

MILC

From Nuno:

The GPU on Jetson (GK20A) with cuda capability 3.2 also supports double precision.

Performance for a 16^4 lattice volume for gauge fixing using MILC+QUDA With overrelaxation code:

GPU: $\text{time_GK20A} / \text{time_980GTX}$

single $\sim 11.2x$

double $\sim 9.7x$

CPU: $\text{time_ARM} / \text{time_hybrid}$

single $\sim 4.5x$

double $\sim 4.9x$

With FFT:

GPU: $\text{time_GK20A} / \text{time_980GTX}$

single $\sim 13.7x$

double $\sim 9.6x$

The GTX980 has $\sim 10x$ more cuda cores than GK20A.

Problems that I found when compiling in Jetson:

- had to remove -m32 from QUDA code
- cannot use cudaHostRegister(), cuda 6.5 toolkit release notes:

"Mapping host memory allocated outside of CUDA to device memory is not allowed on ARM; because of this, cudaHostRegister() is not supported by the CUDA driver on ARM platforms. If required, cudaHostAlloc() with the flag cudaHostAllocMapped can be used to allocate device-mapped host-accessible memory"

- compiling was a bit slow.

Best regards,

Nuno