

Traffic-Aware Clustering and VM Migration in Distributed Data Center

Marco Cello
University of Genoa
Via dell'Opera Pia 13, 16145
Genova, Italy
marco.cello@unige.it

Kang Xi,
Jonathan H. Chao
NYU Polytechnic School of
Engineering
6 MetroTech Center, 11210
New York, USA
kang.xi@nyu.edu,
chao@nyu.edu

Mario Marchese
University of Genoa
Via dell'Opera Pia 13, 16145
Genova, Italy
mario.marchese@unige.it

ABSTRACT

In this paper we propose an algorithmic approach designed to tackle and reduce the congestion events in a Distributed Data Center (DDC). Our solution is based on virtual machines (VMs) migration and, differently from the literature, it analyzes the VMs communication patterns in order to find “tight” clusters of VMs to be migrated.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations; I.5.3 [Pattern Recognition]: Clustering

Keywords

VM placement; congestion control; cloud computing; data center

1. INTRODUCTION AND VM MIGRATION

Nowadays, it is common practice for cloud service providers to run their own geo DDC. There are many advantages of this architecture: 1) Mega data centers are difficult to build, thus delay the time-to-market; 2) the requirements for space and cooling cause building costs to elevate as data centers grow large; 3) physical limitations at a geographic region, such as land size and local energy availability.

Cloud computing tasks are intrinsically highly dynamic making the communications among VMs hardly to be assumed to be perfectly known or predictable. During the cloud service operations, in fact, it is very common that VM-VM communications take place among different DCs. If the data rate among VMs increases, the WAN link connecting the two DCs could experience congestion. This fact has dramatic consequences because moderated level of link

congestion can substantially reduce the QoS experience by the tenants through a high packet loss rate and a high packet delay. Moreover, long-haul inter data center communications are significantly more expensive than intra data center one and a provider may want to keep as much inter-VM communication inside data centers as possible.

In this paper we propose a solution able to reduce the link utilization and to solve the link congestion during the cloud operation in a DDC scenario. Current solutions based on the optimal VMs placement on DCs are impractical either for a high computational complexity (NP-hard) or for the high number of VMs migrated. Heuristic solutions are more practical and useful. Differently to the literature we leverage the concept of VMs cluster and we propose an algorithm that finds a strong connected cluster of VMs inside the DDC and migrate them to another DC to reduce the WAN link utilization and to solve the congestion event.

In particular, the key observation behind our approach is that the migration of single VMs, completely ignores possible information about amount of traffic between the VM and the others within the DCs. Therefore, a more effective method is based on the analysis of the VMs communication patterns in order to find subsets or *clusters* of VMs that are very “tight” among them. The core of our algorithm is then the “clustering algorithm”. We found that the well-known *Single-linkage clustering* algorithm reaches good results.

The algorithm has been implemented using Matlab and tested using the HPC platform of the New York University. Simulations have been conducted using three different scenarios (small, medium and large scenarios) varying the number of DCs and VMs. The traffic matrix of the VMs has been randomly generated using realistic values used in literature. For a small-size scenario, our solution has been compared to the Optimal Placement Problem (OPP) that finds the optimal placement for all the VMs in DDC. Optimal solution reaches a maximum link utilization of 0.79, while our algorithm reaches 0.91 (15% more of the optimal). But the computational time of OPP solution differs from our solution of 5 orders of magnitude making it totally impractical also for very small scenarios. In the large-size scenario, composed by 20 DCs and 2000 VMs, our algorithm performs very well: the utilization on the congested link is reduced by 27.8% over 5 routines of the algorithm and with an average of 2.4 VMs migrated in each routine. Each algorithm routine takes an average of 22.19 seconds to be completed.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

DCC'14, August 18, 2014, Chicago, Illinois, USA.

ACM 978-1-4503-2992-7/14/08.

<http://dx.doi.org/10.1145/2627566.2627582>