


Heuristic for resources allocation on utility computing infrastructures

João Nuno Silva, Luís Veiga, Paulo Ferreira
Distributed Systems Group – INESC-ID
Lisbon, Portugal

technology
from seed



- Lots of computation demanding tasks
 - Parameter sweep simulations, ...
 - Image rendering, ...
 - Everywhere ...
- Most of them are Bag of Tasks (BoT)
 - Algorithmically complex
 - Simple w.r.t. task coordination
 - Easily parallelizable

Bag-of-Tasks Problems (BoT)



technology
from seed

- Where to run them?
 - Single computer (bad performance)
 - Local Cluster (ownership / operation cost)
 - Grid (requires institutional membership)
 - Utility computer infrastructure
 - Mostly efficient
 - Cheap
 - Unrestricted public access

Utility computing (UC)



technology
from seed

- Provider owns clusters of computers
 - Public use
- Users launch virtual machines
 - On demand
 - No reservation mechanisms
- No complex contracts
 - Like utilities, user pays what he spends

BoT on Utility computing Amazon EC2



technology
from seed

- User picks a task distribution middleware
 - Install it on virtual appliance
- Configures data download mechanisms
- Launches virtual machines
 - Deployed on the utility computing infrastructure
- Waits for completion
 - Pay only for the execution time

BoT on Utility computing Amazon EC2

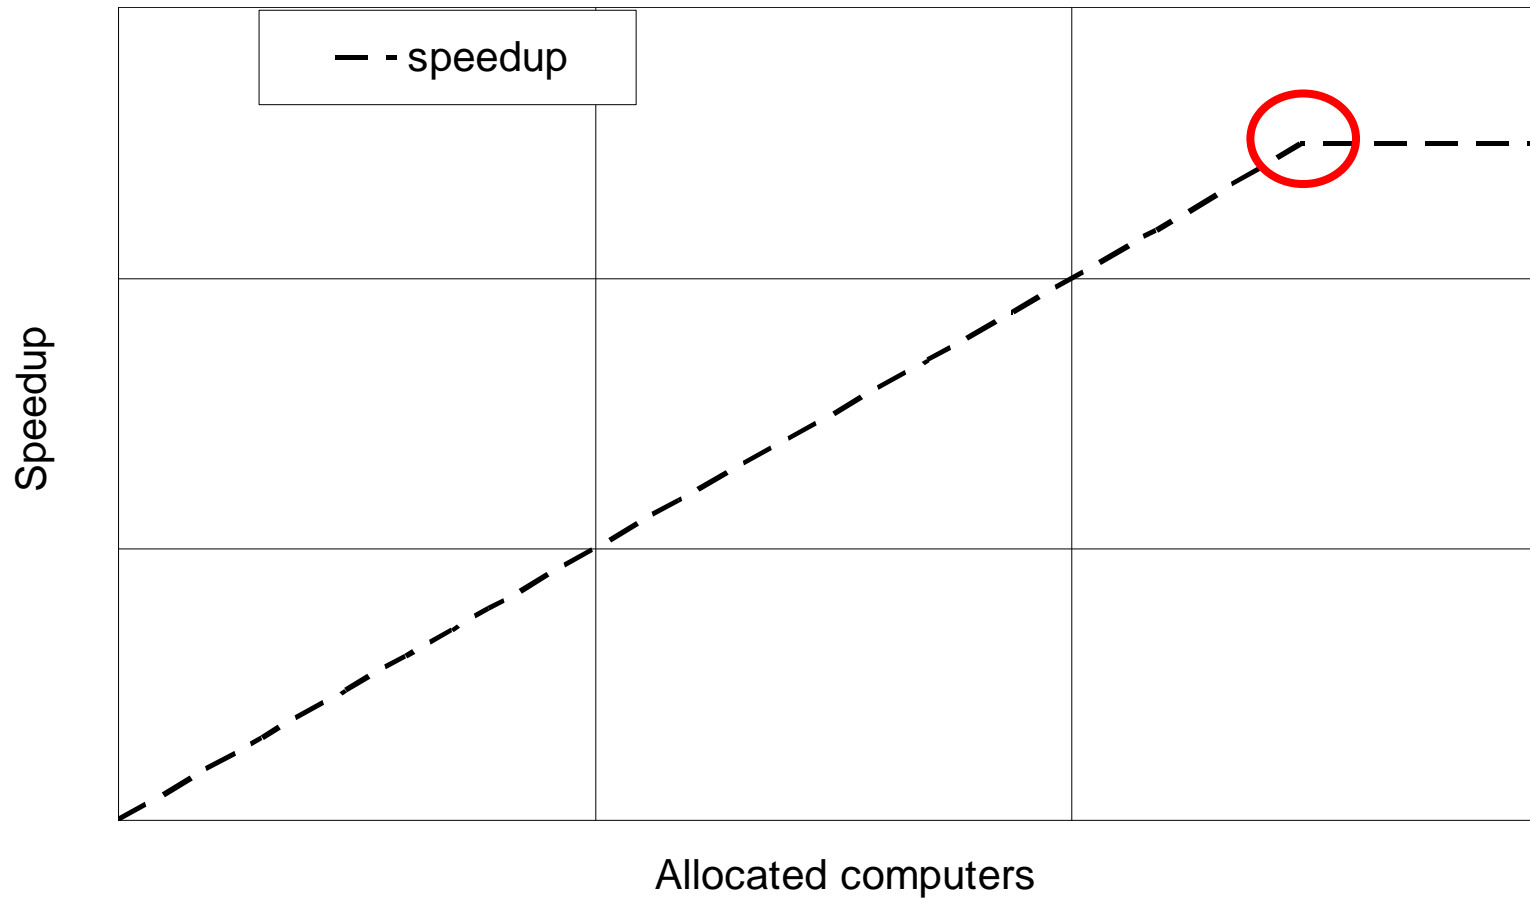


technology
from seed

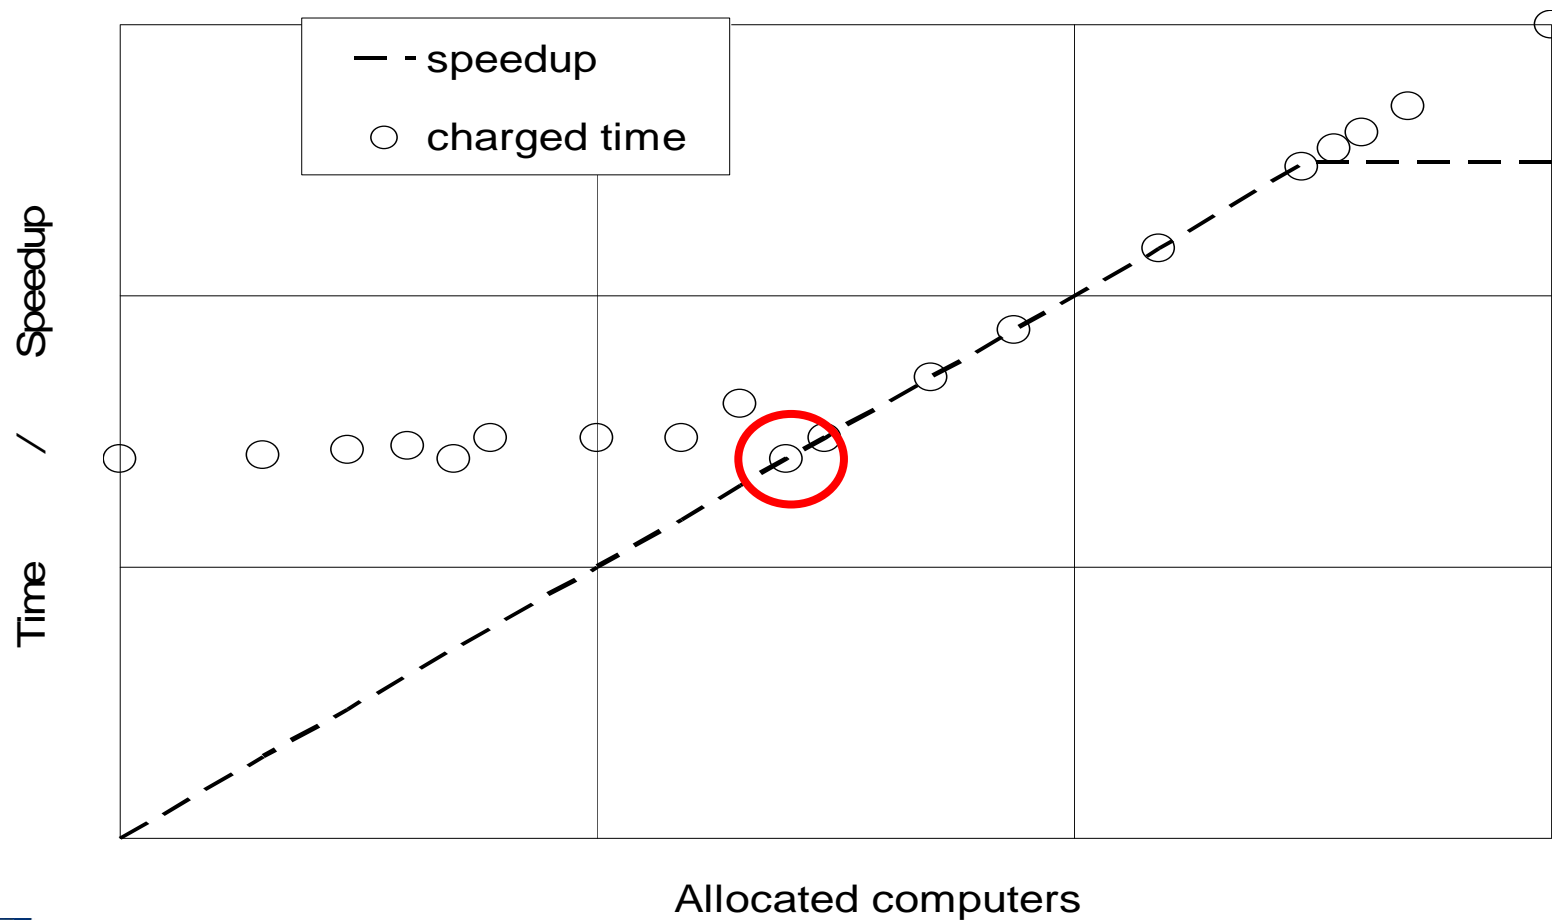
- 32 Tasks / 15m each / 8h total
 - User charged by the hour

# VM's	Cost (charge units)	Wall Time	Speedup
1	$1 \cdot 8h = 8$	8h	1
4	$4 \cdot 2h = 8$	2h	4
8	$8 \cdot 1h = 8$	1h	8
12	$12 \cdot 1h = 12$	45m	10.6
16	$16 \cdot 1h = 16$	30m	16
32	$32 \cdot 1h = 32$	15m	32
64	$64 \cdot 1h = 64$	15m	32

How many computers?



How many computers?



- Development of heuristic for
 - VM allocation on UC infrastructures
 - tailored to BoT problems
 - Reducing waste of charged VM time
 - Guaranteeing best speedups
 - For a fixed payment decided by the user

Resource management on Utility Computing



technology
from seed

- System centralized backend
 - Aimed at global efficiency
 - Physical CPU allocation, node idleness, network traffic, heat dissipation
- User frontend scheduling
 - Simple API to handle
 - VM allocation, accounting, monitoring
 - No optimization mechanisms

Shortcomings of Current Utility Computing



technology
from seed

- User frontend scheduling
 - Unable to determine appropriate/optimal
 - Number of VMs to allocate
 - Number of cores on each VM
 - Amount of memory
 - Network behavior
- Cost-efficient use of Utility Computing?
 - No framework, guidelines, principled approach

How many computers?



technology
from seed

- Should the user decide?
 - Execution times may not be known
 - The less knowledge the user has
 - More prone to error he is
- Automatic allocation
 - Run-time calculation
 - Guarantees:
 - optimal usage of the charged time

- User decides how much to pay
 - *As-if-serial* processing time
 - Minimum fixed cost
 - Relative urgency overhead
 - e.g. minimum+50%
- System guarantees:
 - Best speedups for that amount

- How to reduce wasted time

$$\#VMs = \frac{averageTime \times ntasks}{chargedTimeUnit}$$

- But *averageTime* not known till the end

Heuristic




technology
from seed

- Heuristic should use
 - Average time of finished tasks
 - nTasks not completed
 - Available time on Running hosts


$$\#VMs = \frac{\textit{averageTime} \times \textit{remainingTasks} - \textit{remainingTime}}{\textit{chargedTimeUnit}}$$

- If first tasks take too long?
 - Too many VMs are allocated
- If first tasks are too fast ?
 - Too few VMs are allocated
- Approach: select tasks randomly
 - Disperses locality complexity
 - Evens-out variations

- Tasks Average processing time




- Tasks Average processing time
- Predicted VMs





- Use of correction factor (creationRatio)
 - To reduce over allocation
 - On first estimations
 - Initially lower than 1
 - Fewer VMs than predicted are created
 - Converges to 1
 - At a given ratio (increaseRatio)

- Predicted VMs
- Corrected number of VMs





- Calculates needed host VMs
 - Initial corrections mask variability
- Creates host VMs
 - If more hosts are needed
- Turns hosts VMs off
 - Before increases in payment

- Allocate optimal number of host VMs
 - Ensure payment is minimum
 - getting maximal speedups
 - Normal distribution
- Image rendering

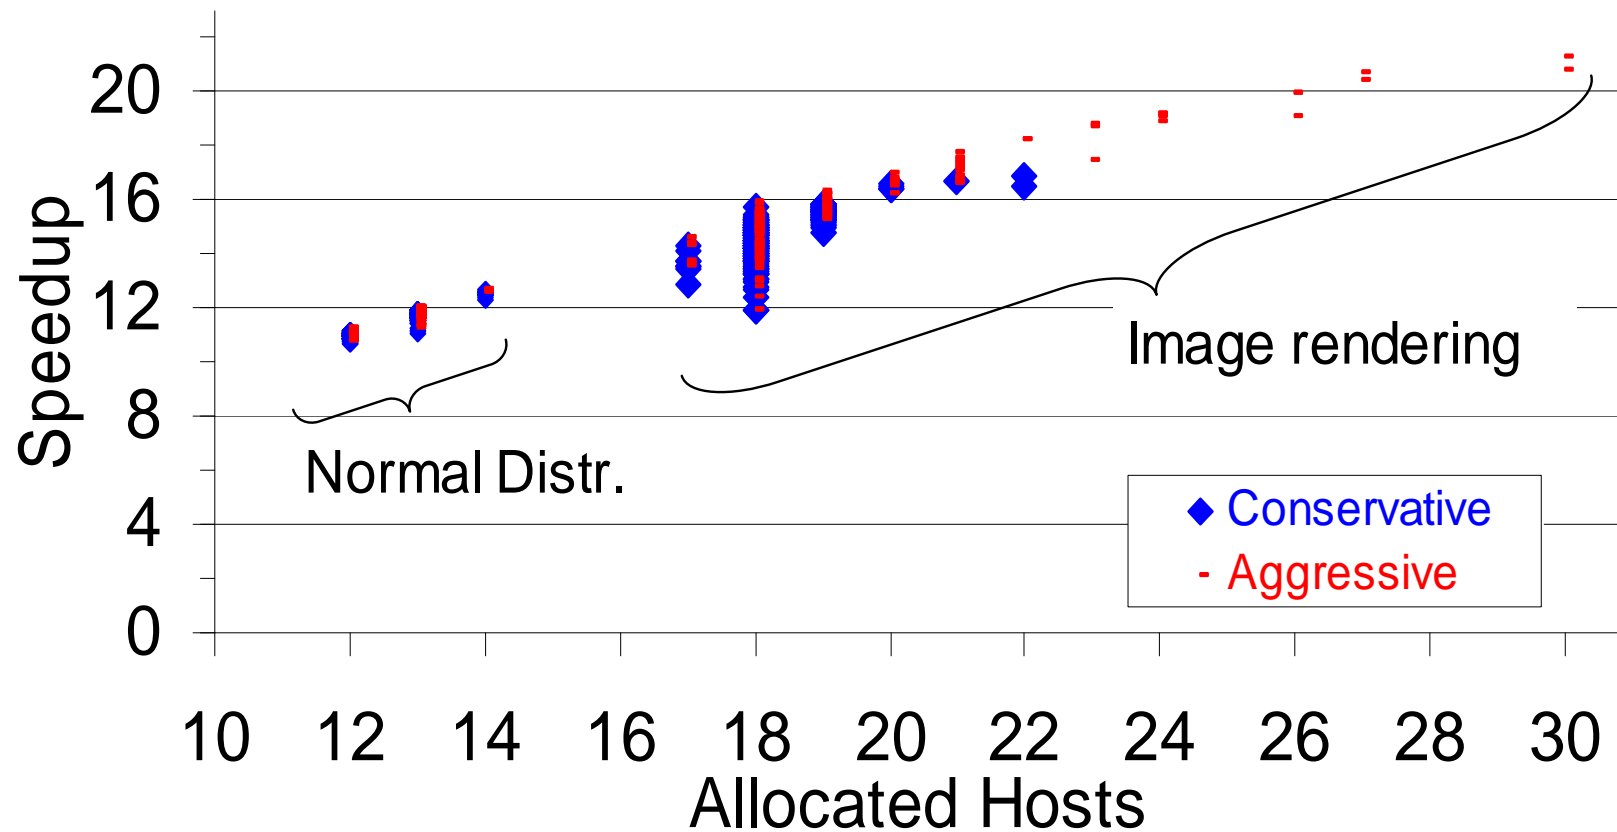


- Allocated hosts and speedups
 - Normal distribution
 - Image rendering



- Allocated hosts vs. speedups
 - Conservative
 - creationRatio 0.5
 - increaseRatio 0.5
 - Aggressive
 - creationRatio 0.75
 - increaseRatio 0.75

- Allocated hosts vs. speedups



Conclusions




technology
from seed

- Heuristic to use utility computing infrastructures efficiently
- Guarantees a desired cost
 - Based on equivalent-serial processing time
 - Achieves high speedups
 - Without user intervention
- Protects user from overshooting penalties
 - Prevents over-allocation beyond optimal point

- User hints
 - number of initially allocated VMs
 - In order to get better results
- Use previous execution times
 - Based on application, data, user, etc.
 - To predict current execution times
 - Optimize task distribution

Questions?
www.gsd.inesc-id.pt/~jnos

**technology
from seed**



**Heuristic for resources allocation on utility
computing infrastructures**