Tunable Approximations to Control Time-to-Solution in an HPC Molecular Docking Mini-App

Davide Gadioli, Gianluca Palermo, Stefano Cherubin, Emanuele Vitali, Giovanni Agosta, Candida Manelfi, Andrea R. Beccari, Carlo Cavazzoni, Nico Sanna, and Cristina Silvano

Abstract—The drug discovery process involves several tasks to be performed *in vivo*, *in vitro* and *in silico*. Molecular docking is a task typically performed *in silico*. It aims at finding the three-dimensional pose of a given molecule when it interacts with the target protein binding site. This task is often done for virtual screening a huge set of molecules to find the most promising ones, which will be forwarded to the later stages of the drug discovery process. Given the huge complexity of the problem, molecular docking cannot be solved by exploring the entire space of the ligand poses. State-of-the-art approaches face the problem by sampling the space of the ligand poses to generate results in a reasonable time budget. In this work, we improve the geometric approach to molecular docking by introducing tunable approximations. In particular, we analyzed and enriched the original implementation with tunable software knobs to explore and control the performance-accuracy tradeoffs. We modeled time-to-solution of the virtual screening task as a function of software knobs, input data features, and available computational resources. Therefore, the application can autotune its configuration according to a user-defined time budget. We used a Mini-App derived by LiGenDock – a state-of-the-art molecular docking application – to validate the proposed approach. We run the enhanced Mini-App on an HPC system by using a very large database of pockets and ligands. The proposed approach exposes a time-to-solution interval spanning more than one order of magnitude with accuracy degradation up to 30%, more in general providing different accuracy levels according to the needs of the virtual screening campaign.

Index Terms—Approximate Computing, Application Autotuning, Molecular Docking, Virtual Screening.

✦

1 INTRODUCTION

The end of Dennard scaling pushed towards new sources of computing efficiency. In this direction, one of the most recent and promising trends is the introduction of approximations. Approximate computing wants to save unnecessary computation efforts while finding results whose quality is considered good enough by the end-user. The reasons behind the increasing popularity of this trend lie in the changing nature of the workloads driving the computing demand [1], [2]. Data mining, computer vision, machine learning, audio and video processing, dynamic simulations, and other classes of applications exhibit an intrinsic resilience to approximations. In these cases, the purpose of the computing system is not to calculate a precise numerical solution. Indeed, the *correctness* of such systems is defined as a constraint on a metric that represents the quality of the results in the continuum, i.e. a *good enough* result. The behavior of these applications is mainly related to the type of input/output data, the absence of a unique result, the complexity of finding the correct answer, and the way used to find the results. More specifically, approximate computing covers the following cases: (i) Applications designed to deal with noisy input data (such as data coming from

sensors), or applications whose end-users are not capable to differentiate among precise results (such as human eyes or ears); (ii) A unique solution does not exist and a range of solutions are equally acceptable, e.g. recommendation systems or web searches; (iii) Although a unique solution exists, the search algorithm is not guaranteed to find it due to the complexity of the solution space (such as in the cases of heuristic searches and machine learning applications); (iv) Applications based on patterns, such as iterative-refinement, which guarantee that the precise version is eventually reached (i.e. Montecarlo methods).

Approximate computing broadly refers to techniques that exploit the intrinsic resilience of applications to realize improvements in efficiency at different layers of the computing stack [3]. However, in this paper, we focus on the context of software approximations.

The quest for computational effectiveness is not only pushing new software domains towards HPC infrastructures, but also towards legitimizing approximate results in domains traditionally bound to absolute exactness. One example comes from the pharmaceutical domain and, more specifically, from the drug discovery process. The problem consists of finding new drugs starting from a huge exploration space of possible molecules. Drug discovery is typically composed of several steps, one of those is the virtual screening procedure to find the most promising molecules to interact with a target protein. The intrinsic characteristics of the virtual screening let us to classify it as a heuristic search, which is one of the aforementioned categories well-suited for approximate computing.

In this context, the domain-expert belonging to a pharmaceutical company defines the input set for the virtual screening application and the time-to-solution allocated to

- D. Gadioli, G. Palermo, S. Cherubin, E. Vitali, G. Agosta and C. Silvano are with Dipartimento di Elettronica, Infomazione e Bioingegneria, Politecnico di Milano, Italy.
- C. Manelfi and A. Beccari are with Dompe' Farmaceutici SpA Research Center, L'Aquila, Italy.
- C. Cavazzoni and N. Sanna are with the Supercomputing Innovation and Application Department, CINECA, Bologna/Roma, Italy.

the job. Typically, the task of defining the input set fits well the domain-knowledge of the company, while the task of defining the time-to-solution is more complex. Usually, the company defines a time budget set by the experiment cost and it relies on a domain-expert to tune the size of the molecule database accordingly. This workflow limits the exploration space without any guarantee neither to find a global optimum nor to find a good local optimum. Therefore, a reduction in the time spent on evaluating a single pair molecule-protein enables the end-user to explore a larger set of candidates, thus increasing the probability to find an interesting solution.

In this paper, we focus on *GeoDock-MA*, a molecular docking Mini-App for High-Performance Computing (HPC) systems, based on the LiGenDock module [4], to capture the geometrical features only. More in general, Mini-Apps represent an important aid for evaluating the architecture and algorithm design space to be explored in the early stages of the code development. *GeoDock-MA* attempts to capture key computation kernels of the molecular docking application for the drug discovery implemented in LiGenDock and exploiting only geometrical features. By developing *GeoDock-MA* in parallel with the new version of LiGenDock, application developers can work with system architects and domain-experts to evaluate alternative algorithms to better satisfy the end-user constraints or to better exploit the architectural features. *GeoDock-MA* enables us a faster performance analysis and the optimization of the key kernels.

The main goal of the proposed approach is to provide tunable approximations to explore performance-accuracy tradeoffs during the docking task, and autotuning [5] to support the optimization phase. In this paper, tradeoffs are enabled enhancing *GeoDock-MA* by exposing five tunable software knobs. Then, we autotune them with the support of a predictive model to control time-to-solution in the virtual screening task. In particular, the main contributions of this paper can be summarized as follows:

- We analyze *GeoDock-MA* to properly introduce approximate computing techniques on the most significant kernels;
- We enable performance/accuracy tradeoffs by exposing tunable software knobs to drive the *GeoDock-MA* approximations;
- We present the *GeoDock-MA* performance model based on the exposed software knobs and input data features for estimating time-to-solution in a virtual-screening process;
- We enhanced *GeoDock-MA* with an autotuning layer to satisfy the user-defined time budget according to the workload characteristics.

The paper is organized as follows. In Section 6, we present the background on drug discovery, virtual screening and molecular docking needed to understand the context of the work. Section 3 introduces the proposed methodology for parameterizing, approximating and autotuning the application implementation. We introduce the experimental setup in Section 4, while in section 5 we quantitatively evaluate the proposed approach. Section 6 presents the related work on molecular docking techniques, approximation methods and autotuning. Finally, Section 7 summarizes the overall findings and concludes the paper.

2 BACKGROUND

The goal of a drug discovery process is to find new drugs starting from a huge exploration space of candidate molecules. Typically, this process involves several *in vivo*, *in vitro* and *in silico* tasks ranging from chemical design to toxicity analysis. Molecular docking represents one stage of this process [6], [7]. It aims at estimating the three-dimensional pose of a given molecule, named *ligand*, when it interacts with the target protein. The ligand is much smaller than the target protein, therefore we focus on a small region of the target protein (or receptor), named *pocket* (or binding site). Given the three-dimensional pose of the ligand within the pocket, we can estimate the strength of the chemical and physical interactions between the ligand and the pocket by computing a geometric fitting score.

The evaluation of the pose of each ligand is independent of the evaluation of all the other candidates. In drug discovery, being the number of ligands to be analyzed above the billion of units, the problem can be considered embarrassingly parallel. However, to find the three-dimensional pose of the ligand when it interacts with the pocket, we have to deal with a large number of degrees of freedom. While the target pocket is represented as a rigid structure, the ligand is represented as a flexible set of atoms bound together by chemical bonds, i.e. sharing electron pairs between atoms (covalent bond). From a purely geometrical point of view, it is possible to identify a subset of bonds – named *rotamers* – which can split the ligand into two disjoint non-empty fragments when they are removed. We can independently rotate each of those fragments without altering the chemical connectivity of the ligand. Therefore, we have to consider changes in the shape of the ligand that can be obtained through the rotation of all its rotatable fragments.

Evaluating the chemical and physical interactions between the ligand and the pocket is a computationally-intensive problem, therefore state-of-the-art approaches [8], [9], [10] suggest to separate the pose prediction task from the virtual screening task. The pose prediction task focuses on providing the best pose for a given ligand in a given binding site, whereas the virtual screening task aims at selecting in a huge database of candidates a small set of promising ligands which best fit the given binding site. The structure of the two tasks are very similar to each other. Several industrial applications [4], [11] provide both functionalities in the same software module. A remarkable difference between the pose prediction and the virtual screening task lies in the approach used for the estimation of the chemical and physical interactions between the ligand and the pocket. It is possible to estimate such interactions by using either a geometrical or a pharmacophoric approach. The geometrical approach estimates the ligand-pocket interactions by using information related only to the shape and volume, while the pharmacophoric approach evaluates the actual chemical and physical interactions. The latter approach is the most computational-intensive one and it is regularly exploited on the pose prediction task. Although the best solution

according to a pharmacophoric score implies a good geometrical score, there are no guarantees on the opposite case. The best solution according to a geometrical score might be either a non-valid solution or a poor solution from a pharmacophoric perspective. Therefore, there is always the need to apply a pharmacophoric docking process. The geometrical approach is used for filtering all the solutions that cannot geometrically fit the target pocket.

In this paper, we exploit the application parameterization targeting approximation and autotuning for controlling the time-to-solution in a virtual screening campaign. This is done according to the target database of molecules (i.e. ligands) to be docked and to the characteristics of the protein binding site (i.e. pocket).

3 METHODOLOGY

This section first introduces *GeoDock-MA* by describing the algorithm and highlighting the application hot spots. Then, we describe the functional analysis used to drive the approximations of the application enabling the accuracy-throughput tradeoffs. Finally, we describe how end-users might leverage the exposed tradeoff for autotuning the application according to the time-to-solution constraints and to the workload characteristics.

3.1 Application Description

In the context of the LiGen toolflow [12], LiGenDock [4] is the module for docking one or more ligands into a target protein. It can be used for both the pose prediction and the virtual screening tasks. It exploits chemical and geometrical features to dock the ligand through an iterative algorithm. In particular, LiGenDock uses chemical features to set the initial pose of the ligand and to drive the docking process at each iteration. However, it also uses geometrical features to optimize the pose of the ligand at each iteration, by taking into account all the degrees of freedom of the problem space.

The optimization of the ligand pose is the most computationally-intensive part of LiGenDock during the virtual screening task. *GeoDock-MA* includes all the functionalities of LiGenDock that optimize the ligand pose by exploiting the geometric approach. *GeoDock-MA* takes as input the target pocket and a database of ligands, and it generates as output the score of each pocket-ligand pair, after optimizing the pose of the ligand.

GeoDock-MA performs the virtual screening task by using the geometric approach. It estimates the pocket-ligand interactions with the similarity between the shape of the ligand and the three-dimensional shape of the pocket in PASS format [13], which is generated by LiGen PASS [12]. Actually, *GeoDock-MA* scores each ligand by using the *overlap score* function. The overlap score, as defined in Equation 1, is the reciprocal of the minimum square distance between the ligand and the pocket:

$$o = \frac{l}{\sum_{i=0}^{l} \min_{j=0}^{p} d^2(L[i], P[j])} \quad (1)$$

where o is the overlap score, l is the number of atoms in the ligand L, p is the number of 3D points in the pocket
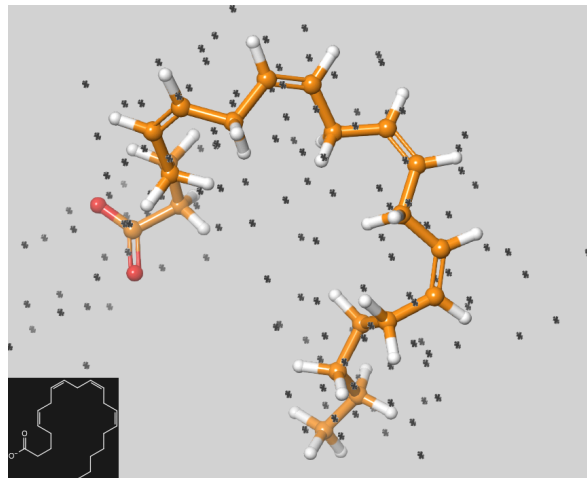

Fig. 1. 3D visualization of a docked ligand (connected structure) inside a PASS version of the target pocket (dark spots).

P, and d^2 represents the squared distance between the i-th atom of the ligand and the j-th point of the pocket. Higher overlap means better geometric compatibility between the pocket and the ligand.

Figure 1 shows an example of docking a ligand inside a pocket (i.e. 1cvu [14]). The ligand structure is clearly visible in the 3D-image and its planar representation is highlighted in the bottom left corner. The larger bubbles are the atoms L of the ligand, while the connections between atoms are the molecule bonds. The dark spots in Figure 1 are the points P representing the PASS version of the target pocket. These points are the centre of the spheres used to model the binding site.

3.2 Analysis of *GeoDock-MA*

GeoDock-MA is designed to target an HPC platform and it leverages machine-level parallelism by using the MPI master/slave paradigm. In particular, the master process reads the ligands' input database and dispatches a ligand to be evaluated to any available slave. Each slave docks geometrically the ligand in the binding site of the target molecule, which is constant for all the evaluated ligands. Finally, the slave computes the final geometrical score, sends its value to the master process, and waits for the next ligand to be evaluated. *GeoDock-MA* (similar to the original LiGenDock) leverages the parallelism at data-level, where each slave evaluates a pocket-ligand pair independently. Given the huge number of ligands to be evaluated, the benefit of this approach is twofold. On one side, the docking algorithm is serial, therefore it avoids any synchronization mechanism in the evaluation of the ligand-pocket pair. On the other side, the application can efficiently leverage the embarrassingly parallel nature of the problem at a high-level.

We profiled the application to identify the critical sections of the code by using GPROF[1]. Figure 2 shows the Call Graph report: It groups the individual functions by the caller.
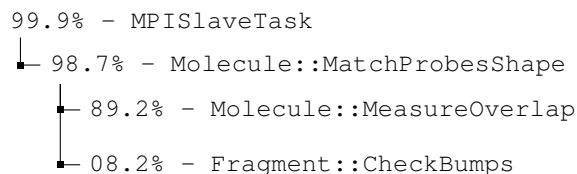
1. GNU gprof https://sourceware.org/binutils/docs/gprof/

```
99.9% - MPISlaveTask
├── 98.7% - Molecule::MatchProbesShape
    ├── 89.2% - Molecule::MeasureOverlap
    ├── 08.2% - Fragment::CheckBumps
```

Fig. 2. Application Call Graph profile. Functions taking less than 2% of the overall execution time are omitted
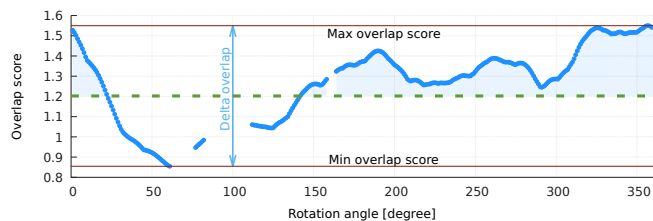

Fig. 3. Changes in the overlap score by rotating a fragment of the ligand. The x-axis represents the angle of the rotation, while the y-axis represents the overlap score of the ligand.

Data: the pocket and the 3D structure of the ligand
Result: the overlap score of the ligand

1 get the list of rotamers;
2 **foreach** *rotamer* **do**
3   grow the right and left fragment;
4   **for** *angle in 0-360 degrees with step 1 degree* **do**
5   rotate left fragment to *angle*;
6   **if** *the ligand shape is feasible* **then**
7   measure the overlap of the ligand;
8   check if the overlap is improved
9   **end**
10   **end**
11   set the left fragment to best angle found;
12   **for** *angle in 0-360 degrees with step 1 degree* **do**
13   rotate right fragment to *angle*;
14   **if** *the ligand shape is feasible* **then**
15   measure the overlap of the ligand;
16   check if the overlap is improved
17   **end**
18   **end**
19   set the right fragment to best angle found;
20 **end**
21 **return** *the overlap score of the ligand;*

Algorithm 1: Pseudo-code of the `MatchProbesShape` kernel, which changes the shape of the ligand to maximize the overlap score.

A large fraction of the execution time of the application is spent on `MatchProbesShape`. This kernel is responsible for the optimization of the shape of the ligand by using a steepest descent algorithm to deal with all the internal degrees of freedom of the ligand. In this paper, we focus on the introduction of possible software knobs through approximation techniques to tune the time-to-solution of this functionality.

Algorithm 1 shows the pseudocode of the target kernel. First, the algorithm identifies the set of *rotamers* (line 1), thus selecting all the possible sources of flexibility in the ligand shape. Then, it iterates over the set of these bonds to find the best shape of the ligand (lines from 2 to 20). In particular, the body of the algorithm grows a left and a right ligand fragment, with respect the two extremes of the bond (line 3). The left and the right fragments are rotated independently. The first to be processed is the left fragment. It is rotated step-by-step up to a full angle (lines from 4 to 5); At each step, we check whether the ligand shape is valid. There is a non-null possibility of internal bumping of the molecule (line 6), which invalidates the shape of the ligand. If the ligand shape is valid, the overlap score of the ligand

is considered during the check for possible improvements (lines from 7 to 9). At the end of the 360-degree exploration, we rotate once more the left fragment to match the angle that maximizes the overlap score (line 11). The same procedure is applied to the right fragment (lines from 12 to 19).

The kernel applies the same computation of the left and right fragments of each rotamer. For this reason, in the rest of the paper, we do not differentiate the two fragments.

Figure 2 also shows that the computation of the overlap score of each pose (`Molecule::MeasureOverlap`) represents the most expensive operation. The implementation of `Molecule::MeasureOverlap` is relatively simple and a lot of effort has already been spent in the past in terms of performance tuning. Our contribution aims at reducing the number of invocations performed by its caller. In particular, we want to avoid the computations which are very likely to do not lead to any improvement.

Figure 3 shows an example of how the rotation of a fragment affects the overlap score of the ligand. The x-axis represents the rotation space, while the blue line shows the overlap score of the ligand according to the position of the fragment. The empty spaces are due to the fact that some of the generated poses of the ligands are not valid because of the internal bumps of the ligand atoms. We define as *delta overlap* the difference between the minimum and maximum overlap of a single rotation. We define as *peak* the set of contiguous and valid rotation angles whose overlap is higher than 50% of the delta.

To further analyze the functional behaviour, we performed an experimental campaign by using a chemical library composed of 113K ligands. This analysis aims at finding patterns that can be leveraged for reducing the number of evaluations for each fragment rotation, thus creating possible application knobs (see Figure 4).

Figure 4a correlates the size of a fragment with its impact on the final score of the ligand. In particular, the x-axis represents the relative size of the fragment with respect to the size of the ligand, while the y-axis represents the delta overlap normalized with respect to the final score of the ligand. It is worth noticing that small fragments have small deltas, which means that such fragments usually have a limited impact on the numeric value of the final score of the ligand.
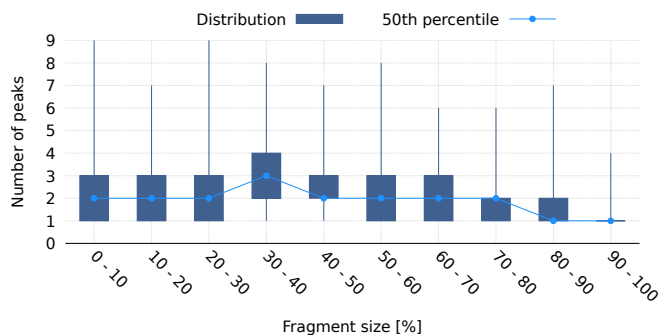
Figure 4b correlates the width of a *peak* (in degree) with its height normalized with respect to the delta overlap. From this plot, we can notice that the peaks that contain the maximum overlap are usually greater than 50 degrees, while narrow peaks rarely reach the maximum height. We can
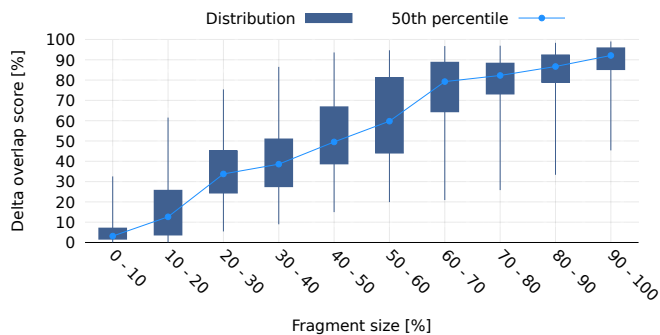
(a) Distribution of the delta overlap.
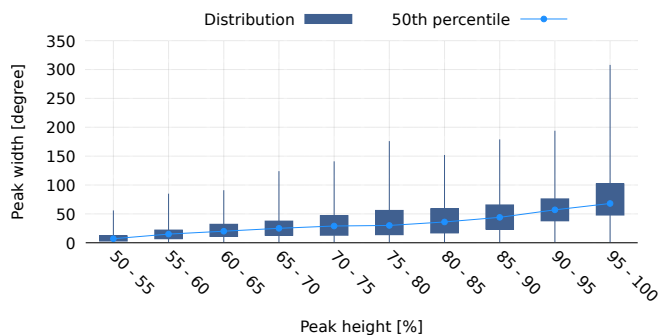
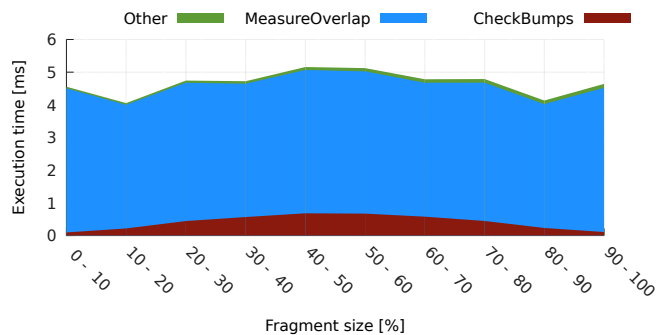(b) Distribution of the peak geometry.

(c) Distribution of the number of peaks.

Fig. 4. Analysis on the peaks of overlap across different fragments. Each plot shows the minimum value, the 25th, 50th, 75th percentile and the maximum value.

(a) `MatchProbesShape` execution time composition.

(b) Frequency distribution of the fragment sizes.

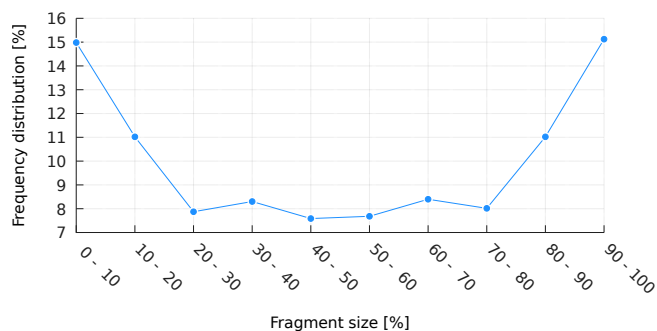Fig. 5. Analysis of execution time and frequency distribution of fragments, grouped by their size.

conclude that the behaviour of the overlap score is rather smooth because small peaks are also narrow.

Figure 4c shows on the y-axis the number of peaks which are contained in a fragment by changing the fragment size on the x-axis. We can notice how larger fragments usually have only one peak, while smaller ones tend to have more peaks.

Besides the functional behaviour of the overlap score, Figure 5a shows the detailed composition of the time spent by the application to find the best rotation angle of a fragment of the ligand (y-axis) according to the size of the fragment (x-axis). From the execution time, we highlight the time spent by measuring the overlap score (Measure-Overlap) and the time spent by checking if the evaluated angle is admissible (CheckBumps). From the plot, we see that computing the overlap score is independent of the size of the fragment. This result is expected because it involves

the evaluation of the whole ligand.

Finally, Figure 5b shows the frequency distribution of the size of a fragment. Due to the definition of the ligand database, smaller and larger fragments are slightly more frequent with respect to the other fragment sizes.

3.3 Exposing Tunable Application Knobs

In the original LiGenDock [4], the authors listed several parameters that can alter the behaviour of the docking algorithm. However, most of them are chemical-specific parameters that do not impact the execution time of *GeoDock-MA*. The only exception is a constant parameter which performs loop perforation [15] on the loops that rotates a fragment of the ligand (lines from 4 to 5 in Algorithm 1). In particular, the baseline version uses a step of 1, while it is possible to increase the step size to skip iterations, thus reducing the number of evaluations, thus increasing the performance of the application.

In addition to this first step, based on the analysis done in the previous section, we can exploit domain-knowledge to define more aggressive software-knobs to approximate the *GeoDock-MA* results. In particular, we know from Figure 4a that small fragments have a limited impact on the delta overlap. Therefore, instead of applying a flat loop perforation (as done in the original application), we introduce a *parametric loop perforation* which let us focus on the most important fragments of the ligand. Whenever the size of the current fragment is below a given THRESHOLD, we use a coarse-grain rotation step (LOW-PRECISION STEP). Otherwise, we use a fine-grain rotation step (HIGH-PRECISION STEP).
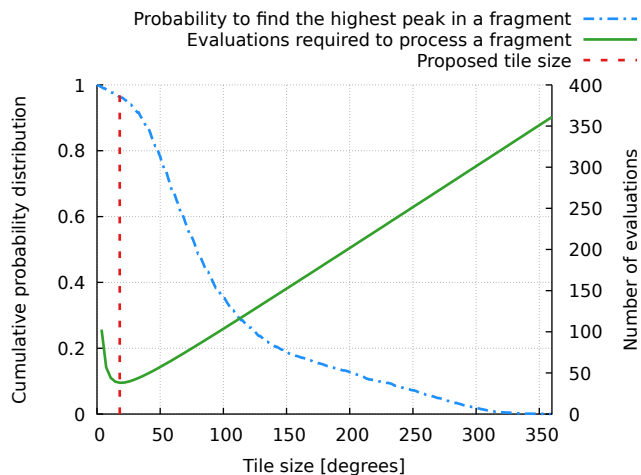
Fig. 6. For each tile size (x-axis), the relation between the number of evaluated rotations (y2-axis) and the probability that the width of the best peak is greater than the given size (y1-axis).

Given that `MatchProbesShape` is a greedy algorithm, we might improve the overlap score by repeating the whole procedure, thus considering multiple time each fragment. In particular, the more we repeat the procedure, the more we increase the probability to find a better pose for the target ligand. Therefore, we define the tunable software knob REP-ETITIONS as the number of times to repeat the procedure. This step seems to be counter-intuitive, however, we argue that it is better to run more times `MatchProbesShape` with aggressive approximations instead of running it only once with fewer approximations.

Furthermore, we can extract other important information about the overlap score from the peak analysis discussed in the previous section. In particular, we can rely on the smoothness of the overlap score through the entire rotation space, which means that each fragment has a limited number of peaks and that the most important peak is usually wide (the median is $68°$). Therefore, we can further reduce the number of evaluations by introducing an *iterative refinement* technique based on the concept of loop tiling and grid-based approaches. For each fragment above THRESHOLD we partition the $360°$ rotation space into several tiles of fixed-size x. Then, we peel and evaluate only one element for each tile (the central one). In the next iteration, we evaluate only the tile corresponding to the most promising peeling element by using HIGH-PRECISION STEP. Given this policy, the number of evaluated rotations (y) is a function of the tile size (x) and HIGH-PRECISION STEP, as described in Eq. 2.

$$y = \frac{360°}{x} + \frac{x}{\text{HIGH-PRECISION STEP}} \qquad (2)$$

We are interested in minimizing the number of pose evaluation while preserving a high probability to identify the most important peak. The minimization of Eq. 2 has a unique solution, its value $\hat{x}$ is defined in Eq. 3.

$$\hat{x} = 6 * \sqrt{10} * \sqrt{\text{HIGH-PRECISION STEP}} \qquad (3)$$

For example, if we set HIGH-PRECISION STEP at the same accuracy of the original algorithm ($1°$), the optimal tile size is $18°$, which means that we have a high probability to

identify the most important peak with the peeling element. In particular, Figure 6 shows, for each tile size (x-axis), the probability that the width of the most important peak is greater than the evaluated tile size (y1-axis, blue line) and the number of evaluated iterations (y2-axis, green line). The red line highlights the optimal value. As a consequence of Eq. 3, we observe that an increment of HIGH-PRECISION STEP implies an increment in the value of the optimal tile size and a decrement of the probability of finding the best peak.

To summarize, starting from the original algorithm described in Algorithm 1, we have introduced five tunable software-knobs (HIGH-PRECISION STEP, LOW-PRECISION STEP, THRESHOLD, REPETITIONS and ENABLE REFINEMENT) to approximate the application by reducing the number of ligand evaluations. The main idea is to spend more time when the computation is more promising.

In particular, Algorithm 2 shows the parametric algorithm of `MatchProbesShape`. All the original algorithm is contained in the outer loop (Line 2) which repeats the pose optimization according to REPETITIONS. The optimization of the pose of left fragment is described between Line 5 and Line 16. According to the relative size of the fragment and to THRESHOLD (Line 5), we run either a coarse grained exploration by using LOW-PRECISION STEP (Line 6) or a fine grained exploration (lines 9-15). In the latter case, we either perform a two-step optimization by using iterative refinements or we perform a flat exploration using HIGH-PRECISION STEP, according to ENABLE REFINEMENT. The two-step optimization evaluates the peeling elements of the rotation (Line 10) and then it refines the exploration of the most promising tile by using HIGH-PRECISION STEP (Line 11). Due to the symmetry of the problem, the same procedure is applied to the right fragment (lines 17-28).

3.4 Application Autotuning

The software knobs defined in the previous section aim at reducing the exploration space of ligand poses, decreasing the time-to-solution of the application and the accuracy of the results as well. However from the end-user point of view, a manual selection of the application configuration it is a non-trivial task. Application autotuning [5] is a well-known field in literature and there are several available tools [16], [17], [18] that can select the most suitable configuration according to application requirements. In this section, we explain how tuning can be automatically done by the application itself according to an execution time budget allocated by the end-user. We used the mARGOt [19] framework to select the configuration of the software knobs that maximizes the accuracy of the result given the time budget.

To select the most suitable configuration, the autotuner should be able to predict the performance of the configurations for the actual input [20]. As the accuracy is platform-independent and it is used to sort the configurations in terms of software knobs, it is possible to run an error profiling campaign only once, averaging the results over a representative set of pockets and ligands.

To complete the execution of the application in the given time budget, we have to estimate the time-to-solution once

Data: the pocket and the 3D structure of the ligand
Result: the overlap score of the ligand

1   get the list of rotamers;
2   **for** *the number of* REPETITIONS **do**
3   **foreach** *rotamer* **do**
4   grow the right and left fragment;
5   **if** *relative size of left fragment* ≤ THRESHOLD **then**
6   place the left fragment in the best angle found with step LOW-PRECISION STEP;
7   **end**
8   **else**
9   **if** ENABLE REFINEMENT **then**
10   evaluate the peeling element for each tile;
11   place the left fragment in the best angle found in the best tile using step HIGH-PRECISION STEP;
12   **end**
13   **else**
14   place the left fragment in the best angle found with step HIGH-PRECISION STEP;
15   **end**
16   **end**
17   **if** *relative size of right fragment* ≤ THRESHOLD **then**
18   place the right fragment in the best angle found with step LOW-PRECISION STEP;
19   **end**
20   **else**
21   **if** ENABLE REFINEMENT **then**
22   evaluate the peeling element for each tile;
23   place the right fragment in the best angle found in the best tile using step HIGH-PRECISION STEP;
24   **end**
25   **else**
26   place the right fragment in the best angle found with step HIGH-PRECISION STEP;
27   **end**
28   **end**
29   **end**
30   **end**
31   **return** *the overlap score of the ligand;*

Algorithm 2: The tunable pseudo-code of the `MatchProbesShape` kernel.

the target architecture and the actual input dataset (pocket and ligands database) are given. Being the target problem embarrassing parallel, without the need of synchronization, the overhead of the MPI environment is negligible even scaling over a large set of nodes. Therefore, assuming homogeneous resources, we predicted the time-to-solution of the serial case, then we re-calculate according to the allocated resources.

Considering only a configuration for the software knobs, it is possible to use input data features to estimate the time-to-solution of the given input. To this end, we modelled the entire database as a set of ligands with the same *average* data features. In particular, we used a multivariate linear regression with interaction to estimate the time-to-solution t_{la} for the average ligand. The vector of predictors $\overline{x}$ is composed of the number of 3D points of the pocket xp_p, the average number of atoms in a ligand xl_a, the average number of rotamers in a ligand xl_r, and all the possible interactions among them (i.e. $xp_p \cdot xl_a$, $xp_p \cdot xl_r$, $xl_a \cdot xl_r$, and $xp_p \cdot xl_a \cdot xl_r$). Thus, the target model is simply composed of $t_{la} = \overline{\alpha} \cdot \overline{x} + \beta$, where $\overline{\alpha}$ is the vector of predictor coefficient, while β is the intercept.

To generalize the approach, we considered the parameters of the regression as a function of the proposed software knobs, given that the impact of the data features on the execution time is strongly dependent on the configuration. By using this information, we can build a model to estimate the time-to-solution according to Eq. 4, where $\overline{k}$ is the vector of software knobs and ν is the number of ligands to dock, in the input database.

$$t = \nu \cdot (\overline{\alpha}(\overline{k}) \cdot \overline{x} + \beta(\overline{k})) \qquad (4)$$

Differently from the accuracy characterization, the performance model should be trained every time the computing platform is changed. However, in both cases, the experiment described in Section 5.1 suggests that a small database of ligands is enough to define the accuracy-performance behaviour.

To recap, we enhanced the original algorithm of the application by exposing software knobs that enable performance-accuracy trade-offs. We used an application autotuner to configure automatically the application according to simple user-oriented parameters: the number of available computational resources and the available time-budget. The data features of the actual input can be either included by the user or directly extracted by a preliminary input data analysis.

4 EXPERIMENTAL SETUP

To assess the benefits of the approximation techniques described in this paper, the experimental setup consists of the data sets used in the experiments, the metrics of interest and the target platform to execute the application.

4.1 Data Sets

To evaluate the functional and extra-functional performance of the proposed approximation techniques, we used a database of ligands composed of 113K ligands. The molecules are different in terms of the number of atoms (between 28 and 153) and rotamers (between 2 and 53). We used 6 pockets protein-ligand complexes derived from the RCSB Protein Data Bank (PDB) [14]: 1b9v, 1c1v, 1cvu, 1c2, 1dh3, 1fm9. In particular, the PASS [13] version of the pockets has also been used together with the database of ligands. Differently from the classical grid representation of the pocket, the PASS (Putative Active Sites with Spheres) version uses spheres to represent binding sites. This solution has been widely used in the context of fast docking [13].

4.2 Metrics of Interest

To measure the performance of *GeoDock-MA*, we considered its throughput and the time-to-solution. In particular, the application throughput is defined as the number of ligand's atoms processed in one second, while the time-to-solution is the time taken by the application to elaborate the input.

We used the metric *overlap degradation* to quantify the mean loss of accuracy introduced by the approximation techniques with respect to the baseline, which is the configuration that leads, on average, to the better overlap score: HIGH-PRECISION STEP = $1°$, THRESHOLD = 0, REPETITIONS = 3 and ENABLE REFINEMENT = $false$. The *overlap degradation* is defined as described in Eq. 5,

$$score_{degradation} = (1 - \frac{overlap_{approx}}{overlap_{original}}) \times 100 \quad (5)$$

where $overlap_{approx}$ is the mean overlap score of the top 1% ligands of the evaluated configuration, while $overlap_{original}$ is the mean overlap score of the top 1% ligands of the non-approximated version of the application (baseline).

4.3 Target Platform

The platform used to execute the experiments is based on dedicated supercomputer NUMA nodes featuring two Intel Xeon E5-2630 V3 CPUs (@2.8 GHz) with 128 GB of DDR4 memory (@1866 MHz) on a dual channel memory configuration. The experiments are executed by using the GALILEO platform located at CINECA supercomputing center[2].

5 EXPERIMENTAL RESULTS

In this section, we evaluated the benefits of the proposed methodology by using four different experiments. Being a data-dependent application, the first experiment aims at assessing data sensitivity by changing the number of ligands used for evaluating a configuration. The second experiment assesses the benefits of applying the approximation techniques with respect to the original version of the application. We show the obtained accuracy-performance tradeoffs for a virtual screening procedure and we evaluate the effects of the overlap degradation on a single ligand docking procedure. The third experiment validates the accuracy of the time-to-solution model. Finally, the fourth experiment wants to prove the benefits of the proposed approach for the end-user on two different application scenarios.

5.1 Data Dependency Evaluation

In this paper, we enhanced the geometrical docking module of LiGenDock with approximation techniques, to trade off the quality of the results with respect to the throughput. Therefore we are interested in finding the set of configurations in the Pareto front, given by the non-dominated solutions considering both target metrics. However, the application needs to find the most promising ligands across a heterogeneous data set. This means that the performance of the application might depend on which ligands are considered.

2. https://www.cineca.it/en

(a) Throughput per process

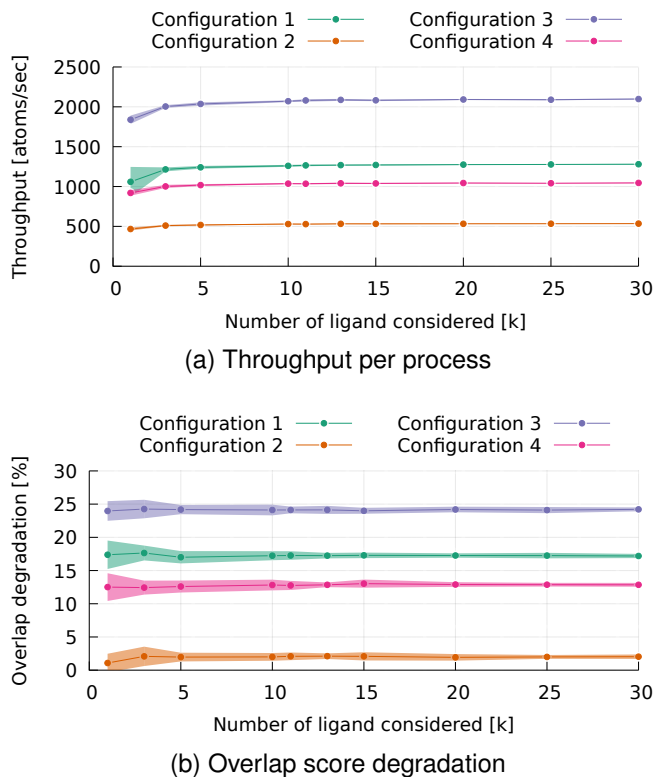

(b) Overlap score degradation

Fig. 7. Application analysis in terms of throughput per process and overlap score degradation by varying the number of ligands. For each configuration we show the average values (dot) and the standard deviation (colored area).

To avoid the profiling phase of the alternative configurations for each database of ligands that we want to evaluate, this experiment wants to assess how much the performance of the application is dependent on the changes in the dataset. To this end, we evaluate four different configurations of the enhanced version of *GeoDock-MA* in terms of tunable knobs. For each configuration, we characterize the application behaviour in terms of throughput and overlap degradation by varying the number of considered ligands. The set of ligands considered to evaluate each configuration has been selected by randomizing 20 times over the full set of 113K elements, thus emulating new datasets.

Figure 7 shows the results of the experiment. In particular, Figure 7a focuses on the application throughput (y-axis), while Figure 7b focuses on the overlap degradation (y-axis). For both of them, each dot represents the mean performance of the evaluated configuration by varying different databases of ligands. The transparent solid area represents the uncertainty of the mean by using the standard deviation of the measures. The x-axis indicates the number of ligands considered in the evaluation.

From these results we see in Figure 7a that the average application throughput has a very limited dependency on both the number of ligands in the target database and the input data (i.e. very small standard deviation). This is an expected result, because the throughput definition we considered is related to the number of atoms of the database instead of the number of ligands, thus providing a normalized measure. On the other hand, Figure 7b shows how the overlap degradation is a little bit more data dependent than
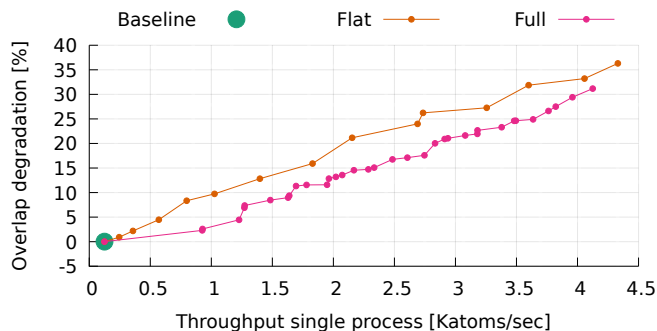
Fig. 8. Pareto front of *GeoDock-MA* in terms of overlap score degradation versus throughput: *Flat* and *Full*
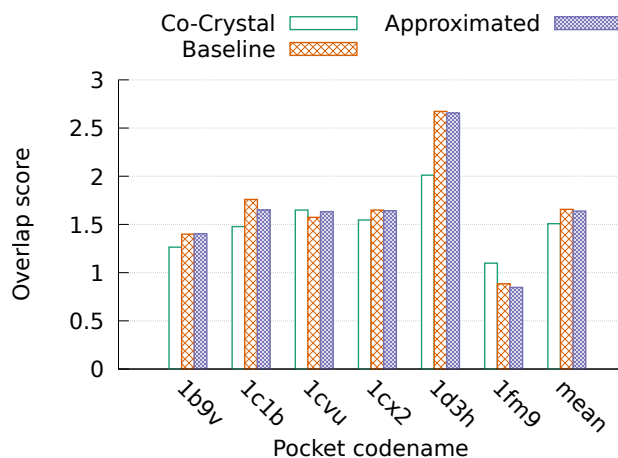

Fig. 9. Overlap score by varying the target pocket for the co-crystallized ligand, for the baseline and proposed approximated version.

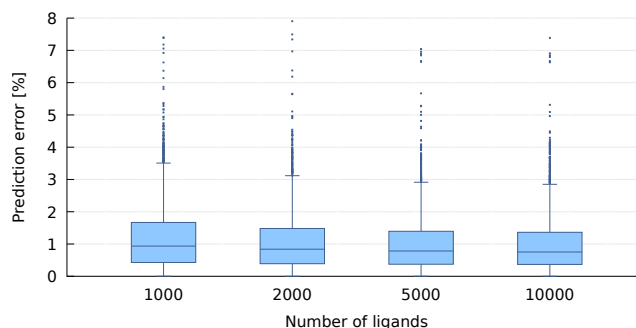Fig. 10. Prediction error of the time-to-solution, by varying the number of ligands to dock in the target pocket.

the throughput. In particular, we need to consider at least $5K$ ligands to have a steady average value. This behaviour is due to the overlap degradation definition that makes the value dependent on the top 1% ligands and therefore on which ligands are selected. However, we can determine that less than 10K ligands are enough to characterize the configurations of the enhanced version of *GeoDock-MA* for both throughput and overlap degradation.

5.2 Trade-off Analysis

This experiment wants to show performance-accuracy tradeoffs when we apply the approximation techniques described in Section 3.3. Figure 8 shows the Pareto front of the design space exploration carried out in a single node of Galileo using 20k ligands, targeting a single pocket.

Figure 8 compares the performance of the application by using a *flat* sampling on the rotation angles, as proposed in the LiGenDock paper [4], with the performance of the application using the *full* set of software knobs proposed in this paper. In particular, the *flat* design space is given by: HIGH-PRECISION STEP $[1°, 2°, 3°, 5°, 10°, 15°, 45°, 60°]$, REPETITIONS $[1, 2, 3]$. The *full* design space (for evaluating all the software knobs proposed in the paper) is given by: HIGH-PRECISION STEP $[1°, 2°, 3°, 5°]$, LOW-PRECISION STEP $[45°, 90°]$, THRESHOLD $[0, 0.3, 0.6, 0.8]$, REPETITIONS $[1, 2, 3]$, ENABLE REFINEMENT $[true, false]$. For both design spaces, we used a full factorial Design of Experiments. We considered as *Baseline* configuration the most precise one, that is exactly the same version that can be derived either from *flat* or *full* – i.e. HIGH-PRECISION STEP=$1°$ and REPETITIONS=3 for the *flat* version, and HIGH-PRECISION STEP=$1°$, LOW-PRECISION STEP=*, THRESHOLD=0, REPETITIONS=3, and EN-ABLE REFINEMENT=$false$ for the *full* version.

As expected, from the results reported in Figure 8, the Pareto front derived by the *full* version dominates the one derived by the *flat* sampling. It is possible to notice how only by enabling the iterative refinement (first configuration on the *full* curve after the *Baseline*), we can greatly improve the throughput of the application (7.4X) with a limited overlap degradation (2.3%) with respect to the *Baseline* version.

The Protein Data Bank (PDB) [14] contains the three-dimensional structural data of biological molecules, providing also the pose of the ligand when co-crystallized within the target pocket (i.e. the actual pose of the best ligand for that pocket). Therefore, we decided to use some

pocket-ligand pairs to better analyze the effects of accuracy degradation. Figure 9 shows for each pocket-ligand pair: a) the overlap score of the crystal, b) the overlap score of the docked ligand using the *baseline* version, and c) the overlap score of the approximated version using only the iterative refinement (i.e. 7.4x speedup). We may notice that the investigated configuration of the enhanced *GeoDock-MA* has a small degradation of the overlap score not only in the average case but also when considering a single target ligand. The overlap score of the co-crystallized ligand is usually lower with respect to the computed ones because the real pose of the ligand takes into consideration also chemical features which are not captured by the geometrical score.

5.3 Time-to-solution Model Validation

This experiment aims at validating the time-to-solution model described in Section 3.4. In particular, the model is defined in the design space of the *full* version described in Section 5.2. To compute the coefficients of the linear regression for each configuration (see Eq. 4), we run the application several times by using $1K$ ligands per pocket for each configuration. The extracted models for each configuration have an average adjusted R^2 value equal to 0.977.

We run an experimental campaign to further validate the accuracy of the time-to-solution prediction of the model, by using the leave-one-out scheme on the pockets and

a different set of ligands with respect to those used for the training. For each pocket and configuration stated in Section 5.2, we executed the application with three different databases randomly selected over the entire set and composed of $1k$, $2k$, $5k$ and $10k$ ligands. We did not use very small sizes for the target database because our goal is to predict the time-to-solution during a virtual-screening process, typically composed of a large number of target ligands. For each run, we stored the predictor value of the input and the observed time-to-solution to extract the prediction error. Figure 10 shows the distribution of the prediction error of our model by varying the number of ligands in the experiment. The average error is below 1% of the observed time-to-solution for the entire range of considered ligand-database size, while the outliers (we validate the application by using more than $15K$ experiments) reach a maximum value of 7.9%, with a trend which is stable by increasing the number of docked ligands.

5.4 Use-case Scenarios

The last experiment wants to evaluate the benefits of the proposed approach for the end-user, i.e. a pharmaceutical company that aims at screening a large set of ligands given a time budget. We envisioned two use cases to exploit the performance-accuracy tradeoff. In the first scenario, we considered a fixed time budget for the computation and we would like to understand what is the effect of incrementing the size of the database to be analyzed, in order to increase the chances to find a better drug. In the second scenario, we fixed the size of the database and we observed the effects of varying the time budget, thus varying the cost of the experiment. We can summarize these two scenarios as attempts to offer to the end-user with two high-level knobs: In the first case the number of ligands to be screened, while in the second case the time budget (i.e. the cost of the experiment). The time-to-solution model will be used to set the right low-level application knobs included in *GeoDock-MA* while satisfying the constraints.

Figure 11 shows the expected behaviour of the application, assuming that the end-user is using 8 nodes of the Galileo supercomputer (see Section 4.3). On the y-axis are represented the expected performance of the application (top 2 plots) and the selected configuration of the software knobs (bottom 5 plots). The performance of the application is defined in terms of the expected completion percentage of the planned ligand database and the related overlap degradation of the result. The completion percentage is the ratio between the number of ligands docked given the time budget and the size of the target ligand database. Each plot includes two lines: one dotted line for the baseline version of *GeoDock-MA* and the other line for the adaptive version proposed in this paper.

The x-axis represents the high-level parameter tuned by the end-user, according to the scenario. In Scenario 1 (Figure 11a), the end-user would like to tune the size of the database of ligands to complete the job given one-day time budget. In Scenario 2 (Figure 11b), the end-user would like to tune the time budget of the application to dock a given database of 500×10^6 ligands.

In both scenarios, we notice that the proposed software knobs enable a swing of several orders of magnitude for the end-user to tune the parameters of the job (i.e. problem size and time-to-solution). Moreover, by using an autotuner together with the time-to-solution model, we are capable to alleviate the burden of the manual selection of the software knobs from the end-user, by exposing simpler parameters. Although the average trend of the application knobs values for both scenarios can be derived from their meaning, the actual values and when to switch among the configurations according to the high-level constraints (i.e. problem size and time-to-solution) is hard to know without any automatic support. A clear example of this is the behaviour of the THRESHOLD-value within the middle range (300–700×10^6) of the problem size considered in the experiment shown by Figure 11a.

In terms of application performance, we can notice how in the first scenario, where we kept constant the time to solution (see Figure 11a), the proposed *adaptive* version is able to process the whole database up to 850×10^6 ligands, the baseline version rapidly decreases the rate to very low values when enlarging the database size. This happens thanks to the capability of the proposed *adaptive* version, not present in the *baseline*, to trade-off accuracy and performance. Moreover, the second plot in Figure 11a also demonstrates how the deep parametrization introduced in the target Mini-App provides a smooth degradation of the application quality. In the second scenario shown in Figure 11b, where we fixed the size of the target database and we varied the allocated time budget, we observed a similar behaviour. In fact, increasing the time budget, the proposed *adaptive* version rapidly reaches (< 1 day) the value of 100% completion at the cost of a low accuracy, while for the baseline we had to allocate more than 20 days.

Finally, in both scenarios we can identify when the ligand database size, or the allocated budget, is set to values that guarantee only 10% of the expected completion for the *baseline* version, the *adaptive* one results with a limited degradation less than 10%. This demonstrates the effectiveness of the extracted low-level knobs in *GeoDock-MA*.

6 RELATED WORK

Being our approach intrinsically interdisciplinary, we partitioned the relevant literature in three sections covering the research areas related to this work: molecular docking, approximation techniques and application autotuning.

Molecular docking techniques. Molecular docking is a well-known research topic addressed in literature from different perspectives. A large set of works addresses the problem by exploiting random-based algorithms, such as genetic algorithms [21], [22] or Monte Carlo simulations [11], [23]. However, a desirable feature of a molecular docking application is the determinism of the solution. Because of the tasks following the *in-silico* step require expensive solution-dependent resources, for several companies having a deterministic and repeatable result is a constraint. Early works in literature, such as [24], generate deterministic solutions. However, they consider only rigid movements of the ligand during the docking procedure. Real case scenarios usually require the rotation of portions of the ligand molecule. Therefore, the limitation of rigid movements is likely to prevent the applicability of the solution in the
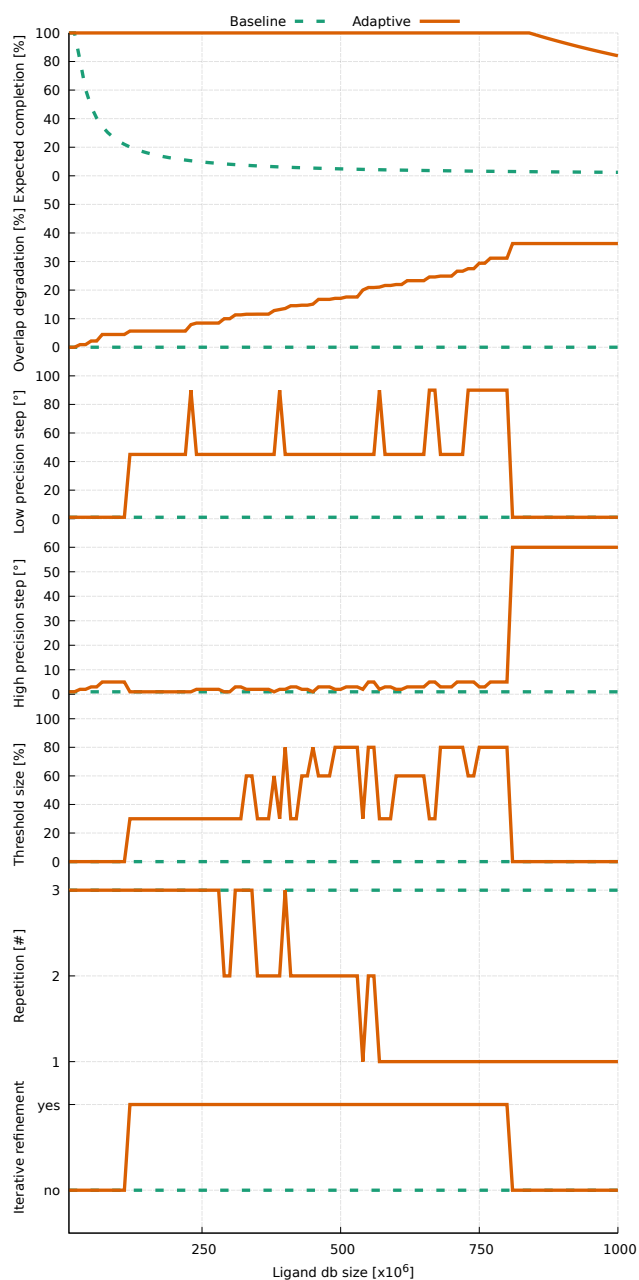
(a) Scenario 1: Varying the size of the ligand database given one-day time budget.
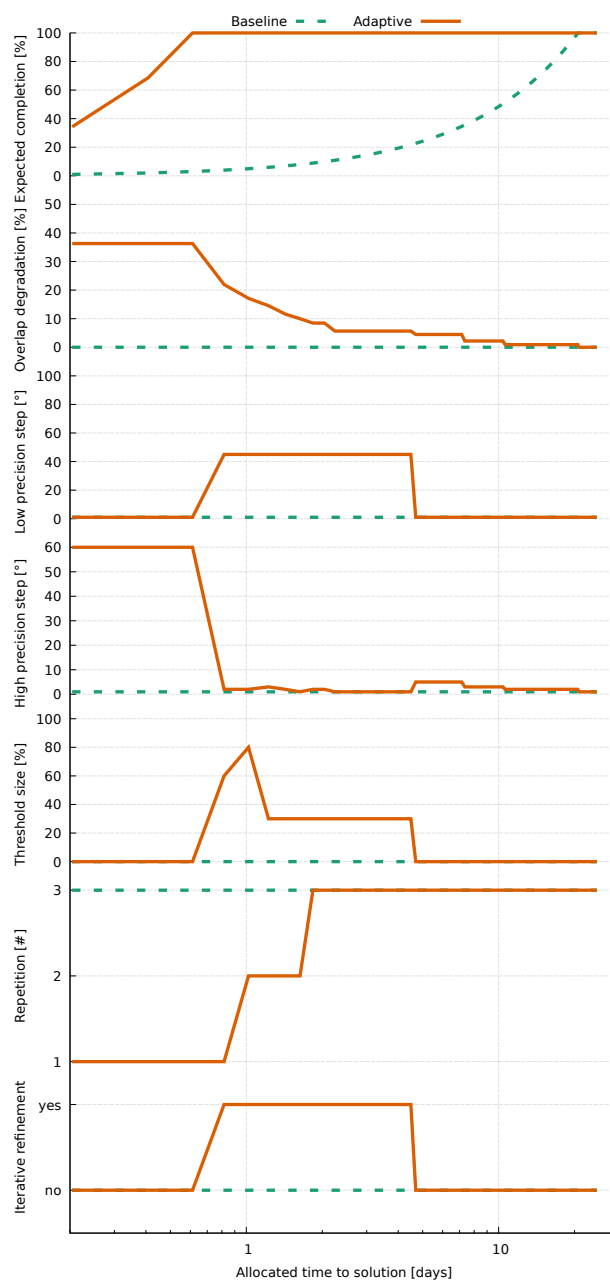
(b) Scenario 2: Varying the allocated time budget given the ligand database size equal to 500×10^6.

Fig. 11. *GeoDock-MA* behavior in terms of expected percentage of ligand database completion, expected overlap degradation, and the selected configuration (a) by varying the size of the input, and (b) the time budget, when using 8 nodes of Galileo supercomputer.

industry. The work of Palma et al. [25] overcomes this issue: they introduced a molecular docking framework to deal with the flexibility of the ligand molecule. It adopts a model of the electrostatic interactions to finalize the docking. Similar works such as DOCK [26], FlexX [27], FlexX-Scan [28] and sur-flex [29] provide deterministic molecular docking of flexible ligands. They enabled the user to exploit several docking algorithms according to the specific use case. All these algorithms rely on both geometric and pharmacophoric properties in their docking algorithms. All these works implement a different docking procedure with respect to LiGenDock. However, no one of them has been designed to expose software knobs to enable explicitly quality-performance tradeoffs. The proposed approach behind *GeoDock-MA* unveil the possibility to tune the docking procedure according to high-level constraints, such as the allocated time budget for a given size of the ligand database to be virtually screened.

Approximation techniques. During the last years, the most promising approaches to achieve performance gains are represented by the trade-off with application accuracy [30], [31]. Algorithm-level approximate computing techniques are well-known in literature [32], [33], [3], and represent also an important challenge for in HPC applica-

tions [34], [35]. In this work, we exploit grid-based optimizations on the geometrical docking kernel. In particular, in computational physics, it is very common to exploit models based on multi-level grids to achieve a fine-grained solution in a restricted area of the whole simulated environment. The size of the grid is a parameter which enables the physicists to trade-off granularity of solution for the increasing/decreasing number of elements to be processed. Nested grids have been exploited for a long time in geophysics applications, such as thermosphere models [36], [37], [38] and ocean flows models [39]. The evolution of nested-grids models made the researchers to abandon regular grids in favour of variable-size grids. Irregular grids have been exploited in climate forecast models to improve the performance of grid-based models. Authors of [40] demonstrate that a variable-resolution stretched grids lead to longer-term climate forecast with the same accuracy of the nested grid models. In physics, variable-size grids are used to discretize geophysical problems such as advection equations [41]. In the field of image rendering, grid processing has been optimized by selecting which tiles need to be processed first and which one later or do not require any processing at all. An element is peeled from each tile, and its value is used to decide how to compute the corresponding tile, e.g. deep peeling [42], [43], [44].

Application autotuning. Autotuning in HPC is a well-known research area [5] where several libraries and tools exist to support the automatic optimization of software knobs. Among the libraries that are tailored to specific tasks some examples are ATLAS [45] for matrix multiplication routines, FTTW [46] for FFT operations, OSKI [47] for sparse matrix kernels, and SPIRAL [48] for digital signal processing. Frameworks such as [49], [16], [18], [17] consider the autotuning problem from a more general perspective. They have been designed to make an application tunable without being bound to a target application field. However, only a few approaches exist to deal with approximation knobs [50], [51], end even fewer frameworks consider data-aware approximation [20], [52], [53]. The work in [20] enhances Petabricks [49] to leverage the accuracy-throughput tradeoffs and to consider input features of the data to be processed, On the other side, [52] and [53] follows the same idea of switching at runtime the configuration according to the input characteristics by applying it to two different application domains, respectively graphics and graph processing. To support the autotuning phase, machine learning techniques have been used in literature [51], [54], [55] to model application metrics and then to predict the best configuration to be applied. This is fundamental when the size of the configuration space is huge, and thus not possible to be entirely profiled, or when the elaboration is heavily data-dependent. The workload we are considering in this paper falls in this second case and these previous works have been used as inspiration to our predictive solution.

7 CONCLUSION

Virtual screening is a crucial task of a drug discovery process as it aims at finding the most promising molecules to evaluate on the later stage of the process. Due to the large number of theoretical molecules that may be evaluated,

any speedup of the task leads to either a reduction of the economic cost of the whole process or the possibility to evaluate a higher number of possible candidates. In this paper, we have analyzed *GeoDock-MA*, which represents the most computational heavy kernels of LiGenDock when it is used to perform a screening on a target database of ligands. From the analysis of *GeoDock-MA*, we identified several approximation possibilities, declined in five software knobs, enabling the computation to gain accuracy only when it is likely to have more impact on the output, thus enabling accuracy-performance tradeoffs. The adaptive version of *GeoDock-MA* provides different accuracy levels according to the evolving needs of the virtual screening experimental campaign. In particular, experimental results demonstrated how, by playing with the introduced software knobs, the proposed tunable application adaptively runs the virtual screening campaign over a ligand database by introducing an accuracy degradation of less than 10% during the same time used by the original version to screen only $1/10$ of the database. Due to the determinism of the proposed algorithm and the software knobs, we defined a prediction model to support the application autotuning. The model is capable to autotune low-level knobs of the application by exposing to the end-user the high-level parameters (such as the time budget, the size of the ligand database and the number of resources). These outcomes represent a huge advantage for pharmaceutical industries in a context where the use of HPC systems and software in drug discovery have become crucial assets to become more and more competitive in the drug discovery process.

REFERENCES

[1] S. Venkataramani, S. T. Chakradhar, K. Roy, and A. Raghunathan. Approximate computing and the quest for computing efficiency. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pages 1–6, June 2015.

[2] Vinay K. Chippa, Srimat T. Chakradhar, Kaushik Roy, and Anand Raghunathan. Analysis and characterization of inherent application resilience for approximate computing. In *Proceedings of the 50th Annual Design Automation Conference*, DAC '13, pages 113:1–113:9, New York, NY, USA, 2013. ACM.

[3] Sparsh Mittal. A survey of techniques for approximate computing. *ACM Comput. Surv.*, 48(4):62:1–62:33, March 2016.

[4] Claudia Beato, Andrea R Beccari, Carlo Cavazzoni, Simone Lorenzi, and Gabriele Costantino. Use of experimental design to optimize docking performance: The case of ligendock, the docking module of ligen, a new de novo design program. *Journal of Chemical Information and Modeling*, 53(6):1503–1517, 2013.

[5] P. Balaprakash, J. Dongarra, T. Gamblin, M. Hall, J. K. Hollingsworth, B. Norris, and R. Vuduc. Autotuning in high-performance computing applications. *Proceedings of the IEEE*, 106(11):2068–2083, Nov 2018.

[6] Evanthia Lionta, George Spyrou, Demetrios K. Vassilatis, and Zoe Cournia. Structure-based virtual screening for drug discovery: Principles, applications and recent advances. *Current Topics in Medicinal Chemistry*, 14(16):1923.

[7] Andrea R. Beccari, Marica Gemei, Matteo Lo Monte, Nazareno Menegatti, Marco Fanton, Alessandro Pedretti, Silvia Bovolenta, Cinzia Nucci, Angela Molteni, Andrea Rossignoli, Laura Brandolini, Alessandro Taddei, Lorena Za, Chiara Liberati, and Giulio Vistoli. Novel selective, potent naphthyl trpm8 antagonists identified through a combined ligand- and structure-based virtual screening approach. In *Scientific reports*, 2017.

[8] Douglas B. Kitchen, Hélène Decornez, John R. Furr, and Jürgen Bajorath. Docking and scoring in virtual screening for drug discovery: methods and applications. *Nature Reviews Drug Discovery*, 3, Nov 2004. Review Article.

[9] Paul D Lyne. Structure-based virtual screening: an overview. *Drug Discovery Today*, 7(20):1047 – 1055, 2002.

[10] Jayashree Srinivasan, Angelo Castellino, Erin K. Bradley, John E. Eksterowicz, Peter D. J. Grootenhuis, Santosh Putta, and Robert V. Stanton. Evaluation of a novel shape-based computational filter for lead evolution: application to thrombin inhibitors. *Journal of Medicinal Chemistry*, 45(12):2494–2500, 2002. PMID: 12036357.

[11] Richard A. Friesner, Jay L. Banks, Robert B. Murphy, Thomas A. Halgren, Jasna J. Klicic, Daniel T. Mainz, Matthew P. Repasky, Eric H. Knoll, Mee Shelley, Jason K. Perry, David E. Shaw, Perry Francis, and Peter S. Shenkin. Glide: a new approach for rapid, accurate docking and scoring. 1. method and assessment of docking accuracy. *Journal of Medicinal Chemistry*, 47(7):1739–1749, 2004. PMID: 15027865.

[12] Andrea R Beccari, Carlo Cavazzoni, Claudia Beato, and Gabriele Costantino. Ligen: a high performance workflow for chemistry driven de novo design. *Journal of Chemical Information and Modeling*, 53(6):1518–1527, 2013.

[13] G Patrick Brady and Pieter FW Stouten. Fast prediction and visualization of protein binding pockets with pass. *Journal of computer-aided molecular design*, 14(4):383–401, 2000.

[14] Helen M. Berman, John Westbrook, Zukang Feng, Gary Gilliland, T. N. Bhat, Helge Weissig, Ilya N. Shindyalov, and Philip E. Bourne. The protein data bank. *Nucleic Acids Res*, 28:235–242, 2000.

[15] Stelios Sidiroglou-Douskos, Sasa Misailovic, Henry Hoffmann, and Martin Rinard. Managing performance vs. accuracy trade-offs with loop perforation. In *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering*, ESEC/FSE '11, pages 124–134, New York, NY, USA, 2011. ACM.

[16] Albert Hartono, Boyana Norris, and Ponnuswamy Sadayappan. Annotation-based empirical performance tuning using orio. In *IPDPS*, pages 1–11. IEEE, 2009.

[17] Ananta Tiwari, Jeffrey K Hollingsworth, Chun Chen, Mary Hall, Chunhua Liao, Daniel J Quinlan, and Jacqueline Chame. Auto-tuning full applications: A case study. *Int. J. High Perform. Comput. Appl.*, 25(3):286–294, August 2011.

[18] R. Mijakovi, M. Firbach, and M. Gerndt. An architecture for flexible auto-tuning: The periscope tuning framework 2.0. In *2016 2nd International Conference on Green High Performance Computing (ICGHPC)*, pages 1–9, Feb 2016.

[19] D. Gadioli, E. Vitali, G. Palermo, and C. Silvano. margot: a dynamic autotuning framework for self-aware approximate computing. *IEEE Transactions on Computers*, pages 1–1, 2018.

[20] Yufei Ding, Jason Ansel, Kalyan Veeramachaneni, Xipeng Shen, Una-May OReilly, and Saman Amarasinghe. Autotuning algorithmic choice for input sensitivity. In *ACM SIGPLAN Notices*, volume 50, pages 379–390. ACM, 2015.

[21] René Thomsen and Mikael H Christensen. Moldock: a new technique for high-accuracy molecular docking. *Journal of medicinal chemistry*, 49(11):3315–3321, 2006.

[22] Gareth Jones, Peter Willett, Robert C Glen, Andrew R Leach, and Robin Taylor. Development and validation of a genetic algorithm for flexible docking. *Journal of molecular biology*, 267(3):727–748, 1997.

[23] Ming Liu and Shaomeng Wang. Mcdock: a monte carlo simulation approach to the molecular docking problem. *Journal of computer-aided molecular design*, 13(5):435–451, 1999.

[24] Fan Jiang and Sung-Hou Kim. soft docking: matching of molecular surface cubes. *Journal of molecular biology*, 219(1):79–102, 1991.

[25] P Nuno Palma, Ludwig Krippahl, John E Wampler, and José JG Moura. Bigger: a new (soft) docking algorithm for predicting protein interactions. *Proteins: Structure, Function, and Bioinformatics*, 39(4):372–384, 2000.

[26] Todd JA Ewing, Shingo Makino, A Geoffrey Skillman, and Irwin D Kuntz. Dock 4.0: search strategies for automated molecular docking of flexible molecule databases. *Journal of computer-aided molecular design*, 15(5):411–428, 2001.

[27] Bernd Kramer, Matthias Rarey, and Thomas Lengauer. Evaluation of the flexx incremental construction algorithm for proteinligand docking. *Proteins: Structure, Function, and Bioinformatics*, 37(2):228–241, 1999.

[28] Ingo Schellhammer and Matthias Rarey. Flexx-scan: Fast, structure-based virtual screening. *PROTEINS: Structure, Function, and Bioinformatics*, 57(3):504–517, 2004.

[29] Ajay N Jain. Surflex-dock 2.1: robust performance from ligand energetic modeling, ring flexibility, and knowledge-based search. *Journal of computer-aided molecular design*, 21(5):281–306, 2007.

[30] A. C. I. Malossi, M. Schaffner, A. Molnos, L. Gammaitoni, G. Tagliavini, A. Emerson, A. Toms, D. S. Nikolopoulos, E. Flamand, and N. Wehn. The transprecision computing paradigm: Concept, design, and applications. In *2018 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1105–1110, March 2018.

[31] Hadi Esmaeilzadeh, Adrian Sampson, Luis Ceze, and Doug Burger. Architecture support for disciplined approximate programming. In *Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS XVII, pages 301–312, New York, NY, USA, 2012. ACM.

[32] J. Han and M. Orshansky. Approximate computing: An emerging paradigm for energy-efficient design. In *2013 18th IEEE European Test Symposium (ETS)*, pages 1–6, May 2013.

[33] Q. Xu, T. Mytkowicz, and N. S. Kim. Approximate computing: A survey. *IEEE Design Test*, 33(1):8–22, Feb 2016.

[34] Yandong Wang, Robin Goldstone, Weikuan Yu, and Teng Wang. Characterization and optimization of memory-resident mapreduce on hpc systems. In *2014 IEEE 28th International Parallel and Distributed Processing Symposium*, pages 799–808, May 2014.

[35] Vassilis Vassiliadis, Charalampos Chalios, Konstantinos Parasyris, Christos D. Antonopoulos, Spyros Lalis, Nikolaos Bellas, Hans Vandierendonck, and Dimitrios S. Nikolopoulos. Exploiting significance of computations for energy-constrained approximate computing. *International Journal of Parallel Programming*, 44:1078–1098, 2016.

[36] TJ Fuller-Rowell. A two-dimensional, high-resolution, nested-grid model of the thermosphere: 1. neutral response to an electric field spike. *Journal of Geophysical Research: Space Physics*, 89(A5):2971–2990, 1984.

[37] TJ Fuller-Rowell. A two-dimensional, high-resolution, nested-grid model of the thermosphere: 2. response of the thermosphere to narrow and broad electrodynamic features. *Journal of Geophysical Research: Space Physics*, 90(A7):6567–6586, 1985.

[38] Wenbin Wang, Tim L Killeen, Alan G Burns, and Raymond G Roble. A high-resolution, three-dimensional, time dependent, nested grid model of the coupled thermosphere–ionosphere. *Journal of Atmospheric and Solar-Terrestrial Physics*, 61(5):385–397, 1999.

[39] Lie-Yauw Oey and Ping Chen. A nested-grid ocean model: With application to the simulation of meanders and eddies in the norwegian coastal current. *Journal of Geophysical Research: Oceans*, 97(C12):20063–20086, 1992.

[40] Michael S. Fox-Rabinovitz, Lawrence L. Takacs, Ravi C. Govindaraju, and Max J. Suarez. A variable-resolution stretched-grid general circulation model: Regional climate simulation. *Monthly Weather Review*, 129(3):453–469, 2001.

[41] Paul A Ullrich and Christiane Jablonowski. An analysis of 1d finite-volume methods for geophysical problems on refined grids. *Journal of Computational Physics*, 230(3):706–725, 2011.

[42] Fbio F. Bernardon, Christian A. Pagot, Joo L. D. Comba, and Cludio T. Silva. Gpu-based tiled ray casting using depth peeling. *Journal of Graphics Tools*, 11(4):1–16, 2006.

[43] Cass Everitt. Interactive order-independent transparency. *White paper, nVIDIA*, 2(6):7, 2001.

[44] Cheng-Kai Chen, Chris Ho, Carlos Correa, Kwan-Liu Ma, and Ahmed Elgamal. Visualizing 3d earthquake simulation data. *Computing in Science & Engineering*, 13(6):52–63, 2011.

[45] R Clint Whaley and Jack J Dongarra. Automatically tuned linear algebra software. In *Proceedings of the 1998 ACM/IEEE conference on Supercomputing*, pages 1–27. IEEE Computer Society, 1998.

[46] Matteo Frigo and Steven G Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005.

[47] Richard Vuduc, James W Demmel, and Katherine A Yelick. OSKI: A library of automatically tuned sparse matrix kernels. In *Journal of Physics: Conference Series*, volume 16, page 521. IOP Publishing, 2005.

[48] Markus Püschel, José MF Moura, Bryan Singer, Jianxin Xiong, Jeremy Johnson, David Padua, Manuela Veloso, and Robert W Johnson. Spiral: A generator for platform-adapted libraries of signal processing alogorithms. *The International Journal of High Performance Computing Applications*, 18(1):21–45, 2004.

[49] Jason Ansel, Yee Lok Wong, Cy Chan, Marek Olszewski, Alan Edelman, and Saman Amarasinghe. Language and compiler support for auto-tuning variable-accuracy algorithms. In *Proceed-*

ings of the 9th Annual IEEE/ACM International Symposium on Code Generation and Optimization, pages 85–96. IEEE Computer Society, 2011.

[50] Henry Hoffmann, Stelios Sidiroglou, Michael Carbin, Sasa Misailovic, Anant Agarwal, and Martin Rinard. Dynamic knobs for responsive power-aware computing. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS XVI, pages 199–212, New York, NY, USA, 2011. ACM.

[51] Xin Sui, Andrew Lenharth, Donald S Fussell, and Keshav Pingali. Proactive control of approximate programs. *ACM SIGOPS Operating Systems Review*, 50(2):607–621, 2016.

[52] Mehrzad Samadi, Janghaeng Lee, D. Anoushe Jamshidi, Amir Hormati, and Scott Mahlke. Sage: Self-tuning approximation for graphics engines. In *Proceedings of the 46th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO-46, pages 13–24, New York, NY, USA, 2013. ACM.

[53] H. Omar, M. Ahmad, and O. Khan. Graphtuner: An input dependence aware loop perforation scheme for efficient execution of approximated graph algorithms. In *2017 IEEE International Conference on Computer Design (ICCD)*, pages 201–208, Nov 2017.

[54] J. Bergstra, N. Pinto, and D. Cox. Machine learning for predictive auto-tuning with boosted regression trees. In *2012 Innovative Parallel Computing (InPar)*, pages 1–9, May 2012.

[55] Elizabeth R. Jessup, Pate Motter, Boyana Norris, and Kanika Sood. Performance-based numerical solver selection in the lighthouse framework. *SIAM J. Scientific Computing*, 38, 2016.