Technical University of Cluj-Napoca
Faculty of Computer Science and Automation
Department of Computer Science

Energy Aware Adaptation Methodology for Improving Service Centers Energy Efficiency

Ioan Salomie, Tudor Cioara, Ionut Anghel, **Daniel Moldovan**

GAMES

SEVENTH FRAMEWORK PROGRAMME

Green Active Management of Energy Efficiency in IT Service Centres

- **Partners**
  - Engineering Ingegneria Informatica (Italy)
  - Politecnico di Milano (Italy)
  - High Performance Computing Centre HLRS (Germany)
  - Technical University of Cluj-Napoca (Romania)
  - IBM Research Laboratory in Haifa (Israel)
  - Christmann Informationstechnik (Germany)
  - EnergoEco (Romania)
  - Enel.si (Italy)

Green Active Management of Energy Efficiency in IT Service Centres

- **Objective**
  - New generation of energy aware self-adapting IT service centers
  - New methodologies, metrics, services and tools for active management of service centers energy efficiency

Outline

- Introduction
- Service Center Energy Awareness Model
  - Energy Aware Context Model (EACM)
  - EACM model processing and management
- Energy Aware Run-Time Adaptivity Methodology
  - Global Control Loop Adaptation technique
  - Local Control Loop Adaptation technique

Introduction

- Service Center Energy Efficiency Issues
  - Environmental challenge
    - Service Center energy consumption will double in next 5 years [1] [2]
    - Worldwide service centers CO_2 emissions are equivalent to half of the total commercial airlines' CO_2 [3]
    - Service center carbon emissions greater than of Argentina and Netherlands

- Problems to be addressed
  - Service center facility cooling
  - Inefficient service center resources utilization
    - ~30% of servers in a service center are unused at any given time [4]

Service Center Energy Awareness Model

Motivation

WHY IS A CONTEXT MODEL NEEDED? → **Dynamic Heterogeneous Environment**

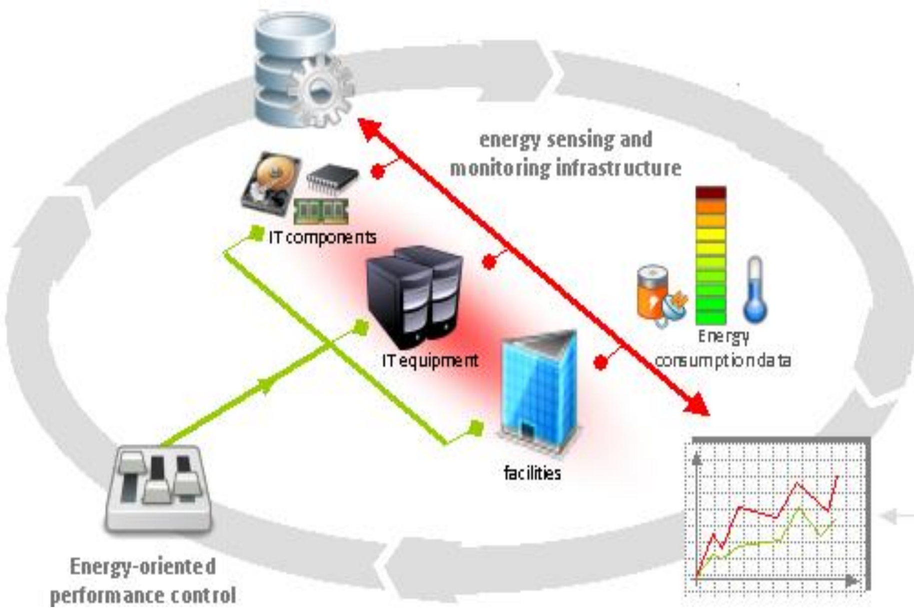

- **Solution**
  - **EACM model - RAP model extension**

Service Center Energy Awareness Model
EACM model

- RAP context model CTX = <R,A,P>
- Context Resources
  - IT Computing Resources
  - IT Facility
  - Activities
- Context Actions
  - Consolidation actions
  - DPM Actions
- Context Policies
  - GPI/KPI policies
  - Low level policies

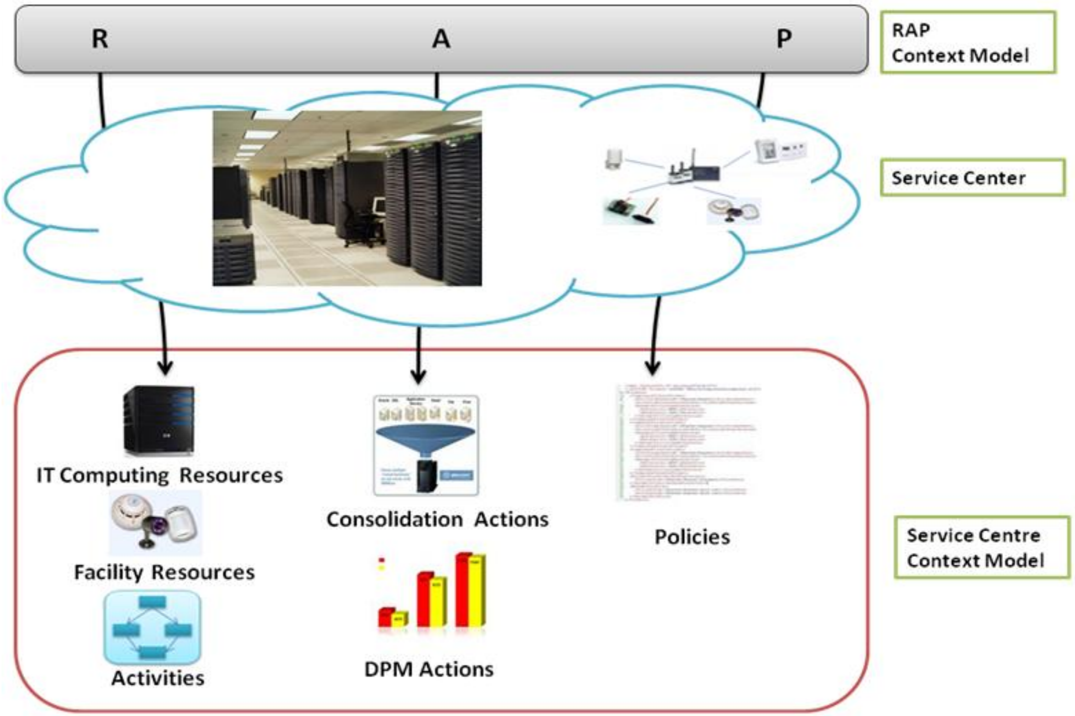

Service Center Energy Awareness Model

EACM model processing

Policies evaluation

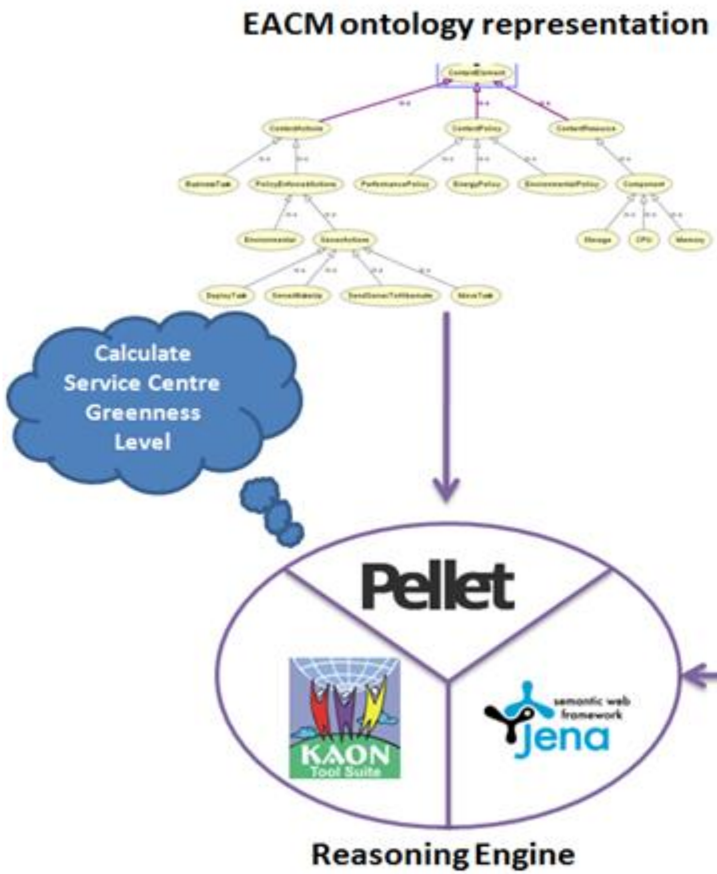

SWRL Rules

```
Server(?ServerRef)
^ name(?ServerRef, "Server1")
^ associatedComponent(?ServerRef, ?CpuRef)
^ CPU(?CpuRef)
^ associatedCore(?CpuRef, ?CoreRef)
^ used(?CoreRef, ?UsedCore)
^ swrlb:greaterThan(?UsedCore, 2400)
^ swrlb:lessThan(?UsedCore, 2600)
^ associatedComponent(?ServerRef, ?MemoryRef)
^ Memory(?MemoryRef)
^ used(?MemoryRef, ?UsedMemory)
^ swrlb:greaterThan(?UsedMemory, 1200)
^ swrlb:lessThan(?UsedMemory, 1600)
^ associatedComponent(?ServerRef, ?StorageRef)
^ Storage(?StorageRef)
^ used(?StorageRef, ?UsedStorage)
^ swrlb:greaterThan(?UsedStorage, 180)
^ swrlb:lessThan(?UsedStorage, 220)
-> respected(EnergyPolicy_1, true)
```

GPI/KPI Policy	Accepted values
EnergyPolicy_ 1	$CPU \in (2400\ MHz, 2600\ MHz)$ $MEM \in (1200\ Mb, 1600\ Mb)$ $HDD \in (180\ Gb, 220\ Gb)$

Energy Aware Run-Time Adaptivity Methodology
Global Control Loop Adaptation Technique

MAPE loop based approach

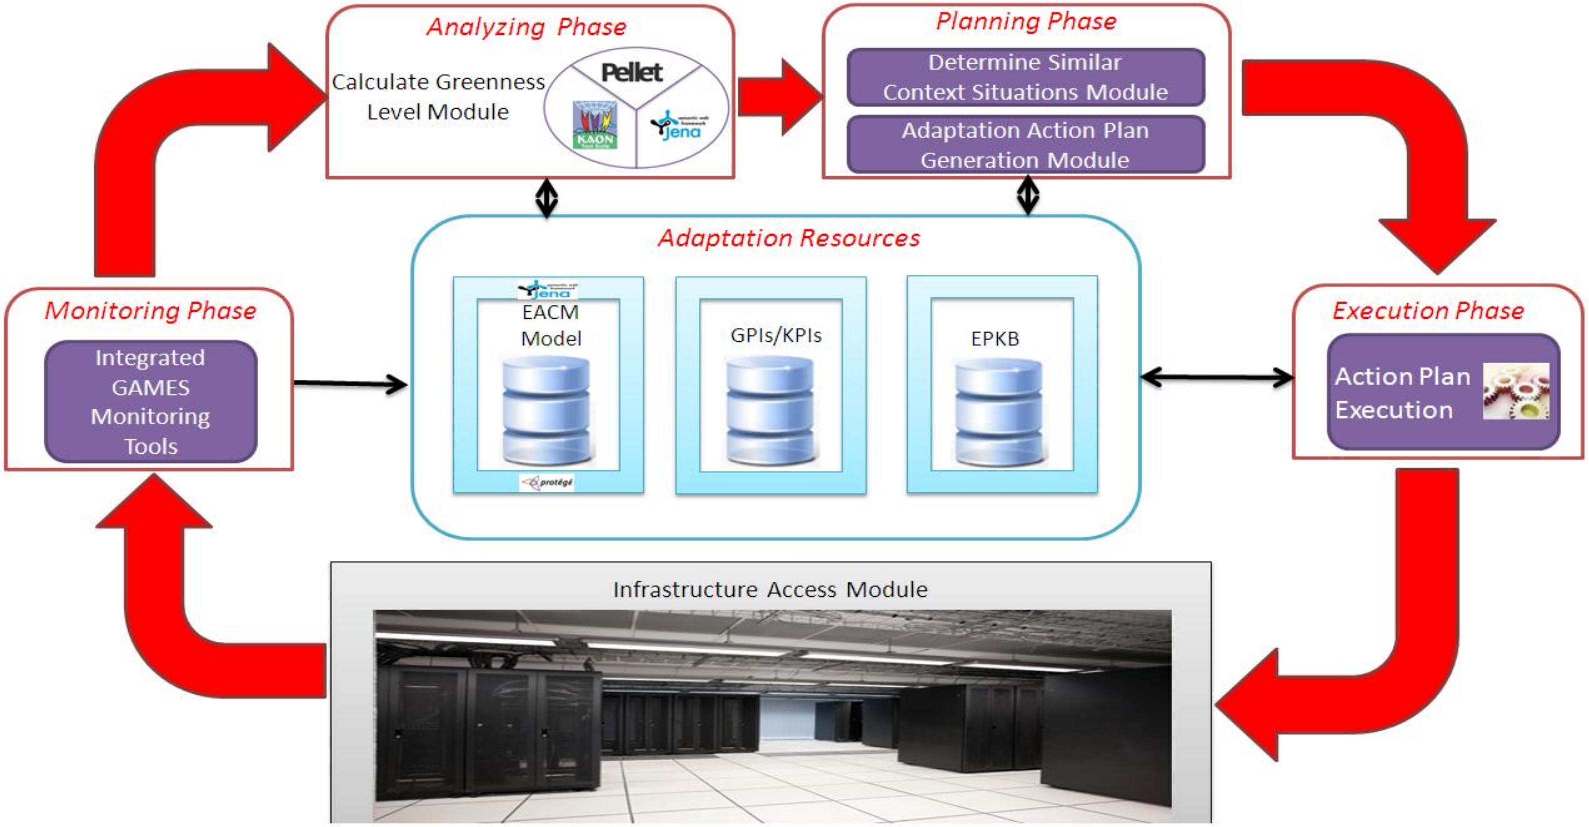

Energy Aware Run-Time Adaptivity Methodology

Global Control Loop – Analysis Phase

- **Service Center Greenness Level Evaluation**
  - EACM model instance entropy

ENTROPY Calculation Formula

$$E_s = \sum_n w_n * EVAL(P_n)$$

W_n - weight of the GPI/KPI policy n

P_n - GPI/KPI policy n

Entropy based adaptation decision

setup entropy threshold value T_e
loop
  calculate model instance entropy E_s
  if ($E_s > T_e$) **fire** adaptation process
  else acceptable greenness level
endloop

Energy Aware Run-Time Adaptivity Methodology

Global Control Loop – Planning Phase Reinforcement Learning

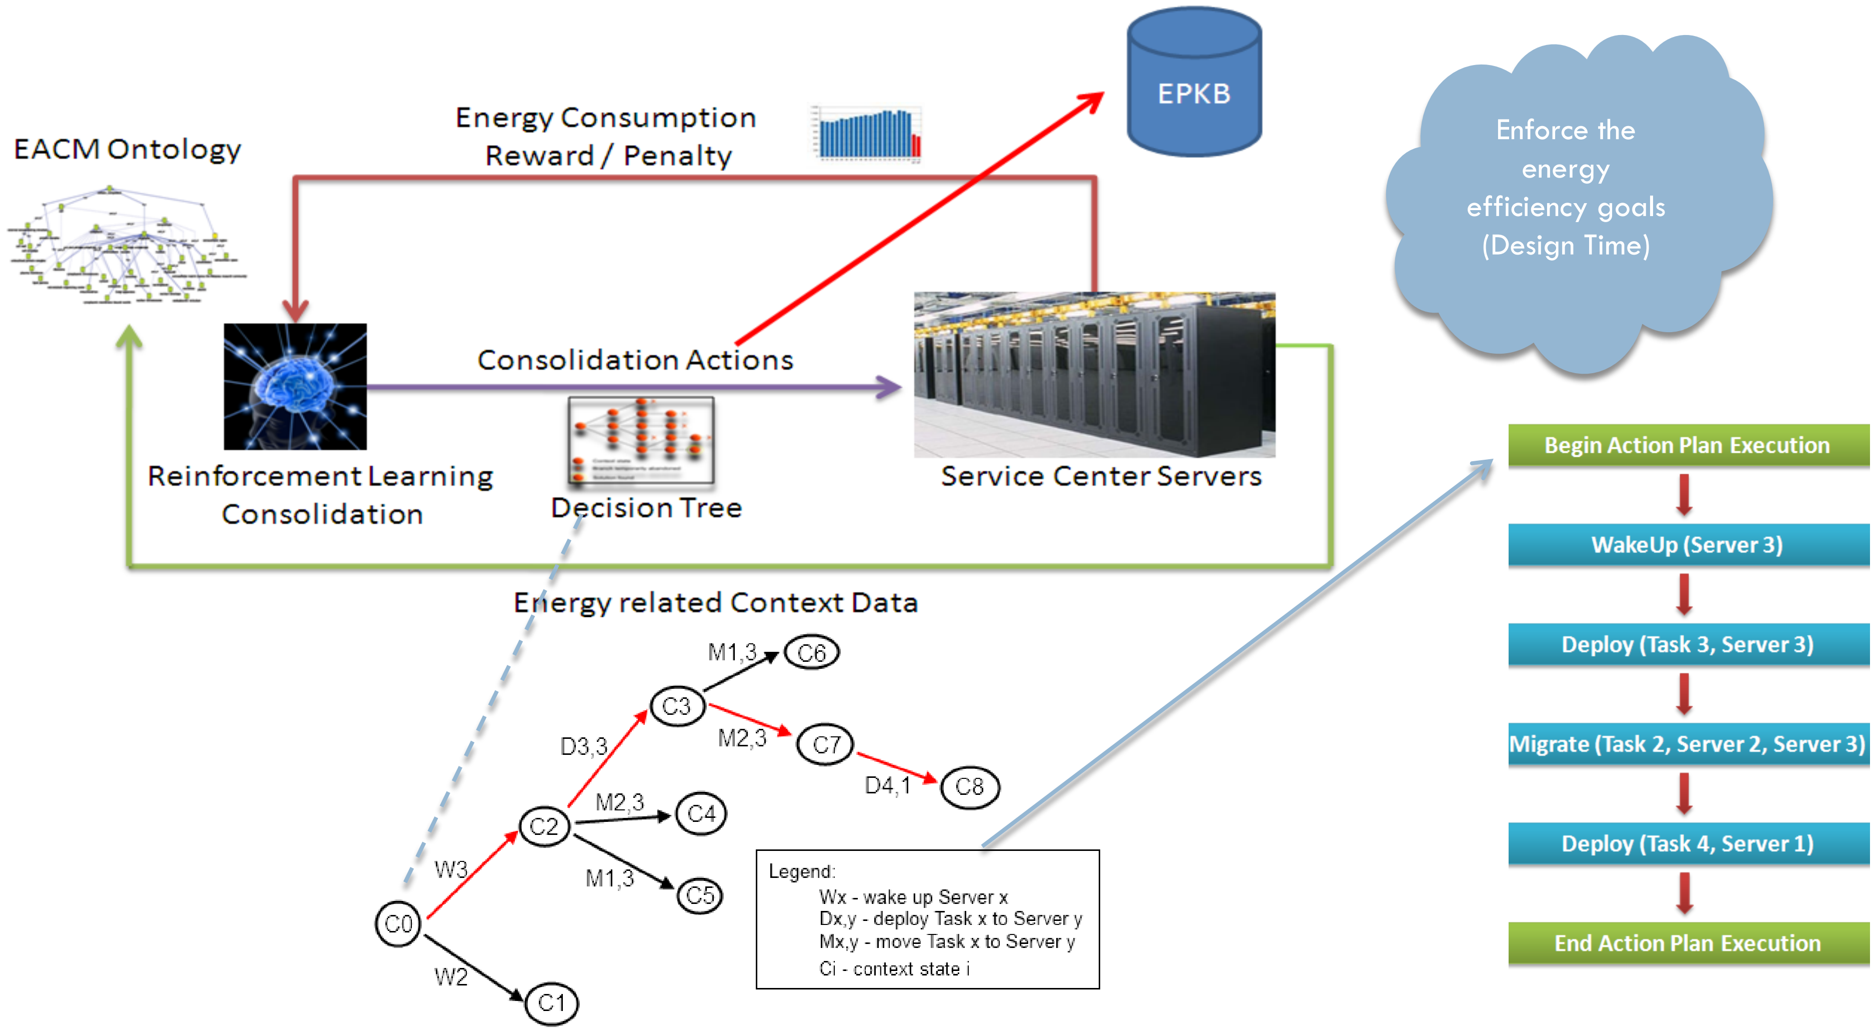

Energy Aware Run-Time Adaptivity Methodology

Global Control Loop case study scenario

- Services pool

Resource name / Service Type	CPU		Memory	
	Importance level [0..1]	Requested cores x [min, max] MHZ	Importance level [0..1]	Requested [min, max] MB
Service 1 (IO Intensive)	0.5	1 x [100, 500]	0.5	[100, 500]
Service 2 (Memory Intensive)	0.2	2 x [100, 200]	0.8	[2000,3000]
Service 3 (CPU Intensive)	0.8	4 x [1500, 2000]	0.2	[100, 500]
Service 4 (Matrix Multiplication)	0.6	3 x [100, 300]	0.4	[500, 1000]
Service 5 (Sort)	0.7	1 x [700, 1500]	0.3	[200, 500]
Service 6 (Search)	0.6	1 x [1000, 1700]	0.4	[400, 800]

- Generated test case workload schedule

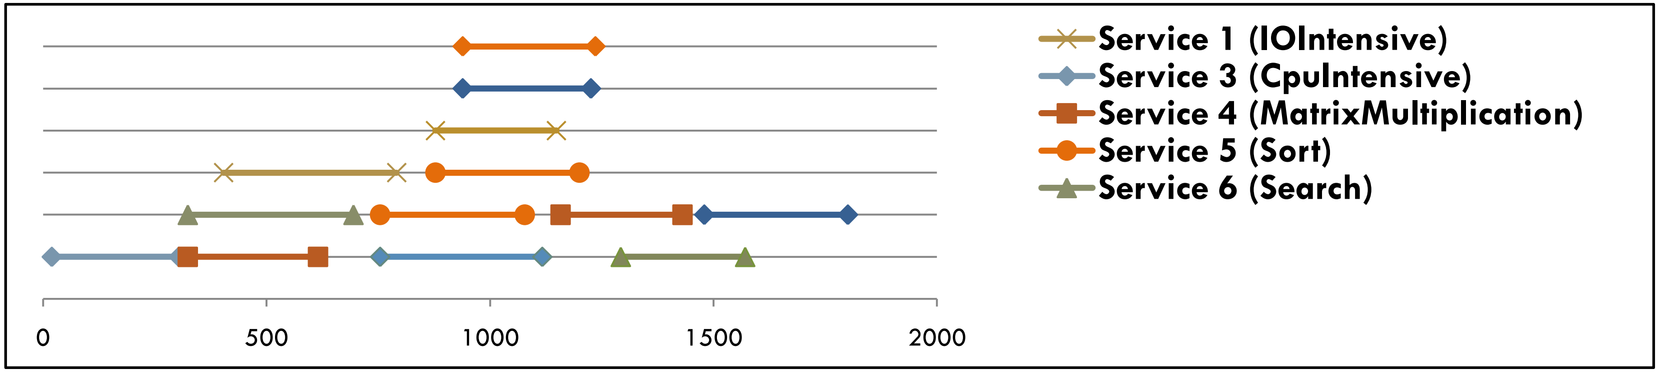

Energy Aware Run-Time Adaptivity Methodology

Global Control Loop case study and experiments

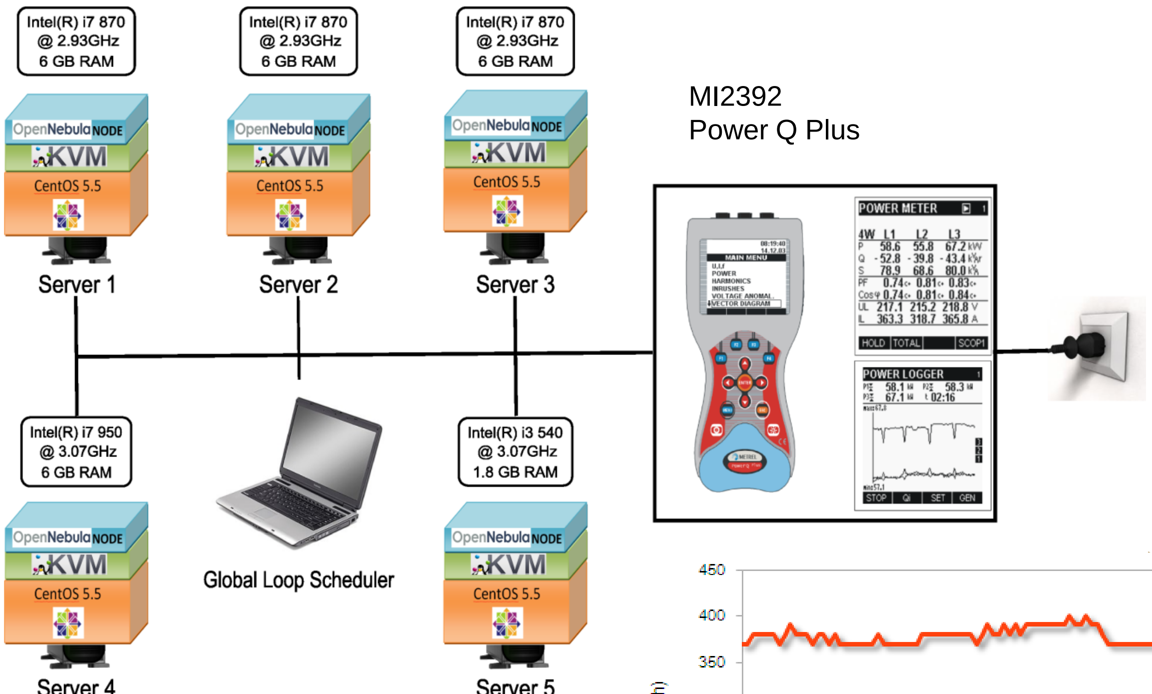

Average Energy Consumption (Service Center Level)

- Dynamic workload generation
- Execution time: 30 min
- GAMES Controller: **133.6Wh**
- OpenNebula Scheduler: **385.66Wh**
- Energy saving: 65.35%

Overhead (measured)

Deploy virtual machine: 5-10Wh
Migrate virtual machine :15-20Wh
Delete virtual machine : 10Wh
OpenNebula Client:100-110Wh
Global Control Loop : 20-30Wh

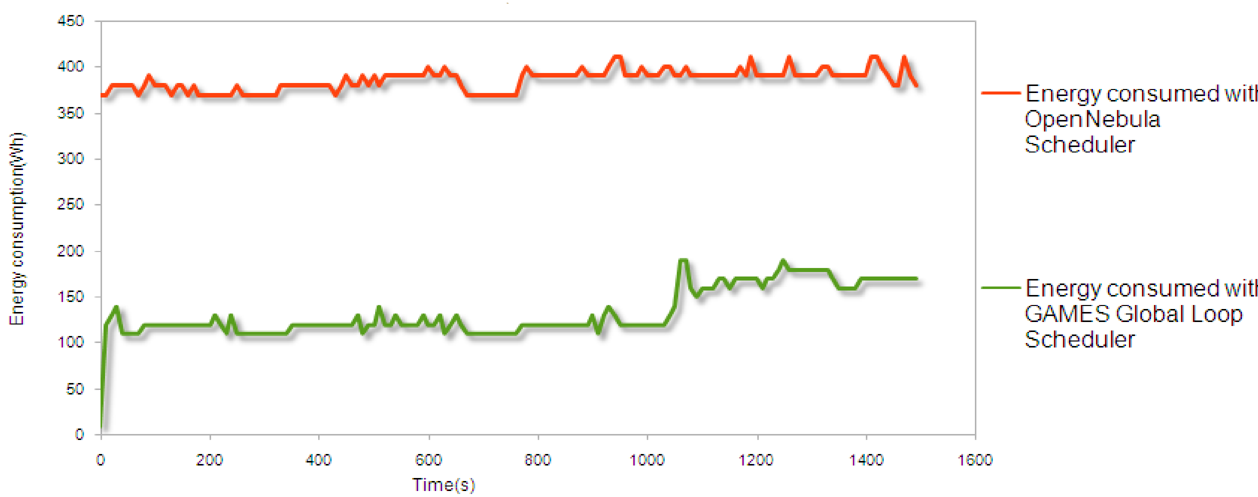

Energy Aware Run-Time Adaptivity Methodology

Local Control Loop Adaptation Technique

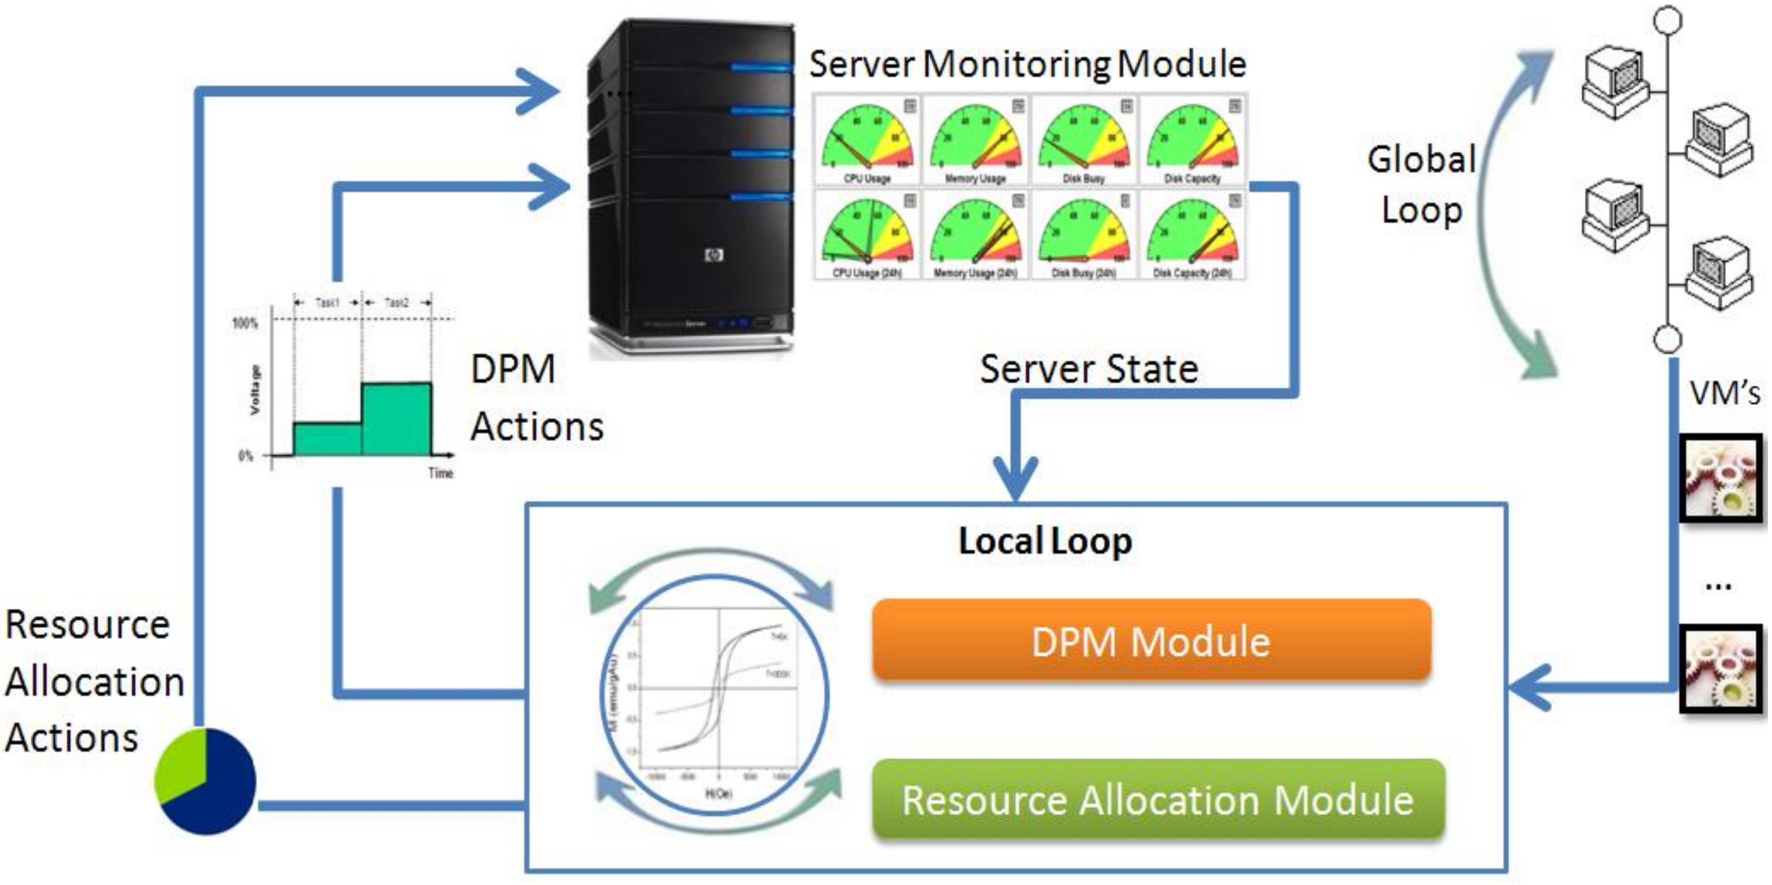

Energy Aware Run-Time Adaptivity Methodology
Global Control Loop – Analysis Phase

Processor adaptation technique
- p-states dynamic management
- Fuzzy logic adaptation technique
  - CPU workload
  - CPU current usage level (p state)
- Features:
  - Noise filtering
  - Progressively adapt to changes

$$fuzzyControl_T = ALPHA * fh(cpuPState, cpuWorkload) + BETA * fl(cpuPState, cpuWorkload) + GAMMA * fm(cpuPState, cpuWorkload) + DELTA * fuzzyControl_{T-1}$$

HDD adaptation technique
- Keep HDDs in low power states as long as possible
- Prediction technique
  - adaptive learning trees, agents, penalty/reward
  - identify sequences of workload followed by long idle periods

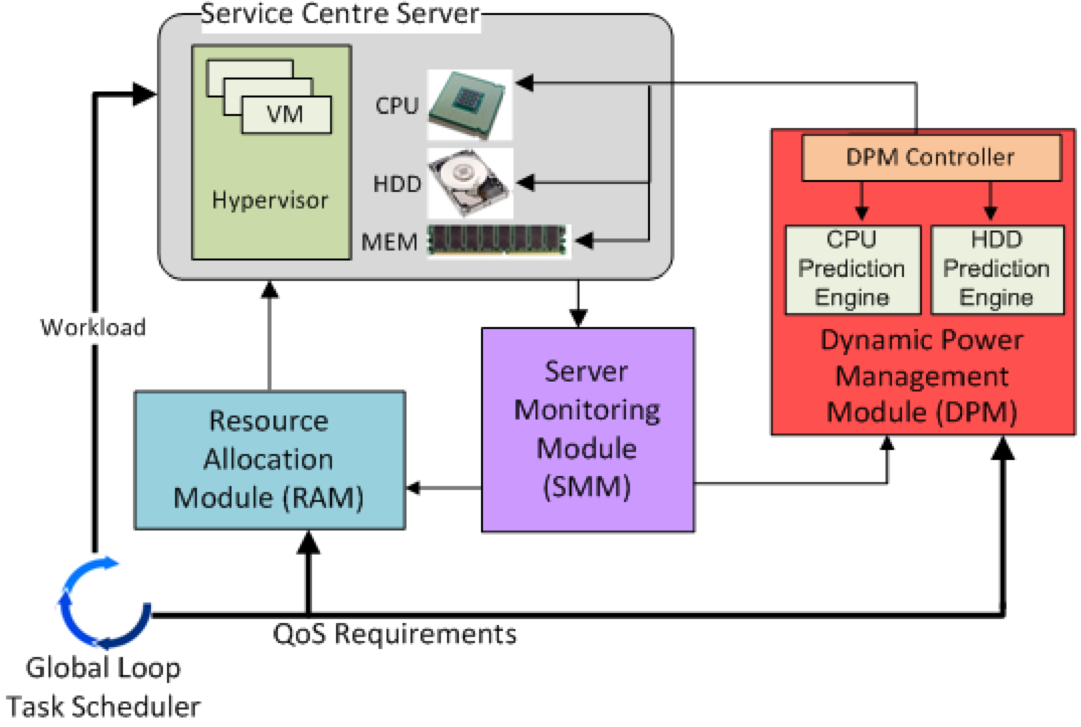

Energy Aware Run-Time Adaptivity Methodology

Local Control Loop Adaptation Technique Test Case 1

Test bed Server
Intel Core 2 CPU E6600@2.40GHz
P-states: p-0 (66% frq), p-1 (100% frq)

Test bed Workload
VM1: 2 Virtual Proc – Medium CPU usage
VM2: 1 Virtual Proc – Alternate High/Low CPU usage

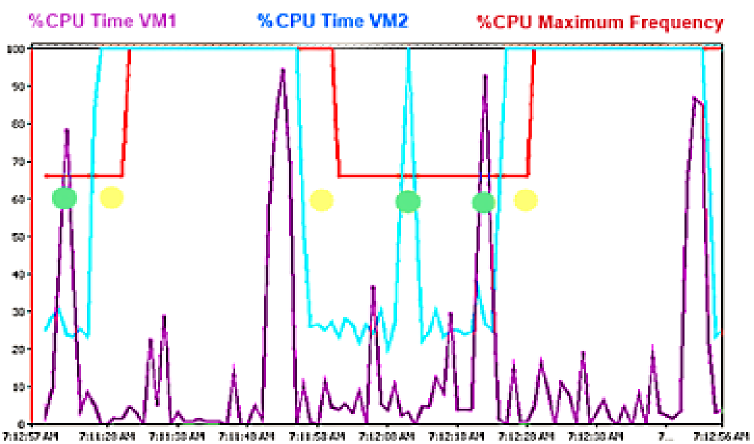

Noise filtering (green points) and progressive adapting to changes

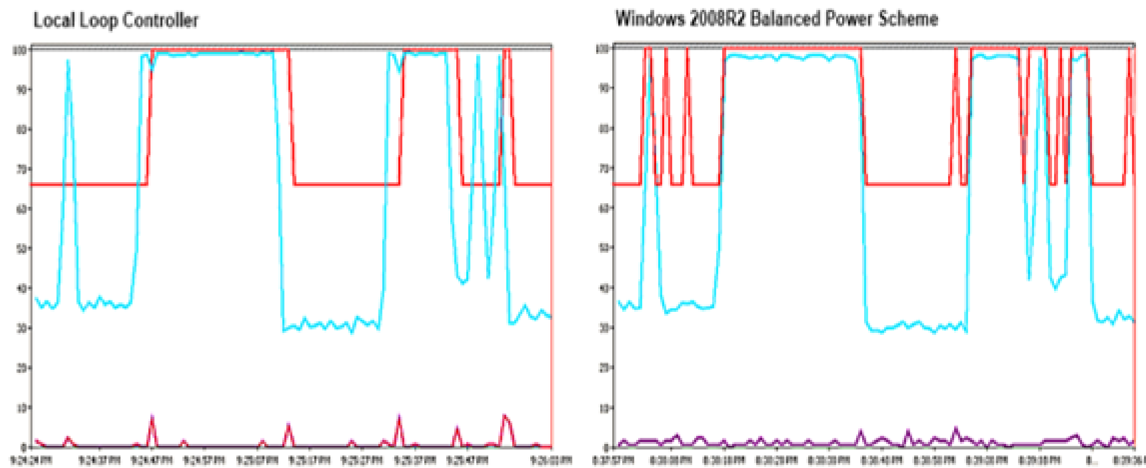

Local Loop Controller vs. Microsoft Balance Power Scheme

Results (LLC vs. BPS for the same workload)

Spikes filtering
- BPS – aggressive processor transitions (workload spikes determine p-state changes)
- BPS emphasis on performance not on energy efficiency

Low power states processor
- MS Balance Power Scheme: 23%
- Local Loop Controller : 39%
- Improvement in resource usage: 16%
- Performance degradation: 4%

Energy Aware Run-Time Adaptivity Methodology
Local Control Loop Adaptation Technique Test Case 2

Test bed Server
Intel® I7 870 @ 2.93 GHz

14 P-states (1.2**,** 1.33, 1.46, 1.6, 1.73, 1.86, 2.0, 2.13, 2.26, 2.39, 2.53, 2.66, 2.79, 2.93 GHz)

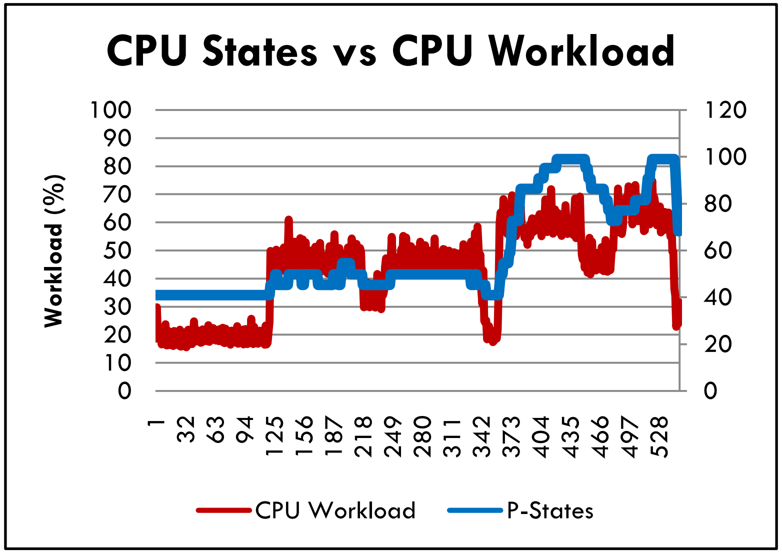

Test CPU Workload

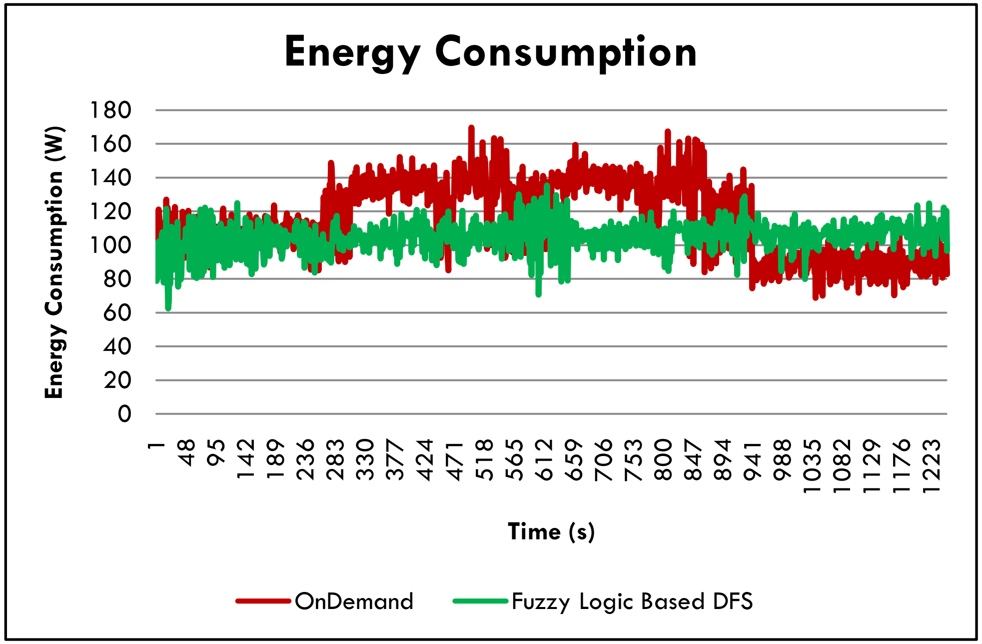

Local Loop Controller vs. CentOS On Demand Power Scheme

Results (LLC vs. On Demand for the same workload)
Energy Efficiency Improvement: 10%
Performance degradation: 15%

Conclusions

- Service center energy awareness model
  - GPI/KPI evaluation by means of reasoning
- Service center level adaptation technique
  - 65.35% estimated energy saving
    - **Note:** IT computing resources energy consumption aprox. 40% from total service center energy consumption
- Server level adaptation technique
  - Server resource usage improvement 16%
  - Performance degradation 4%
  - Energy Efficiency Improvement about 10%

- *Thank You!*

- *Contact*
    - *Distributed Systems Research Laboratory* www.dsrl.coned.utcluj.ro
    - *Daniel Moldovan* daniel.moldovan@cs.utcluj.ro

References

1. U.S. Environmental Protection Agency, ENERGY STAR Program, Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431 (2007)
2. Kaplan, J.M., Forrest, W., Kindler, N.: Revolutioning Data Center Energy Efficiency, McKinsey & Company, Technical Report (2008)
3. M. Poniatowski: Foundations of Green IT: Consolidation, Virtualization, Efficiency, and ROI in the Service center, Prentice Hall, 2009, ISBN-10: 0137043759, ISBN-13: 9780137043750, 978-0137043750, pages xiii
4. Uddin, M., Rahman, A.: Server Consolidation: An Approach to Make Data Centers Energy Efficient & Green. Int. Journal of Scientific & Engineering Research, Vol. 1, Issue 1 (2010).
5. M. Salehie, L. Tahvildari: Self-adaptive software: Landscape and research challenges, ACM Tran. on Aut. and Adaptive Systems, ISSN:1556-4665 (2009).
6. Khargharia, B., Hariri, S., Yousif, M.: Autonomic power and performance management for computing systems, Cluster Computing, Vol. 11 (2), pp. 167 – 181, ISSN:1386-7857 (2008)