

# Distributed QoS-Aware Virtual Machine Scheduling in Support of Cloud Computing Environments

Hadi Salimi  
Ph.D. Thesis Proposal,  
Distributed Systems Laboratory,  
School of Computer Engineering,  
Iran University of Science and Technology,  
[hsalimi@iust.ac.ir](mailto:hsalimi@iust.ac.ir)

Supervisor: Dr. Mohsen Sharifi  
October 2009

## Agenda



- Virtualization Technology (VT)
- Cloud Computing Environment
  - Cloud Definitions
  - Cloud and Grid
  - Cloud Challenges
- VM Scheduling
- Related Works
- Proposed Research Topic
- Project Scheduling
- References

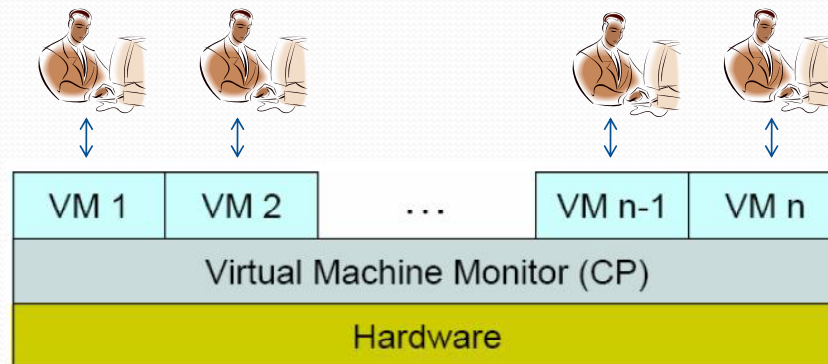
# Virtualization Technology (VT)



- An old technology from late 60's




- Was first coined by IBM [GOL74] to multiplex the power of mainframes



## VT (Cont.)



- Was dormant for decades because of its overhead
- Has become active after recent advanced in hardware [OHL05] and software technologies




- Two main concepts:
  - Virtual Machine (VM)
  - Virtual Machine Monitor (VMM)

# Basic Concepts



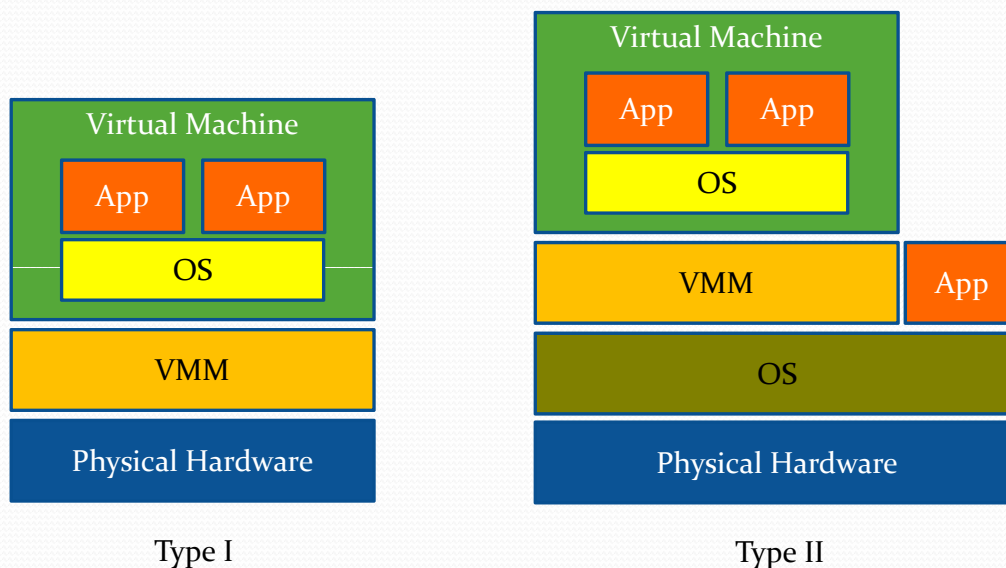
- Virtualization basic concepts [GOL73]:
  - Virtual Machine
  - Virtual Machine Monitor



# VT Categories



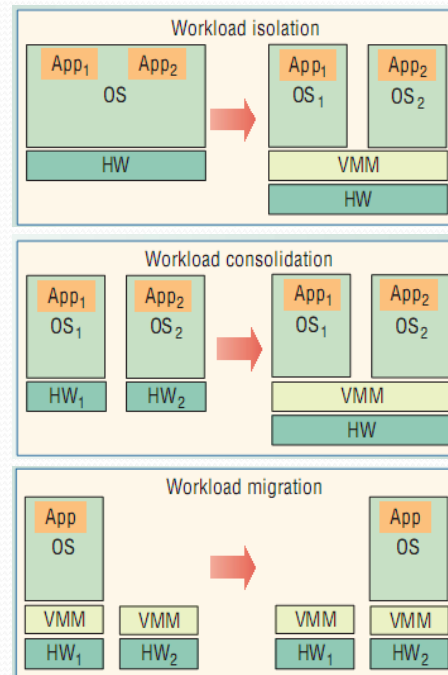
- Types of VMM [GOL74]:



# Advantages of VT



- Three main advantages of VT [UHL05]:
  - Isolation
  - Consolidation
  - Migration



# VT Applications



- Dependability
  - Checkpointing [THE08]
  - System replication [CUL08]
- Security
  - Web testing [YU08]
  - Honey-potting [PRO04]
- Software engineering
  - Software Testing [KOJ05]
  - Deployment [ALP05]

# Using VT on Grids



- VT benefits for grid computing support [FIGo3]:
  - Customization
  - Legacy Support
  - Resource Control
  - High Privileges
- And then, the Cloud buzzword became popular

# Grid and Cloud



Cloud Computing and Grid Computing trends measured by Google

# From Grid to Clouds



- More than twenty different definitions for clouds [TWEo8]
- Grids and Clouds share similar visions [VAQo9] :
  - Reduce computing costs
  - Efficiency
  - Reliability
- Some consider them as the same [DORo9]

# A Cloud Definition



- From our point of view Cloud Computing is:

A scalable distributed computing environment in which a large set of virtualized computing resources, different infrastructures, various development platforms and useful softwares are delivered as a service to customers as a pay-as-you-go manner usually over the Internet.

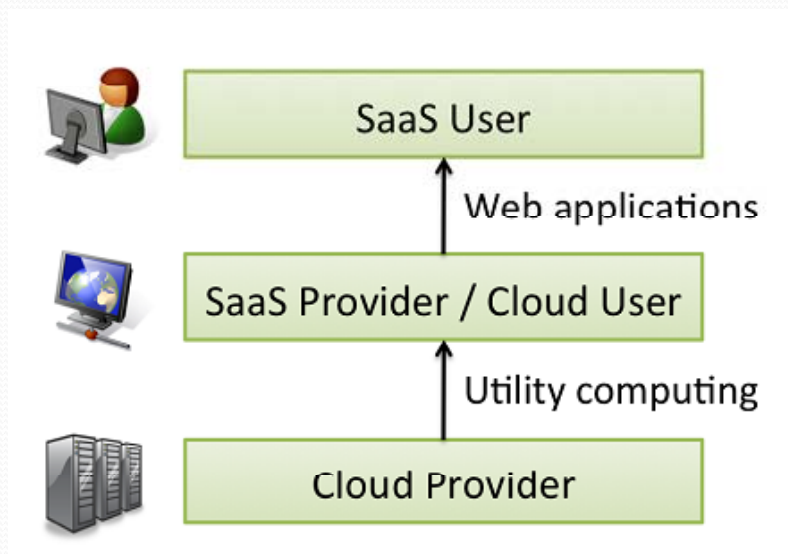
- It is based on a definition on [FOSo8]

# Cloud Main Features



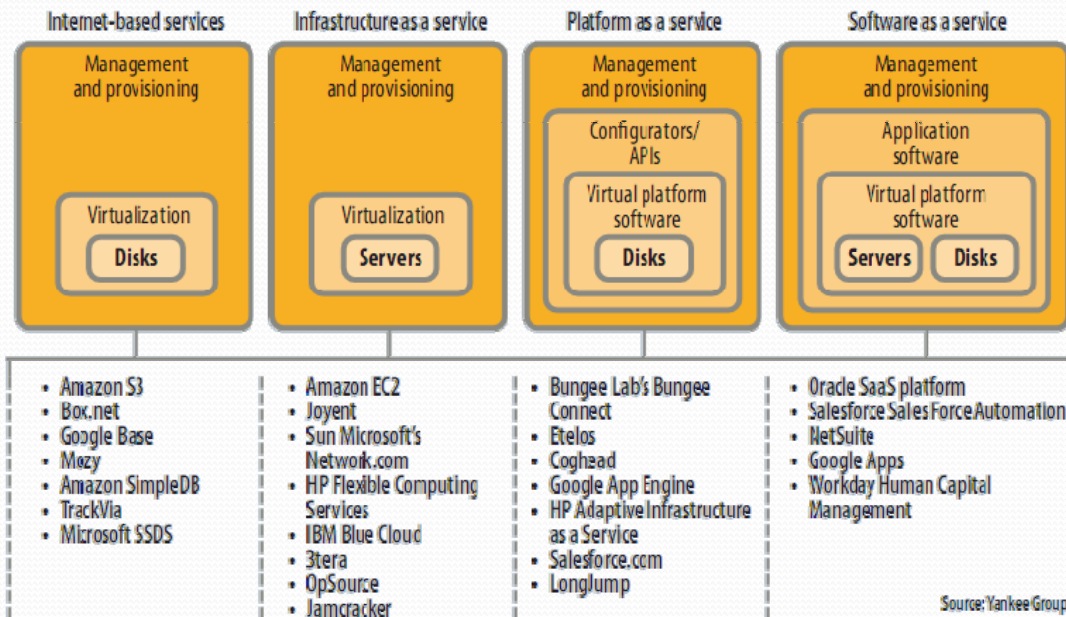
- Main features of a Cloud environment include [GRO09]:
  - Virtualization
  - Scalability
  - Deferent levels of service
    - Infrastructure
    - Platform
    - Software
  - Accessible through web services

# Cloud Architecture



Adopted From [ARM09]

# Cloud Services



Adopted From [LEA09]

# Cloud Challenges



- Main Cloud Computing challenges include [ARM09]:
  - Security
  - Resource Management and Scheduling
  - Automatic Scalability
  - Data Transmission
  - Software Licenses

# VM Scheduling




- Like any other processing unit, VMs need to be scheduled on the cloud in order to:
  - Maximize utilization
  - Do the job faster
  - Consume less energy
- It has its own **benefits** and **challenges**

# VM Scheduling Benefits




- VM scheduling benefits include:

- Easy resource reservation



- VM elasticity

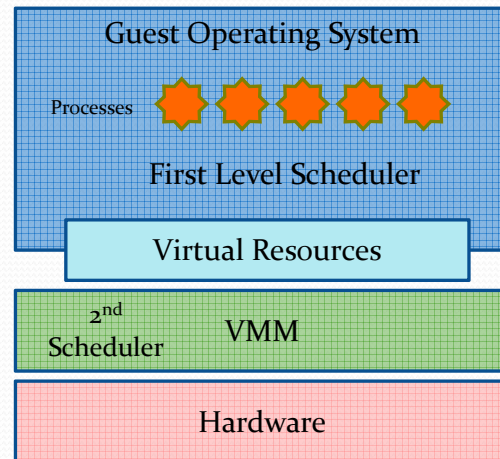


# VM Scheduling Challenges



- VM Scheduling challenges include:

- Two level scheduling
- High level of abstraction
- Unpredictable behavior



# Related Work



- The related works for VM Scheduling can be categorized into:
  - Scheduling VMs on a single system
  - Scheduling VMs on a distributed system

## Related Work (cont.)






- Scheduling VMs on a Single System:
  - Current VMMs like Xen [BAR03]
  - Task-aware VM scheduling [KIM09]
  - User driven VM scheduling [LIN04]
  - Load based VM scheduling [WEN09]
- But, resource allocation among VMs is **poorly understood** [CHE07]

## Related Work (cont.)



- Distributed scheduling of VMs:
  - Open Nebula [BAL09]
  - Eucalyptus [NUR09]
  - Nimbus [KEA08]
  - SnowFlock [LAG09]
  - DVFS-enabled clusters [LAZ09]



# Thesis Proposal

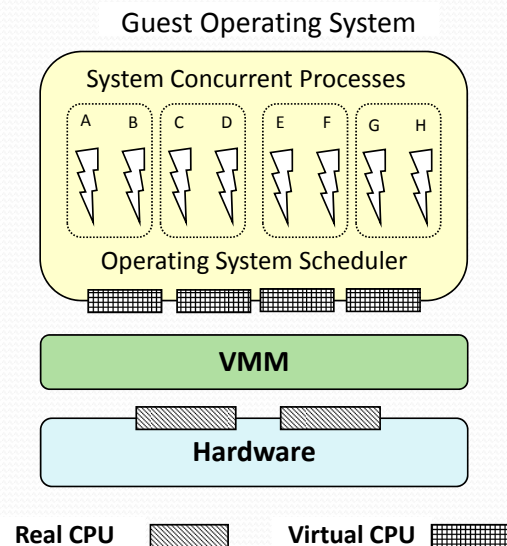


- Virtual Resource Co-Scheduling
- QoS-Aware Distributed VM Scheduling
  - Energy Efficient
  - Utilization Efficient
  - Cost Efficient

# Resource Co-Scheduling



- An example:
  - Four virtual processors
  - Two physical processors
- Which two processors should be co-scheduled?
- How to determine co-working processors?



# Distributed VM Scheduling



- Techniques to schedule VMs on Cloud:
  - Power efficient
  - Utilization efficient
  - Cost efficient -> VM co-location
  
- Building test-beds and simulators for this purpose

# Tools Needed

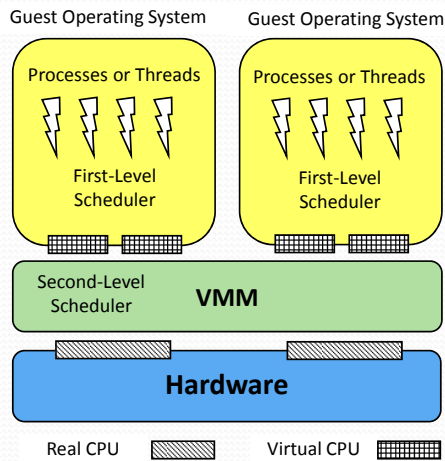


- Cloud Test bed
  - Based on current test beds such as Eucalyptus [NUR09] or using current VMMs like Xen [BAR03] or VMWare [VMW08]
  - Enable to implement different scheduling algorithms on it
  - Customized operating systems for its VMs
  
- Simulator
  - Based on current simulators, such as CloudSim [BUY09] or developing a new one
  - Capable of measuring different QoS aspects such as power consumption

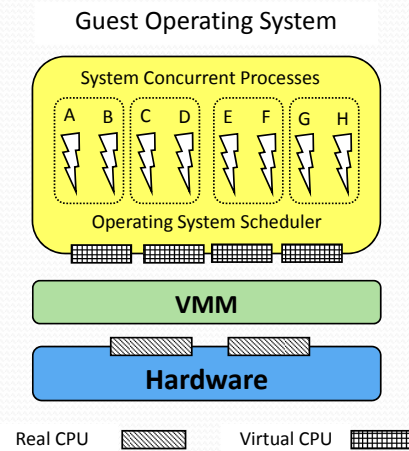
# Works Done



- As a case study, we simulated the effect of virtual processor co-scheduling on an extension of CloudSim [BUY09].

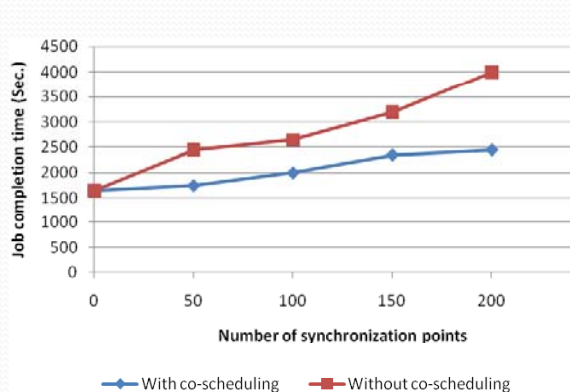


Case Study I

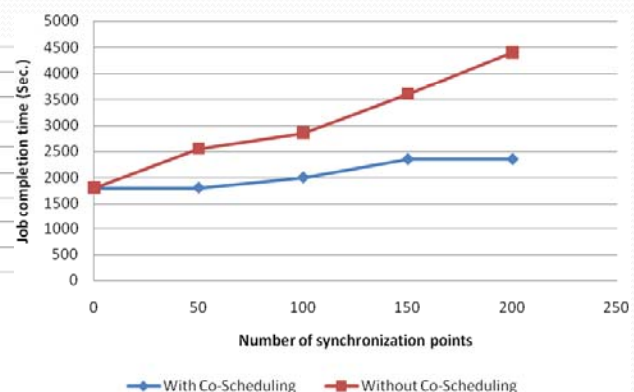


Case Study II

# Evaluation Results



Results of case study I



Results of case study II

H. Salimi, M. Najafzadeh, M. Sharifi, "A Case of Virtual Processor Co-Scheduling on Cloud", Submitted to the Cloud Computing Track of 22<sup>nd</sup> ACM Applied Computing Conference, Sierra, Switzerland, 2010.

# Expected Innovations



- Filling the semantic gap between VMM and VM to schedule resources better
- Distributed VM scheduling algorithms to attain different aspects of QoS such as power consumption
- Design and implementing of simulators and test-beds for evaluation of mentioned algorithms
- Proposing an architecture for QoS-Aware cloud operating systems

# Questions to be Answered



- How filling the semantic gap between VMM and VM could help schedule better?
- Is there any relation between scheduling algorithms used by guest operating system and VMM?
- Which class of distributed scheduling algorithms performs better for scheduling VMs?
- What are the techniques to consume less power energy in data centers?
- What are the main characteristics of an operating system used as an execution environment on cloud?

# Project Scheduling



Activity	Month											
	1	2	3	4	5	6	7	8	9	10	11	12
Implementing VM Co-Scheduling	x	x	X									
Estimation of Resource Dependency			x	X								
Preparing Conference Paper				x	X							
Investigation of Distributed Scheduling Algorithms					x							
Choosing a good implementation test-bed						X						
Deployment of test-bed environment						x	X					
Implementation of Algorithms							x	x	X			
Preparing Journal paper								X	x			
Implementing QoS Aware Algorithms										X	x	
Implementing the Simulator										x	X	x

# Project Scheduling (Cont.)



Activity	Month											
	1	2	3	4	5	6	7	8	9	10	11	12
Evaluation of Algorithms Using Simulator	x	x										
Preparing Conference Papers		x										
Searching for a light node operating system			x	x								
Customization of the so-called operating system				x	x	x	x					
Preparing the journal paper							x					
Preparing thesis								x	x			
Preparing for pre-defense									x	x		
Defending form thesis											x	x

# References



- [DORo8] T. Dorksens, "Cloud Computing – The User-Friendly Version of Grid Computing", *Cloud Computing Journal*, August 2008. <http://cloudcomputing.sys-con.com/node/593313>.
- [TWEo8] "Twenty Experts Define Cloud Computing", SYS-CON Media Inc, [http://cloudcomputing.sys-con.com/read/612375\\_p.htm](http://cloudcomputing.sys-con.com/read/612375_p.htm), 2008.
- [Figo3] R. J. Figueiredo, P. A. Dinda, and A. B. Fortes, "A Case For Grid Computing On Virtual Machines," 23<sup>rd</sup> International Conference on Distributed Computing Systems, Rhode Island, USA, 2003.
- [THEo8] T. Thein, M. Pokharel, S. Chi, J. Park, "A Recovery Model for Survivable Distributed Systems through the use of Virtualization", Fourth International Conference on Networked Computing and Advanced Information Management, Korea, 2008.
- [BALo9] C. V. Blanco, "The OpenNebula Virtual Infrastructure Manager", Xen Directions Europe 2009 joint with LinuxTag, Berlin, Germany, June, 2009.

# References (Cont.)



- [GOL74] R. P. Goldberg, "Survey of Virtual Machine Research," *IEEE Computer*, vol. 7, no. 6, pp. 34-45, 1974.
- [UHL05] R. Uhlig, G. Neiger, D. Rodgers *et al.*, "Intel virtualization technology," *COMPUTER*, vol. 38, no. 5, pp. 48-56, 2005.
- [GOL73] R. Goldberg, "Architecture Principles for Virtual Computer Systems," *PhD Thesis, Harvard University*, 1973.
- [YUo8] Y. Yu, H. Kolam, L. Lam and T. Chiueh, "Applications of a Feather-Weight Virtual Machine", Fourth ACM International Conference on Virtual Execution Environments, WA, USA 2008.
- [CULO8] B. Cully, G. Lefebvre, D. Meyer, M. Feeley, N. Hutchinson, A. Warfield "Remus: High Availability via Asynchronous Virtual Machine Replication", Fifth USENIX Symposium on Networked Systems Design and Implementation, San Francisco, California, 2008, pp. 161-174.

## References (Cont.)



- [PRO04] N. Provos, "A Virtual Honeypot Framework", 13th Conference on USENIX Security Symposium, San Diego, CA, 2004.
- [ALP05] B. Alpern, J. Auerbach, V. Bala, T. Frauenhofer, T. Mummert, and M. Pigott, "PDS: A virtual Execution Environment for Software Deployment," International Conference on Virtual Execution Environments, 2005.
- [VAQ09] L. M. Vaquero, L. Rodero-Merino, J. Caceres, M. Lindner, "A Break in the Clouds: Towards a Cloud Definition", ACM SIGCOMM Computer Communication Review, 2009.
- [FOS08] I. Foster, Y. Zhao, I. Raicu, S. Lu, "Cloud Computing and Grid Computing 360-Degree Compared", Grid Computing Environments Workshop, Austin, 2008.
- [GRO09] R. L. Grossman, "The Case for Cloud Computing", Journal of IT Professional, Vol. 11, Issue 2, 2009.

## References (Cont.)



- [LEA09] N. Leavitt, "Is Cloud Computing Really Ready for Prime Time?" Computer, Vol. 24, No. 1, 2009.
- [KOJ05] T. Koju, S. Takada, N. Doi, "An Efficient and Generic Reversible Debugger using the Virtual Machine based Approach", 1st ACM/USENIX international Conference on Virtual Execution Environments, Chicago, IL, USA, 2005.
- [KIM09] H. Kim, H. Lim, J. Jeong, H. Jo and J. Lee, "Task-Aware Virtual Machine Scheduling for I/O Performance", ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environment, Washington, DC, USA, 2009.
- [LIN04] B. Lin, P. Dinda and D. Lu, "User-Driven Scheduling of Interactive Virtual Machines", 5th IEEE/ACM International Workshop on Grid, Washington, DC, USA, 2004.
- [WEN09] C. Weng, Z. Wang, M. Li and X. Lu, "The Hybrid Scheduling Framework for Virtual Machine Systems", ACM International Conference on Virtual Execution Environments, Washington, DC, USA, 2009.

## References (Cont.)



- [KEAo8] K. Keahey, T. Freeman, "Contextualization: Providing One-Click Virtual Clusters", eScience 2008, Indianapolis, IN. 2008.
- [LAGo9] H. A. Lagar-Cavilla, J. A. Whitney, A. M. Scannell, P. Patchin, S. M. Rumble, E. Lara, M. Brudno and M. Satyanarayanan, "SnowFlock: Rapid Virtual Machine Cloning for Cloud Computing", 4th ACM European Conference on Computer systems, Germany, 2009.
- [LAZo9] G. Laszewski, L. Wang, A. Younge and X. He, "Power-Aware Scheduling of Virtual Machines in DVFS-enabled Clusters", International IEEE Cluster Conference, New Orleans, 2009.
- [BUYo9] R. Buyya, R. Ranjan and R. N. Calheiros, "Modeling and Simulation of Scalable Cloud Computing Environments and the CloudSim Toolkit: Challenges and Opportunities", In Proceedings of the 7th IEEE High Performance Computing and Simulation Conference, Germany, 2009.
- [NURo9] [NURo9] D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, D. Zagorodnov, "The Eucalyptus Open-Source Cloud-Computing System", International Symposium on Cluster Computing and the Grid, China, 2009.

## References (Cont.)



- [BARo3] P. Barham, B. Dragovic, K. Fraser and et al, "Xen and The Art of Virtualization", ACM Symposium on Operating Systems Principles, New York, USA, 2003.
- [VMWo8] VMWare Co., VMWare ESX Server product,  
<http://www.vmware.com/products/esx/index.html>
- [CHEo7] L. Cherkasova , D. Gupta and A. Vahdat,  
"Comparison of the Three CPU Schedulers in Xen", ACM SIGMETRICS Performance Evaluation Review, Vol.35 No.2, 2007.

# Acknowledgement



- My Supervisor: Dr. Sharifi
- Other Professors:
  - Dr. Abdollahi
  - Dr. Minaei
  - Dr. Parsa
- Some DSL members
  - Mr. Omid Kashefi
  - All M.Sc. students who contributed me at their seminar or thesis
    - Arash Rezaei, Ali Hamidi, Ebrahim Asadi, Mahsa Najafzadeh, Somayeh Kafei

## Thanks for Your Attention!

### Q/A