Chapter 4 = 9417 words

1. High Resolution Imaging

2. Infrared Detector Technology

3. Camera

4. COAST

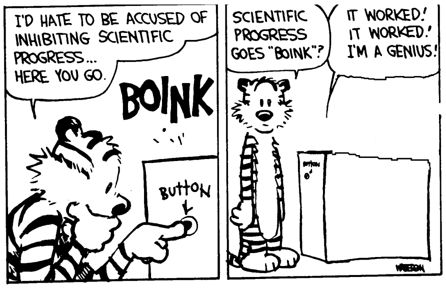

Outline

The techniques of optical interferometry and the advantages of infrared operation have been discussed in chapter one. This chapter describes the existing COAST system and the changes needed to operate in the infrared. The design, construction and performance of the specifically infrared system is examined in detail.

4.1 Description of COAST

The Cambridge Optical Aperture Synthesis Telescope (COAST) is an interferometric array of 4 small telescopes operating on large baselines to produce images at visible wavelengths with an ultimate resolution of 1 milli-arcsecond (mas), Baldwin(1994).

4.1.1 Site layout

The telescopes are arranged in a Y shaped formation, similar to large radio interferometers such as the VLA. This gives good coverage of the u-v plane in a single exposure and allows Earth rotation synthesis to build up a complete picture. The central telescope is fixed about 5 metres in front of the optics lab. The other telescopes can be moved to prepared foundations on three arms extending about 100m from the centre to give a maximum baseline of 160m. The original design of the telescope attempted to minimise any disturbance to the local atmosphere. The site is flat and covered with grass, there are no roadways near the telescopes and the above ground foundations are minimised. The optics lab is half buried and landscaped into the site providing passive temperature stabilisation and reducing wind turbulence.

4.1.2 Number of telescopes

The number of telescopes in a system such as COAST is a compromise. A large number of telescopes can simultaneously measure more baselines but the light from each telescope must then be split more ways to combine with the light from each of the other telescopes. In radio telescopes, a large number of individual antennae are generally used to obtain good instantaneous u-v coverage. This is possible because amplifiers are available that allow the signal to be increased before it is split. At optical and near-infrared wavelengths suitable amplifiers are not available and the shortage of light is a major limiting factor throughout the design of COAST.

A system using N telescopes produces measurements on $\frac{1}{2}N(N-1)$ simultaneous baselines and $\frac{1}{2}(N-1)(N-2)$ independent closure phases. It has been shown theoretically that the signal to noise of a visibility estimate from a single exposure is proportional to $1/\sqrt{N}$ for purely shot noise limited detectors. This suggests that as few telescopes as possible should be used in the array, with four being the minimum number necessary to produce more than one simultaneous closure phase. In addition to the reduced signal levels in the correlator, there is also the extra cost and complexity in building and aligning telescope, path compensator and beam combining optics for a large number of telescopes. These considerations only apply to the number of telescopes actually contributing to the fringe measurement. It may be useful to have extra telescopes in position and aligned so that different baselines can be switched in allowing the array configuration to be changed quickly.

The COAST array will use 4 telescopes giving a total of 6 baselines and 3 independent closure phases.

4.1.3 Telescopes

The telescopes are 0.4m aperture Cassegrains with a modified secondary mirror to produce a parallel output beam of 25 mm diameter at a demagnification of x16. The telescope design is unusual in that the optical axis is horizontal, with the starlight fed from a 0.5m siderostat mounted on computer controlled rotary tables. This maximises the stiffness of the structure. The telescopes sit on kinematic mounts which keep the relative positions of the central point of the siderostat fixed independent of any temperature induced changes in the structure, Boysen(1991). A computer controlled piezo-electric tip-tilt mirror on each telescope removes first order atmospheric phase tilts and feeds the light back to the central telescope through insulated pipes suspended above ground at (painfully) head height. From this telescope the light is reflected into the building.

The optimal telescope size depends on the wavelength and atmospheric conditions. Larger apertures provide more light, but at shorter wavelengths or in poor seeing conditions the greater wavefront corruption across the aperture lowers the fringe visibility and so reduces the overall signal to noise. The optimal aperture size is related to the observing conditions by the Fried length parameter r_0 , defined in chapter one. In the case of photon noise limited observations the optimal aperture is 1.2 r_0 which increases to 2.8 r_0 for tip-tilt corrected wavefronts, Buscher(1988). For observations at visible wavelengths or in poor seeing conditions the telescopes are masked to reduce the aperture to this factor. An estimated r_0 of 10 cm in the visible implies a value of 25 cm in the J band and 50 cm at K band. At the longer infrared wavelengths the full telescope aperture will be used.

4.1.4 Path compensation

To obtain high visibility fringes, the optical path from the star to the beam-combining optics through each of the four telescopes must be identical to within a coherence length of a few wavelengths. When observing objects not directly overhead there is obviously a difference in arrival times for photons from the source to each of the telescope. This is corrected by introducing an extra path delay into the light from each of the telescopes. The path compensation system does this by means of a roof mirror mounted on a trolley moving along 25m of precision rails. This gives 50m of path compensation and allows baselines of up to 150m with objects near the zenith. To achieve the necessary position accuracy the path compensation mirrors are isolated from the trolley by a flexture mount driven by a voice coil. A commercial laser metrology system then measures the position of a corner cube mounted on this roof mirror and provides a correction signal which completes the servo loop. This is shown in the diagram in Figure 4.12. An extra waveform term can be imposed on the calculated trolley position which allows the fringe pattern to be swept past a single element detector.

The path compensator operates in exactly the same way for the infrared as in the existing visible system. The longer observing wavelengths in the infrared mean that the accuracy requirements are relaxed. Unfortunately because of the extra distance to the infrared optical table, the zero path position is different and so it is not possible to use the visible and infrared systems simultaneously.

4.1.5 Autoguider

After leaving the path compensator, the light entering the visible system is first split by a dichroic filter. Here the red light with a wavelength, $\lambda > 650$ nm, is reflected into the beam combining system. The shorter wavelength, blue light is more corrupted by the atmosphere and so is not used for observing. This instead passes to a visible CCD camera. This camera produces an image with a field of view of 20 arc-seconds from each telescope. This is used to acquire the star and correct errors in the telescope pointing. When images have been obtained from all the telescopes, the camera switches into a fast guiding mode. It now reads a small sub region of each image as a simulated quad cell. This provides a correction signal to the tip-tilt mirror on the telescope at about 80 Hz, removing the effects of first order tilts in the wavefront, Cox(1991).

Many of the infrared objects of interest to COAST are very faint in the visible and blue wavebands and so less light is available to the autoguider. However in the infrared the larger atmospheric timescale means that a slower servo loop is required and so a longer exposure time can be used. To observe very faint infrared objects it may be necessary to remove the visible system dichroic, allowing all the visible light to reach the autoguider.

4.1.6 Visible / Infrared selection

One of the major design aims of the infrared system was that it would work in parallel with the visible system and that the construction wouldn't disturb the operation of the existing visible telescope. The system must divide the light between the visible and infrared components. This division should take place after the path compensator since this component can be used for both systems. An ideal solution would be a dichroic filter at the output of the path compensator which would reflect infrared light to the new system while allowing visible light through to the existing beam combiner and autoguider. Unfortunately it proved impossible to design a dichroic filter which could achieve this with high efficiency. These dichroics are based on a stack of coatings, each a fraction of a wavelength thick, deposited on a transparent substrate. It is relatively simple to produce a coating which will transmit long wavelengths while reflecting shorter ones. However the position of the existing visible system, shown in Figure 4.5, means that the opposite arrangement is needed.

An alternative scheme was tested in which a mirror was manufactured on a transparent substrate with an unsilvered central hole covering about 50% of the diameter. This was arranged so that light for the visible system and autoguider would pass through the hole while infrared light is reflected by the rest of the mirror. Since the visible system generally uses masks on the telescope to reduce their aperture, the amount of light it received would not be affected. The larger atmospheric scale length in the infrared means that it can use the full telescope aperture and so would still receive the majority of the light from the outer parts of the telescope. This scheme failed due to difficulties in superimposing the resulting annular pupils in the beam combiner. The diffraction from the central obstruction also produced a larger spot size in the camera which reduced the amount of light collected at each detector.

Eventually a thin film gold coated mirror was produced which reflects about 90% of the infrared light while transmitting 20-40% across the visible bands to the autoguider. This mirror can be removed to allow visible observations and replaced kinematically. Although simultaneous operation is now not possible, the system can be switched between infrared and visible systems in a few seconds.

4.2 COAST beam combination

The beam combiner must interfere the light from each of the telescopes so that an interference pattern is produced. From this the visibility and relative phase of the signal on each baseline can be determined and ultimately a map produced. The main decision in the design of the beam combiner is between combining all the four beams together into a single pattern or combining pairs of beams. The theoretical basis for each system is described in Buscher(1988), which concludes that the choice of system is largely set by practical considerations.

The main difference between COAST in the visible and the infrared is the nature of the detector. The visible system uses photon-counting Avalanche Photo-Diodes (APDs), these are single element devices which record the arrival of each individual photon, Nightingale(1991). An integration is made by incrementing an external counter with each photon event and then periodically reading the value of the counter. The exposure time can thus be made arbitrarily short and there is no noise associated with the process of reading the counter. The major noise source comes from the shot noise on the dark signal. The main disadvantage of the system is a dead-time after each photon arrival which leads to non-linearity at high count rates. The units produced for COAST have a DQE of around 40% and are used at photon rates of up to 1 MHz. The dark signal at an operating temperature of -23 °C is 300-1000 counts/second.

The infrared detector is an array of 256 x 256 separate pixels. A signal is integrated in each pixel for as long as light is falling onto it. At the end of the exposure the whole array or a sub-section is a readout. Each measurement of the signal integrated on a pixel has a noise term independent of the magnitude of the signal. In the camera built for COAST the rms noise on each integration is around 25 e^-/pixel. The pixels are read consecutively with the process of reading a single pixel taking around 50 µs.

4.2.1 Image-plane, all-together

The simplest system combines the beams from all the telescopes into a single fringe pattern which is then imaged directly. In this arrangement path compensated beams are arranged in a line with non redundant spacing as shown in Figure 4.1. A lens would then form a single image on the detector of an Airy pattern from the object, crossed by fringes. The fringes due to interference on each individual baseline can be recovered since they appear at different spatial frequencies due to the non-redundant spacing of the input beams. An exposure would be taken of the whole fringe pattern must be taken in a time short enough to freeze the effects of the atmosphere on the fringe visibility. This is an adaptation of the system used on the William Herschel Telescope (WHT) for aperture masked interferometry, Haniff(1994).

The layout of the system is shown below in Figure 4.1. The beams, each of diameter *D,* arrive from the path compensator. They are then separated in such a way that only a single pair of beams combine to produce fringes at frequency of 1,2,3,4,5 and 6 D. A microscope objective is then used to re-image the fringe pattern at the correct image scale on the detector. In practice an all-mirror system would be used instead of the lens to improve the efficiency.

The array detector naturally suggests such an all-together spatially resolved scheme. This would involve fewer components and so be easier to align as well as having lower losses. A similar system has been proposed for the VLT interferometer, Merkle(1988).

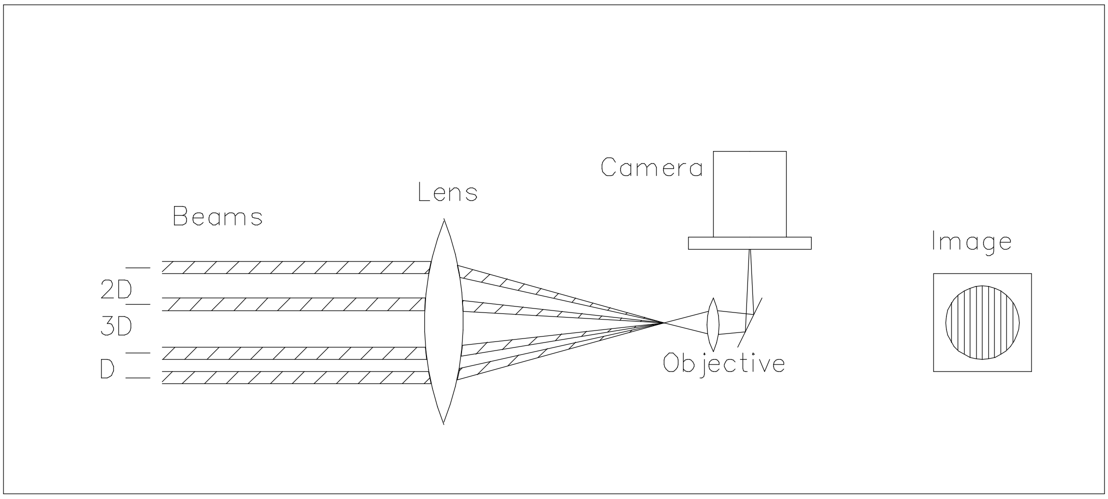

Figure 4.1: All-together, image plane combination scheme

The main problem with a spatial scheme is the number of pixels which must be read out. If a large number of pixels are needed to sample the fringe pattern then the overall noise increases since there is a read noise associated with each pixel. In addition to this, if it takes more than an atmospheric coherence time (t_0) to read out the image then a shutter must be used. In this case contiguous images are not possible and so it becomes difficult to follow changes in atmospheric phase between exposures.

The size of the image on the detector is the major consideration in the all-together scheme. This is determined by the number of pixels needed to adequately sample the fringes. The number of fringes in the image can be easily calculated. The four telescopes form six independent baselines and so produce six sets of fringes across each half of the star image. The minimum spacing between the beams must be twice the fringe frequency to prevent overlaps in the Fourier transform. Assuming 4 pixels/fringe for good sampling, this gives a total of 96 pixels across the diameter of the circular star image. This means that an area containing almost 10,000 pixels must be read-out.. Since there is a rms noise of around 25 electrons associated with each read of a pixel, then in the total image there would be 10^5 electrons of noise. To obtain a reasonable signal to noise a large number of photons would have to be integrated in each exposure. In addition to this it would take around 0.5 seconds to read each image. Since the individual exposures would be only 10 - 20 ms, the overall observing efficiency is low.

These problems can be reduced by compressing the image into fewer pixels on the detector. The aperture masking experiments on the WHT use a CCD camera. This device can use an on-chip binning technique to compress the 2-dimensional star image into a single line before readout. As is explained in chapter 2, this is not possible with the architecture of the NICMOS array. It may be possible to compress the image in one dimension into at best, a few lines of pixels using cylindrical lenses but his would require complex optics inside the dewar.

It is possible to improve on this simple scheme. The number of pixels in the image can be reduced if the incoming beams are arranged in a two-dimensional pattern rather than a line. Then the fringes are produced at different angles across the image. Since the fringes from different baselines can be separated by the different angles, they can be more tightly packed on the chip. It would be possible to recover the same information in the four beams from a region only 32 x 32 pixels at the same 4 pixels/fringe sampling. This system would need to perform a 2d Fourier transform of each exposure which is computationally expensive. A number of optical designs were produced to achieve this scheme.

The simplest arrangement uses a parabolic mirror shown in the diagram below. Light from the beam combiner is steered into the main mirror by a small flat mirror. The parabola brings the beams to a common focus on the detector.

The diagram in Figure 4.2 shows a possible arrangement using a parabolic mirror. If the beam diameter is 25 mm then the main mirror can be either a single 100 mm diameter f/2 mirror or several 25 mm off-axis parabola's. Both of the arrangements are shown in the diagram. The main drawbacks of this scheme are the expense and difficulty in producing the large fast parabolas and the tight space constraints for fitting the steering mirrors.

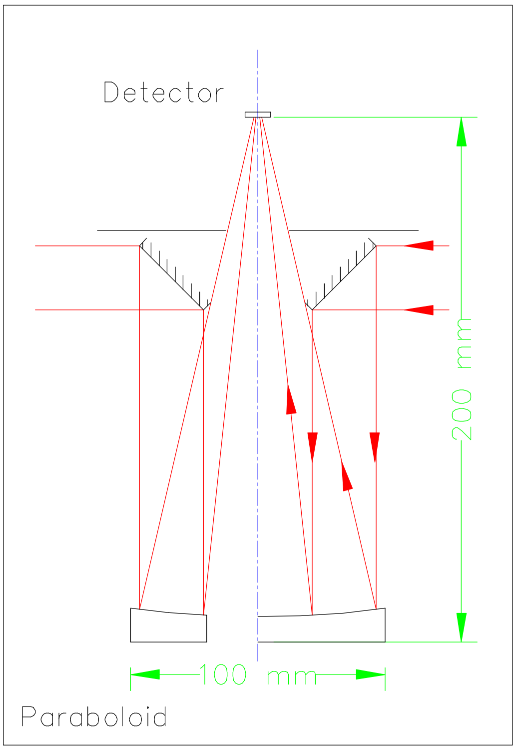

Figure 4.2: Image plane combination with a parabolic mirror

The next variation uses only spherical surfaces. This is a Schwarzschild system where two spherical mirrors are placed on axis with a common centre of curvature. If the radii are in the correct ratio this arrangement produces an image with greatly reduced spherical aberration. The focal length is equal to half the separation of the mirrors, Korsch(1991). The diagram in Figure 4.3 shows the system using either a single 400 mm diameter spherical mirror, or four separate 150 mm diameter mirrors (M3). The light from the path compensator can enter either from the bottom, or from a flat steering mirror (M1).

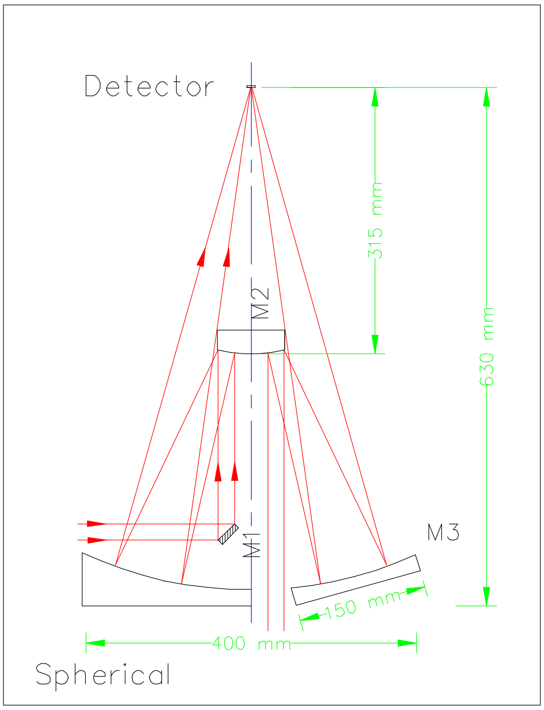

Figure 4.3: Image plane combination with spherical mirrors

A problem common to each of these schemes is that the number of fringes across the image, and so the number of pixels needed to resolve they, are independent of wavelength. However the size of the stellar image on the detector is proportional to the observing wavelength. This means that the optics would need a different focal length for each atmospheric band. Several designs of camera optics were produced which would image the fringes on the detector at the correct scale. Each of the proposed solutions involved complex optical surfaces and would need replacing or re-aligning for different infrared bands.

4.2.2 Pupil plane, pair-wise

The existing visible COAST system uses a pair-wise pupil plane combination scheme. The four path-compensated beams from the telescopes are first combined as two pairs. The two beams in a pair meet at the surface of a beam-splitter after travelling exactly the same path distance. The beam-splitter is designed with a coating which reflects half of the incident light while transmitting the other half. The reflected and transmitted output beams of light now contain information from each telescope in the pair. These beams are then further mixed by pairs in another two beam-splitters. The result of this is that the light from each of the four telescopes is now mixed in the four output beams.

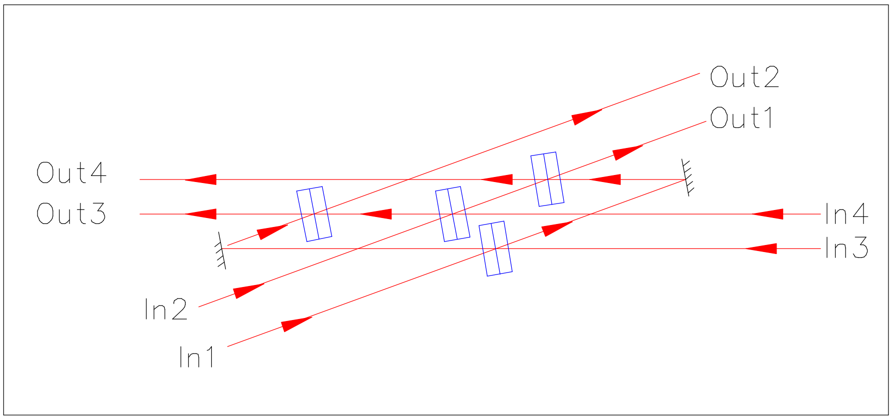

Figure 4.4: Pupil plane combination scheme

Any one of the beams contains all the information about the signal arriving at all the telescopes. Although in practice all four of the beams are measured to provide greater signal to noise. By then subtracting appropriate pairs of output beams any effects common to both telescopes, such as scintillation, can be removed. The fringes are measured by focusing each of the output beams onto a separate single element detector. By imposing a swath waveform on the trolley position it can be scanned backward and forward, through the fringe envelope. A series of maxima and minima are then swept past the detector. The fringe visibility is measured from this time resolved signal on the detector.

The signal due to all baselines is superimposed on each of the outputs. But by scanning the four path trolleys at different non-redundant rates, the fringe pattern due to each baseline can be separated.

The principle disadvantage of this system is that it involves a large number of optical components which must be aligned to a high precision. This scheme is only possible if practical beam-splitter units can be produced. The photons interfere in a beam-splitter which is made from a multi-layer film deposited on a flat substrate. This coating must be capable of reflecting half of the light and transmitting half, independently of the wavelength and polarisation. Since the beam-splitter units are aligned to a high accuracy it would be inconvenient to have to replace them for each wavelength band and so they must be designed to work over a wide range of wavelengths, λ = 1 - 2.5 μm. It was possible to meet these requirements and a beam-splitter unit was built using coatings produced by COMAR Instruments, Cambridge which can operate in the J, H and K bands. The performance of this unit is shown in Figure 4.13.

4.2.3 Fringe tracking

To operate at its highest performance COAST will need to actively track the centre of the fringe envelope as the position wanders due to atmospheric conditions. It will do this by monitoring the change in visibility on each baseline and introduce extra delay in the path compensator to correct for the extra atmospheric path. The pupil plane system has the advantage that each measurement of the fringe visibility is available immediately. This allows the controlling system to follow a single fringe. The image plane scheme only produces an image once every few t_0 and so can only follow slow changes in the fringe position.

4.2.4 Conclusion

The final system uses a pupil-plane , pair-wise, combination scheme. Although the losses are higher than an image plane system, the reduced read noise increases the signal to noise by a large factor. Another major advantage of the pupil-plane scheme is that it uses a very similar design to the existing visible system. This allows the data analysis and control systems to treat the visible and infrared data identically and avoids duplicating the analysis software. The beam combination systems are similar enough that if any problems or unusual effects arise the solutions will be applicable to both.

4.3 COAST structure

This section describes each part of the infrared beam combiner in detail. It examines the problems encountered in the design and constructions and some of the solutions attempted.

4.3.1 General Layout

The diagram in Figure 4.5 shows a plan of the main components inside the COAST optics lab. The light from the telescopes enters at the top of the page. The first table running horizontally across the top of the diagram contain the visible system and the autoguider. The long table on the left carries the path compensation trolleys and extends for 25 m down the building. The infrared system is built on the optical table on the right.

The outlines of the tables and the mirrors are shown in black. The beam-splitters are blue and the dichroic filters green. The beams of white light from the telescopes are shown in white, the red beams are infrared, pink is the red visible light and blue is the blue light used for the autoguider. The components and the beams are drawn at their correct relative size in the instrument.

The light from the four telescopes enters the optics lab at the top of the diagram in Figure 4.5. The beams pass over the top of the beam combiner system and go the path compensation trolleys. Each trolley has two flat mirrors at right angles as shown in Figure 4.12, which reflects the light back toward the beam combiner at the lower level. A commercial laser interferometer measures the position of a corner cube on the front of the trolley. The double mirror arrangement both doubles the path delay possible in the length of rail and makes the system insensitive to vertical tilts in the trolley motion.

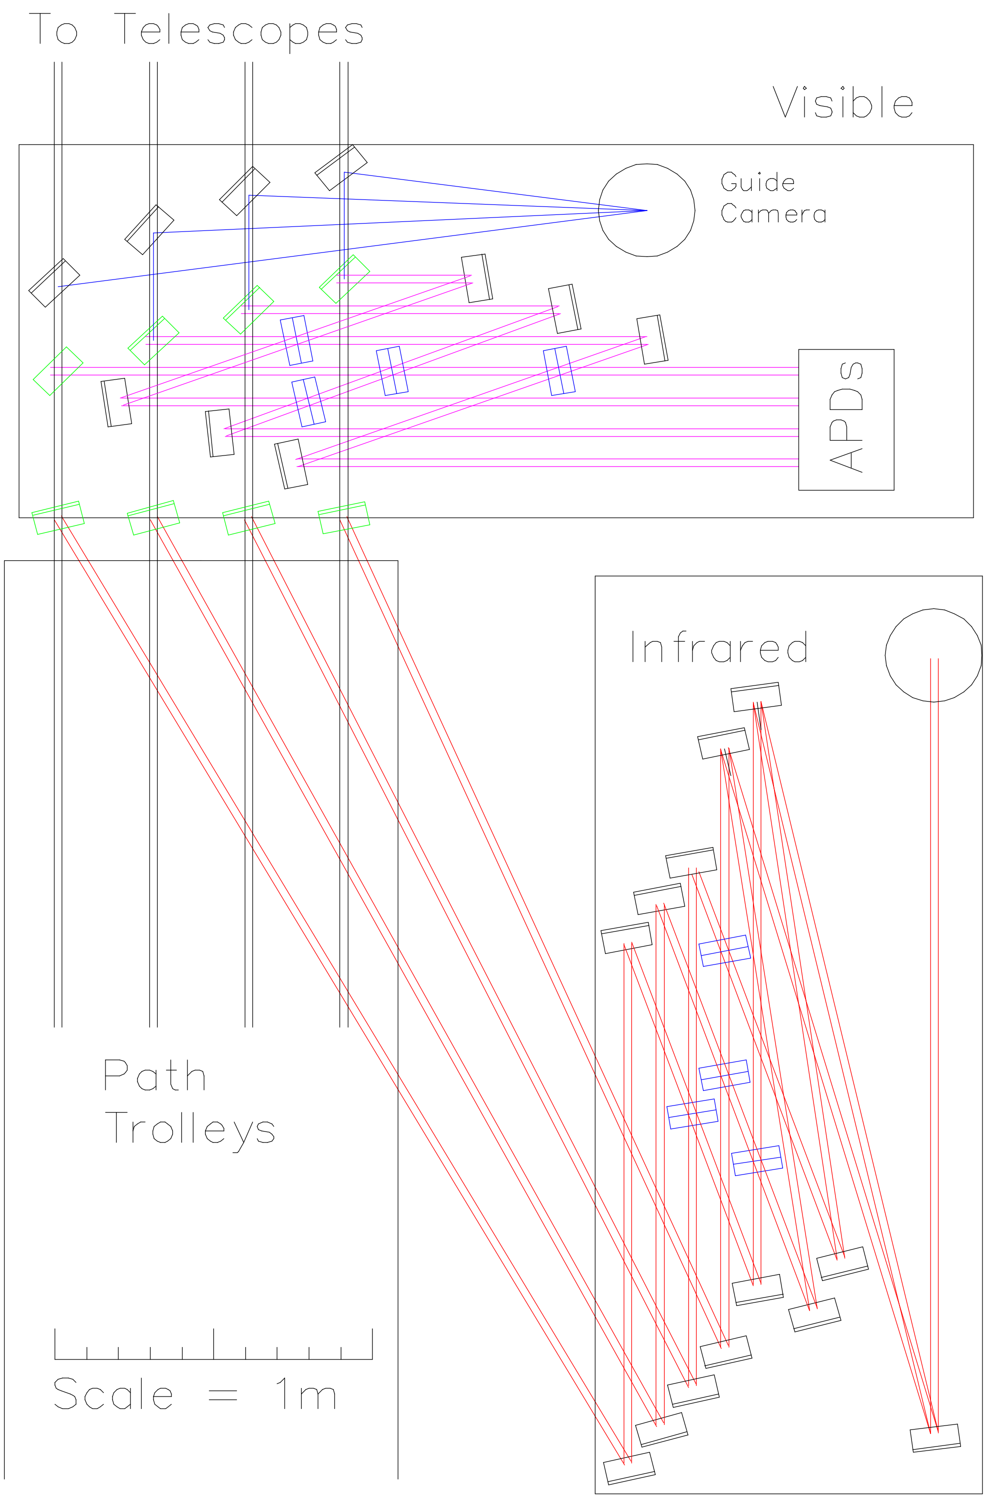

Figure 4.5: Layout of COAST

The white light reflected back toward the beam combiner arrives at a set of mirrors, shown in Figure 4.5 . These have thin film gold coating which reflects 90% of the infrared light from λ = 1 - 2.5 μm while transmitting around 20-40% of the blue light into the visible system. This light is enough to operate the autoguider, but to allow visible observations these four mirrors slide out of the beam.

The light transmitted by the gold mirrors reaches four dichroic filters which transmit blue light, λ < 650 nm, into the autoguider. The red light is reflected into the visible beam combiner. During visible observations only the blue light is used for the autoguider because the red light is needed for the science observations. However the CCD camera used for the autoguider is sensitive to light up to λ = 1μm. For infrared observations, these dichroics could be removed to allow all the visible light to the autoguider.

The majority of the infrared light is reflected by the gold mirrors onto another optical table alongside the path compensation system. Here it is collected by the four mirrors labelled M1-4 in Figure 4.8.

The infrared beam combiner interferes the light from each of the four telescope with the light from the other three telescopes. The four outputs from this unit are then imaged by the infrared camera. The details of this system are described in the next section.

The photographs below show the telescopes in their shortest baseline configuration (Figure 4.6) and a view of the inside of the COAST optics lab (Figure 4.7).

Figure 4.6: View of telescopes

Figure 4.7: Inside the COAST optics lab

4.3.2 Infrared beam combiner

The picture in Figure 4.9 shows a view of the infrared beam combining table, the diagram below shows the paths through the infrared beam combiner in more detail.

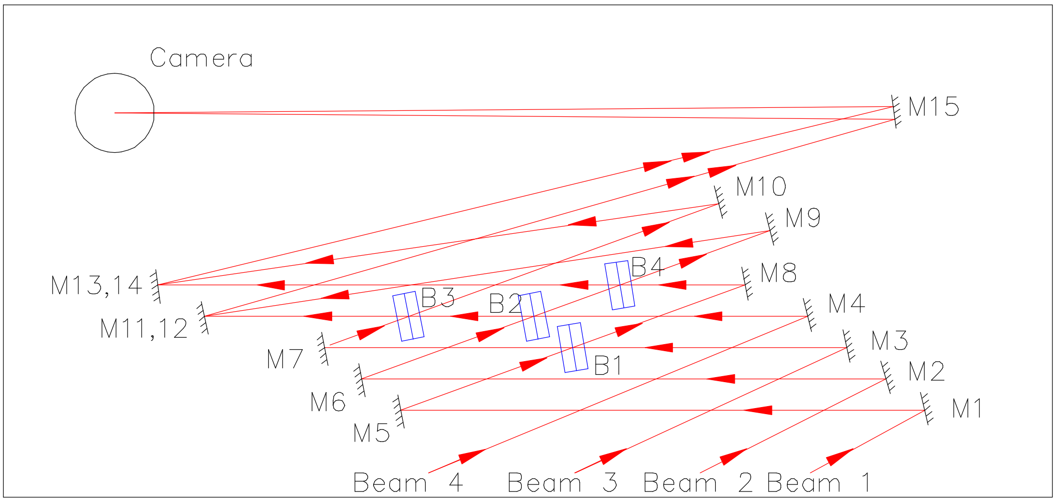

Figure 4.8: Infrared beam combiner

The beams of light from the path compensator arrive at the infrared system and are reflected into the beam combiner by mirrors M1-4. By having two independent mirrors for each beam, one gold dichroic mirror on the visible table and one steering mirror on the infrared table, there are enough degrees of freedom to correct for any difference in the height or alignment of the tables.

The optical tables have a matrix of holes drilled in the top surface, these are used for bolting down the optical components but they can also hold targets used to check the alignment. The four beams enter the beam combiner in parallel each positioned along a line of these holes which allows this part of the alignment to be checked easily. However once the beams have passed through a beam-splitter there is a lateral displacement and this facility is lost. The optical path through the beam-splitter is shown in Figure 4.11.

Figure 4.9: View of Infrared beam combiner

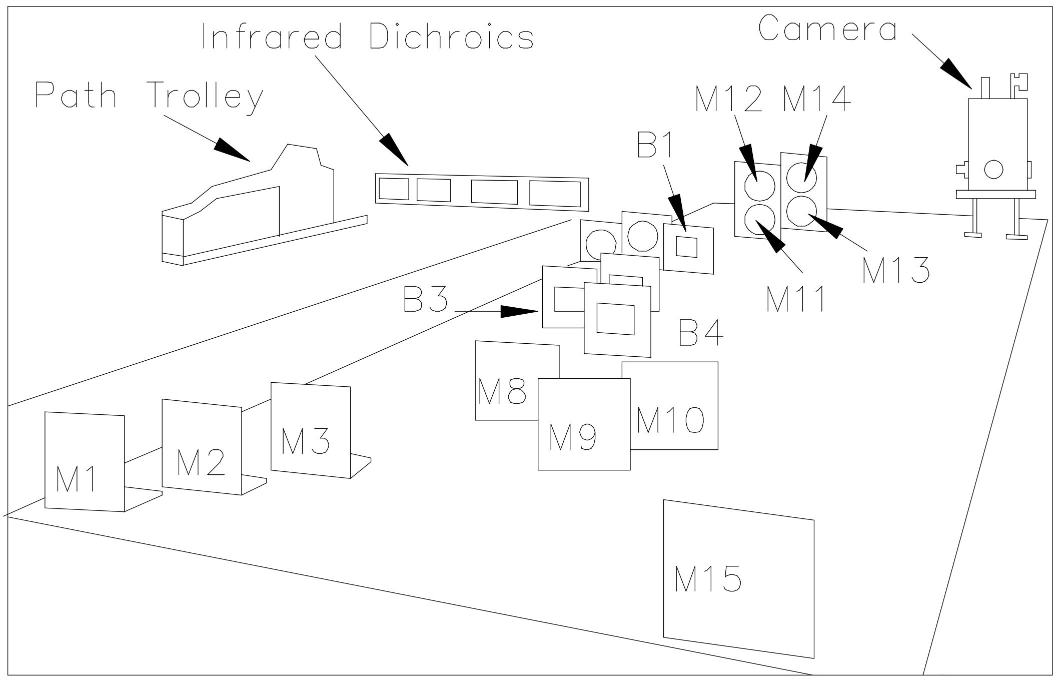

Figure 4.10: Diagram of picture

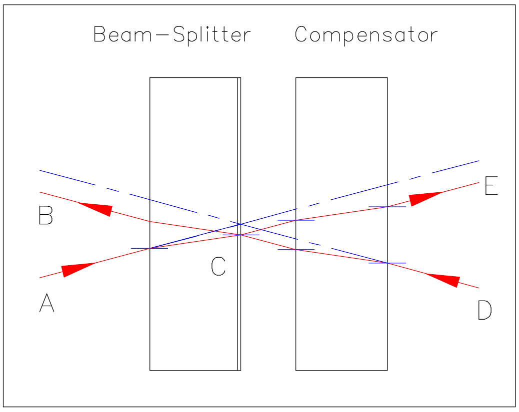

Figure 4.11: Beam-Splitter

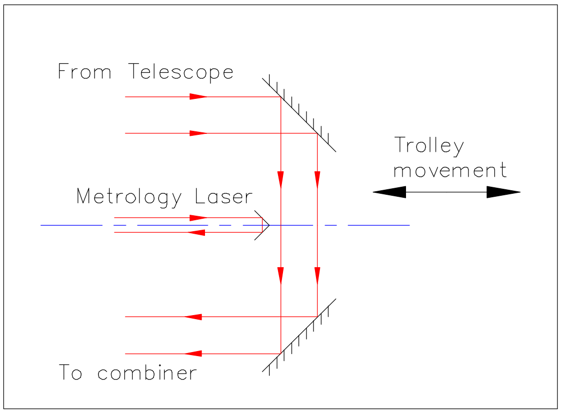

Figure 4.12: Path Trolley

4.3.3 Beam-Splitter design

The diagram in Figure 4.11 shows a COAST beam-splitter unit. It consists of two identical flat parallel sided substrates. The inner face of one component has the beam-splitter coating while all the other faces are anti-reflection coated. Over the wavelength range $\lambda = 1 - 2.5$ μm, the beam-splitter coating is designed to reflect 50% of the incident light and transmit 50% , without effecting the polarisation. The design is only correct for an angle of incidence of 15° and the components are all aligned to within one degree of this. The performance of the beam-splitter and anti-reflection coatings are shown in the plot in Figure 4.13 .

After being reflected from the beam-splitting surface, the ray from A to B would have passed through twice the thickness of the substrate material while the ray D to E would have only passed through air. This would introduce a wavelength dispersion which reduces the fringe visibility. The compensating plate avoids this by having all the beams pass through the same thickness of material. The beam-splitter and compensating plate are mounted in a single unit which allows the angle and position of the beam-splitter to be adjusted by three fine pitched screws. The compensating plate remains fixed since the exact angle between the two is not critical. The beam-splitter holder is designed so that the substrate can be removed and replaced kinematically.

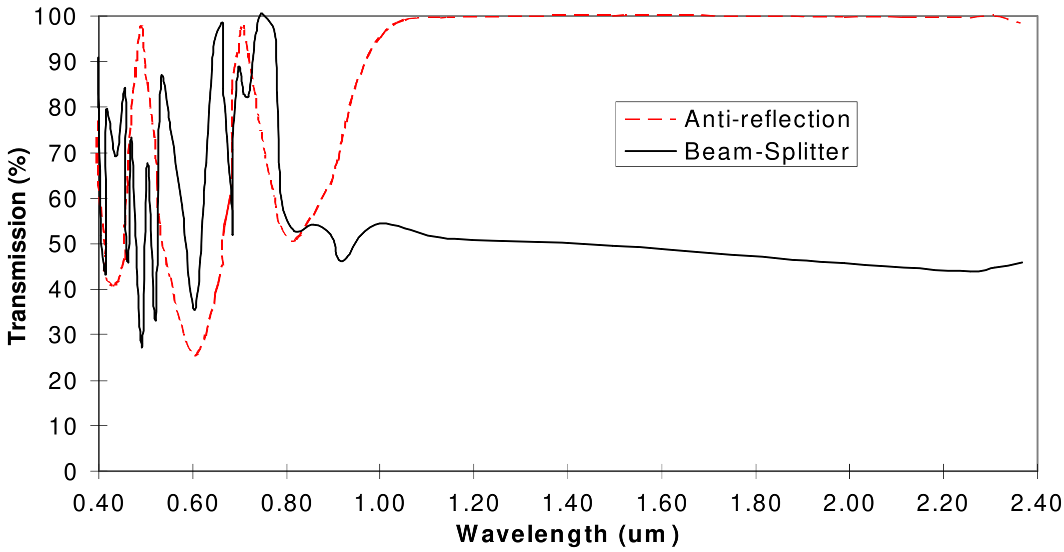

Figure 4.13: Performance of beam-splitter and anti-reflection coatings

4.3.4 Internal calibration and testing

It is convenient to be able to align and test the system during day-light and in all weather conditions. To do this, light an artificial star can be in switched into the system. The light-path through the beam combiner optics is the same in both directions and so if light is introduced back into the system from one of the output beams, then it will appear at all the telescopes. This can be used to align and a test the components. An artificial star is produced from an illuminated pinhole collimated with a lens. A removable mirror between B3 and M11 in Figure 4.8 reflects the light from the artificial star into the combining system where it emerges at each of the telescopes. The telescope siderostats are then driven to a limit position where they reflect the light back into the instrument. Alternatively there is an optical bench inside the optics lab where four mirrors can be installed to reflect the light back before it reaches the telescopes. This allows high visibility fringes to be conveniently produced at any time and under any conditions.

One of the main difficulties in the construction and operation of an interferometer is the need to equalise all the optical paths to within a few wavelengths. The path compensation system removes the large path differences due to the different distances from the telescopes to the star. There is also a path difference due to the position of each the components in the COAST lab. For example, light in beams 1 and 3 first meet at beam-splitter B1 in Figure 4.8 . To reach this point beam 3 is reflected from mirror M3 while beam 1 takes the route M1-M5-B1. This extra path is removed by adding a fixed offset to the position of the trolley in beam 3. Similar corrections are used with each of the other beams. The path error is found by accurately measuring the position of all the components and then calculating the position of the path trolley to remove this. The trolleys are then swept slowly through the expected fringe envelope position while looking for fringes at the output. Since the path through all the components can only be measured to within a centimetre, and the fringe packet is only around 10 wavelengths long, this can take a little time.

The value of these corrections are different for the visible and infrared systems because of the different position of the optical tables. This is one of the problems which prevent simultaneous observations at visible and infrared wavelengths.

4.3.5 Output arrangement

The simplified diagram in Figure 4.4 shows the position of the four output beams from the correlator. Each of these beams is focused onto a separate single detector element. In the visible system this is relatively simple since the detectors have a fibre-optic pigtail which can be conveniently positioned at the output beams. In the infrared scheme the detector elements are individual pixels on a single integrated circuit. Furthermore, the detector chip is contained inside a dewar which can only allow a narrow field of view. The solution is to use a lens immediately in front of the dewar. All the four beams pass through the centre of the lens at different angles and are focused on four separate pixels. Since only a single pixel is read, the signal to noise is maximised if all the light can be focused into that pixel. The beams passing through the lens off-axis would suffer coma and spherical aberration, which would spread the light out on the detector. Both these effects increase with the off-axis angle and so the optimal design has the beams as close to the axis as possible.

The most convenient arrangement is to have the four beams in a square converging on the centre of the lens. This produces the four output spots in the four corners of the detector. Since the parallel beams are to be focused to a spot, the detector must be at the focal plane of the lens. The focal length of the lens is determined by the size of the detector,

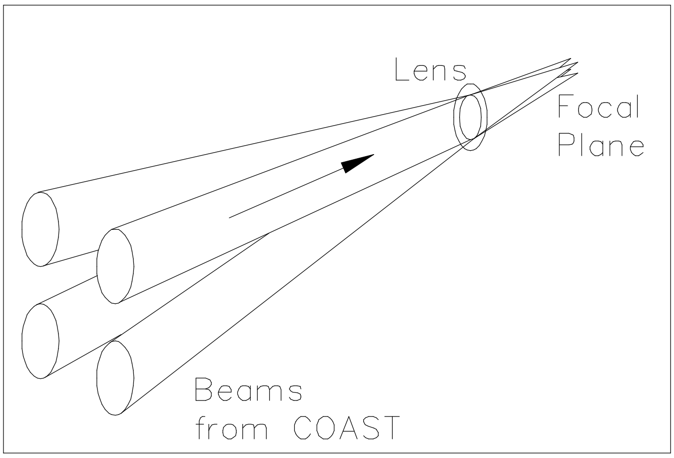

Figure 4.14: Beam imaging arrangement

A longer focal length would allow the lens to be placed further from the detector. The limit on the focal length is limited by diffraction which produces a larger spot size on the detector. At the longest infrared wavelength, $\lambda = 2.2\ \mu m$, a focal length of 180 mm would put the first minimum of the diffraction pattern at the edge of the pixel. A focal length longer than this would reduce the amount light which could be focused into the central pixel. The system uses a focal length of 143 mm which produces outputs spots separated by 90 pixels on the detector. The incoming beams are about 0.7° from the axis. To achieve the correct separation for the beams entering the lens they must be reflected from M15 back along the table to M11-M14. This double pass gives a path of 5 m at which the beams are separated by 150 mm. The mirrors M11-M12 and M13-M14 in Figure 4.8 are double mirrors with M12 and M14 immediately above M11 and M13 respectively. This can be seen more clearly in Figure 4.10. These four separate mirrors allow the position of each of the spots on the detector to be individually adjusted to position it exactly over one pixel. The position of all four spots can be moved together with M15.

4.3.6 Alignment

The system is aligned with the aid of a theodolite mounted on an optical bench just inside the building. This looks into the system along the same line that the light from the telescope would take and views an illuminated target which can be placed at the position of each component. Initially the target is positioned immediately in front of M1 and the angle of the dichroic mirror feeding this beam adjusted. This sets the correct position of the beam on M1. The target is then moved down the table to just in front of M2. Then angle of M1 can now be adjusted until the target is centred again. This process is iterated until the target appears in the same place at both positions. This step is then repeated along every path through the system.

The paths of each beam through the correlator are shown below, the names of the components are from Figure 4.8. In total there are 14 mirrors and four beam-splitters, each time the beam reaches a beam-splitter the number of possible paths through the system doubles. Since the movement of each component generally effects the alignment of several other units the process of aligning the system is time consuming.

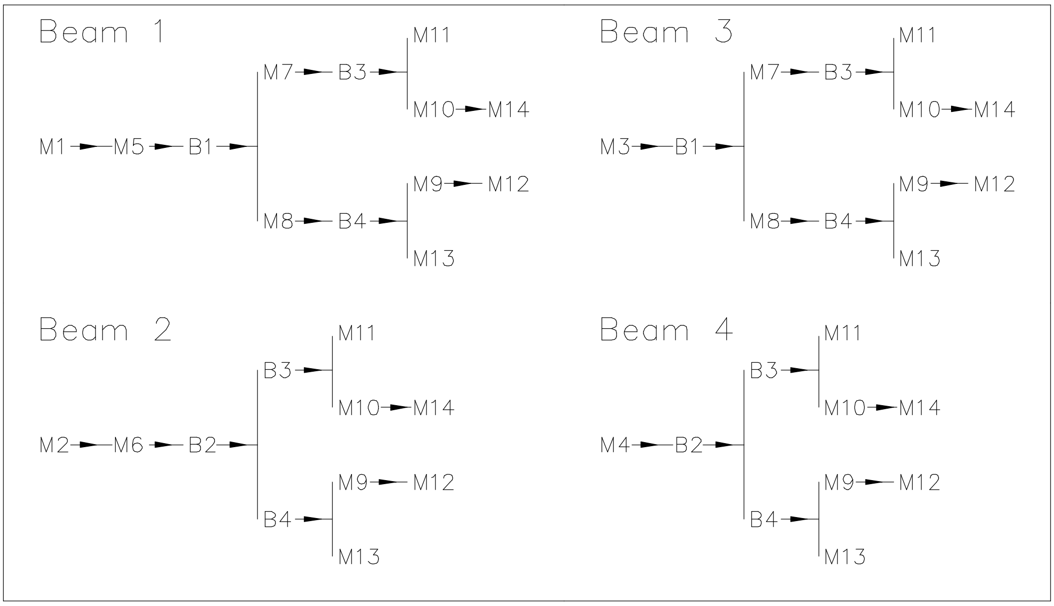

Figure 4.15: List of paths through combiner

The alignment becomes more complicated when the light passes through a beam-splitter. The beam-splitter and anti-reflection coatings were designed for the infrared, λ=1 -2.5 μm, the behaviour of the coatings outside these bands are undefined. Looking through the units in visible light the anti-reflection coatings are partly reflective. This means that when the system was initially aligned by eye, the beams were reflecting from the compensating plate rather than the beam-splitter and all the components were systematically mis-aligned. There are two solutions to this problem.

When the original visible system was built a spare set of beam-splitter plates were made without anti-reflection coatings. The beam-splitters units are designed so that the beam-splitter plate can be replaced kinematically and so the system can be initially aligned with visible light and then the infrared units put back in.

The curves in Figure 4.13 show the performance of the beam-splitter and anti-reflection coatings across the visible and infrared bands. There are a number of wavelengths in the visible region where coincidentally the coatings behaviour almost correctly. Since the beam combiner consists only of flat reflecting surfaces its alignment is independent of wavelength. And so the system can be aligned in one of these narrow wavelength bands. As much less light was now available, a sensitive low-light video camera was used with the theodolite to observe the target.

4.4 Efficiency - Performance

The COAST system ultimately measures the visibility of fringes produced by interfering beams of light from pairs of telescopes. The signal to noise of this measurement depends on the magnitude of the signal and the visibility of the fringes. Both of these quantities depend on the astronomical object and effects in the instrument. To allow measurements to be made of the faintest objects, any effects in the instrument which reduce the flux or visibility must be minimised.

4.4.1 Flux losses.

A major disadvantage of a pupil plane scheme using beam-splitters is the large number of optical surfaces involved. In the COAST system the light from the star encounters over 20 optical surfaces before arriving at the detector. Each reflecting surfaces in the COAST system causes some loss of flux due to absorption or scattering out of the beam. The overall loss through the system can be estimated by simply combining the losses in each component. Since some of the losses are wavelength dependant I will carry out the calculation for each infrared band, and for comparison the visible band as well.
The schematic diagram below summarises the losses in the calculation.

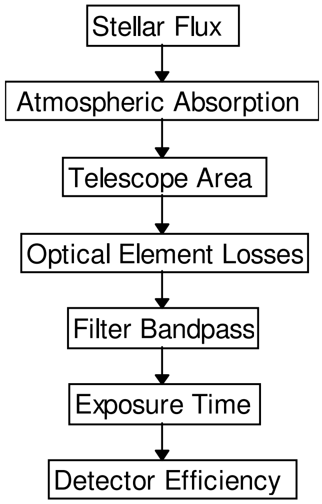

Figure 4.16: Flux losses in the system

4.4.1.1 Stellar Flux

The table below lists the flux at the top of the atmosphere from a star of apparent magnitude 0, from Wamsteker(1981). The flux appears lower at longer infrared wavelengths because the definition of 0 magnitude is based on α Lyr (Vega) which is a relatively blue star.

Band	V	J	H	K
Wavelength (μm)	0.55	1.2	1.6	2.2
Flux (W/m^2/μm)	3.64×10^{-8}	3.18×10^{-9}	1.18×10^{-9}	4.17×10^{-10}
Photons/s/m^2/μm	1.0×10^{11}	2.0×10^{10}	9.8×10^{9}	4.6×10^{9}

Table 4.1: Flux from 0 magnitude star

4.4.1.2 Atmospheric Absorption

The path through the atmosphere reduces the flux by both scattering and absorption. In the near infrared and at a low altitude site, the major loss is from water vapour absorption. It is absorption from the O-H bond in atmospheric water which limits observing to the discrete infrared bands shown above. Within these windows there is still some absorption, but this is highly dependant on water vapour levels and weather conditions and so is difficult to predict. Continuum extinction is normally characterised by an absorption in magnitudes proportional to air mass, where air mass is defined as *Sec*(zenith distance). At visible wavelengths a value of 0.2 mag/air mass is common. In the infrared this value is reduced since Raleigh and aerosol scattering both decrease with wavelength. In practice, to minimise the effects of atmospheric turbulence COAST will observe as near to the zenith as possible, this will also reduce absorption. Since the effects are so variable this study will ignore absorption effects.

4.4.1.3 Telescope Area

The telescope size is a compromise between the light gathering power of larger telescopes and the reduction in fringe visibility caused by atmospheric turbulence across the aperture. Theoretical studies suggest that, in the case of photon-noise limited observations. the optimum aperture is $2.8r_0$ for tip-tilt corrected telescopes, Buscher(1988). When used in the visible bands the 0.4m aperture telescopes are normally operated with masks on the primary mirrors to reduce the aperture to this factor. Using the result that r_0 scales as $\lambda^{6/5}$. and assuming a value of 10 cm for r_0 in the visible, then the telescopes will be smaller than the optimum size in all the infrared bands. The telescopes will be operated at their full aperture of 0.4 m.

Band	V	J	H	K
r_0 (m)	0.10	0.27	0.37	0.53
Optimal aperture (m)	0.29	0.75	1.00	1.50

Table 4.2: Optimal Telescope Aperture

4.4.1.4 Optical element losses.

The COAST interferometer is a complex instrument with a large number of optical surfaces. To reduce flux losses in the instrument it is designed to use reflecting elements where possible. The choice of coating for these surfaces is a compromise between durability and reflectivity.

Band	V	J	H	K
Aluminium	0.87	0.92	0.95	0.96
Silver	0.91	0.95	0.98	0.98
Gold	0.60	0.97	0.98	0.98

Table 4.3: Reflectivity of Mirror Coatings

The values in Table 4.3, are for normal incidence with ideal coatings immediately after deposition, Melles-Griot(1990). In a real instrument the performance of the coating decreases quickly as the surface corrodes. The silver mirrors are especially sensitive to atmospheric sulphur, forming a yellow deposit after only a few days. Silver mirrors with a protective overcoat of magnesium fluoride have been in use on COAST for several years without significant degradation. The values in Table 4.3 do not include scattering losses due to dust on the surface of the optics, this is difficult to quantify but could account for an extra 5% loss from each surface.

The mirrors on the telescope are difficult to replace quickly and are subject to the greatest weathering, these use Aluminium coatings. The components inside the optics lab are better protected, and can be more easily removed for re-coating. Protected silver coatings can be used for these. Gold mirrors could be used in the infrared system but this would make alignment in visible light difficult and prevent components being exchanged between the infrared and visible systems. These mirrors would still need overcoating to provide mechanical protection since the gold surfaces are extremely fragile. Alternatively, dielectric mirrors can be designed to have almost 100% reflectivity at a particular wavelength, but these are less useful for broad band applications.

There are a number of other reflecting surfaces in the COAST system which must be considered separately. The thin film gold mirrors used to split the infrared light from the blue light to the autoguider have a reflectivity of approximately 90% in the infrared bands. The piezoelectric tip-tilt mirror on the telescope uses a proprietary enhanced aluminium coating with a reflectivity of 95% in the visible, the infrared reflectivity is unknown and will be estimated as 90% for this model.

The number of transmission elements has been kept to a minimum, only a single lens, the dewar window and the filter. There is also a window at the entrance to the COAST optics lab. This has an anti-reflection coating designed for the visible which reduces transmission significantly in the infrared. This will probably have to be replaced to achieve the ultimate infrared performance. The components in the camera are a bi-convex lens approximately 6 mm thick and the dewar window which is 2 mm thick fused silica. The filter transmissions are taken from the manufactures curves and include surface reflections.

Band	J	H	K
Lens	0.85	0.85	0.85
Lab Window	0.70	0.65	0.80
Dewar Window	0.85	0.85	0.85
Filter	0.70	0.85	0.75

Table 4.4: External transmission of optical components

The dewar window and lens are not anti-reflection coated and so there is a loss of approx. 6% on each of the four faces. As the design of lens and window is fixed, these components could be given an anti-reflection (AR) coating. If the same AR coating designed for the beam-splitters , see Figure 4.13, was used on the other transmission elements then the losses could be reduced to almost zero.

By considering the component losses in each part of the instrument and the number of optical components, the following losses can be calculated.

Telescope; 2 aluminium mirrors, 1 enhanced aluminium, 1 silver mirror

Path delay; 2 silver mirrors, 1 gold partial mirror

Interferometer; 4 silver mirrors, (assumes beam-splitters are loss-less)

Camera; 2 aluminium mirrors, lens, window, filter

Band	J	H	K
Reflections	0.47	0.57	0.57
Transmission	0.42	0.47	0.50
Total efficiency	0.20	0.28	0.30

Table 4.5: Total efficiency of optical system

4.4.1.5 Detector

Of the photons which arrive at the camera, the number detected depends on the exposure time, the filter bandwidth and the detector quantum efficiency.

The detector quantum efficiency at K band is taken from the manufacture's data Rockwell(1990). The results for shorter bands are extrapolated from response curves published for the device.

Band	J	H	K
Quantum Efficiency	0.50	0.60	0.60

Table 4.6: Quantum efficiency of NICMOS detector

All interferometers are limited by the small telescopes, short exposure time and narrow bandwidths needed to produce high visibility fringes. In the infrared bands COAST can use the full atmospheric window. These values are for the filter set used at COAST which is based on the J, H, K photometry standards.

Band	J	H	K
Bandwidth (μm)	0.26	0.31	0.33

Table 4.7: Infrared filter bandwidth

4.4.1.6 Conclusion - Flux losses

Combining these results we can calculate the flux from a magnitude 0 star measured at the detector. These values are for a single telescope and one of the four outputs of the correlator.

Band	J	H	K
Photons / second	1.6×10^7	1.5×10^7	8.4×10^6

Table 4.8: Predicted photons detected from a magnitude 0 star

Observational test

A series of observations were made to test these predictions. The tests were in two forms, long exposure conventional images were taken of the outputs of the beam combiner. And measurements were made in the normal observing mode of COAST with the camera taking a series of short integrations with a single element.

The measurements were made of Vega (α Lyr) on August 21st 1995. This star was chosen since it is defined as having an apparent magnitude of 0 in all wavelength bands. At the time of the observation the star was almost overhead and so atmospheric absorption was minimised.

The output of the beam combiner consists of four separate beams which are measured by the detector. Ideally all the signal in each beam can be focused onto a single pixel. In practice at the time of the observation the system was not perfectly aligned or focused and so some of the light is spread out into the surrounding pixels. In the photometric measurements the flux from each image is measured over a patch large enough to contain all the light from the star. The proportion of light in the central pixel was also calculated. The data was obtained using 2 outputs of the beam combiner and one telescope. A series of exposures were made at each integration time and the experiment was repeated using each of the three telescopes. The exposure times were 50, 100 and 200 milli-seconds.

Results

	Signal (e^-/s)	Fraction in peak	Signal (e^- / s)	Fraction in peak
Telescope	Output 1		Output 2	
1	5478000	0.22	7240000	0.23
2	7858000	0.27	7208000	0.19
3	6214000	0.35	8841000	0.20
Mean	6.5×10^6	0.28	7.7×10^6	0.21

Table 4.9: Photometry of Vega

The measured signal was 6.5×10^6 and 7.7×10^6 e^-/second, with about $\frac{1}{4}$ of the light in the central pixel. The variation in the flux measured from each telescope, and on each output is largely due to miss-alignments. The predicted flux from Table 4.8, is 16×10^6 e^-/second. The missing signal is due to a combination of miss-alignments in the optics and atmospheric absorption, possible because of high-level cloud. But overall this result suggests that there are no major problems with the instrument calibration.

4.4.2 Visibility losses.

The visibility of the fringes from a particular baseline describes how resolved the source is on that baseline and so is a measurement of the angular size of the object.

Visibility can be interpreted as:

$$V = \frac{I_{Max} - I_{Min}}{I_{Max} + I_{Min}}$$

Where I_{Max} and I_{Min} are the signal at the maximum and minimum of the fringe pattern.

An unresolved source and an ideal instrument would produce a fringe visibility of 100%. In a real telescope the maximum visibility is degraded by wavefront errors introduced by the atmosphere and the optical surfaces.

During an astronomical observation the telescope would observe the target object and then a nearby calibration object which is known to be unresolved. Since the theoretical visibility of the calibrator is 100% the observed visibility on the target will be divided by the visibility of the calibrator to correct for the effects of atmospheric and instrumental visibility loss.

4.4.2.1 Atmosphere

The atmosphere has a large effect on the fringe visibility. Atmospheric turbulence effects introduce phase changes across the originally flat wavefront from the source. The model commonly used to describe these effects, defines a length r_0 over which the phase errors are small. The COAST system uses telescopes with an aperture (D) of approximately this size Larger apertures would increase the signal flux but reduce the visibility of the fringes.

The relationship between decreasing fringe visibility and increasing aperture can be determined by theoretical models. These suggest that for photon noise limited observations there is always a slight gain with increasing telescope size. This is especially true for infrared detectors with a large read noise, where a larger signal flux is needed to achieve photon-noise limited operation.

In practice much of the phase error introduced by the atmosphere is in the form of a mean tilt across the telescope aperture. If these first order tilts are removed by servo controlled mirrors on each telescope, then the usable telescope area for the same loss of visibility can be increased. In the case of tilt corrected apertures there is very little loss of visibility in the region where $D \leq r_0$. As the size of the aperture increases the visibility is reduced, but there is a large enough improvement in the signal to make this worthwhile. When $D/r_0 > 3$ the image breaks up into individual speckles and the advantages of tip-tilt correction are reduced. The relationship between visibility and aperture size is complicated and must be analysed numerically, these values are from Fig. 3.6 in Buscher(1988).

Band	V	J	H	K
r_0 (m)	0.10	0.27	0.37	0.53
D/r_0 (for D=0.4 m)	4	1.481	1.081	0.755
Visibility (with tip-tilt)	0.35	0.80	0.85	0.90
Visibility (no tip-tilt)	0.15	0.42	0.52	0.65

Table 4.10: Atmospheric visibility Loss

4.4.2.2 Optical Surfaces

Irregularities in the optical surfaces add a small random phase error across the wavefront. Assuming that the surface errors in each component are uncorrelated, they will add together in quadrature.

The mirrors and beam-splitters are all figured to λ/20 at 500 nm. Each beam from the telescope undergoes 14 reflections from mirrors. It then passes through 2 beam-splitters with a total of 8 surfaces. The phase error introduced at each surface reduces the visibility by a factor: Buscher(1988)

$$V = 1 - \frac{1}{2}\left(\frac{2\pi\delta}{\lambda}\right)^2$$

Band	V	J	H	K
Mirrors	0.85	0.97	0.98	0.99
Beam-splitters	0.88	0.98	0.99	0.99
Total visibility	0.81	0.96	0.98	0.99

Table 4.11: Instrumental visibility losses

These results suggest that almost perfect visibility fringes can be obtained. In practice using an internal test source a maximum fringe visibility of 60 - 70 % has been achieved. The analysis does not include the visibility loss from miss-aligned components introducing a tilt across the beams.

4.4.2.3 Fringe measurement

The fringe visibility is determined by sweeping the fringe pattern past a single element of the detector array. The path compensator will sweep through the entire fringe envelope in a time much less then the atmospheric coherence time $t_{0.}$ The camera makes a number of measurements of the power in each beam, continually integrating the signal from each beam and then reading the value at regular intervals. If four measurements are made of each fringe the estimate of the visibility is reduced by only 10%, see equation 6.4 Buscher(1988).

4.4.2.4 Conclusion - visibility losses

Combining the atmospheric and instrumental visibility losses described above, the visibility measured from of an unresolved source can be estimated.

It is not necessary to know this value accurately since the true visibility loss will be determined by observing a calibration object immediately after observing the source. It is only necessary that the visibility be has high as possible for increased signal to noise.

Band	J	H	K
Maximum Visibility	0.70	0.75	0.80

Table 4.12: Predicted visibility measurements of a point source

4.4.3 Limiting magnitude

The limiting magnitude of COAST in the infrared can be determined from the minimum number of photons needed to measure a fringe visibility. Using the flux and visibility losses described above and the properties of the detector this minimum number of photons

can be estimated. In this analysis I shall define the limiting magnitude as the ability to measure the visibility modulus from a single fringe packet with a signal to noise of one.
The main source of noise in the infrared system is the detector readout noise, this is an rms uncertainty present in each integration, independent of the signal level.
The form of the signal-to-noise with a real detector has been analysed in Nightingale(1991) for the case of a spatially resolved fringe pattern on an array detector. However, in the COAST infrared system the fringe visibility is measured by sweeping the fringe envelope passed a single pixel and a making a series of integrations. In the region where the detector readout noise dominates, this equation is still a good approximation.

$$SNR = \frac{V^2 N^2}{\sqrt{(N+D)^2 + 2(N+D)N^2V^2 + 2n^2\sigma^4}}$$

Where; N is the number of astronomical photons detected.
V is the fringe visibility at the detector
D is the number of photons from dark current and background sources
n is the number of pixels in the image
σ is the read noise on each integration

The result above assumes that the visibility is measured by taking a Fourier transform of the fringe envelope and measuring the power contained at the fringe frequency. By averaging in the Fourier domain and random phase differences between fringe packets are removed.

This case is probably the most conservative estimate since it uses little information about the fringe packet. Methods for analysing the visibility based on Maximum-Entropy and Bayesian statistical techniques may allow a measurement to be extracted from much lower signal levels.

The dark current of 100 e^-/s is negligible in the very short (1 ms) integrations used at COAST, and so the equation simplifies to:

$$SNR = \frac{V^2 N^2}{\sqrt{N^2 + 2N^3 V^2 + R}}$$

Where R is the readout noise term, $R = 2n^2\sigma^4$

Using typical values for COAST in the infrared the limiting magnitude can be estimated.
The infrared filters have a 20% bandwidth which produces 10 fringes in the centre of the fringe pattern. To measure these fringes without significant loss of visibility, four samples/fringe are needed. This gives a total of 40 integrations in each fringe packet.
The readout noise of the NICMOS camera is approximately 16 electrons/pixel, but two reads are needed for each integration and so this rises to around 25 e^-/pixel.
For an unresolved source, and taking a limiting signal to noise of one, the minimum number of photons needed in each fringe packet is:

Band	J	H	K
Limiting signal/packet	37	32	29

Table 4.13: Limiting signal in a fringe packet

Using the signal received from a 0 magnitude star in Table 4.8, then the limiting magnitude can be derived from the relationship.

$$m_{Limit} = 2.5 Log(\frac{n_{Limit}}{n_{Mag0}})$$

Band	J	H	K
Limit (tip-tilt)	12.1	12.2	11.7
Limit (no tip-tilt)	11.5	12.0	11.3

Table 4.14: Limiting magnitude for COAST in the infrared

4.5 Conclusion

The limiting factor in real observations will almost certainly be the autoguider and tilt correction system. Since the objects observed with COAST are likely to be intrinsically red even less visible light will be available for the autoguider. The limiting magnitude is also estimated for the case where tip-tilt correction is no longer possible. It can be seen from the results in Table 4.14 that the limiting magnitude does not increase as strongly with wavelength as the $\lambda^{4.6}$ suggested in chapter one. This is largely due to the fixed small size of the telescope which could ideally be increased at longer wavelengths. Another factor comes from the definition of the magnitude system which is based on a blue object (α Lyr), many of the objects which are likely to be observed with COAST are much more luminous at longer wavelengths.