Delayed Internet Routing Convergence

David Reddick

Outline

- Background Information
- Importance
- Methodology
- Important Terms
- Results
- Predictive Modeling
- Critique

Autonomous Systems (ASes)

- Collection of routers under a single entity control
- Comprised of internal and boarder routers
  - Include BGP and internal routing protocols

Border Gateway Protocol (BGP)

- Used for inter-domain routing
- Incremental protocol
- Updates whenever there is a change in a route

Border Gateway Protocol (BGP)

- Routing information shared
  - Announcements
  - Withdrawals
    - Explicit
    - Implicit
- Path selection
  - Majority based on ASPath length
  - Able to filter

Distance Vector Routing

- Each node must keep updated distance list
- Any change requires recalculation and sharing
- Slow process after change
- Vulnerable to routing loops
  - Count-to-Infinity

Importance

- Creates packet loss and increased latency
  - Packet loss by factor of 30
  - Latency increase by factor of 4
- Failover effects averaged 3 minutes but could last 30 minutes
  - Significantly slower then PSTN telephone networks
- Convergence greatly effected by number of autonomous systems

Methodology

- Over course of two years
- Geographically and topologically diverse
- Multiple different ISPs

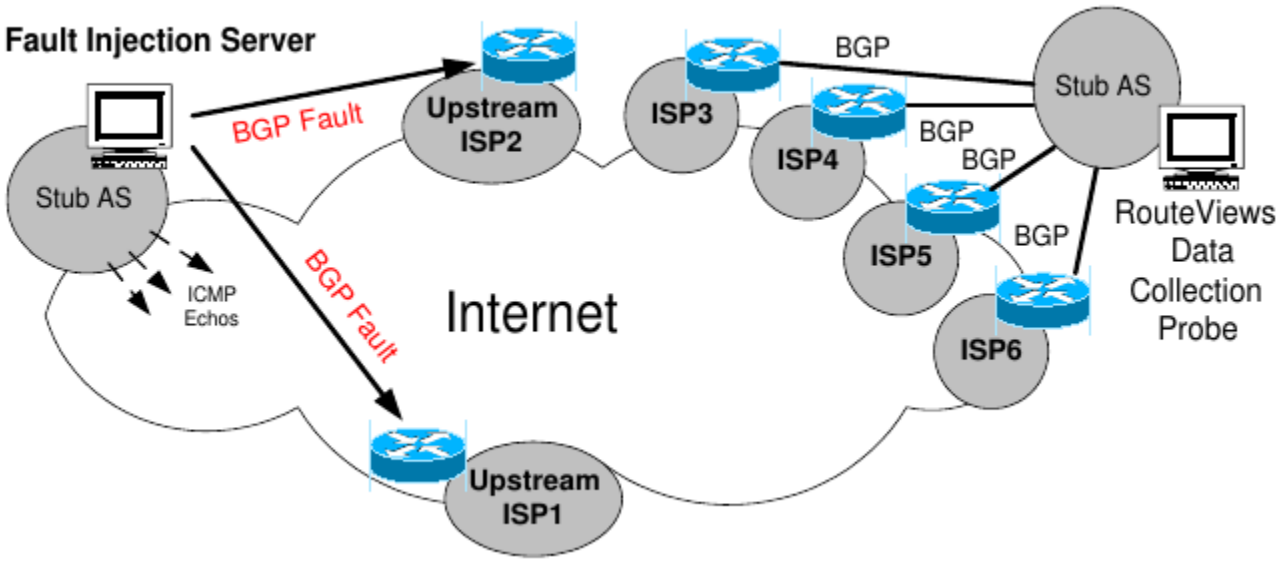

Methodology

- Simulated route failures, repairs, and failover
- 100 randomly selected websites to test against
- Created artificially long path
  - To test that only selected if primary failed

Experimental Terms

- Tup
- Tdown
- Tshort
- Tlong

Routing Results

- Latency

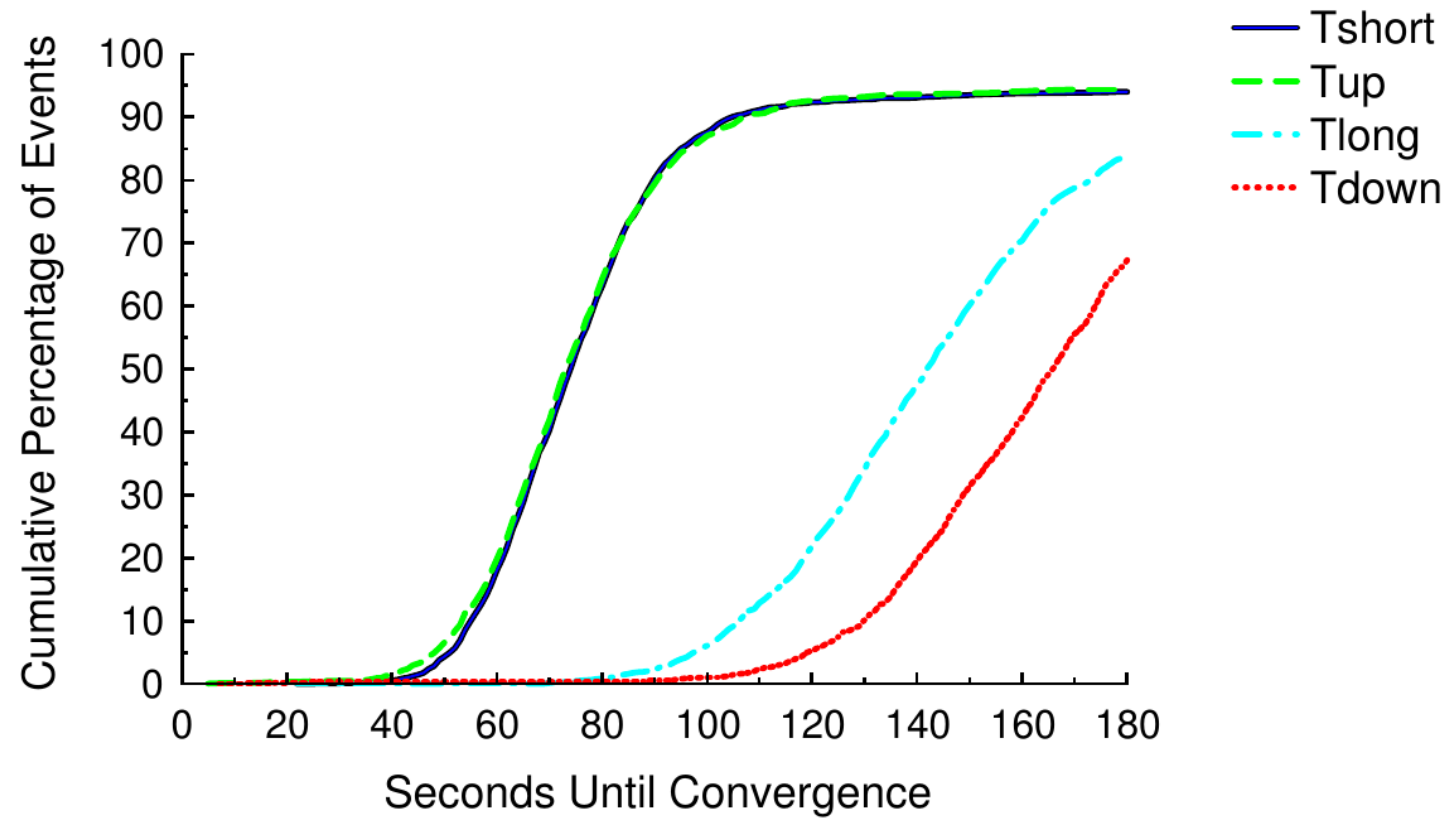

Routing Results

- Updates

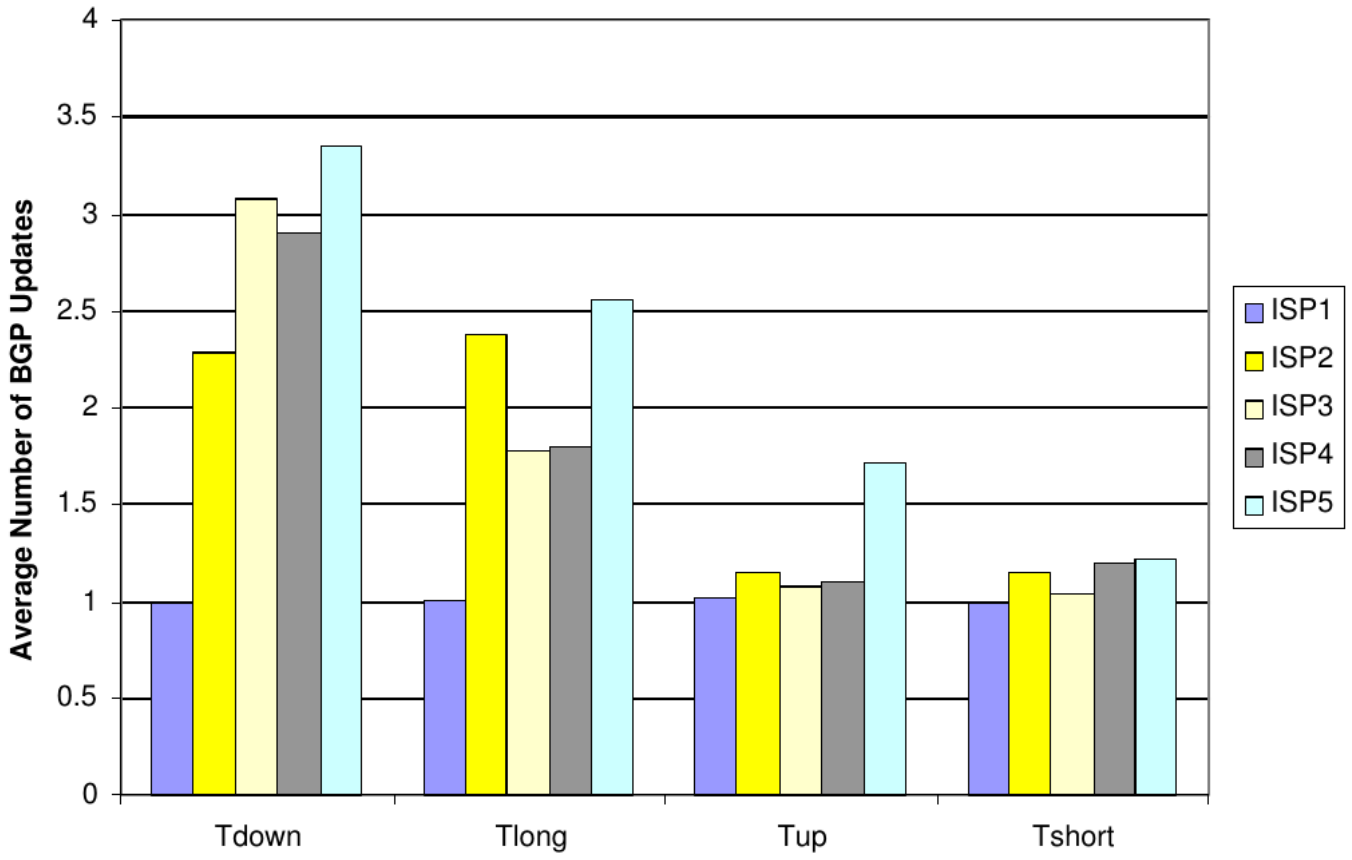

Routing Results

- Tdown

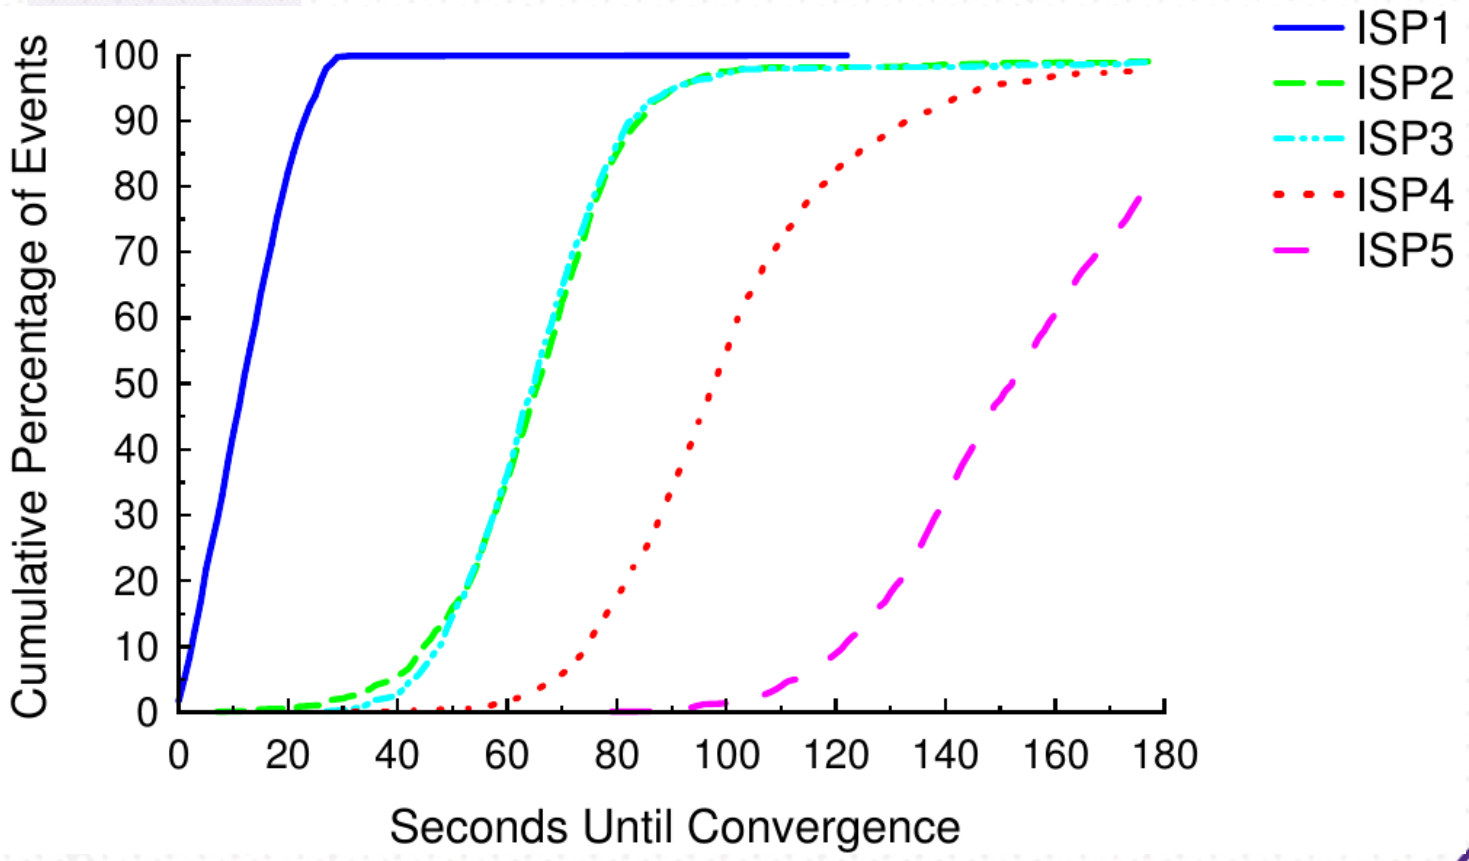

Routing Results

- Tup

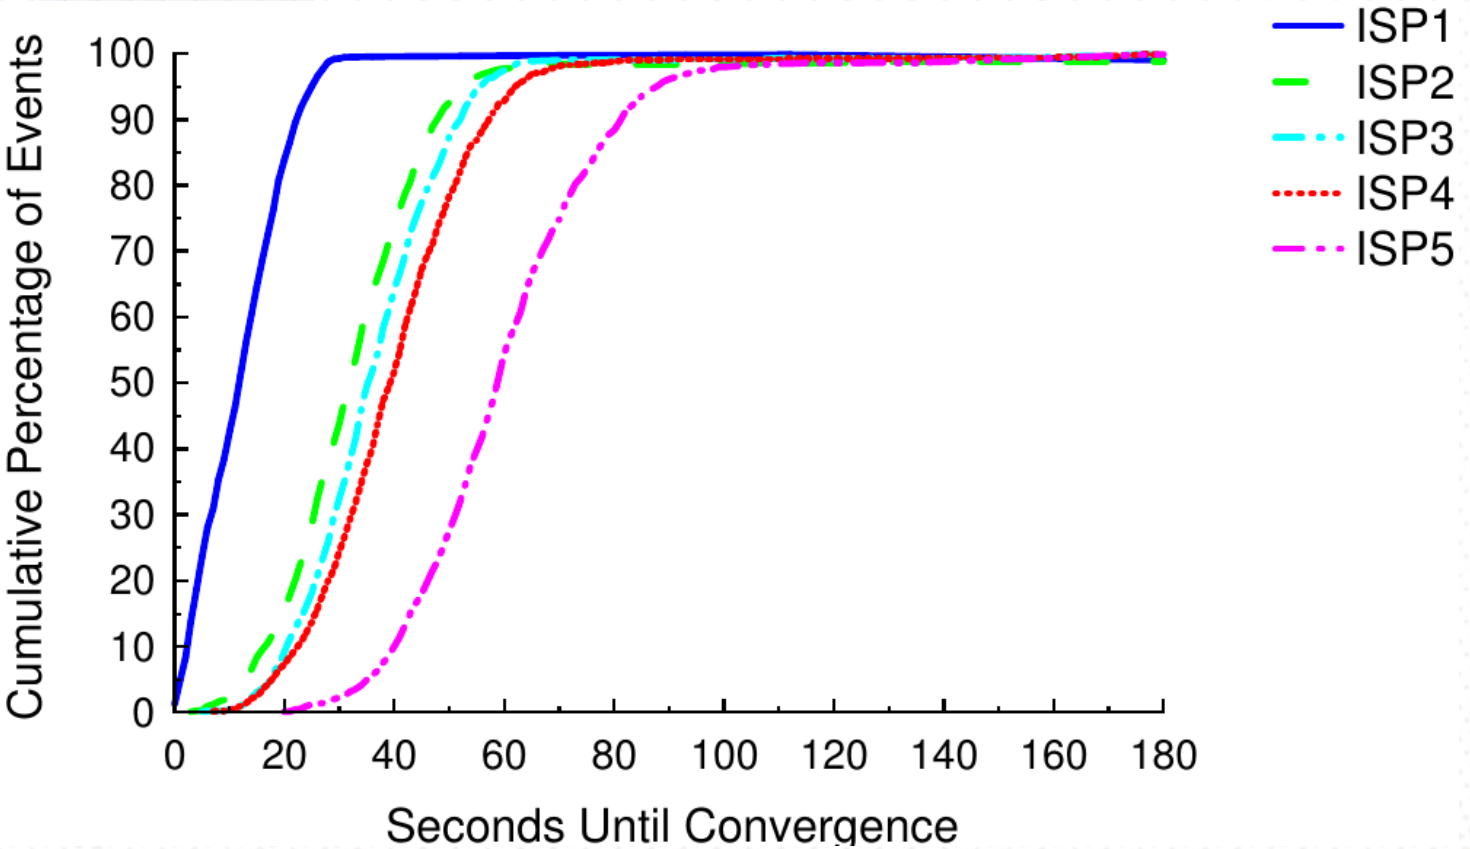

End-to-End Results

- Show even moderate levels of routing table oscillation leads to problems
  - Packet Loss
  - Latency
  - Mis Ordered Packets

End-to-End Results

- Loss

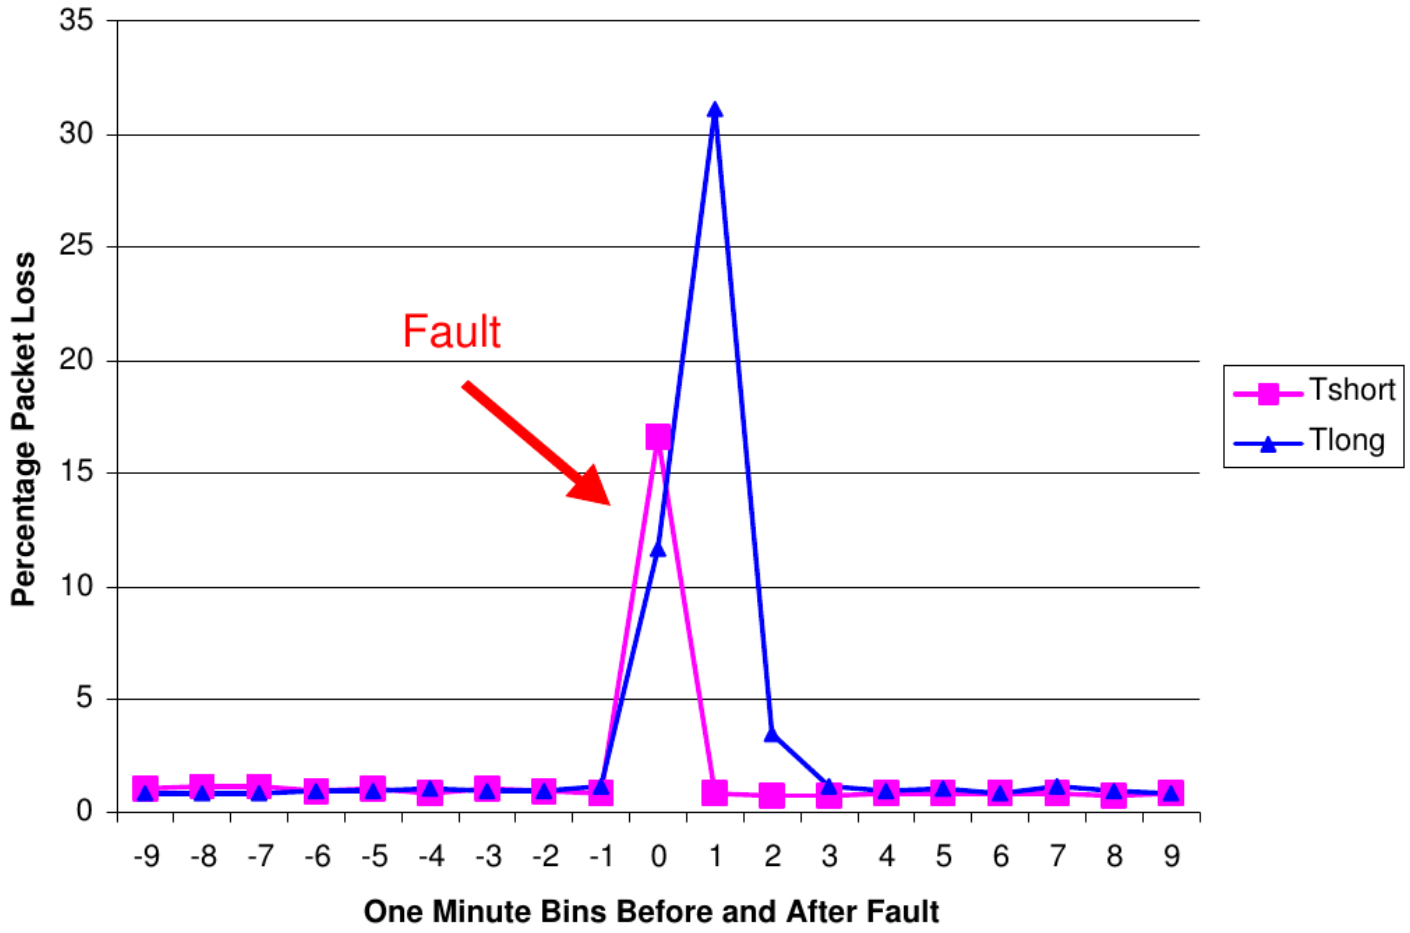

End-to-End Results

- Latency

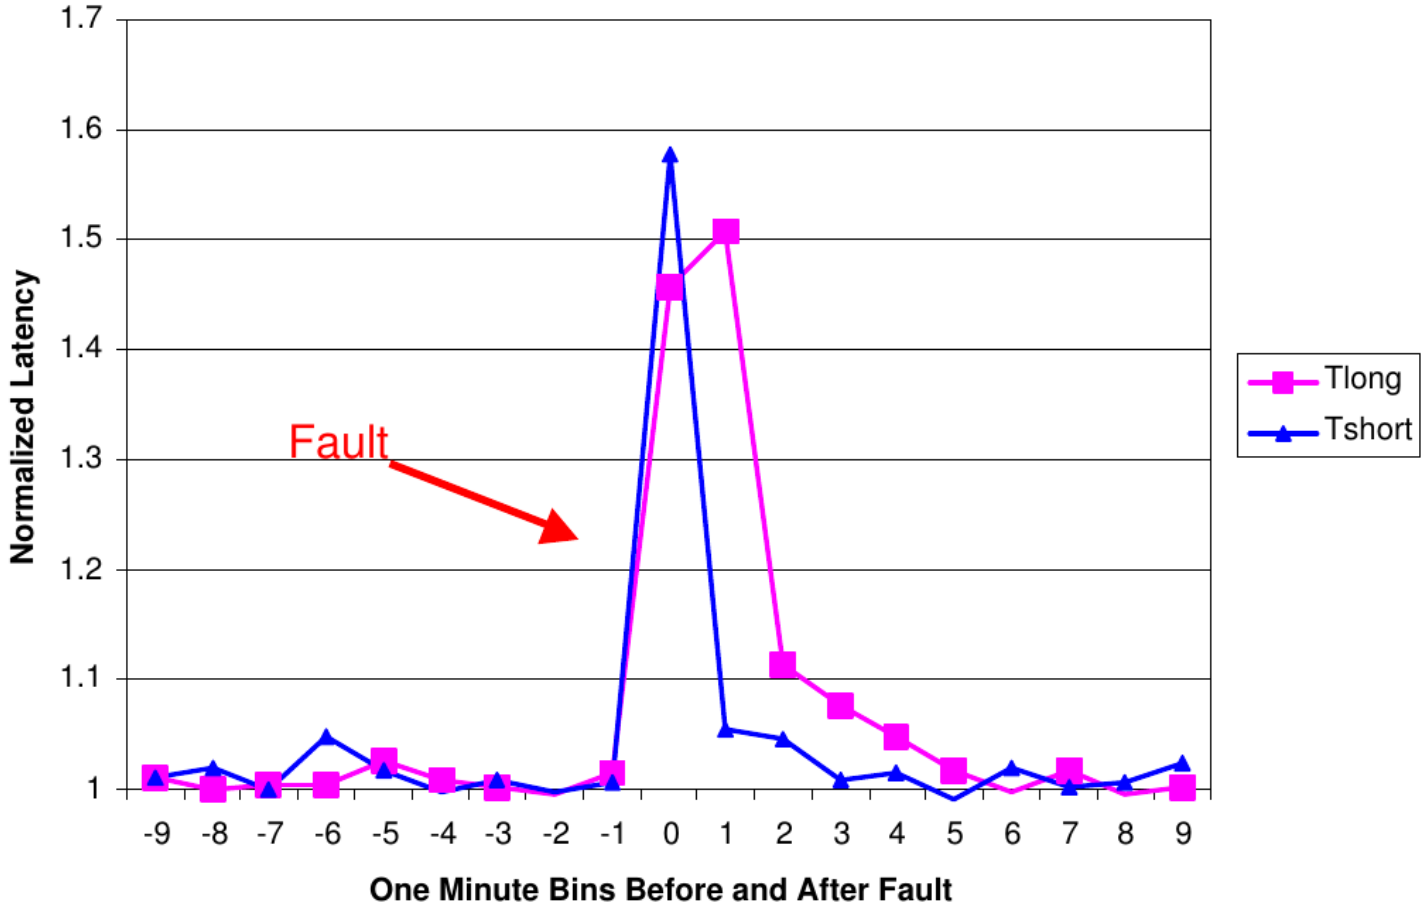

BGP Convergence Model

- Simplified model
  - Full mesh
- Compute theoretical upper and lower bound of BGP convergence
- Use path vector to solve looping problem

Upper Bound Convergence

- Used to calculate worst case convergence
- Assumes when a path is withdrawn, algorithm attempts to find equal or greater length path
- The algorithm will search till all possible paths have been discovered
- Worst case, results in O((n-1)!)
  - However unlikely to ever occur in a real environment

Lower Bound Convergence

- Used to calculate best case convergence
- Assumes bounded message delay
- By using MinRouteAdver timer, able to rapidly converge despite failure
- Best case, results in O(n)

Critique

- Path vector protocol
- Makes the assumption that everything needs reliability
  - TCP is able to add reliability without changing network
- Some of the calculations seemed unnecessary for the paper