

Benchmarking Optimization Solvers and Symmetry Breakers for the Automated Deployment of Component-based Applications in the Cloud

Bogdan David¹, Mădălina Eraşcu¹

¹Faculty of Mathematics and Informatics, West University of Timișoara, Romania

Abstract

Optimization solvers based on methods from constraint programming (OR-Tools, Chuffed, Gecode), optimization modulo theory (Z3), and mathematical programming (CPLEX) are successfully applied nowadays to solve many non-trivial examples. However, for solving the problem of *automated deployment* in the Cloud of *component-based applications*, their computational requirements are huge making automatic optimization practically impossible with the ..current *general* optimization techniques. To overcome the difficulty, we exploited the sweet spots of the underlying problem in order to identify search space reduction methods. We came up with 15 symmetry breaking strategies which we tested in a static symmetry breaking setting on the solvers enumerated above and on 4 classes of problems.

As a result, all symmetry breaking strategies led to significant improvement of the computational time of all solvers, most notably, Z3 performed the best compared to the others. As an observation, the symmetry breaking strategies confirmed that, when applied in a static setting, they may interact badly with the underlying techniques implemented by the solvers.

Keywords

Cloud Computing, resource provisioning, Wordpress, optimization modulo theory, mathematical programming, constraint programming, symmetry breaking, Minizinc, OR-Tools, Chuffed, Gecode, Z3, IBM CPLEX.

1. Introduction

The problem of *automated deployment* in the Cloud of component-based applications received attention due to increased demand of digitalization of businesses. It consists of the following steps: (1) selection of the computing resources, (2) the distribution/assignment of the application components over the available computing resources, and (3) its dynamic modification to cope with peaks of user requests. In paper [3], we tackled only the first two steps of the deployment problem. In particular, our approach was used to synthesize the initial static optimal deployment of the application which consists of an assignment of application components to VMs such that the application functional requirements are fulfilled and costs are minimized.

The contributions of [3] are: (i) we formalized the Cloud deployment problem by abstracting the particularities of four classes of real-world problems; (ii) we proposed a methodology

SC² 2022: The 7th International Workshop on Satisfiability Checking and Symbolic Computation, August 12, 2022, Haifa, Israel

✉ Bogdan.David@@e-uvt.ro (B. David); Madalina.Erascu@@e-uvt.ro (M. Eraşcu)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

analyzing the particularities of the problem with the aim of identifying search space reduction methods (these are methods exploiting the symmetries of the general Cloud deployment problem, respectively methods utilizing the graph representation of the interaction between the components of each application); (iii) we assessed and compared the performance of two types of tools, namely mathematical programming (CPLEX [5]) and computational logic (the optimization modulo theory solver Z3 [6]); (iv) we identified limits in their scalability and applied six search space reduction methods aiming to improve their performance.

This paper extends [3] in the following ways:

1. a new formalization in the Minizinc [7] constraint modeling language;
2. the performance comparison adds the constraint programming solvers OR-Tools [8], Gecode [4] and Chuffed [2] which are available from the Minizinc IDE;
3. the list of the symmetry breakers from [3] is enriched with composition of all possible combinations of single symmetry breakers. These symmetry breakers are tested on the constraint programming solvers OR-Tools [8], Gecode [4] and Chuffed [2], optimization modulo theory solver Z3 [1], and mathematical programming solver CPLEX [5].

The rest of the paper is organized as follows. Section 2 briefly introduces the problem and the particularities of each formalization. Section 3 shows the need for strategies to reduce the search space and, to this aim, introduces symmetry breaking techniques as well as the principles to combine them. Conclusions of the experimental analysis are presented in Section 4.

2. Setting the Scene

Problem Definition The description of the problem first appeared in [3]. We have N interacting components, $C = \{C_1, \dots, C_N\}$, to be assigned to a set of M virtual machines, $V = \{V_1, \dots, V_M\}$. Each component C_i is characterized by a set of requirements concerning the hardware resources. Each virtual machine, V_k , is characterized by a *type*, which is comprised by hardware/software characteristics and leasing price. There are also *structural constraints* describing the interactions between components. The problem is to find:

1. an assignment matrix a with binary entries $a_{ik} \in \{0, 1\}$ for $i = \overline{1, N}$, $k = \overline{1, M}$, which are interpreted as follows: $a_{ik} = 1$ if C_i is assigned to V_k , and 0, otherwise; and
2. the type selection vector t with integer entries t_k for $k = \overline{1, M}$, representing the type (from a predefined set) of each VM leased;

such that: (i) the structural constraints, and (ii) the hardware requirements (capacity constraints) of all components are satisfied; and (iii) the purchasing/ leasing price is minimized.

The *structural constraints* are *application-specific* and derived in accordance with the analysis of the case studies. These are:

- *Conflict*: components in conflict cannot be deployed on the same VM.
- *Co-location*: components in co-location must be deployed on the same VM.
- *Exclusive deployment*: Only one of the components in exclusive deployment must be deployed in the same deployment plan.
- *Require-Provide*: one component requires or provides some functionalities offered, respectively provides, of another. Such an interaction induces constraints on the number of instances corresponding to the interacting components as follows.

- *Full deployment*: components in this relationship must be deployed on all VMs leased, except those which would induce conflicts between components.
- *Deployment with bounded number of instances* occur when the number of instances of deployed components must be equal, greater or less than some values.

General constraints are always considered in the formalization and are related to the: (i) *basic allocation* rules, (ii) *occupancy* criteria, (iii) hardware *capacity* of the VM offers, and (iv) *link* between the VM offers and the components hardware/software requirements.

We stated the problems as a linear constraint optimization problem (COP). We redirect the reader to [3] for a full description of it.

Problem Formalization The formalization for all three types of solvers, that is constraint programming solvers, SMT solvers, and mathematical programming solver, has almost a one-to-one correspondence between linear constraints present in the definition of COP and the implementation. We did not apply optimizations exploiting the particularities of the modelling languages because we wanted to have a fair comparison of the different formalisms. However, in the future work we plan to take advantage of their sweet spots.

Minizinc models The Minizinc models are the ones newly introduced in this paper. They are organized as follows: (i) there are surrogate models for each problem in which the maximum number of needed VMs is computed; (ii) there is a model gathering together all constraints, both general and application specific (the model corresponding to each application instantiates the constraints needed for its modeling) (iii) there is a model gathering together all symmetry breakers developed which are then instantiated based on the tests which want to be performed.

3. Experimental Analysis

The principles of the experimental analysis were introduced in [3]: on one hand, we want to assess the *scalability* of state-of-the-art general CP (Chuffed [2], Gecode [4], OR-Tools [8]), MP (CPLEX [5]) and OMT (Z3 [1]) tools in solving COPs corresponding to realistic case studies. On the other hand, we evaluate the *effectiveness* of various static symmetry breaking techniques in improving the computational time of solving these problems (see Section 3.1). This is because tests (see Tables 1-2) revealed that the naive application of general CP, MP and OMT techniques is not sufficient to solve realistic Cloud deployment applications.

We consider four case studies (Secure Web Container, Secure Billing Email Service, Oryx2, and Wordpress) which exhibit: (i) different hardware characteristics of components and the rich interactions type in between (structural constraints); (ii) the kind of linear constraints used to formalize the problem; and (iii) the kind of solution we are searching for. A full description of these case studies is in [3].

The scalability and effectiveness are evaluated from two perspectives: number of VM offers, respectively number of deployed instances of components. For *Secure Web Container*, *Secure Billing Email* and *Oryx2* applications, we considered up to 500 VM offers. Additionally, for the *Wordpress* application, we considered up to 12 instances of the Wordpress component to be deployed. The set of offers was crawled from the Amazon CPs offers list.

3.1. Experimental Settings

3.1.1. Selected Symmetry Breaking Strategies

Aiming to reduce the search space size, a set of strategies have been selected in order to exploit the particularities of the problem: (i) the VMs needed for application deployment might have different characteristics; (ii) applications components might be in conflict hence conflict-type constraints can be exploited; (iii) the number of instances to be deployed is unknown.

Our approach is incremental and experimental: we start with traditional symmetry breakers that have been used for other problems related to bin-packing and combine them with the aim of further search space reduction.

3.1.1.1. Simple symmetry breakers

Price-based ordering (PR). This strategy aims to break symmetry by ordering the vector containing the types of used VMs decreasingly by price, i.e. $p_k \geq p_{k+1}$, $k = \overline{1, M-1}$. This means that the solution will be characterized by the fact that the columns of the assignment matrix will be ordered decreasingly by the price of the corresponding VMs.

Lexicographic ordering (LX). This corresponds to the traditional strategy aiming to break column-wise symmetries. The constraints to be added aiming to ensure that two columns, k and $(k+1)$ are in a decreasing lexicographic order, i.e. $a_{*k} \succ_{lex} a_{*(k+1)}$, are $\bigwedge_{l=1}^{i-1} (a_{lk} = a_{l(k+1)}) \implies a_{ik} \geq a_{i(k+1)}$, $\forall i = \overline{1, N}$.

Load-based ordering (L). This is a column-wise symmetry breaker which orders decreasingly the columns by the number of the component instances they accommodate: $\sum_{i=1}^N a_{ik} \geq \sum_{i=1}^N a_{i(k+1)}$, $k = \overline{1, M-1}$.

Fixed values (FV). The search space can be reduced also by fixing the values of some variables starting from the application specific constraints, in particular conflict constraints. More precisely, the graph composed by the components being in conflict is used to identify components which must be placed on different machines and hence the values of the corresponding decision variables are fixed. The identification of these components is done by constructing the clique with maximum deployment size.

3.1.1.2. Composed symmetry breakers The symmetry breakers above can be composed leading to the following symmetry breakers:

- FV-PR, FV-L, FV-LX, PR-L, PR-LX, L-PR, L-LX,
- FV-PR-L, FV-PR-LX, FV-L-PR, FV-L-LX, PR-L-LX, L-PR-LX,
- FV-PR-L-LX, FV-L-PR-LX

These symmetry breakers are so the subsequent breaks ties of the former. For example, FV-PR fixes on separate machines the decision variables corresponding to the component instances being in the clique with maximum deployment size and the machines left unoccupied are ordered decreasingly by price. In the case of PR-L-LX, the machines are ordered decreasingly

by price, those with the same price are ordered decreasingly by the number of components they host and those with the same number of instances are ordered lexicographically.

It is worth noticing that the symmetry breakers involving FV must apply FV the very first. This is because FV is used as a preprocessing step which has a positive impact on the solvers as it introduces equalities.

3.1.2. Software and Hardware Settings

We used Minizinc v0.7.0 as the constraint modeling language. We mention that the Minizinc models follow the formalization and no optimizations were performed because we wanted to be as close as possible to the OMT and CPLEX formalizations in order to have a fair computational comparison between the newly considered solvers and the existing results. The CP solvers used (Chuffed, Gecode, OR-Tools) are called from Minizinc IDE with the default values for parameters. The OMT formalization is done using the Z3 Python API and uses quantifier-free linear integer arithmetic. Z3 was used with the default values of the parameters. In the case of the mathematical programming solver CPLEX, we used the Python API with the no symmetry breaking option manually activated.

The source code and the experimental results are available online at <https://github.com/BogdanD02/Cloud-Resource-Provisioning>, release v1.0.0. All reported timings are in seconds. They only include the actual solving time of the optimization problem and not the pre-processing steps.

All tests in this paper were performed on an Intel(R) Core (TM) i5-9400F CPU @ 3.90GHz using Chuffed v0.10.4, Gecode v6.3.0, OR-Tools v9.0.0, CPLEX v12.9.0 and Z3 v4.10.2.

3.2. Results

Tables 1-4 include the results obtained without using symmetry breaking strategies. The list of offers (columns #o) was crawled from the Amazon site¹. Each list of VM offers covers the main instance types, for example, `small`, `medium`, `large`. The list of offers can be viewed as a containment hierarchy (i.e. the list of 20 offers is included in the list of 40 offers etc.). The tables include only those cases for which we obtained a result in a 40 minutes timeframe. The missing values (-) mean that no solution is returned in this timeframe.

One can observe that CPLEX scales the worst for all case studies while OR-Tools the best. However, none of the tools scale for Wordpress with more than 6 instances and several dozens of offers.

To overcome the lack of scalability issue, we applied the symmetry breaking strategies described in Section 3.1.1.

4. Discussion of the Results

We conducted various tests involving 5 solvers and 15 symmetry breakers. Because of lack of space, we can not include all of them here and we direct the reader to check

¹<https://aws.amazon.com/>

Table 1

Scalability tests for Wordpress with different instances (#i). Time values are expressed in seconds.

#i	#o=20	#o=40	#o=250	#o=500	#o=20	#o=40	#o=250	#o=500
OR-Tools					CPLEX			
3	3.49	8.38	96.05	191.04	9.66	-	-	-
4	23.25	56.07	501.43	987.91	121.96	-	-	-
5	149.47	425.03	-	-	446.02	-	-	-
6	493.39	1173.46	-	-	664.68	-	-	-
Gecode					Chuffed			
3	2.13	2.19	-	-	2.05	3.7	45.88	447.56
4	14.84	23.83	-	-	23.73	114.18	1866.61	-
5	162.13	-	-	-	531.19	2278.76	-	-
Z3								
3	2.82	4.13	103.19	391.87				
4	46.46	275.81	-	-				

Table 2

Scalability tests for Oryx2. Time values are expressed in seconds.

#o=20	#o=40	#o=250	#o=500	#o=20	#o=40	#o=250	#o=500
OR-Tools				CPLEX			
0.95	1.12	4.25	5.79	0.16	0.54	-	-
Gecode				Chuffed			
70.99	104.54	234.27	465.72	128.67	154.11	294.25	396.58
Z3							
13.35	15.36	453.2	717.99				

Table 3

Scalability tests for Secure Billing Email Service. Time values are expressed in seconds.

#o=20	#o=40	#o=250	#o=500	#o=20	#o=40	#o=250	#o=500
OR-Tools				CPLEX			
0.85	1.22	8.04	18.68	0.73	4.73	378.47	-
Gecode				Chuffed			
0.73	1.42	76.63	119.32	0.79	0.82	1.35	3.11
Z3							
0.34	0.70	3.37	8.29				

them in the folder PlotData from release v1.0.0 at <https://github.com/BogdanD02/Cloud-Resource-Provisioning>.

We draw the following remarks:

1. Using the 15 symmetry breakers, out of the 5 solvers, the best one from the computational time point of view is Z3.
2. For the virtual best solver, that is Z3, the best symmetry breaker is FVPR. This is because

Table 4

Scalability tests for Secure Web Container. Time values are expressed in seconds.

#o=20	#o=40	#o=250	#o=500	#o=20	#o=40	#o=250	#o=500
OR-Tools				CPLEX			
1.50	4.85	28.79	57.33	2.68	49.64	-	-
Gecode				Chuffed			
10.15	43.84	76.63	779.44	0.89	1.01	2.64	4.54
Z3							
0.45	1.14	8.28	17.48				

we have many Wordpress files to be analyzed (40 files corresponding to Wordpress with 3 up to 12 instances and 20, 40, 250, 500 offers), compared to the other applications (4 files corresponding to 20, 40, 250, 500 offers for each of the other applications) for which reduction methods which exploit the graph representation, that is FV, is of benefit.

3. We also considered the best symmetry breaker for Z3 for each of the case studies. In case of Secure Web Container and Secure Billing Email Service applications the best is FVL, while for Oryx2 is FVLX. There is no surprise that symmetry breakers involving FV give best results, however for more reliable results we should consider more test cases for Secure Billing, Secure Web and Oryx2, since now there are only 4 files analyzed corresponding to different number of offers. We plan to run more tests for each of the case studies for a more accurate analysis.
4. One would expect that the best symmetry breaker is one composing a higher number of individual symmetry breakers as more symmetries are broken so the search space is significantly reduced. However, this is not true: FVPR, composing 2 symmetry breakers, is better than those composing 3 or 4. An explanation for this is, on one hand the number of added constraints which influence the solving time, on the other hand, when using static symmetry breaking, the symmetry breakers can interact badly with the SMT solvers which we used as black box.

References

- [1] Bjørner, N., Phan, A., Fleckenstein, L.: νZ - An Optimizing SMT Solver. In: Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, London, UK, April 11-18, 2015. Proceedings. pp. 194–199 (2015)
- [2] Chu, G., Stuckey, P.J., Schutt, A., Ehlers, T., Graeme Gange, K.F.: Chuffed, a lazy clause generation solver (2022), <https://github.com/chuffed/chuffed>, last accessed 20 August 2022
- [3] Erascu, M., Micota, F., Zaharie, D.: Scalable optimal deployment in the cloud of component-based applications using optimization modulo theory, mathematical programming and symmetry breaking. J. Log. Algebraic Methods Program. **121**, 100664 (2021)
- [4] Gecode Team: Gecode: Generic Constraint Development Environment (2006), available from <http://www.gecode.org>
- [5] IBM: IBM ILOG CPLEX Optimization Studio CPLEX User's Manual. Version 12 (2016)

- [6] de Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
- [7] Nethercote, N., Stuckey, P.J., Becket, R., Brand, S., Duck, G.J., Tack, G.: Minizinc: Towards a standard CP modelling language. In: Bessiere, C. (ed.) Principles and Practice of Constraint Programming - CP 2007, 13th International Conference, CP 2007, Providence, RI, USA, September 23-27, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4741, pp. 529–543. Springer (2007)
- [8] Perron, L., Furnon, V.: OR-Tools, <https://developers.google.com/optimization/>