Journal of Computational Science 27 (2018) 91–106

Contents lists available at ScienceDirect

Journal of Computational Science

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

Accelerated scale-bridging through adaptive surrogate model evaluation

Kenneth W. Leiter[a,*], Brian C. Barnes[b], Richard Becker[c], Jaroslaw Knap[a]

[a] Simulation Sciences Branch, RDRL-CIH-C, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005-5066, USA
[b] Energetic Materials Science Branch, RDRL-WML-B, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005-5066, USA
[c] Impact Physics Branch, RDRL-WMP-C, U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005-5066, USA

ARTICLE INFO

Article history:
Received 11 December 2017
Received in revised form 13 March 2018
Accepted 11 April 2018
Available online 21 April 2018

Keywords:
Multiscale modeling
Surrogate modeling
Gaussian process regression
Energetic materials

ABSTRACT

Multiscale modeling is a systematic approach for the development of high-fidelity models of complex systems. However, multiscale models are often extremely computationally demanding, which precludes their use for practical applications. In this article, we introduce a computational framework for scale-bridging combined with an algorithm to automatically and adaptively replace at-scale models within a multiscale model hierarchy with surrogate models in order to reduce the computational cost of multiscale simulations. A standalone module is introduced and it is responsible for the on-the-fly construction and evaluation of surrogate models within the framework. Such an approach allows multiscale models to easily incorporate surrogate models with minimal code modifications. We employ the framework to construct a multiscale model of 1,3,5-trinitrohexahydro-s-triazine, in which a continuum finite element macroscale solver acquires equation of state through evaluation of a microscale dissipative particle dynamics model. We utilize the model for the simulation of a Taylor impact experiment and demonstrate that the error in the solution incurred by the dynamic use of surrogate models is controllable. Furthermore, we show that the use of surrogate models leads to a reduction in computational cost of between 1/20 and 1/5000 compared to a simulation evaluated without the surrogate modeling approach. In addition, we present a high-resolution simulation of a Taylor impact experiment, which is intractable without surrogate models. We illustrate how the dynamic nature of surrogate model evaluation in these simulations, while reducing computational cost, also increases load imbalance. Finally, we end with a discussion on how the inherent variability in these simulations may constitute a challenge for the current high performance computer systems given their static nature.

Published by Elsevier B.V.

1. Introduction

In the last two decades, multiscale modeling has become a dominant paradigm for building high-fidelity models of complex systems. Multiscale modeling is a divide-and-conquer approach: it strives to systematically identify relevant phenomena occurring at individual scales, both spatial and temporal, and link appropriate at-scale models into a single model hierarchy. The ability of multiscale modeling to yield high-fidelity models has been demonstrated in applications as diverse as materials modeling [1], biophysics [2–4], astrophysics [5] and weather modeling [6].

Multiscale modeling is by and large a computational endeavor. However, only lately, have computational aspects of multiscale modeling become a focus of vibrant research efforts (c.f. [7] for a review of recent developments). These efforts have been now primarily directed towards the advancement of organizing frameworks for systematic construction of multiscale models. Examples include the multiscale universal interface for partitioned-domain methods [8] and work by Borgdorff and colleagues on distributed multiscale computing for hierarchical multiscale modeling [9–11]. Recently, we have formulated a computational methodology for scale-bridging in hierarchical multiscale modeling allowing for seamless assembly of multiscale models [12,13] and employed it to build a two-scale model of the mechanical response of energetic materials [14]. Yet, even with such a relatively simple multiscale model, its computational cost has turned out to be staggering. Moreover, since multiscale models routinely encountered in practice incorporate many computationally expensive at-scale models, the computational cost of such models is bound to become even

* Corresponding author.
E-mail address: kenneth.w.leiter2.civ@mail.mil (K.W. Leiter).

https://doi.org/10.1016/j.jocs.2018.04.010
1877-7503/Published by Elsevier B.V.

higher, possibly casting doubt on the feasibility of using multiscale modeling in practical applications.

In order to significantly reduce the computational cost of complex models, surrogate models, sometimes also referred to as metamodels, have been widely employed for design optimization [15–17]. A surrogate model is a cheaper-to-evaluate approximation of a model, constructed from direct observations of the model. In design optimization, where a model is evaluated repeatedly during the optimization process, replacement of an expensive model with a surrogate model can dramatically reduce the computational cost. Many surrogate modeling methods have been introduced over the years. A comprehensive survey of surrogate modeling techniques in design optimization, including polynomial regression, kriging, multivariate adaptive regression splines, polynomial stochastic collocation, adaptive stochastic collocation, and radial basis functions can be found in [18,19].

Surrogate models are also an attractive approach to substantially lower the computational cost of multiscale models since many multiscale models rely on repeated evaluation of individual at-scale models within a multiscale model hierarchy. The pioneering work on the use of surrogate models for physics-based modeling is due to Pope [20], who employed in situ adaptive tabulation to expedite simulations involving combustion chemistry. The use of surrogate models in multiscale modeling of materials has been explored by Knap et al. [21] and Barton et al. [22] who replaced the evaluation of a crystal plasticity based constitutive law with dynamically constructed kriging-based surrogate models. Recently, Roehm et al. [23] utilized a similar approach in an elastodynamics model. In atomistic models, Li et al. [24] developed an approach to adaptively replace the computation of atomic forces from quantum mechanical models with forces obtained via interpolation of previously stored calculations. For a multiscale biomechanics model, Wirtz et al. [25] investigated the use of kernel surrogate models to learn the interface between macroscale and microscale models.

In this article, we describe a computational methodology for the dynamic construction and evaluation of surrogate models to reduce the computational cost of hierarchical multiscale simulations. The methodology allows for incorporation of surrogate models into any multiscale model with minimal changes. We formulate our methodology in Section 2 and detail its application to construct a two-scale model of an energetic material in Section 3.

2. Multiscale computational framework

2.1. A modular computational framework for scale-bridging

The point of departure for our work is the computational framework for scale-bridging in multiscale simulations of Knap et al. [12], which allows to combine individual at-scale model components together to form a multiscale model. The simplest multiscale model that can be assembled in the framework consists of two at-scale models, the macroscale model F and the microscale model f (c.f. Fig. 1). The macroscale model is a mapping $F : I \times D \mapsto R$, where I is a collection of microscale models, domain $D \subset \mathbb{R}^N$, and range $R \subset \mathbb{R}^M$. Similarly, the microscale model is a mapping $f : \hat{D} \mapsto \hat{R}$ where $\hat{D} \subset \mathbb{R}^n$ and $\hat{R} \subset \mathbb{R}^m$ denote the domain and range of f, respectively. In addition, the framework includes two mappings to transform data between at-scale models. The mapping $G : \tilde{D} \mapsto \hat{D}$, where $\tilde{D} \subset \mathbb{R}^{\tilde{n}}$ is the set of intermediate values derived from values in D by F. Henceforth, we refer to G as the "input filter" since it generates the input to f in the set $\hat{D}$. Likewise, the mapping $g : \hat{R} \mapsto \tilde{R}$, where $\tilde{R} \subset \mathbb{R}^{\tilde{m}}$, is referred to as the "output filter" as it extracts relevant data from the microscale model output to be passed to the macroscale model. More complex multiscale models can, of course, be formed through assemblies of multiple two-scale model building blocks.

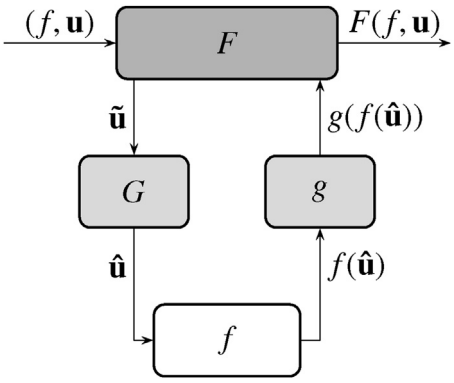

Fig. 1. A two-scale model consisting of macroscale model F and microscale model f. Two mappings transform data between scales: the input filter G which transforms data into an appropriate form for the microscale model and the output filter g which extracts relevant data from the microscale model to inform the macroscale model.

The computational framework closely follows the heterogeneous multiscale method for multiscale model development in which a macroscale model obtains missing required data through the constrained evaluation of a microscale model [26].

The essential component of the scale-bridging framework is the Evaluation Module, which coordinates the interaction between F and f. We refer to the act of sending of $\tilde{\mathrm{u}} \in \tilde{D}$ from F to the Evaluation Module as an "evaluation request" owing to the fact that this action initiates the process of computing the value of (evaluating) $f(\hat{\mathrm{u}})$ for $\hat{D} \ni \hat{\mathrm{u}} = G(\tilde{\mathrm{u}})$. The Evaluation Module serves to: (1) collect requests for evaluation of f from F; (2) apply the input filter to the evaluation requests to prepare input data for microscale models; (3) schedule evaluation requests on available resources; (4) monitor progress of evaluations to detect completion and handle failures; and (5) apply the output filter to extract relevant data from completed f evaluations to return to F. The use of the Evaluation Module offers practical benefits for the construction of multiscale models. First and foremost, microscale models can be incorporated into a multiscale model without the need to modify their computer implementations. Similarly, minimal changes are required to the implementation of a macroscale model and amount to merely adding a function call to request an evaluation of the microscale model. Second, the separation between the two at-scale models allows for their asynchronous evaluation on potentially disparate computational resources. For example, the macroscale model may execute on a conventional processor, while the microscale model on a co-processor. An illustration of a potential layout for a multiscale application employing Evaluation Module on a parallel computer is given in Fig. 2a.

2.2. Automatic surrogate model construction for scale-bridging

Many microscale models routinely encountered in practice are very computationally demanding. For example, dissipative particle dynamics models containing a modest number of particles may often require hours of compute time [27]. Further, when chemical reactions are treated within dissipative particle dynamics models, days of compute time may be necessary per a single simulation [28]. Under such circumstances, the computational cost of a multiscale model may render its use utterly impractical, especially when a repeated evaluation of a costly microscale model is needed. However, in many applications the macroscale model confines the microscale model to a region of its configurational space. Then, it may be advantageous to utilize previous evaluations of the microscale model to construct an approximation to it: a surrogate model. Hereafter, we refer to a surrogate model corresponding to the microscale as f_ε.