Slow Growth for Gauss Legendre Sparse Grids

John Burkardt, Clayton Webster

April 4, 2014

Abstract

A sparse grid for multidimensional quadrature can be constructed from products of 1D rules. For multidimensional quadrature, it is highly desirable to construct families of sparse grids whose point growth is controlled. The properties of the sparse grid are determined by the indexed family of 1D rules that are selected. In this regard, both the precision and nestedness of the 1D rules can be crucial. Thus, two frequent bases for sparse grids are the nested Clenshaw Curtis family, and the high precision Gauss Legendre family. In many cases, sparse grids of comparable precision and comparable point count may be formed from either family. However, it is possible to slightly modify the usual Gauss Legendre family in such a way that a significant reduction in point count occurs, resulting in sparse grids that are noticeably more efficient.

1 Introduction

The multidimensional quadrature problem requires the estimation, over a region Ω, of the integral of a function $f(x)$ whose argument is an m-vector.

If Ω is a product region, a suitable quadrature rule can be developed as a product of a sequence of 1D quadrature rules. The resulting rule is easy to define, evaluate, and analyze, but requires an excessive number of function evaluations for relatively modest precision requirements.

Smolyak's sparse grid procedure [4] can produce rules that "replace" the corresponding product rule with a weighted sum of lower-order product rules. The sparse grid can match the precision of the straightforward product rule, while using far fewer points.

Implementations of the Smolyak procedure often use product rules derived from a family of 1D Clenshaw Curtis rules, which are popular because their nested structure tends to reduce the point count of the sparse grids. An alternative choice is based on a family of Gauss rules, such as the Gauss Legendre rules defined over the interval $[-1, +1]$ with weight function $w(x)$. Unlike Clenshaw Curtis rules, Gauss rules have little or no nesting, but do have about double the precision. Even for 1D quadrature, there continue to be disputes about the relative merits of Clenshaw Curtis and Gauss Legendre rules [5].

If an indexed family of Gauss rules is to be selected for use in a sparse grid, each rule can have about half the order of the corresponding Clenshaw Curtis rule. Since the sparse grid construction process is somewhat elaborate, it is thus not necessarily obvious whether a sparse grid of given index and spatial dimension will have more points using the Clenshaw Curtis or Gauss Legendre family as its basis.

While it might seem that no further improvement can be made to the Gauss Legendre family, it turns out that there is a simple modification which can be made at almost no cost, which preserves the precision of the resulting sparse grids while significantly reducing the point count.

We will present this modification, explain its properties, and compare the size and precision of sparse grids created with and without this change.

2   Quadrature Rules for the 1D Problem

Supposing that Ω is some fixed interval in $\mathbb{R}$, and $f(x)$ is a real-valued function defined on Ω, the quadrature problem seeks an estimate $Q(f)$ for the value of the integral $I(f) = \int_\Omega f(x)\,dx$.

One approach to this problem constructs a *quadrature rule*, which is a pair of n points x_i and weights w_i, from which an approximate integral value is computed:

$$Q(f) = \sum_{i=1}^n w_i f(x_i) \approx I(f) = \int_\Omega f(x)\,dx$$

A quadrature rule is said to have polynomial precision p if, whenever the integrand $f(x)$ is a polynomial of degree p or less, the integral estimate is exact.

A standard procedure for the integration problem is to produce an indexed family $\mathcal{Q}$ of quadrature rules, with a typical element Q^i, whose precisions p_i form an increasing sequence. For a given integrand f, the behavior of the sequence of quadrature estimates $Q^i(f)$ is a practical measure of convergence.

We will assume the quadrature region Ω is actually the closed interval $[-1, +1]$. Using a simple linear mapping, a quadrature problem posed on any general closed interval $[a, b]$ can be reduced to this case.

The *Gauss Legendre* family is not an interpolatory rule. Instead of allowing the abscissas to be freely chosen, they are regarded as additional degrees of freedom. The Gauss Legendre quadrature rule of order n_i will have precision $2 * n_i - 1$, rather than the precision of $n_i - 1$ expected for a standard interpolatory rule.

The Gauss-Legendre family does not include any subsequence of nested rules. Because of this, if we wish to use this family to form a sparse grid quadrature rule, we will not enjoy the reduction in point count associated with the Clenshaw Curtis rules. On the other hand, the higher accuracy of the Gauss-Legendre family will mean that we can economize in other ways.

3   Product Rules For High Dimensional Quadrature

We can develop tools for the multidimensional quadrature problem out of those designed for the 1D case, if the integration region Ω can be represented as a product of 1D regions $\Omega = \bigotimes_{i=1}^m \omega_i$. In that case, we can form a quadrature rule $Q = (X, W)$ for Ω by selecting quadrature rules $q_i = (x_i, w_i)$ for each factor space ω_i and combining them in a standard way to form a product quadrature rule.

If the order of factor q_i is n_i, then the resulting product rule will use $N = \prod_{i=1}^m n_i$ points.

The same question of precision arises for any multidimensional quadrature rule. Such a rule is said to have precision P if its integral estimate is exact whenever the integrand is a polynomial of *total degree p* or less. If the precision of factor q_i is p_i, then the resulting product rule Q has precision $P = \min\{p_i\}$.

In a given dimension m, it is natural to consider assembling a family of product quadrature rules of increasing precision, and in fact the product rule approach allows this to be done in an orderly and analyzable procedure. If we use the same factor rule in all dimensions, then it is easy to see that our product rule family will have orders $N = \{n_0^m, n_1^m, n_2^m, ...\}$. Thus, for example, if we assume that in 1D a 2 point rule has precision 1, then to construct the corresponding product rule of precision 1 in 10, 20 or 30 dimensions will required about a thousand, million, or a billion points. This phenomenal increase in point count shows that the straightforward product rule approach rapidly exceeds any computational budget, given even moderate dimensionality and precision requests.

Multidimensional quadrature problems can always be handled by a Monte Carlo approach, but the drawback there is that the convergence is much slower than that observed for approaches based on polynomial interpolation. There is a serious need for quadrature techniques that can produce a sequence of polynomially precise rules whose point growth is not exponential.

4   The Sparse Grid Construction Rule

We suppose we are interested in quadrature rules for integrands defined over an m-dimensional region.

We assume that we have an indexed family of 1D quadrature rules q^i. These quadrature rules are indexed by the variable i, which we will call the *level* of the 1D rule. The index i begins at 1, and the corresponding first 1D quadrature rule is often taken to be the midpoint rule. It is understood that as the index i increases, so do both n_i, the order or number of points, and p_i, the polynomial accuracy of the rule; the exact nature of these relationships is not necessary to specify yet; it is enough to understand that accuracy increases with i.

We can specify a product rule for an m-dimensional space by creating a *level vector*, that is, an m-vector $\mathbf{i}$, each of whose entries is the 1D level which indexes the rule to be used in that spatial coordinate. The resulting product rule may be written as $q^{i_1} \otimes \cdots \otimes q^{i_m}$. We say that this product rule has a *product level* of $|\mathbf{i}| = \sum_{k=1}^{m} i_k$. Since the lowest value of a 1D level is taken to be 1, the lowest value of $|\mathbf{i}|$ is m. Thus, the first product rule in m dimensions has a product level of m, and is formed by m factors of the midpoint rule, each having 1D level of 1.

A product rule is a method for approximating an integral; a sparse grid is a refined method for approximating integrals that uses weighted combinations of product rules. A sparse grid rule can be indexed by a variable called the *sparse grid level*, here symbolized by ℓ. The lowest value of ℓ is taken to be m. The sparse grid is formed by weighted combinations of those product rules whose product level $|\mathbf{i}|$ falls between $\ell - m + 1$ and ℓ.

The formula for the sparse grid has the form:

$$\mathcal{A}(\ell, m) = \sum_{\ell-m+1 \leq |\mathbf{i}| \leq \ell} (-1)^{\ell - |\mathbf{i}|} \binom{m-1}{\ell - |\mathbf{i}|} (q^{i_1} \otimes \cdots \otimes q^{i_m})$$

The condition under the summation sign can be rewritten as

$$0 \leq \ell - |\mathbf{i}| \leq m - 1$$

so we select just those $\mathbf{i}$ for which the combinatorial coefficient makes sense.

The first sparse grid, of sparse grid level $\ell = m$, will be formed of all product grids with product levels between 1 and m. The lowest possible product level is actually m, and is attained only by the m-dimensional product of the midpoint rule (or whatever the first 1D quadrature rule is). So this one-point rule will be the first sparse grid in the series.

5   Precision of a Sparse Grid

The Smolyak construction has shown how to build a sparse grid rule $\mathcal{A}(\ell, m)$ as a weighted sum of particular product rules, which in turn are formed from 1D quadrature rules. It is not immediately clear from the construction whether or not the resulting sparse grid rule actually has any utility as a quadrature rule. More precisely, if we know the precisions of the 1D quadrature rules, what can we then say about the precision of $\mathcal{A}(\ell, m)$?

In [2], Novak and Ritter considered this question for the special case in which the Clenshaw Curtis family was used as the factors in every dimension. The particular family chosen had the exponential growth behavior, that is $n_1 = 1$, and $n_i = 2^{i-1} + 1$ for $i > 2$. A formula was derived for the precision of the resulting sparse grid rules:

$$P(\mathcal{A}(\ell, m)) = 2\ell - 1$$

and it was further shown that this relationship would hold for any sparse grid whose 1D factor rules satisfied the corresponding condition

$$p(q^l) \geq 2l - 1, \ l = 1, 2, \ldots$$

This result suggests a natural procedure for developing an indexed family of sparse grids of increasing polynomial precision.

6 "GL": Gauss Legendre Linear Growth Family

For any order i, we can define a Gauss Legendre rule over the interval [-1,1] with weight function $w(x) = 1$. In order to construct sparse grids, we will select a sequence of Gauss Legendre rules for our indexed 1D family, which we designate GL. The precision of a Gauss Legendre rule of order i is $2i - 1$. This high precision means that we can choose our 1D family to be the full set of Gauss Legendre rules, indexed by their order. Thus GL^l will denote the Gauss Legendre rule of order l, for $l = 1, 2, \ldots$.

The rules in the GL family just satisfy the sparse grid precision requirement; since a quadrature rule of precision $2i - 1$ must use at least i points, each GL rule has the smallest possible order for the desired precision. This has great benefits in reducing the number of points in resulting sparse grids. However, other families, such as the Clenshaw Curtis family CC, take lower precision rules that have a nested structure, thus offering an alternative method of point reduction. In many cases, a CC sparse grid will use fewer points than a GL sparse grid of the same precision.

7 "GLS": Gauss Legendre Slow Growth Family

Since a sparse grid is the logical sum of many product grids, there is an advantage to using 1D rules that are nested, or which have points in common, because this means that the resulting product rules will also have repeated points, which can be merged when the sparse grid is formed, thus reducing the number of function evaluations required.

We note that the rules in the GL family have no nesting except for the fact that the odd rules all include the abscissa 0. Since no rule can beat the individual GL rules on precision, we consider whether we can improve the GL family's nesting behavior to make it more competitive.

Consider what happens as we ascend through the GL rules of index 1 through 5. During this sequence, we encounter $1 + 2 + 2 + 4 + 4 = 13$ distinct abscissas. Consider a modified family, designated GLS, which only differs from the GL family by replacing every even index rule by a copy of the next higher odd order rule, so that the first five GLS rules would be $\{GL^1, GL^3, GL^3, GL^5, GL^5\}$. Then if we consider advancing through the GLS rules of indices 1 through 5, we encounter $1 + 2 + 0 + 4 + 0 = 7$ distinct abscissas. Thus, we have gained the advantage of a slightly more precise rule on even indices, along with a mild form of nesting.

The sparse grid construction does not require that the rules at distinct levels must be different, nor that they have a particular precision. The precision requirement only demands that the rule at index l have *at least* precision $2 * l - 1$. Thus, our approach does not violate any of the sparse grid conventions. A similar approach was considered by Petras for the case of a Gauss Kronrod family [3].

But whether there is a significant advantage to this method depends on how much this reduction in 1D rule point counts is reflected in the point count of the final sparse grid rule. This effect, of course, depends on the spatial dimension and the sparse grid level.

8 Comparison of Sparse Grid Point Counts

To judge the advantages of the GLS family over the GL family when constructing sparse grids, we compute the point counts for each family, over a range of sparse grid levels and spatial dimensions, and a table of the point count ratios.

We note several patterns in the ratio table:

- Except for dimension 1, the ratio is never less than 1;
- For a fixed dimension, the ratio increases with the level l;
- For a fixed level, the ratio tends to rise at first and then decrease.

M:	1	2	3	4	5	10
L = 1	1	1	1	1	1	1
2	2	5	7	9	11	21
3	3	13	25	41	61	221
4	4	29	69	137	241	1,581
5	5	53	165	385	781	8,761
6	6	89	351	953	2,203	40,405
7	7	137	681	2,145	5,593	162,025
8	8	201	1,233	4,481	13,073	581,385
9	9	281	2,097	8,785	28,553	1,904,465

Table 1: Sparse Grid Point Counts for GL Family

M:	1	2	3	4	5	10
L = 1	1	1	1	1	1	1
2	3	5	7	9	11	21
3	3	9	19	33	51	201
4	5	17	39	81	151	1,201
5	5	33	87	193	391	5,281
6	7	45	153	409	933	19,165
7	7	81	273	777	1,973	61,285
8	9	97	465	1,481	4,013	177,525
9	9	161	705	2,537	7,693	474,885

Table 2: Sparse Grid Points Counts for GLS Family

M:	1	2	3	4	5	10
L = 1	1.00	1.00	1.00	1.00	1.00	1.00
2	0.67	1.00	1.00	1.00	1.00	1.00
3	1.00	1.44	1.32	1.24	1.20	1.10
4	0.80	1.70	1.77	1.69	1.60	1.32
5	1.00	1.61	1.90	1.99	2.00	1.66
6	0.86	1.98	2.29	2.33	2.36	2.11
7	1.00	1.69	2.49	2.76	2.83	2.64
8	0.89	2.07	2.65	3.03	3.26	3.28
9	1.00	1.75	2.97	3.46	3.71	4.01

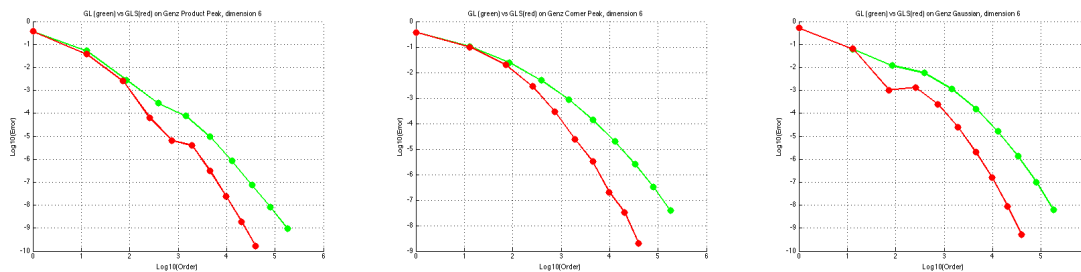Table 3: Sparse Grid Points Count Ratio GL/GLS Families

Figure 1: **GL (green) vs GLS (red) on Genz Functions**
Genz product peak, corner peak, and gaussian integrands, M = 6.

9 Quadrature Tests

Genz[1] prescribed a number of test integrands for multidimensional quadrature. In order to demonstrate the relative efficiency of the GLS family for sparse grids, we used it and the GL family to estimate the integrals of several of the Genz tests, going from sparse grid level 1 to 10. Figure [1] displays the results of these computations.

10 Conclusion

The high precision of Gauss quadrature rules (whether for Legendre, Laguerre, or Hermite weight functions) has made them attractive as a choice for the 1D quadrature family used in constructing sparse grids. Since Gauss rules generally have very little or no nesting structure, however, they often produce sparse grids with point counts that exceed those of less precise, but nested, rules, such as the Clenshaw Curtis family.

   We have shown a simple adjustment to the usual Gauss-Legendre 1D family GL which slightly improves the precision and significantly improves the nesting properties. When the standard sparse grid construction process is carried out using the improved GLS family, the point count for almost all sparse grids decreased noticeably, in some cases to a third or a fourth of the GL count. Sparse grids constructed with the GLS family were shown to have the same precision as those using the GL family; meanwhile, the approximation of some simple test integrands showed a clear efficiency advantage when using the GLS family.

References

[1] ALAN GENZ, A package for testing multiple integration subroutines, in Numerical Integration: Recent Developments, Software and Applications, edited by Patrick Keast, Graeme Fairweather, Reidel, 1987, pages 337-340.

[2] ERICH NOVAK, KLAUS RITTER, High dimensional integration of smooth functions over cubes, Numerische Mathematik, Volume 75, Number 1, November 1996, pages 79-97.

[3] KNUT PETRAS, Smolyak cubature of given polynomial degree with few nodes for increasing dimension, Numerische Mathematik, Volume 93, Number 4, February 2003, pages 729-753.

[4] SERGEY SMOLYAK, Quadrature and interpolation formulas for tensor products of certain classes of functions, Doklady Akademii Nauk SSSR, Volume 4, 1963, pages 240-243.

[5] LLOYD TREFETHEN, Is Gauss quadrature better than Clenshaw-Curtis?, SIAM Review, Volume 50, Number 1, March 2008, pages 67-87.