Internship offer:

Probabilistic rounding error analyses of matrix factorizations

Contact: Theo Mary (theo.mary@lip6.fr), Stef Graillat (stef.graillat@lip6.fr), Fabienne Jézéquel (fabienne.jezequel@lip6.fr)

Context of the internship:

Rounding error bounds of linear algebra computations involving matrices or vectors of size n tend to grow at least linearly with n. Historically, the role of n in the error bounds has been neglected by numerical analysts. In traditional error analysis, terms only depending on the dimensions of the data are called "constants", which reflects the fact that historically n was in practice not large enough for its role in the growth of the error to be deemed significant. However, this view is no longer valid with the extreme size of matrices that are nowadays involved with exascale computing, and with the emergence of low precision arithmetics, which possess a large unit roundoff u. For example, with half precision (16-bit) arithmetic, error bounds of order nu cannot guarantee even a single correct digit for matrices of order a few thousands, and so become quite unsatisfying.

Fortunately, these are worst-case bounds and, in practice, the error often exhibits a much weaker dependence on n. A famous conjecture by Wilkinson from 1961 states that n can be replaced by its square root $\sqrt{n}$ due to statistical effects in the rounding errors: assuming they behave somewhat randomly, positive and negative rounding errors should cancel eath other to some extent. This rule of thumb, which remained formally unproven for 60 years, was proven, under a well defined model of rounding errors, by Higham and Mary in 2019 [1]. Probabilistic rounding error analysis has since then been a very active research field, and several important open problems remain.

Main objective:

The goal of this internship is to tackle one of these open problems: for matrix factorizations (e.g., Cholesky or LU), in practice the average case error is observed to be *independent* of the matrix size n, that is, even the probabilistic $\sqrt{n}$ bound is pessimistic in this case. This surprising behavior makes it possible to solve accurately extreme scale linear systems in very low precisions. The goal will be to provide a formal explanation for this behavior, by developing a probabilistic rounding error analysis specially tailored to matrix factorizations. We will begin by the Cholesky case, which is simpler than LU because of the absence of growth of the matrix coefficients.

Depending on the results of the avenues raised by this internship, a PhD thesis in the same field and in the same team could be proposed.

Other details:

- Internship location: LIP6 laboratory, Sorbonne University, in the city center of Paris.
- Duration: 5/6 months.
- Remuneration: to be specified according to the training.

References:

[1] Higham and Mary, A New Approach to Probabilistic Rounding Error Analysis, *SIAM J. Sci. Comput.* (2019). https://hal.archives-ouvertes.fr/hal-02311269