

Modeling the effects of data locality

András Leitereg¹, Dániel Berényi²

¹ Department of Programming Languages and Compilers, Faculty of Informatics, Eötvös Loránd University, Budapest, Hungary

² Wigner GPU Laboratory, Wigner Research Centre for Physics, Hungarian Academy of Sciences, Budapest, Hungary

Linear algebraic operations are vital in various areas of scientific computing, thus optimizing them is crucial. One aspect of these optimizations is focusing on data locality, that is, utilizing the hierarchical memory layout of the given computing hardware. In the case of linear algebra this is related to the order of operations, indices and access patterns. On different hardware different optimizations give the best performance and these details are usually captured in cost models. We are exploring an alternative way of learning this information with a recursive neural network (Tree Net) trained on benchmarking sample expressions.