Fast cadence speckle-free high-contrast imaging: SFADI and SFI

G. Li Causi, M. Stangalini, S. Antoniucci, F. Pedichini, M. Mattioli, et al.

G. Li Causi, M. Stangalini, S. Antoniucci, F. Pedichini, M. Mattioli, V. Testa, R. Piazzesi, "Fast cadence speckle-free high-contrast imaging: SFADI and SFI," Proc. SPIE 10703, Adaptive Optics Systems VI, 107032U (10 July 2018); doi: 10.1117/12.2312614

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2018, Austin, Texas, United States

Fast cadence speckle-free high-contrast imaging: SFADI and SFI

G. Li Causi[a,b,c], M. Stangalini[b,c], S. Antoniucci[b,c], F. Pedichini[b,c],
M. Mattioli[b], V. Testa[b], R. Piazzesi[b]

[a]INAF-Istituto di Astroficia e Planetologia Spaziali, Via Fosso del Cavaliere, 100, Roma, Italia;
[b]INAF-Osservatorio Astronomico di Roma, Via Frascati 33, Monteporzio Catone, Roma, Italia;
[c]ADONI Adaptive Optics National Lab of Italy

ABSTRACT

We present the R&D status of the Speckle-Free Angular Differential Imaging method (SFADI), that we developed for the SHARK-VIS high-contrast imager for the LBT telescope.

The technique bases on the acquisition of kHz frame-rate image sequences, which we combine in post-processing after speckle identification and suppression in each frame. With respect to the standard angular differential imaging, this method reaches a much smoother residual background and hence higher detection contrast at a given signal-to-noise ratio. Furthermore, it can reveal faint extended sources around bright central stars, and can use de-rotated images as well as quick second-lasting sequences.

We reached a contrast of around 10^{-5} for integration times of the order of tens of minutes at 100 mas for a 5.7^{th} magnitude star, as we demonstrated on both a real-sky acquisition and at the SHARK-VIS laboratory test bench. Such long sequences though produces a large amount of data (around a million frames every 15 minutes) that we manage to processed in a reasonable computation time with the described implementation scheme.

Keywords: Adaptive optics, High-contrast imaging, Speckle-free angular differential imaging, Exoplanets imaging, SHARK-VIS, Large Binocular Telescope

1. INTRODUCTION

A large number of scientific topics in astronomy would benefit from the employ of "high-contrast imaging" techniques, i.e. the capability to image faint objects around bright compact sources. This is the case for instance of exoplanets direct imaging (e.g. [8], [29], [24], [7]), protostellar disk observations (e.g. [14], [13], [3]), and analysis of jets from stellar sources (e.g. [2], [25]). However, this is not a simple task, due to the large fraction of light spread from the central bright source into the surrounding area by the combination of atmospheric turbulence, optical aberrations, and diffraction. In fact, even with the use of the most advanced extreme adaptive optics (EXAO)[23], the intensity of such spread light remains much higher than the faint objects we want to observe.

This is why a number of post-processing techniques has been proposed in the literature, with the aim of modeling and subtracting this background disturbance, like e.g. the angular differential imaging (ADI)[20], the principal component analysis (PCA)[1], the locally optimized combination of images (LOCI)[16], and many others. However, the concept of these methods, namely to create a model of the spread light intensity distribution from a multiple EXAO images acquisition and to reveal the faint object by subtracting this model from each frame, cannot improve the contrast beyond a certain level. This happens because the spread of light is non-stationary (especially in the visible bands), being highly variable both on small spatial scale and short temporal intervals of a few mas and milliseconds respectively, as we described in [28].

To improve this contrast limit, we recently proposed a different approach, the Speckle-Free Angular Differential Imaging (SFADI)[17], which bases on the EXAO acquisition of kHz frame-rate image sequences, which we combine in post-processing after speckle identification and suppression in each frame.

Adaptive Optics Systems VI, edited by Laird M. Close, Laura Schreiber,
Dirk Schmidt, Proc. of SPIE Vol. 10703, 107032U · © 2018 SPIE
CCC code: 0277-786X/18/$18 · doi: 10.1117/12.2312614

Here we describe some details of the SFADI implementation, and the tests that we performed with both a real-sky acquisition and with the optical test bench of the SHARK-VIS high contrast imager that we are building for the LBT telescope[21].

2. SFADI HIGH-CONTRAST IMAGING OF POINT SOURCES

The typical 8-meter class telescope post-EXAO point spread function (PSF) is a rapidly variable shape, consisting of a narrow central peak, encircling the majority of the source light within a radius close to the ideal λ/D of the Airy diffraction image, surrounded by a large number of fainter residual speckles of the same size (Figure 1, left).

We measured the evolution time-scale of such instantaneous PSF at the LBT telescope in the R band, by acquiring a 20 min image sequence at one kHz frame rate (SHARK-VIS Forerunner experiment[22]), during which the seeing varies in the range of 0.8÷1.5 arcsec FWHM. The measured time scales range from the typical lifetimes of the 2 to 5 ms of the AO residual speckles up to the ~70 ms decorrelation time, after which the whole speckles pattern completely renews, as we described in [23]. More, to this intrinsic shape variability, the wind-shake of the telescope adds a ~10÷15 mas jitter, with a 99% power contribution below a frequency of 300 Hz, as measured on this sequence, which rigidly shift the image from frame to frame. The correct acquisition scheme for such evolving object is thus a ~1 kHz frame sequence of less than λ/D sampled images, able to both freeze the seeing and jitter movements, and to resolve the speckles crowding.

The advantages of a fast acquisition like this are multiple. First, it allows a post-facto jitter compensation based on Gaussian fit location of the PSF core, which maximizes the peak contrast of a possible planet. Secondly, the speckles are frozen, leaving large areas of pure sky background, free from any light spread from the central source, at least at the intensity level of one digital count (Figure 1, left). More, it allows a measurement of the frame-by-frame flux variation in the central peak, which we use to renormalize the frames and compute the correct weights for their combination. At last, we can perform a frames selection by instantaneously measuring the shape of the PSF core, in order to reject those frames with a double, elongated, or defocused peak, which amounts to ~10% in the bad-seeing conditions of the Forerunner sequence.

The SFADI method, that we proposed in [17], builds on these useful properties of the fast cadence acquisition scheme, implementing the concept of gathering the planet's flux only when it is not contaminated by the speckle's light, as illustrated in Figure 1.

In fact, the SFADI works by recognizing all the speckles in each re-aligned selected frame, by means of an algorithm (SWAMIS) [11], which performs feature recognition by joining spatial shape information with temporal shape evolution across time-adjacent frames. The SFADI process then suppresses all the identified speckles, by masking the involved pixels in each frame. If the frame sequence is acquired with no tracking of field-rotation (which is the baseline scheme in high-contrast imaging), the unavoidable faint residual of light spread (due to the speckles not recognizable in the single frame) can be retrieved by a median stacking of the masked frames (Figure 2, left panel) and subtracted from each of them. Finally, these masked and subtracted frames are digitally de-rotated and combined by a weighted average in order to obtain the final high-contrast result.

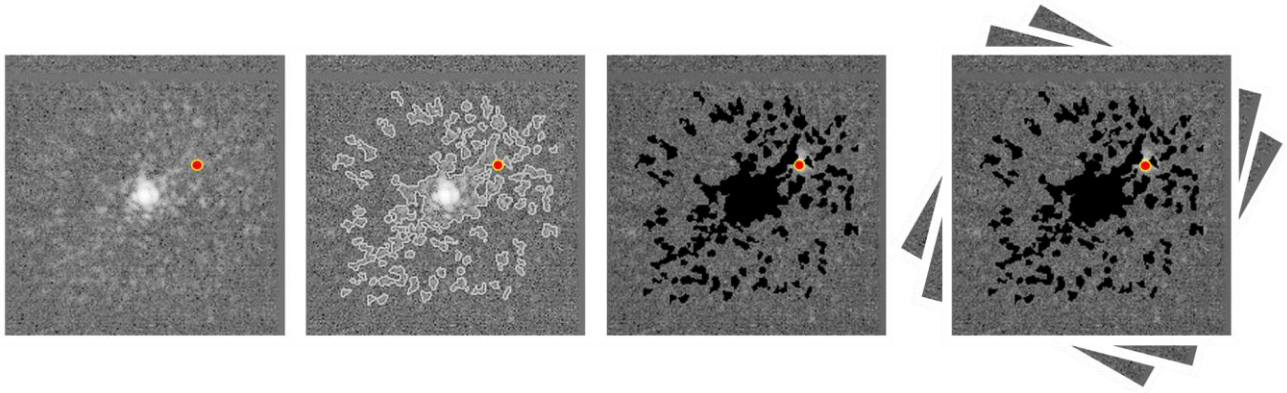

Figure 1: Illustration of the SFADI concept: planet's photons in each frame are collected only when no speckle covers planet's position (red spot). To obtain this, speckles are first identified and suppressed in each frame, then the speckle-suppressed frames are de-rotated and combined.

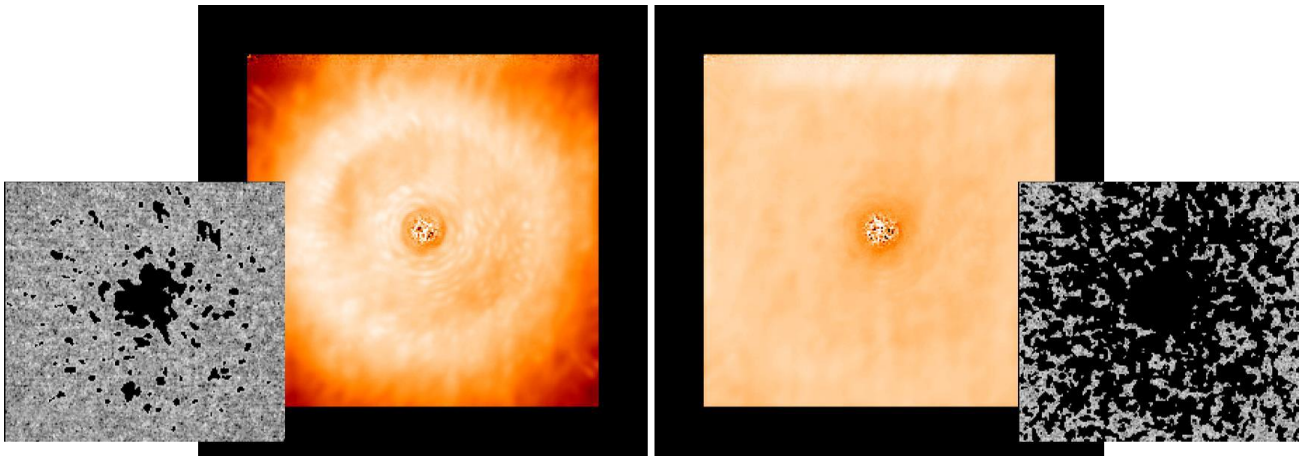

Figure 2: Median stacking of speckle-masked frames with a SWAMIS threshold of 3 (left), and 1.5 (right). Corresponding masked areas in a single frame are shown in the small pictures.

The described speckle suppression is of course only suitable for detecting faint planets, with a peak intensity lower than the adopted SWAMIS threshold, and not for bright ones, that could be recognized and self-masked. In the Forerunner this limit corresponds to a contrast of ~10^{-3}, planets fainter than this are not detectable in the single 1 ms exposures, and thus are not masked. We have found anyway, an intensity attenuation for such faint planets, even if not directly recognizable by the SWAMIS code, which is due to a subtle effect that we describe hereafter in APPENDIX 1.

It's worth noting that for SFADI we do not apply any frames selection based on peak flux, as commonly done in classical lucky imaging. In fact, when peak flux decreases speckles flux increases, so that more speckles are recognized and suppressed, finally resulting in a Signal to Noise Ratio (SNR) always increasing with the number of frames, independently on seeing variation along the acquisition, as we shown in [17]. Indeed, we can consider the SFADI like a sort of "spatial Lucky imaging", for which the classical Lucky imaging rule of thumb that "a very small number of very best frames is always better" does not hold.

In the final SFADI result each pixel has a different noise, because masking pixels also means decreasing the effective number of contributing frames, and hence the effective integration time is different for each pixel. Nonetheless the resulting image is characterized by a very low and smooth residual background, and in the case of the Forerunner dataset (5.7 mag central star and ~70° field rotation) it allows to clearly image synthetic planets as faint as $2 \cdot 10^{-5}$ at 100 mas

from the central star and detect 10^{-5} ones, while 10^{-4} planets are unambiguously detected as close as 50 mas (Figure 3, right panel). Comparison with classical ADI (Figure 3, left panel) shows that the SFADI has a more uniform background and less residual, especially at close distances from the star, and a higher SNR at same planet contrast.

It's worth noting that we injected synthetic planets as integer pixel values, following the un-biased algorithm described in APPENDIX 2, in order to realistically represent the case of a real observation containing true planets.

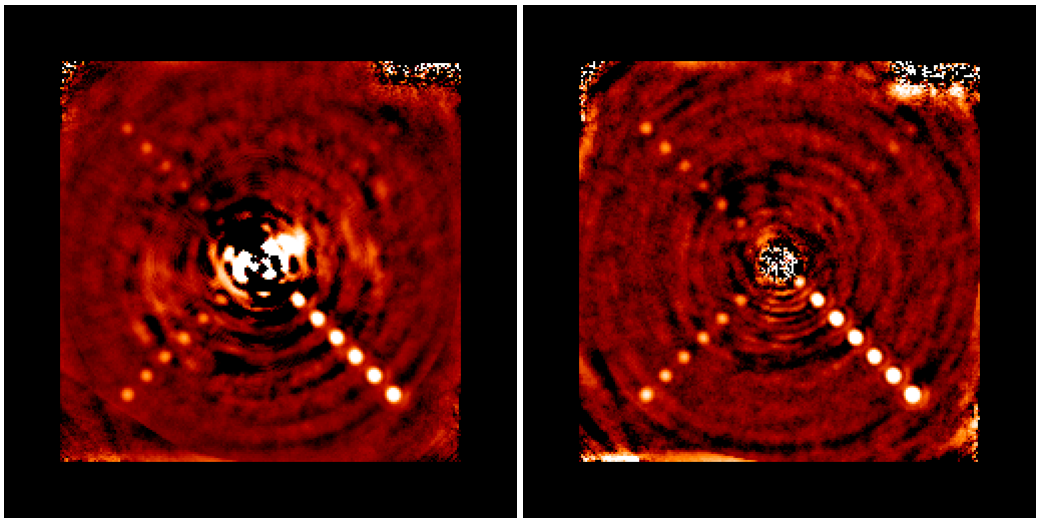

Figure 3: Results from classical ADI (left) compared to SFADI (right) on the described frames selection out of the $1.2 \cdot 10^{6}$ frames sequence of the Forerunner experiment. Injected planets are 10^{-5}, $2 \cdot 10^{-5}$, $5 \cdot 10^{-5}$, 10^{-4} contrast in the four quadrants, respectively.

We can judge the efficacy of the speckle suppression by directly inspecting the variation of the median stack upon variation of the SWAMIS identification threshold. As we see in right panel of Figure 2, a hard threshold of 1.5 ADU, which masks around 50% of the frames across the 100 to 200 mas radii, does not leave any trace of the PSF shape, when applied to the Forerunner dataset. On the contrary, the milder threshold of 3 ADU, used for left panel of Figure 2, masks less than 30% of the frames in the same region, and still leaves a clear fingerprint of the average PSF. A low threshold hence could seem to be the best choice, but this is not true because reducing the number of involved frames rises the image noise correspondingly. Hence, there is always an optimal SWAMIS threshold, which produces the best compromise between the reduction of speckle residuals and the increase of detector noise, as shown in Figure 4.

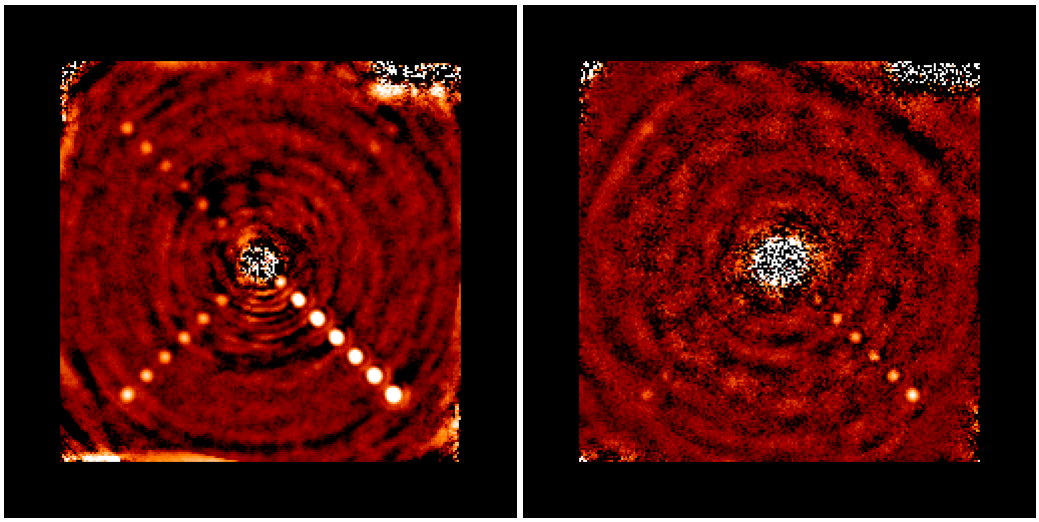

Figure 4: SFADI result using a SWAMIS masking threshold of 3 (left) and 1.5 (right), corresponding to left and right panels of Figure 2.

Maximum contrast in SFADI is thus ultimately limited from detector noise, given that we could almost completely remove the speckles, as right panel of Figure 2 demonstrates. To verify this sentence, we applied the SFADI algorithm to a one second sequence of kHz frame simulations of the upcoming SOUL adaptive optics system for LBT[23], computed by the SHARK Simulation Code (SSC)[27] comparing the acquisition of the Zyla 5.3 air cooled CMOS camera and the acquisition of a photon counting First Light Ocam2k EMCCD. The EMCCD adopt a very high gain (of the order of one thousand) to work as a photon-counting device with an equivalent zero read out noise. Figure 5 depicts the simulation results for a set of 10^{-3} and 10^{-4} planets at different distances, the latter almost undetectable with the CMOS, but clearly visible with the EMCCD.

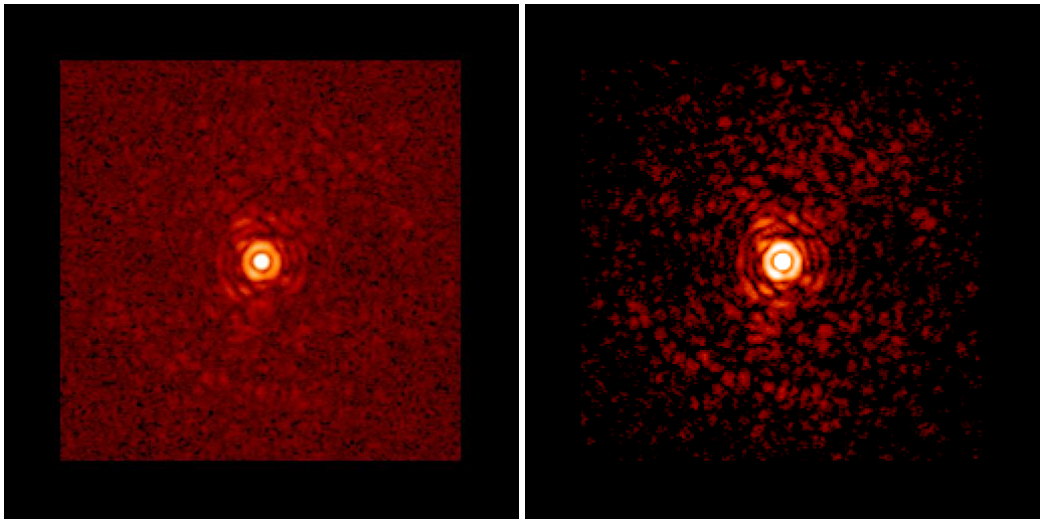

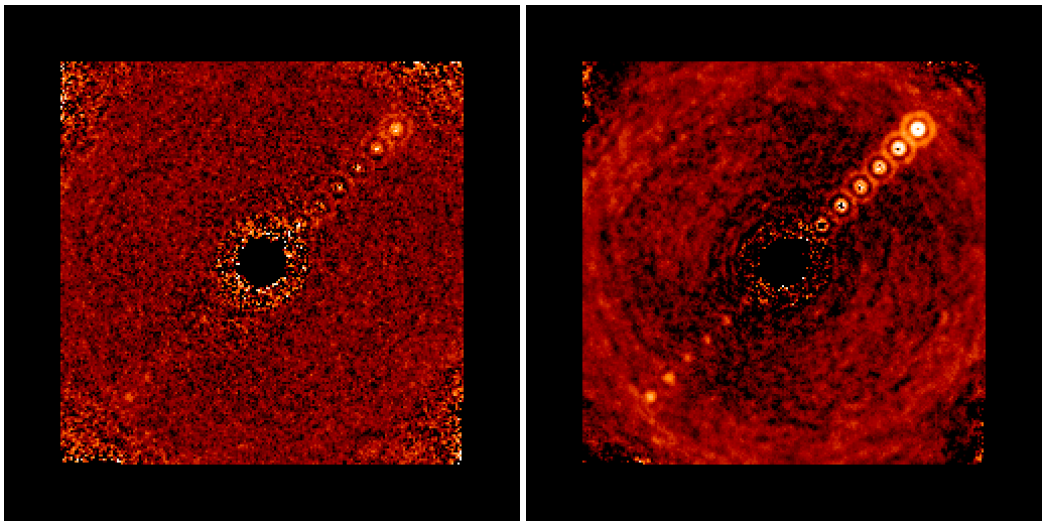

Figure 5: One of the 1000 simulated frames (top) and the corresponding SDADI results (bottom) with CMOS (left) and EMCCD (right) detectors, for a set of 10^{-3} (first quadrant) and 10^{-4} (third quadrant) contrast planets.

3. SFI HIGH-CONTRAST IMAGING OF EXTENDED SOURCES

If the faint source under study is not a point source, but extends for a size comparable to the arc subtended by field rotation at its position, we cannot use any angular differential methods, like ADI and SFADI. In fact, the median background model of Figure 2 would contain an angularly smoothed ghost of the source itself, thus leading to self-subtraction in the final combination.

In the SFADI process though, the step of median background subtraction is always a minimal correction, even negligible if the identification threshold is sufficiently low as in right panel of Figure 2. This means that we can conveniently neglect this step when we are interested in extended sources, thus removing the "angular differential" part of the process, which becomes a direct Speckle-Free Imaging (SFI). Should the extended flux be enough to be visible in the single frame, a local median filtering, with a size larger than λ/D, can be applied to remove the background for a better speckle recognition by the SWAMIS algorithm.

Figure 6 depicts the high-contrast imaging performance of the SFI process for extended sources. We injected the same flux-scaled Z CMa model that we used in [17] (derived from deconvolution of SPHERE/ZIMPOL observations[1]) into a subset of 5000 frames equispaced across the Forerunner sequence. Then we applied the SFI process, with the same frames calibrations described in Sec. 2, and obtained a result very similar to the ideal image, correctly retrieving the wide outflow from the primary and the collimated jet from the companion, whereas both the simple average and the ADI algorithm cannot reveal them without artifacts (Figure 6).

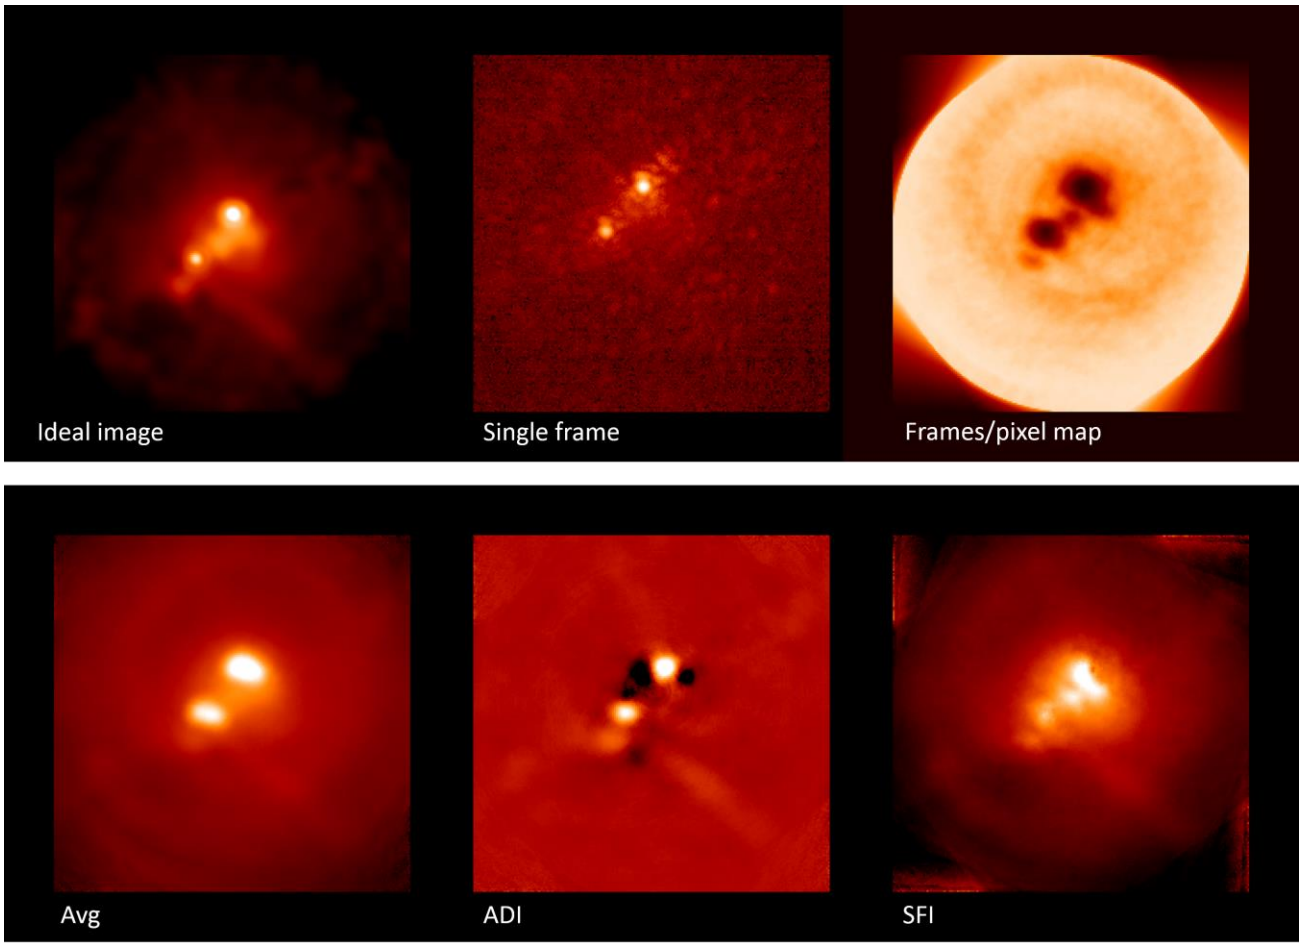

Figure 6: Extended circumstellar structure of the Z CMa binary scaled to have a $5 \cdot 10^{-2}$ contrast (within in a λ/D on the jet) w.r.t. the brighter star. Top-left: ideal image, i.e. convolved with average Forerunner PSF; top-middle: single frame with real speckle pattern from the Forerunner sequence; top-right: map of the effective frames per pixel in the final result; bottom-right: final SFI result, to be compared with the ideal image. Simple average and ADI are shown in left and middle bottom panels, for comparison.

4. FAST CADENCE IMAGING AT THE SHARK-VIS TEST BENCH

In order to evaluate the real contrast limits of the fast acquisition concept, we used the optical test bench of the SHARK-VIS high contrast camera that we are building for the LBT telescope at the Rome Astronomical Observatory[21]. This optical bench physically creates on a focal plane a tunable bright point source and a faint fixed contrast companion moving along a circle centered on the bright one, reproducing both the numerical aperture, the image scale, the fluxes, and the timings of a real on-sky observation with the Forerunner experiment, or with the new SHARK-VIS instrument. At present, speckles are not yet introduced in the optical path, and we only tested the optimal procedures for detector bias subtraction and noise limitation.

Thanks to these tests, we discovered that a simple subtraction of a "master bias", computed as the median, or average, of a dark frames sequence, is not enough to remove the bias contribution at a sufficient level. In fact, detector bias voltage presents low frequency temporal variations, for which the bias pattern at acquisition time is different from the bias pattern at time of dark acquisitions. Though small, these differences propagate through frames de-rotation and stacking, producing a series of dark and bright arc-shaped residuals in the final high-contrast image (see Figure 7, left panels).

The way that we found to successfully avoid this problem is to use a field stop which physically masks the upper and lower rows of the region read from detector, so that the instantaneous bias pattern can be measured, and subtracted, from each frame, resulting in an artifact-free stacked image (Figure 7, right panels).

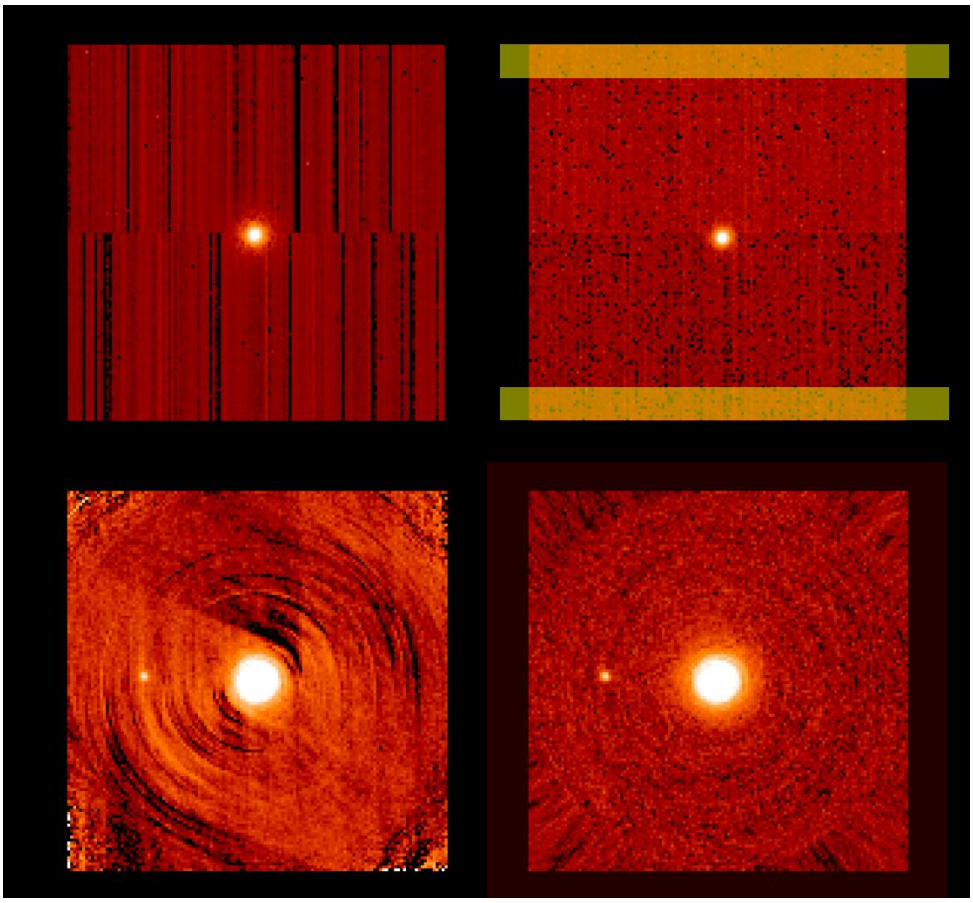

Figure 7: Lab tests for bias subtraction. Left: residuals left by master bias subtraction on direct (top) and de-rotated (bottom) frames stacking. Right: artifact-free de-rotated stack (bottom) using ten top and bottom physically masked rows (yellow stripes in top panel).

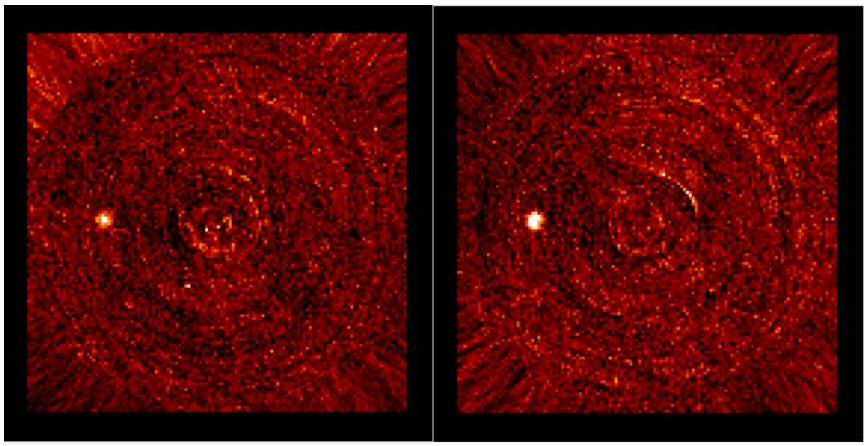

Figure 8: De-rotated frames stacking for the laboratory tests in fast (left) and slow (right) acquisition mode of the Zyla camera. A $5 \cdot 10^{-5}$ contrast planet is acquired in these 1 min sequences, where the central star has been blocked to only measure the planet to detector noise SNR, which results to be 7.0 for the fast mode and 10.5 for the slow mode.

Finally, we tested the sensitivity of the Zyla 5.3 air cooled CMOS camera with different readout velocities, and we found that the slow reading mode provides a 50% gain in planet's SNR. The drawback is that the minimum frame exposure in slow readout mode rises to ~1.5 ms, but still it is faster than the shortest speckle lifetimes that we measured in the Forerunner data, so that we expect to see a similar gain at the sky. This also demonstrates that the detector-only contrast limit, in slow mode, corresponds to a SNR > 10 detection of a $5 \cdot 10^{-5}$ planet around a 5.5 mag artificial star in 1 min exposure, considering the Forerunner optics that we reproduced in this test (Figure 8).

We plan to add a movable phase screen in the optical path, which will physically reproduce the speckles and their variations with the same spatial and temporal properties of the upcoming SOUL adaptive optics system[23], which will replace the current FLAO system[12] used in the Forerunner experiment. With such device, we will test the whole SFADI process in a controlled way, and quantify the effectiveness of the reference rows bias subtraction and slow readout schemes in the real case of crowded speckles.

5. FAST CADENCE IMAGING IMPLEMENTATION

Performing a kHz rate observation of faint sources could mean working with millions or tens of millions ($\sim 4 \cdot 10^6$ frames per hour) of images, which require large disk storage. In fact, the majority of high contrast algorithms, and in particular the ones we have currently under study, namely SFADI[17], Recurrence Quantification Analysis (RQA)[26], Multi-Frame Blind Deconvolution (MFBD), and Incremental Principal Component Analysis (iPCA), do not work on the single frames but on the whole dataset, or on large chunks of frames taken together. This means that it is not possible to process the data in real-time, using each acquired frame on the fly without storing them, though this can be done in SFI mode for a rough quick look with no calibration.

Moreover, the acquired frames sequence, usually coded as 16-bits unsigned integer arrays larger than ~200×200 pixels, typically needs to be cleaned and co-registered, resulting in a 32-bits floating point dataset which occupies up to twice the original data size, and for SFADI a parallel mask sequence is also produced for each chosen threshold. Due to the huge computation time that each of these processes needs and to the required parameter optimization, it is not convenient to delete these derived datasets once used, so that a large storage is mandatory.

In order to run the SFADI process on million frames in a reasonable time, we use data chunking and multi-threading in the SHARK-VIS prototype pipeline, written in IDL language and run on a 12-CPU Intel Core i7 @3.3GHz with 64Gb RAM, which also stores the acquired frames in data cubes of 5000 frames each. Two ways of chunking can be used to segment the dataset, which we call "sequence chunking" and "frame chunking". In the first case a given number of frames, possibly located in different data cubes, are processed as a 3D array and a running partial result is continuously accumulated (Figure 9, left), a method possible for calibrations, averages, histograms, and speckles masking. Other

calculations, like e.g. the median stacking shown in Figure 2, needs all the frames together, but they can be larger than the available memory, so that we segmented the dataset vertically by considering a piece of the frame at each time (see Figure 9, right). On a multi-CPU computer different chunks can be processed in parallel, thus contemporary using 100% CPU resources and 100% RAM allocation thus getting the maximum available speed. With such scheme, we currently obtain around one-day per million frames full process time (~2 hours for peak analysis, ~7 hours for calibration/alignment plus planet injection, ~17 hours for SWAMIS masks, ~1.5 hours for SFADI combination).

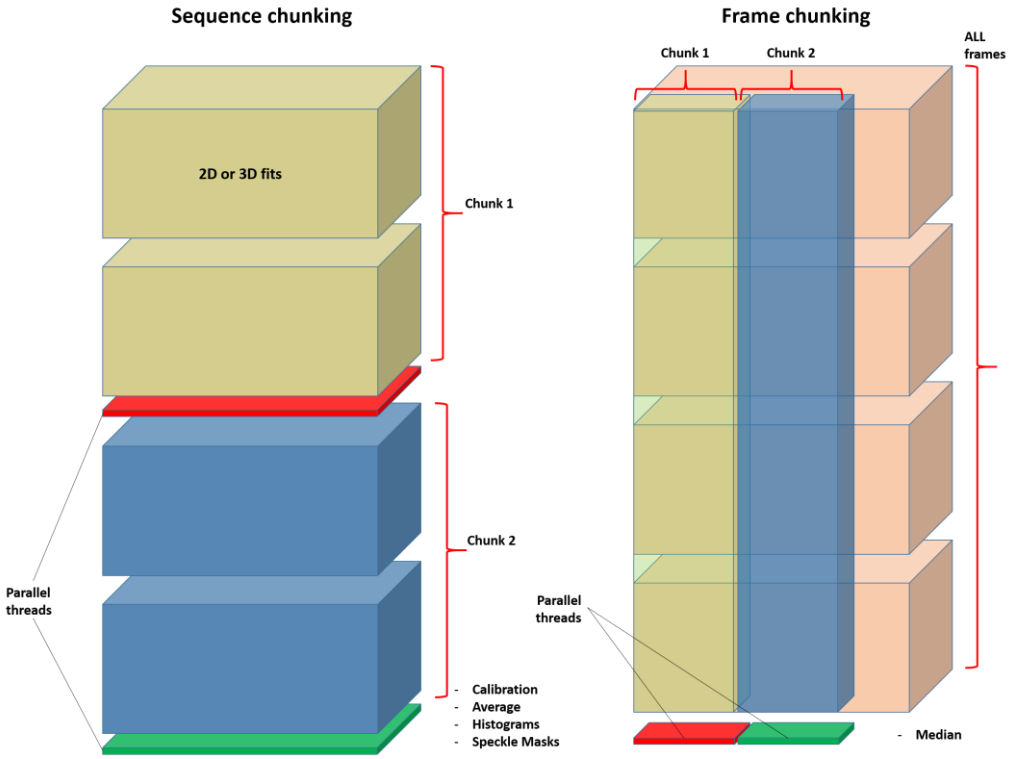

Figure 9: The two ways of segmenting the frames sequence (time is vertical): in "sequence chunking" (left) a large number of chosen frames are read from the frames (or frame-cubes) and processed together, and a running partial result is continuously updated. In "frame chunking" (right) the same sub-region is read from all the frames (or frame-cubes) and processed, and the partial results are then tiled to form the final image.

REFERENCES

[1] Amara, A., & Quanz, S. P. 2012, MNRAS, 427, 948
[2] Antoniucci, S., Podio, L., Nisini, B., et al. 2016, A&A, 593, L13
[3] Avenhaus et al. 2018 in press (arXiv180310882A)
[4] Bailey, V. P., Hinz, P. M., Puglisi, A. T., et al. 2014, Proc. SPIE, 9148, 914803
[5] Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2006, Msngr, 125, 29
[6] Cavarroc, C., Boccaletti, A., Baudoz, P., Fusco, T., & Rouan, D. 2006, A&A, 447, 397
[7] Chauvin et al. 2017 A&A, 605L, 9C
[8] Close et al. 2014 ApJ, 781, L30
[9] Close, L. M., Follette, K., Males, J. R., et al. 2014, in IAU Symp. 299, Exploring the Formation and Evolution of Planetary Systems, ed. M. Booth, B. C. Matthews, & J. R. Graham (Cambridge: Cambridge Univ. Press), 32
[10] Davies, R., & Kasper, M. 2012, ARA&A, 50, 305
[11] DeForest, C. E., Hagenaar, H. J., Lamb, D. A., Parnell, C. E., & Welsch, B. T., 2007, ApJ, 666, 576
[12] Esposito, S., Riccardi, A., Fini, L., et al. 2010, Proc. SPIE, 7736, 773609

[13] Garufi et al. 2016 A&A, 588, A8
[14] Ginski et al. 2016 A&A, 595A, 112G
[15] Kasper, M. 2012, Proc. SPIE, 8447, 84470B
[16] Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É 2007, ApJ, 660, 770
[17] Li Causi, G., Stangalini, M., Antoniucci, et al. 2017, ApJ, 849, 85
[18] Macintosh, B., Graham, J. R., Ingraham, P., et al. 2014, PNAS, 111, 12661
[19] Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, Proc. SPIE, 7015, 701518
[20] Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556
[21] Mattioli, M., Pedichini F., Antoniucci, S., et al 2018, Proc. SPIE, 10702, 10702-160
[22] Pedichini, F., Stangalini, M., Ambrosino, A., et al. 2017, AJ, 154, 74
[23] Pinna, E., Esposito, S., Hink, P., et al. 2016, Proc. SPIE, 9909, 99093V
[24] Sallum et al. 2015 A&A Nature, 527, 342S
[25] Schmid et al. A&A, 602A, 53S
[26] Stangalini, M., Li Causi, G., Pedichini, F., et al. 2018, submitted to ApJ
[27] Stangalini, M., Pedichini, F., Centrone, M., et al. 2014, Proc. SPIE, 9147, 91478F
[28] Stangalini, M., Pedichini, F., Pinna, E., et al. 2017, JATIS, 3, 025001
[29] Zurlo et al. 2016 A&A 587, A57

APPENDIX 1: SELF-ATTENUATION EFFECT IN SPECKLE MASKING

As we told in Sec. 2, the SWAMIS feature recognition code can recognize a speckle only if it is detectable above the image noise and is brighter than a given intensity threshold. In fact, if we do a frame-by-frame comparison of the SWAMIS masks with the original frames, after 10^{-4} and 10^{-5} synthetic planets are injected, we never see a mask spot covering a known planet's position if no speckles are visible in the corresponding. This means, as we expected, that such faint planets are never directly masked in this process.

In our original SFADI presentation in [17], this fact led us to save computation time by computing the speckle masks for the Forerunner sequence once forever on the original dataset, instead of doing it after planets injection. But later we recognized that instead the speckle mask is influenced by the presence of a planet even if it is fully undetectable in the single frame. Figure 10, where we compare the SFADI results using speckle masks computed with and without injected planets, demonstrates this fact. As we see, the 10^{-5} contrast planets, which are well visible if speckle masks are computed without planets (right panel), are barely detectable if speckle masks are computed with planets (left panel). Moreover, close planets at 50 mas would be detectable at all contrasts if this planets attenuation effect were not present.

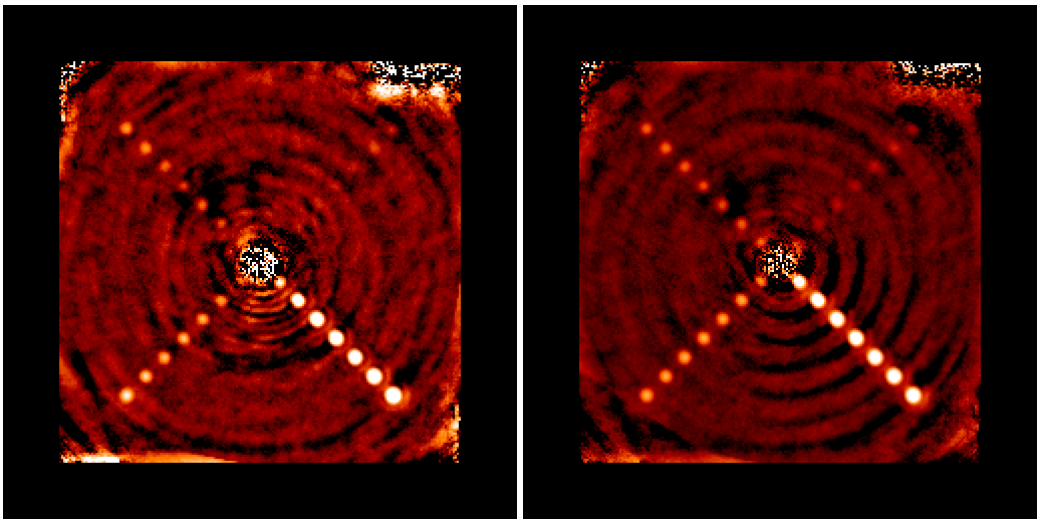

Figure 10: Planets attenuation effect. Left: SFADI with speckle mask computed with injected planets. Right: SFADI with speckle mask computed without injected planets.

We discovered that such self-attenuation is due to single-pixel differences of the speckle masks in the two cases. In particular, we found that when a speckle border falls on a planet's position, the mask for that speckle is, in some frames, one pixel wider than the same mask computed if the planet is not present.

Our explanation is that this happens, with a frequency proportional to the brightness of the planet, when a planet's photon hits a pixel adjacent to a speckle border (as defined by SWAMIS if planet was not present) and, at the same time, the detector noise fluctuation in that pixel is high. In this case, the speckle recognition algorithm sees a wider connected region and extends the mask correspondingly (see Figure 11).

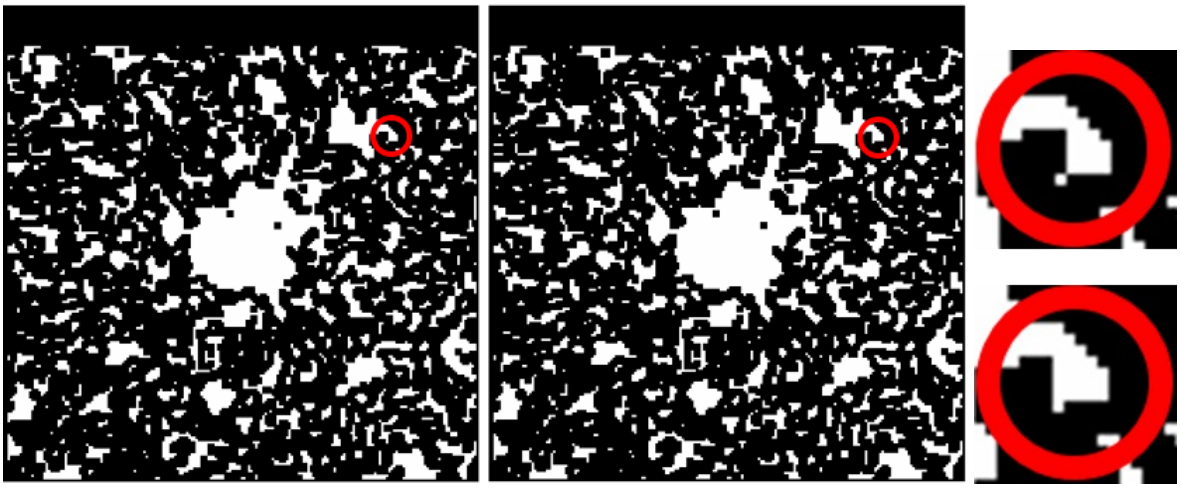

Figure 11: Comparison of the speckle masks computed with the same threshold and on the same frame when planets are not present (left) and after they are injected (center). The two masks are identical apart from a few pixels, like shown in the enlargements of one speckle falling on a planet's position (right), where we see that the mask is larger in the second case.

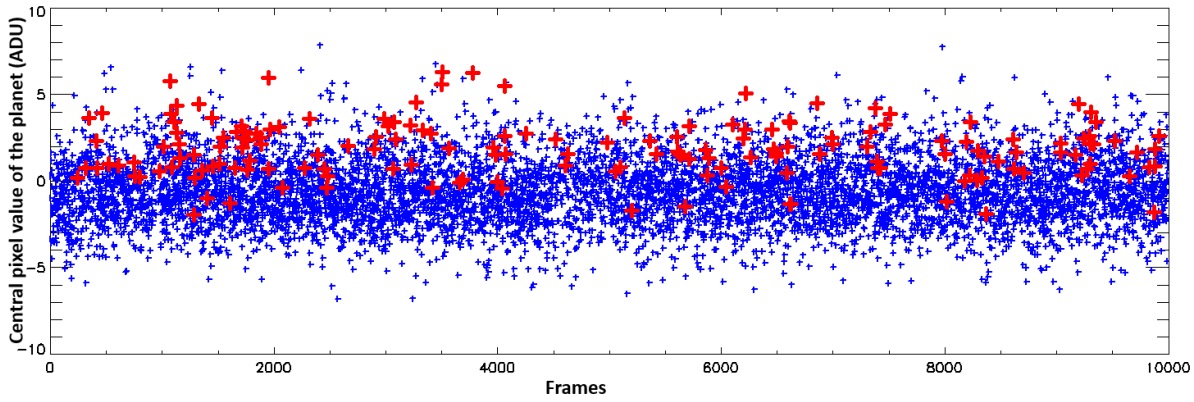

Figure 12: Distribution of the pixel value corresponding to a planet's center, after speckle masking computed without the planet (all points), and with the planet (blue points). Red crosses indicates the extra pixels selectively masked in the latter case and not in the former. The consequence of asymmetrically masking the red points pixels is that the average of the distribution is lowered, thus causing a flux attenuation in all pixels of the combined image where there is a planet.

This extra masking however, only happens for the high tail of a planet's intensity distribution, which thus becomes asymmetric, with the consequence that the average planet flux, after frames stacking, results attenuated (see Figure 12).

We can recovery the correct photometry of the attenuated planets by computing a flux correction, computed via simulations, but so far we have not found a way to avoid the attenuation effect and get back the detection performance expected by ideal speckle suppression. In fact, we tried various morphological filters and other processing but none of them results in a compensation of the asymmetry in the flux distribution. The problem thus remains open.

APPENDIX 2: ACCURATE PLANET INJECTION FOR UNBIASED PHOTOMETRY

During this work, we realized that a particular care must be used when injecting synthetic planets into a real frames sequence. In fact real images are stored from the camera as integer numbers (16-bit unsigned, in the Forerunner case), so that we must quantize the inserted planets model, but if we only consider Poisson statistics for photons and ADU rounding after gain conversion we get a biased photometry, as shown in central panel of Figure 13.

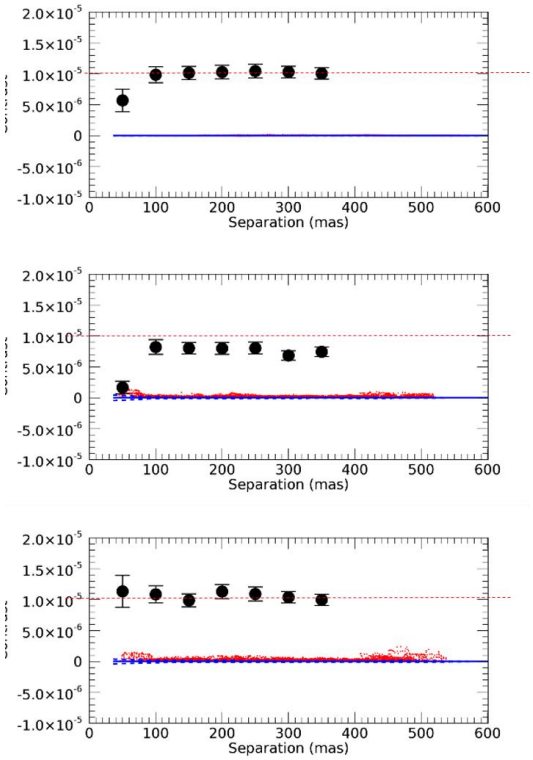

Figure 13: Photometry of an injected planet after SFADI stacking when simulating synthetic planet with floating-point values (top), integer photons and ADU quantization (center), and quantization corrected with Gaussian perturbation as described in the text (bottom).

We saw that this bias is produced, in case of non-integer detector gain, by the rounding error happening after multiplication of integer electrons by gain at the step of ADU computation. An error that propagates until 20% bias in the photometry for fluxes around one photons per frame in a given pixel.

Due that information is lost in the single-frame ADU rounding operation, the best solution that we found to avoid this problem and still keep integer pixel values is to modify their frame-to-frame statistics by adding a random perturbation before rounding, thus retrieving the correct average photometry, as shown in Figure 13, bottom panel.

To be more clear, our injection algorithm consists in the following steps:

1. compute mean floating-point photon image for the injected planets: $\boldsymbol{\gamma}$
2. convert to mean floating-point electrons by quantum efficiency: $\boldsymbol{e^- = \gamma \cdot Q_{eff}}$
3. compute a random perturbation with a standard deviation of $\sigma = 0.5$: $\boldsymbol{n = RND \cdot 0.5}$
4. add the perturbation to non-zero values only: $\boldsymbol{e^- = e^- + n}$
5. convert to floating-point ADU by detector gain: $\boldsymbol{ADU = e^- \cdot GAIN}$
6. convert to nearest integer ADU by rounding: $\boldsymbol{ADU = ROUND(ADU)}$
7. subtract previously added floating-point perturbation (times gain): $\boldsymbol{ADU = ADU - n \cdot GAIN}$
8. add this planets image to the original integer frame: $\boldsymbol{frame = frame + ADU}$