ServiceWave 2010 CONFERENCE
13TH-15TH DECEMBER

Service Provisioning on the Cloud: Distributed Algorithms for Joint Capacity Allocation and Admission Control

Danilo Ardagna, Carlo Ghezzi, Barbara Panicucci, Marco Trubian

Dipartimento di Elettronica e Informazione
Politecnico di Milano

Problem overview

- Computing systems are becoming increasingly virtual
- Moving into a world in which SaaS may be discovered at run-time
- Service composition may occur dynamically
  - Need to provide QoS guarantees

Problem overview

- Recent advances in cloud computing are pushing into another dimension of virtuality:
    - Emerging paradigm, on-demand provisioning of IT services (SaaS, IaaS, PaaS) through the Internet
    - Users do not need to implement and administer the IT infrastructure directly
    - Service Providers obtain economies of scale reducing also energy costs

Our work

- Resource management becomes more complex
- We take the perspective of a SaaS provider:
  - Maximize the revenues from SLAs, while minimizing the cost of use of IaaS resources
  - Distributed algorithm that jointly addresses the capacity allocation and admission control
  - Non-linear programming problem and solved with decomposition techniques

Research assumptions

- Multiple transactional WSs and each service represents a different application
- The hosted WSs can be heterogeneous (resource demands, workload intensities and QoS requirements), and will be modelled as independent WS classes
- SLA contract associated with each WS class:
  - Full revenue iff response time lower than a given threshold

Research assumptions

- Applications are deployed in VMs
- Each VM hosts a single Web service application
- Multiple VMs implementing the same WS class can run in parallel running in homogeneous VMs
- Services can be located on multiple sites
- IaaS providers charge the SaaS on a hourly basis

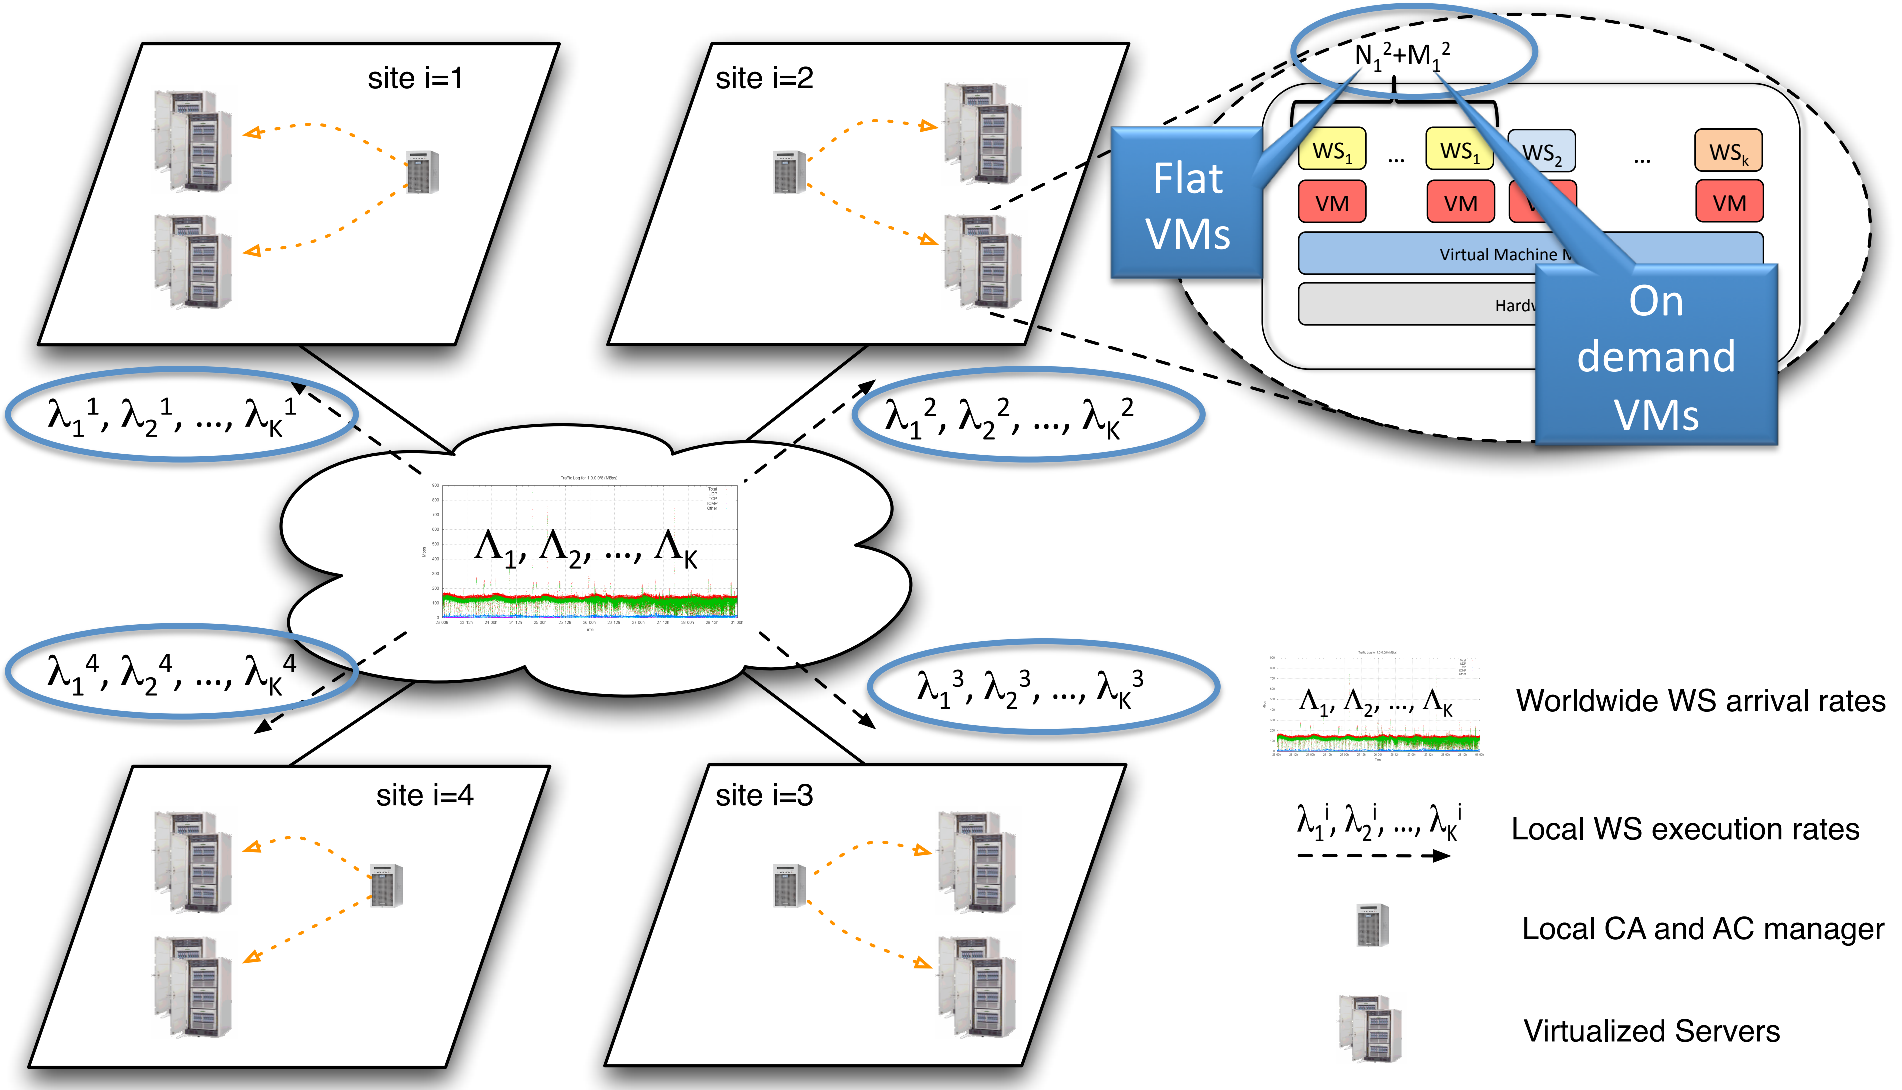
site i=1
site i=2
site i=3
site i=4
$N_1^2+M_1^2$
WS_1
…
WS_1
WS_2
…
WS_k
VM
VM
VM
VM
Virtual Machine M
Hardw
Flat VMs
On demand VMs
$\lambda_1^1, \lambda_2^1, \ldots, \lambda_K^1$
$\lambda_1^2, \lambda_2^2, \ldots, \lambda_K^2$
$\lambda_1^3, \lambda_2^3, \ldots, \lambda_K^3$
$\lambda_1^4, \lambda_2^4, \ldots, \lambda_K^4$
$\Lambda_1, \Lambda_2, \ldots, \Lambda_K$
$\Lambda_1, \Lambda_2, \ldots, \Lambda_K$
Worldwide WS arrival rates
$\lambda_1^i, \lambda_2^i, \ldots, \lambda_K^i$
Local WS execution rates
Local CA and AC manager
Virtualized Servers

Problem notation

- I: Set of IaaS geographically distributed sites
- K: Set of WS request classes

- T: Control time horizon
- C_i: Capacity of VMs at site i
- μ_k: Maximum service rate of a capacity 1 VM
- ρ_k: Utilization upper bound for class k VMs

- ω_k^i: SaaS revenue for a single WS class k request at site i
- $\bar{c}^i$: Time unit cost for flat VMs at site i
- $\hat{c}^i$: Time unit cost for on demand VMs at site i

Problem formulation

- Each WS class hosted in a VM is modelled as an M/G/1/PS queue
- Lower bound on the probability that the response time of a class k request at site *i*, R_k^i, exceeds a threshold $\bar{R}_k$:

$$P[R_k^i \leq \overline{R}_k] \geq 1 - e^{-(C_i\,\mu_k - \frac{\lambda_k^i}{N_k^i + M_k^i})\,\overline{R}_k}$$

Problem formulation

(P)

$$\max \sum_{k\in K} \sum_{i\in I} \left(\omega_k^i \, \lambda_k^i \, (1 - e^{-(C_i \, \mu_k - \frac{\lambda_k^i}{N_k^i + M_k^i}) \, \overline{R}_k})\right) T - \sum_{i\in I} \overline{c}^i \sum_{k\in K} N_k^i - \sum_{i\in I} \hat{c}^i \sum_{k\in K} M_k^i$$

$$\lambda_k^i < \rho_k \, C_i \, \mu_k \, (N_k^i + M_k^i)$$

$$\sum_{i\in I} \lambda_k^i \le \Lambda_k$$

$$\sum_{k\in K} N_k^i \le N^i$$

Problem analysis and solution

- Problem (P) is a general nonlinear optimization problem
- Decomposition techniques support a distributed protocol for the joint capacity allocation and admission control
- Each site solves its problem using both local information and information received from other sites
- Iterative method, at each iteration λ^{i}_{k} is the only information shared among sites
- Optimization duality theory

Problem analysis and solution

- (P) is concave:
  - The eigenvalues of the objective function are non positive
  - Affine constraints
- Strong duality holds
- Solve the primal via the dual
- Lagrangian relaxation allows separating (P) into /I/ subproblem

Duality theory

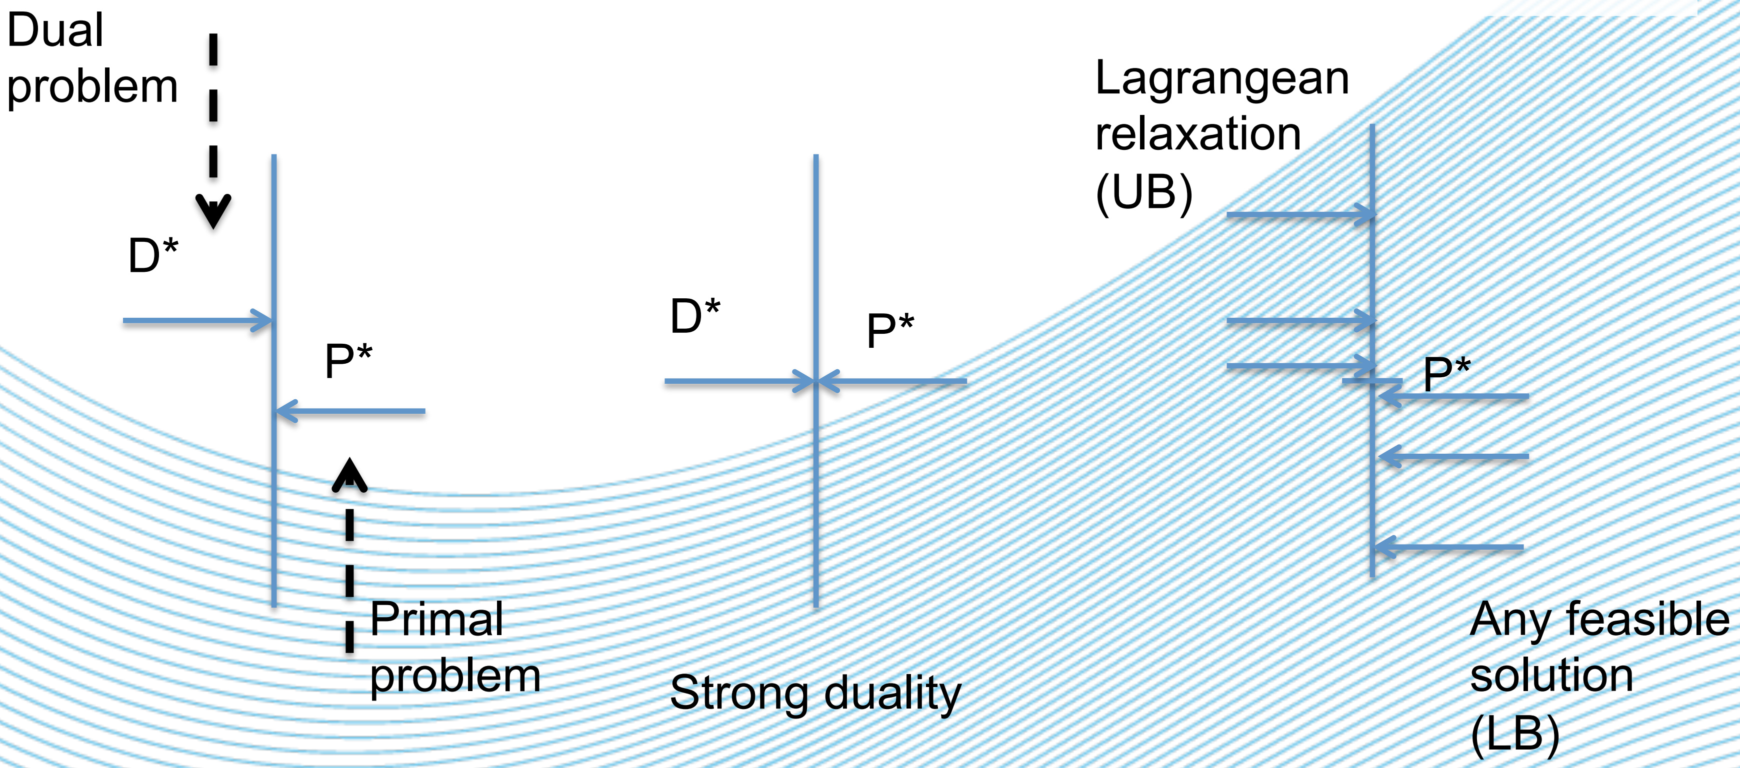

Dual problem

D*

P*

Primal problem

D*

P*

Strong duality

Lagrangean relaxation (UB)

P*

Any feasible solution (LB)

Lagrangean relaxation

(PR)

$$\mathcal{L}(r) = \max \sum_{k \in K} \sum_{i \in I} \big(\omega_k^i \, \lambda_k^i \, (1 - e^{-(C_i \, \mu_k - \frac{\lambda_k^i}{N_k^i + M_k^i}) \, \overline{R}_k})\big) \, T - \sum_{i \in I} \overline{c}^i \sum_{k \in K} N_k^i +$$

$$- \sum_{i \in I} \hat{c}^i \sum_{k \in K} M_k^i - \sum_{k \in K} \big(\sum_{i \in I} \lambda_k^i - \Lambda_k \big) \, r_k,$$

$$\lambda_k^i < \rho_k \, C_i \, \mu_k \, (N_k^i + M_k^i)$$

$$\sum_{k \in K} N_k^i \leq N^i$$

Dual decomposition

- For a given vector of multipliers r≥0, the dual decomposition results in solving at each site i:

$$(SUB_i)$$

$$\max \sum_{k\in K} \left(\omega_k^i\, T\, (1 - e^{-(C_i\,\mu_k - \frac{\lambda_k^i}{N_k^i+M_k^i})\,\overline{R}_k}) - r_k\right) \lambda_k^i - \overline{c}^i \sum_{k\in K} N_k^i - \hat{c}^i \sum_{k\in K} M_k^i + \sum_{k\in K} \Lambda_k\, r_k$$

$$\lambda_k^i < \rho_k\, C_i\, \mu_k\, (N_k^i + M_k^i) \qquad \sum_{k\in K} N_k^i \le N^i$$

- To define the dual function *$L_i(r)$,* we consider, for all i, the optimal value of *(SUB_i)* for a given r≥0

Dual decomposition

- The dual problem is then given by:

$$\min_r \mathcal{L}(r) = \sum_i \mathcal{L}_i(r)$$

- The dual problem admits decentralized solution and it can be solved by using a sub-gradient method:

$$r_k(t+1) = r_k(t) - \alpha_t \left(\mathit{\Lambda}_k - \sum_{i \in I} \lambda_k^i\right)$$

Joint CA and AC procedure

```
while not $STOP$ do
    for $i \in I$ do
        Solve $(SUB_i)$; Let $N_k^i$, $\lambda_k^i$, $M_k^i$ the solution; Let $UB$ the function value;
        Broadcast $\lambda_k^i$;
    end
    for $k \in K$ do
        if $\sum_{i \in I} \lambda_k^i > \Lambda_k$ then
            for $i \in I$ do
                $\hat{\lambda}_k^i := \lambda_k^i \frac{\Lambda_k}{\sum_{i \in I} \lambda_k^i}$;
            end
        end
    end
    Update $LB$ (objective function value of (P) evaluated with $\hat{\lambda}_k^i$ ) if better
    for $k \in K$ do
        $r_k := max(0, r_k - \alpha \frac{UB-LB}{\sum_{k \in K} (\Lambda_k - \sum_{i \in I} \lambda_k^i)^2} (\Lambda_k - \sum_{i \in I} \lambda_k^i))$;
    end
    Update $\alpha$;
    if $(UB - LB)/LB <$ precision threshold OR the maximum number of
    iterations is reached then
        STOP=true;
    end
end
```

Experimental results

- All tests have been performed on VMWare virtual machine running on an Intel Nehalem dual socket quad-core system with 32 GB of RAM
- The virtual machine has a physical core dedicated with guaranteed performance and 4 GB of memory reserved
- KNITRO 6.0 solver

Experimental results

- Large set of randomly generated instances:
  - $|I|$ in [20, 60]
  - $|K|$ in [100, 1000]
  - ρ_k=0.6
  - μ_k has been varied in [0.1,1]

- Stopping criterion: relative gap between the current value of the upper and the lower bound

Experimental results

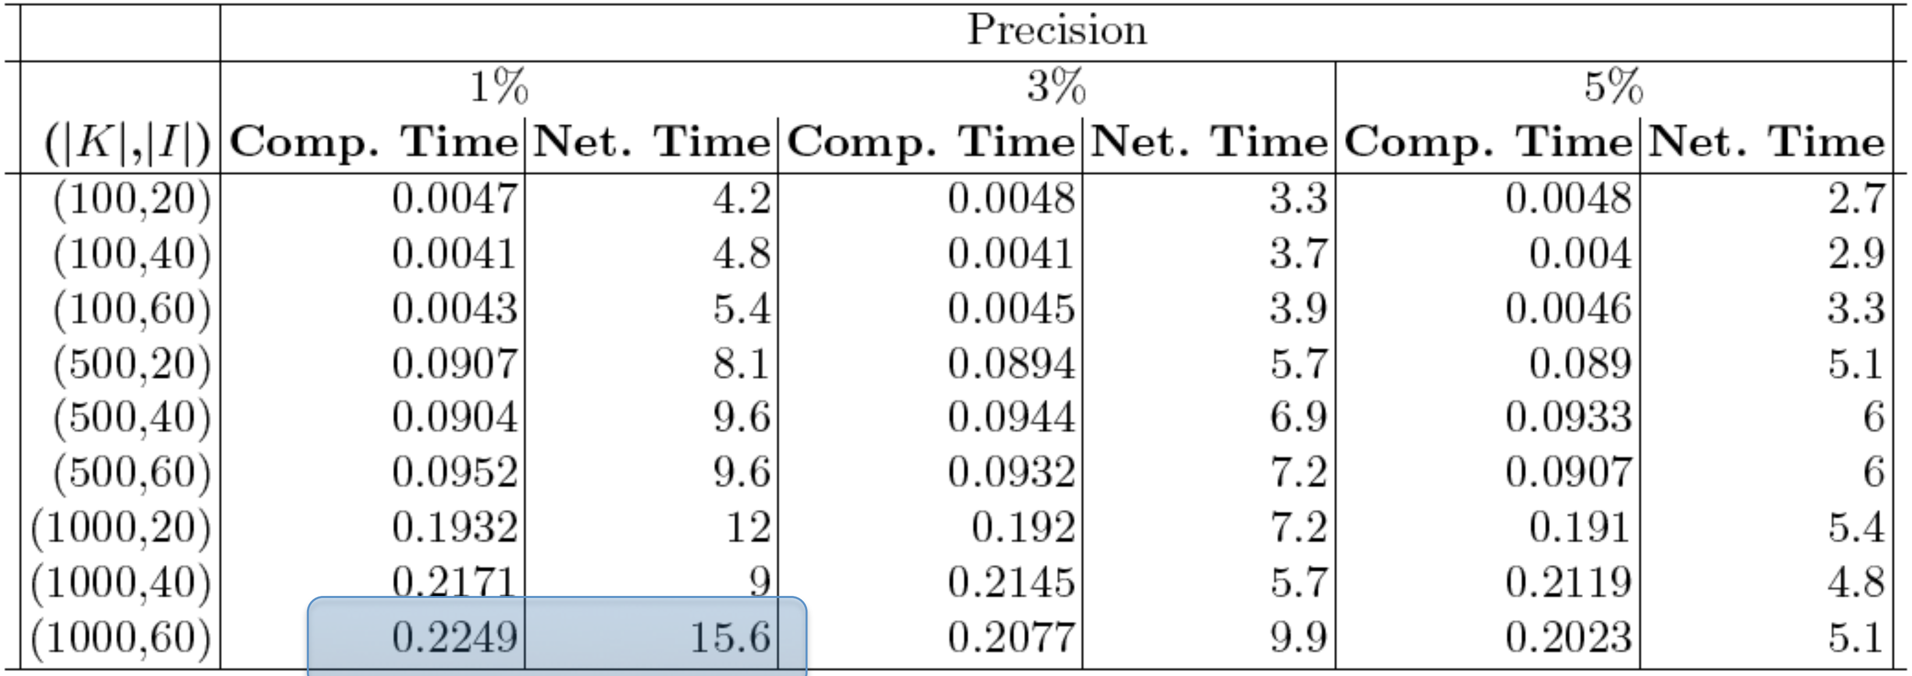

	Precision					
	1%		3%		5%	
$(\lvert K\rvert,\lvert I\rvert)$	Comp. Time	Net. Time	Comp. Time	Net. Time	Comp. Time	Net. Time
(100,20)	0.0047	4.2	0.0048	3.3	0.0048	2.7
(100,40)	0.0041	4.8	0.0041	3.7	0.004	2.9
(100,60)	0.0043	5.4	0.0045	3.9	0.0046	3.3
(500,20)	0.0907	8.1	0.0894	5.7	0.089	5.1
(500,40)	0.0904	9.6	0.0944	6.9	0.0933	6
(500,60)	0.0952	9.6	0.0932	7.2	0.0907	6
(1000,20)	0.1932	12	0.192	7.2	0.191	5.4
(1000,40)	0.2171	9	0.2145	5.7	0.2119	4.8
(1000,60)	0.2249	15.6	0.2077	9.9	0.2023	5.1

Joint CA and AC procedure execution time (s)

Experimental results

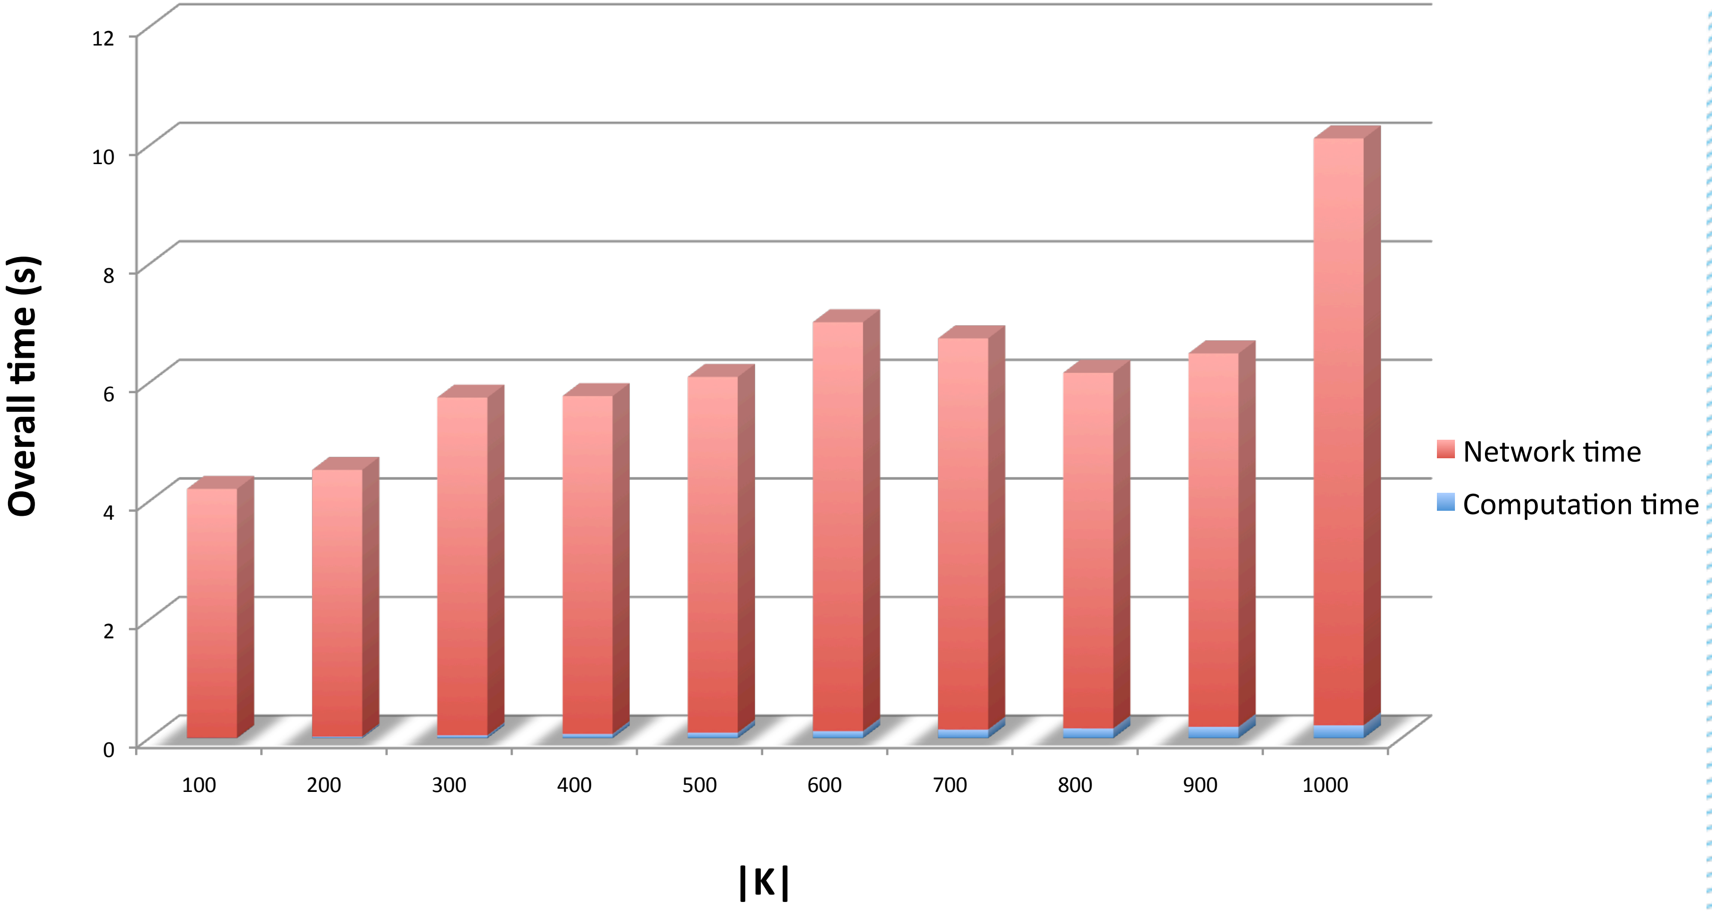

Computation and network times for |*I*| = 60 with varying number of WS Classes at 3% precision

Conclusions

- We proposed a distributed algorithm for the joint admission control and capacity allocation of SaaS cloud
- We developed a scalable solution to the problem providing a-priori fixed optimality guarantees
- Future work will extend the validation by performing experiments in real cloud environments
- Game theoretical formulation of the problem

Thanks for your attention, any questions?

ServiceWave 2010 CONFERENCE

13TH-15TH DECEMBER

GHENT

SILVER SPONSOR

LABS hp

PLATINUM SPONSOR

SIEMENS

GOLD SPONSOR

THALES

ENGINEERING

www.servicewave.eu

ORGANISED BY

European and National Technology Platforms

NESSI

INES Iniciativa Española de Software y Servicios

eMobility (Connect to the network community)

EPoSS European Technology Platform on Smart Systems Integration

NEM Networked & Electronic Media

European Networks of Excellence and Projects

S-CUBE

IT-Tude.com

SOFI

coordinated by

Universität Duisburg Essen

PALUNO The Ruhr Institute for Software Technology

in collaboration with

ICSOC Int'l Conference on Service Oriented Computing

hosted by

ibbt connect.innovate.create