Optical interferometry in practice

ONTHEFRINGE School

Astrometry and Imaging with the VLTI

Keszthely, Hungary,
2nd June – 13th June 2008

Chris Haniff

Astrophysics Group, Cavendish Laboratory,
University of Cambridge, UK

3rd June 2008

This is a “school” – please ask questions

Asking questions doesn’t mean you are stupid, it means you are open to learning new things.

Yesterday we talked about "the best of all possible worlds"

- Telescopes sample the fields at r_1 and r_2.
- Optical train delivers the radiation to a laboratory.
- Delay lines assure that we measure when $t_1=t_2$.
- The instruments mix the beams and detect the fringes.

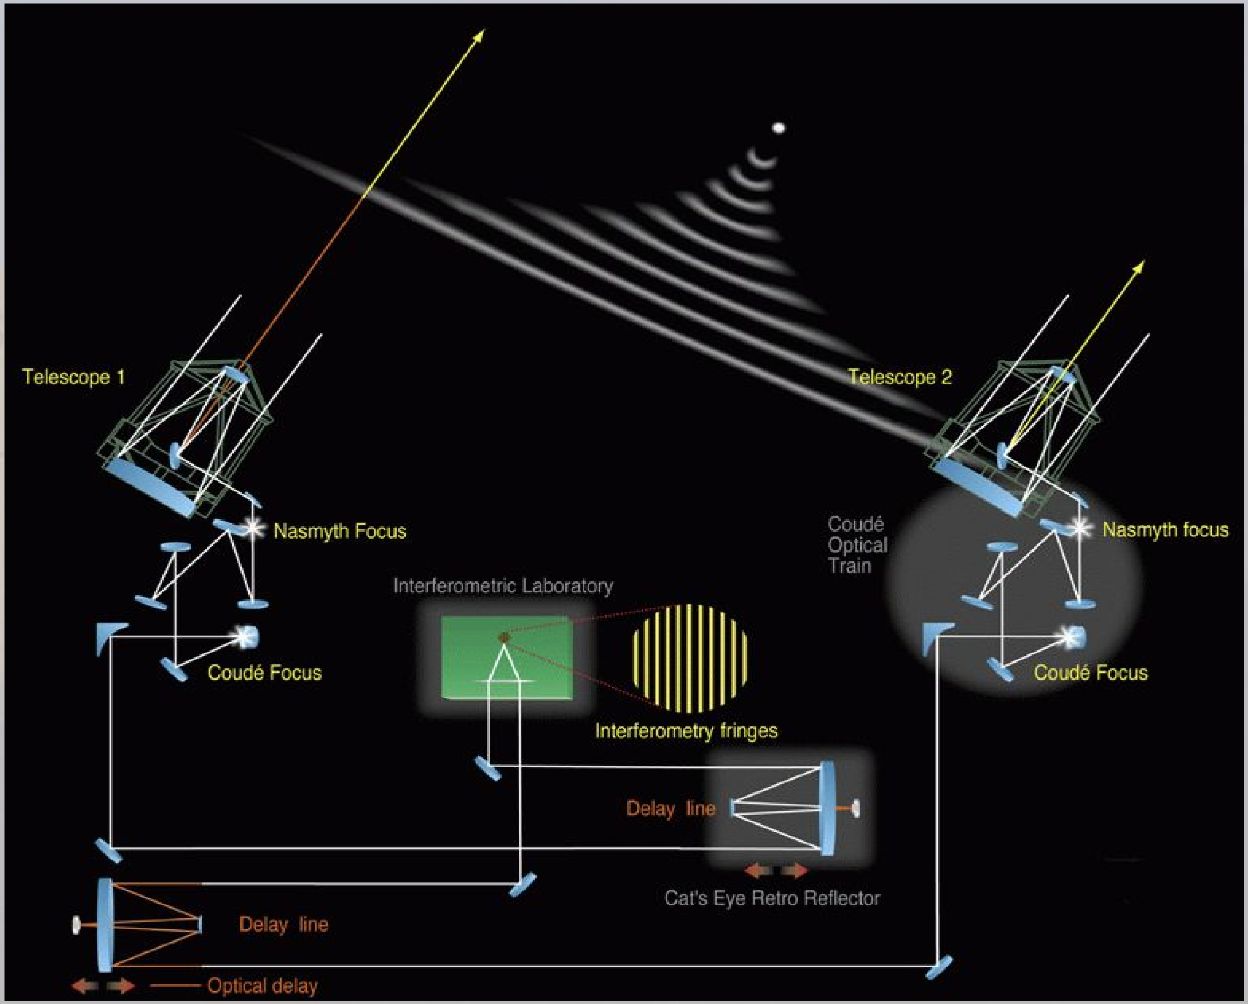

And we learnt what – in principle – interferometry involves

- We choose to measure the Fourier transform of the source rather than the brightness distribution.
- We use an interferometer to make sequential measurements of the Fourier transform.
- Each complex measurement is encoded in the "fringe" pattern seen in the interferometric output.

Detected intensity, $I = \langle \Psi \Psi^* \rangle \propto 2 + 2\cos(k\,[\hat{s}.B + d_1 - d_2])$

$\propto 2 \times [1+\cos(kD)]$, where $D = [\hat{s}.B + d_1 - d_2]$

- We take these measurements of the FT of the source brightness – this is the so-called visibility function – and we inverse Fourier transform them.
- This recovers the true source brightness distribution convolved with the interferometric PSF.
- We deconvolve the PSF and get our final map of the sky.

Today we will talk about the real world

How is this "perfect" process constrained by the real world?

What are the strategies we can use to help overcome these hurdles to success?

Outline

- Sampling of the *(u, v)* plane
- Beam relay
- Delay compensation
- Beam combination

Hardware

- Spatial wavefront fluctuations
- Temporal wavefront fluctuations

Atmospheric

- Sensitivity
- Calibration

Fundamental

How easy is it to sample the *uv* plane?

- Sampling of the *(u, v)* plane
- Beam relay
- Delay compensation
- Beam combination
- Spatial wavefront fluctuations
- Temporal wavefront fluctuations
- Sensitivity
- Calibration

We know where we want to measure – how easy is it to realise this?

What do the VLTI telescope locations look like?

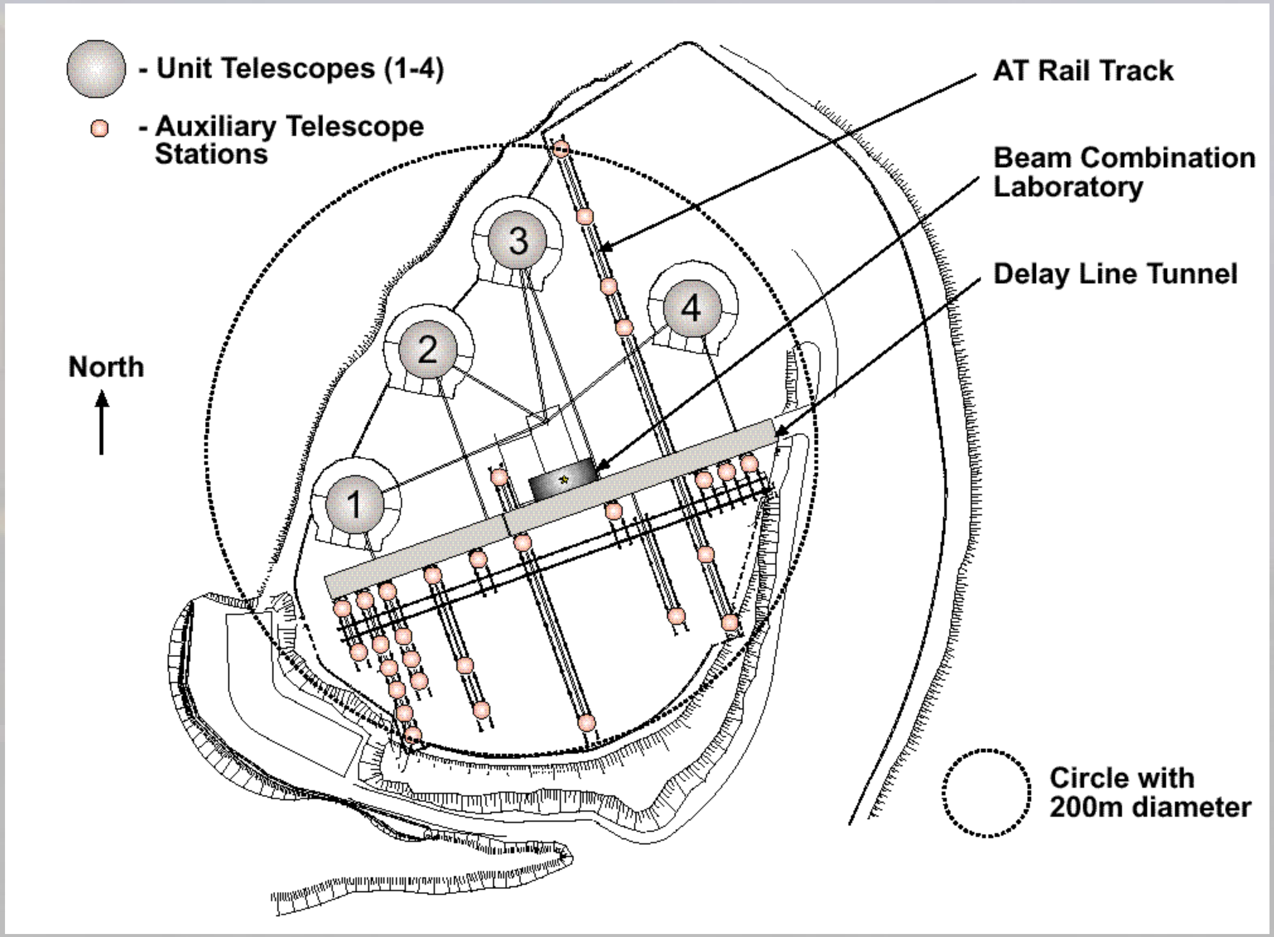

How are those locations related to the *uv* coverage?

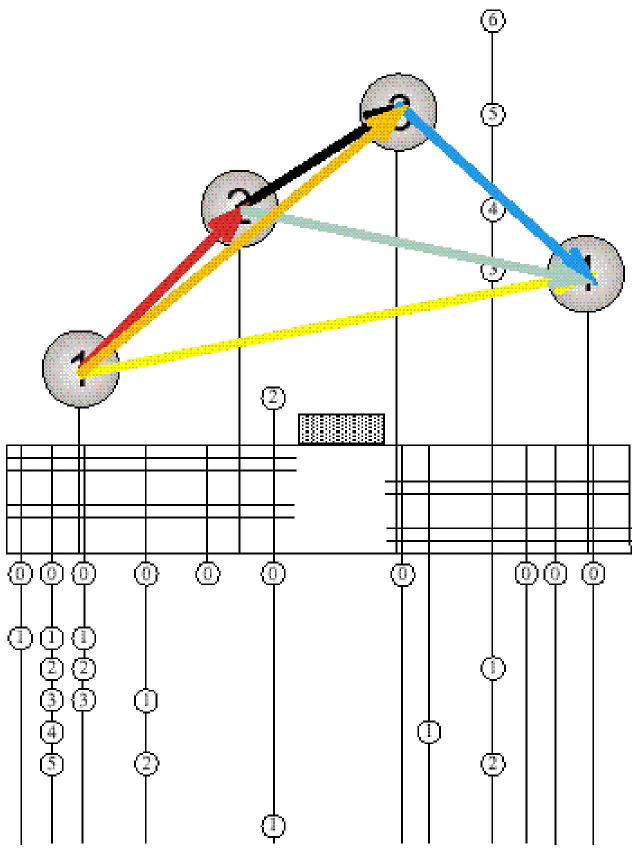

This is the uv-plane:

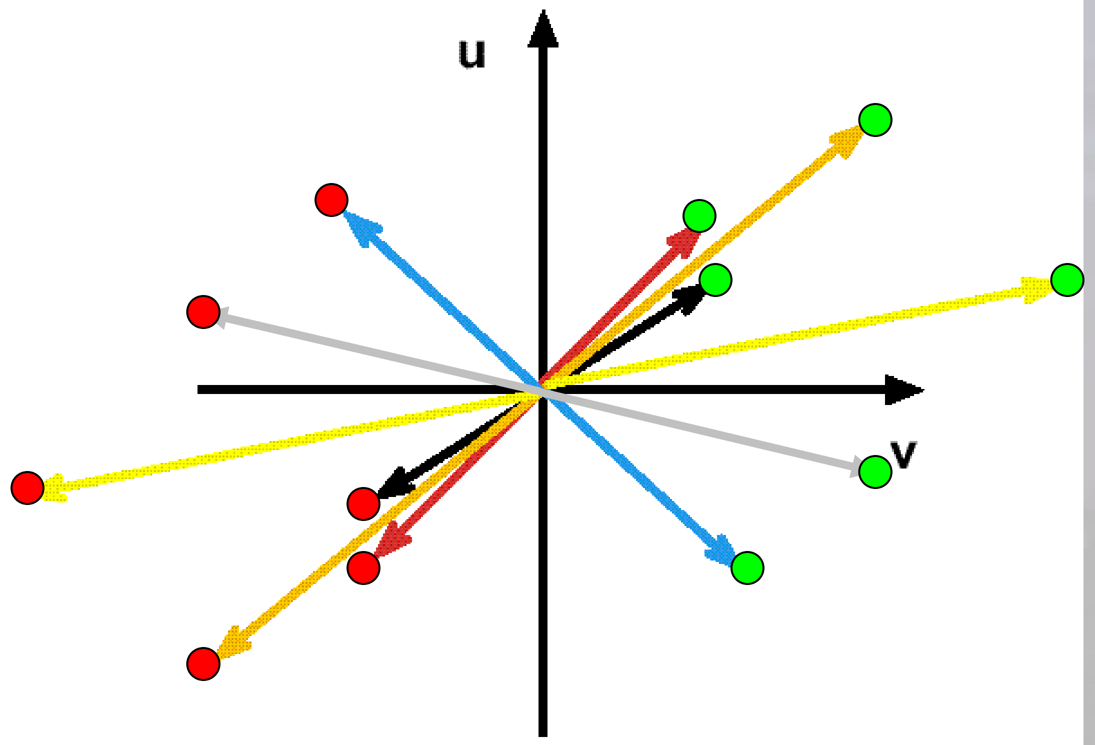

Note: This is the uv-plane for an object at zenith.
In general, the projected baselines have to be used.

OK – but what really matters?

- Rather than re-locate the telescopes to measure different spatial frequencies, we take advantage of the Earth's rotation. In this case the tips of the *uv* vectors sweep out ellipses.

- The properties of these tracks in the Fourier plane will be governed by:
  - The hour angles of the observation.
  - The declination of the source.
  - The stations being used.

- We also need to worry about all of the following:
  - Is there any shadowing of the telescopes by each other?
  - The allowed range of the delay lines - are they long enough?
  - The zenith distance - will the seeing be too poor at low elevations?
  - Can the interferometer fringe-track ok?

Examples of Fourier plane coverage

Dec -15

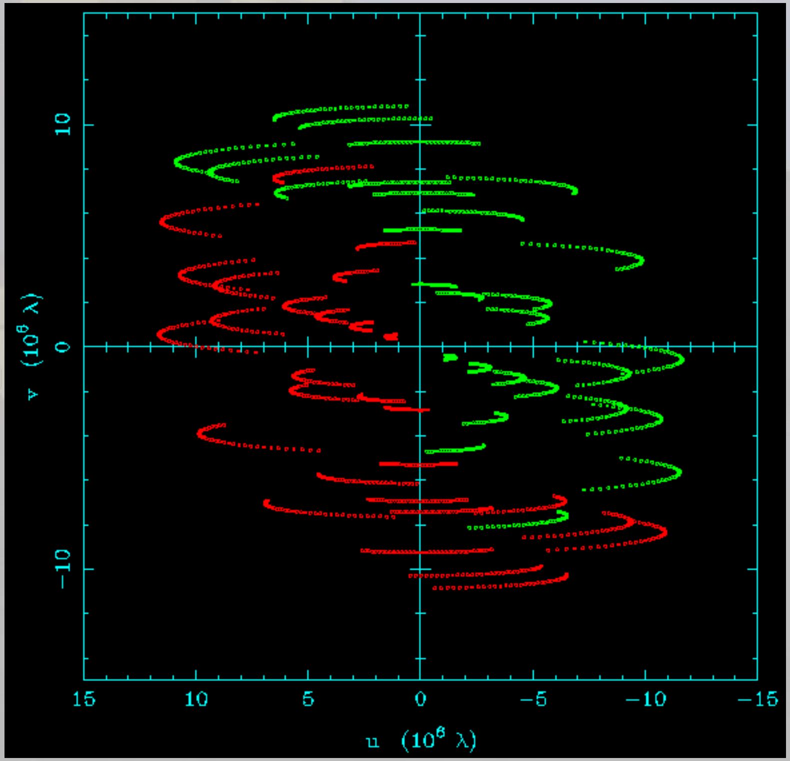

Dec -65

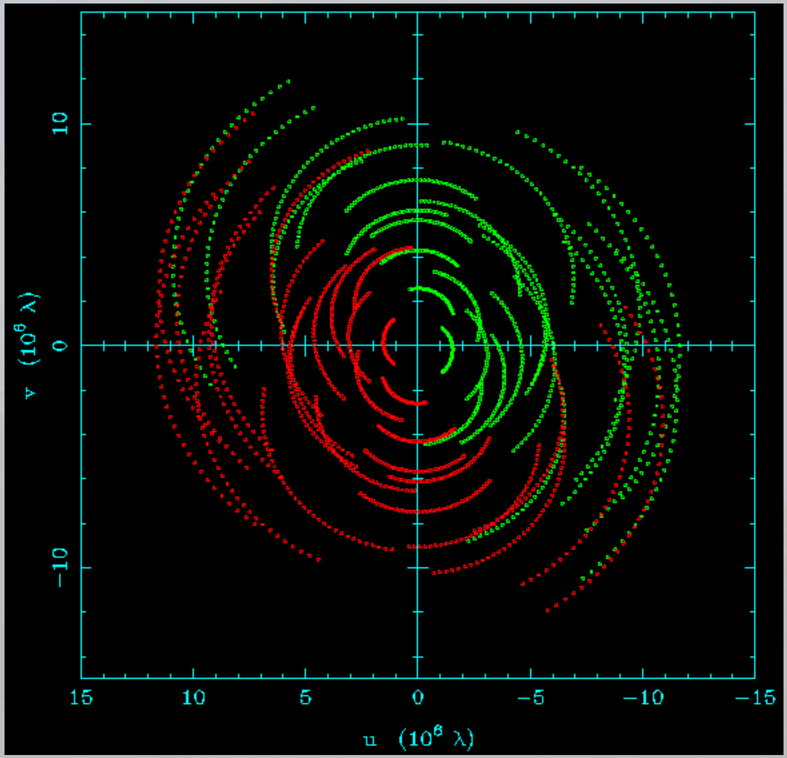

Whatever these look like, don’t forget the rules of thumb we learnt yesterday!

How does the uv-plane coverage impact imaging?

Model

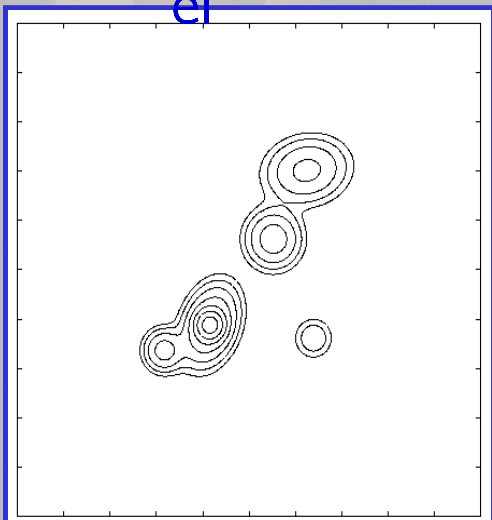

4 telescopes, 6 hours

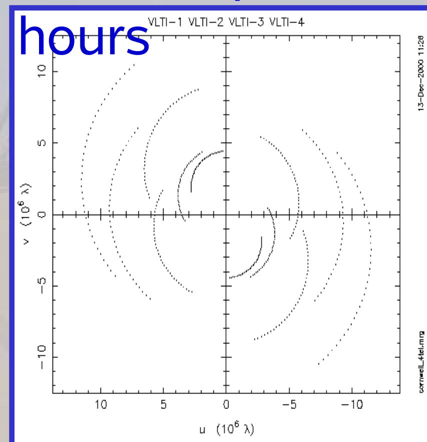

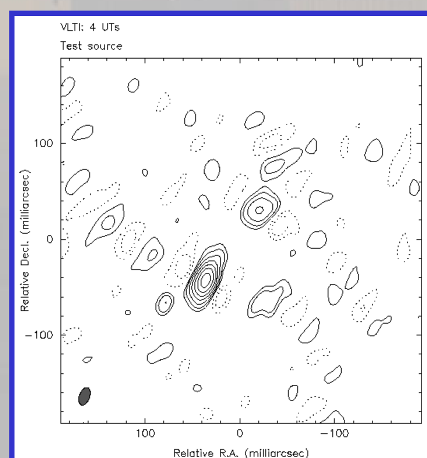

8 telescopes, 6 hours

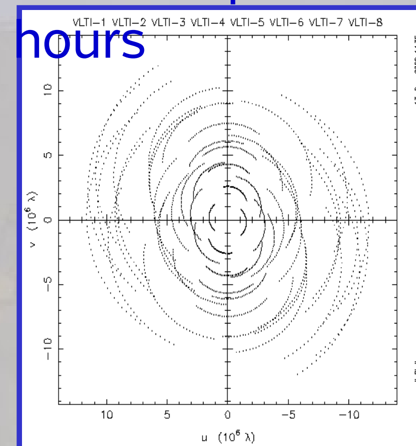

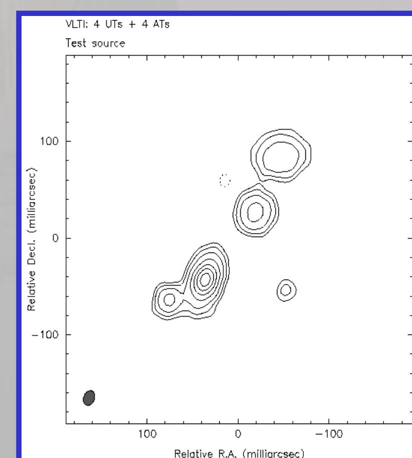

Recap & questions

We may know where we want to sample in the Fourier (*uv*) plane, but unless we can relocate our telescopes at will, this will always be a challenge.

Fixed telescopes can be a significant penalty – in many cases the physical baselines we have may be too long.

If imaging is the goal, then uneven sampling of the *uv* plane is not easily compensated for after the experiment.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the *(u, v)* plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

Multiple wavelengths

Polarization

How do we get the light from the telescopes?

- At radio wavelengths coherent waveguides are often used.

- In the optical/IR we can emulate this in two ways:
  - Free-space propagation in air/vacuum.
  - Guided propagation in an optically denser medium, e.g. using an optical fibre.

Whatever we do – we need to address these issues

- Longitudinal dispersion i.e. the possibility of a mismatch in the optical paths in each interferometer arm due to the wavelength-dependence of the refractive index of the propagating medium.

- The polarization state of the propagating beams: this needs to be matched if the beams are to interfere successfully.

- How can a range of wavelengths be transported in parallel?

- How lossy is the propagation overall?

Current preferences

- Currently, the most favoured approach has been to use free-space propagation in air-filled or evacuated pipes:
  - Few problems exist with longitudinal dispersion and turbulence if the relay pipes are evacuated.
  - For air-filled pipes a small beam diameter can help to limit wavefront fluctuations.

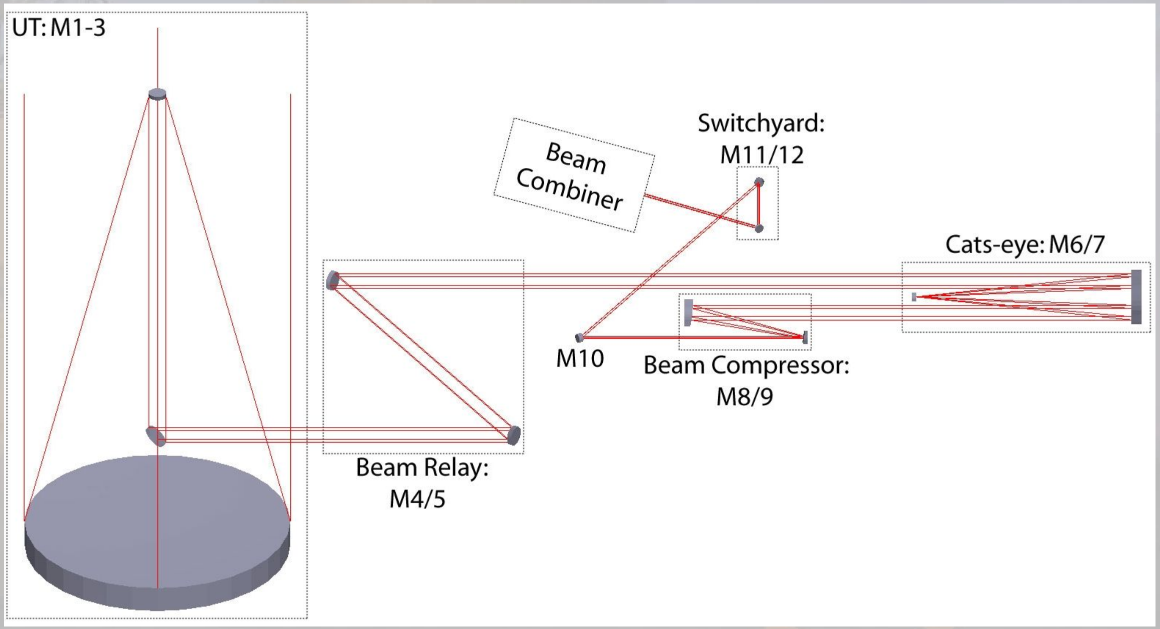

  - Generally a beam diameter $>(\lambda z)^{1/2}$ is used, where z is the propagation length, to as to mimimize diffraction losses.

A question for the "experts"

- When thinking about polarization, in general, there are two issues to deal with.
  1. Can we control the polarization state sufficiently to interfere the beams from the telescopes?
  2. Can we control the polarization state so as to make images of the sky in any polarization state we desire?

  In most optical/IR implementations, problem (1) is always addressed, whereas problem (2) is generally left for "future generations" to think about!

- The key issue to take home is:
  - Non-normal reflections will lead to differential amplitude and phase changes to be introduced into orthogonal polarization states.
    - As long as these are "identical" in each interferometer arm, the beams will interfere

Quiz: is this a sufficient condition to address issue 2?

Recap & questions

Getting the light from the telescopes to the beam combining laboratory is a challenging implementation issue.

Few interferometers have done this “optimally”.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the *(u, v)* plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

We know what they are supposed to do, but what makes this difficult?

Let's remind ourselves what the task is

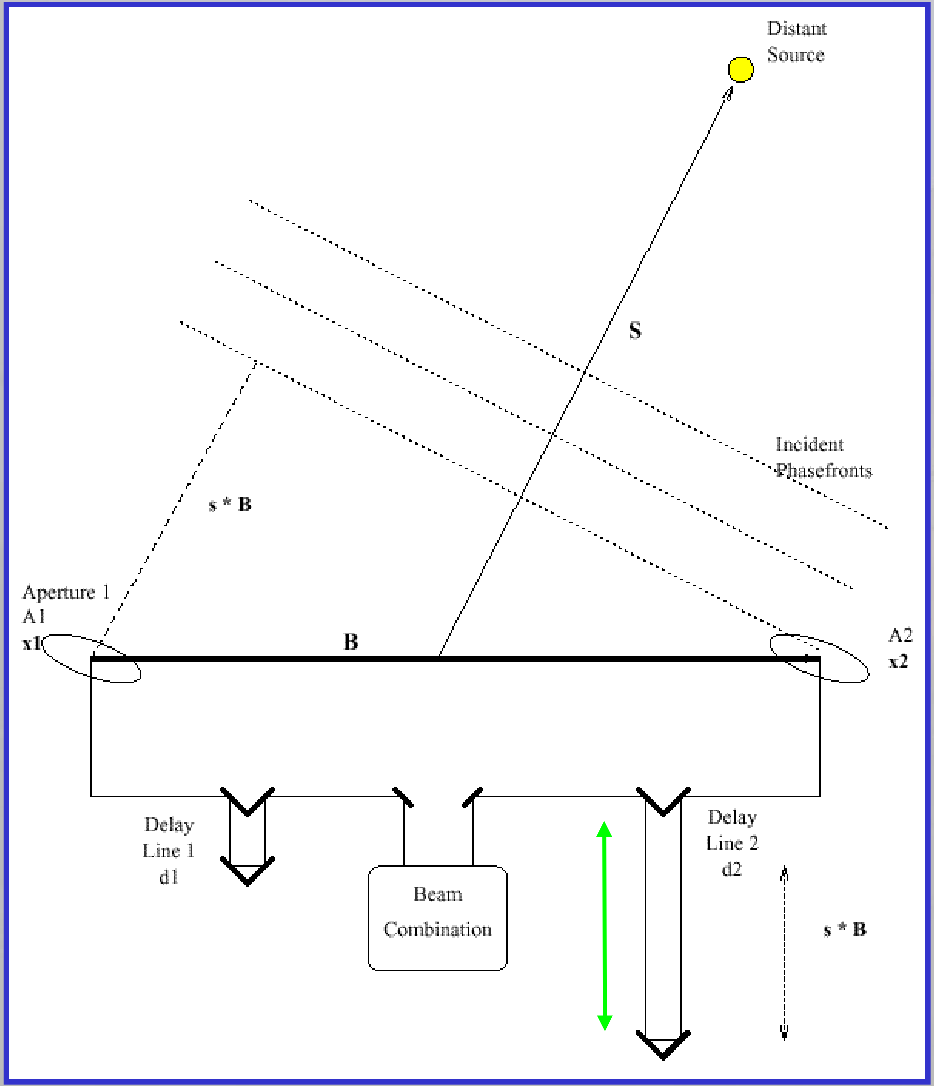

- We need to correct for the potentially very large geometric delay.
- We need to accommodate the time varying nature of the geometric delay.
- We have a large diameter beam of light to deal with.

What does a delay line look like?

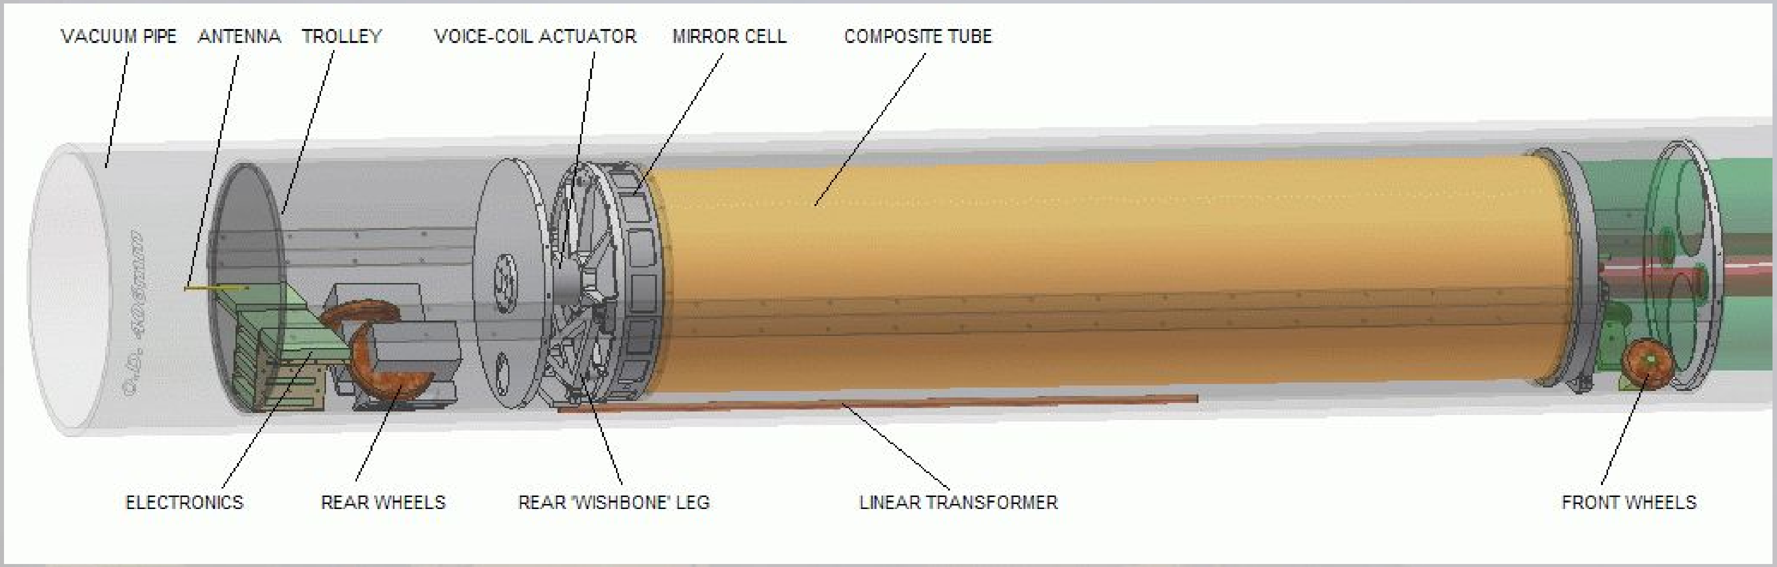

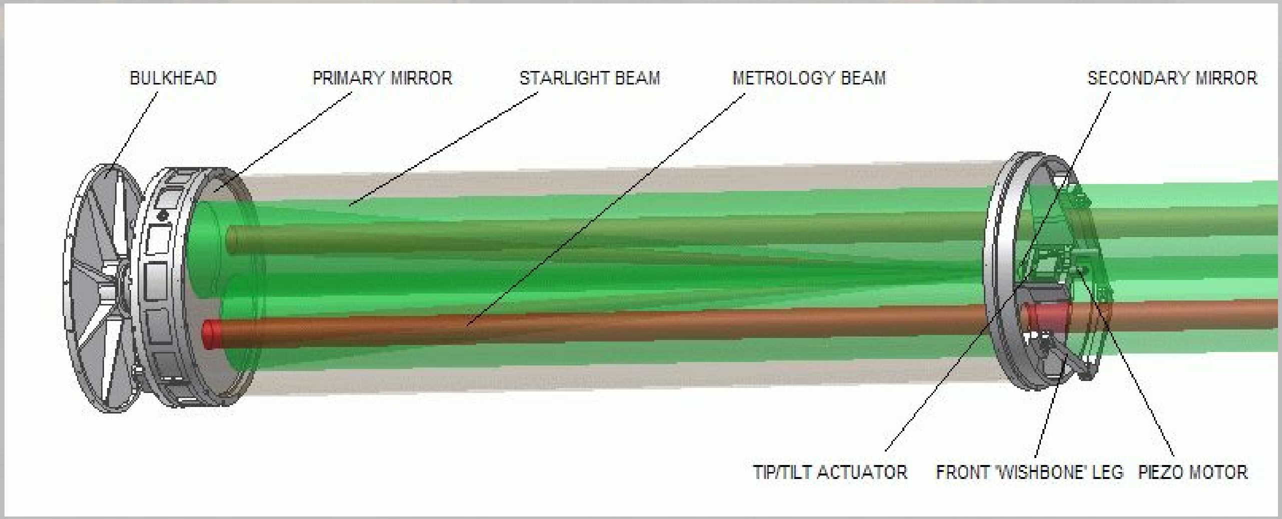

So – how well do delay lines have to perform?

- The OPD added can be as large as the maximum baseline:
  - VLTI has opd_{max}~120m.
- The OPD correction varies roughly as B cos(θ) dθ/dt, with θ the zenith angle.
  - VLTI has v_{max}~ 0.5cm/s (though the carriages can move much faster than this).
- The correction has to be better than $l_{coh} \sim \lambda^2/\Delta\lambda$.
  - A typical stability is ≤ 14nm rms over an integration time.

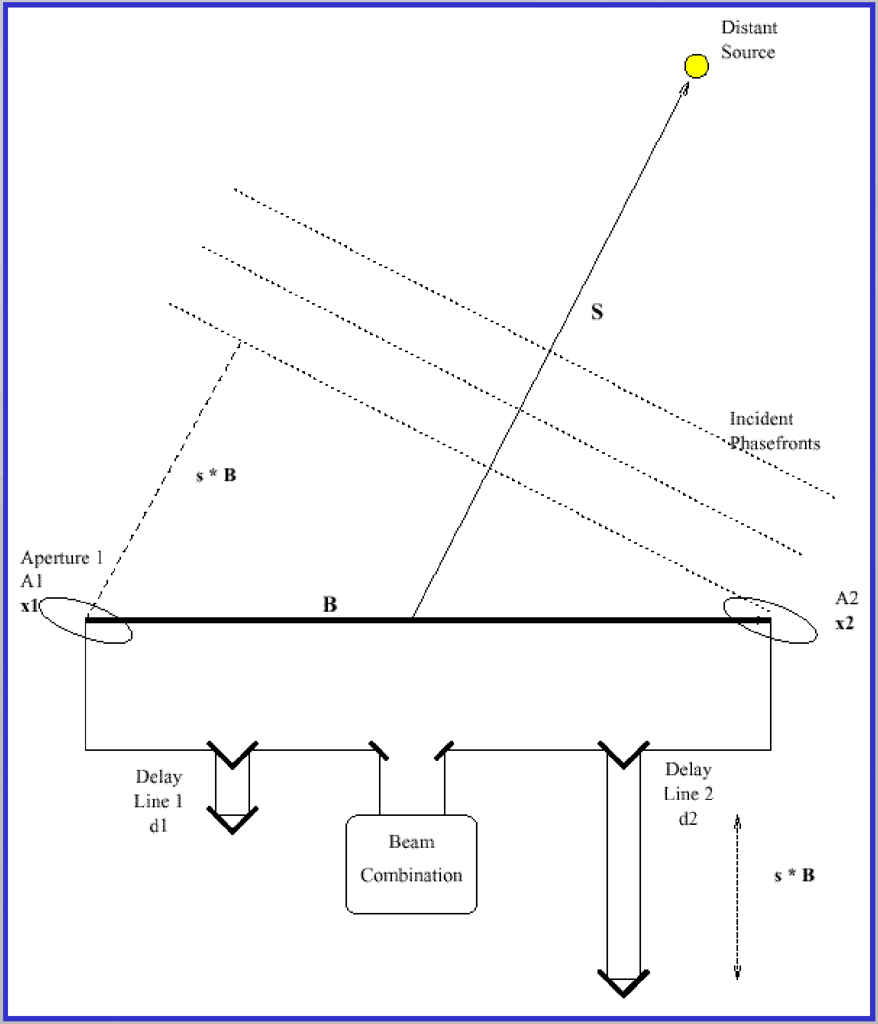

But what do the experts "hide" from you?

- Unless very specialized beam-combining optics are used it is only possible to correct the OPD for a single direction in the sky.
  – This gives rise to a FOV limitation: $\theta_{max} \leq [\lambda/B][\lambda/\Delta\lambda]$.

- For an optical train in air, the OPD is actually different for different wavelengths since the refractive index $n = n(\lambda)$.
  – This longitudinal dispersion implies that different locations of the delay line carts will be required to equalize the OPD at different wavelengths!
  – For a 100m baseline and a source 50^o from the zenith this ΔOPD corresponds to ~10μm between 2.0-2.5μm.
  – More precisely, this implies the use of a spectral resolution, $R > 5$ (12) to ensure good fringe contrast (>90%) in the K (J) band.

Recap & questions

Compensating for the geometric day is a challenging implementation issue.

All implementations usually involve making compromises in some areas of performance.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the (u, v) plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

Yesterday we talked about "adding the fields" – but how do we actually do this?

Lets start by recapping what we heard yesterday

- Add the E fields, E_1+E_2, and then detect the time averaged intensity:

$$\langle (E_1+E_2)\times(E_1+E_2)^* \rangle = \langle |E_1|^2 \rangle + \langle |E_2|^2 \rangle + \langle E_1E_2^* \rangle + \langle E_2E_1^* \rangle$$
$$= \langle\, |E_1|^2 \rangle + \langle\, |E_2|^2 \rangle + \langle\, 2|E_1||E_2| \cos(\varphi)\, \rangle\, ,$$

where φ is the phase difference between E_1 and E_2.

- In practice there are two straightforward ways of doing this:
  - Image plane combination:
    - AMBER and aperture masking experiments.
  - Pupil plane combination:
    - MIDI and systems using fibre couplers (VINCI).

Image plane combination

- Mix the signals in a focal plane as in a Young's slit experiment:

  - In the focused image the transverse co-ordinate measures the delay between the beams.

  - Fringes encoded by use of a non-redundant input pupil.

  - Possible to use dispersion prior to detection in the direction perpendicular to the fringes. Allows measurement of the visibility function at multiple λ.

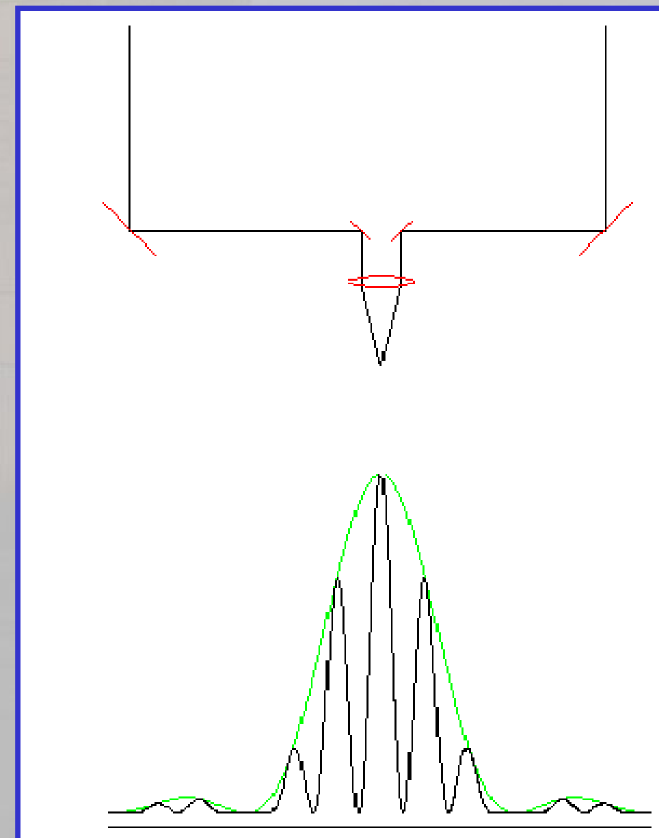

Pupil plane combination

- Mix the signals by superposing afocal (collimated) beams at a beam splitter plate:

  Then focus superposed beams onto a single element detector to measure I.

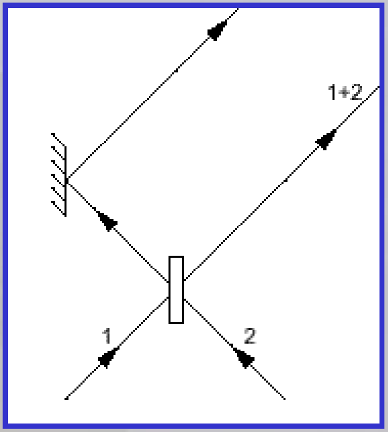

  - Fringes are visualised by measuring intensity versus time.
  - Fringes encoded by use of a non-redundant modulation of delay of each beam.
  - Can use spectral dispersion prior to detection to measure at multiple wavelengths.

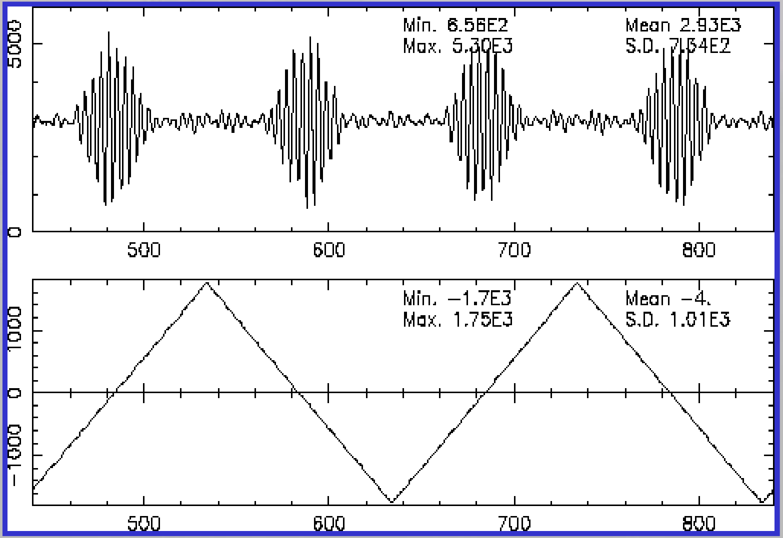

Recap & questions

We can either combine the beams in a focal plane of by superposing afocal beams.

In a focal plane combiner I(kD) is visualised for multiple D at the same time.

In a pupil plane combiner, I(kD) is sampled sequentially in time for different D.

More than two beams can be interfered on the same detector as long as the different fringe patterns can be distinguished without ambiguity.

Timeout

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the *(u, v)* plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

How does the atmosphere change what we described yesterday?

So, how can we understand what the atmosphere does?

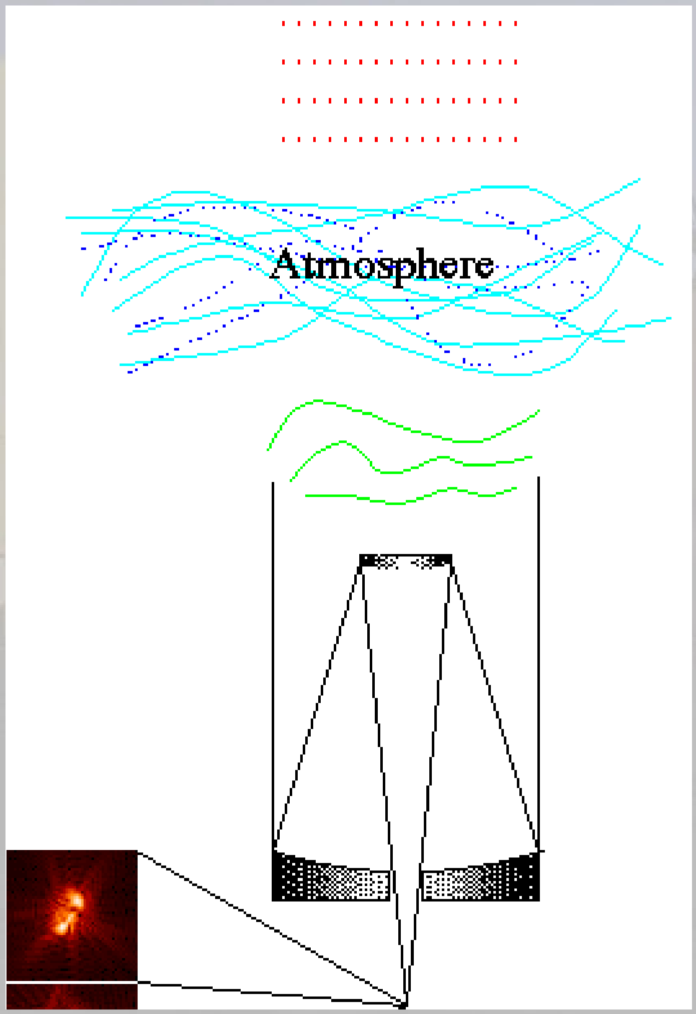

We visualise the atmosphere altering the phase (but not amplitude) of the incoming wavefronts.

We know that this impacts the instantaneous image.

But we need to understand how this impacts the fringe contrast and phase, which are what we are actually interested in measuring.

Characterizing spatial fluctuations in the wavefront

- We use a quantity called the Fried's parameter, r_0.
  - The circular aperture size over which the mean square wavefront error is approximately 1 radian^2.
  - This scales as $\lambda^{6/5}$.
  - Tel. Diameters > or < r_0 delimit different regimes of instantaneous image structure:
    - $D < r_0 \Rightarrow$ quasi-diffraction limited images with image motion.
    - $D > r_0 \Rightarrow$ high contrast speckled (distorted) images.
  - Median r_0 value at Paranal is 15cm at 0.5μm.

If you use a telescope with $D \sim r_0$ in diameter, spatial fluctuations are not very important. No one frequently does this though!

How do these spatial corrugations affect the instantaneous fringe contrast?

– Reduces the rms visibility (-----) amplitude as D/r_0 increases.
– Leads to increased fluctuations in V.
– Both imply a loss in sensitivity.
– Calibration becomes less reliable.

– Moderate improvement is possible with tip-tilt correction (———).
– Higher order corrections improve things but more slowly.

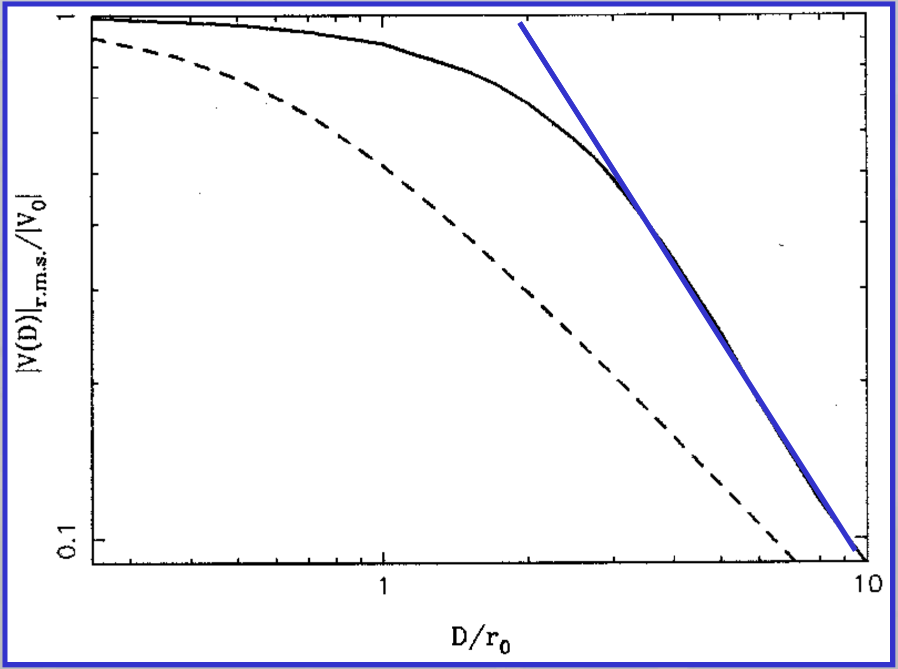

The mean fringe contrast now is a function of both the source structure and the atmospheric conditions

How can we mitigate against these perturbations?

- In principle, there are two approaches to deal with spatial fluctuations for telescopes of finite size:

λ/μm	1.25	1.65	2.2	3.5	5.0
D/r0 for ATs	3.8	2.7	1.9	1.1	0.7
D/r0 for UTs	16.7	11.9	8.4	4.8	3.1

- Use an adaptive optics system correcting higher order Zernike modes:
  - Can use either the source or an off-axis reference star to sense atmosphere.
  - But need to worry about how bright and how far off axis is sensible.

- Instead, spatially filter the light arriving from the collectors:
  - Pass the light through either a monomode optical fibre or a pinhole.
  - This trades off a fluctuating visibility for a variable throughput.

So how does spatial filtering work?

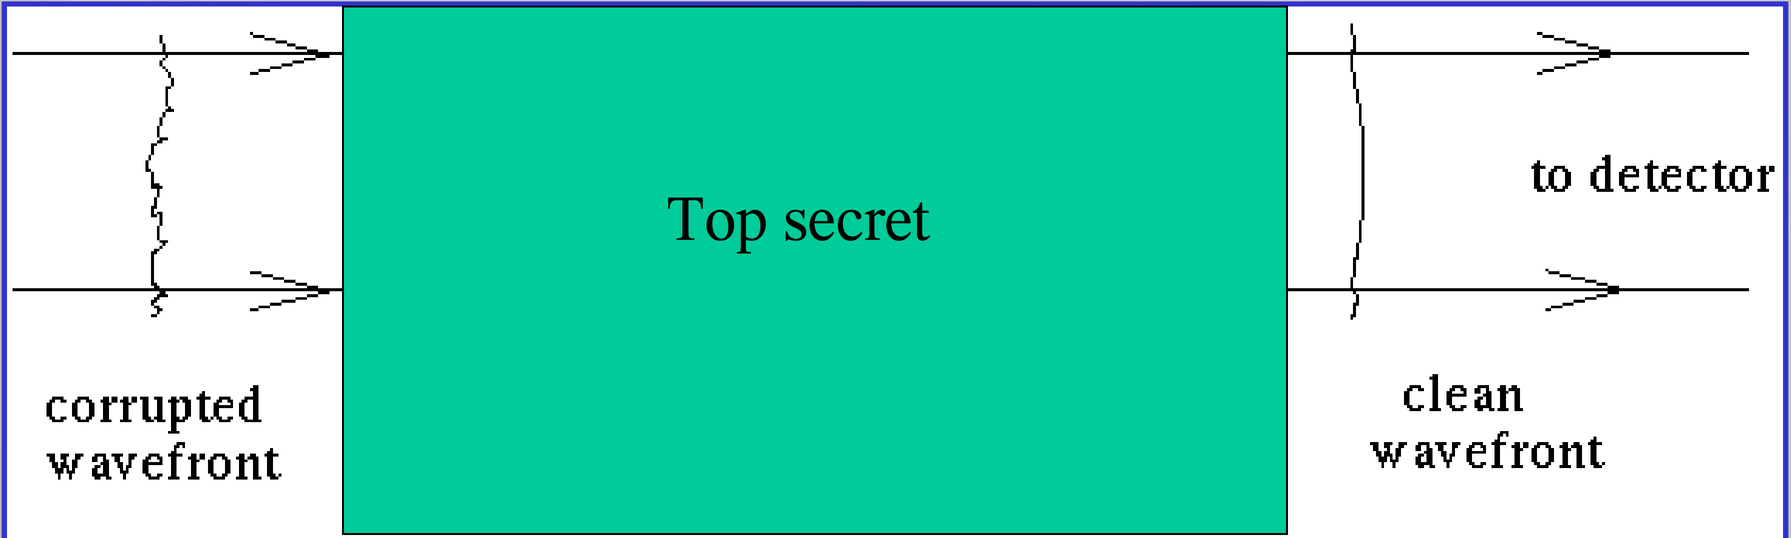

So how does spatial filtering work?

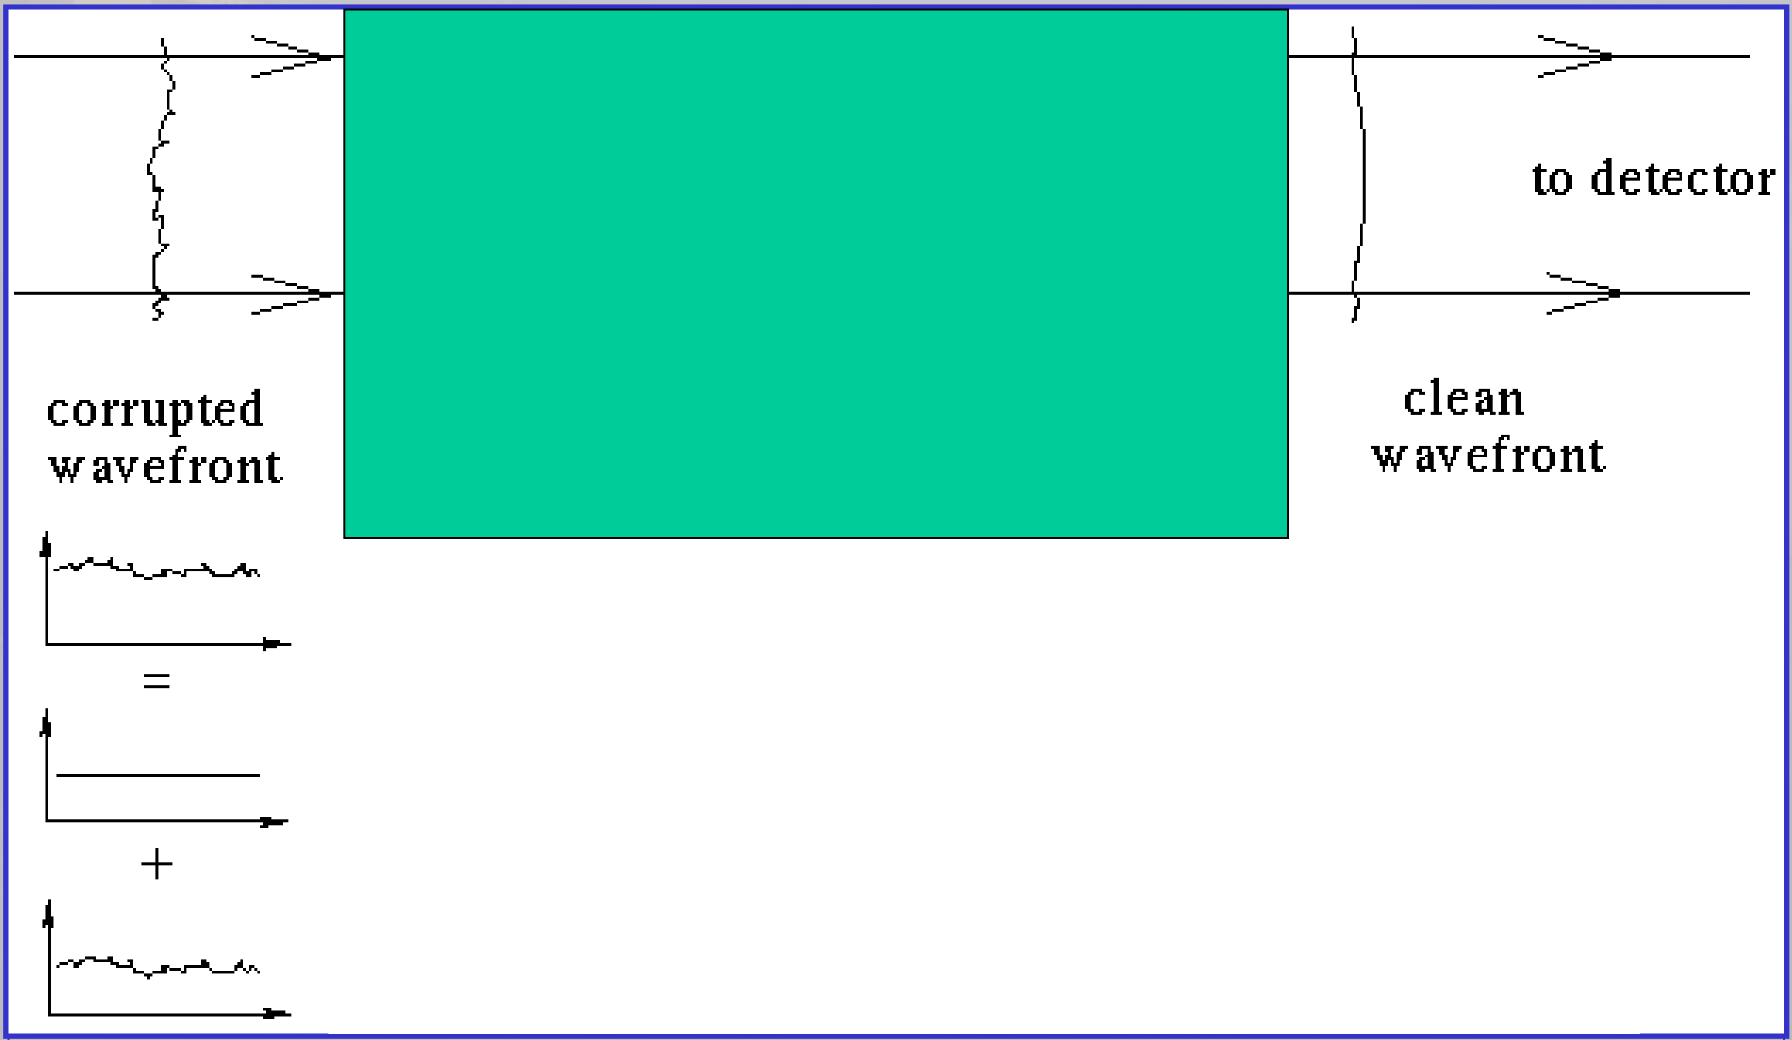

So how does spatial filtering work?

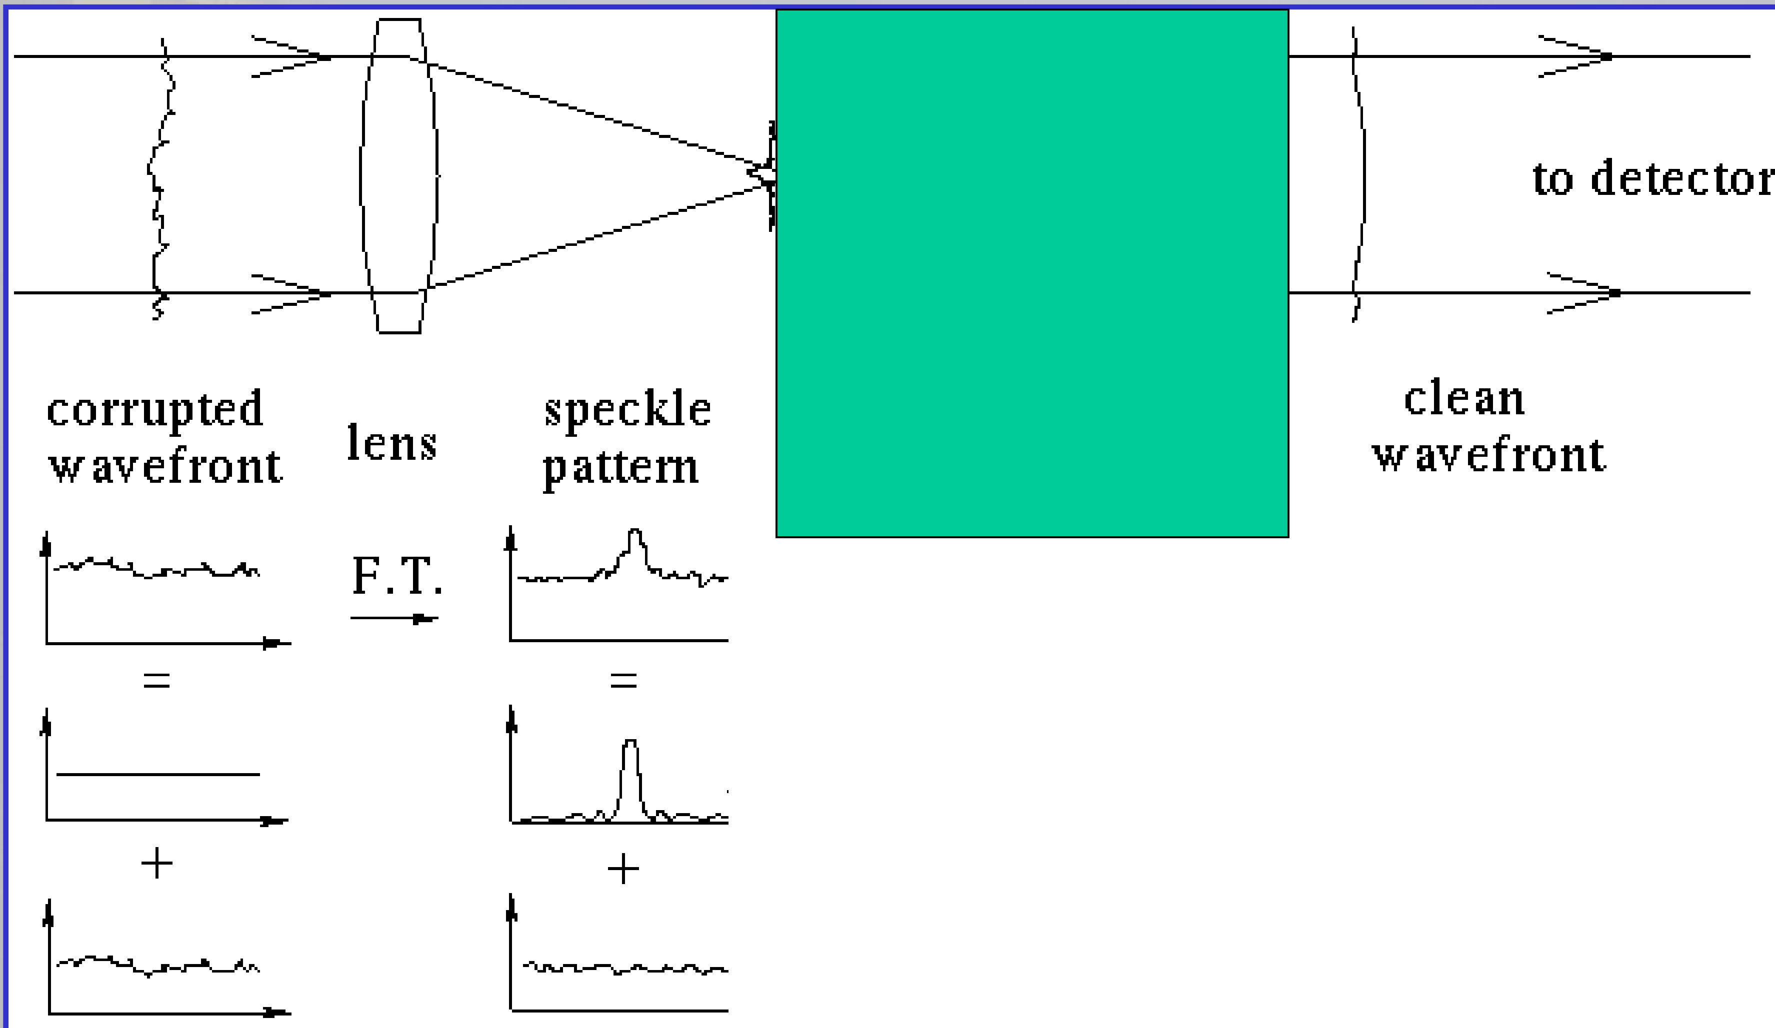

So how does spatial filtering work?

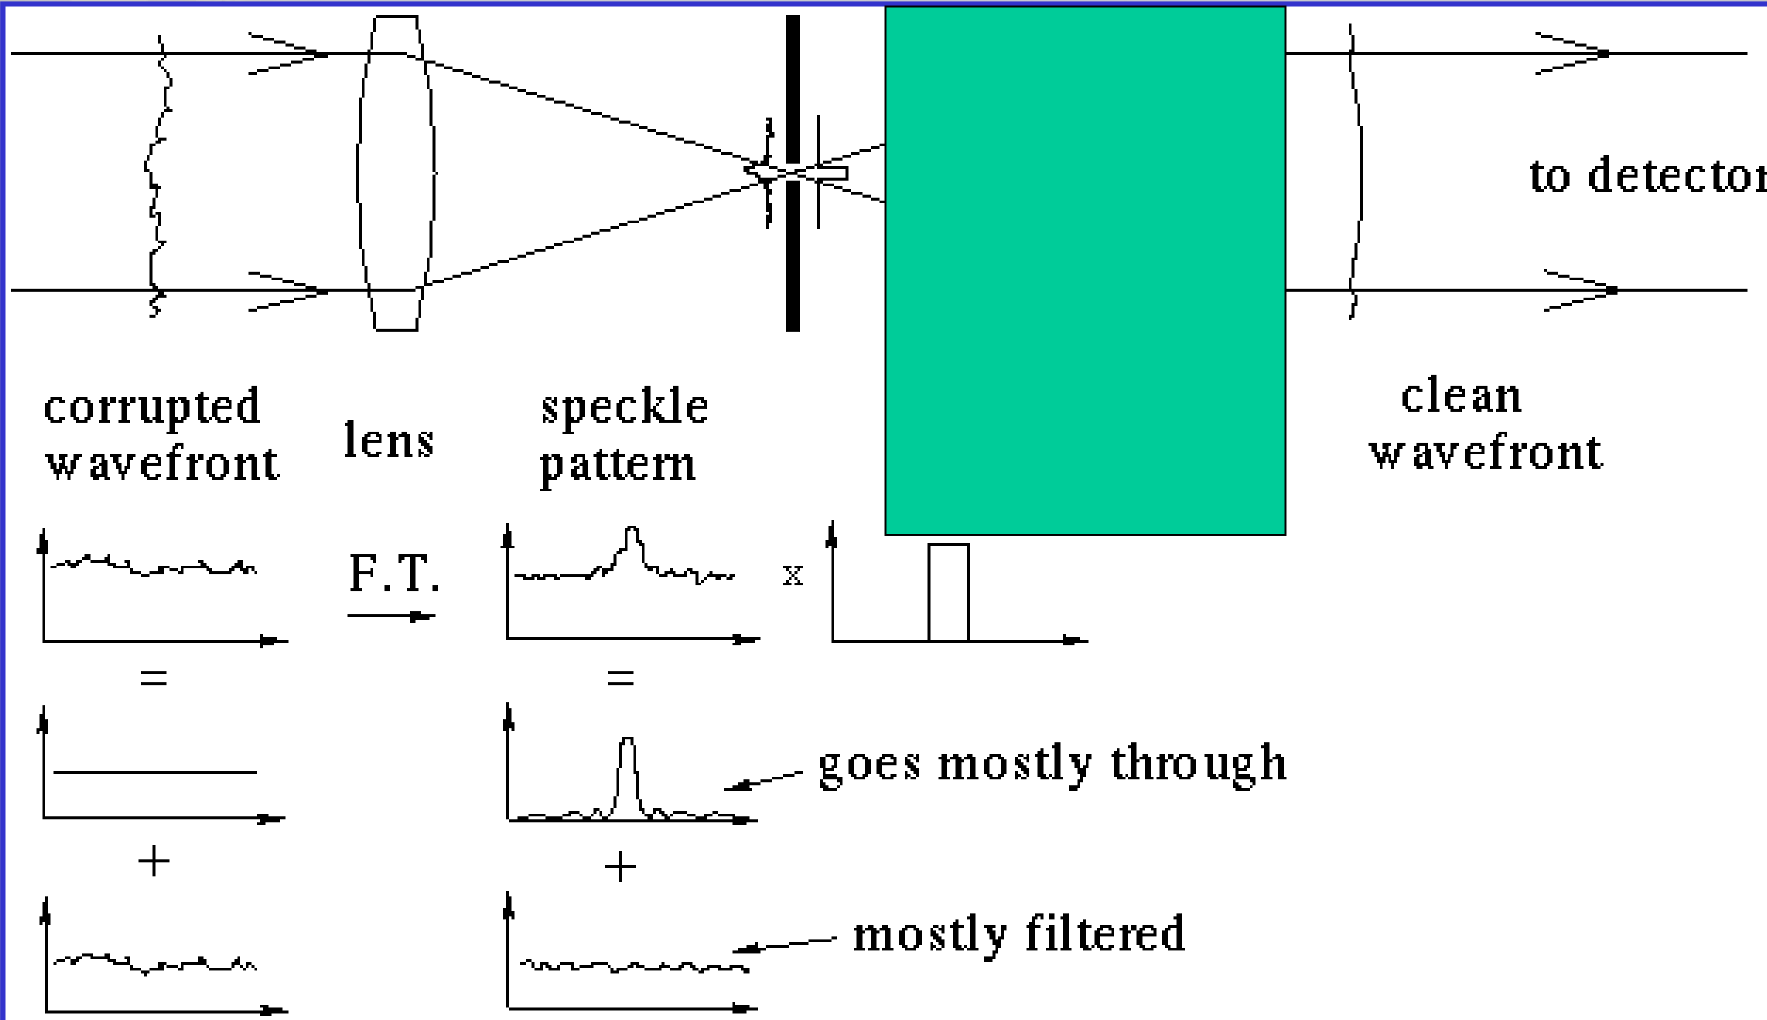

So how does spatial filtering work?

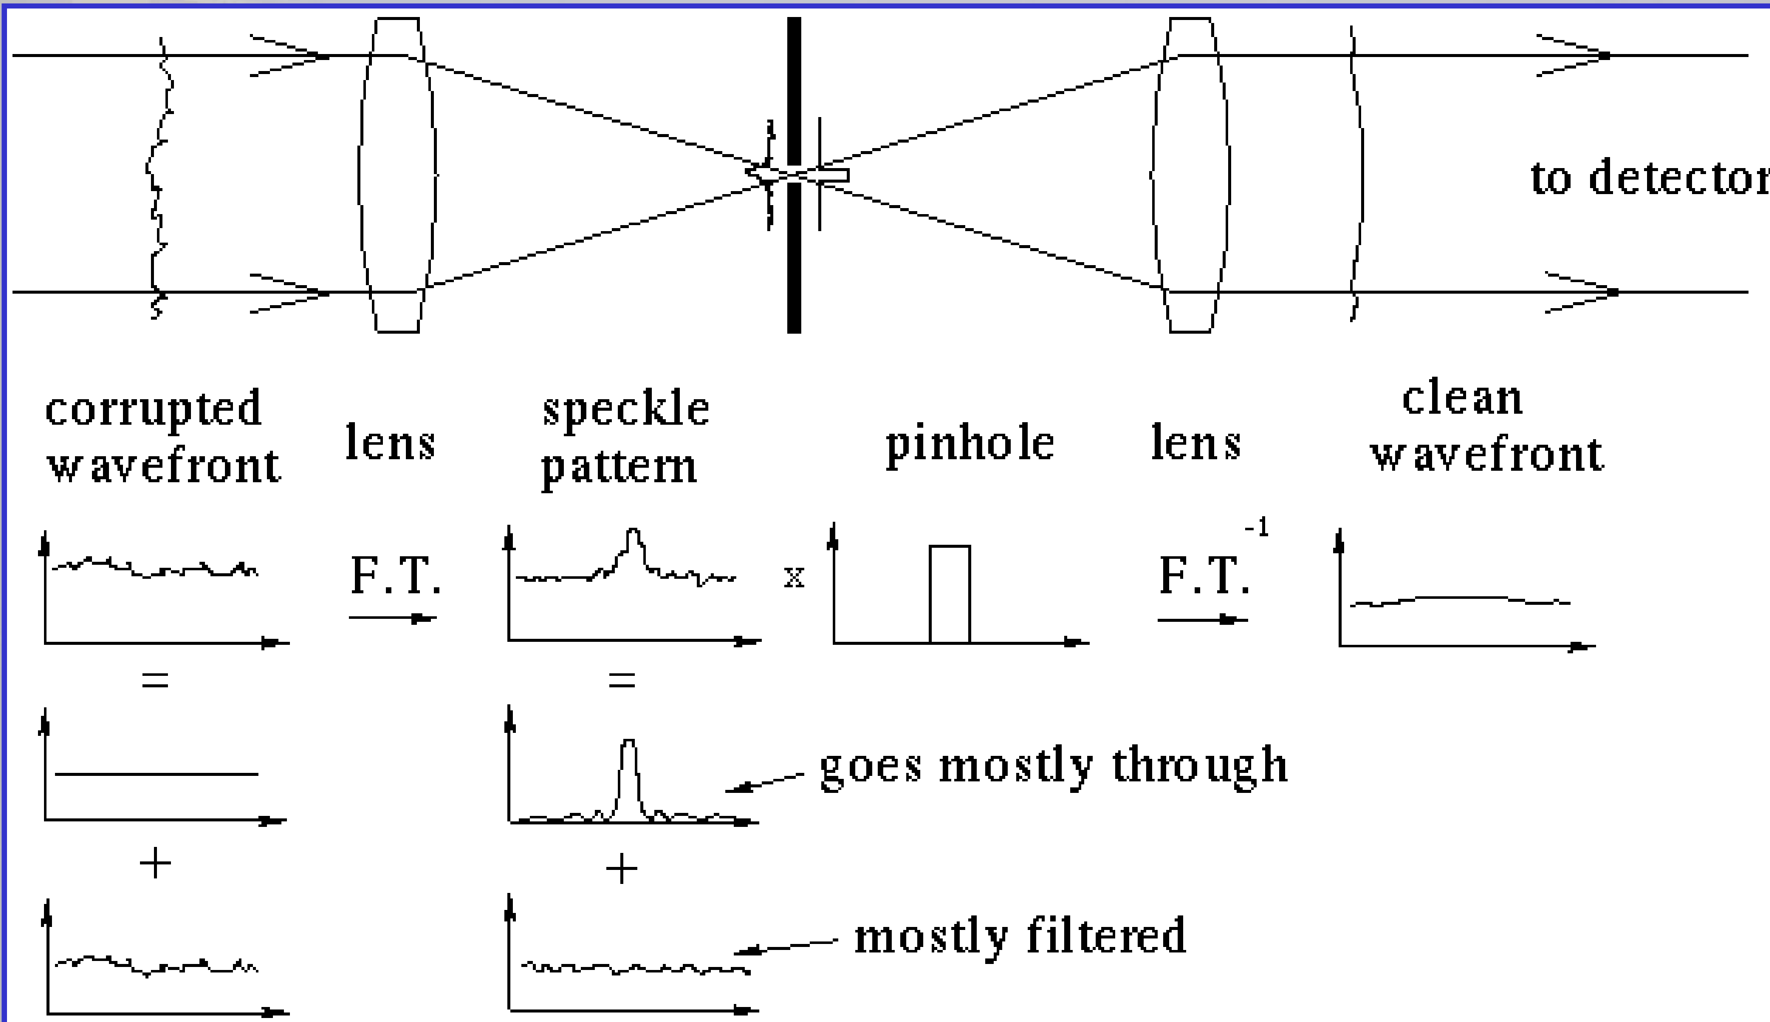

Many interferometers use both strategies

Lets look at how the S/N for fringe contrast measurement scales with telescope size

- Solid = with spatial filter
- Dashed = without spatial filter.
- Different curves are for 2, 5 and 9 Zernike mode correction.

- Implications are:
  - For perfect wavefronts S/N $\propto$ D.
  - Spatial filtering always helps.
  - Can work with large D/r_0 (e.g. ≤ 10).
  - If D/r_0 is too large for the AO system, make D smaller.

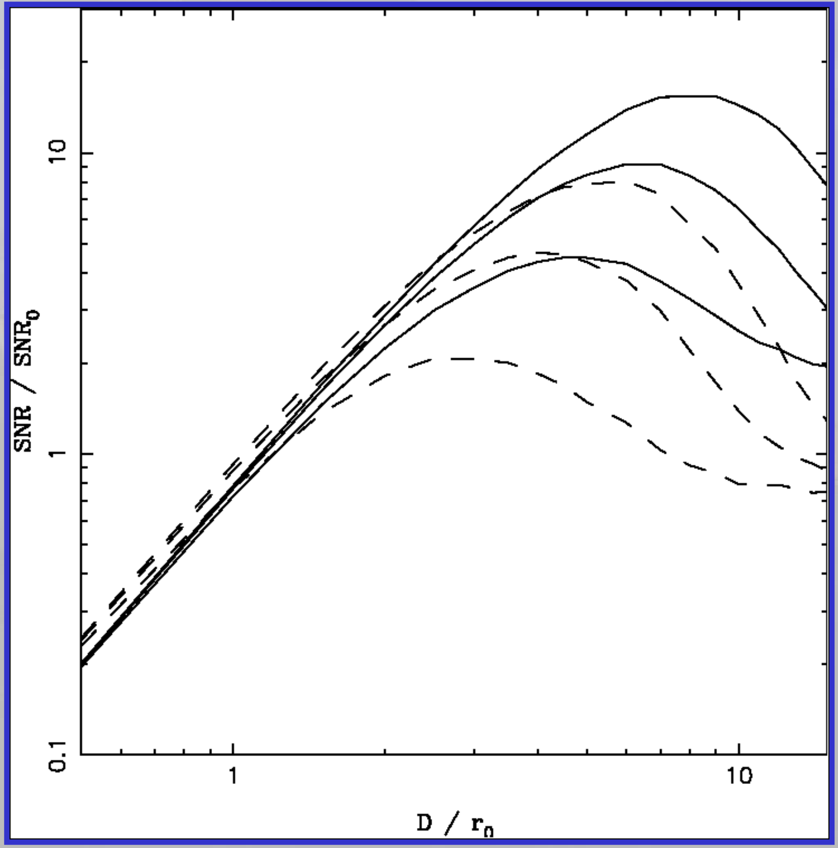

BUT – don’t forget the following for NGS AO

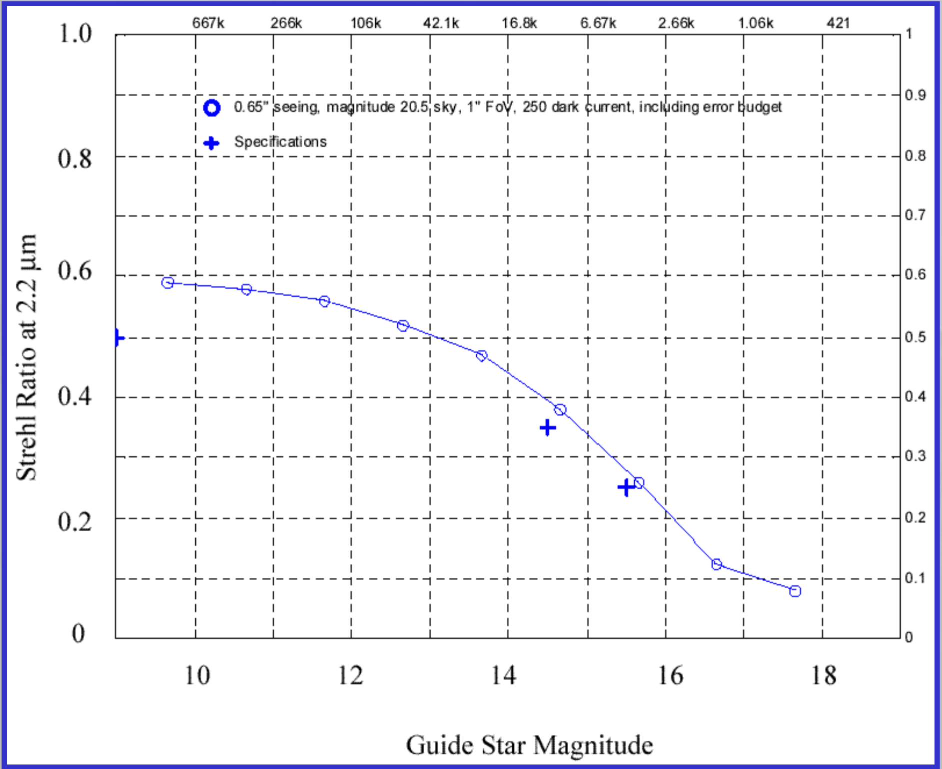

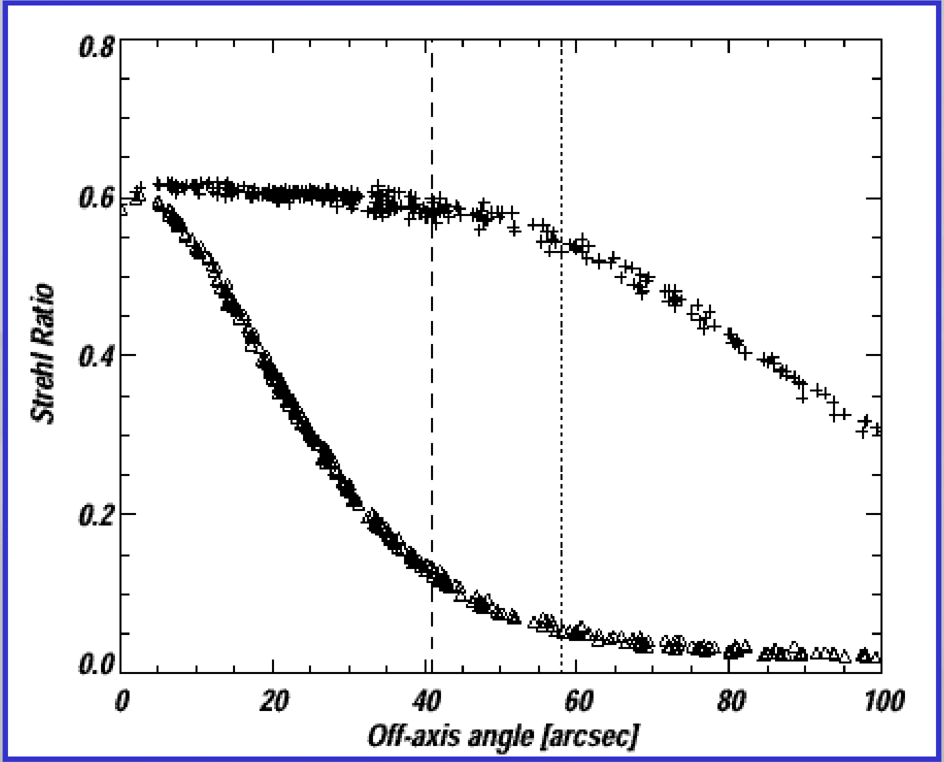

- Influence of guide-star magnitude. This is for MACAO at the VLTI.
- Influence of off-axis angle. This is for a generic 8m telescope at M. Kea.

NGS systems basically offer modest improvement in sky coverage, but are vital in allowing photons to be collected faster for bright sources.

Recap & questions

Spatial fluctuations in the atmosphere lead to a reduction in the mean fringe contrast from the intrinsic source-dependent value.

We have to calibrate this effect by looking at source whose visibility we know a priori – usually an unresolved target.

We have to rely upon the instrumental and atmospheric characteristics being identical (in a statistical sense) for the two observations.

It is possible to moderate the effects of the atmosphere using AO and spatial filtering.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the (u, v) plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

Now we look at how the temporal variation of the atmosphere alters what we heard yesterday.

The Earth's atmosphere – temporal effects

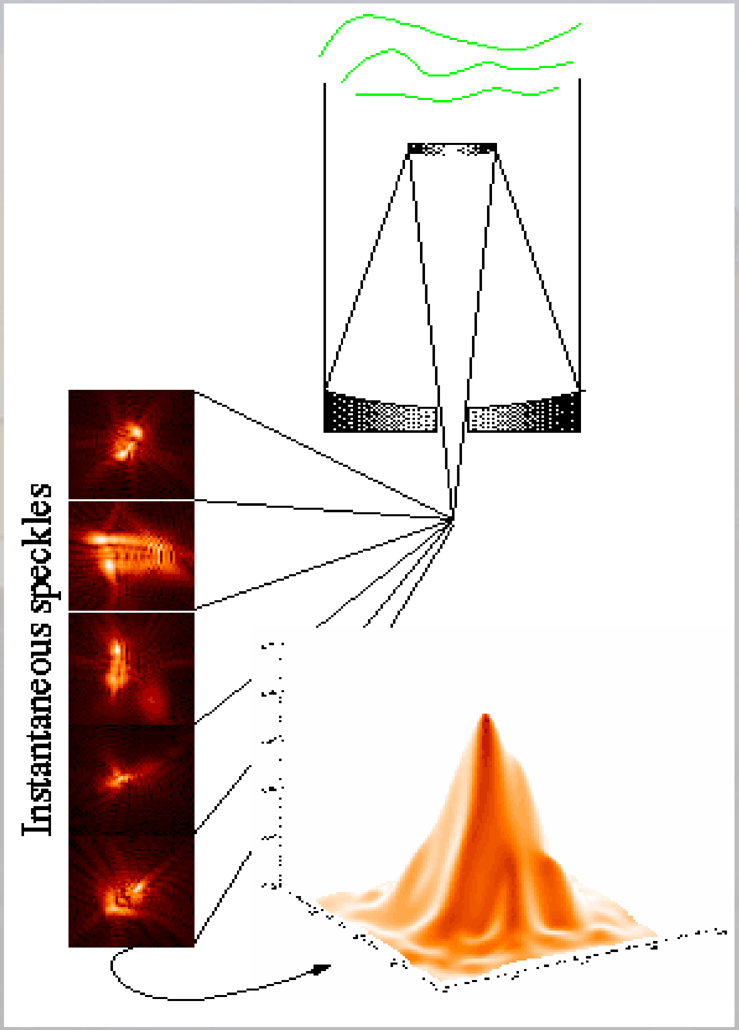

We again visualise the atmosphere altering the phase (but not amplitude) of the incoming wavefronts.

We know that these perturbations change with time.

Again, we need to understand how this impacts the fringe contrast and phase.

Characterizing temporal wavefront fluctuations

- We do this with a quantity called the coherence time, t_0.
  - Heuristically this is the time over which the wavefront phase changes by approximately 1 radian.
- Related to spatial scale of turbulence and windspeed:
  - If we assume that Taylor's "frozen turbulence" hypothesis holds, i.e. that the timescale for evolution of the wavefronts is long compared with the time to blow past your telescope.
  - Gets us a characteristic timescale $t_0 = 0.314\ r_0/v$, with v a nominal wind velocity. Scales as $\lambda^{6/5}$.
- Typical values can range between 3-20ms at 0.5µm.
  - Recent data from Paranal show median value of ~20ms at 2.2µm.

We have to make measurements of the fringe parameters in a time comparable to (or shorter) than t_0.

So how do these temporal fluctuations affect things?

- Temporal fluctuations provide a fundamental limit to the sensitivity of optical arrays.
  - Short-timescale fluctuations blur fringes:
    - Need to make measurements on timescales shorter than ~t_0.
  - Long-timescale fluctuations move the fringe envelope out of measurable region.
    - Fringe envelope is few microns
    - Path fluctuations tens of microns.
    - Requires dynamic tracking of piston errors.

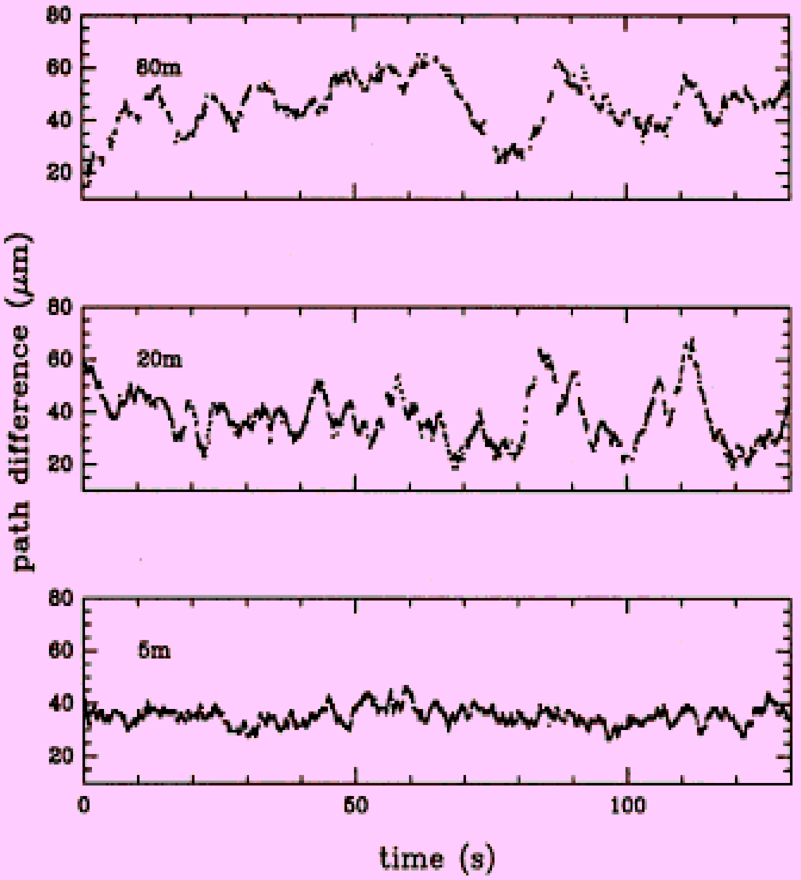

Figure 1. Path offset from the equal optical path position as a function of time for three baselines.

Perturbations to the amplitude and phase of V

- Apart from forcing interferometric measurements to be made on short timescales, the other key problem introduced by temporal wavefront fluctuations is that they alter the phase of the measured visibility function.
- Note also, that if the "exposure time" is too long, the atmosphere reduces the amplitude of the measured visibility too.

Simple Fourier inversion of the coherence function becomes impossible.

- How do we get around the problem of "altered" phases?
  - Dynamically track the atmospheric (and instrumental) OPD excursions at the sub-wavelength level
    - Phase is then a useful quantity.
  - Measure something useful that is independent of the fluctuations.
    - Differential phase.
    - Closure phase.

What you need to know about dynamic fringe tracking

- We can identify several possible fringe-tracking systems:
  - Those that ensure we are close to the coherence envelope.
  - Those that ensure we remain within the coherence envelope.
  - Those that lock onto the white-light fringe motion with high precision.

- The first two of these still need to be combined with short exposure times for any data taking.
- Only the last of these allows for direct Fourier inversion of the measured visibility function.

- As an aside, the second of these is generally referred to as "envelope" tracking or coherencing, while the third is often called "phase" tracking.

Envelope tracking – a coarse method (ask me later)

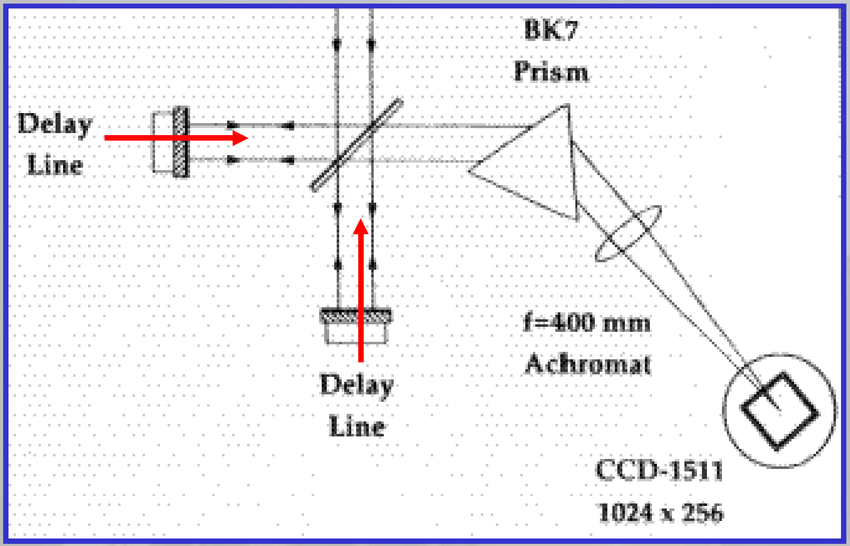

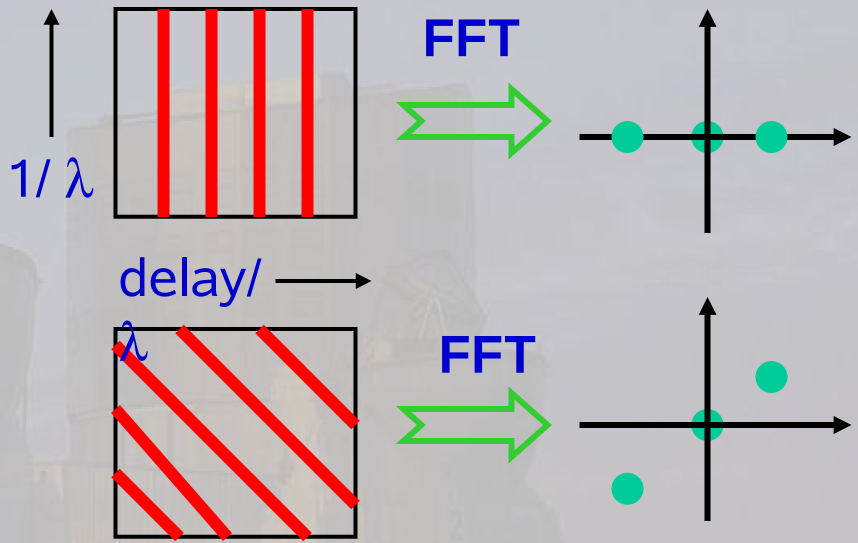

- Fringe envelope tracking methods - e.g. group delay tracking.
  - Observe fringes in dispersed light.
  - Dispersed fringes are tilted when OPD non-zero
  - Recover fringe envelope position using 2-D power spectrum.
  - Can integrate for several seconds – high sensitivity.

Phase tracking – an accurate method

- The "easy" way – this uses the target itself:
  - Use a broad-band fringe tracking channel and measure the position of the white-light fringe in real time.
  - Follow the fringe motion and sample it fast enough so that fringe motion between samples is << 180 degrees.
  - Use these broad-band measurements to phase-reference other narrow-band channels:
    - Increases effective coherence time to seconds.
    - Equivalent to self-referenced adaptive optics on the scale of the array.
  - Because it's a high precision technique it has ~2.5 mag poorer sensitivity than group-delay tracking.

Alternatively, use an off-axis reference target

- The "complex" way: dual-feed operation. This is what PRIMA aims to deliver:
  - Use bright off-axis reference star to monitor the atmospheric perturbations in real-time.
  - Feed corrections to parallel delay-lines observing science target.
  - Use a metrology system to tie two optical paths together.

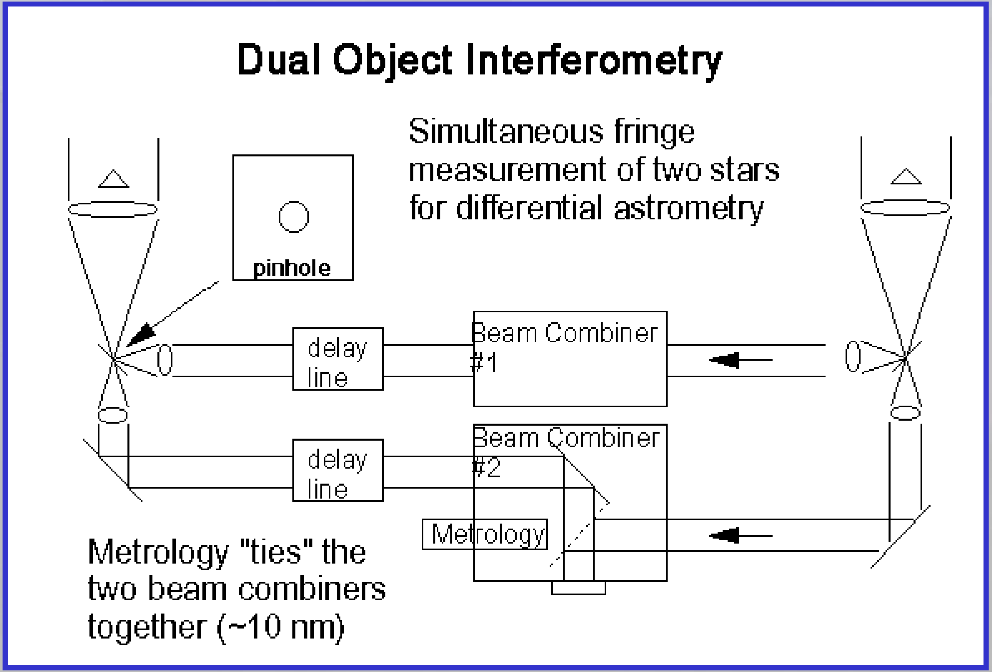

As before, this can extend the effective coherence time by orders of magnitude if the white-light fringe is successfully tracked.

Dual-feed interferometry (cont'd)

- Practical issues:
  - Off-axis wavefront perturbations become uncorrelated as field angle increases and λ decreases.
  - With 1′ field-of-view <1% of sky has a suitably bright reference source (H<12).
  - Metrology is non-trivial.
  - Laser guide stars are not suitable reference sources.

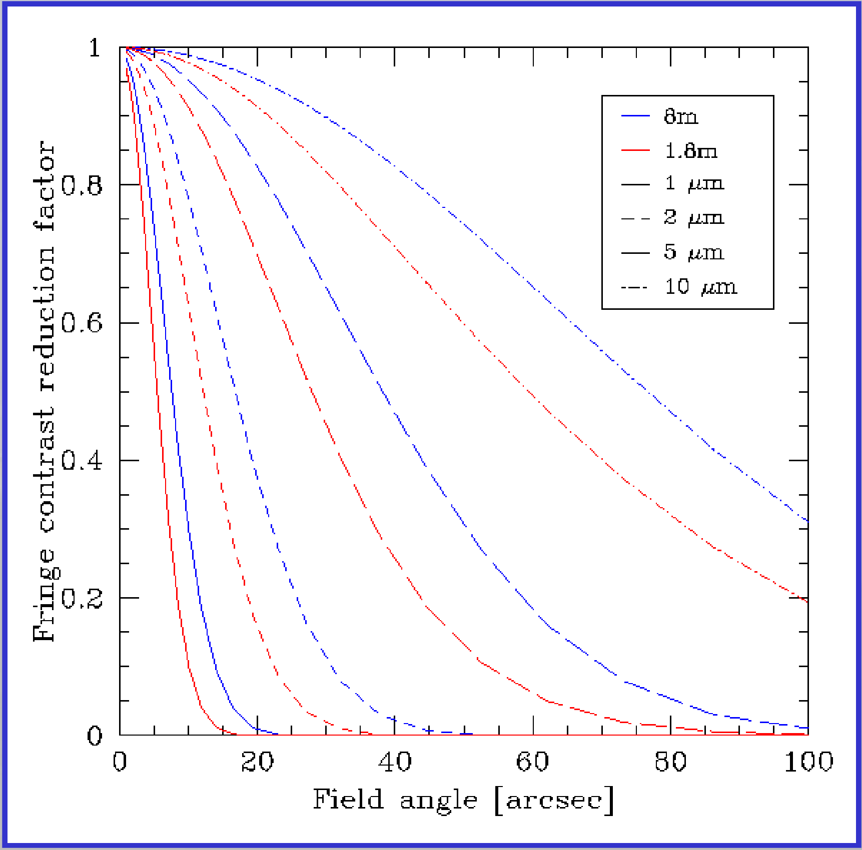

Off-axis reduction in mean visibility for the VLTI site as a function of D and λ.

An alternative strategy is to make measurements of “phase-like” quantities that are not affected so badly by the atmosphere.

This is how you design good experiments in the lab!

Good observables

- In the absence of a PRIMA-like system, optical/IR interferometrists have had to rely upon measuring phase-type quantities that are immune to atmospheric fluctuations.
- These are self-referenced methods - i.e. they use simultaneous measurements of the source itself:
  - Reference the phase to that measured at a different wavelength - differential phase:
    - You make measurements at two wavelengths simultaneously.
    - Depends upon knowing the source structure at some wavelength.
    - Need to know atmospheric path and dispersion.
  - Reference the phase to those on different baselines - closure phase:
    - Independent of source morphology.
    - Need to measure many baselines at once.

Closure phases – what are these?

- Measure visibility phase (Φ) on three baselines simultaneously.
- Each is perturbed from the true phase (ϕ) in a particular manner:

$$\Phi_{12} = \phi_{12} + \varepsilon_1 - \varepsilon_2$$
$$\Phi_{23} = \phi_{23} + \varepsilon_2 - \varepsilon_3$$
$$\Phi_{31} = \phi_{31} + \varepsilon_3 - \varepsilon_1$$

- Construct the linear combination of these:

$$\Phi_{12}+\Phi_{23}+\Phi_{31} = \phi_{12} + \phi_{23} + \phi_{31}$$

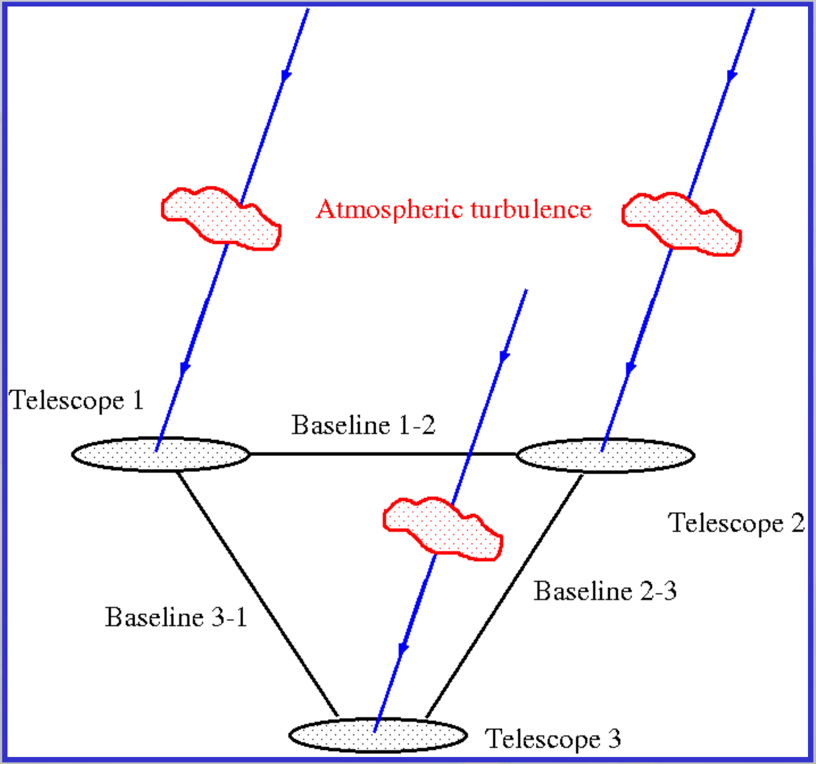

The error terms are antenna dependent – they vanish in the sum.
The source information is baseline dependent – it remains.
We still have to figure out how to use it!

So how do we use these “good” observables?

- Average them (properly) over many realizations of the atmosphere.
- Differential phase, if we are comparing with the phase at a wavelength at which the source is unresolved, is a direct proxy for the Fourier phase we need.
  - Can then Fourier invert straightforwardly.
- Closure phase is a peculiar linear combination of the true Fourier phases:
  - In fact, it is the argument of the product of the visibilities on the baselines in question, hence the name triple product (or bispectrum).

$$V_{12}V_{23}V_{31} = |V_{12}|\,|V_{23}|\,|V_{31}|\exp(I[\Phi_{12} + \Phi_{23} + \Phi_{31}]) = T_{123}$$

  - So we have to use the closure phases as additional constraints in some nonlinear iterative inversion scheme.

Recap & questions

Temporal fluctuations in the atmosphere lead to a reduction in the mean fringe contrast from the intrinsic source-dependent value if the exposure time is too long.

This cannot easily be calibrated, so we try to avoid it.

Much more importantly, temporal fluctuations change the measured fringes phases in a random manner.

One solution is to monitor (and possibly correct) the atmospheric perturbations in real time.

A second strategy is to try to measure combinations of phases that are immune to these perturbations.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the *(u, v)* plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

I will just say a few words here: these are thoughts for you to ponder in the bathtub!

Thinking about interferometric sensitivity

- We have mentioned earlier that sensitivity in an interferometric context really means two things:
  - It must be possible to stabilize the array in real time against atmospheric-induced fluctuations of the OPD.
  - Once this is satisfied, we need to be able to build up enough signal-to-noise on the astronomical fringe parameters of interest.

- The essential implication of this is that the "instantaneous" fringe detection S/N has to be high enough to "track" fringes.

- This signal/noise ratio basically scales as:

$$S/N \propto [VN]^2 \Big/ [(N+N_{dark})^2 + 2(N+N_{dark})N^2V^2 + 2(N_{pix})^2(\sigma_{read})^4]^{1/2}$$

with V = apparent visibility, N = detected photons, N_{dark} = dark current, N_{pix} = number of pixels, σ_{read} = readout noise/pixel.

This formula bears some consideration

$$S/N \propto [VN]^2 / [(N+N_{dark})^2 + 2(N+N_{dark})N^2V^2 + 2(N_{pix})^2(\sigma_{read})^4]^{1/2}$$

with V = apparent visibility, N = detected photons, N_{dark} = dark current, N_{pix} = number of pixels, σ_{read} = readout noise/pixel.

- It depends on the signal coming from the target.
- It depends on the amount of read noise on the detector.
- It depends on the amount of dark/thermal background.
- It depends on the source visibility, i.e. its structure.

This last point is not too unusual really.

It just means that what matters is not the integrated brightness of the target but the surface brightness, i.e. the brightness per unit solid angle on the sky.

Recap and questions

Need to have enough V^2N to stabilize the array.

Then need enough integration time to build up a useful signal-to-noise on the science signal.

The problem is that many interesting targets have low V (and N)!

Solutions include: (a) use an off-axis reference star for stabilization (PRIMA), (b) decompose all long baselines into shorter ones where V is not so low.

Outline

- What are the things that make interferometry less than straightforward in practice?
  - Sampling of the (u, v) plane
  - Beam relay
  - Delay compensation
  - Beam combination
  - Spatial wavefront fluctuations
  - Temporal wavefront fluctuations
  - Sensitivity
  - Calibration

Phew – we are almost done.

Calibration

- The basic observables we wish to estimate are fringe amplitudes and phases.
- In practice the reliability of these measurements is generally limited by systematic errors, not the S/N we have just discussed.
- Hence there is a crucial need to calibrate the interferometric response:
  - Measurements of sources with known amplitudes and phases:
    - Unresolved targets close in time and space to the source of interest.
  - Careful design of instruments:
    - Spatial filtering.
  - Measurement of quantities that are less easily modified by systematic errors:
    - Phase-type quantities.

Examples of real data

- Measurements with the NPOI

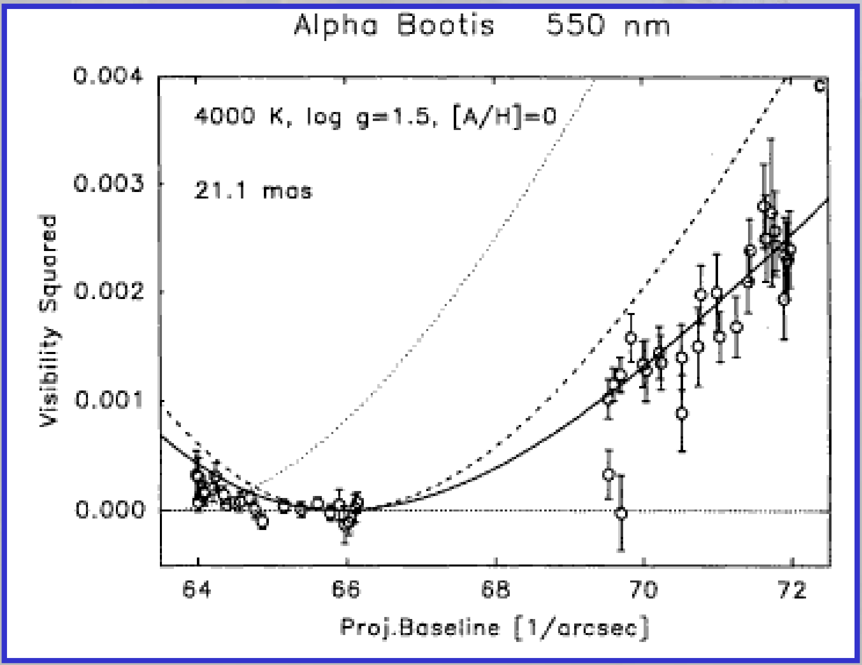

- Measurements with FLUOR

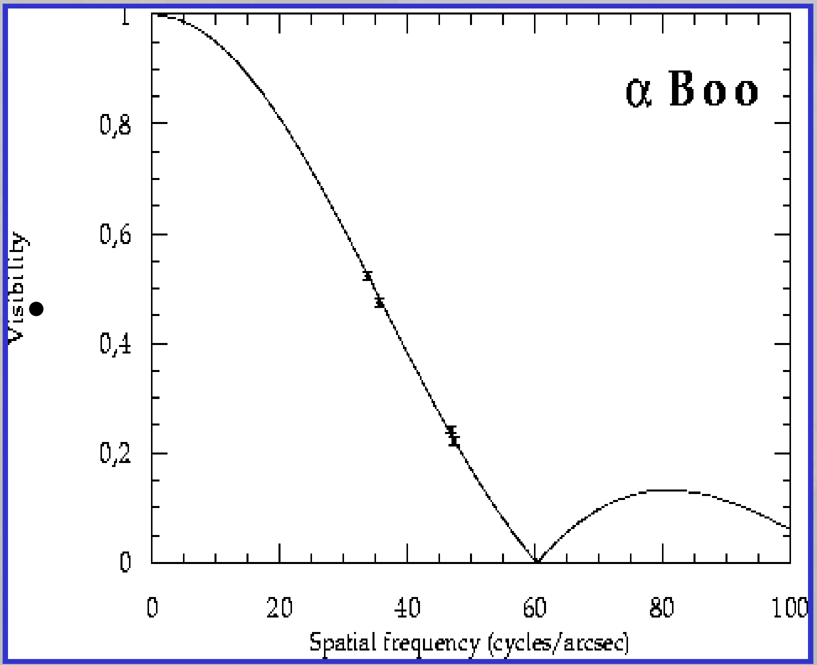

Perrin et al, AA, 331 (1998)

Notice how small these error bars are: compare these to what you get with AMBER and MIDI today.

Reconstructions from different quality "fake" data

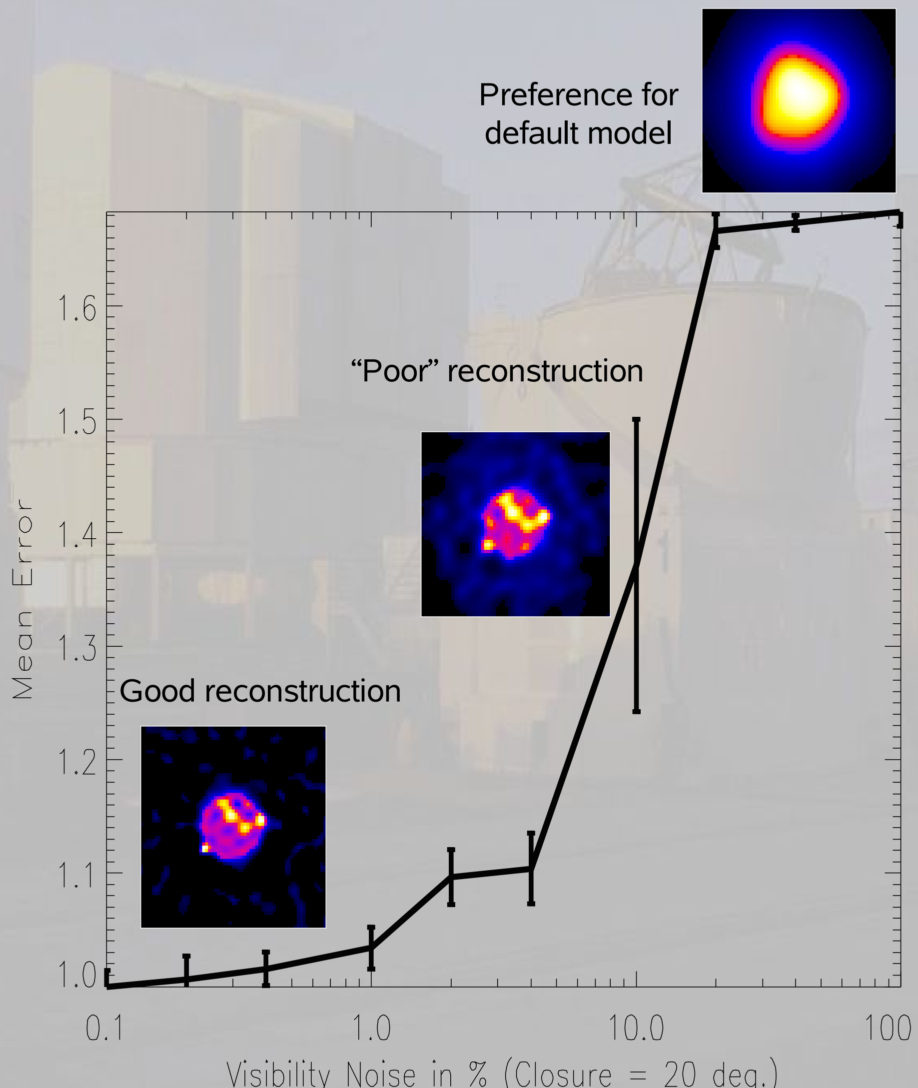

Interestingly, even when the data seem quite poor, you can still learn a lot.

So you need to know what is required for the science.

Key lessons to take away

- Sampling of the *(u, v)* plane
  - What is needed for the scientific questions being addressed.
  - Will the array operate satisfactorily on these baselines.
- Beam relay
  - Maximum efficiency, polarization, stability.
- Delay lines
  - Intrinsic performance, dispersion at long baselines.
- Spatial fluctuations
  - Impact on sensitivity and fringe contrast, potential limitations of AO.
- Temporal fluctuations
  - Impact on sensitivity and fringe phase, need for fringe tracking.
  - Good observables and how these are used.
- Sensitivity
  - An appropriate measure of this in terms of stabilizing the array.
  - V^2N scaling.
- Calibration
  - Importance of matching this to the science desired.

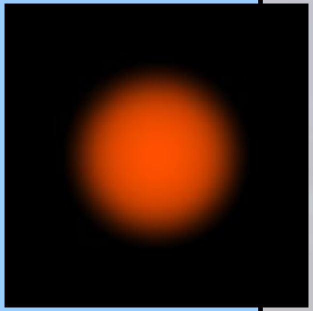

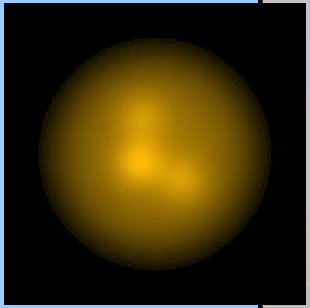

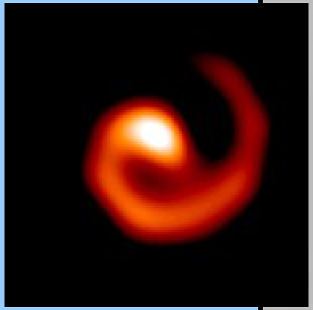

Final quiz

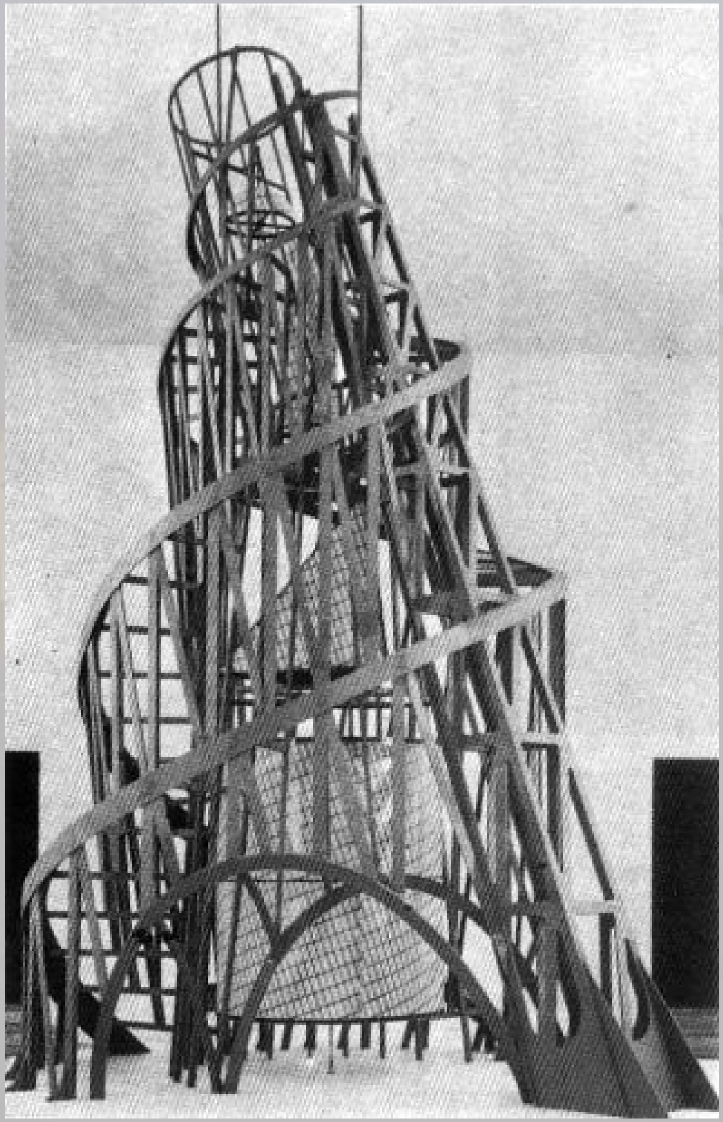

What is this?