Chapter 6

Interferometry: The basic principles

We have seen that the size of the telescope sets a limit on the spatial resolution of our images. There is a practical limit to telescope sizes, which would mean that we would never be able to achieve resolutions beyond that limit. The solution to this technical problem is to use the technique of *interferometry*.

In this chapter we will discuss the basic concepts of interferometry and its various incarnations. We will also discuss some of the basics of coherence of light. The topic of interferometry is, however, too broad to be covered entirely in one chapter. We therefore refer to the literature in the list below for further details.

I (CPD) owe thanks to Tom Herbst (MPIA) for his extremely fruitful and eye-opening 15 minute coffee-break explanation of interferometry with the LBT on 19 May 2010, which resulted in Sections 6.12 and 6.13.

Literature:

Lecture notes of the IRAM summer school on interferometry: `http://iram.fr/IRAMFR/IS/school.htm`.
Lecture notes about interferometry from NRAO: `http://www.cv.nrao.edu/course/astr534/PDFnew.shtml`.
Högbom, J.A., *Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines*, 1974, Astronomy & Astrophysics Supplements Vol. 15, p. 417
Book by Thompson, A.R., Moran, J.M. and Swenson, G.W. Jr., *Interferometry and synthesis in radio astronomy*, 1986, John Wily and Sons. ISBN 0-471-80614-5

6.1 Fizeau interferometry

Perhaps the most basic technique of interferometry is *Fizeau interferometry*, named after Hippolyte Fizeau (1819-1896), a French physicist who first suggested to use interferometry to measure the sizes of stars. The idea is simple: Just take the light of all your telescopes and project them, using a series of well-placed mirrors, all on the *same* image plane, as if the mirrors were all part of one huge mirror. If this is done such that the light from each of the telescopes arrives at the image plane exactly at the same time, the beams of light from all telescopes combined produce a point-spread-function (PSF) that is the Fourier transform of the *combined* apertures of the telescopes. This is as if we actually have a huge mirror, but put black paint on it everywhere except for two or more circular regions.

We know from chapter 3 (see Fig. 3.7) that one can compute this "combined PSF" easily using a Fast Fourier Transform (FFT) algorithm applied to the "image" of the aperture. Let us, for now, take the example of a pair of telescopes next to each other, like the *Large*

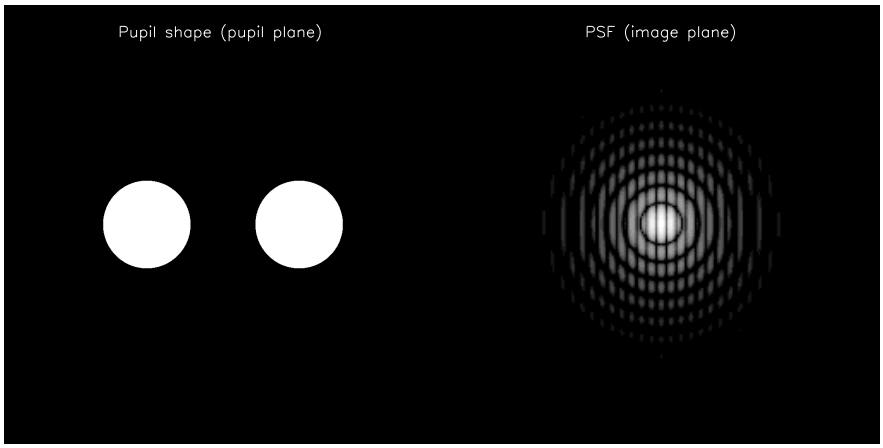

Figure 6.1: The PSF of a pair of two apertures next to each other, computed in exactly the same way as in Fig. 3.7. Left: The aperture pair. Right: The PSF of this pair, computed as the amplitude-squared of the Fourier transform of the pair of apertures (brightness is plotted logarithmically). This is the principle of a Fizeau interferometer, and this example is inspired by the LINC-NIRVANA interferometer to be installed on the Large Binocular Telescope.

Binocular Telescope (LBT) on Mount Graham in Arizona, which has participation from the Max-Planck-Institute for Astronomy in Heidelberg. Instead of having a single circular aperture, we have two circular apertures next to each other. If we compute the Fourier transform of this pair of apertures we get a PSF that is a product of the PSF of each single telescope and a wave-shaped modulation on top, called *fringes*. This is shown in Fig. 6.1. The result of such an interferometer is an image of the sky with a PSF which has the shape shown in Fig. 6.1. Each star on the image thus looks like an Airy disk multiplied by a fringe pattern (wavy modulation). Since the light from the two mirrors is combined on the image plane, Fizeau interferometers are called *focal plane interferometers*. At the Max-Planck-Institute for Astronomy in Heidelberg they are currently building the LINC-NIRVANA instrument for the LBT, which will do exactly this kind of interferometry.

If we would join more than just two mirrors, more modulations of the original single-mirror PSF would appear, and the PSF gets more and more centrally concentrated with fewer "sidelobes". It starts to look more and more like a PSF of a single huge mirror. This technique has been used on the *Multi-Mirror Telescope* (MMT) in Arizona, which featured six 1.8-meter mirrors on a single mounting. Note however that today the MMT no longer has six mirrors: it now features a single 6.5 m mirror. However, the experience gained with the MMT is now put to use for the LBT and for the design and future operation of the LINC-NIRVANA interferometer.

Let us now, as a Gedankenexperiment, add so many mirrors that they all touch each other, and we fill also the space in between with mirror. We have then created one single huge mirror. The modulations now all conspire to produce the Airy pattern for the huge mirror, i.e. much smaller in angular scale than the single-telescope PSF. This is what it should be, because we have now indeed created a true huge mirror.

Now let us go back to the case of the LBT with its future LINC-NIRVANA interferometer: the PSF has this unusual shape of a single-mirror PSF multiplied by a wave. If we would use such a camera to image a field on the sky with extended sources (like e.g. a

galaxy or a spatially resolved image of a solar system planet), then we would still get image degradation (smearing) over a typical scale of the original single-mirror PSF, though with a central peak. It seems that we have gained only partially in terms of angular resolution. The solution lies in the following:

1. Make multiple images over one night, to allow the field-rotation caused by the Earth's rotation to rotate the baseline of the LBT over your object so that the wavy modulation (fringe pattern) occurs in different directions for each image.

2. Use computer software to "deconvolve" the combined set of images into a single high spatial resolution image. The trick here is to ask the computer to find the image which, when convolved with each of the PSFs of each of the observed images, yields the best match to the actual observed images[1].

Fizeau interferometers are particularly useful for cases where the telescope mirrors are close together. This is because in that way a relatively large part of the hypothetical huge mirror is accounted for by the real mirrors. The PSF fringes are, in this case, just a factor of two or three times as narrow as the single-mirror PSF, which helps the computer software very much to find the high-resolution image. If one would use Fizeau interferometry for mirrors that are very far away from each other (measured in units of the mirror size), then each single-mirror PSF would be modulated by a very fine fringe pattern. While this formally yields an angular resolution of a single fringe bump, it contains too little information for computer software to "deconvolve" the image to the sharpness of the single fringe bumps. This is not surprising since the two mirrors, even when accounting for field rotation, would cover only a fraction of the surface area of the hypothetical huge mirror. One can thus not reasonably expect to be able to create images anywhere near what one would obtain with the huge mirror.

The solution would lie in using many such mirrors, each subjective to field rotation. In fact, this is what is done in radio interferometry and in many areas of infrared interferometery. However, the above computer technique to retrieve the high-resolution image would then quickly become impractical: one would have to work with thousands of images, and the methods becomes prohibitively slow. Instead, for such cases one resorts to different techniques. But to understand those, we will have to have a closer look at the concept of *coherence*.

6.2 Coherence in time

In reality astrophysical radiation rarely behaves as a perfect monochromatic wave front. With the exception of laser/maser emission, most astrophysical emission is *incoherent*: radiation emitted at one location does not care about radiation that has been emitted at another location. Yet, the fact that light behaves as a wave shows that even if astrophysical radiation is largely incoherent, some degree of coherence must automatically arise.

Let us do a simple thought experiment which is not mathematically strict, but illustrative. Consider a long tube of hot gas. At the leftmost end an atom emits radiation through the tube. This radiation is a wave, the electric field of which we denote by $E_1(x,t)$. In this

[1] A nice paper describing such computer methods to be used for LINC-NIRVANA is: Bertero & Boccacci (2000) A&A Suppl. 144, 181.

analysis we ignore the vector character of the electromagnetic field: we assume that it is perpendicular to the direction of propagation ($\vec{E}\vec{k} = 0$, where $\vec{k}$ is the wave vector), and we focus on just one of the two polarization components. Next to it another atom emits radiation, given by $E_2(x,t)$, and next to that $E_3(x,t)$ etc., until we reach atom N, the final atom, which emits $E_N(x,t)$. We assume that all these radiation sources emit radiation of the same amplitude A, but mutually non-coherent, so therefore all these waves have completely arbitrary phase compared to each other. In the complex plane this means we add N vectors of equal length but arbitrary direction. Assume now also that over a long time the phases of all these emitting particles gradually shift randomly, so that the sum vector assumes various incarnations of the sum of N randomly oriented vectors of length A. According to standard statistics we know then that the average length of the resulting vector $\sqrt{\langle E^*E\rangle} = \sqrt{N}A$. The intensity $I = \langle E^*E\rangle = NA^2$. The intesity therefore scales *linearly* with N, precisely what one would expect for incoherent emitters, and precisely what one would expect for the particle-description of radiation: Each atom emits its own photons, not caring about any other photons around; the photon chain obeys Poisson statistics[2]. This is an example of how incoherent radiation still produces a wave which, in itself, has some degree of coherence.

The above thought experiment only works, however, if the phases of the emitting particles gradually drift, so that on average we get the $\sqrt{N}$ behavior of the resulting radiation. The question is: for how long will a wave keep its phase? In other words: what is the distance along a wave for which one can assume phase stability? To study this, let us introduce the *autocovariance* $\Gamma(\tau)$ of the wave, as measured at a given location for time τ:

$$\Gamma(\tau) \equiv \langle E^*(t)E(t+\tau)\rangle \equiv \frac{1}{T}\int_0^T E^*(t)E(t+\tau)dt \tag{6.1}$$

where t is time and T is sufficiently large that it can be considered "semi-infinite", i.e. a long time compared to one wave period. For visible light a measurement of one second already amounts to about 10^{15} wave periods – "infinitely many" for most practical purposes. The flux F in units of erg cm^{-2}s^{-1}, for the case of a (nearly) plane wave, is then

$$F = \frac{c}{4\pi}\langle E^*(t)E(t)\rangle = \frac{c}{4\pi}\Gamma(0) \tag{6.2}$$

This is a real quantity, because $\Gamma(0)$ has no imaginary component. However, $\Gamma(\tau)$ is in general a complex quantity. Define the *autocorrelation* $R(\tau)$ as

$$R(\tau) = \frac{\Gamma(\tau)}{\Gamma(0)} \tag{6.3}$$

Its amplitude tells how much the signal stays correlated over a time span of τ, i.e. how much "phase memory" the signal has over that time span. The autocorrelation function therefore allows us to define a typical coherence time τ_c for which $|R(\tau < \tau_c)| \geq 1/e$.

Now coming back to our original question: What is the value of this coherence time? Interestingly, the answer lies as much in the technique of observation as it does in the process of emission. This is because the minimal coherence time is inversely proportional

[2]Keep in mind, however, that this example breaks down if the medium becomes optically thick, since then stimulated emission becomes important even if no masering occurs.

to the wavelength bandwidth at which we observe our astronomical signal. And for an emission line it is also inversely proportional to the line width.

To see this, let us start with the utopic case of a perfectly monochromatic wave. In such a case the phase memory is by definition perfect, and the amplitude of the autocorrelation function $R(\tau)$ is 1 for all values of τ. This is because the Fourier transform of a delta function is a perfect wave. However, if we would modify the wave slightly by introducing a gradual loss of phase memory over the coherence time scale τ_c, the Fourier transform is no longer a perfect delta-function. It will be a slightly broadened peak with a width:

$$\Delta\nu \simeq \frac{1}{\tau_c}. \tag{6.4}$$

In other words: a certain degree of incoherence is automatically related to a certain degree of non-monochromaticness. By using a narrow wavelength filter in our telescope detector we automatically increase the minimal coherence time. However, if we measure an emission line which is narrower than our filter, then the actual coherence time is longer than that set by our filter: it is set by the line width.

6.3 Coherence in time and space

Let us now generalize the notion of coherence to time *and* space. Consider two spatial locations $\vec{r}_1$ and $\vec{r}_2$ where we measure the radiation field. We can now define the quantity

$$\Gamma_{12}(\tau) \equiv \langle E_1^*(t) E_2(t+\tau) \rangle \tag{6.5}$$

where E_1 is a shorthand for $E(\vec{r}_1)$. Let us further define

$$R_{12}(\tau) = \frac{\Gamma_{12}(\tau)}{\sqrt{\Gamma_{11}(0)\Gamma_{22}(0)}} \tag{6.6}$$

If $\vec{r}_1$ and $\vec{r}_2$ lie along the wave vector $\vec{k}$ (i.e. along the propagation of the light), then the spatial coherence and the time coherence are in fact the same thing. Let is, in this particular case, write $R_{12}(\tau)$ as $R(l,\tau)$. We then have

$$|R(c\tau, 0)| = |R(0,\tau)| \tag{6.7}$$

One can therefore, in this case, define the coherence length l_c directly from the coherence time τ_c: $l_c = c\tau_c$.

However, if $\vec{r}_1$ and $\vec{r}_2$ are not along the $\vec{k}$ vector, then things become more complicated and we need to visit the van Cittert-Zernike theorem.

6.4 Van Cittert - Zernike theorem

The van Cittert - Zernike theorem is named after Pieter Hendrik van Cittert (1889-?) and Frits Zernike (1888-1966). Zernike was a professor in Groningen, the Netherlands, and obtained the Nobel prize for physics in 1953 for the invention of the phase contrast microscope. In its simplified version we will focus on here, the van Cittert - Zernike theorem addresses the question of the degree of spatially coherence of emission from some object(s)

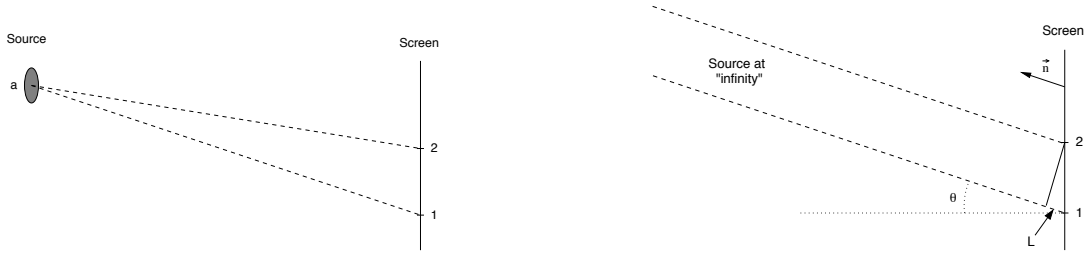

Figure 6.2: Pictograms used for the derivation of the van Cittert - Zernike theorem. Left: a source of a size smaller than the wavelength emitting pseudo-monochromatic radiation, which is detected at two positions on the screen: position 1 and 2. Right: same setup, but with the source moved away to very large distance so that the rays are effectively parallel.

on the sky impinging on a screen. The situation is pictographically shown in Fig. 6.2-left. In this pictogram a small area located at the point marked "a" emits radiation, which is received at locations "1" and "2". We assume that the size of the emitting region at "a" is roughly one wavelength of the emitted radiation. Later we will integrate over zillions of regions "a", but for now we just take a single emitting region. Let us write the distance between the source and point "1" as d_1 and define d_2 likewise. We assume that $d_{1,2} \gg |\vec{r}_2 - \vec{r}_1|$, so that for many (but not all!) purposes we can use the average distance $d = (d_1 + d_2)/2$.

Let us now assume that the emitting region "a" emits perfectly monochromatic radiation at wavelength ν. The electric field at the location "a" can then be written as

$$E_a(t) = A_a e^{-2\pi i \nu t} \tag{6.8}$$

where A_a is a complex number that does not change in time. The electric field emerging from this source falls off, as we know from electrodynamics theory, as $|E| \sim 1/d$, i.e. inversely proportional to the distance from "a". In addition, there is a phase lag between points "a" and "1" (or equivalently "2") because of the fact that light emitted at time t at point "a" will reach point "1" at time $t + d_1/c$. The electric field at point "1" can thus be written as

$$E_1(t) \simeq \frac{\sqrt{\Delta S_a}}{d} A_a \exp\left[-2\pi i \nu \left(t - \frac{d_1}{c}\right)\right] \tag{6.9}$$

and likewise for $E_2(t)$. Here ΔS_a is the surface area of the emitting region at "a". We used $d_1 \simeq d$ in the denominator of the above expression, but in the exponent we kept d_1. But why did we use the square-root of ΔS_a? This is because when we later integrate over many emitting regions, we must account for the fact that these emitting regions are mutually non-coherent. Analogous to the description in Section 6.2, when we add the emission of N non-coherent regions, we get – on average – an electric field amplitude of $\sqrt{N}$ times that of a single region. The field in Eq. (6.9) therefore scales with $\sqrt{\Delta S_a}$.

Our purpose was to study the spatial coherence between points "1" and "2", so let us compute

$$\langle E_1^*(t) E_2(t) \rangle = \frac{1}{T} \int_0^T \frac{\Delta S_a}{d^2} A_a^* A_a \exp\left[2\pi i \nu \left(\frac{d_2 - d_1}{c}\right)\right] dt \tag{6.10}$$

$$= \frac{\Delta S_a}{d^2} A_a^* A_a \exp\left[2\pi i \nu \left(\frac{d_2 - d_1}{c}\right)\right] \tag{6.11}$$

For the special case of $\vec{r}_2 = \vec{r}_1$ we obtain the flux emerging from this emitting region and measured at the screen:

$$F = \frac{c}{4\pi}\langle E_1^*(t)E_1(t)\rangle = \frac{c}{4\pi}\frac{\Delta S_a}{d^2}A_a^*A_a \tag{6.12}$$

But if $\vec{r}_1 \neq \vec{r}_2$ we must find a useful expression for the distance difference $d_2 - d_1$. For this, let us remind ourselves that $d_{1,2} \gg |\vec{r}_2 - \vec{r}_1|$, so that the two rays are essentially parallel, as shown in Fig. 6.2-right. The quantity L shown in the figure is the value of $d_2 - d_1$ that we need. It is given by $L = d_1 - d_2 = |\vec{r}_2 - \vec{r}_1| \sin\theta$, where θ is the angle toward the source, measured from the normal of the screen (see figure). We can write this more practically using vector notation. The vector $\vec{n}$ shown in the figure is a unit vector pointing toward the source. We then have

$$d_1 - d_2 = \vec{n} \cdot (\vec{r}_2 - \vec{r}_1) \tag{6.13}$$

so that we can write the coherence as

$$\langle E_1^*(t)E_2(t)\rangle = \frac{\Delta S_a}{d^2}A_a^*A_a \exp\left[-2\pi i\nu\left(\frac{\vec{n}\cdot(\vec{r}_2 - \vec{r}_1)}{c}\right)\right] \tag{6.14}$$

The inner products are taken in 3-D. But if the screen is purely 2-D, we can cast this inner product into a 2-D version. Let us take the z-axis to be *perpendicular* to the screen. If we define $\vec{r} \equiv \vec{r}_2 - \vec{r}_1$ we thus have

$$\vec{n} \cdot (\vec{r}_2 - \vec{r}_1) \equiv \vec{n} \cdot \vec{r} = n_x r_x + n_y r_y \tag{6.15}$$

because $r_z = 0$ as $\vec{r}$ lies in the plane of the screen. We can now interpret (n_x, n_y) as a 2-D vector describing the angular position of our emitting object on the sky and (r_x, r_y) as a 2-D baseline on the screen. We obtain

$$\langle E_1^*(t)E_2(t)\rangle = \frac{\Delta S_a}{d^2}A_a^*A_a \exp\left[-2\pi i\nu\left(\frac{n_x r_x + n_y r_y}{c}\right)\right] \tag{6.16}$$

This is, for the case of a single emitting area "a", our final result for the moment. It shows that the radiation fields between points "1" and "2" on the screen are perfectly correlated, but have a phase shift of

$$\delta\phi(r_x, r_y) = \nu\frac{n_x r_x + n_y r_y}{c} \tag{6.17}$$

So far this is not really surprising, as it is what one would expect from a plane-parallel wave impinging on the screen at an orientation $\vec{n}$.

But now let us integrate this result over an entire region on the sky, i.e. a continuous series of regions "a". To do this, we have to relate steps in n_x and n_y to the surface area ΔS:

$$dS_a = d^2 dn_x dn_y \tag{6.18}$$

We thus get

$$\langle E_1^*(t)E_2(t)\rangle = \int A^*(n_x, n_y)A(n_x, n_y) \exp\left[-2\pi i\nu\left(\frac{n_x r_x + n_y r_y}{c}\right)\right] dn_x dn_y \tag{6.19}$$

where $A(n_x, n_y)$ now depends on n_x and n_y, as we now integrate over different regions on the sky. Looking carefully at the above equation you may notice that this is in fact a

2-D Fourier integral with r_x/λ and r_y/λ as the wave numbers in x and y direction (with $\lambda = c/\nu$). The integrand that is Fourier transformed is $A^*(n_x, n_y)A(n_x, n_y)$, which is in fact the *intensity* $I(n_x, n_y)$. The correlation between points "1" and "2" is therefore the Fourier transform of the image on the sky. This is the van Cittert - Zernike theorem, and stands at the basis of interferometry.

6.5 Delay lines and the uv-plane

To reconstruct the image on the sky from interferometry measurements we must, according to the van Cittert - Zernike theorem, measure the correlation of the electric field between between as many pairs of points in the "pupil plane" as possible. With "pupil plane" is meant the plane parallel to the wave front of the object we wish to observe. In Fig. 6.3 this is indicated with the dotted line. However, in practice our telescopes are usually not located in the same plane, except for the LBT where both mirrors are on the same mounting. One can solve this problem by effectively delaying the signal received by the telescope closest to the source (the right one in Fig. 6.3) in a *delay line*. For optical interferometers this is simply a tunnel with mirrors where the light is diverted into, such that the path length of the light from both telescopes to the instrument (correlator) is equal. Since the Earth is rotating and thus the relative projected distance between the telescopes (called the *projected baseline*) changes accordingly, the length of the required delay path in the delay line must be continuously adapted. At the VLT this is done with mirrors mounted onto a "train" that moves steadily along a rails at a very precisely tuned velocity. This must be a very high-precision device, since the delay path must be accurate to within 1 μm over projected baselines of up to 200 m!

If we have N telescopes, then we need $N-1$ delay lines to bring them all effectively into the same pupil plane.

One may ask: why do we have to shift all telescopes into the same pupil plane? If we have a perfect plane wave hitting our telescopes, we would have interference also if the telescopes are several periods of the wave out of phase. A perfect plane wave always interferes with itself since $\sin(\phi) = \sin(\phi + 2\pi n)$, even if $n = 100000$ or more. The answer is related to the coherence length. As we saw in Section 6.2, if you measurement is done in a finite-width wavelength band, then there is a finite coherence time τ_c, corresponding to a finite coherence length along the direction of propagation $l_c = c\tau_c$. If two telescopes are not brought to withing a distance l_c from their mutual pupil plane, then the signals of the two telescopes become decoherent, and interferometry is not possible. Taking too narrow a band would mean that you receive a very weak signal, which could kill your interferometry attempt. But the broader your band, the smaller l_c becomes and the closer you have to bring the two telescopes into their common projected plane.

Let us now define a common pupil plane for all N telescopes. The precise position of this plane along the pointing direction (the direction perpendicular to the plane) is not important (and in fact does not affect our measurements), so we just choose one.

Now define a coordinate system in this plane: (x', y'). Again, the precise location of their zero-points are not important. We now want to measure $\langle E_1^*(t)E_2(t)\rangle$ in the entire pupil plane, and thus obtain a two-dimensional function

$$\mathrm{Corr}_E(x_1', y_1', x_2', y_2') = \langle E_1^*(t)E_2(t)\rangle \tag{6.20}$$

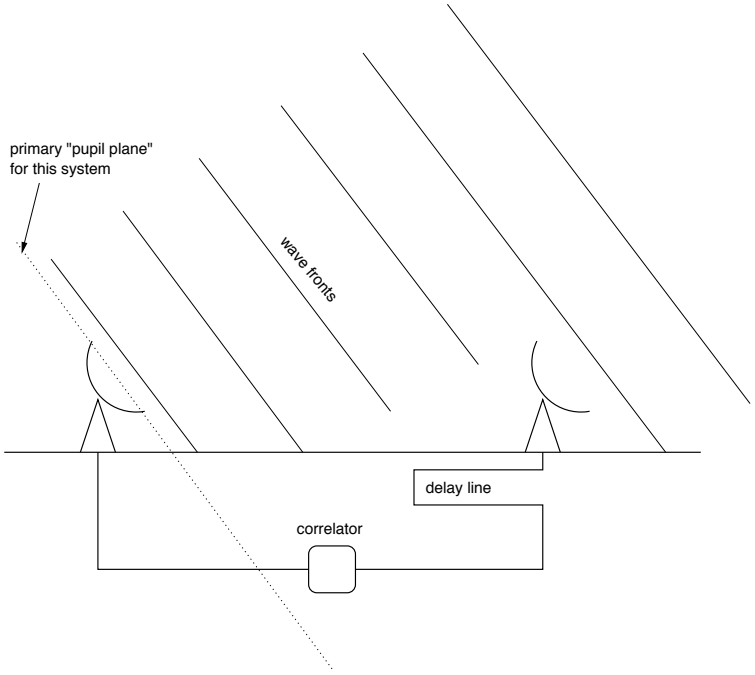

Figure 6.3: Principle of an interferometer with delay lines to effectively "shift" the telescopes into a common "pupil plane". In this example the right telescope is, by the delay line, effectively shifted back to the dotted plane.

Since $\langle E_1^*(t)E_2(t)\rangle$ will depend not on the absolute positions of our telescopes, but only on the relative positions $x'_{12} = x'_2 - x'_1$ and $y'_{12} = y'_2 - y'_1$:

$$\mathrm{Corr}_E(x'_{12}, y'_{12}) = \langle E_1^*(t)E_2(t)\rangle \tag{6.21}$$

It will turn out to be more convenient to measure these coordinates in units of the wavelength:

$$u := \frac{x'_2 - x'_1}{\lambda} \qquad \text{and} \qquad v := \frac{y'_2 - y'_1}{\lambda} \tag{6.22}$$

This is what is called the *uv-plane*. It is actually nothing else than the pupil plane, scaled such that one wavelength corresponds to unity.

If we want to be able to reconstruct an image of an object on the sky, then we must measure

$$\mathrm{Corr}_E(u, v) = \langle E_1^*(t)E_2(t)\rangle \tag{6.23}$$

at as many (u, v) points as possible. In other words, we must have a good *uv-coverage*.

If we have N telescopes measuring simultaneously, then we have $N(N-1)/2$ independent baselines. Note that the projected baseline corresponding to (u, v) is the same as the one corresponding to $(-u, -v)$. Now, the *projected* baseline on the sky changes with time, because the Earth rotates. This means that each baseline describes an elliptic curve in the uv-plane, and therefore you get multiple (u, v)-points (multiple projected baselines) for a single two-telescope baseline on the ground. If you measure over a substantial part of the night, you therefore get a much better uv-coverage as when you do a single short measurement.

Ideally we must cover the uv-plane perfectly, but this is never possible. We will discuss how to make sense out of measurements with imperfect uv-coverage in Section 6.9. But let us first have a look at how to measure $\langle E_1^*(t)E_2(t)\rangle$ in the first place.

6.6 Measuring "fringes": The "visibility" of an interferometric signal

Once we have, using delay lines, effectively shifted a pair of telescopes into their common pupil plane, we can try to measure the correlation of their signals. While in optical interferometry each telescope receives a full image, in radio interferometry a telescope typically receives just a single analog signal (i.e. one single "pixel"). Let us, in this section and the next, focus on the case where we are dealing just with a single pixel and we have to measure the correlation between these signals, as they are received by two or more telescopes.

Strictly speaking we could record the exact wave pattern of $E_1(t)$ and $E_2(t)$, where 1 and 2 stand for telescope 1 and telescope 2, and t is such that they are measured in the same pupil plane (i.e. with delay included). Having both signals, we could then calculate $\langle E_1^* E_2\rangle$ directly using Eq. (B.38) with $\tau = 0$. If we would do this (which is impractical, and for optical interferometry even physically impossible), then we would have at any given time t *absolute phase information*: we know the exact value of E_1 and E_2 at time t. This information is, however, useless, because with time the phase changes rapidly (it changes 2π for each time interval $\Delta t = 1/\nu$). We are more interested in the *relative phase* between the two telescopes. This gives information about the exact position of a source on the sky. In fact, this is exactly one of the two pieces of information in the complex number $\langle E_1^* E_2\rangle$: writing this as $Ae^{i\phi}$ the ϕ is this relative phase. If one can measure this, then one can do astrometry at ultra-high precision.

In practice one never records the precise wave functions $E_1(t)$ and $E_2(t)$. One uses other techniques. One often used technique (in particular in radio interferometry, but also in infrared interferometry sometimes) is the technique of *heterodyne interferometry*, which we will deal with in Section 6.11. But to understand the principles of *long-baseline interferometry* (here defined as interferometry with telescopes that are not on the same mounting), we keep it a bit more simple for the moment.

The simplest way to measure the correlation between the two signals is to simply let them interfere: We redirect both signals onto a single device that measures the square amplitude of the sum of the signals:

$$S_{12} = \langle [E_1^*(t) + E_2^*(t)][E_1(t) + E_2(t)]\rangle \tag{6.24}$$

This is what you would get in optical interferometry if you simply let the two signals interfere on a CCD camera: you measure the intensity of the signal, i.e. the amplitude-squared of the sum of the signals. It is also what you get if you measure the energy output of two electric signals from two radio telescopes linked together. We can measure such a signal with standard technology (e.g. a CCD for optical interferometry). Note that Eq. 6.24 can be regarded as the spatial structure function of $E(x', y', t)$ where (x', y') are the two spatial coordinates in the pupil plane, and $E_1(t) \equiv E(x'_1, y'_1, t)$ and $E_2(t) \equiv E(x'_2, y'_2, t)$.

To get from the measured quantity S_{12} to the desired quantity $\langle E_1^* E_2 \rangle$, we write:

$$\begin{aligned} S_{12} &= \langle [E_1^*(t) + E_2^*(t)][E_1(t) + E_2(t)] \rangle \\ &= \langle E_1^*(t) E_1(t) \rangle + \langle E_2^*(t) E_2(t) \rangle + \langle E_1^*(t) E_2(t) \rangle + \langle E_2^*(t) E_1(t) \rangle \end{aligned} \tag{6.25}$$

The first two terms are simply the intensity of the object measured by each telescope: $S_1 = \langle E_1^*(t) E_1(t) \rangle$, $S_2 = \langle E_2^*(t) E_2(t) \rangle$. If both telescopes have the same diameter, these are the same, and we simply write $S = \langle E^* E \rangle$. We can measure this quantity simply by shutting out one of the telescopes and measuring the signal, which is then S. Eq. 6.25 thus becomes

$$S_{12} = 2S + \langle E_1^*(t) E_2(t) \rangle + \langle E_2^*(t) E_1(t) \rangle \tag{6.26}$$

The last two terms are the correlation and its complex conjugate. This is the quantity we need. So if we evaluate $S_{12} - 2S$ we obtain $\langle E_1^*(t) E_2(t) \rangle + \langle E_2^*(t) E_1(t) \rangle$ which is nearly what we need. To obtain exactly what we need (the *complex number* $\langle E_1^*(t) E_2(t) \rangle$, not just its real part) we can write

$$\langle E_1^*(t) E_2(t) \rangle =: A_{12} e^{i\phi_{12}} \tag{6.27}$$

where A_{12} is a real number (the amplitude of the correlation) and ϕ is the relative phase. We obtain

$$\langle E_1^*(t) E_2(t) \rangle + \langle E_2^*(t) E_1(t) \rangle = 2A_{12} \cos(\phi_{12}) \tag{6.28}$$

We see that Eq. 6.26 becomes

$$S_{12} = 2S + 2A_{12} \cos(\phi_{12}) \tag{6.29}$$

If the two signals are totally decoherent, then $A_{12} = 0$, and we measure $S_{12} = 2S$. If they are perfectly coherent (if we measure a point source on the sky), then $A_{12} = S$. In that case we measure $S_{12} = 2S[1 + \cos(\phi_{12})]$.

So how do we get from Eq. 6.29 to the complex number $\langle E_1^*(t) E_2(t) \rangle$ that we need? This is a bit subtle. If we do a single measurement of S_{12} and S, then we have one real(!) equation (Eq. 6.29) for two real unknowns (A_{12} and ϕ_{12}). So we do not have enough information to fully reconstruct the complex value of $\langle E_1^*(t) E_2(t) \rangle$.

To solve this impasse, we can do a trick: we slightly change the length of the delay line to induce an extra phase difference between the two telescopes. Let us call this artifically induced phase difference $\delta\phi$. If we scan over a few wavelength by changing $\delta\phi$ smoothly, we will see that $S_{12}(\delta\phi)$ will follow a cosine-like curve with an offset:

$$S_{12}(\delta\phi) = 2S + 2A_{12} \cos(\phi_{12} + \delta\phi) \tag{6.30}$$

From this so-called *fringe pattern* we can read off the offset (which should be, and will be exactly $2S$) and the amplitude of the cosine, which is the value of $2A_{12}$. In formulae: we read off the maximum value of $S_{12}(\delta\phi)$, and call it $S_{\max}$, and the minimum of $S_{12}(\delta\phi)$, and call it $S_{\min}$, and we then have

$$A_{12} = \frac{S_{\max} - S_{\min}}{4} \tag{6.31}$$

Strictly speaking, we can now go back to the real pupil plane ($\delta\phi = 0$), and then we obtain the relative phase ϕ_{12} by solving

$$\cos(\phi_{12}) = \frac{S_{12}/2 - S}{A_{12}} \tag{6.32}$$

For any given domain of length 2π this has two solutions, but since we have, by varying $\delta\phi$, scanned the fringe pattern, we can figure out which of these two we must take. In this we we have, at least in principle, found the phase difference between telescopes 1 and 2. But in most practical circumstances this an extremely unreliable quantity. The reason is that it is hard to calibrate an interferometer so accurately that we can find the exact pupil plane, or in other words: to find the exact length of the delay line required to bring both telescopes exactly to a known projected distance from each other. Remember that for IR interferometry we would have to do this to within a fraction of a micron, while the telescopes have distances of up to 200 meters for the VLT! In principle this would be not impossible, but in practice this is very hard. But an even more dramatic problem is the turbulence of the atmosphere: as we have seen in Chapter 4, turbulence induces phase shifts. These will be different (and mutually uncorrelated) for the different telescopes, because the telescope separation is usually larger than the Fried length r_0.

In most circumstances one tries to find the pupil plane (i.e. find the right delay length) by making an initial good guess and then tuning it (i.e. moving $\delta\phi_{12}$ over many, many wavelengths) until you find a fringe pattern. Once you found a fringe pattern, you know that you are within one coherence length from the pupil (see Section 6.3). That is usually the best you can do. Keeping this (unknown) distance from the pupil plane fixed can be done much more accurately. Special devices that continuously fine-tune the mirrors in the delay line (using piezo-electric element) are called *fringe trackers*.

In conclusion: what we have measured is A_{12}, but not ϕ.

The quantity A_{12} is what is called the *correlated flux* and obeys

$$0 \leq A_{12} \leq S \tag{6.33}$$

One can thus split the flux S into a correlated part (A_{12}) and an uncorrelated part ($S-A_{12}$). Professional interferometrists also often use what they call the *visibility* V_{12} (nevermind the confusing name), which is

$$V_{12} = \frac{S_{\max} - S_{\min}}{4S} = \frac{A_{12}}{S} \tag{6.34}$$

(see Eq. 6.31). All these quantities can be directly measured in the manner described above.

Also often the *complex visibility* is used, which is defined as

$$\mathcal{V}_{12} = \frac{\langle E_1^*(t) E_2(t) \rangle}{S} = \frac{A_{12} e^{i\phi_{12}}}{S} \tag{6.35}$$

which is, as we now know, not directly measurable with the techniques in this section: Only its amplitude A_{12} is measurable. But it is nevertheless a very useful quantity, as we shall see in Section 6.8.

6.7 Examples of visibility signals of simple objects

TO BE DONE

6.8 The "Closure phase"

By comparing predicted visibilities of simple models with measured visibilities at different (u, v) we can already obtain some information about what the object on the sky looks like

(Section 6.7). But can we also reconstruct an entire image on the sky from measurements of the visibility in the uv-plane? The answer is: no. We would need the full *complex* visibility $\mathcal{V}(u,v)$ instead of only its amplitude $V(u,v)$. The amplitude information $V(u,v)$ tells us which Fourier components are in our image, but not their position (phase), and thus we cannot do the inverse Fourier transform.

So how to solve this problem? Remember that we can *in principle* measure some phase for each baseline using Eq. 6.32, but that we rejected it because it was fraught with uncertainty. But if we study this uncertainty more carefully, we can do a trick to overcome this problem at least partially.

Suppose we have three telescopes (1, 2 and 3). Let us quantify the (unknown) phase error of each telescope with ϵ_1, ϵ_2 and ϵ_3, respectively. If we measure the phases ϕ_{12}, ϕ_{23} and ϕ_{31} using e.g. Eq. 6.32 we obtain the following measurements:

$$\phi_{12}^{\rm measured} = \phi_{12} + \epsilon_1 - \epsilon_2 \tag{6.36}$$
$$\phi_{23}^{\rm measured} = \phi_{23} + \epsilon_2 - \epsilon_2 \tag{6.37}$$
$$\phi_{31}^{\rm measured} = \phi_{31} + \epsilon_3 - \epsilon_1 \tag{6.38}$$

where ϕ_{12}, ϕ_{23} and ϕ_{31} are the real phases (assuming no atmosphere and no instrument errors). If we would know ϵ_1, ϵ_2 and ϵ_3, then we would be able to retrieve these real phases from the measured phases. But since we do not know these phase errors, we cannot do this. This is in fact what we concluded in Section 6.6 and was the reason why we rejected the use of Eq. 6.32 to find the phases.

However, if we take the sum of Eqs. 6.36, 6.37, 6.38, then we obtain

$$\phi_{123} \equiv \phi_{12} + \phi_{23} + \phi_{31} = \phi_{12}^{\rm measured} + \phi_{23}^{\rm measured} + \phi_{31}^{\rm measured} \tag{6.39}$$

This is called the *closure phase*. As one can see: all unknown errors have dropped out, so it is a reliable quantity. The closure phase contains information about depature from point-symmetry. Any object on the sky that is perfectly point-symmetric (e.g. an ellipse) will have zero closure phase.

The closure phase has a very important property: It is not affected by phase shifts due to the atmosphere. Atmospheric turbulence has therefore no influence on this measurement.

For an interferometer with N baselines, we can determine the closure phase ϕ_{ijk} between each triple of telescopes i, j and k. Now suppose we *assume* the phase between 1 and 2 and between 1 and 3 to be zero (assuming that the telescopes 1, 2 and 3 are *not* aligned along the same line as projected on the sky), then by measuring the closure phase ϕ_{123} we can calculate $\phi_{23} = \phi_{123}$. Now add one telescope, number 4, and measure ϕ_{124} and ϕ_{234}. Since we know ϕ_{12} and ϕ_{23} we obtain the equations

$$\phi_{12} + \phi_{24} + \phi_{41} = \phi_{124} \tag{6.40}$$
$$\phi_{23} + \phi_{34} + \phi_{42} = \phi_{234} \tag{6.41}$$

which is two equations with two unknowns (ϕ_{41} and ϕ_{42}). So we now know again the relative phases at all baselines. In fact, since we now have $N = 4$ telescopes, we have 6 baselines, and for all these baselines we have the amplitude *and* phase information, albeit that we made an assumption for two of the phases. This assumption for these two initial relative phases amounts to an assumption for the position of our object on the sky (which is two coordinates). So apart from this freedom to shift our image anywhere on the sky, we

have now obtained all amplitude and phase information we need. In other words: we have (excluding the two assumed phases): $N(N-1)/2$ visibility amplitudes and $N(N-1)/2-2$ pieces of phase information.

Since the value of S is the same for each baseline pair, we can say that an interferometer array of N telescopes gives $N(N-1)/2$ values of the complex visibility $\mathcal{V}$ (of which two of the phases have been chosen ad-hoc) and one value of the total flux F (which is S divided by the telescope aperture size). Of course, in this calculation it is assumed that none of the baselines are duplicates of each other, which is not generally guaranteed. For instance, the Westerbork telescope or the VLA have telescopes arranges in regular spacing. If you have for instance three telescopes in a row, and the distance between telescope 1 and 2 is the same as that between 2 and 3, then in effect you have not 3, but just 2 independent baselines. Also, keep in mind that $\mathcal{V}(u,v) = \mathcal{V}^*(-u,-v)$, so that when you obtain $N(N-1)$ complex visibilities, you get the ones oppose of the origin for free. In that sense one has in fact $N(N-1)$ complex visibility points, but only half of them are independent.

6.9 Image reconstruction: Aperture synthesis and the CLEAN algorithm

When we do interferometry in the above described way, by measuring visibilities and closure phases, and thus constructing a set of $N(N-1)/2$ independent complex visibility points $\mathcal{V}(u,v)$ plus their mirror copies $\mathcal{V}(-u,-v)$, we may have the information we need in order to apply van Cittert-Zernike theory to reconstruct the image, but in practice this is not so trivial. The problem is that we never have a perfect coverage of the uv-plane. And in order to perform the integrals of the Fourier transformation, we would in fact need a full coverage.

If we would have a nice and regular uv-coverage, for instance a perfectly rectangular grid in (u,v), then we could use the Fast Fourier Transform algorithm (see Section A.5) to obtain our image on the sky. But we rarely have the complex visibilities measured on such a regular grid in the uv-plane. In practice we have irregular spacings and certain uv-areas where there is little or no sampling at all. The complex visibility in these "holes" could be anything; we simply have no information about that. We could put the complex visibility in these "holes" to zero and linearly interpolate between nearby measurements in regions with sufficient sampling. What we then get, after Fourier transformation, is an image of the sky that is quite messy. Such an image is called a *dirty image*. It is in fact the true image convolved with the PSF corresponding to the uv-coverage. This PSF is called the *dirty beam*, because this PSF, due to the "holes" in the uv-coverage, has many sidelobes. A single point source on the sky will thus appear as a point source surrounded by weaker point sources, some of which can actually be rather far away from the actual point source. For reasonable uv-coverage these sidelobes are usually substantially weaker than the main peak, but they still mess up the image pretty much.

People have tried many different ways to "guess" the visibility between measurement points in order to get the best possible image out of it after Fourier transformation. But in practice most of these methods have drawbacks that can be substantial. A radically different approach was proposed in 1974 by Högbom (see literature list). In this paper it was proposed to make a dirty image using one's best guess of the inter-measurement

complex visibility values. Call this image A. Also make an image B which we initially make empty. Now do the following procedure, which is called the `CLEAN` algorithm:

1. Find the highest peak in the image A
2. Fit the peak of the dirty beam to this peak
3. Subtract this properly normalized dirty beam from the image A
4. Add a clean beam with the same strength at the same position in image B
5. Go back to 1 until we hit the noise level

The *clean beam* is a simple Gaussian-shaped peak with the same width as the main peak of the dirty beam. Also the shape of the clean beam can be made elliptic so as to account for the different spatial resolution one typically has in different directions (e.g. if a source is near the horizon). The above scheme is a slightly simplified version of the actual `CLEAN` algorithm, but it shows the principle.

The `CLEAN` algorithm works very well if the object on the sky is a collection of point sources. If it is, on the other hand, a rather smooth configuration it may work less well.

Also note that the typical lengths of the projected baselines determines to which typical sizes of features of your object on the sky you are sensitive to. If you have only very long baselines, then you may see small-scale details, but you miss large scale structures. The `CLEAN` procedure, by adding one clean beam at a time to your image until you reach noise level, will thus simply not add any large scale structure to your image. If the object *does* have large scale structure, then the integrated flux of the image the `CLEAN` procedure produces will be less than the measured integrated flux (measured from the single telescope signal S). By comparing the flux retrieved in the clean image with the single telescope flux you can estimate how much large scale structure is missing.

The full cycle of observing the complex visibilities, filling the uv-plane and reconstructing the image in the image plane using e.g. the `CLEAN` algorithm is called *aperture synthesis* or *synthesis imaging*.

6.10 Primary beam and coherent field of view

If we point our interferometer at the sky we have a spatial resolution of $\theta_i = \lambda/L$, where L is our largest projected baseline. However, we cannot know for sure that we are measuring exactly a single object: the object we want to observe. We may be accidently picking up also another nearby object or objects. We must therefore go a bit more into detail of what we are actually seeing with our interferometer.

First of all, we have the field of view of each of the telescopes individually. Let us define the radius of this field of view as θ_f. For infrared or optical telescopes this is typically much larger than the telescope resolution, $\theta_b = \lambda/D$, though this may vary between instruments used on that telescope, and even between different modes of the same instrument on the same telescope. At any rate: the field of view is larger than the size of the PSF: $\theta_f \gg \theta_b$.

With radio telescopes, on the other hand, the PSF of a single telescope (here often called the "primary beam", hence our use of the index b in θ_b) is usually so wide, that not more than a single "pixel" is in an "image". That is: each telescope simply measures

one signal, not an image with multiple pixels. For millimeter-wave telescopes we are in an intermediate regime where a single telescope still has a small enough primary beam that low-resolution "images" can be made. For instance, the SCUBA-2 instrument on the 15-meter James Clerk Maxwell telescope (JCMT) has a 32x40 detector array, sensitive to 450 μm and 850 μm radiation. Each pixel on the detector has a size of 0.11 cm, not much more than the wavelength of the radiation it is measuring. However, for radio telescopes with much longer wavelengths the "pixels" get also proportionally larger. For instance, for λ =21 cm (the neutral hydrogen hyperfine structure line) the minimum size of a resolution element on the focal plane is of the order of 21 cm, but in practice even larger. For that reason, rather than trying to make multi-pixel images with a single radio telescope, an image is usually obtained through scanning: each pointing of the telescope delivers a single "pixel" of the image. The field of view of each individual pointing is then in fact the same as the beam size: $\theta_f \simeq \theta_b$.

In the analysis of Sections 6.6, 6.8 and 6.9 we measured the interference between the primary beam of telescope 1 with the primary beam of telescope 2. Both beams must be pointing at the same object to within an accuracy of at least θ_b, of course, otherwise no interference can be expected. But in addition the the object we are interested in, any other objects in this primary beam may also participate in the interference. The question is: will the flux from some additional source at an angular distance θ from our source of interest (with θ smaller than the beam size) constructively interfere with our object, or will it simply dilute the coherence? Or equivalently: if we have two point sources A and B, separated by θ but both inside the primary beam, will they still interfere with each other, or are they too far apart to interfere? This is a concrete question, because there is an angular distance θ_c beyond which no interference is possible: the signals become decoherent. This is related to the fact that if we have a projected baseline length of L and we use e.g. source A to define our "pupil plane" (phase difference 0), then the wave front from source B has an angle with respect to the pupil plane that translates into a distance difference of $l = L\theta$. If that distance is larger than the coherence length, then the two sources A and B cannot interfere with each other. Thus $\theta_c = l_c/L$. With Eq. (6.4) and $l_c = c\tau_c$ we thus obtain

$$\theta_c = \frac{c}{L\Delta\nu} = \frac{\lambda}{L}\frac{\nu}{\Delta\nu} = \frac{\lambda^2}{L\Delta\lambda} \equiv \theta_i \frac{\lambda}{\Delta\lambda} \equiv \theta_b \frac{D}{L}\frac{\lambda}{\Delta\lambda} \tag{6.42}$$

Let us call this the *coherent field of view*. According to Eq. (6.42) this coherent field of view can be smaller than the primary beam if:

$$\frac{\Delta\nu}{\nu} \equiv \frac{\Delta\lambda}{\lambda} \gtrsim \frac{D}{L} \tag{6.43}$$

For sufficiently narrow bandwidth and sufficiently large dishes (or small baselines) we need not worry about the coherent field of view, as long as our source (or multi-source object) is smaller than the primary beam. But for very broad bandwidth or very large baselines we may need to worry about this. We can also see from Eq. (6.42) that the coherent field of view θ_c is larger than the interferometer resolution θ_i by a factor

$$\frac{\theta_c}{\theta_i} = \frac{\nu}{\Delta\nu} = \frac{\lambda}{\Delta\lambda} \tag{6.44}$$

So if we, for example, wish to perform interferometery in the N band, and taking the full N band (from 8 to 13 μm roughly) for our interferometry to gain sensitivity, then we have

$\lambda/\delta\lambda \simeq 2$, meaning that our coherent field of view is just twice as large as the resolution of our interferometer!

6.11 Heterodyne receivers

In radio interferometry the technology of *heterodyne receivers* is usually used. The idea is to convert the frequency of the incoming signal to a lower frequency that can be better handled, more easily amplified etc. This technique is also common in everyday life: typically radio or television receivers use this technology.

So how do we convert a signal at frequency ν_s to some intermediate frequency $\nu_0 \ll \nu_s$? We follow the discussion in the book by Thompson here (see literature list). Let us write the source signal at frequency ν_s as $V_s(t)$, which stands for the electric voltage in our receiver induced by the source on the sky that our antenna is receiving. This is obviously a very tiny voltage, as we wish to observe weak sources on the sky. Let us assume that this is a cosine wave:

$$V_s(t) = V_{s,0}\cos(2\pi\nu_s t + \phi_s) \tag{6.45}$$

where ϕ_s is simply some offset phase.

Now introduce a *local oscillator*: a device that produces a voltage $V_{\rm lo}(t)$ at some (different) frequency $\nu_{\rm lo}$:

$$V_{\rm lo}(t) = V_{\rm lo,0}\cos(2\pi\nu_{\rm lo} t + \phi_{\rm lo}) \tag{6.46}$$

Typically this voltage is much larger in amplitude than $V_s(t)$. But it has (or better: should have) a very high phase stability. Typically $\nu_{\rm lo}$ is chosen to be fairly close to ν_s, i.e. $|\nu_{\rm lo} - \nu_s| \ll \nu_s$.

If we now add $V_{\rm lo}(t)$ to our source signal $V_s(t)$ (this is called *mixing*) we get a signal that behaves like a modulated cosine: A cosine with an amplitude "envelope" that changes at a beat frequency $\nu_{\rm beat} = |\nu_s - \nu_{\rm lo}|$. While this beat frequency is now much smaller (and thus much more easily managable) than the original ν_s, this does not yet help us much, because this modulation is only an apparent wave, not a real one. To get a real signal at this beat frequency one must introduce some non-linearity in our electric circuit.

Suppose we put our mixed voltage signal on one end of a diode for which we know that the resulting current I is a highly non-linear function of the input voltage V (in its most extreme case $I = 0$ for $V < 0$ and $I = V/R$ for $V > 0$, where R is some resistance in front or behind the diode). Let us use a power series to model this non-linearity:

$$I(t) = a_0 + a_1 V(t) + a_2 V(t)^2 + a_3 V(t)^3 + \cdots \tag{6.47}$$

The first two terms are simply the linear response, and will just yield the amplitude-modulated cosine wave of the mixed signal. This is not interesting for us. But the quadratic term *is* interesting:

$$\begin{aligned}
V(t)^2 &= [V_s(t) + V_{\rm lo}(t)]^2 && (6.48)\\
&= [V_{s,0}\cos(2\pi\nu_s t + \phi_s) + V_{\rm lo,0}\cos(2\pi\nu_{\rm lo} t + \phi_{\rm lo})]^2 && (6.49)\\
&= V_{s,0}^2\cos^2(2\pi\nu_s t + \phi_s) + V_{\rm lo,0}^2\cos^2(2\pi\nu_{\rm lo} t + \phi_{\rm lo}) + && (6.50)\\
&\quad 2V_{s,0}V_{\rm lo,0}\cos(2\pi\nu_s t + \phi_s)\cos(2\pi\nu_{\rm lo} t + \phi_{\rm lo}) && (6.51)\\
& && (6.52)
\end{aligned}$$

The first two terms are signals at the original ν_s and $\nu_{\rm lo}$ frequencies and are not interesting in our quest for an intermediate frequency conversion of our source signal. But the last term contains the multiplication of our two waves, and here something interesting happens. Let us write out this last term:

$$\begin{aligned} \text{last term} &= 2V_{s,0}V_{\rm lo,0}\cos(2\pi\nu_s t+\phi_s)\cos(2\pi\nu_{\rm lo}t+\phi_{\rm lo}) && (6.53)\\ &= V_{s,0}V_{\rm lo,0}\cos\left[2\pi(\nu_s+\nu_{\rm lo})t+\phi_s+\phi_{\rm lo}\right]+ && (6.54)\\ &\quad V_{s,0}V_{\rm lo,0}\cos\left[2\pi(\nu_s-\nu_{\rm lo})t+\phi_s-\phi_{\rm lo}\right] && (6.55)\end{aligned}$$

So here we see that we now obtain the sum of two cosine waves at frequencies $\nu_s+\nu_{\rm lo}$ and $\nu_s-\nu_{\rm lo}$ respectively. These are not modulations, but actual signals at these frequencies. If $|\nu_{\rm lo}-\nu_s|\ll\nu_s$, then only the second of these terms (the one with $\nu_s-\nu_{\rm lo}$) will be interesting to us, because this gives a signal at a much lower frequency than the original frequency ν_s.

We now put this signal through a frequency filter that selects the frequency $\nu_0=|\nu_s-\nu_{\rm lo}|$ (which is traditionally called the *intermediate frequency*) with some bandwidth $\Delta\nu_0$ and damps out all other modes (including the local oscillator and the signal itself, as well as the $\nu_s+\nu_{\rm lo}$ frequency signal and any other signal not falling in the range $\nu_0\pm\Delta\nu_0/2$). We then obtain the desired frequency-reduced signal:

$$I_0=a_2V_{s,0}V_{\rm lo,0}\cos\left[2\pi(\nu_s-\nu_{\rm lo})t+\phi_s-\phi_{\rm lo}\right] \qquad (6.56)$$

For interferometry this equation has a very important property: The phase ϕ_s of the input signal is still conserved. If we later want to interfere this signal with a signal from another telescope that has a phase ϕ_t, then the phase difference between the signals, $\phi_t-\phi_s$, is still the original one, in spite of the mixing with the local oscillator and the down-grading of the frequency from ν_s to the intermediate frequency $\nu_0\ll\nu_s$, provided that our local oscillator has a very high phase stability so that we can remove the phase difference in $\phi_{\rm lo}$ from this final phase difference. The conservation of phase means that if we do interferometry with the intermediate frequency signal, we still get the same interference! We can thus safely use the intermediate frequency signal for any further interferometry dealings.

We are now almost there, but there is one more thing we have to clarify. Typically we do not receive a single-frequency signal from the sky. We receive a spectrum. The question is now: given some local oscillator frequency $\nu_{\rm lo}$ and some filter at frequency ν_0 and bandwidth $\Delta\nu_0$, which signal frequencies ν_s are we going to be sensitive to? The answer is: any ν_s for which $\nu_0=|\nu_s-\nu_{\rm lo}|$. This gives two solutions:

$$\nu_s=\nu_{\rm lo}\pm\nu_0 \qquad (6.57)$$

each with bandwidth $\Delta\nu_s=\Delta\nu_0$. Since $\nu_0\ll\nu_{\rm lo}$ this creates two sensitive bands that are close together in frequency. These are called the *lower sideband* and the *upper sideband*. This means that the intermediate frequency signal we work with is a mixture of the signals at two nearby frequencies. This is usually not what we want. Thus we need to filter out one of the sidebands *before* the mixer, so that we are only sensitive to the other sideband.

From here on, everything we learned about long baseline interferometry can now be applied to his intermediate frequency signal.

6.12 When telescope size is not negligible compared to the baseline

If we perform true *long baseline interferometry*, where $b \gg D$, we can ignore the finite size of the individual telescope apertures, and regard our visibility measurement as a measurement at a single (u, v) position. If, however, we put our telescopes so close to each other that b is only a few times their diameter (possibly even down to $b = D$ where the mirrors touch each other, but typically at somewhat larger distance), then we have to have a closer look at what we are doing. We already saw in Section 6.1 an example of such a situation: The LBT in Arizona, where two 8.4 meter telescope are arranged in a baseline of 14.6 meters, i.e. there is merely 6.2 meters between the edges of the mirrors. If we combine the light of the two telescope at the image plane (Fizeau interferometry), we have already seen that we get a PSF which is a single-mirror PSF with a fringe pattern on top. We obtained this pattern simply by Fourier transforming a wave front passing through our double-pupil:

$$\text{Image} = |\mathcal{F}[E](\vec{x})|^2 \tag{6.58}$$

where $E(x, y, \text{pupil}) = 0$ there where (x, y) is not part of the double-pupil. If the wave passing through the double-pupil is a plane wave, the function $E(x, y, \text{pupil})$ depicts the shape of the double-pupil (see Fig. 6.1-left), and the image is the PSF (see Fig. 6.1-right).

We can now link this "Fizeau inteferometry language", where we think in terms of the PSF in the image-plane, to "long baseline interferometry language", where we think in terms of measuring the autocorrelation of the radiation field in the pupil-plane. We do this by using the spatial version of the Wiener-Khinchin theorem (Eq. B.10 of Appendix B.1), applied to the wave:

$$\text{Image} = |\mathcal{F}[E](\vec{x})|^2 = \mathcal{F}[\text{Corr}[E, E]](\vec{x}) \tag{6.59}$$

In other words: the image on the image plane is in fact the Fourier transform of the autocorrelation function in the pupil plane! If we compare this to what we learned from the van Cittert-Zernike theorem, this makes perfect sense: the image plane is indeed a projection of the sky, and we already knew that the Fourier transform of the autocorrelation function in the pupil plane represents the image on the sky.

By applying an inverse Fourier transform to Eq. 6.59 we get

$$\text{Corr}[E, E](\vec{u}) = \mathcal{F}\left[|\mathcal{F}[E](\vec{x})|^2\right] = \mathcal{F}[\text{Image}] \tag{6.60}$$

If we have a plane wave hitting our pupil, then the "Image" is the PSF, so the above equation then gives the autocorrelation of the pupil. This function is called the *modulation transfer function* (MTF). Remember, by the way, the telescope transfer function, Eq. 4.54. It is the same thing. Here we just apply it to the two-mirror system as a whole.

Let us apply Eq. 6.60 to the PSF we obtained for the LBT (Fig. 6.1-right). So we Fourier transform the PSF (which itself is the absolute-value-squared of the Fourier transform of the pupil) to the pupil plane. What we get is shown in Fig. 6.4. Note that we Fourier transform the image-plane electric field squared $\langle E^* E \rangle$, not the image plane field E itself, otherwise we would have obtained Fig. 6.1-left back. Instead, we now obtain a triplet of blobs, partly overlapping each other, and each being the "brightest" in their center and dropping off toward their edge. The middle blob corresponds to all possible vectors

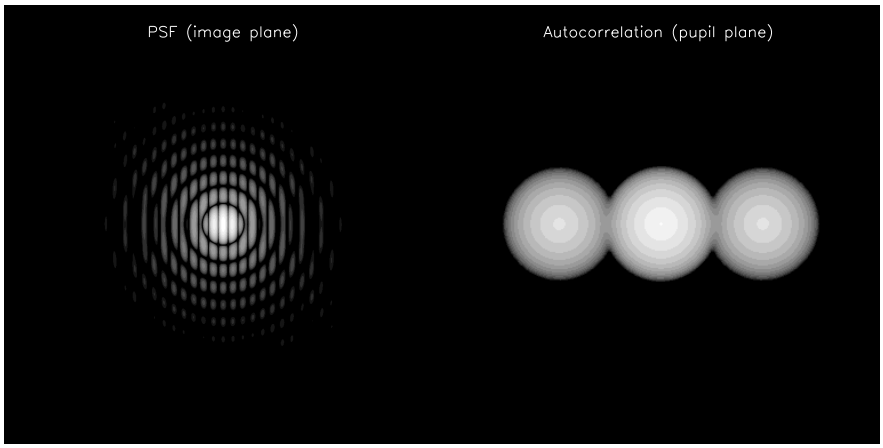

Figure 6.4: Left: the PSF of the Large Binocular Telescope (identical to Fig. 6.1-right). Right: The Fourier transform of this PSF back to the pupil plane.

$\vec{r}$ connecting two points on the *same* mirror, while the two sidelobes correspond to all possible vectors $\vec{r}$ connecting one point on one of the mirrors with another point on the other mirror.

Contrary to the pupil shape function shown in Fig. 6.1-left, the autocorrelation function in Fig. 6.4-right is not a constant function over some area and dropping to zero outside of that area. Instead, the blobs are centrally bright and fade toward their periphery. This has a simple explanation: some vectors $\vec{r}$ connecting two points on the double-pupil may "fit" to a much more limited set of pairs of points on the pupils than others. For instance for a single LBT mirror (8.4 meter diameter) a separation of 8.4 meter ($|\vec{r}| = 8.4$m) is only represented by pairs of points on the telescope edge that are diametrically opposite, while if ($|\vec{r}| = 1$cm) almost every point $\vec{x}$ on the mirror has a point $\vec{x} + \vec{r}$ that is also on the mirror. This means that the autocorrelation function is much better sampled at small $\vec{r}$ than at large $\vec{r}$. A similar reasoning holds for the $\vec{r}$ vectors that link a point $\vec{x}$ on one mirror to a point $\vec{x} + \vec{r}$ on the other mirror: hence also the "sidelobes" in Fig. 6.4-right fade toward their periphery.

Now here we see a difference between the Fizeau-style measurements of $\mathrm{Corr}[E, E]$ and the long-baseline-style measurements of $\mathrm{Corr}[E, E]$; a difference that becomes apparent when the mirrors are close together. For Fizeau interferometry for closeby telescopes the *true* $\mathrm{Corr}[E, E]$ function (before entering the aperture of our pair of telescopes) is strongly modulated by the finite size of the telescope. For ideal long-baseline interferometry (zero telescope diameter) we actually measure the true $\mathrm{Corr}[E, E]$ function without this modulation. In principle the measurement of the fringes of the image-plane image of our Fizeau interferometer represent a measurement of $\mathrm{Corr}[E, E]$ in the pupil plane, but instead of measuring a single (u, v) point, we measure a non-trivial integral of $\mathrm{Corr}[E, E]$ values modulated by the *modulation transfer function* shown in Fig. 6.4-right. If we now push the telescopes apart (of course keeping their diameter fixed), then the blobs in Fig. 6.4-right also move apart (while staying the same size). By measuring the fringes, we in fact measure the $\mathrm{Corr}[E, E]$ in the two "side-lobes" of Fig. 6.4-right, which now become more and more point-like compared to their distance from the origin. We thus approach the long-baseline interferometry-style measurement of $\mathrm{Corr}[E, E]$.

6.13 Pupil-plane versus image-plane interferometry

The modulation of the autocorrelation function with the MTF, discussed in Section 6.12, applies to the case when we combine the light of the two telescopes directly on the image plane. However, one can also combine the light of the two telescopes at the pupil plane. The idea is to create an *exit pupil* for each telescope, which is a miniature version of the telescope *entry pupil* (aperture), and now overlay the two exit pupils of the two telescope precisely over each other by using a beam combiner. The result is a wave at the combined exit pupil that is the direct sum of the waves in the two entry pupils.

$$E_{\text{combined}}(\vec{x}') \propto E(\vec{x}) + E(\vec{x} + \vec{r}) \tag{6.61}$$

where $\vec{x}'$ is the scaled-down exit-pupil version of the $\vec{x}$ on the primary mirror (entry pupil) and $\vec{r}$ is the baseline vector connecting the two telescopes. In other words: by combining the beams in the pupil plane instead of in the image plane, you get interference always between a position on one primary mirror with the equivalent position on the other primary mirror. This means that we are now measuring interferometry truly at a single (u, v) point, in spite of the non-negligible size of the telescope. And since both mirrors are of the same shape and size, there is no modulation of the correlation function by the telescope system. If we observe a single unresolved point source on the sky, what we observe on the image plane is that the entire PSF fades and brightens as you make slight changes of the delay between the two telescopes. This is indeed the fringe pattern we encountered in the long-baseline interferometry and we can thus measure the visibility.

Pupil-plane beam combining and image-plane beam combining both have their advantages and disadvantages. The nice thing about image-plane interferometry is that one gets true images, if one manages to "deconvolve" the measured images with the Fizeau PSF. It does not suffer from the field-of-view limitation due to decoherence. But it requires very good adaptive optics to work, and it only works well for telescope pairs that are close, i.e. the gain in spatial resolution is moderate. The nice thing about pupil-plane interferometry is that it does not suffer from the non-trivial MTF of the pair of telescopes, and it can handle very large baselines and thus obtain incredible spatial resolution. But the field of view is limited by decoherence if the wavelength band $\delta\lambda$ is large (which is usually only a concern in infrared and optical interferometry, less so in millimeter/radio interferometry). And also each measurement is just a single (u, v) point, and one needs many baselines to construct a decent image.