

SLAC-PUB-1839 (Revised)
November 1976
Revised April 1977
(T/E)

A NEW ALGORITHM FOR ADAPTIVE MULTIDIMENSIONAL INTEGRATION*

G. Peter Lepage

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

ABSTRACT

A new general purpose algorithm for multidimensional integration is described. It is an iterative and adaptive Monte Carlo scheme. The new algorithm is compared with several others currently in use, and shown to be considerably more efficient than all of these for a number of sample integrals of high dimension ($n \gtrsim 4$).

(Submitted to J. Comput. Phys.)

*Work supported by the Energy Research and Development Administration.

I. INTRODUCTION

One of the most important computational problems facing physicists today is the evaluation of multidimensional integrals with complicated and sometimes very poorly behaved integrands. To cite but a single example drawn from elementary particle theory, integrals over four or more variables are routinely encountered in the computation of scattering amplitudes using Feynman perturbation theory [1]. In this note a new algorithm for multidimensional integration is described. It employs an iterative and adaptive Monte Carlo scheme.

Among the characteristics of this algorithm are:

- a. A reliable error estimate for the integral is readily computed.
- b. The integrand need not be continuous for the algorithm to function and, in particular, step functions pose no problem. Thus integration over hypervolumes of irregular shape is straightforward.
- c. The convergence rate is independent of the dimension of the integral.
- d. The algorithm is adaptive. It automatically concentrates evaluations of the integrand in those regions where the integrand is largest in magnitude.

Characteristics (a) - (c) are common to all Monte Carlo methods [2-4]. Characteristic (d) is the single most important feature of this algorithm. The major problem in multidimensional integration is the exponential growth with increasing dimension of the integration volume over which the integrand must be sampled. For example, applying any 5-point integration rule (e.g., Gauss-Legendre integration) along each coordinate axis of a 9-dimensional integral

requires $5^9 \sim 2 \times 10^6$ evaluations of the integrand—a formidable number for all but the simplest integrands. Thus any general purpose integration algorithm for use in high dimensions should be adaptive.

In Section II we briefly review the general features of Monte Carlo integration [2-4] and previous work using this approach. In Section III we describe the new algorithm for integration over a single variable. The method is generalized for use in n dimensions in Section IV. Finally in Section V, we compare the performance of this algorithm with a number of others currently in use. It is shown to be considerably more efficient for several sample integrals of high dimension ($n \gtrsim 4$). In the Appendix we discuss some simple modifications of the basic approach which may improve it in particular applications.

II. MONTE CARLO INTEGRATION

Consider the integral of a function of n variables $\vec{x} = (x_1, \dots, x_n)$ over a volume Ω :

$$I = \int_{\Omega} d^n \vec{x} f(\vec{x})$$

If M points \vec{x} are randomly selected from a distribution of points in Ω with density $p(\vec{x})$, it is easily shown that the integral is approximated by

$$S^{(1)} = \frac{\Omega}{M} \sum_{\vec{x}} \frac{f(\vec{x})}{p(\vec{x})} \tag{1}$$

$\rightarrow I \quad \text{as } M \rightarrow \infty$

where the probability density function is normalized to unity:

$$\int_{\Omega} d\vec{x} p(\vec{x}) = 1$$

The quantity $S^{(1)}$ is expected to fluctuate about the true value of the integral as different sets of M random points are chosen. The variance of this fluctuation is given by

$$\sigma^2 = \frac{\int_{\Omega} d\vec{x} \frac{f^2(\vec{x})}{p(\vec{x})} - \left[\int_{\Omega} d\vec{x} f(\vec{x}) \right]^2}{M}$$

For M large, this quantity is approximated by

$$\sigma^2 \approx \frac{S^{(2)} - (S^{(1)})^2}{M-1} \quad (2)$$

where

$$S^{(2)} = \frac{\Omega^2}{M} \sum_{\vec{x}} \left(\frac{f(\vec{x})}{p(\vec{x})} \right)^2$$

The standard deviation σ indicates the accuracy of $S^{(1)}$ as an estimate of I . Note that reliable estimates of the variance are possible only if the integral

$$\int_{\Omega} d\vec{x} \frac{f^2(\vec{x})}{p(\vec{x})}$$

is finite, though even if it is not, $S^{(1)}$ may still approximate I on the average.

There are a number of techniques used to reduce the variance σ^2 for fixed M . Two of the more popular methods used in multidimensional integration are:

- a. Importance Sampling: Here the density $p(\vec{x})$ is varied to reduce σ^2 . As is well known, the optimal variance is obtained when

$$p(\vec{x}) = \frac{|f(\vec{x})|}{\int_{\Omega} d\vec{x} |f(\vec{x})|} \quad (3)$$

Thus when using importance sampling function evaluations are concentrated where the integrand is largest in magnitude (whether or not it is flat there).

- b. Stratified Sampling: To reduce the variance, the integration volume can be subdivided into N smaller volumes of varying sizes. Then a Monte Carlo integration is performed in each subvolume using M/N random points. The variance is varied by changing the relative sizes and locations of the subvolumes and is minimized when the contributions to σ^2 from each subvolume are identical ($=\sigma^2/N$). Thus when using stratified sampling function evaluations are concentrated where the potential error is largest—i. e. , where the integral is both large and rapidly changing.

These and other methods of variance reduction appear to be inappropriate for general purpose integration algorithms as they require detailed knowledge of the integrand's behavior prior to implementation. However Sheppey [5] has devised an iterative algorithm which uses information generated about the integrand during a Monte Carlo integration to reduce the variance in subsequent integrations. He employs stratified sampling. Initially the algorithm divides the n dimensional integration volume (a hypercube) into N^n identical hypercubes using a uniform rectangular grid. A 2-point Monte Carlo integration is performed in each hypercube generating a contribution to the total integral and to the variance. The variances from the hypercubes are then used to define new grid spacings along each axis for use in the next iteration, keeping the total number of hypercubes constant. Thus over a number of iterations hypercubes can be gradually concentrated where the variance was initially largest, and the variance reduced.

This algorithm has enjoyed widespread usage in theoretical physics (see [1] for example) and is quite successful for many applications in two or more

dimensions. The procedure's greatest strength lies in its ability to adapt to the integrand being considered. However the extent to which it can adapt is determined by the number of grid increments along each axis (N), and N in turn is limited by the total number of integrand evaluations (M) allowed per iteration (n = dimensionality):

$$M = 2N^n$$

This limitation proves to be a serious handicap in high dimensions (e.g., $M < 10^6$ for $n=9 \Rightarrow N \leq 4$), as will be illustrated in Section V.

A number of other techniques similar to Monte Carlo methods have also been suggested for multidimensional integration [2, 3, 6, 7]. These quasi-Monte Carlo techniques generally converge faster than $1/M^{1/2}$ for analytic integrands. However they are not adaptive and so may require prohibitively large numbers of function evaluations in high dimensions when the integrand has sharp peaks. Furthermore it is not clear that they are able to integrate functions which are neither analytic nor even continuous. Also no simple estimate of the error is available in most cases.

III. THE ALGORITHM IN ONE DIMENSION

The restrictive relation between the numbers of increments and of function calls in Shephey's method can be avoided using importance sampling rather than stratified sampling to reduce the variance. Intuitively importance sampling seems the inferior of the two methods [3]. However in practice the ability to adapt is the overriding consideration for high dimensions.

Like Shephey's, the algorithm described here is iterative. To illustrate consider a one-dimensional integration:

$$I = \int_0^1 dx f(x)$$

Initially an M-point Monte Carlo integration is performed with a uniform probability density ($p(x)=1$). Besides providing estimates of the integral and the possible error (Eqs. (1) and (2)), the M integrand evaluations can also be used to define an improved probability density for use in the next (M-point) Monte Carlo integration. In this fashion an empirical variance reduction can be gradually introduced over several iterations.

There exist standard numerical techniques for generating lists of evenly distributed 'pseudo-random' numbers. It is more difficult, however, to generate numbers from an arbitrary distribution with density $p(x)$. For present purposes it is natural that $p(x)$ be a step function with N steps. The probability of a random number being chosen from any given step is defined to be a constant, equal to $1/N$ for all steps:

$$p(x) = \frac{1}{N\Delta x_i} \quad x_i - \Delta x_i \leq x < x_i$$
$$i = 1, \dots, N$$

where

$$\sum_{i=1}^N \Delta x_i = 1$$

The probability distribution is tailored to particular integrands simply by adjusting the increment sizes Δx_i . In practice N is limited by the computer storage space available and must be held constant from iteration to iteration (N=50 to 100 typically).

Given M integrand evaluations, the probability distribution or, equivalently, the increment density is refined by subdividing each increment Δx_i into m_i+1 subintervals where

$$m_i = K \frac{f_i \Delta x_i}{\sum_j f_j \Delta x_j} \quad (4)$$

and

$$\begin{aligned} \bar{f}_i &\equiv \sum_{x \in x_i - \Delta x_i}^{x_i} |f(x)| \\ &\propto \frac{1}{\Delta x_i} \int_{x_i - \Delta x_i}^{x_i} dx |f(x)| \end{aligned}$$

Thus each interval is subdivided into as many as $K+1$ subintervals (K is fixed at 1000 typically), and its contribution to the weight function increased in proportion to its contribution to the integral of $|f(x)|$, as required by (3). As it is desirable to restore the number of increments to its original value ($=N$), groups of the new intervals must be amalgamated into larger intervals, the number of subintervals in each group being constant (to preserve the increment density). The net effect is to alter the increment sizes, while keeping the total number constant, so that the smallest increments occur where $|f(x)|$ is largest. The new grid is used and further refined in subsequent iterations until the optimal grid has been obtained (i. e., $m_i = m_j \quad i, j = 1, \dots, N$).

In practice it is best to damp the subdivision algorithm thereby avoiding rapid, destabilizing changes in the grid from iteration to iteration. For example, this can be done by using [5]

$$m_i = K \left\{ \frac{\left[\frac{\bar{f}_i \Delta x_i}{\sum_j \bar{f}_j \Delta x_j} - 1 \right]}{\log \left[\frac{\bar{f}_i \Delta x_i}{\sum_j \bar{f}_j \Delta x_j} \right]} \right\}^\alpha$$

in place of (4). The parameter α determines the rate of convergence and is typically set between 0.2 and 2.

The values \bar{f}_i must be discarded after each iteration because of storage limitations. However a cumulative estimate of the integral and its error can

be made which uses every evaluation of the integrand:

$$\bar{I} = \sigma_{\bar{I}}^2 \sum_i \frac{I_i}{\sigma_i^2} \tag{5}$$

$$\sigma_{\bar{I}} = \left[\sum_i \frac{1}{\sigma_i^2} \right]^{-1/2}$$

Here I_i and σ_i^2 are the integral and standard deviation estimated in iteration i using Eqs. (1) and (2). When the integrand has high, narrow peaks, I_i and σ_i are sometimes badly underestimated in the earliest iterations (before the algorithm has adapted). To partially correct for this effect, it may be better to replace Eqs. (5) by

$$\bar{I} = \frac{\sum_i I_i \left(\frac{I_i^2}{\sigma_i^2} \right)}{\sum_i \frac{I_i^2}{\sigma_i^2}} \tag{6}$$

$$\sigma_{\bar{I}} = \bar{I} \left[\sum_i \frac{I_i^2}{\sigma_i^2} \right]^{-1/2}$$

In any event the χ^2

$$\chi^2 \approx \sum_i \frac{(I_i - \bar{I})^2}{\sigma_i^2} \quad \text{or} \quad \sum_i \frac{(I_i - \bar{I})^2}{\bar{I}^2} \frac{I_i^2}{\sigma_i^2} \tag{7}$$

should not greatly exceed the number of iterations (minus one). The algorithm cannot be trusted when it does.

The number of iterations and the number of integrand evaluations per iteration needed clearly depend upon the complexity of the integrand and the accuracy being sought. In general it is best to use as few integrand evaluations per

iteration as are required by the algorithm to converge smoothly to the optimal grid, at least until it has converged. Once the optimal grid has been found (approximately), the uncertainty (σ_I) in the integral becomes roughly proportional to $1/M^{1/2}$, M being the number of integrand evaluations.

IV. THE ALGORITHM IN n DIMENSIONS

The algorithm described above is easily generalized to handle integrals of arbitrary dimensionality. To illustrate the modifications consider

$$I = \int_0^1 dx \int_0^1 dy f(x, y)$$

A separable probability density function is adopted to limit data storage requirements:

$$p(x, y) = p_x(x) p_y(y)$$

In this case the optimal densities are easily shown to be

$$p_x(x) = \frac{\left[\int_0^1 dy \frac{f^2(x, y)}{p_y(y)} \right]^{1/2}}{\int_0^1 dx \left[\int_0^1 dy \frac{f^2(x, y)}{p_y(y)} \right]^{1/2}}$$

with a similar result for $p_y(y)$. Thus the one dimensional algorithm can be applied along each axis but with \bar{f}_i (in Eq. (4)) defined by

$$\begin{aligned} (\bar{f}_i)^2 &= \sum_{x \in x_i - \Delta x_i}^{x_i} \sum_y \frac{f^2(x, y)}{p_y^2(y)} \\ &\propto \frac{1}{\Delta x_i} \int_{x_i - \Delta x_i}^{x_i} dx \int_0^1 dy \frac{f^2(x, y)}{p_y(y)} \end{aligned}$$

for the x axis and with an analogous definition for the y axis. The generalization to arbitrary dimension is obvious.

V. NUMERICAL EXAMPLES

We have encoded this algorithm in a Fortran IV program called VEGAS and compared its performance with a number of other multidimensional integration methods. A sampling of these results is presented in Tables I-V.

The test integral for Tables I and II has a spherically symmetric Gaussian placed in the center of the integration region:

$$I_n = \left(\frac{1}{a\pi^{1/2}}\right)^n \int_0^1 d^n x \exp\left(-\sum_{i=1}^n \frac{(x_i - \frac{1}{2})^2}{a^2}\right) \quad (8)$$

where $a=0.1$ and $n=4, 9$. Partial results from several iterations of VEGAS are shown, illustrating the convergence of the algorithm ($\alpha=2.0, 1.0$ for $n=4, 9$). For comparison estimates of I_n using crude Monte Carlo ($p(x)=1$) and using Gauss-Legendre formulae along each axis are also exhibited. Of course VEGAS is far superior to crude Monte Carlo. VEGAS is also considerably more efficient than Gauss-Legendre integration (especially when $n=9$), even though the integrand and all of its derivatives are analytic and bounded throughout the integration region. These examples illustrate the importance of using adaptive algorithms for integration in high dimension.

In Table III VEGAS is compared with a code (SHEP) written by Sheppey and employing the algorithm described in Section II. Here the test integral has two spherically symmetric Gaussians equally spaced along the diagonal of the integration volume:

$$I_n = \frac{1}{2} \left(\frac{1}{a\pi^{1/2}}\right)^n \int_0^1 d^n x \left\{ \exp\left(-\sum_{i=1}^n \frac{(x_i - \frac{1}{3})^2}{a^2}\right) + \exp\left(-\sum_{i=1}^n \frac{(x_i - \frac{2}{3})^2}{a^2}\right) \right\} \quad (9)$$

The integration was done for $n=2, 4, 7, 9$ with $a=0.1$. The 'Optimal Standard Deviation' quoted in this table is the standard deviation computed on any

iteration after the optimal grid has been achieved. Results averaged over 15 iterations (Eqs. (6)) are also presented.

For high dimensions ($n \geq 4$), VEGAS converges with far fewer function evaluations than SHEP. The number of function evaluations is independent of the number of increments per axis in VEGAS, and so it is better able to adapt to the integrand in high dimension than is SHEP. Note that when the algorithms are equally adaptive ($n \leq 4$), SHEP performs as well as or better than VEGAS—as expected.

In Table IV we compare VEGAS with the quasi-Monte Carlo methods suggested by Tsuda [7] and by Haselgrove [6]. The test integral is the most difficult considered by Tsuda [7]:

$$I = \int_0^1 dy_1 \cdots \int_0^1 dy_8 f(y_1 \cdots y_8) \tag{10}$$

$$f(y_1 \cdots y_8) = \prod_{i=1}^8 \frac{c}{c+1} \left(\frac{c+1}{c+y_i} \right)^2$$

where c is chosen such that $f(0 \dots 0) = 10^4$. Neither of the quasi-Monte Carlo methods allows a simple error estimate and so this must be inferred from the convergence of the estimates of I . VEGAS is the superior algorithm though Tsuda's method is almost as efficient when many function evaluations are employed.

Finally we present the results of a seventh order calculation (in Feynman perturbation theory) of part of the decay rate of orthopositronium into three photons. The contribution presented here is from the ladder kernel (Γ_g^S , in Ref. [8]). No analytic result is known; the best numerical estimate is $(-5.90 \pm 0.07) \alpha \Gamma_0 / \pi$. Note that the integral contains a step function $\theta(1-x_1-x_2)$ and is therefore zero throughout half the integration volume.

VI. CONCLUSIONS

It is clear that any general purpose multidimensional integration method should be adaptive. We have demonstrated that adaptive Monte Carlo methods are quite effective, especially in high dimensions or with nonanalytic integrands where simple generalizations of the well known one dimensional methods are not so generally applicable. Of course if the integrand is analytic and smooth, techniques such as Gauss-Legendre or quasi-Monte Carlo integration may frequently be superior to adaptive Monte Carlo integration, though for such integrands the latter is more than sufficient to obtain the three or four significant digits usually required in scientific applications.

The author is indebted to Dr. S. J. Brodsky for his comments and suggestions. He also thanks Drs. K. J. F. Gaemers, W. E. Caswell and J. Friedman for many discussions.

APPENDIX

The use of importance sampling as described in Section III is formally equivalent to making the variable transformation $x=g(y)$

$$I = \int_0^1 dx f(x) = \int_0^1 dy f(g(y)) g'(y) \quad (11)$$

with $g'(y) = p(g(y))^{-1}$ and performing a Monte Carlo integration with uniformly distributed points $\{y_i\}$. The algorithm presented in Section III adjusts $g'(y)$ such that $g'(y) \propto |f(g(y))|^{-1}$ thereby improving the Monte Carlo estimate of the integral. For some integrands it is possible to further reduce the variance through use of other well known methods (stratified sampling, antithetic variates, quasi-random numbers, . . . [2]) applied to the smoothed integrand $f(g(y)) g'(y)$ in (11). For example, rather than choosing M random points uniformly distributed on $0 \leq y \leq 1$, the sampling can be stratified by selecting two points in each of $M/2$ subintervals of $[0, 1]$.

The optimal definition of $g'(y)$ may differ from that in Section III when additional variance reduction is used. Thus for stratified sampling, $g'(y)$ should perhaps be chosen such that the variation (rather than the value itself) of $g'(y) f(g(y))$ in each subinterval is uniform over all subintervals.

Trial runs of a modified VEGAS, incorporating stratified sampling, suggest considerable gains in efficiency in low dimensions (where the modified algorithm is very similar to Sheppey's) but very little improvement in higher dimensions. Work is still in progress on this and similar modifications of VEGAS.

REFERENCES

1. J. Aldins et al., Phys. Rev. D 1 (1970), 2378;
P. Pascual and E. de Rafael, Nuovo Cimento Lett. 4 (1970), 1144;
G. Grammar, Jr. and T. Kinoshita, Nucl. Phys. B80 (1974), 461.
2. J. M. Hammersley and D. C. Handscomb, "Monte Carlo Methods,"
Chapters 3,5, Methuen, London, 1964.
3. A. H. Stroud, "Approximate Calculation of Multiple Integrals," Chapter 6,
Prentice-Hall, Englewood Cliffs, New Jersey, 1971.
4. J. H. Halton, SIAM Review 12 (1970), 1.
5. A computer program using Sheppey's method is described in A. Dufner,
Proceedings of the Colloquium on Computational Physics, Marseille, 1970,
p I-46 (unpublished);
B. Lautrup, Proceedings of the Colloquium on Computational Physics,
Marseille, 1971, p I-57 (unpublished).
6. C. B. Haselgrove, Math. Comp. 15 (1961), 323.
7. T. Tsuda, Numer. Math. 20 (1973), 377.
8. W. E. Caswell et al., Phys. Rev. Lett. 38 (1977), 488.

Table 1. Integration of a single Gaussian in 4 dimensions.

The exact result is 1.

VEGAS

Iteration	Per Iteration		Cumulative		No. of Integrand Evaluations
	I_i	σ_i	\bar{I}	$\sigma_{\bar{I}}$	
1	0.790	0.313	0.790	0.313	1000
3	0.970	0.063	0.977	0.059	3000
5	0.952	0.023	0.975	0.018	5000
10	1.008	0.016	0.994	0.007	10 000
Crude Monte Carlo			0.890	0.149	10 000

GAUSS-LEGENDRE Integration

No. of Points/Axis	Integral	No. of Integrand Evaluations
5	6.664	625
6	0.164	1296
10	0.892	10 000
13	1.008	28 561

Table II. Integration of a single Gaussian in 9 dimensions.

The exact result is 1.

VEGAS

Integration	Per Iteration		Cumulative		No. of Integrand Evaluations
	I_i	σ_i	\bar{I}	$\sigma_{\bar{I}}$	
1	0.007	0.005	0.007	0.005	10^4
3	0.643	0.070	0.612	0.064	3×10^4
5	1.009	0.041	0.963	0.034	5×10^4
10	1.003	0.008	1.001	0.005	10^5
Crude Monte Carlo			0.843	0.360	10^5

GAUSS-LEGENDRE Integration

No. of Points/Axis	Integral	No. of Integrand Evaluations
5	71.364	2.0×10^6
6	0.017	1.0×10^7
10	0.774	10^9
15	1.002	3.8×10^{10}

Table III. Integration of a double Gaussian in n dimensions. The exact result is 1.

VEGAS					SHEP				
n=2					n=7				
n=4					n=7				
n=7					n=7				
n=9					n=9				
No. Function Evaluations/Iteration	20 000	20 000	32 000	160 000	100 000	50	0.02	0.991	± .007
No. Increments/Axis	50	50	50	50	50	0.06	0.96	± .04	
Optimal Standard Deviation	0.007	0.02	0.05	0.991	0.96	0.06	0.96	± .04	
Cumulative Result After 15 Iterations	± .002	± .006	± .015	± .007	± .04				
n=2					n=7				
n=4					n=7				
n=7					n=7				
n=9					n=9				
No. Function Evaluations/Iteration	20 000	20 000	32 768	170 000	3 906 250	5	0.25	0.90	± .04
No. Increments/Axis	100	10	4	5	5	0.25	0.90	± .04	
Optimal Standard Deviation	0.0008	0.03	1.0 a)	0.25	0.25	0.25	0.90	± .04	
Cumulative Result After 15 Iterations	1.0001	1.004	*	0.90	0.90	0.90	0.90	± .04	
n=2					n=7				
n=4					n=7				
n=7					n=7				
n=9					n=9				

a) Algorithm would not converge.

b) Too large to be tried.

Table IV. Integrand (10) using VEGAS, Tsuda's algorithm [7], and Haselgrove's algorithm [6]. The exact result is 1. The numbers for Tsuda's method are from Table I in Ref. [7].

No. of Integrand Evaluations	VEGAS	TSUDA (n=40)	HASELGROVE (s_2)
1000	1.083 $\pm .085$	1.248	11.019
5000	1.003 $\pm .004$	0.982	2.946
10 000	1.000 $\pm .002$	1.006	1.974
20 000	1.000 $\pm .001$	1.002	1.484

Table V. Contribution of ladder graph to decay rate of orthopositronium (in units of $\alpha\Gamma_0/\pi$).

	VEGAS n=5	SHEP n=5
No. Function Evaluations/Iteration	33 000	33 000
No. Increments/Axis	50	7
Optimal Standard Deviation	0.73	1.21
Cumulative Result After 15 Iterations	-5.82 $\pm .22$	-6.34 $\pm .35$