Journal of Computational Physics 229 (2010) 2061–2092

Contents lists available at ScienceDirect

Journal of Computational Physics

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

Adaptive spacetime method using Riemann jump conditions for coupled atomistic–continuum dynamics

B. Kraczek [a,1], S.T. Miller [b], R.B. Haber [b,*], D.D. Johnson [a,b,c]

[a] Dept. of Physics, University of Illinois at Urbana–Champaign, 1110 West Green Street, Urbana, Illinois 61801, USA
[b] Dept. of Mechanical Science and Engineering, University of Illinois at Urbana–Champaign, 1304 West Green Street, Urbana, Illinois 61801, USA
[c] Dept. of Materials Science and Engineering, University of Illinois at Urbana–Champaign, 1206 West Green Street, Urbana, Illinois 61801, USA

ARTICLE INFO

Article history:
Received 13 August 2007
Received in revised form 11 November 2009
Accepted 12 November 2009
Available online 26 November 2009

Keywords:
Multiscale modeling
Discontinuous Galerkin
Spacetime finite element
Molecular dynamics
Atomistic–continuum coupling

ABSTRACT

We combine the Spacetime Discontinuous Galerkin (SDG) method for elastodynamics with the mathematically consistent Atomistic Discontinuous Galerkin (ADG) method in a new scheme that concurrently couples continuum and atomistic models of dynamic response in solids. The formulation couples non-overlapping continuum and atomistic models across sharp interfaces by weakly enforcing jump conditions, for both momentum balance and kinematic compatibility, using Riemann values to preserve the characteristic structure of the underlying hyperbolic system. Momentum balances to within machine-precision accuracy over every element, on each atom, and over the coupled system, with small, controllable energy dissipation in the continuum region that ensures numerical stability. When implemented on suitable unstructured spacetime grids, the continuum SDG model offers linear computational complexity in the number of elements and powerful adaptive analysis capabilities that readily bridge between atomic and continuum scales in both space and time.

A special trace operator for the atomic velocities and an associated atomistic traction field enter the jump conditions at the coupling interface. The trace operator depends on parameters that specify, at the scale of the atomic spacing, the position of the coupling interface relative to the atoms. In a key finding, we demonstrate that optimizing these parameters suppresses spurious reflections at the coupling interface without the use of non-physical damping or special boundary conditions.

We formulate the implicit SDG–ADG coupling scheme in up to three spatial dimensions, and describe an efficient iterative solution scheme that outperforms common explicit schemes, such as the Velocity Verlet integrator. Numerical examples, in $1d \times$ time and employing both linear and nonlinear potentials, demonstrate the performance of the SDG–ADG method and show how adaptive spacetime meshing reconciles disparate time steps and resolves atomic-scale signals in the continuum.

1. Introduction

The majority of numerical methods used in materials simulations fall into one of two categories – those derived solely from continuum theory and those derived from atomistic models. These disparate mathematical models represent different

* Corresponding author.
E-mail addresses: bkraczek@utexas.edu (B. Kraczek), scott.miller@psu.edu (S.T. Miller), r-haber@illinois.edu (R.B. Haber), duanej@illinois.edu (D.D. Johnson).
[1] Present address: Institute for Computational Engineering and Sciences, University of Texas, Austin, Texas 78712, USA.

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.11.023

aspects of physical response associated with distinct scales in a given material. However, some physical phenomena involve coupling across these scales and are not directly accessible to methods from either category alone. The goal of coupled atomistic–continuum (AtC) modeling is to address these more difficult problems by combining atomistic and continuum models in a single simulation. The main challenge lies in finding an efficient coupling scheme that captures the interplay between the models while avoiding numerical artifacts. The SDG–ADG method we introduce here uses novel methods that allow energy and momentum to flow freely between the component models, in both directions, as is required in non-equilibrium problems.

1.1. Atomistic–continuum modeling

Atomistic methods typically employ molecular dynamics (MD) or Monte Carlo techniques. Many materials properties can be computed atomistically with a computationally accessible number of atoms, where aspects of the bulk response, such as temperature, particle density and strain state, are controlled by initial values and boundary conditions. However, atomistic simulations are severely limited in the number of atoms that can be considered, the complexity of the atomic interactions and the length of time that can be simulated due to the necessity of treating all atomistic degrees of freedom (d.o.f.). The largest simulations use about a billion atoms, far too few to capture many complex phenomena encountered in real-world materials behavior.

Continuum models typically use finite element, finite-volume or boundary integral methods. These are better suited to model larger volumes of material and longer time intervals than atomistic methods because the number of d.o.f. in a discrete continuum model can be varied to achieve a reasonable balance between computational cost and accuracy. Continuum calculations succeed in a wide range of problems where homogenized or empirical models capture the micro-scale physics reasonably well. However, they break down in situations where continuum theory is unable to capture the relevant material behavior, such as when atomic-scale defects govern the response.

A number of AtC methods have been developed to address multi-scale problems, by combining the strengths of atomistic and continuum methods. This section presents a selective review of existing AtC methods for solids that are relevant to our present work, and we refer the reader to review articles [1–4] for a more complete discussion. It is useful to distinguish AtC modeling problems where the objective is to model the detailed dynamics and kinetics of atomic-scale features, such as dislocations and phase boundaries, as they interact with larger-scale phenomena from those problems where the goal is to replace empirical continuum constitutive relations with relations derived directly from atomistic models. E and co-workers [5,6] refer to these as type-A and type-B problems, respectively.

For type-B problems, it suffices to use hierarchical coupling between macroscopic continuum and microscopic atomistic problems that coexist at a given location in the problem domain. State information from the macroscopic solution delivers boundary conditions and constraints that drive the microscopic atomistic model, while flux and/or energetic information from the microscopic solution provide constitutive information for the macroscopic problem. Depending on the time scales and the number of parameters involved, hierarchical coupling may be either concurrent or serial. Concurrent coupling requires the macro- and microscopic problems to be solved simultaneously, but under certain conditions, it is sufficient to use serial coupling wherein the microscopic response is pre-computed and stored for use in subsequent macroscopic solutions. Serially coupled AtC schemes include coarse-graining of the classical Hamiltonian through constraints [7] or renormalization group methods [8]. Concurrently coupled methods include E and Engquist's use of classical homogenization methods in [5] and the computation of constitutive relations by sampling MD cells constrained to follow the continuum deformation, as implemented with a finite-volume method by Li and E in [6] and a discontinuous Galerkin method by Wang et al. in [9]. Peridynamics theory models macroscopic continuum response with a non-local constitutive relation that mimics the integrated effects of interatomic forces [10,11]. Although it is a continuum method, peridynamics presents an alternative to type-B AtC models in some situations.

Type-A dynamic problems are more demanding because an effective coupling method must transfer momentum and energy between the atomistic and continuum models without spurious reflections or non-physical damping. Any atomic-scale information that is not resolved by the continuum model must be accounted for, also without reflection.

Methods for type-A problems generally employ some form of domain decomposition, where a purely atomistic model holds in one part of the problem domain and a macroscopic model governs elsewhere. The macroscopic model typically uses the material parameters of a suitable continuum constitutive theory, presumably matched to the atomistic model, but a type-B method may be employed instead (*cf.* [12,5]). Suitable conditions defined on an interface, in the form of either a sharp boundary or an overlap region, provide the coupling between the models. Most authors define the domain decomposition *a priori*, keeping the partition fixed throughout the simulation, but some propose adaptive methods to determine and adjust the decomposition during the course of a simulation [6,13].

The construction of reflection-free interfaces is a major theme in the literature on AtC coupling methods for type-A problems. Most methods enforce some form of kinematic compatibility condition between the atomistic and continuum models and remove fine-scale atomistic modes to suppress spurious reflections. To date, this has been achieved through some form of damping. One of the most popular approaches uses the generalized Langevin equation (GLE), a wave equation which damps fine-scale modes in a manner consistent with the fluctuation–dissipation theorem (atomic-scale fluctuations dissipate to an equilibrium state) [14–16]. The GLE-based methods, including the Bridging Scale Method (BSM) of Wagner and Liu [17,13] and related non-reflecting boundary methods of others [18–20], employ time-history kernels to calculate the

damping. The absorbing boundary approach of E and Huang [21,22] achieves similar ends by computing weighted averages, over the current and previous time steps, for the displacements of atoms near the interface. A third approach is used in the Bridging Domain Method (BDM) of Xiao and Belytschko [23], in which unresolved modes are damped by adjusting the atomistic velocities to match continuum-scale constraints.

1.2. Energy and momentum transport

Careful treatment of local transport of energy and momentum, within the component models and in both directions across coupling interfaces, is essential in schemes intended for problems with full coupling between atomic-scale and continuum-scale dynamics. Typically, these problems are not in thermal equilibrium, so methods that use thermostats or damping to suppress spurious reflections are ill-suited to this particular modeling task.

As an example, consider atomic-scale phenomena in the neighborhood of a dynamically propagating crack tip, where defects nucleate and generate localized heating and vibrations in response to far-field loads. The localized heating affects the on-going crack propagation and defect kinetics, while defects interact with other defects through long-range strain fields and through emission and scattering of fine-scale vibrations. When atomistically-modeled defects interact at a distance, a continuum zone must be introduced between them to obtain a tractable problem size. At the same time, the crack propagation and defect kinetics interact with the far-field dynamics via bi-directional energy and momentum transfer. To meet this modeling challenge, we require component continuum and atomistic methods that conserve energy and momentum, either intrinsically or through adaptive error control, and a non-dissipative, non-reflective coupling method. Continuum models that do not require dissipative limiters to suppress oscillations around sharp wavefronts are preferred.

It is commonly recognized that unsupported atomic-scale modes should be treated as thermal modes in the continuum domain (see *e.g.,* [7]), so several methods couple the atomic-scale modes to a continuum temperature field. The phonon method of Karpov et al. [20] employs the GLE to damp all unresolved modes passing from the atomistic to the continuum region. Thermal modes are transferred back to the atomistic domain through the random forcing term in the GLE, in a manner consistent with the thermal-equilibrium assumption. Within the BDM [23], Xiao and Belytschko control the amount of energy absorbed at the boundary according to computed temperature differences in the atomistic and continuum domains. In their method for type-A problems, Li and E [6] match the temperature of atoms overlapping the continuum to the local continuum temperature via use of a thermostat. While the latter two approaches allow for a non-uniform temperature field in the continuum, they are similar to the phonon method of Karpov et al. in that only long-wavelength modes transport energy from the atomistic region to the continuum. Their use of damping or thermostats prevents all of these methods from balancing energy and momentum between the atomistic and continuum models.

1.3. The SDG–ADG method

In this paper, we introduce the *SDG–ADG method*, a coupled AtC method for dynamic, type-A problems in solids that couples disjoint continuum and atomistic domains across a sharp interface. A common mathematical framework, shared between the continuum, atomistic and coupling–interface components, unifies the model. In contrast to previous methods, where the coupling constraints are purely kinematic, the SDG–ADG formulation uses Riemann jump conditions on the coupling interface to simultaneously enforce local momentum balance and kinematic compatibility while maintaining consistency with causality (*i.e.*, the coupling conditions preserve the characteristic structures of the continuum model and the long-wavelength limit of the atomistic model). The data for the Riemann problem includes an *atomistic traction field* and an *atomistic trace operator* that describe, respectively, effective momentum flux and velocity on the coupling boundary. An energy-equivalent mapping of the atomistic tractions to discrete atomic forces closes the coupling model.

We model the continuum domain with the Spacetime Discontinuous Galerkin (SDG) method for elastodynamics [24,25], and for the atomistic region, we introduce a new time-discontinuous, spatially discrete version of the SDG method called the Atomistic Discontinuous Galerkin (ADG) method. Each component method is unconditionally stable, features local momentum balance properties and preserves characteristic structure within its domain; no limiters are required to suppress spurious oscillations in the SDG continuum model; *cf.*, for example, [9]. A high-resolution continuum solution and a detailed specification of the atomic-scale system geometry suffice to suppress spurious reflections at the coupling interface. This contrasts with the reliance of other methods on damping as a means to suppress reflections.

In the following section, we formulate the continuum SDG, the atomistic ADG and the coupled SDG–ADG models for arbitrary spatial dimension, $d \in \{1, 2, 3\}$. The SDG method can be used with a conventional implicit solver, but its most efficient implementation uses a causality-based solver with computational cost that scales linearly with the number of spacetime elements [26–29] (examples in this paper demonstrate both solution strategies). The ADG model requires an implicit solve in each time step. We present an iterative ADG solution method that scales linearly with the number of atoms and, for this application, outperforms the explicit Velocity Verlet (VV) integrator, a popular choice for molecular dynamics codes [30].[2]

[2] Following the MD literature, we refer to the method as "Velocity Verlet". This numerical method is used in various fields under a variety of names (cf. [31,32]).

Combining the causality-based SDG solver in the continuum domain with the iterative solver in the atomistic region and on the coupling interface, we obtain an overall AtC scheme with linear scaling properties.

A monolithic discontinuous Galerkin framework spans the continuum, atomistic and coupling parts of our model. This enables precise enforcement of local balance properties between the component models and facilitates consistent selection of bases and numerical quadrature schemes to support arbitrarily high-order SDG–ADG models with optimal convergence rates. In contrast to finite-volume models, where the stencil expands as the order of accuracy increases, our SDG stencils are compact and do not expand with increasing polynomial order. Spacetime adaptive meshing effectively bridges between continuum and atomistic scales and allows efficient resolution of sharp pulses in the continuum model down to scales below the physical limits of the continuum theory. Thus, in models at finite temperature, the separation of thermal and mechanical response in the continuum can be based solely on physical considerations, rather than be dictated by limited numerical resolution in the continuum model. Our adaptive scheme restricts the spacetime mesh refinement to the trajectories of pulses and other sharp solution features. This circumvents the increased computational expense and numerical error incurred by methods that impose a uniform global time step over the spatial domain; *cf.* [33]. Some AtC methods, such as the BSM and the BDM, use subcycling to enable longer time steps in the continuum domain, but these techniques do not match the linear scaling properties, efficiency or accuracy of the adaptive SDG model.

The SDG formulation uses spacetime control volumes to develop the conservation laws, rather than the spatial control volumes used in AtC methods based on finite-volume and spatial discontinuous Galerkin methods [6,9]. This difference supports unstructured partitions of the spacetime analysis domain that enable the scalable causal solution scheme and the spacetime adaptive mesh refinement in the SDG model. Momentum balances to within machine precision and the energy error is provably dissipative for each SDG continuum element and for each segment of the coupling interface. The (modified, *cf.* Section 2.2.1) ADG model, on the other hand, conserves both momentum and energy for each atom over every time step to within the accuracy of the force integration. Thus, the stability of the coupled model is guaranteed locally as well as globally. It is feasible to reduce the continuum and coupling dissipation errors to the level of machine precision through the use of *h*- or *hp*-adaptive analysis methods.

We introduce a new method for suppressing spurious reflections at the coupling interface that does not rely on damping, time averaging or overlap between the continuum and atomistic zones. We demonstrate through simple models that, even when the impedances of the continuum and atomistic models are perfectly matched and the continuum solution is fully resolved, spurious reflections will occur if the macroscopic description of the continuum boundary geometry and the microscopic description of the atomic positions are chosen independently. The effect is similar to the reflections that occur when an arbitrary gap is left between two identical chains of atoms. In a key finding, we show that proper registration of the coupling interface with respect to the atomic positions suffices to virtually eliminate spurious reflections in our one-dimensional model. The optimal registration can be pre-computed for a given AtC system, so that the optimization has minimal impact on the overall solution cost. Thus, we realize significant reductions in algorithmic complexity and computational expense relative to methods that rely on time-history kernels to suppress spurious reflections. Overall, we obtain a reflection-free, sharp–interface coupling model that does not disturb the balance of energy in the coupled system and that circumvents the physical ambiguity and the smeared response that are inherent to overlapped models.

We present numerical results in Section 3 that demonstrate the accuracy and performance of the SDG–ADG method. In particular, we study energy error and spurious reflectance as narrow pulses traverse a sharp continuum–atomistic interface at zero temperature. The results indicate that spurious reflections diminish by several orders of magnitude when there is proper registration between the coupling interface and the atoms. Additional examples illustrate advanced capabilities of the SDG–ADG model, including nonlinear interatomic potentials, the causal SDG solution scheme and unstructured, adaptive spacetime meshing.

Our numerical implementation is restricted to one spatial dimension ($d = 1$). However, the proposed coupling method, as reflected in its formulation, is extensible to higher dimensions, and a numerical implementation for $d > 1$ is under development. Our focus in this initial study is on reflection-free transmission of pulses in materials at zero temperature, so we do not address the important issue of segregating thermal and mechanical modes at the coupling interface. We discuss prospects for implementations in higher spatial dimensions, methods for treating thermal modes and other extensions of the SDG–ADG method in Section 4.

2. Formulations of continuum, atomistic and coupled methods

The SDG method provides a consistent mathematical framework for formulation of the component continuum and atomistic models, as well as the AtC coupling strategy itself. The ADG method can be understood as the restriction of a continuum SDG method to the set of discrete atomic positions, and our AtC coupling method weakly enforces the same jump conditions on the coupling interface that apply on all interelement boundaries in the SDG method. Thus, we begin this section with a 3-field SDG formulation for linearized elastodynamics, originally presented in [25], in which displacement, velocity and strain appear as independent fields, and show that the 3-field formulation recovers the single-field displacement formulation presented in [24] when kinematic relations are enforced to eliminate velocity and strain as independent fields. We continue by specializing the 3-field continuum method to a spatially discrete 2-field ADG method. We then present the SDG–ADG coupling scheme, including the introduction of an atomistic trace operator that supports the construction of auxiliary velocity

and traction fields on the boundary of the atomistic domain. We show that the SDG–ADG scheme balances linear and angular momentum over individual elements and atoms while dissipating energy only in the continuum region and at the coupling interface. The numerical dissipation is relatively small, and the amount of dissipation can be controlled through adaptive analysis procedures.

2.1. The SDG method for continuum elastodynamics

The notions of *causal* and *non-causal boundaries*[3] play important roles in the SDG formulation. A causal spacetime boundary is one for which all characteristic directions of the governing hyperbolic system have the same orientation relative to the boundary. Thus, a causal boundary separates the spacetime dynamic domain of influence from the domain of dependence for every point on the boundary. This asymmetric dependency structure has been used to construct efficient $\mathcal{O}(N)$ solution schemes for SDG methods based on element-wise or patch-wise causal spacetime grids; see [24,28,26] for a more complete discussion. The causal portion of the piecewise smooth boundary of a spacetime region can be classified as *inflow* or *outflow*, depending on whether the characteristic directions on that part of the boundary are all inward to or outward from the region. *Horizontal* boundaries, *i.e.*, those on which the time coordinate is uniform, are always causal. Any spacetime boundary that is not causal is classified as non-causal. Thus, non-causal boundaries separate regions on which mutually-dependent solutions must be computed simultaneously, a property that increases the solution expense. *Vertical* boundaries parallel to the time axis are always non-causal. Non-vertical boundaries with non-uniform time coordinates are called *inclined* boundaries; these can be either causal or non-causal. Inclined boundaries imply grid motion relative to the reference coordinate frame.

The SDG formulations in [24,25] support spacetime grids that contain any mix of horizontal, vertical and inclined boundaries. To support this level of generality, they use differential forms and the exterior calculus on manifolds to circumvent certain technical problems that arise from the lack of a natural spacetime inner product in classical (non-relativistic) mechanics. However, here we introduce certain simplifying assumptions to rule out inclined boundaries. This restriction eliminates the need for the differential forms notation and allows us to use the more familiar tensor notation. It does not, however, limit the applicability of the proposed coupling scheme, which is fully compatible with the general formulations in [24,25]. The resulting simplified formulation suffices to develop the coupling model in full generality, since coupling interfaces are vertical. It supports all of the numerical results in this paper, with the exception of the results reported in Section 3.3 that highlight the efficient $\mathcal{O}(N)$ causal solution scheme and the powerful adaptive analysis capabilities that are possible on fully unstructured spacetime grids. Please refer to [24,25] for a detailed description of the general continuum formulations, which are the basis of our numerical implementation. The simplifying assumptions for the present development are:

- The spacetime domain is a cylinder defined as the Cartesian product of a time-independent spatial domain and a time interval. *Thus, the spacetime domain boundary has no inclined parts.*
- We use partially-structured, spacetime grids that are defined as the Cartesian product of a time-invariant, possibly unstructured partition of the spatial domain with a conventional partition of the analysis time interval. *Thus, the elements in the spacetime mesh have no inclined faces.*

Although our spacetime formulation is general for any spatial dimension, $d \in \{1,2,3\}$, we only present numerical results for coupled AtC models with $d = 1$ in this paper. Accordingly, we present component expressions for the case $d = 1$ to support our numerical results and to exemplify the tensor formulation. Numerical results for pure continuum models with $d = 2$ can be found in [24–26]. Section 4 includes a discussion of the prospects for implementing the proposed coupling scheme in up to three spatial dimensions.

2.1.1. Spacetime solution domain

Given a *spatial domain*, $\mathcal{D}_{\mathbf{x}} \subset \mathbb{E}^d$ with coordinates $\mathbf{x}$, for which $\mathbb{E}^d$ is Euclidean d-space with covariant and contravariant bases, $\{\mathbf{e}_i\}_{i=1}^d$ and $\{\mathbf{e}^i\}_{i=1}^d$, and an open *time interval*, $\mathcal{I} =]t_0, t_N[\subset \mathbb{R}$ with coordinate t, let $\mathcal{D} \subset \mathbb{E}^d \times \mathbb{R}$ with coordinates $(\mathbf{x}, t)$ be the *spacetime domain* defined by $\mathcal{D} := \mathcal{D}_{\mathbf{x}} \times \mathcal{I}$. Let $\mathcal{P}_{\mathbf{x}} = \{\mathcal{Q}_m\}_{m=1}^M$ be a regular partition of $\mathcal{D}_{\mathbf{x}}$ into open *space elements* $\mathcal{Q}_m$, and let $\mathcal{P}_{\mathcal{I}} = \{\mathcal{I}_n\}_{n=1}^N$ be a partition of $\mathcal{I}$ into open *time intervals* $\mathcal{I}_n =]t_{n-1}, t_n[$. Then $\mathcal{P} := [\{\mathcal{Q}_{(m,n)}\}_{m=1}^M]_{n=1}^N$ in which $\mathcal{Q}_{(m,n)} := \mathcal{Q}_m \times \mathcal{I}_n$ is a partition of $\mathcal{D}$ into *MN spacetime elements*.

We equip the boundaries of the spatial domains, $\partial\mathcal{D}_{\mathbf{x}}$ and $\partial\mathcal{Q}_m$, with outward unit normal vectors $\mathbf{n}_{\mathbf{x}} \in \mathbb{E}^d$ and introduce the unit vector in the positive time direction $\mathbf{e}_t$. Although no natural spacetime inner product exists for classical mechanics, from here on we adopt the convention,

$$\mathbf{a} \cdot \mathbf{e}_t = \mathbf{e}_t \cdot \mathbf{a} = 0 \quad \forall \mathbf{a} \in \mathbb{E}^d. \tag{1}$$

The boundaries of the spacetime domain and the spacetime elements are then equipped with unit normal vectors $\mathbf{n}$ as follows:

[3] Similar to *space-like* and *time-like* separation of events used in special relativity.

$$\mathbf{n}|_{(\mathbf{x},t)\in\partial\mathcal{D}} = \begin{cases} \mathbf{n}_{\mathbf{x}}|_{\mathbf{x}\in\partial\mathcal{D}_{\mathbf{x}}} & \forall(\mathbf{x},t)\in\partial\mathcal{D}_{\mathbf{x}}\times\mathcal{I}, \\ \mathbf{e}_t & \forall(\mathbf{x},t):\mathbf{x}\in\mathcal{D}_{\mathbf{x}},\ t=t_N, \\ -\mathbf{e}_t & \forall(\mathbf{x},t):\mathbf{x}\in\mathcal{D}_{\mathbf{x}},\ t=t_0, \end{cases} \tag{2a}$$

$$\mathbf{n}|_{(\mathbf{x},t)\in\partial\mathcal{Q}_{(m,n)}} = \begin{cases} \mathbf{n}_{\mathbf{x}}|_{\mathbf{x}\in\partial\mathcal{Q}_m} & \forall(\mathbf{x},t)\in\partial\mathcal{Q}_m\times\mathcal{I}_n, \\ \mathbf{e}_t & \forall(\mathbf{x},t):\mathbf{x}\in\partial\mathcal{Q}_m,\ t=t_n, \\ -\mathbf{e}_t & \forall(\mathbf{x},t):\mathbf{x}\in\partial\mathcal{Q}_m,\ t=t_{n-1}. \end{cases} \tag{2b}$$

The domain boundary is partitioned into *non-causal*, *causal inflow* and *causal outflow* parts, respectively, according to

$$\partial\mathcal{D}^{\mathrm{nc}} = \partial\mathcal{D}_{\mathbf{x}}\times\mathcal{I}, \tag{3a}$$

$$\partial\mathcal{D}^{\mathrm{ci}} = \{(\mathbf{x},t)\in\partial\mathcal{D}:\mathbf{n}(\mathbf{x},t)=-\mathbf{e}_t\}, \tag{3b}$$

$$\partial\mathcal{D}^{\mathrm{co}} = \{(\mathbf{x},t)\in\partial\mathcal{D}:\mathbf{n}(\mathbf{x},t)=\mathbf{e}_t\}. \tag{3c}$$

A similar partition of the spacetime element boundaries $\partial\mathcal{Q}_{(m,n)}$ into $\partial\mathcal{Q}^{\mathrm{nc}}_{(m,n)}$, $\partial\mathcal{Q}^{\mathrm{ci}}_{(m,n)}$ and $\partial\mathcal{Q}^{\mathrm{co}}_{(m,n)}$ applies. Fig. 1 shows the spacetime domain, a typical spacetime element and the orientation of the normal vectors on their boundaries for the case $d=2$.

2.1.2. Continuum fields and governing equations

We solve for up to three independent vector and tensor fields in $\mathbb{E}^d$ on $\mathcal{D}$: *displacement* $\mathbf{u}$, *velocity* $\mathbf{v}$ and *linearized strain* $\mathbf{E}$. In particular, our continuum model uses either a 1- or 3-field formulation involving $\{\mathbf{u}\}$ or $\{\mathbf{u},\mathbf{v},\mathbf{E}\}$. The kinematic compatibility relations for $\{\mathbf{u},\mathbf{v},\mathbf{E}\}$ on $\mathcal{D}$ are

$$\dot{\mathbf{u}}-\mathbf{v}=0, \tag{4a}$$

$$\mathrm{sym}\nabla\mathbf{v}-\dot{\mathbf{E}}=0, \tag{4b}$$

where for any field f, ∇f and $\dot{f}$ denote the gradient with respect to spatial position $\mathbf{x}$ and the partial derivative with respect to time. If the initial data impose compatible strain and displacement fields, then satisfaction of (4) implies satisfaction of the strain–displacement relation, $\mathbf{E}=\mathrm{sym}\nabla\mathbf{u}$, on $\mathcal{D}$. From here on, we assume compatible initial data. Strong enforcement of (4) delivers a reduced 1-field formulation in $\{\mathbf{u}\}$.

All partial derivatives are distributional derivatives in this formulation, so they include a jump part wherever the underlying fields suffer a discontinuity. In the context of our discontinuous Galerkin formulation, all solution fields are assumed to be continuous and suitably differentiable on the interiors of the spacetime elements. However, they might exhibit jumps across the element boundaries. Therefore, enforcing (4) on $\mathcal{D}$ is equivalent to enforcing for every element $\mathcal{Q}\in\mathcal{P}$ the system (4) on the interior of $\mathcal{Q}$ with the additional jump conditions on $\partial\mathcal{Q}$:

$$(\mathbf{u}^*-\mathbf{u})|\mathbf{n}\cdot\mathbf{e}_t|=\mathbf{0}, \tag{5a}$$

$$(\mathbf{v}^*-\mathbf{v})(1-|\mathbf{n}\cdot\mathbf{e}_t|)=\mathbf{0}, \tag{5b}$$

$$(\mathbf{E}^*-\mathbf{E})|\mathbf{n}\cdot\mathbf{e}_t|=\mathbf{0}. \tag{5c}$$

The *target values* $\{\mathbf{u}^*,\mathbf{v}^*,\mathbf{E}^*\}$ are fields defined on $\partial\mathcal{Q}$, as specified below. They provide a unified framework for enforcing prescribed boundary and initial data on $\partial\mathcal{D}$ as well as solution-dependent values that preserve the characteristic structure across element boundaries on the interior of $\mathcal{D}$.

We strongly enforce the constitutive relations for linear elastic response to define the dependent fields, *momentum density* and *stress*, as

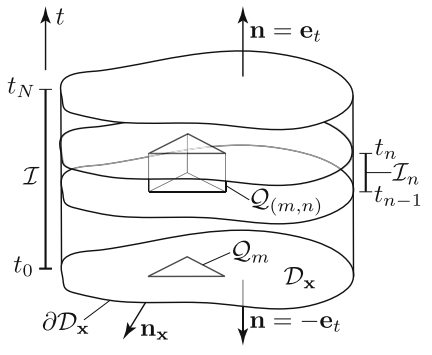

Fig. 1. Spacetime geometry for $d=2$ of solution domain, $\mathcal{D}=\mathcal{D}_{\mathbf{x}}\times\mathcal{I}$, and typical element, $\mathcal{Q}_{(m,n)}=\mathcal{Q}_m\times\mathcal{I}_n$.

$$\mathbf{p} = \rho \mathbf{v}, \tag{6a}$$
$$\boldsymbol{\sigma} = \mathbf{C}(\mathbf{E}), \tag{6b}$$

where ρ is the mass density and $\mathbf{C}$ is the elasticity tensor. We also allow for an external body force per unit mass, $\mathbf{b}$. The equation of motion expresses localized momentum balance on $\mathcal{D}$:

$$\nabla \cdot \boldsymbol{\sigma} - \dot{\mathbf{p}} + \rho \mathbf{b} = \mathbf{0}. \tag{7}$$

Similar to the compatibility relations, it suffices for every element $\mathcal{Q} \in \mathcal{P}$ to enforce (7) on the interior of $\mathcal{Q}$ subject to the additional jump condition on $\partial\mathcal{Q}$,

$$(\boldsymbol{\sigma}^* - \boldsymbol{\sigma})(\mathbf{n}) - (\mathbf{p}^* - \mathbf{p})(\mathbf{n} \cdot \mathbf{e}_t) = \mathbf{0} \tag{8}$$

in which, according to (1), $\boldsymbol{\sigma}(\mathbf{e}_t) = \boldsymbol{\sigma}^*(\mathbf{e}_t) = \mathbf{0}$.

Next we specify the target values to close the system. We partition the non-causal domain boundary according to $\partial\mathcal{D}^{\text{nc}} = \overline{\partial\mathcal{D}^{\mathbf{v}} \cup \partial\mathcal{D}^{\boldsymbol{\sigma}}}$; $\partial\mathcal{D}^{\mathbf{v}} \cap \partial\mathcal{D}^{\sigma} = \emptyset$, where $\partial\mathcal{D}^{\mathbf{v}}$ and $\partial\mathcal{D}^{\sigma}$ are the prescribed-velocity and prescribed-stress boundaries, respectively. The target values are then given by

$$\mathbf{u}^*, \mathbf{E}^*, \mathbf{p}^* = \begin{cases} \mathbf{u}, \mathbf{E}, \mathbf{p} & \text{on } \partial\mathcal{Q}^{\text{co}}, \\ \mathbf{u}^+, \mathbf{E}^+, \mathbf{p}^+ & \text{on } \partial\mathcal{Q}^{\text{ci}} \setminus \partial\mathcal{D}^{\text{ci}}, \\ \bar{\mathbf{u}}, \bar{\mathbf{E}}, \bar{\mathbf{p}} & \text{on } \partial\mathcal{Q} \cap \partial\mathcal{D}^{\text{ci}}, \end{cases} \tag{9a}$$

$$\mathbf{v}^*, \boldsymbol{\sigma}^* = \begin{cases} \mathbf{v}^+, \boldsymbol{\sigma}^+ & \text{on } \partial\mathcal{Q}^{\text{ci}} \setminus \partial\mathcal{D}^{\text{ci}}, \\ \bar{\mathbf{v}}, \boldsymbol{\sigma} & \text{on } \partial\mathcal{Q} \cap \partial\mathcal{D}^{\mathbf{v}}, \\ \mathbf{v}, \bar{\boldsymbol{\sigma}} & \text{on } \partial\mathcal{Q} \cap \partial\mathcal{D}^{\boldsymbol{\sigma}}, \\ \mathbf{v}^{\text{R}}, \boldsymbol{\sigma}^{\text{R}} & \text{on } \partial\mathcal{Q}^{\text{nc}} \setminus \partial\mathcal{D}^{\text{nc}} \end{cases} \tag{9b}$$

in which an undecorated quantity denotes the trace on $\partial\mathcal{Q}$ from the interior of $\mathcal{Q}$, a supercript '+' denotes the trace on $\partial\mathcal{Q}$ from the interior of a spacetime element adjacent (in time or space) to $\mathcal{Q}$, an overbar denotes prescribed initial or boundary data on $\partial\mathcal{D}$, and a superscript 'R' denotes a solution-dependent Riemann value on the common boundary between $\mathcal{Q}$ and an adjacent element.

Enforcing the jump conditions with respect to the Riemann values on non-causal boundaries preserves the characteristic structure of the governing equations between elements. Let $\Gamma_{\alpha\beta} := \partial\mathcal{Q}_\alpha^{\text{nc}} \cap \partial\mathcal{Q}_\beta^{\text{nc}}$ be the common boundary between spatially adjacent elements $\mathcal{Q}_\alpha$ and $\mathcal{Q}_\beta$. The Riemann values are uniquely defined as functions of traces of solution fields from both sides of $\Gamma_{\alpha\beta}^{\text{nc}}$ and are given in the Appendices of [24,25] for spatial dimensions $d = 1, 2$ (see also Section 2.1.4).

Our formulation does not directly enforce local balance of energy, which is written as

$$\dot{\mathcal{E}}_0 - \nabla \cdot [\boldsymbol{\sigma}(\mathbf{v})] - \dot{\mathbf{u}} \cdot \rho \mathbf{b} = 0, \tag{10}$$

where $\mathcal{E}_0$ is the energy density,

$$\mathcal{E}_0 = \frac{1}{2}[\mathbf{v} \cdot \mathbf{p} + \mathbf{E}(\boldsymbol{\sigma})]. \tag{11}$$

The associated jump condition on $\partial\mathcal{Q}$ is

$$(\mathcal{E}_0^* - \mathcal{E}_0)(\mathbf{n} \cdot \mathbf{e}_t) - \mathbf{v}^* \cdot \boldsymbol{\sigma}^*(\mathbf{n}) + \mathbf{v} \cdot \boldsymbol{\sigma}(\mathbf{n}) = 0 \tag{12}$$

in which the target values are computed according to (9).

2.1.3. Weighted residual and weak formulations

We next construct a finite element method based on momentum balance and kinematic compatibility. We obtain either a 3-field or 1-field formulation, depending on whether we strongly enforce the kinematic relations in (4). Let $\mathcal{U}, \mathcal{V}, \mathcal{E}$ be the discrete SDG solution spaces for the independent fields $\mathbf{u}, \mathbf{v}, \mathbf{E}$ on $\mathcal{D}$. Typically, these are piecewise-continuous spacetime polynomials of a specified order that are continuous on the interior of every spacetime element $\mathcal{Q} \in \mathcal{P}$, but that might suffer jumps across the spacetime element boundaries. Let $h_{\mathcal{Q}}$ be the spatial diameter of element $\mathcal{Q}$. We use the constant $k_{\mathcal{Q}} = \|\mathbf{C}\|/h_{\mathcal{Q}}^2$, in which $\|\mathbf{C}\|$ is the operator norm of $\mathbf{C}$, to maintain dimensional consistency.

Problem 1 (*Three-field, weighted residual form*). Find $(\mathbf{u}, \mathbf{v}, \mathbf{E}) \in \mathcal{U} \times \mathcal{V} \times \mathcal{E}$ such that for each $\mathcal{Q} \in \mathcal{P}$

$$\int_{\mathcal{Q}} \left\{ \hat{\mathbf{v}} \cdot (\nabla \cdot \boldsymbol{\sigma} - \dot{\mathbf{p}} + \rho \mathbf{b}) + \left(\nabla \mathbf{v} - \dot{\mathbf{E}} \right)(\hat{\boldsymbol{\sigma}}) + k_{\mathcal{Q}}(\dot{\mathbf{u}} - \mathbf{v}) \cdot \hat{\mathbf{u}} \right\} d\Omega + \int_{\partial\mathcal{Q}^{\text{ci}}} \{ \hat{\mathbf{v}} \cdot (\mathbf{p}^* - \mathbf{p}) + (\mathbf{E}^* - \mathbf{E})(\hat{\boldsymbol{\sigma}}) - k_{\mathcal{Q}}(\mathbf{u}^* - \mathbf{u}) \cdot \hat{\mathbf{u}} \} d\Sigma$$
$$+ \int_{\partial\mathcal{Q}^{\text{nc}}} \{ \hat{\mathbf{v}} \cdot (\boldsymbol{\sigma}^* - \boldsymbol{\sigma})(\mathbf{n}) + (\mathbf{v}^* - \mathbf{v}) \cdot \hat{\boldsymbol{\sigma}}(\mathbf{n}) \} d\Sigma = 0 \quad \forall (\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{E}}) \in \mathcal{U} \times \mathcal{V} \times \mathcal{E} \tag{13}$$

in which $\hat{\boldsymbol{\sigma}} = \mathbf{C}(\hat{\mathbf{E}})$ and where $d\Omega$ and $d\Sigma$ are the spacetime volume and surface differentials.

Problem 1 can be divided into two subproblems that can be solved sequentially. The first subproblem, defined on $\mathcal{V}\times\mathcal{E}$, is independent of the displacement solution. The second subproblem, defined on $\mathcal{U}$, can be solved in sequence using the velocity solution from the first subproblem. This sequential solution procedure is often advantageous in practice.

Integration by parts yields the weak form of the 3-field problem:

Problem 2 (*Three-field weak form*). Find $(\mathbf{u},\mathbf{v},\mathbf{E})\in\mathcal{U}\times\mathcal{V}\times\mathcal{E}$ such that for each $\mathcal{Q}\in\mathcal{P}$

$$\int_{\mathcal{Q}}\Big\{-\nabla\hat{\mathbf{v}}(\boldsymbol{\sigma})+\dot{\hat{\mathbf{v}}}\cdot\mathbf{p}+\hat{\mathbf{v}}\cdot\rho\mathbf{b}-\mathbf{v}\cdot(\nabla\cdot\hat{\boldsymbol{\sigma}})+\mathbf{E}(\dot{\hat{\boldsymbol{\sigma}}})-k_{\mathcal{Q}}(\mathbf{u}\cdot\dot{\hat{\mathbf{u}}}+\mathbf{v}\cdot\hat{\mathbf{u}})\Big\}d\Omega+\int_{\partial\mathcal{Q}^{\mathrm{nc}}}[\hat{\mathbf{v}}\cdot\boldsymbol{\sigma}^{*}(\mathbf{n})+\mathbf{v}^{*}\cdot\hat{\boldsymbol{\sigma}}(\mathbf{n})]d\Sigma$$
$$+\int_{\partial\mathcal{Q}^{\mathrm{ci}}}(\hat{\mathbf{v}}\cdot\mathbf{p}^{*}+\mathbf{E}^{*}(\hat{\boldsymbol{\sigma}})-k_{\mathcal{Q}}\mathbf{u}^{*}\cdot\hat{\mathbf{u}})d\Sigma+\int_{\partial\mathcal{Q}^{\mathrm{co}}}(-\hat{\mathbf{v}}\cdot\mathbf{p}-\mathbf{E}(\hat{\boldsymbol{\sigma}})+k_{\mathcal{Q}}\mathbf{u}\cdot\hat{\mathbf{u}})d\Sigma=0\quad\forall(\hat{\mathbf{u}},\hat{\mathbf{v}},\widehat{\mathbf{E}})\in\mathcal{U}\times\mathcal{V}\times\mathcal{E}. \tag{14}$$

Strong enforcement of (4) and the implied relation $\mathbf{E}=\mathrm{sym}\nabla\mathbf{u}$ on the interiors of all elements $\mathcal{Q}\in\mathcal{P}$ reduces the 3-field formulation to a 1-field model with $\mathbf{v}$, $\mathbf{p}$, $\boldsymbol{\sigma}$ and $\mathbf{E}$ dependent on the independent field $\mathbf{u}$. The jump parts of the kinematic compatibility relations, (5), are not enforced *a priori* and are retained in the weighted residual statement below. The interdependence of the weighting functions $\hat{\mathbf{u}}$, $\hat{\mathbf{v}}$ and $\hat{\boldsymbol{\sigma}}$ in the 1-field formulation can prevent the terms in the integrand on $\partial\mathcal{Q}^{\mathrm{ci}}$ from vanishing independently, an outcome that would invalidate the proof of element-wise momentum balance in [24]. We circumvent this problem by modifying the weighting on the residual of the displacement jump condition. We replace $\hat{\mathbf{u}}$ with $\hat{\mathbf{u}}_0$, the projection of $\hat{\mathbf{u}}$ into the zero-energy subspace of $\mathcal{U}$, *i.e.*, the subspace where $\dot{\hat{\mathbf{u}}}_0=\mathbf{0}$ and $\mathrm{sym}\nabla\hat{\mathbf{u}}_0=\mathbf{0}$ on all elements $\mathcal{Q}\in\mathcal{P}$ (*cf.* [24]).

Problem 3 (*One-field weighted residual form*). Find $\mathbf{u}\in\mathcal{U}$ such that for each $\mathcal{Q}\in\mathcal{P}$

$$\int_{\mathcal{Q}}\hat{\mathbf{v}}\cdot(\nabla\cdot\boldsymbol{\sigma}-\dot{\mathbf{p}}+\rho\mathbf{b})d\Omega+\int_{\partial\mathcal{Q}^{\mathrm{nc}}}\{\hat{\mathbf{v}}\cdot(\boldsymbol{\sigma}^{*}-\boldsymbol{\sigma})(\mathbf{n})+(\mathbf{v}^{*}-\mathbf{v})\cdot\hat{\boldsymbol{\sigma}}(\mathbf{n})\}d\Sigma$$
$$+\int_{\partial\mathcal{Q}^{\mathrm{ci}}}\{\hat{\mathbf{v}}\cdot(\mathbf{p}^{*}-\mathbf{p})+(\mathbf{E}^{*}-\mathbf{E})(\hat{\boldsymbol{\sigma}})-k_{\mathcal{Q}}(\mathbf{u}^{*}-\mathbf{u})\cdot\hat{\mathbf{u}}_0\}d\Sigma=0\quad\forall\hat{\mathbf{u}}\in\mathcal{U} \tag{15}$$

in which $\hat{\mathbf{v}}=\dot{\hat{\mathbf{u}}}$ and $\hat{\boldsymbol{\sigma}}=\mathbf{C}(\nabla\hat{\mathbf{u}})$.

An integration by parts delivers the discrete, 1-field weak problem:

Problem 4 (*One-field weak form*). Find $\mathbf{u}\in\mathcal{U}$ such that for each $\mathcal{Q}\in\mathcal{P}$

$$\int_{\mathcal{Q}}\Big(-\nabla\hat{\mathbf{v}}(\boldsymbol{\sigma})+\dot{\hat{\mathbf{v}}}\cdot\mathbf{p}+\hat{\mathbf{v}}\cdot\rho\mathbf{b}\Big)d\Omega+\int_{\partial\mathcal{Q}^{\mathrm{nc}}}\{\hat{\mathbf{v}}\cdot\boldsymbol{\sigma}^{*}(\mathbf{n})+(\mathbf{v}^{*}-\mathbf{v})\cdot\hat{\boldsymbol{\sigma}}(\mathbf{n})\}d\Sigma$$
$$+\int_{\partial\mathcal{Q}^{\mathrm{ci}}}\{\hat{\mathbf{v}}\cdot\mathbf{p}^{*}+(\mathbf{E}^{*}-\mathbf{E})(\hat{\boldsymbol{\sigma}})-k_{\mathcal{Q}}(\mathbf{u}^{*}-\mathbf{u})\cdot\hat{\mathbf{u}}_0\}d\Sigma-\int_{\partial\mathcal{Q}^{\mathrm{co}}}\hat{\mathbf{v}}\cdot\mathbf{p}\;d\Sigma=0\quad\forall\hat{\mathbf{u}}\in\mathcal{U} \tag{16}$$

in which $\hat{\mathbf{v}}=\dot{\hat{\mathbf{u}}}$ and $\hat{\boldsymbol{\sigma}}=\mathbf{C}(\nabla\hat{\mathbf{u}})$.

The three-field and one-field formulations both balance linear and angular momentum over each spacetime element $\mathcal{Q}\in\mathcal{P}$ and are provably dissipative with respect to energy balance [24,25]. The latter property implies unconditional stability, and the relatively small energy error is effectively controlled by the *h*-adaptive implementation described in [26,28], as demonstrated in Section 3.3.

2.1.4. Specialization to 1d× time

Here we give the component forms of the 3-field and 1-field weak formulations for the case of an elastic rod, $d=1$. These are the specific models used in the numerical examples in Section 3. The volume and surface differentials, $d\Omega$ and $d\Sigma$, are here replaced by specific expressions in dx and dt. Dimensional consistency requires a cross-sectional area A in one-dimensional models. We assume that A is uniform, and it is factored out in the statements below. The cross-section faces generate the non-causal boundary, $\partial\mathcal{Q}^{\mathrm{nc}}=A\times\mathcal{I}^{n}$, where the spatial component of the outward normal vector is $n_x:=\mathbf{n}\cdot\mathbf{e}_x$ in which $\mathbf{e}_x$ is the unit vector in the x-direction.

Problem 5 (*Three-field weak form; d = 1*). Find $(u,v,E)\in\mathcal{U}\times\mathcal{V}\times\mathcal{E}$ such that for each $\mathcal{Q}\in\mathcal{P}$

$$\int_{\mathcal{Q}}\Big\{-\hat{v}_{,x}\sigma+\dot{\hat{v}}p+\hat{v}\rho b-v\hat{\sigma}_{,x}+E\dot{\hat{\sigma}}-k_{\mathcal{Q}}(u\dot{\hat{u}}+v\hat{u})\Big\}dx\,dt+\int_{\partial\mathcal{Q}^{\mathrm{nc}}}(\hat{v}\sigma^{*}+v^{*}\hat{\sigma})n_x\,dt+\int_{\partial\mathcal{Q}^{\mathrm{ci}}}(\hat{v}p^{*}+E^{*}\hat{\sigma}-k_{\mathcal{Q}}u^{*}\hat{u})dx$$
$$+\int_{\partial\mathcal{Q}^{\mathrm{co}}}(-\hat{v}p-E\hat{\sigma}+k_{\mathcal{Q}}u\hat{u})dx=0\quad\forall(\hat{u},\hat{v},\widehat{E})\in\mathcal{U}\times\mathcal{V}\times\mathcal{E}. \tag{17}$$

Problem 6 (*One-field weak form; d = 1*). Find $u \in \mathcal{U}$ such that for each $\mathcal{Q} \in \mathcal{P}$

$$\int_{\mathcal{Q}} \left(-\hat{v}_{,x}\sigma + \dot{\hat{v}}p + \hat{v}\rho b\right)dx\,dt + \int_{\partial\mathcal{Q}^{\mathrm{nc}}} \{\hat{v}\sigma^* + (v^* - v)\hat{\sigma}\}n_x\,dt + \int_{\partial\mathcal{Q}^{\mathrm{ci}}} \{\hat{v}p^* + (E^* - E)\hat{\sigma} - k_{\mathcal{Q}}(u^* - u)\hat{u}_0\}dx - \int_{\partial\mathcal{Q}^{\mathrm{co}}} \hat{v}p\,dx = 0 \quad \forall \hat{u} \in \mathcal{U} \tag{18}$$

in which $\hat{v} = \dot{\hat{u}}$ and $\hat{\sigma} = C\hat{u}_{,x}$.

The target values are defined in (9), in which the Riemann values on the common boundary $\Gamma_{\alpha\beta}$ between adjacent elements $\mathcal{Q}_\alpha$ and $\mathcal{Q}_\beta$ are (*cf.* [24]):

$$v^{\mathrm{R}} = \frac{1}{2}\left\{(v^\alpha + v^\beta) + \frac{c}{C}(\sigma^\beta - \sigma^\alpha)n_x^\alpha\right\}, \tag{19a}$$

$$\sigma^{\mathrm{R}} = \frac{1}{2}\left\{(\sigma^\alpha + \sigma^\beta) + \frac{C}{c}(v^\beta - v^\alpha)n_x^\alpha\right\} \tag{19b}$$

in which C is Young's modulus, c is the elastic wave speed,

$$c = \sqrt{C/\rho} \tag{20}$$

and we have assumed uniform material properties in $\mathcal{D}$. Superscripts α and β denote traces from the interiors of elements $\mathcal{Q}_\alpha$ and $\mathcal{Q}_\beta$ or the normal vector component on $\partial\mathcal{Q}_\alpha^{\mathrm{nc}}$.

2.2. Time finite element method for molecular dynamics

Here we present a time integration scheme for molecular dynamics that derives directly from the above SDG formulation for elastodynamics. We call this the *Atomistic Discontinuous Galerkin* (ADG) method, as it is a time-discontinuous Galerkin finite element method for the spatially discrete atomistic initial value problem. This provides a unified mathematical framework for our continuum and atomistic models and facilitates our subsequent development of the new AtC coupling strategy. Such time discontinuous methods are not new, and are known to have good stability properties at the expense of dissipating energy, as discussed in [34 and 35]. However, a simple modification of the ADG method eliminates numerical dissipation to within the accuracy of the force integration. Although it is implicit, we demonstrate that the ADG method, when combined with a suitable iterative solution scheme, outperforms the velocity Verlet method, a popular explicit integrator.

2.2.1. Atomistic Discontinuous Galerkin method

We identify the spacetime reference domain for the ADG formulation with the set of atoms, $G = \{\gamma\}$, where we simultaneously use γ to index the set of atoms and to denote the γth individual atom. The trajectory of atom γ is described by the displacement–velocity pair $(\mathbf{u}_\gamma, \mathbf{v}_\gamma)$, where $\mathbf{u}_\gamma = \mathbf{x}_\gamma - \bar{\mathbf{x}}_\gamma$ is the displacement of γ from its reference position $\bar{\mathbf{x}}_\gamma$ (here the equilibrium position) to its current position $\mathbf{x}_\gamma$, and $\mathbf{v}_\gamma$ is the velocity of γ. We discretize the atomic trajectories with spatially-uniform time intervals, $\mathcal{I}_n =]t_{n-1}, t_n[$, where the discrete times t_n define the partition of the overall time interval $\mathcal{I}$. Let $\widetilde{\mathcal{U}}$ and $\widetilde{\mathcal{V}}$ denote the discrete ADG solution spaces for the atomic displacements and velocities, $\{(\mathbf{u}_\gamma, \mathbf{v}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}$, such that the components are piecewise-continuous polynomials of specified orders that might suffer jumps at the discrete times $\{t_n\}$. The displacement and velocity histories on each open interval $\mathcal{I}_n$ are single-valued and require no special notation. However, the solution can jump at the discrete times $\{t_n\}$, so we use $\{(\mathbf{u}_\gamma^m, \mathbf{v}_\gamma^m)\}|_{t=t_n}$ to denote the trace of the solution from interval $\mathcal{I}_m$ at time t_n.

The spatially discrete ADG formulation can be derived from the continuum SDG model by representing the mechanics fields as sums of Dirac delta functions centered on the atomic positions, as considered by Irving and Kirkwood [36], and by dropping the strain and stress fields which have no direct atomistic counterparts. The continuum effects of the strain and stress fields are replaced by their atomistic antecedents, the relative displacements of atoms and the non-local force interactions between atoms (*i.e.*, spatially discrete body forces). With this perspective in mind, we introduce the following constitutive model in lieu of (6).

We assume that the potential energy of the atomistic system V^{tot} depends only on atomic positions, so that in the absence of external forces it can be written in terms of the individual atomic displacements:

$$\overline{V}^{\mathrm{tot}}(\{\mathbf{x}_\gamma\}) = \overline{V} + V(\{\mathbf{u}_\gamma\}) \tag{21}$$

in which $\overline{V} = V(\{\bar{\mathbf{x}}_\gamma\})$ is the total potential energy of the system in the reference configuration, chosen to be in static equilibrium. Thus, the internal force acting on atom γ is given by

$$\mathbf{F}_\gamma^{\mathrm{int}}(\{\mathbf{u}_{\gamma'}\}) = -\nabla_{\mathbf{u}_\gamma} V(\{\mathbf{u}_{\gamma'}\}), \tag{22}$$

while the momentum of atom γ is given by

$$\mathbf{p}_\gamma = m_\gamma \mathbf{v}_\gamma \quad \text{(no sum)} \tag{23}$$

in which m_γ is the mass of γ. Let $\mathbf{F}_\gamma$, the net force acting on atom $\gamma \in G$, be decomposed according to

$$\mathbf{F}_\gamma = \mathbf{F}_\gamma^{\text{int}} + \mathbf{F}_\gamma^{\text{ext}} \tag{24}$$

in which $\mathbf{F}_\gamma^{\text{ext}}$ is the resultant of any external forces acting on γ. Then the governing equations of kinematic compatibility and the equation of motion for each atom $\gamma \in G$ reduce to

$$\dot{\mathbf{u}}_\gamma - \mathbf{v}_\gamma = \mathbf{0}, \tag{25a}$$
$$\mathbf{F}_\gamma - \dot{\mathbf{p}}_\gamma = \mathbf{0}, \tag{25b}$$

with jump parts at the start of each interval $\mathcal{I}_n$ given by

$$\left(\mathbf{u}_\gamma^n - \mathbf{u}_\gamma^{n-1}\right)\Big|_{t=t_{n-1}} = \mathbf{0}, \tag{26a}$$
$$\left(\mathbf{p}_\gamma^{n-1} - \mathbf{p}_\gamma^n\right)\Big|_{t=t_{n-1}} = \mathbf{0}. \tag{26b}$$

Following a procedure similar to the one for the continuum case, we introduce the dimensional factor, $k_\gamma = \partial^2 \overline{V}/(\partial \mathbf{x}_\gamma)^2|_{\mathbf{x}_\gamma = \bar{\mathbf{x}}_\gamma}$, and form a weighted residual statement for the ADG problem that weakly enforces (25) and (26).

Problem 7 (*ADG method: weighted residual form*). Find the atomic trajectories $\{(\mathbf{u}_\gamma, \mathbf{v}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}$ such that for each time interval $\mathcal{I}_n$

$$\sum_\gamma \int_{\mathcal{I}_n} \left[\hat{\mathbf{v}}_\gamma \cdot (\mathbf{F}_\gamma - \dot{\mathbf{p}}_\gamma) + k_\gamma(\dot{\mathbf{u}}_\gamma - \mathbf{v}_\gamma)\cdot \hat{\mathbf{u}}_\gamma\right] dt + \sum_\gamma \left[\hat{\mathbf{v}}_\gamma^n \cdot \left(\mathbf{p}_\gamma^{n-1} - \mathbf{p}_\gamma^n\right) + k_\gamma\left(\mathbf{u}_\gamma^n - \mathbf{u}_\gamma^{n-1}\right)\cdot \hat{\mathbf{u}}_\gamma^n\right]\Big|_{t=t_{n-1}} = 0 \quad \forall \{(\hat{\mathbf{u}}_\gamma, \hat{\mathbf{v}}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}. \tag{27}$$

We integrate (27) by parts to obtain the weak problem statement.

Problem 8 (*ADG method: weak form*). Find the atomic trajectories, $\{(\mathbf{u}_\gamma, \mathbf{v}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}$, such that for each time interval $\mathcal{I}_n$,

$$\begin{aligned}\sum_\gamma \int_{\mathcal{I}_n} \left[\hat{\mathbf{v}}_\gamma \cdot \mathbf{F}_\gamma + \dot{\hat{\mathbf{v}}}_\gamma \cdot \mathbf{p}_\gamma - k_\gamma(\mathbf{v}_\gamma \cdot \hat{\mathbf{u}}_\gamma + \mathbf{u}_\gamma \cdot \dot{\hat{\mathbf{u}}}_\gamma)\right] dt + \sum_\gamma \left(\hat{\mathbf{v}}_\gamma^n \cdot \mathbf{p}_\gamma^{n-1} - k_\gamma \mathbf{u}_\gamma^{n-1} \cdot \hat{\mathbf{u}}_\gamma^n\right)\Big|_{t=t_{n-1}} \\ + \sum_\gamma \left(-\hat{\mathbf{v}}_\gamma^n \cdot \mathbf{p}_\gamma^n + k_\gamma \mathbf{u}_\gamma^n \cdot \hat{\mathbf{u}}_\gamma^n\right)\Big|_{t=t_n} = 0 \quad \forall \{(\hat{\mathbf{u}}_\gamma, \hat{\mathbf{v}}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}.\end{aligned} \tag{28}$$

Let $\mathbf{c}$ be a time-independent vector field on $\mathcal{I}_n$. Since (28) holds over $\mathcal{I}_n$ for all $\{(\hat{\mathbf{u}}_\gamma, \hat{\mathbf{v}}_\gamma)\} \in \widetilde{\mathcal{U}} \times \widetilde{\mathcal{V}}$, it must be true for all weightings of the form, $\hat{\mathbf{u}}_\gamma = \mathbf{0}$, $\hat{\mathbf{v}}_\gamma = \mathbf{c}$ and $\hat{\mathbf{u}}_\kappa, \hat{\mathbf{v}}_\kappa = \mathbf{0}$ $\forall \kappa \neq \gamma$, for each atom γ in turn in each time step. Consideration of all such weightings implies balance of linear momentum for individual atoms and for the overall system to within the accuracy of the integral evaluations in each time step. Balance of angular momentum is not guaranteed for general, finite atomic motions in this formulation. However, one could argue that angular momentum balances for infinitesimal atomic excursions away from the equilibrium configuration, similar to the linearized theory of elastodynamics used in the continuum SDG formulation. Nonetheless, the formulation is consistent, so that any imbalance in angular momentum can be controlled with adaptive mesh refinement. For the overall system, linear momentum is conserved to within the accuracy of the linearized theory if $\mathbf{F}^{\text{ext}} = \mathbf{0}$ and the potential V generates self-equilibrating internal forces so that $\sum_\gamma \mathbf{F}_\gamma^n = \mathbf{0}$.

Energy analysis shows that the basic ADG method is guaranteed to be diffusive at the level of individual atoms and, therefore, stable. Setting $\hat{\mathbf{v}}_\gamma = \mathbf{v}_\gamma$, $\hat{\mathbf{u}}_\gamma = \mathbf{0}$ and $\hat{\mathbf{v}}_\kappa, \hat{\mathbf{u}}_\kappa = \mathbf{0}$ for $\kappa \neq \gamma$ in (28), we obtain after some manipulation,

$$\frac{1}{2}\left(\mathbf{v}_\gamma^n \cdot \mathbf{p}_\gamma^n\right)\Big|_{t=t_n} = \frac{1}{2}\left(\mathbf{v}_\gamma^{n-1} \cdot \mathbf{p}_\gamma^{n-1}\right)\Big|_{t=t_{n-1}} + \int_{\mathcal{I}_n} \mathbf{v}_\gamma \cdot \mathbf{F}_\gamma \, dt - \frac{1}{2}\left(\llbracket \mathbf{v}_\gamma \rrbracket \cdot \llbracket \mathbf{p}_\gamma \rrbracket\right)\Big|_{t=t_{n-1}} \tag{29}$$

in which for any quantity f, $\llbracket f \rrbracket|_{t_n} := f^{n-1}|_{t_n} - f^n|_{t_n}$. Since $\llbracket \mathbf{v} \rrbracket \cdot \llbracket \mathbf{p} \rrbracket = \llbracket \mathbf{v} \rrbracket \cdot m \llbracket \mathbf{v} \rrbracket \geqslant 0$, the ADG method is stable, and the last term in (29) is the numerical dissipation over time step $\mathcal{I}_n$ associated with atom γ.

A simple modification of the ADG method allows the dissipation jump term in (29) to vanish, so that energy balances to within the accuracy of the force integral evaluation. Specifically, we obtain per-atom energy and momentum conservation in the ADG method by replacing the weighting functions $\hat{\mathbf{u}}_\gamma$ and $\hat{\mathbf{v}}_\gamma$ in the integral of (27) with their L^2 projections into subspaces containing polynomial functions of order $p-1$, where p is the polynomial order of the parent spaces, $\widetilde{\mathcal{U}}$ and $\widetilde{\mathcal{V}}$. No projection is applied to the weighting functions that enforce the jump conditions at time t_{n-1}. This modification relaxes the weak enforcement of the governing equations in (27) to the minimum level required to support the proofs of momentum and energy balance and forces the inflow jump conditions to vanish independently. Numerical dissipation is eliminated and energy and momentum balance are exact to within the accuracy of the force integral evaluations (if an approximate quadrature is used) or the machine precision, whichever is larger. In practice, it is most efficient to choose a numerical quadrature scheme that matches the accuracy of the force integration to the order of accuracy of the underlying ADG discretization.

This modified ADG (mADG) method is closely related to the energy and momentum-conserving methods discussed by Borri and Bottasso [35] and later by Gross et al. [37] and references therein, for conforming (time-continuous) Galerkin

projections. As shown by Bottasso [38] (see also work by Hulme [39]), some classes of time finite element methods are equivalent to symplectic implicit Runge–Kutta methods, which have been of great interest in the numerical integrator community (see e.g. [32,40]). Similar to what Borri and Bottasso found, but in contrast to the work by Gross et al., the discontinuous basis and weak enforcement of temporal continuity of displacement and velocity (momentum) in the ADG formulation are critical to its successful coupling with the SDG continuum model. Coupling the conforming models in [37] to the efficient adaptive SDG solution scheme would destroy the energy and momentum balance properties.

2.2.2. Iterative solution procedure for the ADG method

The ADG model offers a number of attractive features: arbitrarily high-order accuracy with increasing polynomial order, unconditional stability that supports arbitrarily large time steps, per-atom balance of momentum and energy to within the accuracy of the force integration, and, in the present context, compatibility for coupling to the adaptive SDG continuum model. The ADG method is an implicit integrator, and due to the non-local force interactions between atoms, this implies that all atoms must be solved simultaneously within each time step. The cost of solving simultaneous nonlinear equations for very large systems of atoms is generally thought to outweigh the advantages of implicit models, and this explains the dominance of explicit integrators in MD software.

We next introduce a simple iterative solution scheme for ADG models that scales linearly with problem size. Although atomic force interactions are non-local, their range is limited in a computational setting by either an explicit cut-off radius or by the distance beyond which their magnitudes fall below the machine precision (several lattice spacings in a typical solid crystal). We exploit this property in the design of a simple iterative solution scheme in which the number of iterations required per time step is relatively small and independent of problem size. We show that this property, in combination with other favorable properties inherent to the ADG model, yields an iterative implicit integrator that is competitive with the Velocity Verlet (VV) method, a popular explicit integrator used in MD simulations. AtC coupling in solids requires accurate modeling of mechanical pulses in both the continuum and atomistic domains. In this particular context, the iterative ADG model clearly outperforms the VV method.

The iterative solution scheme begins with a non-disjoint partition of the atomistic spatial domain into a set of M overlapping regions called *bins*, $\{b_\alpha\}_{\alpha=1}^{M}$. Each bin b_α contains $N_\alpha \leqslant \overline{N}$ atoms, where $\overline{N}$ is the maximum bin size, and every atom is in at least one bin. Due to the overlaps between bins, some atoms are in more than one bin. Atoms that fall on a bin boundary are included in that bin. The width of overlap between bins must be at least equal to the effective cut-off radius for atomic force interactions, and $\overline{N}$ is chosen small enough to ensure that a direct solve of the ADG equations for the atoms within any single bin is not too expensive. In the case of solid crystals without diffusion, the assignment of individual atoms to bins can be performed once at the start of a simulation and held fixed from there on.

At the start of each new time step, we extrapolate data from the previous step to generate an initial estimate of the solution for each atom in the current time step. Then we loop over the bins in what amounts to an overlapping block-Gauss–Seidel or multiplicative Schwarz iterative solution scheme [41]. For each current bin b_α, we perform one Newton–Raphson iteration for the ADG equations to update the solutions for the N_α atoms in b_α, while fixing the solutions for all the other atoms in the system. In view of the bin-size limit $\overline{N}$, a direct method is efficient for solving the bin-wise Newton updates. We check convergence of the global ADG residual vector after each complete loop over the bins. We continue looping over the bins until the norm of the residual vector vanishes to within machine precision.

The overlap between bins ensures that all atomic force interactions are included in at least one bin-level Newton update during each loop over the bins. This hastens the propagation of information and dramatically reduces the number of iterations required for convergence relative to schemes with non-overlapping bins. In contrast to large systems arising from elliptic and parabolic models, where the number of Gauss–Seidel iterations required for convergence grows with problem size, the hyperbolic structure of our problem, the finite cut-off radius for atomic force interactions and the finite wave speeds of mechanical pulses in MD models limit the distance information must propagate in each time step before convergence is achieved. Let n_L be the number of loops over the bins required to attain convergence of the ADG model to within machine precision. In practice, n_L is relatively small and independent of problem size.

Although the implicit mADG formulation is unconditionally stable, the iterative solution scheme has a finite radius of convergence that imposes a maximum time-step size. The use of overlapping bins significantly increases the method's radius of convergence relative to methods with non-overlapping bins. As demonstrated in numerical studies below, the iterative solver's radius of convergence is sufficiently large that the iterative mADG scheme competes effectively against the explicit VV method.

We present numerical results for an example in one spatial dimension (extensions to higher dimensions are straightforward) to demonstrate the efficiency of the proposed iterative solution strategy. We consider a one-dimensional, periodic chain containing 256 atoms with atomic spacing a. The time interval is (0,256). The atoms interact through a modified Lennard-Jones potential, U^{mLJ}, given by [40]

$$U^{\mathrm{mLJ}}(r) = U^{\mathrm{LJ}}(r) - U^{\mathrm{LJ}}(r_c) - (r - r_c)\left.\frac{dU^{\mathrm{LJ}}}{dr}\right|_{r=r_c}, \tag{30a}$$

$$U^{\mathrm{LJ}}(r) = 4\epsilon\left[\left(\frac{R}{r}\right)^{12} - 2\left(\frac{R}{r}\right)^{6}\right] \tag{30b}$$

in which U^{LJ} is the standard Lennard-Jones potential, r is the distance between any two atoms, r_c is the cut-off radius beyond which atomic interactions vanish, R sets the potential length scale, ϵ sets the potential depth and we enforce C^1 continuity at $r = r_c$.[4] We choose a cut-off radius, $r_c = 2.5a$, that includes up to second-nearest-neighbor force interactions between atoms. The initial conditions for displacement and velocity correspond to a traveling pulse:

$$u(x,t) = \begin{cases} B\cos^3(kx - \omega t - \phi) & x \leqslant \left|\frac{1}{k}\left(\phi \pm \frac{\pi}{2}\right)\right|, \\ 0 & \text{otherwise} \end{cases} \tag{31}$$

in which the parameter ϕ determines the initial position of the pulse. Here we prescribe the amplitude, $B = a/40$. We solve with the iterative mADG integrator described in Section 2.2.1, so the energy balances to within the accuracy of the force integration for every atom in every time step.

We use the VV method as a reference explicit method to study the relative efficiency of the iterative mADG integrator. As one might expect, the explicit VV method delivers weaker conditions for linear momentum balance than those obtained with the implicit mADG method. It is easy to show that the VV method balances momentum globally over a system of atoms for vanishing external forces [31]. However, in contrast to the ADG and mADG methods, it generally does not balance momentum for individual atoms, due to its use of averaged interatomic forces within each time step. The VV model also does not offer the possibility of high-order accuracy, a feature that is available in (m)ADG models with variable polynomial order. Stability considerations limit the maximum step size. However, the VV integrator requires only one force evaluation per time step, while the number of force evaluations in the iterative mADG method increases with polynomial order and depends on the number of iterations per time step required to solve the implicit equations. Thus, for (m)ADG to compete against VV, we must exploit the (m)ADG method's unconditional stability by using larger step sizes and depend on the method's high-order accuracy to overcome the extra cost of multiple force evaluations per time step.

Since the evaluation of interatomic forces dominates the cost of MD simulations, we use the total number of force evaluations required to complete a simulation, denoted by N_f, as the measure of cost in our efficiency studies. An efficient integrator reduces the ratio between some error measure and the cost N_f. A suitable choice of the error measure depends on the nature and purpose of the simulation problem at hand. The error in global energy balance is a commonly used measure, but taken alone, it is not sufficient in AtC coupling problems where the propagation of mechanical pulses is a central focus. Direct measures of trajectory error and of the residuals of the governing equations can be more meaningful in these problems. Therefore, we consider multiple error measures in this efficiency study.

First, we consider efficiency relative to the global energy error. Since there are no external energy sources, the total system energy should be constant in time and equal to the energy determined by the initial data, E_0. We sample the energy of the computed solution at N discrete times t_i, and define the overall energy error as

$$\epsilon_e := \sqrt{\frac{1}{N(N-1)}\sum_{i=1}^{N}(E_i - E_0)^2}, \tag{32}$$

where E_i is the total energy at the sample time t_i. Fig. 2 depicts the energy error versus the number of force evaluations for VV and iterative ADG simulations with polynomial orders ranging from 1 to 5. We use the mADG model, so the dominant source for energy error in the ADG results is quadrature error associated with the nonlinear interatomic forces. As the time steps become shorter, the ADG errors attain their asymptotic convergence rates. For $p \geqslant 2$, the ADG errors drop below the VV error at costs where the VV error ranges from 10^{-7} to 10^{-8} and $N_f \approx 2050$. This corresponds to a time-step size equal to one eighth of the VV stability limit. Beyond this point, the energy errors for the higher-order ADG methods drop rapidly to machine-precision levels, while the VV error decreases at a much slower rate. Although a normalized energy error smaller than 10^{-7} might appear to be of no practical interest, we shall see that this level of accuracy for energy balance does not ensure accurate renderings of propagating mechanical pulses.

We define normalized displacement and velocity error measures as

$$\epsilon_{\mathbf{u}} = \frac{\|\mathbf{u} - \overline{\mathbf{u}}\|_{\mathbf{L^2}}}{\|\overline{\mathbf{u}}\|_{\mathbf{L^2}}}, \quad \epsilon_{\mathbf{v}} = \frac{\|\mathbf{v} - \overline{\mathbf{v}}\|_{\mathbf{L^2}}}{\|\overline{\mathbf{v}}\|_{\mathbf{L^2}}}, \tag{33}$$

where $\overline{\mathbf{u}}$ and $\overline{\mathbf{v}}$ are reference solutions and the L^2 norm is given by

$$\|(\cdot)\|_{L^2} := \left[\sum_{n=1}^{N}\left[\sum_{\gamma}\left[\int_{\mathcal{I}_n} |\cdot|^2\, dt\right]\right]\right]^{\frac{1}{2}} = \left[\sum_{\gamma}\int_{t_0}^{t_N} |\cdot|^2\, dt\right]^{\frac{1}{2}}. \tag{34}$$

Since no exact solution is available, we use the $p = 5$ mADG method with $\Delta t = 0.0078125$ (or 32,768 time steps) to generate $\overline{\mathbf{u}}$ and $\overline{\mathbf{v}}$. Figs. 3 and 4 show the normalized displacement and velocity errors as functions of the number of force evaluations. In both cases, the mADG errors for $p > 2$ drop below the VV error at costs where the VV error still exceeds 1%, an error level

[4] In general a C^i truncation can be constructed by subtracting from the potential the first $i+1$ terms of its Taylor series, expanded about the desired cut-off radius. Afterward the potential depth can be re-scaled, as necessary.

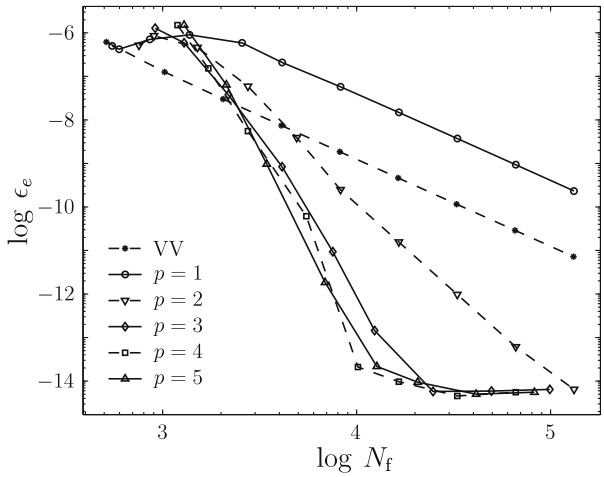

Fig. 2. Energy error ϵ_e versus cost N_f. Results are for Velocity Verlet (VV) algorithm and for iterative mADG (variable p) solution method with polynomial order p. The time-step size at the left-most VV data point corresponds to one half the VV stability limit.

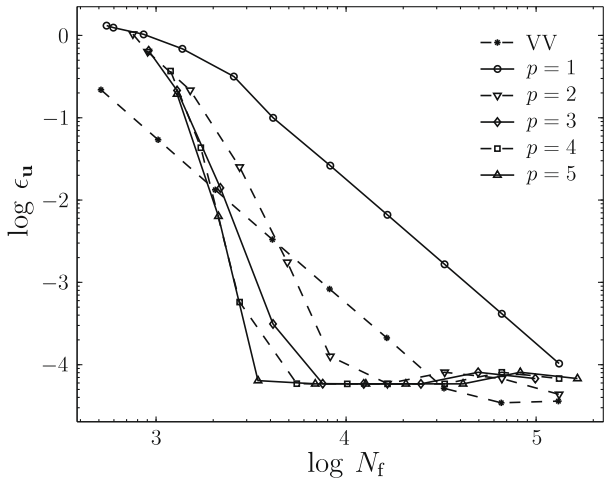

Fig. 3. Normalized displacement error $\epsilon_{\mathbf{u}}$ versus cost N_f for mADG (variable p) and VV methods.

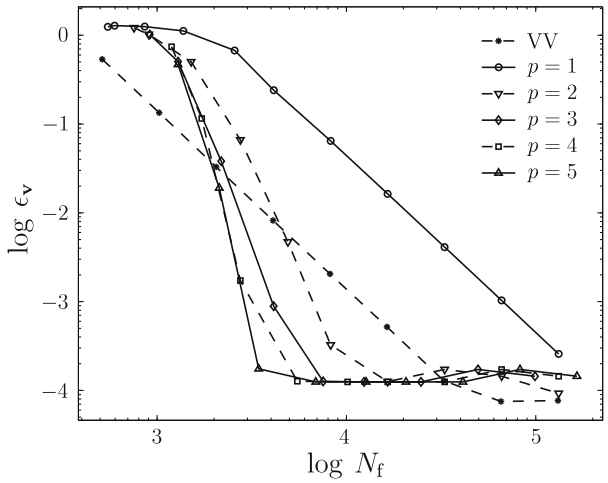

Fig. 4. Normalized velocity error $\epsilon_{\mathbf{v}}$ versus cost N_f for mADG (variable p) and VV methods.

that is too large to reliably model effects of physical interest, such as dispersion of mechanical pulses. For larger values of p, only a modest increment of cost is needed to reduce the error to machine-precision levels.

We introduce a normalized residual error measure to monitor how well the solution satisfies the equations of motion and compatibility:

$$\epsilon_{\mathrm{r}} := \left[\left(\frac{\|\dot{\mathbf{p}} - \mathbf{F}\|_{L^2}}{\|\mathbf{F}\|_{L^2}} \right)^2 + \left(\frac{\|\dot{\mathbf{u}} - \mathbf{v}\|_{L^2}}{\|\mathbf{v}\|_{L^2}} \right)^2 \right]^{\frac{1}{2}}. \tag{35}$$

Fig. 5 shows the normalized residual error as a function of cost. Again, we see that the iterative mADG method with $p \geqslant 2$ equals or outperforms the VV algorithm in the range where the residual error is below 1%.

Overall, we conclude that the high-order mADG method, when implemented with an overlapping block-Gauss–Seidel iterative solver, outperforms the VV algorithm for accuracy ranges of practical interest for AtC coupling problems in solids. Furthermore, higher-order implementations of the ADG method outperform lower-order versions. The low-order ($p = 1$) ADG method shows the same convergence rates as the VV algorithm, but in terms of efficiency, it underperforms VV due to the increased cost of solving the implicit equations. However, ADG methods of all polynomial orders share a common mathematical structure with the continuum SDG model that is critical to the success of our coupling strategy.

2.3. Dynamic response of coupled heterogeneous systems

We consider the dynamic behavior of heterogeneous coupled systems at zero temperature, comprised of some combination of continuum and atomistic zones. We focus on systems in which the impedances of the component models are matched, so that the ideal response is zero reflectance at coupling interfaces, and examine the requirements for passing waves across the interfaces without spurious reflections. We demonstrate through simple models that, in addition to matching the impedances of the component models, the details of the coupling method itself must also be carefully selected to avoid reflections. These studies motivate the sharp–interface AtC coupling model proposed in Section 2.4, in which spurious reflections are effectively suppressed, without damping, by adjusting the position of the coupling interface relative to the atomic positions.

2.3.1. Matching properties of the continuum and atomistic domains

Consider two linearly elastic rods with the same uniform, cross-section area that are coupled end-to-end, as in Fig. 6(a). The condition for non-reflection of mechanical waves at the coupling interface is that the impedances of the materials in the two rods be the same, where the impedance Z_i of material i is the product of its mass density ρ_i and wave speed c_i [42]. Together with (20), this implies

$$\sqrt{\rho_2 C_2} = \sqrt{\rho_1 C_1}. \tag{36}$$

We seek to match the physical properties of the component domains in the atomistic–continuum system in Fig. 6(b), so we employ the stronger condition that the wave speeds and mass densities are as close to one another as possible. This implies that the impedances will be likewise matched. We use a coarse-grained mass density on the atomistic side to match the

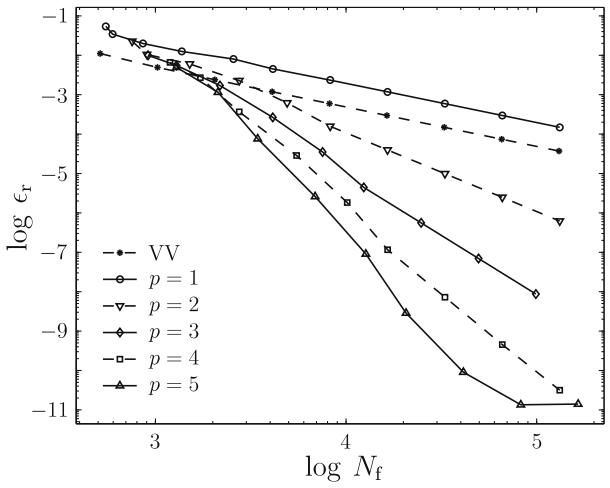

Fig. 5. Normalized residual error ϵ_{r} versus cost N_{f}. The higher convergence rates of the implicit mADG method for $p > 1$ are evident, as is the improved efficiency of higher-order mADG methods relative to the explicit VV method.

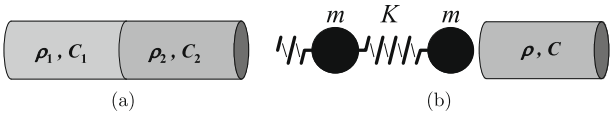

Fig. 6. Representative coupled, heterogeneous systems for $d = 1$. (a) continuum–continuum system comprised of two elastic rods joined end-to-end, and (b) atomistic–continuum system with unspecified coupling.

mass densities, as explained below. Matching the wave speeds is more difficult, since the atomistic wave speed is dispersive, while the continuum wave speed is not.

The wave speeds in an atomistic system are obtained through its dispersion relation [43]. Consider a monatomic mass–spring system with atomic mass m and uniform interatomic spacing a, interacting through linear springs out to the Sth-nearest neighbor. The dispersion relation for this system is

$$\omega^2 = \frac{1}{m}\sum_{s=1}^{S} K^s \sin^2\left(\frac{ksa}{2}\right), \tag{37}$$

where K^s is the spring constant of the sth-nearest-neighbor spring and for a pure sine wave of frequency ν and wavelength λ in either an infinite or a periodic chain of atoms, $\omega = 2\pi\nu$ and $k = 2\pi/\lambda$. The phase velocity is the velocity of each Fourier mode and is given by $c^k = \omega/k$. In the long-wavelength limit, *i.e.*, $a \ll \lambda$, the wave speed is

$$c^{\mathrm{lw}} = a\sqrt{\frac{1}{m}\sum_{s=1}^{S} s^2 K^s}. \tag{38}$$

In many cases, the difference between the dispersive atomistic solution and the non-dispersive continuum solution is small, and, as will be seen in the numerical examples, it is often sufficient to match the continuum wave speed to c^{lw}. To match the impedances of one-dimensional atomistic and continuum systems, as depicted in Fig. 6(b), we assign an effective cross-section area $A^{\mathrm{at}} = A$ to the atomistic model, in which A is the cross-section area of the continuum domain (taken to be uniform here for simplicity). The coarse-grained atomistic mass density, $\rho^{\mathrm{at}} = m/(aA)$, combines with (38) to give the atomistic impedance Z^{at} as

$$Z^{\mathrm{at}} = \sqrt{\frac{m}{A}\sum_{s=1}^{S} s^2 K^s}. \tag{39}$$

After matching the impedances and mass densities of the atomistic and continuum systems, (39) completes the relationships between the physical parameters of the one-dimensional continuum domain and the mass–spring system:

$$c = c^{\mathrm{lw}}, \tag{40a}$$

$$\rho = \frac{m}{aA}, \tag{40b}$$

$$C = \frac{a}{A}\sum_{s=1}^{S} s^2 K^s, \tag{40c}$$

where any two of these relations imply the third. These relations can be extended to address linear multi-atom interactions when the continuum interacts with the atomistic region only in the low-frequency, long-wavelength limit, which is the linear portion of the acoustic branch of the dispersion relation [43].

Although system (40) does not hold, in general, for nonlinear material models, it is applicable to AtC systems characterized by nonlinear interatomic potentials if the loading amplitude is small enough to ensure that the motions of atoms sufficiently removed from defects are limited to small oscillations about their reference positions [44]. Thus, we can reasonably couple a linearly elastic continuum model and an atomistic model having a nonlinear potential, provided that any defects or large-amplitude motions away from the reference configuration are kept within the atomistic domain and safely away from the coupling interface. The continuum and coupling parameters can then be based on a linearized potential [45,18]. For nonlinear interactions where it is either impossible or undesirable to use a linearly elastic continuum model, another criterion must be adopted to determine the bulk properties of the nonlinear continuum model, such as the widely-used Cauchy–Born hypothesis [46,47].

To illustrate the extraction of spring constants from the linearization of a nonlinear potential, consider a chain of atoms interacting through the Lennard-Jones potential (30) with reference positions $\{\bar{x}_\gamma\}$, consistent with an undeformed, equilibrium lattice constant $a = \bar{x}_{\gamma+1} - \bar{x}_\gamma$. Linearizing the system about the reference configuration, we obtain the equivalent spring constants between atoms γ and $\gamma + s$,

$$K^s_\gamma = \partial^2 U^{\mathrm{LJ}}/\partial r^2\Big|_{r=|\bar{x}_{\gamma+s}-\bar{x}_\gamma|}. \tag{41}$$

We note that K_γ^s is not affected by substituting U^{mLJ} for U^{LJ}. We use the stationary potential energy condition to compute $\{\bar{x}_\gamma\}$ and a, and apply (41) to obtain the spring constants, the first three of which are $(K_\gamma^1, K_\gamma^2, K_\gamma^3) = (1.0, -0.00427, -0.00017)$ for first, second and third-nearest neighbors of γ, where ϵ is specified such that $K_\gamma^1 = 1.0$ in suitable units of [energy]/[length]^2. The longer-ranged springs have successively smaller, negative spring constants and will often fall outside the cut-off radius imposed on the potential, so relatively few neighbors need to be included when determining C. Since the lattice parameter a determines the relative positions of all pairs of atoms in the reference configuration of the chain, all of the spring stiffnesses are fully determined by this single parameter and the potential U^{LJ}.

2.3.2. Coupling heterogeneous systems with zero reflectance

Care must be taken when devising coupling strategies, as any inhomogeneity in the overall system due to the coupling scheme itself will induce spurious reflections. In particular, we demonstrate that the atomic positions in any atomistic component of a coupled heterogeneous system must be carefully registered with respect to the geometry of neighboring components to avoid spurious reflections. Apparently, this atomic-scale registration is either ambiguous or left to chance in previous AtC coupling strategies. The following studies motivate, in part, the SDG–ADG coupling strategy described in Section 2.4 in which spurious reflections are suppressed, without damping and without invoking special boundary conditions, solely by adjusting the positions of continuum boundaries relative to the positions of atoms in adjacent atomistic domains.

To illustrate the importance of proper registration between adjacent domains, we consider coupling between a pair of distinct monatomic, uniform mass–spring systems, each interacting through first-nearest-neighbor springs. The physical parameters for system i are the number of atoms $|G_i|$, the lattice constant a_i, the effective area A_i, the atomic mass m_i and the spring constant K_i. Rather than specify K_i directly, it suits our purposes to instead specify $\overline{F}_i$, the spring force at 100% coarse-grained strain, so that $K_i = \overline{F}_i/a_i$. For $R := a_1/a_2$ and $A_1 = A_2 = A$, the two systems will have the same coarse-grained mass densities, $\rho_i = m_i/(a_i A_i)$, and the same long-wavelength wave speeds, c_i^{lw}, if

$$m_2 = \frac{1}{R} m_1, \tag{42a}$$

$$\overline{F}_2 = \overline{F}_1 \iff K_2 = R K_1. \tag{42b}$$

Eq. (42b) ensures that all springs carry the same force when the entire coupled system is subjected to a uniform strain state. Without this constraint, it would be impossible to satisfy equilibrium at the coupling interface under static, constant-strain conditions – a necessary, but not sufficient, condition for reflection-free coupling.

The paired mass–spring systems do not imply a unique means to specify the coupling. Possible coupling methods include a coupling mass m_c, as represented in Fig. 7(a), and a coupling spring, characterized by lattice parameter a_c and subject to

$$\overline{F}_c = \overline{F}_1 = \overline{F}_2, \tag{43}$$

so that the coupling spring constant is $K_c = \overline{F}/a_c$, *cf.* Fig. 7(b). As shown in Fig. 8, arbitrary choices for a_c or m_c produce reflections at the coupling boundary that are qualitatively similar to reflections of a pulse off a knot in a string; they all vanish in the long-wavelength limit [48].

If the interface is a feature of the coupling model, but not of the material we are representing, then the reflection is said to be *unphysical* or *spurious*, and we seek some means to suppress it without otherwise disturbing the solution. For example, suppose we start a traveling pulse in an otherwise quiescent domain 1, and let $\tilde{t}$ be a time when the pulse has passed completely out of domain 1 into domain 2. Since domain 1 should return to a quiescent state at time $\tilde{t}$, minimizing the l^2-norm of the atomic displacements in region 1 with respect to either m_c or a_c (with $\overline{F}_c = \overline{F}_1 = \overline{F}_2$ held fixed) is equivalent to minimizing the spurious reflection. Thus, the optimal coupling lattice parameter and coupling mass are given by

$$a_c^* = \arg\min_{a_c} \sqrt{\sum_{\gamma \in G_1} u_\gamma^2\big|_{\tilde{t}}} = \frac{a_1 + a_2}{2}, \tag{44a}$$

$$m_c^* = \arg\min_{m_c} \sqrt{\sum_{\gamma \in G_1} u_\gamma^2\big|_{\tilde{t}}} = \frac{m_1 + m_2}{2}. \tag{44b}$$

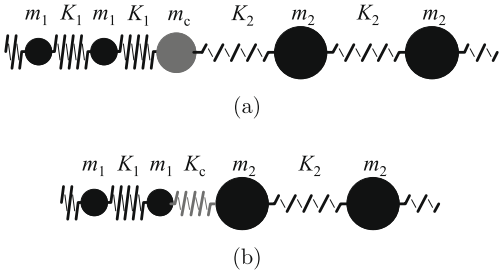

Fig. 7. Coupled monatomic mass–spring systems using (a) coupling mass m_c and (b) coupling spring with constant $K_c := \overline{F}_c/a_c$.

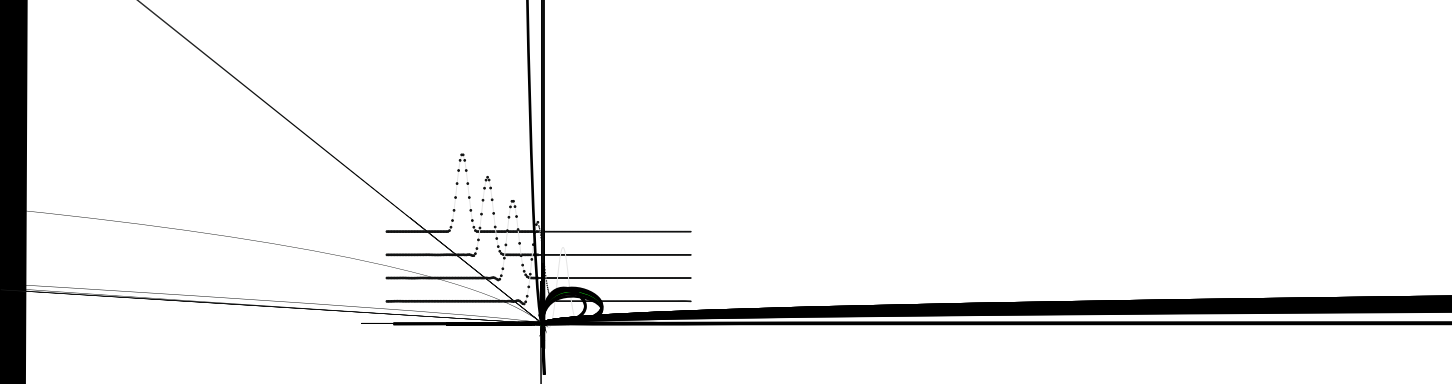

We investigate the conditions for zero reflectance in a coupled, atomistic–continuum system by taking the continuum limit of one of the atomistic mass–spring subsystems in the heterogeneous atomistic model. Specifically, for system 2, we hold the domain length, the coarse-grained mass density ρ_2 and the long-wavelength wave speed c_2^{lw} fixed, but take the infinitesimal limit of the equilibrium lattice parameter, $a_2 \to 0$. The conditions for reflectance-free coupling in the spring-coupling model (*cf.* Fig. 7(b)) remain valid under the limiting process, so (44a) gives the limiting value of the optimal coupling lattice parameter as $\lim_{a_2 \to 0} a_{\mathrm{c}}^* = a_1/2$. Thus, we obtain minimum-reflectance coupling between a continuum model with material parameters determined by (40), in which $S = 1$ and the atomistic parameters come from system 1 with $K_1^1 = \overline{F}_1/a_1$, if the coupling lattice parameter is taken as $a_{\mathrm{c}} = a_1/2$, or equivalently, the coupling spring constant is taken as $K_{\mathrm{c}} = 2\overline{F}_1/a_1$. On the other hand, we expect significant spurious reflections for short-wavelength pulses for any other choice of a_{c}, even if the bulk properties of the continuum and atomistic domains are perfectly matched.

As discussed in Section 1.1, several coupling methods minimize the reflected energy by damping reflected modes, paying special attention to modes that the discrete continuum model cannot resolve. However, we have shown that spurious reflections can arise, even when both domains are fully capable of resolving the given wave form. Therefore, coupling methods that rely on damping to suppress spurious reflections are expected to cause undesirable attenuation of the transmitted signal in these cases. The above analysis suggests that spurious reflections can be suppressed, without sustaining undesirable energy losses due to damping, by adjusting the relative positions of atoms and the AtC coupling interface. We next propose a new AtC coupling scheme that uses this concept to suppress spurious reflections while balancing momentum to within machine precision and energy to within the accuracy of the component continuum and atomistic models.

2.4. Coupled SDG–ADG model

This subsection presents the details of a new AtC coupling strategy that links the continuum SDG and atomistic ADG models. The coupling strategy can also be used to join SDG models with explicit atomistic models, albeit with some compromise in the balance properties. We assume that the material properties of the continuum match the effective properties of the atomistic model in the long-wavelength limit. We further assume that all coupling interfaces are material interfaces. That is, there is no flux of mass or atoms across the coupling interfaces, which are therefore taken to be vertical spacetime manifolds. This assumption precludes modeling certain processes, such as diffusion of solute species through a crystal and solidification. Generalizations to loosen this restriction appear to be possible, but are beyond the scope of the present study. The coupling strategy is intended for use in up to three spatial dimensions, and for the most part we describe it in a dimension-independent fashion. However,

some lower bound, on the order of the atomic spacing a, where the assumptions of continuum theory cease to be valid. In practice, discretization error in numerical continuum models often renders some modes above this physical limit unresolvable. Using an adaptive implementation of the SDG model and an $\mathcal{O}(N)$ causality-based solution scheme in the continuum zone, we are able to remove this limitation and resolve wavelengths all the way down to, and indeed below, the physical limit of continuum theory. Thus, the decomposition into resolvable and unresolvable modes can be based on physical, rather than numerical, criteria.

We focus on maintaining balance of momentum and energy within and between the component models and on suppressing spurious reflections at the coupling interface without introducing non-physical damping. This avoids draining energy from the unresolvable modes which should, in combination with thermal response in the continuum region, participate in the overall energy balance of the coupled system. We use non-overlapping continuum and atomistic zones to enable precise treatment of the balance laws, and show that the coupled SDG–ADG model enforces momentum balance to within machine precision over individual atoms and elements as well as clusters of the same. The coupling model is dissipative, with energy balance maintained to within the accuracy of the component methods. For now, we assume noise-free waveforms at zero temperature under the assumption that all modes are resolvable in the continuum. Methods for distinguishing and treating unresolvable modes are left for future research.

The main components of the coupling scheme are:

- A two-scale geometry model comprised of a macroscopic *continuum region* $\mathcal{D}$, a non-overlapping *atomistic zone* populated by a set of atoms identified by their reference positions, an atomic-scale coupling interface Γ that separates the continuum region from the atomistic zone, and *atomic-scale offsets* that define the position of Γ relative to $\partial\mathcal{D}$ at the scale of the atomic spacing.
- An independent *atomistic stress field* on Γ that models the momentum flux out of the atomistic zone.
- An *atomistic trace operator* that maps the discrete velocities of nearby atoms into a velocity field on Γ. The trace operator also generates an adjoint mapping of the atomistic stress field into discrete, work-equivalent forces acting on individual atoms.
- Weak enforcement of jump conditions on Γ that describe kinematic compatibility and momentum balance between the traces of the continuum fields and the constructed atomistic stress and velocity fields. The jump conditions are written with respect to target Riemann values, as in the SDG method, to preserve the common characteristic structure of the continuum equations and the long-wavelength limit of the atomistic equations.

Numerical results presented in Section 3 demonstrate that optimizing the atomic-scale geometry parameters in the atomistic trace operator reduces spurious reflections at the coupling interface to negligible levels without resorting to non-physical damping.

2.4.1. Geometry of the coupling model

Fig. 9 shows the spatial layout of the two-scale coupling model. The geometry of the continuum spatial domain $\mathcal{D}_{\mathbf{x}} \subset \mathbb{E}^d$ is defined using macroscopic continuum coordinates $\mathbf{x}$. As before, the continuum boundary $\partial\mathcal{D}_{\mathbf{x}}$ is equipped with a unit outward normal vector $\mathbf{n}$, but it now has the disjoint partition, $\partial\mathcal{D}_{\mathbf{x}} = \partial\mathcal{D}^{\mathbf{v}}_{\mathbf{x}} \cup \partial\mathcal{D}^{\sigma}_{\mathbf{x}} \cup \partial\mathcal{D}^{c}_{\mathbf{x}}$, in which $\partial\mathcal{D}^{c}_{\mathbf{x}}$ is the coupling part of the continuum boundary that interacts with the atomistic model.

The atomistic region, with interatomic length scale a, is populated by a set of atoms $G = \{\gamma\}$ identified by their spatial coordinates, $\mathbf{x}_\gamma$. Although the atomistic region has no naturally-defined boundary, our coupling strategy nonetheless requires a coupling interface, denoted by $\Gamma_{\mathbf{x}}$, on which to construct effective atomistic stress and velocity fields that enter the coupling jump conditions described below. The continuum coupling boundary $\partial\mathcal{D}^{c}_{\mathbf{x}}$ provides a natural basis for defining the cou-

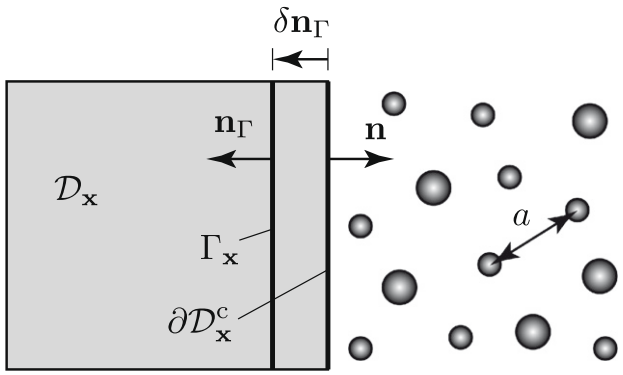

Fig. 9. Spatial geometry of atomistic–continuum coupling model for $d = 2$. An atomic-scale shift of the continuum coupling boundary $\partial\mathcal{D}^{c}_{\mathbf{x}}$ defines the spatial atomistic coupling boundary $\Gamma_{\mathbf{x}}$. A positive value of the shift δ indicates a normal perturbation outward from the atomistic zone.

pling boundary, but the macroscopic coordinates that define $\partial\mathcal{D}^{\mathrm{c}}_{\mathbf{x}}$ are not directly comparable to the atomic-scale coordinates, $\mathbf{x}_\gamma$. Therefore, we introduce an *atomic-scale shift*, $\delta : \partial\mathcal{D}^{\mathrm{c}}_{\mathbf{x}} \to \mathbb{R}$, where δ is comparable in magnitude to the atomic length scale a, and define $\Gamma_{\mathbf{x}} = \{\mathbf{x} = \mathbf{y} - \delta\mathbf{n}(\mathbf{y}) : \mathbf{y} \in \partial\mathcal{D}^{\mathrm{c}}_{\mathbf{x}}\}$ so that coordinates $\mathbf{x} \in \Gamma_{\mathbf{x}}$ are directly comparable to the atomic positions $\mathbf{x}_\gamma$. The coupling boundary $\Gamma_{\mathbf{x}}$ serves as a proxy boundary for the atomistic region, so we equip it with a normal vector that points outward from the atomistic zone: $\mathbf{n}_\Gamma(\mathbf{y} - \delta\mathbf{n}(\mathbf{y})) = -\mathbf{n}(\mathbf{y})\ \forall \mathbf{y} \in \partial\mathcal{D}^{\mathrm{c}}_{\mathbf{x}}$. The values of the shift δ are chosen to minimize spurious reflections, as described in Section 2.4.3. Fig. 9 shows δ as a uniform field, but in general it may vary with position on $\partial\mathcal{D}^{\mathrm{c}}_{\mathbf{x}}$.

We again have the spatial partition of $\mathcal{D}_{\mathbf{x}}$ into space elements: $\mathcal{P}_{\mathbf{x}} = \{\mathcal{Q}_m\}_{m=1}^{M}$, and the temporal partition of the analysis interval $\mathcal{I}$ into time steps: $\mathcal{P}_{\mathcal{I}} = \{\mathcal{I}_n\}_{n=1}^{N}$. We generate the spacetime domain and spacetime mesh as the Cartesian products, $\mathcal{D} = \mathcal{D}_{\mathbf{x}} \times \mathcal{I},\ \mathcal{P} = \mathcal{P}_{\mathbf{x}} \times \mathcal{P}_{\mathcal{I}},\ \Gamma = \Gamma_{\mathbf{x}} \times \mathcal{I}$, and $\Gamma_n = \Gamma_{\mathbf{x}} \times \mathcal{I}_n$. While the time steps determined by $\mathcal{P}_{\mathcal{I}}$ must be taken as uniform across all of the atoms, an unstructured partition of $\mathcal{D}$ can be useful, as exemplified in Section 3.3. The shift of the coupling boundary to Γ is only used in the trace operator for the atomistic velocity field (*cf.* Section 2.4.2), where Γ must be described in consistent atomic-scale coordinates. Integrals that weakly enforce the coupling jump conditions, on the other hand, are macroscopic in character, so for these we write

$$\int_{\Gamma_n} f(\mathbf{x}, t; \mathbf{n}_\Gamma)\, d\Sigma = \int_{\partial\mathcal{D}^{\mathrm{c}}_n} f(\mathbf{x}, t; -\mathbf{n}) d\Sigma. \tag{45}$$

2.4.2. Atomistic stress field and velocity trace operator

This subsection describes the construction of the effective atomistic stress and velocity fields on the atomistic coupling boundary that enter the coupling jump conditions described in Section 2.4.3. The atomistic stress, a second-order tensor field on Γ denoted as $\boldsymbol{\sigma}^{\mathrm{at}}$, represents the flux of linear momentum into the atomistic region due to its interaction with the continuum region.[5] It appears as an independent solution field in the SDG–ADG coupling model, with a piecewise-continuous polynomial structure on Γ. A momentum- and energy-conserving mapping from $\boldsymbol{\sigma}^{\mathrm{at}}$ to discrete forces acting on individual atoms is defined below.

The atomistic velocity field on Γ is a dependent field written as

$$\mathbf{v}^{\mathrm{at}}(\mathbf{x}, t) = \mathrm{tr}^{\mathrm{at}}\big(\mathbf{x}; \{\mathbf{v}_\gamma(t)\}, \delta\big) \tag{46}$$

in which $\mathrm{tr}^{\mathrm{at}}$ is the *atomistic trace operator for velocity*, and $\{\mathbf{v}_\gamma(t)\}$ denotes the velocities of the atoms in G at time t. The primary function of the trace operator is to construct a homogenized continuum velocity field on Γ as a function of the discrete atomic velocities. This allows us to write jump conditions, similar to those in the SDG formulation, that couple the continuum and atomistic response. A secondary purpose is to provide a spatial filter that removes high-frequency modes associated with continuum thermal response from the atomistic velocity field on Γ. Temporal filtering can also be introduced by substituting time-averaged quantities for the unfiltered atomic velocities, $\{\mathbf{v}_\gamma(t)\}$. However, we do not pursue either form of filtering here, since we are concerned with systems at zero temperature where our goal is to pass the entire signal across the coupling interface.

The construction of the atomistic trace operator involves a smooth and continuous atomic-scale fit, $\tilde{\mathbf{v}}(\mathbf{x}, t)$, to the velocities of atoms near the coupling interface:

$$\tilde{\mathbf{v}}(\mathbf{x}, t) = \sum_{\gamma \in G_\Gamma} h_\gamma(\mathbf{x}) \mathbf{v}_\gamma(t), \tag{47}$$

where we require $\sum_{\gamma \in G_\Gamma} h_\gamma(\mathbf{x}) = 1\ \forall \mathbf{x}$, and in which $G_\Gamma \subset G$ is a specified collection of atoms in the vicinity of the coupling interface, typically those within the cut-off radius of the interatomic potential. Extrapolation and restriction to the coupling interface Γ completes the definition of the trace operator:

$$\mathrm{tr}^{\mathrm{at}}\big(\mathbf{x}; \{\mathbf{v}_\gamma(t)\}, \delta\big) := \tilde{\mathbf{v}}(\mathbf{x}, t)|_{\Gamma(\delta)}. \tag{48}$$

The position of Γ relative to the atoms varies with the shift parameter, so varying δ affects the action of the atomistic trace operator.

The coupling stress $\boldsymbol{\sigma}^{\mathrm{at}}$ describes the momentum flux due to the action of the continuum region an the atomistic zone. We need a mapping of $\boldsymbol{\sigma}^{\mathrm{at}}$ into time-dependent forces acting on individual atoms, $\gamma \in G_\Gamma$, to close the model. Let $\mathbf{F}_\gamma$, the net force acting on atom $\gamma \in G$, be decomposed according to

$$\mathbf{F}_\gamma = \mathbf{F}^{\mathrm{int}}_\gamma + \mathbf{F}^{\boldsymbol{\sigma}}_\gamma + \mathbf{F}^{\mathrm{ext}}_\gamma \tag{49}$$

in which $\mathbf{F}^{\mathrm{int}}_\gamma$ is the force acting on γ due to interatomic forces within the atomistic zone, as defined in (22), $\mathbf{F}^{\boldsymbol{\sigma}}_\gamma$ is the coupling force acting on γ derived from $\boldsymbol{\sigma}^{\mathrm{at}}$ and due to interactions with the continuum region, and $\mathbf{F}_\gamma{}^{\mathrm{ext}}$ is the resultant of any non-coupling external forces acting on γ. Clearly, $\mathbf{F}^{\boldsymbol{\sigma}}_\gamma = \mathbf{0}$ for atoms $\gamma \notin G_\Gamma$. We require the coupling forces to satisfy the work-equivalence constraint,

[5] If the coupling boundary is a vertical (material) interface, as we assume in this study, then we only require the normal component of the atomistic stress field, *i.e.,* a traction field on Γ. However, for generality, we continue to use "stress" and $\boldsymbol{\sigma}^{\mathrm{at}}$ to refer to this field.

$$\sum_{\gamma\in G_\Gamma}\int_{\mathcal{I}_n}\hat{\mathbf{v}}_\gamma\cdot\mathbf{F}^\sigma_\gamma\,dt=\int_{\Gamma_n}\hat{\mathbf{v}}^{\mathrm{at}}\cdot\boldsymbol{\sigma}^{\mathrm{at}}(\mathbf{n}_\Gamma)\,d\Sigma\quad\forall\{\hat{\mathbf{v}}_\gamma\}\in\widetilde{\mathcal{V}}\tag{50}$$

in which $\hat{\mathbf{v}}^{\mathrm{at}}(\mathbf{x},t;\delta)=\sum_{\gamma\in G_\Gamma}h_\gamma(\mathbf{x})\hat{\mathbf{v}}_\gamma(t)|_{\Gamma(\delta)}$. After some manipulation, we find that (50) implies the mapping,

$$\mathbf{F}^\sigma_\gamma(t)=\begin{cases}\int_{\Gamma_{\mathbf{x}}}h_\gamma(\mathbf{x})\boldsymbol{\sigma}^{\mathrm{at}}(\mathbf{n}_\Gamma)|_t\,d\mathbf{x} & \text{for }\gamma\in G_\Gamma,\\ 0 & \text{otherwise}\end{cases}\tag{51}$$

and that the coupling stresses $\boldsymbol{\sigma}^{\mathrm{at}}$ and the coupling forces $\mathbf{F}^\sigma_\gamma$ deliver the same momentum input to the atomistic zone. Taking $\{\hat{\mathbf{v}}_\gamma\}=\{\mathbf{v}_\gamma\}$ in (50), we find that they also deliver the same energy input. The stresses $\boldsymbol{\sigma}^{\mathrm{at}}$ can be interpreted as Lagrange multipliers that enforce a holonomic kinematic compatibility constraint on Γ [49].

Since the examples reported in Section 3 involve problems in $1d\times$ time with smooth initial data (zero temperature), we use one-dimensional Lagrange interpolation polynomials for $h_\gamma(x)$. Thus, we write the n-atom trace as

$$v^{\mathrm{at}}(t)=\sum_{\gamma\in G_\Gamma}\tilde{h}_\gamma v_\gamma(t)\tag{52}$$

in which $\tilde{h}_\gamma:=h_\gamma(\mathbf{x})|_{\Gamma_{\mathbf{x}}}$ are restrictions of Lagrange interpolation polynomials in x of degree $n-1$, and where the set G_Γ contains the n atoms nearest to Γ (*cf.* Fig. 10 in which Γ is at $x=0$). Note that Lagrange interpolation polynomials satisfy $\sum_{\gamma\in G_\Gamma}\tilde{h}_\gamma(x)=1\ \forall x$.

2.4.3. Coupled SDG–ADG formulation

The SDG–ADG coupling strategy involves *coupling Riemann values*, $\boldsymbol{\sigma}^{\mathrm{Rc}}$ and $\mathbf{v}^{\mathrm{Rc}}$, that depend on the traces of the continuum stress and velocity fields on $\partial\mathcal{D}^{\mathrm{c}}$ and on $\boldsymbol{\sigma}^{\mathrm{at}}$ and $\mathbf{v}^{\mathrm{at}}$ on Γ (*cf.* (45)). The coupling Riemann values for $d=1$ are

$$v^{\mathrm{Rc}}=\frac{1}{2}\left[(v+v^{\mathrm{at}})+\frac{c}{C}(\sigma^{\mathrm{at}}-\sigma)n_x\right],\tag{53a}$$

$$\sigma^{\mathrm{Rc}}=\frac{1}{2}\left[(\sigma+\sigma^{\mathrm{at}})+\frac{C}{c}(v^{\mathrm{at}}-v)n_x\right]\tag{53b}$$

in which undecorated values v, σ and n_x refer to continuum traces and the spatial normal component on $\partial\mathcal{D}^{\mathrm{c}}$. The parameters c and C are impedance matched according to (40). Our numerical experience, as reported in the examples in Section 3, suggests that Riemann values based on the long-wavelength limit of the atomistic model provide an adequate basis for coupling continuum and atomistic response to general, complex wave forms.

We complete the coupling model by modifying the definition of the target values (9) in the SDG formulation to include the coupling Riemann values on $\partial\mathcal{D}^{\mathrm{c}}$, and by including the coupling forces $\mathbf{F}^\sigma_\gamma$ in an augmented ADG model that includes the auxiliary field $\boldsymbol{\sigma}^{\mathrm{at}}$ and jump conditions for $\mathbf{v}^{\mathrm{at}}$ and $\boldsymbol{\sigma}^{\mathrm{at}}$ on Γ. Specifically, we replace (9b) with

$$\mathbf{v}^*,\boldsymbol{\sigma}^*=\begin{cases}\mathbf{v}^+,\boldsymbol{\sigma}^+ & \text{on }\partial\mathcal{Q}^{\mathrm{ci}}\setminus\partial\mathcal{D}^{\mathrm{ci}},\\ \bar{\mathbf{v}},\boldsymbol{\sigma} & \text{on }\partial\mathcal{Q}\cap\partial\mathcal{D}^{\mathbf{v}},\\ \mathbf{v},\bar{\boldsymbol{\sigma}} & \text{on }\partial\mathcal{Q}\cap\partial\mathcal{D}^{\boldsymbol{\sigma}},\\ \mathbf{v}^{\mathrm{Rc}},\boldsymbol{\sigma}^{\mathrm{Rc}} & \text{on }\partial\mathcal{Q}\cap\partial\mathcal{D}^{\mathrm{c}},\\ \mathbf{v}^{\mathrm{R}},\boldsymbol{\sigma}^{\mathrm{R}} & \text{on }\partial\mathcal{Q}^{\mathrm{nc}}\setminus\partial\mathcal{D}^{\mathrm{nc}}\end{cases}\tag{54}$$

use (49) and (51) in (25b), and append to the ADG formulation jump conditions for kinematic compatibility and momentum balance on Γ:

$$(\mathbf{v}^*-\mathbf{v}^{\mathrm{at}})(1-|\mathbf{n}_\Gamma\cdot\mathbf{e}_t|)=0,\tag{55a}$$

$$(\boldsymbol{\sigma}^*-\boldsymbol{\sigma}^{\mathrm{at}})(\mathbf{n}_\Gamma)|_\Gamma=0\tag{55b}$$

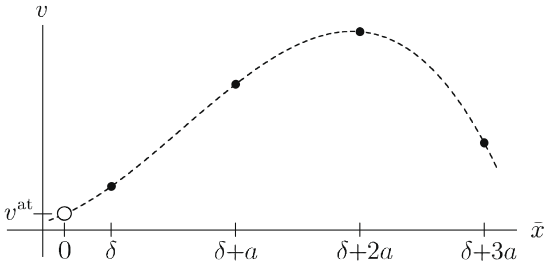

Fig. 10. Representation of the four-atom trace in the reference configuration with the coupling boundary Γ at $\bar{x}=0$. A Lagrange polynomial is fit to the velocities at time t of the four atoms nearest Γ. The trace v^{at} is the restriction of the velocity interpolant to Γ. Because the atomic positions are fixed relative to one another, δ is the only free parameter available to minimize reflections; it alters the trace by moving Γ relative to the atoms.

in which $\mathbf{v}^* = \mathbf{v}^{\mathrm{Rc}}$ and $\boldsymbol{\sigma}^* = \boldsymbol{\sigma}^{\mathrm{Rc}}$. Then, combining the weak forms for the modified SDG and ADG methods, we obtain the weak formulation of the coupled problem. Alternative methods, based on the weighted residual forms and/or the 1-field continuum formulation, are also possible.

Problem 9 (*Coupled SDG–ADG*). Find the continuum solution $(\mathbf{u},\mathbf{v},\mathbf{E}) \in \mathcal{U}\times\mathcal{V}\times\mathcal{E}$, the atomic trajectories $\{(\mathbf{u}_\gamma,\mathbf{v}_\gamma)\} \in \widetilde{\mathcal{U}}\times\widetilde{\mathcal{V}}$, and the boundary traction field $\boldsymbol{\sigma}^{\mathrm{at}} \in \widetilde{\mathcal{S}}$ such that for each 3 $\mathcal{Q} \in \mathcal{P}$

$$\int_{\mathcal{Q}} \Big\{-\nabla\hat{\mathbf{v}}(\boldsymbol{\sigma}) + \dot{\hat{\mathbf{v}}}\cdot\mathbf{p} + \hat{\mathbf{v}}\cdot\rho\mathbf{b} - \mathbf{v}\cdot(\nabla\cdot\hat{\boldsymbol{\sigma}}) + \mathbf{E}(\dot{\hat{\boldsymbol{\sigma}}}) - k_{\mathcal{Q}}(\mathbf{u}\cdot\dot{\hat{\mathbf{u}}} + \mathbf{v}\cdot\hat{\mathbf{u}})\Big\}d\Omega + \int_{\partial\mathcal{Q}^{\mathrm{nc}}} [\hat{\mathbf{v}}\cdot\boldsymbol{\sigma}^*(\mathbf{n}) + \mathbf{v}^*\cdot\hat{\boldsymbol{\sigma}}(\mathbf{n})]\, d\Sigma$$
$$+ \int_{\partial\mathcal{Q}^{\mathrm{ci}}} (\hat{\mathbf{v}}\cdot\mathbf{p}^* + \mathbf{E}^*(\hat{\boldsymbol{\sigma}}) - k_{\mathcal{Q}}\mathbf{u}^*\cdot\hat{\mathbf{u}})d\Sigma + \int_{\partial\mathcal{Q}^{\mathrm{co}}} (-\hat{\mathbf{v}}\cdot\mathbf{p} - \mathbf{E}(\hat{\boldsymbol{\sigma}}) + k_{\mathcal{Q}}\mathbf{u}\cdot\hat{\mathbf{u}})d\Sigma = 0 \;\forall(\hat{\mathbf{u}},\hat{\mathbf{v}},\widehat{\mathbf{E}}) \in \mathcal{U}\times\mathcal{V}\times\mathcal{E} \tag{56}$$

and for each time interval $\mathcal{I}_n$

$$\sum_\gamma \int_{\mathcal{I}_n} \Big[\hat{\mathbf{v}}_\gamma\cdot\mathbf{F}_\gamma + \dot{\hat{\mathbf{v}}}_\gamma\cdot\mathbf{p}_\gamma - \mathbf{v}_\gamma\cdot k_\gamma\hat{\mathbf{u}}_\gamma - \mathbf{u}_\gamma\cdot k_\gamma\dot{\hat{\mathbf{u}}}_\gamma\Big]dt + \sum_\gamma \left[\left(\hat{\mathbf{v}}_\gamma^n\cdot\mathbf{p}_\gamma^{n-1} - \mathbf{u}_\gamma^{n-1}\cdot k_\gamma\hat{\mathbf{u}}_\gamma^n\right)\Big|_{t=t_{n-1}} + \left(-\hat{\mathbf{v}}_\gamma^n\cdot\mathbf{p}_\gamma^n + \mathbf{u}_\gamma^n\cdot k_\gamma\hat{\mathbf{u}}_\gamma^n\right)\Big|_{t=t_n}\right]$$
$$+ \int_{\Gamma_n} \{(\mathbf{v}^* - \mathbf{v}^{\mathrm{at}})\cdot\hat{\boldsymbol{\sigma}}^{\mathrm{at}}(\mathbf{n}_\Gamma) + \hat{\mathbf{v}}^{\mathrm{at}}\cdot(\boldsymbol{\sigma}^* - \boldsymbol{\sigma}^{\mathrm{at}})(\mathbf{n}_\Gamma)\}dt = 0 \;\forall\big(\{(\hat{\mathbf{u}}_\gamma,\hat{\mathbf{v}}_\gamma)\},\hat{\boldsymbol{\sigma}}^{\mathrm{at}}\big) \in \Big(\widetilde{\mathcal{U}}\times\widetilde{\mathcal{V}}\Big)\times\widetilde{\mathcal{S}} \tag{57}$$

in which $\hat{\boldsymbol{\sigma}} = C\widehat{\mathbf{E}}$.

Next, we consider the momentum and energy balance properties of the coupled SDG–ADG formulation. First, we note that the proofs that linear and angular momentum balance with respect to $\boldsymbol{\sigma}^*$ and $\mathbf{p}^*$ to within machine precision on a per-element basis in the uncoupled SDG method are unaffected by the modified target values in (54); *cf.* Section 2.1.3. The proofs of per-atom balance of momentum with respect to the atomic forces $\{\mathbf{F}_\gamma\}$ in the uncoupled ADG method are similarly unaffected by the redefinition of the net atomic forces in (49) and (51); *cf.* Section 2.2.1. We have also shown that the coupling forces $\{\mathbf{F}_\gamma^\sigma\}$ are momentum equivalent to the coupling stresses $\boldsymbol{\sigma}^{\mathrm{at}}$ on Γ. It only remains to show that $\boldsymbol{\sigma}^{\mathrm{Rc}}$ and the coupling stresses $\boldsymbol{\sigma}^{\mathrm{at}}$ are momentum equivalent. Consideration of all atomic velocity weightings in (57) of the form,

$$\hat{\mathbf{v}}_\gamma = \begin{cases} \mathbf{c} + \mathbf{S}(\mathbf{x}_\gamma) & \text{for } \gamma \in G_\Gamma, \\ \mathbf{0} & \text{otherwise}, \end{cases} \tag{58}$$

with $\{\hat{\mathbf{u}}_\gamma\} = \{\mathbf{0}\}$ and $\hat{\boldsymbol{\sigma}}^{\mathrm{at}} = \mathbf{0}$, leads to the desired result for an n-atom trace, if $n > 1$. Thus, the coupled SDG–ADG solution balances linear and angular momentum overall, as well as on a per-element and per-atom basis.

The energy error across the coupling interface is due solely to discrepancies in the jump terms involving $(\mathbf{v}^{\mathrm{Rc}},\boldsymbol{\sigma}^{\mathrm{Rc}})$ and either $(\mathbf{v}^{\mathrm{at}},\boldsymbol{\sigma}^{\mathrm{at}})$ or the continuum traces $(\mathbf{v},\boldsymbol{\sigma})$. Since $(\mathbf{v}^{\mathrm{at}},\boldsymbol{\sigma}^{\mathrm{at}})$ on Γ are energy-equivalent to $(\{\mathbf{v}_\gamma\},\{\mathbf{F}_\gamma^\sigma\})$ on the atoms, we conclude that the coupling part of the SDG–ADG method has the same stability and energy dissipation properties as the component SDG and ADG methods.

2.4.4. Determining the atomic-scale shift

Specific values for the atomic-scale shift δ must be determined in the SDG–ADG coupling method. Our studies show that large spurious reflections generally result when the macroscopic continuum geometry and the microscopic atomic positions are specified independently. However, spurious reflections are virtually eliminated when the relative positions of the coupling interface and the atoms are properly adjusted at the scale of the atomic spacing. The shift parameter δ in the atomistic trace operator closes the atomic-scale geometry model, which would be ambiguous and incomplete without it.

Ideally, δ would be determined by a systematic procedure as an automatic part of the solution process. To date, however, we rely on a heuristic procedure for determining δ as the value that minimizes spurious reflections in a test problem involving a pulse crossing an atomistic–continuum interface. The details of the calibration procedure are described, in the context of concrete examples, as part of the numerical results presented in the following section. Typically, the optimal value of δ is sharply defined and easy to identify. Although the optimal value does depend on the number of atoms in the trace operator, it is relatively insensitive to the shape of the pulse and the direction of wave travel (from atomistic to continuum or from continuum to atomistic). Once a given model is calibrated, δ can be set for all future runs with that model. Prospects for a more systematic method for determining δ are discussed in Section 4.

2.4.5. Iterative solution scheme for SDG–ADG method

We extend the overlapping block-Gauss–Seidel iterative solution scheme to address the coupled SDG–ADG system by introducing a special *coupling bin* $b_C \in \{b_i\}_{i=1}^M$ such that $G \subset b_C$. That is, we place all of the atoms involved in the atomistic trace operator in a common bin. In the iterative solution process, solutions for the atoms in b_C are updated simultaneously with the continuum solution adjacent to the coupling interface. When a causal mesh is used to discretize the continuum, the coupled update is localized to include only one layer of continuum elements along the coupling interface. The remainder of the bins $(\{b_i\}_{i=1}^M \setminus b_C)$ are updated in the same manner as described in Section 2.2.2. For each time step, iterative updates are performed until the solution to Problem 9 converges to within a specified tolerance.

3. Numerical results

This section presents numerical results that demonstrate the effectiveness and the convergence properties of the SDG–ADG coupling scheme as well as the details of a heuristic method for optimizing the shift parameter δ in the atomistic trace operator. Unless otherwise noted, we use first-nearest-neighbor linear spring interactions in the atomistic model (second-nearest-neighbor springs and the nonlinear Lennard-Jones potential are treated in some examples). All units are given in terms of the atomistic quantities, a, K, $\overline{F}$ and $c = c^{\mathrm{lw}}$. In all cases, we match the mass densities of the component domains and the wave speeds in the long-wavelength limit, so that the domains have the same impedance. One half of the problem domain is comprised of continuum elements, the other half of atoms, with periodic boundary conditions and independent coupling fields v^{at} and σ^{at} at each AtC boundary. Thus, there is one coupling boundary at the center of the domain and one that implements the periodic boundary condition. We again specify initial conditions for displacement and velocity corresponding to a traveling pulse, as in (31), and we define the pulse width w as the support of the initial displacement data; that is, $w = \pi/k$. Adjusting the pulse width reveals a range of behaviors. A narrow pulse generates larger relative motion between adjacent atoms and stronger dispersive effects, while a broad pulse reduces the relative motion and the dispersion.

In this section, we fix the level of h- and p-enrichment in the SDG and ADG models in order to focus on the influence of various physical model parameters on spurious reflections at the coupling interface. We choose element sizes and polynomial orders to ensure sufficiently well-resolved solutions that discretization error has negligible impact on the coupling model. Instead, we focus on the influence of parameters such as direction of wave motion, the pulse width, the value of the shift parameter δ and the number of atoms n participating in the trace operator for the atomistic velocity.

In Sections 3.1 and 3.2, we use the unmodified ADG model and the 1-field SDG method on a Cartesian-product spacetime domain in $1d \times$ time, as described in Section 2.1.1, with a constant time step $\Delta t = 0.5a/c$. The polynomial bases for the rectangular spacetime elements in the continuum domain are tensor products of polynomial bases in x and t. The ADG displacements (velocities) are interpolated with cubic (quadratic) polynomials. The SDG displacement field is third-order in time and fourth-order in space. The coupling boundary traction is also cubic in time. As we are interested in obtaining the best possible coupling between the atomistic and continuum methods, and in view of the adaptive SDG method's ability to resolve very small length scales, we set the spatial density of atomistic and continuum d.o.f. to be the same, giving the spatial diameter $h = 5a$ for the continuum elements.

3.1. Performance of the SDG–ADG coupling method

We consider a coupled SDG–ADG model with 20 continuum elements and 100 atoms and with pulses traveling in either direction. The initial mechanical pulse starts at the center of the atomistic (continuum) domain and propagates to the center of the continuum (atomistic) domain. We denote a pulse passing in this direction as A–C (C–A). For comparison purposes, we also consider pure continuum and pure atomistic models for the same domain. Fig. 11 shows sample displacement profiles for pure ADG, pure SDG and coupled SDG–ADG simulations based on 1- and 3-atom trace operators. The most prominent feature in these plots is the spurious reflection off the coupling interface in the ADG–SDG model with the trivial 1-atom trace, where v^{at} is simply the velocity of the atom nearest the coupling interface. Since the distinct systems have perfectly matched impedances, the spurious reflection is solely due to the physically inconsistent coupling. Comparing Fig. 11(c) and (d), we see that the presence of the spurious reflection is largely independent of the direction of wave travel, whether A–C or C–A. This response is qualitatively similar to that of a coupled mass–spring system with an overly stiff coupling spring, as described in Section 2.3. The optimized 3-atom trace, on the other hand, largely attenuates the reflection. Another notable feature in these plots is the accumulation of dispersion effects as the pulse traverses the atomistic domain. Although an existing dispersion tail might follow the pulse as it crosses the continuum zone, dispersion effects will not continue to accumulate. The contrast between the dispersive atomistic model and the non-dispersive continuum model is most pronounced in Fig. 11(a) and (b), in which the wave form of the atomistic profile gradually changes due to dispersion.

System energy and momentum are conserved for SDG–ADG coupling to within the limits of the component methods. In the case of momentum, the normalized, per time-step change in system momentum, remains at or below 10^{-13}. This is within machine precision, given that the condition numbers of the individual atom and element submatrices are estimated to be $\mathcal{O}(10^3) - \mathcal{O}(10^4)$. We define the *dissipative energy error* as $\epsilon_{\mathrm{e}}(t) = [E_{\mathrm{i}} - E(t)]/E_{\mathrm{i}}$ in which E_{i} and $E(t)$ are the initial and current total energies. The total energy is dissipative, as expected, since each component method is likewise dissipative. However, the dissipation is quite small, as seen in Fig. 12, which displays time histories of the dissipative energy error for three separate calculations: A–C, C–A and a pure continuum calculation. The dissipation is larger in the continuum zone for the current model parameters, so we can track the evolution of the pulse as it moves through the domain in the time histories. The smooth transition in the energy as the pulse passes between the atomistic and continuum zones reflects the mathematical compatibility of the SDG and ADG models.

3.2. Optimization of the parameter δ to minimize spurious reflections

The shift parameter δ in the n-atom trace operator can be optimized to minimize spurious reflections, in some suitable measure, after the pulse entirely exits one domain and enters the other. Once again, the computation begins with the pulse

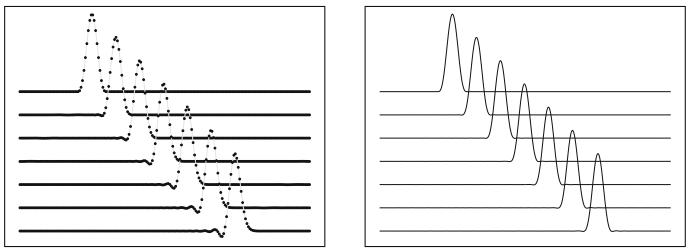

(a) Pure atomistic (b) Pure continuum

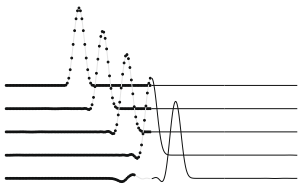

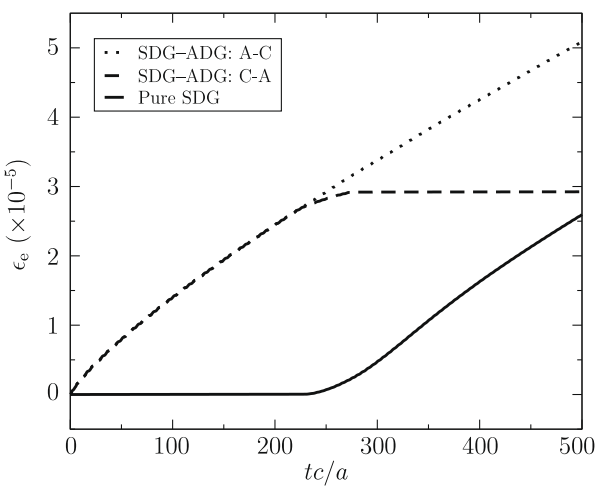

Fig. 12. Time histories of normalized dissipative energy error as pulse travels A–C or C–A with the SDG–ADG model. Results from a pure SDG continuum simulation are also shown for comparison. Here we use 501 atoms, 100 continuum elements, pulse width $w = 50a$ and a 3-atom trace operator.

centered in the atomistic or continuum region; it ends when the pulse reaches the center of the other region. Thus for an A–C calculation, we seek the value of δ that minimizes either the total energy, the l^2-norm of the displacement error or the l^2-norm of the velocity error in the atomistic domain at the end of the calculation, all of which should vanish for perfectly coupled methods:

$$\min_{\delta} \left[\frac{1}{2}\left(\sum_{\gamma} m_\gamma v_\gamma^2\right) + V(\{u_\gamma\})\right], \tag{59a}$$

$$\min_{\delta} \sqrt{\sum_{\gamma} u_\gamma^2}, \tag{59b}$$

$$\min_{\delta} \sqrt{\sum_{\gamma} v_\gamma^2}. \tag{59c}$$

Similarly for the C–A case, we can find the δ that minimizes any of

$$\min_{\delta} \left[\frac{1}{2}\int_{\partial D^{\mathrm{co}}} (vp + \sigma E)\, dx\right], \tag{60a}$$

$$\min_{\delta} \sqrt{\int_{\partial D^{\mathrm{co}}} u^2\, dx}, \tag{60b}$$

$$\min_{\delta} \sqrt{\int_{\partial D^{\mathrm{co}}} v^2\, dx}, \tag{60c}$$

the total energy, the L^2-norm of the displacement error or the L^2-norm of the velocity error in the continuum domain at the end of the calculation.

As shown in Fig. 13, the optimal value of δ depends on the number of atoms in the trace operator. All of the proposed error measures work well and yield the same optimal value of δ to within an order of accuracy of 10^{-5} in the wide-pulse (long-wavelength) limit. Notably, in this limit, the results are the same regardless of whether the pulse travels from the atomistic to the continuum domain (A–C) or from the continuum to the atomistic domain (C–A). Indeed, the results are remarkably consistent through most of the range of pulse widths, with significant deviations arising only when the pulse width approaches the lower-scale limit of continuum theory. Thus, for a given number of trace atoms, it appears that there is a well-defined optimal value for δ in the range of pulse widths where continuum modeling is meaningful. Therefore, it seems reasonable to pre-compute and store the optimal value of δ for reuse in a simulation, particularly if high-frequency continuum response is to be addressed by some other method. Fig. 14 shows reflectance as a function of δ for a pulse width of $50a$ and for various values of n, the number of atoms in the trace operator. The optimal value of δ that minimizes the reflected energy is sharply defined for each value of n. The minima appear smooth when viewed on a non-log plot at higher magnification. In the best case, reflectance is reduced by a factor greater than 10^4 relative to zero shift, and by still larger factors when compared to non-optimal choices of δ. The optimal values of δ are not affected by the direction of wave travel.

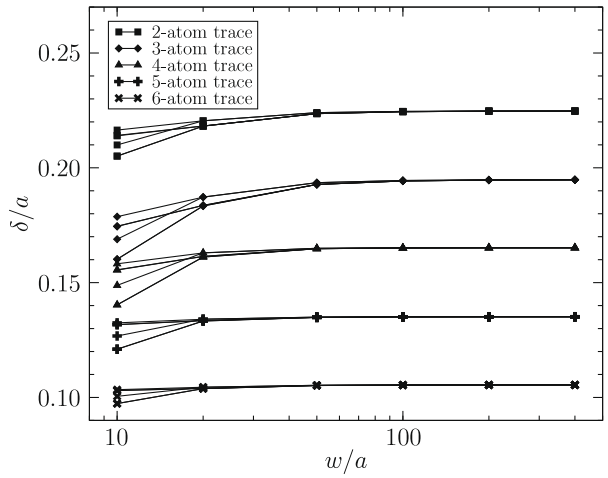

Fig. 13. Optimal values of shift parameter δ for SDG–ADG method with 501 atoms and 100 elements with varying pulse width w in initial data. For each pulse width and for each n-atom trace, the plot shows six different minimizations of reflections. For the pulse passing completely from the atomistic region to the continuum (continuum region to the atomistic), the plot shows results for minimization of energy as well as for minimizations of l^2-norms (L^2-norms) of the displacement and velocity in the atomistic (continuum) region. For specified w and n, the results for reflected energy and for l^2-norm (L^2-norm) of the velocity are indistinguishable at the scale of these plots.

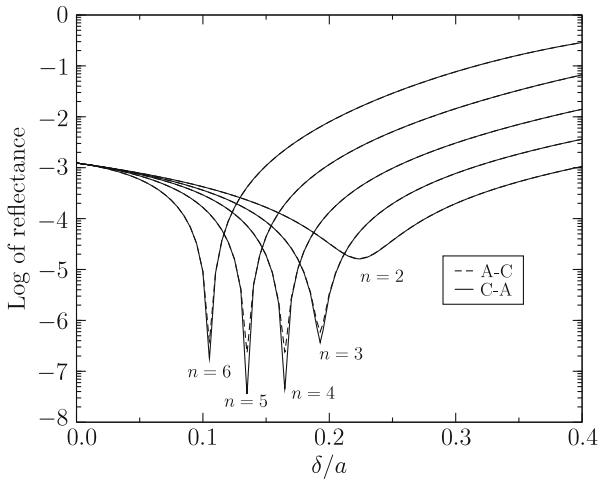

Fig. 14. Reflectance (fraction of energy reflected) for SDG–ADG system, with 501 atoms, 100 continuum elements and pulse width $w = 50a$, for varying shift parameter δ. The values of n indicate the number of atoms participating in the atomistic trace operator. The minima appear smooth in non-log plots of the data.

Fig. 15 shows the minimal reflected energy for different values of n and for various pulse widths, with the pulse traveling either A–C or C–A. In the narrow-pulse limit, a distinction arises between the A–C and C–A results that can be attributed to the effects of dispersion. The dispersion in the atomistic model is more pronounced for a narrow pulse, so a narrow pulse arriving at the coupling boundary from the atomistic domain has a more oscillatory structure than a similar pulse arriving from the continuum domain with zero dispersion.

The values for K^s given in Section 2.3.1 imply a length scale for multi-atom potentials and suggest that, given the relative sizes of the spring constants, treating first- and second-nearest-neighbor spring interactions is sufficient. Fig. 16 displays optimal values of δ in a second-nearest-neighbor model as a function of the spring-constant ratio, K^2/K^1, for $n = 2, 3, 4$. Fig. 17 shows the fraction of reflected energy at the optimal value of δ as a function of the spring-constant ratio for $n = 1, 2, 3, 4$. The most consistent results and the most effective suppression of spurious reflectance are for $n > 2$, where the smallest reflectance occurs in the vicinity of the first-nearest-neighbor system ($K^2/K^1 = 0$). Overall, the SDG–ADG model with optimized shift parameter and $n > 1$ robustly suppresses spurious reflections in second-nearest-neighbor systems, with reflectance below 10^{-5} throughout the range of this study.

3.3. SDG–ADG coupling with causal adaptive meshing in the continuum zone

Ideally, the continuum model should be able to handle mechanical signals all the way down to the atomic limit, without spurious energy losses due to non-physical damping. Some form of dynamic adaptive analysis and a scalable solution

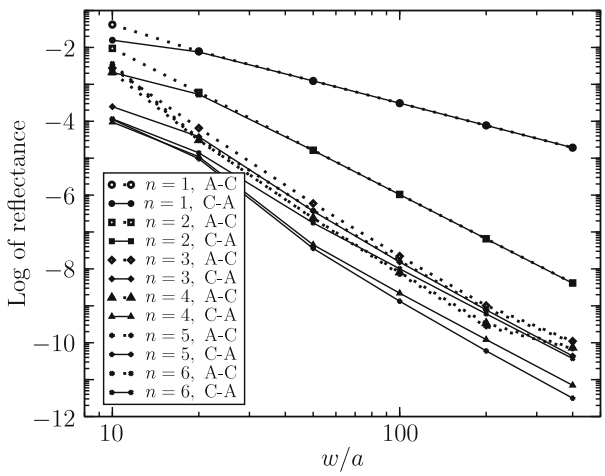

Fig. 15. Study of reflectance versus pulse width for models with 501 atoms, lattice parameter a, 100 continuum elements, optimized δ and n-atom trace operators, for $n = 1, 2, \ldots, 6$.

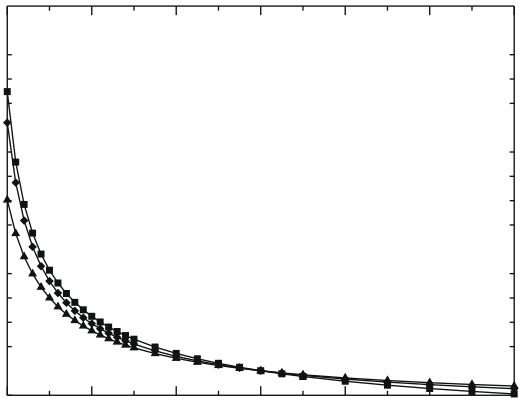

scheme are almost certainly needed to realize this objective. Here we demonstrate that the SDG–ADG AtC coupling method, when implemented with causal h-adaptive meshing in the continuum region, is able to meet this requirement. It is as yet unclear whether any other competing methods can do this. The adaptive meshing generates unstructured spacetime grids that include inclined boundaries. Therefore, our implementation uses the general formulation of the SDG method, as described in [24,26]. However, the AtC coupling boundary is always vertical in spacetime, so the coupling model described in this paper remains valid for the adaptive solution. Our results demonstrate that the adaptive SDG–ADG method resolves undamped atomic-scale signals in the continuum zone. It matches the atomistic and continuum time increments at the coupling interface but does not impose this constraint globally on the continuum discretization.

We use *Tent Pitcher* [26–29], an advancing-front mesh generation procedure, to construct simplicial spacetime meshes that satisfy a patch-wise *causality constraint*. A 'patch' is a small cluster of adjacent elements, and a patch boundary is a collection of oriented *facets* in spacetime. A facet is *causal* if the characteristics all flow in the same direction across it. Tent Pitcher meshes satisfy a causality constraint that requires all patch facets to be causal, except for those that lie on $\partial\mathcal{D}^{\mathrm{nc}}$. Any patch that satisfies this requirement is called a *causal patch*. No global time step restriction is applied; however, the causality constraint imposes a local condition similar to a CFL condition that limits the temporal duration of each patch. Inter-element boundaries within a patch are permitted to be non-causal, so the elements within a patch must be solved simultaneously. However, the solution within a patch depends only on solution data from earlier patches and on prescribed

initial and boundary data. This asymmetric dependency implies a partial ordering of patches that enables an $\mathcal{O}(N)$ patch-by-patch solution procedure wherein the mesh generation and solution processes are interleaved. Our adaptive meshing technique dynamically refines and coarsens the spacetime mesh based on a patch-level *a posteriori* error indicator that measures energy dissipation. The adaptive meshing simultaneously adjusts the temporal and spatial element diameters, while circumventing the need for global remeshing operations. Adaptive remeshing is accomplished through special spacetime element geometries that, unlike conventional remeshing procedures, preserve high-order accuracy by circumventing the need for solution projections during remeshing. In summary, the causal adaptive solution procedure dramatically improves efficiency by combining a scalable solution scheme with highly effective adaptive meshing.

Coupling along the boundary Γ for time interval $\mathcal{I}_n$ involves a causal patch of SDG elements with a vertical boundary on Γ that spans $\mathcal{I}_n$. The causal property uncouples this patch from all other unsolved SDG continuum elements, eliminating the need for large implicit solves that involve the entire continuum domain. Only the causal patches adjacent to Γ need be solved simultaneously with the atoms in the coupling bin b_C in the iterative solution scheme; *cf.* Section 2.4.5. In the adaptive SDG–ADG coupling model, either the continuum or the atomistic region can control the common time-step size on Γ, based on accuracy requirements or causality constraints. The pulse data in (31) contains arbitrarily small wave lengths, so allowing the continuum model to request time steps smaller than those required by the atomistics enables increased accuracy in the present example. However, strong refinement on Γ is only required when a signal is passing from the ADG zone into the SDG region.

Our first example demonstrates the use of *h*-adaptive causal meshing in the SDG–ADG method. We use the 3-field continuum SDG method, which is the mathematically consistent counterpart of the 2-field ADG method. The example uses the same linear spring interactions used in Fig. 11 and the initial conditions given by (31) with $B = a/400$. We use causal adaptive meshing in the continuum region. All solution fields, continuum and atomistic, are interpolated with cubic ($p = 3$) polynomials. The wave pulse is carried through the atomistic domain into the continuum and back into the atomistic domain through periodic boundary conditions. Fig. 18 shows the displacement solution at equally spaced time intervals of $\Delta t = 20a/c$ overlaid on the computational mesh. The default atomistic time-step size is reduced, as indicated by the grey band in the atomistic domain, only when the pulse is transmitted into the continuum region. In general, refinement is required only along the pulse trajectory where its high-frequency components must be resolved. No time-step reduction occurs when the pulse re-enters the atomistic domain because numerical dissipation within the continuum domain suppresses the pulse's highest frequency modes. However, this does not indicate that the effectiveness of our coupling model depends

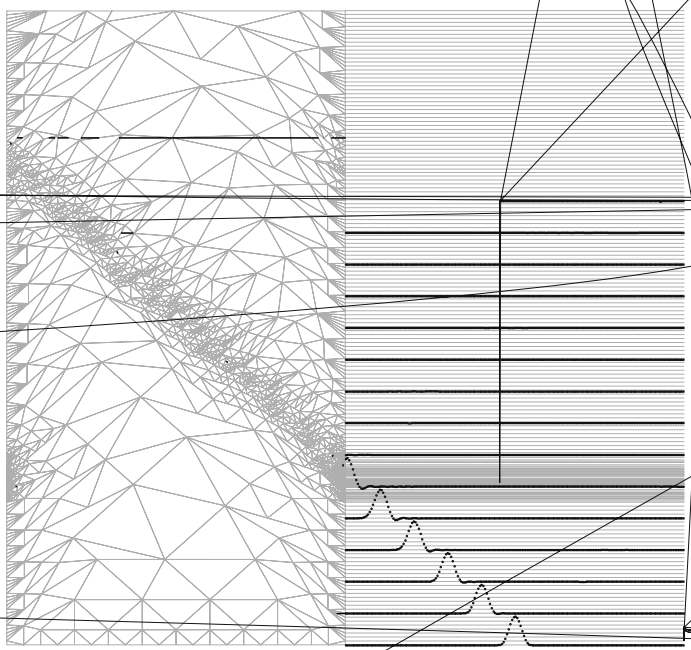

on numerical dissipation, since we have virtually reflection-free transmission when the pulse exits the atomistic zone with its high-frequency modes intact.

The 1-field SDG formulation yields results (not shown) that are qualitatively the same as Fig. 18. However, all independent fields in the 3-field SDG formulation converge as $\mathcal{O}(h^{p+1})$ for smooth solutions, compared with the reduced convergence rate of $\mathcal{O}(h^p)$ for stress and velocity in the 1-field formulation. To illustrate the impact of this difference on computational efficiency, we consider the same problem used in Fig. 18 and use the final dissipative energy error as a simple error metric: $\epsilon_e = (E_i - E_f)/E_i$ in which E_i and E_f are the total energies at the initial and final times of a simulation of one full cycle of the pulse through the periodic domain. We study the relative efficiency of the 1-field and 3-field SDG formulations by varying the element-wise error tolerance that controls the level of adaptive mesh refinement and plotting the error ϵ_e against the CPU time taken for each simulation.

Fig. 19, compares the computational efficiency of the 1-field and 3-field continuum models in the SDG–ADG coupling scheme with $p = 3$. The optimal convergence rates of the 3-field formulation deliver superior efficiency once the solution reaches the range of asymptotic convergence. The 1-field formulation with the same polynomial order requires more adaptive refinement to attain the same error level, making it more expensive for accurate simulations despite the fact that it has three times fewer d.o.f. per element.

In our final example, we demonstrate the ability of the SDG–ADG coupling method to handle asymmetrical, fully nonlinear potentials. We use a modified Lennard-Jones potential (30) with $r_c = 2.5a$, generating interactions out to two nearest-neighbors. C^1 continuity of the potential is enforced at the cut-off. The parameters R and ϵ have been impedance matched to the continuum domain. The nonlinear solution for each atomistic time step is computed with an iterative Newton–Raphson scheme that covers all the atoms plus the continuum SDG elements immediately adjacent to the coupling interfaces at either end of the atomistic domain.

In contrast to the nearest-neighbor linear-spring case, atoms within r_c of the coupling boundary are not in force equilibrium when the system is in its natural stress-free configuration unless we add *balancing forces* to atoms near the coupling boundary to represent force interactions with the 'missing' atoms that have been replaced by the continuum model. These are sometimes referred to as *ghost forces* [50]. These balancing forces are determined such that the entire coupled system is in equilibrium in the absence of dynamic disturbances. These forces are taken to be time-independent for the results reported here. We use the same initial data and geometry as in the previous example, and simulate one complete cycle of wave propagation through the periodic domain. Fig. 20 shows the displacement solution at equally spaced time steps of $\Delta t = 20a/c$ overlaid on the computational mesh. Due to the small amplitude pulse, this figure is directly comparable to the results in Fig. 18.

We optimized the shift parameter δ independently for different values of the adaptive tolerances and found that the level of mesh adaptivity has very little effect on the optimal value of δ. Generally, once δ is determined (given a reasonable initial adaptive tolerance), varying the error tolerance has very little impact on the reflectance. In fact, varying the adaptive tolerance by over two orders of magnitude while holding δ fixed yields changes in the normalized reflected energy of $\mathcal{O}(10^{-5})$, an amount that can be attributed to changes in the discrete solution space. Furthermore, the differences between these optimal values of δ and the comparable values for the structured meshes employed in Section 3.2 are of the order $\mathcal{O}(10^{-5}a)$. We conclude that adaptive spacetime meshing allows the continuum model to resolve fully the signal it receives from the atomistic model.

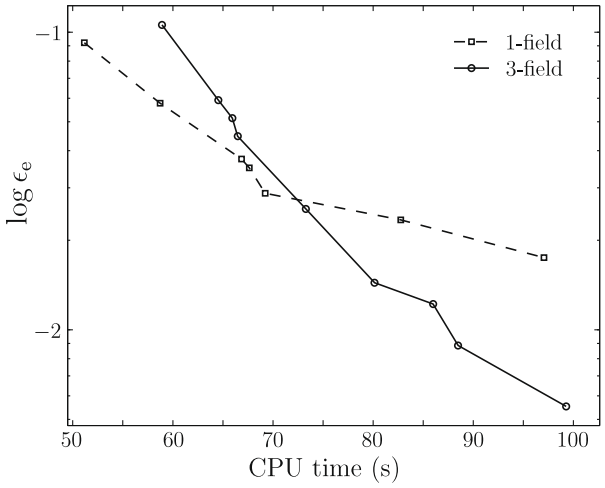

Fig. 19. Dissipative energy error ϵ_e versus CPU time for complete simulation up to $t = 400a/c$ for SDG–ADG model using 1- and 3-field SDG methods in the continuum zone. We use $p = 3$ for all fields. The adaptive tolerance for the per-element dissipation starts at 1×10^{-10} and increases by 2.5×10^{-10} for each new datum. CPU time reflects time to completion of a simulation for a fixed value of the adaptive error tolerance.

(black) is overlain on the computational mesh (grey) at $t = 100a/c$ represents the region where the continuum controls time stepping to resolve information as the pulse enters. The displacement solution

potential (30) is used, with initial conditions given by (31) with $\beta = a/400$. Horizontal grey lines in the atomistics depict time steps. The grey band centered

simulation uses 201 atoms with atomics spacing a (represented as dots on right-hand side) and a continuum length of $200a$. The modified Lennard-Jones

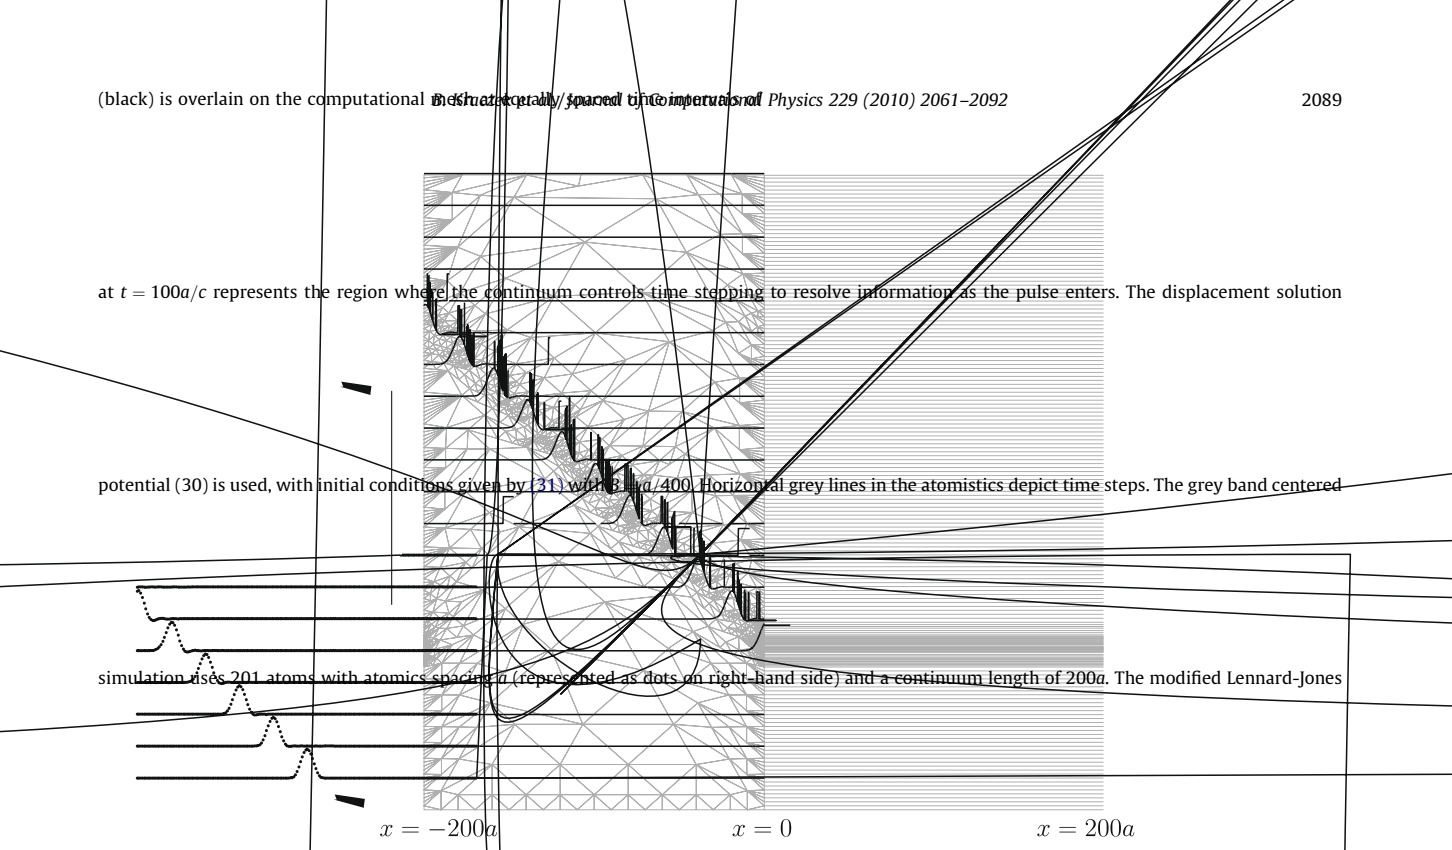

Fig. 20. Time history of SDG–ADG (3-field/2-field with $p = 3$) with causal adaptive meshing in the continuum domain (left-hand side) in $1d \times$ time. The

$\Delta t = 20a/c$.

4. Conclusions and future work

4.1. Summary of results and findings

We have presented a novel method for concurrent coupling of atomistic and continuum models of dynamic mechanical response in solids. Our approach combines the Spacetime Discontinuous Galerkin method for elastodynamics with its mathematically consistent discrete counterpart, the Atomistic Discontinuous Galerkin method. The SDG–ADG coupling strategy derives directly from considerations of momentum balance and kinematic compatibility. Optimization of an atomic-scale configuration parameter, a necessary ingredient of an unambiguous description of the atomic-scale geometry, dramatically reduces spurious reflections at the coupling interface without resorting to overlap regions or artificial damping. The model balances momentum, over individual elements and atoms and over the entire coupled system, to within the machine precision.

The coupled and component models are unconditionally stable. Although they are dissipative, the amount of numerical dissipation can be effectively controlled through adaptive solution procedures. They both offer the possibility of arbitrarily high-order implementations on compact computational stencils.

A modification to the ADG formulation, described in Section 2.2.1, enforces per-atom energy balance to within the accuracy of the numerical quadrature of the force integral. The ADG method is an example of an implicit integrator, and these are typically avoided in large-scale atomistic calculations due to the expense of solving a large coupled system of equations at every time step. However, we have shown that the iterative ADG solution procedure introduced in Section 2.2.2 is a scalable algorithm, since it solves the implicit ADG equations with linear computational complexity in the number of atoms. For the class of problems under consideration here, involving the propagation of mechanical pulses in solids, the iterative ADG scheme (with polynomial order greater than one) outperforms the explicit Velocity Verlet algorithm in terms of computational efficiency. Whether the iterative ADG method can be competitive with existing methods for more general MD applications is a question that requires, and merits, further investigation.

An *atomistic trace operator* defines an effective atomistic velocity on the coupling interface as a function of the velocities of nearby atoms. The trace operator depends on parameters that specify the position of the coupling interface relative to the atoms on the atomic scale. In addition, an auxiliary *atomistic boundary stress* is introduced that provides a bridge between continuum stresses and atomistic forces and a means to enforce momentum balance between the distinct models. These atomistic fields enter the coupling model through the weak enforcement of jump conditions written with respect to Riemann solutions.

When the SDG method is implemented with unstructured, causal meshing we obtain an $\mathcal{O}(N)$, patch-by-patch solution procedure. Spacetime mesh adaptation is handled locally without global remeshing operations and without error-inducing projections of the solution from the old mesh onto the new mesh. Numerical examples in Section 3.3 show that h-adaptive SDG continuum models can fully resolve atomistic signals, including dispersion effects, at scales approaching the atomistic limit. The disparity in time and length scales between continuum and atomistic regions is reconciled through localized, spacetime mesh refinement without imposing small time steps on the entire continuum domain. Although the adaptive SDG model can resolve dynamic information down to sub-atomic length scales, there is a minimum length scale below which continuum theory breaks down and fails to capture dispersion and other important features of the atomistic response. Thus, the applicability of our elastodynamic model is restricted by the physical assumptions of continuum theory, but not by limited numerical resolution.

The fact that the atomic-scale configuration parameter δ can be optimized to suppress spurious reflections without introducing non-physical damping is a key finding of this work. We have shown that the optimal value of this parameter depends on certain details of the coupling scheme, such as the number of atoms involved in the trace operator, but it is invariant with respect to other model parameters, such as the direction of information propagation across the interface, pulse width and the level of mesh refinement. It is widely understood that spurious reflections will occur in coupled AtC models unless the continuum and atomistic impedances are well matched and the continuum model is able to resolve the high-frequency content of signals emanating from the atomistic zone. One of the more interesting findings in this study is that spurious reflections can occur even when these requirements are met if the coupling itself is not properly optimized. Qualitatively, there appears to be an effective impedance associated with the coupling interface itself that must be adjusted to match the rest of the model.

With hindsight, it is perhaps surprising that one would attempt to seamlessly couple continuum and atomistic models while leaving the atomic-scale gap between them to chance. For example, introducing a random gap in an otherwise uniform atomic lattice will produce a reflective surface. This holds for real physics as well as for numerical MD models, as demonstrated in Section 2.3.2. If a random gap causes *physical* reflections in a purely atomistic model, why should it not do the same in a coupled AtC model? Thus, optimization of the shift parameter δ is a physical, rather than numerical, solution to the problem of spurious reflections.

4.2. Extensions and future research

A number of extensions to the coupled Spacetime Discontinuous Galerkin and Atomistic Discontinuous Galerkin methods are under investigation in continuing work. The continuum model should be upgraded to address finite-deformation and nonlinear material response. Extensions of both the continuum and atomistic models are required to address material response at finite temperature. A hyperbolic heat conduction model is essential to capture continuum thermal response at the length and time scales under consideration. An SDG model for hyperbolic conduction is described in [51], and a fully hyperbolic thermomechanical model has been implemented and will be reported in a forthcoming paper. Atomistic-to-continuum coupling at finite temperature would utilize spatial and/or temporal filtering to separate longer-wavelength atomistic modes that are well-resolved by the continuum model from those that are not. Energy associated with the non-resolvable atomistic modes would enter the continuum model as a heat flux. Jump conditions for energy balance, written with respect to Riemann solutions of the thermomechanical problem and in combination with the Generalized Equipartition Theorem [52], would provide the necessary thermomechanical AtC coupling.

The dissipative nature of the continuum SDG method is a disadvantage relative to recent energy–momentum formulations [37] and to the modified ADG formulation in Section 2.2.2. We are currently investigating similar modifications to the SDG method that would deliver per-element balance of energy and momentum in continuum elastodynamics. If successful, these would extend existing energy–momentum methods to support unstructured, causal spacetime grids and to deliver per-element energy and momentum balance.

An AtC coupling model in which a conventional explicit MD integrator is coupled to the continuum via the ADG model might be of value in certain contexts where, for example, it is desirable to reuse an existing explicit MD code base. We developed a low-order ($p = 1$) SDG–ADG–VV coupling method along these lines. Although the method is viable, we found it to be less attractive than the SDG–ADG model due to its relative complexity and because the inconsistent mathematical structures of the ADG and VV algorithms corrupted the balance properties of the component methods. Once we established the superior efficiency of the iterative, higher-order ADG methods relative to the VV algorithm, we found little reason to continue this line of research.

While our implementation is so far limited to one spatial dimension, problems of physical interest must be modeled minimally in two and typically in three spatial dimensions. Since our formulation is dimension-independent, the challenges for implementing SDG–ADG coupling in higher dimensions are of a technical, rather than fundamental, nature. For example, our ADG implementation is general for $d = 1, 2, 3$. Adaptive spacetime meshing in the continuum domain is the most significant challenge. Robust, adaptive meshing software is available to support SDG models in 2d $\times$ time, and only modest modifications are needed to accommodate special meshing and solution requirements along AtC coupling interfaces. A two-dimensional implementation of SDG–ADG coupling is currently under development. For problems posed in three spatial dimensions, our non-adaptive spacetime meshing algorithms are provably robust, but certain parts of the theory that

support adaptive spacetime meshing require further investigation. Thus, additional progress, in spacetime meshing theory and technology, is required to realize a fully-adaptive SDG analysis capability in $3d \times$ time.

At present, we use a heuristic technique to determine the position of the AtC coupling interface that minimizes spurious reflections. Although we have good results using a fixed value for the shift parameter δ, as determined in a calibration study, it is possible that the optimal value of δ might vary over time for problems with stronger nonlinearities. Thus, a more systematic approach that automatically determines the shift as an intrinsic part of the solution procedure is desirable. The fact that δ is a configuration parameter that must be specified to determine fully the problem geometry, suggests that its optimal value might be determined by the Principle of Stationary Action. The smooth variation of the reflectance in the vicinity of the optimal values of δ, *cf.* Fig. 14, is also suggestive of a variational principle. We are investigating a variational reformulation of the SDG–ADG method based on Hamilton's principle in order to explore this possibility, and perhaps, to combine the advantages of multi-symplectic methods with those of the SDG–ADG solution technology.

Adaptive identification of regions requiring atomistic modeling is another direction for continuing investigation. While some existing techniques, such as the quasi-continuum method [12], provide a means to automatically determine where the continuum and atomistic models should hold as a function of time, our current implementation of the SDG–ADG method assumes that this decision is made *a priori* and is time-independent. However, in the context of other application domains [53], we have demonstrated the ability of our adaptive spacetime meshing capabilities to nucleate new boundaries and to track the evolution of interfaces in response to an evolving solution. Thus, the extension of the SDG–ADG method to automatically nucleate atomistic domains and to track moving AtC interfaces is an interesting direction for further development.

Acknowledgments

Support for this research was provided in part by the National Science Foundation for the Center for Process Simulation and Design (CPSD, Grant DMR 01-2169), the Materials Computation Center (MCC, Grant DMR 03-25939), and Grant DMR 07-05089 at the University of Illinois at Urbana-Champaign.

References

[1] R.E. Miller, Direct coupling of atomistic and continuum mechanics in computational materials science, Int. J. Multiscale Comput. Eng. 1 (2003) 57–72.
[2] W.K. Liu, E.G. Karpov, S. Zhang, H.S. Park, An introduction to computational nanomechanics and materials, Comput. Methods Appl. Mech. Eng. 193 (2004) 1529–1578.
[3] G. Lu, E. Kaxiras, An overview of multiscale simulations of materials, in: M. Rieth, W. Shommers (Eds.), Handbook of Theoretical and Computational Nanotechnology, vol. 10, American Scientific Publishers, 2005.
[4] R.E. Miller, E.B. Tadmor, A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods, Modell. Simul. Mater. Sci. Eng. 17 (2009) 053001.
[5] W. E, B. Engquist, The heterogeneous multiscale methods, Comm. Math. Sci. 1 (2003) 87–132.
[6] X. Li, W. E, Multiscale modeling of the dynamics of solids at finite temperature, Comput. Methods Appl. Mech. Eng. 53 (2005) 1650–1685.
[7] R.E. Rudd, J.Q. Broughton, Coarse-grained molecular dynamics and the atomic limit of finite elements, Phys. Rev. B 58 (1998) R5893.
[8] S. Curtarolo, G. Ceder, Dynamics of an inhomogeneously coarse grained multiscale system, Phys. Rev. Lett. 88 (2002) 255504.
[9] W. Wang, X. Li, C.-W. Shu, The discontinuous Galerkin method for the multiscale modeling of dynamics of crystalline solids, Multiscale Model. Simul. 7 (2008) 294–320.
[10] S.A. Silling, Reformulation of elasticity theory for discontinuities and long-range forces, J. Mech. Phys. Solids 48 (2000) 175–209.
[11] M.L. Parks, R.B. Lehoucq, S.J. Plimpton, S.A. Silling, Implementing peridynamics within a molecular dynamics code, Comput. Phys. Comm. 179 (2008) 777–783.
[12] E.B. Tadmor, M. Ortiz, R. Phillips, Quasicontinuum analysis of defects in solids, Phil. Mag. A 73 (1996) 1529–1563.
[13] W.K. Liu, H.S. Park, D. Qian, E.G. Karpov, H. Kadowaki, G.J. Wagner, Bridging scale methods for nanomechanics and materials, Comput. Methods Appl. Mech. Eng. 195 (2006) 1407–1421.
[14] S.A. Adelman, J.D. Doll, Generalized Langevin equation approach for atom/solid-surface scattering: collinear atom/harmonic chain model, J. Chem. Phys. 61 (1974) 4242–4245.
[15] S.A. Adelman, J.D. Doll, Generalized Langevin equation approach for atom/solid-surface scattering: general formulation for classical scattering off harmonic solids, J. Chem. Phys. 64 (1976) 2376–2388.
[16] P.M. Chaikin, T.C. Lubensky, Principles of Condensed Matter Physics, Cambridge University Press, Cambridge, UK, 1995.
[17] G.J. Wagner, W.K. Liu, Coupling of atomistic and continuum simulations using a bridging scale decomposition, J. Comput. Phys. 190 (2003) 249–274.
[18] W. Cai, M. de Koning, V.V. Bulatov, S. Yip, Minimizing boundary reflections in coupled-domain simulations, Phys. Rev. Lett. 85 (2000) 3213–3216.
[19] E.G. Karpov, G.J. Wagner, W.K. Liu, A Green's function approach to deriving non-reflecting boundary conditions in molecular dynamics simulations, Int. J. Numer. Methods Eng. 62 (2005) 1250–1262.
[20] E.G. Karpov, H.S. Park, W.K. Liu, A phonon heat bath approach for the atomistic and multiscale simulation of solids, Int. J. Numer. Methods Eng. 70 (2007) 351–378.
[21] W. E, Z. Huang, Matching conditions in atomistic–continuum modeling of materials, Phys. Rev. Lett. 87 (2001) 135501.
[22] W. E, Z. Huang, A dynamic atomistic–continuum method for the simulation of crystalline materials, J. Comput. Phys. 182 (2002) 234–261.
[23] S.P. Xiao, T. Belytschko, A bridging domain method for coupling continua with molecular dynamics, Comput. Methods Appl. Mech. Eng. 193 (2004) 1645–1669.
[24] R. Abedi, B. Petracovici, R.B. Haber, A spacetime discontinuous Galerkin method for linearized elastodynamics with element-wise momentum balance, Comput. Methods Appl. Mech. Eng. 195 (2006) 3247–3273.
[25] S.T. Miller, B. Kraczek, R.B. Haber, D.D. Johnson, Multi-field spacetime discontinuous Galerkin methods for linearized elastodynamics, Comput. Methods Appl. Mech. Eng. (2009), Published online; doi:10.1016/j.cma.2009.09.012.
[26] R. Abedi, R.B. Haber, S. Thite, J. Erickson, An *h*-adaptive spacetime-discontinuous Galerkin method for linear elastodynamics, Revue européenne de mécanique numérique 15 (2006) 619–642.
[27] J. Erickson, D. Guoy, J. Sullivan, A. Üngör, Building space–time meshes over arbitrary spatial domains, Engineering with Computers 20 (4) (2005) 342–353.
[28] R. Abedi, S.-H. Chung, J. Erickson, Y. Fan, M. Garland, D. Guoy, R. Haber, J. Sullivan, S. Thite, Y. Zhou, Spacetime meshing with adaptive refinement and coarsening, in: Proc. 20th Ann. ACM Symp. on Comput. Geometry, 2004, pp. 300–309.

[29] A. Üngör, A. Sheffer, Tent-pitcher: a meshing algorithm for space–time discontinuous Galerkin methods, in: Proc. 9th Int. Meshing Roundtable, 2000, pp. 161–171.
[30] W.C. Swope, H.C. Andersen, P.H. Berens, K.R. Wilson, A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: application to small water clusters, J. Chem. Phys. 76 (1982) 637–649.
[31] M.P. Allen, D.J. Tildesley, Computer Simulation of Liquids, Oxford University Press, Oxford, Great Britain, 1989.
[32] E. Hairer, C. Lubich, G. Wanner, Geometric Numerical Integration, second ed., Springer-Verlag, Heidelberg, 2006.
[33] E. Bécache, P. Joly, J. Rodríguez, Space–time mesh refinement for elastodynamics. numerical reuslts, Comput. Methods Appl. Mech. Eng. 194 (2005) 355–366.
[34] K. Eriksson, D. Estep, P. Hansbo, C. Johnson, Computational Differential Equations, Cambridge University Press, Cambridge, UK, 1996.
[35] M. Borri, C. Bottasso, A general framework for interpreting time finite element formulations, Comput. Mech. 13 (1993) 133–142.
[36] J.H. Irving, J.G. Kirkwood, The statistical mechanical theory of transport processes IV, J. Chem. Phys. 18 (1950) 817–829.
[37] M. Groß, P. Betsch, P. Steinmann, Conservation properties of a time FE method. Part IV: higher order energy and momentum conserving schemes, Int. J. Numer. Methods Eng. 63 (2005) 1849–1897.
[38] C.L. Bottasso, A new look at finite elements in time: a variational interpretation of Runge–Kutta methods, Appl. Numer. Math. 25 (1997) 355–368.
[39] B.L. Hulme, One-step piecewise polynomial Galerkin methods for initial value problems, Math. Comput. 76 (1982) 637–649.
[40] B. Leimkuhler, S. Reich, Simulating Hamiltonian Dynamics, Cambridge University Press, Cambridge, UK, 2004.
[41] James W. Demmel, Applied Numerical Linear Algebra, SIAM, Philadelphia, PA, 1997.
[42] K.F. Graff, Wave Motion in Elastic Solids, Clarendon Press, Oxford, 1975.
[43] N.W. Ashcroft, N.D. Mermin, Solid State Physics, Harcourt Brace, Orlando, 1976.
[44] F.F. Abraham, R. Walkup, H. Gao, M. Duchaineau, T. Diaz de la Rubia, M. Seager, Simulating materials failure by using up to one billion atoms and the world's fastest computer: work-hardening, Proc. Natl. Acad. Sci. USA 99 (2002) 5783–5787.
[45] M. Moseler, J. Nordiek, H. Haberland, Reduction of the reflected pressure wave in the molecular-dynamics simulation of energetic particle–solid collisions, Phys. Rev. B 56 (1997) 15439–15445.
[46] M. Born, Dynamik der Kristallgitter, B.G. Teubner, Leipzig-Berlin, 1915.
[47] J.L. Eriksen, The cauchy and born hypothesis for crystals, in: M. Rieth, W. Shommers (Eds.), Phase Transformations and Material Instabilities in Solids, Academic Press, 1984.
[48] C.A. Coulson, Alan Jeffrey, Waves: A Mathematical Approach to the Common Types of Wave Motion, second ed., Longman, London, 1977.
[49] Herbert Goldstein, Classical Mechanics, second ed., Addison-Wesley, Reading, Massachusetts, 1980.
[50] V.B. Shenoy, R. Miller, E.B. Tadmor, D. Rodney, R. Phillips, M. Ortiz, An adaptive finite element approach to atomic-scale mechanics – the quasicontinuum method, J. Mech. Phys. Solids 47 (2002) 611.
[51] S.T. Miller, R.B. Haber, A spacetime discontinuous Galerkin method for hyperbolic heat conduction, Comput. Methods Appl. Mech. Eng. 198 (2008) 194–209.
[52] Kerson Huang, Statistical Mechanics, Wiley and Sons, Inc., New York, 1963.
[53] R. Abedi, S.-H. Chung, M.A. Hawker, J. Palaniappan, R.B. Haber, Modeling evolving discontinuities with spacetime discontinuous Galerkin methods, in: A. Combescure, R. de Borst, T. Belytschko (Eds.), IUTAM Symposium on Discretization Methods for Evolving Discontinuities, IUTAM Book Series, vol. 5, Springer, 2007, pp. 59–87.