Generating hierarchical overlays for exponential topologies

P. Verghelet, E. Mocskos

Laboratorio de Sistemas Complejos
Departamento de Computación
FCEN - UBA

Information Resources in Large Distributed Systems

- Public collaboration across many countries and networks, with geographically distributed and heterogeneous resources
- Scheduling requires **up-to-date information** about the state of the system

Information Resources in Large Distributed Systems

- Public collaboration across many countries and networks, with geographically distributed and heterogeneous resources
- Scheduling requires **up-to-date information** about the state of the system
- Dynamics of resource information cannot be captured using a **static hierarchy (defacto implementation)** in medium-to-large environments
- New policies needed. Growing interest in the interaction with **P2P paradigm**, pushing towards **automated and scalable** solutions
- In this scenario, results essential to characterize the performance and behavior of resource information distribution policies, being the hierarchical policy one of the most used

Information Resources in Large Distributed Systems

- Public collaboration across many countries and networks, with geographically distributed and heterogeneous resources
- Scheduling requires **up-to-date information** about the state of the system
- Dynamics of resource information cannot be captured using a **static hierarchy (defacto implementation)** in medium-to-large environments
- New policies needed. Growing interest in the interaction with **P2P paradigm**, pushing towards **automated and scalable** solutions
- In this scenario, results essential to characterize the performance and behavior of resource information distribution policies, being the hierarchical policy one of the most used

Objectives

- Finding a criteria for selecting the root is essential to allow its usage in large scale generated scenarios.
- This work proposes a procedure to automatically select the root of a generated hierarchy obtaining reasonably good performance for the hierarchical policy in a wide range of analyzed scenarios.

Evaluation of system performance

LIR: captures the amount of information that a particular host has from all the entire system in a single moment. For the host k, LIR_k is:

$$LIR_k = \frac{\sum_{h=1}^{N} f(age_h, expiration_h) \cdot resourceCount_h}{totalResourceCount}$$

Evaluation of system performance

LIR: captures the amount of information that a particular host has from all the entire system in a single moment. For the host k, LIR_k is:

$$LIR_k = \frac{\sum_{h=1}^{N} f(age_h, expiration_h) \cdot resourceCount_h}{totalResourceCount}$$

GIR: captures the amount of information that the whole system knows about its resources, calculated as the mean value of every node's LIR.

Testing environment

- To evaluate different scenarios and policies we used GridMatrix, an open source tool based on SimGrid.

Testing environment

- To evaluate different scenarios and policies we used GridMatrix, an open source tool based on SimGrid.
- Using a static scenario where nodes have different amount of resources, but all resources have the same computing power and all links have the same attributes.

Testing environment

- To evaluate different scenarios and policies we used GridMatrix, an open source tool based on SimGrid.
- Using a static scenario where nodes have different amount of resources, but all resources have the same computing power and all links have the same attributes.
- Three different topologies are analyzed: Ring, Clique y Exponential.

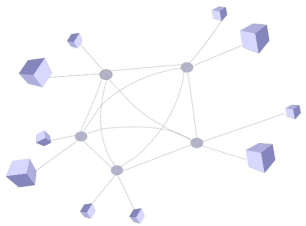
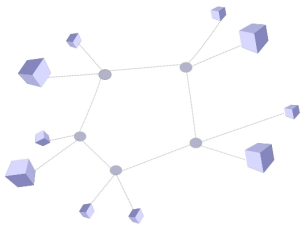
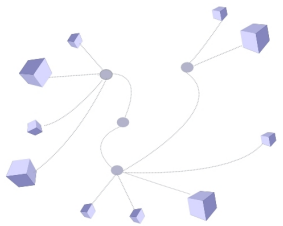

Testing environment

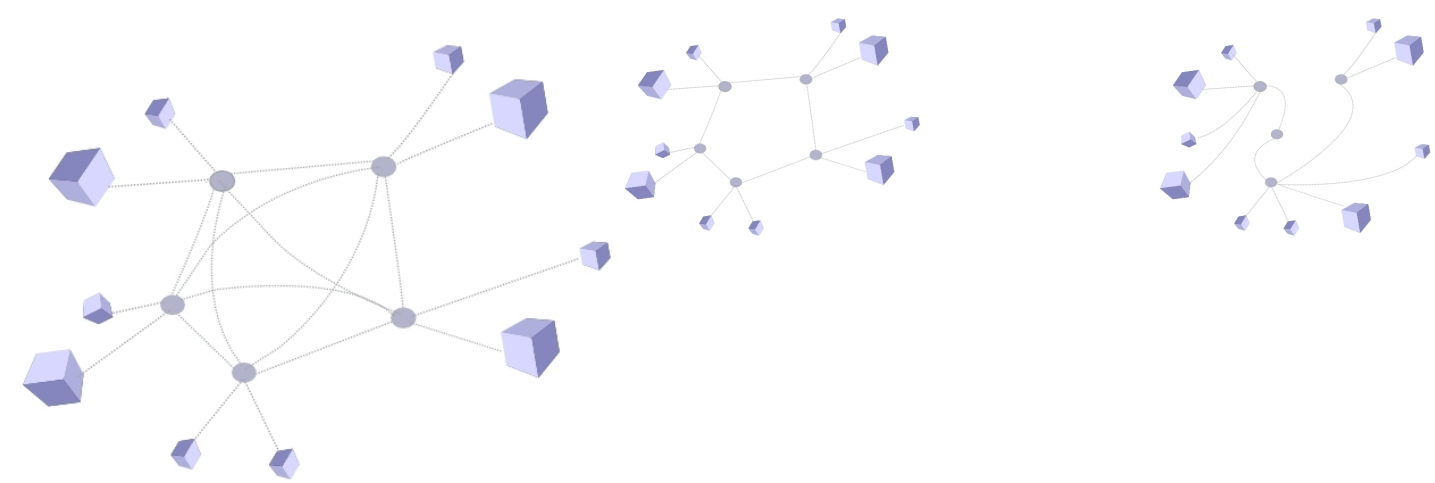

Clique

Clique topology with 5 routers and 10 nodes.
All nodes have the same degree and every node is connected with all other nodes.

Testing environment

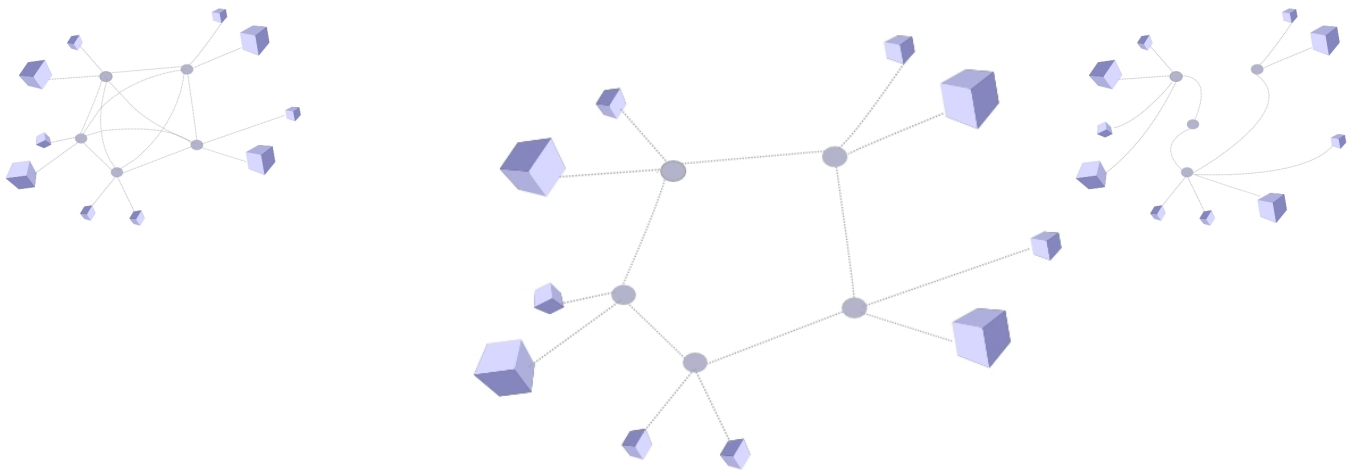

Ring

Ring (or circle) topology with 5 routers and 10 nodes.
All nodes have the same degree (i.e. 2).

Testing environment

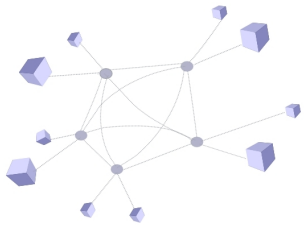

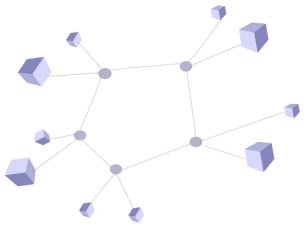

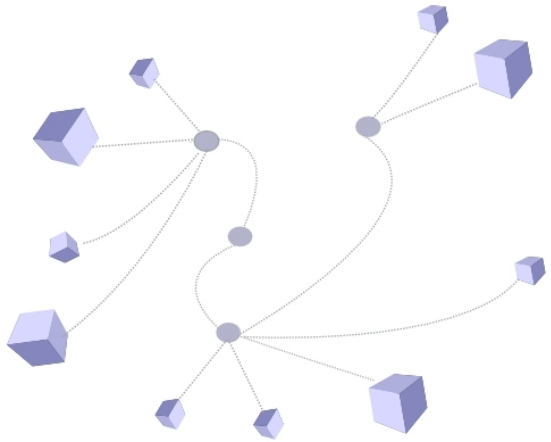

Exponential

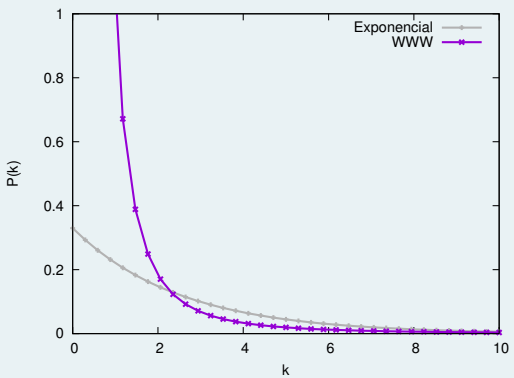

Exponential topology with 5 routers and 10 nodes. This topology resembles the structure of World Wide Web (Internet) and some social networks.
Distributions for nodes' degree:
$p(k) = k^{-\gamma}$ for WWW.
$p(k) = (1 - e^{-\frac{1}{\gamma}})e^{-\frac{k}{\gamma}}$ for exponential topology

Performance

- Evolution of GIR in a **Clique** topology with 15 routers and 30 nodes.

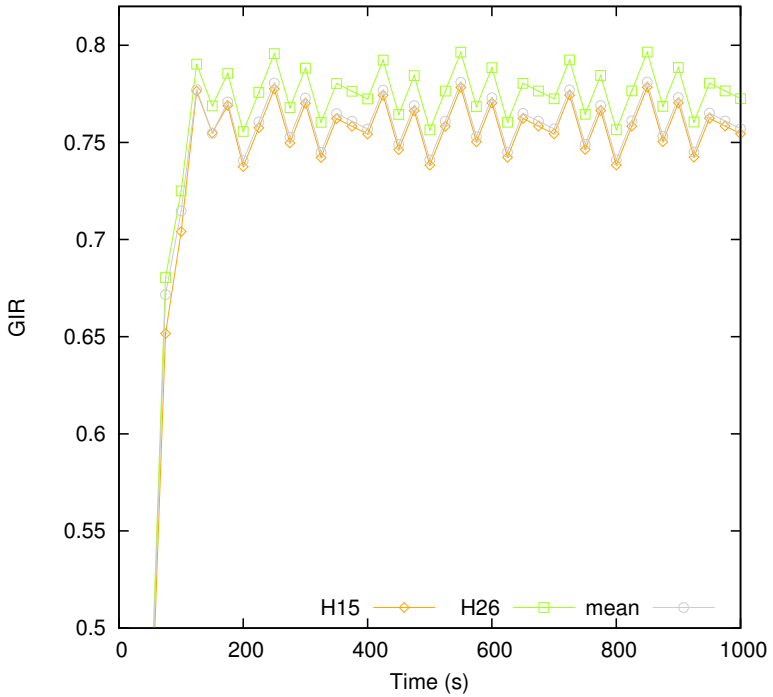

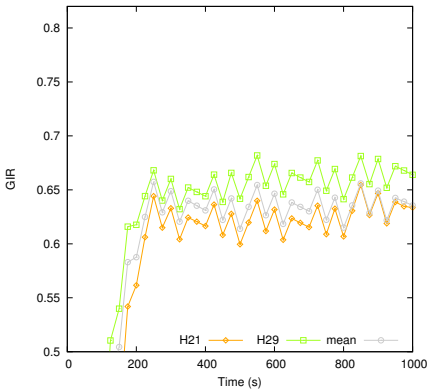

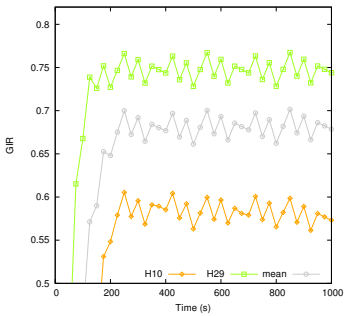

Performance

- Evolution of GIR in a **Ring** topology with 15 routers and 30 nodes.

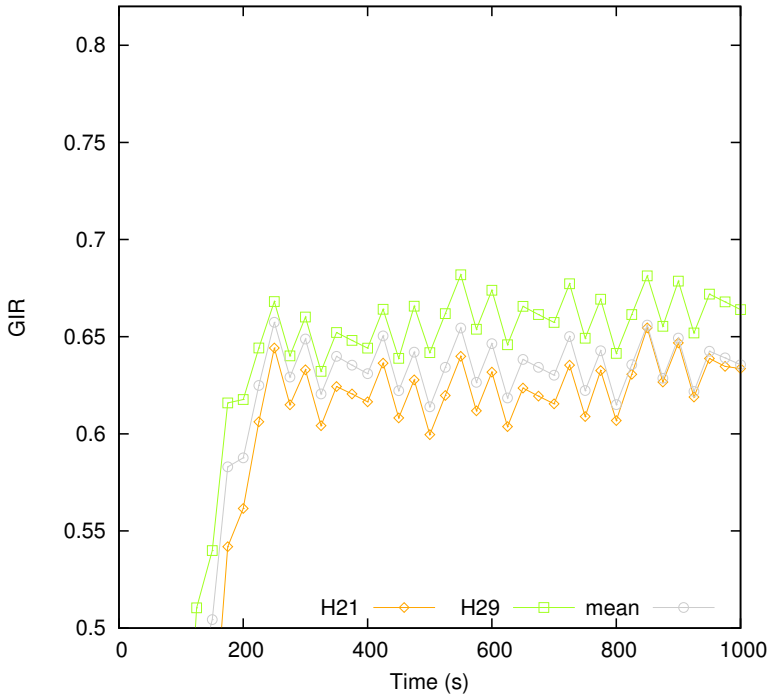

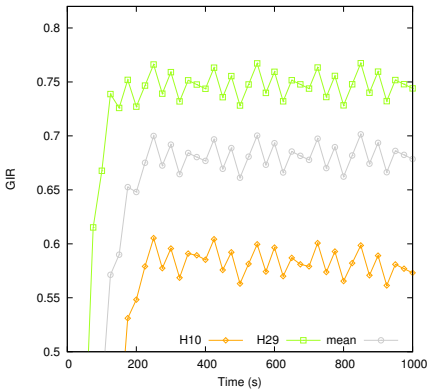

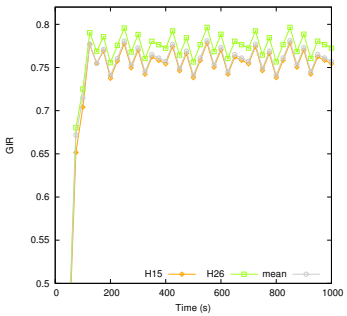

Performance

- Evolution of GIR in an **Exponential** topology with 15 routers and 30 nodes.

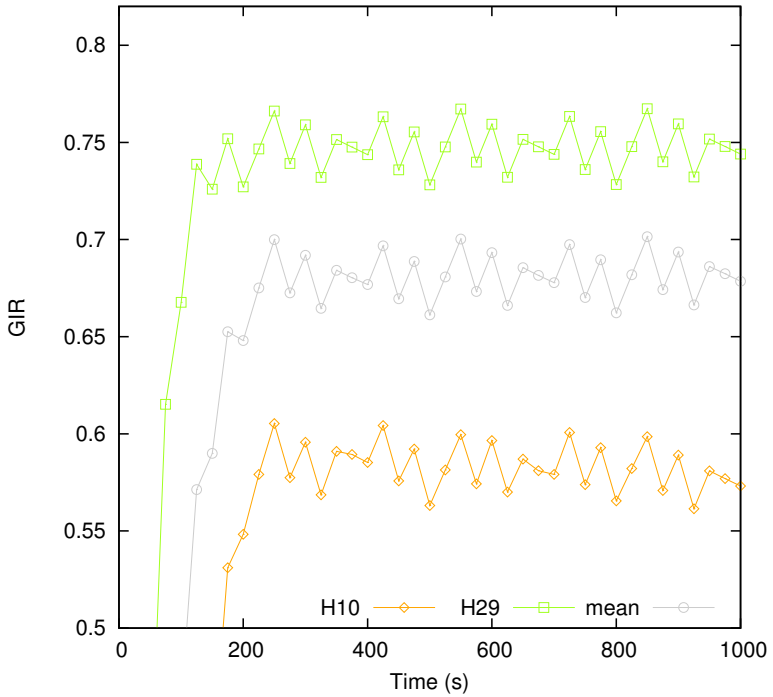

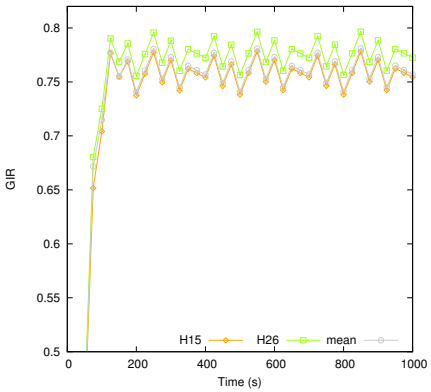

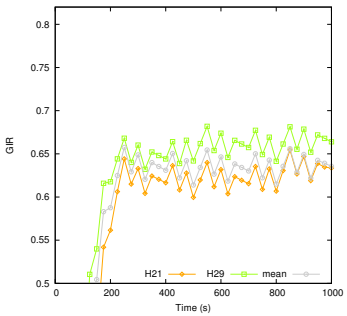

Hierarchical's performance?

- Different nodes selected as root of the hierarchy produces strong differences in the obtained GIR.

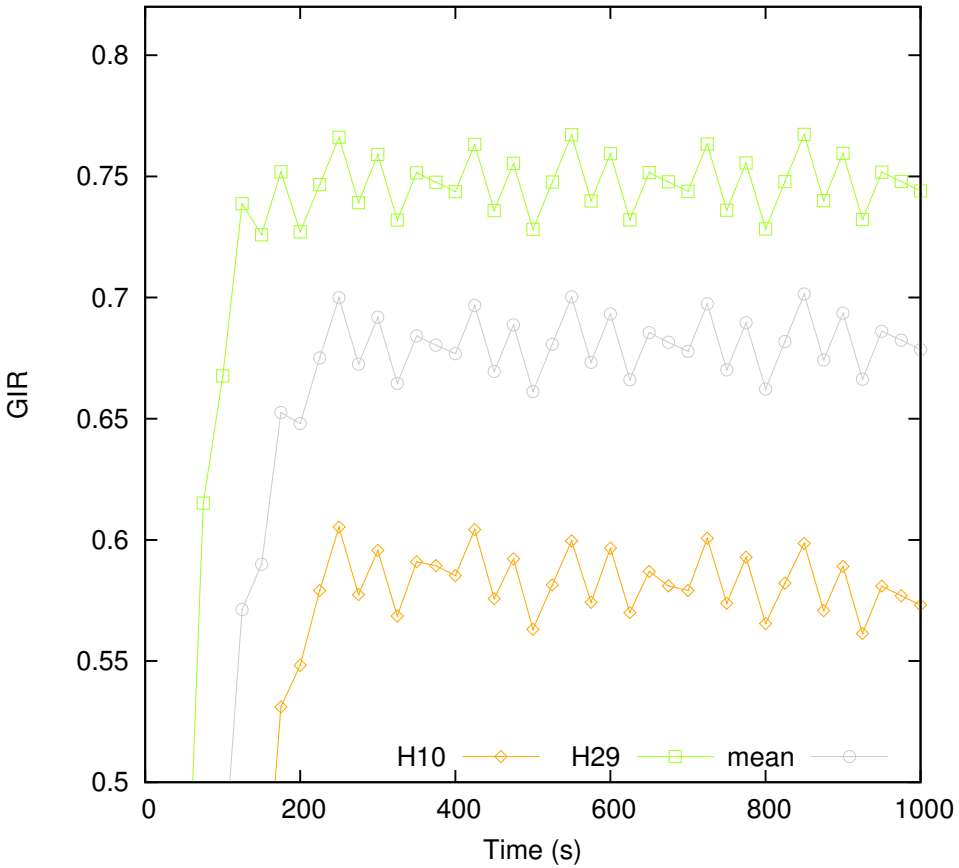

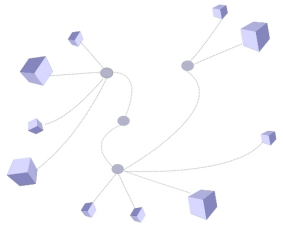

Figure: Evolution of GIR in an exponential topology with 30 nodes.

Hierarchical's performance?

We analyze some factors that may explain these results:

Hierarchical policy overlay
(orange dot lines)

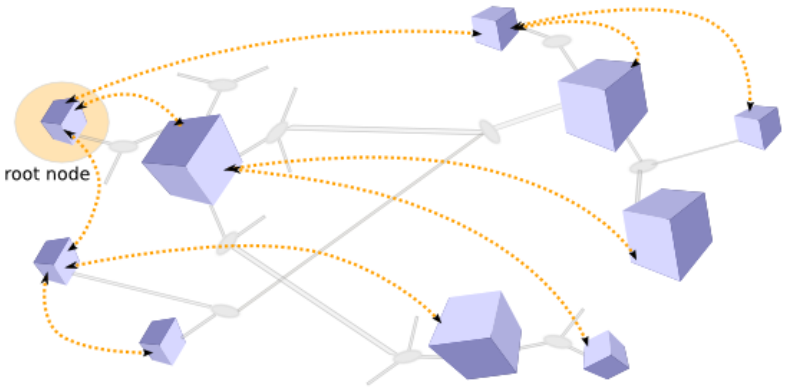

- 3-level hierarchical overlay
- Only the root is in the 1st level
- Starting from the levels on physical network, the algorithm folds levels until the amount of levels is less than 3
- For each level assign nodes in lower level to nearest node with least assigned nodes

Hierarchical's performance?

We analyze some factors that may explain these results:

Hierarchical policy overlay
(orange dot lines)

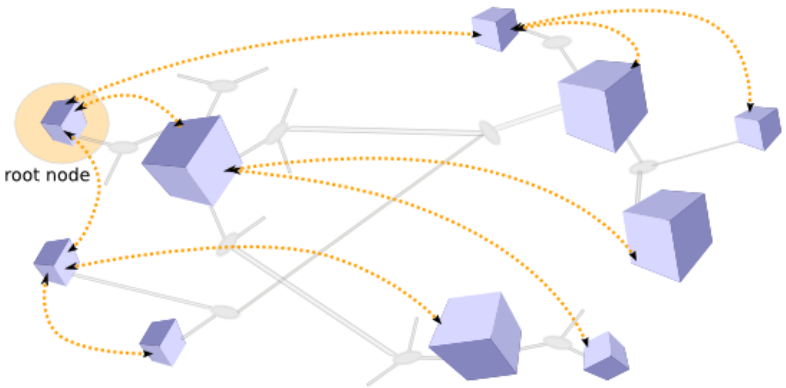

- 3-level hierarchical overlay
- Only the root is in the 1st level
- Starting from the levels on physical network, the algorithm folds levels until the amount of levels is less than 3
- For each level assign nodes in lower level to nearest node with least assigned nodes

The creation of the hierarchical overlay is similar to building the Minimum Broadcast Tree (mbt) [5,6]. Finding the mbt is a NP problem, exploring heuristics to treat it can be useful for constructing the hierarchy, but it escapes the scope of this work.

Hierarchical's performance?

Techniques were implemented to select the root

- i) the higher degree node
- ii) the node with the shortest maximum path to every other one (i.e. minimum eccentricity).

None of them presented satisfactory results (i.e. better or equal than selecting any other node).

Hierarchical's performance?

Techniques were implemented to select the root

- i) the higher degree node
- ii) the node with the shortest maximum path to every other one (i.e. minimum eccentricity).

None of them presented satisfactory results (i.e. better or equal than selecting any other node).

Centerum

The node with the minimum sum of distances to any other node is call *centerum*.

Centerum

- The observed results for exponential topologies of different sizes suggest that having more nodes in the second level of the hierarchy present better GIR (gray dots).

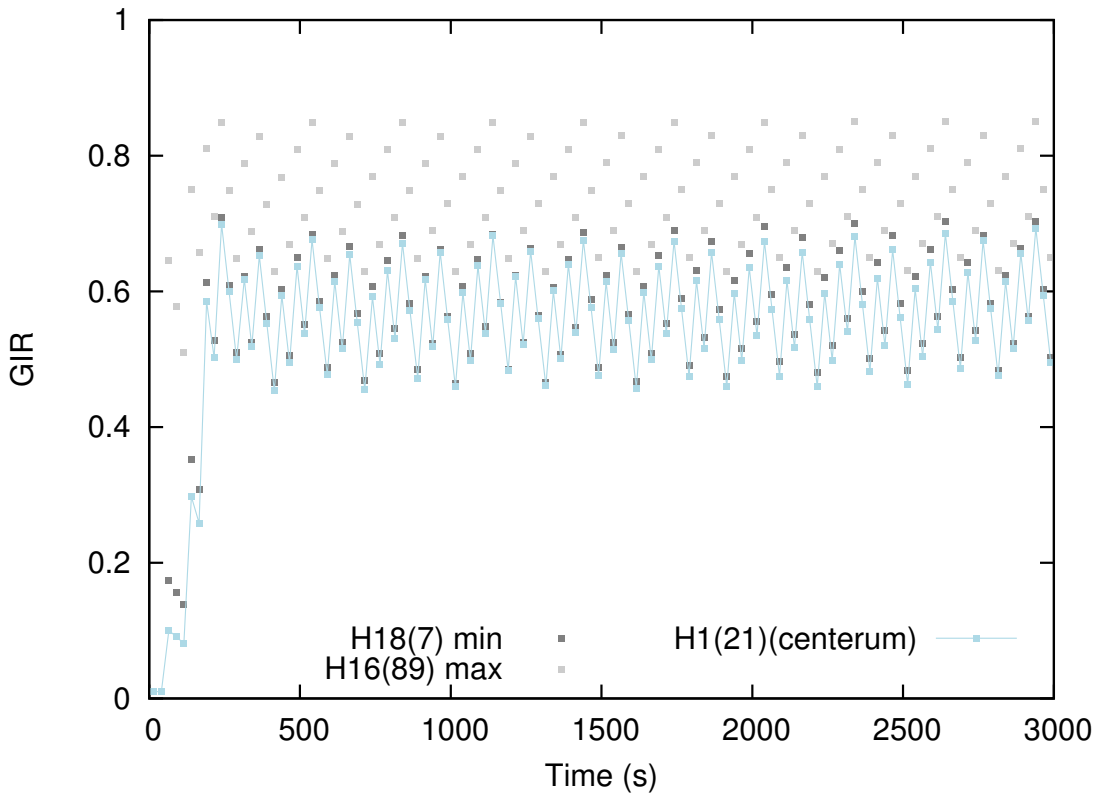

Centerum

- The observed results for exponential topologies of different sizes suggest that having more nodes in the second level of the hierarchy present better GIR (gray dots).
- Only a small fraction of the nodes generates this type of overlay.

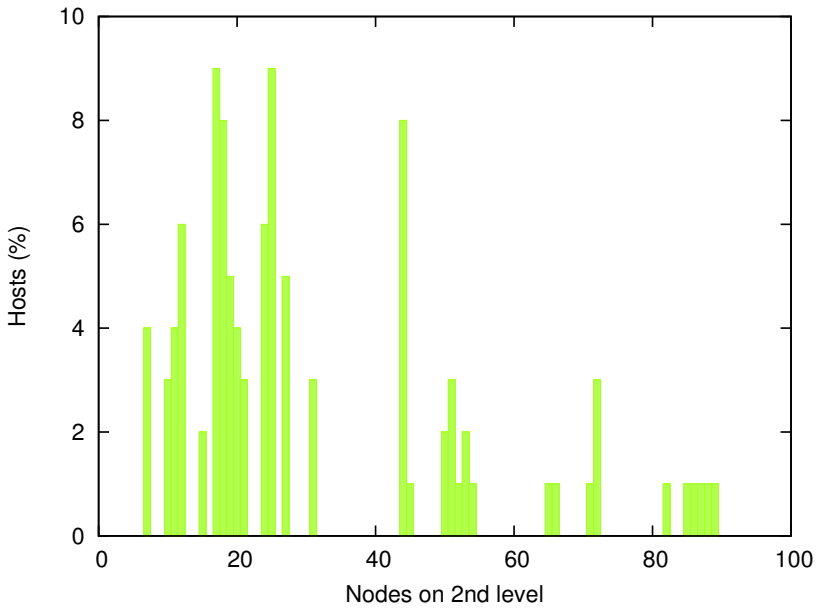

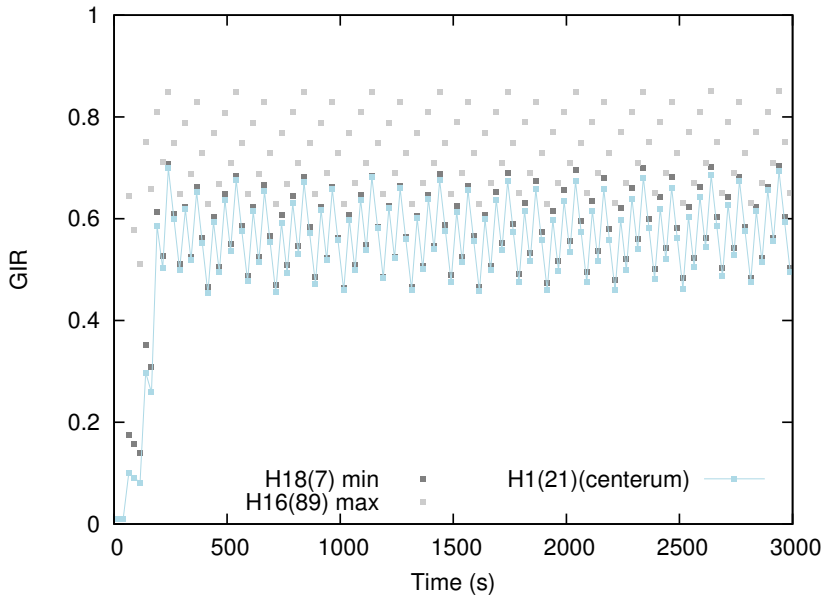

Hierarchical's performance

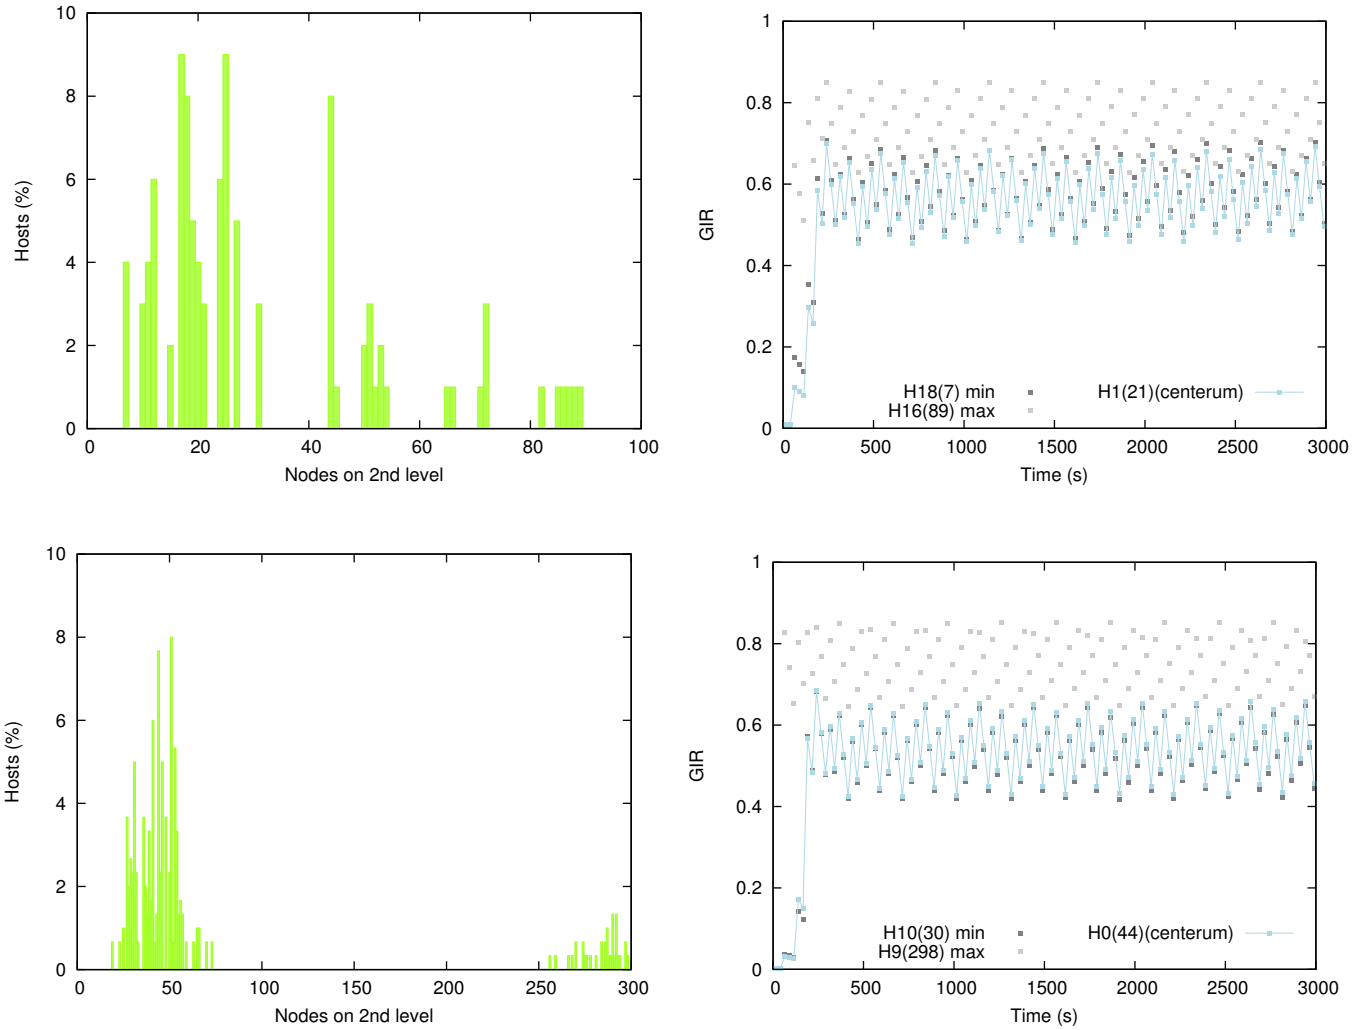

Figure: Evolution of GIR in an exponential topology with 100 nodes (first row) and 300 nodes (second row)

Hierarchical's performance

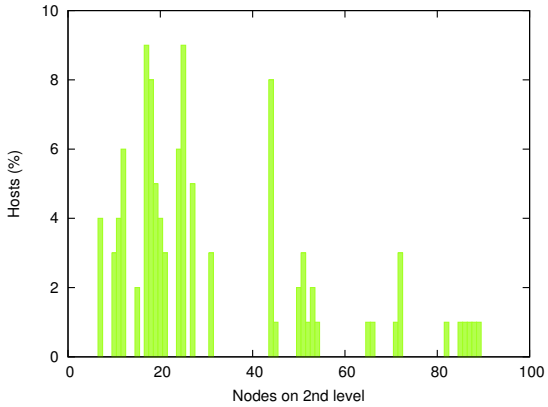

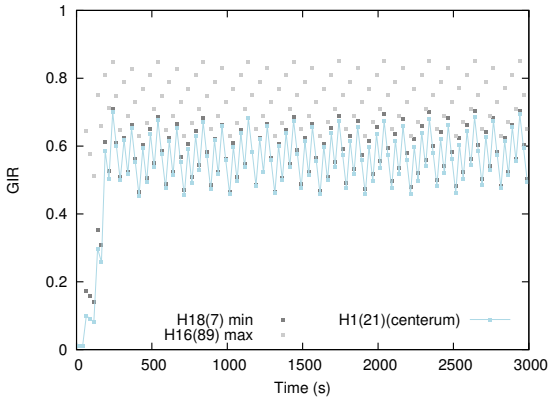

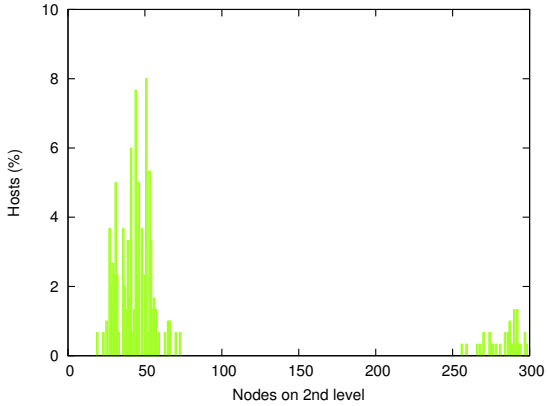

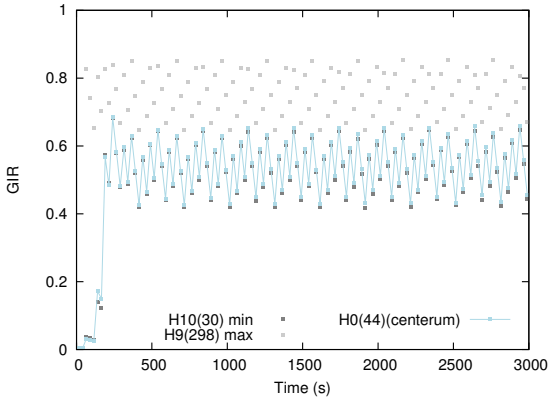

The **Centerum** allows to obtain a representative overlay to be used as comparison for fully distributed policies.

Conclusions

1. Selecting the root for the hierarchy does not impact for clique and ring topologies.
2. The selection of the root for exponential topology severely impacts the system behavior.

Conclusions

1. Selecting the root for the hierarchy does not impact for clique and ring topologies.
2. The selection of the root for exponential topology severely impacts the system behavior.

- Having more nodes in the second level of the hierarchy seems to improve the system performance (resembling a centralized topology).

Conclusions

1. Selecting the root for the hierarchy does not impact for clique and ring topologies.
2. The selection of the root for exponential topology severely impacts the system behavior.

- Having more nodes in the second level of the hierarchy seems to improve the system performance (resembling a centralized topology).
- The centerum allows to obtain a representative overlay to be used as comparison for fully distributed policies.

Conclusions

1. Selecting the root for the hierarchy does not impact for clique and ring topologies.
2. The selection of the root for exponential topology severely impacts the system behavior.

- Having more nodes in the second level of the hierarchy seems to improve the system performance (resembling a centralized topology).
- The centerum allows to obtain a representative overlay to be used as comparison for fully distributed policies.

Thank you !!!
Questions ???