Foreword

In order to achieve high performance on Intel® Many Integrated Core (MIC) processors, the code has to be cache friendly and must express sufficient heterogeneity, parallelism and vector capability. This can be overwhelming for those new to the Xeon Phi coprocessors. Fortunately, many matrix algorithms can be successfully parallelized and vectorized. Furthermore, in the process of preparing the MIC versions of these algorithms, there is no need to reinvent the wheel. For example, the recent releases of Intel® Math Kernel Library contain versions of matrix algorithms optimized specifically for the Xeon Phi. Using the Automatic Offload mode, one can use the MKL code prepared for standard processors to obtain good performance on the MIC. The same MKL code, without any change, can be easily recompiled and executed on Xeon Phi in native mode. The most flexible Compiler Assisted Offload mode needs some essential syntax extensions, but allows for explicit control of data transfer and heterogeneous computing, taking advantage of both multicore host and manycore coprocessors.

The purpose of this text is to present one more possibility. We shall demonstrate how to use the corresponding open source library MAGMA MIC, which allows for efficient use of Xeon Phi in matrix computations without requiring MIC-specific expertise. MAGMA MIC is a C/C++ linear algebra library, which is a successor of LAPACK and SCALAPACK, specially developed for heterogeneous MIC-based architectures by Innovative Computing Laboratory (ICL) from University of Tennessee, Knoxville (http://icl.cs.utk.edu/magma/software/index.html#license).

We propose a practical, hands-on approach. The detailed hypertext table of contents allows for easy navigation through over 200 extensively commented code samples, which can be read in an arbitrary order and used as a reference guide. To make a meaningful comparison possible, all code samples have both MAGMA MIC and MKL versions, to suit the users of both libraries.
We believe that the presented text is a valuable addition to the existing MKL and MAGMA MIC documentation. The reader interested in deeper insight into matrix computations on Xeon Phi is referred to the position [MAGMA MIC] of the bibliography.