Linear Dark Field Control

Milestones White Paper

Milestone #2

Spectral LDFC demonstration 10x raw contrast stability gain at <10^{-7} raw contrast

Milestone #3

LDFC PSF calibration: 10x contrast gain to 1e-8 raw contrast

Team:
Olivier Guyon (PI), Kyohoon Ahn, Ruslan Belikov, Michael Bottom, Mark Chun, Thayne Currie, Vincent Deo, Richard Frazin, Tyler Groffl, Sebastiaan Haffert, Brian Kern, Julien Lozi, Jared Males, Ben Mazin, Bijan Nemati, Barnaby Norris, Eugene Pluzhnyk, Dan Sirbu, Hari Subedi

Apr 26, 2022

Approvals

Released by:

Apr 26, 2022

Olivier Guyon — Date

Principal Investigator, University of Arizona & Subaru Telescope

Approved by:

E-SIGNED by Nicholas Siegler
on 2022-04-27 13:43:31 GMT

Nicholas Siegler — Date

Exoplanet Exploration Program Chief Technologist, NASA JPL

E-SIGNED by Brendan Crill
on 2022-04-27 11:34:56 GMT

Brendan Crill — Date

Exoplanet Exploration Program Deputy Chief Technologist, NASA JPL

E-SIGNED by Doug Hudgins
on 2022-04-27 14:23:28 GMT

Douglas Hudgins — Date

Exoplanet Exploration Program Scientist, NASA HQ

Introduction, Background

The milestones described in this document are part of the NASA-funded effort "Linear Wavefront Control for High Contrast Imaging", which is aimed at improving the efficiency, sensitivity and reliability of wavefront control for exoplanet imaging.

We will develop and demonstrate in the laboratory a new approach to focal plane wavefront control based on linear sensing and control of bright regions of the diffracted and scattered starlight halo rather than the current non-linear iterative techniques which require probing with a DM. Our approach, referred to as Linear Dark Field Control (LDFC), will simultaneously improve the speed and sensitivity of wavefront control and provide the telemetry required for PSF post-processing calibration.

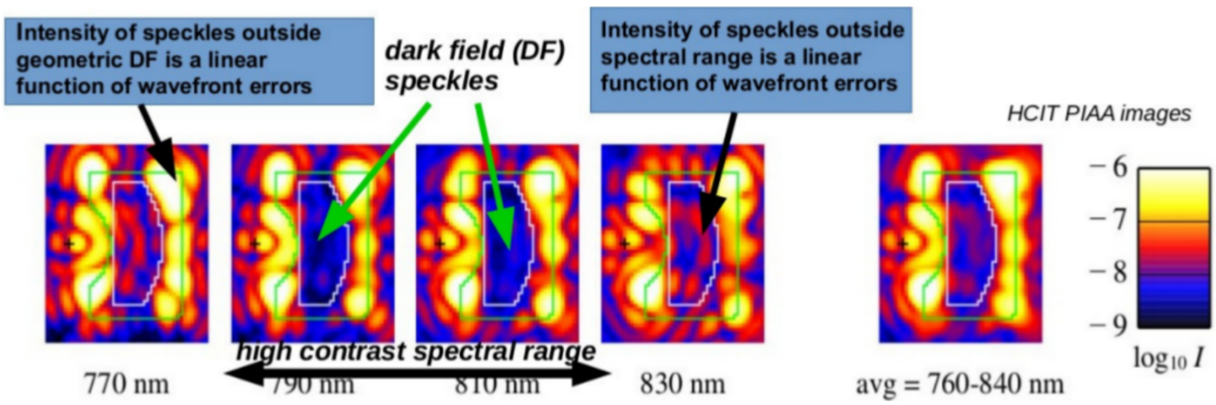

Fig 1: Laboratory coronagraph images illustrating spectral LDFC signal. The images were acquired at the High Contrast Imaging Testbed (HCIT) with a Phase-Induced Amplitude Apodization (PIAA) coronagraph. The contrast ratio is deepest within the control band, and quickly increases outside the spectral band and outside the field mask. The associated light responds linearly to small wavefront perturbations, and is the signal input for LDFC.

LDFC uses starlight that is falling outside the high contrast region (Fig 1) to drive a linear control loop aimed at "freezing" the wavefront state. This bright light (referred to as the ***bright field – BF***) is abundant relative to the much smaller amount of light in the high contrast region (referred to as the ***dark field – DF***). It responds linearly to small wavefront errors, allowing a linear control loop to operate at higher speed and sensitivity than possible with conventional wavefront control schemes relying on measurements in the high contrast region.

The BF can consist of light that is outside the DF in either spatial or spectral dimensions, corresponding to **spatial LDFC** or **spectral LDFC** respectively. Both can also be combined in the **spatial+spectral LDFC** approach. These configurations are illustrated in Fig 2. The same BF signal that drives LDFC's active wavefront control can also be used for post-processing, where the starlight component of the DF component is derived from simultaneous BF measurements, and then subtracted from the raw measurements to reveal faint companions.

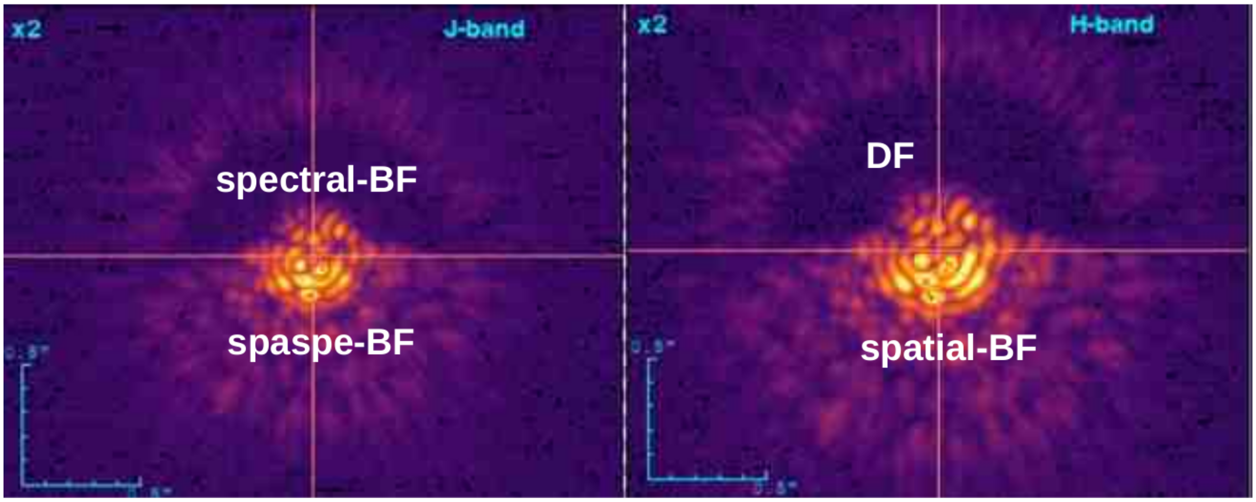

*Fig 2: Laboratory coronagraph images illustrating LDFC signal, acquired on the SCExAO testbed in the near-IR. Left: J-band (1.2 micron). Right: H-band (1.5 micron). The wavefront control system is driven here to achieve high contrast in H-band in the top part of the PSF (dark field region **DF**). In spatial and spectral LDFC, the H-band bright field **spatial-BF** and J-band region **spectral-DF** are used respectively for wavefront stabilization. In LDFC postprocessing, any combination of regions **spatial-BF**, **spectral-BF**, and **spaspe-BF** are used to estimate starlight distribution in region **DF**.*

1. LDFC Algorithm Description

LDFC milestones #2 and #3 are aimed at validating active wavefront stabilization, and postprocessing respectively. The corresponding algorithms are described in this section.

1.1. Active Wavefront Stabilization with LDFC

The Linear Dark Field Control (LDFC) wavefront stabilization loop for high contrast imaging is defined by the following key properties :

- [*LDFC-requ1*] Does not require runtime DM probing for wavefront sensing.
- [*LDFC-requ2*] Relies on focal plane images for wavefront sensing.
- [*LDFC-requ3*] Uses a linear relationship between sensing input and control output
- [*LDFC-requ4*] Uses pixels outside (in spatial and/or spectral dimensions) the high contrast area for wavefront sensing input

LDFC uses light outside the high contrast dark hole for sensing (*LDF-requ4*), as this component (referred to as the bright field) responds linearly to small perturbations in wavefront errors. Here “outside” may refer to spatial or wavelength dimensions, or a combination of both. For example, the spatial LDFC approach uses light that is outside the dark hole geometrical area, and is expected to be especially valuable in single-sided dark hole configuration. Spectral LDFC uses light at wavelengths shorter and/or longer than the dark hole’s spectral band.

LDFC is a wavefront stabilization technique, but is not suitable to iteratively build a dark hole, so it is foreseen that the two approaches would be used alongside. Once a dark hole is established, LDFC may be used as the sole control loop, or can run in addition to other control loop(s).

The technique's strengths and limitations are listed below, compared to the better established DM probing approaches.

LDFC strengths
- **Sensitivity**: LDFC can use more light than available within the spatial and spectral extent of the dark hole, resulting in improved sensitivity
- **DM probing-free**: Since no DM perturbations are required for the loop to operate (LDFC-requ1), science acquisitions can be done at full duty cycle.
- **Ease of calibration**: As LDFC is a linear control technique (LDFC-requ3), it uses derivatives of the pixel intensities relative to DM actuation for calibration. This calibration can be measured by DM probing in a reasonably short amount of time, so the technique is not as sensitive to modeling errors as EFC-like approaches that require a numerical model of the coronagraph system.
- **Linearity**: The linear control loop (LDFC-requ3) is fast to execute, and common linear analysis techniques can be deployed for optimization and analysis.
- **Scalability**: Multiple sensors (cameras at multiple wavelengths) can easily be integrated in a LDFC control scheme, as no dark hole is required and no DM probing is required.

LDFC limitations
- LDFC is a **differential** sensing technique that cannot by itself drive the system to a high contrast state. It is only a wavefront stabilization technique.
- A **null space** may exist: not all wavefront modes can be sensed, and some of the unseen modes can negatively impact contrast.
- **Non-stationarity** of the relationship between bright field, wavefront state and dark hole illumination can build up over time and render LDFC's calibration stale.

A key goal of our activity will be to quantify the null space and calibration non-stationarity, as these two effects are likely to limit LDFC's performance, and they are currently poorly understood.

1.2. LDFC for PSF Post-Processing

The BF component can also be used to estimate light distribution in the DF. The approach is an extension of the LDFC algorithm that focuses on active wavefront stabilization, and is meant to operate in parallel to it. As the relationship between BF and DF is quadratic, we refer to the algorithm as ***quadratic dark field calibration (QDFC)*** to distinguish it from LDFC.

Quadratic Dark Field Calibration (QDFC) for high contrast imaging is defined by the following key properties, many of them shared with LDFC :

- [QDFC-requ1] Speckle calibration does not require runtime DM probing
- [QDFC-requ2] Relies on focal plane images for calibration
- [QDFC-requ3] Leverages the quadratic relationship between sensing input (BF) and calibration region (DF)
- [QDFC-requ4] Uses pixels outside (in spatial and/or spectral dimensions) the high contrast area for calibration input

QDFC uses light outside the high contrast dark hole for sensing (QDFC-requ4). Here “outside” may refer to spatial or wavelength dimensions, or a combination of both, as shown in Fig. 2. This component (referred to as the bright field) responds linearly to small perturbations in wavefront errors, and the DF responds quadratically to wavefront complex amplitude. In the small aberration regime, there is therefore a quadratic relationship between BF and DF, in the form:

$$DF = |A\ (BF-BF_0)|^2 + DF_0$$

Where BF_0 is the reference bright field image (in the absence of wavefront aberrations), A is a complex amplitude matrix, and DF_0 is the reference dark field image.

2. MS#2: Spectral LDFC demonstration 10x raw contrast stability gain at $<10^{-7}$ raw contrast

2.1. Milestone Description and Rationale

Demonstrate a 10x suppression of injected disturbances by use of spectral LDFC stabilization in a dark hole with area covering at least 10 sq-λ/D and reaching a raw contrast (post-LDFC) level below 1e-7.

This is the second of a series of milestones aimed at validating the LDFC approach. The first milestone (MS1) demonstrated with spatial LDFC a 10x gain in raw contrast at a raw contrast level below 1e-5. In this second milestone, we aim to extend these results to **contrast level 100x deeper than specified MS#1** (and ~10x deeper than achieved for MS1 completion), and we also **extend the technique to include spectral LDFC**.

The 1e-7 contrast for MS#2 was chosen to be a significant improvement (100x) over MS#1 contrast, but is still well within the contrast routinely achieved on multiple high-contrast imaging testbeds in both vacuum and air. This raw contrast goal will allow the work to be focused on demonstrating contrast stabilization with LDFC rather than reaching deep contrast, which is achieved by separate techniques (such as electric field conjugation). The expected raw contrast for this milestone is slightly above the raw contrast level reached by the testbed in the absence of dynamic wavefront errors: MS#2 should therefore exceed the required 1e-7 raw contrast level in vacuum tests, while in-air tests should operate near 1e-7 raw contrast.

Measurements will be performed in the visible to near-IR wavelength range within 400-1700 nm for compatibility with vacuum and air testbeds. The DF and BF spectral bands will be non-overlapping, and sufficiently separated to ensure high efficiency spectral LDFC operation.

Other requirements for this milestone (10x gain factor, 10 sq-λ/D control area) are unchanged from MS#1.

2.2. Stability Conditions and Algorithm

For all measurements, the raw contrast in the scoring area shall be <1e-7 (average) both without dynamic errors prior to LDFC, and with dynamic errors plus LDFC. Wavefront errors may need to be applied to demonstrate the LDFC gain, in which case the exact same sequence of aberrations shall be introduced for the ON and OFF sequences to be compared.

In the OFF sequence, dynamic wavefront errors are added to the deformable mirror, and no LDFC active correction is performed. In the ON sequence, the same input wavefront disturbances are added, and LDFC wavefront control is performed. The required 10x contrast stabilization gain is measured between the two sequences.

2.3. Statistical Requirements

This section describes **requirements** associated with the measurement.

The raw contrast shall be measured as the surface brightness averaged over a fixed area covering at least 10 squ-λ/D. The scoring area shall be within the primary control region of the deformable mirror(s).

The measurement statistical requirements are as follows:
- Measurements must be performed over a set of >100 consecutive focal plane images.
- The reported ***per-frame raw contrast value*** will be the average raw contrast within the scoring area.
- The measurement noise on the per-frame raw contrast value, due to photon and readout noise, must be <1e-8 raw contrast in the LDFC ON frames. We note that the corresponding noise may be higher in the LDFC OFF frames due to added photon noise.
- The reported ***measurement set raw contrast value*** will be the average of the per-frame raw contrast between all 100 frames. No frame can be excluded from the measurement.

The >100 frames constitute a ***measurement set***.

A minimum of three measurement sets meeting the 1e-7 raw contrast requirement shall be acquired. For each set, a separate set must be taken in the same conditions with LDFC=OFF to measure the comparison contrast. For each of the 3 pairs, the ratio of average raw contrast values between OFF and ON shall exceed 10x.

2.4. Calibration and Measurement Steps

This section describes a likely (non-binding) sequence of calibrations and measurements to meet the statistical requirements detailed in the previous section. It is provided here for information, but is not part of MS#2 requirements.

Initial setup & Optical Configuration

The high contrast imaging system (HCIS) will be driven to produce a high contrast area in the focal plane at the *high contrast wavelength*, referred to as the *dark hole (DH),* using a non-LDFC iterative approach (EFC, speckle nulling).
The HCIS includes at least one deformable mirror, and includes optical masks necessary to reach high contrast. At least one focal plane imaging camera will acquire the required images. One or several WFC loop(s) may be deployed, independently of LDFC, to maintain the desired

raw contrast level. For example, a coronagraphic low-order WFS may be required to maintain alignment.

Images will be acquired at two or more wavelengths, sequentially (same camera, source wavelength or camera filter changed to switch wavelength) or simultaneously (by a dichroic splitter sending different wavelength channels to separate cameras or pixel areas). One of the imaging wavelengths will be the high contrast wavelength in which the dark hole is achieved. The other wavelength(s) will be used for wavefront sensing.

Contrast calibration

Images will be contrast-calibrated so that each frame can be scaled to a raw contrast map. Raw contrast is defined here as image surface brightness relative to the un-occulted peak surface brightness of the PSF core. In a Lyot-type coronagraph architecture, the un-occulted PSF core used as raw contrast reference is obtained with the Lyot stop(s) in and the focal plane mask out. Contrast calibration during the LDFC measurement sequences will be performed by measuring the surface brightness of fiducials, which could be added (by means of a pupil plane phase modulation) if the original coronagraphic PSF does not include sufficiently bright fiducials.

LDFC calibration

LDFC calibration consists of a LDFC reference and a LDFC response matrix. The reference is the nominal bright field (BF) intensity image that the LDFC loop is driven to maintain. In spectral LDFC, this BF is the image(s) acquired at wavelength(s) other than the high contrast imaging wavelength.
The response matrix encodes the derivative of the BF against each DM actuator. LDFC calibration is measured by poking the DM with modes, and recording the corresponding BF images.
LDFC calibration may be obtained as follows:

- A basis of LDFC DM control modes is defined. The number of such modes is ***n***.
- The LDFC reference array, a vector of ***m*** elements (number of pixels in the BF), is initialized to zero
- For each DM control mode index k=0...n-1, two BF images are acquired: one image with the DM control mode poked with a small positive amplitude +a, and one with an amplitude -a.
- For each DM control mode, the linear response is computed as (I(+a)-I(-a))/(2a) and written in the k-th column of the LDFC response matrix. The average of the two images (I(+a)+I(-a))/(2n) is additively added to the LDFC reference vector.

LDFC calibrations will be acquired with the system is in a high contrast state, after EFC or similar algorithm establishing a dark hole, and with no injected wavefront disturbances. Unintentional small wavefront variations due to inherent testbed stability, if problematic, will be mitigated by acquiring multiple calibrations and averaging them, and optimizing the calibration speed to mitigate drifts. Spectral LDFC calibrations will be acquired at the non-high contrast wavelength(s).

We have learned from previous work on spatial LDFC (including MS#1 demonstration) that efficiency and accuracy of calibrations are essential, and have developed and improved algorithms to do so. A key finding has been that calibrations should take into account the temporal response of testbed hardware (DM(s), camera(s)) to minimize calibration time, which is now part of the calibration acquisition.

LDFC control matrix computation

A regularized pseudo-inverse of the response matrix (RM) is computed and used as the linear control matrix (CM) for spectral LDFC operation.

LDFC loop operation

The BF is extracted from each camera image, and multiplied by the CM to produce an estimate of the DM displacement that should be additively applied to drive the BF to its reference. The corresponding DM command is multiplied by a control loop gain and then added to the DM. These steps are repeated in a loop of at least 100 iterations.

Injected Disturbances

LDFC milestone #1 tests have revealed that the performance is a function of injected disturbances: some disturbances map well with the LDFC control modes, while others overlap with LDFC's measurement null space (set of aberrations that create no LDFC signal). We will perform spectral LDFC demonstrations with several types of disturbances:

(a) Single LDFC control mode
(b) Multiple LDFC control modes (from 2 to the total number of modes)
(c) Single spatial frequency (speckle in focal plane)
(d) Random phase screen following power law

Disturbances (a) and (b) are expected to yield the highest performance and successful completion of MS#2 will require the stated performance goal to be reached for these injected disturbances.

Injected disturbances will be static between consecutive LDFC loop iterations. To test LDFC with multiple injected disturbances, we will change the disturbance prior to running a set number of LDFC loop iterations, chosen sufficiently large to ensure convergence and reach the 10x suppression factor goal. The measurement sequence envisioned for the test is shown in Fig 3.

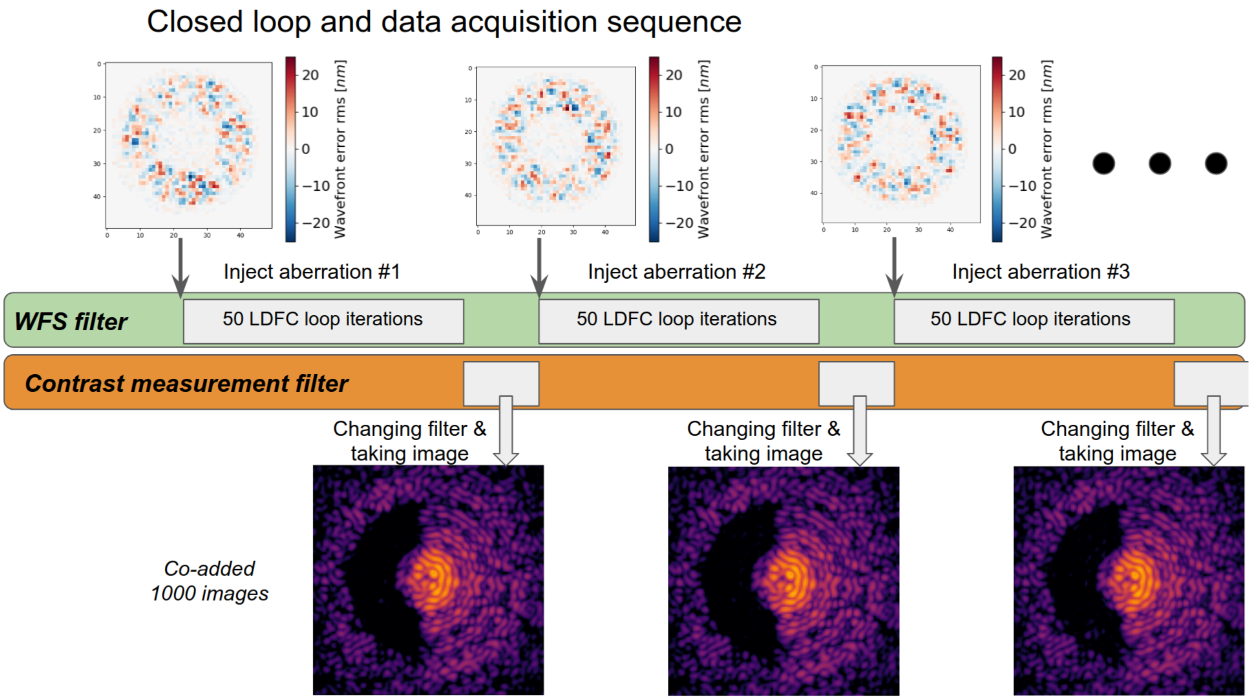

Fig 3: Spectral LDFC measurement sequence. Wavefront disturbances are injected and then corrected with a set number of LDFC loop iterations (50 iterations in this example). Performance measurement requires switching filter bandpass.

LDFC performance analysis and metrics

The performance metric for Milestone #2 completion is a suppression factor (10x) of scattered starlight. Temporal characteristics of the LDFC loop will also be measured and documented to help generalize MS#2 findings to a wider range of conditions:

- **Temporal characteristics of the injected disturbances**: rate of change in the sequence of wavefront aberrations (Fig 3, top row).
- **LDFC control loop parameters**: number of LDFC iterations performed between changes in the injected aberration and contrast measurements (N=50 iterations in the example shown in Fig 3), and the LDFC loop gain g.

These parameters will be varied to measure the dynamic rejection of wavefront aberrations provided by LDFC, and compare to expected values. Specifically, under ideal noiseless conditions, each injected wavefront aberration should be reduced by a factor $(1\text{-}g)^N$.

The temporal performance analysis will be included as part of the MS#2 completion report.

2.6. Test Plan and Schedule

Spectral LDFC tests will be conducted both in air and in vacuum. The milestone report will document performance achieved both in air and in vacuum, focusing on demonstration(s) that reach or exceed the milestone goals.

The in-air testbeds offer hardware flexibility and continuous access, allowing for the optical configuration to be optimized for this milestone, including simultaneous acquisition of multiple spectral bands. In-air tests are being conducted continuously through the period of effort, currently (Feb 2022) achieving a 6x contrast gain at 1e-6 contrast over a 500 sq l/D area (see Fig 1). The vacuum test, scheduled for summer 2022 at JPL, will probe a deeper contrast regime, and will implement a sequential approach where the source spectral band is alternating between the sensing band and the contrast measurement band.

Algorithms and software for this milestone are developed and tested on in-air testbed first, before being deployed for the vacuum test. The Hawaii (SCExAO) and Arizona (MagAO-X) testbeds already share a common software framework and data format, so LDFC code is common between the two systems.

The schedule for the vacuum test is as follows:
March-April 2022 : Migration of existing software and algorithms to JPL
May-July 2022 : Vacuum measurements (see section 2.4 for details)

3. MS#3: LDFC PSF calibration: 10x contrast gain to 1e-8 raw contrast

This milestone focuses on the PSF post-processing part of the LDFC effort, and seeks to demonstrate that the BF (spatial-BF, spectral-BF and/or spaspe-BF in Fig 1) can be used to estimate the DF component, such that the resulting estimate can then be numerically subtracted from the measured DF to improve detection limits.

3.1. Milestone Description and Rationale

The milestone definition is as follows:

Demonstrate a 10x gain in contrast by post-processing using LDFC telemetry, reaching a post-processed contrast level below 1e-8.

This is the third of a series of milestones aimed at validating the LDFC approach. Previous milestones focused on active wavefront control to improve raw contrast by spatial LDFC (MS#1) and spectral LDFC (MS#2). The goal of this third milestone is to validate use of LDFC signals for post-processing of a high contrast dark field.

The 10x ratio will bring a significant gain in exoplanet imaging, and can also relax some of the most challenging optics stability requirements imposed on high contrast imaging systems. Demonstrating that residual speckles can be calibrated will relax the mission raw contrast requirement, allowing for a wider spectral band to be used toward science acquisition, resulting in more efficient observing.

Contrast shall be evaluated in an **area of at least 10 sq-λ/D** (similarly to LDFC MS#1 and MS#2), at a wavelength within the 400 nm to 1700 nm range.

3.2. Telemetry for Post-Processing

The post-processing contrast gain is achieved by estimating the starlight component within the DF, and subtracting it from the measured DF. The ***Telemetry*** refers to input measurements used for this estimation. MS#3 will use the BF (spectral and/or spatial) as telemetry.

3.3. Definition of Contrast and Post-Processed Contrast

The ***contrast*** and ***post-processed contrast*** terms in the milestone definition correspond to a detection limit in the speckle noise regime, instead of the raw contrast level corresponding to physical surface brightness. In the context of this milestone, they are meant to quantify the effect of speckle noise on exoplanet detection.

Contrast is defined as the temporal standard deviation of raw contrast, measured over a sufficiently large set of images. This definition accurately captures how speckle noise limits exoplanet detection in individual exposures. For example, if the PSF raw contrast is 1e-6, and the residual starlight varies by 1% (relative), the contrast is 1e-8. In the extreme case where residual starlight is perfectly stable, contrast is 0, as there is no speckle noise.

This choice implicitly assumes that some form of PSF subtraction will be performed for exoplanet detection, either between images of the same star at different epochs (such as what is done for angular differential imaging), or between images of different stars (reference differential imaging).

Post-processed contrast is the temporal standard deviation of individually processed images. In the context of this milestone, a postprocessed image is a raw contrast image to which an estimate of the starlight component in the dark field is subtracted. The estimate is computed solely from the BF component(s) of the same image.

Important caveats:

- Contrast values, as defined here, do not accurately represent exoplanet detection limits in datasets consisting of a large number of images. Statistical properties, such as averaging to a common central limit between the science and reference images, are then highly relevant. We choose to only consider standard deviation over a set of images, assuming that a 10x gain in this standard deviation will translate into a similar gain in detection limit.
- Photon noise is not considered. Our experiments will be performed with sufficiently large flux to be in the speckle-noise limited regime. When required for the purpose of this milestone, we will remove (subtract) known photon noise variance to optimally quantify speckle noise.
- We do not explore or consider alternative post-processing schemes such as spectral differential imaging or polarimetric differential imaging.

3.4. Statistical Requirements and Success Criteria

Contrast shall be measured over a fixed scoring region covering at least 10 sq-λ/D. The scoring region shall be within the primary control region of the deformable mirror(s). Contrast, as defined in section 3, is computed for each pixel in the scoring region, and then averaged over all pixels in the scoring region, yielding the ***dataset contrast value.*** The **dataset post processed contrast value** is computed following the same steps, but by replacing input raw images with their post processed counterparts.

The dataset shall consist of >1000 focal plane images to ensure reliable estimation of standard deviation values.

MS#3 will be successfully completed if :

- The ratio of **dataset contrast value / dataset post processed contrast value** is **> 10**
- The **dataset post processed contrast value** is **< 1e-8**

3.5. Measurement Steps

This section describes the planned sequence of calibrations, measurements and algorithmic steps to meet the statistical requirements detailed in the previous section.

Initial setup

The high contrast imaging system (HCIS) will be driven to produce a high contrast area in the focal plane, referred to as the dark field (DF), using a non-LDFC iterative approach (EFC, speckle nulling). The HCIS includes at least one deformable mirror, and includes optical masks necessary to reach high contrast. At least one focal plane imaging camera will acquire images.

Contrast calibration

Images will be contrast-calibrated so that each frame can be scaled to a raw contrast map. Raw contrast is defined here as image surface brightness relative to the un-occulted peak surface brightness of the PSF core. In a Lyot-type coronagraph architecture, the un-occulted PSF core used as raw contrast reference is obtained with the Lyot stop(s) in and the focal plane mask out. Contrast calibration during the LDFC measurement sequences will be performed by measuring the surface brightness of fiducials, which could be added (by means of a pupil plane phase modulation) if the original coronagraphic PSF does not include sufficiently bright fiducials.

QDFC Algorithm Approach

The QDFC-based post-processing contrast calibration relies on a mapping from BF to DF, such that an estimate of the DF can be produced from each realization of the BF. This mapping will be implemented as a lookup table between input BF and output DF, to be populated by acquiring a training set of simultaneous BF+DF realizations.

Building the BF to DF mapping from a training set

A series of focal plane images (including both DF and BF) is recorded in the presence of wavefront aberrations. This is the training set, from which the mapping between BF and DF is measured and stored as a lookup table.
Due to the high dimensionality of the BF and DF images, dimensionality reduction approach(es) is(are) required to yield a practical representation of the mapping. One such approach is to use a clustering algorithm to group together similar BF images (as documented in “High contrast imaging at the photon noise limit with self-calibrating WFS/C systems”, Guyon et al, 2021, Proceedings of the SPIE).

The mapping is computed by running a high-speed clustering algorithm to group the training set BF images in small clusters. Each cluster consists of ~100 images with similar BF components. The average, over the cluster size, of the BF and DF are computed and correspond to a single entry in the mapping table between BF and DF. This approach is currently (Feb 2022) in use for QDFC tests in air, providing a 5.5x gain in contrast with random injected disturbances exploring the full dimensionality of the optical system (~1000 modes). The approach is suitable to handle a finite (~<10) number of modes with high efficiency for the MS#3 performance goals. Algorithm improvements aimed at reducing the training set size and more efficiently handling a large number of modes will be tested, but are probably not required to meet MS#3 goals. These include local linearization within clusters and among nearby clusters, and local derivation of the quadratic relationship between BF and DF.

QDFC validation dataset

The QDFC validation dataset will be separate from the training set. A series of >1000 focal plane images is recorded in the presence of dynamic wavefront aberrations. Wavefront aberrations may be added to the DM to increase contrast prior to pre-processing in order to meet the milestone requirements. If so, the added wavefront aberrations must be randomized as follows:

- A set of N>3 wavefront modes is adopted, common to all wavefront realizations
- A wavefront amplitude a is adopted, common to all wavefront realizations
- For each wavefront realization, N random values c_i uniformly distributed in the -a to +a interval are drawn.
- The wavefront realization is the linear sum of the N modes, with coefficients c_i.

The best dataset for the milestone will be acquired without LDFC control to ensure that wavefront variations are sufficiently large to demonstrate the required 10x post-processing gain. We will also process LDFC-controlled datasets and measure the post processing gain. We expect that QDFC will address the following shortcomings of a LDFC-maintained DF:

(a) LDFC closed loop suffers from temporal latency, and is unable to compensate for wavefront errors that evolve on a frame-by-frame timescale
(b) LDFC cannot sense and control aberrations that are not part of the calibration. These includes amplitude aberrations due to optics (or atmosphere) not conjugated to the DM used for LDFC calibration

3.6. Test Plan and Schedule

QDFC tests will be performed both in air and in vacuum. The milestone report will document performance achieved both in air and in vacuum, focusing on demonstration(s) that reach or exceed the milestone goals.

In-air tests are currently (Feb 2022) ongoing, having reached a 5.5x contrast gain (see “High contrast imaging at the photon noise limit with self-calibrating WFS/C systems”, Guyon et al, 2021, Proceedings of the SPIE). Ongoing improvements focused on handling the dynamical

range between bright field and dark field will further improve the in-air QDFC contrast performance (currently at 1e-7) and contrast gain. In-air testing will continue through CY2022, with a goal of reaching the MS#3 performance before the end of the calendar year.
Vacuum tests will be conducted in summer 2022 at JPL, concurrently with MS#2 tests and using the same dataset. Vacuum testing provides access to higher contrast, while air tests allow for testing with a wider choice of optical configurations and easier acquisition of large training sets for QDFC.

4. Reporting Requirements

The milestone data package will include a narrative report discussing how each element of the milestone was met, including appropriate tables and summary charts.

All raw data and calibrations will be made publicly available for independent review, for a period covering at least the duration of this effort. All analysis tools will also be publicly available for review and for use by other groups. Source code will be released under (a) open-source license(s) to the extent possible, in compliance with US regulations.