New Algorithms for Efficient Taylor Model Operations including Arbitrary Precision

Martin Berz and Kyoko Makino

Abstract

Within the framework of Taylor Model implementations, the multiplication of Taylor models is the most time intensive operation, and thus benefits most from performance improvements. We note that in the mulplication, every coefficient is utilized repeatedly in various products, and so it is possible to do limited pre-processing of the coefficients without performance penalty due to the preprocessing, if it leads to significant performance gains in the coefficient-by-coefficient operations.

As a first step, we observe that it is advantageous to represent the "right" Taylor model as a difference between one Taylor model with purely positive coefficients, and another one with purely negative ones. In the TM multiplication, individual coefficient multiplications are carried out separately for positive and negative results. This greatly simplifies the floating point error estimates, which now depend only on the floating point sums of the positive and negative result contributions, so that each multiplication and addition to accumulated coefficient contributiosn can be carried out without roundoff error consideration.

Next we present a method to naturally extend Taylor models to allow arbitrary precision based on the well known two-sum and two-product algorithms which allow the sum and products of two floating point numbers to be represented exactly in terms of unevaluated sums of floating point numbers of staggered magnitude. Within the Taylor expansion framework, another dummy variable is introduced which accounts for different "powers" of machine accuracy. An interesting consequence is that this approach automatically leads to the fact that only those coefficients that need it are represented in higher precision, substantially increasing the efficiency.

However, significant additional gains can be realized if one makes a pre-processing sweep in which all coefficients that are sufficiently above the cutoff threshold are split into two as in the two-product algorithm. This entails that whenever such coefficients are multiplied, the result is exact, and no consideration of the floating point errors is necessary.

Overall, these algorithms lead to efficient implementations of rigorous Taylor model operations. In the conventional double precision case, overhead for rigorous error accounting is nearly negligible for the high-dimension multivariate case. For the case of higher precion up to a few multiples of conventional double accuracy, further significant gains compared to evaluating each product as a conventional two-product can be realized.