



A Lightweight Approach to Network Positioning

Kalavri Vasiliki
EMDC - DS course
2 Mar 2011

Outline

- Introduction
- Virtual Coordinates vs. Meridian
- The Meridian Framework
- Meridian's approach to Positioning Problems
- Simulation Results
- Conclusions

Network Positioning

- Not a problem in small systems
 - We can measure latency
 - We can connect to the closest server
- In large-scale distributed systems:
 - Collecting global information is infeasible
 - Lack of centralized servers

Applications

- File Sharing
- Content Distribution Networks
- Backup Systems
- Communication Networks
- Pub-Sub
- Multi-player Online Games

Outline

- Introduction
- **Virtual Coordinates vs. Meridian**
- The Meridian Framework
- Meridian's approach to Positioning Problems
- Simulation Results
- Conclusions

Network Embedding

- Estimate network distances
- Model the Internet as a geometric space
- Map high-dimensional measurements into a smaller space
- Use of classical distance function and relative error minimization

Drawbacks

- Accuracy
- Parameter-dependent errors
- Coordinates change over time
- Complexity
- Bandwidth Overhead

Meridian Approach (1)

- P2P Overlay
- Design Goals
 - Accurate
 - Scalable
 - Lightweight
- Uses direct measurements


Meridian Approach (2)

- Each node keeps a fixed number of peers
- Peers are organized in concentric rings
- No attempt for a global view
- Suitable for on-demand queries

Outline

- Introduction
- Virtual Coordinates vs. Meridian
- **The Meridian Framework**
 - **Multi-Resolution rings**
 - **Ring Membership Management**
 - **The Gossip Protocol**
- Meridian's approach to Positioning Problems
- Simulation Results
- Conclusions

Multi-Resolution Rings (1)



- Concentric
- Non-overlapping
- Finite number of rings
- Upper limit of k nodes

Multi-Resolution Rings (2)

- Why rings?

Multi-Resolution Rings (2)

- Why rings?
 - The ring structure favors nearby neighbors
 - Each node needs pointers to remote regions

Multi-Resolution Rings (2)

- Why rings?
 - The ring structure favors nearby neighbors
 - Each node needs pointers to remote regions
- Why exponentially increasing radius?

Multi-Resolution Rings (2)

- Why rings?
 - The ring structure favors nearby neighbors
 - Each node needs pointers to remote regions
- Why exponentially increasing radius?
 - Nearby peers are more useful than faraway
 - Manageably small number of rings per node

Ring Membership (1)

- k represents a trade-off: accuracy vs. overhead
- Performance depends on node choice
- Need to promote geographic diversity
- Periodically replace ring members with better alternatives

Ring Membership (2)

- k primary ring members and l secondary organized in a FIFO pool
- Definition of a local coordinate space
- Process runs in the background in linear time
- In case of detection of node failure:
 - The node is replaced by a random secondary
 - Temporal quality degradation

Gossip Protocol - Goals

- Loosely-connected
- Robust
- Up-to-date
- Provide a sufficient diverse set of peers

Gossip Protocol - Steps

- For every ring
 - A sends a gossip packet to random B containing a random node from each ring
- B measures its latency to A and the nodes of the gossip packet by sending probes
- Newly discovered nodes become B's secondaries

Gossip Protocol - Join

- A new node contacts one Meridian node
- Retrieves the entire list of peers
- Places the peers in its own rings
- Participates in the gossip protocol

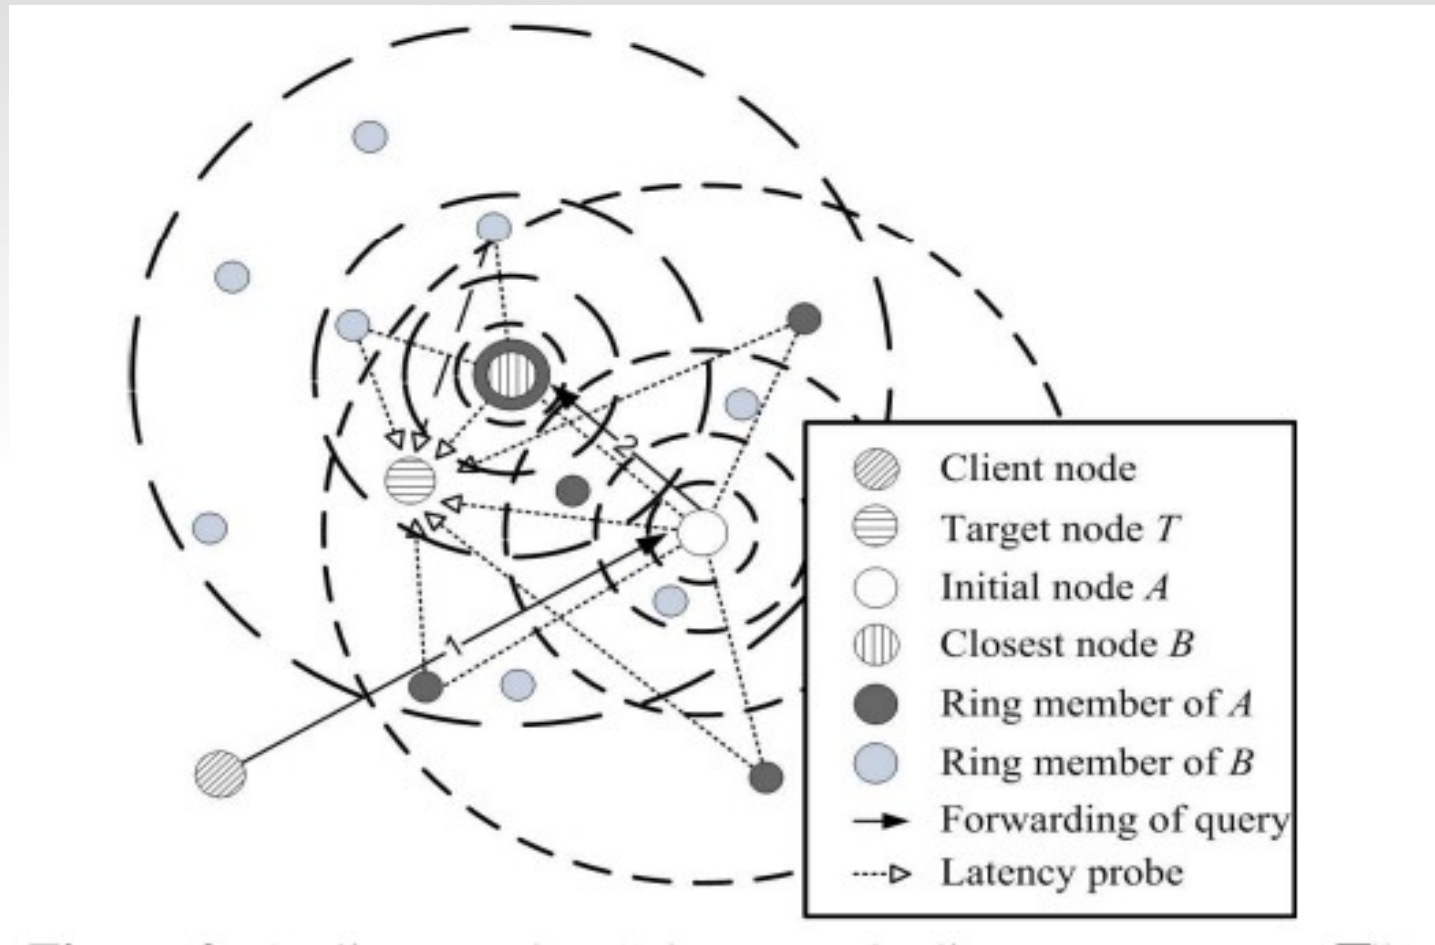
Outline

- Introduction
- Virtual Coordinates vs. Meridian
- The Meridian Framework
- **Meridian's approach to Positioning Problems**
 - **Closest Node Discovery**
 - **Central Leader Election**
 - **Multi-Constraint System**
- Simulation Results
- Conclusions

Closest Node Discovery


- Multi-hop search
- Similar to structured P2P networks
- Physical latencies instead of numerical
- Target nodes need not be part of the overlay

Closest Node Discovery - Steps




Central Leader Election

- Extends the Closest Node Discovery protocol
- Selection based on average latency



Multi-Constraint System

- Node selection where constraints define the boundaries of the solution region




Outline

- Introduction
- Virtual Coordinates vs. Meridian
- The Meridian Framework
- Meridian's approach to Positioning Problems
- **Simulation Results**
- Conclusions


Simulation

- 2000 Meridian nodes
- 500 target nodes
- $k = 16$ nodes per ring
- 9 rings per node
- 8 days-experiment


Scalability




Accuracy - Multi-Constraint



Comparison - Closest Node




Central Leader Election




Anything unexpected?

Central Leader Election




Error drops as group size increases!

Closest Node Discovery



What happens with more nodes per ring?

Closest Node Discovery



Latency decreases with more nodes per ring!

Outline

- Introduction
- Virtual Coordinates vs. Meridian
- The Meridian Framework
- Meridian's approach to Positioning Problems
- Simulation Results
- **Conclusions**

Conclusions (1)

- 3 design goals
 - lightweight
 - scalable
 - accurate
- Uses direct measurements
- On-demand queries - Not a global service

Conclusions (2)

- Meridian Applications:
 - Content Distribution Networks
 - Leader Election
 - SLA fulfillment
 - Computational Grids



A Lightweight Approach to Network Positioning

Kalavri Vasiliki
EMDC - DS course
2 Mar 2011