Real-time Control for the E-ELT

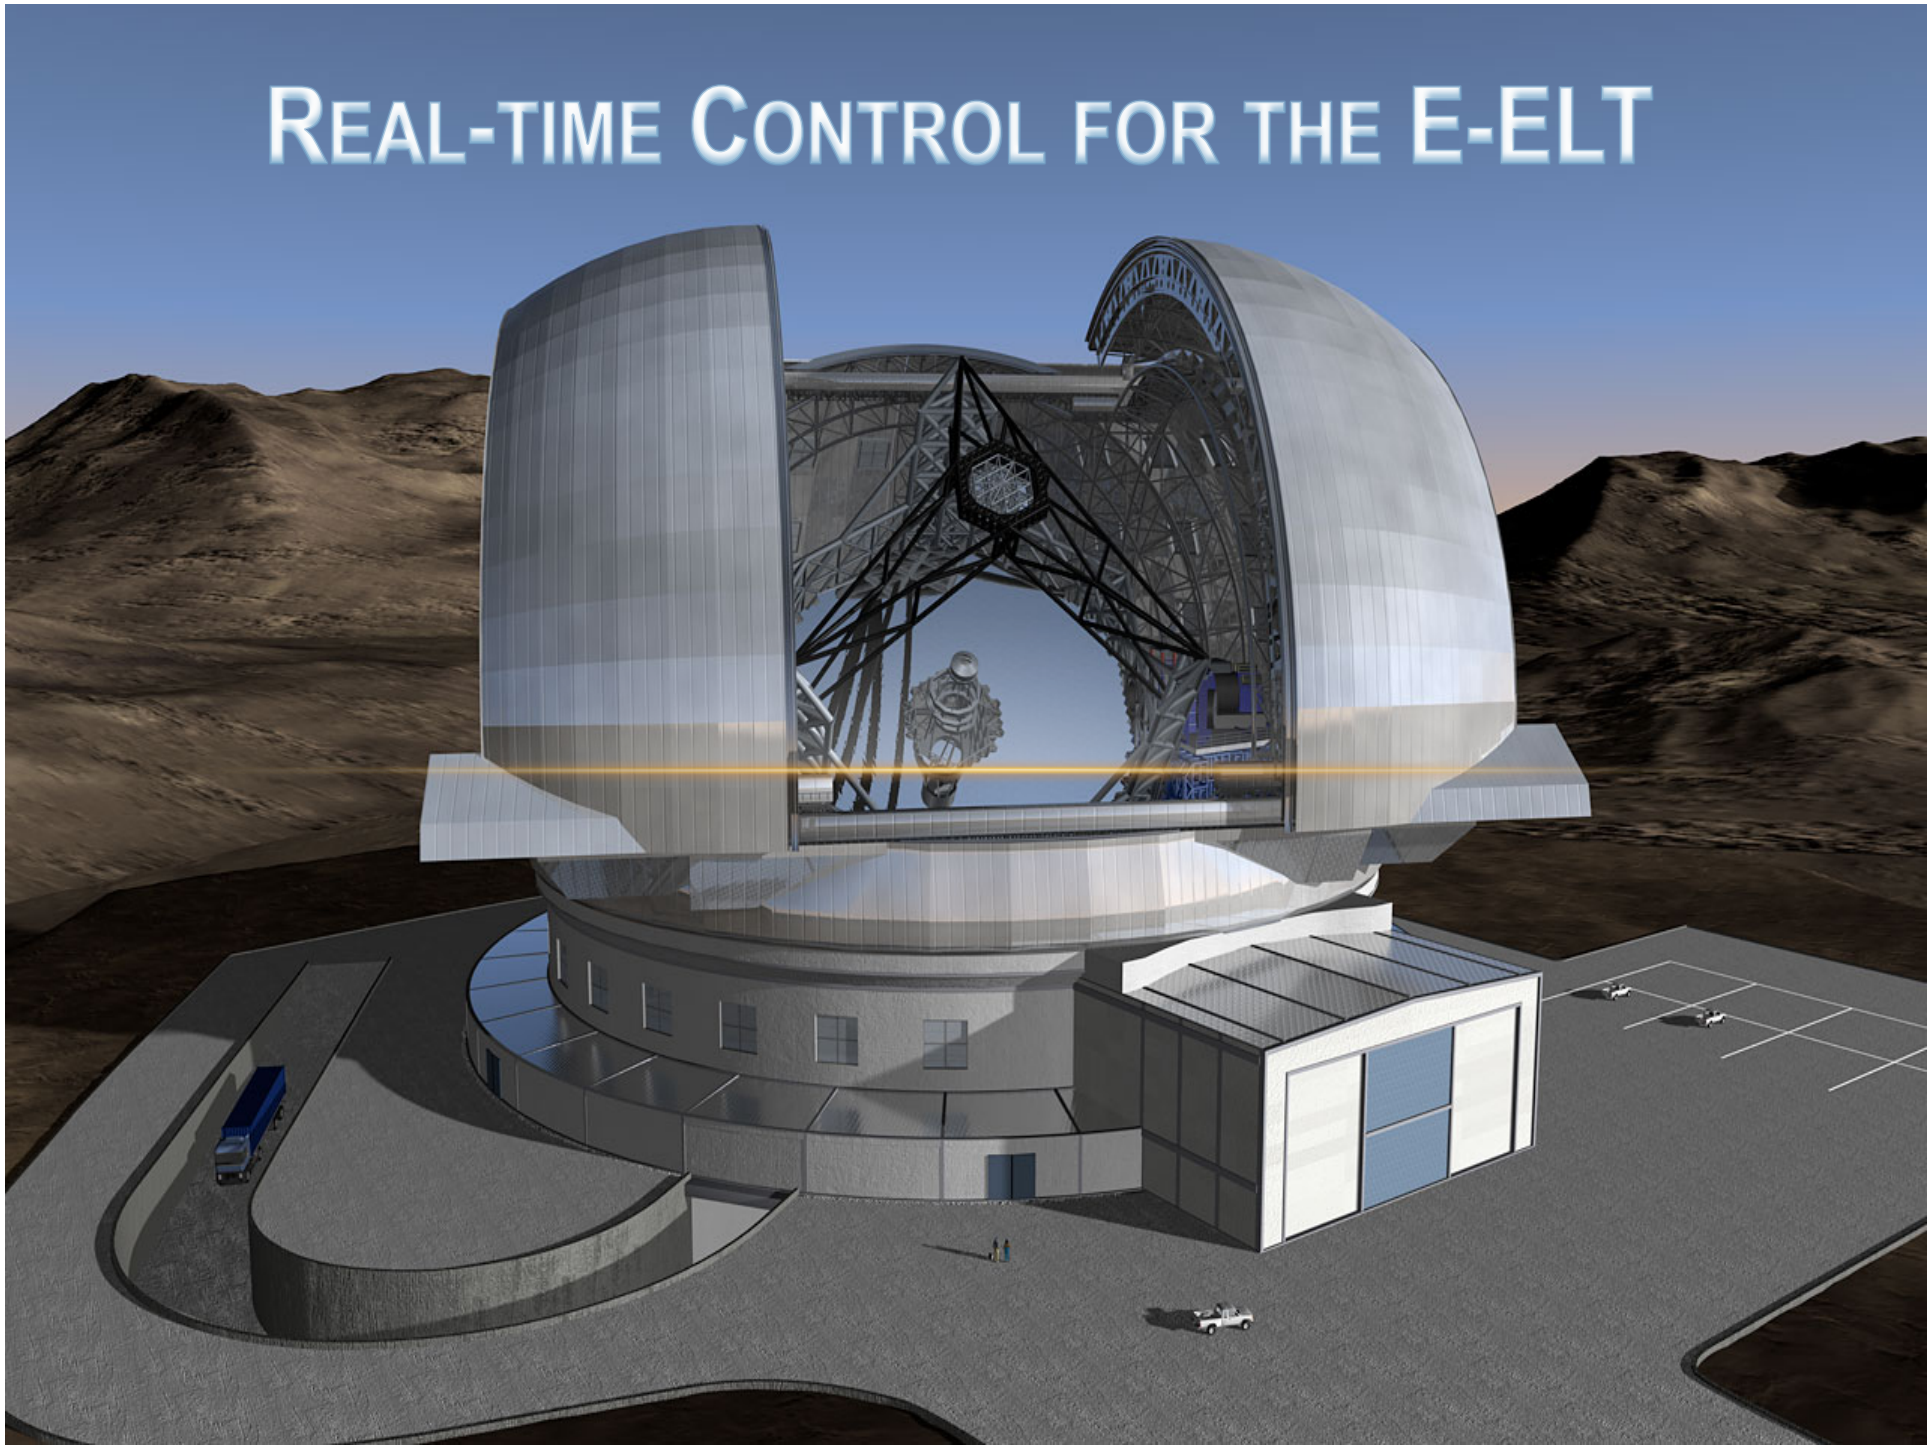

Talk Outline

- Overview of Adaptive Optics at Durham
- AO RTC – What we have done in the past
- What should we do in the future? – Green Flash
- Some random thoughts about new concepts…

AO Projects and Collaboration

- Turbulence monitoring (James talk on Tuesday)
- AO end-to-end simulation (DASP)
- DRAGON – Lab test bench – Using DARC RTC
  - Test new AO concepts; Test new RTC concepts
- CANARY – On-sky test bench (4m telescope) – Using DARC RTC
  - Test new AO concepts; Test new AO algorithms (Urban talk Tuesday)
- Real-time Control – Green Flash (Damien talk on Monday)
  - Demonstrate real-time HPC for AO
- Latter 2 in highly successful collaboration with LESIA, Obs de Paris

The immediate future of AO…

- New techniques for the VLT (AOF) and the E-ELT:
- Improved sky coverage – LGS
- Wider field of view
- Tomography – Multiple LGS and NGS
- Multi-object AO – Open loop systems - MOSAIC
- New algorithms – CuRe, LQG, predictive systems …

- All of this program assumes we have enough real-time computing power
- Do we?? If so, what form should it take?

THE PAST: AN AO RTC IN 2000

- NAOMI facility AO system on the 4.2m WHT.
- Deformable mirror with ~300 actuators
- 128x128 pixel CCDs read out at about 200 frames per second
- Inner control loop run at 1 KHz
- No commodity computer could do this. We needed to handle the data in parallel
- Solution: Multiple Texas C40 DSPs with fast interconnect.

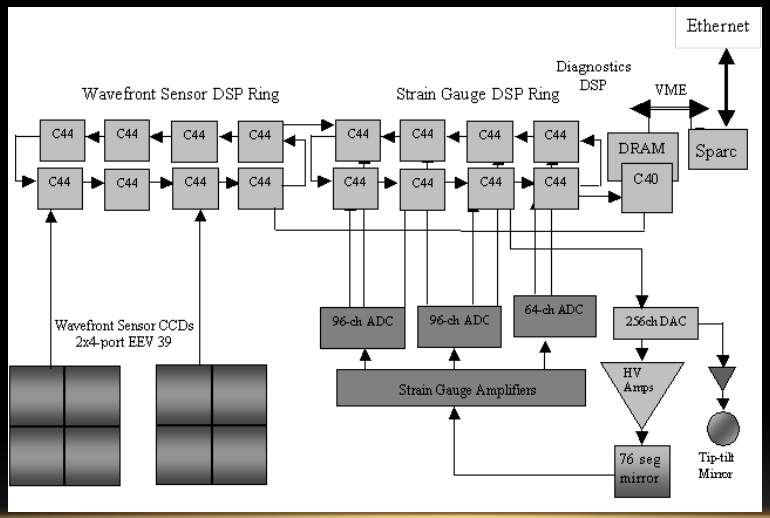

The present:
An AO RTC system in use

- AO on the 8m Very Large Telescope (Note the date!)
- ESO SPARTA – A heterogeneous parallel processing approach (Marcos talk - Tue)
- Required for example for exo-planet finding with SPHERE
- 240x240 pixel detector framed at 1.5 KHz. Low latency, low jitter
- FPGAs for parallel pixel handling
- DSPs for linear algebra wave-front reconstruction
- Will also be used for the Adaptive Optics Facility – Deformable secondary
- Problem: Upgrade path

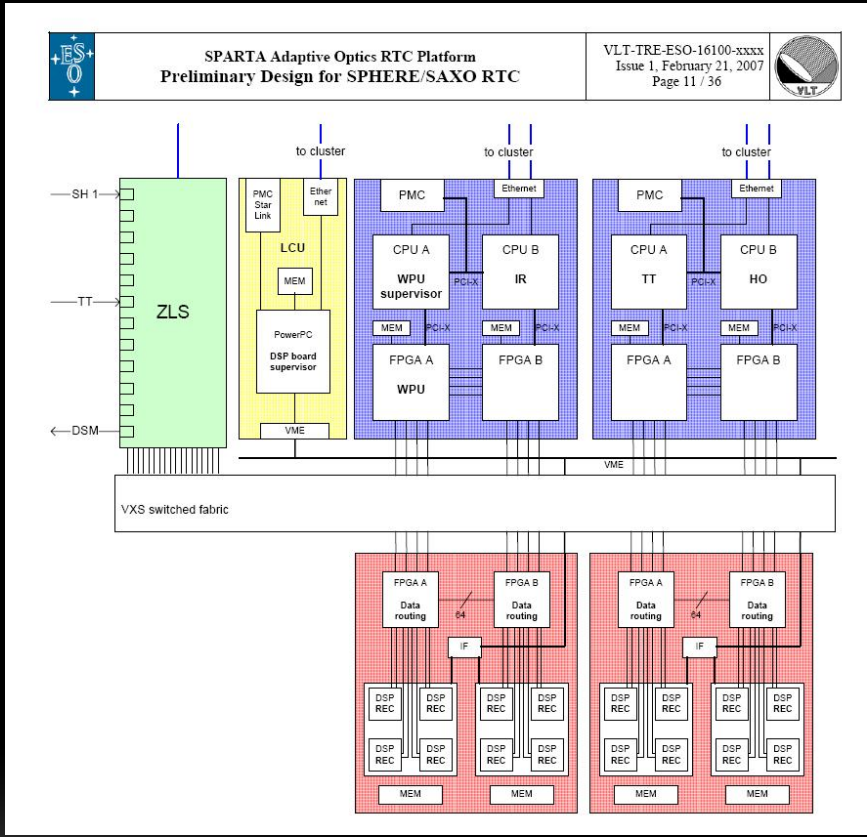

The Future: ???

- We have always required more power and throughput than COTS computer systems can give us.
- Is this still true?
- Maybe only partly…

- How much 'custom stuff' do we need?
- There are huge cost, maintenance and upgrade path benefits to using COTS

- We are not the only people needing real-time HPC
- Keep an eye on what Google and others are doing…

GREEN FLASH

- Damien talk - Monday
- An EU funded H2020 project under HPC call: FETHPC-2014
- "Energy efficient high performance computing for real-time science"
- Demonstrate European expertise in real-time HPC
- A case study: A real-time control computer for the E-ELT
- Develop a working prototype to meet the MCAO requirements by Sept 2018

Top Level Architecture

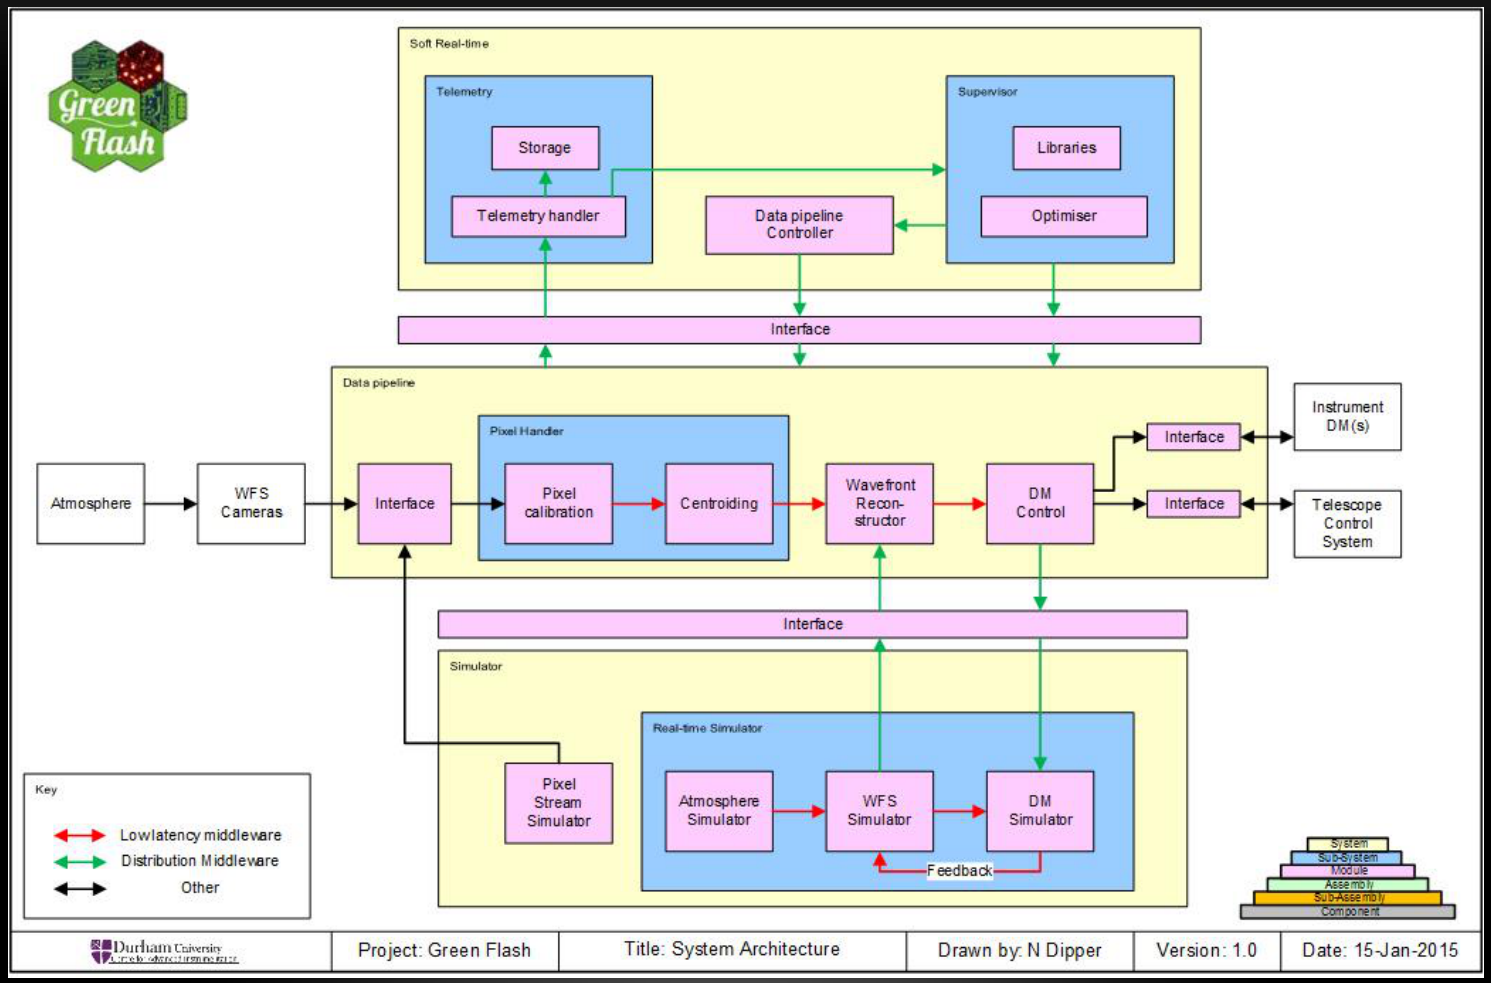

Durham GF Work-packages

- Accelerators – Urban Bitenc
  - Phi (David Jenkins), FPGA (Deli Geng)
- Simulator – Andrew Reeves
  - End-to-end
  - 10G camera simulator (FPGA)
- Ecosystem – Eddy Younger
  - FPGA development environment: PLDA - Quickplay
  - Middleware and Abstraction

PORTABLE PERFORMANCE

- We are looking not only at hardware options but also at software techniques.

- RTC will be an abstracted set of software that assumes that there will be sufficient parallel multi-core HPC hardware to meet the performance requirements
- Designing software with sustainability and portable performance. (Pierre talk)
  - Directives approach: OpenMP, OpenACC
  - Many super-computer applications are 5% efficient!! (cf Hatem talk)
  - Libraries (EG MAGMA) will however be very efficient on multi-core systems
- The E-ELT RTC systems in 2024(?) will NOT use the hardware selected in 2017
  - They will use 'Futureware' (EG Skylake + FPGA – Intel talk)

ABSTRACTION

- Hardware Abstraction
  - OpenMP
  - OpenACC
  - OpenCL …
- Software abstraction is equally important

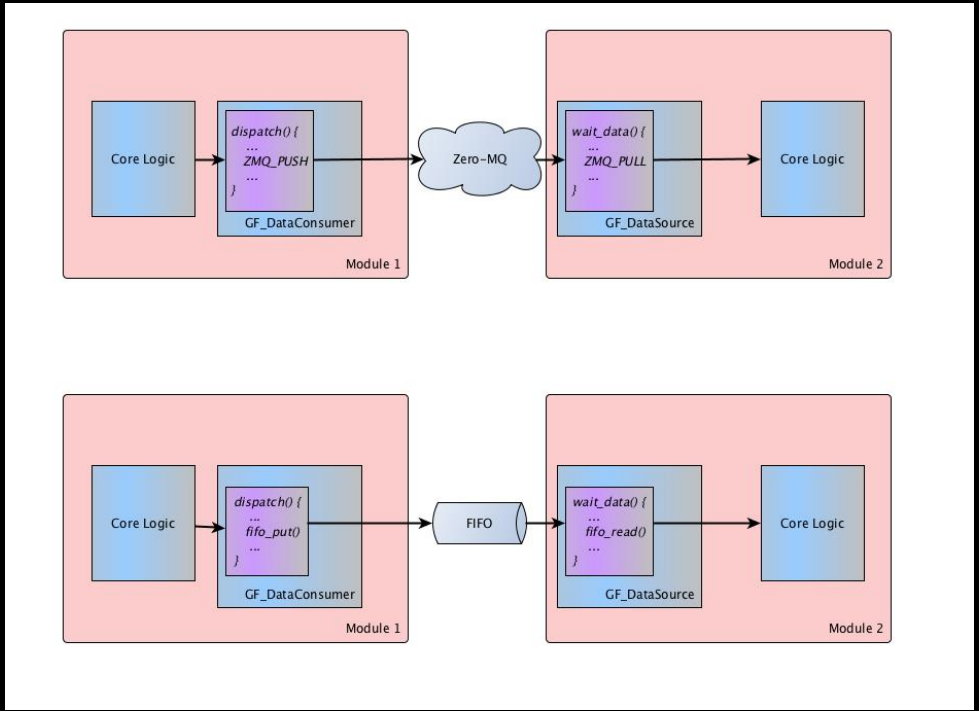

FPGAs

- Try to keep the system hardware agnostic
- The exception is (maybe) the FPGA (Different programming model)
- FPGAs should provide the perfect technology
  - Deterministic and can be customised to the application
- The problems are:
  - Memory bandwidth – This is changing with HBM
    - EG Microgate cards with HMC
  - Ease of programming – This is changing with PLDA QuickPlay …

PROGRAMMING FPGAS

- The holy grain (for years!) has been: C to HDL
- A very impressive new product: QuickPlay / QuickStore
  - IP development using QuickPlay by Lazar Staycov
- PLDA are members of Green Flash
  - We use Quickplay wherever possible in FPGA development
- There will be FPGAs in ELT RTC systems – What will they do?
  - Doing everything in FPGA might require custom cards…
  - What are they best at:
    - Data handling and interconnection– EG Camera pixel data
    - Handling data 'on the fly' – The smart interconnect concept in GF
    - Implementing algorithms that do not often require changing

Possible architecture – 1: MCAO

Durham University

CfAI Centre for Advanced Instrumentation

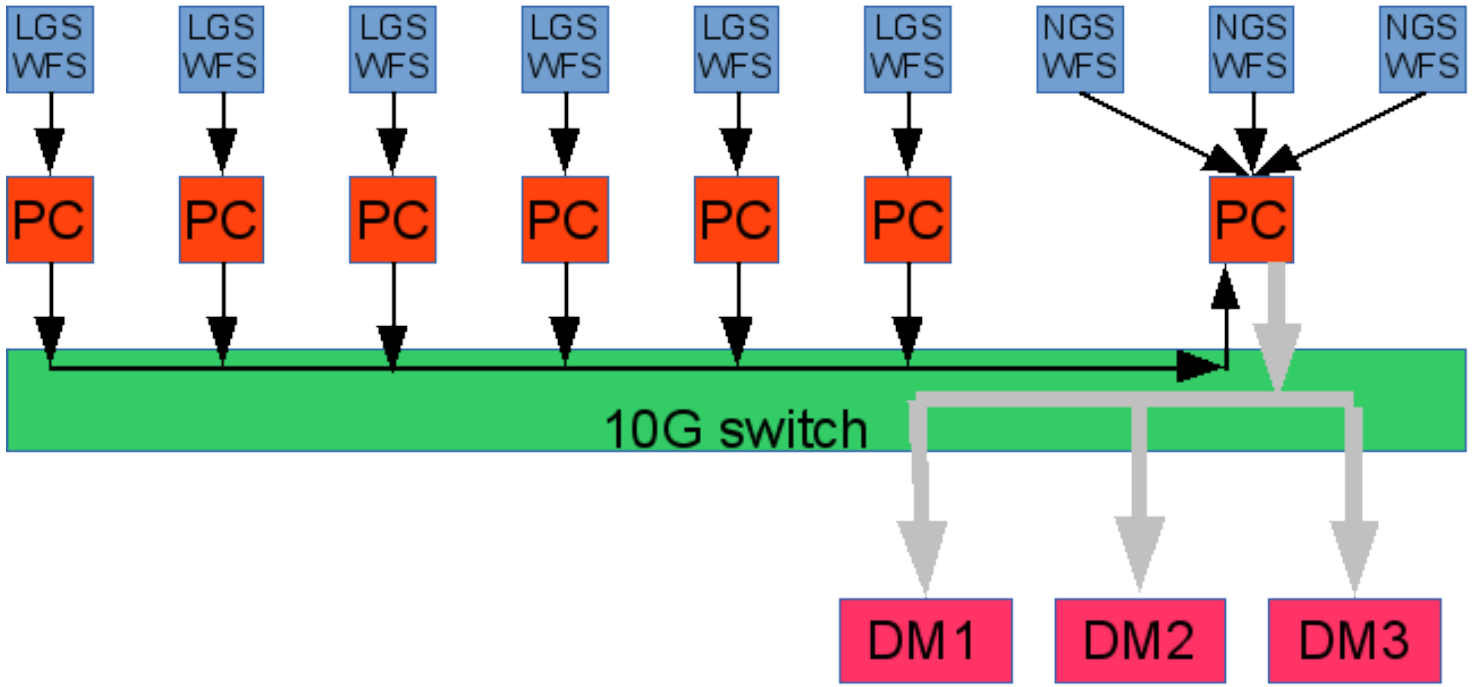

Possible architecture – 2: MOAO

Durham University

CfAI Centre for Advanced Instrumentation

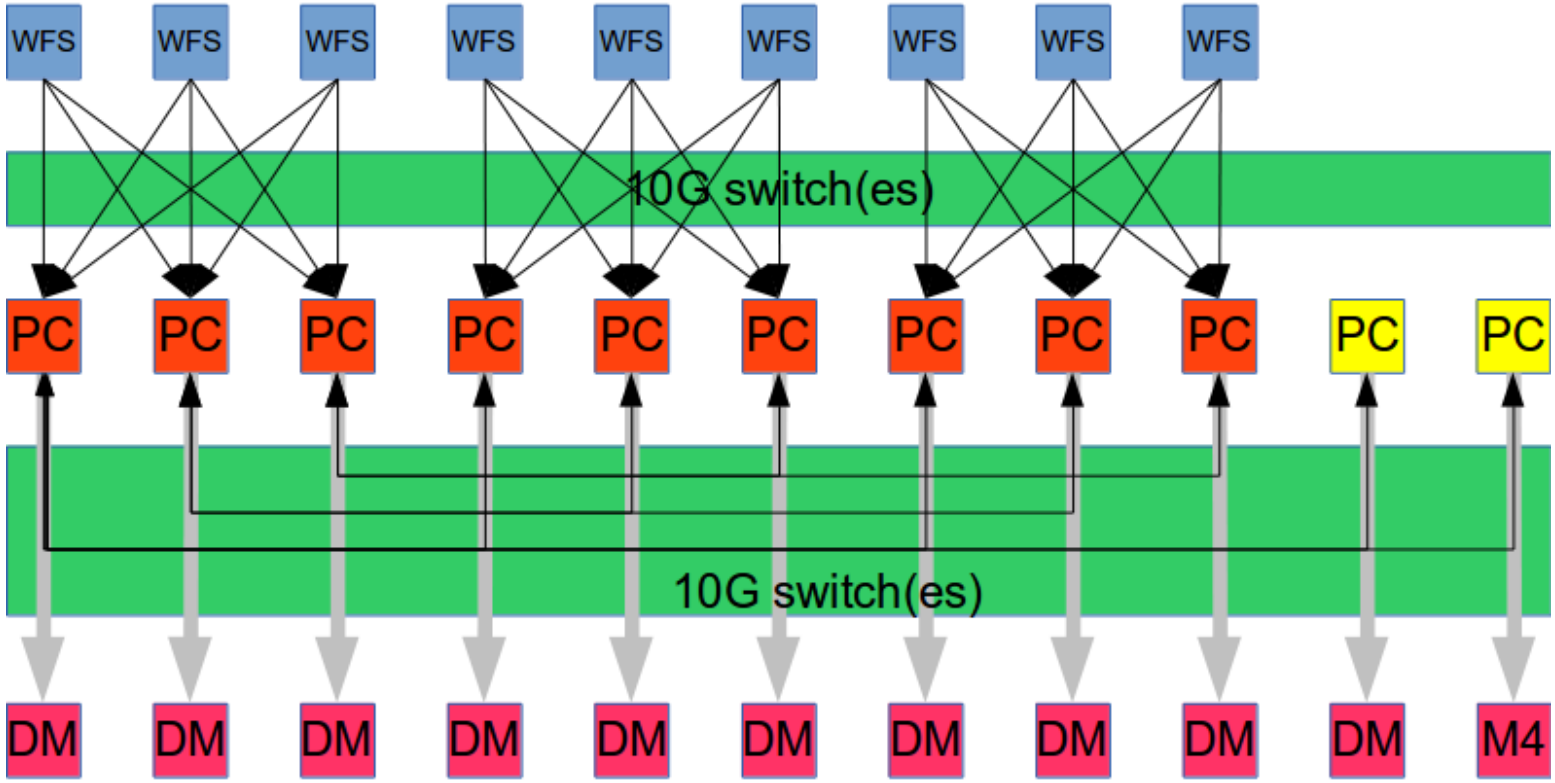

PC = Phi system taking 3 WFS, sharing slopes with 2 others, reconstructing 70k x 1k @250Hz

PC = Phi system receiving 70k slopes, reconstructing 70k x 3k @250Hz

New concepts - 1: Smart CMOS/FPGA Camera

- LGS spot elongation – Large numbers of pixels needed for ELT LGS WFS (cf CANARY)
- Custom chip development is expensive.
- Can we use a cheaper COTS chip closely coupled to an FPGA?
- EG. A professional Ultra HD video-imaging sensor
- Fast readout is achieved via a 32 lane 3Gb/sec interface which is suitable interfacing to a current generation FPGA system.
- This bandwidth is only achievable in a close-coupled detector/FPGA system over short PCB tracks rather than over medium range cabling
- THAT is something an FPGA is good at. (+ Full control of O/P protocol)
- We have a program to investigate this at Durham. This is currently at TRL 0!
- Also applicable in Nuclear Fusion (ITER)

New Concepts – 2: Neural networks

- An application agnostic technique
- ANNs are good at data mining, learning and predicting
  - Used extensively in commercial real-time analysis of huge data sets
- Neural networks as an alternative wave-front reconstructor
  - Tried on CANARY – It works!
- Given adequate opportunities to learn, can a neural net predict turbulence and vibrations?
- Can we remove the requirement for updating huge control matrices in real-time?
- Can we reduce the RTC latency and make the system more stable?
- Can we provide a 'pseudo-deterministic' system
- We are investigating the possibilities using (~~smart interconnects~~) smart masters students

Conclusions

- Priorities:
  - Software techniques
    - These can be demonstrated on various hardware options
  - Software and hardware abstraction
  - These are precepts of Green Flash
- FPGAs will have a role to play
  - Depends on firmware development environments

- We can build systems NOW that meet the E-ELT requirements
  - But they must have a clear upgrade path
  - System could be substantially COTS devices (and libraries)
  - Green Flash will provide a design based on parallel processing technology