


Provided by the author(s) and University College Dublin Library in accordance with publisher policies. Please cite the published version when available.

Title	MILP for the Multi-objective VM Reassignment Problem
Author(s)	Saber, Takfarinas; Ventresque, Anthony; Marques-Silva, Joao; Thorburn, James; Murphy, Liam, B.E.
Publication date	2015-11-11
Publication information	Proceedings of the 2015 IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI)
Conference details	27th IEEE International Conference on Tools with Artificial Intelligence (ICTAI), Vietri Sul Mare, Italy, 9-11 November, 2015
Publisher	IEEE
Item record/more information	http://hdl.handle.net/10197/7205
Publisher's statement	© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works
Publisher's version (DOI)	http://dx.doi.org/10.1109/ICTAI.2015.20

Downloaded 2018-04-23T11:51:18Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa) 

Some rights reserved. For more information, please see the item record link above.



MILP for the Multi-objective VM Reassignment Problem

Takfarinas Saber*, Anthony Ventresque*[†], Joao Marques-Silva[‡], James Thorburn[§], Liam Murphy*

*Lero@UCD, School of Computer Science and Informatics, University College Dublin, Ireland

Email: takfarinas.saber@ucdconnect.ie, {anthony.ventresque, liam.murphy}@ucd.ie

[†]IBM Research, Damastown Industrial Estate, Dublin, Ireland

[‡]Instituto Superior Técnico, Lisboa, Portugal, Email: jpms@ist.utl.pt

[§]IBM Software Group, Toronto, Canada, Email: jthorbur@ca.ibm.com

Abstract—Machine Reassignment is a challenging problem for constraint programming (CP) and mixed integer linear programming (MILP) approaches, especially given the size of data centres. The multi-objective version of the Machine Reassignment Problem is even more challenging and it seems unlikely for CP or MILP to obtain good results in this context. As a result, the first approaches to address this problem have been based on other optimisation methods, including metaheuristics. In this paper we study under which conditions a mixed integer optimisation solver, such as *IBM ILOG CPLEX*, can be used for the Multi-objective Machine Reassignment Problem. We show that it is useful only for small or medium scale data centres and with some relaxations, such as an optimality tolerance gap and a limited number of directions explored in the search space. Building on this study, we also investigate a hybrid approach, feeding a metaheuristic with the results of CPLEX, and we show that the gains are important in terms of quality of the set of Pareto solutions (+126.9% against the metaheuristic alone and +17.8% against CPLEX alone) and number of solutions (8.9 times more than CPLEX), while the processing time increases only by 6% in comparison to CPLEX for execution times larger than 100 seconds.

Keywords—Multi-objective Optimisation; VM/Machine Reassignment; Mixed Integer Linear Programming; Hybrid-Metaheuristics.

I. INTRODUCTION

Background and Research Challenge: Optimisation of data centres, through reassignment of virtual machines (VMs), is largely seen as one of the biggest challenges in data centre management [1]: not only servers are underutilised [2] and the potential savings are important, but the problem has a lot of constraints and is difficult to solve [3]. Constraint programming (CP) and mixed¹ integer programming (MILP) are known to be inefficient for such large scale problems with a limited execution time [4], [5], and usually researchers focus on other optimisation techniques (e.g., local search [6], [7] or greedy algorithms [8]) or mix CP or MILP with some other optimisation solutions (e.g., local search [5]).

In this paper, we address a slightly different and more relevant problem for the industry: *the Multi-objective VM Reassignment Problem*. ‘Optimising a data centre’ seems to suggest that there is something like a *best* reassignment of VMs, but in an enterprise there is no best placement: it is all about which objectives are favoured (by whom and when). It

¹Note that while the constraints can be expressed by binary variables, the objectives (especially the reliability cost, see Section II) require real variables.




Figure 1: Possible reassignments in 2-dimensions (in red, the good ones) and an organisation’s policies (in grey).

is not hard to imagine that different capital allocators (CA, the managers of data centres) may have different perspectives on the best way of making the system better. For instance some CAs may consider that energy consumption is the most important element, while for others it can be the cost of licensing; or some CAs see the reliability as the key element (for instance if they run critical applications), while other CAs have a strict policy regarding response time and then collocation of VMs. These preferences, or policies, may evolve or be in competition: for instance if CAs have two policies for their data centre (e.g., “electricity has to decrease by $x\%$ ” and “management cost has to be limited to $y\%$ ”) and some reassignment solutions can serve both, which one to favour? In this context, CP and MILP do not seem promising approaches, as the search space is large and the constraints hard and complex. As far as we know, the only related work tackles the problem using some other optimisation techniques. One of the challenges here is that execution time is limited: even if the reassignment is done on a monthly or a quarterly basis, as it often happens, the decision process is complex and CAs cannot wait more than a few hours or days: they verify and modify the solutions to suit their needs before making any decision. In this paper, as it is commonly accepted by practitioners and in the literature [5], [9] we use a time limit for the Multi-objective Reassignment Problem.

However, given that CPLEX’s solutions are generally of better quality than the ones of other optimisation techniques, and that there are several relaxation mechanisms in CPLEX and the multi-objective problem itself, we study in this paper whether a MILP solver, such as CPLEX, can be used for the Multi-objective Machine Reassignment Problem.

Example: As a motivating example, consider Figure 1 that shows a reassignment problem with two objectives: the electricity consumption and the reliability (the lower the better for both). Points a to f are all the possible reassignments, satisfying all the constraints of the data centre. The good reassignments, i.e., those that are better than any other in a particular combination of objectives, are in red. Solution e, for instance, is not a good reassignment as it is worse than b on both objectives - same applies to f wrt. c. The decision makers may define some policies, i.e., some rules that describe the desired optimisations: here a maximal value for the electricity consumption and the reliability, defining the area of good assignments that the decision makers can select (grey area on Figure 1). There are still two possible reassignments, b and c. The decision makers can then evaluate them locally and make a relative decision among them, such as “c gives me a better electricity cost but it is not a huge gain, while b’s reliability is higher – so we favour b”. Any technique addressing this problem needs to explore the research space in all directions in order to find varied and good solutions – which is expensive given the size of the problem.

Contributions: In this paper we make the following contributions: (i) we give a thorough study of how suited is a MILP solver (CPLEX) for the Multi-objective VM Reassignment Problem (the problem definition is given in Section II). We show that it is useful only for rather small or medium scale data centres and with some relaxations: a certain tolerance gap and a limited number of directions explored in the search space (Section IV); (ii) we propose an algorithm based on the combination of a MILP solver (CPLEX) and a hybrid-metaheuristic to improve the performance of the solver while keeping the execution time acceptable. This is detailed in Section V.

Hybridization is not a novel idea as such (e.g., see [10]), however combining a solver (which aim is to find ‘proper’ solutions and not only to relax the problem) and a meta-heuristic to benefit from both worlds in a multi-objective context is novel to our best knowledge. CP and MILP provide better solutions (when they get any, as they are expensive/slow) so our goal here is to help CPLEX as much as we can (i.e., increasing the optimality gap and varying the number of vectors) to get some solutions to feed in the meta-heuristic to cover the search space. Note that in [5] Mehta et al. use CP on a relaxed problem, not the original VM placement one, which leads us to think that CP is inefficient for our problem – and anyway we also tried ILOG CP Optimizer and noticed extremely poor performances.

II. PROBLEM DEFINITION

In this section, we give a Multi-objective extension of the VM Reassignment Problem originally proposed by Mehta et al. in [5]. We first describe briefly the different elements of data centres, followed by the constraints of the problem, and we finish with the different objectives that we believe are the most relevant (note that this approach is agnostic to objectives and would work with any other linear objective function).

A. Problem Description and Notation

A data centre is composed of a set \mathcal{M} of physical machines (PMs or servers). Each machine $m_i \in \mathcal{M}$ has a finite amount

$Q_{m_i,r}$ of resource $r \in \mathcal{R}$ (e.g., CPU, RAM, storage). Machines $m_i \in \mathcal{M}$ belonging to a same rack are linked with high network connections, and thus considered as being in the same neighbourhood $N(m_i) \in \mathcal{N}$. Resources are of two different types: (i) transient resources ($r \in \mathcal{TR} \subseteq \mathcal{R}$, e.g., RAM or storage) that are consumed at the original host and also at the destination one during a migration process, or (ii) non-transient ($r \in \overline{\mathcal{TR}}$). The data centre is in charge of a set of VMs $v \in \mathcal{V}$ with resource requirements $d_{v,r}$ for every $r \in \mathcal{R}$. $M_0(v)$ and $M(v)$ respectively indicate the initial and the final host of the VM v during the reassignment. VMs composing a same multi-tier application are usually gathered by services $\mathcal{S} = \{s_1, \dots, s_p\}$, with $s_p = \{v_p^1, \dots, v_p^q\}$.

B. Constraints

We present here the linear constraints of our problem. We follow the problem’s linear constraints provided in [5].

1) *Reassignment Constraints:* Consider a binary variable $x_{v,m}$ for every VM $v \in \mathcal{V}$ and for each machine $m \in \mathcal{M}$, which is set to 1 if $M(v) = m$ and 0 otherwise. Constraints (1) ensure that every VM is reassigned to one and only one machine:

$$\forall v \in \mathcal{V}, \sum_{m \in \mathcal{M}} x_{v,m} = 1 \quad (1)$$

2) *Capacity Constraints:* There are two ways of computing resource utilisation $U_{m,r}$ of a machine $m \in \mathcal{M}$ and a resource $r \in \mathcal{R}$: (2) for non-transient resources and (3) for transient resources.

$$\forall r \in \overline{\mathcal{TR}}, \forall m \in \mathcal{M}, U_{m,r} = \sum_{v \in \mathcal{V}} d_{v,r} \cdot x_{v,m} \quad (2)$$

$$\forall r \in \mathcal{TR}, \forall m \in \mathcal{M}, U_{m,r} = \sum_{v \in \mathcal{V} | M_0(v)=m} d_{v,r} \cdot x_{v,m} + \sum_{v \in \mathcal{V} | M_0(v) \neq m} d_{v,r} \cdot x_{v,m} \quad (3)$$

The total resource utilisation of a machine $m \in \mathcal{M}$ should not exceed its capacity $Q_{m,r}$ for every $r \in \mathcal{R}$:

$$\forall r \in \mathcal{R}, \forall m \in \mathcal{M}, U_{m,r} \leq Q_{m,r} \quad (4)$$

3) *Conflict Constraints:* VMs which belong to a same service have to be reassigned to different machines (e.g., for replication purposes).

$$\forall s \in \mathcal{S}, \forall m \in \mathcal{M}, \sum_{v \in \mathcal{V} | M_0(v)=m} x_{v,m} \leq 1 \quad (5)$$

4) *Dependency Constraints:* Services sometimes depend on each other. VMs of these services need to be close to each other in order to achieve high performance (e.g., in the case of multi-tier applications). Let \mathcal{D} be the set of service dependencies such that $\mathcal{D} = \{(s_i, s_j) | s_i, s_j \in \mathcal{S} \text{ and } s_i \text{ depends on } s_j\}$, then:

$$\forall s_i, s_j \in \mathcal{S}, (s_i, s_j) \in \mathcal{D} \implies \forall v_a \in s_i, \exists v_b \in s_j | N(M(v_a)) = N(M(v_b)) \quad (6)$$

To give a linear definition of constraints (6), we introduce the binary variables $y_{s,n}$ for every service $s \in \mathcal{S}$ and for each neighbourhood $n \in \mathcal{N}$. Constraints (7) and (8) ensure that each variable $y_{s,n}$ is set to 1 if at least one VM from the service s is hosted by a machine in the neighbourhood $n \in \mathcal{N}$ and to 0 otherwise:

$$\forall s \in \mathcal{S}, \forall n \in \mathcal{N} \quad \sum_{v \in \mathcal{S}} \sum_{m \in \mathcal{N}} x_{v,m} \leq |\mathcal{N}| \cdot |\mathcal{S}| \cdot y_{s,n} \quad (7)$$

$$\forall s \in \mathcal{S}, \forall n \in \mathcal{N} \quad \sum_{v \in \mathcal{S}} \sum_{m \in \mathcal{N}} x_{v,m} \geq y_{s,n} \quad (8)$$

If a service s_i depends on s_j , constraints (9) guarantee that there is not any VM from s_i reassigned to a machine in a neighbourhood $n \in \mathcal{N}$ without having at least one VM from s_j reassigned to that neighbourhood:

$$\forall (s_i, s_j) \in \mathcal{D}, \forall n \in \mathcal{N} \quad y_{s_i,n} \leq y_{s_j,n} \quad (9)$$

5) *Spread Constraints*: Data centres are often decentralised and we assume here that they are composed of different sites \mathcal{L} , each machine belonging to a location $l \in \mathcal{L}$. For reliability and security reasons, some services need to be spread over at least *spreadNumber* locations. Let us introduce a binary variable $z_{s,l}$ for every service $s \in \mathcal{S}$ and for each location $l \in \mathcal{L}$. Constraints (10) and (11) ensure that $z_{s,l}$ gets the value 1 if the service s has at least one VM running in a machine at the location l :

$$\forall s \in \mathcal{S}, \forall l \in \mathcal{L} \quad \sum_{v \in \mathcal{S}} \sum_{m \in \mathcal{L}} x_{v,m} \leq |\mathcal{N}| \cdot |\mathcal{S}| \cdot z_{s,l} \quad (10)$$

$$\forall s \in \mathcal{S}, \forall l \in \mathcal{L} \quad \sum_{v \in \mathcal{S}} \sum_{m \in \mathcal{L}} x_{v,m} \geq z_{s,l} \quad (11)$$

Constraints (12) force every service s to run on at least *spreadNumber_s* number of locations:

$$\forall s \in \mathcal{S}, \quad \sum_{l \in \mathcal{L}} z_{s,l} \geq \text{spreadNumber}_s \quad (12)$$

Definition 1 (Machine Reassignment): An assignment A of VMs to machines is a mapping: $A : \mathcal{P} \mapsto \mathcal{M}$, such that $A(v, \mathcal{M}) \rightarrow m$, which satisfies the constraints (1–5 and 7–12). A reassignment is a function: $ReA : A \mapsto A$ which returns a new assignment for a given initial assignment of VMs to machines.

C. Objectives

We focus here on three objectives: electricity cost, VM migration cost and reliability cost, as they are recognised in the literature [11], [12], [13] and make sense in practice. The multi-objective variant of the Machine Reassignment (see Definition 1) consists of minimising the cost functions defined by the objectives, the ones we present here or any others that would be relevant.

1) *Reliability Cost*: For each machine $m \in \mathcal{M}$ and each resource $r \in \mathcal{R}$, we define the safety capacity $SC_{m,r}$ as the amount of resource that it is ‘safe’ to allocate without overloading m – this is similar to the resource buffer that placement algorithms often assume [14]. The risk of failure $R_{m,r}$ is then the difference between the actual utilisation of resource r and the safety capacity, and the reliability cost R_{cost} represents the ‘non reliability’ over the full data centre.

$$\forall m \in \mathcal{M}, \forall r \in \mathcal{R}, \quad R_{m,r} \geq U_{m,r} - SC_{m,r} \geq 0 \quad (13)$$

$$R_{cost} = \sum_{r \in \mathcal{R}} \sum_{m \in \mathcal{M}} R_{m,r} \quad (14)$$

2) *Electricity Cost*: For each machine $m \in \mathcal{M}$, we introduce a binary variable o_m to be set by constraints (15) to 1 if the machine m is switched on (i.e., hosts at least one VM) and 0 otherwise.

$$\forall m \in \mathcal{M}, \quad o_m \leq \sum_{v \in \mathcal{V}} x_{v,m} \leq |\mathcal{V}| \cdot o_m \quad (15)$$

The electricity cost E_{cost} is composed of the electricity consumption of each machine $m \in \mathcal{M}$ multiplied by its price γ_m . The electricity consumption of a machine m is often modelled as a linear function of its CPU utilisation [15], [16], with α_m being its electricity consumption at idle and β_m its consumption per unit of CPU usage.

$$E_{cost} = \sum_{m \in \mathcal{M}} \gamma_m (\alpha_m \cdot o_m + \beta_m \cdot U_{m,CPU}) \quad (16)$$

3) *Migration Cost*: For each VM $v \in \mathcal{VM}$, we introduce a binary variable mig_v to be set by constraints (17) to 1 if v is reassigned and 0 otherwise.

$$\forall v \in \mathcal{V}, \quad \sum_{m \in \mathcal{M} \setminus \mathcal{M}_0(v)} x_{v,m} = mig_v \quad (17)$$

The migration cost M_{cost} concerns all migrating VMs. For each migrating VM $v \in \mathcal{V}$, the migration cost is the time it takes to: (i) prepare the VM $\mu_1(v)$, (ii) to transfer its image from its initial placement to its new host m $\mu_2(v, \mathcal{M}_0(v), m)$ and (iii) to deploy it in the new host $\mu_3(v)$:

$$M_{cost} = \sum_{v \in \mathcal{V}} \left([\mu_1(v) + \mu_3(v)] \cdot mig_v + \sum_{m \in \mathcal{M}} \mu_2(v, \mathcal{M}_0(v), m) \cdot x_{v,m} \right) \quad (18)$$

III. EXPERIMENTAL SETUP

In the current section we present our evaluation setup, i.e., the description of several instances of the problem (data centres to optimise) inspired from a well-known data set [9] and the metrics used to judge the proposed systems. All the algorithms described below have been developed in C++. Experiments are done on a computing cluster with 24 cores 2.0GHz Intel Ivy Bridge CPU and 128GB of RAM.

A. Data set

Google and the ROADEF society (i.e., the French OR society) released a few years ago a data set, now widely used in the OR community, for the evaluation of VM reassignment solutions [9]. This data set represents various data centres, of different sizes and characteristics (e.g., various number of resources), with a large number of constraints. This data set does not provide a multi-objective formulation though and we had to adapt the instances (note that the participants of the challenge optimised only one single weighted sum of the costs proposed – hence there is no possible comparison between our work and others using ROADEF). Our instances aim to model realistic scenarios as we observe them in large companies. For our evaluation, we take the 14 first instances, leaving out only the largest ones (see Table I). Two of the objectives we define are present in ROADEF as cost functions: safety/reliability and migration, and we add electricity. We randomly generate electricity consumption constants (α , β) for every machine $m \in \mathcal{M}$ and also the electricity cost γ for every location $l \in \mathcal{L}$. The data set also comes with a time limit representing the maximum time allowed for the resolution of the instance (300 seconds). We changed the time limit to: 30s for a_1_1 , 1h for a_1_2-5 , 2h for a_2_1-3 , and 10h for the other instances, which is considered realistic by large companies in the context of optimisation performed in a regular basis, e.g., monthly or quarterly.

Instance	# Resources	# Machines	# Services	# Processes
a_1_1	2	4	79	100
a_1_2	4	100	980	1,000
a_1_3	3	100	216	1,000
a_1_4	3	50	142	1,000
a_1_5	4	12	981	1,000
a_2_1	3	100	1,000	1,000
a_2_2	12	100	170	1,000
a_2_3	12	100	129	1,000
a_2_4	12	50	180	1,000
a_2_5	12	50	153	1,000
b_1	12	100	2,512	5,000
b_2	12	100	2,462	5,000
b_3	6	100	15,025	20,000
b_4	6	500	1,732	20,000

Table I: Characteristics of the different instances used in our evaluation.

B. Metrics

The comparison of algorithms optimising a multi-objective problem is complex as their results can be evaluated from different perspectives, such as: the number of solutions on the Pareto front, the variety of solutions [17], or the spread of these solutions. The comparison is even more complex in our case as the problem is large, the Pareto front is not known in most cases and objectives cannot be considered separately. In this paper, we only consider unary metrics, i.e., returning a single value based on solutions found by every algorithm, making a comparison of several algorithms easier.

1) *Number of solutions*: We use the number of *non-dominated (efficient) solutions* in our experiments as our first metric. We refer to it as the quantity of found solutions. This metric is highly important for data centre capital allocators as it gives them more choices. It also provides them with backup solutions if the preferred one appears difficult or technically impossible to implement.




Figure 2: Metrics: number of Pareto solutions (red dots) and hypervolume (grey area).

2) *Hypervolume*: The *hypervolume* [18] (a.k.a., S metric) refers to the quality of a set of solutions. The hypervolume is used [19] in the multi-objective optimisation community for comparing different sets of non-dominated solutions. For every set of solutions, this metric measures the space between the efficient solutions and a reference point far from them. The reference point is defined as a point in the space having the worst objective values and must be identical for all the algorithms, but may be different for every instance. Figure 2 shows a two-dimensional reassignment, with the non-dominated solutions in red and the other (not interesting) solutions in light blue. The initial placement is one of the solutions, generally not on the Pareto front. The hypervolume is the grey area in this 2D space.

IV. CPLEX FOR THE VM REASSIGNMENT PROBLEM

The goal of this section is to study both performance and scalability of a MILP solver (i.e., CPLEX) on the Multi-objective VM Reassignment Problem. First, we explore how CPLEX performs on the (Mono-objective) VM Reassignment Problem, i.e., the original problem from ROADEF. We show how difficult the problem is for CPLEX, even in this simpler version. Then we explore the performance of CPLEX for the Multi-objective VM Reassignment and show that it can be tackled under some restrictions, such as, a limited number of directions (*vectors*) explored in the search space and an optimality tolerance gap.

A. CPLEX for the Mono-objective VM Reassignment Problem

Portal et al. [20] show that the VM Reassignment Problem is too difficult for a MILP solver like CPLEX when applied on the ROADEF instances. CPLEX could only solve 3 instances out of the 14 within the time limit (300s) fixed during the challenge. However CPLEX allows defining an *optimality gap tolerance*, a trigger that stops CPLEX when the current feasible solution falls within a certain percent of CPLEX best estimation of the lower bound (a value smaller or equal to the actual optimal value). For instance, a 5% gap means that any solution that is 5% away from the estimated lower bound is accepted and stops CPLEX. In addition, we have here a larger time budget to solve the instances.

Table II shows the execution time (in seconds) of CPLEX for the resolution of one single vector (the identity vector, i.e., with weights equal 1 for the three objectives). As a reminder, we also add the time limit (last column) that we set for each instance. First, we notice that CPLEX only solves instances

	50%	20%	10%	5%	1%	0.5%	0.1%	0.05%	0.01%	Limit (s)
a_1_1	0.08	0.08	0.08	0.08	0.08	0.12	0.14	0.23	0.25	30
a_1_2	186	183	185	185	1,490	6,321	–	–	–	3,600
a_1_3	27	27	27	37	625	1,691	–	–	–	3,600
a_1_4	50	50	51	51	98	1,682	–	–	–	3,600
a_1_5	9	9	9	9	9	9	10	19	26	3,600
a_2_1	54	55	159	3,670	–	–	–	–	–	7,200
a_2_2	2,511	2,580	2,580	2,736	–	–	–	–	–	7,200
a_2_3	71	71	71	71	5,816	–	–	–	–	7,200
a_2_4	20,445	20,502	20,655	–	–	–	–	–	–	36,000
a_2_5	21,877	22,492	22,513	–	–	–	–	–	–	36,000
b_1	3,482	6,913	6,913	7,094	–	–	–	–	–	36,000
b_{2,3,4}	–	–	–	–	–	–	–	–	–	36,000

Table II: Execution time (s) of CPLEX for the resolution of the identity vector (all objectives have weights equal 1) depending on the optimality tolerance gap. The symbol ‘–’ means that no solution was found in the time limit.

a_{1_1} and a_{1_5} in the time limit (note that 0.01% is the default tolerance gap for CPLEX). This supports the general claim that a MILP solver is inefficient for our problem, even in this simple case with only one vector. We then notice that CPLEX gets a solution with a gap of 0.5% for all small instances, 10% for all medium instances (1% or 5% for some), and solved only one of the hard instances (b_1 , with a 5% gap), even with a 50% gap. We also observe that often CPLEX finds a first good solution (e.g., a_{1_2} , a_{1_3} and a_{1_4} have a solution for 5% quickly, as evidenced by the same time for 50%, 20%, 10% and 5%) but it is then difficult for CPLEX to improve it. As a conclusion, CPLEX does not seem able to scale to large instances but with a proper gap CPLEX finds good solutions.

B. CPLEX for the Multi-objective VM Reassignment Problem

Once we know that CPLEX finds it difficult to solve one weight vector in the search space, we would like to explore what needs to be relaxed in order to help CPLEX optimise more vectors and hence address a proper multi-objective problem. In the current section we look at three elements: (i) given that optimising one weight vector is already difficult to CPLEX and the more vectors we optimise the more we explore the search space, what is the most reasonable number of weight vectors to be optimised? (ii) getting a tiny optimality gap increases exponentially the execution time, therefore what is the best value for this parameter? (iii) what is the best way to use CPLEX? Would it be better to collect all intermediate (feasible) solutions found by CPLEX instead of only keeping the best ones for each weight vector?

Figures 3 show the execution time CPLEX needs to reach different optimality gaps given several maximally spread weight vectors. These vectors are built on the assumption that to explore a maximum of the space the solver needs to target widely spread directions. In our 3-dimensional space (3 objectives), we successively use these vectors: (1,1,1), (0.6, 0.3, 0.1), (0.3, 0.1, 0.6), (0.1, 0.6, 0.3), (0.45, 0.45, 0.1), (0.45, 0.1, 0.45) and (0.1, 0.45, 0.45). We notice that running different vectors increases the final execution time – as it can be expected. It also confirms what we have seen earlier in Table II that for large optimality gaps, there is no large difference in terms of execution time (CPLEX quickly finds good solutions), however the more we decrease the gap, the more important is the increase in execution time. We also see that due to the time budget limitation, we cannot run all the possible optimisations with the different vectors for some optimality gaps (e.g., for a_{1_2} we could only use 2 vectors out of the 7 possible ones with an optimality gap of 1%). Therefore, a decision has to

be made on which values should be set for both variables: the optimality gap and the number of vectors.

According to Figures 3, two patterns emerge: (i) on small instances: asking CPLEX for an optimality gap smaller than 5% increases its execution time drastically, and (ii) on medium instances: this gap drops to 5 – 10%. Therefore, to keep our optimisation in a reasonable execution time, the larger/more complex is the instance the larger the optimality gap we consider. Unlike what we might think, CPLEX execution time is very heterogeneous and varies a lot from a vector to another (execution time curves are not linear). Thus, knowing CPLEX execution time on the first vector does not give any indication/prediction on the execution time for other vectors. Because of the lack of knowledge of the execution time, we have to restrict the number of vectors as much as possible. We are even more constrained regarding some instances such as a_{2_4} and a_{2_5} where we could only run CPLEX on one vector.

Another element of the resolution process that might be interesting to look at is the intermediate (feasible) solutions generated by CPLEX. The way CPLEX works is iterative: it first finds a feasible solution that is either discarded/improved if it is not optimal or kept if it is optimal. In our multi-objective context, those intermediate solutions, while not optimal in a particular combination of objectives (remember that CPLEX solves vectors of weights for the objectives), may sometimes carry some interesting reassignments of the VMs, for instance wrt. some single objective – and hence improve the hypervolume. Figures 4 show the hypervolume obtained on the instances a_{1_1} , a_{1_2} and a_{1_3} by CPLEX when increasing the number of vectors for different optimality gaps going from 50% to 0.01%. The three figures at the left side show the hypervolume obtained when only the best solutions are kept in the solution set while the three figures at the right side give the hypervolume for the same experiments with all solutions in the solution set. We notice from the graphs at the left that optimising the objectives using a vector of weights with a small gap does not imply getting a better hypervolume for the multiple objectives (e.g., on the instance a_{1_3} , two vectors and a gap of 5% get a better hypervolume than using a gap of 1%). This is mainly due to the fact that optimising a compromise of objectives using their linear combination may lead to optimising one objective at the expense of the others. This is also caused by the fact that the objectives are in different units and of different scales (e.g., the electricity cost has a larger scale than the migration cost). Collecting all the feasible solutions during the optimisation of every vector may then be a good improvement: see the three graphs at the right

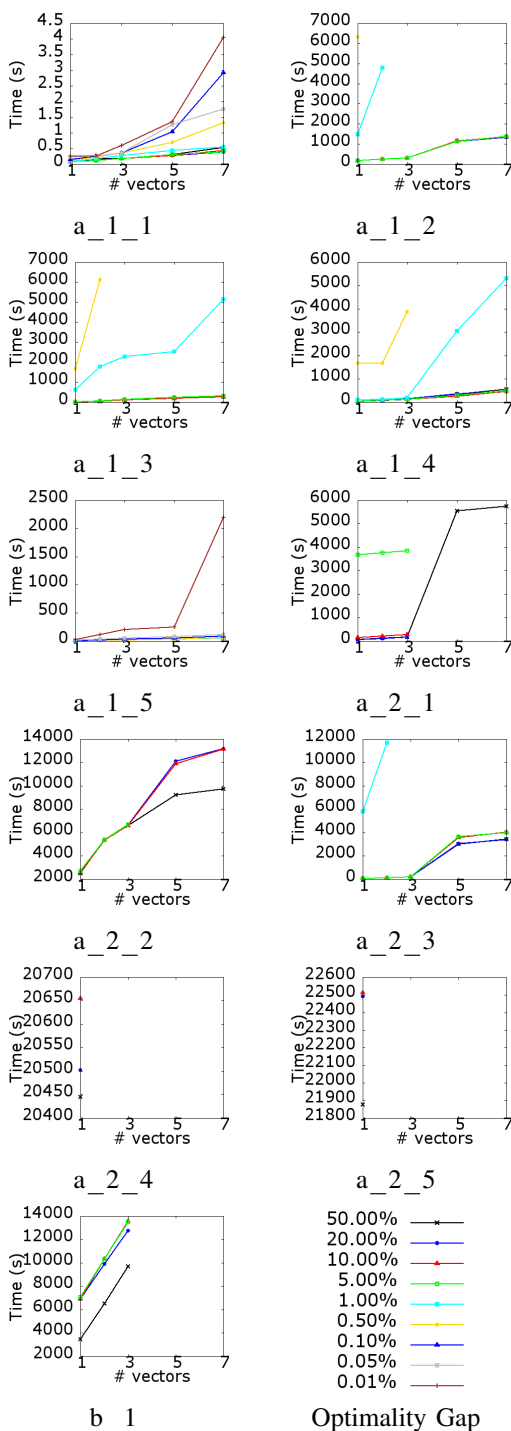


Figure 3: Execution time (s) of CPLEX on ROADEF instances, with maximally spread objective weight vectors and using different optimality gaps.

of Figure 4. We notice that using a small optimality tolerance gap gives us better results than using a larger one. We also notice that we get an improvement in terms of hypervolume. This is an interesting behaviour especially since we did not add any noticeable extra computation time (CPLEX already collects the intermediary solutions, and filtering/removing the dominated solutions requires a negligible execution time). In the rest of our evaluations we collect all intermediate (good, i.e., non-dominated) solutions.




Figure 4: Hypervolume obtained on instances $a_{1_1_1}$, $a_{1_1_2}$ and $a_{1_1_3}$ using CPLEX for different numbers of weight vectors with different optimality gaps by either only considering the best solution found or by collecting all the feasible solutions during the optimisation.

Figure 5 shows the hypervolumes obtained with different number of vectors and different tolerance gaps, and this can be read together with Figure 3 to figure what composition of number of vectors and tolerance gap gives the best time-hypervolume trade-off. We notice that globally the hypervolume increases with every new vector. This improvement is more noticeable during the 5 first vectors. It seems to stagnate between the 5th and 7th (whenever CPLEX reaches them in the time constraint), especially when the optimality gap is tiny. This leads us to think that running several CPLEX optimisations with a large number of vectors (larger than 5) would not be as beneficial as we might think. It would increase the execution time without improving the hypervolume. This idea can obviously be withdrawn if the managers of the data centre are ready to spend more time to achieve better hypervolume results.

We see from Figures 3 and 5 that using an optimality gap of 5% and 3 weight vectors for small instances and an optimality gap of 10% and 1 weight vector for medium instances allow us to achieve a good hypervolume while keeping the execution time reasonable for all the instances.

V. COMBINING A SOLVER AND A METAHEURISTIC

We saw earlier that CPLEX gets good results on small and medium instances, but, it becomes really hard to improve those results without increasing either the number of vectors or the optimality gap, and thus dramatically impacting the execution time. In this section, we want to see whether combining a solver with a metaheuristic allows us to get better results in a reasonable time. A comparison of different metaheuristics has already been performed [21] and a scalable

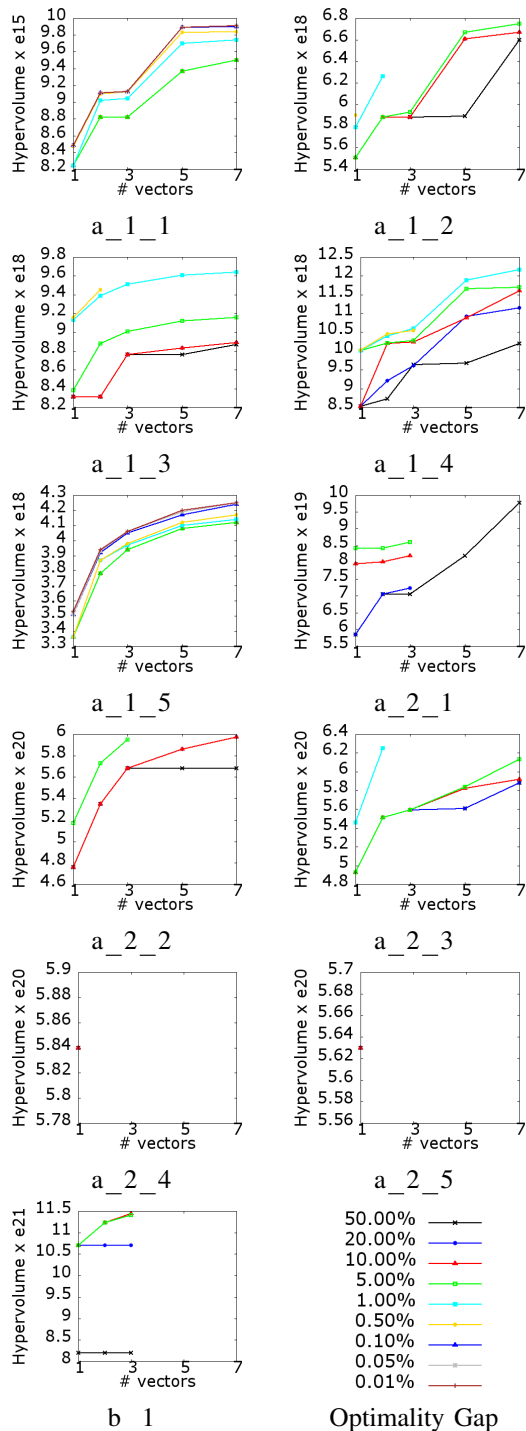


Figure 5: Hypervolume obtained with CPLEX (all feasible solutions) on ROADEF instances, using maximally spread weight vectors and different optimality gaps.

hybrid-metaheuristic (GeNePi) proposed to optimise the Multi-objective VM Reassignment Problem. GeNePi outperforms state-of-the-art algorithms on both quantity and quality of solutions. GeNePi successively applies three metaheuristics: (i) a greedy algorithm (i.e., GRASP [8]) to quickly find good solutions representing the research space, (ii) a genetic algorithm (i.e., NSGA-II [22]) by combining the previous solutions to get better solutions optimising different objective trade-offs, and (iii) a local search (i.e., PLS [23]) to refine the

Pareto front and find more non-dominated solutions.

We take the previous implementation of CPLEX and we give its results to GeNePi, non-dominated solutions found using CPLEX being the initial population. It might happen that CPLEX does not find enough solutions to fill out an entire population (in our case, a population of size 20). In this case the original greedy algorithm in GeNePi is applied to fill this gap and compensate this lack of solutions. In our implementation, GeNePi applies 10 iterations of its second phase (i.e., NSGA-II) in order to evolve the initial population into a fitter one, by getting better and more scattered solutions (spread over the research space). At the end, GeNePi refines the Pareto front by applying a unique iteration of Pareto Local Search (PLS) on the 10 most isolated solutions. Although this step does not bring a large improvement in terms of hypervolume, it is important as it provides decision-makers with more implementation choices. Beside these choices, we use the same parameters as in the GeNePi paper [21].

Table III shows the results obtained on the modified ROADEF instances from a_1_1 to b_1 , in terms of hypervolume, number of non-dominated solutions and execution time. Results are obtained using 10 runs of GeNePi alone (the average is taken), CPLEX alone, and our new metaheuristic (i.e., CPLEX combined with GeNePi) while respectively defining the optimality gap and the number of weight vectors to 5% and 3 for small instances (i.e., a_1_x), and 10% and 1 for medium and large ones (i.e., a_2_x and b_1). It also shows the execution time of each algorithm for every instance. Table III confirms that GeNePi succeeds in improving the hypervolume and getting a large number of non-dominated solutions while keeping the execution time relatively low. We also notice that CPLEX outperforms GeNePi in terms of hypervolume in 8 instances out of 11 with an average improvement of 102%. However, GeNePi gets on average 63 times more non-dominated solutions (note that the number of solutions while interesting for the decision makers, is not as important as the hypervolume, and anyway too many solutions makes them difficult to explore). We see that CPLEX is slightly better in terms of execution time (given its parallel implementation and the fact that it runs on more powerful memory intensive machines). However we see that CPLEX struggles to scale to large instances (GeNePi gets a better hypervolume on a_2_4 and a_2_5 with an execution time of respectively 342s and 347s vs. 21,173s and 22,548s for CPLEX). Compared to CPLEX, we clearly see that adding GeNePi to CPLEX helps to improve the hypervolume (an increase of 17.84% on average), and also to get more non-dominated solutions (8.9 times more solutions on average), while keeping the execution time low (an average execution time increase of 31.10%, but of only 6% for execution times larger than 100s). We also see that CPLEX + GeNePi outperforms the hypervolume obtained by GeNePi alone (with an average increase of 126.96%) and that unlike CPLEX alone, CPLEX + GeNePi gets better hypervolumes than GeNePi in all instances. CPLEX + GeNePi also gets a fairly reasonable number of non-dominated solutions. However, GeNePi still gets a larger number of solutions (5 times more on average).

To summarise, we can say that CPLEX is good at getting few solutions with a good quality but does not scale well to large instances, unlike GeNePi and other (hybrid-

Instance	Initial Hyp	GeNePi			CPLEX			CPLEX + GeNePi		
		Hyp	#Sol	Time	Hyp	#Sol	Time	Hyp	#Sol	Time
a_1_1	6.16E+15	9.71E+15	225	1.58	8.82E+15	3	0.18	9.74E+15	30	0.38
a_1_2	4.75E+18	5.69E+18	160	518	5.93E+18	8	303	6.17E+18	42	319
a_1_3	6.57E+18	7.22E+18	132	441	9.01E+18	5	151	9.06E+18	21	167
a_1_4	8.10E+18	9.19E+18	136	309	1.03E+19	6	138	1.09E+19	35	158
a_1_5	2.42E+18	3.15E+18	283	332	3.94E+18	4	21	4.04E+18	54	34
a_2_1	4.57E+19	5.89E+19	231	284	7.95E+19	5	159	8.17E+19	39	166
a_2_2	3.33E+20	4.36E+20	197	600	4.76E+20	7	2,580	5.01E+20	42	2,694
a_2_3	2.55E+20	3.69E+20	202	695	4.93E+20	2	71	5.13E+20	32	154
a_2_4	3.21E+20	6.19E+20	253	342	5.84E+20	3	20,655	6.68E+20	43	21,173
a_2_5	4.96E+20	5.81E+20	220	347	5.63E+20	2	22,513	5.86E+20	23	22,548
b_1	8.20E+21	8.74E+21	244	14,991	1.07E+22	2	8,913	1.09E+22	29	9,452

Table III: Summary of results obtained with GeNePi, CPLEX and CPLEX combined with GeNePi using a gap of 5% and 3 vectors for small instances, and a gap of 10% and a unique vector for medium and large instances.

metaheuristics. Also, knowing that CPLEX execution time on one vector is low does not provide much information on the time to run on several other vectors. Combining CPLEX with GeNePi seems to be a good solution to improve CPLEX's results in both hypervolume and number of non-dominated solutions with a relatively low increase in the execution time. Another advantage of our approach is that it can be adapted to the size of the problem: when the size increases, we have the option to either continue relaxing some of the parameters used for CPLEX (i.e., less vectors or a larger optimality gap) or to replace CPLEX with another 'bootstrapping' technique (e.g., greedy algorithm) to feed in the meta-heuristic.

VI. CONCLUSION

This paper assesses whether a MILP solver, such as CPLEX, can be applied to the Multi-objective VM Reassignment Problem, a large and difficult problem with a lot of complex constraints: given some reassignment objectives (e.g., electricity cost, migration cost, reliability cost), find the best set of reassignment solutions in a limited time – the limit being quite large (10 hours for large instances). We show that CPLEX can be used with some relaxations: allowing an optimality tolerance gap (which stops CPLEX when the solutions found are close to the optimal) and limiting the number of directions explored in the search space (giving CPLEX only few vectors of weighted objectives to explore). We also propose to combine the results of CPLEX to the execution of a metaheuristic (GeNePi) and we compare CPLEX alone, the metaheuristic alone and the combination of both. We observe that CPLEX is better than the metaheuristic (an improvement of the hypervolume of 102% in comparison to GeNePi) while the combination CPLEX+GeNePi outperforms both in terms of hypervolume (an average increase of 126.9% vs. GeNePi and 17.8% vs. CPLEX), while the execution time remains acceptable (an increase of only 6% on average in comparison to CPLEX for execution times larger than 100s). As future work, we would like to study the usage of algorithms combining metaheuristics and a mathematical resolution for very large instances as a scalable substitution for the MILP solver.

Acknowledgement: This work was supported, in part, by Science Foundation Ireland grants 10/CE/I1855 and 13/RC/2094, and by Science Foundation Ireland Industry Fellowship grant 13/IF/12789.

REFERENCES

- [1] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," *FGCS*, pp. 755–768, 2012.
- [2] R. W. Ahmad, A. Gani, S. H. A. Hamid, M. Shiraz, A. Yousafzai, and F. Xia, "A survey on virtual machine migration and server consolidation frameworks for cloud data centers," *JNCA*, pp. 11–25, 2015.
- [3] A. Corradi, M. Fanelli, and L. Foschini, "Vm consolidation: A real case based on openstack cloud," *FGCS*, pp. 118–127, 2014.
- [4] F. Hermenier, S. Demassej, and X. Lorca, "Bin repacking scheduling in virtualized datacenters," in *CP*, pp. 27–41, 2011.
- [5] D. Mehta, B. O'Sullivan, and H. Simonis, "Comparing solution methods for the machine reassignment problem," in *CP*, pp. 782–797, 2011.
- [6] H. Gavranović, M. Buljubašić, and E. Demirović, "Variable neighborhood search for google machine reassignment problem," *ENDM*, pp. 209–216, 2012.
- [7] F. Brandt, J. Speck, and M. Völker, "Constraint-based large neighborhood search for machine reassignment," *Ann Oper Res*, pp. 1–29, 2012.
- [8] M. Gabay and S. Zaourar, "A GRASP approach for the machine reassignment problem," in *EURO*, 2012.
- [9] "Google/roaDEF challenge'12." <http://challenge.roaDEF.org/2012/en/>.
- [10] P. Shaw, "Using constraint programming and local search methods to solve vehicle routing problems," in *CP*, pp. 417–431, 1998.
- [11] D. Filani, J. He, S. Gao, M. Rajappa, A. Kumar, P. Shah, and R. Nagappan, "Comparing vm-placement algorithms for on-demand clouds," in *Dynamic data center power management*, 2008.
- [12] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation," in *CloudCom*, pp. 254–265, 2009.
- [13] B. Schroeder and G. A. Gibson, "A large-scale study of failures in high-performance computing systems," *TDSC*, pp. 337–351, 2010.
- [14] X. Li, A. Ventresque, J. Omana Iglesias, and J. Murphy, "Scalable correlation-aware virtual machine consolidation using two-phase clustering," in *HPCS*, 2015.
- [15] J. Xu and J. A. Fortes, "Multi-objective virtual machine placement in virtualized data center environments," in *GreenCom*, pp. 179–188, 2010.
- [16] C.-H. Lien, Y.-W. Bai, and M.-B. Lin, "Estimation by software for the power consumption of streaming-media servers," *TIM*, pp. 1859–1870, 2007.
- [17] E. Zitzler, M. Laumanns, L. Thiele, C. M. Fonseca, and V. G. da Fonseca, "Why quality assessment of multiobjective optimizers is difficult," in *GECCO*, pp. 666–673, 2002.
- [18] E. Zitzler and L. Thiele, "Multiobjective optimization using evolutionary algorithms. a comparative case study," in *PPSN*, pp. 292–301, 1998.
- [19] K. Bringmann and T. Friedrich, "Approximation quality of the hypervolume indicator," *Artificial Intelligence*, pp. 265–290, 2013.
- [20] G. M. Portal, M. Ritt, L. M. Borba, and L. S. Buriol, "Simulated annealing for the machine reassignment problem," *AOR*, pp. 1–22, 2012.
- [21] T. Saber, A. Ventresque, X. Gandibleux, and L. Murphy, "Genepi: A multi-objective machine reassignment algorithm for data centres," in *HM*, pp. 115–129, 2014.
- [22] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *TEVC*, pp. 182–197, 2002.
- [23] A. Alsheddy and E. E. Tsang, "Guided pareto local search based frameworks for biobjective optimization," in *CEC*, 2010.