Speckle Interferometry: Reality vs. Simulation

Authors:*A. Tokovinin*
Version: 1
Date: 2023-11-09
File: prj/HARI/specultation/Report.tex

Abstract: Performance of the highly productive speckle interferometry instrument at the 4.1 m SOAR telescope is analyzed and compared to simulations, including such effects as finite exposure time, wide spectral bandwidth, and all kinds of noise. The overall agreement is reassuring, while the current signal processing algorithm is close to optimum. This gives a solid base for predicting the data quality for future observations and automating the process as much as feasible. An upgrade to a quantum CMOS detector with a sub-electron readout noise offers a gain in sensitivity by at least one magnitude.

1 Introduction

Speckle interferometry, introduced by Labeyrie (1970), has become a standard method of high-resolution observations at optical wavelengths, complementing adaptive optics in the infrared. Recording and processing of a large number of images with short exposures and fine pixel scales presented a formidable technical challenge at that time. With the advent of modern computers and electron multiplication (EM) CCDs, however, speckle interferometry has become a technically feasible and efficient technique. Nowadays, appearance of low-noise CMOS detectors challenges the dominant role of EM CCDs in speckle interferometry and opens new horizons.

The High Resolution Camera (HRCam) has been used at the Southern Astrophysical Research (SOAR) 4.1-m telescope in Chile since 2008. It has been a very productive and low-maintenance instrument for speckle interferometry. In 2017, the detector has been upgraded to an iXon EM CCD with a high quantum efficiency and a low noise (Tokovinin, 2018). HRCam uses 10-15 nights per year and typically covers 300 targets in a night. To date (2023 October), there are 38,781 accumulated speckle observations (measurements of binary stars and non-resolutions) — the largest speckle data set in the world.

So far, the performance of HRCam received little attention, all efforts have been devoted to the observations. Here, we want to investigate and understand the quality of HRCam data by comparing them to simulations. The primary motivation is to evaluate potential gain in sensitivity offered by an upgrade to a high-end CMOS detector. Our second goal is to quantify the quality of the speckle data by such parameters as speckle contrast and signal to noise ratio (SNR) and to study their dependence on the observing conditions and instrument parameters. Estimates of the expected SNR will help in planning future observations by setting acceptable conditions for each target. Eventually, the data quality prediction and control will help to automate the observations which, so far, largely depend on the observer's expertise.

In Section 2, the requisite information on HRCam is assembled, including the detector parameters, data processing, and the photometric calibration. Simulations are covered in Section 3, quantifying

the impact of finite exposure time and bandwidth and the role of the detector noise. The simulation code is described in the Appendix. Based on the simulations, the benefit of upgrading to a CMOS detector is explored in Sect. 4. The perspective of automating speckle interferometry in the future is briefly discussed in Sect. 5, and the results are summarized in Sect. 6.

2 Performance of HRCam at SOAR

A brief summary of HRCam is given here for consistency, with emphasis on the performance and algorithms. The reader is referred to Tokovinin (2018) for further details.

2.1 Detector Characteristics

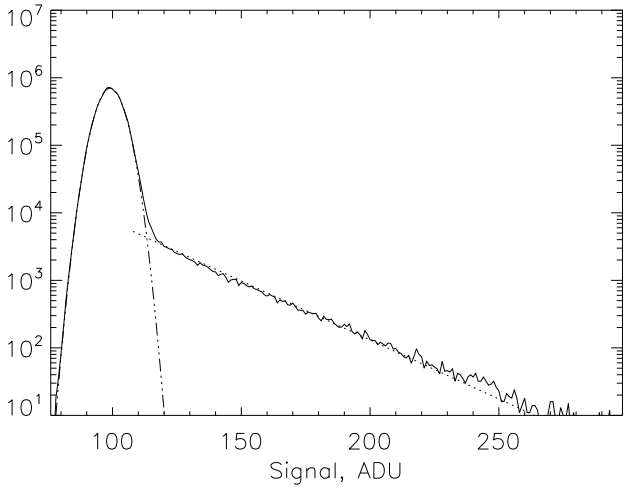

Figure 1: Histogram of the HRCam dark signal with a EM gain of $E_G = 300$ recorded on October 26, 2023 (full line). The fitted Gaussian and negative-exponential components are plotted by the dash-dot and dotted lines, respectively. Compare to a similar plot in (Tokovinin, 2018, Fig. 3).

Since 2017, the EM CCD detector iXon X3 888 from Andor has been installed in HRCam. The optical magnification has been adjusted to project its 13 μm square pixels at a scale of $p = 0.01575''$ on the sky, to sample the diffraction-limited speckles adequately. For a telescope diameter $D = 4.1$ m, the critical sampling of $\lambda/(2D)$ corresponds to a pixel scale of $0.0135''$ and $0.0207''$ at a wavelength λ of 540 and 822 nm, respectively. The speckle image is slightly under-sampled in the y filter and over-sampled in the I filter by 0.76 times.

The iXon camera has a conversion factor of $g = 10.1$ electrons per ADU, as specified by the vendor and confirmed by our measurements in 2016 ($q = 8.8$ was estimated in 2023). The electron multiplication (EM) gain E_g setting corresponds to the actual signal amplification, to within a few per cent. The rms readout noise (RON) is 4.5 ADU or 45 electrons, and its impact becomes negligibly small at $E_G > 100$. A typical histogram of the dark signal (Fig. 1) shows a Gaussian distribution with the width that characterizes the RON and an exponential tail produced by the single-electron events. These clock-induced charge (CIC) events do not depend on the exposure time (the dark current is negligibly small at -60° C), and their rate is defined by the fractional area of the exponential tail; it is 0.016 for the data shown in Fig. 1. The decrement of the exponent is 24.9 ADU or 250 electrons. It

is less than $E_g = 300$, likely because the CIC events are amplified slightly less than the photon events. Parameters of the detector, re-measured in 2023, show no degradation relative to our measurements in 2016, despite its intensive use during six years.

2.2 Data Acquisition and Processing

Most HRCam data are acquired in the standard mode, with a 200×200 pixels region of interest (ROI), 400 images per data cube, without binning. The field of view, 3.15″, is large enough to capture seeing-limited images without truncation. The minimum exposure time in this mode is 24.4 ms, and the effective exposure time (interval between successive images) is 27.9 ms. Acquisition of a single data cube thus takes 12 s, and two cubes per observation are normally recorded with $E_g = 100$. These parameters are adopted in the simulations. Faint targets are observed in the I filter; its transmission and the detector spectral response define the bandwidth with an average wavelength of 822 nm and a bandwidth of 140 nm. Note, however, that both parameters depend on the stellar temperature, and the effective response becomes "redder" for red stars. The y filter (543/22 nm) is used for observations of bright and close pairs with maximum angular resolution.

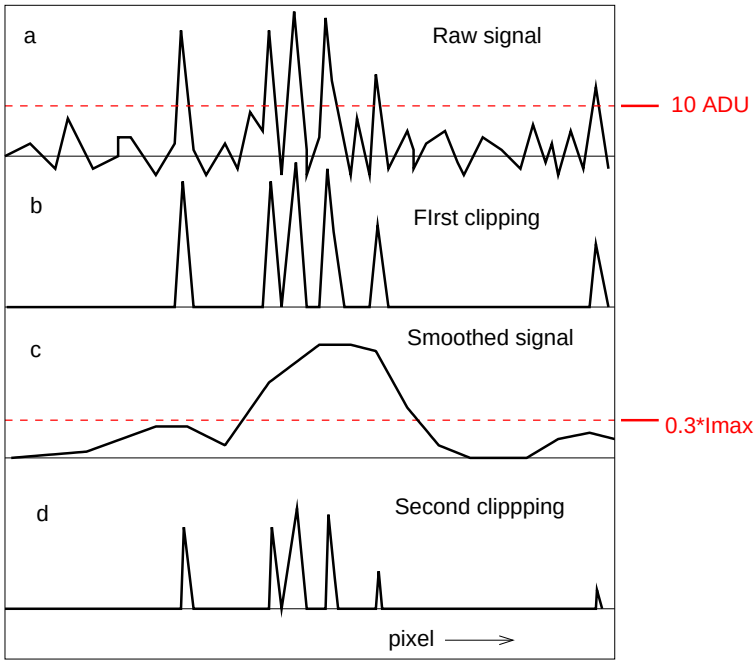

Figure 2: Illustration of the standard signal processing in HRCam. A single line of the bias-subtracted image at various processing stages is plotted by thick black lines, the thresholds are shown in red; see the text.

The algorithm of the data-cube processing has been "trained" on the real data to optimize detection of speckle signal from faint sources. It is illustrated in Fig. 2. The average number of CIC events in a 200×200 frame, 800, becomes comparable to the number of detected stellar photons $N_{\rm ph}$ for faint stars and seriously degrades the sensitivity. After subtraction of the bias, most pixels are empty, containing only the readout noise (a), while some pixels contain the speckle and CIC signals. To reduce the effect of RON, the signal is clipped at the level of 10 ADU, and all pixels below this level are set to zero (b). The first clipping still transmits the CIC events, which bias the calculation of the image centroid and flux. So, the clipped signal is smoothed with a 5-pixel square kernel, and the second threshold equal

to the 0.3 fraction of its maximum is defined (c). It is subtracted from the smoothed signal, negative pixels are set to zero, and the clipped smoothed image is used for the centroid calculation. The second threshold is also subtracted from the clipped signal (b), and the resulting non-negative image (d) is used to calculate the speckle power spectrum (PS) and the "lucky" image centered on the brightest pixel. The integral of the doubly clipped image (d) is much (by a factor of $\sim$7) less than the original signal, biasing the estimated fluxes.

The simulations described below confirm that the current algorithm produces near-optimum results for faint stars. However, calculation of the PS does not need the second clipping, and the image (b) can be used instead. It is important, however, to subtract from it the average background because the PS is normalized by its value at zero spatial frequency f, which equals the square of the total flux. For correct evaluation of the speckle power, only the stellar flux must be used for normalization, so the average background must be subtracted from the clipped image in (b) before calculation of the Fourier transform (FT) and summation of its square modulus over all frames in the cube. Otherwise, the CIC background contributes to the flux, producing a spike in the PS at $f = 0$. The second clipping used in the standard algorithm reduces the background effect, but not completely. So, the algorithm of PS calculation will be modified to use only single clipping with background subtraction. Other data products (re-centered and "lucky" images) will still be derived from the doubly-clipped images.

2.3 Photometric Calibration of HRCam

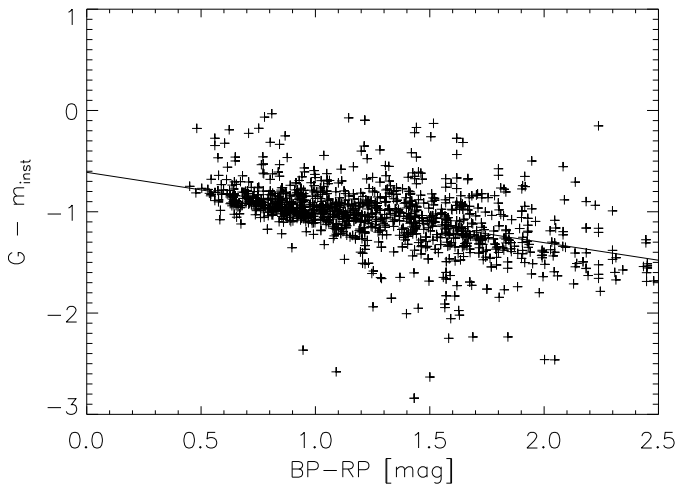

Figure 3: Photometric calibration of HRCam vs. Gaia. The line is a linear fit. Negative outliers are explained by poor transparency, positive outliers are wrongly identified sources.

During 2022-2023, a large number of red dwarfs from the 100-pc Gaia Catalog of Nearby Stars were observed with HRCam (Tokovinin, 2023). Parameters of 1325 such observations (flux, exposure time, EM gain, and coordinates) were retrieved from the general speckle database and matched to the GCNS sources, adding their photometry (magnitudes in the G, BP, and RP bands). Fluxes recorded in the database were strongly biased by the image clipping (see above), so they were recomputed from the saved average images.

Given the measured flux F in ADU, the conversion factor $g = 10.1$ electrons per ADU, the exposure

time t in seconds, and the EM gain E_g, the instrumental magnitudes of HRCam are computed as

$$m_{\rm inst} = 25 - 2.5\log_{10}[gF/(tE_g)]. \quad (1)$$

The offset of 25 mag is arbitrary. The difference between instrumental magnitudes and G is plotted vs. $BP - RP$ color in Fig. 3. The spectral response of HRCam in the I filter is redder compared to the Gaia G band, and the linear fit is

$$m_{\rm inst} \approx G - 0.61 - 0.35\ (BP - RP). \quad (2)$$

So, the photometric zero point of HRCam in the instrumental I system is 25.6 mag, and such star gives a flux of 1 electron/s. The instrumental magnitudes can be estimated as $G - 0.35(BP - RP)$.

3 Simulations

Simulation of speckle data is implemented in IDL (Appendix A). It is split in two parts. First, a noiseless data cube is generated. Then, the effect of all noise sources is simulated. This allows us to study the noise using the same input speckle pattern. Real data on a single bright star can be used as well for simulating noise and optimizing the data processing algorithm.

3.1 Exposure Time and Spectral Bandwidth

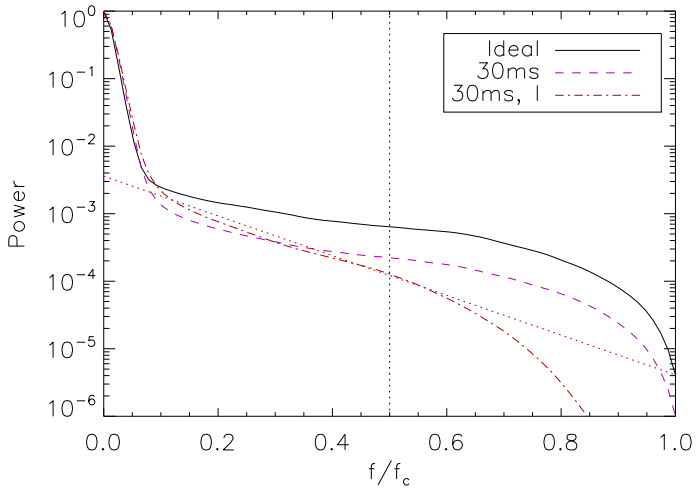

Figure 4: Azimuthally averaged PS of simulated noiseless speckle cubes for a 0.8″ seeing: an ideal (instantaneous and monochromatic) speckle at 800 nm, a monochromatic exposure of 30 ms (dashed line), and imaging in the I band (dash-dot). The black dotted line marks $0.5f_c$, the red dotted line is a linear approximation of the I-band PS.

The PS of speckle images has two components: the seeing-limited core at low spatial frequencies and the high-frequency part extending up to the cutoff frequency $f_c = D/\lambda$. Only the high-frequency part corresponding to the speckle signal is of interest here. The PS is normalized to one at zero frequency; it is rotationally symmetric. According to the theory, in an ideal speckle pattern the high-frequency part of the PS $P(f)$ should be approximately equal to $0.435(r_0/D)^2T_0(f)$, where r_0 is the Fried radius at the imaging wavelength and $T_0(f)$ is the diffraction-limited transfer function;

Table 1: Speckle contrast PS slope in simulations

Case	S	p_0	p_1	FWHM
Ideal	-3.204	-2.805	-0.776	0.688
30 ms	-3.671	-3.194	-0.924	0.628
30 ms, I	-4.033	-2.444	-2.943	0.609

$T_0(0.5) = 0.39$. For $D/r_0 = 20$, we get $P = 4.3 \times 10^{-4}$. It is convenient to use logarithmic quantities, so the speckle signal is characterized by the parameter $S = \log_{10} P(0.5 f_c)$; in the above numerical example, $S = -3.37$.

Figure 4 shows how the speckle the contrast is reduced by finite exposure time and wide spectral bandwidth. The same seeing of $0.8''$ was simulated using three different codes of increasing complexity (see Appendix A). The $\log_{10} P(f/f_c)$ of the real data is approximated by a straight line between $0.2 f_c$ and $0.8 f_c$, and these parameters (intercept p_0 and slope p_1) are determined for the simulated data as well. Obviously, $S \approx p_0 + 0.5 p_1$. Table 1 lists the PS parameters corresponding to Fig. 4. The Full Width at Half Maximum (FWHM) of the simulated average re-centered images is given in the last column in arcseconds. It is consistent between the simulations and less than the input seeing because the image motion has been compensated by centering. We note that the finite exposure time degrades the speckle contrast almost uniformly at all frequencies, and the slope p_1 becomes only slightly steeper; the PS shape resembles $T_0(f)$, as in the monochromatic case. On the other hand, a wide spectral bandwidth degrades mostly the high-frequency signal, and the PS becomes much steeper compared to the monochromatic speckle.

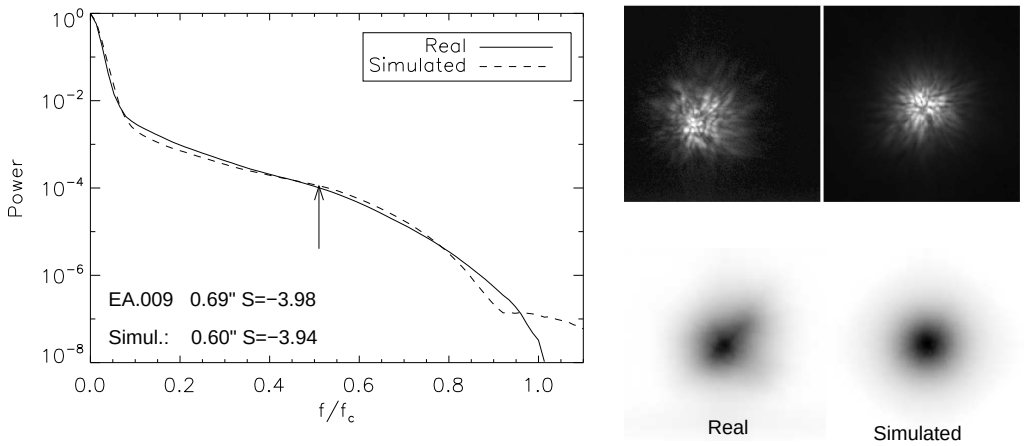

Figure 5: Comparison between real and simulated data in the I filter. Left: radially averaged PS, right: individual speckle images (top) and the average images in negative rendering (bottom). Note that the linear fit between 0.2 and 0.8 f_c (not plotted) is a poor approximation in this case, while it acceptable in some other cases.

For comparison with simulations, a real image cube EA.009 recorded on January 28, 2021 was selected. This is a bright reference star observed in the I filter with an effective exposure time of

27.8 ms. The seeing was relatively good, and the average image is almost round with a FWHM of $0.69''$. Figure 5 compares the real and simulated speckle PS. The values of S match almost exactly. The FWHM of the average simulated (re-centered and co-added) image is $0.60''$; it is smaller than the simulated seeing of $0.80''$ owing to the tilt compensation. The real image is slightly larger, but it shows signs of residual aberrations (note a small "tail" on the upper-right side).

To optimize the exposure time, we need to know the correlation between speckle patterns in successive frames of the data cube. For a lag of m frames, the correlation $C'(m)$ is defined as

$$C'(m) = \frac{1}{N_z - m} \sum_{i,pixels} I_i I_{i+m} / \sum_{i,pixels} I_i^2, \tag{3}$$

where I_i is the ith image. This correlation tends to 1/2 at large m, so a more meaningful parameter is $C = 2(C' - 0.5)$: it ranges from one for a perfect correlation to zero for uncorrelated images.

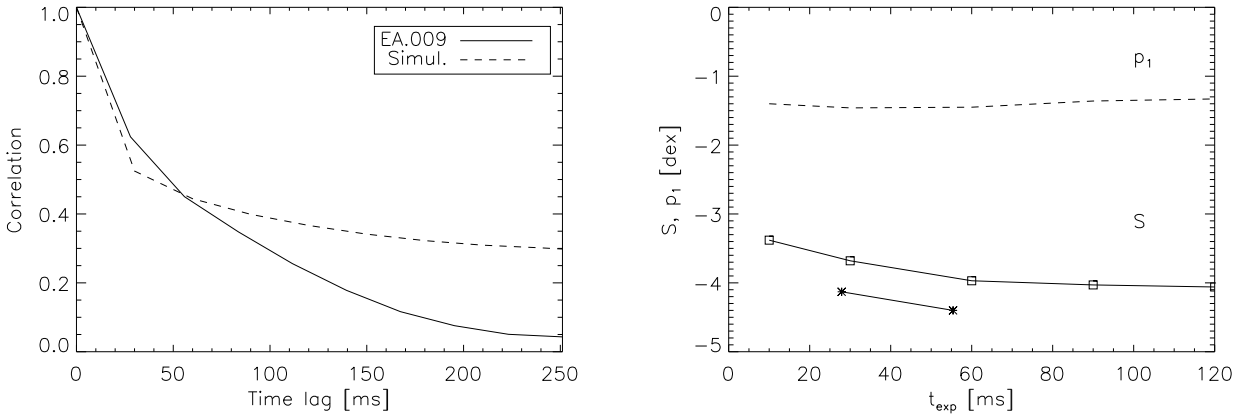

Figure 6: **Left:** correlation C between successive speckle images in the real image cube EA.009 (full line) and in the simulated monochromatic cube with a 30 ms exposure (dashed line). **Right:** dependence of the speckle signal S and PS slope p_1 on the exposure time. Asterisks show results for the real image cube EA.009.

The left panel of Fig. 6 shows correlation between frames in the real data cube EA.009 and in the simulated cube. At the largest lag of 0.25 s, the simulated correlation is still quite large, 0.29, while in the real cube it is only 0.05. This is a consequence of the crude two-layer representation of the atmosphere in simulations. Increasing the lower-layer wind speed makes the simulated correlation at 0.25 s smaller. On the other hand, the fast drop of the simulated correlation at $m = 1$ is qualitatively similar to the real data.

In Fig. 6, right, the effect of increasing the exposure time is explored. An exposure increase from 30 to 60 ms reduces the speckle signal S by 0.3 dex (a factor of two). This reduction is almost uniform at all spatial frequencies, and the slope p_1 remains approximately constant. Asterisks show results for the original image cube EA.009 and for the cube with frames binned (averaged) pairwise. Binning reduces the speckle signal from -4.13 to -4.40 dex, also by a factor of two.

Figure 7 shows how the level of the speckle signal in the I filter degrades with increased seeing. The slope p_1, on the other hand, remains relatively stable.

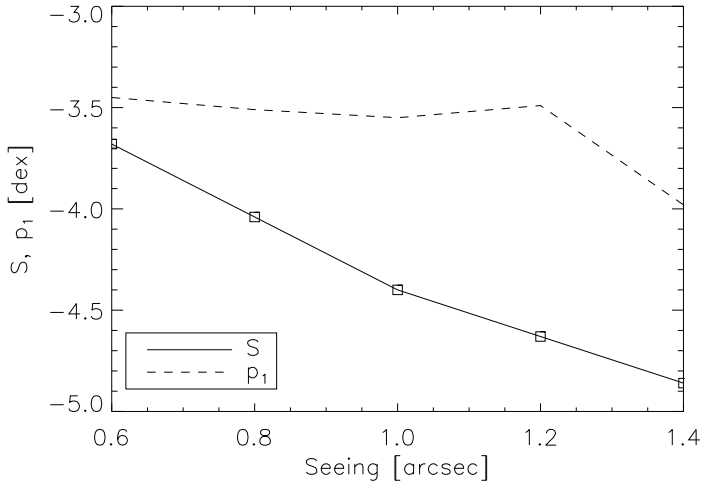

Figure 7: Dependence of the speckle signal S (solid line) and the PS slope p_1 (dashed line) on seeing. Simulations with a 30 ms exposure time and an I filter.

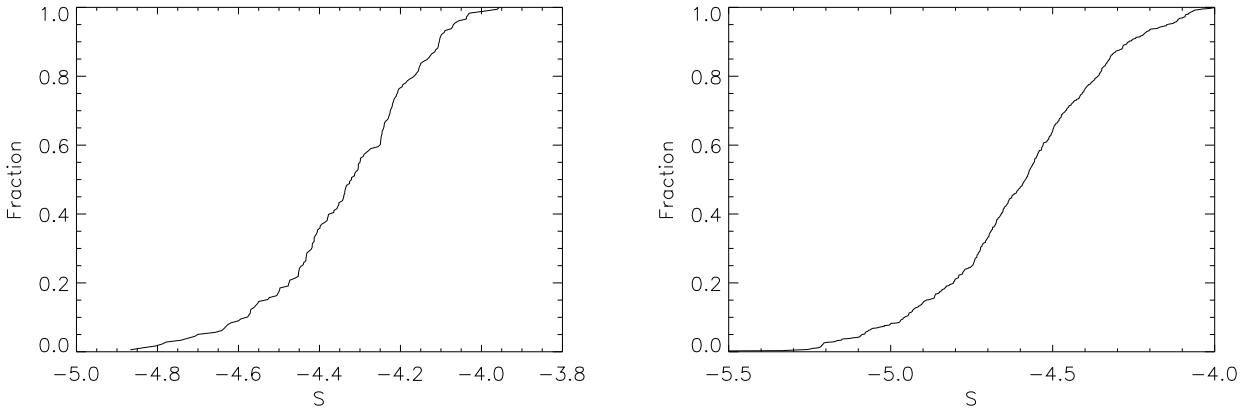

Figure 8: Cumulative histograms of the speckle signal S in the I filter without binning. Left: run 2021a, $N = 178$, median -4.32; right: run 2023a, $N = 620$, median -4.58. The median values of the PS slope p_1 are -3.78 and -3.62, respectively.

Cumulative histograms of the actual speckle signal S in two observing runs are shown in Fig. 8. Comparing with Fig. 7, the median values of S correspond to a seeing slightly better than $1''$. Note, however, that the standard processing algorithm with double clipping slightly biases S to smaller values, and this bias depends on $N_{\rm ph}$, as demonstrated by processing simulated data by two alternative methods; on the other hand, for bright stars double clipping has little effect. The values of slope p_1 match simulations reasonably well. As we saw above, p_1 is sensitive mostly to the spectral bandwidth.

Compared to the simulations, the speckle contrast in the real data may be reduced by additional factors such as optical aberrations (e.g. imperfect focusing) and telescope vibration. The reasonably good match between real data and simulations is encouraging, suggesting that these additional degrading factors play a minor role. The atmospheric parameters adopted in the simulations (e.g. the

wind speed) differ from the actual (unknown) parameters. Approximate agreement with the real data justifies the use of our simulations for analysis of the HRCam performance.

3.2 Signal to Noise Ratio and Limiting Magnitude

The noise simulator reads a noiseless image cube from a fits file, generates pixel values affected by the detector noise and random numbers of photons, computes the power spectrum, and determines its relevant parameters. The amplification noise in a EM CCD is also simulated: each photon generates a random signal distributed as a negative exponent with a decrement A which equals the average amplitude ($A = 24.5$ in Fig. 1), and the amplified signals from all photons are summed up. The photon numbers in each pixel follow the Poisson distribution with an average value equal to the sum of the stellar signal and CIC. Simulating the HRCam detector with $E_g = 100$, we adopt $A = 10$ and a CIC rate of 0.02. A threshold of 10 ADU is used in the PS calculation.

If a perfect detector records $N_{\rm ph}$ photon events per frame on average, the PS increases by $P_{\rm bias} = 1/N_{\rm ph}$ at all frequencies. This theoretical result has been reproduced in the simulation. The value of $P_{\rm bias}$ is evaluated by averaging the PS over frequencies beyond f_c, where the speckle power is zero. The bias is subtracted from the PS before its azimuthal averaging and model-fitting. Additional noise sources increase the PS bias, reducing the effective number of photons, and this reduction factor $1/(P_{\rm bias}N_{\rm ph})$ is a convenient dimensionless characteristic of the sensitivity loss caused by a noisy detector in comparison with an ideal one.

For bright stars, the photon noise should dominate, and the relation $P_{\rm bias} = 1/N_{\rm ph}$ should hold. So, the measured values of $P_{\rm bias}$ can be used to estimate $N_{\rm ph}$ and, by comparing it with the recorded flux, to check g, the conversion factor of the detector. Analysis of the HRCam data revealed that the flux recorded in the database was strongly biased by double clipping. Using correct fluxes verifies the expected linear relation between $P_{\rm bias}$ and $1/N_{\rm ph}$ and leads consistently to the gain factor $g = 4.7$ el/ADU, two times less than measured. The reason of this apparent discrepancy is the amplification noise. It doubles the signal variance, compared to a pure Poisson distribution, and effectively halves $N_{\rm ph}$. Simulation of a EM CCD with amplification noise indeed shows that $P_{\rm bias} = 2/N_{\rm ph}$. So, the HRCam data are consistent with simulations. For bright stars, where the speckle noise dominates anyway, the SNR loss due to the amplification noise is not detrimental, but it matters for faint stars, as well as CIC.

For one frame, the variance of the FT square modulus at each spatial frequency equals its mean value (this is a consequence of the negative-exponential distribution of the square modulus). Averaged over N_z frames, the relative rms fluctuations of the PS equal $1/\sqrt{N_z}$. Let 10^S be the value of the speckle power at half of the cutoff frequency, which we identify with the useful signal. The total PS signal is $10^S + P_{\rm bias}$, so the signal to noise ratio (SNR) in one element of the PS is

$$SNR = \sqrt{N_z}\ 10^S/(10^S + P_{\rm bias}). \tag{4}$$

For bright stars, $10^S \gg P_{\rm bias}$, the speckle noise dominates, and SNR tends to 20 for $N_z = 400$. For faint stars, $P_{\rm bias}$ becomes the dominant term, and its fluctuations in the average PS determine the SNR. Transition between the bright-star and faint-star regimes around $N_{\rm ph} \sim 10^4$ depends of the level of the speckle signal S.

Optionally, a binary star can be simulated by adding a shifted copy of the speckle pattern scaled by the binary intensity ratio, before simulating the noise. Figure 9 shows fringes in the PS of a binary

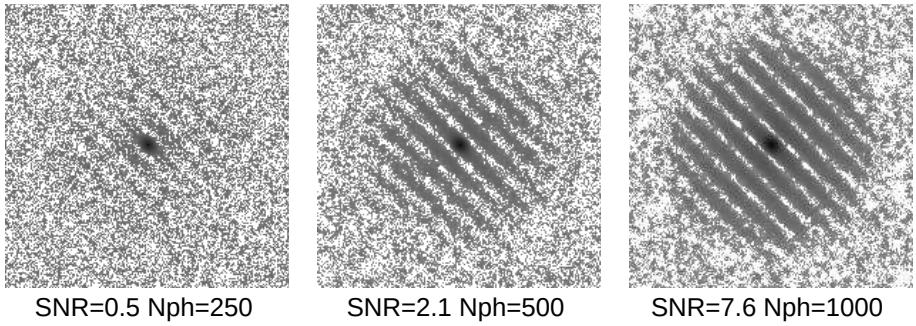

Figure 9: Power spectra of a simulated binary star with equal components displayed on a negative logarithmic scale. The SNR and $N_{\rm ph}$ are indicated. Simulated EM CCD with CIC=0.02, Ampl=10, and threshold=10. An ideal speckle data cube was used at input.

star at different photon fluxes. Note that the SNR parameter here refers to single stars (fringes reduce the average level of the spectrum and further decrease the SNR). At SNR=0.5, the fringes are barely visible because the speckle signal is lost in the fluctuations of $P_{\rm bias}$.

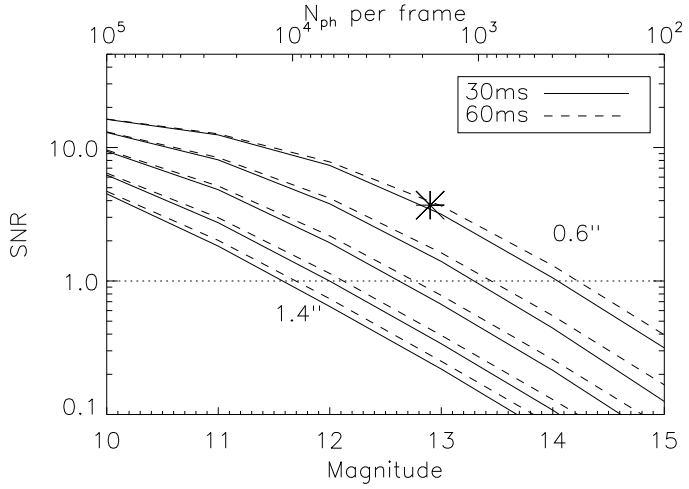

Figure 10: Signal to noise ratio in the simulated HRCam i with a EM CCD detector vs. stellar magnitude in the instrumental system I. Full lines and upper axis correspond to the 30 ms exposure, dashed lines to the 60 ms exposure, both with 400 frames per data cube. The curves are computed for seeing values of 0.6, 0.8, 1.0, 1.2, and 1.4 arcsec. The asterisk shows the observation of V1311 Ori D in 2021.7983 (see text).

Using realistic simulations of the noisy speckle data allows us to explore the combined effect of the seeing variation and the source flux on the SNR. Results of these simulations are plotted in Fig. 10. The instrumental magnitudes of stars that reach S/N=1 range from 11.5 to 14.5 mag, depending on the seeing. These estimates are slightly optimistic because, compared to the simulation, the speckle contrast is additionally reduced by imperfect optics (e.g. a focus error) and by telescope vibration. On the other hand, a slower wind in the upper atmosphere would increase S and the magnitude limit relative to the simulations. Doubling the exposure time from 30 to 60 ms doubles the flux $N_{\rm ph}$, but reduces the speckle power by the same amount, so the small net SNR gain is due to secondary factors like CIC; it comes at a cost of doubling the acquisition time.

As a real example, consider observation of the faint red dwarf V1311 Ori D reported in (Tokovinin,

2022). On that date, 2021.7983, the seeing was very good. The FWHM of the centered images in the data cube EK.026 (exposure time 50 ms) is 0.63″, the estimated speckle signal is $S = -4.38$, and the SNR is 3.7. With the instrumental magnitude of 12.92, estimated from the $G = 13.92$ mag and the color ($BP - RP \approx 2.4$), the experimental point (asterisk) falls near the upper curves in Fig. 10. Note that for very red stars the effective spectral response differs from the response assumed in the simulation.

4 Comparison between EM CCD and CMOS Detectors

Recently, the noise level of CMOS detectors has improved to the point where they become competitive with EM CCDs. Amateur CMOS detectors, e.g. from ZWO, are very cheap and readily available, while their typical noise is about 1 electron. Hamamatsu developed recently a scientific qCMOS camera ORCA-Quest with a sub-electron RON.[1] Its indicative cost is $ 50K. This Hamamatsu camera has been installed in the speckle interferometer of the 2.5-m telescope and demonstrated improvement in limiting magnitude (Strakhov et al., 2023). A CMOS detector, even with a 1-el. RON, eliminates such problems of EM CCDs as CIC and amplification noise. It is assumed here that the quantum efficiency of all detectors (back-illuminated silicon) is similar.

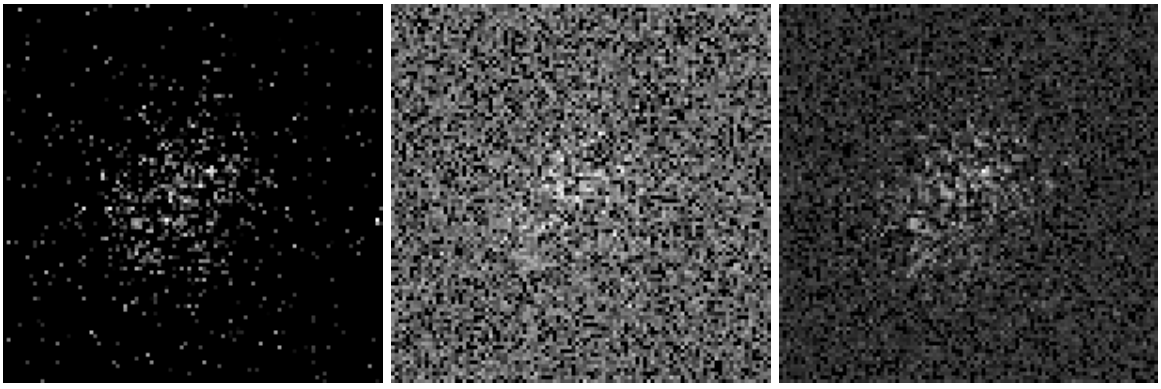

Figure 11: Single noisy frames with $N_{\rm ph} = 1000$ (about 2.5 photons per speckle) corresponding to a EM CCD (left), a CMOS with RON=1 (center) and a qCMOS with RON=0.3 (right). Ideal speckle images with $D/r_0 = 13.5$ are used. The central region of 100×100 pixels is shown in the square-root stretch.

These detectors are compared below by additional simulations. The amplification noise and CIC are both absent for CMOS. We assume a RON of 1 el for a cheap CMOS camera and 0.3 el for a qCMOS. The EM CCD parameters are same as above. Figure 11 compares simulated speckle images of a faint source as they would look in three detectors, with the same $N_{\rm ph}$.

Signal clipping in the calculation of the PS is still needed for a CMOS to reduce the impact of RON in "empty" pixels that do not contain useful signal. Figure 12 plots the effective flux loss vs. threshold for a low-noise qCMOS camera and two levels of photon flux. The optimum threshold appears to be near 3×RON. For a CMOS with RON of 1 el., the same 3-sigma threshold is adopted. It seems that

[1] https://www.hamamatsu.com/us/en/product/cameras.html

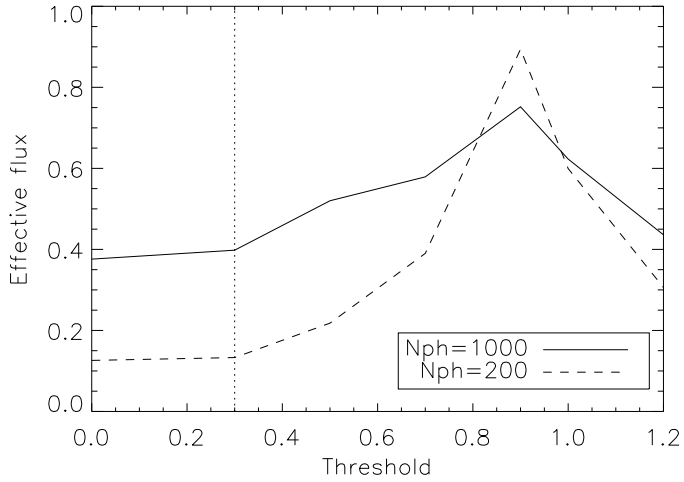

Figure 12: Effective loss of flux $1/(P_{\rm bias}N_{\rm ph})$ vs. threshold for a qCMOS with RON=0.3 (marked by the vertical dotted line) and for two levels of the flux.

clipping has not been used by Strakhov et al. (2023), who found a smaller gain in sensitivity when switching from a EM CCD to a qCMOS. Note also that their EM CCD has a higher level of CIC events than at SOAR (0.04 vs. 0.02).

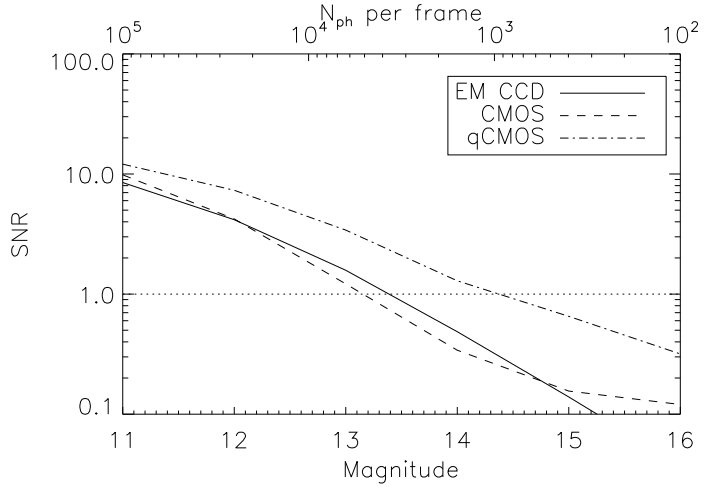

Figure 13: Signal to noise ratio (SNR) in the PS at half of the cut-off frequency vs. magnitude in the I band for a $0.8''$ seeing and three simulated detectors: an EM CCD, a scientific CMOS (RON of 1 el), and a qCMOS (RON of 0.3 el). The upper axis gives the number of photons per frame in a 30 ms exposure, the SNR assumes 400 frames per data cube.

Figure 13 compares the SNR vs. magnitude plot for the three detectors. It assumes equal $N_{\rm ph}$ and an optimized threshold in the PS calculation. A realistic level of the speckle signal $S = -4$ dex corresponding to a $0.8''$ seeing and an I filter is adopted (see Fig. 7). An upgrade from a EM CCD to a qCMOS offers a sensitivity gain of $\sim$1 mag. Note that the SNR is computed for $N_z = 400$. Increasing the number of acquired frames by a factor of ten pushes the dotted line down by a factor of three, and 16th mag stars become accessible with a qCMOS.

The SNR plot in Fig. 13 refers to the speckle signal at $0.5f_c$. In a wide spectral band like I, the PS has a steep slope, and its value at $0.2f_c$ is larger by almost an order of magnitude, $S \sim -3$ dex

(Fig. 4). So, speckle observations of faint stars are feasible at a reduced spatial resolution, and there is an obvious trade-off between resolution and sensitivity.

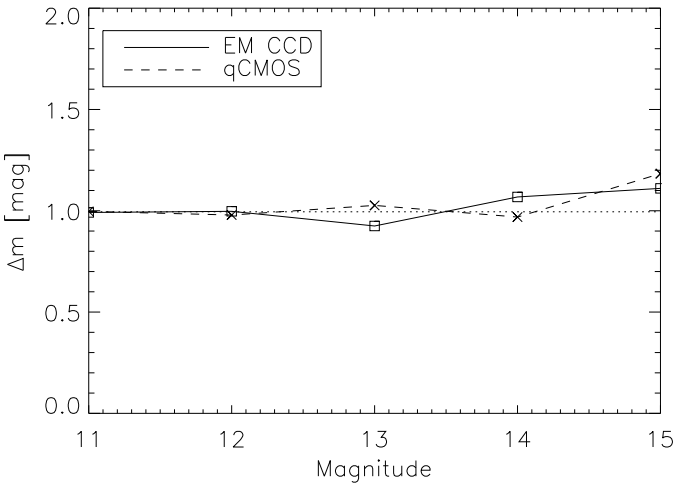

Figure 14: Magnitude difference of a simulated binary measured by fitting fringes in the noisy PS. The input value $\Delta m = 1$ is marked by the dotted line. The simulated detectors are a EM CCD (solid line) and a qCMOS (dashed line).

Although CCDs and CMOS detectors are linear, the speckle signal processing involves two non-linear operations: clipping and PS calculation. The PS is proportional to the square of the signal, and an additive background such as CIC events is no longer additive in the PS. A numerical experiment was conducted to evaluate the effect of signal processing on the relative photometry of a binary with a magnitude difference of 1 mag (Fig. 14). The result is reassuring. Even in the conditions of SNR<1, when the fringes are barely seen (see Fig. 9), no systematic trend in the estimated Δm is present with both a EM CCD and a qCMOS.

5 Towards Automated Observations?

Speckle observations is a complex process. It starts with preparation of the common program for each observing run. Merging all active projects into a common program has many advantages and increases the overall efficiency, compared to the classical by-program telescope use. A common set of calibrators, for example, ensures consistent data for all projects.

The outcome of each observation depends on the current conditions, mostly on the seeing and transparency. Such factor as telescope shake is important on windy nights, restricting the pointing direction. So, the selection of targets is managed flexibly in real time, balancing between priorities, visibility, and current conditions, while optimizing the telescope slews. Quick online calculation of the PS helps the observer to evaluate the performance and to adjust the target choice accordingly. For example, if a source is resolved into a new tight binary, observation of a reference star immediately after is needed for a correct data reduction.

Considering this inherent complexity and the need to take real-time decisions based on several variables, it will be difficult to substitute an experienced observer by an automatic process. The study done here helps this task in several ways, allowing to replace subjective evaluations of the observing conditions and the results by quantitative metrics such as $N_{\rm ph}$, S, and SNR. Simulations will help

us to define reasonable performance goals and acceptable conditions for each target; this will set the stage for automating the observations. The automation should be gradual, starting from sequencing routine actions and progressing towards robotic operation under human supervision.

Processing of the observations acquired on a speckle night starts by running a pipeline that computes the PSs and the associated data products, as described in (Tokovinin, 2018). Manual inspection of the data and fitting binary parameters is relatively fast: it takes less than a half of the telescope time used to acquire the data; about as much time is spent for the subsequent analysis of the results. The results are incorporated into the observing program and used to decide on the next observations. For example, detection of a rapid orbital motion may call for a repeated measurement within a year or even sooner. Thus, the process program-observations-program is circular, with a rapid feedback. This differs from the classical open-ended approach where the data acquisition is disconnected from the data analysis, which may be done years later.

6 Summary

So far, operation of HRCam relies on the human experience and on the quick-look evaluation of the PS immediately after acquisition of the image cube. Faint targets are observed in the wide I band, maximizing the detected flux; the exposure time is also often doubled from its standard value of 27.9 ms. In contrast, bright stars are observed mostly with the narrow y filter to maximize the resolution. Targets are selected flexibly, depending on the observing conditions.

This study compares the actual level of the speckle signal S with realistic simulations that model the effects of seeing, finite exposure time, and spectral bandwidth; a reasonable agreement has been found (see Fig. 5). Determination of the photometric zero point establishes the relation between the number of detected photons $N_{\rm ph}$ and the Gaia photometry. Taken together, these results allow calculation of the SNR as a quantitative measure of the data quality. Estimation of the SNR will help in planning future observations and in their optimization in real time, and lays the ground for potential automating speckle interferometry in the future.

In the process of confronting real HRCam data with simulations, certain aspects of the data processing have been clarified. The current algorithm uses double clipping in the calculation of the PS. The resulting fluxes are strongly biased, and the double clipping is not really necessary (it is needed, however, for correct centroiding of faint stars). Simulations indicate that the current algorithm, tuned empirically using real data, gives near-optimum results, and changing its parameters such as the threshold does not lead to a noticeable SNR improvement.

Evaluation of the sensitivity gain from an upgrade to a qCMOS detector has been the main motivation of this study. Simulations demonstrate that even a cheap CMOS with a 1 el. RON can be competitive with the current EM CCD, if suitable clipping is applied the data processing. A qCMOS with 0.3 el. RON offers a gain of at least one magnitude in the sensitivity (Fig. 13). A CMOS detector overcomes two major limitations of a EM CCD, namely the CIC background and the amplification noise.

A Simulation Code

The IDL code `simspec4.pro` simulates a realistic speckle data cube for the 4.1-m SOAR telescope. Monochromatic image of a star at wavelength λ is computed as a FT of the light-wave amplitude at the pupil. The size of the computing grid N, 200×200, and the angular pixel scale $p = 0.01575''$, are chosen to match the real data. The physical size of the grid in the pupil space equals $L = \lambda/p$ (10.5 m for 0.8 μm, 0.05 m per spatial pixel). This sets the pupil radius in pixels; the central obstruction of 0.25 is also emulated in constructing the pupil mask. Figure 15 helps to visualize the geometry of simulations.

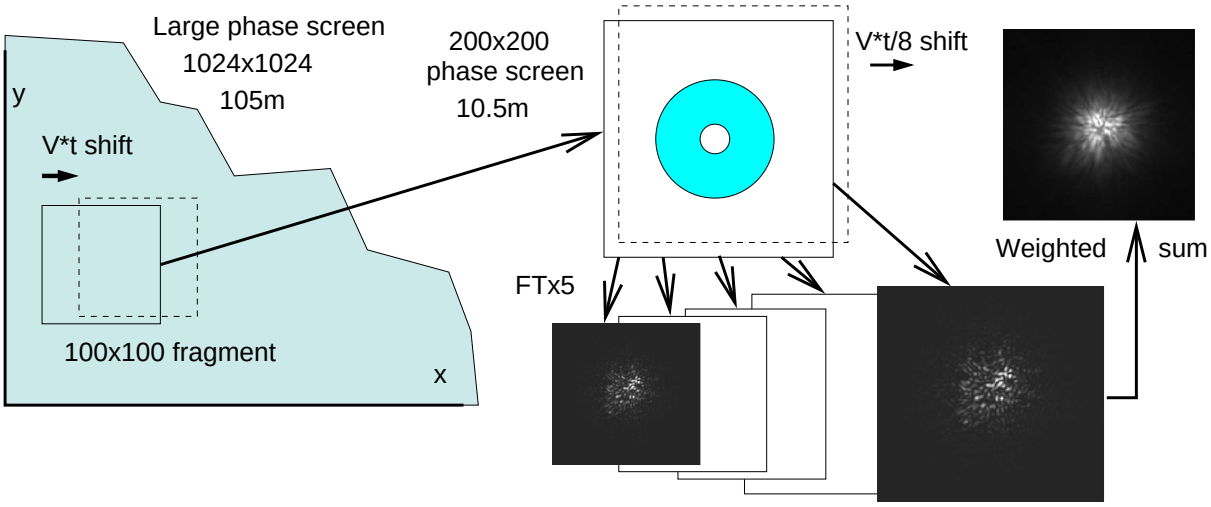

Figure 15: Scheme of the speckle simulation code (see text)

Atmospheric phase screens obeying Kolmogorov statistics are generated using the standard Fourier method, given the Fried parameter $r_0 = 0.98\lambda/\epsilon$ for the seeing ϵ radian. In the simplest version of the code, the same pupil-space grid is used to generate the phase screens. However, to simulate effects of the finite exposure, the phase screens of 1024×1024 with twice larger pixels (physical size 105 m) are generated. The pupil grid corresponds to a 100×100 fragment of such screen, and it is interpolated on a finer 200×200 grid for the image calculation.

To simulate temporal evolution of the speckle pattern, at each successive frame the origin of the selected fragment of the large phase screen is translated horizontally (in X) by Vt, where V simulates the wind speed and t is the exposure time (e.g. 1.2 m for $V = 40$ m/s and $t = 0.03$ s). The translation is rounded to an integer number of pixels. Orthogonal shifts are applied at every 10th translation in order to sample the full large screen, so the fragment's motion is in fact tilted with respect to X or Y axes by 1/10 (5.7°). When the fragment reaches the edge of the large screen, it "rolls over" in both coordinates without discontinuity (the screens are doubly periodic owing to their generation method).

Phase perturbation at the telescope pupil is produced by several turbulent layers moving with different speeds, and the temporal evolution of speckle is governed mostly by their changing sum rather than by the overall translation over the pupil. So, two phase screens are generated. One is translated in X with a speed of 8 m/s, another in Y with a speed of 40 m/s. The turbulence intensity is equally distributed between these screens. For a seeing of 0.8" this layout corresponds to the

atmospheric time constant of 1.3 ms (at 500 nm). All parameters in the code can be easily modified.

To reproduce the image smearing during exposure, the latter is split into 8 steps. The two phase-screen fragments selected for a given exposure and re-binned on a 200×200 grid are shifted during the exposure by small steps of $Vt/8$ in the orthogonal directions. At each step, the speckle image is computed, and each frame of the image cube is the average of these eight images. For the following frame, new fragments are cut out from the large phase screens.

Simulation of speckle in a wide bandwidth adds another layer of complexity. The spectral response of HRCam in the I filter is modeled by a combination of 5 wavelengths λ_i from 0.75 to 0.90 μm with relative weights of [0.724, 0.694, 0.539, 0.347, 0.190] — product of the filter transmission and the detector response. The effective wavelength is 0.822 μm. The actual response depends on the spectrum of the star, of course.

For each of the 8 sub-frames, we compute five images corresponding to λ_i and combine them with relative weights defined by the spectral response. So, calculation of one speckle frame requires 8×5=40 FTs, and simulation of the image cube takes about 45 s The phase perturbation, defined for the reference wavelength λ_0, is scaled by the factor $f = \lambda_0/\lambda_i$. However, this is not sufficient because the pupil size and the computing grid are dimensioned for λ_0, not for λ_i. This is accounted for by stretching each monochromatic image by the factor $1/f$ (at longer wavelengths, the speckles become larger). The same effect could be achieved by shrinking the phase pattern and the pupil function by a factor f before calculating the image by FT. As a check, suppose that the phase aberration is linear (pure tilt). Scaling it by f reduces the tilt and moves the image closer to the field center. Stretching the image by $1/f$ times moves it back to the original position, so the tilt becomes achromatic. Speckles in the simulated polychromatic images are extended radially, as in the real images, while the image moves as a whole owing to random tilts.

The noiseless image cube (either simulated or real) serves as input to the general-purpose noise simulator `noisesimul.pro`. As the simulator does not "know" the wavelength, telescope diameter, and pixel scale, the cutoff frequency f_c (in pixels) must be provided at input. The number of photons per frame $N_{\rm ph}$ sets the detected flux from the star. The noise simulation algorithm is described below; it is implemented by the following piece of code relevant to one frame:

```
  tmp = cube[*,*,i] ; noiseless speckle image
  if keyword_set(binary) then tmp += binary[2]*shift(tmp, binary[0], binary[1])
  tmp = tmp/total(tmp)*Nph ; normalize by average photon number
; Poisson distribution of events in each pixel
  for k=0,nx-1 do for l=0,nx-1 do begin ; loop over pixels
    npix =  randomn(seed, poisson=tmp[k,l]+CIC)  ; number of events in pixel
    if (ampl gt 1) then begin ; EM CCD, amplification noise
      s = 0
      if (npix gt 0) then for j=0,npix-1 do s += -alog(randomu(seed))
      ncube[k,l,i] = ampl*s + ron*randomn(seed)
    endif else ncube[k,l,i] = npix + ron*randomn(seed)  ; CMOS
  endfor  ; pixel loop
```

Each frame `tmp` is normalized to a unit sum and multiplied by $N_{\rm ph}$ to get the expected photon number per pixel. This number serves to generate the actual random number of photons `npix` after adding the CIC probability (the Poisson random number generator is used). For a CMOS detector,

CIC=0, and the pixel value equals `npix` plus a Gaussian readout noise, which must be specified in electrons. For a EM CCD, each photon is "amplified" to an average level `Ampl` (in ADU) with a negative-exponential distribution, and the resulting pixel signal is a sum of all randomly amplified photons. This models the EM amplification noise. The readout noise is also simulated, but for a EM CCD it should be specified in ADUs rather than in electrons. The noise-amplification step is skipped for a CMOS by setting `Ampl=1`. Simulated noisy speckle images are illustrated in Fig. 11.

References

Labeyrie, A. Attainment of Diffraction Limited Resolution in Large Telescopes by Fourier Analysing Speckle Patterns in Star Images. 1970, A&A, 6, 85

Strakhov, I. A., Safonov, B. S., Cheryasov, D. V. Speckle Interferometry with CMOS Detector. 2023, Astroph. Bull., 78, 2345

Tokovinin, A. Ten years of speckle interferometry at SOAR. 2018, PASP, 130, 5002

Tokovinin, A. The family of V1311 Ori: a young sextuple system or a mini-cluster? 2022, AJ, 163, 127

Tokovinin, A. Exploring Thousands of Nearby Hierarchical Systems with Gaia and Speckle Interferometry. 2023, AJ, 165, 180