Length Scales in Bayesian Automatic Adaptive Quadrature

Gh. Adam[1,2,a] and S. Adam[1,2,b]

[1]*Laboratory of Information Technologies, Joint Institute for Nuclear Research,*
6, Joliot Curie St., 141980 Dubna, Moscow Region, Russia
[2]*Horia Hulubei National Institute for Physics and Nuclear Engineering (IFIN-HH),*
30, Reactorului St., Măgurele - Bucharest, 077125, Romania

Abstract. Two conceptual developments in the Bayesian automatic adaptive quadrature approach to the numerical solution of one-dimensional Riemann integrals [Gh. Adam, S. Adam, Springer LNCS 7125, 1–16 (2012)] are reported. First, it is shown that the numerical quadrature which avoids the overcomputing and minimizes the hidden floating point loss of precision asks for the consideration of three classes of integration domain lengths endowed with specific quadrature sums: *microscopic* (trapezoidal rule), *mesoscopic* (Simpson rule), and *macroscopic* (quadrature sums of high algebraic degrees of precision). Second, sensitive diagnostic tools for the Bayesian inference on macroscopic ranges, coming from the use of Clenshaw-Curtis quadrature, are derived.

1 Introduction

The derivation of efficient and reliable algorithms implementing floating point computation within the Bayesian automatic adaptive quadrature (`BAAQ`) approach [1] to the solution of definite Riemann integrals rests on the correct identification and solution of the existing underwater stones which were not previously considered neither within the rigorous mathematical theory of the quadrature [2], nor in the so far proposed automatic adaptive quadrature (`AAQ`) algorithms and codes (see, e.g., [3–6]. Since the publication of the `BAAQ` overview in [1], we have identified two kinds of such matters which are addressed in the present paper.

First, the discovery that, in floating point computations, the algebraic degree of precision of an interpolatory quadrature sum is to be abandoned in favour of its floating point degree of precision [7] naturally raises the question concerning the *appropriateness* of a given interpolatory quadrature sum over a given integration range. Second, the correct assessment of the computational complexity of the Riemann definite integral over a given integration range enables the definition of decision paths which avoid the overcomputing coming from unsuitable activation of the Bayesian inference machinery.

Following this general plan, the discussion is organized in three sections. Section 2 provides a description of the general frame of the `BAAQ` approach to the numerical solution of the one-dimensional Riemann integrals. Section 3 discusses the derivation of consistent numerical solutions which avoid the overcomputing and minimize the hidden floating point precision losses. It is shown that this

---

[a]e-mail: adamg@jinr.ru,adamg@theory.nipne.ro
[b]e-mail: adams@jinr.ru,adams@theory.nipne.ro

asks for the consideration of *three distinct classes of integration domain lengths: microscopic, mesoscopic, macroscopic* with corresponding quadrature sums given respectively by the trapezoidal rule, the Simpson rule, and quadrature sums of high algebraic degrees of precision. Section 4 discusses new improved analysis tools for the assessment of the computational complexity of the integrals over macroscopic integration ranges in the frame of the Clenshaw-Curtis quadrature implementation described in the related paper [8]. The paper ends with concluding remarks in section 5.

2 General frame of the Bayesian automatic adaptive quadrature

We consider the `BAAQ` numerical solution of the (proper or improper) Riemann integral

$$I \equiv I[f] = \int_a^b g(x) f(x) \, dx, \tag{1}$$

under the assumption that the real valued *integrand function* $f(x)$ is continuous almost everywhere on $[a, b]$ such that (1) exists and is finite. The *weight function* $g(x)$ either absorbs an analytically integrable difficult factor in the integrand (e.g., endpoint singularity or oscillatory function), or else $g(x) \equiv 1, \ \forall x \in [a, b]$.

The automatic adaptive quadrature (`AAQ`) solution of (1) rests on the use of an *interpolatory quadrature sum* to get an approximation $Q \equiv Q[f]$ to $I[f]$. The *meaningfulness* of $Q[f]$ is assessed by deriving a *bound* $E \equiv E[f] > 0$ to the remainder $R[f] = I[f] - Q[f]$. For a prescribed accuracy τ requested at input, the approximation Q to I is assumed to end the computation provided

$$|R[f]| < E < \tau. \tag{2}$$

The specification of τ needs two parameters: the absolute accuracy ε_a and the relative accuracy ε_r, such that [3]

$$\tau = \max\{\varepsilon_a, \varepsilon_r |I|\} \simeq \max\{\varepsilon_a, \varepsilon_r |Q|\}. \tag{3}$$

If the condition (2) is not satisfied, the `AAQ` approach to the solution attempts at decreasing the error E by the subdivision of the integration domain $[a, b]$ into subranges using bisection. This procedure builds a *subrange binary tree* the evolution of which is controlled by an associated *priority queue*.

In this way, an *adaptive partition* Π_N of $[a, b]$ is generated,

$$\Pi_N[a, b] = \{a = x^0 < x^1 < \cdots < x^i < \cdots < x^N = b | N \geq 1\}. \tag{4}$$

Over each subrange $[x^{i-1}, x^i] \subset [a, b]$, *local* pairs $\{q_i, e_i > 0\}$ are computed for the approximation of the integral and its associated error and *global* outputs $\{Q_N, E_N > 0\}$ are got by summing the results obtained over the N existing subranges in (4). After each partition update, the termination criterion (2) is checked until it gets fulfilled.

The existing strict mathematical bounds to $R[f]$ [2] are of little use in the implementation of practical codes. The derivation of a practical bound $e > 0$ to the error associated to q rests on *probabilistic arguments* the validity of which is always subject to doubt.

The `BAAQ` advancement to the solution incorporates the rich `AAQ` accumulated empirical evidence (see, e.g., [3–6, 9–12]) into a general frame based on the Bayesian inference [13]. Essentially, the probabilistic character of the `AAQ` approach to the derivation of local error estimates $e > 0$ is preserved. However, each step of the gradual advancement to the solution is scrutinized based on a set of hierarchically ordered criteria which enable decision taking in terms of the established diagnostics.

3 Three classes of integration domain lengths

The so far reported implementations of `BAAQ` or `AAQ` algorithms rested on the use of interpolatory quadrature sums characterized by *high algebraic degrees of precision* (see, e.g., [3–6]).

The algebraic degree of precision is an invariant feature of a quadrature sum over the field $\mathbb{R}$ of the real numbers: its value remains constant irrespective of the extent and the localization of the current integration domain over the real axis.

However, this feature gets lost under floating point computations. In this case, the suitable corresponding quantity for the characterization of an interpolatory quadrature sum is its *floating point degree of precision* [7] which *significantly varies with the extent and the localization of the floating point integration domain representation within the set of the machine numbers*.

Since the floating point degree of precision was found to decrease dramatically under the decrease of the length of the integration domain, the importance of the quadrature sums of high algebraic degrees of precision will correspondingly diminish at the small integration domain lengths.

Reversing the argument, we may formulate the inverse problem: find the family of the integration ranges $[\alpha, \beta]$ over which the floating point degree of precision cannot exceed a prescribed value d. The answer follows from the definition 2 of [7]:

1^0. *Assume that the floating point computations are done over a set of machine numbers characterized by a t-bit significand.*

2^0. *Let $fl(a)$ denote the floating point approximation of $a \in \mathbb{R}$, and let*

$$X = \max\{fl(|\alpha|), fl(|\beta|)\}, \quad X > 0, \quad \rho = fl(|\beta - \alpha|/X), \quad 0 < \rho \le 2, \tag{5}$$

be defined over the finite integration range $[\alpha, \beta] \subset \mathbb{R}$.

3^0. *Then the length of the integration domain $[\alpha, \beta] \subset \mathbb{R}$ over which there can be defined a quadrature sum of the floating point degree of precision at most d is bounded from above by*

$$\min(X, \rho) \le 2^{-t/d} . \tag{6}$$

At small d values, this bound characterizes integration range boundaries staying very close to each other, hence significant cancellation by subtraction comes at the very beginning of the computation from the numerical assessment of the integration domain length. Under the use of a quadrature sum of high algebraic degree of precision, the error in the floating point computation of the integration domain length will entail hidden supplementary loss of precision associated to the computation of the integrand values at the error prone quadrature knots. Moreover, within the existing `AAQ` and `BAAQ` implementations, further enhancement of the precision loss at these two stages of the computation will result as a consequence of the need of subrange subdivisions.

Therefore, the quadrature sums to be implemented over ranges specified by (6) will have to satisfy the strict requirement to restrict the principal source of the floating point precision loss to the single unavoidable step stemming from the computation of the integration domain length.

There are two remarkable small d values corresponding to quadrature sums which satisfy to this quest [2]: $d = 2$, implemented by the trapezoidal rule, and $d = 4$, implemented by the Simpson rule.

The actual maximum domain lengths at which these two rules are to be implemented will exceed the threshold (6). The two empirical thresholds, τ_μ and τ_m, defined below, secure convenient definitions of the *optimality ranges* of the various quadrature sums at hand:

- *Microscopic ranges*, characterized by the threshold condition

$$0 < \min(X, \rho) \le \tau_\mu = 2^{-22} . \tag{7}$$

The *trapezoidal rule* is to be used over a microscopic range.

- *Mesoscopic ranges*, characterized by the threshold condition

$$\tau_\mu = 2^{-22} < \min(X, \rho) \leq \tau_m = 2^{-10} \, . \tag{8}$$

The *Simpson rule* is to be used over a mesoscopic range.

- *Macroscopic ranges*, characterized by the threshold condition

$$\min(X, \rho) > \tau_m = 2^{-10} \, . \tag{9}$$

Quadrature sums of high algebraic degrees of precision can be used over macroscopic ranges.

Given a microscopic (or mesosocopic) integration range $[\alpha, \beta] \subset \mathbb{R}$, the minimization of the round-off errors within the trapezoidal rule (or Simpson rule respectively) is secured as follows.

The integration range $[\alpha, \beta]$ is mapped onto $[0, 1]$ by the substitution (floating point representations and floating point operations with the involved quantities are assumed) $x = \alpha + hy$, $h = \beta - \alpha$, and the current Riemann integral (1) over $[\alpha, \beta]$ is transformed accordingly to get

$$I[\varphi] = h \int_0^1 g(\alpha + hy)\, \varphi(y)\, dy, \quad \varphi(y) = f(\alpha + hy).$$

This step associates the unavoidable round-off cancellation error coming from the computation of the integration domain length h.

Besides the minimum number of integrand evaluations asked by the corresponding quadrature rule, additional requested integrand evaluations are performed at suitable newly added *machine number* reduced abscissas inside $[0, 1]$ in terms of which all the newly involved subtraction operations in the resulting composite rules are done *exactly*.

The next section shows that, over macroscopic integration ranges, the Clenshaw-Curtis (CC) quadrature results in efficient Bayesian inferences concerning the computational complexity of the input integral (1) at the first attempt to solve it over the whole integration domain $[a, b]$.

This makes possible the definition of improved decision paths concerning the continuation of the integration under subrange subdivision with substantial decrease of the computing effort.

4 Improved diagnostic tools for Bayesian inference

The elucidation of the mathematical structure of the Chebyshev expansion coefficients entering the CC quadrature and the subsequent scrutiny of the outputs obtained for the even and odd rank Chebyshev coefficient subsets evidenced [8] the capability to precisely forecast, from the variation of these coefficients within the respective subsets, the quality of the output of the still to be computed CC quadrature sums as fast converging, hopefully converging, or biased by integrand ill-conditioning, respectively.

The quantitative characterization of the expected quadrature sum output as fast converging or hopefully converging (hence the expected occurrence of an *easy* integral (1) the solution of which does not need the activation of the Bayesian inference machinery) depends on the parity properties of the weight function $g(x)$ in (1).

Assume first that $g(x)$ is a superposition of even and odd terms with respect to the centre c of $[a, b]$. Then let $\{c_{\max}^{(e)}, c_{\min}^{(e)}\}$, respectively $\{c_{\max}^{(o)}, c_{\min}^{(o)}\}$, denote the largest, respectively the smallest absolute magnitudes within the even rank and odd rank subsets and let $c_{\max} = \max\{c_{\max}^{(e)}, c_{\max}^{(o)}\}$. Further, let $c_{\min} = \min\{c_{\min}^{(e)}, c_{\min}^{(o)}\}$.

A preliminary condition to be checked for these quantities is whether there is a *dominant* subset or not within the pair of the even and odd rank subsets.

The even rank subset is said to be dominant if

$$c_{\max}^{(o)} < p c_{\max}^{(e)}. \tag{10}$$

Conversely, the odd rank subset is said to be dominant provided

$$c_{\max}^{(e)} < p c_{\max}^{(o)}. \tag{11}$$

Here, $p = 2^{-8} \simeq .39 \times 10^{-2}$ denotes an empirically defined threshold pointing to two figure accuracy.

If neither (10) nor (11) are satisfied, the fulfilment of the condition

$$c_{\min} < p c_{\max} \tag{12}$$

provides a first necessary convergence criterion of the Chebyshev series expansion.

The second necessary convergence criterion is got in terms of the decreasing rates of the Chebyshev coefficients under the increase of their labels inside each of the even and odd rank Chebyshev coefficient subsets. Convergence is said to be achieved if either the highest label Chebyshev coefficient in the corresponding subset has reached the floating point round-off threshold, or a monotonic decrease of the absolute coefficient magnitudes inside the highest label quarter subset was observed.

Under the occurrence of a dominant subset, the above two convergence conditions are checked for that subset only.

Assume next that the weight function $g(x)$ is either even or odd. In this case, the convergence evaluation involves a supplementary check for that Chebyshev coefficient subset showing the parity of $g(x)$, namely, *the assessment of the precision loss coming from the cancellation by subtraction*. For the sake of the argument, let us assume that $g(x)$ is even. The argument for an odd $g(x)$ is similar.

A convenient measure of the expected precision loss is provided by the ratio

$$\rho = c_{\max}^{(e)}/c_{\max}. \tag{13}$$

Two issues are possible.

First, $\rho > p$, with p defined as above. Then the precision loss is assumed to be not important. In this case, both the condition (12) and the decreasing rate of the Chebyshev coefficients are checked within the even coefficient subset only.

Second, under $\rho \leq p$, the precision loss is important. In this case, the requested relative error accuracy ε_r at input in the termination criterion (3) may be unrealistic. The best accuracy expected from a `BAAQ` solution cannot exceed the threshold

$$\bar{\varepsilon}_r = \max\{\varepsilon_r, \varepsilon_0/\rho\}, \tag{14}$$

which is to replace ε_r in the termination criterion (3). Here, ε_0 denotes the machine epsilon with respect to the addition and the quantity ρ is given by (13).

As a rule, significant precision loss occurs under heavy integrand oscillations inside the integration domain $[a, b]$. Therefore, if the diagnostic of catastrophic cancellation by subtraction, which is asking to stop the computation, was not set, then the evaluation of the integral is continued under subrange subdivision, hopefully at a discretization abscissa where the integrand vanishes.

5 Concluding remarks

The growth of the Bayesian automatic adaptive quadrature (`BAAQ`) approach to the numerical solution of the Riemann integrals beyond the status report given in [1] needs new conceptual developments. Two such concepts were elaborated in the present paper.

The first concept addresses the numerical solution of the Riemann integrals within a `BAAQ` approach which avoids the overcomputing and secures the minimization of the direct and hidden floating point loss of precision. This was shown to be possible provided the manifold of the nonvanishing integrand domain lengths is split into three submanifolds of distinct extension ranges endowed with specific quadrature sums: *microscopic – trapezoidal rule*, *mesoscopic – Simpson rule*, and *macroscopic – quadrature sums of high algebraic degrees of precision*.

The second concept addresses the Bayesian inference over macroscopic integration ranges. It is shown that the Clenshaw-Curtis (CC) quadrature provides fast and sensitive diagnostic tools which promote it as the best `BAAQ` candidate to the provision of the principal quadrature sum.

The scrutiny of the variation of the Chebyshev expansion coefficients within the CC quadrature enabled quantitative characterization of their properties. This allowed early Bayesian inference concerning the integrand conditioning and the expected output quality: (i) well-conditioned integrand, typical for an easy (or hopefully converging) integral within the standard automatic adaptive quadrature approach; (ii) heavily oscillatory integrand asking for the scrutiny of the possible redefinition of the attainable output accuracy within the `BAAQ` approach; (iii) highly probable integrand ill-conditioning asking for the activation of the integrand profile analysis procedure for the inference of precise conditioning diagnostics.

Acknowledgements

References

[1] Gh. Adam and S. Adam, in *Mathematical Modeling and Computational Science* (MMCP2011), Gh. Adam, J. Buša, and M. Hnatič (Eds.) (Springer, Heidelberg, LNCS 7125, 2012) pp.1–16

[2] H. Brass and K. Petras, *Quadrature Theory. The Theory of Numerical Integration on a Compact Interval* (Math. Surv. and Monogr. vol. 178, American Mathematical Society, Providence, 2011)

[3] R. Piessens, E. de Doncker-Kapenga, C. W. Überhuber, and D. K. Kahaner, `QUADPACK`, *a subroutine package for automatic integration* (Springer Verlag, Berlin, 1983)

[4] P.J. Davis and P. Rabinowitz, *Methods of Numerical Integration* (2-nd Ed., Academic Press, Orlando (Fla), USA, 1984)

[5] A.R. Krommer and C.W. Ueberhuber, *Computational Integration* SIAM, Philadelphia, 1998.

[6] `NAG` *Library Manual*, Mark 24, section D01: http://www.nag.co.uk/numeric/fl/nagdoc_fl24/html/D01/d01conts.html (The Numerical Algorithms Group Ltd., Oxford, UK, 2012)

[7] S. Adam and Gh. Adam, in *Mathematical Modeling and Computational Science* (MMCP2011), Gh. Adam, J. Buša, and M. Hnatič (Eds.) (Springer, Heidelberg, LNCS 7125, 2012) pp.189–194

[8] S. Adam and Gh. Adam, *Summation Paths in Clenshaw-Curtis Quadrature*, related paper at this Conference

[9] J.N. Lyness, SIAM Rev. **25**, 63–87 (1983)

[10] Gh. Adam and A. Nobile, IMA J. Numer. Analysis **11**, 271–296 (1991)

[11] Gh. Adam, Rom. Journ. Phys. **38**, 527–538 (1993)

[12] Gh. Adam, S. Adam, and N.M. Plakida, Comp. Phys. Comm. **154**, 49–64 (2003)

[13] P. Bessiere, E. Mazer, J.M. Ahuactzin, and K. Mekhnacha, *Bayesian Programming* (Chapman and Hall/CRC, 2013)