Les Cahiers du GERAD ISSN: 0711–2440

Parallel versions of the mesh adaptive direct search algorithm

S. Le Digabel, A. Lesage-Landry, S. Mendoza, C. Tribes

G–2026–35

July 2026

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

Citation suggérée : S. Le Digabel, A. Lesage-Landry, S. Mendoza, C. Tribes (Juillet 2026). Parallel versions of the mesh adaptive direct search algorithm, Rapport technique, Les Cahiers du GERAD G–2026–35, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2026-35) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Suggested citation: S. Le Digabel, A. Lesage-Landry, S. Mendoza, C. Tribes (July 2026). Parallel versions of the mesh adaptive direct search algorithm, Technical report, Les Cahiers du GERAD G–2026–35, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2026-35) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2026 – Bibliothèque et Archives Canada, 2026

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2026 – Library and Archives Canada, 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Parallel versions of the mesh adaptive direct search algorithm

Sébastien Le Digabel [a, c]

Antoine Lesage-Landry [b, c]

Samuel Mendoza [b, c]

Christophe Tribes [a, c]

[a] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

[b] *Department of Electrical Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

[c] *GERAD, Montréal (Qc), Canada, H3T 1J4*

sebastien.le-digabel@polymtl.ca
antoine.lesage-landry@polymtl.ca
samuel.mendoza@polymtl.ca
christophe.tribes@polymtl.ca

July 2026
Les Cahiers du GERAD
G–2026–35

Copyright © 2026 Le Digabel, Lesage-Landry, Mendoza, Tribes

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : This work surveys the different parallel variants of the mesh adaptive direct search (MADS) algorithm for constrained blackbox optimization. These problems can inherently imply high computational costs due to the possible large number of variables and multi-modality of the search space. In addition, the potential time-intensive nature and time heterogeneity of the blackboxes defining the problem prompts the need for efficient implementations. Parallelism emerges as an actionable solution to mitigate computation time, as modern computer systems rely on multi-core architecture. The reviewed methods employ diverse levels of parallelism and distinct parallel strategies to effectively tackle each aspect outlined above. The manuscript details the practical implementations, provides computational results, and offers insights into the advantages and limitations of each MADS parallel method.

Keywords : Blackbox optimization; derivative-free optimization; mesh adaptive direct search; parallel computing

Résumé : Ce travail présente une revue des différentes variantes parallèles de l'algorithme de recherche directe sur treillis adaptatifs (MADS) pour l'optimisation de boîtes noires avec contraintes. Ces problèmes peuvent impliquer des coûts de calcul élevés en raison du nombre potentiellement grand de variables et de la multi-modalité de l'espace de recherche. De plus, le temps d'évaluation possiblement important et l'hétérogénéité temporelle des boîtes noires définissant le problème motivent le besoin d'implémentations efficaces. Le parallélisme s'impose comme une solution concrète pour réduire les temps de calcul, les systèmes informatiques modernes reposant sur des architectures multicoeurs. Les méthodes recensées exploitent divers niveaux de parallélisme et des stratégies parallèles distinctes afin de traiter efficacement chacun des aspects mentionnés ci-dessus. Le manuscrit détaille les implémentations pratiques, fournit des résultats numériques et offre un éclairage sur les avantages et les limites de chaque méthode parallèle de MADS.

Mots clés : Optimisation de boîte noire; optimisation sans dérivées; recherche directe sur treillis adaptatifs; calcul parallèle

Acknowledgements: This work is supported by the NSERC Alliance-Mitacs Accelerate grant ALLRP 571311–21 ("Optimization of future energy systems") in collaboration with Hydro-Québec.

1 Introduction

This work considers optimization problems of the form

$$\min_{\mathbf{x}\in\Omega} f(\mathbf{x}), \tag{$\mathbb{P}$}$$

where $\Omega = \{\mathbf{x} \in \mathcal{X} \mid c_j(\mathbf{x}) \leq 0, j \in \mathcal{J} = \{1, 2, \ldots, m\}\} \subseteq \mathbb{R}^n$, $m, n \in \mathbb{N}$, is the feasible set and $\mathcal{X}$ is a subset of $\mathbb{R}^n$, typically defined by bound constraints. The functions $f : \mathcal{X} \to \mathbb{R} \cup \{\infty\}$ and $c_j : \mathcal{X} \to \mathbb{R} \cup \{\infty\}$, $j \in \mathcal{J}$, are typically evaluated through a computer simulation with no available derivatives, and is considered as a blackbox. We consider that evaluating such a blackbox is time-consuming and that only a limited budget of evaluations is available. We further assume that the blackbox may be subject to heterogeneity in the time required to evaluate a point, as illustrated in Figure 1.

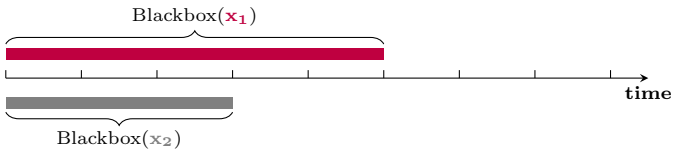

Figure 1: Illustration of a time-heterogeneous blackbox.

Problem ($\mathbb{P}$) is a blackbox optimization (BBO) problem, for which derivative-free optimization (DFO) algorithms are considered. These techniques are described in [10, 19], and some practical applications are illustrated in [2].

With DFO methods, a rule of thumb is to allow a budget of the order of $1{,}000n$ evaluations to obtain satisfactory results. This is not conceivable when blackboxes are time-consuming because this process would lead to impractical resolution times. The situation worsens when dealing with large-scale problems, typically with more than 50 variables, as the evaluations budget escalates significantly.

Parallel computing appears as a solution because most of computational resources exploits multi-core architectures available through various levels, e.g., workstations, servers, and supercomputers. In addition, with the end of Dennard scaling, parallelization has become much more important than before, as the alternative of increasing processor clock speed is no longer possible due to heat dissipation and power consumption issues.

An obvious way of using parallelism is to "open" the blackbox and to modify it to exploit the available parallel resources. However, this is not always possible due to many reasons: The source code of the blackbox may be unavailable, too complicated, too old, etc. Hence, in this work, we focus on the parallelization of DFO algorithms, and in particular direct search (DS) methods.

In an effort to take advantage from modern multi-core infrastructures, multiple DFO methods have been extended to leverage parallel computations through their resolution process. In particular, DFO methods can benefit from parallelism to reduce the time resolution, improve solution quality, and increase the size of BBO problems that can be practically solved. Specifically, DS methods, a branch of DFO methods, provides a framework for the effective use of parallel computing.

Throughout this work, we consider the mesh adaptive direct search (MADS) [7] algorithm and its C++ implementation, `NOMAD` version 3 [44], and we review its different parallel versions. We measure the performance of these parallel algorithms by introducing metrics for benchmarking in a context of parallel computing, and we give recommendations to guide users on what kind of parallelism should be used depending on the different aspects of the problem at hand. This study is conducted with MADS, but remains applicable to other DS methods.

This work is structured as follows: Section 2 reviews the literature and provides an overview of parallel methods for BBO, Section 3 describes MADS, Section 4 presents its parallel versions, and

Section 5 reports computational experiments with recommendations. Finally, concluding remarks are provided in Section 6.

2 Parallelism in BBO

A guiding principle for the efficient use of parallelization to solve (ℙ) is to consider that the time spent to evaluate the objective or constraint functions is much higher than the time required by an optimization algorithm to generate trial points. Hence, the parallelization that is the most likely to be beneficial, i.e., which yields a large reduction in overall execution time, consists of performing several concurrent evaluations of the blackbox and/or to perform an internal parallelization of the blackbox [49]. As mentioned above, we consider that the latter is not available in this work. Parallelization strategies, i.e., functional parallelism, domain decomposition, multi-search parallelism, and the hybridization of mutually compatible aforementioned strategies, are described in details in [55]. It also provides a literature review on integrative framework for parallel computational optimization in operations research. While these strategies are formulated in an operations research context, the scope of applicability can be straightforwardly extended to other optimization algorithm classes, such as DS methods, as described in [21].

Parallelism in BBO and specifically for DS methods, is a longstanding topic of interest. The first DS methods to employ parallelism with proofs of convergence are based on the generalized pattern search (GPS) [61] and the generating set search (GSS) [41] algorithms. The resulting method is the asynchronous parallel pattern search [26, 27, 35, 40, 42]. Implementation examples can be found in the `APPSPACK` [24] and `HOPSPACK` [52] solvers. Reference [45] combines instances of MADS with different line search methods which are run in parallel to accelerate the resolution process. A specific parallel search for MADS is later introduced in the technical report [58]. Other DS methods include the Nelder-Mead [48] whose first parallel version is described in [20], and revisited more recently in [50]. The DIRECT algorithm [37] is hybridized with GSS in [25] and its parallelized versions are analyzed in [29–32]. Other DFO methods include model-based algorithms utilizing a trust region and a model such as polynomial or radial basis functions. In this case, parallelization is not as inherent as for DS methods. Nonetheless, examples of parallel DFO can be found in [14, 15, 36, 53, 56, 64]. Parallel surrogate-assisted methods are also described in [16, 23, 28]. Bayesian optimization methods based on the efficient global optimization framework [38] are parallelized, for example, in [62, 65, 66]. There exists also multiple other parallel heuristics such as [63]. More generally, the survey [3] reviews different kinds of parallel heuristics.

Finally, parallel DS methods are employed in specific applications such as hyperparameters tuning [6], local optimum enumeration [43], multiobjective optimization [60], and small (or moderate) dimensional engineering design problems driven by computationally expensive simulations, where large-batch parallel sampling can outperform more elaborate techniques [17, 51].

3 MADS: Mesh Adaptive Direct Search

MADS is a direct search method introduced in [7] and refined in several articles, including its last major evolution in [12]. The version given in Algorithm 1 corresponds to the one provided in [10].

Blackbox evaluations occur at two different instances: The search and the poll. The points to be evaluated are gathered in two sets $\mathcal{S}^k \subset \mathbb{R}^n$ and $\mathcal{P}^k \subset \mathbb{R}^n$ and are evaluated opportunistically, i.e., evaluations are interrupted as soon as a new best point is found. When evaluations are computationally costly, as in BBO, this strategy is effective because it can reduce the number of blackbox evaluations which is usually subject to a budget constraint. In a sequential programming environment, opportunism occurs at an individual point level. Thus, a minimum of one evaluation can be computed in order to terminate the current iteration if a simple decrease condition is reached by an evaluation

Algorithm 1: MADS **with Extreme Barrier**

Initialization
  Set $\Delta^0 > 0$ and $\mathbf{x}^0 \in \Omega$.
  Start the w worker processes for evaluation.
  Evaluate $\{f(\mathbf{x}^0), c_1(\mathbf{x}^0), c_2(\mathbf{x}^0), \ldots, c_m(\mathbf{x}^0)\}$ and initialize cache $\mathcal{C}$.
for $k = 0, 1, \ldots$ **do**
  Mesh (construction of $\mathcal{M}^k$ (1))
    Set $\mathbf{D} \in \mathbb{R}^{n \times p}$ a positive spanning matrix.
    Set $\delta^k = \min\{\Delta^k, (\Delta^k)^2\}$.
  Search (optional)
    Build the search set $\mathcal{S}^k \subset \mathcal{M}^k$.
    Evaluate $\{f(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), \ldots, c_m(\mathbf{x})\}$ for all $\mathbf{x} \in \mathcal{S}^k$ (opportunistically)
    if $f_\Omega(\mathbf{x}) < f_\Omega(\mathbf{x}^k)$ for some $\mathbf{x} \in \mathcal{S}^k$ **then**
      set $\mathbf{x}^{k+1} \leftarrow \mathbf{x}$, $\Delta^{k+1} \leftarrow \tau^{-1}\Delta^k$
      go to **Termination**.
    end
  Cache Search (optional and only for $k > 0$)
    if $f_\Omega(\mathbf{x}^*) < f_\Omega(\mathbf{x}^k)$ for some $\mathbf{x}^* \in \mathcal{C}$ **then**
      set $\mathbf{x}^{k+1} \leftarrow \mathbf{x}^*$, $\Delta^{k+1} \leftarrow \tau\Delta^*$
      go to **Termination**.
    end
  Poll
    Build the poll set $\mathcal{P}^k \subset \mathcal{M}^k$.
    Perform evaluations and compute success
    if *Success* **then**
      go to **Termination**.
    end
    Set $\mathbf{x}^{k+1} \leftarrow \mathbf{x}^k$, $\Delta^{k+1} \leftarrow \tau\Delta^k$ (no success).
  Termination
    if *Stopping criteria* **then**
      End the algorithm.
      End the w workers.
    end
end

point. In contrast, in a parallel environment, evaluations are dispatched concurrently to several computing resources. The minimum number of evaluations that will be performed before the premature termination of an iteration is, therefore, equal to the number of computing resources being used in parallel. As a result, the evaluation budget is spent more quickly. This is why a computation time budget is a better-suited constraint in a parallel environment.

MADS relies on a special structure called the mesh: at iteration k, it is a discretization of the space of variables defined by

$$\mathcal{M}^k = \left\{ \mathbf{x}^k + \delta^k \mathbf{D}\mathbf{y} \;\middle|\; \mathbf{y} \in \mathbb{N}^p \right\} \subset \mathbb{R}^n, \tag{1}$$

where $\delta^k > 0$ is the mesh size parameter, $\mathbf{D} \in \mathbb{R}^{n \times p}$ is a positive spanning matrix, and $p \in \mathbb{N}$. MADS selects the candidate trial points on the mesh.

The notion of locality is parameterized by $\Delta^k > 0$, the frame size parameter, which defines the maximal distance between the incumbent solution $\mathbf{x}^k$ and the poll points $\mathbf{x} \in \mathcal{P}^k$ with

$$\|\mathbf{x} - \mathbf{x}^k\|_\infty \leq b\Delta^k,$$

where $b = \max_{d' \in \mathbb{D}} \|d'\|_\infty$ and $\mathbb{D}$ is the set of columns of matrix $\mathbf{D}$.

At every iteration k, δ^k and Δ^k are resized via $\tau \in (0, 1) \cap \mathbb{Q}$, the mesh size adjustment parameter, such that $0 < \delta^k \leq \Delta^k$ for all k. By doing so, a variety of directions that grow densely in the unit sphere is guaranteed to be generated.

All these conditions lead to the poll step being rigidly defined, but they ensure the convergence of the method. As detailed in [10], upon some assumptions about the problem, global convergence of MADS to a local optimum is guaranteed.

The optional search step is more flexible as it allows any strategy to be used for the construction of the sets $\mathcal{S}^k$. It can be generic or specific to a problem, and only needs to generate points on the current mesh. Examples of search strategies can be found in [5, 13, 18].

MADS is often run with a cache to provide a history of the evaluations that have been carried out. The main purpose of the cache is to ascertain whether the point has been previously evaluated, thus preventing redundant evaluations. Optionally, the cache can be used during a special search step only when trial points are added to the cache asynchronously (see Section 4.1.2). If a point $\mathbf{x}^*$ in the cache is better than $\mathbf{x}^k$, it would be used as a new incumbent to start an iteration, and the corresponding frame size Δ^* is adopted. This step is disabled in a sequential environment.

In practice, the stopping criteria indicated in the algorithm can be multiple. Typically, they include a maximum number of blackbox evaluations, a time budget, and/or a minimal mesh or frame size.

MADS provides two distinct approaches to deal with inequality constraints. First, a straightforward extension with extreme barrier (EB), which simply defines $f_\Omega(\mathbf{x}) = f(\mathbf{x})$ if $\mathbf{x} \in \Omega$, and $f_\Omega(\mathbf{x}) = \infty$ otherwise. Algorithm 1 is based on the EB. The second approach is referred to as the progressive barrier (PB) and is defined in [8]. It considers a function $h(\mathbf{x})$ that measures the constraint violations at $\mathbf{x}$ and ranks the iterates in a graph (called a filter) where compromises in terms of f (objective) and h (constraints) are defined.

Finally, MADS is available in the open-source solver `NOMAD` [44] at `https://www.gerad.ca/en/software/nomad/`. This solver will be used for our numerical comparisons.

4 Parallel versions of MADS

This section describes the four parallel methods of MADS, namely pMADS-S for "Parallel MADS-Synchronous" (Section 4.1.1), pMADS-A for "Parallel MADS-Asynchronous" (Section 4.1.2), COOP-MADS for "Cooperative MADS" (Section 4.2), and PSD-MADS for "Parallel Space Decomposition MADS" (Section 4.3).

These methods are defined such that the convergence properties of MADS also hold. While PSD-MADS has been the subject of [9], the other three methods have never been described or benchmarked in the literature. These methods follow the master-worker paradigm typically implemented with the `Message Passing Interface` (`MPI`) library [57]. Following the parallel strategies described in Section 2, pMADS-A and pMADS-S are classified as functional parallelism, while PSD-MADS is classified as domain decomposition parallelism. Lastly, COOP-MADS is classified as multi-search parallelism.

4.1 pMADS

The first two parallel versions of MADS are both labelled pMADS and follow Algorithm 1. In these methods, the evaluation of trial points $\mathcal{P}^k$ is performed in parallel in a master-worker fashion (see Figure 3). In doing so, the master is responsible for sending evaluation points to the workers, who are then tasked with evaluating them (see Figure 2). The two versions of pMADS are the synchronous variant (pMADS-S) and the asynchronous variant (pMADS-A). They differ in the management of the evaluations by the w workers (shaded area in Algorithm 1). Starting and ending the workers are managed by the master process when the algorithm ends.

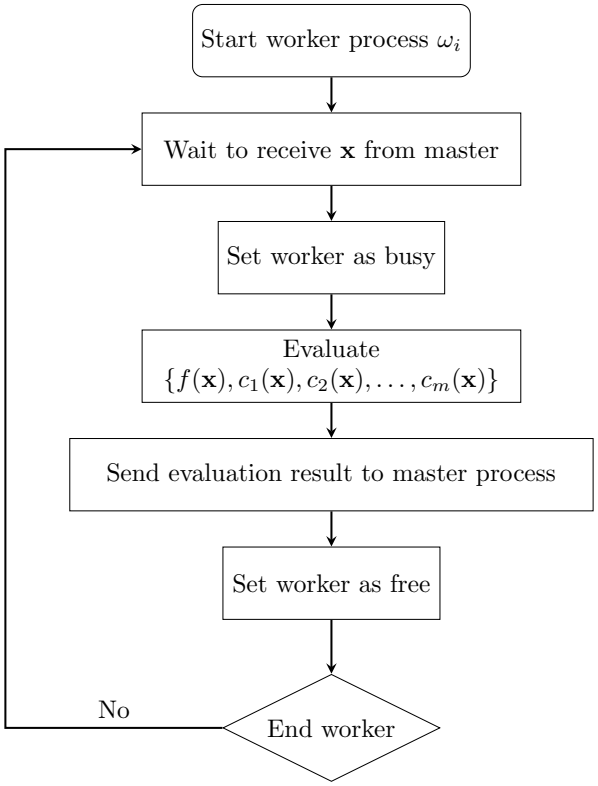

Figure 2: Flow chart of a single worker process for pMADS.

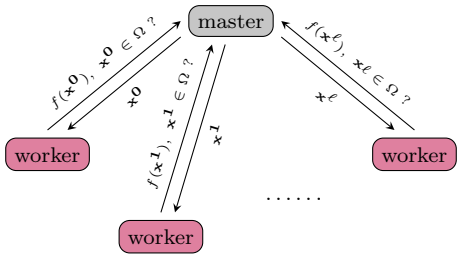

Figure 3: Master-worker communication organization for pMADS.

4.1.1 pMADS-S

pMADS-S has a deterministic behaviour, but at the cost of introducing a synchronization barrier which restricts the master to end an iteration only when all evaluations in progress are terminated. Consequently, this leads to idle computing resources at each iteration and hence to a loss in efficiency. The synchronization barrier eliminates the need for the cache search step.

An opportunistic block evaluation approach, using a block of size w, is achieved by executing the **Termination** step upon success, as shown in Figure 4.

4.1.2 pMADS-A

As mentioned in [44], pMADS-A is a simplified version of the asynchronous parallel pattern search algorithm (`APPSPACK`) [24]. The master is permitted to terminate an iteration and to generate new trial points as soon as a new success is obtained by a worker, even if evaluations are still in progress on other workers.

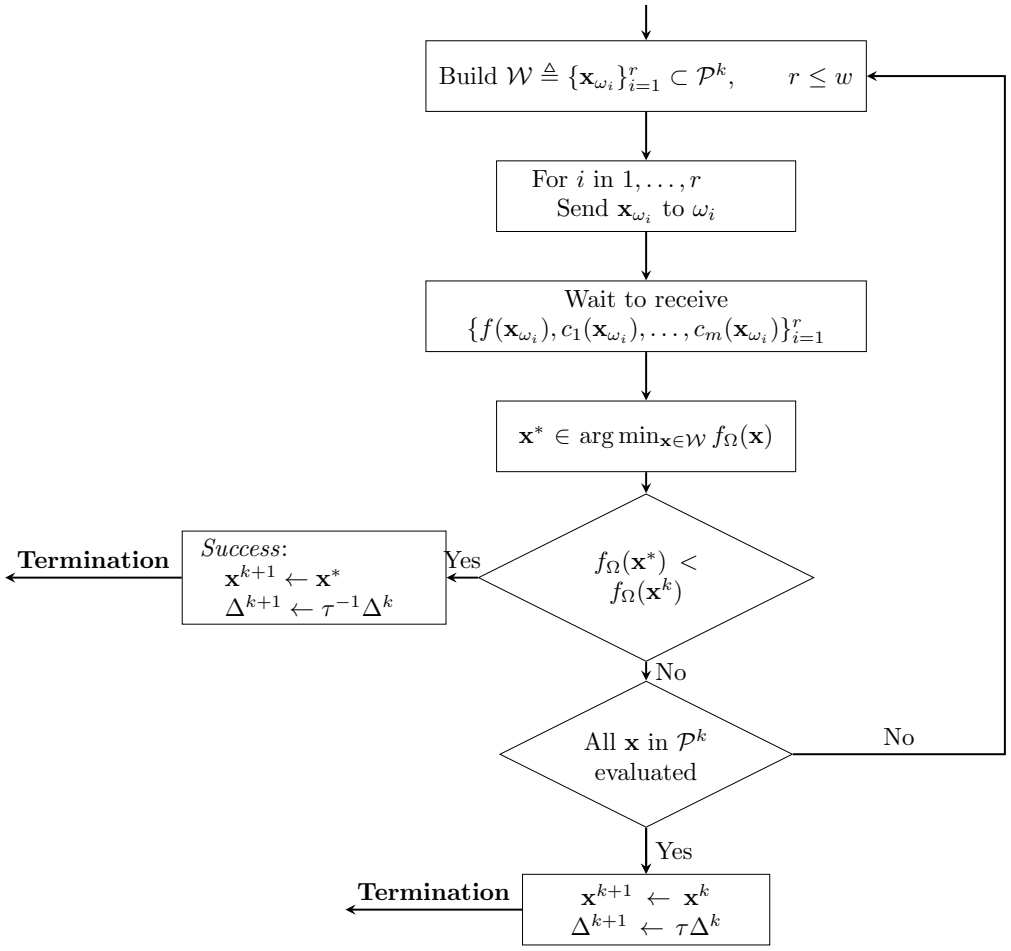

Figure 4: Master process flow chart of pMADS-S to manage parallel evaluations by the w parallel workers allocated to NOMAD, where ω_i represents their respective indices. Details the shaded area in Poll step of Algorithm 1.

In the NOMAD implementation of pMADS-A, evaluations currently in progress are never cancelled (see Figure 5). But once all $\mathbf{x}$ in $\mathcal{P}^k$ have been sent to workers for evaluation and no success has been obtained among the received evaluations, the current incumbent solution and the frame size are updated. In the event that an evaluation that took place in an iteration prior to the current one finds the best solution, NOMAD backtracks and sets this solution as the incumbent one during a cache search step. Therefore, there is no synchronization barrier between MPI processes during the Poll step. This makes this mode effective for time-varying blackboxes because no MPI processes need to wait for the most time-consuming one to finish evaluating its solution. It is imperative to avoid idling computing resources, as this leads to poor overall performance and induces bottlenecks in a parallel programming context. This assertion is elaborated upon in Section 5 where numerical results are provided and discussed to support this claim. The main downside of the asynchronous mode is pMADS-A's non-deterministic behaviour, which limits the reproducibility of the results. We remark that pMADS-A is the default mode used in NOMAD when parallelism is enabled. Also, one can note that there still exists a synchronization barrier between the Search and Poll steps that can affect the algorithm effectiveness. Notice that, to account for the possibility that evaluations may still be in progress during the **Termination** step, the stopping criteria of Algorithm 1 has to check if all evaluations are terminated.

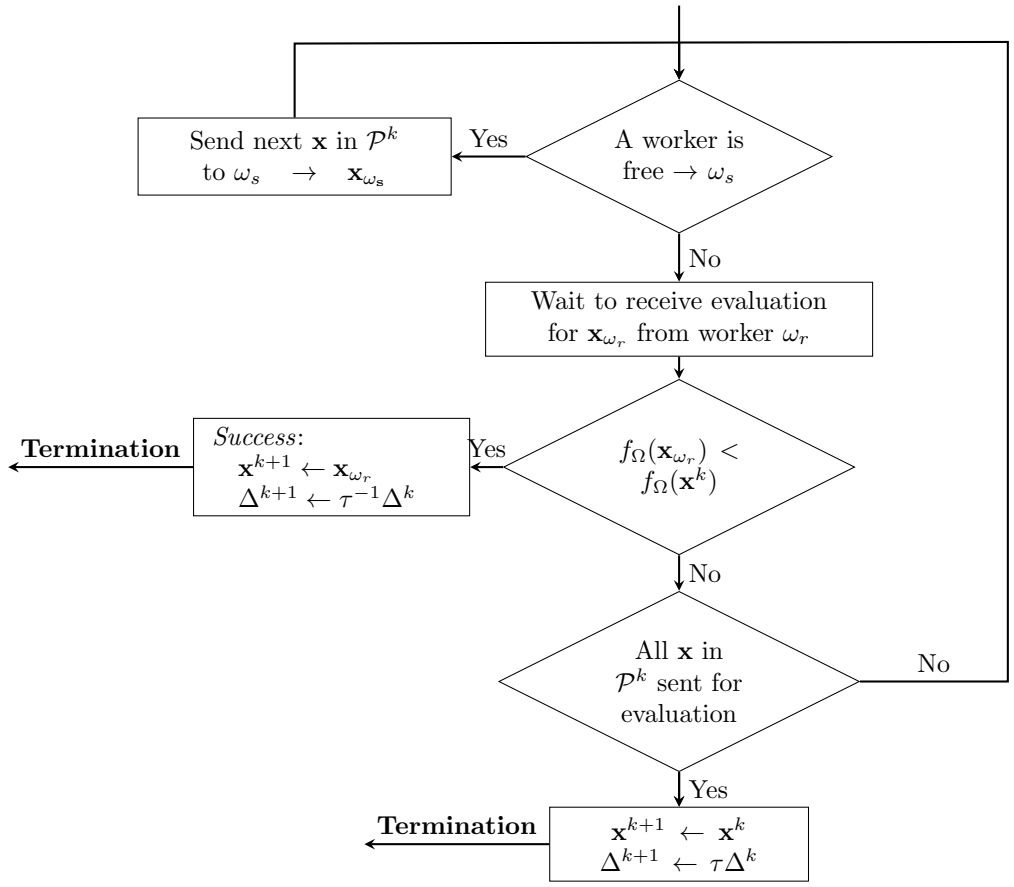

Figure 5: Flow chart of the master process for pMads-A. Details the shaded area in Poll step of Algorithm 1.

4.2 COOP-MADS

To accelerate the exploration of $\mathcal{X}$ and to prevent a fast convergence of MADS to a local optimum, multiple paths can be generated by running c different instances of MADS in parallel. This can be particularly beneficial for multi-modal problems where the aim is to converge at a global optimum. Each MADS instance is initialized with a unique random seed, which leads to a varied selection of directions from OrthoMADS [1]. Hence, it yields distinct behavioural patterns in trial points generation. These MADS instances operate independently and without any synchronization mechanisms between them. To balance the evaluation effort, a dynamic distribution of the global evaluation budget is implemented between instances. Note that a static budget allocation which consists in distributing evaluations equally between instances often results in the resource under-utilization due to the varying convergence rates of individual instances.

In COOP-MADS, the only information exchange between MADS instances is via the shared cache to prevent redundant evaluations in an instance and between the instances. The optional cache search could also be enabled to periodically synchronize the incumbent solution between the instances. This option was not used in this work. Also, for simplicity, each MADS instance runs as a single process to perform the algorithm and the evaluations. The master-worker paradigm used in pMADS is not considered for COOP-MADS.

4.3 PSD-MADS

Solving high-dimensional BBO problems is a computationally demanding task. PSD-MADS [9] stands out as an algorithm tailored specifically to large-scale BBO problems. It specifically leverages a parallel space decomposition approach to increase MADS' scalability. The core concept of PSD-MADS is centred around the manager process generating optimization subproblems $\mathbb{P}_p$ by randomly selecting variable subspaces $\mathcal{N}_p$ from the original extensive variable space $\mathcal{N}$, i.e., $\mathcal{N}_p \subseteq \mathcal{N}$. These subproblems are then distributed among workers by the manager. Each worker applies MADS on their assigned subproblem and aims to improve the shared incumbent solution $\mathbf{x}^*$. Upon completing its task a worker receives a new subproblem $\mathbb{P}_p$ from the manager. A subproblem is defined by the incumbent solution given as an initial point, the initial and minimum mesh sizes Δ_0^p and $\Delta_{\min}^p$ and the fixed variables in $\mathcal{N}$ to define the variable subspace $\mathcal{N}_p$.

Note that the PSD-MADS operates asynchronously as no synchronization step is present. The algorithm introduces four roles, namely, the manager, the pollster, the regular worker, and the cache server. The pollster is a special worker which solves the original problem ($\mathbb{P}$) with a single-poll direction MADS algorithm instead of the conventional $2n$ or $n+1$ poll directions. The pollster is crucial to ensure the convergence of PSD-MADS. The manager assumes the responsibility of selecting the optimization subproblems $\mathbb{P}_p$ and performs adjustments to the main mesh and continuously updates the value of the incumbent solution. The cache server stores all evaluated points to avoid unnecessary, multiple, and expensive function evaluations. Finally, a regular worker is a MADS instance that solves a subproblem $\mathbb{P}_p$.

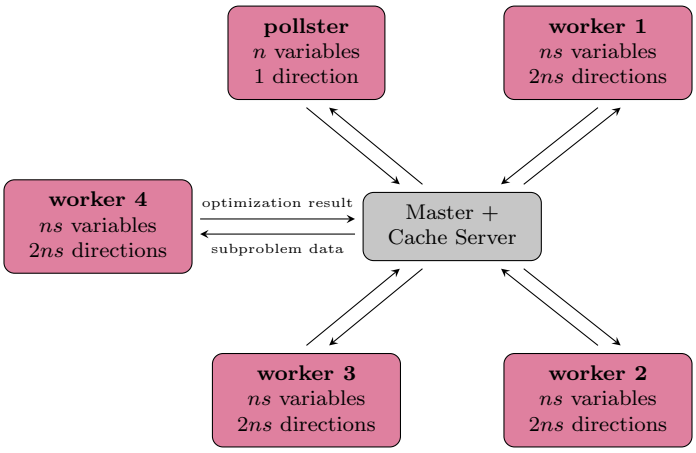

Figure 6: Master-worker communication organization for PSD-MADS, where for each worker p, the optimization result corresponds to the final mesh size Δ_{end}^p and the final solution $\mathbf{x}_p$, while the subproblem data corresponds to the starting solution $\mathbf{x}_0$, the set of variables $\mathcal{N}_p$, and the initial and minimum mesh sizes Δ_0^p and $\Delta_{\min}^p$.

5 Computational study

In this section, we conduct numerical experiments to evaluate the performance of the described algorithms. We use different sets of test problems each of which are targeting specific blackbox characteristics, to make recommendations on when an algorithm should be used. All the experiments are run on Intel Core i7-12700 @ 2.10GHz processors using Version `3.9.1` of `NOMAD` [44] with `MPICH 4.3.0`. We disable the use of quadratic models because this feature is not available in PSD-MADS. Our tests use OrthoMADS $2n$ due to the lack of an OrthoMADS $n+1$ implementation in pMADS-A and PSD-MADS. The number of parallel resources used is set according to the dimension of the given problem and the number of parallel workers that can compute an evaluation given an algorithm. Hence, depending on the algorithm, the number of parallel workers changes. For pMADS, a single process out of the total number provided to `MPI` is used as the master, leading to the use of $n+1$ processes. For COOP-MADS,

we consider $n+1$ MADS instance processes for the algorithm. For PSD-MADS, we use a master and a cache server that do not compute evaluations, leading then to a total of $n+2$ processes. Blackbox evaluations are performed in batch mode, with input and output files managed by `NOMAD`.

5.1 Scalability analysis

The performance of a parallel implementation is mainly characterized by its capacity to reduce a task execution time based on the computing units it uses, which refers to scalability. We use the following two metrics to measure scalability: speedup and efficiency. Let $S(\eta, p)$ denote the speedup for a task of size η using p computing units, defined as:

$$S(\eta, p) = \frac{T_1(\eta)}{T_p(\eta)},$$

where $T_1(\eta)$ and $T_p(\eta)$ are the execution times of the sequential and parallel implementations, respectively. Let the efficiency $E(\eta, p)$ be defined as:

$$E(\eta, p) = \frac{S(\eta, p)}{p}.$$

The efficiency relates the speedup to the number of parallel processing units used to achieve the task. In our case, the task to parallelize is not predetermined and it is the result of a dynamic aggregation of smaller tasks, namely, the blackbox functions evaluations, as the optimization algorithm proceeds. Hence, efficiency and speedup depend on communication overheads, blackbox evaluation times, and possible synchronization barriers in the algorithm. Maximum speedup and efficiency occur when communication and synchronization overheads stay small for a relatively large number of computing units p.

A scalability analysis is performed with a fixed size η of 500 blackbox evaluations. A comparison between pMADS-A runs on a blackbox from [39]' collection set with different artificial evaluation times (`tBBE`) is presented in Figure 7. These results illustrate that parallel implementation significantly accelerates the computation time. Even at the lowest measured efficiency, a speedup of 43 is achieved using 64 processors. When subject to an evaluation time of 30 seconds, the speedup increases from 43 to 54. Hence, for more time-consuming blackbox functions, the parallel performance improves further, as the impact of communication overheads becomes negligible.

5.2 Heterogeneous test

Next, we illustrate the impact of the synchronisation barrier over the computation time by comparing pMADS-S and pMADS-A. The solar simulator [4] implements in `C++` a concentrated solar power thermal plant which relies on numerical methods such as Monte Carlo simulation, Newton's method, kernel smoothing, and other iterative methods. The heterogeneous nature of the blackbox comes from combining these methods. The blackbox solar offers ten instances, from which we use solar4.1 because it presents the largest dimension for a single-objective constrained optimization problem, that is, 29 variables and 16 constraints.

Figure 8 illustrates that as more processors are involved, the run time difference between pMADS-S and pMADS-A increases. This is due to the presence of the poll synchronization barrier in pMADS-S because, e.g., it only takes one evaluation with a significantly higher computation time than the others running in parallel to cause substantial computation resource idling. Increasing the number of cores then leads to more evaluations being launched in parallel, thereby increasing the likelihood of encountering a more costly evaluation at each iteration. The potential slow down is particularly sensitive to the blackbox evaluation time distribution but is completely avoided in the case of pMADS-A which steadily decreases its run time.

Due to its clear limitation, for the rest of this work, pMADS-S will be omitted.

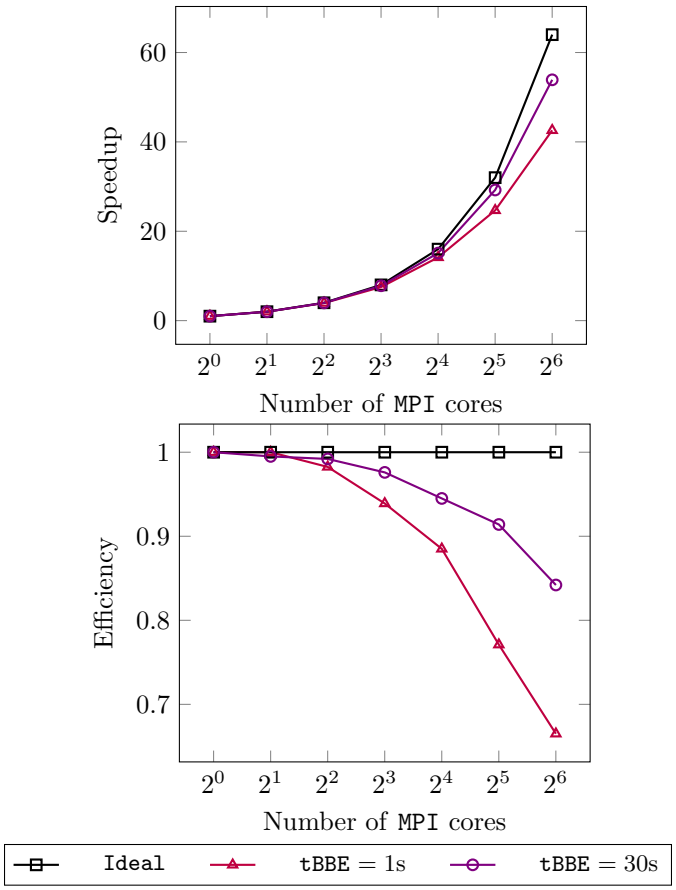

Figure 7: Parallel speedup and efficiency of pMADS-A measured over an optimization problem of 32 variables with a budget of 500 evaluations and a blackbox evaluation time (tBBE) of 1s and 30s. The reported number of MPI cores represents the worker processors.

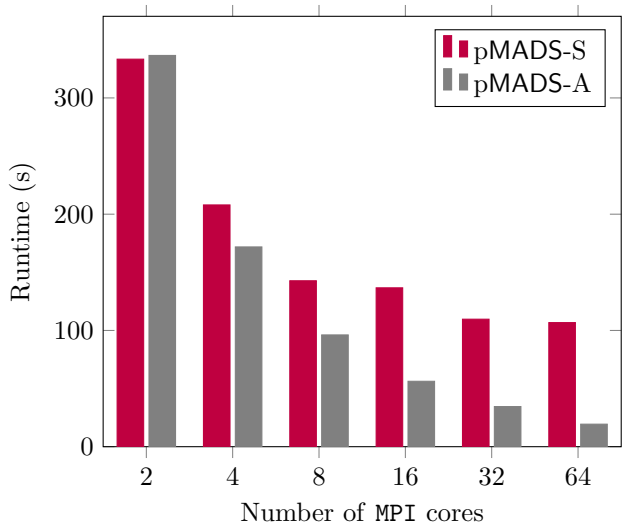

Figure 8: Running time in seconds of pMADS-A and pMADS-S for different numbers of MPI cores when solving solar4.1 with a 500 blackbox evaluation budget. The reported number of MPI cores represents both the master and worker processors.

5.3 Benchmarking with data profiles

Data profiles are used for benchmarking algorithms [47]. Comparisons are done in terms of computation time and number of evaluations scaled with problem dimensions n_p. The data profile of an algorithm is a series of values $d_a(k)$ that gives the proportion of τ-solved problems, with k the groups of $n_p + 1$ function evaluations or the time [11]. The accuracy of an evaluation point is obtained at a given time or evaluation number for a given problem using the best and the initial objective values.

In this work, we have considered the cross-instances best objective value f^* obtained by all algorithms on all run instances of a given problem. The initial points of problems are given and fixed. Different run instances are obtained by providing fixed seeds to run algorithms.

5.4 Moré-Wild tests

In this section, the parallel versions of MADS, are compared on the Moré-Wild [47] problems. The Moré-Wild collection is a well-know test set for benchmarking unconstrained DFO algorithms in low dimensions. We select the SMOOTH subset which provides 53 unconstrained optimization problems where the objective function is twice continuously differentiable. For each parallel version of MADS and each problem, runs are conducted with 10 different random seeds that affect the trial points generation. This gives a total of 530 run instances per algorithm tested.

In Figure 9, for $\tau = 10^{-5}$ and below, MADS and COOP-MADS solve more problems than the pMADS and PSD-MADS for the given evaluation budgets. We note that evaluation profiles for pMADS and PSD-MADS are still increasing towards the end. This probably indicates that more problem could have been solved by giving a higher evaluation budget. When looking at portion of solved instances with respect to time, pMADS and PSD-MADS achieve their best results faster, with flatter profiles.

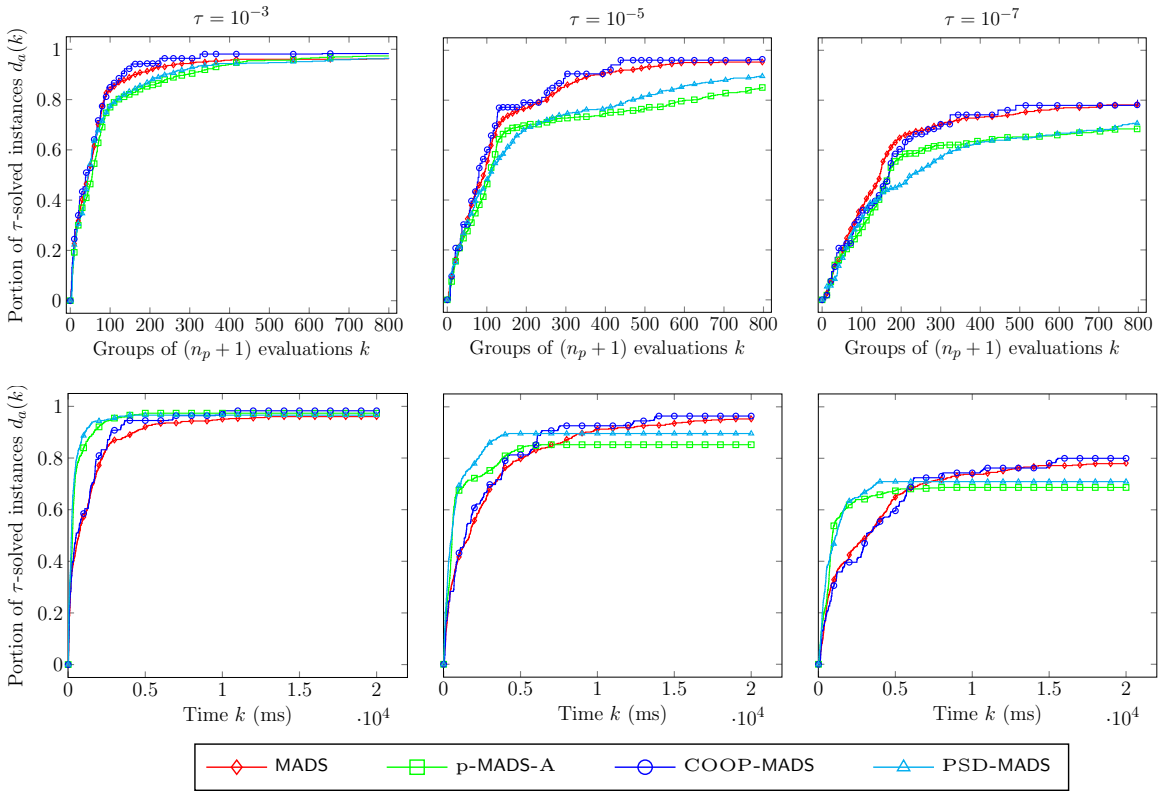

Figure 9: Data profiles on the Moré-Wild SMOOTH problem set.

Table 1: Description of the selected continuous constrained test set. The superscript § denotes the use of multiple starting points for the given problem.

#	Name	Ref.	n	m	Bnds
1	CHENWANG-F2§	[67]	8	6	yes
2	CHENWANG-F3§	[67]	10	8	yes
3	CRESCENT	[8]	10	2	yes
4	DISK	[8]	10	1	no
5	FLOUDAS	[22]	3	3	yes
6	G2	[33]	10	2	yes
7	G9	[33]	7	4	yes
8	HS83	[34]	5	6	yes
9	HS108	[34]	9	13	yes
10	HS114	[34]	9	6	yes
11	MAD6§	[67]	5	7	yes
12	OPTENG-RBF	[8]	3	4	yes
13	PENTAGON	[46]	6	15	no
14	SPRING§	[54]	3	4	yes
15	TAOWANG-F2§	[59]	7	4	yes
16	ZHAOWANG-F5	[67]	13	9	yes

5.5 Constrained test

Lastly, we consider a set of continuous constrained problems and use parallel versions of MADS with progressive barrier [8]. The progressive barrier offers a more sophisticated approach to handling inequality constraints compared to the extreme barrier described in Algorithm 1. As the default strategy in `NOMAD`, the progressive barrier influences both the detection of success and the management of incumbents, allowing for the identification of two incumbent solutions; a feasible and an infeasible one. Despite its added complexity, it has demonstrated greater effectiveness than the extreme barrier, while maintaining comparable parallelization capabilities.

Table 1 lists and references the problem instances. Ten seeds are again considered to increase the problem collection. Data profiles for different tolerances τ are compared in Figure 10. PSD-MADS do not perform well compared with the other algorithms. The strategy of space decomposition seems not well adapted on small dimension problems with inequality constraints. pMADS-A performance is higher on smaller computation time and as good as COOP-MADS and MADS in terms of evaluations.

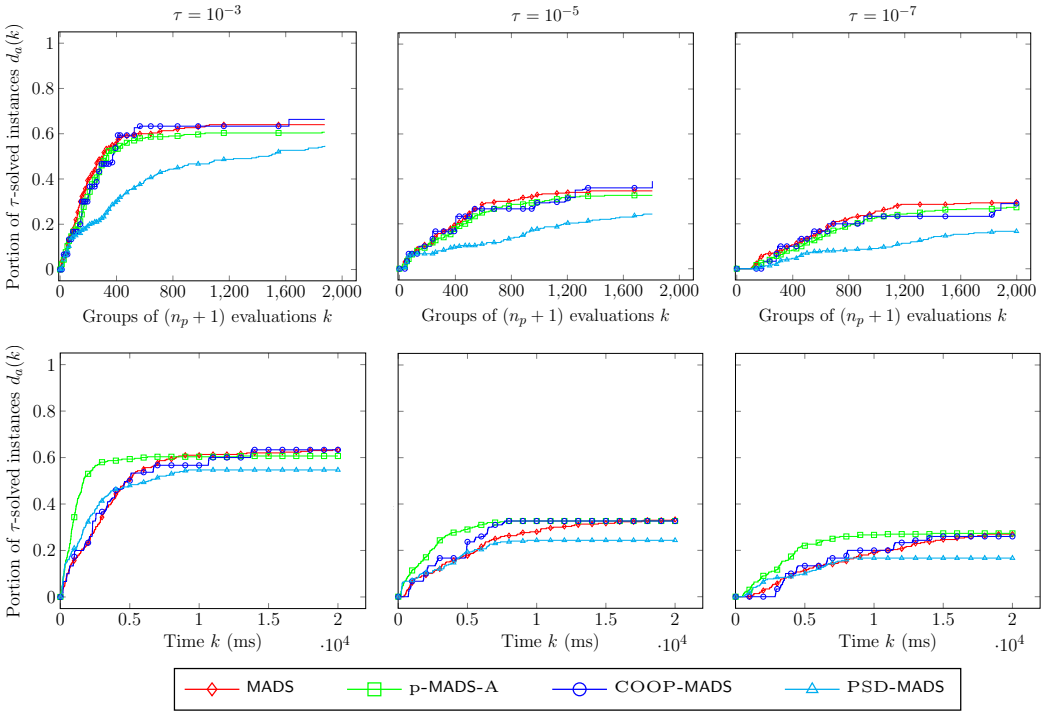

Figure 10: Data profiles on the inequality-constrained problem set of Table 1.

5.6 Discussion

The previous subsections numerically illustrates the effectiveness of parallel variants of the MADS algorithm. The results highlight the critical importance of employing an asynchronous parallel strategy when subject to heterogeneous evaluation times among blackbox functions. Lastly, the pMADS implementations demonstrates strong scalability performance with its efficient utilization of up to 64 computing cores. This in turn confirms its suitability for parallel computing contexts.

6 Closing remarks

In this work, we discuss and numerically compare the parallel extensions of the mesh adaptive direct search (MADS) algorithm for blackbox optimization (BBO). First, variants with a synchronization barrier, parallel MADS-Synchronous (pMADS-S) and without, parallel MADS-Asynchronous (pMADS-A), between evaluations are presented. While omitting the barrier improves computational efficiency, it also leads to the non-deterministic behaviour of pMADS-A. Next, the cooperative-MADS (COOP-MADS) algorithm, where several MADS instances are run in parallel with a shared cache, is outlined. Lastly, we survey the parallel space decomposition MADS (PSD-MADS), which specifically targets high-dimensional BBO problems by employing MADS on variable subspaces. The aforementioned extensions are designed such that the local convergence analysis of MADS holds for them as well. Finally, the benefits of each variant, e.g., in terms of scalability, efficiency, running time and/or number of evaluations, is exemplified in numerical simulations. In both unconstrained and constrained problems, p-MADS-A stands out as a time-effective extension for BBO on multiple cores. In the future, our study is to be extended to the `NOMAD` 4 implementation of MADS. While, the resolution performance should be similar, the computation time and resource scalability may be affected by the its shared memory implementation.

Data availability statement

The open-source solver `NOMAD` [44] is available at https://www.gerad.ca/en/software/nomad/.

Use of AI statement

The authors used AI-assisted tools (ChatGPT, OpenAI) solely to improve the language and readability of the manuscript; all scientific content and conclusions are entirely the authors' own.

References

[1] M.A. Abramson, C. Audet, J.E. Dennis, Jr., and S. Le Digabel. OrthoMADS: A Deterministic MADS Instance with Orthogonal Directions. SIAM Journal on Optimization, 20(2):948–966, 2009.

[2] S. Alarie, C. Audet, A.E. Gheribi, M. Kokkolaras, and S. Le Digabel. Two decades of blackbox optimization applications. EURO Journal on Computational Optimization, 9:100011, 2021.

[3] E. Alba, G. Luque, and S. Nesmachnow. Parallel metaheuristics: recent advances and new trends. International Transactions in Operational Research, 20(1):1–48, 2013.

[4] N. Andrés-Thió, C. Audet, M. Diago, A.E. Gheribi, S. Le Digabel, X. Lebeuf, M. Lemyre Garneau, and C. Tribes. Solar: a solar thermal power plant simulator for blackbox optimization benchmarking. Optimization and Engineering, 26(3):1815–1861, 2025.

[5] C. Audet, V. Béchard, and S. Le Digabel. Nonsmooth optimization through Mesh Adaptive Direct Search and Variable Neighborhood Search. Journal of Global Optimization, 41(2):299–318, 2008.

[6] C. Audet, C.-K. Dang, and D. Orban. Efficient use of parallelism in algorithmic parameter optimization applications. Optimization Letters, 7(3):421–433, 2013.

[7] C. Audet and J.E. Dennis, Jr. Mesh Adaptive Direct Search Algorithms for Constrained Optimization. SIAM Journal on Optimization, 17(1):188–217, 2006.

[8] C. Audet and J.E. Dennis, Jr. A Progressive Barrier for Derivative-Free Nonlinear Programming. SIAM Journal on Optimization, 20(1):445–472, 2009.

[9] C. Audet, J.E. Dennis, Jr., and S. Le Digabel. Parallel Space Decomposition of the Mesh Adaptive Direct Search Algorithm. SIAM Journal on Optimization, 19(3):1150–1170, 2008.

[10] C. Audet and W. Hare. "Derivative-Free and Blackbox Optimization". Springer Series in Operations Research and Financial Engineering. Springer, Cham, Switzerland, 2nd edition, 2026.

[11] C. Audet, W. Hare, and C. Tribes. A summary of benchmarking constrained, multi-objective and surrogate-assisted optimization methods. Optimization Letters, 2026.

[12] C. Audet, S. Le Digabel, and C. Tribes. The Mesh Adaptive Direct Search Algorithm for Granular and Discrete Variables. SIAM Journal on Optimization, 29(2):1164–1189, 2019.

[13] C. Audet and C. Tribes. Mesh-based Nelder-Mead algorithm for inequality constrained optimization. Computational Optimization and Applications, 71(2):331–352, 2018.

[14] F.V. Berghen. CONDOR: A Constrained, Non-Linear, Derivative-Free Parallel Optimizer for Continuous, High Computing Load, Noisy Objective Functions. PhD thesis, Université Libre de Bruxelles, Belgium, 2004.

[15] F.V. Berghen and H. Bersini. CONDOR, a new parallel, constrained extension of Powell's UOBYQA algorithm: Experimental results and comparison with the DFO algorithm. Journal of Computational and Applied Mathematics, 181:157–175, 2005.

[16] G. Briffoteaux. Parallel surrogate-based algorithms for solving expensive optimization problems. PhD thesis, Université de Lille; Université de Mons, 2023.

[17] G. Campos, W. da Silva Pereira, R. Vercellino, J. Mueller, and M. Mann. Parallel derivative-free optimization for simulation-based design of behind-the-meter energy systems. Computers and Chemical Engineering, page 109422, 2025.

[18] A.R. Conn and S. Le Digabel. Use of quadratic models with mesh-adaptive direct search for constrained black box optimization. Optimization Methods and Software, 28(1):139–158, 2013.

[19] A.R. Conn, K. Scheinberg, and L.N. Vicente. "Introduction to Derivative-Free Optimization". MOS-SIAM Series on Optimization. SIAM, Philadelphia, 2009.

[20] J.E. Dennis, Jr. and V. Torczon. Direct Search Methods on Parallel Machines. SIAM Journal on Optimization, 1(4):448–474, 1991.

[21] J.E. Dennis, Jr. and Z. Wu. Parallel continuous optimization, pages 649–670. Sourcebook of parallel computing. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

[22] C.A. Floudas. "Nonlinear and mixed-integer optimization: fundamentals and applications". Oxford University Press, New York, 1995.

[23] J.C. García-García, R. García-Ródenas, and E. Codina. A surrogate-based cooperative optimization framework for computationally expensive black-box problems. Optimization and Engineering, 21(3):1053–1093, 2020.

[24] G.A. Gray and T.G. Kolda. Algorithm 856: APPSPACK 4.0: Asynchronous parallel pattern search for derivative-free optimization. ACM Transactions on Mathematical Software, 32(3):485–507, 2006.

[25] J.D. Griffin and T.G. Kolda. Asynchronous parallel hybrid optimization combining DIRECT and GSS. Optimization Methods and Software, 25(5):797–817, 2010.

[26] J.D. Griffin and T.G. Kolda. Nonlinearly-constrained optimization using heuristic penalty methods and asynchronous parallel generating set search. Applied Mathematics Research eXpress, 25(5):36–62, 2010.

[27] J.D. Griffin, T.G. Kolda, and R.M. Lewis. Asynchronous parallel generating set search for linearly-constrained optimization. SIAM Journal on Scientific Computing, 30(4):1892–1924, 2008.

[28] R.T. Haftka, D. Villanueva, and A. Chaudhuri. Parallel surrogate-assisted global optimization with expensive functions – a survey. Structural and Multidisciplinary Optimization, 54(1):3–13, 2016.

[29] J. He, A. Verstak, M. Sosonkina, and L.T. Watson. Performance Modeling and Analysis of a Massively Parallel DIRECT–Part 2. International Journal of High Performance Computing Applications, 23(1):29–41, 2009.

[30] J. He, A. Verstak, L.T. Watson, and M. Sosonkina. Design and implementation of a massively parallel version of DIRECT. Computational Optimization and Applications, 40(2):217–245, 2007.

[31] J. He, A. Verstak, L.T. Watson, and M. Sosonkina. Performance Modeling and Analysis of a Massively Parallel DIRECT–Part 1. International Journal of High Performance Computing Applications, 23(1):14–28, 2009.

[32] J. He, L.T. Watson, and M. Sosonkina. Algorithm 897: VTDIRECT95: Serial and parallel codes for the global optimization algorithm DIRECT. ACM Transactions on Mathematical Software, 36(3):1–24, 2009.

[33] A.-R. Hedar and M. Fukushima. Derivative-free filter simulated annealing method for constrained continuous global optimization. Journal of Global Optimization, 35(4):521–549, 2006.

[34] W. Hock and K. Schittkowski. "Test Examples for Nonlinear Programming Codes", volume 187 of Lecture Notes in Economics and Mathematical Systems. Springer, Berlin, Germany, 1981.

[35] P.D. Hough, T.G. Kolda, and V. Torczon. Asynchronous Parallel Pattern Search for Nonlinear Optimization. SIAM Journal on Scientific Computing, 23(1):134–156, 2001.

[36] P.D. Hough and J.C. Meza. A Class of Trust-Region Methods for Parallel Optimization. SIAM Journal on Optimization, 13(1):264–282, 2002.

[37] D.R. Jones, C.D. Perttunen, and B.E. Stuckman. Lipschitzian optimization without the Lipschitz constant. Journal of Optimization Theory and Application, 79(1):157–181, 1993.

[38] D.R. Jones, M. Schonlau, and W.J. Welch. Efficient Global Optimization of Expensive Black Box Functions. Journal of Global Optimization, 13(4):455–492, 1998.

[39] N. Karmitsa. Test problems for large-scale nonsmooth minimization. Technical Report B. 4/2007, Department of Mathematical Information Technology, University of Jyväskylä, Jyväskylä, Finland, 2007.

[40] T.G. Kolda. Revisiting Asynchronous Parallel Pattern Search for Nonlinear Optimization. SIAM Journal on Optimization, 16(2):563–586, 2005.

[41] T.G. Kolda, R.M. Lewis, and V. Torczon. Optimization by direct search: New perspectives on some classical and modern methods. SIAM Review, 45(3):385–482, 2003.

[42] T.G. Kolda and V. Torczon. On the Convergence of Asynchronous Parallel Pattern Search. SIAM Journal on Optimization, 14(4):939–964, 2004.

[43] J. Larson and S.M. Wild. Asynchronously Parallel Optimization Solver for Finding Multiple Minima. Mathematical Programming Computation, 10(3):303–332, 2018.

[44] S. Le Digabel. Algorithm 909: NOMAD: Nonlinear Optimization with the MADS algorithm. ACM Transactions on Mathematical Software, 37(4):44:1–44:15, 2011.

[45] G. Liuzzi and K. Truemper. Parallelized hybrid optimization methods for nonsmooth problems using NOMAD and linesearch. Computational and Applied Mathematics, 37(3):3172–3207, 2018.

[46] L. Lukšan and J. Vlček. Test Problems for Nonsmooth Unconstrained and Linearly Constrained Optimization. Technical Report V-798, ICS AS CR, 2000.

[47] J.J. Moré and S.M. Wild. Benchmarking Derivative-Free Optimization Algorithms. SIAM Journal on Optimization, 20(1):172–191, 2009.

[48] J.A. Nelder and R. Mead. A Simplex Method for Function Minimization. The Computer Journal, 7(4):308–313, 1965.

[49] P.-M. Olsson. Methods for Network Optimization and Parallel Derivative-Free Optimization. PhD thesis, Linköping University, 2014.

[50] Y. Ozaki, S. Watanabe, and M. Onishi. Accelerating the Nelder-Mead method with predictive parallel evaluation. 6th ICML Workshop on Automated Machine Learning, 185:186, 2019.

[51] M. Pang and C.A. Shoemaker. Comparison of parallel optimization algorithms on computationally expensive groundwater remediation designs. Science of The Total Environment, 857:159544, 2023.

[52] T.D. Plantenga. HOPSPACK 2.0 User Manual. Technical Report SAND2009-6265, Sandia National Laboratories, Livermore, CA, October 2009.

[53] R.G. Regis and C.A. Shoemaker. Parallel radial basis function methods for the global optimization of expensive functions. European Journal of Operational Research, 182(2):514–535, 2007.

[54] J.F. Rodríguez, J.E. Renaud, and L.T. Watson. Trust Region Augmented Lagrangian Methods for Sequential Response Surface Approximation and Optimization. Journal of Mechanical Design, 120(1):58–66, 1998.

[55] G. Schryen. Parallel computational optimization in operations research: A new integrative framework, literature review and research directions. European Journal of Operational Research, 287(1):1–18, 2020.

[56] C.A. Shoemaker and R.G. Regis. "mapo: using a committee of algorithm-experts for parallel optimization of costly functions". In "SPAA '03: Proceedings of the Fifteenth Annual ACM Symposium on Parallel Algorithms and Architectures", pages 242–243. ACM, 2003.

[57] M. Snir, S.W. Otto, S. Huss-Lederman, D.W. Walker, and J. Dongarra. "MPI: The Complete Reference". The MIT Press, Cambridge, Massachusetts, 1995.

[58] B. Talgorn, S. Alarie, and M. Kokkolaras. Parallel Surrogate-based Optimization Using Mesh Adaptive Direct Search. Technical Report G–2020–38, Les cahiers du GERAD, 2020.

[59] J. Tao and N. Wang. DNA Double Helix Based Hybrid GA for the Gasoline Blending Recipe Optimization Problem. Chemical Engineering and Technology, 31(3):440–451, 2008.

[60] S. Tavares, C.P. Brás, A.L. Custódio, V. Duarte, and P. Medeiros. Parallel strategies for Direct Multi-search. Numerical Algorithms, 92(3):1757–1788, 2023.

[61] V. Torczon. On the Convergence of Pattern Search Algorithms. SIAM Journal on Optimization, 7(1):1–25, 1997.

[62] A. Tran, J. Sun, J.M. Furlan, K.V. Pagalthivarthi, R.J. Visintainer, and Y. Wang. pBO-2GP-3B: A batch parallel known/unknown constrained Bayesian optimization with feasibility classification and its applications in computational fluid dynamics. Computer Methods in Applied Mechanics and Engineering, 347:827–852, 2019.

[63] S. Vázquez, M.J. Martín, B.B. Fraguela, A. Gómez, A. Rodríguez, and B. Elvarsson. Novel parallelization of simulated annealing and Hooke & Jeeves search algorithms for multicore systems with application to complex fisheries stock assessment models. Journal of Computational Science, 17:599–608, 2016.

[64] W. Xia and C.A. Shoemaker. GOPS: efficient RBF surrogate global optimization algorithm with high dimensions and many parallel processors including application to multimodal water quality PDE model calibration. Optimization and Engineering, 22(4):2741–2777, 2021.

[65] D. Zhan, J. Qian, and Y. Cheng. Pseudo expected improvement criterion for parallel EGO algorithm. Journal of Global Optimization, 68(3):641–662, 2017.

[66] D. Zhan and H. Xing. Expected improvement for expensive optimization: a review. Journal of Global Optimization, 78(3):507–544, 2020.

[67] J. Zhao and N. Wang. A bio-inspired algorithm based on membrane computing and its application to gasoline blending scheduling. Computers and Chemical Engineering, 35(2):272–283, 2011.