High-angular resolution imaging of disks and planets

Gaspard Duchêne [a,b]

[a]*Laboratoire d'Astrophysique de Grenoble – Université Joseph Fourier – BP 53 – 38041 Grenoble cedex 9 – FRANCE*

[b]*Department of Astronomy – 601 Campbell Hall – University of California Berkeley – Berkeley CA 94720-3411 – USA*

Abstract

During the last two decades, the advent of high-angular resolution imaging techniques on large optical and infrared telescopes has generated an empirical revolution in many aspects of astrophysics. In this contribution, I describe some of the techniques developed in recent years to obtain high-angular resolution images and present their advantages and limitations. I then illustrate the key role played by these techniques by presenting some of the most important results obtained in the study of extrasolar planets and planet-forming disks around stars other than our Sun.

Key words:

1 Foreword: why care about high angular resolution?

Since Galileo pointed the very first telescope towards the sky 400 years ago, astronomers continuously pushed technical limits to build telescopes of increasing size. We have now reached an era in which 8–10 m telescopes are the norm and 30–40 m telescopes are less than a decade away. Astronomers' quest for increasingly larger telescopes was first motivated by the collecting area advantage provided by a large telescope. While this is still a crucial argument, a second factor is now becoming equally important in the mind of telescope-builders, namely angular resolution.

Email address: `gaspard.duchene@obs.ujf-grenoble.fr` (Gaspard Duchêne).

Geometrical optics teaches us that the image of a point source (e.g., a star at infinity) observed through a telescope whose primary mirror is a perfect circle is a so-called Airy pattern, which is characterized by a diffraction-limited core and a series of concentric annuli. The typical size of the Airy pattern, measured as its full width at half maximum (FWHM) [1] , is $\theta_{diff} \approx \lambda/D$ (expressed in radians) where λ is the observing wavelength and D the primary mirror diameter (both expressed in meters). All objects observed through that telescope will produce images that are the result of the convolution of their intrinsic intensity distribution ("image") by that Airy pattern. Therefore, anything substantially smaller than θ_{diff} will remain indistinguishable from a point source and its image will be a perfect Airy pattern.

In order to resolve small details in the image an astronomical object, say tiny craters on the Moon, the irregular surface of giant stars, or the inner structure of distant galaxies, *one needs a telescope that is large enough so that* $\theta_{diff} < \theta_{detail}$. Let us now consider some numbers to understand the implications of this requirement. First of all, the diffraction limit can be re-expressed as $\theta_{diff} \approx 0\rlap{.}{''}2 \times \lambda(\mu\mathrm{m})/D(\mathrm{m})$. Unfortunately, when observing from the ground, light from astronomical objects has to travel through the Earth's atmosphere before reaching the telescope. As we will see in Section 2, the resolution of visible and near-infrared images is limited to $\theta_{atm} \approx 0\rlap{.}{''}5$–$2''$ due to atmospheric turbulence, so that on all large telescopes, images have a spatial resolution that is dictated by the atmosphere, not by diffraction off the telescope mirror. In other words, the atmosphere prevents large telescope from yielding their full resolutive power.

Despite these limitations, astronomy has made huge progress with normal observations using large telescope, i.e., with turbulence-limited observations. So one may wonder whether it is important to devise elaborate techniques or instruments to achieve the best possible resolution with a large telescope. To answer that question, one needs to address another one: How does θ_{atm} compare with the apparent size of interesting astronomical objects? Of course, objects of all sizes exist in the sky. Here is a list that goes all the way from our Solar System to distant galaxies:

– Radius of Jupiter's Galilean moons at the time of Earth-Jupiter opposition, as seen from Earth: $\sim 1\rlap{.}{''}5$;
– Radius of Betelgeuse (α Ori), as seen from Earth: $0\rlap{.}{''}055$;
– Sun-Jupiter distance, as seen from $10\,\mathrm{pc}$ away: $0\rlap{.}{''}5$;
– Radius of a typical protoplanetary disk, as seen from $150\,\mathrm{pc}$: $\sim 2''$;
– Half-light radius of M31 globular clusters, as seen from the Earth: $\lesssim 1''$;

---

[1] The radius of the first null of the Airy function, $r_{null} = 1.22\lambda/D$, is also frequently referred to as the "diffraction limit". This is the formal Rayleigh criterion for diffraction, which is comparable to θ_{diff} adopted here for simplicity.
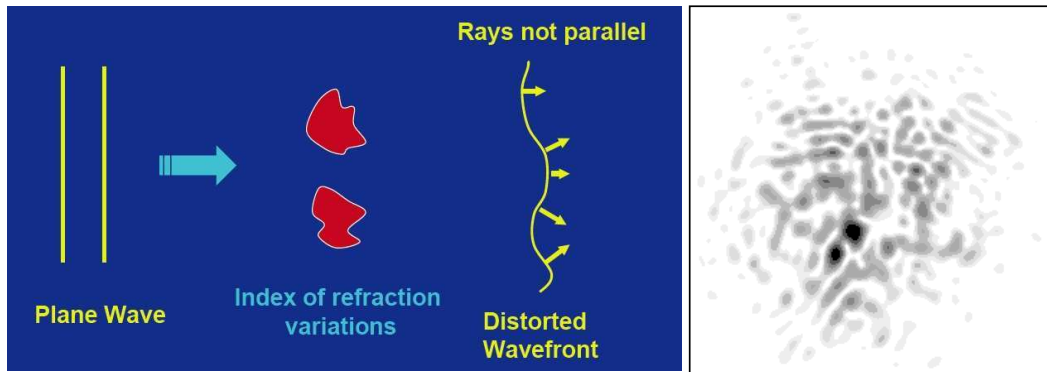
Fig. 1. *Left:* Effect of the patchy atmosphere on incoming light from a star at infinity. The wavefront is planar before entering the atmosphere, but it is corrugated after traveling through it. Credit: Claire E. Max (UCSC). *Right:* Resulting short-exposure image ($t_{int} \ll 1\,\mathrm{s}$), composed of many independent diffraction-limited speckles. Each speckle has a FWHM equal to λ/D but they are spread over an area whose FWHM is the direct imaging "seeing".

– Radius of typical galactic disks at z=0.5: $\sim 0\!''\!.25$.

If one hopes to resolve some structure in these (and many other) objects, diffraction-limited images from large telescopes are required. This is the reason why astronomers have developed several techniques to overcome the limitations imposed by atmospheric turbulence. These techniques, whose aim is to recover the diffraction limit that a telescope offers in theory, are generally called "high-angular resolution techniques". They will be described in detail in Section 3, together with a preview of the next generation techniques, before presenting some of the major advances they have already permitted in the field of extrasolar planets and protoplanetary disks in Sections 4 and 5.

Throughout this contribution, we focus on a range of wavelength spanning the visible and the near- (and mid-) infrared. Observations at other wavelengths are either impossible from the ground (γ rays, X rays, ultraviolet, far-infrared) or with such different detection techniques that the role of atmospheric turbulence is totally different and beyond the scope of this review (radio regime).

2 The Earth's turbulent atmosphere

2.1 The atmosphere's influence on astronomical images

Let us consider a star at infinity. Its light waves reach the Earth as plane waves, which is what a telescope has been designed to make an image of. This plane wave, as it travels through the atmosphere, encounters regions
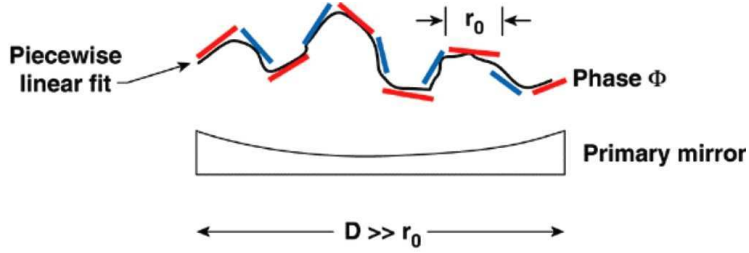
Fig. 2. Idealization of the incoming corrugated wavefront as a series of small, inclined planar wavefronts. Credit: Claire E. Max (UCSC).

of different temperature and pressure. Because the air has a temperature-dependent refractive index, each layer of the atmosphere deflects the incoming starlight differently. Worse, each layer in the atmosphere is actually made of independent patches of air, with different temperature and physical conditions. As illustrated in Fig. 1, the incoming wavefront is therefore distorted before it reaches the ground. In other words, the patchy atmospheric layer acts as a phase screen on the incoming starlight, in the sense that some parts of the wavefront reach the telescope mirror slightly after/before others.

When this corrugated wavefront reaches the telescope, it does not create a clean Airy pattern image, as it should, because the wavefront is not planar (or perfectly cophased) anymore, as illustrated in Fig. 2. In a first approximation, we can consider this wavefront as a collection of small singe-phase subsections. Since constructive light interference requires cophasing, only the subset of all subsections corresponding to a given phase can interfere together. Since they are spread over the entire primary mirror, they will create an Airy pattern whose resolution corresponds to the diffraction limit of the mirror, albeit with a reduced intensity with respect to the image resulting from a plane wave. The corrugated wavefront, which is the sum of many independent such collections of subsections, will create a large number of Airy functions, each shifted from one another because of their different phase (or, equivalently, their common local wavefront slope). These individual Airy functions are called "speckles"; the right panel in Fig. 1 shows the result of a short integration with a turbulence-dominated telescope.

The location of the speckles is random around the image center, with a probability distribution that is roughly Gaussian. Each short integration therefore results in a different "speckle cloud". This is a consequence of wind, which moves around the patches of air above us and, to a smaller extent, of convection in the atmosphere. A familiar consequence of these motions is the scintillation of star in the night sky that we can see with the naked eye. If we integrate for a "long" time with an astronomical camera (we will see later that 1 second is a long time), the speckles eventually completely overlap each other, resulting in a smooth, broad Gaussian profile, whose FWHM is the atmospheric "seeing", θ_{atm}. All diffraction-limited information is lost in such a
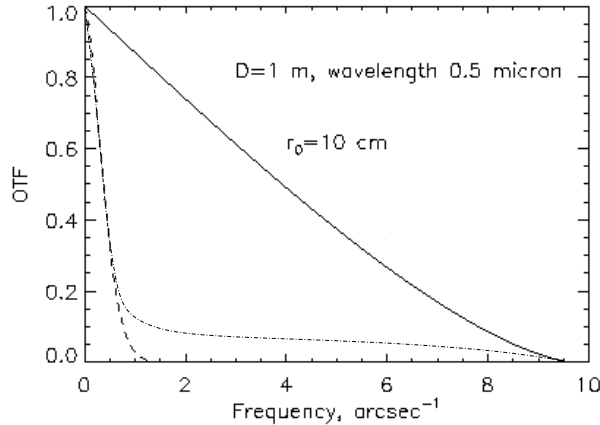
Fig. 3. Optical transfer function of a perfect telescope (solid line) and of the same telescope observing through a turbulent atmosphere (dashed line: long integration; dot-dashed line: short integration). Credit: Andrei Tokovinin (CTIO).

long exposure. The seeing, usually defined at $0.5\,\mu$m, varies from $0''\!.5$ to several arcseconds, depending on the quality of the observatory site.

In essence, the atmosphere acts on the incoming image as low-pass filter on the spatial frequencies. Fig. 3 illustrates the optical transfer function of an atmosphere-free telescope (solid curve) and of a "normal" telescope in both short- (dot-dashed curve) and long-integration modes (dashed curve). A telescope observing without atmospheric turbulence has a non-zero response to all spatial frequencies up to D/λ, while the same telescope affected by turbulence has a zero-response for frequencies above r_0/λ for long integrations (see Section 2.2 for a definition of the Fried parameter, r_0). In a short exposure, although the high spatial frequencies are strongly damped, they are not completely cut off and it is still possible to study fine details in the objects, as we will see in Section 3; this is a consequence of the presence of diffraction-limited speckles in the short exposure image shown in Fig 1. The optical transfer function of a telescope observing through a turbulent atmosphere was estimated by Fried (1966) in a classical work in which he defined the r_0 parameter.

2.2   Characteristics of atmospheric turbulence

What are the basic properties of the atmosphere and its turbulence? Answering this question will provide us with the key information needed to devise ways of circumventing the "spatial resolution problem" inherent to all ground-based observations.
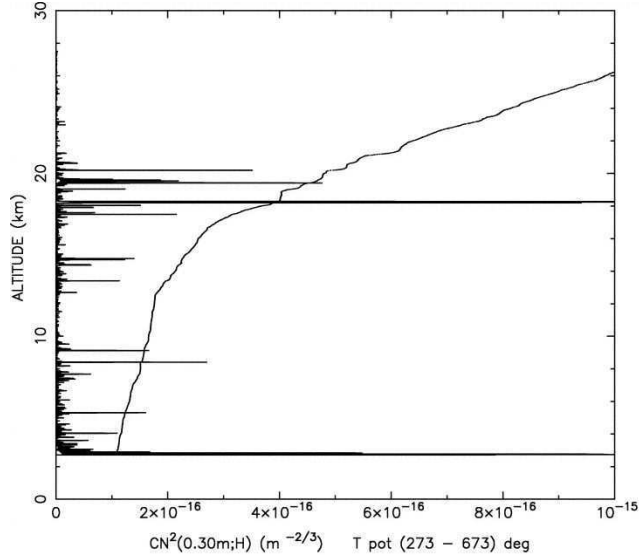
5

Fig. 4. Vertical profile of the atmospheric turbulence above Cerro Pachón, site of the Gemini telescope in Chile, shown as the horizontal histogram. C_n^2 is the measure of local strength of the turbulence. Most of the turbulence is concentrated in 2 separate layers, as is the case for the site of most observatories: the ground layer, related to mountain- and dome-induced non-laminar air flows, and a high-elevation layer (at about 18 km above sea level in this case). The solid curve shows the variation in air temperature (declining towards the right-hand side of the plot). From Azouit & Vernin (2005).

First of all, what are the typical spatial/angular scales associated to the atmosphere irregular composition/temperature? The first and foremost parameter that we can define is the so-called Fried parameter, r_0, which is the coherence length of the atmosphere (see Fig. 2). This is the size of the subsections of the corrugated wavefront that can be considered as cophased. In the best astronomical sites, this quantity is on order $r_0 \approx 10$–20 cm when considering visible light (all quantities quoted here correspond to $\lambda = 0.5\,\mu$m, unless otherwise stated). This means that if we use a telescope whose diameter is larger than r_0, the images it produces are dominated by atmospheric turbulence, since the incoming wavefront is globally decorrelated. The "seeing" is nothing more than the diffraction limit associated to this coherence length: $\theta_{atm} \approx \lambda/r_0$. Here we see that even for modest telescope sizes, we have to worry about atmospheric turbulence.

A second key parameter is how fast the phase screen changes. The timescale associated with the loss of coherence of the atmosphere is called the coherence timescale, τ_0. At first order, this is dictated by the wind speed in the turbulent layer. Typical windspeed are in the range $V \approx 10$–$50\,\mathrm{m\,s}^{-1}$, so that after just a few millisecond, the phase screen has lost its coherence and the speckle cloud is completely independent from the previous one. This is why a 1-second exposure can be considered as a "long integration" for visible images. In a short exposure, the number of independent speckles roughly scales as $t_{int}/\tau_0(D/r_0)^2$.
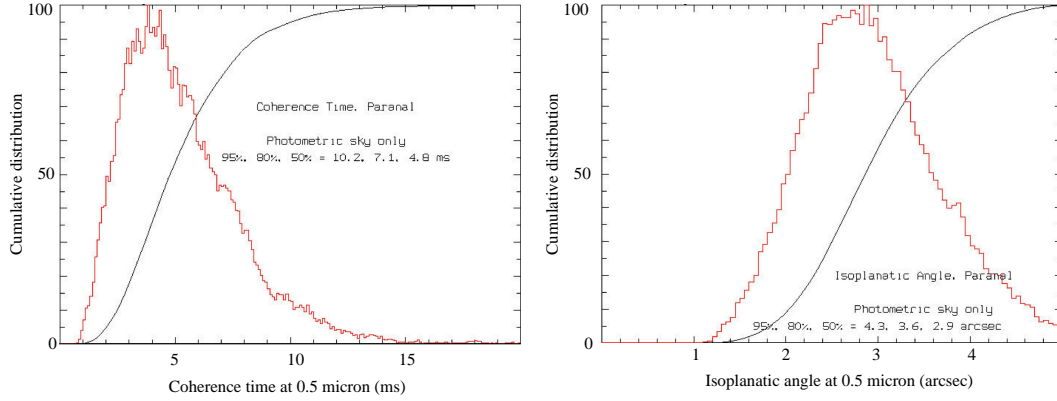
6

Fig. 5. *Left:* Histogram and cumulative distribution of the atmosphere's coherence timescale (τ_0, at $0.5\,\mu$m) at Cerro Paranal, site of ESO's VLT in Chile, over the year 2003. *Right:* Histogram and cumulative distribution of the atmosphere's isoplanatic angle (θ_0, at $0.5\,\mu$m) for the same period. Credit: ESO.

A third important quantity is the so-called isoplanatic angle, θ_0, which is the "coherence angle", or the angular size over which the turbulence can be considered as coherent. This can be derived from r_0 and the elevation of the turbulent layer. In principle, the atmosphere is turbulent at all elevations above the telescope. In practice, however, this turbulence is mostly located in two layers at most observatory sites: the ground layer (this is a consequence of the non-laminar flow induced by the mountain and the dome itself) and a certain high-elevation layer, several km above (see Fig. 4). Focusing on the latter, we see that the isoplanatic angle is typically a few arcseconds at Cerro Paranal. Fig. 5 shows histograms and cumulative distributions for the measured coherence time and the isoplanatic angle for the year 2003. There are some nights where the conditions can be great ($\tau_0 \geq 10\,\text{ms}$, $\theta_0 \geq 4''$), but most of the time, the turbulence is both very fast and associated with small spatial scales.

The last important aspect that describes atmospheric turbulence is its wavelength dependence. Indeed, the refractive index of the atmosphere becomes gradually more uniform at longer wavelength, so the deflections of incoming waves are smaller and the resulting wavefront is less distorted. Under the (reasonable) approximation that atmospheric turbulence can be described with a Kolmogorov model, it can be demonstrated that $r_0(\lambda) \propto \lambda^{6/5}$. So, if $r_0 = 15\,\text{cm}$ at $0.5\,\mu$m (corresponding to a $0''\!.7$ seeing), the coherence length is about $90\,\text{cm}$ at $2.2\,\mu$m and $5.5\,\text{m}$ at $10\,\mu$m. In other words, even the largest telescopes on Earth are almost diffraction-limited in the mid-infrared. Since both τ_0 and θ_0 are linearly correlated to r_0, one realizes that *fighting atmospheric turbulence is harder in the visible than in the near-infrared for three converging reasons: the number of speckles becomes much higher, they randomly move much faster, and any type of correction will apply to an increasingly smaller patch of the sky.* This has direct consequences on the methods used to fight turbulence, as we will see later on.
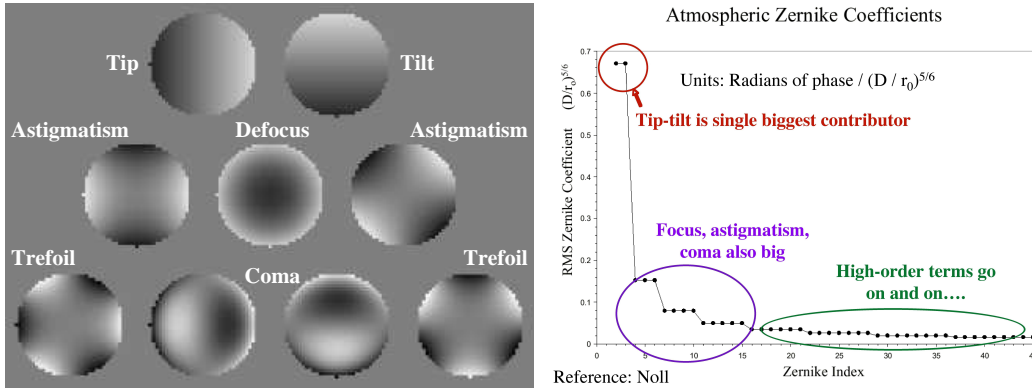
7

Fig. 6. *Left:* The first Zernike modes: atmospheric turbulence transforms an initially horizontal planar wavefront into a corrugated wavefront that can be described as a linear combination of Zernike modes. *Right:* The relative power of the Zernike modes under the assumption that atmospheric turbulence follows a Kolmogorov spectrum. Credit: Claire E. Max (UCSC), based on Noll (1976).

As we have discussed above, atmospheric turbulence transforms an incoming planar wavefront into a distorted one. One way to describe the latter is to decompose it into independent "modes". The most convenient and frequently used projection system are the so-called Zernike modes, the first few of which are illustrated in Fig. 6. The top two modes, collectively referred to as "tip-tilt", correspond to overall slopes in the wavefront which simply result in image displacement. The right panel of Fig. 6 shows that distorted wavefronts are usually well described by the first 10–20 Zernike modes, with decreasing power associated to modes of increasing order. In reality, there is one more mode, the zeroth-order mode, which is called "piston". It is an overall phase shift, constant over the entire telescope diameter, from the original plane wave. It has no consequence for image formation for a single telescope since it acts as a simple temporal delay on the wavefront. It is not discussed further here, but one must remember that it plays a crucial role in long baseline interferometry (because different telescopes suffer from different piston values).

3   High-angular resolution imaging techniques

In this section, we discuss three major techniques used for fighting atmospheric turbulence on single telescopes, in order of complexity. We will conclude with a rapid comparison of the performances offered by these techniques. We do not discuss the principles of long-baseline interferometry, which are discussed elsewhere in this volume, but remind the reader that interferometers using large individual telescopes (such as VLTI and Keck) are usually fed with adaptive optics-corrected beams, so these two techniques are not completely independent.
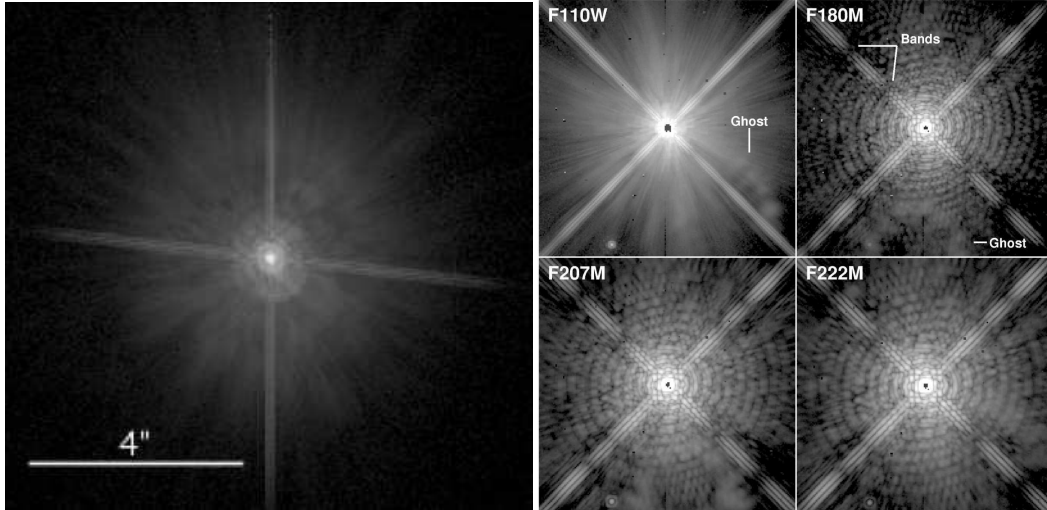
8

Fig. 7. Examples of images of point sources observed with HST, using ACS (*left*, $\lambda = 0.6\,\mu$m) and NICMOS (*right*, with several filters). A logarithmic stretch has been used to highlight the faint structure at large radii; the central parts of the NICMOS images are saturated, resulting in small dark disks. The fact that the telescope cannot be approximated by a perfect circular mirror (central obscuration, secondary mirror support, ...), as well as small optical defects, account for the low-level departure from perfect Airy functions. A faint companion is seen in the bottom left part of the NICMOS images.

3.1 Going above the atmosphere

In principle, the simplest method to beat atmospheric turbulence is to place one's telescope above it. This is one reason why most modern ground-based observatories are located at high elevations. However, as high as one can safely place an observatory (e.g., with sufficient atmospheric pressure for workers to actually be able to breathe while they build it!), the main turbulent layer remains higher than the mountain itself (see Fig. 4). In addition, the "ground layer" of turbulence is, by definition, always located where the telescope is. Therefore *it is not possible to entirely get rid of the turbulence by simply choosing a higher mountain as observing site.* Although turbulence is lower at higher elevations, ground-based telescopes are located at high elevation mostly in order to 1) reduce the water vapor content in the atmosphere upward, for mid-infrared observations in particular, 2) to stay as far away from light pollution as possible, and 3) to be located above the main layers of low-elevation clouds.

Since no mountain is high enough, let us place our telescope in space. This is precisely the idea behind a telescope like the Hubble Space Telescope (hereafter, HST). Obviously, atmospheric turbulence is not a factor anymore and the images obtained by HST are close to being diffraction-limited at all wavelengths (see Fig. 7). One just needs to build cameras that have fine enough a spatial sampling to take full advantage of it. In addition, placing the telescope
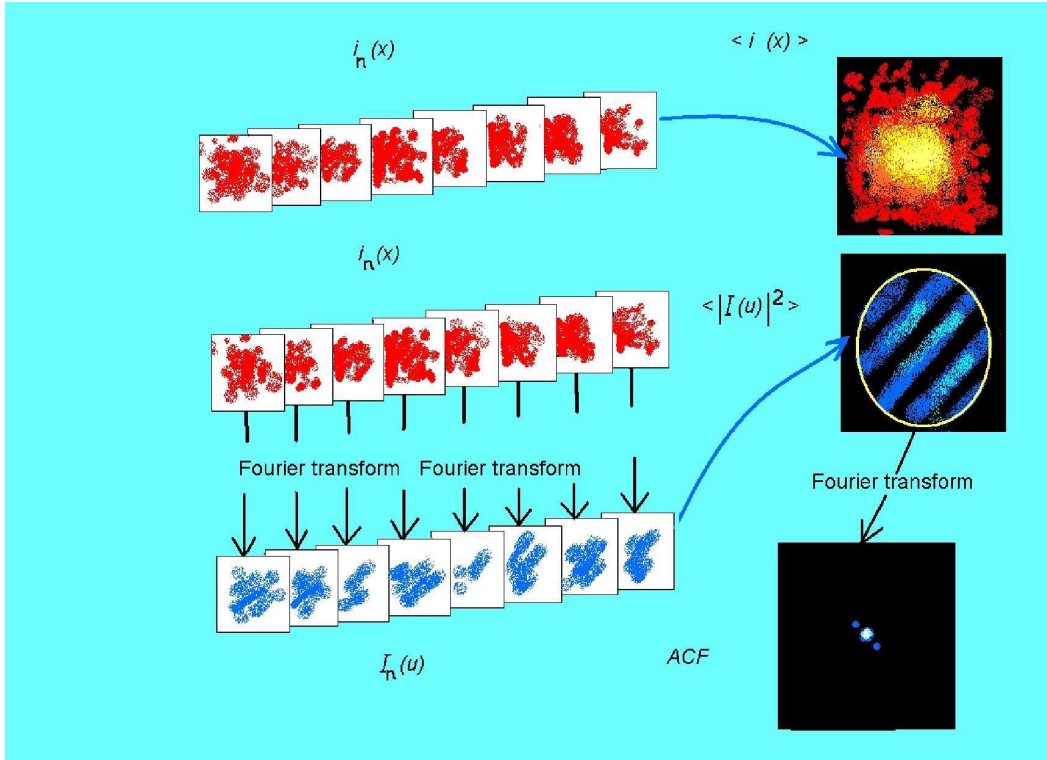
9

Fig. 8. Illustration of the bispectrum speckle interferometry method. Many short exposure are obtained, each of them is Fourier transformed. The Fourier power spectra are averaged to a single one, which is then inverse-Fourier transformed to obtain the final image. Credit: Yuri Balega (SAO, Russia).

above the atmosphere opens up the possibility to conduct observations in otherwise unattainable wavelength regimes. For instance, ultraviolet observations are impossible from the ground, and HST has long been a key observatory for this regime (until the demise of the STIS instrument). Similarly, the Spitzer Space Telescope (as well as IRAS and ISO before it) allows diffraction-limited in the mid- to far-infrared regime which is almost completely blocked by our Earth's atmosphere.

Of course, these advantages come at a price. First of all, there is an actual economical price, since space missions are always extremely expensive. In the case of HST, the price is even higher since human space-flight, aboard the Space Shuttle, was needed to launch the observatory and is still needed for the so-called servicing missions during which key elements and instruments can be changed to improve the capabilities of HST. Beyond the economical aspect, it is currently simply impossible to send in space an 8–10 m telescope, which results in a limited achievable spatial resolution (HST has a diameter of only 2.4 m, and Spitzer is an 85 cm-telescope). The largest optical telescopes in the world are bound to be ground-based, below atmospheric turbulence. The obvious other limitation associated with space-based telescopes is that it is difficult, if at all possible, to go and repair the telescope when something
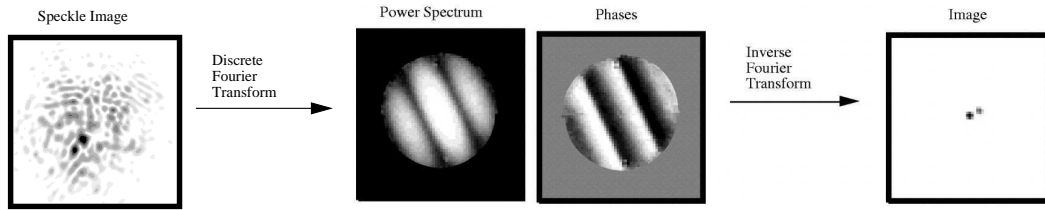
Fig. 9. Illustration of the bispectrum speckle interferometry method for a $0\overset{''}{.}14$ binary systems observed at $2.2\,\mu\mathrm{m}$ from the $5\,\mathrm{m}$ Hale telescope on Mount Palomar. The binary system, which is clearly detected in the post-processed image, is undetected in the raw speckle frames. Adapted from Patience et al. (1998).

goes wrong. As robust as they can be made, space telescope eventually will encounter a problem that could take just a few minutes/hours to repair if it were on the ground but that may compromise the entire mission in space. Also, to survive the rough, vibration-filled, journey to space, the telescope and its instruments have to be built with a well-proven, sometimes older technology. This sometimes results in smaller and noisier detectors than are available for ground-based instruments. In the particular case of HST, this problem is somewhat alleviated by the regular servicing missions, but this is impossible for most space telescopes. Overall, these limitations explain why astronomers keep building large telescopes on the ground, which offer nice a complement to space-based telescopes.

3.2   Speckle interferometry

To beat the atmosphere when observing from the ground, the easiest thing to do is to *take images very fast so that the various independent speckle do not have time to incoherently overlap with each other to produce a seeing halo*. Clearly, the images shown on the right hand side of Fig 1 contains information at the highest spatial frequencies (see also Fig 3); it is just a matter of finding an appropriate post-processing method to take full advantage of it. Typically, modern cameras can obtain images with integrations on the order of $0.1\,\mathrm{s}$ in the near-infrared and as short as $10\,\mathrm{ms}$ in the visible, effectively "freezing" the atmospheric turbulence.

To reach a satisfying signal-to-noise ratio, one usually obtains many (several thousands, typically) successive speckle frames. The data analysis can be performed in two ways. The easiest one consists in identifying the brightest speckle in each frame and to shift-and-add all frames to that reference point. With this method, simply called "shift-and-add", one obtains an image whose core is diffraction-limited since it is the superposition of many speckles at the very same location. However, the outer parts of the image correspond to the overlap of all those speckles that were not centered in the individual frames. There is therefore a seeing halo that prevents high-contrast imaging, although the
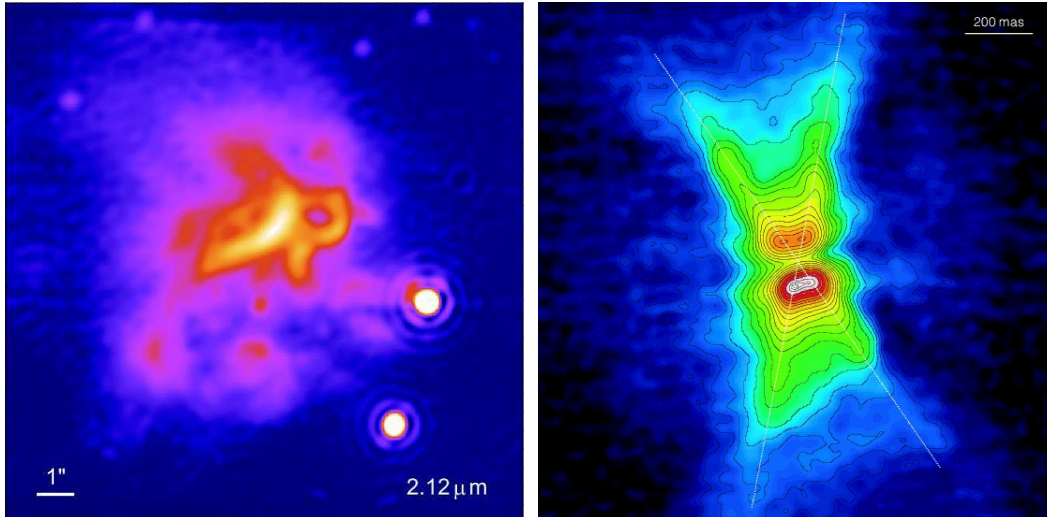
Fig. 10. Speckle interferometry imaging of two complex objects: the massive protostellar outflow source IRAS 23151+5912 (*left*: 2.2 μm image from Weigelt et al., 2006) and the Red Rectangle planetary nebula (*right*: 2.2–3.3 μm false color image from Tuthill et al., 2002).

achievable contrast depends both on sky conditions and instrumental setup. The resulting images are quite similar to those presented in Fig. 16, although the latter have been obtained with an adaptive optics system.

A slightly modified version of the shift-and-add method is called "Lucky Imaging". Quite simply, this method relies on "luck" or, more precisely, on statistical deviations far away from the mean. In each individual frames, the location of the speckles is random, with a density probability that is roughly Gaussian. The number of speckles per frame also varies randomly around its average value. Every once in a while, there are images in which the number of speckles is remarkably small and/or all the speckles are located within a much smaller area than the seeing halo. In these images, there seems to be a single bright speckle and only few faint ones around it. Therefore, one can go through thousands of speckle frames and select the best ones, maybe just a few percent of all frames. Shift-and-adding these frames will result in images that have a strongly reduced seeing halo and, therefore, a much higher contrast than regular shift-and-add. Obviously, this comes at the cost of throwing away 90% of the images or more, so that there is a trade-off between the achievable contrast and the signal-to-noise ratio in the final image.

Both of these methods, however, do not take full advantage of the information available in the raw speckle frames, as they end up merging together many speckles into the speckle halo. The loss of high spatial frequency information stems from the fact that these speckles are combined incoherently. A more elaborate method to overcome this problem, called speckle interferometry, was first proposed by Labeyrie (1970); this method is illustrated in
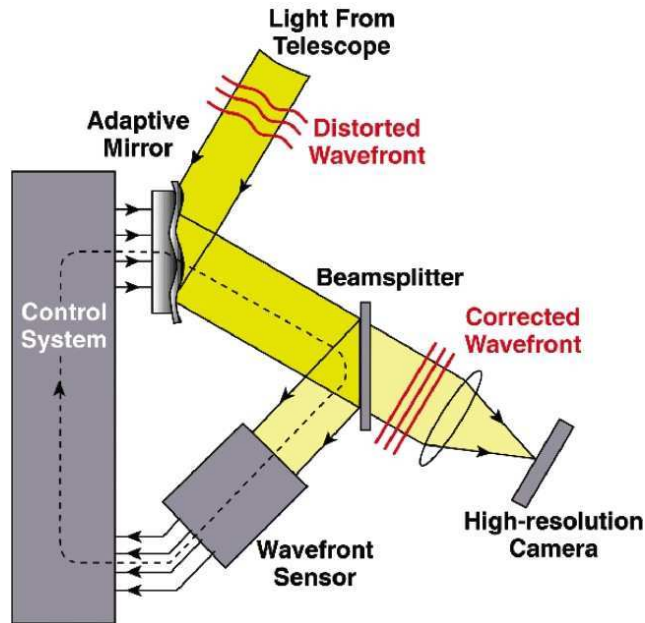
12

Fig. 11. Schematics of an AO system. Credit: Claire E. Max (UCSC).

Fig. 8. To allow for coherent addition of all speckles, the basic idea consists in computing the Fourier transform of raw speckle frames. The key point here is that all speckles are identical: they all are diffraction-limited images of the object (i.e., a convolution of their intrinsic image with the telescope optical transfer function), and they only differ by their location in the image plane or, equivalently, by their phase in the Fourier plane. If one observes just before/after the target an unresolved reference star, its Fourier transform will be a measure of the "atmospheric speckle transfer function", mostly a map of the phase deviations induced by the atmosphere. Dividing the target's Fourier transform by that of the reference star, one obtains a power spectrum, i.e., the square of the Fourier transform amplitude, that is therefore independent of atmospheric turbulence since all speckles have contributed to it *in a coherent way*. Taking the inverse Fourier transform of this power spectrum recreates an image of the object that is free of turbulence, meaning that it has a resolution that is set by the diffraction of the telescope.

In the simple case of a binary system, the expected power spectrum is very simple: it is a 2-dimensional sinusoid whose wavelength is λ/ρ, where ρ is the binary separation, and whose amplitude is proportional to the binary flux ratio (for equal-flux binaries, the minimum intensity of the power spectrum is zero). When such a morphology is found in the object's power spectrum, it is advisable to perform a fit in the Fourier plane to recover the binary parameter, rather than Fourier transform it back to the image plane, for noise propagation reasons. For more complex objects with no prior analytical model, the inverse Fourier transform is the only way to proceed. The speckle interferometry method is illustrated for the simple case of a binary star in Fig. 9 and for
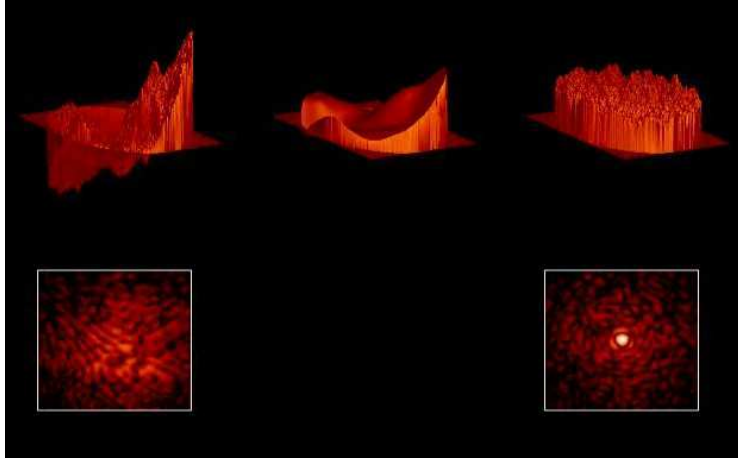
13

Fig. 12. *Top left:* Simulated incoming wavefront. *Top center:* Corresponding DM configuration. *Top right:* Simulated corrected wavefront. *Bottom:* The resulting images without and with AO correction (*left* and *right*, respectively). This simulation corresponds to a Strehl ratio of 40%. Credit: Jean-Luc Beuzit (Obs. Grenoble).

more complex objects in Fig. 10.

All methods based on the acquisition of very short integrations suffer from a common problem: sensitivity. Analyzing several thousand frames helps improve the final signal-to-noise but this is not necessarily enough. In a small fraction of a second, only bright objects provide a sufficient number of photons to record speckles with sufficient signal-to-noise to allow the Fourier transform computation. For instance, on an 8–10 m telescope, speckle observations are limited to objects whose infrared magnitude is about $K \lesssim 10$. Also, because atmospheric turbulence has a much stronger effect at shorter wavelengths, visible speckle imaging requires integration times as short as 10 ms, placing stringent constraint on the instrument design and required object brightness.

On the other hand, *speckle interferometry can allow to resolve objects below the diffraction limit under some circumstances.* Particularly, if one observes a tight binary system, the sinusoidal wave pattern in the power spectrum can be fit accurately as soon as its minimum is reached within the sampled spatial frequency range. In other word, if the binary separation is $\lambda/2D \leq \rho \leq \lambda/D$, one can determine the binary parameters. This is a factor of 2 below the diffraction limit. This phenomenon is also true for "normal" diffraction-limited images in principle: if the binary separation is larger than half the diffraction limit, one can tell that the source is not a point source and can therefore fit a double point source to the image. However, performing this fit in the image plane is much less accurate than in the Fourier plane. Indeed, in the Fourier plane, the fit is performed over many pixels and the key quantity to measure, the sinusoid wavelength, can be measured with great accuracy. This "super-resolution" effect only works if one knows the exact (analytical) model that needs to be fit to the image and is not always reachable. However, this
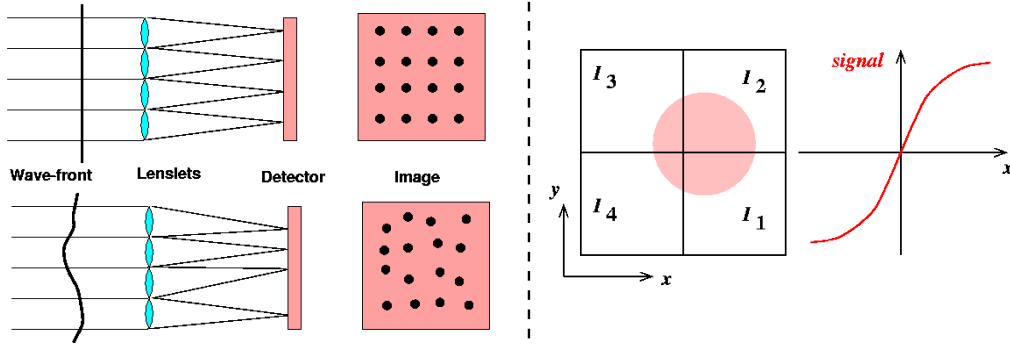
14

Fig. 13. *Left:* Illustration of the Shack-Hartmann WFS principle. *Right:* Measurement of the displacement of an individual spot with 4-pixel photometry. Credit: Andrei Tokovinin (CTIO).

is one of the key advantages of speckle interferometry over other high-angular resolution methods.

3.3 Adaptive Optics

What if a space telescope like HST is not sufficient (for instance because it is "only" a 2.4 m telescope) and one's favorite target is not observable with speckle interferometry (because it is a faint galaxy, say)? The only hope of obtaining high-resolution images of that target is to observe it with a telescope equipped with an adaptive optics (AO) system. The basic idea behind AO systems, first proposed by Babcock (1953) but only achieved by astronomers in the early 1990s (2 decades after the military developed such systems...), consists in *analyzing in real time the incoming corrugated wavefront and to correct it using a deformable mirror (DM) that can be adjusted rapidly enough to beat the fast-changing atmospheric turbulence.* If this can be done, then it becomes possible to integrate for really long exposures, several minutes or more, allowing for observations of very faint targets.

Fig. 11 illustrates the principles of AO and Fig. 12 shows simulated wavefronts before and after bouncing off the DM, whose shape is shown in the center. Before reaching the science camera, the incoming light from a "guide star" encounters a beamsplitter that divides the flux in 2 parts. In many systems, the beamsplitter is a dichroic, which lets the infrared light reach the camera and reflects the visible light towards the wavefront sensor (WFS). In some systems, however, the beamsplitters simply lets through a fraction of the incoming light and reflects the remainder with both beams working at the same wavelength. This is for instance the case for the infrared WFS in the NACO system installed at ESO's VLT.
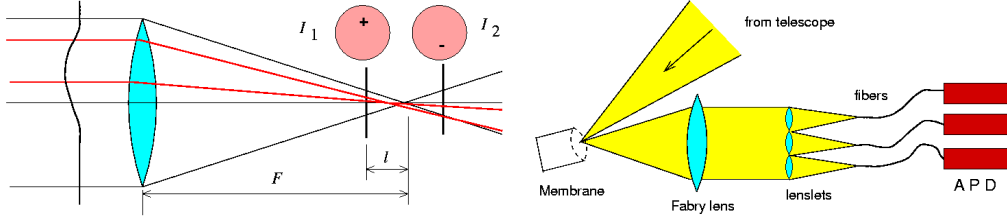
15

Fig. 14. *Left:* Effect of the wavefront curvature: it results in a slightly displaced focus from the nominal focal plane. *Right:* Illustration of the curvature WFS. Credit: Andrei Tokovinin (CTIO).

An AO systems consists of three key elements: a WFS, a DM, and a fast computer to link these two elements in real-time. The role of the WFS is to reconstruct an approximate map of the incoming wavefront, in which the continuous curves of the wavefront are simplified to a set of many independent inclined planar wavefronts (see Fig. 2). To probe atmospheric turbulence properly, the WFS should decompose the wavefront in as many as possible independent measurements. Generally, the wavefront is therefore split in many independent "apertures" whose equivalent size on the primary mirror should be as small as r_0. Similarly, the shape of the DM can be modified by a number of small actuators to recreate the shape of the wavefront. It is normally placed in a plane that is conjugate with the primary mirror and the equivalent actuator spacing should also be on the order of r_0 in order to reproduce the fine scale structure induced by atmospheric turbulence.

The most common design for the WFS is the so-called Shack-Hartmann sensor, in which the distorted wavefront is first split in many small parts by a lenslet array (see Fig. 13). If the incoming wavefront was perfectly planar, then we would obtain a 2D grid of regularly-spaces images of a point source. If, on the contrary, the wavefront is corrugated, then each little image is shifted from its nominal location, the amplitude and direction of the shift being dictated by the local slope of the wavefront at the entrance of the lenslet. By measuring the 2D offset for each spot, it is straightforward to reconstruct the shape of the entire wavefront. The offsets are generally measured through a linear combination of the fluxes of 4 adjacent pixels of the WFS (see right part of Fig. 13). Other WFS designs, such as the pyramid sensor, are quite similar in that they measure locally the slope of the wavefront through the displacement of the photocenter associated to a certain lenslet.

A completely different design for the WFS is the "curvature system" invented by Roddier (1988); it is illustrated in Fig. 14. To understand its principles, let us first consider a planar wavefront. The telescope will form an image that is in the focal plane, where the WFS is located. If we now displace the WFS by equal amount in front and behind the focal plane, we measure equal intensities per surface area since the beam is symmetric about the focal point. Let us now
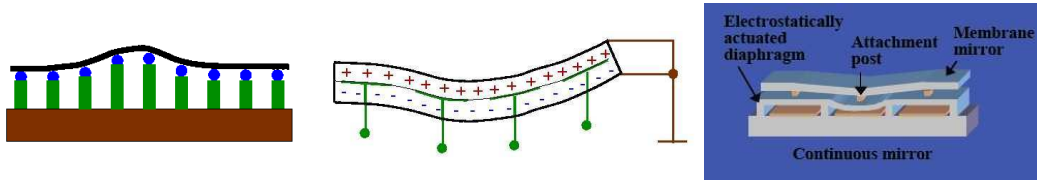
16

Fig. 15. *Left:* Schematics of a piezo-electric DM. *Center:* Schematics of a bimorph DM. *Right:* Schematics of a MEMS (micro-electro-mechanical systems) DM. Credit: Andrei Tokovinin (CTIO) and Claire E. Max (UCSC).

consider a curved wavefront. Because it is not planar, the telescope will form an image at a point that is in front or behind the nominal focal point. If we now repeat our back-and-forth motion of the WFS, the two measured intensities per surface area are not equal anymore. Measuring the difference in flux on the WFS as we regularly move it in front and behind the nominal focal point therefore allows us to measure the curvature of the wavefront. In a curvature WFS (which was first installed at the CFHT telescope), a lenslet array feeds a series of photodiodes and the system also contains a vibrating membrane to periodically change the location of the focal plane. From the measured curvatures, one can reconstruct the actual shape of the wavefront. This type of WFS gives access to the second derivative of the wavefront (rather than the first derivative for the Shack-Hartmann design), from which the wavefront can be reconstructed with many less measurements if the ratio D/r_0 is not too large and a few tens of lenslet are sufficient to provide excellent AO corrections in the near-infrared on a 3.6 m telescope. Guyon (2005) compared the relative efficiency of the various WFS designs in the context of planet detection on large ground-based telescopes.

The other key element of an AO system is the DM[2]. Here again, there are several possible technologies, which are illustrated in Fig. 15. The two most common designs are the piezo-electric and bimorph DMs. In the former, the mirror is a thin reflective layer supported by many mechanical actuators whose length can be adjusted by applying a certain voltage on it. As voltages change, the shape of the DM surface is modified to match the incoming wavefront. The bimorph mirror is a two-layer mirror in which at least one of the layer is piezo-electric. Again, applying appropriate voltages, it is possible to morph the coated, reflective layer into the required shape. More modern developments bring DMs into the realm of micro-mechanical systems (MEMS). These systems, developed originally for telecommunications purposes, are tiny flat mirrors that can be adjusted by 1–3 supports each. It is possible to create DMs made of thousands of such mirrors, as is planned for the Gemini "Planet Imager" next generation AO systems. The key advantage of MEMS device, besides offering the prospective of being very cheap since produced industri-

---

[2] This mirror is dedicated to the correction of high-order Zernike modes: the tip-tilt deformations (inclination of the wavefront) is dealt with by a dedicated mirror.
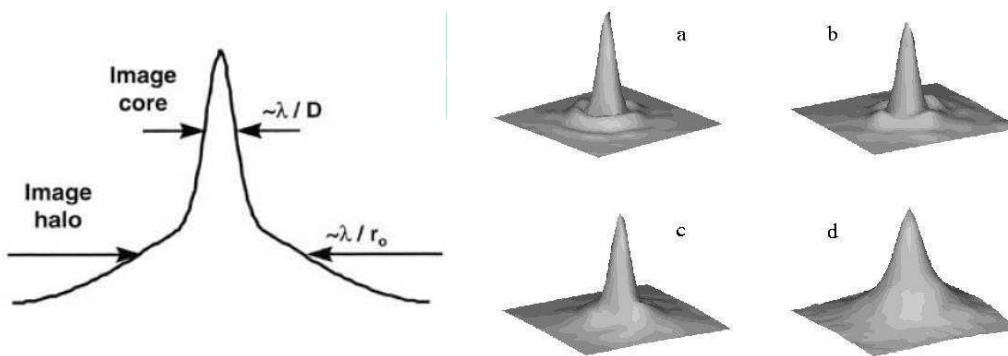
Fig. 16. *Left:* Spatial cut through an AO-corrected image of a point source. A diffraction-limited core is superimposed on an uncorrected seeing halo. Credit: Claire E. Max (UCSC). *Right:* AO images obtained with the CFHT Adaptive Optics Bonnette. The Strehl ratio in the images decreases from over 50% (top left) down to about 10% (lower right).

ally, is that they can be very small. DMs with thousands of actuators are expected to fit in just a few centimeters, which will permit to easily fit them into the optical path of any AO system. The first MEMS-based AO system is currently being brought to the sky at the Lick 1 m telescope (Gavel , 2007). Other, more innovative, designs exist for DMs but they are not discussed here.

To be effective, an AO system must fulfill several important requirements. First of all, the entire system must be fast: an entire cycle (probing the incoming wavefront, computing the required correction and applying it to the DM) must be shorter than the turbulence coherence time at the working wavelength of the WFS. *Since most systems use visible light for the wavefront sensing, this implies a working frequency of $\sim 500\,Hz$ or higher.* As a consequence, the WFS camera must be operated at a very high read-out rate while maintaining as low an electric noise as possible. This is why most current systems have their WFS working with visible light, taking advantage of the advent of extremely low-noise CCDs. The low-noise detector in the WFS is also a key feature: the incoming light is spread over many detector elements (see below), so if one wants to observe stars that are relatively faint, it is critical to have a very efficient optical system. A factor of 3 improvement in the detector noise allows to use stars that are 1 mag fainter as AO guide star.

Another important factor for AO systems is the number of actuators behind the DM. *On an 8–10 m telescope, the required number of actuators to match the r_0 sampling is a few thousand if the observing wavelength is in the visible.* The number of simultaneous measurements from the DM must be at least as large in order to allow for a determination of the command to be sent to the DM. The conversion of WFS measurements into DM commands is usually performed as a matrix multiplication after decomposing the wavefront shape
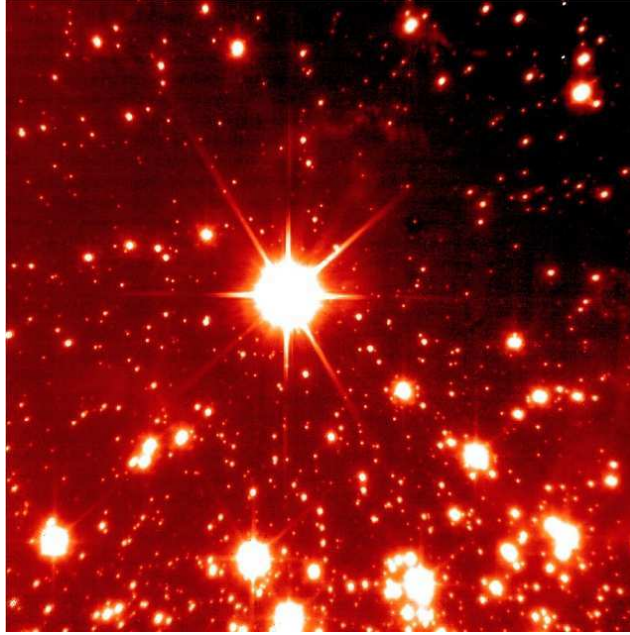
Fig. 17. Image of a dense stellar field obtained at $2.2\,\mu$m with NACO at VLT. The bright star in the center was used as the AO guide star. The stars located in the outskirts of this image, located about $20''$ away suffer from a poorer correction (lower Strehl ratio and core elongation) due to anisoplanetism. Credit: ESO.

into Zernike modes, which means that the number of basic operations to be computed by the real-time computer become rapidly overwhelming, even if "smart" algorithms are used. And these operations must be performed within a very short time, on order $1\,$ms or so. Therefore, it is mandatory to have a very powerful real-time computer to run an effective AO system.

Current AO systems have DMs that have a few hundred degrees of freedom, instead of the optimal few thousand; only one system in the world (the AEOS system) currently has more than 500 actuators. As a consequence, the correction provided by existing system is only partial, with only the lowest-order Zernike modes being sampled; nonetheless, these represent the major part of the turbulent power (see Fig. 6), so the correction can still be pretty good (Roddier et al., 1991). Partial correction is illustrated in Fig 12, as well as in Fig. 16. Similar to the results of shift-and-add speckle frames, AO images consist of a diffraction-limited core superposed on a much wider, uncorrected seeing halo. To measure the quality of the AO correction, astronomers frequently refer to the so-called Strehl ratio. This is a number that is in the 0–100% range, with 100% implying perfect correction, i.e., an Airy pattern image for point-like sources. The Strehl ratio is measured as follows: consider a certain telescope and observing wavelength. The perfect point spread function (PSF), i.e. the response of the telescope+instrument system to a point-like object, can be computed analytically or numerically (this is usually not exactly an Airy function, but something very similar). If one scales this perfect PSF to
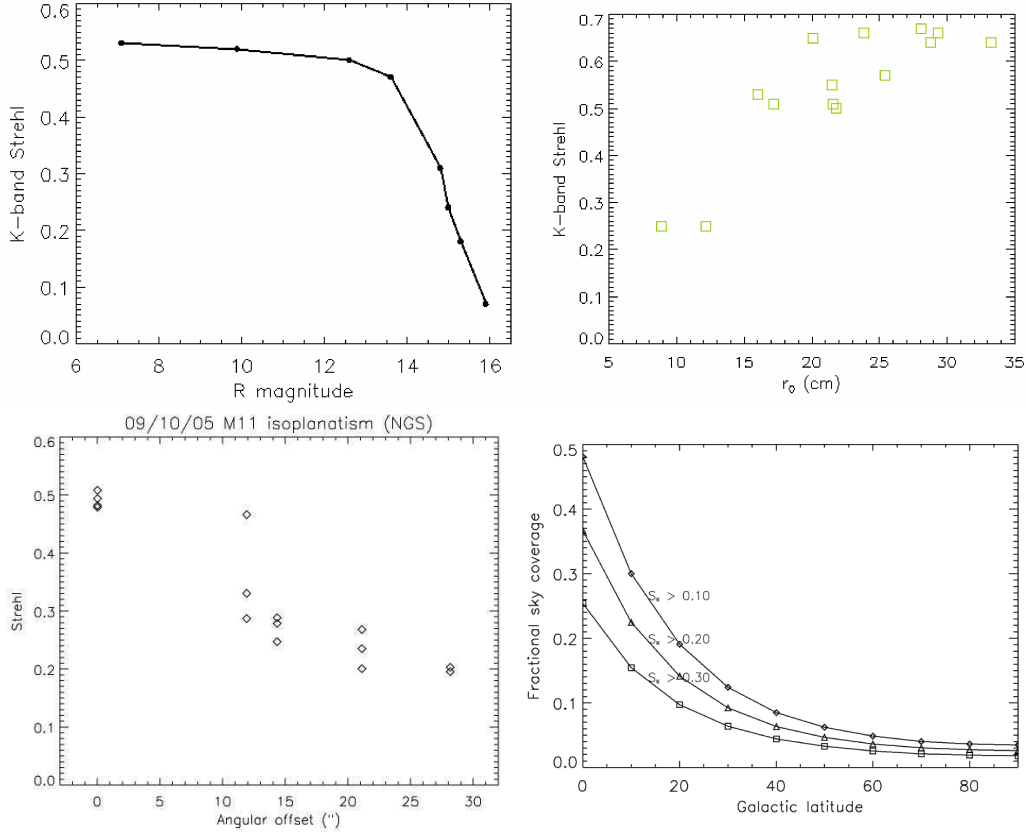
Fig. 18. Dependence of AO Strehl ratio on guide star magnitude (*top left*), Fried's parameter (*top right*) and distance to the guide star (*bottom left*). The latter two plots assume a bright guide star ($R \leq 8$). Also shown in the overall sky coverage as a function of galactic latitude for three fiducial Strehl ratios (*bottom right*). Credit: W. M. Keck Observatory.

have the same *total flux* as the observed image of a point source, then the Strehl ratio is the ratio of the *peak flux* of the observed point source to that of the perfect PSF. Current AO typically reach Strehl ratios in the 40–60% range at $2.2\,\mu$m for bright guide stars. Note that as soon as the Strehl ratios reaches about 15% the measured FWHM of images is almost equal to the theoretical diffraction limit of the telescope. Higher values of the Strehl ratio mean that the seeing halo is further reduced, providing higher contrast in the 2–$10\,\lambda/D$ angular separation range.

Although AO systems allow in principle the observation of faint targets, they nonetheless require substantial amounts of light for its guiding, especially for Shack-Hartmann sensors which split the wavefront with several hundred lenslets. The guide star must be bright enough that its photocenter can be accurately measured with a 10–20 cm aperture and an exposure time on the order of 1 ms. However, the guide star does not need to be the science target: it can be slightly offset from it. Basically, an AO system corrects for atmospheric
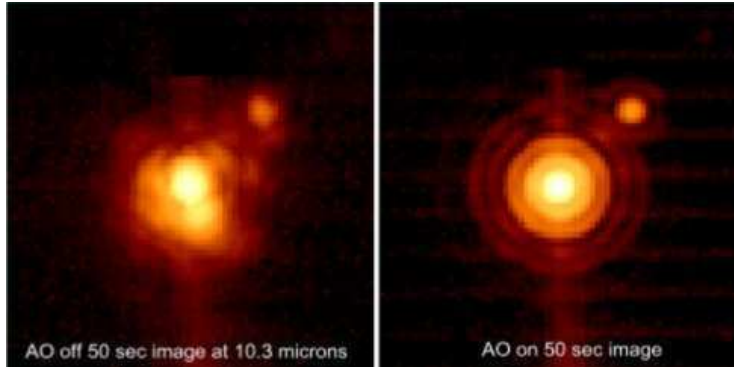
Fig. 19. Example of $10\,\mu$m images without (*left*) and with (*right*) AO correction on the 6.5 m MMT. The uncorrected image is almost diffraction-limited whereas the AO-enhanced images have a Strehl ratio of $\sim$98%. Credit: Laird Close (U. Arizona).

turbulence as measured in the direction of the guide star. Given the properties of turbulence, *the AO correction can only be valid over an area around that star whose angular size is roughly the isoplanatic angle introduced earlier.* This is just a few arcsecond in the visible (see Fig. 5), but it is on the order of 20–30″ in the near-infrared, another reason why AO systems primarily feed near-infrared cameras. So, even if a science target is too faint to be used as a guide star, there is a fair chance that a nearby star can be used for guiding. Of course, the closer the guide star, the better the correction, and it is sometimes preferable to use a fainter but closer guide star than a bright star located 30″ away from the science target. As can be seen in Fig. 17, the AO correction gradually worsens as the distance to the guide star increases.

The AO correction depends on many factors, particularly the guide star brightness, its distance to the science target, the turbulence properties (most importantly r_0 and τ_0) as well as telescope- and instrument-specific design features. The guide star brightness is probably the single most important factor: to reach a given accuracy on the photocenter measurements, a certain signal-to-noise, hence a certain flux, per WFS aperture is required. If the star is too faint, the system must be tuned to a slower update frequency, and the applied correction is not quite fast enough to track atmospheric turbulence. Alternatively, one has to settle for poorer accuracy measurement and correct for less Zernike modes. Either way, the resulting correction is necessarily poorer. Fig. 18 illustrates the capabilities of the AO system at the W. M. Keck Observatory; these are representative of currently available systems although there are some important system-to-system differences. The bottom right panel of that figure indicates the probability of finding an adequate guide star as a function of galactic latitude (and required Strehl ratio). Clearly, there are many less bright stars high above the galactic plane and it is increasingly difficult to find guide stars as one goes to higher and higher latitudes.

Other factors can limit the efficiency of a guide star. Most importantly, it is
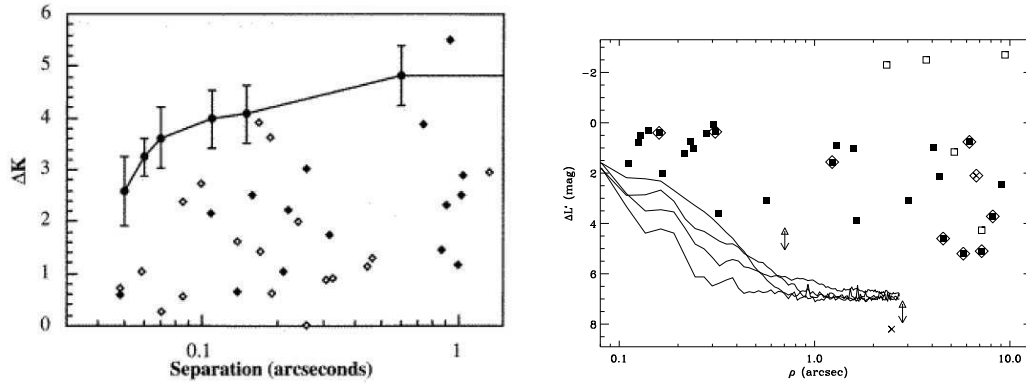
Fig. 20. *Left:* Detection limit for companions at $\lambda = 2\,\mu$m with speckle interferometry on the Palomar 5 m telescope (Patience et al., 1998). *Right:* Detection limit for companions at $\lambda = 3.8\,\mu$m with AO on the VLT 8.2 m telescope; the different curves correspond to different realizations of the Strehl ratio (Duchêne et al. , 2007). The two curves are directly comparable since the λ/D ratio is almost the same for the two instrumental set-ups and both projects conducted surveys of similar size.

assumed that the guide star is point-like. It is possible use an extended source as guide star, as long as it is possible to measure its photocenter in the WFS. With typical pixel sizes of 1–3″, the guide star should nonetheless be smaller than this. For observations of planets in our Solar System, such as Jupiter, Saturn or Netpune, astronomers usually use their larger satellites, whose size is on the order of 1″. A particularly difficult situation occurs when the guide star is an almost equal flux binary star whose separation is slightly larger than the seeing. In this case, the system is likely to oscillate from one photocenter to the other, resulting in an unwanted double-peaked PSF.

Finally, a key factor in AO imaging is the observing wavelength: because the atmospheric coherence time is much shorter at visible wavelengths, current AO systems provide only very partial corrections shortwards of $1\,\mu$m. A few AO systems have been demonstrated to provide substantial corrections in the visible, albeit without reaching the diffraction limit (Doel et al., 2000; Roberts & Neyman, 2002). Conversely, AO systems provide almost perfect correction (Strehl ratios well above 90%) in the mid-infrared, as shown in Fig. 19. At these wavelengths, even direct-imaging instruments are close to being diffraction-limited. However, an AO system working in this regime, as the MMT system (e.g. Close et al., 2003), provides a much more stable PSF and allows high contrast imaging, which can be key to one's scientific goals.

Assuming an appropriate guide star can be found near the science target, AO allows to obtain diffraction-limited images *and* to use long integration times. The latter point is critical for spectroscopy purposes, for instance: it is almost impossible to combine speckle imaging with spectroscopy (see Genzel et al., 1997, for a rare example of the combination of speckle interferometry and spectroscopy). By now, most major observatories have a facility AO system.
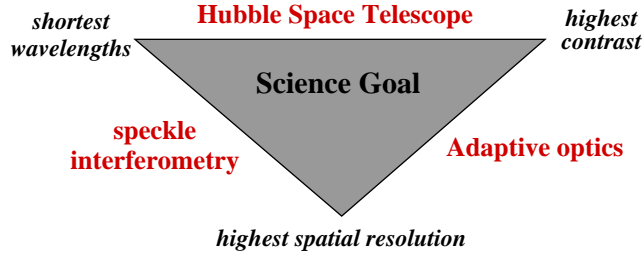
22

Fig. 21. Decision-making trade-off: to decide which high-angular resolution technique is most appropriate, one must consider what the science goals imply in term of technical capabilities. The final answer is not unique and other factors can play important roles.

Along with space telescopes, they are the high-angular resolution devices that have provided the vast majority of datasets in the field of extrasolar planets and protoplanetary disks, as we will see in Sections 4 and 5.

3.4   Performance comparison

So, what is the best high-angular resolution imaging technique: space telescopes, speckle interferometry or AO? Well, none of them. Or, rather, all of them! In the end, it really depends on what the science goals are. Let us consider a situation in which we want to use a high-angular technique at infrared wavelengths to search for very close companions to some objects. HST is probably not a good answer because of its limited diameter, hence resolution (only about $0\farcs18$ at $2.2\,\mu$m). If we are interest in companions in the $0\farcs05$-$0\farcs2$ separation range, then speckle interferometry is our best friend. Indeed, these systems provide almost uniform detection limits for companions at all separations down to $\lambda/2D$ or so (see left panel of Fig. 20). On the other hand, if our interest is in very faint companions at separations on the order of $0\farcs5$, then AO, which enables long exposures and higher contrasts, is our best choice (see right panel of Fig. 20). Also, AO and space telescopes are the only method allowing observing modes beyond regular imaging: spectroscopy, polarimetry, or coronography are almost impossible in speckle interferometry.

To decide which technique to use for a given project, Fig. 21 offers a decision chart that can be used as a general guideline. *The final answer is a trade-off between the observing wavelength, the required contrast and the need to reach the highest possible resolution.* As pointed out above, other factors can come into play and no absolute answer can be given. Nonetheless, this chart may be of interest to disentangle the various methods available.
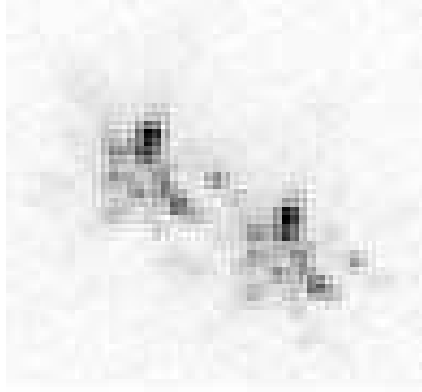
Fig. 22. A binary star (ζ Bootis) observed with speckle interferometry (Tubbs, 2003). Each component has its associated speckle cloud, both of which are exactly identical. If the separation of the binary is wide enough (this particular system is not quite wide enough since there is overlap between the speckle clouds), it is possible to use the power spectrum of one component as a reference to obtain the power spectrum of the other one. The simultaneousness results in a much cleaner final power spectrum and an enhanced ability to detect faint, tight companions.

3.5 Hardware and software improvements

Once a high-angular resolution image has been obtained, what can we do to increase our chances to detect a faint companion close to it? This is a key question for many projects. Fortunately it has some answers. There are indeed methods to improve on the techniques presented above. Some consists in using some dedicated post-processing methods (together with an appropriate observing plan), while others require a hardware upgrade to be implemented. Some of the most frequently used methods are discussed here, with a particular interest to improving the contrast and detectability of faint objects (e.g., planets or disks) close to a star. Notably, coronography is beyond the scope of this contribution; this powerful technique, which yields much improved contrasts around bright targets, can be used in combination with AO, for instance.

3.5.1 Improving on speckle interferometry

Speckle interferometry requires the observation of a "reference", a point-like star that is used to provide a "speckle map" (or, more exactly, its Fourier transform). Because the atmosphere varies quite rapidly, it changes somewhat between the acquisition of speckle frames on the science target and on the reference star. Therefore, computing the ratio of the two power spectra does not perfectly remove all effects of atmospheric turbulence. To improve on speckle interferometry, one can try and use "speckle holography". This technique relies on the presence, in the same field-of-view, of another bright star whose speckle cloud is recorded *at the same time* as that of the science target. Fig. 22

shows a speckle frame for a binary star: clearly, both components have exactly the same pattern of speckles. If the separation between the two objects is large enough, it is possible to completely isolate their speckle clouds and, therefore, to use the "other" star as the speckle reference. This method has been used, in particular, to study some components of binary stars at a very high spatial frequency (Koresko, 1998, 2002). Unfortunately, not many interesting objects in the sky have a bright star located 1–$3''$ away, so this remains a method that is not frequently used. It offers a quite substantial improvement over regular speckle imaging whenever it can be used, however.

3.5.2   Improving on HST and AO imaging

In AO images and, to a smaller extent, in HST images, the main problem when searching for faint "things" around a star lies in the imperfections of the PSF. Fig. 7 and 12 illustrate the many low-intensity defects that could easily be misunderstood as a faint point source in the vicinity of a star. To disentangle a real companion from PSF artifacts (particularly residual speckles), it is necessary to determine the PSF at the time of observation. This can then be used either to subtract it from the image of the science target or to use some deconvolution algorithm. While it is sometimes possible to use post-processing software to reconstruct the PSF[3], the ideal method consists in using a star that is located within the instrument's field-of-view and that is bright enough that the wings of the PSF have a good signal-to-noise. As in speckle holography, *the contemporaneity of the observations of the target and reference star ensures that they are associated to the same (imperfect) PSF*. Binary or multiple systems are great for this purpose. Recently Sheehy et al. (2006) developed a method to determine the PSF in a crowded field where the wings of each star PSF overlaps with its neighbors. Unfortunately, these particular configurations do not apply to all objects in the sky. The next best thing to do is to obtain an independent image of an object which is known/supposed to be a point source in the very same conditions as the science target.

As simple as this may sound, obtaining a good PSF for ground-based AO systems is quite a feat. Indeed, to match the AO correction quality, the PSF star should have the same brightness in the visible (usual domain of the AO WFS) and in the near-infrared (to match science camera settings). It should also be observed through the same section of the atmosphere. Ideally, the PSF star should be observed just before or after the science target from which it

---

[3] The TINYTIM software is available online for constructing synthetic HST PSFs (`http://www.stsci.edu/software/tinytim/tinytim.html`). For curvature AO systems, Véran et al. (1997) have proposed an algorithm to reconstruct the PSF from the WFS measurements.
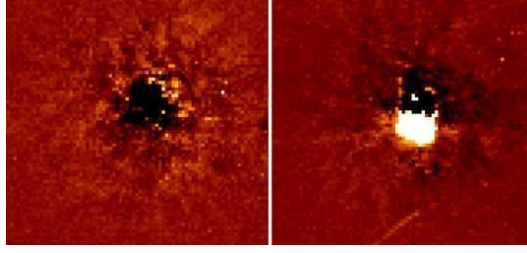
Fig. 23. Example of PSF subtraction on non-simultaneous images. Each image is the result of the subtraction of two Keck AO 2.2 μm images separated by about 10 second. Both images are shown on the same linear stretch and have a field of view of about 1″. Clearly, in the second case, there is a strong residual (whose peak reaches about 15% of the peak value in the individual images) below the primary star that could have been confused with a companion in the absence of the other subtracted image.

should be very close in the sky (to be observed at the same airmass). However, the key problem is that the turbulence just does not stand still and in the few minutes it takes to switch from one star to another, it is unlikely that the observing conditions will be exactly matched, resulting in imperfect PSF matching. Even if the PSF was quite good a match to the science target images, it probably would only apply to a limited subset of the science target observations: if one obtains a series of images extending over a few minutes, the change in PSF from the beginning to the end of the sequence is likely to be important. Fig. 23 shows the resulting of subtracting two images *of the same star* separated by only 10 s. The large residuals in the difference images, typically spread over a 1″-radius, is a reminder of the very fast and (sometimes) strong variability of AO PSFs. In fact, *it is quite frequent that an image obtained with the same instrumental set-up several months before provides a better PSF than one's own PSF observations just 5 minutes before/after the science target.* It is therefore not a bad idea to search the archive for potential PSFs, even if one has obtained dedicated images of point sources.

A more subtle effect related to AO PSF variability is that it is not the same everywhere and that the corresponding noise is not random. Of course, noise is stronger closer to the central star, but the key point is that it does not follow a simple Poisson statistics that has the nice advantage of averaging to zero. If Poisson statistics applied to the AO artifacts, one would just obtain many images and average them out. As shown in Fig. 24 (right panel, dashed lines), this simply does not work: *the final signal-to-noise ratio is almost independent of the number of frames that are combined.* This is in large part due to the fact that AO speckles are "pinned", i.e., their location is not random. Residual speckles are much more likely to be found at the location of Airy rings than in between them (Bloemhof et al., 2001). This effect, together with other "quasi-static speckles", which evolve on timescales of (tens of) minutes (see left panel of Fig. 24), explains why one can find speckles that are stable on timescales
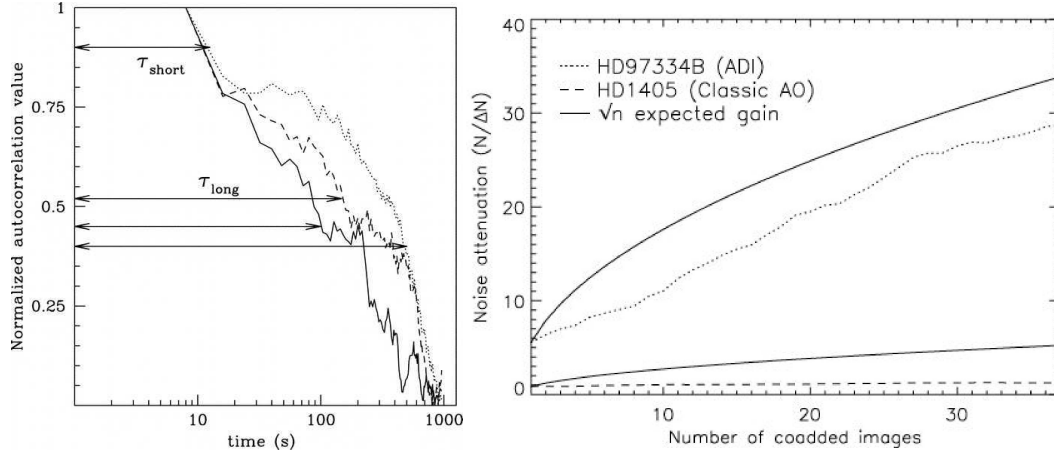
26

Fig. 24. *Left:* Temporal auto-correlation function for successive AO images from the AEOS 3.6 m telescope, measured at different distances from the central star (0.''1 [solid], 0.''19 [dashed], 0.''23 [dotted]). One can see, particularly at the larger distances, how the correlation remains important, even beyond 1 minute, due to quasistatic speckles. From Hinkley et al. (2007). *Right:* Increase in noise attenuation (or signal-to-noise ratio for point source detection) as a function of the number of coadded images. The dashed curve is using AO images without further processing, whereas the dotted curve uses the ADI technique. In both cases, the solid lines indicate the corresponding $\sqrt{N}$ extrapolation assuming all images are independent and have similar intensities. Note how the dashed curve does not show any improvement as N is increased. From Marois et al. (2006).

of minutes at these locations. Observing the same object for an hour may not be sufficient to make sure that what seems to be a real companion is not a "pinned speckle". *In the end, the best method to confirm the nature of a companion found in AO images, is to observe it with different filters or, even more convincing, with different telescopes.*

When dealing with AO images, quasi-static speckles are usually the limiting factor in searching for very faint point sources near a bright star (e.g. Sivara-makrishnan et al., 2002). There are several methods to attempt to remove these speckles that go beyond a basic PSF fit/subtraction. First of all, it is wise to filter out much of the PSF halo, which is generally axisymmetric and very smooth (i.e., whose power spectrum mostly lies at low spatial frequencies). *The most basic approach consists in removing an axisymmetric profile, based on the average radial intensity profile in the image.* More refined filters can be used to get somewhat improved results. The left panel of Fig. 25 shows an image after subtraction of an axisymmetric profile. Clearly, this method can only be employed for point source detection: extended structures, such as circumstellar disks, would be severely affected by such a filter.

When searching for faint companions, however, one has to recognize that all speckles have the same appearance as a faint point source: low-intensity Airy functions. *It is therefore both difficult to disentangle a real companion from a*
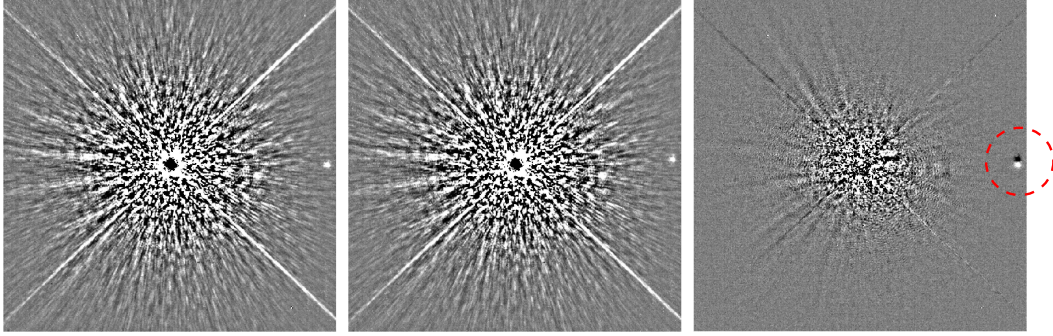
27

Fig. 25. Illustration of the ADI technique with 2.2 μm Keck AO images. *Left and center:* Two images that have been obtained after a rotation of a few degrees of the sky with respect to the detector and have been post-processed by subtracting an azimuthally averaged intensity profile. *Right:* The difference image, shown on the same grayscale, clearly reveals the faint companion located just off the right edge of the image as a positive-negative "double object" along the ortho-radial direction because of field rotation. Note also the much reduced speckle noise at closer separations. Credit: Christian Marois (LLNL).

fluke speckle and impossible to use numerical filters to remove speckles without removing true companions as well. Random speckles can be averaged out and taking a series of images of the same target readily identifies them. These are usually not a big problem, as long as many images are taken. The real difficulty lies in quasi-static speckle. Those look like point sources, do not move, and have roughly constant intensity over time (at least over a few minutes). To remove quasi-static speckles, differential techniques are required. There are three main such approaches: Angular Differential Imaging (ADI), Simultaneous Differential Imaging (SDI), and Polarimetric Differential Imaging (PDI). As discussed below, ADI and SDI are more appropriate when searching for point-like objects, such as faint companions, whereas PDI is the best technique when studying extended circumstellar material in scattered light.

The ADI technique is an AO counterpart to HST's "roll subtraction" technique. With HST, the PSF is very stable and all its defects remain constant over time, with a fixed orientation with respect to the telescope (since they are due to scattering off fixed telescope optics). The telescope being in space, it is possible to simply rotate the telescope and obtain a new image. Any companion (or disk) will rotate by the same amount, but the PSF defects remain at the same locations. Subtracting one image from the other results in an (almost) perfect cancellation of the central star whereas the companion can now be seen in two locations, as a pair of positive/negative images. This method has been used for many years to improve the sensitivity to faint companions/structures around bright stars (Schneider et al., 1998; Lowrance et al., 1999). Exporting this method to ground-based AO images is not trivial, as telescopes or instruments cannot always rotate. On some alt-azimuthal telescopes, it is possible to use a field rotator to maintain the PSF structure
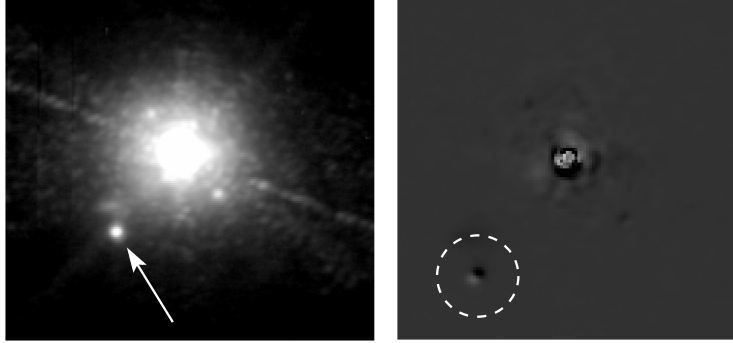
Fig. 26. Illustration of the SDI technique. *Left:* a raw VLT 1.6 μm AO image where a companion can be seen but may be confused with the many surrounding speckles. *Right:* Result of the SDI subtraction technique. The speckles have canceled out almost perfectly whereas the faint companion appears as a positive-negative "double object" along the radial direction because of the applied wavelength stretch. Credit: Guillaume Montagnier (Obs. Grenoble).

fixed while the star rotate during its nightly motion. If the star passes high enough in the sky, it is possible to register a large rotation in an hour or so around transit, permitting to use the ADI method. This method is not as powerful as roll subtraction with HST because of the intrinsic AO PSF time variability, but the improvement over regular AO imaging is very substantial (see Fig. 25).

SDI, the second main differential technique, was first proposed by Racine et al. (1999) to take advantage of the very nature of speckles. These are diffraction features off the telescope/instrument/AO bench. Therefore they obey diffraction laws. Most importantly, if a speckle appears at a certain location at a given wavelength, corresponding speckles would be seen with similar intensities in images obtained at nearby wavelengths, albeit at slightly different location. Indeed, *for all speckles in this "family", the distance to the star is directly proportional to* λ. The basic idea of SDI therefore consists in obtaining simultaneous images in several filters with adjacent bandpasses. The simultaneity ensures that the AO correction is identical for all images, so that the same (quasi-static and random) speckles are present in all of them, to the extent that non-common optical path errors are negligible. Data processing includes realigning all images and scaling them based on the central star flux. Most importantly, one must "stretch" radially the image obtained in the shorter wavelength filter by a λ_2/λ_1 ratio to realign the location of all speckles (if $\lambda_2 > \lambda_1$). Only real companions, whose position does *not* depend on wavelength, will not be aligned after this stretch is applied. Subtracting one image from another then provides a vastly improved reduction in PSF halo, enhancing the detectability of faint point sources, which again appear as a positive/negative pair similar to the ADI technique (see Fig. 26). The SDI technique works with any set of filters, as long as they are close in wavelength, but it can be further improved by using customized filters to match spectral
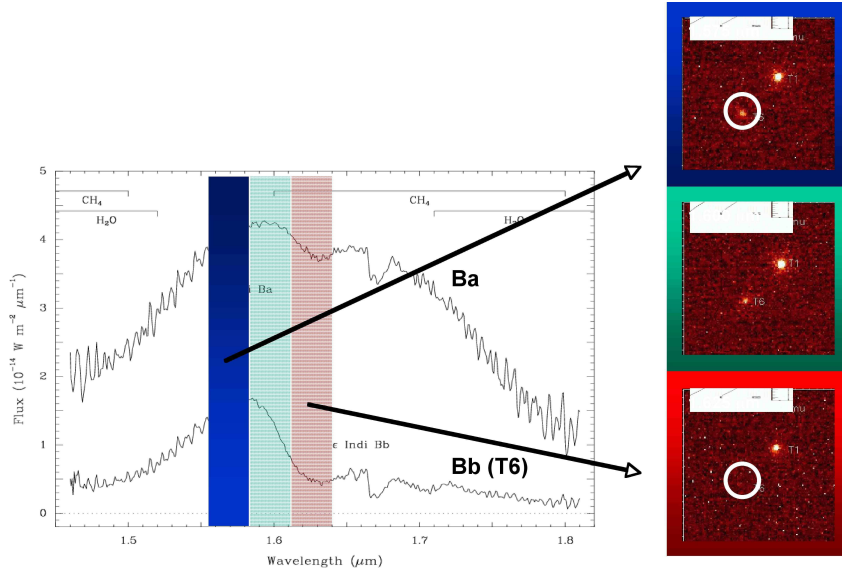
Fig. 27. Illustration of the SDI technique with $1.6\,\mu$m VLT AO images. Three images are obtained in adjacent narrow-band filters that have been selected to probe a deep methane absorption band, which is characteristic of very cool objects (low-mass brown dwarfs). In this example, the companion is cool enough that it is almost undetected in the longest wavelength image whereas it is clearly detected in the other two filters. Adapted from McCaughrean et al. (2004).

features of objects of interest. In Fig. 27, the three selected filters span the methane band which is seen in absorption in very cool objects. In this example, the "disappearance" of the companion in the longest wavelength filter shows that it is indeed a low-mass brown dwarf.

Both SDI and ADI were primarily devised to search for faint companions in the immediate vicinity of bright stars, and they may therefore not be adequate for all science purposes. ADI, or some alternate rotation-based subtraction, has also been successfully used to image circumstellar disks (Fitzgerald et al., 2007). However, in case of approximately axisymmetric circumstellar structure, ADI should not be used, as it would result in self-subtraction of the feature of interest. A more appropriate technique to study extended circumstellar material is PDI. As described in Kuhn et al. (2001) and illustrated in Fig 28, this technique consists in simultaneously obtaining two orthogonally polarized images of the target (with a Wollaston prism) and to subtract one from the other. The central star and its uncorrected speckles, all assumed to be unpolarized, will subtract off perfectly as a result of the simultaneousness of the acquisition. Any polarized signal around the star, such as scattered light off a disk or an envelope, will be highlighted, on the other hand. By rotating a waveplate in front of the instrument, it is possible to select various polarization angles. Using 0°/90° and 45°/135° pairing, one obtains Stokes Q and U images, respectively. To construct a linear polarization map of the circumstellar material, these images can be compared to a total intensity image, either con-
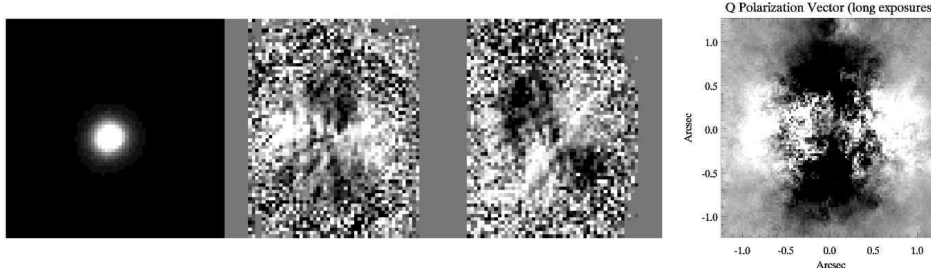
30

Fig. 28. Illustration of the PDI technique. *Left:* 1.6 μm tip-tilt-corrected UKIRT images of HD 169142, from Kuhn et al. (2001); from left to right, the Stokes parameters I, Q and U. Note how the disk is impossible to detect in the Stokes I image but is readily apparent in linear polarization maps. *Right:* 2.2 μm VLT AO Stokes Q image of TW Hya, from Apai et al. (2004). Note the same pattern of positive/negative lobes, the characteristic "butterfly" pattern induced by axisymmetric linearly polarized circumstellar material.

structed as the sum of the all 4 images to create the Stokes Q and U images, or obtained beforehand/afterwards by removing the Wollaston prism. In either case, one has to assume that the PSF has not changed between independent exposures, which is not necessarily a good assumption when using AO: only the Stokes Q and U images are free of time variability. If at all possible, it is better to compare models of the source with these latter images rather than with a linear polarization map.

To improve even more on these high-contrast techniques, it is possible to create even more refined techniques that combine SDI, ADI and PDI, or either of them with numerical filters, for instance. In the end, there is no "best technique": the advantages and drawbacks of each method must be weighted, keeping in mind both the feasibility of these techniques at a particular observatory and a specific science goal. The main thing to remember, however, is that *it is possible (and highly recommended) to do much better than simply accumulating many images and averaging them all together!*

3.5.3 Hardware improvements on AO

AO has now become a standard on large telescopes worldwide, and the number of results they have led to keeps rising, keeping pace with HST. There are some major limitations to AO, though, the most important one being the need for a relatively bright star in order to sample the wavefront with sufficient accuracy. For extragalactic projects, this is a major problem, as most bright stars are located in the Galactic disk, through which we generally cannot study distant galaxies. The limited sky coverage of AO is illustrated in the bottom right panel of Fig. 18.

To circumvent this problem, it has been proposed to create an artificial star
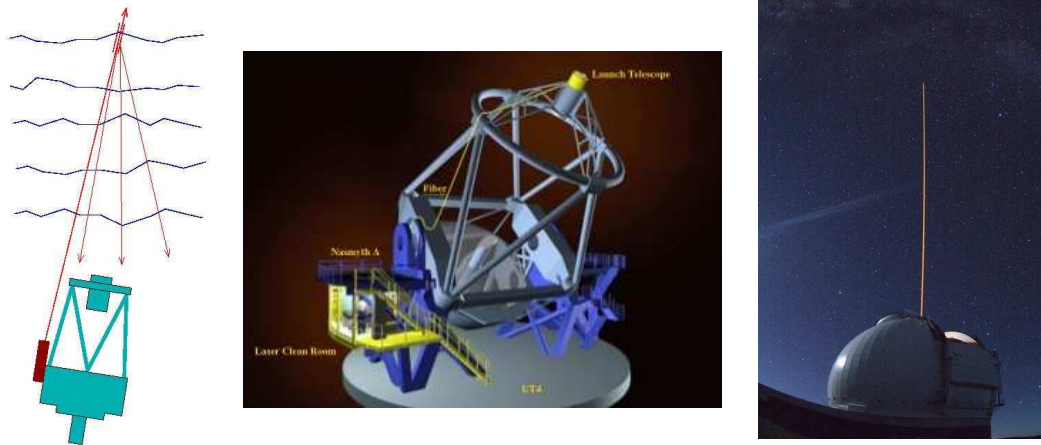
31

Fig. 29. *Left:* Basic principle of LGS: creating an artificial "star" high up in the atmosphere. *Center:* Implementation design of the LGS module at VLT. *Right:* The Rayleigh-scattering trace of the LGS at the Keck II telescope. Credit: Andrei Tokovinin (CTIO), ESO and W. M. Keck Observatory.

just next to the science target, using a powerful laser that would be launched in parallel to the telescope pointing axis. When launching a laser skybound, one can see from the ground its trace, due to Rayleigh scattering (see right panel of Fig. 29). This could be used as a reference star, but most of its intensity is coming from the lower layers of the atmosphere, so that only a fraction of the turbulence is sampled (and therefore can be corrected). Using a fined-tuned laser, it is also possible to excite sodium atoms into a desexcitation cascade producing a strong and quasi-monochromatic 589 nm visible light. This is the method used in most laser-fed AO systems. There is a layer of sodium in the Earth's upper atmosphere, and it is usually located about 90 km above sea level. The artificial "laser guide star" (LGS) is therefore located above pretty much all atmospheric turbulence and the wavefront we receive from it suffers the same disturbances as those of distant stars. This artificial star can therefore be fed into the WFS to close the AO loop. Interestingly, if one uses a near- or mid-infrared science camera, it is possible to exactly align the artificial star with the science target, since the LGS emits only at visible wavelengths. This way, one can be sure to sample the most relevant part of atmospheric turbulence. LGS modules are now in function or planned at most large telescopes.

The use of an artificial star appears like the ultimate answer to the main limitation of AO, sky coverage. However, this technique also has its own limitations. The first one is that the elevation of the sodium layer in our atmosphere is slowly changing, so that the WFS must be readjusted regularly to always be in focus. Furthermore, this layer has an intrinsic thickness of a few kilometers. The artificial star is therefore not really a "point" source, but rather a "tube" source. When observing at large zenithal angles, this can be a problem, as each
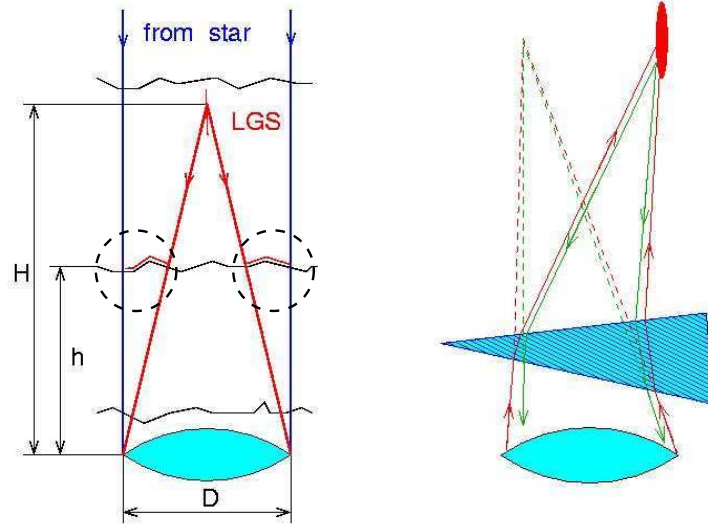
32

Fig. 30. Illustration of the two main limitations of LGS for AO. *Left:* Because of the finite elevation H of the guide star, only the central parts of the turbulence layer (at elevation h) are sampled whereas its outermost regions remain unprobed; this is called the "cone effect". *Right:* As light travels through the atmosphere to and from the LGS, the path of photons is identical in both directions. Any deflection induced by refraction in the atmosphere therefore cancels out and cannot be probed with the LGS. With this technique alone, tip-tilt cannot be measured. Credit: Andrei Tokovinin (CTIO).

spot on the WFS becomes elongated and, consequently, its centroid is harder to determine accurately. Beyond these two effects, there are two physical effects, which are illustrated in Fig. 30, that are intrinsic limitations to an LGS AO system. The first problem is called the "cone effect". The artificial star, as high as it is in the atmosphere, is not at infinity. The light cone defined by the artificial star and the telescope primary mirror only sample a fraction of the turbulence that is traversed by the wavefront from the science target. The outermost "rings" of these layers remain completely unprobed by the artificial star photons. Therefore, the AO correction is only partial and can never be perfect.

The second main problem stems from the fact that the upward laser propagates through the same turbulent layers as the artificial star photons that reach the telescope. Since light is diffracted in the same way in both directions, any overall wavefront inclination (tip-tilt) is lost and the WFS cannot know about it (see right panel of Fig. 30). In other words, the AO system will correct all but the two Zernike modes that carry the most power! The resulting image would have a high Strehl ratio but bounce all over the place as the phase screen changes in front of the target. To overcome this problem, it is necessary to observe a "normal" natural guide star *in addition to the artificial star* to sample, and correct for, the tip-tilt motion. Although it may seem that
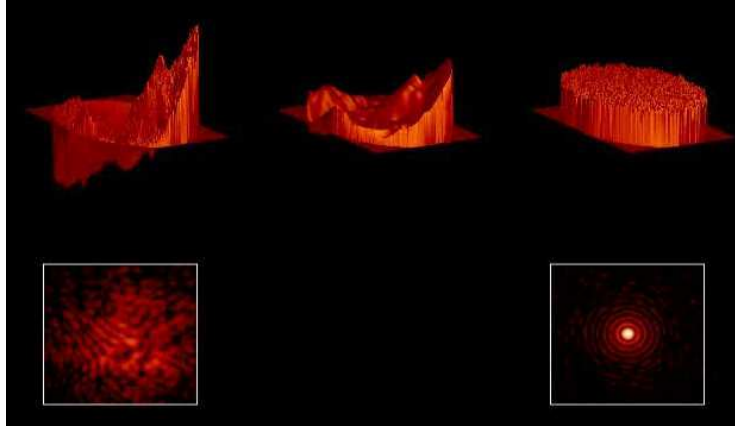
Fig. 31. Same as Figure 12 but for the future "extreme" AO system at VLT, SPHERE. The simulation corresponds to a Strehl ratio of about 96%. Note the much improved wavefront residuals and the associated reduction in speckle noise. Credit: Jean-Luc Beuzit (Obs. Grenoble).

the advantage of LGS is not that important after all, it must be realized that the tip-tilt modes can be characterized by only two numbers (the wavefront slope along two orthogonal axes), instead of the several hundred measurements needed to run a DM. It is therefore possible to use much fainter natural reference stars in Laser Guide Star mode than in Natural Guide Star mode. Typically, the limiting magnitude for operating the AO system is improved by ∼5 magnitude, resulting in an accessible sky coverage that is vastly improved.

Looking beyond existing systems, the next generation of AO systems, devoted to the search for planets around nearby stars (see Section 4), will provide major improvements. The currently developed systems (SPHERE at VLT and GPI at Gemini), which are categorized as "extreme AO systems", are based on the same general ideas as current systems. In short, they will be faster systems, running with DMs that have many more degrees of freedom, and they hope to reach Strehl ratios of 90–95% at $2\,\mu$m. Fig. 31 illustrates the incoming and corrected wavefronts, as well as the resulting image. It should be compared to Fig 12, which is a similar simulation for a currently existing system.

AO systems will also be key elements of the next generation of large ground-based telescopes. With diameters of 30–40 m, the effect of turbulence will be so important that it will not be possible to operate instruments without feeding them with AO-corrected beams. The use of an LGS will be necessary so as to point these telescope towards cosmologically interesting fields. Unfortunately, the large size of the telescopes will result in even more dramatic problems due to the cone effect described above. To overcome this issue, the AO systems of next generation telescopes will be "multi-conjugate AO" (MCAO) systems. Instead of using a single reference star, a single WFS and a single DM, an MCAO system uses several (laser and/or natural) guide stars, the same number of WFSs, as well as 2 or more DMs which are conjugated to different
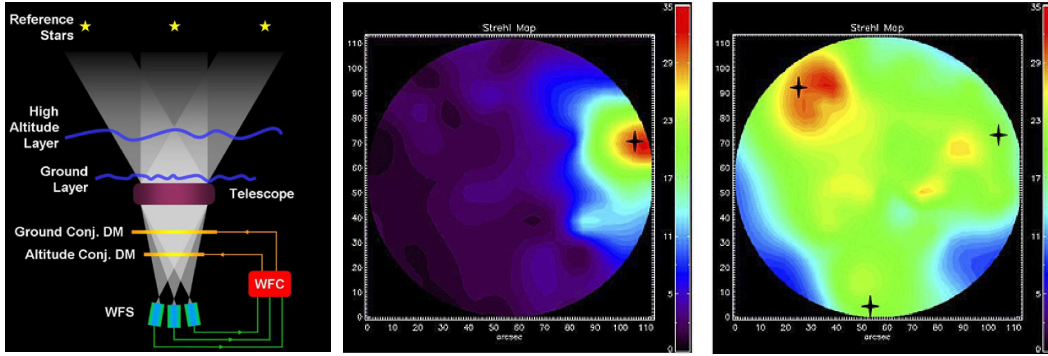
Fig. 32. *Left:* Basic concept of an MCAO system. More than 1 guide star are used in order to probe the turbulence of several layers in the atmosphere, resulting in a much wider sky area where the AO correction is good. *Center and Right:* On-sky performance of the first MCAO system (MAD at VLT): the plots show a map of the measured Strehl ratio with only one guide star (center, conventional AO system) and with three guide stars (right), with the same color scale. A good correction (Strehl ratio of about 20%) is obtained over most of the 2′-diameter field-of-view. Crosses indicate the location of the guide stars. Credit: ESO.

altitudes, to better correct layered atmospheric turbulence. By spreading the guide stars in the sky, it is possible to sample the turbulence over a wider region of the sky, alleviating the cone effect of LGS. The first MCAO system has recently been installed at VLT for demonstration of the method. On a single 8 m-telescope, the main advantage of MCAO is not so much to correct for the cone effect than to allow for good AO corrections over a much larger field-of-view. As illustrated in Fig. 32, instead of having a good correction close to the guide star but essentially no correction outside of the isoplanatic angle, *an MCAO system with 3 guide stars produces a roughly uniform correction over a region that is several times larger* (Tallon & Foy, 1990). The gain in corrected field-of-view offers the possibility to map large areas at high-angular resolution very efficiently with AO.

LGS, extreme AO systems, MCAO systems: these are the main avenues for improving on the current AO systems. Most current and all future large telescopes will be equipped with these, allowing for high contrast, high sky coverage, large field-of-view imaging at high-angular resolution. Even though space telescopes and, with a much more limited range of feasibility, speckle interferometry will likely continue to produce important results, it is a fair bet that AO will remain central in the future of high-angular resolution astronomy.

4   Imaging extrasolar planets: where we stand

Our current knowledge of the physical and chemical properties of planets is essentially limited to Solar System objects, and it is unclear whether the same
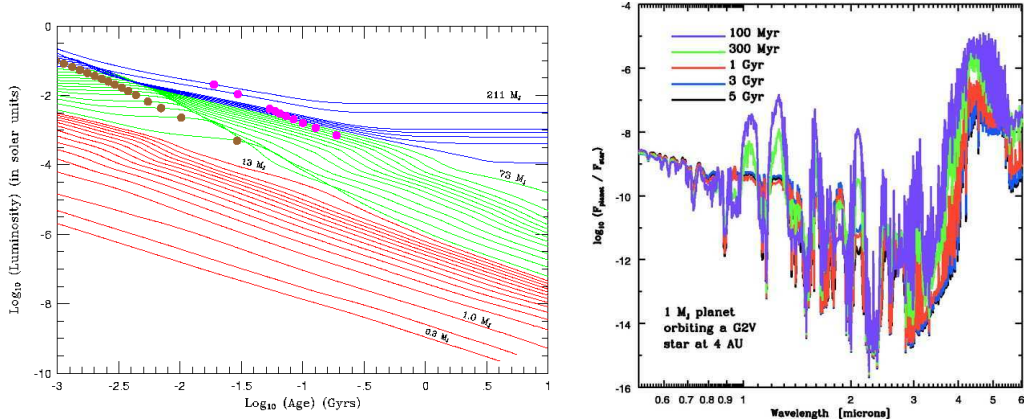
Fig. 33. *Left:* Temporal evolution of luminosity for stars (blue tracks), brown dwarfs (green tracks) and planets (red tracks). Very low-mass objects are much brighter and present smaller contrasts with respect to stars when they are young. Credit: Adam Burrows (U. of Arizona). *Right:* Predicted spectra (normalized to the stellar spectrum) for a $1\,M_{Jup}$ planet located at $4\,\mathrm{AU}$ from a parent G dwarf as a function of time (Burrows et al., 2004). Note the change in strength of some spectral features.

trends apply to other planetary systems. The search for extrasolar planets is one of the most rapidly advancing fields in astrophysics. More than 250 planets have been discovered around nearby Main Sequence stars in just about 12 years, most of them being identified through radial velocity measurements of the reflex motion of their parent star. We can infer their mass distribution, as well as the distribution of their orbital elements such as semi-major axis and eccentricity (see N. Santos' contribution). These properties, however, are more directly related to their formation process than to their physical properties, such as composition, density, ... Although highly successful, the radial velocity method is an indirect method, in that the presence of the planet is not inferred from the detection of its photons. It is therefore not possible with this approach to obtain a spectrum of the planet, a key observable to determine some of its basic physical properties, such as temperature and radius. The only case in which physical parameters can currently be determined corresponds to systems where the planet transits in front of its host star. The radius of the planet can be inferred and compared to planets in the Solar System. While these systems offer a first opportunity to understand the chemical and physical status of these objects, it can only be used for a small subset of extrasolar planets. Although the number of such systems is rapidly rising and statistically significant trends may soon be ascertained, further studies will require much larger samples to be analyzed, which, in turn, implies the need to obtain a spectrum of non-transiting planets.

For many years now, groups developing evolutionary models for stars and brown dwarfs have extended these into the mass regime of planets, below the
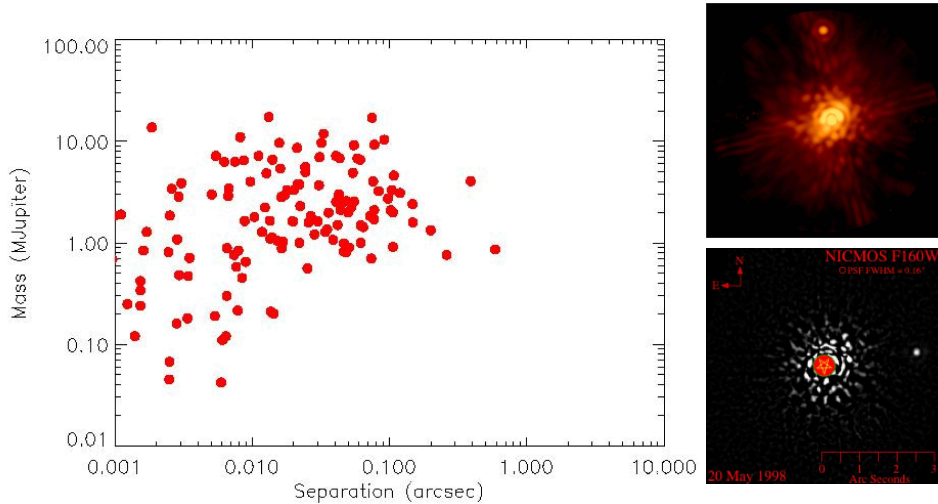
Fig. 34. *Left:* Plots of know extra-solar planets as a function of their mass and projected semi-major axis. *Right:* Examples of faint candidate companions detected at 1.6 μm with the AEOS AO system (*top*; separation $\sim 2''$, $10^{6.1}$ contrast ratio) and HST/NICMOS (*bottom*; separation $\sim 2''.5$, $10^{5.3}$ contrast ratio). Credit: Jean-Luc Beuzit (Obs. Grenoble), Ben Oppenheimer (AMNH, New York) and Glenn Schneider (U. Arizona).

deuterium mass limit ($M \approx 13 M_{Jup}$). The left panel of Fig. 33 shows one set of models, placed in a luminosity-time diagram. While stars eventually reach the equilibrium phase of the Main Sequence, brown dwarfs and planets keep cooling down as they age. Many physical properties, such as radius, density, surface temperature, ..., are predicted by these models, which however also depend on some unknown factors, such as cloud coverage or the size distribution of the cloud particles. Now that the existence of extrasolar planets is firmly establish, astronomers wish to measure such detailed physical properties in order to test the underlying physical assumptions, such as composition, equation of state, cloud properties. To reach this goal, one must first obtain *direct images* of exoplanets, and later on their spectra. Synthetic models already make testable predictions regarding the spectra of extrasolar planets (see right panel of Fig. 33).

What are the expected properties of planets? Based on our own Solar System and on known exoplanets, we expect planets to be located within 10–20 AU, and as close as 0.03 AU, from their parent star. Given a typical distance of 10 pc, this means that their projected separation will be a fraction of an arcsecond, at best, except maybe for distant ice giants such as Uranus and Neptune. The left panel in Fig. 34 shows the known exoplanets (as of Oct. 2005) in a mass-projected separation diagram. Direct imaging under a $1''$-seeing would not resolve *any* of these systems. *High-angular resolution techniques are therefore mandatory to image exoplanets.* Unfortunately, a close projected separation is not the only hurdle that has to be overcome. Planets are exceedingly
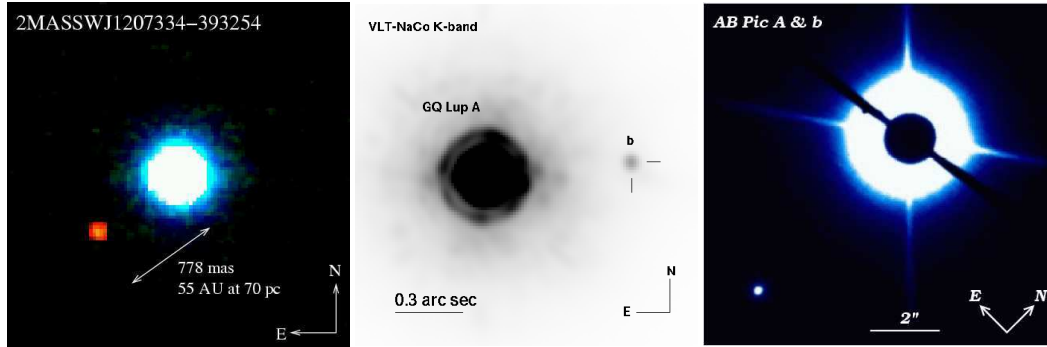
37

Fig. 35. AO Images of giant planet candidates around nearby young stars. From left to right: 2MASSWJ 1207334-39325 (Chauvin et al., 2004), GQ Lup (Neuhäuser et al., 2005) and AB Pic (Chauvin et al., 2005). Further analysis of the companion to GQ Lup resulted in its classification as a brown dwarf (see text).

faint objects. In the visible and near-infrared regime, exoplanets shine by reflecting their parent star's light. The corresponding flux ratio depends on the planet size, albedo and distance to the star. In our Solar System, the Sun-Earth flux ratio is larger than 10^9 and the Sun-Jupiter flux ratio is about a factor of 10 smaller. Contrast ratios are more favorable in the wavelength regime of the planet's own thermal emission. For the Earth, with $T \sim 300\,\mathrm{K}$, this occurs around $10\,\mu\mathrm{m}$; reaching the 10^{6-7} contrast is however almost as hard due to the strong thermal background for ground-based instruments. Focusing on the visible/near-infrared regime, the conclusion is that we must be able to reach contrasts of at least 10^9 within a fraction of an arcsecond of the parent star to be able to obtain the image of an exoplanet.

Currently, the highest contrast to very faint companion can only be obtained with deep HST or high quality ground-based AO images. The right panel of Fig. 34 shows two of the highest contrasts candidate companions detected to date. Current instrumentation does not provide the required combination of high contrast and very close projected separation needed to detect known exoplanets. More powerful instruments are being built (see Section 3.5), that will likely provide the first images of known planets within the next decade. In the meantime, we can use the high quality of current instruments to search for planets in a slightly different, more favorable context. Most importantly, as readily seen in Fig. 33, planets are much brighter when they are young since their luminosity is in large part fueled by gravitational contraction. Searching for planets around 1 Myr-old stars requires a minimum contrast that is about 1000 times smaller than searching for Main Sequence stellar systems. It is therefore not surprising that the first planetary-mass candidate companions (shown in Fig. 35) have been found around very young (2–30 Myr-old) low-mass stars and brown dwarfs. Interestingly, these objects, whose mass range from 8 to 25 M_{Jup}, are on wide orbits around their primary, with projected separations ranging from 55 to 260 AU. It seems improbable that these objects form through the same processes as the known population of extrasolar
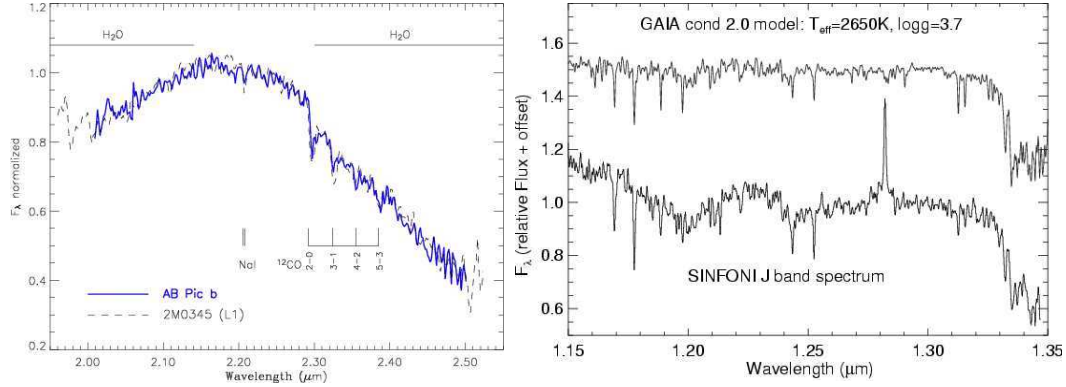
Fig. 36. Spectra of giant planet candidates around nearby young stars. *Left:* AB Pic b (Chauvin et al., 2005). *Right:* GQ Lup b (Seifahrt et al., 2007). The spectra are compared either to another very cool object of known effective temperature or to a model for such an object.

planets, whose semi-major axis never exceed 5 AU.

Assessing the basic physical properties of a low-mass companion, such as its mass, on the sole basis of near-infrared images is not an easy task, since it is necessary to compare observed fluxes to (uncertain) evolutionary models. Consequently, one can only propose a relatively wide range of possible masses for the companions. To better constrain the companions' properties, the acquisition of spectra is necessary. Once a (candidate) planet has been imaged, it is just a matter of integrating long enough to obtain a spectrum with sufficiently high signal-to-noise (see Fig. 36). A spectrum allows to determine the companion effective temperature, which allows for more precise determination of its mass. In the case of the companion to GQ Lup, the "possibly planetary-mass companion", as suggested as the time of discovery, is now undoubtedly classified as a brown dwarf companion, with $M \sim 25 M_{Jup}$ (McElwain et al., 2007; Marois et al., 2007), placing it in the same category as the Gl 229 and ϵ Indi systems (Nakajima et al., 1995; Scholz et al. , 2003). While it may not be as exciting as an extrasolar planet, this system, together with the others presented in Fig. 35, nonetheless demonstrate the power of AO to obtain images and subsequent spectra of very faint companions. With extreme AO systems, we can expect to obtain similar datasets at much higher contrasts, possibly reaching into the realm of known exoplanets, namely Jupiter-like planets on orbits with semi-major axes of 1–5 AU. This will finally allow direct comparison of our own Solar System with other planetary systems.

As a final note about direct detection of very low-mass companions, the example of GQ Lup b must be further discussed to illustrate the observational difficulties involved with imaging searches for exoplanets. As a bright, nearby, young low-mass star, this object has long been a favorite target for early faint companion surveys. Indeed, GQ Lup B was unknowingly "discovered" as early as 1994 with a first generation AO system (Come-On Plus, at ESO's
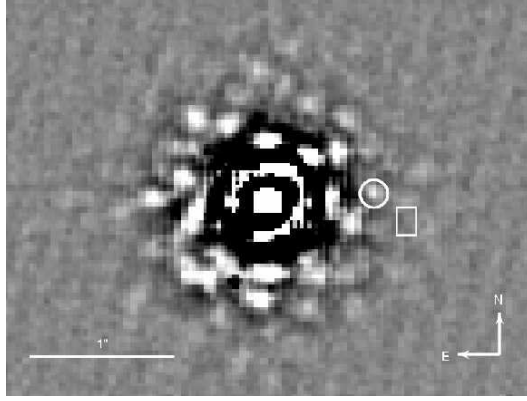
Fig. 37. Image of the GQ Lup very low-mass companion obtained with the Come-On Plus AO system on ESO's 3.6 m telescope in 1994. If the companion detected recently by Neuhäuser et al. (2005) were a background object, it would have been located in the white square. Rather it was then located in the white circle, where it should be if it were physically bound to GQ Lup. Unfortunately, in this early dataset, it was not possible to convincingly prove that this was not another uncorrected speckle due to the Earth's atmosphere. From Janson et al. (2006).

3.6 m-telescope). Unfortunately, the companion was lost in the uncompensated speckle cloud, as can be seen in Fig. 37. It is only after the "discovery" of the companion with NACO on VLT that this early dataset could be used to confirm the bound companion through common proper motion. This example illustrates the difficulty of detecting faint point sources within the speckle cloud that "normal" AO systems do not correct. With modern systems, and even with future generation systems, this occurs at fainter levels than in these early images, but the problem is nonetheless present and must be kept in mind. As discussed in Section. 3.5, strategies to circumvent this problem exist and it is important to use them to be able to confirm companions right at the detection limit.

Overall, our success in imaging extrasolar planets is very limited, yet very promising, so far. Only a couple of such objects have been found, but they have permitted to refine both the observing strategies to detect them and the type of follow-up observations required to determine their physical properties. Within a decade or so, we expect future instruments to allow detection of many new extra-solar planets, around both young and old stars. These instruments will also take advantage of elaborate coronographic techniques to further enhance the achievable contrast.

5 Imaging disks: structure and dust content

Over the last decade, the search for extrasolar planets has been paralleled by more and more detailed observations of protoplanetary disks. These are the birth site of planets and we will likely learn as much about planet formation by studying their properties as we will by determining the statistical properties of mature planetary systems. Circumstellar disks come in different flavors and the means to study them are many. We will briefly discuss these two aspects before focusing on a subset of recent results regarding protoplanetary disks.

5.1 Scientific goals and methods

Circumstellar disks are ubiquitous to star formation. As the prestellar core collapses to form the central star, angular momentum conservation results in the formation of a flattened structure that eventually settles into a slightly sub-Keplerian rotating disk. The disk is mostly made of gas, although 1% or so of it mass is made of small dust grains. Gas viscosity induces accretion on the central star while angular momentum is transported outwards. In this environment, dust grains experience several competing mechanisms (see Dominik et al., 2007, for a recent review). First of all, low-velocity encounters between grains tend to form larger-size aggregates. When the relative velocities are too large, however, large particles are shattered, replenishing the large reservoir of small grains (Weidenschilling, 1977, 1980, 1984). At the same time, drag forces induced by the gas tend to modify the spatial distribution of dust particles. Gravity tends to force large grains into the disk midplane whereas pressure gradients lead them to drift inwards (e.g. Barrière-Fouchet et al., 2005). As time goes by, the pristine disk made of sub-micron (interstellar-like) dust grains is transformed into a disk where mm-sized and larger particles are gathered into a thin midplane layer. There the density can become large enough to trigger the formation of km-sized bodies whose gravity is intense enough to sweep all dusty material in their surrounding, on their way to form planetesimals. At intermediate stages, one expect a stratified structure for disks, with increasingly larger grains as one peers deeper into the disk. Considering the large uncertainties associated with numerical models predicting the evolution of dust grains in disks, *a major challenge for modern astronomy consists in determining empirically the properties of disks.* G. Lodato's and R. Alexander's contributions provide more details on the theoretical/numerical studies concerning disk evolution.

Disks surrounding young stars are generally classified in two categories. Young disks around T Tauri and Herbig Ae/Be stars (low- and intermediate-mass Pre-Main Sequence stars, respectively) represent an early phase, in which most of
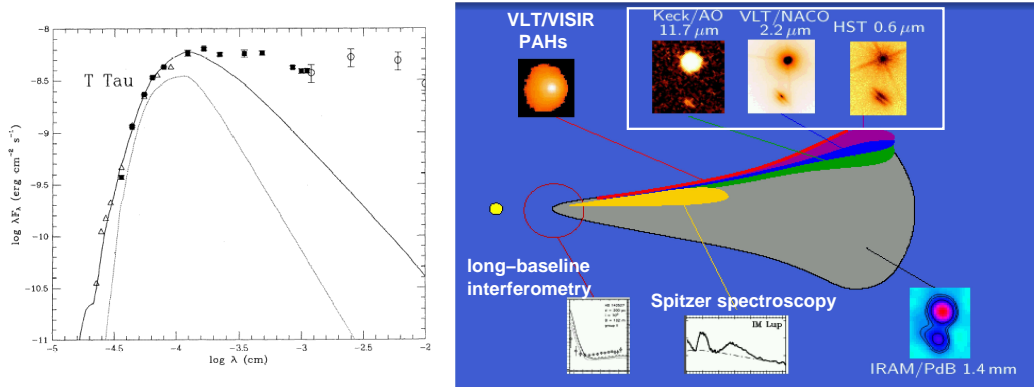
41

Fig. 38. *Left:* The "old" method to study circumstellar disk: gather a complete SED and fit a model to it (Bertout et al., 1988). *Right:* The "modern" way to study circumstellar disk: use many different techniques, which all probe different dust populations. Credit: Christophe Pinte (U. of Exeter).

the mass is still in the gaseous phase and the dust population still contains copious amounts of submicrons grains (Kenyon & Hartmann, 1987; Bertout et al., 1988). On the other hand of the spectrum, debris disks are gas-free, and their "dust" content is dominated by large (km-sized and larger) bodies. Frequent collision between large bodies produce a population of small grains that are rapidly swept away due to radiative and wind pressures or accreted onto the central star (Backman & Paresce, 1993; Meyer et al., 2007). While gaseous disks are optically thick to the parent star's and their own radiation, debris disks are optically thin. Based on strong asymmetries observed in debris disks, it is believed that these systems already host planetary systems. Since primordial disks survive for at most 10 Myr, the planet formation process must proceed in just a few million years. In this contribution, we only discuss young gaseous disks in the interest of space. Moreover, we focus on the study of the dust component of young disks: even though it only represents a minor component in terms of total mass, it plays a vital role in the formation of planets.

There are many observational approaches to study young disks. The first methods to be employed were spatially unresolved measurements. With a typical size of 100–300 AU and a distance of 140 pc or more, the apparent size of disks, is on the order of $1''$, and it is safe to consider them as unresolved in seeing-limited datasets. The spectral energy distribution (SED) of a disk (left panel of Fig. 38) offered the first robust evidence that young stars were indeed surrounded by circumstellar disks: the strong near- to far-infrared excesses found around a majority of young stars is understood as the thermal emission of a continuous, flared, passive disk [4]. Beyond the SED, the next level

---

[4] A passive disk is heated by the central starlight only; accretion luminosity viscously deposited within the disk is generally negligible in terms of heating except for the very high density inner disk midplane, which is so embedded that its emission
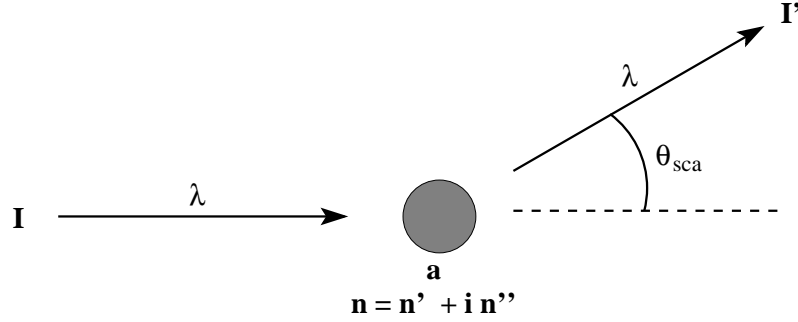
Fig. 39. Schematics of the scattering of a photon off a spherical dust grain. There is a certain probability that the photon is absorbed rather than scattered, so that, on average, $I' < I$, with $I'/I = albedo$. In a scattering event, the photon wavelength λ remains unchanged, whereas its direction of propagation is modified by an angle θ_{sca}. The probability distribution for this angle is the so-called "phase function". The key parameters that determine *albedo* and the phase function are: the grain size (a), the photon wavelength λ and the grain refractive index (n).

in observational refinement is spectroscopy. In the regime of the dust thermal emission, it can provide some insight on the disk properties. The shape of the $10\,\mu$m silicate emission feature, believed to come from the optically thin surface layer of disks, can be used to infer dust grain size: small ($\sim 0.1\,\mu$m) grains produce a strong, narrow peak around $10\,\mu$m whereas large ($\gtrsim 2\,\mu$m) grains produce such a broad feature that it is hard to detect over the continuum flux (Bouwman et al., 2001; Kessler-Silacci et al., 2006). In the near-infrared, spectroscopy has been used to demonstrate that the innermost regions of disks consist of a directly heated vertical wall whose temperature reaches the dust sublimation temperature (Muzerolle et al., 2003). While it is possible to reproduce the SED and infrared spectra of a disk with radiative transfer models, it is difficult to constrain most of the relevant physical parameters, however.

To go further, resolved images must be obtained. This is now possible over a very broad range of wavelengths. In the (sub)-millimeter regime, long baseline interferometers, such as IRAM's Plateau de Bure Interferometer or (in a few years) ALMA, produce maps of the dust thermal emission. Because disks are optically thin in this long wavelength regime, such observations are sensitive to the bulk of the disk mass. It is therefore possible to infer the disk total dust mass, as well as to constrain their dust size distribution. This is discussed in more detail in L. Testi's contribution. Near- and mid-infrared long baseline interferometry, as they provide an exquisite spatial resolution of a few milliarcsecond, now provide a direct access to the innermost regions of circumstellar disks. Results from this technique are discussed in R. Akeson's contribution.

---

is entirely reprocessed by higher layers before reaching the observer. An exception to this configuration is the FU Ori-like phenomenon, in which accretion luminosity surpasses by up to an order of magnitude or more the luminosity of the central star.
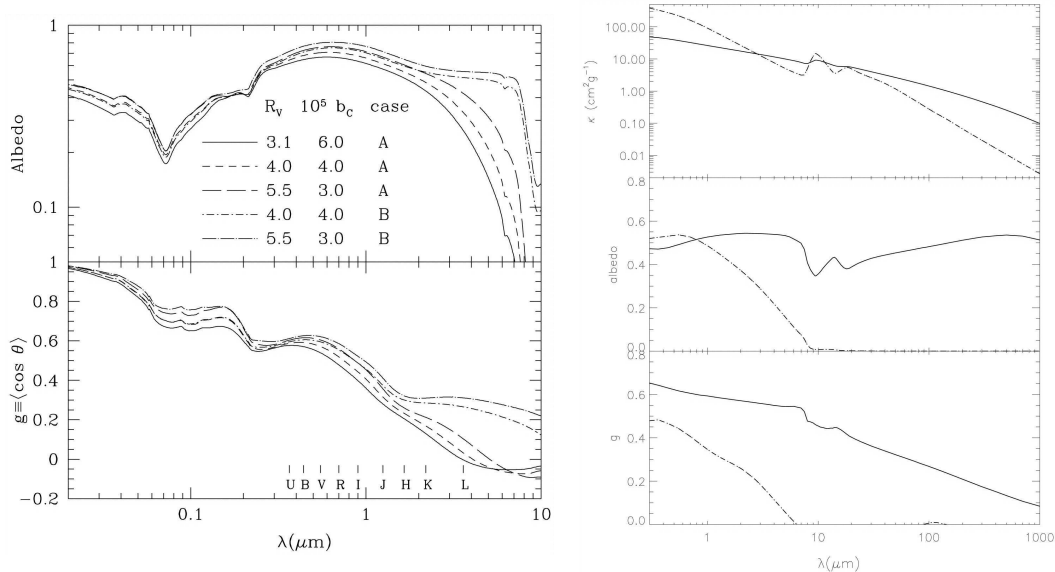
Fig. 40. *Left:* Albedo and scattering asymmetry parameter as a function of wavelength for several dust grain size distributions (Weingartner & Draine, 2001). Case B models contain larger grains, resulting in a higher albedo and more asymmetric phase function in the mid-infrared. *Right:* Total dust opacity, albedo and scattering asymmetry parameter for two dust mixtures: interstellar dust (dashed), and protoplanetary disk dust (solid). The latter contains much larger grains than the interstellar medium (Wood et al., 2002).

The present contribution focuses on scattered light images of young disks. When incoming visible and near-infrared photons interact with a dust grain, they can be either absorbed or scattered. The probability of being scattered rather than absorbed, which depends on the grain properties (composition, size, shape) and the photon wavelength, is called the albedo. If the photon is scattered, it is deviated by a random angle whose probability distribution is called the scattering phase function, which is frequently characterized by the so-called "asymmetry parameter" $g = <I(\theta)\cos\theta> / <I(\theta)>$ (with $0 \le g \le 1$). In short, large grains tend to scatter preferentially forward (small-angle deviations, $g \to 1$) whereas small grains scatter isotropically ($g \to 0$). Fig. 39 summarizes a scattering event. For spherical homogeneous grains, scattering can be considered analytically: the Mie theory predicts the scattering phase function as well as the grain albedo. The composition of the grain is described by its (complex) refractive index, n. For non-homogeneous or non-spherical grains, analytical treatments do not always exist, but simplifying assumptions (random orientation of irregularly shaped particles, "effective medium" theory) can be made. Because young disks are optically thick, the stellar photons interact with dust grains high up in the disk, close to its "surface". Some photons are scattered even more deeply into the disk, helping heat the inner regions of the disk, while other scatter upwards, away from the disk and to the observer. *The scattered photons we received therefore have only interacted with the uppermost layers of the disk.*
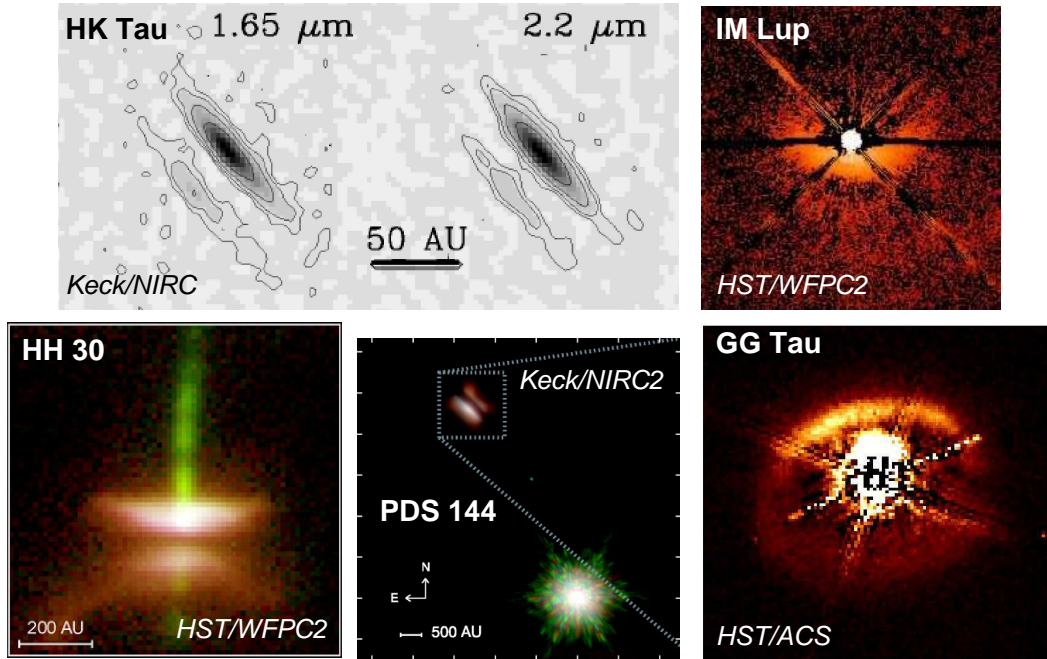
44

Fig. 41. A family portrait of young gaseous disks observed in scattered light with high-angular resolution imaging techniques. References: HK Tau (speckle holography; Koresko, 1998), IM Lup (HST; credit: Christophe Pinte, U. of Exeter), HH 30 (HST; Burrows et al., 1996), PDS 144 (AO; Perrin et al., 2006) and GG Tau (HST; Krist et al., 2005). The total field-of-view of the each image range from $1''$ to $6''$.

Using scattered light images, in combination with other techniques, it is possible to constrain the wavelength dependence of the dust properties: opacity, albedo and scattering asymmetry parameter. Fig. 40 shows some predicted behaviors for several dust populations. It is then possible empirically constrain the basic properties of the dust population (composition, size distribution) and, in turn, to better constrain the disk properties. Indeed, *because young disks are optically thick, their appearance in scattered light images is a complex combination of the dust properties with the disk geometrical structure. Only once one of the two factors is clearly determined, it is possible to study the other.* A key factor in this approach, as demonstrated in Section 5.2, is the combination of datasets covering a wide range of wavelengths. Only wavelength-dependent factors come into play in explaining the differences in disk images, so it is possible to disentangle geometrical effects from dust properties. Moreover, dust opacity always decreases as the wavelength increases (see right panel of Fig. 40). An immediate implication is that visible, near-infrared and mid-infrared scattered photons have interacted with the disk at different elevations above the midplane: longer wavelength photons need to travel deeper inside the disk before interacting with a dust grain. Using longer wavelength imaging therefore allows to probe deeper into the disk, a feature that can be used to test stratified disk structures as discussed in Section 5.2, for instance.
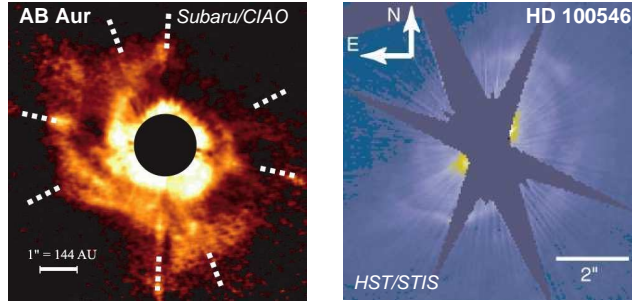
Fig. 42. Example of asymmetric disks seen in scattered light. Both images were obtained using coronographic technique in addition to high-angular resolution methods. Images from Fukagawa et al. (2004, AO) and Grady et al. (2001, HST).

5.2   Results from high-angular resolution imaging

The number of spatially resolved images of young protoplanetary disks is rapidly increasingly. Many of them have been resolved in scattered light using high-angular resolution techniques. The contrast between a young star and its surrounding disk in the visible or near-infrared is quite large, and coronography is necessary in many cases to detect the disks. There are however some systems in which the disk can be detected without this additional technique. Fig. 41 presents a limited sample of such disks. HH 30, HK Tau B and PDS 144 B share a common particular geometry: these disks are seen close to edge-on, within 5° of perfect alignment. The immediate interpretation of such images is that these disks are optically thick and have a flared geometry (consistent with hydrostatic equilibrium). Therefore, the disk blocks our line of sight to the central star and we only detect photons that come from the central star and have scattered off the disk surface at the outer edge. In this situation, there is no contrast problem. Furthermore, in binary systems like HK Tau and PDS 144, it is possible to use the bright nearby companion as a good AO guide star or speckle holography reference star. The GG Tau system is particular in that it is a binary system surrounded by a large dusty ring. With this geometry, the ring scattered light is constrained to a region located 1–2″ away from the stars, and it is then possible to perform a clean PSF subtraction and reveal it. The IM Lup disk is a rare example of a scattered light image of a young disk around a single star that is not observed in an edge-on configuration. The stability of the HST PSF is a key asset in this case to overcome the strong star-to-disk contrast of this more general configuration.

From a single scattered light image, several basic factors can be inferred. First of all, the disk radius can be estimated quite readily (in some cases, the disk may extend further than the scattered light nebulosity, though). In binary/multiple systems, the disk (inner and) outer radii can be compared to the system's projected separations to determine whether the disks have been truncated by the orbital motion of the components (Stapelfeldt et al.,
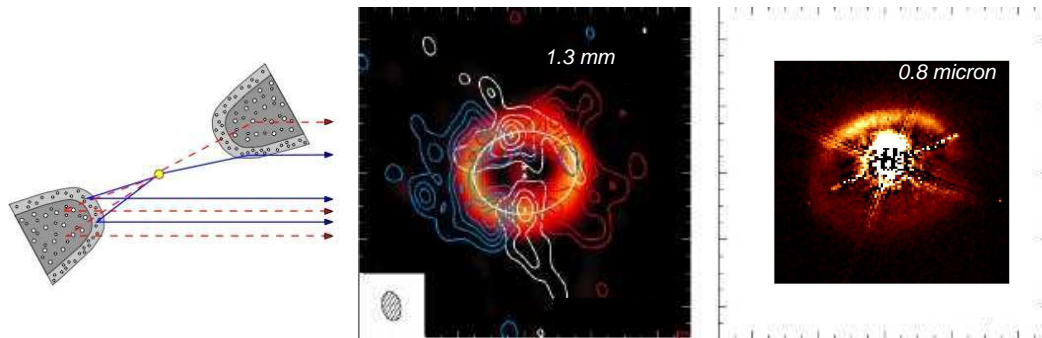
46

Fig. 43. *Left:* Sketch of the geometry of the circumbinary ring of GG Tau; the observer is located on to the right. The top portion, which is closest, is seen in scattered light with a smaller scattering angle than the bottom (back) part of the ring. *Right:* Interferometric millimeter emission map of the ring, showing its azimuthal symmetry, aligned with a visible (HST/ACS 0.7 μm) scattered light image of the ring, with the top/front part much brighter than the bottom/back part due to scattering anisotropy. Adapted from Guilloteau et al. (1999) and Krist et al. (2005).

1998; Monin & Bouvier, 2000). When sufficient knowledge is available, it is even possible to revert the reasoning and use the observed radii of disks to place some constraints on the orbital motion of systems whose orbital period is far too long to be measurable (Beust & Dutrey, 2006). The disk inclination can usually be determined to within a few degrees with a single image. This is straightforward in edge-on disks ($i \gtrsim 85°$). For other systems, the inclination can be inferred assuming that the disk is axisymmetric via the aspect ratio of elliptical isophotal contours. Note, however, that scattering is angle-dependent, so that isophotal ellipses do not necessarily geometrically correspond to a circle centered on the star and running along the disk surface. Still, this can provide a first estimate of the disk inclination. Finally, scattered light images can reveal strong asymmetries, as can be seen in Fig. 42. Because young disks are optically thick, one must be cautious about the interpretation of such asymmetries. It may be tempting to conclude that they trace spiral density waves in the disk, possibly generated by a forming protoplanet, for instance. However, since the disk is optically thick, increasing locally its density does not result in an enhanced scattered light flux. Rather, such structures can be interpreted as evidence for locally higher scale heights in the disk. A warped disk structure can also help in producing asymmetric images. Both explanations share the basic idea of modifying the disk *geometry*, rather than its density structure, to produce local enhancements of scattered light. It is nonetheless possible that such geometrical distortions are the consequences of an underlying density wave structure, but *it is necessary to obtain (optically thin) millimeter thermal emission maps to confirm this scenario* (Piétu et al., 2005).

Among the key disk properties that one would like to determine is the total dust disk mass (remember that the dust represents only 1% of the total disk
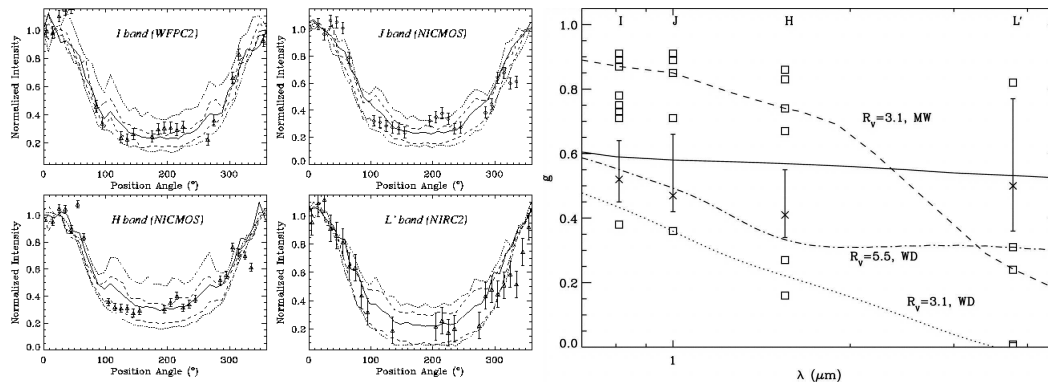
47

Fig. 44. *Left:* Measured normalized azimuthal intensity profiles along the GG Tau ring at several wavelengths, along with modeled profiles from Monte Carlo simulations. All profiles are almost identical, indicating that the scattering phase function hardly varies with wavelength. *Right:* Variation of the scattering asymmetry parameter with wavelength based on fits to the intensity profiles (crosses and associated errorbars). Note the almost constant behavior of g with wavelength, to be compared with Fig. 40. Empty squares are inappropriate models. The curves correspond to interstellar ($R_V = 3.1$) and molecular cloud ($R_V \approx 5.5$) dust populations. The solid line is the HH 30 dust model from Wood et al. (2002). From Duchêne et al. (2004).

mass, however). In the millimeter regime, this is straightforward since disks are optically thin. It is just a matter of assuming a certain dust opacity and a disk temperature, and one can derive analytically the total dust mass. With visible near-infrared scattered light images, the situation is not nearly as simple. The disks are optically thick and the photons we receive from the disk have barely scratched the surface. What lies beneath it is of no consequence to images and, therefore, we cannot infer the total mass of a disk based on a single image alone. In an edge-on geometry, the optical depth to the central star is proportional to the product if the disk mass (actually, the dust column density along the line of sight) and the dust opacity. While this seems to provide a constraint on the disk mass, it really yields a lower limit to the total dust mass, furthermore affected by at least a factor of 10 uncertainty from the assumed dust opacity (see Fig. 40). To go beyond this simple lower limit, one has to assume a certain disk geometry to convert a line of sight column density into a total disk dust mass. This derived disk mass is therefore model-dependent. In conclusion, *disk masses can only be determined with reasonable accuracy from optically thin measurements in the millimeter regime.*

As discussed in Section 5.1, combining scattered light images with other datasets is the best approach to constrain the dust properties. This is illustrated in Fig. 43, in which a millimeter interferometric map is compared to an HST scattered light image of the ring. The ring is uniform in the long-wavelength regime, indicative of a smooth, unperturbed, density structure. From that dataset, the inner and outer radii of the ring, as well as its inclination can be determined with high accuracy. The scattered light image, on other hand,
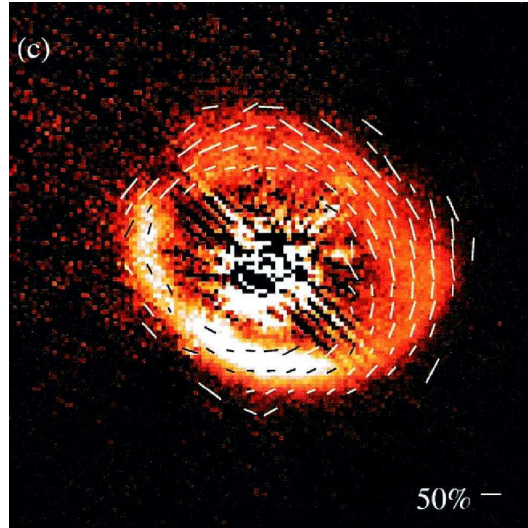
48

Fig. 45. Scattered light image of the GG Tau ring obtained at $1\,\mu$m with HST/NICMOS (Silber et al., 2000). Superposed are the linear polarization vectors. Note that the polarization rate is low where the ring intensity is high and highest close to the ring semi-major axis, as expected from single scattering off an inclined ring.

reveals two sides with very different properties: the northern side is bright and narrow whereas the southern side is faint and wide. This is explained by the geometry of the ring, sketched in the left panel of Fig. 43, as the northern part is inclined towards the observer. We then receive forward-scattered photons from that region, whereas the southern region produces back-scattered photons. Since dust populations generally favor forward-throwing scattering, the northern section is brighter. In this system, where the geometry is well determined, it is possible to use the images to reconstruct the scattering phase function and compare it to dust models.

This analysis has now been conducted over a broad wavelength range (see Fig. 44), revealing an almost achromatic phase function that is incompatible with interstellar and molecular cloud dust populations (McCabe et al., 2002; Duchêne et al., 2004; Krist et al., 2005). A strongly forward-scattering at a wavelength as long as $3.8\,\mu$m implies the presence of a substantial amount of grains large than $1\,\mu$m in radius, which are virtually absent of the interstellar medium. Clearly, this is evidence that dust grains have started to grow to larger sizes. It is interesting to note that the dust model devised by Wood et al. (2002) to match the observed SED of HH 30, which includes grains as large as $1\,$mm, produces a very slowly declining asymmetry parameter wavelength-dependence, roughly consistent with the observations of the GG Tau ring (right panel of Fig. 44). This particular model can, however, be excluded on the basis of another high-resolution observation of the ring: an HST/NICMOS $1\,\mu$m linear polarization map (see Fig. 45). Scattering off dust grains induces a linear polarization in the originally unpolarized starlight. The degree of linear
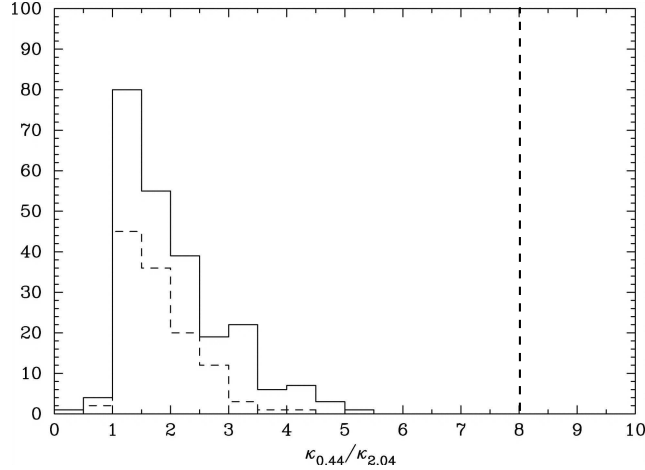
49

Fig. 46. Probability distribution of the ratio of dust opacities at 0.44 μm and 2.04 μm in the HH 30 disk. The vertical dashed line indicates the expected value for a dust population representative of molecular clouds; the corresponding value for interstellar dust is even higher, beyond the plot limit. Adapted from Watson & Stapelfeldt (2004).

polarization depends on the geometry of the scattering, with a maximum where the scattering angle approaches 90°. Scattering at low angles, on the other hand, induces low levels of linear polarization. Furthermore, the degree of linear polarization depends on the dust grain size: small grains tend to polarize very efficiently, as opposed to large dust grains. As a result, the maximum 1 μm polarization induced by the dust model proposed by Wood et al. (2002) is on the order of 40%. On the other hand, observations of the GG Tau ring have found linear polarization of 50% at scattering angle $\sim$ 110°, implying that small grains are proportionally more abundant in that particular system (Silber et al., 2000). Similarly, higher polarization rates have been found in the disk surrounding TW Hya at 1.6 μm than at 1.2μm (Hales et al., 2006), a behavior that is expected for small grains-dominated dust populations.

Another evidence of grain growth has been unearthed by a multi-wavelength analysis of the HH 30 edge-on disk. While the analysis of a single image is sensitive to the product κM_d (see above), simultaneously analyzing images obtained at different wavelengths readily constrains the wavelength dependence of the dust opacity. The longer the wavelength, the smaller the dust opacity and the closer the two scattered light nebulae. While in the case of the HV Tau C disk, the opacity law is consistent with interstellar extinction, the dust population of the HH 30 disk clearly has a much flatter ("grayer") opacity law, indicative of the presence of larger grains than in the interstellar medium or than in molecular clouds (see Fig. 46).

While scattered light images of young circumstellar disks imply the presence of larger grains than in the interstellar medium or in molecular clouds, it is difficult to infer a precise size distribution. To convert the observed scatter-
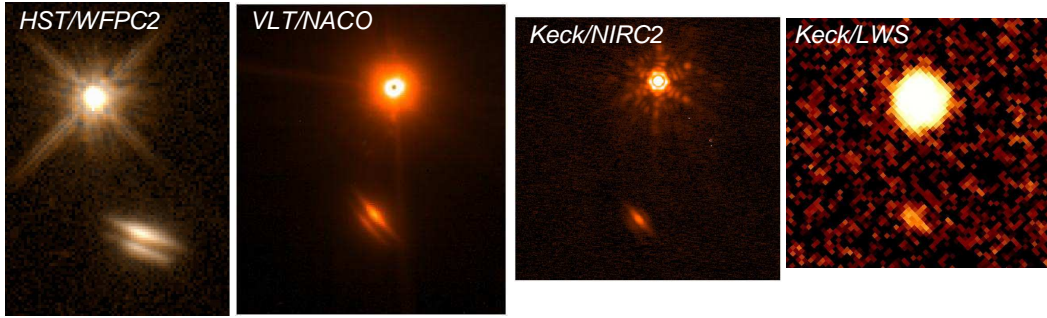
Fig. 47. Mosaic of scattered light images of the HK Tau binary. From left to right, images obtained at 0.6 μm (HST; Stapelfeldt et al., 1998), 2.2 μm and 3.8 μm (AO; credit: Caer McCabe, IPAC) and 11.3 μm (direct imaging; McCabe et al., 2003). The images are not all rotated exactly in the same way.

ing/absorption properties into grain sizes, several questionable assumptions have to be made. In particular, spherical grains are generally assumed. Also, some models assume that the grains are aggregates of smaller inclusions of different composition while other models keep the grains of different composition physically separated. In most models, dust size distributions are parametrized, sometimes with a simple power law, sometimes with more complex analytical forms. As a consequence, the quoted values of a_{max}, the maximum size of dust grains, are not easy to compare from one study to another and, generally speaking, are very much model-dependent.

Despite these limitations, high-resolution scattered light images clearly offer evidence that grain grow within circumstellar disks. Indeed, in the mid-infrared, *only* grains of a few microns in size can scatter efficiently. Because of the much larger background emission in the mid-infrared, it is virtually impossible to obtain high-contrast images in this regime, however. On the other hand, atmospheric turbulence is much less of a problem, since direct images on large telescopes are frequently close to being diffraction-limited. It is therefore possible to detect mid-infrared scattered light in the case of (low contrast) edge-on disks, provided they are sufficiently optically thick to remain opaque at these longer wavelengths. This has already been successfully done for the HK Tau B disk, which has now been images over a very wide range of wavelengths (0.6–11.3 μm). The detection of this disk at 11.3 μm in scattered light not only requires that the disk be optically thick, but also that it contains copious amounts of grains larger than 1 μm. Interestingly, the discovery image, obtained in the visible, was consistent with interstellar-like dust, seemingly excluding the presence of larger grains. Indeed, a continuous power-law size distribution extending beyond 1 μm is inconsistent with visible and near-infrared images (McCabe et al., 2003). A solution to this problem could be that the dust population is stratified, as suggested by models of disk evolution. At longer wavelengths, we would receive scattered light from deeper disk layers, that contain larger grains than the disk surface. Such a stratified structure for this particular disk is also reinforced by the analysis of its mil-

51

limeter emission map (Duchêne et al., 2003). It has also been inferred in the GG Tau ring, for which it is possible prove geometrically that scattering at $3.8\,\mu$m arises from a layer located twice as close from the midplane than the visible scattered light (Duchêne et al., 2004). The presence of larger grains in deeper layers is therefore qualitatively consistent with the observations. Indeed, a recent hydrodynamical simulation of the GG Tau ring has proven that the expected amount of vertical settling is consistent with observations, a first direct comparison of model predictions with observations as far as disk stratification is concerned (Pinte et al., 2007).

Scattered light images can therefore prove critical in determining both basic geometrical properties of disks as well as in providing constraints on their dust content. They offer insight that other techniques, such as SED modeling or thermal emission mapping cannot provide. It is therefore important to consider obtaining such images when studying a particular disk, something that current high-angular resolution techniques offer in many (but not all) cases. More importantly, this section has illustrated the fact that *interpreting datasets of different nature with a single model is the more powerful tool and, indeed, the only way to constrain with some degree of certainty many of the disk and dust physical parameters*. Linear polarization mapping is a natural complement to scattered light images. Other key observables are thermal emission maps, silicate emission feature spectroscopy, entire SED modeling and long baseline near-infrared interferometry, discussed in other contributions to this book.

Finally, it must be emphasized that interpreting scattered light images (and other type of observations of disks) is not a matter of measuring a flux in an image and to use an analytical formula to derive a maximum grain size or the exact degree of flaring of a disk, for instance. *Synthetic disk images (and SEDs) must be computed and compared to observations.* Numerical codes that treat the radiative transfer in the specific case of circumstellar disks (flared geometry, treatment of large optical depths) are now available in the community and it is important to use them to place conclusions about the physical structure and content of disks on more secure grounds. All results discussed here have been obtained by using at least one of these radiative transfer codes. S. Wolf's contribution summarizes both the principles, successes and limitations of existing codes in this framework.

6 Summary and perspectives

There is a wide array of high-angular resolution techniques. In this contribution, the focus has been placed on single-telescope diffraction limited imaging, which is accessible via space-based telescope or using speckle interferometry or AO on ground-based telescopes. Each of these techniques has its advan-

tages and limitations, so the best choice of instrument is dependent on the specifics of the science target and goals. These techniques are already in use to study both extrasolar planets and protoplanetary disks. In some cases, such as the upcoming search for mature planetary systems or the ongoing analysis of optically thin debris disks, the use of coronography is necessary to reach the necessary contrasts. The robustness of these techniques is now well established and the number of important results they enabled keeps increasing.

Long-baseline interferometry opens a new window to study both planets and circumstellar disks by providing access to an even higher spatial resolution. It is important to remember, however, that it will be possible to interpret such observations only with other datasets at hand, such as high-angular images as described here. Only a combination of observational techniques can yield a clear view of a particular object.

In the upcoming future, new generation AO devices (LGS AO, extreme AO, MCAO) will be installed on large telescopes, opening a new avenue to study exoplanets and disks, thanks to higher contrasts in particular. They will also be key to the next generation of large telescopes. Together with a larger space telescope (JWST), we can expect a new harvest of high-angular observations of planets and disks, and their related slew of discoveries concerning the structure and dust content of disks on one hand and the frequency and properties of extrasolar planets on the other hand.

As a final note, astronomy has constantly been driven by the ability to build larger telescope and to search for finer details in order to better understand the physics of stars, and the techniques described here will be central in continuing this tradition in the upcoming decades.

A   Related websites

In addition to the references provided throughout this contribution, I list below a few websites that provide useful additional information and references. This is just a limited list that can be used as a springboard for the interested reader towards many more informative sites and documents, offering a much deeper view of many details that were only touched upon in this contribution.

Speckle interferometry websites:

– `http://www.cis.rit.edu/research/thesis/bs/2000/hoffmann/thesis.html`
– `http://www.mpia.de/homes/tubbs/thesis/contents.html#full_contents`

Adaptive optics websites:

- `http://www.eso.org/projects/aot/`
- `http://cfao.ucolick.org/`
- `http://www.ctio.noao.edu/ãtokovin/tutorial/intro.html`

Extrasolar planets websites:

- `http://exoplanet.eu/`
- `http://exoplanets.org/`
- `http://obswww.unige.ch/ũdry/planet/planet.html`

Circumstellar disks websites:

- `http://www.circumstellardisks.org/`
- `http://www.disksite.com/`

References

Apai, D., Pascucci, I., Brandner, W., Henning, T., Lenzen, R., Potter, D. E., Lagrange, A.-M., & Rousset, G. 2004, A&A, 415, 671

Azouit, M., & Vernin, J. 2005, PASP, 117, 536

Babcock, H. W. 1953, PASP, 65, 229

Backman, D. E., & Paresce, F. 1993, in *Protostars and Planets III*, 1253

Barrière-Fouchet, L., Gonzalez, J.-F., Murray, J. R., Humble, R. J., & Maddison, S. T. 2005, A&A, 443, 185

Bertout, C., Basri, G., & Bouvier, J. 1988, ApJ, 330, 350

Beust, H., & Dutrey, A. 2006, A&A, 446, 137

Bloemhof, E. E., Dekany, R. G., Troy, M., & Oppenheimer, B. R. 2001, ApJL, 558, L71

Bouwman, J., Meeus, G., de Koter, A., Hony, S., Dominik, C., & Waters, L. B. F. M. 2001, A&A, 375, 950

Burrows, C. J., et al. 1996, ApJ, 473, 437

Burrows, A., Sudarsky, D., & Hubeny, I. 2004, ApJ, 609, 407

Chauvin, G., Lagrange, A.-M., Dumas, C., Zuckerman, B., Mouillet, D., Song, I., Beuzit, J.-L., & Lowrance, P. 2004, A&A, 425, L29

Chauvin, G., et al. 2005, A&A, 438, L29

Close, L. M., et al. 2003, ApJL, 598, L35

Doel, A. P., Dunlop, C. N., Buscher, D. F., Myers, R. M., Sharples, R. M., & Major, J. V. 2000, New A., 5, 223

Dominik, C., Blum, J., Cuzzi, J. N., & Wurm, G. 2007, Protostars and Planets V, 783

Duchêne, G., Ménard, F., Stapelfeldt, K., & Duvert, G. 2003, A&A, 400, 559

Duchêne, G., McCabe, C., Ghez, A. M., & Macintosh, B. A. 2004, ApJ, 606, 969

Duchêne, G., Bontemps, S., Bouvier, J., André, P., Djupvik, A. A., & Ghez, A. M. 2007, A&A, 476, 229

Fitzgerald, M. P., Kalas, P. G, Duchêne, G., Pinte, C., & Graham, J. R. 2007, ApJ, 670, 556

Fried, D. L. 1966, Journal of the Optical Society of America, 56, 1372

Fukagawa, M., et al. 2004, ApJL, 605, L53

Gavel, D. 2007, SPIE Proc., vol., 6467, in press

Genzel, R., Eckart, A., Ott, T., & Eisenhauer, F. 1997, MNRAS, 291, 219

Grady, C. A., et al. 2001, AJ, 122, 3396

Guilloteau, S., Dutrey, A., & Simon, M. 1999, A&A, 348, 570

Guyon, O. 2005, ApJ, 629, 592

Hales, A. S., Gledhill, T. M., Barlow, M. J., & Lowe, K. T. E. 2006, MNRAS, 365, 1348

Hinkley, S., et al. 2007, ApJ, 654, 633

Janson, M., Brandner, W., Henning, T., & Zinnecker, H. 2006, A&A, 453, 609

Kenyon, S. J., & Hartmann, L. 1987, ApJ, 323, 714

Kessler-Silacci, J., et al. 2006, ApJ, 639, 275

Koresko, C. D. 1998, ApJL, 507, L145

Koresko, C. D. 2002, AJ, 124, 1082

Krist, J. E., et al. 2005, AJ, 130, 2778

Kuhn, J. R., Potter, D., & Parise, B. 2001, ApJL, 553, L189

Labeyrie, A. 1970, A&A, 6, 85

Lowrance, P. J., et al. 1999, ApJL, 512, L69

Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556

Marois, C., Macintosh, B., & Barman, T. 2007, ApJL, 654, L151

McCabe, C., Duchêne, G., & Ghez, A. M. 2002, ApJ, 575, 974

McCabe, C., Duchêne, G., & Ghez, A. M. 2003, ApJL, 588, L113

McCaughrean, M. J., Close, L. M., Scholz, R.-D., Lenzen, R., Biller, B., Brandner, W., Hartung, M., & Lodieu, N. 2004, A&A, 413, 1029

McElwain, M. W., et al. 2007, ApJ, 656, 505

Meyer, M. R., Backman, D. E., Weinberger, A. J., & Wyatt, M. C. 2007, Protostars and Planets V, 573

Monin, J.-L., & Bouvier, J. 2000, A&A, 356, L75

Muzerolle, J., Calvet, N., Hartmann, L., & D'Alessio, P. 2003, ApJL, 597, L149

Nakajima, T., Oppenheimer, B. R., Kulkarni, S. R., Golimowski, D. A., Matthews, K., & Durrance, S. T. 1995, Nature, 378, 463

Neuhäuser, R., Guenther, E. W., Wuchterl, G., Mugrauer, M., Bedalov, A., & Hauschildt, P. H. 2005, A&A, 435, L13

Noll, R. J. 1976, Journal of the Optical Society of America, 66, 207

Patience, J., Ghez, A. M., Reid, I. N., Weinberger, A. J., & Matthews, K. 1998, AJ, 115, 1972

Perrin, M. D., Duchêne, G., Kalas, P., & Graham, J. R. 2006, ApJ, 645, 1272

Piétu, V., Guilloteau, S., & Dutrey, A. 2005, A&A, 443, 945

Pinte, C., Fouchet, L., Ménard, F., Gonzalez, J.-F., & Duchêne, G. 2007, A&A, 469, 963

Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, PASP, 111, 587

Roberts, L. C., Jr., & Neyman, C. R. 2002, PASP, 114, 1260

Roddier, F. 1988, Appl. Optics, 27, 1223

Roddier, F., Northcott, M., & Graves, J. E. 1991, PASP, 103, 131

Schneider, G., Thompson, R. I., Smith, B. A., & Terrile, R. J. 1998, Proc. SPIE, 3356, 222

Scholz, R.-D., McCaughrean, M. J., Lodieu, N, & Kuhlbrodt, B. 2003, A&AL, 398, L29

Seifahrt, A., Neuhäuser, R., & Hauschildt, P. H. 2007, A&A, 463, 309

Sheehy, C. D., McCrady, N., & Graham, J. R. 2006, ApJ, 647, 1517

Silber, J., Gledhill, T., Duchêne, G., & Ménard, F. 2000, ApJL, 536, L89

Sivaramakrishnan, A., Lloyd, J. P., Hodge, P. E., & Macintosh, B. A. 2002, ApJL, 581, L59

Stapelfeldt, K. R., Krist, J. E., Menard, F., Bouvier, J., Padgett, D. L., &

Burrows, C. J. 1998, ApJL, 502, L65

Tallon, M., & Foy, R. 1990, A&A, 235, 549

Tubbs, R. N. 2003, Ph.D., Cambridge University

Tuthill, P. G., Men'shchikov, A. B., Schertl, D., Monnier, J. D., Danchi, W. C., & Weigelt, G. 2002, A&A, 389, 889

Véran, J.-P., Rigaut, F., Maître, H., & Rouan, D. 1997, J. Opt. Soc. Am. A, 14, 3057

Watson, A. M., & Stapelfeldt, K. R. 2004, ApJ, 602, 860

Weidenschilling, S. J. 1977, MNRAS, 180, 57

Weidenschilling, S. J. 1980, Icar., 44, 172

Weidenschilling, S. J. 1984, Icar., 60, 553

Weigelt, G., Beuther, H., Hofmann, K.-H., Meyer, M. R., Preibisch, T., Schertl, D., Smith, M. D., & Young, E. T. 2006, A&A, 447, 655

Weingartner, J. C., & Draine, B. T. 2001, ApJ, 548, 296

Wood, K., Wolff, M. J., Bjorkman, J. E., & Whitney, B. 2002, ApJ, 564, 887