



SCION: Data Plane Overview

Adrian Perrig

Network Security Group, ETH Zürich

SCION Data Plane Overview


- Data plane: How to send packets
[Chapter 2.2, Chapter 8]
 - Path lookup
 - Path combination
 - Path encoding in packet

Path Lookup

- Steps of a host to obtain path segments
 - Host contacts RAINS server with a name
H → RAINS: www.scion-architecture.net
RAINS → H: ISD X, AS Y, local address Z
 - Host contacts local path server to query path segments
H → PS: ISD X, AS Y
PS → H: up-path, core-path, down-path segments
 - Host combines path segments to obtain end-to-end paths, which are added to packets


Path Lookup: Local ISD

- Client requests path segments to $\langle \text{ISD}, \text{AS} \rangle$ from local path server
- If down-path segments are not locally cached, local path server send request to core path server
- Local path server replies
 - Up-path segments to local ISD core ASes
 - Down-path segments to $\langle \text{ISD}, \text{AS} \rangle$
 - Core-path segments as needed to connect up-path and down-path segments

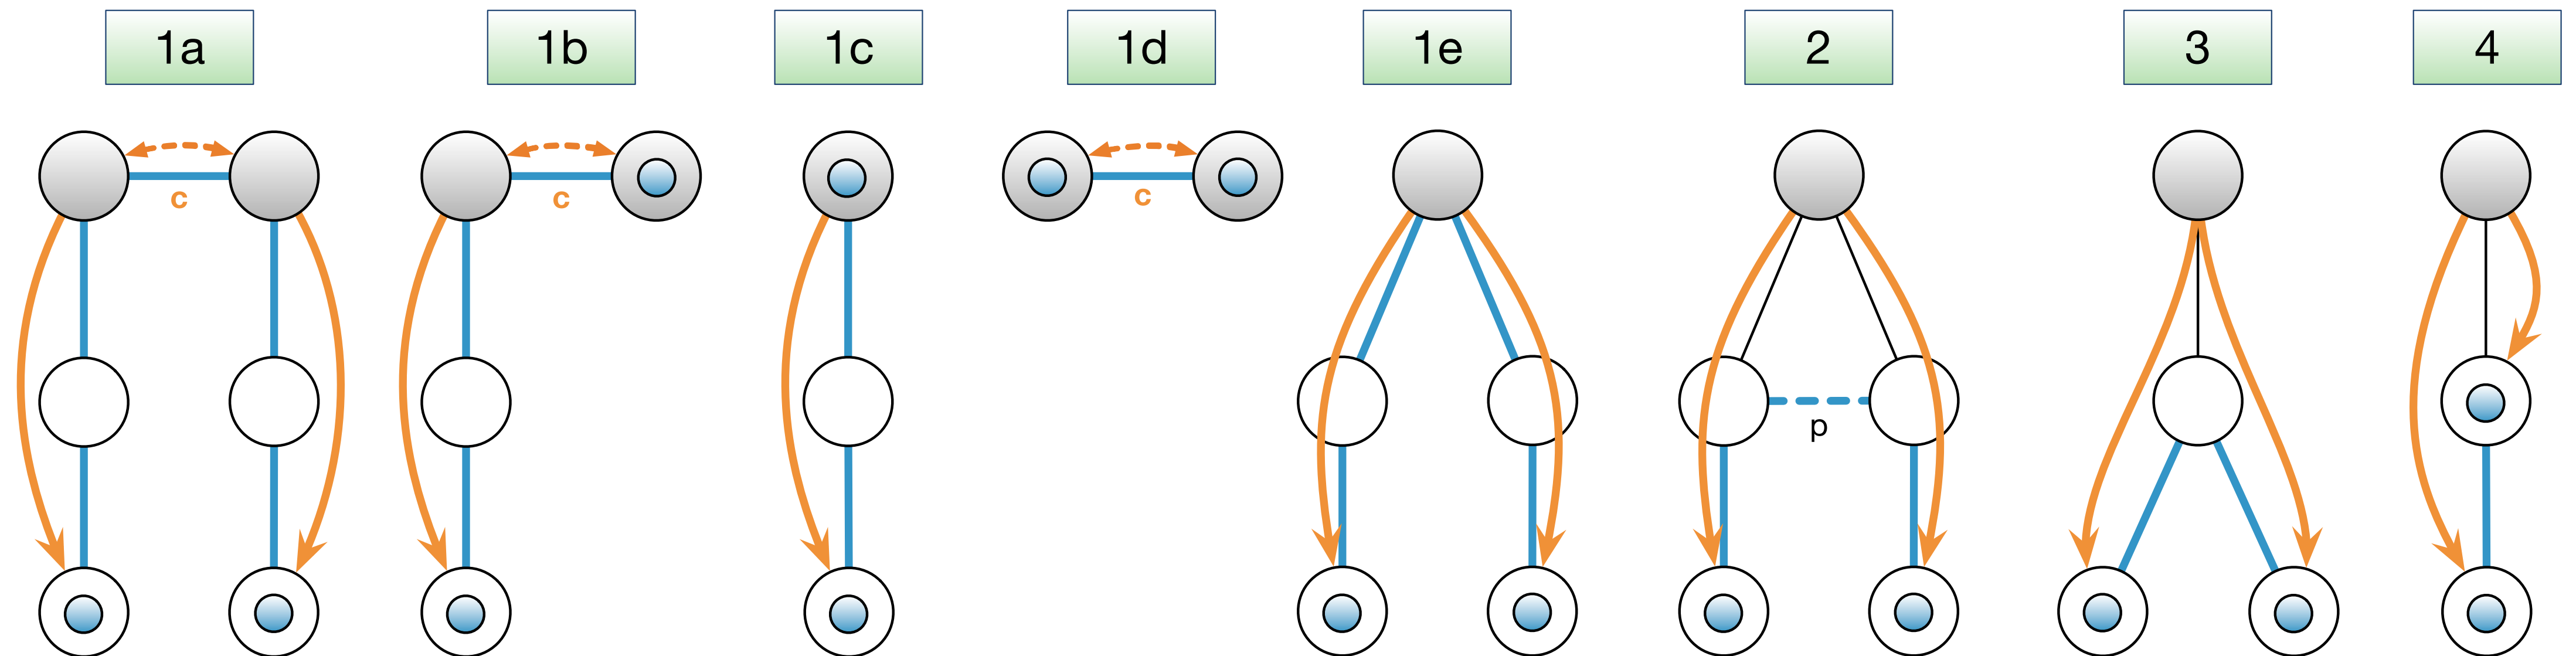


Path Lookup: Remote ISD

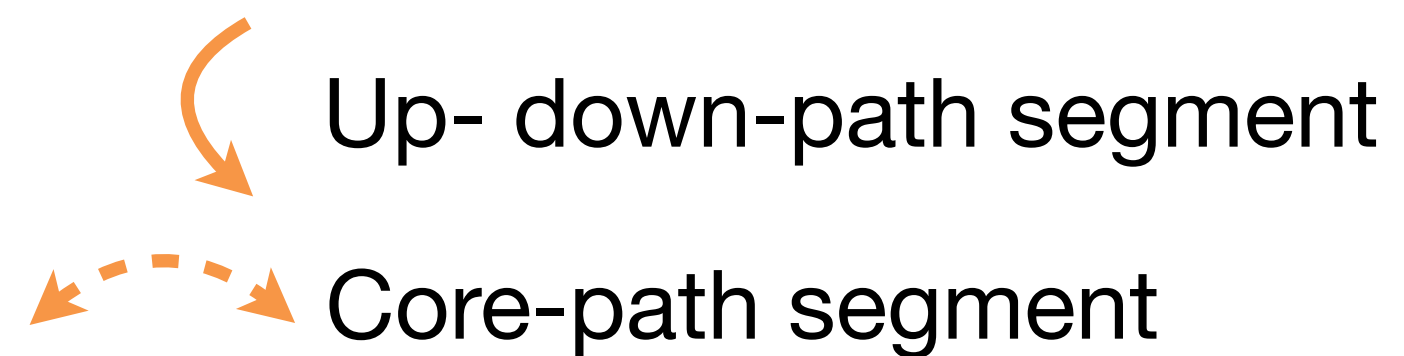
- Host contacts local path server requesting $\langle \text{ISD}, \text{AS} \rangle$
- If path segments are not cached, local path server will contact core path server
- If core path server does not have path segments cached, it will contact remote core path server
- Finally, host receives up-, core-, and down-segments



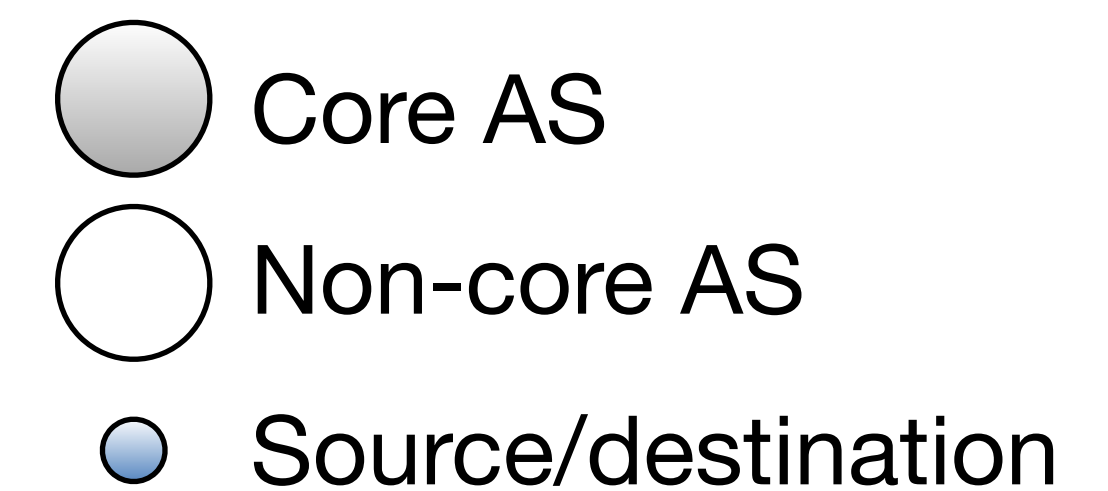
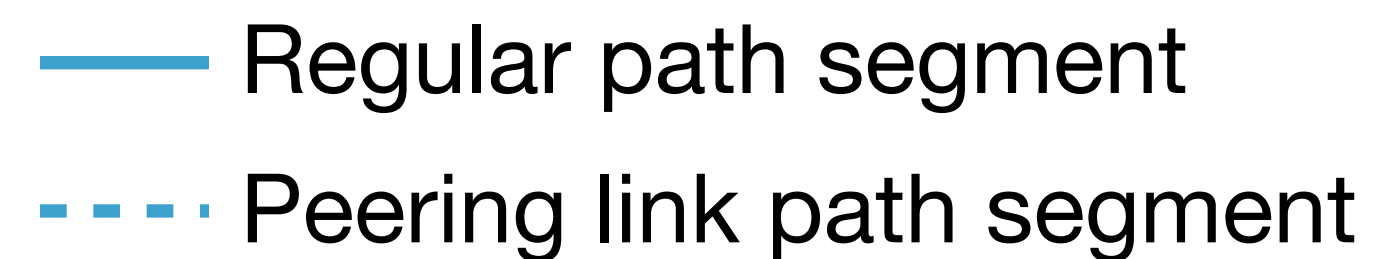
Path Combination



Control-plane path segments:





Data-plane paths:





Path Combination Example (1)

- Core-segment combination:
Up-path segment +
core-path segment +
down-path segment





Path Combination Example (2)

- Peering shortcut: up-path segment and down-path segment offer same peering link





Path Combination Example (3)

- Peering shortcut: up-path segment and down-path segment offer same peering link

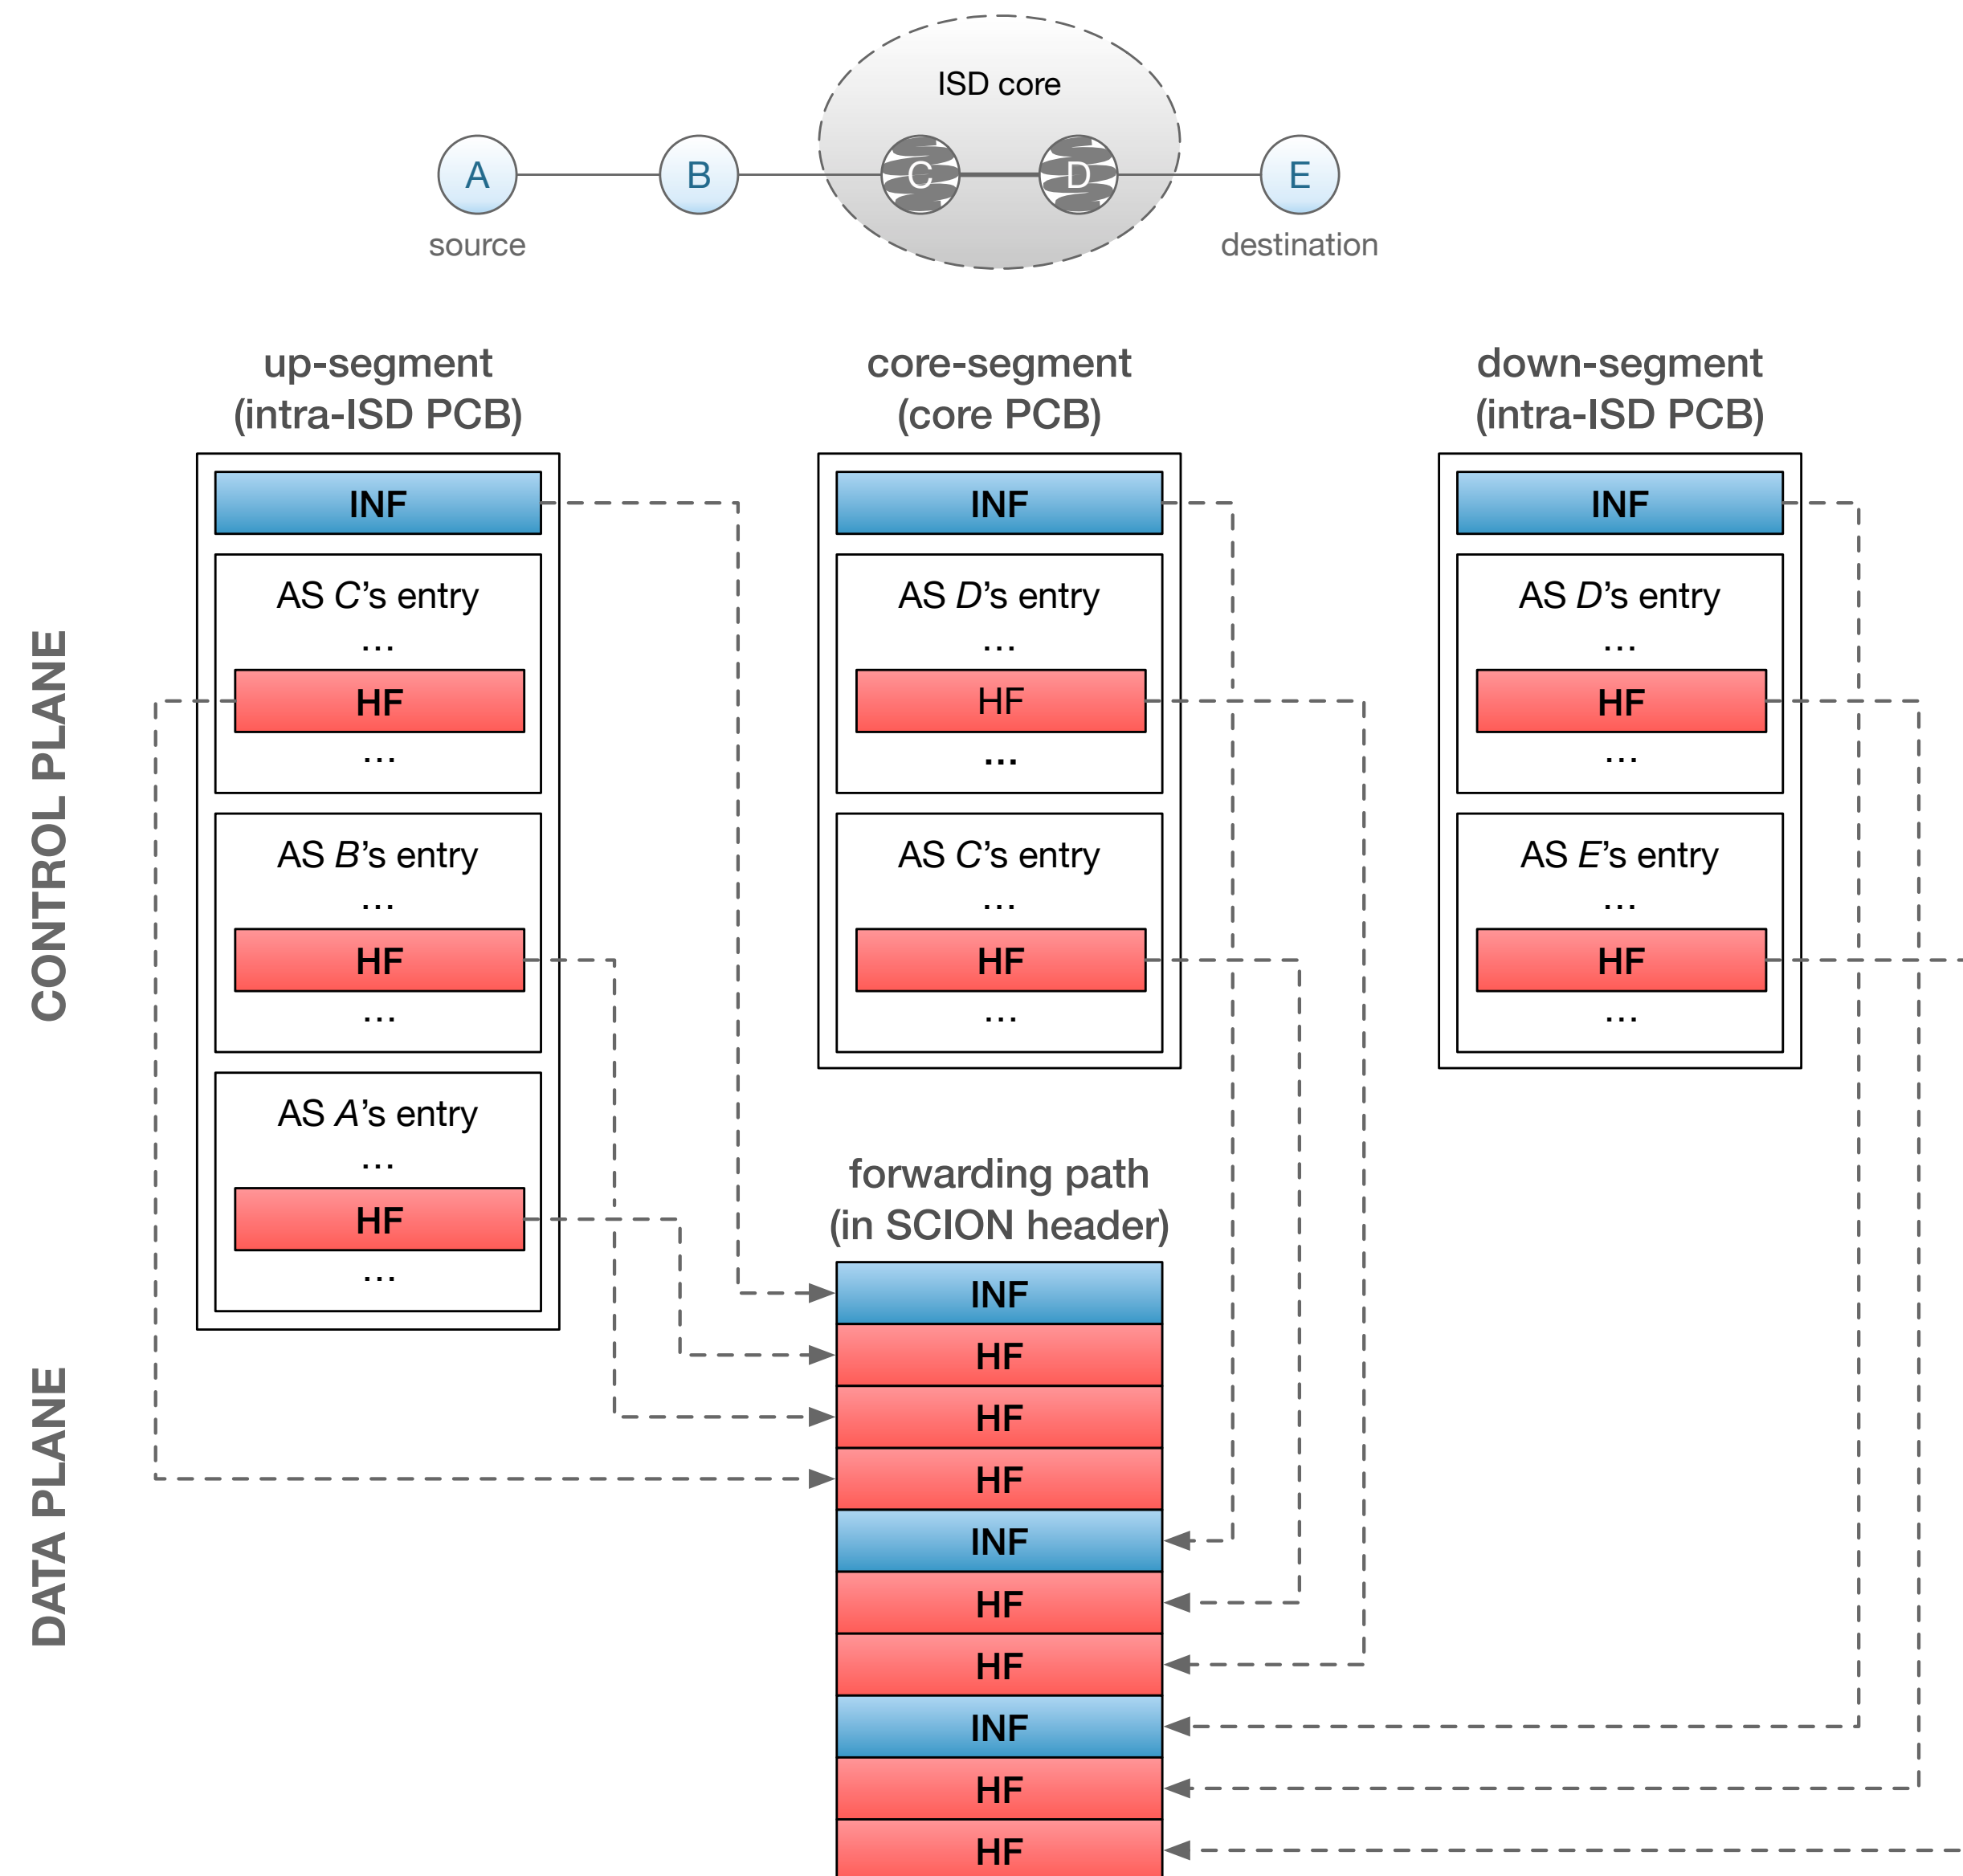


Path Combination Example (4)

- AS shortcut path through common AS on up-path and down-path segment



Path Construction



SCION Packet Header

- SCION common header encodes:

- Version
- Destination and Source address types
- Total packet and header length
- Pointer to current info and hop field
- Next header type field

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Version				DstType						SrcType							TotalLen														
HdrLen								CurrINF								CurrHF								NextHdr							

- SCION source and destination address encoding

- ISD-AS of source and destination are listed first to simplify parsing (constant offset)
- Destination local address is also at a fixed location
- Source local address is at a variable location

0	11	31	43	63
DstISD	DstAS		SrcISD	SrcAS
DstHostAddr (IPv6)				
SrcHostAddr (IPv4)			Padding	

Info and Hop Field Contents


- An **info field** provides information about a path segment, which consists of one or multiple **hop fields**

- An info field contains



- Flags: PEER, SHORTCUT, UP
 - Timestamp containing the creation time
 - ISD identifier
 - Path segment length


- A hop field contains



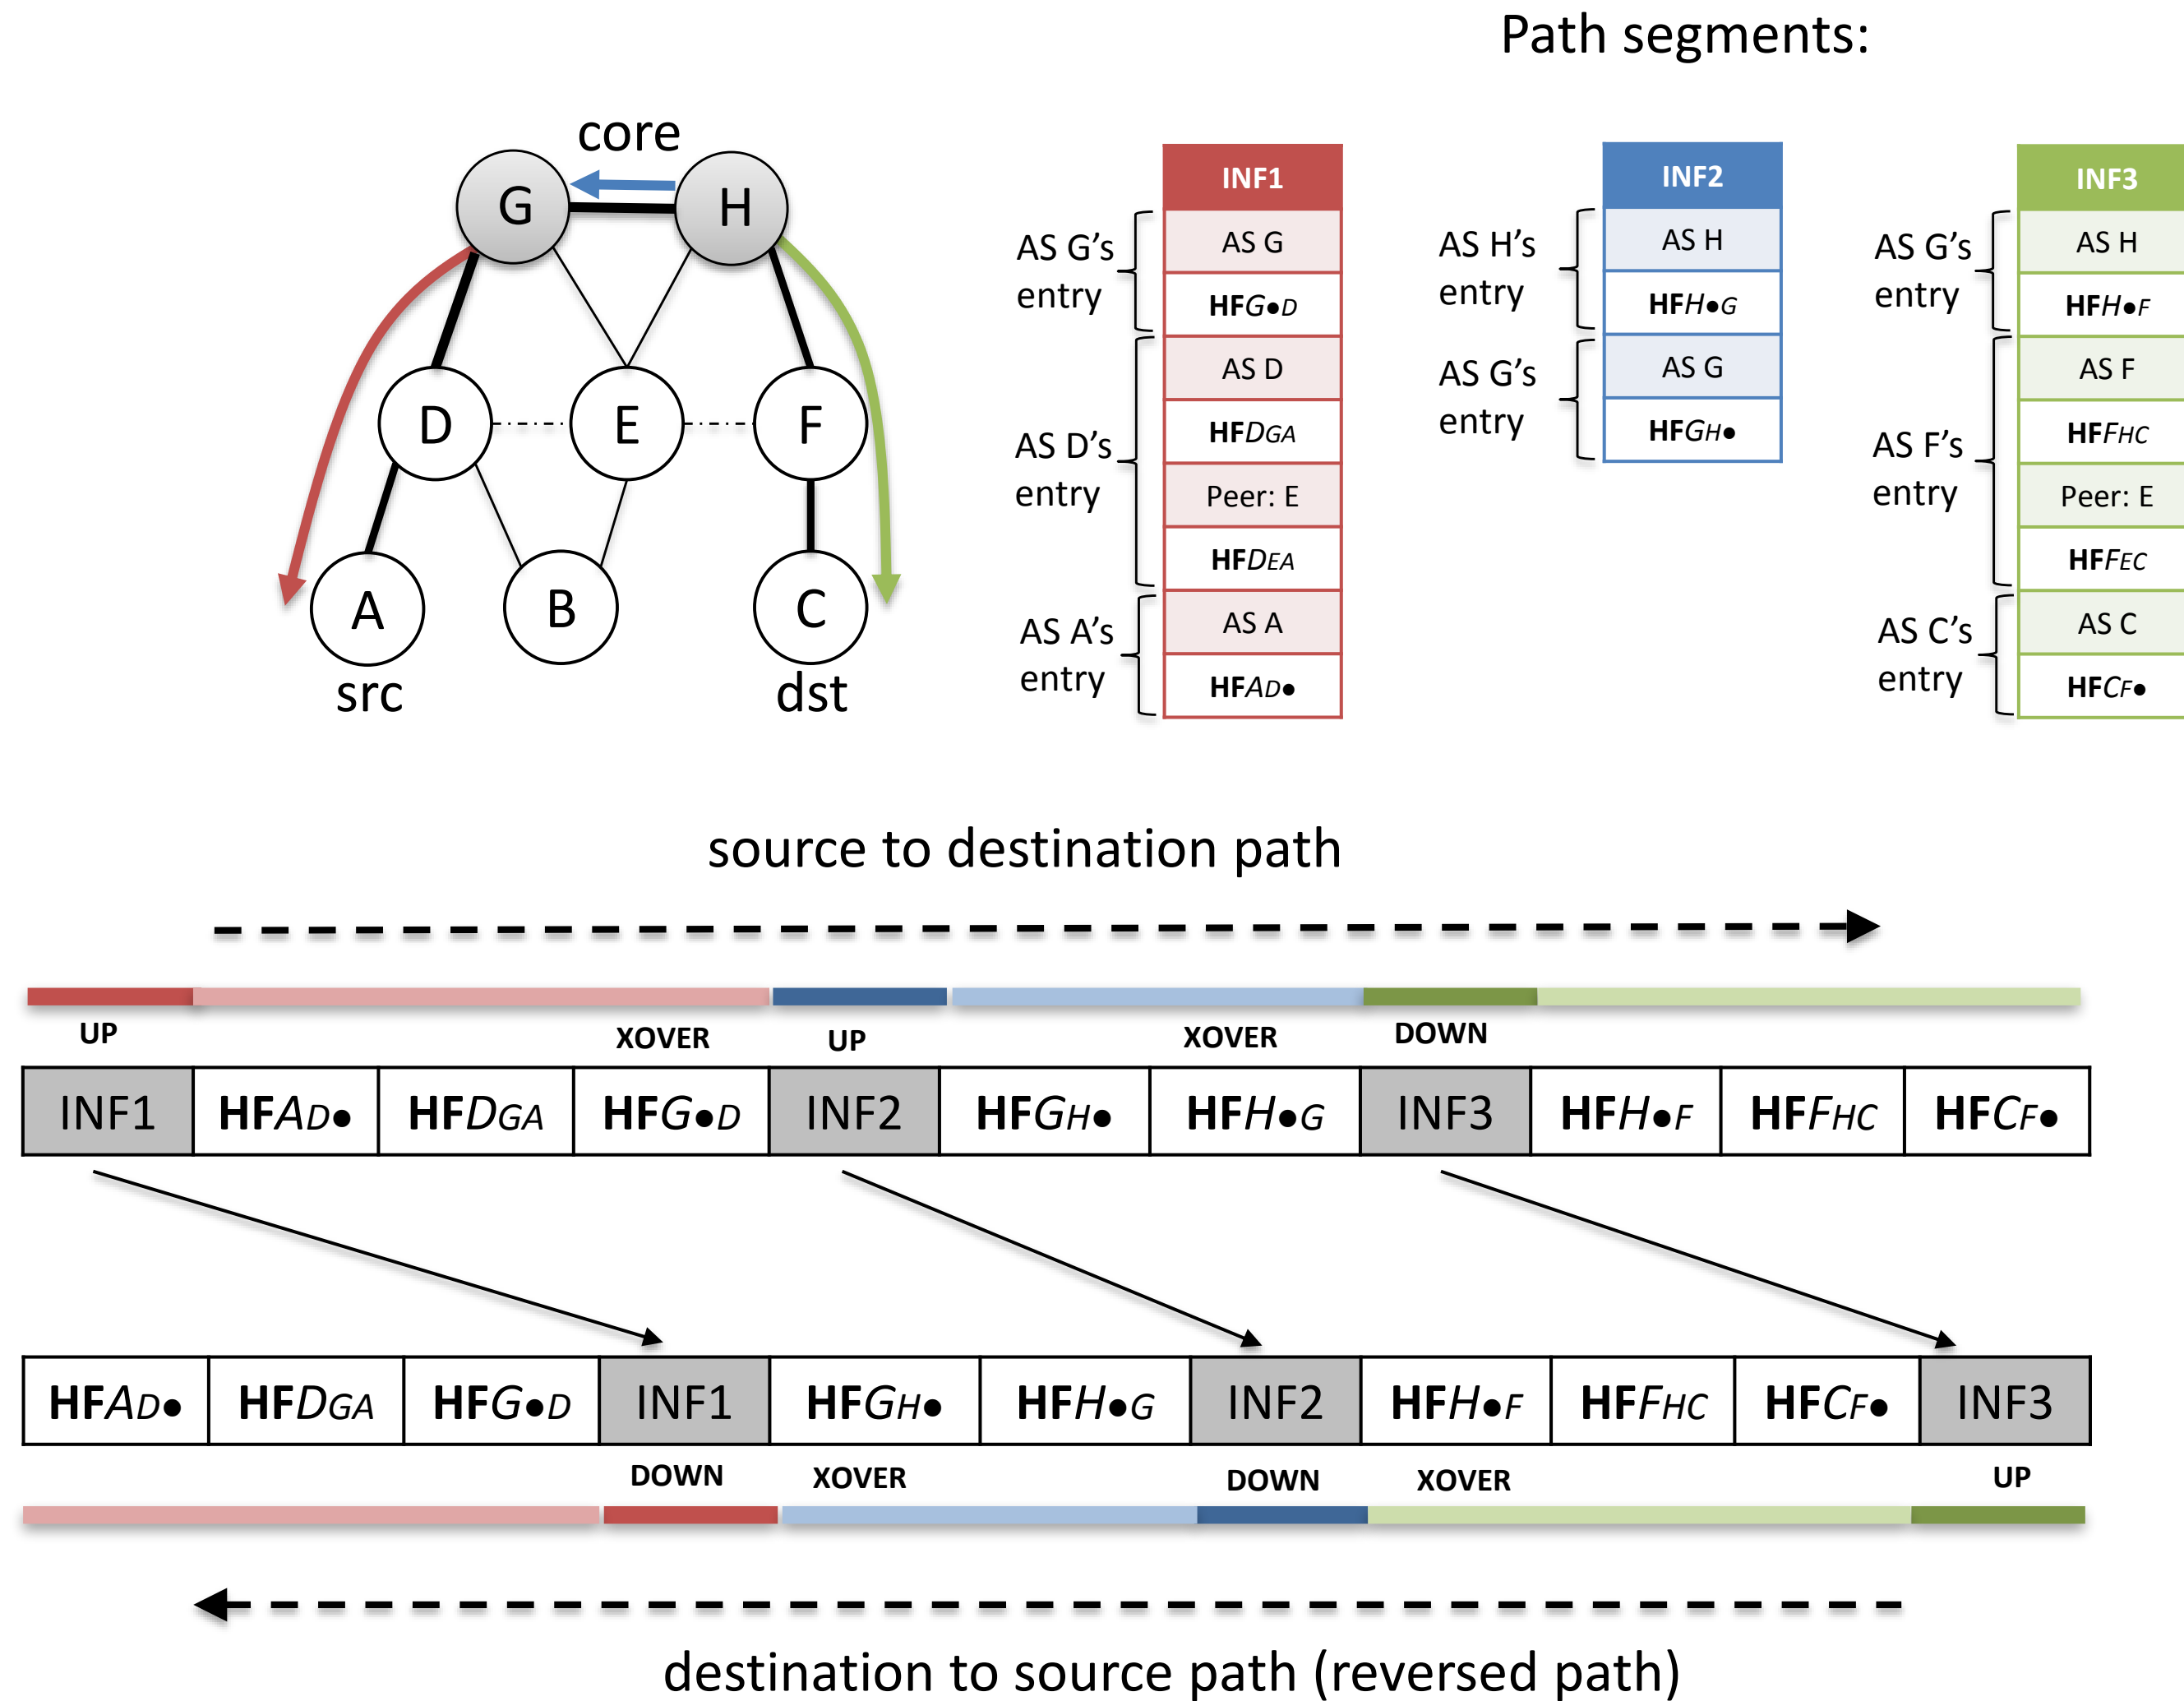
- Flags: CONTINUE/STOP, FWD-ONLY, VRFY-ONLY, XOVER
 - Expiration time, relative to timestamp in info field
 - Ingress and egress interface identifiers
 - Message Authentication Code (MAC)

Ingress and Egress Interface Identifiers

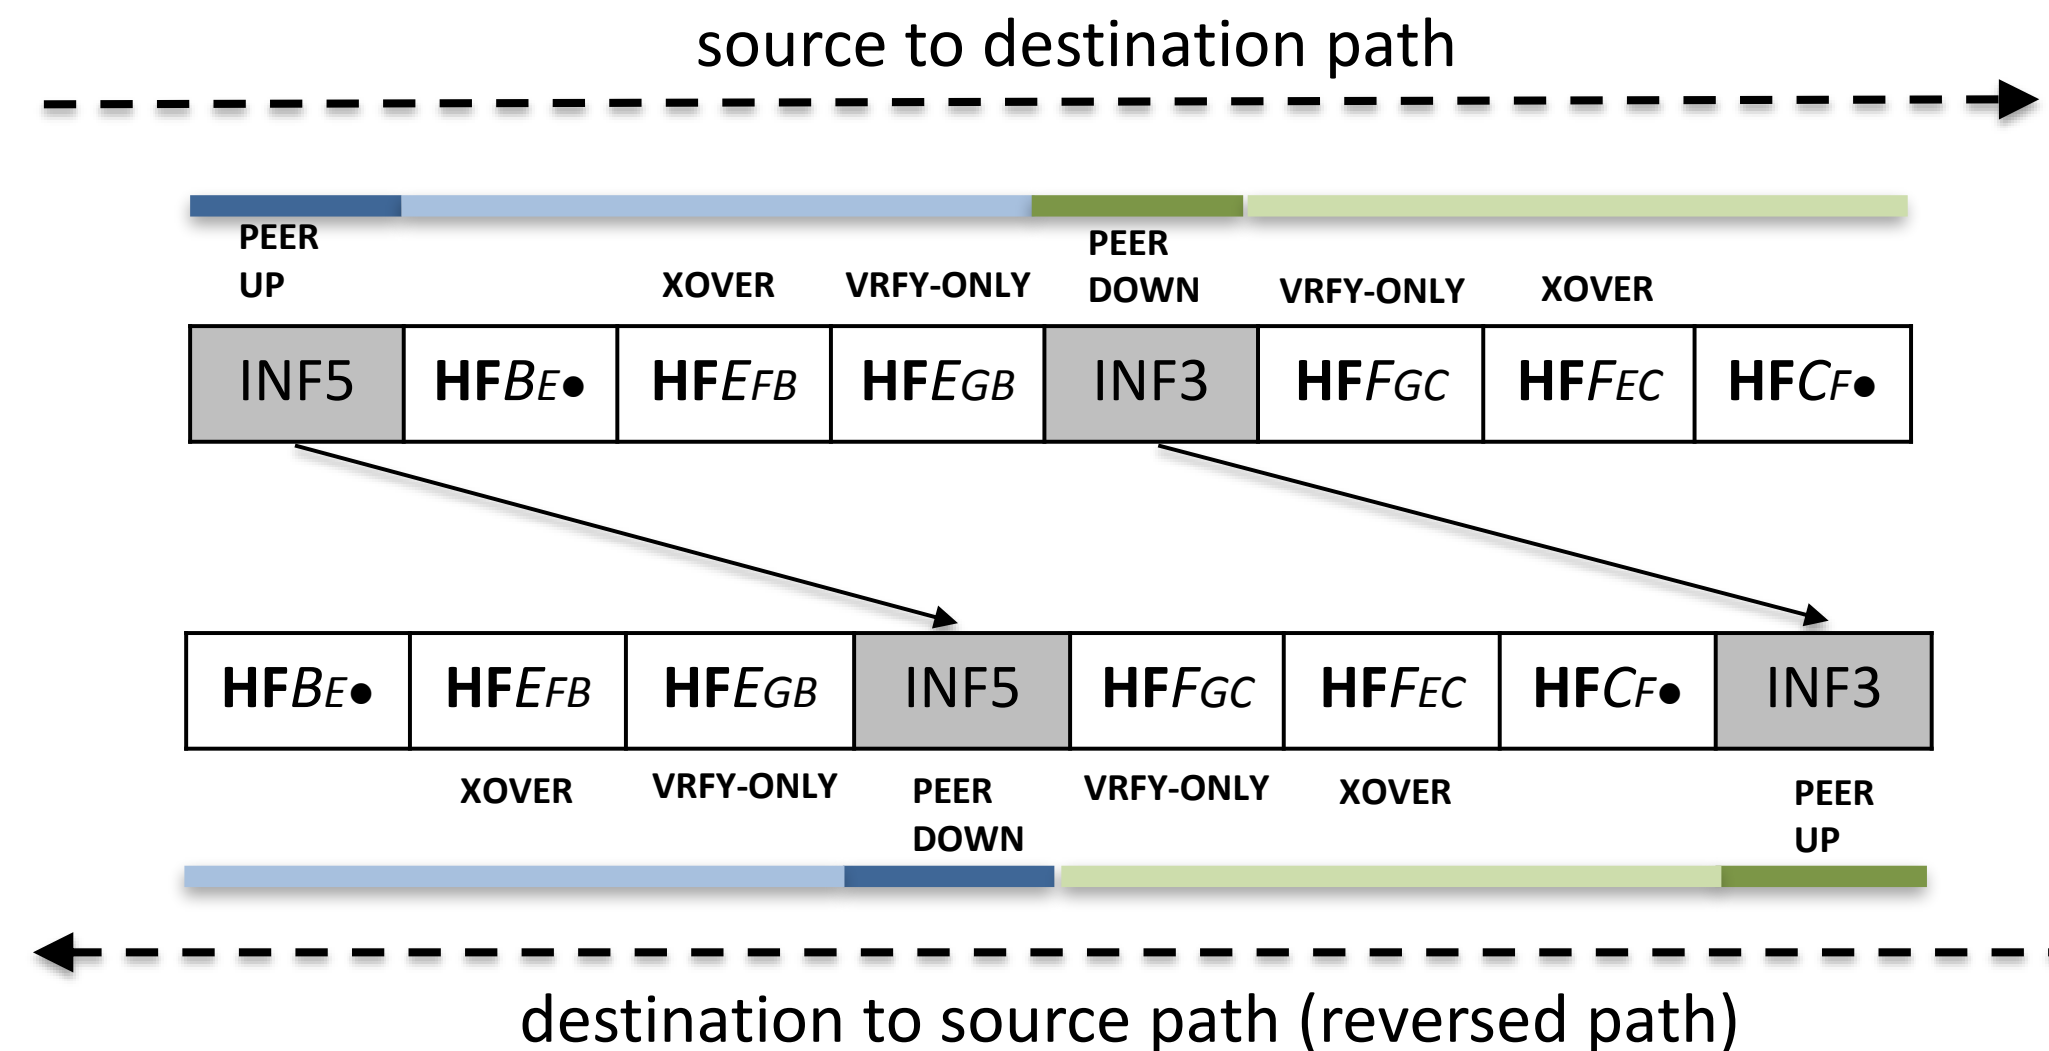
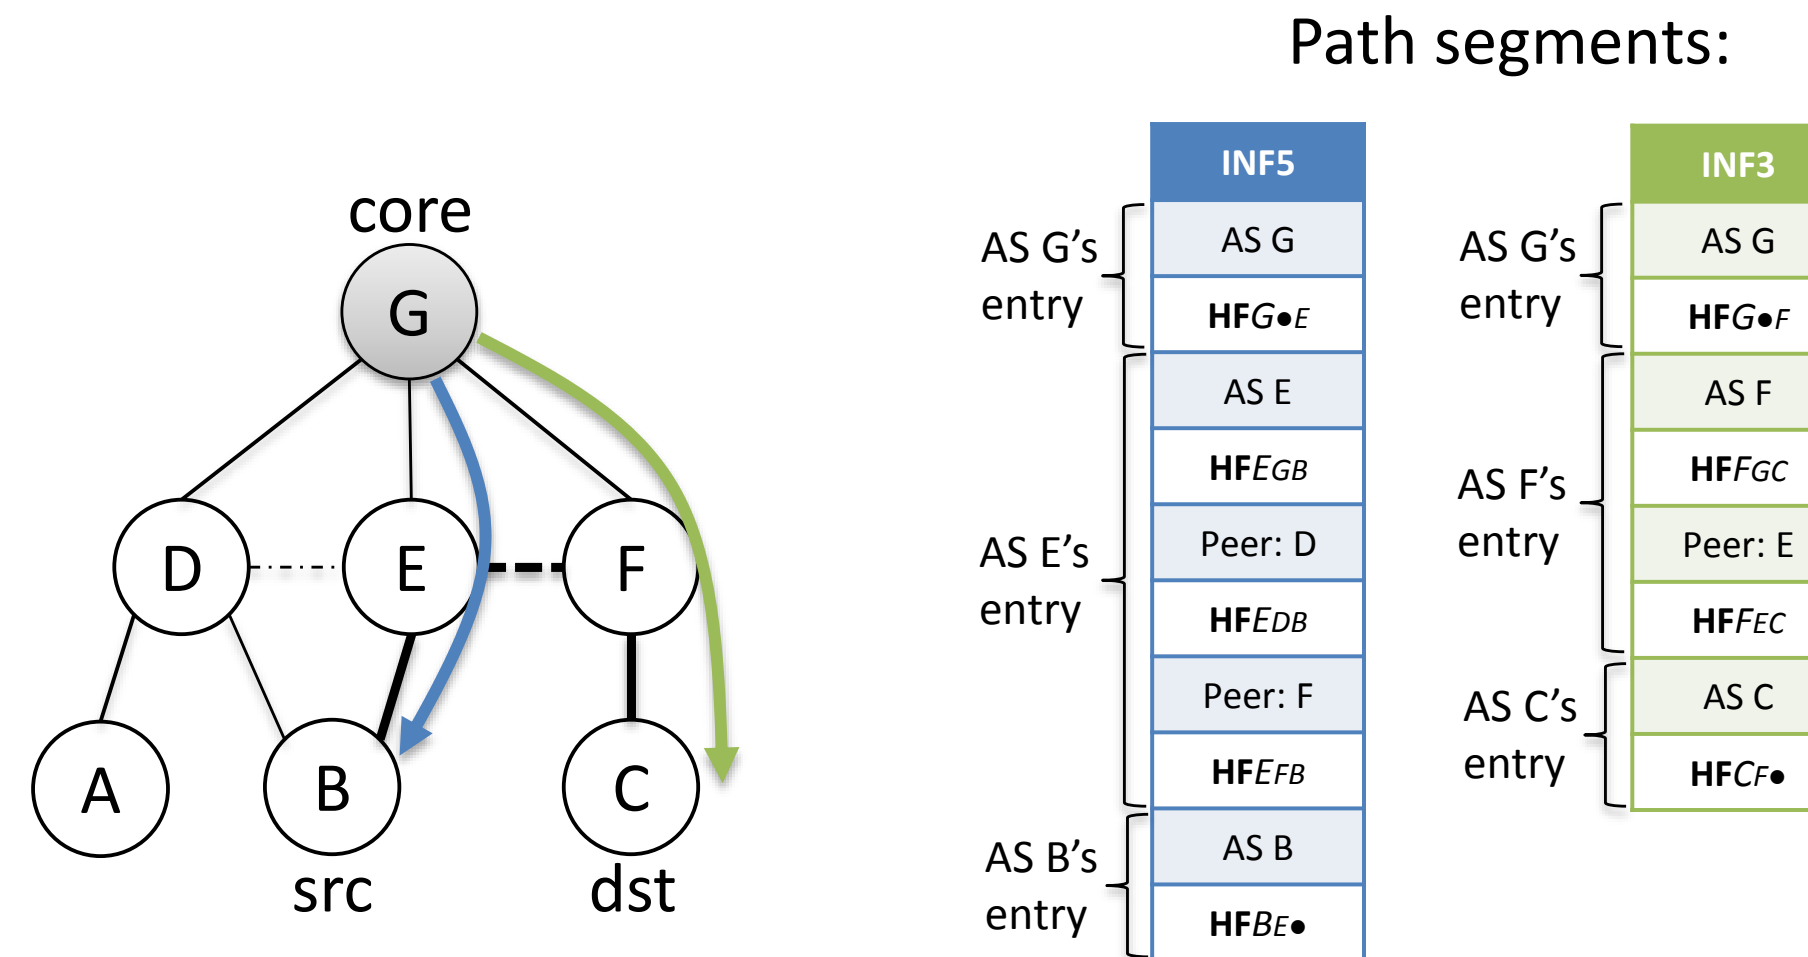
- Each AS assigns a unique integer identifier to each interface that connects to a neighboring AS
- The interface identifiers identify ingress/egress links for traversing AS
- ASes use internal routing protocol to find route from ingress SCION border router to egress SCION border router
- Examples
 - Yellow path: L:4, O:3,6, R:1
 - Orange path: L:5, O:2,6, R:1



Path Encoding in Packet



Path Encoding in Packet



Hop Field MAC Verification

- Message Authentication Code (MAC) computation and verification of Hop Field MAC value based on local AS secret key
 - Key is not shared with any external entity
- Computation: $\text{MAC}_K(\text{Timestamp}, \text{Flags}'_{\text{HF}}, \text{ExpTime}, \text{Ingress}, \text{Egress}, \text{HF}'))$
 - HF' is hop field of previous AS
- In most cases, HF' size is 8 bytes, so MAC computation can be done over 128 bits: with CMAC and AES, only a single encryption operation is needed
- With AESni HW crypto, only ~50 cycles are needed to compute MAC!
 - Note that a DRAM memory lookup takes ~200 cycles
 - AES operation requires less energy than TCAM lookup
 - Thus, **SCION forwarding can be faster and require less energy than IP forwarding**

For More Information ...

- ... please see our web page:
www.scion-architecture.net
- Chapter 8 of our book “SCION: A secure Internet Architecture”
 - Available from Springer this Summer 2017
 - PDF available on our web site