

RESEARCH CENTRE

**Inria Centre at the University of
Lille**

IN PARTNERSHIP WITH:

Université de Lille

2024

ACTIVITY REPORT

Project-Team

BONUS

Big Optimization aNd Ultra-Scale Computing

IN COLLABORATION WITH: Centre de Recherche en Informatique, Signal
et Automatique de Lille

DOMAIN

**Applied Mathematics, Computation and
Simulation**

THEME

**Optimization, machine learning and
statistical methods**

Inria

Contents

Project-Team BONUS	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Presentation	3
3 Research program	3
3.1 Decomposition-based Optimization	3
3.2 Machine Learning-assisted Optimization	5
3.3 Ultra-scale Optimization	6
4 Application domains	8
4.1 Introduction	8
4.2 Big optimization for complex scheduling	8
4.3 Big optimization for engineering design	8
5 Social and environmental responsibility	9
6 New software, platforms, open data	9
6.1 New software	9
6.1.1 pBB	9
6.1.2 ParadiseO	9
6.1.3 pyparadiseo	10
6.1.4 pySBO	10
6.1.5 moead-framework	11
6.1.6 Zellij	11
6.2 New platforms	11
6.2.1 SLICES-FR/GRID'5000 testbed: major achievements in 2024	11
7 New results	12
7.1 Decomposition-based optimization	12
7.1.1 Multi-Objective Reinforcement Learning Based on Decomposition	12
7.1.2 Multi-objective Quantum Approximate Optimiser based on decomposition	13
7.1.3 Landscape analysis of heterogeneous multi-objective problems using decomposition	13
7.1.4 Two-Level Approach for Simultaneous Component Assignment and Layout Optimiz- ation with Applications to Spacecraft Optimal Layout	14
7.1.5 Multi-objective optimization for complex and high-dimensional problems	14
7.2 ML-assisted optimization	15
7.2.1 Neural Architecture Search and Hyperparameter Optimization	15
7.2.2 Hybrid surrogate-based optimization	16
7.2.3 Machine learning and quantum algorithms	16
7.3 Ultra-scale parallel optimization	17
7.3.1 Parallel tree-based search algorithms	17
7.3.2 Parallel hyperparameter optimization of spiking neural networks	18
7.3.3 Large-scale parallel graybox optimization	19
8 Bilateral contracts and grants with industry	19
8.1 Bilateral grants with industry	19

9 Partnerships and cooperations	20
9.1 International initiatives	20
9.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	20
9.1.2 Participation in other International Programs	20
9.2 International research visitors	21
9.2.1 Visits of international scientists	21
9.2.2 Visits to international teams	21
9.2.3 Other european programs/initiatives	22
9.3 National initiatives	22
9.3.1 ANR	22
9.4 Regional initiatives	24
10 Dissemination	24
10.1 Promoting scientific activities	24
10.1.1 Scientific events: organisation	24
10.1.2 Scientific events: selection	25
10.1.3 Journal	25
10.1.4 Invited talks	25
10.1.5 Leadership within the scientific community	26
10.1.6 Scientific expertise	26
10.1.7 Research administration	26
10.2 Teaching - Supervision - Juries	26
10.2.1 Teaching	26
10.2.2 Supervision	27
10.2.3 Juries	27
10.3 Popularization	28
10.3.1 Others science outreach relevant activities	28
11 Scientific production	28
11.1 Major publications	28
11.2 Publications of the year	29
11.3 Cited publications	32

Project-Team BONUS

Creation of the Project-Team: 2019 June 01

Keywords

Computer sciences and digital sciences

- A1.1.11. – Quantum architectures
- A3.4.4. – Optimization and learning
- A3.4.5. – Bayesian methods
- A6.2.6. – Optimization
- A6.2.7. – High performance computing
- A7.1.4. – Quantum algorithms
- A8.2.1. – Operations research
- A8.2.2. – Evolutionary algorithms
- A9.6. – Decision support
- A9.7. – AI algorithmics

Other research topics and application domains

- B3.1. – Sustainable development
- B3.1.1. – Resource management
- B7. – Transport and logistics
- B8.1.1. – Energy for smart buildings

1 Team members, visitors, external collaborators

Research Scientist

- Abdelmoiz Zakaria Dahi [INRIA, ISFP]

Faculty Members

- Bilel Derbel [Team leader, UNIV LILLE, Professor Delegation, until Aug 2024]
- Bilel Derbel [Team leader, UNIV LILLE, Professor, from Sep 2024]
- Nouredine Melab [UNIV LILLE, Professor]
- El-Ghazali Talbi [UNIV LILLE, Professor]

PhD Students

- Mahmoud El Mehdi El Khadiri [INRIA, from Oct 2024]
- Thomas Firmin [UNIV LILLE, until Sep 2024]
- Guillaume Helbecque [UNIV LILLE, until Sep 2024]
- Bohdan Ivaniuk Skulskyi [Vinci, CIFRE]
- Julie Keisler [EDF, CIFRE]
- Houssem Ouertatani [IRT SYSTEM X, until Sep 2024]
- David Redon [UNIV LILLE, ATER, from Sep 2024]
- David Redon [UNIV LILLE, until Aug 2024]
- Jerome Rouze [UNIV MONS]
- Ivan Tagliaferro De Oliveira Tezoto [UNIV LILLE, from Oct 2024]

Interns and Apprentices

- Nathan Bouvier [ENS DE LYON, Intern, from Jun 2024 until Jul 2024]
- Nathan Davouse [INRIA, Intern, from Sep 2024]

Administrative Assistants

- Nathalie Bonte [INRIA]
- Karine Lewandowski [INRIA]

Visiting Scientists

- Cosijopii Garcia Garcia [CONACYT, until Apr 2024]
- Mohammed Thousif Pagala [Indian Institute Of Science Bangalore]
- Francesco Zito [University of Catania, from Apr 2024 until May 2024]

2 Overall objectives

2.1 Presentation

Solving an optimization problem consists in optimizing (minimizing or maximizing) one or more objective function(s) subject to some constraints. This can be formulated as follows:

$$\begin{aligned} \text{Min/Max } \mathbf{F}(\mathbf{x}) &= (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_m(\mathbf{x})) \\ \text{subject to } \quad \mathbf{x} &\in \Omega, \end{aligned}$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is the decision variable vector of dimension n , Ω is the domain of \mathbf{x} (decision space), and $\mathbf{F}(\mathbf{x})$ is the objective function vector of dimension $m \geq 1$. The objective space is composed of all values of $\mathbf{F}(\mathbf{x})$ corresponding to the values of \mathbf{x} in the decision space.

Nowadays, in many research and application areas we are witnessing the emergence of the big era (big data, big graphs, etc). In the optimization setting, the problems are increasingly big in practice. Big optimization problems (BOPs) refer to problems composed of a large number of environmental input parameters and/or decision variables (high dimensionality), and/or many objective functions that may be computationally expensive. For instance, in smart grids, many optimization problems may involve a large number of consumers (appliances, electrical vehicles, etc.) and multiple suppliers with various energy sources. In the area of engineering design, the optimization process must often take into account a large number of parameters from different disciplines. In addition, the evaluation of the objective function(s) often consist(s) in the execution of an expensive simulation of a black-box complex system. This is for instance typically the case in aerodynamics where a CFD-based simulation may require several hours. On the other hand, to meet the high-growing needs of applications in terms of computational power in a wide range of areas including optimization, high-performance computing (HPC) technologies have known a revolution during the last decade (see Top500 international ranking ([Edition of November 2022](#))). Indeed, HPC is evolving toward ultra-scale supercomputers composed of millions of cores supplied in heterogeneous devices including multi-core processors with various architectures, GPU accelerators and MIC coprocessors.

Beyond the “big buzzword” as some say, solving BOPs raises at least four major challenges: (1) tackling their high dimensionality in the decision space; (2) handling many objectives; (3) dealing with computationally expensive objective functions; and (4) scaling up on (ultra-scale) modern supercomputers. The overall scientific objectives of the BONUS project consist in addressing efficiently these challenges. On the one hand, the focus will be put on the design, analysis and implementation of optimization algorithms that are scalable for high-dimensional (in decision variables and/or objectives) and/or expensive problems. On the other hand, the focus will also be put on the design of optimization algorithms able to scale on heterogeneous supercomputers including several millions of processing cores. To achieve these objectives raising the associated challenges a program including three lines of research will be adopted (Fig. 1): decomposition-based optimization, Machine Learning (ML)-assisted optimization and ultra-scale optimization. These research lines are developed in the following section.

From the software standpoint, our objective is to integrate the approaches we will develop in our [ParadisEO \[3, 50\]](#) framework in order to allow their reuse inside and outside the BONUS team. The major challenge will be to extend [ParadisEO](#) in order to make it *more collaborative* with other software including machine learning tools, other (exact) solvers and simulators. *From the application point of view*, the focus will be put on two classes of applications: *complex scheduling and engineering design*.

3 Research program

3.1 Decomposition-based Optimization

For the large-scale optimization problems we consider (wrt variables, objectives), their decomposition into simplified and loosely coupled or independent subproblems is essential to raise the challenge of scalability. The first line of research is to *investigate the decomposition approach in the two spaces (decision and objective) and their combination, as well as their implementation on ultra-scale architectures*. The motivation of the decomposition is twofold: first, the decomposition allows the parallel resolution of the

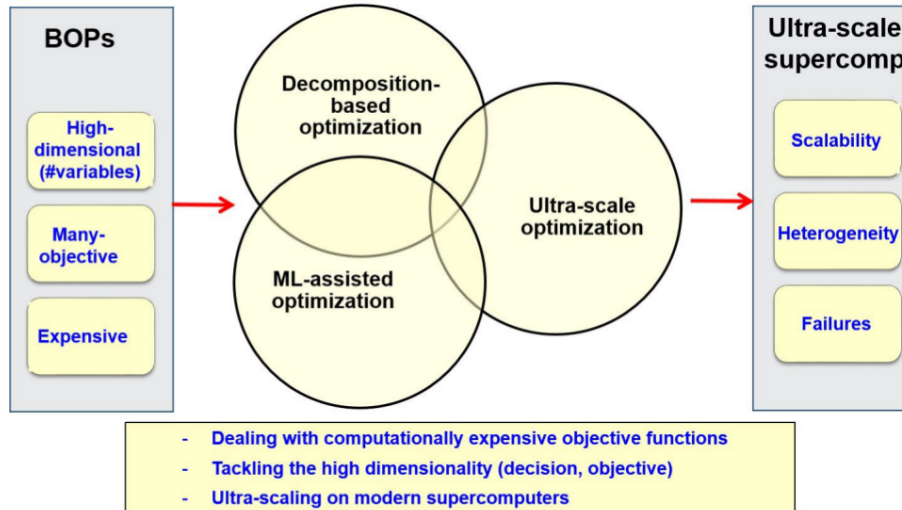


Figure 1: Research challenges/objectives and lines

resulting subproblems on ultra-scale architectures. Here also several issues will be addressed: the definition of the subproblems, their coding to allow their efficient communication and storage (checkpointing), their assignment to processing cores, etc. Second, decomposition is necessary for solving large problems that cannot be solved (efficiently) using traditional algorithms. Indeed, for instance with the popular NSGA-II algorithm the number of non-dominated solutions¹ increases drastically with the number of objectives leading to a very slow convergence to the Pareto Front². Therefore, decomposition-based techniques are gaining a growing interest. The objective of BONUS is to *investigate various decomposition schemes and cooperation protocols between the subproblems* resulting from the decomposition to generate efficiently global solutions of good quality. Several challenges have to be addressed: (1) how to define the subproblems (decomposition strategy), (2) how to solve them to generate local solutions (local rules), and (3) how to combine these latter with those generated by other subproblems and how to generate global solutions (cooperation mechanism), and (4) how to combine decomposition strategies in more than one space (hybridization strategy)?

The *decomposition in the decision space* can be performed following different ways according to the problem at hand. Two major categories of decomposition techniques can be distinguished: the first one consists in *breaking down the high-dimensional decision vector* into lower-dimensional and easier-to-optimize blocks of variables. The major issue is how to define the subproblems (blocks of variables) and their cooperation protocol: randomly *vs.* using some learning (e.g. separability analysis), statically *vs.* adaptively, etc. *The decomposition in the decision space can also be guided by the type of variables i.e. discrete vs. continuous.* The discrete and continuous parts are optimized separately using cooperative hybrid algorithms [57]. *The major issue of this kind of decomposition is the presence of categorical variables in the discrete part [53].* The BONUS team is addressing this issue, rarely investigated in the literature, within the context of vehicle aerospace engineering design. The second category consists in the *decomposition according to the ranges of the decision variables* (search space decomposition). For continuous problems, the idea consists in iteratively subdividing the search (e.g. design) space into subspaces (hyper-rectangles, intervals, etc.) and select those that are most likely to produce the lowest objective function value. *Existing approaches meet increasing difficulty with an increasing number of variables and are often applied to low-dimensional problems. We are investigating this scalability challenge (e.g. [11]). For discrete problems, the major challenge is to find a coding (mapping) of the search space to a decomposable entity.* We have proposed an interval-based coding of the permutation space for solving big permutation problems. The approach opens perspectives we are investigating [8], in terms of ultra-scale parallelization, application to multi-permutation problems and hybridization with metaheuristics.

¹A solution x dominates another solution y if x is better than y for all objectives and there exists at least one objective for which x is strictly better than y .

²The Pareto Front is the set of non-dominated solutions.

The *decomposition in the objective space* consists in breaking down an original many-objective problem (MaOP) into a set of cooperative single-objective subproblems (SOPs). The decomposition strategy requires the careful definition of a scalarizing (aggregation) function and its weighting vectors (each of them corresponds to a separate SOP) to guide the search process towards the best regions. Several scalarizing functions have been proposed in the literature including weighted sum, weighted Tchebycheff, vector angle distance scaling, etc. These functions are widely used but they have their limitations. For instance, using weighted Tchebycheff might do harm diversity maintenance and weighted sum is inefficient when it comes to deal with nonconvex Pareto Fronts [48]. Defining a scalarizing function well-suited to the MaOP at hand is therefore a difficult and still an open question being investigated in BONUS [5, 7]. Studying/defining various functions and in-depth analyzing them to better understand the differences between them is required. Regarding the weighting vectors that determine the search direction, their efficient setting is also a key and open issue. They dramatically affect in particular the diversity performance. Their setting rises two main issues: how to determine their number according to the available computational resources? when (statically or adaptively) and how to determine their values? *Weight adaptation is one of our main concerns that we are addressing especially from a distributed perspective.* They correspond to the main scientific objectives targeted by our bilateral ANR-RGC BigMO project with City University (Hong Kong). The other challenges pointed out in the beginning of this section concern the way to solve locally the SOPs resulting from the decomposition of a MaOP and the mechanism used for their cooperation to generate global solutions. To deal with these challenges, our approach is to design the decomposition strategy and cooperation mechanism keeping in mind the parallel and/or distributed solving of the SOPs. Indeed, we favor the local neighborhood-based mating selection and replacement to minimize the network communication cost while allowing an effective resolution [5]. The major issues here are how to define the neighborhood of a subproblem and how to cooperatively update the best-known solution of each subproblem and its neighbors.

To sum up, the objective of the BONUS team is to come up with scalable decomposition-based approaches in the decision and objective spaces. In the decision space, a particular focus will be put on high dimensionality and mixed-continuous variables which have received little interest in the literature. We will particularly continue to investigate at larger scales using ultra-scale computing the interval-based (discrete) and fractal-based (continuous) approaches. We will also deal with the rarely addressed challenge of mixed-continuous variables including categorical ones (collaboration with ONERA). In the objective space, we will investigate parallel ultra-scale decomposition-based many-objective optimization with ML-based adaptive building of scalarizing functions. A particular focus will be put on the state-of-the-art MOEA/D algorithm. This challenge is rarely addressed in the literature which motivated the collaboration with the designer of MOEA/D (bilateral ANR-RGC BigMO project with City University, Hong Kong). Finally, the joint decision-objective decomposition, which is still in its infancy [59], is another challenge of major interest.

3.2 Machine Learning-assisted Optimization

The Machine Learning (ML) approach based on metamodels (or surrogates) is commonly used, and also adopted in BONUS, to assist optimization in tackling BOPs characterized by time-demanding objective functions. The second line of research of BONUS is focused on ML-aided optimization to raise the challenge of expensive functions of BOPs using surrogates but also to assist the two other research lines (decomposition-based and ultra-scale optimization) in dealing with the other challenges (high dimensionality and scalability).

Several issues have been identified to make efficient and effective surrogate-assisted optimization. First, infill criteria have to be carefully defined to adaptively select the adequate sample points (in terms of surrogate precision and solution quality). The challenge is to find the best trade-off between exploration and exploitation to efficiently refine the surrogate and guide the optimization process toward the best solutions. The most popular infill criterion is probably the *Expected Improvement* (EI) [52] which is based on the expected values of sample points but also and importantly on their variance. This latter is inherently determined in the kriging model, this is why it is used in the state-of-the-art *efficient global optimization* (EGO) algorithm [52]. However, such crucial information is not provided in all surrogate models (e.g. Artificial Neural Networks) and needs to be derived. In BONUS, we are currently investigating

this issue. Second, it is known that surrogates allow one to reduce the computational burden for solving BOPs with time-consuming function(s). However, using parallel computing as a complementary way is often recommended and cited as a perspective in the conclusions of related publications. Nevertheless, *despite being of critical importance parallel surrogate-assisted optimization is weakly addressed in the literature*. For instance, in the introduction of the survey proposed in [51] it is warned that because the area is not mature yet the paper is more focused on the potential of the surveyed approaches than on their relative efficiency. *Parallel computing is required at different levels that we are investigating*.

Another issue with surrogate-assisted optimization is related to high dimensionality in decision as well as in objective space: it is often applied to low-dimensional problems. *The joint use of decomposition, surrogates and massive parallelism is an efficient approach to deal with high dimensionality. This approach adopted in BONUS has received little effort in the literature*. In BONUS, we are considering a generic framework in order to enable a flexible coupling of existing surrogate models within the state-of-the-art decomposition-based algorithm MOEA/D. This is a first step in leveraging the applicability of efficient global optimization into the multi-objective setting through parallel decomposition. Another issue which is a consequence of high dimensionality is the mixed (discrete-continuous) nature of decision variables which is frequent in real-world applications (e.g. engineering design). *While surrogate-assisted optimization is widely applied in the continuous setting it is rarely addressed in the literature in the discrete-continuous framework*. In [53], we have identified different ways to deal with this issue that we are investigating. Non-stationary functions frequent in real-world applications (see Section 4.1) is another major issue we are addressing using the concept of deep Gaussian Processes.

Finally, as quoted in the beginning of this section, ML-assisted optimization is mainly used to deal with BOPs with expensive functions but it will also be investigated for other optimization tasks. Indeed, ML will be useful to assist the decomposition process. In the decision space, it will help to perform the separability analysis (understanding of the interactions between variables) to decompose the vector of variables. In the objective space, ML will be useful to assist a decomposition-based many-objective algorithm in dynamically selecting a scalarizing function or updating the weighting vectors according to their performances in the previous steps of the optimization process [5]. Such a data-driven ML methodology would allow us to understand what makes a problem difficult or an optimization approach efficient, to predict the algorithm performance [4], to select the most appropriate algorithm configuration [9], and to adapt and improve the algorithm design for unknown optimization domains and instances. Such an autonomous optimization approach would adaptively adjust its internal mechanisms in order to tackle cross-domain BOPs.

In a nutshell, to deal with expensive optimization the BONUS team will investigate the surrogate-based ML approach with the objective to efficiently integrate surrogates in the optimization process. The focus will especially be put on high dimensionality (e.g. using decomposition) with mixed discrete-continuous variables which is rarely investigated. The kriging metamodel (Gaussian Process or GP) will be considered in particular for engineering design (for more reliability) addressing the above issues and other major ones including mainly non stationarity (using emerging deep GP) and ultra-scale parallelization (highly needed by the community). Indeed, a lot of work has been reported on deep neural networks (deep learning) surrogates but not on the others including (deep) GP. On the other hand, ML will be used to assist decomposition: importance/interaction between variables in the decision space, dynamic building (selection of scalarizing functions, weight update, etc.) of scalarizing functions in the objective space, etc.

3.3 Ultra-scale Optimization

The third line of our research program that accentuates our difference from other (project-)teams of the related Inria scientific theme is the ultra-scale optimization. *This research line is complementary to the two others, which are sources of massive parallelism* and with which it should be combined to solve BOPs. Indeed, ultra-scale computing is necessary for the effective resolution of the large amount of subproblems generated by decomposition of BOPs, parallel evaluation of simulation-based fitness and metamodels, etc. These sources of parallelism are attractive for solving BOPs and are natural candidates for ultra-scale supercomputers³. However, their efficient use raises a big challenge consisting in managing

³In the context of BONUS, supercomputers are composed of several massively parallel processing nodes (inter-node parallelism) including multi-core processors and GPUs (intra-node parallelism).

efficiently a massive amount of irregular tasks on supercomputers with multiple levels of parallelism and heterogeneous computing resources (GPU, multi-core CPU with various architectures) and networks. Raising such challenge requires to tackle three major issues: scalability, heterogeneity and fault-tolerance, discussed in the following.

The *scalability* issue requires, on the one hand, the definition of scalable data structures for efficient storage and management of the tremendous amount of subproblems generated by decomposition [55]. On the other hand, achieving extreme scalability requires also the optimization of communications (in number of messages, their size and scope) especially at the inter-node level. For that, we target the design of asynchronous locality-aware algorithms as we did in [49, 58]. In addition, efficient mechanisms are needed for granularity management and coding of the work units stored and communicated during the resolution process.

Heterogeneity means harnessing various resources including multi-core processors within different architectures and GPU devices. The challenge is therefore to design and implement hybrid optimization algorithms taking into account the difference in computational power between the various resources as well as the resource-specific issues. On the one hand, to deal with the heterogeneity in terms of computational power, we adopt in BONUS the dynamic load balancing approach based on the Work Stealing (WS) asynchronous paradigm⁴ at the inter-node as well as at the intra-node level. We have already investigated such approach, with various victim selection and work sharing strategies in [58], [8]. On the other hand, hardware resource specific-level optimization mechanisms are required to deal with related issues such as thread divergence and memory optimization on GPU, data sharing and synchronization, cache locality, and vectorization on multi-core processors, etc. These issues have been considered separately in the literature including our works [10]. Actually, in most of existing works related to GPU-accelerated optimization only a single CPU core is used. This leads to a huge resource wasting especially with the increase of the number of processing cores integrated into modern processors. Using jointly the two components raises additional issues including data and work partitioning, the optimization of CPU-GPU data transfers, etc.

Another issue the scalability induces is the *increasing probability of failures* in modern supercomputers [56]. Indeed, with the increase of their size to millions of processing cores their Mean-Time Between Failures (MTBF) tends to be shorter and shorter [54]. Failures may have different sources including hardware and software faults, silent errors, etc. In our context, we consider failures leading to the loss of work unit(s) being processed by some thread(s) during the resolution process. The major issue, which is particularly critical in exact optimization, is how to recover the failed work units to ensure a reliable execution. Such issue is tackled in the literature using different approaches: algorithm-based fault tolerance, checkpoint/restart (CR), message logging and redundancy. The CR approach can be system-level, library/user-level or application-level. Thanks to its efficiency in terms of memory footprint, adopted in BONUS [2], the application-level approach is commonly and widely used in the literature. This approach raises several issues mainly: (1) which critical information defines the state of the work units and allows to resume properly their execution? (2) when, where and how (using which data structures) to store it efficiently? (3) how to deal with the two other issues: scalability and heterogeneity?

The last but not least major issue which is another roadblock to exascale is the programming of massive-scale applications for modern supercomputers. *On the path to exascale, we will investigate the programming environments and execution supports able to deal with exascale challenges: large numbers of threads, heterogeneous resources, etc.* Various exascale programming approaches are being investigated by the parallel computing community and HPC builders: extending existing programming languages (e.g. DSL-C++) and environments/libraries (MPI+X, etc.), proposing new solutions including mainly Partitioned Global Address Space (PGAS)-based environments (Chapel, UPC, X10, etc.). It is worth noting here that our objective is not to develop a programming environment nor a runtime support for exascale computing. Instead, we aim to collaborate with the research teams (inside or outside Inria) having such objective.

To sum up, we put the focus on the design and implementation of efficient big optimization algorithms dealing jointly (uncommon in parallel optimization) with the major issues of ultra-scale computing

⁴A WS mechanism is mainly defined by two components: a victim selection strategy which selects the processing core to be stolen and a work sharing policy which determines the part and amount of the work unit to be given to the thief upon WS request.

mainly the scalability up to millions of cores using scalable data structures and asynchronous locality-aware work stealing, heterogeneity addressing the multi-core and GPU-specific issues and those related to their combination, and scalable GPU-aware fault tolerance. A strong effort will be devoted to this latter challenge, for the first time to the best of our knowledge, using application-level checkpoint/restart approach to deal with failures.

4 Application domains

4.1 Introduction

To validate the designed techniques, use standard benchmarks to facilitate the comparison with related works. In addition, we also target real-world applications in the context of our collaborations and industrial contracts. From the *application* point of view two classes are targeted: *complex scheduling* and *engineering design*. The objective is twofold: proposing new models for complex problems and solving efficiently BOPs using jointly the three lines of our research program. In the following, are given some use cases that are the focus of our current industrial collaborations.

4.2 Big optimization for complex scheduling

Three application domains are targeted: energy and transport & logistics. In the **energy** field, with the smart grid revolution (multi-)house energy management is gaining a growing interest. optimize the multi-house energy consumption taking into account (different designs of) the energy market

The key challenge is to optimize the multi-house energy consumption taking into account (different designs of) the energy market. *This kind of demand-side management will be of strategic importance for energy companies in the near future.* In collaboration with the EDF energy company we are working on the formulation and solving of optimization problems on demand-side management in smart micro-grids for single- and multi-user frameworks. These complex problems require taking into account multiple conflicting objectives and constraints and many (deterministic/uncertain, discrete/continuous) parameters. A representative example of such BOPs that we are addressing is the scheduling of the activation of a large number of electrical and thermal appliances for a set of homes optimizing at least three criteria: maximizing the user's comfort, minimizing its energy bill and minimizing peak consumption situations. On the other hand, we investigate the application of parallel Bayesian optimization for efficient energy storage in collaboration with the energy engineering department of University of Mons.

4.3 Big optimization for engineering design

The focus is for now put on the aerospace vehicle design, a complex multidisciplinary optimization process, we are exploring in collaboration with ONERA. The objective is to find the vehicle architecture and characteristics that provide the optimal performance (flight performance, safety, reliability, cost etc.) while satisfying design requirements [47]. A representative topic we are investigating, and will continue to investigate throughout the lifetime of the project given its complexity, is the design of launch vehicles that involves at least four tightly coupled disciplines (aerodynamics, structure, propulsion and trajectory). Each discipline may rely on time-demanding simulations such as Finite Element analyses (structure) and Computational Fluid Dynamics analyses (aerodynamics). Surrogate-assisted optimization is highly required to reduce the time complexity. In addition, the problem is high-dimensional (dozens of parameters and more than three objectives) requiring different decomposition schemas (coupling *vs.* local variables, continuous *vs.* discrete even categorical variables, scalarization of the objectives). Another major issue arising in this area is the non-stationarity of the objective functions which is generally due to the abrupt change of a physical property that often occurs in the design of launch vehicles. In the same spirit than deep learning using neural networks, we use Deep Gaussian Processes (DGPs) to deal with non-stationary multi-objective functions. Finally, the resolution of the problem using only one objective takes one week using a multi-core processor. The first way to deal with the computational burden is to investigate multi-fidelity using DGPs to combine efficiently multiple fidelity models. This approach has been investigated this year within the context of the PhD thesis of A. Hebbal. *In addition, ultra-scale computing is required at different levels to speed up the search and improve the reliability which is a major*

requirement in aerospace design. This example shows that we need to use the synergy between the three lines of our research program to tackle such BOPs.

Finally, we recently started to investigate the application of surrogate-based optimization in the epidemiologic context. Actually, we address in collaboration with Monash University (Australia) the contact reduction and vaccines allocation of Covid-19 and Tuberculosis.

5 Social and environmental responsibility

Optimization is ubiquitous to countless modern engineering and scientific applications with a deep impact on society and human beings. As such, the research of the BONUS team contributes to the establishment of high-level efficient solving techniques, improving solving quality, and addressing applications being more and more large-scale, complex, and beyond the solving ability of standard optimization techniques.

Furthermore, BONUS has performed technology transfer actions using different ways: open-source software development, transfer-to-industry initiatives, and teaching.

Our team has also initiated a start-up creation project. Specifically, G. Pruvost who did his Ph.D thesis within BONUS (defended on Dec. 2021), co-founded the OPTIMO Technologies start-up (2021-2023) with the support of Inria Startup Studio, dealing with sustainable mobility issues (e.g. sustainable, personalized and optimized itinerary planning). Although the startup could not continue due to a lack of necessary fundin, it demonstrates the impactful potential of our team and the significant value our research can generate for both the economic and social environment.

6 New software, platforms, open data

6.1 New software

6.1.1 pBB

Name: Permutation Branch-and-Bound

Keywords: Optimisation, Parallel computing, Data parallelism, GPU, Scheduling, Combinatorics, Distributed computing

Functional Description: The algorithm proceeds by implicit enumeration of the search space by parallel exploration of a highly irregular search tree. pBB contains implementations for single-core, multi-core, GPU and heterogeneous distributed platforms. Thanks to its hierarchical work-stealing mechanism, required to deal with the strong irregularity of the search tree, pBB is highly scalable. Scalability with over 90% parallel efficiency on several hundreds of GPUs has been demonstrated on the Jean Zay supercomputer located at IDRIS.

URL: <https://gitlab.inria.fr/jgmyp/permutationbb>

Publication: [hal-03689608](https://hal.archives-ouvertes.fr/hal-03689608)

Contact: Nouredine Melab

Participants: Jan Gmys, Nouredine Melab, Imen Chakroun, Mohand Mezmaz, Rudi Leroy

6.1.2 ParadisEO

Keyword: Parallelisation

Scientific Description: ParadisEO (PARallel and DIStributed Evolving Objects) is a C++ white-box object-oriented framework dedicated to the flexible design of metaheuristics. Based on EO, a template-based ANSI-C++ compliant evolutionary computation library, it is composed of four modules: * Paradiseo-EO provides tools for the development of population-based metaheuristic (Genetic algorithm, Genetic programming, Particle Swarm Optimization (PSO)...) * Paradiseo-MO provides

tools for the development of single solution-based metaheuristics (Hill-Climbing, Tabu Search, Simulated annealing, Iterative Local Search (ILS), Incremental evaluation, partial neighborhood...)

* Paradiseo-MOEO provides tools for the design of Multi-objective metaheuristics (MO fitness assignment schemes, MO diversity assignment schemes, Elitism, Performance metrics, Easy-to-use standard evolutionary algorithms...)

* Paradiseo-PEO provides tools for the design of parallel and distributed metaheuristics (Parallel evaluation, Parallel evaluation function, Island model)

Furthermore, Paradiseo also introduces tools for the design of distributed, hybrid and cooperative models:

* High level hybrid metaheuristics: coevolutionary and relay model

* Low level hybrid metaheuristics: coevolutionary and relay model

Functional Description: Paradiseo is a software framework for metaheuristics (optimisation algorithms aimed at solving difficult optimisation problems). It facilitates the use, development and comparison of classic, multi-objective, parallel or hybrid metaheuristics.

URL: <https://gitlab.inria.fr/paradiseo/paradiseo>

Contact: El-Ghazali Talbi

Partners: CNRS, Université de Lille

6.1.3 pyparadiseo

Keywords: Optimisation, Framework

Functional Description: pyparadiseo is a Python version of Paradiseo, a C++-based open-source white-box framework dedicated to the reusable design of metaheuristics. It allows the design and implementation of single-solution and population-based metaheuristics for mono- and multi-objective, continuous, discrete and mixed optimization problems.

URL: <https://gitlab.inria.fr/paradiseo/pyparadiseo>

Contact: Nouredine Melab

Participant: Jan Gmys

6.1.4 pySBO

Name: Python library for Surrogate-Based Optimization

Keywords: Parallel computing, Evolutionary Algorithms, Multi-objective optimisation, Black-box optimization, Optimisation

Functional Description: The pySBO library aims at facilitating the implementation of parallel surrogate-based optimization algorithms. pySBO provides re-usable algorithmic components (surrogate models, evolution controls, infill criteria, evolutionary operators) as well as the foundations to ensure the components inter-changeability. Actual implementations of sequential and parallel surrogate-based optimization algorithms are supplied as ready-to-use tools to handle expensive single- and multi-objective problems. The illustrated documentation of pySBO is available on-line through a dedicated web-site.

URL: <https://pysbo.readthedocs.io/en/latest/>

Publication: tel-03853862

Contact: Nouredine Melab

Participants: Guillaume Briffoteaux, Pierre Tomenko, François Géréme

6.1.5 moead-framework

Name: multi-objective evolutionary optimization based on decomposition framework

Keywords: Evolutionary Algorithms, Multi-objective optimisation

Scientific Description: Moead-framework aims to provide a python modular framework for scientists and researchers interested in experimenting with decomposition-based multi-objective optimization. The original multi-objective problem is decomposed into a number of single-objective sub-problems that are optimized simultaneously and cooperatively. This Python-based library provides re-usable algorithm components together with the state-of-the-art multi-objective evolutionary algorithm based on decomposition MOEA/D and some of its numerous variants.

Functional Description: The package is based on a modular architecture that makes it easy to add, update, or experiment with decomposition components, and to customize how components actually interact with each other. A documentation is available online. It contains a complete example, a detailed description of all available components, and two tutorials for the user to experiment with his/her own optimization problem and to implement his/her own algorithm variants.

URL: <https://github.com/moead-framework>

Publication: hal-03818749

Contact: Geoffrey Pruvost

Participants: Geoffrey Pruvost, Bilel Derbel, Arnaud Liefoghe

6.1.6 Zellij

Keywords: Global optimization, Partitioning, Metaheuristics, High Dimensional Data

Functional Description: The package generalizes a family of decomposition algorithms by implementing four distinct modules (geometrical objects, tree search algorithms, exploitation and exploration algorithms such as Genetic Algorithm, Bayesian Optimization or Simulated Annealing). The package is divided into two versions, a regular and a parallel one. The main target of Zellij is to tackle HyperParameter Optimization (HPO) and Neural Architecture Search (NAS). Thanks to to this framework, we are able to reproduce various decomposition based algorithms, such as DIRECT, Simultaneous Optimistic Optimization, Fractal Decomposition Algorithm, FRACTOP.. Future works will focus on multi-objective problems, NAS, distributed version and a graphic interface for monitoring and plotting.

URL: <https://github.com/ThomasFirmin/zellij>

Contact: Thomas Firmin

6.2 New platforms

6.2.1 SLICES-FR/GRID'5000 testbed: major achievements in 2024

Participants: Bilel Derbel (*contact person*), Hugo Dominois.

- **Keywords:** Experimental testbed, large-scale computing, high-performance computing, GPU computing, cloud computing, big data

- **Functional description:** Grid'5000 is a project initiated in 2003 by the French government and later supported by different research organizations including Inria, CNRS, the french universities, Renater which provides the wide-area network, etc. The overall objective of Grid'5000 was to build by 2007 a mutualized nation-wide experimental testbed composed of at least 5000 processing units and distributed over several sites in France (one of them located at Lille). From a scientific point of view, the aim was to promote scientific research on large-scale distributed systems. Beyond BONUS, Grid'5000 is highly important for the HPC-related communities from our three institutions (ULille, Inria and CNRS) as well as from outside.

Within the framework of CPER contract "Data", the equipment of Grid'5000 at Lille has been renewed in 2017-2018 in terms of hardware resources (GPU-powered servers, storage, PDUs, etc.) and infrastructure (network, inverter, etc.). The renewed testbed has been used extensively by many researchers from Inria and outside. Half-day trainings have been organized with the collaboration of BONUS to allow the newcomer users to get started with the use of the testbed. A new IA-dedicated CPER contract "CornellIA" has been accepted (2021-2027).

Since late 2023, B. Derbel took over as the scientific leader of N. Melab. More importantly, GRID'5000 has evolved to merge with the FIT platform in order to evolve towards the SLICES-FR European experimental infrastructure. As such, B. Derbel is the site leader of the Lille site at SLICES-FR. He is strongly involved in the site leader committee, as well as on the managing aspects of the SLICES-FR site in Lille. During 2024, two clusters have been renewed and are now available for the SLICES-FR users.

- URL: [Grid'5000/SLICES-FR](#)

7 New results

During the year 2024, we have addressed different issues/challenges related to the three lines of our research program. The major contributions are summarized in the following sections. Besides, alongside these contributions we came out with other contributions [44, 26, 17, 13], including general-purpose surveys and taxonomies on advanced search techniques, that are not discussed here-after to keep the presentation more focused.

7.1 Decomposition-based optimization

We report five major contributions related to decomposition-based and multi-objective optimization. The first contribution [42] concerns the investigation of multi-objective decomposition in the context of reinforcement learning. The second one [15] concerns the integration of decomposition based techniques to leverage a well-known quantum optimization algorithm for multi-objective optimization problems. The third contribution concerns the design and analysis of new decomposition-based features for multi-objective landscape analysis when facing objective heterogeneity. The last contributions are focused on solving difficult and high-dimensional problems coming from complex application domains, respectively in Spacecraft Optimal Layout [21], Arc routing [25] and Federated Learning [24]. These contributions are discussed in more details in the following.

7.1.1 Multi-Objective Reinforcement Learning Based on Decomposition

Participants: Florian Felten (*SnT, University of Luxembourg*), El-Ghazali Talbi (*contact person*), Grégoire Danoy (*SnT, University of Luxembourg*).

Multi-objective reinforcement learning (MORL) extends traditional RL by seeking policies making different compromises among conflicting objectives. The recent surge of interest in MORL has led to diverse studies and solving methods, often drawing from existing knowledge in multi-objective optimization based on decomposition (MOO/D). Yet, a clear categorization based on both RL and MOO/D is lacking

in the existing literature. Consequently, MORL researchers face difficulties when trying to classify contributions within a broader context due to the absence of a standardized taxonomy. To tackle such an issue, we introduce in [20] multi-objective reinforcement learning based on decomposition (MORL/D), a novel methodology bridging the literature of RL and MOO. A comprehensive taxonomy for MORL/D is presented, providing a structured foundation for categorizing existing and potential MORL works. The introduced taxonomy is then used to scrutinize MORL research, enhancing clarity and conciseness through well-defined categorization. Moreover, a flexible framework derived from the taxonomy is introduced. This framework accommodates diverse instantiations using tools from both RL and MOO/D. Its versatility is demonstrated by implementing it in different configurations and assessing it on contrasting benchmark problems. Results indicate MORL/D instantiations achieve comparable performance to current state-of-the-art approaches on the studied problems. By presenting the taxonomy and framework, we offer a comprehensive perspective and a unified vocabulary for MORL. This not only facilitates the identification of algorithmic contributions but also lays the groundwork for novel research avenues in MORL.

7.1.2 Multi-objective Quantum Approximate Optimiser based on decomposition

Participants: Zakaria Abdelmoiz Dahi (*contact person*), Francisco Chicano (*Univ. Malaga, Spain*), Gabriel Luque (*Univ. Malaga, Spain*), Bilel Derbel (*contact person*), Enrique Alba (*Univ. Malaga, Spain*).

Quantum computation uses quantum mechanical principles to reach beyond-classical computational power. This has endless applications, especially in optimisation-problems' solving. Most of today's quantum optimisers, more specifically, Quantum Approximate Optimisation Algorithm (QAOA), were originally designed to solve single-objective problems, although real-life scenarios include generally dealing with multiple objectives. Very preliminary literature with design/implementation limitations has been done in this sense. This makes dealing with such limitations and expanding the QAOA applicability to multi-objective optimisation an important step towards advancing quantum computation. To do so, we present in [29] a decomposition-based Multi-Objective QAOA (MO-QAOA) able to solve multi-objective problems. The proposal's design explores QAOA's features considering the error-prone and limited nature of today's quantum computers as well as the costly quantum simulation. This work's contributions stand in designing both, (I) sequential and parallel MO-QAOA, based on (II) weighted-sum and Tchebycheff scalarisation, by (III) exploring the QAOA's parameters' transference. The validation has been done using 2, 3 and 4-objectives problems of several sizes/complexities/types, using up to 2000 slaves/jobs running quantum computer simulators, as well as three real IBM 127-qubits' quantum computers. The results show up to 89% execution-time decrease, which supports the applicability/reliability of the proposal in today's time-constrained and error-prone quantum computers.

7.1.3 Landscape analysis of heterogeneous multi-objective problems using decomposition

Participants: Raphaël Cosson, Roberto Santana (*University of Basque Country*), Bilel Derbel (*contact person*), Arnaud Liefoghe.

The heterogeneity among objectives in multi-objective optimization can be viewed from several perspectives. In [16], we are interested in the heterogeneity arising in the underlying landscape of the objective functions, in terms of multi-modality and search difficulty. Building on recent efforts leveraging the so-called single-objective NK-landscapes to model such a setting, we conduct a three-fold empirical analysis on the impact of objective heterogeneity on the landscape properties and search difficulty of bi-objective optimization problems. Firstly, for small problems, we propose two techniques based decomposition with the aim of studying the distribution of the solutions in the objective space. Secondly, for large problems, we investigate the ability of existing landscape features to capture the degree of heterogeneity among the two objectives. Thirdly, we study the behavior of two state-of-the-art multi-objective evolutionary algorithms, namely MOEA/D and NSGA-II, when faced with a range of problems with different degrees

of heterogeneity. Although one algorithm is found to consistently outperform the other, the dynamics of both algorithms vary similarly with respect to objective heterogeneity. Our analysis suggests that novel approaches are needed to understand the fundamental properties of heterogeneous bi-objective optimization problems and to tackle them more effectively.

7.1.4 Two-Level Approach for Simultaneous Component Assignment and Layout Optimization with Applications to Spacecraft Optimal Layout

Participants: Juliette Gamot (*contact person*), Mathieu Balesdent (*ONERA DTIS, Palaiseau*), Romain Wuilbercq (*ONERA DTIS, Palaiseau*), Arnault Tremolet (*ONERA DTIS, Palaiseau*), Nouredine Melab.

In [21], we investigated decomposition in the context of component assignment and layout optimization. Optimal layout problems consist in positioning a given number of components in order to minimize an objective function while satisfying geometrical or functional constraints. Such kinds of problems appear in the design process of aerospace systems such as satellite or spacecraft design. These problems are NP-hard, highly constrained and dimensional. This paper describes a two-stage algorithm combining a genetic algorithm and a quasi-physical approach based on a virtual-force system in order to solve multi-container optimal layout problems such as satellite modules. In the proposed approach, a genetic algorithm assigns the components to the containers while a quasi-physical algorithm based on a virtual-force system is developed for positioning the components in the assigned containers. The proposed algorithm is experimented and validated on the satellite module layout problem benchmark. Its global performance is compared with previous algorithms from the literature.

7.1.5 Multi-objective optimization for complex and high-dimensional problems

Participants: Daniel Porumbel (*CEDRIC CS Laboratory, CNAM*), Igor Machado Coelho (*Universidade Federal Fluminense, Brazil*), El-Ghazali Talbi (*contact person*), José Á. Morell, Zakaria Abdelmoiz Dahi (*contact person*), Francisco Chicano (*University of Malaga*), Gabriel Luque (*University of Malaga*), Enrique Alba (*University of Malaga*).

In this section we describe two different contributions that have in common the use of multi-objective optimization techniques to tackle high dimensional complex problems.

In [25], we address the bi-objective Capacitated Arc Routing Problem (CARP) by considering two levels of solution interpretation: implicit and explicit solutions. An algorithm that translates implicit solutions into explicit solutions is called a decoder. The decoder takes as input a permutation of the required edges and generates a Pareto frontier of CARP solutions. While bi-objective CARP was our main focus and starting point, we could also use the proposed framework to solve a bi-objective version of the traveling salesman problem by plugging-in a different decoder. The bi-objective CARP asks to service (the demands of) a set of required edges using a fleet of vehicles of limited capacity so as to minimize: (i) the total travelled distance and (ii) the length of the longest route. Any permutation of the required edges constitutes an implicit CARP solution. The decoder constructs all non-dominated explicit solutions that service the edges in the order indicated by, i.e., the decoder is an exact algorithm that returns the optimal Pareto frontier subject to the service order. To achieve competitive CARP results it is also important to reinforce the decoder using a local search operator that acts on explicit routes (and that may change the service order). The resulting algorithm was even able to find a new total-cost upper bound, improving upon the best solutions reported in the (considerably larger) mono-objective CARP literature. This shows that (some of) the proposed ideas can also be useful in single objective optimization, since the second objective can be seen as a guide for the mono-objective search process.

In [24], we consider the application of multi-objective optimization techniques to tackle a seemingly different problem arising in Federated learning. Federated Learning is a paradigm that proposes pro-

protecting data privacy by sharing local models instead of raw data during each iteration of model training. However, these models can be large, with many parameters, provoking a substantial communication cost and having a notable environmental impact. Reducing communication overhead is paramount but conflictual to maintaining the model's accuracy. In our work, we explore the add-in that multi-objective evolutionary algorithms can provide for solving the communication overhead problem while achieving high accuracy. We do this by 1) realistically modelling and formulating this task as a multi-objective problem by considering the devices' heterogeneity, 2) including all the communication-triggering aspects, and 3) applying a multi-objective evolutionary algorithm with an intensification operator to solve the problem. A simulated client-server architecture of four devices with four different processing speeds is studied. Both fully connected and convolutional neural network models are investigated with 33,400 and 887,530 weights, respectively.

7.2 ML-assisted optimization

In this axis, we describe our contributions on ML-assisted optimization techniques following three directions: (1) the optimization of deep neural architectures and hyperparameters, (2) the efficient building of surrogates and their integration into optimization algorithms to deal with expensive black-box objective functions, and (3) the synergy between machine learning and quantum optimization algorithms. Our contributions in each direction are discussed in more details in the following.

7.2.1 Neural Architecture Search and Hyperparameter Optimization

Participants: Julie Keisler (*contact person*), Housseem Ouertatani (*contact person*), Co-sijopii Garcia-Garcia (*contact person*), El-Ghazali Talbi (*contact person*), Bilel Derbel (*contact person*), Sandra Claudel (*EDF Lab Paris-Saclay*), Gilles Cabriel (*EDF Lab Paris-Saclay*), Cristian Maxim (*IRT SystemX*), Smaïl Niar (*UPHF, Valenciennes*), Alicia Morales-Reyes (*INAOE, Mexico*), Hugo Jair Escalante (*INAOE, Mexico*).

Deep neural networks (DNNs), particularly convolutional neural networks (CNNs), have garnered significant attention in recent years for addressing a wide range of challenges in image processing and computer vision. Neural architecture search (NAS) and hyperparameter optimization (HPO) have emerged as crucial fields aiming to automate the design and configuration of CNN models. IN this respect, we contributed the following.

In [23], we propose DRAGON (for DiRected Acyclic Graph OptimizationN), an algorithmic framework to automatically generate efficient deep neural networks architectures and optimize their associated hyperparameters. The framework is based on evolving Directed Acyclic Graphs (DAGs), defining a more flexible search space than the existing ones in the literature. It allows mixtures of different classical operations: convolutions, recurrences and dense layers, but also more newfangled operations such as self-attention. Based on this search space we propose neighbourhood and evolution search operators to optimize both the architecture and hyper-parameters of our networks. These search operators can be used with any metaheuristic capable of handling mixed search spaces. We tested our algorithmic framework with an asynchronous evolutionary algorithm on a time series forecasting benchmark. The results demonstrate that DRAGON outperforms state-of-the-art handcrafted models and AutoML techniques for time series forecasting on numerous datasets. DRAGON has been implemented as a python open-source package.

In [36, 37], we designed a fast NAS method based on Bayesian optimization (BO). BO is a black-box search method particularly valued for its sample efficiency. It is especially effective when evaluations are very costly, such as in hyperparameter optimization or Neural Architecture Search (NAS). While Gaussian Processes underpin most BO approaches, we instead use deep ensembles. This allows us to construct a unified and improved representation, leveraging pretraining metrics and multiple evaluation fidelities, to accelerate the search. More specifically, we use a simultaneous pretraining scheme where multiple metrics are estimated concurrently. Consequently, a more general representation is obtained. A novel multi-fidelity approach is proposed, where the unified representation is improved both by high and low

quality evaluations. These additions significantly accelerate the search time, finding the optimum on NAS-Bench-201 in an equivalent time and cost to performing as few as 50 to 80 evaluations. The accelerated search time translates to reduced costs, in terms of computing resources and energy consumption. As a result, applying this NAS method to real-world use cases becomes more practical and not prohibitively expensive. We demonstrate the effectiveness and generality of our approach on a custom search space. Based on the MOAT architecture, we designed a search space of CNN-ViT hybrid networks. The search method yields a better-performing architecture than the baseline in only 70 evaluations.

Finally, in [32], we propose a novel strategy to speed up the performance estimation of neural architectures by gradually increasing the size of the training set used for evaluation as the search progresses. We evaluate this approach using the CGP-NASV2 model, a multi-objective NAS method, on the CIFAR-100 dataset. Experimental results demonstrate a notable acceleration in the search process, achieving a speedup of 4.6 times compared to the baseline. Despite using limited data in the early stages, our proposed method effectively guides the search towards competitive architectures. This study highlights the efficacy of leveraging lower-fidelity estimates in NAS and paves the way for further research into accelerating the design of efficient CNN architectures.

7.2.2 Hybrid surrogate-based optimization

Participants: Guillaume Briffoteaux (*contact person*), Nouredine Melab (*contact person*), Mohand Mezmaç (*University of Mons*), Daniel Tuyttens (*University of Mons*), Jan Gmys (*Univ. Lille*).

Surrogate models are built to produce computationally efficient versions of time-complex simulation-based objective functions so as to address expensive optimization. In surrogate-assisted evolutionary algorithms (SAEA), the surrogate model evaluates and/or filters candidate solutions produced by evolutionary operators. In surrogate-driven optimization (SDO), the surrogate is used to define the objective function of an auxiliary optimization problem whose resolution generates new candidates. IN this context, we have contributed the following results.

In [14], hybridization of the before mentioned acquisition processes is investigated with a focus on robustness with respect to the computational budget and parallel scalability. A new hybrid method based on the successive use of acquisition processes during the search outperforms competing approaches regarding these two aspects on the Covid-19 contact mitigation problem. To further improve the generalization to larger ranges of search landscapes, another new hybrid method based on the dispersion metric is proposed. The integration of landscape analysis tools in surrogate-based optimization seems promising according to the numerical results reported on the CEC2015 test suite.

Moreover, Parallel Surrogate-Based Optimization (PSBO) is an efficient approach to deal with black-box time-consuming objective functions. According to the available computational budget to solve a given problem, we investigated and opposed in [22] three classes of algorithms: Surrogate-Assisted Evolutionary Algorithms (SAEAs), Bayesian Optimization Algorithms (BOAs), and Surrogate-free Evolutionary Algorithms (EAs). A large set of benchmark functions and engineering applications are considered with various computational budgets. As such, we come up with guidelines for the choice between the three categories. According to the computational expensiveness of the objective functions and the number of processing cores, we identify a threshold from which SAEAs should be preferred to BOAs. Based on this threshold, we derive a new hybrid Bayesian/Evolutionary algorithm that allows one to tackle a wide range of problems without prior knowledge of their characteristics.

Finally, we also investigated multi-fidelity surrogated-based optimization for scheduling of pumped hydro energy storage in [19] and multi-objective surrogate-Assisted optimization for the modeling of machining in [18].

7.2.3 Machine learning and quantum algorithms

Participants: Djaafar Zouache, Adel Got, Deemah Alarabiat, Laith Abualigah, El-Ghazali Talbi (*contact person*), Zakaria Abdelmoiz Dahi (*contact person*), Francisco Chicano, Gabriel Luque.

In this section, we describe two seemingly different contributions being the intersection of machine learning and quantum algorithms. In the first contribution [60], we are interested in using quantum-inspired algorithm to tackle the well-known feature selection problem. In the second one [28], we are interested in designing machine learning enhanced algorithm for quantum algorithm transpilation. These two contributions are described in more details in the following.

In [60], we present a novel algorithm for tackling the feature selection problem modeled as a multi-objective problem. Our approach draws inspiration from quantum computing and combines the strengths of the Firefly Algorithm (FA) and the Particle Swarm Optimizer (PSO). Leveraging quantum computing enhances solution distribution, while the cooperative nature of FA and PSO facilitates effective exploration of the feature space. Additionally, we introduce two fixed-size external archives, dedicated to storing the best solutions. The archive sizes are controlled using the epsilon dominance relation. We evaluate the efficiency of our algorithm through an extensive comparison against both single and multi-objective feature selection algorithms that enjoy high regard in the field. Furthermore, we propose a high-performance detection system that harnesses our algorithm alongside three Convolutional Neural Network Algorithms. This system demonstrates its potential in accurately identifying COVID-19 disease from X-ray images.

In [28], we propose an evolutionary deep neural network that learns the qubits' layout initialisation of the most advanced and complex IBM heuristic used in today's quantum machines. The aim is to progressively replace weakly scalable transpilation heuristics with machine learning models. Previous work using machine learning models for qubits' layout initialisation suffers from some shortcomings in the proposal's correctness and generalisation as well as benchmarks diversity, utility, and availability. Our work solves those flaws by (I) devising a complete Machine Learning pipeline including the ETL component and the evolutionary deep neural model using the linkage learning algorithm P3, (II) a modelling applicable to any quantum algorithm with a special interest to both optimisation and machine learning ones, (III) diverse and fresh benchmarks using calibration data of four real IBM quantum computers collected over 10 months (Dec. 2022 and Oct. 2023) and training dataset built using four types of quantum optimisation and machine learning algorithms, as well as random ones. The proposal has been proven to be more efficient and simple than state-of-the-art deep neural models in the literature.

7.3 Ultra-scale parallel optimization

During the year 2024, we have made contributions with respect to four main research directions in our parallel optimization axis: (1) ultra-scale Parallel tree-based exact algorithms, (2) parallel hyperparameter optimization of spiking neural networks, (3) large-scale parallel graybox optimization. Our contributions in each research direction are discussed in more details in the following.

7.3.1 Parallel tree-based search algorithms

Participants: Guillaume Helbecque (*contact person*), Nouredine Melab (*contact person*), Tiago Carneiro (*IMEC, Belgium*), Pascal Bouvry (*University of Luxembourg*), Engin Kayraklioglu, Jan Gmys (*Univ. Lille*), Ezhilmathi Krishnasamy.

As evidenced by the Top500 ranking, modern supercomputers continue to grow in scale and complexity, the design and implementation of efficient algorithms for BOPs have become increasingly challenging. One key aspect of this challenge lies in the development of effective data structures and workload distribution mechanisms that can operate efficiently across diverse computational environments. In this context, we came out with several contributions to parallel and distributed algorithms aimed at solving BOPs, focusing on tree-based search strategies. Specifically, we present a PGAS data structure, DistBag-DFS,

designed for unbalanced tree-based algorithms, and demonstrate its scalability and performance in both single-node and large-cluster settings. Additionally, we delve into GPU-accelerated optimization using Chapel, emphasizing the importance of portability across different hardware platforms while maintaining high performance in tree-search-based optimization tasks. These contributions are discussed in more details in the following paragraphs.

Firstly, the design and implementation of algorithms for increasingly large and complex modern supercomputers requires the definition of data structures and workload distribution mechanisms in a productive and scalable way. In [33, 45], we have proposed a PGAS data structure, termed DistBag-DFS, along with a Work-Stealing mechanism for the class of parallel tree-based algorithms that explore unbalanced trees using the depth-first search strategy. The contribution has been implemented and packaged as an open-source module in the Chapel PGAS language. The experimentation of the contribution in a single-node setting using backtracking applied to fine-grained Unbalanced Tree-Search (UTS) benchmark shows that 68% of the linear speed-up can be achieved. In addition, the scalability of the contribution has been evaluated using the Branch-and-Bound algorithm to solve big instances of the Flowshop Scheduling problem on a large cluster. The reported results reveal that 50% of strong scaling efficiency is achieved using 400 computer nodes (51,200 processing cores).

Secondly, in [35] and [34], we explored the issue of GPU vendor-agnosticism in the context of tree-based optimization. In [35, 39], we revisited the design and implementation of a generic, multi-pool GPU-accelerated tree-search algorithm using Chapel. This algorithm was instantiated using the backtracking method and tested on the N-Queens problem. For performance evaluation, the Chapel-based approach was compared to the low-level counterparts in Nvidia CUDA and AMD HIP. The results indicated that, in a single-GPU setup, Chapel's high-level GPU abstraction resulted in only an 8% (resp. 16%) performance loss compared to CUDA (resp. HIP). In a multi-GPU setting, up to 80% (resp. 71%) of the baseline speed-up was achieved for coarse-grained problem instances on Nvidia (resp. AMD) GPUs.

Thirdly, in [34, 46], we focused on the design and implementation of a multi-GPU Branch-and-Bound algorithm in Chapel using the DistBag-DFS at the intra-node level. While CPU cores were used for parallel tree exploration, GPU devices were leveraged to accelerate the bounding phase, which is particularly compute-intensive. Extensive experiments on the Permutation Flowshop Scheduling Problem demonstrated that the proposed approach achieved strong scaling efficiencies of up to 63% and an average of 75% when using GPU-powered processing nodes, including 8 NVIDIA A100 devices and AMD MI50 GPUs. These results highlight the efficiency of our approach in solving large combinatorial optimization problems, while ensuring portability across different hardware platforms.

Finally, In [27], we investigate the viability of using the Chapel high-productivity language as a tool to achieve both code and performance portability in large-scale tree-based search. As a case study, we implemented a distributed backtracking for solving permutation combinatorial problems. Extensive experiments conducted on big N-Queens problem instances, using up to 512 NVIDIA GPUs and 1024 AMD GPUs on Top500 supercomputers, reveal that it is possible to scale on the two different systems using the same tree-based search written in Chapel. This trade-off results in a performance decrease of less than 10% for the biggest problem instances.

7.3.2 Parallel hyperparameter optimization of spiking neural networks

Participants: Thomas Firmin (*contact person*), El-Ghazali Talbi, Pierre Boulet.

The work described here-after is tightly related to our second research axis on ML-assisted optimization. We in fact consider the design of Spiking Neural Network (SNN). SNNs are peculiar networks based on the dynamics of timed spikes between fully asynchronous neurons. Their design is complex and differs from usual artificial neural networks as they are highly sensitive to their hyperparameters. As such, Hyperparameter optimization of SNNs is a difficult task which has not yet been deeply investigated in the literature. In [42], we designed a scalable constrained Bayesian based optimization algorithm that prevents sampling in non-spiking areas of an efficient high dimensional search space. These search spaces contain infeasible solutions that output no or only a few spikes during the training or testing phases, we call such a mode a “silent network”. Finding them is difficult, as many hyperparameters

are highly correlated to the architecture and to the dataset. We leverage silent networks by designing a spike-based early stopping criterion to accelerate the optimization process of SNNs trained by spike timing dependent plasticity and surrogate gradient. We parallelized the optimization algorithm asynchronously, and ran large-scale experiments on heterogeneous multi-GPU Petascale architecture. Results show that by considering silent networks, we can design more flexible high-dimensional search spaces while maintaining a good efficacy. The optimization algorithm was able to focus on networks with high performances by preventing costly and worthless computation of silent networks.

In [31], to accelerate the hyperparameter optimization of SNNs trained by surrogate gradient, we propose to leverage silent networks and multi-fidelity. We designed an asynchronous black-box constrained and cost-aware Bayesian optimization algorithm to handle high-dimensional search spaces containing many silent networks, considered as infeasible solutions. Large-scale experimentation was computed on a multi-nodes and multi-GPUs environment. By considering the cost of evaluations, we were able to quickly obtain acceptable results for SNNs trained on a small proportion of the training dataset. We can rapidly stabilize the inherent high sensitivity of the SNNs' hyperparameters before computing expensive and more precise evaluations. We have extended our methodology for search spaces containing 21 and up to 46 layer-wise hyperparameters. Despite an increased difficulty due to the higher dimensional space, our results are competitive, even better, compared to their baseline. Finally, while up to 70% of sampled solutions were silent networks, their impact on the budget was less than 4%. The effect of silent networks on the available resources becomes almost negligible, allowing to define higher dimensional, more general and flexible search spaces.

7.3.3 Large-scale parallel graybox optimization

Participants: Lorenzo Canonne, Bilel Derbel (*contact person*), Miwako Tsuji, Mitsuhiro Sato.

In [15], we design, develop and analyze parallel variants of a state-of-the-art graybox optimization algorithm, namely Drils (Deterministic Recombination and Iterated Local Search), for attacking large-scale pseudo-boolean optimization problems on top of the large-scale computing facilities offered by the supercomputer Fugaku. We first adopt a Master/Worker design coupled with a fully distributed Island-based model, ending up with a number of hybrid OpenMP/MPI implementations of high-level parallel Drils versions. We show that such a design, although effective, can be substantially improved by enabling a more focused iteration-level cooperation mechanism between the core graybox components of the original serial Drils algorithm. Extensive experiments are conducted in order to provide a systematic analysis of the impact of the designed parallel algorithms on search behavior, and their ability to compute high-quality solutions using increasing number of CPU-cores. Results using up to 1024×12-cores NUMA nodes, and NK-landscapes with up to binary variables are reported, providing evidence on the relative strength of the designed hybrid cooperative graybox parallel search. This contribution is also tightly related to our first research axis since the heuristic search algorithms considered here are based on the concept of variable decomposition in the decision space.

8 Bilateral contracts and grants with industry

8.1 Bilateral grants with industry

Our current industrial granted projects are as follows.

- EDF (2021-2024, Paris): this joint project with EDF, a major electrical power player in France, targets the automatic design and configuration of deep neural networks applied to the energy consumption forecasting. A budget of 62K€ is initially allocated, in the context of the **PGMO** programme of Jacques Hadamard foundation of mathematics. A budget of 150K€ is then allocated for funding a PhD thesis (CIFRE).

- Confiance.ai project (2021-2024, Paris): this joint project with the **SystemX Institute of Research and Technology (IRT)** and Université Polytechnique Hauts-de-France is focused on multi-objective automated design and optimization of deep neural networks with applications to embedded systems. A Ph.D student (H. Ouertatani) has been hired in Oct. 2021 to work on this topic.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

AnyScale

Title: Parallel Fractal-based Chaotic optimization: Application to the optimization of deep neural networks for energy management

Duration: 2022 – 2025

Coordinator: El-Ghazali Talbi

Partners: • Ecole Mohammadia d'Ingénieurs Rabat (Maroc)

Inria contact: El-Ghazali Talbi

Summary: Many scientific and industrial disciplines are more and more concerned by big optimisation problems (BOPs). BOPs are characterised by a huge number of mixed decision variables and/or many expensive objective functions. Bridging the gap between computational intelligence, high performance computing and big optimisation is an important challenge for the next decade in solving complex problems in science and industry.

The goal of this associated team project is to come up with breakthrough in nature-inspired algorithms jointly based on any-scale fractal decomposition and chaotic approaches for BOPs. Those algorithms are massively parallel and can be efficiently designed and implemented on heterogeneous exascale supercomputers including millions of CPU and GPU (Graphics Processing Units) cores. The convergence between chaos, fractals and massively parallel computing will represent a novel computing paradigm for solving complex problems.

From the application and validation point of view, we target the automatic design of deep neural networks, applied to the prediction of the electrical energy consumption and production.

9.1.2 Participation in other International Programs

MoU RIKEN R-CCS / Japan

Participants: Bilel Derbel, David Redon.

Title: Memoremndum of Understanding

Partner Institution(s): • RIKEN Center of Computational Science, Japan

Date/Duration: 2021 – 2026

Additional info/keywords: This MoU aims at strengthening the research collaboration with one of the world-wide leading institute in HPC targeting the solving of computing-intensive optimization problems on top of the japanese Fugaku supercomputer facilities(ranked in TOP500).

9.2 International research visitors

9.2.1 Visits of international scientists

Other international visits to the team

- R. Ellaia (EMI, Univ. of Rabat, Morocco)
- E. Alba (University of Malaga, Spain)
- G. Danoy (Univ. Luxembourg, Luxembourg)
- J. Lopez Espin (Univ. of Elche, Spain)

9.2.2 Visits to international teams

Research stays abroad

Bilel Derbel

Visited institution: Shinshu University

Country: Japan

Dates: July 2024

Context of the visit: CIMO Workshop

Zakaria Abdelmoiz DAHI

Visited institution: Ludwig-Maximilians-Universität München

Country: Germany

Dates: June-September 2024

Context of the visit: Scientific collaboration on quantum algorithms

Zakaria Abdelmoiz DAHI

Visited institution: University of Malaga

Country: Spain

Dates: April, October, November 2024

Context of the visit: Scientific collaboration on evolutionary and quantum optimization

El-Ghazali Talbi

Visited institution: University of Luxembourg

Country: Luxembourg

Dates: October 2024

Context of the visit: Scientific collaboration and international networking

El-Ghazali Talbi**Visited institution:** Univ. Elche**Country:** Spain**Dates:** April 2024**Context of the visit:** Scientific collaboration and international networking**El-Ghazali Talbi****Visited institution:** Univ. Catania**Country:** Italy**Dates:** July 2024**Context of the visit:** Scientific collaboration and international networking**El-Ghazali Talbi****Visited institution:** ISG (Institut Supérieur de Gestion), Tunis**Country:** Tunisia**Dates:** March 2024**Context of the visit:** Scientific collaboration and international networking**El-Ghazali Talbi****Visited institution:** EMI Rabat**Country:** Morocco**Dates:** September 2024**Context of the visit:** Scientific collaboration and international networking**9.2.3 Other european programs/initiatives**

- ERC Generator "Exascale Parallel Nature-inspired Algorithms for Big Optimization Problems", supported by University of Lille call (2023-2025, Total: 99K€). The goal of this project is to come up with breakthrough in nature-inspired algorithms jointly based on fractal decomposition and chaotic optimization approaches for BOPs. Those algorithms are massively parallel and can be efficiently designed and implemented on heterogeneous exascale supercomputers including millions of CPU/GPU cores, and neuromorphic accelerators composed of billions of spiking neurons. E.-G. Talbi is the leader of this project.

9.3 National initiatives**9.3.1 ANR**

- Bilateral ANR-NSF France/USA PRCI TunnelOPT (2024-2027, Grant: 562K€, PI: B. Derbel) in collaboration with Colorado State University (Co-PI: D. Whitley).
New optimization algorithms developed over the last two decades can efficiently solve a wide-range of combinatorial optimization problems. Nevertheless, existing combinatorial optimization techniques still struggle to efficiently handle the unprecedented complexity of the problems encountered in modern engineering, scientific, and numerical applications. Often these problems are multi-objective; hence implying other degrees of difficulty. Achieving scalability is a major

concern, specifically with respect to the number of variables and the number of objectives; but also with respect to modern parallel and distributed resources, including massively parallel multi-core and multi-GPU based resources. In this France/USA bilateral ANR PRCI project, we focus on the design and the fundamental understanding of innovative stochastic heuristic search algorithms empowered by graybox optimization methods. In fact, new graybox formulations allow us to compute the eigenvectors of the search neighborhood for local search methods that apply to a range of fundamental combinatorial problems such as logical satisfiability (e.g. MAXkSAT) and routing (e.g. the Travelling Salesman Problem). Furthermore, it becomes possible to tunnel between local optima in linear time. By describing how local optima (Pareto or not) are organized into regular hypercube subspaces that form non-planar lattices; we propose to set up the foundations of a tunneling engine to navigate in parallel over multiple lattices in an efficient and effective manner. Such a tunneling engine is by-product of fundamental investigations from fitness landscape analysis, local search hybridized with graybox genetic operators, general-purpose adaptive stochastic search heuristics, multi-objective evolutionary optimization, as well as, parallel and distributed optimization models. The ultimate goal of this work is to lead to a flexible, yet powerful and scalable framework for attacking complex graybox combinatorial optimization problems.

- ANR PRCI EVARISTE (2024-2028, Grant: 493K€, WP PI: B. Derbel) in collaboration with Université Angers (ANR PI: A. Goëffon), Université de Rennes (EPE), CNRS Laboratoire d'Informatique de l'Ecole Polytechnique (X).

The EVARISTE project proposes a new approach to optimize solution exploration strategies used by exact resolution methods dedicated to solve combinatorial constraint satisfaction problems. These methods generally rely on building a decision tree that gradually constructs a solution, satisfying the various constraints of the problem and potentially optimizing an objective. By using concepts and methodologies from evolutionary algorithms and fitness landscapes analysis, the goal is to develop more effective order heuristics for the decision variables of the problem and the selection of their values for classical tree-based exploration of the solution space. This involves shifting the focus from solving combinatorial problems in the initial search space to exploring the space of heuristics with appropriate metrics, leading to the discovery of new strategies for solvers. The fundamental challenge of determining an optimal sequence will be intricately connected to a challenging optimization issue known as the "distance geometry problem," which will play a central role in our approach. Ultimately, this work seeks to provide an alternative and explanatory approach to constraint solvers, which are frequently treated as black-box systems, using analytical tools and identified characteristics.

- ANR PEPR IA - Participant, project Emergences (Prof. El-Ghazali Talbi) (2023-2027 Grant: 586K€)

The expected scientific results for the Emergences project are mainly focused on performance in term of accuracy and energy efficiency of near-physics embedded AI models. These will be studied under three aspects: on emerging AI models, innovative training algorithms and the use of the physics of components. Other metrics will also be studied such as latency, tolerance to noise, suitability to process input data in various forms etc. Thus, three types of models will be explored: spiking neural networks and event-based models, disruptive physics-inspired models and near-physics design for ML. The objective at the end of this project is to be able to provide guidance towards a choice of model, a training algorithm and a given hardware solution on a per use-case basis.

- Bilateral ANR-FNR France/Luxembourg PRCI UltraBO (2023-2026, Grant: 207K€ for BONUS, PI: N. Melab) in collaboration with University of Luxembourg (Co-PI: G. Danoy).

According to Top500 modern supercomputers are increasingly large (millions of cores), heterogeneous (CPU-GPU) and less reliable (MTBF < 1h) making their programming more complex. The development of parallel algorithms for these ultra-scale supercomputers is in its infancy especially in combinatorial optimization. Our objective is to investigate the MPI+X and PGAS-based approaches for the exascale-aware design and implementation of hybrid algorithms combining exact methods (e.g. B&B) and metaheuristics (e.g. evolutionary algorithms) for solving challenging optimization problems. We will address in a holistic (uncommon) way three roadblocks on the

road to exascale: locality-aware ultra-scalability, CPU-GPU heterogeneity and checkpointing-based fault tolerance. Our application challenge is to solve to optimality very hard benchmark instances (e.g. Flow-shop ones unsolved for 25 years). For the validation, various-scale supercomputers will be used, ranging from petascale platforms, to be used for debugging, including Jean Zay (France), ULHPC (Luxembourg), SILECS/Grid'5000 (CPER CornellIA) and MesoNet (PIA Equipex+) to exascale supercomputers, to be used for real production, including the two first supercomputers of Top500 (Frontier *via* our Georgia Tech partner, Fugaku *via* our Riken partner) as well as the two EuroHPC coming ones.

- ANR PEPR Numpex/Axis Exa-MA (2022-2027, Grant: Total: 6,5M€).
The goal of the high-performance Digital for Exascale (Numpex) program, dedicated to both scientific research and industry, is twofold: (1) designing and developing the software building-blocks for the future exascale supercomputers, and (2) preparing the major application areas aimed at fully harnessing the capabilities of these latter. Numpex is composed of 5 axes including Exa-MA, which stands for *Exascale computing: Methods and Algorithms* and is organized in 7 WPs including *Optimize at Exascale* (WP5). The overall goal of WP5 consists in the design and implementation of exascale algorithms to efficiently and effectively solve large optimization problems. The research topics of the BONUS team are perfectly in line with the framework of WP5. E-G. Talbi and N. Melab are respectively the leader of and a contributor to this work-package.
- ANR PIA Equipex+ MesoNet (2021-2027, Grant: Total: 14,2M€, For ULille: 1,4M€).
The goal of the project is to set up a distributed infrastructure dedicated to the coordination of HPC and AI in France. This inclusive and structuring project, supported by GENCI partners (MESRI, CNRS, CEA, CPU, INRIA), aims to integrate at least one mesocenter by region making them regional references and relays. The infrastructure, fully integrated with the European Open Science Cloud (EOSC) initiative, should have a significant impact on the appropriation by researchers of the national and regional public HPC and AI facilities. Coordinated by GENCI, MesoNet gathers 22 partners including the mesocenter located at ULille, for which N. Melab is the co-PI. The MesoNet infrastructure is highly important for the research activities of BONUS and many other research groups including those of Inria. In addition to the funding dedicated to hardware equipment including nation-wide federated supercomputer and storage, funding will be devoted to research engineers, one of them for ULille (4,5 years), and a PhD for BONUS as well.

9.4 Regional initiatives

- CPER CornellIA (2021-2027, Grant: 820K€): this project aims at strengthening the research and infrastructure necessary for the development of scientific research in responsible and sustainable Artificial Intelligence at the regional (Hauts-de-France) level. The scientific leader in Lille is in charge of the management and the renewal of the hardware equipment of Grid'5000/SLICES-FR nationwide experimental testbed and hiring an engineer for its system and network administration and user support and development. B. Derbel took over N. Melab the responsibility of the infrastructure management and its coordination with other partners starting from late 2023. He is member of the CornellIA executive board.

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- E.-G. Talbi (Steering committee Chair): Intl. Conf. on Optimization and Learning (OLA).
- E.-G. Talbi (Steering committee): IEEE Workshop Parallel Distributed Computing and Optimization (IPDPS/PDCO).

- E.-G. Talbi (Steering committee): Intl. Conf. on Metaheuristics and Nature Inspired Computing (META).
- B. Derbel (special session co-chair): Advances in Decomposition based Evolutionary Multi-objective Optimization (ADEMO), special session at CEC/WCCI 2024.
- Z. Dahi and B. Derbel (special session co-chairs): special session on Quantum AI at CEC/WCCI 2024.
- N. Melab (Seminar chair): 11th edition of the seminar series related to Simulation and HPC at the University of Lille. The edition includes 4 seminars from IMEC (Belgium), University of Luxembourg, Safran Tech and Airbus.

10.1.2 Scientific events: selection

Member of the conference program committees

- The Parallel Problem Solving from Nature international conference (PPSN).
- The ACM Genetic and Evolutionary Computation Conference (GECCO).
- The IEEE Congress on Evolutionary Computation (CEC).
- European Conference on Evolutionary Computation in Combinatorial Optimization (EvoCOP).
- International Conference on Evolutionary Multi-criterion Optimization (EMO).
- Intl. Conf. on Optimization and Learning (OLA).
- QAI workshop co-located with IJCAI.

10.1.3 Journal

Member of the editorial boards

- E.-G. Talbi (Editorial board member) : ACM Transactions on Evolutionary Learning and Optimization (TELO), since 2023.
- N. Melab (Associate Editor): ACM Computing Surveys, since 2019.

Reviewer - reviewing activities

- IEEE Transactions on Evolutionary Computation (TEVC).
- Applied Soft Computing (Elsevier).
- ACM Transactions on Evolutionary Learning and Optimization.
- ACM Computing Surveys.

10.1.4 Invited talks

- E.-G. Talbi. Neural architecture search: a unified view, Keynote speaker, February 2024, Workshop on Intelligence in Business and Industry, Tunis, Tunisia.
- E.-G. Talbi. Metaheuristics for the automated design and configuration of deep neural networks, Keynote speaker, June 2024, ECCO'2024 37th Conference of the European Chapter on Combinatorial Optimization, Gent, Belgium.
- B. Derbel. Big Optimization and Ultra-Scale Computing, November 2024, KAIST, Seoul, Korea.
- B. Derbel. Heterogeneous multi-objective optimization, July 2024, Shinshu University, Japan.
- N. Melab, T. Carneiro and G. Helbecque. Big Optimization using Ultra-scale Computing, March 2024, University of Malaga, Spain.

10.1.5 Leadership within the scientific community

- B. Derbel: Scientific leader of Grid'5000/SLICES-FR HPC testbed at Lille, since 2024.
- E.-G. Talbi: Co-president of the working group “META: Metaheuristics - Theory and applications”, GDR RO and GDR MACS.
- B. Derbel: Chair of the IEEE CIS Task Force on Decomposition-based Techniques in Evolutionary Computation (DTEC).
- E.-G. Talbi: Co-Chair of the IEEE Task Force on Cloud Computing within the IEEE Computational Intelligence Society.

10.1.6 Scientific expertise

- E.-G. Talbi: External project expert, Fonds de la Recherche Scientifique (F.R.S-FNRS), Belgium, March 2024.
- E.-G. Talbi: External project expert, Fonds de la Recherche Scientifique (F.R.S-FNRS), Belgium, October 2024.
- N. Melab: ANR PRCE, CE46 – Calcul haute performance, Modèles numériques, simulation, applications, April 2024.

10.1.7 Research administration

- B. Derbel: Coordinator of the research theme (GT)OPTIMA at the CRISAL UMR laboratory, since late 2023.
- B. Derbel: Member of the Scientific Board of the CRISAL UMR laboratory, since late 2023.
- B. Derbel: Member of the Scientific Board for the MADIS doctoral school at the University of Lille, since 2022.
- N. Melab: Member of the Scientific Board for the Inria Lille research center, from Feb. 2019 to Feb. 2024.
- N. Melab: Chargé de Mission of High Performance Computing and Simulation at Université de Lille, since 2010.

10.2 Teaching - Supervision - Juries

10.2.1 Teaching

Taught courses

- Bachelor: Z. Dahi, Algorithms and Data Structures, 36h, University of Lille, France.
- International Master lecture: N. Melab, Supercomputing, 45h ETD, M2, University of Lille, France.
- Master lecture: N. Melab, Operations Research, 60h ETD, M1, University of Lille, France.
- Master: B. Derbel, Algorithms and Complexity, 35h, M1, University of Lille, France.
- Master: B. Derbel, Optimization and machine learning, 24h, M1, University of Lille, France.
- Engineering school: E.-G. Talbi, Advanced optimization, 36h, Polytech'Lille, University of Lille, France.
- Engineering school: E.-G. Talbi, Data mining, 36h, Polytech'Lille, University of Lille, France.
- Engineering school: E.-G. Talbi, Operations research, 60h, Polytech'Lille, University of Lille, France.
- Engineering school: E.-G. Talbi, Graphs, 25h, Polytech'Lille, University of Lille, France.

Teaching responsibilities

- Head of the international relations: E.-G. Talbi, Polytech'Lille, Université de Lille, France.
- Head of the international relations: B. Derbel, Computer Science Department, Faculty of Science and Technology, Université de Lille, France.
- Master leading: N. Melab, Co-head (with O. Goubet) of the international Master 2 of High-performance Computing and Simulation, Université de Lille, France.

10.2.2 Supervision

- PhD defense [40]: Maxime Gobert, Contributions to the Analysis and Design of Parallel Batched Bayesian Optimization Algorithms. Supervisors: N. Melab (Université de Lille) and D. Tuytens (Université de Mons, Belgium). Defended on June 21st, 2024.
- PhD (co-tutelle) defense [41]: Guillaume Helbecque, PGAS-based Parallel Branch-and-Bound for Ultra-Scale GPU-powered Supercomputers. Supervisors: N. Melab (Université de Lille) and P. Bouvry (Université du Luxembourg). Defended on Jan. 10th, 2025.
- PhD defense: Housseem Ouertatani, Multi-objective optimization of deep neural networks for embedded applications. Supervisors: E.-G. Talbi and S. Niar (Université Polytechnique Hauts-de-France). Defended on Dec. 9th, 2024.
- PhD in progress: Mehdi El Khadiri, Exascale optimization using fractal-based decomposition. Supervisor: E.-G. Talbi. Started in 2024.
- PhD in progress: Bohdan Ivaniuk, Automated design and multi-objective optimization of parallel deep networks for automatic detection in real-time. Supervisor: E.-G. Talbi. Started in 2023.
- PhD in progress: Jérôme Rouzé, Parallel Hybrid Metaheuristics for Noisy Intermediate-Scale Quantum Machines. Supervisors: N. Melab (Université de Lille) and D. Tuytens (Université de Mons, Belgium). Started in Nov. 2023.
- PhD (co-tutelle) in progress: Ivan Tagliaferro de Oliveira Tezoto, Exascale Exact Optimization based on the MPI+X Approach. Supervisors: N. Melab (Université de Lille) and G. Danoy (Université du Luxembourg). Started in Oct. 2024.
- PhD in progress: Thomas Firmin, Pulse neuron networks and parameter optimization for massively parallel GPU-powered clusters. Supervisors: E.-G. Talbi and P. Boulet (Emeraude Team, CRISAL lab). Started in Oct. 2021, defense scheduled in Januray 2025.
- PhD in progress: Julie Keisler, Réseaux de neurones profonds pour la prédiction de séries spatio-temporelles. Supervisor: E.-G. Talbi, CIFRE with EDF. Started in Oct. 2022.
- PhD in progress: David Redon, Enabling Large Scale Computational Intelligence with HPC. Supervisors: B. Derbel and P. Fortin (Université de Lille). Started in Oct. 2020.

10.2.3 Juries

- E.-G. Talbi (Reviewer / rapporteur): PhD thesis of Francesco Zito, In what ways do data-driven AI algorithms impact complex real-world systems and decision making?, University of Catania (Italy), defended on December 2024.
- E.-G. Talbi (Reviewer / rapporteur): PhD thesis of Arcadi Llanza Carmona, Design of fractal decomposition based algorithms: Application to computer vision problems, Université Paris-Est Créteil, Defended on December 2024.
- B. Derbel (Reviewer / Rapporteur): PhD thesis of Khaoula Zeiman, Optimisation du déploiement de capteurs dans les bâtiments intelligents, Université de Haute-Alsace, Defended on June 2024.

- N. Melab (Jury president): PhD thesis of Soufiane Mallem, Design of Meta-Learning and Label Correction Algorithms Based on Bilevel Optimization for Learning with Noisy Labels, Université Paris-Est Créteil, defended on Dec. 2024.
- N. Melab (Reviewer): PhD thesis of Xinwei Jin, A Personalized, Uncertainty-aware Trustworthy Algorithm for Effective Pain Assessment using Biosignals, The University of Sydney, Australia, defended June 2024.
- N. Melab (Reviewer): PhD thesis of Zineb Ziani, AI and HPC Convergence for Enhanced Anomaly Detection, Université Paris-Saclay, Defended Jan. 2025 (Reviewed in Dec. 2024).

10.3 Popularization

10.3.1 Others science outreach relevant activities

- B. Derbel: Welcoming of a group of international visiting (school) students at Inria, April 2024, presentation and discussion about being a researcher.
- B. Derbel and Z. Dahi: Welcoming of a young (school) student performing an internship at Inria, January 2024, non-expert introduction to the optimization field and general presentation of the team's research activities.

11 Scientific production

11.1 Major publications

- [1] O. Abdelkafi, L. Idoumghar and J. Lepagnot. 'A Survey on the Metaheuristics Applied to QAP for the Graphics Processing Units'. In: *Parallel Processing Letters* 26.3 (2016), pp. 1–20.
- [2] A. Bendjoudi, N. Melab and E.-G. Talbi. 'FTH-B&B: A Fault-Tolerant Hierarchical Branch and Bound for Large Scale Unreliable Environments'. In: *IEEE Trans. Computers* 63.9 (2014), pp. 2302–2315 (cit. on p. 7).
- [3] S. Cahon, N. Melab and E.-G. Talbi. 'ParadisEO: A Framework for the Reusable Design of Parallel and Distributed Metaheuristics'. In: *J. Heuristics* 10.3 (2004), pp. 357–380 (cit. on p. 3).
- [4] F. Daolio, A. Liefooghe, S. Verel, H. Aguirre and K. Tanaka. 'Problem Features versus Algorithm Performance on Rugged Multiobjective Combinatorial Fitness Landscapes'. In: *Evolutionary Computation* 25.4 (2017) (cit. on p. 6).
- [5] B. Derbel. 'Contributions to single- and multi- objective optimization: towards distributed and autonomous massive optimization'. HDR dissertation. Université de Lille, 2017 (cit. on pp. 5, 6).
- [6] B. Derbel, A. Liefooghe, Q. Zhang, H. Aguirre and K. Tanaka. 'Multi-objective Local Search Based on Decomposition'. In: *Parallel Problem Solving from Nature - PPSN XIV - 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings*. 2016, pp. 431–441.
- [7] B. Derbel, G. Pruvost, A. Liefooghe, S. Verel and Q. Zhang. 'Walsh-based surrogate-assisted multi-objective combinatorial optimization: A fine-grained analysis for pseudo-boolean functions'. In: *Applied Soft Computing* 136 (Mar. 2023), p. 110061. DOI: [10.1016/j.asoc.2023.110061](https://doi.org/10.1016/j.asoc.2023.110061). URL: <https://hal.science/hal-04073811> (cit. on p. 5).
- [8] J. Gmys, M. Mezmaiz, N. Melab and D. Tuytens. 'IVM-based parallel branch-and-bound using hierarchical work stealing on multi-GPU systems'. In: *Concurrency and Computation: Practice and Experience* 29.9 (2017) (cit. on pp. 4, 7).
- [9] A. Liefooghe, F. Daolio, S. Verel, B. Derbel, H. Aguirre and K. Tanaka. 'Landscape-aware performance prediction for evolutionary multi-objective optimization'. In: *IEEE Transactions on Evolutionary Computation* 24.6 (2020), pp. 1063–1077. DOI: [10.1109/TEVC.2019.2940828](https://doi.org/10.1109/TEVC.2019.2940828). URL: <https://hal.archives-ouvertes.fr/hal-02294201> (cit. on p. 6).
- [10] T. V. Luong, N. Melab and E.-G. Talbi. 'GPU Computing for Parallel Local Search Metaheuristic Algorithms'. In: *IEEE Trans. Computers* 62.1 (2013), pp. 173–185 (cit. on p. 7).

- [11] A. Nakib, S. Ouchraa, N. Shvai, L. Souquet and E.-G. Talbi. ‘Deterministic metaheuristic based on fractal decomposition for large-scale optimization’. In: *Appl. Soft Comput.* 61 (2017), pp. 468–485 (cit. on p. 4).
- [12] T.-T. Vu and B. Derbel. ‘Parallel Branch-and-Bound in Multi-core Multi-CPU Multi-GPU Heterogeneous Environments’. In: *Future Generation Computer Systems* 56 (20th Mar. 2016), pp. 95–109. DOI: [10.1016/j.future.2015.10.009](https://doi.org/10.1016/j.future.2015.10.009). URL: <https://inria.hal.science/hal-01067662>.

11.2 Publications of the year

International journals

- [13] M. Becerra-Rozas, J. Lemus-Romani, B. Crawford, R. Soto and E.-G. Talbi. ‘Intelligent decision-making for binary coverage: Unveiling the potential of the multi-armed bandit selector’. In: *Expert Systems with Applications* 251 (Oct. 2024), p. 124112. DOI: [10.1016/j.eswa.2024.124112](https://doi.org/10.1016/j.eswa.2024.124112). URL: <https://hal.science/hal-04901029> (cit. on p. 12).
- [14] G. Briffoteaux, N. Melab, M. Mezmaç and D. Tuyttens. ‘Investigating surrogate-based hybrid acquisition processes. Application to Covid-19 contact mitigation’. In: *Applied Soft Computing* 151 (Jan. 2024), p. 111134. DOI: [10.1016/j.asoc.2023.111134](https://doi.org/10.1016/j.asoc.2023.111134). URL: <https://hal.science/hal-04383812> (cit. on p. 16).
- [15] L. Canonne, B. Derbel, M. Tsuji and M. Sato. ‘Large-scale and cooperative graybox parallel optimization on the supercomputer Fugaku’. In: *Journal of Parallel and Distributed Computing* 191 (Sept. 2024), p. 104921. DOI: [10.1016/j.jpdc.2024.104921](https://doi.org/10.1016/j.jpdc.2024.104921). URL: <https://hal.science/hal-04867018> (cit. on pp. 12, 19).
- [16] R. Cosson, R. Santana, B. Derbel and A. Liefoghe. ‘On bi-objective combinatorial optimization with heterogeneous objectives’. In: *European Journal of Operational Research* 319.1 (16th Nov. 2024), pp. 89–101. DOI: [10.1016/j.ejor.2024.06.029](https://doi.org/10.1016/j.ejor.2024.06.029). URL: <https://hal.science/hal-04692894> (cit. on p. 13).
- [17] T. Dokeroglu, T. Kucukyilmaz and E.-G. Talbi. ‘Hyper-heuristics: A survey and taxonomy’. In: *Computers & Industrial Engineering* 187 (Jan. 2024), p. 109815. DOI: [10.1016/j.cie.2023.109815](https://doi.org/10.1016/j.cie.2023.109815). URL: <https://hal.science/hal-04901015> (cit. on p. 12).
- [18] F. Ducobu, N. Kugalur-Palanisamy, G. Briffoteaux, M. Gobert, D. Tuyttens, P. Arrazola and E. Rivière-Lorphèvre. ‘Identification of the Constitutive and Friction Models Parameters via a Multi-Objective Surrogate-Assisted Algorithm for the Modeling of Machining—Application to Arbitrary Lagrangian Eulerian Orthogonal Cutting of Ti6Al4V’. In: *Journal of Manufacturing Science and Engineering* 146.6 (1st June 2024). DOI: [10.1115/1.4065223](https://doi.org/10.1115/1.4065223). URL: <https://hal.science/hal-04801841> (cit. on p. 16).
- [19] P. Favaro, M. Gobert and J.-F. Toubeau. ‘Multi-fidelity optimization for the day-ahead scheduling of Pumped Hydro Energy Storage’. In: *Journal of Energy Storage* 103 (Dec. 2024), p. 114096. DOI: [10.1016/j.est.2024.114096](https://doi.org/10.1016/j.est.2024.114096). URL: <https://hal.science/hal-04801851> (cit. on p. 16).
- [20] F. Felten, E.-G. Talbi and G. Danoy. ‘Multi-Objective Reinforcement Learning Based on Decomposition: A Taxonomy and Framework’. In: *Journal of Artificial Intelligence Research* 79 (26th Feb. 2024), pp. 679–723. DOI: [10.1613/jair.1.15702](https://doi.org/10.1613/jair.1.15702). URL: <https://hal.science/hal-04901041> (cit. on p. 13).
- [21] J. Gamot, M. Balesdent, R. Wuilbercq, A. Tremolet and N. Melab. ‘Two-Level Approach for Simultaneous Component Assignment and Layout Optimization with Applications to Spacecraft Optimal Layout’. In: *Applied Sciences* 14.18 (10th Sept. 2024), p. 8120. DOI: [10.3390/app14188120](https://doi.org/10.3390/app14188120). URL: <https://hal.science/hal-04809409> (cit. on pp. 12, 14).
- [22] M. Gobert, G. Briffoteaux, J. Gmys, N. Melab and D. Tuyttens. ‘Observations in applying Bayesian versus evolutionary approaches and their hybrids in parallel time-constrained optimization’. In: *Engineering Applications of Artificial Intelligence* 137 (Nov. 2024), p. 109075. DOI: [10.1016/j.engappai.2024.109075](https://doi.org/10.1016/j.engappai.2024.109075). URL: <https://hal.science/hal-04707543> (cit. on p. 16).

- [23] J. Keisler, T. El-Ghazali, C. Sandra and C. Gilles. ‘An Algorithmic Framework for the Optimization of Deep Neural Networks Architectures and Hyperparameters’. In: *Journal of Machine Learning Research* 25.201 (2024). URL: <https://hal.science/hal-04901056> (cit. on p. 15).
- [24] J. Á. Morell, Z. A. Dahi, F. Chicano, G. Luque and E. Alba. ‘A multi-objective approach for communication reduction in federated learning under devices heterogeneity constraints’. In: *Future Generation Computer Systems*. Lecture Notes in Computer Science 155 (June 2024), pp. 367–383. DOI: [10.1016/j.future.2024.02.022](https://doi.org/10.1016/j.future.2024.02.022). URL: <https://hal.science/hal-04900917> (cit. on pp. 12, 14).
- [25] D. Porumbel, I. Coelho and E.-G. Talbi. ‘Using an exact bi-objective decoder in a memetic algorithm for arc-routing (and other decoder-expressible) problems’. In: *European Journal of Operational Research* 313.1 (Feb. 2024), pp. 25–43. DOI: [10.1016/j.ejor.2023.07.021](https://doi.org/10.1016/j.ejor.2023.07.021). URL: <https://hal.science/hal-04901021> (cit. on pp. 12, 14).
- [26] E.-G. Talbi. ‘Metaheuristics for variable-size mixed optimization problems: A unified taxonomy and survey’. In: *Swarm and Evolutionary Computation* 89 (Aug. 2024), p. 101642. DOI: [10.1016/j.swevo.2024.101642](https://doi.org/10.1016/j.swevo.2024.101642). URL: <https://hal.science/hal-04901075> (cit. on p. 12).

International peer-reviewed conferences

- [27] T. Carneiro, E. Kayraklioglu, G. Helbecque and N. Melab. ‘Investigating Portability in Chapel for Tree-based Optimization on GPU-powered Clusters’. In: *Euro-Par 2024: Parallel Processing, 30th International European Conference on Parallel and Distributed Computing*. Madrid, Spain, 26th Aug. 2024, pp. 386–399. DOI: [10.1007/978-3-031-69583-4_27](https://doi.org/10.1007/978-3-031-69583-4_27). URL: <https://hal.science/hal-04709120> (cit. on p. 18).
- [28] Z. A. Dahi, F. Chicano and G. Luque. ‘An Evolutionary Deep Learning Approach for Efficient Quantum Algorithms Transpilation’. In: *Proceedings, Part II of the 27th European Conference, EvoApplications 2024, Held as Part of EvoStar 2024, Aberystwyth, UK, April 3–5, 2024*. EvoApplications 2024 - 27th European Conference. Vol. 14635. Lecture Notes in Computer Science. Aberystwyth, United Kingdom: Springer Nature Switzerland, 21st Mar. 2024, pp. 240–255. DOI: [10.1007/978-3-031-56855-8_15](https://doi.org/10.1007/978-3-031-56855-8_15). URL: <https://hal.science/hal-04900875> (cit. on p. 17).
- [29] Z. A. Dahi, F. Chicano, G. Luque, B. Derbel and E. Alba. ‘Scalable Quantum Approximate Optimiser for Pseudo-Boolean Multi-objective Optimisation’. In: *LNCS. Parallel Problem Solving from Nature – PPSN XVIII*. Vol. 15151. Lecture Notes in Computer Science. Hagenberg, Austria: Springer Nature Switzerland, 7th Sept. 2024, pp. 268–284. DOI: [10.1007/978-3-031-70085-9_17](https://doi.org/10.1007/978-3-031-70085-9_17). URL: <https://hal.science/hal-04867099> (cit. on p. 13).
- [30] Z. A. Dahi, J. Galindo and D. Galindo. ‘On the Effect of Teaching with Videos: Case Study of Nested Loops in Computer Programming in Industrial Engineering Degree’. In: *2024 XVI Congreso de Tecnología, Aprendizaje y Enseñanza de la Electrónica (XVI Technologies Applied to Electronics Teaching Conference) (TAEE)*. 2024 XVI Congreso de Tecnología, Aprendizaje y Enseñanza de la Electrónica (XVI Technologies Applied to Electronics Teaching Conference) (TAEE). Malaga, Spain: IEEE, 26th June 2024, pp. 1–4. DOI: [10.1109/TAEE59541.2024.10604987](https://doi.org/10.1109/TAEE59541.2024.10604987). URL: <https://hal.science/hal-04900933>.
- [31] T. Firmin, P. Boulet and E.-G. Talbi. ‘Asynchronous Multi-Fidelity Hyperparameter Optimization of Spiking Neural Networks’. In: *2024 International Conference on Neuromorphic Systems (ICONS)*. Arlington, France: IEEE, 30th July 2024, pp. 55–62. DOI: [10.1109/ICONS62911.2024.00016](https://doi.org/10.1109/ICONS62911.2024.00016). URL: <https://hal.science/hal-04904504> (cit. on p. 19).
- [32] C. Garcia-Garcia, B. Derbel, A. Morales-Reyes and H. J. Escalante. ‘Speeding up the Multi-objective NAS Through Incremental Learning’. In: *LNCS LNAI. 23rd Mexican International Conference on Artificial Intelligence*. Vol. 15247. Lecture Notes in Computer Science. Puebla, Mexico: Springer Nature Switzerland, 17th Oct. 2025, pp. 3–15. DOI: [10.1007/978-3-031-75543-9_1](https://doi.org/10.1007/978-3-031-75543-9_1). URL: <https://hal.science/hal-04867049> (cit. on p. 16).

- [33] G. Helbecque, T. Carneiro, N. Melab, J. Gmys and P. Bouvry. ‘PGAS Data Structure for Unbalanced Tree-Based Algorithms at Scale’. In: *Computational Science – ICCS 2024*. ICCS 2024 - 24th International Conference on Computational Science. Vol. 14834. Malaga, Spain, 29th June 2024, pp. 103–111. DOI: [10.1007/978-3-031-63759-9_13](https://doi.org/10.1007/978-3-031-63759-9_13). URL: <https://hal.science/hal-04636184> (cit. on p. 18).
- [34] G. Helbecque, E. Krishnasamy, T. Carneiro, N. Melab and P. Bouvry. ‘A Chapel-based Multi-GPU Branch-and-Bound Algorithm’. In: *Euro-Par 2024: Parallel Processing Workshops*. 22nd International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms. Madrid, Spain, 27th Aug. 2024. URL: <https://hal.science/hal-04709106> (cit. on p. 18).
- [35] G. Helbecque, E. Krishnasamy, N. Melab and P. Bouvry. ‘GPU-Accelerated Tree-Search in Chapel versus CUDA and HIP’. In: 14th IEEE Workshop Parallel / Distributed Combinatorics and Optimization (PDCO 2024). San Francisco, United States, 31st May 2024. DOI: [10.1109/IPDPSW63119.2024.00156](https://doi.org/10.1109/IPDPSW63119.2024.00156). URL: <https://hal.science/hal-04551856> (cit. on p. 18).
- [36] H. Ouertatani, C. Maxim, S. Niar and E.-G. Talbi. ‘Accelerated NAS via pretrained ensembles and multi-fidelity Bayesian Optimization’. In: *Conference proceedings are published by Springer in Lecture Notes in Computer Science*. 33rd International Conference on Artificial Neural Networks (ICANN). Lugano, Switzerland, 17th Sept. 2024. DOI: [10.1007/978-3-031-72332-2_17](https://doi.org/10.1007/978-3-031-72332-2_17). URL: <https://hal.science/hal-04611343> (cit. on p. 15).
- [37] H. Ouertatani, C. Maxim, S. Niar and E.-G. Talbi. ‘Neural Architecture Tuning: A BO-Powered NAS Tool’. In: International Conference in Optimization and Learning (OLA). Dubrovnik, Croatia, 13th May 2024. DOI: [10.1007/978-3-031-77941-1_7](https://doi.org/10.1007/978-3-031-77941-1_7). URL: <https://hal.science/hal-04540630> (cit. on p. 15).
- [38] J. Rouzé, N. Melab and D. Tuytens. ‘A Parallel Genetic Algorithm for Qubit Mapping on Noisy Intermediate-Scale Quantum Machines’. In: International Conference in Optimization and Learning - OLA 2024. Optimization and Learning. Dubrovnik, Croatia: Springer, 9th Mar. 2025. URL: <https://hal.science/hal-04916922>.

Conferences without proceedings

- [39] G. Helbecque, E. Krishnasamy, N. Melab and P. Bouvry. ‘GPU Computing in Chapel: Application to Tree-Search Algorithms’. In: International Conference on Optimization and Learning (OLA 2024). Dubrovnik, Croatia, 13th May 2024. URL: <https://hal.science/hal-04551844> (cit. on p. 18).

Doctoral dissertations and habilitation theses

- [40] M. Gobert. ‘Contributions to the Analysis and Design of Parallel Batched Bayesian Optimization Algorithms’. Université de Mons (Belgique), 21st June 2024. DOI: [10.1115/1.4065223\(2024\)](https://doi.org/10.1115/1.4065223(2024)). URL: <https://hal.science/tel-04801888> (cit. on p. 27).
- [41] G. Helbecque. ‘PGAS-based Parallel Branch-and-Bound for Ultra-Scale GPU-powered Supercomputers’. Université de Lille; Université du Luxembourg, 10th Jan. 2025. URL: <https://hal.science/tel-04902137> (cit. on p. 27).

Reports & preprints

- [42] T. Firmin, P. Boulet and E.-G. Talbi. *Parallel Hyperparameter Optimization Of Spiking Neural Networks*. Dec. 2024. DOI: [10.1016/j.neucom.2024.128483](https://doi.org/10.1016/j.neucom.2024.128483). URL: <https://hal.science/hal-04464394> (cit. on pp. 12, 18).
- [43] D. Redon, P. Fortin, B. Derbel, M. Tsuji and M. Sato. *Massively parallel CMA-ES with increasing population*. 16th Sept. 2024. URL: <https://hal.science/hal-04698736>.
- [44] P. E.-G. Talbi. *Metaheuristics for (Variable-Size) Mixed Optimization Problems: A Unified Taxonomy and Survey*. Université de Lille; INRIA Lille Nord Europe; CRISTAL; BONUS, 8th Jan. 2024, p. 55. URL: <https://inria.hal.science/hal-04381240> (cit. on p. 12).

Software

- [45] [SW] G. Helbecque, J. Gmys, T. Carneiro, N. Melab and P. Bouvry, *P3D-DFS: Productivity- and Performance-aware Parallel Distributed Depth-First Search* version v1.0, 25th Sept. 2024. LIC: BSD 3-Clause License. DOI: [10.5281/zenodo.7328540](https://doi.org/10.5281/zenodo.7328540), HAL: [hal-04709165](https://hal.science/hal-04709165), URL: <https://github.com/Guillaume-Helbecque/P3D-DFS> (cit. on p. 18).
- [46] [SW] G. Helbecque, E. Krishnasamy, N. Melab and P. Bouvry, *GPU-accelerated tree-search in Chapel* 25th Sept. 2024. LIC: BSD 3-Clause License. DOI: [10.5281/zenodo.10786275](https://doi.org/10.5281/zenodo.10786275), HAL: [hal-04709186](https://hal.science/hal-04709186), URL: <https://github.com/Guillaume-Helbecque/GPU-accelerated-tree-search-Chapel> (cit. on p. 18).

11.3 Cited publications

- [47] M. Balesdent, L. Brevault, N. B. Price, S. Defoort, R. Le Riche, N.-H. Kim, R. T. Haftka and N. Bérend. ‘Space Engineering: Modeling and Optimization with Case Studies’. In: ed. by G. Fasano and J. D. Pintér. Springer International Publishing, 2016. Chap. Advanced Space Vehicle Design Taking into Account Multidisciplinary Couplings and Mixed Epistemic/Aleatory Uncertainties, pp. 1–48. DOI: [10.1007/978-3-319-41508-6_1](https://doi.org/10.1007/978-3-319-41508-6_1). URL: http://dx.doi.org/10.1007/978-3-319-41508-6_1 (cit. on p. 8).
- [48] B. Derbel, D. Brockhoff, A. Liefoghe and S. Verel. ‘On the Impact of Multiobjective Scalarizing Functions’. In: *Parallel Problem Solving from Nature - PPSN XIII - 13th International Conference, Ljubljana, Slovenia, September 13-17, 2014. Proceedings*. 2014, pp. 548–558 (cit. on p. 5).
- [49] B. Derbel, A. Liefoghe, G. Marquet and E. Talbi. ‘A fine-grained message passing MOEA/D’. In: *IEEE Congress on Evolutionary Computation, CEC 2015, Sendai, Japan, May 25-28, 2015*. 2015, pp. 1837–1844 (cit. on p. 7).
- [50] J. Dreo, A. Liefoghe, S. Verel, M. Schoenauer, J. J. Merelo, A. Quemy, B. Bouvier and J. Gmys. ‘Paradiseo: From a Modular Framework for Evolutionary Computation to the Automated Design of Metaheuristics’. In: *GECCO 2021 - Genetic and Evolutionary Computation Conference*. 2021 Genetic and Evolutionary Computation Conference Companion. ACM Sigev. Lille / Virtual, France: ACM, July 2021, pp. 1522–1530. DOI: [10.1145/3449726.3463276](https://doi.org/10.1145/3449726.3463276). URL: <https://hal-pasteur.archives-ouvertes.fr/pasteur-03220556> (cit. on p. 3).
- [51] R. Haftka, D. Villanueva and A. Chaudhuri. ‘Parallel surrogate-assisted global optimization with expensive functions – a survey’. In: *Structural and Multidisciplinary Optimization* 54(1) (2016), pp. 3–13 (cit. on p. 6).
- [52] D. Jones, M. Schonlau and W. Welch. ‘Efficient Global Optimization of Expensive Black-Box Functions’. In: *Journal of Global Optimization* 13(4) (1998), pp. 455–492 (cit. on p. 5).
- [53] J. Pelamatti, L. Brevault, M. Balesdent, E.-G. Talbi and Y. Guerin. ‘How to deal with mixed-variable optimization problems: An overview of algorithms and formulations’. In: *Advances in Structural and Multidisciplinary Optimization, Proc. of the 12th World Congress of Structural and Multidisciplinary Optimization (WCSMO12)*. Springer, 2018, pp. 64–82. DOI: [10.1007/978-3-319-67988-4_5](https://doi.org/10.1007/978-3-319-67988-4_5). URL: http://dx.doi.org/10.1007/978-3-319-67988-4_5 (cit. on pp. 4, 6).
- [54] F. Shahzad, J. Thies, M. Kreutzer, T. Zeiser, G. Hager and G. Wellein. ‘CRAFT: A library for easier application-level Checkpoint/Restart and Automatic Fault Tolerance’. In: *CoRR* abs/1708.02030 (2017). URL: <http://arxiv.org/abs/1708.02030> (cit. on p. 7).
- [55] N. Shavit. ‘Data Structures in the Multicore Age’. In: *Communications of the ACM* 54.3 (2011), pp. 76–84 (cit. on p. 7).
- [56] M. Snir and al. ‘Addressing Failures in Exascale Computing’. In: *Int. J. High Perform. Comput. Appl.* 28.2 (May 2014), pp. 129–173 (cit. on p. 7).
- [57] E.-G. Talbi. ‘Combining metaheuristics with mathematical programming, constraint programming and machine learning’. In: *Annals OR* 240.1 (2016), pp. 171–215 (cit. on p. 4).

-
- [58] T. Vu and B. Derbel. ‘Parallel Branch-and-Bound in multi-core multi-CPU multi-GPU heterogeneous environments’. In: *Future Generation Comp. Syst.* 56 (2016), pp. 95–109 (cit. on p. 7).
- [59] X. Zhang, Y. Tian, R. Cheng and Y. Jin. ‘A Decision Variable Clustering-Based Evolutionary Algorithm for Large-Scale Many-Objective Optimization’. In: *IEEE Trans. Evol. Computation* 22.1 (2018), pp. 97–112 (cit. on p. 5).
- [60] D. Zouache, A. Got, D. Alarabiat, L. Abualigah and E.-G. Talbi. ‘A novel multi-objective wrapper-based feature selection method using quantum-inspired and swarm intelligence techniques’. In: *Multimedia Tools and Applications* 83.8 (2024), pp. 22811–22835. DOI: [10.1007/s11042-023-16411-9](https://doi.org/10.1007/s11042-023-16411-9) (cit. on p. 17).