Optical Design Group

MWI PERFORMANCE DOCUMENT

René Doyon
Jean-Francois Lavigne
David Lafrenière
Simon Thibault

Revision No.	Author & Date	Approval & Date	Description
Revision 1	Simon Thibault 2004-10-12	Simon Thibault 2004-10-12	Creation
Revision 2	Simon Thibault 2004-10-19	Simon Thibault 2004-10-19	Mid-term document

TABLE OF CONTENTS

1 PURPOSE

The MWI Performance Document describes:

- Performances simulation with Zemax Physical Optic Progragation of the MWI.
- Speckle noise suppression capabilities of the MWI and IFU in the absence of differential phase or amplitude errors between different wavelengths.

2 APPLICABLE DOCUMENTS

Document ID	Source	Title

3 PERFORMANCES SIMULATION WITH ZEMAX

3.1 Zemax Physical Optics Propagation (POP) tool

The different simulations realized in this chapter were done with the Physical Optics Propagation (POP) tool in Zemax. This tool takes in consideration both the optical aberrations of the system and its diffraction effects. Simulations of the performances of our camera need such a tool since those two effects can have a non-negligible influence on the results.

The POP tool will be used here with its Fresnel propagation to go from one surface to another. This will allow us to know the impact of the wavefront deformation introduced by any surface along the optical path.

3.2 Specifications

The following specifications have been defined for the MWI (Doyon, 2004):

- **Detector:** Hawaii-2RG 2040x2040, 18 µm pixel
- **Microlens array:** square shape, 54 µm pitch (= 3 detector pixels assuming unit magnification), f/6 output corresponding to $\lambda f/D$~9 µm or half a detector pixel.
- **Input focal ratio:** f/90
- **FOV**: 5.3"x5.3"
- **Spectral resolution:** ~50
- **Wavelengths:** 1.52, 1.58, 1.64 and 1.70 µm

The possibility to use a microlens array with a 72 µm pitch so that four detector pixels covers its area will also be studied. The input focal ratio would then be f/120 to keep the same sampling in the focal plane. This setting will be referred to as the 4x4 pixels case while the previous one will be referred to as the 3x3 pixels case. The FOV used will also be smaller than 5.3” x 5.3” in the simulations presented in this report so that the matrix size stays reasonable.

3.3 Optical system definition

Our optical system starts with an 8-m diameter paraxial lens with a 720 m or a 960 m focal length to have respectively a f/90 or a f/120 input. The lenslet array is placed in that focal plane. Since the radius of curvature of the microlenses necessary to have the required f/# is dependent on the wavelength, an average over our waveband is taken. This can easily be calculated with the lensmaker equation applied to a plano-convex lens:

$$\frac{1}{f} = (n-1)\left(-\frac{1}{R}\right) .$$

Some relay can then be added after the microlenses focal plane to simulate one channel of the camera.

3.4 POP settings

One has to be careful in the simulation of the microlens array with the POP tool. The phase reference for the wavefront propagation has to be set to a plane instead of the best fit that is automatically chosen by Zemax. This comes from the fact that Zemax takes the diameter of the lenslet array instead of the diameter of a single lenslet for its calculation of the Fresnel number. Then, its calculations are done as if it was in the near field when in fact it is in the far field. The way to correct this is to force Zemax to take a plane reference for the phase as it would have been the case in the far field (see Zemax manual June 19, 2003 p.456-457 and p.469 for more details).

The next step is to set the proper parameters for the propagation of the wavefront with the POP tool represented in figure 1. The first one to be set is the x and y sampling. To have a proper sampling of each microlens, to have a reasonable field of view and a tolerable computation time, the sampling is set to 1024 x 1024 pixels. In order to diminish errors due to aliasing from the Fourier transforms, only a matrix of 512 x 512 pixels is kept on the detector.

The width of the area sampled in the pupil plane in mm has then to be set. This is decided by the size of a pixel in the focal plane that is needed. From there, the width, W, or FOV of the pupil plane can easily be calculated with the following equations:

	Pupil Plane	Focal Plane
FOV	W	$N_{px}\lambda F/W$
Pixel size	W/N_{px}	$\lambda F/W$

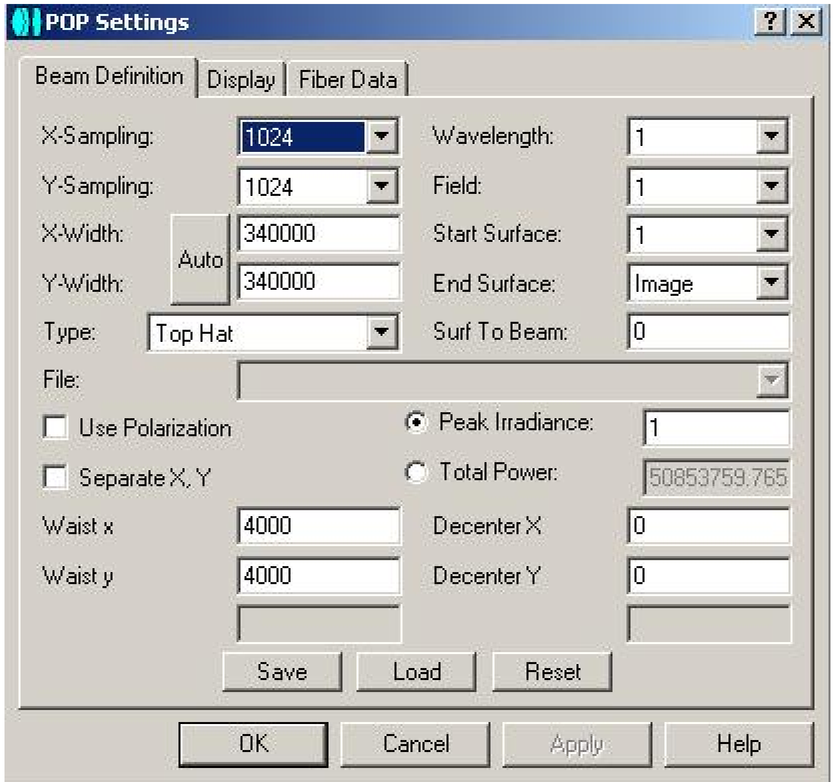

Figure 1 : Example of the POP settings for a propagation at 1.70 microns

N_{px} is the sampling or number of pixels, λ is the wavelength used and F is the focal distance of the primary mirror. The size of a pixel is set to be 3.6 x 3.6 microns so that 15 x 15 pixels are on one micro-lens. This pixel size is arbitrary and the error introduced by different sampling of a micro-lens is still to be investigated. In the example showed in figure 1, a wavelength of 1.70 microns is considered and a FOV of 340 000 mm is found. The FOV for each wavelength can be found in table 1.

λ (µm)	FOV 3x3 px (mm)	FOV 4x4 px (mm)
1.52	304 000	405 333
1.58	316 000	421 333
1.64	328 000	437 333
1.70	340 000	453 333

Table 1 : Pupil FOV necessary to keep a pixel size of 3.6 microns in the focal plane

3.5 Transmission through a perfect relay

Figures 2, 3, 4 and 5 show respectively the sag surface of a 33 x 33 microlens array, the PSF in the telescope focal plane, the PSF obtained in the microlens focal plane and the

PSF obtained after a perfect relay with the settings showed in figure 1. Figure 6 is a comparison of both PSFs with a cut through the x-axis. This figure shows that some small errors are introduced by the propagation. To know the real impact of such errors in our simulations it will be imperative to compare the reconstructed PSFs at each wavelength. To do so, the PSFs in the telescope focal plane will have to be computed outside Zemax and used as a source file in the POP propagation. This problem is explained in more details in section 3.7.

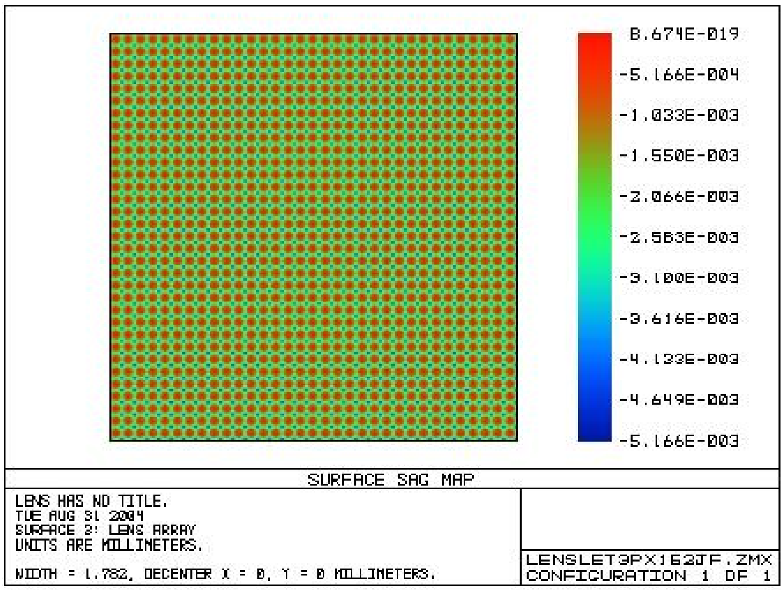

Figure 2 : sag surface of the microlens array

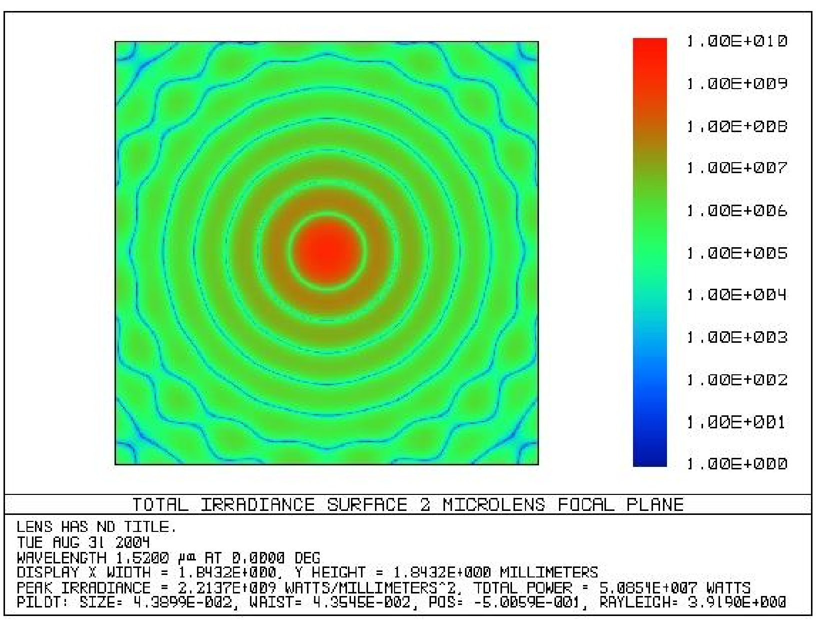

Figure 3 : PSF in the telescope focal plane

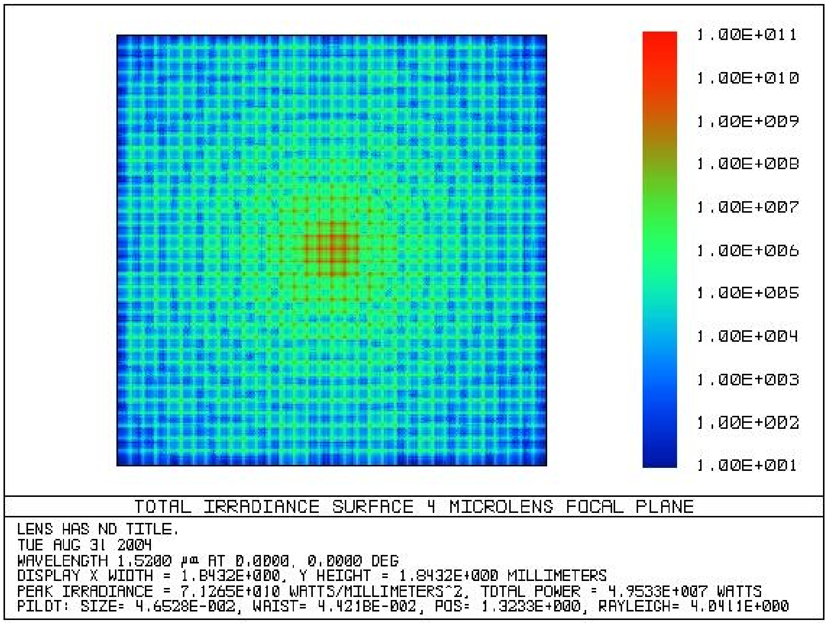

Figure 4 : PSF inthe microlens focal plane

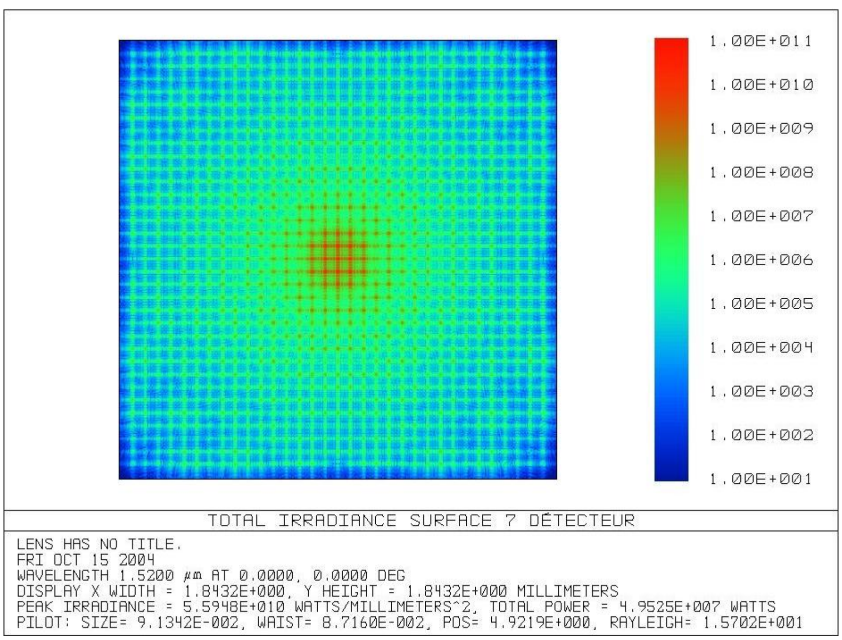

Figure 5 : PSF after a perfect relay

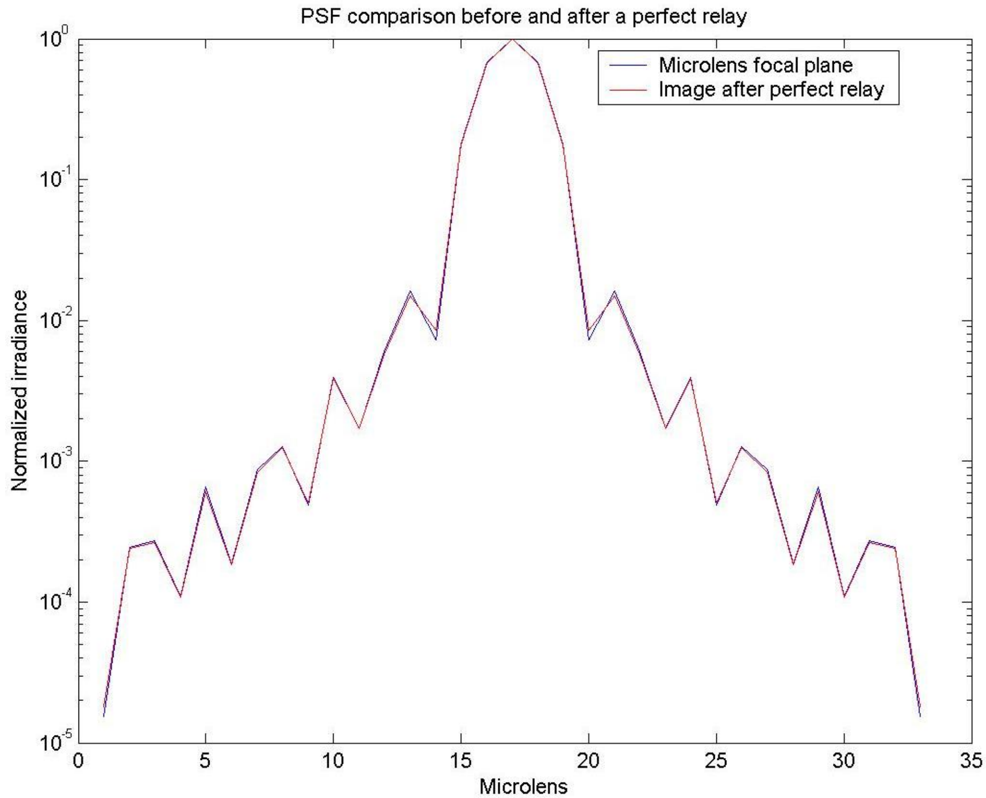

Figure 6 : Cross x cut of the PSF in the microlens focal plane and after a perfect relay

3.6 Sensitivity to optical aberrations

It is important to characterize the influence of the different optical aberrations on the camera performances in order to do the best optical design possible. To do so, a Zernike fringe phase surface is applied to an image of a microlens which is then refocused on the detector. In this setting, only one microlens is illuminated by a top hat function. The metric chosen is the ensquared energy. Though, the absolute value of the ensquared energy in a given square calculated by Zemax fluctuates with different sampling and is then not reliable. The metric will then be the relative variation of the ensquared energy of a given aberration compared to the case with no aberration for a constant sampling. It is considered that there is one guard pixel and that the PSF is centered in the middle of the remaining pixels. Three different cases are investigated: the 3x3 pixels case, the 4x4 pixels case and the 4x4 pixels case with the microlens array having a tilt of 45° relative to the detector. The aberrations investigated were defocus, astigmatism in x and in y, coma and spherical aberration. The results for each case are presented in graphics 1, 2 and 3 and a representation of each aberration is presented in figure 7.

It can be seen from the graphics that the ensquared energy is very sensible to spherical aberration in the three cases. Also, one can conclude that the simulation of a tilted lenslet array could be difficult with Zemax as it is shown in figure 7. The validity of the results for this case presented here are then plus or less reliable.

In order to see if one setting does better then the others their spherical aberration is compared in graphic 4. The three curves are about the same so for the rest of the simulations the initial setting with 3x3 pixels per microlens will be kept since it allows a bigger field of view on the detector. The maximal loss of ensquared energy is set to be of 10% in which case the optical system would need to have aberrations less than 0.3λ peak-to-valley. This sets the limit for the optical system image quality.

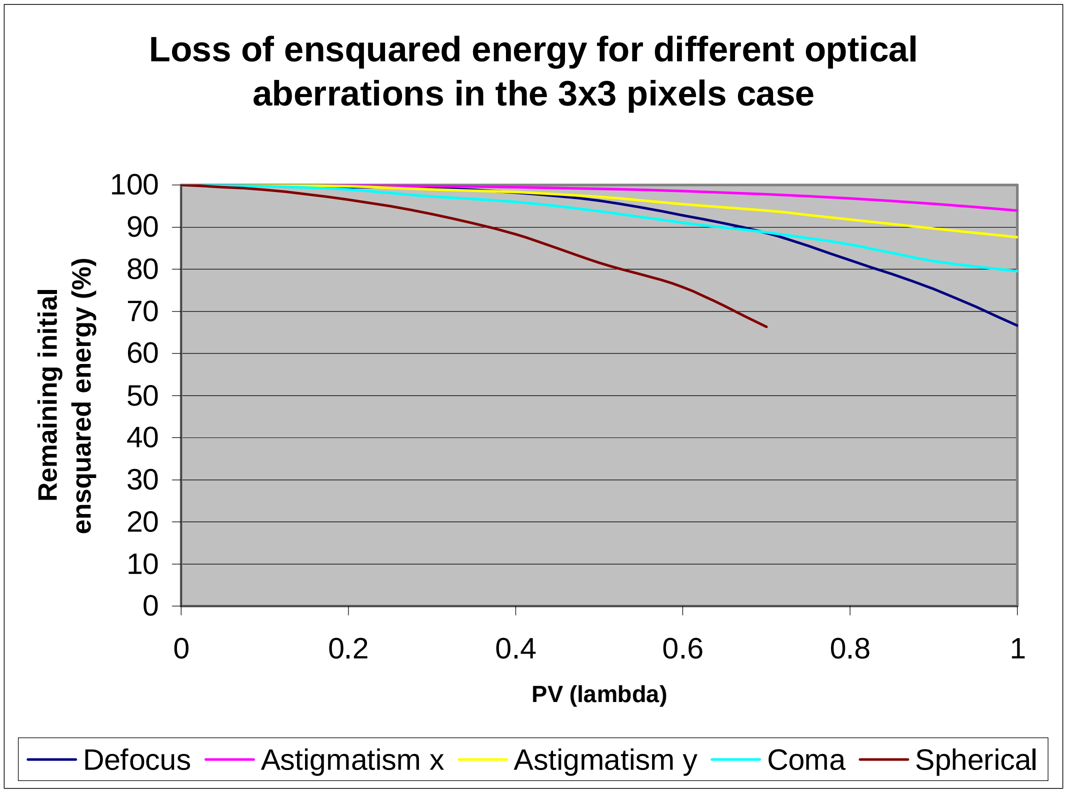

Graphic 1

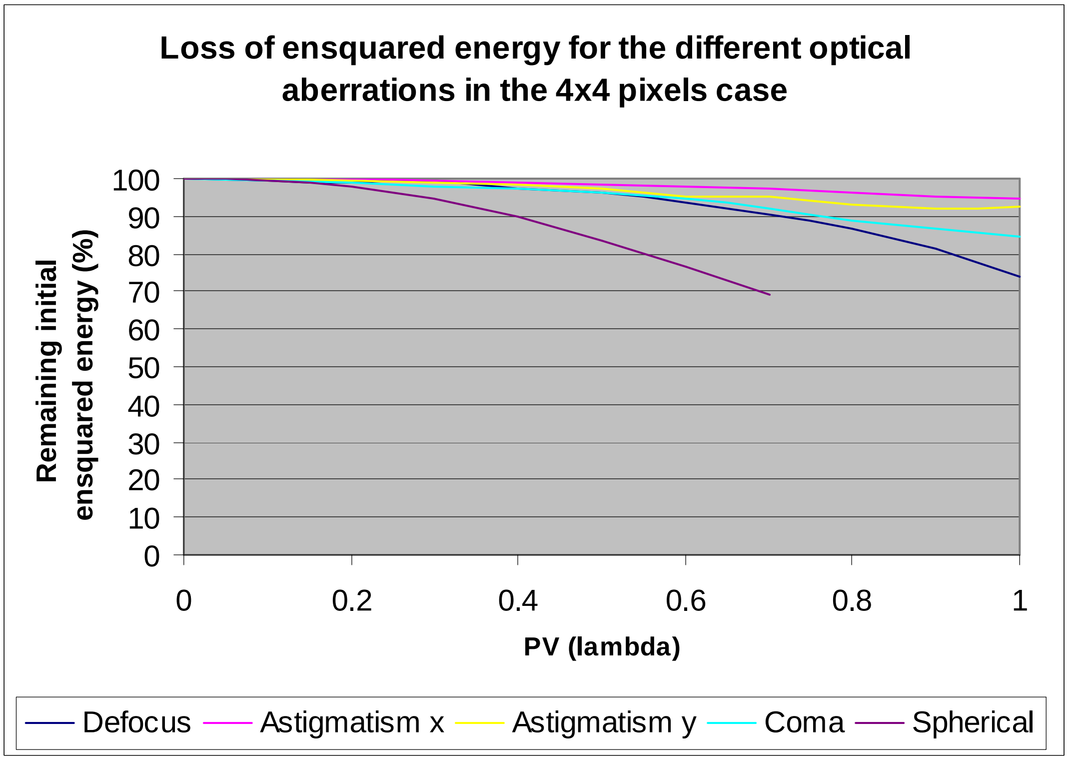

Graphic 2

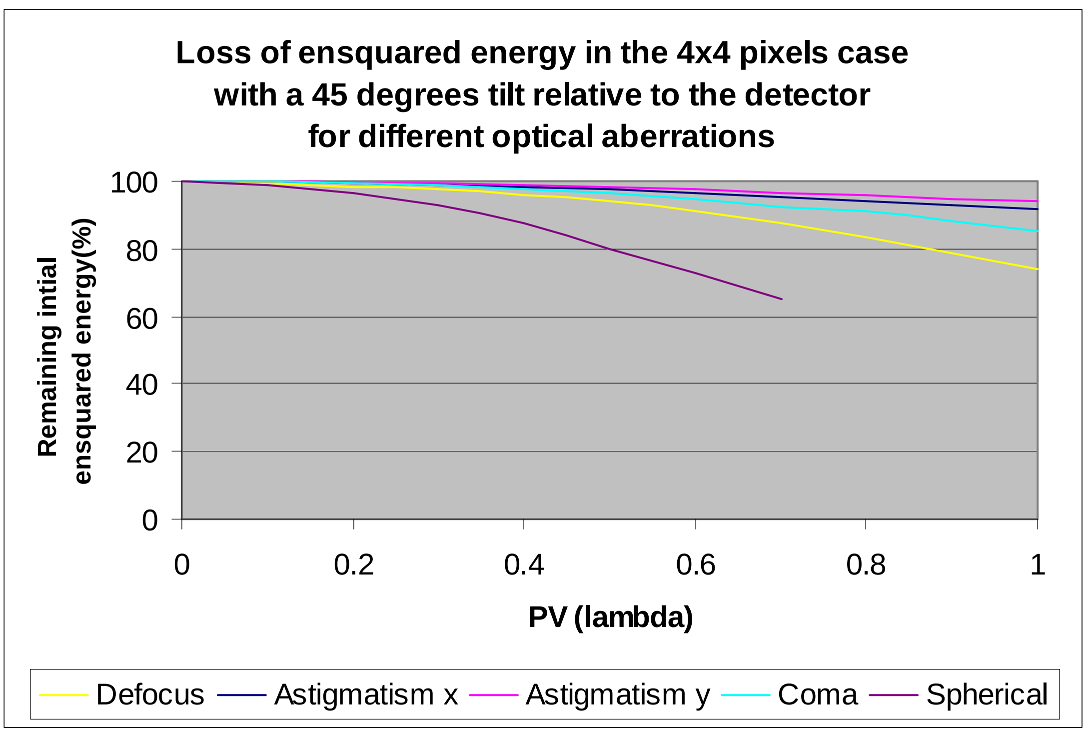

Graphic 3

3.7 Transmission through real relays

It can be seen from the report on the optical design that each channel of the instrument will have a different optical system associated to it. Different optical elements means different aberrations in each channel. In order to calibrate the impact of these different aberrations, the four PSFs will be transmitted trough their respective channel, reconstructed and subtracted one from the other. The remaining error will then set the performances attainable with the camera.

Though, when such simulations are done, some difficulties that were not present previously appears. First, the fact that the FOV in the pupil plane changes with the wavelength (see table 1) and that the sampling stays the same make the pixel size in the pupil plane change with the wavelength. This has the effect of changing the pixelisation of the primary mirror in the pupil plane from wavelength to wavelength which in turns introduces some numerical errors in the focal plane. This could normally be avoided by setting both the pixel size in the focal plane and in the pupil plane as constant for each wavelength and changing the sampling so that the FOV in the pupil plane showed in table 1 is obtained. Though, Zemax allowing sampling with matrices of a power of two only, this cannot be done. This error is shown in figure 8 for each wavelength used. To characterize the effect of this error, a radial average of each PSF is taken and compared with the others. It is then found that the error is not tolerable outside approximately the first ring of the PSF. This could be avoided by generating a source file containing a perfect PSF for each wavelength. The possibility of doing so will be studied in the near future.

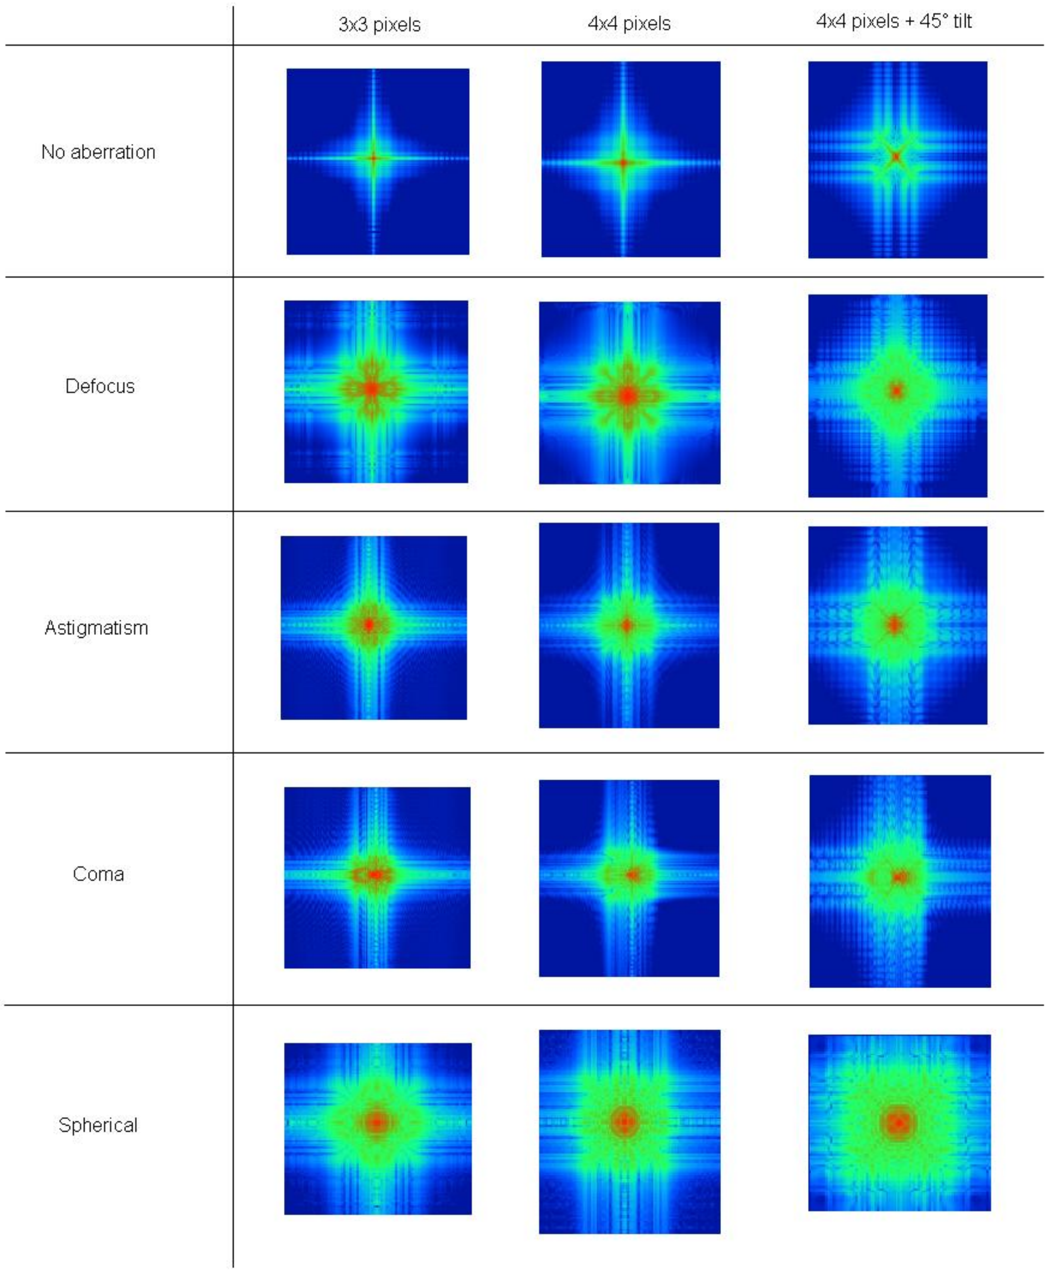

Figure 7 : PSFs examples for different aberrations in the three cases studied

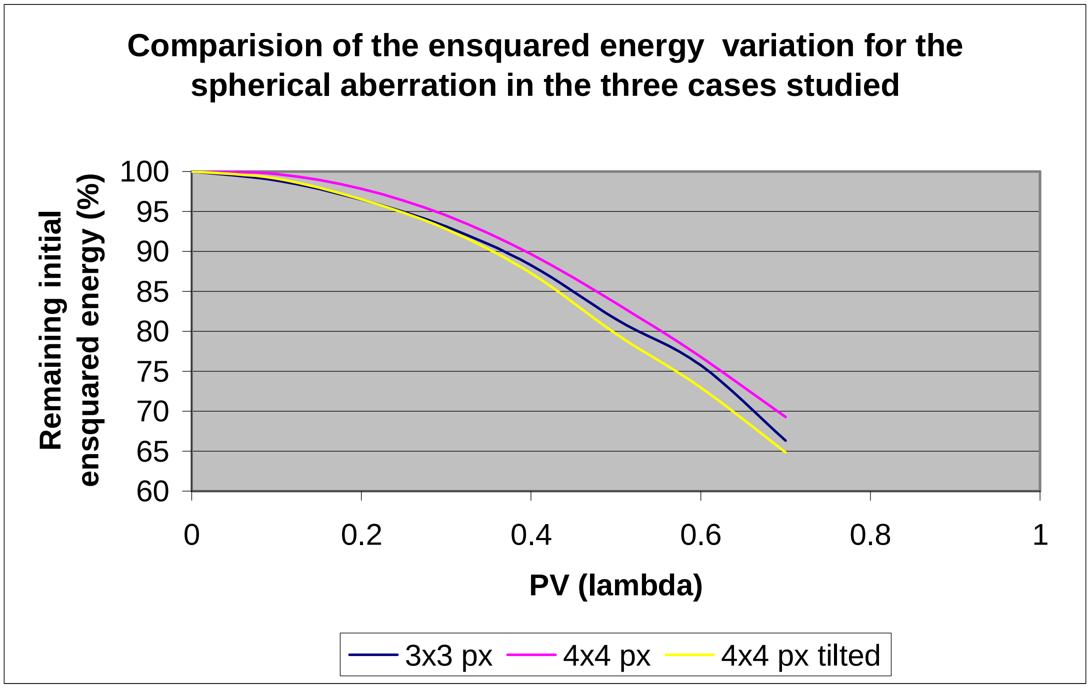

Graphic 4

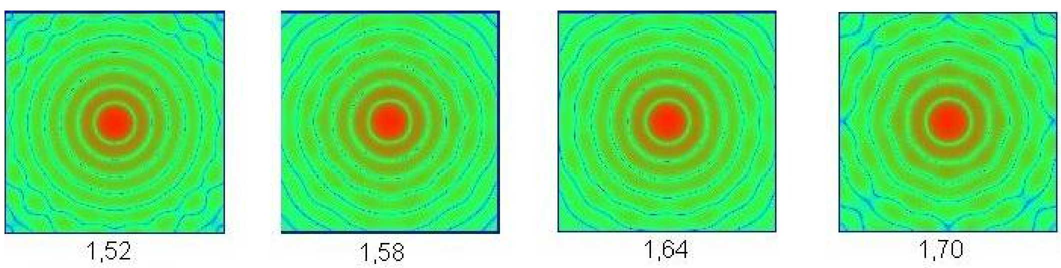

Figure 8 : PSFs in the focal plane for the different wavelength in microns. Their difference can be explained by the different pixelisation of the primary mirror in the pupil plane.

For now, one PSF has been generated in the telescope focal plane at a given wavelength and has been used as a source file for our four channels. This ensures that each channel starts with the same PSF. Even though we agree that some weird effects could come from the fact that the initial PSF is created at a given wavelength and then transmitted at another one, this is only used here for a representation purpose. The accurate answer would come from the use of analytically generated PSFs for each wavelength as mentioned above.

Figure 9 and 11 represent some preliminary results of the reconstructed PSFs with a cut along the x-axis for each channel without and with a coronagraph and figure 10 is the PSF obtained after the transmission through a real relay. The reconstruction algorithm used for these results was extremely simple and did not consider the possibility of using a flat field. In future simulations, this will be taken into account to get more accurate results. Also, the propagation done by Zemax through the real relays is not well understood yet. Some work still has to be done to ensure that the results are accurate.

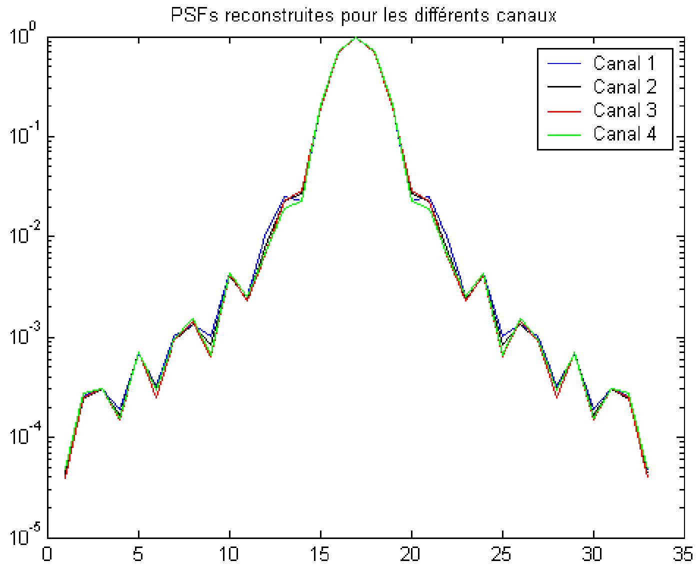

Figure 9 : Reconstructed PSFs for each channel of the relay

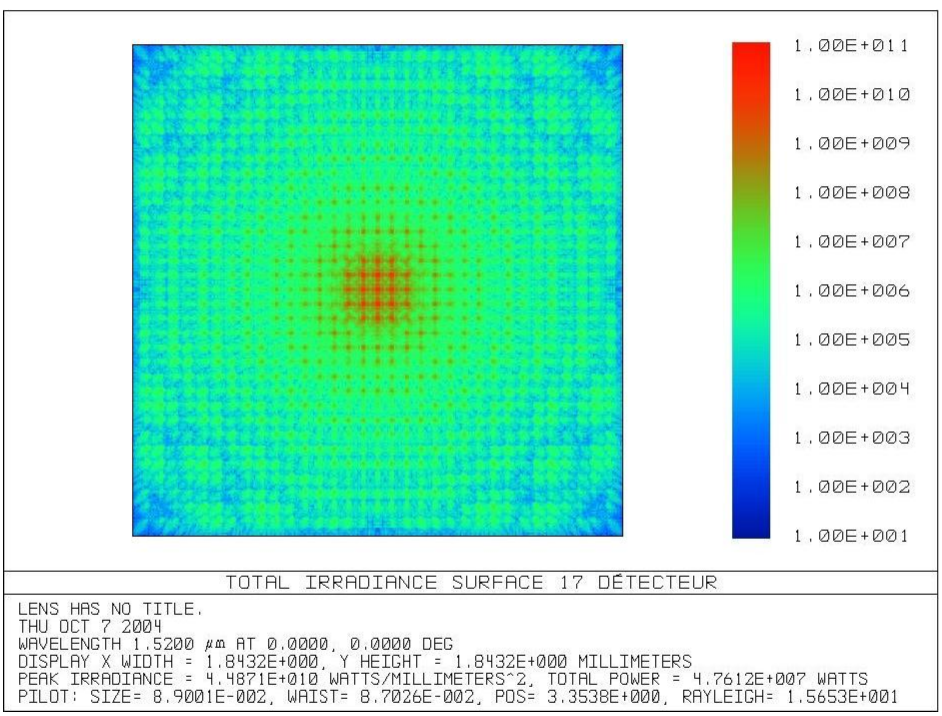

Figure 10 : PSF after the channel at 1.52 microns

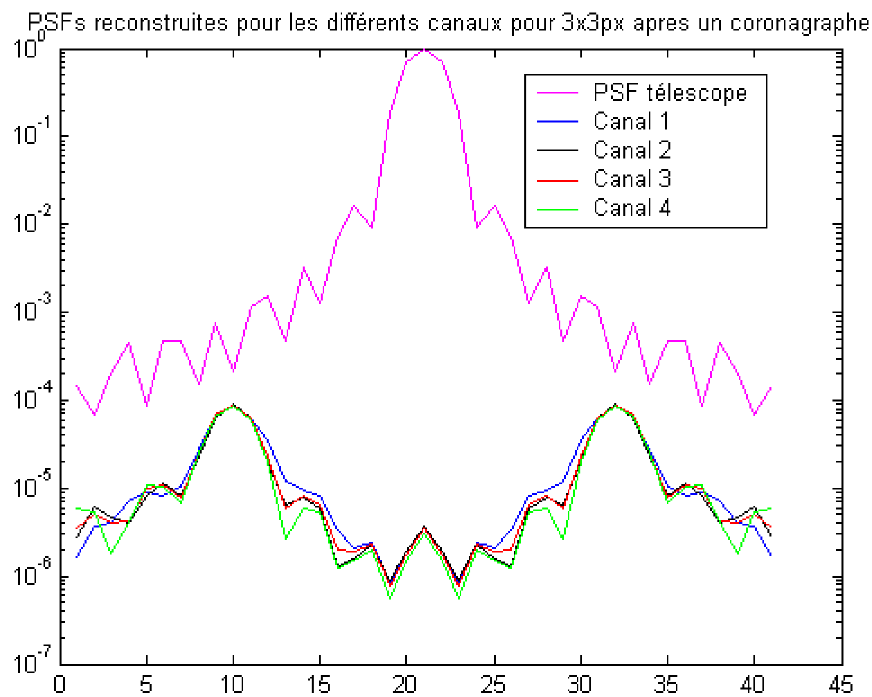

Figure 11 : Reconstructed coronagraphic PSFs for each channel of the relay

3.8 Future work

The first thing to be done will have to be the implementation of a source file containing a perfect PSF for each wavelength. In an ideal world, the reconstructed PSFs would then subtract perfectly one from the other. Although the errors introduced by the propagation through a perfect relay seem to be minimal, they will have to be evaluated more accurately. Then, the propagation done by Zemax through real relays will have to be better understood. Some unexpected features have been observed and still have to be explained.

Once this is done, we will have a powerful tool that will allow us to test a lot of different aspects of the camera such as the performances of different optical systems, the real impact of using four instead of three pixels per microlens and the advantage of using a guard pixel. We will also be able to generate different PSFs including phase errors closer to the ones that will be obtained at the entrance of the real camera and see how the instrument is dealing with them.

4 SPECKLES SUPPRESSION & SIGNAL RECOVERY

4.1 Introduction

Planet detection in the close vicinity of the PSF will likely be limited by static or slowly evolving speckles produced by small residual phase or amplitude errors. Both science cameras (MWI and IFU) proposed for this project have as a goal to suppress this speckle noise by some amount to improve the sensitivity of the instrument. This document looks at the speckle noise suppression capabilities of the MWI and IFU in the absence of differential phase or amplitude errors between different wavelengths.

At this stage, this work serves two purposes. First it verifies that simple algorithms can be used with the data product of the MWI and IFU to suppress speckles by sufficiently large factors. Second, it verifies that it is possible to suppress speckles without suppressing (completely) the signal of a companion of arbitrary spectrum.

4.2 Speckle suppression

Given N simultaneous images of a PSF, each at different but close wavelengths (a PSF cube), and **in the absence of differential aberrations**, then the same pattern of speckle is present in all images although at a different scale (diffraction scale $\propto \lambda/D$). Re-scaling images to a common scale brings all speckles to the same spatial location in all images and only a slight variation of the speckle intensity with wavelength remains. This evolution of the speckle intensity with wavelength is what I will refer to as the PSF chromatic evolution. The re-scaled images can be used to subtract the speckles using one of two general techniques.

Throughout this document, the speckle attenuation is defined as the ratio, in one annulus, of the standard deviation of the residual signal integrated over a disc of one FWHM divided by the standard deviation of the signal of the original PSF less an azimuthally averaged profile integrated over the same disc.

4.2.1 Images differences

Suppose that image n is of interest and that is it desired to suppress speckles in this image. The simplest operation to perform is to subtract image n+1 from it, this is called a simple difference (SD):

$$SD = I_n - I_{n+1}$$

This will subtract speckles to some level but will leave residuals due to the PSF chromatic evolution, the speckle attenuation factor obtained with an SD is $\sim\Delta\lambda/\lambda$, where λ is the wavelength of image n and $\Delta\lambda$ is the wavelength spacing of the two images. Smaller residuals are obtained if similar residuals from (I_{n-1} - I_n) are subtracted; this constitutes a double difference (DD):

$$DD = 0.5 \times (SD_n - SD_{n-1}) = (I_n - I_{n+1})/2 - (I_{n-1} - I_n)/2 = I_n - I_{n-1}/2 - I_{n+1}/2$$

The factor 0.5 is needed to normalize the signal properly to the signal in image n. The DD can attenuate speckles by a factor ~ $(\Delta\lambda/\lambda)^2$. In the same fashion, a double double difference, which can attenuate speckles by ~ $(\Delta\lambda/\lambda)^3$, is defined:

$$DDD = \frac{2}{3}(DD_n - DD_{n-1}) = I_n - I_{n-1} + I_{n-2}/3 - I_{n+1}/3$$

The SD, DD and DDD can be separately applied to all N images, in this case they are referred to as running SD, running DD and running DDD.

Suppose that a companion has a spectrum that is so sharply peaked over a very narrow wavelength range that it is present in only one of the N images, say in image n. A cold methane dwarf approaches this limit in the *H* band when observed at low resolution. In this case, the above operations will remove the speckles from image n but will leave the signal of the companion as it is absent in all other images used to do the subtraction.

If the companion is present in all images, then it will be at a different separation in each of the re-scaled images: a companion is moved radially by $r\Delta\lambda/\lambda$ when two images are brought to a common scale, where r is the original separation of the companion, λ is the wavelength of one of the two images and $\Delta\lambda$ is the wavelength spacing between the two images. If the displacement of the companion between images is greater than ~2λ/D (diameter of first dark ring), then effectively at a given separation in the re-scaled image the companion is present in a single image and the above considerations apply, namely speckles can be subtracted and the signal of the companion will be preserved.

If the displacement is less than 2λ/D, then at a given separation in the re-scaled images the companion is present in one image and partially present in a few other images. So the above procedure will subtract a fraction of the companion signal. Since the displacement is proportional to separation, companions at smaller separations will be suppressed more. A remedy for this is to use images more widely separated in wavelengths, when possible, to make the displacement between images larger.

4.2.2 Polynomial fit

An alternative way to subtract the speckles is to subtract a fitted spectrum from each "spectral pixel" of the re-scaled cube rather than subtracting images, this is the approach of Sparks & Ford 2002. In a re-scaled cube, the intensity of a pixel (or a speckle) varies smoothly with wavelength (chromatic evolution) and is easily fitted by a low order polynomial, which can be used to subtract the PSF contribution to that pixel (the speckle). This technique has to be implemented cautiously however, since a fit will be biased by the presence of a companion.

If a companion is present in only one of the N images, then the spectrum of a pixel containing the companion will have a point that will deviate from the smooth trend of the spectrum. This point will pull the polynomial fit and, as a result, the fit will subtract part of the companion signal. The solution is to ignore the point containing the companion signal when calculating the fit.

If the companion is present in all images and the displacement of the companion between images is greater than $2\lambda/D$, then effectively at a given separation in the re-scaled image the companion is present in a single image and ignoring a single point from the fit as above does the trick. If the displacement is less than $2\lambda/D$, then at a given separation in the re-scaled images the companion is present in one image and partially present in a few other images. So more than one point can pull on the fit and part of the companion signal will be subtracted. The solution is to ignore more points for the fit. This idea is illustrated in figure 4.1.

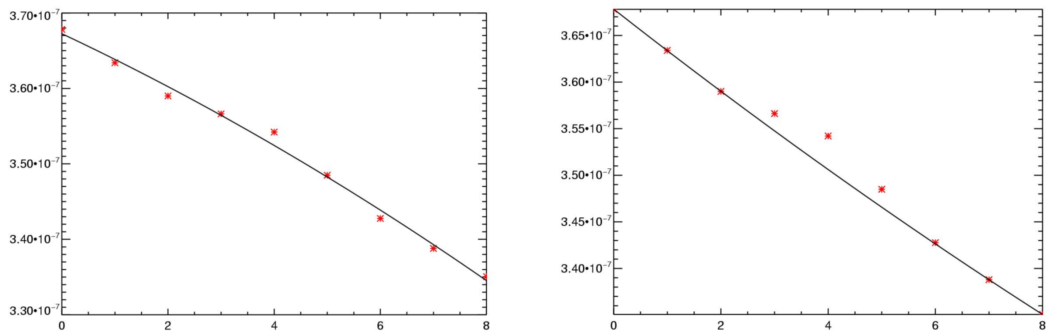

Figure 4.1: Spectrum of a pixel at a separation of 0.92" in a re-scaled cube of resolution R=50 including a planet at an original separation of 1", the planet appears at a separation of 0.92" in channel four in the re-scaled cube. (*left*) 2nd degree polynomial fit on all points (*right*) 2nd polynomial fit ignoring channels 3-6.

Generally, in a cube of logarithmically spaced spectral samples at a resolving power R, the companion is present in $2R/r$ images of the re-scaled cube, where r is the separation of the companion expressed in units of λ/D. This is the optimal number of points to neglect from the fit to ensure that the fit is not biased by the presence of the companion and that the companion is not subtracted. This is also twice the spacing between channels to use for SD, DD or DDD to ensure that the companion is not subtracted.

After the speckles have been subtracted, the residual cube can be scaled back to its original scale and the companion piles up again at all wavelengths. The "spectral pixel" at the position of the companion now contains the spectrum of that companion. For detection, the cube can be collapsed to improve the signal-to-noise.

4.3 Simulations

To verify the speckles suppression of the above algorithms and the preservation of the companion signal, a generic PSF cube was constructed. A single phase screen including 80 nm of static aberrations with a PSD $\propto f^{-3}$ was used to generate some level of speckles (no atmosphere, and no correction). The specific shape and structure of the PSF are not critical to consider only speckle suppression. The cube was originally generated at a resolution of R=1000 from 1.5 to 1.8 μm and was then properly binned to produce the data products of the MWI and IFU. The MWI data cube consists of 4 spectral channels (1.52 μm, 1.58 μm, 1.64 μm and 1.70 μm), each of bandwidth 2%. The IFU data cube, at R=50, consists of 9 contiguous spectral channels covering the wavelength range 1.5 to 1.8 μm. No source of noise was included in the PSF cube. It was later verified that the presence of noise (photon noise, flat-field noise, read noise, …) does not degrade the attenuation performance above the limit imposed by that noise.

For polynomial fits, it is always necessary to ignore at least one point to calculate the fit, otherwise the planet (if a planet is present) would be significantly subtracted. It is important to realize that we have no prior knowledge of the location of a planet (if present) and therefore we do not know a priori which point to ignore. The subtraction algorithm has to determine by itself if it should exclude a point, and if so, which one. This is accomplished iteratively by looking at the residual spectrum of the pixel after subtraction of a fit on all points. The situation is different for a methanated companion, essentially present in a single (or a few) channel (~1.58 μm); in this case it is fine to systematically ignore this channel (or a few) when calculating all fits.

Table 1 lists the different algorithms investigated with various options.

Subtraction type	Option	explanation
SD, DD & DDD	spacing=*N* (default *N*=1)	Subtraction made with images separated by *N* channels
Fit, degree 1 or 2	ignore=*c1*, *c2*, …	Ignore systematically channels *c1*, *c2*, ... for fit
Fit, degree 1 or 2	ignore=*N*	Ignores at least one and up to *N* channels for fit: iterative algorithm that determines which channels to ignore
Fit, degree 1 or 2	robust	Like ignore=*N*, but in this case $N=2R/r$ is function of separation

Table 1: PSF subtraction algorithms

4.4 MWI Case

4.4.1 Speckle suppression

Figure 4.2 shows the mean attenuation for one image of the MWI data cube and different subtraction algorithms. These curves corresponds to the maximum attenuation that can be achieved using the different algorithms and are set by the chromatic evolution of the PSF. In the presence of noise, the real attenuation will be set by that noise (up to the limiting curves shown). This figure shows that a DD, a DDD and a polynomial fits of degree 1 and 2 can all provide attenuations below 0.01. For polynomial fits of degree 1 ignoring 2 points and of degree 2 ignoring 1 point, the attenuation is calculated in the channel(s) ignored, since in the absence of noise the attenuation is infinite in the other channel(s).

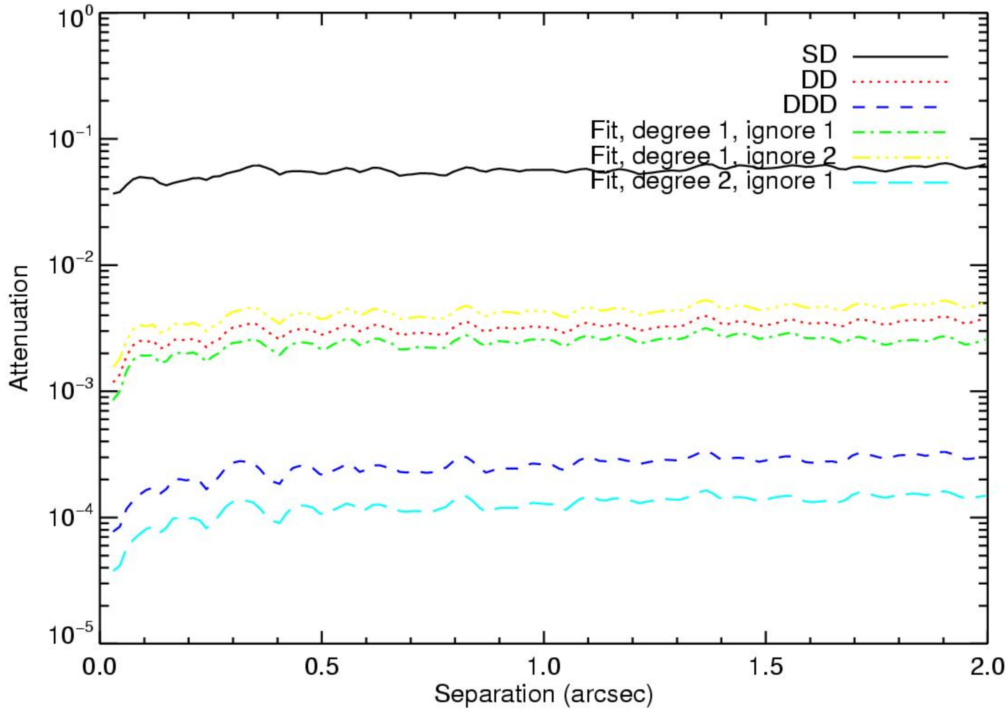

Figure 4.2: Mean speckles attenuation in one image of the MWI data cube for different subtraction algorithms.

4.4.2 Companion recovery efficiency

To determine the fractional signal remaining after the subtraction of the speckles using different algorithms, virtual planets were implemented in the original data cube and the subtraction was carried through. The residual cube was then scaled back to its original scale and collapsed over the spectral dimension to produce a “broadband” image. The signal recovered for each planet inside an aperture of diameter λ/D was compared to that of the implanted planet. The ratio of these two quantities is the recovery efficiency. The exercise was repeated for virtual planets having a flat spectrum and for planets having the spectrum of a T8 dwarf. For planets with a T8 dwarf spectrum, the cube was not collapsed but only the channel at 1.58 μm was retained, because collapsing the entire cube results in a lower signal-to-noise.

Figure 4.3 shows the recovery efficiency for different algorithms, for companions with flat spectra, and the effective speckle attenuation for each algorithm, defined as the attenuation divided by the recovery efficiency. All algorithms maintain an effective attenuation below 0.01 down to ~0.5”, while the 2nd degree polynomial fit ignoring 1 point reaches closer in, but provided that noise is at a level of ~10^{-4} at short separations.

Figure 4.4 shows the corresponding curves for companions with T8 dwarf spectra. The benefit of such a spectrum is clear as the recovery efficiency remains close to one at all separations and all algorithms maintain attenuations below 0.01 at all separations.

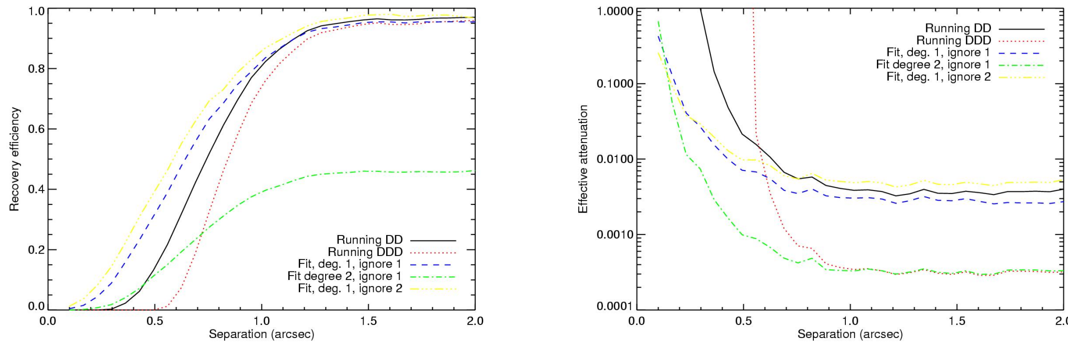

Figure 4.3: (*left*) Companion recovery efficiency for a companion with a flat spectrum (*right*) Effective speckle attenuation.

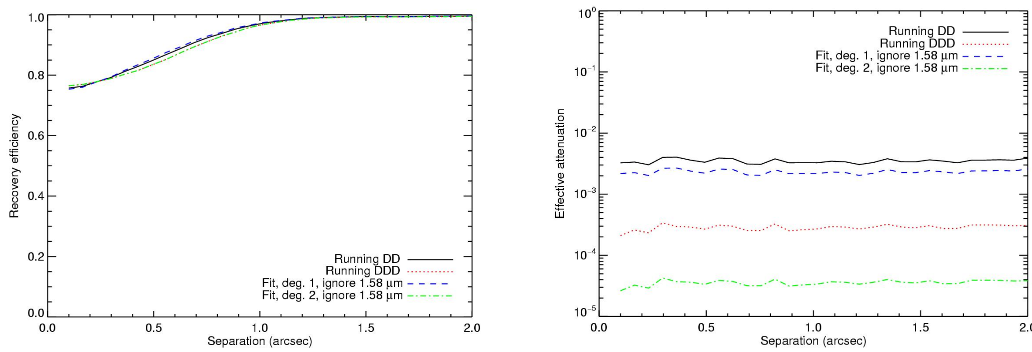

Figure 4.4: (*left*) Companion recovery efficiency for a companion with a T8 dwarf spectrum (*right*) Effective speckle attenuation.

4.5 IFU Case

4.5.1 Speckle suppression

Figure 4.5 shows the mean attenuation for one image of the IFU data cube and different subtraction algorithms, all of which permit attenuations below 0.01.

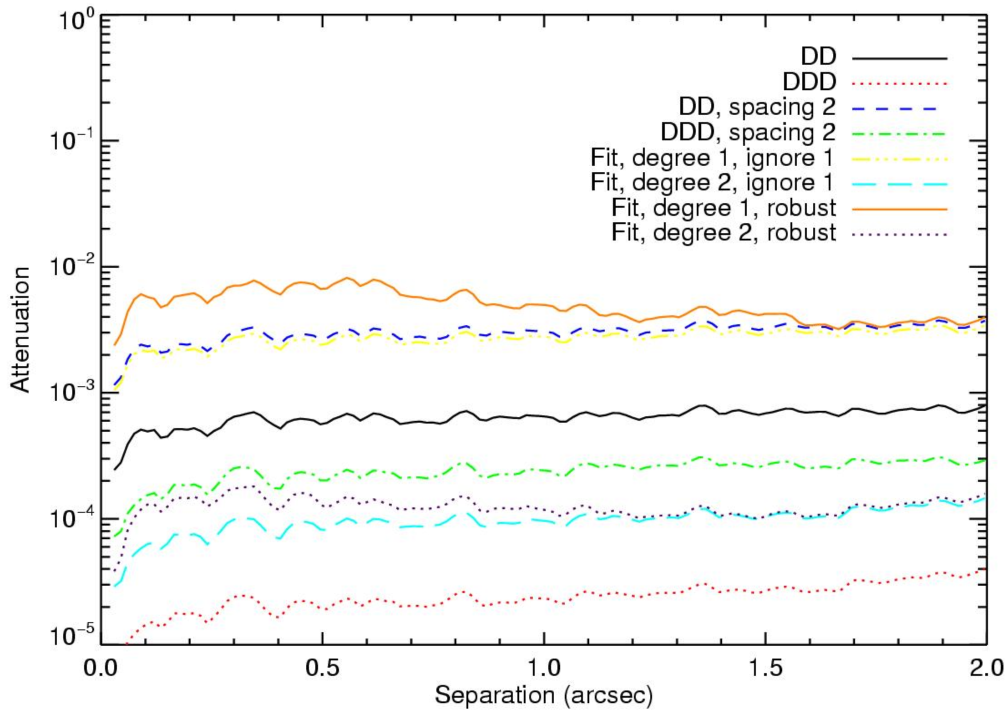

Figure 4.5: Mean speckles attenuation in one image of the IFU data cube for different subtraction algorithms.

4.5.2 Recovery efficiency

The procedure is the same as in section 4.2. Here, for planets with a T8 dwarf spectrum, the cube was collapsed only over the three channels in which the planet is bright, collapsing the entire cube results in a lower signal-to-noise.

Figure 4.6 shows the recovery efficiency for different algorithms, for companions with flat spectra, and the effective speckle attenuation for each algorithm. The polynomial fit of degree 2 maintains an effective attenuation below 0.01 down to ~0.2" (~5 λ/D), while most other algorithms maintain attenuations below 0.01 down to ~0.5".

Figure 4.7 shows the corresponding curves for companions with T8 dwarf spectra. Again, all algorithms maintain attenuations below 0.01 at all separations, and polynomial fits that ignore three channels centered on the peak of emission have good recovery efficiency. It has also been verified that the "robust" polynomial fits give the same results as the fits ignoring the three channels centered on the peak of emission.

In a general case, a "robust" polynomial fit of degree two is the most attractive.

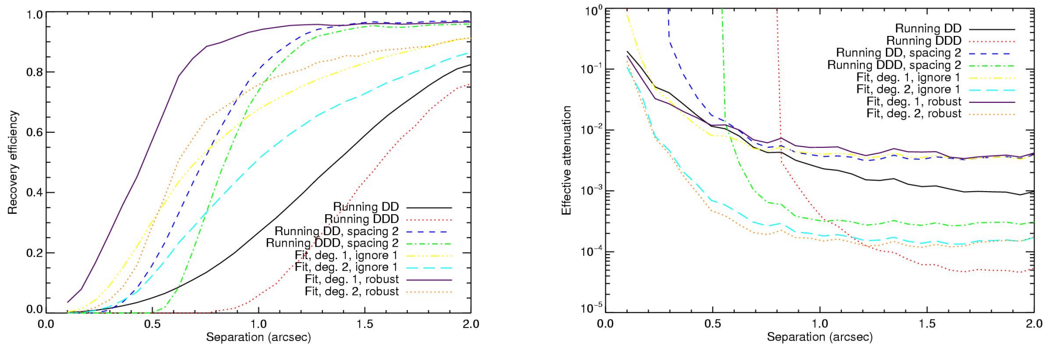

Figure 4.6: (*left*) Companion recovery efficiency for a companion with a flat spectrum (*right*) Effective speckle attenuation.

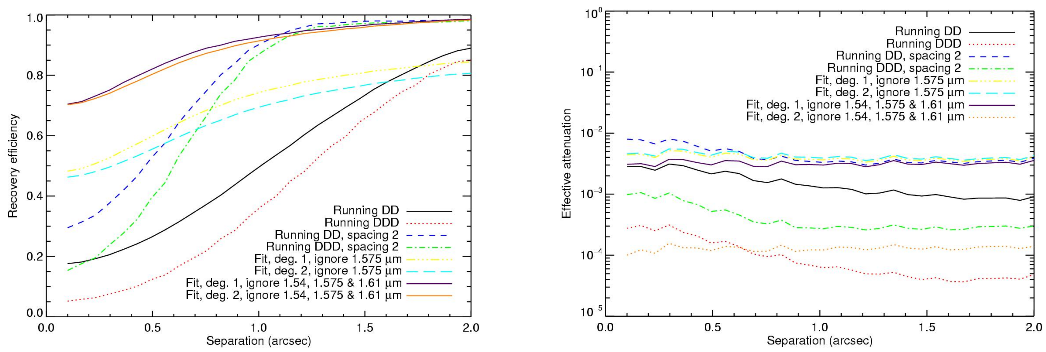

Figure 4.7: (*left*) Companion recovery efficiency for a companion with a T8 dwarf spectrum (*right*) Effective speckle attenuation.

4.6 Image examples

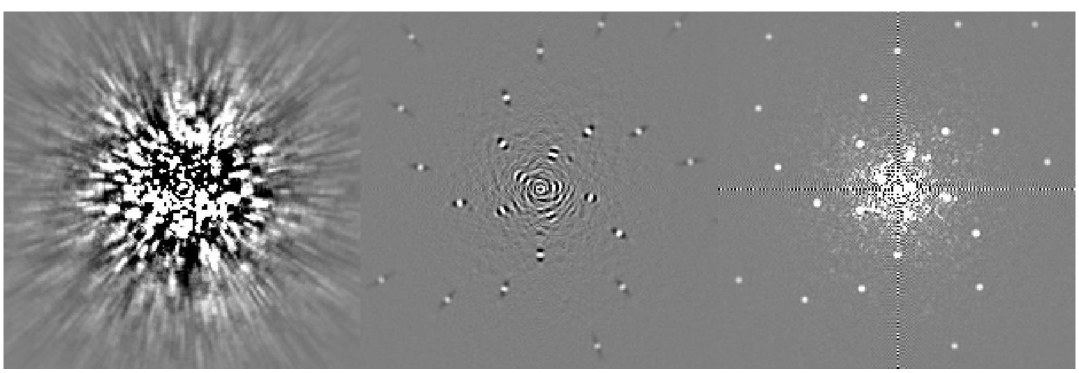

Figure 4.8: (*left*) Collapsed cube including virtual planets with flat spectra, an azimuthally averaged profile has been subtracted (*middle*) Residual collapsed cube after subtraction with polynomial fit of degree 1 including all points, the negative signal is due to biases in the fit introduced by the presence of the planets (*right*) Residual collapsed cube after subtraction with a "robust" polynomial fit of degree 1, the negative signal has disappeared and the core of the planets are brighter. Images are 3" on a side. Stretch of image on the left is 10 times that of the other two images.

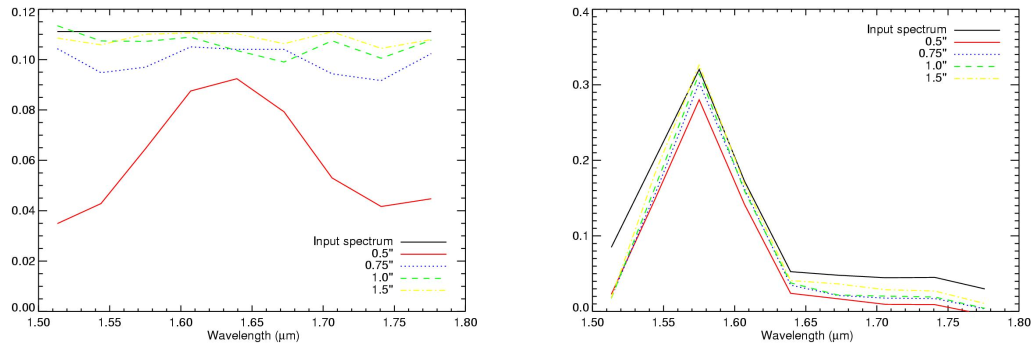

Figure 4.9: Recovered spectra compared with input spectra for the IFU at various separation using the polynomial fit of degree 1 (*left*) flat spectrum (*right*) T8 spectrum.

4.7 Summary

The ExAOC performances will dictate which algorithms to use since they will determine the level of photon noise (or some other limiting noise) compared to that of speckle noise. For a real system, it is this limiting noise curve that has to be divided by the recovery efficiency. Assuming that photon noise will be at about 10^{-2} times the speckle noise or less, then all algorithms are able to reach the photon noise. **In this case it is the recovery efficiency alone that determines the best algorithm to use: this algorithm is the polynomial fit of degree 1 ignoring two points in the MWI case and its robust version in the IFU case.** The figure below shows the effective attenuations of these algorithms with the MWI and IFU in a case in which the photon noise is 10^{-2} times the speckle noise, for flat and methanated spectra.

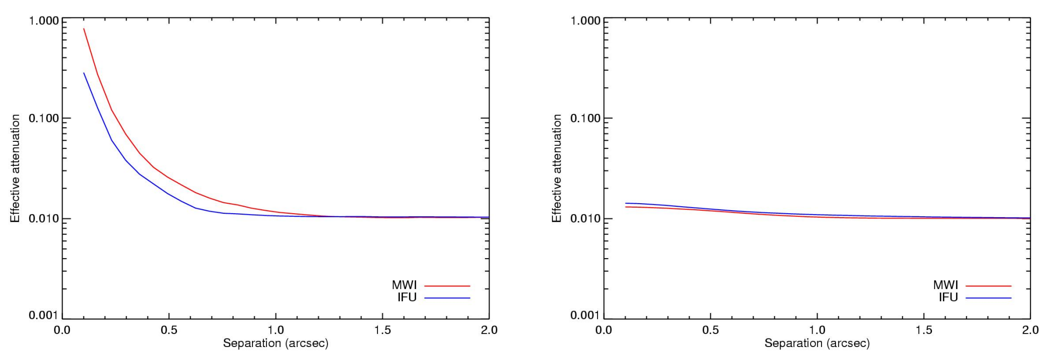

Figure 4.10: Effective attenuation in the presence of photon noise at a level of 10^{-2} times the level of speckle noise using a polynomial fit of degree 1 ignoring 2 channels (MWI case, flat spectrum) and ignoring the 1.58 µm channel (MWI case, T8 spectrum) and a robust polynomial fit of degree 1 (IFU case) (*left*) Flat spectrum (*right*) T8 spectrum

It would be interesting to repeat this exercise with PSF that resembles more the PSFs that will be delivered by the ExAOC to make sure that the algorithms perform as well. It would also be useful to get an estimate of the level of photon noise compared to that of speckle noise for a typical one-hour exposure. It would also be worthwhile to repeat the exercise for an IFU with broader wavelength coverage to see the effect on the recovery efficiency for flat spectrum planets.