

CUBA – a library for multidimensional numerical integration

T. Hahn
Max-Planck-Institut für Physik
Föhringer Ring 6, D-80805 Munich, Germany

April 5, 2004

Abstract

The CUBA library provides new implementations of four general-purpose multidimensional integration algorithms: Vegas, Suave, Divonne, and Cuhre. Suave is a new algorithm, Divonne is a known algorithm to which important details have been added, and Vegas and Cuhre are new implementations of existing algorithms with only few improvements over the original versions. All four algorithms can integrate vector integrands and have very similar Fortran, C/C++, and Mathematica interfaces.

1 Introduction

Many problems in physics (and elsewhere) involve computing an integral, and often enough this has to be done numerically, as the analytical result is known only in a limited number of cases. In one dimension, the situation is quite satisfactory: standard packages, such as QUADPACK [1], reliably integrate a broad class of functions in modest CPU time. The same is unfortunately not true for multidimensional integrals.

This paper presents the CUBA library with new implementations of four algorithms for multidimensional numerical integration: Vegas, Suave, Divonne, and Cuhre. They have a C/C++, Fortran, and Mathematica interface each and are invoked in a very similar way, thus making them easily interchangeable, e.g. for comparison purposes. All four can integrate vector integrands. Cuhre is a deterministic algorithm, the others use Monte Carlo methods.

Vegas is the simplest of the four. It uses importance sampling for variance reduction, but is only in some cases competitive in terms of the number of samples needed to reach a prescribed accuracy. Nevertheless, it has a few improvements over the original algorithm [2, 3] and comes in handy for cross-checking the results of other methods.

Suave is a new algorithm which combines the advantages of two popular methods: importance sampling as done by Vegas and subregion sampling in a manner similar to Miser [4]. By dividing into subregions, Suave manages to a certain extent to get around

Vegas' difficulty to adapt its weight function to structures not aligned with the coordinate axes.

Divonne is a further development of the CERNLIB routine D151 [5]. Divonne works by stratified sampling, where the partitioning of the integration region is aided by methods from numerical optimization. A number of improvements have been added to this algorithm, the most significant being the possibility to supply knowledge about the integrand. Narrow peaks in particular are difficult to find without sampling very many points, especially in high dimensions. Often the exact or approximate location of such peaks is known from analytic considerations, however, and with such hints the desired accuracy can be reached with far fewer points.

Cuhre* employs a cubature rule for subregion estimation in a globally adaptive subdivision scheme [6]. It is hence a deterministic, not a Monte Carlo method. In each iteration, the subregion with the largest error is halved along the axis where the integrand has the largest fourth difference. Cuhre is quite powerful in moderate dimensions, and is usually the only viable method to obtain high precision, say relative accuracies much below 10^{-3} .

The new algorithms were coded from scratch in C, which is a compromise of sorts between C++ and Fortran, combining ease of linking to Fortran code with the availability of reasonable memory management. The declarations have been chosen such that the routines can be called from Fortran directly. The Mathematica versions are based on the same C code and use the MathLink API to communicate with Mathematica.

2 Vegas

Vegas is a Monte Carlo algorithm that uses importance sampling as a variance-reduction technique. Vegas iteratively builds up a piecewise constant weight function, represented on a rectangular grid. Each iteration consists of a sampling step followed by a refinement of the grid. The exact details of the algorithm can be found in [2, 3] and shall not be reproduced here.

Changes with respect to the original version are:

- Sobol quasi-random numbers [7] rather than pseudo-random numbers are used. Empirically, this seems to accelerate convergence quite a bit, most noticeably in the early stages of the integration.

From theoretical considerations it is of course known (see e.g. [8]) that quasi-random sequences yield a convergence rate of $\mathcal{O}(\log^{n_d} n_s/n_s)$, where n_d is the number of dimensions and n_s the number of samples, which is much better than the usual $\mathcal{O}(1/\sqrt{n_s})$ for ordinary Monte Carlo. But these convergence rates are meaningful only for large n_s and so it came as a pleasant surprise that the gains are considerable already at the beginning of the sampling process. It shows that quasi-Monte Carlo methods blend well with variance-reduction techniques such as importance sampling.

*The D from the original name was dropped since the CUBA library uses double precision throughout.

Similarly, it was not clear from the outset whether the statistical standard error would furnish a suitable error estimate since quasi-random numbers are decidedly non-random in a number of respects. Yet also here empirical evidence suggests that the standard error works just as well as for pseudo-random numbers.

- The present implementation allows the number of samples to be increased in each iteration. With this one can mimic the strategy of calling Vegas with a small number of samples first to ‘get the grid right’ and then using an alternate entry point to perform the ‘full job’ on the same grid with a larger number of samples.
- The option to add simple stratified sampling on top of the importance sampling, as proposed in the appendix of [2], has not been implemented in the present version. Tests with the Vegas version from [9], which contains this feature, showed that convergence was accelerated only when the original pseudo-random numbers were used and that with quasi-random numbers convergence was in fact even slower in some cases.

Vegas’ major weakness is that it uses a separable (product) weight function. As a consequence, Vegas can offer significant improvements only as far as the integrand’s characteristic regions are aligned with the coordinate axes.

3 Suave

Suave uses Vegas-like importance sampling combined with a globally adaptive subdivision strategy: Until the requested accuracy is reached, the region with the largest error at the time is bisected in the dimension in which the fluctuations of the integrand are reduced most. The number of new samples in each half is prorated for the fluctuation in that half.

A similar method, known as recursive stratified sampling, is implemented in Miser [4]. Miser always samples a fixed number of points, however, which is somewhat undesirable since it does not stop once the prescribed accuracy is reached.

Suave first samples the integration region in a Vegas-like step, i.e. using importance sampling with a separable weight function. It then slices the integration region in two, as Miser would do. Suave does not immediately recurse on those subregions, however, but maintains a list of all subregions and selects the region with the largest absolute error for the next cycle of sampling and subdivision. That is, Suave uses global error estimation and terminates when the requested relative or absolute accuracy is attained.

The information on the weight function collected in one Vegas step is not lost. Rather, the grid from which the weight function is computed is stretched and re-used on the subregions. A region which is the result of $m - 1$ subdivisions thus has had m Vegas iterations performed on it.

The improvements over Vegas and Miser come at a price, which is the amount of memory required to hold all the samples. Memory consumption is not really severe on

modern hardware, however. The component that scales worst is the one proportional to the number of samples, which is

$$8(n_d + n_c + 1)n_s \text{ bytes},$$

where n_d is the number of dimensions of the integral, n_c the number of components of the integrand, and n_s the number of samples. For a million samples on a scalar integrand of 10 variables, this works out to 96 megabytes – not all that enormous these days.

3.1 Description of the algorithm

As Suave is a new algorithm, the following description will be fairly detailed. For greater notational clarity, n_c -dimensional vectors are denoted with a vector arrow (\vec{f}) and n_d -dimensional vectors with boldface letters (\mathbf{x}) in the following, where n_d is the dimension of the integral and n_c the number of components of the integrand.

The essential inputs are ε_{rel} and ε_{abs} , the relative and absolute accuracies, n_s^{new} , the number of samples added in each iteration, n_s^{max} , the maximum number of samples allowed, and p , a flatness parameter described below.

Suave has a main loop which calls a Vegas-like sampling step. The main loop is responsible for subdividing the subregions and maintaining the totals. The sampling step does the actual sampling on the subregions and computes the region results.

3.1.1 Main loop

1. Initialize the quasi-random-number generator and allocate a data structure for the entire integration region. Initialize its Vegas grid with equidistant bins.
2. Sample the entire integration region with n_s^{new} points. This gives an initial estimate of the integral \vec{I}_{tot} , the variance $\vec{\sigma}_{\text{tot}}^2$, and $\vec{\chi}_{\text{tot}}^2$.
3. Find the component c for which $r_c = \sigma_{c,\text{tot}} / \max(\varepsilon_{\text{abs}}, \varepsilon_{\text{rel}} I_{c,\text{tot}})$ is maximal.
If none of the r_c 's exceeds unity, indicate success and return.
4. If the number of samples spent so far equals or exceeds n_s^{max} , indicate failure and return.
5. Find the region r with the largest σ_c^2 .
6. Find the dimension d which minimizes $F_c(r_L^d) + F_c(r_R^d)$, where $r_{L,R}^d$ are the left and right halves of r with respect to d . $F_c(r_{L,R}^d)$ is the fluctuation of the samples that fall into $r_{L,R}^d$ and is computed as

$$F_c(r_{L,R}^d) = \left[\left\| 1 + \tilde{F}_c(\mathbf{x}_i \in r_{L,R}^d) \right\|_p \right]^{2/3} = \left[\sum \left| 1 + \tilde{F}_c(\mathbf{x}_i \in r_{L,R}^d) \right|^p \right]^{2/3p}, \quad (1)$$

where all samples \mathbf{x}_i that fall into the respective half are used in the norm/sum and the single-sample fluctuation \tilde{F}_c is defined as

$$\tilde{F}_c(\mathbf{x}) = w(\mathbf{x}) \left| \frac{f_c(\mathbf{x}) - I_c(r)}{I_c(r)} \right| \frac{|f_c(\mathbf{x}) - I_c(r)|}{\sigma_c(r)}.$$

This empirical recipe combines the relative deviation from the region mean, $(f - I)/I$, with the χ value, $|f - I|/\sigma$, weighted by the Vegas weight w corresponding to sample \mathbf{x} . Note that the I_c and σ_c values of the entire region r are used.

Samples strongly contribute to F the more they lie away from the predicted mean *and* the more they lie out of the predicted error band. Tests have shown that large values of p are beneficial for ‘flat’ integrands, whereas small values are preferred if the integrand is ‘volatile’ and has high peaks. p has thus been dubbed a flatness parameter. The effect comes from the fact that with increasing p , F becomes more and more dominated by ‘outliers,’ i.e. points with a large \tilde{F} .

The power $2/3$ in Eq. (1) is also used in Miser, where it is motivated as the exponent that gives the best variance reduction ([9], p. 315).

7. Refine the grid associated with r , i.e. incorporate the information gathered on the integrand in the most recent sample over r into the weight function. This is done precisely as in Vegas (see [2]), with the extension that if the integrand has more than one component, the marginal densities are computed not from f^2 but from the weighted sum[†]

$$\overline{f^2} = \sum_{c=1}^{n_c} \frac{f_c^2}{I_{c,\text{tot}}^2}.$$

8. Bisect r in dimension d :

Allocate a new region, r_L , and copy to r_L those of r ’s samples falling into the left half. Compute the Vegas grid for r_L by appropriately “stretching” r ’s grid, i.e. by interpolating all grid points of r with values less than $1/2$.

Set up r_R for the right half analogously.

9. Sample r_L with $n_L = \max\left(\frac{F_c(r_L)}{F_c(r_L) + F_c(r_R)} n_s^{\text{new}}, n_s^{\text{min}}\right)$ and r_R with $n_R = \max(n_s^{\text{new}} - n_L, n_s^{\text{min}})$ points, where $n_s^{\text{min}} = 10$.

[†]It is fairly obvious that scale-invariant quantities must be used in the sum, otherwise the component with the largest absolute scale would dominate. It is less clear whether $\eta_0 = (\int f_c d\mathbf{x})^2 = I_{c,\text{tot}}^2$, $\eta_1 = (\int |f_c| d\mathbf{x})^2$, or $\eta_2 = \int f_c^2 d\mathbf{x}$ (or any other) make the best weights. Empirically, η_0 turns out to be both slightly superior in convergence and easier to compute than η_1 and η_2 and has thus been chosen in Suave.

A possible explanation for this is that in cases where there are large compensations within the integral, i.e. when $\int f_c d\mathbf{x} \ll \int |f_c| d\mathbf{x}$, it is particularly necessary for the overall accuracy that component c be sampled accurately, and thus be given more weight in $\overline{f^2}$, and this is better accomplished by dividing f_c^2 by the “small” number η_0 than by the “large” number η_1 or η_2 .

10. To safeguard against underestimated errors, supplement the variances by the difference of the integral values in the following way:

$$\sigma_{c,\text{new}}^2(r_{R,L}) = \sigma_c^2(r_{R,L}) \left(1 + \frac{\Delta_c}{\sqrt{\sigma_c^2(r_L) + \sigma_c^2(r_R)}} \right)^2 + \Delta_c^2$$

for each component c , where $\vec{\Delta} = \frac{1}{4}|\vec{I}(r_L) + \vec{I}(r_R) - \vec{I}(r)|$.

This acts as a penalty for regions whose integral value changes significantly by the subdivision and effectively moves them up in the order of regions to be subdivided next.

11. Update the totals: Subtract r 's integral, variance, and χ^2 -value from the totals and add those of r_L and r_R .
12. Discard r , put r_L and r_R in the list of regions.
13. Go to Step 3.

3.1.2 Sampling step

The function which does the actual sampling is a modified Vegas iteration. It is invoked with two arguments: r , the region to be sampled and n_m , the number of new samples.

1. Sample a set of n_m new points using the weight function given by the grid associated with r . For a region which is the result of $m - 1$ subdivisions, the list of samples now consists of m sets of samples.
2. For each set of samples, compute the mean \vec{I}_i and variance $\vec{\sigma}_i^2$.
3. Compute the results for the region as

$$I_c = \frac{\sum_{i=1}^m w_{i,c} I_{i,c}}{\sum_{i=1}^m w_{i,c}}, \quad \sigma_c^2 = \frac{1}{\sum_{i=1}^m w_{i,c}}, \quad \chi_c^2 = \frac{1}{\sigma_c^2} \left[\frac{\sum_{i=1}^m w_{i,c} I_{i,c}^2}{\sum_{i=1}^m w_{i,c}} - I_c^2 \right],$$

where the inverse of the set variances are used as weights, $w_{i,c} = 1/\sigma_{i,c}^2$. This is simply Gaussian error propagation.

For greater numerical stability, χ_c^2 is actually computed as

$$\chi_c^2 = \sum_{i=1}^m w_{i,c} I_{i,c}^2 - I_c \sum_{i=1}^m w_{i,c} I_{i,c} = \sum_{i=2}^m w_{i,c} I_{i,c} (I_{i,c} - I_{1,c}) - I_c \sum_{i=2}^m w_{i,c} (I_{i,c} - I_{1,c}).$$

4 Divonne

Divonne uses stratified sampling for variance reduction, that is, it partitions the integration region such that all subregions have an approximately equal value of a quantity called the spread \vec{s} , defined as

$$\vec{s}(r) = \frac{1}{2}V(r)\left(\max_{\mathbf{x} \in r} \vec{f}(\mathbf{x}) - \min_{\mathbf{x} \in r} \vec{f}(\mathbf{x})\right), \quad (2)$$

where $V(r)$ is the volume of region r . What sets Divonne apart from Suave is that the minimum and maximum of the integrand are sought using methods from numerical optimization. Particularly in high dimensions, the chance that one of the previously sampled points lies in or even close to the true extremum is fairly small.

On the other hand, the numerical minimization is beset with the usual pitfalls, i.e. starting from the lowest of a (relatively small) number of sampled points, Divonne will move directly into the local minimum closest to the starting point, which may or may not be close to the absolute minimum.

Divonne is a lot more complex than Suave and Vegas but also significantly faster for many integrands. For details on the methods used in Divonne please consult the original references [5]. New features with respect to the CERNLIB version (Divonne 4) are:

- Integration is possible in dimensions 2 through 33 (not 9 as before). Going to higher dimensions is a matter of extending internal tables only.
- The possibility has been added to specify the location of possible peaks, if such are known from analytical considerations. The idea here is to help the integrator find the extrema of the integrand, and narrow peaks in particular are a challenge for the algorithm. Even if only the approximate location is known, this feature of hinting the integrator can easily cut an order of magnitude out of the number of samples needed to reach the required accuracy for complicated integrands. The points can be specified either statically, by passing a list of points at the invocation, or dynamically, through a subroutine called for each subregion.
- Often the integrand subroutine cannot sample points lying on or very close to the integration border. This can be a problem with Divonne which actively searches for the extrema of the integrand and homes in on peaks regardless of whether they lie on the border. The user may however specify a border region in which integrand values are not obtained directly, but extrapolated from two points inside the ‘safe’ interior.
- The present algorithm works in three phases, not two as before. Phase 1 performs the partitioning as outlined above. From the preliminary results obtained in this phase, Divonne estimates the number of samples necessary to reach the desired accuracy in phase 2, the final integration phase. Once the phase-2 sample for a particular subregion is in, a χ^2 test is used to assess whether the two sample averages are consistent with each other within their error bounds. Subregions which fail this test

move on to phase 3, the refinement phase, where they can be subdivided again or sampled a third time with more points, depending on the parameters set by the user.

- For all three phases the user has a selection of methods to obtain the integral estimate: a Korobov [10] or Sobol [7] quasi-random sample of given size, and the cubature rules of Genz and Malik [11] of degree 7, 9, 11, and 13 that are also used in Cuhre. The latter are embedded rules and hence provide an intrinsic error estimate (that is, an error estimate not based on the spread). When this independent error estimate is available, it supersedes the spread-based error when computing the total error. Also, regions whose spread-based error exceeds the intrinsic error are selected for refinement, too.

In spite of these novel options, pseudo-random sampling and the cubature rules of the original Divonne algorithm were not implemented.

Due to its complexity, the new Divonne implementation was painstakingly tested against the CERNLIB routine to make sure it produces the same results before adding the new features.

5 Cuhre

Cuhre is a deterministic algorithm which uses one of several cubature rules of polynomial degree in a globally adaptive subdivision scheme. The subdivision algorithm is similar to Suave's (see Sect. 3.1.1) and works as follows:

While the total estimated error exceeds the requested bounds:

- 1) choose the region with the largest estimated error,
- 2) bisect this region along the axis with the largest fourth difference,
- 3) apply the cubature rule to the two subregions,
- 4) merge the subregions into the list of regions and update the totals.

Details on the algorithm and on the cubature rules employed in Cuhre can be found in the original references [6]. The present implementation offers only superficial improvements, such as an interface consistent with the other CUBA routines and a slightly simpler invocation, e.g. one does not have to allocate a workspace.

In moderate dimensions Cuhre is very competitive, particularly if the integrand is well approximated by polynomials. As the dimension increases, the number of points sampled by the cubature rules rises considerably, however, and by the same token the usefulness declines. For the lower dimensions, the actual number of points that are spent per invocation of the basic integration rule are listed in the following table.

number of dimensions	4	5	6	7	8	9	10	11	12
points in degree-7 rule	65	103	161	255	417	711	1265	2335	4433
points in degree-9 rule	153	273	453	717	1105	1689	2605	4117	6745

6 Download and Compilation

The source code of the CUBA library can be downloaded as a gzipped tar file from the Web site <http://www.feynarts.de/cuba>. The archive unpacks into a directory `Cuba-1.0`. Change into this directory and type “make” to build the library `libcuba.a` and the MathLink executables `Vegas`, `Suave`, `Divonne`, and `Cuhre`. If Mathematica and/or the `mcc` MathLink compiler are not available, type “make lib” to build just the library.

The distribution contains two demonstration programs in Fortran and C, as well as the test suite used in Sect. 8, which is written in Mathematica.

The code is C99 compliant and compiles flawlessly with the GNU C compiler, versions 2.95 and higher. Other C compilers may have difficulties with inline functions and variable-size arrays, which are C99 extensions. In a pinch, edit the makefile and uncomment the line

```
CFLAGS += -DNDIM=8 -DNCOMP=2
```

This fixes the size of all internal arrays at compile time but, of course, at most 8-dimensional integrals of at most 2-component integrands can now be integrated.

Linking Fortran or C/C++ code that uses one of the algorithms is straightforward, just add `-lcuba` (for the CUBA library) and `-lm` (for the math library) to the compiler command line, as in

```
f77 -o myexecutable mysource.f -lcuba -lm
cc -o myexecutable mysource.c -lcuba -lm
```

7 User Manual

7.1 Usage in Fortran

Although written in C, the declarations have been chosen such that the routines are directly accessible from Fortran, i.e. no wrapper code is needed. In fact, `Vegas`, `Suave`, `Divonne`, and `Cuhre` can be called as if they were Fortran subroutines respectively declared as

```
subroutine vegas(ndim, ncomp, integrand,
&   epsrel, epsabs, flags, mineval, maxeval,
&   nstart, nincrease,
&   neval, fail, integral, error, prob)
```



```
subroutine suave(ndim, ncomp, integrand,
&   epsrel, epsabs, flags, mineval, maxeval,
&   nnew, flatness,
&   nregions, neval, fail, integral, error, prob)
```

```

      subroutine divonne(ndim, ncomp, integrand,
&      epsrel, epsabs, flags, mineval, maxeval,
&      key1, key2, key3, maxpass,
&      border, maxchisq, mindeviation,
&      ngiven, ldxgiven, xgiven, nexttra, peakfinder,
&      nregions, neval, fail, integral, error, prob)

      subroutine cuhre(ndim, ncomp, integrand,
&      epsrel, epsabs, flags, mineval, maxeval,
&      key,
&      nregions, neval, fail, integral, error, prob)

```

7.1.1 Common Arguments

- **integer** `ndim` $\langle in \rangle$, the number of dimensions of the integral.
- **integer** `ncomp` $\langle in \rangle$, the number of components of the integrand.
- **external** `integrand` $\langle in \rangle$, the integrand. The external subroutine which computes the integrand is expected to be declared as

```

      subroutine integrand(ndim, x, ncomp, f)
      integer ndim, ncomp
      double precision x(ndim), f(ncomp)

```

- **double precision** `epsrel`, `epsabs` $\langle in \rangle$, the requested relative and absolute accuracies. The integrator tries to find an estimate \hat{I} for the integral I which for every component c fulfills $|\hat{I}_c - I_c| \leq \max(\varepsilon_{\text{abs}}, \varepsilon_{\text{rel}} I_c)$.
- **integer** `flags` $\langle in \rangle$, flags governing the integration:
 - Bits 0 and 1 encode the verbosity level, i.e. 0 to 3.
Level 0 does not print any output, level 1 prints ‘reasonable’ information on the progress of the integration, level 2 also echoes the input parameters, and level 3 further prints the subregion results (if applicable).
 - Bit 2 = 0, all sets of samples collected on a subregion during the various iterations or phases contribute to the final result.
 - Bit 2 = 1, only the last (largest) set of samples is used in the final result.

To select e.g. last samples only together with verbosity level 2, pass $6 = 4 + 2$ for the flags.

- **integer** `mineval` $\langle in \rangle$, the minimum number of integrand evaluations required.

- `integer maxeval` $\langle in \rangle$, the (approximate) maximum number of integrand evaluations allowed.
- `integer nregions` $\langle out \rangle$, the actual number of subregions needed (not present in Vegas).
- `integer neval` $\langle out \rangle$, the actual number of integrand evaluations needed.
- `integer fail` $\langle out \rangle$, an error flag:
 - `fail` = 0, the desired accuracy was reached,
 - `fail` = -1, dimension out of range,
 - `fail` > 0, the accuracy goal was not met within the allowed maximum number of integrand evaluations. While Vegas, Suave, and Cuhre simply return 1, Divonne can estimate the number of points by which `maxeval` needs to be increased to reach the desired accuracy and returns this value.
- `double precision integral(ncomp)` $\langle out \rangle$, the integral of `integrand` over the unit hypercube.
- `double precision error(ncomp)` $\langle out \rangle$, the presumed absolute error of `integral`.
- `double precision prob(ncomp)` $\langle out \rangle$, the χ^2 -probability (not the χ^2 -value itself!) that `error` is not a reliable estimate of the true integration error[‡].

7.1.2 Vegas-specific Arguments

- `integer nstart` $\langle in \rangle$, the number of integrand evaluations per iteration to start with.
- `integer nincrease` $\langle in \rangle$, the increase in the number of integrand evaluations per iteration.

7.1.3 Suave-specific Arguments

- `integer nnew` $\langle in \rangle$, the number of new integrand evaluations in each subdivision.
- `double precision flatness` $\langle in \rangle$, the parameter p in Eq. (1), i.e. the type of norm used to compute the fluctuation of a sample. This determines how prominently ‘outliers,’ i.e. individual samples with a large fluctuation, figure in the total fluctuation, which in turn determines how a region is split up. As suggested by its name, `flatness` should be chosen large for ‘flat’ integrands and small for ‘volatile’ integrands with

[‡]To judge the reliability of the result expressed through `prob`, remember that it is the null hypothesis that is tested by the χ^2 test, which is that `error` is a reliable estimate. In statistics, the null hypothesis may be rejected only if `prob` is fairly close to unity, say `prob` > .95.

high peaks. Note that since `flatness` appears in the exponent, one should not use too large values (say, no more than a few hundred) lest terms be truncated internally to prevent overflow.

7.1.4 Divonne-specific Arguments

- **integer** `key1` *<in>*, determines sampling in the partitioning phase:
`key1 = 7, 9, 11, 13` selects the cubature rule of degree `key1`. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.
For other values of `key1`, a quasi-random sample of $n_1 = |\text{key1}|$ points is used, where the sign of `key1` determines the type of sample,
 - `key1 > 0`, use a Korobov quasi-random sample,
 - `key1 < 0`, use a Sobol quasi-random sample.
- **integer** `key2` *<in>*, determines sampling in the final integration phase:
`key2 = 7, 9, 11, 13` selects the cubature rule of degree `key2`. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.
For other values of `key2`, a quasi-random sample is used, where the sign of `key2` determines the type of sample,
 - `key2 > 0`, use a Korobov quasi-random sample,
 - `key2 < 0`, use a Sobol quasi-random sample,
and $n_2 = |\text{key2}|$ determines the number of points,
 - $n_2 \geq 40$, sample n_2 points,
 - $n_2 < 40$, sample $n_2 n_{\text{need}}$ points, where n_{need} is the number of points needed to reach the prescribed accuracy, as estimated by Divonne from the results of the partitioning phase.
- **integer** `key3` *<in>*, sets the strategy for the refinement phase:
`key3 = 0`, do not treat the subregion any further.
`key3 = 1`, split the subregion up once more.
Otherwise, the subregion is sampled a third time with `key3` specifying the sampling parameters exactly as `key2` above.
- **integer** `maxpass` *<in>*, controls the thoroughness of the partitioning phase: The partitioning phase terminates when the estimated total number of integrand evaluations (partitioning plus final integration) does not decrease for `maxpass` successive iterations.

A decrease in points generally indicates that Divonne discovered new structures of the integrand and was able to find a more effective partitioning. `maxpass` can be understood as the number of ‘safety’ iterations that are performed before the partition is accepted as final and counting consequently restarts at zero whenever new structures are found.

- **double precision border** $\langle in \rangle$, the width of the border of the integration region. Points falling into this border region will not be sampled directly, but will be extrapolated from two samples from the interior. Use a nonzero **border** if the integrand subroutine cannot produce values directly on the integration boundary.
- **double precision maxchisq** $\langle in \rangle$, the maximum χ^2 value a single subregion is allowed to have in the final integration phase. Regions which fail this χ^2 test and whose sample averages differ by more than **mindeviation** move on to the refinement phase.
- **double precision mindeviation** $\langle in \rangle$, a bound, given as the fraction of the requested error of the entire integral, which determines whether it is worthwhile further examining a region that failed the χ^2 test. Only if the two sampling averages obtained for the region differ by more than this bound is the region further treated.
- **integer ngiven** $\langle in \rangle$, the number of points in the **xgiven** array.
- **integer ldxgiven** $\langle in \rangle$, the leading dimension of **xgiven**, i.e. the offset between one point and the next in memory.
- **double precision xgiven(ldxgiven,ngiven)** $\langle in \rangle$, a list of points where the integrand might have peaks. Divonne will consider these points when partitioning the integration region. The idea here is to help the integrator find the extrema of the integrand in the presence of very narrow peaks. Even if only the approximate location of such peaks is known, this can considerably speed up convergence.
- **integer nextra** $\langle in \rangle$, the maximum number of extra points the peak-finder subroutine will return. If **nextra** is zero, **peakfinder** is not called and an arbitrary object may be passed in its place, e.g. just 0.
- **external peakfinder** $\langle in \rangle$, the peak-finder subroutine. This subroutine is called whenever a region is up for subdivision and is supposed to point out possible peaks lying in the region, thus acting as the dynamic counterpart of the static list of points supplied in **xgiven**. It is expected to be declared as

```
subroutine peakfinder(ndim, b, n, x)
integer ndim, n
double precision b(2,ndim)
double precision x(ldxgiven,n)
```

The bounds of the subregion are passed in the array **b**, where **b(1, *d*)** is the lower and **b(2, *d*)** the upper bound in dimension *d*. On entry, **n** specifies the maximum number of points that may be written to **x**. On exit, **n** must contain the actual number of points in **x**.

In contrast to the other algorithms, Divonne passes the integrand one more argument, i.e. the integrand subroutine is really declared as

```
subroutine integrand(ndim, x, ncomp, f, phase)
integer ndim, ncomp, phase
double precision x(ndim), f(ncomp)
```

The fifth argument, **phase**, indicates the integration phase:

- 0, sampling of the points in **xgiven**,
- 1, partitioning phase,
- 2, final integration phase,
- 3, refinement phase.

This information might be useful if the integrand takes long to compute and a sufficiently accurate approximation of the integrand is available. The actual value of the integral is only of minor importance in the partitioning phase, which is instead much more dependent on the peak structure of the integrand to find an appropriate tessellation. An approximation which reproduces the peak structure while leaving out the fine details might hence be a perfectly viable and much faster substitute when **phase .lt. 2**.

In all other instances, **phase** can be ignored and it is entirely admissible to declare the integrand with only four arguments.

7.1.5 Cuhre-specific Arguments

- **integer key** *<in>*, chooses the basic integration rule:
key = 7, 9, 11, 13 selects the cubature rule of degree **key**. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.
For other values, the default rule is taken, which is the degree-13 rule in 2 dimensions, the degree-11 rule in 3 dimensions, and the degree-9 rule otherwise.

7.2 Usage in C/C++

Being written in C, the algorithms can of course be used in C/C++ directly. The declarations are as follows:

```

typedef void (*integrand_t)(const int *, const double [],
    const int *, double []);

void Vegas(const int ndim, const int ncomp, integrand_t integrand,
    const double epsrel, const double epsabs,
    const int flags, const int mineval, const int maxeval,
    const int nstart, const int nincrease,
    int *neval, int *fail,
    double integral[], double error[], double prob[])

void Suave(const int ndim, const int ncomp, integrand_t integrand,
    const double epsrel, const double epsabs,
    const int flags, const int mineval, const int maxeval,
    const int nnew, const double flatness,
    int *nregions, int *neval, int *fail,
    double integral[], double error[], double prob[])

void Divonne(const int ndim, const int ncomp, integrand_t integrand,
    const double epsrel, const double epsabs,
    const int flags, const int mineval, const int maxeval,
    const int key1, const int key2, const int key3,
    const int maxpass, const double border,
    const double maxchisq, const double mindeviation,
    const int ngiven, const int ldxgiven, double xgiven[],
    const int nextra,
    void (*peakfinder)(const int *, const double [], int *, double []),
    int *nregions, int *neval, int *fail,
    double integral[], double error[], double prob[])

void Cuhre(const int ndim, const int ncomp, integrand_t integrand,
    const double epsrel, const double epsabs,
    const int flags, const int mineval, const int maxeval,
    const int key,
    int *nregions, int *neval, int *fail,
    double integral[], double error[], double prob[])

```

These prototypes are contained in `cuba.h` which should (in C) or must (in C++) be included when using the CUBA routines. The arguments are as in the Fortran case, with the obvious translations, e.g. `double precision = double`. Note, however, the declarations of the integrand and peak-finder functions, which expect pointers to integers rather than integers. This is required for compatibility with Fortran.

For convenience, the Divonne prototype glosses over the fact that Divonne passes an optional fifth argument to the integrand (see end of Sect. 7.1.4). Usually the integrand

is declared with only four arguments since this extra information is not needed. With the ‘correct’ prototype, the compiler would only generate unnecessary warnings (in C) or errors (in C++). In the rare cases where the integrand really has five arguments, an explicit typecast to `integrand_t` must be used in the invocation of `Divonne`.

7.3 Usage in Mathematica

The Mathematica versions are based on essentially the same C code and communicate with Mathematica via the MathLink API. When building the package, the executables `Vegas`, `Suave`, `Divonne`, and `Cuhre` are compiled for use in Mathematica. In Mathematica one first needs to load them with the `Install` function, as in

```
Install["Divonne"]
```

which makes a Mathematica function of the same name available. These functions are used almost like `NIntegrate`, only some options are different. For example,

```
Vegas[x^2/(Cos[x + y + 1] + 5), {x,0,5}, {y,0,5}]
```

integrates a scalar function, or

```
Suave[{Sin[z] Exp[-x^2 - y^2],  
Cos[z] Exp[-x^2 - y^2]}, {x,-1,1}, {y,-1,3}, {z,0,1}]
```

integrates a vector. As is evident, the integration region can be chosen different from the unit hypercube. Innermore boundaries may depend on outermore integration variables, e.g. `Cuhre[1, {x,0,1}, {y,0,x}]` gives the area of the unit triangle.

The functions return a list which contains the results for each component of the integrand in a sublist {integral estimate, estimated absolute error, χ^2 probability}. For the `Suave` example above this would be

```
{{1.1216, 0.000991577, 0.0000104605},  
 {2.05246, 0.00146661, 0.00920716}}
```

The other parameters are specified via the following options. Default values are given on the right-hand sides of the rules.

7.3.1 Common Options

- `PrecisionGoal` \rightarrow 3, the number of digits of relative accuracy to seek, that is, $\varepsilon_{\text{rel}} = 10^{-\text{PrecisionGoal}}$.
- `AccuracyGoal` \rightarrow 12, the number of digits of absolute accuracy to seek, that is, $\varepsilon_{\text{abs}} = 10^{-\text{AccuracyGoal}}$. The integrator tries to find an estimate \hat{I} for the integral I which for every component c fulfills $|\hat{I}_c - I_c| \leq \max(\varepsilon_{\text{abs}}, \varepsilon_{\text{rel}} I_c)$.

- **MinPoints** -> 0, the minimum number of integrand evaluations required.
- **MaxPoints** -> 50000, the (approximate) maximum number of integrand evaluations allowed.
- **Verbose** -> 1, how much information to print on intermediate results, can take values from 0 to 3.

Level 0 does not print any output, level 1 prints ‘reasonable’ information on the progress of the integration, level 2 also echoes the input parameters, and level 3 further prints the subregion results (if applicable). Note that the subregion boundaries in the level-3 printout refer to the unit hypercube, i.e. are possibly scaled with respect to the integration limits passed to Mathematica. This is because the underlying C code, which emits the output, is unaware of any scaling of the integration region, which is done entirely in Mathematica.

- **Final** -> All, whether only the last (largest) or all sets of samples collected on a subregion during the various iterations or phases contribute to the final result.
- **Regions** -> False, whether to return the tessellation of the integration region (thus not present in Vegas, which does not partition the integration region).

If **Regions** -> True is chosen, a two-component list is returned, where the first element is the list of regions, and the second element is the integration result as described above. Each region is specified in the form **Region**[x_{ll} , x_{ur} , res , df], where x_{ll} and x_{ur} are the multidimensional equivalents of the lower left and upper right corner, res is the integration result for the subregion, given in the same form as the total result but with the χ^2 value instead of the χ^2 probability, and df are the degrees of freedom corresponding to the χ^2 values.

Cuhre cannot state a χ^2 value separately for each region, hence the χ^2 values and degrees of freedom are omitted from the **Region** information.

- **Compiled** -> True, whether to compile the integrand function before use. Note two caveats:
 - The function values still have to pass through the MathLink interface, and in the course of this are truncated to machine precision. Not compiling the integrand will thus in general not deliver more accurate results.
 - Compilation should be switched off if the compiled integrand shows unexpected behaviour. As the Mathematica online help points out, “the number of times and the order in which objects are evaluated by **Compile** may be different from ordinary Mathematica code.”

7.3.2 Vegas-specific Options

- **NStart** -> 1000, the number of integrand evaluations per iteration to start with.
- **NIncrease** -> 500, the increase in the number of integrand evaluations per iteration.

7.3.3 Suave-specific Options

- **NNew** -> 1000, the number of new integrand evaluations in each subdivision.
- **Flatness** -> 50, the parameter p in Eq. (1), i.e. the type of norm used to compute the fluctuation of a sample. This determines how prominently ‘outliers,’ i.e. individual samples with a large fluctuation, figure in the total fluctuation, which in turn determines how a region is split up. As suggested by its name, **Flatness** should be chosen large for ‘flat’ integrands and small for ‘volatile’ integrands with high peaks. Note that since **Flatness** appears in the exponent, one should not use too large values (say, no more than a few hundred) lest terms be truncated internally to prevent overflow.

7.3.4 Divonne-specific Options

- **Key1** -> 47, an integer which governs sampling in the partitioning phase:
Key1 = 7, 9, 11, 13 selects the cubature rule of degree **Key1**. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.
For other values of **Key1**, a quasi-random sample of $n_1 = |\mathbf{Key1}|$ points is used, where the sign of **Key1** determines the type of sample,
 - **Key1** > 0, use a Korobov quasi-random sample,
 - **Key1** < 0, use a Sobol quasi-random sample.
- **Key2** -> 1, an integer which governs sampling in the final integration phase:
Key2 = 7, 9, 11, 13 selects the cubature rule of degree **Key2**. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.
For other values of **Key2**, a quasi-random sample is used, where the sign of **Key2** determines the type of sample,
 - **Key2** > 0, use a Korobov quasi-random sample,
 - **Key2** < 0, use a Sobol quasi-random sample,and $n_2 = |\mathbf{Key2}|$ determines the number of points,
 - $n_2 \geq 40$, sample n_2 points,

- $n_2 < 40$, sample $n_2 n_{\text{need}}$ points, where n_{need} is the number of points needed to reach the prescribed accuracy, as estimated by Divonne from the results of the partitioning phase.
- **Key3** \rightarrow 1, an integer which sets the strategy for the refinement phase:
Key3 = 0, do not treat the subregion any further.
Key3 = 1, split the subregion up once more.
Otherwise, the subregion is sampled a third time with **Key3** specifying the sampling parameters exactly as **Key2** above.
- **MaxPass** \rightarrow 5, the number of passes after which the partitioning phase terminates. The partitioning phase terminates when the estimated total number of integrand evaluations (partitioning plus final integration) does not decrease for **MaxPass** successive iterations.
A decrease in points generally indicates that Divonne discovered new structures of the integrand and was able to find a more effective partitioning. **MaxPass** can be understood as the number of ‘safety’ iterations that are performed before the partition is accepted as final and counting consequently restarts at zero whenever new structures are found.
- **Border** \rightarrow 0, the width of the border of the integration region. Points falling into this border region are not sampled directly, but are extrapolated from two samples from the interior. Use a nonzero **Border** if the integrand function cannot produce values directly on the integration boundary.
The border width always refers to the unit hypercube, i.e. it is not rescaled if the integration region is not the unit hypercube.
- **MaxChisq** \rightarrow 10, the maximum χ^2 value a single subregion is allowed to have in the final integration phase. Regions which fail this χ^2 test and whose sample averages differ by more than **MinDeviation** move on to the refinement phase.
- **MinDeviation** \rightarrow .25, a bound, given as the fraction of the requested error of the entire integral, which determines whether it is worthwhile further examining a region that failed the χ^2 test. Only if the two sampling averages obtained for the region differ by more than this bound is the region further treated.
- **Given** \rightarrow {}, a list of points where the integrand might have peaks. A point is a list of n_d real numbers, where n_d is the dimension of the integral.
Divonne will consider these points when partitioning the integration region. The idea here is to help the integrator find the extrema of the integrand in the presence of very narrow peaks. Even if only the approximate location of such peaks is known, this can considerably speed up convergence.

- **NExtra** \rightarrow 0, the maximum number of points that will be considered in the output of the **PeakFinder** function.
- **PeakFinder** \rightarrow (**{}**&), the peak-finder function. This function is called whenever a region is up for subdivision and is supposed to point out possible peaks lying in the region, thus acting as the dynamic counterpart of the static list of points supplied with **Given**. It is invoked with two arguments, the multidimensional equivalents of the lower left and upper right corners of the region being investigated, and must return a (possibly empty) list of points. A point is a list of n_d real numbers, where n_d is the dimension of the integral.

7.3.5 Cuhre-specific Options

- **Key** \rightarrow 0, chooses the basic integration rule:

Key = 7, 9, 11, 13 selects the cubature rule of degree **Key**. Note that the degree-11 rule is available only in 3 dimensions, the degree-13 rule only in 2 dimensions.

For other values, the default rule is taken, which is the degree-13 rule in 2 dimensions, the degree-11 rule in 3 dimensions, and the degree-9 rule otherwise.

8 Tests and Comparisons

Four integration routines may seem three too many, but as the following tests show, all have their strengths and weaknesses. Fine-tuning the algorithm parameters can also significantly affect performance.

In the following, the test suite of Genz [12] is used. Rather than testing individual integrands, Genz proposes the following six families of integrands:

1. Oscillatory: $f_1(\mathbf{x}) = \cos(\mathbf{c} \cdot \mathbf{x} + 2\pi w_1),$
 2. Product peak: $f_2(\mathbf{x}) = \prod_{i=1}^{n_d} \frac{1}{(x_i - w_i)^2 + c_i^{-2}},$
 3. Corner peak: $f_3(\mathbf{x}) = \frac{1}{(1 + \mathbf{c} \cdot \mathbf{x})^{n_d+1}},$
 4. Gaussian: $f_4(\mathbf{x}) = \exp(-\mathbf{c}^2(\mathbf{x} - \mathbf{w})^2),$
 5. C^0 -continuous: $f_5(\mathbf{x}) = \exp(-\mathbf{c} \cdot |\mathbf{x} - \mathbf{w}|),$
 6. Discontinuous: $f_6(\mathbf{x}) = \begin{cases} 0 & \text{for } x_1 > w_1 \vee x_2 > w_2, \\ \exp(\mathbf{c} \cdot \mathbf{x}) & \text{otherwise.} \end{cases}$
- (3)

Parameters designated with w are non-affective, they vary e.g. the location of peaks, but should in principle not affect the difficulty of the integral.

Parameters designated with c are affective and in a sense “define” the difficulty of the integral, e.g. the width of peaks are of this kind. The c_i are positive and the difficulty increases with $\|\mathbf{c}\|_1 = \sum_{i=1}^{n_d} c_i$.

The testing procedure is thus: Choose uniform random numbers from $[0, 1)$ for the c_i and w_i . Renormalize \mathbf{c} for a given difficulty. Run the algorithms with the integrands thus determined. Repeat this procedure 20 times and take the average.

For comparison, Mathematica’s **NIntegrate** function was included in the test. Unfortunately, when a maximum number of samples is prescribed, **NIntegrate** invariably uses non-adaptive methods, by default the Halton–Hammersley–Wozniakowski quasi-Monte Carlo algorithm. The comparison may thus seem not quite balanced, but this is not entirely true: Lacking an upper bound on the number of integrand evaluations, **NIntegrate**’s adaptive method in some cases ‘locks up’ (spends an inordinate amount of time and samples) and the user can at most abort a running calculation, but not extract a preliminary result. The adaptive method could reasonably be used only for some of the integrand families in the test, and it was felt that such a selection should not be done, as the comparisons should in the first place give an idea about the *average* performance of the integration methods, without any fine-tuning.

Table 1 gives the results of the tests as described above. This comparison chart should be interpreted with care, however, and serves only as a rough measure of the performance of the integration methods. Many integrands appearing in actual calculations bear few or no similarities with the integrand families tested here, and neither have the integration parameters been tuned to ‘get the most’ out of each method.

The Mathematica code of the test suite is included in the downloadable CUBA package.

9 Summary

The CUBA library offers a choice of four independent routines for multidimensional numerical integration: Vegas, Suave, Divonne, and Cuhre. They work by very different methods, summarized in the following table:

Routine	Basic integration method	Algorithm type	Variance reduction
Vegas	Sobol quasi-random sample	Monte Carlo	importance sampling
Suave	Sobol quasi-random sample	Monte Carlo	globally adaptive subdivision
Divonne	Korobov quasi-random sample	Monte Carlo	stratified sampling,
	<i>or</i> Sobol quasi-random sample	Monte Carlo	aided by methods from
	<i>or</i> cubature rules	deterministic	numerical optimization
Cuhre	cubature rules	deterministic	globally adaptive subdivision

All four have a C/C++, Fortran, and Mathematica interface and can integrate vector integrands. Their invocation is very similar, so it is easy to substitute one method by another for cross-checking. For further safeguarding, the output is supplemented by a χ^2 probability which quantifies the reliability of the error estimate.

$n_d = 5$

j	Vegas	Suave	Divonne	Cuhre	NIntegrate
1	162000 \pm 0	127300 \pm 32371	21313 \pm 11039	819 \pm 0	218281 \pm 0
2	11750 \pm 1795	13500 \pm 1539	17353 \pm 3743	56238 \pm 40917	218281 \pm 0
3	16125 \pm 2411	11500 \pm 1000	17208 \pm 2517	1174 \pm 444	218281 \pm 0
4	56975 \pm 11372	20100 \pm 4745	19636 \pm 6159	22577 \pm 31424	218281 \pm 0
5	14600 \pm 3085	15250 \pm 2337	21675 \pm 4697	150423 \pm 0	218281 \pm 0
6	19750 \pm 4999	23850 \pm 2700	39694 \pm 14001	1884 \pm 215	218281 \pm 0

$n_d = 8$

j	Vegas	Suave	Divonne	Cuhre	NIntegrate
1	153325 \pm 20274	124350 \pm 35467	28463 \pm 31646	3315 \pm 0	212939 \pm 13557
2	12650 \pm 1987	21050 \pm 4594	22030 \pm 3041	91826 \pm 58513	218281 \pm 0
3	24325 \pm 3753	29350 \pm 3588	67104 \pm 16906	18785 \pm 22354	218281 \pm 0
4	38575 \pm 16169	29250 \pm 8873	24849 \pm 5015	62322 \pm 44328	218281 \pm 0
5	15150 \pm 2616	25500 \pm 6444	32885 \pm 5945	151385 \pm 0	218281 \pm 0
6	18875 \pm 2512	40900 \pm 7196	116744 \pm 32533	9724 \pm 9151	218281 \pm 0

$n_d = 10$

j	Vegas	Suave	Divonne	Cuhre	NIntegrate
1	156050 \pm 21549	129800 \pm 30595	32176 \pm 30424	7815 \pm 0	214596 \pm 16481
2	14175 \pm 2672	24800 \pm 5464	25684 \pm 7582	144056 \pm 25983	218281 \pm 0
3	30275 \pm 6296	51150 \pm 15608	139737 \pm 18505	109150 \pm 58224	218281 \pm 0
4	29475 \pm 10277	34050 \pm 10200	27385 \pm 8498	105763 \pm 49789	218281 \pm 0
5	16150 \pm 2791	31400 \pm 7715	44393 \pm 18654	153695 \pm 0	218281 \pm 0
6	22100 \pm 3085	74900 \pm 32203	136508 \pm 17067	73200 \pm 64621	218281 \pm 0

Test parameters:

- number of dimensions: $n_d = 5, 8, 10$,
- requested relative accuracy: $\varepsilon_{\text{rel}} = 10^{-3}$,
- maximum number of samples: $n_s^{\text{max}} = 150000$,
- integrand difficulties: $\frac{\text{Integrand family } j}{\|\mathbf{c}_j\|_1} \begin{array}{|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 \\ \hline 6.0 & 18.0 & 2.2 & 15.2 & 16.1 & 16.4 \\ \hline \end{array}$

Table 1: The number of samples used, averaged from 20 randomly chosen integrands from each integrand family j defined in Eq. (3). Values in the vicinity of n_s^{max} generally indicate failure to converge. **NIntegrate** seems not to be able to stop at around the limit of **MaxPoints** $\rightarrow n_s^{\text{max}}$, but always samples considerably more points.

The source code is available from <http://www.feynarts.de/cuba> and compiles with gcc, the GNU C compiler. The C functions can be called from Fortran directly, so there is no need for adapter code. Similarly, linking Fortran code with the library is straightforward and requires no extra tools.

A parallel version of CUBA is in preparation, in particular Divonne has strategically very obvious points for parallelization.

The routines in the CUBA library have all been carefully tested, but it would of course be folly to believe they are completely error-free. The author welcomes any kind of feedback, in particular bug and performance reports, at hahn@feynarts.de.

Acknowledgements

I thank A. Hoang for involving me in a discussion out of which the concept of the Mathematica interface was born and T. Fritzsche, M. Rauch, and A.M. de la Ossa for testing.

References

- [1] R. Piessens, E. de Doncker, C. Überhuber, D. Kahaner, *QUADPACK* – a subroutine package for automatic integration, Springer-Verlag, 1983.
- [2] G.P. Lepage, *J. Comp. Phys.* **27** (1978) 192.
- [3] G.P. Lepage, Report CLNS-80/447, Cornell Univ., Ithaca, N.Y., 1980.
- [4] W.H. Press, G.R. Farrar, *Comp. in Phys.* **4** (1990) 190.
- [5] J.H. Friedman, M.H. Wright, *ACM Trans. Math. Software* **7** (1981) 76;
J.H. Friedman, M.H. Wright, SLAC Report CGTM-193-REV, CGTM-193, 1981.
- [6] J. Berntsen, T. Espelid, A. Genz, *ACM Trans. Math. Software* **17** (1991) 437;
J. Berntsen, T. Espelid, A. Genz, *ACM Trans. Math. Software* **17** (1991) 452;
TOMS algorithm 698.
- [7] P. Bratley, B.L. Fox, *ACM Trans. Math. Software* **14** (1988) 88;
TOMS algorithm 659.
- [8] H. Niederreiter, Random number generation and quasi-Monte Carlo methods, SIAM, 1992.
- [9] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical recipes in Fortran, 2nd edition, Cambridge University Press, 1992.

- [10] N.M. Korobov, Number theoretic methods in approximate analysis (in Russian), Fizmatgiz, Moscow, 1963.
A comprehensive English reference on the topic of good lattice points (of which the Korobov points are a special case) is H.L. Keng, W. Yuan, Applications of number theory to numerical analysis, Springer-Verlag, 1981.
- [11] A. Genz, A. Malik, *SIAM J. Numer. Anal.* **20** (1983) 580.
- [12] A. Genz, A package for testing multiple integration subroutines, in: P. Keast, G. Fairweather (eds.), Numerical Integration, Kluwer, Dordrecht, 1986.