# Symphony

## Distributed Hashing In A Small World

Paper by Gurmeet Manku, Mayank Bawa, and Prabhakar Raghavan (2003)

Presentation by Andrew Berns

# What is Symphony?

- P2P Distributed Hashing Protocol
- O(1) links per node, $O(\log^2 n)$ routing
- Combines several ideas:
  - **DHT**
  - **Kleinberg's Small World model**
  - **Other improvments**

# Symphony DHT

- Ring topology
- Each node chooses an ID from [0,1)
- Each object is hashed to a m-bit key K, and is located at the node with lowest ID greater than $K / 2^m$

# Kleinberg's Small World and Symphony

- Like in Kleinberg's Small World, each node contains *k long distance links (k is not fixed)*
  - *Chosen using a harmonic distribution PDF:*

$$\frac{1}{x \ln n}$$

# Estimating system size

- For the PDF, knowledge of the system size is important
  - Can be done using a central server
  - Decentralized method is to use estimation protocol:

$$\frac{s}{X_s}$$

# Long Distance Links

- *Number of incoming links bounded by 2k*
- *If a chosen long distance link is over capacity, try to find another one a bounded number of times*

***Using these long distance links results in $O(1/k \log^2 n)$ routing!***

# What about dynamic networks?

- Joining a network:
  - Find an existing member
  - Select an ID
  - Locate two neighbors, using them to estimate the number of nodes in the system
  - Select long distance links

*Requires looking up k links*

$O(\log^2 n)$

# What about dynamic networks?

- Leaving a network:
  - Snap all long distance incoming and outgoing links
    - Others will find new outgoing links
  - Neighbors establish a link between themselves

*Requires updating k old incoming links*

$O(\log^2 n)$

# Symphony Protocol Review

- Ring topology with IDs in [0, 1)
- Links to immediate neighbors, as well as *k long distance links, chosen with a harmonic distribution*
- *$O(1/k \log^2 n)$ routing*
- *$O(\log^2 n)$ joining and leaving*

# Practical Considerations and Improvements

- Number of long distance links
- Bidirectional Routing
- Lookahead

# How many long distance links?

- More long distance links show diminishing returns

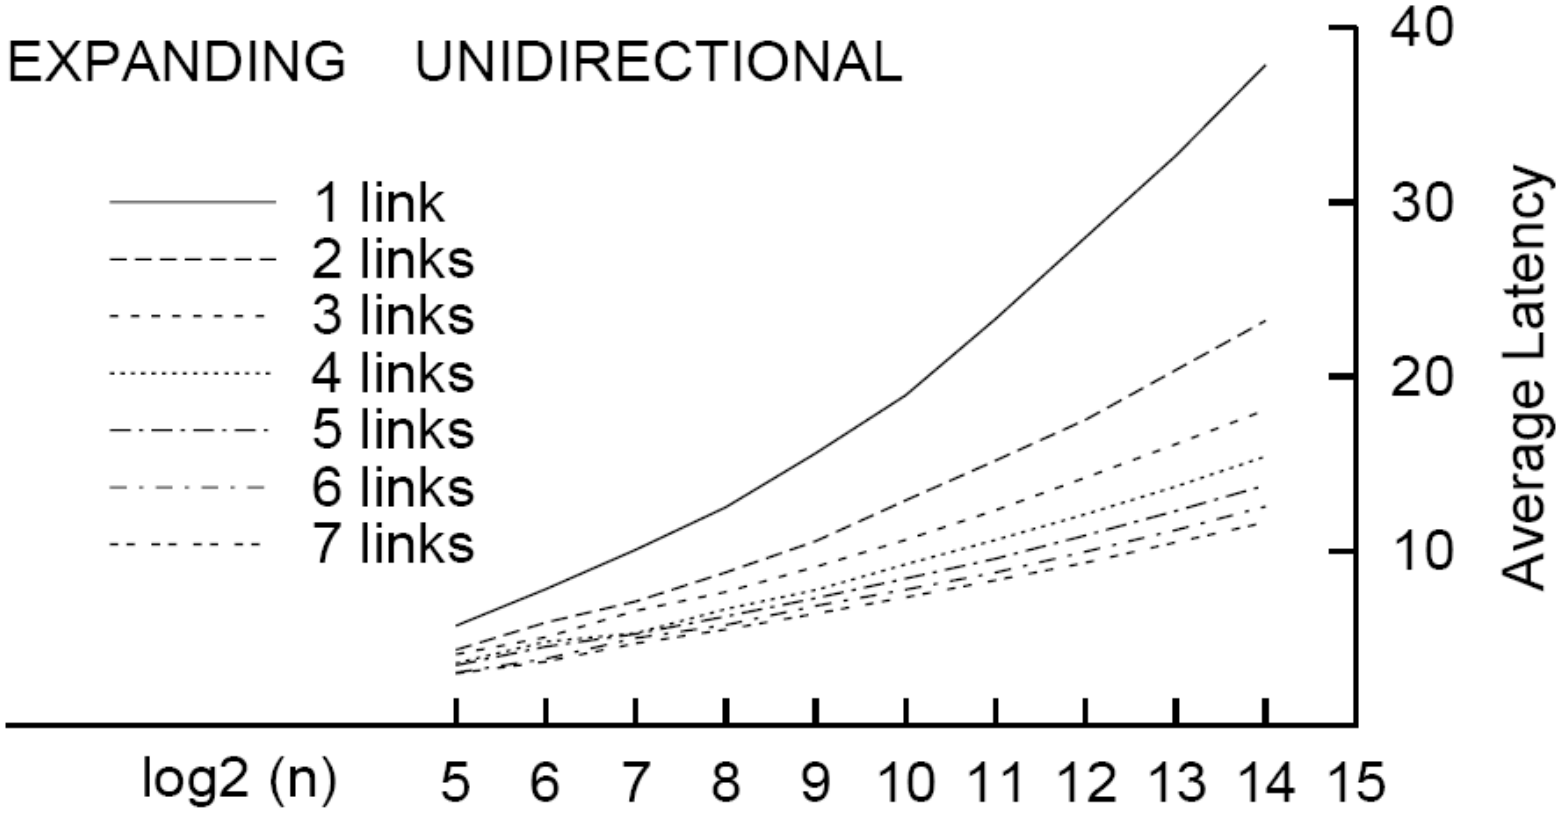

# Bidirectional Routing

- Bidirectional routing offers about a 30% improvement

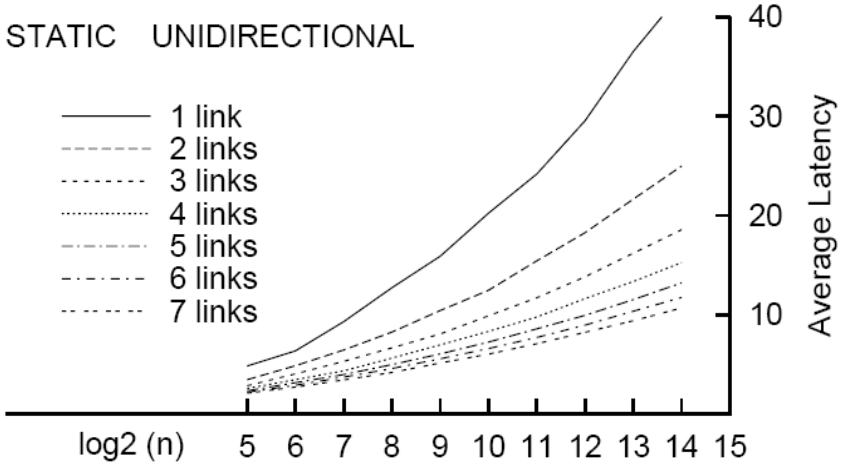


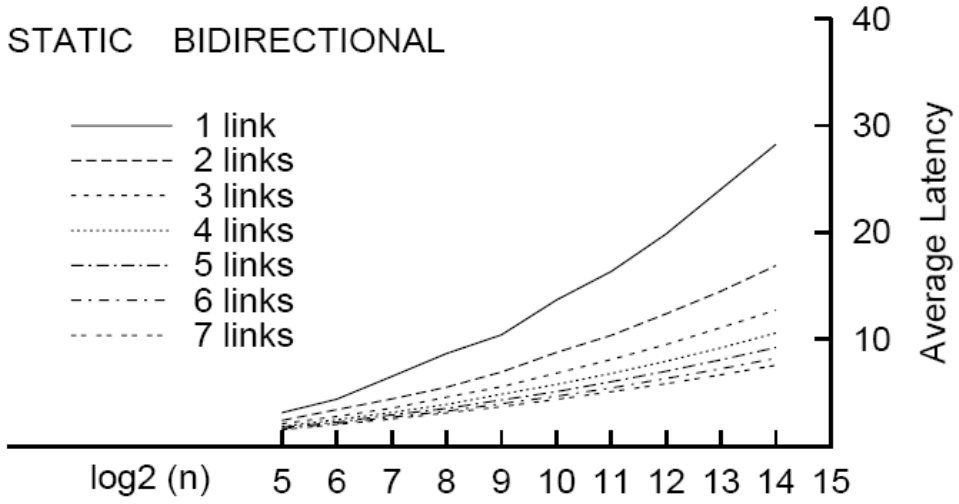

# Routing with Lookahead

- Using Lookahead (NoN) offers about a 40% improvement

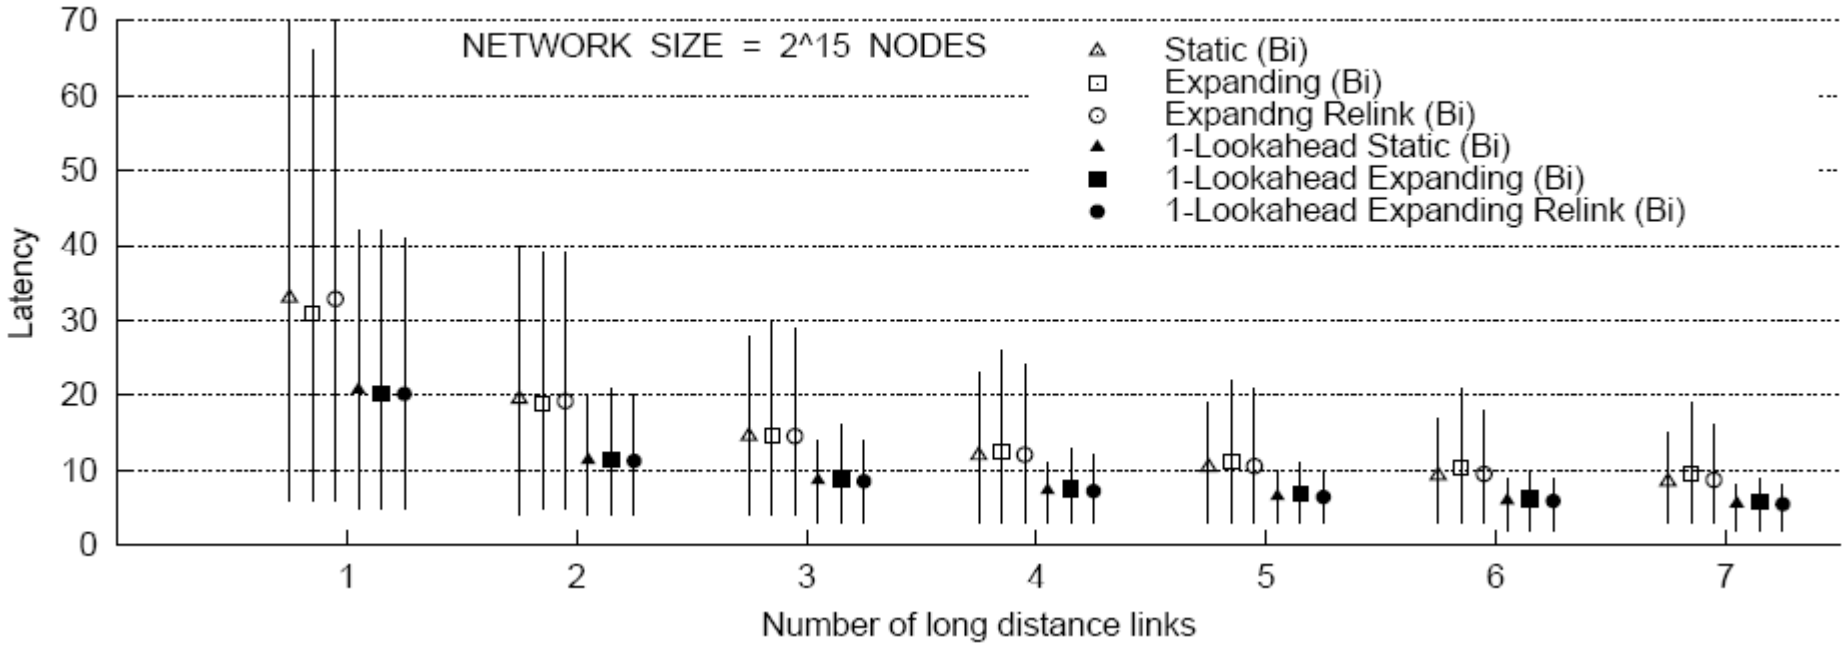

# Symphony's Benefits

- Low State Maintenance
- Fault Tolerance
- Degree vs. Latency Tradeoff
  - Support for Heterogeneous Nodes
  - Incremental Scalability
  - Flexibility

# Conclusion

- Symphony uses Kleinberg's Small World idea to create an efficient P2P network
  - *$O(\log^2 n)$ routing*
  - *$O(1)$ links*

- Future work includes implementing a system with heterogeneous nodes and factoring in network proximity

# References

- All figures taken from the Symphony paper:

G. S. Manku, M. Bawa, and P. Raghavan. Symphony: Distributed hashing in a small world. *Proc. 4th USENIX Symposium on Internet Technologies and Systems (USITS)* 2003.