Highly Parallel, High-Precision Numerical Integration

D.H. Bailey[*]   J.M. Borwein[†]

April 1, 2008

Abstract

This paper describes schemes for rapidly computing numerical values of definite integrals to very high accuracy (hundreds to thousands of digits) on highly parallel computer systems. Such schemes are of interest not only in computational physics and computational chemistry, but also in experimental mathematics, where high-precision numerical values of definite integrals can be used to numerically discover new identities. This paper presents performance results for 1-D and 2-D integral test suites on highly parallel computer systems. Results are also given for certain problems that derive from mathematical physics. One of these results confirms a conjecture to 20,000 digit accuracy. The performance rate for this calculation is 690 Gflop/s on 1024 CPUs of a state-of-the-art parallel system. Other results, which range in precision from 120 to 500 digits, and for 1-D, 2-D, 3-D and 4-D integrals, derive from Ising theory. The largest of these calculations required 28 hours on 256 CPUs. We believe that these are the first instances of evaluations of nontrivial 3-D and 4-D integrals to multi-hundred-digit accuracy.

1. Introduction

Numerical integration (often termed "numerical quadrature") has numerous applications in applied mathematics, particularly in fields such as mathematical physics and computational chemistry. Recently such techniques have found application in the emerging discipline of experimental mathematics, namely the application of high-performance computing to research questions in mathematics. In particular, high-precision numerical values of certain definite integrals, when combined with integer relation detection algorithms, can be used to discover previously unknown analytic evaluations (i.e., closed-form formulas)

for certain integrals, and to provide strong numerical confirmation that such computer-discovered identities are valid.

An "integer relation detection" scheme is a numerical algorithm which, given an n-long vector (x_i) of high-precision floating-point values, can recover the integer coefficients (a_i) such that $a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = 0$ (to within available precision), or else determine that there are no such integers less than a certain size. The precision required in these computations is typically 100–500 digits, although occasionally thousands of digits are required. In one instance, 50,000-digit arithmetic was needed to find the underlying relation [9]. The best-known integer relation algorithm is the "PSLQ" algorithm [8].

Here are some examples of this approach. In 2002, the present authors and Greg Fee of Canada were inspired by a recent problem in the *American Mathematical Monthly* [2]. We found, by using an early version of the integration scheme described in this paper, together with a PSLQ program, that if $Q(a)$ is defined by

$$Q(a) = \int_0^1 \frac{\arctan(\sqrt{x^2 + a^2})\, dx}{\sqrt{x^2 + a^2}(x^2 + 1)},$$

then

$$
\begin{aligned}
Q(0) &= \pi \log 2/8 + G/2 \\
Q(1) &= \pi/4 - \pi\sqrt{2}/2 + 3\sqrt{2}\arctan(\sqrt{2})/2 \\
Q(\sqrt{2}) &= 5\pi^2/96.
\end{aligned}
$$

Here $G = \sum_{k \geq 0} (-1)^k/(2k+1)^2 = 0.91596559417\ldots$ is Catalan's constant. These specific experimental results then led to the following general result, which now has been rigorously established, and several others:

$$\int_0^\infty \frac{\arctan(\sqrt{x^2 + a^2})\, dx}{\sqrt{x^2 + a^2}(x^2 + 1)} = \frac{\pi}{2\sqrt{a^2 - 1}}\left[2\arctan(\sqrt{a^2 - 1}) - \arctan(\sqrt{a^4 - 1})\right].$$

As a second example, the present authors empirically determined that

$$
\begin{aligned}
\frac{2}{\sqrt{3}} \int_0^1 \frac{\log^6(x)\arctan[x\sqrt{3}/(x-2)]}{x+1}\, dx = {}& \frac{1}{81648}\big[-229635 L_3(8) \\
& + 29852550 L_3(7)\log 3 - 1632960 L_3(6)\pi^2 + 27760320 L_3(5)\zeta(3) \\
& - 275184 L_3(4)\pi^4 + 36288000 L_3(3)\zeta(5) - 30008 L_3(2)\pi^6 \\
& - 57030120 L_3(1)\zeta(7)\big],
\end{aligned}
$$

where $L_3(s) = \sum_{n \geq 1}\left[1/(3n-2)^s - 1/(3n-1)^s\right]$ and where $\zeta(s) = \sum_{n \geq 1} 1/n^s$ is the Riemann zeta function. Based on these experimental results, general results of this type have been conjectured but few have yet been rigorously established.

2

As a third example, one of the present authors (Borwein) and British physicist David Broadhurst conjectured that

$$\frac{24}{7\sqrt{7}} \int_{\pi/3}^{\pi/2} \log \left| \frac{\tan t + \sqrt{7}}{\tan t - \sqrt{7}} \right| dt \overset{?}{=} L_{-7}(2) =$$ (1)

$$\sum_{n=0}^{\infty} \left[\frac{1}{(7n+1)^2} + \frac{1}{(7n+2)^2} - \frac{1}{(7n+3)^2} + \frac{1}{(7n+4)^2} - \frac{1}{(7n+5)^2} - \frac{1}{(7n+6)^2} \right].$$

In this paper we will describe how we strongly confirmed this identity, which originates in quantum theory, by evaluating the left-hand integral to 20,000-digit accuracy and showing it to be in full agreement with a *Mathematica* evaluation of the right-hand side to the same precision.

As a final example, in some research performed jointly with Richard Crandall, the present authors examined the following integrals, which arise in the Ising theory of mathematical physics [5]:

$$C_n := \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{1}{\left(\sum_{j=1}^n (u_j + 1/u_j) \right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}.$$ (2)

We were able to show that these can be written as one-dimensional integrals:

$$C_n = \frac{2^n}{n!} \int_0^\infty p K_0^n(p) \, dp,$$ (3)

where K_0 is the *modified Bessel function* [1]

$$K_0(p) := \int_0^\infty e^{-p \cosh t} dt.$$

When we computed C_n for fairly large n, e.g.

$$C_{1024} = 0.63047350337438679612204019271087890435458707871273234\ldots$$

we found that these rather quickly approached a limit, numerically close to $2e^{-2\gamma}$ (they agree to over 300 places). Indeed, we were then able to prove that

$$\lim_{n \to \infty} C_n = 2e^{-2\gamma}.$$

The above examples are ordinary one-dimensional integrals, but 2-D, 3-D and 4-D integrals are also of interest. Along this line, in the same study of Ising integrals, the present authors and Richard Crandall determined that if

$$E_n := 2 \int_0^1 \cdots \int_0^1 \prod_{i<j} \left(\frac{u_i - u_j}{u_i + u_j} \right)^2 \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}.$$ (4)

3

then after the evaluation of an equivalent (but *much* more complicated) triple integral to 250-digit accuracy, we found that

$$E_5 \overset{?}{=} 42 - 1984 \operatorname{Li}_4(1/2) + \frac{189\pi^4}{10} - 74\zeta(3) - 1272\zeta(3)\log 2 + 40\pi^2\log^2 2$$
$$-\frac{62\pi^2}{3} + \frac{40\pi^2\log 2}{3} + 88\log^4 2 + 464\log^2 2 - 40\log 2,$$

where the notation $\overset{?}{=}$ is employed to emphasize that we do not yet have a formal proof for this evaluation. See [7] for some additional examples.

2. The Tanh-Sinh Quadrature Algorithm

In a previous paper [11], one of the present authors and two co-authors investigated several numerical integration schemes, suitable for high-precision usage, and exhibited results for computer runs with 400- and 1000-digit precision. We concluded the "tanh-sinh" quadrature scheme holds the best promise for very high-precision usage, particularly for functions with infinite derivatives or blow-up singularities at endpoints, although Gaussian quadrature also has merit in those cases where the integrand function is well-behaved in the full closed interval.

The tanh-sinh scheme is based on the observation, rooted in the Euler-Maclaurin summation formula, that for certain bell-shaped integrands, a simple block-function approximation to the integral is remarkably accurate [3, pg. 180]. This principle is exploited in the tanh-sinh scheme by transforming an integral of a given function $f(x)$ on a finite interval such as $[-1, 1]$ to an integral on $(-\infty, \infty)$, by using the change of variable $x = g(t)$, where $g(t) = \tanh(\pi/2 \cdot \sinh t)$. The function $g(t)$ has the property that $g(x) \to 1$ as $x \to \infty$ and $g(x) \to -1$ as $x \to -\infty$, and also that $g'(x)$ and all higher derivatives rapidly approach zero for large positive and negative arguments. Thus one can write, for $h > 0$,

$$\int_{-1}^{1} f(x)\,dx \;\; = \;\; \int_{-\infty}^{\infty} f(g(t))g'(t)\,dt \;\; = \;\; h\sum_{j=-\infty}^{\infty} w_j f(x_j) + E(h),$$

where $x_j = g(hj)$ and $w_j = g'(hj)$ and $E(h)$ is an error term. In many cases, even where $f(x)$ has an infinite derivative or an integrable singularity at one or both endpoints, the resulting integrand $f(g(t))g'(t)$ is a smooth bell-shaped function for which the Euler-Maclaurin argument applies. In these cases, the error $E(h)$ decreases very rapidly with h.

In short, one can approximate the integral of $f(t)$ on $[-1, 1]$ by the sum $\sum_{j=-N}^{N} w_j f(x_j)$, where the abscissas are given by $x_j = g(hj)$ and the weights are given by $w_j = g'(hj)$. In our implementation, the parameter h is set to 2^{-k}, where k is the "level" of the quadrature calculation, and N is chosen large enough that terms beyond N are smaller than the "epsilon" of the arithmetic precision being used. Successively larger levels reduce h in half, double the number of abscissa-weight pairs (and thus double the number of function evaluations required in a quadrature calculation), but also approximately double

4

the number of correct digits in the result, in many cases. The abscissa-weight pairs of one level are the even-indexed pairs for the next level. Full details are given in [11]. The basic tanh-sinh integration scheme was first introduced by Takahasi and Mori [17].

3. Error Estimation

We should add that some highly accurate error estimation schemes have recently been established for tanh-sinh quadrature, estimates that can be calculated in a very similar manner to the quadrature result itself [4]. For example,

$$E_2(h, m) \quad := \quad h(-1)^{m-1} \left(\frac{h}{2\pi} \right)^{2m} \sum_{j=a/h}^{b/h} D^{2m}[g'(t)f(g(t))](jh) \qquad (5)$$

yields extremely accurate estimates of the error, even in the simplest case $m = 1$. Further, these error estimates can be used to obtain rigorous certificates on the values of quadrature results. For example, by using these methods we were able to establish rigorously that the experimental identity

$$\frac{24}{7\sqrt{7}} \int_{\pi/3}^{\pi/2} \log \left| \frac{\tan t + \sqrt{7}}{\tan t - \sqrt{7}} \right| dt \quad \stackrel{?}{=} \quad L_{-7}(2),$$

which was mentioned in the introduction, holds to within 3.82×10^{-49}.

Several other examples of such certificates are given in [4]. These certificate calculations are typically rather expensive, but the highly parallel implementation techniques described here can be used to obtain these results in reasonable run time.

4. High-Precision Arithmetic

The Arbitrary Precision (ARPREC) computation library was used to perform the required high-precision arithmetic computations described in this paper [10]. This software library is written in C++, but it includes both C++ and Fortran-90 translation modules, so that existing C++ and Fortran-90 application programs can utilize this library by making only very minor changes to the source code. In most cases, it is only necessary to change type statements and input/output statements of the variables that one wishes to be treated as arbitrary precision, and all other operations are automatically performed by the library. One fortunate feature of high-precision numerical quadrature, as described in this paper, is that individual high-precision arithmetic operations and transcendental function evaluations can be performed locally on a single processor. Thus it is not necessary to invoke parallel processing within the ARPREC library itself, at least for the problems considered below.

5. 1-D Parallel Implementation

The tanh-sinh quadrature scheme is very well suited for implementation on a highly parallel computer system, using either a MPI or OpenMP programming model, although there are several details that must be observed to avoid major reductions in parallel performance.

Two approaches were considered for a parallel implementation: (1) computing abscissa-weight pairs in parallel, distributing all resulting pairs to all processors, and then parceling out function evaluations to processors in some evenly distributed manner; and (2) calculating the abscissa-weight pair indexed j only on processor $\rho(j)$ (for some processor assignment function $\rho(j)$), and performing only those function evaluations corresponding to pairs indexed j on processor $\rho(j)$. (Some additional details, such as the error estimation procedure, are not mentioned in this brief sketch of the parallelization strategy.) Either way, function-weight products are summed locally, then these sums are combined onto a single processor and added together using high-precision arithmetic to produce a global sum. Option (1) requires substantial interprocessor communication, and in the tests below is scalable only to about 256 processors. Option (2) requires much less communication and only a modest amount of memory on each node, but suffers from severe load imbalances (and reductions in scalability) unless the processor assignment function $\rho(j)$ is chosen carefully.

One straightforward processor assignment scheme is a simple cyclic scheme, namely $\rho(j) = j \bmod m$, where m is the total number of processors. One difficulty with this scheme derives from the fact that different integration problems require different numbers of abscissa-weight pairs to achieve a given accuracy target. Even among the problems described in the next section, some achieve full 2000-digit accuracy with only nine levels of abscissa-weight pairs ($h = 2^{-9}$), while others require 13 levels of abscissa-weight pairs ($h = 2^{-13}$), which means 16 times as many pairs and 16 times as many function evaluations. But if one has pre-computed 13 levels of abscissa-weight pairs, then when only nine levels are used, the abscissa-weight array is accessed with a stride (interval) of 16. Such power-of-two strides result in catastrophic load imbalances when a cyclic assignment function is employed—some processors have a large fraction of the pairs and corresponding function evaluations, while other processors literally have none.

Table 1 gives the results of tests of five different assignment schemes. In these tests, 70,000 indices (the approximate number of abscissa-weight pairs actually generated in the quadrature computations described in the Section 7) were assigned to processors according to the five schemes. The smallest and largest number of indices assigned to any processor by a given scheme, for various processor numbers and strides, are shown in the table. The more nearly equal these max and min figures are, the better the assignment scheme. The five schemes are:

1. CYC, a cyclic scheme: $\rho(j) = \mathrm{mod}(j, m)$

2. BC1, a block-cyclic scheme: $\rho(j) = \mathrm{mod}(\lfloor j/16 \rfloor, m)$

3. BC2, a block-cyclic scheme: $\rho(j) = \mathrm{mod}(\lfloor j/17 \rfloor, m)$

4. MCBC, a mixed cyclic, block-cyclic scheme: $\rho(j) = \mathrm{mod}(j + \lfloor j/16 \rfloor, m)$

5. RAND, a pseudo-random scheme: $\rho(j) = \lfloor z_j m \rfloor$, where z_j is a uniform generator on $(0, 1)$.

		CYC		BC1		BC2		MCBC		RAND	
Proc.	Stride	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max
16	1	4375	4375	4368	4384	4369	4386	4374	4376	4267	4505
16	2	0	4375	2184	2192	2056	2322	2187	2188	2133	2247
16	4	0	4375	1092	1096	1028	1290	1093	1094	1037	1135
16	8	0	4375	546	548	514	774	546	547	500	596
16	16	0	4375	273	274	257	516	273	274	246	305
64	1	1093	1094	1088	1104	1088	1105	1093	1095	1022	1169
64	2	0	1094	544	552	512	585	545	548	505	600
64	4	0	1094	272	276	256	325	272	276	242	315
64	8	0	1094	136	138	128	195	136	138	112	176
64	16	0	1094	68	69	64	130	68	69	48	87
256	1	273	274	272	288	272	289	272	274	234	320
256	2	0	274	136	144	128	153	136	137	106	168
256	4	0	274	68	72	64	85	68	69	45	93
256	8	0	274	34	36	32	51	34	35	17	51
256	16	0	274	17	18	16	34	17	18	6	29
1024	1	68	69	64	80	68	85	67	69	42	97
1024	2	0	69	32	40	32	45	32	36	16	54
1024	4	0	69	16	20	16	25	16	20	6	30
1024	8	0	69	8	10	8	15	8	10	0	19
1024	16	0	69	4	5	4	10	4	5	0	12

Table 1: Min/max processor counts for five assignment functions (70,000 indices)

It can be seen from the results in Table 1 that of the five schemes mentioned, the one named "mixed cyclic, block-cyclic" (MCBC) is the best. It provides a virtually perfect load balance across a large range of processors (up to 1024) and strides (up to stride 16). This is the scheme that was used in the computations described below.

6. 1-D Test Problems

The following 14 integrals are taken from the suite used in the earlier paper [11]. They are typical of the integrals that have been encountered in experimental math research, except that in each of these cases an analytic result is known, as shown below, facilitating the checking of results:

- 1–4: Continuous functions on finite intervals.

- 5–6: Continuous functions on finite intervals, but with an infinite derivative at an endpoint.

- 7–10: Functions on finite intervals with an integrable singularity at an endpoint.

- 11–13: Functions on an infinite interval.

- 14: An oscillatory function on an infinite interval.

$$1 \quad : \quad \int_0^1 t \log(1 + t)\, dt \;=\; 1/4$$

$$2 \quad : \quad \int_0^1 t^2 \arctan t\, dt \;=\; (\pi - 2 + 2\log 2)/12$$

$$3 \quad : \quad \int_0^{\pi/2} e^t \cos t\, dt \;=\; (e^{\pi/2} - 1)/2$$

$$4 \quad : \quad \int_0^1 \frac{\arctan(\sqrt{2 + t^2})}{(1 + t^2)\sqrt{2 + t^2}}\, dt \;=\; 5\pi^2/96$$

$$5 \quad : \quad \int_0^1 \sqrt{t} \log t\, dt \;=\; -4/9$$

$$6 \quad : \quad \int_0^1 \sqrt{1 - t^2}\, dt \;=\; \pi/4$$

$$7 \quad : \quad \int_0^1 \frac{\sqrt{t}}{\sqrt{1 - t^2}}\, dt \;=\; 2\sqrt{\pi}\,\Gamma(3/4)/\Gamma(1/4)$$

$$8 \quad : \quad \int_0^1 \log^2 t\, dt \;=\; 2$$

$$9 \quad : \quad \int_0^{\pi/2} \log(\cos t)\, dt \;=\; -\pi \log(2)/2$$

$$10 \quad : \quad \int_0^{\pi/2} \sqrt{\tan t}\, dt \;=\; \pi\sqrt{2}/2$$

$$11 \quad : \quad \int_0^\infty \frac{1}{1 + t^2}\, dt \;=\; \pi/2$$

$$12 \quad : \quad \int_0^\infty \frac{e^{-t}}{\sqrt{t}}\, dt \;=\; \sqrt{\pi}$$

$$13 \quad : \quad \int_0^\infty e^{-t^2/2}\, dt \;=\; \sqrt{\pi/2}$$

$$14 \quad : \quad \int_0^\infty e^{-t} \cos t\, dt \;=\; 1/2$$

7. 1-D Performance Results

The results of the 1-D parallel quadrature tests are given in Table 2. The first line gives the run time, in seconds, for the initialization process (calculating all abscissa-weight pairs). The second column gives the levels of abscissa-weight pairs (see Sections 2 and 4) required to achieve the target accuracy (10^{-2000}) for the individual problems. The corresponding number of abscissa-weight pairs is roughly 8.7×2^k, where k is the number of levels shown in the table. The target accuracy was achieved in each problem, except in Problem 14 where the accuracy was 10^{-1972}. When 14 levels are used, or if a slightly smaller value of h is used, the error target is achieved here also.

8

Problem	Levels	Processors					
Number	Required	1	4	16	64	256	1024
Init		4329.04	1085.34	271.87	68.88	17.73	5.02
1	10	480.07	101.63	25.55	6.45	1.65	0.53
2	10	1403.63	294.32	74.04	18.83	4.99	1.54
3	10	1421.99	317.01	79.69	20.42	5.24	1.83
4	10	1553.24	328.73	82.13	20.84	5.52	1.63
5	9	236.68	51.62	12.90	3.30	0.93	0.30
6	10	26.31	5.62	1.42	0.36	0.11	0.05
7	10	52.65	11.46	2.87	0.72	0.20	0.10
8	9	234.06	50.98	12.85	3.26	0.90	0.27
9	10	1552.38	333.24	83.60	21.34	5.44	1.84
10	10	1138.78	245.45	61.39	15.73	3.99	1.44
11	11	25.30	5.17	1.30	0.33	0.09	0.04
12	12	655.03	161.99	40.71	10.20	2.65	0.80
13	13	871.99	216.50	54.13	13.65	3.52	0.97
14	13	8291.43	1826.02	457.02	114.84	29.48	7.87
Total		22272.58	5035.08	1261.47	319.15	82.44	24.23
Speedup		1.00	4.42	17.66	69.79	270.17	919.22

Table 2: Parallel run times (in seconds) and speedup ratios for 1-D problems

These runs were made on "System X," an Apple G5-based parallel supercomputer at Virginia Technical University, using the IBM xlf90 Fortran-90 compiler (for the parallel quadrature application program) and the IBM xlC C++ compiler (for the ARPREC library). The parallel program performs all functions of the single processor code, including the calculation of an estimated error in the result [11]. The one-processor timings shown in Table 2 are for an efficient single-processor version of this program, with no parallel constructs. Thus the speedup ratios shown in the table are true comparisons to single-processor performance. Note that these are not scaled speedup figures—each run, including the 1-CPU run, is for the same full-sized problem.

Note also that these timings exhibit super-linear speedup up to 256 CPUs, but drop back a bit for 1024 CPUs. The reason for these super-linear speedups is that the array space required on individual processors, for large processor counts, drops to the level that all required data can be contained in first or second level cache, thus permitting faster execution times. Only at 1024 CPUs does the performance loss due to very high parallelism overcome the savings from favorable cache behavior.

8. Two-Dimensional Quadrature
The tanh-sinh scheme described above can be generalized to two or more di-

mensions. In particular, a 2-D iterated integral can be approximated as follows:

$$\int_{-1}^{1}\int_{-1}^{1} f(x, y)\, dx\, dy \;=\; \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(g(s), g(t)) g'(s) g'(t)\, ds\, dt$$

$$=\; h \sum_{k=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} w_j w_k f(x_j, x_k) + E(h),$$

where $g(t) = \tanh(\pi/2 \cdot \sinh(t))$ as in the 1-D case, and where x_j and w_j are the 1-D abscissas and weights. This same approach can easily be extended to numerically evaluate more general integrals of the form

$$\int_{a}^{b}\int_{c(y)}^{d(y)} f(x, y)\, dx\, dy.$$

As before, the Euler-Maclaurin formula asserts that for a certain class of functions $f(x, y)$, including many with infinite derivatives and blow-up singularities at the boundaries of the rectangle, the error $E(h)$ in the above approximation goes to zero faster than any power of h. As a result, 2-D tanh-sinh quadrature, like the 1-D version, often achieves quadratic convergence, wherein each additional level of abscissa-weight pairs yields twice as many correct digits in the result.

However, 2-D quadrature inherently is much more expensive than 1-D quadrature, because the number of function evaluations in a 2-D array, assuming the same overall spacing, is many times larger than in a 1-D problem. Millions of function evaluations may be required to obtain, say, 100-digit accuracy in the result. Also, it has been found that 2-D tanh-sinh scheme is more sensitive to anomalies such infinite derivatives or blow-up singularities at boundaries. In such cases, each additional level typically yields only about 1.4 times as many correct digits, whereas in 1-D quadrature, problems with similar anomalies typically exhibit quadratic convergence (each additional level approximately doubles the number of correct digits). What's more, in 2-D quadrature, each additional level quadruples the computational cost instead of merely doubling the cost, since four times as many function evaluations are required.

9. 2-D Parallel Implementation

The parallel implementation of the 2-D scheme again relies crucially on a carefully chosen scheme for allocating processors to the abscissa array for function evaluations. The program assigns a batch of 16 consecutively numbered processors to each column, and then assigns the function evaluations in this column among these 16 processors. In this way, the program exploits available parallelism in both dimensions. The particular assignment scheme used by the program is as follows: the 16 processors p that satisfy $\lfloor p/16 \rfloor = \mathrm{mod}(j + j/16, n/16)$ are assigned to column j of the 2-d array of abscissas. Then within column j, location (i, j) is assigned to the processor p that satisfies $\mathrm{mod}(p, 16) = \mathrm{mod}(i + i/16, 16)$. Note that both rows and columns employ a mixed cyclic, block-cyclic scheme, which provides an even load balance for function evaluations, yet avoids difficulties with power-of-two strides.

10. 2-D Test Problems

The serial and parallel implementations of this 2-D tanh-sinh quadrature scheme have been tested on a suite of eight test problems. Because of the much higher computational cost of 2-D quadrature, a more modest goal of 100-digit accuracy was established in these problems, using 120-digit working precision. As before, this set includes some rather difficult examples, including one problem with a non-differentiable point at a boundary (Problem 1), two problems with a blow-up singularity at a boundary (Problems 4 and 6), two problems where the inner integral boundary is not merely an interval but instead bounded by two functions (Problems 7 and 8), and one problem on an infinite interval (Problem 5).

$$1 \quad : \quad \int_0^1 \int_0^1 \sqrt{s^2 + t^2}\, ds\, dt \;=\; \sqrt{2}/3 - \log(2)/6 + \log(2 + \sqrt{2})/3$$

$$2 \quad : \quad \int_0^1 \int_0^1 \sqrt{1 + (s - t)^2}\, ds\, dt \;=\; -\sqrt{2}/3 - \log(\sqrt{2} - 1)/2 + \log(\sqrt{2} + 1)/2 + 2/3$$

$$3 \quad : \quad \int_{-1}^1 \int_{-1}^1 (1 + s^2 + t^2)^{-1/2}\, ds\, dt \;=\; 4\log(2 + \sqrt{3}) - 2\pi/3$$

$$4 \quad : \quad \int_0^\pi \int_0^\pi \log[2 - \cos s - \cos t]\, ds\, dt \;=\; 4\pi G - \pi^2 \log 2$$

$$5 \quad : \quad \int_0^\infty \int_0^\infty \sqrt{s^2 + st + t^2}\, e^{-s-t}\, ds\, dt \;=\; 1 + 3/4 \cdot \log 3$$

$$6 \quad : \quad \int_0^1 \int_0^1 (s + t)^{-1}[(1 - s)(1 - t)]^{-1/2}\, ds\, dt \;=\; 4G$$

$$7 \quad : \quad \int_0^1 \int_0^t (1 + s^2 + t^2)^{-1/2}\, ds\, dt \;=\; -\pi/12 - 1/2 \cdot \log 2 + \log(1 + \sqrt{3})$$

$$8 \quad : \quad \int_0^\pi \int_0^t (\cos s \sin t) e^{-s-t}\, ds\, dt \;=\; 1/4 \cdot (1 + e^{-\pi})$$

11. 2-D Performance Results

Performance results for the 2-D quadrature program are shown in Table 2. In each problem over 100-digit accuracy was achieved, except for Problems 4 and 6, where the errors were 10^{-86} and 10^{-80}, respectively. No results are shown in this table for four processors, since the parallel program assumes a minimum of 16 processors.

It is clear from these results that, unlike the 1-D case, there is a large difference in run times between well-behaved integrands and those with singularities at a corner or boundary. Those without such anomalies can be evaluated to over 100-digit accuracy with just six levels, requiring only a few minutes run time. For those problems that do exhibit such anomalies, nine levels are needed, requiring many more function evaluations. Indeed, for Problems 4 and 6, even nine levels of abscissa-weight pairs evidently were not sufficient—it appears that

Problem Number	Levels Required	Processors				
		1	16	64	256	1024
1	9	1246.26	96.42	24.66	7.05	3.33
2	6	19.03	1.52	0.46	0.27	0.73
3	7	82.79	6.56	1.91	0.64	1.17
4	9	15310.44	1194.52	305.11	81.88	24.40
5	9	2209.86	170.84	44.38	12.23	4.62
6	9	1552.87	120.86	30.80	8.67	3.37
7	6	21.79	1.72	0.54	0.28	0.73
8	6	113.04	8.90	2.87	1.08	1.51
Total		20556.08	1601.34	410.73	112.10	39.86
Speedup		1.00	12.84	50.05	183.37	515.71

Table 3: Parallel run times (in seconds) and speedup ratios for 2-D problems

one additional level would be required in each case to achieve over 100 digit accuracy, which would multiply the run times by a factor of four.

As in the 1-D case, the parallel run times are real-time measurements in seconds. Also, as in the 1-D case, the run times shown for 1 CPU are true single-processor runs without any parallel constructs. Thus the speedup ratios are true parallel speedups, and not "scaled speedups" or other ratios.

The parallel speedups for 2-D quadrature are not nearly as high as for 1-D quadrature, in part because features such as estimated error calculation are significantly more complicated than in the 1-D case, and also because it is more difficult to allocate tasks evenly in 2-D. Also, much of this reduction in scalability occurs in the shorter-running problems, whose modest computational work cannot be as efficiently distributed among 1024 processors.

12. Application: Confirmation of a Conjectured 1-D Identity
In the introduction, we briefly mentioned the conjectured identity

$$\frac{24}{7\sqrt{7}} \int_{\pi/3}^{\pi/2} \log \left| \frac{\tan t + \sqrt{7}}{\tan t - \sqrt{7}} \right| dt \quad \overset{?}{=} \quad L_{-7}(2) \quad = \tag{6}$$

$$\sum_{n=0}^{\infty} \left[\frac{1}{(7n+1)^2} + \frac{1}{(7n+2)^2} - \frac{1}{(7n+3)^2} + \frac{1}{(7n+4)^2} - \frac{1}{(7n+5)^2} - \frac{1}{(7n+6)^2} \right].$$

This integral arose out of some studies in quantum field theory, in analysis of the volume of ideal tetrahedra in hyperbolic space. Note that the integrand function has a nasty singularity at $t = \arctan(\sqrt{7})$ (see Figure 1). The question mark is used because no formal proof is yet known. We note that Richard Crandall [14] has observed that the right-hand expression $L_{-7}(2)$ is also given by the integral

$$L_{-7}(2) \quad = \quad -\int_0^1 \frac{1 + 2u + u^2 + 2u^3 + u^4}{1 + u + u^2 + u^3 + u^4 + u^5 + u^6} \log u \, du. \tag{7}$$

Figure 1: Integrand function with singularity

Because of the interest expressed by researchers in the above conjecture and some related conjectures [12], we decided to calculate the integral

$$\frac{24}{7\sqrt{7}} \int_{\pi/3}^{\pi/2} \log\left|\frac{\tan t + \sqrt{7}}{\tan t - \sqrt{7}}\right| dt \quad = \quad 1.15192547054449104710169239732054996\ldots$$

to 20,000-digit accuracy (which approaches the limits of presently feasible computation) and compare with a 20,000-digit evaluation of the six-term infinite series on the right-hand side of (6). This integral was evaluated by splitting it into two integrals, the first from $\pi/3$ to $\arctan(\sqrt{7})$, and the second from $\arctan(\sqrt{7})$ to $\pi/2$, and then applying the 1-D tanh-sinh scheme to each part. This test was successful—the numerical value of the integral on the left-hand side of (6) agrees with the numerical value of the six-term infinite series on the right-hand side of (6) to at least 19,995 digits. The infinite series in was evaluated in approximately five hours on a personal computer using *Mathematica*.

Efficiently performing any numerical computation to 20,000-digit accuracy requires advanced techniques. The key idea here is to note that high-precision multiplications can be evaluated using a linear convolution scheme, which in turn can be performed using fast Fourier transforms. High-precision divisions can then be done using Newton iterations. Quadratically convergent algorithms (where every iteration roughly doubles the number of correct digits) are known for most transcendental functions. These advanced algorithms are incorporated

13

CPUs	Init	Integral #1	Integral #2	Total	Speedup
1	*190013	*1534652	*1026692	*2751357	1.00
16	12266	101647	64720	178633	15.40
64	3022	24771	16586	44379	62.00
256	770	6333	4194	11297	243.55
1024	199	1536	1034	2769	993.63

Table 4: Parallel run times (in seconds) and speedup ratios for the 20,000-digit problem

in the ARPREC multi-precision package and are automatically invoked when a numeric precision of greater than about 2,000 digits is specified [10]. Aside from precision level, the parallel integration was performed exactly as described in Section 5. Thirteen levels (approximately 85,000) abscissa-weight pairs were required to achieve the target accuracy in this problem.

Performance results on the Virginia Tech system are shown in Table 4. Multi-CPU timings include barrier waits and communication operations, as before. The one-CPU timings (noted by asterisks) are the sums of individual process timings in a 64-CPU run, in which only local computation was timed. Based on tests on other problems and problem sizes, these summed timings are accurate estimates of the timings of a true 1-CPU run, which would have taken 32 days in this case. In these runs, as well as the others reported in the paper, timing variations of up to 6% have been noted in tests. The performance rate for the 1024-CPU run is 690 Gflop/s, based on a measurement of floating-point operation count done on the Seaborg system at LBNL, using a hardware performance monitoring tool.

As a separate test of our computer program, we also evaluated Crandall's integral (7), which, as noted above, is equal to $L_{-7}(2)$, or in other words the right-hand side of the conjectured identity (6). On 1024 CPUs, our program was able to evaluate this integral to 20,000 digits (which completely agrees with the results above) in 416 seconds runtime, including initialization.

We believe that these two evaluations are the highest-precision non-trivial numerical integrations performed to date.

13. Application: Identifying Ising Integrals

As mentioned above, the present authors, together with Richard Crandall, have recently investigated certain integrals that arise in the Ising theory of

mathematical physics [5]. The three classes we studied are:

$$C_n := \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{1}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n} \tag{8}$$

$$D_n := \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{\prod_{i<j} \left(\frac{u_i - u_j}{u_i + u_j}\right)^2}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n} \tag{9}$$

$$E_n := 2 \int_0^1 \cdots \int_0^1 \prod_{i<j} \left(\frac{u_i - u_j}{u_i + u_j}\right)^2 \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}. \tag{10}$$

As we noted above, C_n can be transformed to a fairly simple 1-D integral:

$$C_n = \frac{2^n}{n!} \int_0^\infty p K_0^n(p) \, dp,$$

where K_0 is the modified Bessel function [1]. We have computed C_n to 500-digit accuracy, for $2 \le n \le 8$, and for all powers of two up to $n = 1024$. These numerical values are available at [6]. As we noted in the introduction, this data alerted us to the fact that C_n converge to a limiting value, namely $2e^{-2\gamma}$. This and numerous other facts are established in [5]

Computing D_n and E_n to high precision is substantially more difficult, because there is no known transformation to a simple 1-D integral as in the case of C_n. However, we were able to reduce the dimension of these integrals by two, so that, for instance, D_5 and E_5 are given by tractable but very complicated triple integrals. In this way, we have been able to calculate D_n and E_n to precision levels between 120 and 500 digits for n up to six.

Needless to say, these computations are very demanding, both of hardware and software. Defining the integrand function for E_5, for instance, requires 60 lines of high-level computer code, even after some simplification. The integrand for D_6 requires 700 lines of code, most of which have over 60 characters per line. For these runs, where the integrands are well-behaved at boundaries, we were able to use multi-dimensional Gaussian quadrature. We could have used tanh-sinh quadrature here, but the run times would have been somewhat longer. The computer runs themselves were performed on the Bassi system, an IBM Power5-based parallel computer system at the Lawrence Berkeley National Laboratory, and on System X, an Apple G5-based parallel computer system at the Virginia Institute of Technology. The computation of D_5 to 500 digits required 18 hours on 256 CPUs; the computation of E_6 to 250 digits required 28 hours on 256 CPUs.

These computations lie at the edge of presently available numerical techniques and computing technology. Indeed, we are not aware of any other instance of a successful multi-dimensional quadrature of a nontrivial function to several-hundred-digit accuracy. Full details of the background and methodology are given in [5]. The numerical values that we have obtained are available at [6].

15

In at least one instance, the numerical value we obtained in our computation led directly to a conjectured evaluation. For instance, using PSLQ we discovered that

$$E_5 \stackrel{?}{=} 42 - 1984 \operatorname{Li}_4(1/2) + \frac{189\pi^4}{10} - 74\zeta(3) - 1272\zeta(3)\log 2 + 40\pi^2 \log^2 2$$
$$- \frac{62\pi^2}{3} + \frac{40\pi^2 \log 2}{3} + 88 \log^4 2 + 464 \log^2 2 - 40 \log 2.$$

As before, the notation $\stackrel{?}{=}$ is employed to emphasize that we do not yet have a formal proof for this evaluation. However, this experimental detection is quite strong—190 orders of magnitude beyond the level that could reasonably be ascribed to numerical round-off error or any other artifact.

References

[1] Milton Abramowitz and Irene A. Stegun, *Handbook of Mathematical Functions,* Dover, NY, 1970.

[2] Zafar Ahmed, "Definitely an Integral," *American Mathematical Monthly*, vol. 109 (2002), no. 7, pg. 670–671.

[3] Kendall E. Atkinson, *Elementary Numerical Analysis*, John Wiley and Sons, 1993.

[4] David H. Bailey and Jonathan M. Borwein, "Effective Error Bounds in Euler-Maclaurin-Based Quadrature Schemes," *Proceedings of the 2006 Conference on High-Performance Computing Systems*, IEEE Computer Society, 2006, available at `http://crd.lbl.gov/~dhbailey/dhbpapers/hpcs06.pdf`.

[5] David H. Bailey, Jonathan M. Borwein and Richard Crandall, "Integrals of the Ising Class," 2006, available at `http://crd.lbl.gov/~dhbailey/dhbpapers/ising.pdf`.

[6] David H. Bailey, Jonathan M. Borwein and Richard E. Crandall, "Ising Data," 2006, *Journal of Physics A: Mathematical and General*, vol. 39 (2006), pg. 12271–12302.

[7] David H. Bailey, Jonathan M. Borwein, Vishaal Kapoor and Eric Weisstein, "Ten Problems in Experimental Mathematics: A Challenge," *American Mathematical Monthly*, vol. 113, no. 6 (Jun 2006), pg. 481-409.

[8] David H. Bailey and David Broadhurst, "Parallel Integer Relation Detection: Techniques and Applications," *Mathematics of Computation*, vol. 70, no. 236 (2000), pg. 1719–1736.

[9] David H. Bailey and David J. Broadhurst, "A Seventeenth-Order Polylogarithm Ladder," manuscript, 1999, available at `http://crd.lbl.gov/~dhbailey/dhbpapers/ladder.pdf`.

[10] David H. Bailey, Yozo Hida, Xiaoye S. Li and Brandon Thompson, "ARPREC: An Arbitrary Precision Computation Package," technical report LBNL-53651, software and documentation available at `http://crdl.bl.gov/~dhbailey/mpdist`.

[11] David H. Bailey, Xiaoye S. Li and Karthik Jeyabalan, "A Comparison of Three High-Precision Quadrature Programs," *Experimental Mathematics*, vol. 14 (2005), no. 3, pg. 317–329.

[12] J. Borwein and D. Broadhurst, "Determination of rational Dedekind Zeta Invariants of Hyperbolic Manifolds and Feynman Knots and Links," available at `http://arxiv.org/hep-th/9811173`.

[13] J. Borwein, J. Zucker and J. Boersma, "Evaluation of Character Euler Double Sums," *Ramanujan Journal*, `http://www.springerlink.com/content/x24044144434w803/fulltext.pdf`.

[14] Richard E. Crandall, private communication, April 2005.

[15] William Gropp, Ewing Lusk, Anthony Skjellum, *Using MPI: A Portable Parallel Programming with the Message-Passing Interface*, MIT Press, Cambridge, MA, 1996.

[16] Yozo Hida, Xiaoye S. Li and David H. Bailey, "Algorithms for Quad-Double Precision Floating Point Arithmetic," *15th IEEE Symposium on Computer Arithmetic*, IEEE Computer Society, 2001, pg. 155–162.

[17] H. Takahasi and M. Mori, "Double Exponential Formulas for Numerical Integration," Publications of RIMS, Kyoto University, vol. 9 (1974), pg. 721–741.