Stable Node Positioning

Michal Szymaniak

michal@cs.vu.nl

Context

- Large-scale distributed system
  - Globule CDN
- Latency-driven content placement and client redirection
  - Goal: optimization of client access latencies
- Geometric model of latencies
  - GNP == Global Network Positioning
  - Internet == N-dimensional geometric space
  - Nodes == Positions, Latencies ~ Distances

2D Snapshot

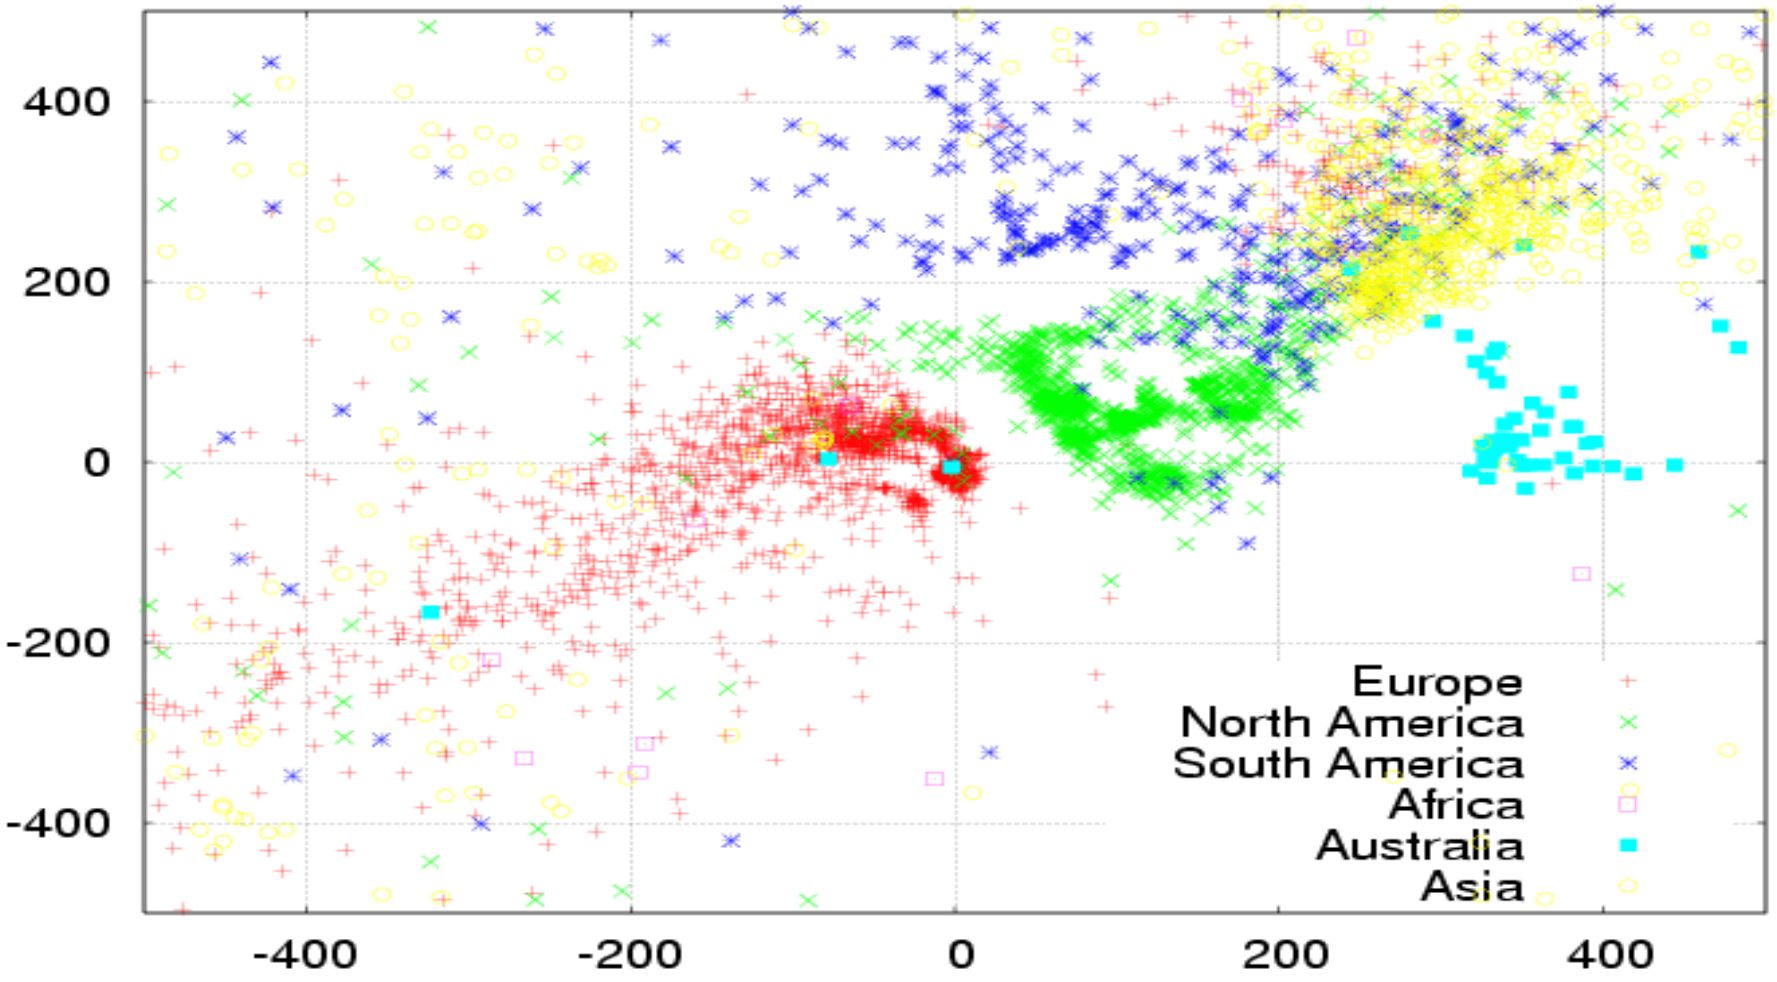

- How stable are these node positions over time?

Dynamics of Positions

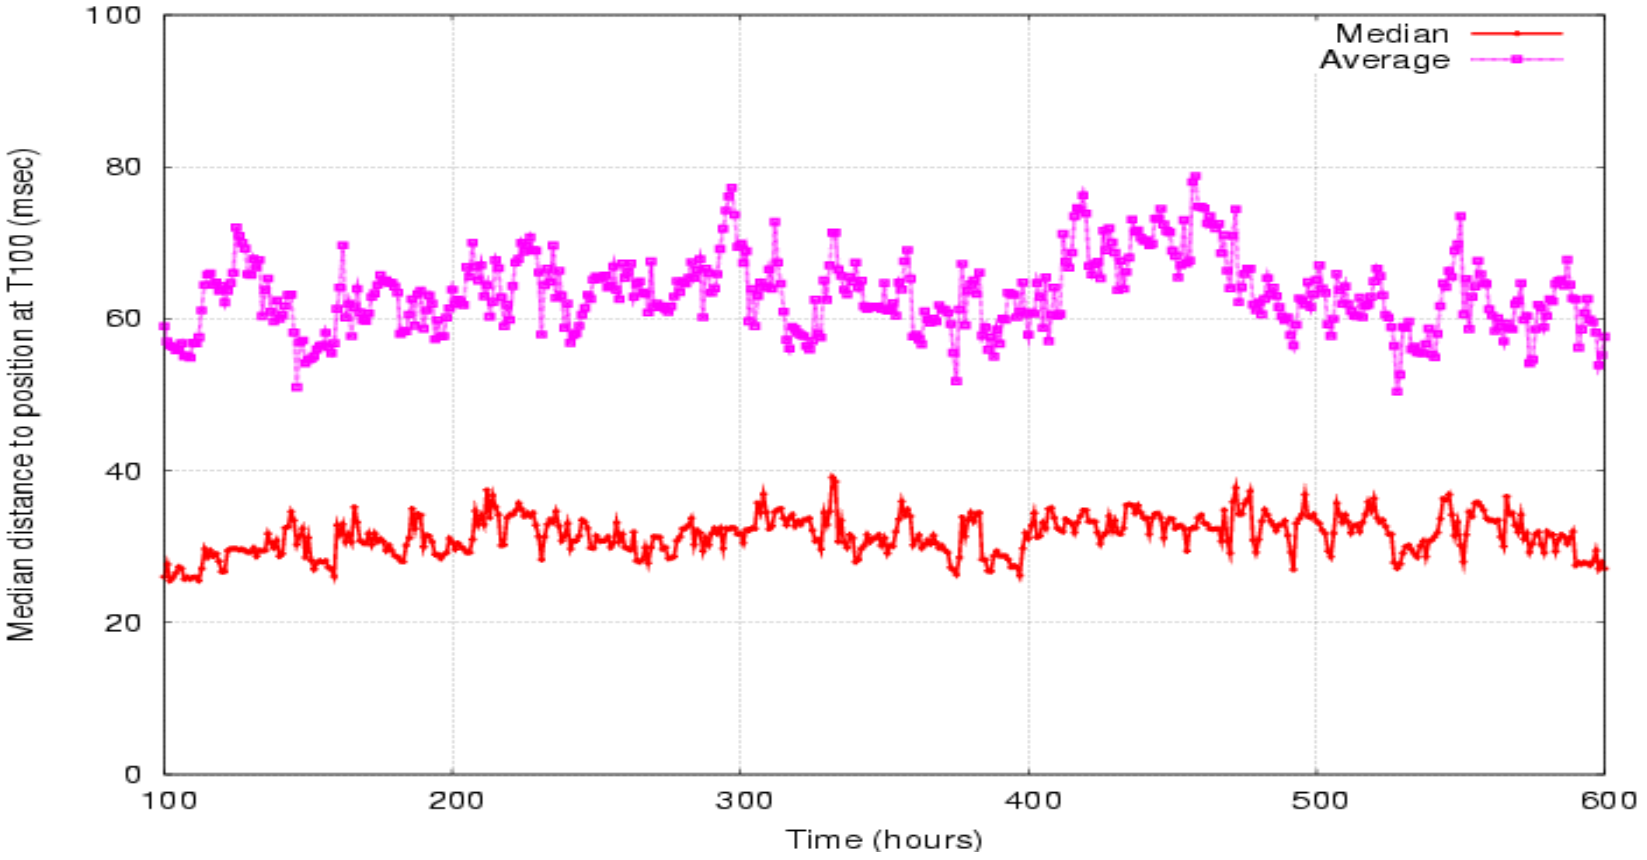

- Positions computed 1 hour later are often completely different
- Is it possible to obtain more stable positions?

Agenda

- Positioning Algorithm
- Latency Analysis
- Latency Stabilization
- Latency-Estimation Accuracy
- Case Study: ADSL Client
- Conclusion

Positioning Algorithm

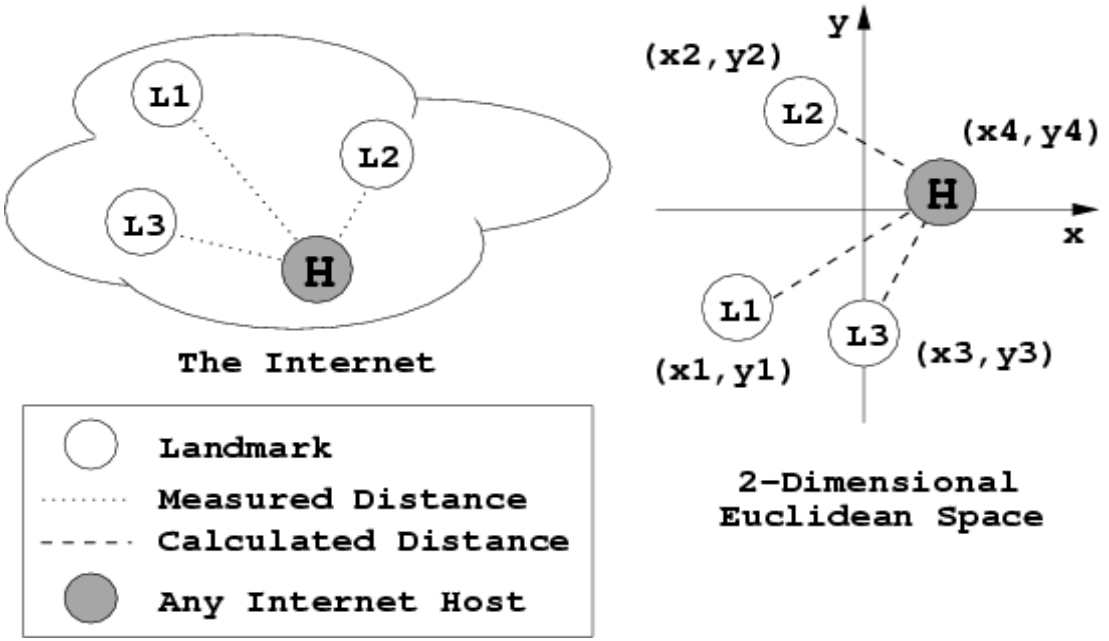

- Positions computed relative to so-called Landmarks
  - Landmark positions are already known
  - Distances to landmarks ~ Latencies to landmarks
- Repositioning caused by changes in node-to-landmark latencies

Latency Analysis

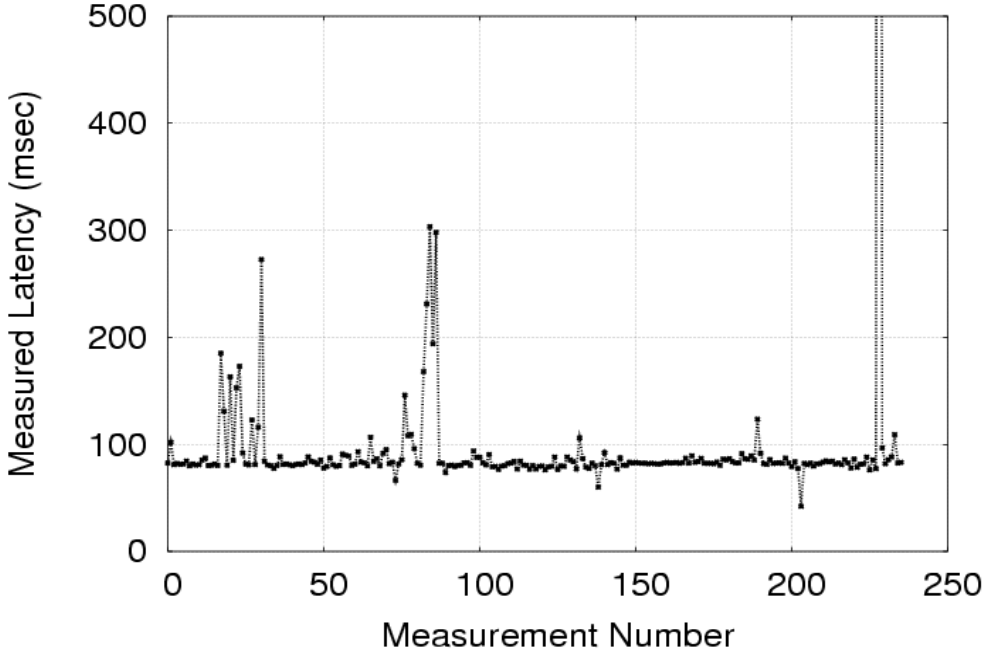

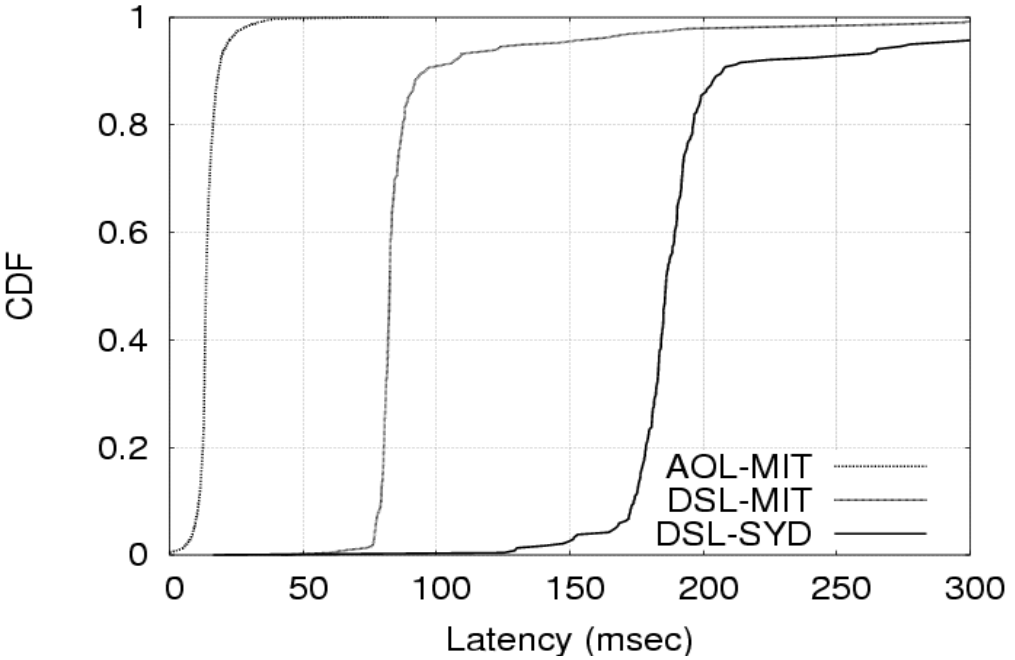

- Latency traces usually contain certain dominant values
  - Higher-than-dominant values – network congestion (noise)
  - Lower-than-dominant values – overloaded landmark (noise)
- Dominant values ~ empty-wire latencies; what if we use only them?

Latency Stabilization

- Wire-latency discovery:
  - Take median of the N most-recent measurements (N tunable)
- Result:

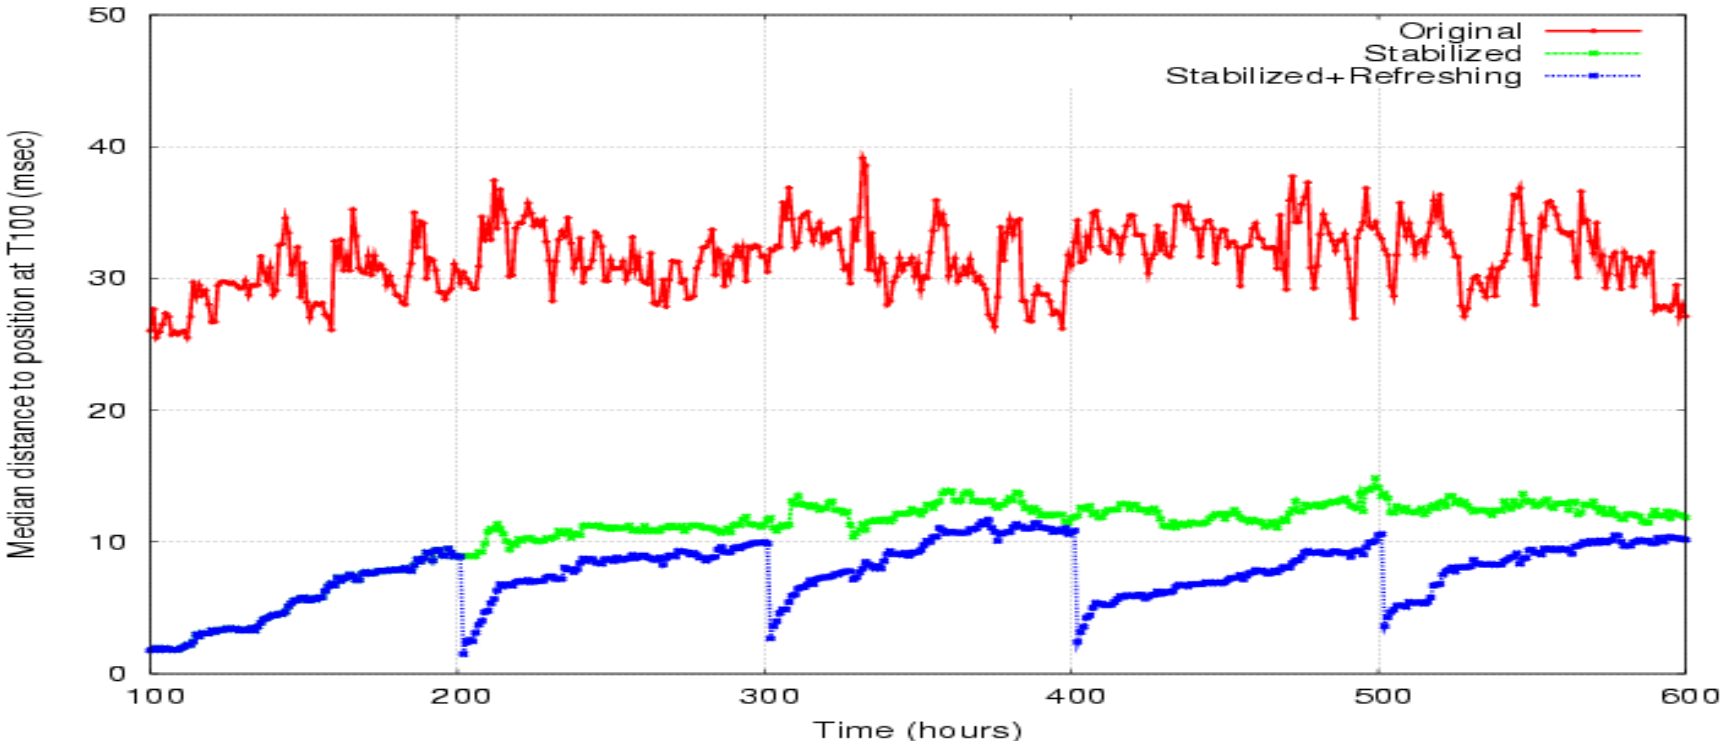

Latency-Estimation Accuracy

- Using wire latencies slightly improves estimation accuracy
  - Intuition: we eliminate accidental repositioning caused by noise

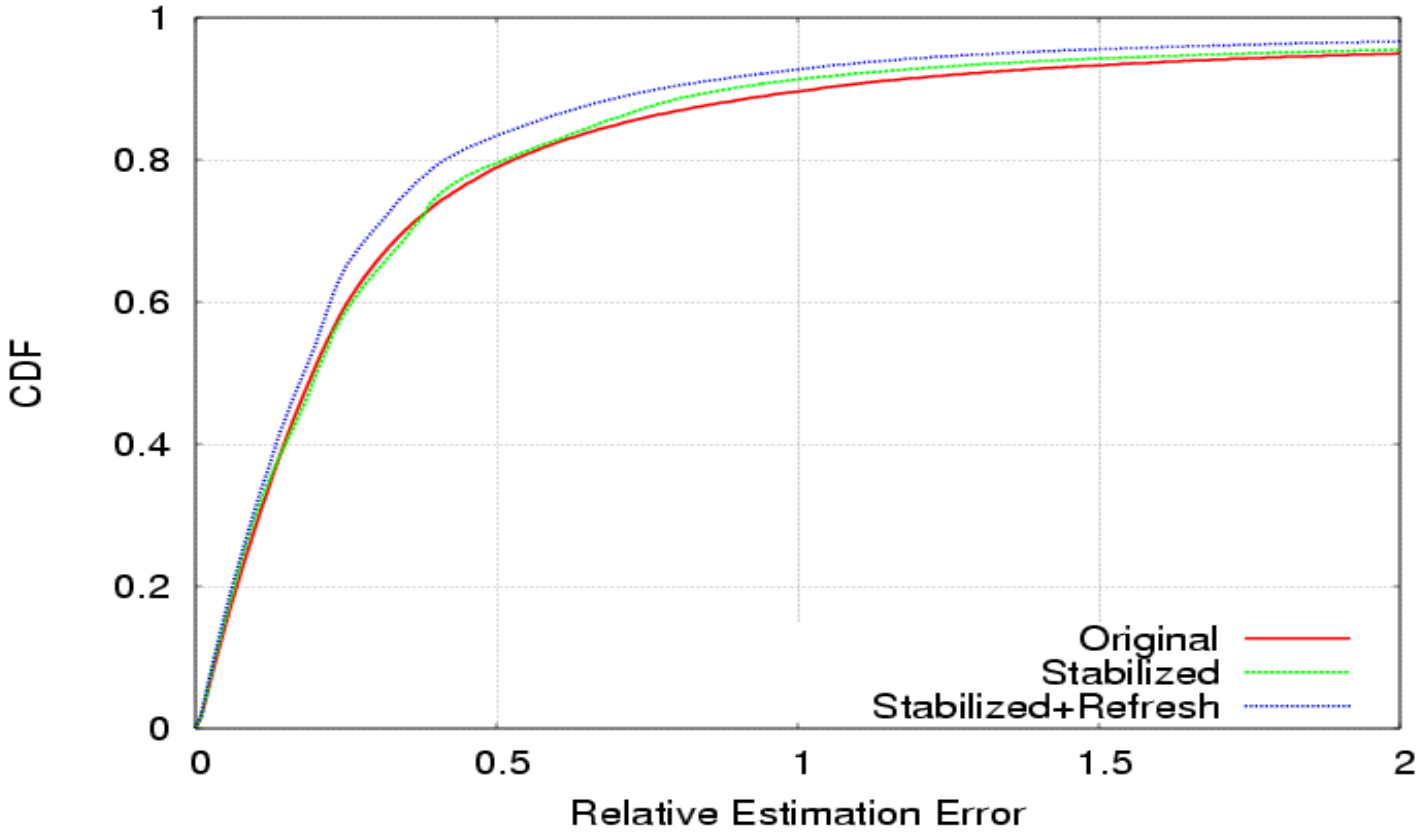

Case Study: ADSL Client

- Benefits visible even better for a single client:

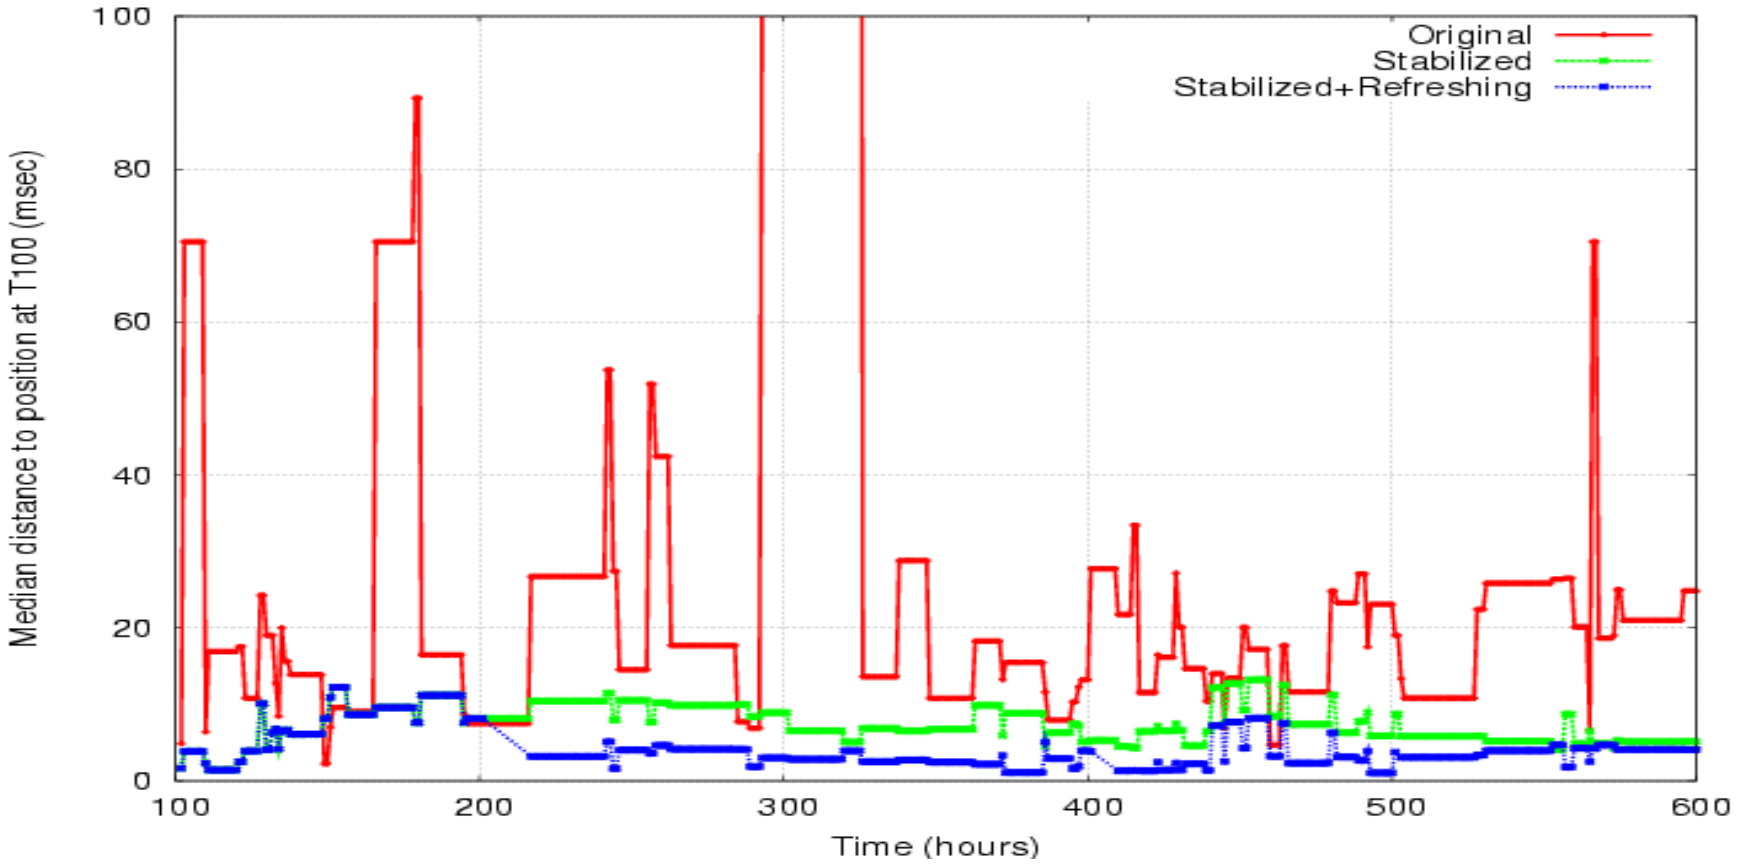

Conclusion

- Node positions calculated by GNP can be stabilized
  - Stable positions calculated based on empty-wire latencies
  - Empty-wire latencies ~ medians over N most-recent latencies
- Benefits:
  - Reduced number of positioning operations
  - Reliable identification of node clusters
  - Side effect: higher accuracy of latency estimation