

On distributed wavefront reconstruction for large-scale adaptive optics systems

de Visser, Coen; Brunner, Elisabeth; Verhaegen, Michel

DOI

[10.1364/JOSAA.33.000817](https://doi.org/10.1364/JOSAA.33.000817)

Publication date

2016

Document Version

Accepted author manuscript

Published in

Journal of the Optical Society of America A: Optics and Image Science, and Vision

Citation (APA)

de Visser, C., Brunner, E., & Verhaegen, M. (2016). On distributed wavefront reconstruction for large-scale adaptive optics systems. *Journal of the Optical Society of America A: Optics and Image Science, and Vision*, 33(5), 817-831. <https://doi.org/10.1364/JOSAA.33.000817>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

On Distributed Wavefront Reconstruction for Large Scale Adaptive Optics Systems

Cornelis C. de Visser^{1,*}, Elisabeth Brunner², and Michel Verhaegen²

¹*Department of Control & Simulation, Delft University of Technology,
Kluyverweg 1, 2629 HS, Delft, The Netherlands*

²*Delft Center for Systems and Control, Delft University of Technology,
Mekelweg 2, 2628 CD, Delft, The Netherlands*

**Corresponding author: c.c.devisser@tudelft.nl*

The D-SABRE (Distributed Spline based ABeration REconstruction) method is proposed for distributed wavefront reconstruction with application to large scale adaptive optics systems. D-SABRE decomposes the wavefront sensor domain into any number of partitions and solves a local wavefront reconstruction problem on each partition using multivariate splines. D-SABRE reconstruction accuracy is within 1% of a global approach with a speedup that scales quadratically with the number of partitions. The D-SABRE is compared to the CuRe-D method in open-loop and closed-loop simulations using the YAO adaptive optics simulation tool. D-SABRE reconstruction accuracy exceeds Cure-D for low levels of decomposition, and D-SABRE proved to be more robust to variations in the loop gain.

© 2016 Optical Society of America

OCIS codes: 010.1080, 010.7350, 010.1285, 000.3860, 000.4430, 350.1260.

1. Introduction

In the coming decade, a new generation of extremely large scale optical telescopes will see first light. It is well known that increasing the size of the telescope aperture is only beneficial if the adaptive optics (AO) system, which compensates for turbulence induced wavefront aberrations, scales accordingly [16]. In particular, the total number of wavefront sensors (WFS) and deformable mirror (DM) actuators required in an AO system to obtain a given image quality is determined by the telescope diameter and the Fried coherence length [6]. To

update the DM actuator commands such that the incoming aberrated wavefront is corrected the global wavefront phase has to be reconstructed from the WFS measurements at kHz range frequencies.

Most conventional wavefront reconstruction (WFR) methods, like the finite difference (FD) method [6, 10, 19], can be formulated as a matrix vector multiplication (MVM) in which a reconstruction matrix is first constructed offline, and then multiplied online with a vector of WFS measurements to obtain the unknown wavefront phase. The computation of the reconstruction matrix in its most naive form is an $O(N^3)$ operation, where N is the number of unknown phase samples. Applying the real-time MVM scales with $O(N^2)$. For the extreme-AO (XAO) system of the future European Extremely Large Telescope (E-ELT) the number of unknowns N is in the range of $10^4 - 10^5$. Current real-time performance of standard methods fails to meet the required update frequency for such systems. Hence, there has been a focus on improving the computational efficiency of the WFR operation.

Currently, one of the most computationally efficient zonal reconstruction methods is the Cumulative Reconstruction (CuRe) method, which is of complexity order between $O(12N)$ and $O(19N)$, depending on the implementation [17]. Fourier domain methods are an efficient alternative to the MVM methods discussed above, because the partial differential equations that relate the wavefront slopes to the wavefront itself can be reduced to division schemes in the complex plane. The Fourier domain method presented by Poyneer et al. in [13] has a time efficiency of $O(N \log N)$. The Haar wavelet based reconstruction method by Hampton et al. reach efficiencies of the order $O(10N)$ [9]. In the class of iterative methods, the complexity of a multigrid PCG algorithm presented by [7] scales with $O(N \log N)$. The Fractal Iterative Method (FrIM) by Thiébaud and Tallon [20] provides a minimum variance solution of linear complexity with the computational cost depending on the number of iterations.

While the state of the art in high performance WFR methods reach linear computational complexity orders, it is important to note that these numbers are WFS-array wide, or *global*, numbers. Even with a linear complexity order, the limit in single CPU core performance will be reached at some point. The disadvantage of current wavefront reconstruction methods is that they have not been designed specifically for parallel processing architectures. While certain computational operations can be parallelized straightforwardly (e.g. MVM multiplications can easily be distributed over multiple CPU cores), their underlying construction is that of a non-distributed global method. As a direct result of this, the reconstructor can only be calculated at a single central point at which all WFS measurements must be gathered and at which the DM influence matrix is calculated. A tell-tale sign of any non-distributed global method is that the global reconstruction matrix is a dense matrix.

Recently, an extension of the CuRe algorithm [18] was proposed which combines the original line integral approach with domain decomposition to tackle the high noise propagation

of the centralized algorithm. This distributed version of CuRe, called CuRe-D, scales with $O(20N)$ and has been shown to be suitable for parallel implementation.

The main contribution of this paper is a distributed wavefront reconstruction method which is designed in particular for use in XAO systems. This new method is the Distributed Spline Based Aberration Reconstruction (D-SABRE) method. The D-SABRE method is an extension of the recently introduced SABRE method which firstly used multivariate simplex B-splines to locally model wavefront aberrations on non-rectangular WFS arrays [5]. The D-SABRE is based on a decomposition of the global wavefront sensor domain into any number of triangular partitions, where each partition supports a local SABRE model which depends only on local WFS measurements.

The D-SABRE method is a two-stage method in which each stage is a distributed operation. In the first stage, a local wavefront reconstruction problem is solved in parallel on each triangular partition using local WFS slopes, resulting in a local SABRE model. Each SABRE model has an unknown piston mode, and as a result, there is no continuity between SABRE models on neighboring partitions. In the second stage of the D-SABRE, a continuous global wavefront is reconstructed by equalizing the piston modes of the local SABRE models using a new distributed piston mode equalization (DPME) algorithm. Additionally, the dual ascent method from [1] is reformulated into a new efficient distributed form by exploiting the inherent sparseness of the D-SABRE reconstruction and constraint matrices.

The advantages of the D-SABRE over current distributed wavefront reconstruction methods can be summarized as follows:

1. D-SABRE is based on a local least-squares estimates and has in this sense locally optimal noise rejection.
2. D-SABRE does *not* suffer from noise propagation, and the accuracy and noise resilience actually improve on a global scale as the size of the partitions increases and as the global WFS array increases in size.
3. Certain parallel hardware such as a GPU fulfill their potential speedup only for a sufficiently large computational task per processor i.e. partitions cannot be chosen too small. This requirement is completely in line with the fact that D-SABRE accuracy and noise resilience increases with increasing partition size.
4. The D-SABRE wavefront is an analytic solution to the wavefront reconstruction problem. Without any further interpolation leading to additional approximation errors phase estimates are available at any location in the WFS domain. This is an advantage in case of misalignments between the actuator and subaperture arrays or if a different actuator distribution is given.

The D-SABRE currently presented should be seen as a baseline, or simplest possible version. Future versions will include higher degree (e.g. cubic) splines which will require fewer subapertures and further increase accuracy, more advanced sensor models like that presented in [8], and more advanced estimators such as the minimum variance estimator that exploit a-priori knowledge of the turbulence and wavefront statistics.

This paper is outlined as follows. First, we provide brief preliminaries on the SABRE method for wavefront reconstruction in Sec. 2. We introduce the new D-SABRE method in Sec. 3, and also provide analyses of algorithm convergence and computational complexity. Additionally, we present a tutorial example of the D-SABRE in Sec. 3. In Sec. 4, we present the results from a numerical validation where we compare the D-SABRE to the global SABRE method and the distributed CuRe-D method [18] in open- and closed-loop simulations. Finally, conclusions are provided in Sec. 5.

2. Preliminaries on the SABRE method for wavefront reconstruction

The D-SABRE method is an extension of the recently introduced SABRE method for wavefront reconstruction. In order to aid the reader in the understanding of the theory, preliminaries on the SABRE method will be provided in this section. The SABRE was first introduced in [5], to which we refer for a more in-depth coverage of the matter.

2.A. Wavefront reconstruction from slopes and curvature

The relationship between the slopes of the wavefront phase and the wavefront phase can be described in the form of the following system of first order partial differential equations [10]:

$$\sigma_x(x, y) = \frac{\partial \phi(x, y)}{\partial x}, \quad (1a)$$

$$\sigma_y(x, y) = \frac{\partial \phi(x, y)}{\partial y}, \quad (1b)$$

with $\phi(x, y)$ the unknown wavefront, and with $\sigma_x(x, y)$ and $\sigma_y(x, y)$ the wavefront slopes at location (x, y) in the directions x and y , respectively.

2.B. The SABRE method on a single triangle

The local model elements of a SABRE model are defined on individual triangles, rather than on the rectangular elements used by FD methods. In [5] it is shown that on a single triangle, denoted t , the wavefront phase $\phi(x, y)$ is approximated with a SABRE model of degree d as follows:

$$\phi(x, y) \approx \mathbf{B}^d(b(x, y)) \cdot \mathbf{c}^t, \quad d \geq 1, \quad (x, y) \in t \quad (2)$$

with $\mathbf{B}^d(b(x, y))$ the vector of basis polynomials and with \mathbf{c}^t the vector of B-coefficients. The values in the vector of basis polynomials $\mathbf{B}^d(b(x, y))$ depend only on the geometry of

the sensor array and the polynomial degree d . Therefore, at any given time instant, the particular form of the SABRE model is determined by the B-coefficients \mathbf{c}^t . Consequently, wavefront reconstruction with the SABRE essentially consists of estimating the values of the B-coefficients given a set of WFS measurements.

It is shown in [5] that (2) leads to the following SABRE slope sensor model on a single triangle:

$$\sigma_x(x, y) = \frac{d!}{(d-1)!} \mathbf{B}^{d-1}(b(x, y)) \mathbf{P}^{d,d-1}(\mathbf{a}_x) \cdot \mathbf{c}^t + n_x(x, y), \quad (3a)$$

$$\sigma_y(x, y) = \frac{d!}{(d-1)!} \mathbf{B}^{d-1}(b(x, y)) \mathbf{P}^{d,d-1}(\mathbf{a}_y) \cdot \mathbf{c}^t + n_y(x, y), \quad (3b)$$

with $d \geq 1$ the degree of the SABRE model, with $\mathbf{B}^{d-1}(b(x, y))$ the basis polynomials of degree $d-1$, with \mathbf{c}^t the vector of B-coefficients from (2), and with $n_x(x, y)$ and $n_y(x, y)$ in (3) residual terms which contain both sensor noise and modeling errors.

The matrices $\mathbf{P}^{d,d-1}(\mathbf{a}_x)$ and $\mathbf{P}^{d,d-1}(\mathbf{a}_y)$ in (3) are the *de Casteljau matrices* in the (Cartesian) directions \mathbf{a}_x and \mathbf{a}_y which are essential to the SABRE method [4]. In essence, the de Casteljau matrices allow a reformulation of the PDE's from (1) into a set of algebraic equations in terms of the B-coefficients \mathbf{c}^t .

A number of possible sensor geometries were introduced in [5]. In this work we use the Type-1 and Type-2 sensor geometries, which should be seen as baseline geometries. Recently, a more advanced sensor model for the SABRE was introduced by Guo et al in [8]. This sensor model is better suited for use with real-world SH sensors which provide a spatial average of the wavefront slopes instead of point-wise local spatial derivatives of the wavefront.

2.C. The SABRE method on a complete triangulation

The SABRE method can be used with large scale wavefront sensor arrays by combining any number triangles into a triangulation. The full-triangulation, or *global*, SABRE model has a predefined continuity order r between the local models which means that the r^{th} order directional derivatives of neighboring local SABRE models match exactly on the triangle edges. In [5] it is shown that the wavefront phase can be approximated at any point (x, y) in the WFS domain with a SABRE model as follows:

$$\phi(x, y) \approx \mathbf{B}^d \mathbf{c}, \quad (4)$$

with \mathbf{B}^d the global B-form regression matrix and with \mathbf{c} the global vector of B-coefficients [5]. Given (4), the global WFR problem is constructed from rows of the form (3) as follows:

$$\sigma = \mathbf{Dc} + \mathbf{n}, \quad (5a)$$

$$0 = \mathbf{Ac}, \quad (5b)$$

with $\boldsymbol{\sigma} = [\boldsymbol{\sigma}_x^\top \ \boldsymbol{\sigma}_y^\top] \in \mathbb{R}^{2K \times 1}$ the vector of measured wavefront slopes, with \mathbf{n} a residual noise vector, and with \mathbf{A} the global constraint matrix. The spline regression matrix \mathbf{D} in (5a) is constructed as follows:

$$\mathbf{D} = d\mathbf{B}^{d-1}\mathbf{P}_{\mathbf{u}}^{d,d-1}, \quad (6)$$

with \mathbf{B}^{d-1} the global basis function matrix of degree $d - 1$, and with $\mathbf{P}_{\mathbf{u}}^{d,d-1}$ the full-triangulation de Casteljau matrix which is constructed as shown in. Note the slight difference in notation of \mathbf{D} in (6) compared to that presented in [5]; it will become apparent in following sections why this change was made.

The global constraint matrix \mathbf{A} in (5b) is constructed as follows:

$$\mathbf{A} := \begin{bmatrix} \mathbf{H} \\ \mathbf{h} \end{bmatrix} \in \mathbb{R}^{(EV+1) \times J\hat{d}}, \quad (7)$$

with $\mathbf{H} \in \mathbb{R}^{(EV) \times J\hat{d}}$ the full-rank *smoothness matrix*. The vector $\mathbf{h} = [1 \ 0 \ \dots \ 0] \in \mathbb{R}^{1 \times J\hat{d}}$ is the *anchor vector* which was first introduced in [5], and which is used to fix the piston mode (the unknown integration constant) to a predetermined constant.

For low degree ($d \leq 2$) basis polynomials, the resulting system (5) will be fully determined given a Type-1 or Type-2 SH sensor geometry and the constraints in (7). However, when moving towards higher degree ($d \geq 3$) basis polynomials (5) will be underdetermined, and will not lead to a unique solution. Future work on obtaining higher degree D-SABRE models will therefore be focused on imposing specific (e.g. "do-nothing" boundary conditions [14]) boundary conditions on the external edges of the D-SABRE submodels.

In [5] it was shown that the global wavefront reconstruction problem can be formulated as an equality constrained least squares optimization problem:

$$\arg \min \frac{1}{2} \|\boldsymbol{\sigma} - \mathbf{D}\mathbf{c}\|_2^2 \quad \text{subject to } \mathbf{A}\mathbf{c} = 0, \quad (8)$$

with $\boldsymbol{\sigma}$ the slopes from the wavefront sensor, with \mathbf{D} from (6), with \mathbf{c} the global vector of B-coefficients, and with \mathbf{A} the smoothness matrix from (7).

This constrained optimization problem can be reduced into a unconstrained problem by introducing the Lagrangian for (8):

$$\mathcal{L}(\mathbf{c}, \mathbf{y}) = \frac{1}{2} \|\boldsymbol{\sigma} - \mathbf{D}\mathbf{c}\|_2^2 + \mathbf{y}^\top \mathbf{A}\mathbf{c}, \quad (9)$$

with \mathbf{y} a vector of Lagrangian multipliers.

The minimum of (9) is:

$$\frac{\partial \mathcal{L}(\mathbf{c}, \mathbf{y})}{\partial \mathbf{c}} = -\mathbf{D}^\top (\boldsymbol{\sigma} - \mathbf{D}\mathbf{c}) + \mathbf{A}^\top \mathbf{y} \stackrel{!}{=} 0, \quad (10)$$

with which the following B-coefficient estimator is derived:

$$\hat{\mathbf{c}} = (\mathbf{D}^\top \mathbf{D})^{-1} (\mathbf{D}^\top \boldsymbol{\sigma} - \mathbf{A}^\top \mathbf{y}). \quad (11)$$

The problem with (11) is that the undefined piston mode causes $\mathbf{D}^\top \mathbf{D}$ to be rank deficient, and therefore not invertible. In [5] this issue was solved by immediately projecting the spline regressors on the nullspace of the constraint matrix \mathbf{A} .

Here we require a more explicit formulation:

$$\begin{aligned} \hat{\mathbf{c}} &= \mathcal{N}_A (\mathcal{N}_A^\top \mathbf{D}^\top \mathbf{D} \mathcal{N}_A)^{-1} \mathcal{N}_A^\top \mathbf{D}^\top \boldsymbol{\sigma}, \\ &= \mathbf{R}_N \mathbf{D}^\top \boldsymbol{\sigma}, \\ &= \mathbf{Q}_N \boldsymbol{\sigma}, \end{aligned} \quad (12)$$

with \mathcal{N}_A a basis for $\text{null}(\mathbf{A})$, with $\mathbf{R}_N = \mathcal{N}_A (\mathcal{N}_A^\top \mathbf{D}^\top \mathbf{D} \mathcal{N}_A)^{-1} \mathcal{N}_A^\top$ and with $\mathbf{Q}_N = \mathbf{R}_N \mathbf{D}^\top$ the SABRE reconstruction matrix from [5].

3. Distributed WFR with simplex B-splines

In this section the D-SABRE distributed wavefront reconstruction method is introduced. The D-SABRE consists of 2 stages; with Stage-1 the distributed local WFR stage, and Stage-2 the combined distributed piston mode equalization and distributed dual ascent post-smoothing stage. A schematic of the D-SABRE algorithm is shown in Fig. 1.

3.A. D-SABRE Stage-1: Distributed local WFR

The first stage of D-SABRE involves the decomposition of the global WFR problem from (8) into a set of local sub-problems. For this we make use of the domain decomposition method from [12]. First, we decompose the global triangulation \mathcal{T} into a set of G sub-triangulations as follows:

$$\mathcal{T} = \bigcup_{i=1}^G \mathcal{T}_i \quad (13)$$

with each \mathcal{T}_i containing J_i triangles.

Every sub-triangulation \mathcal{T}_i in turn consists of two parts; a core part Ω_i and an overlap part Ξ_i :

$$\mathcal{T}_i = \Omega_i \cup \Xi_i, \quad \Omega_i \cap \Xi_i = \emptyset \quad (14)$$

with Ω_i containing J_{Ω_i} triangles and Ξ_i containing J_{Ξ_i} triangles such that $J_i = J_{\Omega_i} + J_{\Xi_i}$.

The purpose of Ξ_i is to overlap neighboring sub-triangulation core parts \mathcal{T}_j in order to increase numerical continuity between neighboring partitions, see Fig. 2. In this paper we use the term ‘*Overlap-Level*’ (OL) to define the size of Ξ_i . The OL is a scalar, which determines how many layers of simplices from the core partition are included in Ξ_i , see Fig. 2.




Fig. 1: Outline of the D-SABRE algorithm.

We can now introduce the *overlap overhead* ρ as follows:

$$\rho = \frac{J_i}{J_{\Omega_i}}, \rho \geq 1 \quad (15)$$

with $\rho = 1$ indicating no overlap. It is not trivial to provide a relation between overlap level and overlap overhead because it strongly depends on the geometry of the triangulation and on the location of a partition within the global triangulation. Nevertheless, for a simplex Type-1 geometry, the minimum and maximum ρ can easily be determined using basic geometry rules:

$$\rho_{Type-1} = \begin{cases} 1 + \frac{OL(6OL+8\sqrt{J_{\Omega_i}/2})}{J_{\Omega_i}} & (\text{max overhead}), \\ 1 + \frac{OL(OL+4\sqrt{J_{\Omega_i}/2})}{J_{\Omega_i}} & (\text{min overhead}), \end{cases} \quad (16)$$

with $OL \geq 0$ the overlap level.


We now assume that the wavefront can be approximated locally on each sub triangulation \mathcal{T}_i as follows:

$$\phi_i(x, y) \approx s_{r_i}^d(x, y) = \mathbf{B}_i^d \mathbf{c}_i, \quad 1 \leq i \leq G \quad (17)$$

with $\phi_i(x, y)$ the local wavefront phase, with $s_{r_i}^d(x, y)$ a D-SABRE partition, with \mathbf{B}_i^d the local matrix of B-form regressors, and with \mathbf{c}_i the set of local B-coefficients.

The goal now is to determine \mathbf{c}_i for all G local models. For each sub-triangulation \mathcal{T}_i we

4 Partitions with OL-0 overlap



4 Partitions with OL-2 overlap (gray)




Fig. 2: An OL-0 partitioning without overlap (left) and an OL-2 partitioning with 2 levels of overlap between partitions (right) using the same initial triangulation containing 200 triangles. In this case $J_{\Omega_i} = 50$ for all partitions, while $J_{E_1}, J_{E_4} = 44$ and $J_{E_2}, J_{E_3} = 48$ for the OL-2 partitioning

can decompose (8) into a set of local sub-problems as follows:

$$\arg \min \|\sigma_i - \mathbf{D}_i \mathbf{c}_i\|_2^2 \quad \text{subject to } \mathbf{A}_i \mathbf{c}_i = 0, \quad 1 \leq i \leq G \quad (18)$$

with σ_i local WFS slopes, and with \mathbf{A}_i a *local* constraint matrix, and with \mathbf{D}_i the local version of (6) as follows:

$$\mathbf{D}_i = d\mathbf{B}_i^{d-1} \mathbf{P}_{\mathbf{u}_i}^{d,d-1} \quad (19)$$

with $\mathbf{P}_{\mathbf{u}_i}^{d,d-1}$ the local de Casteljau matrix.

Note that the local constraint matrix \mathbf{A}_i cannot be obtained by decomposing the global constraint matrix \mathbf{A} into G blocks because it is not block diagonal. Instead, \mathbf{A}_i is constructed from a new local smoothness matrix \mathbf{H}_i and a new local anchor constraint h_i :

$$\mathbf{A}_i := \begin{bmatrix} \mathbf{H}_i \\ \mathbf{h}_i \end{bmatrix} \in \mathbb{R}^{(E_i V_i + 1) \times J_i \hat{d}}, \quad (20)$$

where it is important to note that \mathbf{H}_i does *not* contain any smoothness conditions linking a partition i to any other partition. For the anchor vector we have $\mathbf{h}_i = [1 \ 0 \ \dots \ 0] \in \mathbb{R}^{1 \times J_i \hat{d}}$.

The local WFR problem (18) can be solved in the same fashion as (12):

$$\begin{aligned} \hat{\mathbf{c}}_i^{loc} &= \mathcal{N}_{A_i} (\mathcal{N}_{A_i}^\top \mathbf{D}_i^\top \mathbf{D}_i \mathcal{N}_{A_i})^{-1} \mathcal{N}_{A_i}^\top \mathbf{D}_i^\top \sigma_i, \\ &= \mathbf{R}_{N_i} \mathbf{D}_i^\top \sigma_i, \\ &= \mathbf{Q}_{N_i} \sigma_i, \end{aligned} \quad (21)$$

with \mathcal{N}_{A_i} a basis for $\text{null}(\mathbf{A}_i)$, and with $\mathbf{R}_{N_i} = \mathcal{N}_{A_i}(\mathcal{N}_{A_i}^\top \mathbf{D}_i^\top \mathbf{D}_i \mathcal{N}_{A_i})^{-1} \mathcal{N}_{A_i}^\top$.

Note that (21) does not depend on information from any other partition, and as a result, each of the G local reconstruction problems can be solved in parallel. In Fig. 3 an example is shown of the results of the first distributed stage of the D-SABRE. Clearly, the local models are disconnected and do not accurately approximate the global wavefront. This is a direct result of the anchor constraint in each partition being independent from that of neighboring partitions.

3.B. Distributed piston mode equalization

In order to equalize the phase offsets (i.e. unknown integration constants) between neighboring D-SABRE partitions, a process called piston mode equalization (PME) is introduced. PME effectively resolves the difference between the estimated per-partition (unknown) integration constants. PME offsets an entire D-SABRE partition with a single constant k_i such that the maximum phase offset between it and a neighboring partition is minimized:

$$\tilde{s}_{r_i}^d(x, y) = \mathbf{B}_i^d \mathbf{c}_i + k_i, \quad 1 \leq i \leq G, \quad k_1 = 0. \quad (22)$$

where the tilde indicates that an offset k_i as been applied to $s_{r_i}^d(x, y)$.

The first partition (i.e. $s_{r_1}^d(x, y)$) has $k_1 = 0$ per definition. This partition is indicated as the *master partition*, relative to which all other partitions are equalized. The D-SABRE user is free to choose the location of the master partition, but a smart choice is a partition that is located as close as possible to the center of the global triangulation.

In order to calculate k_i a definition of the concept of neighboring partitions is required. For this, we define the neighbors of $s_{r_i}^d(x, y)$ as all partitions $s_{r_m}^d(x, y)$ with $1 \leq m \leq G$ which share at least one vertex $v_{i,m}$ in the core parts of their triangulations. We then define \mathcal{M}_i as the set that contains the indices of all neighbors of partition i , with $G_i = |\mathcal{M}_i|$ the total number of neighbor partitions.

The phase offset between partition i and m is calculated from the phase offset at the shared vertex $v_{i,m}$. A unique property of the simplex B-splines is that spline function value on a vertex is equal to the value of the B-coefficient located at that vertex [5, 11]. On the shared vertex $v_{i,m}$ we have $s_{r_i}^d(v_{i,m}) = \mathbf{c}_{\Omega_i}$ and $s_{r_m}^d(v_{i,m}) = \mathbf{c}_{\Omega_m}$. Hence, the estimated phase offset $\hat{k}_{i,m}$ between partition i and m is:

$$\hat{k}_{i,m} = \mathbf{c}_{\Omega_i} - \mathbf{c}_{\Omega_m}, \quad m \in \mathcal{M}_i. \quad (23)$$

Note that in the presence of noise, the estimate of $k_{i,m}$ can easily be improved by taking the average of the phase offsets at any number of shared vertices. If partition i has a total of G_i direct neighbors, the PME constant k_i for partition i is determined as follows:

$$k_i = \max \{ \hat{k}_{i,m} \}, \quad \forall (m > i) \in \mathcal{M}_i \quad (24)$$




Fig. 3: Original wavefront (top); The D-SABRE model after completion of the distributed local reconstruction stage.




Fig. 4: 4-step (A,B,C,D) DPME operation on 6 partitions with G_1 the master partition. Gray & black arrows indicate information flow, with black arrows the actual information used in a DMPE step.

where it should be noted that k_i is the *maximum* offset between partition i and m with $m > i$; this asymmetry is required for the PME operation to converge.

In the form of (23), PME is a sequential operation. However, it is straightforward to modify (23) into a distributed consensus problem form we indicate as distributed piston mode equalization (DPME):

$$\hat{k}_{i,m}(l+1) = \mathbf{c}(l)_{\Omega_{i,m}} - \mathbf{c}(l)_{\Omega_{m,i}}, \quad m \in \mathcal{M}_i, \quad l = 1, 2, \dots, L, \quad (25a)$$

$$k_i(l+1) = \max \{ \hat{k}_{i,m}(l+1) \}, \forall (m > i) \in \mathcal{M}_i, \quad l = 1, 2, \dots, L. \quad (25b)$$

which converges when $l = L$.

In Fig. 4 the concept of D-PME is demonstrated. In essence, each partition continuously adapts its offset based on data obtained from its direct neighbors which are themselves continuously adapting their offsets.

At each iteration, the set $\mathcal{K}_i^{(l)}(m)$ combines all G_i phase offsets between partition i and partitions m for iteration l :

$$\mathcal{K}_i^{(l)}(m) = \{ \hat{k}_{i,m}(l) \}, \quad 1 \leq i \leq G, \forall m \in \mathcal{M}_i. \quad (26)$$

This set will prove to be instrumental during the distributed dual ascent stage of D-SABRE.

In Fig. 5 the DPME method has been applied to the discontinuous reconstructed wavefront from Fig. 3.




Fig. 5: D-SABRE model after completion of DPME stage, with remaining discontinuities exaggerated (top), see also **Visualization 1**; D-SABRE model after completion of the distributed post-smoothing stage (bottom), see also **Visualization 2**.

3.C. Distributed post-smoothing

DPME minimizes the unknown phase offsets between local D-SABRE partitions, but in general this does not result in a smooth reconstruction of the wavefront, especially in the presence of (sensor) noise. Smoothness can be improved by increasing the overlap level as discussed in Sec. 3.A, but only up to some point and at the cost of reduced computational efficiency. In order to obtain a smooth reconstruction a distributed dual ascent (DDA) method is used to enforce smoothness between local D-SABRE partitions.

The derivation of the DDA smoother starts with the decomposition of the global Lagrangian from (9) into sub-Lagrangians according to [1]:

$$\begin{aligned}\mathcal{L}(\mathbf{c}, \mathbf{y}) &= \sum_{i=1}^G \mathcal{L}_i(\mathbf{c}_i, \mathbf{y}), \\ &= \sum_{i=1}^G \left(\frac{1}{2} \|\sigma_i - \mathbf{D}_i \mathbf{c}_i\|_2^2 + \mathbf{y}^\top \mathbf{G}_i \mathbf{c}_i \right),\end{aligned}\quad (27)$$

with \mathbf{y} the *global* vector of Lagrange duals, \mathbf{c}_i a sub-vector of local B-coefficients, and with \mathbf{G}_i a naive partitioning of the global constraint matrix \mathbf{A} :

$$\mathbf{A} = \begin{bmatrix} \mathbf{G}_1 & \mathbf{G}_2 & \cdots & \mathbf{G}_G \end{bmatrix}. \quad (28)$$

Dual decomposition of (27) leads to the following algorithm [1]:

$$\mathbf{c}_i(l+1) = \operatorname{argmin}_{\mathbf{c}_i} \mathcal{L}_i(\mathbf{c}_i, \mathbf{y}(l)), \quad (29a)$$

$$\mathbf{y}(l+1) = \mathbf{y}(l) + \alpha(l) \mathbf{A} \mathbf{c}(l), \quad (29b)$$

with $\alpha(l)$ an iteration dependent constant and with l the dual decomposition iterator.

The dual update step in (29b) clearly is a centralized operation as it requires a gathering and then broadcasting of global variables (i.e. $\mathbf{y}(l)$ and $\mathbf{c}(l)$).

By exploiting the sparseness structure of the global constraint matrix \mathbf{A} a more efficient formulation can be obtained. The structure of \mathbf{A} is as follows [2, 3]:

$$\begin{bmatrix} \mathbf{A}_{1,1} & 0 & 0 & \cdots & \mathbf{A}_{1,\star} & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_{2,2} & 0 & \cdots & \mathbf{A}_{2,\star} & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \mathbf{A}_{3,1} & 0 & 0 & \cdots & 0 & \mathbf{A}_{3,\star} & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & & & & \ddots & & & & & & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & \mathbf{A}_{EQ,\star} & 0 & \cdots & \mathbf{A}_{EQ,J} \\ \mathbf{h} & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \\ \mathbf{c}_3 \\ \vdots \\ \mathbf{c}_J \end{bmatrix} = 0, \quad (30)$$

with each block $\mathbf{A}_{q,j} \in \mathbb{R}^{Q \times \hat{d}}$ containing all the smoothness conditions for a single edge q and triangle j . Note that for each edge q , there are exactly two blocks $\mathbf{A}_{q,j}$; one for each triangle j containing the mutual edge q .

By exploiting the sparseness of \mathbf{A} we can derive two new submatrices \mathbf{H}_i and $\mathbf{H}_{i,M}$ for each partition i . The submatrix \mathbf{H}_i contains all smoothness conditions that define continuity inside partitions i . The submatrix $\mathbf{H}_{i,M}$ on the other hand contains all smoothness conditions that govern not only continuity inside partition i , but also all continuity between partitions i its neighbors. In Appendix 6 the details of both decompositions are presented.

Using the submatrices \mathbf{H}_i and $\mathbf{H}_{i,M}$ a local distributed version of (29) is obtained:

$$\mathbf{c}_i(l+1) = \hat{\mathbf{c}}_i^{loc} + \mathbf{R}_{N_i} \mathbf{H}_i^\top \mathbf{y}_i(l) \quad (31a)$$

$$\mathbf{y}_i(l+1) = \mathbf{y}_i(l) + \alpha(l) \mathbf{H}_{i,M} \mathbf{c}_{i,M}(l+1) \quad (31b)$$

with $\hat{\mathbf{c}}_i^{loc}$ the locally estimated B-coefficients from (21) which during the iteration are constant, with \mathbf{R}_{N_i} the matrix from (21), and with \mathbf{y}_i the local dual vector. The dual update step (31b) uses $\mathbf{c}_{i,M}$ which is the vector of all B-coefficients in $s_{r_i}^d$ as well as all B-coefficients in $s_{r_m}^d$ that are subject to continuity conditions on $\mathcal{T}_i \cap \mathcal{T}_m, \forall m \in \mathcal{M}_i$.

The resulting method is fully distributed in the sense that each model partition only shares information with its direct neighboring model partitions.

3.D. D-SABRE Stage-2: Merging DPME with DDA smoothing

Both the DPME and DDA operations introduced in the previous sections are iterative. These operations could be executed sequentially, starting with DPME to remove the unknown phase offsets, and ending with DDA to post-smooth the wavefront. This would, however, defeat the purpose of D-SABRE as a distributed WFR method because it would require all partitions to complete the DPME operation before starting the DDA operation. The reason for this is that if DDA is naively merged with DPME the DDA smoother will blend neighboring partitions regardless of convergence of the DPME operation. This results in significant undesirable and unphysical artifacts in the reconstructed wavefront.

The solution to this is a new concept indicated as virtual DPME (V-DPME). The idea is that even though DPME has not yet converged, each partition has at all times exact information on the phase offset between it and its direct neighbors. In V-DPME this information is used to virtually minimize all piston modes between a partition i and its G_i neighbors. After V-DPME, the DDA algorithm can safely smooth the locally piston mode equalized partitions without introducing artifacts into the reconstructed wavefront.

The information required for V-DPME is obtained during the DPME operation in the form of the set $\mathcal{K}_i^{(l)}(m)$ from (26) which contains all phase offsets between partition i and its neighbors. For V-DPME a new vector of virtual offsets is created from elements of $\mathcal{K}_i^{(l)}(m_j)$:

$$\begin{aligned} \mu_{i,M}(l) &= [0_{1 \times J_{\Omega_i} \hat{a}} \quad \mathcal{K}_i^{(l)}(m_1) \cdot \mathbf{1}_{1 \times J_{m_1} \hat{a}} \quad \mathcal{K}_i^{(l)}(m_2) \cdot \mathbf{1}_{1 \times J_{m_2} \hat{a}} \quad \cdots \quad]^\top, \\ &1 \leq i \leq G, \forall m_j \in \mathcal{M}_i. \end{aligned} \quad (32)$$

with J_{m_j} the total number of triangles in $\Omega_i \cap \Omega_{m_j}$.

Substitution of (32) in (31b), together with the DPME from (25b) results in the final formulation of Stage-2 of the D-SABRE :

$$\mathbf{c}_i(l+1) = \hat{\mathbf{c}}_i^{loc} + \mathbf{R}_{N_i} \mathbf{H}_i^\top \mathbf{y}_i(l) \quad (33a)$$

$$\mathbf{y}_i(l+1) = \mathbf{y}_i(l) + \alpha(l) \mathbf{H}_{i,M} (\mathbf{c}_{i,M}(l+1) + \mu_{i,M}(l)) \quad (33b)$$

which converges when $l = L$, after which time the actual PME value is added resulting in the final B-coefficient estimate:

$$\mathbf{c}_i = \mathbf{c}_i(L) + k_i(L) \quad (34)$$

3.E. Algorithm Convergence

The distributed piston mode equalization (DPME) and distributed dual ascent (DDA) steps of D-SABRE Stage-2 are both iterative in nature. In this section we provide an analysis of convergence of Stage-2 as a whole. It is important to note that all parallel operations in Stage-2 are allowed to run asynchronously, i.e. the parallel processors do not require a global clock-tick.

However, optimal convergence of the DPME operation does depend on synchronicity between neighboring partitions in the sense that failure of one parallel processor to calculate an initial local WF reconstruction in Stage-1 or an accurate PME update in Stage-2 at a specific global clock-tick will influence all partitions that through inter-partition continuity depend on these results. The DDA step on the other hand can be run completely asynchronously, because V-DPME negates any inter-partition offsets. For this analysis, we will assume that local wavefront reconstruction and DPME iterations run synchronously on the parallel processors.

DMPE convergence depends solely on the maximum distance, counted in number of partitions, between the master partition and any other partition. If this distance is R , then DPME converges in exactly R iterations because it takes at most R steps for information from the master partition to reach the partition furthest away from the master partition. Each step requires a single DPME iteration leading to exactly R iterations. Note that the 'max' operation in (25b) only influences the path the information follows, and not the total number of steps.

We base our analysis of convergence of the DDA step on the theory presented in [1]. First, let ϵ^* be the optimal value of the optimization problem (18):

$$\epsilon_i^* = \min \|\boldsymbol{\sigma}_i - \mathbf{D}_i \mathbf{c}_i\|_2^2 \quad \text{subject to } \mathbf{A}_i \mathbf{c}_i = 0, \quad 1 \leq i \leq G, \quad (35)$$

also, let

$$r_i(l+1) = \mathbf{A}_i \mathbf{c}_i(l+1) \quad (36)$$

be the primal residual at iteration $l + 1$.

Then [1] proves that

$$\epsilon_i^* - \epsilon_i(l + 1) \leq (y^*)^\top r(l + 1), \quad l \rightarrow \infty, \quad (37)$$

with $\epsilon_i(l + 1)$ the optimal value of (18) at iteration $l + 1$, and with y^* the optimal Lagrange dual. The right-hand side of (37) approaches zero as $l \rightarrow \infty$ since $r(l)$ approaches zero when the constraints are met. Therefore, we have $\lim_{l \rightarrow \infty} \epsilon_i(l + 1) = \epsilon_i^*$ implying convergence.

In practice we find that our DDA algorithm converges in at most a few tens of iterations, with adequate continuity achieved in $L = 10$ iterations, which agrees with the rules of thumb provided in [1]. In this case the norm of the residual of the smoothness constraints (see (7)) is $\|\mathbf{H}\mathbf{c}\|_2^2 < 1e - 3$ for $\text{SNR} \geq 0\text{dB}$ and reaches $\|\mathbf{H}\mathbf{c}\|_2^2 < 1e - 8$ for $\text{SNR} \geq 20\text{dB}$.

We also find that initializing the DDA iteration with a smooth initial estimate for $\hat{\mathbf{c}}_i^{\text{loc}}$ significantly improves convergence of the DDA step. Finally, the rate of convergence of the DDA stage depends strongly on the value of $\alpha(l)$ in (33). In our algorithm, a value of $\alpha(l) = 0.5$ leads to fast convergence.

3.F. Computational complexity analysis

Key to the utility of the D-SABRE method is its computational performance. In this section, the results from a theoretical analysis of computational complexity of the D-SABRE method are presented. The scope of this analysis is limited to theoretically required compute performance per parallel processor. Specific hardware dependent issues such as transport latency, cache size, and available instruction sets are not included in this analysis. In our analysis we focus on the real-time reconstruction operation, and not on algorithm initialization. In the analysis, we indicate R and L as the total number of iterations in respectively the DPME and DDA steps of Stage-2. We compare our results to that of the global SABRE method. The results are summarized in Table 1.

For the global SABRE we find for the complexity of the matrix-vector operation in (12):

$$U_{\text{SABRE}} = \mathcal{O}\left(2N J \hat{d}\right), \quad (38)$$

with N the total number of WFS, and with \hat{d} and J respectively the total number of B-coefficients and triangles in the global SABRE model.

For Stage-1 of the D-SABRE we find for the complexity of the matrix-vector operation in (21):

$$U_{\text{stage-1}} = \mathcal{O}\left(2N_i J_i \hat{d}\right), \quad (39)$$

with N_i the total number of WFS in partition i , with $J_i = J_{\Omega_i} + J_{\Xi_i}$ the total number of triangles (including the overlap triangles) in partition i .

Determining the complexity of Stage-2 is somewhat more involved because it is an iterative algorithm consisting of a number of operations. In the following, we shall assume that all necessary simplifications in terms of static matrix pre-calculation have been made. We then find for the complexity of the matrix-vector and vector-vector operations associated with DPME and DDA in (33):

$$U_{stage-2} = \mathcal{O} \left(R J_{\Omega_i} \hat{d} + L \left(J_{\Omega_i} \hat{d} (1 + r_H) + c_H (1 + r_H) \right) \right), \quad (40)$$

with R the total number of DPME iterations, J_{Ω_i} the total number of triangles in the core triangulation of partition i (see (14)) and with $r_H = \text{row}(\mathbf{H}_{i,M})$ and $c_H = \text{col}(\mathbf{H}_{i,M})$. It should be noted here that the complexity of calculating of the V-DPME and DPME offsets is negligible compared to the matrix-vector and vector-vector operations in (40).

In general (for $OL > 0$), we have $r_H \gg 1$, $c_H \gg 1$, $r_H < J_{\Omega_i} \hat{d} < c_H < J_i \hat{d} = \rho J_{\Omega_i} \hat{d}$ with ρ from (15). Additionally, we can assume that $2N_i = J_i = \rho J_{\Omega_i}$ for a Simplex Type-I WFS geometry [5]. With these (conservative) assumptions and under the introduction of ρ , we can simplify and combine (39) and (40):

$$\begin{aligned} U_{DSABRE} &= \mathcal{O} \left(\rho^2 J_{\Omega_i}^2 \hat{d} + R J_{\Omega_i} \hat{d} + L \left((J_{\Omega_i} \hat{d})^2 + \rho (J_{\Omega_i} \hat{d})^2 \right) \right), \\ &= \mathcal{O} \left(\rho^2 J_{\Omega_i}^2 \hat{d} + R J_{\Omega_i} \hat{d} + L (1 + \rho) (J_{\Omega_i} \hat{d})^2 \right), \\ &= \mathcal{O} \left(J_{\Omega_i} \hat{d} (\rho^2 J_{\Omega_i} + R + L \hat{d} (1 + \rho) J_{\Omega_i}) \right), \end{aligned} \quad (41)$$

from which it can be concluded that the contribution of DPME to the total computational complexity can be neglected for as long as R is small compared to the other terms in brackets, which holds for most realistic partitionings.

If all parallel processors operate synchronously, and no transport latency is present, the speedup factor over the global SABRE is obtained by combining (38) with (41):

$$S_{DSABRE} = \mathcal{O} \left(\frac{2N \cdot J \hat{d}}{J_{\Omega_i} \hat{d} (\rho^2 J_{\Omega_i} + R + L \hat{d} (1 + \rho) J_{\Omega_i})} \right). \quad (42)$$

If all partitions have the same number of triangles in the core of the partition (which is desirable), we have $G = J/J_{\Omega_i}$ with G the total number of partitions.

In addition, if we assume that use is made of a Simplex Type-I geometry (in which case $N = J/2$) then using the result $J = G J_{\Omega_i}$ reduces (42) to:

$$\begin{aligned} \tilde{S}_{DSABRE} &= \mathcal{O} \left(\frac{2J/2G J_{\Omega_i} \hat{d}}{J_{\Omega_i} \hat{d} (\rho^2 J_{\Omega_i} + R + L \hat{d} (1 + \rho) J_{\Omega_i})} \right), \\ &= \mathcal{O} \left(\frac{JG}{\rho^2 J_{\Omega_i} + R + L \hat{d} (1 + \rho) J_{\Omega_i}} \right), \\ &= \mathcal{O} \left(\frac{G^2}{\rho^2 + R/J_{\Omega_i} + L \hat{d} (1 + \rho)} \right), \end{aligned} \quad (43)$$

From (43) it can be concluded 1) the influence of the DPME step on the computational complexity is negligible as long as $R/J_i \ll \rho^2$; 2) the overlap overhead reduces the speedup by a factor ρ^2 ; and 3) that the DDA step is the dominant term in the computational efficiency if $L\hat{d} > \rho$.

As a rule of thumb, a partitioning can (and should) be designed such that the overlap overhead $1 < \rho < 2$. For that overlap level, $L = 10$ produces adequate results in most cases, showing that the DDA step in the form of the variable L indeed has a dominant influence on the speedup factor. In fact, when the DDA step is not used, and if we (reasonably) assume $R/J_i \ll \rho^2$, we obtain a speedup factor of $\frac{1}{\rho^2}G^2$ for the linear D-SABRE . In Table 1 the results of the complexity analysis are summarized.

Finally, it must be noted here that the efficiency of DDA step can be improved significantly when only inter-partition smoothness is enforced in (33) instead of full-partition smoothness. In that case r_H and c_H will be significantly smaller to the point that the overlap overhead will have a dominant contribution to the complexity.

	Global SABRE	D-SABRE (per core)	Speedup (per core)
WFR	$\mathcal{O}(\hat{d}J^2)$	$\mathcal{O}(\rho^2 J_{\Omega_i}^2 \hat{d})$	$\mathcal{O}(G^2/\rho^2)$
DPME	0	$\mathcal{O}(RJ_{\Omega_i} \hat{d})$	$\mathcal{O}(1/(RJ_{\Omega_i} \hat{d}))$
DDA	0	$\mathcal{O}(L(1+\rho)(J_{\Omega_i} \hat{d})^2)$	$\mathcal{O}(1/(L(1+\rho)(J_{\Omega_i} \hat{d})^2))$
Total	$\mathcal{O}(J^2 \hat{d})$	$\mathcal{O}(J_{\Omega_i} \hat{d}(\rho^2 J_{\Omega_i} + R + L\hat{d}(1+\rho)J_{\Omega_i}))$	$\mathcal{O}\left(\frac{G^2}{\rho^2 + R/J_{\Omega_i} + L\hat{d}(1+\rho)}\right)$

J: number of triangles in global triangulation; *J_i* number of triangles per partition
G: total number of partitions; *Ď*: B-coefficients per triangle (*Ď* = 3 for linear D-SABRE)
R: DPME iterations ($R \ll \hat{d}J_i$); *L*: DDA iterations, $\rho = J_i/J_{\Omega_i}$

Table 1: Comparison of theoretical computational complexity of global SABRE and D-SABRE , where a Simplex Type-1 WFS geometry (i.e. $N = J/2$) is assumed. All partitions run on a separate (perfect) parallel processor.

As an example of Table 1, consider first a 64×64 WFS array such that $N = 4096$, and assume that a linear spline model (i.e. $\hat{d} = 3$) is used on a Type-1 sensor geometry. In this case, the global SABRE would require $\mathcal{O}(J^2 \hat{d}) = \mathcal{O}(2 \cdot 10^8)$ flops for a single reconstruction. For the D-SABRE with 64 partitions $OL = 2$, $R = 4$ and no DDA ($L = 0$), we have $J_{\Omega_i} = 128$ and subsequently $1.53 \leq \rho \leq 2.19$ with (16). In this case, the contribution of DPME can be neglected ($R/J_{\Omega_i} \ll 1$), and we have as a worst case $\mathcal{O}(\rho^2 J_{\Omega_i}^2 \hat{d}) = \mathcal{O}(2 \cdot 10^5) = \mathcal{O}(57N)$ flops per D-SABRE core per reconstruction. For a 200×200 array and a 20×20 partitioning ($N = 40000$), $OL = 2$, $R = 10$ and no DDA ($L = 0$) we have $J_{\Omega_i} = 200$ and subsequently $1.42 \leq \rho \leq 1.92$. Again, the contribution of DPME can be neglected, and we find for the worst case $\mathcal{O}(4 \cdot 10^5) = \mathcal{O}(11N)$ flops per D-SABRE core per reconstruction.

Even though hardware dependent issues are not scope of this paper, we want to mention that communication latency between the partitions for the iterative, distributed Stage-2 of the D-SABRE method is not considered to be an issue in a planned parallel GPU implementation. For the reference case of a 200×200 SH array and a 20×20 partitioning and a standard of the shelf GPU, all necessary data, e.g. reconstruction matrices, of the D-SABRE method can be stored directly on the GPU memory and no communication over the slow GPU to CPU connection will be necessary.

4. Simulations with the D-SABRE

In this section, the results from numerical experiments with the D-SABRE are presented. These experiments are aimed at validating the D-SABRE method by a simulation study.

First, the reconstruction accuracy and noise resilience of the D-SABRE are compared to that of the centralised SABRE method in Sec. 4.A and Sec. 4.B. As in [5], a von Karman turbulence model was used to simulate phase screens which are then measured by a Fourier optics based Shack-Hartmann (SH) WFS simulator which generates wavefront phase slopes for a square grid of varying numbers of subapertures. Slope measurements of different signal to noise ratio (SNR) levels are created by adding gaussian distributed white noise.

The second part of this chapter contains a comparison of D-SABRE to the distributed Cumulative Reconstructor (CuRe-D) method which provides like D-SABRE parallelizable wavefront reconstruction of linear complexity and was published in [18]. Both methods were implemented for YAO [15], a Monte-Carlo simulation tool for astronomical AO systems, to test their reconstruction performance on a reference simulator. The main goal of this study was to obtain a comparison of the two methods under equal conditions to understand if the D-SABRE in its most basic form reaches the same level of performance as the CuRe-D method for different decomposition levels. This baseline version of the D-SABRE applies a spline model of polynomial degree $d = 1$ and continuity order $r = 0$; it further performs only the D-PME step in Stage-2. After an open-loop analysis of reconstruction accuracy and noise resilience under the influence of additive white noise, we tested both methods in a closed-loop environment for a SCAO setting considering different levels of photon-shot noise. Herefore, an additional projection step of the reconstructed phase onto the deformable mirror was implemented to allow the correction in combination with a simple integrator control provided by YAO.

4.A. Validation of D-SABRE reconstruction accuracy

The first experiment is aimed at comparing the D-SABRE with the global SABRE method from [5]. For this a Fourier optics based SH lenslet array is used to obtain wavefront slopes from a set of 100 simulated wavefronts observing at a wavelength of 699 nm . For the von




Fig. 6: Comparison of average Strehl ratio of the global SABRE method with that of various D-SABRE variants as function of the signal to noise ratio where 0dB corresponds with equal magnitudes of noise and signal. All D-SABRE variants use a decomposition of the global 50×50 WFS grid (2500 WFS) into 25 partitions.




Fig. 7: Left: Average Strehl ratio as a function of the total number of WFS and a fixed number of 100 partitions. Right: Average Strehl ratio as a function of signal to noise ratio for a varying partition count on a fixed 50×50 WFS grid where 0dB corresponds with equal magnitudes of noise and signal. In both cases an OL-1 overlap with DDA smoothing is used.

Karman turbulence model a Fried coherence length of $0.2 m$ was used with a turbulence outer scale of $50 m$ and the telescope diameter is assumed to be $2 m$. The SH lenslet array consists of 2500 lenslets, laid out in a 50×50 grid with full illumination of all subapertures. Noise affecting the SH wavefront sensor measurements is simulated through white noise added to the slope measurements, with the signal-to-noise ratio (SNR) provided by the ratio of the slope and noise variance in the logarithmic decibel scale, where a SNR of 0dB implies that the magnitude of the noise is equal to that of the signal.

In Fig. 6, the average Strehl ratio is plotted as a function of SNR for different variants of the D-SABRE and compared to that of the global SABRE method. For the D-SABRE the global WFS domain is decomposed into 25 partitions. For each SNR setting a total of 100 reconstructions are performed. The results show that the D-SABRE with OL-4 and DDA approximates the global SABRE reconstruction within 1% in terms of the Strehl ratio for signal to noise ratio's (SNR) $\geq 20dB$, within 5% for SNR's $\geq 10dB$, and within 10% for SNR $\geq 5dB$. In addition Fig. 6 shows that including the DDA step significantly improves D-SABRE reconstruction accuracy for higher noise cases.

4.B. Investigation of D-SABRE PME error propagation

An important issue with current state-of-the-art distributed WFR methods is the increase of the reconstruction error as the total size of the WFS array increases [18]. In contrast to these methods, D-SABRE is not subject to noise propagation within partitions because an optimal least-squares estimator is used. However, D-SABRE is subject to inaccuracies in the least squares estimates of the piston mode offsets that are propagated between partitions. In principle, these inaccuracies will only occur if the model residue is non-white or if the sample of the piston offset on shared vertices is too small. In fact, as the absolute size of the partitions increases (i.e. more triangles per partition), the piston mode offset estimation will become more accurate as the total number of shared vertices used to calculate the offset will increase.

In this paragraph, we present the results of an investigation into PME error propagation issues with the D-SABRE using numerical simulations.

In the left plot of Fig. 7, the results from the simulations of WFS array sizing are shown. In this case, the WFS array size is increased from 10^2 to 10^4 sensors, while the total number of partitions is fixed at 100. For each WFS array a set of 100 reconstructions are conducted at a low SNR of 6.67 dB and a high SNR of 30 dB. From Fig. 7 it can be observed that the Strehl ratio actually increases with increasing WFS array size for both the SABRE and D-SABRE method for all SNR values. Also, it demonstrates that when the absolute size of partitions increases, the propagation of PME errors *decreases*.

In the right plot of Fig. 7, the results from simulations with different partitionings of a 50×50 grid are shown. In this case, the global WFS array is decomposed into 25, 100, and 625 local partitions. At each SNR level 100 reconstructions are performed after which the result is averaged. This time, there is a clear positive influence on the Strehl ratio of *reducing* the total number of partitions on the fixed 50×50 WFS array.

From Fig. 7 it can be concluded that the D-SABRE is subject to propagation of PME errors between partitions, but is not subject to noise propagation. Also, the PME error actually reduces in size as the total number of shared vertices is increased.

4.C. Comparison with CuRe-D in YAO open-loop simulations

The linear D-SABRE method is compared with the distributed Cumulative Reconstructor (CuRe-D) in terms of reconstruction accuracy and noise propagation for an open-loop configuration.

4.C.1. Open-loop configurations

The AO simulation tool YAO was used to create 100 wavefront realizations with a Fried parameter $r_0 = 0.18$ cm and simulate the corresponding diffraction based slope measurements




Fig. 8: Triangulation and partitioning example for D-SABRE. Left: Global Type II triangulation for a 32×32 SH array constructed with vertices located on SH center locations (cyan) and added vertices (blue) to form the regular Type II triangulation and to cover the entire pupil domain. Center: Partitioning of the global triangulation in 16 partitions on a 4×4 partition array. Right: Partitioning in 16 partitions with applied overlap level of $OL=2$.

of a Shack Hartmann wavefront sensor with a 64×64 subaperture array on a circular telescope pupil of 30 m diameter. In YAO, the slope measurements are computed, after full propagation and subaperture image formation, on a 10×10 pixel array using the center of mass algorithm. Subapertures of illumination higher or equal than 50% were taken into account. In this section, zero noise is assumed in the simulation of the slope measurements and an observation wavelength of 650 nm is considered.

Due to its hierarchical structure, the CuRe-D method is limited to partition arrays of size $N_p \times N_p$ with N_p as powers of two, and it provides wavefront estimates at the corners of the SH subapertures. Though such partitioning is not necessary for D-SABRE, in order to compare the two methods, the D-SABRE was configured accordingly and the obtained spline estimate of the wavefront is evaluated at the subaperture corner positions located within the pupil domain. Note that the D-SABRE method computes an analytical solution of the wavefront reconstruction problem providing phase estimates at any point in the triangulation which covers the telescope pupil. In Fig. 8, we show the example of a Type II triangulation for a 32×32 SH subaperture array. The two graphics on the right depict a decomposition of the global partition on the left into a 4×4 partitioning with an overlap level of $OL=2$.

Analogous to Sec. 4.A and Sec. 4.C, noise effects in the sensing process are simulated with additive white noise on the slope measurements provided by YAO. Low signal to noise ratios of $SNR < 12\text{ dB}$ are considered in this section to investigate noise resilience as well as reconstruction accuracy of the two distributed methods. To do so, the relative RMS values of




Fig. 9: Reconstruction accuracy and noise resilience for different levels of decomposition on a 64×64 SH array of D-SABRE and CuRe-D . In the D-SABRE case different levels of partition overlap are considered. The average relative residual RMS is plotted as a function of signal to noise ratio.

the residual wavefronts are compared for varying signal to noise ratios. The data provided is the mean and its standard deviation computed from RMS errors of wavefront reconstructions for 100 phase realizations.

4.C.2. Open-loop results

In Fig. 9, the results from simulations with different partitionings of the 64×64 SH array are shown. For both methods, we consider $N_p \times N_p$ partition arrays for $N_p = 2, 4, 8$ resulting in square partitions each containing 32×32 , 16×16 and 8×8 subapertures respectively. In case of D-SABRE, a Type-II triangulation was used and overlap levels of 1, 2, $64/2N_p$ and $64/N_p$ applied.

In comparison with CuRe-D, the D-SABRE shows an advantage in reconstruction accuracy for lower levels of decomposition. For the example of a 2×2 partitioning of the domain, D-SABRE clearly outperforms CuRe-D even for minimal amount of partition overlap. For higher level of decomposition, equal or better accuracy can only be achieved by increasing the overlap level. For the 4×4 partitioning, $OL = 2$ leads to performance comparable to that of CuRe-D, whereas for 8×8 partitioning an overlap of $OL=64/N_p$ is necessary to achieve better results than CuRe-D.

In terms of noise resilience, a similar trend can be observed. For lower levels of decomposition, D-SABRE proves more resilient to an increase in the amount of additive white noise than CuRe-D. We considered the intervals between SNR values at which relative RMS errors are provided in Fig. 9 and computed the average rate of change of the relative RMS error over these SNR intervals for both methods through finite differences. Averaged over the different SNR intervals, the rate with which the accuracy of D-SABRE for the 2×2 partitioning decreases for increasing noise is only 56% of the rate observed for CuRe-D. For the 4×4 partitioning, the rate of accuracy decay for D-SABRE increases to 70%, for the 8×8 partitioning to 98% of the rate observed for CuRe-D. These values were obtained for an overlap level of $OL=64/2N_p$ in case of the D-SABRE method.

The results presented in this section can be explained with the very different noise propagation behaviors of the two methods. It has been shown in [17] and [18], that the cumulative approach of the CuRe-D algorithm leads to accumulation of noise for larger numbers of subapertures per partition. Hence, increasing the level of decomposition for the distributed CuRe-D improves the noise propagation properties of the method. It can be seen in Fig. 9 that this reduction of noise propagation within partitions outweighs the negative effect of piston mode estimation errors introduced through the decomposition. Contrary, as discussed in Sec. 4.C, the D-SABRE does not suffer from noise propagation within partitions which contain a large number of subapertures. However, D-SABRE is subject to PME error propagation between partitions, which results in loss of performance for an increased level of

decomposition given a fixed size WFS array. This effect can be counteracted by increasing the overlap level OL, or increasing the size of the WFS array. For strong decomposition, the D-SABRE in the tested baseline version meets the performance of CuRe-D only for sufficient amount of overlap OL, which decreases the speed up as shown in (43).

As discussed in Sec. 2.B, more advanced sensor models can further improve the quality of the local D-SABRE reconstructions and reduce the inaccuracies in the piston mode estimate which are currently propagated between partitions. This would allow higher levels of decomposition and less partition overlap. Current work on a C-implementation of the D-SABRE method will make it feasible to test the method for larger WFS arrays of $\approx 10^4 - 10^5$ subapertures as will be found in XAO system.

4.D. Comparison with CuRe-D in YAO closed-loop simulations

To conclude this chapter, D-SABRE and CuRe-D are compared in a closed-loop environment regarding their response to photon shot noise. Since D-SABRE and CuRe-D provide estimates of the residual phase and not directly actuator commands for its correction, an additional step mapping the phase onto the deformable mirror had to be implemented which was then combined with the simple integrator control law provided by YAO.

4.D.1. Closed-loop configurations

As in Sec. 4.C, we evaluate the D-SABRE phase estimates at the subaperture corners, i.e. the locations where phase values are provided by CuRe-D. We want to stress at this point that D-SABRE provides an analytical solution to the wavefront reconstruction problem and therefore phase estimates over the whole pupil plane. Missalignments between subaperture and actuator grid can be taken into account without additional approximation error e.g. introduced through further interpolation. Further, the B-spline basis function matrix introduced in (4), which evaluates the D-SABRE phase estimates, can be precomputed. This allows to apply the deformable mirror fit as a direct mapping between the B-coefficients and the actuator commands.

We consider a SCAO system with a stackarray deformable mirror defined through the shape of the actuator influence functions which are provided by YAO. Based on the influence function values at the subaperture corners, the actuator commands were computed in order to fit the deformable mirror optimally in a least-squares sense to the phase values. This fitting step is performed in a centralized manner and was not optimized for speed. To allow the computation of meaningful results which require a certain number of closed-loop iteration for several turbulence realizations, the density of the actuators is set lower than the density of the 64×64 subaperture array. We opted for a 32×32 actuator grid where the actuator positions are conjugated to every second subaperture corner in the SH array. Again the




Fig. 10: Long exposure Strehl ratio comparing closed-loop performance and noise resilience of D-SABRE and CuRe-D for different of levels photon shot noise in a system with a 64×64 SH array and a 32×32 actuator array. Left: Turbulence with Fried parameter $r_0 = 18\text{ cm}$; right: $r_0 = 25\text{ cm}$. A 4×4 partitioning is constructed for both distributed methods. Two loop gains are considered for each method respectively optimized for low (solid) and high (dashed) noise levels. In the D-SABRE case $OL = 64 / (2N_p)$ is used. The Strehl ratio is averaged for 10 turbulence realizations and plotted as a function of the guide star magnitude.

focus was to compare the behavior of the methods for changes in parameters, like photon-shot noise and loop gain, under identical circumstances, rather than aiming for a simulation of a XAO dimensioned system. The coupling between the actuators is set to 20% and the normalized response threshold (in WFS signal) below which an actuator will not be kept as valid to 30%.

To avoid initial transient, the long exposure Strehl ratio is accumulated starting with the 50th iteration and the loop is evaluated for 1 *s* with a sampling time of 2 *ms*. The results are given for a wavelength of 650 *nm* and 10 turbulence realizations. The loop gain of the integrator control was tuned for both methods to perform optimally averaged for the same set of turbulence realizations used in the following experiment. Pairs of optimal loop gains for D-SABRE and CuRe-D were computed for two different noise levels: for a low noise scenario simulated with a high flux guide star of magnitude 5 and for a high noise scenario simulated with a low flux guide star of magnitude 11. The distributed methods D-SABRE and CuRe-D were applied for a 4×4 partitioning of the 64×64 SH array; in the case of D-SABRE, an overlap level OL=8 was chosen. As in the previous section, the D-SABRE was performed in its baseline version.

4.D.2. Closed-loop results

In Fig. 10, the mean of the closed-loop Strehl ratios is provided with the standard deviation of the mean for different levels of photon shot noise indicated by the magnitude of the guide star. The solid lines show the Strehl ratios obtained by the respective method if applying loop gains optimized for guide star magnitude 5 (i.e. high flux), the dashed lines show the results obtained for loop gains optimized for guide star magnitude 11 (i.e. low flux). The closed-loop experiment was performed for the set of turbulence realizations with Fried parameter $r_0 = 18$ *cm* used for the open-loop tests in Sec. 4.C. For the loop gains 0.7 and 0.3 obtained by tuning D-SABRE and CuRe-D respectively for high flux, the D-SABRE outperforms CuRe-D for all considered guide star magnitudes. The advantage which the applied configuration of D-SABRE has shown for the open-loop environment with additive white noise (see Fig. 9) is preserved in the closed loop experiment under exposure of photon shot noise, despite not with the same margin. Decreasing the loop gain for CuRe-D to 0.2, the optimal loop gain for low flux, boosts its performance for guide star magnitudes > 9 . The analog adjustment of the D-SABRE loop gain to 0.6, results in a less significant improvement for the low flux range and D-SABRE does not reach the performance of CuRe-D for the resulting levels of photon shot. However in the high flux range, D-SABRE proves in this experiment greater robustness to the variation of the loop gain. D-SABRE preserves its performance for guide star magnitudes < 7 , whereas CuRe-D is more sensitive to changes of the loop gain and suffers a significant drop in Strehl ratio. The experiment was repeated for a set of 10 turbulence realizations with

Fried parameter $r_0 = 25 \text{ cm}$. The results are depicted in the right plot of Fig. 10 and confirm the findings made in this section. For sufficient amount of overlap and a moderate level of decomposition, the baseline D-SABRE meets the performance of the CuRe-D method also in a closed-loop setting and under the impact of photon shot noise.

5. Conclusions

In this paper a new distributed multivariate spline based wavefront reconstruction method is introduced. This new method, the D-SABRE (Distributed Spline based ABeration REconstruction), is an extension of the SABRE method for wavefront reconstruction. The D-SABRE is aimed at large scale wavefront reconstruction problems (e.g. 200x200 WFS grids) where global reconstructions are not realistic.

The D-SABRE method partitions the global WFS domain into any number of partitions, and then solves the local reconstruction problem in 2 distributed stages. In the first stage the the local wavefront reconstruction problem is solved for every partition; this stage is completely distributed in the sense that no communication between partitions is necessary. In the second stage, the local partitions are smoothed into a single continuous global wavefront. This stage requires only communication between direct partition neighbors. The blending consists of 2 operations; the distributed piston mode equalization (D-PME) and the distributed dual ascent (DDA) operation. D-PME is aimed at equalizing the unknown piston modes between partitions, DDA is aimed at post-smoothing the resulting partitions. Both operations are iterative, and convergence of the operations is proved. Without D-PME no accurate global wavefront can be obtained. The DDA operation on the other hand can be considered optional, and should only be applied when a continuous global wavefront reconstruction is desired because of its significant negative influence on computational complexity.

An analysis of computational complexity showed that a speedup factor over the SABRE can be obtained that is quadratically proportional to the total number of partitions. The most dominant factor for the complexity is the local reconstruction stage (State-1) and the optional distributed post-smoothing (DDA) step in Stage-2. The D-PME step in Stage-2 is cheap in terms of computational complexity.

Numerical experiments with the D-SABRE shows that reconstruction accuracy approximates the global wavefront reconstruction within 1% in terms of the Strehl ratio for signal to noise ratio's (SNR) $\geq 20dB$, within 5% for SNR's $\geq 10dB$, and within 10% for very low SNR $\geq 5dB$. The most important factors determining reconstruction accuracy are the overlap level between partitions, and the absolute partition size, where a larger partition leads to more accurate results.

The D-SABRE is not subject to noise propagation in the same sense as the CuRe and CuRe-D methods. Instead it is subject to propagation of piston mode offset estimations

between partitions. The magnitude of this PME error propagation decreases as the absolute size of the partitions increases. In that sense, the D-SABRE favors very large WFS arrays, that are partitioned into partitions that are as large as possible given a particular hardware setup.

Finally, a comparison of the D-SABRE with the CuRe-D method is given through open- and closed-loop simulations which are performed with the AO simulation tool YAO. The open-loop analysis of the performance under the influence of additive white noise shows a clear advantage of D-SABRE for low levels of domain partitioning. If the number of partitions is increased and the size of the partitions decreases, overlap between the partitions has to be applied in order to outperform CuRe-D. In the tested closed-loop SCAO configurations, D-SABRE and CuRe-D showed comparable behavior for increasing levels of photon shot noise. D-SABRE has proven to be more robust to variations in the gain of the control loop.

We conclude the D-SABRE will be most useful on very large scale ($> 200 \times 200$) WFS arrays, where the computational hardware favors low levels of domain partitioning resulting in partitions that are as large as possible. This is in contrast with CuRe-D, which is useful on large scale WFS arrays but which requires an as high as possible domain partitioning to prevent noise propagation.

Future work on the D-SABRE method will be focused on extending the approach 1) to higher degree polynomials (i.e. $d \geq 3$), which requires imposing explicit boundary conditions on the local models; 2) by implementing more advanced sensor models such as that introduced in [8] that exploit the integrative nature of the SH sensors; 3) by implementing more advanced estimation schemes that exploit knowledge of turbulence statistics; and finally 4) by implementing D-SABRE on a real-world parallel hardware (GPU) setup.

6. Appendix

By exploiting the sparseness of \mathbf{A} in (30) we can derive two new submatrices \mathbf{H}_i and $\mathbf{H}_{i,m}$ for each partition i . For this we first introduce $\mathcal{E} = 1, 2, \dots, E$ as the global index set of *all* triangle edges in the global triangulation \mathcal{T} , and $\mathcal{C} = 1, 2, \dots, J \cdot \hat{c}$ as the global index set of *all* B-coefficients in the global spline model s_r^d .

The first submatrix of \mathbf{A} , indicated as the *inner constraint matrix* \mathbf{H}_i , contains blocks that only influence B-coefficients inside the partition i :

$$\mathbf{H}_i := \mathbf{A}(\mathcal{E}_{\Omega_i}, \mathcal{C}_{\Omega_i}). \quad (44)$$

with $\mathcal{E}_{\Omega_i} \subset \mathcal{E}$ the set of indices of all triangle edges inside the core part Ω_i of the sub-triangulation \mathcal{T}_i from see (13), and with $\mathcal{C}_{\Omega_i} \subset \mathcal{C}$ the set of all B-coefficient indices in Ω_i .

The second submatrix $\mathbf{H}_{i,m}$ contains all blocks in \mathbf{A} that influence B-coefficients inside the partition i *as well as* B-coefficients in neighboring partitions m through the action of

the continuity conditions:

$$\mathbf{H}_{i,m} := \mathbf{A}(\mathcal{E}_{i,m}, \mathcal{C}_{i,m}). \quad (45)$$

with $\mathcal{E}_{i,m} = \mathcal{E}_{\Omega_i} \cup \mathcal{E}_{i \rightarrow m}$ the indices of all edges in \mathcal{T}_i that influence $s_{r_i}^d$. Finally, we define $\mathcal{C}_{i,m} \subset \mathcal{C}$ as the indices of all B-coefficients in $s_{r_i}^d$ as well as all B-coefficients in $s_{r_m}^d$ that are subject to continuity conditions on $\mathcal{T}_i \cap \mathcal{T}_m$.

We define $\mathbf{y}_i = \mathbf{y}(\mathcal{E}_{\Omega_i})$ as the local dual vector, with $y_{i,m} = \mathbf{y}(\mathcal{E}_{i,m})$ the inter-partition dual vector, with $\mathbf{c}_i = \mathbf{c}(\mathcal{C}_{\Omega_i})$ the local B-coefficient vector, and with $\mathbf{c}_{i,m} = \mathbf{c}(\mathcal{C}_{i,m})$ all B-coefficients in $s_{r_i}^d$ as well as all B-coefficients in $s_{r_m}^d$ that are subject to continuity conditions on $\Omega_i \cap \Omega_m$.

References

1. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3:1–122, 2010.
2. C. C. de Visser. *Global Nonlinear Model Identification with Multivariate Splines*. PhD thesis, Delft University of Technology, 2011.
3. C. C. de Visser, Q. P. Chu, and J. A. Mulder. A new approach to linear regression with multivariate splines. *Automatica*, 45(12):2903–2909, 2009.
4. C. C. de Visser, Q. P. Chu, and J. A. Mulder. Differential constraints for bounded recursive identification with multivariate splines. *Automatica*, 47:2059–2066, 2011.
5. C.C. de Visser and M. Verhaegen. Wavefront reconstruction in adaptive optics systems using nonlinear multivariate splines. *Journal of the Optical Society of America*, 30:82–95, 2013.
6. D. L. Fried. Least-square fitting a wave-front distortion estimate to an array of phase-difference measurements. *Journal of the Optical Society of America*, 67:370–375, 1977.
7. L. Gilles, C. R. Vogel, and B. L. Ellerbroek. Multigrid preconditioned conjugate-gradient method for large-scale wave-front reconstruction. *Journal of the Optical Society of America*, 19:1817–1822, 2002.
8. S. Guo, R. Zhan, J. Li, J. Zou, R. Xu, and C. Liu. Sequential deconvolution from wave-front sensing using bivariate simplex splines. *Optics Communications*, 342(0):73–78, 2015.
9. P. J. Hampton, P. Agathoklis, and C. Bradley. A new wave-front reconstruction method for adaptive optics systems using wavelets. *IEEE Journal of Selected Topics in Signal Processing*, 2:781–792, 2008.
10. J. Herrmann. Least-squares wave front errors of minimum norm. *Journal of the Optical Society of America*, 70:28–35, 1980.

11. M. J. Lai and L. L. Schumaker. *Spline Functions on Triangulations*. Cambridge University Press, 2007.
12. M. J. Lai and L. L. Schumaker. A domain decomposition method for computing bivariate spline fits of scattered data. *SIAM Journal on Numerical Analysis*, 47:911–928, 2009.
13. L. A. Poyneer, D. T. Gavel, and J. M. Brase. Fast wave-front reconstruction in large adaptive optics systems with use of the Fourier transform. *Journal of the Optical Society of America*, 19:2100–2111, 2002.
14. R. Rannacher, S. Turek, and J. G. Heywood. Artificial boundaries and flux and pressure conditions for the incompressible navier-stokes equations. *International Journal for Numerical Methods in Fluids*, 22:325–352, 1996.
15. F. Rigaut and M. van Dam. Simulating astronomical adaptive optics systems using yao. In *Third AO4ELT Conference - Adaptive Optics for Extremely Large Telescopes*, 2013.
16. F. Roddier. Curvature sensing and compensation: a new concept in adaptive optics. *Applied Optics*, 27:1223–1225, 1986.
17. M. Rosensteiner. Cumulative reconstructor: fast wavefront reconstruction algorithm for extremely large telescopes. *Journal of the Optical Society of America*, 28:2132–2138, 2011.
18. M. Rosensteiner. Wavefront reconstruction for extremely large telescopes via cure with domain decomposition. *Journal of the Optical Society of America*, 29:2328–2336, 2012.
19. W. H. Southwell. Wave-front estimation from wave-front slope measurements. *Journal of the Optical Society of America*, 79:998–1006, 1980.
20. E. Thiébaud and M. Tallon. Fast minimum variance wavefront reconstruction for extremely large telescopes. *Journal of the Optical Society of America*, 27:1046–1059, 2010.