Research Repository UCD

Title	iVMp: an Interactive VM Placement Algorithm for Agile Capital Allocation
Authors(s)	Li, Xi, Ventresque, Anthony, Stokes, Nicola, Thorburn, James, Murphy, John
Publication date	2013-06-03
Publication information	Li, Xi, Anthony Ventresque, Nicola Stokes, James Thorburn, and John Murphy. "iVMp: An Interactive VM Placement Algorithm for Agile Capital Allocation." Institute of Electrical and Electronic Engineers (IEEE), 2013.
Conference details	IEEE 6th International Conference on Cloud Computing, Santa Clara Marriott, CA, USA, 28 June - 3 July, 2013
Publisher	Institute of Electrical and Electronic Engineers (IEEE)
Item record/more information	http://hdl.handle.net/10197/7144
Publisher's statement	© © 2013 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)	10.1109/CLOUD.2013.4

Downloaded 2023-04-07T10:52:29Z

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

iVMp: an Interactive VM Placement Algorithm for Agile Capital Allocation

Xi Li*, Anthony Ventresque*, Nicola Stokes*, James Thorburn†, John Murphy*

*Lero and School of Computer Science and Informatics

University College Dublin, Dublin, Ireland

Email: xi.li@ucdconnect.ie; {anthony.ventresque, nicola.stokes, j.murphy}@ucd.ie

†IBM Software Group, Toronto, Canada

Email: jthorbur@ca.ibm.com

Abstract— Server consolidation is an important problem in any enterprise, where capital allocators (CAs) must approve any cost saving plans involving the acquisition or allocation of new assets and the decommissioning of inefficient assets. Our paper describes *iVMp* an interactive VM placement algorithm, that allows CAs to become ‘agile’ capital allocators that can interactively propose and update constraints and preferences as placements are recommended by the system. To the best of our knowledge this is the first time that this interactive VM placement recommendation problem has been addressed in the academic literature. Our results show that the proposed algorithm finds near optimal solutions in a highly efficient manner.

Keywords—Interactive VM Placement; Server Consolidation; Capital Allocation

I. INTRODUCTION

Virtual Machine (VM) placement is a major problem in large-scale computing environments, as it encompasses various crucial concerns ranging from optimisation of energy consumption, the minimisation of new purchasing, load-balancing and server consolidation, to name but a few. The latter, server consolidation, aims to minimise the number of physical machines (PMs) required for a data-centre to operate by decommissioning inefficient PMs. The static and off-line instance of this problem, i.e. when the set of PMs and VMs is known in advance (off-line) and does not evolve over time (static), is of major concern to large companies, especially since servers in such environments, in the absence of this solution, tend to be utilised at sub-optimal capacity. In this consolidation scenario, any new placement is subject to approval from various capital allocators (CAs) in the organisation and any global plan is iteratively updated as placements are accepted or rejected. This requires an interactive solution, which can adapt to user feedback. The key challenge then becomes one of the timeliness of generating VM placement recommendations.

This interactive scenario has not been addressed in the academic literature, which largely focusses on techniques that provide optimised placement using various operation research methods (i.e. constraint programming, linear programming) [1][2] that tend to be very effective, but are too inefficient to provide a practical solution. Industry products also fall short of a solution for this specific case of server

consolidation. These products focus on either load-balancing concerns internal to a specific cluster of the enterprise, or individual migration of VMs. The focus of this research is on the automatic placement of VMs on PMs within the global context, encompassing many clusters, of many varieties. In short, there is no automatic VM placement approach for a heterogeneous large IT system that takes account of user feedback and the complex nature of this static offline server consolidation problem.

In this paper we present *iVMp*, an interactive heuristic based VM placement algorithm which can scale to large heterogeneous environments, and find a near optimal placement solution quickly. We test *iVMp* on real-traces from a data centre in IBM Toronto Canada with respect to effectiveness (quality) and efficiency (time). Our results show that *iVMp* generates a close to optimal solution in less than one second.

II. SYSTEM AND ALGORITHM DESCRIPTION

iVMp (see Figure 1) is an agile solution where CAs interact with plans recommended by the system. An enterprise IT infrastructure is composed of various virtual clusters (VCs), i.e. subsets of the whole infrastructure. In each of these VCs, CAs have specific permissions, obligations or special preferences such as a limit in the number of VMs that can run on a single PM (e.g., for licensing reasons), or that they do not allow specific VM-PM mappings (e.g., for cost allocation reasons), etc.

Figure 1. Overview of our Agile Capital Allocation

iVMp iteratively provides new plans as agile CAs validate or give feedback to the system (e.g., new constraints). In an

Table I
EXPERIMENTAL RESULTS WITH EMPTY LOAD (EL) AND ARTIFICIAL LOAD (AL)

Methods	No. of Hosts		CPU Wastage (GHz)		RAM Wastage (GB)		Resource Wastage Metric		Execution Time (sec)	
	EL	AL	EL	AL	EL	AL	EL	AL	EL	AL
FFD_RAM	20	29	740.8	969.94	3240.5	3525.9	4.563	6.189	0.059	0.064
FFD_CPU	74	82	15940.5	13724.7	6695.0	2843.9	48.857	41.440	0.059	0.059
DotProduct	20	28	811.2	1279.8	3240.5	3014.3	4.845	6.374	0.352	0.337
CPLEX	17	21	229.1	64.2	256.1	186.5	1.048	0.585	194.867	222.3
iVMp	17	22	88.2	77.2	1386.8	1185.8	1.030	1.874	0.627	0.711

enterprise, this is a long process that is undertaken on a regular basis, and rarely done for a global system.

The algorithm we implemented in iVMp is inspired by the Dot-Product heuristic [3]. Both Dot-Product and iVMp algorithm try to balance the different dimensions (e.g., RAM, CPU, network, disk) of servers, to limit the resource wastage. When they select a VM to be mapped on a PM, they chose the VM whose resource requirements are complementary to the residual capacity of the PM. However, while the Dot-Product is based on the scalar product and tends to be driven by the size of the VMs, iVMp uses the cosine and considers that the ‘shape’ of the VM is more important. Like many other heuristics, iVMp processes the PMs one by one in decreasing order of their size. For each PM, it first places the biggest VM, and then fills the residual capacity with other high ranked VMs (w.r.t. cosine value).

III. EVALUATION

The experiments presented in this paper consider two types of resources, CPU and RAM. However, our algorithm can run with more dimensions than these if necessary. We test the iVMp algorithm against two versions of the well known first-fit-decreasing heuristic (according to RAM or CPU), Dot-Product and IBM ILOG CPLEX, one of the leading optimisation problem solving tool. We consider the latter as the *optimal* solution. The efficiency parameters are the number of PMs required for the placement, the total CPU wastage (sum of remaining CPU on selected PMs), the total RAM wastage (ditto for RAM), and the resource wastage metric [4] (sum of difference between the normalized remaining resources of two dimensions on each PM). The smaller the wastage metrics, the more balanced the resource utilization between two dimensions. The dataset consists of real system traces from IBM Toronto, Canada: 342 PMs and 2,869 VMs, which is a large dataset compared to most of the experiments in the literature. The actual load values on the PMs were unavailable, so we explore two scenarios: the first *empty-load* (EL) where PMs are initially empty and the other called *artificial-load* (AL) with a 0%-30% random load on the PMs.

Table I shows that the efficiency of iVMp is very close to the optimal solution and outperforms all the other heuristics, while the execution time is very good (less than a second) and definitely orders of magnitude better than CPLEX.

Figure 2 compares the efficiency of iVMp and CPLEX when the number of VMs that need to be allocated varies.

For this evaluation we duplicated the set of 2,869 VMs. Figure 2 shows that iVMp scales up very well and always provides a solution in a few seconds regardless of the size of the dataset.

Figure 2. Scale-up, logarithmic scale: CPLEX and iVMp.

IV. CONCLUSION

We presented a system and an algorithm for interactive VM placement that generates near optimal solutions in real time for a real data centre. Historically, Capital Allocators have had to predict future requirements and then make purchasing decisions based on those projections - the ability to govern the allocation of capital stopped when the machines go into service. With iVMp, however, capital planning becomes agile capital allocation, opening the door to a massive improvement in the optimal use of computing resources to maximize ROI.

ACKNOWLEDGMENT

This work was supported, in part, by Science Foundation Ireland grant 10/CE/I1855 to Lero and by Enterprise Ireland Innovation Partnership in cooperation with IBM and University College Dublin under grant IP/2010/0061.

REFERENCES

- [1] F. Hermenier, X. Lorca, J. Menaud, G. Muller, and J. Lawall, Entropy: A Consolidation Manager for Clusters, in VEE, 2009.
- [2] E. Feller, L. Rilling, and C. Morin, Energy-aware Ant Colony Based Workload Placement in Clouds, in GRID, 2011.
- [3] S. Lee, R. Panigrahy, V. Prabhakaran, V. Ramasubramanian, K. Talwar, L. Uyeda, and U. Wieder, Validating Heuristics for Virtual Machines Consolidation, Microsoft Research Technical Report, 2011.
- [4] J. Xu and J. a. B. Fortes, Multi-Objective Virtual Machine Placement in Virtualized Data Center Environments, Green-Com & CPSCOM, 2010.