohesive zone models with the thick level set approach to fracture. *Eng Fract Mech.* 2018;193:214-247.
14. van der Meer FP, Sluys LJ. The thick level set method: sliding deformations and damage initiation. *Comput Method Appl M.* 2015;285:64-82.
15. Latifi M, van der Meer FP, Sluys LJ. Fatigue modeling in composites with the thick level set interface method. *Compos Part A-Appl S.* 2017;101:72-80.
16. Kerfriden P, Gosselet P, Adhikari S, Bordas SPA. Bridging proper orthogonal decomposition methods and augmented newton-krylov algorithms: an adaptive model order reduction for highly nonlinear mechanical problems. *Comput Method Appl M.* 2011;200:850-866.
17. Hernández JA, Caicedo MA, Ferrer A. Dimensional hyper-reduction of nonlinear finite element models via empirical cubature. *Comput Method Appl M.* 2017;313:687-722.
18. Galland F, Gravouil A, Malvesin E, Rochette M. A global model reduction approach for 3d fatigue crack growth with confined plasticity. *Comput Method Appl M.* 2011;200:699-716.
19. Kerfriden P, Gouy O, Rabczuk T, Bordas SPA. A partitioned model order reduction approach to rationalise computational expenses in nonlinear fracture mechanics. *Comput Method Appl M.* 2013;256:169-188.
20. Ryckelynck D. A priori hyperreduction method: an adaptive approach. *J Comput Phys.* 2005;202:346-366.
21. Ryckelynck D, Missoum Benziane D, Cartel S, Besson J. A robust adaptive model reduction method for damage simulations. *Comput Mater Sci.* 2011;50:1597-1605.
22. Kerfriden P, Passieux JC, Bordas SPA. Local/global model order reduction strategy for the simulation of quasi-brittle failure. *Int J Numer Methods Eng.* 2012;89:154-179.
23. Rocha IBCM, van der Meer FP, Sluys LJ. An adaptive domain-based POD/ECM hyper-reduced modeling framework without offline training. *Comput Method Appl M.* 2020;358:28pp.
24. Liu Z, Bessa M, Liu WK. Self-consistent clustering analysis: an efficient multi-scale scheme for inelastic heterogeneous materials. *Comput Method Appl M.* 2016;306:319-341.
25. Li S, Thouless MD, Waas AM, Schroeder JA, Zavattieri PD. Use of a cohesive-zone model to analyze the fracture of a fiber-reinforced polymer-matrix composite. *Compos Sci Technol.* 2005;65:537-549.

26. van der Meer FP, Moës N, Sluys LJ. A level set model for delamination – modeling crack growth without cohesive zone or stress singularity. *Eng Fract Mech.* 2012;79:191-212.
27. Moës N, Belytschko T. Extended finite element method for cohesive crack growth. *Eng Fract Mech.* 2002;69:813-833.
28. Burman E, Claus S, Hansbo P, Larson MG, Massing A. Cutfem: discretizing geometry and partial differential equations. *Int J Numer Methods Eng.* 2015;104:472-501.
29. Rocha IBCM, van der Meer FP, Sluys LJ. Efficient micromechanical analysis of fiber-reinforced composites subjected to cyclic loading through time homogenization and reduced-order modeling. *Comput Method Appl M.* 2019;345:644-670.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Rocha IBCM, van der Meer FP, Mororó LAT, Sluys LJ. Accelerating crack growth simulations through adaptive model order reduction. *Int J Numer Methods Eng.* 2020;121:2147–2173.
<https://doi.org/10.1002/nme.6303>