Stability Issues in OSPF Routing

Anindya Basu

Jon G Riecke

SIGCOMM'01

Srikant Sharma

Goal

- Study OSPF stability
- Stability indicators
  - Network Convergence Times
  - Routing Load on Processors
  - Route Flaps

Goal…

- Scenarios
  - OSPF with TE
  - Node failures
  - Message Loss
  - Trigger thresholds
  - Subsecond HELLO timers
  - Alternate strategies for refreshing link-state information

Overview

- The study is carried out in the form of simulations over a real network topology consisting of 292 nodes and 765 links
- Previous simulations usually incorporated around 10s of nodes
- Uses a very detailed processor model (not discussed here)

OSPF-TE

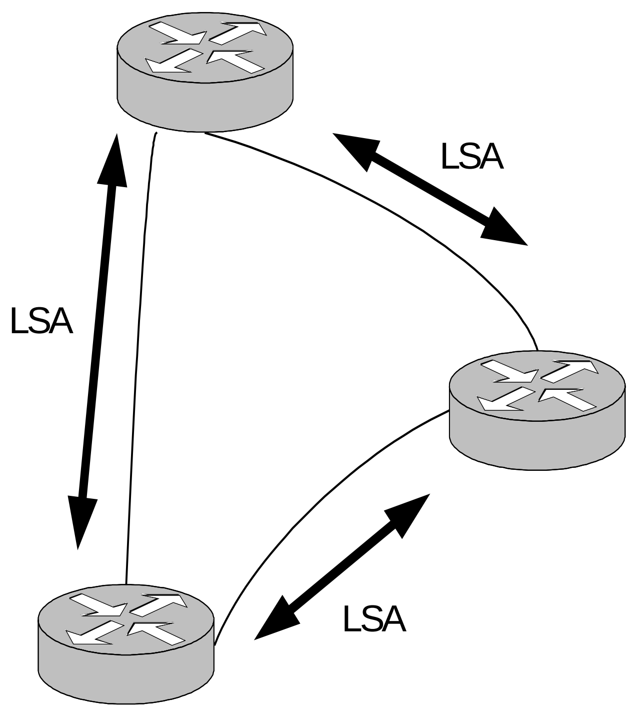

- Open Shortest Path First protocol with Traffic Engineering extensions
- Uses Opaque LSA to convey TE information
- TLV, triggering

Network Convergence Time

- Time taken to stabilize Routing changes
- Depends on the diameter of the network
- With TE there are additional Complexities
  - Extra time for propagating changes and re-reserving bandwidth

Routing Load and Route Flaps

- High load indicates instability
- Steady State
  - HELLO packets and periodic LSA refresh
- Under perturbations
  - Flooding of LSAs to update the status everywhere
- Route flaps are changes to routing tables: indicate the intensity of perturbation

Experiments with OSPF-TE

- Setup
  - 292 nodes, 765 links
  - 59 hub nodes and 233 POP spanning 77 cities
  - Hubs are connected using T3 to OC-12 links
  - Latencies varied from 10 microseconds to 113 millisecods

Node Failures

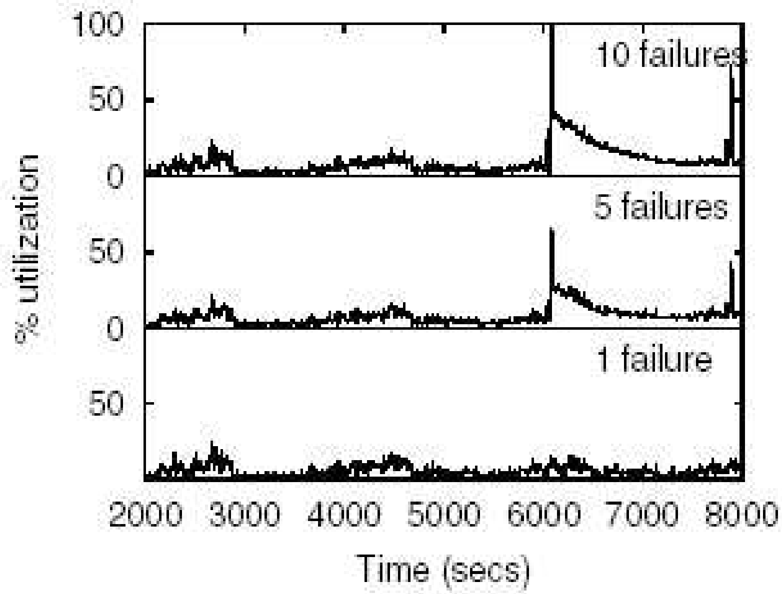

(a) Processor Utilization

- Processor utilization
  - No effect on single failures
  - Around 50% with 5 node failure
  - Around 100% with 10 node failure

Node Failures…

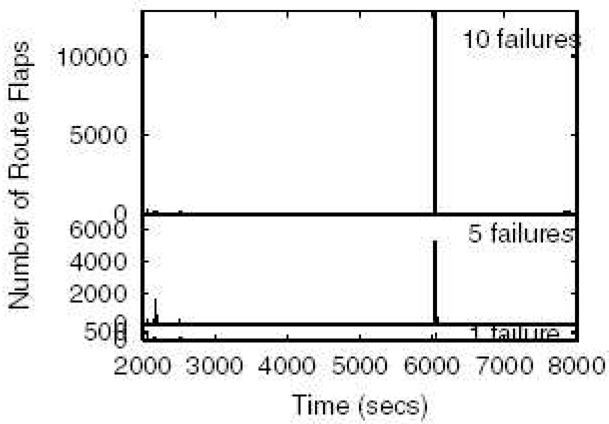

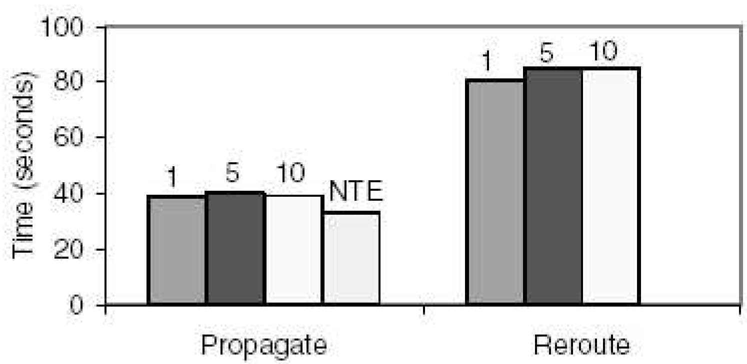

- Route Flaps
  - As little as 5 failures create problems
- Convergence times
  - OSPF is robust
- TE has 20% overhead

Message Loss

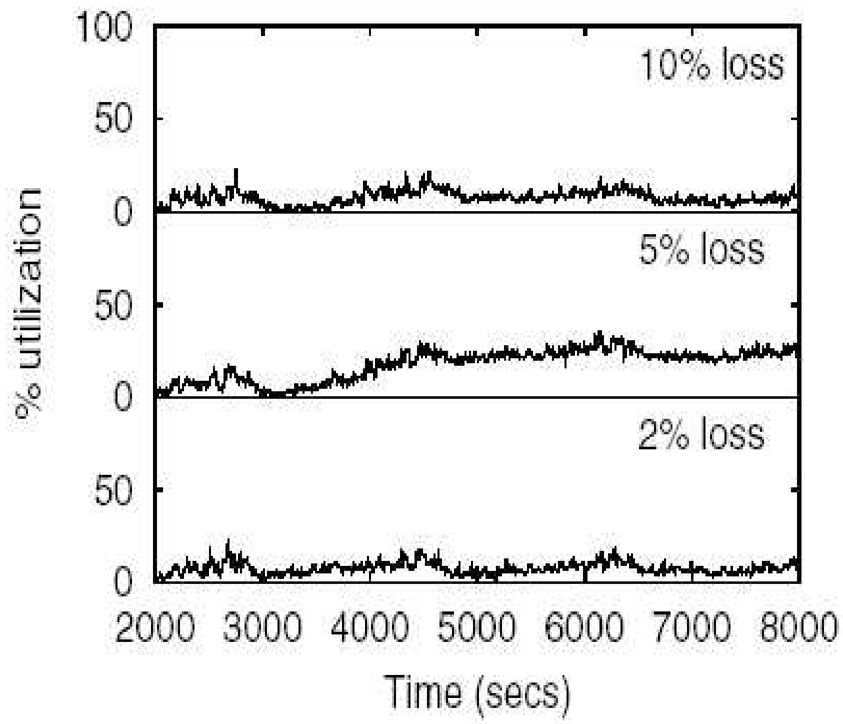

Processor Utilization

- 2% & 10% do not have much effect, reason being that OSPF thinks that a node has fewer neighbors and hence traffic goes down
- 5% utilization is affected

Message Loss…

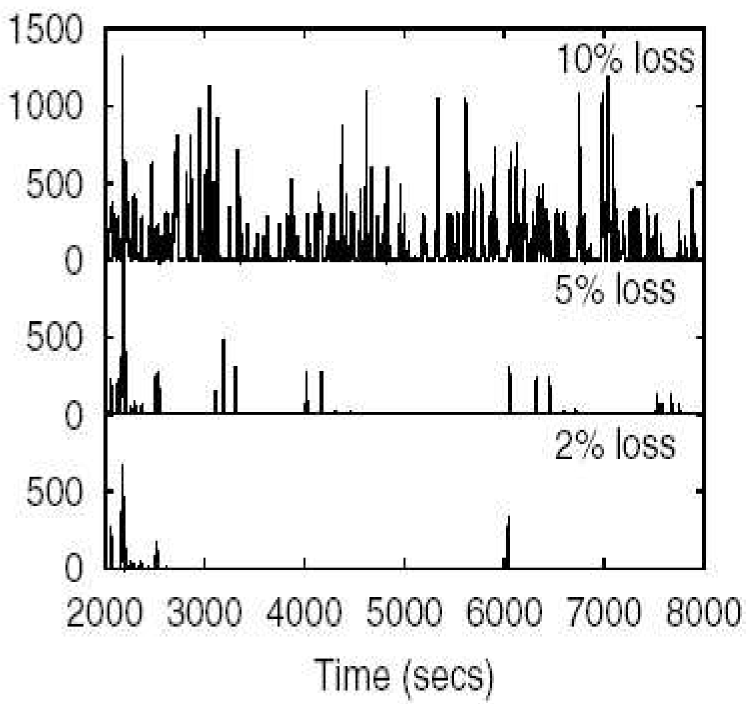

- Route Flaps
  - Not much flapping at 2% & 5%
  - Large amount of flapping at 10%

Convergence Times

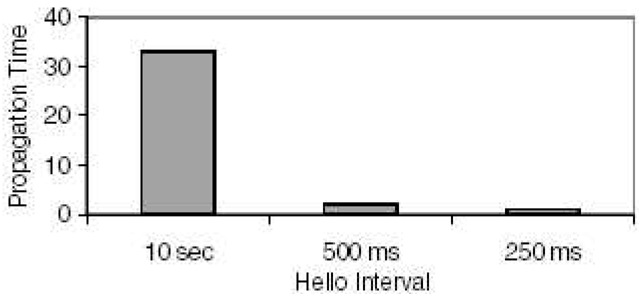

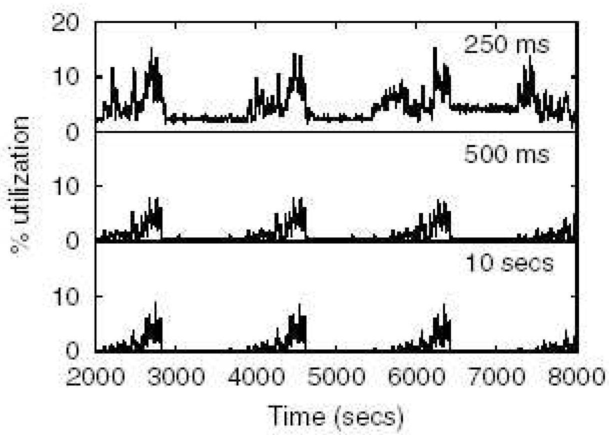

- Hello messages are usually sent every 10 seconds
- Propagation time given Hello timer of
  - 10 sec → 33 sec
  - 500ms → 3 sec
  - 250ms → 1 sec
- Processor load for subsecond Hellos is basically the same (as long as over 275 ms)
- Route Flaps are huge if <275ms

Refresh Synchronization

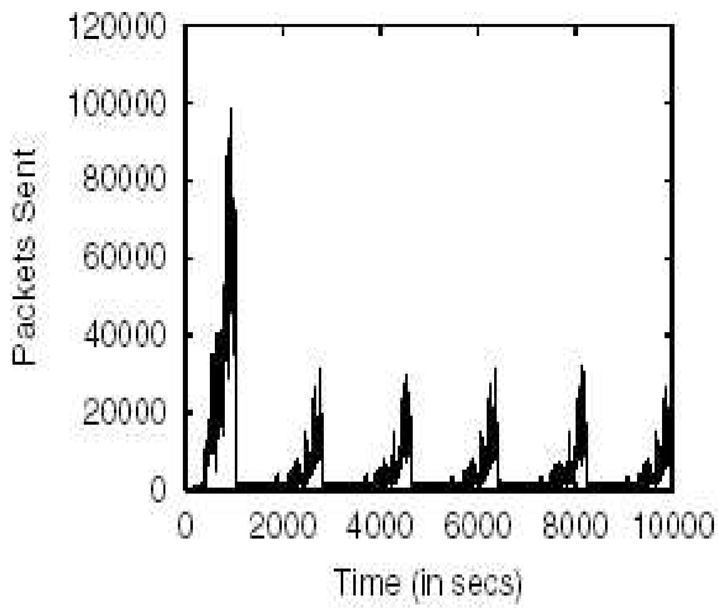

(a) Standard OSPF

- Based on the LSA timers, periodic bursts of traffic can develop
- Proposal:
  - randomization
  - Zinin mechanism

Refresh Synchronization

- Zinnin Proposal
  - Refresh messages are grouped and sent back to back
  - The groups are spaced apart in time
- Randomization
  - Refresh messages are sent at a random interval

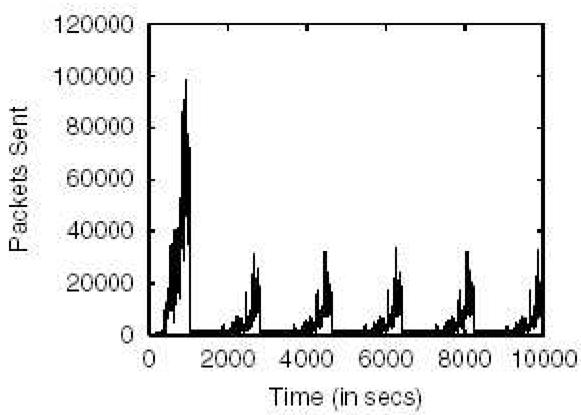

(b) Zinin

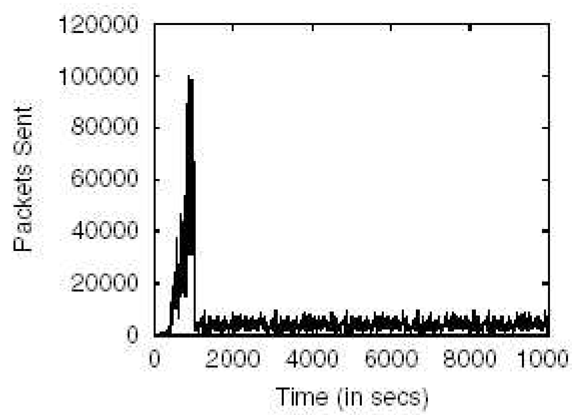

(c) Randomized

Conclusion

- OSPF-TE is robust
- TE has little effect on convergence time
- OSPF is resilient to message loss
- Setting subsecond HELLO timers is beneficial
- Randomization can prevent bursts of LSA traffic