Hands-On — CUDA SDK - Libraries, Numerical Accuracy

Siegfried Höfinger

ASC Research Center, TU Wien

October 20, 2025

→ https://tinyurl.com/cudafordummies/ii/ho4/notes-ho4.pdf

Outline

Hands-On — CUDA SDK - Libraries, Numerical Accuracy

Exercise

Q1) *Given the following matrix,*

$$\mathbf{A} = \begin{pmatrix} 0.30 & -0.61 & 0.40 & 0.37 & -0.49 \\ 0.51 & -0.29 & -0.41 & 0.36 & 0.61 \\ 0.08 & -0.38 & -0.66 & -0.50 & -0.40 \\ 0.00 & -0.45 & 0.46 & -0.62 & 0.46 \\ 0.80 & 0.45 & 0.17 & -0.31 & -0.16 \end{pmatrix}$$

how could we quickly check with the help of CUBLAS that this is actually a matrix consisting of only eigenvectors of a symmetric matrix M ?

20 min

A1) *Since the eigenvectors of a symmetric matrix M are orthogonal to each other, it follows that the inverse, A^{-1}, is simply the transposed, A^t, making the matrix matrix multiplication, $A^t \times A$ result in the unit matrix E. Matrix matrix multiplication is a straightforward case for CUBLAS (see below version for download).*

→ https://tinyurl.com/cudafordummies/ii/ho4/chck_ev.cu

Exercise

Q2) *Could we make use of CUDA managed unified memory, i.e. cudaMallocManaged(), when calling CUBLAS, for example when modifying the previous case ?*

10 min

A2) *Yes we can, as long as we don't forget to synchronize the device after the CUBLAS call (see below version for download).*

→ https://tinyurl.com/cudafordummies/ii/ho4/chck_ev_v3.cu
→ https://stackoverflow.com/questions/65501537/cudamallocmanaged-unified-memory-with-cublas