

### Optimally Efficient Multicast in Structured Peer-to-Peer Networks

Dirk Bradler, TU Darmstadt



Jussi Kangasharju, University of Helsinki



Max Mühlhäuser, TU Darmstadt



# Agenda

**Introduction**

**DTC Construction**

**Evaluation**

**Outlook**

**Multicast is performed by a distributed tree construction (DTC) algorithm**

- **Tree uses only local information**
- **No peer receives any message twice**
- **Tree depth comparable to related multicast approaches**

**Many application domains:**

- **Streaming**
- **Prefix search**
- **Range queries**
- **Location based services**

# Agenda


**Introduction**

**DTC Construction**

**Evaluation**

**Outlook**

# DTC In Content Addressable Network




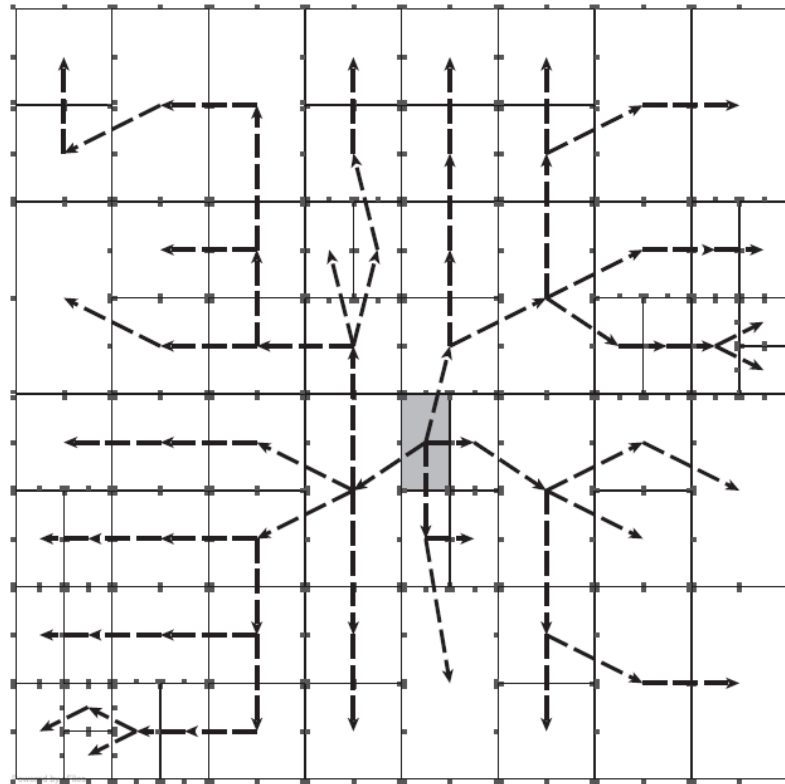
Origin O starts DTC and sends its own position with the multicast message.

## DTC Construction Rule

When a peer X receives the query, it computes for every one of its neighbors the vector from the center of the root's zone to the center of the neighbor's zone. If that vector intersects the common border surface between X and the neighbor, then X should add that neighbor as its child.

Assumption: Size of neighbor's zone is known locally

# DTC Tree



Subset of CAN overlay. Multicast is started from marked node and forms a tree.

# Agenda

**Introduction**

**DTC Construction**

**Evaluation**

**Outlook**

# Average number of messages received per node

Messages Received	DTC CAN	Simple Flooding	ALM
1	2000	1	1224
2-3	0	2	684
4-5	0	6	78
6-7	0	22	12
8-9	0	67	2
10-11	0	247	0
12-13	0	925	0
>14	0	Sum 7851	0
<b>Sum</b>	<b>2000</b>	<b>26106</b>	<b>3168</b>


CAN network:

- 2000 nodes
- 10 dimensions
- Averaged over 30 simulations





# ALM In Content Addressable Network



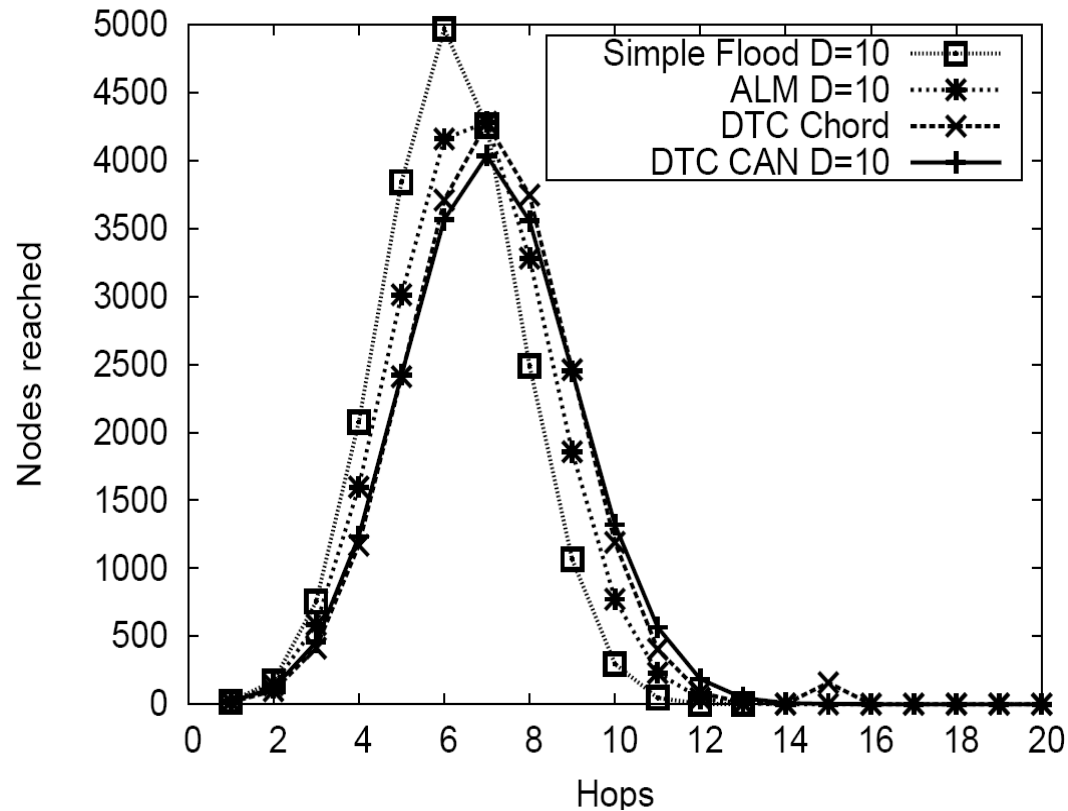
Multicast approach by  
Ratnasami:

Start from an arbitrary  
node.

Keep sending north and  
south until end of ID  
space.

Then send to your  
left/right neighbour until  
end of ID space

# Query Depth

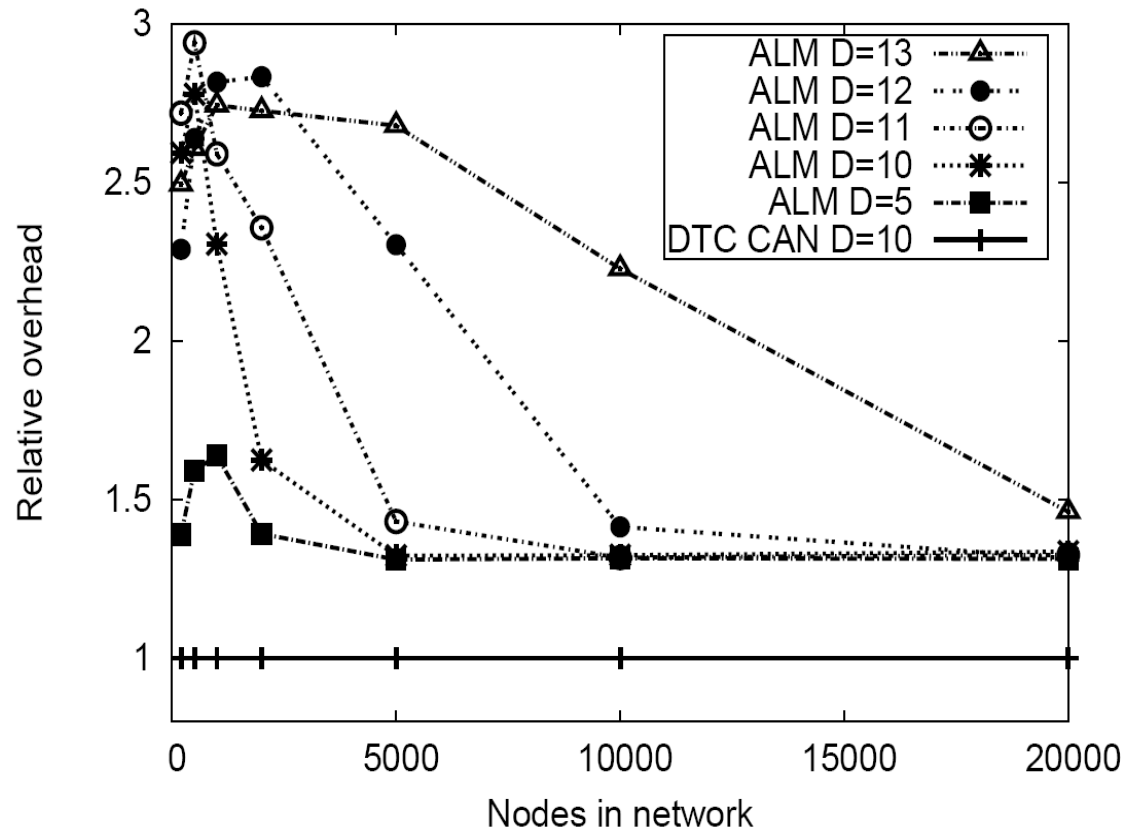


## Network properties:

- 20000 nodes
- 10 dimensions
- Averaged over 30 simulations

Simple Flood shows optimal multicast depth. ALM and DTC perform similar.


# Average overhead of ALM compared to DTC



ALM overhead varies between 1.4 and 3 times the optimal number of needed messages. On large numbers of peers the overhead tends to converge to 1.4.

Overhead tends to stay constant for a high number of peers.

# Malicious Peers



DTC is sensitive to attacks.

In case of 20% malicious peers only 50% of the network are reachable with a multicast message.



# Agenda

**Introduction**

**DTC Construction**

**Evaluation**

**Outlook**

# Outlook


Range queries and Prefix search are in an early testing phase.

Technical report: <http://arxiv.org/abs/0808.1207>

Basic idea for prefix search:

- Replace hash function with a quad tree
- Perform DTC in a defined subset of the network
- Query results must be located in the subset of the network

Showstopper: Bottlenecks may occur



---

# Questions

---



?

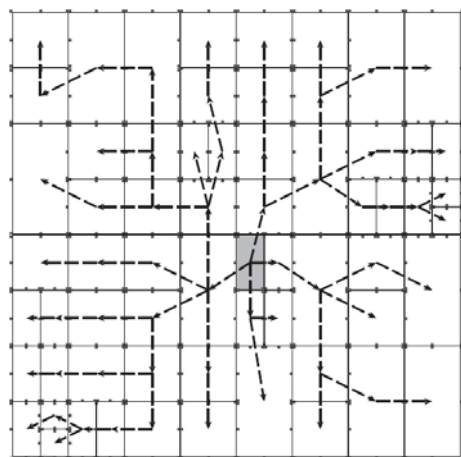
?

?


**Thank you**

# Optimally Efficient Multicast and Prefix Search in Structured Peer-to-Peer Networks

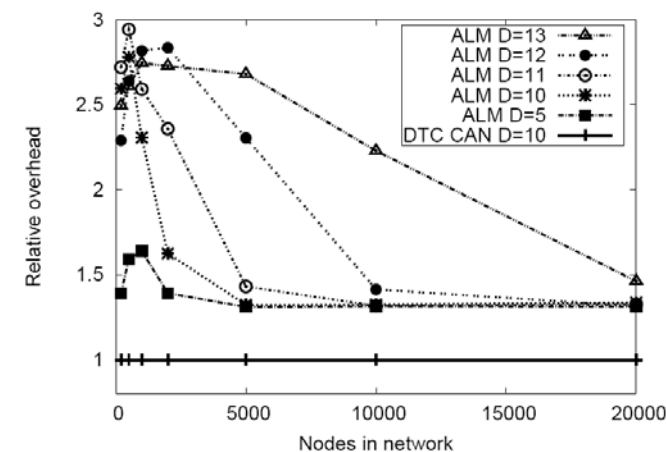
- Application Level Multicast without any duplicates → Minimal amount of messages
- Enables Prefix search in DHTs (tested with Chord and CAN)



Subset of CAN overlay.  
Multicast is started from marked node and forms a tree.



Peers calculate virtual line between origin (O) and their neighbors.  
If an intersection between the peers own zone and its neighbors zone exists the multicast message is forwarded.


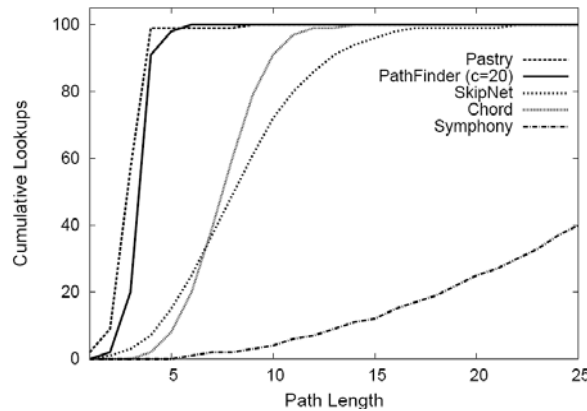
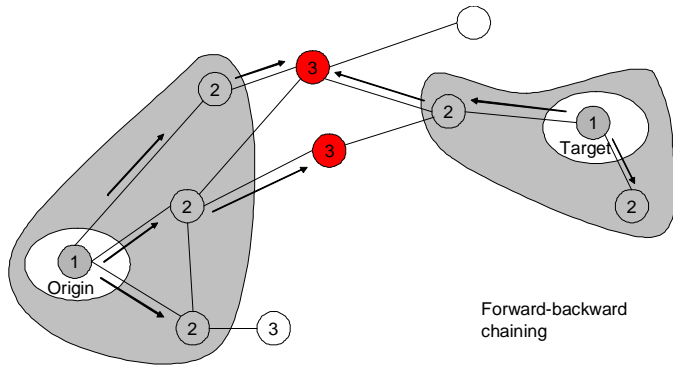


Related solutions tend to have a message overhead of >30%.  
More dimensions cause higher overhead.



# PathFinder: Efficient Lookups and Efficient Search in Peer-to-Peer Networks

- PathFinder provides low latency DHT lookups with fix number of neighbors.
- Desired random graph properties are preserved.



Idea: Precalculate a path in an Erdős Rényi random graph. It is build using a global known construction rule. Forward-backward chaining finds a match in <30ms for 1mio. nodes.

Cumulative lookups of PathFinder compared to established DHT overlays with 20k nodes. PathFinder latency comparable to Pastry at 20k nodes but number of neighbors < 1/3.

Recovery time for massive failures (10%-50%) with 10k peers.

DHT is established in less than 60 seconds up to 50% node failures