On the Effects of Pre-computed Route Caching in Multiple Peer Group ATM-PNNI Networks

by

Kamalesh S Kalarickal

I T T C

Outline of the Presentation

- Introduction
- Pre-computed Routing Strategy
- Mechanics of Experimentation
- Preliminary Results
- Pre-computation for Multiple Peer Groups
- Route Cache Update Heuristics
- Conclusion

Overview of PNNI Routing

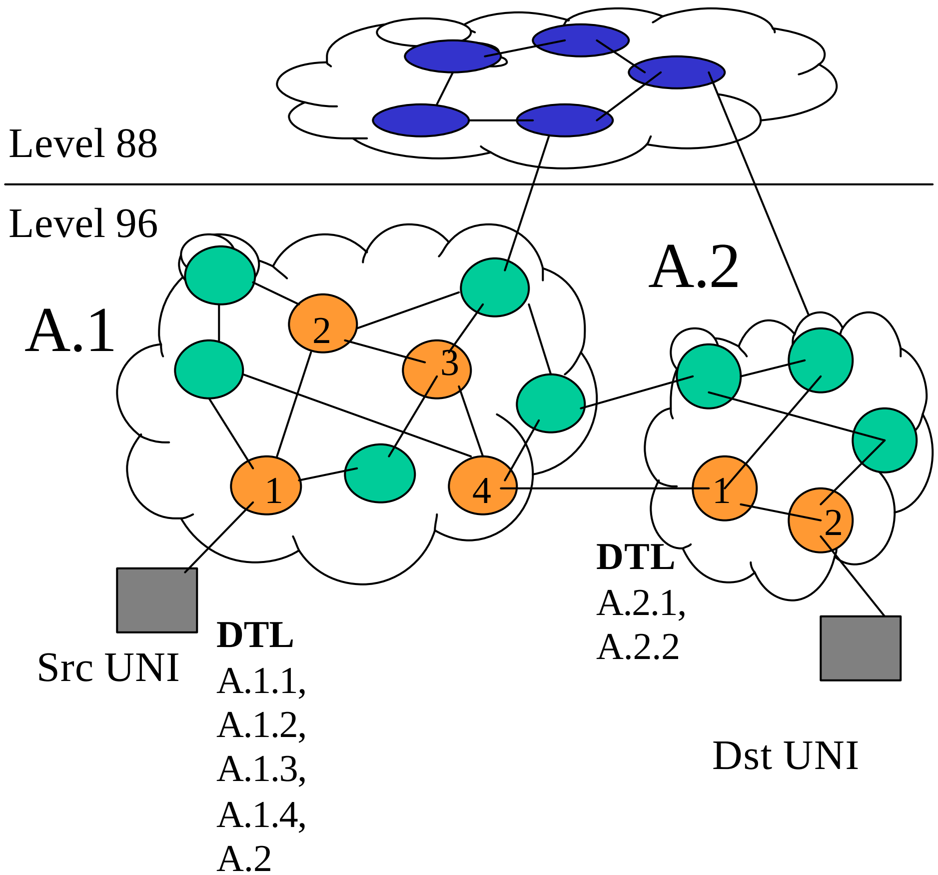

- PNNI *hierarchical* addressing.
- Topology *database* of network resources.
- Link state *flooding* via PTSEs.
- PNNI uses *source routing.*
- Route re-computation at border node per peer group.
- *Crankback* on failures.

Stages in On-Demand Routing

1. Representation of network as a graph.
2. Generic Call Admission Control to prune the graph based on QoS requirements of the call.
3. Dijkstra’s Single Shortest Path Algorithm, minimizing a selected routing constraint.
4. Designated Transit List (DTL) is created, containing route to the destination.

Problem Statement

- On-Demand Routing in ATM PNNI networks involves expensive computation.
- This is carried out for *every* incoming call.
- It is also redone at the ingress node of every other peer group in the path.
- This results in increased call setup time.

Proposed Solution

- Pre-Compute routes to destinations in advance and store in a route cache.
- When a call arrives at the source node, the cache is searched to find an appropriate route.
- If a route is found, it is used.
- Otherwise on-demand routing is carried out and the route is inserted into the route cache.

Pre-Computed Routing

- Reduction in time for route discovery as compared to on-demand routing, results in *lower setup time*.
- But this occurs at the expense of the quality of routes due to aging of route cache entries.
- Hence update the route cache to reflect changed network state.

Comparison

On-Demand Routing.

Advantages:

- More up to date information for routing decision. Hence route failures are lower.

Disadvantage:

- Increased average call setup time due to overhead of route computation for each call.

Pre-Computed Routing.

Advantage:

- Lower average routing time leads to lower average call setup time.

Disadvantage:

- Decreased call acceptance, due to aging of information used to pre-compute routes.

Pre-computed Routing Strategy

Implementation and preliminary results

Stages in Pre-Computed Routing

1. Representation of network as a graph.
2. Max-Min B/W range is divided into equivalence classes of bandwidths.
3. For each quantized QoS level, do GCAC.
4. For every unique node in the database, run Dijkstra's Single Shortest Path Algorithm, and compute route.
5. DTL, if obtained, is stored along with the *effective bandwidth* of the route, in the cache.

Route Cache Policy

- Routes are pre-computed and stored in the cache during initialization and also during a route cache update.
- Pre-computed route processing is assumed to occur on a dedicated co-processor.
- While the route cache is being updated, it is not available for pre-computed route lookup.

Route Cache Policy ..(cont'd)

- Selecting from Route Cache.
  - First, a check is made for the destination node address.
  - If the set of DTLs is not empty, it is searched for a route that will satisfy the QoS requirements of the call.
  - If more than one such route exists, the route that leaves maximum residual bandwidth (load balancing) is chosen.

Route Cache Policy ..(cont'd)

- Invalidation of Cache Entries.
  - Can be set to occur when there is a Crankback event indicating a bad routing decision.
- Replacement of Cache entries.
  - Occurs when an invalidated entry is filled by a subsequent on-demand route computation.
  - Occurs when the cache is entirely rewritten during an update.

Mechanics of Experimentation

Metrics, Traffic, Evaluation, Analysis

I T T C

Metrics & Parameters

- Average Bandwidth Acceptance Ratio.
- Average Call Setup Time.
- Location of Failures.
- Network Topology.
- Peer group size.
- Traffic Characteristics.
- Cache Update Policy

Flat Network Topology

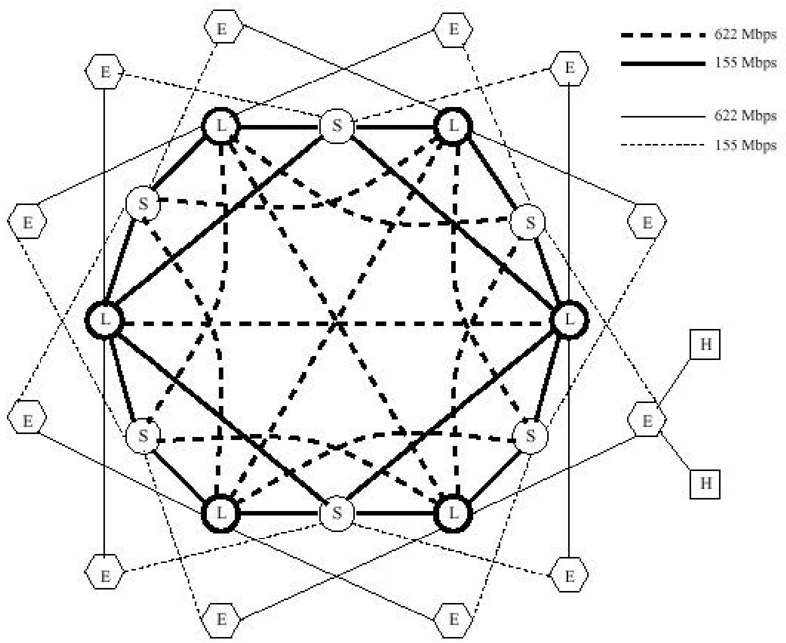

Hierarchical Network Topology

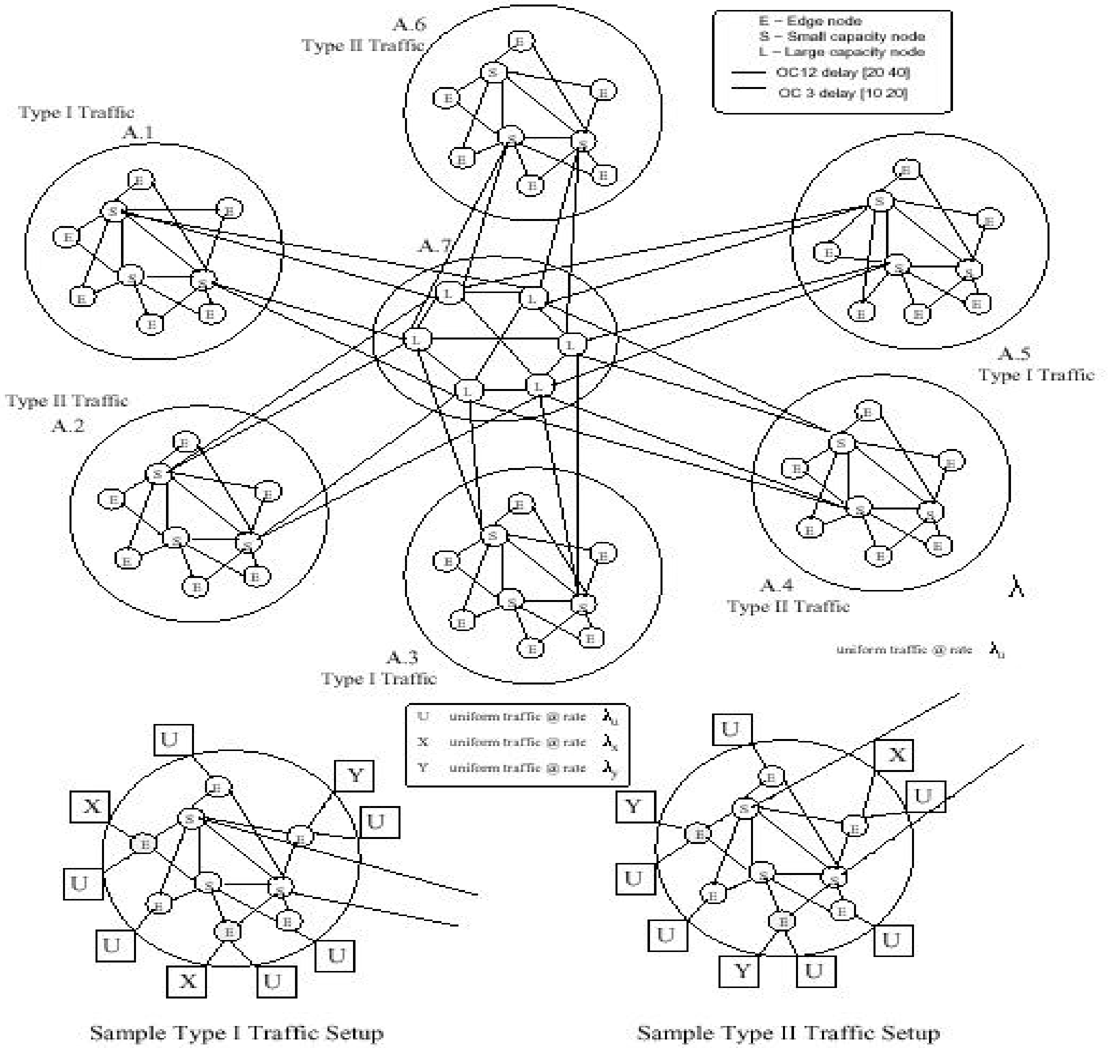

Traffic Parameters

Node Type	**Destination Choice**	**Arrival Distribution**
Type U	*uniform* among 35 other Type U nodes	*Poisson* [U]
Type X	*uniform* among 9 other Type X nodes	*Poisson* [X]
Type Y	*uniform* among 9 other Type Y nodes	*Poisson* [Y]

Application	**Bandwidth Range (uniform distribution)**	**Mean Duration (exponential)**	**%**
Low	64Kbps - 1.5Mbps	1/30s	30%
Medium	1.5Mbps –3.6Mbps	1/45s	30%
High	5Mbps - 15Mbps	1/60s	40%

Experimental Evaluation Setup

- *Warm-up* ratio of 10% for all runs.
- Default PNNI Parameter values:
  - Aggregation Policy: Symmetric Star.
  - Routing Policy: Widest-Min.hop.

Statistical Analysis

- Multiple runs with different random seeds, to compute average values.
- Test for statistical significance.
  - For the 95 % confidence interval (alpha = 0.05), the data values lie on average within 2 % of the population mean.
- Hypothesis testing of difference of means using the Student’s t-distribution (for small sample sizes).

Establishing Value of Route Caching

- As a flat "single peer group" network scales, the average setup time should decrease for pre-computed routing as compared to on-demand routing.
- For a fixed topology, the average call acceptance ratio should decrease, when there are NO updates to the route cache.

Preliminary Results

Effect of Topology Scaling on Average Call Setup Times

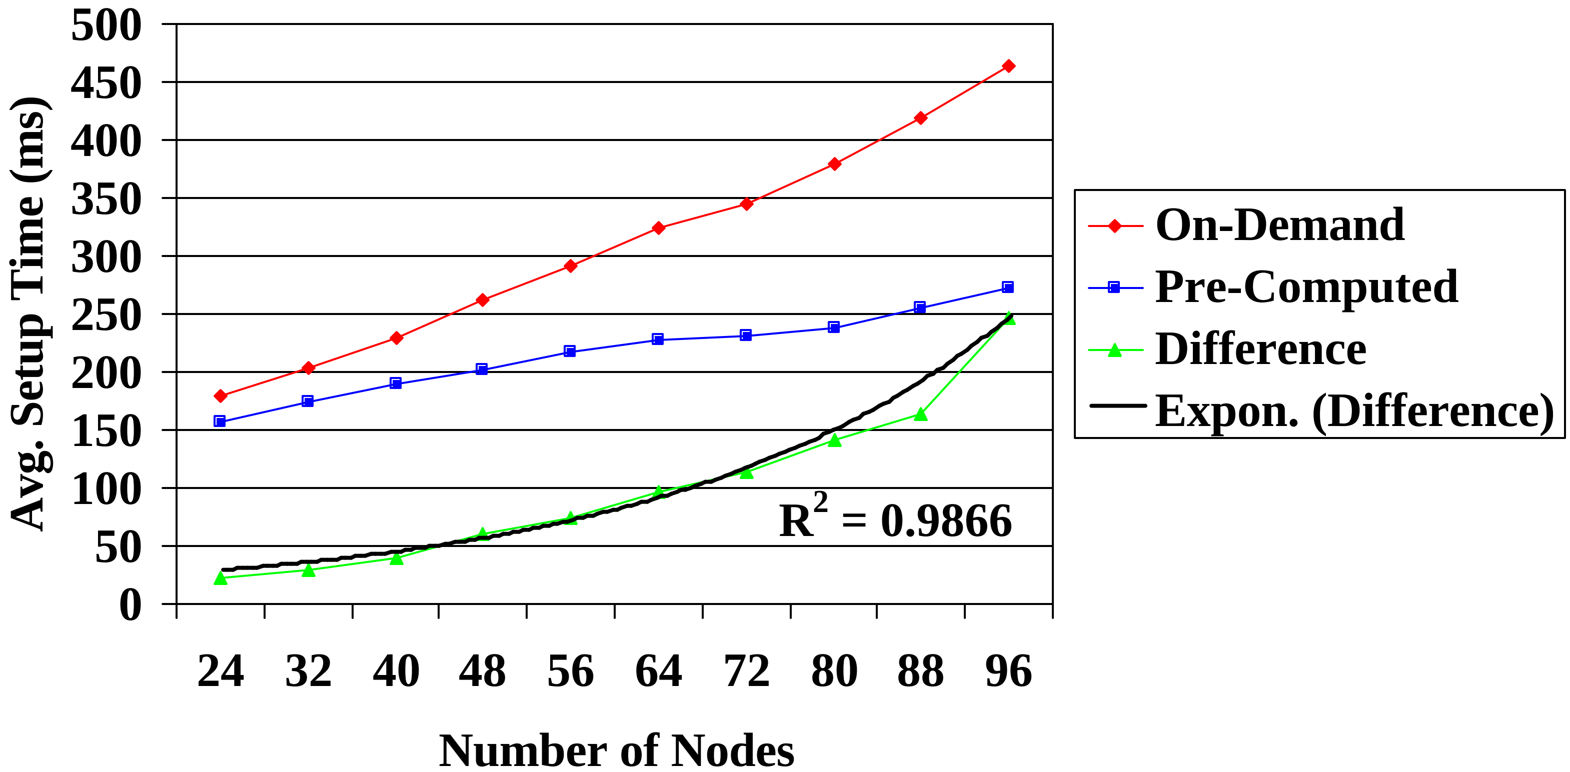

Percentage Difference in Avg. Setup Times for Flat Network

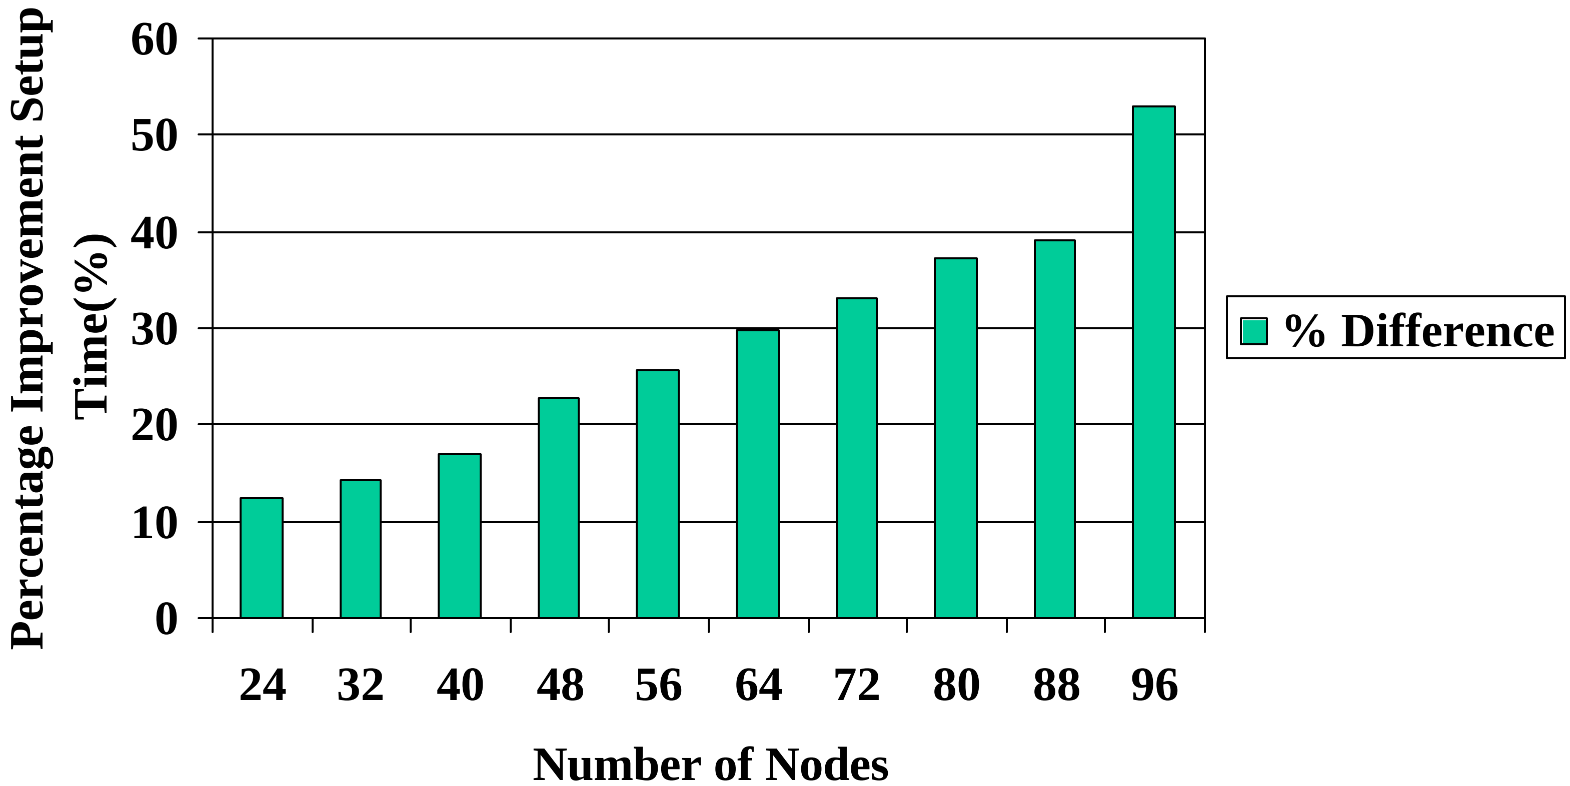

Average Bandwidth Acceptance Ratio with Increasing Load

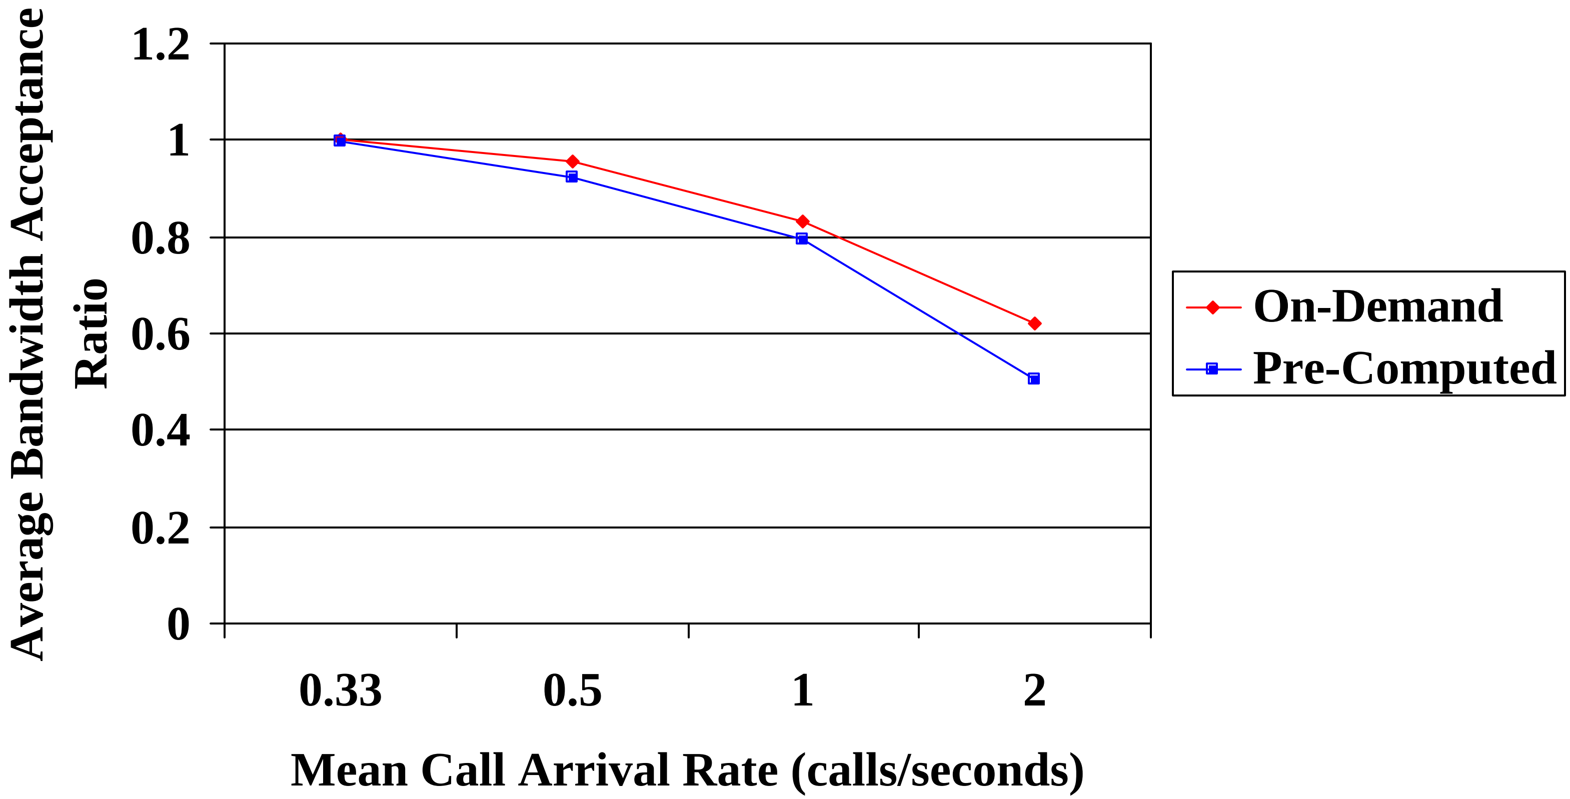

Pre-computed routing for Multiple Peer Group topologies

Motivation, Evaluation, and Results

Motivation

- Multiple peer grouping is the way by which PNNI handles the problem of network scaling.
- No previous study on pre-computed routing for MPG to our knowledge.
- Effect of PTSE aging and its repercussions on the quality of the routes more interesting within a MPG hierarchy.

Hierarchical MPG Topology

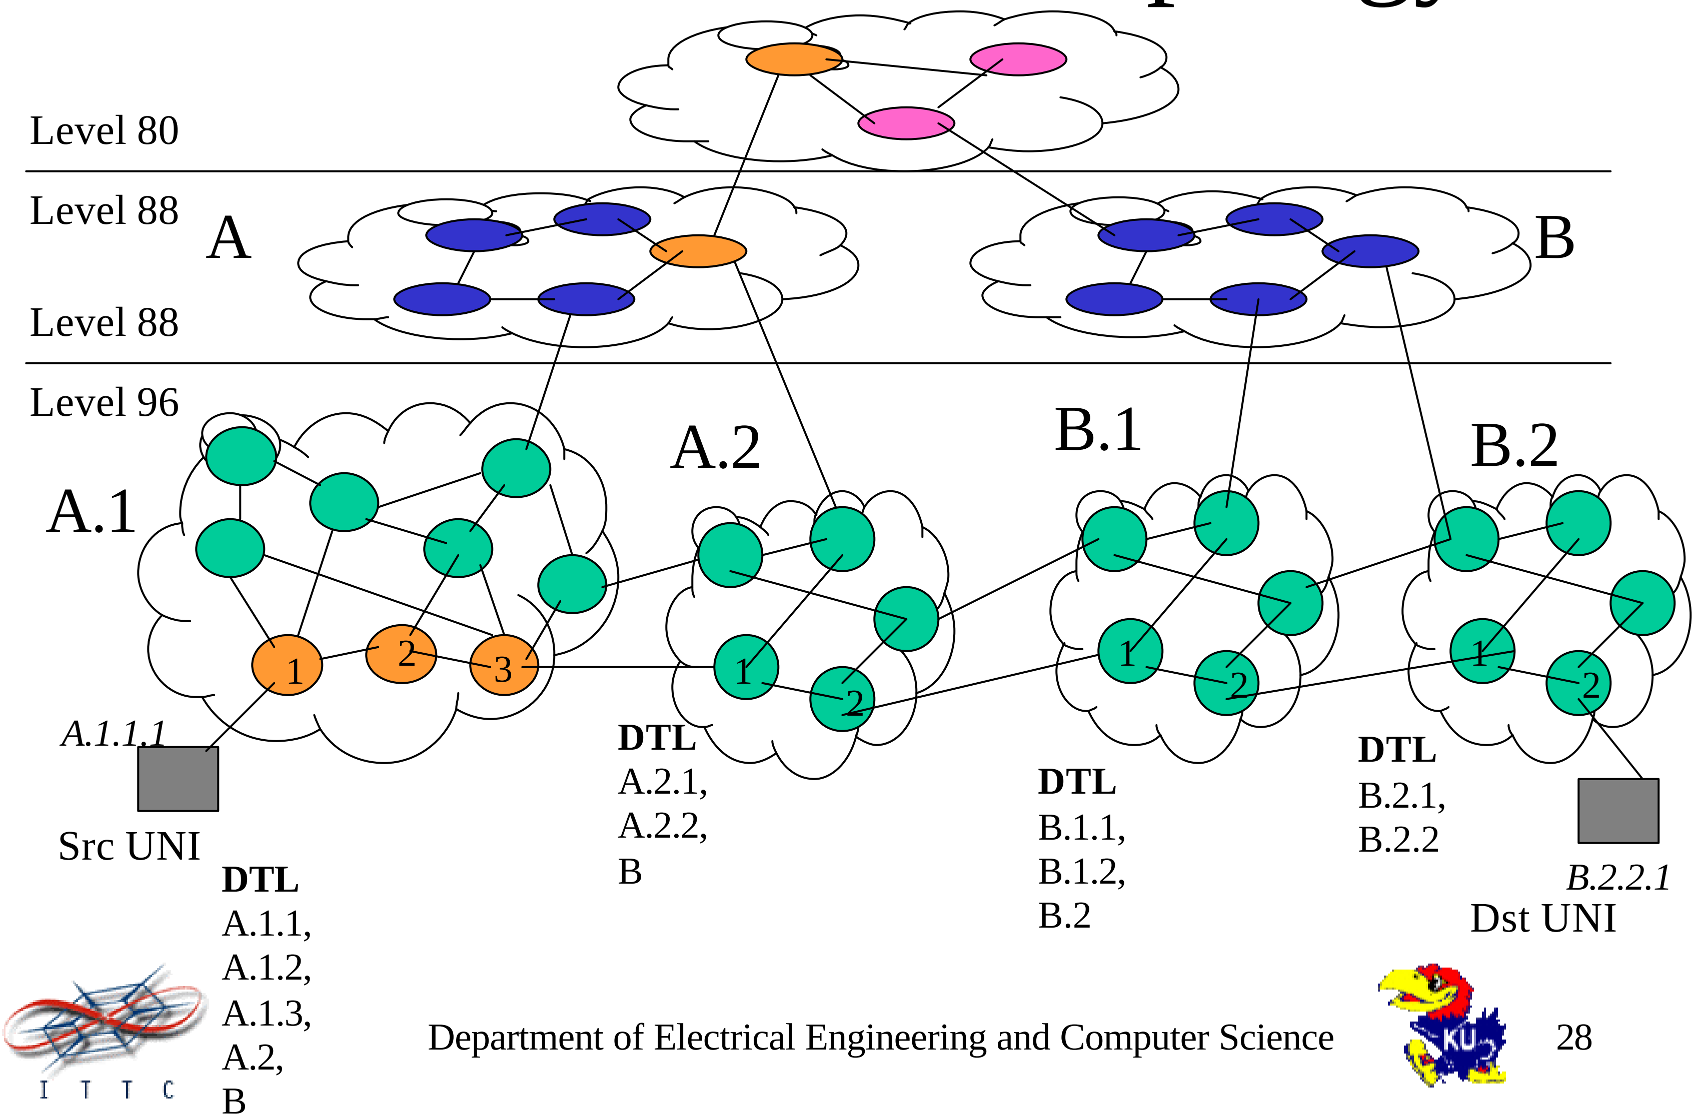

Route Caching in MPG networks

- Effect of Number of Peer Groups on average call setup time.
- Effect of Number of Peer Groups on average call acceptance ratio.
- Location of call failures as the Number of Peer Groups is changed.

Average Setup Time with Peer Group Size (no cache updates)

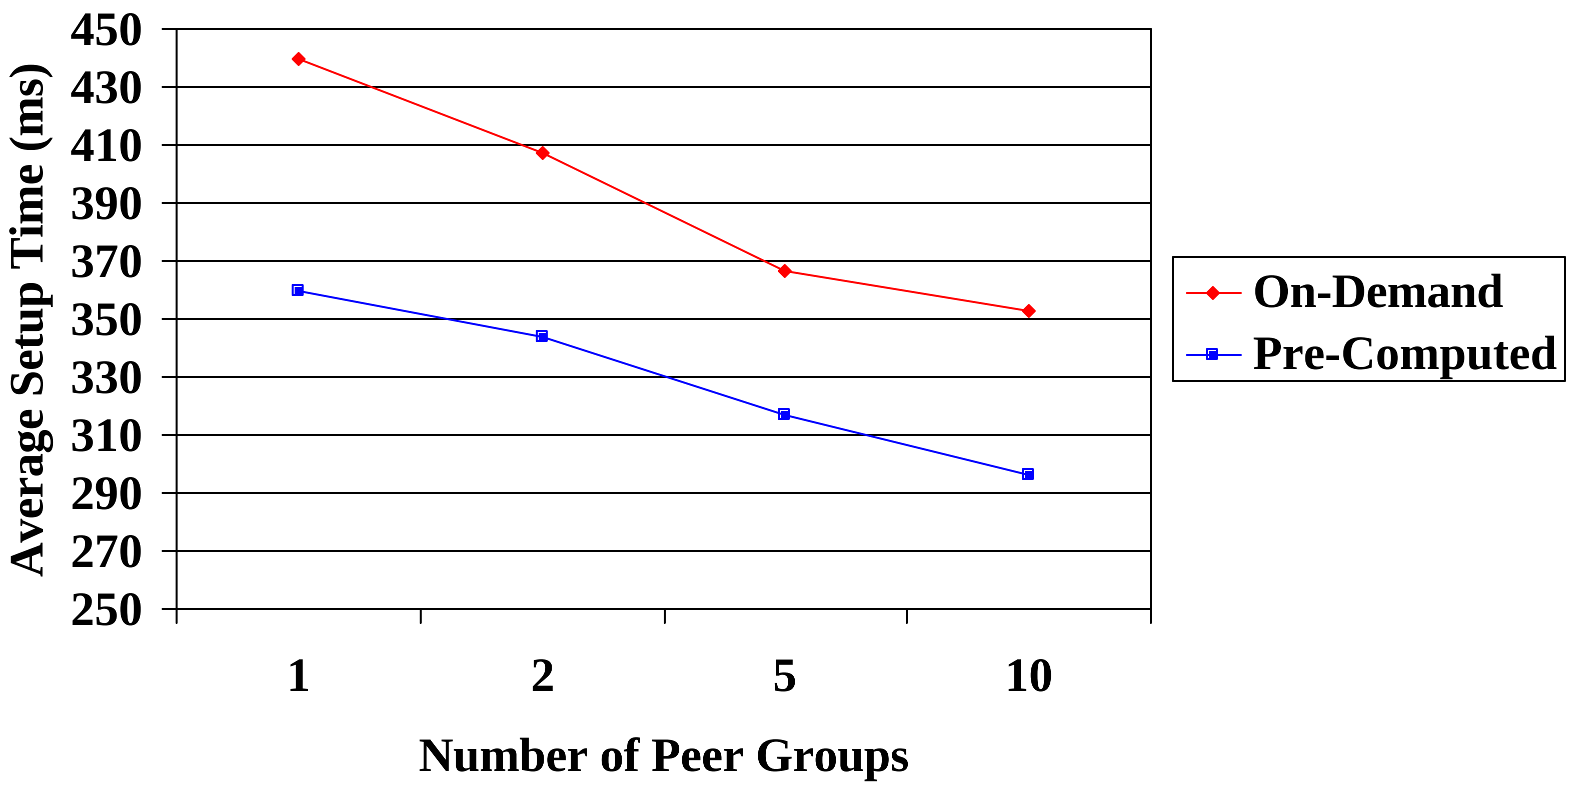

Percentage Difference in Avg. Setup Times for MPG

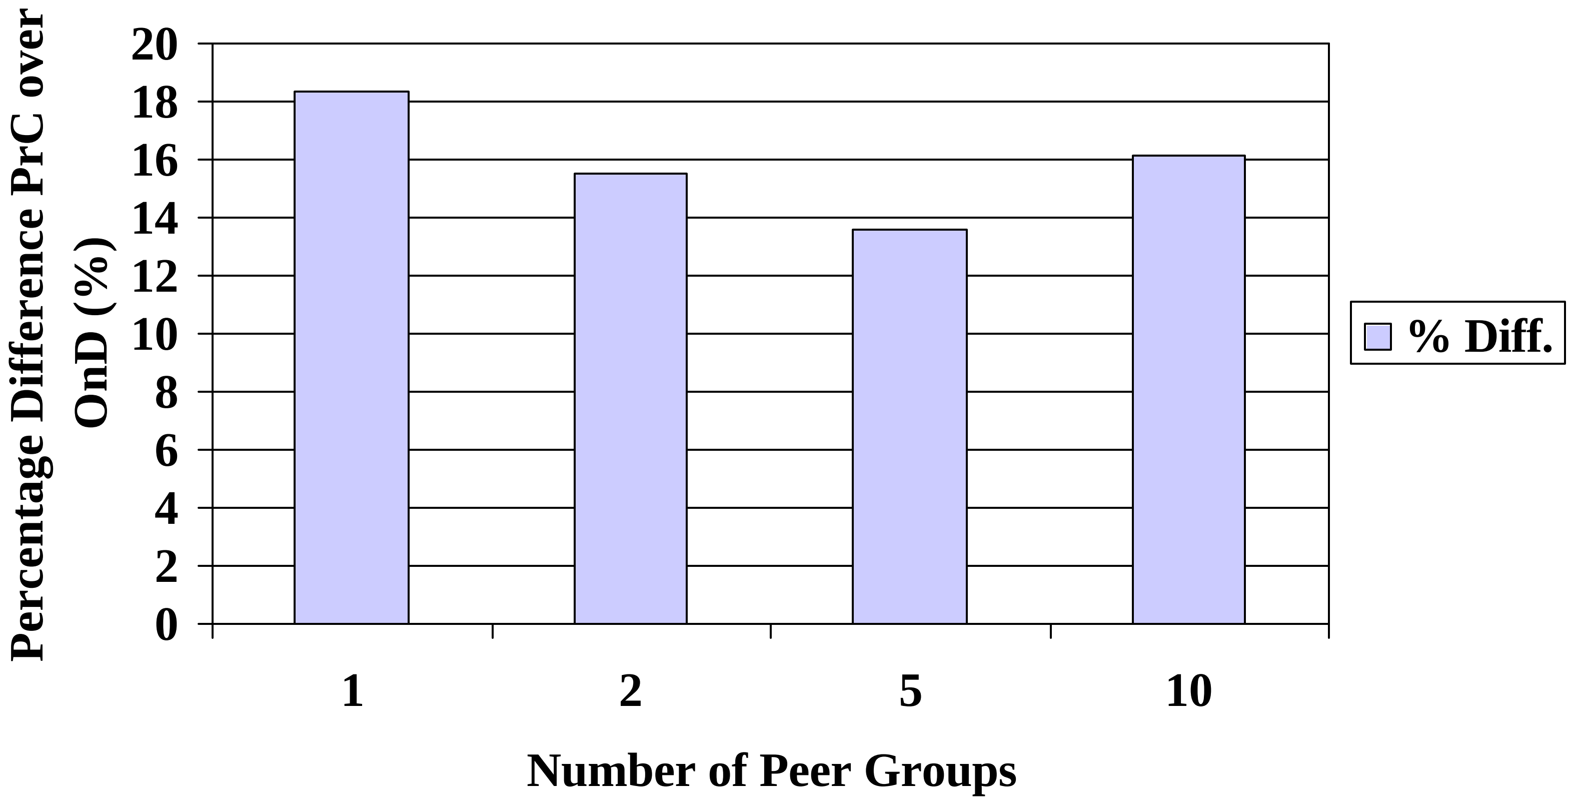

Average Bandwidth Acceptance with Peer Group Size

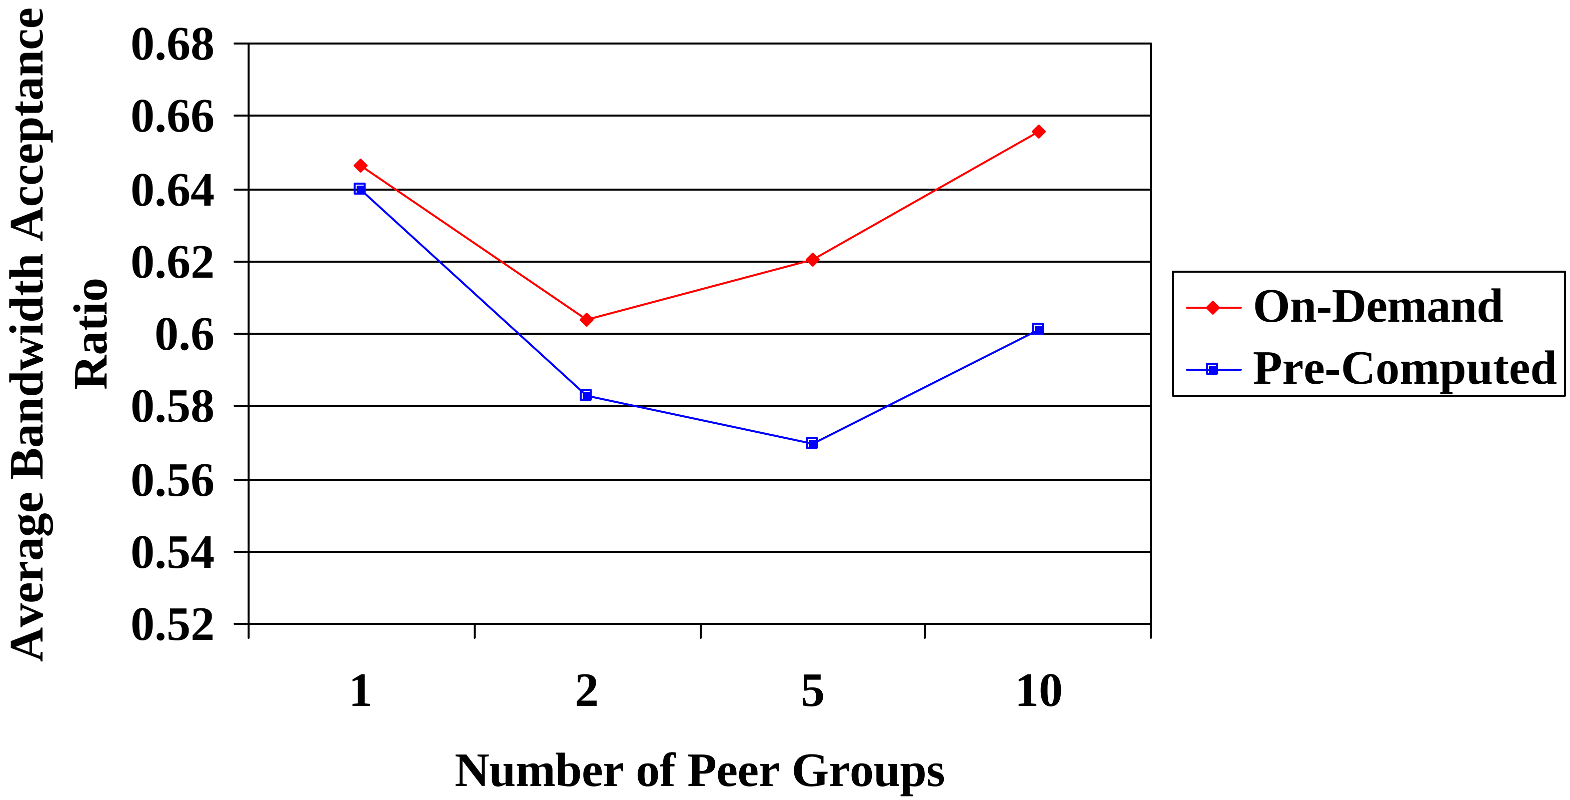

Average Database Size Vs. Peer Group Size

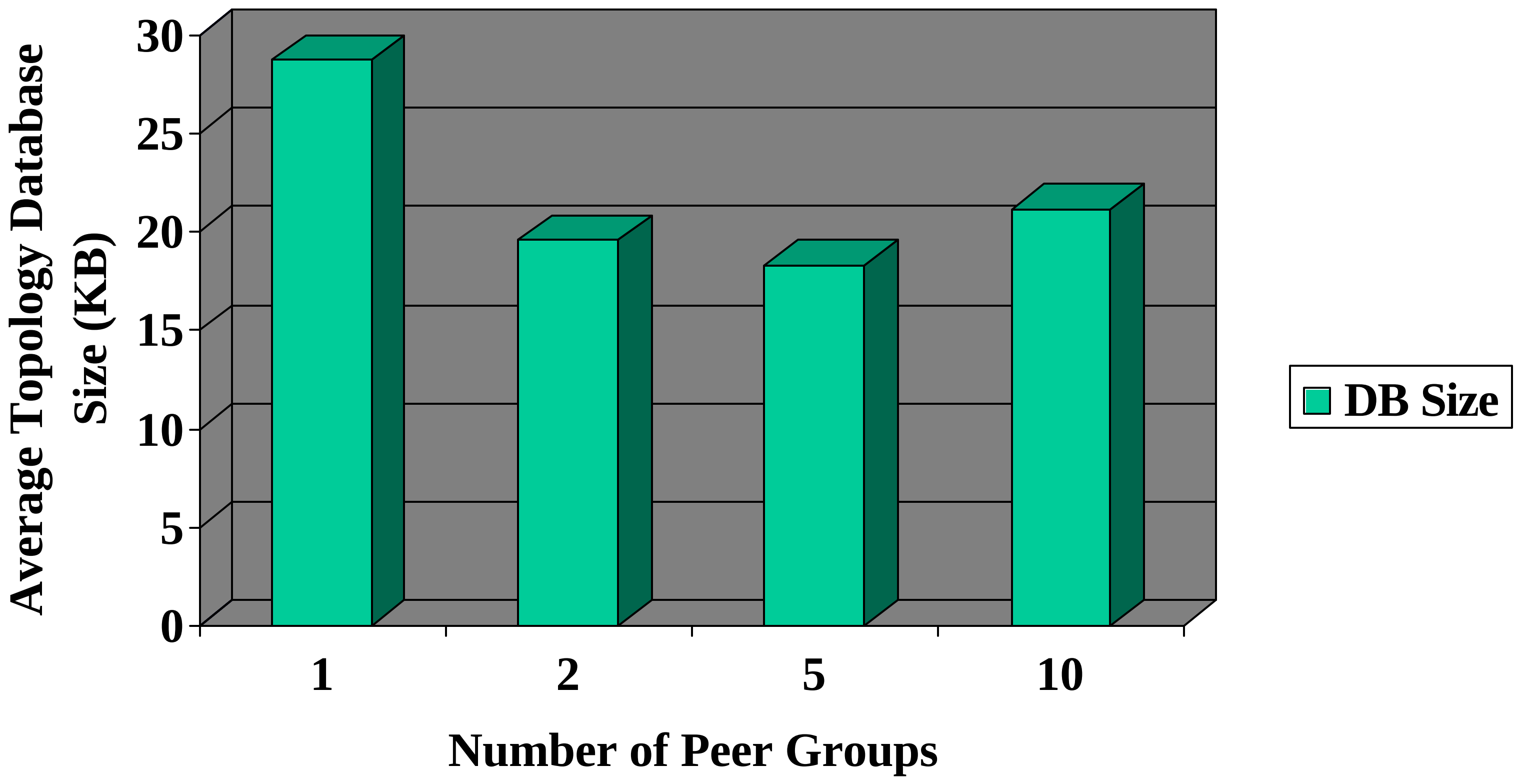

Location of Call Failures with Peer Group Size (On-Demand)

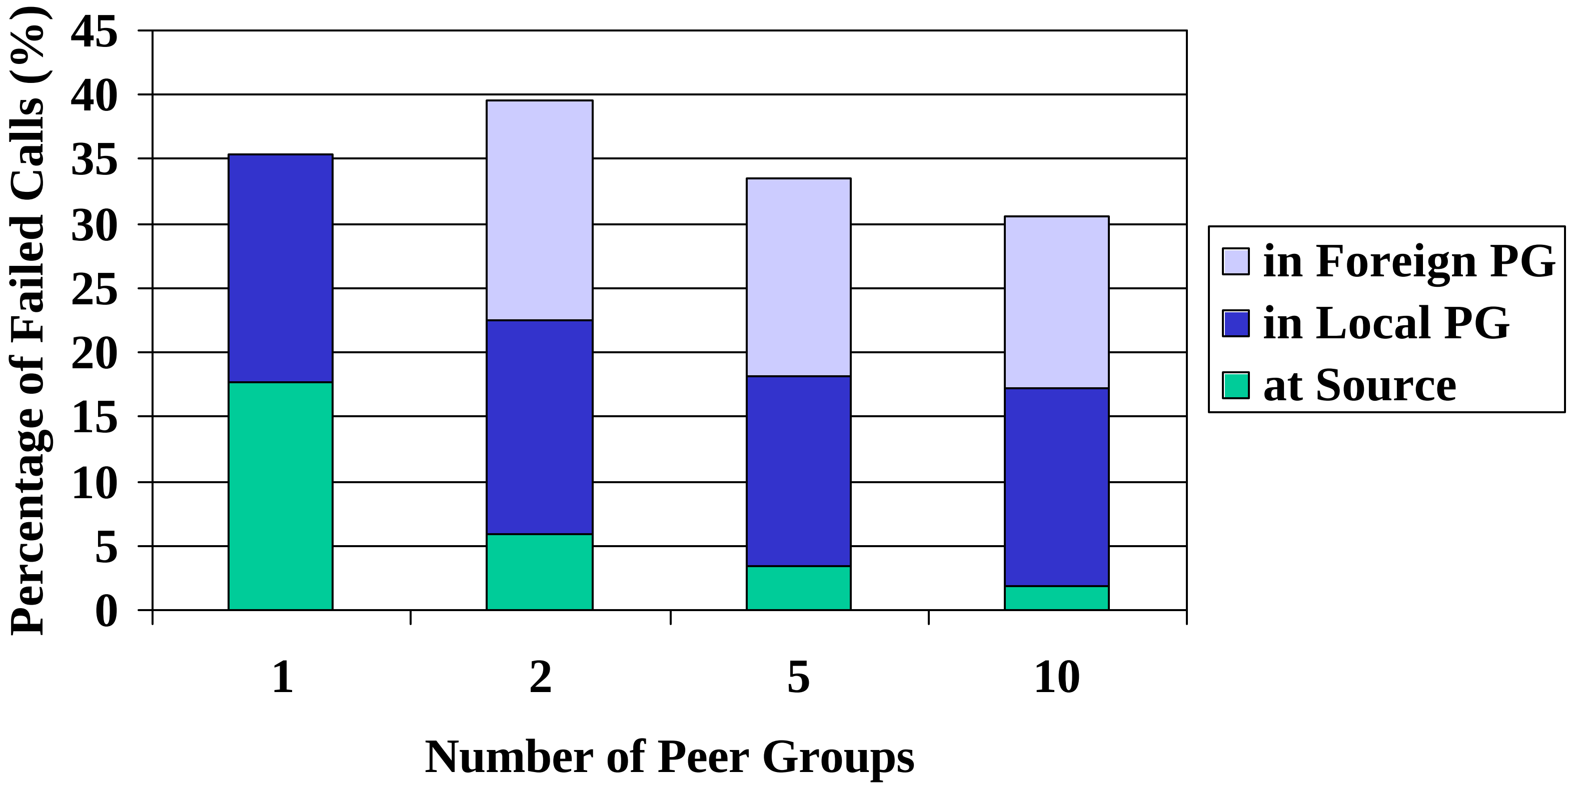

Location of Call Failures with Peer Group Size (Pre-Computed)

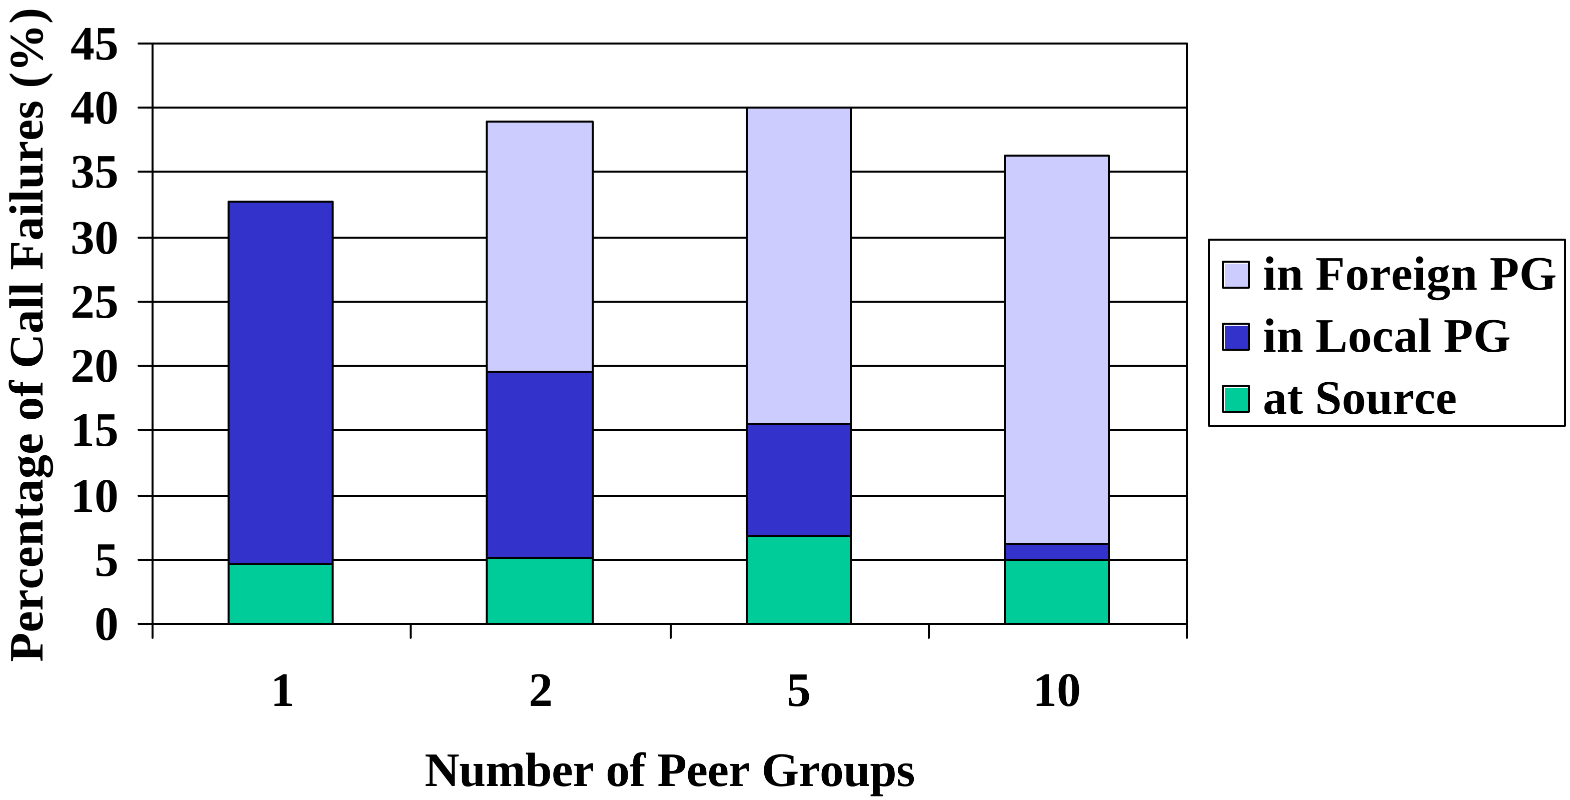

Comparison of Failure Locations (Pre-Computed Vs On-Demand)

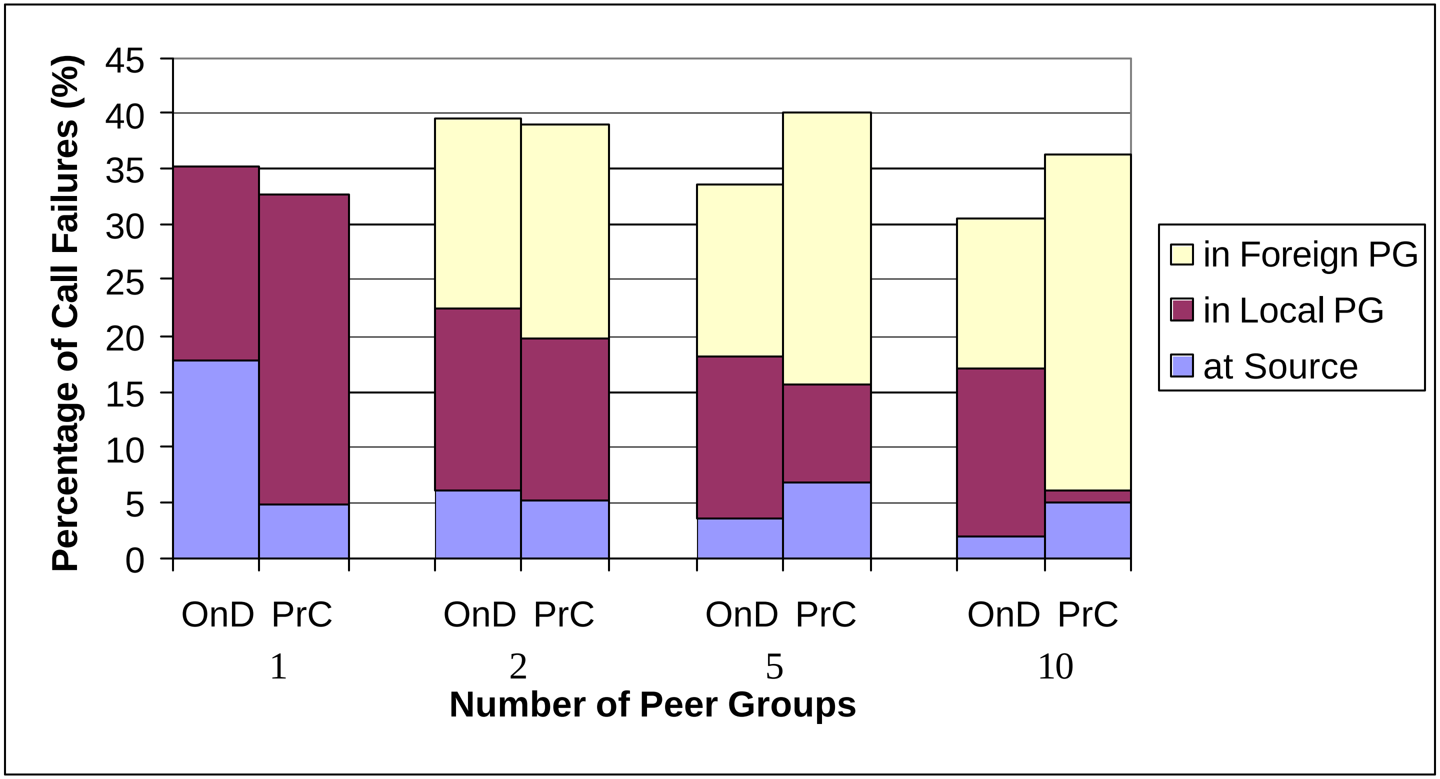

Route Cache Update Heuristics

Characteristics of a Route Cache Update Heuristic

- Intended to refresh the pre-computed cache with routes based on new topology data.
- The cache update policy should try and follow the network state as closely as possible.
- It should aim to retain the lower setup times obtained due to caching and attain the call success ratio of on-demand routing.

Types of Heuristics

- No updates (trivial case).
- Timer based updates.
- Crankback based invalidation.
- PTSE count based updates.
- Combinations of above policies: ptse_timer, ptse_crankback, crankback_timer, ptse_timer_crankback.

Comparison of Heuristics for Updating the Cache

Metrics:

- Average call setup time.
- Average call acceptance ratio.
- Combined metric = average call rejection ratio * average call setup time.

Timer based updates

- Simplest and least intuitive of all update heuristics.
- Ideal value of timer??
- Hypothesis: For an MPG topology, synchronizing the timer with the re-aggregation rate will yield the best setup times.

Average Setup Time Vs. Update Timer Period

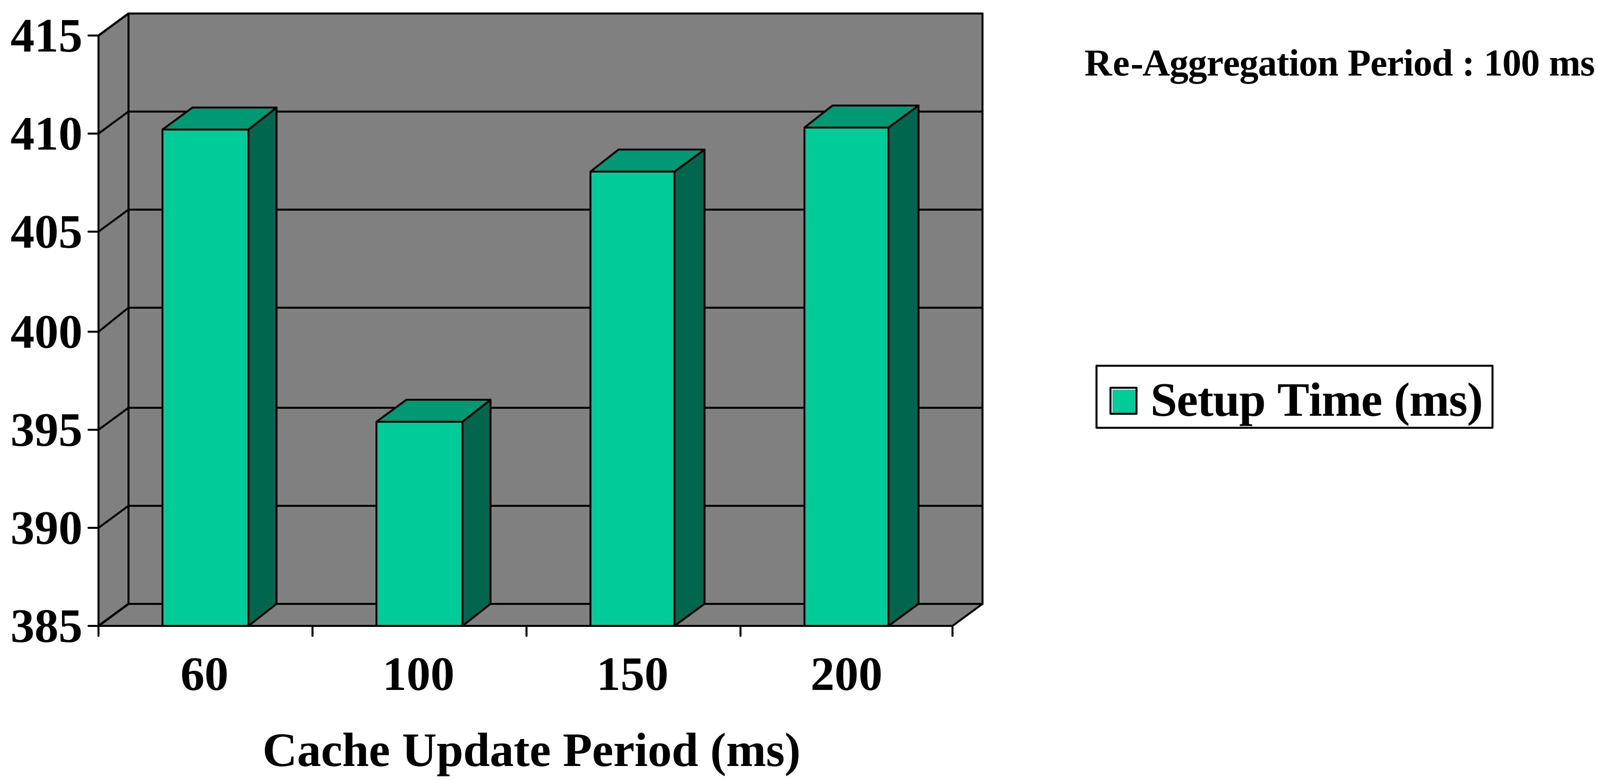

Average Bandwidth Acceptance Ratio Vs.Update Timer Period

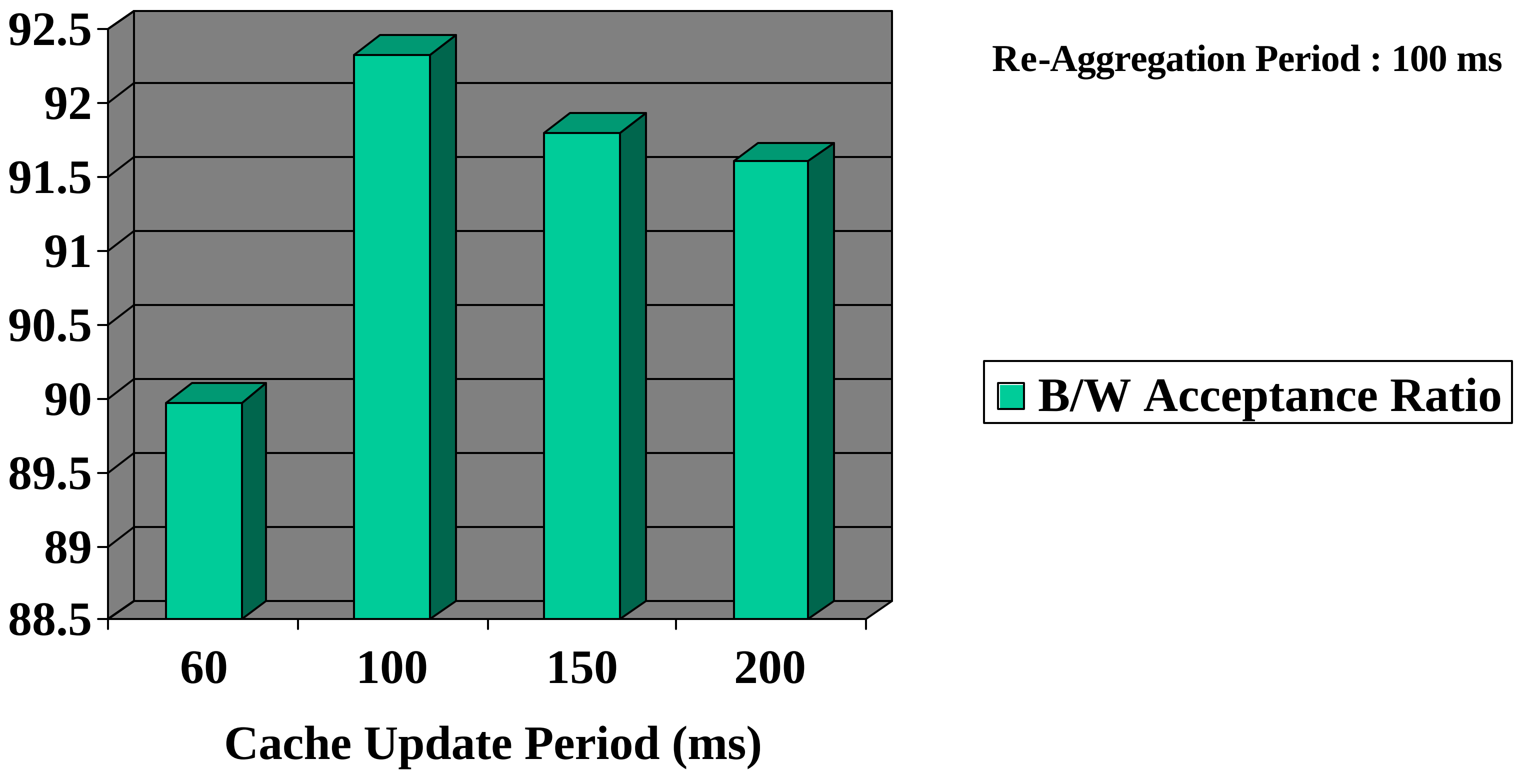

Crankback Based Invalidation

- Logical update condition, since a crankback happens when a route failure occurs.
- Indication of inaccurate information or staleness of cache entry.
- Results in invalidation of cache entry.
- Subsequent on-demand routing and re-filling of the cache entry.

PTSE Count Based Update

- PTSEs are flooded when there is significant change and/or after sufficiently long idle period.
- Counting PTSE's received for any Peer Group/ Logical Level, is a good indicator of state change in that Peer Group/ Level.
- This metric is associated with network specific components, and is available essentially for free.
- PTSE count, upon reaching a specified value, is used to initiate a route cache update.

Combined Heuristics

- Each individual heuristic detects network changes by considering mutually independent information.
- Combining them should lead to a better heuristic for initiating route cache updates.

Avg. Setup Time for Different Cache Update Heuristics

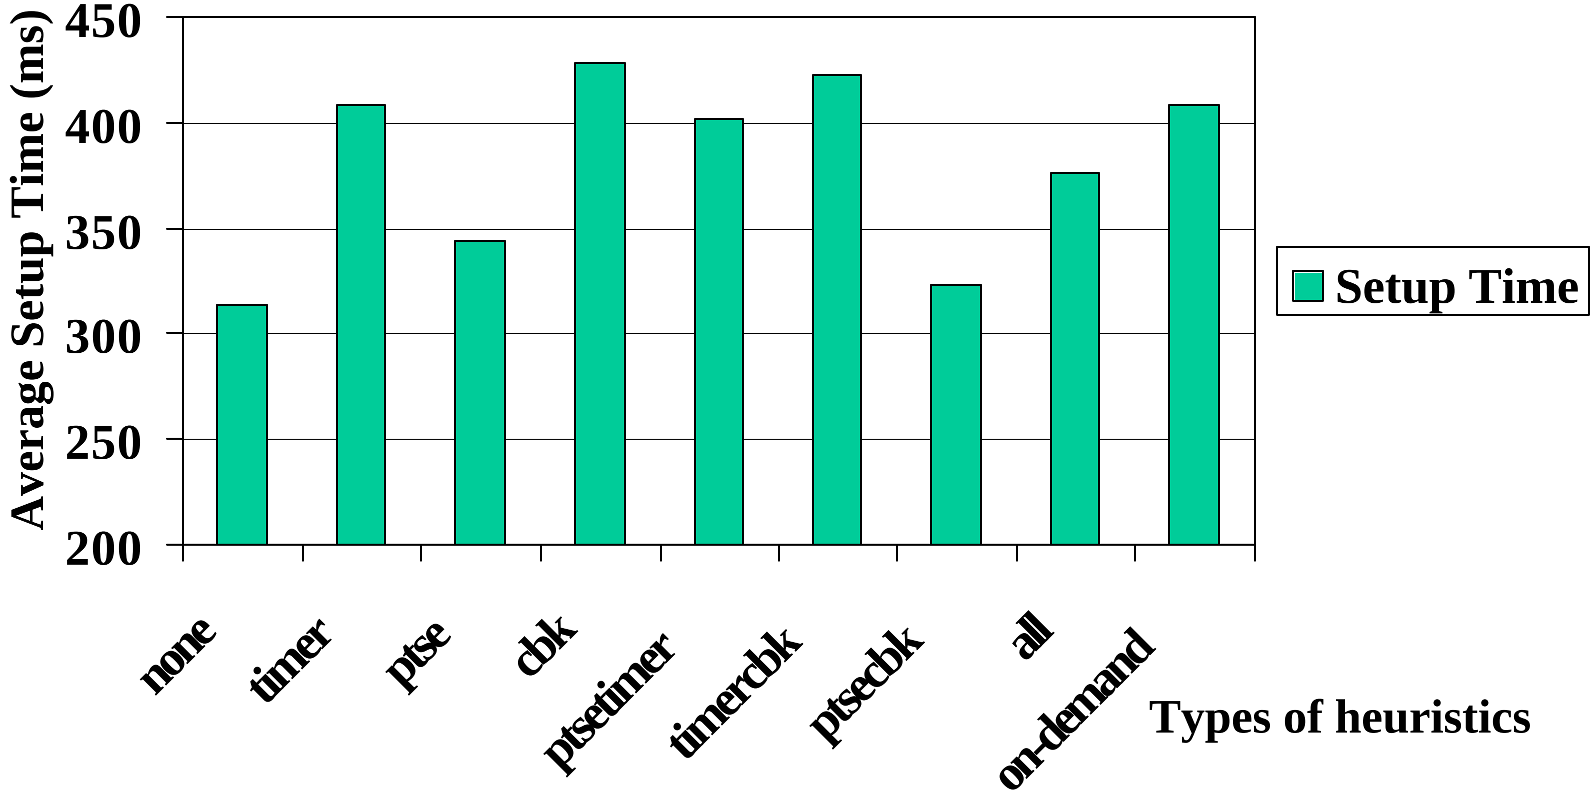

Avg. Call Acceptance Ratio for Diff. Cache Update Heuristics

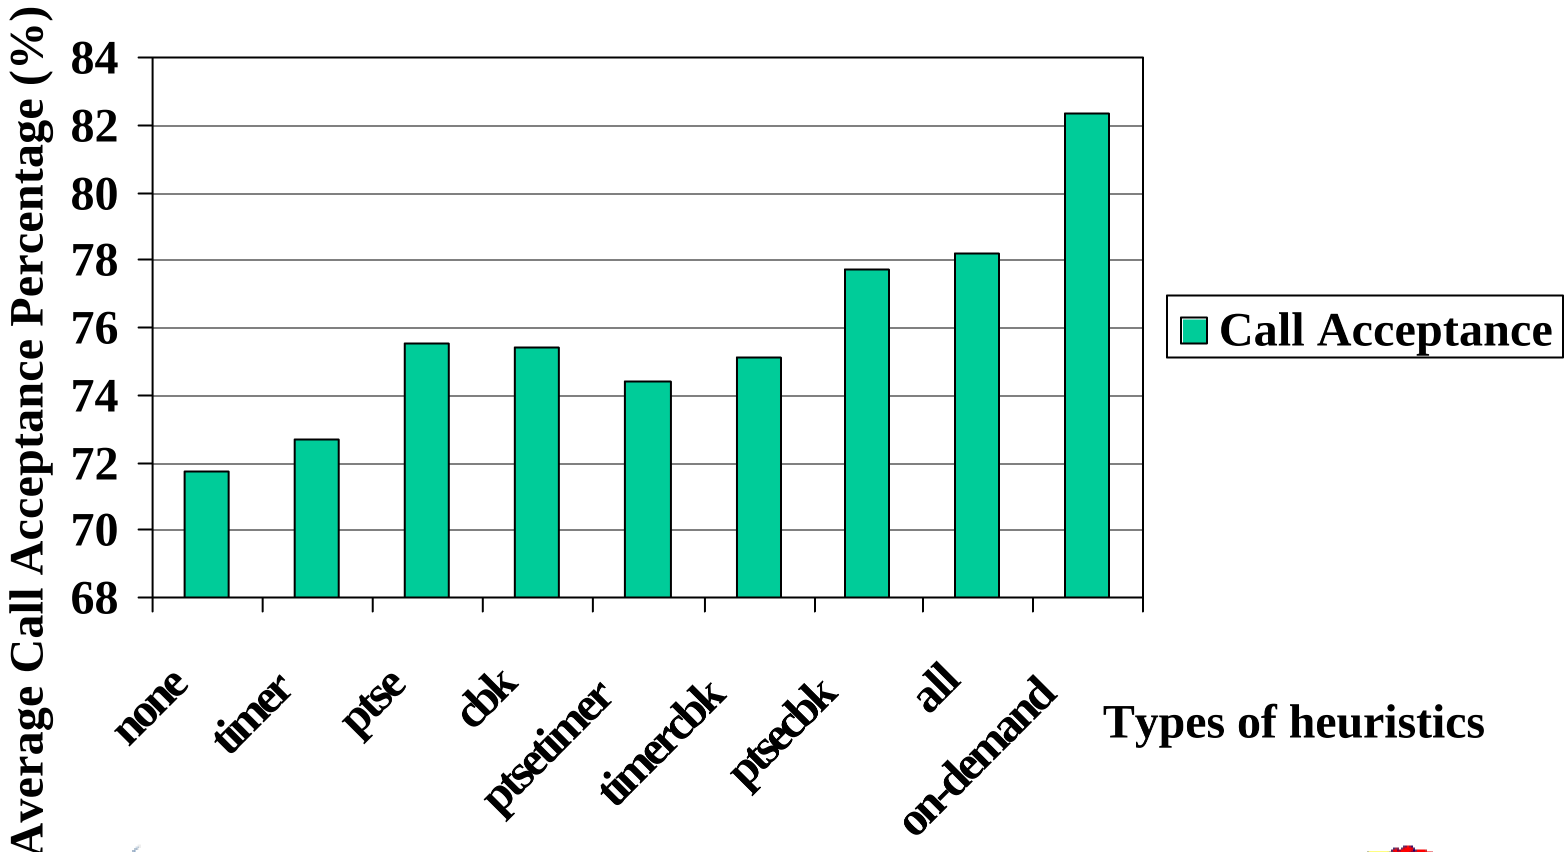

Performance of Route Cache Update Heuristics

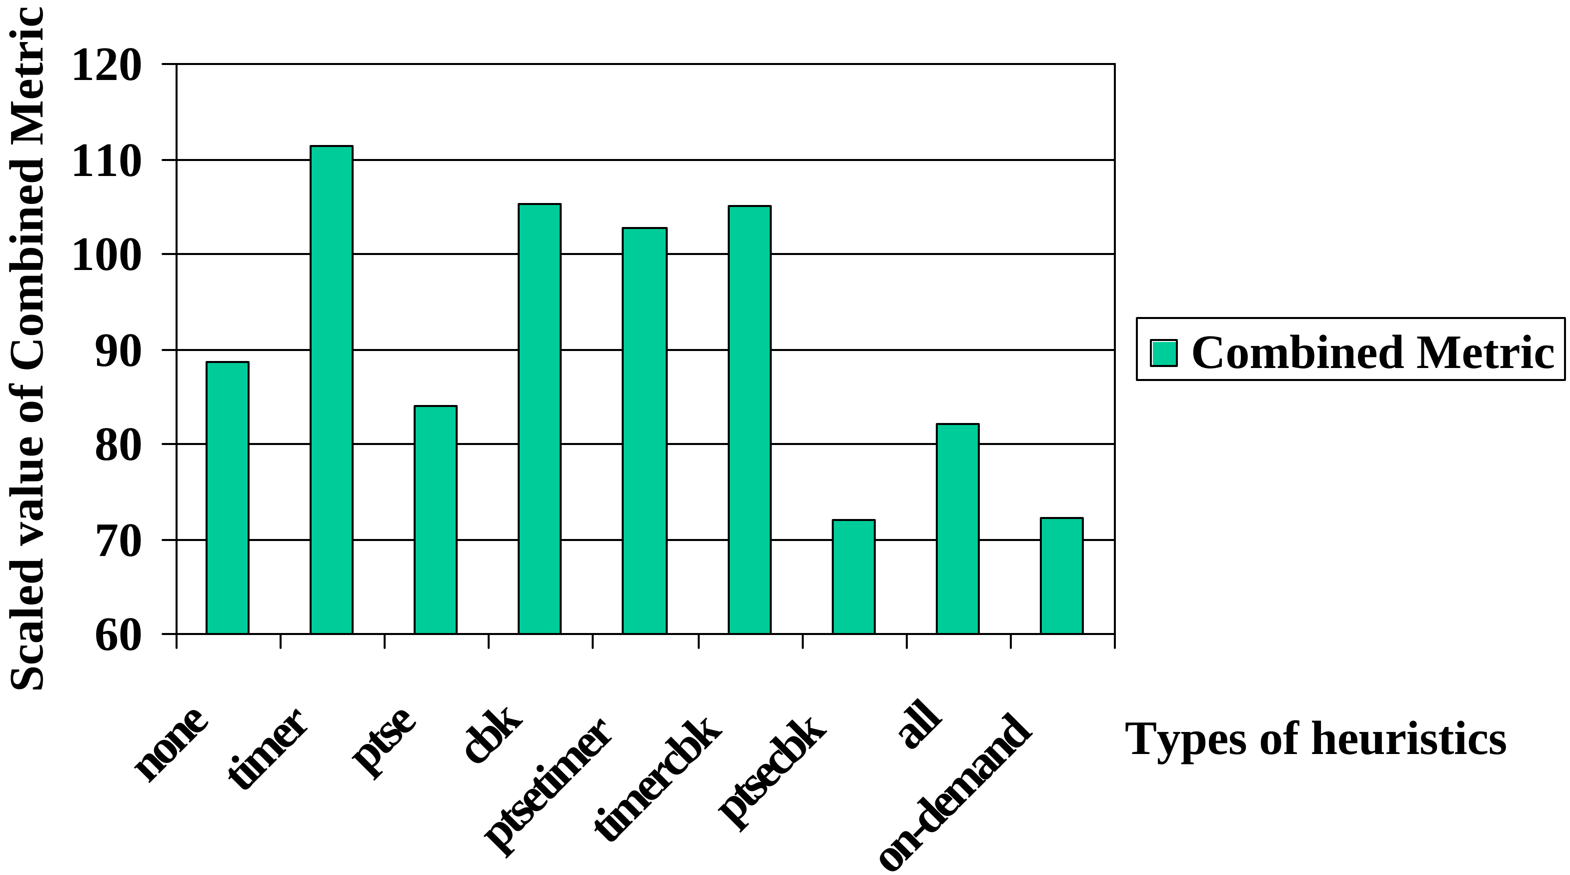

Revalidating results

- Applying ptse_crankback heuristic to the chosen MPG topology.
- Study effect of update heuristic on setup times and call acceptance ratio.
- Investigate location of failures.

Setup Time in Multiple Peer Group with Route Cache Updates

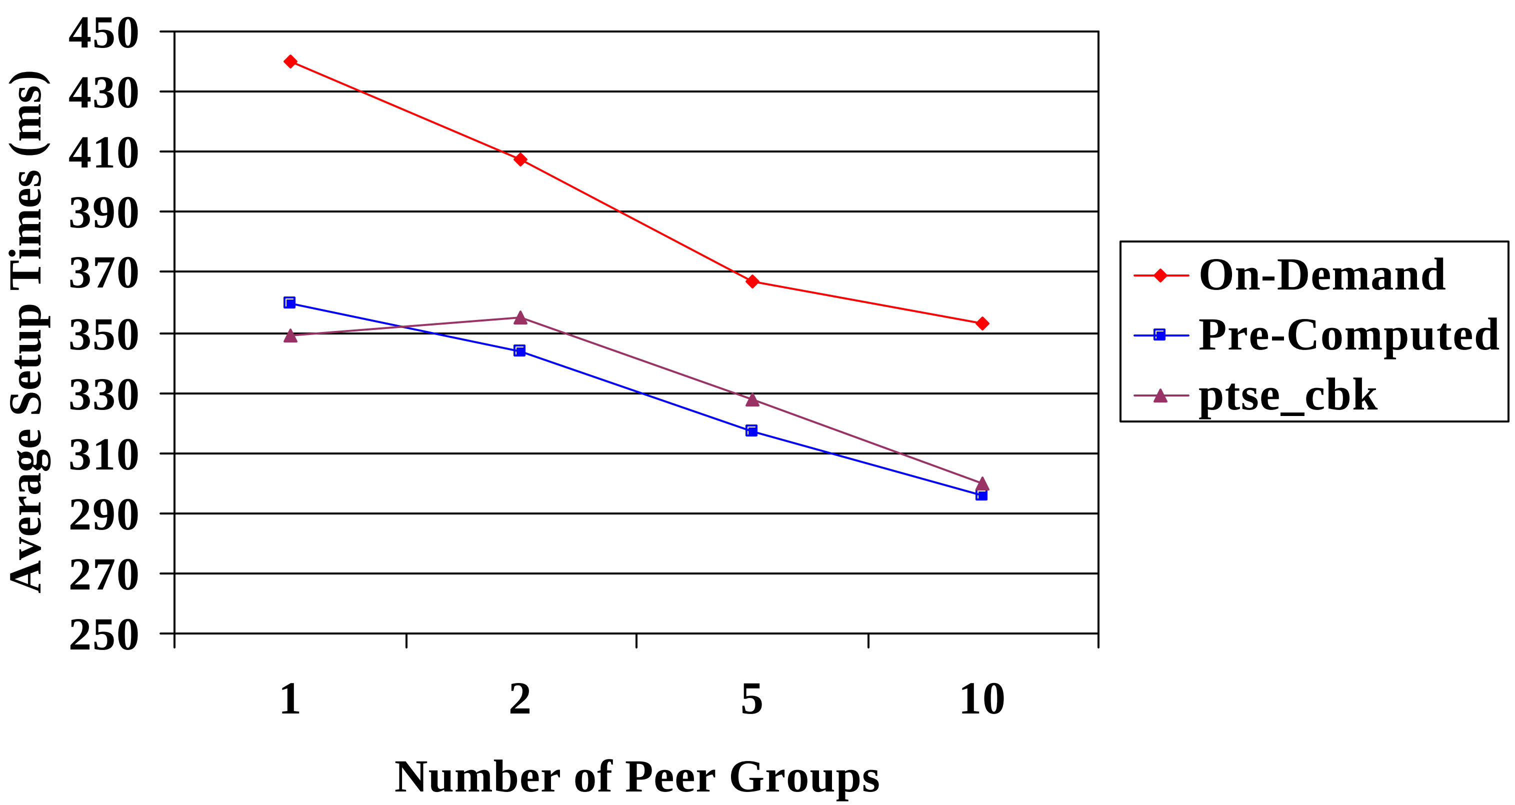

B/W Acceptance in Multiple Peer Group with Route Cache Updates

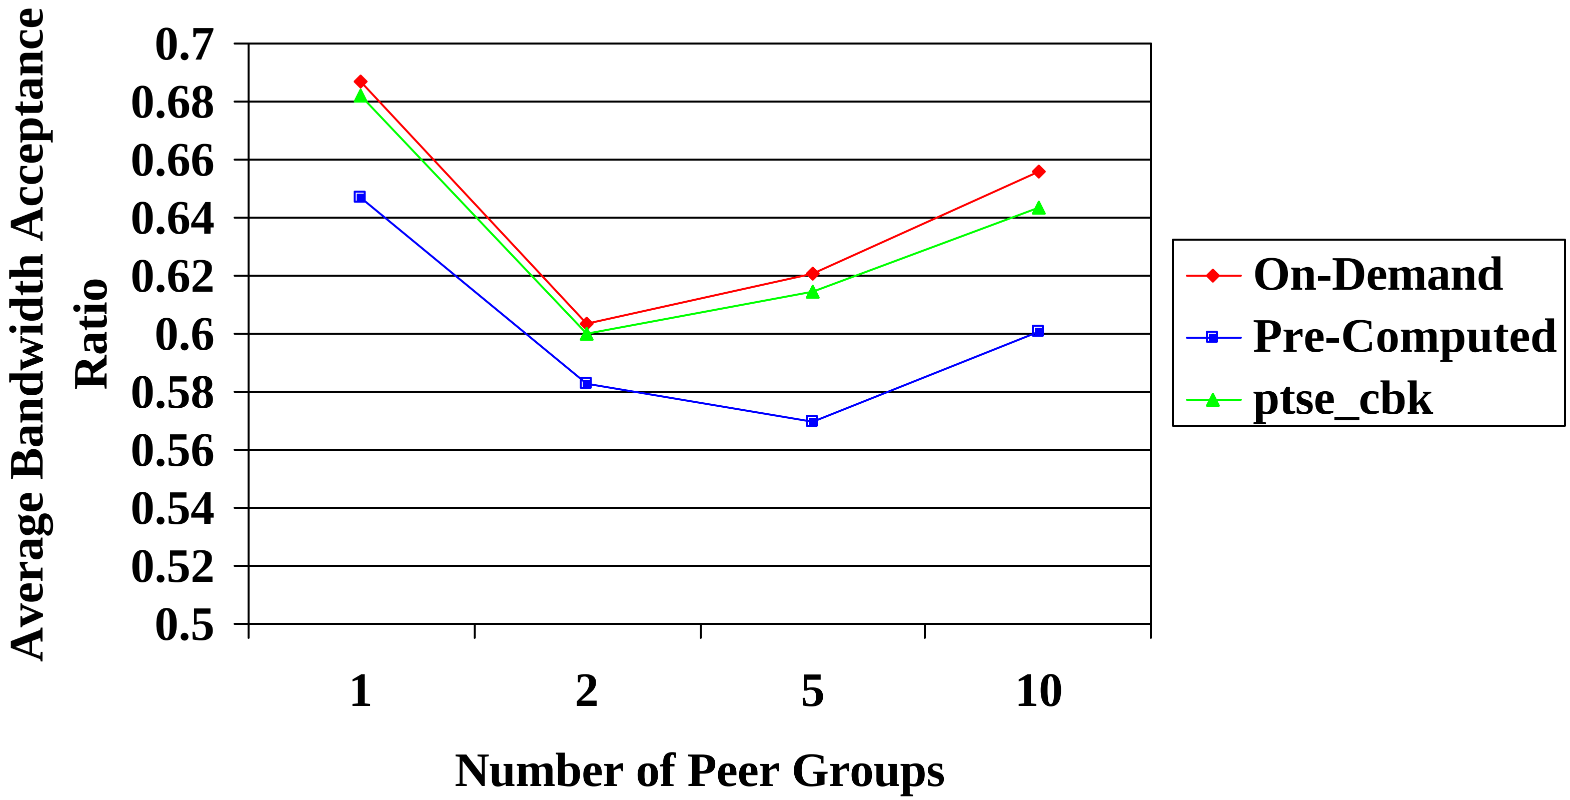

Best of Both Worlds !!

Location of failures

- Hypothesis: With route cache updates, the number of failures in foreign peer groups should be a smaller fraction of the total number of failures, as compared to no route cache updates.

Location of Call Failures (no updates)

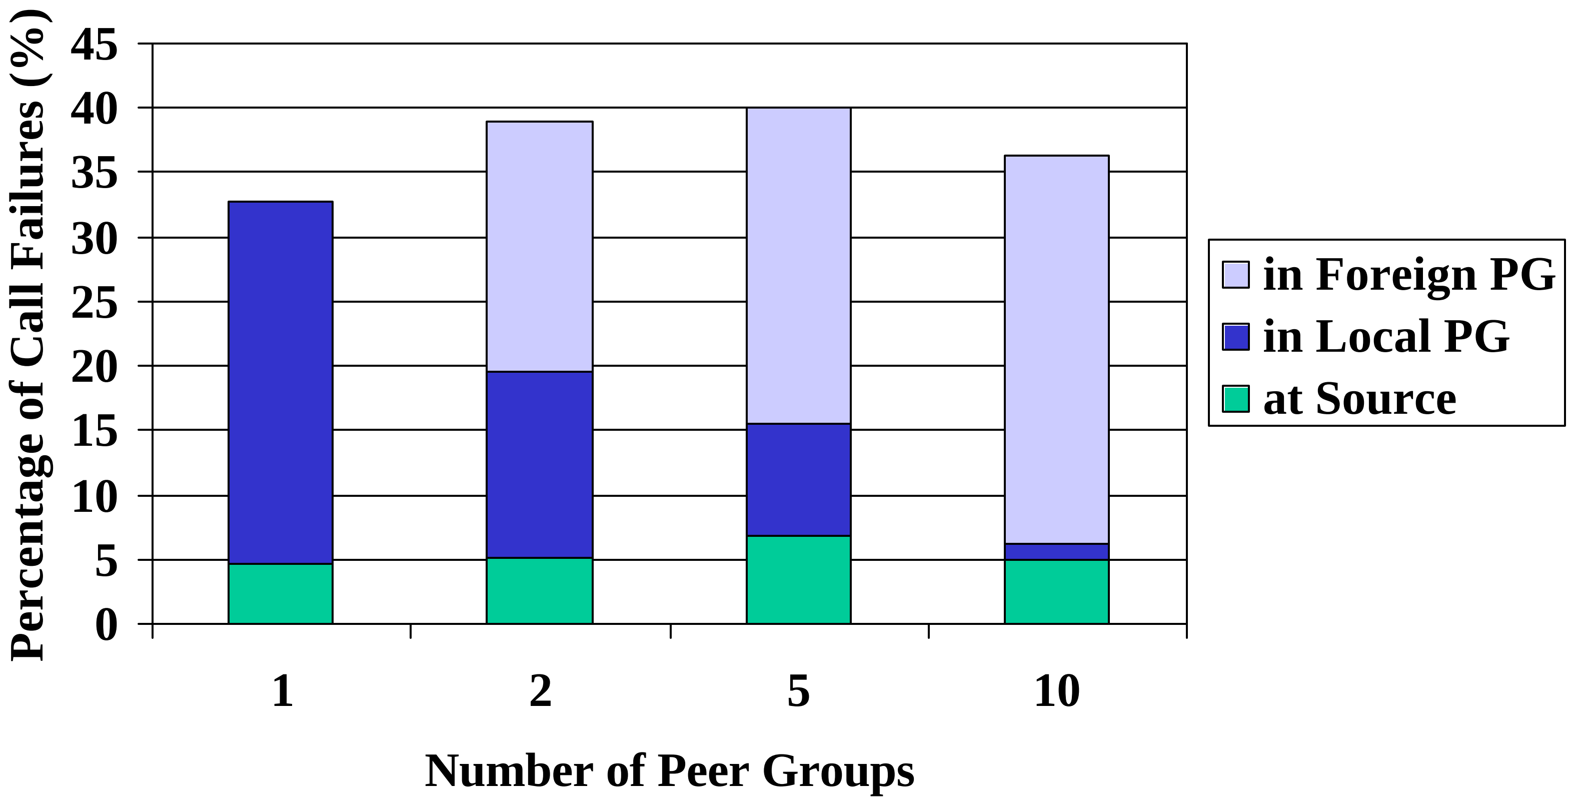

Location of Failures (with ptse_crankback heuristic)

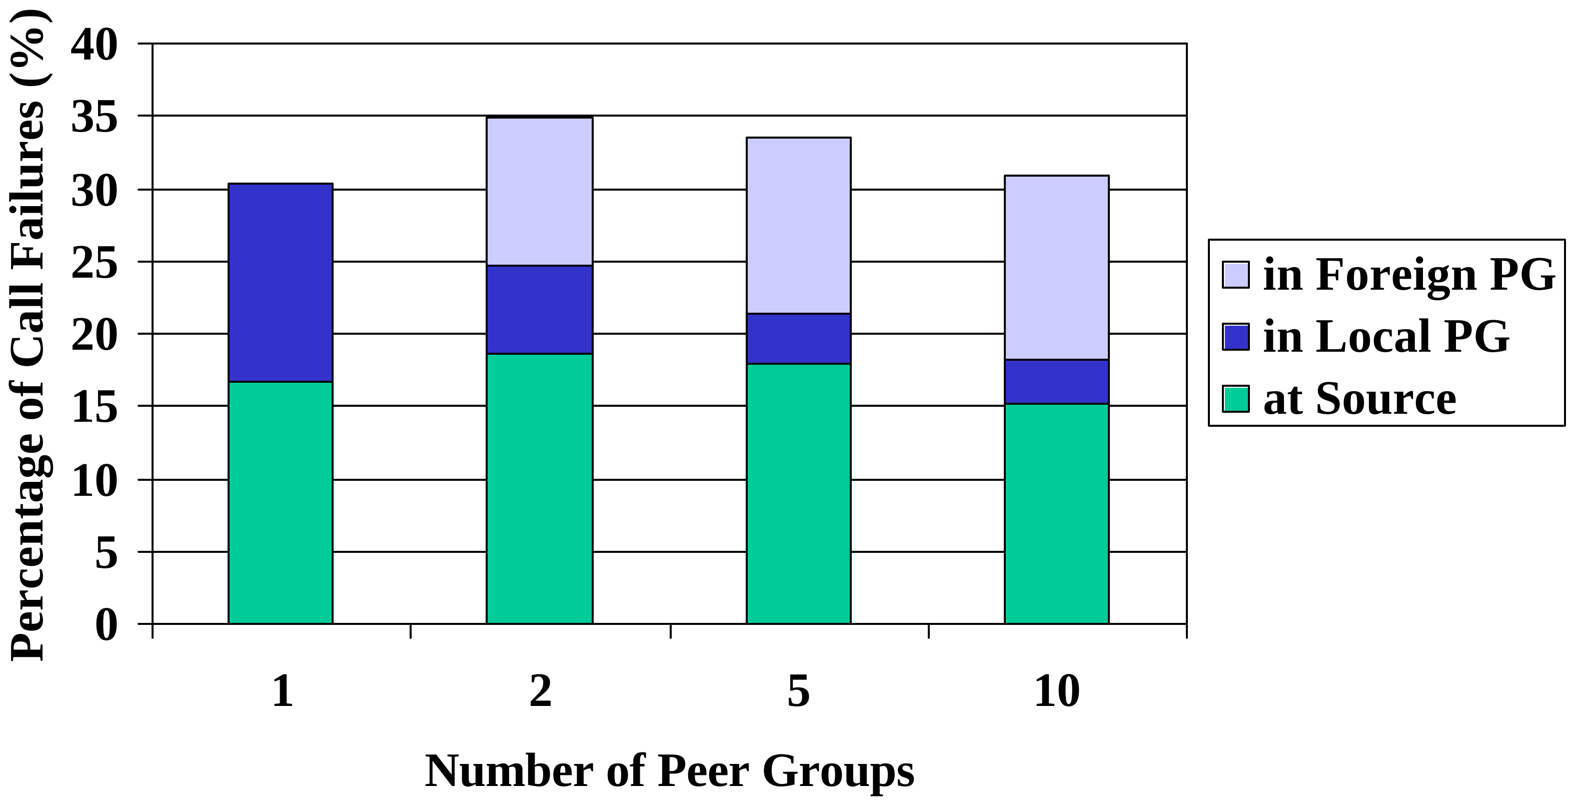

Comparison of Failure Locations (No Update Vs. Ptse_cbk)

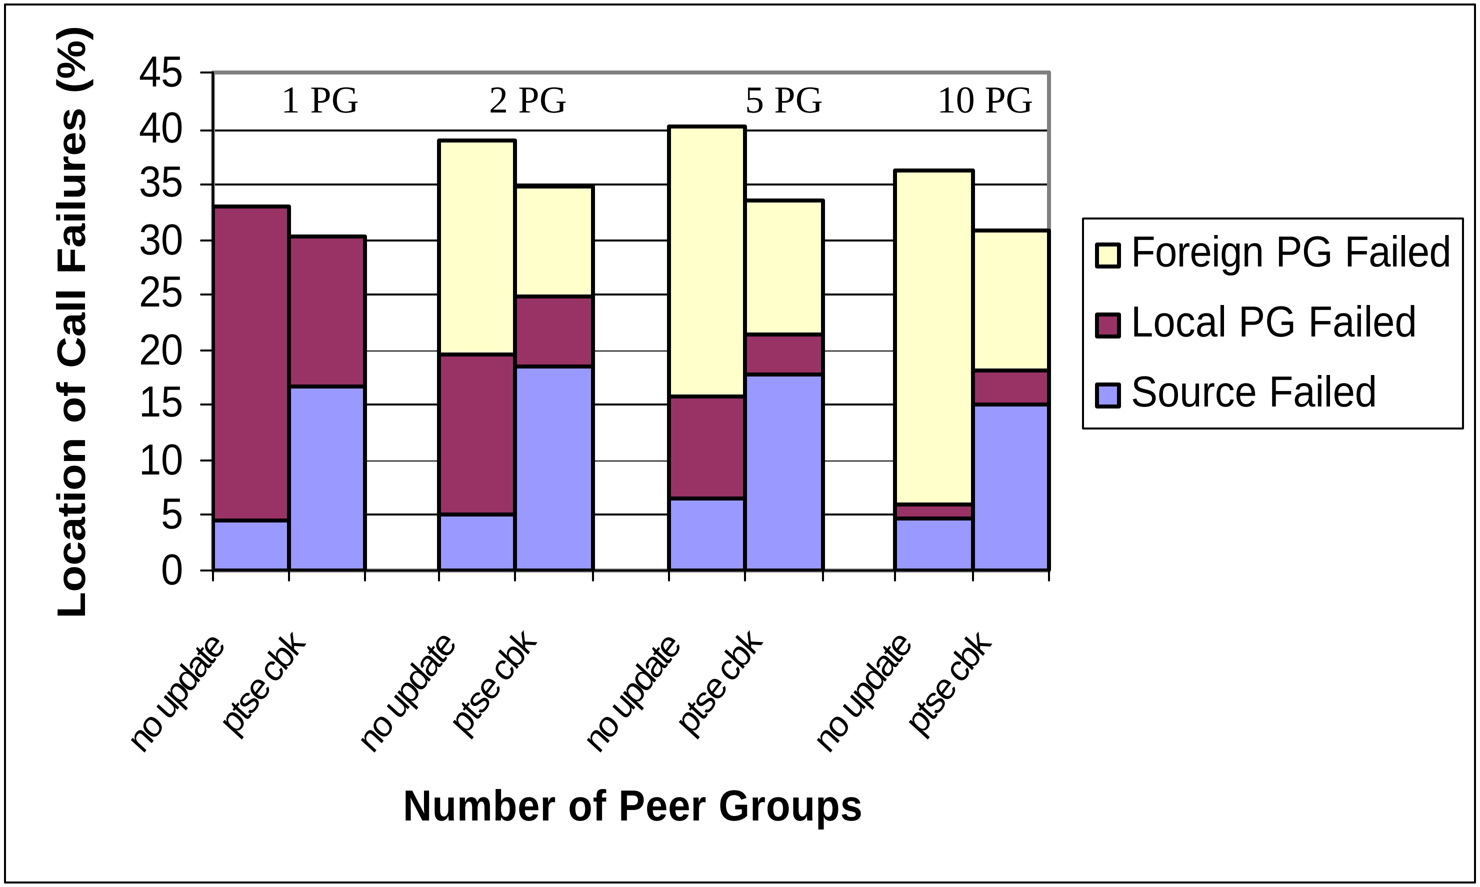

On-Demand Processing Times

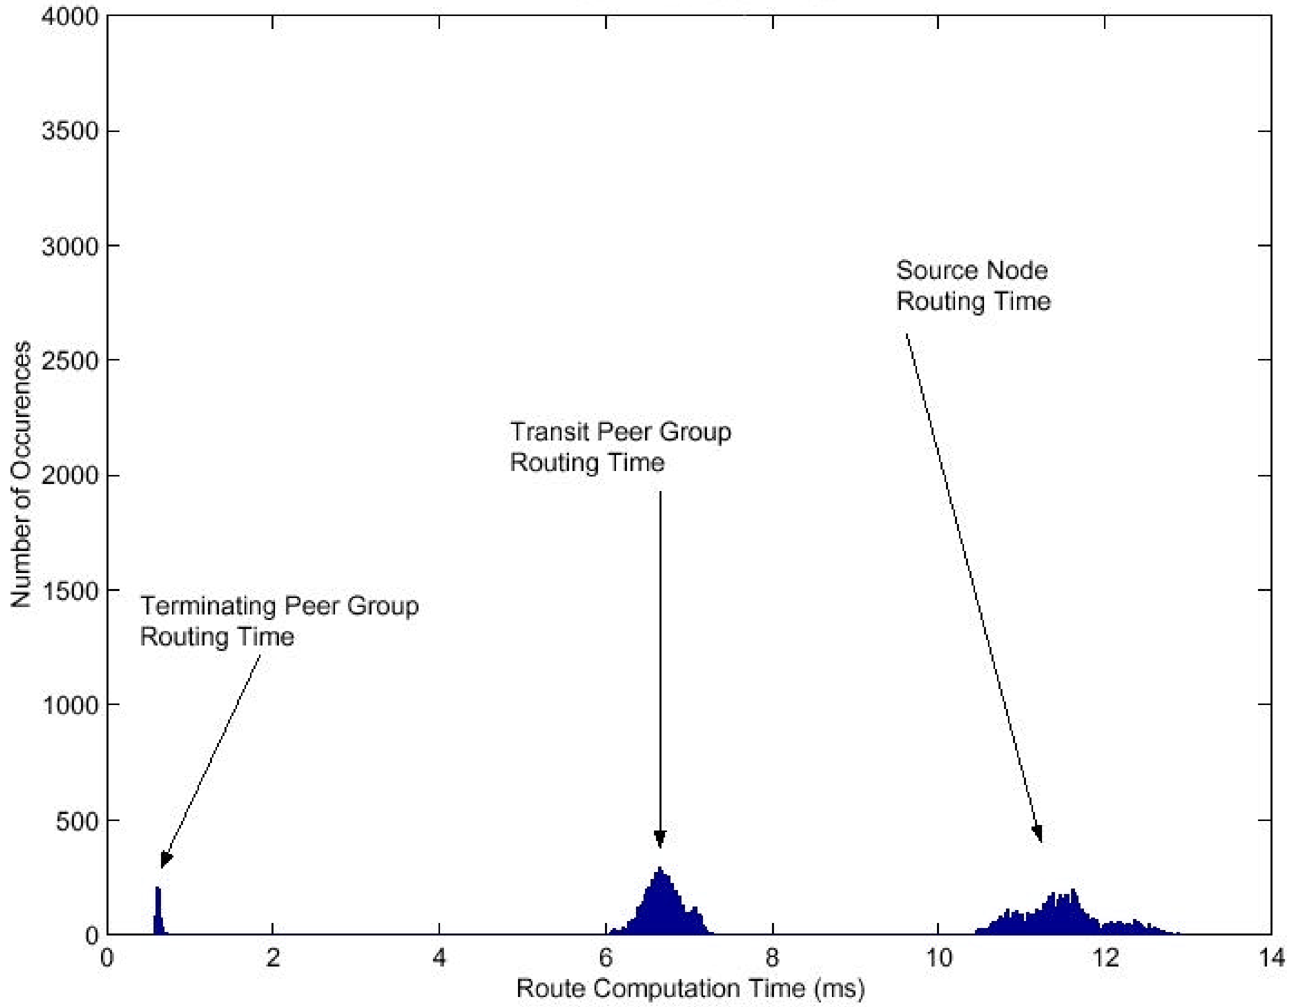

Pre-Computed Processing Times

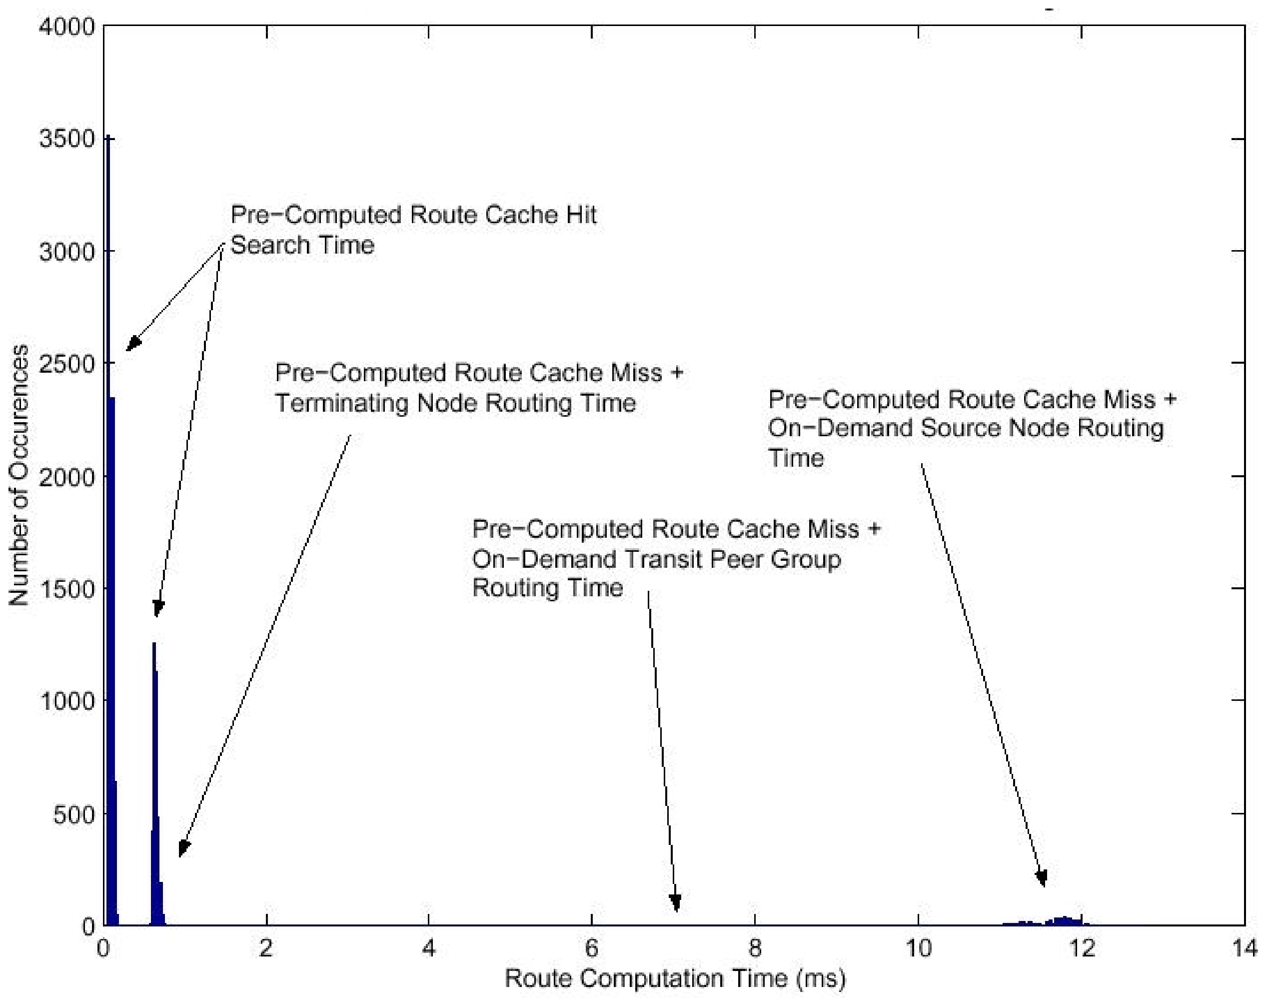

On-Demand Processing Time

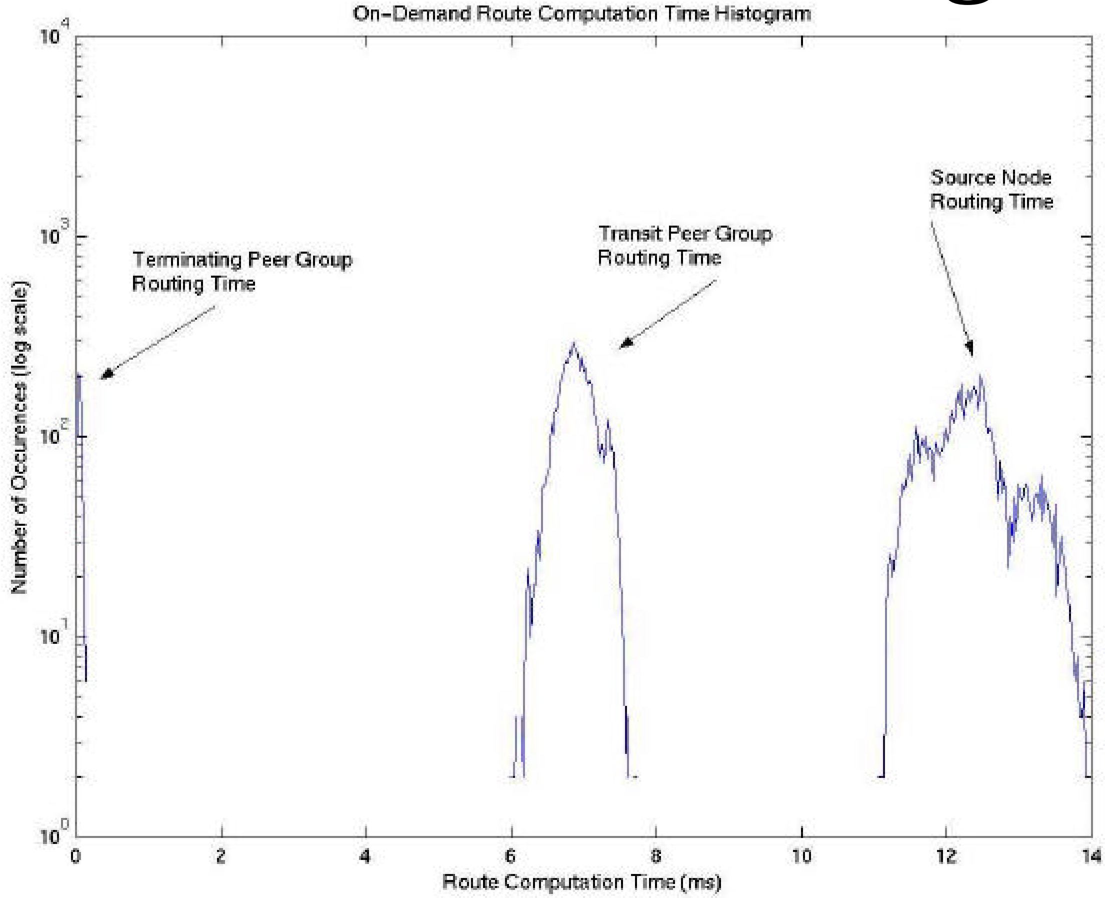

Pre-Computed Processing Time

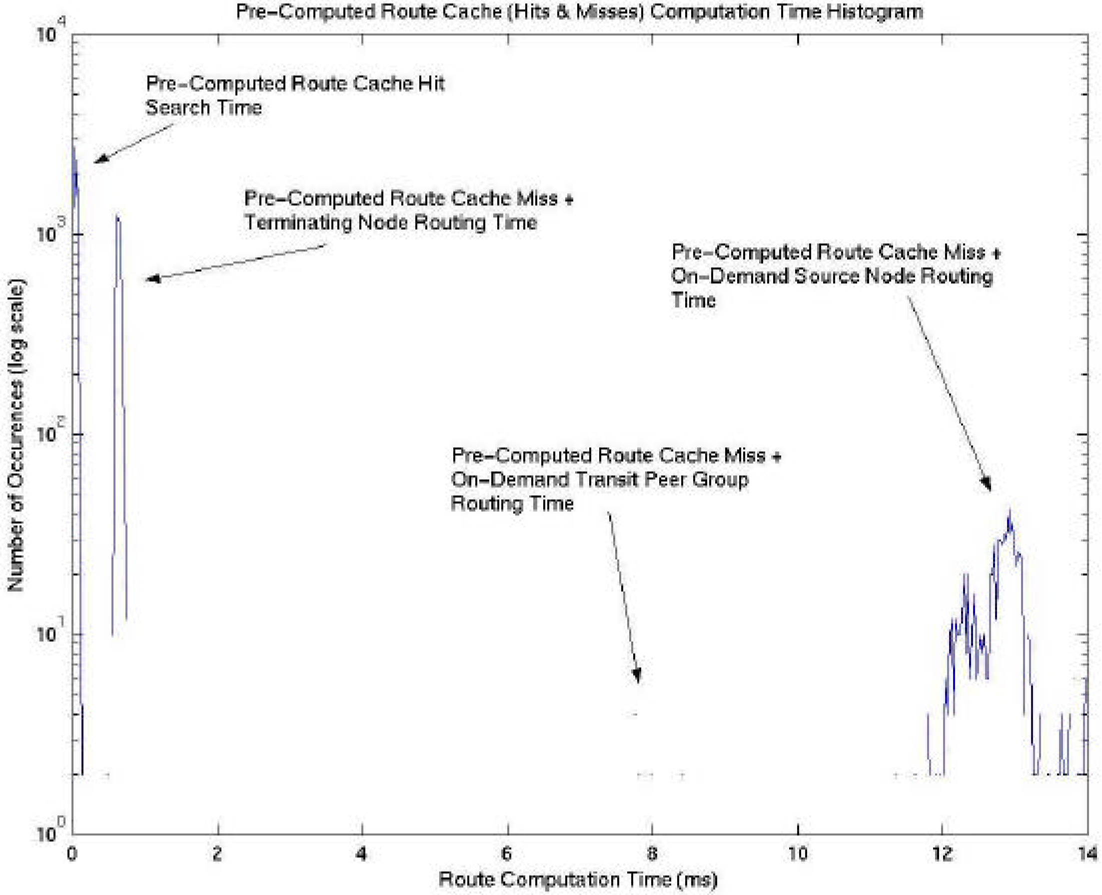

Conclusion

Lessons Learned, Future Work, Q&A

Lessons Learned

- Pre-computed route caching is effective in reducing the average call setup time as compared to on-demand routing, even as the network scales.
- On-demand routing gives better call acceptance as compared to pre-computed routing, for increasing traffic load.
- Route cache updates helps in increasing the call acceptance ratio of pre-computed routes, while achieving lower call setup times.

Lessons Learned

- In Multiple Peer Group topologies, periodic timer based route cache updates are most effective when synchronized with the re-aggregation timer.
- PTSE count heuristic is an effective means of tracking network state.
- Combining ptse count and crankback initiated invalidation shows best overall performance among the chosen heuristics.

Future Work

- Route Cache Policy enhancements:
  - Efficient route pre-computation.
  - Better cache data structures for storage of routes.
  - Check relevance of a chosen cache entry before it is used for connection establishment.
  - Investigating effects of crankback retries and alternate routing.

Future Work (cont'd)

- Hybrid Routing algorithm:
  - on-demand computations to first determine the foreign peer groups to traverse in a route,
  - then consult a route cache for crossing the source peer group to reach the first border node.
- Use of PTSE count to indicate when to carry out re-aggregation within a peer group.

Q&A

Thank You!!