LOW-COST SEARCH IN TREE-STRUCTURED P2P OVERLAYS: THE NULL-BALANCE BENEFIT

Fraunhofer
IML

RUB

Peter Detzner | Jana Gödeke

Steffen Bondorf

06. October 2021

MOTIVATION

Peer-to-Peer (P2P) Networks

- Peer-to-Peer (P2P) networks are well-known and established in the research community
- Various applications have been established, e.g.
  - Open-source money (Bitcoin, …)
  - Online games
  - …
- P2P networks can be classified in
  - Unstructured
  - Structured

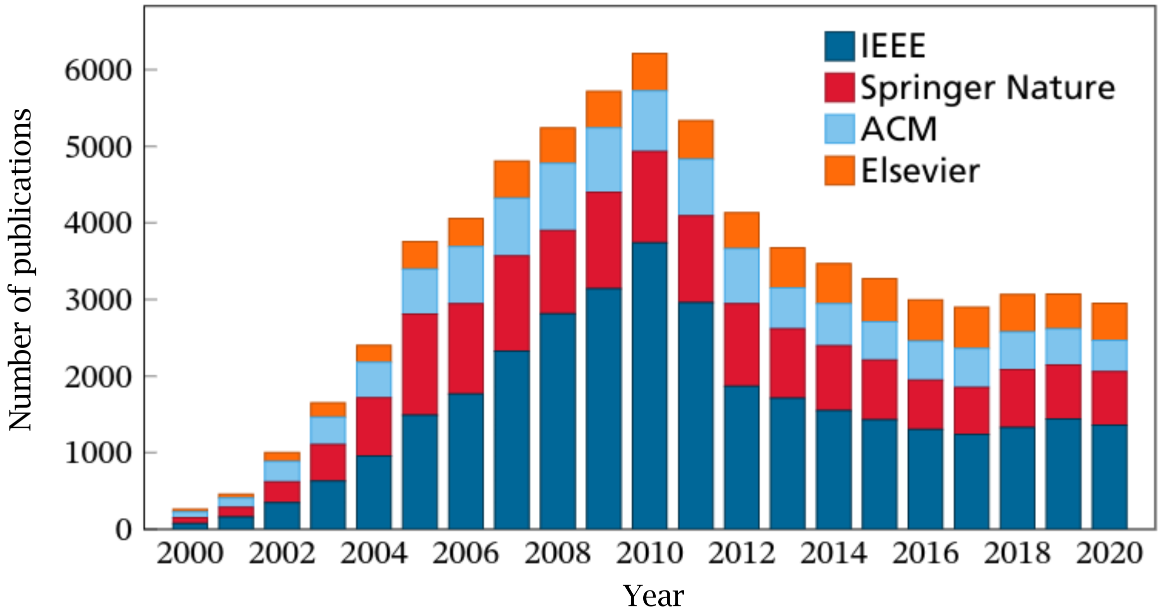

Figure 1: Number of publications over time

INTRODCUTION

BAlanced Overlay m-ary Tree Overlay Network* (BATON*)

- BATON* [1]
  - Based on BATON[2]
  - m-ary height-balanced tree:
    - *… at any node in the tree, the height of any two subtrees of its children differ by at most 1 …*
- System Model
  - Represented by (level:number), e.g. 4:8
  - Maintains links
    - Parent, up to m-children,
    - RoutingTable (RT) = LRT + RRT
    - Left- and right adjacent

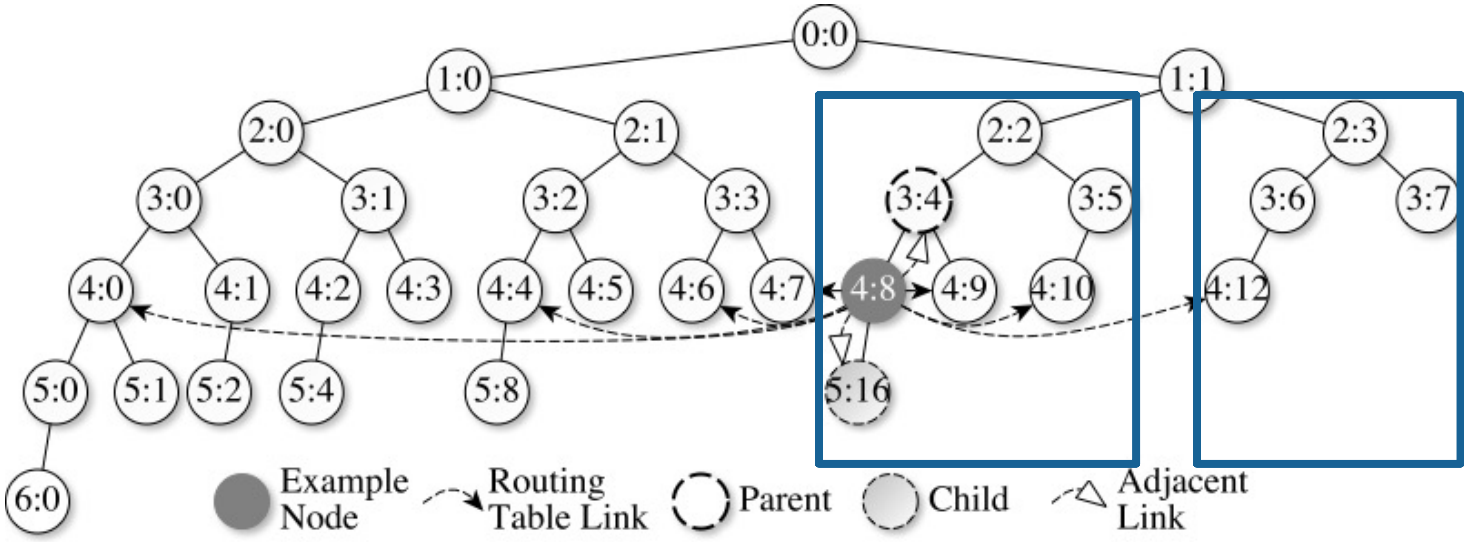

Figure 2: The BATON* tree has a height of 7. Since node (4:8) has links to neighbors (4:0) and (4:4), links to the children are (5:0), (5:1), (5:8)

[1] Jagadish, H.V., Ooi, B.C., Tan, K.-L., Vu, Q.H., Zhang, R., 2006. Speeding up search in peer-to-peer networks with a multi-way tree structure, in: Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data - SIGMOD '06. Presented at the the 2006 ACM SIGMOD international conference, ACM Press, Chicago, IL, USA, p. 1. https://doi.org/10.1145/1142473.1142475

[2] Jagadish, H.V., Ooi, B.C., Vu, Q.H., 2005. BATON: a balanced tree structure for peer-to-peer networks, in: Proceedings of the 31st International Conference on Very Large Data Bases, VLDB '05. VLDB Endowment, Trondheim, Norway, pp. 661–672.

Fraunhofer IML

An Analysis of BATON*

- Distributed Hash Table might require Load Balancing
- Rotation can lead to invalid RT entries
  - SearchCost will be infinite
  - Not deterministic
- Mutual Exclusion during join [3] during network restructuring
- Not an optimal height

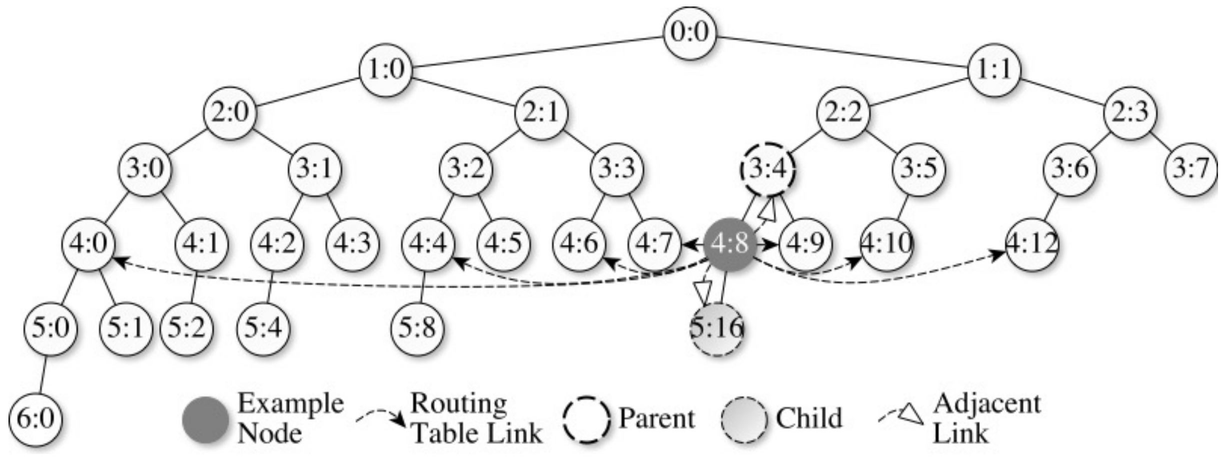

Figure 3: The BATON* tree has a height of 7. Since node (4:8) has links to neighbors (4:0) and (4:4), links to the children are (5:0), (5:1), (5:8).

[3] Masinde, N., Graffi, K., 2020. Peer-to-Peer-Based Social Networks: A Comprehensive Survey. SN COMPUT. SCI. 1, 299. https://doi.org/10.1007/s42979-020-00315-8

Fraunhofer IML

NULL-BALANCED M-ARY TREE OVERLAY NETWORKS

nBATON* - System Model

- Null-balanced[4]:
  - *Definition: "An m-ary tree is null-balanced if any two leaf nodes differ in level by 1"*

- Challenge:
  - Find a free spot
  - Minimizing the tree height
  - Keep the tree balanced

- RQ: Has the height an impact on search?

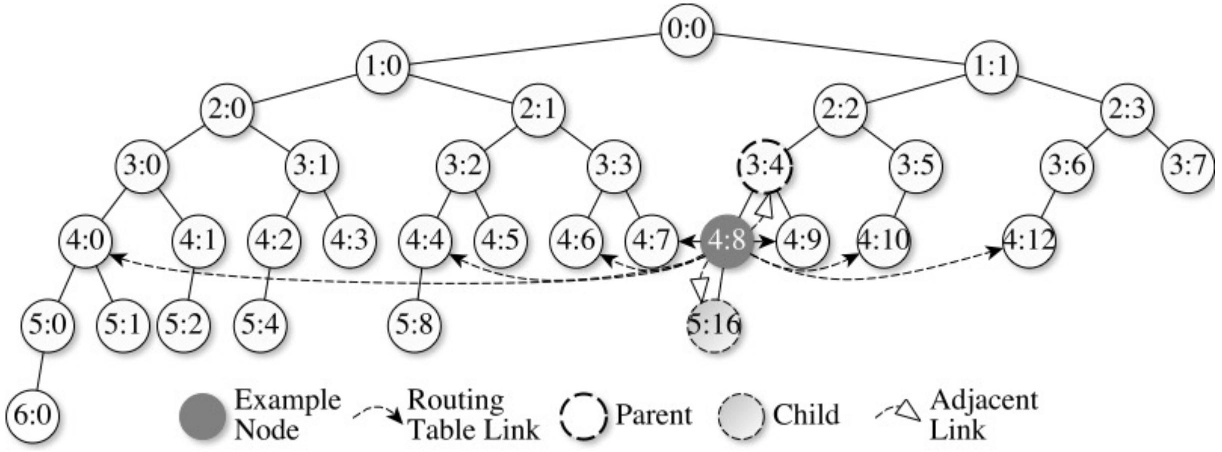

Figure 4: The BATON* tree has a height of 7 and 34 nodes in total

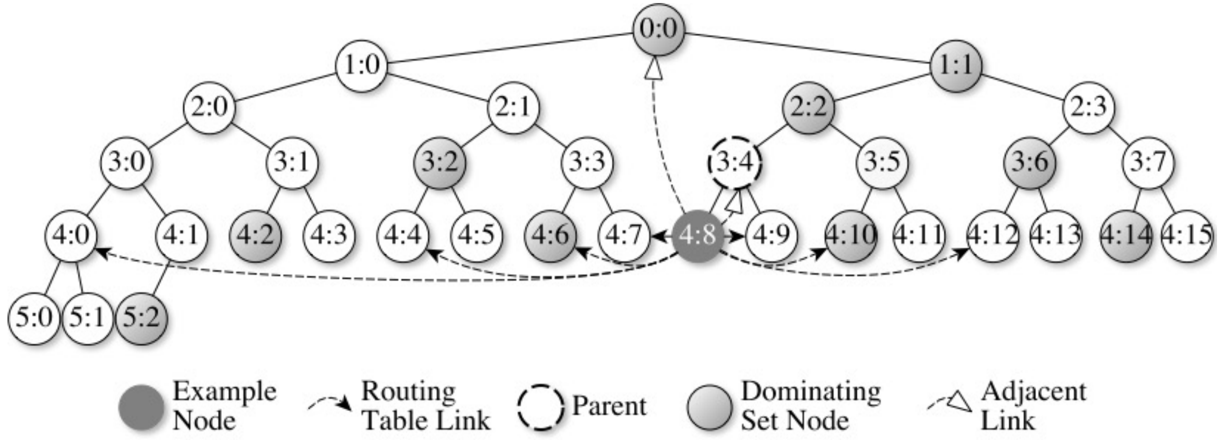

Figure 5: The nBATON* tree has a height of 6 and 34 nodes in total

[4] Cha, S.-H., 2012. On integer sequences derived from balanced k-ary trees, in: Proceedings of the 6th WSEAS International Conference on Computer Engineering and Applications, and Proceedings of the 2012 American Conference on Applied Mathematics, AMERICAN-MATH'12/CEA'12. World Scientific and Engineering Academy and Society (WSEAS), Stevens Point, Wisconsin, USA, pp. 377–381.

Fraunhofer IML

nBATON* - Adjustments to Join/Leave Operation

- Null-balanced constraints
  - Only nodes on the level below the highest level can accept a new node
  - Nodes from levels below the highest level need to find a successor from the highest level

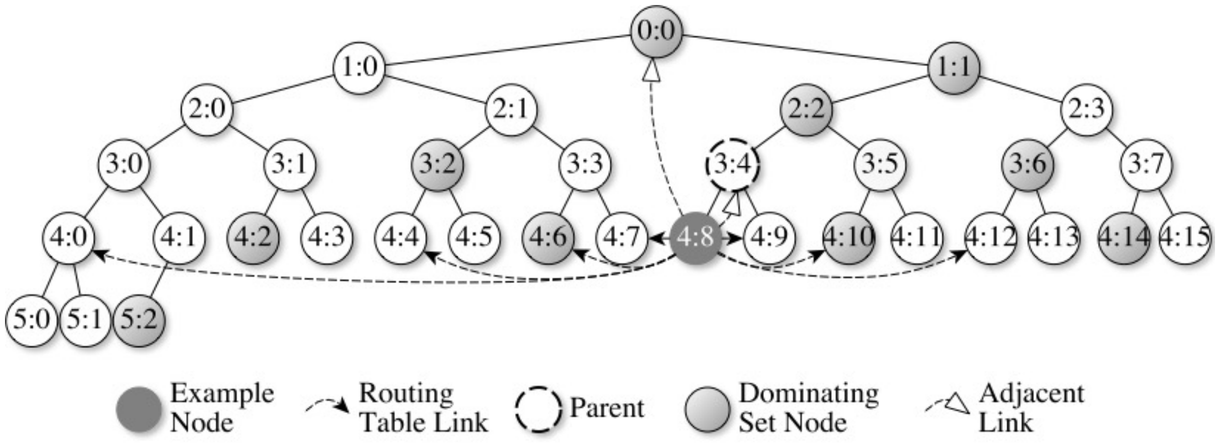

Figure 6: The nBATON* tree has a height of 6 and 34 nodes in total.

- Utilizing of Dominating Set Nodes [5] to find a free spot
  - Based on the construction of the RT

$$DSN_m(l) = \{i \cdot m \mid i = 2k+1, k \in \mathbb{N}_0, i \cdot m < m^l\}$$

$$DSN_m(0) = \{0\} \qquad DSN_m(1) = \{\lceil \frac{m}{2} \rceil\}$$

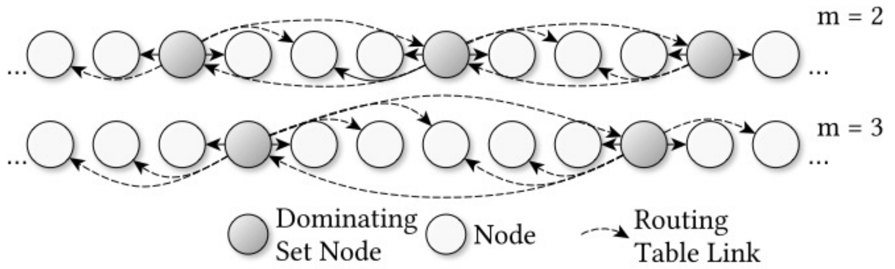

Figure 7: Example of routing neighbors (in-level) links for a dominating set node (DSN), for fanout m = 2 and m = 3

[5] Chunlin Yang and Xiuqi Li. 2005. 'Dominating-set-based' searching in peer-to-peer networks. *Int. J. High Perform. Comput. Netw.* 3, 4 (December 2005), 205–210. DOI:https://doi.org/10.1504/IJHPCN.2005.008562

Fraunhofer IML

nBATON* - Adjustments to Join/Leave Operation

- Utilizing the DS:
  - few nodes have an overview of the level
- Join/Leave Operation
  - Traverse the tree over the DSNs

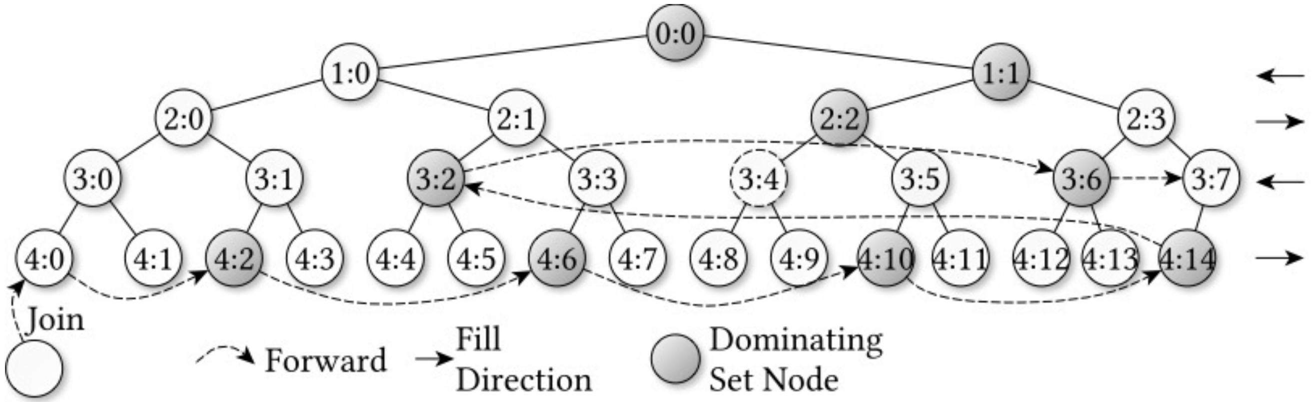

Figure 8: A node sends a join request to node (4:0). This join request gets forwarded through the dominating set nodes according to the swipe direction

Fraunhofer IML

Experimental Setup

- Large-scale networks using ns-3 [6]
  - C++ development
  - Network-size from 1,000 .. 10,000 nodes
  - fanouts m = {2, 4, 6, 8, 10}
  - 1,000 search exact queries
  - 10 times with 10 different seeds

- Dataset is publicly available on GitHub [7]

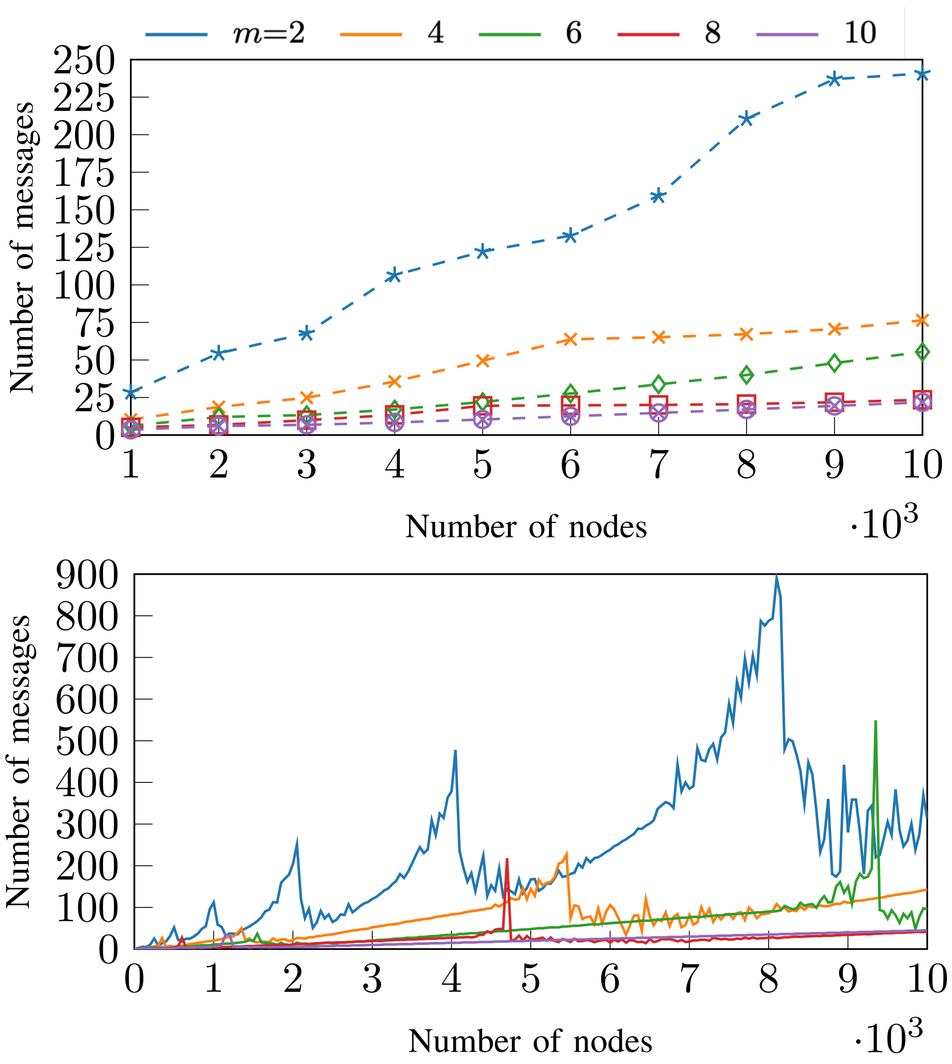

Figure 9: Averaged routing messages for a join (top) and the number of routing hops for a single node (bottom)

[6] Riley, G.F., Henderson, T.R., 2010. The ns-3 Network Simulator, in: Wehrle, K., Güneş, M., Gross, J. (Eds.), Modeling and Tools for Network Simulation. Springer, Berlin, Heidelberg, pp. 15–34. https://doi.org/10.1007/978-3-642-12331-3_2
[7] https://github.com/iml130/dataset-lcn2021 [8] https://github.com/iml130

Fraunhofer IML

Evaluation of the Cost for SearchExact Queries

- Impact of height on the search
- BATON* gets outperformed up to 50% by nBATON* in terms of routing hops

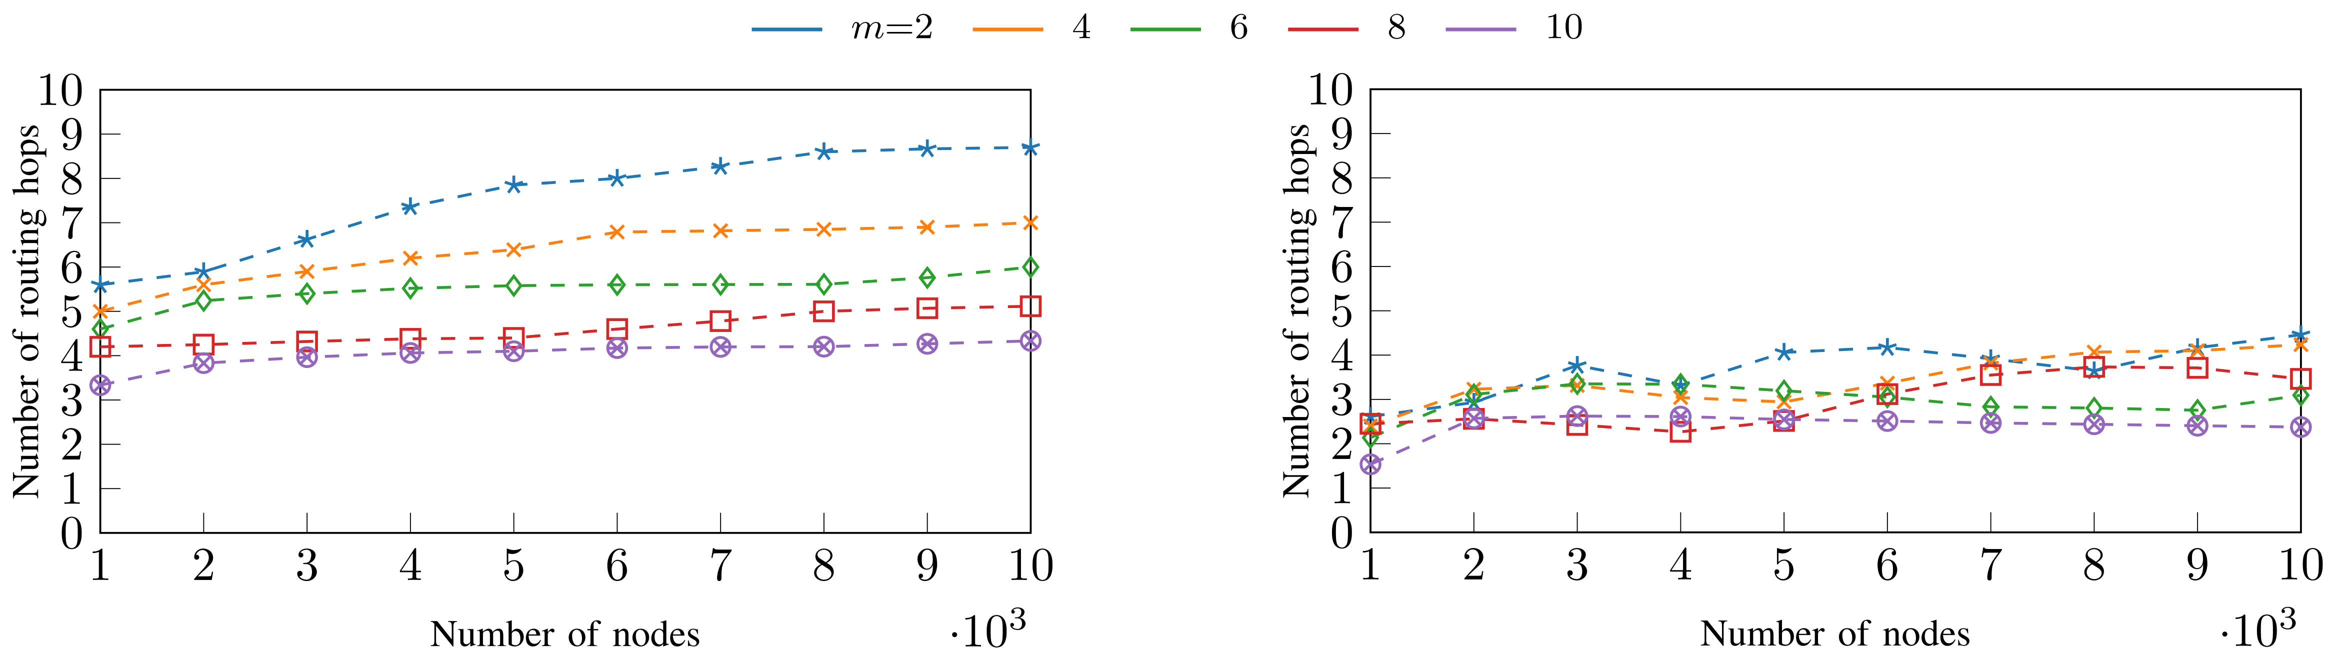

Figure 10: Costs for SearchExact queries: BATON* [6] (left) vs. nBATON* (right)

Fraunhofer IML

nBATON* Amortization – Is it worth it?

- Break-even point BATON vs. nBATON*
  - Assumption: search are much more frequent than updates (join/ leave)
  - Cost-Model based on probability for search and update operation

$$C = \alpha \cdot S + (1 - \alpha) \cdot U, \text{ with } 0 \leq \alpha \leq 1.$$

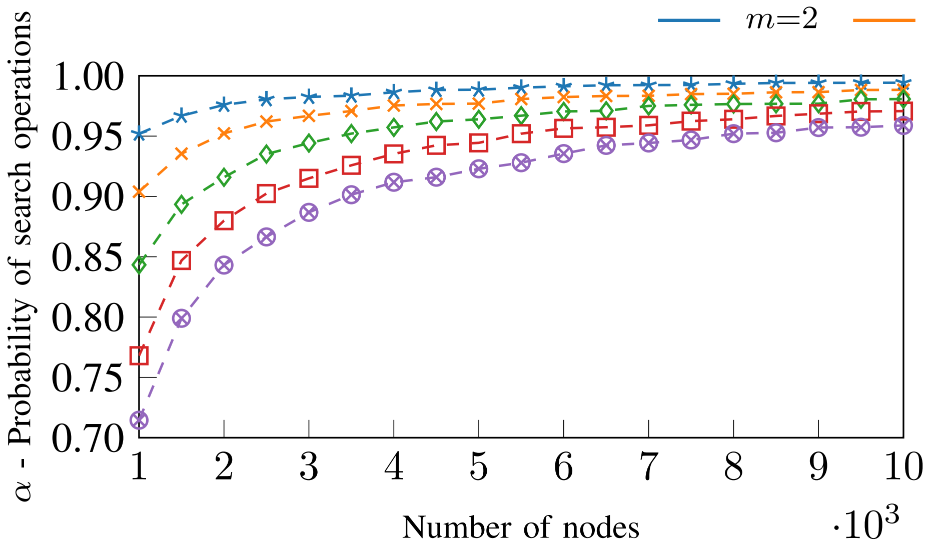

Figure 11: Required probability of search operations for a given initial network size

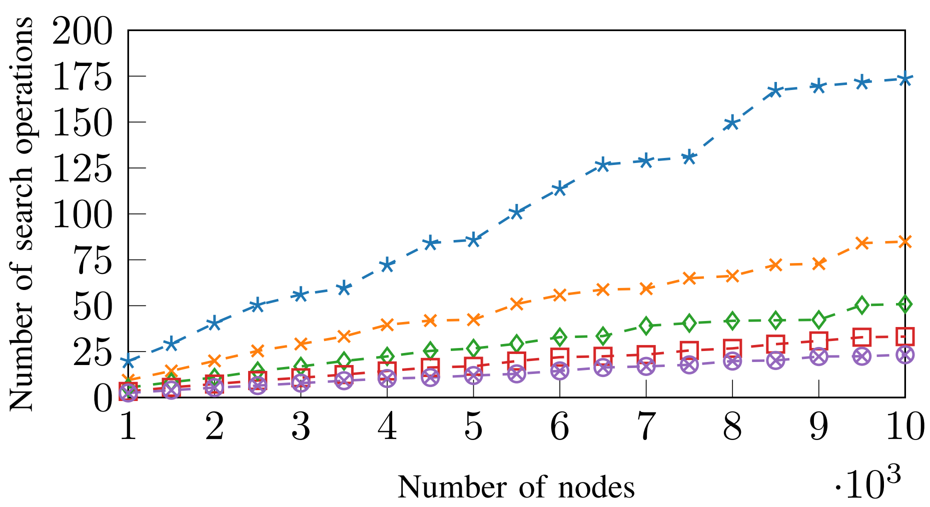

Figure 12: Comparison of nBATON* with BATON* based on the probability for search and update (join or leave) operations

Conclusion – We can achieve even more

- Null-Balanced m-ary Overlay Tree Network nBATON*
  - Minimal height
  - Utilizes Dominating Set Nodes
  - Outperforms BATON* in search queries by up to 50%

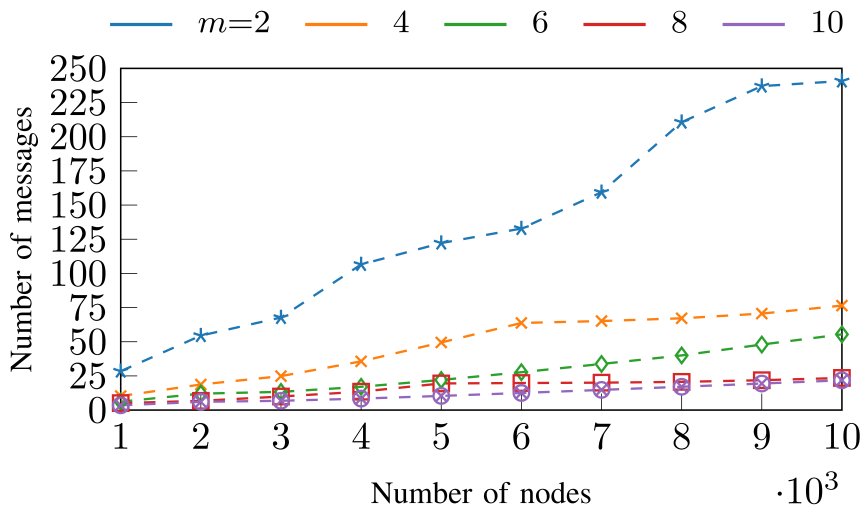

- Future Work:
  - Reducing the number of hops during join/leave
    - Minimizing/ reducing the Dominating Set
  - Making nBATON* publicly available