Challenges in real-time control of adaptive optics systems for high resolution ground based astronomy

Donald Gavel and Marc Reinig
Laboratory for Adaptive Optics
University of California Observatories
UC Santa Cruz

Accelerator-based Computing and Manycore Meeting
SLAC, Stanford, CA
November 30, 2009

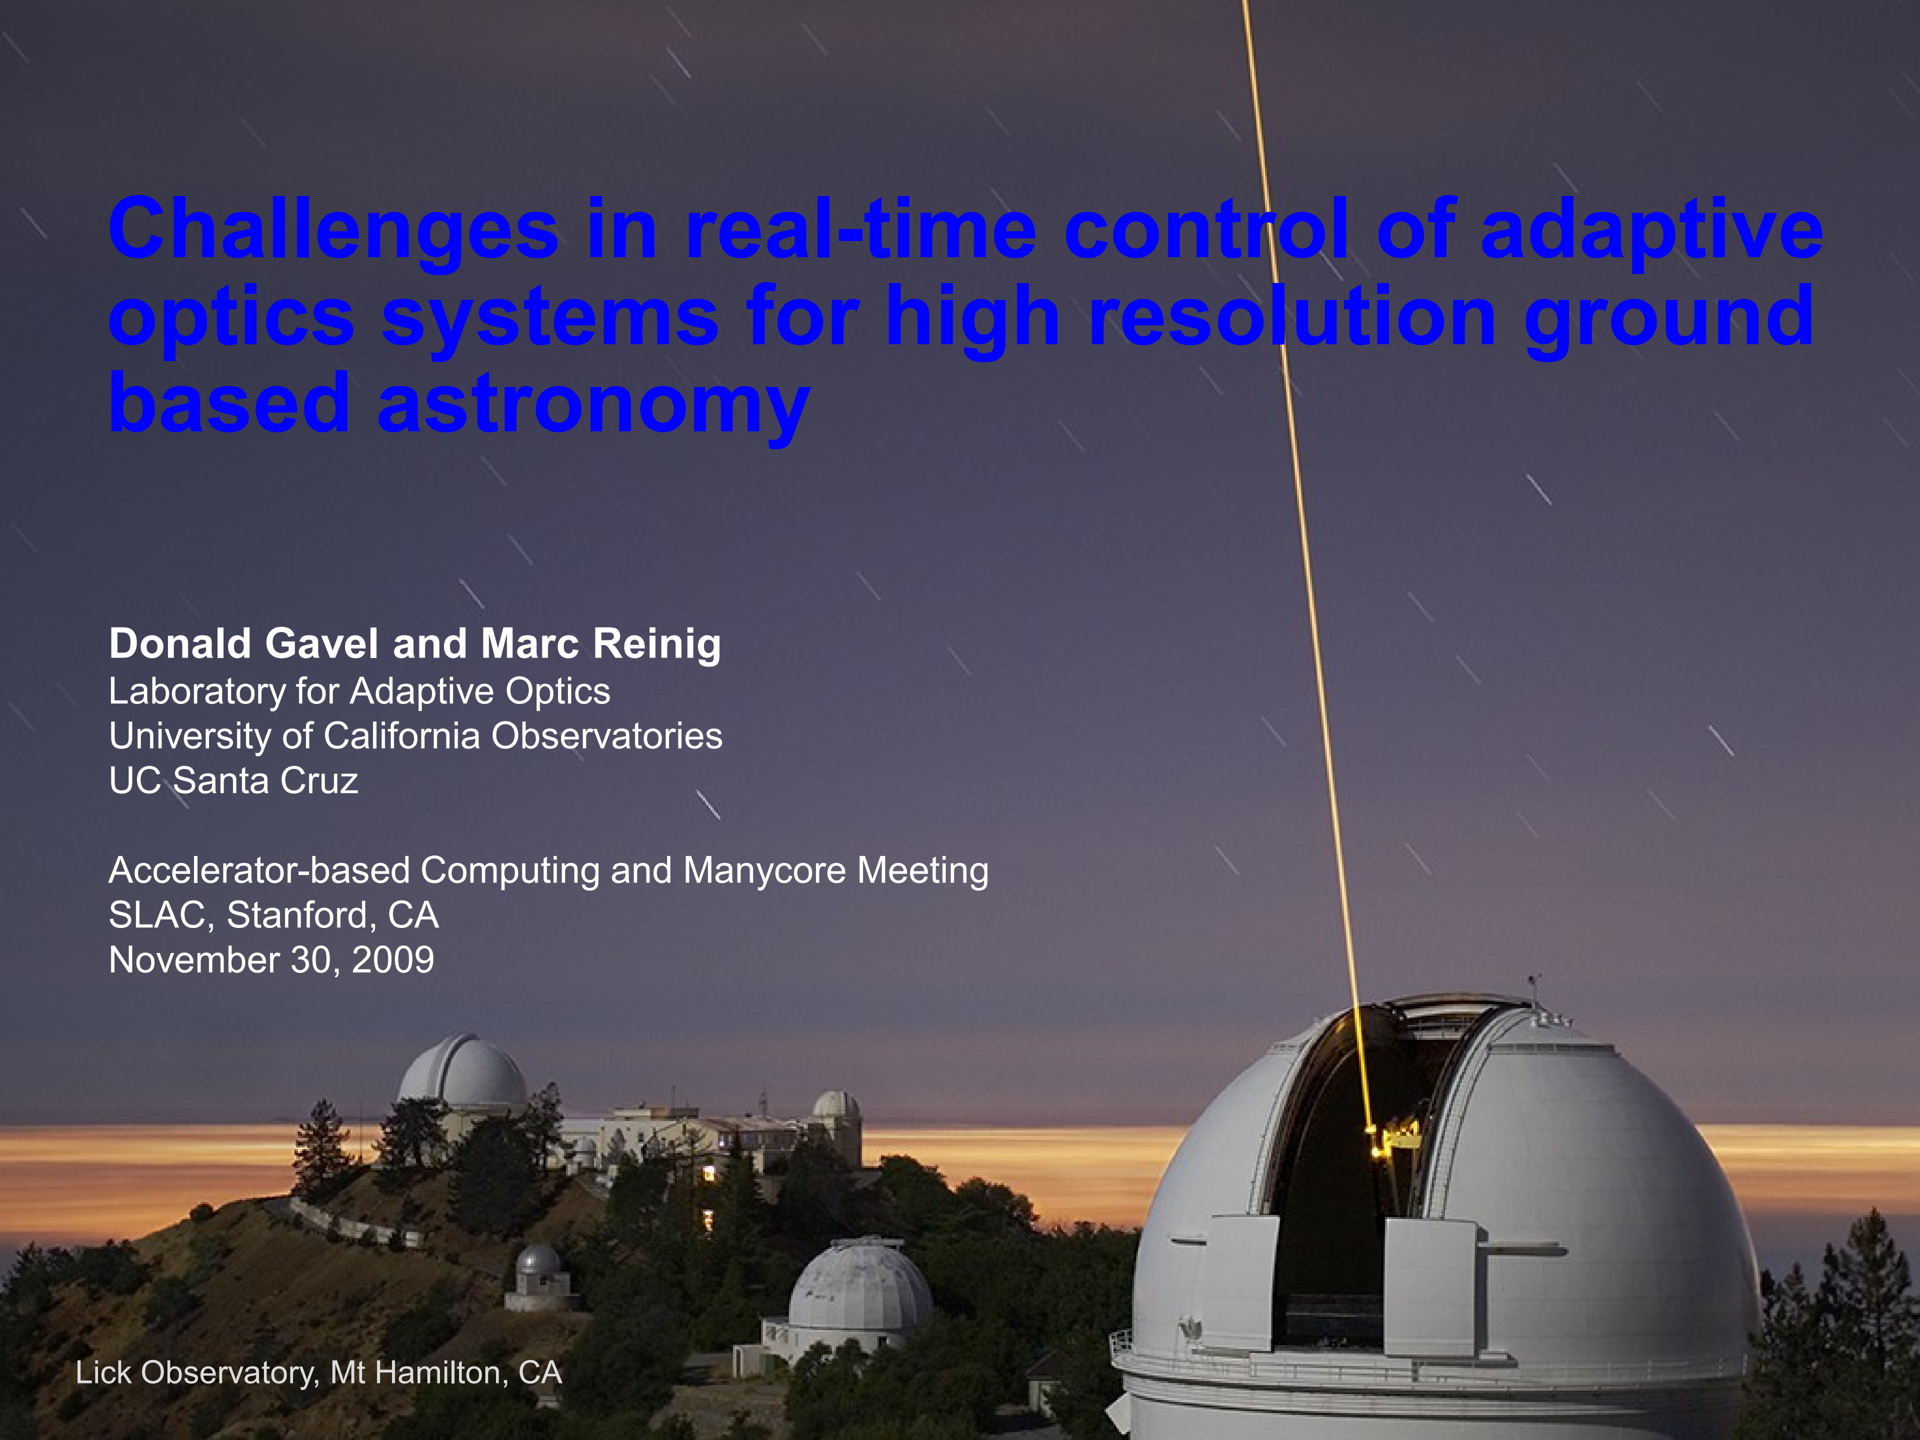

Lick Observatory, Mt Hamilton, CA

The Atmosphere Blurs Astronomical Images

- Temperature fluctuations in small patches of air cause changes in index of refraction (like many little lenses)
- Light rays are refracted many times (by small amounts)
- When they reach telescope they are no longer parallel
- Hence rays can't be focused to a point:

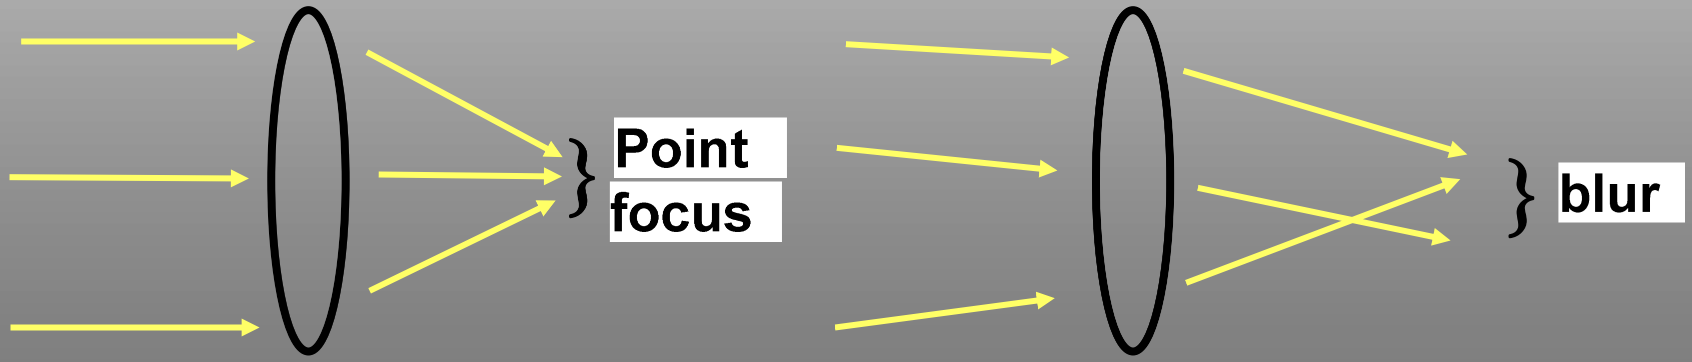

Parallel light rays

Light rays affected by turbulence

Adaptive Optics inverts the wavefront aberration with an “anti-atmosphere” (deformable mirror)

Feedback loop:
next cycle
corrects the
(small) errors
of the last cycle

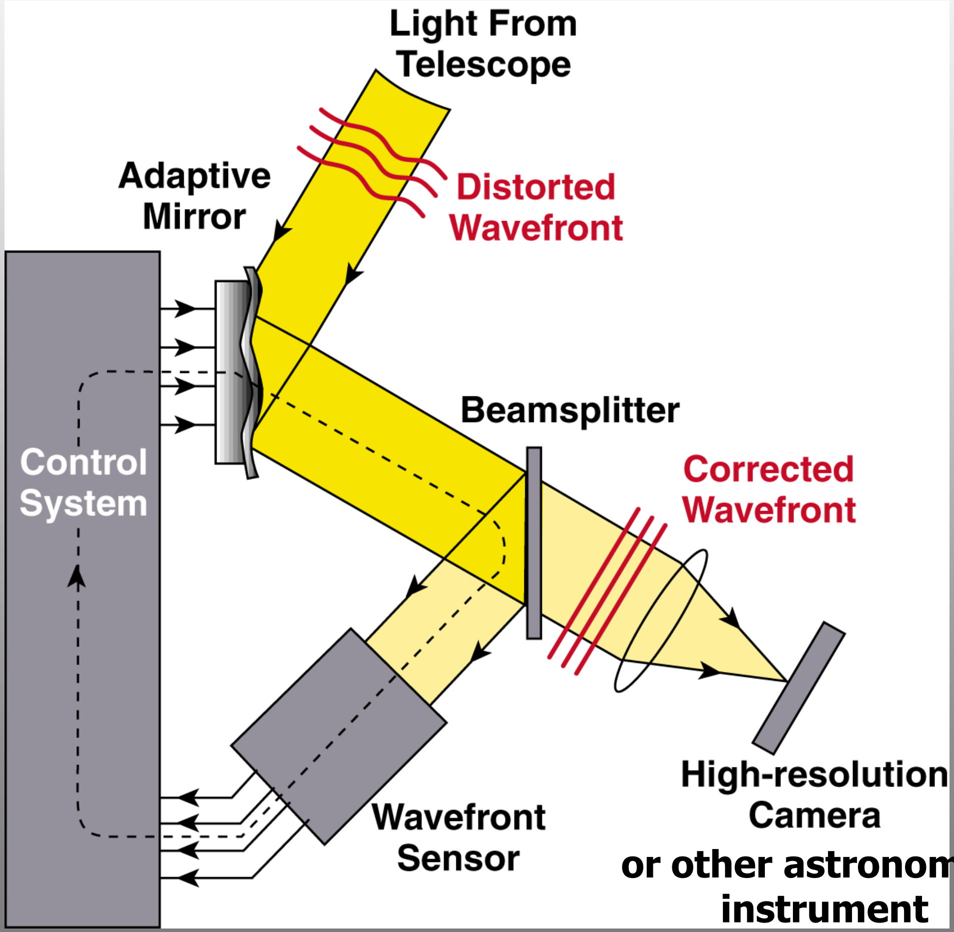

or other astronomical instrument

Astronomy with Adaptive Optics: AO on the Keck Telescope brings the Galactic Center into focus

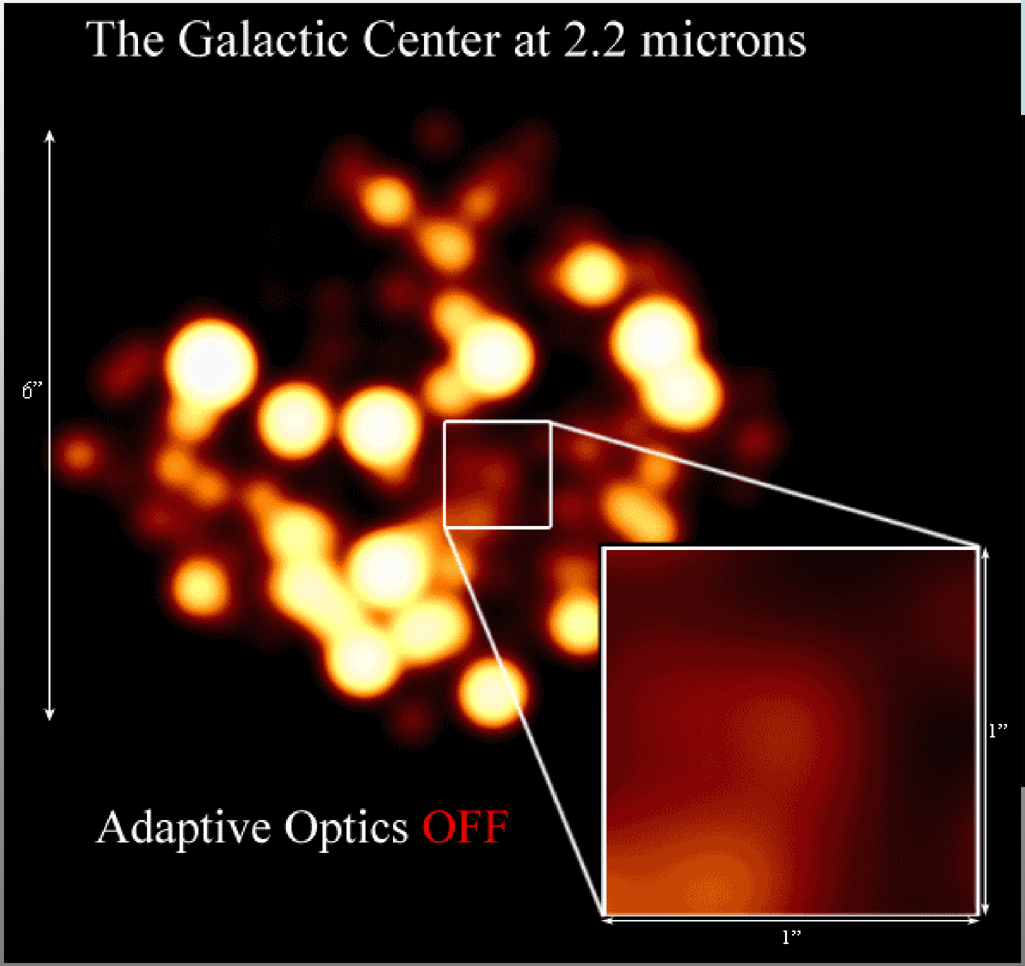

Prof. Andrea Ghez,
UCLA Galactic Center Group

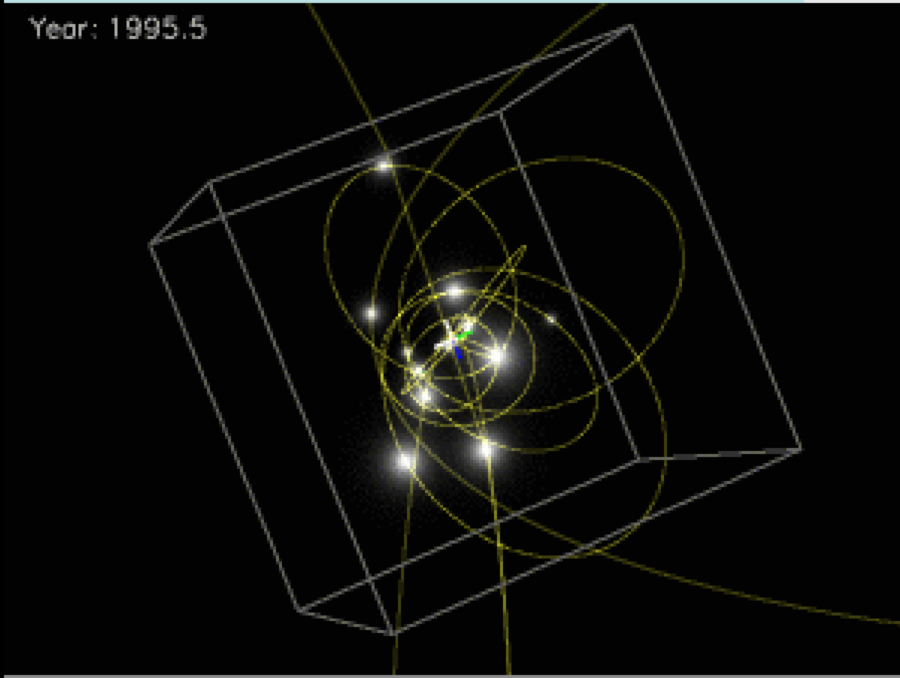

AO correction needs to keep up with atmospheric turbulence: ~ 1000 updates / second

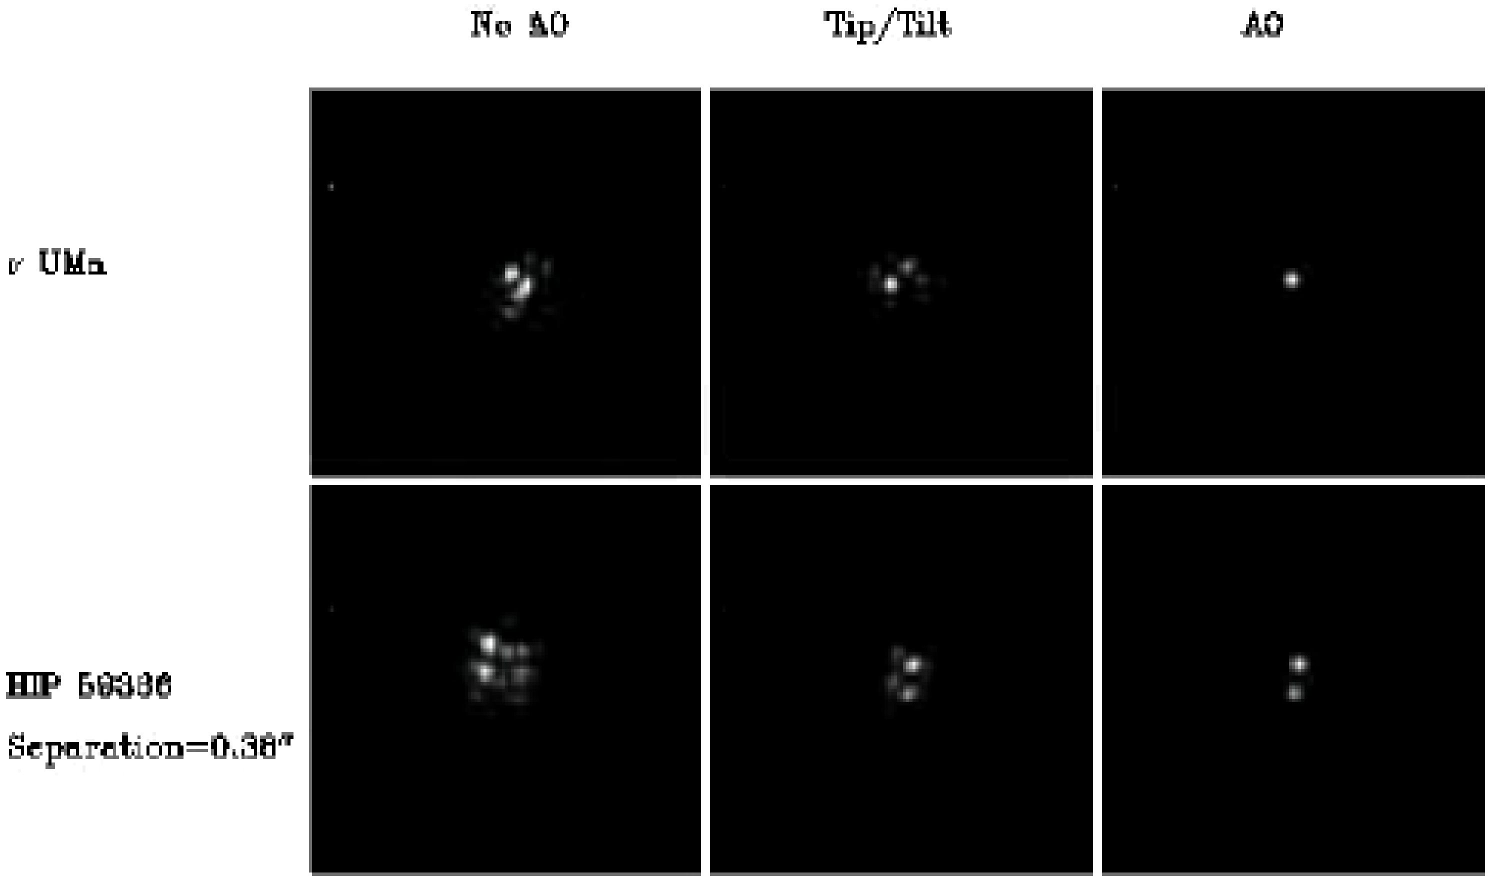

Closed loop Strehl=0.74, 2.2μm, r_0=18cm at 6500Å

57ms exposures, 4.8" field of view

If there is no nearby star, make your own "star" using a laser

Concept

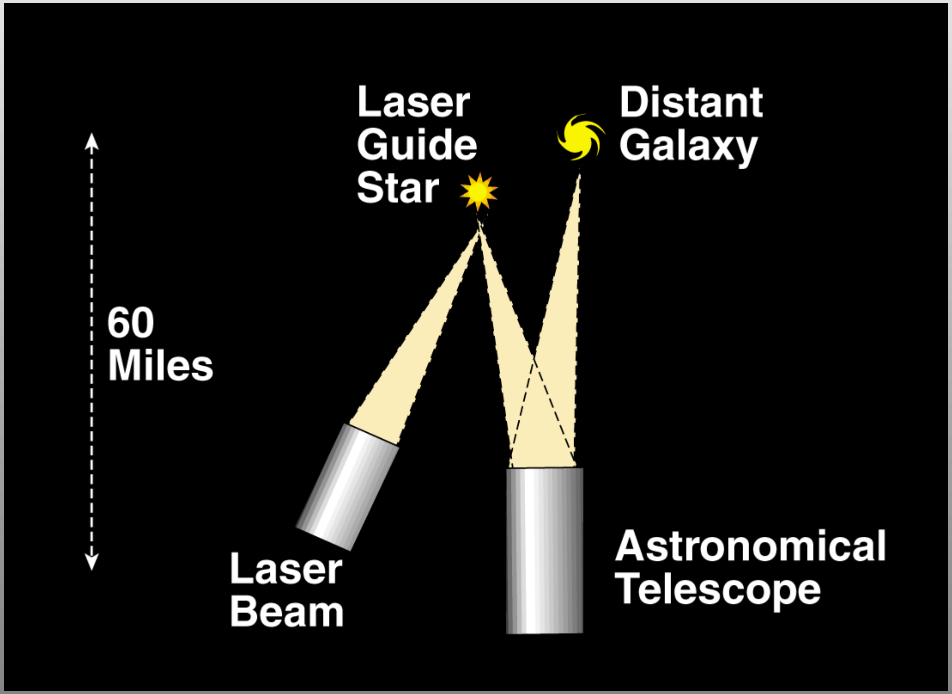

Implementation

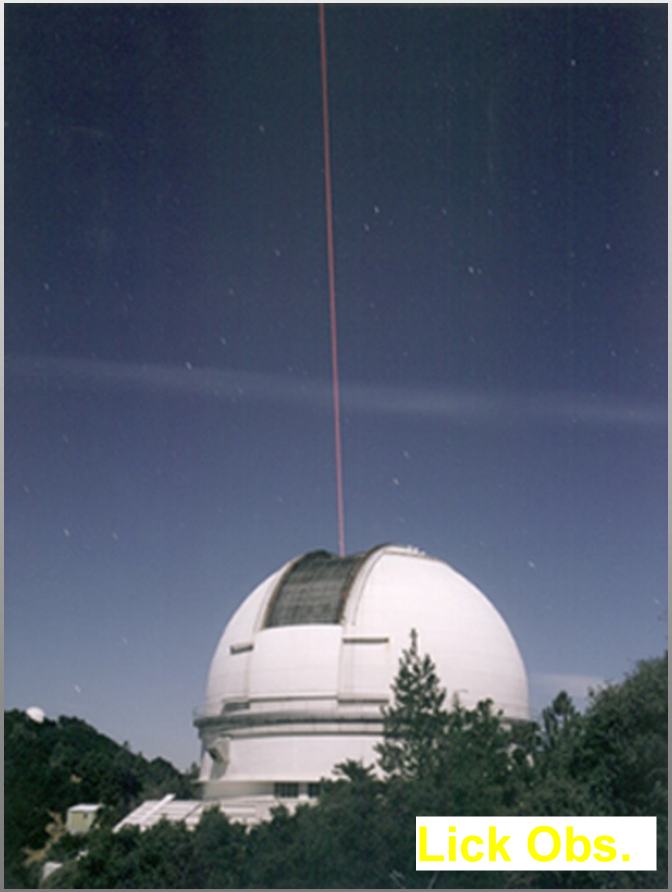

Anatomy of a Laser Guide Star

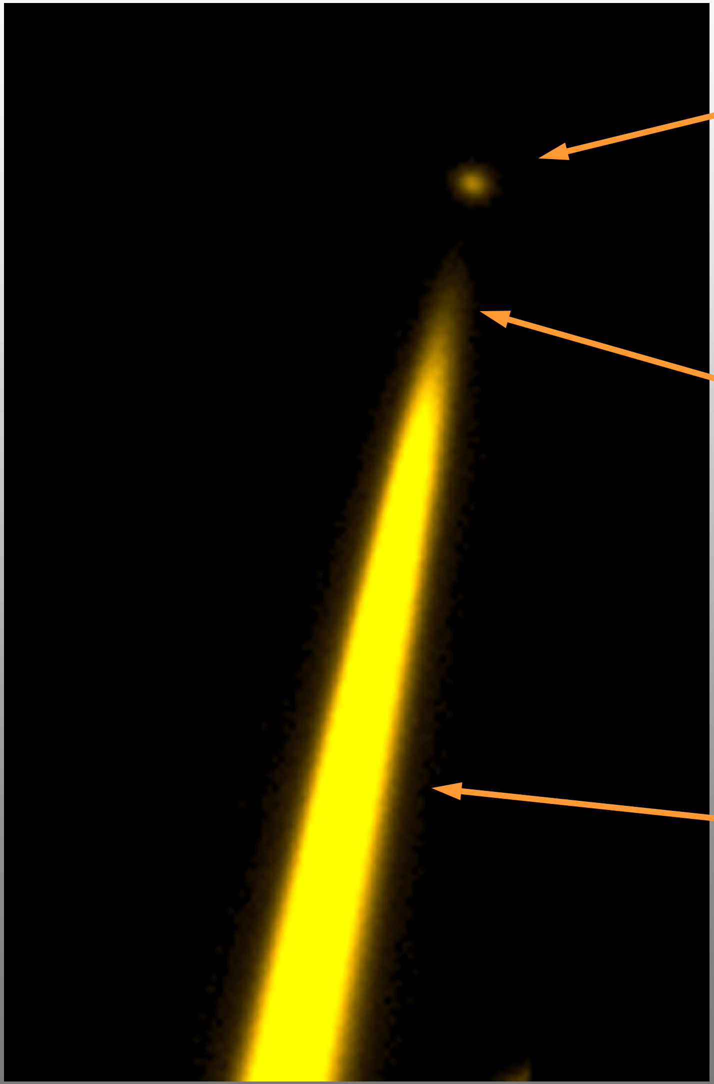

The Guide Star:
Fluorescent scattering by the mesospheric Sodium layer at ~ 95 km

Maximum altitude of (unwanted) backscatter from the air ~ 35 km

Back scatter from air molecules

Wavefront phase is corrected with a deformable mirror

MEMS deformable mirror with electrostatic actuators

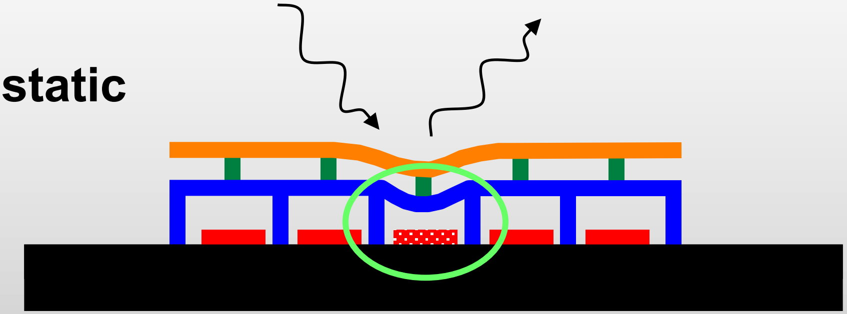

Simplified actuator model:

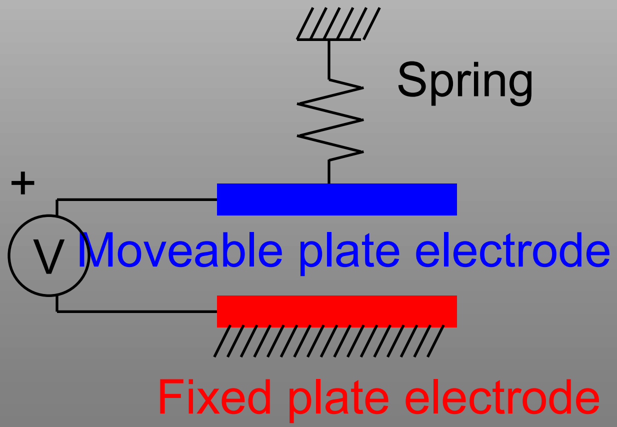

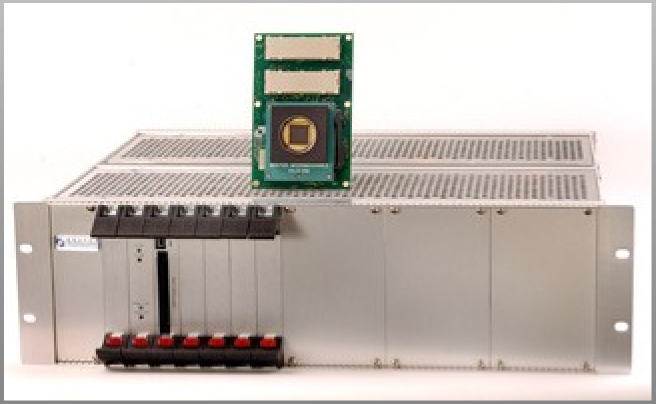

Diagrams and photo courtesy Boston Micromachines Corporation

Adaptive Optics control needs are expanding

- Larger telescopes
  - Spatial sampling set by the atmosphere -> number of samples grows with D^2
  - D=10 meter today, D=30-40 meter within the next decade
- Shorter wavelength science bands (moving from IR to Visible λ)
  - More precise correction needed (fraction of λ)
  - More samples, both spatial and temporal
- Wider field of view
  - Multiple laser guide stars – Tomography
- All of this points to higher speed computation on increased amounts of data

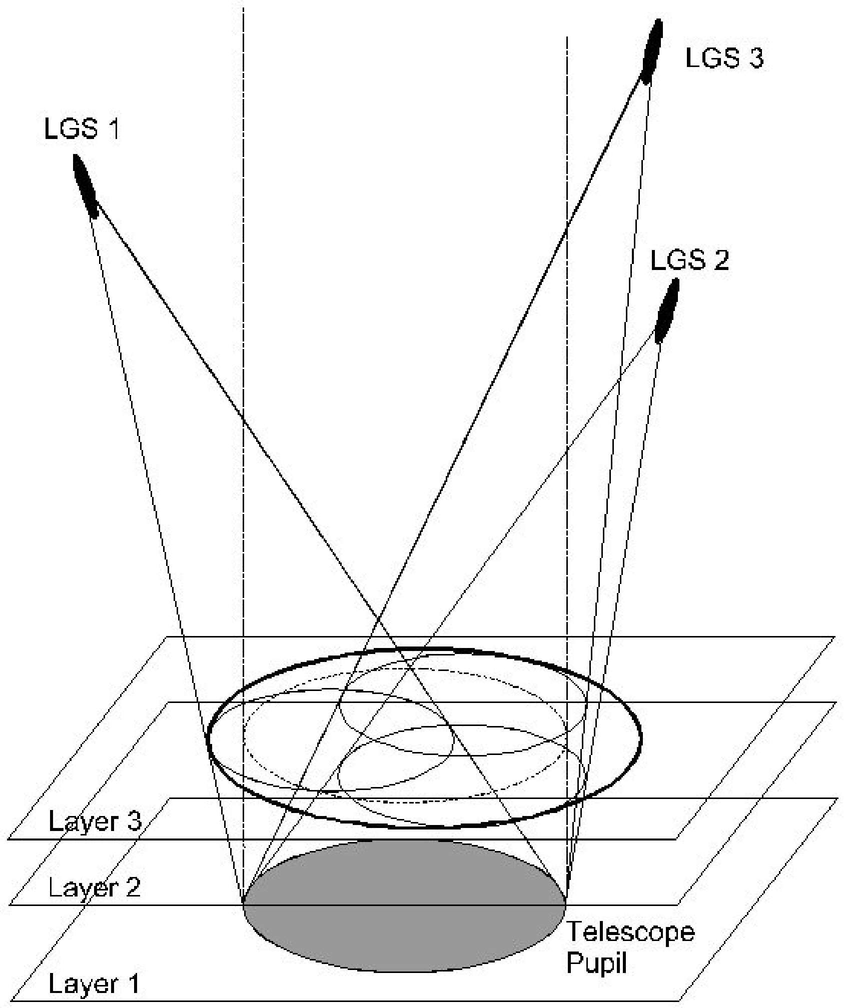

Tomographic Wavefront Reconstruction:
a quick summary

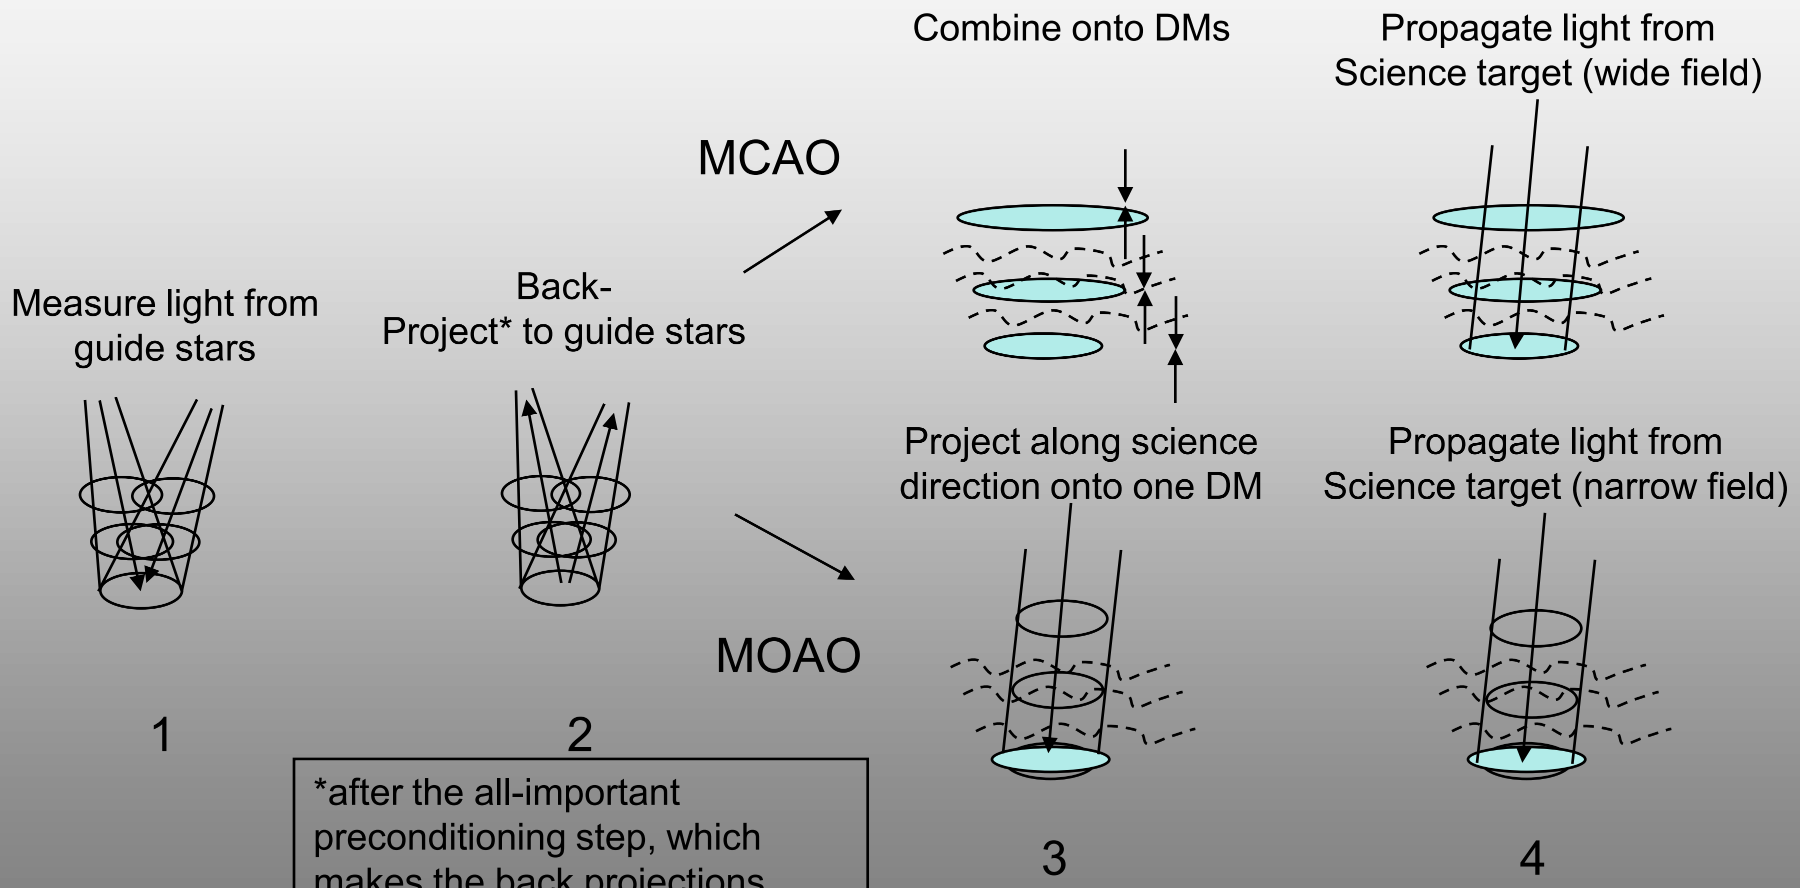

Measure light from guide stars

1

Back-Project* to guide stars

2

MCAO

Combine onto DMs

Propagate light from Science target (wide field)

MOAO

Project along science direction onto one DM

3

Propagate light from Science target (narrow field)

4

*after the all-important preconditioning step, which makes the back projections consistent with all the data

Tomography AO control architecture is a mixture of pipelined and massively-parallel elements

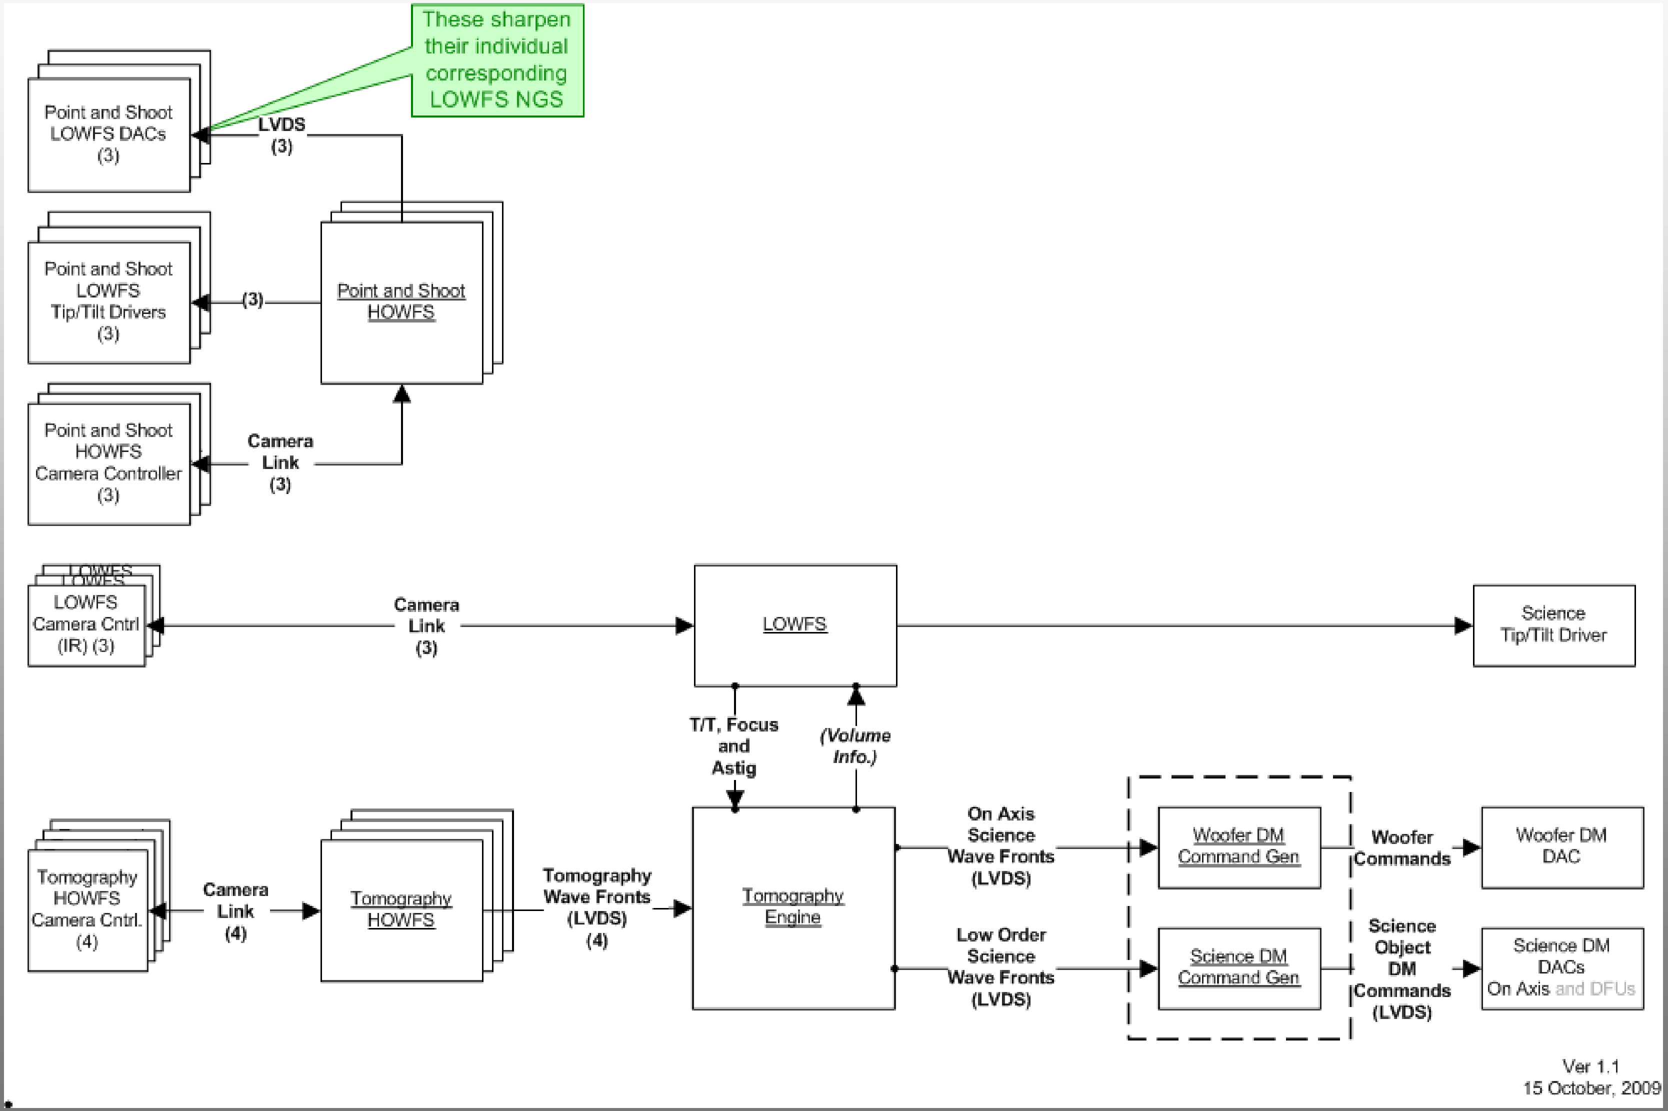

These sharpen their individual corresponding LOWFS NGS

Ver 1.1
15 October, 2009

The tomography engine processor array maps to the atmospheric volume

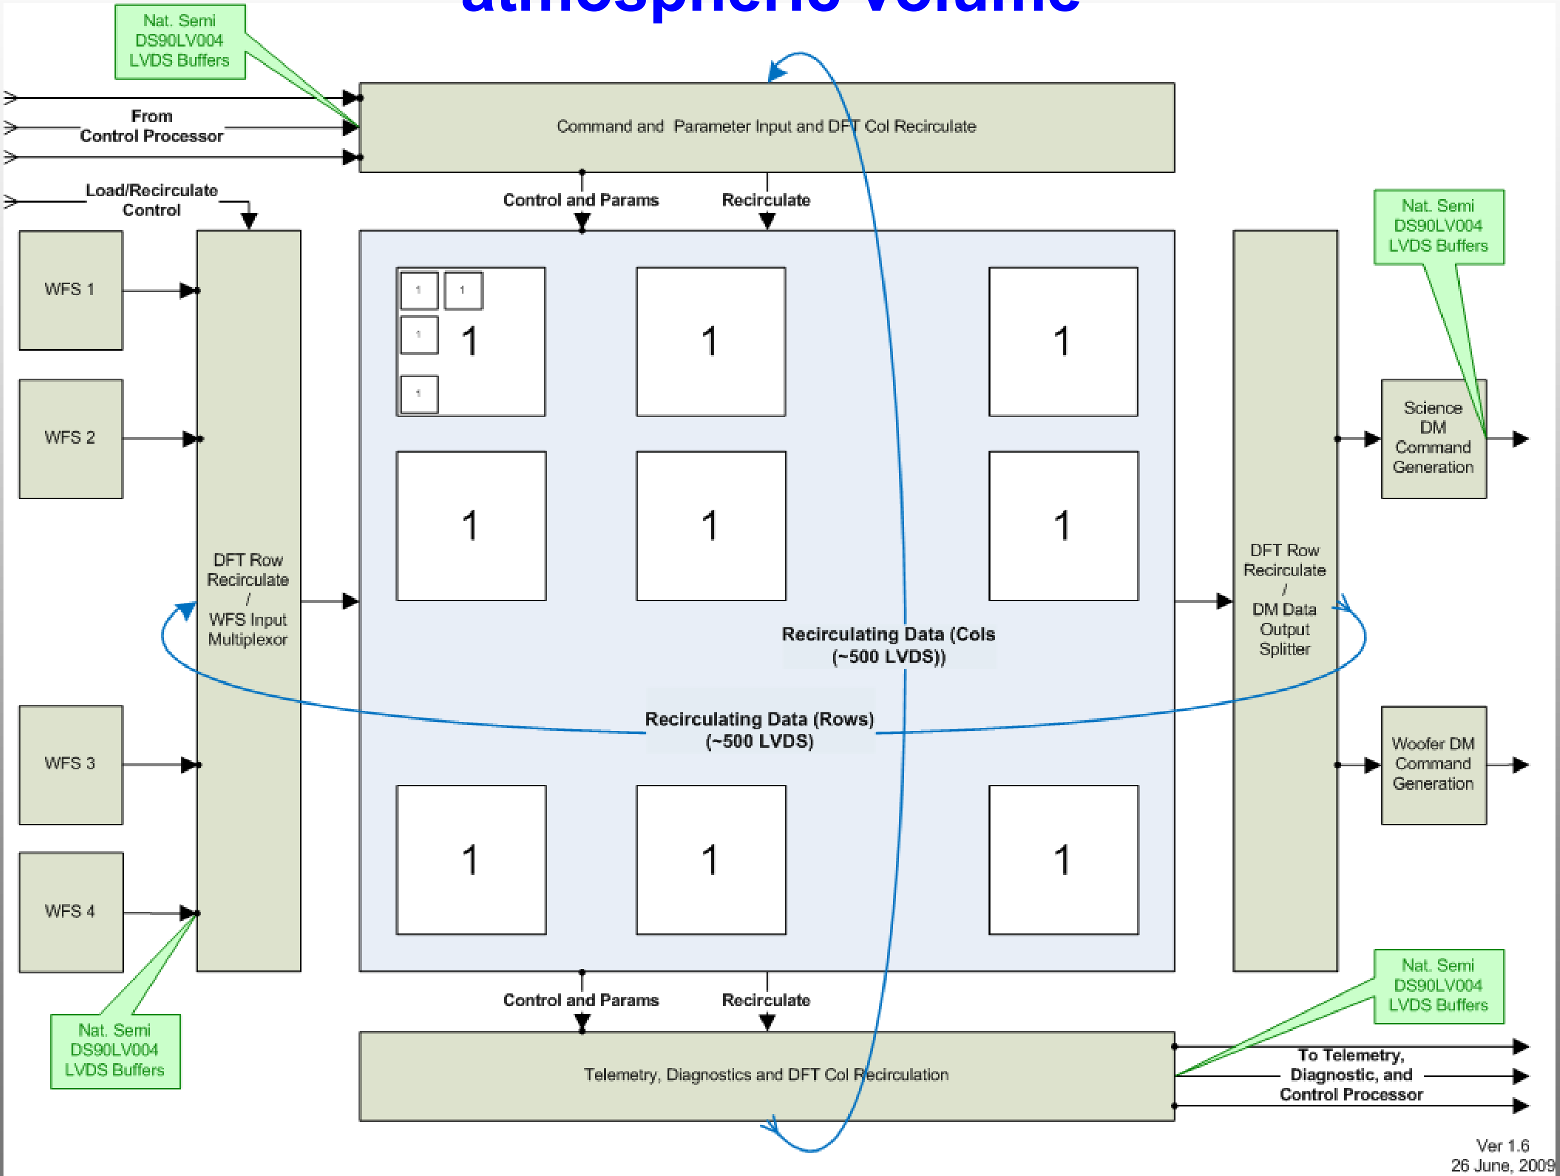

Details of processing element connectivity

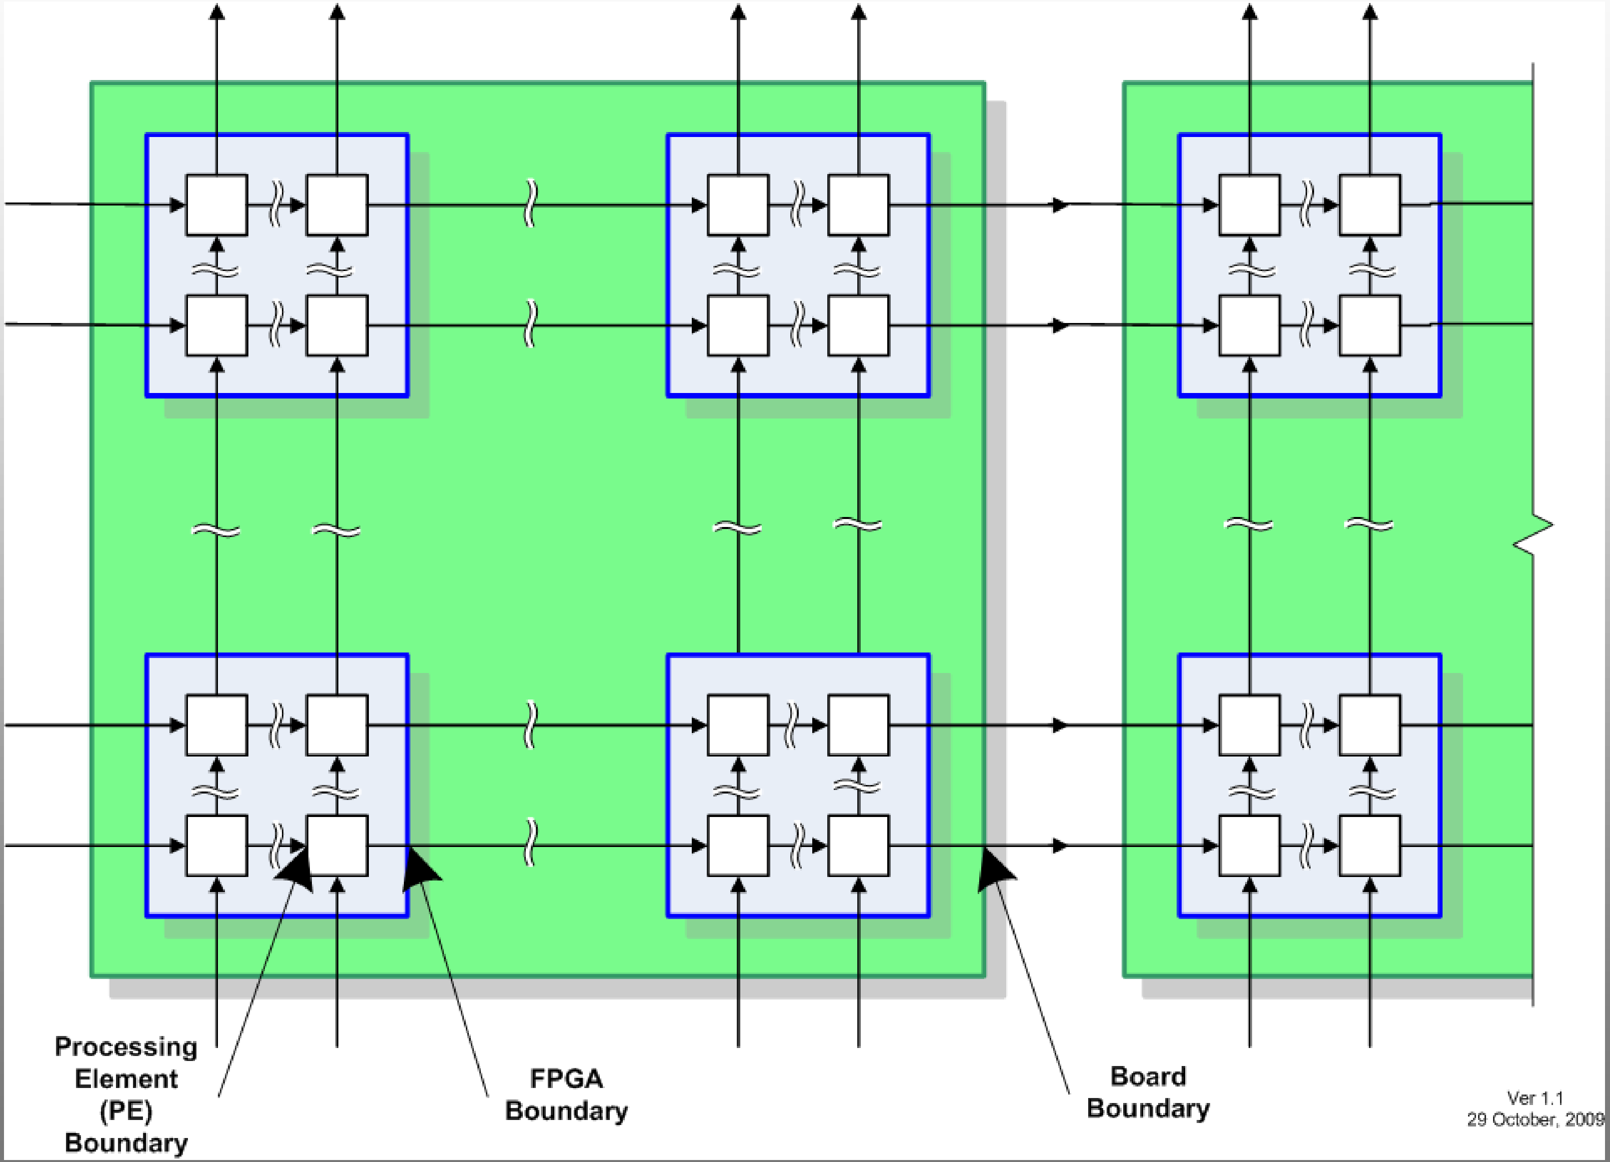

Lots of speedup from parallelization, but serial steps demand low communication latency

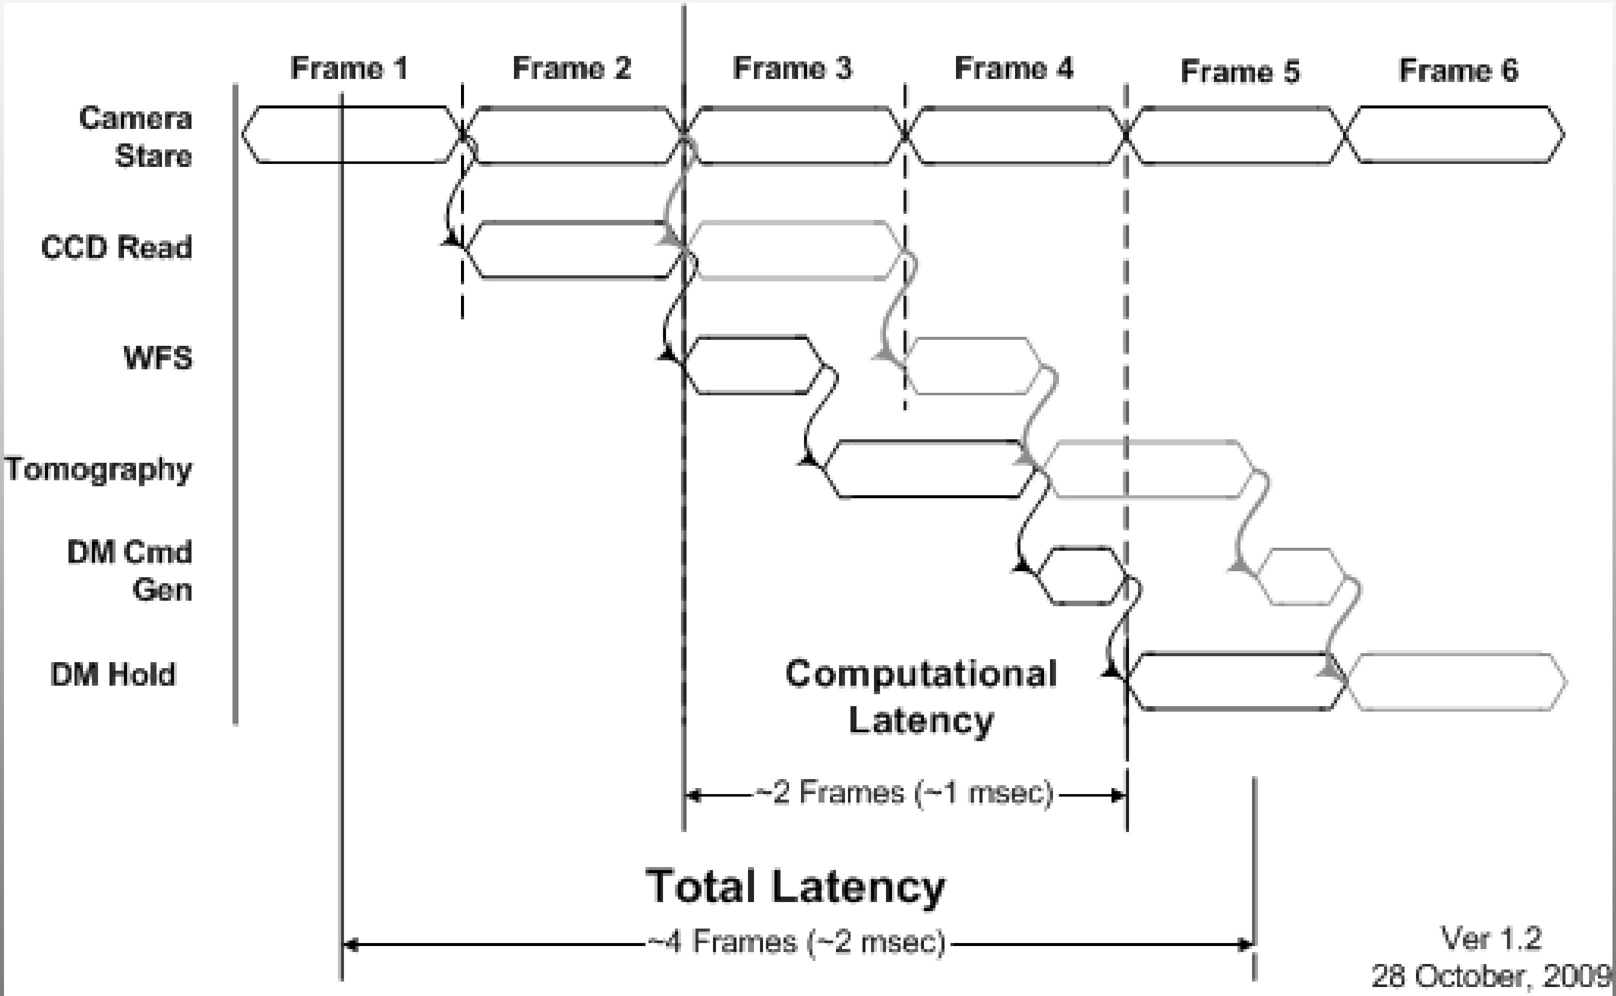

At various stages in the algorithm, compute elements represent:

- 3-D spatial sample points in the atmospheric volume
- 2-D spatial sample points on the aperture associated with each wavefront sensor
- 2-D spatial sample points on the aperture associated with each deformable mirror
- 2-D Fourier domain sample points in each layer of atmosphere
- 2-D Fourier domain sample points on the aperture associated with each wavefront sensor
- 2-D Fourier domain sample points on the aperture associated with each deformable mirror

The 3-D systolic array performs elemental operations:

- Lateral distortion-correct (“stencil”)
- Lateral shift and scale
- Z sum (forward propagation)
- Z distribute (back propagation)
- Filtering (massively parallel in Fourier domain)
- Masking (massively parallel in spatial domain)
- Fourier transform

A single processing element (FPGA architecture)

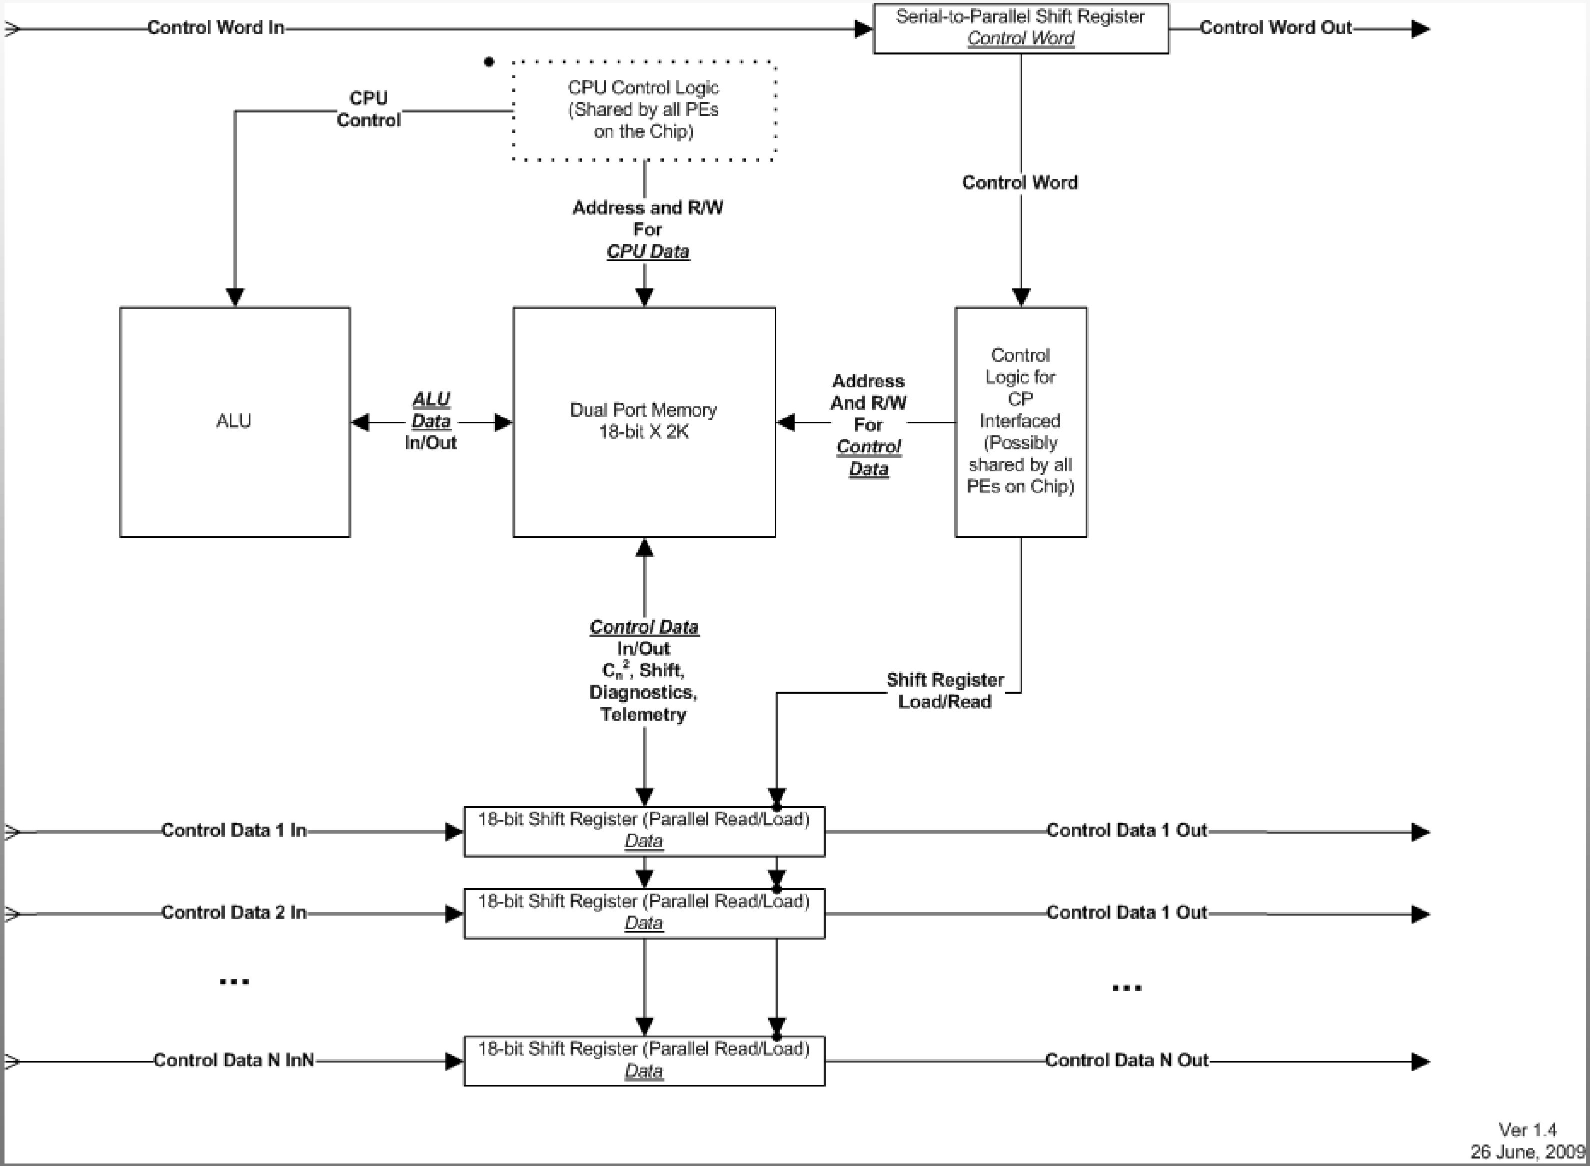

Control Word In

Serial-to-Parallel Shift Register
Control Word

Control Word Out

CPU Control Logic
(Shared by all PEs
on the Chip)

CPU
Control

Address and R/W
For
CPU Data

Control Word

ALU

ALU
Data
In/Out

Dual Port Memory
18-bit X 2K

Address
And R/W
For
Control
Data

Control
Logic for
CP
Interfaced
(Possibly
shared by all
PEs on Chip)

Control Data
In/Out
C_n^2, Shift,
Diagnostics,
Telemetry

Shift Register
Load/Read

Control Data 1 In
18-bit Shift Register (Parallel Read/Load)
Data
Control Data 1 Out

Control Data 2 In
18-bit Shift Register (Parallel Read/Load)
Data
Control Data 1 Out

...

...

Control Data N InN
18-bit Shift Register (Parallel Read/Load)
Data
Control Data N Out

Ver 1.4
26 June, 2009

Key issues limit scalability

- Low latency: data in to data out time has direct impact on AO performance
    - Processors are I/O bound (both FPGA and GPU) – data transmission is as expensive as data computation
- Fourier transforms:
    - Essential to the AO tomography algorithm
    - Dominant source of computational delay
    - Fast in the GPU, but I/O bound (starved pipeline) and favors larger arrays than used in AO
    - FFT (Cooley-Tukey) is not the fastest implementation when distributed on multi processing units.
    - Fastest DFT is O(N) rather than O(1) – forcing increase in communication and processing speed $\propto D/\lambda$

Calculating the DFT with an array of compute elements

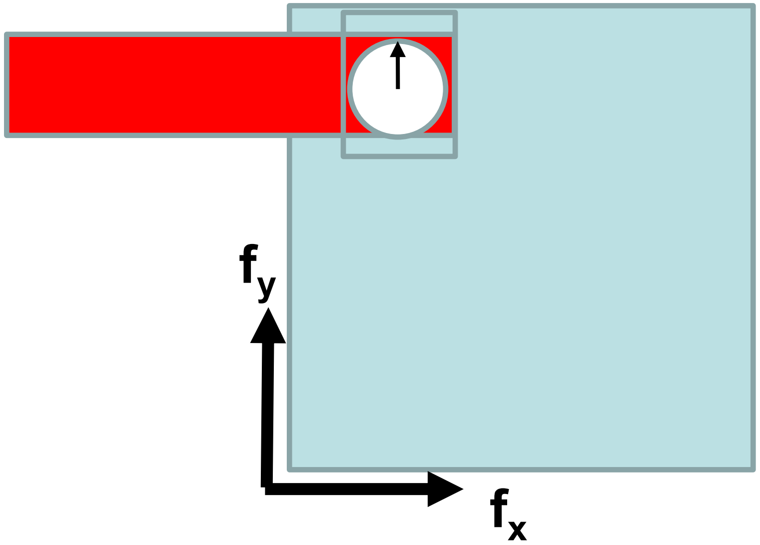

$$\Phi_1(f_x,y)=\Sigma_k\ \phi(x,y)\ \exp\{-2\pi i f_x x_k\}$$

$$\Phi(f_x,f_y)=\Sigma_k\ \Phi_1(f_x,y)\ \exp\{-2\pi i f_y y_k\}$$

○ = $\exp\{2\pi i f x\}$

🟥 = row of data $\phi(x)$

☐ = compute element

Calculating the DFT with an array of compute elements

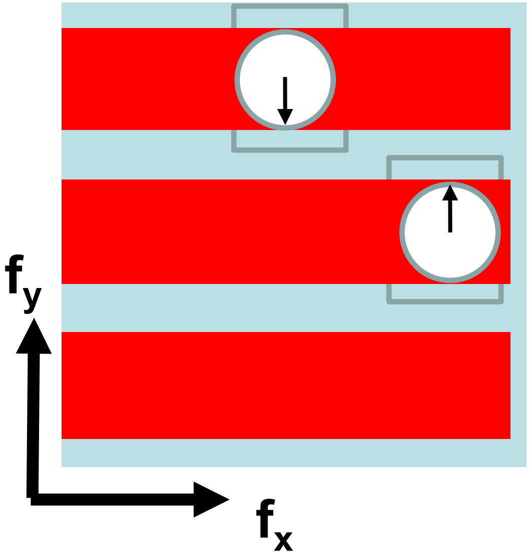

$$\Phi_1(f_x,y)=\Sigma_k\ \phi(x,y)\ \exp\{-2\pi i f_x x_k\}$$

$$\Phi(f_x,f_y)=\Sigma_k\ \Phi_1(f_x,y)\ \exp\{-2\pi i f_y y_k\}$$

○ = $\exp\{2\pi i f x\}$

🟥 = row of data $\phi(x,y)$

☐ = compute element

Calculating the DFT with an array of compute elements

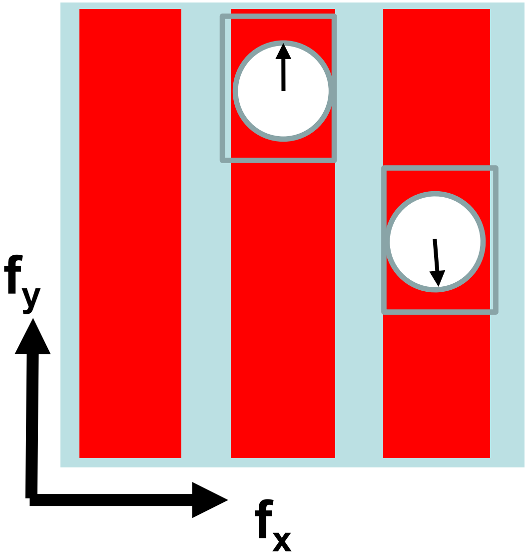

$$\Phi_1(f_x,y)=\Sigma_k\ \phi(x,y)\ \exp\{-2\pi i f_x x_k\}$$

$$\Phi(f_x,f_y)=\Sigma_k\ \Phi_1(f_x,y)\ \exp\{-2\pi i f_y y_k\}$$

◯ = exp{2πi*fx*}

🟥 = column of data $\phi(f_x,y)$

☐ = compute element

Conclusions

- AO real-time processors are transitioning from fast single CPU solutions to the massively-parallel domain
- Key AO multi-processor architecture needs are not a clean match to the market driven needs
- Even with massive parallelization, the AO algorithm (as we now understand it) is not O(N) speed-up – and so is not sustainable with increasing D/λ

Acknowledgement: This work was supported by the Gordon and Betty Moore Foundation, Keck Observatory, and the National Science Foundation in support the Keck Observatory Next Generation Adaptive Optics Program