



HAL
open science

Multi-Objective and Cooperative Power Planning for Datacenter With On-Site Renewable Energy Sources

Léo Grange, Patricia Stolf, Georges da Costa, Paul Renaud-Goud

► **To cite this version:**

Léo Grange, Patricia Stolf, Georges da Costa, Paul Renaud-Goud. Multi-Objective and Cooperative Power Planning for Datacenter With On-Site Renewable Energy Sources. *IEEE Access*, 2022, 10, pp.104067-104092. 10.1109/ACCESS.2022.3210523 . hal-03907179

HAL Id: hal-03907179

<https://hal.science/hal-03907179>

Submitted on 3 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

RESEARCH ARTICLE

Multi-Objective and Cooperative Power Planning for Datacenter With On-Site Renewable Energy Sources

LÉO GRANGE, PATRICIA STOLF^{ID}, GEORGES DA COSTA^{ID}, AND PAUL RENAUD-GOUD

IRIT, University of Toulouse, 31062 Toulouse, France

Corresponding author: Georges Da Costa (georges.da-costa@irit.fr)

This work was supported in part by Agence Nationale de la Recherche (ANR) through the Context of the Project DATAZERO under Grant ANR-15-CE25-0012, and in part by Project DATAZERO2 under Grant ANR-19-CE25-0016.

ABSTRACT To reduce the ecological impact of datacenters and to operate them in areas with unreliable electrical grid, the use of on-site renewable energy sources (RESs) is actively studied nowadays. The main barrier to their wide adoption is related to the intermittent nature of common RESs, such as solar panels or wind turbines. Several works already demonstrated the suitability of energy storage devices and energy consumption scheduling to mitigate this issue. We propose to abstract such an infrastructure by two independent black-box systems: A power producer takes care of the RESs and storage devices; a power consumer manages the datacenter itself. An optimization module cooperates with these two systems to find the best power production and consumption plans for an upcoming time window. Each system having different goals, we tackle the multi-objective problem using a new evolutionary algorithm to find a set of ideal trade-offs. Evaluating the outcome of a power plan for an electrical or computing system is however usually costly. To reduce the number of evaluations of these objective functions, a common method is to use a cheap surrogate. Using time series properties, a novel surrogate method is presented. The performance of our approach is evaluated with maximization of quality of service and minimization of greenhouse emissions as respective objectives. Experiments are performed first using a simplified datacenter model with computable Pareto front, with the main findings replicated in a realistic model and a datacenter scheduler from the literature.

INDEX TERMS Renewable energy, datacenter, multi-objective optimization, surrogate method.

I. INTRODUCTION

The environmental impact of information and communication technologies is a major concern nowadays. With datacenters becoming more and more important through the increasing use of online platforms and more generally of various kinds of cloud computing, multiple type of research have been performed during this last decade to reduce their energy consumption. The global greenhouse gas (GHG) reduction goal along with the continuous decrease of cost of renewable energy sources (RESs) [16] also explain the recent scientific and industrial interest in using RESs to power datacenters [20], [25].

The associate editor coordinating the review of this manuscript and approving it for publication was Ramani Kannan^{ID}.

The main challenge of integrating RESs in computing infrastructures is to deal with the intermittent aspect of power produced by commonly used sources, such as photovoltaic panels or wind turbines. The global idea behind existing works is to bring the power consumption as close as possible from the production, by managing the workload and power state of the machines, by taking advantage of energy storage devices (ESDs) (e.g. batteries) or spatial location of geo-distributed datacenters [25]. The work presented in this paper is studying offline optimization of a single datacenter with on-site RESs and ESDs. However, it addresses the problem in a novel way by two means: separating energy production concerns from computing ones and encompassing its multi-objective nature by using Pareto-dominance guided optimization.

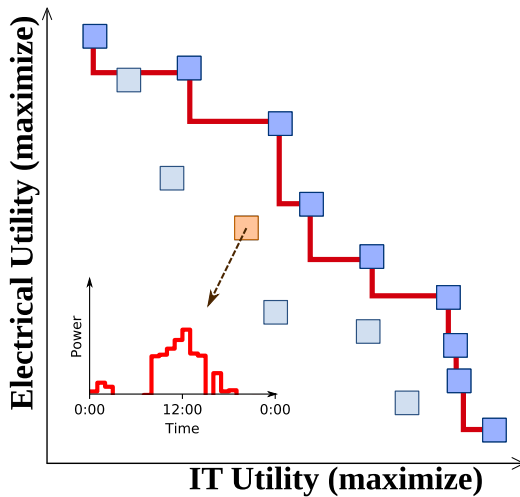


FIGURE 1. Illustration of the multi-objective optimization of power profiles. Each resulting solution consists in a power availability along a time window used as a constraint to find its objectives values.

Hence we propose to split the overall infrastructure of such a datacenter in two distinct parts viewed as *black-boxes* for an optimization process in-between. On the one hand, the electrical infrastructure management is in charge of responding to a power demand through the various power sources and storage. On the other hand, the computing infrastructure is in charge of scheduling the tasks and controlling various power leverage of the hardware. Each of these black-boxes has its own, internal objective. By taking a planning of power over an upcoming time window as an input, called *power profile* thereafter, each black-box outputs the corresponding objective value according to its optimization method and internal model. This black-box design is similar to the infrastructure considered along the ANR Datazero project [35] and leads to several benefits, detailed in section II.

The overall goal of the work presented here is to find possible time series of power (power profiles) which maximize the objective of both black-boxes (*i.e.* computing and electrical supply parts). Optimization of black-boxes functions is nevertheless challenging, especially in the case of required solution space (large time series). Additionally, the black-boxes studied here rely on complex models and are expected to be costly, such as solving a task scheduling optimization under power constraints. Hence, limiting the amount of costly evaluation of these black-boxes is desired.

Another challenge is the articulation of the objectives of the two black-boxes. The objective of the electrical part is expected to be related to the ecological impact of providing energy and the objective of computing part to be some quality of service metric or economical income.¹ Even if the precise meaning and unit of these objectives are known, minimizing

¹It is possible to use any other kind of objectives, but this is the original use case of the presented work and therefore is used as an example along this paper.

the ecological footprint and maximizing the quality of service are two antagonists or at least conflicting objectives.

Instead of reducing the problem to a mono-objective formulation similarly to existing works, we consider the two concerns are hardly miscible. Hence, we propose to tackle the bi-objective problem in terms of Pareto dominance. Figure 1 illustrates the result of the multi-objective optimization described here. Each solution is a possible power profile during a time window, which leads to a specific trade-off between the two objectives. With this multi-objective approach, we aim at providing a set of solutions that offers datacenter operators the freedom to choose the right trade-off between carbon footprint and immediate income according to their green commitment.

Several contributions are presented in this article.

- we design USPEA2, a multi-objective evolutionary algorithm integrating well with objective approximation techniques
- a novel surrogate method for black-box functions defined on time series is proposed, using Haar transform to perform dimensionality reduction
- an extensive performance evaluation is performed over several scenarios of renewable production based on realistic data from the National Renewable Energy Laboratory (NREL) [31], using a simplified infrastructure model
- the approach is validated using a datacenter model and scheduling algorithm from the literature [9], along with a workload based on the Google cluster traces

A description of the problem addressed by our approach are detailed in section II, followed by a description of related works and methods in section III. Our multi-objective optimization approach, USPEA2, is presented in section IV along with the two surrogate methods for time series defined objective functions. Section V covers the methodology of the experiments and describes the simplified model used for extensive evaluation. The results of the experiments are split between section VI for preliminary results used for tuning some parameters and section VII for the core results. They are interpreted and discussed along section VIII and finally section IX concludes this paper and explores future works.

II. PROBLEM STATEMENT

The optimal management of computing resources and workload with variable power constraint is known to be a hard problem. The same goes for the optimal engagement of multiple electrical sources and storage devices to maintain a stable power supply. Despite the high complexity of these optimization problems, a lot of works proposed more reasonable heuristics which give good enough results in practice, for either the computing or the management of the electrical sources. The biggest challenge to provide efficient solutions to the joint optimization probably lies directly in the complexity of both problems.

Instead of proposing new methods to solve the overall problem, the proposed approach keeps the two independent

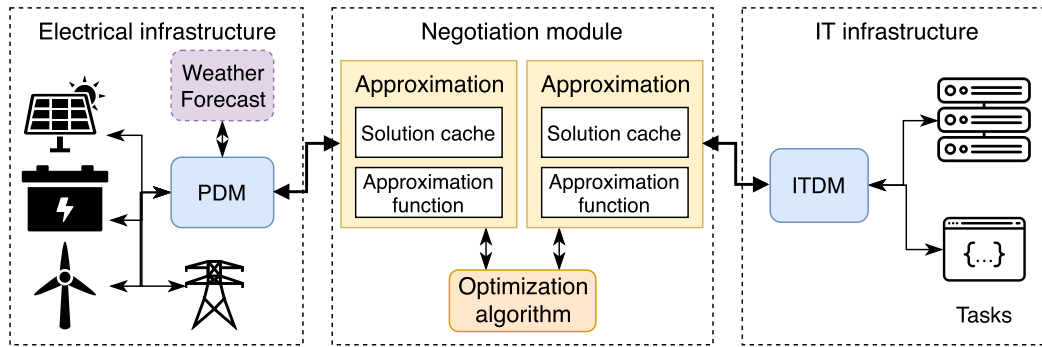


FIGURE 2. Overview of the datacenter and renewable sources infrastructure, using the proposed multi-objective optimization approach. The two infrastructures are represented along with their decision module, Power Decision Module (PDM) and IT Decision Module (ITDM).

sub-problems, which can be implemented using existing methods with minimal modifications. We propose a generic optimization algorithm, which is in charge of driving the two parts responsible of these sub-problem optimizations in order to find solutions that satisfy both sides.

An overview of the whole datacenter infrastructure is shown in Figure 2. Each *decision module* represents both the optimization algorithm, internal model, and data for the electrical part or the computing part of the infrastructure. The decision modules are black-boxes for the overall optimization mechanism (negotiation module in the figure). These modules are assumed to be designed and developed by experts in their own field, either electrical engineering or datacenter management. Their purpose is to evaluate possible power consumption or production plans represented by *power profiles*, which are sequences of power values at each time step in a given time window. Such evaluation is done by an internal optimization, leading to associate to a power profile a value representing its objective with arbitrary unit and scale. We denote as *utility* such a value, used as a preference indicator, assuming only that higher values are representing better outcomes.

This black-box is interesting for multiple reasons. In addition to help to reuse existing works and to make easier for experts to focus on the problem of their own field, this avoids ending with a huge, centralized optimization problem with many joint constraints. From a more technical and practical perspective, being independent of each other is expected to facilitate the design process and make it possible to modify one part of the infrastructure without requiring rethinking of the other nor involving deep changes of the overall optimization method.

In a power profile, the set of possible values at each time step is noted $\mathcal{W} \subseteq \mathbb{R}$. A power profile for a window of T time steps is therefore an element of \mathcal{W}^T . Utility value may be any number in \mathbb{R} with the only assumption that preferred solutions lead to greater values. We can abstract the evaluation of the power profiles by one of the decision modules as a function $f : \mathcal{W}^T \rightarrow \mathbb{R}$. This utility evaluation function is noted f_{itdm} for the IT decision module and f_{pdm} for the electrical one.

The generic optimization problem here consists therefore of finding a power profile p in \mathcal{W}^T such as the associated utilities $f_{itdm}(p)$ and $f_{pdm}(p)$ are both as good as possible. Each utility value represents however a distinct objective. Furthermore, these objectives are expected to be antagonists in most cases, as consuming more energy should favor the computing part at the expense of the electrical part.

The whole problem should therefore be considered as a multi-objective one. Reducing the problem to a single-objective, for instance by assigning an economical cost to the ecological footprint simplifies the formulation of the optimization goal but raises many other issues. Letting apart the difficulties of setting a reasonable cost for ecological impact, it makes the solutions very sensitive to the chosen value. This also ignores the fact that minimizing the ecological footprint may be a separate goal for some institutions, or even for big companies that desire to leverage such an effort in their brand image.

Without additional information on the relative importance of each objective, solutions may only be compared by the way of Pareto dominance [14]. Two solutions are therefore equally good if no one dominates the other by having better values for all the objectives. The set of non-dominated solutions, forming the *Pareto front*, represents all the optimal trade-offs between the different objectives. By knowing the Pareto front of a multi-objective problem, a decision maker has as much information as possible to select one of these optimal trade-offs. We will not discuss here the way this final decision is taken, which may be either a human decision maker or an automated process. Instead, the approach aims at providing the necessary information to take this decision. For this purpose, having the complete and exact Pareto front is ideal, but not always feasible. When not feasible easily, as in the problem presented here, finding a good approximation in a reasonable time is already valuable.

The optimization problems solved by each decision module are complex and we expect good quality heuristics with realistic models to be both resource- and time-consuming. Consequently, the cost of evaluating the utility associated to a power profile is far from being negligible. The time

available to take a decision is however limited, as it concerns an upcoming time window and becomes useless if taken too late. To take this into account, the number of possible evaluations of power profiles is constrained by a budget, denoted $eval_{max}$. Without more prior knowledge of the time required for PDM and ITDM, the two are considered in the same budget, such as the evaluation of a solution by the two decision modules counts as 2 costly evaluations. The complete problem may therefore be formulated as: find the best Pareto front approximation set \mathcal{A} by using at most $eval_{max}$ power profile evaluations (represented by either f_{itdm} or f_{pdm}). This budget constraint is one of the characteristics of approaches dealing with optimization of expensive black-boxes, such as Regis [36].

To give an idea of the time required for a given budget, simple hypotheses are proposed. We suppose first that the total time is largely dominated by the costly function evaluation and that a given profile is evaluated by both the ITDM and the PDM simultaneously, as if each decision module optimization process is running on a different machine. With a given budget $eval_{max}$ and an average execution time of the slowest decision module of d , the total duration is $eval_{max} * d/2$. Therefore, with a budget and $eval_{max} = 100$ evaluations and a costly evaluation duration of $d = 1$ minute, the total duration is close to 1 hour (50 minutes). For this reason it is critical to keep the budget in a smaller range compared to some other multi-objective evolutionary algorithm (MOEA) applications where days of computation may be acceptable. The target budget here is between 100 and 400 evaluations, which may both be interesting values depending on the duration of a single costly evaluation and the time horizon of the optimization.

III. RELATED WORKS

While it initially appeared as a specific case of energy-aware datacenter management, the integration of RESs and ESDs is now explored by a wild variety of works [25]. A number of them focus on off-site sources, by the way of concepts such as demand-response in the smart grid [33]. When on-site RESs sources are studied, two main categories are usually distinguished. On the one hand, those leveraging the geographical distribution of multiple datacenters in order to dynamically manage the load across each place depending on the local production, such as [24] and [26]. On the other hand, works considering only a single datacenter with on-site RESs, often use the temporal flexibility of the workload or degradation of the quality of service to leverage variable power production [19], [28], [32]. On this aspect, our work belongs to the second category, with a single datacenter powered by on-site RESs and ESDs.

All these approaches have in common to use a global formulation of the scheduling and energy supply problem. Additionally, only a handful of works in the whole area, such as [27], propose a purely multi-objective solution with Pareto front approximation, but with rather different models and goals. As the black-box design presented in the previous

section has not been explored by existing works in renewable-aware datacenter management, the rest of the section focuses on available methods for solving such a problem, and specifically in a multi-objective context.

Finding an approximation of the Pareto front for a multi-objective problem is well studied in the field of MOEAs. The core of evolutionary algorithms is the exploration of the solution space by many creation and modification of individuals (representing possible solution) and the evaluation of their objective values. Thousands of possible solutions are usually evaluated by these approaches. Hence, when the evaluation of a single solution is computationally expensive or the time available is limited, this amount of evaluation is prohibitive.

To reduce the time spent evaluating solutions by evolutionary algorithms, it is common to use an *approximation function* (or surrogate function) to substitute some of the evaluations by cheap approximations [7], [22].

A. INTEGRATION OF MOEA WITH SURROGATE MODELS

The integration of surrogate models with MOEAs may be done in different ways, with three main categories in the existing works. A first design, referred to here as *synchronous* integration, consists of applying the chosen MOEA using the real objective functions for a small number of iterations, followed by its application with the surrogate objective functions for more iterations [5], [30]. This process is repeated multiple times, using the last solution set as an initial population for the next phase and optionally updating the surrogate model after each non-surrogate phase. Such approaches do not require modifications of the MOEA used internally and can use surrogate method with relatively costly construction cost, but are easily impacted by a surrogate performing badly during a single phase and not well adapted for a low budget of costly objective evaluations (a few hundred).

The second category, *asynchronous* surrogate integration, mixes evaluated and surrogate-approximated individuals in the same MOEA population, performing a costly evaluation only if an approximated individual is expected to improve either the solution set or the surrogate model [23]. A few modifications in existing MOEA are usually required, or completely novel approaches are proposed, but such methods may be applied even with a low evaluation budget and the continuous verification of potentially good solutions reduces the impact of badly approximated individuals.

In addition to the synchronous or asynchronous integration, a few works also tried to integrate the surrogate model more deeply into the population selection and evolution process [2]. In this case, the main goal of the surrogate model is to allow the discovery of interesting areas of the solution space, either expected to improve the overall solution quality or to explore unknown parts of the space. Contrary to the two other categories, the surrogate model is not used simply as a substitute for evaluating possible solutions, but to perform local or global search in the solution space. Some solutions located in the area of interest are then evaluated using the

costly objective function, which in turn improves the knowledge of the surrogate model.

We propose in this article a new method, which belongs to the *asynchronous* category and is adapted to low evaluation budgets. Contrary to [23], the individuals are evaluated based on the number of generations during which they survive as good solutions. In addition, it embeds a mechanism to perform a costly evaluation when surrogate model considers to lack local information to accurately approximate a solution.

B. SURROGATE METHODS FOR BLACK-BOXES

A surrogate of a costly function $f : \mathcal{S} \rightarrow \mathbb{R}$ is a function $f' : \mathcal{S} \rightarrow \mathbb{R}$ which gives, for each input $s \in \mathcal{S}$ a value as close as possible of $f(s)$ with lower cost (in terms of time or memory for instance).

When f is a black-box, regression techniques may be used to build a surrogate by sampling some of the solution space. We focus here on what is called *online* surrogate building, *i.e.* methods suited for learning a function progressively. Each decision module is a costly black-box function for which a surrogate is desired. It is expected that their implementation depends on both environmental factors (weather forecast or planned workload submission) and internal state (current state of charge of the batteries or running tasks). For these reasons, building a surrogate model of these functions offline, in a static way, is not adapted. Online methods may be used instead, which are progressively learning the shape of the utility functions as new solutions are evaluated. In addition, because the solution space consists in time series of T time steps, with up to hundreds of time steps considered, surrogate methods should handle many-dimensions inputs.

Common surrogate methods used in evolutionary algorithms include Kriging models, artificial neural networks and Radial Basis Functions (RBF) approximation [7], [22], [40], [41]. Some of them, such as Kriging, are known to be prohibitively costly for many-dimensions problem spaces [41]. Two methods from the literature, which are known to scale better with the number of dimensions, are considered in this article: multilayer perceptrons (a form of artificial neural network) and RBF. We do not aim to enter too much into the details of these methods here, but instead highlight their most important aspects.

As other artificial neural networks, Multilayer Perceptrons (MLPs) requires a training phase, with different solvers used in the literature. Artificial neural networks in general are notoriously known for the many model parameters, called hyperparameters, to carefully select in order to obtain good results [39]. In addition to the choice of a solver and its internal parameters (learning rate, number of iterations and many others depending on the solver), the shape of the hidden layers (number of layers and neurons in each) and the *activation function* deeply affects the performances of MLPs.

Regression methods based on radial-basis functions (RBF) are similar to a single-layer artificial neural network using a specific family of activation function. RBF of reasonable sizes are easier and cheaper to train than the general case of

MLP and are controlled by fewer hyperparameters. We will focus only on the cubic activation function, as it leads to good results in many surrogate-based optimization works, including high dimensional cases (≥ 100 solution space dimensions) and does not requires to tune an additional hyperparameter [36].

We believe that these methods, while performing well in many problems, are not well adapted for simultaneously having large solution space and few known samples caused by the low evaluation budget we are targeting. Two new methods, named Average Distance (AD) and Multiresolution Haar Transform (MHT) are proposed in the next section, which take the time-series aspect into consideration.

IV. APPROACH DESCRIPTION

The approach presented here consists of two parts: a MOEA adapted for using surrogate models with low evaluation budget and the surrogate model itself. First, we propose a novel method, adapted to low evaluation budgets and based on Strength Pareto Evolutionary Algorithm 2 (SPEA2) [48], a classical multi-objective genetic algorithm. Then, two surrogate models, adapted to online usage and the high-dimensional nature of the time series are presented

A. ADAPTING SPEA2 FOR SURROGATE OBJECTIVES FUNCTIONS

SPEA2 [48] is a genetic algorithm, designed to approximate Pareto front of multi-objective black-box problems. Its individual selection function is designed to favor both the improvement of the solutions towards the Pareto front and the diversity of the solutions to cover as much of the front as possible.

Similarly to Karakasis and Giannakoglou [23], a few additions are proposed to allow the MOEA to handle a mix of individuals with either evaluated (by the way of the costly objective functions) or approximated objective vectors. Each time a costly evaluation of an individual is performed, it is added to a set of known data points \mathcal{P} used by the surrogate models. Contrary to existing approaches however, the surrogate methods are not considered to be reliable, in the sense that they may fail to give an estimation of the objective value. The causes of a failure, called thereafter a *miss*, are fully controlled by the surrogate method but may consist for instance of a lack of local data near the estimated solution.

To take into account the possibility of an approximation miss, surrogate function f' objective space is augmented by a special value $f' : \mathcal{S} \rightarrow \mathbb{R} \cup \{Miss\}$. Handling surrogate model misses is expected to be beneficial to the overall optimization, by avoiding taking into account approximations that are known to be of low quality. In such situation, it may be better to evaluate the solution immediately, which will additionally feed the knowledge base of surrogate model.

Two sets of modifications are proposed on the original SPEA2 algorithm. The first, presented as the *naive integration*, only consists in minimal changes to allow approximation of objectives of individuals and their evaluation with




FIGURE 3. Flowchart of the proposed algorithm, replacing the usual individual objectives evaluation in SPEA2 to handle approximation method.

costly objective functions when some criterion is met. Then a set of improvements are detailed to handle mix of evaluated and approximated individuals, forming a new algorithm named USPEA2.

1) NAIVE INTEGRATION

A simplified version of the SPEA2 algorithm with a few modifications is given in algorithm 2. The fitness assignment function and selection operator (*setSpea2Fitness()* and *spea2Selection()* in the listing) are not detailed here, but are the same as described in the original SPEA2 article [48]. The initialization of the population, denoted by the *initialPopulation()* function in the algorithms, is detailed later in section V-D3.e. Modifications to the original algorithm are highlighted in blue. They provide an integration of the utility approximation functions in SPEA2. The flowchart in Figure 3 gives an overview of the approximation or evaluation process replacing the usual individual objectives evaluation, corresponding roughly to the detailed algorithm 1. The utility values of each individual are approximated if possible, otherwise directly evaluated by interrogating each decision module.

When an individual is approximated using the surrogate models, a *lifetime* is associated to it with an initial value denoted lf_{base} . At each generation, the lifetime of all individuals is decremented. When it reaches 0, a complete evaluation of the individual objectives is performed. If any of the surrogate models fails to approximate the utility value (miss), the individual is immediately evaluated. Each time an individual is evaluated with the costly objective functions, it is also added along with its real utility to the knowledge base \mathcal{P} of each surrogate. The number of generations is not the stopping criterion here, contrary to many genetic algorithms. Instead, the number of costly evaluations (sum of both decision modules) performed from the beginning is used. To ensure the time spent by the algorithm is negligible compared to the evaluation functions, a mechanism forces at least one individual to be evaluated at each generation, as shown in algorithm 2. If no evaluation has been performed by the *evaluateRequired()* function, the best approximated individual of the population according to the SPEA2 fitness value is directly evaluated.

The offspring is generated first by selecting pairs of individuals using binary tournament with replacement. For each pair, individuals are either mated (with a probability pb_{cross}) or both mutated (probability $1 - pb_{cross}$). The mating function

Algorithm 1 Addition to SPEA2 for Mixed Approximation and Evaluations

Function *evaluateRequired*(*population*, *archive*)

```

foreach ind ∈ population ∪ archive do
  ind.lifetime ← ind.lifetime - 1 ;
  if not ind.valid then
    approximate(ind) ;
  end
  if ind.lifetime = 0 then
     $u_{pdm} \leftarrow f_{pdm}(\text{ind.value})$  ;
     $u_{idm} \leftarrow f_{idm}(\text{ind.value})$  ;
    ind.utility = ( $u_{pdm}$ ,  $u_{idm}$ ) ;
    ind.lifetime ← ∞ ;
  end
  ind.valid ← true ;
end
  
```

Function *approximate*(*ind*)

```

 $u_{pdm} \leftarrow f'_{pdm}(\text{ind.value})$  ;
 $u_{idm} \leftarrow f'_{idm}(\text{ind.value})$  ;
ind.utility ← ( $u_{pdm}$ ,  $u_{idm}$ ) ;
if  $u_{pdm} = \text{Miss} \vee u_{idm} = \text{Miss}$  then
  ind.lifetime ← 0 ;
else
  ind.lifetime ←  $lf_{base}$  ;
end
  
```

used here is a one-point crossover. A mutation consists of an independent probability pb_{indmut} for each time step of the individual to be replaced by a random value from a uniform distribution in the range of acceptable values \mathcal{W} . Finally, each new individual is marked as having undefined utility values ($\text{ind.valid} \leftarrow \text{false}$).

2) USPEA2

SPEA2 was designed to use reliable utility functions. In our approach, the utilities of an individual are either evaluated directly by the decision modules or by a surrogate function. While the utility is reliable in the first case, in the second case the approximated utility is not, as it may change if the individual is evaluated later. Such approximated utility is therefore considered *unreliable*, as well as the individual to which the approximation is associated.

Unreliable utilities may lead SPEA2 to perform badly. An individual with approximated utility values may dominate an evaluated, reliable individual, replacing it in the archive. If the estimated individual is finally evaluated with lower utilities than the eliminated one, this results in a decrease of the overall individual pool quality and potentially in a deterioration of the Pareto front approximation.

To overcome this problem, we propose USPEA2, for Unreliability-SPEA2. Figure 4 illustrates the algorithm, which is detailed in algorithm 3, with modifications compared to SPEA2 highlighted. It consists in the introduction of a second archive of individuals, called *reliable archive*,

Algorithm 2 Simplified Algorithm of SPEA2. Lines in Blue Are Additions to Provide a Naive Support of the Utility Approximation Functions

```

Function optimizeSpea2()
  archive  $\leftarrow \emptyset$  ;
  population  $\leftarrow$  initialPopulation(populationSize) ;
  ended  $\leftarrow$  false ;
  while not ended do
    ended  $\leftarrow$  endingCondition() ;
    if ended then
      population  $\leftarrow$  {ind | ind  $\in$  population, ind.lifetime =  $\infty$ } ;
      archive  $\leftarrow$  {ind | ind  $\in$  archive, ind.lifetime =  $\infty$ } ;
    else
      evaluateRequired(population, archive) ;
      if no evaluation performed then
        evaluateBestIndividual() ;
      end
    end
    setSpea2Fitness(population, archive) ;
    archive  $\leftarrow$  spea2Selection(population, archive, archiveSize) ;
    if not ended then
      population  $\leftarrow$  generateOffspring(archive) ;
    end
  end
  return archive ;

```

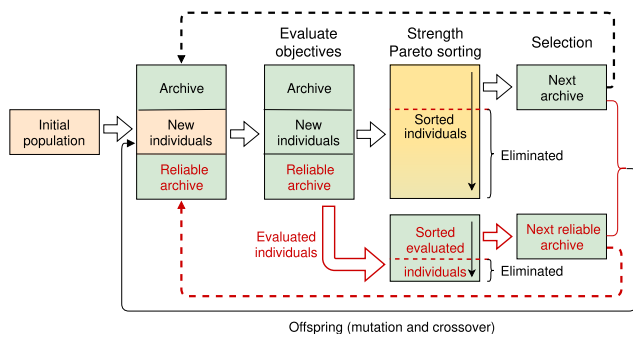


FIGURE 4. Schematic view of USPEA2 algorithm with modifications compared to SPEA2 highlighted in red.

containing only evaluated solutions. Individuals from this second archive allow maintaining a set of good and reliable solutions even if some of them are temporarily dominated by unreliable individuals. At each generation, after the fitness assignment and selection of the next archive in accordance with the classic SPEA2 scheme, an additional selection is done only on evaluated solutions, using the same selection operator. To avoid the introduction of a large number of duplicated individuals, which is not addressed efficiently by SPEA2, uniqueness of individuals is ensured when the two multi-set archives are merged.

B. SURROGATE MODELS FOR TIME SERIES

Evaluating the quality of a potential solution $p \in \mathcal{W}^T$ is achieved by the way of the black-box objective functions

$f_{idm}(p)$ and $f_{pdm}(p)$. To approximate them, a surrogate model has the data base of previously evaluated solutions \mathcal{P} , with their power values normalized in $[0, 1]$ for practical reasons.

A power profile may be viewed as an arbitrary vector of T dimensions. Hence, classic surrogate methods known to work with high-dimensional solution space, such as RBF and MLP can be used. However, they are not designed to leverage the *time series* aspect of the power profiles.

To the best of our knowledge, no existing work already focused on optimizing costly functions defined on time series space. Despite the many areas where time series are studied, such as in economics and meteorology, their use cases are usually quite different, such as finding similarities between several time series [1] or predicting future values of time series [4]. Therefore, we propose new approaches to perform utility approximation adapted for long profiles, inspired by works on signal and time series processing. These methods, contrary to existing surrogate methods, may fail to return an approximation for some inputs (*miss*). As argued previously, we believe that forcing a costly evaluation because the surrogate method considers lacking local data to be accurate is not necessarily bad. Quite the contrary, it helps improving the quality of the surrogate around this new solution and prevent to return very inaccurate values which could impact negatively the overall optimization.

1) AVERAGE-DISTANCE METHOD

By assuming that small differences between power profiles (high similarity) generally imply small differences in their

Algorithm 3 Simplified Algorithm of the Unreliability-SPEA2 (USPEA2) Approach. Changes From SPEA2 Are Highlighted in Blue

```

Function optimizeUspea2()
  archive  $\leftarrow \emptyset$ ;
  reliableArchive  $\leftarrow \emptyset$ ;
  population  $\leftarrow$  initialPopulation(populationSize);
  ended  $\leftarrow$  false;
  while not ended do
    ended  $\leftarrow$  endingCondition();
    if not ended then
      evaluateRequired(population, archive);
      if no evaluation performed then
        evaluateBestIndividual();
      end
    end
    uniqueArchive  $\leftarrow$  unique(archive  $\cup$  reliableArchive);
    setSpea2Fitness(population, uniqueArchive);
    archive  $\leftarrow$  spea2Selection(population, uniqueArchive, archiveSize);
    reliablePopulation  $\leftarrow$  {ind | ind  $\in$  population, ind.lifetime =  $\infty$ };
    uniqueReliableArchive  $\leftarrow$  {ind | ind  $\in$  uniqueArchive, ind.lifetime =  $\infty$ };
    setSpea2Fitness(reliablePopulation, uniqueReliableArchive);
    reliableArchive  $\leftarrow$  spea2Selection(reliablePopulation, uniqueReliableArchive, archiveSize);
    if not ended then
      population  $\leftarrow$  generateOffspring(archive);
    end
  end
  return reliableArchive;

```

utility values, a distance function may be used to achieve utility approximation.

The proposed method, named simply Average Distance (or AD) is similar in some aspects to a weighted k -nearest neighbors method [17]. It consists, first, in selecting profiles among the known and already evaluated ones which are closer than a given distance r_{close} from the unknown profile to estimate. The approximation is then performed by doing a weighted average of utility values of the close solutions, with a weight depending on their relative closeness. Based on an initial study, it appeared that euclidean distance was not always appropriate to compare profiles. In order to penalize some important local differences (e.g. a peak in one of the power profiles), a different measurement is used, based on the average of the square differences at each time step. As this is equivalent to mean square error (MSE) defined in statistics, it is used here and qualified by abuse of language as the distance between two profiles.

The complete algorithm is detailed in algorithm 4. The approximation of an individual p is considered as a miss when less than n_{close} are known in a radius of r_{close} around p .

2) MULTIREOLUTION HAAR TRANSFORM DISTANCE METHOD

While there is no existing work usable directly for building surrogate models of function defined on time series,

Algorithm 4 Average-Distance (AD) Approximation Algorithm, for a Given *target* Profile and Using the Known Solutions \mathcal{P} With Their Utility Values Given by *objective*

```

Function distanceUtilityApproximation(target, rclose, nclose, P, lfbase)
   $\mathcal{P}_{close} \leftarrow \{p \in \mathcal{P}, \text{mse}(p, \text{target}) < r_{close}\}$ ;
  if  $|\mathcal{P}| \geq n_{close}$  then
    approx
       $\leftarrow$  distanceWeighted(target,  $r_{close}$ ,  $\mathcal{P}_{close}$ , mse)
    ;
    return approx;
  else
    return Miss;
  end
Function distanceWeighted(target, radius, Pclose, distance)
  totalWeight  $\leftarrow \sum_{p \in \mathcal{P}_{close}} \text{radius} - \text{distance}(p, \text{target})$ ;
  return
     $\sum_{p \in \mathcal{P}_{close}} \text{objective}(p) \frac{\text{radius} - \text{distance}(p, \text{target})}{\text{totalWeight}}$ ;

```

similarity detection of time series shares some challenges with our problem. Several works already demonstrated




FIGURE 5. Illustration of the transformation of a time series $X = (8, 6, 1, 5)$ using Haar wavelets, resulting in flatten representation $HC_2 = (5, 2, 1, -2)$.

some qualities of wavelet transform for similarity detection and similarity indexing of time series and more specifically of Haar wavelets [1], [10], [11]. Wavelet transforms provide an elegant way to achieve dimensionality reduction in time series, by resulting in a multiresolution representation.

We first need to introduce some notations and properties of the Haar wavelet transform, as well as a definition of a partial distance between Haar representations. Then, the surrogate method itself is described.

a: OVERVIEW OF THE HAAR TRANSFORM

Transforming a series of $N = 2^k$ values using Haar wavelet gives a new series of the same size, containing the coefficients of the wavelet for different combinations of scale and shift. These coefficients are commonly grouped by *sequence* [42], each containing the coefficients for a given scale (or frequency). Figure 5 shows the application of Haar wavelet transform, with a few simplifications on the actual resulting values for illustration purpose.

A total of $k + 1$ groups are obtained this way, denoted by H_0, \dots, H_k . The first, H_0 , is a single value proportional to the global average of the series. The other groups $H_i, i \in \{1, k, \dots\}$ have a size given by $|H_i| = 2^{i-1}$. Each sequence H_i contains information on the original series with a temporal resolution doubled compared to H_{i-1} .

Let X, Y be two series of size $N = 2^k$. The concatenation of the sequence H_0, \dots, H_i of the Haar transform of X , called here its *flatten notation*, is denoted by $HC_i(X)$.

Haar transform is fast to compute, having a time complexity in $O(n)$ [11], with n the length of the signal. Also, the first components of the Haar representation depend only on the *low frequency* features, while each successive depth increase

the high frequency details, allowing efficient dimensionality reduction.

b: HAAR-BASED DISTANCE FUNCTION

The Haar transform of signals may be used to define a new partial distance function. The Euclidean distance between two signals is the same as the Euclidean distance between the flattened representation of the Haar transform of those signals [11]. With $dist$ being the Euclidean distance function, and $X, Y \in \mathbb{R}^l$ two signals of length $l = 2^k$, we have:

$$dist(X, Y) = dist(HC_k(X), HC_k(Y)) \quad (1)$$

We introduce a *Haar distance* function, denoted $hdist_d$, which gives the distance between the first d sequences of two Haar transform:

$$hdist_d(X, Y) = dist(HC_d(X), HC_d(Y))$$

The Euclidean distance is monotonically increasing when new elements are added to each vector, such as:

$$dist(u_0 \dots u_{l-1}, v_0 \dots v_{l-1}) \leq dist(u_0 \dots u_l, v_0 \dots v_l)$$

Two properties must be highlighted here. For two signals $X, Y \in \mathbb{R}^l$, with $l = 2^k$:

- 1) with $d = k$, the maximum depth, the Euclidean distance and the Haar distance are equal (Equation (1))
- 2) Haar distance is monotonically increasing with the value of d (Equation (2))

$$\forall i, j \in \{0, \dots, 2^k\}, \quad i \leq j, \quad hdist_i(X, Y) \leq hdist_j(p, q) \leq dist(p, q) \quad (2)$$

These properties are exploited in a new surrogate model, leveraging both dimensional reduction by using Haar distance and local precision around known solutions by using the monotonic convergence towards the real Euclidean distance.

c: HAAR-BASED SURROGATE

We propose here a new surrogate method, named Multiresolution Haar Transform (MHT). It uses the same basic idea of Average Distance method, but exploit the Haar transform of the solutions at multiple temporal *resolutions*. The full algorithm is detailed in algorithm 5. Basically, it looks for the profiles closer than a fixed radius r_{close} to the one to approximate according to the Haar distance at a given depth. This depth, starting from 0, is progressively increased, hence reducing the number of profiles close enough to be considered. It stops when the increase of depth results in a close solutions set \mathcal{P}_{close} smaller than a given threshold n_{close} . The approximation is then calculated based on the average of the objective values of the close solution set, weighted by their closeness. As the Haar wavelets are only defined on series of length being a power of 2, the power profiles are filled with 0 up to the nearest valid length.

To be easier to understand, the full algorithm given in algorithm 5 is simplified and not optimized. Some optimizations may be used to achieve low average complexity.

Algorithm 5 Haar-Based Utility Approximation Algorithm, Approximating the Objective Value for *target* Profile and Using the Known Solutions \mathcal{P} . The Function *distanceWeighted()* Is Defined in algorithm 4

Function *haarUtilityApproximation*(*target*, r_{close} , n_{close} , \mathcal{P} , lf_{base})

```

depth ← 0 ;
 $\mathcal{P}_{close} \leftarrow \mathcal{P}$  ;
 $\mathcal{P}_{accepted} \leftarrow \emptyset$  ;
while  $depth < k \wedge |\mathcal{P}_{close}| \geq n_{close}$  do
   $\mathcal{P}_{close} \leftarrow \{p | p \in \mathcal{P}_{close}, hdist_{depth}(p, target) < r_{close}\}$  ;
  if  $|\mathcal{P}_{close}| \geq n_{close}$  then
     $\mathcal{P}_{accepted} \leftarrow \mathcal{P}_{close}$  ;
    depth ← depth + 1 ;
  end
end
if  $\mathcal{P}_{accepted} \neq \emptyset$  then
  approx ← distanceWeighted(
    target,  $r_{close}$ ,  $\mathcal{P}_{accepted}$ ,  $hdist_{depth}$ ) ;
  return approx ;
else
  return Miss ;
end

```

Mainly, profiles (known solutions and a target profile) may be stored along with their Haar transform result, avoiding recalculation at each distance measurement. This makes Haar distance computing complexity $O(n)$, with n the profiles length.

Therefore, the complexity of Haar surrogate algorithm is dominated by its main loop having a maximum of $\log_2(n)$ iterations. Each iteration requires to evaluate the Haar distance with a subset of the known profiles set, for a worst-case time in $O(n \cdot m)$, m being the number of known solutions. The overall, worst case time complexity is thus $O(n \log(n) \cdot m)$.

However, we can expect better performance on average. The set of the considered solution at iteration i of the main loop is always a subset of the close solutions found at iteration $i - 1$. In addition, the depth of the Haar distance is increasing with each generation, starting from constant-time complexity (distance between two 1-dimension vectors) and doubling at each iteration. Thus, the most costly distance computing occurs with the smaller set of retained solutions.

V. EVALUATION METHODOLOGY

To evaluate our approach, a practical optimization problem must be targeted. While some multi-objective functions are commonly used in the MOEA field for evaluation purpose, such as ZDT [47] or DTLZ [15], they do not capture some characteristics of the problem exposed here. Particularly, they are not defined on time series, which is our primary target.

Our approach is designed to use black-box IT and electrical decision modules. Hence its evaluation requires to specify

the models and algorithms used for the concrete modules. Knowing the real Pareto front is ideal to evaluate a purely multi-objective optimization method, as it is a solid and unquestionable point of comparison.

We target a medium scale datacenter powered partially by on-site RES and batteries. A scheduling algorithm from the literature, such as RECO [9], is usable as the IT decision module. Finding the real Pareto front of the overall problem with such detailed model, however, is hardly feasible.

Hence a simplified model of datacenter and electrical infrastructure is proposed first. It is designed to allow a Mixed-Integer Linear Programming (MILP) formulation of the overall optimization problem with reasonable solving time, providing a very accurate Pareto-optimal set of solutions, while keeping most characteristics as expected in a similar real-world problem. This model is therefore used to perform extensive evaluations of the proposed heuristics. To validate the main findings, additional experiments are performed by using the RECO scheduling algorithm in place of the simplified IT model.

A. COMPLEX IT MODEL

The RECO algorithm [9] consists in a batch task scheduling algorithm under power envelope constraint. The computing infrastructure handles heterogeneous machines, with Dynamic Voltage and Frequency Scaling (DVFS) and power-on/power-off models taking into account the time and power required for booting up or shutting down the machines. It focuses on a workload of batch tasks with due dates, each having different requirements in terms of CPU. Its model nor the RECO algorithm itself are modified here. To represent the utility of a solution, the percentage of due date respected is used.

B. SIMPLIFIED MODEL FOR EVALUATION

The practical optimization problem proposed for the evaluation is composed of two sub-models, tied together only by the way of provided or consumed power denoted p_t for each time $t \in \{0, \dots, T - 1\}$. An instance of the model concerns a single time window of T time steps, of fixed duration δ each. A summary of the symbols used in the model, along with their values as detailed later, is given in table 1.

1) ELECTRICAL MODEL

We consider a set of RESs, producing together, at each time step $t \in \{0, \dots, T\}$, an average power of r_t . In addition, batteries may be used to store some energy. These batteries have a total capacity bat_{max} energy units, with an initial charge of bat_{init} at the beginning of the window, as given by Equation (3a). The remaining charge at the beginning of a time step t is denoted by c_t and is constrained by Equation (3b). The whole electrical system is connected to the electrical grid. The power bought from or sold to the grid to it at a time t is given by g_t . The sign of g_t indicates if the power is bought ($g_t > 0$) or sold ($g_t < 0$). Finally, the maximum power made available by the electrical part is the sum of the

renewable production, battery charge or discharge and grid usage, as given by Equation (3c).

$$\forall t \in [0, T), \quad \begin{cases} c_0 = bat_{init} & (3a) \\ 0 \leq c_t \leq bat_{max} & (3b) \\ p_t \leq r_t + (c_{t+1} - c_t)\delta + g_t & (3c) \end{cases}$$

The utility of the electrical part represents the environmental impact of a decision, calculated through several criteria. A first part, in Equation (4), depends on the grid usage, with each unit of energy bought decreasing the utility by ghg_{buy} and each unit sold increasing it by ghg_{sell} .

$$u_{grid} = \sum_{t=0}^{T-1} \begin{cases} -ghg_{buy} \cdot \delta \cdot g_t & \text{if } g_t \geq 0 \\ -ghg_{sell} \cdot \delta \cdot g_t & \text{otherwise} \end{cases} \quad (4)$$

As battery life highly depends on the amount and the depth of charge and discharge cycle [12], [18], [43], a second part of the utility reflects this additional aging cost, as in Equatio (5).

$$u_{aging} = cost_{aging} \sum_{t=0}^{T-1} |c_{t+1} - c_t| \quad (5)$$

Finally, a third part represents a *potential saving* coming from the energy stored in the battery at the end of the time window presented in Equation (6). The potential gain is calculated as if it saved the same amount of energy from being bought later on the grid, while taking into consideration the aging of batteries caused when the energy will be used.

$$u_{storage} = c_T(ghg_{buy} - cost_{aging}) \quad (6)$$

The total utility, for the electrical part, is therefore given by Equation (7).

$$u_{elec} = u_{storage} - u_{aging} - u_{grid} \quad (7)$$

2) IT MODEL

The IT part is modeled as a set of M identical machines, each consuming $power_{machine}$ when being used. The workload is considered as a fluid mass W , expressed in equivalent run time for a single machine (machine-time unit). This total mass of the workload W can be divided into an arbitrary number of chunks to be run on any machine at any time (Equation (8a)). It can be seen as a *divisible load* with no communication cost and a linear cost on the load. However, each part of the workload requires an entire machine during a full time step.

The number of machines used at a given time t is equal to the amount of tasks executed, noted w_t and constrained by Equation (8b). The minimum power required at every time step is therefore given by Equation (8c).

$$\forall t \in [0, T), \quad \begin{cases} \sum_{t=0}^{T-1} \delta \cdot w_t \leq W & (8a) \\ 0 \leq w_t \leq M & (8b) \\ p_t \geq power_{machine} \cdot w_t & (8c) \end{cases}$$

The utility of the IT part depends on the executed workload. We consider that the benefit for executing any workload is homogeneous through $gain_{task}$:

$$u_{base} = gain_{task} \sum_{t=0}^{T-1} \delta \cdot w_t \quad (9)$$

A second part of the utility is given by an additional gain when a task is submitted early. This gain, for each task, is at most $gain_{early}$ if the task is executed in the first slot, as given in Equation (10).

$$u_{early} = gain_{early} \sum_{t=0}^{T-1} \frac{(T-1-t)w_t}{(T-1)\delta} \quad (10)$$

Maximizing u_{early} is then equivalent to minimizing the flow time of the workload.

Finally, Equation (11) gives the overall IT utility.

$$u_{it} = u_{early} + u_{base} \quad (11)$$

C. DECISION MODULES AND CENTRALIZED PARETO APPROXIMATION

The IT and electrical models can be used to derive a centralized MILP formulation. However the optimization problem has two different objectives: maximization of both u_{itdm} and u_{pdm} . To obtain an approximation of the Pareto front of solutions, two steps are required. In a first phase, the extreme solutions are obtained by optimizing each objective without any constraint on the other. It results in two pairs of utility values, $(u_{pdm}^{min}, u_{itdm}^{max})$ when ITDM utility is optimized and $(u_{pdm}^{max}, u_{itdm}^{min})$ when it is PDM utility. Then, a desired number of solutions distributed along the Pareto front are obtained by maximizing one of the utility subject to an additional constraint on the second:

$$\begin{aligned} & \text{maximize } u_{itdm} \\ & \text{subject to } u_{pdm} \geq u_{pdm}^{min} + \alpha (u_{pdm}^{max} - u_{pdm}^{min}) \end{aligned}$$

By repeating this for different values of $\alpha \in [0, 1]$, a good approximation of the real Pareto front is provided with as many solutions as required.

Each infrastructure model can also be used to implement each decision module. By accounting only for the constraints of the electrical infrastructure (respectively IT infrastructure), the utility of a power profile P is given by maximizing u_{pdm} (respectively u_{itdm}) with additional constraints:

$$\forall t \in [0, T), \quad p_t = P_t$$

D. PARAMETERS FOR EVALUATION

The evaluation of our approach is done using several instances of the electrical and computing infrastructure models. Some parameters, such as the capacity of the battery bat_{max} or the duration of a time step δ , are given by an expression of higher level parameters. This helps to make simpler the definition of the different *scenarios*. The values of the parameters related to GHG emission and renewable power

TABLE 1. Summary of symbols used in the electrical and IT models.

Symbol	Meaning	Value in final evaluation
ghg_{buy}	Cost (in GHG emission) of an energy unit from the grid	0.3 kg CO ₂ -eq/kWh
ghg_{sell}	Resell price of an energy unit to the grid provider	0.15 kg CO ₂ -eq/kWh [†]
$cost_{aging}$	Cost due to battery aging when battery is charged or discharged.	0.0483 kg CO ₂ -eq/kWh
$gain_{task}$	Benefit when a unit of workload is executed in the time window.	1
$gain_{early}$	Maximal additional benefit when a unit of workload is executed early	0.2
T	Number of time steps in a time window	{20, 80 or 320}
δ	Duration of a single time step	72/ T
r_t	Renewable power available at time t	-
c_t	Energy in the batteries at beginning of time t	-
bat_{max}	Total capacity of the batteries	*
bat_{init}	Energy in the batteries at $t = 0$	$bat_{max} \cdot bat_{soc}^*$
g_t	Power bought from ($g_t > 0$) or sold to ($g_t < 0$) the grid at time t	-
$power_{machine}$	Machine consumption when used	200 W
W	Total amount of workload to execute in the time window	*
M	Number of available machines	100
w_t	Amount of workload executed at time t	-
p_t	Power used at time t	-

[†] Except for the specific experiments showing the impact of this value in approach results.

* Calculation detailed in section V-D.

production use existing studies and publicly available data to keep realistic figures.

1) TIME WINDOW AND TIME STEP

For all the experiments, a fixed time window of 72 hours is used, starting at midnight. However, the number of time steps T is set to either 20, 80 or 320. Time step duration δ thus depends on T such as $\delta = 72/T$ hours.

2) WORKLOAD AND MACHINES

The amount of workload considered is based on the maximal work that may be executed by the datacenter if all the machines are used during the whole time window. A factor is then applied, to consider an average load of $\overline{Load} = 70\%$. Therefore, we can define its value using Equation (12) below.

$$W = \overline{Load} \cdot \delta T \cdot M \quad (12)$$

The experiments account for $M = 100$ machines. Each is consuming an average power of $power_{machine} = 200$ W under load, for a datacenter peak consumption of 20 kW. The ITDM utility benefit is fixed to values in arbitrary units, $gain_{task} = 1$ and an additional maximal benefit when executed early $gain_{early} = 0.2$.

3) BATTERY CAPACITY AND INITIAL STATE

The maximum capacity is set to provide enough energy to fully power all the machines during 12 hours. With the

parameters chosen for the IT part, this gives approximately 240 kWh of batteries. This is a relatively high figure, but results in an initial investment of around \$65,000 based on the 2016 values [13]. This seems acceptable, staying in the order of magnitude of the budgets and operating cost of such a medium-scale datacenter.

$$bat_{max} = power_{machine} \cdot M \cdot 12 \quad (13)$$

The initial capacity is set by using a coefficient bat_{soc} , such as $bat_{init} = bat_{max} \cdot bat_{soc}$.

a: PHOTOVOLTAIC PRODUCTION DATA

We used the high quality synthetic photovoltaic production data from the U.S. National Renewable Energy Laboratory (NREL), contained in the Solar Power Data for Integration Studies dataset [31]. It contains, for hundreds of power plant locations, 5-minutes power data during all the year 2006.

We used the data from a power plant located at coordinates N34°85' W120°45', near the city of Santa Maria in California. This is a medium power plant, with a peak capacity of 4 MW (with a maximum of 3.93 MW reached in these data). For the need of our experiments, we choose only a few 3-days periods of interest inside the year-long data, with associated production shown in Figure 6:

- **ExcessRenew**: almost perfectly sunny days, starting **2006-06-01**




FIGURE 6. Power production for the simulated PV power plant in California during the chosen periods.

- **Normal:** noisy days with the second one being very cloudy, starting **2006-06-10**
- **FewRenew:** very cloud days during the winter, starting **2006-01-01**

This power plant is scaled to fit the need of the considered datacenter. More precisely, we assume that the goal was to produce a yearly amount of energy equals to the yearly amount of energy consumed by the datacenter (considering the average load $\overline{Load} = 0.7$, as given above). The yearly average of power production is $\overline{PvProd} = 730.9$ kW. Therefore, the scaling factor used on the PV data is $pvScaling = 1.915 \times 10^{-2}$, as given by Equation (14) below.

$$pvScaling = \frac{M \cdot \overline{Load}}{\overline{PvProd}} \quad (14)$$

b: GAS EMISSION PARAMETERS

The electrical utility for a given profile is an amount of GHG emission avoided the given power during the time window of the profile. These values are given in equivalent kilogram of CO₂. In this work, we do not focus on the total environmental impact of building such a green datacenter, considering it already been built. Hence, even if their manufacturing is a pollution source [3], [34], using energy from photovoltaic panels is considered free of any additional environmental impact.

i) BATTERIES

Technically, using a battery does not cause additional emissions, as its manufacturing is already performed. The battery lifetime is nonetheless greatly affected by its usage and it may be considered that once a battery is considered to be too deteriorated it will be replaced by a new one. Hence it is possible to incur some of the initial battery cost (economical or ecological) to each performed operation, depending on how much it affects its lifetime, as a form of amortization.

We consider Lithium Manganese Oxide (LMO) batteries, from the lithium-ion family. According to Hao *et al.* [21], the emission of GHG caused by their manufacturing, denoted $battery_{GHG}$, is about 96.6 kg CO₂-eq/kWh. While it is not the only parameter affecting the lifetime of batteries, the number of full charge-discharge cycle is a good aging predictor [12], assuming non-extreme usages. The value used in the following is a pessimistic lifetime of $N^{cycle} = 1000$ full charge/discharge cycles [12], [18], [43].

To account for the impact of the charge and discharge depth on the aging [12], a linear model is used. This accounts lifetime degradation, therefore equivalent GHG emission, for each unit of energy either charged or discharged. Applying Equation (15) with the considered values leads to $cost_{aging} = 0.0483$ kg CO₂-eq/kWh.

$$cost_{aging} = \frac{battery_{GHG}}{2 \cdot N^{cycle}} \quad (15)$$

ii) ELECTRICAL GRID

The California Environmental Protection Agency provides an annual GHG emission inventory [8] for multiple sectors, including electrical power generation. Based on the 2016 report, the average annual values for the year 2014, including energy imports, is 0.3 kg CO₂-eq/kWh. This value is used for ghg_{buy} .

Selling an excess of production on the electrical grid is expected to avoid some production of GHG. For one unit of renewable energy sold in the grid, the avoided GHG is likely less than the emission caused by the production of one unit on the overall grid. In addition to the various losses such as the conversion and transport ones, the grid is not always able to absorb excessive energy as the traditional power plants may already satisfy the demand at a given time. To account of this, the avoided GHG emission for selling energy to the electrical grid ghg_{sell} is considered here, on average, to be half of ghg_{buy} .

c: COMPLEX ITDM SETTINGS

The experiments involving the RECO scheduler [9], [38] use the same parameters and workload as in the original paper [9]. The computing infrastructure consists in 30 bi-processor machines of two kinds, with a workload composed of 1029 tasks of various duration and CPU requirements submitted along a 48-hours time window.

To use it as IT objective function in our approach, the power profile evaluated is used as the maximum power envelope and the utility (objective value) is computed as the percentage of tasks scheduled without violation of their due dates. In the experiments involving the complex IT model, we consider the same electrical model as for other experiments. However, to take into account the lower total consumption for the associated workload, the renewable production and battery capacity are scaled to 20% of the value used for the simplified IT infrastructure.

d: SCENARIOS

The number of model parameter combinations is important, but the main goal of the experiments is to evaluate the performance of the approach in a few different conditions. We define therefore a set of 3 scenarios, consisting of variations of the photovoltaic panels' production data and battery settings. Their names and the values used are given in table 2. They are chosen to have both some expected normal situation and edge cases where the problem space may be easier or more difficult to explore.

TABLE 2. Definition of the scenarios used for the experiments.

Name	PV production	bat_{soc}
ExcessRenew	Good	0.5
Normal	Bad	0.5
FewRenew	Worst	0.25

e: META-HEURISTIC SETTINGS

The SPEA2 and USPEA2 algorithms, along with their mutation and crossover functions, are configured by several parameters. The values chosen for probability of crossover pb_{cross} and the independent probability of time step mutation pb_{indmut} are detailed in section VI-A, based on experimental results. Some others are not experimentally studied here.

The size of the archives and of the population, $archiveSize$ and $populationSize$, are both set to 20. This relatively small value is chosen because a low number of costly objective function evaluation budget is aimed in this approach. A higher number of individuals in the population would require more evaluations, which is not desirable here. Lower values tend to limit too much the diversity of the individuals, which is required to have useful crossover.

As described in section IV-A1, the mutation function used causes each time step of the concerned individual to be potentially replaced by a random value following a uniform distribution. By knowing the minimum and maximum power produced or consumed by each decision module (p_{min} and p_{max}), the range of the uniform distribution is given by:

$$[\max(p_{min}^{idm}, p_{min}^{pdm}); \min(p_{max}^{idm}, p_{max}^{pdm})]$$

In practice with the models used, the minimum is bounded to 0 (electrical infrastructure not providing power, computing infrastructure with all machines powered off). The maximum is unbounded for the PDM, but limited by the ITDM, when all machines are used simultaneously.

In the description of SPEA2 and USPEA2 algorithms, the individuals forming the population at the beginning of the first generation are obtained through the $initialPopulation(size)$ function. Two different initial population schemes are studied here.

The primary one is the *Random* scheme. It simply consists of a generation of $size$ individuals for which every time step values is randomly selected using a uniform distribution. The bounds of the distribution are the same as for the mutation function.

The second scheme is called *Best*. While $size - 2$ individuals are generated similarly to the *Random* scheme, 2 individuals are provided by the decision modules themselves. Each module is supposed able to compute an optimal power planning for its own objective (without any constraint related to the other's requirement). Therefore, the best power planning of each module is added to the initial population and is expected to be part of the real Pareto front as it maximizes one of the objectives.

To compare the results of different heuristics, two common multi-objective performance metrics [37] are used here:

hypervolume indicator and generational distance. The hypervolume indicator (abbreviated HV), introduced in [46], gives the space covered between a set of non-dominated solutions and a pessimistic (Nadir) point. It measures both the diversity of the solutions (their distribution along the two objectives) and the accuracy (*i.e.* the proximity with the real Pareto front) of the Pareto front approximation. The same reference point is used for all the comparable results (same scenario and same number of time steps) and consists of the minimum values for each objective across all these results. For ease of interpretation, the normalized hypervolume (nHV) is used, considering the Pareto front obtained by the MILP for the same scenario and time window as a reference.

Initially proposed by Van Veldhuizen and Lamont [44], the generational distance (GD) (also referred to as *average distance to Pareto front*), measures only the accuracy aspect. It is often used to help to interpret the hypervolume values. Normalization is performed on the objective values to scale each dimension of the objective space in $[0, 1]$ before computing the generational distance, for the same reasons as for hypervolume.

VI. PRELIMINARY EXPERIMENTS ON APPROACH PARAMETERS

Several parameters must be chosen for the optimization algorithms (SPEA2 and USPEA2) and the different surrogate methods.

- pb_{cross} , the probability of crossover/mutation
- pb_{indmut} , the probability for each time step value of a mutated individual to be modified
- lf_{base} , the initial lifetime of approximated objective values
- r_{close} , the distance considered as *close* for the Average Distance (AD) approximation and Multiresolution Haar Transform (MHT) approximation methods
- n_{close} , the minimum number of close known solutions to do an approximation (AD and MHT)
- number of neurons and structure of the hidden layers (MLP)
- activation function (MLP)

The behavior of the overall optimization process depends on the interaction between (U)SPEA2 and the surrogate models. For instance, the accuracy of the surrogate models affects the individual that will be preferred by the MOEA to survive, affecting which individuals will be eventually evaluated. In turn, the evaluated individuals affect the database of known solutions, impacting the quality of the approximation performed by the surrogate. Therefore the precise impact of each parameter is expected to be tied to all the others. It would be rather impractical, especially for more complex problems, to evaluate the overall results of such an approach for a large set of parameter combinations. Instead, we propose a simplified methodology to choose them.

The proposed methodology does not aim to find the optimal combination of their values for our precise evaluation scenarios (overtuning the parameters). To enforce that,

only one of the evaluated scenarios is used for setting the parameters. In a first time, the two parameters of (U)SPEA2 (crossover probability and independent mutation probability) are studied, without considering any approximation method. Then for each surrogate approach, the impact of their parameters on the quality of approximation is studied, using a prepared training set of power profiles and associated objective values (*i.e.*, independently from a MOEA). Finally, the initial lifetime of approximated individuals lf_{base} is studied by using all the previously chosen parameters, by using MOEA with surrogate method.

A. SPEA2 AND USPEA2 PARAMETERS

First of all, the core parameters of the genetic algorithm need to be set to reasonable values, leading to good results without any objective approximation method. To evaluate the impact of their value on the overall results, SPEA2 is used with the scenario *Normal* and 80 time steps in the window (T). Their resulting metrics (hypervolume and average distance to Pareto front) are measured after 100, 1600 and 6400 evaluations. The probability of crossover pb_{cross} is chosen between 0 (mutation only) and 1 (crossover only). For the probability of individual gene mutation pb_{indmut} , it is chosen such as the average number of values (time step) mutated in an individual vary between 1 and T following a geometrical progression. As genetic algorithms results have an inherent part of randomness, especially for a low number of iterations, each combination of parameters is repeated 10 times. The values shown below are therefore the average metrics over the 10 runs.

Figure 7 presents all the results in the form of heatmaps. Results for a very low number of iterations are expected to be very noisy. Hence, the results after 100 evaluations are not shown here, with all hypervolume values being around 0.38 ± 0.01 and GD between 0.136 ± 0.001 and without any clear pattern. With longer execution of the genetic algorithm, correlations between performance indicators and hyperparameters appear.

For the hypervolume indicator, the worst results are reached when either pb_{cross} is close to 1.0 (no mutation) or when pb_{indmut} is close to 1.0 (mutation causing all the power profile to change), both for medium amount of evaluations (Figure 7a) and important amount (Figure 7c). The best results for 1600 evaluations are achieved with $0.03 \leq pb_{indmut} \leq 0.173$ and $0 \leq pb_{cross} \leq 0.6$. For a larger amount of evaluations, a lower value of pb_{indmut} seems preferable ($0 \leq pb_{indmut} \leq 0.072$) as well as a value of pb_{cross} closer to 0.5.

When it comes to the average distance to Pareto front (GD), the disparity is less important than hypervolume even for higher number of evaluations as presented in Figure 7b (1600 evaluations) and Figure 7d (6400 evaluations). A noticeable improvement is however observable with crossover rate pb_{cross} between 0.5 and 0.8 along with non-extreme values of pb_{indmut} .

As our experiments will focus on low and moderated amount of evaluations, we chose values in favorable range for both hypervolume and GD at 1600 evaluations, while taking into account the results for 6400 evaluations in order to choose in this acceptable range. In all the following experiments, we use values of $pb_{cross} = 0.5$ and $pb_{indmut} = 0.072$.

B. SURROGATE METHODS PARAMETERS

Four different surrogate methods are explored. The RBF with *cubic* function is parameter-free, but each of the three others has various parameters to tune for them to perform well. The use of surrogate models of a complex objective function is well studied for single-objective evolutionary optimization, especially when the surrogate model is constructed offline. In such cases, the surrogate (or approximation function) may be studied as a prediction function and a lot of common statistical analysis may be performed to evaluate the quality of the approximations. For instance, the R^2 or the root-mean-square deviation (RMSD) of the surrogate may be computed using a well distributed set of samples. It may be assumed that a surrogate model leading to higher R^2 (or lower RMSD) will perform better. In the case of MOEA and/or online methods, it is more difficult to evaluate an approximation method and very little literature exists on the topic.

Furthermore, the two new surrogate methods proposed may either return an approximated value or indicate a “miss”, considering that the known solutions set is not appropriate to perform the approximation. We, therefore, choose to use two different metrics for measuring the quality of an approximation method: the accuracy of the approximations themselves and their *miss rate*. The impact of the parameters of each approach on the accuracy of the PDM and ITDM utilities is evaluated independently. While accounting the inherent limits of independent evaluation once used in a multi-objective setup, it allows choosing reasonable values for each parameter.

1) INDEPENDENT EVALUATION

To obtain a realistic dataset of power profiles and objective values, the individuals generated during SPEA2 runs are recorded and stored. These profiles were obtained using SPEA2 with the settings presented in section VI-A (80 time steps with *Normal* scenario), 1600 evaluations and the chosen values for pb_{cross} and pb_{indmut} . This is repeated for 10 different runs, leading to a set of 8000 known profiles for both ITDM and PDM utilities.

Using this dataset, each surrogate method is evaluated with different parameters. A subset of 100 profiles is randomly selected in the dataset to form the set of *known solutions*, while all the others are used to perform approximations and compare to the real utility. For AD and MHT, the number of these profiles for which the approximation algorithm is not successful is counted as *missed*. This process is repeated 10 times for each combination of parameters in order to have robust results. Two metrics are obtained by averaging results for a combination of parameters: the RMSD between




FIGURE 7. Performance of SPEA2 depending on crossover rate and independent mutation probability, for different evaluation budgets. Higher is better for hypervolume, lower is better for GD.

successful approximation and real utility value of a profile, and the ratio of missed approximations for AD and MHT.

a: MHT AND AD

Our two surrogate methods are both controlled by two parameters with similar meanings, r_{close} and n_{close} . The value of each parameter is selected among 6 possibilities, following a geometrical progression. Values of n_{close} vary from 1 to 32 for both MHT and AD. The distance metric is however not the same between the two approaches: it is a derivative of Euclidean distance for MHT while it is a mean square error (MSE) for AD, which grows much faster than Euclidean distance. Different ranges of values are therefore studied for r_{close} , both obtained as geometric series of ratio 2, from 1/24 to 4/3 for MHT and from 1/144 to 2/9 for AD. Their start terms may appear somehow arbitrary for historical reasons, but are capturing accurately the different behaviors of the surrogate methods.

Figure 8 shows the results for the Multiresolution Haar Transform (MHT) approach. As the miss rates are the same when approximation is performed for either ITDM utility or PDM utility, they are given in a single figure 8c. Small values of r_{close} and/or large values of n_{close} lead to very important miss rates (up to 100%). This is expected, as both cause the requirement for performing the approximation to be more difficult to meet (closer solutions required for large n_{close} and smaller distance to known solutions for small r_{close}).

The quality of approximations is presented in Figure 8a for the PDM utility. As only the successful (*i.e.*, not missed) approximations are considered, it must be highlighted that combinations with miss rates close to 1.0 may have their quality computed with a very small number of samples. The worst values are obtained for large values of r_{close} (≥ 4) and for $n_{close} = 1$. While the scale of utility values (and therefore the RMSD) is not the same, the approximation of ITDM utilities gives very similar patterns, as shown in Figure 8b.

Ideal values for n_{close} and r_{close} would perform an approximation of all the profiles matching exactly the real utility value for this power profile (both miss rate and RMSD as close to 0 as possible). As none of the tested value combinations is perfect, a trade-off must be chosen. The combination selected is the best one for quality with a constraint on miss rate $\leq 3\%$, resulting for the MHT method in $n_{close} = 4$ and $r_{close} = \frac{1}{6}$.

For Average Distance (AD) approximation method, similar results are summarized similarly in Figure 9. The resulting miss rates shown in Figure 9c differ a lot compared to MHT results. This is expected as the distance metric used by both methods is not the same. Low miss rates ($\leq 5\%$) are only achieved with the highest value of r_{close} . A few other combinations result in miss rates lower than 90%, by requiring few close solutions to return an approximation ($n_{close} \leq 2$). Figures 9a and 9b present the quality of successful approximations for PDM and ITDM utility functions respectively.

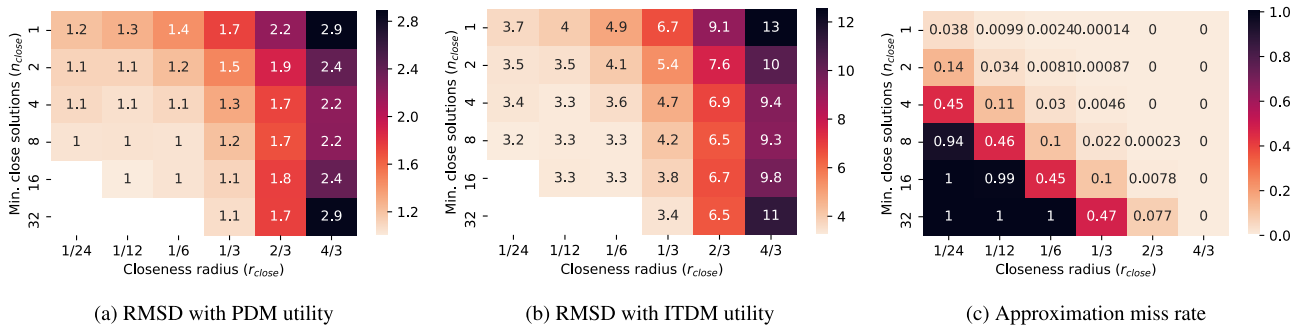


FIGURE 8. Miss rate and root-mean-square deviation (RMSD) for the Multiresolution Haar Transform (MHT) approximation method using 100 known profiles.

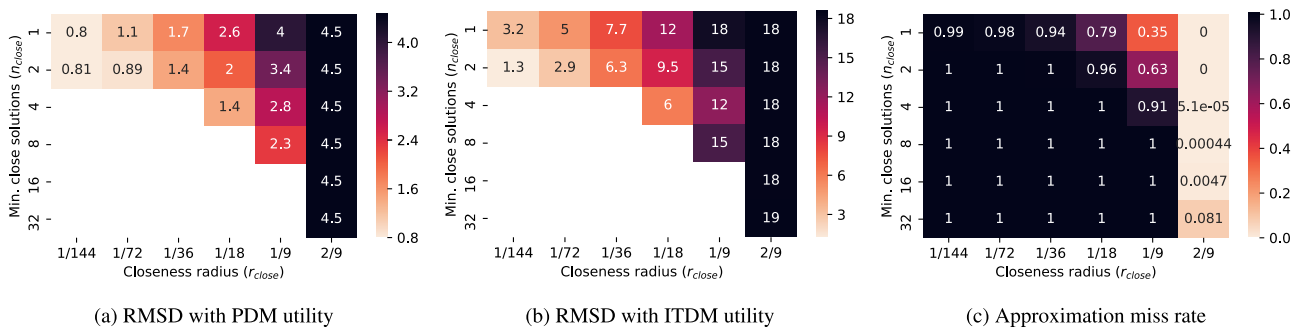


FIGURE 9. Miss rate and root-mean-square deviation (RMSD) for the Average Distance (AD) approximation method using 100 known profiles.

Similarly to MHT results, the RMSD is lower for small values of r_{close} and n_{close} , with a more important deviation for $r_{close} = 8$. However, the values suggest a worse approximation quality for all parameter combinations compared to MHT ones, except for some values which lead to very high miss rates ($\geq 99\%$).

Contrary to MHT results, where it was easy to find values with both a relatively good approximation quality and a low miss rate, no such combination appears for AD. As it seems important to have a relatively good approximation quality, a different trade-off is used to choose the parameters. While leading to a high (80%) miss rate, parameters $n_{close} = 1$ and $r_{close} = \frac{1}{18}$ lead to a relatively good approximation quality.

b: MULTILAYER PERCEPTRON HYPERPARAMETERS

Neural networks in general are known to have many hyperparameters that are difficult to setup without testing many configurations for a given problem. Here, we focus only on a subset of them: the activation function and the shape of the hidden layers. The solver L-BFGS [29] showed better and more stable results in preliminary experiments compared to stochastic methods such as Adam and is therefore used in the presented experiments.

To study their impact, we proceed similarly to the previous analysis of AD and MHT methods, without the need to care about approximation misses as an MLP gives an approximation as long as the network was trained. The results of two

activation functions are detailed, *relu* and *tanh*. Additional experiments with a *logistic* activation were made and led to results slightly different but with many similarities with *tanh*. The structure of the hidden layers is studied by setting the number of layers (between 1 and 16) and the number of neurons on each layer (between 1 and 32).

Figure 10 presents the results for both activation functions, showing RMSD for PDM utility on the left and for ITDM utility on the right. For recall, the scale of ITDM utility values being higher than PDM ones on the studied profiles, higher RMSD are expected for the first.

We clearly see that *tanh* performs badly on this problem, with results on ITDM utility on Figure 10d being significantly worse than other surrogate methods for all tested parameters. Its accuracy for PDM utility, detailed in Figure 10c, is comparatively better but stays behind the ones of *relu*. For *relu* activation function, some better results are observed especially for ITDM utility in Figure 10b. The absolute best accuracy is obtained for maximum values of parameters (64 neurons per hidden layer and 16 layers), but close values are obtained by lowering this high number of layers. For the PDM utility, shown in Figure 10a, better values are reached with 8 or more neurons per hidden layer and 8 or more layers.

A set of parameters that leads to a good trade-off between ITDM and PDM accuracy is using *relu* activation with 8 layers of 32 neurons. These figures seem really high considering our problem, especially the number of layers as many works

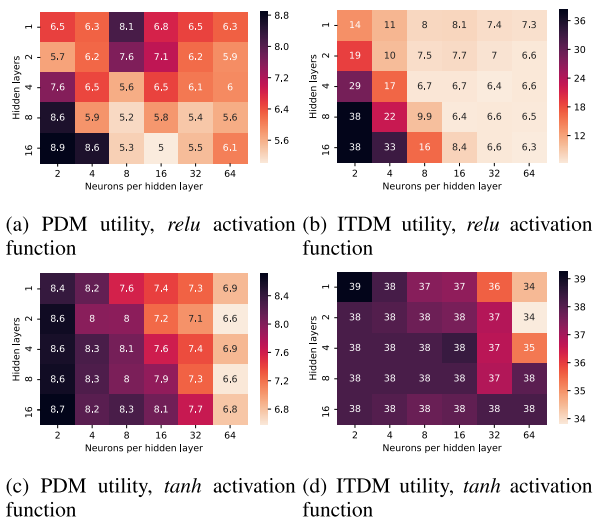


FIGURE 10. Root-mean-square deviation (RMSD) for different hidden layers and activation function of Multilayer Perceptron (MLP), with 100 known profiles.

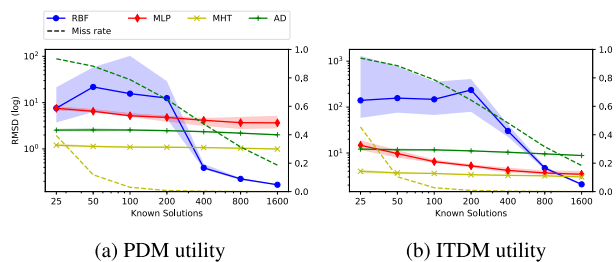


FIGURE 11. Median and spreading (top and bottom deciles) of the root-mean-square deviation of the different selected surrogate, depending on the number of samples of the real utility values (known profiles). Miss rates of MHT and AD are given on the right axis, as dashed lines.

suggest that 2 or 3 hidden layers are generally sufficient for approximating any function. However, a small training set is rarely considered, with in our case only slightly more samples than input neurons.

c: IMPACT OF NUMBER OF KNOWN SOLUTIONS ON ACCURACY

To help understanding better the behavior of the surrogate methods, an additional set of experiments is performed using a similar setup. Using the selected parameters for each approach, the accuracy and miss rates are studied for different sizes of known solution pool (training set). Similarly to the previous experiments, a number of solutions are selected randomly in the large pool of 8000 profiles and the rest is used for evaluation. Each combination of surrogate method and number of known solutions leads to 10 experiments with different training sets.

Results are given in Figure 11 for both ITDM and PDM utilities, as log-log plot and with the spreading displayed as a surface around the median value. It shows first that MHT and AD only slightly improve their accuracy when more

solutions are known. However, their RMSD is quite good compared to the other methods for few known solutions, especially MHT which outperforms all the other methods for up to 200 samples for PDM utility and 800 for ITDM one. By looking at their miss rates, it seems to confirm that the concept of approximation miss helps to keep the accuracy good when the training set is small.

The RBF method exhibits unique behaviors among the different surrogate methods. The first noticeable one is the huge spreading of low number of known solutions (≤ 200). This is a good indication that it is not adapted for such conditions, being highly unstable as small changes in the data may cause high variations of the radial function weights. The second observation is its very high accuracy for large budgets (≥ 400 for PDM and ≥ 1600 for ITDM utility). Past this amount, it largely outperforms the other methods and behave in a stable way.

Finally, the multilayer perceptron method gives unequal results on PDM and ITDM utilities. For the PDM utility, it slightly improves with the number of known solutions, but becomes more and more unstable and never performs better than AD or MHT. This may be a consequence of a badly adapted method for this problem combined with some over-fitting phenomenon [41] (neural network trying to fit “too well” the set of known solutions and resulting in inaccurate approximation for other inputs).

C. LIFETIME OF APPROXIMATION

The last unset parameter of the approach is lf_{base} (initial lifetime of approximated objectives) which is meaningful only when combining both optimization and approximation algorithms. Using the previously detailed values for all the other parameters, USPEA2 with the approximation algorithms are evaluated with end condition set to either 100 or 400 evaluations. The initial population is generated using the *Random* scheme. The experiments are performed 10 times for each settings. A total of 4 values of lf_{base} are chosen between 1 and 16, following a geometrical progression.

Higher values of lf_{base} mean that an individual with approximated objective functions needs to survive more generations in order to be evaluated with the real objective functions. While this allows to reduce the number of required evaluations, it also increases the risk of keeping bad individuals longer.

Figure 12 shows the resulting normalized hypervolume (to maximize) and generational distance (to minimize) of the final solution set for each approach. For a low target of objective functions evaluations ($eval_{max} = 100$), approaches either do not improve significantly (AD and RBF) or give better results with higher lf_{base} , as shown by figures 12a and 12c. This is especially visible for MHT, for which the resulting nHV increases from 0.596 for $lf_{base} = 1$ to 0.661 for $lf_{base} = 16$. Similarly, for multilayer perceptrons surrogate, the average distance to Pareto front (GD) is continuously decreasing down to 0.068.

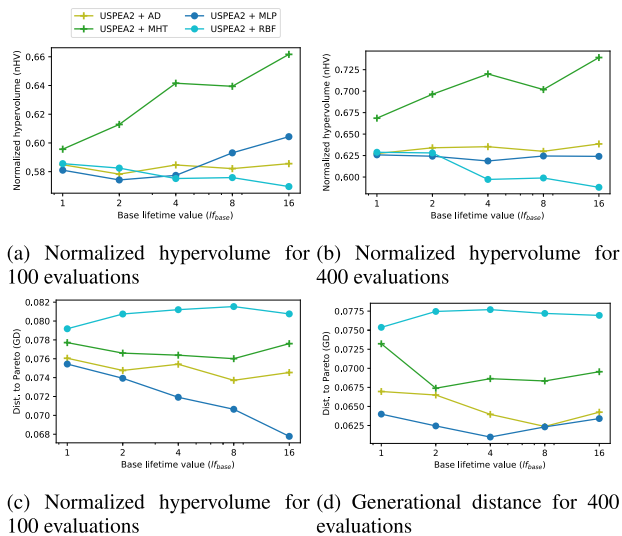


FIGURE 12. Impact of the base lifetime (lf_{base}) value on hypervolume and generational distance of resulting Pareto approximation when using the different surrogate methods.

TABLE 3. Parameters used for SPEA2 and USPEA2.

Symbol	Meaning	Value
pb_{cross}	Probability of crossover	0.5
pb_{indmut}	Probability of a timestep to change during a mutation	0.072
$archiveSize$	Maximum size of the archives	20
$populationSize$	Maximum size of the population	20

Figures 12b and 12d present the results for a higher evaluation budget ($eval_{max} = 400$). In this case, the benefit of a high value of lf_{base} is less clear, with only MHT showing clear improvement on hypervolume, with a small decrease at $lf_{base} = 8$. It is however not significant according to its standard deviation (0.045 for this value of lifetime). Contrary to the results obtained with a lower budget, the generational distance is almost not improved for $lf_{base} \geq 2$ and is even slightly impacted negatively for $lf_{base} \geq 8$.

Using a value of $lf_{base} \geq 8$ appears to give good results when USPEA2 is used. However, increasing the value of lf_{base} also tends to increase the number of generations of the genetic algorithm required to reach a solution, as more generations are needed for an approximated individual to be evaluated. Even if the execution time is expected to be bound by the evaluation of real objective functions themselves, the initial lifetime of approximated individuals should not be chosen uselessly too high. Increasing lf_{base} from 8 to 16 present an improvement for $eval_{max} = 400$ only possibly for MHT and also for MLP for lower budget ($eval_{max} = 100$). Therefore, the value used in the rest of the experiments presented here is $lf_{base} = 8$.

VII. EXPERIMENTAL RESULTS

In order to validate the proposed algorithms, several series of experiments are conducted, using the parameters defined along the previous section and summarized in tables 3 and 4. To evaluate the robustness of the algorithms with respect to

TABLE 4. Parameters used for each objective function surrogate method.

Method	Parameter	Value
MHT	r_{close}	1/6
	n_{close}	4
AD	r_{close}	1/18
	n_{close}	1
MLP	Neurons per layer	32
	Hidden layers	8
	Activation function	<i>relu</i>
RBF	Radial function	<i>cubic</i>

the definition of the objective functions, the impact of an important constant of the PDM utility function on the quality of the optimization is studied first. Then, the behavior of each heuristic depending on the initial population used is studied. This shows how different levels of initial knowledge of the solution space may impact the quality of the result for each heuristic. Finally, the results are compared for all the defined scenarios and different sizes of the solution space (number of steps in the time window, T) to show the robustness of the heuristics depending on the problem definition and how they scale for huge solution spaces.

A. IMPACT OF PDM UTILITY FORMULATION

The quality of the objective function approximation may depend on various characteristics of the real utility functions themselves. For instance, functions with clear non-linearity are usually more challenging to approximate, whereas very symmetrical ones may instead lead to better approximations. On the other hand, the formulation of utility functions also impacts directly the shape of the objective space of feasible solutions. This may lead to larger or sparse spaces, more difficult to explore with MOEA in general.

In the PDM utility as formulated in section V-B, the added utility for selling energy on the grid (ghg_{sell}) causes some kind of symmetry, depending on its value, compared to the utility removed for buying on the grid (ghg_{buy}). For instance, when $ghg_{sell} = ghg_{buy}$, any profile with the same amount of energy ($\sum_{t=0}^{T-1} p_t \cdot \delta$) will have the same PDM utility as the renewable energy may be sold directly to the grid when produced and bought at any time with no impact on resulting utility. However, this also extends the range of PDM utility values and consequently the size of the real Pareto front. By looking at the behavior of the approaches for different values of ghg_{sell} , it is possible to have an overview of the impact of such utility function property on the overall approaches. Different values are used, from 0 (selling energy for no benefit) to ghg_{sell} (selling a unit of energy compensates exactly the purchase cost).

Figure 13 gives the resulting normalized hypervolume for each value of ghg_{sell} , with end condition set to a budget of $eval_{max} = 100$ or $eval_{max} = 400$ evaluations of real objective functions. A general trend for all the heuristics is that values closer to ghg_{buy} result in a lower hypervolume, likely because of the increased objective space. The behaviors of each surrogate differ between the two budgets of evaluations.

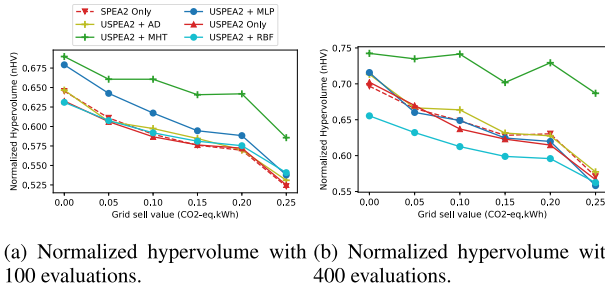


FIGURE 13. Impact of the grid sell utility parameter (ghg_{sell}) for the different approximation methods used with USPEA2.

Particularly, for $eval_{max} = 100$ MLP gives results higher than the baseline SPEA2 without surrogate with $ghg_{sell} = 0$ and becomes closer with increased values. With a higher budget, using multilayer perceptrons surrogate leads to no significant improvement of hypervolume. Similarly, USPEA2 with radial-basis functions (RBF) is very close to the baseline for low budget, but results in smaller values for $eval_{max} = 400$.

The proposed Multiresolution Haar Transform approach gives higher hypervolume in both cases. However, for a low evaluation budget, the relative improvement compared to SPEA2 without surrogate method is similar for all values of ghg_{sell} , whereas it increases with it when $eval_{max} = 400$. The other approaches, USPEA2 without surrogate and with AD, perform similarly to the baseline.

B. MAXIMUM NUMBER OF EVALUATIONS

The problem considered here is a black-boxes optimization with evaluation budget constraint. The more evaluations are allowed, the more resources and time are required to perform the optimization. To be able to take a decision in a reasonable amount of time, the main target is between 100 and 400 evaluations. To understand how this budget affects the quality of the resulting Pareto front approximation, a larger range is evaluated.

The results of each variant are compared for $eval_{max}$ between 50 and 1600 in Figure 14, with scenario *Normal* and 80 time steps in the time window. The variant *SPEA2 + None* (original SPEA2 algorithm without approximation method) is considered as the baseline for comparing the other results. The objective values of each individual from the initial population must be either evaluated or approximated using the surrogate model in the same way as any further new individual. Therefore, it must be highlighted that for small values of $eval_{max}$, a significant amount of evaluations may be spent *before* the first series of crossover and mutations. It may be up to $2 \cdot populationSize$ (one costly evaluation per decision module), or 40 evaluations in case none of them are successfully approximated.

For all the methods using USPEA2, the hypervolume increases with higher number of evaluations, as shown in Figure 14a. The relationship between hypervolume and $eval_{max}$ is almost logarithmic for all of them except RBF,

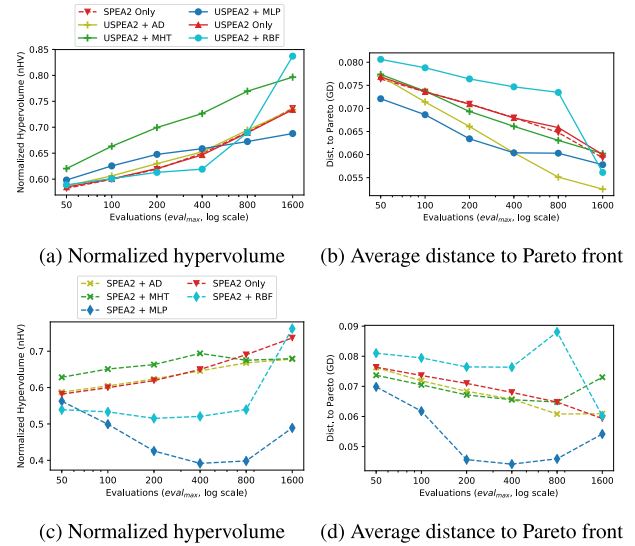


FIGURE 14. Performance of each variants for different evaluation budgets, with random initial individuals.

which increases quickly when $eval_{max}$ exceeds 400. The approximation method MHT lead to the highest hypervolume for all the values of $eval_{max}$ below the maximum tested of 1600, in which RBF gives slightly higher results. The relative improvement of MHT on hypervolume compared to the baseline is between 6.5% (for $eval_{max} = 50$) and 13% (for $eval_{max} = 200$). The other approaches are close to the baseline. The MLP surrogate methods, initially slightly upper the baseline, cross it after $eval_{max} = 400$ and ends up 6.6% below for the highest budget. For all the tested conditions, USPEA2 and SPEA2 without surrogate are giving similar values, with less than 0.5% relative difference.

Figure 14b shows the average distance to Pareto, or generational distance, which must be minimized. The global trend is a logarithmic decrease with the evaluation budget, with RBF diverging here too for $eval_{max} = 1600$. The Average-Distance (AD) and MLP surrogate methods result in lower generational distances than the baseline. The gap is increasing with $eval_{max}$ for AD and globally decreasing for MLP. On this metric, MHT is always close to the baseline (maximum relative difference of 3%). Without using surrogate methods, USPEA2 and SPEA2 result in very close generational distance (difference < 2%).

Results when using SPEA2 together with the different surrogates are presented in Figure 14c for hypervolume and Figure 14d for generational distance. Their behaviors are different from the ones observed with USPEA2, especially for high budgets ($eval_{max} \geq 400$). For all the values of $eval_{max}$, the hypervolume of the Pareto front approximation obtained is lower than with the USPEA2 algorithm, with the most noticeable degradation observed for MLP and RBF. The generational distance is however lower in some conditions, with a significant gap for MLP but also slightly better results for MHT. The explanation of this phenomenon, counter-intuitive at first, is discussed in the next section. As similar

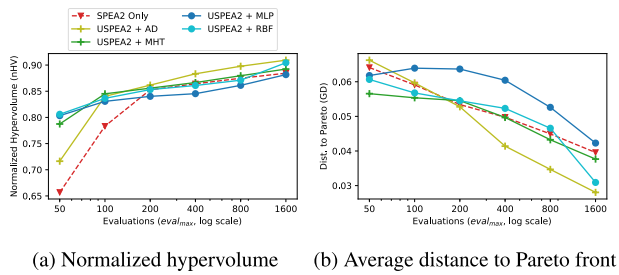


FIGURE 15. Evolution of the performance with increasing evaluation budget and *best* initial individuals.

results are observed in all other experiments, we focus only on USPEA2 when surrogate models are used in the rest of the section. With the same argument, SPEA2 and USPEA2 with no objective approximation method giving comparable results in all the other conditions studied (less than 2.5% of difference), only SPEA2 is given for comparison.

C. INITIAL POPULATION SCHEMES

In genetic algorithms, the choice of initial individuals may greatly impact the convergence speed. As an alternative to the fully *Random* scheme, the *Best* scheme includes an optimal solution of each decision module. As this individual is expected to be optimal for one of the objectives, it must be on the actual Pareto front.

The experiments here are identical to the ones from section VII-B except for the initial population scheme used. Similarly, the normalized hypervolume and the generational distance of each variant are presented in Figure 15 for different values of $eval_{max}$.

Compared to the use of the *Random* scheme in Figure 14, all the methods give higher hypervolume and lower *GD*. Both metrics also improve with the number of costly evaluations allowed $eval_{max}$ for all the USPEA2-based variants. Contrary to what is observed for the *random* scheme, there are important differences between surrogate methods on resulting hypervolume for low number of evaluations, up to 23% of the baseline value for $eval_{max} = 50$. In this setup, RBF and MLP give both the best results for such low evaluation number, followed closely by MHT (+20%) and finally, AD improves the baseline hypervolume by only 9%. This difference between surrogate methods is significantly lowered with higher budget of evaluations ($\pm 3\%$ for $eval_{max} = 1600$).

The generational distance shown in Figure 15b is also globally improved compared to *Random* initial population. The behaviors of MLP and RBF are however almost reversed, with MLP giving the worst (higher) generational distance than SPEA2 without surrogate for $eval_{max} \geq 100$. The AD method still exhibits a lower distance to Pareto front for moderate to high budgets ($eval_{max} \geq 400$).

D. SCENARIOS AND NUMBER OF TIME STEPS

In order to evaluate the behavior of the approximation methods in various situations, experiments are performed for each

power scenario defined in section V-D3.d. Another interesting point to show is the behavior of each approach when the size of the problem changes. As described in section II, the number of time steps in a power profile (T) is also the number of dimensions of the solution space. Therefore, the same 72-hours time window is considered with either 20, 80 or 320 time steps. The two targeted evaluation budgets, 100 and 400 are also studied.

The results for $eval_{max} = 100$ evaluations are shown in Figure 16. For all the algorithms, better results are always achieved for lower number of time steps, both for normalized hypervolume and generational distance. For both metrics, there is also a noticeable difference between each scenario, with *FewRenew* leading to better overall results and *ExcessRenew* to slightly worst ones, especially for *GD*. With the exception of RBF in a few cases, the resulting *GD* for a given scenario and number of time steps is close between all the variants. The combination of USPEA2 with our proposed MHT surrogate gives the highest hypervolume in all the conditions, with a lower gap compared to the other in the *Normal* scenario. Using multilayer perceptrons (MLP) does not improve significantly the hypervolume except for the scenario *FewRenew* and for large time window ($T = 320$).

Figure 17 presents the results for a higher budget of costly evaluation $eval_{max} = 400$. The same observation on the global impact of the scenarios is still accurate for both generational distance and hypervolume. A major difference between the results detailed for lower budget is the behavior of the Radial-Basis Functions (RBF) surrogate. It leads to the higher hypervolume values for all the scenarios when the number of time steps is low ($T = 20$) but inversely to worst results (both on hypervolume and generational distance) than the baseline SPEA2 for larger time windows. On all scenarios, MHT gives the highest hypervolume when medium or large time windows are considered and is below RBF for the short case $T = 20$. The generational distance of all approaches is close to the baseline in most conditions, with only AD giving systematically better values.

E. USING COMPLEX IT MODEL

The previous experiments were performed on a simplified infrastructure model, with therefore objective functions exhibiting less complex behaviors. Results presented here are instead using the RECO scheduler [9] with a more realistic machine and tasks model. The experimental setup consists in the *Normal* scenario, with $T = 80$ and a budget of 100 costly evaluations.

Figure 18 shows the solutions obtained by running the different approaches over a complete optimization run. The approximated Pareto front from each approach is highlighted, but dominated solutions are also plotted as individual points outside the front. While this result alone does not allow a deep analysis, it provides valuable clues. On this run, the capacity of MHT and MLP to spread across the extremum is clearly visible, while AD is only improving a very local area towards the maximum values. MHT is the only method which ends

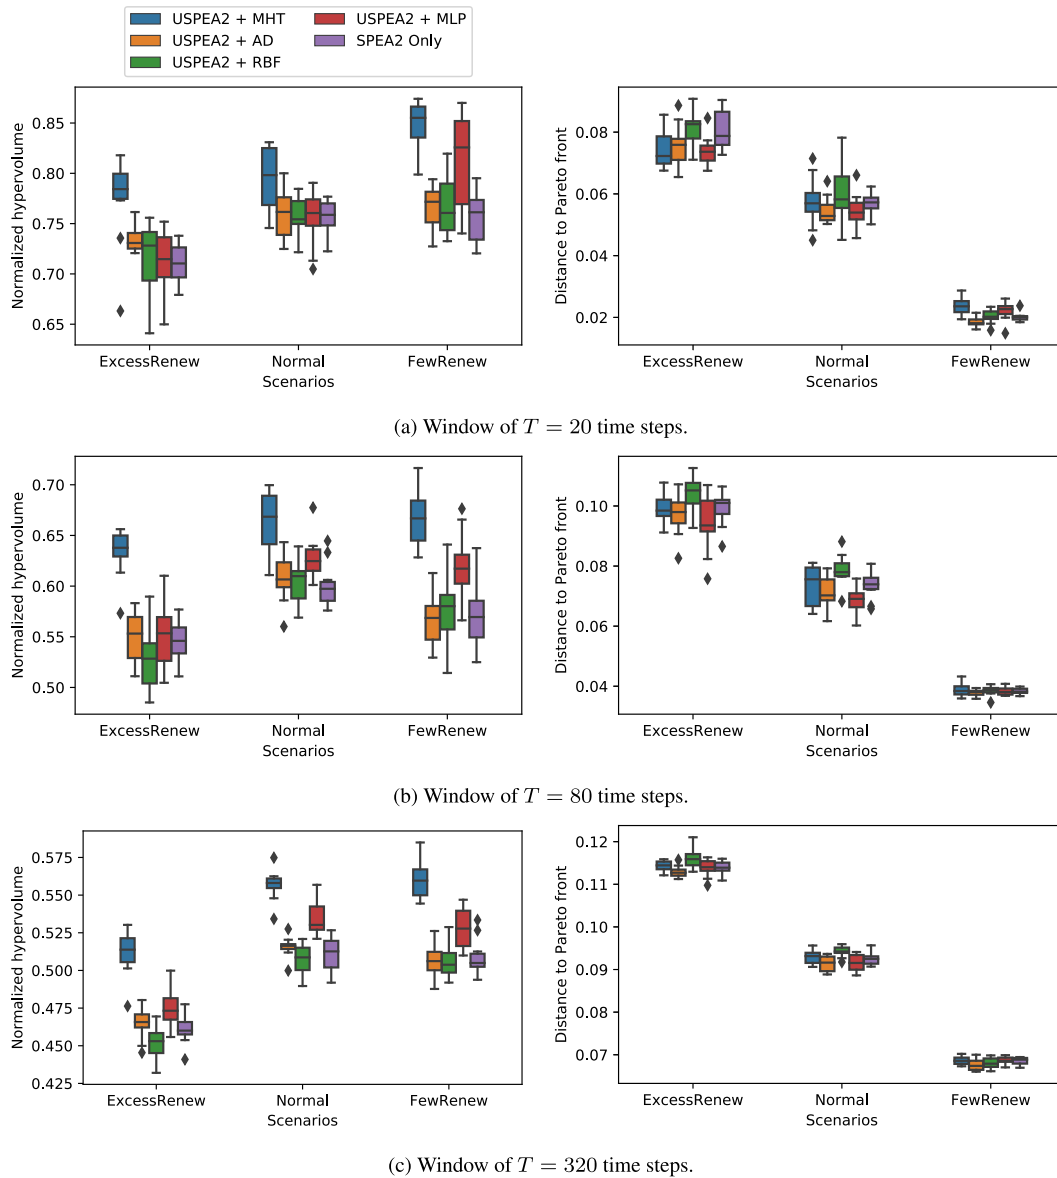


FIGURE 16. Median and quartiles of both metrics (normalized hypervolume and generational distance) for the different heuristics with a budget $eval_{max} = 100$ costly evaluations over several scenarios.

with all its solutions being Pareto-optimal between them. On the contrary, AD only 9 out of 20 solutions fulfill this condition. It can also be noted that AD and MHT both results in most of their solutions dominating the ones of SPEA2 without surrogate method.

As mentioned earlier, the high complexity of this IT model prevents obtaining the real Pareto front of the problem. To allow comparing the approaches with higher level metrics, some modifications of the hypervolume indicator and generational distance are proposed. The hypervolume can still be computed but not normalized with the real Pareto front. To keep its value in meaningful range, it is calculated after normalizing each objective value by considering the minimum and maximum found among all the experiments.

The generational distance requires a Pareto front to which measure the distances from each solution. While this is not commonly done in the literature, we propose for these experiments to extract the Pareto-optimal solutions among all the results to form the *best approximation* of the Pareto front PF^* . Generational distance is then computed as for the other experiments using this approximation instead of the MILP-based front.

These modified metrics are compared between the approaches in Figure 19, over 10 runs with the same parameters. Using the MHT surrogate method gives better results than the other approaches most of the time, for both hypervolume and generational distance metrics. On the hypervolume indicator, using any surrogate method leads to more spread

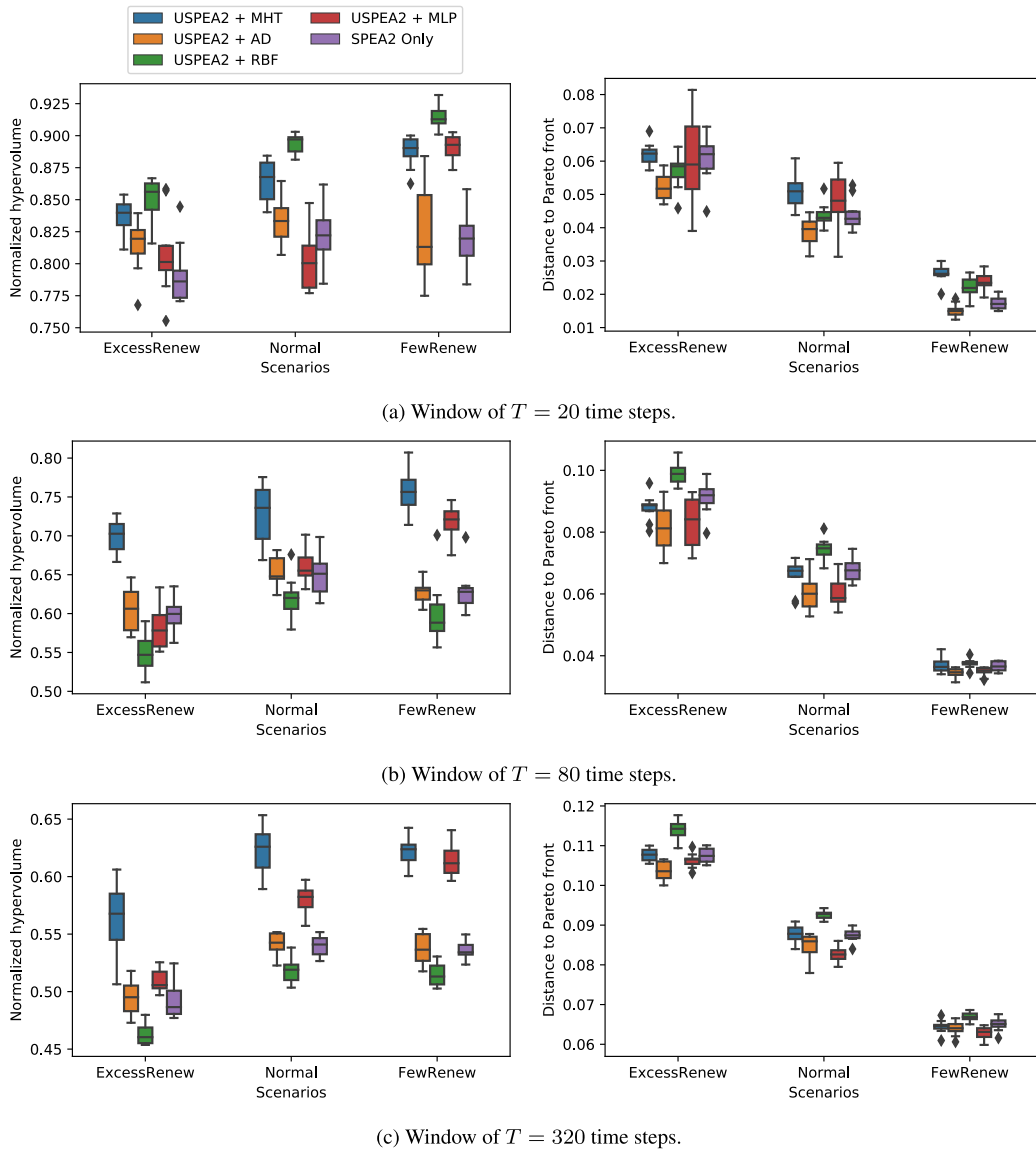


FIGURE 17. Median and quartiles of both metrics (normalized hypervolume and generational distance) for the different heuristics with a budget $eval_{max} = 400$ costly evaluations over several scenarios.




FIGURE 18. Resulting solutions and Pareto front approximation after 100 costly evaluations, using RECO scheduler from the literature and complex IT infrastructure.

results compared to SPEA2 alone. In addition, RBF and MLP appear to give worst results than the baseline at least half of the time. Finally, using USPEA2 with AD surrogate method leads to slightly better hypervolume and generational distance than the baseline in the nominal case.

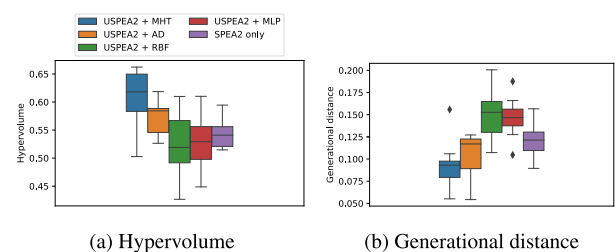


FIGURE 19. High level metrics for all the approaches using the RECO scheduler and complex infrastructure.

VIII. DISCUSSION

Many experiments are presented in the previous sections. Both aspects of the proposed approach, *i.e.* USEPA2 and the surrogate methods, are therefore evaluated across many different scenarios and conditions.

The proposed modifications of SPEA2 to design USPEA2 aim at improving the performance when the individuals consist in a mix of evaluated and approximated solutions (through an online-trained surrogate model). Several desirable properties of USPEA2 must be highlighted. First, no significant difference is observed compared to SPEA2 when no surrogate method is used, as shown by the results in section VII-B. The benefit of the proposed *reliable archive* nevertheless appears when any of the tested surrogates is used. Hence, USPEA2 leads to higher hypervolume in most conditions, especially for medium and high amount of evaluations ($eval_{max} \geq 400$). The impact of USPEA2 on the average distance to the Pareto front (GD) is less clear and depends more on the scenario and surrogate method used.

The average-distance based method (AD) is a trivial approach and appear to perform relatively badly when used alone, as detailed in section VI-B1. On the other hand, the multiresolution Haar wavelet distance method (MHT) is designed specifically to handle time series specificities.

Most experiments using MHT approximation give a better hypervolume. Notably, the results of Figure 14a show it reaches similar or higher hypervolume compared to SPEA2 alone with only a 4th of the costly evaluation budget. However it does not reduce significantly the generational distance in the experiments performed on the simplified model. This is also confirmed by the experiments made with the *Best* initial scheme, shown in Figure 15a, where MHT approximation appear to improve only slightly the hypervolume when the individuals are already well spread. Such results suggest that using MHT prevents costly evaluation of many individuals near to the known Pareto-optimal solutions, allowing a faster spreading, while failing to identify accurately non-dominated individuals among them (which are closer to the real Pareto front from only a small difference).

Compared to MHT, the behavior of the overall optimization is quite different when AD is used. While it leads to few or no improvement of hypervolume in most cases, it results in a slightly better GD for large budget (≥ 400) of costly evaluations. It also performs better than all the other methods when the individuals are already well spread, as shown by the *Best* initial population scheme. This is likely explained by a low ability to prevent individual evaluations close to the current approximation of the Pareto front.

Both RBF and MLP only improve the quality of the optimization in a limited amount of cases. The multilayer perceptrons surrogate gives good results compared to the baseline in a specific scenario (*FewRenew*) and for large solution space (time window with 320 time steps). Oppositely, using a surrogate based on radial-basis functions performs better with small solution space. RBF also gives very good results for both hypervolume and generational distance for the largest budget tested, by knowing 1600 solutions at the end. The two points confirm that RBF is remarkably well-suited when the ratio of the size of the training set over the number of dimensions of the problem is higher than our target.

The most important result is the *stability* of the proposed MHT method over the various tested conditions. First, it is less affected than the other methods by the formulation of the objective functions (scenarios and internal values such as ghg_{sell} used to compute the PDM utility). It is also the only method that gives significant improvements when a more complex IT infrastructure is considered, using an arguably realistic scheduler from the literature in place of the simplified model. Then it grants a significant improvement of hypervolume for any target budget of costly evaluations. Finally, it is clearly robust to the high-dimensionality aspect of large time series, resulting in relative hypervolume improvement similar or greater for values of T from 20 to 320. This new method may seem simpler than approaches such as RBF and MLP. It is however designed deeply for taking the time series aspect into account, mitigating the issues of high-dimensional space, with the results confirming this initial hypothesis.

IX. CONCLUSION AND FUTURE WORKS

In this article, we presented and validated with extensive experiments an approach for multi-objective power planning optimization between a datacenter and an electrical infrastructure with on-site RESs, each having its own objective and being considered as a black-box. Multi-objective heuristics being especially costly, a surrogate-assisted method was explored in order to reduce the amount of evaluation of solutions required to find a set of good power planning trade-offs. Two surrogate methods are proposed, Average Distance (AD) and Multiresolution Haar Transform (MHT), designed to deal efficiently with time series. In addition, a variant of the multi-objective evolutionary algorithm SPEA2, named USPEA2 is presented to handle better a mix of individuals with evaluated objectives and unreliable, approximated values.

By improving the quality and the stability of the results when a surrogate method is used, USPEA2 proved to be a robust MOEA for leveraging online objective function approximation. The proposed MHT surrogate method led to a clear improvement of the quality of the solution set when used together with USPEA2 in all tested conditions compared to other methods and absence original SPEA2 algorithm without surrogate. By taking the time series aspect of the solution space, this method gives better results than techniques known to provide good quality surrogates in the literature, such as multilayer neural networks and radial-basis functions. The evaluation using a more complex and arguably rich datacenter and workload model, with a scheduler from the literature confirms the benefit of the proposed approach in real use cases.

The work presented here is also the first to attempt a multi-objective optimization of datacenter and rich electrical infrastructure while being totally agnostic of the respective models. Such a black-boxes model has tangible benefits in real-world situations, by reducing the coupling between electrical and IT aspects, which are in practice managed by separated entities and evolving independently.

The multi-objective aspect, in a Pareto sense, avoids to decide the articulation of objectives beforehand and allows to use arbitrary objectives for the IT and the electrical parts. In this work, we only focused on 2 objectives, one for each part of the infrastructure. However, the approach itself supports more objectives. A possibly interesting extension is the addition of an objective common to both parts such as the economical aspects.

The choice of a solution among the best trade-offs (Pareto-optimal solutions) has not been studied here. This approximation of the Pareto-front can be used directly by a human decision maker or to implement automated policies. By providing a set of good trade-offs, an automatic decision making algorithm can take into account both *a priori* articulation of objectives and *a posteriori* knowledge (shape of the Pareto-front, knee points [6], [45], extremum solutions). Even with the improvements proposed here, finding the Pareto-front with costly ITDM and PDM is quite time consuming. Hence a possible use case is in a two-phases optimization process such as proposed by Pahlevan et al. [32], in which it can take a good offline decision for a middle-term horizon and is completed by cheap online heuristic to adjust the decision to small deviations from the power planning.

When a longer period than the time window is considered, our approach can be used in a multi-objective sequential decision making problem. At each decision epoch, a solution (power planning for the upcoming time window) must be chosen among a set of Pareto-optimal trade-offs. In the presented work, the black-box model is used to hide these states, making it hard to reason on a time longer than a decision epoch. Future studies could explore how to tackle this long-term decision making problem with different amounts of information on the ITDM and PDM states.

ACKNOWLEDGMENT

Experiments presented in this paper were carried out on Grid5000 testbed, supported by a scientific interest group hosted by Inria, including CNRS, RENATER, and several Universities and organizations (<http://www.grid5000.fr>).

REFERENCES

- [1] S. Aghabozorgi, A. S. Shirkhorshidi, and T. Y. Wah, "Time-series clustering—A decade review," *Inf. Syst.*, vol. 53, pp. 16–38, Oct. 2015, doi: [10.1016/j.is.2015.04.007](https://doi.org/10.1016/j.is.2015.04.007).
- [2] T. Akhtar and C. A. Shoemaker, "Multi objective optimization of computationally expensive multi-modal functions with RBF surrogates and multi-rule selection," *J. Global Optim.*, vol. 64, no. 1, pp. 17–32, Jan. 2016, doi: [10.1007/s10898-015-0270-y](https://doi.org/10.1007/s10898-015-0270-y).
- [3] E. Alsema, "Energy payback time and CO₂ emissions of PV systems," in *Practical Handbook of Photovoltaics*, A. McEvoy, T. Markvart, and L. Castañer, Eds. Boston, MA, USA: Academic, Jan. 2012, pp. 1097–1117, doi: [10.1016/B978-0-12-385934-1.00037-4](https://doi.org/10.1016/B978-0-12-385934-1.00037-4).
- [4] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*. Hoboken, NJ, USA: Wiley, 2015.
- [5] G. Bramerdorfer and A.-C. Zavoianu, "Surrogate-based multi-objective optimization of electrical machine designs facilitating tolerance analysis," *IEEE Trans. Magn.*, vol. 53, no. 8, pp. 1–11, Aug. 2017.
- [6] J. Branke, K. Deb, H. Dierolf, and M. Osswald, "Finding knees in multi-objective optimization," in *Parallel Problem Solving From Nature* (Lecture Notes in Computer Science), X. Yao et al., Eds. Berlin, Germany: Springer, 2004, pp. 722–731.
- [7] A. E. I. Brownlee, J. R. Woodward, and J. Swan, "Metaheuristic design pattern: Surrogate fitness functions," in *Proc. Companion Publication Annu. Conf. Genetic Evol. Comput.*, New York, NY, USA, Jul. 2015, pp. 1261–1264, doi: [10.1145/2739482.2768499](https://doi.org/10.1145/2739482.2768499).
- [8] *California GHG Emission Inventory Edition*, California Environ. Protection Agency, Sacramento, CA, USA, 2016.
- [9] S. Caux, P. Renaud-Goud, G. Rostirolla, and P. Stolf, "IT optimization for datacenters under renewable power constraint," in *Euro-Par 2018: Parallel Processing* (Lecture Notes in Computer Science), M. Aldinucci, L. Padovani, and M. Torquati, Eds. Cham, Switzerland: Springer, 2018, pp. 339–351.
- [10] F. K. Chan, A. W.-C. Fu, and C. Yu, "Haar wavelets for efficient similarity search of time-series: With and without time warping," *IEEE Trans. Knowl. Data Eng.*, vol. 15, no. 3, pp. 686–705, May 2003, doi: [10.1109/tkde.2003.1198399](https://doi.org/10.1109/tkde.2003.1198399).
- [11] K.-P. Chan and A. W.-C. Fu, "Efficient time series matching by wavelets," in *Proc. 15th Int. Conf. Data Eng.*, 1999, pp. 126–133.
- [12] K. H. Chen and Z. D. Ding, "Lithium-ion battery lifespan estimation for hybrid electric vehicle," in *Proc. 27th Chin. Control Decis. Conf. (CCDC)*, May 2015, pp. 5602–5605, doi: [10.1109/ccdc.2015.7161797](https://doi.org/10.1109/ccdc.2015.7161797).
- [13] C. Curry, "Lithium-ion battery costs and market," *Bloomberg New Energy Finance*, vol. 5, pp. 4–6, Jul. 2017.
- [14] K. Deb, "Multi-objective optimization," in *Search Methodologies*. Cham, Switzerland: Springer, 2014, pp. 403–449.
- [15] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable test problems for evolutionary multiobjective optimization," in *Evolutionary Multiobjective Optimization*. Cham, Switzerland: Springer, 2005, pp. 105–145.
- [16] V. Devabhaktuni, M. Alam, S. S. S. R. Depuru, R. C. Green, D. Nims, and C. Near, "Solar energy: Trends and enabling technologies," *Renew. Sustain. Energy Rev.*, vol. 19, pp. 555–564, Mar. 2013, doi: [10.1016/j.rser.2012.11.024](https://doi.org/10.1016/j.rser.2012.11.024).
- [17] S. A. Dudani, "The distance-weighted k-nearest-neighbor rule," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-6, no. 4, pp. 325–327, Apr. 1976, doi: [10.1109/TSMC.1976.5408784](https://doi.org/10.1109/TSMC.1976.5408784).
- [18] B. Dunn, H. Kamath, and J.-M. Tarascon, "Electrical energy storage for the grid: A battery of choices," *Science*, vol. 334, no. 6058, pp. 928–935, 2011, doi: [10.1126/science.1212741](https://doi.org/10.1126/science.1212741).
- [19] N. Goiri, W. Katsak, K. Le, D. Thu Nguyen, and R. Bianchini, "Parasol and GreenSwitch: Managing datacenters powered by renewable energy," in *Proc. 18th Int. Conf. Architectural Support Programming Lang. Operating Syst. (ASPLOS)*, New York, NY, USA, 2013, pp. 51–64, doi: [10.1145/2451116.2451123](https://doi.org/10.1145/2451116.2451123).
- [20] G. Cook, J. Lee, T. Tsai, A. Kong, J. Deans, B. Johnson, and E. Jardim, "Clicking clean: Who is winning the race to build a green internet?" Greenpeace, Washington, DC, USA, Tech. Rep. 5, Jan. 2017.
- [21] H. Hao, Z. Mu, S. Jiang, Z. Liu, and F. Zhao, "GHG emissions from the production of lithium-ion batteries for electric vehicles in China," *Sustainability*, vol. 9, no. 4, p. 504, Apr. 2017, doi: [10.3390/su9040504](https://doi.org/10.3390/su9040504).
- [22] Y. Jin, "A comprehensive survey of fitness approximation in evolutionary computation," *Soft Comput.*, vol. 9, pp. 3–12, Jan. 2005, doi: [10.1007/s00500-003-0328-5](https://doi.org/10.1007/s00500-003-0328-5).
- [23] M. K. Karakasis and K. C. Giannakoglou, "On the use of metamodel-assisted, multi-objective evolutionary algorithms," *Eng. Optim.*, vol. 38, no. 8, pp. 941–957, Dec. 2006.
- [24] A. Khosravi, A. Nadjaran Toosi, and R. Buyya, "Online virtual machine migration for renewable energy usage maximization in geographically distributed cloud data centers," *Concurrency Comput., Pract. Exper.*, vol. 29, no. 18, p. e4125, May 2017, doi: [10.1002/cpe.4125](https://doi.org/10.1002/cpe.4125).
- [25] F. Kong and X. Liu, "A survey on green-energy-aware power management for datacenters," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 1–38, Jan. 2015, doi: [10.1145/2642708](https://doi.org/10.1145/2642708).
- [26] N. Kumar, G. S. Aujla, S. Garg, K. Kaur, R. Ranjan, and S. K. Garg, "Renewable energy-based multi-indexed job classification and container management scheme for sustainability of cloud data centers," *IEEE Trans. Ind. Informat.*, vol. 15, no. 5, pp. 2947–2957, May 2019, doi: [10.1109/TII.2018.2800693](https://doi.org/10.1109/TII.2018.2800693).
- [27] H. Lei, R. Wang, T. Zhang, Y. Liu, and Y. Zha, "A multi-objective co-evolutionary algorithm for energy-efficient scheduling on a green data center," *Comput. Oper. Res.*, vol. 75, pp. 103–117, Nov. 2016, doi: [10.1016/j.cor.2016.05.014](https://doi.org/10.1016/j.cor.2016.05.014).

- [28] Y. Li, A.-C. Orgerie, and J.-M. Menaud, "Balancing the use of batteries and opportunistic scheduling policies for maximizing renewable energy consumption in a cloud data center," in *Proc. 25th Eur. Int. Conf. Parallel, Distrib. Netw.-Based Process. (PDP)*, Mar. 2017, pp. 408–415, doi: [10.1109/PDP.2017.24](https://doi.org/10.1109/PDP.2017.24).
- [29] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Math. Program.*, vol. 45, no. 1, pp. 503–528, 1989, doi: [10.1007/BF01589116](https://doi.org/10.1007/BF01589116).
- [30] S. Z. Martinez and C. A. C. Coello, "Combining surrogate models and local search for dealing with expensive multi-objective optimization problems," in *Proc. IEEE Congr. Evol. Comput.*, Jun. 2013, pp. 2572–2579, doi: [10.1109/cec.2013.6557879](https://doi.org/10.1109/cec.2013.6557879).
- [31] National Renewable Energy Laboratory. (2012). *Solar Power Data for Integration Studies*. [Online]. Available: <https://www.nrel.gov/grid/solar-power-data.html>
- [32] A. Pahlevan, M. Rossi, G. P. Del Valle, D. Brunelli, and D. Atienza, "Joint computing and electric systems optimization for green datacenters," in *Handbook Hardware/Software Codesign*. Dordrecht, The Netherlands: Springer, 2017, pp. 1163–1183, doi: [10.1007/978-94-017-7267-9_35](https://doi.org/10.1007/978-94-017-7267-9_35).
- [33] D. Paul, W.-D. Zhong, and S. K. Bose, "Demand response in data centers through energy-efficient scheduling and simple incentivization," *IEEE Syst. J.*, vol. 11, no. 2, pp. 613–624, Jun. 2017, doi: [10.1109/JSYST.2015.2476357](https://doi.org/10.1109/JSYST.2015.2476357).
- [34] J. Peng, L. Lu, and H. Yang, "Review on life cycle assessment of energy payback and greenhouse gas emission of solar photovoltaic systems," *Renew. Sustain. Energy Rev.*, vol. 19, pp. 255–274, Mar. 2013, doi: [10.1016/j.rser.2012.11.035](https://doi.org/10.1016/j.rser.2012.11.035).
- [35] J.-M. Pierson, G. Baudic, S. Caux, B. Celik, G. Da Costa, L. Grange, M. Haddad, J. Lecuire, J.-M. Nicod, L. Philippe, V. Rehn-Sonigo, R. Roche, G. Rostirolla, A. Sayah, P. Stolf, M.-T. Thi, and C. Varnier, "DATAZERO: Datacenter with zero emission and robust management using renewable energy," *IEEE Access*, vol. 7, pp. 103209–103230, 2019, doi: [10.1109/ACCESS.2019.2930368](https://doi.org/10.1109/ACCESS.2019.2930368).
- [36] R. G. Regis, "Evolutionary programming for high-dimensional constrained expensive black-box optimization using radial basis functions," *IEEE Trans. Evol. Comput.*, vol. 18, no. 3, pp. 326–347, Jun. 2014, doi: [10.1109/TEVC.2013.2262111](https://doi.org/10.1109/TEVC.2013.2262111).
- [37] N. Riquelme, C. Von Lücken, and B. Baran, "Performance metrics in multi-objective optimization," in *Proc. Latin Amer. Comput. Conf. (CLEI)*, Oct. 2015, pp. 1–11.
- [38] G. Rostirolla, "Scheduling cloud data center powered by renewable energy only with mixed phases-based workload," Ph.D. dissertation, Dept. Mathématiques, Informatique, Télécommunications de Toulouse (MITT), Univ. Toulouse, Toulouse, France, 2019.
- [39] N. Schilling, M. Wistuba, L. Drumond, and L. Schmidt-Thieme, "Hyperparameter optimization with factorized multilayer perceptrons," in *Machine Learning and Knowledge Discovery in Databases* (Lecture Notes in Computer Science), A. Appice, P. P. Rodrigues, V. S. Costa, J. Gama, A. Jorge, and C. Soares, Eds. Cham, Switzerland: Springer, 2015, pp. 87–103.
- [40] S. Shan and G. G. Wang, "Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions," *Struct. Multidisciplinary Optim.*, vol. 41, no. 2, pp. 219–241, Mar. 2010, doi: [10.1007/s00158-009-0420-2](https://doi.org/10.1007/s00158-009-0420-2).
- [41] L. Shi and K. Rasheed, "A survey of fitness approximation methods applied in evolutionary algorithms," in *Computational Intelligence in Expensive Optimization Problems, Adaptation Learning and Optimization*, Y. Tenne and C.-K. Goh, Eds. Berlin, Germany: Springer, 2010, pp. 3–28, doi: [10.1007/978-3-642-10701-6_1](https://doi.org/10.1007/978-3-642-10701-6_1).
- [42] R. S. Stanković and B. J. Falkowski, "The Haar wavelet transform: Its status and achievements," *Comput. Electr. Eng.*, vol. 29, no. 1, pp. 25–44, Jan. 2003, doi: [10.1016/S0045-7906\(01\)00011-8](https://doi.org/10.1016/S0045-7906(01)00011-8).
- [43] P. Van den Bossche, F. Vergels, J. Van Mierlo, J. Matheys, and W. Van Autenboer, "SUBAT: An assessment of sustainable battery technology," *J. Power Sources*, vol. 162, pp. 913–919, Jun. 2006, doi: [10.1016/j.jpowsour.2005.07.039](https://doi.org/10.1016/j.jpowsour.2005.07.039).
- [44] A. David Van Veldhuizen and B. Gary Lamont, "Evolutionary computation and convergence to a Pareto front," in *Proc. Conf. Late Breaking Papers Genetic Program.*, 1998, pp. 221–228.
- [45] H. Wang, M. Olhofer, and Y. Jin, "A mini-review on preference modeling and articulation in multi-objective optimization: Current status and challenges," *Complex Intell. Syst.*, vol. 3, no. 4, pp. 233–245, Dec. 2017. [Online]. Available: <http://link.springer.com/10.1007/s40747-017-0053-9>
- [46] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach," *IEEE Trans. Evol. Comput.*, vol. 3, no. 4, pp. 257–271, Nov. 1999, doi: [10.1109/4235.797969](https://doi.org/10.1109/4235.797969).
- [47] E. Zitzler, K. Deb, and L. Thiele, "Comparison of multiobjective evolutionary algorithms: Empirical results," *Evol. Comput.*, vol. 8, no. 2, pp. 173–195, 2000.
- [48] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the strength Pareto evolutionary algorithm," ETH, Swiss Federal Inst. Technol., Zürich, Switzerland, TIK-Rep. 103, 2001.



LÉO GRANGE received the master's degree in computer science, specialized in distributed systems and safety critical software from the University of Toulouse and the Ph.D. degree in computer science from the Institut de Recherche en Informatique de Toulouse (IRIT), University of Toulouse, in 2019. For his Ph.D. studies, he worked on energy-efficient scheduling for cloud computing and more specifically for datacenters powered with on-site renewable energy sources.



PATRICIA STOLF received the Ph.D. degree from INSA, Toulouse, France, in 2004. She is currently an Associate Professor in computer science with the University of Toulouse. She is also a member of the IRIT Laboratory. Her current research interests include large scale distributed systems, such as grid or clouds, distributed algorithms, autonomous computing, resources management, load-balancing, and energy-aware and thermal-aware placement. She has been involved in different research projects, such as the ACTION COST IC0804 "Energy Efficiency in Large Scale Distributed Systems," the European CoolEmAll Project, and the National ANR SOP Project. She is currently working on the ANR DATAZERO Project and studying how to manage the electricity and IT services in a datacenter operated with several green energy sources.



GEORGES DA COSTA received the Ph.D. degree from LIG, Grenoble, France, in 2005, and the Habilitation degree from University Paul Sabatier, Toulouse, France, in 2015. He is currently an Assistant Professor in computer science with the University of Toulouse. He is also a member of the IRIT Laboratory. His current research interests include energy-aware distributed systems, HPC and cloud computing, large scale energy-aware distributed systems, performance evaluation, and ambient systems. He serves on several PCs in the energy-aware systems, grid, and peer-to-peer fields. He was the Chair of the COST1305 working group on "Programming models and runtimes." He is also working in the French funded project, Energumen, on scheduling moldable application on HPC systems and in the French funded project, Datazero, on scheduling in renewable powered datacenters.



PAUL RENAUD-GOUD received the Ph.D. degree from the École Nationale Supérieure de Lyon, France, in 2012. After his postdoctoral research, he is with the Institut Mathématiques de Bordeaux, France, and Chalmers Tekniska Högskola, Sweden. He is currently an Assistant Professor with University of Toulouse, France, and feels close to Camille Noûs' research.