

Dig into MPLS: Transit Tunnel Diversity

Yves Vanaubel, Pascal Mérindol,
Benoit Donnet, Jean-Jacques Pansiot



Agenda

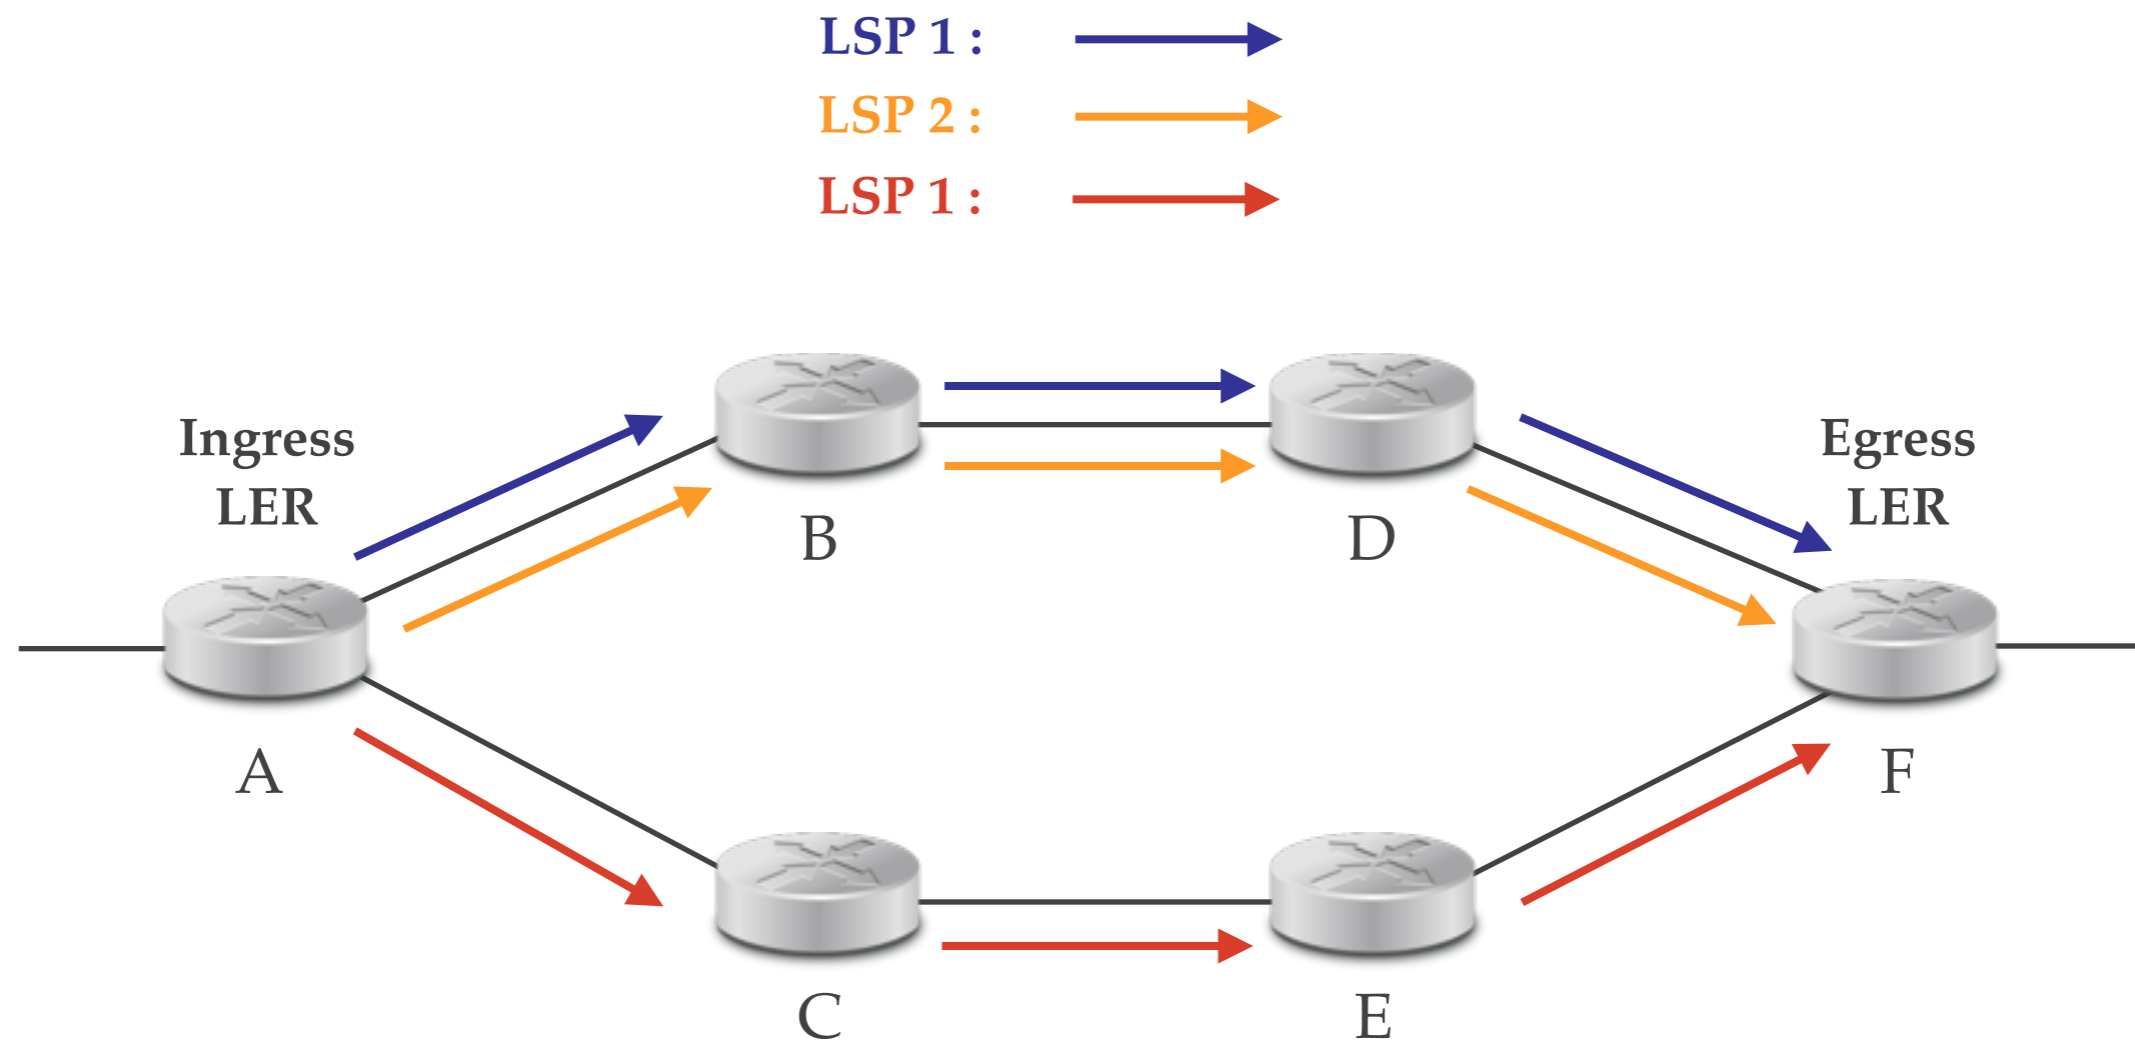
- Motivations
- Measuring MPLS
- LPR
- Evaluation
- Conclusion

Motivations

- Current studies are mainly about MPLS discovery and its impact on topology discovery
 - B. Donnet, M. Luckie, P. Mérindol, J.-J. Pansiot. *Revealing MPLS Tunnels Obscured from Traceroute*. In ACM SIGCOMM Computer Communication Review, 42(2)., pp. 87-93. April 2012
 - J. Sommers, B. Eriksson, P. Barford. *On the Prevalence and Characteristics of MPLS Deployments in the Open Internet*. In Proc. Internet Measurement Conference (IMC). October 2008.
 - T. Flach, E. Katz-Basset, R. Govindan. *Quantifying Violations of Destination-Based Forwarding on the Internet*. Proc. Internet Measurement Conference (IMC). November 2012.
- Actual usage of MPLS by operators not yet studied

Motivations (2)

- Observation
 - several LSPs may exist for a given <Ingress LER, Egress LER>



Motivations (3)


- Those LSPs between <Ingress LER, Egress LER> can represent
 - load balancing (i.e., ECMP) between the ingress and the egress LER
 - traffic-engineering
- We want to distinguish 3 types of MPLS tunnels
 - mono-path
 - transit tunnel without TE
 - transit tunnel with TE
- We focus on
 - *explicit* MPLS tunnels
 - transit tunnels
 - ✓ inter-domain tunnels seems negligible

Measuring MPLS

- The discovery of MPLS can be based on standard active measurement tools ([CCR2012])
 - ping
 - traceroute
- Two options are required
 1. **ICMP extension** ([RFC4950])
 - ✓ if an MPLS router must forge an ICMP `time_exceeded` message, it should quote the MPLS LSE stack in it
 2. **TTL propagate** ([RFC3443])
 - ✓ the ingress LER of an MPLS should initialize the LSE-TTL with the value inside the IP-TTL field (iTTL)
 - ✓ the opposite operation is done by the egress LER (oTTL)

Measuring MPLS (2)

- MPLS Explicit tunnels
 - RFC4950 \wedge RFC3443



Traceroute output:


1. R₁
2. R₂ - *MPLS tag*
3. R₃ - *MPLS tag*
4. R₄ - *MPLS tag*
5. R₅
6. Destination

LPR

- Label Pattern Recognition algorithm
- Allows to distinguish multi-FEC from IP load balancing
- Passive classification method
 - works offline, once the data has been collected
 - requires no additional probing than traceroute
- Recognizes behaviors of LDP vs. RSVP-TE based on MPLS labels distribution and IP addresses
- LPR provides four classes

LPR (2)

- Class 1
 - *MonoLSP*



Trace LSP1:

1. A
2. B - Label L1
3. C - Label L2
4. D

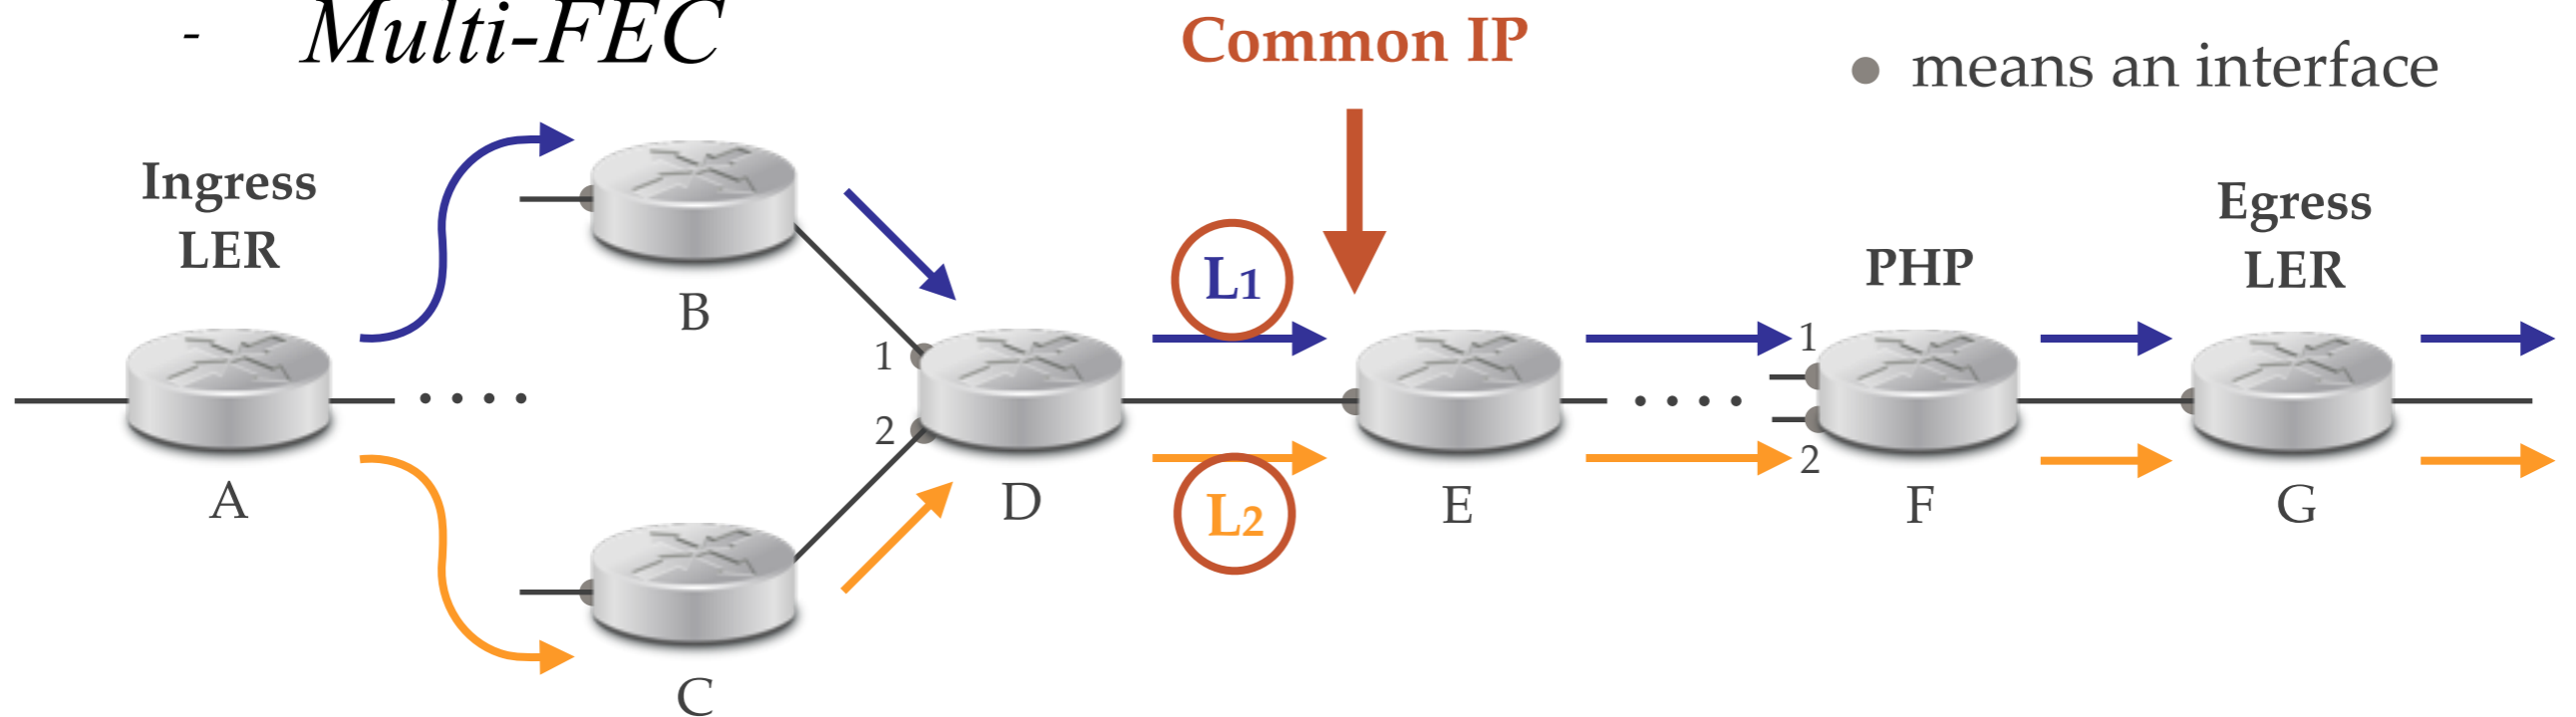
Trace LSP2:

1. A
2. B - Label L1
3. C - Label L2
4. D

Same IP addresses and same Labels

LPR (3)

- Class 2
 - *Multi-FEC*



Trace LSP1:

1. A
2. ...
3. B - Label
4. D₁ - Label
5. E - Label L1
6. ...
7. F₁ - Label
8. G


Trace LSP2:

1. A
2. ...
3. C - Label
4. D₂ - Label
5. E - Label L2
6. ...
7. F₂ - Label
8. G

**Different labels
for at least
1 common IP**

LPR (4)

- Class 3
 - *ECMP Mono-FEC*: disjoint routers
Common IP



Trace LSP1:

1. A
2. ...
3. B - Label
4. D1 - Label L1
5. E - Label L2
6. ...
7. F1 - Label
8. G

Trace LSP2:

1. A
2. ...
3. C - Label
4. D2 - Label L1
5. E - Label L2
6. ...
7. F2 - Label
8. G


Same label
∀ common IPs

LPR (5)

- Class 3 (cont.)

- *ECMP Mono-FEC*: parallel links

Same labels along all the LSPs



Different IPs are aliases!

Trace LSP1:

1. A
2. ...
3. B - Label L1
4. C₁ - Label L2
5. D₁ - Label L3
6. ...
7. E₁ - Label L4
8. F

Trace LSP2:

1. A
2. ...
3. B - Label L1
4. C₂ - Label L2
5. D₂ - Label L3
6. ...
7. E₂ - Label L4
8. F

LPR (6)

- Class 4
 - *Unclassified*
- If PHP is used, the Egress LER does not exhibit labels
- It may happen that LSPs do not intersect on a common IP address
- Those tunnels are arbitrarily tagged as unclassified

Evaluation

- Archipelago platform
- First traceroute cycle of each month since 2010 until December 2014
- 60 cycles


Evaluation (2)

- Numbers of tunnels



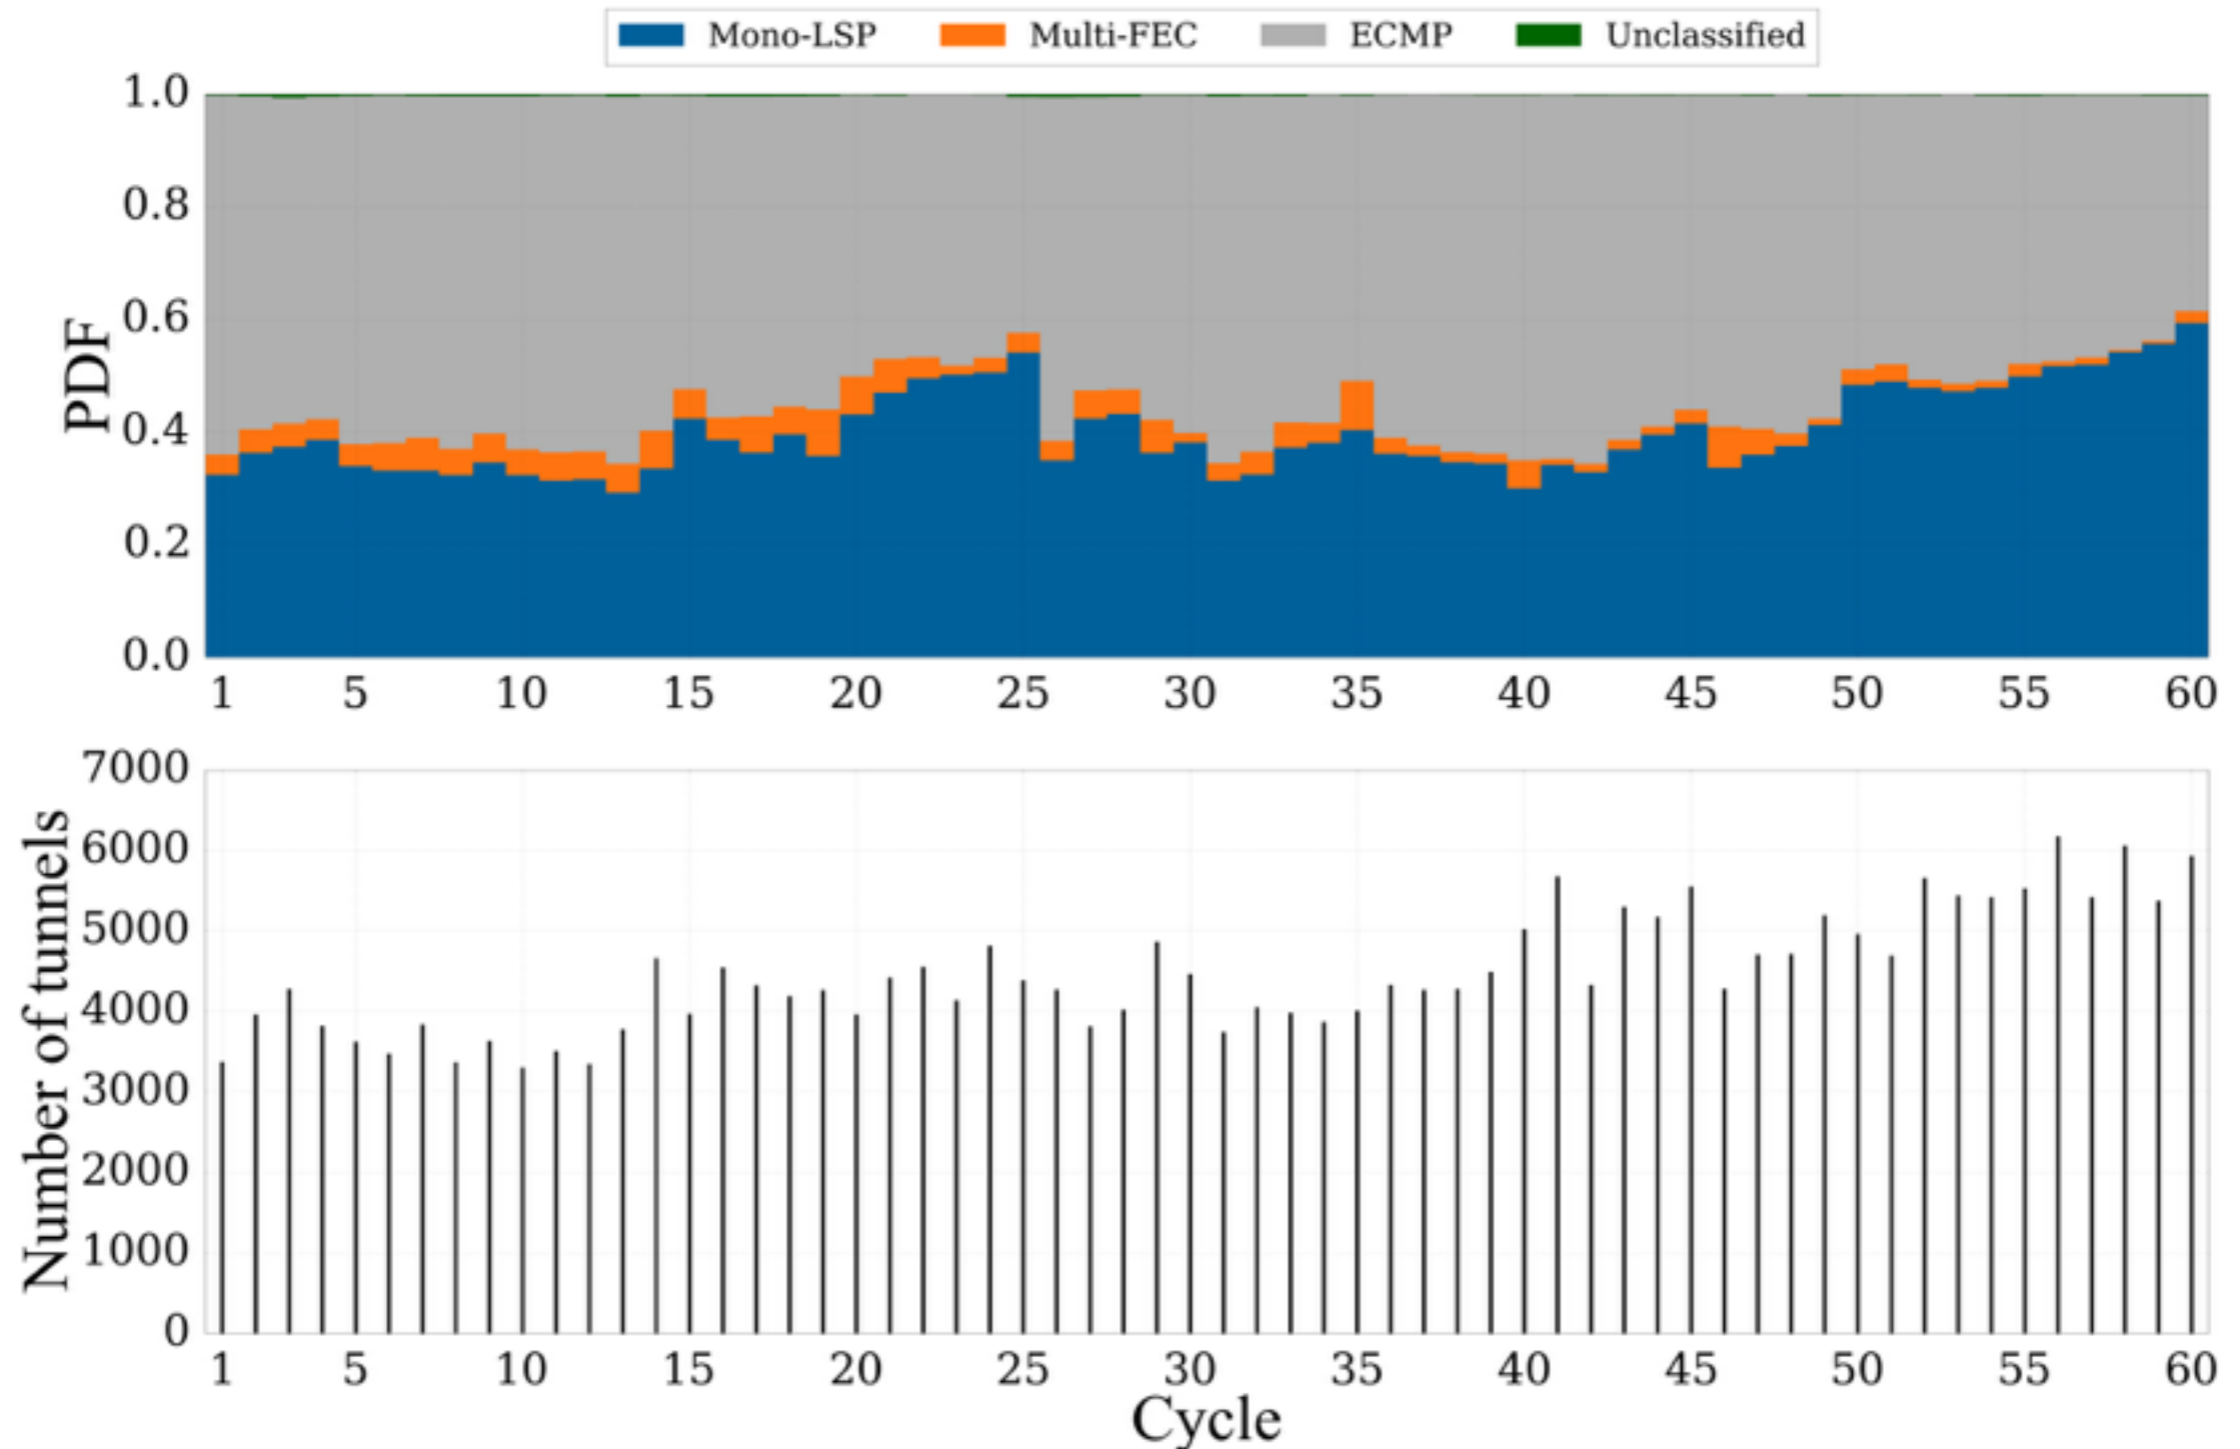
Evaluation (3)

- AS6453



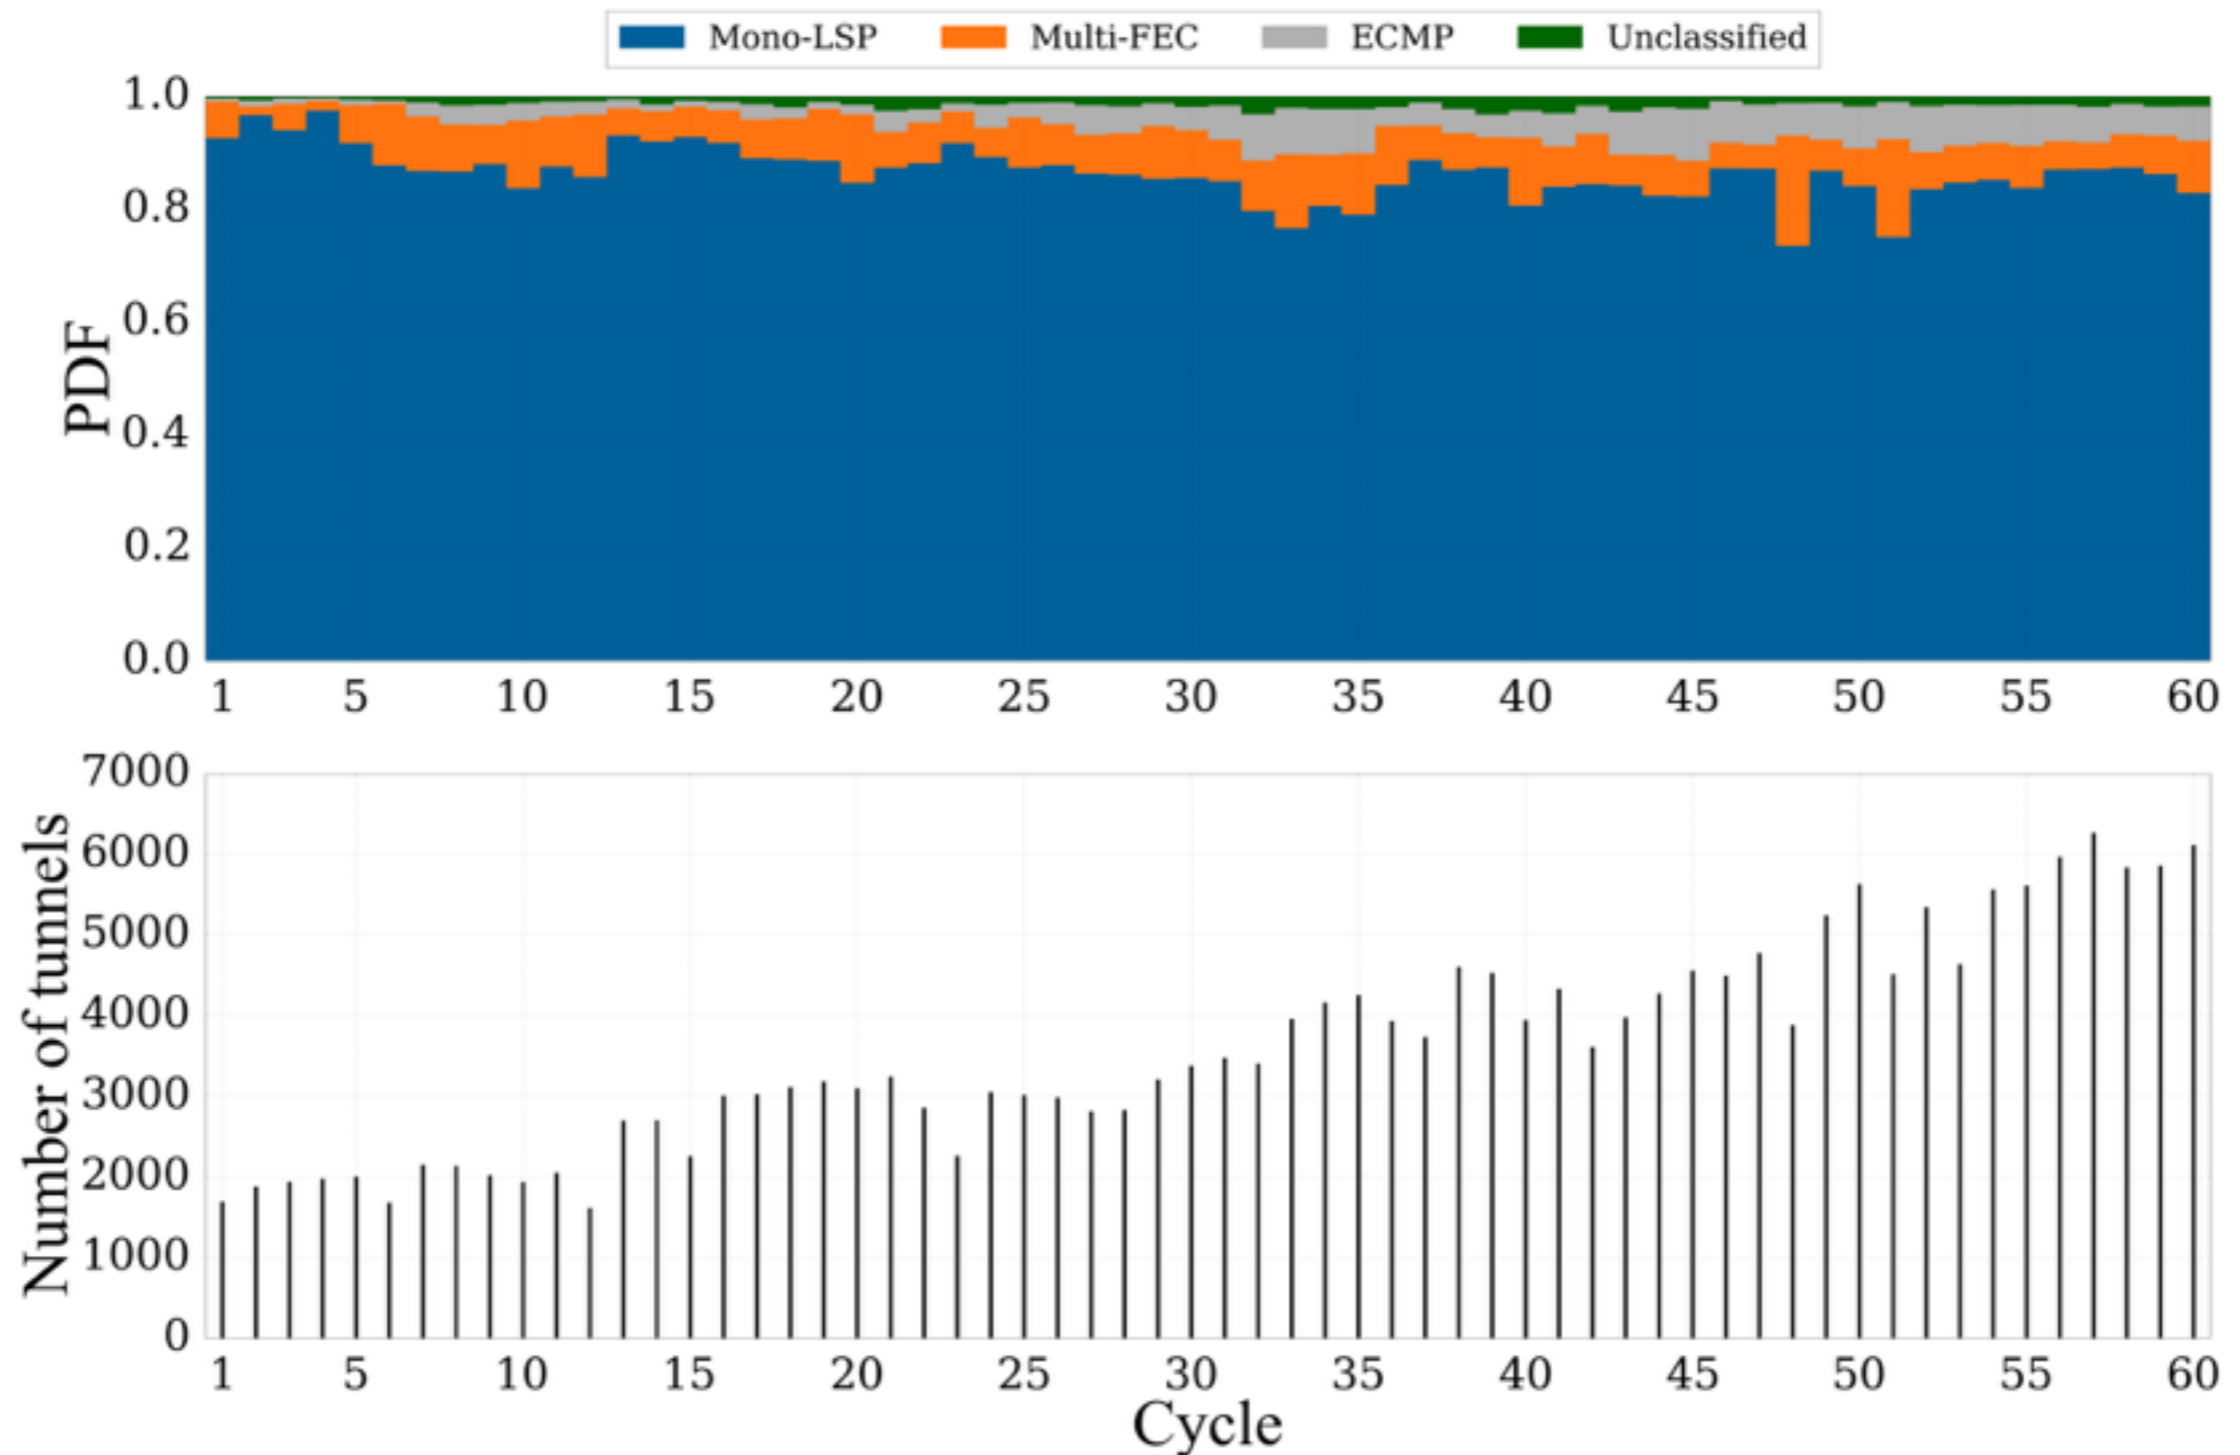
Evaluation (4)

- AS1273




Evaluation (5)

- AS2914



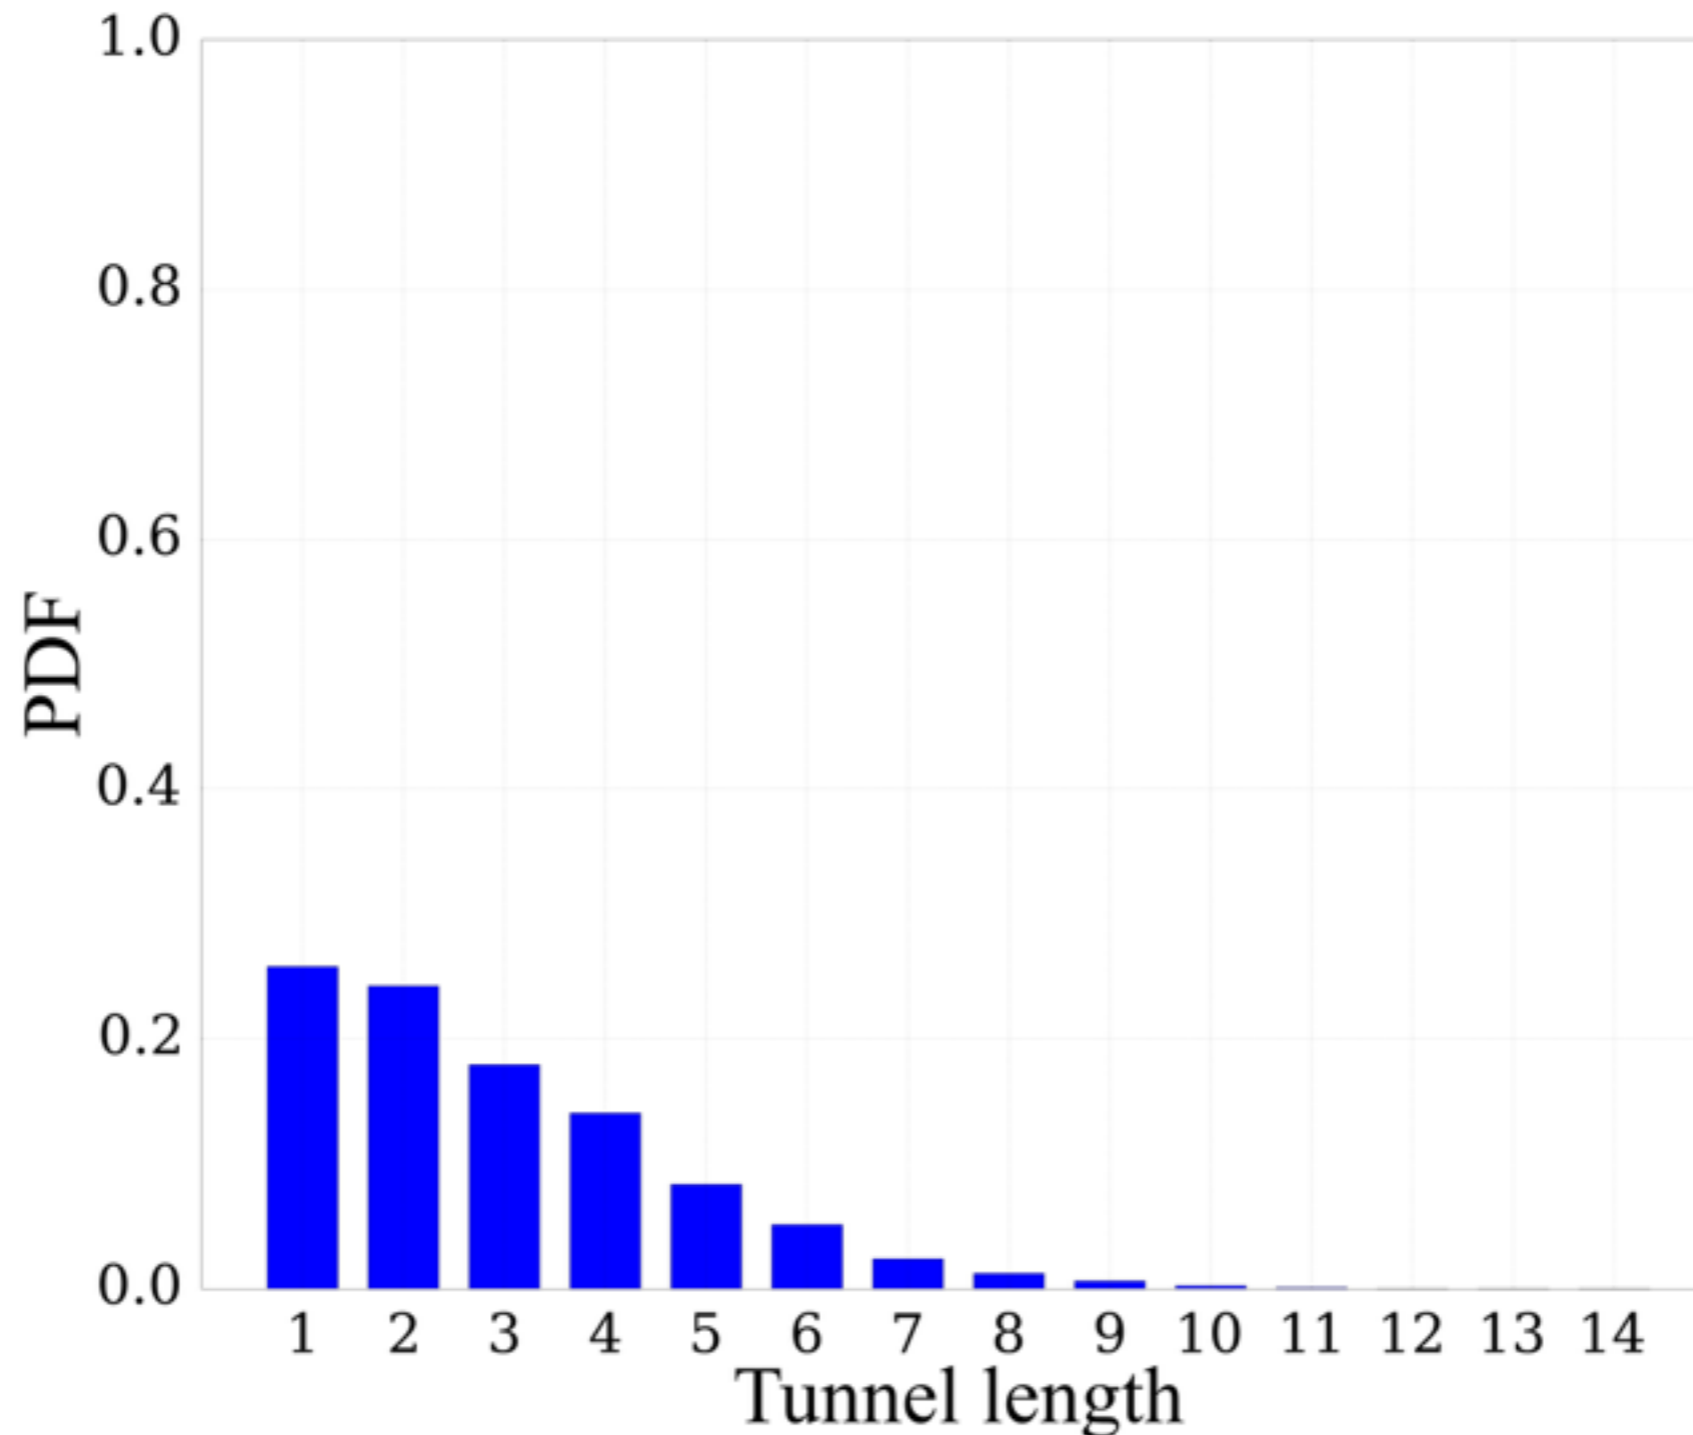
Evaluation (6)

- AS286




Evaluation (7)

- Tunnel *length*



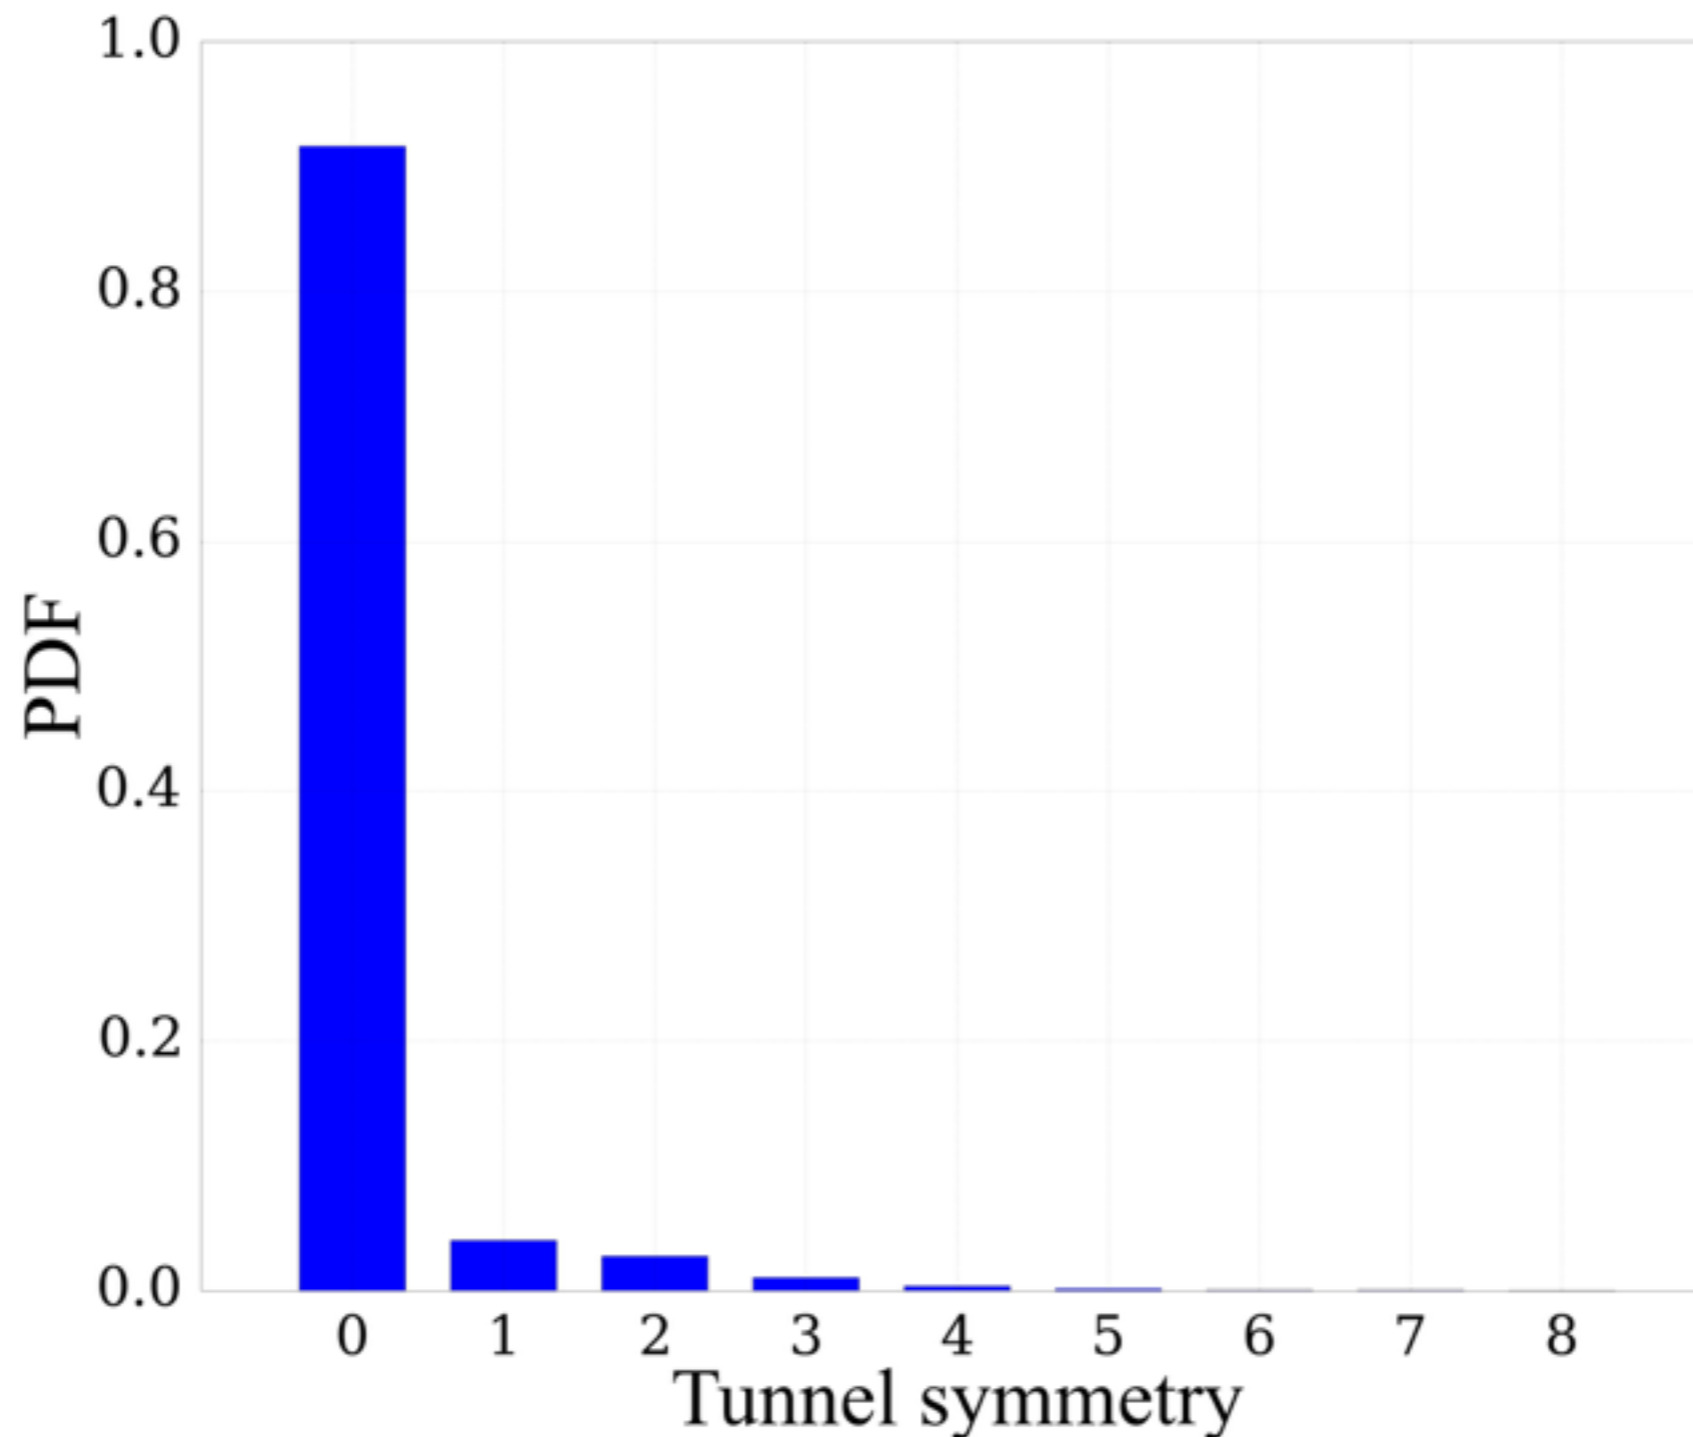
Evaluation (8)

- Tunnel *width*



Evaluation (9)

- Tunnel *symmetry*



Conclusion

- New algorithm to reveal TE usage within ASes
 - label distribution
 - ✓ Mono- or Multi-FEC
 - ECMP load balancing
 - ✓ parallel links or disjoint routers
 - dynamics
 - ✓ temporal evolution of MPLS deployment and usage

Conclusion (2)

- Next steps
 - deeper investigation of Multi-FEC class
 - ✓ high frequency traceroute to observe labels behavior
 - deeper investigation of ECMP class
 - ✓ Paris Traceroute mda mode
 - providing a library to automatically export MPLS tunnels usage from CAIDA dataset