Fast Bayesian automatic adaptive quadrature

Gh. Adam, S. Adam

LIT-JINR Dubna, Russia & IFIN-HH Bucharest, Romania

Invited Lecture at the RO-LCG 2014 conference, November 3-5, 2014

Outline

- Introduction
- Hardware Environment
- Bayesian Steps of the Analysis
- New Priority Queue Perspective
- A Few Case Studies

Motivation

The present talk describes an attempt at implementing a ***robust, reliable, fast,*** *and* ***highly accurate*** *computational tool the main purpose of which is to enable* ***modeling of physical phenomena*** *within* ***numerical experiments*** *asking for the evaluation of large numbers of* ***Riemann integrals*** *by numerical methods.*

For instance, the study of the behavior of a system under ***sudden change*** *of an* ***inner order parameter****, which results in drastic modification of the mathematical properties of the integrand (e.g., in phase transitions or processes involving fragmentation or fusion, nanostructures)* ***cannot be accommodated*** *within the* ***standard automatic adaptive quadrature (AAQ)*** *approach to the numerical solution, due to the impossibility to decide in advance on the correct choice of the convenient library procedure.*

The ***Bayesian automatic adaptive quadrature (BAAQ)*** *approaches the numerical solution of the integrals by* ***merging rigorous mathematical criteria*** *with the* ***reality*** *of the* ***hardware*** *and* ***software environments****, such as to avoid, if at all possible, unreliable outputs originating in the human factor.*

References on the Standard Approach to Automatic Adaptive Quadrature

The standard AAQ was systematically developed in QUADPACK, the *de facto* standard of one-dimensional numerical integration.

See:

- R.Piessens, E. deDoncker-Kapenga, C.W. Überhuber, D.K. Kahaner, *QUADPACK, A Subroutine Package for Automatic Integration*, Springer, Berlin, 1983
- P.J. Davis, P. Rabinowitz, *Methods of Numerical Integration*, Academic, NY, 1984
- A.R. Krommer, C.W. Ueberhuber, *Computational Integration*, SIAM, Philadelphia, 1998
- J.N.Lyness, When not to use an automatic quadrature routine, SIAM Review, 25, 63-87 (1983)
- Gh. Adam, Case studies in the numerical solution of oscillatory integrals, Romanian J. Phys., 38, 527-538 (1993)

References on the Bayesian Approach to Automatic Adaptive Quadrature

See:

- Gh. Adam, S. Adam, Handling accuracy in Bayesian automatic adaptive quadrature, to be published in Journal of Physics: Conference Series (subm. 09 2014)
- Gh. Adam, S. Adam, Bayesian Automatic Adaptive Quadrature: an Overview, in *Mathematical Modeling and Computational Science*, LNCS, Vol. 7125, Eds. Gh. Adam, J. Busa, M. Hnatic, (Springer-Verlag, Berlin, Heidelberg), 2012, pp. 1-16
- S. Adam, Gh. Adam, Floating Point Degree of Precision in Numerical Quadrature, in *Mathematical Modeling and Computational Science*, LNCS, Vol. 7125, Eds. Gh. Adam, J. Busa, M. Hnatic, (Springer-Verlag, Berlin, Heidelberg), 2012, pp. 189-194
- Gh. Adam, S. Adam, Quantitative Conditioning Criteria in Bayesian Automatic Adaptive Quadrature, in *Proceedings of 2012 5th Romania Tier 2 Federation Grid, Cloud & High Performance Computing Science*, UT Press, Cluj-Napoca, Romania (IEEE Conference Series), 2012, pp. 35-38 http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6528239&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxpls%2Ficp.jsp%3Farnumber%3D6528239
- Gh. Adam, S. Adam, Principles of the Bayesian automatic adaptive quadrature, Numerical Methods and Programming: Advanced Computing (RCC MSU) 2009, Vol.10, pp. 391-397 (http://num-meth.srcc.msu.ru)
- Gh. Adam, S. Adam, N.M. Plakida, Reliability conditions in quadrature algorithms, Computer Physics Communications, Vol. 154 (2003) pp.49-64

Standard Input Numerical Problem

Given the (proper or improper) Riemann integral

$$I[a,b]f = \int_a^b w(x) f(x) \, \mathrm{d}\, x, \; -\infty < a < b < +\infty$$

we seek a ***globally adaptive*** numerical solution $\{Q, E > 0\}$

of it within input accuracy specifications $\{\varepsilon_r > 0, \varepsilon_a \geq 0\}$

i.e.,

$$| I[a,b]f - Q | < E < \max\{\varepsilon_r \, | I[a,b]f \, |, \varepsilon_a\}$$
$$\approx \max\{\varepsilon_r \, | Q \, |, \varepsilon_a\}$$

Permanent Features of the Automatic Adaptive Quadrature (AAQ)

- The computation scheme developed within the AAQ approach implements an ***integrand adapted discretization*** of $[a, b]$, which defines a ***partition*** of $[a, b]$,

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

  Over each subrange $[x^{i-1}, x^i] \subseteq [a,b]$, a (possibly subrange dependent) ***local quadrature rule*** yields a local pair $\{q, e\}$ where, q stands for the output of an ***interpolatory quadrature sum*** solving $I[x^{i-1}, x^i]f$, while $e > 0$ stands for the output of a ***probabilistic estimate*** of the ***local error*** associated to q.
- A partition dependent ***global pair*** solving $I[a,b]f$, $\{Q \equiv Q_N, E \equiv E_N > 0\}$, is got by summing up the individual outputs $\{q, e\}$ over the subranges of the partition $\Pi_N[a,b]$. In what follows, to simplify notations, we will consider a ***generic*** subrange $[\alpha, \beta] \subseteq [a,b]$ standing for any subrange $[x^{i-1}, x^i]$ of $\Pi_N[a,b]$.
- The ***number of subranges*** of $\Pi_N[a,b]$ ***starts*** with $N = 1$ and, if necessary, it is increased by ***gradual refinement*** of $\Pi_N[a,b]$ until either the global accuracy criterion is satisfied, or a failure diagnostic is issued.

Standard Approach to the Automatic Adaptive Quadrature (SAAQ)

- Implements the refinement of the partition $\Pi_N[a,b]$ as a ***subrange binary tree*** the evolution of which is controlled by an associated ***priority queue***.
  The binary tree ***initialization*** equates the root with the input integration domain $[a, b]$ over which a first global output $\{Q_1, E_1 > 0\}$ is computed.
  If the termination criterion is not fulfilled, then a ***recursive procedure*** is followed: the priority queue is activated, the resulting root is ***bisected***, local estimates $\{q, e > 0\}$ are computed over each resulting sibling, the global quantities $\{Q_N, E_N > 0\}$ are updated, the end of computation is checked again.
- The standard subdivision scheme can be supplemented with a ***convergence acceleration algorithm*** if the occurrence of an ***integrand singularity*** was heuristically inferred.
- Two remarks: First, **SAAQ *successfully solves*** integrals over continuous integrands. Second, **SAAQ *fails badly*** under the occurrence of inner either zero-measure or singular discontinuities of the integrand.

The BAAQ has to cope with both these circumstances.

Four Enhancements of the Standard Automatic Adaptive Quadrature (ESAAQ) (1)

- ***A more reliable local error estimator within Clenshaw-Curtis*** (CC) ***quadrature*** uses 4-term CC-like error estimator [Gh. Adam & A. Nobile, *IMAJNA*, **11**, 271-96 (1991)]
- ***Stabilization of the local quadrature sums as Riemann sums*** is checked by means of two criteria:

(i) The output generated by each sibling of the current root is ***meaningful*** provided in the pair $\{q, e > 0\}$, q carries out two accurate significant figures at least.
Otherwise, it is ***doubtful*** and the ***conventional values*** $\{q = 0, e = \Omega\}$ are returned, where Ω denotes a very large positive number close to machine overflow.

(ii) Let $\{q_p, e_p\}$ denote the output carried by a meaningful parent brought to the root
position. Let $\{q_d, e_d\}$ denote the sums of the outputs coming from its two siblings which are meaningful according to the criterion (i). If $|q_p - q_d| > e_p + e_d$, then conventional meanings are assigned to both descendants.

Note: The second criterion is effective for highly oscillatory integrands.

Four Enhancements of the Standard Automatic Adaptive Quadrature (ESAAQ) (2)

The check of the local quadrature sums stabilization as Riemann sums brings fundamental modifications to the update of the partition $\Pi_N[a,b]$ and the computation of the global output $\{Q_N, E_N > 0\}$.

- ***The priority queue brings to the root position every conventional subrange before any meaningful subrange.*** Indeed, this is a straightforward consequence of the fact that the place of a subrange in the priority queue list is determined by the magnitude of its local quadrature error estimate.

- ***Conditional activation of the global termination criterion.***

Is a consequence of the fact that the output brought by a subrange marked as conventional is useless for the update of the global pair $\{Q_N, E_N > 0\}$.

An easy implementation of this feature was obtained by keeping record of the number of subranges carrying conventional meaning and by attempting the *check of the global termination criterion only provided this number reaches the floor value zero.*

Optimistic Assumptions of the Standard AAQ Documented to Result in Pitfalls

- ***Guaranteed reliability*** of the ***local*** quadrature rule output $\{q, e\}$ over any subrange $[\alpha,\beta]\subseteq[a,b]$.
- ***Nil effect*** of an ***accidentally*** occurring ***spurious*** local output $\{q, e\}$ on the ***global*** termination decision.
- The ***global*** error estimates $E > 0$ ***always*** provide ***upper bounds*** to the (unknown) errors $|I\text{-}Q|$.
- The ***decision path*** defined inside the AAQ global control chain ***always*** activates the ***right procedure*** for the advancement to the solution.
- **Fact**: The use of ***bisection*** as the ***exclusive*** subrange subdivision strategy results in ***failures*** whenever ***inner*** finite discontinuities or singularities of the integrand (or its first order derivative) are present.
- Absence of ***Gibbs phenomenon***

Critical Issues of the Bayesian Automatic Adaptive Quadrature (BAAQ)

- ***Parameters to be simultaneously accounted for within* BAAQ:**
  - ***Integrand profile*** (exceptional cases: zero, constant, linear; continuous slowly or mildly variable and/or oscillatory; continuous highly variable and/or oscillatory; almost continuous with endpoint or inner zero-measure discontinuities; almost continuous with endpoint or inner singularities; a combination of all the above)
  - ***Extension of the integration domain***
  - ***Occurrence of an explicitly defined weight function***
- ***Fast algorithm convergence*** in case of ***zero-measure or singular discontinuity points is secured provided these are brought to the interval endpoints entering the partition*** $\Pi_N[a,b]$**.** This ***cannot be done*** within the **ESAAQ *approach.***
- ***Almost complete cancellation by subtraction*** *is to be identified and diagnosed along the lines defined within the standard* AAQ (see QUADPACK).
- ***The finite binary floating point arithmetic which is implemented in the hardware results in fundamental specific consequences.***

Outline

- Introduction
- **Hardware Environment**
- Bayesian Steps of the Analysis
- New Priority Queue Perspective
- A Few Case Studies

Range of the Admissible Values of the Relative Error ε_r

An ***insightful*** relative error parameter ε_r satisfies

$$\varepsilon_{r,\max} \geq \varepsilon_r \geq \varepsilon_{r,\min} = \varepsilon_0 / \varepsilon_{r,\max}\,, \;\; \varepsilon_{r,\max} = 2^{-11} \approx 0.5 \cdot 10^{-2}$$

- The ***lower bound*** $\varepsilon_{r,\min}$ stems from the ***finite length*** of the significands of the floating point machine numbers.

  Here, ε_0 stays for the ***machine epsilon*** (the largest relative spacing) denoting the smallest positive machine number ε_0 for which

$$fl(1. + \varepsilon_0) > 1.; \;\; fl(1. + \eta) = 1., \;\; \forall\, \eta \in [0, \varepsilon_0).$$

- With the given empirical choice of $\varepsilon_{r,\min}$, related to the ***upper bound*** $\varepsilon_{r,\max}$, it is ***always*** possible (and necessary) to get a relative precision of the computed output Q carrying out at least two significant accurate decimal digits.

The Hardware Environment

- **64-bit *floating point* approximation of the real numbers.** Is done via a ***finite set of machine numbers*** of the form $(\pm)\ s \times b^{e}$ (the 64-bit definition assumed here asks for a ***radix*** $b = 2$, an ***exponent*** $e = 11$, a ***fixed length significand*** s of 52 bits, and a ***sign bit***)
- **Consequence: *each*** machine number (except for 0,+∞,-∞,NaN) associates an ***infinite*** subset of real numbers surrounding it over some real axis interval
- **Corollary1: *Bayesian decision path inferences are necessary*** for the advancement toward the solution of the Riemann integrals by floating point computations.
- **Corollary2:** A nasty source of ***spurious output*** occurs whenever the ***ordering relationship*** valid for pairs of real numbers is ***falsified*** while going to machine numbers. Specifically, the real number ordering relationships ">, larger" and "<, smaller" ***change*** to the "=, equal" ordering relationship over the machine number set (hence result in ***false Bayesian inferences***) whenever ***two different real numbers*** are approximated by a ***same machine number.***

Distance between neighbouring machine numbers

Check of reliability criterion fulfilment

$$fl(1+(k+1)\varepsilon_0) - fl(1+k\varepsilon_0) = \varepsilon_0$$

$\varepsilon_0 = 2^{-52}$ *is right !*

$\varepsilon_0 = 2^{-53}$ *is wrong !*

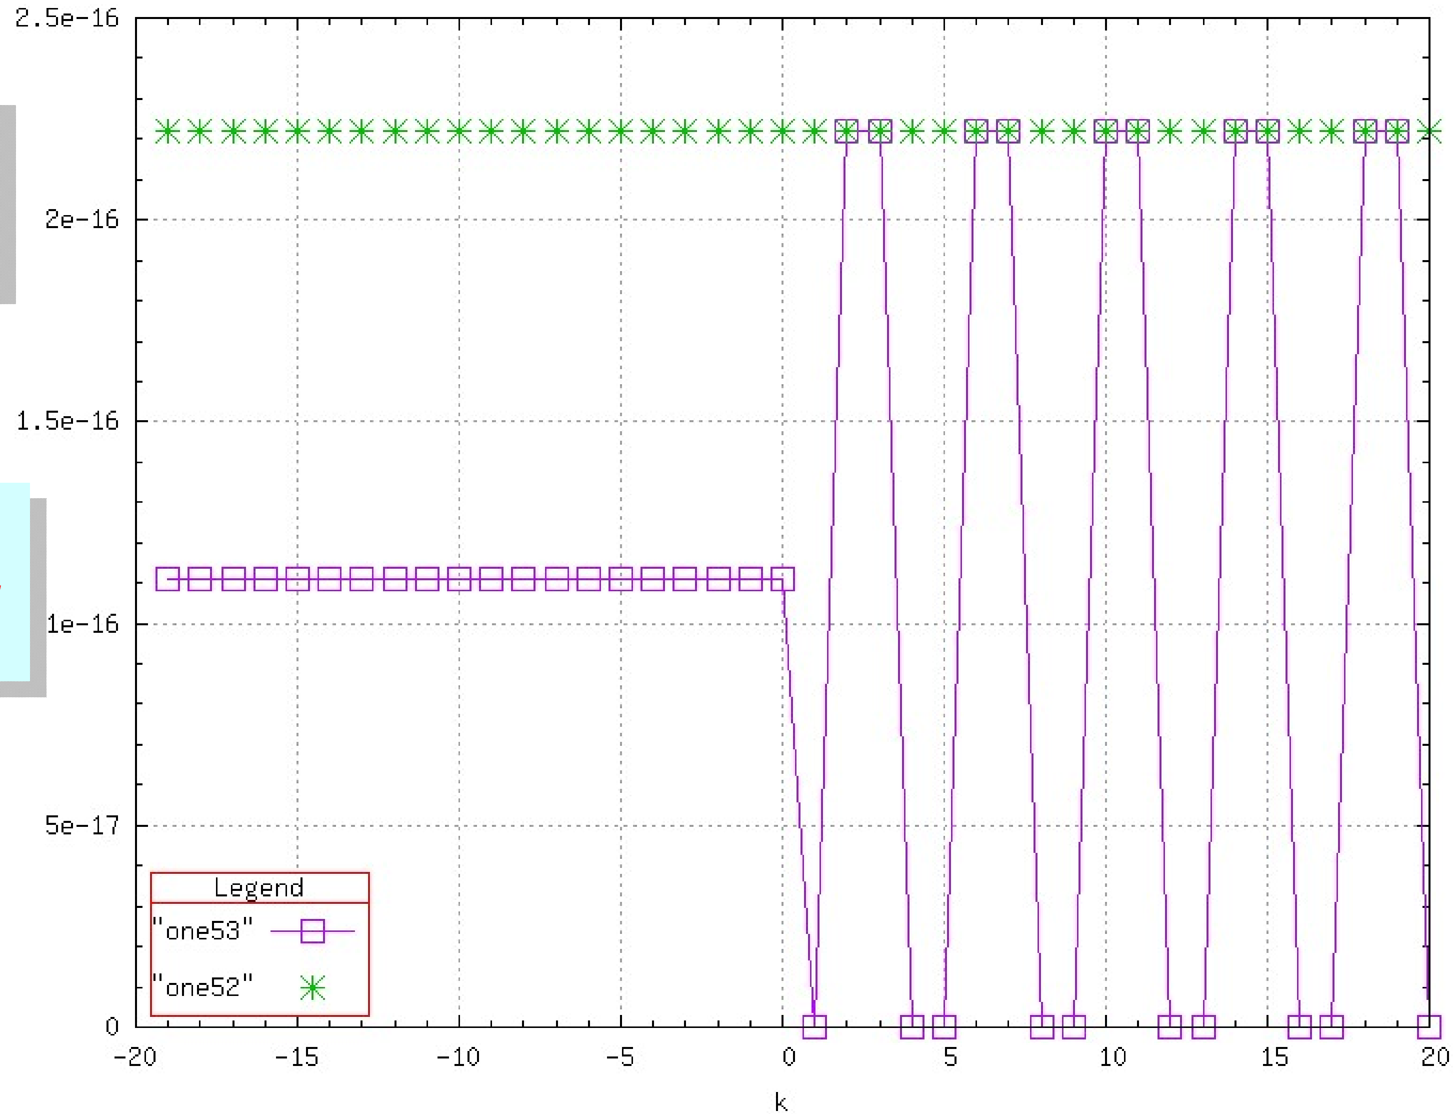

Algebraic Degree of Precision of an Interpolatory Quadrature Sum

- An interpolatory quadrature sum approximation of the Riemann integral $I[\alpha,\beta]f$, $[\alpha,\beta]\subseteq[a,b]$ writes

$$q_{n+1}[\alpha,\beta]f = I[\alpha,\beta]p_n$$

where the interpolatory polynomial $p_n(x)$ values equate those of the integrand function $f(x)$ at a specific set of quadrature knots x_k,

$$p_n(x_k) = f(x_k),\ k = 0, 1, ..., n.$$

- The quadrature sum $q_{n+1}[\alpha,\beta]f$ solves ***exactly*** the polynomial integrals over the ***fundamental power set***,

$$q_{n+1}[\alpha,\beta]x^k = I[\alpha,\beta]x^k,\ \forall k = 0,1,\cdots,d,\ \forall[\alpha,\beta]\subset\Re$$

- The maximum degree d, at which these identities hold, defines the ***algebraic degree of precision*** of the quadrature sum $q_{n+1}[\alpha,\beta]f$.
- Over the field of real numbers, d is a ***specific universal parameter*** of a given interpolatory quadrature sum, irrespective of the ***extent and localization*** of the integration domain on the real axis.

Forward Floating Point Degree of Precision of an Interpolatory Quadrature Sum (1)

In the calculation over $\Re$ of the set of probe integrals

$$\sigma_m = I[0,\beta]\pi_m,\ \pi_m(x) = \sum_{l=0}^{m} x^l,\ \beta > 0,\ m = 0,1,\cdots,d$$

each monomial x^l entering the integrand $\pi_m(x)$ brings a ***distinct, non-negligible***, contribution to σ_m.

In floating point computations, the above property of the monomials x^l of bringing distinct, non-negligible contributions to σ_m may get infringed both at integration limits $\beta << 1$ and $\beta >> 1$.

The maximum degree $d_{\text{fp}} \le d$ at which the identity of the individual monomial contributions is preserved in floating point computations at $\beta << 1$ defines the ***forward floating point degree of precision*** of the quadrature sum.

Its definition is formalized in the next slide.

Forward Floating Point Degree of Precision of an Interpolatory Quadrature Sum (2)

1. Let $I[\alpha,\beta]f$ denote the input integral of interest, defined over a finite integration range $[\alpha,\beta]\subset\Re$.

2. Let $q[\alpha,\beta]f$ denote an interpolatory quadrature sum of algebraic degree of precision d , which solves $I[\alpha,\beta]f$ through floating point computations over a set of machine numbers characterized by a t-bit significand.

3. Let $fl(a)$ denote the floating point approximation of $a\in\Re$ and let

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

4. If $\xi > 0$ stands for either X or ρ, we define

$$d_\xi=\begin{cases} d & \text{iff}\quad \xi\ge x_m \\ \lceil \ln\varepsilon_0/\ln\xi\rceil & \text{iff}\quad \xi< x_m\end{cases}$$

where $\varepsilon_0 = 2^{-t}$, $x_m = fl(\varepsilon_0{}^{1/d})$, and $\lceil a\rceil$ is the ceiling of $fl(a)$.

5. Then the ***floating point degree of precision***, $0<d_{\text{fp}}\le d$ of $q[\alpha,\beta]f$ is

$$d_{\text{fp}} = \min\{d_X, d_\rho\}$$

with d_X and d_ρ computed from **4.** for the terms X and ρ of the pair **3.**

Variation of the Forward Floating Point Degree of Precision with Interval Length

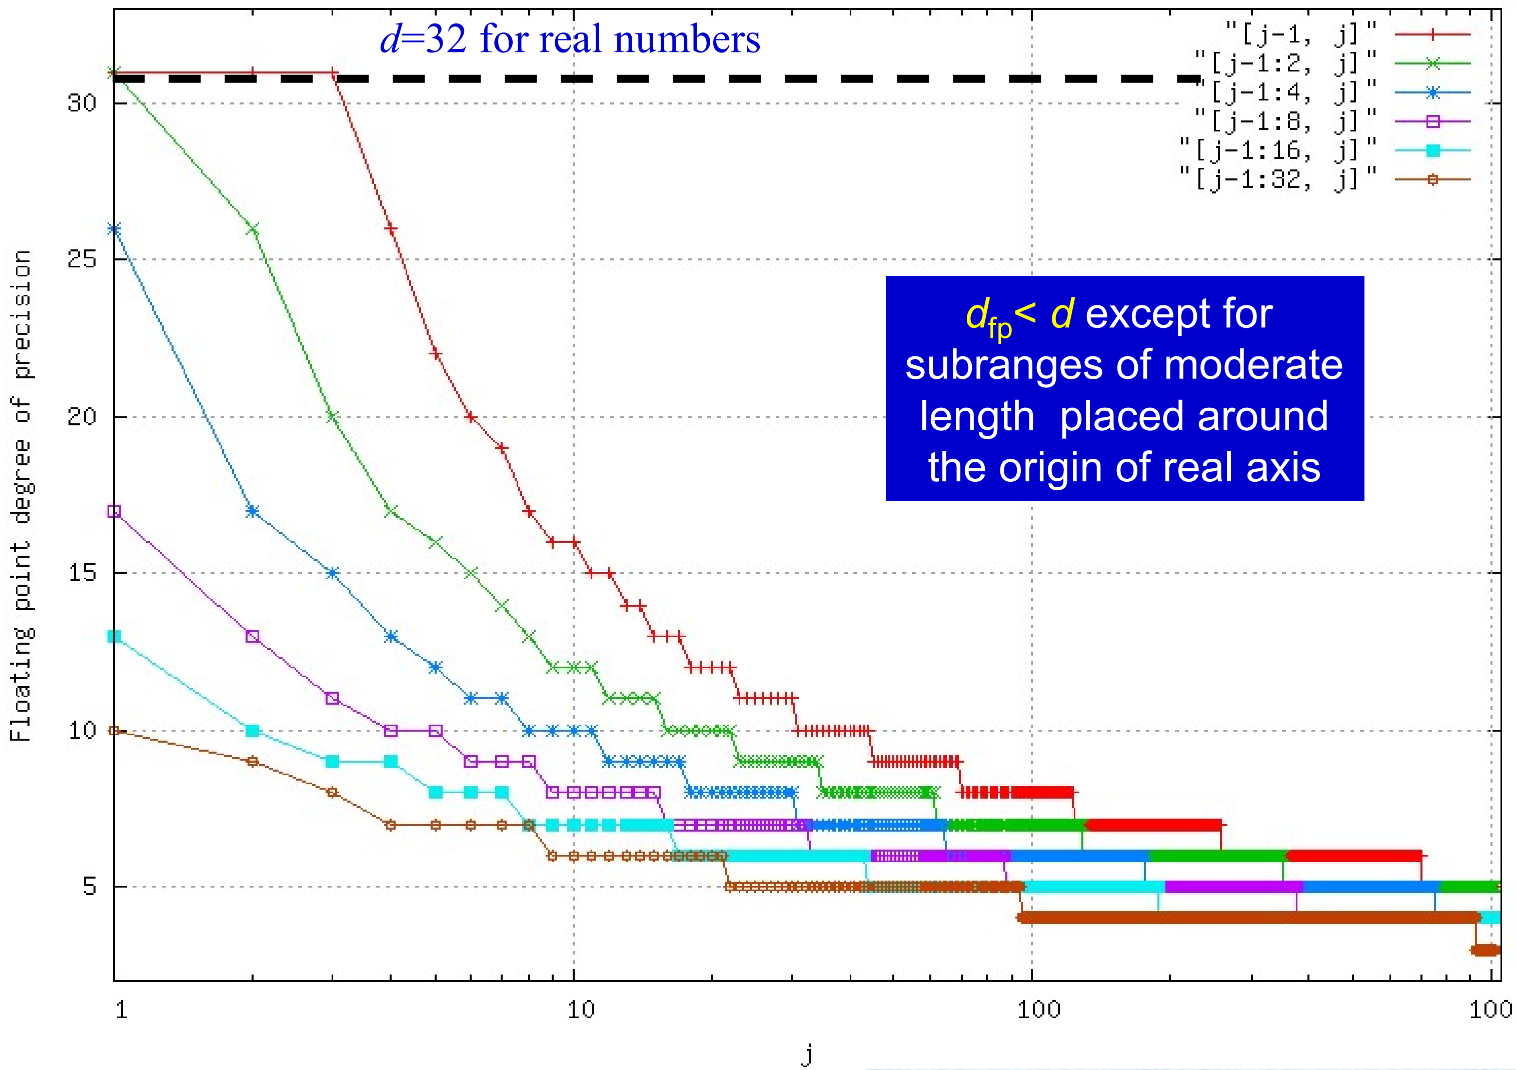

S. Adam, Gh. Adam, LNCS 7125, Springer (2012)

Local Quadrature Rules Coping with the Hardware Environment

Given $[\alpha, \beta] \subset \Re$, let $X = \max\{fl(|\alpha|), fl(|\beta|)\},\ X > 0;\quad \rho = fl(|\beta - \alpha| / X),\ 0 < \rho \le 2.$
For Intel 8087 processor, $\varepsilon_p = 2^{-64}$, $\varepsilon_0 = 2^{-52}$, $\tau_u = u/\varepsilon_0$, $u = 2^{-1022}$,
$\tau_m = (\varepsilon_p/16)^{1/2}$, $\tau_M = 2^{-11}$

Range $[\alpha, \beta]$	Characterization	Reference q (q_1)	Auxiliary q (q_2)
Macroscopic	$X > \tau_M$, $\rho > \tau_M$	33 knot Clenshaw-Curtis (CC-33) $d_{CC} = 32$	21 knot Gauss-Kronrod (GK-21) $d_{GK} = 31$
Mesoscopic	$X > \tau_u$, $\rho > \tau_m$	Simpson	Trapeze/Rectangle
Tiny	Remaining $\{X, \rho\}$ machine numbers	Trapeze	Rectangle

Features:

- ***Macroscopic range -*** **quasi-continuous** *machine number distribution inside*
- ***Mesoscopic range –*** **non-uniform discrete** *machine number distribution inside*
- ***Tiny range –*** **(quasi-)uniform** *machine number distribution inside*

Overview

- ❑ Introduction
- ❑ Hardware Environment
- ➡ Bayesian Steps of the Analysis
- ❑ New Priority Queue Perspective
- ❑ A Few Case Studies

Bayesian Output Assessment

(A) At the *first* attempt to solve the integral over $[a, b]$

- ***End computation*** if an ***exceptional case*** was detected (vanishing integrand, vanishing output q for non-vanishing integrand, catastrophic cancellation by subtraction)
- ***End computation*** if the input ***accuracy criteria were met***
- ***Propose*** the future ***decision path*** to be followed for:
  - ***easy integral*** ($e = |q_1 - q_2| < \tau_0$)
  - ***difficult integral*** ($e \geq \tau_0$)

(B) Path followed in case of an *easy integral*

1. Pick up from the ***priority queue*** the subrange $[\alpha, \beta] \subseteq [a, b]$ characterized by the ***largest*** local error estimate $e > 0$.

2. ***Bisect*** it, ***compute*** pairs $\{q', e'\}$, $\{q'', e''\}$, over subranges, check stabilization as Riemann sums .

3. If (stabilization) then Update global quantities Q, E. Check for ***global termination.***

4. If (.NOT.stabilization .OR. .NOT.termination) then

If ($e > e' + e''$) then

update the priority queue; go to step **1.**

else ***change*** decision path to ***difficult.***

(C) Path followed in case of a *difficult integral*

Cannot be characterized in a few lines. The following discussion is done for ***macroscopic*** integration domains, the most frequently encountered case.

Partition of the Integration Domain $[a, b]$

- The advancement of the computation of a ***difficult integral*** $I[a, b]f$ generates a ***partition*** of $[a, b]$, $\Pi_N = \{a = x^0 < x^1 < \cdots < x^i < \cdots x^N = b\}$, $N \geq 2$
- **Two classes of *subrange endpoints* inside Π_N:**
  - ***Terminal endpoints:***
    - ► The ***ends*** a and b of $[a, b]$
    - ► ***Inner abscissas*** $x^i \in \Pi_N$ at which either $f(x^i - 0) \neq f(x^i + 0)$ or $f'(x^i - 0) \neq f'(x^i + 0)$
  - ***Two sided endpoints*** are ***inner*** x^i at which $f(x^i - 0) = f(x^i + 0)$ and $f'(x^i - 0) = f'(x^i + 0)$
- **Generation of *terminal endpoints***

  Is the primary goal of the Bayesian analysis whenever an ***ill-conditioning*** diagnostic is inferred to hold inside an ***isolated inner region*** of the subrange subject to the analysis.

  The essential feature enabling solutions (under the occurrence of both ***singularities*** and ***finite discontinuities***) is the ***scale invariance property*** of the Bayesian conditioning diagnostics of ***genuinely resolved*** integrand profiles.

  The developed procedures use combined top-down (***sharpening***) and down-up (***localization to machine accuracy***) approaches. Appropriate ***lateral limits*** inside the left and right neighbourhoods of an inner terminal endpoint are generated both for the integrand and its first order derivative.
- **Generation of new *two-sided endpoints***

  Is done if the Bayesian analysis established that either the allowed rate of variation of a ***monotonic*** integrand was exceeded, or a non-vanishing number of integrand ***oscillations*** was missed within the integrand profile of the reference quadrature sum over the region of interest.

Tools for Bayesian Analysis: Integrand Profiles

- For any $[\alpha,\beta]\subseteq[a,b]$ we write the ***symmetric*** decomposition

$$[\alpha,\beta]=[\alpha,\gamma]\cup[\gamma,\beta],\ \ \gamma=(\beta+\alpha)/2,\ h=(\beta-\alpha)/2.$$

- Over the ***left*** (l) and ***right*** (r) halves of $[\alpha, \beta]$, the floating point integrand values entering the quadrature sums are computed respectively as

$$f_k^l=f(\alpha+h\eta_k),\quad f_k^r=f(\beta-h\eta_k),$$

where

$$0\le\eta_0<\eta_1<\cdots<\eta_k<\cdots\eta_n=1,\ \ n\in\{n_{\mathrm{CC}},n_{\mathrm{GK}}\}$$

stay for the floating point values of the ***reduced modified quadrature knots*** associated to the CC and GK quadrature sums.

- Notice that $f_0^l=f(\alpha),\ f_n^l=f_n^r=f(\gamma),\ f_0^r=f(\beta)$ are ***inherited*** from ancestor subranges while at $0<\eta_k<1$, values f_k^l, f_k^r are computed at ***each*** attempt to evaluate $I[\alpha, \beta]f$.
- **Definition.** The ***integrand profiles over half-subranges*** consist of appropriately chosen ***sets of pairs*** $\{\eta_k, f_k^l\}$ and $\{\eta_k, f_k^r\}$ respectively, including those coming from the abscissas pairs $\{\alpha, \gamma\}$ and $\{\gamma, \beta\}$.
- Three kinds of integrand profiles are of interest for Bayesian analysis:
  - those ***involving the union of the CC-32 and GK-21 knots*** (enable Bayesian diagnostics on the conditioning of $f(x)$ over $[\alpha, \beta]$;
  - those ***involving the CC-32 knots*** (enable inferences concerning the quality of q_{CC} as a Riemann integral sum)
  - integrand profile pieces over either lateral or central close proximity neighbourhoods of abscissas $x^i\in\Pi_N$.

Tools for Bayesian Analysis: Finite Differences

- Inferences on the ***integrand conditioning*** are got corroborating theorems of the calculus with features of the ***integrand profile*** concerning:
  - the ***sign*** and ***rate of variation*** of the integrand ***slope***
  - the ***sign*** of the ***integrand curvature.***
- The ***first order divided differences,*** defined over ***pairs*** of ***adjacent*** integrand profile abscissas $\{(x_\lambda, f_\lambda), \lambda = k\text{-}1, k\}$ provide local approximations of the ***integrand slope.***
  A ***denominator-insensitive-to-subtraction*** definition is provided by the use of reduced modified quadrature knots instead of the algebraic abscissas x_λ,
  $$d_{k-1,k} = \delta_{k-1,k} / (\eta_k - \eta_{k-1}); \quad \delta_{k-1,k} = f_k - f_{k-1}.$$
- The ***sign of the curvature*** of $f(x)$ may be inferred from the ***variation*** of the first order divided differences over ***triplets*** of ***adjacent*** abscissas $\{(x_\lambda, f_\lambda), \lambda = k\text{-}1, k, k+1\}$.
  If $\{x_{k-1}, x_k, x_{k+1}\} \in [x^{i-1}, x^i] \subset \Pi_N$, then we get simply
  $$\delta_{k-1,k,k+1} = d_{k,k+1} - d_{k-1,k}.$$
  If, however, $x_k \equiv x^i \in \Pi_N$ and x^i is ***two sided***, then:
  - inside $(x_k, x_{k+1}) : \delta_{k-1,k,k+1} = d_{k,k+1} - \tilde{\rho} d_{k-1,k}$
  - inside $(x_{k-1}, x_k) : \delta_{k-1,k,k+1} = \tilde{\rho}^{-1} d_{k,k+1} - d_{k-1,k}$,

  where $\tilde{\rho} = (x^{i+1} - x^i)/(x^i - x^{i-1})$.

Heuristic Equation to Zero of Negligible Floating Point Differences

- **Implementation steps**
  - ***Given*** the input pair of floating point computed values v_1, v_2;
  - ***Compute*** their difference $d_{12} = fl(v_2 - v_1)$;
  - ***Check*** negligibility criterion: $\mathrm{ngl} = \{|v_1|, |v_2|, u_r\} \cdot \varepsilon_{\mathrm{roff}} - |d_{12}| > 0$;
  - ***Conditional*** equation to zero: <u>If</u> (ngl) $d_{12} = 0$.

  Hardware environment: $\varepsilon_{\mathrm{roff}} = \mu\varepsilon_0 \quad (\mu = 4.)$; $u_r = u/\varepsilon_{\mathrm{roff}}$,
  where ε_0 is the machine epsilon, u is the machine underflow.
- **Necessity.** To make ***easy*** and ***stable*** Bayesian inferences concerning:
  - ► ***Assessment*** of the ***relative magnitudes*** of the coefficients of the Chebyshev series expansions of the integrand within the fast Chebyshev transform with backward summation
  - ► ***Rough localization*** and ***conditioning assessment*** of resolved monotonicity interval ends within integrand profiles
  - ► ***Identification*** and ***extension sharpening*** of monotonicity intervals of vanishing curvature
  - ► ***Sharpening localization*** of the inferred points of discontinuity, either at the ends of or inside the monotonicity intervals.
- **Feasibility:** within the numerical quadrature of ***difficult*** quadrature problems, ***small*** quantities always occur together with ***large*** quantities, which ***dominate*** the output quality assessment and the accuracy of the decision processes.

Hardware and Software Conformity to the IEEE 754 Standard

The* IEEE 754 *Standard governs the *binary floating point arithmetic* (number formats, basic operations, conversions, and exceptional conditions).

Hardware and ***Software*** conformity to the standard allows doing a ***reliable*** and ***portable*** *analysis*.

Critical points where we have found ***deviations*** from the standard:

i. Floating point ***comparison*** operations involving RAM and processor machine numbers of a ***same*** real number which does not equate the approximating floating point numbers;

ii. Unpredictable effects following from ***compiler code optimization***.

Overview

- Introduction
- Hardware Environment
- Bayesian Steps of the Analysis
- New Priority Queue Perspective
- A Few Case Studies

New Priority Queue Perspective

- *The (**enhanced**) **standard automatic adaptive quadrature** is based on a **priority queue** which uses a **simple key***
- *Within the **Bayesian approach** a **hierarchically distributed tree structure** of the terminal subranges is obtained.*
  *It consists of **four classes** of subranges each of which is characterized by different integral processing and priority queues governing the subrange subdivision:*
  1. *subranges in **undefined state**;*
  2. ***endpoint singular** subranges;*
  3. ***well-conditioned** subranges;*
  4. ***ill-conditioned** subranges*

Class: Undefined State

- ***Definition:*** ***Undefined state:*** *Characterizes the* ***input integration domain*** *prior to the attempt to generate quadrature knots over it,* ***or*** *the conditioning check has met a criterion pointing to an* ***insufficiently resolved integrand profile****.*
- ***Features:***
  - ⟶ *Defines a* ***transient state*** *in the priority queue*
  - ⟶ *Each subrange in undefined state has the* ***highest priority***
  - ⟶ *All subranges falling in this class are* ***candidates on equal footing*** *to analysis and validation (possibility for* ***parallel*** *processing)*

Class: State with Endpoint Singularity

- ***Definition:*** ***Endpoint singularity:*** *May arise as a solution of a boundary layer problem under ill-conditioning diagnostic*
- ***Features:***
  - → *Defines a **transient state** in the priority queue*
  - → *Each subrange in singular state is **pending** until all the undefined state subranges are resolved*
  - → *All subranges falling in this class get corroborated with each other local accuracy assignments and **parallel** processing is used to solve them*

Class: Well-Conditioned State

- ***Definition:*** ***Well-conditioned state:*** *All conditioning check criteria ended successfully for undefined state input* ***or*** *the current descendent comes from a well-conditioned parent.*
- ***Features:***

  → *The priority of the well-conditioned subranges comes next after that of the endpoint singular subranges*

  → *The integrand profiles of the well-conditioned subranges coming from undefined state are stored and processed in parallel*

  → 1. *The* ***global*** *pair* $\{Q, E > 0\}$ *is updated and end of computations* (eoc) *is checked*

  → 2. **If** (.NOT. eoc) **then** *the assessment of the* ***significance*** *of the local error estimates is set over the manifold of well-conditioned subranges. The priority queue includes exclusively significant well-conditioned subranges.*

  → 3. ***Parallel*** *processing is activated for:* = subrange *subdivision by* ***bisection;*** = *activation of the* ***local quadrature rules***

  → ***Go to step*** 1.

Class: Ill-Conditioned State

- ***Class rank in the priority queue:***

  *the ill-conditioned (irregular) subrange class occupies the **lowest rank** in the priority queue.*

  A low degree quadrature sum is used for the derivation of an estimate of the value of the integral over the irregular subranges.

- ***Bayesian inference:***

  *the existence of representative subranges of such a class points to the **impossibility** to solve the given integral unless it is further analyzed and reformulated.*

Overview

- Introduction
- Hardware Environment
- Bayesian Steps of the Analysis
- New Priority Queue Perspective
- A Few Case Studies

Errors associated to AAQ solution of fundamental power series integrals

$$\int_0^1 x^n dx = \frac{1}{n+1}, \ n = 0, 1, \cdots, 200$$

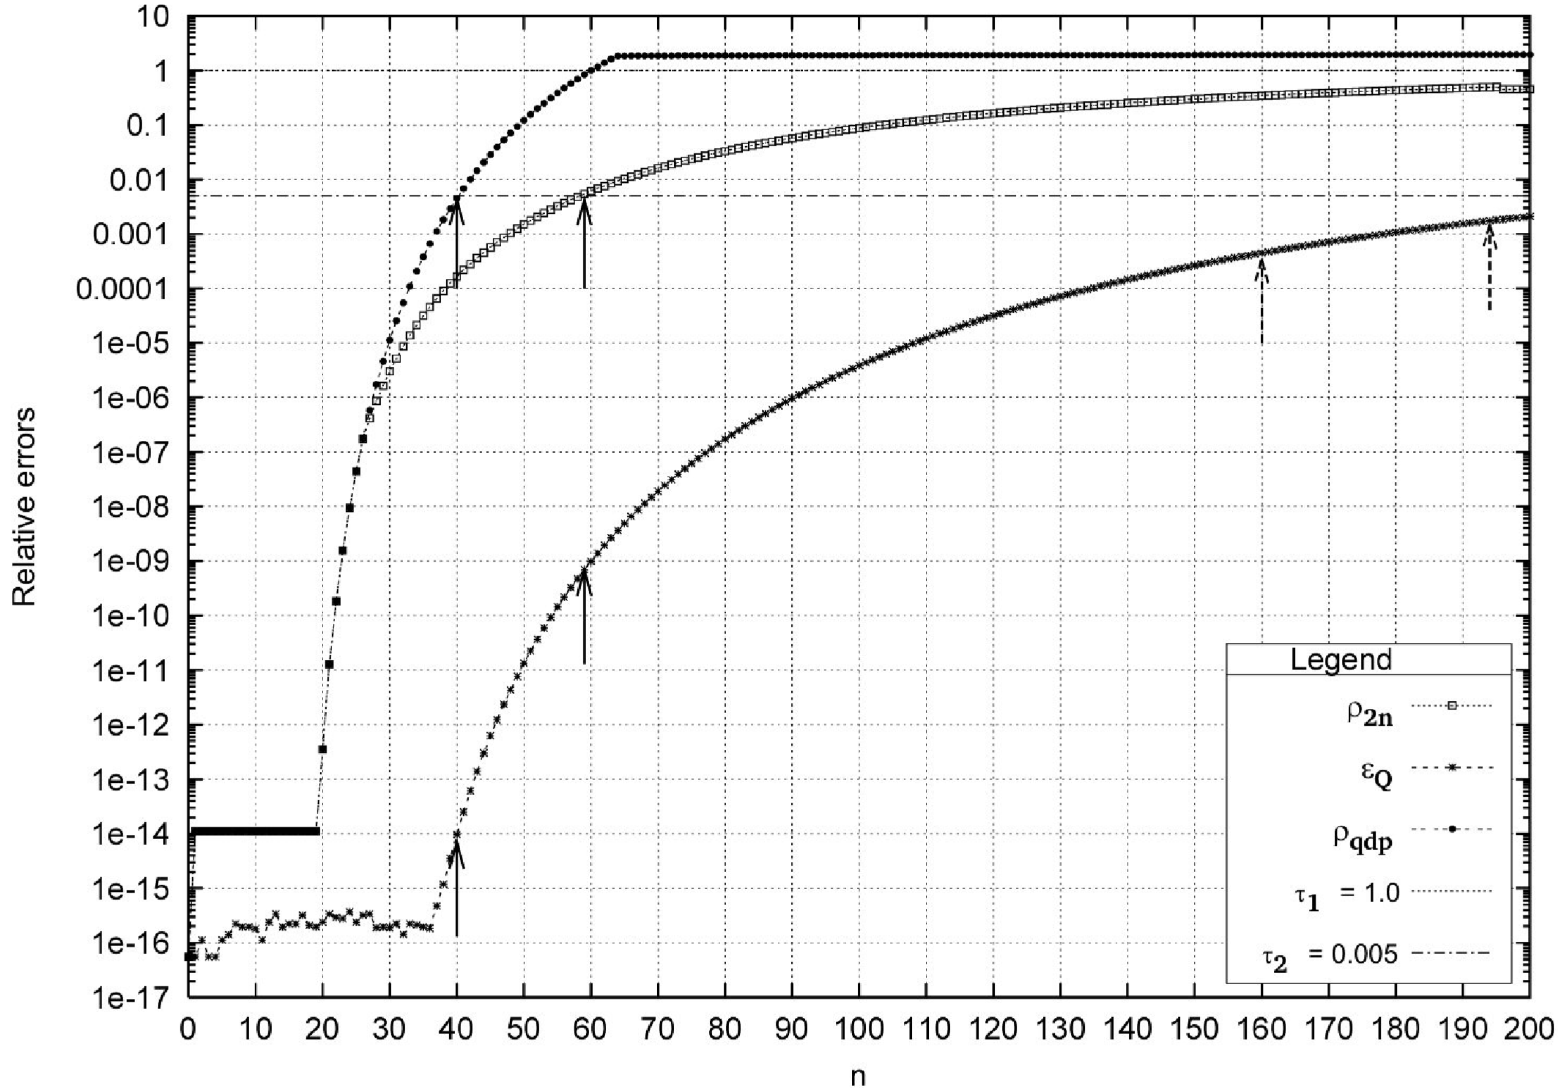

Errors associated to AAQ solution of integrals involving cut centrifugal-like potential

$$\int_0^b \frac{1}{x^2+1}\,dx = \arctan(b),\ \ b \in [1, 260]$$

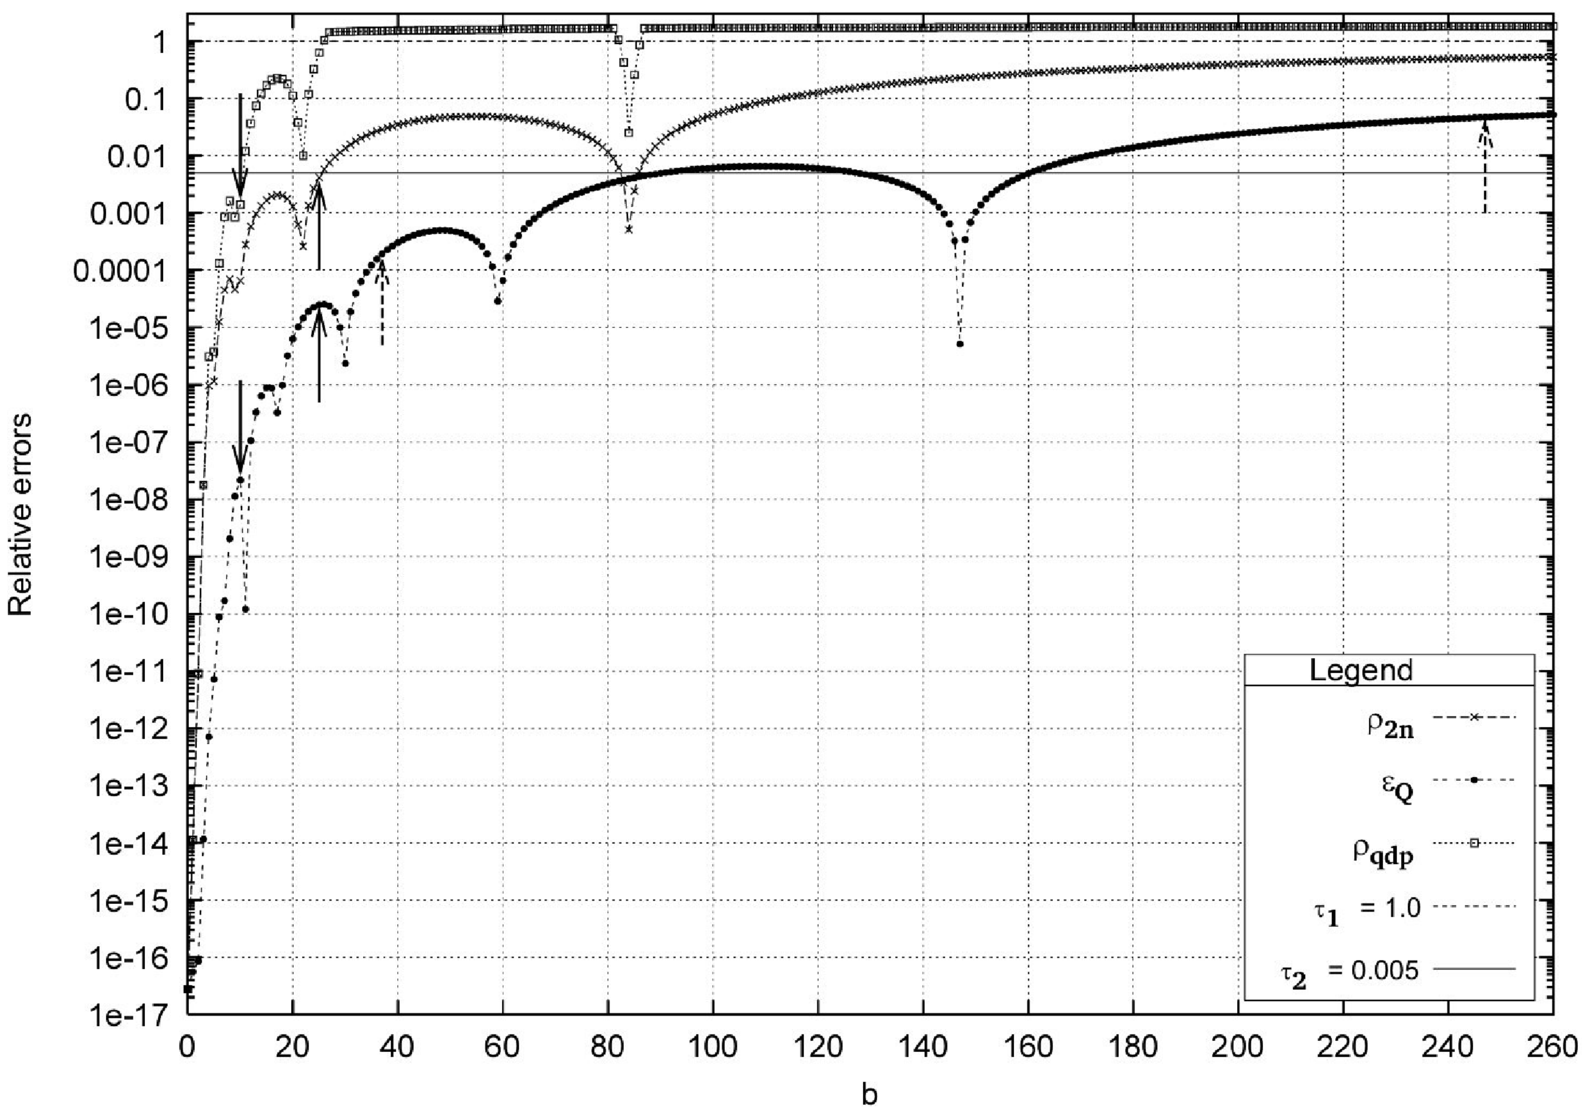

Errors associated to AAQ solution of the family of oscillatory integrals

$$\int_{-1}^{1} e^{p(x-x_0)} \cos(\omega x)\, dx = 2e^{-px_0}[p \sinh(p)\cos(\omega) + \omega \cosh(p)\sin(\omega)]/(\omega^2 + p^2)$$

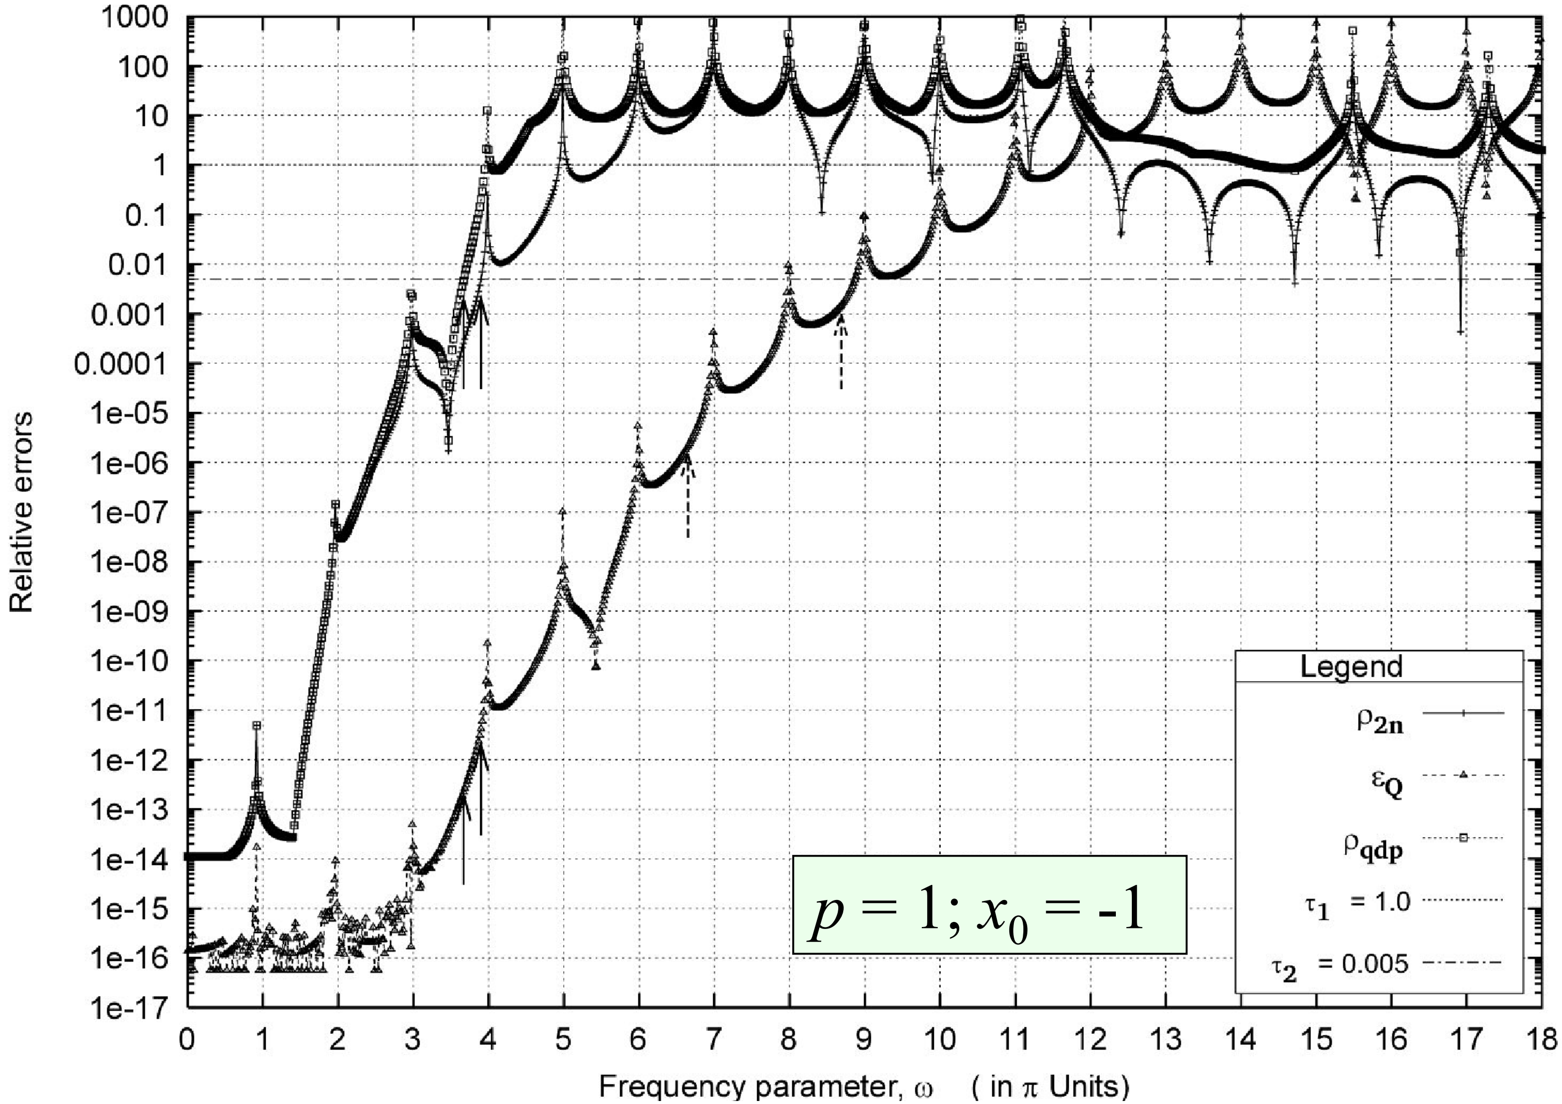

Errors associated to AAQ solution of the family of oscillatory integrals

$$\int_0^1 2e^{-px_0}\cosh(px)\cos(\omega x)\,dx = 2e^{-px_0}[p\sinh(p)\cos(\omega)+\omega\cosh(p)\sin(\omega)]/(\omega^2+p^2)$$

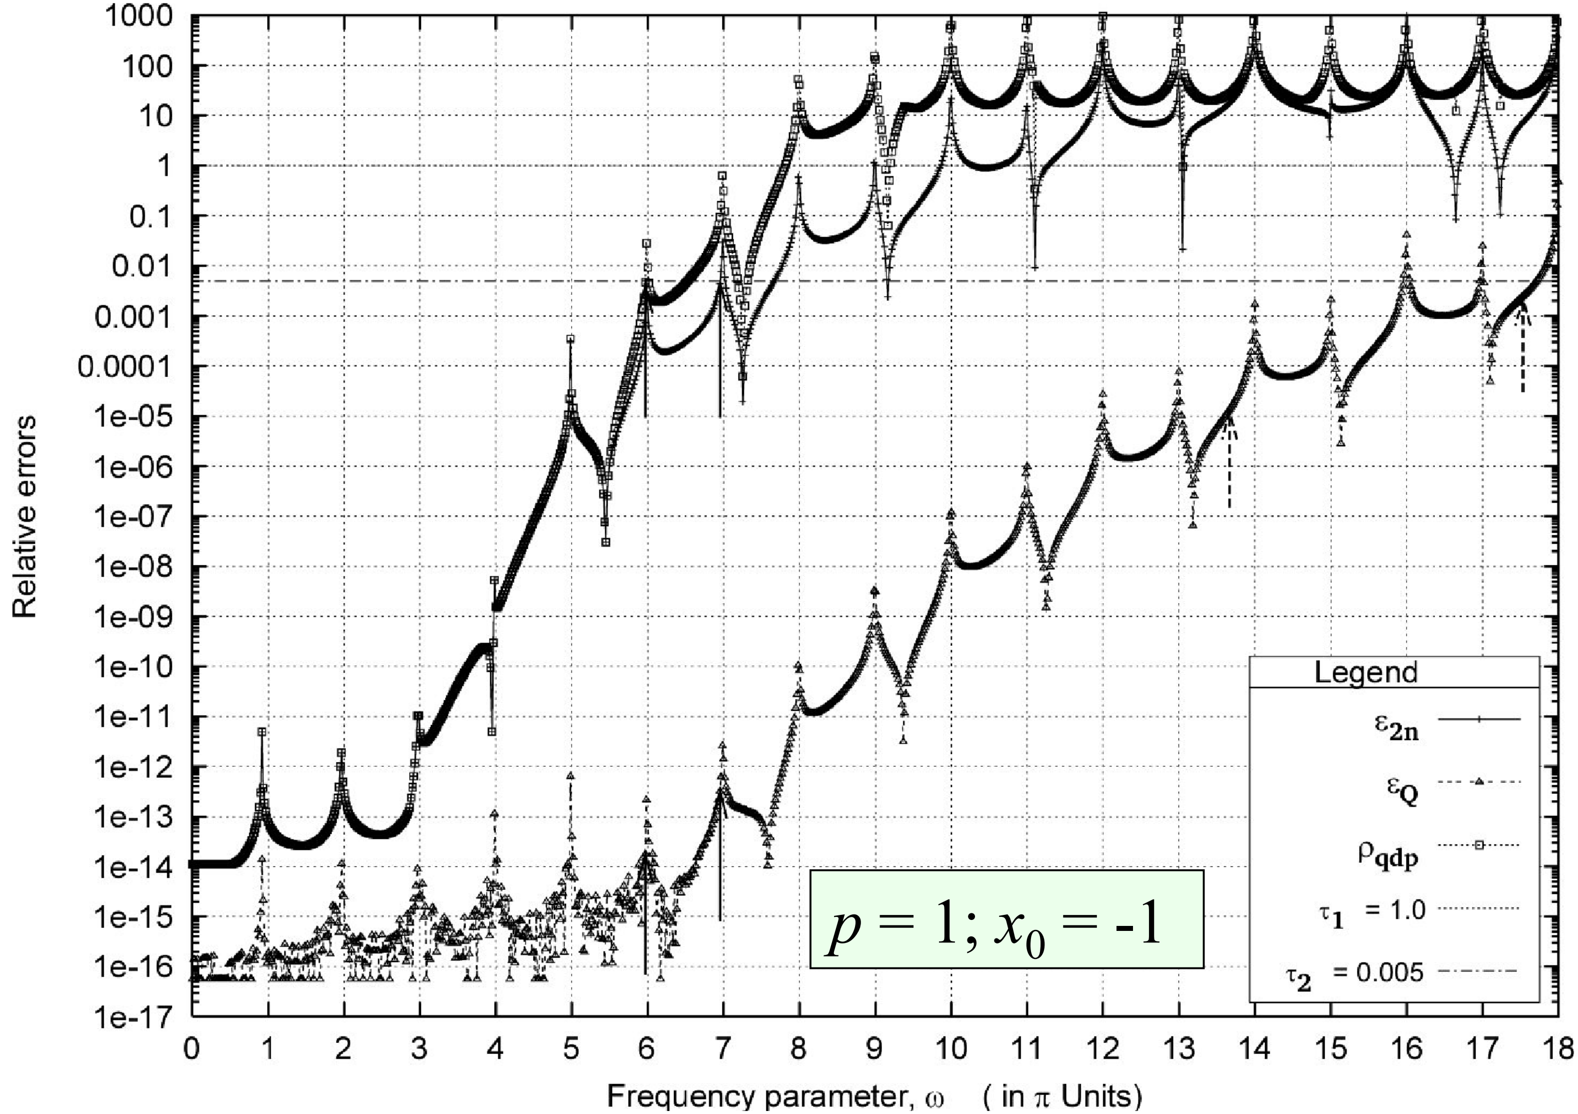

Errors associated to AAQ solution of the family of oscillatory integrals

$$\int_{-1}^{1} e^{p(x-x_0)} \sin(\omega x)\, dx = 2e^{-px_0}[p\cosh(p)\sin(\omega) - \omega\sinh(p)\cos(\omega)]/(\omega^2 + p^2)$$

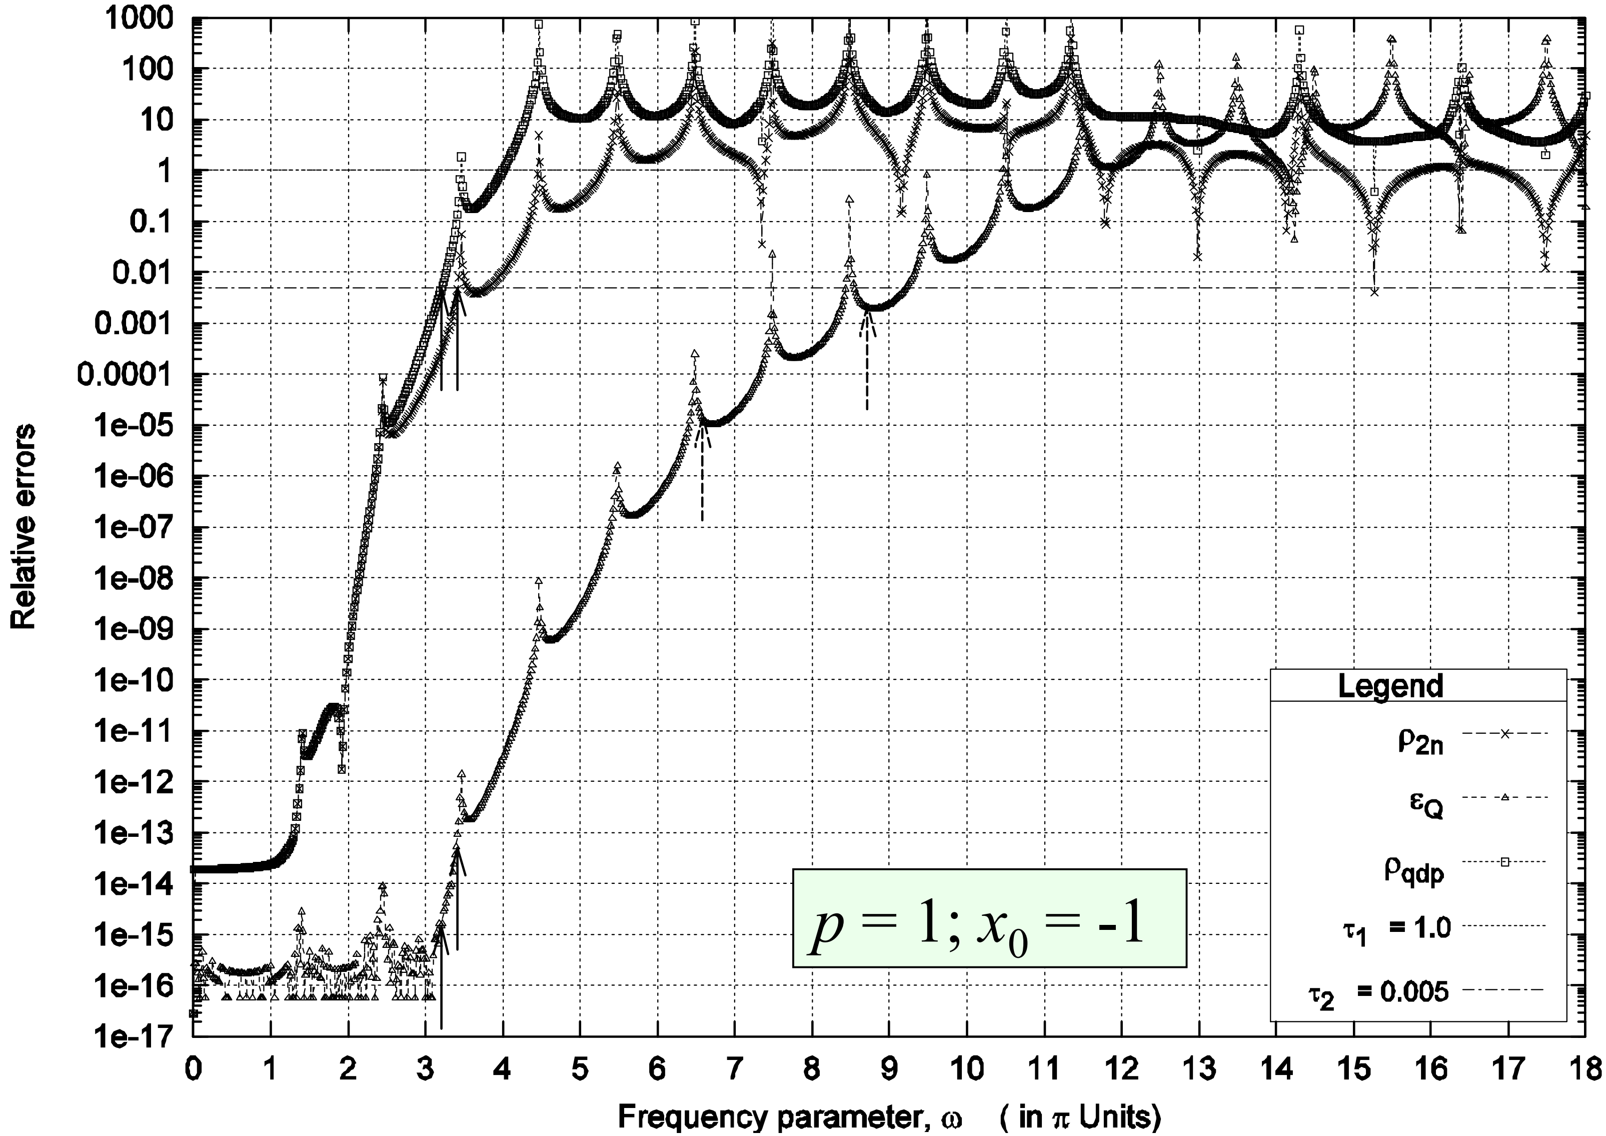

Errors associated to AAQ solution of the family of oscillatory integrals

$$\int_0^1 e^{-px_0}\sinh(px)\sin(\omega x)\,dx = 2e^{-px_0}[p\cosh(p)\sin(\omega)-\omega\sinh(p)\cos(\omega)]/(\omega^2+p^2)$$

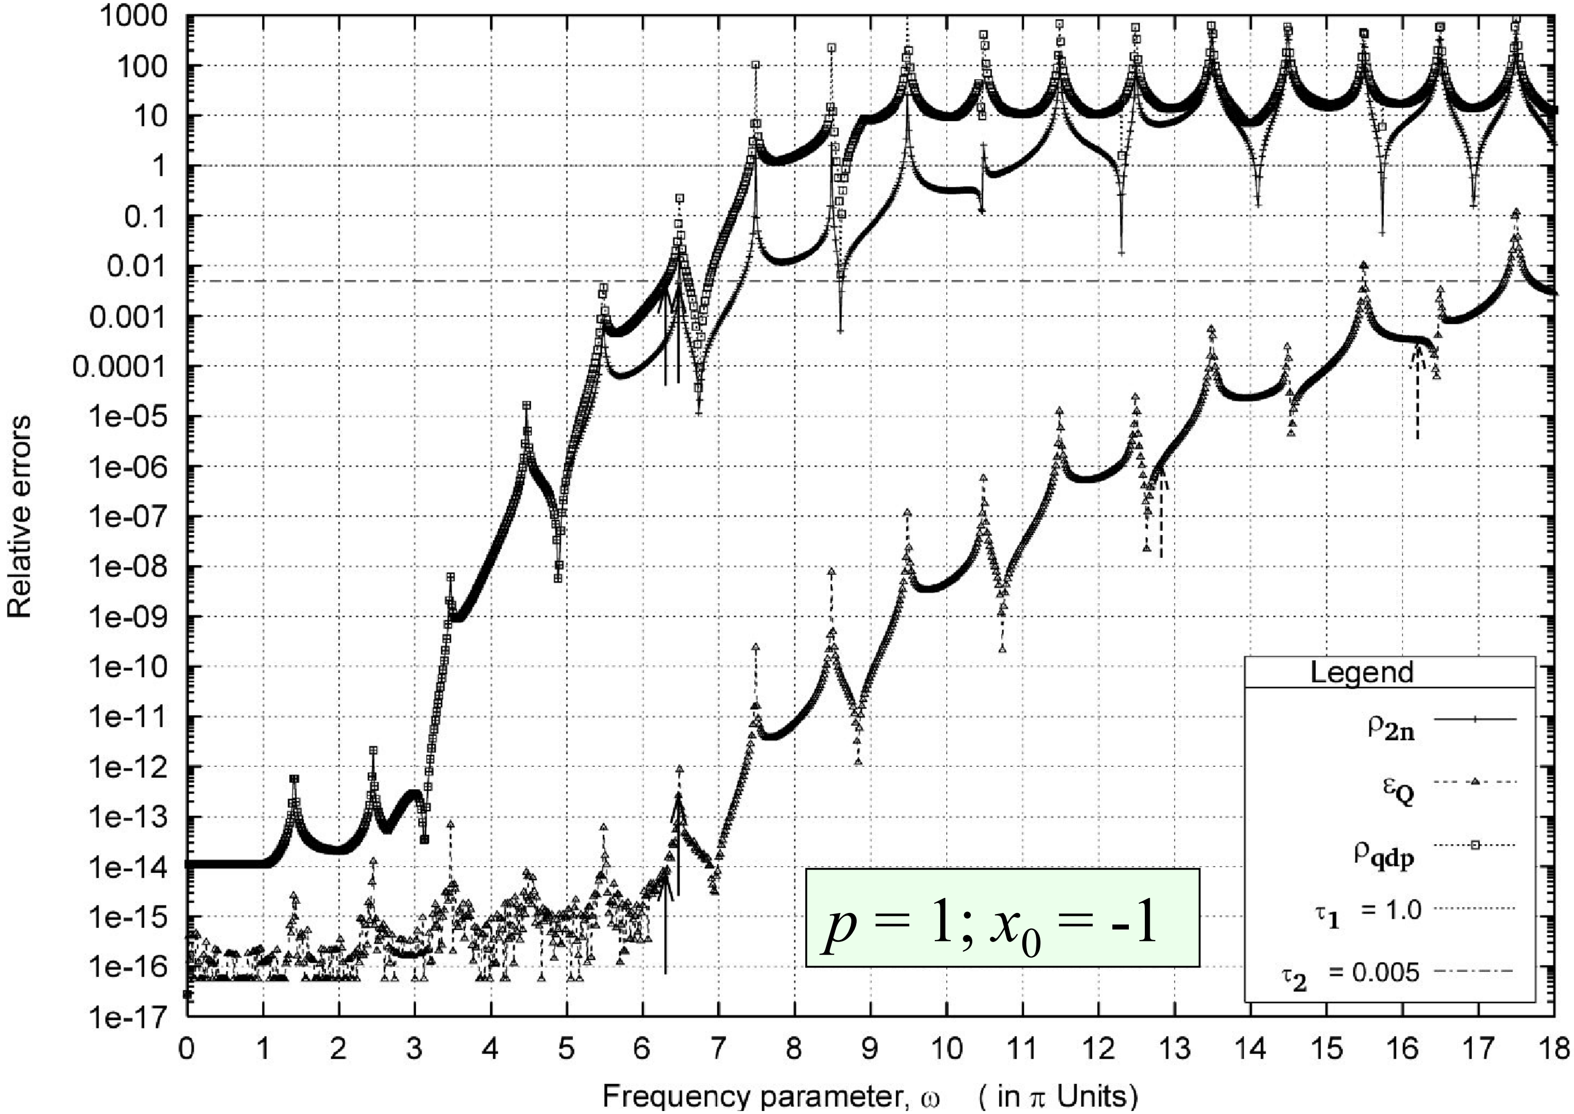

Standard Input Numerical Problem

Given the proper (or improper) Riemann integral

$$I[a,b]f = \int_a^b w(x) f(x)\,\mathrm{d}\,x, \ -\infty < a < b < +\infty$$

we seek a ***globally adaptive*** numerical solution $\{Q, E > 0\}$

of it within input accuracy specifications $\{\varepsilon_r > 0, \varepsilon_a \geq 0\}$

i.e.,

$$|\, I[a,b]f - Q \,| < E < \max\{\varepsilon_r \,|\, I[a,b]f \,|, \varepsilon_a\}$$
$$\approx \max\{\varepsilon_r \,|\, Q \,|, \varepsilon_a\}$$

See Quadpack, NAG, etc.

Reliable Input Numerical Problem

Given the proper (or improper) Riemann integral

$$I[a,b]f = \int_a^b w(x) f(x)\,\mathrm{d}\,x, \; -\infty < a < b < +\infty$$

we seek a ***globally adaptive*** numerical solution $\{Q, E > 0\}$

of it within input accuracy specifications $\{\varepsilon_r > 0, \varepsilon_a \geq 0\}$

i.e.,

$$|I[a,b]f - Q| < E < \max\{\varepsilon_r\,|I[a,b]f|, \varepsilon_a\}$$

$$\approx \min\{\varepsilon_r^{\max}, \max\{\varepsilon_r\,|Q|, \varepsilon_a\}\}$$

Oscillatory Integrand

$$\int_0^b \frac{\cos(\omega x)}{x^2+1}dx = \frac{1}{2}\left(1+\epsilon(b)\right)\pi e^{-\omega}$$

- QUADPACK DQAWO (NAG D01ANF): out of 637 outputs, 505 spurious

Reason: decision of the control routine to activate the ε algorithm.

- Present: *all outputs reliable*

Singular Integrand

$$\int_0^1 \frac{dx}{\sqrt{x^2+2x-2}} = \pi/2 - \arctan(\sqrt{2}/2) + \log(3)/2$$

Singularity at $x_s = \sqrt{3} - 1$.

• QUADPACK DQGKS (using GK10-21): out of 105 outputs, 42 spurious

• Number of integrand evaluations needed to get ten figure accuracy:

– QUADPACK DQGKP (using GK10-21) 294 (requires specification of singularity as input)

– Present (using GK10-21) 391 (from which 97 needed to identify and resolve singularity to machine precision)

Singular Integrand

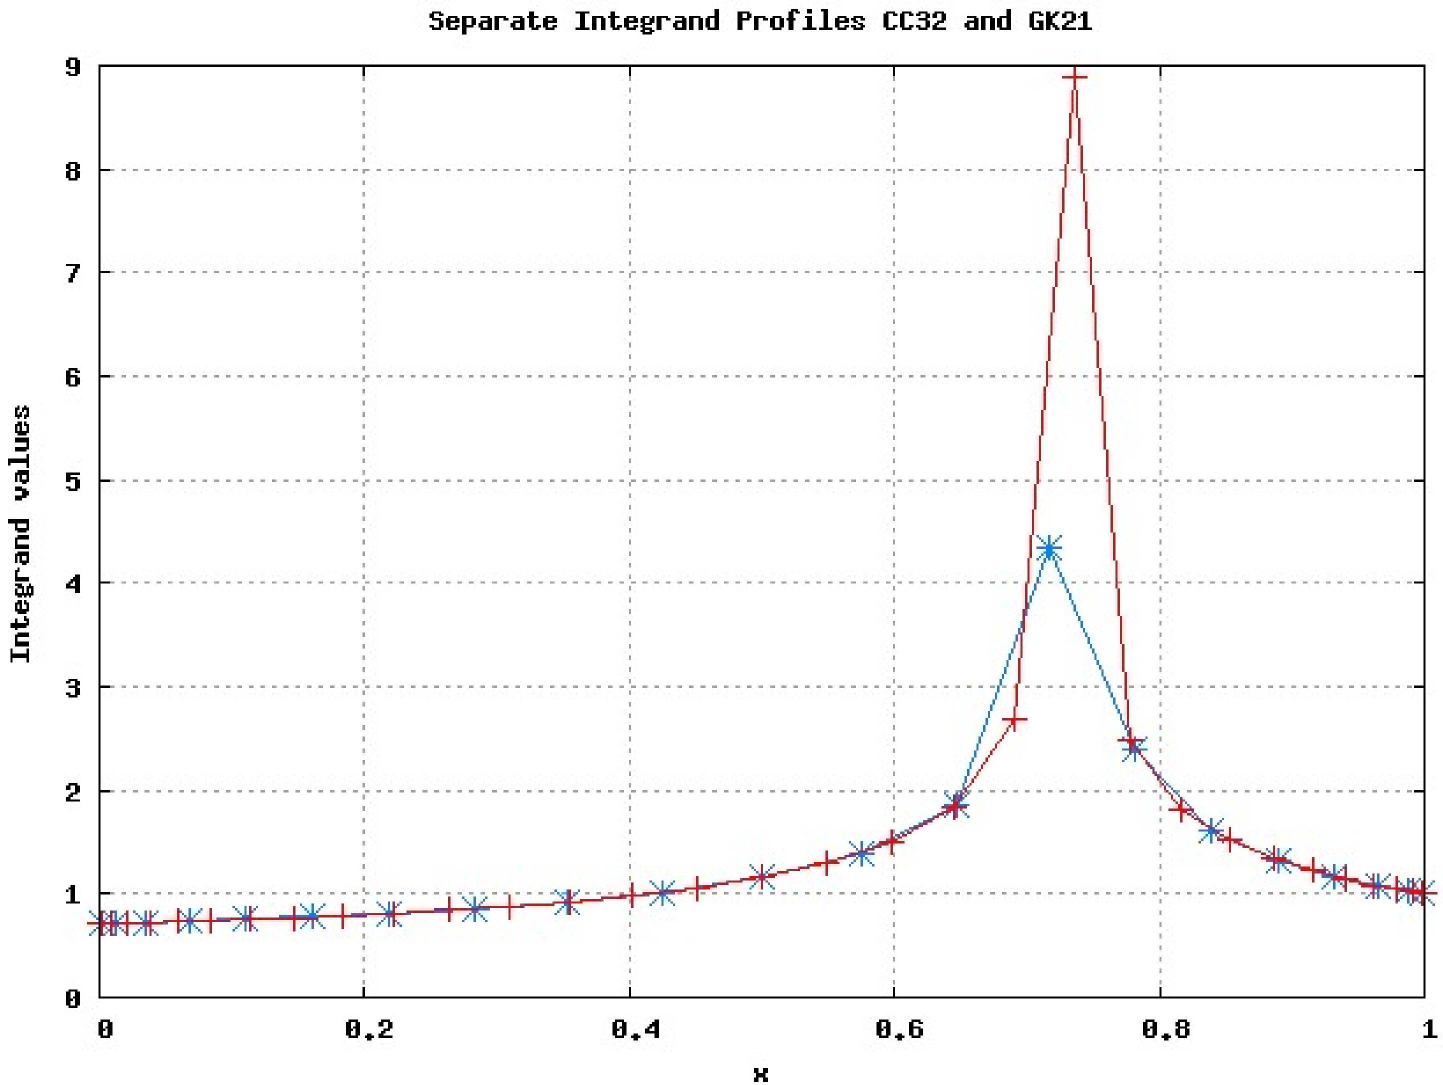

Singular Integrand

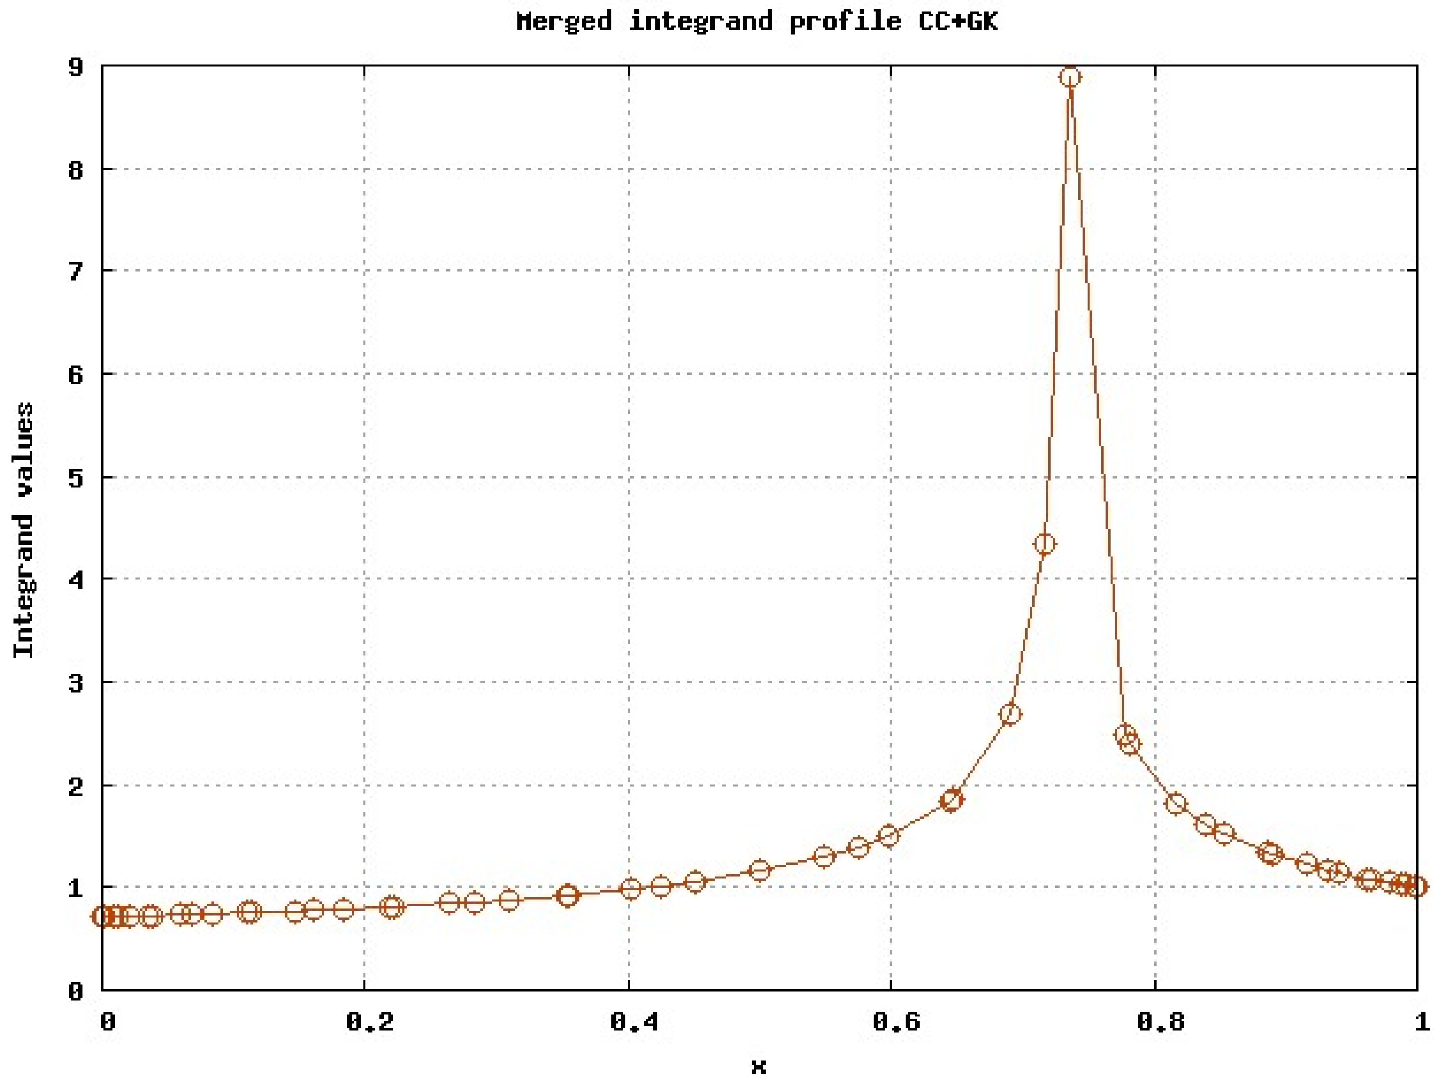

Singular Integrand

• • • *Singularity is not a machine number*:

$$f(x) = \frac{1}{\sqrt{|x^2 + 2x - 2|}}; \quad \text{range} = \ [a, 1]\ ,$$

where a runs around the *singular point* $x_s = \sqrt{3} - 1$. Analysis has to face catastrophic cancellation by subtraction.

• • • Algebraically equivalent problem:

$$f(t) = \frac{1}{\sqrt{|t^2 + 2\sqrt{3}t|}}; \quad \text{range} = \ [a, 1]\ ,$$

where a runs around the *singular point* $t_s = 0$, gets the analysis free of cancellation effects.

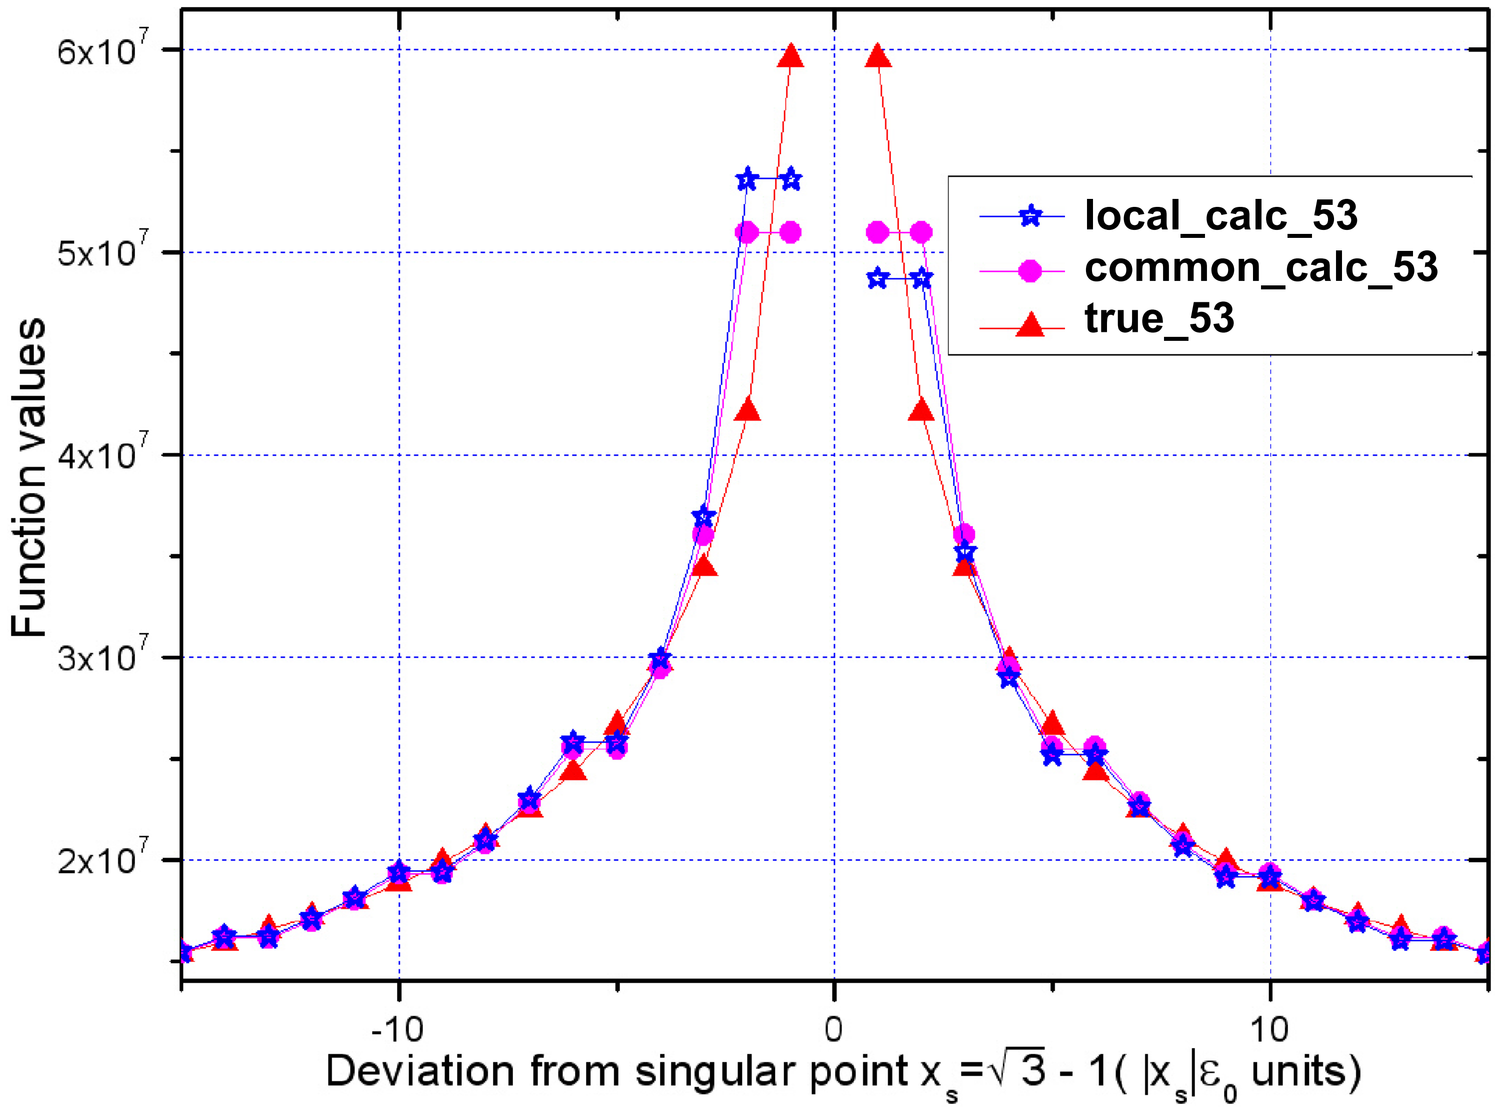

Illustration of inappropriateness of the value $\boldsymbol{\varepsilon_0 = 2^{-53}}$

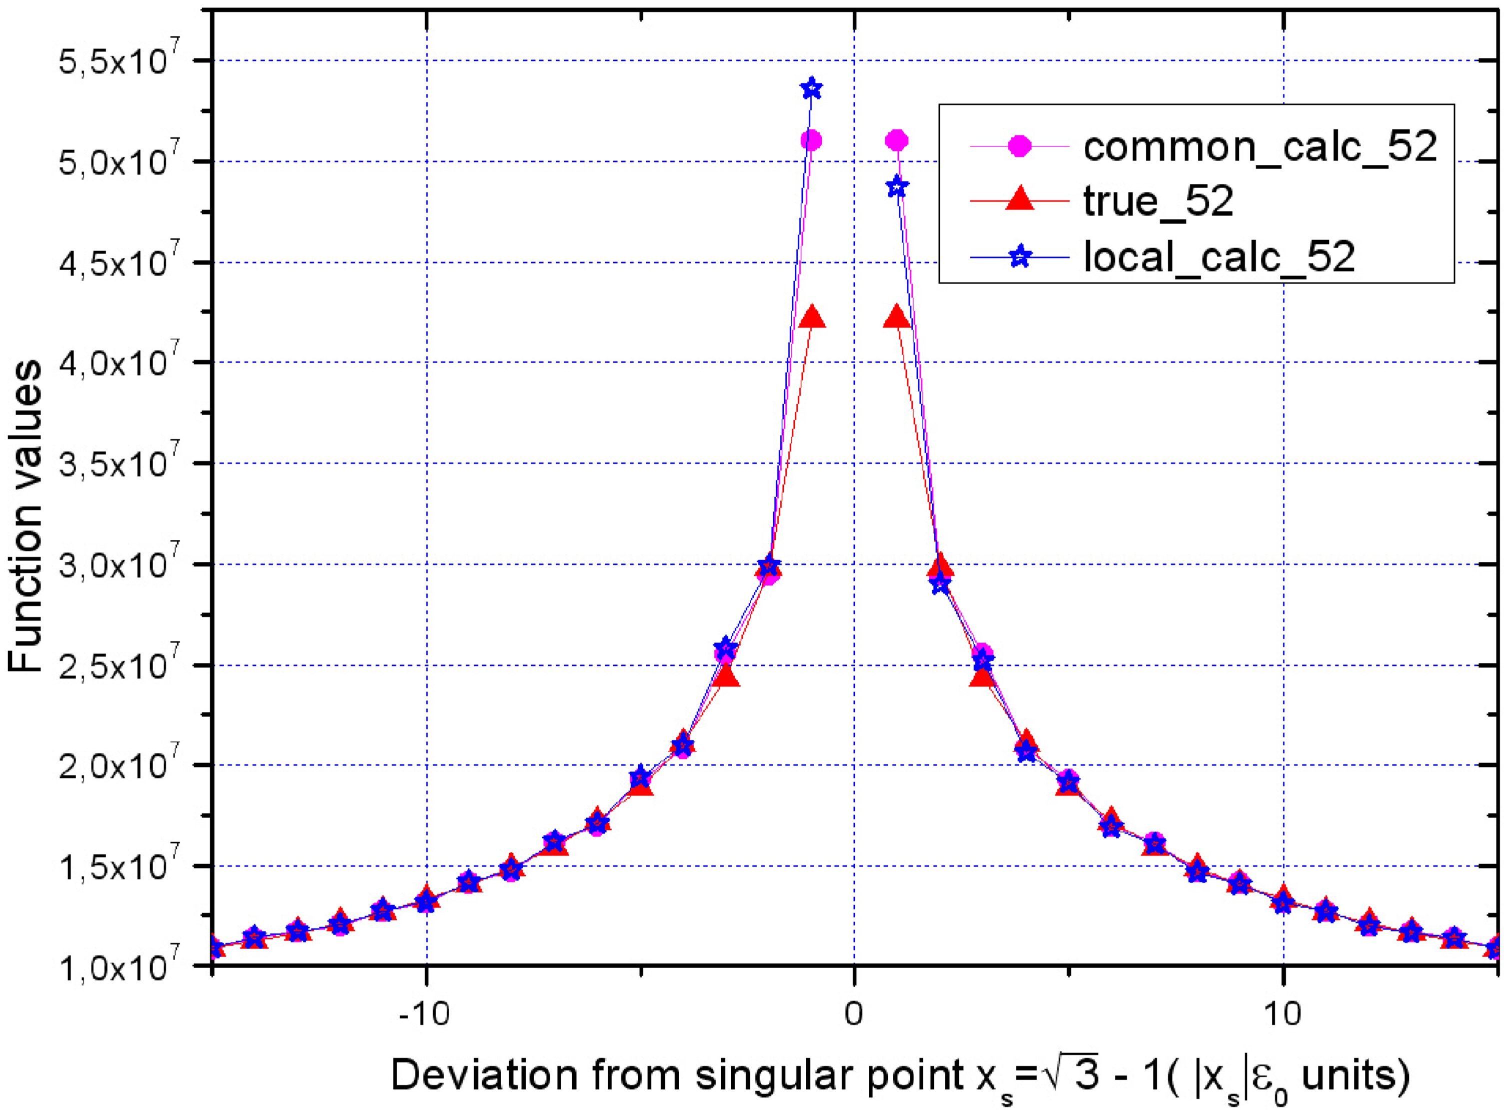

The value $\varepsilon_0 = 2^{-52}$ is right !

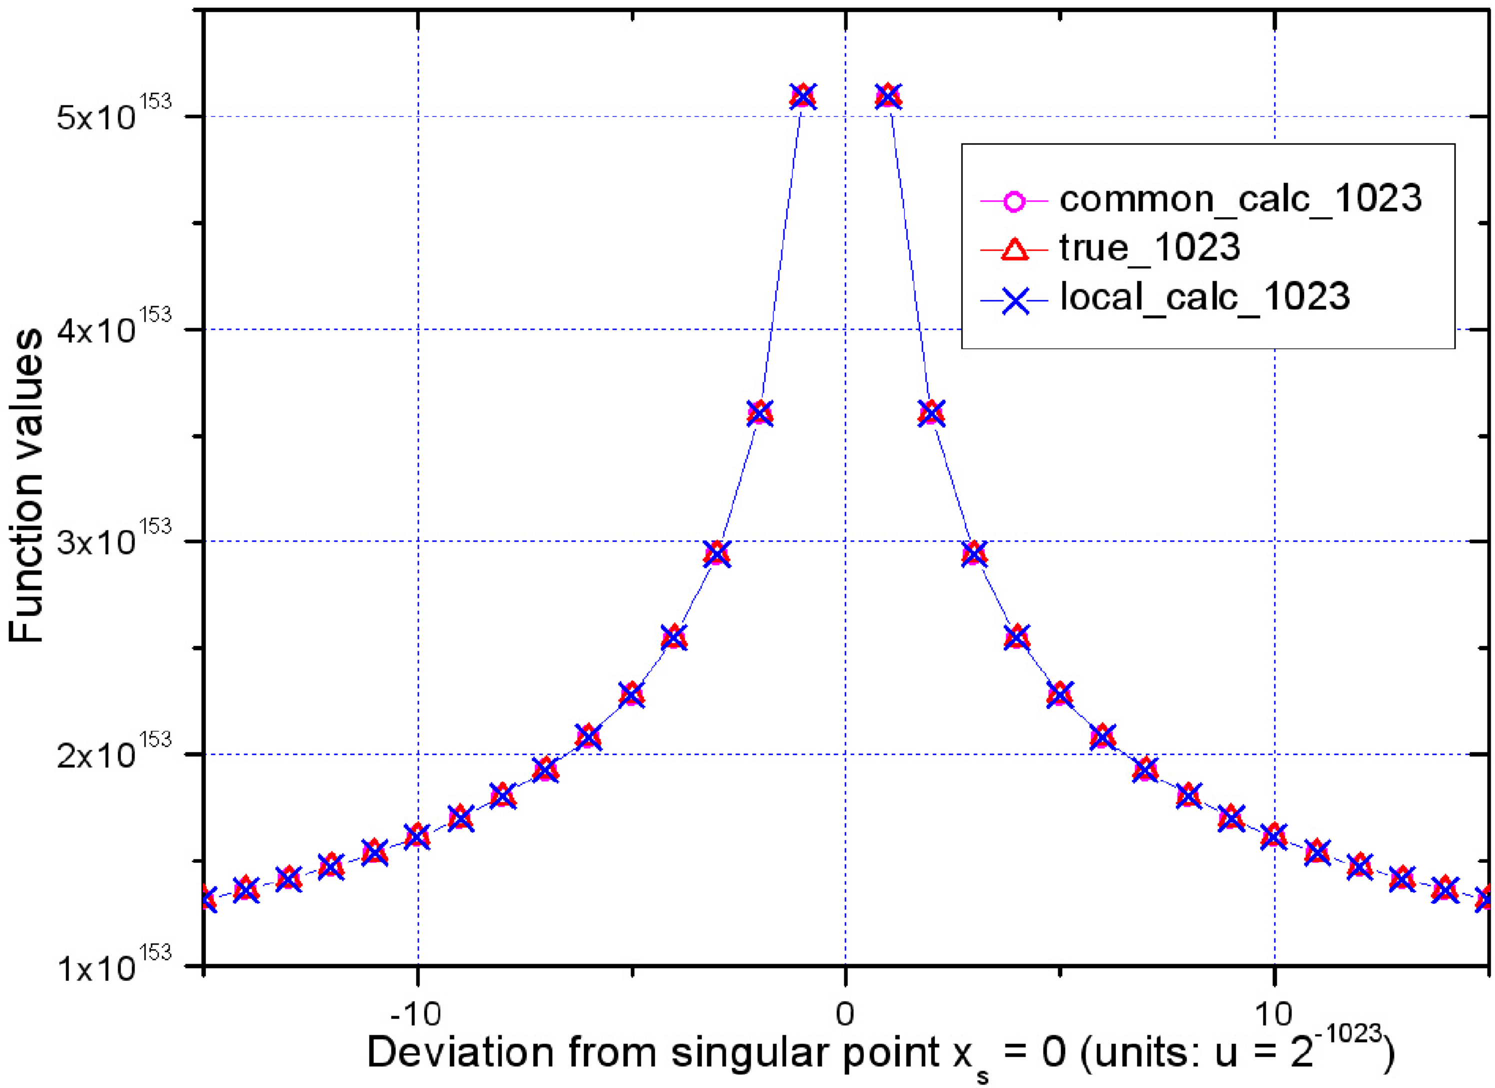

Moving singularity at $x_s = 0$ solves all troubles

Thank you for your attention !