ELT-SCALE REAL-TIME CONTROL ON A XEON PHI

David Jenkins

Durham University



OUTLINE

- The challenges involved in ELT-scale adaptive optics real-time control
- The Knights Landing Xeon Phi standalone processor
- Initial investigation and hardware benchmarking
- Preliminary base-line results for a simple adaptive optics SCAO loop
- Optimisation of DARC for the Xeon Phi and initial results
- Conclusions + Future considerations





WHAT'S THE PROBLEM?

- The problem size for an adaptive optics (AO) RTC scales with the fourth power of telescope diameter for a simple SCAO system
- More complex systems such as MCAO and LTAO have even greater computational requirements
- A large matrix vector multiplication (MVM) is required for many reconstruction techniques and is one of the most computationally demanding operations





WHAT'S THE PROBLEM?

- An RTC not only needs to run at a high frame rate but the frame times need to consistent, there needs to be low jitter
- Using multiple processing nodes and/or accelerator cards can introduce increased latency and jitter due to the offload of data and instructions
- FPGAs can be very deterministic and provide very consistent frame times, however they can have high start up costs, longer development times, and are not as immediately flexible as other options
- GPUs provide very high computational throughput for highly parallelisable tasks, however they can introduce accelerator offload latency and use proprietary drivers and software





THE XEON PHI: KNIGHTS LANDING

- A standalone socketed host processor or PCI accelerator card with 64+ processing cores and 16GB of High Bandwidth memory
- 512bit wide vector registers for 16 simultaneous single precision FP operations
- Socketed host is the target of this investigation to eliminate accelerator off-load latency and allow the use of standard programming tools and languages
- The Knights Landing (KNL) runs standard Linux with the option of an optimised kernel provided by Intel for full functionality





THE XEON PHI: KNIGHTS LANDING

- Is a single Xeon Phi host processor capable of meeting the requirements for ELT scale SCAO?
- Then for systems with more than one wavefront sensor a single Xeon Phi per wavefront sensor should provide the necessary computational performance
- The Xeon Phi has the option of integrated OmniPath fabric interconnect which can help to reduce the latency of distributed workloads







Knights Landing Overview



<u>TILE</u>	2 VPU	СНА	2 VPU		
		1 MR			
	Core	L2	Core		

Chip: 36 Tiles interconnected by 2D Mesh Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW DDR4: 6 channels @ 2400 up to 384GB IO: 36 lanes PCIe Gen3. 4 lanes of DMI for chipset Node: 1-Socket only Fabric: Omni-Path on-package (not shown)

Vector Peak Perf: 3+TF DP and 6+TF SP Flops Scalar Perf: ~3x over Knights Corner Streams Triad (GB/s): MCDRAM : 400+; DDR: 90+

Source Intel: All products, computer systems, dates and figures specified are preliminary based on current expe are subject to change without notice. KNL data are preliminary based on current expectations and are s without notice. 1Binary Compatible with Intel Xeon processors using Haswell Instruction S. numbers are based on STREAM-like memory access pattern when estimated based on internal Intel analysis a





KNIGHT LANDING 7210

CHOOSE YOUR OPTIMIZATION POINT							
inside						RECOMMENDED	
inside	CORES	GHZ	MEMORY	FABRIC	DDR4	POWER ²	CUSTOMER Pricing ³
7210	64	1.3	16GB 6.4 GT/s	% s	38 GB 2133 MHz	215W	\$2438
Best Value ¹ Available beginning in September ² Plus 15W for i ³ Pricing shown is for parts without integrated fabric.	ntegrated fabric Add additional \$278 for integrated fat	pric versions of thes	e parts. Integrated fabric p	No Parts available in Octo	96 GB		(intel)







KNIGHT LANDING OPTIONS

CHOOSE YOUR OPTIMIZATION POINT							
	r Integrated					RECOMMENDED	
inside	CORES	GHZ	MEMORY	FABRIC	DDR4	POWER ²	CUSTUMER PRICING ³
7290 ¹ Best Performance/Node	72	1.5	16GB 7.2 GT/s	Yes	384GB 2400 MHz	245W	\$6254
7250 Best Performance/Watt	68	1.4	16GB 7.2 GT/s	Yes	384GB 2400 MHz	215W	\$4876
7230 Best Memory Bandwidth/Core	64	1.3	16GB 7.2 GT/s	Yes	384GB 2400 MHz	215W	\$3710
7210 Best Value	64	1.3	16GB 6.4 GT/s	Yes	384GB 2133 MHz	215W	\$2438
¹ Available beginning in September ² Plus 15W for integrated fab ³ Pricing shown is for parts without integrated fabric. Add addition	ric al \$278 for integrated fab	ric versions of these	e parts. Integrated fabric p	arts available in Octob	ber.		(intel)





STREAM MEMORY BANDWIDTH

- A tool for measuring the memory bandwidth of a system
- Provided as part of Intel's MICPERF benchmarking utility
- Results for the available 7210:
 - Memory performance: 442.14 GB/s (approx 500 GB/s for top end Xeon Phi)
- Compare to the quoted maximum performance:
 - Memory maximum: 400+ GB/s
- NVidia Tesla P100 quoted maximum:
 - Theoretical memory maximum: up to 732 GB/s
 - Roughly 580 GB/s achievable (estimate)





MATRIX INVERSION

- Important for the calculation of control matrices for most reconstruction techniques
- MOAO will have particularly high requirements due to the number of WFSs and DMs involved in the reconstruction
- LU ('lower upper') decomposition using the Intel Math Kernel Library, MKL
 - $M^T \times (M \times M^T)^{-1}$
- MCAO: for 60,000 x 10,000: 3.1s + 2s + 3.1s = 8.2 s





MATRIX INVERSION & MATRIX MULTIPLICATION









PRIMARY TEST CASE

- MICADO SAMI module for first light SCAO correction, European-ELT
- MICADO will eventually use the MAORY MCAO system
- E-ELT SCAO
 - 80x80 or 74x74 WFS sub-apertures with 8x8 pixels, Shack-Hartman
 - ~5000 actuator DM :- E-ELT high-order adaptive mirror (M4)
- Potential requirements
 - 500Hz framerate
 - <100µs RMS jitter</p>





MVM BENCHMARK

- Simple matrix vector multiplication (MVM) operation using the Intel MKL cblas_sgemv function
- The MVM is the most computationally demanding aspect of least-squares reconstruction
- It's also a major aspect of other reconstruction techniques such as LQG
- Profiling how the MVM performance scales on the KNL with problem size and available processing threads







20/12/2016



SIMPLE RTC LOOP

- Performance profiling of a very simple least-squares SCAO like RTC loop on the Xeon Phi, not pipelined
- Pixel calibration, centre of gravity slope calculation, MVM reconstruction and application of a constant gain to DM commands
- Using OpenMP for multi-threading and MKL for the MVM, timing is done in software using an OpenMP function call
- All written in the standard C programming language with pragma directives to enable OpenMP functionality





SIMPLE RTC LOOP



Problem size:

20/12/2016





SIMPLE RTC LOOP

Currently no real-time kernel. Should help with peaks (outliers)







NEXT: DARC ON THE XEON PHI

- Optimisation of the Durham Adaptive optics Real-time Controller (DARC) for the Xeon Phi KNL
- DARC is written in C and so is immediately able to be compiled and run on the Xeon Phi
- However some optimisation is required to achieve peak performance and these optimisation can be carried over to other standard CPU based systems
- Currently very early preliminary stage with only minor changes carried out





RESULTS: LATENCY AND JITTER





Sub-apertures	Pixels	Mean frame- time (µs)	Mean frame-rate (Hz)
80x80	6x6	1550 +- 39	646 +- 16
	8x8	1650 +- 36	607 +- 13
74x74	6x6	1350 +- 34	743 +- 18
	8x8	1450 +- 33	692 +- 15

74x74, 8x8 percentiles:

99%	1520 µs	660 Hz
99.9%	1550 μs	640 Hz
99.99%	1610 µs	620 Hz





MAIN AREAS OF OPTIMISATION

- Optimise the current multi-threading capabilities of DARC for many more threads
- DARC multi-threading uses POSIX threads (Pthreads), a low level execution model independent from any one language
- As the KNL is a socketed CPU running Linux Pthreads are already supported
- Optimise array operations for vectorisation
- Set-up the BIOS, kernel, and OS for optimal low-latency performance







CONCLUSIONS

- KNL can compute an ELT scale MVM at a rate greater than 1.5 kHz
- A baseline simple RTC loop can perform at around 1.26 kHz
- DARC, full RTC in simulation mode, currently running >600 Hz
- Still some more work needs to be done to achieve optimal performance
- Jitter is really quite good already <40 microseconds for DARC
- Could meet the potential specification for ELT scale SCAO of ~500Hz with <100microsec jitter
- Shows promise also for multiple WFS multiple Xeon Phi MOAO systems





20/12/2016



FUTURE

- More testing and optimisation to be done
- Real-time kernel
- Memory and cache
- More vectorisation
- Pixel calibration can be done as pixels arrive rather than on a subaperture basis
- Other algorithms, LQG, correlation, match filtering
- Real 10G camera data
- Investigate multiple Xeon Phis for more complex systems
- Full MAORY demonstration
- Future generations? Knights Hill ~2X performance





Thank-you for listening

