On The Correlation between Route Dynamics and Routing Loops

Ashwin Sridharan and Sue. B. Moon
and Christophe Diot

Problem Statement

- Identify possible causes of routing loops within the Sprint backbone.
  - Methodology to correlate loops detected in traffic traces with routing events.
  - Any dominant cause(s) ?
  - Analyze impact of routing events on loop characteristics.

Talk Layout

- Routing Loops
  - Classification, causal sources.
- Methodology
  - Collection of data
  - Detection of loops and correlation with events.
- Analysis of data
  - Contribution of various protocols to loop creation.
  - Effectiveness of detection technique.
  - Effect of updates on path length distribution.
- Conclusions

Routing Loops

- Finite speed of propagation causes loops.
  - Routers change state in reaction to event.
  - After update, they broadcast new state.
  - Routing protocols have non-zero convergence time
  - BGP and ISIS routing protocols within Sprint.
- Can be classified based on cause/duration.
  - **Transient:** occur in normal state of operation.
  - **Persistent:** typically associable to anomalies.

An ISIS Loop

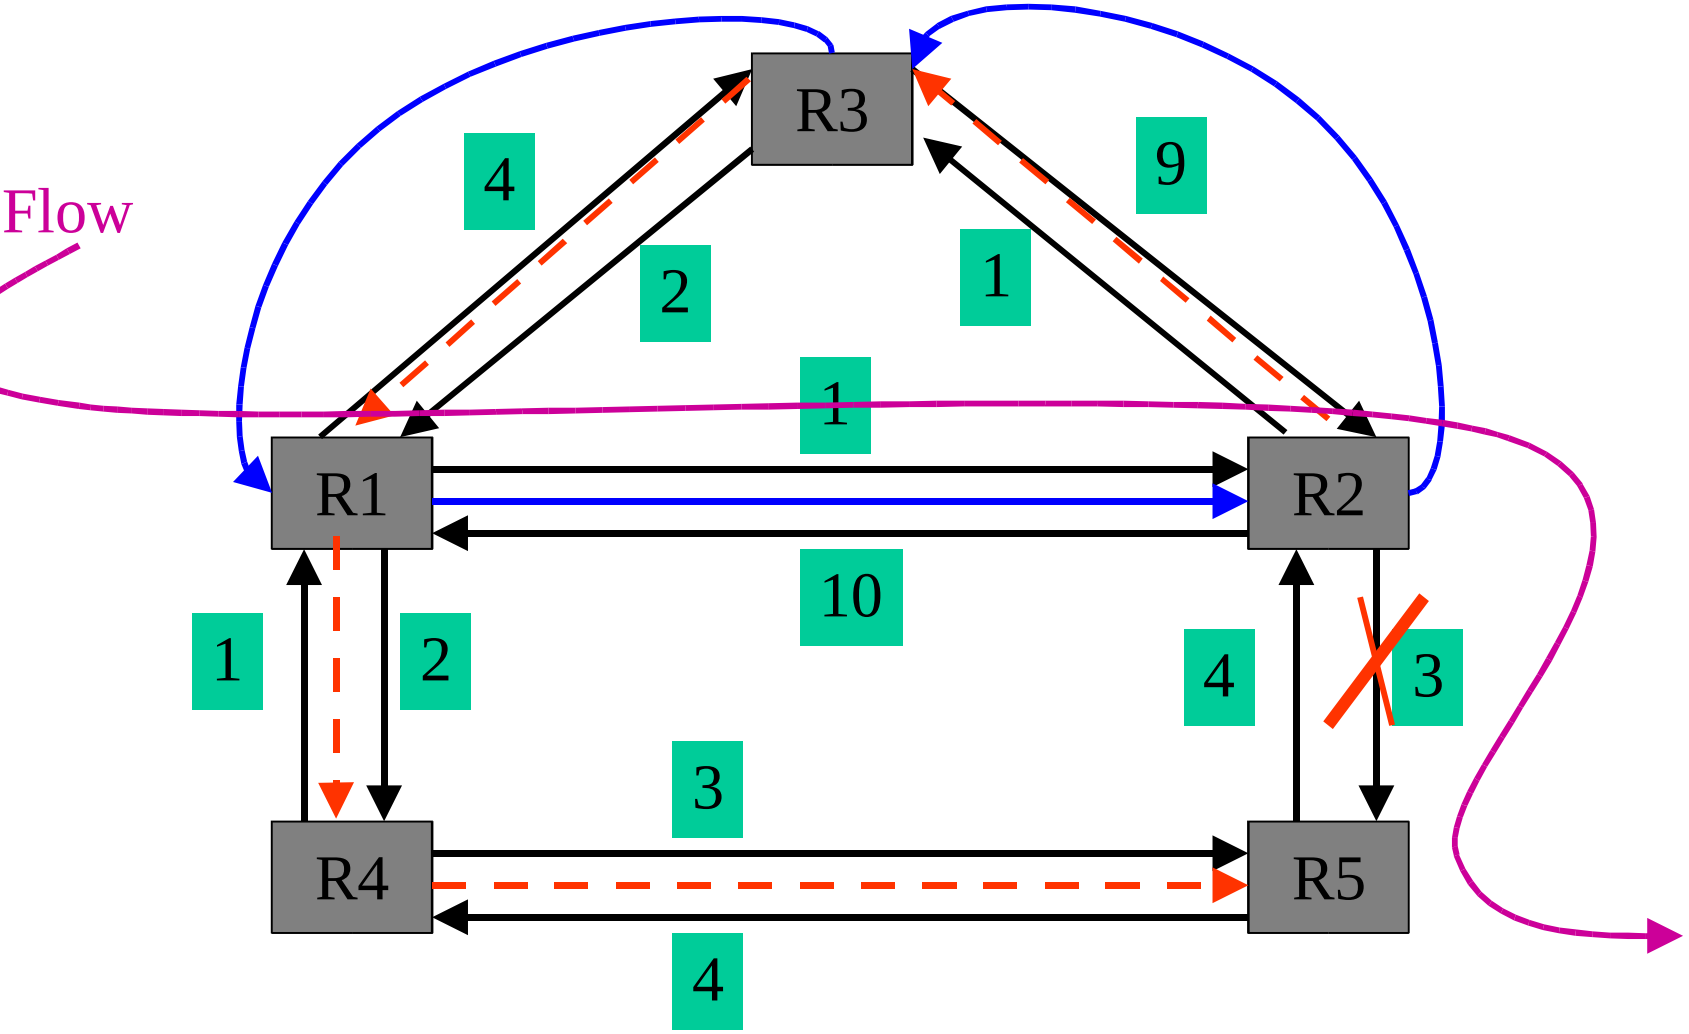

A BGP Loop

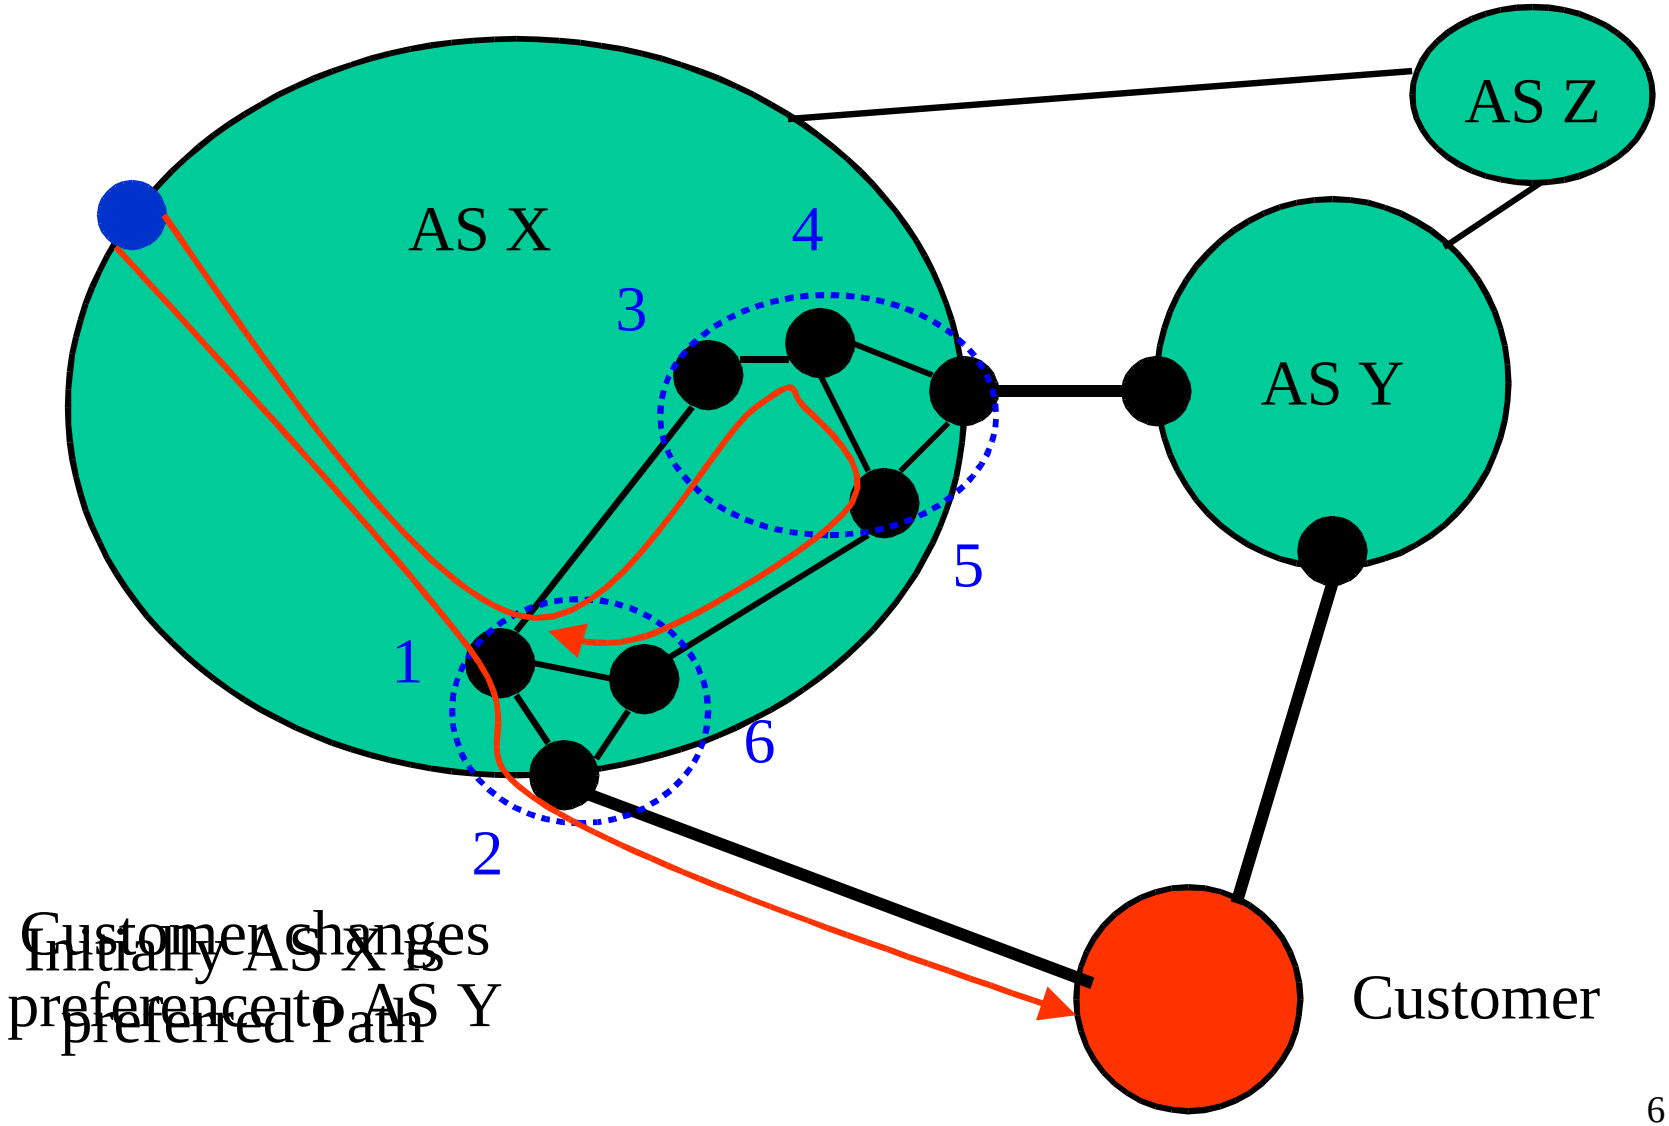
AS Z
AS X
4
3
AS Y
5
1
6
2
Customer changes preference to AS Y
Initially AS X B preferred Path
Customer

Methodology

- Collection of data.
  - Packet Traces.
  - Routing traces.
- Detection of packet loops in traces.
  - [Hengartner et al.]
- Correlation of packet loops with routing events.
  - Correlation with BGP events.
  - Correlation with ISIS events.

Collection of Data

- Collected OC-48 traces from 6 backbone links using Sprint IPMON equipment.
  - Dumps first 44 bytes from each packet.
  - Timestamps packet using GPS.
- BGP updates collected via Zebra BGP daemon peering with a BGP router.
- Pyrt ISIS routing daemon creates adjacency with an ISIS router and collects LSPs.

Detecting Packet Loops

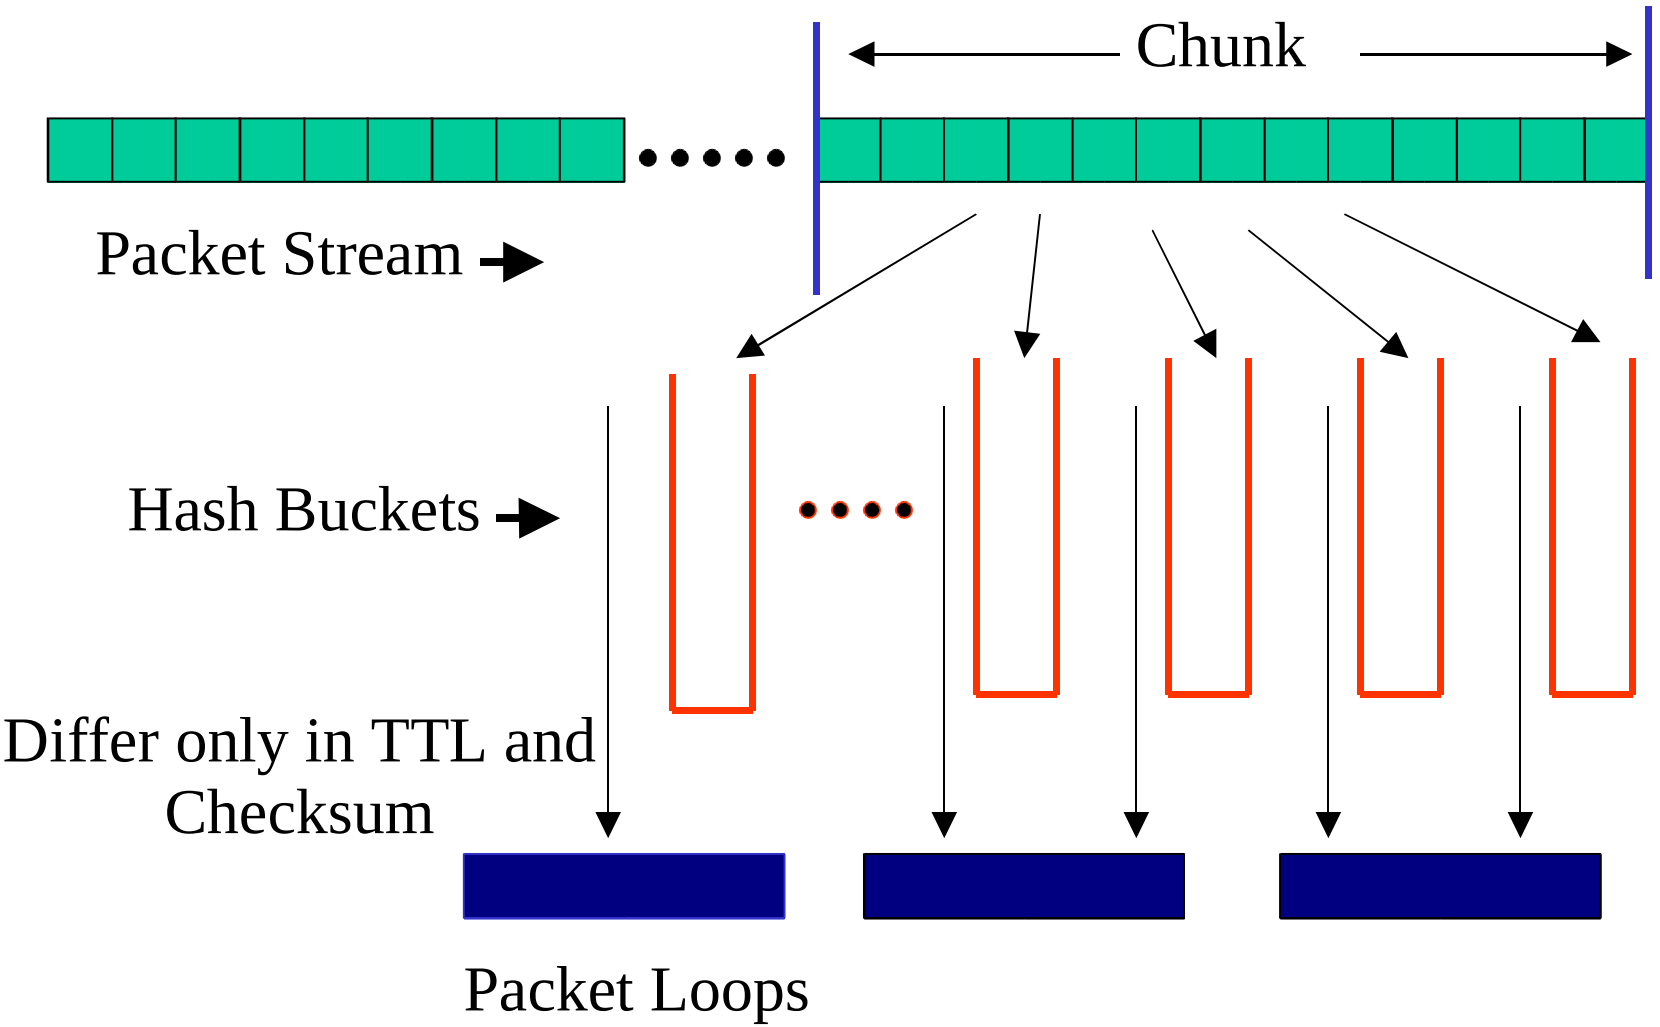

Correlating packet loops and BGP Events

- Feed BGP updates to a Zebra router emulating the BGP decision process.
- For each BGP update
  - Determine changes in next-hop or AS Path for any loop.
  - If change in vicinity of loop origin, assume event responsible for loop.

Correlating packet loops and ISIS Events

- After each LSP is received, compute shortest path from observation node to all destinations.
- For each packet loop
  - Determine any change in forwarding path.
  - Determine if it overlaps with previous path.
  - If event in vicinity of loop, assume event was causal in the creation of the loop.

Analysis of Data

- Do both protocols cause routing loops ?
  - All loops in traces associable only with BGP updates.
- Link state protocols have fast convergence time.
- Extensive use of multiple equal cost paths prevents overlap of ISIS forwarding path.
  - Monitored links were inter-POP links.

Analysis of Data – (2)

- How effective is the detection technique ?
  - Affected by “distance” of source from observation point.
  - Updates related to events in other Ases may get filtered out.

Matching Efficiency

Trace	% Transient & BGP Updates	% Persistent & BGP Updates	% Persistent & no Updates	Total
NYC-20	40.1	0	50.8	90.8
NYC-21	80.2	0	7.5	87.9
NYC-22	18.8	0	80.6	99.4
NYC-23	3.3	0	0	3.3
NYC-24	70.0	0	0	70.0
NYC-25	43.7	15.5	0	59.2

Average AS Path Length

Trace	Avg. AS Path Length
NYC-20	1.34
NYC-21	1.04
NYC-22	0.51
NYC-23	1.74
NYC-24	1.61
NYC-25	1.63

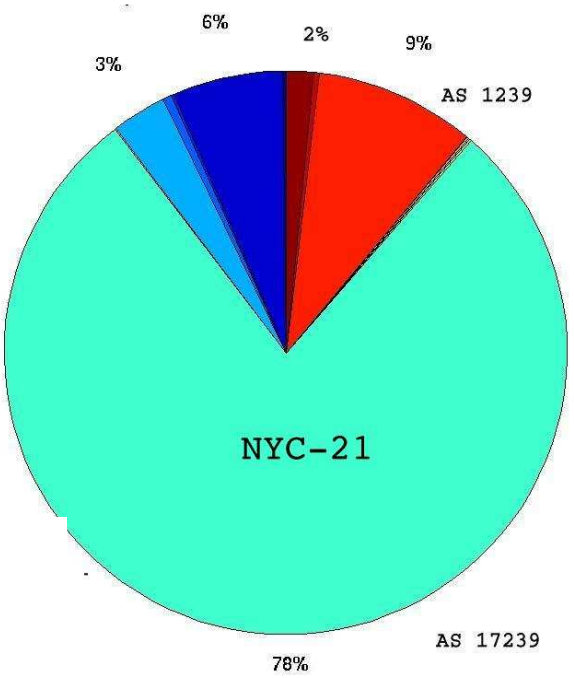
6%
2%
9%
3%
AS 1239
NYC-21
AS 17239
78%

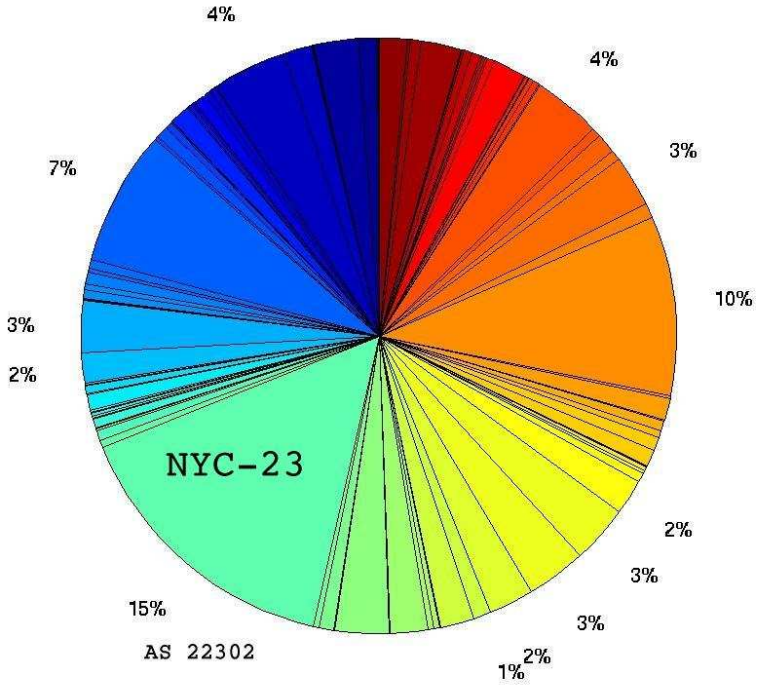
4%
4%
3%
7%
10%
3%
2%
NYC-23
2%
3%
3%
15%
AS 22302
1%
2%

Impact of BGP updates on loop length

- Path Length defined as the number of hops in a loop.
- Relationship between path length distribution and BGP updates.
  - If updates impacts large set of destinations, more likely that path length distribution has a higher variance.

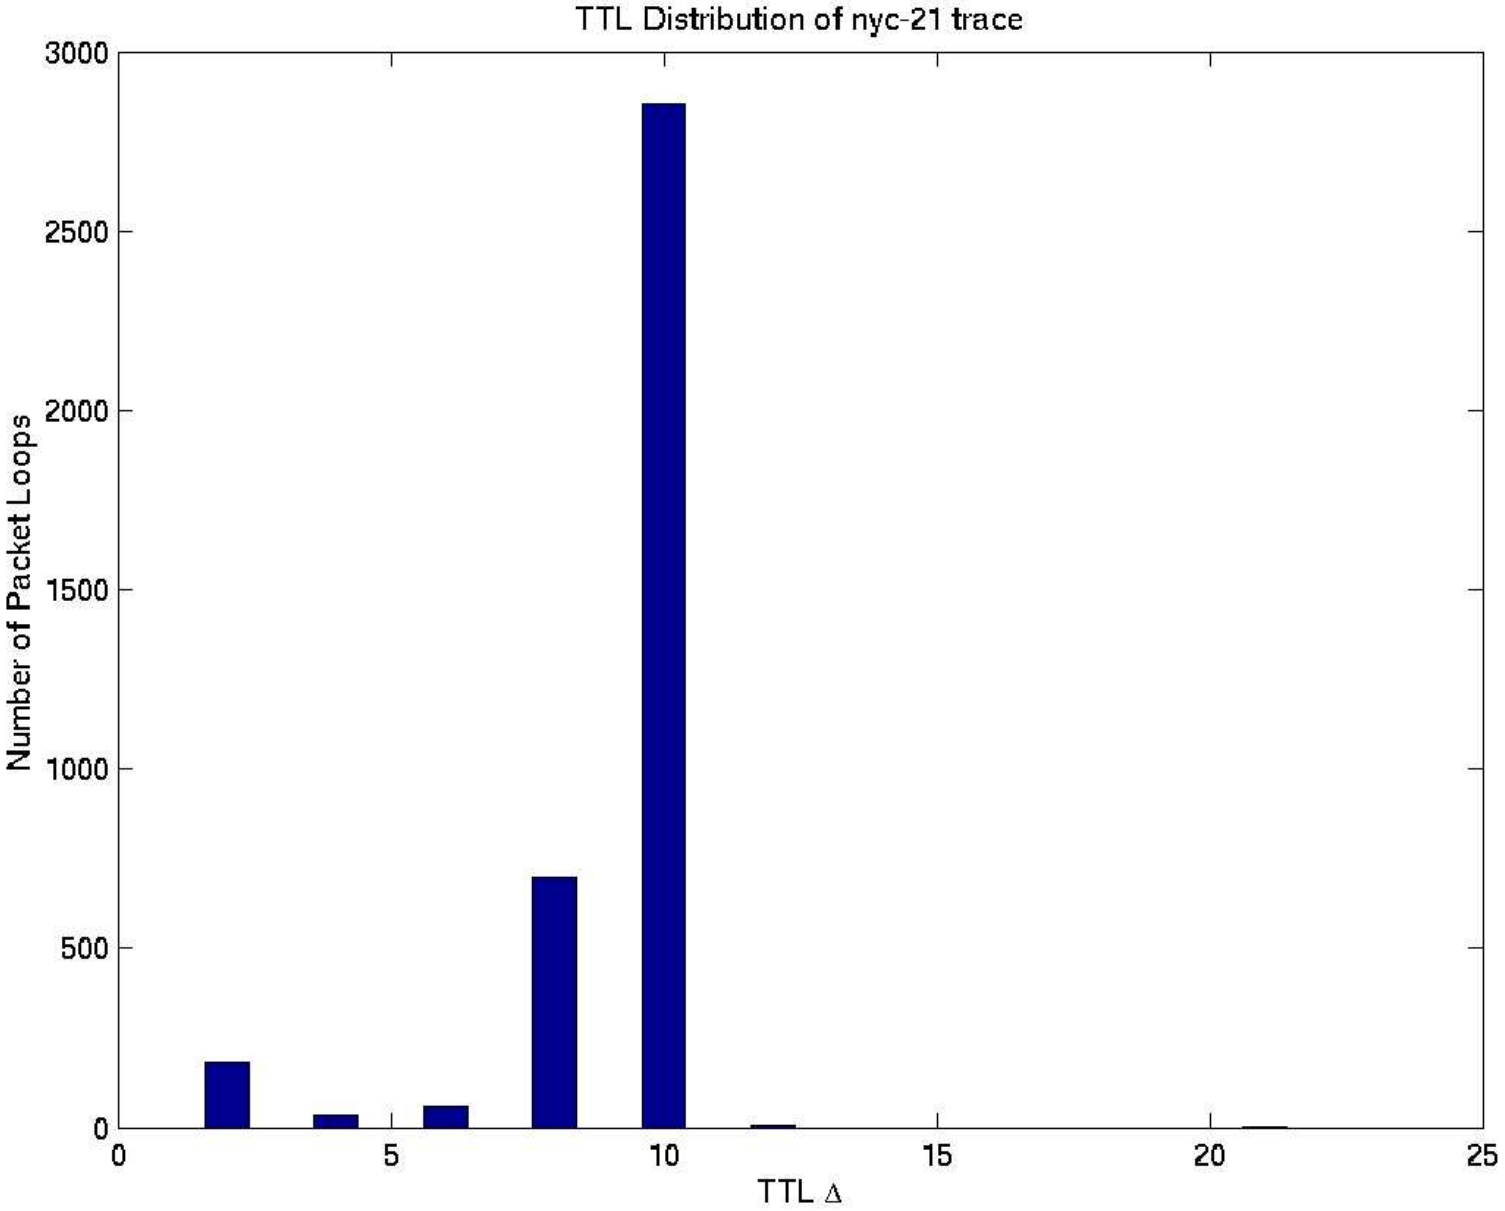
TTL Distribution of nyc-21 trace
Number of Packet Loops
TTL Δ
0
500
1000
1500
2000
2500
3000
0
5
10
15
20
25

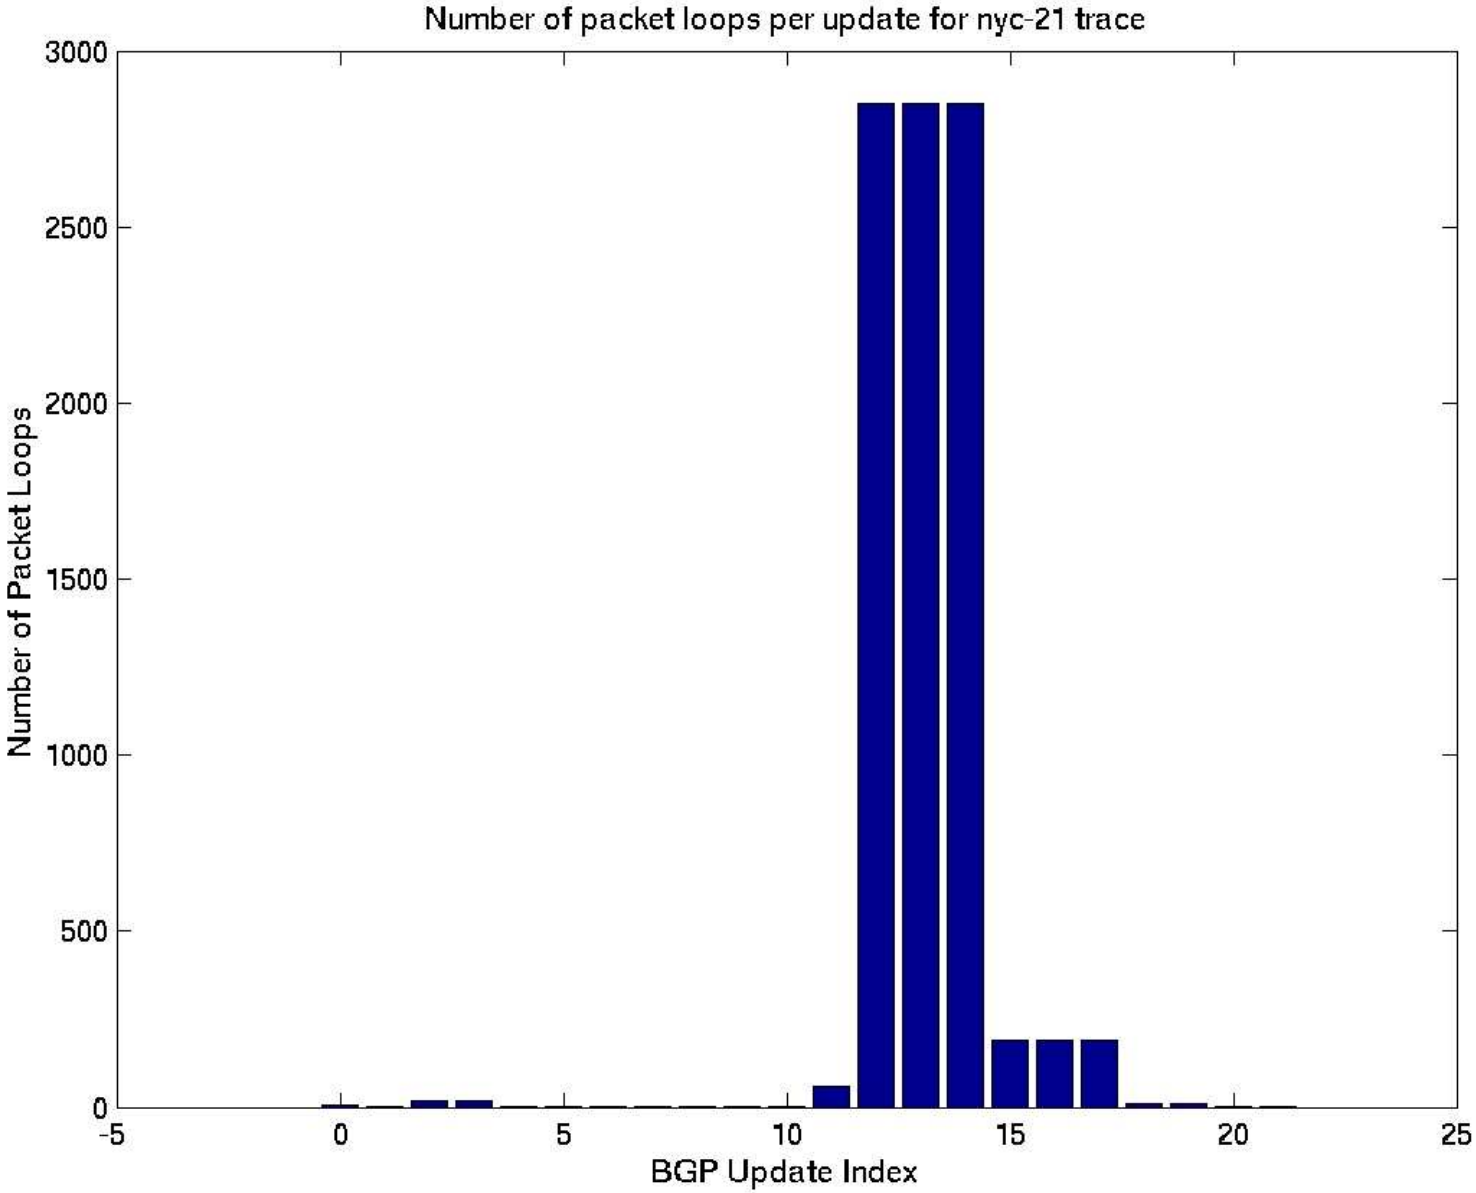
Number of packet loops per update for nyc-21 trace
Number of Packet Loops
BGP Update Index
3000
2500
2000
1500
1000
500
0
-5
0
5
10
15
20
25

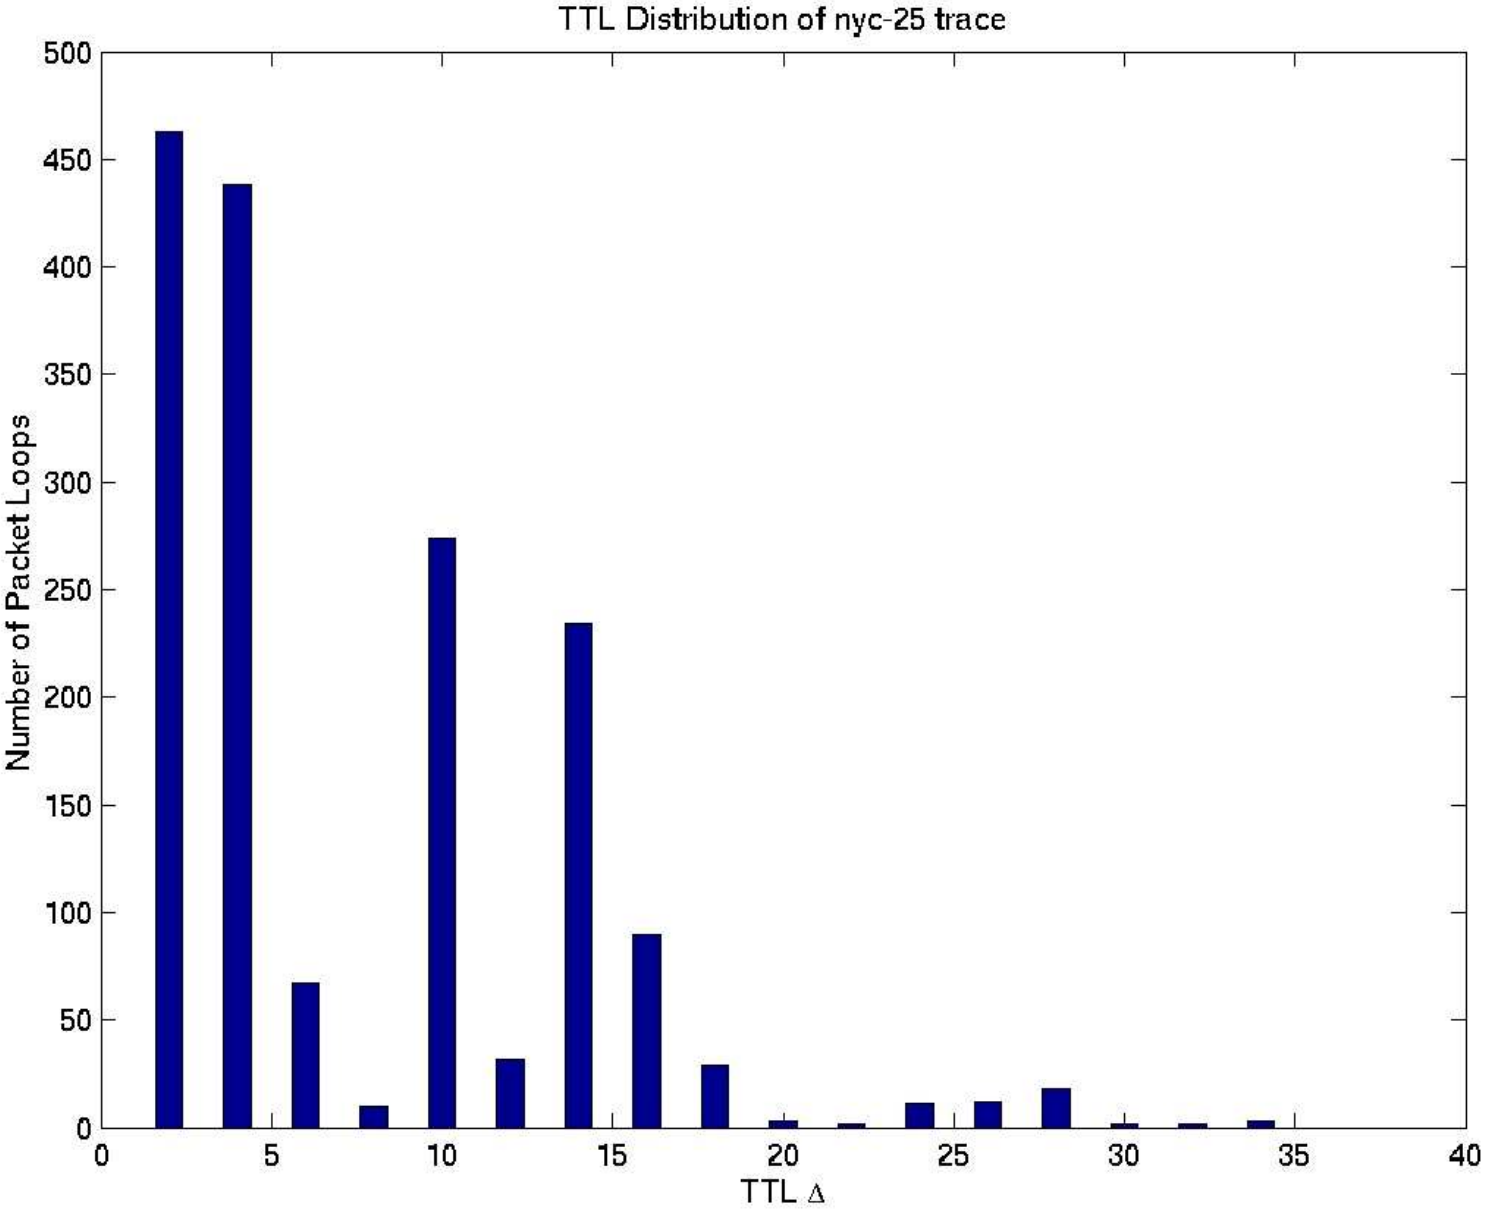
TTL Distribution of nyc-25 trace
Number of Packet Loops
TTL Δ
0
50
100
150
200
250
300
350
400
450
500
0
5
10
15
20
25
30
35
40

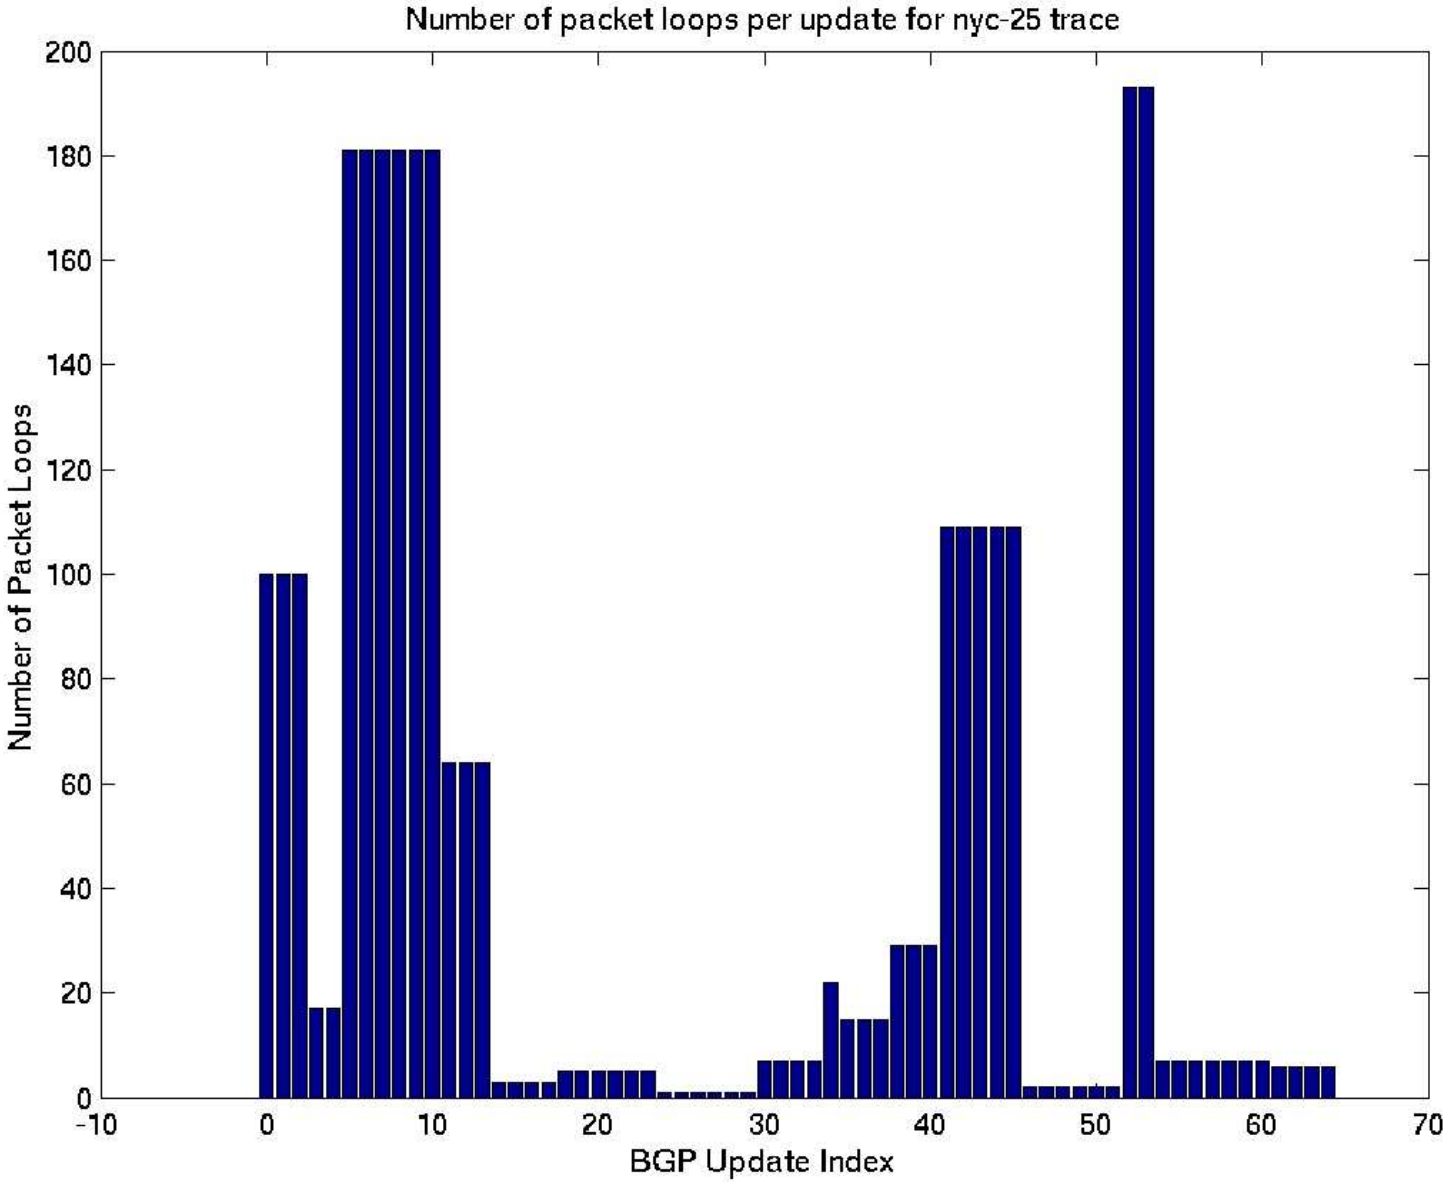
Number of packet loops per update for nyc-25 trace
Number of Packet Loops
BGP Update Index
200
180
160
140
120
100
80
60
40
20
0
-10
0
10
20
30
40
50
60
70

Conclusions

- Methodology to correlate routing events with packet loops.
- BGP updates were almost exclusively responsible for routing loops.
- No loop creation event directly associable with ISIS.
  - Attributable to equal cost multiple paths.
- Correlation between BGP updates and path length distribution.