Energy Efficient VM Scheduling for Cloud Data Centers: Exact allocation and migration algorithms

Chaima Ghribi, Makhlouf Hadji and Djamal Zeghlache

Networks and Mobile Multimedia Services Department
CNRS UMR Samovar, Telecom SudParis
Evry, France

ExtremeGreen workshop, During CCGrid 2013
May 13-16, 2013, Delft, The Netherlands.

Agenda

Research Context

Research Issue

Objectives & Contributions

Proposed Algorithms

Evaluation

Conclusion

Research Context

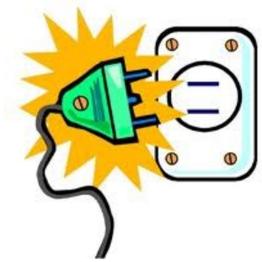

- Energy efficiency is increasingly important for future ICT.

- Rapid growth of the demand for computational power has led to the creation of large-scale datacenters.

- The tremendous rise of electricity consumption is a major contributor to the escalating cost of datacenter ownership and to the increasing atmospheric CO_2 levels.

Research Context

- **Gartner Report :**

Source: Gartner Report, Financial Times, 2007

–IT industry contributes **2%** of world's total CO_2 emissions

–Roughly equivalent to aviation industry

–IT energy usage will double next 4 years

- **McKinsey report :**

Source: McKinsey & Company, “Revolutionizing Data Center Energy Efficiency, 2008”

– A typical data center consumes as much energy as 25,000 households.

–Carbon emission due to Data Centers worldwide is more than both Argentina and the Netherlands emission.

–The total energy bill for data centers in 2010 will be over $11 billion and energy costs in a typical data center doubles every five years.

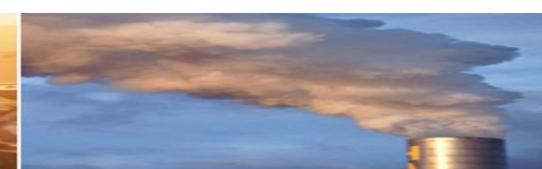

How to Reduce the Energy Consumption in Cloud Data centers ?

Research Issue

- Technologies like free cooling used by modern data centers (e.g. Google, Facebook...) are interesting for reducing the power consumption of cooling.

- Modern data centers are also operating at a raised operating temperature which decreases cooling costs and increases power efficiency.

To reduce the power consumption of cloud data centers we aim to reduce the power consumed by IT equipments.

Research Issue

- Servers are permanently switched on even if they are not used.
- Idle servers (which did not host any virtual machine) consume about 60% of their peak power.
- Idle power wasted when servers run at low utilization.

Objectives

- Proposing an energy aware VM scheduler based on dynamic placement of VMs in cloud data centers.

- Proposing **optimal** algorithms for VM allocation and migration to reduce power consumption in cloud data centers

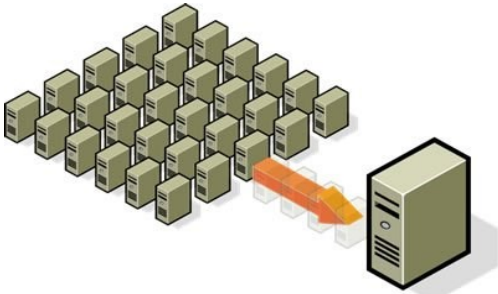

Contributions

- **Energy-aware VM scheduler :**
 Responsible for the optimal energy aware VM placement in the data center.

- **Energy consumption estimator :**
Can rely on an energy estimation tool that uses power models to infer power consumption of VMs or servers from resource usage.

- **Cloud Iaas manager:**
(i.e. OpenStack, OpenNebula and Eucalyptus) control and manage cloud resources and handle clients requests, VM scheduling and fetch and store images in storage spaces.

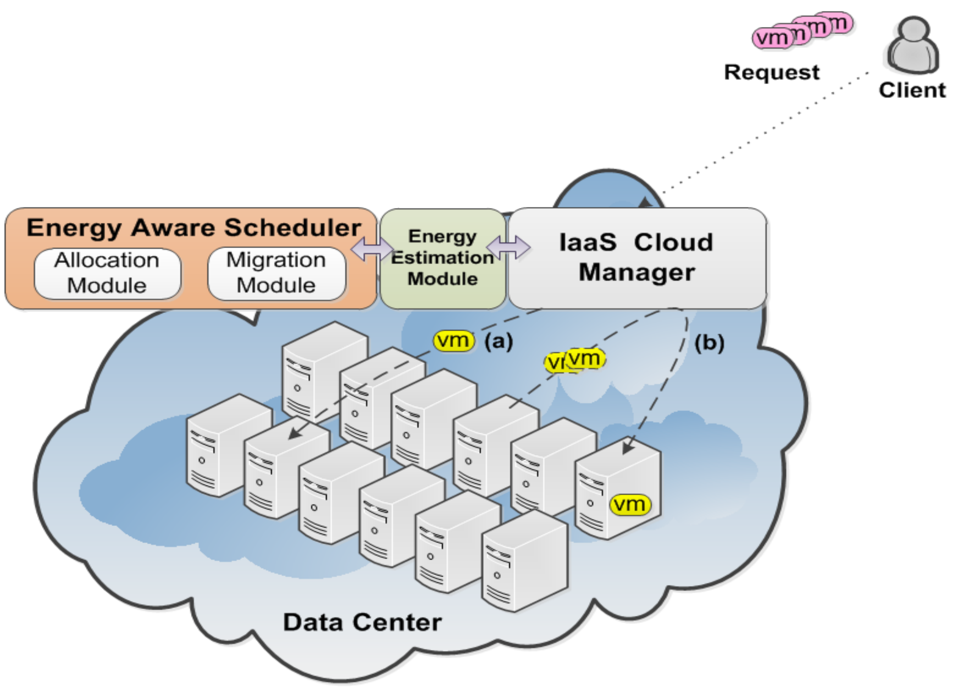

Contributions

- Exact VM Allocation algorithm
- Exact VM Migration algorithm
- Adapted energy aware best fit algorithm

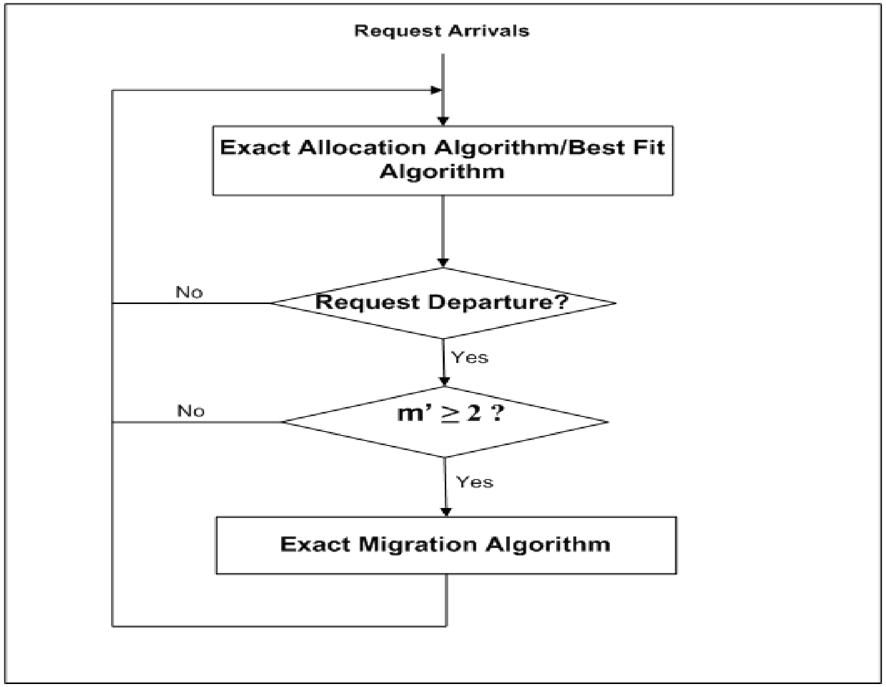

Exact VM Allocation algorithm

- **Objective:**
  Initial VM placement with reducing the number of used servers.

- **Mathematical Programming Formulation**:

❖Modelled as a bin packing problem with a minimum power consumption objective.

$$
\begin{cases}
\min Z = \sum_{j=1}^{m} e_j \\
\sum_{i=1}^{n} p_i x_{ij} \leq P_{j,Max} e_j - P_{j,Current} \\
\sum_{j=1}^{m} x_{ij} = 1; \quad i \in \{1..n\}; \quad j \in \{1..m\} \\
\sum_{j=1}^{m} e_j \geq \left\lceil \frac{\sum_{j=1}^{m} P_{j,current}}{P_{j,Max}} \right\rceil \\
e_j \in \{0,1\}; \quad x_{ij} \in \{0,1\}
\end{cases}
$$

Exact VM Migration algorithm

- **Objective:**
Optimize constantly the data center's power consumption via dynamic VM consolidation.

- **Mathematical Programming Formulation**:

❖ Based on a mathematical model for VM consolidation via migration that relies on a linear integer programming formulation.

$$
\begin{cases}
\max M = \sum_{i=1}^{m'} P_{i,idle} y_i - \sum_{i=1}^{m'} \sum_{j=1}^{m'} \sum_{k=1}^{q_i} p'_k z_{ijk} \\
z_{ijk} + z_{jlk'} \leq 1 \quad j \neq i;\ l \neq j;\ k \neq k' \\
\sum_{\substack{j=1 \\ j \neq i}}^{m'} z_{ijk} \leq 1 \quad i \in \{1..m'\}; j \in \{1..m'\}; l \in \{1..m'\} \\
\qquad k \in \{1..q_i\}; k' \in \{1..q_j\} \\
\sum_{\substack{i=1 \\ i \neq j}}^{m'} \sum_{k=1}^{q_i} p_k z_{ijk} \leq (P_{j,Max} - P_{j,Current})(1 - y_j) \\
\sum_{\substack{j=1 \\ j \neq i}}^{m'} \sum_{k=1}^{q_i} z_{ijk} = q_i y_i; \quad z_{ijk} \in \{0,1\}; y_i \in \{0,1\} \\
\sum_{i=1}^{m'} y_i \leq m' - \left\lceil \frac{\sum_{j=1}^{m'} P_{j,Current}}{P_{j,Max}} \right\rceil \\
z_{ijk} \Delta t_k \geq T_0
\end{cases}
$$

Exact VM Migration algorithm

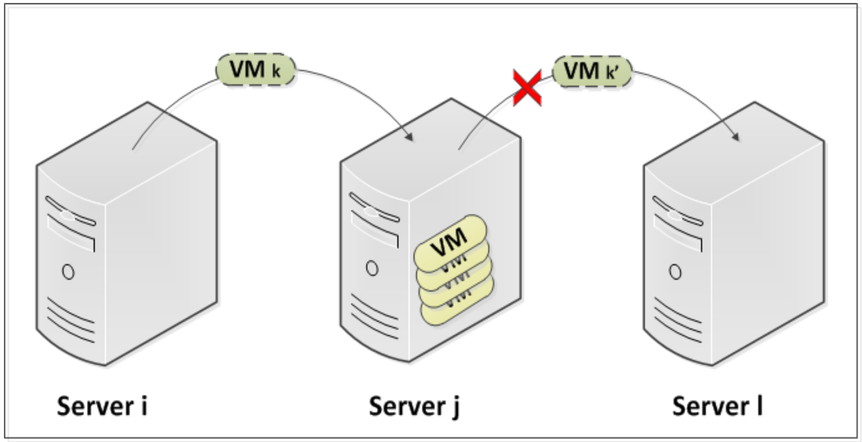

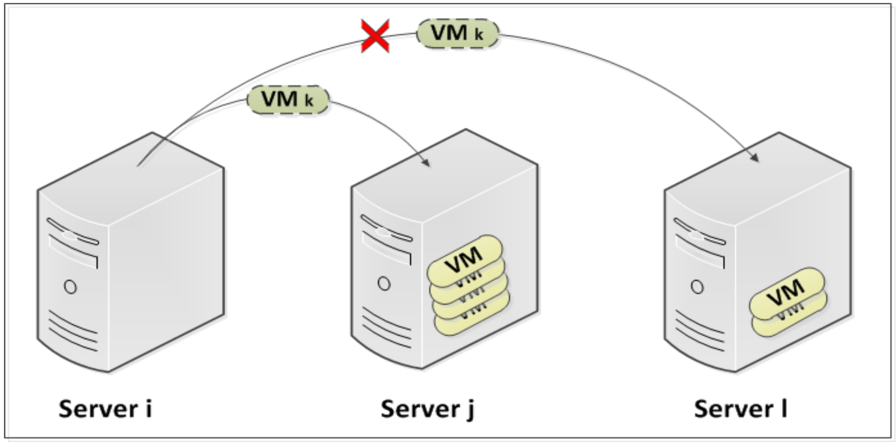

- A server candidate to a migration should not migrate his own VMs
- A VM can not be migrated to many servers in the same time

Adapted energy aware best fit algorithm

Adaptation of the Best-Fit heuristic which consists in two principal steps:

- Sort items (VMs) in a decreasing sequence of their power consumption.

- Place all the sorted VMs by considering the first item (VM) in a server with a minimum rest of power consumption.

Evaluation

- Our proposed algorithms are evaluated through a Java language implementation and the linear solver CPLEX.
- A dedicated simulator is developed.
- The objective is to quantify the percentage of energy savings or power consumption savings that can be expected when combining the exact allocation algorithm and the exact migration algorithm.
- Simulations were conducted and significant results were obtained by combining both algorithms.

Evaluation

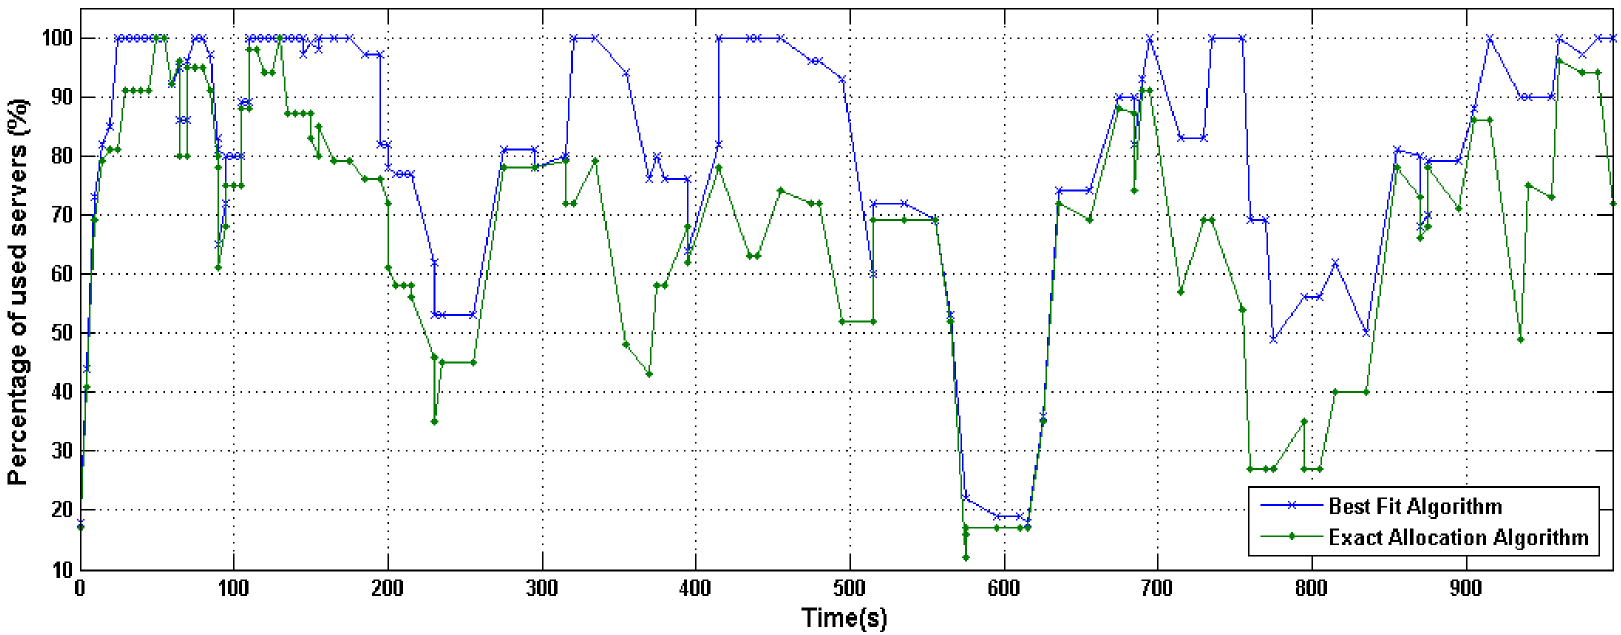

- Comparison between Exact Allocation and Best Fit algorithms

Evaluation

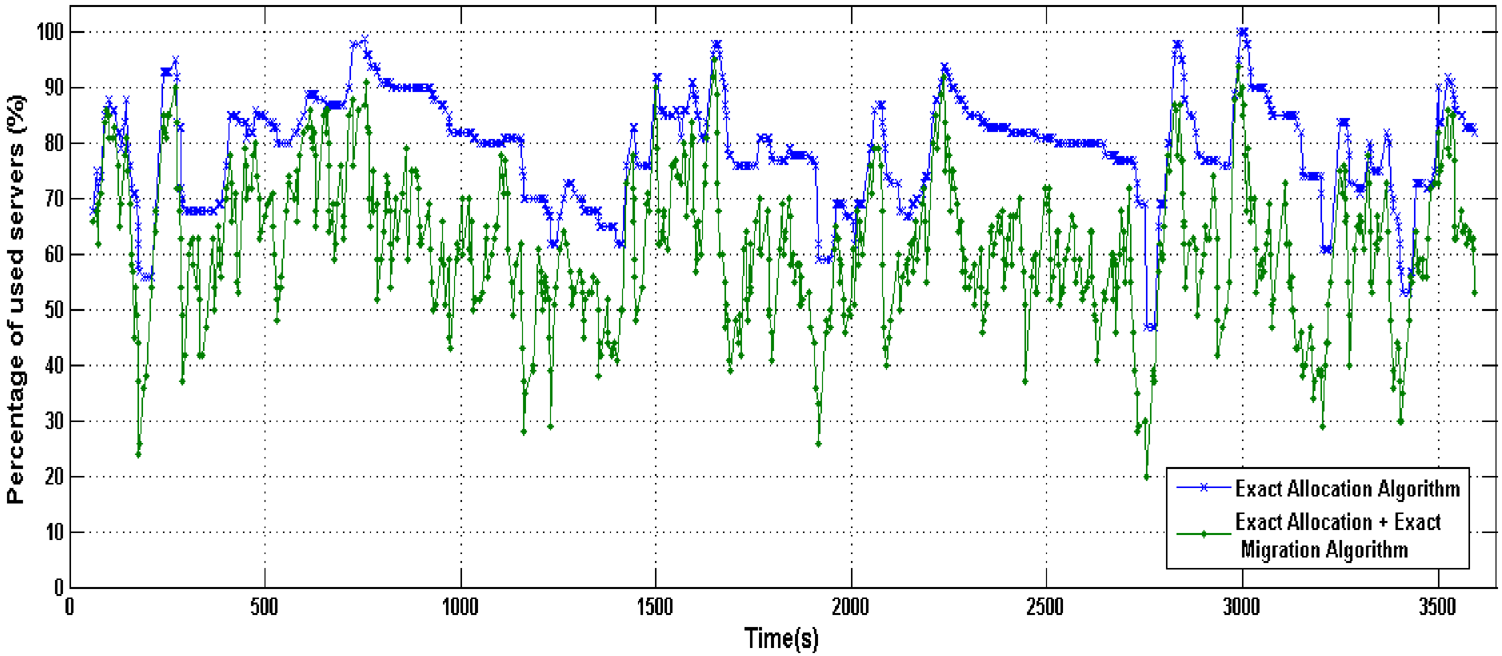

- Performance comparison of the exact allocation algorithm with and without migration

Evaluation

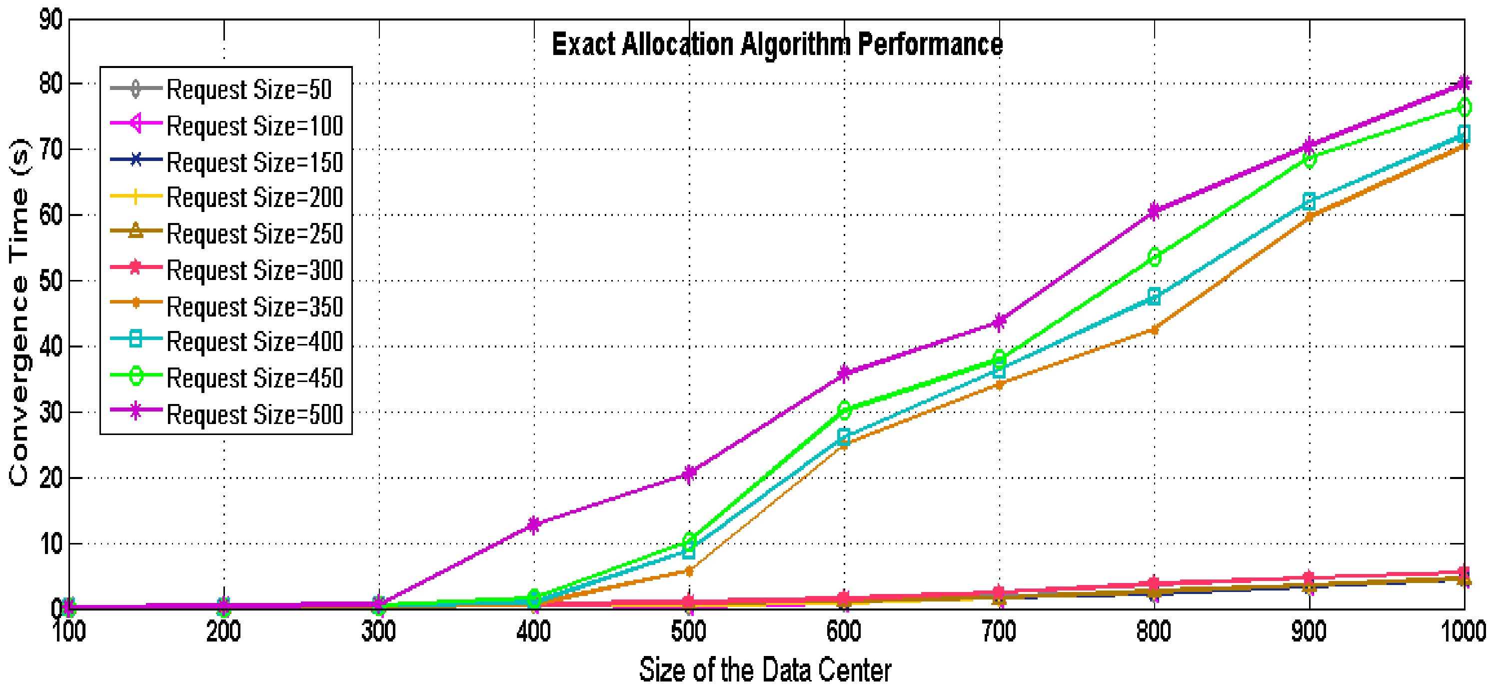

- Convergence time of the Exact Allocation Algorithm

Evaluation

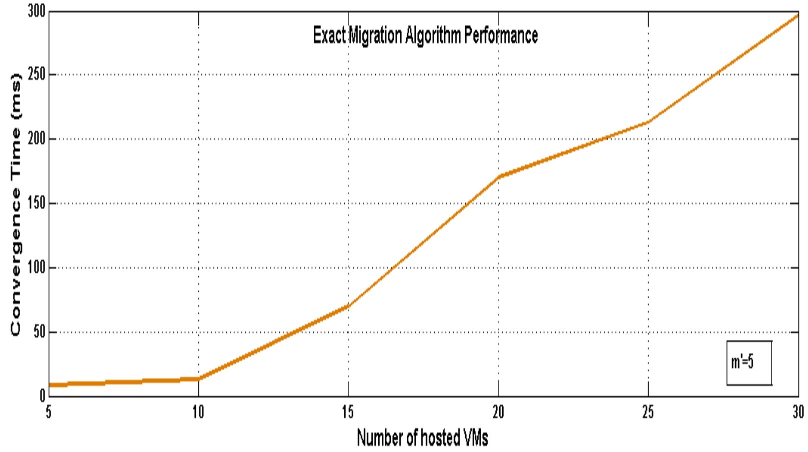

■ Convergence time of the Exact Migration Algorithm (m'=5)

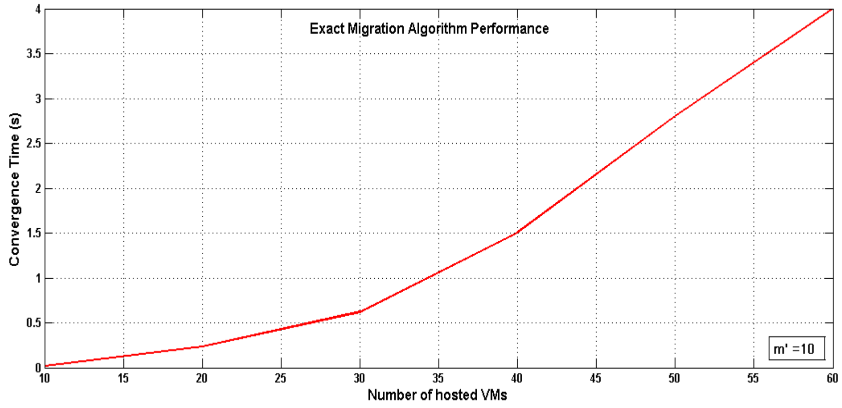

■ Convergence time of the Exact Migration Algorithm (m'=10)

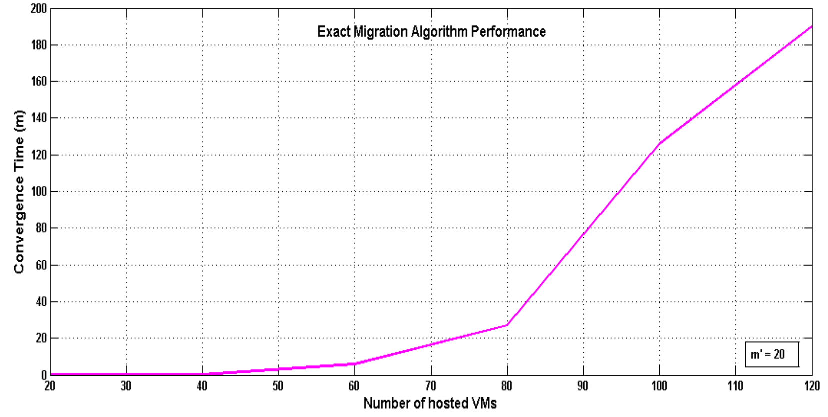

■ Convergence time of the Exact Migration Algorithm (m'=20)

Evaluation

$\mu^{-1}(s)$ \ $\lambda^{-1}(s)$	5	10	15	20	25	30
10	35,55	36,59	00,00	00,00	00,00	00,00
20	27,29	34,00	35,23	38,50	00,00	00,00
30	17,48	27,39	35,21	40,32	**41,89**	36,58
40	16,77	18,85	22,02	32,31	39,90	40,50
50	10,86	16,17	19,85	22,30	39,20	36,52
60	08,63	14,29	18,01	22,13	25,15	30,68
70	08,10	14,00	14,86	15,90	22,91	23,20
80	07,01	10,20	10,91	15,34	17,02	21,60
90	06,80	09,52	10,31	14,70	16,97	19,20
100	05,90	07,50	08,40	12,90	16,00	14,97

- Percentage of gained energy when migration is used

Conclusion

- This paper investigate the idea of using exact algorithms to solve the problem of energy efficient VM scheduling in cloud data centers.

- Exact proposed algorithms can achieve significant energy savings while maintaining feasible runtimes.

Energy Efficient VM Scheduling for Cloud Data Centers: Exact allocation and migration algorithms

Chaima Ghribi, Makhlouf Hadji and Djamal Zeghlache

Networks and Mobile Multimedia Services Department
CNRS UMR Samovar, Telecom SudParis
Evry, France

ExtremeGreen workshop, During CCGrid 2013
May 13-16, 2013, Delft, The Netherlands.