Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope

Hatem Ltaief*, Ali Charara*, Damien Gratadour[†], Nicolas Doucet*[†], Bilel Hadri[‡], Eric Gendron[†], Saber Feki[‡] and David Keyes*

*Extreme Computing Research Center

[‡]Supercomputing Lab

King Abdullah University of Science and Technology, Thuwal, Jeddah 23955

Email: Hatem.Ltaief, Ali.Charara, Bilel.Hadri, Saber.Feki, David.Keyes@kaust.edu.sa

[†]Paris Observatory - LESIA, Paris, France

Email: Damien.Gratadour, Nicolas.Doucet, Eric.Gendron@obspm.fr

Abstract—The European Extremely Large Telescope (E-ELT) is one of today’s most challenging projects in ground based astronomy. Addressing one of the key science cases for the E-ELT, the study of the early Universe, requires the implementation of multi-object adaptive optics (MOAO), a dedicated concept relying on turbulence tomography. We use a novel pseudo-analytical approach to simulate the performance of tomographic reconstruction of the atmospheric turbulence in a MOAO system on real datasets. We simulate simultaneously 4K galaxies in a common field of view on massively parallel supercomputers during a single night of observations. We are able to generate a first-ever high resolution galaxy map at almost a real-time throughput. This simulation scale opens new research horizons in numerical methods for experimental astronomy, some core components of the pipeline standing as pathfinders toward actual operations and future astronomic discoveries on the E-ELT.

I. INTRODUCTION

The European Extremely Large Telescope (E-ELT) is perhaps the most challenging project in ground-based astronomy due its unprecedented diameter size (40 meters) and the complexity of the planned instrument suite. In particular, MOSAIC [1] is one of the most critical instruments due to its intrinsically multiplexed design. Coupled to a novel adaptive optics concept, called multi-object AO (MOAO), the MOSAIC instrument embarks several deformable mirrors used to compensate for the atmospheric turbulence in real-time in order to remove any distortions that may occur on the resulting images for various positions in a large science field.

We describe a simulation of MOSAIC on real turbulence profiles and assess its throughput in terms of image quality in the Field of View (FoV). These image quality estimates are also called Point Spread Functions (PSFs). The crux of the paper is a pseudo-analytic formulation of the error propagation in an AO system, which helps to expose parallelism in the simulation process by providing an approach based on well-established linear algebra software libraries. The simulation

corresponds to a pipeline of successive covariance matrix generations along with compute-intensive operations.

This pipeline can be represented by two nested loops. The outer one is the long exposure and generates a tomographic reconstructor for each position in the science FoV. This tomographic reconstructor is then fed to an inner loop which is in charge of generating the actual short exposure PSFs, with varying turbulence conditions. Averaging these short exposure PSFs over several iterations of the inner and outer loops, emulates the variation of turbulence conditions and lead to an estimate of the system performance along the course of a night of observations.

The actual workload depends on many parameters from the telescope (e.g., the number of wavefront sensors), from the system environment (e.g., atmospheric turbulence), and from the number of targets of interest. While this simulation is critical by itself at this stage of the E-ELT design to define the proper dimensioning for the instrument, it can also be used toward the generation of a high resolution performance map, which may be a key reference for astronomers to predict the scientific return to be expected at the focus of the instrument on a number of science fields. To our knowledge, this is the first-ever generation of such high resolution performance map for such a complex instrument.

The implementation of the simulation pipeline relies on ScaLAPACK [2], the state-of-the-art high performance dense linear algebra software library. ScaLAPACK uses a two-dimensional block cyclic data distribution to scatter the various data structures over the memory of the computational nodes and relies on efficient BLAS implementations to extract performance from the underlying hardware resources. The resulting high-throughput pipeline is able to cope with 4K galaxies in a common field of view on two massively parallel supercomputers and generate a corresponding high resolution performance map for a night of observations in almost real-time, i.e., we generate the PSFs at the same speed at which

we observe the sky.

The remainder of the paper is organized as follows. Section II describes the multi-object adaptive optics (MOAO) technique. Section III recalls the state-of-the-art, details related work, and states the problem. Section IV highlights the MOAO implementation on distributed-memory systems. Section V outlines the hardware specifications of the two leading-edge supercomputers on which the experiments have run. The science performance results and analysis are presented in Section VI using real observation data. HPC performance results and analysis are demonstrated in Section VII. Section VIII presents the impact for the overall MOAO technique moving forward and we conclude in Section IX.

II. THE MULTI-OBJECT ADAPTIVE OPTICS TECHNIQUE

Probing the light from the first, most distant galaxies is key to understanding the epoch of reionisation, the first period in the history of the Universe accessible to our telescopes, during which the hydrogen-dominated intergalactic medium transitioned, under the radiation of the first stars, from the neutral transparent state it acquired 380K years after the Big Bang to the ionized state we observe today, about 14 billion years later. Characterizing the physical properties of these extremely faint high redshift galaxies requires the largest collecting surfaces to gather enough photons to perform spectroscopy. Additionally, in order to achieve sufficient statistics in a reasonable observing time, a large number of objects must be observed in parallel in the largest possible field of view (FoV). The efficiency of this observational technique, called Multi-object spectroscopy (MOS [3]), is further enhanced when coupled to high angular resolution because these distant sources ($z > 7$) are also very compact (radius < 0.5 arcsec). The European Extremely Large Telescope (E-ELT, [4]) will be the largest optical telescope to be built in the next decade, its first light being anticipated in 2024. With a diameter of 39m, almost 5 times greater than the state-of-the-art 8m Very Large Telescope (VLT) units, the collecting area of its primary mirror will allow astronomers to gain access to the early Universe. However, such a gigantic scientific instrument, with an estimated weight of more than 3500 T, cannot be placed in orbit and has to be built on the ground, below the atmospheric turbulence, which limits the achievable resolution to that of a telescope with a diameter of a few tens of centimeters.

Adaptive Optics (AO) is an instrumental technique, used on large astronomical telescopes since the 1990s [5], and intended to compensate for the effect of turbulence on image quality by means of one or several deformable mirrors (DM) flattening the wavefront to provide sharper images, close to the diffraction limit, at the focus of the telescope. It requires the use of one or several wavefront sensors (WFS) capturing the light from guide sources at a high frame rate (1ms), shorter than the coherence time of the turbulence, to measure the wavefront in real-time. The simple concept of AO with a single DM and single WFS is recalled in Figure 1a. In this classical form, so-called single conjugate, AO is not adequate to perform the wide field turbulence correction required to

(a) Single conjugate AO concept

(b) Open-Loop tomography concept.

(c) Observing the GOODS South cosmological field with MOAO.

Fig. 1. The MOAO methodology.

study cosmological fields, and tomographic AO techniques must be implemented for that purpose. This concept for a single science channel is depicted in Figure 1b. Several WFS are used to measure the wavefront in various directions using the light from guide stars. From these measurements, a real-time controller performs a tomographic reconstruction of the wavefront to produce commands to a dedicated DM in the science channel. In theory, a MOAO system can host an arbitrary number of WFS and an arbitrary number of science channel [6].

One of the instruments proposed for the E-ELT is MOSAIC [1], a multi-object integral field (multi-IFU) spectrograph coupled to a novel AO concept, called multi-object AO (MOAO). In this concept, the turbulence compensation is applied only on small *islands of interest* across a larger FoV, in the directions of the scientific targets observed by the MOS. This new

AO technique relies on the use of multiple WFS to perform a tomographic reconstruction of the turbulence induced wavefront perturbations in real-time. The MOAO system of MOSAIC will use not only field stars for the various WFS but also artificial stars, so-called laser guide stars, created in the mesosphere using high intensity lasers tuned to Sodium excitation lines. An example of observations with MOAO is depicted in Figure 1c with a fraction of the GOODS South field [7] spanning several arcminutes in radius. The field stars that can be used for guiding are encircled in blue, the position of the artificial stars to be created in the field is indicated with orange stars and the positions of possible science channel with green rectangles.

The AO system complexity scales with the square of the telescope diameter so the factor 5 foreseen in diameter planned for the E-ELT translates into a factor 25 in complexity. Moreover, in the case of MOAO, this number is multiplexed both by the number of targets that can be observed in parallel by the MOS, which can be as large as 20, and by the number of hours of exposure required to observe the faintest targets, which can span over several nights of observations. This significant scale factor strongly affects our ability to lead large scale numerical simulations of the system behavior in order to assess its scientific performance and take today the right decisions in the design trade-off studies. Indeed, the construction of this multi-million Euros instrument is based on preliminary design studies starting today.

In order to explore a large parameter space in the system dimensioning (i.e., the number of WFS, their position in the FoV or the density of measurements per WFS for instance), and optimize the instrument design to ensure that scientific requirements can be met at the lowest cost, thousands of simulation runs, with adjustable system configurations, each representing hours of observations at the telescope, must be performed. While Monte Carlo approaches, widely used in the astronomical instrumentation community, can adequately quantify the behavior of the system sub-components individually, they tend to be highly inefficient to address the system performance globally due to its high level of complexity and to the large ratio between the length of the observing time (hours) and the turbulence coherence time (few ms).

In recent years, a novel pseudo-analytical approach has been developed for tomographic AO system simulations based on the definition of an accurate error budget during a given observing run. In this approach, in order to obtain an estimate of the final image quality after hours of observations, the ideal output from the optical system is progressively degraded by various error terms, due for instance to measurement noise or to the limited coverage of the FoV by the WFS guide sources. These error terms can be estimated from a number of system and turbulence parameters and it can be shown that, in most cases, the dominant term in the error budget is the tomographic reconstruction error due to variations in the turbulence strength and distribution in altitude during the observations. This new simulation scheme is able to expose parallelism at several levels of the distributed-memory hardware systems.

III. CURRENT STATE-OF-THE-ART

The pseudo-analytical simulation scheme for tomographic adaptive optic (AO) systems has been described in [8] and [9]. It inherits from a number of various simulation techniques, such as Point Spread Functions (PSF) reconstruction, Fourier analysis and end-to-end simulation codes, but also to data analysis tools currently implemented on real systems for operations. It is mainly based on a matrix formalism to compute the covariance matrices of the residual tomographic error and an optimized approach for the computation of the tomographic reconstructor. A critical aspect in the simulation pipeline is the generation of the covariance matrices. These can be computed using matrix operators linking the different modal expansions of the phase between the layer-projected meta-pupils at various altitudes, which makes the computational load proportional to the size of the system. Instead, we propose to compute each covariance coefficient in a simplified analytical way opening the way to parallelization.

Monte Carlo simulations of a multi-object AO (MOAO) system at the scale of the E-ELT have already been performed by other authors [10]. They were able to address the system performance at a global level with varying system parameters and focus on a number of system sub-components limitations. This work allowed them to derive general instrument design rules. However, the equivalent observing time simulated is limited to few minutes as compared to the hours required for a realistic representation of a cosmological field observation. In their simulation, the strength and distribution in altitude of the turbulence is kept fixed which represents only a single snapshot of the conditions over the telescope during a deep long exposure observation.

The authors give very few details on their actual implementation and do not mention execution time for a single Monte Carlo loop nor give an estimate to overall time to solution. Extrapolating execution times from other optimized Monte Carlo simulation tools [11], [12] to this scale of system leads to an estimate of at best few Monte Carlo iterations per second. Simulating one hour of observation of a system running at 250 frames per second would thus take approximately a day in wall-clock time.

At the core of a MOAO system operations is the tomographic reconstructor (ToR) matrix. In a real system, a linear approach is used to compute the commands to be applied on the deformable mirrors (DM) from the wavefront sensors (WFS) measurements. The latter are concatenated into a single vector for each frame, and this vector is multiplied by the ToR matrix to produce a command vector, each component of the latter being the actual command applied to one of the DM's actuators. This is done few hundred times per second, depending on the chosen frame rate. The contribution from the tomographic reconstruction error in the MOAO system error budget thus depends on the ToR and the actual turbulence conditions during the observation. Thus, the efficient computation of the ToR is one of the paramount components, both for the simulations and also for the actual system operations.

Fig. 2. Overview of the process to generate short exposures images for each science channel.

Least squares approaches have been used since the 1980s to reconstruct the wavefront [13], and then for tomography [14], [15]. Briefly summarized, if $\vec{m}$ is a measurement vector (including noise) and $\vec{t}$ the quantity to be retrieved (in our case, the measurements that should be obtained in the direction of the scientific target), the linear application that minimizes the norm of $|\vec{t} - R\vec{m}|^2$ (or more generally the scalar product $(\vec{t} - R\vec{m})^T \cdot \Delta \cdot (\vec{t} - R\vec{m})$) is given by:

$$R = C_{tm} \cdot C_{mm}^{-1} \quad (1)$$

where C_{mm} stands for the covariance matrix between all the measurements of all the WFS of the instrument, and C_{tm} is the covariance matrix between $\vec{t}$, and all the other system measurements.

To compute the ToR in Eq. 1, a Cholesky-based solver is employed to take advantage of the positive-definiteness of the covariance matrix C_{mm} . The typical number of measurements which corresponds to the actual size of C_{mm} targeted for the E-ELT telescope is around 100K for a single FoV. The size of the right hand side C_{tm} depends on how many galaxies are contained in the FoV, since all galaxies share the same C_{mm} . These dense linear algebra (DLA) matrix operations have been well-studied in the literature and broadly available in the open-source shared-memory and distributed-memory libraries, LAPACK [16] and ScaLAPACK [17], respectively, as well as highly optimized vendor numerical libraries (e.g., Intel MKL [18] and Cray LibSci [19]). The first comprehensive high performance MOAO implementation [20] has been optimized on shared-memory systems equipped with hardware accelerators using the PLASMA library associated with the dynamic runtime system OmpSs [21]. However, the number of measurements is limited to 32K and only five galaxies have been investigated in the experiments.

We introduce here the first large scale MOAO implementation with 100K number of measurements for a single FoV, simulating 4K galaxies during a single night of real observations data coming from the CANARY tomographic AO demonstrator [22], installed on the WHT telescope in the Canary islands.

IV. IMPLEMENTATION DETAILS

A. Algorithmic Innovations

1) *The MOAO Pseudo-Analytical Method:* In our simulation approach the tomographic reconstruction error is expressed as a covariance matrix C_{ee} of the tomographic error vector $\vec{e}$ during a given short exposure. A full description of this simulation approach is given in [23] in which the reader will find detailed technical background on AO simulations as well. As recalled by the authors, the tomographic error covariance matrix can be expressed as:

$$C_{ee} = C_{tt} - C_{tm}R^T - RC_{tm}^T + RC_{mm}R^T \quad (2)$$

and we can see that replacing R by its expression given in Eq. 1 shows that the computation of matrix C_{ee} can be based on the covariance matrices C_{tt} representing the true sensor measurements, in addition to C_{tm} and C_{mm} , which are all covariance matrices of wavefront sensors (WFS) measurements.

To derive the actual phase structure function of the tomographic error C_{ee} for a given direction, it is then necessary to project it onto the space of the corresponding actuator of the deformable mirror using the pseudo-inverse of the so-called interaction matrix D . The former calibrates the interaction between the deformable mirrors (DM) and the true sensor (TS). This projected version of the covariance matrix, C_{vv} , is obtained as:

$$C_{vv} = DC_{ee}D^T. \quad (3)$$

The phase structure function is then obtained, under some given hypothesis, by interpolating the C_{vv} covariance function on a grid depending on the DM geometry. The PSF is then obtained by Fourier transforming the exponential of the opposite phase structure function. The validity of this approach has been extensively demonstrated in [23].

It should be noted here that these matrices all depend on atmospheric conditions. In a real experiment, the ToR would be computed using a series of turbulence parameters recorded at a given time during the night and then used on the system while the turbulence continues to evolve for a given time (typically tens of minutes to an hour) before it is updated. Realistic numerical simulations should thus address this dimension in the problem and be able to reproduce both the ToR computation for a given state of the atmosphere and successive short exposure images obtained with changing conditions. These short exposure images are then averaged to obtain one equivalent long exposure image per science channel. This process can then be repeated a number of times to obtain several long exposure images per channel and per night.

In contrast to Monte Carlo simulation approaches, in which this effect can only be simulated sequentially, the formalism of Eq. 1, Eq. 2 and Eq. 3 allows us to naturally benefit from it and expose more parallelism since all the atmospheric turbulence realizations are treated independently. The overall process to obtain a long exposure snapshot for each science channel from several short exposure images is depicted in Figure 2 and our implementation of this pipeline, is able to benefit both from the multiplex provided by the multiple science channels to

be assessed and from the independence of each turbulence realization. In this figure, the covariance matrices for the ToR and the tomographic error computations are generated in the so-called *Covmat* boxes.

The reader has to keep in mind that this pseudo-analytical approach does not include a number of terms in the overall MOAO error budget. In this pipeline, we concentrate on the dominant error term for MOAO: the tomographic error, as already discussed in [23]. An end-to-end simulation scheme, relying on a Monte Carlo approach would provide more accurate results, including all possible sources of errors in an AO system (see [24] for a full description of the typical AO error budget). However, since tomographic error is the main dimensioning parameter for a MOAO system, the results obtained from our pipeline give accurate trends for image quality evolution over varying seeing conditions and guide stars configuration, which are not strictly comparable with the output of a Monte Carlo simulation. This is discussed in more depth in [23]. It is nevertheless the most valuable input for a system design study. Considering the very large multiplex gain it provides, this input makes the pseudo-analytical approach the perfect tool for such system dimensioning studies, which require to span a wide range of values in the parameter space.

2) *Load Imbalance Issues*: Once the various covariance matrices have been generated, the pseudo-analytical model describing the MOAO simulation (Eq. 1 and 2) is mostly composed of dense matrix computations involving dense linear solvers (for the tomographic reconstructor calculation) and successive calls to Level 3 BLAS matrix operations (for the atmospheric turbulence integrations at each snapshot). The various data structures C_{mm} , R , C_{tm} , C_{tt} , C_{ee} and C_{vv} have irregular sizes, ranging from $100K \times 100K$ (due to the number of WFS=14 at full scale) for C_{mm} to $5K \times 5K$ for C_{vv} (due to the number of actuators at full scale). Considering the pipeline, the simulation starts therefore in the outer loop with computation on large data structures and finishes in the inner loop operating on much smaller data structures. This causes load imbalance that may prevent linear scaling on large number of computational nodes. Load balancing is not investigated in this study and will be beneficial to consider in the future.

3) *MOAO Kernel Portability*: One of the most important characteristics of this novel MOAO model is how the BLAS operations play a major role as basic blocks, not only in terms of performance, since BLAS have been vastly adopted by vendor numerical libraries, but also in terms of portability across different hardware architectures. In fact, it is critical to ensure the MOAO simulation can run on various architectures with BLAS-like performance since it is not clear how the hardware landscape will evolve in the next decade. This is of particular importance since this simulation pipeline includes some critical building blocks of the strategy of operations of the instrument, such as the efficient computation of the ToR. This will ensure maintainability of these operational tools over a long term, consistent with the lifetime of the instrument.

Fig. 3. Sketch of the MOAO implementation, showing the long exposure snapshot stage and the nested short exposure images stage.

B. Implementation Innovations

The high performance MOAO implementation consists of two phases: the long exposure snapshot in which occurs the calculation of several tomographic reconstructors (one for each detected science channel) and the short exposure images, which integrate the atmospheric turbulences across science channels, as depicted in Figure 3. The implementation is based on ScaLAPACK, which internally relies on the Message Passing Interface [25] for inter-node communications.

1) *The Long Exposure Snapshot*: All ToRs for each science channel are obtained by performing a Cholesky-based linear solver using ScaLAPACK block algorithms on distributed-memory systems. Using a two-dimensional block cyclic data distribution, the various matrices are scattered among a grid of tightly-coupled $P \times Q$ processors, where P is set to be the number of science channels and Q the number of cores per nodes (one MPI process per core). To avoid redundant calculations, we stack the right-hand sides C_{tm} from each science channel and solve the corresponding dense linear system of equations with a single Cholesky factorization, followed by the usual forward and backward substitutions. Because the science channels of the FoV are tightly-connected, data movement occurs only within a subset of all processes involved for a given long exposure snapshot, and therefore, thanks to the compute-intensiveness of the operation, data movement overheads are limited. The aforementioned mechanism holds when simultaneously operating long exposure snapshots, thanks to the Basic Linear Algebra Communication Subprograms (BLACS [26]) context, which behaves similar to MPI sub-communicators. Once this large solver has finished, each ToR solution will be used to feed the inner loop of the pipeline, i.e., the short exposure images, as described in the next Section.

2) *The Short Exposure Images*: Once this phase has been activated, each ToR for the science channels is read in parallel and used to perform each short exposure image concurrently, where atmospheric turbulence is successively applied to the ToR at each iteration (Eq. 2 and 3).

V. SYSTEM AND ENVIRONMENT SETTINGS

Our experiments have been conducted on two Cray systems. The first one, codenamed *Shaheen-2* from KAUST Supercomputing Lab, is a Cray XC40 system with the Cray Aries network interconnect, which implements a Dragonfly network topology. It has 6174 compute nodes, each with two-sockets Intel Haswell 16 cores running at 2.30GHz and 128GB of DDR3 main memory and we use Intel compiler suites v16.3.3.210. The theoretical peak performance for *Shaheen-2* is 7.27Pflops/s, while its sustained HPL performance has been recorded at 5.54Pflops/s.

The second system, codenamed *Cray-KNL*, is still a Cray XC with Aries interconnect but now featuring compute nodes with Intel Xeon Phi Knights Landing (KNL) processors. Every KNL has a base frequency of 1.4GHz and is equipped with 192GB DDR4-2400 RAM and 16GB MCDRAM. The theoretical peak performance of the *Cray-KNL* system is 1.56Pflops/s, while its sustained HPL performance has been recorded at 780Tflops/s. The *Cray-KNL* system is operated in *quadrant* mode while the memory is in *cache* mode, where all MCDRAM is configured as direct-mapped cache. We use 64 cores out of 68 on every KNL for computation, while 4 cores were dedicated to system services. Furthermore, hugepages were employed. Executables for the KNL target have been generated with the Intel compiler version 17.0.1.132. The work load managers on the *Shaheen-2* and the *Cray-KNL* systems are native SLURM and Moab/TORQUE+ALPS, respectively. On both systems, we rely on the ScaLAPACK implementation from the vendor Cray LibSci.

VI. SCIENCE PERFORMANCE ANALYSIS

In order to assess the scientific output of the simulation, we defined a typical simulation scenario outlined both in Table I and in Figure 4. The wavefront sensors (WFS) guide sources configuration is depicted in the upper left panel of Figure 4, where the position of the field stars used for guiding is indicated with blue dots and the position of Laser guide stars with red dots. The turbulence profiles used to simulate turbulence evolution are shown in the bottom left panel of Figure 4 and have been extracted from experimental data acquired with the CANARY demonstrator [22]. In this sub-figure, the strength of the turbulence is plotted as a function of altitude in the atmosphere (y-axis) and time (x-axis). The MOAO system parameters list in Table I is representative of a full scale instrument with maximum WFS number and density on 40m diameter telescope pupil and, thus, of the most complex simulation task to perform in the context of the MOSAIC design trade-off studies.

The output for the simulation of one observation for several science channels of the instrument is pictured in the right panel of Figure 4. Each image of this panel is a short exposure image (S.E.) in a direction of a given science channel, obtained for a single tomographic reconstructor (ToR) with a turbulence profile extracted from the turbulence profile timeline of the bottom left panel. The variation in the image shape can be explained by the variations in space and altitude of the

TABLE I
LIST OF SYSTEM AND TURBULENCE PARAMETERS.

System Parameters	
Telescope diameter	38.9 m
WFS number	up to 14
WFS density	76x76
Number of measurements for 1 WFS	10240
Total number of measurements	up to 102400
Number of actuators for 1 science channel	5336
Number science channels	up to 64x64
Turbulence Parameters	
Coherence scale (at $0.5\mu\text{m}$)	0.05 to 0.35 m
Number of layers	60
ToR update rate	every hour
Error computation rate	every 10 minutes

turbulence strength distribution between the various science channels direction.

As explained for instance in Section 6 of [9], from such a set of images, that can be as large as 64×64 images, one can extract 2 different kinds of observables for a single short exposure: a map of the Strehl ratio across the science field of view (FoV) and a map of the ensquared energy. Both observables are a measure of the image quality at the output of the instrument. The Strehl number is measured as the ratio between the maximum of the obtained image and the maximum of a diffraction limited image (i.e. as it would be if there was no turbulence) and the ensquared energy is measured as the integral of the image in a box around the maximum. The latter provides a measurement of the amount of power actually concentrated in this box, the size of which being tunable to the actual integral field spectrograph aperture size. This is, therefore, a measure of the actual sensitivity of the instrument.

The main intent of this preliminary experiment is to assess the output of the full simulation pipeline with actual experimental data as an input to emulate a realistic scenario for the evolution of atmospheric turbulence conditions during a night of observing. From this preliminary results, depicted in Figure 5, simulating different short exposures performance maps with varying turbulence conditions shows to be very consistent. First the performance maps peak at the location of the guide sources. This is the typical behavior to be expected for a tomographic AO system, in which turbulence reconstruction works better in the direction of the guide sources. Additionally, the evolution in the maps observed between the various turbulence profiles is also compatible with the content of these profiles. While turbulence weakens, as shown by the value of r_0 (a measure of the inverse of the turbulence strength), the performance increases globally on the map.

The pipeline can produce similar outputs for a tunable number of science channels and can also be called with a tunable number of short exposure computations and long exposure snapshots. While these very preliminary results are very encouraging, our goal is to produce high resolution maps of the instrument performance for a variety of scenario in terms of system dimensioning and turbulence distribution and

Fig. 4. Overview of the simulation input parameters and output short and long exposure images for a single science channel and during a full night.

evolution in the context of MOSAIC MOAO design studies. This will require an extensive use of the simulation pipeline at the largest scale possible.

VII. HPC PERFORMANCE RESULTS

In this section, we demonstrate the performance of our implementation for the large-scale MOAO framework on the two distributed-memory systems *Shaheen-2* and *Cray-KNL*, as described in Section V. Figures 6a and 7a show Gflops/s performance, while Figures 6b and 7b present time to solution when increasing the number of wavefront sensors (WFS), on *Shaheen-2* and *Cray-KNL*, respectively. Gflops/s is an important metric because it assesses the actual software implementation by showing how much of the hardware resources is being used, and in the case of a direct dense linear solver, it is essentially inversely proportional to the actual runtime. Performance and elapsed time increase while augmenting the number of WFSs, since the covariance matrix C_{mm} gets larger (up to a problem size of 102,400, see Table I), as expected. Since our simulation is mostly compute-bound for the most part, the achieved performance is decent on both systems for small to medium node counts compared to the theoretical peak performance, although this is not a tight upper-bound for the workload we consider in the simulation pipeline. However, the performance begins to be crippled due to the process idleness encountered because of load imbalance during the short exposure inner loop, as described in Section IV-A2. This behavior is captured on both systems, starting from 1024 and 512 nodes, on *Shaheen-2* and *Cray-KNL*, respectively. The elapsed time is however very promising since the MOAO simulation pipeline reaches an overall throughput of one Point Spread Function (PSF) generated at full system capacity in less than a second when WFSs is set to full scale, i.e., 14. This throughput is computed by dividing the total execution time by the number of PSFs generated. The gain in time to solution

compared to the traditional Monte Carlo method is quite impressive. Using the state-of-the-art end-to-end simulation pipeline COMPASS [27], considering the same system scale, one Monte Carlo iteration for a single target takes about 200 milliseconds to compute. To generate short exposures PSFs as in our pseudo-analytic pipeline requires about 10k Monte Carlo iterations, thus leading to a time to solution for a single short exposure PSF of about 200 seconds. With such PSF throughput, generating a high resolution performance map, as displayed in Figure 4, would take more than 230 hours or about 10 days. Additionally, beyond the execution time of a single Monte Carlo iteration, a significant overhead has to be added in the Monte Carlo scheme every time the atmospheric conditions are updated.

To better emphasize this load imbalance issue, Figure 8 reports the time breakdown of each computational phase of the overall MOAO simulation for WFS=14, when the system is at full scale, on the *Shaheen-2* system. The three computational stages, i.e., the *matcov*, the *ToR* and the *CeeCvv*, have different scalability, when we increase the number of nodes. While the generation of covariance matrices *matcov* and the *ToR* computation scale appropriately, the computation inside the short exposure inner loop does not and prevents, therefore, the scalability on larger node counts. This is again aligned with the load imbalance issue due to the irregular data structures, as discussed in Section IV-A2.

In summary, there is clearly room for further improvement, as outlined in the next section.

VIII. IMPACTS FOR MOAO MOVING FORWARD

A. Scalability Issue Discussion

There are some implications which have not been addressed in this paper and may alleviate the scalability limitation when scaling up to high core counts. Indeed, multi-object instruments such as MOSAIC are capable of observing multiple

Fig. 5. Ensquared energy (left) and Strehl (right) maps on a 64x64 grid for 4 different turbulence conditions (as reported by r_0) corresponding to 4 sets of short exposures in the science FoV. One ToR was generated per pixel of one image and the same ToR is used for the 4 different turbulence conditions.

(a) Strong scalability.

(b) Time to solution.

Fig. 6. Strong scaling and time to solution of the large scale MOAO simulation up to 6144 nodes (196,608 cores) on *Shaheen-2* system.

(a) Strong scalability.

(b) Time to solution.

Fig. 7. Strong scaling and time to solution of the large scale MOAO simulation up to 512 nodes (32,768 cores) on *Cray-KNL* system.

galaxies in parallel for a given FoV. While the evolution of atmospheric conditions influences globally the quality of the images obtained for every scientific target, it can be sampled in time, generated in common and then separately

applied on each single science channel for the computation of image quality. This may enable two levels of parallelism in the simulation, which each highlight two key ingredients, when moving forward with exascale simulations, i.e.,

Fig. 8. Time breakdown of WFS-14.

communication reduction and synchronization reduction. The former is achieved by tightly coupling the multiple science channels simulations, and therefore, removes the dependencies across multiple turbulence profile samples. Data movement then occurs between adjacent science channels of a given FoV, which share only a common tomographic reconstructor. The latter allows thread synchronizations to take place only within a shared-memory node, removing expensive global synchronizations, seen in the bulk synchronous programming model.

The overall simulations of MOAO for the E-ELT can then be translated into a batch of independent long exposures (LE) computation, where each LE processes a group of science channels for common atmospheric conditions evolution. The influence of the regular atmospheric turbulence changes can then be simulated independently on each science channel in an embarrassingly parallel fashion. The main goal of this is to generate the high resolution performance map. Although critical, this new dimension of parallelism may require to move away from the static data distribution, which is at the core of existing state-of-the-art dense linear algebra libraries, e.g., ScaLAPACK. A more dynamic data distribution may enable to perform work stealing at runtime, and therefore, may cope with the challenging load balancing issues faced at large-scale. This is still an open research problem in the dense linear algebra and runtime communities, since typical dense applications would operate on data structures with homogeneous sizes. In summary, this problem is beyond the scope of this paper and may be addressed in a future work.

B. Moving Forward with MOSAIC

The preliminary design studies of multi-million Euros astronomical instruments such as MOSAIC for the E-ELT require to explore a large parameter space including science channels distribution across the FoV, guide star brightness and positions, laser guide star constellation diameters, atmospheric variability, several observation wavelengths and more. Building high resolution image quality maps to monitor the evolution of scientific performance across this parameters space will require the computation of millions of ToR and tens of millions of short exposure images. Leading such a large scale simulation campaign, in order to find the right trade-off in the instrument design and both meet the scientific requirements and optimize the cost, is highly dependent on

efficient numerical simulations. While Monte Carlo based approaches would require several days to simulate a single long exposure image, considering that one short exposure Point Spread Function (PSF) can be generated at best in few hours, we have shown that a whole night of observation, including multiple long exposure images, for several science channels, can be simulated in about an hour.

Moreover, some building blocks of the simulation pipeline, such as the ToR computation stage, are also core components of the system operational strategy. This simulation scale opens new research horizons in numerical methods for experimental astronomy, with some core developments standing as pathfinders toward actual operations and future astronomical discoveries on the E-ELT. Leveraging these high performance linear algebra techniques has shown to be very efficient in providing the multiplex gain required to accelerate significantly the simulation process, enabling a wide coverage of parameters in the system dimensioning studies. Similar approaches could be leveraged to assist the actual system operations when commissioned and increase its performance and stability on sky, hence its science throughput.

As demonstrated in [20], since the MOAO simulation relies extensively on compute-bound kernels provided by the vendors' optimized numerical BLAS libraries, hardware accelerators appear to be a valid path moving forward and can significantly leverage the performance numbers reported here. This simulation is not architecture-dependent, and therefore can possibly run on various hardware vendors. Energy efficiency is also critical for astronomers since the telescope sites have usually limited power envelope. Energy efficient hardware solutions (e.g., ARM processors) may be considered in the near future and porting this MOAO simulation may not require extensive work, thanks to its strong BLAS dependence.

IX. CONCLUSION

This paper highlights the real-time massively distributed multi-object adaptive optics simulations for the European Extremely Large Telescope. This simulation allows the user to create a high resolution performance map of 4K positions in the science FoV and achieves a throughput of one PSF generated every 4 seconds, which makes it real-time. This simulation is not only key to the design studies for MOSAIC but also to the actual instrument construction, as the core of the simulation, the tomographic reconstructor computation, is also at the core of the instrument operation. The next steps, once we introduce a new level of parallelism, which consists in processing multiple positions simultaneously in the short-exposure inner loop, will be to increase the number of field positions, produce multi-wavelength PSFs, and simulate even larger field of views, which consists in processing multiple positions simultaneously in the short exposure inner loop. While this will certainly increase performance and reduce time to solution, the load imbalance issues may still be present for extreme scale MOAO simulations. Fine-grained computations using task-based programming model associated with dynamic runtime systems [28], [29] may mitigate this load imbalance

overhead by further keeping all processing units busy, while overlapping communications with computations.

ACKNOWLEDGMENT

The authors would like to thank Aniello Esposito from Cray Inc. for his help in running the code and the vendor Cray for systems remote accesses in the context of the Cray Center of Excellence awarded to the Extreme Computing Research Center at KAUST. For computer time, this research used the resources from KAUST Supercomputing Laboratory for *Shaheen-2* core hours allocation.

REFERENCES

- [1] F. Hammer, B. Barbuy, J. G. Cuby, L. Kaper, S. Morris, C. J. Evans, P. Jagourel, G. Dalton, P. Rees, M. Puech, M. Rodrigues, D. Pearson, and K. Disseau, "MOSAIC at the E-ELT: A multi-object spectrograph for astrophysics, IGM and cosmology," in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9147, Aug. 2014, p. 914727.
- [2] L. S. Blackford, J. Choi, A. Cleary, E. F. D'Azevedo, J. W. Demmel, I. S. Dhillon, J. J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. W. Walker, and R. C. Whaley, *ScaLAPACK Users' Guide*. Philadelphia: Society for Industrial and Applied Mathematics, 1997.
- [3] C. J. Evans, et al., "Science case and requirements for the MOSAIC concept for a multi-object spectrograph for the European Extremely Large Telescope," in *Ground-based and Airborne Instrumentation for Astronomy V*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9147, Jul. 2014, p. 914796.
- [4] A. McPherson, J. Spyromilio, M. Kissler-Patig, S. Ramsay, E. Brunetto, P. Dierickx, and M. Cassali, "E-ELT update of project and effect of change to 39m design," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 8444, Sep. 2012.
- [5] R. Davies and M. Kasper, "Adaptive Optics for Astronomy," *Annual Review of Astronomy and Astrophysics*, vol. 50, pp. 305–351, Sep. 2012.
- [6] G. Rousset, et al., "EAGLE MOAO system conceptual design and related technologies," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 7736, Jul. 2010.
- [7] M. Giallisco, et al., "The Great Observatories Origins Deep Survey: Initial Results from Optical and Near-Infrared Imaging," *The Astrophysical Journal*, vol. 600, pp. L93–L98, Jan. 2004.
- [8] A. Charara, H. Ltaief, D. Gratadour, D. E. Keyes, A. Sevin, A. Abdelfattah, E. Gendron, C. Morel, and F. Vidal, "Pipelining Computational Stages of the Tomographic Reconstructor for Multi-Object Adaptive Optics on a Multi-GPU System," in *SC'14*. IEEE, 2014, pp. 262–273. [Online]. Available: <http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=7012142>
- [9] C. Morel, E. Gendron, D. Gratadour, A. Sevin, and G. Rousset, "Pseudo-analytic simulation of woofer-tweeter MOAO system: application to MOSAIC," in *Adaptive Optics Systems V*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9909, Jul. 2016, p. 990976.
- [10] A. G. Basden, N. A. Bharmal, R. M. Myers, S. L. Morris, and T. J. Morris, "Monte Carlo simulation of ELT-scale multi-object adaptive optics deformable mirror requirements and tolerances," *Monthly Notice of the Royal Astronomical Society*, vol. 435, pp. 992–998, Oct. 2013.
- [11] D. Gratadour, M. Puech, C. Véraud, P. Kestener, M. Gray, C. Petit, J. Brulé, Y. Clénet, F. Ferreira, E. Gendron, M. Lainé, A. Sevin, G. Rousset, F. Hammer, I. Jégouzo, M. Paillous, S. Taburet, Y. Yang, J.-L. Beuzit, A. Carlotti, M. Westphal, B. Epinat, M. Ferrari, T. Gautrais, J. C. Lambert, B. Neichel, and S. Rodionov, "COMPASS: an efficient, scalable and versatile numerical platform for the development of ELT AO systems," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9148, Aug. 2014, p. 6.
- [12] A. G. Basden, T. Butterley, M. A. Harrison, T. J. Morris, R. M. Myers, R. W. Wilson, E. Younger, T. Fusco, B. Le Roux, and B. Neichel, "Performance of Monte-Carlo simulation of adaptive optics systems of the EAGLE multi-IFU instrument for E-ELT," in *Adaptive Optics Systems*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 7015, Jul. 2008, p. 70156Y.
- [13] E. P. Wallner, "Optimal wave-front correction using slope measurements," *Journal of the Optical Society of America (1917-1983)*, vol. 73, p. 1771, Dec. 1983.
- [14] T. Fusco, J. Conan, G. Rousset, L. M. Mugnier, and V. Michau, "Optimal wave-front reconstruction strategies for multiconjugate adaptive optics," *Journal of the Optical Society of America A*, vol. 18, pp. 2527–2538, Oct. 2001.
- [15] F. Vidal, E. Gendron, and G. Rousset, "Tomography approach for multi-object adaptive optics," *Journal of the Optical Society of America A*, vol. 27, no. 26, p. A260000, Oct. 2010.
- [16] E. Anderson, Z. Bai, C. Bischof, S. L. Blackford, J. W. Demmel, J. J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. C. Sorensen, *LAPACK User's Guide*, 3rd ed. Philadelphia: Society for Industrial and Applied Mathematics, 1999.
- [17] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley, *ScaLAPACK Users' Guide*. Philadelphia, PA: SIAM, 1997, <http://www.netlib.org/scalapack/slgl/>.
- [18] Intel, "Math Kernel Library," Available at <http://software.intel.com/en-us/articles/intel-mkl/>.
- [19] Cray, "Libsci," <http://docs.cray.com>.
- [20] H. Ltaief, D. Gratadour, A. Charara, and E. Gendron, "Adaptive optics simulation for the world's largest telescope on multicore architectures with multiple gpus," in *Proceedings of the Platform for Advanced Scientific Computing Conference*, ser. PASC '16. New York, NY, USA: ACM, 2016, pp. 9:1–9:12. [Online]. Available: <http://doi.acm.org/10.1145/2929908.2929920>
- [21] A. Duran, E. Ayguadé, R. M. Badia, J. Labarta, L. Martinell, X. Martorell, and J. Planas, "OmpSs: a proposal for programming heterogeneous multi-core architectures," *Parallel Processing Letters*, vol. 21, no. 2, pp. 173–193, 2011. [Online]. Available: <http://dx.doi.org/10.1142/S0129626411000151>
- [22] F. Vidal, E. Gendron, G. Rousset, T. Morris, A. Basden, R. Myers, M. Brangier, F. Chemla, N. Dipper, D. Gratadour, D. Henry, Z. Hubert, A. Longmore, O. Martin, G. Talbot, and E. Younger, "Analysis of on-sky MOAO performance of CANARY using natural guide stars," *Astronomy & Astrophysics*, vol. 569, Jun. 2014.
- [23] É. Gendron, A. Charara, A. Abdelfattah, D. Gratadour, D. Keyes, H. Ltaief, C. Morel, F. Vidal, A. Sevin, and G. Rousset, "A novel fast and accurate pseudo-analytical simulation approach for MOAO," in *Adaptive Optics Systems IV*, ser. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 9148, Aug. 2014, p. 91486L.
- [24] F. Ferreira, E. Gendron, G. Rousset, and D. Gratadour, "Deriving comprehensive error breakdown for wide field adaptive optics systems using end-to-end simulations," in *Adaptive Optics Systems V*, vol. 9909, Jul. 2016, p. 990979.
- [25] M. P. I. Forum, "MPI-3: Extensions to the Message Passing Interface Standard," <http://www.mpi-forum.org/>, July 2015.
- [26] J. Dongarra and R. C. Whaley, "A user's guide to the BLACS v1.1," Department of Computer Science, University of Tennessee, Knoxville, Knoxville, TN 37996, USA, LAPACK Working Note 94, May 1997, updated May 5, 1997 (Version 1.1). [Online]. Available: <http://www.netlib.org/lapack/lawns/lawn94.pdf>; <http://www.netlib.org/lapack/lawns/lawn94.ps>
- [27] D. Gratadour, F. Ferreira, A. Sevin, N. Doucet, Y. Clénet, E. Gendron, M. Lainé, F. Vidal, J. Brulé, M. Puech, C. Véraud, and A. Carlotti, "COMPASS: status update and long term development plan," in *Adaptive Optics Systems V*, vol. 9909, Jul. 2016, p. 990971.
- [28] C. Augonnet, S. Thibault, R. Namyst, and P. Wacrenier, "StarPU: A unified platform for task scheduling on heterogeneous multicore architectures," *Concurrency Computat. Pract. Exper.*, vol. 23, pp. 187–198, 2011, (to appear).
- [29] G. Bosilca, A. Bouteiller, A. Danalis, T. Herault, P. Lemarinier, and J. J. Dongarra, "DAGuE: A generic distributed DAG engine for high performance computing," *Parallel Computing*, 2011, to appear, <http://dx.doi.org/10.1016/j.parco.2011.10.003>.