Combining Heuristics for Optimizing the Placement of IoT applications in the Fog

Ye Xia[1,2]

supervisors: Thierry Coupaye[1] Frédéric Desprez[2]
Xavier Etchevers[1] Loïc Letondeur[1]

[1]Orange Labs
first.last@orange.com

[2]INRIA
first.last@inria.fr

The 33rd ACM/SIGAPP Symposium On Applied Computing
12 / 04 / 2018

Outline

from Cloud to Fog

How to satisfy a time-sensitive IoT application ?

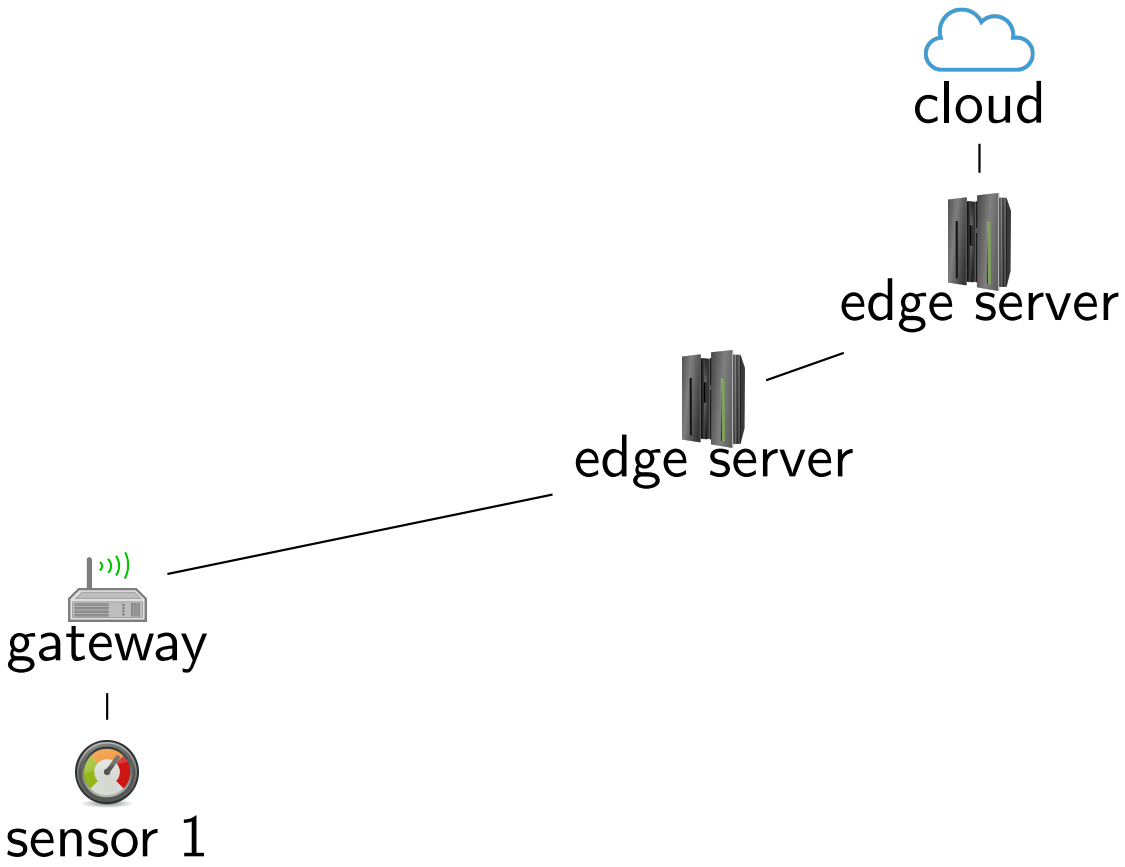

from Cloud to Fog

How to avoid congestions in the core network ?

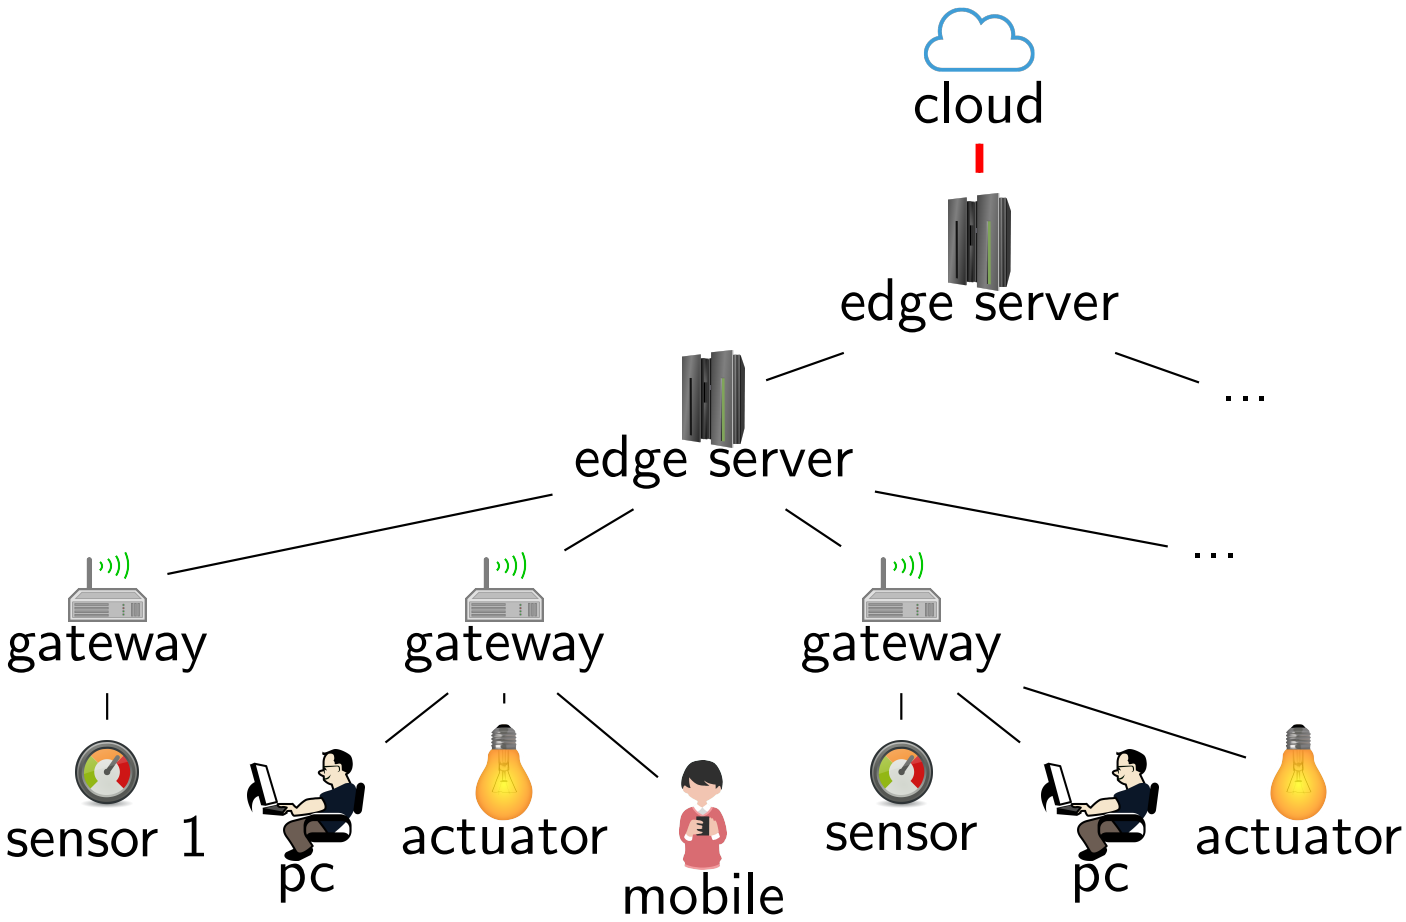

from Cloud to Fog

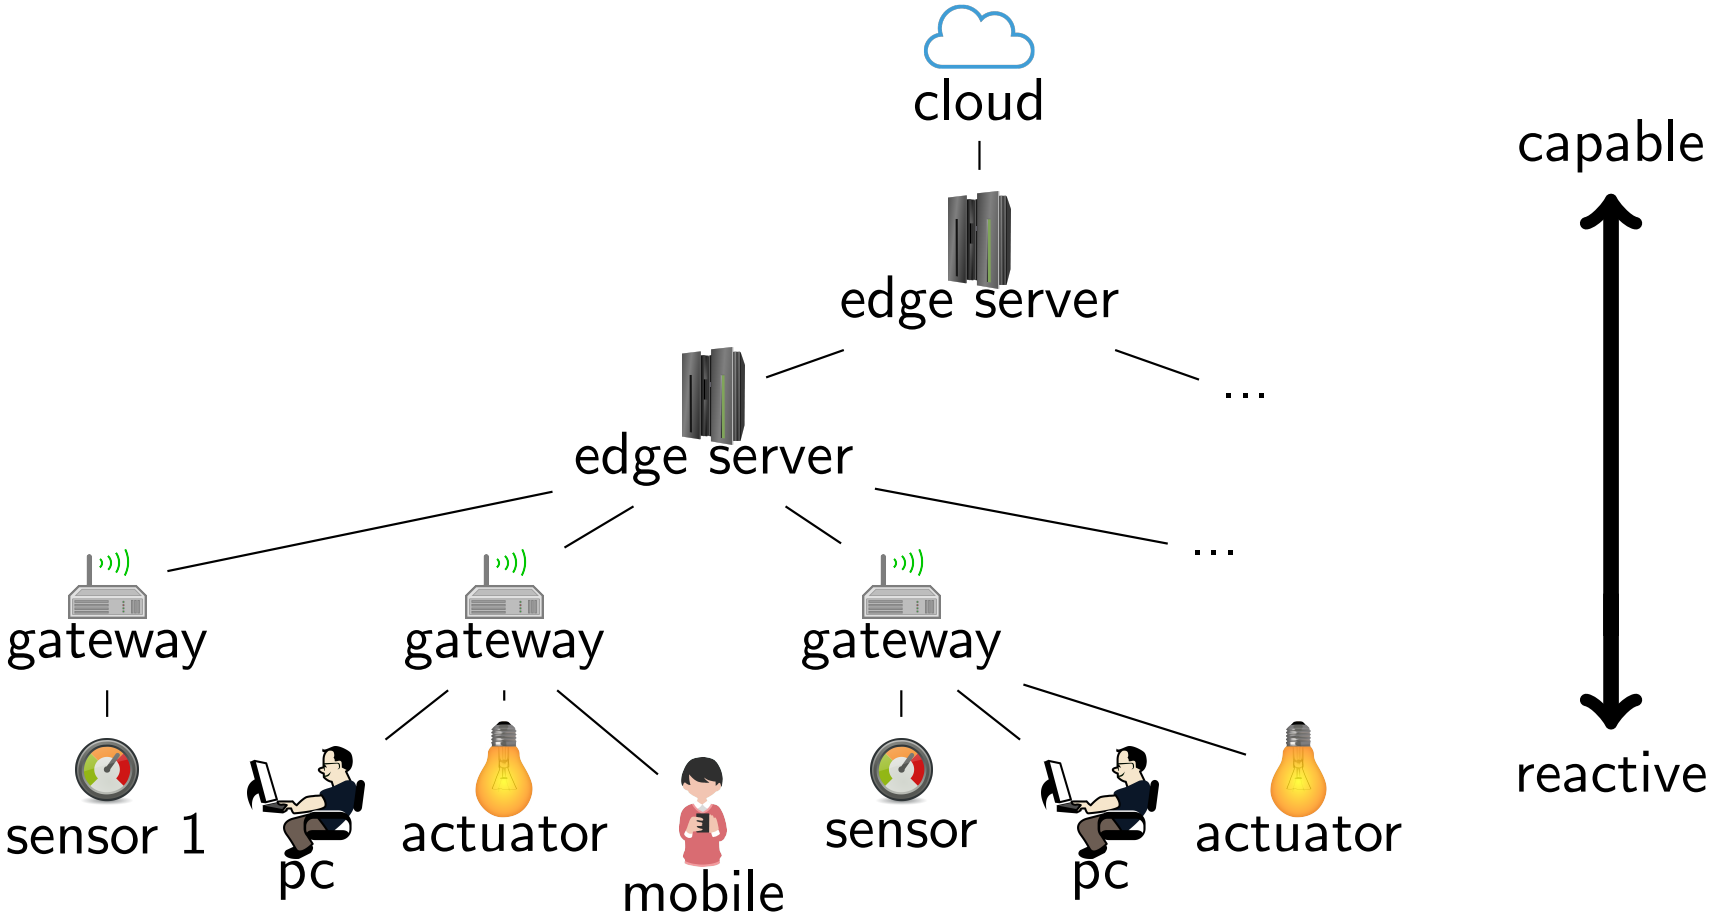

Placement Problem Description

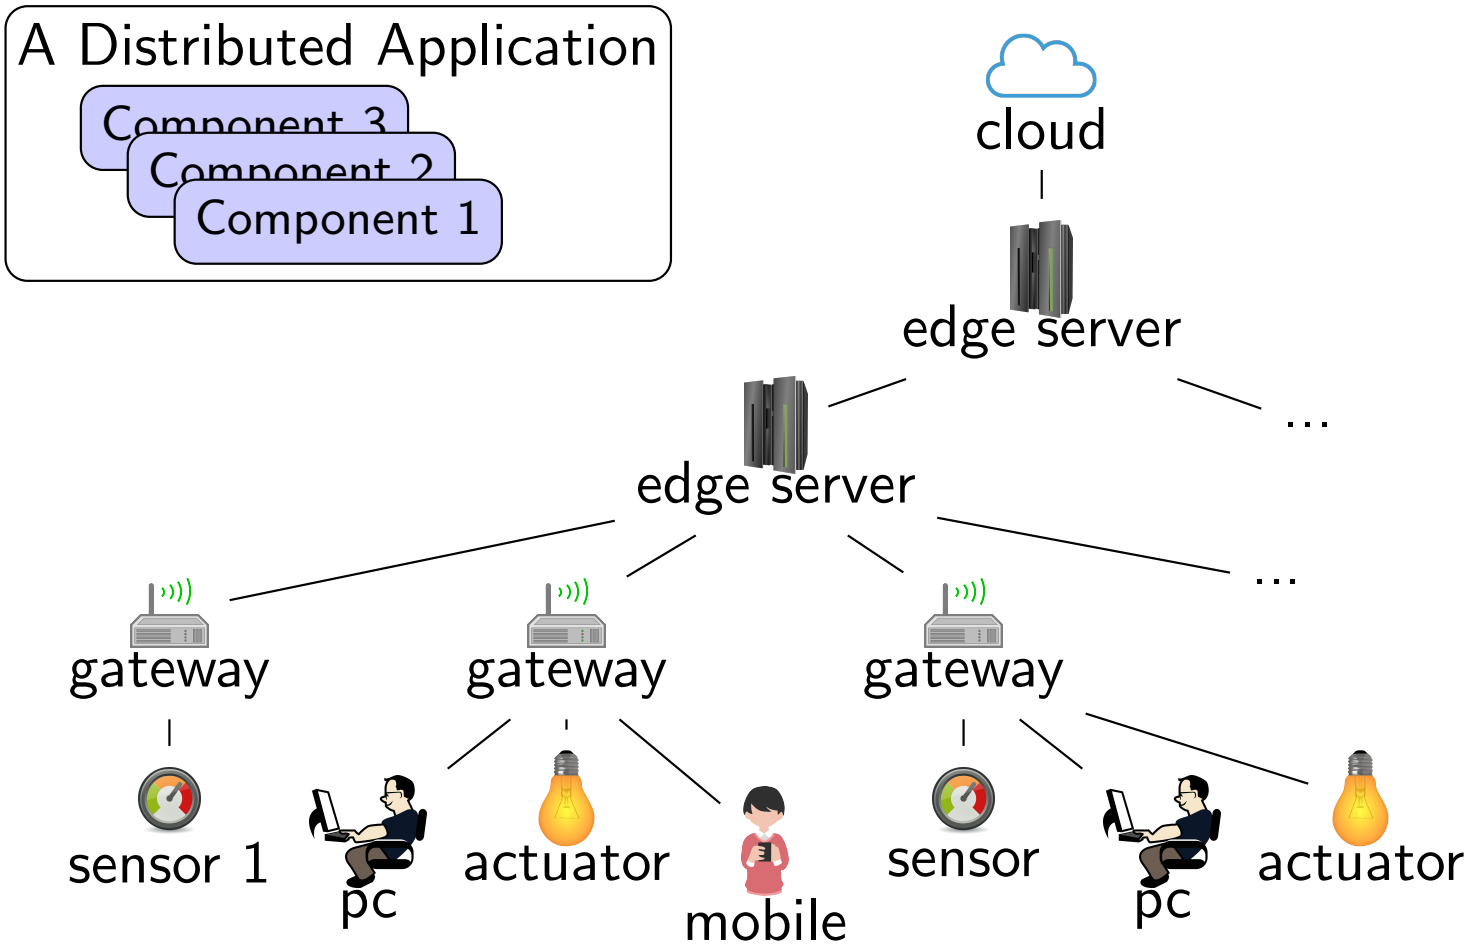

Placement Problem Description

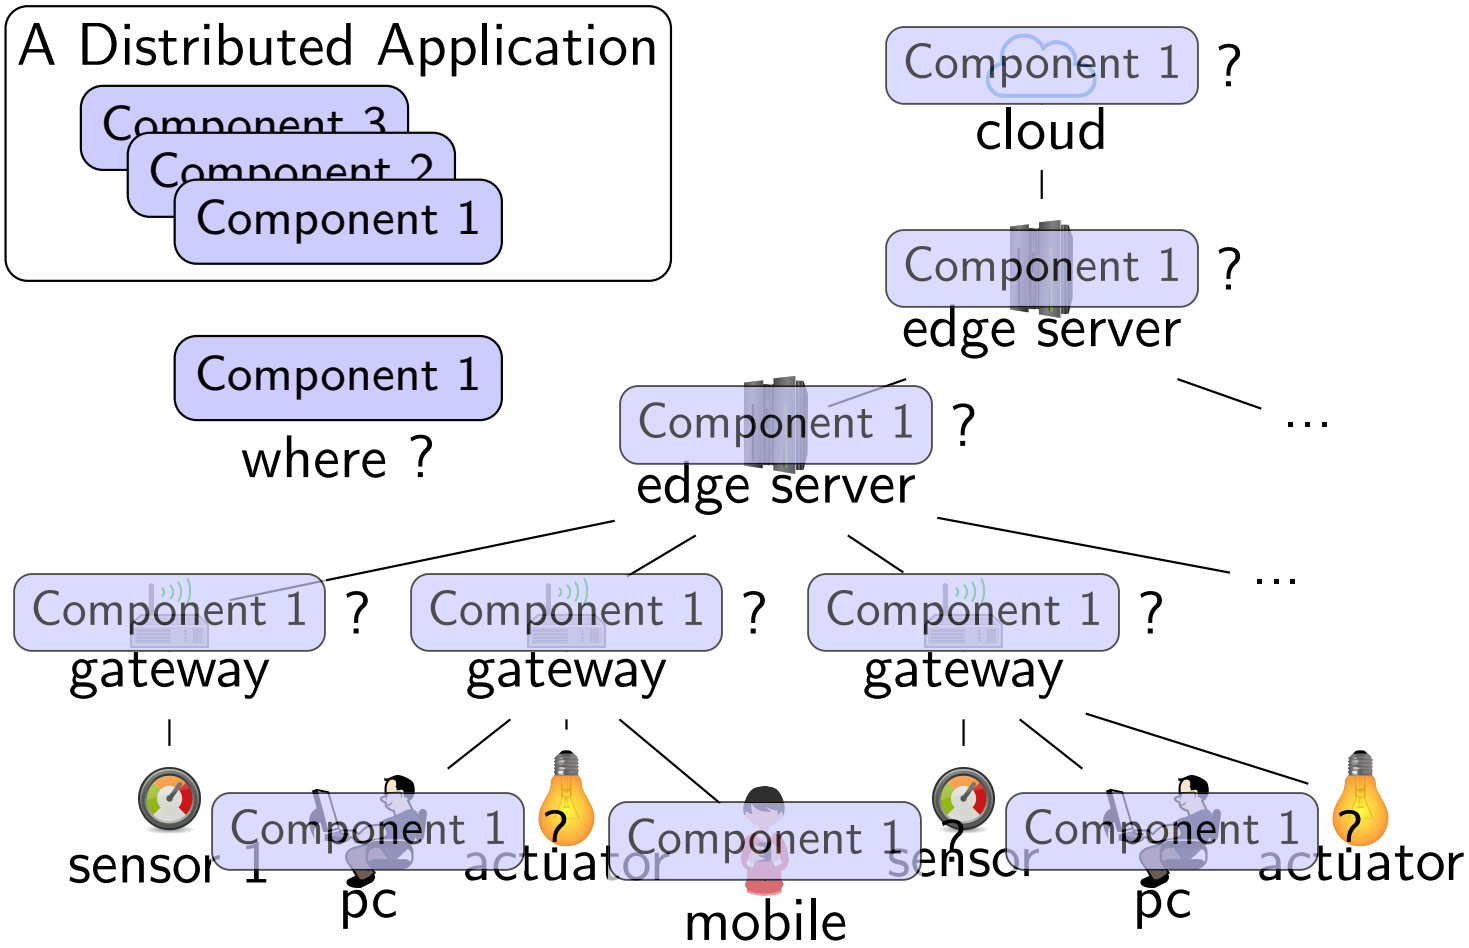

Challenges

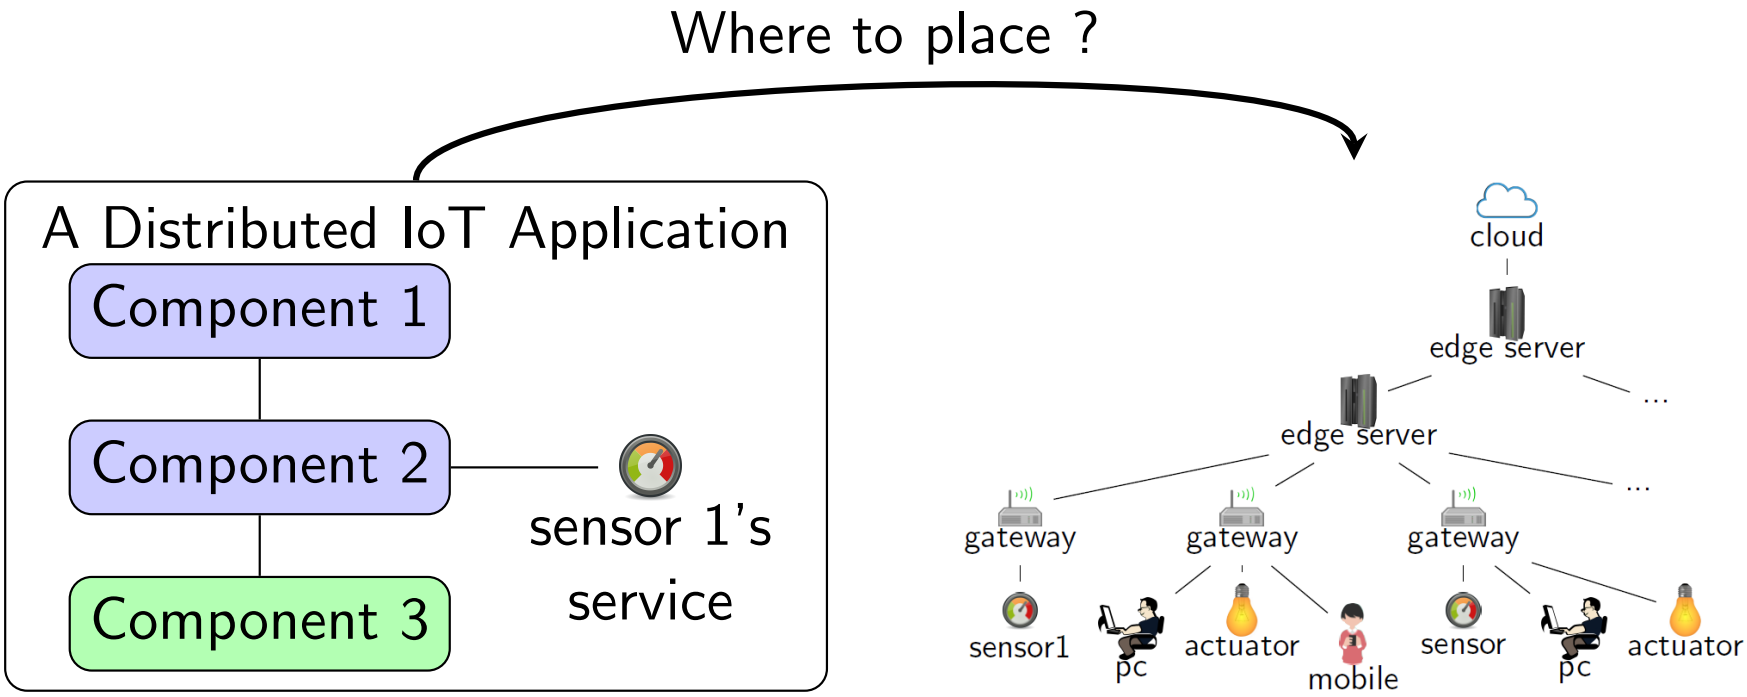

Challenges

- Locality
- Constraint Diversity & Heterogeneity
  - processing, storage, and network resources
  - Component 3 : a certain user's devices, certain OS
- Scalability

Outline

Model — Infrastructure

Infrastructure

- fog node : CPU, RAM, DISK
- appliance
- link : network latency (Lat), bandwidth (BW)

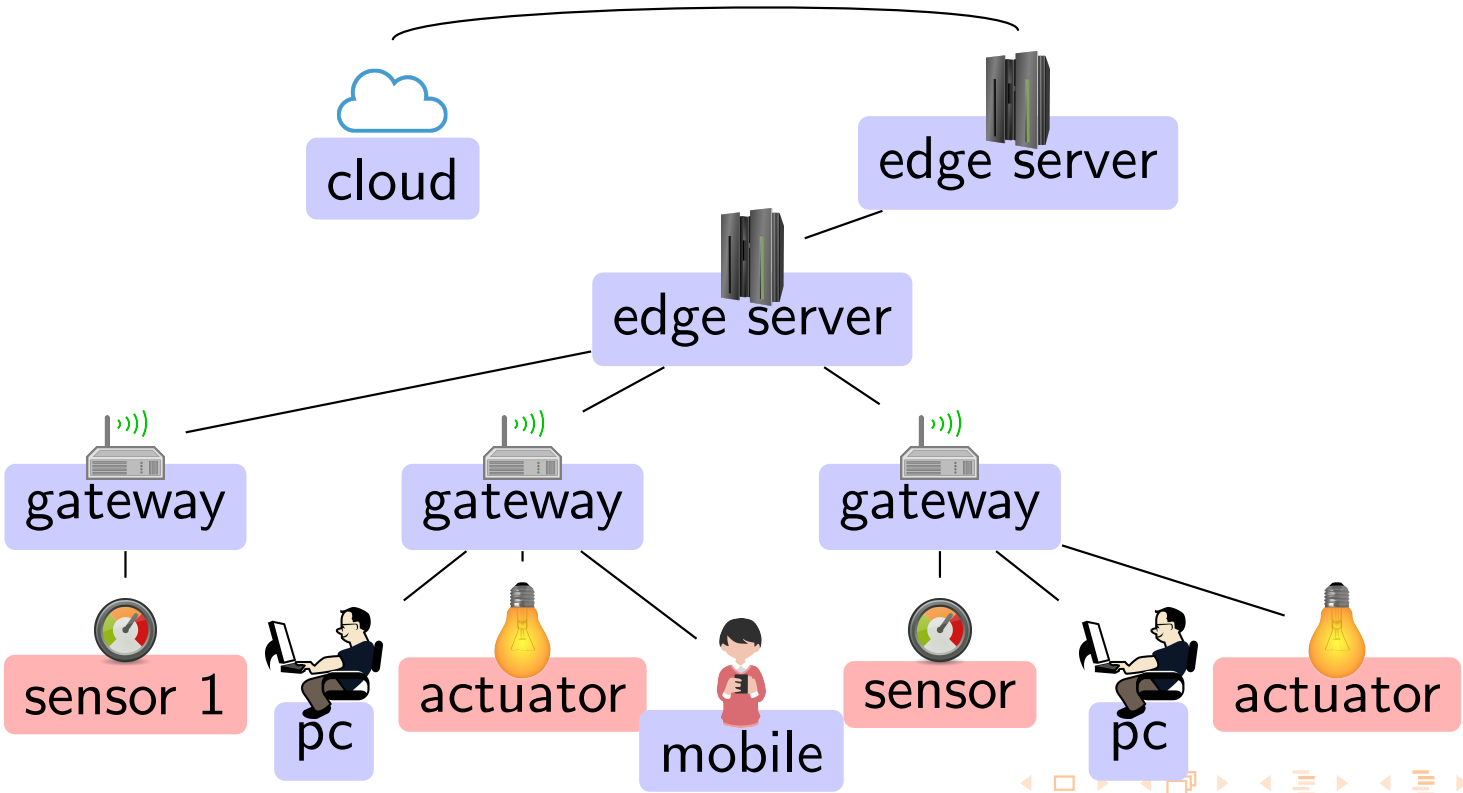

Model — Applications

Apps : the set of applications to place

Application

- component : reqCPU, reqRAM, reqDISK, Dedicated Zone
- connected object
- binding : reqBW, maximal latency

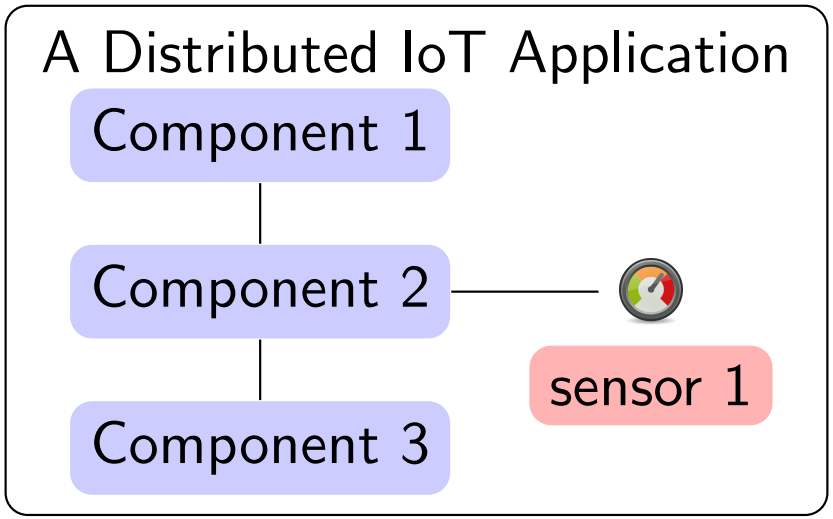

Constraints

Placement

- each component -≻ one fog node

$$a\ placement\ = \begin{pmatrix} comp_1, node_i \\ comp_2, node_j \\ \dots \\ comp_n, node_k \end{pmatrix}$$

A Solution must satisfy :

- Dedicated Zone
- Consumption of CPU / RAM / DISK
- Consumption of BW
- Binding's maximal latency

Solution Selection

How to select among multiple solutions ?
Goal : minimizing applications' average response time

Objective Function

min : Weighted Average Latency (WAL)

$$WAL = \sum_{bind} \frac{bind.reqBW}{total\_BW} \times bind.Lat$$

Objective Function Evaluation

SIMGRID [4]

Correlation : 0.8029574

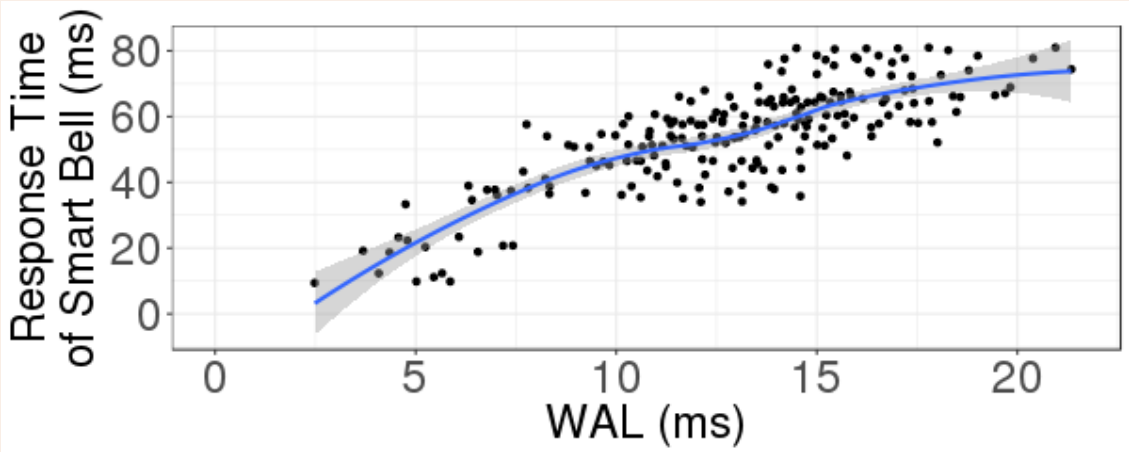

Outline

Example – Backtrack Search

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

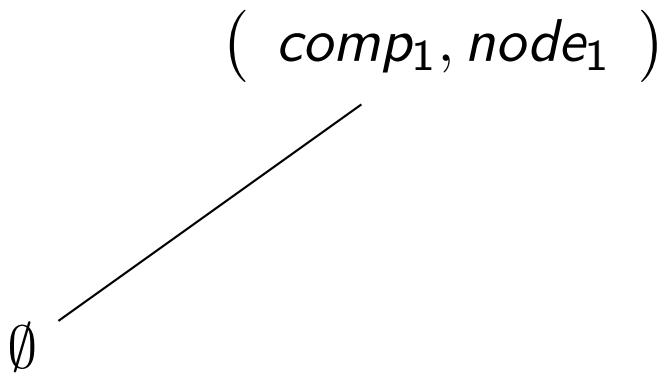

Example – Backtrack Search

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

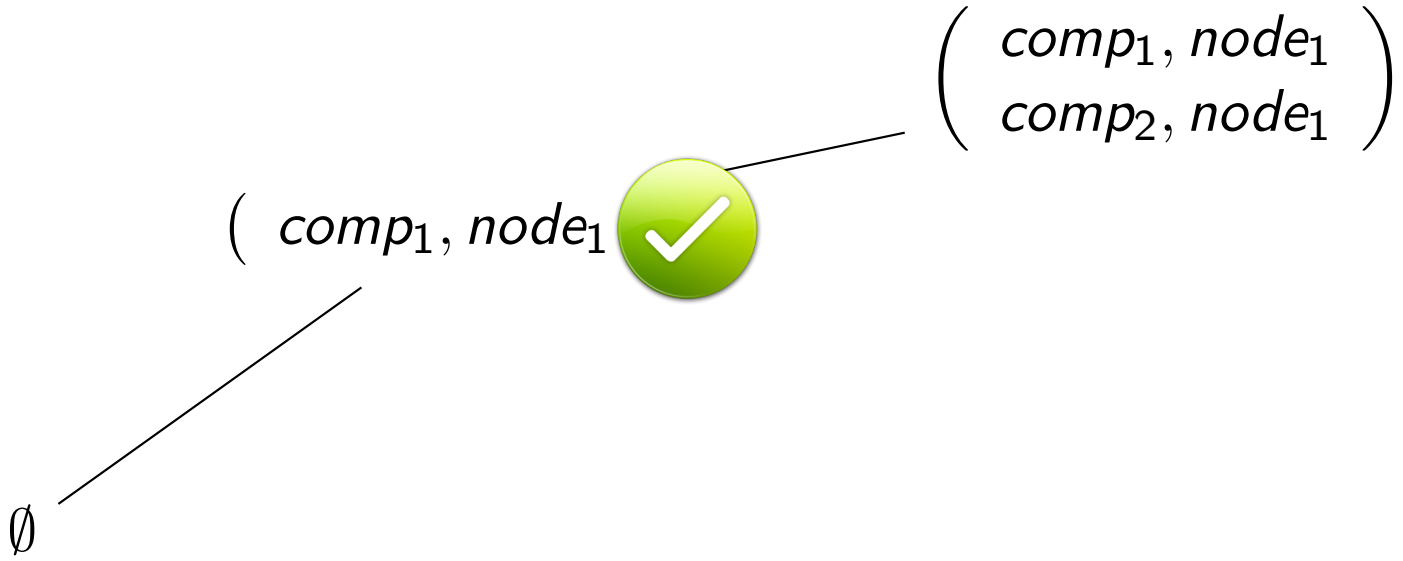

Example – Backtrack Search

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

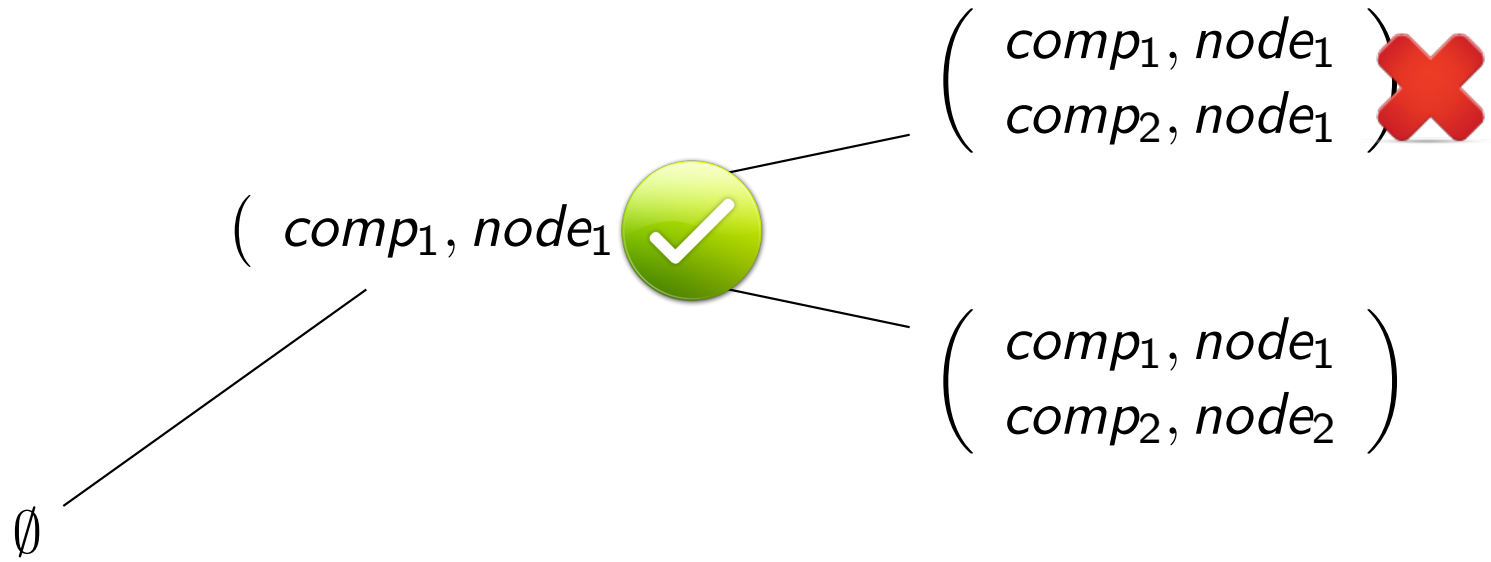

Example – Backtrack Search

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

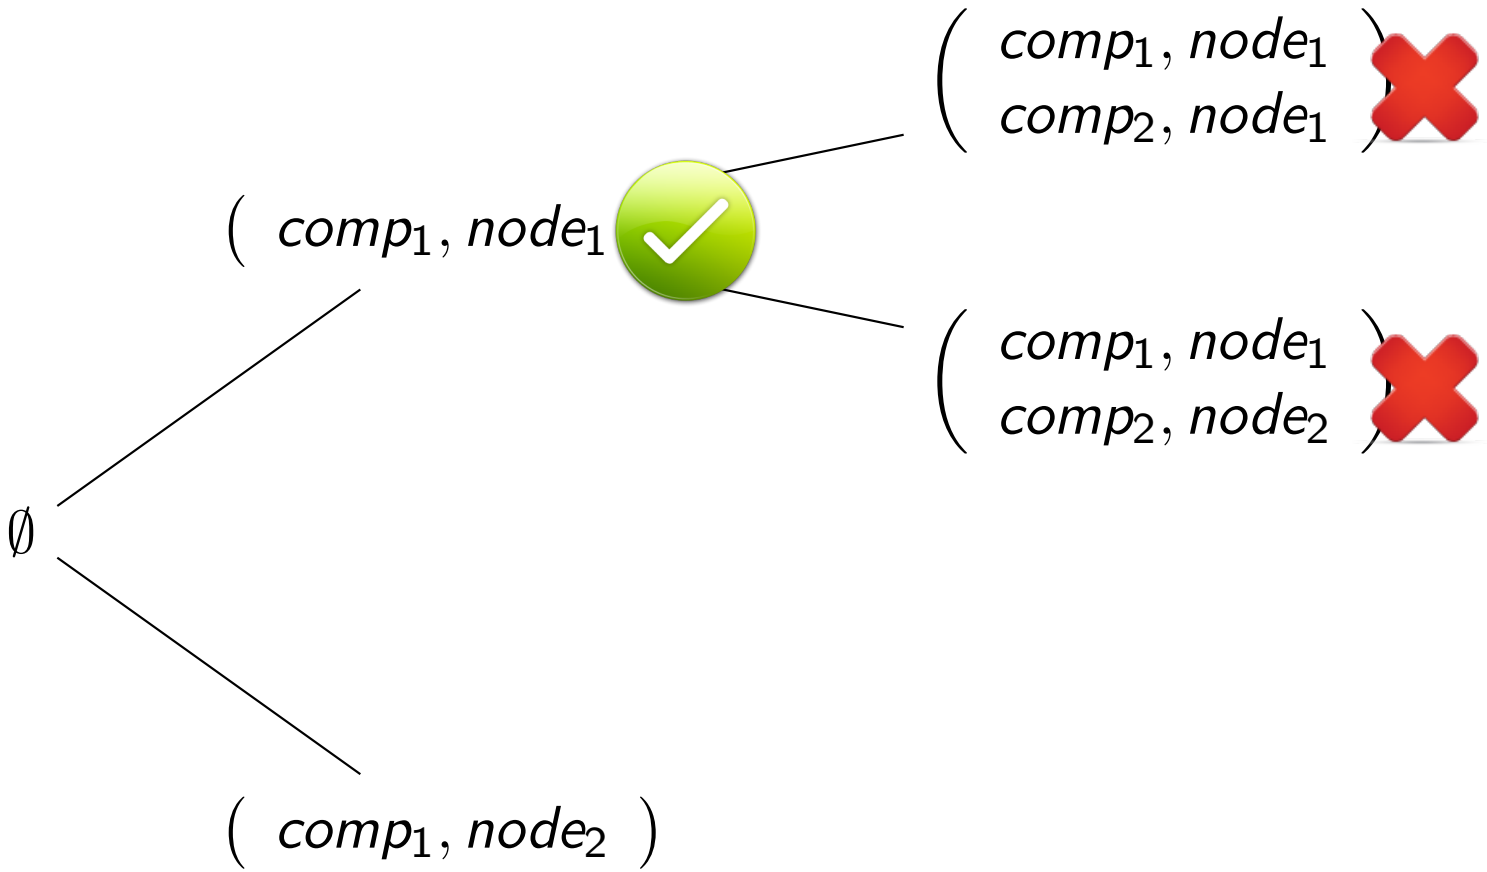

Example – Backtrack Search

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

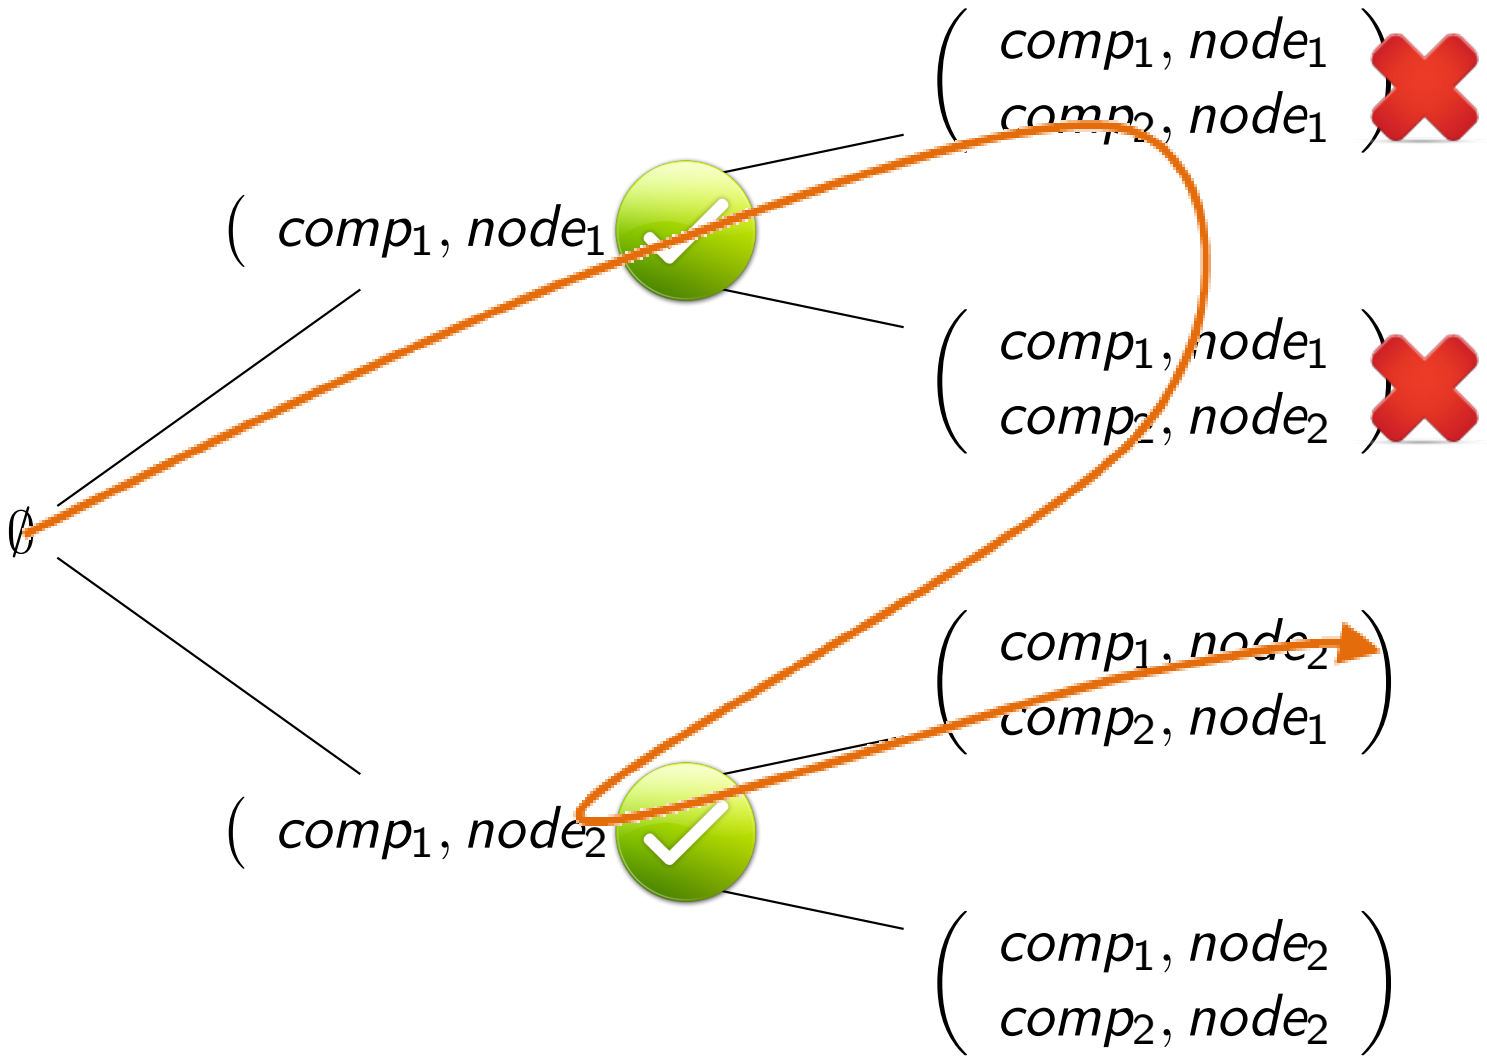

Naive Placement Algorithms

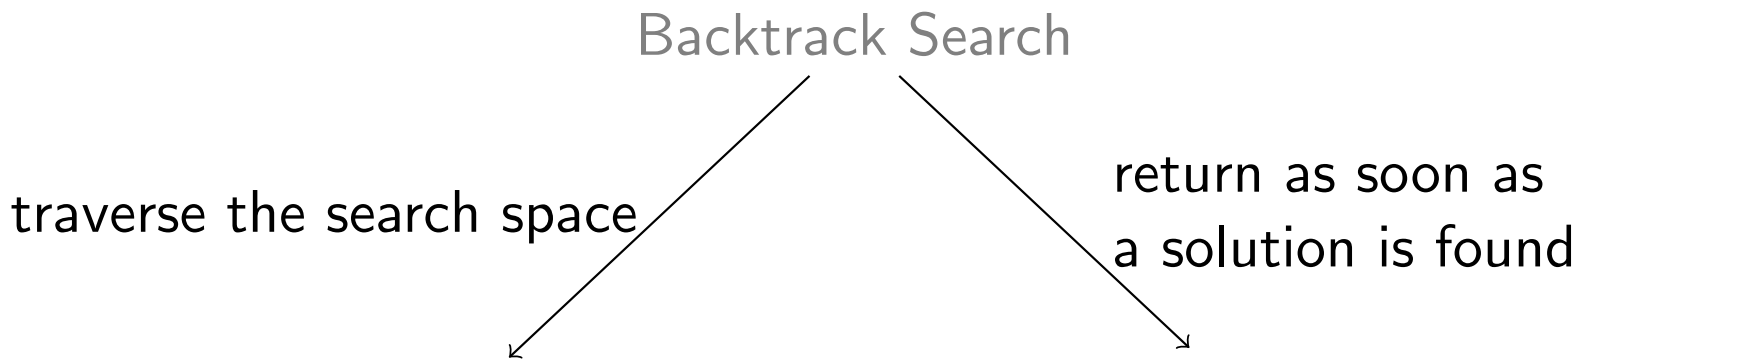

Exhaustive Algorithm
- guarantee to return the optimal solution
- not scalable

FirstFit Algorithm
- more scalable
- no guarantee on returned solution's quality

Example – Impact of Fog Node Order

$\{comp_1, comp_2\}$ -$\succ$ $\{node_1, node_2\}$

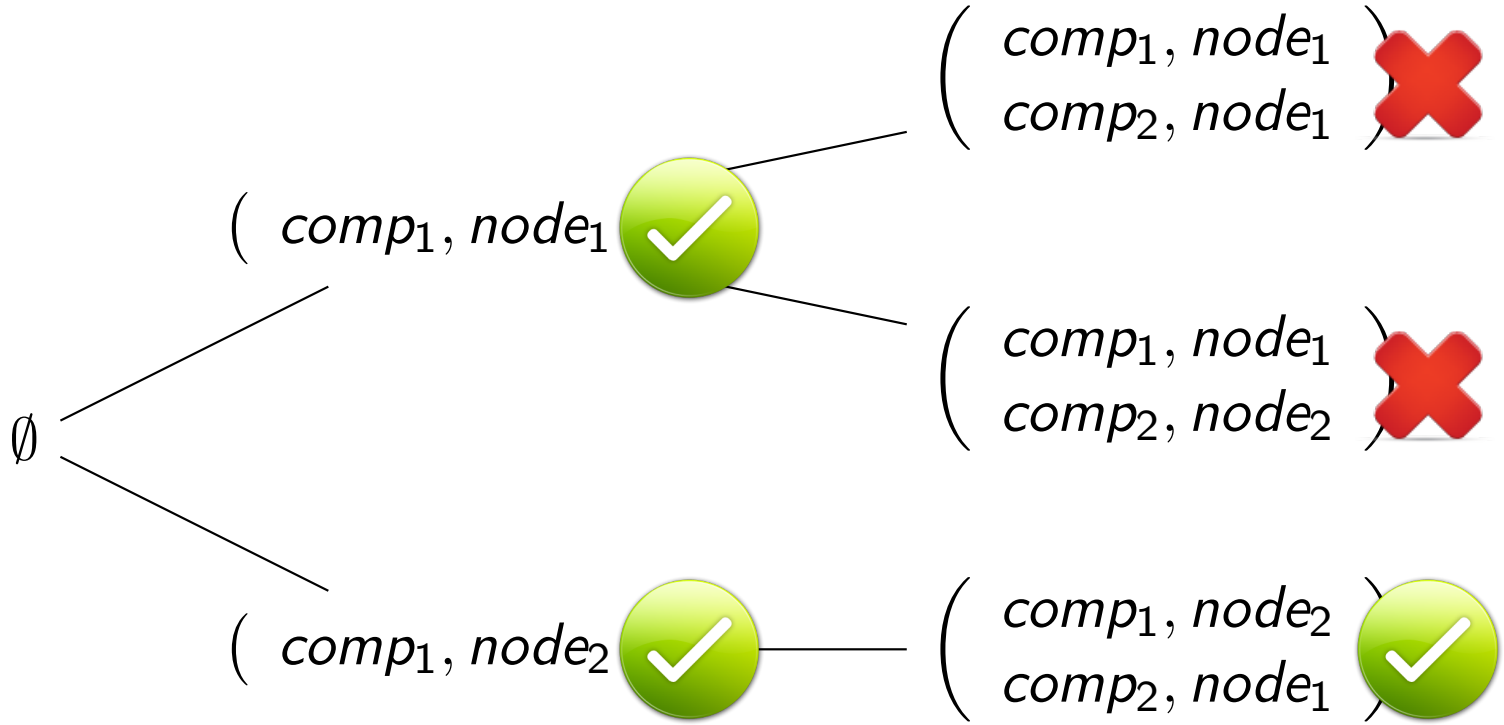

Heuristics

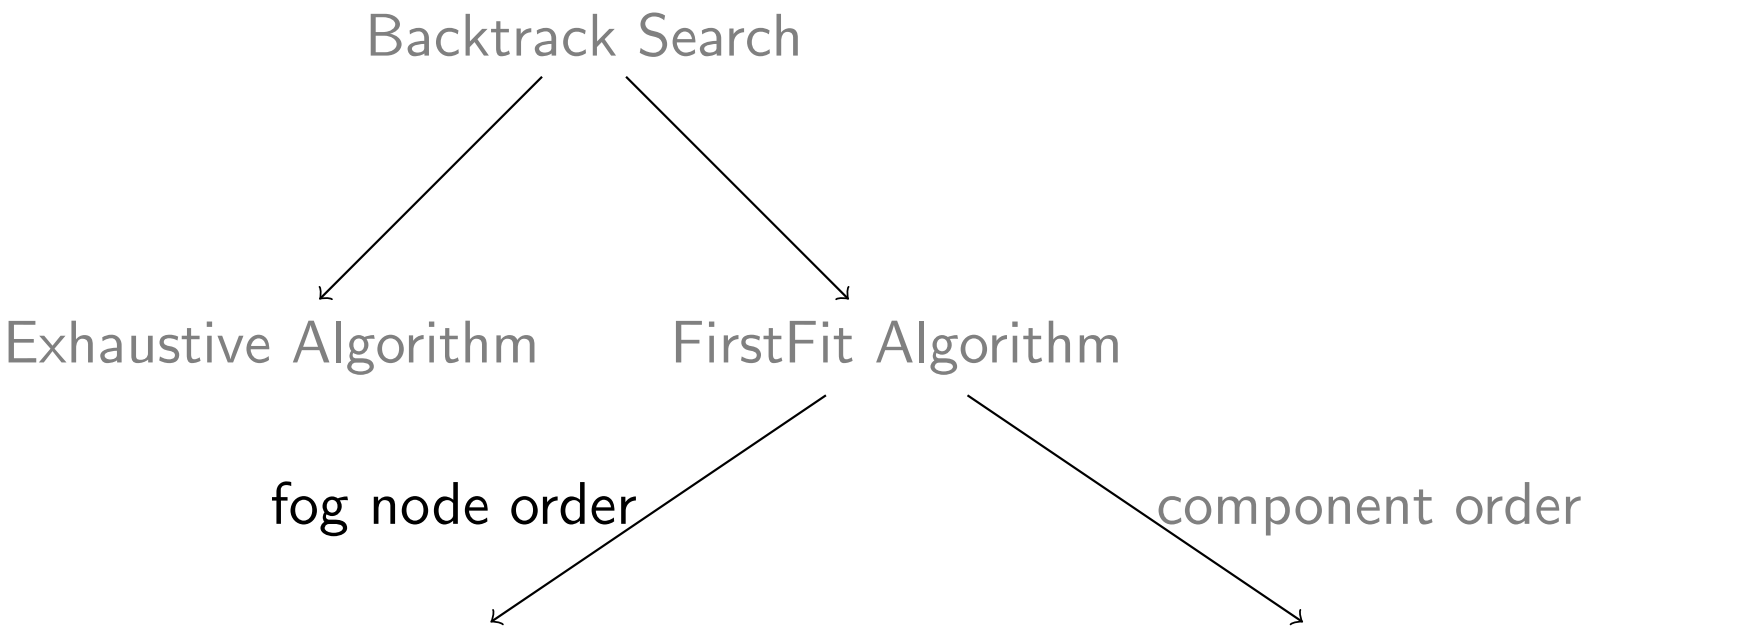
Backtrack Search
Exhaustive Algorithm
FirstFit Algorithm
fog node order
component order
Anchor-based Fog Nodes Ordering
- for lowering WAL & accelerating
Dynamic Components Ordering
- only for accelerating the search

AFNO Process

1. Anchor Calculation

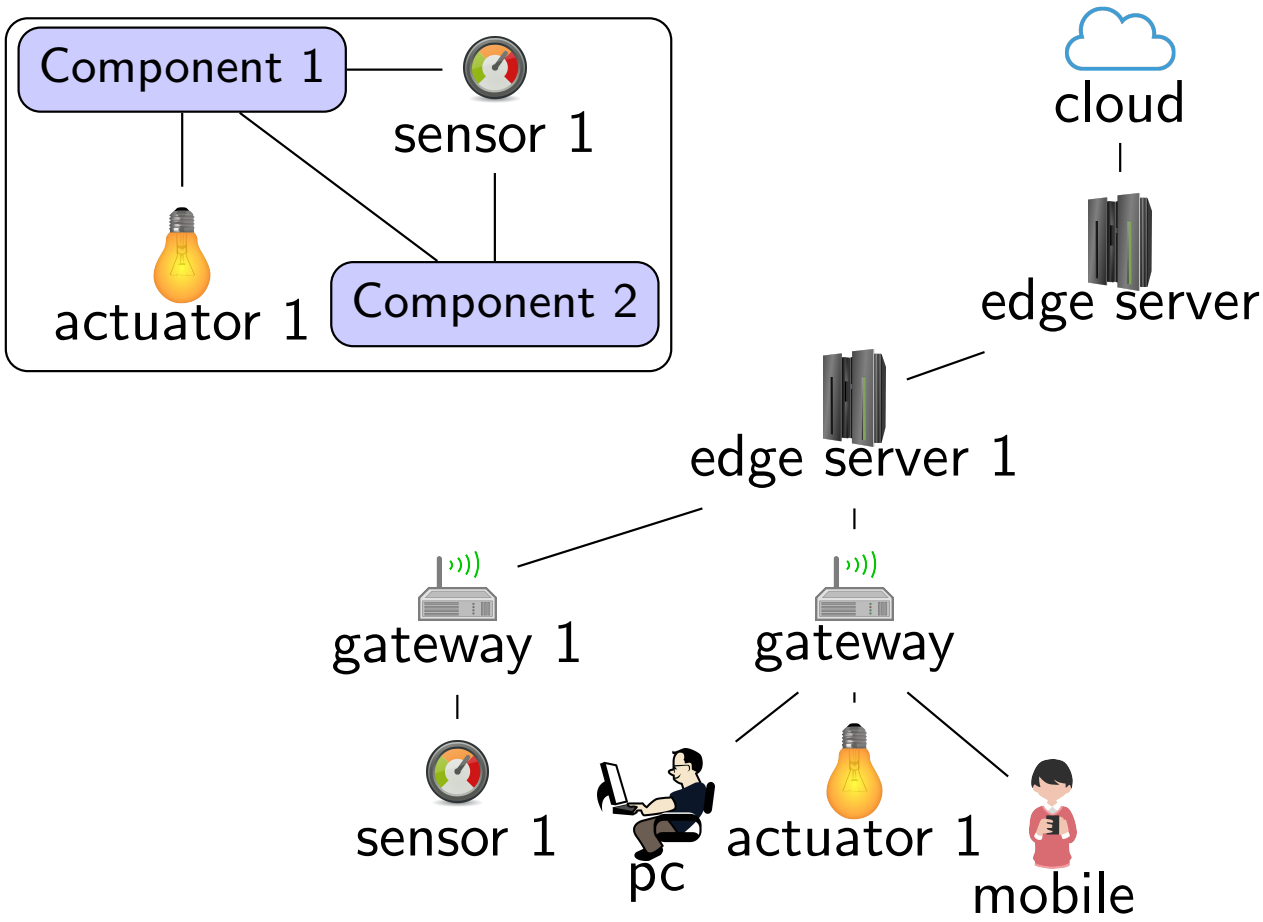

2. Fog Nodes Ordering
3. Search

AFNO Process

1. Anchor Calculation

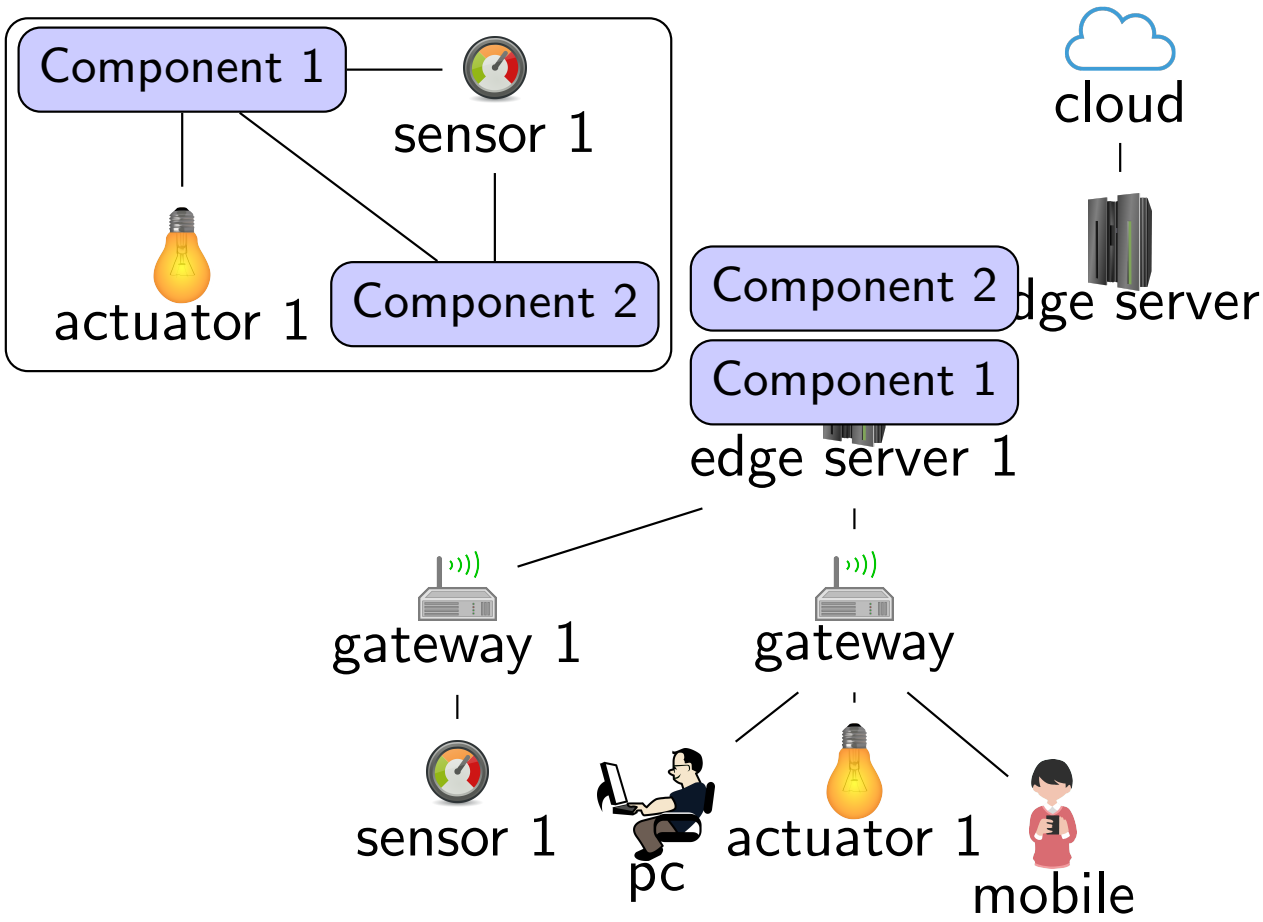

2. Fog Nodes Ordering
3. Search

AFNO Process

1. Anchor Calculation

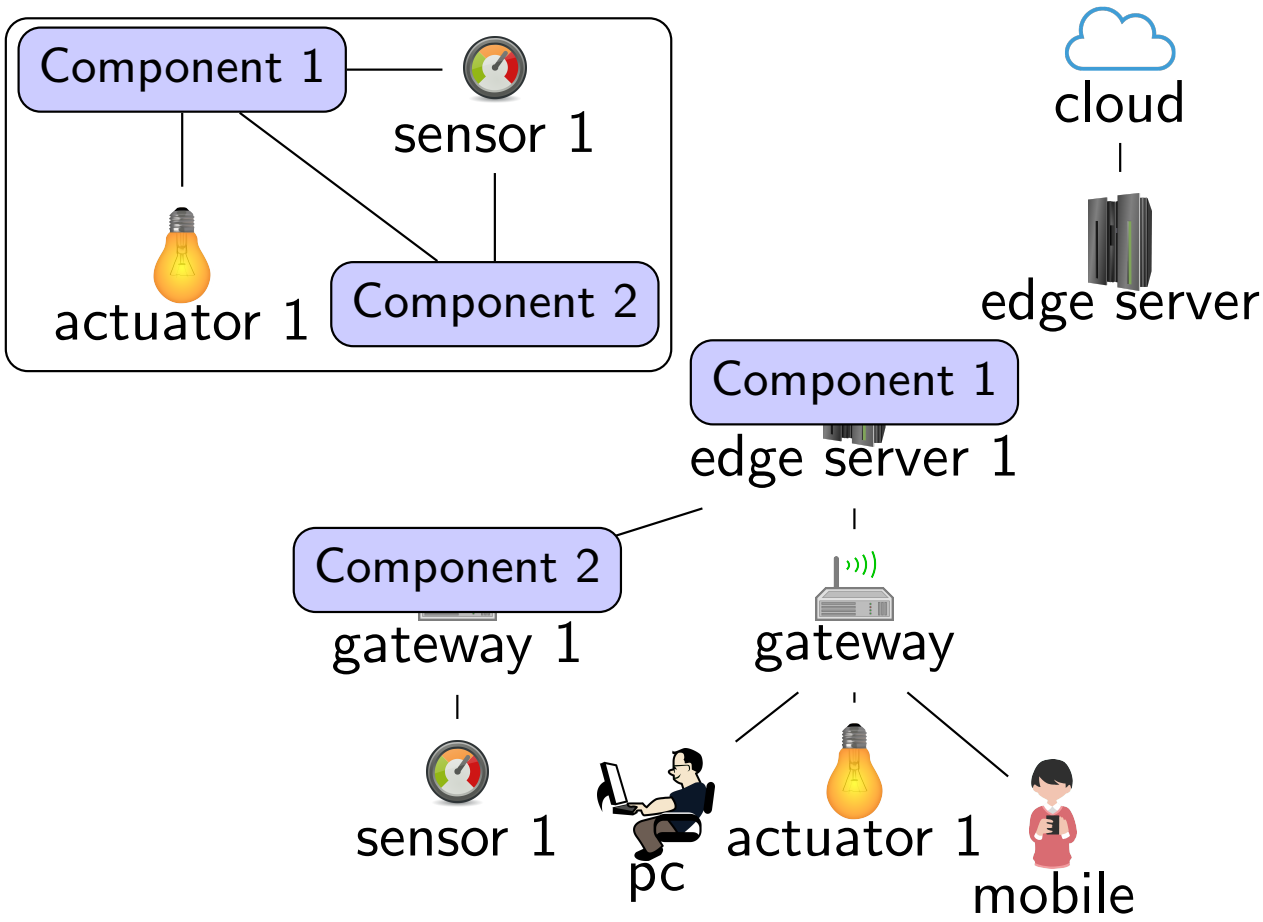

2. Fog Nodes Ordering
3. Search

Example – Impact of Component Order

$\{comp_1, comp_2\}$ -≻ $\{node_1, node_2\}$

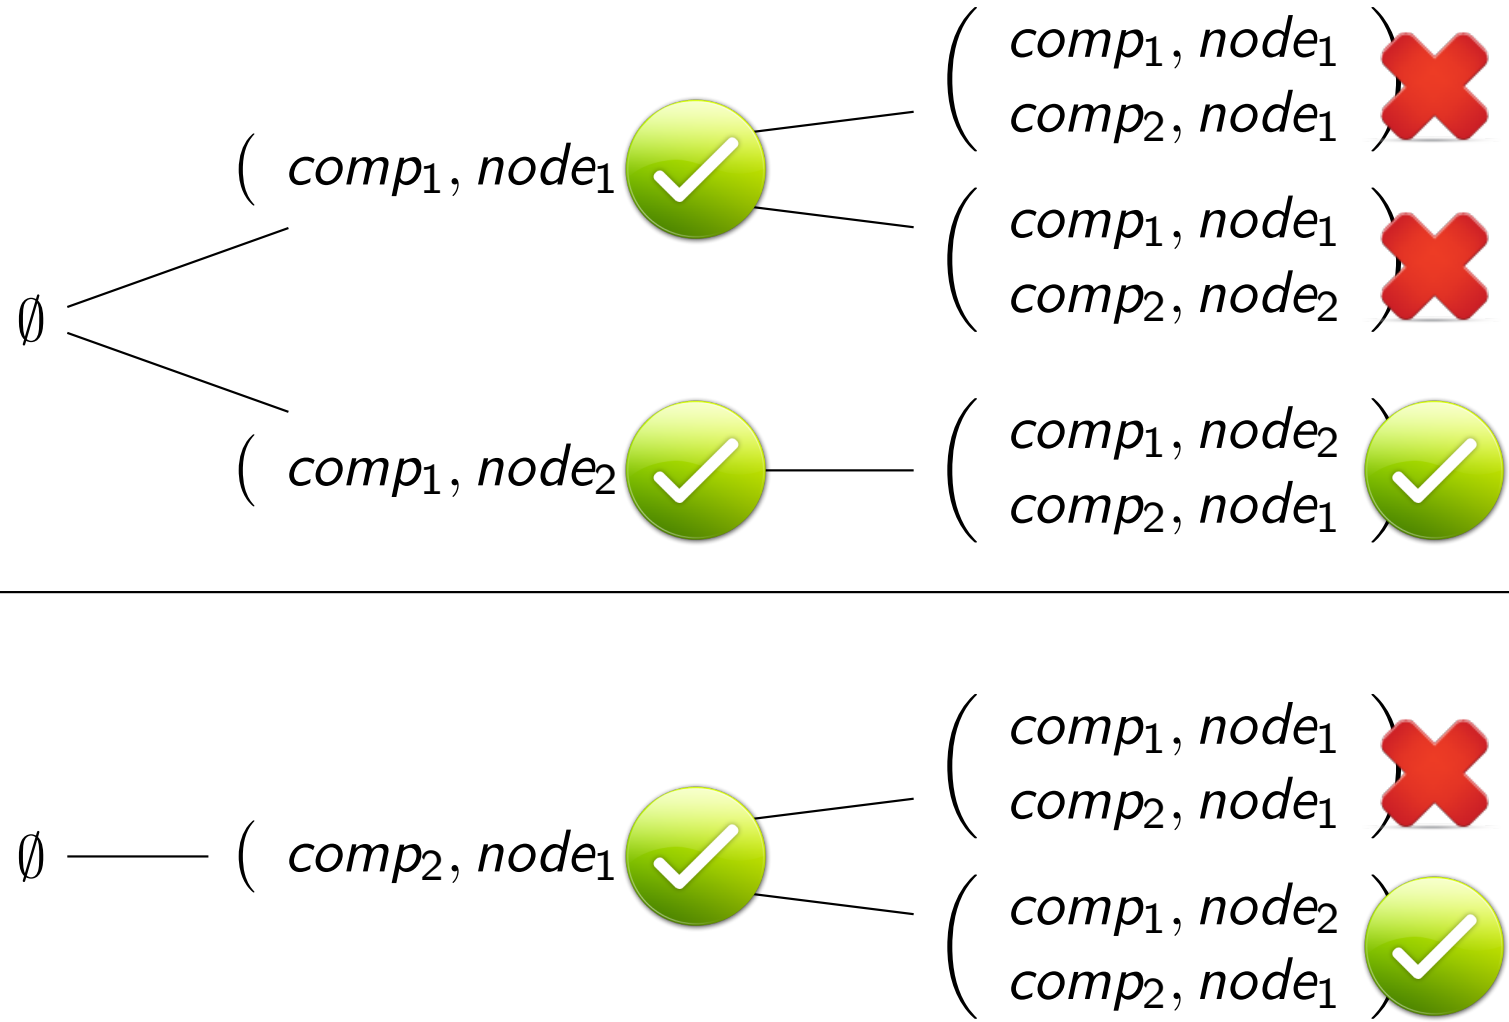

Heuristics

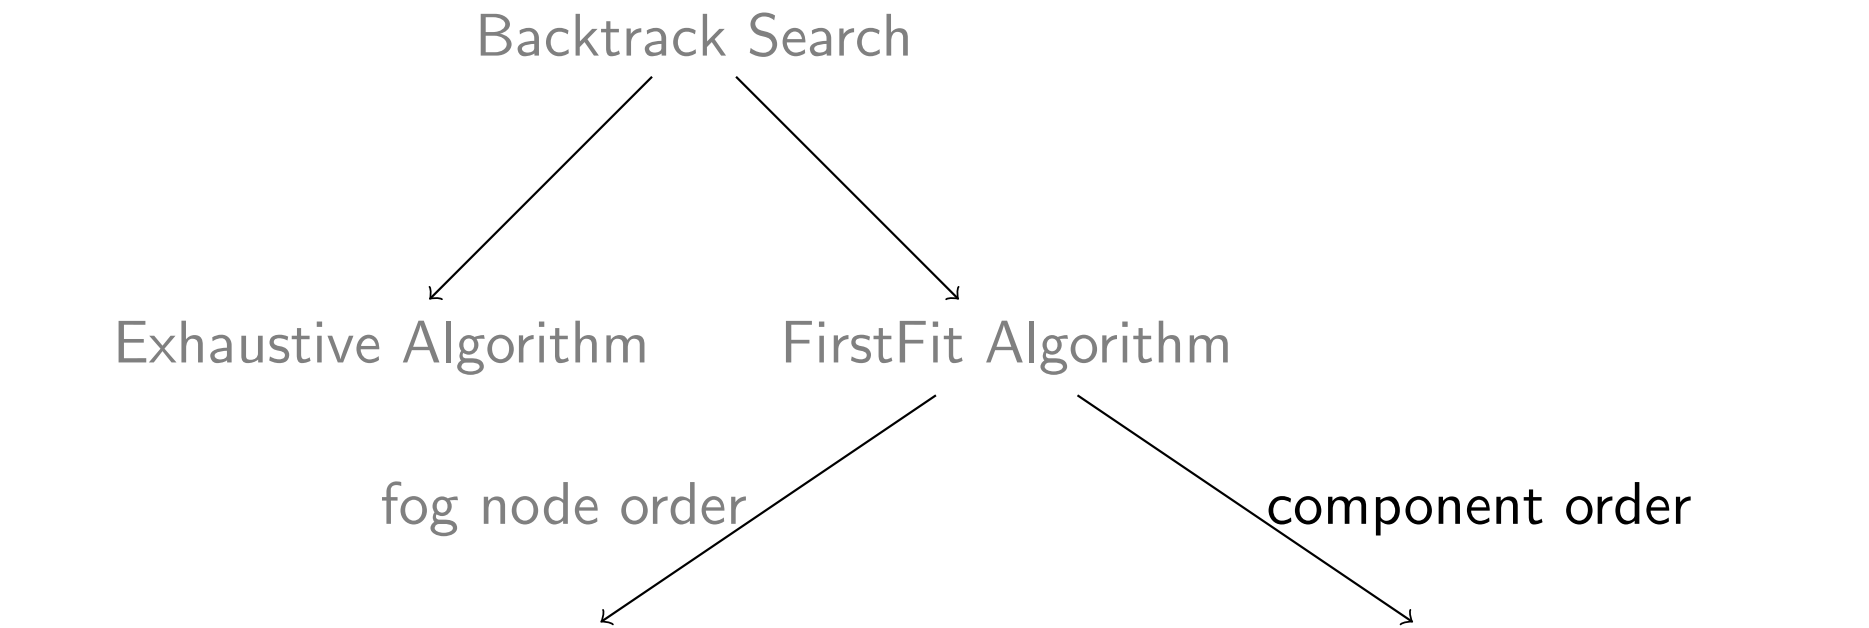
Backtrack Search
Exhaustive Algorithm
FirstFit Algorithm
fog node order
component order
Anchor-based Fog Nodes Ordering
- for lowering WAL & accelerating
Dynamic Components Ordering
- only for accelerating the search

Heuristics' Combination

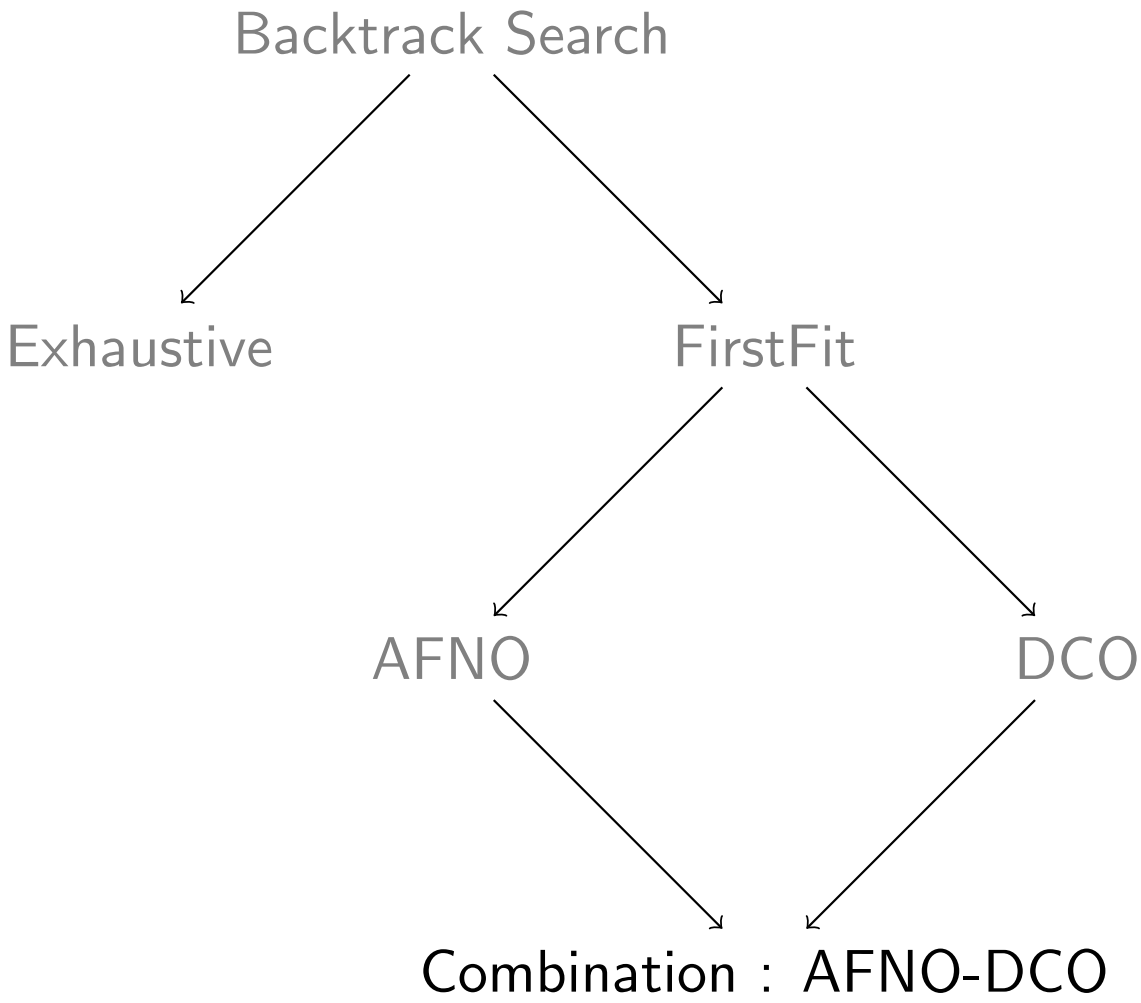

Evaluation — Scalability

timeout : 30 mins

Place 1 application in a growing infrastructure

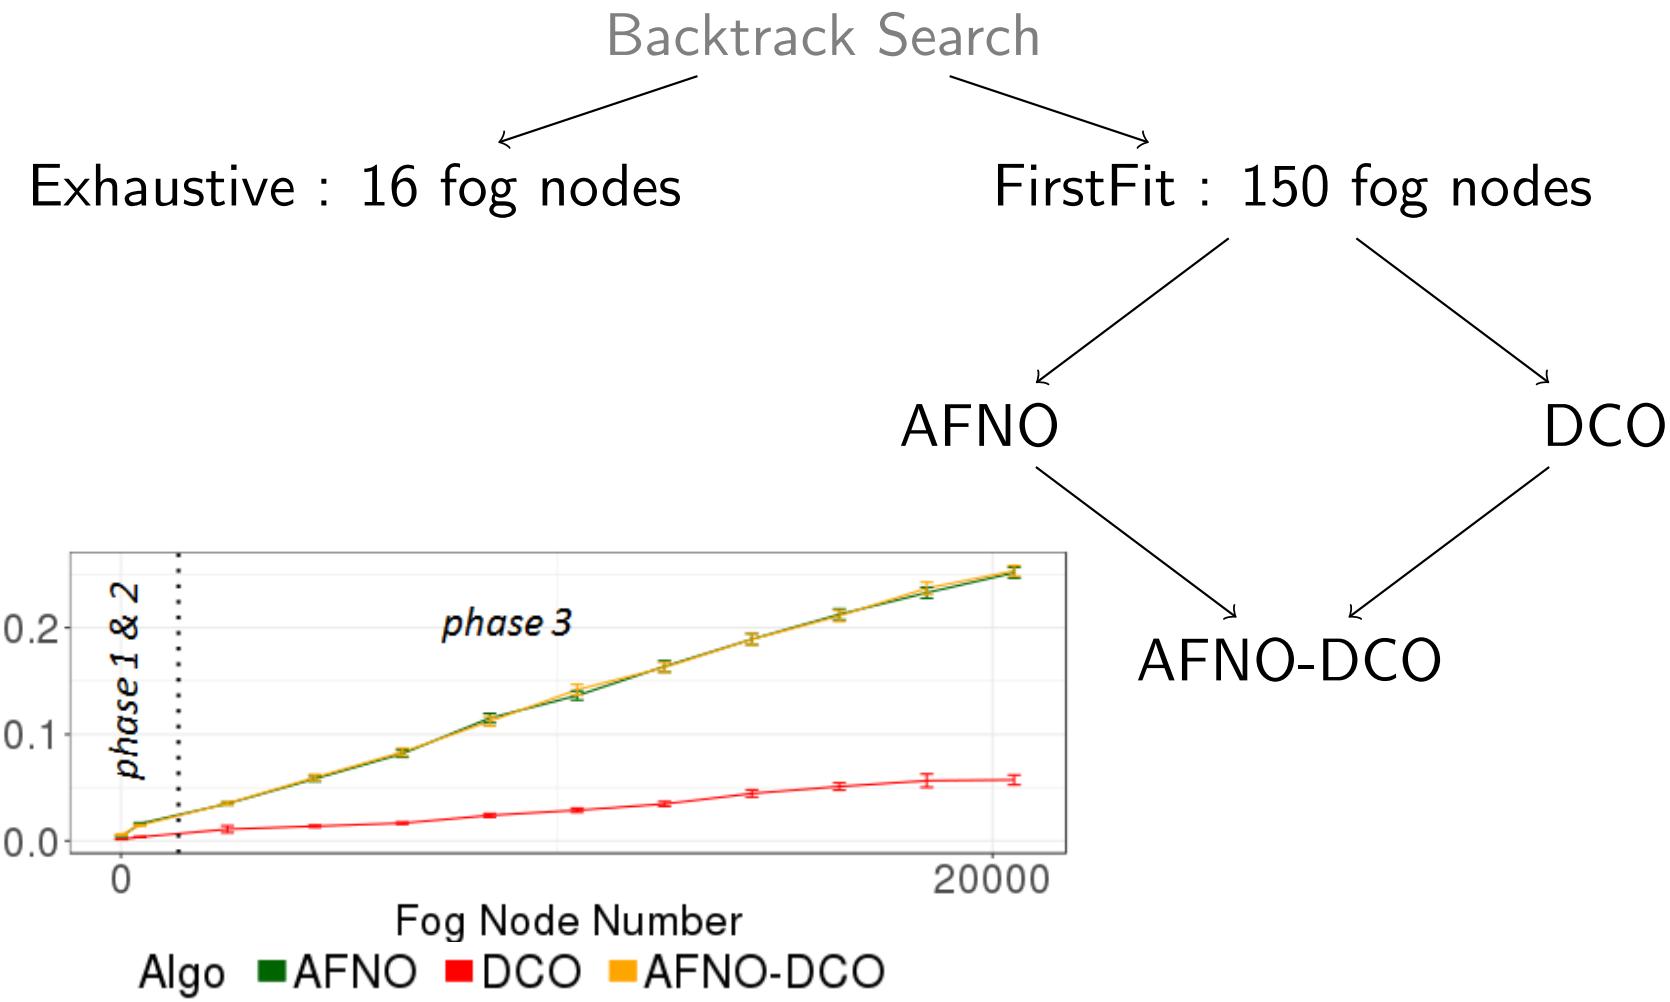

Evaluation — Scalability

timeout : 30 mins

Place more and more applications in an infrastructure with 20495 fog nodes

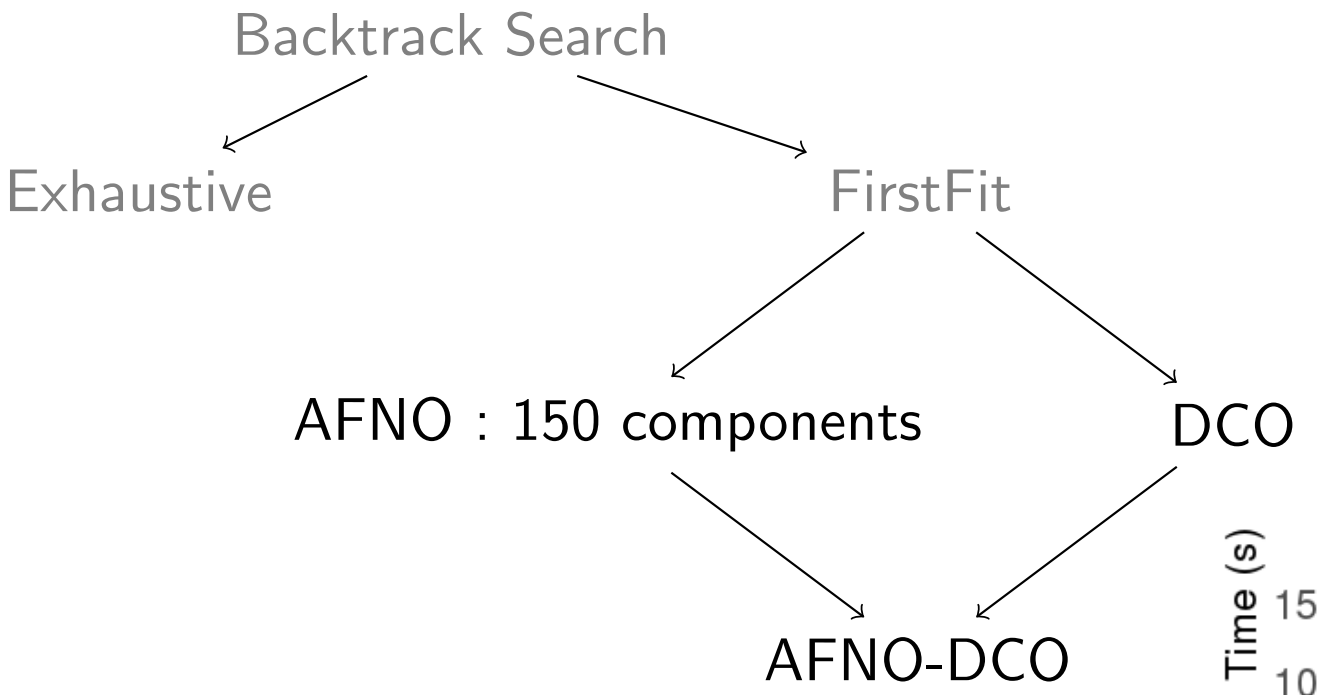

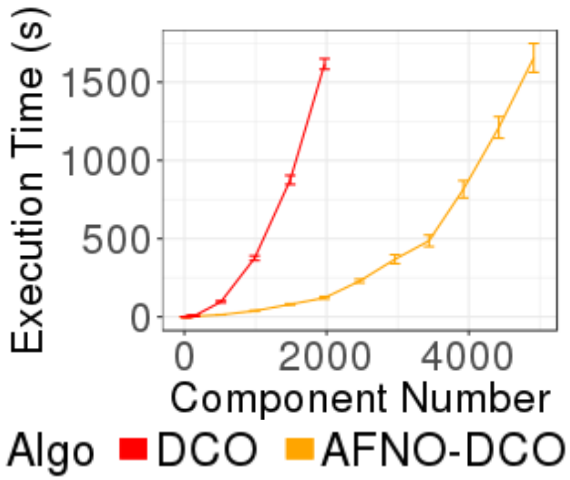

Evaluation — Result Quality

Average WAL

Exhaustive	$=$	$99.6\% \times$ AFNO-DCO
AFNO	$\simeq$	$100\% \times$ AFNO-DCO
DCO	$=$	$140\% \times$ AFNO-DCO
FirstFit	$=$	$142\% \times$ AFNO-DCO

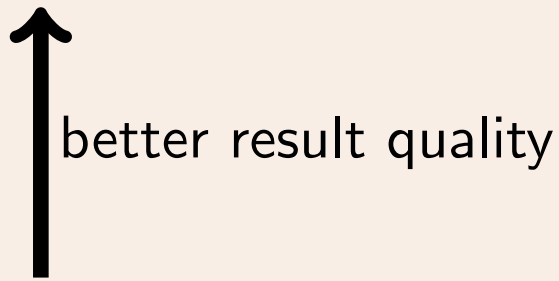

AFNO-DCO : the best compromise in terms of scalability and quality

Outline

Conclusion

Contribution

- Model, Objective Function
- Placement Algorithms (Exhaustive, FirstFit)
- Heuristics (AFNO, DCO)

Future Work

- Infrastructure Dynamicity (e.g., churn & mobility)
- Orange Labs internal testbed[3]

Conclusion

Contribution

- Model, Objective Function
- Placement Algorithms (Exhaustive, FirstFit)
- Heuristics (AFNO, DCO)

Future Work

- Infrastructure Dynamicity (e.g., churn & mobility)
- Orange Labs internal testbed[3]

Thanks for your attention!

References

[1] Ye Xia, Xavier Etchevers, Loïc Letondeur, Thierry Coupaye and Frédéric Desprez
Combining Hardware Nodes and Software Components Ordering-based Heuristics for Optimizing the Placement of Distributed IoT Applications in the Fog.
In: The 33rd ACM/SIGAPP Symposium On Applied Computing. ACM. 2018.

[2] Ye Xia, Xavier Etchevers, Loïc Letondeur, Adrien Lebre, Thierry Coupaye and Frédéric Desprez
Combining Heuristics to Optimize and Scale the Placement of IoT Applications in the Fog.
In: 11th IEEE/ACM International Conference on Utility and Cloud Computing. IEEE/ACM. 2018.

[3] Loïc Letondeur, François-Gaël Ottogalli and Thierry Coupaye.
A demo of application lifecycle management for IoT collaborative neighborhood in the fog - Practical experiments and lessons learned around docker.
In: Fog World Congress. IEEE. 2017.

[4] Casanova Henri, Giersch Arnaud, Legrand Arnaud, Quinson Martin and Suter Frédéric
Versatile, scalable, and accurate simulation of distributed applications and platforms.
In: Journal of Parallel and Distributed Computing 74.10 (2014), pp. 2899–2917.