Ensuring Correctly Rounded Mathematical Functions (M1 internship)

Location. INRIA/LORIA, Nancy, `www.loria.fr`.

Advisor. Paul Zimmermann, research director, team Caramba, `Paul.Zimmermann@inria.fr`.

Head of Inria Nancy - Grand Est. Bruno Lévy, `Bruno.Levy@inria.fr`.

Context. The IEEE-754 standard defines several binary floating-point formats: single, double, and quadruple precision. For basic arithmetic operations (addition for example) it requires *correct rounding*, i.e., the result should be the closest one to the exact value $a + b$ in the target format. However, for mathematical functions like $\sin, \exp, \log$, the IEEE-754 standard does not require correct rounding. As a consequence, current mathematical libraries (GNU libc, Intel Math Library, AMD LibM, Redhat Newlib, OpenLibm, Musl, ...) do not yield the best possible results [1].

Scientific Objectives. The objective of this internship is to design correctly-rounded algorithms (for single, double, and quadruple precision mathematical functions) and to efficiently implement them in the C language. Each function should be correctly-rounded for all four IEEE 754 rounding modes: to nearest, towards zero, towards $-\infty$, towards $+\infty$. The ultimate goal is to obtain code that is as efficient (or even more) than current mathematical libraries, and to submit that code for inclusion into these libraries. Some first implementations have been designed within the CORE-MATH project (`https://homepages.loria.fr/PZimmermann/CORE-MATH/`).

Required skills. This internship requires skills in arithmetic, and a very good knowledge of the C language.

References

[1] Zimmermann, P. Accuracy of mathematical functions in single precision. `https://members.loria.fr/PZimmermann/papers/accuracy.pdf`, 2021. Version of September 7, 22 pages.