THE SHIFTED PROPER ORTHOGONAL DECOMPOSITION: A MODE DECOMPOSITION FOR MULTIPLE TRANSPORT PHENOMENA*

J. REISS[†], P. SCHULZE[‡], J. SESTERHENN[†], AND V. MEHRMANN[‡]

Abstract. Transport-dominated phenomena provide a challenge for common mode-based model reduction approaches. We present a model reduction method, which is suited for these kinds of systems. It extends the proper orthogonal decomposition (POD) by introducing time-dependent shifts of the snapshot matrix. The approach, called shifted proper orthogonal decomposition (sPOD), features a determination of the *multiple* transport velocities and a separation of these. One- and two-dimensional test examples reveal the good performance of the sPOD for transport-dominated phenomena and its superiority in comparison to the POD.

Key words. transport phenomena, model reduction, mode decomposition, shifted proper orthogonal decomposition, proper orthogonal decomposition, principal component analysis

AMS subject classifications. 65Z05, 35Q35

DOI. 10.1137/17M1140571

1. Introduction. Model reduction is an important method for deriving low-dimensional models from experimental or numerical data. These reduced-order models often allow fast simulations commonly used for control, optimization, and parameter studies and are nowadays a necessary tool in many fields. Further, these models yield a better understanding of the dynamical process by identifying the essential dynamics. Formally, the goal is to obtain a low-dimensional description, which approximates the mapping from a set of inputs to a set of outputs. Among others, inputs can be design parameters, system conditions, or controls. Common outputs are performance or physical quantities like measurements or even full flow solutions.

A major class of reduced-order modeling approaches is given by input-output interpolation methods, which do not aim to reduce the internal dynamics but which are only based on the input-output behavior, usually described by a transfer function (see [2, 5, 21, 29] and references therein). These methods are successfully applied to linear systems, but extensions to nonlinear systems are rare and have some drawbacks, such as low computational efficiency (cf. [18]).

In contrast, the most common model reduction techniques for nonlinear systems are based on a superposition of modes describing the system state. Examples are reduced basis methods [15, 16, 20, 38, 44], balanced proper orthogonal decomposition (POD) [43, 46], dynamic mode decomposition [12, 39, 42], and POD [23, 32], which technically reduces to a singular value decomposition (SVD). The SVD provides the best low-rank approximation of a matrix with respect to the 2-norm. Depending on

^{*}Submitted to the journal's Methods and Algorithms for Scientific Computing section July 25, 2017; accepted for publication (in revised form) March 14, 2018; published electronically May 8, 2018.

http://www.siam.org/journals/sisc/40-3/M114057.html

Funding: This work was supported by the Deutsche Forschungsgemeinschaft (DFG) as part of the collaborative research center SFB 1029, Substantial Efficiency Increase in Gas Turbines Through Direct Use of Coupled Unsteady Combustion and Flow Dynamics, Project A02, Development of Reduced-Order Models for Pulsed Combustion.

[†]Institut für Strömungsmechanik und Technische Akustik, TU Berlin, Berlin, Germany (reiss@tnt.tu-berlin.de, sesterhenn@tnt.tu-berlin.de).

 $^{^{\}ddagger}$ Institut für Mathematik, TU Berlin, Berlin, Germany (pschulze@math.tu-berlin.de, mehrmann@math.tu-berlin.de).

the application, the POD has several synonyms, as for instance, principal component analysis [25] or Karhunen–Loeve decomposition [14]. The number of modes needed to obtain a sufficiently accurate approximation of the system in question is crucial for the computational cost of evaluating the reduced-order model (online stage) and accordingly for its usefulness. Furthermore, nonlinear reduced-order models are often combined with hyperreduction methods, e.g., EIM [4] or DEIM [11], to achieve an efficient offline/online decomposition [22].

Transport-dominated phenomena are usually a challenge for mode-based methods, since their dynamical behavior cannot be captured accurately by a few spatial modes in a dyadic structure (cf. [13, 28]). Recently, there have been some efforts to invent model reduction methods suitable for the efficient description of transport-dominated phenomena. Usually some time-dependent shift is introduced to compensate the transport. It is used in the framework of symmetry reduction (cf. [6, 17, 31, 37]), where the translation is accounted for by applying a Lie group action to a symmetry-reduced or frozen solution. In [31], for instance, the framework has been analyzed for nonlinear parameter-dependent evolution equations and applied to a numerical simulation of the Burgers equation. A different approach is presented in [24], where the transport is incorporated by using a coordinate mapping which is related to the solution of Monge–Kantorovich optimal mass transport problems. The approach is illustrated by means of snapshots of shallow water waves and of a hurricane. A methodology based on L^1 -norm minimization has been applied in [1], which shows much better results for hyperbolic problems than the commonly applied L_{2} norm minimization of the error. The minimization is based on a set of dictionaries which are computed in an offline phase. There are various other approaches which aim to efficiently reduce transport phenomena; see, for instance, [9, 10, 19, 30, 40, 41, 45].

Multiple transport velocities are less studied. Just one shift as is used in most of the cited works is not sufficient if different transport velocities are present, as is common in technical applications. An efficient and general model reduction methodology for *multiple* transport phenomena is still missing. With this contribution we aim at improving this situation by introducing the shifted proper orthogonal decomposition (sPOD).

The key tool of the new approach is a time-dependent shift in combination with a procedure to *separate* different transports within the system. The dominant transport velocities are determined by front tracking or by considering the dependence of the singular values of the shifted snapshot matrix on the time-dependent shift. This shift structurally extends the dyadic structure of POD to systems with transport and allows thereby a better approximation. The method is purely data based in contrast to the symmetry reduction methods, offering a wide range of applicability.

In this paper, we focus on obtaining a very low-dimensional representation of the solution. The corresponding low-dimensional subspace can be used to build a reduced-order model, for instance, by an interpolation procedure as in [26] or by a Galerkin projection of the original model. However, the construction of a reduced-order model is not within the scope of this paper.

During the review process we became aware of the work [36], which is similar to this work using SVDs in different velocity frames. The authors of [36] utilize a greedy approach in order to decompose snapshot matrices with multiple transport velocities. They consider the linear wave equation with two transported quantities and their method was shown to yield a decomposition with just a few modes outperforming the POD. However, the greedy approach is not able to describe this linear transport with the minimum number of modes indicated by the analytic solution. This is because structures which are extracted by the first greedy iteration cannot be reattributed to a different velocity frame, which is necessary to obtain low-rank approximations for hyperbolic cases with multiple transports. In contrast, this flexibility is given by the new sPOD algorithm which is able to decompose the same example with the minimum number of modes and thus finds the analytic solution up to any given achievable tolerance (cf. section 2.2).

In the following section, we motivate and derive the method by means of examples in a one-dimensional domain, which are a challenge for common model reduction techniques as the POD. For this, first we consider one moving signal and show how to reduce it. Then, a more complex system with two different transports is considered and a procedure is derived to separate the different velocity components. Then the detection of the velocities is discussed and a case of crossing shocks is investigated. In section 3 we apply the new method to a two-dimensional test case from fluid mechanics with transported developing vortex pairs with nontrivial velocities. Finally, we conclude and give an outlook on ongoing and future work.

2. One-dimensional model problems. In this section, first the idea of the sPOD approach will be developed considering a one-dimensional example. As a test problem, the linear wave equation

~

(1)
$$\begin{aligned} \partial_t \rho + \rho_0 \partial_x u &= 0, \\ \partial_t u + c^2 / \rho_0 \partial_x \rho &= 0, \end{aligned}$$

is considered on a periodic domain x = (0, L]. Here, u is the velocity, ρ is the density (fluctuation), ρ_0 is a reference density, and c is the speed of sound. The general solution can be written as

(2)
$$q(x,t) = \begin{bmatrix} \rho(x,t) \\ u(x,t) \end{bmatrix} = q_+(x-ct) \begin{bmatrix} \rho_0 \\ c \end{bmatrix} + q_-(x-(-c)t) \begin{bmatrix} \rho_0 \\ -c \end{bmatrix}$$

with arbitrary initial conditions $q_{\pm}(x)$ for the two transported quantities, the Riemann invariants. The corresponding transport velocities are $c_{\pm} = \pm c$. By choosing

(3)
$$q_{\pm}(x - (\pm ct)) = \frac{1}{2} \sum_{n=0}^{\infty} \beta_n \cos(k_n (x - (\pm ct)) + \Theta_n)$$

with $k_n = n 2\pi/L$, the solution can be rewritten as

(4)
$$q(x,t) = \sum_{n=0}^{\infty} \beta_n \begin{bmatrix} \rho_0 \cos(k_n \, ct) \cos(k_n \, x + \Theta_n) \\ c \sin(k_n \, ct) \sin(k_n \, x + \Theta_n) \end{bmatrix}$$

with constants β_n and Θ_n . Every nonzero amplitude β_n yields a mode of a vibrating string. For the remainder of this paper, we set $\rho_0 = 1$, c = 1, and L = 1.

As stated in section 1, model reduction methods often build on describing a dynamical system by a superposition of a small number of modes. One of the most popular approaches is the POD, which aims at approximating the solution by a linear combination of orthonormal modes ϕ_l

(5)
$$q(x_i, t_n) \approx \sum_l \alpha_l(t_n) \phi_l(x_i)$$

with time-dependent coefficients α_l . Usually, a snapshot matrix $X_{ij} = q(x_i, t_j)$ is introduced, which is a space-time-discretized solution. A low-dimensional representation minimizing the approximation error in the 2-norm is determined via an SVD of the snapshot matrix X. Often the numerical affort of the SVD is decreased by reducing the set of discrete time and spatial points [3] in a way that does not significantly effect the quality of the modes, i.e., the dynamical behavior of the system still needs to be captured. The squared singular values determine the mean square amplitude of the corresponding modes, which coincides with the kinetic energy if the mode represents the velocity [27]. Consequently, strongly decaying singular values allow a good representation with a few modes.

The ability of the POD to describe the solution of (1) eminently depends on the initial condition. If a vibrating string with a dominant frequency and a few harmonics is chosen, i.e., as in (4) with a few dominant β_n , the POD will find these few modes and, consequently, it will deliver an accurate low-dimensional representation. However, if a transported quantity with high gradient is given, i.e., as in (2) with strongly localized initial conditions q_{\pm} , the singular value decay is rather gradual and one is forced to use a high number of modes to get a reasonable description. This variability makes the example well-suited to develop a method, which can handle transported quantities while including the classical POD as a special case.

2.1. One transported quantity. Setting $q_{-}(x) = 0$ in the solution (2) of the linear wave equation (1) leads to a single transported quantity, which is given by a shift of the initial condition q_{+} with the transport velocity c. To provide a challenging case for POD, a sharp Gaussian pulse of $q_{+}(x) = \exp(-(x-x_{0})^{2}/\delta^{2})$ is chosen with $\delta = L/50$. The analytic solution is shown in Figure 1, left, for the time interval [0, 1.25L/c]. For the discrete snapshot matrix $X_{ij} = q(x_{i}, t_{j})$, we chose 250 equidistant time and 200 space points. In this subsection we only consider the density, if not stated otherwise.

For the considered example, a very slow decay of the singular values is observed (cf. Figure 1, left). Even though the analytic solution can be formulated by only one



FIG. 1. Single transported quantity. Left: Singular value decay (outer picture) and density in the space-time diagram (inner picture). The spectrum shows a very slow decay. Right: If the solution is described in a co-moving frame, i.e., the solution is shifted at every time to compensate the transport, the singular spectrum shows a rapid decay. One mode is sufficient to describe the flow (apart from a small numerical error of the shift procedure).

A1326 REISS, SCHULZE, SESTERHENN, AND MEHRMANN

transported Gaussian pulse, many POD modes are needed for a good representation. The reason for this is that the structure (5) is not adequate to describe a solution of the form (2). Roughly speaking the dyadic products describe rectangular (with respect to the space-time diagram) structures well, while diagonal structures such as in Figure 1, left, cannot be well represented as a dyadic product. This explains the common failure of the mode-based model reduction for transport phenomena (cf. [9, 36]). However, if the velocity is known or can easily be detected, the solution is near at hand. A time-dependent shift which compensates the transport velocity yields a structure that can be written as a single dyadic product, which in turn corresponds to the description by just one mode, Figure 1, right.

To this end, we introduce the continuous shift operator $\mathcal{T}^c f(x,t) := f(x-ct,t)$ acting on a function f(x,t). The discrete analogue of \mathcal{T}^c is denoted with $T^c(\cdot)$ which acts on a snapshot matrix of a time- and space-dependent function and shifts the *i*th column by ct_i in space. Thus, for one transport we seek approximations

(6)
$$q(x_i, t_n) \approx T^c \left(\sum_l \alpha_l(t_n) \phi_l(x_i) \right).$$

Clearly such an approximation can be constructed by applying the inverse shift T^{-c} into the co-moving frame and performing an SVD. For our example of the single transported quantity, we obtain $T^{-c}(q(x,t)) = q(x,0)$, which is described by one mode. Note that this shift demands an interpolation in the general case where ct_i is not an integer multiple of the grid size. Note further that the approach also works with a nonconstant transport velocity, since the shift can be set for each time step separately.

In a similar manner, shifts have been used in other reduction frameworks, e.g., for the model reduction of a combustion [26] and in a symmetry reduction framework [6, 17, 37]. The latter one treats partial differential equations (PDEs) which are equivariant with respect to a group action which in turn induces symmetries in the solution space. Symmetry-reduced surrogate models on the PDE level are obtained by exploiting the equivariance of the original PDE. In contrast, we are looking for a general decomposition of a snapshot matrix by shifted modes without requiring equivariance of the original PDE or symmetries in their solutions.

2.2. Multiple transported quantities. Many relevant systems feature multiple transported quantities. We extend (6) to a *multiframe decomposition* of the form

(7)
$$q(x_i, t_n) \approx \sum_{k=1}^{N_s} T^{c_k} \left(\sum_l \alpha_l^k(t_n) \phi_l^k(x_i) \right),$$

where N_s is the number of transport velocities. We seek to decompose a given spacetime field into this structure. To this end, a best fit of (7) in the vector 2-norm will be used, where the ansatz modes will be constructed by shifting and reducing the field data to identify low-rank structures in each co-moving frame. The reduction and best fit are iterated to decrease the cross-influence between the different transports and to obtain a clear separation.



FIG. 2. The density of two transported quantities, the left and right going waves. MS&R (top left: full-order solution; top right: error of (8); bottom left: right going component of MS&R; bottom right: left going component of MS&R).

To illustrate the new approach for problems with multiple transports, we consider the solution (2) of the wave equation (1) with

$$q_{+} = q_{-} = \exp(-(x - L/2)^2/(L/50)^2),$$

which describes a pressure pulse in a system initially at rest. The analytic solution of the density for these two transported quantities is shown in Figure 2, top left. Due to the solution structure (cf. (2)), it is known that the full information is described by just two modes. Ideally, a model reduction approach should detect these modes. However, while the solution structure is known in this special case, we refrain from using details of it. The desired method should find such modes purely data-based. In this way, such a procedure is expected to work also for nonhyperbolic transport. For the remainder of this subsection, we assume the transport velocities to be known, while in section 2.3 methods of determining transport velocities based on the snapshot data are discussed.

To identify good ansatz modes, recall that in the case of one transported quantity it is well described by the first modes in a co-moving system, while a decomposition of structures in a different velocity frame leads to the need of many modes to describe the dynamics reasonably well. Consequently, the first mode has an inferior contribution to the dynamics, if the shift velocity does not agree with one of the transport velocities. Thus, a naive approach to *decompose* different velocity components is given by the following procedure.

A1327

MULTI-SHIFT & REDUCE (MS&R)

Input:

• $j \times n$ snapshot matrix X with n (temporal) snapshots of j grid points

• transport velocities c_k , $k = 1, \ldots, N_s$, N_s number of shift velocities

Procedure:

1. Shift & Decompose

Compute the SVDs

$$U^k \Sigma^k (V^k)^T = T^{-c_k} (X)$$

with $U^k \in \mathbb{R}^{j,j}$, $\Sigma^k \in \mathbb{R}^{j,n}$, and $V^k \in \mathbb{R}^{n,n}$ for $k = 1, \ldots, N_s$. Here, T^{-c_k} denotes the discrete shift operator; cf. end of section 2.1.

2. Truncate

Approximate the matrix $U^k \Sigma^k (V^k)^T$ by neglecting the singular values σ_{r+1} , ..., $\sigma_{min(k,n)}$, where r is chosen as small as possible but as large as necessary to obtain a good approximation. This leads to the approximate SVDs

$$\tilde{U}^k \tilde{\Sigma}^k (\tilde{V}^k)^T \approx U^k \Sigma^k (V^k)^T$$

with $\tilde{U}^k \in \mathbb{R}^{j,r}$, $\tilde{\Sigma}^k \in \mathbb{R}^{r,r}$, and $\tilde{V}^k \in \mathbb{R}^{n,r}$ for $k = 1, \ldots, N_s$.

Altogether MS&R

$$MS\&R : X \to \tilde{\Sigma}^k, \tilde{U}^k, \tilde{V}^k$$

produces modes containing parts of the field which can be represented well in the respective velocity frame. A naive approximation of the original field would then be given by

(8)
$$X \approx \tilde{X} = \sum_{k=1}^{N_s} T^{c_k} \left(\tilde{U}^k \tilde{\Sigma}^k (\tilde{V}^k)^T \right).$$

If the matrices would contain only information of the associated velocity frame this approximation could be exact. In general a perfect decomposition is not to be expected and instead parts of the field are overrepresented.

The result of applying this procedure to the pressure pulse example is shown in Figure 2, bottom, where we applied MS&R with r = 1 in each velocity frame. MS&R leads to an approximate identification of the respective velocity component, since this dominates the respective first mode, while the other velocity component gives a relatively small contribution. The reconstruction is done according to (8) and the error is shown in Figure 2, top right.

While the main structures are reproduced well, the result is far from perfect. Especially, in the region where the left and right going pulses overlap, a large error is produced. These strongly localized structures with respect to space and time are represented well in both velocity frames, since a point structure is nearly unchanged by a shift and can be described perfectly by a dyadic product. This ambiguity impedes the separation of the two transported quantities. Furthermore, the quality of the solution cannot be improved by adding modes, since this counteracts the separation.

To improve the decomposition, a new algorithm is presented in the following which is referred to as the sPOD algorithm. The MS&R identifies by construction low-rank structures in each velocity frame. These are used to create a set of ansatz modes to optimize the approximation (7) of the field in a least squares sense. To this end, the error or residual is decomposed by MS&R to identify the part of the error, which can be represented in each velocity frame. Finally, the decomposition is iteratively improved. An implementation of the sPOD algorithm and the numerical examples considered in this paper are available from [34].

The sPOD algorithm

Input:

- $j \times n$ snapshot matrix X with n (temporal) snapshots of j grid points
- transport velocities c_k , $k = 1, ..., N_s$, N_s number of shift velocities *Procedure*:

1. Initialize $j\times n$ matrix $\tilde{X}=0,~j\times r$ matrix $\tilde{U}^k=0,~n\times r$ matrix $\tilde{V}^k=0$ loop

2. calculate residual

$$R = X - \tilde{X}$$

3. use MS&R to create ansatz modes

$$MS\&R : R \to \Sigma_r^k, U_r^k, V_r^k$$

4. optimize in a least squares sense

$$\min_{\alpha^k, \alpha^k_r} \left\| \hat{X} - X \right\|_2^2$$

with

(9)
$$\hat{X} = \sum_{k=1}^{N_s} T^{c_k} \left(\tilde{U}^k \alpha^k (\tilde{V}^k)^T + U_r^k \alpha_r^k (V_r^k)^T \right),$$

where α^k and α_r^k are diagonal coefficient matrices of appropriate dimensions 5. calculate new modes by SVD

$$U^k \Sigma^k (V^k)^T = \tilde{U}^k \alpha^k (\tilde{V}^k)^T + U_r^k \alpha_r^k (V_r^k)^T$$

and truncate it as in step 2 of MS&R

$$\tilde{U}^k \tilde{\Sigma}^k (\tilde{V}^k)^T \approx U^k \Sigma^k (V^k)^T$$

6. update approximation

$$\tilde{X} = \sum_{k=1}^{N_s} T^{c_k} \left(\tilde{U}^k \tilde{\Sigma}^k (\tilde{V}^k)^T \right)$$

until $||R||_2$ does not reduce further

Due to the initialization, the first decomposition is MS&R of the original field. In further steps the residual is reduced by providing modes which are constructed by MS&R of the residual and allow removing structure which is multiply accounted for. For example, the strong peaks in the residual in Figure 2 are prominently visible in the residual modes U_r^k and can be used to remove this error to a large degree. The iteration in sPOD offers the possibility to remove it up to any achievable tolerance.

An open question is how to choose the r in the MS&R for each frame in an optimal way. For cases where r is not clear from physical considerations, we propose two different heuristics.

One method is to perform the sPOD decomposition with a large number of modes per frame and select the most important ones afterward. This is done by sorting the singular values of all reference frames in one list and choosing the modes associated with the largest singular values. If a certain approximation error is prescribed, one has to calculate the residual from the reconstruction. It cannot, however, be calculated from the singular values directly, due to the nonorthogonality of modes belonging to different frames.

An alternative way of choosing r is to start with small numbers of modes and successively add modes in a greedy fashion. We explain this in more detail with an example of three reference frames and starting with zero modes for each frame, i.e., $r_0 = [000]$, where r_k denotes the vector containing the numbers of modes of each velocity frame after the kth greedy iteration. At the first iteration, different sPOD approximations are computed for $\mathbf{r} = [100], \mathbf{r} = [010], \text{ and } \mathbf{r} = [001]$. The errors of the three different sPOD approximations are compared and we add only one mode to that frame which corresponds to the smallest error, for instance, $r_1 = [010]$ if the comparison shows the greatest advantage of adding a mode to the second frame. This procedure can be continued until a certain error tolerance is achieved and can be implemented as a loop around the sPOD algorithm. Of course, the greedy algorithm is based on a locally optimal choice which does not guarantee that the resulting vector r is optimal. Furthermore, this approach can be quite costly if the amount of data and the number of frames are big since the sPOD algorithm has to be performed multiple times. Nevertheless, in the considered numerical examples the greedy algorithm yields sPOD approximations with high accuracy while only using a small number of modes per frame.

The sPOD algorithm is tested for the density of the acoustic pulse with one mode per frame. The convergence behavior can be seen from the solid graph of Figure 3, left, where the convergence is indicated by the decreasing mean error over iterations. The considered error measure is given by the 2-norm of the error divided by the 2-norm of the full-order snapshot matrix, i.e.,

(10)
$$\operatorname{mean \ error} = \frac{\left\| X - \tilde{X} \right\|_2}{\|X\|_2}$$

Some remarks follow.



FIG. 3. Two transported quantities. Left: Convergence of the sPOD algorithm (2 modes). The error of the sPOD approximation decreases to less than 5×10^{-13} . Right: Full-order solution of the density, sPOD approximation, and the error (note the color scale).

1. While in the first iterations the sharp structures allow a good separation of the different transport directions, the residual tends to lose this structure. The interference of the components of different velocity create a broad structure, visible as the broad diagonal stripes in Figure 2. This tends to slow the convergence after the first iterations.

2. For the considered example, the analytic solution has the form of two modes, each in a different frame, i.e., $X = T^{c_1}(u^1\sigma^1(v^1)^T) + T^{c_2}(u^2\sigma^2(v^2)^T)$. This allows us to consider the modes from the initial MS&R as the modes of the exact solution perturbed by the component of the respective other velocity frame. For example, for the first frame of reference, the shifted snapshot matrix is

$$T^{-c_1}(X) = u^1 \sigma^1 (v^1)^T + T^{c_2 - c_1} (u^2 \sigma^2 (v^2)^T).$$

Assuming the perturbation $T^{c_2-c_1}(u^2\sigma^2(v^2)^T)$ to be small, the perturbed singular values and vectors can be expressed as a linear perturbation of the unperturbed ones with $\delta\sigma^1 = (u^1)^T \delta X^1 v^1$ and $\delta u^1 \sigma^1 = \delta X^1 v^1 - u^1 \delta\sigma^1$ characterizing the linear change of σ^1 and u^1 . For the considered pressure pulse example, the mode obtained by MS&R is well approximated by $\tilde{u}^1 \approx u^1 + \delta u^1$. This suggests that the exact mode is indeed dominant in its reference frame and the component of the other frame is approximately a linear perturbation. The broad stripes in Figure 2 are thereby explained as an overlap of the exact mode and the mode of the other frame in the respective velocity frame.

3. Broad structures in the residual can be better resolved with more modes. Therefore, increasing the number of modes r in the MS&R for the residual tends to result in an improved convergence. Indeed, we have observed for the wave equation with some less sharp initial conditions that choosing r = 1 prevents the procedure from converging to the analytic solution. For these examples increasing the number of residual modes r improves convergence and results in finding the analytic solution up to any given achievable tolerance. However, there is a trade-off between convergence and computational effort, both of which have to be accounted for when choosing the number of residual modes.

4. The ansatz in (9) can also be extended so as to include cross terms of the form $U^i(V^j)^T$ with $i \neq j$. This can be achieved by optimizing over dense coefficient matrices α^k and α_r^k rather than over diagonal ones. We have observed that this often leads to a faster convergence of the sPOD iteration which is due to the increased number of degrees of freedom for the optimization in step 4 of the sPOD algorithm. Nevertheless, we omit the cross terms for reasons of computational effort. The number of unknowns in the resulting system of linear equations to be solved in step 4 would scale with $N_s r^2$ instead of $N_s r$. This is a minor issue for examples in one dimension but it leads to high computing times when dealing with the two-dimensional example presented in section 3.

5. As an alternative to reducing the different velocity frames simultaneously, one could proceed sequentially by removing the detected structures in a greedy fashion. This is done in the recent work [36]. It is clear from Figure 2 that after the first iteration each of the modes contains structure from the respective other velocity frame. However, due to the greedy character these artifacts are not removed, since the modes from the first iteration remain unchanged. As a consequence, the sequential procedure does not reduce to descriptions with just one mode per velocity frame in the pressure pulse example. Moreover, the obtained modes do not reflect the physics of two moving waves properly, which might also be disadvantageous for model reduction.



FIG. 4. Two transported quantities (only density): Comparison of full-order solution with sPOD and POD approximation (left: full-order solution; middle: sPOD approximation (2 modes); right: POD approximation (2 modes)).

The error of approximation obtained by the sPOD algorithm after 40 iterations is shown in Figure 3, right. As can be seen, the solution is approximated excellently by just two modes. This is the desired low-dimensional representation we have been looking for, agreeing with the analytic solution given by two modes. Note that since only the density was provided, the Riemann invariants cannot be calculated, suggesting that the construction did not implicitly use the hyperbolic structure of the equation.

The good performance of the sPOD algorithm gets more striking when comparing it to the standard POD. For this purpose, Figure 4 depicts the comparison between the full-order solution, the sPOD approximation with two modes, and the POD approximation with two modes. It is obvious that the POD approximation is highly inadequate, whereas the sPOD approximation matches the full solution excellently. To obtain the same accuracy as the two sPOD modes (error less than 3×10^{-14} ; cf. blue, solid graph in Figure 3, left), more than 80 POD modes are required. Furthermore, if only two POD modes are used, the relative mean error, as defined in (10), is almost 1. Of course, this insufficient performance of the POD was to be expected, since deliberately we have chosen an example which provides a big challenge for the POD. Nevertheless, this test case which is generic for many practical problems gives a first impression of the potential of the proposed sPOD.

For all the considerations made so far, the reduction was based on the density alone. When considering the density and the velocity together, the convergence of the sPOD algorithm appears to be faster even though the eventually achieved error is somewhat higher but the order of magnitude is 10^{-13} (see Figure 3, left). Apart from the convergence behavior, the results for the case that density and velocity are considered together are very similar to the results shown before and therefore omitted.

2.3. Determination of the shift velocities. If the shift velocities are unknown, different strategies can be used to obtain good candidates for them. This is largely independent of the sPOD, which is the main focus of this paper, but important for practical use. Often physical insight reveals (some of) the involved velocities, i.e., the flow or sound velocity. The main target of the sPOD is the decomposition of strongly transport-dominated cases. This class of problems is especially interesting since classical methods like the POD fail here most severely; see, e.g., [9]. In these cases the transport velocities can often simply be determined by usage of data-based tracking of peaks or threshold values within the snapshot matrix. Another method is discussed in the following, where the velocity detection is performed via a maximization of singular values. The tracking-based method is applied in the example of



FIG. 5. The singular value spectrum of the density snapshot matrix as a function of the shift velocity. Left: The pressure pulse, which creates two traveling waves with the speeds $c_{\pm} = \pm 1$, which are clearly visible as maxima of the leading singular value. Right: The standing wave. The zero velocity has the biggest leading singular value. The positive and the negative sound speed are also visible as local maxima.

section 2.5, while the velocities of the example in section 3 are determined by means of the singular value maximization.

A purely data-driven method is obtained by examining the singular value spectrum of the shifted snapshot matrix as a function of the shift velocity, i.e., the (constant) velocity the time-dependent shift is based on. In Figure 5, left, the spectrum of shifted snapshot matrices of the example considered in section 2.2 is shown for a range of shift velocities between -1.25 and 1.25. First, it should be noted that, assuming periodic boundaries, the square integral of the solution

$$\int_0^{t_{end}} \int_0^L (q(x,t))^2 \mathrm{d}x \mathrm{d}t$$

does not change by a shift in the x-direction. A numerical approximation of the shift keeps this invariance up to the interpolation error. Consequently, the Frobenius norm, which is directly connected with the singular values by

$$||X||_F^2 = \sum_{i,j} X_{i,j}^2 = \sum_i \sigma_i^2,$$

is also shift-invariant up to the interpolation error. This allows us to directly compare spectra for different shift velocities. If a certain singular value increases by a change of shift velocity, others have to decrease to keep the sum of the squared singular values constant. The transport velocities, which are the positive and the negative sound velocity $c_{+} = \pm 1$, are clearly visible from the maxima of the leading singular value. The classical POD is recovered for a shift velocity of zero, where a slow decay of the singular values can be seen in Figure 5, left. The shift dependence of the pressure pulse case is contrasted with the case of a standing wave, Figure 5, right. Here the zero shift velocity leads to a maximization of the leading singular value. However, the sound velocities are also visible as local maxima. The standing wave can be represented by two traveling waves with velocities $\pm c$, so that these are also reasonable candidates for shift velocities. It should be noted that the spectra are only a first indicator. If multiple transport velocities are involved, each of the transported quantities influences the singular spectra of the others. For instance, the pressure pulse can be expressed by one mode per transport velocity, which is not obvious from Figure 5, left, since each of the transported quantities slows the decay of the singular spectrum of the other one. The simple sampling used here can be replaced by a gradient method, since a change of singular values with respect to a change of the matrix can easily be calculated, similar to the considerations in remark 2 following (10).

2.4. Uniqueness and robustness of the method. For the case $N_s = 1$ and $c_1 = 0$ in (7), the sPOD reduces to the classical POD and hence for this special case an optimal approximation is given due to the optimality of the SVD. A simple counterexample shows that such a global optimality cannot be guaranteed for the sPOD in general. Consider again the pressure pulse example from section 2.2 and add a third velocity $c_3 = 0$ to the sPOD algorithm. Since the solution can be described perfectly by the two moving modes, the zero velocity frame is redundant. However, we can describe the analytic solution also by adding a harmonic mode to the zero velocity frames; cf. (3) and (4). Accordingly, the analytic solution allows different but equivalent low-dimensional representations. Adding a mode in the zero frame and subtracting from the Riemann invariants does not change the composed solution, e.g.,

$$\begin{aligned} q(x_i, t_n) &= T^c \left(\begin{bmatrix} \rho_0 \\ c \end{bmatrix} q_+(x) \right) + T^{-c} \left(\begin{bmatrix} \rho_0 \\ -c \end{bmatrix} q_-(x) \right) \\ &= T^c \left(\begin{bmatrix} \rho_0 \\ c \end{bmatrix} (q_+(x) - \cos(k_1 x)) \right) + T^{-c} \left(\begin{bmatrix} \rho_0 \\ -c \end{bmatrix} (q_-(x) - \cos(k_1 x)) \right) \\ &+ 2T^0 \left(\begin{bmatrix} \rho_0 \cos(k_1 ct) \cos(k_1 x) \\ c \sin(k_1 ct) \sin(k_1 x) \end{bmatrix} \right). \end{aligned}$$

This shows that the separation is not unique and all of these solutions are fix points of the sPOD iteration. Indeed, the numerical experiments show that the sPOD algorithm does not remove the component of the zero velocity frame but instead converges to a solution which is represented as a sum of the three provided frames. The accuracy is still comparable to the sPOD just using the two co-moving subspaces. This means that the sPOD is not optimal in the sense that a minimum total number of modes is not guaranteed. Note that a decomposition based on three modes is still much better than the classical POD in this example. One could think of strategies to enforce removing redundant information, for instance, by augmenting the least squares approach in step 4 of the sPOD algorithm with an ℓ_1 -norm regularization (see, e.g., [8]). However, this exceeds the scope of this paper.

If the transport velocity is not determined correctly but with a small error, an increased number of modes is required to obtain the same accuracy as with the correct transport velocity. However, as can be seen in Figure 5, the leading singular value decreases continuously from the maxima. Hence, small errors from the velocity detection lead to a small increase of the required number of modes. This growth depends essentially on the width of the maximum. A sharp maximum, i.e., a maximum exhibiting a strong decay of the leading singular value close to it, requires a precise velocity determination. On the other hand, sharper maxima are also easier to determine, which suggests that also for these cases a sufficiently accurate determination of the dominant transport velocity is possible with the presented method.

2.5. Two crossing shocks. In this section we consider a more physical and challenging example: The crossing of two shock waves. The example of two crossing shocks is a generic phenomenon which occurs in many applications. The nonconstant shock velocities and amplitudes are inherently nonlinear and change during the crossing of the shocks. However, this is not a problem for the sPOD decomposition which can also work with variable velocities. To this end, the shift operator

THE SHIFTED PROPER ORTHOGONAL DECOMPOSITION A1335



FIG. 6. Full-order solution of two crossing shocks: Pseudocolor plots of the density (left), velocity (middle), and pressure (right).

introduced in section 2.1 is generalized to account for time-dependent velocities by replacing the linear expression ct by a general time-dependent shift coordinate $x^{sh}(t)$, i.e., $\mathcal{T}^{x^{sh}}f(x,t) := f(x - x^{sh}(t), t)$.

The Euler equations for an ideal gas with constant heat capacity are simulated by the skew-symmetric scheme with a shock filter, as described in [35]. The nonlinear shock filter [7] chooses the filter strength on the basis of a shock detector. The initial condition is $\rho = [1.0138, 0.6000, 1.0138]$, u = [-96.2197, -310.0174, -523.8264], p = [1.2720, 0.6000, 0.2720], yielding two shocks colliding with a shock Mach-number of Ma = 1.4. The shocks are at x = [1/3, 2/3]. The initial shock structure is chosen to have a near steady state with the given shock filter. The adiabatic exponent is $\gamma = 1.4$.

The simulated solution is depicted in Figure 6 for the density, the velocity, and the pressure. The depicted space-time diagrams represent the transposes of the respective snapshot matrices which are used as an input of the sPOD algorithm. The snapshots have been normalized such that the highest occurring absolute value of each quantity is equal to one to avoid numerical errors due to the different scales. In contrast to the examples before, the case of two crossing shocks exhibits nonperiodic boundaries. This leads to the situation that we need values from outside of the computational domain when shifting the snapshot matrix. We treat this here by a constant extrapolation over the boundaries, while a general treatment of nonperiodic boundaries within the sPOD framework will be addressed in a future work.

We are looking for a decomposition of the snapshot matrices with a prescribed relative error tolerance of one percent. To apply the sPOD the time-dependent shift needs to be determined. Here, we apply a simple threshold search within the snapshot matrix of the velocity. More precisely, we define two threshold values characterizing the velocity jumps at the two shocks. To this end, the respective threshold value has to be chosen in between the constant values at the left- and the right-hand side of the respective shock. Regarding the velocity snapshots depicted in Figure 6, we have chosen the thresholds -390m/s (normalized -0.74) for the right border of the zone between the shocks (border between cyan and dark blue in the middle plot of Figure 6) and -200m/s (normalized -0.38) for the left border (between yellow and cyan). Based on these thresholds, the time-dependent shifts are determined by searching the last value, which is larger than the corresponding threshold, in each column of the snapshot matrix (each discrete time step).

The shift coordinates obtained from the threshold search do not intersect but rather bounce off each other; see Figure 7(a). In order to account for the crossing of the shock waves, the parts after the time of shortest distance between the two curves are switched and in the intermediate area a linear interpolation is applied. The resulting shift coordinates (cf. Figure 7(b)) provide the basis for the shift matrices



(a) Shift coordinates determined via thresh- (b) Shift coordinates after correction to acold search. count for the crossing.

FIG. 7. Shift coordinates.



FIG. 8. Error of the sPOD approximation of two crossing shocks: Pseudocolor plots of the density (left), velocity (middle), and pressure (right).

applied in the sPOD algorithm. In order to satisfy the given accuracy requirement of one percent, we apply the sPOD algorithm together with a greedy increment of the number of modes as explained in section 2.2.

We end up with seven modes, four for the left going and three for the right going wave. The assembled sPOD approximation is nearly indistinguishable from the simulated solution; the error plots are depicted in Figure 8. The approximation agrees neatly with the full-order solution which can be comprehended by looking at the error whose maximum amplitude is around six percent of the maximum amplitude of the full-order solution. The relative mean error is less than one percent as required. Note that this could be achieved with just seven modes in total, whereas a POD approximation needs 51 modes to attain the same accuracy.

The error plotted in Figure 8 reveals structures. First, the shocks are visible. This is caused by a not strictly constant shock structure due to small variations in the strength of the shock adaptive filter. The dynamics of this numerical artifact is not fully described by the used number of sPOD modes. Second, waves in the density are emerging from the shock crossing and from the initial position of the left shock. The two extra waves travel with the flow velocity. The steady state condition of the left shock is slightly perturbed. Small perturbations travel as characteristic waves, so that this perturbation creates a so-called entropy wave. A similar perturbation is created by a change of shock filter strength during the shock crossing. These structures in the error could be removed by adding frames with the flow speeds to the decomposition.

THE SHIFTED PROPER ORTHOGONAL DECOMPOSITION



FIG. 9. 2D vortex pair: Full-order solution (vorticity contour lines at $-270, -250, \ldots, 270$).

3. A two-dimensional model problem. In the last section we have considered one-dimensional test examples, which provide a big challenge for the classical POD. The sPOD performs excellently, because its structure is adequate for the transport-dominated examples.

In this section we want to explore the capabilities of the introduced sPOD algorithm by considering a two-dimensional, transport-dominated, nonhyperbolic example with nontrivial velocities. We consider a flow governed by the incompressible Navier– Stokes equations with a vortex pair as initial condition. The initial conditions are created from two single vortices with vorticity

$$\omega_{0,i} = \omega_{e,i} \left(1 - (r_i/r_0)^2 \right) \exp\left(-(r_i/r_0)^2 \right)$$

where $i = 1, 2, r_i = \sqrt{(x - x_{0,i})^2 + (y - y_{0,i})^2}$ is the distance from the respective vortex core and r_0 denotes the vortex size. The size of each vortex is chosen to be $r_0 = 0.1$ and the centers are at $(x_{0,1/2}, y_{0,1/2}) = (\pm 0.1, 0)$. The strengths are $\omega_{e,1/2} =$ ∓ 299.5 and the viscosity of the fluid is $\nu = 1/Re = 1/1000$. We consider a periodic domain to avoid further complication of the boundary treatment for the sPOD. The dynamical behavior of the vortex pair is simulated by means of the energy conserving, skew-symmetric scheme described in [33]. The periodic domain is discretized by 512^2 equidistant points with a time step of $\Delta t = 8 \cdot 10^{-4}$ s with a fractional-step time stepper. The solution is depicted in Figure 9 by means of a contour plot for different times. The initial vorticity field induces a movement of the vortex pair in the positive y-direction. Additionally, a secondary, weaker vortex pair moves in the negative y-direction with a smaller transport velocity. Due to this secondary vortex, two different, nontrivial velocities are present. The simulated time is chosen such that the primary and the secondary vortex pair do not meet again in the periodic domain. Compared with the linear wave equation from section 2, this example is significantly more challenging for several reasons. First, we do not know the analytic solution of this test case. Second, the transport velocities are non-constant and unknown a priori. Third, the two-dimensional problem not only leads to larger data sets but tests the applicability in more dimensions.

Before we apply the sPOD algorithm as introduced in section 2, we need to find proper candidates for the dominant transport velocities. We are looking for two velocities which we treat as constant, even though this is not strictly true. For the velocity determination, we again consider the singular values of shifted snapshot matrices where we only shift in the *y*-direction. Since the considered case is twodimensional, we need to reshape the snapshot matrix as it would be done for the POD.



FIG. 10. Two-dimensional vortex pair: Left: Velocity scan in the y-direction. The two marked maxima at -1.1193 and 2.7663 are used for the decomposition. Right: Convergence of the sPOD measured by the norm of the residual related to the norm of the field. In the inset the ratio of the largest singular value of the residual and the smallest singular value of the sPOD approximation in the respective frame is plotted (solid line for the primary, red dashed for the secondary vortex pair). The horizontal line marks the value one below which the resolved singular values are greater than those of the residual. The second inset shows the maximal orthogonality error; cf. (11).

The singular values as a function of the transport velocity are shown in Figure 10, left. The maxima are at -1.1193 and 2.7663, which are used as the transport velocities in the decomposition of the secondary and the primary vortex pair, respectively.

A decent approximation can be obtained by choosing 14 modes for the primary and 10 for the secondary vortex pair, so in total 24 modes. The number of modes is selected by the first method described after the sPOD algorithm in section 2.2. A decomposition with a large number of modes (here 15 per frame) is calculated and the modes connected with the largest singular values from both frames are selected. The decay of the error with this procedure is seen later in Figure 12, right. Moreover, the chosen number of modes varies during the sPOD algorithm. More precisely, we start with one mode per frame and add one mode to each frame in each iteration of the sPOD algorithm until the specified mode number, here 15 per frame, is reached. The resulting convergence of the sPOD for the two-dimensional vortex pair problem is depicted in Figure 10, right. In comparison to the pressure pulse example (cf. Figure 3, left), the convergence is worse but still the relative mean error decreases to a value of slightly less than 1%. Due to the used constant shift velocities and the complexity of the transport phenomenon, a much smaller error cannot be expected. In the insets of Figure 10, right, two further convergence criteria are investigated which are introduced in the following.

Two important properties of the SVD are the orthogonality of the approximation to the residual and that the smallest singular value of the approximation is not smaller than the biggest singular value of the residual. These two properties are essential for the optimality of the POD. We formulate their analogues for the case of multiple moving frames:

(11)
$$\left\langle \tilde{u}_{l}^{k}(\tilde{v}_{l}^{k})^{T}, T^{-c_{k}}(R) \right\rangle_{F} = 0 \quad \text{and} \quad \min(\tilde{\Sigma}^{k}) \ge \max(\Sigma_{r}^{k})$$

for all $k = 1, ..., N_s$. Here $R = X - \tilde{X}$ is the residual between the snapshot matrix Xand the sPOD approximation $\tilde{X} = \sum_{k=1}^{N_s} T^{c_k} (\tilde{U}^k \tilde{\Sigma}^k (\tilde{V}^k)^T)$ as output from the sPOD algorithm (see section 2.2). Furthermore, \tilde{u}_l^k and \tilde{v}_l^k are the columns of \tilde{U}^k and \tilde{V}^k , respectively. The matrix Σ_r^k is diagonal containing the singular values of $T^{-c_k}(R)$ and $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product. The first inset in Figure 10, right, shows the ratio $\max(\Sigma_r^k)/\min(\tilde{\Sigma}^k)$ plotted over the iterations for both the primary and the secondary vortex pair. The criterion $\min(\tilde{\Sigma}^k) \geq \max(\Sigma_r^k)$ is satisfied from iterations 9 and 12 on, respectively. In the second inset, the iterations of the maximum relative orthogonality error

$$\max_{l} \left(\frac{\left| \left\langle \tilde{u}_{l}^{k} \left(\tilde{v}_{l}^{k} \right)^{T}, T^{-c_{k}} \left(R \right) \right\rangle_{F} \right|}{\| \tilde{u}_{l}^{k} \left(\tilde{v}_{l}^{k} \right)^{T} \|_{F} \| T^{-c_{k}} \left(R \right) \|_{F}} \right)$$

are shown. Here, the decay is not as clear as for the other criteria, which is due to a dominance by modes associated with a small weight, while the modes associated with larger singular values have a much lower orthogonality error. After all, the decay of the mean error appears to be the most appropriate criterion for describing the convergence.

The first three modes obtained by the sPOD for each transported vortex pair are shown in Figure 11. The depicted modes correspond to the respective co-moving frames, i.e., \tilde{u}_l^k . For both velocity frames the respective first mode obviously describes the corresponding vortex pair. The higher modes seem to fulfill different roles in both cases. When considering the dynamical numerical solution, we observe that the secondary, down-going pair strongly changes in time; a rotating vorticity distribution is visible at creation. It looks very much like the second mode; thus, the higher modes represent the change of shape of the vortex pair. On the other hand, the numerical simulation results reveal that the shape of the primary, up-going vortex pair changes less in time. Right after the separation the sharp vortex structure emerges. The higher modes of the primary vortex pair have a strong weight at the front and rear



FIG. 11. Two-dimensional vortex pair: First three sPOD modes for each transported vortex pair (top: primary vortex pair; bottom: secondary vortex pair).



FIG. 12. Two-dimensional vortex pair: Left, middle: Time amplitudes for the first three modes for each transported vortex pair. Right: Mean error over number of modes (POD versus sPOD).

edges of the pair. These structures seem to account for a velocity correction, since they induce a translational shift of the primary vortex pair. The velocity correction is caused by the nonconstant propagation velocity of the vortex pair.

The corresponding time amplitudes of the spatial modes are depicted in Figure 12. The amplitudes are given by the right singular vectors \tilde{v}_l^k , describing the temporal development of the respective mode \tilde{u}_l^k , multiplied by the corresponding singular value $\tilde{\sigma}_{l}^{k}$. For large values of t, the first mode dominates the behavior of the primary vortex pair, while for the secondary pair the first two modes have a nearly similar weight. Thus, the first and the second mode of the secondary vortex pair both contribute to its shape. This can be comprehended by considering Figure 11, where it can be seen that the vortex pair appearing in the first mode seems to fit into the cavity formed by the structures of the second mode. This agrees with the observation that the secondary vortex pair depicted in Figure 9 is wider than the vortex pair shown in the first mode; cf. Figure 11, bottom left. The separation of the two vortex pairs is completed at about t > 0.15s. The interplay between the modes of the two vortex pairs is complicated up to that point, but afterward a simpler picture emerges. For the primary vortex pair, the absolute value of the amplitude of the dominating first mode decays gradually in time, which reflects the reduction of the vortex strength by viscosity. The influence of the second mode becomes stronger from roughly t = 0.2on, i.e., after the full separation. It leads to a reduction of the vortex pair transport velocity. Also, the third mode contributes to this velocity decrease. In contrast, the initially negative time amplitude of the first mode of the secondary vortex pair increases and changes the sign over time. After separation all three amplitudes are nearly constant; this corresponds to the plausible observation that the friction has less impact on the secondary than on the primary vortex pair due to the lower velocity gradient.

While the modes give a clear and intuitive description of the flow dynamics, still some unexpected structures are visible. Namely, one can also see some stripe structures in the y-direction, especially at the first mode of the secondary vortex pair; cf. Figure 11. Our first analysis indicates that this is *not* due to an incomplete separation but rather is caused by a nonuniqueness of the decomposition. This nonuniqueness can be comprehended by the following consideration. If we add a constant offset to one of the transported quantities, the combined sPOD approximation remains unchanged if the negative counterpart offset is added to the other transported quantity. Note that offsets are shift-invariant and, hence, well represented in all velocity reference frames. The occurring stripes seem to be localized variants of these offsets along the y-direction. These structures may be removed by an additional, physically motivated

A1340



FIG. 13. Two-dimensional vortex pair: Comparison of full-order solution with sPOD and POD approximation at t = 0.27s (left: full-order solution; middle: sPOD approximation (9 modes for the primary, 6 for the secondary pair); right: POD approximation (15 modes)). Left half of each plot: vorticity contour lines at $-270, -250, \ldots, 270$; right half of each plot: pseudocolor plot.

constraint. Nevertheless, the main structure of the vortex pairs is already captured quite well by the respective first modes.

The spatial modes and their time amplitudes both provide physically meaningful insights into the dynamics of the vortex movements. Finally, the sPOD is to be evaluated quantitatively by comparing it to the full-order solution and the POD approximation. To this end, Figure 13 shows the full-order solution, the sPOD approximation (9+6 modes), and the POD approximation (15 modes), each for a constant time t = 0.27s. It is hard to see a difference between the sPOD approximation and the full-order solution, since they match almost exactly. In contrast, even though the POD approximation captures the main structures of the vortex pairs quite well, it reveals some unphysical structures in the inside as well as a blurred surface of the primary vortex pair (top) and, furthermore, some spurious edges in the secondary vortex pair (bottom). The superiority of the sPOD algorithm can also be comprehended by considering Figure 12, right, where the mean error is plotted over the number of modes for the POD and the sPOD. The sPOD mean error is significantly smaller than the POD mean error for all tested numbers of modes although the difference is not as striking as in the wave equation example; cf. section 2.2. This is not surprising, since the one-dimensional example is ideal for the reduction by the sPOD, while the vortex pair problem provides some additional challenges. Further, the small transport distance compared with the relatively large interaction time at the beginning reduces the difference between the methods.

It should be noted that very high numbers of modes may be problematic for the described method. The separation of velocity frames builds on low-rank approximations so that the algorithm is expected to become ineffective for a large number of modes.

All in all, the sPOD method performs very well despite the nontrivial and nonconstant velocities of the two-dimensional vortex pair problem and further delivers modes which allow a physically meaningful interpretation.

4. Summary and outlook. We have presented the sPOD as a new model reduction approach. It generalizes the common POD by allowing for time-dependent shifts of the modes. More precisely, given a snapshot matrix X and transport velocities c_1, \ldots, c_{N_s} , the sPOD constructs an approximate decomposition of the form

(12)
$$X \approx \sum_{k} T^{c_k} \left(\tilde{U}^k \tilde{S}^k (\tilde{V}^k)^T \right)$$

A1342

with low-rank matrices $\tilde{U}^k \tilde{S}^k (\tilde{V}^k)^T$ and shift operators T^{c_k} which shift every column of their respective arguments by an amount defined by c_k . Due to the shifts, the sPOD is especially well-suited for transport-dominated phenomena. Even problems with multiple shifts can be successfully separated and decomposed based on the iterative sPOD algorithm which is based on truncated SVDs of shifted snapshot matrices and shifted residual matrices. The algorithm can easily be generalized to nonconstant velocities as demonstrated in the example of two crossing shocks.

For the case $N_s = 1$ and $c_1 = 0$ in (7), the common POD is recovered and hence the sPOD delivers good low-dimensional representations for problems where the POD is applied successfully. Moreover, for the linear wave equation we have demonstrated that the sPOD is able to describe transport phenomena exhibiting sharp moving structures with the minimal number of modes and recovering the analytic solution up to machine precision. In contrast, for these examples the common POD performs poorly due to the slow singular value decay. In general, for examples where the snapshot matrix can be perfectly described by a decomposition of the form (12), the analytic solution is a fix point of the sPOD iteration. However, in general there may be more than one fix point and it is a priori unclear to which one the sPOD iteration converges. This nonuniqueness has been investigated numerically for the wave equation with two transported quantities by adding a redundant velocity $c_3 = 0$. Still the sPOD iteration reduces the error up to any given achievable tolerance, although it does not detect that there is a redundancy in the number of modes.

We have tested the sPOD algorithm also for more complex examples, where the analytic solution is unknown. For the example of two crossing shock waves in one dimension and for the vortex pair in two dimensions, the sPOD shows a better performance than the POD in number of modes needed to ensure a certain accuracy. Moreover, in the co-moving frames the sPOD modes lend themselves a clear and intuitive description as opposed to the POD modes in the lab frame.

The application of the sPOD depends on the transport velocities c_k . We have discussed methods for determining good values for the c_k in cases where they are not known from physical considerations, namely, for instance, by peak or front tracking or by a singular value maximization. In these cases the sPOD can be computed purely based on snapshot data. We have also investigated numerically the case when the velocities are slightly under- or overestimated. For the linear wave equation example with two moving quantities, the sPOD needs a few more modes (depending on how strong the velocities are disturbed) but the singular value decay is still much steeper compared with the POD. All in all for practical applications, the sPOD qualifies as an alternative to the POD for transport-dominated phenomena yielding a stronger singular value decay and physically more intuitive modes.

The results obtained so far motivate further development and generalization of the sPOD. In order to ensure local optimality, the decomposition (12) may be directly reformulated in terms of an optimization problem. This is likely to be more expensive than the algorithm proposed here but in return guarantees local optimality. Moreover, an extension to nonperiodic boundaries is necessary to be able to handle more realistic settings, e.g., with reflecting and nonreflecting boundaries. In addition, the considerations made so far only apply to the mode decomposition of the snapshot matrix. How the sPOD modes can be used in order to obtain a dynamic reduced-order model with an efficient offline/online decomposition needs to be investigated in the future.

REFERENCES

- R. ABGRALL, D. AMSALLEM, AND R. CRISOVAN, Robust model reduction by L¹-norm minimization and approximation via dictionaries: Application to nonlinear hyperbolic problems, Adv. Model. Simul. Eng. Sci., 3 (2016).
- [2] A. C. ANTOULAS, C. A. BEATTIE, AND S. GUGERCIN, Interpolatory model reduction of largescale dynamical systems, in Efficient Modeling and Control of Large-Scale Systems, J. Mohammadpour and K. M. Grigoriadis, eds., Springer, New York, 2010, pp. 3–58.
- [3] T. J. BARBER, M. H. AHMED, AND N. Y. ABDEL SHAFI, POD snapshot data reduction for periodic fluid flows, in Proceedings of the 43rd AIAA Aerospace Sciences Meeting and Exhibit, 2005.
- [4] M. BARRAULT, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, An 'empirical interpolation' method: Application to efficient reduced-basis discretization of partial differential equations, C. R. Math. Acad. Sci. Paris, 339 (2004), pp. 667–672.
- C. BEATTIE AND S. GUGERCIN, Interpolatory projection methods for structure-preserving model reduction, Systems Control Lett., 58 (2009), pp. 225–232.
- W. J. BEYN AND V. THÜMMLER, Freezing solutions of equivariant evolution equations, SIAM J. Appl. Dyn. Syst., 3 (2004), pp. 85–116.
- [7] C. BOGEY, N. DE CACQUERAY, AND C. BAILLY, A shock-capturing methodology based on adaptative spatial filtering for high-order non-linear computations, J. Comput. Phys., 228 (2009), pp. 1447–1465.
- [8] S. BOYD AND L. VANDENBERGHE, Convex Optimization, 7th ed., Cambridge University Press, Cambridge, UK, 2009.
- [9] N. CAGNIART, Y. MADAY, AND B. STAMM, Model Order Reduction for Problems with Large Convection Effects, Preprint HAL-01395571, 2016, https://hal.archives-ouvertes.fr.
- [10] K. CARLBERG, Adaptive h-refinement for reduced-order models, Internat. J. Numer. Methods Engrg., 102 (2015), pp. 1192–1210.
- [11] S. CHATURANTABUT AND D. C. SORENSEN, Nonlinear model reduction via discrete empirical interpolation, SIAM J. Sci. Comput., 32 (2010), pp. 2737–2764.
- [12] K. K. CHEN, J. H. TU, AND C. W. ROWLEY, Variants of dynamic mode decomposition: Boundary condition, Koopman, and Fourier analyses, J. Nonlinear Sci., 22 (2012), pp. 887–915.
- [13] P. G. CONSTANTINE AND G. IACCARINO, Reduced Order Models for Parameterized Hyperbolic Conservations Laws with Shock Reconstruction, Annual Research Briefs, Center for Turbulence Research, Stanford, CA, 2012.
- [14] P. DEHEUVELS AND G. V. MARTYNOV, A Karhunen-Loeve decomposition of a Gaussian process generated by independent pairs of exponential random variables, J. Funct. Anal., 255 (2008), pp. 2363–2394.
- [15] M. A. DIHLMANN AND B. HAASDONK, Certified PDE-constrained parameter optimization using reduced basis surrogate models for evolution problems, Comput. Optim. Appl., 60 (2015), pp. 753–787.
- [16] J. L. EFTANG, Reduced Basis Methods for Parametrized Partial Differential Equations, Ph.D. thesis, Norwegian University of Science and Technology, Trondheim, Norway, 2011.
- [17] F. FEDELE, O. ABESSI, AND P. J. ROBERTS, Symmetry reduction of turbulent pipe flows, J. Fluid Mech., 779 (2015), pp. 390–410.
- [18] K. FUJIMOTO AND J. M. A. SCHERPEN, Balanced realization and model order reduction for nonlinear systems based on singular value analysis, SIAM J. Control Optim., 48 (2010), pp. 4591–4623.
- [19] J.-F. GERBEAU AND D. LOMBARDI, Approximated Lax pairs for the reduced order integration of nonlinear evolution equations, J. Comput. Phys., 265 (2014), pp. 246–269.
- [20] M. A. GREPL, Certified reduced basis methods for nonaffine linear time-varying and nonlinear parabolic partial differential equations, Math. Models Methods Appl. Sci., 22 (2012), 1150015.
- [21] S. GUGERCIN, A. C. ANTOULAS, AND C. A. BEATTIE, A rational Krylov algorithm for optimal \mathcal{H}_2 model reduction, in Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems, 2006, pp. 1665–1667.
- [22] B. HAASDONK AND M. OHLBERGER, Efficient reduced models and a posteriori error estimation for parametrized dynamical systems by offline/online decomposition, Math. Comput. Model. Dyn. Syst., 17 (2011), pp. 145–161.
- [23] M. HINZE AND S. VOLKWEIN, Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: Error estimates and suboptimal control, in Dimension Reduction of Large-Scale Systems, P. Benner, D. C. Sorensen, and V. Mehrmann, eds., Springer, Berlin, 2005, pp. 261–306.

- [24] A. IOLLO AND D. LOMBARDI, Advection modes by optimal mass transfer, Phys. Rev. E, 89 (2014), 022923.
- [25] I. T. JOLLIFFE, Principal Component Analysis, Springer, New York, 1986.
- [26] M. LEMKE, A. MIEDLAR, J. REISS, V. MEHRMANN, AND J. SESTERHENN, Model reduction of reactive processes, in Active Flow and Combustion Control 2014, R. King, ed., Springer, Cham, Switzerland, 2015, pp. 234–262.
- [27] D. M. LUCHTENBURG, B. R. NOACK, AND M. SCHLEGEL, An Introduction to the POD Galerkin Method for Fluid Flows with Analytical Examples and MATLAB Source Codes, Technical Report 01/2009, Department of Fluid Dynamics and Engineering Acoustics, TU Berlin, 2009.
- [28] D. J. LUCIA, Reduced Order Modeling for High Speed Flows with Moving Shocks, Ph.D. thesis, Air Force Institute of Technology, 2001.
- [29] A. J. MAYO AND A. C. ANTOULAS, A framework for the solution of the generalized realization problem, Linear Algebra Appl., 425 (2007), pp. 634–662.
- [30] R. MOJGANI AND M. BALAJEWICZ, Lagrangian Basis Method for Dimensionality Reduction of Convection Dominated Nonlinear Flows, preprint, arXiv:1701.04343v1, 2017.
- [31] M. OHLBERGER AND S. RAVE, Nonlinear reduced basis approximation of parameterized evolution equations via the method of freezing, C. R. Math. Acad. Sci. Paris, 351 (2013), pp. 901–906.
- [32] M. RATHINAM AND L. R. PETZOLD, A new look at proper orthogonal decomposition, SIAM J. Numer. Anal., 41 (2003), pp. 1893–1925.
- [33] J. REISS, A family of energy stable, skew-symmetric finite difference schemes on collocated grids, J. Sci. Comput., 65 (2015), pp. 821–838.
- [34] J. REISS AND P. SCHULZE, Code for the Paper "The Shifted Proper Orthogonal Decomposition: A Mode Decomposition for Multiple Transport Phenomena," 2018, https://doi.org/ 10.5281/zenodo.1174271.
- [35] J. REISS AND J. SESTERHENN, A conservative, skew-symmetric finite difference scheme for the compressible Navier-Stokes equations, Comput. & Fluids, 101 (2014), pp. 208–219.
- [36] D. RIM, S. MOE, AND R. J. LEVEQUE, Transport Reversal for Model Reduction of Hyperbolic Partial Differential Equations, preprint, arXiv:1701.07529, 2017.
- [37] C. W. ROWLEY, I. G. KEVREKIDIS, J. E. MARSDEN, AND K. LUST, Reduction and reconstruction for self-similar dynamical systems, Nonlinearity, 16 (2003), pp. 1257–1275.
- [38] G. ROZZA, D. B. P. HUYNH, AND A. T. PATERA, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, Arch. Comput. Methods Eng., 15 (2008), pp. 229–275.
- [39] P. J. SCHMID AND J. L. SESTERHENN, Dynamic mode decomposition of numerical and experimental data, in Proceedings of the 61st Annual Meeting of the APS Division of Fluid Dynamics, 2008.
- [40] J. SESTERHENN AND A. SHAHIRPOUR, A Lagrangian Dynamic Mode Decomposition, preprint, arXiv:1603.02539, 2016.
- [41] T. TADDEI, S. PEROTTO, AND A. QUARTERONI, Reduced basis techniques for nonlinear conservation laws, ESAIM Math. Model. Numer. Anal., 49 (2015), pp. 787–814.
- [42] J. H. TU, Dynamic Mode Decomposition: Theory and Applications, Ph.D. thesis, Princeton University, Princeton, NJ, 2013.
- [43] J. H. TU AND C. W. ROWLEY, An improved algorithm for balanced POD through an analytic treatment of impulse response tails, J. Comput. Phys., 231 (2012), pp. 5317–5333.
- [44] K. URBAN AND A. T. PATERA, A new error bound for reduced basis approximation of parabolic partial differential equations, C. R. Math. Acad. Sci. Paris, 350 (2012), pp. 203–207.
- [45] G. WELPER, Transformed Snapshot Interpolation, preprint, arXiv:1505.01227v1, 2015.
- [46] K. WILLCOX AND J. PERAIRE, Balanced model reduction via the proper orthogonal decomposition, AIAA J., 40 (2002), pp. 2323–2330.