

MODEL REDUCTION WITH MAPREDUCE-ENABLED TALL AND SKINNY SINGULAR VALUE DECOMPOSITION*

PAUL G. CONSTANTINE[†], DAVID F. GLEICH[‡], YANGYANG HOU[‡], AND JEREMY TEMPLETON[§]

Abstract. We present a method for computing reduced-order models of parameterized partial differential equation solutions. The key analytical tool is the singular value expansion of the parameterized solution, which we approximate with a singular value decomposition of a parameter snapshot matrix. To evaluate the reduced-order model at a new parameter, we interpolate a subset of the right singular vectors to generate the reduced-order model's coefficients. We employ a novel method to select this subset that uses the parameter gradient of the right singular vectors to split the terms in the expansion, yielding a mean prediction and a prediction covariance—similar to a Gaussian process approximation. The covariance serves as a confidence measure for the reduced-order model. We demonstrate the efficacy of the reduced-order model using a parameter study of heat transfer in random media. The high-fidelity simulations produce more than 4TB of data; we compute the singular value decomposition and evaluate the reduced-order model using scalable MapReduce/Hadoop implementations. We compare the accuracy of our method with a scalar response surface on a set of temperature profile measurements and find that our model better captures sharp, local features in the parameter space.

Key words. model reduction, simulation informatics, MapReduce, Hadoop, tall-and-skinny SVD

AMS subject classification. 65F99

DOI. 10.1137/130925219

1. Introduction and motivation. High-fidelity simulations of partial differential equations are typically too expensive for design optimization and uncertainty quantification, where many independent runs are necessary. Cheaper reduced-order models (ROMs) that approximate the map from simulation inputs to quantities of interest may replace expensive simulations to enable such parameter studies. These ROMs are constructed with a relatively small set of high-fidelity runs chosen to cover a range of input parameter values. Each evaluation of the ROM is a linear combination of basis functions derived from the outputs of the high-fidelity runs; ROM constructions differ in their choice of basis functions and method for computing the coefficients of the linear combination. Projection-based methods project the residual of the governing equations (e.g., a Galerkin projection) to create a relatively small system of equations for the coefficients; see the recent preprint [4] for a survey of projection-based techniques. Alternatively, one may derive a closely related optimization problem to compute the coefficients [9, 10, 14]. These two formulations can provide a measure

*Received by the editors June 17, 2013; accepted for publication (in revised form) February 7, 2014; published electronically October 30, 2014. This work was supported by Sandia's Computer Science Applied Research (CSAR) and the Advanced Simulation and Computing (ASC) programs.

<http://www.siam.org/journals/sisc/36-5/92521.html>

[†]Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO 80401 (paul.constantine@mines.edu).

[‡]Computer Science, Purdue University, West Lafayette, IN 47907 (dgleich@purdue.edu, hou13@purdue.edu).

[§]Sandia National Laboratories, Livermore, CA 94551 (jatempl@sandia.gov). Sandia National Laboratories is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

of confidence along with the ROM. However, they are often difficult to implement in existing solvers since they need access to the equation’s operators or residual.

To bypass implementation difficulties, one may use black-box methods—e.g., collocation [3, 33] or Gaussian process regression [30]—which are essentially interpolation methods applied to the high-fidelity outputs. They do not need access to the differential operators or residuals and are, therefore, relatively easy to implement. However, measures of confidence are more difficult to formulate and compute. Several works have explored using interpolation instead of projection or optimization to compute the ROM coefficients. This approach is justified when the governing equations are unknown or only a set of PDE solutions are available [25]. It is also useful when nonlinearities in the governing equations prohibit a theoretically sound Galerkin projection [1, 2]. For these reasons, it has been applied to several problems in aerospace engineering [8, 23, 18, 28].

In this paper, we extend these ideas by equipping an interpolation-based ROM with a novel parameter-dependent confidence measure. In the taxonomy of Eldred, Giunta, and Collis [17], this may be considered a hybrid of data-fitting and projection-based ROMs. We first view the ROM from the perspective of the singular value expansion (SVE) of the parameterized PDE solution, where the left singular functions depend on space and time, and the right singular functions depend on the input parameters. We approximate these functions with the components of a singular value decomposition (SVD) of a tall, dense matrix of solutions computed at a set of input parameter values, i.e., parameter *snapshots*. Many reduced basis methods use the left singular vectors of the snapshot matrix for the ROM basis, where each snapshot represents a spatially varying solution. In contrast, each column of our snapshot matrix contains the full spatio-temporal solution for a given parameter. We have observed interesting behavior in the right singular vectors in several applications: as the index of the singular vector increases, its components—when viewed as evaluations of a parameter-dependent function—become more oscillatory. This is consistent with the common use of the phrase “higher order modes” to describe the singular vectors with large indices. More importantly, the rate that the singular vectors become oscillatory depends on the parameter. In particular, the singular vectors become oscillatory faster in regions of the parameter space where the PDE solution changes rapidly with small parameter perturbations.

We exploit this observation to devise a heuristic for choosing the subset of the left singular vectors comprising ROM; instead of selecting all left singular vectors whose corresponding singular value is above a chosen threshold, we examine the gradients of the right singular vectors at the parameter value where we wish to evaluate the ROM. After some critical index, the right singular vectors are too irregular to safely interpolate. For each ROM evaluation, the right singular vectors are divided into two categories: (i) those that are sufficiently smooth for interpolation, and (ii) those that oscillate too rapidly. The first category identifies the left singular vectors used in the ROM, and the coefficients are computed with interpolation. The remaining left singular vectors are used to compute a measure of confidence similar to the prediction variance in Gaussian process regression. The number of left singular vectors in each category may be different for different ROM evaluations depending on the irregularity of the right singular vectors at the interpolation point; we explore this in the numerical examples. The heuristics we employ to categorize the right singular vectors are based on the work of Hansen in the context of ill-posed inverse problems [20]. We describe the ROM methodology in section 2.

In section 4, we demonstrate the ROM and its confidence measure with a parameter study of heat transfer in random media. A brick is heated on one side, and we measure how much heat transfers to the opposite side given the parameterized thermal conductivity of the material. The high-fidelity simulations use Sandia National Laboratories’ finite element production code Aria [27] on its capacity cluster with a mesh containing 4.2 million elements. The study uses 8192 simulations, which produce approximately 4 terabytes (TB) of data. We study the effectiveness of the reduced-order model and compare its predictions with a response surface on two relevant scalar quantities of interest computed from the full temperature distribution.

Given the data-intensive computing requirements of this and similar applications, we have chosen to implement the ROM in the popular Hadoop distribution [12] of MapReduce [15]. By expressing each step of the ROM construction in the MapReduce framework, we take advantage of its built-in parallelism and fault tolerance. MapReduce’s scalability enables us to compute the SVD of a snapshot matrix with 64 columns and roughly five billion rows—approximately 2.3TB of data—without the custom hard disk I/O that would be necessary to use standard parallel routines such as Scalapack [7]. The algorithm we use for the SVD in MapReduce is based on the communication-avoiding QR decomposition [16] as described in our previous work [13]. We present the implementation of the ROM in section 3.

2. A reduced-order modeling approach. Let $f = f(x, s)$ be the solution of a partial differential equation (PDE), where $x \in \mathcal{X} \subset \mathbb{R}^4$ are the space and time coordinates (three spatial dimensions and a temporal dimension), and $s \in \mathcal{S} \subset \mathbb{R}$ is an input parameter. We restrict our attention to models with a single scalar parameter to keep the presentation simple. The ROM employs both interpolation and approximations of derivatives in the parameter space \mathcal{S} . While these operations are possible with more than one parameter, challenges arise in higher dimensions—e.g., the choice of interpolation nodes and the accuracy of derivatives—that we will avoid here. If one is willing to address the difficulties of interpolation and approximating derivatives in multiple dimensions, then our approach can be extended.

The interpretation of f as the solution of a PDE is important for two reasons. First, the solution of many PDE models can be shown to be a smooth function of both the space/time variables and the parameter, and we restrict our attention to sufficiently smooth solutions. Second, computational tools will compute f at all values of x given an input s . In other words, we cannot evaluate f at a specific x without evaluating it for every x . We will not discuss refinement of the ROM—i.e., which parameter values to use for a new set of high-fidelity runs to best improve an initial ROM—but our heuristic offers a natural criterion for such a selection. We assume that computing $f(x, s)$ for a particular s is computationally expensive. We want to use the outputs from a few expensive computations at a chosen set of input parameters to approximate $f(x, s)$ at some other s in a manner that is less computationally expensive than solving the differential equation.

We assume that f is continuous and square-integrable ($\int_{\mathcal{X}} \int_{\mathcal{S}} f^2 ds dx < \infty$). In practice, the techniques we use will perform better if f is smooth, e.g., admits continuous derivatives up to some order. Since f is continuous, it admits a uniformly convergent series representation known as the *singular value expansion* (SVE); see [19] for more details on the SVE:

$$(2.1) \quad f(x, s) = \sum_{k=1}^{\infty} \mu_k u_k(x) v_k(s).$$

The singular functions $u_k(x)$ and $v_k(s)$ are continuous and orthonormal,

$$(2.2) \quad \int_{\mathcal{X}} u_{k_1} u_{k_2} dx = \int_{\mathcal{S}} v_{k_1} v_{k_2} ds = \delta_{k_1, k_2}.$$

The singular values are positive and ordered in decreasing order,

$$(2.3) \quad \mu_1 \geq \mu_2 \geq \dots \geq 0.$$

Hansen discusses methods for approximating the factors of the SVE using the SVD. We will employ his construction [19, section 5], which ultimately uses point evaluations of the function f to construct a matrix suited for the SVD.

Let x_1, \dots, x_M with $x_i \in \mathcal{X}$ be the points of a discretization of the spatio-temporal domain. A run of the PDE solver produces an approximate solution at these points in the domain for a given input s . We assume that the spatio-temporal discretization is sufficient to produce an accurate approximation of the PDE solution for all values of s ; in practice, such an assumption should be verified. Let s_1, \dots, s_N with $s_j \in \mathcal{S}$ be a set of input parameters where the approximate PDE solution will be computed; we call these the *training* runs. The number N is the budget of simulations, and we expect that $N \ll M$ for most cases. In other words, we assume that the number of nodes in the spatio-temporal discretization is much larger than the budget of simulations.

From these approximate solutions, we construct the tall, dense matrix

$$(2.4) \quad \mathbf{F} = \begin{bmatrix} f(x_1, s_1) & \cdots & f(x_1, s_N) \\ \vdots & \ddots & \vdots \\ f(x_M, s_1) & \cdots & f(x_M, s_N) \end{bmatrix}.$$

Next we compute the thin SVD,

$$(2.5) \quad \mathbf{F} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad \mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_N),$$

where, following [19], we treat $\sigma_k \approx \mu_k$ and

$$(2.6) \quad \mathbf{U} \approx \begin{bmatrix} u_1(x_1) & \cdots & u_N(x_1) \\ \vdots & \ddots & \vdots \\ u_1(x_M) & \cdots & u_N(x_M) \end{bmatrix}, \quad \mathbf{V} \approx \begin{bmatrix} v_1(s_1) & \cdots & v_N(s_1) \\ \vdots & \ddots & \vdots \\ v_1(s_N) & \cdots & v_N(s_N) \end{bmatrix}.$$

In other words, we treat the entries of the left and right singular vectors as evaluations of the singular functions at the points x_i and s_j , respectively.

2.1. Oscillations in the singular vectors. We will leverage the work of Hansen [21, 20] on computational methods for linear, ill-posed inverse problems to develop the heuristics for the ROM. He observed that, for a broad class of integral equation kernels found in practice, the singular functions become more oscillatory (i.e., cross zero more frequently) as the index k increases. Rigorous proofs of this observation are available for some special cases. However, it is easy to construct kernels whose singular functions do not behave this way. For example, take a kernel whose singular functions become more oscillatory with increasing k and shuffle the singular functions. Such counterexamples offer evidence that a general statement is difficult to formulate.

We have observed similar phenomena for functions $f = f(x, s)$ coming from parameterized partial differential equations. This observation is corroborated by many

studies in coherent structures based on the closely related proper orthogonal decomposition [24]. Additionally, we have observed that these oscillations may not increase uniformly over the parameter domain. In particular, the rate of increasing oscillations may be greater in regions of the parameter space where f has large parameter gradients; we provide two illustrative examples below.

The components of the singular vectors inherit the observed oscillating behavior of the singular functions. In particular, the oscillations increase as the index k increases, and they increase more rapidly in regions corresponding to large differences in the elements of the data matrix \mathbf{F} . These rapid oscillations manifest as an increase with k in the magnitude of the difference between entries of the singular vectors corresponding to evaluations of the singular functions that are nearby in parameter space. For example, the difference between $\mathbf{V}_{j,k+1} \approx v_{k+1}(s_j)$ and $\mathbf{V}_{j',k+1} \approx v_{k+1}(s_{j'})$ with s_j neighboring $s_{j'}$ will be greater than the difference between $\mathbf{V}_{j,k} \approx v_k(s_j)$ and $\mathbf{V}_{j',k} \approx v_k(s_{j'})$. (Note that when the model contains more than one parameter, the notion of *neighboring* becomes more complicated.) However, since there is finite resolution in the parameter space, there is typically some k after which the discretization is insufficient to represent the oscillations, and this pattern breaks down. The phenomenon is similar to approximating a sequence of sine waves with increasing frequency using the same grid. We formalize this notion in the following assumption.

Assumption 1. Let $\mathcal{S} = [s_1, s_N]$ be a closed interval with a discretization $s_i = s_1 + (i - 1)\Delta s$, where $\Delta s = (s_N - s_1)/(N - 1)$, and let \mathbf{V} be defined as in (2.6). There is an $R = R_i \leq N$ such that the sequence of difference magnitudes between neighboring right singular vector entries will increase for k from 1 to R , i.e.,

$$(2.7) \quad |\mathbf{V}_{i+1,k+1} - \mathbf{V}_{i,k+1}| > |\mathbf{V}_{i+1,k} - \mathbf{V}_{i,k}|.$$

For $k > R$, the relationship becomes unpredictable due to the finite resolution in the parameter space \mathcal{S} .

Note our restriction to a single parameter and a uniform discretization of the parameter space. These restrictions can be relaxed with appropriate discretizations of a multivariate space. We will use Assumption 1 to justify a heuristic that distinguishes between singular functions that can be resolved and those that cannot given the discretization.

Next we give two concrete examples of the observed behavior in the right singular vectors. The first is a steady state advection-diffusion type boundary value problem,

$$(2.8) \quad \frac{df}{dx} + s \frac{d^2 f}{dx^2} = -1, \quad x \in [-10, 10], \quad s \in [2, 20],$$

with homogeneous boundary conditions. The solution is given by

$$(2.9) \quad f(x, s) = \frac{\exp(20s^{-1})(x - 10) + 20 \exp(s^{-1}(10 - x)) - x - 10}{1 - \exp(20s^{-1})}.$$

The parameter s represents the ratio of diffusion to advection. Figure 1 shows the results of the SVD approximation to the SVE factors for an overresolved model (1999 points in the discretization of the parameter space) and an underresolved model (15 points in the parameter space). Observe how the first seven singular functions scaled by their respective singular values become oscillatory at different rates in different regions of the parameter space. In particular, the singular functions oscillate more rapidly in regions of the parameter space corresponding to more advection. Also note

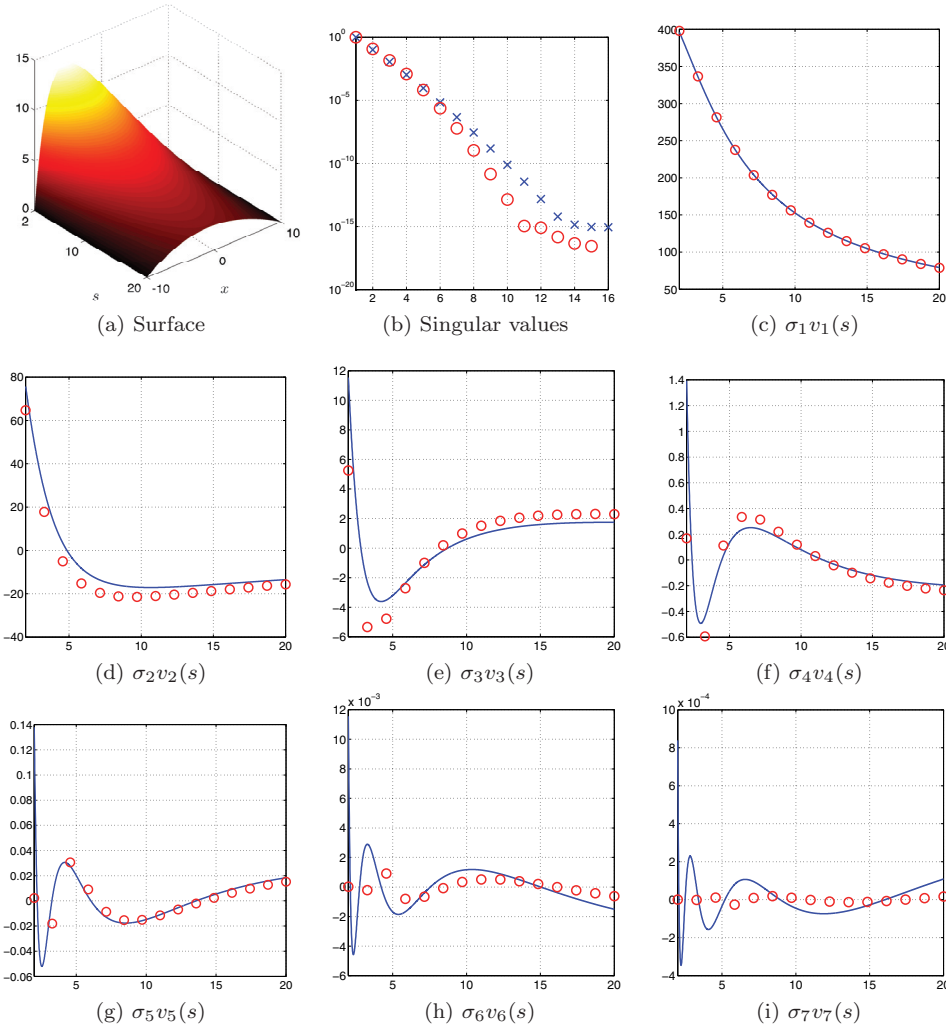


FIG. 1. The top left figure shows the solution $f(x, s)$ to the advection-diffusion equation (2.8). The figure to its right shows the singular values of a finely resolved \mathbf{F} with 1999 columns (blue 'x's) and a coarsely resolved \mathbf{F} with 15 columns (red 'o's). Each are scaled by the maximum singular value from each set. The remaining figures show the approximations of the singular functions $v_1(s)$ through $v_7(s)$ scaled by the respective singular values for \mathbf{F} with 1999 columns (blue lines) and 15 columns (red 'o's'). (Colors are visible in the electronic version.)

how the underresolved approximations deviate from the overresolved approximations in regions of high oscillations.

The second example is another second order boundary value problem with spatially varying coefficients,

$$(2.10) \quad -\frac{d}{dx} \left(a \frac{df}{dx} \right) = 1, \quad x \in [0, 1],$$

with homogeneous boundary conditions, and

$$(2.11) \quad a = a(x, s) = 1 + 4s(x^2 - x), \quad s \in [0.1, 0.9].$$

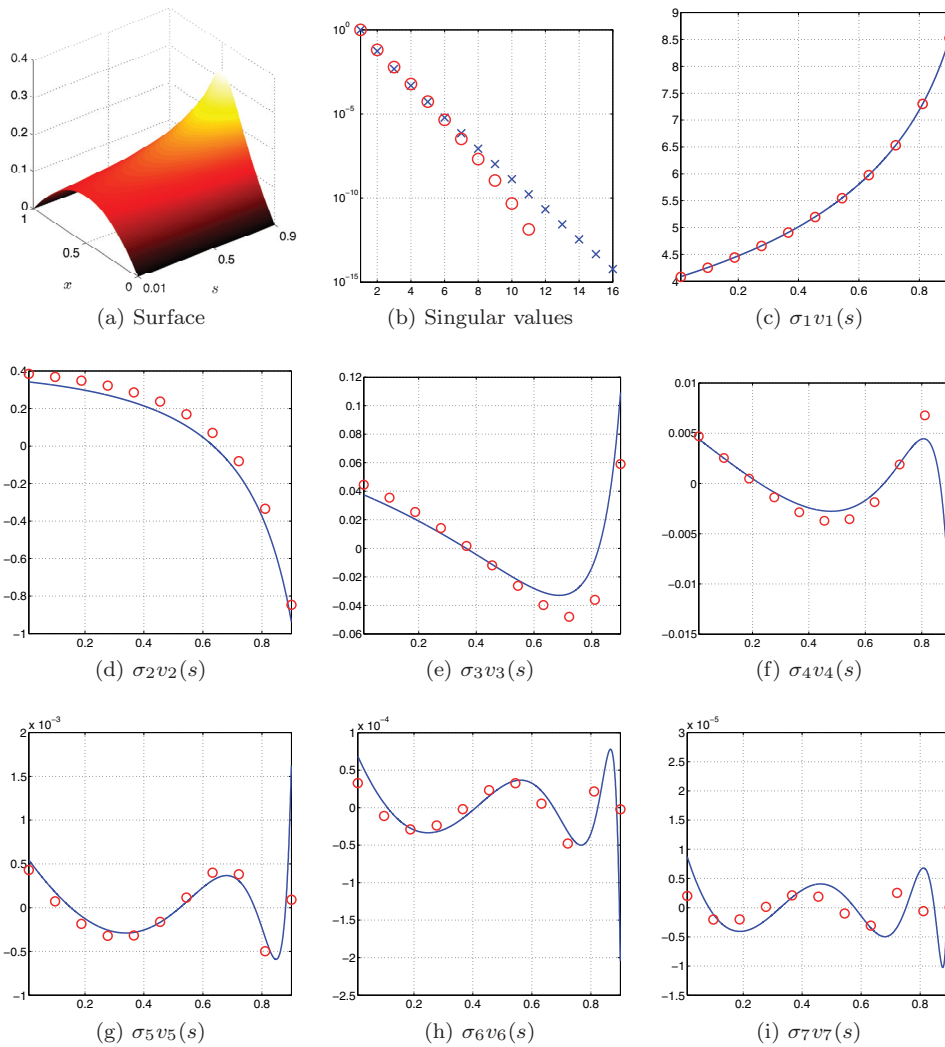


FIG. 2. The top left figure shows the solution $f(x, s)$ to the boundary value problem (2.10). The figure to its right shows the singular values of a finely resolved \mathbf{F} with 1999 columns (blue x 's) and a coarsely resolved \mathbf{F} with 11 columns (red o 's). Each are scaled by the maximum singular value from each set. The remaining figures show the approximations of the singular functions $v_1(s)$ through $v_7(s)$ scaled by the respective singular values for \mathbf{F} with 1999 columns (blue lines) and 11 columns (red o 's). (Colors are visible in the electronic version.)

The solution is

$$(2.12) \quad f(x, s) = -\frac{1}{8s} \log(1 + 4s(x^2 - x)),$$

which is plotted in the top left of Figure 2. Outside the domain, f has a singularity at $(x = 0.5, s = 1)$, which causes f to grow rapidly along the line $x = 0.5$ near the boundary $s = 0.9$. This local feature of the solution results in more rapid oscillations of the singular functions $v_k(s)$ near the boundary $s = 0.9$. The first seven singular functions, scaled by the singular values, are plotted in Figure 2. The rapid oscillations near the parameter boundary $s = 0.9$ are clearly visible. In those same figures, we plot

the components of the corresponding singular vectors, scaled by the singular values, for a data matrix \mathbf{F} with columns computed at eleven equally spaced parameter values in the interval $[0.1, 0.9]$. Notice how the components of the singular vectors deviate from the singular functions as k increases, particularly in the regions of rapid oscillations.

In the next sections, we will exploit the observation of nonuniformly increasing oscillations in the right singular vectors to devise a heuristic for the ROM.

2.2. Constructing the reduced-order model. Recall that the goal is to approximate $f(x, s)$ for some input s that was not used to compute a training run. We can use the existence of the SVE to justify the following approach. Since we treat the components of the right singular vectors \mathbf{V} as evaluations of the singular functions $v_k(s)$, we can interpolate between the singular vector components to approximate the singular functions at new values of s . More precisely, define

$$(2.13) \quad \tilde{v}_k(s) = \mathcal{I}(s; v_k(s_1), \dots, v_k(s_N)),$$

where \mathcal{I} is an interpolation operator that takes a value of s and the components of the singular vector as arguments. The form of the interpolant may depend on the selection of the points s_j . For example, if these points are the Chebyshev points or the nodes of a Gaussian quadrature rule, then high order global polynomial interpolation is possible. If the points are uniformly spaced, then one may use piecewise polynomials or radial basis functions.

Unfortunately, the increasingly oscillatory character of the functions $v_k(s)$ as k increases combined with the fixed discretization s_j causes concern for any chosen interpolation procedure as k approaches N . In other words, the smoothness of $v_k(s)$ decreases as k increases, which diminishes confidence in the interpolation accuracy. Therefore, we seek to divide the right singular vectors into two groups: those that are smooth enough to accurately interpolate and those that are not. Specifically, we seek an $R = R(s)$ with $R \leq N$ such that for $k \leq R$ we have confidence in the accuracy of the interpolant $\tilde{v}_k(s)$. We treat the remaining interpolations with $k > R$ as unpredictable, and we model them with a random variable. We will discuss the choice of R in the next section.

Given R , we model the PDE output at the space-time coordinate x_i for the new parameter value s as

$$(2.14) \quad \tilde{f}(x_i, s) = \sum_{k=1}^R \sigma_k u_k(x_i) \tilde{v}_k(s) + \sum_{k=R+1}^N \sigma_k u_k(x_i) \eta_{k-R},$$

where η_k are uncorrelated random variables with mean zero and variance one; these represent the uncertainty in the interpolation procedure for increasingly oscillatory functions. Under this construction, the vector of values $\tilde{f}(x_i, s)$ is a random vector with mean and covariance,

$$(2.15) \quad \begin{aligned} \mathbb{E} [\tilde{f}(x_i, s)] &= \sum_{k=1}^R \sigma_k u_k(x_i) \tilde{v}_k(s), \\ \text{Cov} [\tilde{f}(x_i, s), \tilde{f}(x_j, s)] &= \sum_{k=R+1}^N \sigma_k^2 u_k(x_i) u_k(x_j). \end{aligned}$$

The reduced-order model we propose is the mean of this random vector,

$$(2.16) \quad f(x_i, s) \approx \mathbb{E} \left[\tilde{f}(x_i, s) \right].$$

The diagonal components of the covariance matrix provide a measure of confidence for the reduced-order model at each x_i , similar to the prediction variance of a Gaussian process regression model [30].

Next we show that the reduced-order model is equivalent to applying the interpolation procedure independently to the rows of a low rank approximation of the matrix \mathbf{F} . To set up the notation, partition

$$(2.17) \quad \mathbf{U} = [\mathbf{U}_1 \quad \mathbf{U}_2], \quad \mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_1 & \\ & \mathbf{\Sigma}_2 \end{bmatrix}, \quad \mathbf{V} = [\mathbf{V}_1 \quad \mathbf{V}_2],$$

where \mathbf{U}_1 , $\mathbf{\Sigma}_1$, and \mathbf{V}_1 contain R columns. Then

$$(2.18) \quad \begin{aligned} \mathbf{F} &= \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \\ &= \mathbf{U}_1\mathbf{\Sigma}_1\mathbf{V}_1^T + \mathbf{U}_2\mathbf{\Sigma}_2\mathbf{V}_2^T \\ &= \mathbf{F}_1 + \mathbf{F}_2 \\ &= \begin{bmatrix} f^{(1)}(x_1, s_1) & \cdots & f^{(1)}(x_1, s_N) \\ \vdots & \ddots & \vdots \\ f^{(1)}(x_M, s_1) & \cdots & f^{(1)}(x_M, s_N) \end{bmatrix} + \begin{bmatrix} f^{(2)}(x_1, s_1) & \cdots & f^{(2)}(x_1, s_N) \\ \vdots & \ddots & \vdots \\ f^{(2)}(x_M, s_1) & \cdots & f^{(2)}(x_M, s_N) \end{bmatrix}. \end{aligned}$$

Then we have the following proposition.

PROPOSITION 1. *If \mathcal{I} from (2.13) is a linear operation, then*

$$(2.19) \quad \begin{aligned} \mathbb{E} \left[\tilde{f}(x_i, s) \right] &= \mathcal{I} \left(s; f^{(1)}(x_i, s_1), \dots, f^{(1)}(x_i, s_N) \right), \\ \text{Cov} \left[\tilde{f}(x_i, s), \tilde{f}(x_j, s) \right] &= \sum_{k=R+1}^N f^{(2)}(x_i, s_k) f^{(2)}(x_j, s_k). \end{aligned}$$

Proof. For a function $g = g(s)$ with evaluations $g(s_j)$, the linear interpolation can be written

$$(2.20) \quad \mathcal{I}(s; g(s_1), \dots, g(s_N)) = \sum_{j=1}^N w_j g(s_j)$$

for some set of weights $w_j = w_j(s)$. Then,

$$\begin{aligned}
 \mathbb{E} \left[\tilde{f}(x_i, s) \right] &= \sum_{k=1}^R \sigma_k u_k(x_i) \tilde{v}_k(s) \\
 &= \sum_{k=1}^R \sigma_k u_k(x_i) \left(\sum_{j=1}^N w_j v_k(s_j) \right) \\
 (2.21) \quad &= \sum_{j=1}^N w_j \left(\sum_{k=1}^R \sigma_k u_k(x_i) v_k(s_j) \right) \\
 &= \sum_{j=1}^N w_j f^{(1)}(x_i, s_j) \\
 &= \mathcal{I} \left(s; f^{(1)}(x_i, s_1), \dots, f^{(1)}(x_i, s_N) \right),
 \end{aligned}$$

as required. The covariance expression is easily proved using the linear algebra notation. Define the $N \times N$ matrix $\mathbf{C}_{ij} = \text{Cov}[\tilde{f}(x_i, s), \tilde{f}(x_j, s)]$. Then by the orthogonality of the columns of \mathbf{V}_2 ,

$$(2.22) \quad \mathbf{C} = \mathbf{U}_2 \mathbf{\Sigma}_2^2 \mathbf{U}_2^T = \mathbf{U}_2 \mathbf{\Sigma}_2 \mathbf{V}_2^T \mathbf{V}_2 \mathbf{\Sigma}_2^T \mathbf{U}_2^T = \mathbf{F}_2 \mathbf{F}_2^T$$

as required. \square

2.3. Choosing R . We must still choose R that determines the split between smooth and nonsmooth singular vectors. We will exploit the observation of the oscillating singular variation vectors from section 2.1 and make use of Assumption 1. We define the following *variation metric*:

$$(2.23) \quad \tau(r, s) = \sum_{k=1}^r \left| \frac{\mathbf{V}_{j+1,k} - \mathbf{V}_{j,k}}{\Delta s} \right| \quad \text{for } s_j \leq s < s_{j+1}.$$

By Assumption 1, τ is an increasing function of r up to some $R = R(s)$. Generally, if τ is too large, then we have entered the range of k where interpolations of $v_k(s_j)$ are not to be trusted. We will quantify this with a threshold $\bar{\tau}$. Given $\bar{\tau}$, we choose $R = R(s, \bar{\tau})$ to be the largest r such that $\tau(r, s) \leq \bar{\tau}$.

To determine the appropriate threshold $\bar{\tau}$, we use a set of PDE evaluations $f_\ell = f(x, s_\ell)$ with $\ell = 1, \dots, L$ for testing, where s_ℓ is not in the training set (i.e., $s_\ell \neq s_j$ for any ℓ or j). We choose a set of candidate thresholds $\bar{\tau}_m$. For each testing model and each candidate threshold, we compute the relative error

$$(2.24) \quad \mathcal{E}(s_\ell, \bar{\tau}_m) = \left[\sum_{i=1}^M \left(f(x_i, s_\ell) - \sum_{k=1}^R \sigma_k u_k(x_i) \tilde{v}_k(s_\ell) \right)^2 \right]^{1/2} \left/ \left[\sum_{i=1}^M f(x_i, s_\ell)^2 \right]^{1/2} \right.,$$

where $R = R(s_\ell, \bar{\tau}_m)$. These errors can be visualized, and the final threshold is chosen so that the error in the testing set is relatively small.

We demonstrate this process using the boundary value problem from (2.10). The training models consist of solutions computed at eleven equally spaced values of the parameter s in the range $[0.1, 0.9]$. We compute a test model at the midpoint of each




FIG. 3. The log of the relative error in the mean prediction of the ROM as a function of s and $\bar{\tau}$. (Colors are visible in the electronic version.)

interval $[s_j, s_{j+1}]$, where s_j was used to compute the training models. The range of the variation metric τ from (2.23) for these testing sites is roughly 0.1 to 42.8. We choose 20 candidate thresholds $\bar{\tau}_m$ in this range and compute the error in the reduced-order model at the testing sites (see (2.24)) for each candidate threshold. These errors are displayed in Figure 3. We want to choose the split R such that the ROM uses the fewest left singular vectors with the maximal accuracy; using fewer singular vectors reduces the computational work and produces a simpler model. The errors in Figure 3 show that the reduced-order model is as accurate as possible for each testing site after the fourth candidate threshold, which is roughly $\bar{\tau} = 6.84$. For this threshold, Table 1 displays the split R between the right singular vectors that admit an accurate interpolant and those that are modeled with a random variable for each of the testing sites. Notice that the number of smooth right singular vectors is smaller for testing sites near the boundary $s = 0.9$ of the parameter domain, which is precisely what we would have expected.

We can compare the splitting strategy based on $R = R(s)$ with a standard truncation strategy based on the magnitudes of the singular values of \mathbf{F} . The mean of the random vector (2.16) is equivalent to interpolating a truncated SVD approximation of the data matrix \mathbf{F} , as shown in Proposition 1. However, the magnitudes of the singular values provide no insight into the uncertainty in the interpolation procedure. Our splitting strategy chooses a different truncation for the mean (2.16) for each s based on the capability of the interpolation procedure to accurately approximate the right singular functions $v_k(s)$ at the point s . The singular values that are not in the mean contribute to the covariance-based confidence measure from (2.15). A global truncation based on the singular values would create the same prediction variance for

TABLE 1

The split and the corresponding ROM error for $\bar{\tau} = 6.84$ and different values of s .

s	$R(s, \bar{\tau})$	$\mathcal{E}(s, \bar{\tau})$
0.0545	4	0.0006
0.1435	5	0.0007
0.2325	6	0.0008
0.3215	6	0.0010
0.4105	5	0.0013
0.4995	6	0.0017
0.5885	4	0.0023
0.6775	5	0.0034
0.7665	3	0.0059
0.8555	2	0.0129

every s , and it would always be on the order of the largest truncated singular value. In other words, it provides no information on how the confidence in the prediction changes as s varies.

3. Implementation in Hadoop. Constructing the ROM for highly resolved simulations (i.e., large M) requires significant data processing. We have implemented the construction in the MapReduce framework, which enables us to take advantage of Hadoop for large-scale distributed data processing. To construct the reduced-order model, the outputs from the high-fidelity simulations are sent to and stored in the Hadoop cluster. Our implementation then proceeds in three steps:

1. Create the tall-and-skinny matrix \mathbf{F} from the simulation data.
2. Compute the SVD of \mathbf{F} .
3. Generate the coefficients of the reduced-order model and evaluate solutions from the ROM.

In what follows, we give a very brief overview of the MapReduce framework, and then we describe each step of the implementation in Hadoop.

3.1. MapReduce/Hadoop. Google devised MapReduce because of the frustration programmers experienced as they constantly juggled the complexity of developing distributed, fault-tolerant data computational algorithms [15]. Early data-intensive computing at Google was a complex mix of ad hoc scripts. Their solution was the MapReduce computation model: a simple, general interface for a common template behind their data analysis tasks that hides the details of the parallel implementations from the programmer. Due to its generality, the MapReduce computation model has also been an effective paradigm for parallelizing tasks on GPUs [22], multicore systems [31], and traditional HPC clusters [29].

The MapReduce model consists of two elements inspired by functional programming: a *map* operation to transform the input into a key/value pair and a *reduce* operation to process information by using the same key. The user provides both functions, which cannot have any side effects. A MapReduce implementation executes the *map* function on the entire dataset in parallel; see Figure 4. A canonical MapReduce dataset is a TB-sized text file split into individual lines. In this case, each *map* function receives only a few thousand lines from the enormous file, which it processes and sends via *shuffle* to the appropriate *reduce* function. The *shuffle* operation groups *map* outputs by the key, and the *reduce* function processes all outputs with the same key—e.g., counting the number of times a word appears in a large collection of documents.

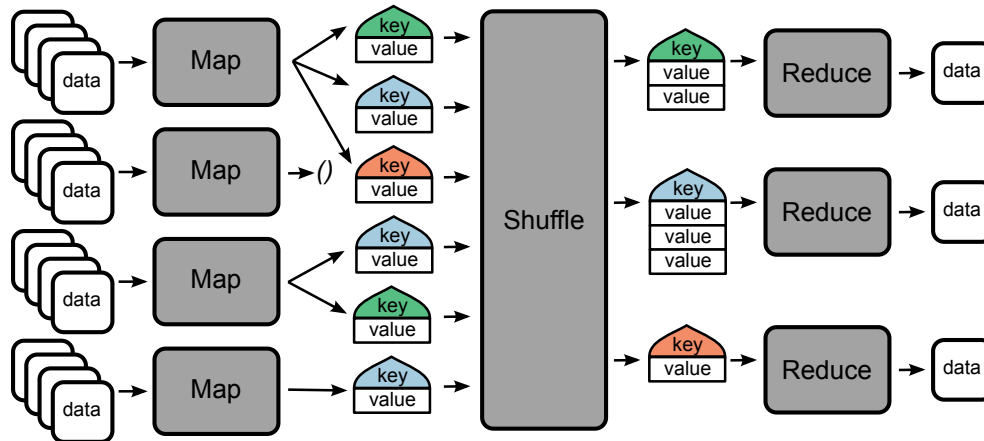


FIG. 4. The MapReduce system by Google is inspired by a functional programming paradigm and consists of three phases: transform data (map), aggregate results (shuffle), and compute and store (reduce). All map and reduce functions are independent, which allows the system to schedule them in parallel.

Google’s implementation of MapReduce is proprietary. An alternative, open source implementation named Hadoop has become the industry standard for large-scale data processing. More information on Hadoop can be found at the Cloudera website [32]. The Hadoop Distributed File System (HDFS) is a fault-tolerant, replicated, block file system designed to run using inexpensive, consumer grade hard disk drives on a large set of nodes.

Map($key = \text{simulation id}$, $value = \text{empty}$)

Read simulation data based on the simulation id. Emit each point in a simulation as a record where the key is the row of the matrix \mathbf{F} —constructed from the spatial location and time step—and the $value$ contains both the column id in the matrix—given by the value of the parameter s —and the value from the simulation data.

Reduce($key = \text{row id}$, $values = \{\text{column id}, F_{ij}\}$)

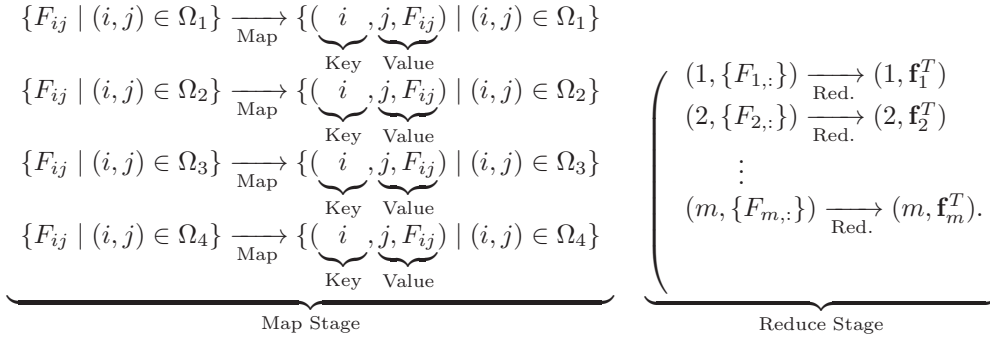
Read all of the values, and emit the combined row as a record where the key is the row id and the $value$ is the array \mathbf{f}_i^T .

FIG. 5. Map and Reduce functions to assemble the matrix \mathbf{F} from simulation data.

3.2. Assembling the matrix from simulation outputs. The first step in the construction of the ROM is to reorganize the data into the tall-and-skinny matrix \mathbf{F} . This step is particularly communication intensive as it requires reorganizing the data from the columns (the natural outputs of the simulations) into rows for the tall-and-skinny SVD routine. To do this in Hadoop we create a text file where each line is a path to a file containing the outputs of one simulation stored in HDFS. Then the *map* function reads simulation data from HDFS and outputs the data keyed on the row of matrix \mathbf{F} . The *reduce* function aggregates all the entries in the row and outputs the realized row. The outcome of this first MapReduce iteration is the matrix \mathbf{F} stored by rows on the distributed file system. More explicit descriptions of the functions are given in Figure 5.

From a matrix perspective, each mapper processes a subset of the entries of \mathbf{F} . For example, assume the disjoint index sets Ω_1 , Ω_2 , Ω_3 , and Ω_4 contain the indices of

F . Then the following diagram shows four mappers processing the simulation data:

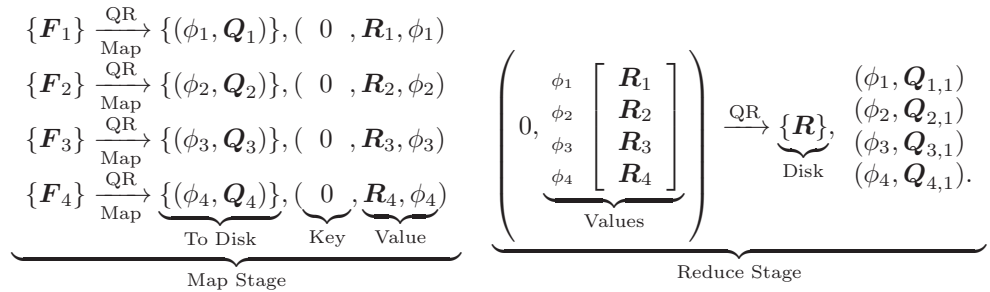


Hadoop assigns the reducers randomly to nodes of the cluster. If we have four nodes, then the output from those four reducers will all be stored together:

$$F = \{\{F_1\}, \{F_2\}, \{F_3\}, \{F_4\}\},$$

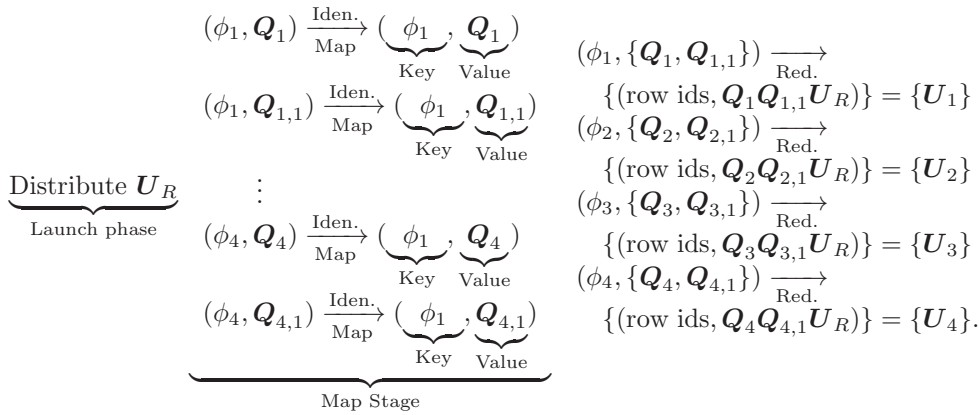
where each F_i is a random subset of rows of the original matrix. Each of these blocks F_i also stores the id's of each row it contains.

3.3. TSQR and SVD in Hadoop. Once we have the matrix F stored by rows on disk, we compute its tall-and-skinny QR (TSQR) factorization [5]. The basis for the MapReduce TSQR algorithm is the communication-avoiding QR factorization [16]. The strategy is to divide the tall-and-skinny matrix F into many smaller, tall matrices to be decomposed independently via the QR factorization. This process is repeated on the new matrix formed by all the R factors computed in the previous step until there is only a single R left. This algorithm was shown to have superior numerical stability compared to a large Householder-style procedure [26]. The HDFS stores the matrix in small chunks according to its internal splitting procedure. Each *map* function reads a small submatrix and computes a QR factorization. To record the association between this QR factorization and all other QR factorizations computed in the Map stage, we create a small tag ϕ that is a universally unique identifier. The *map* function then writes the Q factor back to disk with this small tag ϕ . Finally, it outputs the R factor and the same tag ϕ with key 0. All map functions output their R factor with that same key. Because of this, the outputs all go to the same reducer. This single reducer was not a limitation for our applications, but a recursive procedure in [5] is possible if the *reduce* becomes burdensome on a single node. The diagram below and the description that follows demonstrates the procedure when F is split into four blocks:



There are two types of outputs represented in the diagram above: (i) those surrounded by curly braces $\{\dots\}$ are written to disk for processing in the future and (ii) those

surrounded by parentheses (...) are used in the next stage. The tag ϕ_i uniquely identifies each *map* function. The output from this first MapReduce job is the matrix \mathbf{R} . The identifiers and factors $(\phi_i, \mathbf{Q}_{i,1})$ are fed into the next MapReduce job that computes the matrix \mathbf{U} of left singular vectors. To form \mathbf{U} , we follow the R-SVD procedure described in [11]. We compute the SVD of the small matrix $\mathbf{R} = \mathbf{U}_R \mathbf{\Sigma} \mathbf{V}^T$ on one node and store $\mathbf{\Sigma}$ and \mathbf{V} on disk. With another *map* and *reduce*, we distribute the matrix \mathbf{U}_R to all tasks and combine it with both the input (ϕ_i, \mathbf{Q}_i) and the stored output $(\phi_i, \mathbf{Q}_{i,1})$ from the last stage. The shuffle moves all data with the same key to the same reducer, which uses the tags ϕ_i to align the blocks of \mathbf{Q}_i computed in the first map stage with those output in the first reduce stage. Then in the *reduce*, we read all outputs with the same tag ϕ_i and compute the products $\mathbf{Q}_i \mathbf{Q}_{i,1} \mathbf{U}_R$ to get \mathbf{U}_i . The picture is



This is a numerically stable computation of \mathbf{U} in the SVD of \mathbf{F} stored in HDFS. For more details about the *map* and *reduce* functions, see [5]. The codes for computing the TSQR and SVD can be found at <http://www.github.com/arbenson/mrtsqr>.

3.4. Evaluating the reduced-order model in MapReduce. Next we describe the procedure for evaluating the ROM in MapReduce for a given parameter value s . If one needs to evaluate the ROM at many values of s , this can be done in parallel with our existing codes.

There are two steps involved in evaluating the ROM. The first step is evaluating the interpolated function $\tilde{v}_k(s)$ at s . Since \mathbf{V} is small, this step is executed on a single node. The second step is estimating the ROM prediction and its variance via

$$(3.1) \quad \mathbb{E} \left[\tilde{f}(x_i, s) \right] = \sum_{k=1}^R u_k(x_i) \sigma_k \tilde{v}_k(s), \quad \text{Var} \left[\tilde{f}(x_i, s) \right] = \sum_{k=R+1}^N \sigma_k^2 u_k(x_i)^2.$$

Recall that $R = R(s)$ is the splitting of the singular value expansion at the point s described in section 2.3. Further, recall that the matrix \mathbf{U} computed in the SVD of \mathbf{F} holds the coefficients $u_k(x_i)$. We can evaluate the ROM at all points x_i by computing the matrix-vector product:

$$(3.2) \quad \mathbf{f}(s) = \mathbf{U}(:, 1:R) \tilde{\mathbf{v}}(s),$$

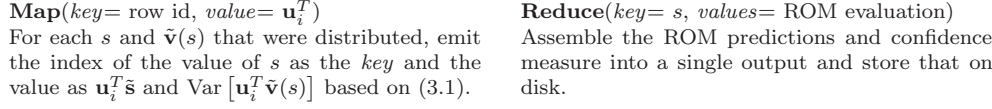


FIG. 6. Map and Reduce functions to compute the interpolants.

where

$$(3.3) \quad \mathbf{f}(s) = \begin{bmatrix} \mathbb{E}[\tilde{f}(x_1, s)] \\ \vdots \\ \mathbb{E}[\tilde{f}(x_M, s)] \end{bmatrix}, \quad \tilde{\mathbf{v}}(s) = \begin{bmatrix} \sigma_1 \tilde{v}_1(s) \\ \vdots \\ \sigma_R \tilde{v}_R(s) \end{bmatrix}.$$

Since the matrices \mathbf{F} and \mathbf{U} are tall, we can compute such matrix-vector products in Hadoop by distributing the small vector $\tilde{\mathbf{v}}$ to a *map* function that will compute a subset of the entries of the matrix-vector product. A subsequent *reduce* collects each submatrix-vector product into a single file; see Figure 6. Viewed schematically for \mathbf{U} stored in four blocks,

$$(3.4) \quad \underbrace{\begin{matrix} \{U_1\} \\ \{U_2\} \\ \{U_3\} \\ \{U_4\} \end{matrix}}_{\text{Distribute } \tilde{\mathbf{v}}(s)} \xrightarrow{\text{Map}} \underbrace{\begin{matrix} (s, \mathbf{f}_1(s), \text{Var}[\mathbf{f}_1(s)]) \\ (s, \mathbf{f}_2(s), \text{Var}[\mathbf{f}_2(s)]) \\ (s, \mathbf{f}_3(s), \text{Var}[\mathbf{f}_3(s)]) \\ (s, \mathbf{f}_4(s), \text{Var}[\mathbf{f}_4(s)]) \end{matrix}}_{\text{Map Stage}}$$

Key
Value

$$(3.5) \quad \underbrace{\begin{pmatrix} \mathbf{f}_1(s), \mathbf{f}_2(s), \mathbf{f}_3(s), \mathbf{f}_4(s) \\ s, \text{Var}[\mathbf{f}_1(s)], \text{Var}[\mathbf{f}_2(s)], \\ \text{Var}[\mathbf{f}_3(s)], \text{Var}[\mathbf{f}_4(s)] \end{pmatrix}}_{\text{Reduce Stage}} \xrightarrow[\text{Reduce}]{\text{Join}} (\mathbf{f}(s), \text{Var}[\mathbf{f}(s)]).$$

Although we illustrate this function with a single interpolation point s , which results in a single reduce with one key, our implementations are designed to handle around one-to-two thousand points s simultaneously. Thus, we actually distribute s and $\tilde{\mathbf{v}}(s)$ for all of the one-to-two thousand points simultaneously. In this case, we would have one reducer for each value s . Our codes for manipulating the simulations and performing the interpolation are found at <http://www.github.com/dgleich/simform> in the branch `simform-sisc`.

4. Numerical experiment. In this section we apply the model reduction method to a parameter study with a large-scale heat transfer model in random heterogeneous media. In what follows, we describe the physical model, the parameter study, and the construction of the ROM. We compare the ROM’s predictions for the quantity of interest with a standard response surface approach. We close this section with some remarks on the computational issues we encountered working with 4TB of data from the simulations. The simulation data set for the following study is available on the Open Science Data Cloud at <https://www.opensciencedatacloud.org/publicdata/heat-transfer/>.




FIG. 7. Representative temperature fields at time $t_f = 2000$ s for bubble radius $s = 0.390$ cm (left) and $s = 2.418$ cm (right). Distortions in the temperature field due to the presence of bubbles are clearly visible in the latter. (Colors are visible in the electronic version.)

4.1. Heat transfer model. We consider a partial differential equation model for bulk conductive heat transfer of a temperature field $T = T(x, t)$:

$$(4.1) \quad \frac{\partial}{\partial t}(\rho c_p T) = \nabla \cdot (\kappa \nabla T), \quad (x, y, z) \in \mathcal{D}, \quad t \in [0, t_f].$$

The spatial domain \mathcal{D} is a rectangular brick of size $20 \times 10 \times 10$ cm³ centered at the origin. The brick contains two materials roughly corresponding to foam ($\rho = 319$ kg/m³) and stainless steel ($\rho = 7900$ kg/m³) chosen for their contrast in thermal conductivity. The values for temperature-dependent specific heat c_p and conductivity κ are shown in Tables 4 and 3 in the appendix. The distribution of the two materials within the brick is realized by the following procedure: (i) set the material in the entire brick to steel, (ii) choose 128 locations within the brick uniformly at random, (iii) for each location, find all points in the brick within a given radius s and set the material to foam. The result is a steel brick with randomly distributed, potentially overlapping foam bubbles of radius s .

A given brick begins ($t = 0$) at room temperature $T = 298$ °K with a Dirichlet boundary condition of 1000 °K on the face $x = 10$ cm. The temperature field advances to the final time $t_f = 2000$ s, and the quantities of interest are measured at the face opposite the prescribed boundary condition (the *far face*) at $x = -10$ cm. We are interested in two quantities: (i) the average temperature on the far face and (ii) the proportion of the temperature on the far face that exceeds 475 °K.

The finite element simulation uses Sandia Labs' SIERRA Multimechanics Module: Aria [27] with a regular mesh of $256 \times 128 \times 128$ elements constructed with CUBIT [6]. One simulation takes approximately four hours on an 8-core node of Sandia's Red Sky capacity cluster (Dual Xeon 5500-series, 2.93 GHz, 2GB per core). The simulation outputs contain the full temperature field at nine uniformly spaced times between $t = 1200$ s and $t_f = 2000$ s, which are stored in the Exodus II binary file format. Each simulation output file is approximately 500MB. Two representative temperature fields are shown in Figure 7.

4.2. Parameter study. We use the heat transfer model to study the effects of the bubble radius parameter s on the temperature distribution on the far face via the quantities of interest. Intuitively, as s increases, more of the brick becomes foam, and we expect a lower temperature on the far face $x = -10$ cm due to foam's lower conductivity.




FIG. 8. Monte Carlo estimates of the mean of the average far face temperature as a function of the bubble radius (left). The mean of the proportion of the temperature that exceeds 475°K as a function of the bubble radius (right). (Colors are visible in the electronic version.)

To address the variability in random media, we choose 128 random realizations of the 128 locations for the bubble centers. For a given radius s , we run 128 simulations—one for each realization of the bubble locations. We use these simulations to compute Monte Carlo estimates of the mean of each quantity of interest. Note that there are 384 random variables (three components per location) characterizing the locations of the bubbles, so Monte Carlo estimates are the only feasible option.

For each realization of the bubble locations, we run 64 simulations varying s uniformly between 0.039 cm and 2.496 cm. This results in a total of 8192 simulations—128 bubble realizations \times 64 values for s . For each simulation, we compute the two quantities of interest: (i) the average temperature over the far face and (ii) the proportion of the far face temperature that exceeds 475°K . We then approximate the mean over the bubble realizations with Monte Carlo. Finally, we plot the estimates of the mean as a function of the radius s . With 128 realizations, the 95% confidence intervals are within 1% of the mean, so we do not plot them. These results are shown in Figure 8. As expected, the propagation of the temperature decreases as s increases. However, the decrease is qualitatively different for the two quantities of interest: the average far face temperature decreases smoothly as a function of s while the proportion of the temperature above a threshold has a dramatic change as a function of s . In the next section, we test the ability of the reduced-order model to reproduce these results.

4.3. Approximating the SVE. Before testing the ROM, we use all of the simulation data to study the components of the SVE of the parameterized temperature field. In the notation of section 2, $f = f(x, s)$ now carries the following interpretation: f is the temperature; x contains the three spatial coordinates, the time coordinate, and an identifier for the realization of the random bubble locations; and s is the bubble radius. The matrix \mathbf{F} from (2.4) has 64 columns corresponding to the 64 values of the bubble radius. Each column contains 128 independent simulations—one for each realization of the bubble locations. This translates to 4926801024 rows (257 x nodes \times 129 y nodes \times 129 z nodes \times 9 times \times 128 bubble locations), which is approximately 2.3TB of data.

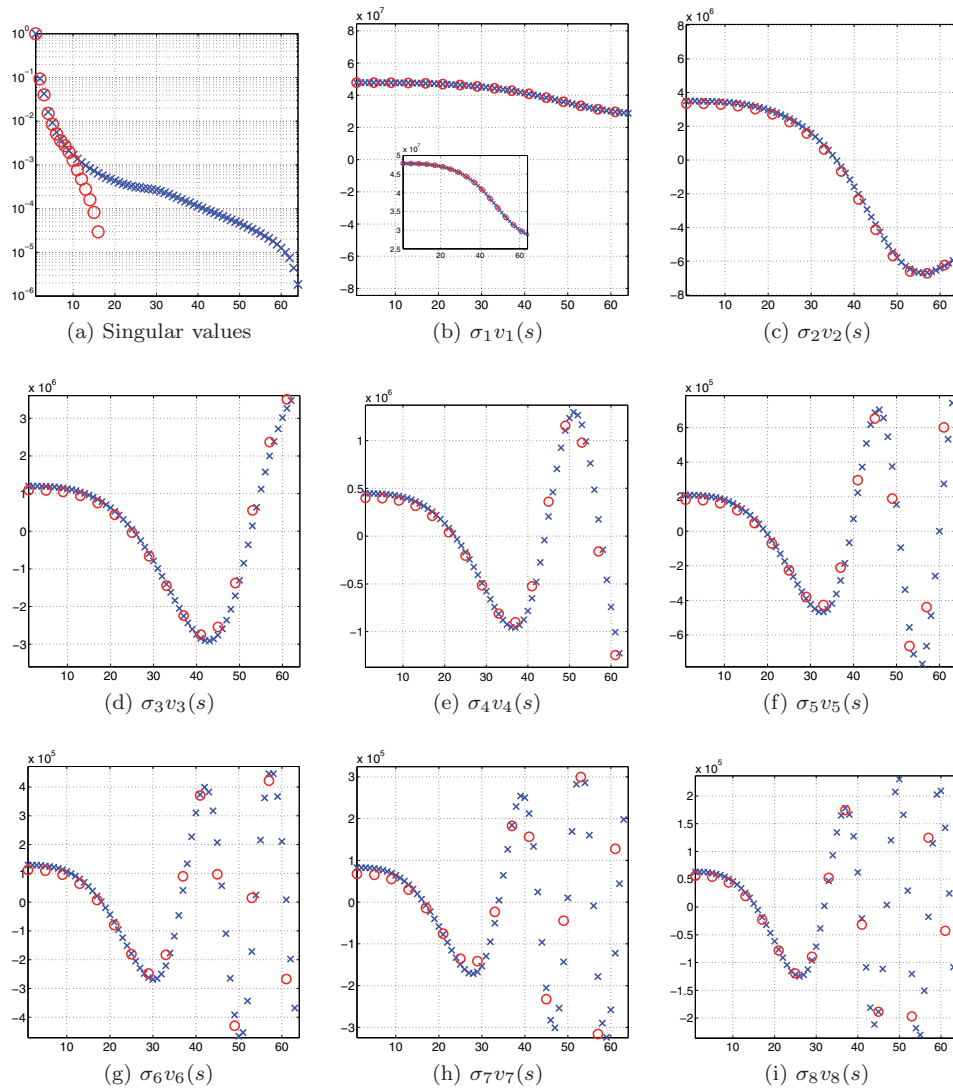


FIG. 9. The top left figure shows the singular values of a finely resolved \mathbf{F} with 64 columns (blue x 's) and a coarsely resolved \mathbf{F} with 16 columns (red o 's). Each are scaled by the maximum singular value from each set. The remaining figures show the approximations of the singular functions $v_1(s)$ through $v_8(s)$ scaled by the respective singular values for \mathbf{F} with 64 columns (blue x 's) and 16 columns (red o 's). (Colors are visible in the electronic version.)

The singular values normalized by the largest singular value and the first eight right singular vectors scaled by the singular values are shown in Figure 9 with the blue x 's. The rapid decay of the singular values indicates the tremendous correlation among components of the temperature fields as the radius varies. More importantly, we see the more rapid increase in the oscillations of the right singular vectors for larger values of s . This reflects the fact that temperature fields with similar bubble radii have larger differences for larger values of the radius s . This also means we expect a more accurate ROM for smaller values of s .

We computed this SVD and all other MapReduce-based analyses of this dataset

on a 10-node Hadoop cluster at Stanford University’s Institute for Computational and Mathematical Engineering. Each node in the cluster has six 2TB hard drives, one Intel Core i7-960, and 24GB of RAM. The 2.3TB matrix took approximately 12 hours. Below, we discuss the time required for additional pre- and post-processing work.

4.4. Construction and validation of the ROM. We use a subset of the simulations as the training set for the ROM. In particular, we choose $s_j = 0.039j$ for $j = 1, 5, 9, \dots, 61$ as the values of s whose corresponding simulations are used to build the ROM. Thus, the matrix \mathbf{F} for constructing the ROM contains 16 columns and the same roughly five billion rows, i.e., the same locations in the domain of space, time, and the locations of the bubble centers. This matrix contains approximately 600GB of data, and the SVD step took approximately eight hours on the Hadoop cluster. The singular values normalized by the largest singular value and the components of the first eight right singular vectors scaled by their respective singular values of \mathbf{F} are plotted in Figure 9 with red o’s.

To choose the threshold $\bar{\tau}$ that defines the splitting described in section 2.3, we use a subset of the simulations as a testing set. In particular, we choose $s_j = 0.039j$ with $j = 3, 7, 11, \dots, 59$ as the values of s whose simulations we will use for testing. Note that these correspond to the midpoints of the intervals defined by the values of s used for training. Figure 10 shows the errors as a function of the bubble radius s and the variation threshold $\bar{\tau}$. After $\bar{\tau} = 0.55$, the approximation does not improve with more terms (i.e., larger R in (2.14)), so we choose $\bar{\tau} = 0.55$ since we want a ROM with the fewest terms. Table 2 displays the splitting R and the associated error \mathcal{E} for the different values of s using variation threshold $\bar{\tau}$.

Finally, we visually compare the error in the ROM with the space-time varying confidence measure. Figure 11 displays the ROM error and the confidence measure at the final time t_f for one realization of the bubble locations and two values of the bubble radius, $s = 0.39$ cm and $s = 1.95$ cm. Both measures are larger near the bubble boundaries and larger near the face containing the heat source. Visualizing these measures enables such qualitative observations and comparisons.

4.5. Comparison with a response surface. One question arises frequently in the context of reduced-order modeling: if one is only interested in a scalar quantity of interest from the full PDE solution, then what is the advantage of approximating the full solution with a reduced-order model? Why not just use a scalar response surface to approximate the quantity of interest as a function of the parameters? To address this question, we compare two approaches for the parameter study in section 4.2:

1. Use a response surface to interpolate the means of each of the two quantities of interest over a range of bubble radii. We use the quantities of interest at bubble radii $s_j = 0.039j$ for $j = 3, 7, 11, \dots, 59$ to decide the form of the response surface: piecewise linear, nearest neighbor, cubic spline, or piecewise cubic Hermite interpolation (PCHIP). The response surface form with the lowest testing error is constructed from the mean quantities of interest for bubble radii $s_j = 0.039j$ for $j = 1, 5, 9, \dots, 61$ —the same values whose simulations are used to construct the ROM. The response surface prediction is then computed for $j = 1, 2, 3, \dots, 61$.
2. Use the ROM to approximate the temperature field on the far face at the final time for each realization of the bubble location. Then compute the two quantities of interest for each approximated far face temperature distribution, and compute a Monte Carlo approximation of the mean (i.e., a simple average).




FIG. 10. The log of the relative error in the mean prediction of the ROM as a function of s and the threshold $\bar{\tau}$. (Colors are visible in the electronic version.)

TABLE 2
The split and the corresponding ROM error for $\bar{\tau} = 0.55$ and different values of s .

s	$R(s, \bar{\tau})$	$\mathcal{E}(s, \bar{\tau})$
0.08	16	1.00e-04
0.23	15	2.00e-04
0.39	14	4.00e-04
0.55	13	6.00e-04
0.70	13	8.00e-04
0.86	12	1.10e-03
1.01	11	1.50e-03
1.17	10	2.10e-03
1.33	9	3.10e-03
1.48	8	4.50e-03
1.64	8	6.50e-03
1.79	7	8.20e-03
1.95	7	1.07e-02
2.11	6	1.23e-02
2.26	6	1.39e-02

The results of this study are shown in Figure 12. For the first quantity of interest (the average temperature over the far face), the cubic spline response surface approach adequately captures the behavior as a function of the bubble radius due to the relative smoothness of the response. However, the PCHIP response surface approximation of the second quantity of interest (the proportion of far face temperature that exceeds

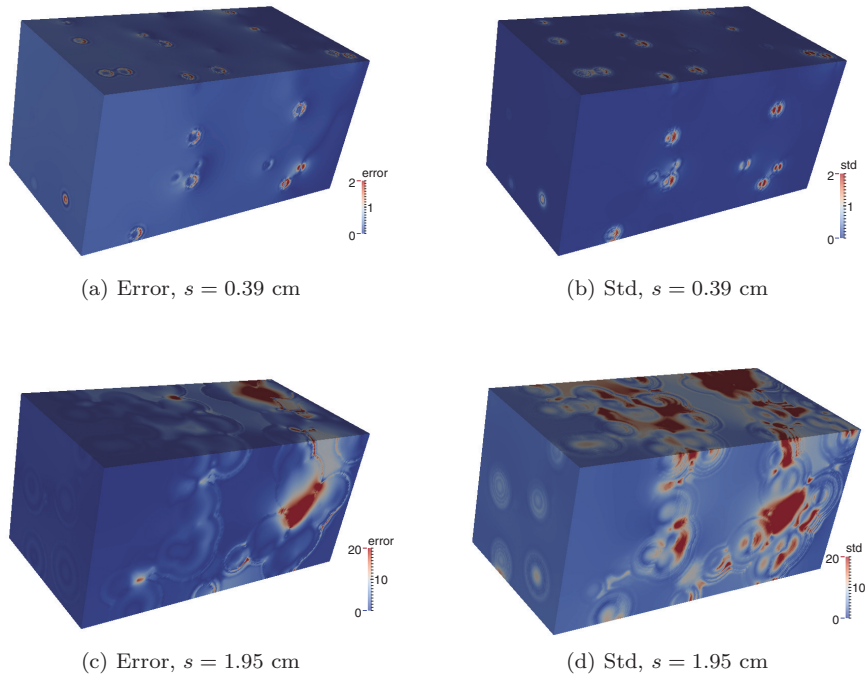


FIG. 11. Absolute error in the reduce order model compared with the prediction standard deviation for one realization of the bubble locations at the final time for two values of the bubble radius, $s = 0.39$ and $s = 1.95$ cm. (Colors are visible in the electronic version.)

475 °K) is substantially less accurate than the ROM due to the sharp transition and the low resolution in the parameter space. A global polynomial or radial basis function approximation would fare worse (e.g., exhibit Gibbs oscillations near the transition) due to the global nature of the basis. We conclude that the choice of ROM versus response surface depends on the quantity of interest. Broadly speaking, if the quantity of interest is a smooth function of the parameters, then a response surface is likely sufficient. However, if the quantity of interest is a highly nonlinear or discontinuous function of the full PDE solution, then computing such a quantity from the ROM approximation will yield better results. While it was not pertinent to this example, the ROM-based approach also provides the 95% confidence bounds for the Monte Carlo estimates, since the temperature distribution for each realization of the bubble locations is approximated.

4.6. Computational considerations. We end this section with a few notes on the experience of running 8192 large-scale simulations, transferring them to a Hadoop cluster, and building the reduced-order model.

Each heat transfer simulation took approximately four hours using eight processors on Sandia's Red Sky. Communication times were negligible, but some runs required multiple tries due to occasional network failures. Also, some runs had to be duplicated due to node failures and changes in the code paths. Each mesh with its own conductivity field took approximately twenty minutes to construct using Cubit after substantial optimizations. Unreliable network data transmissions and bursty data write patterns (e.g., one hundred jobs on Red Sky simultaneously transferring data to the Stanford cluster) forced us to write custom codes to validate the data transfer.

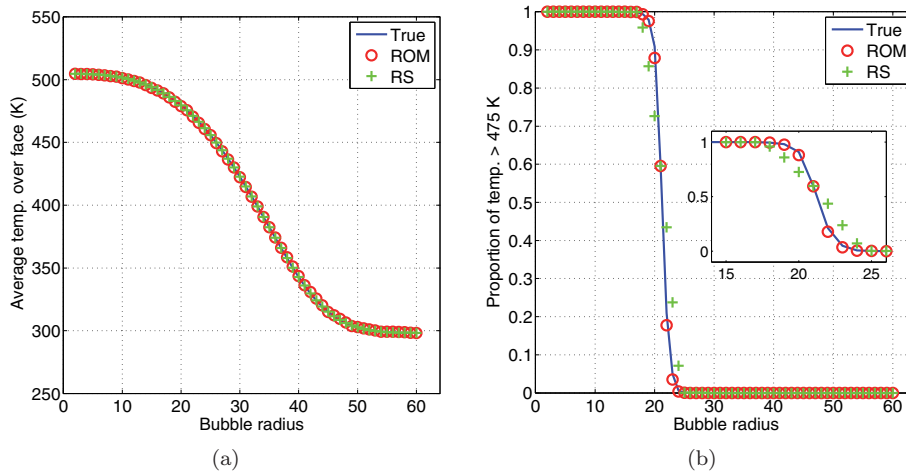


FIG. 12. Comparing ROM with response surface in the UQ study. (Colors are visible in the electronic version.)

Working with the simulation data involved a few pre- and postprocessing steps, such as interpreting 4TB of Exodus II files from Aria. The preprocessing steps took approximately 8–15 hours. We collected precise timing information, but we do not report it as these times are from a multitenant, unoptimized Hadoop cluster where other jobs with sizes ranging between 100GB and 2TB of data sometimes ran concurrently. Also, during our computations, we observed failures in hard disk drives and issues causing entire nodes to fail. Given that the cluster has 40 cores, these calculations consumed at most 2400 cpu-hours—compared with the 262144 cpu-hours it took to compute 8192 heat transfer simulations on Red Sky. Thus, evaluating the ROM was about 100 times faster than computing a full simulation.

We did not compare our Hadoop implementation with an MPI implementation. The dominant bottleneck in the evaluation of the ROM is the data I/O involved in processing 4TB of simulation data into a 2.3TB matrix, computing its SVD, writing it to disk, computing the interpolants, and writing those outputs back to Exodus II files. We expect that any MPI implementation would take at least 3–4 hours based on pure I/O considerations (assuming around 1GB/sec sustained data transfer speeds). It would also be substantially more complicated to implement. We used Hadoop primarily for the ease of use of its programming model.

Although MapReduce is an appealing paradigm for computing factorizations of tall matrices, the Hadoop MapReduce ecosystem has not developed simple tools for working with large databases of spatio-temporal data. For instance, writing ad hoc utilities to extract data from Exodus II files and utilities for simply stated queries like *retrieve all values of temperature with $x = -0.1$ and $t = 2000$* occupied much of the second and third authors' time (see the `mr-exodus2farface.py` code in the online code for this particular script). We see many opportunities for high-impact software designed to address the mundane details of manipulating and analyzing large databases of spatio-temporal simulation data.

5. Summary and conclusions. We presented a method for building a reduced-order model of the solution of a parameterized PDE. The method is based on approximating the factors of a SVE of the solution using the elements of an SVD of a

matrix whose columns are spatio-temporally discretized solutions at different parameter values. The SVD step compares to reduced basis methods, which project the governing equations with a subset of the left singular vectors to create a small system whose solution yields the coefficients of the reduced-order model. In contrast, our method interpolates the right singular vectors in the parameter space to compute the coefficients of the ROM. By examining the gradient of the right singular vectors as the index increases, we determine a separation between factors we can accurately interpolate and those whose oscillations are too rapid to be represented on the parameter grid. This separation yields a mean prediction and a prediction covariance for each point in the spatio-temporal domain—similar to Gaussian process regression models. We use Hadoop/MapReduce to implement the ROM including the communication-avoiding, tall-and-skinny SVD, which enables the computation to scale to outputs from large-scale high-fidelity models.

We tested the model reduction method on a parameter study of large-scale heat transfer in random media. We compared the results of the ROM with a standard response surface method for approximating the scalar quantities of interest, and we found that while the cheaper response surface was appropriate for a smooth quantity of interest, the ROM was better at approximating a quantity of interest with a sharp transition in the parameter space. The 8192 heat transfer simulations used in the study generated approximately 4TB of data. In the course of the study, we applied the MapReduce-based SVD computation to matrices with approximately 600GB and 2.2TB of data. We found that existing MapReduce tools for working with such large-scale simulation data lack robustness and generality. There is an opportunity in computational science to create better tools to further simulation-based scientific exploration.

6. Appendix.

Tables 3 and 4 contain material properties for foam and steel.

TABLE 3

Thermal conductivity κ . Linear interpolation is used between specified values. Constant values are used outside of the bounds.

Foam ($\rho = 319 \text{ kg/m}^3$)		Steel ($\rho = 7900 \text{ kg/m}^3$)	
T (K)	κ (W/mK)	T (K)	κ (W/mK)
303	0.0486	273	13.4
523	0.0706	373	16.3
		773	21.8
		973	26.0

TABLE 4

Specific heat capacity c_p . Linear interpolation is used between specified values. Constant values are used outside of the bounds.

Foam ($\rho = 319 \text{ kg/m}^3$)		Steel ($\rho = 7900 \text{ kg/m}^3$)	
T (K)	c_p (J/kgK)	T (K)	c_p (J/kgK)
296	1269	273	502
323	1356	673	565
373	1497		
423	1843		
473	1900		
523	2203		

Acknowledgments. We thank the anonymous reviewers for helpful comments and suggestions. We also thank Margot Gerritsen at Stanford’s Institute for Computational and Mathematical Engineering for procurement of and access to the Hadoop cluster. We thank Austin Benson at Stanford for his superb code development for the TSQR and TSSVD. We thank Joe Ruthruff at Sandia for his efforts in developing the infrastructure to run the Aria cases. Finally, we thank David Rogers at Sandia. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

REFERENCES

- [1] C. AUDOUZE, F. DE VUYST, AND P. B. NAIR, *Reduced-order modeling of parameterized PDEs using time-space-parameter principal component analysis*, Internat. J. Numer. Methods Engrg., 80 (2009), pp. 1025–1057.
- [2] C. AUDOUZE, F. DE VUYST, AND P. B. NAIR, *Nonintrusive reduced-order modeling of parametrized time-dependent partial differential equations*, Numer. Methods Partial Differential Equations, 29 (2013), pp. 1587–1628.
- [3] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., 45 (2007), pp. 1005–1034.
- [4] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A Survey of Model Reduction Methods for Parametric Systems*, Tech. report, Max Planck Institute, Magdeburg, Germany, 2013.
- [5] A. BENSON, D. F. GLEICH, AND J. DEMMEL, *Direct tall-and-skinny QR factorizations in MapReduce architectures*, in Proceedings of the First IEEE Conference on Big Data, 2013, pp. 264–272.
- [6] T. D. BLACKER, W. J. BOHNHOFF, AND T. L. EDWARDS, *CUBIT Mesh Generation Environment. Volume 1: Users Manual*, Technical report, Sandia National Labs., Albuquerque, NM, 1994.
- [7] L. S. BLACKFORD, *ScaLAPACK User’s Guide*, Vol. 4, SIAM, Philadelphia, 1997.
- [8] T. BUI-THANH, K. WILLCOX, AND M. DAMODARAN, *Applications of Proper Orthogonal Decomposition for Inviscid Transonic Aerodynamics*, Technical report, MIT, Cambridge, MA, 2003, <http://hdl.handle.net/1721.1/3694>.
- [9] T. BUI-THANH, K. WILLCOX, AND O. GHATTAS, *Model reduction for large-scale systems with high-dimensional parametric input space*, SIAM J. Sci. Comput., 30 (2008), pp. 3270–3288.
- [10] K. CARLBERG, C. FARHAT, J. CORTIAL, AND D. AMSALLEM, *The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows*, J. Comput. Phys., 242 (2013), pp. 623–647.
- [11] T. F. CHAN, *An improved algorithm for computing the singular value decomposition*, ACM Trans. Math. Software, 8 (1982), pp. 72–83.
- [12] CLUDERA, *Hadoop version 0.20.2 in Cludera Hadoop Distribution Version CDH3U4*, available online at <http://www.cludera.com>, 2012.
- [13] P. G. CONSTANTINE AND D. F. GLEICH, *Tall and skinny QR factorizations in MapReduce architectures*, in Proceedings of the Second International Workshop on MapReduce and Its Applications, ACM, New York, 2011, pp. 43–50.
- [14] P. G. CONSTANTINE AND Q. WANG, *Residual minimizing model interpolation for parameterized nonlinear dynamical systems*, SIAM J. Sci. Comput., 34 (2012), pp. A2118–A2144.
- [15] J. DEAN AND S. GHEMAWAT, *MapReduce: Simplified data processing on large clusters*, in Proceedings of the 6th Symposium on Operating Systems Design and Implementation (OSDI2004), 2004, pp. 137–150.
- [16] J. DEMMEL, L. GRIGORI, M. HOEMMEN, AND J. LANGOU, *Communication-optimal parallel and sequential QR and LU factorizations*, SIAM J. Sci. Comput., 34 (2012), pp. A206–A239.
- [17] M. ELDRDRE, A. GIUNTA, AND S. COLLIS, *Second-order corrections for surrogate-based optimization with model hierarchies*, in Proceedings of the 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, American Institute of Aeronautics and Astronautics, 2004.
- [18] J. GOSS AND K. SUBBARAO, *Inlet shape optimization based on POD model reduction of the Euler equations*, in Proceedings of the 12th Annual AIAA ISSMO Multidisciplinary Analysis and Optimization Conference, AIAA 2008–5809, 2008.

- [19] P. C. HANSEN, *Computation of the singular value expansion*, Computing, 40 (1988), pp. 185–199.
- [20] P. C. HANSEN, *Discrete inverse problems: Insight and algorithms*, Fundam. Algorithms 7, SIAM, Philadelphia, 2010.
- [21] P. C. HANSEN, M. E. KILMER, AND R. H. KJELDSSEN, *Exploiting residual information in the parameter choice for discrete ill-posed problems*, BIT, 46 (2006), pp. 41–59.
- [22] B. HE, W. FANG, Q. LUO, N. K. GOVINDARAJU, AND T. WANG, *Mars: A MapReduce framework on graphics processors*, in Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques, PACT '08, ACM, New York, 2008, pp. 260–269.
- [23] K. LEE, T. NAM, C. PERULLO, AND D. N. MAVRIS, *Reduced-order modeling of a high-fidelity propulsion system simulation*, AIAA J., 49 (2011), pp. 1665–1682.
- [24] J. L. LUMLEY, G. BERKOOZ, AND C. W. ROWLEY, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, UK, 2012.
- [25] H. V. LY AND H. T. TRAN, *Modeling and control of physical processes using proper orthogonal decomposition*, Math. Comput. Modelling, 33 (2001), pp. 223–236.
- [26] D. MORI, Y. YAMAMOTO, AND S.-L. ZHANG, *Backward error analysis of the AllReduce algorithm for Householder QR decomposition*, Jpn. J. Ind. Appl. Math., 29 (2012), pp. 111–130.
- [27] P. K. NOTZ, S. R. SUBIA, M. M. HOPKINS, H. K. MOFFAT, AND D. R. NOBLE, *Aria 1.5: User Manual*, Technical Report SAND2007-2734, Sandia National Laboratories, Albuquerque, NM and Livermore, CA, 2007.
- [28] Z. OSTROWSKI, R. A. BIALECKI, AND A. J. KASSAB, *Estimation of constant thermal conductivity by use of proper orthogonal decomposition*, Comput. Mech., 37 (2005), pp. 52–59.
- [29] S. J. PLIMPTON AND K. D. DEVINE, *MapReduce in MPI for large-scale graph algorithms*, Parallel Comput., 37 (2011), pp. 610–632.
- [30] C. E. RASMUSSEN AND C. WILLIAMS, *Gaussian processes for machine learning*, MIT Press, Cambridge, MA, 2006.
- [31] J. TALBOT, R. M. YOO, AND C. KOZYRAKIS, *Phoenix++: modular MapReduce for shared-memory systems*, in Proceedings of the Second International Workshop on MapReduce and its Applications, MapReduce '11, ACM, New York, 2011, pp. 9–16.
- [32] VARIOUS, *Hadoop version 0.20*, Cloudera CDH3, <http://hadoop.apache.org>, <http://cloudera.com>, 2010.
- [33] D. XIU AND J. S. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., 27 (2005), pp. 1118–1139.