Multi-Objective VM Reassignment for the Enterprise

Takfarinas Saber[1], Anthony Ventresque[1], Liam Murphy[1], and El-Ghazali Talbi[2]

[1] Lero@UCD, School of Computer Science and Informatics,
University College Dublin, Ireland
`takfarinas.saber@ucdconnect.ie`, {`anthony.ventresque, liam.murphy`}`@ucd.ie`
[2] Laboratoire d'Informatique Fondamentale de Lille,
INRIA Lille Nord Europe, Université de Lille1
`El-Ghazali.Talbi@lifl.fr`

1 Introduction

Virtual Machine placement, that is assigning Virtual Machines (VMs) to Physical Machines (PMs) in a data centre is an important field of research as it is a difficult problem (NP-Hard) while better placements can generate large savings for companies. In particular, we have shown in a previous work [1] that decision makers can take advantage of a multi-objective presentation to make better reassignments that correspond to their preferences. In the context of the Enterprise, the problem is even more complex as the Entreprise is not monolithic but is a composition of several, somewhat autonomous, entities (we call them VCs for Virtual data Centres) through acquisitions and reorganisations. VM placement here needs to consider the *preferences* or *objectives* of the VCs in terms of usage of their computing resources [2]. For instance, some groups may want to keep some free resources in their own data centres (e.g., a testing group having the objective to patch quickly any defects found in production); while another group may want its VMs to have priority on a specific resource, wherever its VMs are hosted (e.g., a R&D department running CPU intensive processes). This is to some extent a reverse version of the broking problem in multi-clouds [3] with VCs and capital allocators instead of providers and brokers, but with two main differences: (i) it is a mix of cooperation and competition. (ii) exact costs of placements are not known in advance but only after VCs try to place the VMs that they have been given.

We present here[3] a preliminary study of the multi-objective VM reassignment problem for the Entreprise. We want to give decision makers a large panel of good and accurate solutions covering the different objectives that make sense for them: in this study we consider electricity consumption, VM migration and reliability costs; while letting each entity evaluate and modify possible reassignments according to its own preferences. We propose E-GeNePi, an adaptation of GeNePi [1] to the Enterprise context. The preliminary results show that E-GeNePi finds in average +114% solutions and gets the best hypervolume for 11 experiments out of 12.

2 Problem Definition

An Enterprise's data centre is composed of a set $\mathcal{C}$ of VCs: $\mathcal{C} = \{c_1, \dots c_n\}$. Each c_i is given a set $\mathcal{M}_i \in \mathcal{M}$ of PMs, $\mathcal{M}_i = \{m_1, \dots, m_n\}$. Each m_j has several resources $r \in \mathcal{R}$ (e.g., RAM, CPU, disk), in limited capacities $Q_{m_j,r}$. The Enterprise manages a set $\mathcal{V}$ of VMs, $\mathcal{V} = \{v_i, \dots, v_l\}$ hosted on PMs. The quantity of resource r that every v_k needs is fixed to $d_{k,r}$. Several constraints apply here and readers interested can refer to the abundant literature for more details [1, 4]. An important element to keep in mind though is that some constraints are defined at a global level while others apply only to a particular VC (e.g., specific capacity limitation for PMs in a VC).

We define an assignment A of VMs to PMs as a mapping: $A : \mathcal{V} \mapsto \mathcal{M}$, such that $A(v, \mathcal{M}) \to m$, which satisfies all the constraints of the model.A reassignment is then a function that modifies an initial assignment: $ReA : A \mapsto A$ and gives a new assignment of processes to machines.

In the context of this paper we are considering that every VC has a particular objective, defined as a combination of three simple objectives: electricity consumption, VM migration and reliability costs. At the decision making level of the Enterprise, the goal is to find a set of non-dominated solutions (reassignment of VMs), respecting the constraints and minimising the three objectives.

---

[3] **Acknowledgement**: This work was supported, in part, by Science Foundation Ireland grant 10/CE/I1855 to Lero - the Irish Software Engineering Research Centre.

3 E-GeNePi: GeNePi for the Enterprise

E-GeNePi is composed of two main optimisation functions (see Figure 1): one for decision makers and one for VCs' capital allocators. Decision makers run GeNePi, an hybrid algorithm for VM reassignment in data-centres [1]. GeNePi runs successively three steps with the objective of (i) exploring quickly the search space (this first step is inspired from GRASP [5]), (ii) introducing some variety and quality in the solutions found (using NSGA-II [6]) and (iii) increasing the number of solutions (through a PLS-based heuristic [7]). GeNePi has proven to be more effective than other classical solutions for multi-objective machine reassignment.

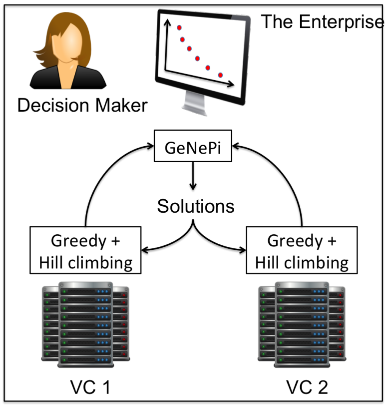

Fig. 1. E-GeNePi

Every time GeNePi finds a non-dominated solution, it forwards it on to every VC. VCs know their own resources better than the capital allocators of the Enterprise, and they can determine if there is a feasible assignment of the VMs they are given (and what is the best placement). To check the feasibility of the solution, VCs use a greedy algorithm and a reparation (if needed). If no feasible assignment is found, GeNePi is informed and this solution is discarded - whatever the other VCs have found. Whenever they find a feasible solution, VCs run a hill climbing algorithm to improve it, until a time-out is reached or no further improvement is found. If all VCs manage to reach this point, values given by the VCs are sent back to GeNePi which updates the values of the assignement w.r.t. the objectives.

4 Evaluation and Discussion

We evaluate our method using 12 instances of the ROADEF Challenge 2012 [4], modifying them as follows: (i) each *location* from the original problem definition becomes a VC – note that our approach does not require VCs to be contained in only one location though; (ii) we add electricity consumption (see [1]); (iii) we remove some of the constraints (spread, dependability) to make the problem more tractable; (iv) we give each VC a triplet of weights for the three objectives, representing the VC's preference. We observe that E-GeNePi performs better than the other solutions (namely, E-GRASP, E-NSGA-II and E-PLS when we use the corresponding heuristics instead of GeNePi in E-GeNePi): in average +114% solutions found, and the best hypervolume for 11 instances out of 12. Let's take instance a_2_2 as an example: this instance has 25 VCs, a total of 100 PMs and 1,000 VMs, each of them defined on 12 resources. E-GeNePi finds 53 solutions with an hypervolume of 3.704×10^{20}, while E-NSGA-II finds 41 solutions (hypervolume=3.506×10^{20}), and E-GRASP and E-PLS get 30 solutions (hypervolume=3.266×10^{20}).

We gave an initial definition of the Multi-Objective VM Reassignment problem for the Enterprise. In the future we plan to enable the exchange of information between VCs to decrease the number of non-feasible solutions and improve the quality of solutions. E-GeNePi shows encouraging preliminary results and we expect to improve them through a study of the interactions between the three steps composing GeNePi.

References

1. Saber, T., Ventresque, A., Gandibleux, X., Murphy, L.: Genepi: a multi-objective machine reassignment algorithm for data centres. In: Hybrid Metaheuristics. (2014)
2. Li, X., Ventresque, A., Stokes, N., Thorburn, J., Murphy, J.: ivmp: An interactive vm placement algorithm for agile capital allocation. In: IEEE CLOUD. (2013) 950–951
3. Kessaci, Y., Melab, N., Talbi, E.G.: A Pareto-based genetic algorithm for optimized assignment of virtual machines requests in a multi-cloud environment. In: CEC. (2013)
4. ROADEF/EURO: Challenge 2012, machine reassignment. (http://challenge.roadef.org/2012/en/)
5. Feo, T.A., Resende, M.G.: Greedy randomized adaptive search procedures. JGO (1995) 109–133
6. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. TEVC (2002) 182–197
7. Alsheddy, A., Tsang, E.E.: Guided pareto local search based frameworks for biobjective optimization. In: CEC. (2010)