EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

FACULTY OF
SCIENCE
Communication Networks

P4-Based Implementation of BIER and BIER-FRR for Efficient Multicast

Daniel Merling, **Steffen Lindner**, Michael Menth

http://kn.inf.uni-tuebingen.de

The Need for BIER

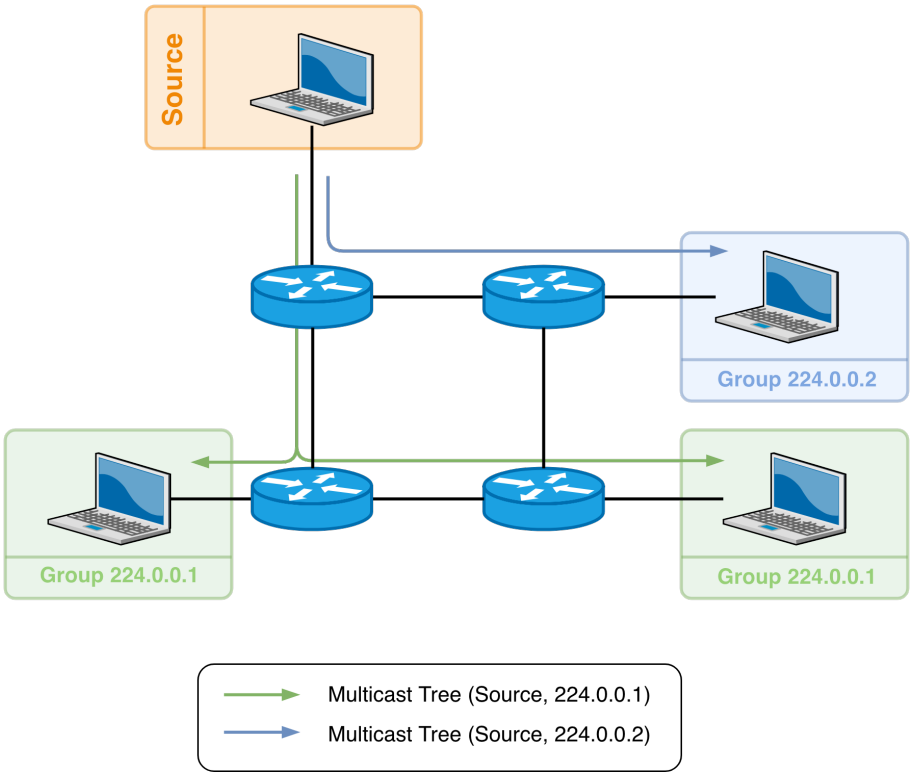

- IP multicast is widely used to address multiple receivers
- Hosts can subscribe and unsubscribe leveraging IGMP
- Connected routers propagate this information using multicast routing protocols, e.g., PIM

- Problem
  - Intermediate routers need to store (S, G) state
  - All routers need to recalculate their states upon group changes

Solution: Bit Index Explicit Replication (BIER)

BIER – Concept

- Domain concept

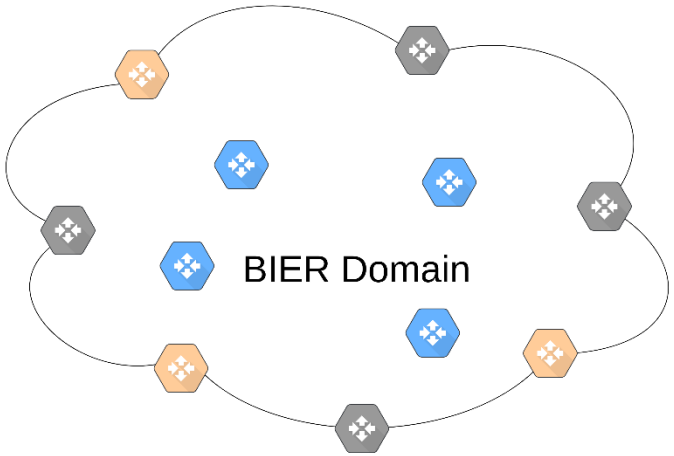

BIER – Concept

- Domain concept
  - Ingress nodes (BFIR)
    - Add BIER header
      - Contains all destinations of the packet

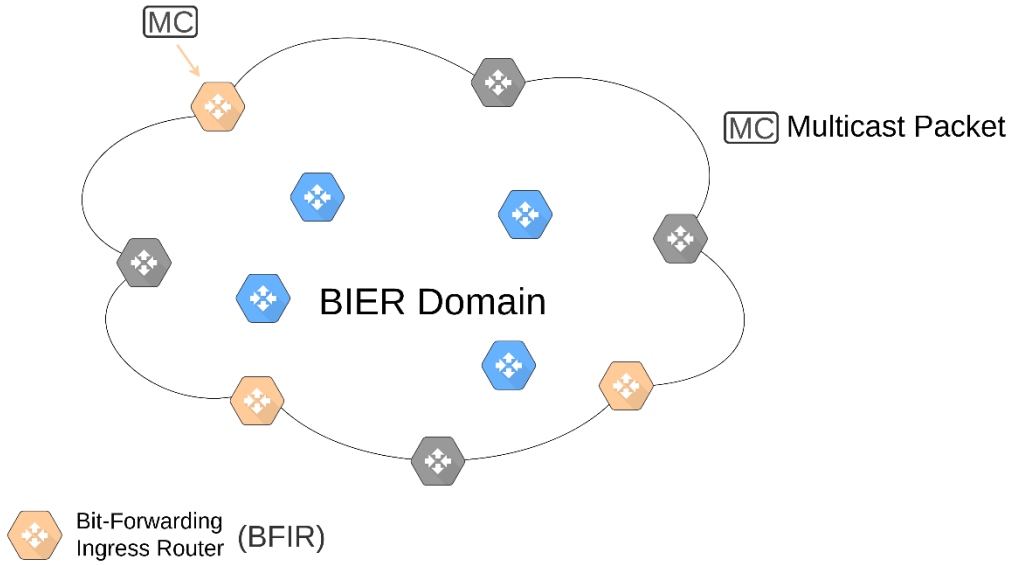

BIER – Concept

- Domain concept
  - Ingress nodes (BFIR)
    - Add BIER header
      - Contains all destinations of the packet
  - Core nodes (BFR)
    - Forward and replicate packets on paths from the routing underlay (e.g., IGP)
    - Distribution on tree structure

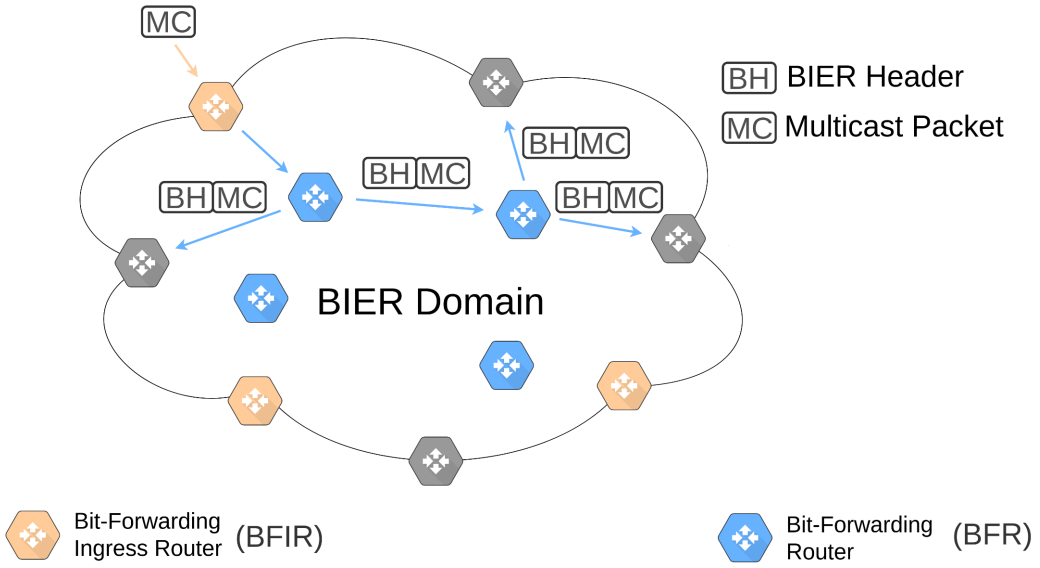

BIER – Concept

- Domain concept
  - Ingress nodes (BFIR)
    - Add BIER header
      - Contains all destinations of the packet
  - Core nodes (BFR)
    - Forward and replicate packets on paths from the routing underlay (e.g., IGP)
    - Distribution on tree structure
  - Egress nodes (BFER)
    - Remove BIER header

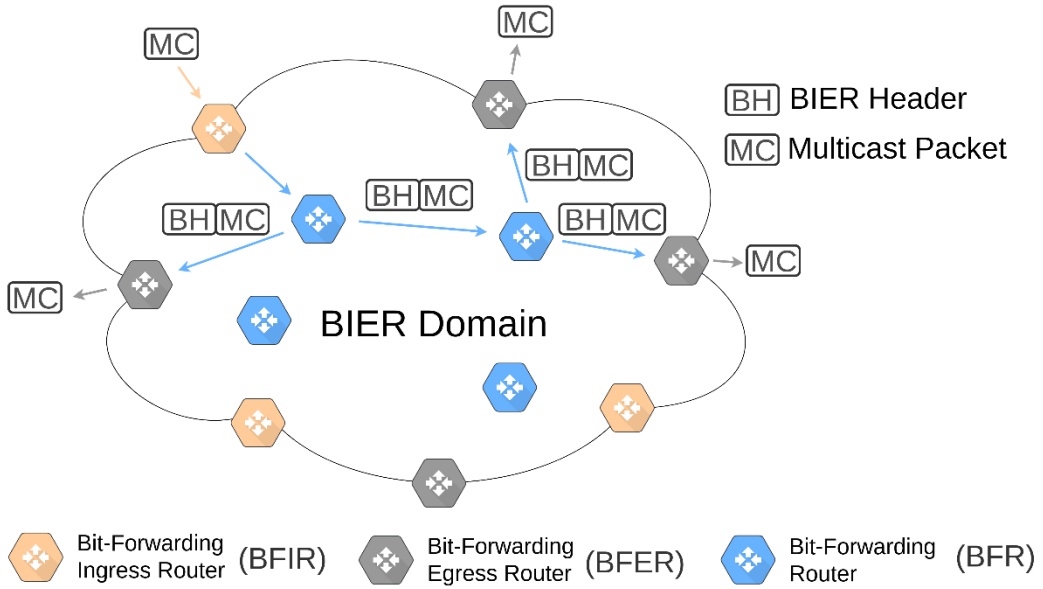

BIER Forwarding

- BIER header
  - Contains bit string where each BFER is assigned to a position
    - If BFER should receive a packet copy, its bit is activated in the packet header
  - Individual BIER header for each IPMC group
  - Before a BFRs forwards a packet to a NH, it clears bits of BFERs that are reached via other NHs from the packet header to avoid duplicates

BIFT of BFR 1

BFER	Next-Hop	Forwarding Bitmask (F-BM)
2	2	100110
3	3	100110
4	4	011000
5	4	011000
6	2	100110

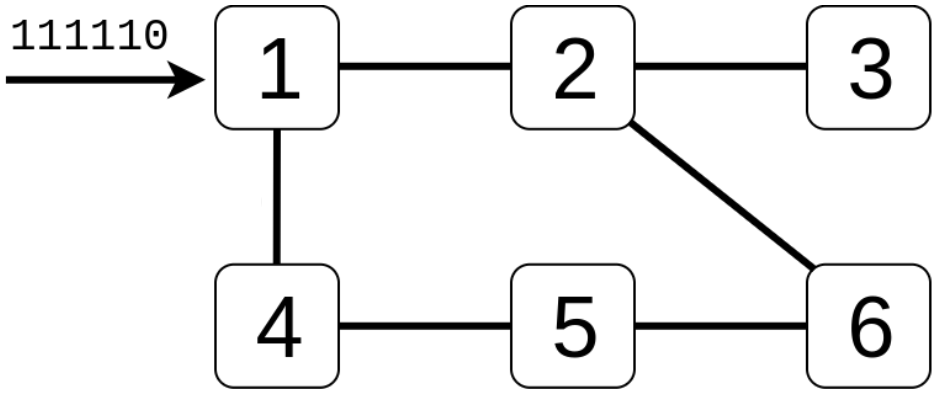

BIER Forwarding

- BIER header
  - Contains bit string where each BFER is assigned to a position
    - If BFER should receive a packet copy, its bit is activated in the packet header
  - Individual BIER header for each IPMC group
  - Before a BFRs forwards a packet to a NH, it clears bits of BFERs that are reached via other NHs from the packet header to avoid duplicates

BIFT of BFR 1

BFER	Next-Hop	Forwarding Bitmask (F-BM)
2	2	100110
3	3	100110
4	4	011000
5	4	011000
2	4	100110

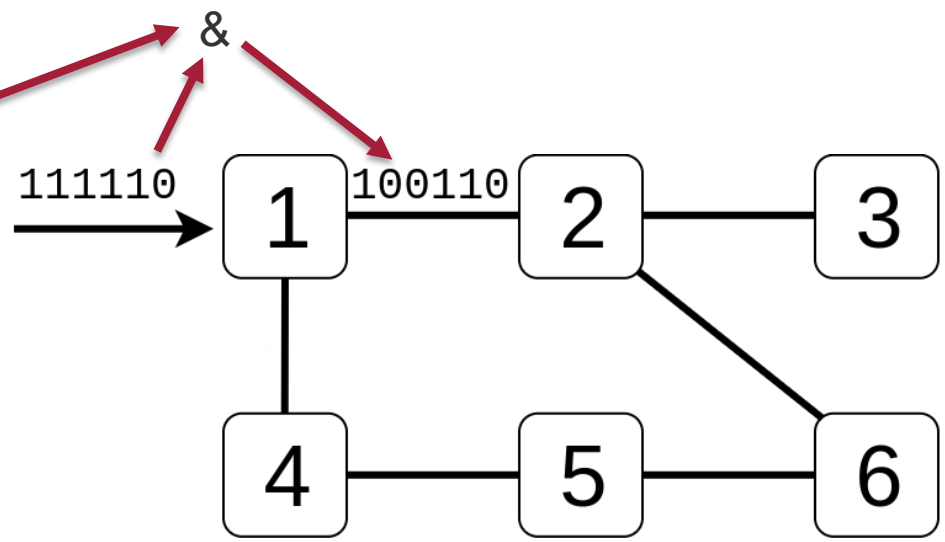

BIER Forwarding

- BIER header
  - Contains bit string where each BFER is assigned to a position
    - If BFER should receive a packet copy, its bit is activated in the packet header
  - Individual BIER header for each IPMC group
  - Before a BFRs forwards a packet to a NH, it clears bits of BFERs that are reached via other NHs from the packet header to avoid duplicates

BIFT of BFR 1

BFER	Next-Hop	Forwarding Bitmask (F-BM)
2	2	100110
3	3	100110
4	4	011000
5	4	011000
6	2	100110

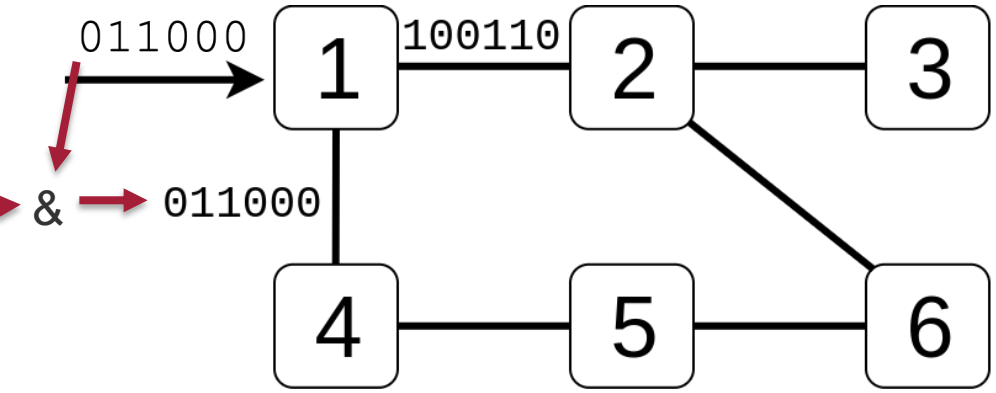

BIER - Properties

- BIER header
  - Contains bit string where each BFER is assigned to a position
    - If BFER should receive a packet copy, its bit is activated in the packet header
  - Individual BIER header for each IPMC group
  - Before a BFRs forwards a packet to a NH, it clears bits of BFERs that are reached via other NHs from the packet header to avoid duplicates
- Iterative forwarding mechanism
  - In each iteration one packet copy is forwarded to a NH
- Scalability
  - No dynamic state within core devices
  - Signaling only to BFIRs

BIER Implementation (I)

- Motivation for BIER implementation
  - Solves issues of traditional multicast (state, signaling, …)
  - Significant support from several global players (Cisco, Juniper, Nokia, Google, …)
  - Yet, there is no working BIER implementation!
    - Evaluation, operability, …
- Implemented with P4 on Intel Tofino ASIC
  - Runs at 100 Gb/s
  - Published in IEEE Access

BIER Implementation (II)

- Why P4?
  - Efficient implementation with legacy devices not possible
  - P4 offers required flexibility
    - Define new header
    - Define processing pipeline
    - Iterative forwarding procedure
  - Easier management with separated control plane

BIER Implementation (III)

- In each iteration, one next-hop is served
  - Clone packet & recirculate

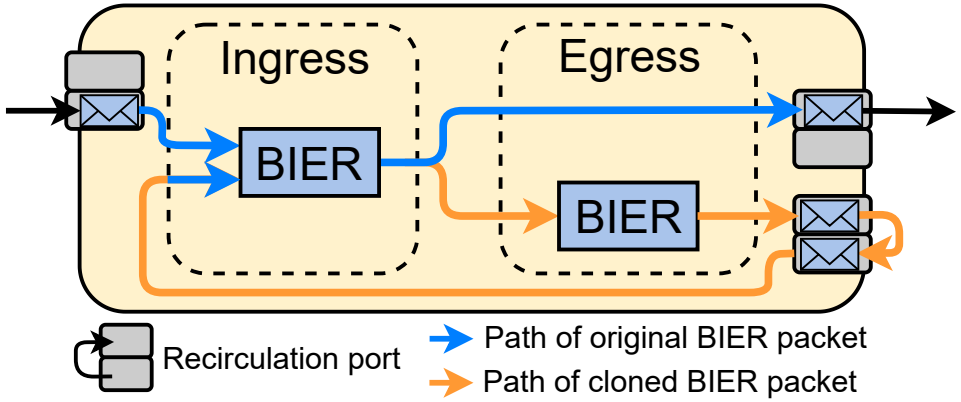

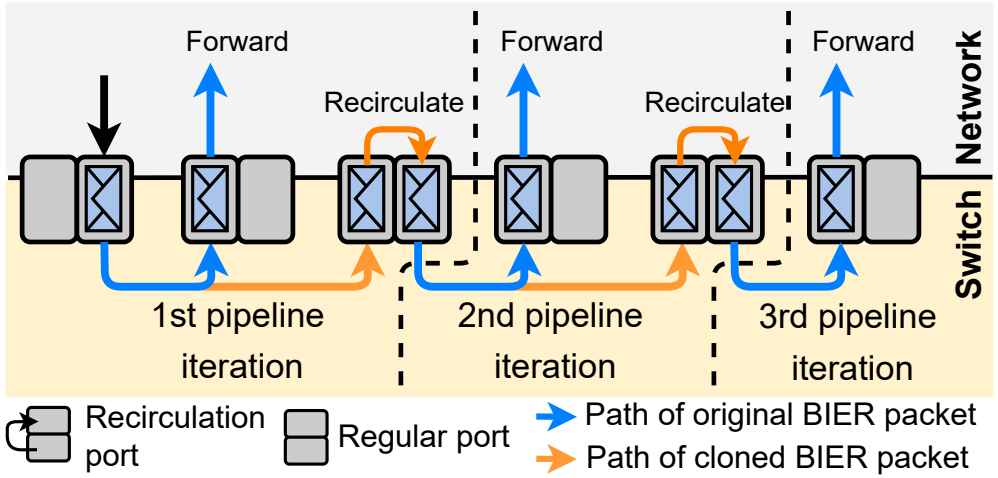

BIER Implementation (IV)

Problem: Recirculation requires capacity

- 100 Gbit/s multicast traffic with 5 next-hops results in 400 Gbit/s recirculation traffic
- Solution: Add dedicated recirculation ports to increase recirculation capacity

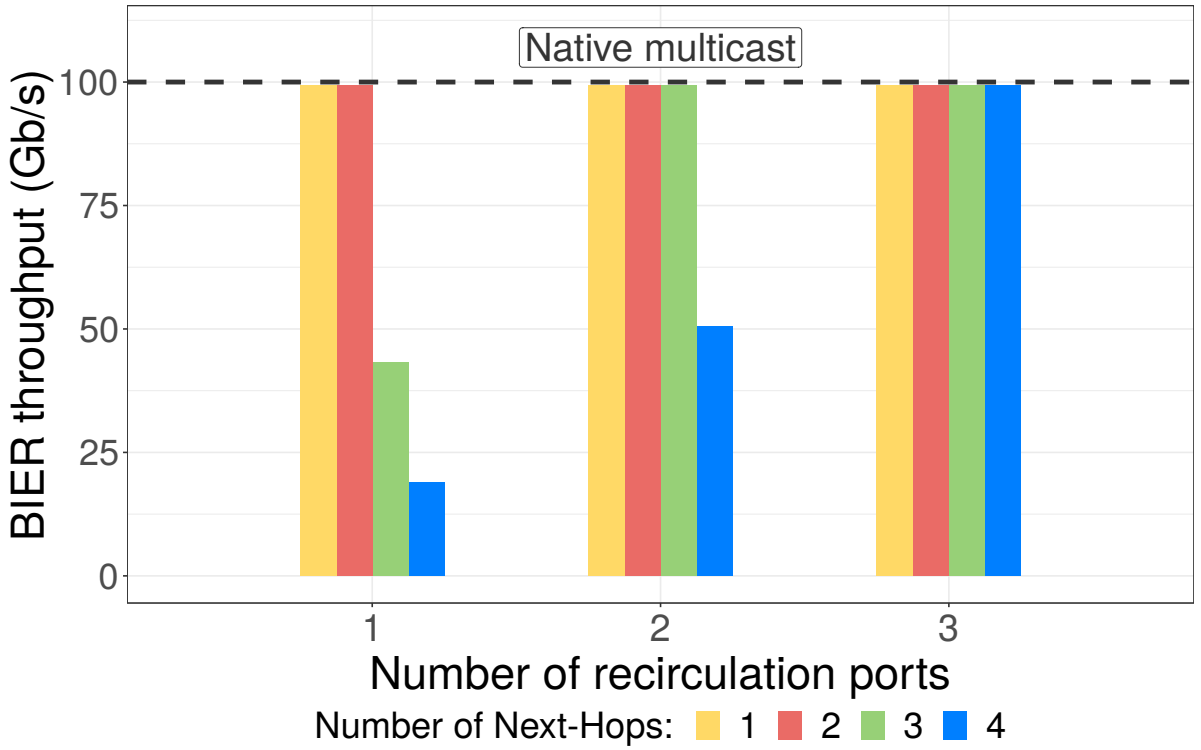

BIER-FRR

- Fast Reroute (FRR) deviates traffic around a local failure
  - Link failure
  - Node failure

- No native FRR support in P4
  - Tofino generates a special packet when ports are up/down
  - We store this information in registers to apply FRR

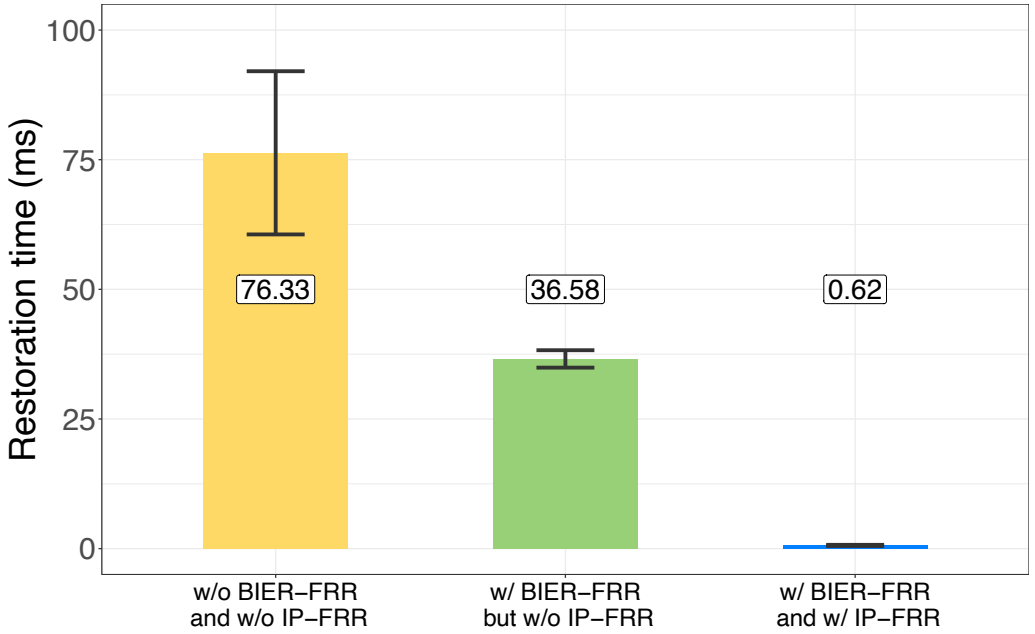

Steffen Lindner

University of Tuebingen
Faculty of Science
Department of Computer Science
Chair of Communication Networks