

Connectivity-Aware Virtual Machine Placement in 60GHz Wireless Cloud Centers

Linghe Kong¹, Linsheng Ye¹, Bowen Wang¹, Xiaofeng Gao¹, Fan Wu¹,
Guihai Chen¹, and M. Shamim Hossain²

¹ Shanghai Key Laboratory of Scalable Computing and Systems,
Department of Computer Science and Engineering,
Shanghai Jiao Tong University, Shanghai 200240, China.

² College of Computer and Information Sciences,
King Saud University, Riyadh 11543, Saudi Arabia.
`linghe.kong@sjtu.edu.cn`

Abstract. Benefiting from the 60GHz technology, physical machines in advanced cloud centers are connected by not only the conventional wired links but also the wireless communications. The 60GHz *millimeter-wave* (mmWave) introduces valuable advantages into cloud centers including flexibility, scalability and high rate. Nevertheless, mmWave is constrained by directional communications, i.e., a wireless link is connected if and only if two directional antennas face to each other. This constraint introduces a new problem in cloud service: the *virtual machine* (VM) placement should consider the real-time connectivity if communications are required between VMs. Otherwise, rotating the antenna costs additional delay, resulting in performance degradation. To address this problem, we propose a novel *connectivity-aware VM placement* (CAVMP) specialized for 60GHz wireless cloud center. The core of CAVMP is to dynamically place VMs in order to improve the utilization and avoid overloads while taking the connectivity state into account. We build a 2-rack cloud to measure the connectivity feature of mmWave communications. In addition, we conduct extensive simulations to evaluate CAVMP. Performance results demonstrate that CAVMP significantly outperforms existing VM placement schemes in wireless cloud center.

1 Introduction

Recent years, both industry and academia pay great attention to cloud computing. In industrial field, plenty of cloud centers are built all over the world to provide promising cloud services such as Amazon EC2, Microsoft Azure and Alibaba Cloud. In academic field, many efforts have been contributed on cloud computing from different directions, e.g., FairCloud [12] and ElasticSwitch [13].

With the development of 60GHz *millimeter-wave* (mmWave) technology, wireless cloud centers [4] are available from lab to market, in which racks are able to transmit data by wireless communications. Compared with the conventional wired links, wireless communications show three merits. (i) The topology is flexible by adjusting the direction of wireless transmission; (ii) Getting rid of the complex cabling system, the cloud center is easily scalable by deploying directional antenna by putting antennas on racks' roofs. (iii) The 60GHz band

has up to 7GHz channel bandwidth. With a suitable modulation, the data rate can easily achieve multi-gigabit level. It is promising that using wireless communications as a complementary for wired links in cloud centers.

However, different from the omni-directional transmissions in conventional WiFi, mmWave communications require directional transmissions due to its high attenuation in the air. Field test [6] shows a 7-degree and 15-meter communication range of the mmWave device in real prototype. During data transmissions, it is required that the antennas in two mmWave transceivers face to each other. In addition, rotating the mmWave antenna introduce extra delay no matter by motor or beamforming. This feature causes a new problem, connectivity awareness, in *virtual machine* (VM) placement in cloud computing for minimizing the time consumption on antenna adjustment.

In cloud centers, VM placement is a fundamental problem, which dynamically allocates the resources of physical machines (PMs) including CPU and memory according to the applications. Inappropriate VM placement results in over-provision or under-provision issues. The over-provision wastes unnecessary energy and resource. And the under-provision degrades the performance.

Great efforts have been contributed in VM placement in literature. In [1], the VM placement is developed to manage the service level agreement (SLA) violations. Then, [8] proposes the joint-VM-provisioning, which considers the correlation among VMs. Furthermore, MPT [15] is designed to combine the utilization ratio as well as the energy efficiency. Communication traffic are considered in [7]. These works optimize different metrics for VM placement. However, none of them involves in the connectivity.

In this paper, we propose a new *connectivity-aware VM placement* (CAVMP) scheme that dynamically allocates the VMs according to the connectivity state to enhance the utilization of cloud center. There are four major components in CAVMP. First, the *resource monitor* senses the real-time utilization and connectivity state. Second, based on the sensory and historical results, the *utilization estimator* predicts the future utilization of existing VMs because the utilization of these VMs is time-varying. Third, the *VM requestor* reports the request of new VMs from users. Fourth, the *strategy manager* makes the decision on new VM placement and antenna adjustment.

The contributions of this work are three-fold:

- To the best of our knowledge, this is the first work to study the problem of connectivity-aware VM placement in 60GHz wireless cloud center.
- To tackle this problem, we propose a novel scheme, named CAVMP. In this scheme, the connectivity tracking is newly added into the resource monitor module; ARIMA and GARCH models are adopted to estimate the future utilization and its volatility; and a new algorithm is developed to make the decision of VM placement.
- We build a wireless cloud prototype including 2 racks equipped with mmWave devices. Based on the communication features, we conduct extensive simulation to evaluate CAVMP. Performance results demonstrate that CAVMP improves up to 8% average utilization ratio while reducing 95% time consumption on antenna adjustment, compared with the state-of-the-art scheme.




Fig. 1. Illustrations of directional transmissions in existing wireless clouds.

2 Preliminary of 60GHz Wireless Cloud Center

In cloud computing, 60GHz mmWave communication technology is exploited to supplement the wired links, which has the following advantages:

First, the topology of wireless cloud is flexible by rotating antennas. The wireless cloud Flyways [6] was implemented as an incremental overlay network to the wired cloud. The mmWave devices HXI with horn antenna are placed on top of racks to generate wireless links as shown in Fig. 1(a). Then, Zhou et al. [17] proposed 3D Beamforming to establish the indirect link between two racks by reflecting the signal using the ceiling as shown in Fig. 1(b). Both Flyways and 3D Beamforming can adjust their antenna directions to change the topology.

Second, it is not easy to expand a wired cloud center due to the complicated cabling workload, which wastes huge amount of time and manpower. However, the wireless cloud center is able to get rid of the cabling procedure and easy to be scaled up. In addition, the cloud center can decrease the construction and maintenance cost without wired cables.

Third, mmWave communication also enables high-speed data transmissions, which satisfies the transmission requirements of cloud centers. The 60GHz spectrum was set as unlicensed band by the FCC in 2001 and the available 7GHz channel bandwidth supports multi-gigabit wireless communications.

However, the limitation of mmWave is its directional transmission, so that two mmWave devices are connected only if their antenna beams can cover each other. Hence, when one device needs to link a new device, the antenna directions need to be adjusted, which leads to additional time consumption.

3 Design of Connectivity-Aware VM Placement

The objective of the virtual machine placement is to improve the average utilization of physical machines (PMs). Without considering the limitation of directional transmission, existing VM placement schemes cannot perform well in 60GHz cloud center. Since some applications require data transmission among multiple VMs, (i) the performance will dramatically degrade if there is no link establishment; (ii) the utilization ratio will be reduced if time is consumed on antenna adjustment. To address this new problem in wireless cloud, we plan to construct a new *connectivity-aware VM placement (CAVMP)* scheme.

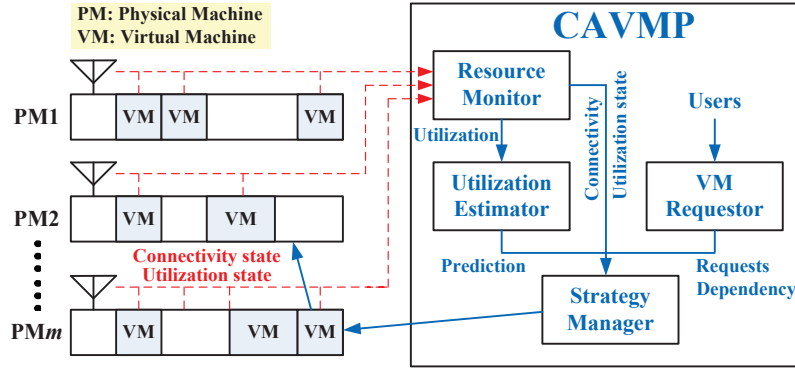


Fig. 2. System architecture of CAVMP.

3.1 Design Overview

The idea of CAVMP is to place VMs according to the connectivity state. It is better that dependent VMs (data exchange are required among them) are placed in two already connected PMs to minimize the time for antenna adjustment.

In the system, we assume that there are thousands of PMs in the wireless cloud; a rack is composed of multiple PMs; VM placement happens periodically every Δt ; the time consumption of antenna adjustment cannot be ignored, which is proportional to the given angle speed; the duration of updated period exceeds the interval Δt so that the VM placement keeps pace with the demand changes.

We design the architecture of CAVMP as shown in Fig. 2. CAVMP has four key modules: resource monitor, utilization estimation, VM requestor, and strategy manager.

First, the *resource monitor* tracks the utilization (including CPU, memory, and storage) of all PMs and the real-time connectivity state.

Second, the historical and real-time utilization are fed into the *utilization estimator* to predict the future utilization of existing VMs in the next Δt time. The estimator predicts not only the expected demand, but also the volatility, which indicates the fluctuating degree of the demand.

Third, the *VM requestor* gets the new VM requests from users and translates these information into utilization request and link dependency.

Finally, the *strategy manager* takes the results from all other modules as inputs, and figures out the placement strategy scheme about where each VM should be placed. The final output is a binary matrix $X_{(t+1)} = [x_{mn}]$ for the placement in the next time slot.

In the following subsections, we introduce four components in details.

3.2 Resource monitor

The resource monitor collects the states of existing VMs in PMs, including three kinds of data: utilization state, connectivity state, and deployment matrix.

Utilization state. For a VM n during $[t, t + \Delta t)$, the utilization state is denoted by U_t^n , which includes the real-time computing and storage resources. The resource monitor collects the utilization every Δt and log all history.

Connectivity state. The resource monitor also collects the real-time antenna directions in order to imply the connectivity state, denoted by $C_t = [c_{ij}]$, where $c_{ij} = 1$ indicates that rack i and j are connected by 60GHz mmWave.

Deployment matrix. The resource monitor collects the real-time VM deployment in PMs, which is denoted by $X_t = [x_{mn}]$, where $x_{mn} = 1$ indicates that VM n is placed on PM m .

The transmission area of a 60GHz mmWave transceiver can be expressed as a cone model with the transmission range λ and the transmission angle α as shown in Fig. 1(a). With the transmission area, the locations of all antennas, and the height of ceiling, the connectivity state can be calculated.

3.3 Utilization Estimator

The resource consumption of existing VMs in PMs are dynamic. This module is used to estimate the future utilization. Conventional works usually exploit the average utilization or historical data as inputs to operate the VM placement scheme. However, we propose to predict not only the expectation of the future utilization, but also its variance as an fluctuation indicator. Benefitting from this variance, the quality assurance can be taken into account, whose goal is:

$$Pr\{L_m > R_m\} < \epsilon, \forall m \quad (1)$$

where ϵ is a small positive constant, R_m is the capacity of PM m , and L_m is the load of PM m . To achieve SLA, this Equation should be satisfied.

Based on the historical data and the real-time monitor results, we leverage the *AutoRegressive Integrated Moving Average* (ARIMA) model [3] to estimate the mean of future utilization and the *Generalized AutoRegressive Conditional Heteroskedasticity* (GARCH) model [2] to estimate the variance.

ARIMA-based mean estimation. Let U_t^n as the utilization of VM n during $[t, t + \Delta t)$ and L as the lag operator, where $LU_t^n = U_{t-1}^n$. Then, the lag difference ∇ is expressed by

$$\begin{cases} L^d U_t^n = U_{t-d}^n, \\ \nabla^d U_t^n = \nabla(\nabla^{d-1} U_t^n), \end{cases} \quad (2)$$

where $d \geq 1$ and $\nabla^0 = 1$.

In order to eliminate the impact of periodicity, we define $\nabla_j U_t^n = U_t^n - U_{t-j}^n$. Hence, the stationary time series of utilization can be formulated by ARIMA(p, q) model. Using the classic Box-Jenkins methodology, the values of p and q are easily to be determined. Thus, the mean and k -step-ahead prediction of U_t^n can be determined.

GARCH-based variance estimation. Besides the mean estimation, the variance estimation is significant to optimize the VM placement. We leverage the advanced GARCH(p, q) model to estimate the standard deviation by

$$Z_{i\tau} = \sqrt{h_{i\tau}} e_\tau, \quad (3)$$

and

$$h_{i\tau} = \gamma_{i0} + \sum_{j=1}^P \gamma_{ij} Z_{i\tau-j}^2 + \sum_{j=1}^Q \beta_{ij} h_{i\tau-j}, \quad (4)$$

where $e_\tau \sim \mathcal{N}(0, 1)$, $Z_{i\tau}$ is a zero-mean, Gaussian process, and $h_{i\tau}$ is the time-varying conditional variance. Based on training, parameters in Eq. (4) can be decided using method proposed in [3]. Thus, the GARCH process can forecast the conditional variance.

3.4 VM Requestor

In a cloud center, new requests of VMs are always proposed by users. These VMs need to be added into PMs. The VM requestor module delivers the new VMs information from users to the Strategy Manager module, including the number of new VMs and their dependency.

Number of new VMs. A user requires computing and storage resources from cloud center, which are described by several VMs. We assume that all VMs have the same utilization at the beginning. Thus, all new requests could be transformed into the number of new VMs.

Dependency. Partial VMs need to work together for one user's application, so links are required for these VMs, which is so-called dependency in this paper. The dependency is denoted by D , where D is a $n \times n$ matrix and $D_{ij} = 1$ implies that VM i and j is dependent.

3.5 Strategy Manager

To achieve high utilization without SLA violations, the strategy manager needs to match the future demand. In addition, to minimize the delay of antenna adjustment, the strategy manager needs to consider the connectivity state.

Three performance metrics are adopted to evaluate the effectiveness of the proposed strategy manager:

- Overload ratio V . This is the ratio of time periods where the reserved resource is lower than the actual demand over all the periods.
- Average utilization \bar{U} . This is the average utilization of the allocated resource over all the periods.
- Average delay of antenna adjustment \bar{T} . This is the average time consumed by adjusting the antennas' directions over all the periods.

Our objective is to keep a low overload ratio V while achieving a high average utilization \bar{U} and a low average adjustment time \bar{T} . To fulfill our goal, the resource monitor collects the utilization of each VM and the connectivity statement. The utilization estimator predicts the mean μ , the variance σ , the correlation ρ and their historical data. With μ, σ, ρ and the connectivity state C_t , the strategy manager determines the VM placement without violating SLA.

Based on the above analysis, we formulate the VM placement problem as follows:

$$\left\{ \begin{array}{l} \text{Objective} \quad \min (L_{max} - L_{min}) \\ \text{Subject to} \quad Pr(L_m > R_m) < \epsilon \\ \quad \sum_{i=1}^m \sum_{j=1}^n |c_{ij(t+1)} - c_{ijt}| < \epsilon \\ \quad x_{mn} \in \{0, 1\} \\ \quad \sum_{i=1}^n x_{mi} = 1 \end{array} \right. \quad (5)$$

where ϵ is a small positive value. Suppose that L follows a normal distribution. Hence, $Pr(L_m > R_m) < \epsilon$ can transform to

$$R_m \geq E[L_m] + c_p(\mu, \sigma^2) \sqrt{var[L_m]}, \quad (6)$$

where $c_p(\mu, \sigma^2)$ follows the normal distribution with mean μ and variance σ , $E[L_m]$ is the sum of expectations of utilization of all the hosted VMs and $var[L_m]$ is the variance of the workload.

Based on the problem formulation, we design our CAVMP algorithm using the modern portfolio theory as the following pseudo-code.

Algorithm 1: CAVMP Algorithm

Input: $C_t, X, \mu_t, \sigma, U, D$
Output: VM placement strategy

- 1 Sort PM by load in increasing order;
- 2 Sort new VM by utilization in decreasing order;
- 3 **foreach** new VM n **do**
- 4 **foreach** PM m **do**
- 5 **if** $\sum_{i=1}^m \sum_{j=1}^n |c_{ij(t+1)} - c_{ijt}| < \epsilon$ **then**
- 6 **if** $E[L_m] + c_p(\mu, \sigma^2) \sqrt{var[L_m]} < R_m$ **then**
- 7 Place new VM n to PM m ;
- 8 Adjust the position of PM m in PM list to ensure increasing order;
- 9 $R_m = R_m - E[L_m] - c_p(\mu, \sigma^2) \sqrt{var[L_m]}$;
- 10 $x_{mn} = 1$;
- 11 **break**;

The CAVMP algorithm places a VM with large U into the PM with small load that can hold the requirement of both SLA violation and connectivity state. Although CAVMP cannot guarantee an optimal solution due to multiple constraints but it is still efficient to find a better solution than existing solutions. Since VM placement problem is polynomial time many-to-one reducible to multiple knapsack problem, which is a known NP-optimization problem, the proposed



Fig. 3. The prototype of mmWave radios are equipped on the racks in cloud center.

CAVMP provides a greedy-like solution for this NP problem with tolerant approximation errors and time complexity $O(mn)$.

4 Performance Evaluation

In this section, we build a 2-rack prototype of 60GHz wireless cloud center and evaluate the time consumption of antenna adjustment. Based on the time consumption, we conduct simulations to evaluate the performance of CAVMP.

4.1 Prototype and Field Test

To measure the transmission feature of 60GHz wireless cloud center, we build a prototype of mmWave radio as shown in Fig. 3. This radio is supported by the liftable and rotatable cranks, so its height and direction could be arbitrarily adjusted by motor. The radio frontend consists of an mmWave transceiver to provide 4Gbps-bitrate transmission in 60GHz band and a customized cylinder metal waveguide as the antenna to form the signal into a beam. Then, the beam can be considered as the cone model with the angle α .

We conduct field test of a pair of such radios by HD video transmission as shown in Fig. 3. The transmission angle α is nearly 9° and the communication range is about 13m. Especially, we vary the distance between two radios with a step of 1m and the 4 Gbps communication link is maintained without obvious lag from 1 to 10m. The rotation speed is $60^\circ/\text{s}$ and the lift speed is $0.5\text{m}/\text{s}$.

4.2 Simulation

Then, we evaluate the proposed CAVMP algorithm by extensive simulations. We randomly generate 1000 utilization traces with different means and variations. All these utilization traces are not independent, thus their placement requires some communications among VMs. The connectivity state is also randomly generated in a 8×8 -rack cloud center. Each trace or state contains 500 historical data. We use the first 200 traces to train the time series model and the remaining data to test our algorithm and compare with existing algorithms.

In Fig. 4, we compare CAVMP with two VM placement algorithms Static and MPT about the cumulative distribution function (CDF) of utilization. Static is always provisioning for the peak. Static allocates resources based on historical data without any prediction. MPT is a state-of-the-art VM placement algorithm




Fig. 4. CDF of resource utilization.




Fig. 5. Delay of antenna adjustment.

with the correlation among VMs. However, it does not consider the directional transmission of mmWave antenna. As shown in the figure, Static has the most machines running with low utilization. Since Static is not predictive and dynamic, it cannot respond to time-varying demand well. The result of MPT shows that almost all the VMs is running with utilization larger than 80%. CAVMP performs better than Static algorithm and improves about 8% compared with MPT. This result demonstrates the proposed CAVMP can achieve a high resource utilization in 60GHz wireless cloud center.

The comparison on average delay of antenna adjustment is shown in Fig. 5. We find that CAVMP significantly outperform the other two algorithms, which is always close to zero due to the connectivity-aware design. In contrast, Static and MPT algorithms only consider the utilization and ignore the connectivity state, so their average delays are usually larger than 2s. Moreover, with the increase of size, the dependency of connectivity is more complicated. Hence, the trend of delay increases with the size.

5 Related Work

The related work can be classified into two categories:

Network performance prediction: The volatility prediction in bandwidth reservation is proposed in [10] to build an auto-scaling system that dynamically books the minimum bandwidth resources from multiple data centers for the VoD provider in [11]. In [16], the robust dynamic approach is designed to periodically identifies bandwidth allocation to virtual networks.

VM placement: Clark et al. [5] present the design, implementation, and evaluation of high-performance OS migration built on top of the Xen virtual machine monitor. A fast and transparent application migration system is proposed in [9]. And the traffic-aware virtual machine placement is studied in [7] to improve the network scalability.

However, the connectivity-aware solution is still a vacancy in VM placement.

6 Conclusion

In this paper, we investigate the connectivity-aware VM placement solution for 60GHz wireless cloud center. Time series prediction techniques and algorithm design are combined to provide better performance in minimizing the delay of antenna adjustment and maximizing the resource utilization. We build a prototype and evaluate CAVMP using extensive simulations. Performance results show that the effectiveness of CAVMP.

Acknowledgment: This research was supported by China 973 Project 2014CB340303, NSFC Grant No. 61672349, 61303202, 61672353, and 61472252.

References

1. Bobroff, N., Kochut, A., Beaty, K.: Dynamic placement of virtual machines for managing sla violations. In IFIP/IEEE IM, (2007).
2. Bollerslev, T.: Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, Vol. 31, No. 3, pp. 307–327, (1986).
3. Box, G. E., Jenkins, G. M., Reinsel, G. C., Ljung, G. M.: *Time series analysis: forecasting and control*. John Wiley & Sons, (2015).
4. Chen, M., Jin, H., Wen Y., Leung V.: *Enabling Technologies for Future Data Center Networking: A Primer*. *IEEE Network*, Vol. 27, No. 4, pp. 8–15, (2013).
5. Clark, C., Fraser, K., Hand, S., Hansen, J. G., Jul, E., Limpach, C., Warfield, A.: Live migration of virtual machines. In *USENIX NSDI*, (2005).
6. Halperin, D., Kandula, S., Padhye, J., Bahl, P., Wetherall, D.: Augmenting data center networks with multi-gigabit wireless links. In *ACM SIGCOMM Computer Communication Review*, Vol. 41, No. 4, pp. 38–49, (2011).
7. Jiang, J. W., Lan, T., Ha, S., Chen, M., Chiang, M.: Joint VM placement and routing for data center traffic engineering. In *IEEE INFOCOM*, (2012).
8. Meng, X., Isci, C., Kephart, J., Zhang, L., Bouillet, E., Pendarakis, D.: Efficient resource provisioning in compute clouds via vm multiplexing. In *ACM ICAC*, (2010).
9. Nelson, M., Lim, B. H., Hutchins, G.: Fast Transparent Migration for Virtual Machines. In *USENIX Annual*, (2005).
10. Niu, D., Li, B., Zhao, S.: Understanding demand volatility in large vod systems. In *ACM NOSSDAV*, (2011).
11. Niu, D., Xu, H., Li, B., Zhao, S.: Quality-assured cloud bandwidth auto-scaling for video-on-demand applications. In *IEEE INFOCOM*, (2012).
12. Popa, L., Kumar, G., Chowdhury, M., Krishnamurthy, A., Ratnasamy, S., Stoica, I.: FairCloud: sharing the network in cloud computing. In *ACM SIGCOMM*, (2012).
13. Popa, L., Yalagandula, P., Banerjee, S., Mogul, J. C., Turner, Y., Santos, J. R.: ElasticSwitch: practical work-conserving bandwidth guarantees for cloud computing. *ACM SIGCOMM Computer Communication Review*, Vol. 43, No. 4, pp. 351–362, (2013).
14. Verma, A., Dasgupta, G., Nayak, T. K., De, P., Kothari, R.: Server workload analysis for power minimization using consolidation. In *USENIX Annual*, (2009).
15. Wei, W., Wei, X., Chen, T., Gao, X., Chen, G.: Dynamic correlative VM placement for quality-assured cloud service. In *IEEE ICC*, (2013).
16. Zhang, M., Wu, C., Yang, Q., Jiang, M.: Robust dynamic bandwidth allocation method for virtual networks. In *IEEE ICC*, (2012).
17. Zhou, X., Zhang, Z., Zhu, Y., Li, Y., Kumar, S., Vahdat, A., Zheng, H.: Mirror mirror on the ceiling: flexible wireless links for data centers. *ACM SIGCOMM Computer Communication Review*, Vol. 42, No. 4, pp. 443–454, (2012).