

Use of an MPLS LSE as an Ancillary Data Pointer

draft-bryant-mpls-aux-data-pointer-00

Stewart Bryant (sb@stewartbryant.com),

Alexander Clemm (ludwig@clemm.org),

Toerless Eckert (tte@cs.fau.de)

A General Point

- Although what is presented is a proposed solution to a specific MPLS problem, it is useful to consider the added value of pointers in packet design.
- Pointers allow specified rather than implicit data/action association/direction.
- draft-bcx-rtgwg-tcr-00 (not MPLS) gives a core complete idea of what can be achieved using pointers.

Problem Statement

- Certain use cases benefit from ancillary data processed/accessed as part of a forwarding decision
- Problem : How to add information to MPLS packets in a way that is:
 - Suitable for efficient high speed forwarding.
 - Easy for the more modern existing h/w to add the feature.
 - Backwards compatible in terms of basic forwarding with legacy h/w

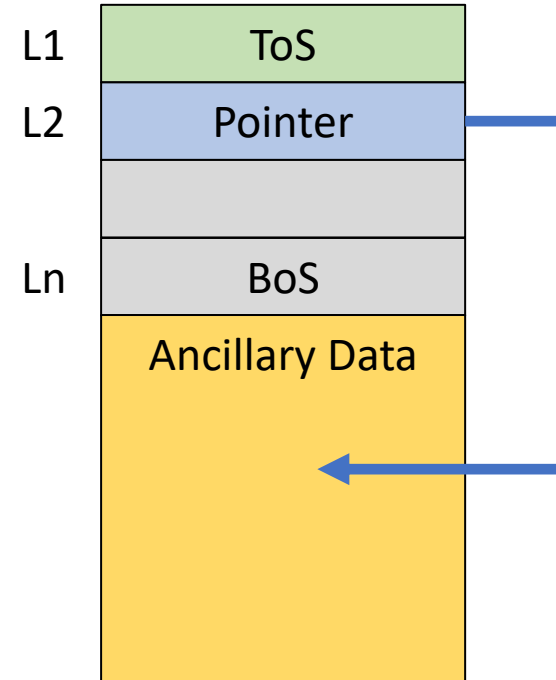
Approaches Proposed So Far

- The approaches that have been proposed so far* rely on the forwarder:
 - Finding out if there is applicable ancillary data below BoS
 - Deducing which of the ancillary data applies to this hop
- Some methods make it easier to deduce if there is data, but not where the data is.
- None of the proposed methods deal well with case of ancillary data that is different at different hops.
- This approach builds on the observation in draft-kompella-mpls-mspl4fa that if an LSE is not ToS the TC and TTL bits are not used.

* Except draft-stein-srtsn which puts the information in the stack

Core Idea

- Use the “spare” non-ToS fields as a pointer to the ancillary data applicable at this hop.
- Semantic: “Process as described by L1, using the information pointed to by L2”
- Forwards normally when L2 not a pointer or when LSR does not understand the pointer mechanism



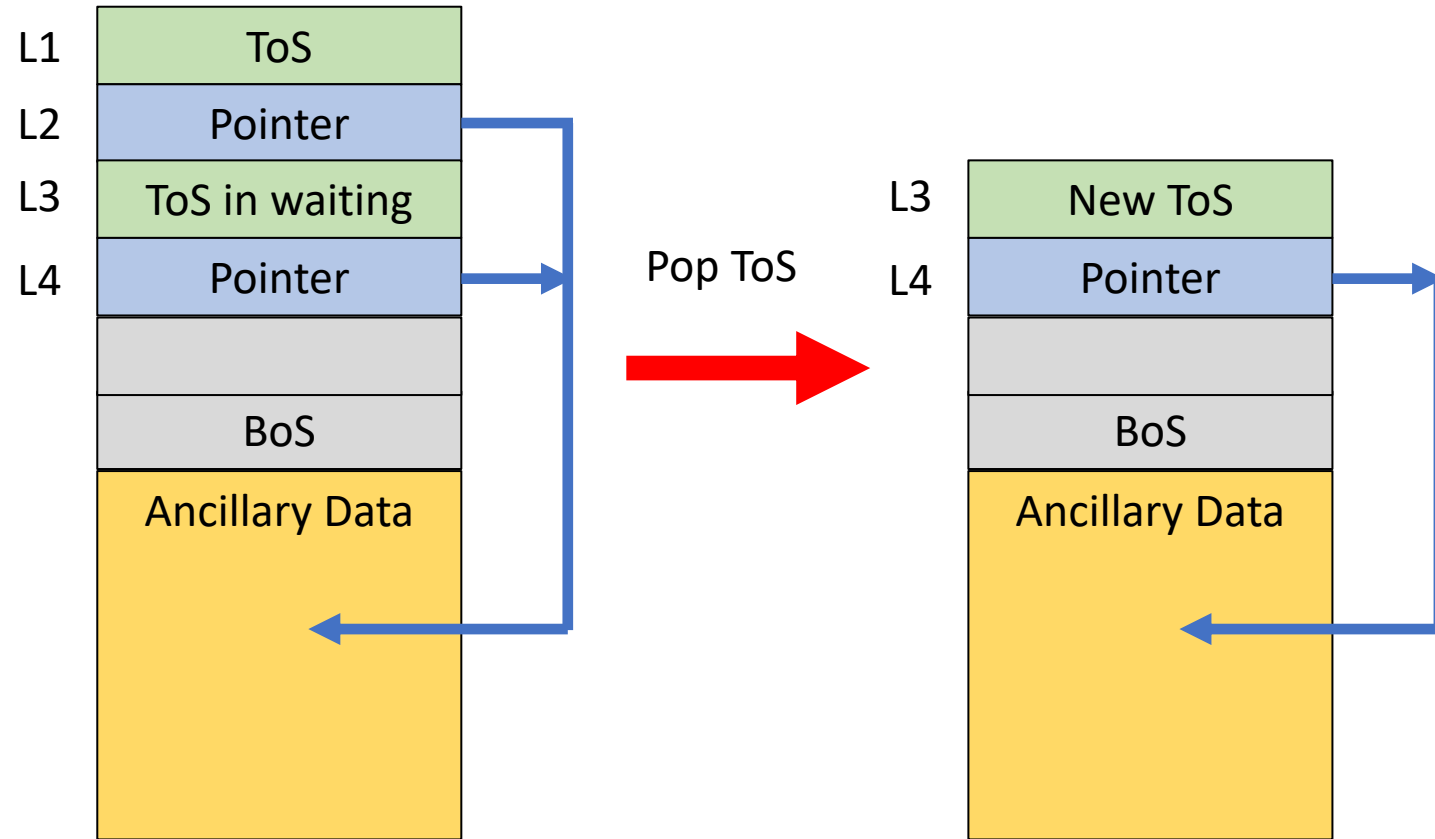
Advantages

- Ability to find the ancillary data without reading the whole of the stack
 - Speculative processing can be expensive
- Ability to specify which ancillary data is applicable applicable to which forwarding label
- Simplifies packet parser as no deduction or discursion needed
- Inherently general and extensible.

Special Purpose Label

- Assumption: pointer will an SPL of some sort.
- Could make ToS indicate pointer follows, BUT that means
 - We need to change the FEC of the ToS label
 - We no longer have legacy compatibility
 - We will need more labels in the global label table.
 - Increased cost of distribution and management
 - Some LSRs (particularly PE LSRs) are already saturating the global label table.
- We investigate another approach later in draft.

Single Pointer From Multiple LSEs

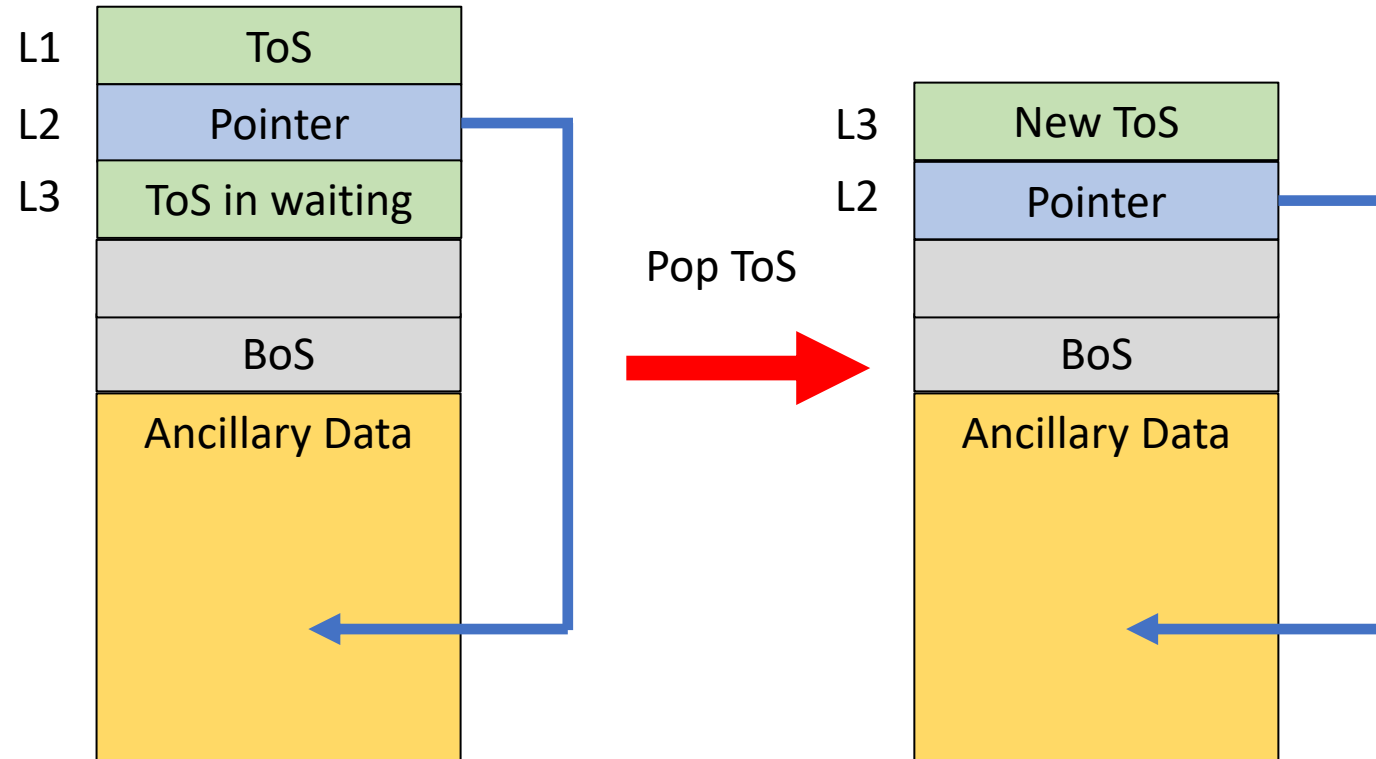


- Applicable to any label stack
- Particularly applicable to SR

BUT

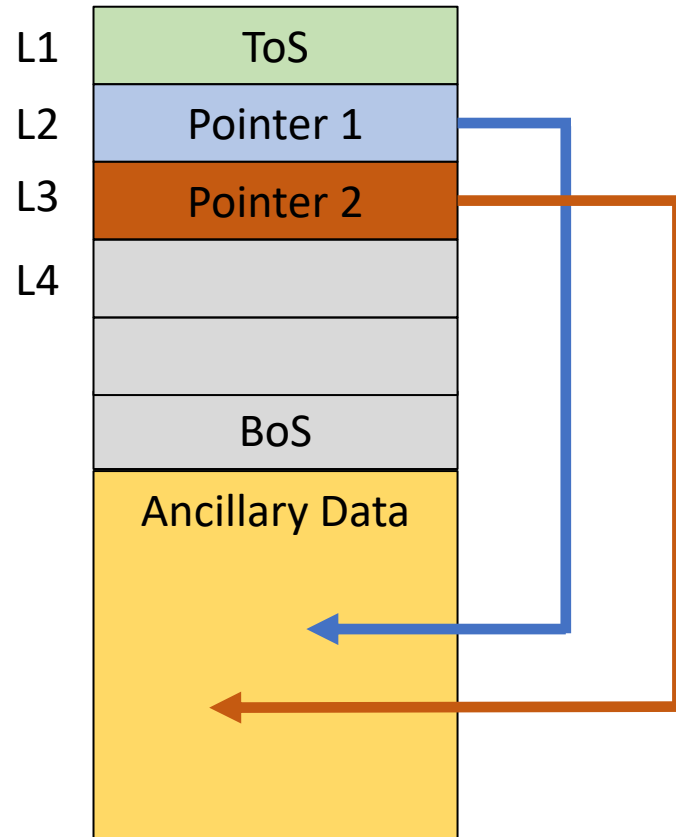
- Not efficient in use of stack space to duplicate pointers

A More Efficient Approach



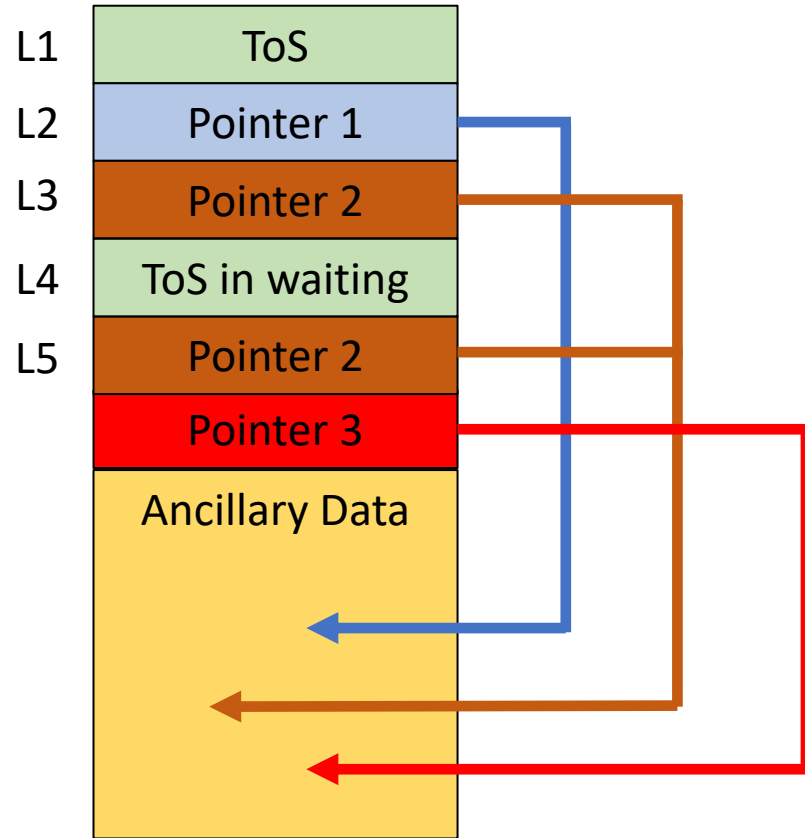
- Use an in-stack move (swap)
- More efficient in stack space
- More complex than simple Pop
- Need to correct pointer – subtract 4 bytes
- Problem – how to know when to stop propagating the pointer LSE?
 - “TTL” in pointer ?
 - Bit in L3 ?
 - FEC of L3 ?
 - Pointer pop SPL?

Multiple Pointers



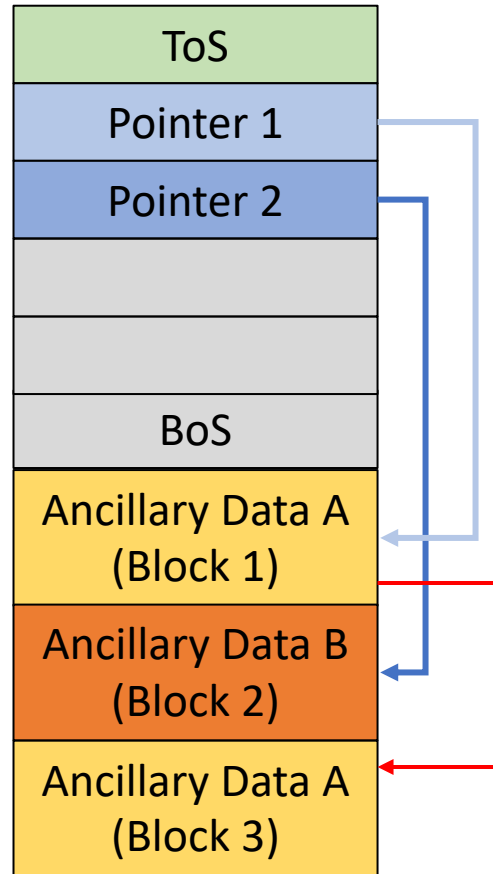
- There are times when multiple pointers are needed, for example iOAM and LBF
- L1, L2 and L3 are a group of pointers for pop and “in-stack swap” operations
- The semantic is “Process as described by L1, using the information pointed to by L2 and L3”

Multiple Pointers cont



- Pointer groups can include pointers to objects in common with other pointer groups and unique pointers.

Variable-length data and data insertion (in -01)




- Ancillary Data allocated as fixed-size Ancillary Data Blocks.
- To add Ancillary Data, allocate another Block and have previous Block link to its successor
- Allows to grow (and shrink) ancillary data without affecting pointers to it
 - e.g. growing A does not affect pointer to B
 - read/write operations to Ancillary Data across Blocks can be transparent to applications
- Compare with Disk Operating Systems: variable file sizes and fixed sector sizes

Disposition of Ancillary Data

- Ancillary data needs to be removed before the payload is passed out of the MPLS domain.
- This can be a lot more complicated than just disposing of n bytes.
- Some methods:
 - FEC of BoS LSE (as in PW or MPLS VPN)
 - SPL at BoS
 - BoS LSE can point to ancillary data that describes the disposition. This is a powerful approach.

Pointer LSE



SPL or Regular Label?

- SPLs are in short supply
- ESPLs need twice the stack space
- Could we use a regular label as a pointer label?
- We are not talking global labels here we are talking of a small number of network wide agreed labels to be specially recognized by the forwarder.
- We do not need a common label block in the normal sense (as needed by SR) since these labels will not appear in the FIB.
- We will have to modify the label manager in LSRs that could have these labels in the FIB to exclude these labels. That is work but should not be “hard”.
- Applicable to pointers and other types of indicator label.

Questions?