MIT Open Access Articles

Adaptive Smolyak Pseudospectral Approximations

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Conrad, Patrick R., and Youssef M. Marzouk. "Adaptive Smolyak Pseudospectral Approximations." SIAM Journal on Scientific Computing 35, no. 6 (January 2013): A2643–A2670. © 2013, Society for Industrial and Applied Mathematics

As Published: <http://dx.doi.org/10.1137/120890715>

Publisher: Society for Industrial and Applied Mathematics

Persistent URL: <http://hdl.handle.net/1721.1/86020>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Terms of use text: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

ADAPTIVE SMOLYAK PSEUDOSPECTRAL APPROXIMATIONS*

PATRICK R. CONRAD[†] AND YOUSSEF M. MARZOUK[‡]

Abstract. Polynomial approximations of computationally intensive models are central to uncertainty quantification. This paper describes an adaptive method for nonintrusive pseudospectral approximation, based on Smolyak’s algorithm with generalized sparse grids. We rigorously analyze and extend the nonadaptive method proposed in [P. G. Constantine, M. S. Eldred, and E. T. Phipps, *Comput. Methods Appl. Mech. Engrg.*, 229–232 (2012), pp. 1–12], and compare it to a common alternative approach for using sparse grids to construct polynomial approximations, direct quadrature. Analysis of direct quadrature shows that $\mathcal{O}(1)$ errors are an intrinsic property of some configurations of the method, as a consequence of internal aliasing. We provide precise conditions, based on the chosen polynomial basis and quadrature rules, under which this aliasing error occurs. We then establish theoretical results on the accuracy of Smolyak pseudospectral approximation, and show that the Smolyak approximation avoids internal aliasing and makes far more effective use of sparse function evaluations. These results are applicable to broad choices of quadrature rule and generalized sparse grids. Exploiting this flexibility, we introduce a greedy heuristic for adaptive refinement of the pseudospectral approximation. We numerically demonstrate convergence of the algorithm on the Genz test functions, and illustrate the accuracy and efficiency of the adaptive approach on a realistic chemical kinetics problem.

Key words. Smolyak algorithms, sparse grids, orthogonal polynomials, pseudospectral approximation, approximation theory, uncertainty quantification

AMS subject classifications. 41A10, 41A63, 47A80, 65D15, 65D32

DOI. 10.1137/120890715

1. Introduction. A central issue in the field of uncertainty quantification is understanding the response of a model to random inputs. When model evaluations are computationally intensive, techniques for *approximating* the model response in an efficient manner are essential. Approximations may be used to evaluate moments or the probability distribution of a model’s outputs, or to evaluate sensitivities of model outputs with respect to the inputs [21, 43, 33]. Approximations may also be viewed as *surrogate models* to be used in optimization [22] or inference [23], replacing the full model entirely.

Often one is faced with black box models that can only be evaluated at designated input points. We will focus on constructing multivariate polynomial approximations of the input-output relationship generated by such a model; these approximations offer fast convergence for smooth functions and are widely used. One common strategy for constructing a polynomial approximation is interpolation, where interpolants are conveniently represented in Lagrange form [1, 39]. Another strategy is projection, particularly orthogonal projection with respect to some inner product. The results of such a projection are conveniently represented by the corresponding family of orthog-

*Submitted to the journal’s Methods and Algorithms for Scientific Computing section September 10, 2012; accepted for publication (in revised form) September 11, 2013; published electronically November 19, 2013. This work was supported by the Scientific Discovery through Advanced Computing (SciDAC) program funded by the US Department of Energy, Office of Science, Advanced Scientific Computing Research under award DE-SC0007099.

<http://www.siam.org/journals/sisc/35-6/89071.html>

[†]Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 (prconrad@mit.edu). This author’s work was supported by a Department of Defense NDSEG Fellowship and an NSF graduate fellowship.

[‡]Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 (ymarz@mit.edu).

onal polynomials [4, 21, 40]. When the inner product is chosen according to the input probability measure, this construction is known as the (finite-dimensional) polynomial chaos expansion (PCE) [14, 32, 8]. Interpolation and projection are closely linked, particularly when projection is computed via discrete model evaluations. Moreover, one can always realize a change of basis [9] for the polynomial resulting from either operation. Here we will favor orthogonal polynomial representations, as they are easy to manipulate and their coefficients have a useful interpretation in probabilistic settings.

This paper discusses *adaptive Smolyak pseudospectral approximation*, an accurate and computationally efficient approach to constructing multivariate polynomial chaos expansions. Pseudospectral methods allow the construction of polynomial approximations from point evaluations of a function [4, 3]. We combine these methods with *Smolyak's algorithm*, a general strategy for sparse approximation of linear operators on tensor product spaces, which saves computational effort by weakening the assumed coupling between the input dimensions. Gerstner and Griebel [13] and Hegland [16] developed adaptive variants of Smolyak's algorithm for numerical integration and illustrated the effectiveness of on-the-fly heuristic adaptation. We extend their approach to the pseudospectral approximation of functions. Adaptivity is expected to yield substantial efficiency gains in high dimensions—particularly for functions with anisotropic dependence on input parameters and functions whose inputs might not be strongly coupled at high order.

Previous attempts to extend pseudospectral methods to multivariate polynomial approximation with sparse model evaluations employed ad hoc approaches that are not always accurate. A common procedure has been to use sparse quadrature, or even dimension-adaptive sparse quadrature, to evaluate polynomial coefficients directly [41, 21]. This leads to at least two difficulties. First, the truncation of the polynomial expansion must be specified independently of the quadrature grid, yet it is unclear how to do this, particularly for anisotropic and generalized sparse grids. Second, unless one uses excessively high-order quadrature, significant aliasing errors may result. Constantine, Eldred, and Phipps [6] provided the first clear demonstration of these aliasing errors and proposed a Smolyak algorithm that does not share them. That work also demonstrated a link between Smolyak pseudospectral approximation and an extension to Lagrange interpolation called *sparse interpolation*, which uses function evaluations on a sparse grid and has well characterized convergence properties [24, 2].

The first half of this work performs a theoretical analysis, placing the solution from [6] in the broader context of Smolyak constructions, and explaining the origin of the observed aliasing errors for general (e.g., anisotropic) choices of sparse grid and quadrature rule. We do so by using the notion of polynomial exactness, without appealing to interpolation properties of particular quadrature rules. We establish conditions under which tensorized approximation operators are exact for particular polynomial inputs, then apply this analysis to the specific cases of quadrature and pseudospectral approximation; these cases are closely related and facilitate comparisons between Smolyak pseudospectral algorithms and direct quadrature. Section 2 develops *computable* one-dimensional and tensorized approximations for these settings. Section 3 describes general Smolyak algorithms and their properties, yielding our principal theorem about the polynomial exactness of Smolyak approximations, and then applies these results to quadrature and pseudospectral approximation. Section 4 compares the Smolyak approach to conventional direct quadrature. Our error analysis of direct quadrature shows why the approach goes wrong and allows us to draw an important conclusion: in almost all cases, direct quadrature is not an appro-

priate method for constructing polynomial expansions and should be superseded by Smolyak pseudospectral methods.

These results provide a rigorous foundation for *adaptivity*, which is the second focus of this paper. Adaptivity makes it possible to harness the full flexibility of Smolyak algorithms in practical settings. Section 5 introduces a fully adaptive algorithm for Smolyak pseudospectral approximation, which uses a single tolerance parameter to drive iterative refinement of both the polynomial approximation space and the corresponding collection of model evaluation points. As the adaptive method is largely heuristic, section 6 demonstrates the benefits of this approach with numerical examples.

2. Full tensor approximations. Tensorization is a common approach for lifting one-dimensional operators to higher dimensions. Not only are tensor products computationally convenient, but they provide much useful structure for analysis. In this section, we develop some essential background for computable tensor approximations, then apply it to problems of (i) approximating integrals with numerical quadrature, and (ii) approximating projection onto polynomial spaces with pseudospectral methods. In particular, we are interested in analyzing the errors associated with these approximations and in establishing conditions under which the approximations are *exact*.

2.1. General setting. Consider a collection of one-dimensional linear operators $\mathcal{L}^{(i)}$, where (i) indexes the operators used in different dimensions. In this work, $\mathcal{L}^{(i)}$ will be either an integral operator or an orthogonal projector onto some polynomial space. We can extend a collection of these operators into higher dimensions by constructing the tensor product operator

$$(2.1) \quad \mathcal{L}^{(\mathbf{d})} := \mathcal{L}^{(1)} \otimes \cdots \otimes \mathcal{L}^{(d)}.$$

The one-dimensional operators need not be identical; the properties of the resulting tensor operator are constructed independently from each dimension. The bold parenthetical superscript refers to the tensor operator instead of the constituent one-dimensional operators.

As the true operators are not available computationally, we work with a convergent sequence of computable approximations, $\mathcal{L}_m^{(i)}$, such that

$$(2.2) \quad \|\mathcal{L}^{(i)} - \mathcal{L}_m^{(i)}\| \rightarrow 0 \text{ as } m \rightarrow \infty$$

in some appropriate norm. Taking the tensor product of these approximations provides an approximation to the full tensor operator, $\mathcal{L}_{\mathbf{m}}^{(\mathbf{d})}$, where the level of the approximation may be individually selected in each dimension, so the tensor approximation is identified by a multi-index $\mathbf{m}$. Typically, and in the cases of interest in this work, the tensor approximation will converge in the same sense as the one-dimensional approximation as all components of $\mathbf{m} \rightarrow \infty$.

An important property of approximation algorithms is whether they are *exact* for some inputs; characterizing this set of inputs allows us to make useful statements about finite order approximations.

DEFINITION 2.1 (exact sets). *For an operator $\mathcal{L}$ and a corresponding approximation $\mathcal{L}_m$, define the exact set as $\mathcal{E}(\mathcal{L}_m) := \{f : \mathcal{L}(f) = \mathcal{L}_m(f)\}$ and the half exact set $\mathcal{E}_2(\mathcal{L}_m) := \{f : \mathcal{L}(f^2) = \mathcal{L}_m(f^2)\}$.*

The half exact set will help connect the exactness of a quadrature rule to that of the closely related pseudospectral operators. This notation is useful in proving the

following lemma, which relates the exact sets of one-dimensional approximations and tensor approximations.

LEMMA 2.2. *If a tensor approximation $\mathcal{L}_{\mathbf{m}}^{(\mathbf{d})}$ is constructed from one-dimensional approximations $\mathcal{L}_{m_i}^{(i)}$ with known exact sets, then*

$$(2.3) \quad \mathcal{E}(\mathcal{L}_{m_1}^{(1)}) \otimes \cdots \otimes \mathcal{E}(\mathcal{L}_{m_d}^{(d)}) \subseteq \mathcal{E}(\mathcal{L}_{\mathbf{m}}^{(\mathbf{d})}).$$

Proof. It is sufficient to show that the approximation is exact for an arbitrary monomial input $f(\mathbf{x}) = f^{(1)}(x^{(1)})f^{(2)}(x^{(2)}) \cdots f^{(d)}(x^{(d)})$, where $f^{(i)}(x^{(i)}) \in \mathcal{E}(\mathcal{L}_{m_i}^{(i)})$, because we may extend to sums by linearity:

$$\begin{aligned} \mathcal{L}_{\mathbf{m}}^{(\mathbf{d})}(f^{(1)} \cdots f^{(d)}) &= \mathcal{L}_{m_1}^{(1)}(f^{(1)}) \otimes \cdots \otimes \mathcal{L}_{m_d}^{(d)}(f^{(d)}) \\ &= \mathcal{L}^{(1)}(f^{(1)}) \otimes \cdots \otimes \mathcal{L}^{(d)}(f^{(d)}) = \mathcal{L}^{(\mathbf{d})}(f). \end{aligned}$$

The first step uses the tensor product structure of the operator and the second uses the definition of exact sets. $\square$

2.2. Multi-indices. Before continuing, we must make a short diversion to multi-indices, which provide helpful notation when dealing with tensor problems. A multi-index is a vector $\mathbf{i} \in \mathbb{N}_0^d$. An important notion for multi-indices is that of a *neighborhood*.

DEFINITION 2.3 (neighborhoods of multi-indices). *A forward neighborhood of a multi-index $\mathbf{k}$ is the multi-index set $n_f(\mathbf{k}) := \{\mathbf{k} + \mathbf{e}_i : \forall i \in \{1 \dots d\}\}$, where the $\mathbf{e}_i$ are the canonical unit vectors. The backward neighborhood of a multi-index $\mathbf{k}$ is the multi-index set $n_b(\mathbf{k}) := \{\mathbf{k} - \mathbf{e}_i : \forall i \in \{1 \dots d\}, \mathbf{k} - \mathbf{e}_i \in \mathbb{N}_0^d\}$.*

Smolyak algorithms rely on multi-index sets that are *admissible*.

DEFINITION 2.4 (admissible multi-indices and multi-index sets). *A multi-index $\mathbf{k}$ is admissible with respect to a multi-index set $\mathcal{K}$ if $n_b(\mathbf{k}) \subseteq \mathcal{K}$. A multi-index set $\mathcal{K}$ is admissible if every $\mathbf{k} \in \mathcal{K}$ is admissible with respect to $\mathcal{K}$.*

Two common admissible multi-index sets with simple geometric structure are *total-order* multi-index sets and *full tensor* multi-index sets. One often encounters total-order sets in the sparse grids literature and full tensor sets when dealing with tensor grids of points. The total-order multi-index set $\mathcal{K}_n^t$ comprises those multi-indices that lie within a d -dimensional simplex of side length n :

$$(2.4) \quad \mathcal{K}_n^t := \{\mathbf{k} \in \mathbb{N}_0^d : \|\mathbf{k}\|_1 \leq n\}.$$

The full tensor multi-index set $\mathcal{K}_{\mathbf{n}}^f$ is the complete grid of indices bounded termwise by a multi-index $\mathbf{n}$:

$$(2.5) \quad \mathcal{K}_{\mathbf{n}}^f := \{\mathbf{k} \in \mathbb{N}_0^d : \forall i \in \{1 \dots d\}, k_i < n_i\}.$$

2.3. Integrals and quadrature. Let $X^{(i)}$ be an open or closed interval of the real line $\mathbb{R}$. Then we define the weighted integral operator in one dimension as follows:

$$(2.6) \quad \mathcal{I}^{(i)}(f) := \int_{X^{(i)}} f(x) w^{(i)}(x) dx,$$

where $f : X^{(i)} \rightarrow \mathbb{R}$ is some real-valued function and $w^{(i)} : X^{(i)} \rightarrow \mathbb{R}^+$ is an integrable weight function. We may extend to higher dimensions by forming the tensor product integral $\mathcal{I}^{(\mathbf{d})}$, which uses separable weight functions and Cartesian product domains.

FIG. 2.1. Consider, one-dimensional Gaussian quadrature rule with three points, $\mathcal{Q}_3^{(i)}$, which is exact for fifth degree polynomials. This diagram depicts the exact set, $\mathcal{E}(\mathcal{Q}_3^{(i)})$, and half exact set, $\mathcal{E}_2(\mathcal{Q}_3^{(i)})$, of this quadrature rule.

Numerical quadrature approximates the action of an integral operator $\mathcal{I}^{(i)}$ with a weighted sum of point evaluations. For some family of quadrature rules, we write the “level m ” quadrature rule, comprised of $p^{(i)}(m) : \mathbb{N} \rightarrow \mathbb{N}$ points, as

$$(2.7) \quad \mathcal{I}^{(i)}(f) \approx \mathcal{Q}_m^{(i)}(f) := \sum_{j=1}^{p^{(i)}(m)} w_j^{(i)} f(x_j^{(i)}).$$

We call $p^{(i)}(m)$ the growth rate of the quadrature rule, and its form depends on the quadrature family; some rules only exist for certain numbers of points and others may be tailored, for example, to produce linear or exponential growth in the number of quadrature points with respect to the level.

Many quadrature families are exact if f is a polynomial of a degree $a^{(i)}(m)$ or less, which allows us to specify a well-structured portion of the exact set for these quadrature rules:

$$(2.8) \quad \mathbb{P}_{a^{(i)}(m)} \subseteq \mathcal{E}(\mathcal{Q}_m^{(i)}),$$

$$(2.9) \quad \mathbb{P}_{\lfloor a^{(i)}(m)/2 \rfloor} \subseteq \mathcal{E}_2(\mathcal{Q}_m^{(i)}),$$

where $\mathbb{P}_a$ is the space of polynomials of degree a or less. It is intuitive and useful to draw the exact set as in Figure 2.1. For this work, we rely on quadrature rules that exhibit polynomial accuracy of increasing order, which is sufficient to demonstrate convergence for functions in L^2 [4].

Tensor product quadrature rules are straightforward approximations of tensor product integrals that inherit convergence properties from the one-dimensional case. The exact set of a tensor product quadrature rule includes the tensor product of the constituent approximations’ exact sets, as guaranteed by Lemma 2.2 and depicted in Figure 2.2.

2.4. Polynomial projection. A polynomial chaos expansion approximates a function with a weighted sum of orthonormal polynomials [38, 40]. Let $\mathcal{H}^{(i)} := L^2(X^{(i)}, w^{(i)})$ be the separable Hilbert space of square-integrable functions $f : X^{(i)} \rightarrow \mathbb{R}$, with inner product defined by the weighted integral $\langle f, g \rangle = \mathcal{I}^{(i)}(fg)$, and $w^{(i)}(x)$ normalized so that it may represent a probability density. Let $\mathbb{P}_n^{(i)}$ be the space of univariate polynomials of degree n or less. Now let $\mathcal{P}_n^{(i)} : \mathcal{H}^{(i)} \rightarrow \mathbb{P}_n^{(i)}$ be an orthogonal projector onto this subspace, written in terms of polynomials $\{\psi_j^{(i)}(x) : j \in \mathbb{N}_0\}$

FIG. 2.2. Consider a two-dimensional quadrature rule constructed from three point Gaussian quadrature rules, $\mathcal{Q}_{(3,3)}^{(2)}$. This diagram depicts the exact set, $\mathcal{E}(\mathcal{Q}_{(3,3)}^{(2)})$, and half exact set, $\mathcal{E}_2(\mathcal{Q}_{(3,3)}^{(2)})$.

orthonormal with respect to the inner product of $\mathcal{H}^i$:

$$(2.10) \quad \mathcal{P}_n^{(i)}(f) := \sum_{j=0}^n \langle f(x), \psi_j^{(i)}(x) \rangle \psi_j^{(i)}(x) = \sum_{i=0}^n f_j \psi_j^{(i)}(x).$$

The polynomial space $\mathbb{P}_n^{(i)}$ is of course the *range* of the projection operator. These polynomials are dense in $\mathcal{H}^{(i)}$, so the polynomial approximation of any $f \in \mathcal{H}^{(i)}$ converges in the L^2 sense as $n \rightarrow \infty$ [4, 40]. If $f \in \mathcal{H}^{(i)}$, the coefficients must satisfy $\sum_{i=0}^{\infty} f_i^2 < \infty$.

Projections with finite degree n omit terms of the infinite series, thus incurring *truncation error*. We can write this error as

$$(2.11) \quad \|f - \mathcal{P}_n^{(i)}(f)\|_2^2 = \left\| \sum_{j=n+1}^{\infty} f_j \psi_j^{(i)} \right\|_2^2 = \sum_{j=n+1}^{\infty} f_j^2 < \infty.$$

Hence, we may reduce the truncation error to any desired level by increasing n , removing terms from the sum in (2.11) [4, 17].

The d -dimensional version of this problem requires approximating functions in the Hilbert space $\mathcal{H}^{(d)} := \mathcal{H}^{(1)} \otimes \cdots \otimes \mathcal{H}^{(d)}$ via a tensor product basis of the univariate polynomials defined above:

$$(2.12) \quad \mathcal{P}_{\mathbf{n}}^{(d)}(f) = \sum_{i_1=0}^{n_1} \cdots \sum_{i_d=0}^{n_d} \langle f, \Psi_{\mathbf{i}} \rangle \Psi_{\mathbf{i}},$$

where $\Psi_{\mathbf{i}}(\mathbf{x}) := \prod_{j=1}^d \psi_{i_j}^{(j)}(x^{(j)})$. The multi-index $\mathbf{n}$ tailors the range of the projection to include a rectangular subset of polynomials.

As in the one-dimensional case, truncation induces error equal to the sum of the squares of the omitted coefficients, which we may similarly reduce to zero as $n_i \rightarrow \infty \forall i$. The multivariate polynomial expansion also converges in an L^2 sense for any $f \in \mathcal{H}^{(d)}$ [4].

2.5. Aliasing errors in pseudospectral approximation. The inner products defining the expansion coefficients above are not directly computable. Pseudospectral approximation provides a practical nonintrusive algorithm by approximating these

inner products with quadrature. Define the pseudospectral approximation in one dimension as

$$\begin{aligned} \mathcal{S}_m^{(i)}(f) &:= \sum_{j=0}^{q^{(i)}(m)} \mathcal{Q}_m^{(i)} \left(f \psi_j^{(i)} \right) \psi_j^{(i)}(x) \\ (2.13) \qquad &= \sum_{j=0}^{q^{(i)}(m)} \tilde{f}_j \psi_j^{(i)}(x), \end{aligned}$$

where $q^{(i)}(m)$ is the polynomial truncation at level m , to be specified shortly [4, 17]. Pseudospectral approximations are constructed around a level m quadrature rule, and are designed to include as many terms in the sum as possible while maintaining accuracy. Assuming that $f \in L^2$, we can compute the L^2 error between the pseudospectral approximation and an exact projection onto the same polynomial space:

$$(2.14) \qquad \left\| \mathcal{P}_{q^{(i)}(m)}^{(i)}(f) - \mathcal{S}_m^{(i)}(f) \right\|_2^2 = \left\| \sum_{j=0}^{q^{(i)}(m)} f_j \psi_j^{(i)} - \sum_{k=0}^{q^{(i)}(m)} \tilde{f}_k \psi_k^{(i)} \right\|_2^2 = \sum_{j=0}^{q^{(i)}(m)} (f_j - \tilde{f}_j)^2.$$

This quantity is the *aliasing error* [4, 17]. The error is nonzero because quadrature in general only approximates integrals; hence, each $\tilde{f}_i$ is an approximation of f_i . The pseudospectral operator also incurs truncation error, as before, which is orthogonal to the aliasing error. We can expand each approximate coefficient as

$$\begin{aligned} \tilde{f}_j &= \mathcal{Q}_m^{(i)} \left(f \psi_j^{(i)} \right) \\ &= \sum_{k=0}^{\infty} f_k \mathcal{Q}_m^{(i)} \left(\psi_j^{(i)} \psi_k^{(i)} \right) \\ (2.15) \qquad &= \sum_{k=0}^{q^{(i)}(m)} f_k \mathcal{Q}_m^{(i)} \left(\psi_j^{(i)} \psi_k^{(i)} \right) + \sum_{l=q^{(i)}(m)+1}^{\infty} f_l \mathcal{Q}_m^{(i)} \left(\psi_j^{(i)} \psi_l^{(i)} \right). \end{aligned}$$

The first step substitutes in the polynomial expansion of f , which we assume is convergent, and rearranges using linearity. The second step partitions the sum around the truncation of the pseudospectral expansion. Although the basis functions are orthonormal, $\langle \psi_j^{(i)}, \psi_k^{(i)} \rangle = \delta_{jk}$, we cannot assume in general that the approximation $\mathcal{Q}_m^{(i)}(\psi_j^{(i)} \psi_k^{(i)}) = \delta_{jk}$. Now substitute (2.15) back into the aliasing error expression:

$$\begin{aligned} (2.16) \qquad &\sum_{j=0}^{q^{(i)}(m)} (f_j - \tilde{f}_j)^2 \\ &= \sum_{j=0}^{q^{(i)}(m)} \left(f_j - \sum_{k=0}^{q^{(i)}(m)} f_k \mathcal{Q}_m^{(i)} \left(\psi_j^{(i)} \psi_k^{(i)} \right) - \sum_{l=q^{(i)}(m)+1}^{\infty} f_l \mathcal{Q}_m^{(i)} \left(\psi_j^{(i)} \psi_l^{(i)} \right) \right)^2. \end{aligned}$$

This form reveals the intimate link between the accuracy of pseudospectral approximations and the polynomial accuracy of quadrature rules. All aliasing is attributed to the inability of the quadrature rule to determine the orthonormality of

the basis polynomials, causing one coefficient to corrupt another. The contribution of the first two parenthetical terms on the right of (2.16) is called *internal aliasing*, while the third term is called *external aliasing*. Internal aliasing is due to inaccuracies in $\mathcal{Q}(\psi_j^{(i)}\psi_k^{(i)})$ when *both* $\psi_j^{(i)}$ and $\psi_k^{(i)}$ are included in the expansion, while external aliasing occurs when *only one* of these polynomials is included in the expansion. For many practical quadrature rules (and for all those used in this work), if $j \neq k$ and $\psi_j^{(i)}\psi_k^{(i)} \notin \mathcal{E}(\mathcal{Q})$, and hence the discrete inner product is not zero, then $\|\mathcal{Q}(\psi_j^{(i)}\psi_k^{(i)})\|_2$ is $\mathcal{O}(1)$ [34]. As a result, the magnitude of an aliasing error typically corresponds to the magnitude of the aliased coefficients.

In principle, both types of aliasing error are driven to zero by sufficiently powerful quadrature, but we are left to select $q^{(i)}(m)$ for a particular quadrature level m . External aliasing is $\mathcal{O}(1)$ in the magnitude of the truncation error, and thus it is driven to zero as long as $q^{(i)}(m)$ increases with m . Internal aliasing could be $\mathcal{O}(1)$ with respect to the function of interest, meaning that the procedure neither converges nor provides a useful approximation. Therefore, the obvious option is to include as many terms as possible while setting the internal aliasing to zero.

For quadrature rules with polynomial exactness, we may accomplish this by setting $q^{(i)}(m) = \text{floor}(a^{(i)}(m)/2)$. This ensures that the internal aliasing of $\mathcal{S}_m^{(i)}$ is zero, because $\forall j, k \leq q^{(i)}(m), \psi_j^{(i)}\psi_k^{(i)} \in \mathcal{E}(\mathcal{Q}_m^{(i)})$. Equivalently, a pseudospectral operator $\mathcal{S}_m^{(i)}$ using quadrature $\mathcal{Q}_m^{(i)}$ has a range corresponding to the half exact set $\mathcal{E}_2(\mathcal{Q}_m^{(i)})$. Alternatively, we may justify this choice by noting that it makes the pseudospectral approximation exact on its range, $\mathbb{P}_{q^{(i)}(m)}^{(i)} \subseteq \mathcal{E}(\mathcal{S}_m^{(i)})$.

Given this choice of $q^{(i)}(m)$, we wish to show that the pseudospectral approximation converges to the true function, where the magnitude of the error is as follows:

$$(2.17) \quad \|f - \mathcal{S}_m^{(i)}(f)\|_2^2 = \sum_{j=0}^{q^{(i)}(m)} \left(\sum_{k=q^{(i)}(m)+1}^{\infty} f_k \mathcal{Q}(\psi_j^{(i)}\psi_k^{(i)}) \right)^2 + \sum_{l=q^{(i)}(m)+1}^{\infty} f_l^2.$$

The two terms on the right-hand side comprise the external aliasing and the truncation error, respectively. We already know that the truncation error goes to zero as $q^{(i)}(m) \rightarrow \infty$. The external aliasing also vanishes for functions $f \in L^2$, as the truncated portion of f likewise decreases [34]. In the case of Gaussian quadrature rules, a link to interpolation provides precise rates for the convergence of the pseudospectral operator based on the regularity of f [4].

As with quadrature algorithms, our analysis of pseudospectral approximation in one dimension is directly extensible to multiple dimensions via full tensor products. We may thus conclude that $\mathcal{S}_m^{(d)}$ converges to the projection onto the tensor product polynomial space in the same sense. The exact set follows Lemma 2.2, and hence the tensor product approximation inherits zero internal aliasing if suitable one-dimensional operators are used.

3. Smolyak algorithms. Thus far, we have developed polynomial approximations of multivariate functions by taking tensor products of one-dimensional pseudospectral operators. Smolyak algorithms avoid the exponential cost of full tensor products when the input dimensions are not fully coupled, allowing the use of a telescoping sum to blend different lower-order full tensor approximations.

Example 3.1. Suppose that $f(x, y) = x^7 + y^7 + x^3y$. To construct a polynomial expansion with both the x^7 and y^7 terms, a full tensor pseudospectral algorithm would

estimate all the polynomial terms up to x^7y^7 , because tensor algorithms fully couple the dimensions. This mixed term is costly, requiring, in this case, an 8×8 point grid for Gaussian quadratures. The individual terms can be had much more cheaply, using 8×1 , 1×8 , and 4×2 grids, respectively. Smolyak algorithms help realize such savings in practice.

This section reviews the construction of Smolyak algorithms and presents a new theorem about the exactness of Smolyak algorithms built around arbitrary admissible index sets. We apply these results to quadrature and pseudospectral approximation, allowing a precise characterization of their errors.

3.1. General Smolyak algorithms. As in section 2, assume that we have for every dimension $i = 1 \dots d$ a convergent sequence $\mathcal{L}_k^{(i)}$ of approximations. Let $\mathcal{L}$ denote the collection of these sequences over all the dimensions. Define the difference operators

$$(3.1) \quad \Delta_0^{(i)} := \mathcal{L}_0^{(i)} = 0,$$

$$(3.2) \quad \Delta_n^{(i)} := \mathcal{L}_n^{(i)} - \mathcal{L}_{n-1}^{(i)}.$$

For any i , we may write the exact or “true” operator as the telescoping series

$$(3.3) \quad \mathcal{L}^{(i)} = \sum_{k=0}^{\infty} \mathcal{L}_k^{(i)} - \mathcal{L}_{k-1}^{(i)} = \sum_{k=0}^{\infty} \Delta_k^{(i)}.$$

Now we may write the tensor product of the exact operators as the tensor product of the telescoping sums, and interchange the product and sum:

$$(3.4) \quad \begin{aligned} \mathcal{L}^{(1)} \otimes \dots \otimes \mathcal{L}^{(d)} &= \sum_{k_1=0}^{\infty} \Delta_{k_1}^{(1)} \otimes \dots \otimes \sum_{k_d=0}^{\infty} \Delta_{k_d}^{(d)} \\ &= \sum_{\mathbf{k}=0}^{\infty} \Delta_{k_1}^{(1)} \otimes \dots \otimes \Delta_{k_d}^{(d)} \end{aligned}$$

Smolyak’s idea is to approximate the tensor product operator with truncations of this sum [31]:

$$(3.5) \quad A(\mathcal{K}, d, \mathcal{L}) := \sum_{\mathbf{k} \in \mathcal{K}} \Delta_{k_1}^{(1)} \otimes \dots \otimes \Delta_{k_d}^{(d)}.$$

We refer to the multi-index set $\mathcal{K}$ as the *Smolyak multi-index set*, and it must be admissible for the sum to telescope correctly. Smolyak specifically suggested truncating with a total-order multi-index set, which is the most widely studied choice. However, we can compute the approximation with any admissible multi-index set. Although the expression above is especially clean, it is not the most useful form for computation. We can reorganize the terms of (3.5) to construct a weighted sum of the tensor operators:

$$(3.6) \quad A(\mathcal{K}, d, \mathcal{L}) = \sum_{\mathbf{k} \in \mathcal{K}} c_{\mathbf{k}} \mathcal{L}_{k_1}^{(1)} \otimes \dots \otimes \mathcal{L}_{k_d}^{(d)},$$

where $c_{\mathbf{k}}$ are integer *Smolyak coefficients* computed from the combinatorics of the difference formulation. One can compute the coefficients through a simple iteration

over the index set and use (3.5) to determine which full tensor rules are incremented or decremented. In general, these coefficients are nonzero near the leading surface of the Smolyak multi-index set, reflecting the mixing of the most accurate constituent full tensor approximations.

If each sequence of one-dimensional operators converges, then the Smolyak approximation converges to the tensor product of exact operators as $\mathcal{K} \rightarrow \mathbb{N}_0^d$. For the isotropic simplex index set, some precise rates of convergence are known with respect to the side length of the simplex [36, 37, 35, 29, 30]. Although general admissible Smolyak multi-index sets are difficult to study theoretically, they allow detailed customization to the anisotropy of a particular function.

3.2. Exactness of Smolyak algorithms. In the one-dimensional and full tensor settings, we have characterized approximation algorithms through their exact sets—those inputs for which the algorithm is precise. This section shows that if the constituent one-dimensional approximations have nested exact sets, Smolyak algorithms are the ideal blending of different full tensor approximations from the perspective of exact sets; that is, the exact set of the Smolyak algorithm contains the union of the exact sets of the component full tensor approximations. This result will facilitate subsequent analysis of sparse quadrature and pseudospectral approximation algorithms. This theorem and our proof closely follow the framework provided by Novak and Ritter [25, 26] and Barthelmann, Novak, and Ritter [2], but include a generalization to arbitrary Smolyak multi-index sets.

THEOREM 3.2. *Let $A(\mathcal{K}, d, \mathcal{L})$ be a Smolyak algorithm composed of linear operators with nested exact sets, i.e., with $m \leq m'$ implying that $\mathcal{E}(\mathcal{L}_m^{(i)}) \subseteq \mathcal{E}(\mathcal{L}_{m'}^{(i)})$ for $i = 1 \dots d$, where $\mathcal{K}$ is admissible. Then the exact set of $A(\mathcal{K}, d, \mathcal{L})$ contains*

$$\begin{aligned} \mathcal{E}(A(\mathcal{K}, d, \mathcal{L})) &\supseteq \bigcup_{\mathbf{k} \in \mathcal{K}} \mathcal{E}(\mathcal{L}_{k_1}^{(1)} \otimes \dots \otimes \mathcal{L}_{k_d}^{(d)}) \\ (3.7) \quad &\supseteq \bigcup_{\mathbf{k} \in \mathcal{K}} \mathcal{E}(\mathcal{L}_{k_1}^{(1)}) \otimes \dots \otimes \mathcal{E}(\mathcal{L}_{k_d}^{(d)}). \end{aligned}$$

Proof. We begin by introducing notation to incrementally build a multi-index set dimension by dimension. For a multi-index set $\mathcal{K}$ of dimension d , let the restriction of the multi-indices to the first i dimensions be $\mathcal{K}^{(i)} := \{\mathbf{k}_{1:i} = (k_1, \dots, k_i) : \mathbf{k} \in \mathcal{K}\}$. Furthermore, define subsets of $\mathcal{K}$ based on the i th element of the multi-indices, $\mathcal{K}_j^{(i)} := \{\mathbf{k}_{1:i} : \mathbf{k} \in \mathcal{K}^{(i)} \text{ and } k_{i+1} = j\}$. These sets are nested, $\mathcal{K}_j^{(i)} \supseteq \mathcal{K}_{j+1}^{(i)}$, because $\mathcal{K}$ is admissible. Also let $k_i^{\max}$ denote the maximum value of the i th component of the multi-indices in the set $\mathcal{K}$.

Using this notation, one can construct $\mathcal{K}$ inductively,

$$(3.8) \quad \mathcal{K}^{(1)} = \{1, \dots, k_1^{\max}\},$$

$$(3.9) \quad \mathcal{K}^{(i)} = \bigcup_{j=1}^{k_i^{\max}} \mathcal{K}_j^{(i-1)} \otimes j, \quad i = 2 \dots d.$$

It is sufficient to prove that the Smolyak operator is exact for an arbitrary f with tensor structure $f = f_1 \times \dots \times f_d$. Suppose there exists a $\mathbf{k}^*$ such that $f \in \mathcal{E}(\mathcal{L}_{\mathbf{k}^*}^{(d)})$. We will show that if $\mathcal{K}$ is an admissible multi-index set containing $\mathbf{k}^*$, then $A(\mathcal{K}, d, \mathcal{L})$ is exact on f . We do so by induction on the dimension i of the Smolyak operator and the function.

First, consider the $i = 1$ case. $A(\mathcal{K}^{(1)}, 1, \mathcal{L}) = \mathcal{L}_{k_1^{\max}}^{(1)}$, where $k_1^{\max} \geq k_1^*$. Hence $\mathcal{E}(A(\mathcal{K}^{(1)}, 1, \mathcal{L})) = \mathcal{E}(\mathcal{L}_{k_1^{\max}}^{(1)})$.

For the induction step, we construct the $(i + 1)$ -dimensional Smolyak operator in terms of the i -dimensional operator:

$$(3.10) \quad A(\mathcal{K}^{(i+1)}, i + 1, \mathcal{L}) = \sum_{j=1}^{k_{i+1}^{\max}} A(\mathcal{K}_j^{(i)}, i, \mathcal{L}) \otimes (\mathcal{L}_j^{(i+1)} - \mathcal{L}_{j-1}^{(i+1)}).$$

This sum is over increasing levels of accuracy in the $(i + 1)$ th dimension. We know the level required for the approximate operator to be exact in this dimension; this may be expressed as

$$(3.11) \quad \mathcal{L}_j^{(i+1)}(f_{i+1}) = \mathcal{L}_{j-1}^{(i+1)}(f_{i+1}) = \mathcal{L}^{(i+1)}(f_{i+1}) \text{ when } j - 1 \geq k_{i+1}^*.$$

Therefore the sum (3.10) can be truncated at the k_{i+1}^* term, as the differences of higher terms are zero when applied to f :

$$(3.12) \quad A(\mathcal{K}^{(i+1)}, i + 1, \mathcal{L}) = \sum_{j=1}^{k_{i+1}^*} A(\mathcal{K}_j^{(i)}, i, \mathcal{L}) \otimes (\mathcal{L}_j^{(i+1)} - \mathcal{L}_{j-1}^{(i+1)}).$$

Naturally, $\mathbf{k}_{1:i}^* \in \mathcal{K}_{k_{i+1}^*}^{(i)}$. By nestedness, $\mathbf{k}_{1:i}^*$ is also contained in $\mathcal{K}_j^{(i)}$ for $j \leq k_{i+1}^*$. The induction hypothesis then guarantees

$$(3.13) \quad f_1 \otimes \cdots \otimes f_i \in \mathcal{E}(A(\mathcal{K}_j^{(i)}, i, \mathcal{L})) \quad \forall j \leq k_{i+1}^*.$$

Applying the $(i + 1)$ -dimensional Smolyak operator to the truncated version of f yields

$$(3.14) \quad \begin{aligned} & A(\mathcal{K}^{(i+1)}, i + 1, \mathcal{L})(f_1 \otimes \cdots \otimes f_{i+1}) \\ &= \sum_{j=1}^{k_{i+1}^*} A(\mathcal{K}_j^{(i)}, i, \mathcal{L})(f_1 \otimes \cdots \otimes f_i) \otimes (\mathcal{L}_j^{(i+1)} - \mathcal{L}_{j-1}^{(i+1)})(f_{i+1}). \end{aligned}$$

Since each of the i -dimensional Smolyak algorithms is exact, by the induction hypothesis, we replace them with the true operators and rearrange by linearity to obtain

$$(3.15) \quad \begin{aligned} A(\mathcal{K}^{(i+1)}, i + 1, \mathcal{L})(f_1 \otimes \cdots \otimes f_{i+1}) &= \mathcal{L}^{(i)}(f_1 \otimes \cdots \otimes f_i) \otimes \sum_{j=1}^{k_{i+1}^*} (\mathcal{L}_j^{(i+1)} - \mathcal{L}_{j-1}^{(i+1)})(f_{i+1}) \\ &= \mathcal{L}^{(i)}(f_1 \otimes \cdots \otimes f_i) \otimes \mathcal{L}_{k_{i+1}^*}^{(i+1)}(f_{i+1}). \end{aligned}$$

The approximation in the $(i + 1)$ th dimension is exactly of the level needed to be

(a) The exact set for a level four Smolyak quadrature in two dimensions, based on *linear* growth Gaussian quadrature rules.

(b) The exact set for a level three Smolyak quadrature in two dimensions, based on *exponential* growth Gaussian quadrature rules.

FIG. 3.1. The exact set diagram for two Smolyak quadrature rules, and the corresponding basis for a Smolyak pseudospectral approximation. $\mathcal{E}(\mathcal{Q})$ is shown in a solid line, $\mathcal{E}_2(\mathcal{Q})$ is the dashed line. The staircase appearance results from the superposition of rectangular full tensor exact sets.

exact on the $(i + 1)$ th component of f . Then (3.15) becomes

$$(3.16) \quad \mathcal{L}^{(i)}(f_1 \otimes \cdots \otimes f_i) \otimes \mathcal{L}^{(i+1)}(f_{i+1}) = \mathcal{L}^{(i+1)}(f_1 \otimes \cdots \otimes f_{i+1}).$$

Thus the Smolyak operator is precise for f , and the claim is proven. $\square$

3.3. Smolyak quadrature. We recall the most familiar use of Smolyak algorithms, sparse quadrature. Consider a family of one-dimensional quadrature rules $\mathcal{Q}_k^{(i)}$ in each dimension $i = 1 \dots d$; denote these rules by $\mathcal{Q}$. The resulting Smolyak algorithm is written as:

$$(3.17) \quad A(\mathcal{K}, d, \mathcal{Q}) = \sum_{\mathbf{k} \in \mathcal{K}} c_{\mathbf{k}} \mathcal{Q}_{\mathbf{k}}^{(d)}.$$

This approximation inherits its convergence from the one-dimensional operators. The set of functions that are exactly integrated by a Smolyak quadrature algorithm is described as a corollary of Theorem 3.2.

COROLLARY 3.3. *For a sparse quadrature rule satisfying the hypothesis of Theorem 3.2,*

$$(3.18) \quad \mathcal{E}(A(\mathcal{K}, d, \mathcal{Q})) \supseteq \bigcup_{\mathbf{k} \in \mathcal{K}} \mathcal{E}(\mathcal{Q}_{\mathbf{k}}^{(d)}).$$

Quadrature rules with polynomial accuracy do have nested exact sets, as required by the theorem. An example of Smolyak quadrature exact sets is shown in Figure 3.1.

3.4. Smolyak pseudospectral approximation. Applying Smolyak's algorithm to pseudospectral approximation operators yields a sparse algorithm that converges under similar conditions as the one-dimensional operators from which it is constructed. This algorithm is written as

$$(3.19) \quad A(\mathcal{K}, d, \mathcal{S}) = \sum_{\mathbf{k} \in \mathcal{K}} c_{\mathbf{k}} \mathcal{S}_{\mathbf{k}}^{(d)}.$$

The Smolyak algorithm is therefore a sum of different full tensor pseudospectral approximations, where each approximation is built around the polynomial accuracy of a single full tensor quadrature rule. It is not naturally expressed as a set of formulas for the polynomial coefficients, because different approximations include different polynomials. The term $\Psi_{\mathbf{j}}$ is included in the Smolyak approximation if and only if $\exists \mathbf{k} \in \mathcal{K} : \Psi_{\mathbf{j}} \in \mathcal{E}_2(\mathcal{Q}_{\mathbf{k}}^{(\mathbf{d})})$. Here, $\mathcal{Q}_{\mathbf{k}}^{(\mathbf{d})}$ is the full tensor quadrature rule used by the full tensor pseudospectral approximation $\mathcal{S}_{\mathbf{k}}^{(\mathbf{d})}$. As in the full tensor case, the half exact set of a Smolyak quadrature rule defines the range of the Smolyak pseudospectral approximation.

Once again, the Smolyak construction guarantees that the convergence of this approximation is inherited from its constituent one-dimensional approximations. Our choices for the pseudospectral operators ensure nestedness of the constituent exact sets, so we may use Theorem 3.2 to ensure that Smolyak pseudospectral algorithms are exact on their range.

COROLLARY 3.4. *If the constituent one-dimensional pseudospectral rules have no internal aliasing and satisfy the conditions of Theorem 3.2, then the resulting Smolyak pseudospectral algorithm has no internal aliasing.*

We additionally provide a theorem that characterizes the external aliasing properties of Smolyak pseudospectral approximation, which the next section will contrast with direct quadrature.

THEOREM 3.5. *Let $\Psi_{\mathbf{j}}$ be a polynomial term included in the expansion provided by the Smolyak algorithm $A(\mathcal{K}, d, \mathcal{S})$, and let $\Psi_{\mathbf{j}'}$ be a polynomial term not included in the expansion. There is no external aliasing of $\Psi_{\mathbf{j}'}$ onto $\Psi_{\mathbf{j}}$ if either of the following conditions is satisfied: (a) there exists a dimension i for which $j'_i < j_i$, or (b) there exists a multi-index $\mathbf{k} \in \mathcal{K}$ such that $\Psi_{\mathbf{j}}$ is included in the range of $\mathcal{S}_{\mathbf{k}}^{(\mathbf{d})}$ and $\Psi_{\mathbf{j}'} \Psi_{\mathbf{j}} \in \mathcal{E}(\mathcal{Q}_{\mathbf{k}}^{(\mathbf{d})})$, where $\mathcal{Q}_{\mathbf{k}}^{(\mathbf{d})}$ is the quadrature rule used in $\mathcal{S}_{\mathbf{k}}^{(\mathbf{d})}$.*

Proof. If condition (a) is satisfied, then $\Psi_{\mathbf{j}}$ and $\Psi_{\mathbf{j}'}$ are orthogonal in dimension i , and hence that inner product is zero. Every quadrature rule that computes the coefficient $f_{\mathbf{j}}$ corresponding to basis term $\Psi_{\mathbf{j}}$ is accurate for polynomials of at least order $2\mathbf{j}$. Since $j'_i + j_i < 2j_i$, every rule that computes the coefficient can numerically resolve the orthogonality, and therefore there is no aliasing. If condition (b) is satisfied, then the result follows from the cancellations exploited by the Smolyak algorithm, as seen in the proof of Theorem 3.2. $\square$

These two statements yield extremely useful properties. First, any Smolyak pseudospectral algorithm, regardless of the admissible Smolyak multi-index set used, has no internal aliasing; this feature is important in practice and not obviously true. Second, while there is external aliasing as expected, the algorithm uses basis orthogonality to limit which external coefficients can alias onto an included coefficient. The Smolyak pseudospectral algorithm is thus a practically “useful” approximation, in that one can tailor it to perform a desired amount of work while guaranteeing reliable approximations of the selected coefficients. Computing an accurate approximation of the function only requires including sufficient terms so that the truncation and external aliasing errors are small.

4. Comparing direct quadrature to Smolyak pseudospectral approximation. The current uncertainty quantification literature often suggests a *direct quadrature* approach for constructing polynomial chaos expansions [42, 21, 7, 18]. In this section, we describe this approach and show that, in comparison to a true Smolyak algorithm, it is inaccurate or inefficient in almost all cases. Our comparisons

will contrast the theoretical error performance of the algorithms and provide simple numerical examples that illustrate typical errors and why they arise.

4.1. Direct quadrature polynomial expansions. At first glance, direct quadrature is quite simple. First, choose a multi-index set $\mathcal{J}$ to define a truncated polynomial expansion

$$(4.1) \quad f \approx \sum_{\mathbf{j} \in \mathcal{J}} \tilde{f}_{\mathbf{j}} \Psi_{\mathbf{j}}.$$

The index set $\mathcal{J}$ is typically admissible, but need not be. Second, select any d -dimensional quadrature rule $\mathcal{Q}^{(\mathbf{d})}$, and estimate every coefficient as

$$(4.2) \quad \tilde{f}_{\mathbf{j}} = \mathcal{Q}^{(\mathbf{d})}(f \Psi_{\mathbf{j}}).$$

Unlike the Smolyak approach, we are left to choose $\mathcal{J}$ and $\mathcal{Q}^{(\mathbf{d})}$ independently, giving the appearance of flexibility. In practice, this produces a more complex and far more subtle version of the truncation trade-off discussed in section 2. Below, we will be interested in selecting $\mathcal{Q}$ and $\mathcal{J}$ to replicate the quadrature points and output range of the Smolyak approach, as it provides a benchmark for achievable performance.

Direct quadrature does not converge for every choice of $\mathcal{J}$ and $\mathcal{Q}^{(\mathbf{d})}$; consider the trivial case where $\mathcal{J}$ does not grow infinitely. It is possible that including far too many terms in $\mathcal{J}$ relative to the polynomial exactness of $\mathcal{Q}^{(\mathbf{d})}$ could lead to a nonconvergent algorithm. Although this behavior contrasts with the straightforward convergence properties of Smolyak algorithms, most reasonable choices for direct quadrature do converge, and hence this is not our primary argument against the approach.

Instead, our primary concern is aliasing in direct quadrature and how it reduces performance at finite order. Both internal and external aliasing are governed by the same basic rule, which is just a restatement of how we defined aliasing in section 2.5.

Remark 4.1. For a multi-index set $\mathcal{J}$ and a quadrature rule $\mathcal{Q}^{(\mathbf{d})}$, the corresponding direct quadrature polynomial expansion has no aliasing between two polynomial terms if $\Psi_{\mathbf{j}} \Psi_{\mathbf{j}'} \in \mathcal{E}(\mathcal{Q}^{(\mathbf{d})})$.

The next two sections compare the internal and external aliasing with both theory and simple numeric examples.

4.2. Internal aliasing in direct quadrature. As an extension of Remark 4.1, direct quadrature has no internal aliasing whenever every pair $\mathbf{j}, \mathbf{j}' \in \mathcal{J}$ has no aliasing. We can immediately conclude that for any basis set $\mathcal{J}$, there is some quadrature rule sufficiently powerful to avoid internal aliasing errors. In practice, however, this rule may not be a desirable one.

Example 4.2. Assume that for some function with two inputs, we wish to include the polynomial basis terms $(a, 0)$ and $(0, b)$. By Remark 4.1, the product of these two terms must be in the exact set; hence, the quadrature must include at least a full tensor rule of accuracy (a, b) . Although we have not asked for any coupling, direct quadrature must assume full coupling of the problem in order to avoid internal aliasing.

Therefore we reach the surprising conclusion that direct quadrature inserts significant coupling into the problem, whereas we selected a Smolyak quadrature rule in hopes of leveraging the absence of that very coupling—making the choice inconsistent. For most sparse quadrature rules, we cannot include as many polynomial terms as the Smolyak pseudospectral approach without incurring internal aliasing, because the quadrature is not powerful enough in the mixed dimensions.

FIG. 4.1. The exact set and polynomials included in the direct quadrature construction from Example 4.3.

FIG. 4.2. Numerical results for Example 4.3; each color square indicates the log of the coefficient magnitude for the basis function at that position. The circle identifies the correct nonzero coefficient.

Example 4.3. Let $\mathbf{X}$ be the two-dimensional domain $[-1, 1]^2$. Select a uniform weight function, which corresponds to a Legendre polynomial basis. Let $f(x, y) = \psi_0(x)\psi_4(y)$. Use Gauss–Legendre quadrature and an exponential growth rule, such that $p^{(i)}(m) = 2^{m-1}$. Select a sparse quadrature rule based on a total order multi-index set $\mathcal{Q}_{K_5^t}^2$. Figure 4.1 shows the exact set of this Smolyak quadrature rule (solid line) along with its half exact set (dashed line), which encompasses all the terms in the direct quadrature polynomial expansion.

Now consider the $\mathbf{j} = (8, 0)$ polynomial, which is in the half exact set. The product of the $(0, 4)$ and $(8, 0)$ polynomial terms is $(8, 4)$, which is not within the exact set of the sparse rule. Hence, $(0, 4)$ aliases onto $(8, 0)$ because this quadrature rule has limited accuracy in the mixed dimensions.

Using both the Smolyak pseudospectral and direct quadrature methods, we numerically compute the polynomial expansion for this example. The resulting coefficients are shown in Figure 4.2. Even though the two methods use the same information and project f onto the same basis, the Smolyak result has no internal aliasing while direct quadrature shows significant internal aliasing. Although both methods correctly compute the $(0, 4)$ coefficient, direct quadrature shows aliasing on $(8, 0)$ as

predicted, and also on $(10, 0)$, $(12, 0)$, and $(14, 0)$. In this case, direct quadrature is unable to determine the order of the input function or even whether the input is a function of x_1 or x_2 . Alternating terms are computed correctly because of the parity of the functions.

The $\mathcal{O}(1)$ errors observed in this simple example demonstrate why it is crucial to eliminate internal aliasing in the construction of one-dimensional pseudospectral approximations and to ensure that the full tensor and Smolyak algorithms inherit that property. More complex functions demonstrate the same type of error, except that the errors resulting from multiple source terms are superimposed. Examples of the latter are given by Constantine, Eldred, and Phipps [6].

There are some choices for which direct quadrature has no internal aliasing: full tensor quadrature rules and, notably, Smolyak quadrature constructed from one-dimensional Gaussian quadrature rules with $p^{(i)}(m) = m$, truncated according to an isotropic total-order multi-index set. However, many useful sparser or more tailored Smolyak quadrature rules, e.g., based on exponential growth quadrature rules or adaptive anisotropic Smolyak index sets, will incur internal aliasing if the basis selection matches the range of the Smolyak algorithm. This makes them a poor choice when a comparable Smolyak pseudospectral algorithm uses the same evaluation points and produces an approximation with the same polynomial terms, but is guaranteed by construction to have zero internal aliasing. Alternatively, it is possible to select a sufficiently small polynomial basis to avoid internal aliasing, but this approach requires unnecessary conservatism that could easily be avoided with a Smolyak pseudospectral approximation.

4.3. External aliasing. The difference in external aliasing between direct quadrature and Smolyak pseudospectral approximation is much less severe. Both methods exhibit external aliasing from terms far outside the range of the approximation, as such errors are a necessary consequence of using finite-order quadrature. Since the methods are constructed from similar constituent one-dimensional quadrature rules, aliasing is of similar magnitude when it occurs.

Comparing Theorem 3.5, condition (b), and Remark 4.1, we observe that if the direct quadrature method has no external aliasing between two basis terms, the equivalent Smolyak pseudospectral algorithm will not either. Yet the two methods perform differently because of their behavior on separable functions. Condition (a) of Theorem 3.5 provides an additional condition under which external aliasing will not occur under a Smolyak pseudospectral algorithm, and thus it has strictly less external aliasing in general.

Example 4.4. If we repeat Example 4.3 but choose f to be a polynomial outside the approximation space, $f = \psi_6(x)\psi_6(y)$, we obtain the results in Figure 4.3. Now every nonzero coefficient is the result of external aliasing. Direct quadrature correctly computes some terms because of either parity or the few cases where Remark 4.1 is satisfied. However, the Smolyak approach has fewer errors because the terms not between $(0, 0)$ and $(6, 6)$ are governed by condition (a) of Theorem 3.5, and hence have no external aliasing.

This example is representative of the general case. Direct quadrature always incurs at least as much external aliasing as the Smolyak approach, and the methods become equivalent if the external term causing aliasing is of very high order. Although both methods will always exhibit external aliasing onto coefficients of the approximation for nonpolynomial inputs, the truncation can in principle be chosen to include all the important terms, so that the remaining external aliasing is acceptably small.

FIG. 4.3. Numerical results for Example 4.4; each color square indicates the log of the coefficient magnitude for the basis function at its position. The circle indicates the correct nonzero coefficient. The Smolyak pseudospectral approach has fewer terms corrupted by external aliasing in this case.

4.4. Summary of comparison. Compared to the Smolyak pseudospectral approach, direct quadrature yields larger internal *and* external aliasing errors. Because of these aliasing errors, direct quadrature is essentially unable to make efficient use of most sparse quadrature rules. The Smolyak pseudospectral approach, on the other hand, is guaranteed never to have internal aliasing if the one-dimensional pseudospectral operators are chosen according to simple guidelines. We therefore recommend against using direct quadrature. The remainder of the paper will focus on extensions of the basic Smolyak pseudospectral approach.

5. Adaptive polynomial approximations. When constructing a polynomial approximation of a black-box computational model, there are two essential questions: first, which basis terms should be included in the expansion; and second, what are the coefficients of those basis terms? The Smolyak construction allows detailed control over the truncation of the polynomial expansion and the work required to compute it. Since we typically do not have a priori information about the functional relationship generated by a black-box model, we develop an adaptive approach to tailor the Smolyak approximation to this function, following the dimension-adaptive quadrature approaches of Gerstner and Griebel [13] and Hegland [16].

The Smolyak algorithm is well suited to an adaptive approach. The telescoping sum converges in the same sense as the constituent one-dimensional operators as the index set grows to include $\mathbb{N}_0^d$, so we can simply add more terms to improve the approximation until we are satisfied. We separate our adaptive approach into two components: local improvements to the Smolyak multi-index set and a global stopping criterion.

5.1. Dimension adaptivity. Dimension adaptivity is responsible for identifying multi-indices to add to a Smolyak multi-index set in order to improve its accuracy. A standard refinement approach is simply to grow an isotropic simplex of side length n . [13] and [16] instead suggest a series of greedy refinements that customize the Smolyak algorithm to a particular problem.

The refinement used in [13] is to select a multi-index $\mathbf{k} \in \mathcal{K}$ and to add the forward neighbors of $\mathbf{k}$ that are admissible. The multi-index $\mathbf{k}$ is selected via an error indicator $\epsilon(\mathbf{k})$. We follow [13] and assume that whenever $\mathbf{k}$ contributes strongly to the result of the algorithm, it represents a subspace that likely needs further refinement.

Let $\mathbf{k}$ be a multi-index such that $\mathcal{K}' := \mathcal{K} \cup \mathbf{k}$, where $\mathcal{K}$ and $\mathcal{K}'$ are admissible multi-index sets. The triangle inequality (for some appropriate norm, see section 5.3)

bounds the change in the Smolyak approximation produced by adding $\mathbf{k}$ to $\mathcal{K}$, yielding a useful error indicator:

$$(5.1) \quad \|A(\mathcal{K}', d, \mathcal{L}) - A(\mathcal{K}, d, \mathcal{L})\| \leq \|\Delta_{k_1}^1 \otimes \cdots \otimes \Delta_{k_d}^d\| =: \epsilon(\mathbf{k}).$$

Conveniently, this error indicator does not change as $\mathcal{K}$ evolves, so we need only compute it once. At each adaptation step, we find the $\mathbf{k}$ that maximizes $\epsilon(\mathbf{k})$ and that has at least one admissible forward neighbor. Then we add those forward neighbors.

5.2. Termination criterion. Now that we have a strategy to locally improve a multi-index set, it is useful to have a global estimate of the error of the approximation, ϵ_g . We cannot expect to compute the exact error, but even a rough estimate is useful. We follow Gerstner and Griebel's choice of global error indicator

$$(5.2) \quad \epsilon_g := \sum \epsilon(\mathbf{k}),$$

where the sum is taken over all the multi-indices that are eligible for local adaptation at any particular step (i.e., that have admissible forward neighbors) [13]. The algorithm may be terminated when a particular threshold of the global indicator is reached, or when it falls by a specified amount.

5.3. Error indicators and work-considering algorithms. Thus far we have presented the adaptation strategy without reference to the problem of polynomial approximation. In this specific context, we use the $L^2(\mathbf{X}, w)$ norm in (5.1), because it corresponds to the convergence properties of pseudospectral approximation and thus seems an appropriate target for greedy refinements. This choice is a heuristic to accelerate performance—albeit one that is simple and natural, and has enjoyed success in numerical experiments (see section 6). Moreover, the analysis of external aliasing in Theorem 3.5 suggests that, in the case of pseudospectral approximation, significant missing polynomial terms alias onto some of the *included* lower-order coefficients, giving the algorithm a useful indication of which direction to refine. This behavior helps reduce the need for smoothness in the coefficient pattern. Section 6.1 provides a small fix that further helps with even or odd functions.

One is not required to use this norm to define $\epsilon(\mathbf{k})$, however, and it is possible that other choices could serve as better heuristics for some problems. Unfortunately, making definitive statements about the properties or general utility of these heuristic refinement schemes is challenging. The approach described above is intended to be broadly useful, but specific applications may require experimentation to find better choices.

Beyond the choice of norm, a commonly considered modification to $\epsilon(\mathbf{k})$ is to incorporate a notion of the computational effort required to refine $\mathbf{k}$. Define $n(\mathbf{k})$ as the amount of work to refine the admissible forward neighbors of $\mathbf{k}$, e.g., the number of new function evaluation points. [13] discusses an error indicator that provides a parameterized sliding scale between selecting the term with highest $\epsilon(\mathbf{k})$ and the lowest $n(\mathbf{k})$:

$$(5.3) \quad \epsilon_{w,1}(\mathbf{k}) = \max \left\{ w \frac{\epsilon(\mathbf{k})}{\epsilon(\mathbf{1})}, (1-w) \frac{n(\mathbf{1})}{n(\mathbf{k})} \right\}.$$

Here $w \in [0, 1]$ is the tuning parameter, and $\epsilon(\mathbf{1})$ and $n(\mathbf{1})$ are the indicator and cost of the first term. Putting $w = 0$ considers only the standard error indicator and $w = 1$ considers only the cost. A different indicator with a similar intent is

$$(5.4) \quad \epsilon_{w,2}(\mathbf{k}) = \epsilon(\mathbf{k}) - \tilde{w}n(\mathbf{k}),$$

where $\tilde{w} > 0$ describes a conversion between error and work. Both of these methods will sometimes select terms of low cost even if they do not appear to provide immediate benefit to the approximation. Yet we find both methods to be challenging to use in practice, because of the difficulty of selecting the tuning parameter. One can remove this particular tuning requirement by taking a ratio

$$(5.5) \quad \epsilon_{w,3}(\mathbf{k}) = \frac{\epsilon(\mathbf{k})}{n(\mathbf{k})}.$$

This indicator looks for “efficient” terms to refine—ones that are expected to yield greater error reduction at less cost—rather than simply the highest-error directions. We performed some numerical experiments with these methods, but none of the examples demonstrated significant improvement. Furthermore, poor choices of tuning parameters can be harmful because they can essentially make the algorithm revert to a nonadaptive form. We do not give detailed results here for those experiments because they are not particularly conclusive; some types of coefficient patterns may benefit from work-considering approaches, but this remains an open problem.

On a similar note, using ϵ_g as a termination criteria is also a heuristic. As our experiments in section 6 will show, for most smooth functions ϵ_g is an excellent estimate of the approximation accuracy. In other cases, the indicator can be quite poor; hence one should not rely on it exclusively. In practice, we typically terminate the algorithm based on a combination of elapsed wall clock time, the global error indicator, and an error estimate computed from limited ad hoc sampling.

6. Numerical experiments. Our numerical experiments focus on evaluating the performance of different quadrature rules embedded within the Smolyak pseudospectral scheme, and on evaluating performance of the adaptive Smolyak approximation strategy. Aside from the numerical examples of section 4, we do not investigate the performance of direct quadrature any further. Given our theoretical analysis of aliasing errors and the numerical demonstrations in [6], one can conclude without further demonstration that destructive internal aliasing indeed appears in practice.

This section begins by discussing practical considerations in the selection of quadrature rules. Then we evaluate convergence of Smolyak pseudospectral approximation schemes (nonadaptive and adaptive) on the Genz test functions. Next, we approximate a larger chemical kinetic system, illustrating the efficiency and accuracy of the adaptive method. Finally, we evaluate the quality of the global error indicator on all of these examples.

6.1. Selection of quadrature rules. Thus far we have sidestepped practical questions about which quadrature rules exist or are most efficient. Our analysis has relied only on polynomial accuracy of quadrature rules; all quadrature rules with a given polynomial accuracy allow the same truncation of a pseudospectral approximation. In practice, however, we care about the cumulative cost of the adaptive algorithm, which must step through successive levels of refinement.

Integration over a bounded interval with uniform weighting offers the widest variety of quadrature choices, and thus allows a thorough comparison. Table 6.1 summarizes the costs of several common quadrature schemes. First, we see that linear-growth Gaussian quadrature is asymptotically much less efficient than exponential growth in reaching any particular degree of exactness. However, for rules with fewer than about ten points, this difference is not yet significant. Second, Clenshaw–Curtis shows efficiency equivalent to exponential-growth Gaussian: both use n points to

TABLE 6.1

The cost of four quadrature strategies as their order increases: linear-growth Gauss–Legendre quadrature (Lin. G), exponential-growth Gauss–Legendre quadrature (Exp. G), Clenshaw–Curtis quadrature (C-C), and Gauss–Patterson quadrature (G-P). We list the number of points used to compute the given rule (p), the polynomial exactness (a), and the total number of points used so far (t). For nested rule, (p) = (t), so the total column is omitted.

Order	Lin. G			Exp. G			C-C		G-P	
	p	a	t	p	a	t	p	a	p	a
1	1	1	1	1	1	1	1	1	1	1
2	2	3	3	2	3	3	3	3	3	5
3	3	5	6	4	7	7	5	5	7	10
4	4	7	10	8	15	15	9	9	15	22
5	5	9	15	16	31	31	17	17	31	46
6	6	11	21	32	63	63	31	31	63	94
m	m	$2m-1$	$m^2-m/2$	2^{m-1}	2^m-1	2^m-1	$2^{m-1}+1$	$2^{m-1}+1$		

reach n th-order polynomial exactness [5]. However, their performance with respect to external aliasing differs: Clenshaw–Curtis slowly loses accuracy if the integrand is of order greater than n , while Gaussian quadrature gives $\mathcal{O}(1)$ error even on $(n+1)$ -order functions [34]. This may make Clenshaw–Curtis Smolyak pseudospectral estimates more efficient. Finally, we consider Gauss–Patterson quadrature, which is nested and has significantly higher polynomial exactness—for a given cumulative cost—than the other types [27]. Computing the quadrature points and weights in finite-precision (even extended-precision) arithmetic has practically limited Gauss–Patterson rules to 255 points, but we recommend them whenever this is sufficient.

For most other weights and intervals, there are fewer choices that provide polynomial exactness, so exponential-growth Gaussian quadrature is our default choice. In the specific case of Gaussian weight, Genz has provided a family of Kronrod extensions, similar to Gauss–Patterson quadrature, which may be a useful option [12].

If a linear growth rule is chosen and the domain is symmetric, we suggest that each new level include at least two points, so that the corresponding basis grows by at least one even and one odd basis function. This removes the possibility of unexpected effects on the adaptive strategy if the target function is actually even or odd.

6.2. Basic convergence: Genz functions. The Genz family [10, 11] comprises six parameterized functions, defined from $[-1, 1]^d \rightarrow \mathbb{R}$. They are commonly used to investigate the accuracy of quadrature rules and interpolation schemes [2, 20]. The purpose of this example is to show that different Smolyak pseudospectral strategies behave roughly as expected, as evidenced by decreasing L^2 approximation errors as more function evaluations are employed. These functions are as follows:

$$\begin{aligned}
 \text{oscillatory: } f_1(x) &= \cos \left(2\pi w_1 + \sum_{i=1}^d c_i x_i \right); \\
 \text{product peak: } f_2(x) &= \prod_{i=1}^d (c_i^{-2} + (x_i - w_i)^2)^{-1}; \\
 \text{corner peak: } f_3(x) &= \left(1 + \sum_{i=1}^d c_i x_i \right)^{-(d+1)}; \\
 \text{Gaussian: } f_4(x) &= \exp \left(- \sum_{i=1}^d c_i^2 (x_i - w_i)^2 \right);
 \end{aligned}$$

$$\begin{aligned} \text{continuous: } f_5(x) &= \exp \left(- \sum_{i=1}^d c_i^2 (|x_i - w_i|)^2 \right); \\ \text{discontinuous: } f_6(x) &= \begin{cases} 0 & \text{if } x_1 > w_1 \text{ or } x_2 > w_2, \\ \exp \left(\sum_{i=1}^d c_i x_i \right) & \text{otherwise.} \end{cases} \end{aligned}$$

Our first test uses five isotropic and *nonadaptive* pseudospectral approximation strategies. The initial strategy is the isotropic full tensor pseudospectral algorithm, based on Gauss–Legendre quadrature, with order growing exponentially with level. The other four strategies are total-order expansions of increasing order based on the following quadrature rules: linear-growth Gauss–Legendre, exponential-growth Gauss–Legendre, Clenshaw–Curtis, and Gauss–Patterson. All the rules were selected so that the final rule would have around 10^4 points.

We consider 30 random realizations of each Genz function in $d = 5$ dimensions; random parameters for the Genz functions are drawn uniformly from $[0, 1]$, then normalized so that $\|\mathbf{w}\|_1 = 1$ and $\|\mathbf{c}\|_1 = b_j$, where j indexes the Genz function type and the constants b_j are as chosen in [2, 20]. This experiment only uses the first four Genz functions, which are in C^∞ , as pseudospectral methods have well-known difficulties on functions with discontinuities or discontinuous derivatives [4]. Each estimate of L^2 approximation error is computed by Monte Carlo sampling with 10^4 samples. Figure 6.1 plots L^2 error at each stage, where each point represents the mean error over the 30 random functions.

Relatively simple conclusions can be drawn from these data. All the methods show fast convergence, indicating that the internal aliasing issues have indeed been resolved. In contrast, one would expect direct quadrature to suffer from large aliasing errors for the three superlinear growth rules. Otherwise, judging the efficiency of the different rules is not prudent, because differences in truncation and the structure of the test functions themselves obscure differences in efficiency. In deference to our adaptive strategy, we ultimately do not recommend this style of isotropic and function-independent truncation anyway.

To test our *adaptive* approach, Figure 6.2 shows results from a similar experiment, now comparing the convergence of an adaptive Smolyak pseudospectral algorithm with that of a nonadaptive algorithm. To make the functions less isotropic, we introduce an exponential decay, replacing each c_i with $c_i e^{i/5}$, where the c_i are generated and normalized as above. For consistency, both algorithms are based on Gauss–Patterson quadrature. As we cannot synchronize the number of evaluations used by the adaptive algorithm for different functions, we plot individual errors for the 30 random functions instead of the mean error. This reveals the variability in difficulty of the functions, which was hidden in the previous plot. We conclude that the adaptive algorithm also converges as expected, with performance comparable to or better than the nonadaptive algorithm. Even though we have included some anisotropy, these functions include relatively high degrees of coupling; hence, in this case the nonadaptive strategy is a fairly suitable choice. For example, the “product peak” function shows little benefit from the adaptive strategy. Although omitted here for brevity, other quadrature rules produce similar results when comparing adaptive and nonadaptive algorithms.

6.3. Adaptivity: Chemical kinetics. To further illustrate the benefits of an adaptive Smolyak approach, we build a surrogate for a realistic simulation of a com-

FIG. 6.1. Mean L^2 convergence of the nonadaptive isotropic total-order Smolyak pseudospectral algorithm with various quadrature rules, compared to the full tensor pseudospectral algorithm, on the Genz test functions.

bustion kinetics problem. Specifically, we consider the autoignition of a methane-air mixture given 14 uncertain rate parameters. Governing equations for this process are a set of stiff nonlinear ODEs expressing conservation of energy and of chemical species [19]. The uncertain rate parameters represent activation energies of reactions governing the conversion of methane to methyl, each endowed with a uniform distribution varying over $[0.8, 1.25]$ of the nominal value. These parameters appear in Arrhenius expressions for the species production rates, with the reaction pathways and their nominal rate parameters given by the GRIMech 3.0 mechanism [15]. The output of interest is the logarithm of the ignition time, which is a functional of the trajectory of the ODE system, and is continuous over the selected parameter ranges. Simulations were performed with the help of the TChem software library [28], which provides convenient evaluations of thermodynamic properties and species production rates, along with Jacobians for implicit time integration.

Chemical kinetics are an excellent testbed for adaptive approximation because, by the nature of detailed kinetic systems, we expect strong coupling between some inputs and weak coupling between others, but we cannot predict these couplings a priori. We test the effectiveness of adaptive Smolyak pseudospectral methods based on the four quadrature rules discussed earlier. As our earlier analysis suggested that Gauss–Patterson quadrature should be most efficient, our reference solution is a non-

FIG. 6.2. L^2 convergence of the adaptive and nonadaptive Gauss-Patterson Smolyak pseudospectral algorithm. Individual results for 30 random instances of the Genz functions are shown.

adaptive Gauss-Patterson total-order Smolyak pseudospectral expansion. We ran the nonadaptive algorithm with a total-order index set truncated at $n = 5$ (which includes monomial basis terms up through $\psi_{23}^{(i)}$), using around 40,000 point evaluations and taking over an hour to run. We tuned the four adaptive algorithms to terminate with approximately the same number of evaluations.

Figure 6.3 compares convergence of the five algorithms. The L^2 errors reported on the vertical axis are Monte Carlo estimates using 10^4 points. Except for a small deviation at fewer than 200 model evaluations, all of the adaptive methods significantly outperform the nonadaptive method. The performance of the different quadrature rules is essentially as predicted in section 6.1: Gauss-Patterson is the most efficient, exponential-growth Gauss-Legendre and Clenshaw-Curtis are nearly equivalent, and linear-growth Gauss-Legendre performs worse as the order of the polynomial approximation increases. Compared to the nonadaptive algorithm, adaptive Gauss-Patterson yields more than two orders of magnitude reduction in the error at the same number of model evaluations. Linear growth Gaussian quadrature is initially comparable to exponential growth Gaussian quadrature, because the asymptotic benefits of exponential growth do not appear while the algorithm is principally using very small one-dimensional quadrature rules. At the end of these experiments, a reasonable number of higher order quadrature rules are used and the difference becomes visible.

FIG. 6.3. L^2 convergence of ignition delay in a 14-dimensional chemical kinetic system; comparing a nonadaptive isotropic total-order Gauss–Patterson-based Smolyak pseudospectral algorithm to the adaptive algorithm with various quadrature rules.

We conclude by illustrating that the adaptive algorithm is effective because it successfully focuses its efforts on high-magnitude coefficients—that is, coefficients that make the most significant contributions to the function. Even though the nonadaptive expansion has around 37,000 terms and the final adaptive Gauss–Patterson expansion only has about 32,000 terms, the adaptive expansion exhibits much lower error because most of the additional terms in the nonadaptive expansion are nearly zero. By skipping many near-zero coefficients, the adaptive approach is able to locate and estimate a number of higher-order terms with large magnitudes. Figure 6.4 depicts this pattern by plotting the difference between the numbers of included terms in the final adaptive Gauss–Patterson and nonadaptive expansions. The adaptive algorithm does not actually add any higher-order monomials; neither uses one-dimensional basis terms of order higher than $\psi_{23}^{(i)}$. Instead, the adaptive algorithm adds mixed terms of higher total order, thus capturing the coupling of certain variables in more detail than the nonadaptive algorithm. The figure shows that terms through 30th order are included in the adaptive expansion, all of which are products of nonconstant polynomials in more than one dimension.

6.4. Performance of the global error indicator. To evaluate the termination criterion, we collected the global error indicator during runs of the adaptive algorithm for all of the test functions described above, including the slowly converging nonsmooth Genz functions omitted before. The discontinuous Genz function does not include the exponential coefficient decay because the discontinuity already creates strong anisotropy. Results are shown for Gauss–Patterson quadrature. The relationship between the estimated L^2 error and the global error indicator ϵ_g is shown in Figure 6.5. For the smooth test functions, ϵ_g is actually an excellent indicator, as it is largely within an order of magnitude of the correct value and essentially linearly related to it. However, the nonsmooth Genz functions illustrate the hazard of relying too heavily on this indicator: although the adaptive algorithm does decrease both the errors and the indicator, the relationship between the two appears far less direct.

FIG. 6.4. The plot depicts the difference between the number of coefficients of a particular magnitude and order in the final adaptive and nonadaptive Gauss-Patterson-based expansions. The horizontal axis is the order of the term and the vertical axis specifies $\log_{10}$ of the coefficient value. The color represents $\log_{10}$ of the difference between the two methods, where positive values indicate more terms in the nonadaptive expansion. Hence, the dark blue at $(6, -10)$ indicates that the nonadaptive expansion includes around 3,000 extra terms of magnitude 10^{-10} and the dark red at $(10, -8)$ indicates that the adaptive expansion includes about 1,000 extra terms of magnitude 10^{-8} . Grey squares are the same for both expansions and white squares are not present in either.

FIG. 6.5. Relationship between the termination criterion (5.2) and the estimated L^2 error for every function tested.

7. Conclusions. This paper gives a rigorous development of Smolyak pseudospectral algorithms, a practical approach for constructing polynomial chaos expansions from point evaluations of a function. A common alternative approach, direct quadrature, has previously been shown to suffer from large errors. We explain

these errors as a consequence of internal aliasing and delineate the exact circumstances, derived from properties of the chosen polynomial basis and quadrature rules, under which internal aliasing will occur. Internal aliasing is a problem inherent to direct quadrature approaches, which use a single (sparse) quadrature rule to compute a set of spectral coefficients. These approaches fail because they substitute a numerical approximation for only a portion of the algorithm, i.e., the evaluation of integrals, without considering the impact of this approximation on the entire construction. For almost all sparse quadrature rules, internal aliasing errors may be overcome only through an inefficient use of function evaluations. In contrast, the Smolyak pseudospectral algorithm computes spectral coefficients by assembling tensor product pseudospectral approximations in a coherent fashion that avoids internal aliasing by construction; moreover, it has smaller external aliasing errors. To establish these properties, we extend the known result that the exact set of a Smolyak pseudospectral approximation contains a union of the exact sets of all its constituent tensor product approximation operators to the case of arbitrary admissible Smolyak multi-index sets. These results are applicable to any choice of quadrature rule and generalized sparse grid, and are verified through numerical demonstrations; hence, we suggest that the Smolyak pseudospectral algorithm is a superior approach in almost all contexts.

A key strength of Smolyak algorithms is that they are highly customizable through the choice of admissible multi-index sets. To this end, we describe a simple alteration to the adaptive sparse quadrature approaches of [13, 16], creating a corresponding method for adaptive pseudospectral approximation. Numerical experiments then evaluate the performance of different quadrature rules and of adaptive versus non-adaptive pseudospectral approximation. Tests of the adaptive method on a realistic chemical kinetics problem show multiple order-of-magnitude gains in accuracy over a nonadaptive approach. Although the adaptive strategy will not improve approximation performance for every function, we have little evidence that it is ever harmful and hence widely recommend its use.

While the adaptive approach illustrated here is deliberately simple, many extensions are possible. For instance, as described in section 5.3, measures of computational cost may be added to the dimension refinement criterion. One could also use the gradient of the L^2 error indicator to identify optimal directions in the space of multi-indices along which to continue refinement, or to avoid adding all the forward neighbors of the multi-index selected for refinement. These and other developments will be pursued in future work.

A flexible open-source C++ code implementing the adaptive approximation method discussed in this paper is available at <https://bitbucket.org/mituq/muq/>.

Acknowledgments. The authors would like to thank Paul Constantine for helpful discussions and for sharing a preprint of his paper, which inspired this work. We would also like to thank Omar Knio and Justin Winokur for many helpful discussions, Tom Coles for help with the chemical kinetics example, and the reviewers for their thoughtful comments.

REFERENCES

- [1] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., 45 (2007), pp. 1005–1034.
- [2] V. BARTHELMANN, E. NOVAK, AND K. RITTER, *High dimensional polynomial interpolation on sparse grids*, Adv. Comput. Math., 12 (2000), pp. 273–288.

- [3] J. P. BOYD, *Chebyshev and Fourier Spectral Methods*, 2nd rev. ed., Dover, Mineola, NY, 2001.
- [4] C. CANUTO, M. Y. HUSSAINI, A. QUARTERONI, AND T. A. ZHANG, JR., *Spectral Methods: Fundamentals in Single Domains*, Springer, Berlin, 2006.
- [5] C. CLENSHAW AND A. CURTIS, *A method for numerical integration on an automatic computer*, Numer. Math., 2 (1960), pp. 197–205.
- [6] P. G. CONSTANTINE, M. S. ELDRED, AND E. T. PHIPPS, *Sparse pseudospectral approximation method*, Comput. Methods Appl. Mech. Engrg., 229–232 (2012), pp. 1–12.
- [7] M. S. ELDRED AND J. BURKARDT, *Comparison of non-intrusive polynomial chaos and stochastic collocation methods for uncertainty quantification*, in 47th AIAA Aerospace Sciences Meeting, Orlando, FL, 2009, AIAA paper 2009-0976, (2009).
- [8] O. G. ERNST, A. MUGLER, H.-J. STARKLOFF, AND E. ULLMANN, *On the convergence of generalized polynomial chaos expansions*, ESAIM: Math. Model. Numer. Anal., 46 (2012), pp. 317–339.
- [9] W. GANDER, *Change of basis in polynomial interpolation*, Numer. Linear Algebra Appl., 12 (2005), pp. 769–778.
- [10] A. GENZ, *Testing Multidimensional integration routines*, in Tools, Methods, and Languages for Scientific and Engineering Computation, B. Ford, J. Rault, and F. Thomasset, eds., North-Holland, Amsterdam, 1984, pp. 81–94.
- [11] A. GENZ, *A package for testing multiple integration subroutines*, in Numerical Integration: Recent Developments, Software and Applications, P. Keast and G. Fairweather, eds., D Reidel, Dordrecht, The Netherlands, 1987, pp. 337–340.
- [12] A. GENZ, *Fully symmetric interpolatory rules for multiple integrals over infinite regions with Gaussian weight*, J. Comput. Appl. Math., 71 (1996), pp. 299–309.
- [13] T. GERSTNER AND M. GRIEBEL, *Dimension-adaptive tensor-product quadrature*, Computing, 71 (2003), pp. 65–87.
- [14] R. GHANEM AND P. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, 1991.
- [15] Z. Q. GREGORY P. SMITH, D. M. GOLDEN, M. FRENKLACH, N. W. MORIARTY, B. EITENEER, M. GOLDENBERG, C. T. BOWMAN, R. K. HANSON, S. SONG, W. C. GARDINER, JR., V. V. LISSIANSKI, *GRI-MECH 3.0*, http://www.me.berkeley.edu/gri_mech.
- [16] M. HEGLAND, *Adaptive sparse grids*, ANZIAM J, 44 (2003), pp. 335–353.
- [17] J. S. HESTHAVEN, S. GOTTLIEB, AND D. GOTTLIEB, *Spectral Methods for Time-Dependent Problems*, Cambridge University Press, Cambridge, 2007.
- [18] X. HUAN AND Y. M. MARZOUK, *Simulation-based optimal Bayesian experimental design for nonlinear systems*, J. Comput. Phys., 232 (2013), pp. 288–317.
- [19] R. J. KEE, M. E. COLTRIN, AND P. GLARBORG, *Chemically Reacting Flow: Theory and Practice*, Wiley-Interscience, Hoboken, NJ, 2003.
- [20] A. KLIMKE AND B. WOHLMUTH, *Algorithm 847: Spinterp: Piecewise multilinear hierarchical sparse grid interpolation in MATLAB*, ACM Trans. Math. Software, 31 (2005), pp. 561–579.
- [21] O. P. LE MAITRE AND O. M. KNIO, *Spectral Methods for Uncertainty Quantification*, Springer, Dordrecht, The Netherlands, 2010.
- [22] A. MARCH AND K. WILLCOX, *Constrained multifidelity optimization using model calibration*, Struct. Multidisciplin. Optim., 46 (2012), pp. 93–109.
- [23] Y. M. MARZOUK, H. N. NAJM, AND L. A. RAHN, *Stochastic spectral methods for efficient Bayesian solution of inverse problems*, J. Comput. Phys., 224 (2007), pp. 560–586.
- [24] F. NOBILE, R. TEMPONE, AND C. G. WEBSTER, *A sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM J. Numer. Anal., 46 (2007), p. 2309–2345.
- [25] E. NOVAK AND K. RITTER, *High dimensional integration of smooth functions over cubes*, Numer. Math., 75 (1996), pp. 79–97.
- [26] E. NOVAK AND K. RITTER, *Simple cubature formulas with high polynomial exactness*, Constr. Approx., 15 (1999), pp. 499–522.
- [27] T. N. L. PATTERSON, *The optimum addition of points to quadrature formulae*, Math. Comp., 22 (1968), pp. 847–856.
- [28] C. SAFTA, H. NAJM, AND O. KNIO, *TChem*, <http://www.sandia.gov/tchem/>.
- [29] W. SICKEL AND T. ULLRICH, *The Smolyak algorithm, sampling on sparse grids and function spaces of dominating mixed smoothness*, East J. Approx., 13 (2007), pp. 387–425.
- [30] W. SICKEL AND T. ULLRICH, *Tensor products of Sobolev-Besov spaces and applications to approximation from the hyperbolic cross*, J. Approx. Theory, 161 (2009), pp. 748–786.
- [31] S. A. SMOLYAK, *Quadrature and interpolation formulas for tensor products of certain classes of functions*, Dokl. Akad. Nauk SSSR, 4 (1963), pp. 240–243.

- [32] C. SOIZE AND R. GHANEM, *Physical systems with random uncertainties: Chaos representations with arbitrary probability measure*, SIAM J. Sci. Comput., 26 (2004), pp. 395–410.
- [33] B. SUDRET, *Global sensitivity analysis using polynomial chaos expansions*, Reliability Eng. System Safety, 93 (2008), pp. 964–979.
- [34] L. N. TREFETHEN, *Is Gauss quadrature better than Clenshaw–Curtis?*, SIAM Rev., 50 (2008), pp. 67–87.
- [35] G. W. WASILKOWSKI AND H. WOZNAKOWSKI, *Explicit cost bounds of algorithms for multivariate tensor product problems*, J. Complexity, 11 (1995), pp. 1–56.
- [36] G. W. WASILKOWSKI AND H. WOZNAKOWSKI, *Weighted tensor product algorithms for linear multivariate problems*, J. Complexity, 15 (1999), pp. 402–447.
- [37] G. W. WASILKOWSKI AND H. WOZNAKOWSKI, *Polynomial-time algorithms for multivariate linear problems with finite-order weights: Worst case setting*, Found. Comput. Math., 5 (2005), pp. 451–491.
- [38] N. WIENER, *The homogeneous chaos*, Amer. J. Math., 60 (1938), pp. 897–936.
- [39] D. XIU AND J. S. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., 27 (2005), pp. 1118–1139.
- [40] D. XIU AND G. E. KARNIADAKIS, *The Wiener–Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., 24 (2002), pp. 619–644.
- [41] D. XIU, *Efficient collocational approach for parametric uncertainty analysis*, Comm. Comput. Phys., 2 (2007), pp. 293–309.
- [42] D. XIU, *Fast numerical methods for stochastic computations: A review*, Comm. Comput. Phys., 5 (2009), pp. 242–272.
- [43] D. XIU, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, Princeton, NJ, 2010.