Mathematical Modeling and Computational Physics (MMCP 2017) *Dubna, July* 3 - 7, 2017

Disentangling Complexity in Bayesian Automatic Adaptive Quadrature

Gh. Adam, S. Adam

LIT JINR-Dubna & IFIN-HH Bucharest

Work supported within JINR topic 05-6-1119-2014/2019 and the Romania-LIT Hulubei-Meshcheryakov Programme, JINR Orders 218/10.04.2017, pp. 22, 23 and 219/10.04.2017, pp. 93-95

Overview

- **General frame of the Bayesian automatic adaptive quadrature (BAAQ)**
- Three classes of integration domain lengths
- Early Bayesian problem conditioning diagnostics
- Conclusions

Mathematical Problem

We consider the BAAQ numerical solution of the Riemann integral

$$I \equiv I[f] = \int_a^b g(x)f(x)dx, \qquad -\infty < a < b < \infty,$$

under the assumption that the real valued ***integrand function*** $f(x)$ is continuous almost everywhere on $[a, b]$ such that I 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 I rests on the use of the m-***panel rule approach*** (use of a fixed quadrature rule over ranges of different lengths) (*).

Given the integration domain $[a, b]$, an ***interpolatory quadrature sum*** is used to get an approximation $Q \equiv Q[f]$ to $I[f]$.

The ***meaningfulnes***s of $Q[f]$ is assessed by deriving a bound $E \equiv E[f]>0$ to the remainder $R[f] = I[f] - Q[f]$.

(*) R. Piessens et al., QUADPACK, A Subroutine Package for Automatic Integration, Springer, 1983, chap.II

For a prescribed accuracy τ requested at input, the approximation Q to I is assumed to end the computation provided

$$|R[f]| < E < \tau.$$

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

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

If the remainder boundedness condition is not satisfied, the AAQ m-***panel rule approach*** to the solution attempts at decreasing the error E by the ***subdivision*** of the integration domain $[a,b]$ ***into subranges*** using ***bisection*** and the computation of a ***local pair*** $\{q, e > 0\}$ over each newly defined subrange $[\alpha, \beta] \subset [a, b]$ by repeated use of a same quadrature rule.

This procedure builds a ***subrange binary tree*** the evolution of which is controlled by an associated ***priority queue***.

Local pairs $\{q_i, e_i > 0\}$ are computed over the i-th subrange of $[a, b]$ and ***global*** outputs $\{Q_N, E_N > 0\}$ are got by summing the results obtained over the N existing subranges in $[a, b]$.

After each subrange binary tree update, the termination criterion is checked until it gets fulfilled, or it is decided that the given integral cannot be solved.

The existing strict mathematical bounds to $R[f]$ are of little use in the implementation of practical codes.

The derivation of a ***practical bound*** $e > 0$ to q rests on ***probabilistic arguments*** the validity of which is always subject to doubt.

The BAAQ advancement to the solution rests on ***Bayesian inference*** based on four pillars:

- theory of the Riemann integral
- theory of the numerical integration (quadrature)
- features of the floating point computation
- accumulated empirical evidence

Essentially, the probabilistic character of the AAQ approach 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.

Overview

- General frame of the Bayesian automatic adaptive quadrature
- **Three classes of integration domain lengths**
- Early Bayesian problem conditioning diagnostics
- Conclusions

The so far reported automatic adaptive quadrature algorithms implemented interpolatory quadrature sums characterized by ***high algebraic degrees of precision***. 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*** which ***significantly varies*** with the ***extent*** and the ***localization*** of the ***floating point representation*** of the ***ends*** of the ***integration domain within the set of the machine numbers***.

Since the floating point degree of precision decreases dramatically under the decrease of the (***absolute/relative***) length of the integration domain, the use of quadrature sums of high algebraic degrees of precision gets inappropriate at small ***absolute/relative*** (whichever is the smallest) integration domain lengths.

[(*)] S. Adam, Gh. Adam, Springer LNCS 7125, 189-194 (2012)

Features of the Floating Point Degree of Precision

- *Gliding integration range* [0,1] *on the real axis*

The following plot gives outputs for the family of 1023 integration ranges

$$\{[j\text{-}1, j];\quad j = 1, 2, \ldots, 1023\}$$

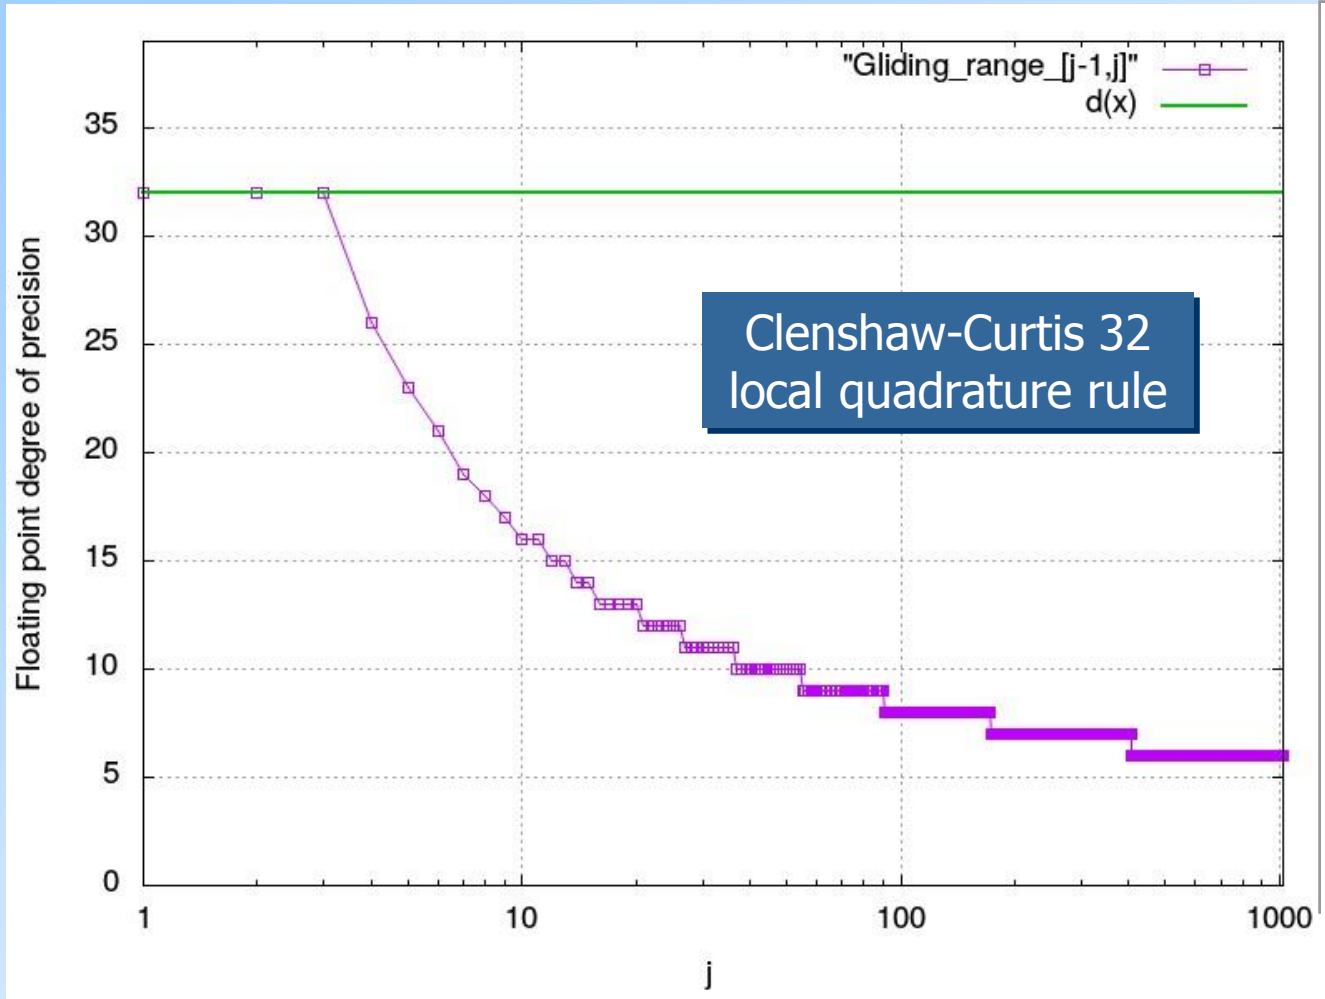

Variation of the floating point degree of precision of the CC-32 local quadrature rule over the gliding range **[0, 1]** ***versus its distance j from the origin.*** It is shown that $d_{fp} = d = 32$ at low j values ($j = 1, 2, 3$), then d_{fp} abruptly decreases at larger but small enough j, to show slower decreasing rates under the displacement of [0,1] far away from the origin, reaching a bottom value $d_{fp} = 6$ at $407 \leq j \leq 1023$.

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

$$X = \max\{fl(|\alpha|), fl(|\beta|)\}, \mathrm{X} > 0, \rho = fl(|\beta - \alpha|/X), 0 < \rho \leq 2$$

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

It was argued (*) that there can be defined three classes of quadrature sums, each of which is appropriate over integration domain ranges separated by two empirical thresholds, τ_μ and τ_m , as follows:

- ***Microscopic ranges***, characterized by the threshold conditions

$$0 < \min(X, \rho) \leq \tau_\mu = 2^{-20}.$$

(Composite) ***Simpson*** and ***trapezoidal rules*** can be used.

- ***Mesoscopic ranges***, characterized by the threshold conditions

$$\tau_\mu < \min(X, \rho) \leq \tau_m = 2^{-6}.$$

(Composite) ***4-interval Newton*** and ***Simpson rules*** can be used.

- ***Macroscopic ranges***, characterized by the threshold condition

$$\min(X, \rho) \leq \tau_m.$$

Quadrature sums of high algebraic degrees of precision can be used. Of special interest is the ***Clenshaw-Curtis quadrature*** using as knots the extremal points of the 32-nd degree Cebyshev polynomial of the first kind.

(*) Gh. Adam, S. Adam, EPJ-WoC vol. 108, 02002 (2016)

Minimizing the Precision Loss over Integration Ranges of Microscopic or Mesosocopic Lengths

- ***Standard Input*: Ends of the integration range. Then $[\alpha, \beta] \subset \mathbb{R}$ 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 over $[\alpha, \beta]$ is transformed to

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

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***.

- ***Alternative*: provide *α* and *h* as inputs** defined to machine accuracy.

No precision loss happens in this case!

Well-Conditioned Integration Rules over Mesosocopic Ranges

An n-interval *equally spaced* partition of $[a,b]$ is performed: $x_k = a + k \cdot h$ where $h = b - a$ denotes the integration domain length. At $n = 2$ we have the *Simpson rule*

$$q_S[a,b] = h\left(f_1 + \frac{1}{6}c_1\right)$$

while at $n = 4$ we have the 4-*interval Newton rule*

$$q_{N,4}[a,b] = \frac{b-a}{2}\left[(f_1 + f_3) + \frac{2}{15}(c_1 + c_3) + \frac{1}{45}(c_1 + c_2 + c_3)\right]$$

Here $c_k = f_{k-1} - 2f_k + f_{k+1}$ is proportional to the *integrand curvature* at k-th abscissa.

From the 4-interval Newton rule and the *composite Simpson rule* over its corresponding abscissa set, we get an error estimate associated to the 4-interval Newton rule:

$$e_{N,4}[a,b] = \frac{b-a}{180}[3|c_1 + c_3| + 2|c_1 + c_2 + c_3|].$$

Clenshaw-Curtis Quadrature over Macroscopic Ranges

The *computation of the Riemann definite integral*

$$I = \int_a^b g(x)f(x)dx, \qquad -\infty < a < b < \infty,$$

by Clenshaw Curtis (CC) *quadrature interpolates the reduced integrand,*

$$\phi(y) = f(c + hy), c = (b + a)/2\,, h = (b - a)/2\,, y \in [-1,1],$$

at the set of (*n*+1) CC *quadrature knots, spanned by the extremal abscissas of the Chebyshev polynomial of the first kind and degree n,*

$$y_j^n = \cos j\pi/n; j = 1,2,\cdots,n$$

by the truncated Chebyshev series expansion

$$L_n^{\varphi} = \sum_{k=0}^{n}{}'' \, b_k^n T_k(y)$$

where the superscript " shows that the first and the last terms of the sum are halved.

The derivation of the CC quadrature sum is done substituting $f(x)$ in I by its approximation and solving analytically the resulting integral. A weight-function-$g(x)$-dependent quadrature sum is obtained.

The most ***computer intensive task*** *within the CC quadrature comes from the* ***computation of the coefficients*** b_k^n . *The straightforward solution expresses the column vector* $B = ({}^n/_2)(b_0^n, \cdots, b_n^n)^\top$ *by the matrix product* $B = T\Phi$ where T is a symmetric matrix which collects together the values of the Chebyshev polynomials at the CC knots,

$$T_{kj} = T_{jk} = T_k\left(y_j^n\right), \qquad k, j = 0,1, \cdots, n$$

while Φ is the column vector of the computed integrand values,

$$\Phi = \left(\phi(y_0^n), \cdots, \phi(y_n^n)\right)^\top$$

Since each line of the matrix T runs over the set $\{y_j^n\}$, the occurrence of vanishing and ∓ 1 elements enables a T *matrix irreducible block decomposition which results in a significantly reduced computational cost of the numerical evaluation of the coefficient set* B. Its derivation heavily depends on the properties of the set of the CC quadrature knots.

Following ***the original convention*** of labelling the CC quadrature knots in the ***increasing order of the arguments of the cosine function***, the elements in the knot set monotonically decrease with the label j from y_0^n=1 to $y_j^n = -1$ and are symmetrically distributed around the origin of [-1, 1]: $y_{n-j}^n = y_j^n$.

The CC quadrature knots are characterized by ***inheritance***, $y_{2j}^{2n} = y_j^n$.

Moreover, under even n, the middle knot vanishes, $y_{n/2}^n = 0$.

These properties show that, for $n = 2^m$, which is assumed henceforth, there are $n = 2^{m-1} - 1$ ***non-vanishing knots*** inside the open interval (0, 1) and these ***can be ordered within*** a ***knot complete binary tree*** of root $y_1^{(2)} = y_{2^{m-2}}^{(m)} = \sqrt{2}/2$ and depth *m*-2. ***In simplified notation, here and in what follows,*** **(*m*)** ***stays for*** $\mathbf{2^m}$. At the *l*-th depth level inside the tree there are 2^l knots,

$$y_{2j-1}^{l+2} = y_{2^{m-2-l}\cdot(2j-1)}^{(m)} = \cos[(2j-1)\pi/2^{l+2}],\ j = 1, 2, \cdots, 2^l,\ l = 0, 1, 2, \cdots, m-2\,.$$

The ***binary tree structure*** and its ***ordering key*** at given *m* will be referred to as H_m . The occurrence of special arguments in knot column vector ***allows the computation of all the knots to machine accuracy*** using the square root function and recurrence. Indeed, for the root tree we have $y_1^{(2)} = \sqrt{2}/2$.
At each depth level $l = 0, 1, 2, \cdots, m-2,$ each pair of "genetically related descendants", of arguments α and $\pi/2 - \alpha$ respectively, are obtained from the values of $\cos(2\alpha)$ (the reference father of argument 2α) and of $\sin(2\alpha)$ (its "genetically related sibling" of argument $\pi/2 - 2\alpha$). The root is its own "genetically related sibling".

Within the m-panel approach, the fundamental CC-quadrature sum over macroscopic ranges is **CC-32**, ***the knots of which over the reduced range*** [-1,1] ***are given by the extremal points of the Chebyshev polynomial of the first kind and degree*** 32.

The ***parity properties*** of the Chebyshev polynomials

$$T_k(y), k = 0,1,\cdots,n;\ \ \mathrm{y} \in [-1,1]$$

result in corresponding *parity properties* of the *Chebyshev series expansion coefficient set* b: ***odd*** – $b_{2k+1}^{(m)}, \mathrm{k} = 1,2,\ldots,2^{m-1}$ and

even – $b_{2k}^{(m)}, \mathrm{k} = 0,1,\ldots,2^{m-1}$, respectively.

The CC quadrature sums involve even-rank *b* *coefficients under an even weight function* $g(x)$ and **odd-rank *b* *coefficients under an odd weight function* $g(x)$**.

Nevertheless, if the up-norm of the missing coefficient subset ***exceeds*** that of the included coefficient subset, this entails ***precision loss*** in the computed quadrature sums.

Assume that the integrand function can be ***analytically separated*** into even and odd components ***over the original integration domain***. Then, since the computation of the Chebyshev series expansion coefficients involves exclusively integrand components matching the parity of the weight function, the Clenshaw-Curtis quadrature gets ***free from precision loss due to cancellation by subtraction***.

Overview

- General frame of the Bayesian automatic adaptive quadrature
- Three classes of integration domain lengths
- **Early Bayesian problem conditioning diagnostics**
- Conclusions

More Features of Clenshaw-Curtis Quadrature

- Assume CC-32 *Clenshaw-Curtis quadrature* ***with modified reduced abscissas***.
- *Modified reduced abscissas* (MRA) *are defined as distances from the standard reduced abscissas* (SRA) *to the nearest integration domain ends*.

Advantage: *the computation of MRA is done to machine accuracy as well, while the computation of the distances between neighboring* MRA *avoids the precision loss by subtraction.*

CC-32 with MRA *is characterized by a set of* 17 MRA:

$$0 = \eta_0 < \eta_1 < \cdots < \eta_{16}=1,\ \eta_{k+\mu} - \eta_{k+\mu-1} > \eta_k - \eta_{k-1},\ \forall \mu > 0.$$

- *Generation of the integrand profile* (IP) *over subranges*

Given $f:[a,b] \to \mathbb{R}$, with *the integration domain of half-width* $h = (b-a)/2$, computation of the integrand values for CC-32 *is done over the left and right halves* and stored in two separate vectors:

$$f_k^{\,l} = f(a + h\eta_k);\ f_k^{\,r} = f(b - h\eta_k);\ k = 0, 1, \cdots, 16$$

such that $f(a+h) = f(b-h)$.

- *Default input assumption*: $I[a,b]$ is a ***proper*** Riemann integral, i.e., *integration domain* (ID) *of finite length* and *integrand boundedness*.

Flow Chart of Early Bayesian Inference

The early Bayesian inference enables the accommodation of the standard automatic adaptive quadrature approach within the Bayesian automatic adaptive quadrature.

- Step1: ***Integrand boundedness check*** over the initial IP enables:
  - **Definition of extremal (max,min) integrand values, together with their location inside the IP**
  - **End of computation (EOC) under detection of *exceptional cases*:**
    = (computationally) constant integrand
    = odd integrand with respect to the ID centre $c = (b+a)/2$
- Step2: If(.NOT. EOC) Computation of ***Riemann sums over sublattices***
  - •• ***Two* CC-32 *sublattices* are defined respectively by:**
    - the 17 *inherited* CC-16 reduced abscissas (CC16)
    - the 16 *newly added* Fejer reduced abscissas (FJ16)
  - •• ***Pairs of Riemann sums over sublattices*:**
    $q_{CC16}[(*)]$ and $q_{FJ16}[(*)]$ where $(*)$ stays either for f or $|f|$
  - •• ***Riemann sum quadrature rule outputs for* CC-32 IP:**
    - *trapezoidal rule* quadrature sums $q_{CC32}[(*)] = \left(q_{CC16}[(*)] + q_{FJ16}[(*)]\right)/2$
    - *rough error estimates* $e_{CC32}[(*)] = \left(|q_{CC16}[(*)] - q_{FJ16}[(*)]|\right)$
- Step3: EOC under detection of ***catastrophic cancellation by subtraction***:
  $|q_{CC32}[f]| < \tau \cdot q_{CC32}[|f|]$, τ close to the machine epsilon with respect to addition.

- Step4: If (.NOT. EOC) definition of ***problem adapted*** accuracy parameters $\overline{\varepsilon_a}$, $\overline{\varepsilon_r}$
- Step5: If (.NOT. EOC) the default is the ***grey diagnostics*** (GD) [postponed decision].

  An ***ill-conditioning*** (IC) diagnostic is decided iff either

$$|q_{CC32}[f]| < 2.0 \cdot e_{CC32}[f]$$

or $q_M > t_i \cdot q_m$ *where* $t_i \sim 100$ *is an empirical threshold for the comparison of*

$$q_M = \max\{|q_{CC16}[f]|, |q_{FJ16}[f]|\};\ q_m = \min\{|q_{CC16}[f]|, |q_{FJ16}[f]|\}$$

- Step6: GD ***diagnostic may be changed by the analysis of the coefficients of the Chebyshev series expansion as follows***:
  - EOC iff ***negligible*** highest label even-rank and odd-rank CC32 coefficients
  - IC iff ***well-conditioning criteria are infringed*** for suitably chosen binary tree structure dependent subsets of CC32 coefficients:
    - = *monotonicity*
    - = *fast enough absolute magnitude decay rate*
- Step7: ***Path to subrange subdivision***:

  If (IC) then
  - *refine* IC diagnostic at conveniently chosen IP abscissas
  - *subdivide current subrange* into diagnostic-dependent finer subranges

  elseif (GD)
  - proceed along the *standard* AAQ scheme

A Few Illustrative Examples

- Remember the four pillars of the Bayesian inference:
  - theory of the Riemann integral
  - theory of the numerical integration (quadrature)
  - features of the floating point computation
  - accumulated empirical evidence
- The role of the empirical evidence in the implementation of a robust, reliable, and efficient Bayesian inference path is illustrated by a few selected numerical examples
- Three typical examples are shown:
  - *integral with inner singularity integrand* [QUADPACK, pag. 110]:

$$\int_0^1 |x^2 + 2x - 2|^{-1/2} \mathrm{d}x - \text{singularity at } \sqrt{3} - 1$$

  - *integral with inner derivative singularity integrand* [Clenshaw&Curtis, Num.Math. 1960]:

$$\int_{-1}^1 |x + 0.5|^{1/2} \mathrm{d}x - \text{derivative singularity at } -0.5$$

  - *integral with oscillatory integrand* [Adam&Nobile, IMAJNA 1991]:

$$\int_{-1}^1 e^{p(x-x_0)} \sin(\omega x) \mathrm{d}x - \text{for various } p,\ x_0,\ \text{and } \omega$$

Integral with singular integrand – typical cases

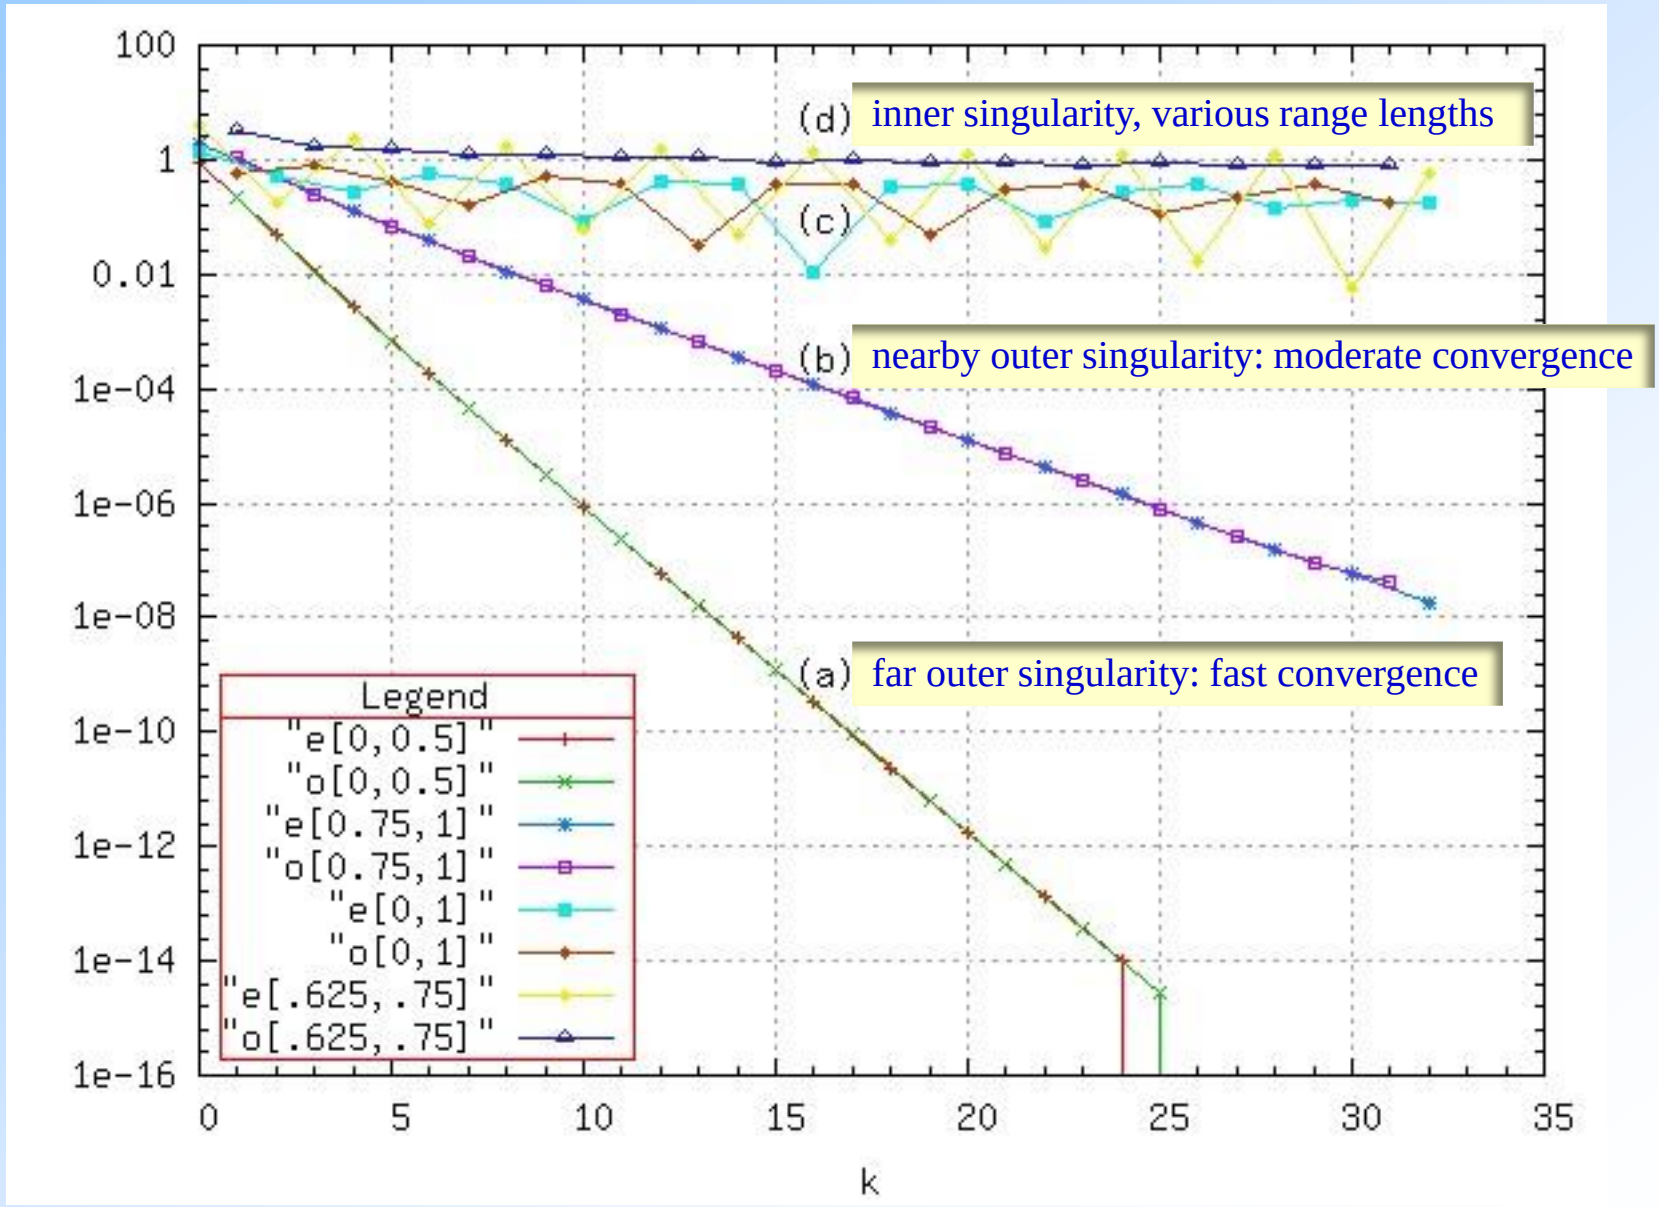

Integral with singular integrand – algebraic values of CC coefficients

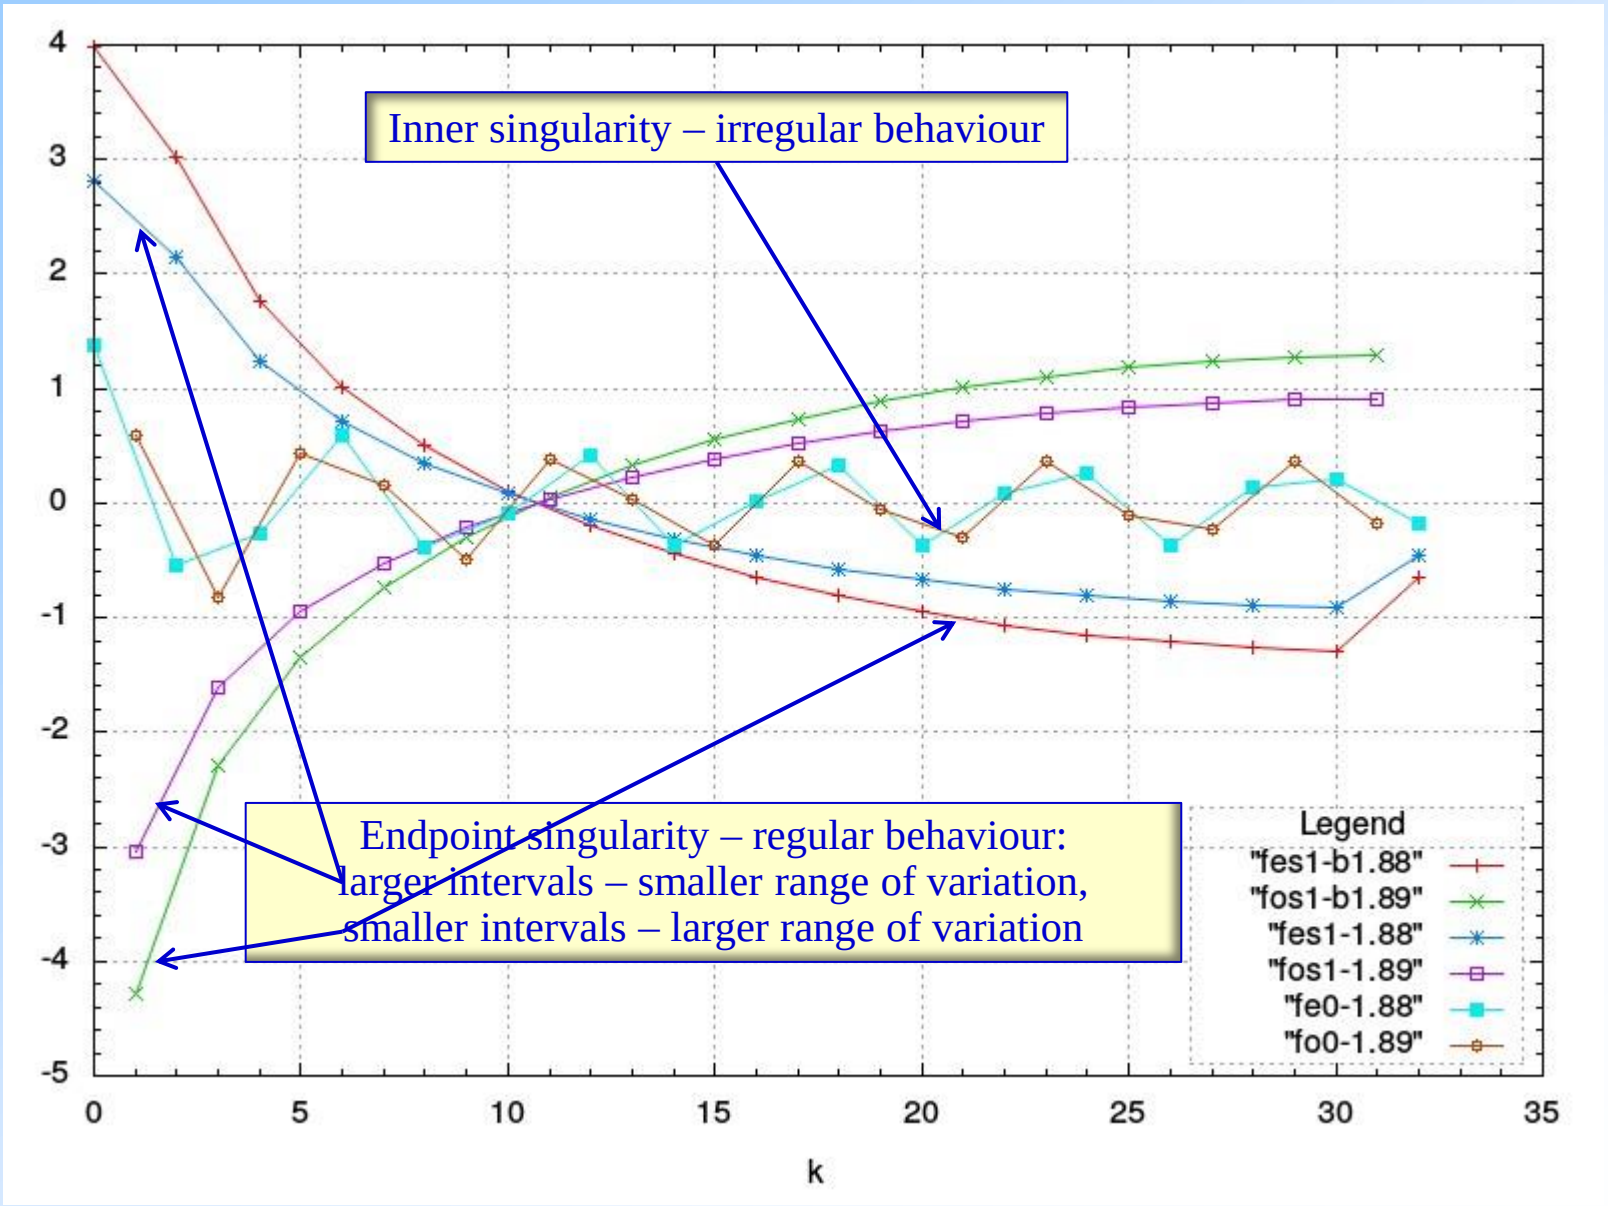

Integral with singular integrand – moduli of CC coefficients

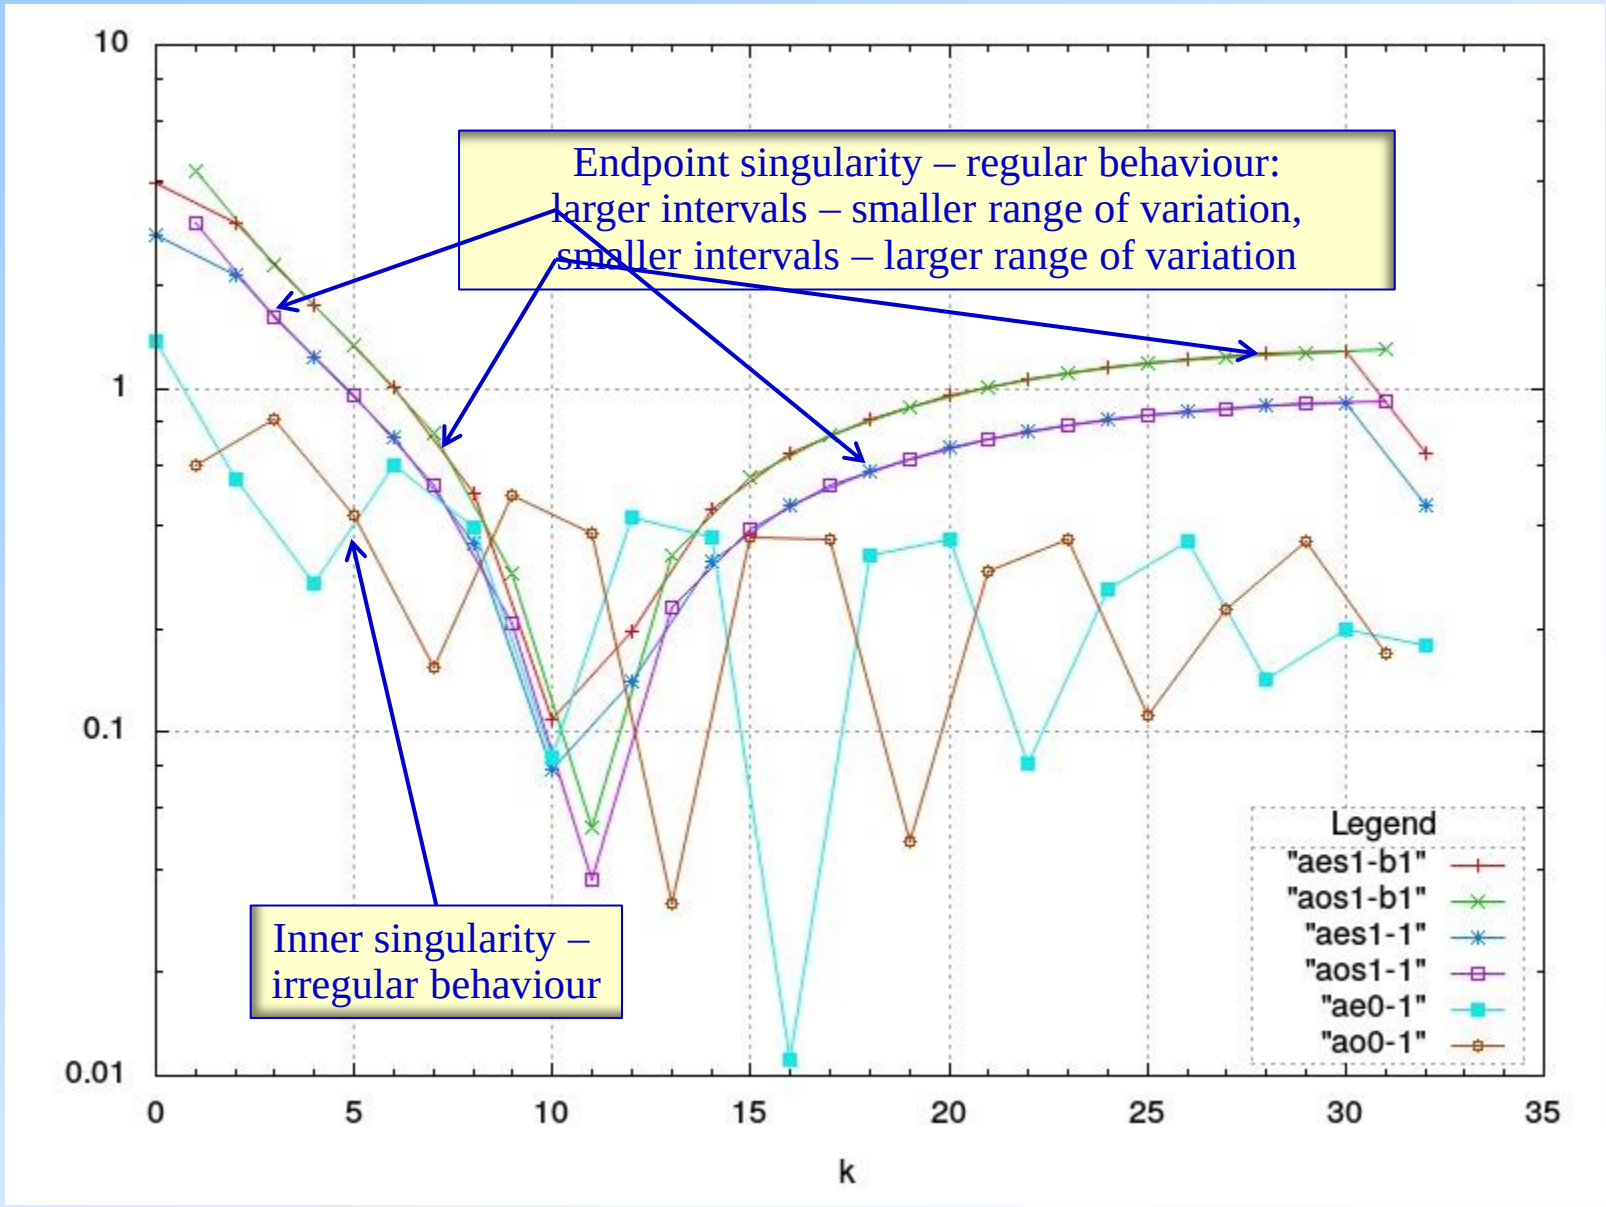

Integral with derivative singularity integrand – typical cases

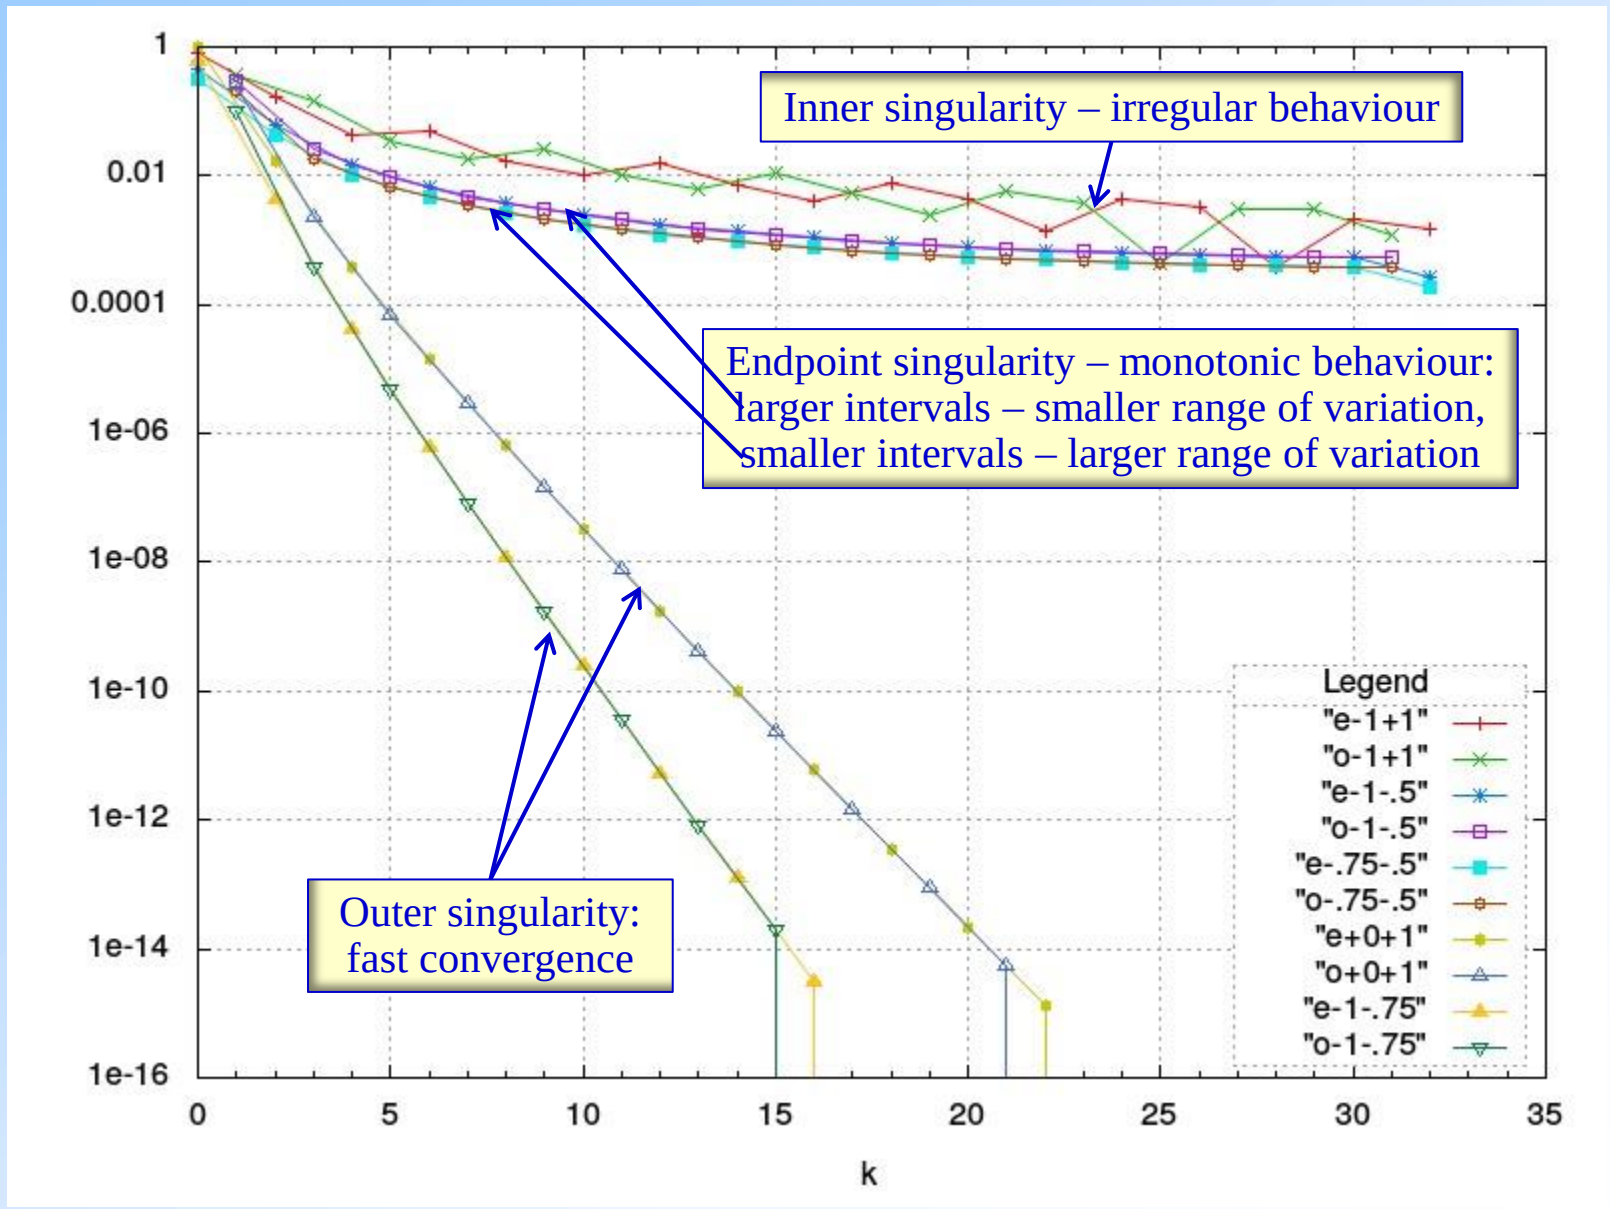

Integral with oscillatory integrand – $p = 5,\ x_0 = -1$, *various* ω

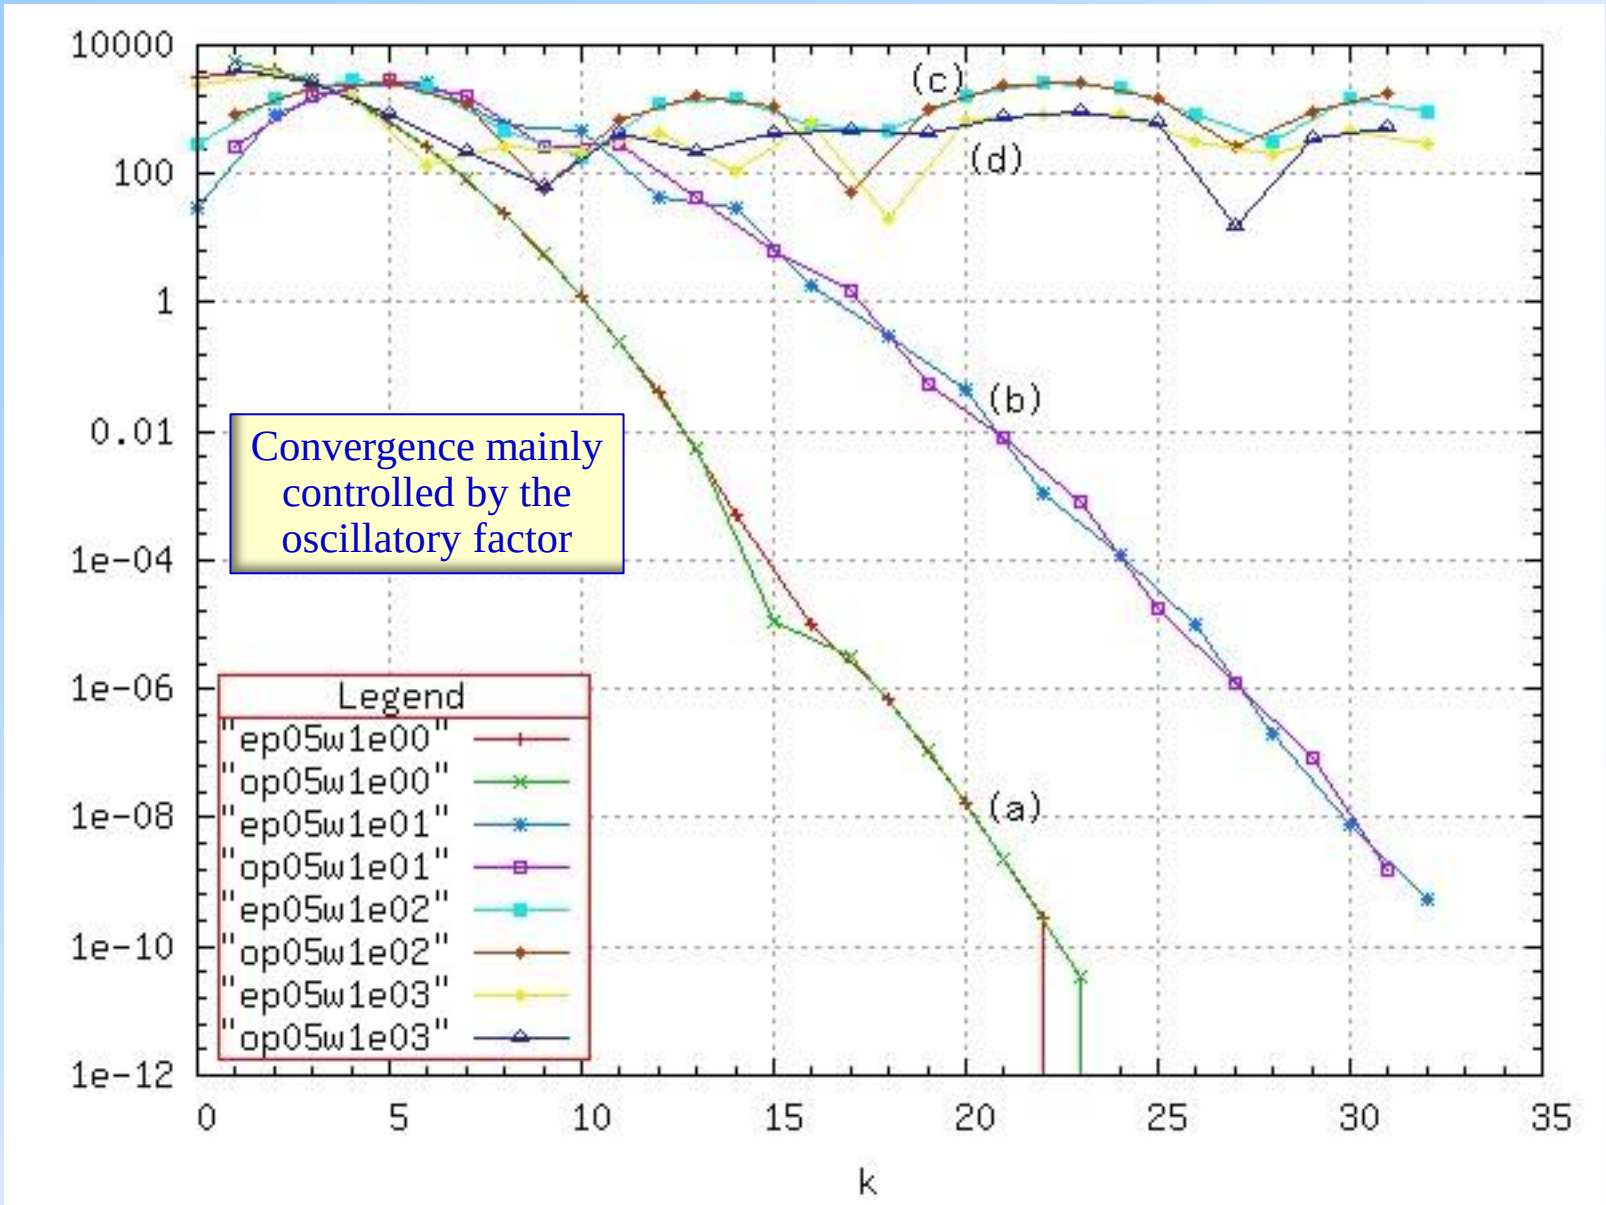

Integral with oscillatory integrand – $p = 40,\ x_0 = -1,$ *various* ω

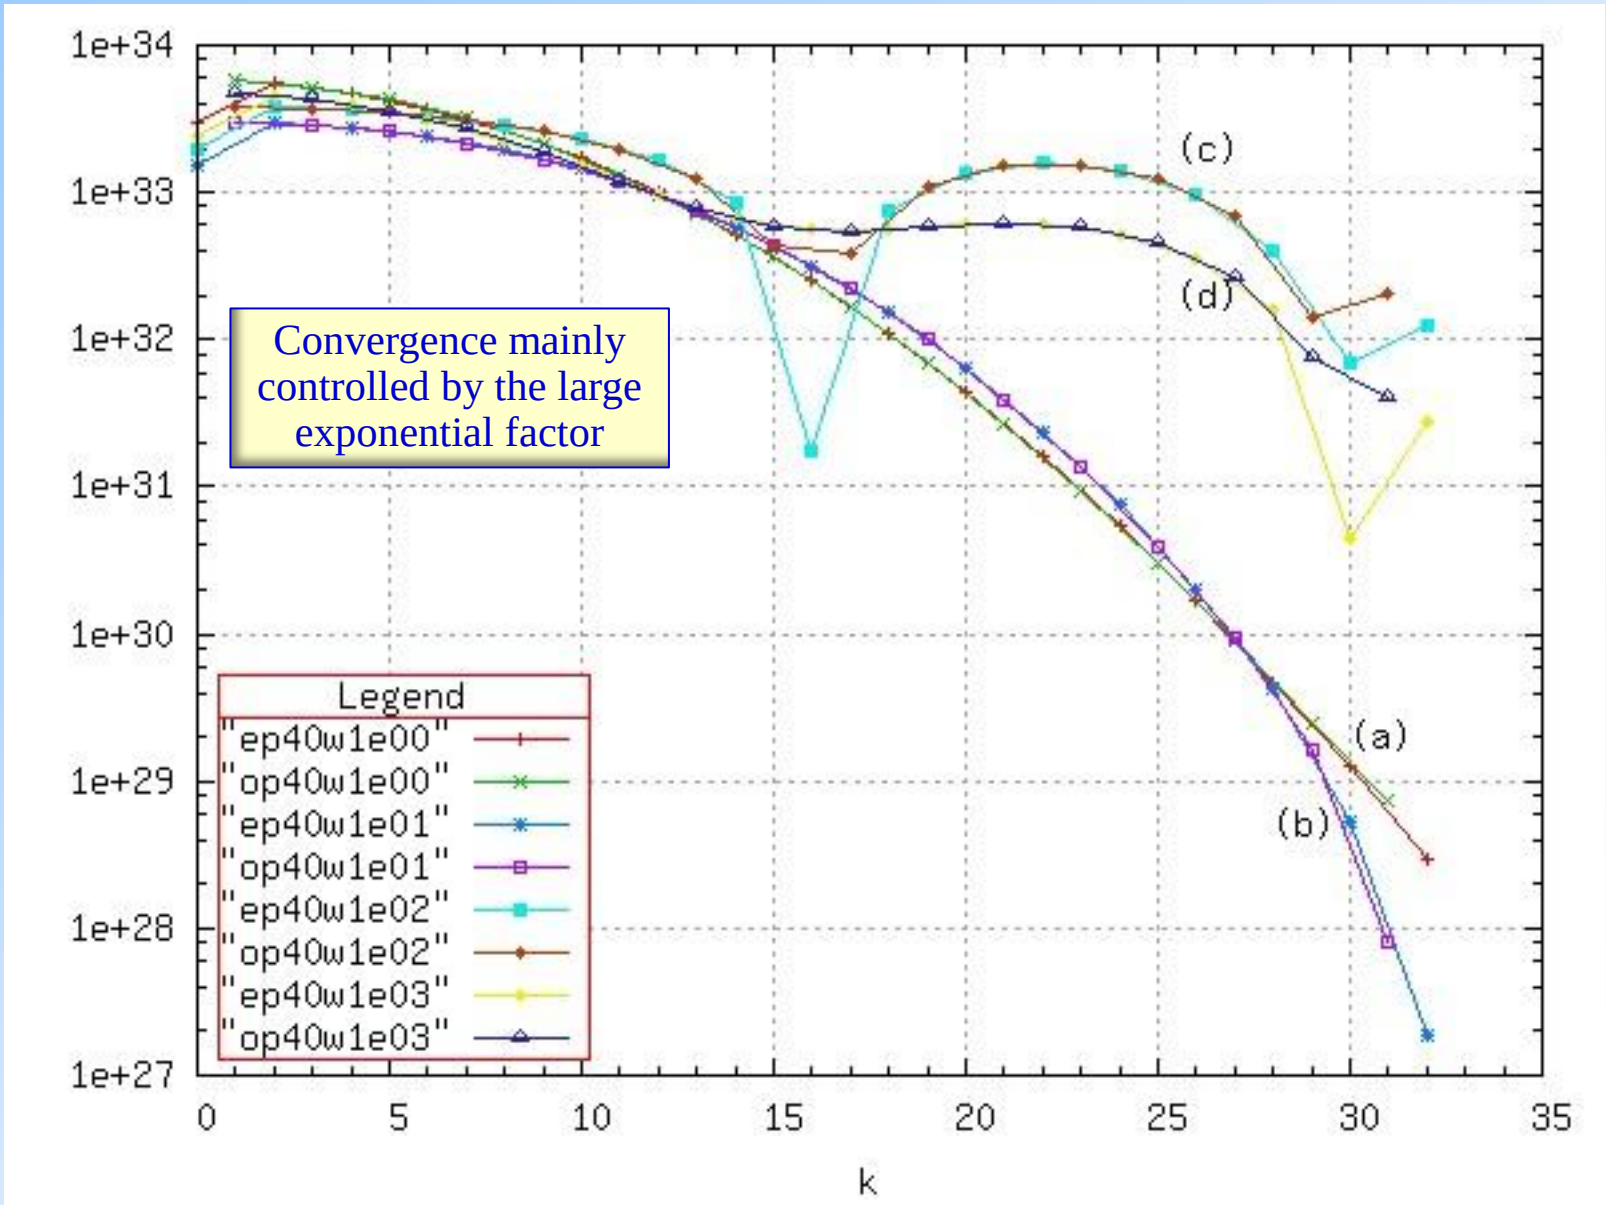

Overview

- General frame of the Bayesian automatic adaptive quadrature
- Three classes of integration domain lengths
- Early Bayesian problem conditioning diagnostics
- **Conclusions**

- *The present report discusses a Bayesian automatic adaptive quadrature (BAAQ) solution for numerical integration which is simultaneously **robust**, **reliable**, and **efficient**, yielding **maximum possible output accuracy** in numerical experiments under arbitrary behavior of the integrand function*
- *An essential ingredient of the solution is the **multiscale approach***
- *An **early decision path** to the integrand profile (IP) scrutiny enables the **identification** of **trivial** or **manifestly unsolvable** problems as well as the need to **relax** the user requested **accuracy parameters***
- *Within the **Clenshaw-Curtis quadrature** over macroscopic ranges, the scrutiny of the Chebyshev expansion coefficients enable further identification of unresolved **ill-conditioned features**.*
- *We are thus left either with a **hopefully well-conditioned integral**, for which the standard automatic adaptive quadrature can be used, or with a **manifestly ill-conditioned problem** for which an improved version of the full BAAQ machinery is activated.*

Thank you for your attention !