Xeon Phi

David Jenkins, Ali Basden, Matt Townson

The Xeon Phi

A many-core CPU (64-72 cores, depending on model)

A standard CPU - self hosted and self booting

- Runs standard Linux, standard compiler tools
  - We currently use Centos
- Not an accelerator (in accelerators work package for historical reasons)
  - Though an accelerator version will be made available

A key design decision (important for AO RTC): 16GB High bandwidth memory

- 320GB/s theoretical, xxx GB/s measured at Durham

Wide vectorisation unit: 16 floats operated on simultaneously in each core

RTC pipeline operations

First step - a simple test using simple code performing the basic RTC operations

Image calibration, Slope calculation, Reconstruction

Using standard Intel BLAS libraries where appropriate

To get best performance

For ELT SCAO case (74x74 sub-apertures):

~800us computation time -> 1.2kHz

Therefore, this becomes the maximum goal for a real RTCS

DARC

Durham AO Real-time Controller

- The RTC for CANARY
- C-based
- Compiles without modification on the Phi
  - But not performance optimised

Improvements to DARC code

Vectorisation and memory alignment

- Previously mostly aligned to 16 byte boundaries
- Now more aligned to 64 byte boundaries

Thread synchronisation

- Previously well optimised for smaller numbers of cores
- Now, mutexes are replaced by spin locks and spin waiting

Thread affinity, count and priority

- Has always been user-tunable
- We have now investigated optimal affinity, count and priority for the Phi

Current state

Runs at about 800Hz for ELT SCAO, jitter O(10) us rms

Bottlenecks: MVM (as expected), DM command addition, calibration

A full on-sky tested code

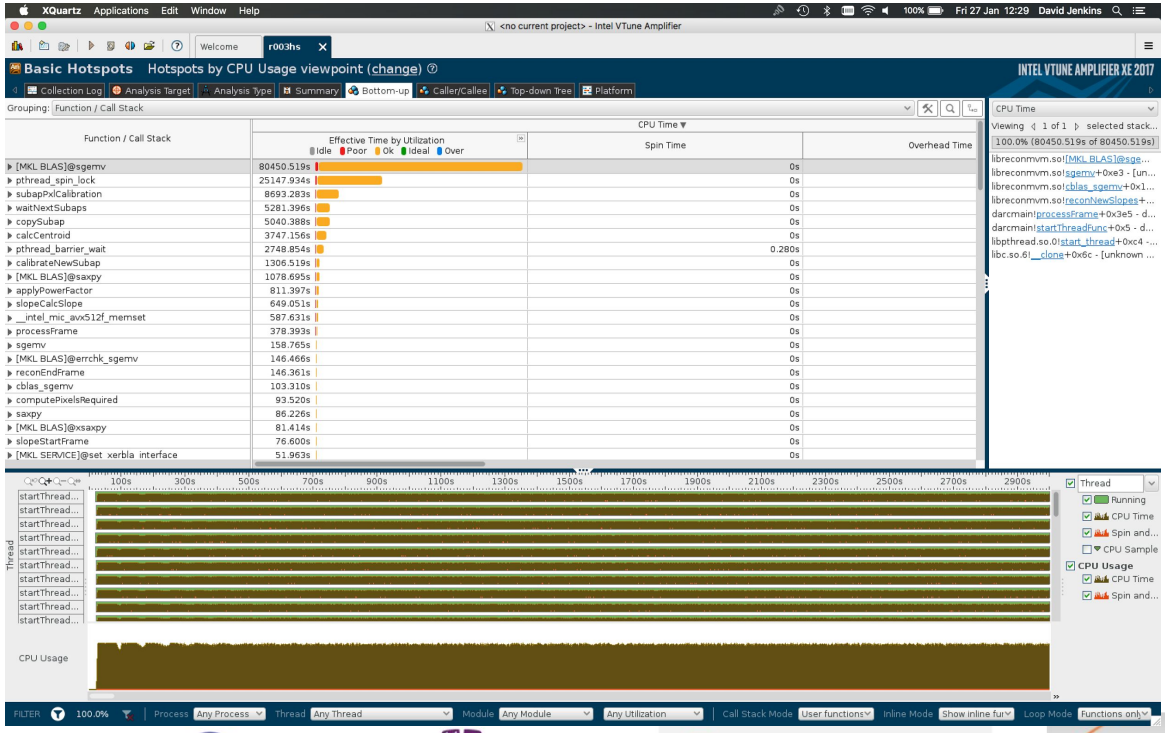

l'Observatoire de Paris LESIA
Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique

Durham University

MICROGATE

PLDA

Ongoing work

Tree-addition for DM commands (rather than linear)

    Allowing removal of global barriers

Investigation of half-precision float storage

Investigation of calibration points

Attaching a 10G camera (now attached)

Real-time kernel

AO Supervisor

To provide calibration and control

Non-real-time

But requires regular update of some parameters

Specifically, the large reconstructor

- Matrix inversion and matrix-matrix multiplication being key operations
- Such operations are well supported on the Phi
- Many such libraries exist, including extension to multiple Phi's
- (standard HPC stuff!)

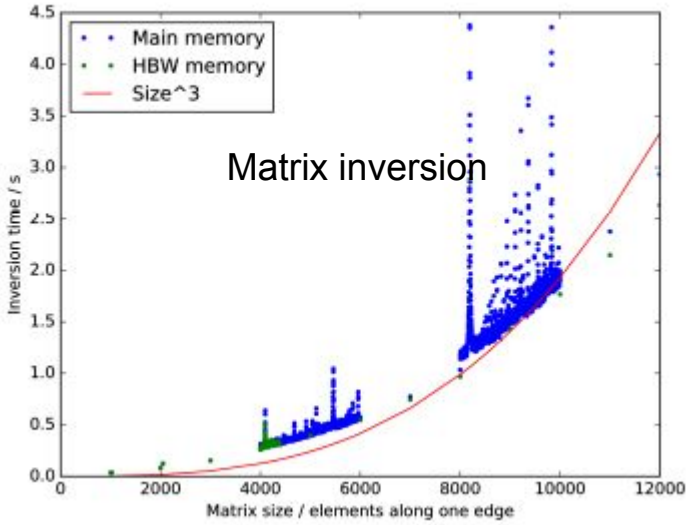

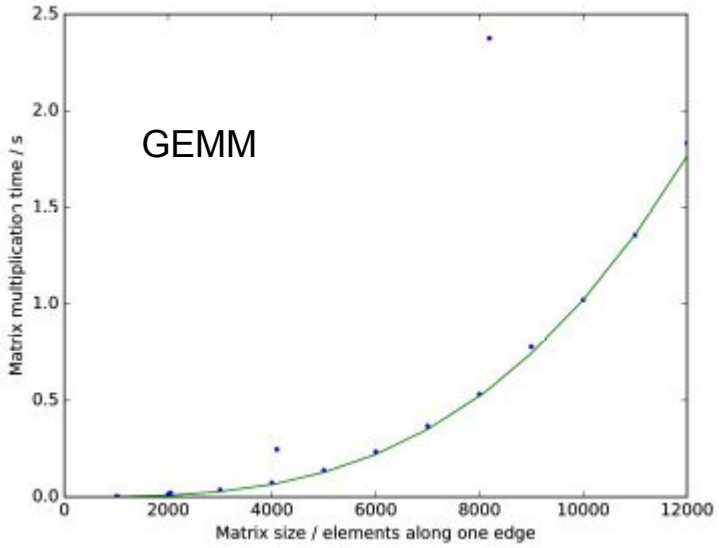

Potential hardware design

4U rack system, cost approx £40k

Other options could be:

- 1 LGS into multiple KNL
  - Allowing more advanced processing algorithms
- Each LGS multicast into another server dedicated to saving of pixel data
- KNL (or GPU) accelerator within each node for MVM

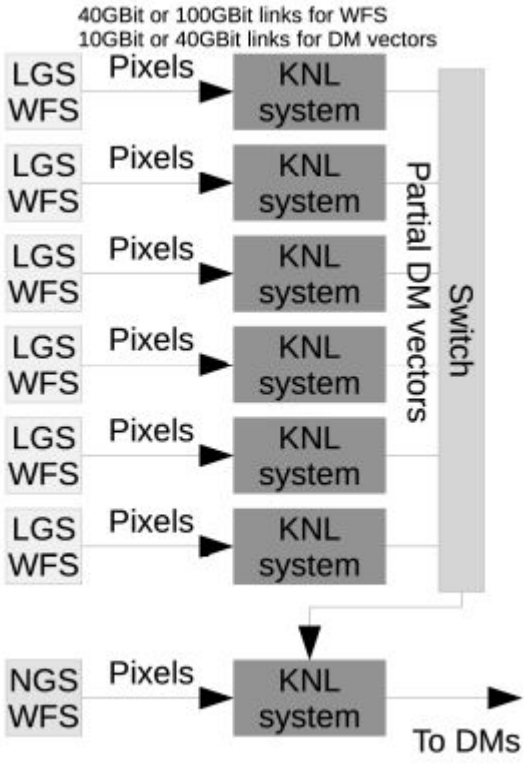

Conclusion

Xeon Phi technology is highly suited to ELT-scale AO RTC

- Low cost
- Low barrier to adoption
- Standard computer software techniques
  - We can use standard HPC libraries
- Code fully portable to other platforms:
  - Power8
  - ARM
  - AMD/Intel CPU
  - I486, Pentium, possibly i386 (with old Linux kernel)
  - A fully generic solution