Energy-Efficient Virtual Machine Replication and Placement in a Cloud Computing System

Hadi Goudarzi and Massoud Pedram

Presented by: Payman Khani

Overview:

- INTRODUCTION
- SYSTEM MODEL
- PROBLEM FORMULATION
- PORPOSED ALGORITHM
- SIMULATION RESULTS
- CONCLUSION
- FUTURE WORK

INTRODUCTION

- By utilizing Virtual Machines (VM) and doing server consolidation in a datacenter, a cloud provider can reduce the total energy consumption for servicing his clients with little performance degradation.
- Placing multiple copies of a VM on different servers and distributing the incoming requests among these VM copies can reduce the resource requirement for each VM copy and help the cloud provider utilize the servers more efficiently.

INTRODUCTION

- **Server consolidation:**
  Enables the assignment of multiple virtual machines (VMs) to a single physical server. By this action, some of the available servers can be turned off or put into some deep sleep state, thereby, lowering power consumption of the computing system.
- Modern servers tend to consume 50% or so of their peak power in idle state.
- Consolidation involves performance-power tradeoff.
- The IT infrastructure provided by the datacenter owners/operators must meet various Service Level Agreements (SLAs) established with the clients.

INTRODUCTION

- SLAs :
  - ✓ Resource related (e.g., amount of computing power, memory/storage space, network bandwidth).
  - ✓ performance related (e.g., service time or throughput).
  - ✓ Quality of service(Qos) related (24-7 availability, data security, percentage of dropped requests.)

- To minimize the energy consumption using consolidation, these SLA constraints should be considered.

SYSTEM MODEL

➢ Assumptions and system configuration:

Servers of a given type are modeled by:

- ✓ Processing capacity = CPU cycle
- ✓ Memory BW= The rate that data can read or store into memory by processor.
- ✓ Energy cost

Symbol name	***Definition***
c_i^m, c_i^p	Required memory BW and total processing capacity for the i[th] client
L_i	Max. # of servers allowed to serve the i[th] client
s_k	Set of servers of type k
C_j^p, C_j^m	Total CPU cycle and memory BW of the j[th] server, shorthand notation for $C_{S_k}^p$ and $C_{S_k}^m$
P_j^0	Constant power consumption of the j[th] server operation in the active mode. Same as $P_{S_k}^0$
P_j^p	Power of operating the j[th] server which is proportional to the utilization of processing resources, shorthand notation for $P_{S_k}^p$
T_e	Duration of a *decision epoch* in seconds
x_j	A pseudo-Boolean integer to determine if the j[th] server is ON (1) or OFF (0)
y_{ij}	A pseudo-Boolean integer to determine if the i[th] VM is assigned to the j[th] server (1) or not (0)
ϕ_{ij}^p, ϕ_{ij}^m	Portion of the processing and memory BW resources of the j[th] server that is allocated to the i[th] client
ϕ_j^p, ϕ_j^m	Portion of the processing and memory BW resources of the j[th] server that is allocated to any client

SYSTEM MODEL

- Energy cost = P * T
- P= P_*^0 + P_*^p (utilization of the server)
- If multiple copies of a VM are placed on different servers, the following constraints should be satisfied:

1) $\sum_j \phi_{ij}^p C_j^p = c_i^p$
2) $\phi_{ij}^m y_{ij} C_j^m = c_i^m$

- Constraint (1) enforces the summation of the reserved CPU cycles on the assigned servers to be equal to the required CPU cycles for client i.
- Constraint (2) enforces the provided memory BW on assigned servers to be equal to the required memory BW for the original VM.
- **This constraint enforces the cloud provider not to sacrifice the Quality of Service (QoS) for its clients.**

SYSTEM MODEL

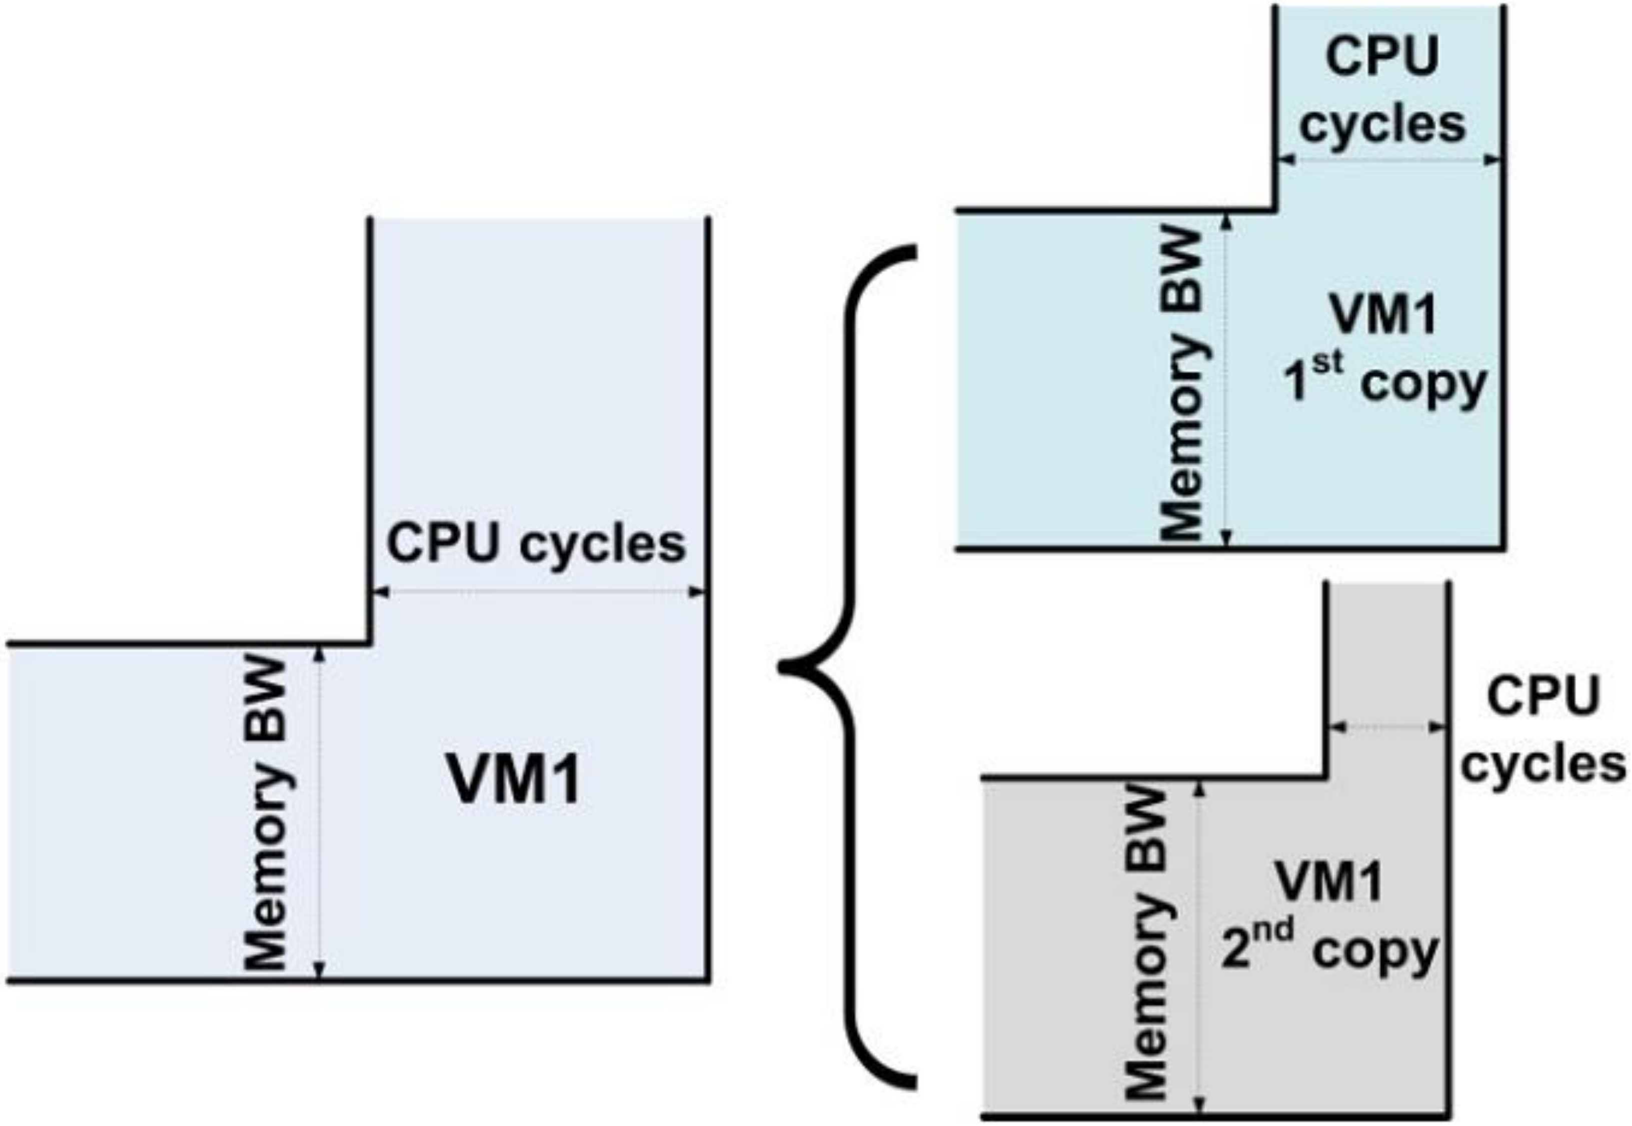

SYSTEM MODEL

- VM controller (VMC) : responsible for determining the resource requirements of the VMs and placing them on servers.
- The VMC performs these tasks based on two different optimization procedures:
  - ✓ Dynamic optimization: performs whenever it is needed.
  - ✓ Semi-static optimization: performs periodically (at periods of Te).

- The role of the semi-static optimization procedure in the VMC is to determine whether to create multiple copies of VMs on different servers and assign VMs to servers.
- The goal of this optimization is to minimize the energy cost of the active servers in datacenter.

SYSTEM MODEL

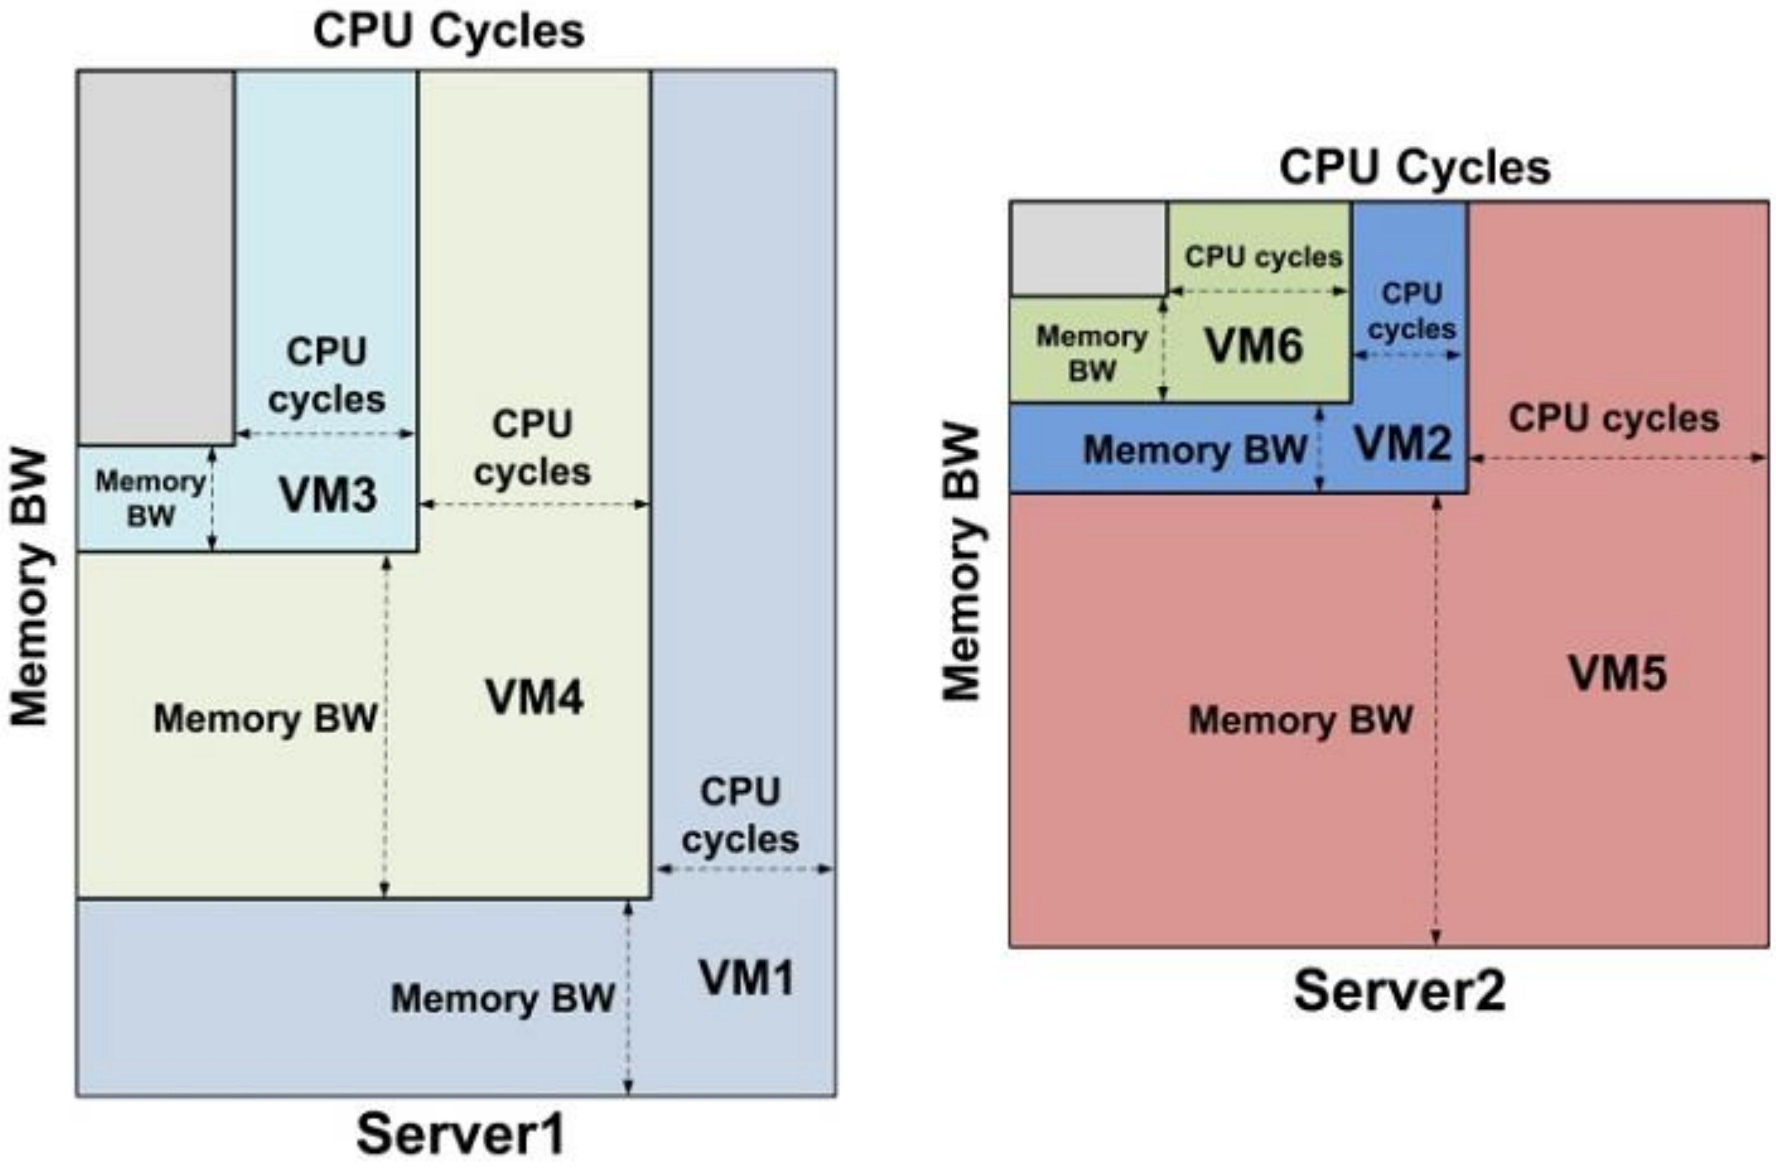

PROBLEM FORMULATION

- The objective function is the summation of the energy cost of the ON servers based on a fixed power factor and a variable power term linearly related to the server utilization.
- MERA for Multi-dimensional Energy-efficient Resource Allocation

$$Min \quad T_e \sum_j x_j \left(P_j^0 + P_j^p \sum_i \phi_{ij}^p \right)$$

PROBLEM FORMULATION

- subject to:

$$y_{ij} \in \{0,1\},\ x_j \in \{0,1\},\ \phi_{ij}^p \geq 0,\ \phi_{ij}^m \geq 0$$

$$\phi_j^p = \sum_i \phi_{ij}^p \leq 1$$

$$\phi_j^m = \sum_i \phi_{ij}^m \leq 1$$

$$\sum_j C_j^p \phi_{ij}^p = c_i^p$$

$$\phi_{ij}^m y_{ij} C_j^m = c_i^m$$

$$\sum_i y_{ij} \leq L_i$$

$$x_j \geq \sum_i \phi_{ij}^p$$

$$y_{ij} \geq \phi_{ij}^p$$

PORPOSED ALGORITHM

- Energy-efficient VM Replication and Placement algorithm- EVRP
- Clients are ordered (non-increasing) based on their processing requirement.
- Based on this ordering, the optimal numbers of copies of the VMs are determined and these copies are placed on servers using dynamic programming.
- local search method: servers are turned off based on their utilization and VMs are placed on the rest of the servers (if possible) to minimize the energy consumption as much as possible.

PORPOSED ALGORITHM

- **Energy Efficient VM Placement Algorithm:**
- ✓ φ_j^p and φ_{j}^m for each server are set to zero.
- ✓ For each VM, a method based on DP is used to determine the number of copies placed on different servers.
- ✓ Energy cost of assigning a copy of the i_{th} VM to a server from server type k is calculated based on equation:

$$c_{ik}(\alpha) = \phi_{ij}^p P_j^p + P_j^0 c_i^m / C_j^m$$

- ✓ where α(between 1 and Li) denotes the size of the assigned VM to the server. φ_{ij}^p is calculated from equation:

$$\phi_{ij}^p = (\alpha c_i^p / L_i) / C_j^p$$

PORPOSED ALGORITHM

$$\phi_{ij}^{p} = (\alpha c_i^p / L_i)/C_j^p$$

- For example, in case of Li=4 if half of the CPU cycle requirement of the VM is provided by a copy of the VM, α is equal to 2 and φ_{ij}^{p} is equal to $0.5 c_i^p / C_j^p$.

$$c_{ik}(\alpha) = \phi_{ij}^{p} P_j^p + P_j^0 c_i^m / C_j^m$$

- The first term is the cost related to the CPU utilization of the server.
- The second term is the replacement of the constant energy cost of the active server.
- For each VM, this equation is calculated for each server type and different values of α(between 1 and Li).
- Moreover for each server type, Li active servers and Li inactive servers that can service at least the smallest copy of the VM are selected as candidate hosts.
- For active servers, the value of cost is decremented by ε to select them over inactive servers in an equal energy scenario.

PORPOSED ALGORITHM

- After calculating cost for each possible assignment, the problem is reduced to

$$Min \ \sum_{j \in P} y_{ij}^{\alpha} c_{ij}(\alpha)$$

- Subject to:

$$\sum_{j \in P} \alpha y_{ij}^{\alpha} = L_i$$

- Where y_{ij}^{α} denotes the assignment parameter for j_{th} server with VM with size of α(1 if assigned and 0 otherwise).
- Moreover, *P* denotes the set of candidate servers for this assignment.

PORPOSED ALGORITHM

- After finding the assignment solution, φ_j^p and φ_j^m of the selected servers are updated. Then, the next VM is chosen and this procedure is repeated until all VMs are placed.
- **Local Search method:**
  - ✓ To improve the results of the proposed VM placement algorithm.
  - ✓ To minimize the total energy consumption in the system, all servers with utilization less than a threshold are examined.
  - ✓ Utilization of a server is defined as the maximum resource utilization in different resource dimensions in the server.
  - ✓ To examine these under-utilized servers, each of them is turned off one by one and total energy consumption is found by placing their VMs on other active servers using the proposed DP placement method.

SIMULATION RESULTS

- min Power Parity (mPP):Based on first fit
- EVRP(Energy-efficient VM Replication and Placement algorithm)-Li = 5
- Baseline: EVRP – Li=1

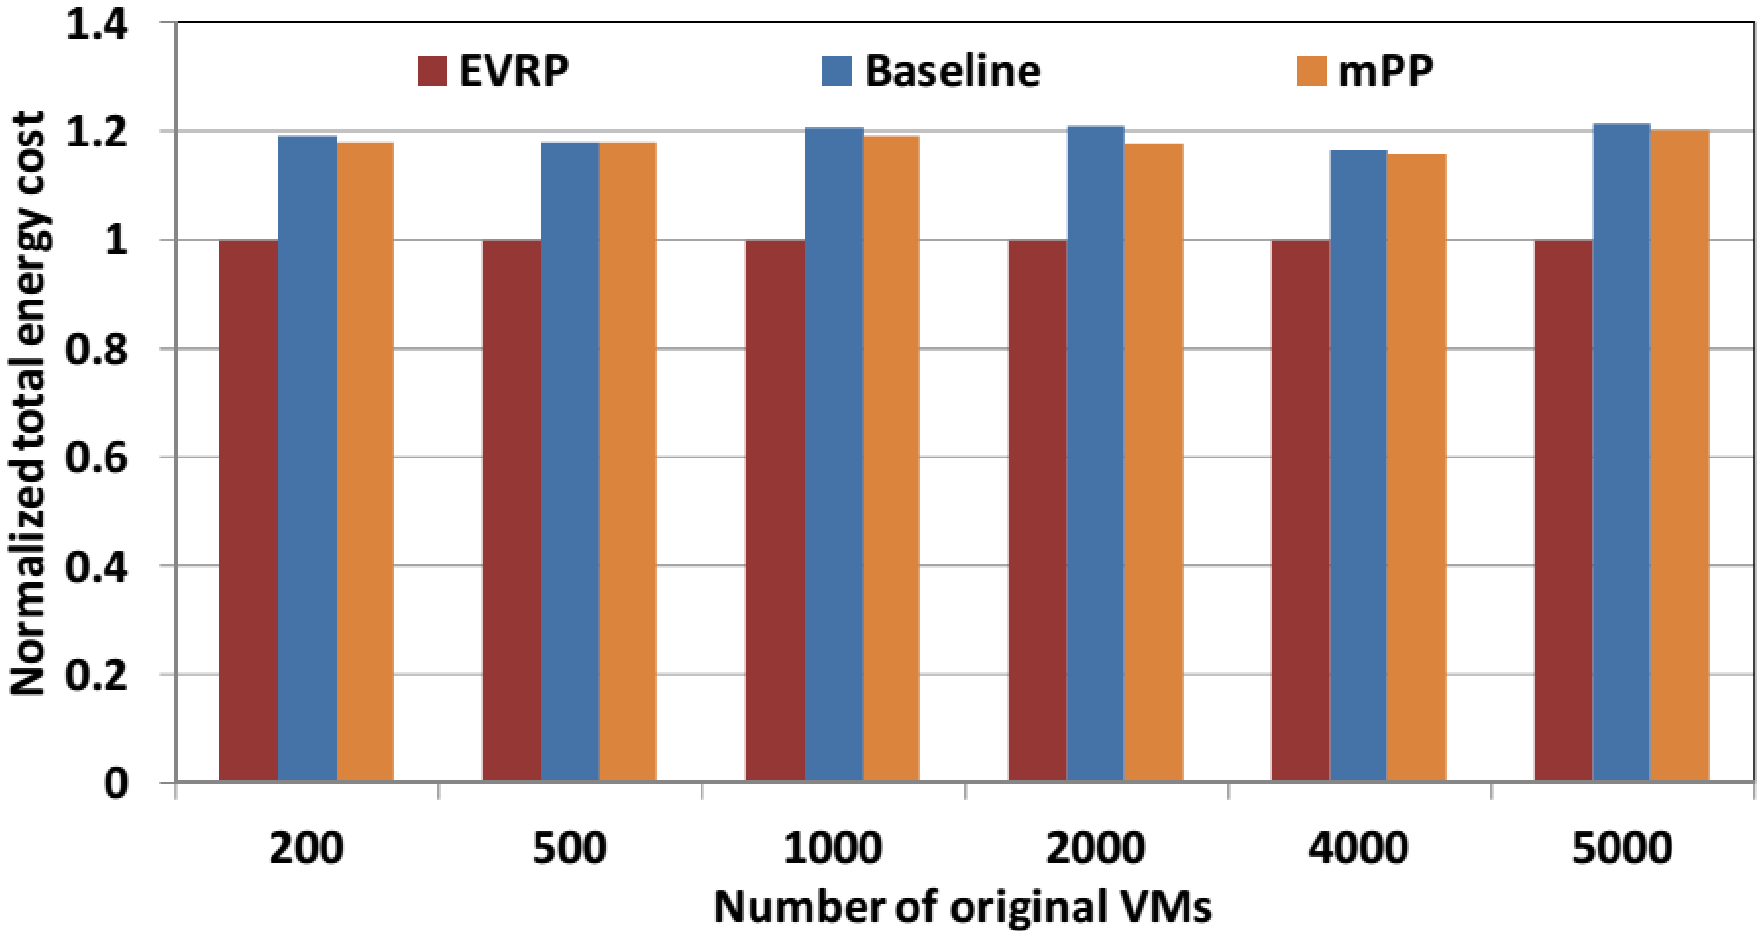

SIMULATION RESULTS

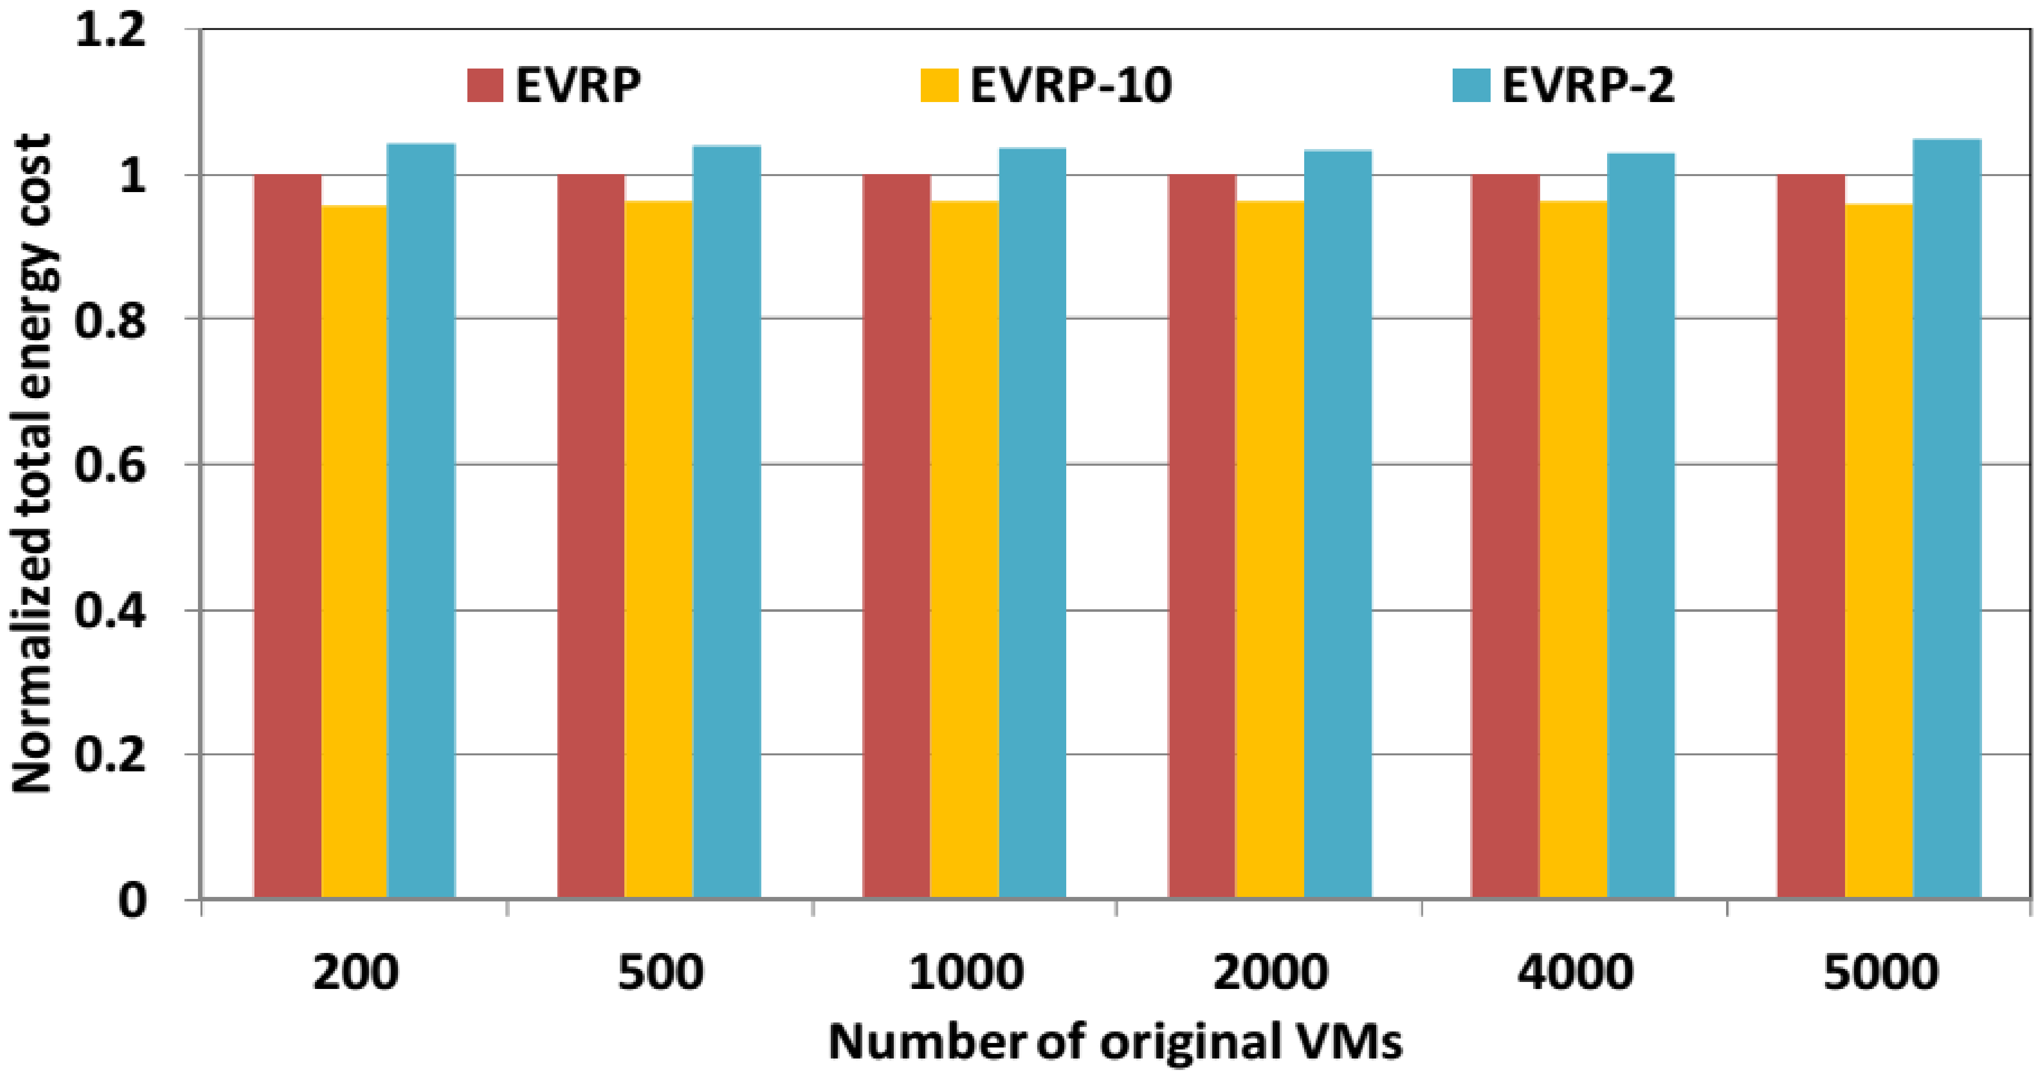

SIMULATION RESULTS

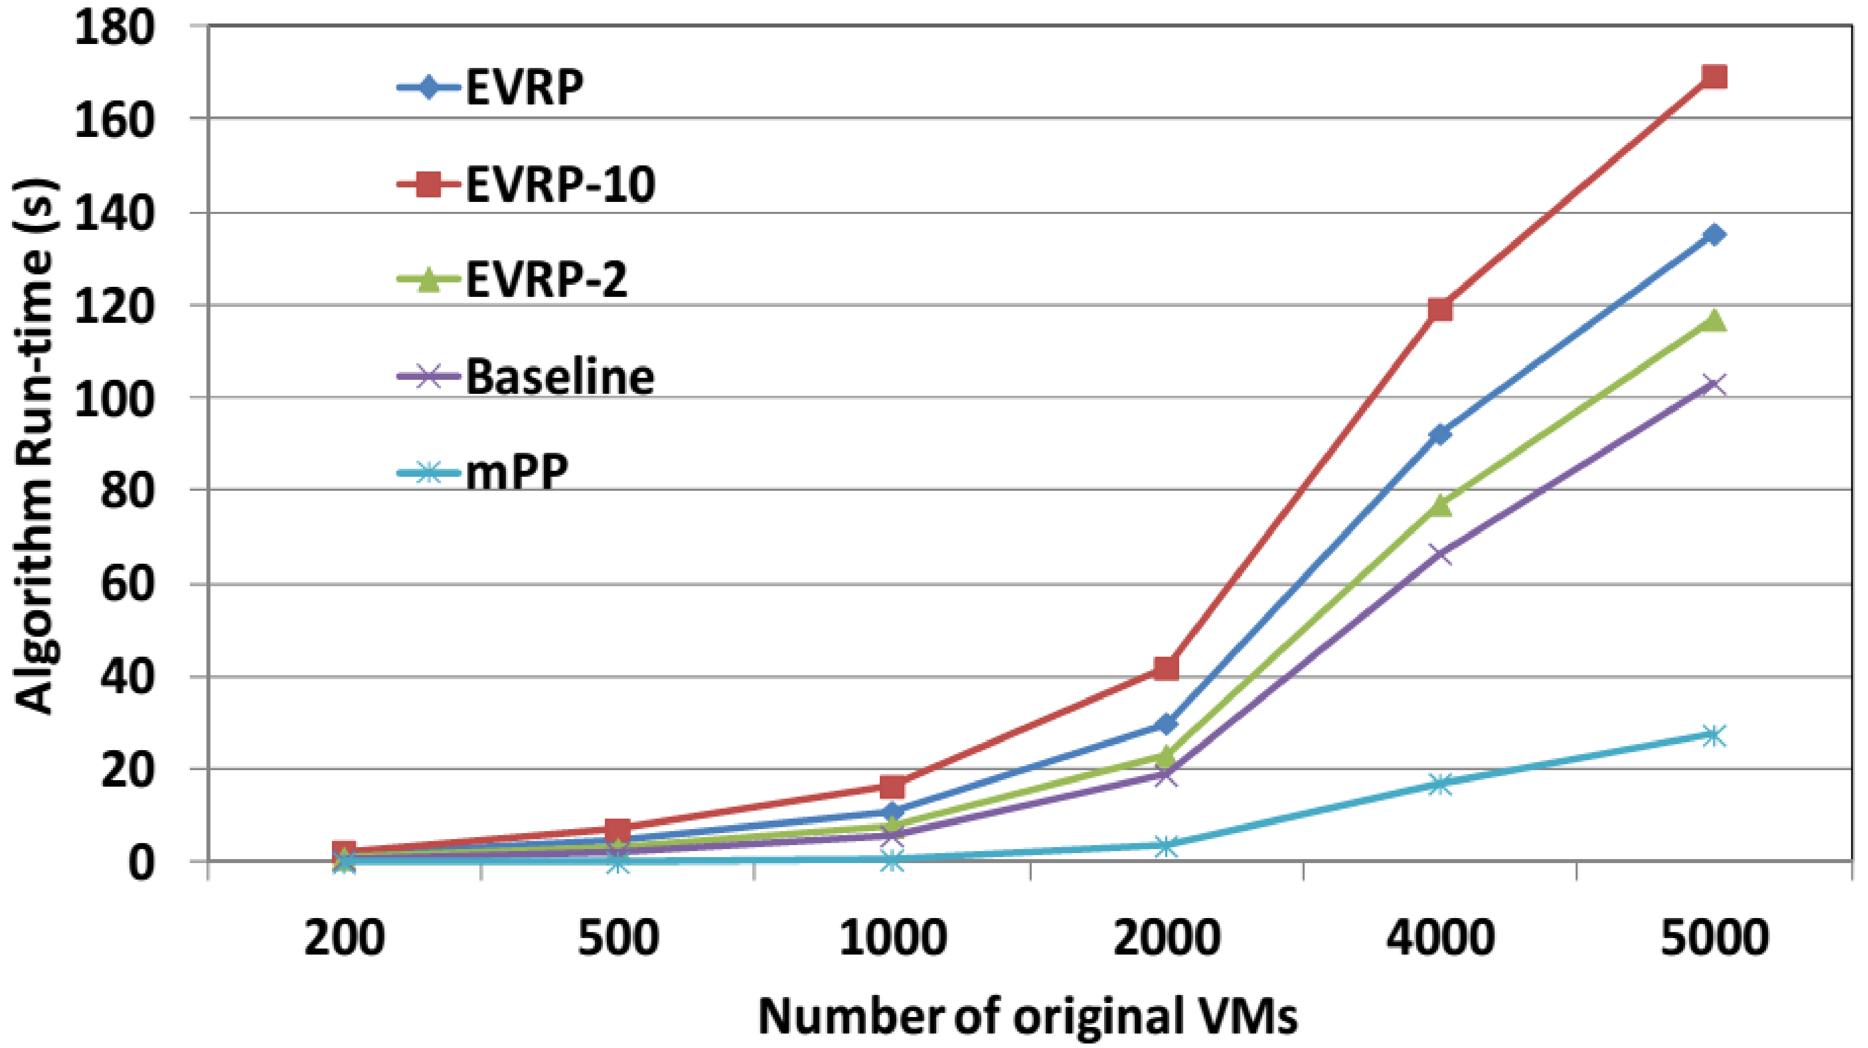

CONCLUSION

- Using this approach we generate multiple copies of VMs without sacrificing the QoS.(fixed BW & Li)
- An algorithm based on dynamic programming and local search was provided to determine the number of VM copies, and then place them on the servers to minimize the total energy cost in the cloud computing system.
- This approach reduces the energy cost by up to 20% with respect to prior VM placement techniques.

THANKS FOR YOUR ATTENTION