IS-IS Flooding Speed advertisement

draft-decraene-lsr-isis-flooding-speed-01

Bruno Decraene	(Orange)
Chris Bowers	(Juniper)
Jayesh J	(Juniper)
Tony Li	(Arista)
Gunter Van de Velde	(Nokia)

IGP flooding: introduction

IGP flooding is paramount for Link State IGP

Slow flooding directly translates to

- Delayed network reaction to failure hence packet loss
- LSDB inconsistency hence routing inconsistencies and micro-loops.

IGP flooding is hard

- Needs to be fast to have LSDB in sync
  - A single node failure → N LSPDUs → need to fast flood N LSPDUs
- Needs to be slow enough to not overload the routers
  - Not overloading my adjacent neighbors is enough

IS-IS flooding: status

IS-IS has no signaling to control flooding speed.

Current flooding speed is extremely conservative

- Default delay between LSPDUs: 33ms or 100ms
- → 30pps / 340**k**bit/s per neighbor
- vs BGP policer set to 20 000pps/ 224Mbits/s

Same flooding speed for the last 20 years

- While CPU and interfaces speed improved a lot

Proposed evolution

1 TLV in IS-IS Hello: Flooding Speed TLV

Advertises my LSPDU receiving speed

to my upstream flooding neighbors.

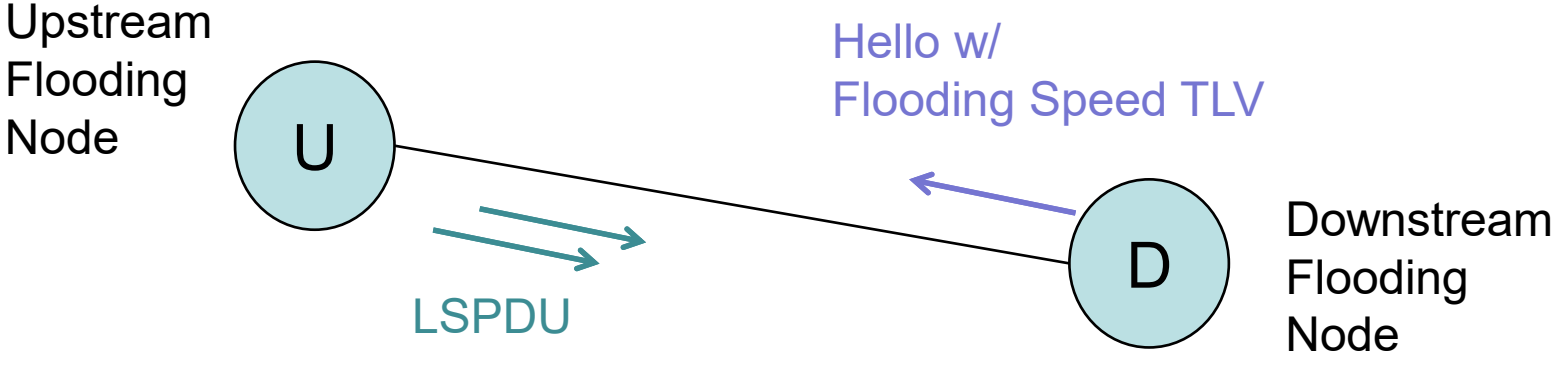

Speed Parameters

Parameters based on existing implementations /knobs

- Reduces implementation effort
- Increases adoption likelihood
- Avoids changing touchy parts of the code

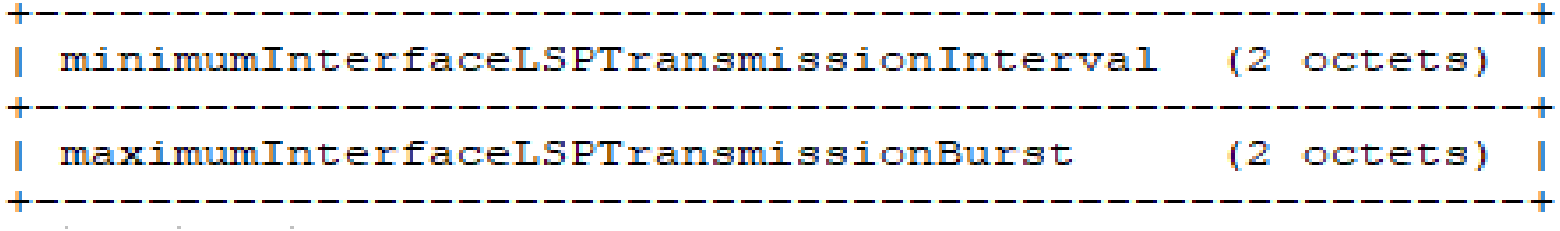

Figure 1: Flooding Speed TLV

Delay between 2 LSPDUs

minimumInterfaceLSPTransmissionInterval

- The minimum interval, in milliseconds, between two consecutive LSPDUs
- Matches (all) existing flooding implementations
  lsp-interval (C, J), lsp-pacing-interval (N, H), lsp tx interval (A)
- Akin to a CPU/processing performance

Fast Flood N LSP

maximumInterfaceLSPTransmissionBurst

- Number of (un-acknowledged) LSPDUs which may be sent in a burst
- Matches some implementations

  fast-flood (H), DDoS filtering (bandwidth-limit; burst-size-limit) (J)
- Akin to a memory/buffer size

Toward flow control

Acknowledging the reception of LSPDUs

using existing ISO specification (PSNP, CSNP)

provides dynamic flow control

- maximumInterfaceLSPTransmissionBurst serves as the static transmission window
- xSNP serves as dynamic acknowledgement

Summary (1)

Improving IS-IS flooding speed.

By advertising two static parameters

- from the flooding downstream node
- to the flooding upstream node

Matching existing implementations behavior

Summary (2)

Draft is short but allows for a choice of 3 levels of improvements on the upstream flooding node:

1. Coordination of lsp-pacing
   - No change on flooding implementation
   - Significant improvement in real life (e.g. x10)
2. Use of a burst size
   - Significant improvement for the advertisement of a set of LSPs (node failure)
3. Dynamic flow control
   - Significant improvement compared to lsp-pacing

Local choice/behavior of the upstream node when receiving the (same) TLV.