nano-tera.ch
swiss scientific initiative in health / security / environment systems

YINS RTD 2013

Multi-Objective System-Level Management of Geo-Distributed Data Centers

Ali Pahlevan, Pablo Garcia del Valle, David Atienza

Embedded Systems Laboratory (ESL), EPFL

Cloud Computing

- Cloud computing is pervasive

- IT companies business

- DCs huge power & annual electricity bills

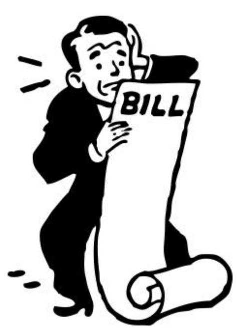

Data Centers (DCs) electricity consumption will reach 8% of the worldwide electricity by 2020

Green Geo-Distributed DCs

- Multiple DCs in different geographical locations
- Renewable energy sources
- Network connection
- Electricity price diversities

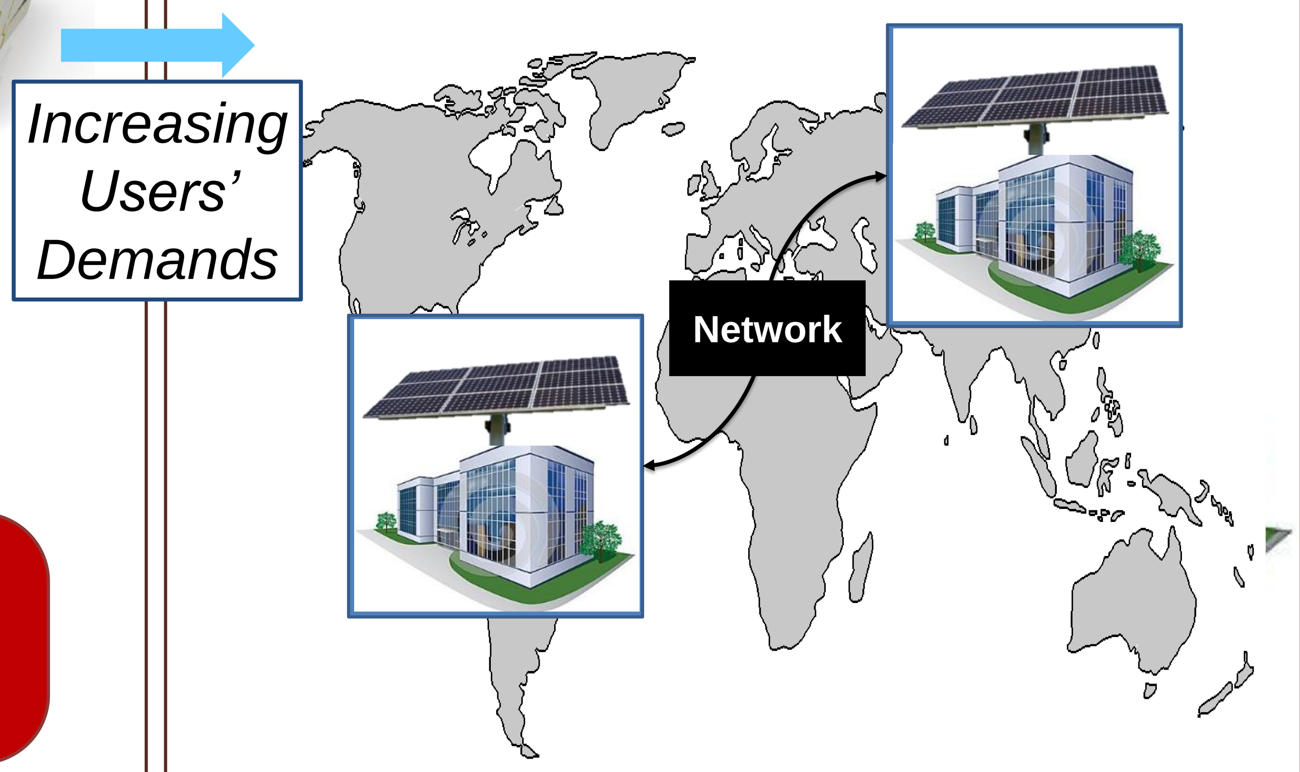

Problem Description

- Multi objective DC optimization
  - Operational cost
  - Energy consumption
  - Performance (Response time)

To improve the VM allocation & DC efficiency

Exploit Application Characteristics

- DCs host heterogeneous applications
  - *CPU-load correlation*: maximum VMs' utilizations coincide at the same time
  - *Data correlation*: amount of data that VMs need to exchange

Energy-performance trade-off

Solution: Two-Phase Multi-Objective VM Placement Exploiting VMs' Characteristics

A VMs Clustering

VM utilization peaks do not coincide along with lower data communication

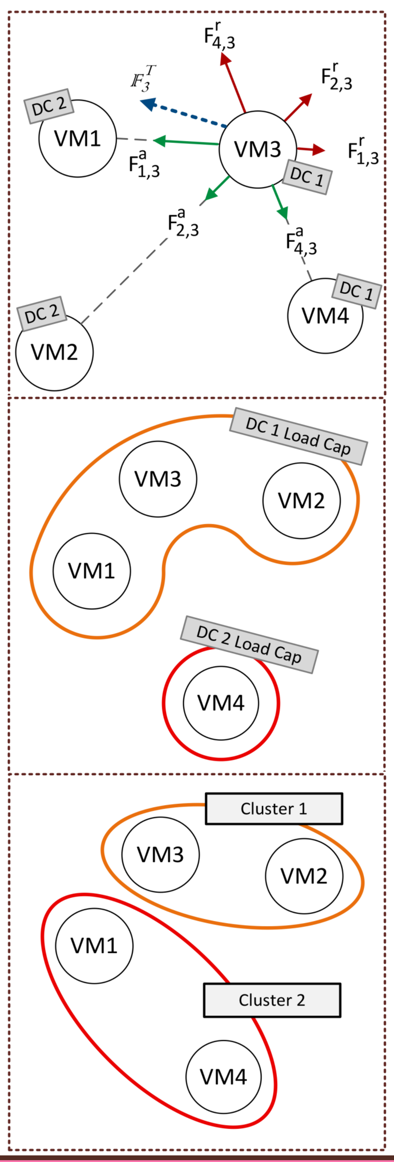

Step 1

Repulsion & attraction forces based on CPU-load & data correlations

Step 2

Modified K-MEANS clustering per each DC

- Load cap based on energy sources

Step 3

K-MEANS output correction based on QoS

- VMs selection for migration

Two-phase VM placement (global clustering + local allocation) + green controller

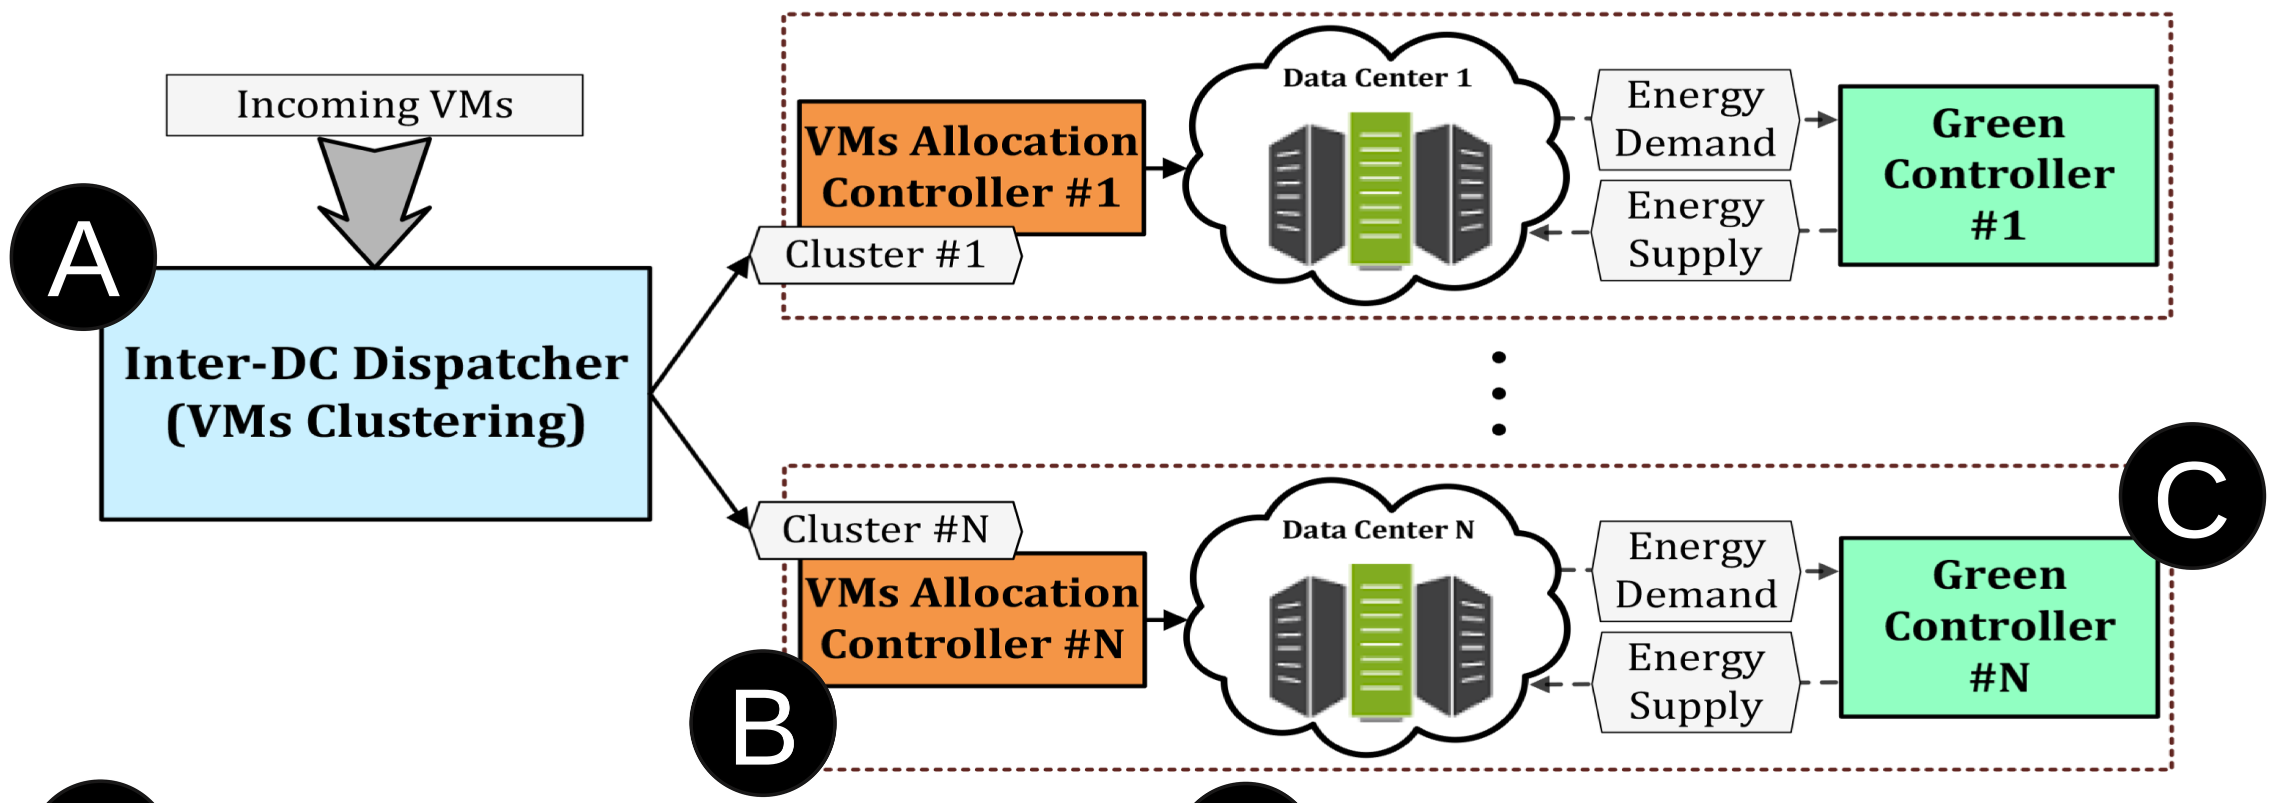

B VMs Allocation [1]

- Find VMs with lowest CPU-load correlation
- First-Fit-Decreasing algorithm for consolidation

C Green Controller

- Low complexity rule-based controller to manage renewable & DC energy
- Compensating difference between real & forecasted information

Experimental Results

Experimental Setup

DC	Number of Servers	PV Capacity (KWp)	Battery Capacity (KWh)
DC1	1500	150	960
DC2	1000	100	720
DC3	500	50	480

Proposed algorithm compared to state-of-the-art solutions:

- Energy-aware VM allocation (Ener-aware) [1]
- Cost-aware VM placement (Pri-aware) [2]
- Network-aware VM placement (Net-aware) [3]

Total Cost, Energy & Performance

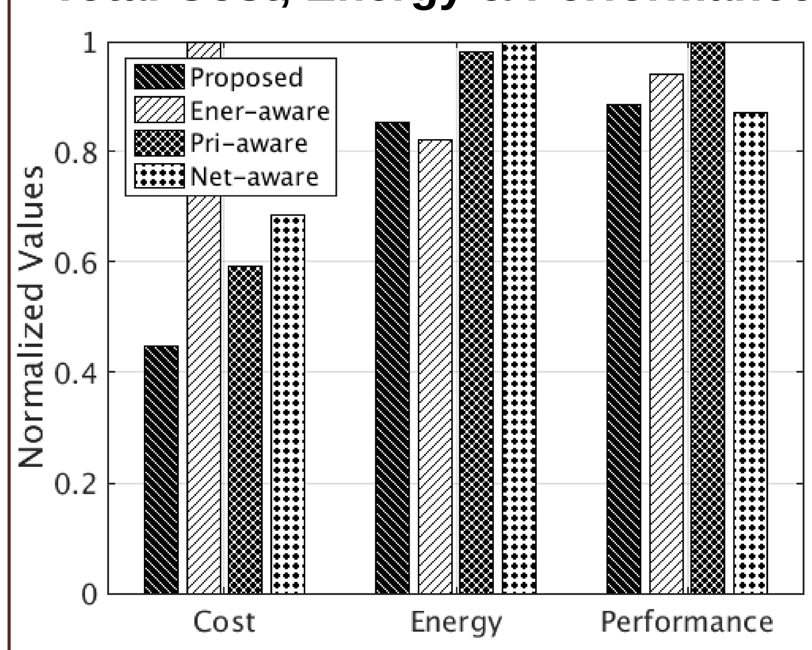

Cost- & Energy-Performance Trade-offs

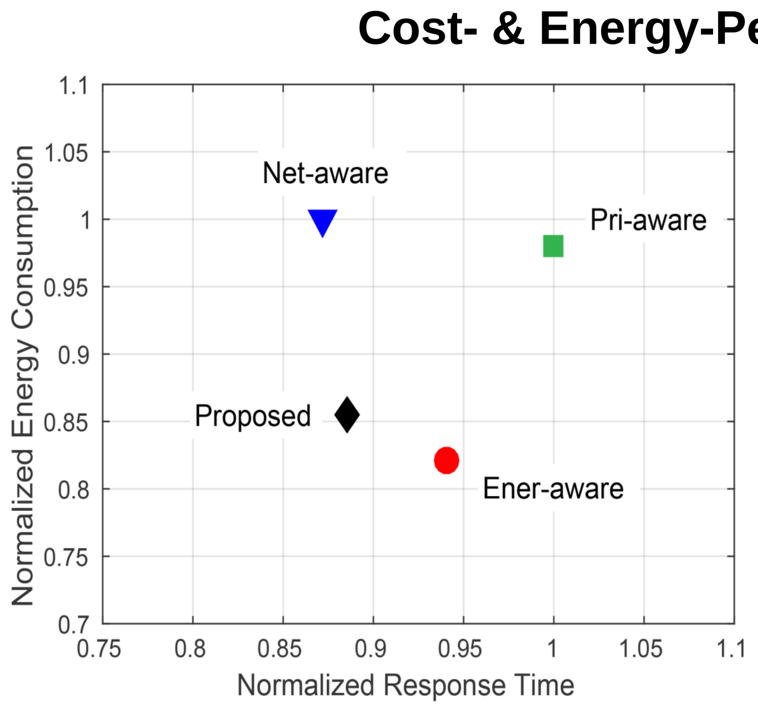

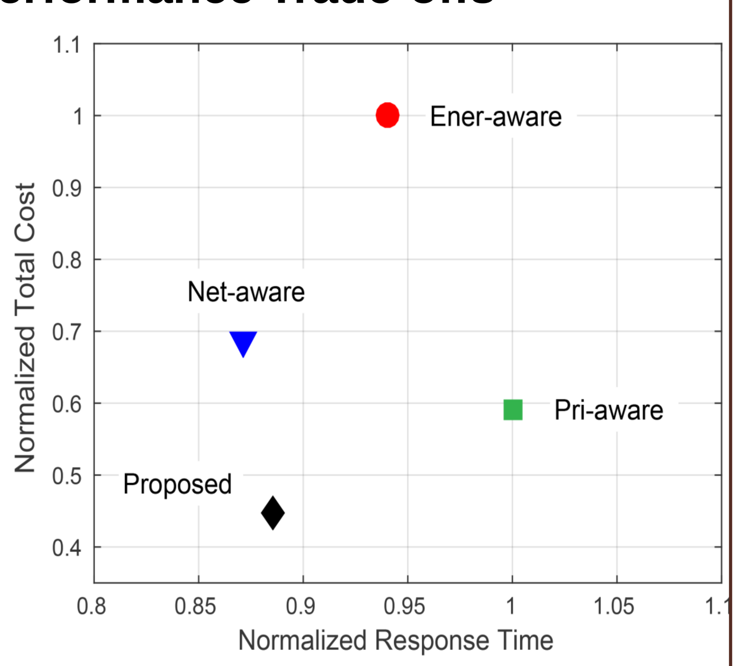

Conclusion

- Cost optimization & energy-performance trade-off in green geo-distributed DCs
- Two-phase multi-objective VM placement exploiting VMs' characteristics
  - **Up to 55% operational cost savings**
  - **Up to 15% energy savings** with only 3% overhead compared to Ener-aware
  - **Up to 12% performance improvements** with only 2% degradation compared to Net-aware

References

[1] J. Kim et al., "Correlation-Aware VM Allocation For Energy-Efficient Datacenters," DATE, 2013

[2] L. Gu et al., "Joint Optimization of VM Placement and Request Distribution For Electricity Cost Cut in Geo-Distributed Data Centers," ICNC, 2015

[3] O. Biran et al., "A Stable Network-Aware VM Placement For Cloud Systems," CCGRID, 2012