

Joint Optimization of Server and Network Resource Utilization in Cloud Data Centers

Biyu Zhou¹, Jie Wu², Lin Wang³, Fa Zhang¹, and Zhiyong Liu¹

¹Institute of Computing Technology, Chinese Academy of Sciences

²Temple University ³Technische Universitat Darmstadt

Outline

- **Backgrounds**
- Problem
- Solutions
- Evaluation

Backgrounds

- Virtual machine placement
 - A key component of cloud resource management.
 - Improve resource utilization, reduce operational costs.
- Challenges
 - **Jointly** optimize the **server** and **network** resource utilization.
 - Design **fast & performance-guaranteed** algorithms.
- Our focus
 - Design efficient algorithms to maximize the **overall** resource utilization in multiple dimensions without violating resource capacity constraints.

Outline

- Background
- Problem
- Solutions
- Evaluation

Abstracting the problem

- Problem
 - Given a set of VMs, how to pack VMs into servers such that **overall** resource utilization is **maximized** while no constraints on resources are violated.
- Definition of size
 - The size of a VM is the **product** of the resource demands on all its resource dimensions.
 - Portray the level of load along multiple dimensions in a **unified** manner.
- Solving the problem is **NP-hard**.

Outline

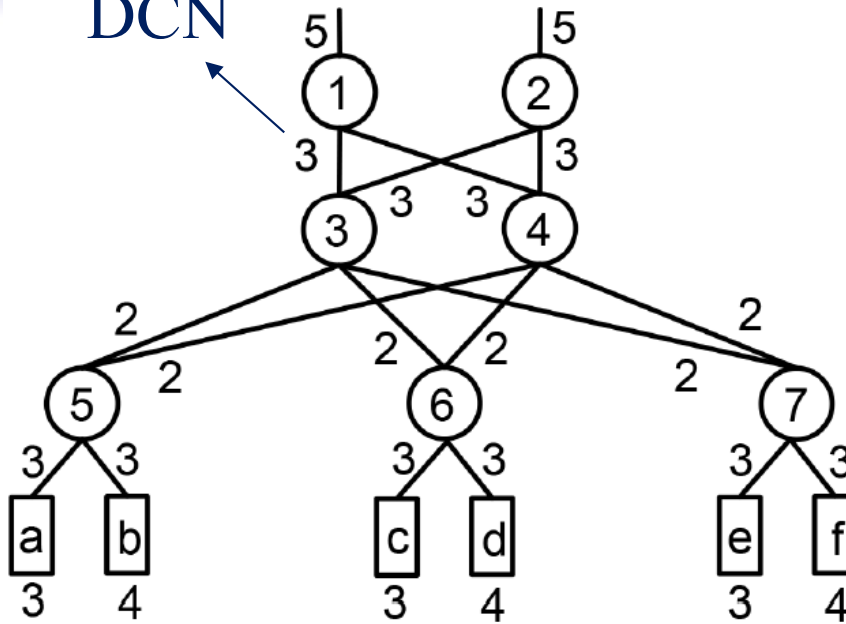
- Background
- Problem
- Solutions
- Evaluation

Analysis

- Main difficulty
 - Networks are oversubscribed.
 - Bandwidth is shared among servers.
 - Server & network.
- Total resource capacity
 - The sum of the possible size of all its servers.
 - An upper bound.



Motivation solutions

DCN



Solution 1:
 $(3*3+4*1)*2$
 $+3*2=32$

Solution 2:
 $(3*1+4*3)*2$
 $+4*2=38$



Analysis

- Observation
 - An efficient solution will allocate the server that has a larger CPU capacity with a larger bandwidth capacity.
- A two-staged solution
 - Bandwidth allocation using **min-cost max-flow algorithm**.
 - **Two-dimensional item packing**.

Solutions

- Offline scenario
 - For m identical servers, and an accommodating input, there exists an offline algorithm that achieves $1/4$ approximation in linear time.
 - For an accommodating input, there exists an offline algorithm that achieves $1/4r$ approximation in linear time.

Solutions

- Online scenario
 - SortFirstFit: always places the arriving VM into the server with maximum CPU capacity which has enough residual CPU capacity and bandwidth
 - SortWorstFit: always places the arriving VM into the server with maximum residual size.
- The maximum possible bandwidth in both algorithms are computed by solving a min-cost max-flow problem.

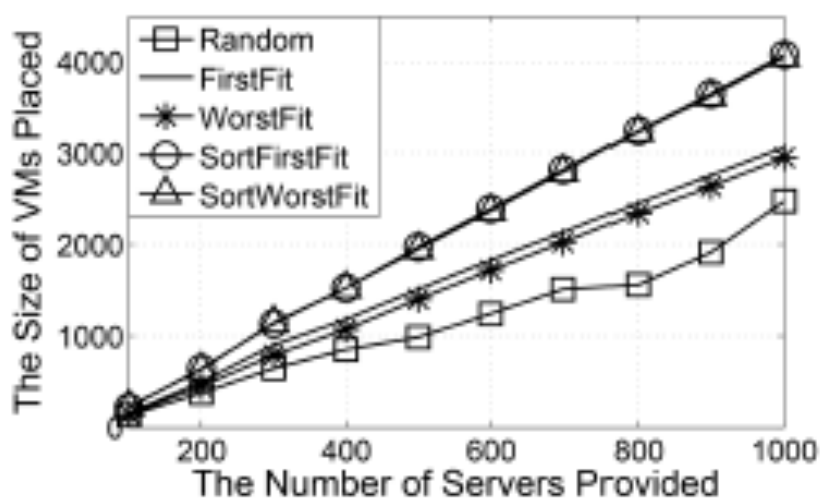
Outline

- Background
- Problem
- Solutions
- Evaluation


The evaluation setting

- Benchmarks
 - FirstFit (as less server as possible)
 - WorstFit (as balance as possible)
 - Random
- Metric
 - The size of VMs placed.

The impact of arriving VM sequence



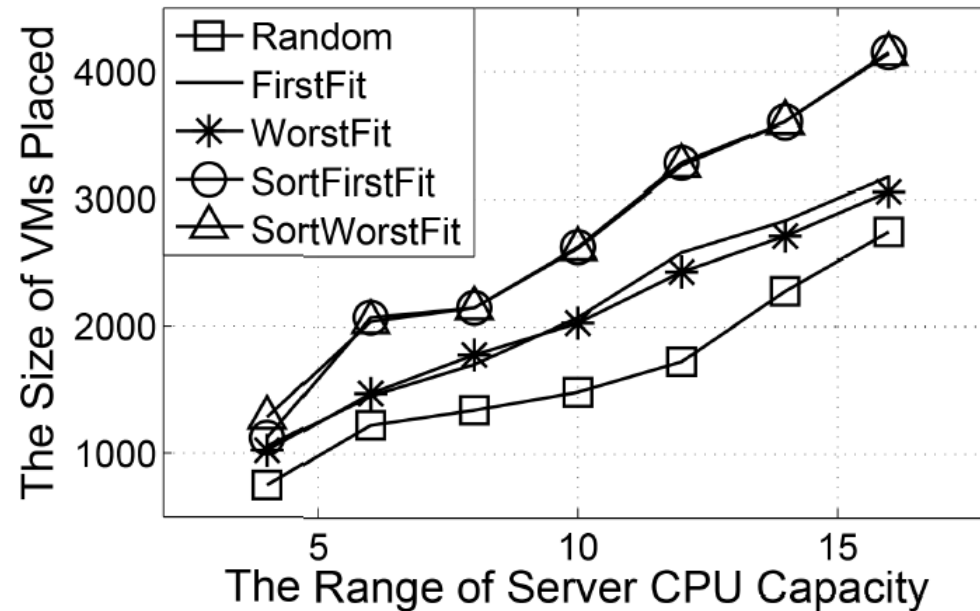
(a)



(b)

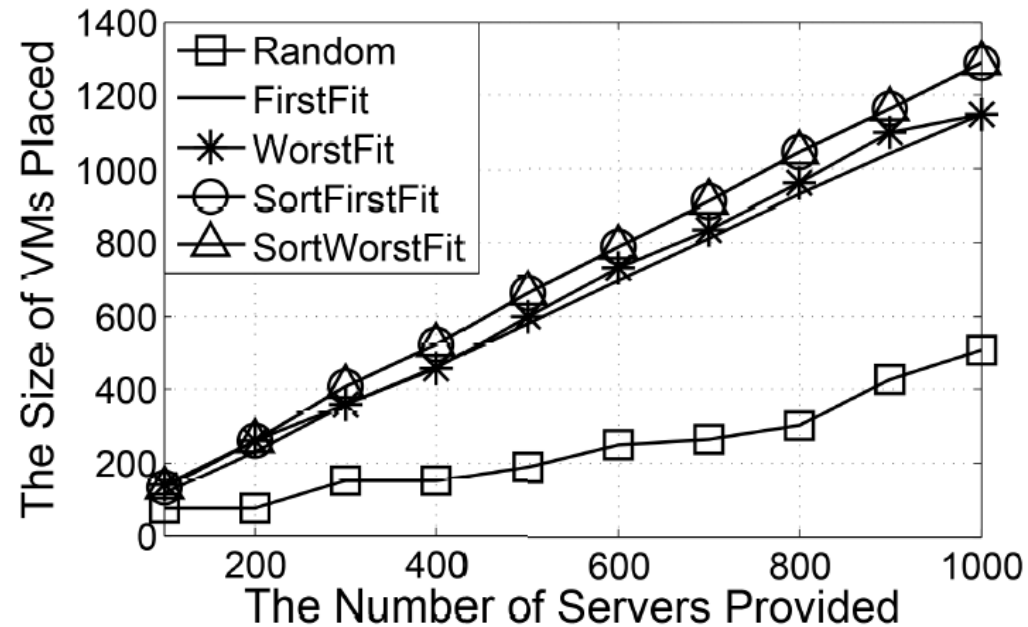
- (a) our heuristics works much better than other heuristics under both the arriving VM sequence that consists of a sequence of small VMs followed by a sequence of large VMs and the arriving VM sequence that consists of a sequence of large VMs followed by a sequence of small VMs.
- (b) we can conclude that processing sorting before placement may make the placement more efficient.

The impact of hardware heterogeneity



- (a) the advantages of sorting will be more significant when the difference of resources of servers in the network are larger.

Testing with real traces



- The results in this situation verify that the performances of the five heuristics using trace-driven data sets are in line with those using the above three generated data sets.

Thanks for your attention!