

High-Performance Kernels for Exact Linear Algebra

Jeremy Johnson¹, Tze Meng Low², Matthew Lambert³, Peter Oostema², B. D. Saunders³

¹ Drexel University, Philadelphia PA, USA, jjohnson@cs.drexel.edu

² Carnegie Mellon University, Pittsburgh PA, USA, lowt,poostema@andrew.cmu.edu

³ University of Delaware, Newark DE, USA, lambert,saunders@udel.edu

High-performance linear algebra libraries are typically built on top of fast matrix-matrix multiplication kernels. Significant effort, by the numerical linear algebra community, has been devoted to the implementation and optimization of these kernels on a wide variety of computer architectures [1, 3, 4].

The computer algebra community has taken advantage of this work in [5], avoiding duplication of effort by calling numeric kernels with block size chosen so that overflow is guaranteed not to occur and exact results are provided. Other efforts have been devoted to specialized coefficients domains, such as GF(2) [6, 7] and GF(3) [8], where domain specific optimizations, such as bit packing, bit slicing and table lookup have been used together with domain specific algorithms such as four Russians. These efforts have tended to focus on the domain specific optimizations and not necessarily memory hierarchy and architecture specific optimizations that have been the focus of the numeric linear algebra community.

The BLIS (BLAS-like Library Instantiation Software) framework [2] is an effort to provide easy access to the optimizations used in fast matrix kernels. By rewriting a few key kernels, the user can take advantage of the framework for efficient use of the memory hierarchy and other architectural features. In this presentation we report on an investigation of the use of BLIS to develop matrix-matrix multiplication kernels over various exact coefficient domains.

References

- [1] R. Clint Whaley, Antoine Petitet and Jack J. Dongarra, *Automated Empirical Optimization of Software and the ATLAS Project*, Parallel Computing, 27, pp. 3–35 (2001).
- [2] Field G. Van Zee and Robert A. van de Geijn, *BLIS: A Framework for Rapidly Instantiating BLAS Functionality*, ACM Transactions on Mathematical Software, 41(3), pp. 1–33 (2015).
- [3] Kazushige Goto and Robert A. van de Geijn, *High-performance implementation of the level-3 BLAS*, ACM Transactions on Mathematical Software, 35(1), pp. 1–14 (2008).
- [4] Kazushige Goto and Robert A. van de Geijn, *Anatomy of High-Performance Matrix Multiplication*, ACM Transactions on Mathematical Software. 34(3), pp. 1–25 (2008).
- [5] Jean-Guillaume Dumas and Pascal Giorgi and Clément Pernet, *Dense Linear Algebra over Word-Size Prime Fields: the FFLAS and FFPACK Packages*, ACM Transactions on Mathematical Software, 35(3), pp. 1–42 (2008).

- [6] Martin Albrecht and Gregory Bard, *The M4RI Library – Version 20121224*, <http://m4ri.sagemath.org> (2012).
- [7] Martin Albrecht, Gregory Bard, William Hart, *Algorithm 898: Efficient Multiplication of Dense Matrices over $GF(2)$* , ACM Transactions on Mathematical Software, 37 (1), pp. 1–14 (2010).
- [8] "J-G. Dumas and T. Gautier and M. Giesbrecht and P. Giorgi and B. Hovinen and E. Kaltofen and B. D. Saunders and W. Turner and G. Villard, *Linbox: A Generic Library for Exact Linear Algebra*, in ICMS'02, pp. 40–50 (2002).