Optimizing Linear Algebraic Operations for Improved Data-Locality

Dániel Berényi[1], András Leitereg[2], Gábor Lehel[2]

[1]Wigner Research Centre for Physics of the Hungarian Academy of Sciences, Budapest, Hungary

[2]Department of Informatics, Eötvös Loránd University, Budapest, Hungary

In this talk we present a selection of higher-order functions that were chosen to compactly represent dense linear algebraic expressions (like matrix multiplication or tensor contraction) and are suitable for optimization. First, we show that these primitives have desirable algebraic properties and are closed under certain fusion laws. Second, we present subdivision rules for splitting the primitives, and exchange rules for the cases where a higher-order function is passed as a function argument to another. Such rules make it possible to start from naive forms of linear algebraic expressions, then generate permutations and different block rearrangements which have much better data locality and therefore performance compared to the naive version.