6-29-2018

A Comparison Study of Different Algorithms for Energy-Aware Placement of Virtual Machines

Dr.Fatos Xhafa Professor

3rd IISA-2018 Conference, Hong Kong, China

A Comparison Study of Different Algorithms for Energy-Aware Placement of Virtual Machines

29th June 2018

Alejandro Olvera, Fatos Xhafa

Department of Computer Science, UPC, BarcelonaTech, Spain
WEB: http://www.cs.upc.edu/~fatos/
Email: fatos@cs.upc.edu

Agenda of this talk

- **Energy Consumption in Cloud Data Centers**
- **Heuristics for Optimising Energy Consumption in Cloud Data Centers.**
- **Experimental Study, Simulation and Evaluation**
- **Conclusions and Outlook**

Energy Consumption in Cloud Data Centers

Cloud Data Centers: Key IT infrastructure

- Delivering computing resources on-demand over the Internet.
- Hundreds of thousand of servers worldwide and increasing...
- High energy consumption and CO2 emissions.
  - A datacenter can generate 170 million tons of CO2 per year.
  - Excessive heat discharge leads to the reduced lifetime of IT infrastructure.
- Virtualization: Virtual Machines (VMs) –each VM runs in an independent OS and applications.

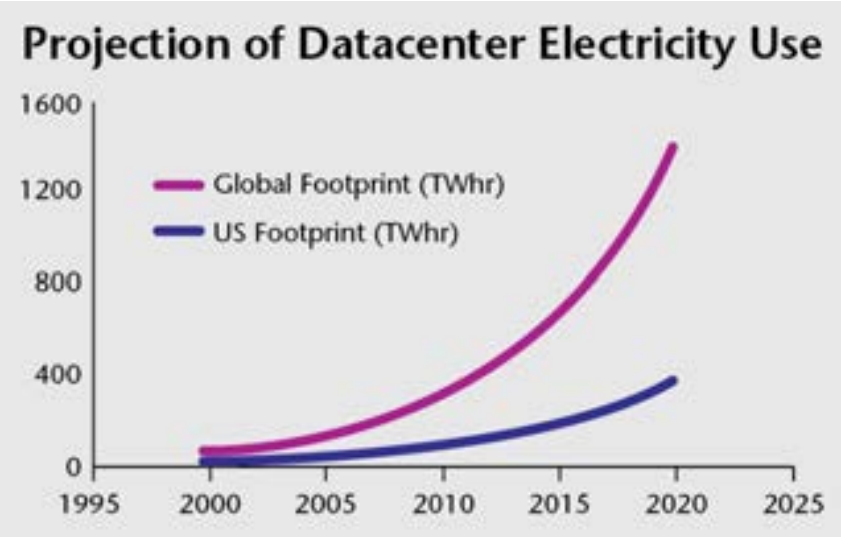

Three-Layer Data Center Architecture

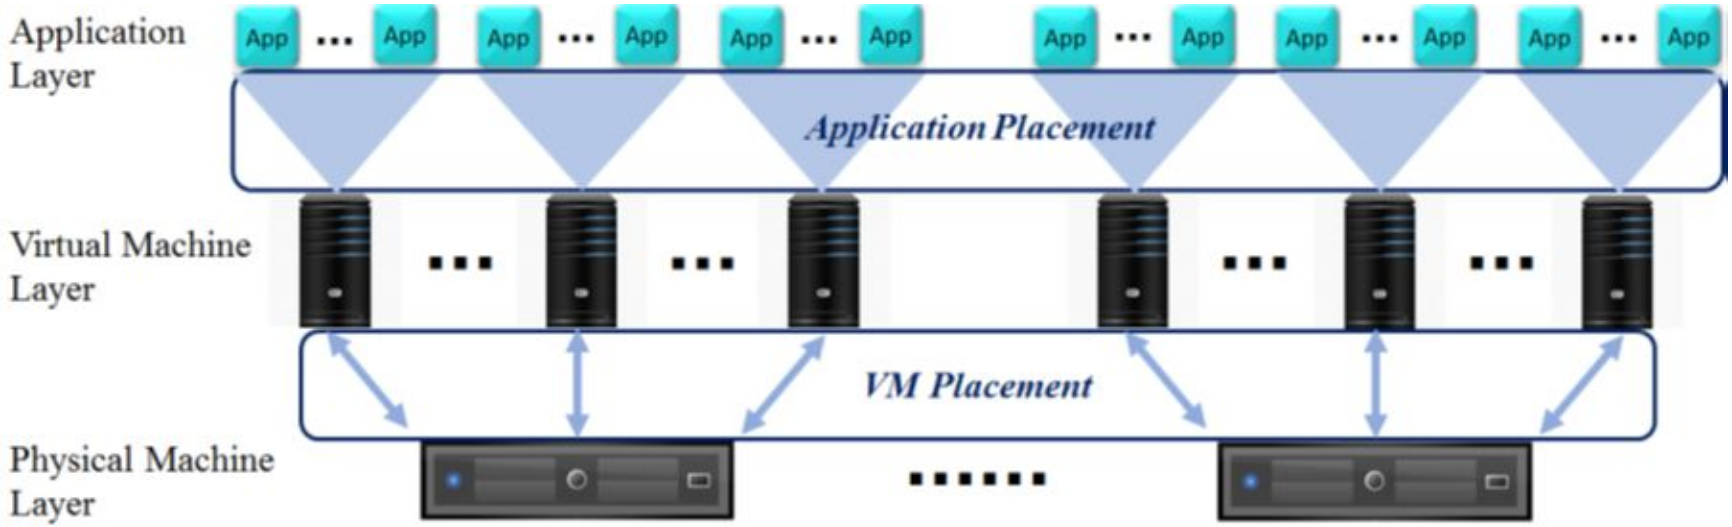

Application Layer: Assign incoming applications to VMs for execution.

VM Layer: VM Placement, migrations, executions, creation (i.e. management).

PM Layer: ON / OFF operations, sleep cycles, cooling and dynamic voltage frequency scaling (DVFS).

Energy Consumption: Problem formulation (I)

- Consider a cloud data center:
  - k physical machines (**servers**) $\{M_1, M_2, \ldots, M_k\}$
  - Each M_i has a number of **VMs** running: $\{VM_1^i, \ldots, VM_p^i\}$
  - **VMs** are characterised by a number of parameters (CPU, memory, network bandwidth, MIPS, etc.) –useful for task allocation. ***Analytics*** about VM's state are also known.
  - T**asks/applications** are submitted by users independently.
  - Tasks/applications are allocated to VMs.
  - VMs can be **migrated** from one M_i to another M_j.
  - **SLA requirements** (Service Level Agreement) should be satisfied (or violated as little as possible).
- Assign tasks/applications to VMs (under VM migration) so as to **minimise energy consumption** of the data center.

Energy Consumption: Bin packing formulation

$$\text{minimize} \quad z = \sum_{i=1}^{n} y_i \qquad w_j = \textit{weight} \text{ of item } j,$$

$$c = \textit{capacity} \text{ of each bin,}$$

$$\text{subject to} \quad \sum_{j=1}^{n} w_j x_{ij} \le c y_i, \quad i \in N = \{1, \ldots, n\},$$

$$\sum_{i=1}^{n} x_{ij} = 1, \quad j \in N,$$

$$y_i = 0 \text{ or } 1, \quad i \in N,$$

$$x_{ij} = 0 \text{ or } 1, \quad i \in N, j \in N,$$

$$y_i = \begin{cases} 1 & \text{if bin } i \text{ is used;} \\ 0 & \text{otherwise,} \end{cases}$$

$$x_{ij} = \begin{cases} 1 & \text{if item } j \text{ is assigned to bin } i; \\ 0 & \text{otherwise.} \end{cases}$$

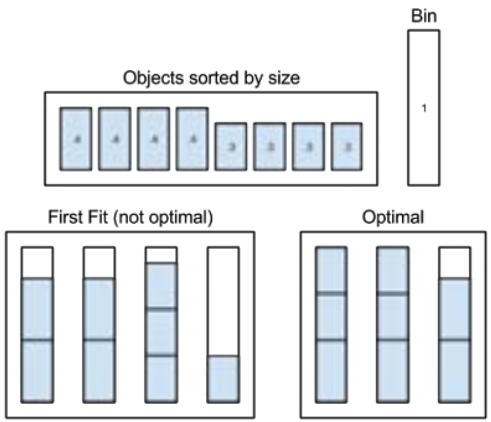

Assign each *item* to one *bin* that the total *weight* of *items* in each *bin* does not exceed *c* and the number of *bins* is minimum.

Application Assignment
weight = MIPS, RAM, CPUs...
item = application
bin = VM

VM Placement
weight = MIPS, RAM, storage, bw..
item = VM
bin = host

Optimisation Process: The Full Picture

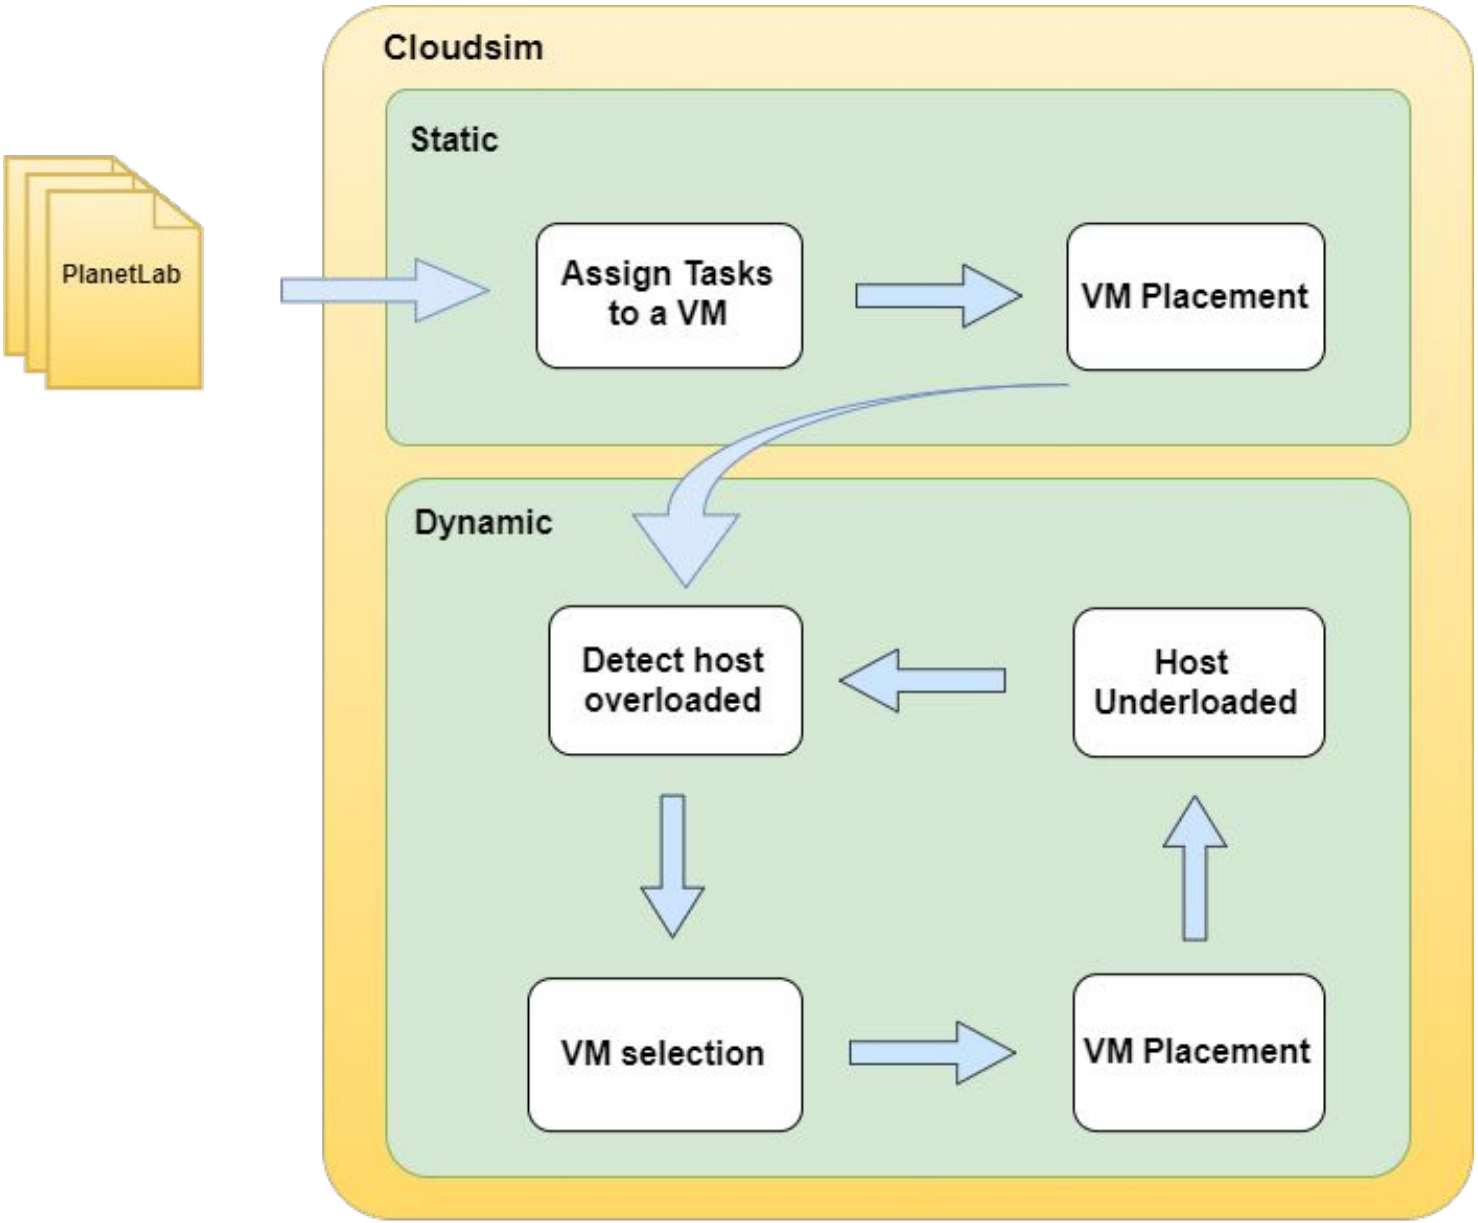

Optimisation Process: Static Case

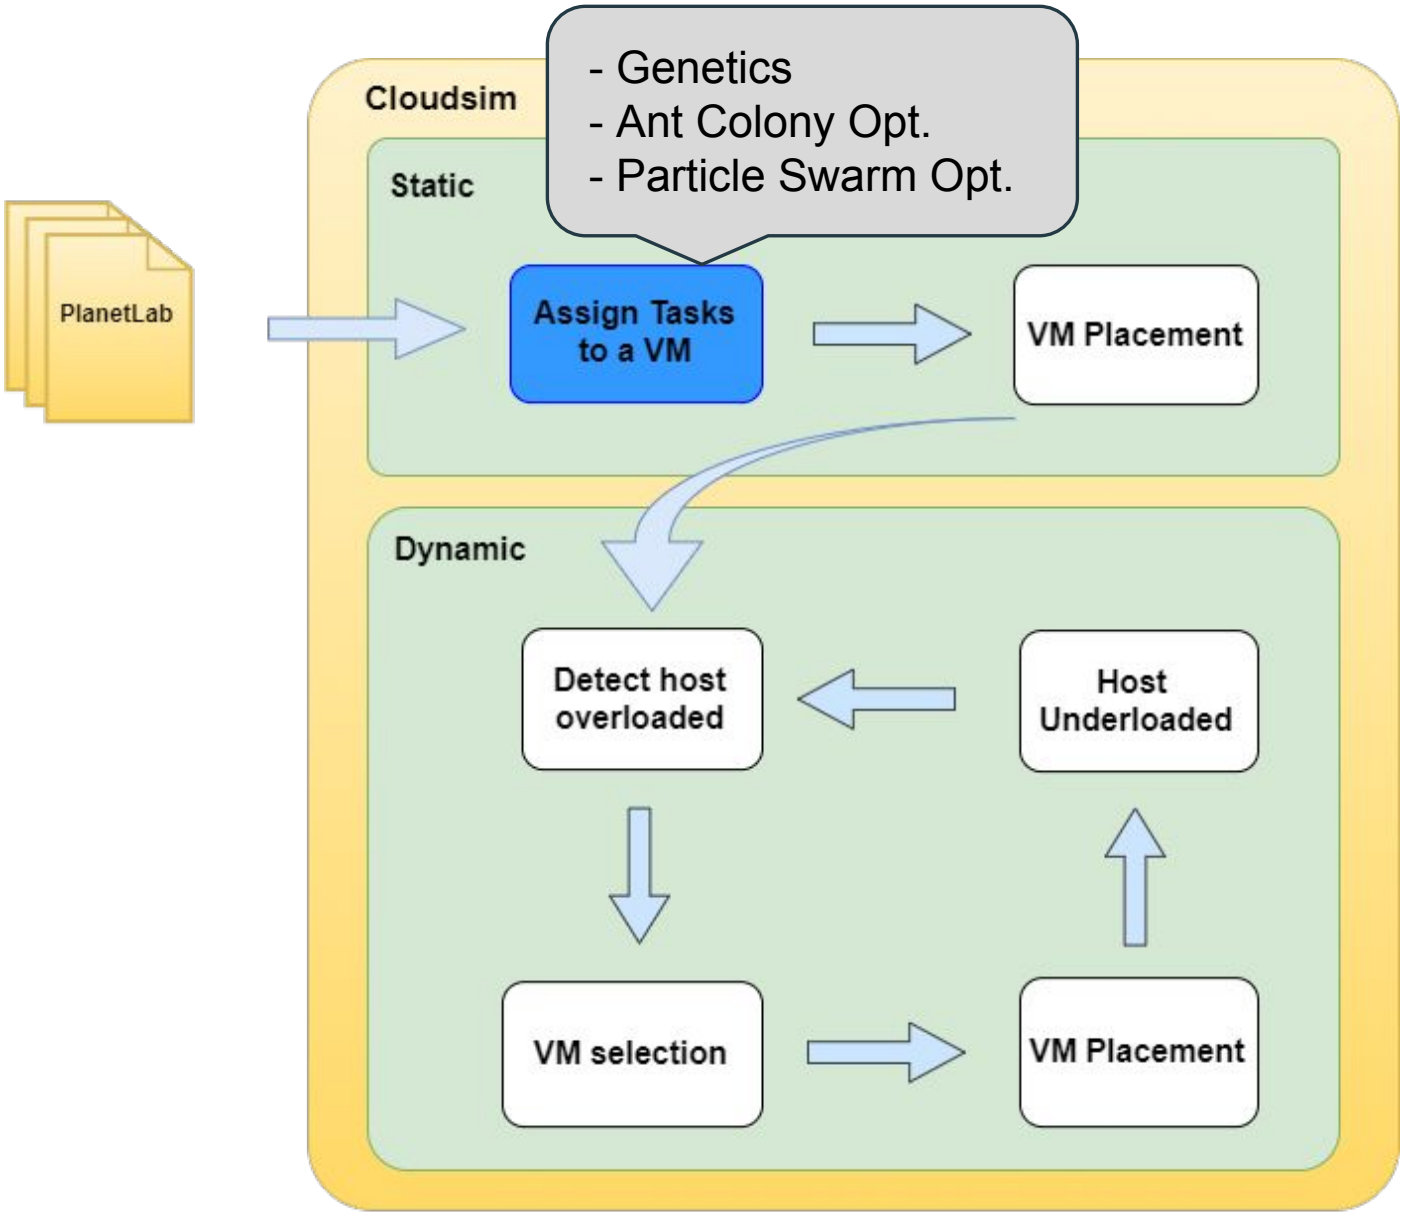

Static Case: Repairing Genetic Algorithm (RGA)

- **RGA**: minimize the energy consumption and makespan by maximizing resource utilization.
- **Key features**:
  - **Chromosome**: a vector of length = number of applications (i.e. each position is the application id 1, 2,...), and each position value is the VM associated to this application.
  - **Initial population**:
    - Longest Cloudlet Faster Processor (LCFP): the longest applications to the VMs with more MIPS capacity.
  - **Infeasible-Solution Repairing Procedure**: convert infeasible solutions to feasible solutions.
    - **Violated VM**: one of its resource constraints surpassed;
    - **Move applications to other VMs** w.r.t resource constraints.
    - If all VMs has been solved, the infeasible solution is feasible.
  - Rest of ingredients...
- Faster and better convergence than standard GA

Optimisation Process: VM Placement Heuristics

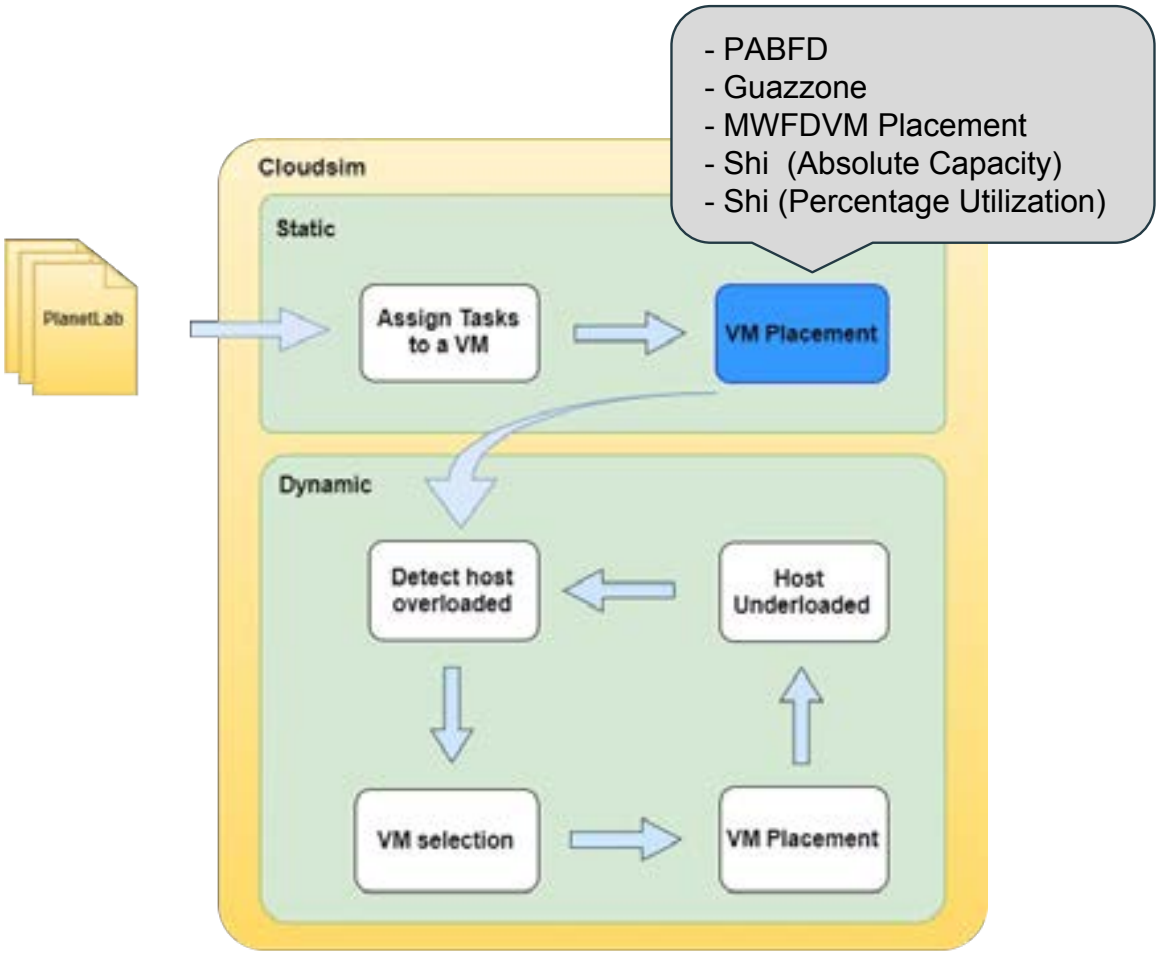

- PABFD: Power Aware Best Fit Decreasing -the default CloudSim algorithm.
- Guazzone *et al.*: Best-Fit Decreasing strategy.
- MWFDVP: Modified Worst Fit Decreasing VM Placement.

Optimisation Process: Dynamic Case

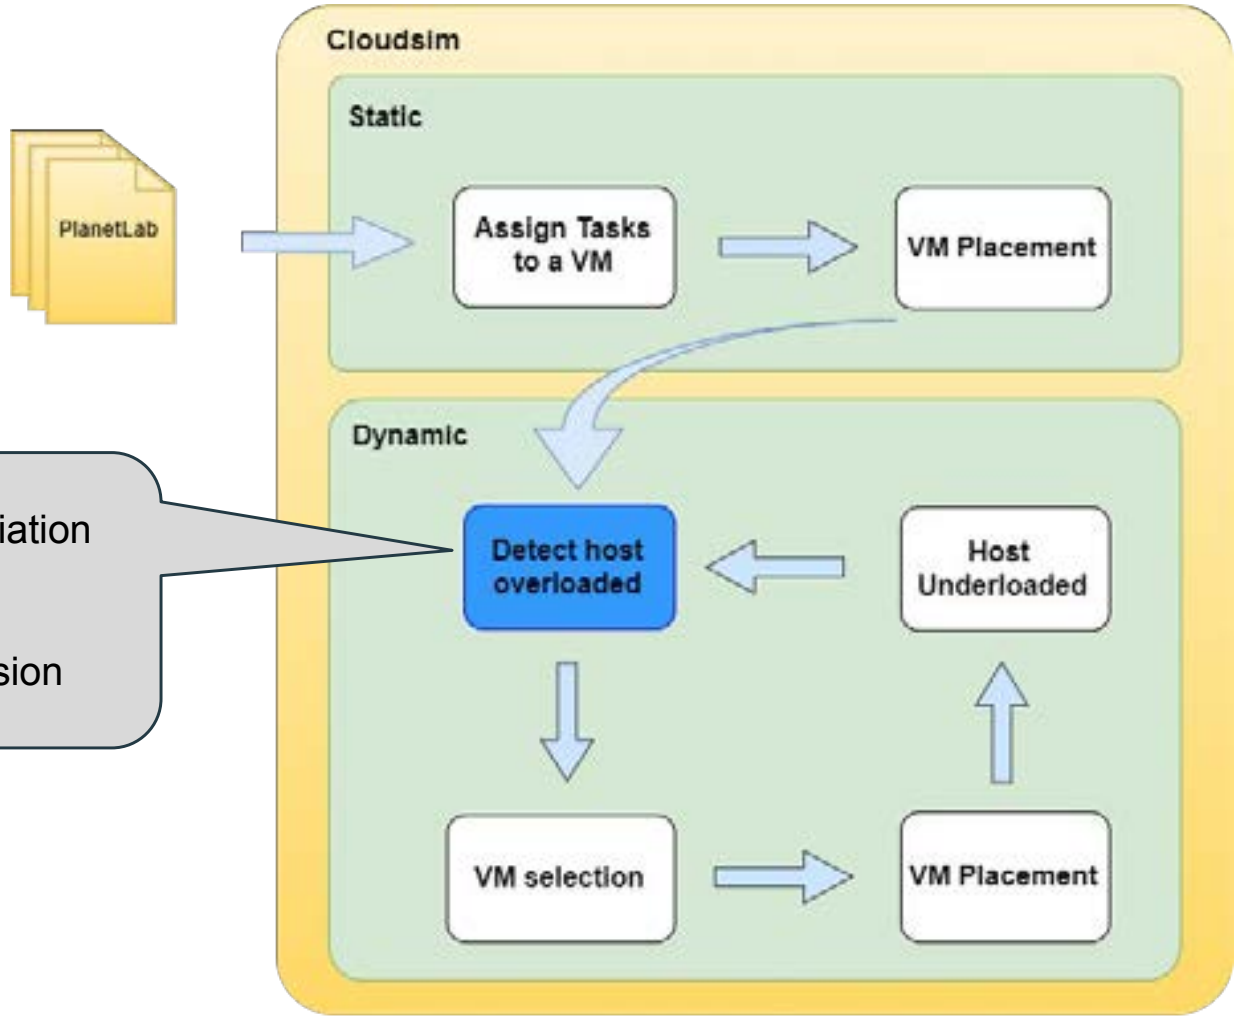

- Median Absolute Deviation
- Interquartile Range
- **Local Regression**
- Robust Local Regression

Optimisation Process: Dynamic Case (cont'd)

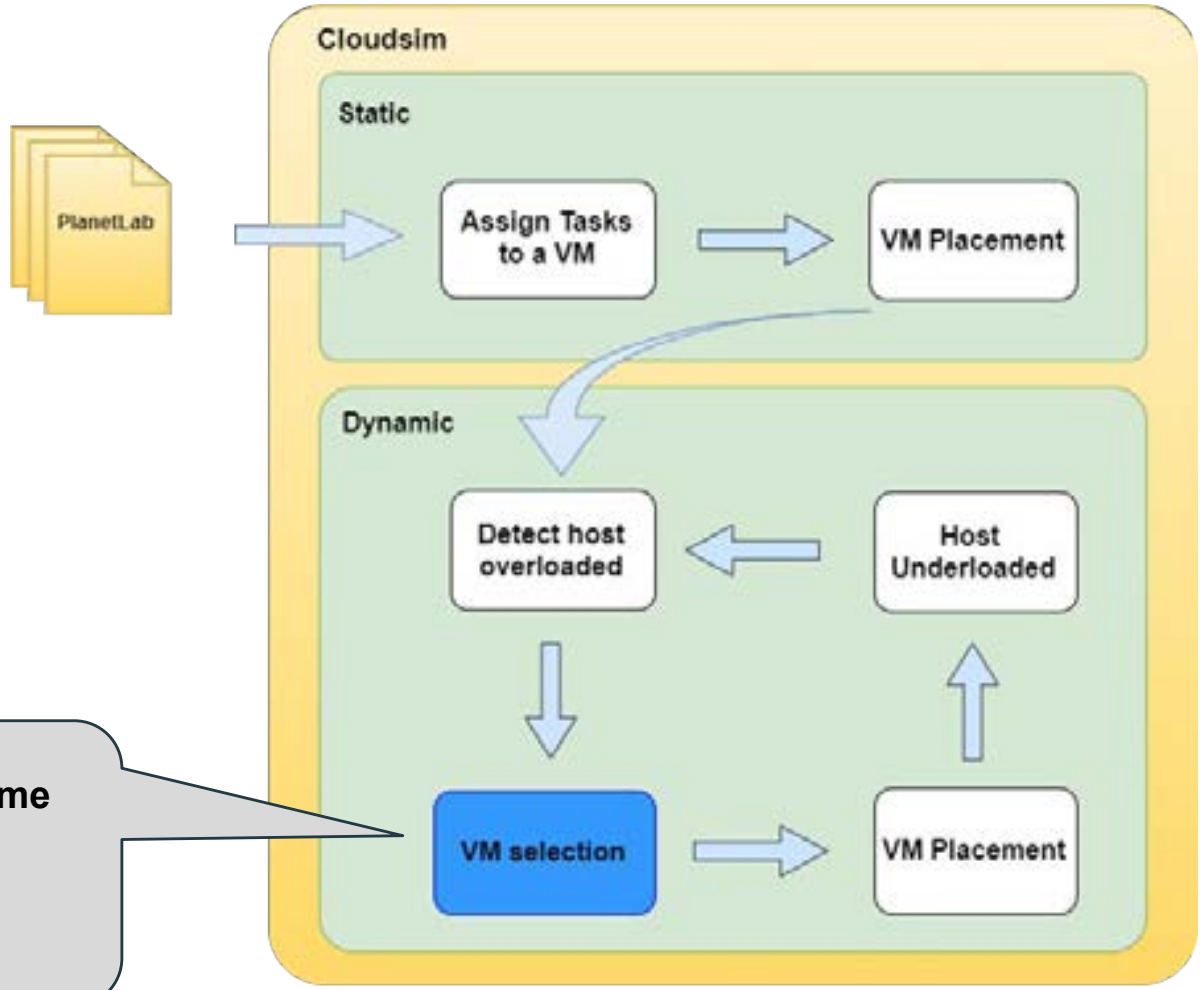

- **Minimum Migration Time**
- Random Choice
- Maximum Correlation
- Minimum Utilization

Optimisation Process: Dynamic Case (cont'd)

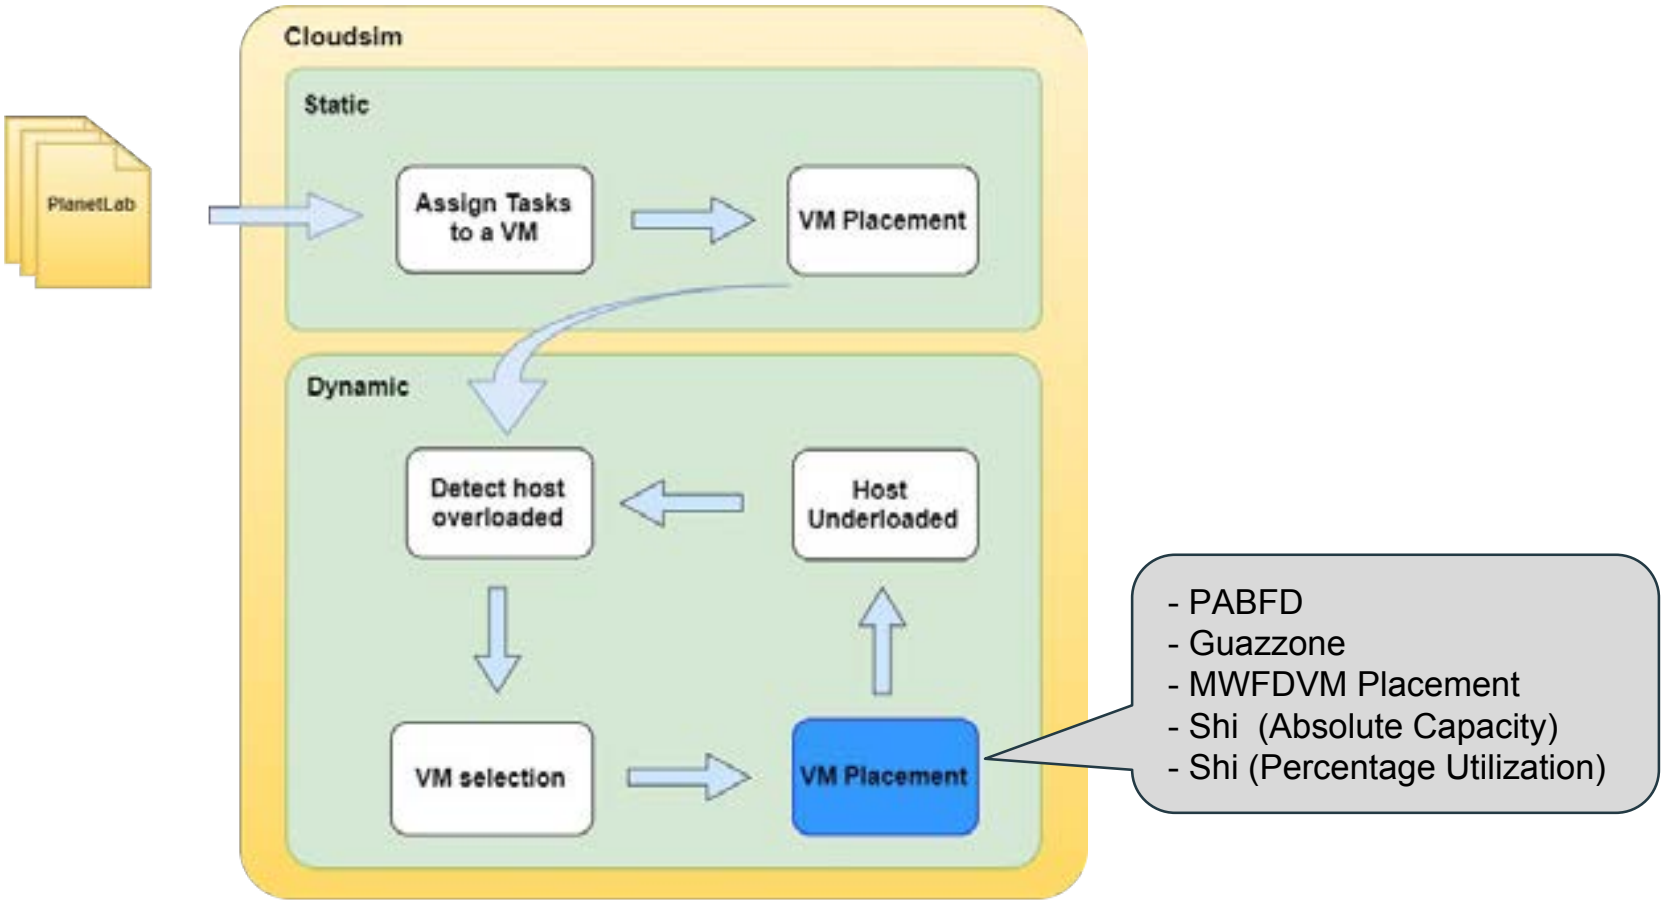

Optimisation Process: Dynamic Case (cont'd)

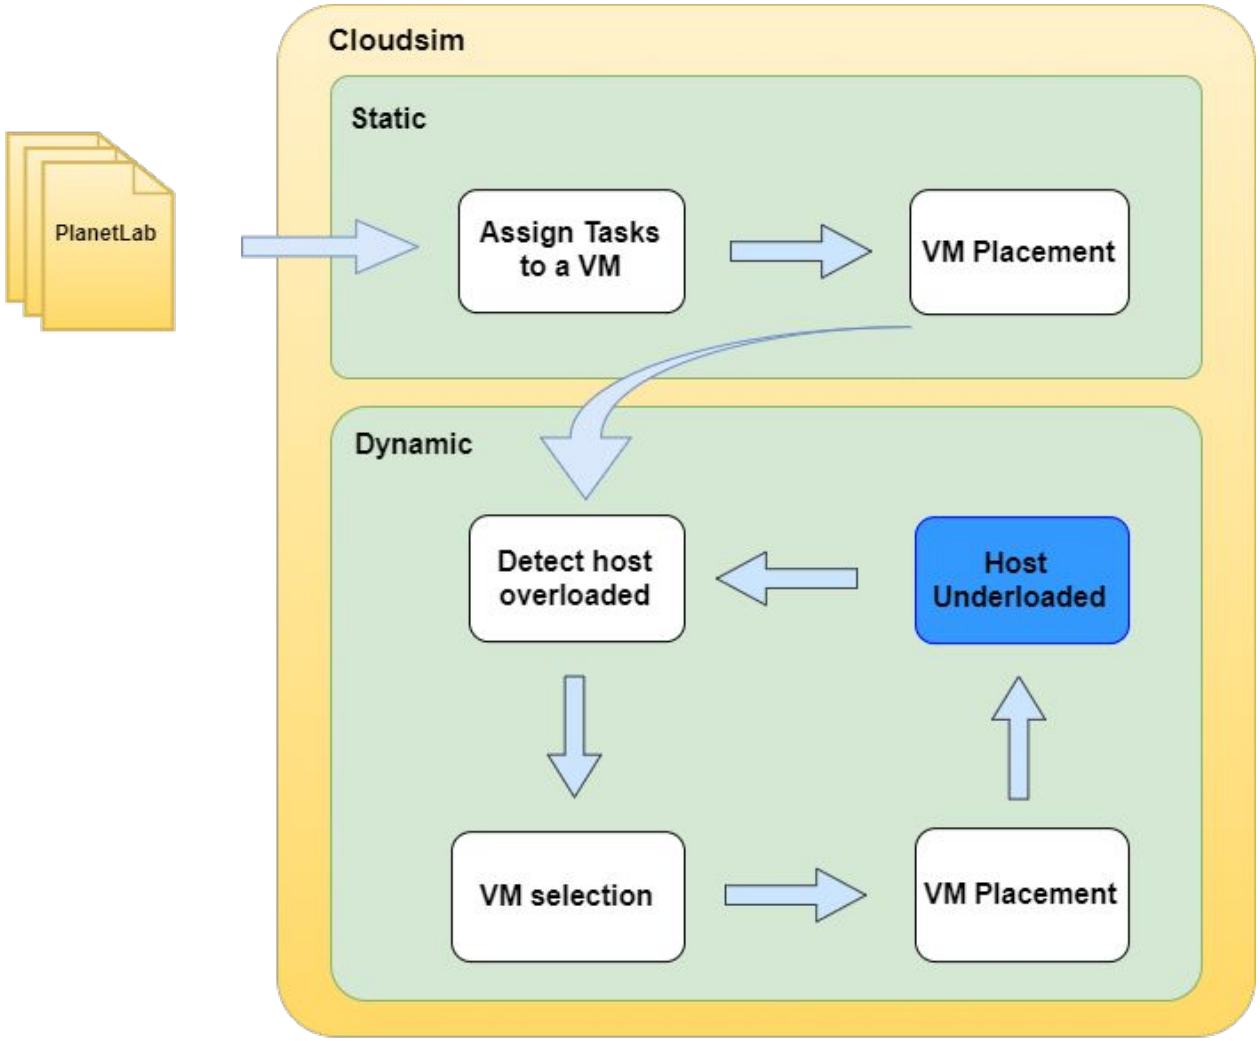

- Default CloudSim algorithm.

Experimental Design: Data sets and CloudSim

- Cloud Computing Dataset: **PlanetLab files**
- **CloudSim**:
  - CloudSim is a **data center simulator** described as a modelling, simulation and experimentation tool-kit for clouds, data centers and VMs.
  - To setup **cloud environments to test applications performance**.
  - Different **cloud deployment models** (public, private, hybrid), different types of servers, VMs and applications, scheduling policies, etc.
  - **New implementations** via extension of its classes.
  - It **avoids low-level details in power consumption model** (e.g. it doesn't include the power consumption used by switches and routers.)

Experimental Design: CloudSim Power Consumption Model

- CloudSim considers mainly the **CPU usage** for calculating the **power consumption**.
  - Each server type has its own power model consumption.
  - In increments of 10% of CPU usage, power consumption changes.
  - Linear interpolation between two energy consumptions variations.
- CloudSim: **Quality of Service**
  - Performance degradation due to VMs migrations (PDM metrics)
  - **SLA** based on response time
  - **SLATAH** (SLA violation per Time of Active Host)

Conclusions and Outlook

- **Energy Consumption** in Cloud Data Centers is a critical problem! It is a cross-layer problem at the Cloud stack...
- As an optimisation problem at application layer (i.e. assign incoming applications to VMs for execution):
  - It is very complex due to **multi-objective & conflicting criteria**.
  - **CloudSim** can be used to implement different algorithms in the optimization process of energy consumption (**single objective**).
  - **VMs migration** is an interesting technique to optimise energy consumption but goes in detriment of QoS / SLA
  - **Repairing GA** showed efficiency and good quality.
  - Other algorithms showed a variety of performance in different scenarios (from empty to fully overloaded data centers).
- **Future**: Variety of Cloud Models + **Prediction**

UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH

WORK BASED ON

- **MSc Thesis**:

Alejandro Olvera Anton, Implementation and Evaluation of Profile-based Prediction for Energy consumption in a Cloud Platform, FIB, UPC, Barcelona, 2017

Interested to publish in International Book Series?

- **Lecture Notes in Data Engineering and Communication Technologies, Springer.**
http://www.springer.com/series/15362

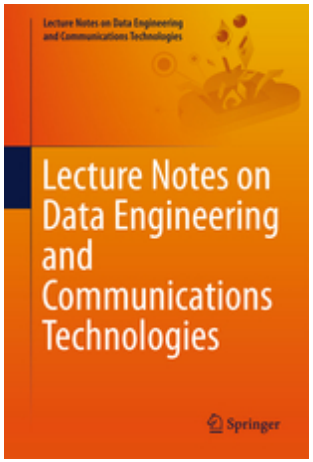

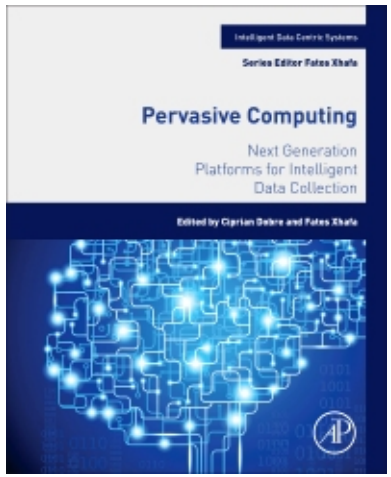

- **Intelligent Data Centric Systems, Elsevier**
https://www.elsevier.com/books/book-series/intelligent-data-centric-systems-sensor-collected-intelligence

New Elsevier Journal

- **Internet of Things; Engineering Cyber Physical Human Systems.** https://www.journals.elsevier.com/internet-of-things

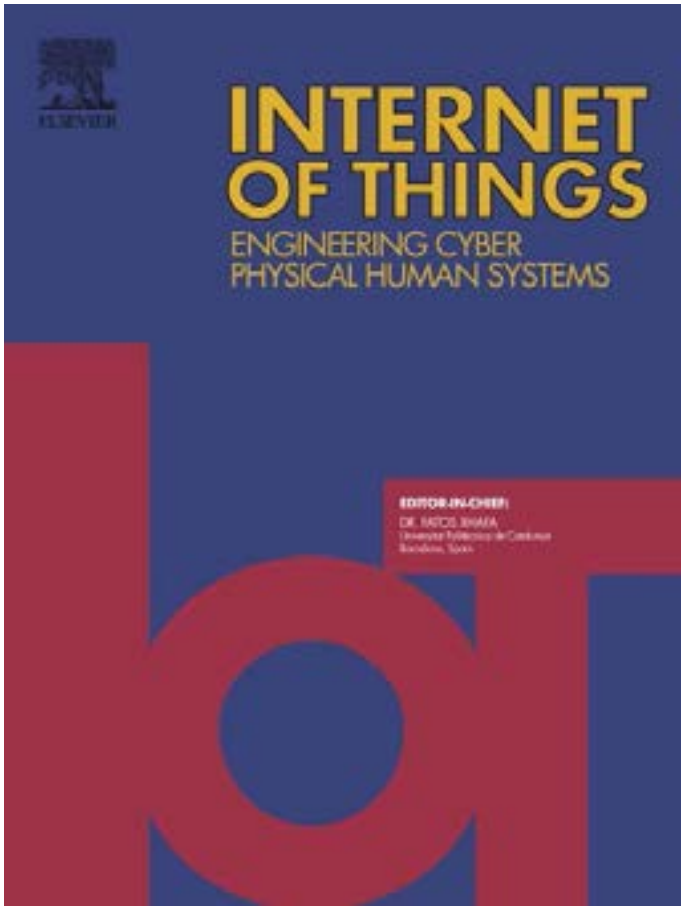

UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH

THANK YOU!

Q?

Email: fatos@cs.upc.edu

Web: http://www.cs.upc.edu/ fatos/

DBLP: https://dblp.uni-trier.de/pers/hd/x/Xhafa:Fatos