
Approximate solutions to (generalized) matrix chains

Paolo Bientinesi^{*1}, Francisco Lopez¹, and Lars Karlsson¹

¹Umeå University – Sweden

Abstract

Thanks to convenient high-level interfaces, many high-level languages and libraries (e.g., Matlab/Octave, Julia, R, Armadillo) make it possible for users to write programs that involve not only scalars, but also vectors and matrices. The translation of such programs to a sequence of computational kernels (such as those in the BLAS and LAPACK libraries) is a problem, known as the "Linear Algebra Mapping Problem" (LAMP), that all those high-level languages must solve. If the input expression is a matrix chain (i.e., a product of three or more matrices), then the problem is equivalent to finding a parenthesization that minimizes a given cost function (e.g., number of floating point operations). For "short" chains, this can be done quickly, even at runtime. Surprisingly, most languages fail to solve even this simple setup. For more complex inputs, the problem becomes NP-complete and it is not feasible to look for an optimal solution. In this talk, I will present theoretical and practical results for the approximate solution to matrix chains and generalized matrix chains.

^{*}Speaker