

Magnolia: a novel DHT architecture for keyword-based search


Ashish Gupta, Manan Sanghi
 Peter A. Dinda, Fabian Bustamante
 Department of Computer Science,
 Northwestern University
 {ashish,manan, pdinda, fabianb}@cs.northwestern.edu




Why Magnolia ?

DHT based systems → big improvement over unstructured systems
 (1) $O(\log n)$ routing and lookup
 (2) Effective Load balancing

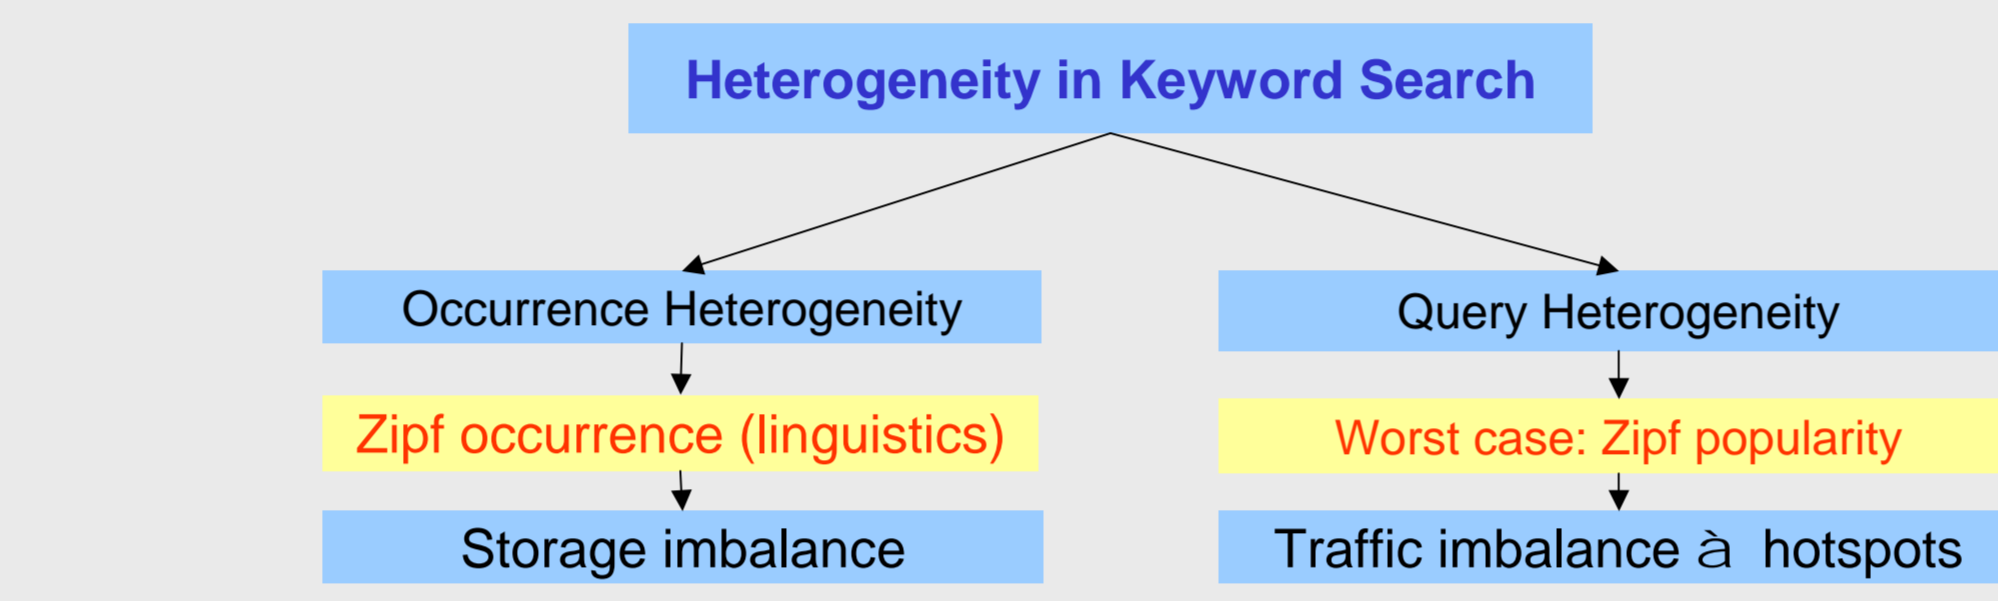
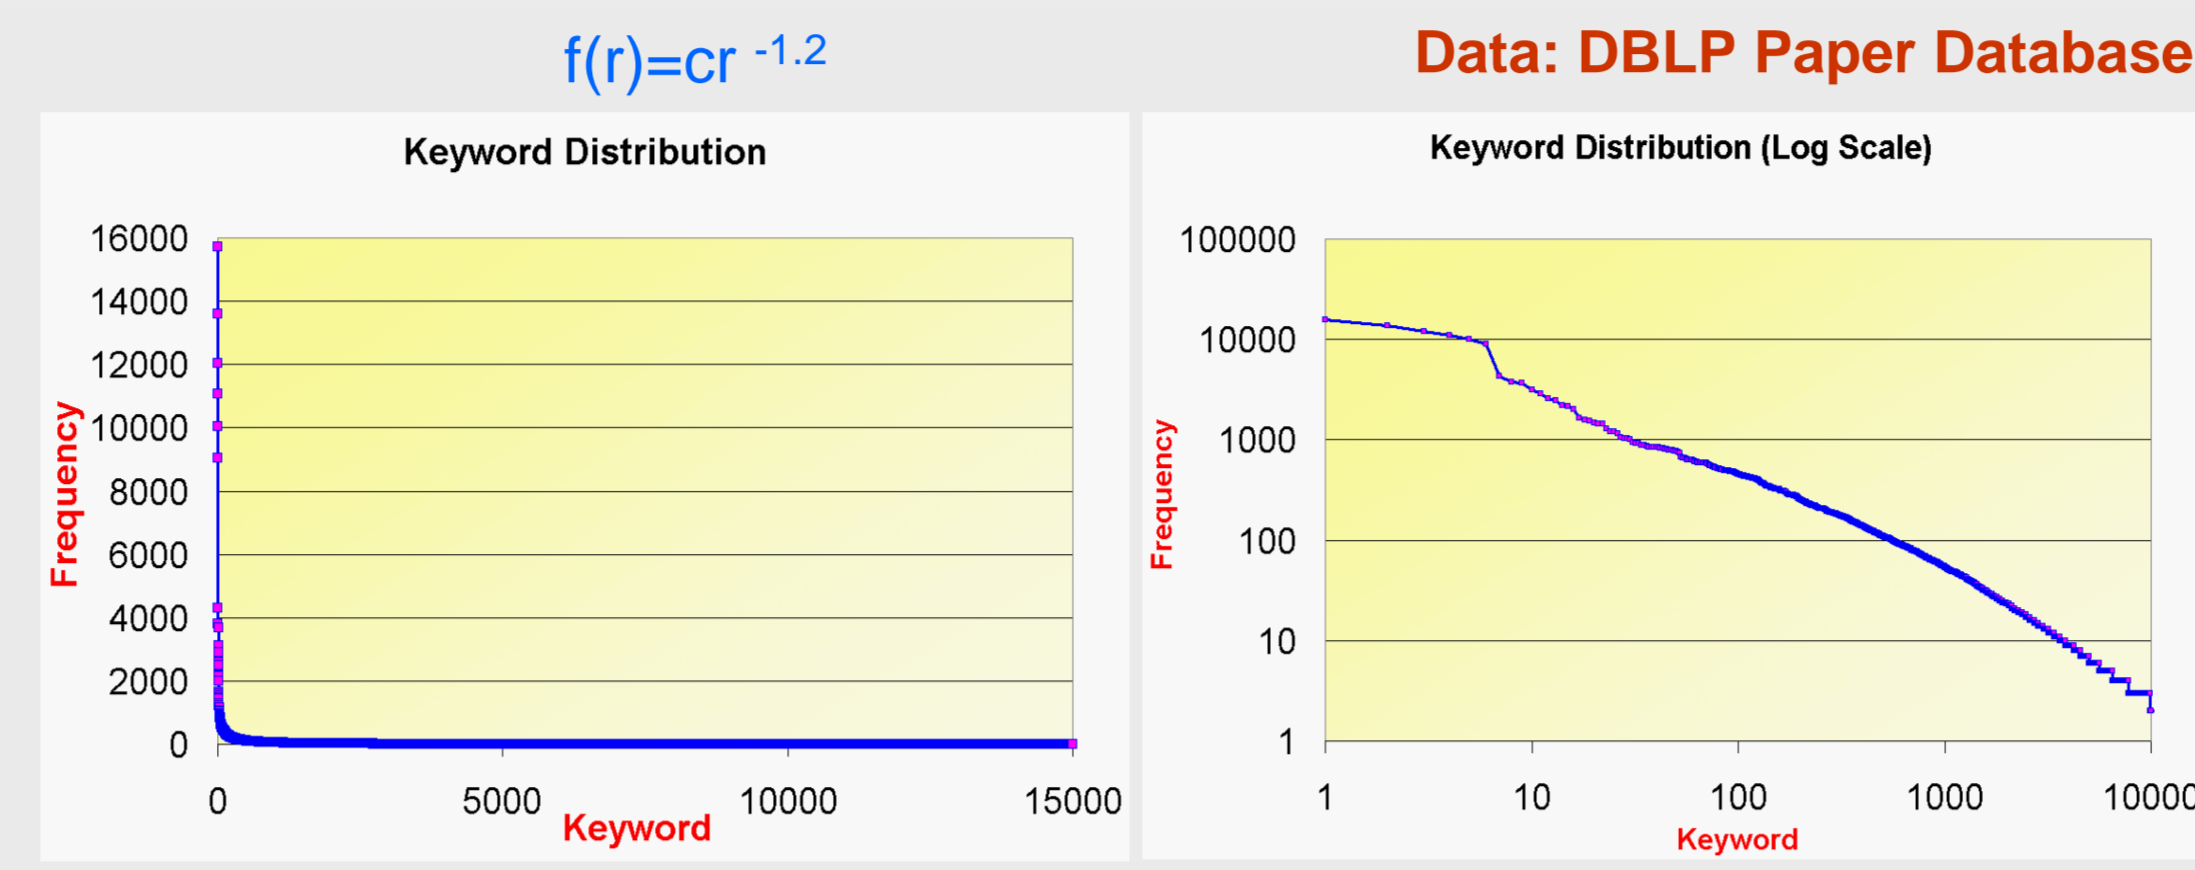
Problem → Keyword search difficult because of hashing based lookup



Current approach: Hash each keyword separately and store pointers at $h(\text{keyword})$

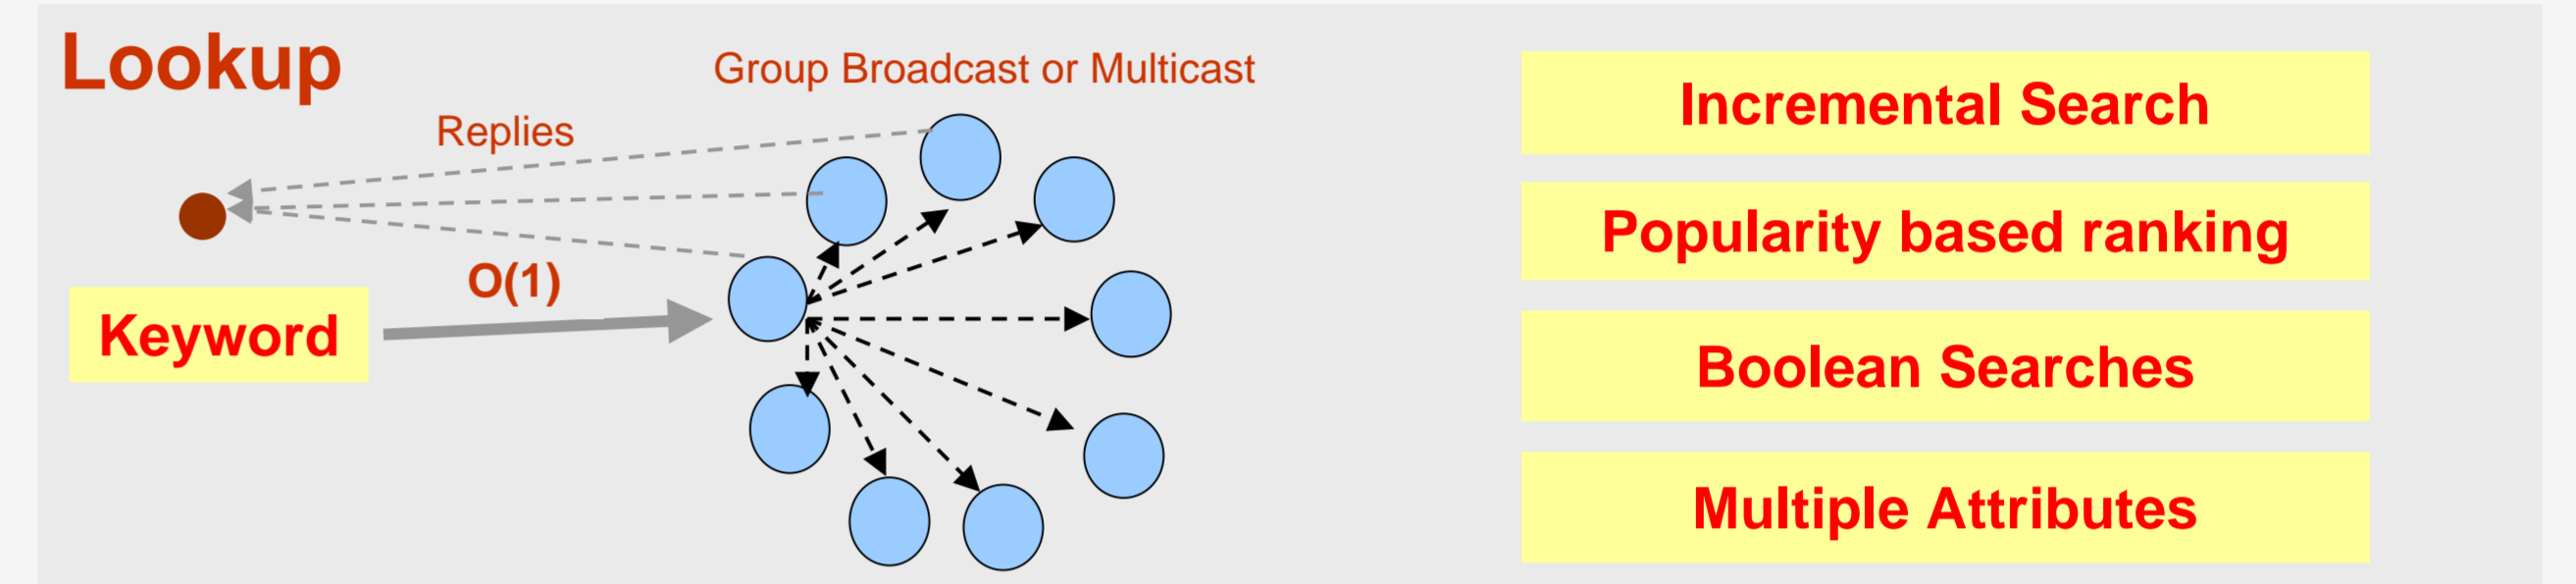
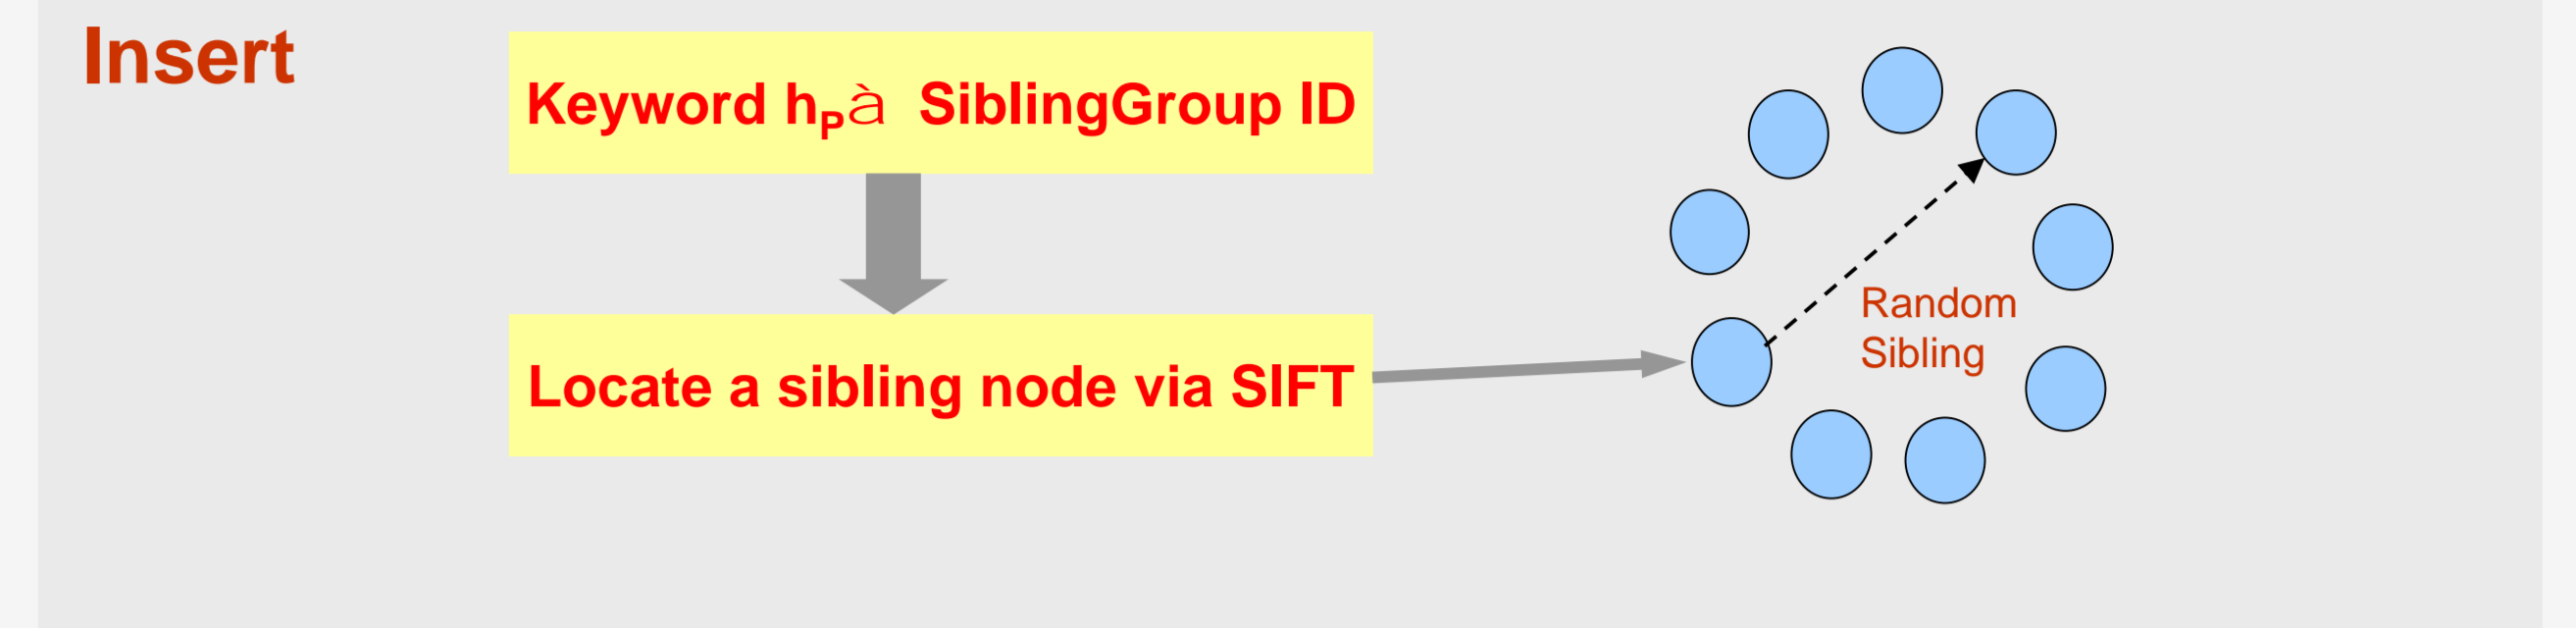


Need to balance out load and provide scalability, without sacrificing DHT properties (routing efficiency and state)



1. Node overload with keywords
2. Popular nodes fail : entire keyword set gone
3. Routing and Query hotspots for popular keywords

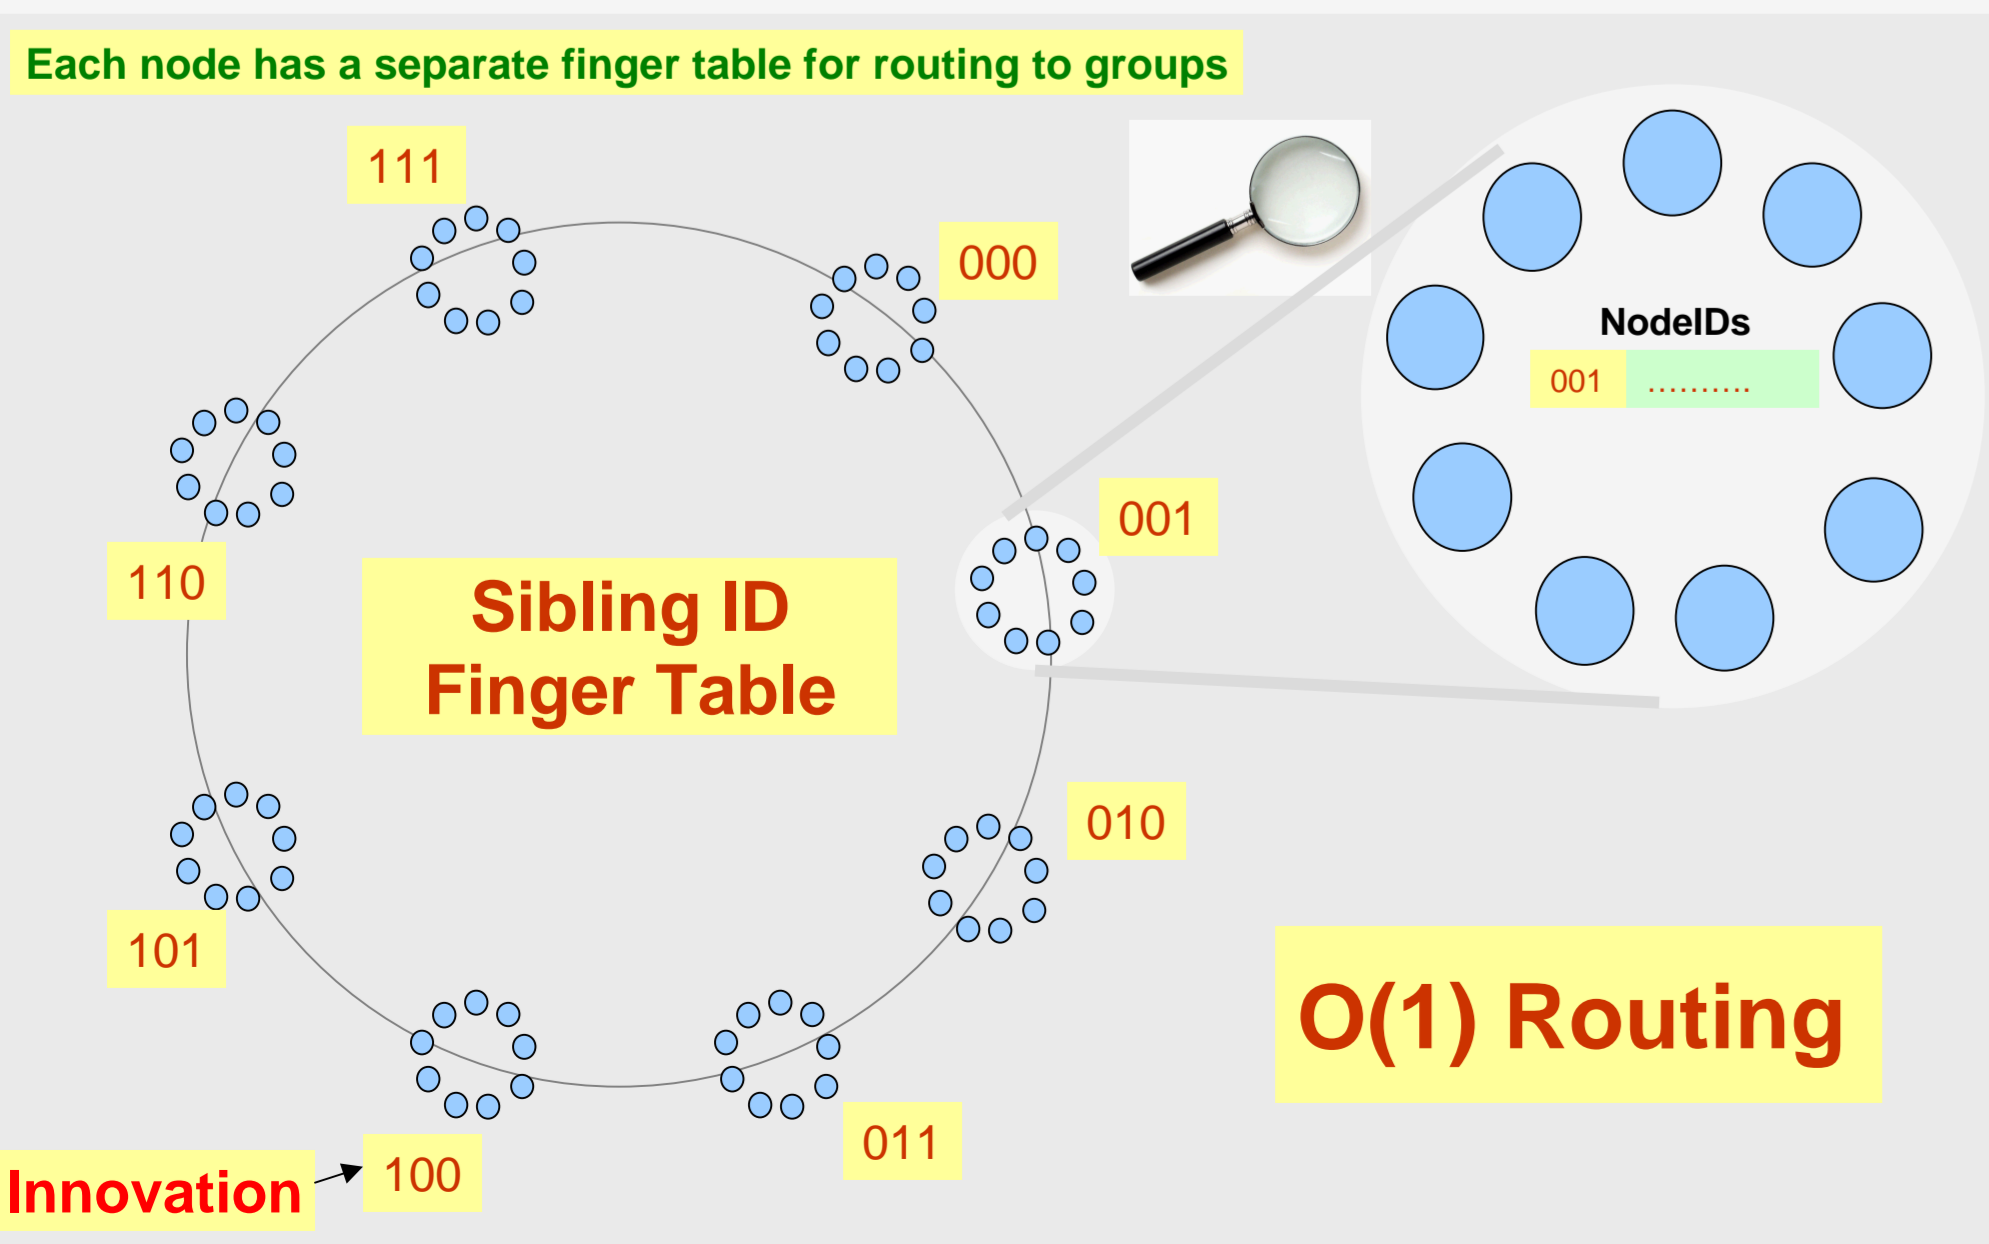
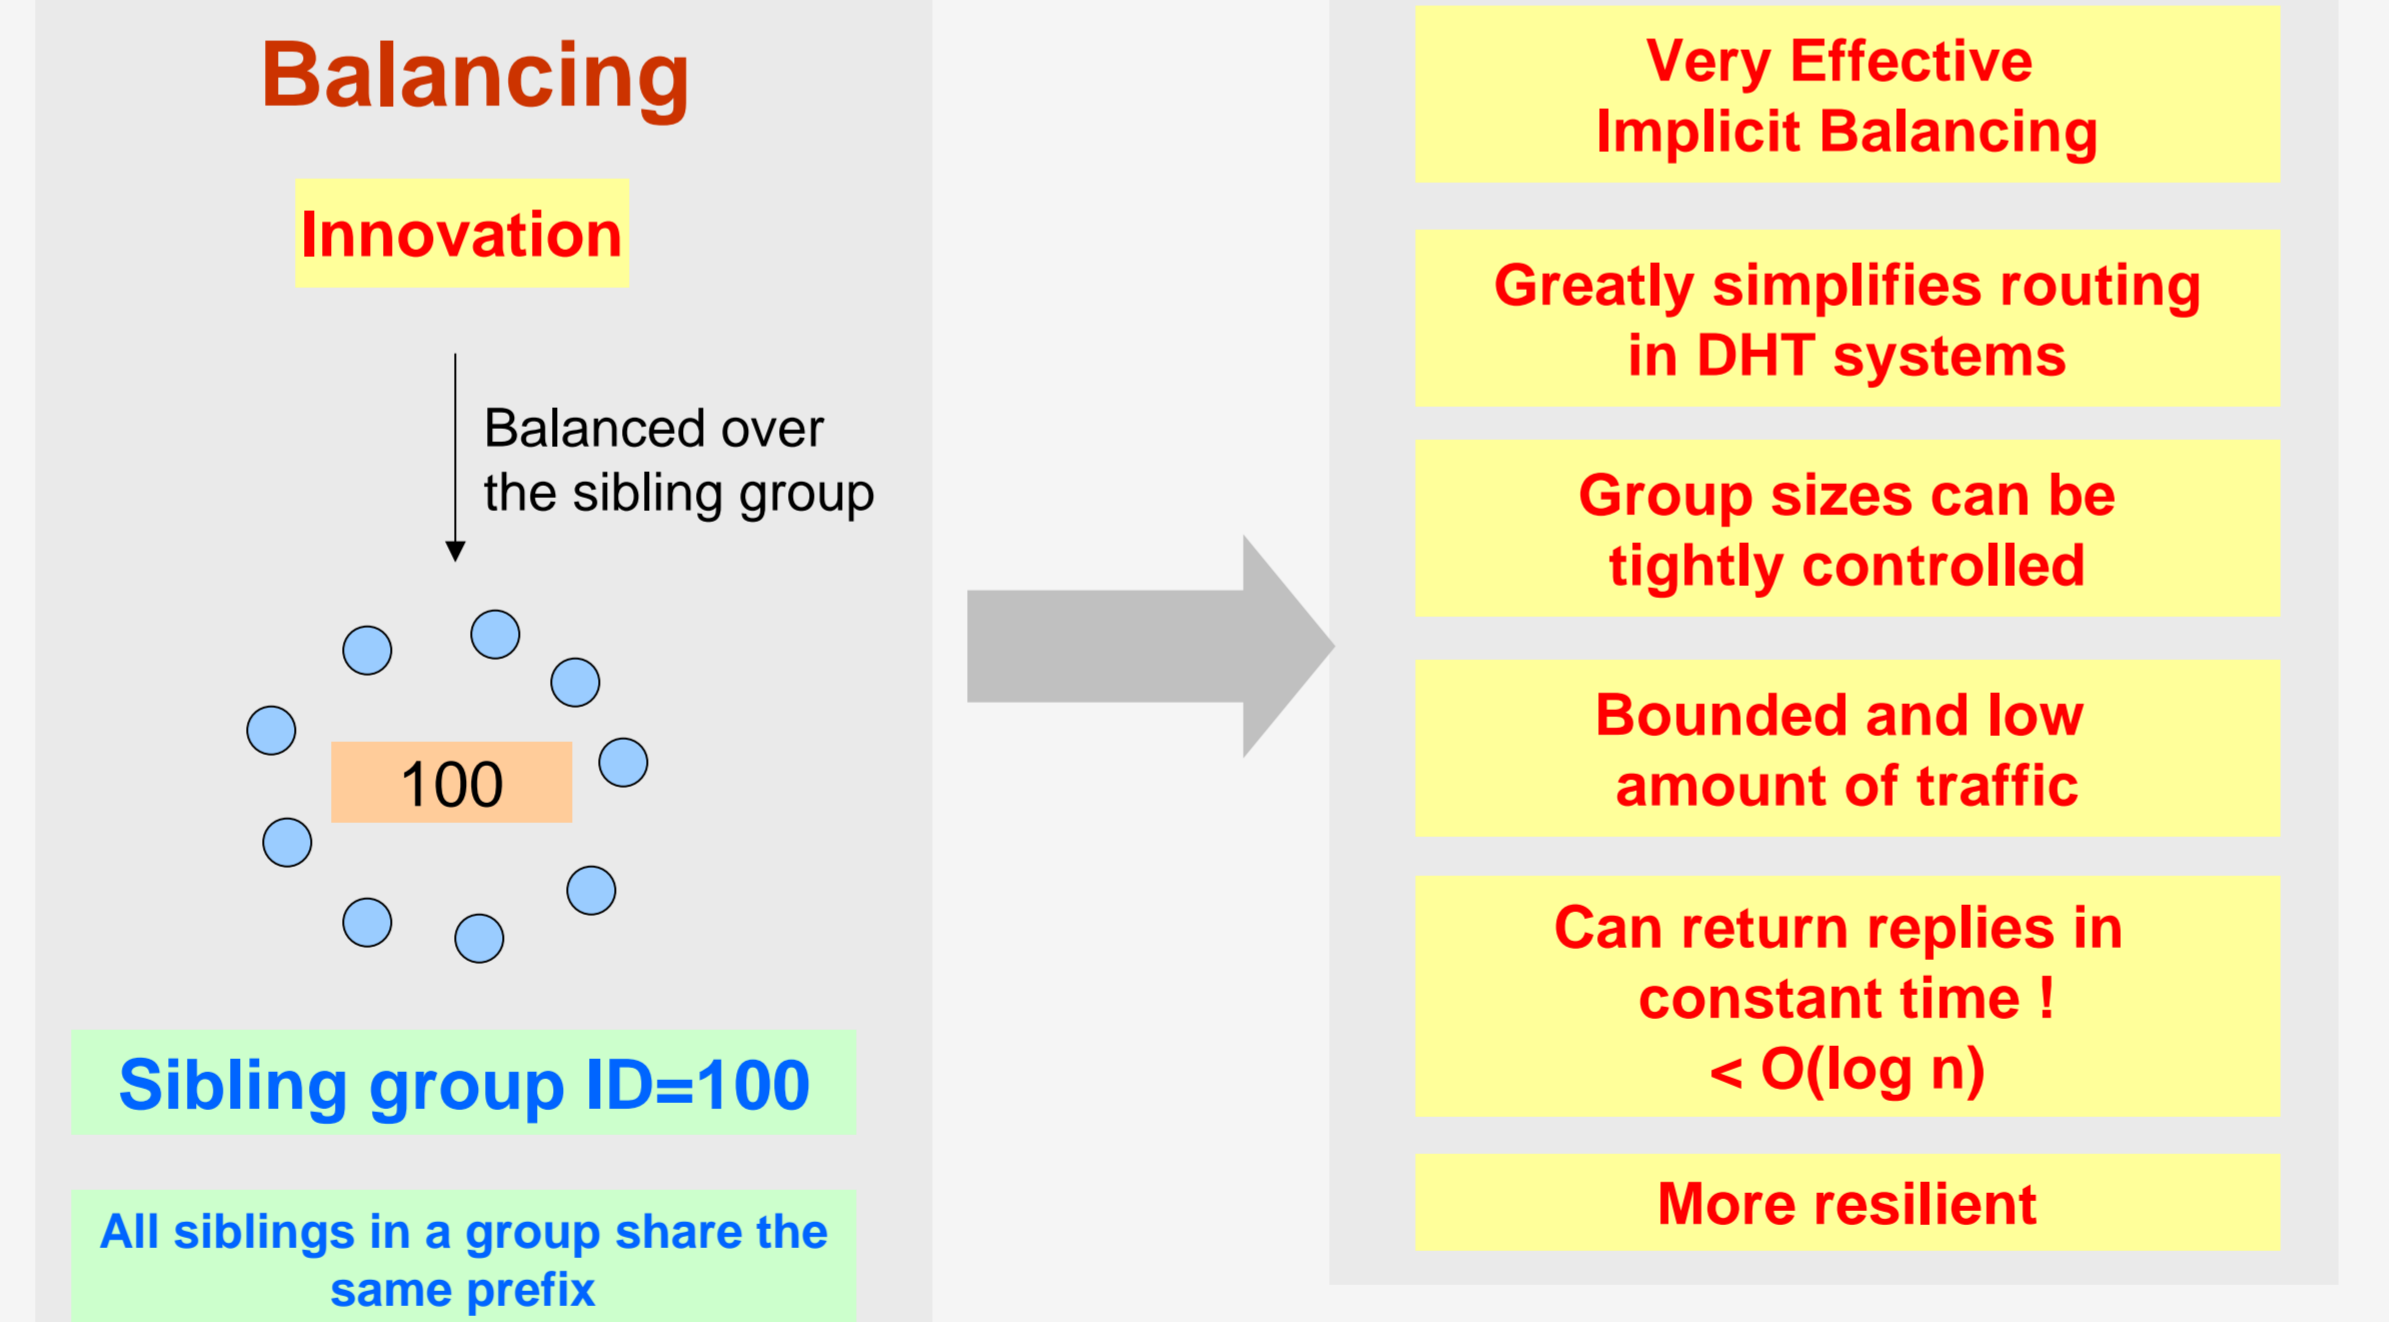
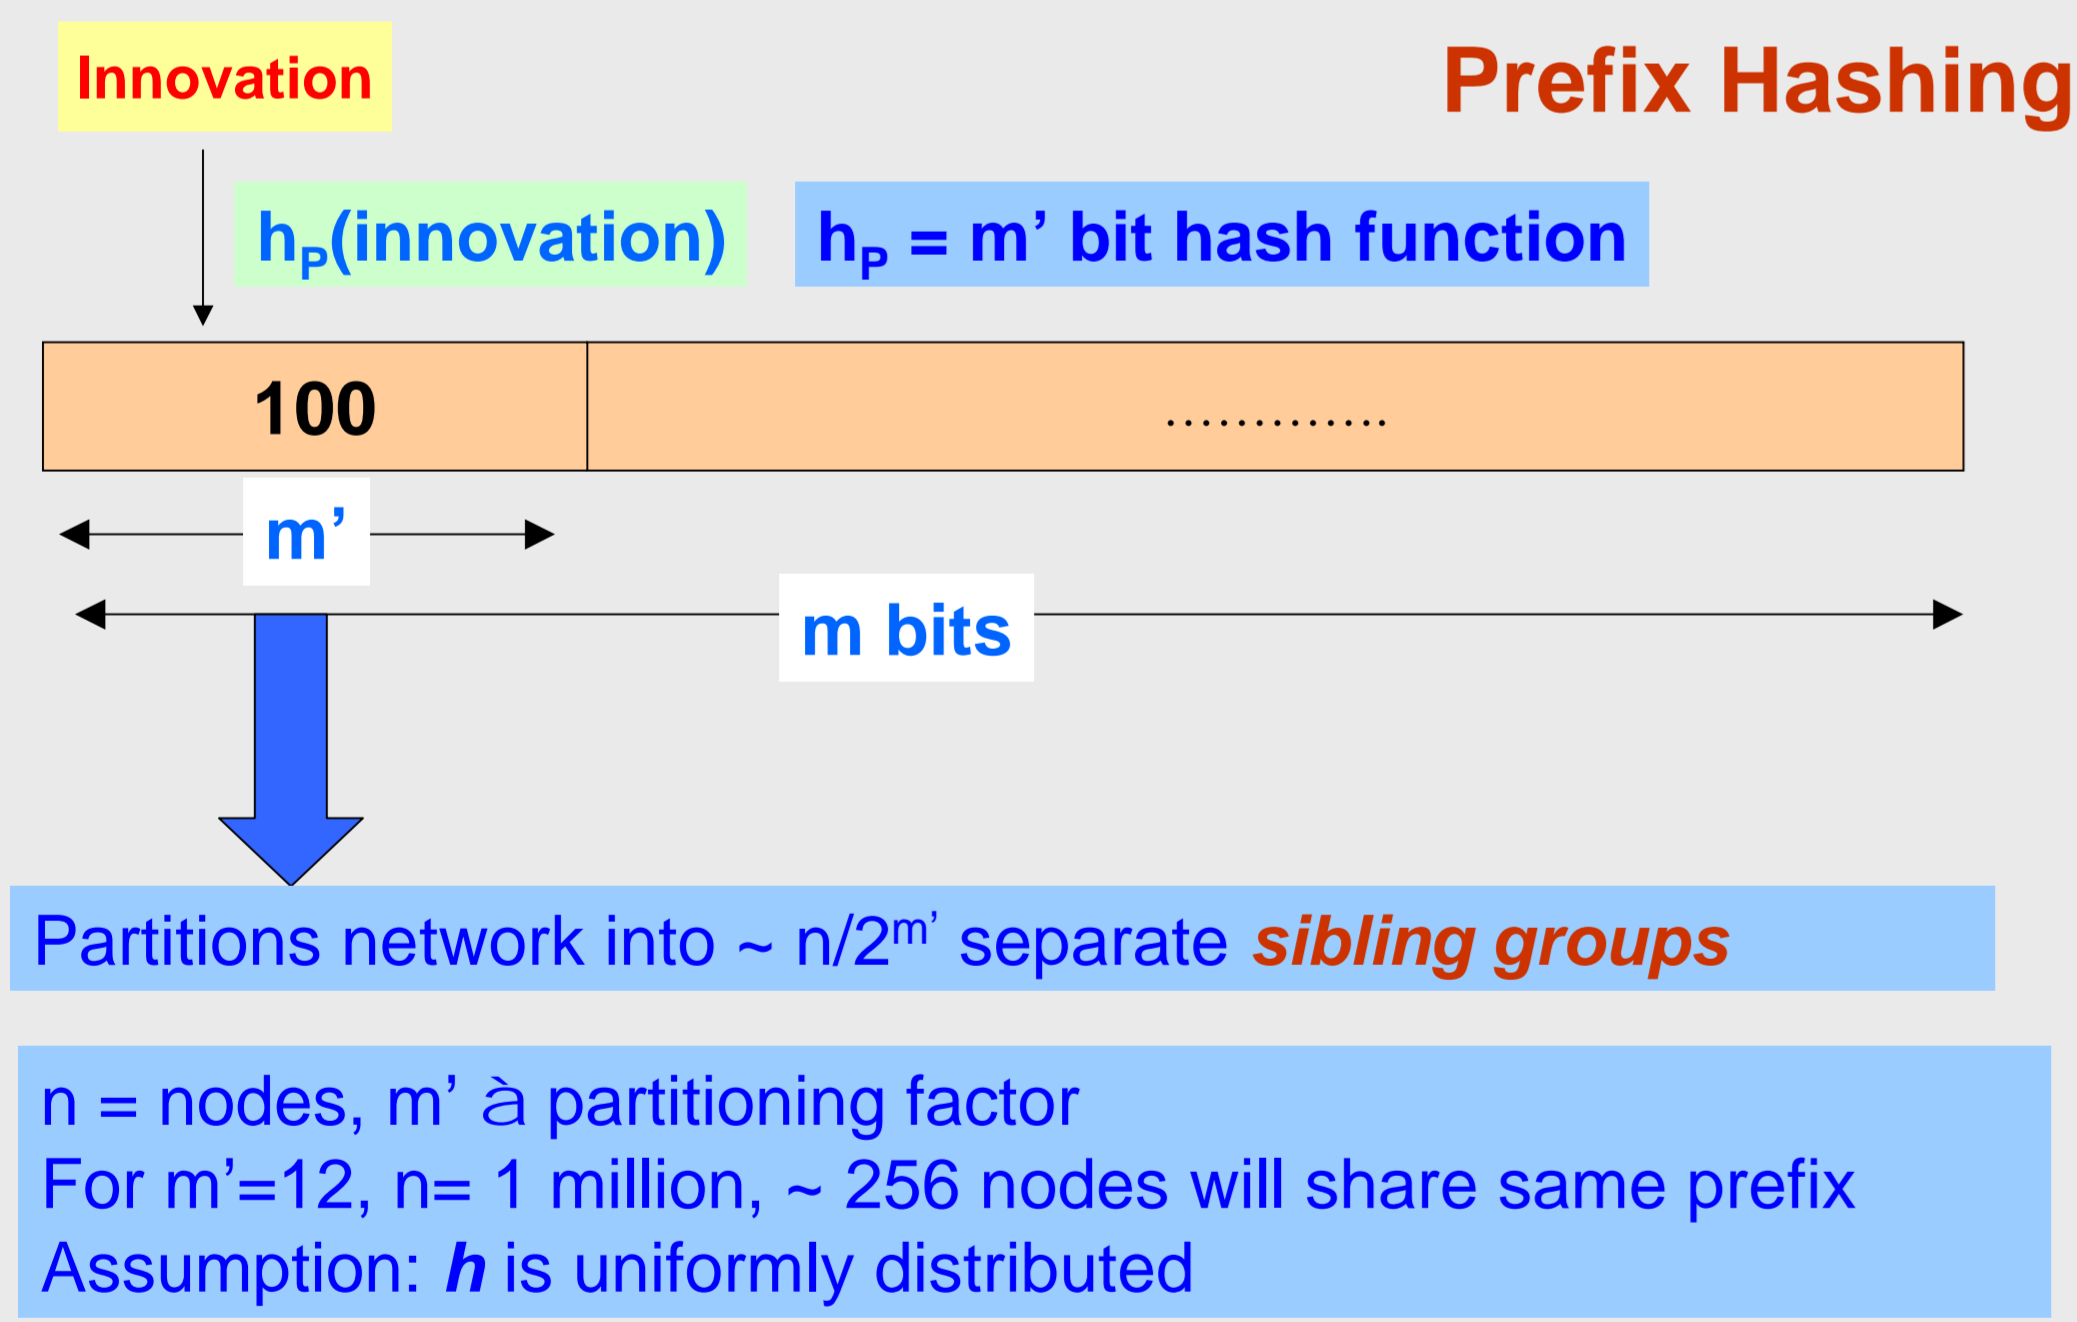
Storage and Lookup



Analytical Bounds m' (sibling ID length) = 12, n (number of nodes) = 1 million

- Routing and Lookup = $O(1)$
- Numbers of Nodes visited = $O(m' + n/2^{m'}) + r$
 ~ 268 nodes for complete query
- Total Traffic Generated = $O(m' + n/2^{m'}) + r$
 Traffic ~ 268 + r units for complete query
- Routing state per node = $O(\log n + 2^{m'} + n/2^{m'})$
 ~40 Kbytes

Prefix-based Hashing and Sibling Partitioning



Good Balancing Properties

Question: What is the variance of load over all the nodes ?

For a Zipf keyword distribution, final load variance can be computed from original distribution variance

$$V_{\text{final}} = \frac{V_{\text{orig}} \cdot k \cdot 2^{m'}}{n^2}$$

V_{orig} = Keyword occurrence variance in input
 V_{final} = Load variance over all nodes
 k = number of keywords

Simulation Results

