

9

On tables and graphs

In the *Chapters* 10 to 19, we study several families of special functions and their various series and continued fraction representations. Only a small number of these representations is also found in [AS64]. The latter are marked with the symbol $\mathfrak{A}\mathfrak{S}$ in the margin.

The collected formulas are further illustrated numerically and graphically. We now explain how to interpret and use the tables and graphs. In the sequel we consistently use z for a complex argument and x for a real argument.

9.1 Introduction

While we mention the domain of convergence with every continued fraction in the next chapters, the precise convergence behaviour is not described. Since, in practice, it is the initial convergence behaviour that matters and not the asymptotic one, we illustrate the convergence rate empirically.

This is done,

- either numerically, in tables, where we evaluate different continued fraction representations for a large range of arguments,
- or graphically, by presenting level curves of significant digits, or graphing the evolution of the approximants' accuracy.

The former is detailed in *Section* 9.2 and the latter in *Section* 9.3. All tables and graphs are labelled and preceded by an extensive caption.

9.2 Comparative tables

In the next chapters all formulas which are evaluated in one of the tables, are marked with the symbol \mathfrak{T} in the right margin. For formulas that are not marked in that way no numerical illustration of their behaviour is given. All tables are composed in the same way. The two leftmost columns contain the function argument and the function value. The function value is the correctly rounded mathematical value, verified in a variety of programming

environments. In case the function value $f(z)$ is a complex value, only its *signed modulus*

$$|f(z)|_s = \operatorname{sgn}(\Re f(z)) |f(z)| \quad (9.2.1)$$

is returned. The sign of $\Re(f(z))$ indicates whether the complex value $f(z)$ lies in the right or the left half-plane. The other columns contain the relative truncation error

$$\left| \frac{f(z) - f_n(z)}{f(z)} \right| \quad (9.2.2)$$

incurred when using a certain partial sum or continued fraction approximant $f_n(z)$ instead of the function $f(z)$ under investigation. The continued fraction approximant $f_n(z)$ can be either a classical approximant $f_n(z; 0)$ or a modified approximant $f_n(z; w_n)$.

The evaluation of the special function for the selected arguments is exactly rounded to 7 decimal digits and the truncation errors are upward rounded to 2 decimal digits. Since the modulus of the truncation error (9.2.2) is always positive, the sign is omitted here.

The approximant number n doesn't appear in the table but is mentioned in the caption. By tabulating the truncation error for different n , also the speed of convergence is illustrated.

The function arguments are selected in the intersection of the domains associated with each of the formulas evaluated in the table (with a slight exception for some series representations). The resulting set of arguments is traversed in the following way, if applicable: from the positive real axis over the first quadrant to the positive imaginary axis, then through the second quadrant of the complex plane to the negative real axis and so on. As a rule templates of all possible function arguments are tabulated, for increasing modulus, except when function evaluations for different arguments are related by symmetry relations. The numerical illustration of the elementary functions forms an exception: since these are thoroughly illustrated graphically, evaluations in the tables are restricted to real arguments only.

When evaluating the approximants of a limit periodic continued fraction $K_{m=1}^{\infty}(a_m/1)$ of which the partial numerators do not tend to zero, use of one or more modifications may be appropriate. In that case the evaluations without modification and with use of the different modifications are tabulated. We clearly indicate in the caption of the table which column in the table illustrates which modification.

When the upward rounded relative truncation error satisfies

$$C = \Delta \left(\left| \frac{f(z) - f_n(z)}{f(z)} \right| \right) \leq 5 \times 10^{-s}, \quad s \in \mathbb{N}, \quad (9.2.3)$$

then the approximation $f_n(z)$ guarantees s significant decimal digits compared to the exact value $f(z)$. When $C \simeq 10^k$ with $k \geq 0$, care must be taken in interpreting the quality of the approximation $f_n(z)$. For $k > 1$ we find $|f_n(z)| \simeq 10^k |f(z)|$, while for $k = 0$ we can very well have $|f_n(z)| \ll |f(z)|$. In both cases $f_n(z)$ can be way off, even missing to predict the magnitude of $f(z)$. In general

$$|f(z) - f_n(z)| \leq C|f(z)| \implies |f_n(z)| \in |f(z)| [1 - C, 1 + C].$$

All printed values in the tables are verified and therefore reliable. Where IEEE 754 arithmetic was insufficient because of overflow or underflow, multiprecision interval arithmetic or high precision computer algebra implementations were used.

EXAMPLE 9.2.1: Consider

$${}_2F_1(1/2, 1; 3/2; z) = \frac{1}{2\sqrt{z}} \operatorname{Ln} \left(\frac{1 + \sqrt{z}}{1 - \sqrt{z}} \right),$$

which has the regular C-fraction representation (15.3.7) given by

$$z {}_2F_1(1/2, 1; 3/2; z) = \mathbf{K}_{m=1}^{\infty} \left(\frac{c_m z}{1} \right), \quad z \in \mathbb{C} \setminus [1, +\infty),$$

$$c_1 = 1, \quad c_m = \frac{-(m-1)^2}{4(m-1)^2 - 1}, \quad m \geq 2.$$

The function also has the M-fraction representation (15.3.12) given by

$$\frac{1/2}{1/2 + z/2} - \frac{z}{3/2 + 3z/2} - \frac{4z}{5/2 + 5z/2} - \dots, \quad |z| < 1$$

and the so-called Nörlund fraction representation (15.3.17) given by

$$\frac{1}{1-z} + \frac{z(1-z)}{3/2 - 5/2z} + \mathbf{K}_{m=2}^{\infty} \left(\frac{m(m-1/2)z(1-z)}{(m+1/2) - (2m+1/2)z} \right), \quad \Re z < 1/2.$$

The intersection of the domain of $f(z) = {}_2F_1(1/2, 1; 3/2; z)$ with the convergence domains of the three continued fractions is the set

$$(\{z : |z| < 1\} \cap \{z : \Re z < 1/2\}) \setminus \{z : \operatorname{Arg} z = \pi\}.$$

So we can choose arguments:

- on the positive real axis in the interval $[0, 1/2)$,
- in all four quadrants as long as we remain inside the unit circle and have the real part less than $1/2$,
- and on the imaginary axis in the interval $(-i, i)$.

TABLE 9.2.1: Because of the symmetry property $f(x + ix) = f(x - ix)$, we can restrict ourselves to the upper half-plane, which we traverse in counterclockwise direction as explained. We evaluate the 20th classical approximant $f_{20}(z; 0)$ of each fraction and compare it to the function evaluation at the argument. Remember that for complex arguments only the signed modulus $|f(z)|_s$, as defined in (9.2.1), is displayed instead of the complex function value $f(z)$.

| x | ${}_2F_1(1/2, 1; 3/2; x)$ | (15.1.4) | (15.3.7) | (15.3.12) | (15.3.17) |
|-----|---------------------------|----------|----------|-----------|-----------|
| 0.1 | 1.035488e+00 | 2.5e-23 | 4.0e-32 | 1.5e-20 | 1.5e-20 |
| 0.2 | 1.076022e+00 | 5.6e-17 | 1.3e-25 | 1.5e-14 | 1.7e-13 |
| 0.3 | 1.123054e+00 | 3.0e-13 | 1.4e-21 | 4.8e-11 | 7.6e-09 |
| 0.4 | 1.178736e+00 | 1.4e-10 | 1.8e-18 | 1.4e-08 | 5.0e-05 |

| x | $ {}_2F_1(1/2, 1; 3/2; x + ix) _s$ | (15.1.4) | (15.3.7) | (15.3.12) | (15.3.17) |
|-----|------------------------------------|----------|----------|-----------|-----------|
| 0.1 | 1.033684e+00 | 3.6e-20 | 3.9e-29 | 1.5e-17 | 1.4e-17 |
| 0.2 | 1.066938e+00 | 8.0e-14 | 1.1e-22 | 1.6e-11 | 9.3e-11 |
| 0.3 | 1.097258e+00 | 4.2e-10 | 8.4e-19 | 5.0e-08 | 1.5e-06 |
| 0.4 | 1.121184e+00 | 1.8e-07 | 5.5e-16 | 1.6e-05 | 1.4e-03 |

| x | $ {}_2F_1(1/2, 1; 3/2; ix) _s$ | (15.1.4) | (15.3.7) | (15.3.12) | (15.3.17) |
|-----|--------------------------------|----------|----------|-----------|-----------|
| 0.1 | 9.985628e-01 | 2.3e-23 | 1.4e-32 | 1.6e-20 | 1.8e-21 |
| 0.3 | 9.875589e-01 | 2.4e-13 | 3.6e-23 | 5.5e-11 | 2.9e-12 |
| 0.5 | 9.678199e-01 | 1.0e-08 | 5.8e-19 | 1.5e-06 | 2.1e-08 |
| 0.7 | 9.425900e-01 | 1.1e-05 | 2.4e-16 | 1.3e-03 | 3.1e-06 |
| 0.9 | 9.147830e-01 | 2.1e-03 | 1.6e-14 | 1.9e-01 | 6.8e-05 |

| x | $ {}_2F_1(1/2, 1; 3/2; x - ix) _s$ | (15.1.4) | (15.3.7) | (15.3.12) | (15.3.17) |
|------|------------------------------------|----------|----------|-----------|-----------|
| -0.1 | 9.673650e-01 | 3.2e-20 | 5.4e-30 | 1.6e-17 | 2.8e-19 |
| -0.3 | 9.077224e-01 | 2.9e-10 | 2.8e-21 | 6.1e-08 | 2.4e-11 |
| -0.5 | 8.563213e-01 | 1.2e-05 | 1.3e-17 | 1.8e-03 | 2.3e-08 |
| -0.7 | 8.123036e-01 | 1.3e-02 | 2.1e-15 | 3.9e+00 | 9.9e-07 |

We see that, for the real argument $x = 0.3$, the 20th approximant $f_{20}(x; 0)$ of the C-fraction (15.3.7) ensures 21 significant decimal digits because

$$\left| \frac{{}_2F_1(1/2, 1; 3/2; x) - f_{20}(x; 0)}{{}_2F_1(1/2, 1; 3/2; x)} \right| \leq 1.4 \times 10^{-21}, \quad x = 0.3.$$

Clearly, here the C-fraction delivers the better approximant. The evaluation of (15.3.7) can further be improved with the use of the modifications (15.3.5), $w(z) = \frac{1}{2}(\sqrt{1-z} - 1)$, and (15.3.6),

$$w_n^{(1)}(z) = w(z) + \frac{c_{n+1}z + z/4}{1 + 2w(z)}.$$

TABLE 9.2.2: The approximant $f_{20}(x)$ of (15.3.7) is first evaluated without modification and subsequently with the mentioned modifications. Note that the first truncation error column equals the first truncation error column of Table 9.2.1 for real arguments x , since both concern the classical approximant $f_{20}(x; 0)$ of the C-fraction.

| x | ${}_2F_1(1/2, 1; 3/2; x)$ | (15.3.7) | (15.3.7) | (15.3.7) |
|-----|---------------------------|----------|----------|----------|
| 0.1 | 1.035488e+00 | 4.0e-32 | 2.6e-35 | 6.5e-38 |
| 0.2 | 1.076022e+00 | 1.3e-25 | 8.8e-29 | 4.8e-31 |
| 0.3 | 1.123054e+00 | 1.4e-21 | 1.1e-24 | 9.6e-27 |
| 0.4 | 1.178736e+00 | 1.8e-18 | 1.4e-21 | 1.9e-23 |

With use of the modifications, the truncation error in $x = 0.3$ decreases to

$$\left| \frac{{}_2F_1(1/2, 1; 3/2; x) - f_{20}(x; w(x))}{{}_2F_1(1/2, 1; 3/2; x)} \right| \leq 1.1 \times 10^{-24}, \quad x = 0.3,$$

$$\left| \frac{{}_2F_1(1/2, 1; 3/2; x) - f_{20}(x; w_{20}^{(1)}(x))}{{}_2F_1(1/2, 1; 3/2; x)} \right| \leq 9.6 \times 10^{-27}, \quad x = 0.3,$$

respectively ensuring 21, 24 and 26 significant decimal digits.

EXAMPLE 9.2.2: The ratio of parabolic cylinder functions $U(1, x)/U(0, x)$ has the C-fraction representation (16.5.7),

$$\frac{U(1, x)}{U(0, x)} = \frac{1}{x} + \mathbf{K}_{m=2}^{\infty} \left(\frac{m - \frac{1}{2}}{x} \right), \quad x > 0.$$

Since the partial numerators tend to infinity, use of the modification

$$\tilde{w}_n(x) = \frac{-x + \sqrt{4(n + 1/2) + x^2}}{2}$$

is recommended when evaluating the approximants.

TABLE 9.2.3: We tabulate the relative error of the 5th approximants $f_5(x)$ which are first evaluated without modification and then with the modification $w_5(x)$.

| x | $U(a, x)/U(a - 1, x)$ | (16.5.7) | (16.5.7) |
|--------|-----------------------|----------|----------|
| 0.25 | 8.329323e-01 | 1.4e+00 | 2.0e-02 |
| 0.75 | 6.485192e-01 | 1.8e-01 | 4.4e-03 |
| 1.25 | 5.211635e-01 | 4.0e-02 | 1.0e-03 |
| 5.25 | 1.813514e-01 | 7.3e-06 | 1.2e-07 |
| 20.25 | 4.920381e-02 | 2.6e-11 | 5.9e-14 |
| 50.25 | 1.988869e-02 | 3.1e-15 | 1.2e-18 |
| 100.25 | 9.973574e-03 | 3.2e-18 | 3.1e-22 |

The tables in the handbook are not discussed, only presented, because the conclusions are obvious most of the times. The speed of convergence can be observed. The variation in the magnitude of the truncation error throughout the complex plane is clear. Specific observations, such as extremely slow convergence, are confirmed in the existing literature.

Other counter-intuitive behaviour can be understood by taking a closer look at the formulas involved. For instance, while T-fractions correspond to series developments both for small and large z (around 0 and ∞), this behaviour is not confirmed (at first sight) in the evaluation of (17.1.48). Take a look at the approximation of $J_\nu(x)$ in the *Tables* 17.1.1 and 17.1.2. For real-valued arguments the Bessel function is real-valued. But here the T-fraction (exceptionally) introduces an imaginary part in the approximation of $J_{\nu+1}(x)/J_\nu(x)$. Of course, this disturbs the quality of the approximation on the real axis. Because of (17.2.2) a similar observation can be made for the Bessel function $I_\nu(ix)$ evaluated on the imaginary axis. In other parts of the complex plane the behaviour is as expected.

9.3 Reliable graphs

For graphical illustrations of the specific behaviour of special functions in subsets of their domain, we refer to [AS64] and [SO87]. Also several websites are devoted to the visualisation of special functions, both for real and complex variables. Our interest is in the approximation power of series representations and continued fraction representations and hence in the visualisation of the truncation error incurred when using these approxima-

tions. We therefore show *level curves* of s such that

$$\left| \frac{f(z) - f_n(z)}{f(z)} \right| \leq 5 \times 10^{-s}, \quad z \in \mathbb{C}, \quad s \in \mathbb{N} \quad (9.3.1)$$

or

$$\left| \frac{f(x) - f_n(x)}{f(x)} \right| = 5 \times 10^{-s}, \quad x \in \mathbb{R}, \quad s \in \mathbb{R}_0^+ \quad (9.3.2)$$

where $f_n(z)$ is the n^{th} approximant of a series or continued fraction representation for the function $f(z)$. Continued fraction representations depicted in graphs are marked throughout the chapters with \boxtimes in the right margin.

In all our plots the grid lines are drawn one unit apart, both in the horizontal and the vertical direction (aspect ratio 1 for each unit square). If the x - or y -axis belong to the plot, they are shown as a solid line. For instance, in the area $[1, 19] \times [0, 9]$ only the x -axis is shown, while in the area $[-2, 2] \times [-2, 2]$ both axes appear. Since the axes do not always appear in the picture, we have preferred not to label them. The caption provides sufficient information on the plotted area.

So all our graphs show implicit relations $R_{s,t}(x, y) = 0$ involving at most two real unknowns x and y and some real parameters s and/or t . Given that this kind of graphing problem has been discussed for centuries, it is unsurprising that there is an abundance of (partial) solutions to it. It is, however, surprising that until recently [Tup91] there was no computer method capable of reliably solving this problem.

The algorithm implemented in **GrafEq** (pronounced “graphic”) correctly graphs mathematical formulas involving the basic operations, inequalities and known elementary functions [Tup04]. When applied to a difficult formula that is beyond its capabilities, the algorithm clearly marks the pixels that it cannot decide to belong to the graph, as uncertain. At no point does the algorithm use any approximations that may cause it to produce an incorrect graph. We summarise the internal workings of **GrafEq** below. In order for **GrafEq** to be useful in the context of the continued fraction handbook, two extensions were developed [BCJ⁺05]:

- since expressions in a complex variable z may get quite complicated when manually converted to a relation in $x = \Re z$ and $y = \Im z$, it is necessary to add the direct handling of complex variables to **GrafEq**’s interface;
- since **GrafEq** only has implementations of the elementary functions and none of the special functions, we need to be able to dynamically extend the list of functions known by **GrafEq** by providing their implementation.

Internal working of GrafEq. Any formula $R(x, y)$, when evaluated with specific real numbers x and y , is always either false (F) or true (T). Given a mathematical formula $R(x, y)$ and a rectangular region $[L, R] \times [B, T]$ of the Cartesian plane \mathbb{R}^2 , GrafEq produces an illustration that consists of a $W \times H$ rectangular array of pixels. Each pixel herein represents a closed rectangular region of the plane. Since no algorithm can produce correct black and white graphs, black meaning that there is at least one solution of $R(x, y)$ within the pixel and white meaning that there are no solutions within the pixel, we allow to color some pixels “uncertain”, meaning that there may or may not be solutions of $R(x, y)$ within the pixel.

Even if the bounds L, R, B and T of the graphing area are given as machine numbers, the bounds of individual pixels may not be representable exactly. Therefore GrafEq uses inner and outer bounds of the rectangular region corresponding to the pixel. The inner bounds are used to show the guaranteed existence of solutions and the outer bounds to show the absence of solutions. Further GrafEq makes use of interval arithmetic with boolean values to represent and process the result of formula evaluations. Three boolean intervals are possible, $\langle F, F \rangle, \langle F, T \rangle, \langle T, T \rangle$ with $F < T$. The boolean intervals provide:

- domain tracking, by keeping track of whether or not a quantity such as \sqrt{x} with $x < 0$, is well-defined,
- continuity tracking, by providing information on whether a quantity is continuous or not within the given bounds,
- branch tracking, by tracing to which branch each piece belongs when breaking a discontinuous evaluation apart into pieces.

Plotting special functions. Additional function implementations must also return interval evaluations and support the domain and continuity boolean intervals required by GrafEq’s internal engine. To guarantee the reliability of the function evaluations, the results in [CVW06] on the implementation of special functions are used.

EXAMPLE 9.3.1: We show a simple illustration of (9.3.1) for a T-fraction approximant of the exponential function, namely

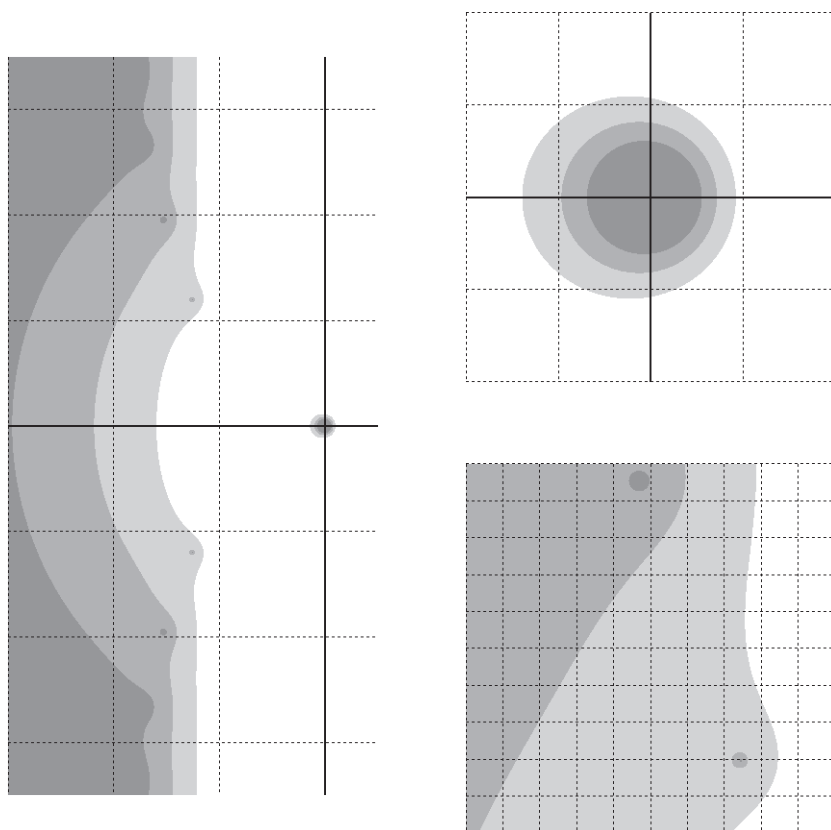
$$f(z) = \exp(z) - 1,$$

$$f_n(z) = \frac{z}{1-z} + \mathbf{K}_{m=2}^n \left(\frac{(m-1)z}{m-z} \right), \quad n = 8,$$

$$\frac{|f(z) - f_n(z)|}{|f(z)|} \leq 5 \times 10^{-s}, \quad s = 6, 7, 8.$$

T-fractions have the property that, besides being useful for relatively small z , they approximate well for small $1/z$ at the same time.

FIGURE 9.3.1: We consider the region $-30 \leq \Re z \leq 5, |\Im z| \leq 35$ (exceptionally, in this figure the grid-lines are 10 units apart) and zoom in on the regions $|\Re z| \leq 2, |\Im z| \leq 2$ and $-20 \leq \Re z \leq -10, 10 \leq \Im z \leq 20$ (in these figures the grid-lines are 1 unit apart, as usual). The regions corresponding to $s = 6, 7, 8$ are respectively coloured light-grey, medium-grey and dark-grey, respectively. Axes and grid-lines are coloured black.



Note the very small isolated regions in the left half-plane of the larger drawing, which are impossible to locate without a reliable graphing method. The resolution of the small insets is 192×192 pixels while that of the larger figure is 192×384 pixels. With the same number of 73728 plot-points for the latter, the computer algebra system Maple (version 10) is unable to produce the correct graph for $s = 8$.

EXAMPLE 9.3.2: The following is an illustration of (9.3.2) taken from *Chapter 14*. In contrast to the above example where $s \in \mathbb{N}$, here s can take on all positive real values. Let the exponential integrals $E_n(z)$ be defined in $\Re z > 0$ by

$$E_n(z) := \int_1^\infty \frac{e^{-zt}}{t^n} dt, \quad n \in \mathbb{N}.$$


Analytic continuation of $E_n(z)$ to the cut plane $|\arg z| < \pi$ extends the definition and yields a single-valued function. Except on the negative real axis, they can be represented by the series representation (14.1.11), the asymptotic series expansion (14.1.13), the S-fraction (14.1.16),

$$\begin{aligned} e^z E_n(z) &= \frac{1}{z} + \frac{n}{1+z} + \frac{1}{z} + \frac{n+1}{1+z} + \frac{2}{z} + \frac{n+2}{1+z} + \dots, \quad n \in \mathbb{N}, \\ &= \frac{1/z}{1} + \mathbf{K}_{m=2}^\infty \left(\frac{a_m/z}{1} \right), \quad a_{2k} = n+k-1, \quad a_{2k+1} = k, \\ & \quad \quad \quad |\arg z| < \pi \end{aligned}$$

and the real J-fraction representation (14.1.23),

$$e^z E_n(z) = \frac{1}{n+z} + \mathbf{K}_{m=2}^\infty \left(\frac{(2-m)(n+m-2)}{n+2m-2+z} \right), \quad |\arg z| < \pi.$$

FIGURE 9.3.2: On the vertical axis we display the value s in (9.3.2) in the range $0 \leq s \leq 12$, where $f_n(z)$ respectively equals the 5th partial sum of the series development (14.1.11) (lightest), the 5th partial sum of the asymptotic series (14.1.13) (second lightest), the 5th approximant of the S-fraction (second darkest) and the 5th approximant of the real J-fraction (darkest) of $f(z) = E_3(z)$, all for real $z = x$ with $1 \leq x \leq 25$.



It is easy to see how the continued fraction representations outperform the series developments. From *Table 14.1.1* one also sees that in the interval $[1, 25]$ the 5th modified approximant of the S-fraction guarantees on average one more significant decimal digit than the classical 5th approximant.

Besides the graphs shown in *Example 9.3.1*, for which we needed reliable graphing software, it is also interesting to take a look at the next figure, in which we show the number of significant digits delivered by successive approximants $f_n(x)$ of the Thiele interpolating continued fraction (11.3.9)

$$\tan(z) = \frac{z}{1} + \frac{-4\pi^{-2}z^2}{1} + \mathbf{K}_{m=1}^{\infty} \left(\frac{m^4 - 4\pi^{-2}m^2z^2}{2m+1} \right),$$

$$z \in \mathbb{C} \setminus \{\pi/2 + k\pi : k \in \mathbb{Z}\}$$

for real arguments x , where from light to dark $n = 5, 6$ and 7 . At the interpolation points $\pm m\pi/2$, $m \in \mathbb{N}_0$, the accuracy and hence the number of significant digits is infinite, but in the neighbourhood of the interpolation points the peaks in the graph are so steep that the traditional device independent graphing tools miss each of them. Fortunately **GrafEq** does not!

