Experimental Results on IS-IS Flooding

Sarah Chen

Tony Li

Two Drafts on IS-IS Flooding

- draft-ginsberg-lsr-isis-flooding-scale
  - Flow control using PSNP as an indicator
- draft-decraene-lsr-isis-flooding-speed
  - Feedback from RX to help flow control
- Purpose of this experiment:
  - How are we doing?
  - How do we improve?

Experiment Setup

- 4 Routers, p2p links
  - Bring R3-R4 adjacency down
  - Inject 2K LSPs to R4
  - Bring R3-R4 adjacency up
- Measures:
  - Number of LSP retransmissions
  - Flooding duration

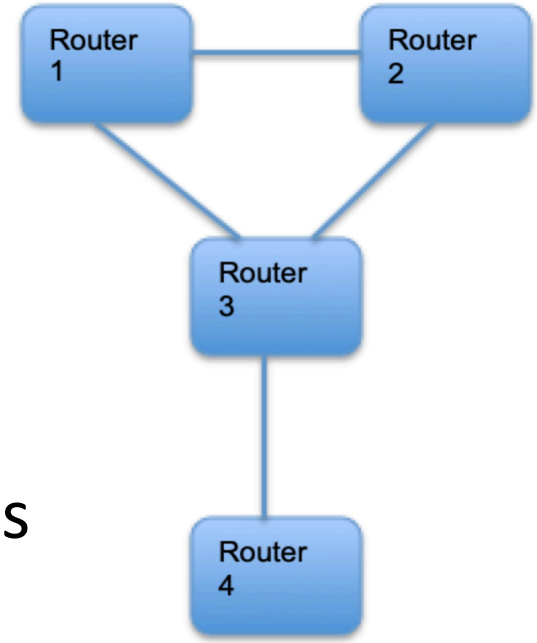

Initial Results

- LSP Tx interval: 33ms
  - Time between two LSP updates.
- PSNP interval: 2s
  - Time between two PSNPs.
- Result:
  - Flooding duration: 94.94s >> 33ms * 2K (66s)
  - Number of retransmissions: 803 >> 0

Reduce LSP Tx Interval

- Reducing LSP Tx Interval can speed up flooding

Tx Interval (ms)	Flooding Duration (s)	# of retransmissions
20	75.24	1387
10	35.88	1081
5	24.51	977
1	11.81	1638

  - Flooding duration >> expected (txInterval * 2K)
  - Huge number of retransmissions.
  - Huge delay in LSP acknowledgement.

TX Flow Control

- draft-ginsberg-lsr-isis-flooding-scale

```
Algorithm - run once/sec
------------------------

if (CurrentUackLSP > (CurrentLSPTxMax * UackSafe)) {
  CurrentLSPTxMax =
      max(MinLSPTx, (CurrentLSPMaxTx*UpdateBackoff))
} else { // CurrentUackLSP is at a safe level
  CurrentLSPTxMax = min(MaxLSPTx,
              CurrentLSPTxMax*((100 + UpdateIncrement)/100))
}
```

- Not quite effective:
  - Tunable parameters depend on receiver.
  - Receiver's PSNP approach is unknown.
  - Too conservative if the receiver is slow sending PSNPs.

Prompt PSNP

- Prompt PSNP can reduce LSP retransmissions
- draft-ginsberg-lsr-isis-flooding-scale

3.2. Rate of LSP Acknowledgments

On point-to-point networks, PSNP PDUs provide acknowledgments for received LSPs. [ISO10589] suggests that some delay be used when sending PSNPs. This provides some optimization as multiple LSPs can be acknowledged in a single PSNP.

If faster LSP flooding is to be used safely, it is necessary that LSPs be acknowledged more promptly as well. This requires a reduction in the delay in sending PSNPs.

As PSNPs also consume link bandwidth and packet queue space and protocol processing time on receipt, the increased sending of PSNPs should be taken into account when considering the rate at which LSPs can be sent on an interface.

  - How much delay?

Prompt PSNP

- draft-decraene-lsr-isis-flooding-speed

```
The way LSPs are acknowledged faster is a local decision on the
receiving IS.  Without limiting the possibilities, there are at least
two options:

o  Reduce partialSNPInterval.  Possibly reduce it even further when
   the IS-IS adjacency initially transitions to the UP state, when a
   large number of LSPs need to be received quickly, until the LSDB
   has been synchronized.  The choice of this lower value is a local
   choice.  It may depends on the (available) processing power of the
   node, the number of adjacencies been brought up at the same time,
   the requirement to synchronize the LSDB more quickly.

o  Track the number of received and un-acknowledged LSP per interface
   and level, and send a PSNP when the number reaches 90 (max per
   PSNP) or is significant enough e.g., 10 or
   InterfaceLSPReceiveWindow/2.
```

  - Threshold-based PSNP. What is a good number for threshold?

PSNP Tradeoff

- Too slow: cause retransmissions
- Too fast: waste bandwidth, queue space, …

- Which PSNP interval/threshold to use?

- Drafts should tell us how fast to send PSNP.

Reduce PSNP Interval

- Reducing PSNP interval can reduce LSP retransmissions and flooding duration.

PSNP interval (s)	Flooding duration (s)	# of retransmissions
1	14.25	1291
0.5	10.39	218
0.1	2.97	0

Use Threshold-Based PSNP

- Threshold-based PSNP is effective to handle LSP bursts

PSNP threshold	Flooding duration (s)	# of retransmissions
1 *	3.05	0
15 *	2.95	0
30	3.06	0
60	3.11	0
90 *	3.17	0

  - Choice of threshold below 90 does not matter much

Hybrid PSNP Approach

- Two approaches can co-exist:
  - Interval-based approach for infrequent LSPs
  - Threshold-based approach for bursty LSPs

	Flooding Duration (s)	# of Re-transmissions
Initial (33ms/2s)	94.94	803
Faster Tx (1ms/2s)	11.81	1638
Faster PSNP (1ms/0.1s)	2.97	0
Hybrid PSNP (1ms/2s/th=15)	2.95	0

Further Reduce LSP Tx Interval

- With the hybrid PSNP approach, we can reduce the LSP Tx interval more aggressively:
  - PSNP interval 2s, threshold 15.

Lsp Tx Interval (ms)	Flooding Duration (s)	# of retransmissions
1	2.95	0
0.5	1.57	0
~0 (2K LSPs in one burst)	N/A	0

  - No LSP retransmissions

Summary

- Threshold-based PSNP is effective.
- Hybrid PSNP approach
  - Slow PSNP for infrequent LSP updates.
  - Fast PSNP for LSP bursts.
- With prompt PSNPs, LSP Tx interval can be set more aggressively.
- Difficulties in implementing TX flow control.