A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database

Citation: Winokur, Justin, Patrick Conrad, Ihab Sraj, Omar Knio, Ashwanth Srinivasan, W. Carlisle Thacker, Youssef Marzouk, and Mohamed Iskandarani. "A Priori Testing of Sparse Adaptive Polynomial Chaos Expansions Using an Ocean General Circulation Model Database." Computational Geosciences 17, no. 6 (July 4, 2013): 899–911.

As Published: http://dx.doi.org/10.1007/s10596-013-9361-3

Publisher: Springer Netherlands

Persistent URL: http://hdl.handle.net/1721.1/106563

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Comput Geosci (2013) 17:899–911
DOI 10.1007/s10596-013-9361-3

ORIGINAL PAPER

A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database

Justin Winokur · Patrick Conrad · Ihab Sraj · Omar Knio · Ashwanth Srinivasan · W. Carlisle Thacker · Youssef Marzouk · Mohamed Iskandarani

Received: 20 December 2012 / Accepted: 12 June 2013 / Published online: 4 July 2013

Abstract This work explores the implementation of an adaptive strategy to design sparse ensembles of oceanic simulations suitable for constructing polynomial chaos surrogates. We use a recently developed pseudo-spectral algorithm that is based on a direct application of the Smolyak sparse grid formula and that allows the use of arbitrary admissible sparse grids. The adaptive algorithm is tested using an existing simulation database of the oceanic response to Hurricane Ivan in the Gulf of Mexico. The a priori tests demonstrate that sparse and adaptive pseudo-spectral constructions lead to substantial savings over isotropic sparse sampling in the present setting.

J. Winokur · I. Sraj · O. Knio
Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA

P. Conrad · Y. Marzouk
Department of Aeronautics and Astronautics, Massachusetts Insitute of Technology, Cambridge, MA 02139, USA

A. Srinivasan · M. Iskandarani (✉)
Rosenstiel School of Marine and Atmospheric Science, University of Miami, 4600 Rickenbacker Causeway, Miami FL 33149, USA
e-mail: miskandarani@rsmas.miami.edu

W. C. Thacker
Cooperative Institute of Marine and Atmospheric Sciences, Miami, FL 33149, USA

W. C. Thacker
Atlantic Oceanographic and Meteorological Laboratory, Miami, FL 33149, USA

Keywords Uncertainty quantification · Polynomial chaos · Sparse Smolyak quadrature · Adaptive sampling · Ocean modeling

1 Introduction

Polynomial chaos (PC) methods are series-based approaches developed in recent years to quantify the uncertainties in the output of numerical simulations due to uncertainties in their input data and to explicitly represent the dependence of specific quantities of interests (QoIs) on this uncertain data (see [1] and references therein). One of the main advantages of such representations is that the response of a highly complex, computationally expensive model to random inputs can be reasonably well represented by a simple functional form. The PC representation can then be used as a surrogate for the full model, thus enabling efficient sampling, e.g., for the purpose of inference [2–4] or sensitivity analysis [1, 5, 6]. Recent applications of PC methods to oceanic simulations can be found in [7, 8] and [9].

There are two traditional methods to construct PC representations. The so-called intrusive method [1, 10] is based on injecting the polynomial expansion into the model equations and on applying a Galerkin-type formalism to derive model equations for the unknown coefficients. In contrast, nonintrusive methods rely on individual model realizations, i.e., deterministic runs corresponding to particular (fixed) values of the random inputs, to create the desired PC representations of selected QoIs. One of the advantages of the nonintrusive technique in the context of complex models is that it avoids the need to modify existing or legacy codes.

In this work, we shall specifically focus on a particular class of nonintrusive methods, namely one relying on

spectral projection methodology [1]. In this approach, determination of PC coefficients of a specific QoI amounts to performing a multidimensional quadrature. Unfortunately, traditional quadrature approaches suffer the so-called curse of dimensionality, whereby the number of realizations scale exponentially with the number of dimensions. Sparse quadratures [11, 12] provide one avenue to mitigate this high computational cost.

In prior work [8], we relied on a classical Smolyak sparse grid tensorization [13] to study the Hybrid Coordinate Ocean Model (HYCOM)-simulated sea surface temperature's (SST) dependence on the combined uncertainties in subgrid-scale mixing and wind drag parameters. To obtain accurate statistics, Alexanderian et al. [8] relied on a systematic isotropic refinement of the sparse grid, namely by increasing the resolution level along all dimensions of the random parameter space. As further described below, this necessitated the generation of a relatively large database involving hundreds of realizations.

The present study is motivated by the desire to explore, in the context of short-duration Ocean General Circulation Model predictions under extreme forcing events, whether adaptive sparse quadrature can afford the construction of PC surrogates with equal or improved fidelity to isotropic refinement, based on a substantially smaller number of realizations and consequently at a small fraction of the computational cost. It is also motivated by ongoing investigations [9] of the oceanic circulation in larger domains than was considered in our previous Gulf of Mexico (GOM) analysis [8] and for longer simulation times. Since the cost of such extended analyses effectively limits the number of realizations that can be simulated, it proved highly desirable to analyze the performance of adaptive sampling schemes prior to launching them to dynamically sample large-scale computations. Consequently, in this paper, we exploit the preexisting database to investigate the properties of a recent pseudo-spectral algorithm [14, 15] that is based on a direct application of the Smolyak sparse grid formulas. The advantage of this generalized approach is that it affords the use of arbitrary admissible sparse grids and it maximizes the number of polynomial coefficients that can be computed without internal aliasing.

The layout of this paper is as follows: In Section 2, we discuss polynomial chaos expansions and the use of nonintrusive spectral projection (NISP) to obtain the coefficients. In Section 3, we discuss the computational model and the previously computed database. We also discuss the enrichment of the preexisting database in light of the global sensitivity analysis in [8]. In Section 4, we introduce the Smolyak pseudo-spectral projection method and apply it to the realization database. In Section 5, we analyze the performance of the adaptive algorithms and demonstrate an order-of-magnitude reduction in required realizations. Finally, in Section 6, we discuss potential improvements on our methodology and applications to other problems.

2 Spectral approximation

In this work, we restrict our attention to the case where the random model inputs can be parametrized in terms of a canonical, d-dimensional random vector $\boldsymbol{\xi}$ where d is the number of uncertain parameters. Let $U(\boldsymbol{\xi}, \boldsymbol{x}, t)$ denote a $\boldsymbol{\xi}$-dependent QoI. The truncated PC expansion of U takes the form

$$U(\boldsymbol{\xi}, \boldsymbol{x}, t) \doteq \sum_{k=0}^{P} U_k(\boldsymbol{x}, t)\Psi_k(\boldsymbol{\xi}) \tag{1}$$

where $U_k(\boldsymbol{x}, t)$ is the series coefficients and the $\Psi_k(\boldsymbol{\xi})$ form an orthogonal basis with respect to the weighted inner product

$$\langle U, V\rangle = \int_{\Omega} U(\boldsymbol{\xi})V(\boldsymbol{\xi})w(\boldsymbol{\xi})\mathrm{d}\boldsymbol{\xi}, \text{ so that} \langle \Psi_i, \Psi_j\rangle = \delta_{i,j}\|\Psi_i\|^2. \tag{2}$$

The choice of basis functions is largely dictated by the weight function, $w(\boldsymbol{\xi}) > 0$. When the components of $\boldsymbol{\xi}$ are independent and uniformly distributed, as in the present case, the basis functions are multidimensional Legendre polynomials [16]. PC representations allow us to represent smooth QoIs with rapidly converging expansions [1], readily determine statistical moments [1] such as the mean and variance of QoIs, compute model output sensitivity to input parameters [5, 8], and substantially speed up the solution of inverse problems [4, 9, 17].

The main computational burden of PC methods is the determination of the coefficients U_k. As previously mentioned, in this work, we focus exclusively on NISP, which aims at minimizing the L_2 norm of the difference between the function $U(\boldsymbol{\xi})$ and its truncated series representation (1). This is achieved by taking the inner product of the series with each basis polynomial, invoking the orthogonality of the basis, and replacing the ensuing integrals with numerical quadrature to get

$$\begin{aligned} U_k &= \frac{\langle U, \Psi_k\rangle}{\langle \Psi_k^2\rangle} \approx \frac{\langle U, \Psi_k\rangle_Q}{\langle \Psi_k^2\rangle} \text{ with } \langle U, \Psi_k\rangle \approx \langle U, \Psi_k\rangle_Q \\ &\equiv \sum_{q=1}^{Q} U(\boldsymbol{\xi}_q)\Psi_k(\boldsymbol{\xi}_q)\omega_q \end{aligned} \tag{3}$$

where $\boldsymbol{\xi}_q$ and ω_q are appropriate multidimensional quadrature points and weights.

The determination of U_k thus becomes a multidimensional quadrature problem. Traditionally, multidimensional

quadrature is built using tensor products of 1D quadrature rules, such as Gaussian quadrature. This approach, however, suffers from the curse of dimensionality as the number of realizations scales exponentially with dimension, namely as N^d where N is the number of nodes in the underlying 1D quadrature rule and d is the number of stochastic dimensions. As N and d become large, it is imperative to devise alternate quadrature rules with smaller computational cost. An optimal quadrature would maximize accuracy for a fixed number of function evaluations (realizations), would be adaptive and nested in order to allow incremental improvements while capitalizing on previous work, and would exhibit a slower growth pattern as the number of dimensions increases. Finally, it is important to point out that any proposed quadrature rule must have minimum exactness requirements so as to preserve the orthogonality of the basis functions retained in the series and, in so doing, avoid the deleterious effects of internal aliasing (see [15] and Appendix A).

Two avenues for improving the traditional tensor product-based rules are the use of alternative tensorization and/or the use of nested 1D rules such as the Clenshaw-Curtis, Fejèr, or Gauss–Kronrod–Patterson rules [11, 18]. A well-known approach is the Smolyak quadrature [11, 13, 19]. The key idea is to represent the multidimensional quadrature as a telescoping sum of lower level quadrature rules. Let Q_ℓ represent the 1D quadrature rule operator where ℓ refers to the quadrature level, and let $Q_\ell^d U = \left(Q_{\ell_1} \otimes ... \otimes Q_{\ell_d}\right) U$ be the corresponding d-dimensional rule. Furthermore, define the difference operator $\Delta_\ell = Q_\ell - Q_{\ell-1}$ with $\Delta_0 = Q_0 = 0$. Rewriting Q_ℓ in terms of Δ_ℓ, inserting the resulting expression in the tensor product,

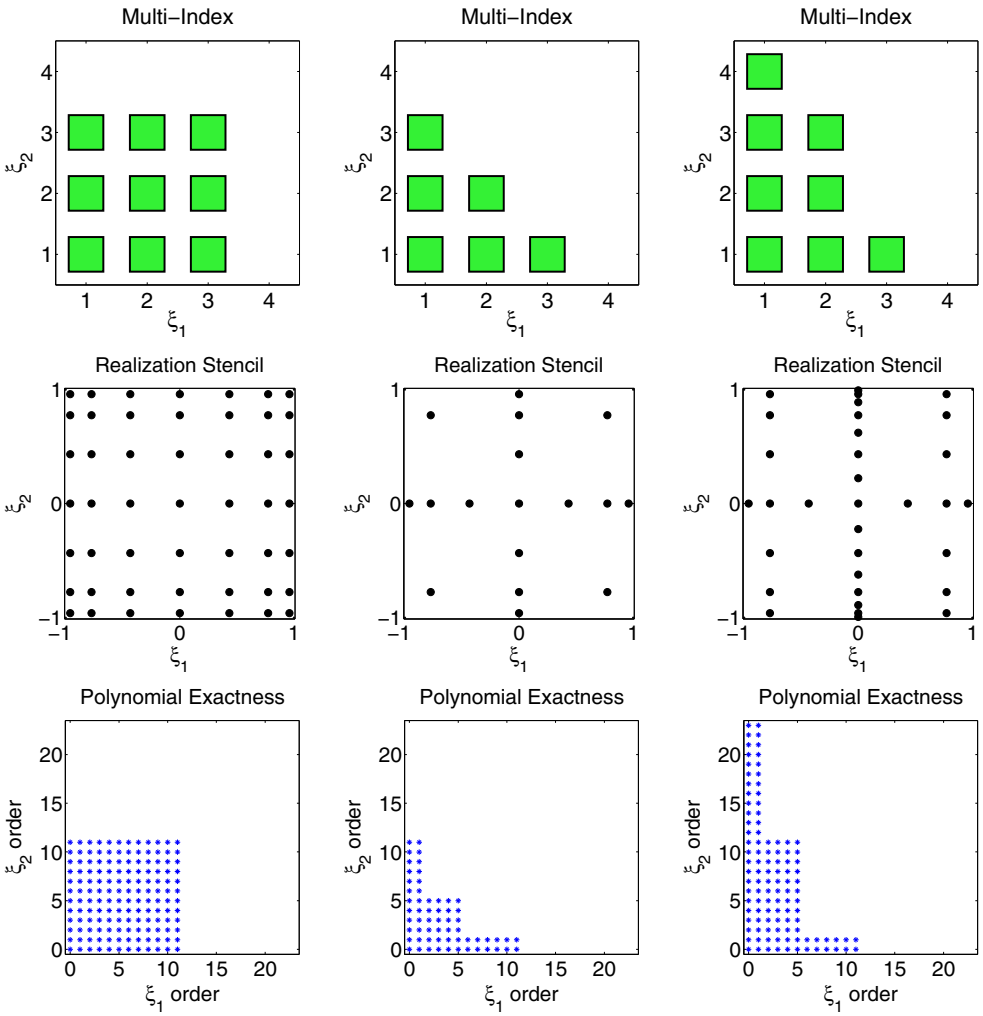

(a) "Full-Tensor" 49 Realizations

(b) ClassicSmolyak 17 Realizations

(c) ArbitraryMulti-Index 33 Realizations

Fig. 1 Examples of 2D tensorizations, showing the multi-index construction (*top*), the realization stencil (*middle*), and the resulting polynomial exactness (*bottom*). The realization stencil and polynomial exactness correspond to Gauss–Kronrod–Patterson 1D rules [11, 18]. **a** An *isotropic* full-tensor representation of level $\ell = 3$ multi-index such that $\boldsymbol{\mathcal{F}} = \{\boldsymbol{k} : ||\boldsymbol{k}||_\infty \leq \ell\}$. **b** A classical Smolyak representation of level $\ell = 3$ where multi-index is such that $\boldsymbol{\mathcal{F}} = \{\boldsymbol{k} : ||\boldsymbol{k}||_1 \leq \ell + d - 1\}$. **c** General anisotropic multi-index

and switching the order of tensorization and summation, $Q_\ell^d U$ can now be rewritten as a telescoping sum over a multi-index set, $\mathcal{F}$:

$$Q^d U = \sum_{\boldsymbol{k}\in\mathcal{F}} \left(\Delta_{k_1} \otimes \ldots \otimes \Delta_{k_d}\right) U. \tag{4}$$

Here, $\boldsymbol{k} = (k_1, k_2, \ldots, k_d)$ is a vector of integer indices referring to the quadrature level in each direction and such that $k_i \leq \ell_i$. Different types of Smolyak quadrature can be obtained with different choices of the index set $\mathcal{F}$. This is illustrated in Fig. 1 for a simple 2D example. When the quadrature level is allowed to roam the hyper-rectangle $k_i \leq \ell_i$, alternatively $\|\boldsymbol{k}\|_\infty \leq l$, we obtain the traditional full tensorization shown in Fig. 1a. The classical Smolyak quadrature is obtained when $\mathcal{F}$ contains only the multi-indices, satisfying $\|\boldsymbol{k}\|_1 \leq \ell + d - 1$ as shown in Fig. 1b. The last column in Fig. 1 shows an arbitrary (admissible) index set $\mathcal{F}$, demonstrating that multidimensional quadrature need not be isotropic. Not only can different 1D rules be used in different directions but a quadrature with more specific multidimensional order can also be created [19].

3 HYCOM database

The interest in improving the quadrature sampling of polynomial chaos expansions (PCEs) came about during a global sensitivity analysis of the uncertainty in subgrid mixing and wind drag parameters of an ocean general circulation model (HYCOM) [8]. Here, we summarize this prior work and motivate the need for more effective sampling strategies by highlighting the drawbacks of isotropic refinement.

The work in [8] revolved around studying the impact of four uncertain subgrid-scale parameters on the simulated oceanic response to Hurricane Ivan in 2004. Three parameters were connected with the mixed-layer parameterization and one with the wind drag coefficient; the uncertainty of these four parameters is listed in Table 1. The main output uncertainties, or QoIs, correspond to averaged value of sea surface height (SSH), SST, and mixed-layer depth (MLD), within a geographic region whose bottom-left and upper-right corners are at 90W/20N and at 82W/30N, respectively; in addition, the heat flux (Q_{tot}) in a 150-km region surrounding the observed center of Ivan was also investigated. The HYCOM computational grid covered the GOM and a large portion of the Caribbean Sea at a resolution of $1/25^\circ \approx 4$ km and 20 vertical layers; the simulation spanned the entire transit of Ivan through the Gulf of Mexico from 9 September to 14 September 2004. Each realization required 2.5 h on an eight-core node and required 764 megabytes of storage. This large computational burden is the primary motivation behind the present work.

Table 1 The random input parameters for HYCOM

Parameter	Description	Distribution
θ_1	Critical Richardson number	$\mathcal{U}(0.25, 0.7)$
θ_2	Background viscosity (m^2/s)	$\mathcal{U}(10^{-4}, 10^{-3})$
θ_3	Background diffusivity (m^2/s)	$\mathcal{U}(10^{-5}, 10^{-4})$
θ_4	Stochastic wind drag coefficient	$\mathcal{U}(0.2, 1.0)$

Here, $\mathcal{U}[a_i, b_i]$ designates a uniformly distributed probability density function over the interval $a_i \leq \theta \leq b_i$; the canonical stochastic variables $|\xi_i| \leq 1$ are then defined as $\theta_i(\xi_i) = \frac{a_i+b_i}{2} + \frac{b_i-a_i}{2}\xi_i$

The PCE basis consisted of tensor products of Legendre polynomials, and the series was truncated isotropically at fifth order in each direction. The numerical quadrature consisted of a classical Smolyak sparse grid based on the 1D Gauss–Kronrod–Patterson rule. An ensemble of 385 realizations was required to pseudo-spectrally compute polynomials of order ≤ 5. This represented a substantial amount of computational time and storage given the grid's spatial resolution and extent. The adequacy of the PCE representation was checked via a number of error metrics [8]. Global sensitivity analysis revealed that the uncertainty in the QoIs was dominated by the third and fourth parameters (the background diffusivity and wind drag coefficient, respectively) whereas the first two contributed little (see Fig. 2).

Additional insight into the response of the oceanic circulation to the uncertain parameters can be gained from Fig. 3, which depicts projections of the box-averaged SST onto the axes of the random parameter space. The plots reveal an almost linear dependence of SST on ξ_1, ξ_2, and ξ_3 (panels a, b, and c, respectively), which suggests that a lower order PCE along these dimensions would be sufficient. In contrast, Fig. 3d reveals a highly curved response for higher values of ξ_4, which indicates that a higher order expansion is required to suitably represent it.

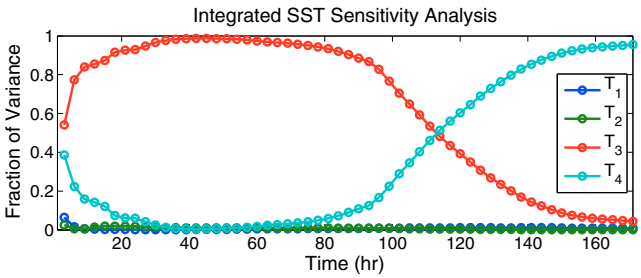

Fig. 2 Time evolution of sea surface temperature global sensitivity indices. The background diffusivity dominates until about hour 80, when Ivan enters the averaging box, whereas the wind drag coefficient becomes the dominant contributor at later times. T_i is the total sensitivity index of parameter θ_i [8]

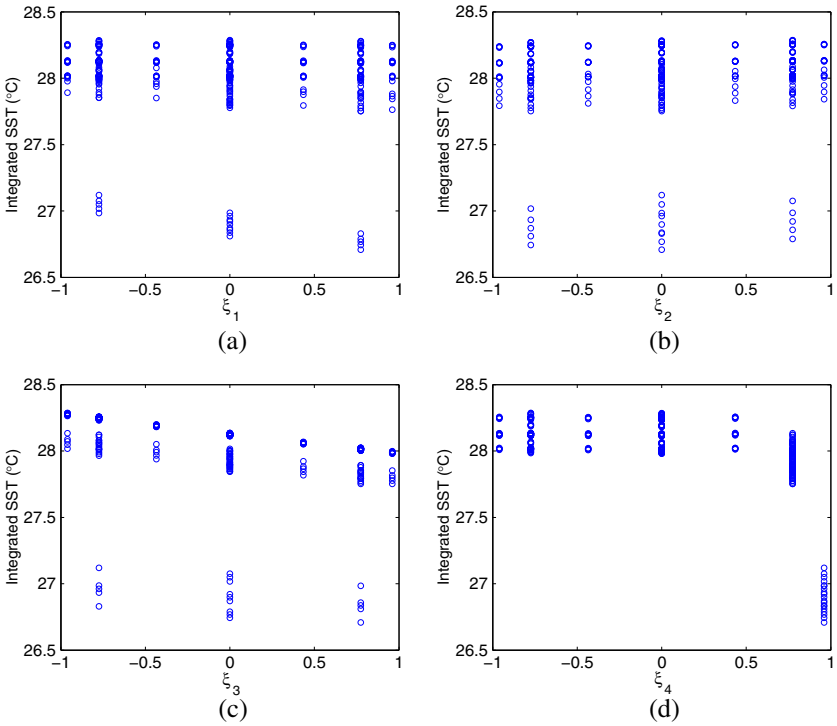

Fig. 3 Projections of the box-averaged SST at $t = 150$ h on canonical axes of the random parameter space: **a** ξ_1, **b** ξ_2, **c** ξ_3, and **d** ξ_4

Figure 4 shows the normalized coefficients, $U_k||\Psi_k||$, labeled by the corresponding multi-index; the zeroth coefficient is excluded. As expected from Fig. 3, the only significant normalized coefficient involving ξ_3 is the linear term. Apart from this term, the normalized coefficients having significant amplitude involve monomials of ξ_4. The slow decay of normalized coefficients involving ξ_4 also highlights the need for increased resolution along that axis.

In preparation for the analysis of adaptive expansions, additional computations were performed using the same model and stochastic setting considered in [8]. On one hand, we generated a separate database consisting of 256-member Latin hypercube sampling (LHS) ensemble. This was deemed desirable for the purpose of evaluating the quality of the PCE representations based on a set of realizations that is independent from that used for their construction.

On the other hand, motivated in large part by the results of the global sensitivity analysis, a preliminary assessment of adaptive refinement was conducted, namely by enriching the isotropic database generated in [8] preferentially along the third and fourth dimensions. Specifically, the database was enriched by including additional levels in the Smolyak quadrature, in such a way as to enable us to accurately capture polynomials of order ≤ 5 in the first two variables and ≤ 7 in the third and fourth variables.

Using the resulting anisotropic data, we analyzed the behavior of the solution for the different anisotropic truncations listed in Table 2. In particular, we used the relative L2 error between the corresponding PC representations and the LHS ensemble as a fidelity metric. Since the relative importance of individual components of $\boldsymbol{\xi}$ may be time-dependent, we plot in Fig. 5 the time evolution of this metric for different refinement levels. The results indicate that the PC representations which include higher order terms in the first two dimensions perform marginally better than those with a lower order truncation, whereas a substantial

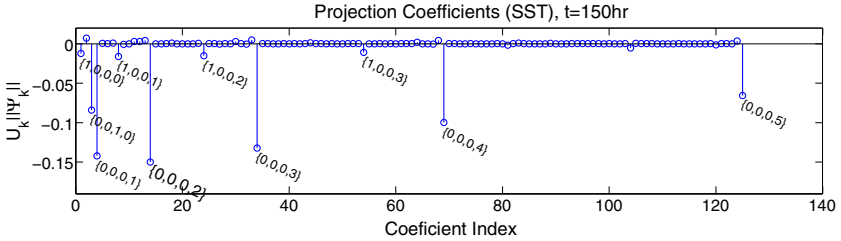

Fig. 4 Normalized coefficients of the box-averaged SST at $t = 150$ h. An isotropic fifth total order expansion is used. The coefficients with the largest amplitudes are labeled with the corresponding multi-index representing the polynomial order along each stochastic direction. The zeroth term is omitted for clarity

Table 2 Table of anisotropic series experiments

Simple truncation	P	Number of realizations
$\mathbf{p} = (5, 5, 5, 5)$	126	385
$\mathbf{p} = (5, 5, 7, 7)$	168	513
$\mathbf{p} = (2, 2, 5, 5)$	36	73
$\mathbf{p} = (2, 2, 7, 7)$	59	169

In the first column, $\mathbf{p}$ is a vector representing the polynomial order of accuracy along the individual dimensions, the second column represents the size of the resulting basis, and the third column indicates the number of realizations in the Smolyak sparse grid. The polynomial basis is described using multi-index set $\mathcal{F} = \left\{\mathbf{k} \in \mathbb{N}_0^d : \sum_{i=1}^d \frac{k_i}{p_i} \leq 1\right\}$

reduction in the error can be observed at later times when the PCE is enriched to seventh order in the third and fourth directions. This suggests that a coarse resolution in the first two directions can be safely adopted and that, using adaptive refinement, a suitable representation can be obtained at a substantially reduced cost compared to isotropic refinement.

In the following, we will investigate how the adaptive pseudo-spectral algorithm can be used to systematically refine the sampling of the stochastic parameter space and the corresponding PC representation and to quantify the performance gains thus achieved. The enriched Smolyak set with $\mathbf{p} = (5, 5, 7, 7)$ will be specifically used as a database for the a priori tests.

4 General adaptive PCE

A disadvantage of the dimensional refinement scheme above is that it is relatively inefficient, potentially leading to the inclusion of a large number of realizations with every refinement step. Another disadvantage concerns the associated truncation scheme which, in order to avoid internal aliasing (Appendix A), severely limits the number of polynomials included in the truncated basis. An attractive avenue to overcome these drawbacks is provided by the adaptive pseudo-spectral algorithm recently proposed in [15]. In this section, we will first outline its construction and then evaluate its performance based on a priori tests that exploit the HYCOM database described in the previous section.

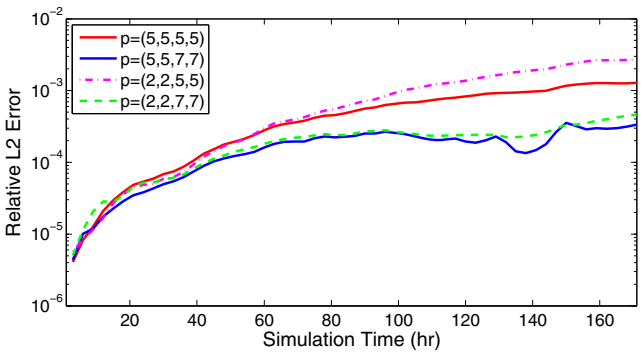

Fig. 5 Relative L2 error between the area-averaged SST and the LHS sample. Curves are generated for different anisotropic refinement levels, as indicated

4.1 Smolyak pseudo-spectral projection

As described in [15], the basis of the pseudo-spectral projection is to apply Smolyak algorithm directly to construct the PCE instead of purely generating the quadrature. Thus, the final projection becomes a weighted sum of aliasing-free subprojections. This represents an extension of the Smolyak tensor construction from quadrature operators to projection operators.

Specifically, the PC representation resulting from the pseudo-spectral algorithm is expressed as a sum of isotropic and anisotropic full-tensor subprojections [15]. Full tensorization leads to a "rectangular" multi-index, as illustrated in Fig. 6a for a 2D example. The dotted rectangle in Fig. 6b depicts the set, $\mathcal{O}$, of polynomials whose coefficients can be determined without internal aliasing. The PC basis in the pseudo-spectral construction contains the orthogonal polynomials whose orders belong to the union of all $\mathcal{O}_k$'s [15].

Note that each subprojection in the pseudo-spectral construction is a full PC representation of U that is free of internal aliasing. As schematically illustrated in Fig. 7, the subprojections are then combined based on the Smolyak construction. A key feature of the pseudo-spectral construction is that it affords a general refinement approach based on successive inclusion of any admissible multi-index, $\mathcal{F}$, of quadrature rules [19] while maintaining the representation free of internal aliasing. Another beneficial feature concerns the ability to retain, using the same realization stencil, a larger number of polynomials than is possible with a classical dimensional truncation/quadrature approach. This, in particular, includes high-order monomial terms that are not accurately estimated using the latter. As a specific example, for the HYCOM database with 513 realizations, using the pseudo-spectral algorithm, we are able to determine an orthogonal basis containing 402 polynomials, whereas only a 168-polynomial basis is obtained using the direct quadrature approach.

4.2 General adaptive driver

The construction of the adaptive driver in [15] follows the general framework introduced by Gerstner and Griebel [19] for adaptive quadrature. We begin with the first-order quadrature in all dimensions, which permits only a zeroth-order expansion. We then add all the forward neighbors and perform the projection. The pseudo-spectral construction of

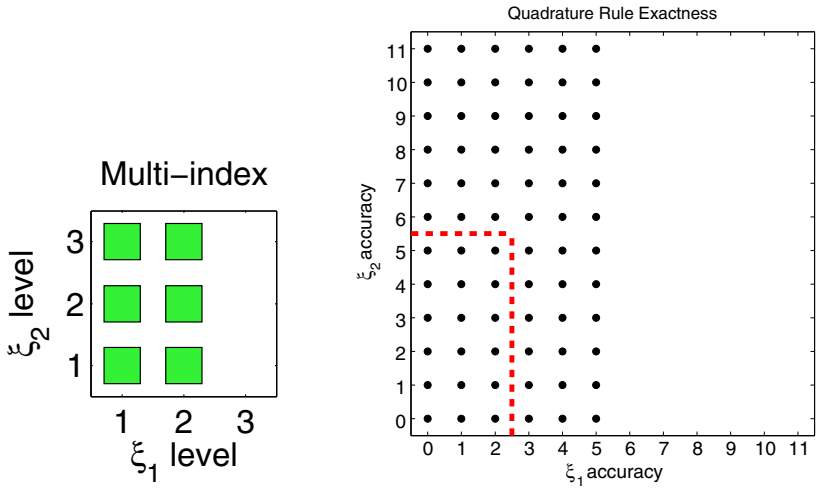

(a) Multi-Index (b) Quadrature Exactness and ResolvableTerms

Fig. 6 **a** Schematic illustration a fully tensorized projection in two dimensions. **b** Schematic of the polynomial exactness of the quadrature rule. The polynomials whose order fall within the *dotted rectangle* are resolvable in the resulting projection without internal aliasing

a function U is built on the tensor product of projections *differences* of the form

$$\left(\Delta_{k_1} \otimes \ldots \otimes \Delta_{k_d}\right) U,$$

which can be interpreted as a PCE of projection differences. The L_2 norm of this difference can be readily used to define an error indicator for multi-index $\boldsymbol{k}$:

$$\epsilon(\boldsymbol{k}) = || \left(\Delta_{k_1} \otimes \ldots \otimes \Delta_{k_d}\right) U ||. \quad (5)$$

The indicator represents the surplus of variance due to the $\boldsymbol{k}$ subprojection. The surplus is computed for each $\boldsymbol{k} \in \mathcal{F}$, and the subprojection with the highest $\epsilon(\boldsymbol{k})$ is selected for subsequent refinement, which generally consists of inclusion of admissible forward neighbors. In the present a priori analysis, the selection of a multi-index $\boldsymbol{k}$ for refinement is subject to the requirement that it has at least one forward neighbor that meets all of the following criteria:

1. The forward neighbor is not already in $\mathcal{F}$,

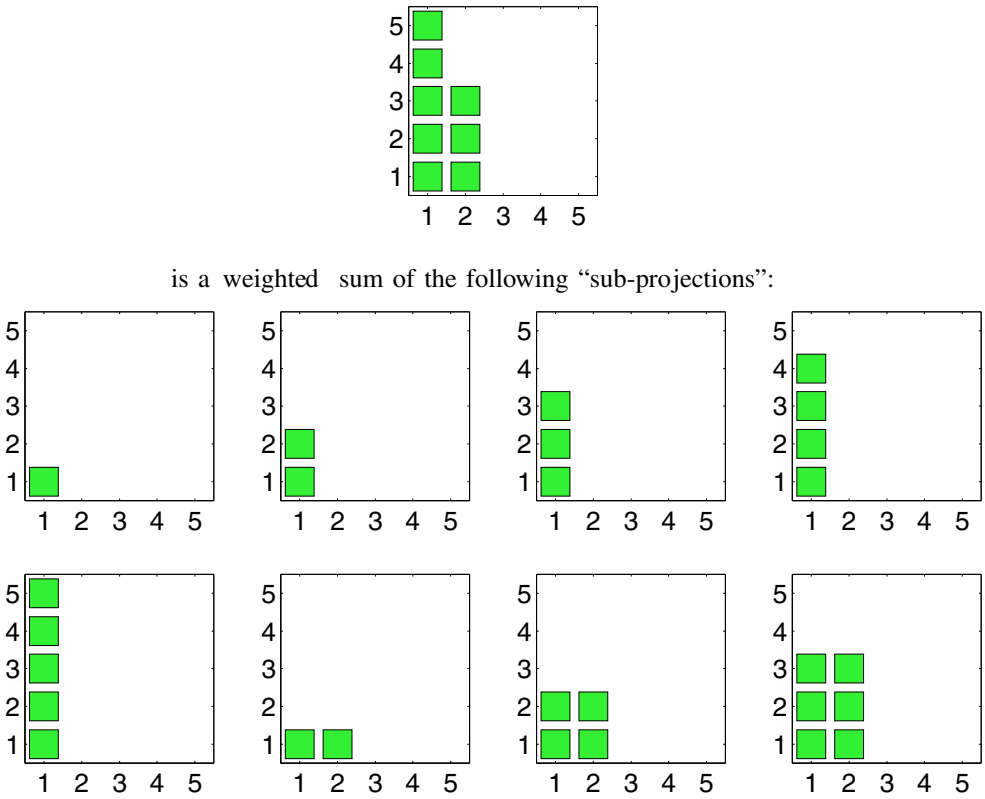

is a weighted sum of the following "sub-projections":

Fig. 7 Example showing how the final projection is a weighted sum of full-tensor (rectangular) subprojections

2. $\mathcal{F}$ remains admissible with the inclusion of the forward neighbor, and
3. The additional realizations associated with the forward neighbor are in the database.

If no forward neighbor meets the criteria above, then the projection with the next highest indicator, ϵ, is chosen. The process repeats until a valid forward neighbor cannot be added or the database is exhausted. All of the forward neighbors of the critical multi-index which meet the criteria above are added to $\mathcal{F}$ in a single refinement step, and a new PCE is calculated using the Smolyak pseudo-spectral projection. See Fig. 8 for an illustration.

Note that when the refinement algorithm is used dynamically, i.e., not in conjunction with a preexisting database, requirement 3 is naturally omitted. Also, note that in all cases, the first two iterations will always be identical. The first iteration consists of a zeroth-order projection, whereas the second iteration includes all forward neighbors in all directions.

5 Results

The adaptive algorithm described in Section 4.2 is first tested using the area-averaged SST as the QoI, and the algorithm's performance is analyzed by computing the relative error between its PCE and the LHS ensemble. The study is performed for $t = 60$ h and $t = 120$ h, i.e., at times when ξ_3 is dominant and when both ξ_3 and ξ_4 contribute substantially to the total variance (see Fig. 2). From Fig. 3, we expect SST to exhibit a steep dependence on ξ_4 at later times, whereas the dependence on ξ_3 is expected to remain essentially linear. Note that at $t = 120$ h, the hurricane is located within the analysis region.

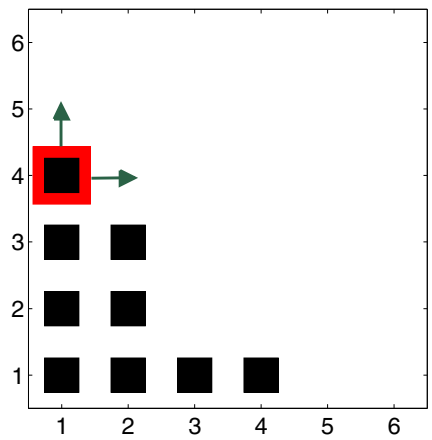

Fig. 8 Schematic illustration of a refinement step showing the addition of two forward neighbors of a critical multi-index

As previously discussed, the adaptive pseudo-spectral algorithm is anticipated to lead to performance gains due to two features: (1) judicious selection of stencils so as to maximize the ability to capture the variance of the selected QoI and (2) maximizing the PC basis, namely by including all polynomials that can be evaluated while avoiding internal aliasing. Figure 9 illustrates the impact of these features by comparing the errors in the area-averaged SST for different basis refinement strategies:

S1. The adaptive pseudo-spectral algorithm,
S2. The pseudo-spectral PCE construction algorithm applied to an isotropically refined sparse grid, and
S3. The isotropic Smolyak quadrature with direct projection.

For S2 and S3, the final point includes the entire database. Thus, the key difference between the second and third constructions concerns the inclusion of a larger basis in the former. In all cases, both the quadrature accuracy *and* the PCE stencil are refined.

Figure 9 shows the error of the adaptive pseudo-spectral algorithm decreasing rapidly and essentially saturating after the fifth iteration. At this stage, the adaptive pseudo-spectral PCE (S1) uses 69 realizations and the corresponding basis has 59 polynomials. In contrast, Fig. 9 indicates that the decay of the relative error is substantially slower with isotropic refinement (S3), though noticeable improvement over Smolyak quadrature is observed when the basis is constructed pseudo-spectrally (S2).

The curves also indicate that the error levels experienced with isotropically refined grids remain substantially larger than with the adaptive pseudo-spectral algorithm, at least until the entire database is used. Combined, the results indicate that the adaptive pseudo-spectral algorithm can lead to substantial gains that arise due to judicious refinement of the realization stencil and of the PC basis. An illustration of the savings afforded by the adaptive pseudo-spectral construction is provided in the right panel of Fig. 9, which contrasts probability density functions (PDFs) of the SST obtained based on the PCE generated with the adaptive scheme at iteration 5 (69 realizations), with results obtained similarly but using the full database. The close agreement between all three curves highlights the computational savings obtained using the adaptive pseudo-spectral algorithm as well as the validity of the corresponding representation.

It is interesting to note that with isotropic sampling, a large number of realizations are added as the refinement level increases, whereas in adaptive sampling, the number of realizations increases at a much slower rate with the iteration count. It is also interesting to note the plateau in the relative RMS is small, about 1 %. This is consistent with the earlier analysis in [8], which indicated that a level 5

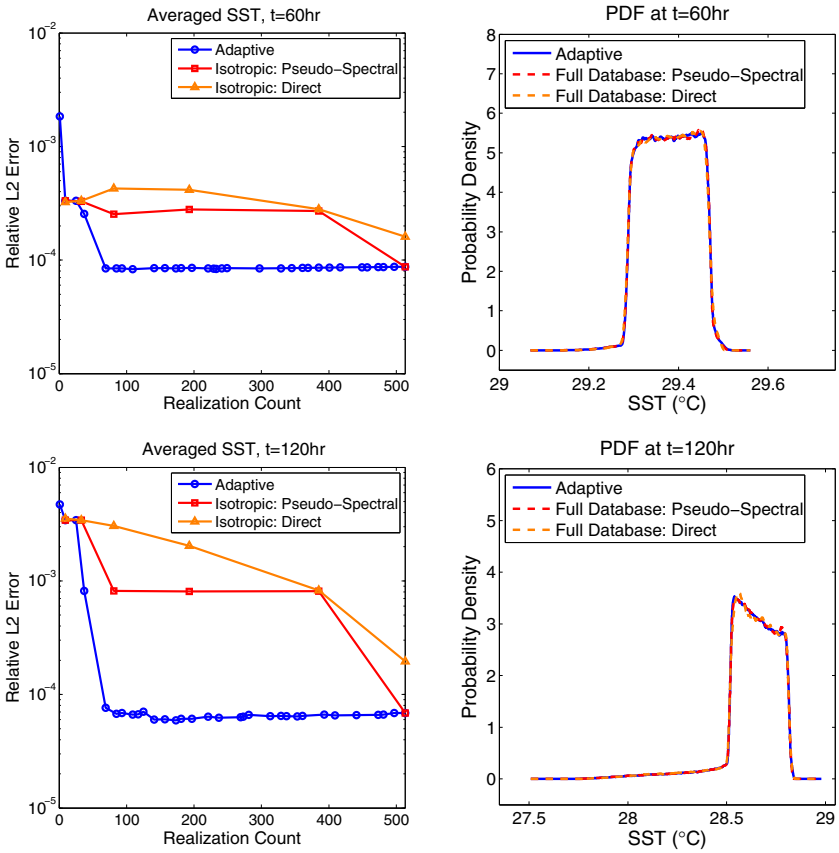

Fig. 9 *Left*: LHS error versus realization count for the three refinement strategies: (S1) the adaptive pseudo-spectral algorithm; (S2) the pseudo-spectral construction applied in combination with isotropic refinement, except for the last point where the full database is used; and (S3) Smolyak sparse quadrature, except for the last iteration for which dimension truncation is used in conjunction with the full database. *Right*: PDFs of SST based on (S1) the PCE corresponding to the adaptive pseudo-spectral algorithm at iteration 5 (69 realizations and 59 polynomials), (S2) the pseudo-spectral PCE using the full database, and (S3) the PCE constructed using Smolyak quadrature using the full database. In all cases, the PDFs are generated using 10^6 samples drawn from the corresponding PCEs. *Top row*: $t = 60$ h; *bottom row*: $t = 120$ h

isotropic refinement leads to a suitable representation of the solution response.

In order to investigate the origin of the plateau in the RMS error in Fig. 9, we have examined the evolution of the realization stencil in the adaptive refinement scheme. This exercise revealed that the error plateau arises when the refinement reaches the limits of the database, i.e., when forward neighbors of the critical multi-index are not available. An illustration of the evolution of the adaptive refinement is provided in Fig. 10, which depicts projections of the stencil on the ξ_3–ξ_4 plane. Recall that the first two iterations are always identical and reflect an isotropic refinement step in all directions. Consistent with the results of the sensitivity analysis [8] and the projections shown in Fig. 3, the refinement first proceeds along ξ_3. Because a linear approximation in ξ_3 effectively captures the corresponding dependence, refinement then proceeds along ξ_4. The steep variation of SST at the higher values of ξ_4 (see Fig. 3) requires a high-order polynomial representation; refinement proceeds along ξ_4 until the limits of the database are reached (iteration 5). This then forces the algorithm to sequentially select the remaining stencils afforded by the database, including the mixed terms that appear at the higher iteration counts.

As noted, the refinement became limited by the realization database at iteration 5 for both $t = 60$ h and $t = 120$ h. Without such limitation, the refinement would have more monomial terms in ξ_4. Once the database was exhausted, the relative L2 error plateaued in Fig. 9 and all additional terms had little impact.

Another interesting observation in Figs. 9 and 10 is that the adaptive refinement proceeds in a similar fashion for the area-averaged SST at $t = 60$ h and $t = 120$ h. This

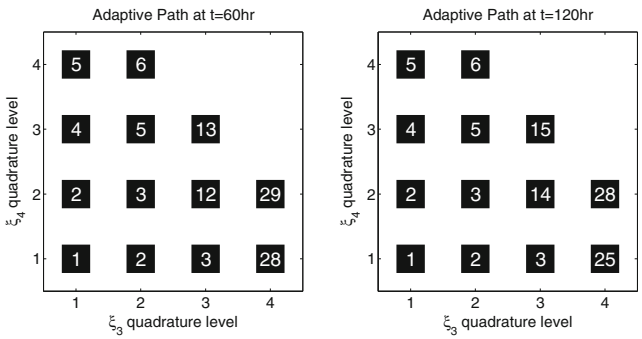

Fig. 10 Schematic illustration of the evolution of the adaptive refinement algorithm. Plotted are projections of the stencils on the ξ_3–ξ_4 plane. The *boxes* indicate the refinement levels, whereas the *numbers within each box* refer to the iteration number at which the box was added to the stencil. Recall that iterations 1 and 2 are always isotropic. Note that refinement occurs first along ξ_3 but then progresses along ξ_4, where the limits of the database are reached. *Left*: $t = 60$ h; *right*: $t = 120$ h

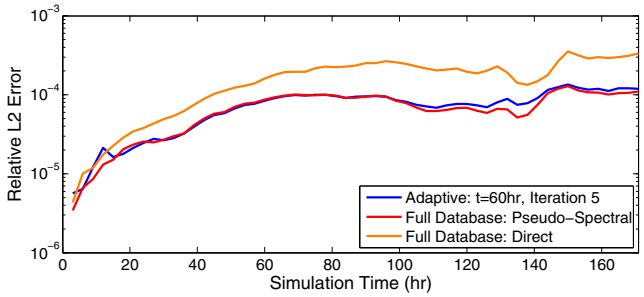

Fig. 11 Relative L2 error between the PCE of the averaged SST and the LHS sample. Plotted are curves generated with (1) the pseudo-spectral algorithm adapted to the solution at $t = 60$ h, (2) the pseudo-spectral algorithm using the full database, and (3) the Smolyak quadrature using the full database. For the adapted solution, the refinement is stopped after iteration 5, leading to 69 realizations and a PCE with 59 polynomials. The full 513 database curves have 402 polynomials for the pseudo-spectral construction and 168 polynomials for the Smolyak quadrature

can be explained by noting that the dependence on ξ_3, the dominant variable at early times, is to a large extent captured by a linear term throughout the simulation. As the hurricane enters into the analysis region, the drag coefficient becomes the dominant contributor to the variance. As depicted earlier, the dependence on ξ_4 is highly nonlinear and evolves with time. However, it is rapidly captured with the next few refinement steps. Though time-dependent, subsequent iterations that lead to the inclusion of mixed terms take place well after the error has decreased dramatically. This should be viewed as a particular feature of the present application and may not be necessarily the case for others.

To further verify the fidelity of the adaptive PCE with time, we plot in Fig. 11 the evolution of the relative LHS error based on the adaptive pseudo-spectral solution truncated at iteration 5 (69 realizations, 59 polynomials), the pseudo-spectral expansion using the full database (513 realizations, 402 polynomials), and the Smolyak quadrature using the full database (168 polynomials). The adaptive solution is based on the box-averaged SST at $t = 60$ h. The results show close agreement between the adapted solution and the pseudo-spectral approximation based on the full database, throughout the computations. Also, note that at larger times, the adaptive solution exhibits error levels that are substantially smaller than those obtained using the Smolyak quadrature, though the latter involves a much larger number of realizations.

It is also instructive to examine whether a PCE resulting from adaptive refinement based on a specific QoI would also prove suitable for the purpose of representing other QoIs and whether adaptive refinement driven by multiple QoIs simultaneously would potentially require a large number of realizations and consequently limit performance gains. Though these issues in general vary from one application to another, it is still worthwhile to address them in the present setting of the oceanic response to hurricane-strength winds.

Figure 12 shows the relative L2 error based on the LHS ensemble for heat flux, Q_{tot} (Section 3), and the box-averaged SST. Plotted are the results obtained using PCEs that are adapted using either (1) the box-averaged SST or (2) Q_{tot} as the sole refinement QoI at $t = 60$ h. The curves show that the decay of the error is not identical in both cases, indicating that the refinement does not follow the same pattern and that optimal sampling and PCE may depend on the selected QoI. This is not surprising since the SST is averaged over a fixed region, whereas Q_{tot} involves an average in a circular area around the center of the hurricane. Thus, Q_{tot} is always affected by the high winds whereas the box-averaged SST experiences both weak winds and high winds. Note, however, that in both cases, a plateau in the error curves is achieved at about the same number of realizations and much earlier than with isotropic sampling. Thus, in the present examples, the performance gains do not appear to be substantially affected by the QoI selected.

An alternative to the single QoI adaptation is to adapt concurrently based on multiple QoIs, wherein the

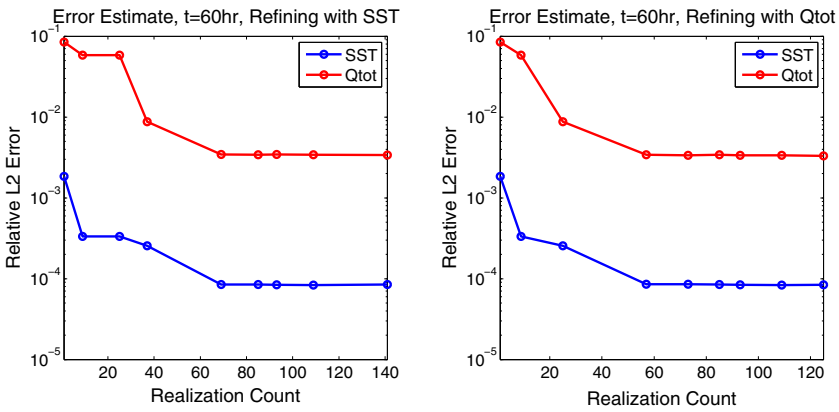

Fig. 12 Relative L2 error between LHS sample and the PCEs for SST and Q_{tot}. The adaptive refinement is based on the critical multi-index for **a** the box-averaged SST and **b** Q_{tot}

refinement considers the union of the forward neighbors corresponding to the critical multi-indices of all QoIs. This strategy potentially involves a larger number of realizations than with single-QoI adaptation but naturally overcomes the drawback identified above. In Fig. 13, we plot curves of the relative L2 LHS error against the realization count for all four QoIs considered, namely SST, SSH, MLD, and Q_{tot}. The results indicate that with the present refinement strategy, all error estimates decay rapidly. Note that as expected, in the present case, the refinement steps generally include a larger number of realizations than in single-QoI adaptation. The error curves tend to become essentially flat when the realization count exceeds 70. Thus, for the present setting, similar performance gains are achieved when considering SST alone or in combination with SSH, MLD, and Q_{tot}. Of course, in other settings, some QoIs may converge slower than others, and it may be advantageous to consider more elaborate algorithms, for instance stopping refinement based on a specific QoI when the corresponding PCE is deemed sufficiently resolved.

We conclude with brief remarks concerning the scope of the present implementations and potential extensions. In the a priori analyses above, the adaptive algorithm was allowed to proceed until the entire database of precomputed realizations was exhausted. In dynamic implementation, this is of course not possible, and suitable stopping criteria must be provided. In this regard, we point out that the adaptive sampling scheme supports a variety of approaches, based for instance on monitoring the saturation of the variance in selected QoIs, providing global tolerance levels for the error indicators, or detecting occurrence of plateaus in the interpolation error. The implementation of such criteria would naturally exploit the variance-based error indicators associated with active multi-indices and may also involve the corresponding number of forward neighbors. As discussed in [19], the latter may also be used for the purpose of introducing suitable cost functions and accordingly optimize the enrichment of the realization stencil.

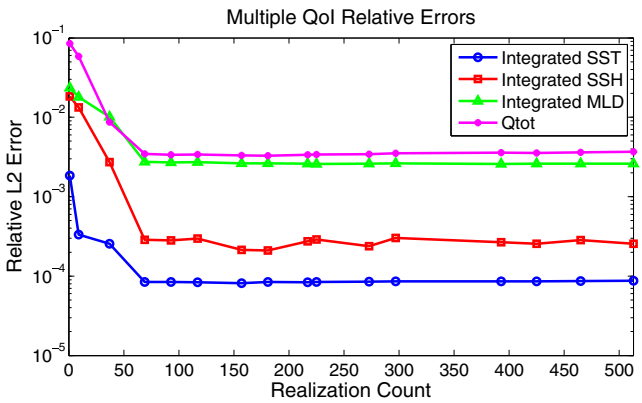

Fig. 13 Relative error versus realization count. Plotted are curves SST, SSH, MLD, and Q_{tot}. The stencil is enriched by the union of all admissible forward neighbors determined for each QoI individually

6 Conclusion

This work has exploited an existing database of HYCOM simulations of the oceanic response in the GOM to Hurricane Ivan to perform an a priori analysis and testing of adaptive methods. The database was adapted from prior work [8], where a sparse isotropic sampling scheme was used to propagate uncertainties in the wind drag coefficient and subgrid mixing parameterizations. Guided by the sensitivity analysis in [8], the database was enriched along the dominant directions. This resulted in an extended database comprising 513 independent realizations, with stochastic coordinates located at the quadrature points of an anisotropic sparse grid. For the purpose of characterizing representation errors, we also generated an independent Latin hypercube sample consisting of 256 realizations.

The adaptive pseudo-spectral algorithm recently developed in [15] was implemented and deployed on the enriched database. It accommodates anisotropic refinement on arbitrary admissible sparse grids and enables the evaluation of a "maximal" basis set in a manner that avoids internal aliasing. Inherent in the adaptive pseudo-spectral construction is an error refinement indicator which can be readily used to drive anisotropic refinement of the sparse grid. This involves inclusion of admissible forward neighbors of the critical multi-index or critical multi-index set.

The performance of the adaptive pseudo-spectral algorithm was also tested a priori using the HYCOM database. The tests included adaptive refinement based on a single QoI and on multiple QoIs. In particular, the tests indicated that suitable PCEs of the uncertain model response can be constructed at a small fraction of the cost of isotropic sampling and that the performance gains arise from the algorithm's ability to refine efficiently, as well as from the inclusion of all polynomials whose coefficients can be determined without internal aliasing.

The present a priori tests thus provide support to the deployment of the adaptive pseudo-spectral algorithm in a dynamic fashion. The adaptive pseudo-spectral algorithm has already been adopted in a study of the oceanic response to Typhoon Fanapi where it is used to build the surrogate needed for uncertainty propagation and parameter inference; results from this study will be reported elsewhere [9].

Acknowledgments This research was supported by the Office of Naval Research, award N00014-101-0498; by the US Department of Energy, Office of Advanced Scientific Computing Research, award numbers DE-SC0007020, DE-SC0008789, and DE-SC0007099; and by the Gulf of Mexico Research Initiative, contract numbers SA1207GOMRI005 (CARTHE) and SA12GOMRI008 (DEEP-C).

Appendix A: Internal aliasing

As discussed in [15], internal aliasing refers to the errors incurred in the numerical evaluation of inner products between elements of the truncated basis. Let $\mathcal{P}_m f$ and $\mathcal{S}_m f$ refer to the exact and approximate projections of a function $f \in L_2$ on the truncated *orthonormal* basis $\{\psi_i\}_{i=0}^m$; we have

$$\mathcal{P}_m f = \sum_{i=0}^{m} \hat{f}_i \psi_i(\xi) \text{ with } \hat{f}_i = \langle f, \psi_i \rangle \tag{6}$$

$$\mathcal{S}_m f = \sum_{i=0}^{m} \tilde{f}_i \psi_i(\xi) \text{ with } \tilde{f}_i = \langle f, \psi_i \rangle_Q = \langle f, \psi_i \rangle + \epsilon_i \tag{7}$$

where $\langle \cdot, \cdot \rangle$ refers to the continuous inner product (exact integration), $\langle \cdot, \cdot \rangle_Q$ refers to its approximation via quadrature, and ϵ_i is the quadrature error. An expression for the error can be obtained by inserting the untruncated series representation $\mathcal{P}_\infty f$ into (7), which results in

$$\tilde{f}_i = \sum_{j=0}^{m} \hat{f}_j \langle \psi_j, \psi_i \rangle_Q + \sum_{j=m+1}^{\infty} \hat{f}_j \langle \psi_j, \psi_i \rangle_Q. \tag{8}$$

The first term in the above equation has been dubbed the internal aliasing error whereas, the second one is the (usual) external aliasing. If the numerical quadrature preserves discrete orthogonality for elements of the truncated basis, i.e., $\langle \psi_j, \psi_i \rangle_Q = \delta_{i,j}$ for $(i, j) \leq m$, we get the familiar result that the difference between the pseudo-spectral coefficients $\tilde{f}_i$ and the spectral coefficients $\hat{f}_i$ is solely due to external aliasing [20].

Note that when the selected quadrature does not preserve discrete orthogonality in the truncated basis, large errors can generally be anticipated, which frequently leads to poor representations. Consequently, it is generally essential to avoid internal aliasing, and various means can be adopted for this purpose. With direct quadrature rules, however, this may lead to severe restriction on the truncated basis used in the expansion. An illustration of this phenomenon is shown in Fig. 14, which schematically depicts a 2D example. In the figure, we have used green boxes to depict elements of a truncated basis and black boxes to depict polynomial orders for which the quadrature would be exact. In this example, the selected quadrature would not be suitable, since inner products involving polynomials of order 4 or larger may not be exactly computed. For instance, the inner product $\langle \psi_5(\xi_1)\psi_5(\xi_2) \rangle$ would not vanish. To avoid internal aliasing using the quadrature depicted in the figure, the truncated basis would need to be further reduced by dropping all polynomials of order 4 or higher. As discussed in Section 4.1, in multidimensional settings, the adaptive pseudo-spectral construction provides an attractive alternative to direct projection, as it generally enables us to retain a larger number of additional polynomials than with direct quadrature while avoiding internal aliasing errors.

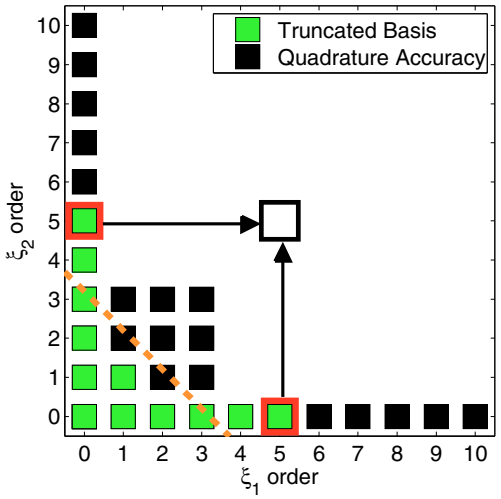

Fig. 14 Schematic demonstration of internal aliasing in 2D. The *green boxes* represent a truncated polynomial basis and the *black boxes* delineate the boundary of the polynomial exactness of direct quadrature. Depicted is an example showing how the inner product of the selected polynomials (*red squares*) falls outside of the region of accuracy of the underlying quadrature. Only polynomials with orders falling below the *orange dotted line* can be computed without internal aliasing

References

1. Le Maitre, O.P., Knio, O.: Spectral Methods for Uncertainty Quantification: with Applications to Computational Fluid Dynamics. Springer Series in Scientific Computation. Springer, Dordrecht (2010)
2. Balakrishnan, S., Roy, A., Ierapetritou, M.G., Flach, G.P., Georgopoulos, P.G.: Uncertainty reduction and characterization for complex environmental fate and transport models: an empirical bayesian framework incorporating the stochastic response surface method. Water Resour. Res. **39**(12), 1350 (2003)
3. Isukapalli, S.S., Balakrishnan, S., Georgopoulos, P.G.: Computationally efficient uncertainty propagation and reduction using the stochastic response surface method. In: 43rd IEEE Conference on Decision and Control, 2004, CDC, vol. 2, pp. 2237–2243. IEEE (2004)
4. Marzouk, Y.M., Najm, H.N.: Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. J. Comput. Phys. **228**(6), 1862–1902 (2009)
5. Crestaux, T. et al.: Polynomial chaos expansion for sensitivity analysis. Reliab. Eng. Syst. Saf. **94**(7), 1161–1172 (2009)
6. Sudret, B.: Global sensitivity analysis using polynomial chaos expansions. Reliab. Eng. Syst. Saf. **93**(7), 964–979 (2008)
7. Thacker, W.C., Srinivasan, A., Iskandarani, M., Knio, O.M., Le Henaff, M.: Propagating boundary uncertainties using polynomial expansions. Ocean Model. **43–44**, 52–63 (2012)
8. Alexanderian, A., Winokur, J., Sraj, I., Srinivasan, A., Iskandarani, M., Thacker, W.C., Knio, O.M.: Global sensitivity analysis in an ocean general circulation model: a sparse spectral projection approach. Comput. Geosci. **16**(3), 1–22 (2012)

9. Sraj, I., Iskandarani, M., Srinivasan, A., Thacker, W.C., Winokur, J., Alexanderian, A., Lee, C.-Y., Chen, S.S., Knio, O.M.: Bayesian inference of drag parameters using Fanapi AXBT data. Mon. Wea. Rev. (2013). doi:10.1175/MWR-D-12-00228.1
10. Ghanem, R.G., Spanos, P.D.: Stochastic Finite Elements: a Spectral Approach. Dover, Mineola (2003)
11. Gerstner, T., Griebel, M.: Numerical integration using sparse grids. Numer. Algorithms **18**(3–4), 209–232 (1998)
12. Petras, K.: On the smolyak cubature error for analytic functions. Adv. Comput. Math. **12**, 71–93 (2000)
13. Smolyak, S.A.: Quadrature and interpolation formulas for tensor products of certain classes of functions. Dokl. Akad. Nauk SSSR **4**(240–243), 123 (1963)
14. Constantine, P.G., Eldred, M.S., Phipps, E.T.: Sparse pseudospectral approximation method. Comput. Methods Appl. Mech. Eng. **229–232**, 1–12 (2012)
15. Conrad, P.R., Marzouk, Y.M.: Adaptive Smolyak pseudospectral approximations. Submitted: CoRR, abs/1209.1406 (2012)
16. Xiu, D.B., Karniadakis, G.E.: The Wiener-Askey polynomial chaos for stochastic differential equations. SIAM J. Sci. Comput. **24**, 619–644 (2002)
17. Marzouk, Y.M., Najm, H.N., Rahn, L.A.: Stochastic spectral methods for efficient Bayesian solution of inverse problems. J. Comput. Phys. **224**, 560–586 (2007)
18. Patterson, T.N.L.: The optimum addition of points to quadrature formulae. Math. Comput. **22**, 847–856 (1968)
19. Gerstner, T., Griebel, M.: Dimension-adaptive tensor-product quadrature. Computing **71**(1), 65–87 (2003)
20. Canuto, C., Hussaini, M.Y., Quarteroni, A., Zang, T.A.: Spectral Methods in Fluid Dynamics. Springer Series in Computational Physics. Springer, New York (1988)