RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

2025 *Activity Report*

RESEARCH CENTRE: **Inria Centre at the University of Lille**
IN PARTNERSHIP WITH: **Université de Lille**

Project-Team

BONUS

Big Optimization aNd Ultra-Scale Computing

In collaboration with Centre de Recherche en Informatique, Signal et Automatique de Lille

Project-Team BONUS

Creation of the Project-Team: 2019 June 01

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

Keywords

Computer sciences and digital sciences

A1.1.11. – Quantum architectures

A6.2.6. – Optimization

A6.2.7. – HPC for machine learning

A7.1.4. – Quantum algorithms

A8.2.1. – Operations research

A8.2.2. – Evolutionary algorithms

A9.2.4. – Optimization and learning

A9.2.5. – Bayesian methods

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

Contents

1 Team members, visitors, external collaborators

Research Scientist

- Abdelmoiz Zakaria Dahi [INRIA, ISFP]

Faculty Members

- Bilel Derbel [Team leader, UNIV LILLE, Professor, HDR]
- Sarah Degaugue [UNIV LILLE, Associate Professor, from Sep 2025]
- Nouredine Melab [UNIV LILLE, Professor, HDR]
- El-Ghazali Talbi [UNIV LILLE, Professor, HDR]

Post-Doctoral Fellows

- Muhammad Junaid Ali [INRIA, Post-Doctoral Fellow, from Apr 2025]
- Jorge Mario Cruz Duarte [UNIV LILLE, Post-Doctoral Fellow]
- Guillaume Helbecque [UNIV LILLE, Post-Doctoral Fellow]

PhD Students

- Lander Argote Garcia [UNIV LILLE, from Oct 2025]
- Francesco Cecere [INRIA, from Sep 2025]
- Mahmoud El Mehdi El Khadiri [INRIA]
- Thomas Firmin [UNIV LILLE, until Jan 2025]
- Bohdan Ivaniuk Skulskyi [VINCI, CIFRE]
- Julie Keisler [EDF, CIFRE, until Feb 2025]
- David Redon [UNIV LILLE, ATER]
- Jerome Rouze [UNIV MONS]
- Ivan Tagliaferro De Oliveira Tezoto [UNIV LILLE, until Nov 2025]
- Jean Philippe Valois [UNIV LILLE, from Oct 2025]

Interns and Apprentices

- Riadh Boulifi [INRIA, Intern, from Apr 2025 until Jul 2025]
- Nathan Davouse [INRIA, Intern, until Feb 2025]
- Jean Philippe Valois [UNIV LILLE, Intern, from Mar 2025 until Aug 2025]

Administrative Assistant

- Karine Lewandowski [INRIA]

Visiting Scientist

- Mohammed Thousif Pagala [INDIAN INSTITUTE OF SCIENCE, BANGALORE]

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:

$$\begin{array}{ll} \text{Min/Max}\ \mathbf{F}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_m(\mathbf{x})) \\ \text{subject to} \qquad \mathbf{x} \in \mathbf{\Omega}, \end{array}$$

where $\mathbf{x} = (x_1, x_2, ..., x_n)$ is the decision variable vector of dimension n, $\mathbf{\Omega}$ 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, 44] 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 resulting subproblems on ultra-scale architectures. Here also several issues will be addressed: the definition of

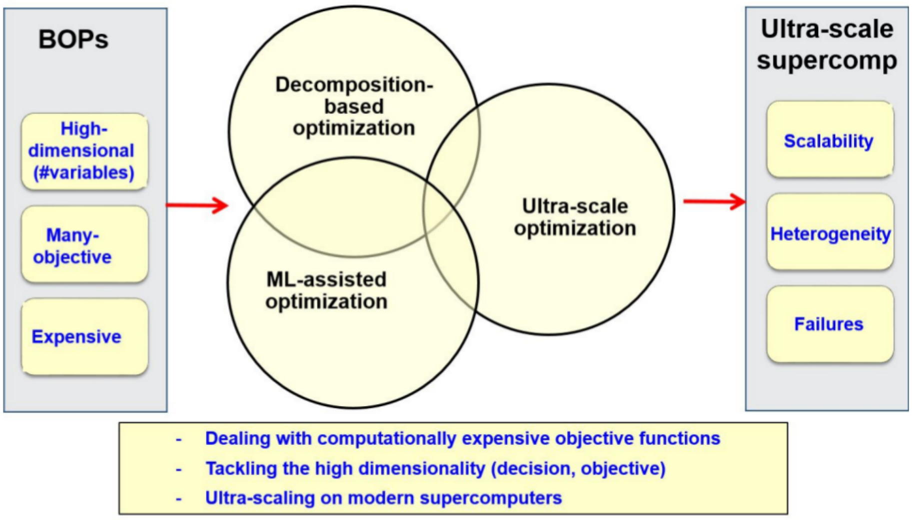

Figure 1: Research challenges/objectives and lines

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 solutions [1] increases drastically with the number of objectives leading to a very slow convergence to the Pareto Front [2]. 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 [51]. *The major issue of this kind of decomposition is the presence of categorial variables in the discrete part [47]. 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.

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

[1] A 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.
[2] The Pareto Front is the set of non-dominated solutions.

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 [42]. 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 categorial 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 [53], 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) [46] 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 [46]. 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 [45] 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 [47], 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 supercomputers [3]. However, their efficient use raises a big challenge consisting in managing 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 [49]. On

[3]In 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).

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 [43, 52]. 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 paradigm [4] 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 [52], [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 [50]. Indeed, with the increase of their size to millions of processing cores their Mean-Time Between Failures (MTBF) tends to be shorter and shorter [48]. 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 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] A 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.

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 confort, minimizing its energy bill and minimzing 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 [41]. 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 categorial 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.

4.4 Big optimization for and using NISQ systems

Beyond classical application domains, we investigate large-scale optimization problems and paradigms arising in the context of Noisy Intermediate-Scale Quantum (NISQ) systems following two complementary lines of research.

On the one hand, current quantum hardware is characterized by limited qubit counts, constrained connectivity, and significant noise levels, which severely restrict the execution of quantum circuits. As a consequence, the compilation and mapping of quantum programs onto NISQ devices give rise to large-scale combinatorial optimization problems that must be solved efficiently on classical high-performance computing (HPC) platforms. In this context, a central problem is qubit allocation, which maps logical qubits onto physical qubits while optimizing objectives such as circuit depth, communication overhead, or error rates. This problem is NP-hard and closely related to quadratic assignment and complex scheduling, requiring advanced algorithmic techniques and massive parallelism. We investigate both exact and heuristic approaches. Exact methods rely on ultra-scale parallel branch-and-bound algorithms to compute reference optimal solutions for moderate-size circuits, while parallel metaheuristics and hybrid AI-based methods address larger instances where exact resolution is infeasible.

On the other hand, emerging quantum computing paradigms offer new opportunities for addressing hard optimization problems and related machine learning tasks. In fact, beyond optimization for NISQ systems, we also explore optimization using NISQ systems, with the long-term objective of integrating quantum devices as accelerators within hybrid quantum–classical workflows. Although current hardware remains limited, this perspective naturally connects combinatorial optimization, HPC, and quantum computing, and prepares the ground for future scalable hybrid optimization frameworks. In particular, quantum annealers and gate-based quantum algorithms provide novel frameworks for solving optimization problems formulated as Quadratic Unconstrained Binary Optimization (QUBO) models. These approaches encompass a broad class of algorithms, ranging from quantum-inspired methods to fully quantum and hybrid classical–quantum techniques, such as the Quantum Approximate Optimization Algorithm (QAOA).

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, Geoffrey 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 funding, it demonstrates the impactful potential of our team and the significant value our research can generate for both the economic and social environment.

6 Highlights of the year

- El-Ghazali Talbi chaired the 38th European Conference on Combinatorial Optimization (ECCO XXXVIII), which was held this year in Marrakech, Morocco. ECCO is affiliated with one of the largest working groups of the Association of European Operational Research Societies (EURO), a regional grouping within the International Federation of Operational Research Societies (IFORS), established in 1975 with the aim of promoting Operational Research throughout Europe.

- Abdelmoiz Zakaria Dahi received an Outstanding Reviewer Award at the flagship ACM GECCO 2025 (Genetic and Evolutionary Computation Conference) in recognition of his high-quality reviews for the Evolutionary Combinatorial Optimization and Metaheuristics (ECOM) track, one of the largest tracks of the conference.

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 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/jgmys/permutationbb

Publication: hal-03689608

Contact: Nouredine Melab

Participants: Jan Gmys, Nouredine Melab, Mohand Mezmaz, 2 anonymous participants

7.1.2 pBB-chpl

Name: Parallel Branch-and-Bound using Chapel

Keywords: Optimization, Parallel computing, Data parallelism, GPU, Combinatorics, Distributed computing

Scientific Description: pBB-chpl is a Chapel-based software platform designed for ultra-scale parallel Branch-and-Bound (B&B) computations. Unlike its predecessor, pBB, which targets permutation problems and follows the MPI+X implementation approach, pBB-chpl is generic and adopts the PGAS (Partitioned Global Address Space) paradigm to enhance productivity.

At its core, pBB-chpl features a scalable data structure called distBag_DFS, combined with a multi-level work-stealing mechanism. This combination is encapsulated within the DistributedBag module and seamlessly integrated into the Chapel programming language, making it well-suited for exascale computing.

pBB-chpl already supports a wide range of optimization problems, including binary knapsack, quadratic assignment (with qubit allocation instantiations), flowshop scheduling, and N-Queens. In addition, pBB-chpl offers multiple parallel B&B skeletons, ensuring flexibility across various architectures, including multi-core and GPU-powered desktops, laptops, commodity clusters, and high-end supercomputers. The platform and its comprehensive documentation are open-source and freely accessible on GitHub.

Functional Description: pBB-chpl is a Chapel-based software platform designed for ultra-scale parallel Branch-and-Bound (B&B) computations. Unlike its predecessor, pBB, which targets permutation problems and follows the MPI+X implementation approach, pBB-chpl is generic and adopts the PGAS (Partitioned Global Address Space) paradigm to enhance productivity.

At its core, pBB-chpl features a scalable data structure called distBag_DFS, combined with a multi-level work-stealing mechanism. This combination is encapsulated within the DistributedBag module and seamlessly integrated into the Chapel programming language, making it well-suited for exascale computing.

pBB-chpl already supports a wide range of optimization problems, including binary knapsack, quadratic assignment (with qubit allocation instantiations), flowshop scheduling, and N-Queens. In addition, pBB-chpl offers multiple parallel B&B skeletons, ensuring flexibility across various architectures, including multi-core and GPU-powered desktops, laptops, commodity clusters, and high-end supercomputers. The platform and its comprehensive documentation are open-source and freely accessible on GitHub.

URL: https://github.com/Guillaume-Helbecque/P3D-DFS

Publications: tel-04902137, hal-05449040, hal-05267434, hal-04165491

Contact: Guillaume Helbecque

7.1.3 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 shemes, MO diversity assignment shemes, 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

7.1.4 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

7.1.5 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 Gérémie

7.1.6 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 Liefooghe

7.1.7 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

7.2 New platforms

7.2.1 SLICES-FR/GRID'5000 testbed: major achievements in 2025

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 "CornelIA" has been accepted (2021-2027).

  Since late 2023, Bilel Derbel took over Nouredine Melab as the scientific leader. 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, Nouredine Melab is member of the SLICES-FR steering committe for the University of Lille. Bilel Derbel is the site leader of the Lille site at SLICES-FR with a strong involvement in the site leader committe, as well as on the mannaging aspects of the SLICES-FR site in Lille. During 2024, two clusters have been renewed and are now available for the SLICES-FR users. In 2025, an engineer-dedicated position was renewed, and a new GPU cluster was acquired and is currently being installed.

- **URL:** Grid'5000/SLICES-FR

8 New results

During the year 2025, we have addressed different issues/challenges related to the three lines of our research program. The major contributions are summarized in the following sections.

8.1 Decomposition-based Optimization

We report two major contributions related to decomposition-based techniques in the objective space targeting respectively: (1) constrained multi-objective continuous optimization problems, and (2) unconstrained multi-objective combinatorial optimization problems.

8.1.1 Combining Penalty-based and Decomposition-based Approaches for Constrained Multi-objective Continuous Optimization

Participants: Saúl Zapotecas-Martínez *(INAOE, Mexico)*, Bilel Derbel *(contact person)*, Néstor García-Rojas *(INAOE, Mexico)*, Miguel Jiménez-Domínguez *(INAOE, Mexico)*, Raquel Díaz-Hernández *(INAOE, Mexico)*, Leopoldo Altamirano-Robles *(INAOE, Mexico)*, Carlos Coello Coello *(CINVESTAV, Mexico)*.

The Multi-Objective Evolutionary Algorithm based on Decomposition (MOEA/D) has emerged as a robust and computationally efficient framework for addressing complex optimization challenges. In recent years, there has been a significant focus on adapting MOEA/D to effectively tackle constrained multi-objective optimization problems. In this work, we introduce a number of enhancements to hybirdize the MOEA/D framework with the integration of penalty functions aimed at improving constraint handling. Specifically, in [27, 26], we introduce a novel dynamic penality function which adapts to current search status. In [25], we propose hybridize the MOEA/D concept with Particle Swarm Optimization ending up with a novel Multi-objective Particle Swarm Optimization (MOPSO) approach based on decomposition framework and integrating additionnal dynamlic constraint-handling mechanisms. To evaluate the effectiveness of our proposed approaches, we conduct extensive experiments using the widely recognized CEC'2009 continuous benchmark problems. Our methodology is rigorously compared against state-of-the-art multi-objective optimization algorithms, allowing for a comprehensive assessment of its performance. The experimental results show that our enhanced decomposition-based approach yields solutions that are not only competitive but, in specific instances, outperform those generated by the leading algorithms in the field. Additionally, we discuss the implications of our findings for future research and practical applications, highlighting the potential of using dynamic penalty functiosn to advance the state of the art in constrained multi-objective optimization. This work is a collaboration with colleagues and PhD Students at the INAOE institute, Mexico.

8.1.2 Combining Local search and Decomposition for Multi-objective Combinatorial Optimization

Participant: Bilel Derbel *(contact person)*.

In this work, we address multi-objective combinatorial optimization problems, which are characterized by the discrete nature of their search spaces and the need to simultaneously optimize several conflicting objective functions. Local search (LS) techniques are widely recognized as a cornerstone in the design of efficient algorithms for combinatorial optimization, due to their ability to exploit problem structure and intensify the search around high-quality solutions. In contrast, evolutionary approaches are particularly well suited to handling multiple objectives, with decomposition-based, dominance-based, and indicator-based paradigms being among the most prominent frameworks.

In [24], we focus on the hybridization of iterated local search (ILS) with decomposition-based evolutionary multi-objective optimization. More specifically, we consider a simple yet effective combination of ILS with the well-established MOEA/D framework, in which the original multi-objective problem is decomposed into a set of scalar subproblems through objective aggregation. This decomposition allows each subproblem to be tackled using a cooperative single-objective LS, while coordination among subproblems by performing perturbation and replacement in a cooperative manner hence promoting diversity and coverage of the Pareto front. The resulting hybrid approach leverages the intensification capabilities of ILS and the structured exploration induced by decomposition.

The proposed method is evaluated against standard evolutionary tehcniques, and its performance is assessed on a broad range of challenging binary MNK-landscape benchmark instances. Experimental results demonstrate the superiority of the hybrid decomposition-based approach, highlighting the benefits of accurately integrating ILS withing the MOEA/D framework.

8.2 ML-Assisted Optimization and Emerging Optimization Approaches

In this research axis, we present our contributions to ML-assisted and alternative hybrid optimization techniques along four main directions: (1) the optimization of deep neural network architectures and hyperparameters; (2) the design of neuromorphic optimization techniques built upon the emerging paradigm of Spiking Neural Networks (SNNs); (3) the development and analysis of novel fitness landscape analysis tools; and (4) the investigation of quantum-based optimization techniques for combinatorial optimization. Our contributions in each of these directions are discussed in detail in the following four subsections, and we conclude with a brief overview of additional related contributions.

8.2.1 Hyperparamter Optimization and Bayesian optimization in Automated Machine Learning

Participants: Francesco Zito *(University of Catania, Italy)*, El-Ghazali Talbi *(contact person)*, Claudia Cavallaro *(University of Catania, Italy)*, Vincenzo Cutello *(University of Catania, Italy)*, Mario Pavone *(University of Catania, Italy)*, Nathan Davouse *(contact person)*.

In this work, we explore the intersection of Automated Machine Learning techniques and optimization techniques, particularly on the optimization of Artificial Neural Networks through hyperparameter tuning. Artificial Neural Networks are in fact widely used across various fields; however, building and optimizing them presents significant challenges. For example, by employing an effective hyperparameter tuning, shallow neural networks might become competitive with their deeper counterparts.

In [20], we highlight the importance of Hyperparameter Optimization (HPO) in enhancing neural network performance. We examine various metaheuristic algorithms employed and, in particular, their effectiveness in improving model performance across diverse applications. Despite significant advancements in this area, a comprehensive comparison of these algorithms across different deep learning architectures remains lacking. This work aims to fill this gap by systematically evaluating the performance of metaheuristic algorithms in optimizing hyperparameters and discussing advanced techniques such as parallel computing to adapt metaheuristic algorithms for use in hyperparameter optimization with high-dimensional hyperparameter search space.

Additionnaly, in [23], we specifically focus on Large Language Models (LLMs). In fact, Fine-tuning these models for domain-specific applications is significantly constrained by the computational costs associated with their training. as such, we propose two complementary approaches to address the HPO challenge in LLM fine-tuning: Bayesian Optimization based on Gaussian Process (BO-GP) and Partition-Based Optimization (PBO). On the one hand, BO efficiently exploits historical knowledge to achieve optimal results within a limited number of evaluations, but its inherently sequential nature poses scalability challenges. On the other hand, PBO enables massive parallelization, making it more scalable but requiring significantly more evaluations to converge. To leverage their complementary strengths for optimizing expensive objective functions, we investigate these methods and propose a hybrid BO-PBO algorithm. This work represents a step toward harnessing the potential of parallel Bayesian Optimization-based algorithms for solving expensive optimization problems in exascale computing environments, which is tightly related to our third research axis.

8.2.2 On Neuromorphic Computing and Optimization

Participants: El-Ghazali Talbi *(contact person)*, Jorge Mario Cruz-Duarte *(contact person)*, Nathan Bouvier.

Neuromorphic computing (NC) introduces a novel paradigm called Spiking Neural Networks (SNNs), representing a major shift from traditional digital computing. NC leverages spiking neurons, adaptive synapses, event-driven processing, and biologically-inspired learning mechanisms to develop efficient, brain-like systems optimized for real-time, parallel processing and low power consumption. In this respect, we conducted a number of investigation at with the aim of leveraging NC for designing innovative optimization algorithms. This is summarized in the following.

- In [39], we investigate the modelling and implementation of optimization algorithms and particularly metaheuristics using the NC paradigm as an alternative to Von Neumann architectures, leading to breakthroughs in solving optimization problems. Notice that our work departs from previous the research trend in NC which has concentrated on machine learning applications and neuroscience simulations. As such, we discuss Neuromorphic-based metaheuristics (Nheuristics) which are supposed to be characterized by low power, low latency and small footprint. Since NC systems are fundamentally different from conventional Von Neumann computers, several challenges are posed to the design and implementation of Nheuristics. A guideline based on a classification and critical analysis is conducted on the different families of metaheuristics and optimization problems they address. We also discuss future directions that need to be addressed to expand both the development and application of Nheuristics.

- In [30], we propose an algorithm that integrates the principles of evolutionary algorithms (EAs) with NC to create efficient and energy-aware metaheuristics. The proposed neuromorphic EA (NEVA) has been mapped on a SNN which involves defining the neuron model, information encoding, network architecture, and learning rules. To our knowledge this is the first EA designed using the NC paradigm. Computational experiments on QUBO, 3-SAT, and knapsack problems show the efficiency of the proposed NEVA algorithm. By designing a neuromorphic memetic algorithm that combines EAs with local search, the results have been improved both in terms of solution quality and search time.

- In [38], we present NeurOptimiser, a fully spike-based optimisation framework that materialises the neuromorphic-based metaheuristic paradigm through a decentralised NC system. The proposed approach comprises a population of Neuromorphic Heuristic Units (NHUs), each combining spiking neuron dynamics with spike-triggered perturbation heuristics to evolve candidate solutions asynchronously. The NeurOptimiser's coordination arises through native spiking mechanisms that support activity propagation, local information sharing, and global state updates without external orchestration. We implement this framework on Intel's Lava platform, targeting the Loihi 2 chip, and evaluate it on the noiseless BBOB suite up to 40 dimensions. We deploy several NeurOptimisers using different configurations, mainly considering dynamic systems such as linear and Izhikevich models for spiking neural dynamics, and fixed and Differential Evolution mutation rules for spike-triggered heuristics. Although these configurations are implemented as a proof of concept, we document and outline further extensions and improvements to the framework implementation. Results show that the proposed approach exhibits structured population dynamics, consistent convergence, and milliwatt-level power feasibility. They also position spike-native MHs as a viable path toward real-time, low-energy, and decentralised optimisation.

8.2.3 On Stochastic Operators and Fitness Landscapes in Combinatorial Optimization

Participants: Brahim Aboutaib, Sébastien Verel, Cyril Fonlupt, Bilel Derbel *(contact person)*, Arnaud Liefooghe, Belaïd Ahiod.

Stochastic operators are the backbone of many optimization algorithms. Besides the existing theoretical analysis that studies the asymptotic runtime of such algorithms, characterizing their performance using fitness landscape analysis is far away. The fitness landscape approach proceeds by considering multiple characteristics to understand and explain an optimization algorithm's performance or the difficulty of an optimization problem. In particular, a landscape-oriented approach can be combined with ML-based appraoches to tackle high-level automated tasks such as algorithm perfromance prediction or algorithm selection.

In [13], we analyze the fitness landscapes of stochastic operators by focusing on the number of local optima and their relation to the optimization performance. The search spaces of two combinatorial problems are studied: the NK-landscape and the Quadratic Assignment Problem, using binary string-based and permutation-based stochastic operators. The classical bit-flip search operator is considered for binary strings, and a generalization of the deterministic exchange operator for permutation representations is devised. We study small instances, ranging from randomly generated to real-like instances, and large instances from the NK-landscape. For large instances, we propose using an adaptive walk process to estimate the number of locally optimal solutions. Given that stochastic operators are usually used within population and single-solution-based evolutionary optimization algorithms, we contrast the performance of the -EA, and an Iterated Local Search, versus the landscape properties of large size instances of the NK-landscapes. Our analysis shows that characterizing the fitness landscapes induced by stochastic search operators can effectively explain the optimization performances of the algorithms under consideration.

8.2.4 On Quantum Optimization

Participants: Zakaria Abdelmoiz Dahi *(contact person)*, Francisco Chicano *(University of Malaga, Spain)*, Gabiel Luque *(University of Malaga, Spain)*, Rodrigo Gil-Merino *(University of Malaga, Spain)*, Iván Delgado Alba *(University of Malaga, Spain)*, Ivica Turkalj *(Fraunhofer Institute of Industrial Mathematics)*, Tom Ewen *(Fraunhofer Institute of Industrial Mathematics)*, Pascal Halffmann *(Fraunhofer Institute of Industrial Mathematics)*, Janik Maciejewski *(Lebensversicherung AG)*, Michael Trebing *(Fraunhofer Institute of Industrial Mathematics)*, Bilel Derbel.

Quantum computers leverage the principles of quantum mechanics to do computation with a potential advantage over classical computers. While a single classical computer transforms one particular binary input into an output after applying one operator to the input, a quantum computer can apply the operator to a superposition of binary strings to provide a superposition of binary outputs, doing computation apparently in parallel. This feature allows quantum computers to speed up the computation compared to classical algorithms. Unsurprisingly, quantum algorithms have been proposed to solve optimization problems in quantum computers. Furthermore, a family of quantum machines called quantum annealers are specially designed to solve optimization problems. In this respect, quantum computing provides a number of possibilities to design new powerful optimization techniques. However, the community still lacks both a practical and theoretical understanding of the strength and weaknesses of quatum optimization techniques. In this repspect, we constributed the following:

- In [14], we provide an introduction to quantum optimization from a practical point of view while specifically focusing on combinatorial optimization domains. We introduce the reader to the use of quantum annealers and quantum gate-based machines to solve optimization problems. Besides, in [33], we present a systematic literature review and a public web repository QoverC of the existing Quantum Computer Simulators (QCS) for quantum computation in general, and the leading ones for optimisation in particular. This can be viewed as the largest QCS study to date, where we include 199 web, desktop, and hybrid simulators, over 22 programming languages. We also provide a comprehensive comparison spanning over 10 metrics.
- In [21], we focus on QAOA a hybrid quantum-classical algorithm to solve optimization problems in gate-based quantum computers. QAOA is based on a variational quantum circuit that can be interpreted as a discretization of the annealing process that quantum annealers follow to find a minimum energy state of a given Hamiltonian. This ensures that QAOA must find an optimal solution for any given optimization problem when the number of layers, p, used in the variational quantum circuit tends to infinity. In practice, the number of layers is usually bounded by a small number. This is a must in current quantum computers of the NISQ era, due to the depth limit of the circuits they can run to avoid problems with decoherence and noise. We show mathematical evidence that QAOA requires exponential time to solve linear functions when the number of layers is less than the number of different coefficients of the linear function n. We conjecture that QAOA needs exponential time to find the global optimum of linear functions for any constant value of p, and that the runtime is linear only if $p \geq n$. We then conclude that we need new quantum algorithms to reach quantum supremacy in quantum optimization and disucss few alternatives.
- In [40], we present methodological improvements to variational quantum algorithms (VQAs) for solving multicriteria optimization problems. First, we reformulate the parameter optimization task

of VQAs as a multicriteria problem, enabling the direct use of classical algorithms from various multicriteria metaheuristics. This hybrid framework outperforms the corresponding single-criteria VQAs in both average and worst-case performance across diverse benchmark problems. Second, we propose a method that augments the hypervolume-based cost function with coverage-oriented indicators, allowing explicit control over the diversity of the resulting Pareto front approximations.

- In [22], we present an ML-based pipeline, allowing users to choose the appropriate moment to perform a given computation based on the estimation of the Jensen-Shannon divergence between the noisy and ideal distributions of quantum sampling. This includes (I) an extract-transform-load data module, (II) an ML unit for quantum features forecasting and error prediction, and (III) a web-based visualisation unit. The pipeline was built/tested using 3.5 months of calibration data from three real 127-qubit IBM quantum machines.

8.2.5 Other Contributions

Participants: R. Ragonnet, A. E. Hughes, D. S. Shipman, M. T. Meehan, A. S. Henderson, G. Briffoteaux, Nouredine Melab *(contact person)*, D. Tuyttens, E. S. Mcbryde, J. M. Trauer, Bohdan Ivaniuk-Skulskyi, Nadiya Shvai, Amir Nakib, El-Ghazali Talbi *(contact person)*, Sune Nielsen, Grégoire Danoy, Wiktor Jurkowski, Roland Krause, Reinhard Schneider, Pascal Bouvry.

In addition to the previous described work, we also contributed the following work which is tightly related to the design of intelligent and learning-based optimization techniques:

- Following our previous joint publications with the Monash University (Australia) on parallel surrogate-based optimization for Covid-19 epidemics control, we extended our contributions in [16], to the study of the impact of school closure. We used a mathematical model to simulate the COVID-19 epidemics of 74 countries, incorporating observed data from 2020 to 2022 and historical school closure timelines. The conclusion of the study is that closures generally reduced peak hospital occupancy and deaths, though a few countries saw increased mortality due to shifts in immunity and infection age distribution.

- Video anomaly detection (VAD) plays a critical role in identifying rare and unusual events in video streams, with applications ranging from surveillance to industrial monitoring. However, the generalization of VAD models to diverse datasets and anomaly types remains a challenge due to the limited amount of training data. In [32], we propose novel generalization techniques for the state-of-the-art transformer-based model, AnomalyClip. Our approach leverages multimodal data mixing, combining external datasets with textual descriptions to generate pseudo-anomaly samples through Adaptive Instance Normalization and Gaussian blending. Experimental evaluations on benchmarks such as ShanghaiTech, UCF-Crime, and XD-Violence demonstrate the efficacy of our techniques, achieving significant improvements in area under the curve metrics. This work highlights the potential of training-focused strategies to improve the robustness and scalability of VAD systems in high-performance computing contexts, which also relates to our third research axis.

- Protein structure prediction is an essential step in understanding the molecular mechanisms of living cells with widespread application in biotechnology and health. The inverse folding problem (IFP) of finding sequences that fold into a defined structure is in itself an important optimization problem at the heart of rational protein design. In [34], a multi-objective genetic algorithm (MOGA) using the diversity-as-objective (DAO) variant of multi-objectivization is presented, which optimizes the secondary structure similarity and the sequence diversity at the same time and hence searches deeper in the sequence solution space. To validate the final optimization results, a subset of the best sequences was selected for tertiary structure prediction. Comparing secondary structure annotation and tertiary structure of the predicted model to the original protein structure demonstrates that relying on fast approximation during the optimization process permits to obtain meaningful sequences.

8.3 Ultra-scale Parallel Optimization

During the year 2025, we have made contributions with respect to three main research directions in our parallel optimization axis: (1) large scale parallel optimization for continuous blackbox problems, (2) Scalable and Portable GPU-Accelerated Branch-and-Bound Algorithms for Heterogeneous Multi-GPU Systems, and (3) Ultra-scale Optimization for Qubit Allocation in NISQ Quantum Systems. Our contributions in each research direction are discussed in more details in the following.

8.3.1 Massively Parallel Continous Optimization with CMA-ES on the Fugaku Supercomputer

Participants: David Redon, Pierre Fortin, Bilel Derbel *(contact person)*, Miwako Tsuji *(RIKEN R-CCS, Japon)*, Mitsuhisa Sato *(RIKEN R-CCS, Japon)*.

The Increasing Population Covariance Matrix Adaptation Evolution Strategy (IPOP-CMA-ES) algorithm is a reference stochastic optimizer dedicated to blackbox continuous optimization, where no prior knowledge about the underlying problem structure is available. In [17], we focus on accelerating IPOP-CMA-ES using high-performance computing and parallelism for solving large-scale optimization problems on large scale compute platforms. In collaboration with the RIKEN R-CCS, Japan, we manage to speeding up both the linear algebra operations and the function evaluations of IPOP-CMA-ES on the Fugaku japenese supercomputer; there-by, contributing the following:

- We first show how the CMA-ES linear algebra operations can be accelerated using BLAS and LAPACK routines. This requires the rewrite of some of these operations in order to introduce more efficient BLAS routines.
- We present two parallel strategies for IPOP-CMA-ES to fully exploit a large number of CPU cores (up to several thousands). Such a number of CPU cores implies multiple compute nodes in distributed memory, each node being composed of multiple cores in shared memory. The goal here is to leverage large-scale parallelism (via multiple nodes) to benefit from the increasing number of (parallel) evaluations in IPOP-CMA-ES. The first strategy performs descents in the same order of population size as the original IPOP-CMA-ES, while the second strategy concurrently processes descents of different population sizes.
- We thoroughly compare hybrid MPI+OpenMP implementations of our two strategies on 6144 cores (128 AFX nodes) of the supercomputer Fugaku in order to determine which one is the most relevant on such a large-scale parallel architecture. We also present results obtained on top of 512 Intel Xeon cores in order to fairly support our findings when using a more conventional HPC compute cluster. Our empirical comparisons are performed using the reference BBOB (Black-Box Optimization Benchmarking) benchmark configured with various dimensions and various function evaluation costs. Accordingly, we conduct a comprehensive analysis to assess the impact of the considered parallel strategies on both performance and solution quality, as a function of the target function, problem dimensionality, and evaluation costs

8.3.2 Scalable and Portable GPU-Accelerated Branch-and-Bound Algorithms for Heterogeneous Multi-GPU Systems.

Participants: Guillaume Helbecque *(contact person)*, Nouredine Melab *(contact person)*, Ezhilmathi Krishnasamy *(Univ. Luxembourg)*, Tiago Carneiro *(IMEC, Belgium)*, Pascal Bouvry *(Univ. Luxembourg)*, Ivan Tagliaferro, Grégoire Danoy *(Univ. Luxembourg)*.

Branch-and-Bound (B&B) algorithms are central to solving exact combinatorial optimization problems, but their irregular and dynamic search patterns make efficient parallelization challenging. Modern high-performance computing platforms are increasingly heterogeneous, relying on GPU accelerators from multiple vendors. Designing scalable B&B algorithms for such systems requires balancing performance, portability, and programmability, while efficiently handling irregular workloads. In [15, 28, 29], we explored GPU-accelerated B&B designs that leverage pool-based parallelism and dynamic load balancing, highlighting the trade-offs between high-level PGAS programming and low-level GPU implementations.

- Portable PGAS-based GPU-accelerated Branch-and-Bound Algorithms at Scale.
  In [15], we proposed a high-level, portable B&B algorithm implemented in the Chapel language using the Partitioned Global Address Space (PGAS) model. By combining a pool-based search strategy with dynamic load balancing, the approach handles the irregular workloads inherent to B&B while maintaining portability across GPU architectures. The algorithm was evaluated on the N-Queens and permutation flowshop scheduling problems, demonstrating strong performance and cross-platform portability. Scaling experiments on a TOP500 pre-exascale supercomputer using up to 1,024 GPUs highlight the potential of high-level PGAS programming to exploit large-scale heterogeneous systems efficiently.
- A Portable Branch-and-Bound Algorithm for Cross-Architecture Multi-GPU Systems.
  Building on the same foundation, we proposed in [28, 29] a low-level C implementation using CUDA and HIP to fully exploit NVIDIA and AMD GPU architectures. It introduces an optimized multi-pool data structure and dynamic load balancing tailored for multi-GPU setups, along with GPU-specific optimizations for peak performance. Experiments on the permutation flowshop scheduling problem with up to 8 GPUs demonstrate significantly improved performance and scalability compared to the PGAS-based implementation, illustrating the trade-offs between portability, programmability, and high efficiency in heterogeneous GPU environments.

Together, these contributions show a progression from high-level, portable designs to low-level, performance-optimized implementations, providing both practical guidance and conceptual insights for building scalable B&B algorithms on modern heterogeneous multi-GPU supercomputers.

8.3.3 Ultra-scale Optimization for Qubit Allocation in NISQ Quantum Systems

Participants: Jean-Philippe Valois *(contact person)*, Guillaume Helbecque *(contact person)*, Nouredine Melab *(contact person)*, Jérôme Rouzé *(contact person)*, Jan Gmys, Daniel Tuyttens.

Qubit allocation and mapping are critical steps in adapting abstract quantum circuits to real quantum hardware, particularly for Noisy Intermediate-Scale Quantum (NISQ) devices with limited qubit connectivity and high error rates. Efficient solutions must address both the combinatorial complexity of these problems and the performance limitations of current computing platforms. In [19, 18], we investigated exact and heuristic approaches to qubit placement, showing how advanced algorithmic techniques combined with parallel computing can enhance scalability, reduce circuit depth, and minimize execution errors.

- Efficient and Scalable Branch-and-Bound Algorithm for Exact Qubit Allocation.
  In [19], we formulated qubit allocation as a permutation-based quadratic assignment problem and develops a branch-and-bound algorithm for its exact resolution. A refined sequential implementation achieves significantly faster runtimes than prior exact approaches, establishing a new state-of-the-art. Building on this foundation, a parallel implementation leverages both intra-node and inter-node parallelism on HPC infrastructures. Experimental results demonstrate near-linear strong scaling within nodes and substantial distributed scalability across multiple nodes. Using this approach, the method produces reference optimal solutions for benchmark circuits of up to 26 qubits—far beyond previously reported limits—showing that large-scale parallelization can significantly extend the reach of exact qubit allocation methods.
- A Parallel Memetic Algorithm for Qubit Mapping on Noisy Intermediate-Scale Quantum Machines.
  Complementing the exact approach, in [18], we introduced a parallel hybrid metaheuristic for qubit mapping (PMA-QM). PMA-QM is a memetic algorithm that combines a genetic algorithm with a local

search metaheuristic to optimize the placement of logical qubits onto physical qubits while respecting hardware connectivity constraints, minimizing circuit depth, and reducing error rates. A fine-tuned parallel model accelerates this computationally intensive hybrid approach, and problem-specific knowledge is incorporated to improve solution quality. Experiments on medium-to-large scale quantum circuits demonstrate that PMA-QM consistently outperforms the SWAP-based BidiREctional (SABRE) algorithm, a reference heuristic for qubit allocation implemented in the IBM Qiskit framework, delivering high-quality solutions where exact methods are computationally infeasible.

Together, these contributions illustrate a complementary optimization strategy for NISQ quantum circuits: ultra-scale exact branch-and-bound methods push the boundaries of optimal qubit allocation for small-to-medium circuits, while parallel memetic heuristics provide practical, high-quality solutions for larger circuits, enabling more efficient and reliable execution on NISQ devices.

8.3.4 Other contributions

Participants: Jean-Philippe Valois *(contact person)*, Nouredine Melab *(contact person)*, Thomas Firmin.

We conclude this section by highlighting our work on the design of a parallel island genetic algorithm for triangle-based Image Reconstruction. Specifically, in [31], we proposed a parallel Island Model Genetic Algorithm (IMGA) that reconstructs images with fixed colored triangles. In our appraoch, multiple subpopulations evolve with periodic migration, improving convergence and achieving up to 50% better reconstruction than standard GAs or prior hybrid methods.

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 enerygy consumption and production.

9.1.2 Participation in other International Programs

MoU RIKEN R-CCS / Japan

Participants: Bilel Derbel, David Redon.

Title: Memoremdum of Understanding

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

Date/Duration: 2021-2026

Additionnal 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

- Daniel Tuyttens (Univ. Mons, Belgium)
- Grégoire Danoy (Univ. Luxembourg, Luxembourg)

9.2.2 Visits to international teams

Zakaria Abdelmoiz Dahi

Visited institution: University of Naples, Quantum Computing and Smart systems laboratory (QUASAR)

Country: Italy

Dates: Nov 2025

El-Ghazali Talbi

Visited institution: Universidad Elche

Country: Spain

Dates: June 2025

El-Ghazali Talbi

Visited institution: EMI - University of Rabat

Country: Morocco

Dates: March 2025, May 2025

9.3 European initiatives

9.3.1 Other european programs/initiatives

Participant: El-Ghazali Talbi.

- 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.

9.4 National initiatives

9.4.1 ANR

- Bilateral ANR-NSF France/USA PRCI TunnelOPT (2024-2027, Grant: 562K€, PI: Bilel Derbel) in collaboration with Colorado State University (Co-PI: Darrell 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 fondamental 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 PRC EVARISTE (2024-2028, Grant: 493K€, WP PI: Bilel Derbel) in collaboration with Université Angers (ANR PI: Adrien 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 (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 (2023-2027, Grant: 207K€ for Bonus, PI: Nouredine Melab) in collaboration with University of Luxembourg (Co-PI: Grégoire 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 CornelIA) 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. El-Ghazali Talbi and Nouredine 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 Nouredine Melab is the co-PI. The MesoNet infrastucture 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.

9.5 Regional initiatives

Participants: Bilel Derbel, Nouredine Melab.

- CPER CornelIA (2021-2027, Grant: 800K€ in 2023/24 and 160K€ in 2025/26): 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. Bilel Derbel took over Nouredine Melab the responsability of the infrastructure managment and its coordination with other partners starting from late 2023. He is member of the CornelIA executive board.

10 Dissemination

10.1 Promoting scientific activities

10.1.1 Scientific events: organisation

General chair, scientific chair

- El-Ghazali Talbi (Conference Chair): European Conference on Combinatorial Optimization ECCO, 2025.
- El-Ghazali Talbi (Steering committee Chair): Intl. Conf. on Optimization and Learning (OLA), 2025.
- El-Ghazali Talbi (Steering committee): IEEE Workshop Parallel Distributed Computing and Optimization (IPDPS/PDCO), 2025.
- El-Ghazali Talbi (Steering committee): Intl. Conf. on Metaheuristics and Nature Inspired Computing (META), 2025.
- Bilel Derbel (workshop co-chair): Decomposition Techniques in Evolutionary Optimization (DTEO), workshop affiliated to ACM GECCO 2025.
- Bilel Derbel (special session co-chair): Advances in Decomposition based Evolutionary Multi-objecvtive Optimization (ADEMO), sepecial session at CEC/WCCI 2025.
- Abdelmoiz Zakaria Dahi (special session co-chairs): sepecial session on Quantum AI at CEC 2025.
- Nouredine Melab (Seminar chair): 12th edition of the seminar series related to Simulation and HPC at the University of Lille. The edition icludes 4 seminars from IMEC (Belgium), University of Luxembourg, Safran Tech and AirBus.

Member of the organizing committees

- Abdelmoiz Zakaria Dahi: member of the organizing committe. ACM GECCO 2025, Malaga, Spain.
- El-Ghazali Talbi: member of the organizing committe. ECCO'2025, European Conference on Combinatorial Optimization (EURO Society), May 2025.
- El-Ghazali Talbi: member of the organizing committe. OLA'2025, Int. Conf. on Optimization and Learning, April 2025.

10.1.2 Scientific events: selection

Member of the conference program committees

- The European Conference on Artificial Intelligence (ECAI).
- The International Joint Conference on Neural Networks (IJCNN).
- Area chair NeurIPS, Thirty-nine Conference on Neural Information Processing Systems.
- 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.
- PAW-ATM 2025 (Parallel Applications Workshop – Alternatives to MPI+X) in conjunction with SIGHPC/IEEE/TCHPC Intl. Conf. for High Performance Computing, Networking, Storage, and Analysis (SC'2025).

10.1.3 Journal

Member of the editorial boards

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

Reviewer - reviewing activities

- IEEE Transactions on Evolutionary Computation (TEVC).
- ACM Computing Surveys, ACM.
- Engineering Applications of Artificial Intelligence (EAAI), Elsevier.
- Future Generation Computer Systems (FGCS), Elsevier.
- Journal of Computational Science (JoCS), Elsevier.
- International Journal of Imaging Systems and Technology (IMA), Wiley.

10.1.4 Invited talks

- El-Ghazali Talbi: "Neuromorphic-based optimization algorithms", ISC'2025 IA meets decision making, Catania, June 2025.
- El-Ghazali Talbi: "Neuromorphic computing and optimization: A two way synergy", Universidad Muguel Hernandez, Elche, Spain, Dec 2025.
- Abdelmoiz Zakaria Dahi: Talk at the Quantum inforation working group of the Univeristy of Lille. Quantum vs Classical Computation: Synergies and Boundaries. Octobre 2025:

10.1.5 Leadership within the scientific community

- Nouredine Melab: Member of the steering committee of the SLICES-FR, a large-scale experimental research infrastructure in computer science, focusing on distributed computing and networking, from wireless and IoT to cloud computing and HPC. Since 2024.
- Nouredine Melab: Member of the General Assembly of the MesoNet Equipex+ project (decision-making body appointing the Scientific, Steering, and User Committees). Since 2021.
- Bilel Derbel: Scientific leader of SLICES-FR testbed at Lille, a large-scale experimental research infrastructure in computer science, focusing on distributed computing and networking, from wireless and IoT to cloud computing and HPC, since 2023.
- El-Ghazali Talbi: Co-president of the working group "META: Metaheuristics - Theory and applications", GDR RO and GDR MACS.
- El-Ghazali Talbi: Co-Chair of the IEEE Task Force on Cloud Computing within the IEEE Computational Intelligence Society.

10.1.6 Scientific expertise

- Bilel Derbel: Expert reviewer for the HCERES - member of the evaluation committe of the UMR IRIT, 2025
- Bilel Derbel: Expert reviewer for the ANR PRCE program, CE23 – Artificial intelligence and Data Sciences, 2025.
- Bilel Derbel: Member of the PhD hiring committee of the UE Marie Skłodowska-Curie Innovative Training Networks (MSC ITN) Generation Quantum (GenQ)
- Nouredine Melab: Expert reviewer for the ANR PRCE program, CE25 – Software Science and Engineering – Multi-purpose communication networks, digital infrastructures, 2025.
- Nouredine Melab: Participation in the evaluation process for the election to the position of Researcher at the Jožef Stefan Institute, Ljubljana, Slovenia, September 2025.
- El-Ghazali Talbi: Expert reviewer, Fonds de la Recherche Scientifique (F.R.S-FNRS), Belgium, 2025

10.1.7 Research administration

- Bilel Derbel: Member of the Departmental Council of Computer Science Department of the Faculty of Science and Technology (FST), University of Lille, since 2025.
- Bilel Derbel: Member of the Scientific Board of the CRIStAL UMR laboratory, since late 2023.
- Bilel Derbel: Member of the Scientific Board for the MADIS doctoral school at the University of Lille, since 2022.
- Bilel Derbel: Coordinator of the research theme (GT) OPTIMA at the CRIStAL UMR laboratory, since late 2023.
- Nouredine Melab: Chargé de Mission of High Performance Computing and Simulation at Université de Lille, from 2010 to 2025.

10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

10.2.1 Teaching

Taught courses

- Master: Abdelmoiz Zakaria Dahi, Data Minining, 30h. Master in computer science, University of Lille, France.
- Bachelor: Abdelmoiz Zakaria Dahi, Algorithms and Data Structures, 36h, University of Lille, France.
- International Master : Nouredine Melab, Supercomputing, 45h ETD, M2, University of Lille, France.
- Master: Nouredine Melab, Operations Research, 60h ETD, M1, University of Lille, France.
- Master: Bilel Derbel, Algorithms and Complexity, 35h, M1, University of Lille, France.
- Master: Bilel Derbel, Optimization and machine learning, 24h, M1, University of Lille, France.
- Bachelor: Bilel Derbel, Algorithms and Data Structures, 36h, University of Lille, France.
- Engineering school: El-Ghazali Talbi, Advanced optimization, 36h, Polytech'Lille, University of Lille, France.
- Engineering school: El-Ghazali Talbi, Data mining, 36h, Polytech'Lille, University of Lille, France.
- Engineering school: El-Ghazali Talbi, Operations research, 60h, Polytech'Lille, University of Lille, France.
- Engineering school: El-Ghazali Talbi, Graphs, 25h, Polytech'Lille, University of Lille, France.

Teaching responsibilities

- Head of the international relations: El-Ghazali Talbi, Polytech'Lille, Université de Lille, France.
- Head of the international relations: Bilel Derbel, Computer Science Department, Faculty of Science and Technology, Université de Lille, France.
- Master leading: Nouredine 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

- HDR defense: Loïc Brevault (ONERA Palaiseau), Methodologies for multidisciplinary design analysis and optimization, and uncertainty quantification with aerospace applications. Supervisor: Nouredine Melab, Defended December 2^{nd}, 2025.
- PhD (cotutelle) defense [36]: Guillaume Helbecque, PGAS-based Parallel Branch-and-Bound for Ultra-Scale GPU-powered Supercomputers: Supervisors: Nouredine Melab (Université de Lille) and P. Bouvry (Université du Luxembourg). Defended January 10^{th}, 2025.
- PhD defense [35]: Thomas Firmin, Pulse neuron networks and parameter optimization for massively parallel GPU-powered clusters. Supervisors: El-Ghazali Talbi and P. Boulet (Emeraude Team, CRIStAL lab). defended Januray 2025.
- PhD defense [37]: Julie Keisler, Réseaux de neurones profonds pour la prédiction de séries spatio-temporelles. Supervisor: El-Ghazali Talbi, CIFRE with EDF. defended Januray 2025.
- PhD in progress: David Redon, Enabling Large Scale Computational Intelligence with HPC. Supervisors: Bilel Derbel and P. Fortin (Université de Lille). Started in Oct. 2020.

- PhD in progress: Jérôme Rouzé, Parallel Hybrid Metaheuristics for Qubit Allocation on NISQ Quantum Systems. Supervisors: Nouredine Melab (Université de Lille) and D. Tuyttens (Université de Mons, Belgium). Started in Nov. 2023.
- PhD in progress: Bohdan Ivaniuk, Automated design and multi-objective optimization of parallel deep networks for automatic detection in real-time. Supervisor: El-Ghazali Talbi. Started in 2023.
- PhD in progress: Mehdi El Khadiri, Exascale optimization using fractal-based decomposition. Supervisor: El-Ghazali Talbi. Started in 2024.
- PhD (cotutelle): Ivan Tagliaferro De Oliveira Tezoto, Exascale Exact Optimization based on the MPI+X Approach. Supervisors: Nouredine Melab (Université de Lille) and G. Danoy (Université du Luxembourg). Oct. 2024 to Nov. 2025.
- PhD in progress: Jean-Philippe Valois, Massively Parallel Exact Optimization for Qubit Allocation in Quantum Systems. Supervisor: Nouredine Melab, Started in Oct. 2025.
- PhD in progress: Francesco Cecere. Large-scale graybox tunneling for multi-objective optimization. Supervisor: Bilel Derbel, started September 2025.
- PhD in progress: Lander Argote, Towards quantum-utility multi-objective variational optimisers. Supervisors, Abdelmoiz Zakaria Dahi and Bilel Derbel. Started October 2025.

10.2.3 Juries

- Bilel Derbel (president of Jury), HDR defense of Raca TODOSIJEVIC: The power of change and simplicity in combinatorial optimization, University Polytechnique Haut-de-France, France. Garant : Prof. Abdelhakim Artiba
- Bilel Derbel (member of Jury), HDR defense of Mahmoud GOLABI: From Classical to Learning-Enhanced Optimization : Advancing Logistics and Interactive Multi-Objective Search, University Haute-Alsace, France. Garant : Prof. Lhassane Idoumghar.
- Nouredine Melab (Reviewer), PhD thesis of Jean-Sébastien Lerat: Design and Distributed Deployment of AI Models for Computer Vision and Industry 4.0 Applications, University of Mons, Belgium, August 2025 (private defense), November 2025 (public defense).
- Nouredine Melab (Reviewer), PhD Thesis of Zineb Ziani: AI and HPC Convergence for Enhanced Anomaly Detection, Université Paris-Saclay, defended January 2025.

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 HierarchicalBranch and Bound for Large ScaleUnreliable Environments'. In: *IEEE Trans. Computers* 63.9 (2014), pp. 2302–2315 (cit. on p. 10).

[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. 6).

[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. 9).

[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. 8, 9).

[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. URL: https://hal.science/hal-04073811 (cit. on p. 8).

[8] J. Gmys, M. Mezmaz, N. Melab and D. Tuyttens. '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. 7, 10).

[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. URL: https://hal.archives-ouvertes.fr/hal-02294201 (cit. on p. 9).

[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. 10).

[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. 7).

[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. URL: https://inria.hal.science/hal-01067662.

11.2 Publications of the year

International journals

[13] B. Aboutaib, S. Verel, C. Fonlupt, B. Derbel, A. Liefooghe and B. Ahiod. 'On stochastic operators, fitness landscapes, and optimization heuristics performances'. In: *Evolutionary Computation*. In Special Collection: CogNet 33.4 (1st Dec. 2025), pp. 459–484. DOI: 10.1162/evco.a.24. URL: https://hal.science/hal-05202575 (cit. on p. 19).

[14] F. Chicano, G. Luque, Z. A. Dahi and R. Gil-Merino. 'Combinatorial optimization with quantum computers'. In: *Engineering Optimization* 57.1 (31st Jan. 2025), pp. 208–233. DOI: 10.1080/0305215X.2024.2435538. URL: https://hal.science/hal-05465036 (cit. on p. 20).

[15] G. Helbecque, E. Krishnasamy, T. Carneiro, N. Melab and P. Bouvry. 'Portable PGAS-based GPU-accelerated Branch-and-Bound Algorithms at Scale'. In: *Concurrency and Computation: Practice and Experience* 37.25-26 (2nd Oct. 2025), e70321. DOI: 10.1002/cpe.70321. URL: https://hal.science/hal-05267434 (cit. on p. 23).

[16] R. Ragonnet, A. E. Hughes, D. S. Shipman, M. T. Meehan, A. S. Henderson, G. Briffoteaux, N. Melab, D. Tuyttens, E. S. Mcbryde and J. M. Trauer. 'Estimating the impact of school closures on the COVID-19 dynamics in 74 countries: A modelling analysis'. In: *PLoS Medicine* 22.1 (21st Jan. 2025), e1004512. DOI: 10.1371/journal.pmed.1004512. URL: https://hal.science/hal-05457913 (cit. on p. 21).

[17] D. Redon, P. Fortin, B. Derbel, M. Tsuji and M. Sato. 'Massively parallel CMA-ES with increasing population'. In: *Concurrency and Computation: Practice and Experience* 37.27-28 (25th Dec. 2025), e70362. DOI: 10.1002/cpe.70362. URL: https://hal.science/hal-04698736 (cit. on p. 22).

[18] J. Rouzé, N. Melab, J. Gmys and D. Tuyttens. 'A parallel memetic algorithm for qubit mapping on noisy intermediate-scale quantum machines'. In: *Engineering Applications of Artificial Intelligence* 161 (Dec. 2025), p. 112081. DOI: 10.1016/j.engappai.2025.112081. URL: https://hal.science/hal-05264970 (cit. on p. 23).

[19] J.-P. Valois, G. Helbecque and N. Melab. 'Efficient And Scalable Branch-and-Bound Algorithm for Exact Qubit Allocation'. In: *Future Generation Computer Systems* 179 (2026), p. 108342. DOI: 10.1016/j.future.2025.108342. URL: https://hal.science/hal-05449040 (cit. on p. 23).

[20] F. Zito, E.-G. Talbi, C. Cavallaro, V. Cutello and M. Pavone. 'Metaheuristics in automated machine learning: Strategies for optimization'. In: *Intelligent Systems with Applications* 26 (June 2025), p. 200532. DOI: 10.1016/j.iswa.2025.200532. URL: https://hal.science/hal-05464598 (cit. on p. 18).

International peer-reviewed conferences

[21] F. Chicano, Z. A. Dahi and G. Luque. 'The Quantum Approximate Optimization Algorithm Can Require Exponential Time to Optimize Linear Functions'. In: GECCO '25 Companion: Genetic and Evolutionary Computation Conference Companion. NH Malaga Hotel Malaga Spain, Spain: ACM, 14th July 2025, pp. 2403–2411. DOI: 10.1145/3712255.3734319. URL: https://hal.science/hal-05465066 (cit. on p. 20).

[22] Z. A. Dahi, F. Chicano, G. Luque and I. Delgado Alba. 'Optimising Quantum Calculations Reliability via Machine Learning: The IBM Case Study'. In: *Optimising Quantum Calculations Reliability via Machine Learning: The IBM Case Study*. IEEE QAI 2025 - IEEE International Conference on Quantum Artificial Intelligence. Naples, Italy, 2nd Nov. 2025. URL: https://hal.science/hal-05472988 (cit. on p. 21).

[23] N. Davouse and E.-G. Talbi. 'Bayesian and Partition-Based Optimization for Hyperparameter Optimization of LLM Fine-Tuning'. In: OLA2025 - International Conference in Optimization and Learning. Vol. 2808. Communications in Computer and Information Science. Dubai, United Arab Emirates: Springer Nature Switzerland, 3rd Jan. 2026, pp. 16–33. DOI: 10.1007/978-3-032-13589-6_2. URL: https://hal.science/hal-05464629 (cit. on p. 18).

[24] B. Derbel. 'A Simple Combination of Local Search and MOEAD for Combinatorial Multi-objective Optimization'. In: GECCO '25 Companion: Genetic and Evolutionary Computation Conference Companion. NH Malaga Hotel Malaga Spain, Spain: ACM, 14th July 2025, pp. 355–358. DOI: 10.1145/3712255.3726688. URL: https://hal.science/hal-05442276 (cit. on p. 17).

[25] N. García-Rojas, M. Jiménez-Domínguez, S. Zapotecas-Martínez, R. Díaz-Hernández, L. Altamirano-Robles and B. Derbel. 'Constrained Multi-Objective Optimization with dMOPSO: An Adaptive Penalty Approach'. In: 2025 IEEE Conference on Artificial Intelligence (CAI). Santa Clara, United States: IEEE, 5th May 2025, pp. 1280–1285. DOI: 10.1109/CAI64502.2025.00223. URL: https://hal.science/hal-05442304 (cit. on p. 17).

[26] N. García-Rojas, S. Zapotecas-Martínez, R. Díaz-Hernández, L. Altamirano-Robles and B. Derbel. 'Towards a Penalty Annealing Approach in MOEA/D for Constrained Multi-Objective Optimization'. In: GECCO '25 Companion: Genetic and Evolutionary Computation Conference Companion. NH Malaga Hotel Malaga Spain, Spain: ACM, 14th July 2025, pp. 1927–1934. DOI: 10.1145/3712255.3734307. URL: https://hal.science/hal-05467511 (cit. on p. 17).

[27] M. Jiménez-Domínguez, N. García-Rojas, S. Zapotecas-Martínez, B. Derbel and C. Coello Coello. 'A Dynamic Penalty Approach in MOEA/D for Constraint Handling in Multi-objective Problems'. In: 24th Mexican International Conference on Artificial Intelligence, MICAI. Vol. 16222. Lecture Notes in Computer Science. Guanajuato, Mexico: Springer Nature Switzerland, 20th Oct. 2026, pp. 73–85. DOI: 10.1007/978-3-032-09044-7_6. URL: https://hal.science/hal-05442314 (cit. on p. 17).

[28] I. Tagliaferro, G. Helbecque, E. Krishnasamy, N. Melab and G. Danoy. 'A Portable Branch-and-Bound Algorithm for Cross-Architecture Multi-GPU Systems'. In: *Euro-Par 2025 - Parallel Processing Workshops in 31th International European Conference on Parallel and Distributed Computing*. 23rd International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms. Dresde, Germany, 26th Aug. 2025. URL: https://hal.science/hal-05454431 (cit. on p. 23).

[29] I. Tagliaferro, G. Helbecque, E. Krishnasamy, N. Melab and G. Danoy. 'Performance and Portability in Multi-GPU Branch-and-Bound: Chapel Versus CUDA and HIP for Tree-Based Optimization'. In: *2025 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 39th IEEE International Parallel & Distributed Processing Symposium. Milano, Italy: IEEE, 13th Aug. 2025, pp. 1293–1295. DOI: 10.1109/IPDPSW66978.2025.00217. URL: https://hal.science/hal-05263968 (cit. on p. 23).

[30] E.-G. Talbi and N. Bouvier. 'NEVA: A Neuromorphic Evolutionary Algorithm'. In: OLA2025 - International Conference in Optimization and Learning. Vol. 2808. Communications in Computer and Information Science. Dubai, United Arab Emirates: Springer Nature Switzerland, 3rd Jan. 2026, pp. 64–79. DOI: 10.1007/978-3-032-13589-6_5. URL: https://hal.science/hal-05464618 (cit. on p. 19).

[31] J.-P. Valois, T. Firmin and N. Melab. 'A Parallel Island Genetic Algorithm for Triangle-based Image Reconstruction'. In: *2025 IEEE/ACS 22nd International Conference on Computer Systems and Applications (AICCSA)*. AICCSA 2025 - ACS/IEEE International Conference on Computer Systems and Applications. Doha, Qatar: IEEE, 5th Jan. 2026, pp. 1–8. DOI: 10.1109/AICCSA66935.2025.11315373. URL: https://hal.science/hal-05458101 (cit. on p. 24).

Conferences without proceedings

[32] B. Ivaniuk-Skulskyi, N. Shvai, A. Nakib and E.-G. Talbi. 'Enhancing Generalization in Video Anomaly Detection through Multimodal Data Mixing'. In: 2025 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). Milano, France: IEEE, 3rd June 2025, pp. 333–342. DOI: 10.1109/IPDPSW66978.2025.00058. URL: https://hal.science/hal-05464588 (cit. on p. 21).

Scientific book chapters

[33] Z. A. Dahi, E. Alba, R. Gil-Merino, G. Luque and B. Derbel. 'QoverC: A Profiler of Quantum Computer Simulators for Quantum Optimisation'. In: *Combinatorial Optimization using Quantum Computing - A Gentle Introduction*. 1st June 2026. URL: https://hal.science/hal-05473000. In press (cit. on p. 20).

[34] S. Nielsen, G. Danoy, W. Jurkowski, R. Krause, R. Schneider, E.-G. Talbi and P. Bouvry. 'Evolutionary Algorithms for the Inverse Protein Folding Problem'. In: *Handbook of Heuristics*. Springer Nature Switzerland, 17th Oct. 2025, pp. 1295–1319. DOI: 10.1007/978-3-032-00385-0_59. URL: https://hal.science/hal-05464577 (cit. on p. 21).

Doctoral dissertations and habilitation theses

[35] T. Firmin. 'Parallel hyperparameter optimization of spiking neural networks'. Université de Lille, 21st Jan. 2025. URL: https://theses.hal.science/tel-05109284 (cit. on p. 31).

[36] 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. 31).

[37] J. Keisler. 'Automated Deep Learning : algorithms and software for energy sustainability'. Université de Lille, 24th Jan. 2025. URL: https://theses.hal.science/tel-05097588 (cit. on p. 31).

Reports & preprints

[38] J. M. Cruz-Duarte and E.-G. Talbi. *NeurOptimisation: The Spiking Way to Evolve*. University of Lille; INRIA; INRIA Lille - Nord Europe, 2025. DOI: 10.48550/arXiv.2507.08320. URL: https://hal.science/hal-05464644 (cit. on p. 19).

[39] E.-G. Talbi. *Neuromorphic-based metaheuristics: A new generation of low power, low latency and small footprint optimization algorithms*. University of Lille; INRIA Lille - Nord Europe, 2025. DOI: 10.48550/arXiv.2505.16362. URL: https://hal.science/hal-05464639 (cit. on p. 19).

[40] I. Turkalj, T. Ewen, P. Halffmann, J. Maciejewski, M. Trebing and Z. A. Dahi. *Enhancing Variational Quantum Algorithms for Multicriteria Optimization*. 2025. DOI: 10.48550/arXiv.2506.22159. URL: https://hal.science/hal-05472980 (cit. on p. 20).

11.3 Cited publications

[41] 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. URL: http://dx.doi.org/10.1007/978-3-319-41508-6_1 (cit. on p. 11).

[42] B. Derbel, D. Brockhoff, A. Liefooghe 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. 8).

[43] B. Derbel, A. Liefooghe, 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. 10).

[44] J. Dreo, A. Liefooghe, 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 Sigevo. Lille / Virtual, France: ACM, July 2021, pp. 1522–1530. DOI: 10.1145/3449726.3463276. URL: https://hal-pasteur.archives-ouvertes.fr/pasteur-03220556 (cit. on p. 6).

[45] 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. 9).

[46] 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. 8).

[47] 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. URL: http://dx.doi.org/10.1007/978-3-319-67988-4_5 (cit. on pp. 7, 9).

[48] 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. 10).

[49] N. Shavit. 'Data Structures in the Multicore Age'. In: *Communications of the ACM* 54.3 (2011), pp. 76–84 (cit. on p. 9).

[50] 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. 10).

[51] 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. 7).

[52] 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. 10).

[53] 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. 8).