Journal of Computational Physics 230 (2011) 8573–8602

Contents lists available at SciVerse ScienceDirect

Journal of Computational Physics

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

An eigen-based high-order expansion basis for structured spectral elements

X. Zheng, S. Dong *

Center for Computational & Applied Mathematics, Department of Mathematics, Purdue University, United States

ARTICLE INFO

Article history:
Received 9 November 2010
Received in revised form 3 August 2011
Accepted 11 August 2011
Available online 26 August 2011

Keywords:
Expansion basis
High order methods
Spectral element
Simultaneous diagonalization

ABSTRACT

We present an eigen-based high-order expansion basis for the spectral element approach with structured elements. The new basis exhibits a numerical efficiency significantly superior, in terms of the conditioning of coefficient matrices and the number of iterations to convergence for the conjugate gradient solver, to the commonly-used Jacobi polynomial-based expansion basis. This basis results in extremely sparse mass matrices, and it is very amenable to the diagonal preconditioning. Ample numerical experiments demonstrate that with the new basis and a simple diagonal preconditioner the number of conjugate gradient iterations to convergence has essentially no dependence or only a very weak dependence on the element order. The expansion bases are constructed by a tensor product of a set of special one-dimensional (1D) basis functions. The 1D interior modes are constructed such that the interior mass and stiffness matrices are simultaneously diagonal and have identical condition numbers. The 1D vertex modes are constructed to be orthogonal to all the interior modes. The performance of the new basis has been investigated and compared with other expansion bases.

1. Introduction

In high-order approaches with spectral elements or *p*-finite elements, the computational domain is first partitioned using a number of elements, much like in low-order finite element methods. Within each element, however, an expansion of the field variables will further be performed, usually in terms of a set of high-order polynomial bases [21,28,17] or rational polynomial bases [24]. The selection of the high-order expansion basis functions within the elements directly influences the conditioning and sparsity of the resultant system matrix after discretization [4,14], which in turn influences the number of iterations to convergence in iterative solvers. The expansion basis therefore intimately influences the numerical efficiency and performance of the high-order methods. An ideal expansion basis would be expected to yield a linear algebraic system with a well-conditioned sparse matrix that can be solved using iterative solvers as efficiently as possible.

One of the earliest set of high-order basis functions aiming to improve the matrix sparsity and conditioning utilizes the integrals of Legendre polynomials [28]. This basis produces a diagonal block in the elemental stiffness matrix involving the interior modes. However, for unstructured elements it leads to an exponential increase in the condition number of the system matrix with respect to the element order [1]. Carnevali et al. [7] employed an orthogonalization process in the basis construction for triangles and tetrahedra, which resulted in local matrices better-conditioned than those of [28]. The growth of the condition number with respect to the element order is however still exponential.

* Corresponding author.
E-mail address: sdong@math.purdue.edu (S. Dong).

0021-9991/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2011.08.009

Jacobi polynomials have been employed in the construction of several of the more recent expansion bases. In [32] the expansion basis for triangular elements involves orthogonal edge functions of the same edge and orthogonal face functions based on Jacobi polynomials. Another basis based on Jacobi polynomials employs generalized tensor products [26,31,17], and has been very successful in computational fluid dynamics (CFD). We will refer to this basis as the Sherwin–Karniadakis basis (or SK-Basis for abbreviation) hereafter in this paper. It provides a unified treatment for all commonly encountered element types in two and three dimensions. For unstructured elements it is based on a modification of the Dubiner's construction of orthogonal polynomials in a triangle [13]. For structured elements this basis is equivalent to a modification of the basis of [28]. The Sherwin–Karniadakis basis is one of the most commonly used high-order approaches in CFD, and more recently also in computational structural mechanics [11]. A recent effort to extend the Sherwin–Karniadakis basis is presented in [5], in which two parameters have been introduced through the integrated Jacobi polynomials. However, the best basis among the extensions, in terms of sparsity and conditioning, corresponds to the original Sherwin–Karniadakis basis. In [6] a set of basis functions based on the tensor products of Jacobi polynomials is presented employing barycentric coordinates, where the indices of the Jacobi polynomials are chosen to improve upon the sparsity pattern of the elemental matrices. These bases have been compared with several existing ones, and are observed to be comparable to the basis of [28] in terms of conditioning. It is also observed in [6] that in terms of conditioning the Sherwin–Karniadakis basis is comparable or markedly better than several other existing bases.

In [2], the basis functions of [28] for triangular and tetrahedral elements are modified through an orthogonalization of the face/interior modes. This reduces the coupling of these modes, and as a result the growth rate of the condition number with respect to the element order is no longer exponential [2], unlike that of [28,7]. In the expansion basis for a triangle presented in [1], mutually orthogonal edge functions of the same edge, and orthogonal face functions, have been employed in barycentric coordinates, and the edge functions are also orthogonalized with respect to the face functions. The basis is observed to be considerably better-conditioned than the basis of [28]. The numerical tests in [1] also show that, the sparsity of the elemental stiffness matrix with the Sherwin–Karniadakis basis has a strong favorable influence on the conditioning of the global matrix, and that the Sherwin–Karniadakis basis often leads to fewer iterations in the iterative solver than other bases.

The expansion bases discussed above are often referred to as modal bases in the literature. One can also construct expansion bases associated with a set of nodal points, that is, the nodal expansion bases. The commonly used nodal points are the Gauss–Lobatto points [21]. For unstructured elements, the nodal points can be chosen by a minimization of an electrostatic potential [16] or the Lebesgue constant (Fekete points) [29]. Certain types of partial differential equations may require vector-valued approximation functions. For high-order vector finite elements, we refer to [3] for a review and see also the references therein.

By distinguishing the boundary modes (i.e. 2D/3D vertex/edge modes, and also face modes in 3D) from the interior modes (also referred to as the bubble modes) in the expansion basis, one can order the degrees of freedom in an appropriate fashion, and perform a Schur-complement operation to condense out the contribution of the interior modes from the system matrix. This process is also referred to as static condensation in the literature. The linear algebraic problem after discretization will therefore involve the solution of the boundary modes with the Schur-complemented system matrix, usually using an iterative solver, followed by a successive local elemental solution of the interior modes on each element. Schur complement significantly improves the property of the linear algebraic system and the performance of the linear equation solvers. It therefore should always be performed by default for element-based high-order methods. We note that this is indeed the case in practical applications [8–10,12]. One also notes that in practical implementations only local elemental Schur complement needs to be performed.

In light of the above, one can appreciate that it is the property (e.g. conditioning, sparsity) of the Schur-complemented system matrix that critically influences the performance of the method. When investigating the performance of expansion bases, it is therefore more pertinent to consider the Schur-complemented system matrix, rather than the full matrix (with interior modes included). This consideration has a significant implication on the construction of expansion bases. One can easily show that, given a set of expansion basis, if one modifies the interior modes only (assuming identical polynomial space) to obtain a set of new basis functions while keeping the same boundary modes as before, then the Schur-complemented system matrix (assuming a bilinear form) will remain identical as before. In other words, a modification of the interior modes only in an expansion basis will not affect the numerical performance of the Schur-complemented system. Therefore, the construction of the boundary modes plays a critical role in the performance of the expansion bases. We can therefore infer that in several existing expansion bases some constructions will have no or only minor effect on the performance of the bases. For example, in [2] the orthogonalization of the face modes for triangular elements will have no effect on the Schur-complemented system matrix. The orthogonalization of the interior modes for the tetrahedral elements therein will not affect the system performance, while the orthogonalization of the face modes for the tetrahedral elements will. The same point also applies to the orthogonalization of the face modes for the triangular elements in [1]. We note that both [2] and [1] have considered the full system matrix in the tests.

High sparsity in the system matrix is a desirable property for expansion bases, and it tends to positively influence the conditioning of the resulting algebraic problem. For example, it is observed from numerical tests that, for the Sherwin–Karniadakis basis, the conditioning of the global matrix is influenced more strongly by the sparsity than by the conditioning of the elemental matrix [1]. Because of sparsity, even though the condition number of the elemental matrix for the Sherwin–Karniadakis basis is higher than that of the basis of [1], the condition number for the global matrix with the

Sherwin–Karniadakis basis is smaller (see [1]). Considerations of sparsity have also prompted the use of orthogonalization in several expansion bases [7,32,2,5,1].

The sparsity and the conditioning of the system matrices can be improved with simultaneous diagonalization techniques. It is known that two symmetric positive definite (SPD) matrices can be simultaneously diagonalized (the SPD condition can be relaxed; see [15]). Simultaneous diagonalization technique has been widely employed in the automatic control community (see e.g. [19,18]), where a contragredient transform simultaneously makes the controllability and observability Gram matrices diagonal. The use of simultaneous diagonalization in solving partial differential equations (PDE) has been explored only very recently. The first use is probably in [23] for a Legendre–Galerkin spectral method, in which the mass and stiffness (Laplace operator) matrices are made simultaneously diagonal. This transform is termed "Fourierization" in [23]. Very recently, simultaneous diagonalization has been used in [27] for generating the interior bubble functions in high-order finite elements.

In this paper, we present a high-order expansion basis for the spectral element approach. The new expansion basis is based on the construction of a special set of high-order basis functions in one dimension (1D). For the construction of the 1D interior modes, we exploit the idea of simultaneous diagonalization and present an algorithm that makes the elemental mass and stiffness matrices for the submatrix involving only the interior modes simultaneously diagonal. The construction algorithm involves a parameter k. In our expansion basis, we choose the 1D interior modes corresponding to the construction with parameter value $k = \frac{1}{2}$. This corresponds to a set of "balanced" eigen-functions, in the sense that the elemental mass and stiffness matrices for the submatrix involving only the interior modes are both diagonal and have identical condition numbers. We note that the interior modal functions of [27] are equivalent to a particular case, $k = 1$, in the current construction, which leads to an elemental stiffness matrix (interior submatrix) that is an identity matrix and a diagonal elemental mass matrix (interior submatrix).

The more crucial construction, in terms of the effects on the performance, lies in the uncommon 1D vertex modes we have adopted in the current expansion basis. These 1D vertex modes are constructed such that they are orthogonal to all the 1D interior modes. This set of 1D basis functions (vertex and interior modes) gives rise to a set of unique boundary modes in higher dimensions. As noted in the foregoing discussions, the construction of the boundary modes plays a crucial role in the conditioning of the Schur-complemented system matrix. The boundary modes in our basis as constructed above result in especially low condition numbers in the global mass matrix, and a well-conditioned Schur-complemented system matrix with clustered eigenvalues.

Since the Sherwin–Karniadakis basis is one of the most commonly used expansion bases and it has been shown by a number of studies to yield very favorable performance compared to other bases (see [6,1,5] and the discussions in foregoing paragraphs), we will mostly compare the performance of the current expansion basis with the Sherwin–Karniadakis basis in this paper. Our results show that the current basis exhibits a significantly superior efficiency in terms of the conditioning of the system matrix and the number of iterations to convergence with the conjugate gradient (CG) iterative solver.

In the current paper we will focus on the construction and performance of the new expansion basis for structured spectral elements. By structured elements we refer to quadrilateral elements in two dimensions and hexahedral elements in three dimensions. Extension of the current construction to the unstructured elements (e.g. triangle, tetrahedron) is highly nontrivial for the boundary modes, especially for the unconventional vertex modes adopted in the current basis. The constructions for unstructured elements will be considered in a subsequent paper.

2. Eigen-based high-order expansion basis

We present the construction of the new expansion basis in this section. We first discuss how to construct the high-order basis functions in one dimension, which constitute the building components for the basis in higher dimensions. Construction for structured spectral elements in two and three dimensions (2D, 3D) and their properties will be discussed afterwards.

2.1. Expansion basis functions in one dimension

Consider the standard element in one dimension, $-1 \leqslant \xi \leqslant 1$, where ξ denotes the coordinate. We distinguish the two vertex modes and the interior modes. By "modes" we interchangeably refer to the basis functions. By convention (see [28,17]), the interior modes vanish at the two end points $\xi = \pm 1$, while a vertex mode takes the unit value at one end point and vanishes at the other. The vertex modes and the interior modes are constructed in different fashions in our basis.

2.1.1. Interior modes

For the construction of the interior modes, we exploit the idea of simultaneous diagonalization (see [19,18,23,27]), such that with the new basis functions the elemental mass and stiffness matrices involving the interior modes (hereafter referred to as the interior matrices) will simultaneously become diagonal.

Consider the polynomial space of the interior modes, $V_I = \{f(\xi) \in P_N : f(-1) = f(1) = 0\}$, where N is the element order and P_N represents the set of polynomials of degree at most N. We start with an arbitrary basis for V_I, e.g. with polynomials of the form $(1-\xi^2)\xi^{p-1}$ $(p = 1,\ldots,N-1)$. Let $\psi_p(\xi)$ $(p = 1,\ldots,N-1)$ denote this set of initial basis functions. The interior mass matrix

with the initial basis is $\mathbf{M} = [M_{pq}]_{(N-1)\times(N-1)}$ with $M_{pq} = \int_{-1}^{1} \psi_p \psi_q d\xi$. The interior stiffness matrix is $\mathbf{K} = [K_{pq}]_{(N-1)\times(N-1)}$ with $K_{pq} = \int_{-1}^{1} \frac{d\psi_p}{d\xi}\frac{d\psi_q}{d\xi} d\xi$. Note that both the mass and stiffness matrices are symmetric positive definite in the polynomial space V_I.

Consider a set of new basis functions $\phi_p(\xi)$ ($p = 1,\ldots,N-1$) given by

$$\phi_p(\xi) = \sum_{q=1}^{N-1} y_{pq}\psi_q(\xi), \tag{1}$$

where y_{pq} are the coefficients to be determined. Under the new basis, the interior mass matrix $\mathbf{M}'$ and the interior stiffness matrix $\mathbf{K}'$ are respectively given by

$$\mathbf{M}' = \mathbf{Y}\mathbf{M}\mathbf{Y}^T, \quad \mathbf{K}' = \mathbf{Y}\mathbf{K}\mathbf{Y}^T, \tag{2}$$

where $\mathbf{Y} = [y_{pq}]_{(N-1)\times(N-1)}$ is the matrix of coefficients and is non-singular, and the superscript in $(\cdot)^T$ denotes transposition. We would like to find $\mathbf{Y}$ such that $\mathbf{M}'$ and $\mathbf{K}'$ are both diagonal.

We construct the coefficient matrix $\mathbf{Y}$ as follows. Since $\mathbf{M}$ is symmetric positive definite, it can be diagonalized,

$$\mathbf{X}^T\mathbf{M}\mathbf{X} = \mathbf{\Lambda}_M, \tag{3}$$

where $\mathbf{X}$ is the orthogonal matrix formed by the eigenvectors of $\mathbf{M}$, and $\mathbf{\Lambda}_M$ is a diagonal matrix of the eigenvalues of $\mathbf{M}$, which are all positive. Next we construct a matrix

$$\mathbf{L} = \left(\mathbf{X}\mathbf{\Lambda}_M^{-\frac{1}{2}}\right)^T \mathbf{K}\left(\mathbf{X}\mathbf{\Lambda}_M^{-\frac{1}{2}}\right), \tag{4}$$

where the exponential in $\mathbf{\Lambda}_M^{-\frac{1}{2}}$ is understood to be element-wise operations on the main diagonal. One can note that the matrix $\mathbf{L}$ is also symmetric positive definite. As a result, $\mathbf{L}$ can be diagonalized,

$$\mathbf{Z}^T\mathbf{L}\mathbf{Z} = \mathbf{\Lambda}_S, \tag{5}$$

where $\mathbf{Z}$ is the orthogonal matrix formed by the eigenvectors of $\mathbf{L}$, and $\mathbf{\Lambda}_S$ is the diagonal matrix of the eigenvalues of $\mathbf{L}$, which are all positive. We set the coefficient matrix $\mathbf{Y}$ as

$$\mathbf{Y} = \left(\mathbf{X}\mathbf{\Lambda}_M^{-\frac{1}{2}}\mathbf{Z}\mathbf{\Lambda}_S^{-\frac{k}{2}}\right)^T, \tag{6}$$

where k is a parameter. It is straightforward to verify that,

$$\mathbf{Y}\mathbf{M}\mathbf{Y}^T = \mathbf{\Lambda}_S^{-k}, \quad \mathbf{Y}\mathbf{K}\mathbf{Y}^T = \mathbf{\Lambda}_S^{1-k}. \tag{7}$$

Therefore, with the coefficient matrix given by Eq. (6) we can construct a set of new basis functions using Eq. (1), and under this set of new basis functions the interior mass and stiffness matrices are both diagonal.

Let us consider the effect of the parameter k in the construction (Eq. (6)) on the resulting matrices. Based on Eq. (7), the construction with $k = 0$ results in an identity interior mass matrix, and $k = 1$ results in an identity interior stiffness matrix. The choice of the parameter value for k, and correspondingly the basis functions $\phi_p(\xi)$, have a significant influence on the condition numbers of the elemental interior matrices (see below), as well as on the condition numbers of the Schur-complemented global matrix in higher dimensions. In our expansion basis, we employ the functions $\phi_p(\xi)$ corresponding to $k = 1/2$ as the 1D interior modes. This choice leads to identical condition numbers for the elemental interior mass and interior stiffness matrices; see Eq. (7). At the same element order, the condition number of our elemental interior matrices is the square root of the condition number of the non-identity elemental interior matrix of the case $k = 0$ or $k = 1$. Our choice corresponds to the "balanced" contragredient transform in the automatic control literature [18].

In the basis of [27] Solin and Vejchodsky has solved a generalized eigenvalue problem, which results in an elemental interior stiffness matrix which is an identity matrix, and a diagonal interior mass matrix. The basis of [27] has a correspondence to the case with parameter $k = 1$ in the current construction.

In Fig. 1 we show the condition number of the elemental interior mass and stiffness matrices (i.e. the submatrix involving only the interior modes) as a function of the element order in 1D for the current basis, the Sherwin–Karniadakis basis [17], and the basis of [27]. Note that for the current basis and the basis of [27] both the interior mass and stiffness matrices are diagonal matrices. For the Sherwin–Karniadakis basis, the interior stiffness matrix is a diagonal matrix while the interior mass matrix is not. One can observe that, for the basis of [27], the interior stiffness matrix (an identity matrix) has a perfect condition number (unit value), while the condition number of the interior mass matrix increases significantly with increasing element order and it is considerably larger than those with the current and the Sherwin–Karniadakis bases. For the current basis, the interior mass and stiffness matrices have identical condition numbers, which increase with respect to the element order quite moderately compared to the interior mass matrix of the basis of [27]. Compared to the Sherwin–Karniadakis basis, the condition numbers with the current basis are lower for element orders 9 or below, and lie between those of the interior stiffness and mass matrices of the Sherwin–Karniadakis basis for element orders beyond 9.

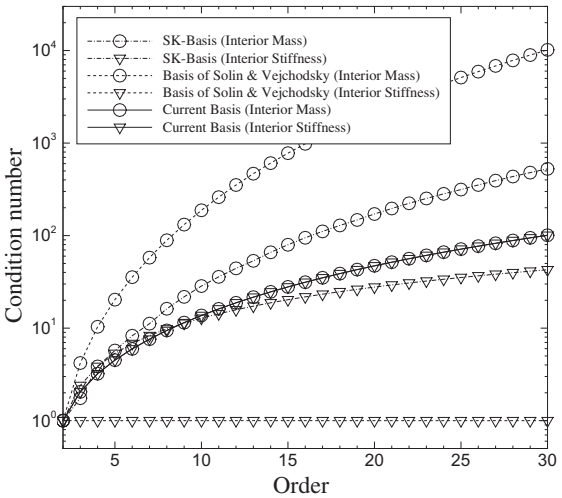

Fig. 1. Comparison of condition numbers of the elemental interior mass and interior stiffness matrices in 1D as a function of the element order with the Sherwin–Karniadakis basis [17], the basis of Solin and Vejchodsky [27], and the current basis. "SK-Basis" stands for the Sherwin–Karniadakis basis.

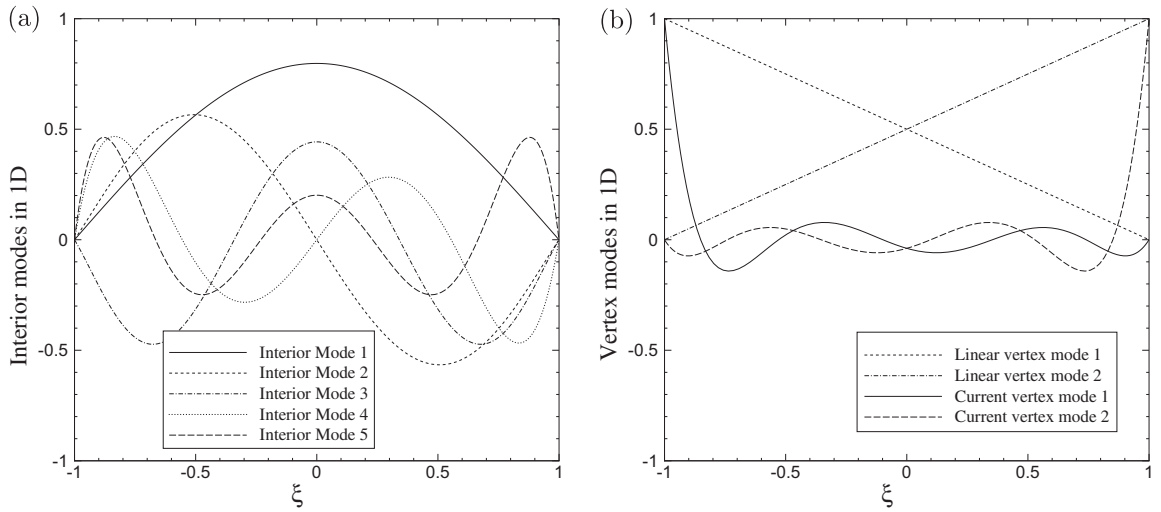

Fig. 2. 1D basis functions (element order 6): (a) Profiles of the 1D interior modes of the current basis. (b) Profiles of the linear vertex modes and the 1D vertex modes of the current basis.

In Fig. 2(a) we plot the distribution of the current interior modes $\phi_p(\xi)$ in the standard element $-1 \leqslant \xi \leqslant 1$ for an element order $N = 6$. We note that these functions are either symmetric or anti-symmetric with respect to the origin $\xi = 0$, which is necessary for and greatly simplifies the implementation of boundary continuity between adjacent elements in higher dimensions. We further note that these functions are all polynomials of degree N or $N-1$, where N is the element order. For example, the profile of the first interior mode in Fig. 2(a) resembles that of a quadratic polynomial, while in actuality this mode is a polynomial of degree 6 restricted to the interval $-1 \leqslant \xi \leqslant 1$. The implication is that the current expansion basis is not a hierarchical basis, unlike the Sherwin–Karniadakis basis.

Remarks. Note that the eigenvectors in the matrices **X** and **Z** of Eq. (6) can not be uniquely determined even after normalization. For example, for distinct eigenvalues the eigenvectors can only be uniquely determined to within a sign, and for a repeated eigenvalue the corresponding eigenvectors can be rotated arbitrarily in the corresponding eigenspace. So the coefficient matrix **Y** (and therefore the basis functions $\phi_p(\xi)$ in Eq. (1)), which makes the interior mass and stiffness matrices simultaneously diagonal, is not unique. In the implementations, once we have computed the set of 1D interior modes by choosing the eigenvectors and determining the matrix **Y**, we will use this set as building blocks for constructing basis functions in higher dimensions throughout the application.

2.1.2. Vertex modes

We construct the vertex modes as follows. Note that the interior modes $\phi_p(\xi)$ ($p = 1,\ldots,N-1$, N being the element order) in Eq. (1) are orthogonal to one another. For construction of the vertex mode $\phi_0(\xi)$ corresponding to $\xi = -1$, we start with the linear vertex function,

$$\psi_0(\xi) = \frac{1-\xi}{2}. \tag{8}$$

Then we perform a Gram–Schmidt process to orthogonalize $\psi_0(\xi)$ with respect to the interior modes $\phi_p(\xi)$ ($p = 1,\ldots,N-1$),

$$\phi_0(\xi) = \psi_0(\xi) - \sum_{p=1}^{N-1} \frac{\langle \psi_0, \phi_p \rangle}{\langle \phi_p, \phi_p \rangle} \phi_p(\xi), \tag{9}$$

where the inner product is defined as $\langle f, g \rangle = \int_{-1}^{1} f(\xi) g(\xi) d\xi$. We employ the resulting function as the vertex mode $\phi_0(\xi)$ in our expansion basis. Similarly, for the construction of the vertex mode $\phi_N(\xi)$ corresponding to $\xi = 1$, we start with the linear vertex function,

$$\psi_N(\xi) = \frac{1+\xi}{2} \tag{10}$$

and perform a Gram–Schmidt process to orthogonalize $\psi_N(\xi)$ with respect to all the interior modes $\phi_p(\xi)$,

$$\phi_N(\xi) = \psi_N(\xi) - \sum_{p=1}^{N-1} \frac{\langle \psi_N, \phi_p \rangle}{\langle \phi_p, \phi_p \rangle} \phi_p(\xi). \tag{11}$$

The resulting function is used as the vertex mode $\phi_N(\xi)$.

In Fig. 2(b) we plot the two 1D vertex modes $\phi_0(\xi)$ and $\phi_N(\xi)$ of our expansion basis (Eqs. (9) and (11)), together with the linear vertex modes $\psi_0(\xi)$ and $\psi_N(\xi)$, in the standard element $-1 \leqslant \xi \leqslant 1$. This is for an element order $N = 6$, in accordance with the interior modes in Fig. 2(a). Note that, unlike the linear vertex functions, the vertex modes in the current basis are both polynomials of degree N.

The 1D vertex modes as constructed above are crucial to the performance of our expansion basis in higher dimensions. These modes directly contribute to the 2D/3D vertex and edge modes (and also the face modes in 3D), and have a significant impact on the conditioning of the Schur-complemented system matrix. As we have noted in Section 1, given a set of expansion basis functions, a change solely in the interior modes (while keeping the boundary modes the same as before) will have no effect on the Schur-complemented system matrix. Our construction of the 1D vertex modes described above results in especially low condition numbers for the mass matrix, and tightly clustered eigenvalues for the Schur-complemented system matrix in higher dimensions, which will be demonstrated in subsequent sections.

The vertex modes $\phi_0(\xi)$ and $\phi_N(\xi)$ in Eqs. (9) and (11), and the interior modes $\phi_p(\xi)$ ($p = 1,\ldots,N-1$) given by Eq. (1), together form the current expansion basis in one dimension. In Fig. 3 we show the sparsity patterns of the elemental mass matrix (Fig. 3(a)) and stiffness matrix (Fig. 3(b)) of the current basis in 1D. The degrees of freedom are ordered such that the two vertex modes are followed by the interior modes. One can observe that the submatrices associated with the interior modes are diagonal for both the mass and stiffness matrices. There exists no coupling between the vertex modes and the interior modes in the mass matrix due to the Gram–Schmidt orthogonalization.

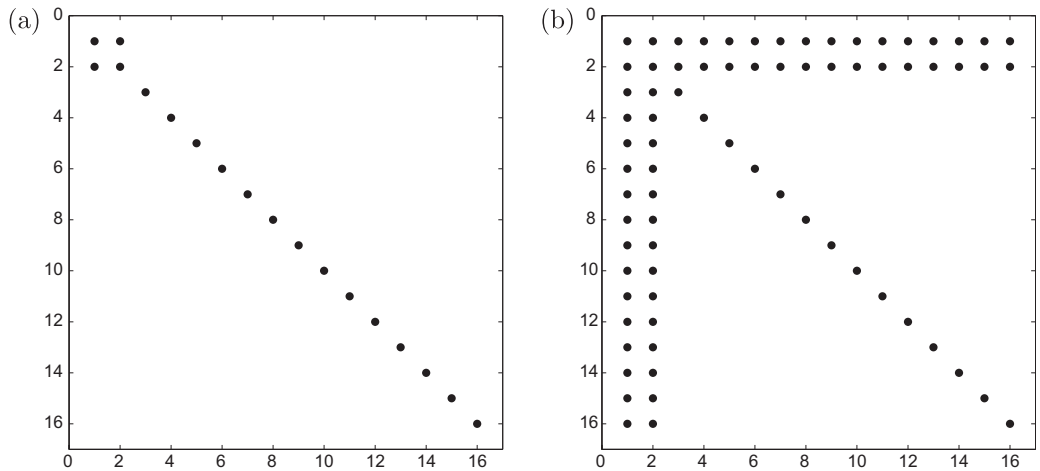

Fig. 3. Sparsity pattern of the 1D elemental mass matrix (a) and stiffness matrix (b) of the current expansion basis for an element order 15.

Remarks. The current expansion basis functions are numerically constructed polynomials, and the explicit analytic expressions for the polynomial forms are unknown. This is quite different from the commonly-used existing expansion bases. This fact, however, does not pose difficulty or induce overhead in the manipulations of these basis functions (e.g. in forward/backward transforms between modal/physical spaces, or evaluation of a function when represented with the new basis) once they are computed. The key is to note that the current basis functions are all polynomials of degree at most N. Therefore Gaussian quadratures equally apply to the current basis. Once the coefficient matrix $\mathbf{Y}$ is computed from Eq. (6), we can compute the values of current 1D basis functions on the 1D quadrature points using Eqs. (1), (9) and (11). Then the arbitrary initial basis that is involved in the computation of the current interior modes can be discarded. Subsequent manipulations of the current basis for the forward/backward transforms between modal/physical spaces and the evaluation of a function represented using the current basis will not involve the coefficient matrix $\mathbf{Y}$ or the translation between the arbitrary initial basis and the current basis.

2.2. Expansion basis for structured elements in higher dimensions

We next construct the expansion bases for structured spectral elements in higher dimensions (quadrilateral elements in 2D and hexahedral elements in 3D) based on the 1D basis functions defined in the previous subsection. The construction is conceptually straightforward, but special care needs to be taken on certain issues, such as the mapping between the physical and standard domains due to the uncommon vertex modes employed in the current basis. Similar to other expansion bases, we will employ a tensor product of the 1D modes to construct the expansion basis functions in higher dimensions.

2.2.1. Definitions of basis functions in 2D and 3D

We define the expansion basis functions only in the standard domain (or standard element). In practice, elements in physical domains will first be mapped to the standard domain, and then the operations involving the basis functions will be performed. In 2D the standard domain Ω_{st}^{2D} for a quadrilateral element is defined as the square

$$\Omega_{st}^{2D} = \{(\xi_1, \xi_2) : -1 \leqslant \xi_1 \leqslant 1,\ -1 \leqslant \xi_2 \leqslant 1\}$$

and in 3D the standard domain Ω_{st}^{3D} for a hexahedral element is defined as the cube

$$\Omega_{st}^{3D} = \{(\xi_1, \xi_2, \xi_3) : -1 \leqslant \xi_1 \leqslant 1,\ -1 \leqslant \xi_2 \leqslant 1, -1 \leqslant \xi_3 \leqslant 1\}.$$

In the above definitions we have used ξ_i ($i = 1,2,3$) to denote the coordinates in the standard domains.

For the quadrilateral element, we consider the polynomial space

$$V^{2D}(\Omega_{st}^{2D}) = \text{span}\left\{\xi_1^p \xi_2^q :\ (\xi_1, \xi_2) \in \Omega_{st}^{2D},\ 0 \leqslant p \leqslant N_1,\ 0 \leqslant q \leqslant N_2\right\},$$

where N_1 and N_2 are the element orders in the two directions (we assume $N_1 = N_2$ for simplicity). We employ a set of expansion basis functions, φ_{pq} ($0 \leqslant p \leqslant N_1, 0 \leqslant q \leqslant N_2$), for V^{2D} that are defined by

$$\varphi_{pq}(\xi_1, \xi_2) = \phi_p(\xi_1)\phi_q(\xi_2), \tag{12}$$

where ϕ_p are the 1D basis functions defined by Eqs. (9), (11) and (1). To facilitate the implementation of boundary conditions, we distinguish the boundary modes (vertex or edge modes) from the interior modes. The vertex modes are identified by φ_{pq} ($p = 0$ or N_1, $q = 0$ or N_2), the edge modes by φ_{pq} ($p = 0$ or N_1, $1 \leqslant q \leqslant N_2 - 1$) or ($q = 0$ or N_2, $1 \leqslant p \leqslant N_1 - 1$), and the interior modes by φ_{pq} ($1 \leqslant p \leqslant N_1 - 1, 1 \leqslant q \leqslant N_2 - 1$).

For the hexahedral element, we consider the polynomial space

$$V^{3D}(\Omega_{st}^{3D}) = \text{span}\left\{\xi_1^p \xi_2^q \xi_3^r :\ (\xi_1, \xi_2, \xi_3) \in \Omega_{st}^{3D},\ 0 \leqslant p \leqslant N_1,\ 0 \leqslant q \leqslant N_2,\ 0 \leqslant r \leqslant N_3\right\},$$

where N_i ($i = 1,2,3$) are the element orders in the three directions (Assume $N_1 = N_2 = N_3$ for simplicity). We employ a set of basis functions, φ_{pqr} ($0 \leqslant p \leqslant N_1, 0 \leqslant q \leqslant N_2, 0 \leqslant r \leqslant N_3$), for V^{3D} defined by

$$\varphi_{pqr}(\xi_1, \xi_2, \xi_3) = \phi_p(\xi_1)\phi_q(\xi_2)\phi_r(\xi_3). \tag{13}$$

Analogous to 2D, we can distinguish the boundary modes (vertex modes, edge modes, and face modes) from the interior modes. The vertex modes are identified by those with all three indices taking the boundary values (0 or N_i, $i = 1,2,3$). The edge modes are those with exactly two indices taking the boundary values. The face modes are those with exactly one index taking the boundary values. The interior modes are identified by those with no index taking the boundary values.

In Fig. 4 we show a diagram (contour plot) of our expansion basis functions for a 2D quadrilateral element with an element order 5. The outer layer of the plots is for the vertex modes (those at the four corners) and the edge modes, and the rest are for the interior modes. The restriction of the edge modes onto the corresponding edge is a curve that is either symmetric or anti-symmetric with respect to the mid-point of the edge, which enables the implementation of the continuity condition across element boundaries by a simple sign change for the anti-symmetric modes when the local coordinates are in opposite directions across the boundary. The current expansion basis therefore ensures globally continuous expansions.

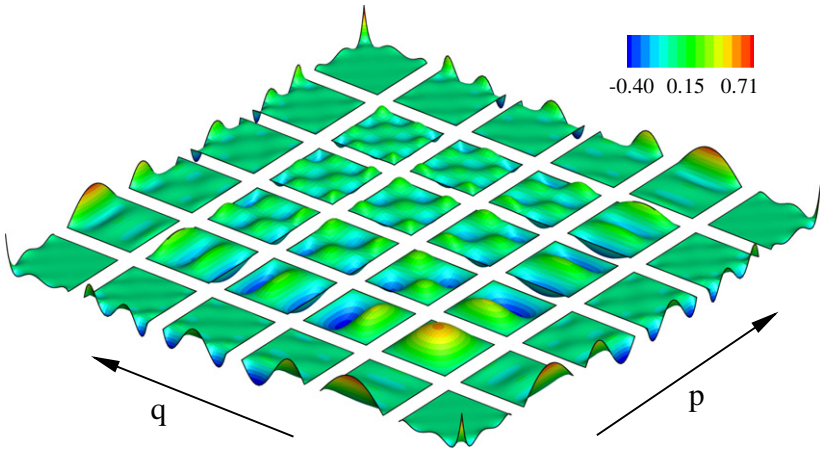

Fig. 4. Profiles (contours) of the current expansion basis functions for a 2D quadrilateral element with an element order 5.

Remarks. Because the current basis is based on tensor products (we consider structured elements here), the fast numerical technique "sum factorization" [20,17] also applies to the current basis. When a function is represented with the current basis, with sum factorization the function evaluation on all the quadrature points in multiple dimensions amounts to several successive matrix multiplications involving the matrix of values of 1D modes on 1D quadrature points and the matrix of expansion coefficients of the function. The cost for function evaluation with the current basis is the same as that of other tensor product-based expansion bases such as the Sherwin–Karniadakis basis.

2.2.2. Mapping between standard and physical elements

The expansion basis functions are defined only in the standard domains. For a general element in physical space we need to map the standard element to the physical element. The choice of the vertex modes (Eqs. (9) and (11)) in our expansion basis presents a slight complication for this mapping.

The mapping between the standard and physical elements must be a one-to-one function. The simplest of one-to-one mappings is a linear function in 1D, and the tensor product of a linear function in higher dimensions. For an isoparametric representation, if the vertex modes are linear functions or tensor products of linear functions, they can be used for the mapping between the standard and physical elements.

In our expansion basis, the vertex modes defined by Eqs. (9) and (11) are not linear functions. With the isoparametric form, in order to represent a linear mapping between the standard and physical domains in 1D, we need to employ all the vertex modes and the interior modes in the representation, that is,

$$x = \frac{x_B - x_A}{2}\xi + \frac{x_A + x_B}{2} = \hat{x}_0\phi_0(\xi) + \hat{x}_N\phi_N(\xi) + \sum_{p=1}^{N-1}\hat{x}_p\phi_p(\xi), \tag{14}$$

where $x \in [x_A, x_B]$ is the coordinate of the physical element. In the above equation $\hat{x}_p$ are the expansion coefficients, $\hat{x}_0 = x_A$, $\hat{x}_N = x_B$, and $\hat{x}_p$ ($p = 1,\ldots,N-1$) are determined by a Galerkin projection. In higher dimensions, this would require the use of all the boundary and interior modes in the isoparametric representation of the standard/physical domain mapping. This will complicate the implementation of the current expansion basis.

In order to simplify the implementation with the current expansion basis, in 1D we will directly use the linear vertex modes in Eqs. (8) and (10) for the mapping between the standard domain and the physical domain, that is,

$$x = \frac{x_B - x_A}{2}\xi + \frac{x_A + x_B}{2} = x_A\psi_0(\xi) + x_B\psi_N(\xi). \tag{15}$$

On the other hand, for the expansion of a function defined on the physical element, we will employ the set of expansion bases ϕ_p ($p = 0,\ldots,N$) defined by Eqs. (9), (11) and (1). In higher dimensions, for elements with straight sides, we will employ the tensor products of the 1D linear vertex functions (Eqs. (8) and (10)) for the mapping between the standard element and the physical element, while for the function expansions we use the set of basis functions defined by Eq. (12) or (13).

For elements with curved sides, we follow a convention from [17], namely, we will allow only one side to be curved per element. Let us use the 2D quadrilateral element for illustration of the domain mapping in this case. For domain mapping with the current basis we use an idea similar to that outlined in [17]. In addition to the tensor products of the linear vertex functions, some special edge modes (or face modes in 3D) corresponding to the curved side will also be involved in the mapping between physical and standard domains. The special edge modes used in the domain mapping is slightly different than the edge modes in the current basis as defined in Eq. (12). The special edge modes corresponding to the curved side are

constructed as the tensor products of the current 1D interior modes (along the direction of curved edge) and the 1D linear vertex mode (along the other direction). These special edge modes used for domain mapping for the current basis are actually the edge modes of another basis which we will refer to as "Basis-LV" and will be defined shortly in the next subsection.

2.2.3. Sparsity of mass and stiffness matrices

We next look into the sparsity of the mass and stiffness matrices resulting from the current basis in two and three dimensions. A typical entry in the mass matrix is defined by $M_{pq} = \int \varphi_p \varphi_q d\mathbf{x}$, where $\varphi_p(\mathbf{x})$ and $\varphi_q(\mathbf{x})$ denote the basis functions in 2D or 3D, and the 2D or 3D indices of the basis functions have been arranged into a 1D linear index, and $p, q = 0, \ldots, N_t - 1$ with N_t being the total degrees of freedom. A typical entry in the stiffness matrix is defined by $K_{pq} = \int \nabla\varphi_p \cdot \nabla\varphi_q d\mathbf{x}$.

Let us first consider the local elemental matrices. In Fig. 5(a) and (b) we show the sparsity patterns of the elemental mass matrix and stiffness matrix, respectively, for a 2D quadrilateral element in the standard domain with the current expansion basis. This is for an element order 10. A non-zero entry in the matrix is represented by a dot in the plots. The degrees of freedom are ordered such that the vertex modes are followed by the edge modes, and then by the interior modes. For comparison, in Fig. 5(c) and (d) we have shown the sparsity patterns of the elemental mass and stiffness matrices for another expansion basis which we will refer to as "Basis-LV". The basis functions in Basis-LV are constructed by the tensor products of a set of 1D basis functions that consist of the 1D interior modes the same as in the current basis (Eq. (1)) and the 1D linear vertex modes (Eqs. (8) and (10)). Note that the difference between the current expansion basis and Basis-LV in two and three dimensions lies in the boundary modes. The current basis employs the 1D Gram–Schmidt orthogonalized vertex modes (Eqs. (9) and (11)), while Basis-LV employs the 1D linear vertex modes in the constructions. The purpose for the comparisons with Basis-LV, here and in subsequent sections, is to demonstrate the crucial effect of the uncommon 1D vertex modes employed in the current expansion basis on the numerical efficiency. Also for comparison, in Fig. 5(e) and (f) we show the sparsity patterns of the elemental mass and stiffness matrices for the Sherwin–Karniadakis basis [17].

One can observe that the current expansion basis results in an extremely sparse elemental mass matrix. The 4×4 non-zero matrix block on the top-left corner of the mass matrix (Fig. 5(a)) corresponds to the 4 vertex modes of the quadrilateral element. For any particular edge mode, the only modes it is non-orthogonal to are the corresponding edge mode residing on the opposite side of the quadrilateral element and itself. This gives rise to a non-zero upper or lower diagonal line in the matrix block corresponding to the edge modes. The mass submatrix corresponding to the interior modes is diagonal. For the current basis, there exist only 169 non-zero entries in the elemental mass matrix, out of a total of 14,641 ($= 121 \times 121$) entries for an element order 10. In contrast, the elemental mass matrices produced by the Basis-LV and the Sherwin–Karniadakis basis are much denser. Basis-LV yields a diagonal submatrix block corresponding to the interior modes, but the 1D linear vertex modes involved in the boundary mode construction result in a considerably larger number of non-zero entries in the rest of the mass matrix (Fig. 5(c)). There exists a total of 2401 non-zero entries in the mass matrix for an element order 10 with Basis-LV. With the Sherwin–Karniadakis basis, no diagonal submatrix block exists in the mass matrix, and there are 1225 non-zero entries with an element order 10.

The current basis also results in a very sparse elemental stiffness matrix for quadrilateral elements (Fig. 5(b)). Note that, with the current basis, a vertex mode is orthogonal to all the interior modes with respect to the 2D stiffness operator, but not to any vertex or edge mode. This results in the full submatrix blocks involving only the vertex modes, and involving the vertex and edge modes, on the top rows and left columns of the stiffness matrix (Fig. 5(b)). With respect to the 2D stiffness operator, an edge mode is non-orthogonal to the four vertex modes, and to the corresponding edge mode on the opposite side of the quadrilateral element. It is also non-orthogonal to the interior modes that share the same 1D interior function in the construction as this edge mode. As a result, in the submatrix block involving only the edge modes, this results in the non-zero main diagonal, and an upper or lower off-diagonal entries. This also results in the non-zero entries that represent the couplings between the edges modes and the interior modes. For an interior mode, with respect to the 2D stiffness operator, it is non-orthogonal only to the four edge modes that share in their construction the same 1D mode in either ξ_1 or ξ_2 direction as this interior mode. In the elemental stiffness matrix the submatrix corresponding to only the interior modes is diagonal. For a quadrilateral element with an element order 10, the current expansion basis results in an elemental stiffness matrix that contains 1105 non-zero entries, out of a total of 14,641 entries. From Fig. 5(d) one can observe that the Basis-LV results in the same sparsity pattern in the elemental stiffness matrix as the current expansion basis. With the Sherwin–Karniadakis basis, the elemental stiffness matrix does not contain a diagonal submatrix block. However, the total number of non-zero entries in its elemental stiffness matrix is less than that of the current basis and Basis-LV. For an element order 10, the elemental stiffness matrix contains 741 non-zero entries with the Sherwin–Karniadakis basis.

We next consider the sparsity of the elemental matrices of a 2D quadrilateral element in a distorted physical domain. Fig. 6(a) shows a sketch of the element. Note that in this case a variable Jacobian (space-dependent) is involved in the computation of the mass and stiffness matrices owing to the mapping between the standard and physical domains, and that in terms of sparsity the stiffness matrix will be more severely affected by the geometric irregularity due to the presence of the gradient operator in its computation. Fig. 6(b), (c) and (d) show sparsity patterns of the mass matrices with the current basis, Basis-LV, and Sherwin–Karniadakis basis, respectively. The stiffness matrices are not shown here because they are full for this element with all three bases. We observe that the interior submatrix with the current basis (and also Basis-LV) is no-longer a diagonal matrix, due to the variable Jacobian, but is still quite sparse. With the current basis and an element order 10 the mass matrix contains 1661 non-zero entries, out of a total of 14,641 entries. On the other hand, with Basis-LV and the Sherwin–Karniadakis basis there are 6321 and 3465 non-zero entries, respectively, in the mass matrices.

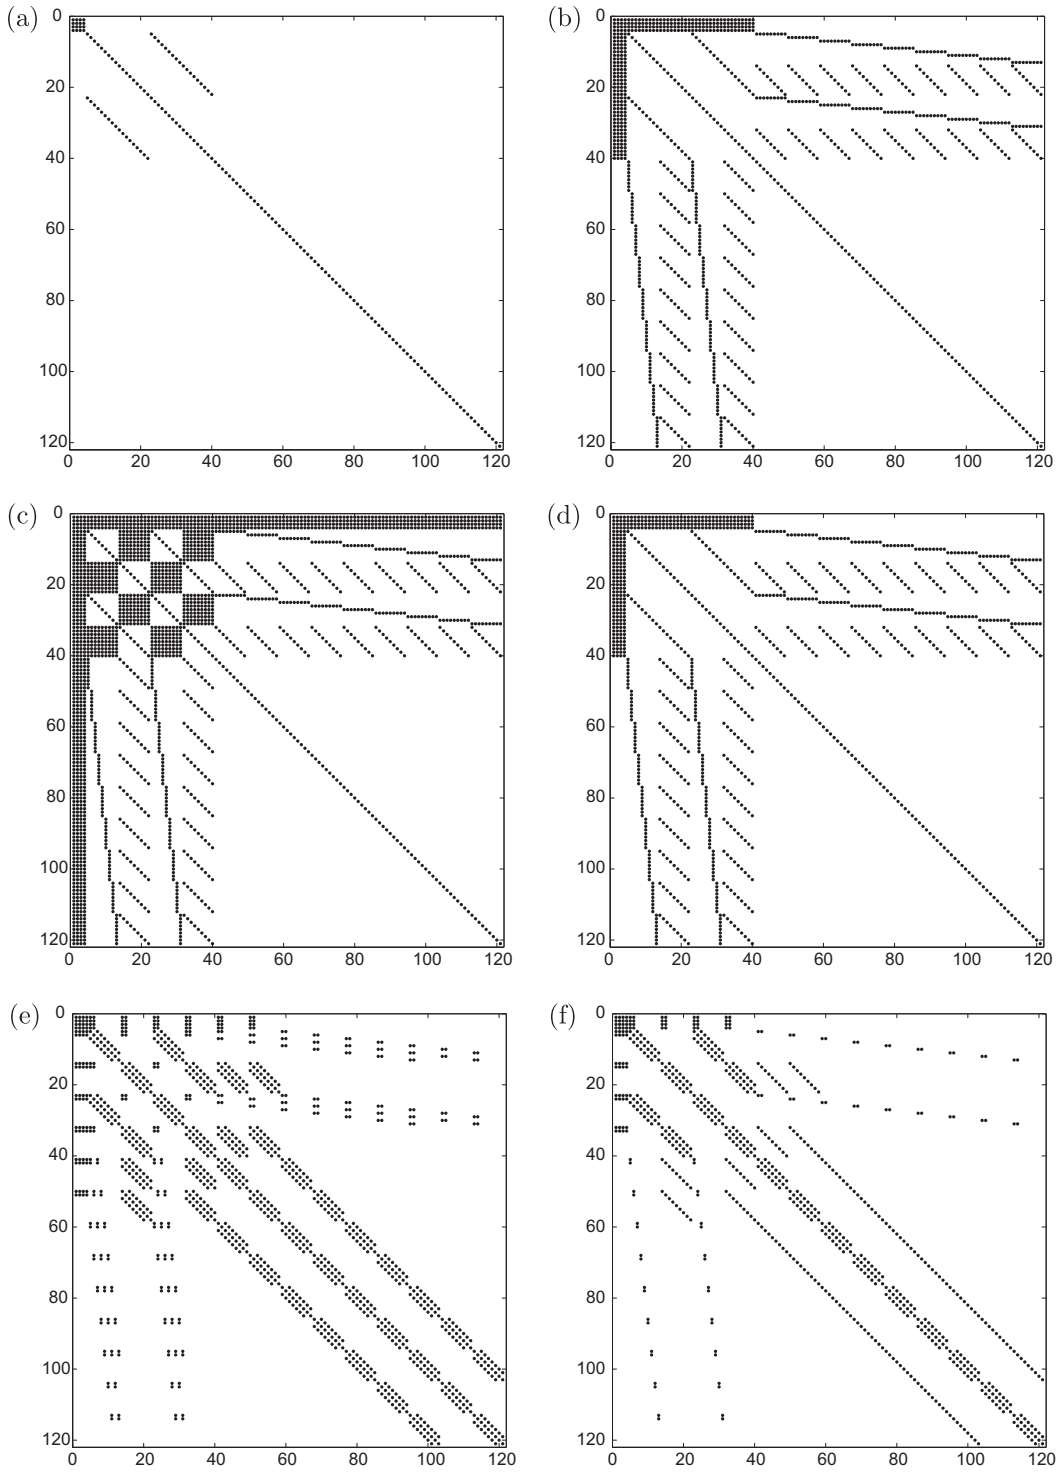

Fig. 5. Sparsity patterns of elemental matrices for a 2D quadrilateral element (order 10, standard domain): left columns, elemental mass matrix; right columns, elemental stiffness matrix; (a) and (b), current expansion basis; (c) and (d), Basis-LV (see text for definition); (e) and (f), Sherwin–Karniadakis basis.

Let us next study the sparsity of the elemental matrices for a 3D hexahedral element. Fig. 7(a) and (b) respectively show the sparsity patterns of the elemental mass and stiffness matrices for a hexahedral element (element order 10) in the standard domain with the current expansion basis. The ordering of the variables is such that the vertex modes are followed by

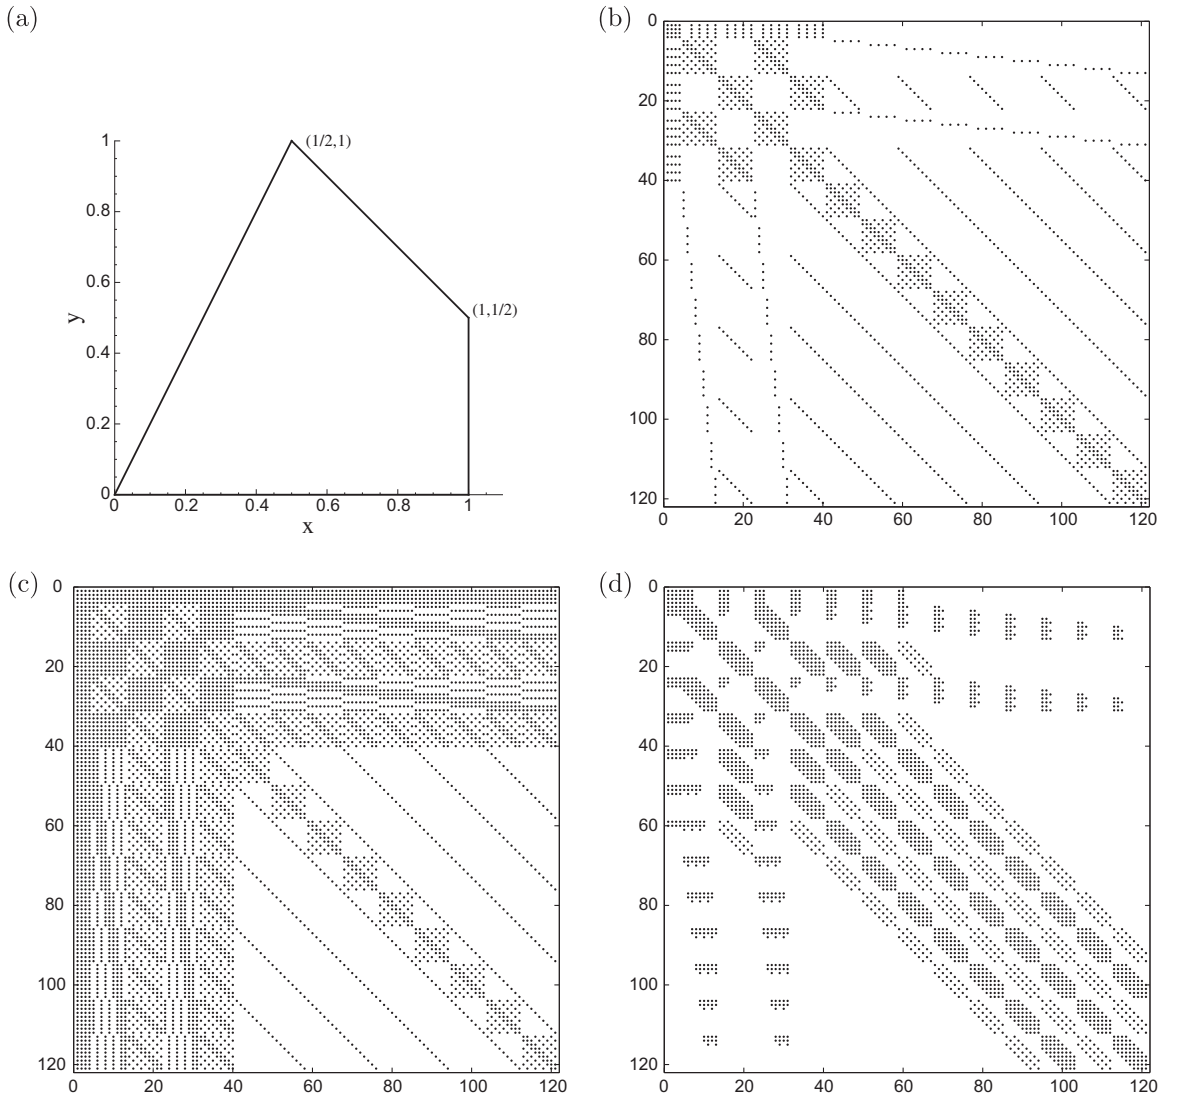

Fig. 6. Sparsity patterns of the elemental mass matrices for one distorted 2D quadrilateral element (element order 10): (a) sketch of 2D distorted quadrilateral element; mass matrices with current basis (b), Basis-LV (c), and Sherwin–Karniadakis basis (d). The stiffness matrices are full for this element with all three bases.

the edge modes, and then by the face modes, and finally by the interior modes. For comparison, in Fig. 7(c) and (d) we show the sparsity patterns of the mass and stiffness matrices obtained with the Basis-LV, and in Fig. 7(e) and (f) we show the sparsity patterns of the elemental matrices with the Sherwin–Karniadakis basis.

It is evident from Fig. 7(a) that the current basis results in an extremely sparse elemental mass matrix for the hexahedral elements. With the current basis a vertex mode is non-orthogonal only to the 8 vertex modes, and is orthogonal to all the other modes of the hexahedral element. For an edge mode associated with a particular edge, it is orthogonal to all the modes except for itself and the three corresponding edge modes associated with the other three edges parallel to the current one. For a face mode, it is only non-orthogonal to itself and the corresponding face mode associated with the face opposite to the current one in the hexahedron. An interior mode is orthogonal to all the modes except itself. For a hexahedral element with an element order 10, the elemental mass matrix from the current basis contains only 2197 non-zero entries, out of a total of 1,771,561 (= 1331×1331) entries. In contrast, the elemental mass matrix associated with the Basis-LV is considerably denser, even though its submatrix block involving only the interior modes is diagonal (Fig. 7(c)). The mass matrix contains 117,649 non-zero entries (for element order 10) with the Basis-LV. With the Sherwin–Karniadakis basis, one can observe that the elemental mass matrix does not contain a diagonal submatrix block (Fig. 7(e)), unlike Basis-LV or the current basis. For an element order 10, the elemental mass matrix with the Sherwin–Karniadakis basis contains 42,875 non-zero entries.

The elemental stiffness matrix for the 3D hexahedral element obtained with the current basis is also very sparse (Fig. 7(b)). In the current basis, with respect to the 3D stiffness operator, a vertex mode is non-orthogonal only to the vertex

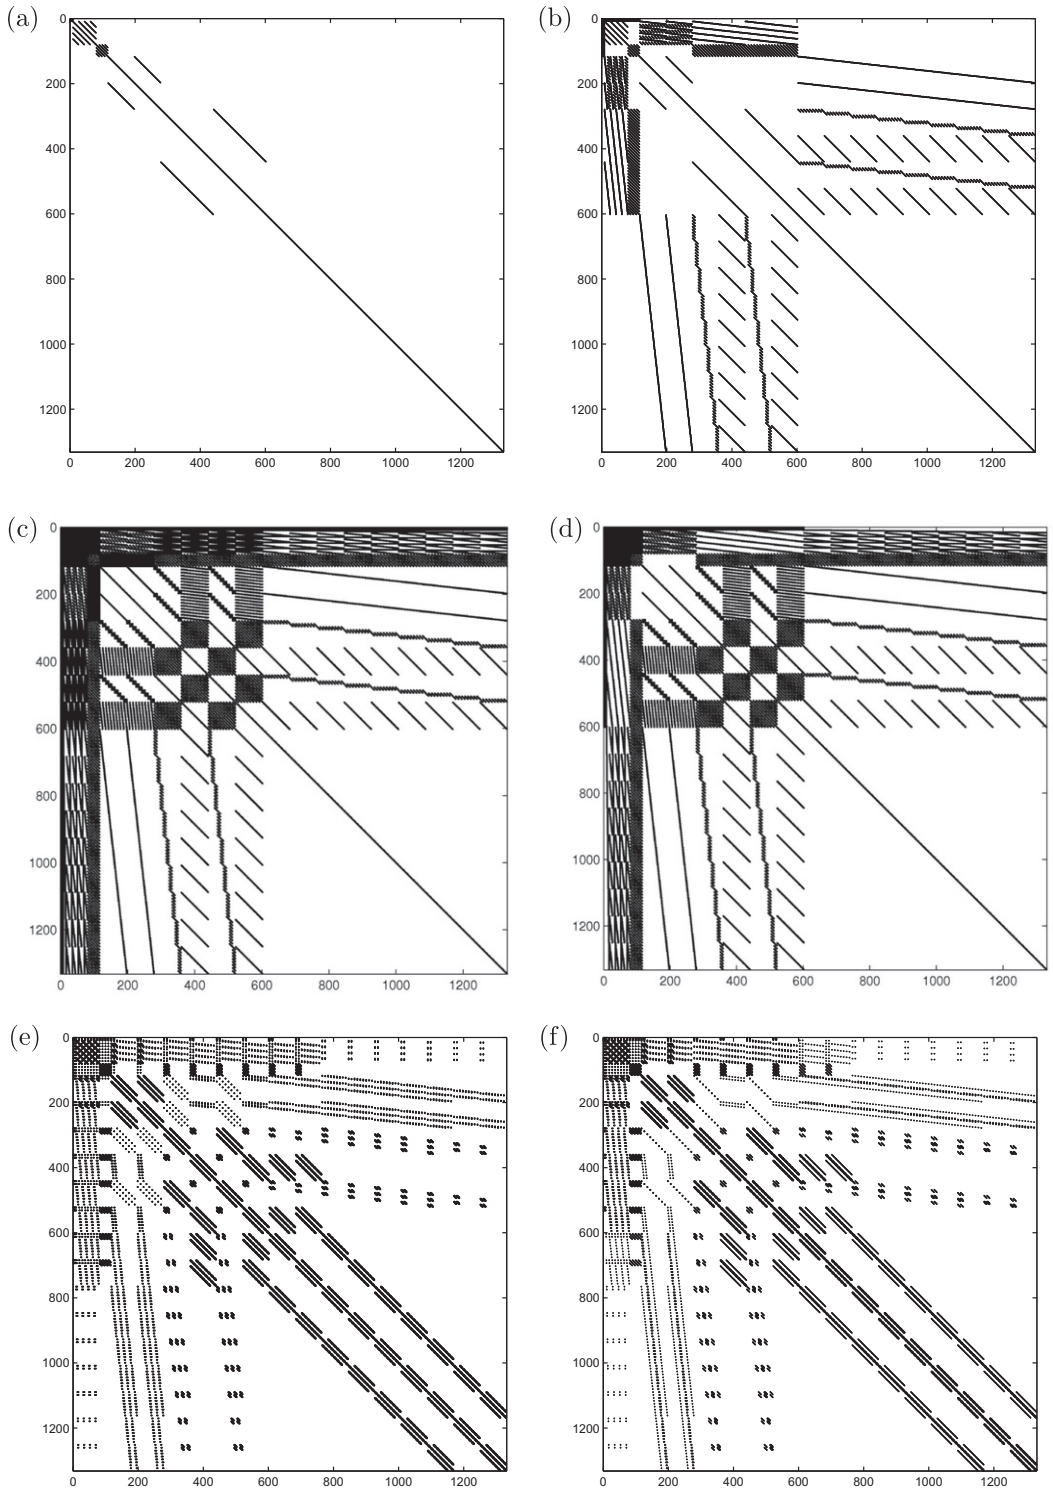

Fig. 7. Sparsity patterns of elemental matrices for a 3D hexahedral element (order 10, standard domain): Left columns, elemental mass matrix; right columns, elemental stiffness matrix; (a) and (b), current expansion basis; (c) and (d), Basis-LV; (e) and (f), Sherwin–Karniadakis basis.

modes and all the edge modes. An edge mode interacts only with the 8 vertex modes, itself and the three other edge modes associated with the edges that are parallel to the current edge, as well as a number of face modes associated with 4 of the 6 faces. With respect to the 3D stiffness operator, a face mode interacts with 8 edge modes, itself and the corresponding face

mode associated with the face opposite to the current one in the hexahedron, and a number of interior modes. On the other hand, an interior mode is non-orthogonal with respect to the 3D stiffness operator only to itself and 6 face modes. The submatrix of the elemental stiffness matrix involving only the interior modes is diagonal. For a hexahedral element with an element order 10, the elemental stiffness matrix resulting from the current basis contains 20,449 non-zero entries, among a total of 1,771,561 entries. In contrast, the elemental stiffness matrices with the Basis-LV and the Sherwin–Karniadakis basis exhibits considerably more complicated patterns and are significantly denser. With an element order 10, there are 70,993 non-zero entries in the elemental stiffness matrix obtained with Basis-LV, and 32,227 non-zero entries with the Sherwin–Karniadakis basis.

We next consider the sparsity of the elemental matrices of a 3D hexahedral element in a distorted physical domain. Fig. 8 shows a sketch of the element, in which the coordinates of all the vertices have been marked. In Fig. 9 we show sparsity patterns of the elemental mass and stiffness matrices (element order 8) with the current basis (plots (a) and (b)), together with those of Basis-LV (plots (c) and (d)) and the Sherwin–Karniadakis basis (plots (e) and (f)). It is evident that all the matrices are significantly denser compared to those of the standard element. For the current basis and Basis-LV, the Jacobian of the mapping between the standard and physical domains has destroyed the diagonality of the interior sub-blocks of the mass/stiffness matrices as observed for the standard element. For the mass matrices, out of a total of 531,441 (= 729 $\times$ 729 for element order 8) entries, there exist 10,659, 132,327 and 64,757 non-zero entries respectively with the current basis, Basis-LV, and Sherwin–Karniadakis basis. For the stiffness matrices, there correspondingly exist 100,187, 255,879, and 185,831 non-zero entries respectively for the three expansion bases.

The sparsity patterns demonstrated above are for the elemental matrices of a single element in 2D or 3D. Let us next investigate the sparsity of the global mass and stiffness matrices with multiple elements. We first consider 2D quadrilateral elements. Fig. 10 shows the sparsity patterns of the global mass and stiffness matrices (with element order 8) obtained from the current expansion basis, together with those of the Basis-LV and the Sherwin–Karniadakis basis for comparison. The problem setting is a 2D square domain, $\{(x,y) : -1 \leqslant x,y \leqslant 1\}$, discretized using 8 equal-sized quadrilateral elements, with 2 elements in one direction and 4 elements in the other. Dirichlet conditions are assumed on the boundaries of the domain. We employ the current expansion basis, as well as the Basis-LV and the Sherwin–Karniadakis basis, and form the global mass matrix and the global stiffness matrix, with the interior modes included (i.e. no Schur complement). The Dirichlet boundary modes are excluded from the global matrices because they are known. The global degrees of freedom are ordered such that the vertex modes are followed by the edge modes, and then successively by the interior modes of each element.

Fig. 10(a) shows the sparsity pattern of the global mass matrix of the current expansion basis. One can observe that the global mass matrix from the current basis is extremely sparse, with 525 non-zero entries out of a total of 216,225 (= 465 $\times$ 465) entries for an element order 8. For comparison, the sparsity patterns of the global mass matrices obtained with the Basis-LV and the Sherwin–Karniadakis basis are shown in Fig. 10(c) and (e), respectively. With Basis-LV, the submatrix involving only the interior modes of the elements is a diagonal matrix. The coupling between the vertex, edge and interior modes is much stronger compared to the current basis. The global mass matrix contains 5117 non-zero entries with Basis-LV. The Sherwin–Karniadakis basis results in multiple non-zero upper/lower diagonals in the submatrix involving only the interior modes. The global mass matrix from the Sherwin–Karniadakis basis contains 4257 non-zero entries.

Fig. 10(b), (d) and (f) compare the sparsity patterns of the global stiffness matrices obtained with the current basis, Basis-LV, and the Sherwin–Karniadakis basis. The sparsity pattern of the current basis is identical to that of Basis-LV, and the submatrix involving only the interior modes is diagonal. Both matrices contain 2765 non-zero entries. The global stiffness

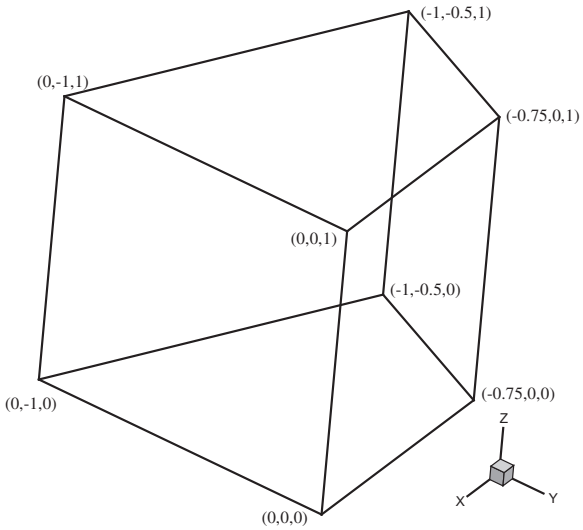

Fig. 8. Geometry of a distorted 3D hexahedral element.

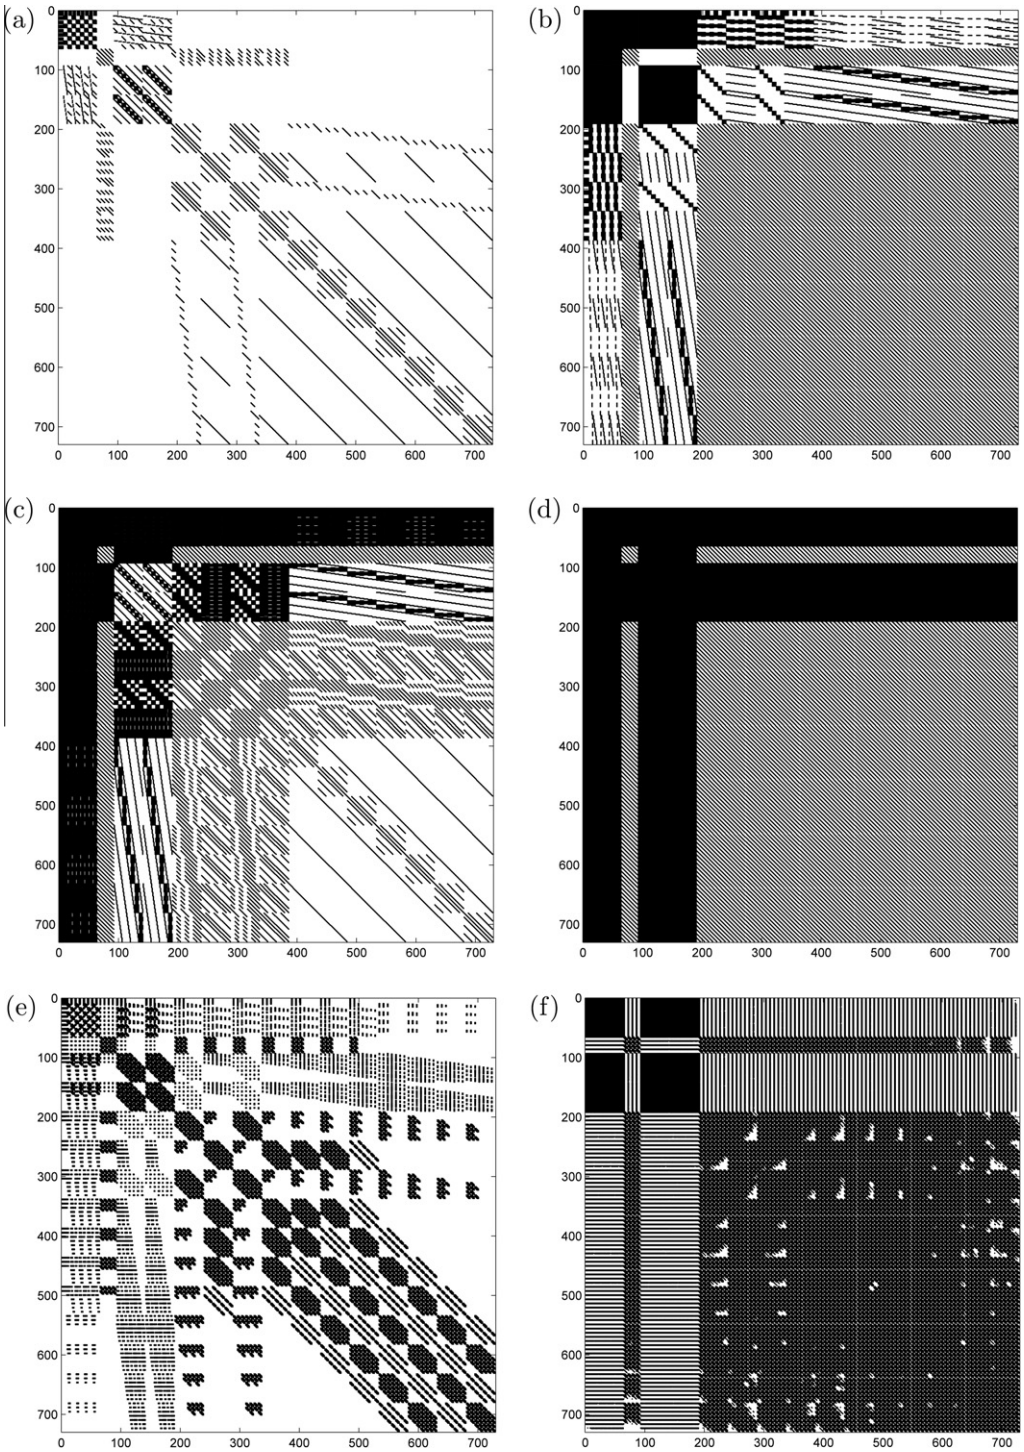

Fig. 9. Sparsity patterns of mass and stiffness matrices for one distorted 3D hexahedral element shown in Fig. 8 (element order 8). Left column, mass matrices; right column, stiffness matrices. (a) and (b), current basis; (c) and (d), Basis-LV; (e) and (f), Sherwin–Karniadakis basis.

matrix from the Sherwin–Karniadakis basis does not contain a diagonal submatrix, unlike the current basis and Basis-LV. It contains 2465 non-zero entries, slightly fewer than that from the current basis and Basis-LV.

We next consider the sparsity of the global mass and stiffness matrices in 3D for an irregular domain. Fig. 11 shows a mesh of four distorted hexahedral elements that is used to discretize this domain. The coordinates of the vertices of these elements have been provided in the figure. We assume Dirichlet conditions on the boundaries of the domain, and form the

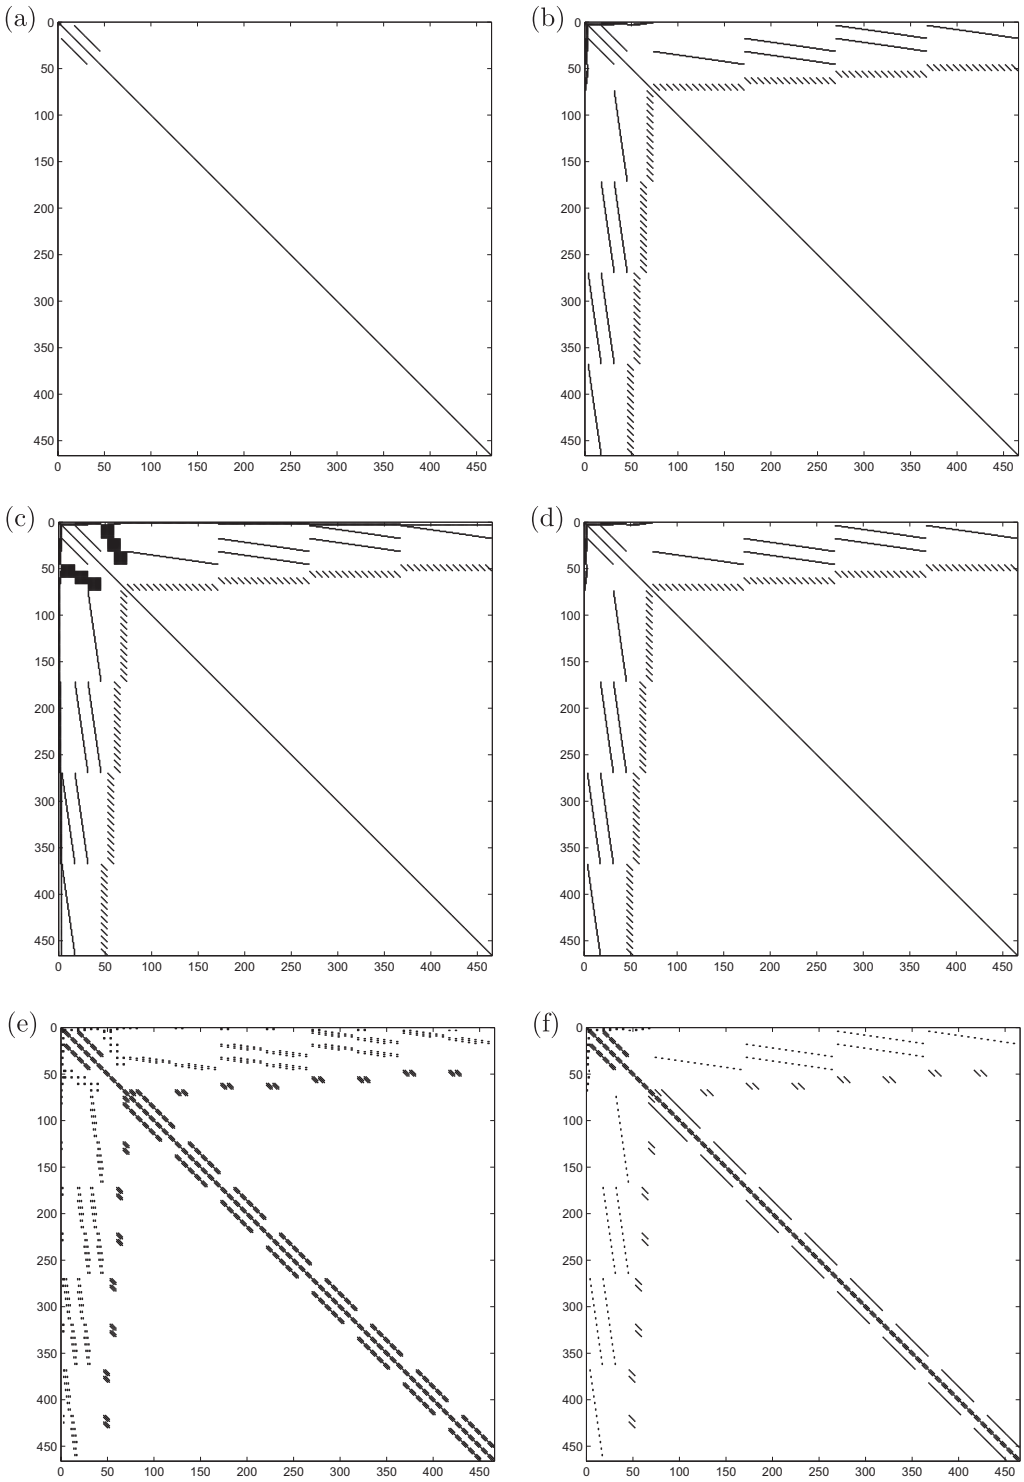

Fig. 10. Sparsity patterns of 2D global matrices (no Schur complement) with 8 quadrilateral elements (element order 8). Left column, global mass matrices, right column, global stiffness matrices. (a) and (b): current expansion basis; (c) and (d): Basis-LV; (e) and (f): Sherwin–Karniadakis basis.

global mass and stiffness matrices, and compute their Schur complements by condensing out the interior modes of all the elements. In Fig. 12 we show the sparsity patterns of the Schur-complemented global mass and stiffness matrices (element order 8) obtained with the current expansion basis (top row), Basis-LV (middle row), and the Sherwin–Karniadakis basis

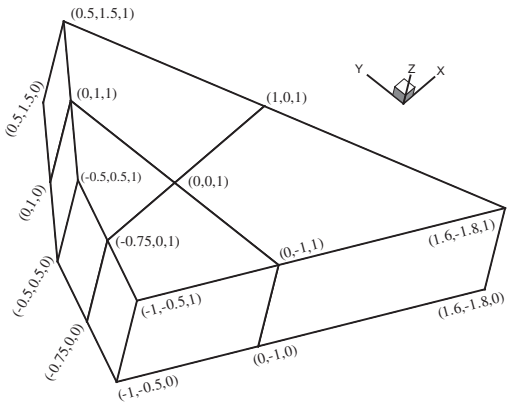

Fig. 11. Mesh of 4 distorted hexahedral elements with irregular 3D domain.

(bottom row). Note that the Dirichlet boundary modes are excluded from the Schur-complemented matrices since they are known. The current basis and Basis-LV exhibit the same sparsity pattern in the Schur-complemented global stiffness matrix, and similar but slightly different patterns in the Schur-complemented global mass matrix. These patterns are quite different than those from the Sherwin–Karniadakis basis. Among a total of 41209 ($= 203 \times 203$) entries, the Schur-complemented global mass matrix contains 3605 non-zero entries with the current basis, 3773 non-zero entries with Basis-LV, and 4051 non-zero entries with Sherwin–Karniadakis basis. The Schur-complemented global stiffness matrix contains 4515 non-zero entries with the current basis and Basis-LV, and 16125 non-zero entries with the Sherwin–Karniadakis basis.

2.2.4. Conditioning of mass and stiffness matrices

We next investigate the conditioning of the global matrices resulting from the current expansion basis in two and three dimensions. Because a Schur complement would always be performed to condense out the interior modes when solving the partial differential equations with the spectral element approach in practice, it is the Schur-complemented matrices that will directly influence the numerical efficiency of the method (see Section 1). We therefore will look into the condition numbers of the Schur-complemented global mass and stiffness matrices.

Fig. 13 demonstrates the conditioning of the current expansion basis for 2D quadrilateral elements. In Fig. 13(a) we show the condition number of the Schur-complemented global mass matrix as a function of the element order for the current basis. This is for a square domain, $\{(x,y) : -1 \leqslant x,y \leqslant 1\}$, discretized with 4 equal-sized quadrilateral elements (2 elements in each direction), assuming Dirichlet conditions on domain boundaries. For comparison, the condition numbers with the Sherwin–Karniadakis basis, and the basis of [27], have also been included in the plot. With the basis of [27], for quadrilateral elements, the interior modes are tensor products of the 1D generalized eigen-functions which result in an identity 1D interior stiffness matrix and a diagonal 1D interior mass matrix. We note that the reference [27] considered only the interior modes, and did not touch on the boundary mode in one or higher dimensions. Here for the boundary modes, we have employed the tensor products of the 1D generalized eigen-functions from [27] and the 1D linear vertex modes (Eqs. (8) and (10)).

One can observe that the Schur-complemented global mass matrix from the current basis has considerably lower condition numbers (by orders of magnitude) than those of the Sherwin–Karniadakis basis and the basis of [27]. The condition numbers with the basis of [27] are comparable to those with the Sherwin–Karniadakis basis. But the condition number with the former basis grows more rapidly with increasing element order. For element orders within 10 the basis of [27] yields smaller condition numbers than the Sherwin–Karniadakis basis, while beyond order 10 its condition numbers are larger.

Fig. 13(b) shows the condition number of the Schur-complemented global stiffness matrix as a function of the element order obtained with the current basis, the Sherwin–Karniadakis basis, and the basis of [27]. The condition numbers with the current basis are markedly smaller than those of the other two bases, except that for very low element orders (orders 2 and 3) they are comparable to those of the basis of [27]. As the element order increases, a more rapid growth in the condition number is observed with the basis of [27]. With an element order below about 12, the basis of [27] exhibits lower condition numbers than the Sherwin–Karniadakis basis, while for element orders beyond 12 it yields notably higher condition numbers.

We next demonstrate the conditioning of the current basis for 3D hexahedral elements with Fig. 14. We consider a cubic domain, $\{(x,y,z):-1 \leqslant x,y,z \leqslant 1\}$, and discretize this domain using 4 equal-sized hexahedral elements, with 2 elements in two of the three directions and one element in the other direction. The same element order has been used for all elements, and the order has been varied systematically in the tests to study its effects on the condition numbers of the global matrices. Dirichlet conditions are assumed on the boundaries of the domain.

In Fig. 14(a) we plot the condition number of the Schur-complemented global mass matrix as a function of the element order obtained with the current expansion basis. We have also included the results obtained with the Sherwin–Karniadakis

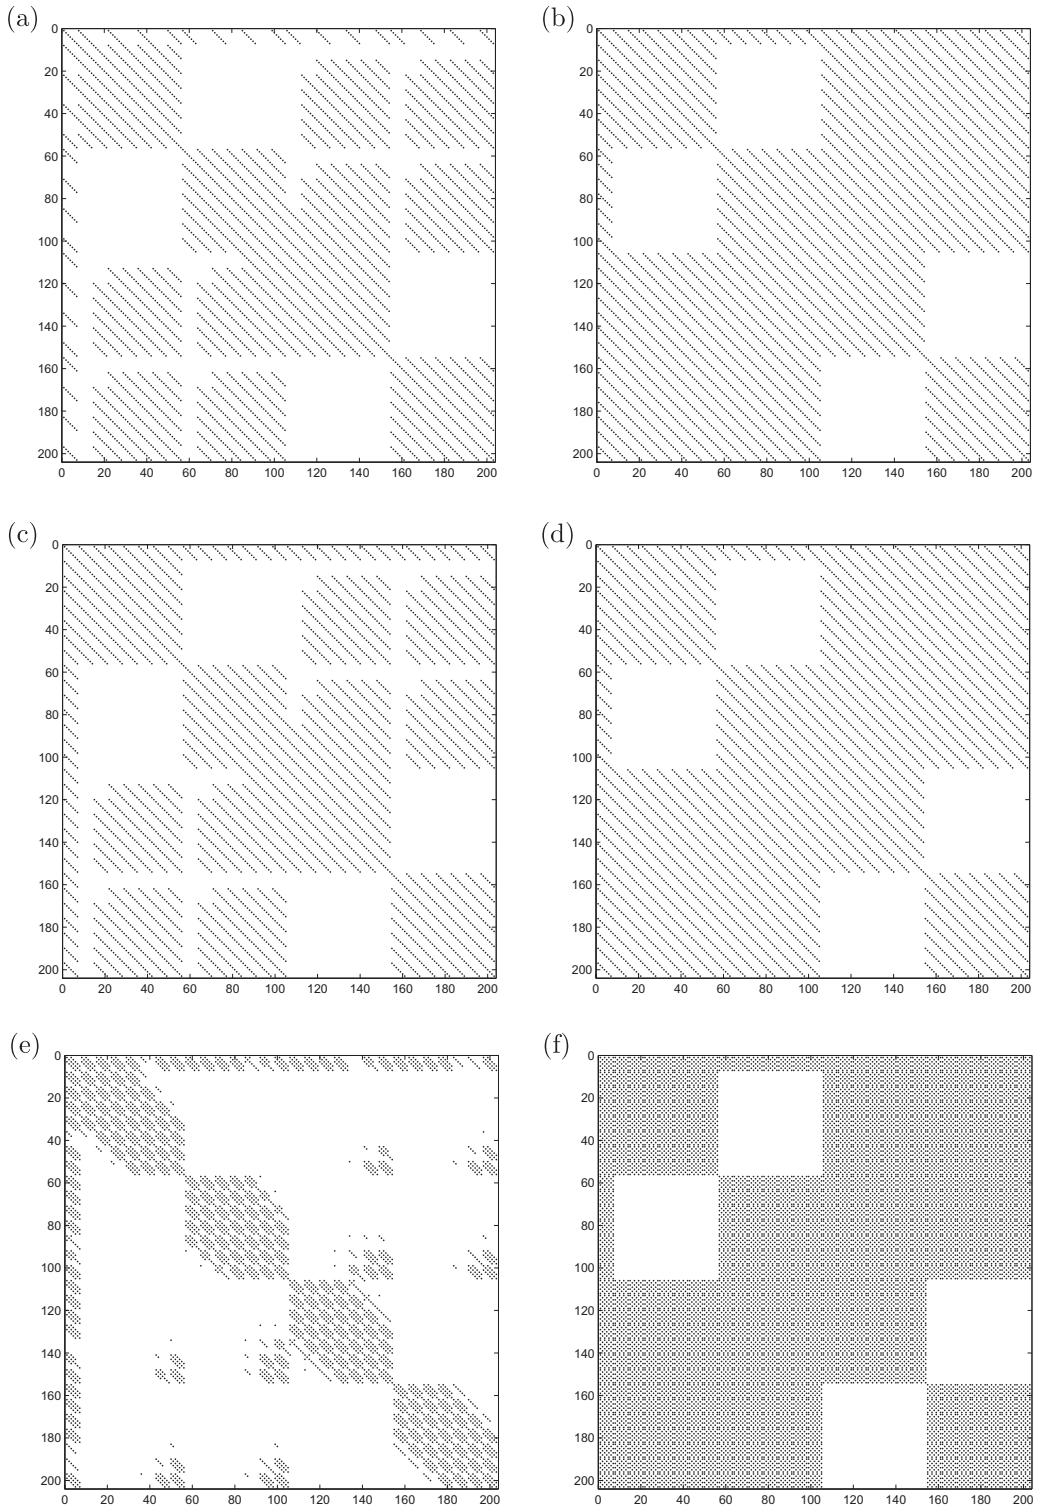

Fig. 12. Sparsity patterns of Schur-complemented global mass matrices (left column) and stiffness matrices (right column) for the 3D mesh in Fig. 11 of four distorted hexahedral elements (element order 8): (a) and (b), current basis; (c) and (d), Basis-LV; (e) and (f), Sherwin–Karniadakis basis.

basis and the Basis-LV for comparison. The mass matrix resulting from the current basis is considerably better-conditioned, with its condition numbers two to three orders of magnitude lower, than those from the other two bases. One can also observe that the condition numbers with Basis-LV are notably lower than those with the Sherwin–Karniadakis basis.

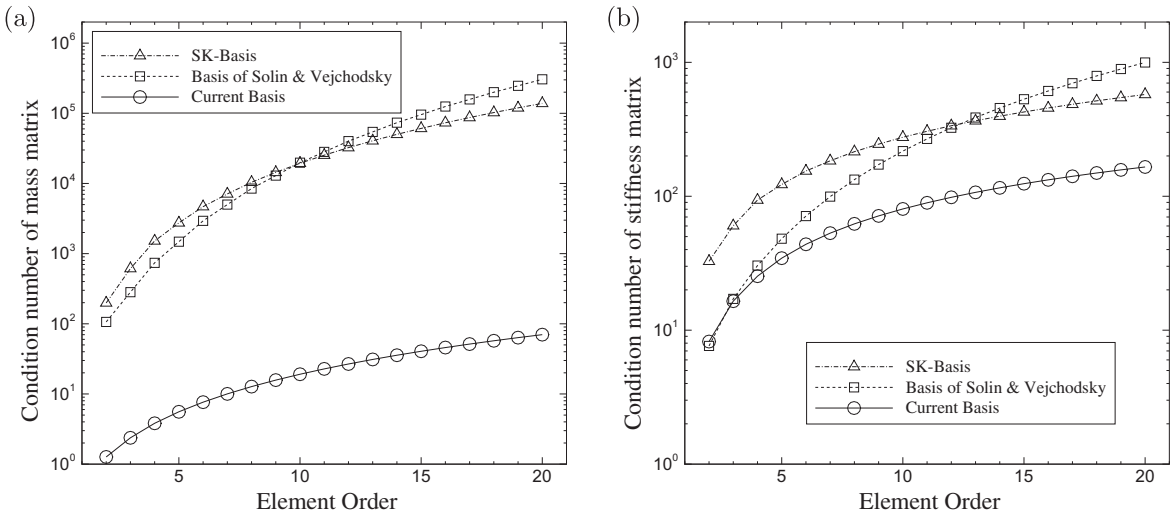

Fig. 13. Condition numbers of the Schur-complemented global mass matrix (a) and stiffness matrix (b) as a function of the element order computed with the current basis, the Sherwin–Karniadakis basis, and the Basis of [27], with four 2D quadrilateral elements.

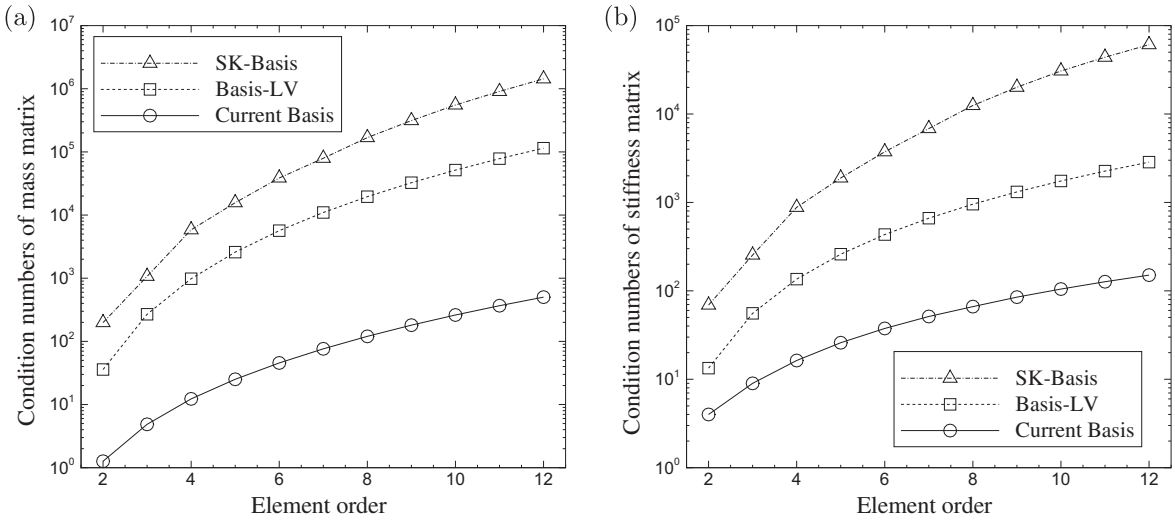

Fig. 14. Condition numbers of the Schur-complemented global mass matrix (a) and stiffness matrix (b) as a function of the element order computed with the current basis, the Sherwin–Karniadakis basis, and Basis-LV, with four 3D hexahedral elements.

In Fig. 14(b) we compare the condition numbers of the Schur-complemented global stiffness matrix as a function of the element order resulting from the current basis, Basis-LV, and the Sherwin–Karniadakis basis. One can again observe the significantly superior conditioning of the current basis. The condition numbers of the Schur-complemented stiffness matrix with the current basis are one to two orders of magnitude smaller than those of the Sherwin–Karniadakis basis, and are also significantly smaller than those with the Basis-LV. The conditioning of the stiffness matrix with Basis-LV is also markedly better compared with that of the Sherwin–Karniadakis basis.

Let us next demonstrate the conditioning of the current expansion basis with an irregular geometric domain. We consider the domain and the mesh of four distorted hexahedral elements depicted in Fig. 11. Note that the sparsity patterns of the Schur-complemented global mass and stiffness matrices for this mesh have been shown in Fig. 12 for an element order 8. Let us now consider the conditioning of these matrices, assuming Dirichlet conditions on the boundaries of the domain. In Fig. 15 we show the condition numbers of the Schur-complemented global mass matrix (Fig. 15(a)) and global stiffness matrix (Fig. 15(b)) as a function of the element order for the current expansion basis. The condition numbers with Basis-LV and the Sherwin–Karniadakis basis have also been included in the plots for comparison. One can again observe that the condition numbers from the current expansion basis are significantly smaller than those from Basis-LV and the Sherwin–Karniadakis basis.

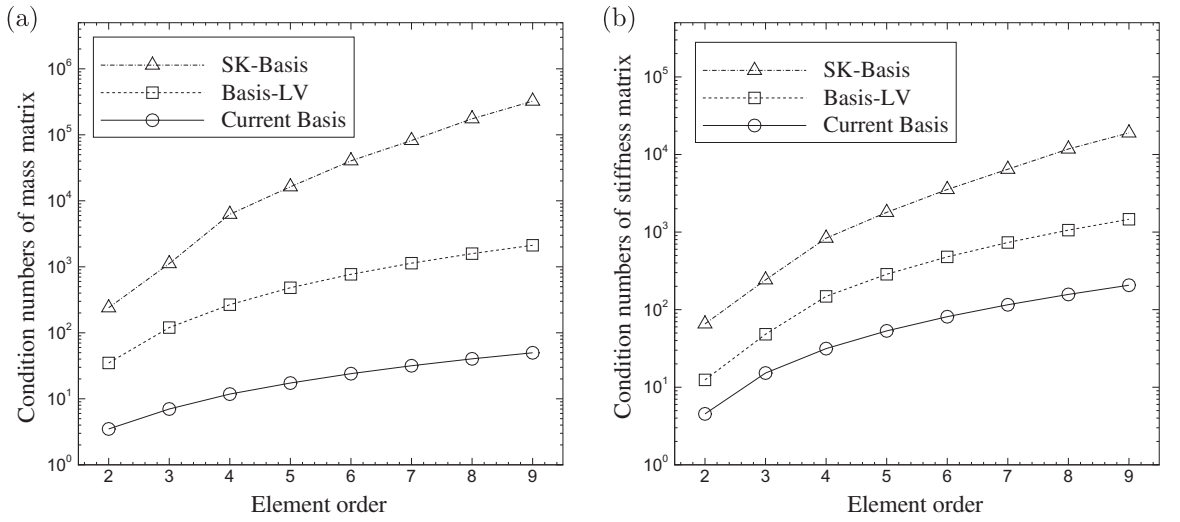

Fig. 15. Condition numbers of the Schur-complemented global mass matrix (a) and stiffness matrix (b) as a function of element order computed with the current basis, Sherwin–Karniadakis basis and Basis-LV, with the four distorted 3D hexahedral elements in Fig. 11.

3. Numerical examples

In this section we demonstrate the efficiency of the proposed expansion basis using several numerical examples. We will concentrate on the Helmholtz equation in two and three dimensions, which is the prototype equation resulting from appropriate discretizations (e.g. with explicit treatment of the convection term) of the incompressible Navier–Stokes equations, heat equations, and wave equations. We look into the conditioning of the resultant coefficient matrices and the number of iterations it takes the conjugate gradient solver to converge to a specified tolerance.

3.1. Helmholtz equation in two and three dimensions

3.1.1. Convergence rate

We first demonstrate the exponential convergence rate of the spectral element discretization employing the current expansion basis for smooth solutions. Consider a domain $\Omega \subset \mathbb{R}^d$ ($d = 2$ or 3) and the Helmholtz equation

$$\nabla^2 u - \lambda u = f, \tag{16}$$

on this domain, where λ is a prescribed constant, f is a known function in 2D or 3D, and u is the scalar field function to be solved. We assume Dirichlet conditions on the boundary of the domain, $u|_{\partial\Omega} = u_0(\mathbf{x})$. By taking the L^2 inner product of Eq. (16) with a test function and integration by part, one obtains the weak form of the equation,

$$\int_\Omega \nabla u \cdot \nabla \phi d\mathbf{x} + \lambda \int_\Omega u\phi d\mathbf{x} = -\int_\Omega f\phi d\mathbf{x}, \tag{17}$$

where the test function ϕ belongs to $H^1(\Omega)$ and vanishes on the Dirichlet boundaries. We partition the domain Ω using a set of structured spectral elements. Let Ω^h denote the discretized domain, and $X^h \subset H^1(\Omega^h)$ denote the approximation space of u. Then the discretized problem is: find $u^h \in X^h$ such that

$$\int_{\Omega^h} \nabla u^h \cdot \nabla \phi^h d\mathbf{x} + \lambda \int_{\Omega^h} u^h \phi^h d\mathbf{x} = -\int_{\Omega^h} f^h \phi^h d\mathbf{x}, \tag{18}$$

where ϕ^h is the discrete version of the test function, $\phi^h \in X^h$, and it vanishes on the Dirichlet boundary $\partial\Omega^h$.

We employ the high-order expansion basis presented in Section 2 in the discretized Eq. (18), and investigate the convergence rate by systematically looking into the errors of the numerical solution against an analytic solution. Fig. 16(a) demonstrates the convergence rate for the 2D Helmholtz equation computed using the current expansion basis. Here we have considered a square domain, $\Omega = \{(x,y) : -1 \leqslant x,y \leqslant 1\}$, discretized it using four equal-sized quadrilateral elements, with 2 elements in each direction. In this case, $\lambda = 100$, and $f = -(\lambda + 2\pi^2)\cos\pi x \cos\pi y$,and the analytic solution to the Helmholtz equation is given by

$$u = \cos\pi x \cos\pi y. \tag{19}$$

Dirichlet conditions are imposed on the domain boundaries according to the form of the analytic solution. The same element order has been used for all elements. We vary the element order systematically, and for each element order compute the L^∞

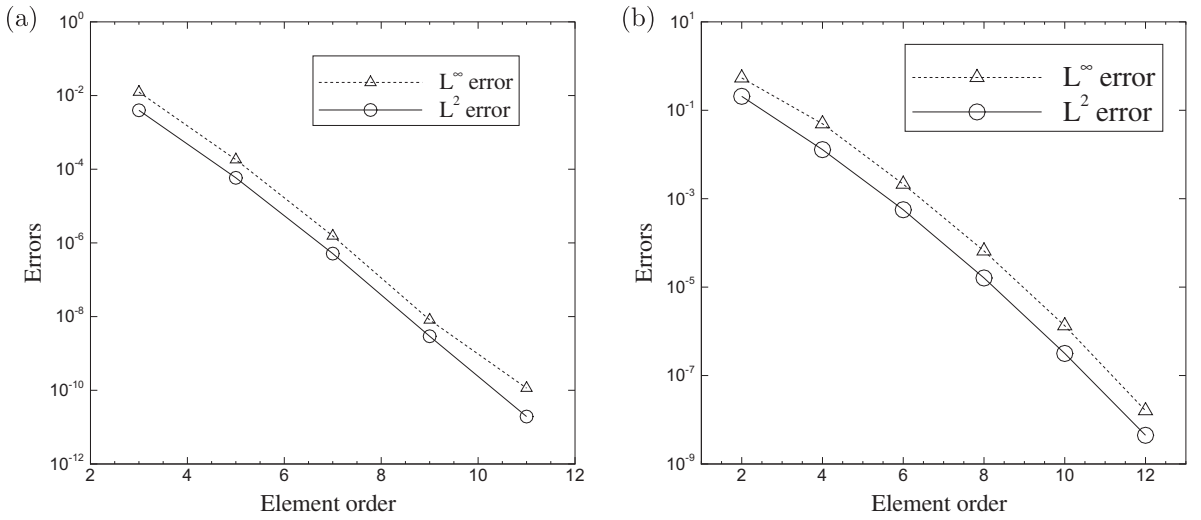

Fig. 16. Exponential convergence rate with the current expansion basis: L^∞ and L^2 errors as a function of the element order for the Helmholtz equation ($\lambda = 100$) in 2D (a) and 3D (b). Four quadrilateral elements are used in (a) and four hexahedral elements are used in (b).

and L^2 errors of the numerical solution against the analytic solution. Fig. 16(a) shows these errors as a function of the element order obtained with the current basis. The errors clearly decrease exponentially with respect to the element order, showing an the exponential convergence rate as the element order increases.

Fig. 16(b) demonstrates the convergence rate for the 3D Helmholtz equation. This is for a cubic domain, $\Omega = \{(x,y,z) : -1 \leqslant x,y,z \leqslant 1\}$, discretized using 4 equal-sized hexahedral elements, with 2 elements in two of the three directions and one element in the other direction. The parameter is $\lambda = 100$, and the function f is given by $f = -(\lambda + 3\pi^2)\cos\pi x\cos\pi y\cos\pi z$. In this case, the Helmholtz equation has an analytic solution given by

$$u = \cos\pi x \cos\pi y \cos\pi z. \tag{20}$$

The figure shows the L^∞ and L^2 errors of the numerical solution against the analytic solution as a function of the element order obtained with the current basis. One can observe the exponential decrease of the numerical errors with respect to the element order for 3D elements.

3.1.2. Numerical efficiency with 2D elements

We next concentrate on the computational efficiency, particularly the conditioning of the global coefficient matrix and the number of iterations to convergence for the conjugate gradient solver, for solving the Helmholtz equation employing the current basis. Let us first consider the 2D Helmholtz equation. We will employ the analytic solution given by Eq. (19). We first consider numerical tests on a square domain, $-1 \leqslant x, y \leqslant 1$, and assume Dirichlet conditions on the domain boundaries. We conduct tests on two spectral element meshes: (1) a smaller mesh of 4 equal-sized quadrilateral elements with 2 elements in each direction, and (2) a larger mesh of 16 equal-sized elements with 4 elements in each direction.

Let us first briefly comment on the solution process. When solving the linear algebraic system resulting from the discretized Eq. (18), we first perform a Schur complement (static condensation) to condense out the interior degrees of freedom from all the elements. This can be accomplished by arranging the unknown global degrees of freedom in a fashion such that the vertex modes are followed by the edge modes, and then by the interior modes of each element. Note that the Shur complement only needs to be performed at the local elemental level; see e.g. [17] for details of this procedure. As a result, solving the linear algebraic system involves two stages: first solve a reduced linear system for the boundary modes with the conjugate gradient solver, involving the Schur-complemented coefficient matrix, followed by a direct solve of a small linear system for the interior modes successively within each element. The second stage involves a simple matrix vector multiplication because the interior coefficient matrix can be pre-computed and pre-factorized. The efficiency of the conjugate gradient solve of the first stage is determined by the properties of the Schur-complemented global coefficient matrix, such as the condition number and the eigenvalue distribution. We will next look into the conditioning, and eigenvalue distribution of the Schur-complemented global coefficient matrix, as well as the number of iterations to convergence in the conjugate gradient solver.

We observe that the current expansion basis exhibits a superior numerical efficiency, in terms of the condition numbers and the number of iterations to convergence for the conjugate gradient solver. We demonstrate this point using Fig. 17, in which we look into the condition numbers of the Schur-complemented global coefficient matrix and the number of conjugate gradient iterations on the two spectral element meshes. Plots (a) and (b) are for the smaller mesh of 4 elements, and plots (c) and (d) are for the larger one with 16 elements.

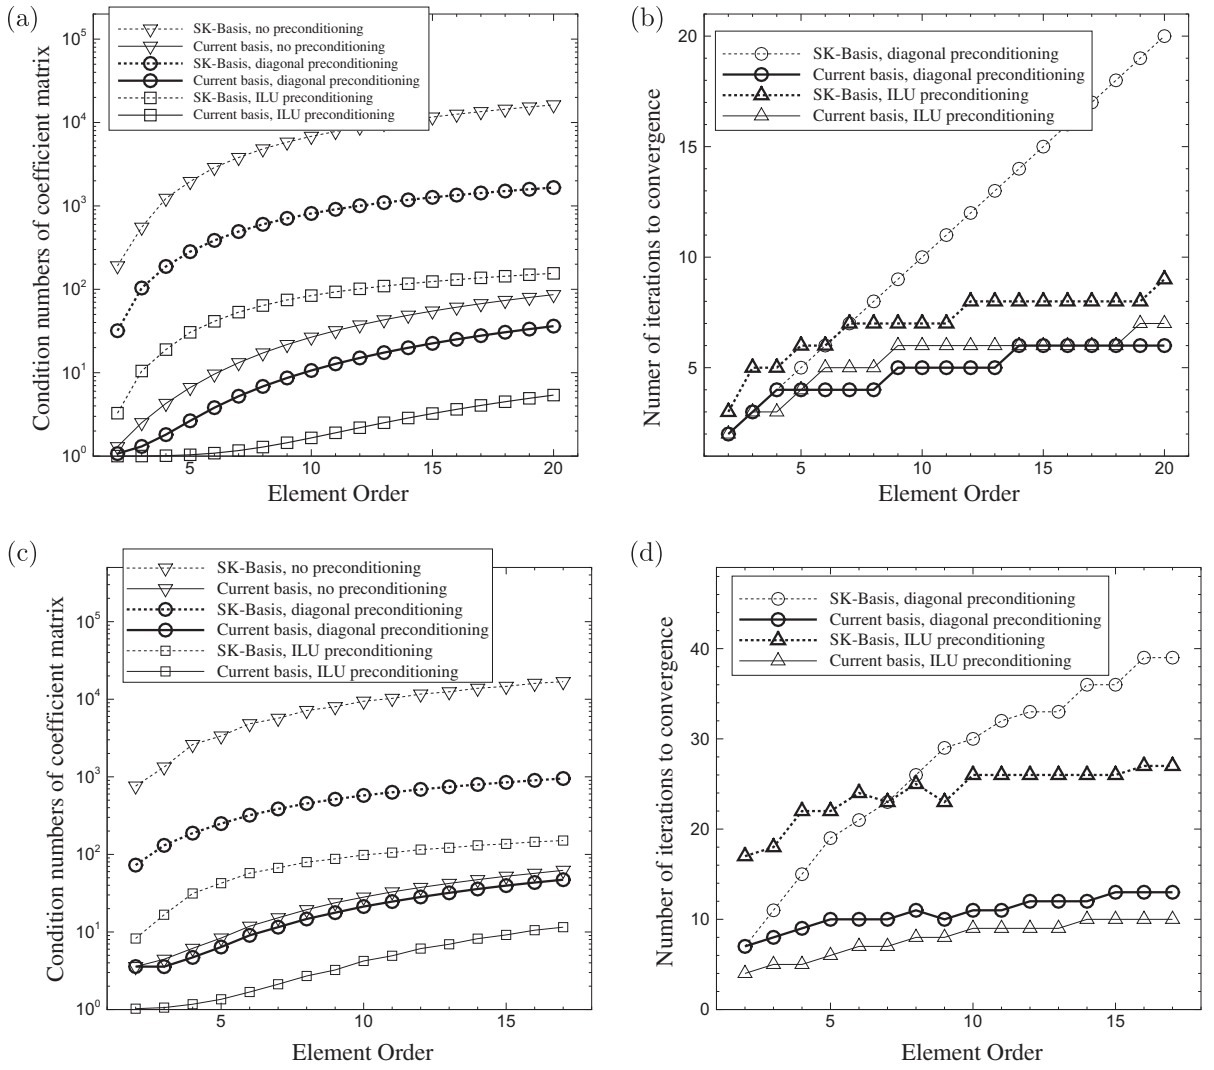

Fig. 17. 2D Helmholtz equation ($\lambda = 1000$): Condition numbers of the Schur-complemented coefficient matrix (a,c) and the number of conjugate gradient iterations to convergence (tolerance 10^{-8}) (b,d) as a function of the element order. Plots (a) and (b) are for a mesh with 4 quadrilateral elements. Plots (c) and (d) are for a mesh with 16 quadrilateral elements.

Fig. 17(a) shows the condition numbers of the Schur-complemented global coefficient matrix of the Helmholtz equation ($\lambda = 1000$) as a function of the element order with the current expansion basis. We compare the cases without preconditioning, with diagonal preconditioning (i.e. Jacobi preconditioner), and with the incomplete LU (ILU) preconditioning [22]. The corresponding condition numbers from the Sherwin–Karniadakis basis for this problem have also been included in the plot for comparison. One can evidently observe that in all cases (with or without preconditioning) the condition numbers resulting from the current basis are considerably smaller, by over an order of magnitude, than those from the Sherwin–Karniadakis basis. We also note that the ILU preconditioner result in notably lower condition numbers than the diagonal preconditioning for both the current and the Sherwin–Karniadakis bases.

In Fig. 17(b) we plot the number of iterations it takes the conjugate gradient solver to converge to a specified tolerance (10^{-8}) as a function of the element order on the smaller mesh of 4 elements. The conjugate gradient solver for computing the boundary modes has been preconditioned with the diagonal preconditioner and the ILU preconditioner. For comparison we have included in Fig. 17(b) the results with both the current basis and the Sherwin–Karniadakis basis. Consider first the results with diagonal preconditioning. One can observe that, with the current basis, the number of iterations remains nearly a constant (with slight increase) with increasing element order for element orders above 3. For example, as the element order increases from 4 to 20 the number of iterations varies from 4 to 6. In contrast, with a diagonal preconditioning the number of iterations with the Sherwin–Karniadakis basis exhibits a linear increase with respect to the element order for the range of values studied here (Fig. 17(b)). With the ILU preconditioner, for the current basis, we do not observe a significant difference in the number of iterations when compared to the diagonal preconditioner. The numbers of iterations with both

preconditioners are essentially the same. Note, however, that the ILU preconditioner has reduced the condition number of the Schur-complemented coefficient matrix significantly more than the diagonal preconditioner (Fig. 17(a)). For the Sherwin–Karniadakis basis, on the other hand, the ILU preconditioner has dramatically improved upon the number of the iterations when compared to the diagonal preconditioning. With the ILU preconditioner, the number of iterations for the Sherwin–Karniadakis basis remains nearly a constant with increasing element order for element orders beyond 7. The number of iterations with the Sherwin–Karniadakis basis is at a higher level than that with the current basis.

Let us next consider the results on the larger mesh of 16 elements. Fig. 17(c) and (d) show the condition numbers of the Schur-complemented global coefficient matrix and the number of conjugate gradient iterations to convergence, respectively, as a function of the element order with the diagonal and ILU preconditioners on the larger mesh. The trend exhibited by the data is similar to that from the smaller mesh. In terms of the condition number, the current basis exhibits a considerably lower value, without preconditioning or with the same type of preconditioner, than the Sherwin–Karniadakis basis. In terms of the number of iterations to convergence, for the current basis, the increase in the number of iterations as the element order increases is very slight with both the diagonal and the ILU preconditioners. For example, as the element order increases from 6 to 17, the number of iterations with the current basis increases from 10 to 13 with the diagonal preconditioning, and increases from 6 to 10 with the ILU preconditioning. With the current basis, the two curves for the number of iterations with the diagonal preconditioning and the ILU preconditioning are essentially parallel to each other (Fig. 17(d)), with slightly fewer iterations (by 2 or 3 iterations) for the ILU preconditioner. The Sherwin–Karniadakis basis exhibits a near-linear increase in the number of iterations versus the element order with the diagonal preconditioner. With the ILU preconditioner, the number of iterations for this basis has been considerably improved, and it exhibits only a slight increase with increasing element order. However, the number of iterations with the Sherwin–Karniadakis basis has a much larger value, by more than two times, than that with the current basis (Fig. 17(d)). For example, with the ILU preconditioner, for an element order 10 the Sherwin–Karniadakis basis requires 26 conjugate gradient iterations to converge, while the current basis requires only 9 iterations.

The number of conjugate gradient iterations is not only influenced by the condition number, but also by the distribution of the eigenvalues of the coefficient matrix. Well-clustered eigenvalues will lead to a faster convergence. The conjugate gradient iterations will converge in at most n steps with n distinct eigenvalues [30]. In Fig. 18 we plot the eigenvalue distributions of the Schur-complemented global coefficient matrix of the 2D Helmholtz equation for the current basis and the Sherwin–Karniadakis basis. The distributions are computed for cases without preconditioning, and with the diagonal and ILU preconditioning. Fig. 18(a) is for the smaller mesh of 4 quadrilateral elements with an element order 14, and Fig. 18(b) is for the larger mesh of 16 elements with an element order 10. A general observation is that preconditioning has caused the eigenvalues to be more clustered compared to cases without preconditioning. With the current expansion basis, the diagonal preconditioner has effectively clustered the eigenvalues around several distinct values or distinct groups of values. For example, on the smaller mesh with an element order 14 and the diagonal preconditioning, there exist only 6 distinct eigenvalues, with 48 out of a total of 53 eigenvalues having the unit value. On the larger mesh with an element order 10, the diagonal preconditioning has clustered the eigenvalues into three distinct groups (around 0.32, 1.0, and 1.39), with 180 eigenvalues out of a total of 225 having the unit value. The favorable eigenvalue distribution of the current basis resulting from the diagonal preconditioning accounts for the fact that, even though the diagonal preconditioner is not as effective in reducing the condition number as the ILU preconditioner, the actual number of conjugate gradient iterations it requires is comparable to that of the ILU preconditioner; see Fig. 17(b) and (d). This indicates that the simple diagonal preconditioner is very effective for the current expansion basis. In contrast, the eigen spectrum of the Sherwin–Karniadakis basis with

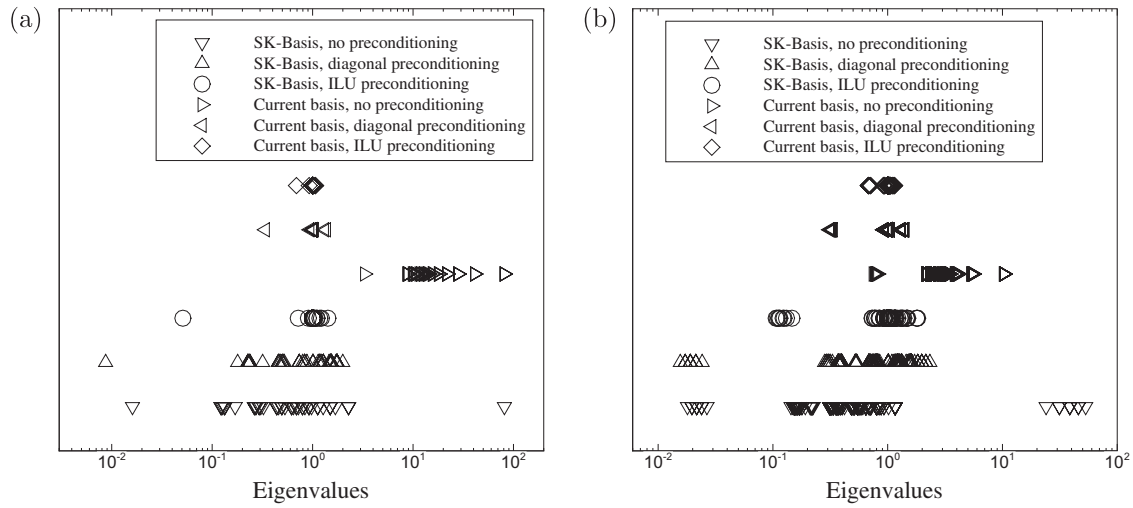

Fig. 18. Eigenvalue distributions of the Schur-complemented Helmholtz coefficient matrix in 2D, (a) with 4 quadrilateral elements and an element order 14, (b) with 16 quadrilateral elements and an element order 10.

diagonal preconditioning is quite scattered on the real axis, thus leading to a large number of iterations to convergence. So the Sherwin–Karniadakis basis does not work effectively with the diagonal preconditioner. The eigenvalue distribution with the ILU preconditioning for the Sherwin–Karniadakis basis is also not as favorable as that for the current basis. The current basis with the ILU preconditioning exhibits the best eigenvalue distribution compared to the other cases considered here.

We next consider the 2D Helmholtz equation on a domain of irregular geometry and with small λ values. Fig. 19(a) shows a sketch of the domain and a mesh of 8 distorted quadrilateral elements used to discretize this domain. The coordinate values for all vertices of the elements have been provided in the figure. We solve the Helmholtz equation with two small λ values ($\lambda = 0$ and 1.0) with this mesh and various element orders, again using the analytic solution in Eq. (19). In Fig. 19(b) we show the number of iterations to convergence (tolerance 10^{-8}) of the conjugate gradient solver as a function of the element order, obtained with the current expansion basis and the Sherwin–Karniadakis basis. A simple diagonal preconditioner has been employed with the conjugate gradient solver in the computations. There exists very little difference in terms of the number of iterations between $\lambda = 0$ and $\lambda = 1$. When the element order is small (below about order 5), the numbers of iterations to convergence are about the same between the current basis and Sherwin–Karniadakis basis. When the element order increases further, the difference in the number of iterations between the two bases becomes more notable. The increase in the number of iterations with the current basis is quite slight with increasing element order, while with the Sherwin–Karniadakis basis the increase in the number of iterations is comparatively more significant.

The above results on the condition numbers, the eigenvalue distribution, and the number of iterations to convergence, suggest that the current expansion basis is very amenable to the diagonal preconditioning. Because of the favorable conditioning and sparsity of the mass and stiffness matrices resulting from the current expansion basis, as demonstrated in Section 2, a simple diagonal preconditioner is very effective with the current basis. A more sophisticated preconditioner such as ILU can further improve the efficiency of the current basis compared to the simple diagonal preconditioning. However, we have observed that in cases the efficiency of the current basis with the diagonal preconditioning can be comparable to that with the more sophisticated preconditioner. On the other hand, the Sherwin–Karniadakis basis does not work effectively with the diagonal preconditioning, and it requires more sophisticated preconditioning techniques for efficient computations. If one takes into account the fact that the simple diagonal preconditioning is computationally extremely cheap and can be trivially parallelized, while more sophisticated preconditioners tend to be considerably more costly and very challenging to parallelize, it can be appreciated that the properties of our expansion basis demonstrated by these results are very advantageous.

3.1.3. Numerical efficiency with 3D elements

We will next investigate the efficiency of the current expansion basis for the Helmholtz equation in 3D. We first consider the analytic solution given by Eq. (20), with $\lambda = 1000$, on a cubic domain $\Omega = \{(x,y,z) : -1 \leqslant x,y,z \leqslant 1\}$. We conduct numerical tests on two spectral element meshes: (1) a smaller mesh of 4 equal-sized hexahedral elements, with 2 elements in two of the three directions and 1 element in the other direction, and (2) a larger mesh of 16 equal-sized hexahedral elements, with 4 elements in one direction and 2 elements in the other two directions. Dirichlet conditions are imposed on the domain boundaries.

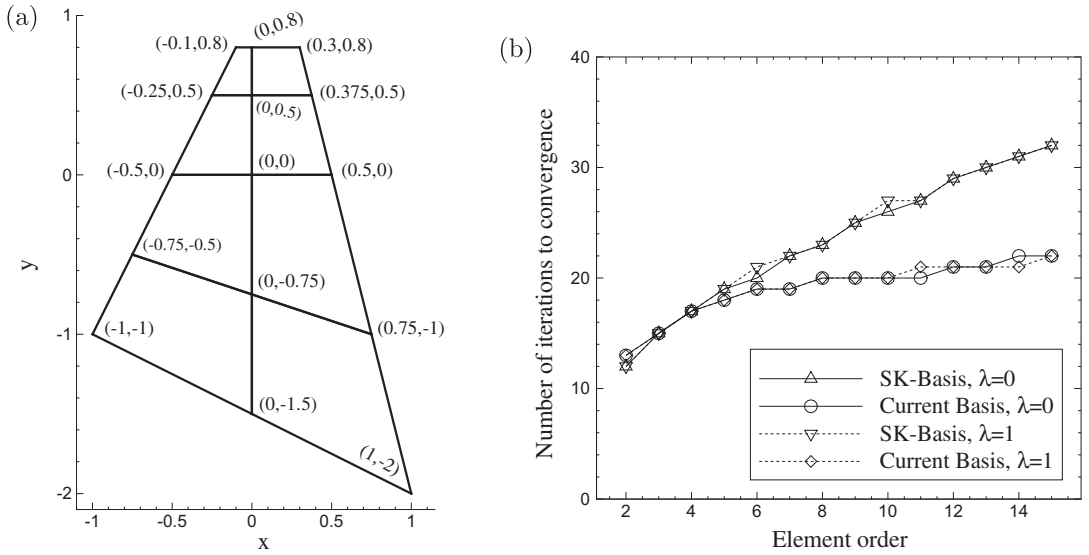

Fig. 19. 2D Helmholtz equation ($\lambda = 0$ and 1.0) on a domain of irregular geometry with 8 distorted quadrilateral elements (a), and the number of iterations to convergence as a function of element order (b) obtained with the current basis and Sherwin–Karniadakis basis (conjugate gradient solver with diagonal preconditioner).

Let us first consider the conditioning of the coefficient matrix and the number of iterations to convergence on the smaller mesh. Fig. 20(a) shows the condition numbers of the Schur-complemented global Helmholtz coefficient matrix, without preconditioning and with the diagonal preconditioning, as a function of the element order for the current basis on the smaller mesh. The condition numbers with Basis-LV and the Sherwin–Karniadakis basis have also been shown in the plot for comparison. It is evident that without preconditioning the condition numbers with the current basis are considerably smaller (by two orders of magnitude) than those with Basis-LV and the Sherwin–Karniadakis basis. The condition numbers with Basis-LV are also smaller than those with the Sherwin–Karniadakis basis. The diagonal preconditioning has reduced the condition numbers by an order of magnitude compared to the cases without preconditioning. However, the relative magnitudes of the condition numbers for the three bases are unchanged from the no-preconditioning case. With diagonal preconditioning, the current basis leads to condition numbers two orders of magnitude smaller than Basis-LV and the Sherwin–Karniadakis basis. At a low element order (order 3) the condition numbers from Basis-LV and the Sherwin–Karniadakis basis are about the same. But as the element order increases the condition numbers with the Sherwin–Karniadakis basis become larger than those with the Basis-LV.

In Fig. 20(b) we plot the number of iterations it takes the conjugate gradient solver to converge to a specified tolerance (10^{-8}) for the 3D Helmholtz equation, with the diagonal preconditioning, as a function of the element order on the smaller mesh of 4 elements for the current basis, Basis-LV, and the Sherwin–Karniadakis basis. One can observe that at a low element order (orders 2 or 3) the numbers of iterations with the three bases are the same. For element orders beyond 4 significant

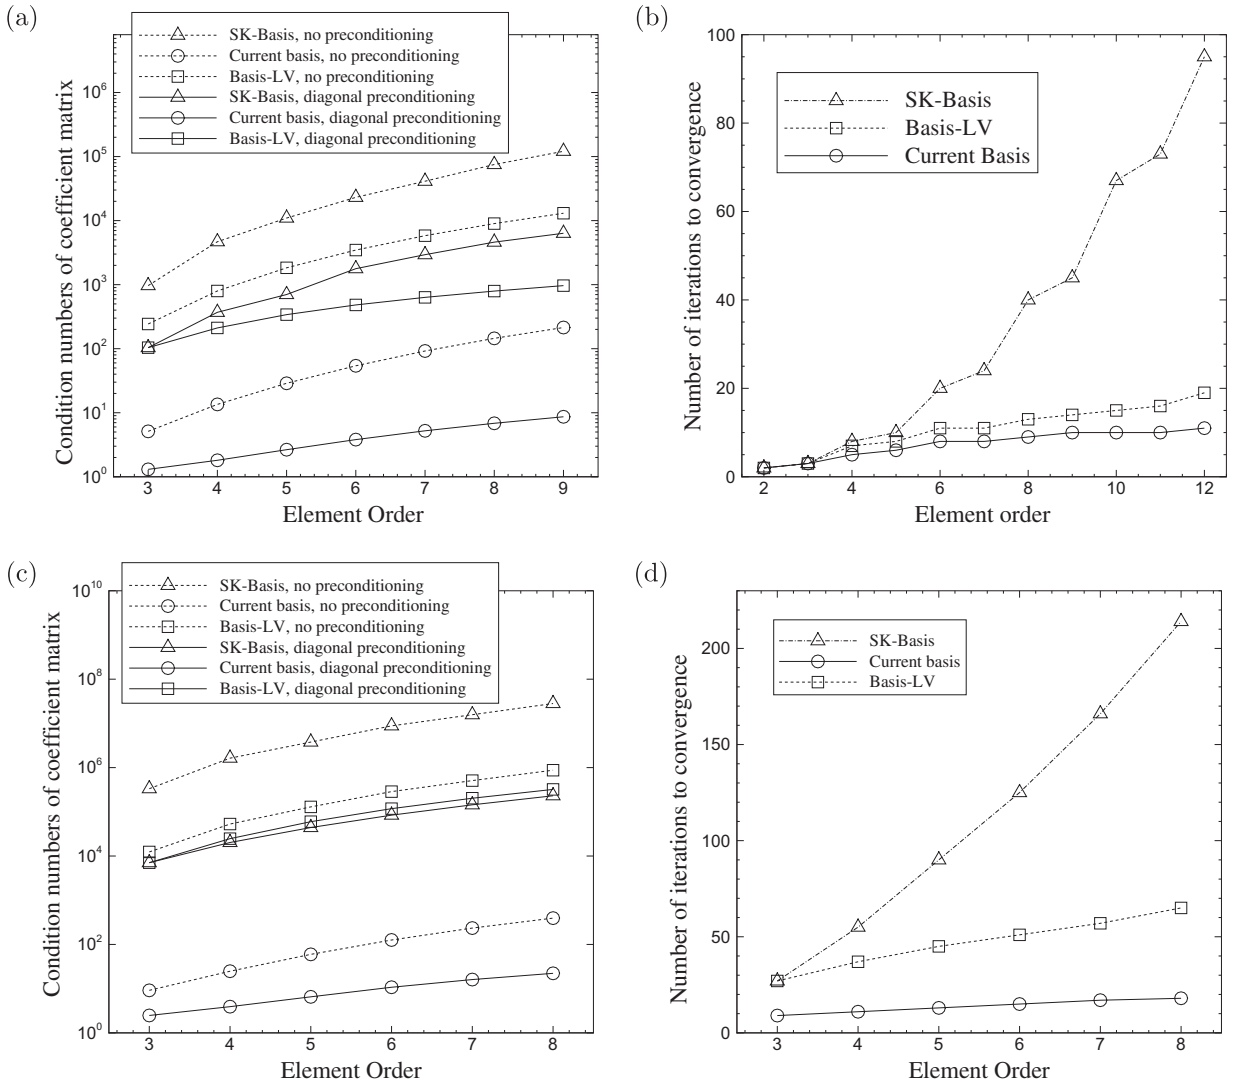

Fig. 20. Helmholtz equation in 3D ($\lambda = 1000$): Condition numbers of the Schur-complemented coefficient matrix (a,c) and the number of conjugate gradient iterations to convergence (tolerance 10^{-8}) (b,d) as a function of the element order. Plots (a) and (b) are for a mesh with 4 hexahedral elements. Plots (c) and (d) are for a mesh with 16 hexahedral elements.

differences can be observed in the number of iterations. The growth in the number of iterations for the Sherwin–Karniadakis basis is superlinear with respect to the element order, indicating that this basis does not work effectively with the diagonal preconditioning as the element order increases in 3D. The number of iterations for Basis-LV grows significantly slower than the Sherwin–Karniadakis basis, but faster than the current basis. For the current basis, only a slight increase in the number of iterations has been observed as the element order increases. For fairly large element orders (e.g. beyond order 8 for this problem) the number of iterations remains nearly a constant with increasing element order.

On the larger mesh of 16 elements we have observed the same trend about the matrix conditioning and the number of iterations as manifested on the smaller mesh. Fig. 20(c) and (d) respectively show the condition numbers of the Schur-complemented global coefficient matrix (without preconditioning and with diagonal preconditioning) and the number of conjugate gradient iterations to convergence (with diagonal preconditioning) as a function of the element order for the mesh of 16 hexahedral elements. The condition numbers of the current basis (with and without preconditioning) are considerably smaller than those of Basis-LV and the Sherwin–Karniadakis basis. The condition numbers of Basis-LV are smaller than those of the Sherwin–Karniadakis basis without preconditioning. But with diagonal preconditioning they are comparable to, and somewhat larger than, those of the Sherwin–Karniadakis basis. In term of the number of iterations, the current basis requires much fewer iterations to converge. For example, with an element order 6 and diagonal preconditioning, it takes 15 conjugate gradient iterations with the current basis to converge to a tolerance 10^{-8}, while it takes 51 and 125 iterations, respectively, with the Basis-LV and the Sherwin–Karniadakis basis to converge to the same tolerance. As the element order increases, only a slight increase in the number of iterations has been observed with the current basis. In contrast, the increase in the number of iterations with the Basis-LV and the Sherwin–Karniadakis basis is very significant with respect to the element order (Fig. 20(d)).

We next consider the 3D Helmholtz equation on a domain with irregular geometry. We will employ the 3D domain and the mesh of four distorted hexahedral elements shown in Fig. 11. We solve the Helmholtz equation with two small λ values ($\lambda = 0$ and 1.0) on this mesh with various element orders, assuming Dirichlet conditions on the boundaries of the domain. Fig. 21 shows the number of iterations to convergence (tolerance 10^{-8}) of the conjugate gradient solver as a function of the element order obtained with the current basis, Basis-LV and the Sherwin–Karniadakis basis. In these tests a diagonal preconditioner has been employed with the conjugate gradient solver. The numbers of iterations for $\lambda = 0$ and $\lambda = 1.0$ are the same with the tests considered here. Regarding the performance of the three expansion bases, we observe a trend similar to that from the previous cases with a regular geometry. As the element order increases, with the current basis we observe only a slight increase in the number of iterations, compared to Basis-LV and Sherwin–Karniadakis basis. On the other hand, the growth in the number of iterations to convergence with the other two bases is notably more significant, especially with the Sherwin–Karniadakis basis.

3.2. Unsteady diffusion equation in two dimensions

In the previous numerical examples, the numbers of elements involved in the problems are quite small and the domain has a relatively simple geometry. Let us next consider a test problem involving relatively larger numbers of elements and a domain with nontrivial geometry. We will consider the unsteady diffusion equation in two dimensions,

$$\frac{\partial u}{\partial t} = \nu \nabla^2 u + f(\mathbf{x}, t), \tag{21}$$

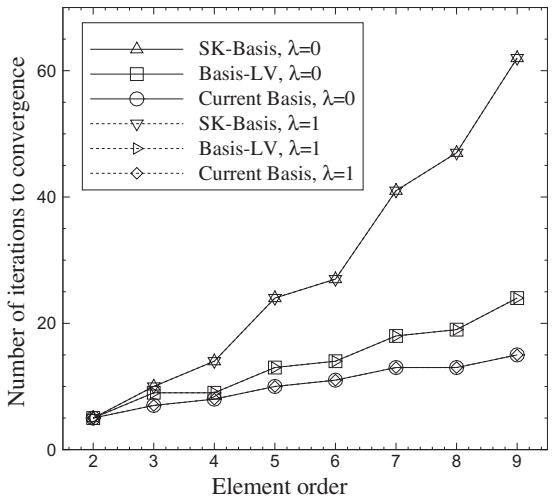

Fig. 21. 3D Helmholtz equation ($\lambda = 0$ and 1.0): Number of iterations to convergence as a function of element order (conjugate gradient solver with diagonal preconditioner), with the 4 distorted hexahedral elements on the irregular domain depicted in Fig. 11.

where $\mathbf{x}$ and t are respectively the spatial coordinate and the time. The constant $v > 0$ is the diffusion coefficient. $u(\mathbf{x},t)$ is the scalar field to be solved for. $f(\mathbf{x},t)$ is a prescribed field function, and we employ

$$f(\mathbf{x},t) = \cos 5x \sin 5y \sin 5t \tag{22}$$

in this test. We consider a rectangular domain in the x–y plane, $-10 \leqslant x \leqslant 15$ and $-10 \leqslant y \leqslant 10$, with a circular disk within (see Fig. 22). The disk has a diameter $D = 1.0$, and its center coincides with the origin of the coordinate system. We assume periodic boundary conditions on the boundaries in both x and y directions, that is,

$$\begin{aligned} u(-10,y,t) &= u(15,y,t), \\ u(x,-10,t) &= u(x,10,t). \end{aligned} \tag{23}$$

On the disk surface, $\Gamma = \{\mathbf{x} : |\mathbf{x}| = 0.5\}$, we impose the Dirichlet boundary condition, $u(\mathbf{x},t)|_{\mathbf{x}\in\Gamma} = 0$.We assume a zero initial condition for the field function, $u(\mathbf{x},0) = 0$.

For temporal discretization of Eq. (21), we employ the second-order backward differentiation formula (BDF2),

$$\frac{1}{\Delta t}\left(\frac{3}{2}u^{n+1} - 2u^n + \frac{1}{2}u^{n-1}\right) = v\nabla^2 u^{n+1} + f^{n+1}, \tag{24}$$

where n is the time step index, Δt is the time step size, and f^{n+1} is the function $f(\mathbf{x},t)$ evaluated at time step $n+1$. Re-arranging the above equation, we obtain a Helmholtz equation for u^{n+1},

$$\frac{3/2}{v\Delta t}u^{n+1} - \nabla^2 u^{n+1} = R, \tag{25}$$

where $R = \frac{1}{v}f^{n+1} + \frac{1}{v\Delta t}(2u^n - \frac{1}{2}u^{n-1})$.

We use the spectral element approach to discretize Eq. (25) in space, and employ the current expansion basis, Basis-LV, and the Sherwin–Karniadakis basis in the spectral element discretizations. When solving the resulting linear algebraic

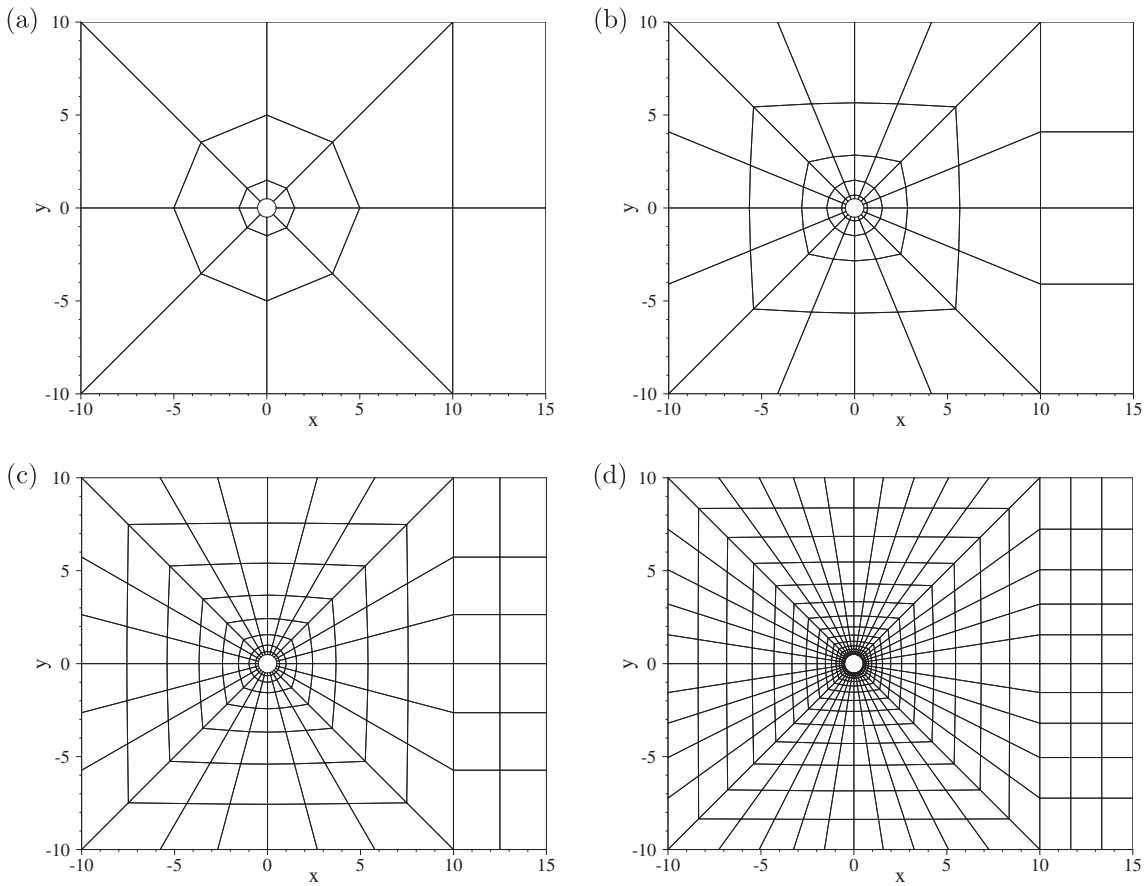

Fig. 22. Spectral element meshes with 26 (a), 84 (b), 204 (c), and 630 (d) quadrilateral elements.

system of the spectral element expansion coefficients, we first perform a Schur complement to condense out the interior degrees of freedom for all elements, as discussed in detail in the previous subsection.

We consider four spectral element meshes (Fig. 22) to discretize the domain, respectively with 26, 84, 204 and 630 quadrilateral elements. For a given mesh, we use the same element order for all the elements. With each spectral element mesh, we systematically increase the element order from 2 to 30, and for each element order we integrate the diffusion equation in time and record the number of iterations it takes the conjugate gradient solver to converge to a prescribed tolerance within a time step. The number of iterations per time step will be compared among different expansion bases.

Fig. 23 shows the number of iterations per time step as a function of the element order for the current expansion basis. The plots (a, b, c, d) are obtained with the four different meshes of 26, 84, 204 and 630 elements. The number of iterations obtained with the Sherwin–Karniadakis basis and the Basis-LV have also been included in the plots for comparison. These results are for a diffusion coefficient $\nu = 0.001$ and a time step size $\Delta t = 0.01$. In the computations the conjugate gradient solver has been preconditioned with the diagonal preconditioner. A tolerance of 10^{-8} has been used for the stopping criterion (reduction in the magnitude of residual relative to the initial residual) in the conjugate gradient solver.

We observe that with the current expansion basis, for a given spectral element mesh, the number of iterations per time step essentially remains a constant as the element order increases (Fig. 23). As the number of elements becomes large in the spectral element mesh, there is a slight increase in the number of iterations per time step with increasing element order for the current basis. But this dependence is very weak. For example, with the mesh of 630 elements the number of iterations per step increases only from 15 to 19 as the element order increases from 17 to 30 (Fig. 23(d)).

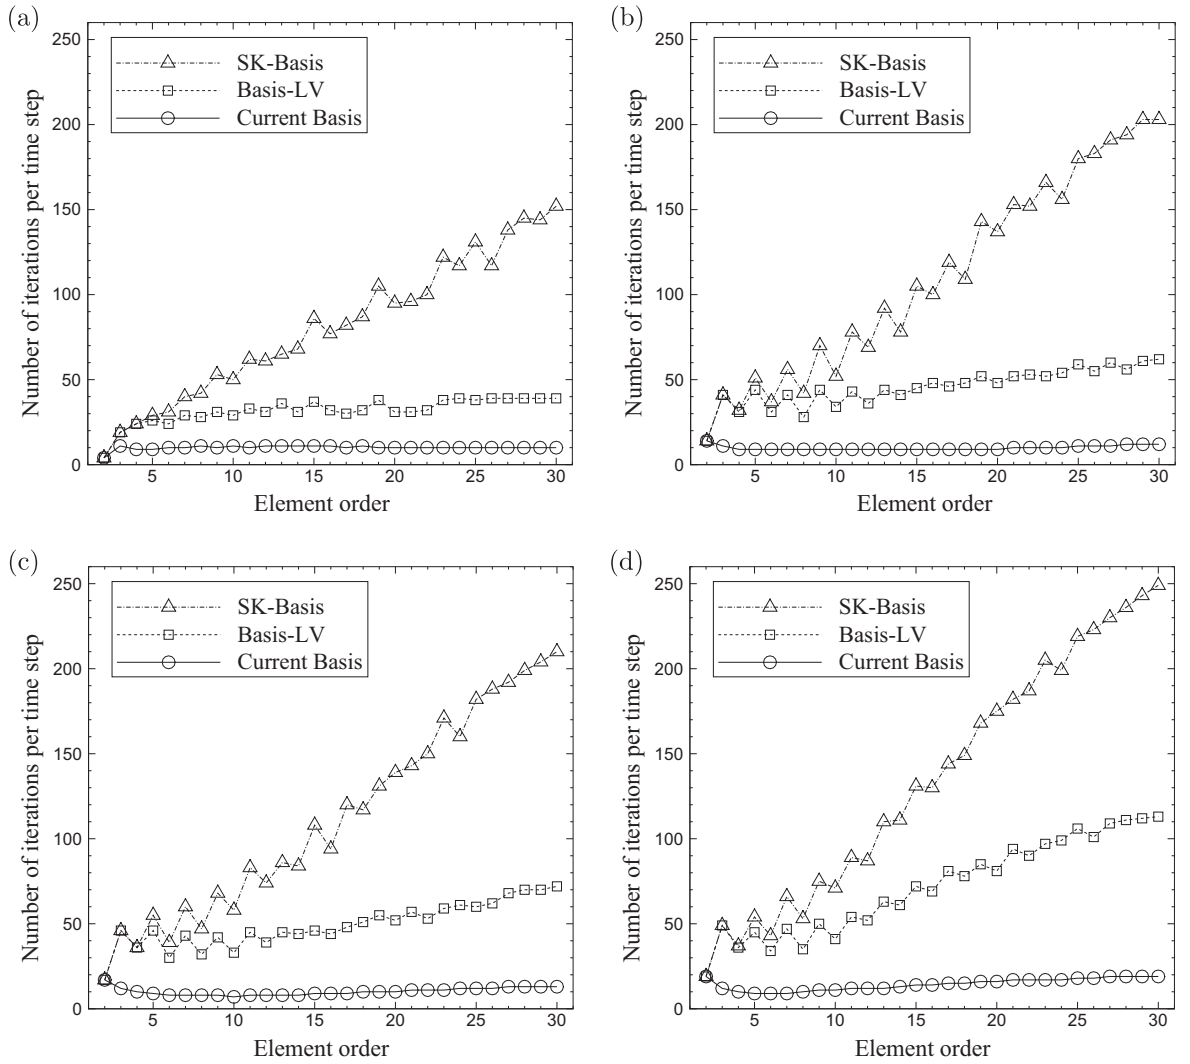

Fig. 23. Number of conjugate gradient iterations per time step as a function of the element order for the unsteady diffusion equation ($\nu = 0.001, \Delta t = 0.01$) discretized with spectral element meshes with 26 (a), 84 (b), 204 (c), and 630 (d) quadrilateral elements.

With the Basis-LV, we observe that, the number of iterations per time step is significantly larger, generally by five-fold, than that with the current expansion basis. For example, with the mesh of 204 elements and an element order 15, it takes 46 conjugate gradient iterations per time step with Basis-LV (Fig. 23(c)) compared to 9 iterations per step with the current expansion basis. When the number of elements in the mesh becomes large, the increase in the number of iterations per step with respect to the element order for Basis-LV becomes very notable. For example, with the mesh of 630 elements, the number of iterations per time step with Basis-LV increases from 81 to 113 as the element order increases from 17 to 30. Note that the difference between Basis-LV and the current basis lies in the 1D vertex modes employed in the construction of the basis functions. In the current basis, the 1D vertex modes are given by Eqs. (9) and (11). In the Basis-LV, 1D linear vertex modes given by Eqs. (8) and (10) are employed in the construction. The difference in the 1D vertex modes in the construction of current basis and Basis-LV has caused an enormous difference in the performance, and the current basis is vastly superior.

With the Sherwin–Karniadakis basis, on the other hand, we observe a near-linear increase in the number of iterations per step with respect to the element order for a given mesh (Fig. 23), and the rate of growth in the number of iterations also appears to increase as the number of elements in the mesh increases. Note that the diagonal preconditioning has been used with the conjugate gradient solver here. This again shows that the Sherwin–Karniadakis basis does not work effectively with the diagonal preconditioning. More sophisticated preconditioners would be required (e.g. [25]) in order for this basis to work efficiently. Note the zigzag pattern in the number of iterations in the curves obtained with the Sherwin–Karniadakis basis and the Basis-LV in the plots. The zigzag pattern has also been observed in other spectral element computations, and is a genuine phenomenon. It is possibly related to the characteristics of the solution, the mesh non-uniformity, element distribution, and the mesh resolutions.

In Fig. 24 we examine the growth in the number of iterations per time step with respect to the size of the mesh for a fixed element order. Here we plot the number of iterations per step as a function of the number of elements in the mesh for a moderate element order 6 and for a fairly large element order 20. One can observe that with the fixed moderate element order 6 , the number of iterations per time step with the current expansion basis remains a constant as the number of elements in the mesh increases. The same observation also applies to the Sherwin–Karniadakis basis and Basis-LV for the meshes with 84 elements or larger. With an element order 20, for the current basis, there is only a slight increase in the number of iterations per step with respect to the number of elements (for the two largest meshes). For the Sherwin–Karniadakis basis and the Basis-LV, the increase in the number of iterations with respect to the number of elements is notably more significant with an element order 20.

In Fig. 25 we compare the numbers of conjugate-gradient iterations per time step obtained with different expansion bases on the mesh with 204 elements, for a set of parameters with increased diffusion coefficient ν and time step size Δt values ($\nu = 0.1$ and $\Delta t = 0.1$). These parameters result in a smaller coefficient in the u^{n+1} term in the resultant Helmholtz equation after temporal discretization; see Eq. (25). A diagonal preconditioner has been employed in the conjugate gradient solver. We observe from the data a trend similar to that of previous cases. One can observe that the current expansion basis requires significantly fewer iterations per time step compared to Basis-LV and the Sherwin–Karniadakis basis, and that the growth in the number of iterations with increasing element order for the current basis is also notably smaller than the other two bases.

Finally we consider the wall time associated with different expansion bases. In Table 1 we summarize the wall-clock time of simulations of the 2D unsteady diffusion equation ($\nu = 0.001, \Delta t = 0.01$) on the mesh of 630 quadrilateral elements, which corresponds to Fig. 23(d), for element orders 9 and 15. The timing was collected on a single processor of an Intel Xeon workstation cluster, and the values in the table are the average over a number of time steps. The current problem involves

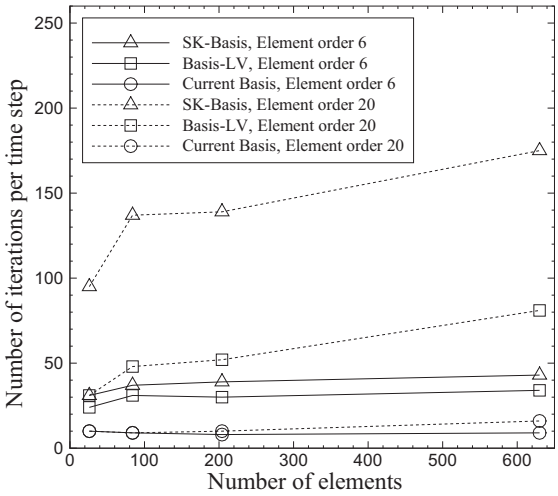

Fig. 24. Number of iterations per time step as a function of the number of elements for two fixed element orders 6 and 20.

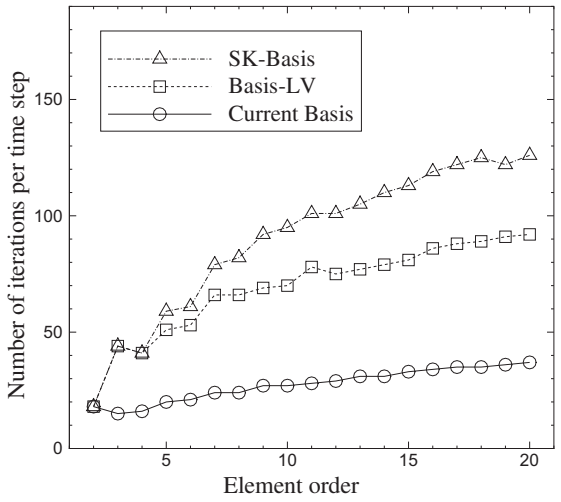

Fig. 25. Number of iterations per time step as a function of element order for 2D unsteady diffusion equation ($\nu = 0.1, \Delta t = 0.1$), discretized with 204 quadrilateral elements. A conjugate gradient solver with a diagonal preconditioner has been used.

Table 1
Wall-time per time step (in seconds) for solving 2D unsteady diffusion equation ($\nu = 0.001, \Delta t = 0.01$) on the mesh of 630 elements (see Fig. 23(d)) with element orders 9 and 15 . The column "Force" refers to wall time used for evaluating the force field (Eq. (22)); the column "Solver" refers to wall time within a time step excluding the force evaluation time.

	Order 9		Order 15	
	Force	Solver	Force	Solver
Current Basis	0.2	0.039	0.51	0.17
Basis-LV	0.2	0.11	0.51	0.40
SK-Basis	0.2	0.16	0.51	0.74

a time-dependent analytic force field function (Eq. (22)), which needs to be evaluated on the quadrature points at every time step. In the current implementation, we have used a parser routine external to the program code for evaluating this force expression, instead of hardwiring the force function into the code to use the internal mathematical functions inherent in the programming language. Such an implementation, although extremely general (can handle any kind of functions without re-compiling the code, not limited to the sinusoidal functions in Eq. (22)), does induce a notable amount of overhead because of the external parser routine. Also note that in simulations of practical problems (such as in computational fluid dynamics) a 3D force field will not likely be involved in most cases. In Table 1, we have therefore separated the wall-time costs used for the force evaluation (column "Force") from the rest of the computation time within a time step (column "Solver"). Note that the wall time in the "Solver" column includes the cost of all computations in a time step (excluding force evaluation), besides the actual time spent in the conjugate gradient solver. From Table 1, one can observe that the force evaluation time is quite significant compared to the rest of solver time within a time step due to the external parser routine. One can also observe that the solver time (excluding force evaluation) with the current expansion basis is significantly less than those with Basis-LV and the Sherwin–Karniadakis basis.

4. Concluding remarks

In this paper we have presented a high-order expansion basis for structured spectral elements. The new basis exhibits a superior numerical efficiency to the commonly-used Jacobi polynomial-based expansion basis [17], in terms of the conditioning of the coefficient matrix and the number of iterations to convergence with the conjugate gradient solver. It results in extremely sparse mass matrices with low condition numbers. A notable advantage of the new basis is that it is very amenable to, and can work effectively with, the diagonal preconditioning. Numerical experiments demonstrate that, with the new expansion basis and a simple diagonal preconditioner, the number of conjugate gradient iterations to convergence has only a very weak dependence on the element order. Because the diagonal preconditioner is numerically very simple and it can be trivially parallelized, one can appreciate that the presented basis enjoys a tremendous advantage in numerical efficiency.

The construction of the new expansion basis in two and three dimensions relies on the numerical construction of a set of special 1D basis functions. The 1D interior modes are constructed in a fashion such that the 1D interior mass matrix and

interior stiffness matrix are simultaneously diagonal, and moreover they have identical condition numbers. The 1D vertex modes are constructed through a Gram–Schmidt process such that they are orthogonal to all the 1D interior modes. The construction of the special 1D vertex modes is critical to the sparsity of coefficient matrices and to the superior conditioning and numerical efficiency enjoyed by the current expansion basis. Much of the efficiency and the matrix sparsity will be lost if the current 1D vertex modes are replaced by the commonly used linear vertex functions.

We have only considered the new expansion basis for structured elements in the current paper. Extension of the current construction to unstructured elements is nontrivial, primarily owing to the desire to retain the special vertex modes, which causes complications to maintaining an appropriate polynomial space for unstructured elements. These issues and a construction for unstructured spectral elements will be the focus of a subsequent paper.

Acknowledgements

The authors gratefully acknowledge the support from ONR and NSF. Computer time was provided by the TeraGrid through a TRAC grant.

References

[1] R. Abdul-Rahman, M. Kasper, Higher order triangular basis functions and solution performance of the CG method, Comput. Methods Appl. Mech. Eng. 197 (2007) 115–127.
[2] S. Adjerid, M. Aiffa, J.E. Flaherty, Hierarchical finite element bases for triangular and tetrahedral elements, Comput. Methods Appl. Mech. Eng. 190 (2001) 2925–2941.
[3] M. Ainsworth, J. Coyle, Hierarchic finite element bases on unstructured tetrahedral meshes, Int. J. Numer. Methods Eng. 58 (2003) 2103–2130.
[4] I. Babuska, M. Griebel, J. Pitkaranta, The problem of selecting the shape functions for a p-type finite element, Int. J. Numer. Methods Eng. 36 (1989) 3759–3779.
[5] S. Beuchler, V. Pillwein, Sparse shape functions for tetrahedral p-FEM using integrated Jacobi polynomials, Computing 80 (2007) 345–375.
[6] M.L. Bittencourt, M.G. Vazquez, T.G. Vazquez, Construction of shape functions for the h- and p-versions of the FEM using tensorial product, Int. J. Numer. Methods Eng. 71 (2007) 529–563.
[7] P. Carnevali, R.B. Morris, Y. Tsuji, G. Taylor, New basis functions and computational procedures for p-version finite element analysis, Int. J. Numer. Methods Eng. 36 (1993) 3759–3779.
[8] S. Dong, Direct numerical simulation of turbulent Taylor–Couette flow, J. Fluid Mech. 587 (2007) 373–393.
[9] S. Dong, Turbulent flow between counter-rotating concentric cylinders: a direct numerical simulation study, J. Fluid Mech. 615 (2008) 371–399.
[10] S. Dong, J. Shen, An unconditionally stable rotational velocity-correction scheme for incompressible flows, J. Comput. Phys. 229 (2010) 7013–7029.
[11] S. Dong, Z. Yosibash, A parallel spectral element method for dynamic three-dimensional nonlinear elasticity problems, Comput. Struct. 87 (2009) 59–72.
[12] S. Dong, X. Zheng, Direct numerical simulation of spiral turbulence, J. Fluid Mech. 668 (2011) 150–173.
[13] M. Dubiner, Spectral methods on triangles and other domains, J. Sci. Comput. 6 (1991) 345.
[14] N.B. Edgar, K.S. Surana, On the conditioning number and the selection criteria for p-version approximation functions, Comput. Struct. 60 (1996) 521–530.
[15] G.H. Golub, C.F. van Loan, Matrix Computations, 3rd ed., Johns Hopkins University Press, 1996.
[16] J.S. Hesthaven, From electrostatics to almost optimal nodal sets for polynomial interpolation in a simplex, SIAM J. Numer. Anal. 35 (1998) 655–676.
[17] G.E. Karniadakis, S.J. Sherwin, Spectral/hp Element Methods for Computational Fluid Dynamics, 2nd ed., Oxford University Press, 2005.
[18] A.J. Laub, M.T. Heath, C.C. Paige, R.C. Ward, Computation of system balancing transformations and other applications of simultaneous diagonalization algorithms, IEEE Trans. Auto. Control AC-32 (1987) 115–122.
[19] B.C. Moore, Principal component analysis in linear systems: controllability, observability, and model reduction, IEEE Trans. Auto. Control AC-26 (1981) 17–32.
[20] S.A. Orszag, Spectral methods for problems in complex geometries, J. Comput. Phys. 37 (1980) 70.
[21] A.T. Patera, A spectral method for fluid dynamics: laminar flow in a channel expansion, J. Comput. Phys. 54 (1984) 468.
[22] Y. Saad, Iterative Methods for Sparse Linear Systems, SIAM, 2003.
[23] J. Shen, L.-L. Wang, Fourierization of the Legendre–Galerkin method and a new space-time spectral method, Appl. Numer. Math. 57 (2007) 710–720.
[24] J. Shen, L.-L. Wang, H. Li, A triangular spectral element method using fully tensorial rational basis functions, SIAM J. Numer. Anal. 47 (2009) 1619–1650.
[25] S.J. Sherwin, M. Casarin, Low-energy basis preconditioning for elliptic substructured solvers based on unstructured spectral/hp element discretization, J. Comput. Phys. 171 (2001) 394–417.
[26] S.J. Sherwin, G.E. Karniadakis, A triangular spectral element method: applications to the incompressible Navier-Stokes equations, Comput. Methods Appl. Mech. Eng. 123 (1995) 189–229.
[27] P. Solin, T. Vejchodsky, High-order finite elements based on generalized eigenfunctions of the Laplacian, Int. J. Numer. Meth. Eng. 73 (2008) 1374–1394.
[28] B.A. Szabó, I. Babuška, Finite Element Analysis, John Wiley & Sons, New York, 1991.
[29] M.A. Taylor, B.A. Wingate, R.E. Vincent, An algorithm for computing Fekete points in the triangle, SIAM J. Numer. Anal. 38 (2000) 1707–1720.
[30] L.N. Trefethen, D. Bau, Numerical Linear Algebra, SIAM, 1997.
[31] T.C. Warburton, S.J. Sherwin, G.E. Karniadakis, Basis functions for triangular and quadrilateral high-order elements, SIAM J. Sci. Comput. 20 (1999) 1671–1695.
[32] J.P. Webb, R. Abouchacra, Hierarchical triangular elements using orthogonal polynomials, Int. J. Numer. Meth. Eng. 38 (1995) 245–257.