

Bayless, Sam; Kodirov, Nodir; Iqbal, Syed M.; Beschastnikh, Ivan; Hoos, Holger H.; Hu, Alan J.

Scalable constraint-based virtual data center allocation. (English) Zbl 07153703

Artif. Intell. 278, Article ID 103196, 26 p. (2020)

Summary: Constraint-based techniques can solve challenging problems arising in highly diverse applications. This paper considers the problem of virtual data center (VDC) allocation, an important, emerging challenge for modern data center operators. To address this problem, we introduce NETSOLVER, a system for VDC allocation that is based on constraint solving. NETSOLVER represents a major improvement over existing approaches: it is sound, complete, and scalable, providing support for end-to-end, multi-path bandwidth guarantees across all the layers of hosting infrastructure, from servers to top-of-rack switches to aggregation switches to access routers. NETSOLVER scales to realistic data center sizes and VDC topologies, typically requiring just seconds to allocate VDCs of 5–15 virtual machines to physical data centers with 1000+ servers, maintaining this efficiency even when the data center is nearly saturated. In many cases, NETSOLVER can allocate 150%–300% as many total VDCs to the same physical data center as previous methods. Finally, we show how NETSOLVER can be extended with additional optimization constraints, such as VM affinity and hotspot minimization, demonstrating the flexibility of our approach.

The performance and flexibility of NETSOLVER are made possible by our formalization of the VDC allocation problem in terms of multi-commodity flows, and the corresponding efficient handling of network flow problems in the underlying constraint solvers. This shows the importance of supporting flow-based constraints, which are more mature in ILP- vs. SMT-based constraint solving.

MSC:

68Txx Artificial intelligence

Keywords:

constraint solver; data center; virtual data center; allocation

Software:

CPLEX; AClib ; Gurobi; KLEE; SMAC ; z3

Full Text: [DOI](#)

References:

- [1] Prasad, M. R.; Biere, A.; Gupta, A., A survey of recent advances in SAT-based formal verification, *Int. J. Softw. Tools Technol. Transf.*, 7, 2, 156-173 (2005)
- [2] Cadar, C.; Dunbar, D.; Engler, D., KLEE: unassisted and automatic generation of high-coverage tests for complex systems programs, (*Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation (2008)*), USENIX Association: USENIX Association Berkeley, CA, USA), 209-224
- [3] Rintanen, J., Madagascar: efficient planning with SAT, (*The 2011 International Planning Competition (2011)*), 61
- [4] Biere, A.; Kröning, D., Sat-based model checking, (*Handbook of Model Checking (2018)*), Springer), 277-303 · [Zbl 1392.68232](#)
- [5] Gurobi, Gurobi Optimizer Reference Manual (2014)
- [6] Bayless, S.; Bayless, N.; Hoos, H.; Hu, A., SAT modulo monotonic theories, (*Proceedings of the 29th AAAI Conference on Artificial Intelligence (2015)*)
- [7] Popa, L.; Kumar, G.; Chowdhury, M.; Krishnamurthy, A.; Ratnasamy, S.; Stoica, I., Faircloud: sharing the network in cloud computing, (*Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM) (2012)*), ACM), 187-198
- [8] Lee, J.; Turner, Y.; Lee, M.; Popa, L.; Banerjee, S.; Kang, J.-M.; Sharma, P., Application-driven bandwidth guarantees in datacenters, *ACM SIGCOMM Comput. Commun. Rev.*, 44, 467-478 (2014)
- [9] Guo, C.; Lu, G.; Wang, H. J.; Yang, S.; Kong, C.; Sun, P.; Wu, W.; Zhang, Y., Secondnet: a data center network virtualization architecture with bandwidth guarantees, (*Proceedings of the 6th International Conference on Emerging Networking Experiments and Technologies. Proceedings of the 6th International Conference on Emerging Networking Experiments and Technologies, Co-NEXT '10 (2010)*), ACM), 15

- [10] Bayless, S.; Kodirov, N.; Beschastnikh, I.; Hoos, H. H.; Hu, A. J., Scalable constraint-based virtual data center allocation, (Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI) (2017))
- [11] Hutter, F.; Hoos, H. H.; Leyton-Brown, K., Sequential model-based optimization for general algorithm configuration, (Proceedings of the 5th Learning and Intelligent Optimization Conference (LION '11), vol. 5 (2011)), 507-523
- [12] Tantawi, A. N., A scalable algorithm for placement of virtual clusters in large data centers, (2012 IEEE 20th International Symposium on Modeling, Analysis & Simulation of Computer and Telecommunication Systems (MASCOTS) (2012), IEEE), 3-10
- [13] Ballani, H.; Costa, P.; Karagiannis, T.; Rowstron, A., Towards predictable datacenter networks, *ACM SIGCOMM Comput. Commun. Rev.*, 41, 242-253 (2011)
- [14] Zhani, M. F.; Zhang, Q.; Simona, G.; Boutaba, R., Vdc planner: dynamic migration-aware virtual data center embedding for clouds, (2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013) (2013), IEEE), 18-25
- [15] Rost, M.; Fuerst, C.; Schmid, S., Beyond the stars: revisiting virtual cluster embeddings, *ACM SIGCOMM Comput. Commun. Rev.*, 45, 3, 12-18 (2015)
- [16] Yu, M.; Yi, Y.; Rexford, J.; Chiang, M., Rethinking virtual network embedding: substrate support for path splitting and migration, *ACM SIGCOMM Comput. Commun. Rev.*, 38, 2, 17-29 (2008)
- [17] Cheng, X.; Su, S.; Zhang, Z.; Wang, H.; Yang, F.; Luo, Y.; Wang, J., Virtual network embedding through topology-aware node ranking, *ACM SIGCOMM Comput. Commun. Rev.*, 41, 2, 38-47 (2011)
- [18] Chowdhury, N. M.K.; Rahman, M. R.; Boutaba, R., Virtual network embedding with coordinated node and link mapping, (Proceedings of the IEEE International Conference on Computer Communications (INFOCOM) (2009), IEEE), 783-791
- [19] Lischka, J.; Karl, H., A virtual network mapping algorithm based on subgraph isomorphism detection, (Proceedings of the 1st ACM Workshop on Virtualized Infrastructure Systems and Architectures (2009), ACM), 81-88
- [20] Huang, T.; Rong, C.; Tang, Y.; Hu, C.; Li, J.; Zhang, P., Virtualrack: bandwidth-aware virtual network allocation for multi-tenant datacenters, (2014 IEEE International Conference on Communications (ICC) (2014), IEEE), 3620-3625
- [21] Yuan, Y.; Wang, A.; Alur, R.; Loo, B. T., On the feasibility of automation for bandwidth allocation problems in data centers, (Formal Methods in Computer-Aided Design (FMCAD), 2013 (2013), IEEE), 42-45
- [22] Meng, X.; Pappas, V.; Zhang, L., Improving the scalability of data center networks with traffic-aware virtual machine placement, (Proceedings of the IEEE International Conference on Computer Communications (INFOCOM) (2010), IEEE), 1-9
- [23] Kakadia, D.; Kopri, N.; Varma, V., Network-aware virtual machine consolidation for large data centers, (Proceedings of the Third International Workshop on Network-Aware Data Management (2013), ACM), 6
- [24] Ballani, H.; Gunawardena, D.; Karagiannis, T., Network Sharing in Multi-Tenant Datacenters (February 2012), Tech. rep.
- [25] LaCurts, K.; Mogul, J. C.; Balakrishnan, H.; Turner, Y., Cicada: Introducing Predictive Guarantees for Cloud Networks, vol. 14, 14-19 (2014)
- [26] Angel, S.; Ballani, H.; Karagiannis, T.; O'Shea, G.; Thereska, E., End-to-end performance isolation through virtual datacenters, (Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI) (2014)), 233-248
- [27] Raiciu, C.; Barre, S.; Pluntke, C.; Greenhalgh, A.; Wischik, D.; Handley, M., Improving datacenter performance and robustness with multipath TCP, *ACM SIGCOMM Comput. Commun. Rev.*, 41, 266-277 (2011)
- [28] Alizadeh, M.; Edsall, T.; Dharmapurikar, S.; Vaidyanathan, R.; Chu, K.; Fingerhut, A.; Matus, F.; Pan, R.; Yadav, N.; Varghese, G.; Conga, G., Conga: distributed congestion-aware load balancing for datacenters, *ACM SIGCOMM Comput. Commun. Rev.*, 44, 503-514 (2014)
- [29] Vahdat, A., A look inside Google's data center network ("2017)
- [30] Guo, C.; Lu, G.; Li, D.; Wu, H.; Zhang, X.; Shi, Y.; Tian, C.; Zhang, Y.; Lu, S., BCube: a high performance, server-centric network architecture for modular data centers, *ACM SIGCOMM Comput. Commun. Rev.*, 39, 4, 63-74 (2009)
- [31] Greenberg, A.; Hamilton, J. R.; Jain, N.; Kandula, S.; Kim, C.; Lahiri, P.; Maltz, D. A.; Patel, P.; Sengupta, S., VL2: a scalable and flexible data center network, *ACM SIGCOMM Comput. Commun. Rev.*, 39, 51-62 (2009)
- [32] TRILL: transparent interconnection of lots of links ("2017)
- [33] Duffield, N. G.; Goyal, P.; Greenberg, A.; Mishra, P.; Ramakrishnan, K. K.; van der Merive, J. E., A flexible model for resource management in virtual private networks, *ACM SIGCOMM Comput. Commun. Rev.*, 29, 95-108 (1999)
- [34] Belbekkouche, A.; Hasan, M. M.; Karmouch, A., Resource discovery and allocation in network virtualization, *IEEE Commun. Surv. Tutor.*, 14, 4, 1114-1128 (2012)
- [35] Fischer, A.; Botero, J. F.; Beck, M. T.; De Meer, H.; Hesselbach, X., Virtual network embedding: a survey, *IEEE Commun. Surv. Tutor.*, 15, 4, 1888-1906 (2013)
- [36] Fischer, A.; Botero Vega, J. F.; Duelli, M.; Schlosser, D.; Hesselbach Serra, X.; De Meer, H., Alevin—a framework to develop, compare, and analyze virtual network embedding algorithms, (Electronic Communications of the EASST (2011)), 1-12
- [37] Bari, M. F.; Boutaba, R.; Esteves, R.; Granville, L. Z.; Podlesny, M.; Rabbani, M. G.; Zhang, Q.; Zhani, M. F., Data center network virtualization: a survey, *IEEE Commun. Surv. Tutor.*, 15, 2, 909-928 (2013)
- [38] Gupta, A.; Kleinberg, J.; Kumar, A.; Rastogi, R.; Yener, B., Provisioning a virtual private network: a network design problem for multicommodity flow, (Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing (2001), ACM), 389-398 · [Zbl 1323.68014](#)
- [39] Szeto, W.; Iraqi, Y.; Boutaba, R., A multi-commodity flow based approach to virtual network resource allocation, (Global Telecommunications Conference (GLOBECOM), vol. 6 (2003), IEEE), 3004-3008

- [40] Even, S.; Itai, A.; Shamir, A., On the complexity of time table and multi-commodity flow problems, (Proceedings of the 16th Annual Symposium on Foundations of Computer Science (1975), IEEE), 184-193
- [41] Sir, M. Y.; Senturk, I. F.; Sisikoglu, E.; Akkaya, K., An optimization-based approach for connecting partitioned mobile sensor/actuator networks, (2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS) (2011), IEEE), 525-530
- [42] Kaufman, D. E.; Nonis, J.; Smith, R. L., A mixed integer linear programming model for dynamic route guidance, *Transp. Res., Part B, Methodol.*, 32, 6, 431-440 (1998)
- [43] IBM, ILOG CPLEX Optimization Studio 12.6.1: multi-commodity flow cuts
- [44] Marchand, H.; Martin, A.; Weismantel, R.; Wolsey, L., Cutting planes in integer and mixed integer programming, *Discrete Appl. Math.*, 123, 1-3, 397-446 (2002) · [Zbl 1130.90370](#)
- [45] Bayless, S., SAT Modulo Monotonic Theories (2017), University of British Columbia, Ph.D. thesis
- [46] Bar-Yehuda, R.; Even, S., A local-ratio theorem for approximating the weighted vertex cover problem, *North-Holl. Math. Stud.*, 109, 27-45 (1985) · [Zbl 0557.90072](#)
- [47] Eén, N.; Sorensson, N., Translating pseudo-Boolean constraints into SAT, *J. Satisf. Boolean Model. Comput.*, 2, 1-26 (2006) · [Zbl 1116.68083](#)
- [48] Sorensson, N.; Een, N., Minisat – a SAT solver with conflict-clause minimization, (SAT 2005, vol. 53 (2005)), 1-2
- [49] De Moura, L.; Bjørner, N., Z3: an efficient SMT solver, (Tools and Algorithms for the Construction and Analysis of Systems (2008)), 337-340
- [50] Hutter, F.; López-Ibáñez, M.; Fawcett, C.; Lindauer, M.; Hoos, H. H.; Leyton-Brown, K.; Stützle, T., Aclib: a benchmark library for algorithm configuration, (International Conference on Learning and Intelligent Optimization (2014), Springer), 36-40
- [51] Bodík, P.; Menache, I.; Chowdhury, M.; Mani, P.; Maltz, D. A.; Stoica, I., Surviving failures in bandwidth-constrained datcenters, (Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (2012), ACM), 431-442
- [52] Andreyev, A., Introducing data center fabric, the next-generation Facebook data center network (“2017)
- [53] Isard, M.; Prabhakaran, V.; Currey, J.; Wieder, U.; Talwar, K.; Goldberg, A., Quincy: fair scheduling for distributed computing clusters, (Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (2009), ACM), 261-276
- [54] Zaharia, M.; Borthakur, D.; Sen Sarma, J.; Elmeleegy, K.; Shenker, S.; Stoica, I., Delay scheduling: a simple technique for achieving locality and fairness in cluster scheduling, (Proceedings of the 5th European Conference on Computer Systems (2010), ACM), 265-278
- [55] Sherry, J.; Hasan, S.; Scott, C.; Krishnamurthy, A.; Ratnasamy, S.; Sekar, V., Making middleboxes someone else’s problem: network processing as a cloud service, *ACM SIGCOMM Comput. Commun. Rev.*, 42, 4, 13-24 (2012)

This reference list is based on information provided by the publisher or from digital mathematics libraries. Its items are heuristically matched to zbMATH identifiers and may contain data conversion errors. It attempts to reflect the references listed in the original paper as accurately as possible without claiming the completeness or perfect precision of the matching.