



BIER Fast ReRoute

H. Chen, M. McBride, **S. Lindner**, M. Menth, A. Wang, G. Mishra, Y. Liu, Y. Fan, L. Liu, X. Liu

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



- ▶ BIER-FRR Architecture
 - Definitions
 - Operation
 - Organization of Backup BIFT(s)
- ▶ Usage
 - Protection Levels
 - Backup Strategies
- ▶ Comparison



- ▶ Definition: classification of connections between BFR neighbors (BFR-NBRs)
 - Direct connection (i.e., via the link layer)
 - Indirect connection (i.e., tunnel via routing underlay)

- ▶ Definition: forwarding actions for FRR
 - Plain: send the mere BIER packet to a BFR-NBR via a direct connection
 - Tunnel: encapsulate the BIER packet with a tunnel header towards a BFR-NBR (indirect connection)
 - Explicit: forward the packet over an explicit path to a BFR-NBR, e.g., through segment routing
 - Segment IDs (in case of SR) must be given

(Needed because some nodes may be reached w/ Plain under failure-free conditions while they require a tunnel to be reached in case of failures.)

► Definition: backup forwarding entries

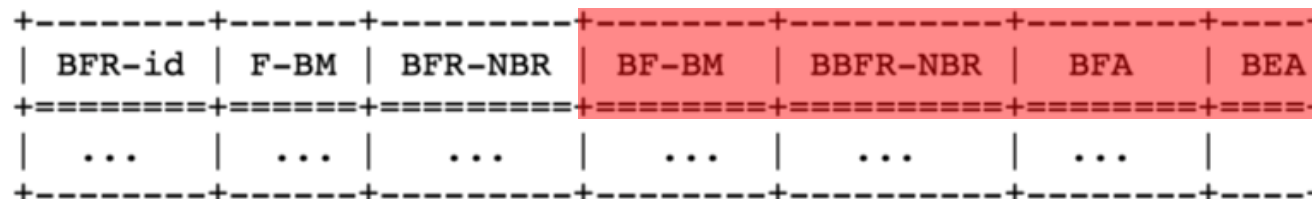
- Backup F-BM (BF-BM)
- Backup BFR-NBR (BBFR-NBR)
- Backup forwarding action (BFA)
- Backup entry active (BEA) flag

► Computation: backup F-BM of a specific BFER indicates

- All BFERs that share the primary and backup BFR-NBR of the BFER and
- All BFERs that have the backup BFR-NBR of the specific BFER as primary BFR-NBR

► Definition: extended BIFT for BIER-FRR

- Backup entries depend on operation mode





- ▶ Activation and Deactivation of FRR
 - Upon failure detection (e.g., via BFD)
 - BEA flag set in backup forwarding entry for affected BFER
 - If the primary BFR-NBR is reachable again, the BEA flag is deactivated
 - Failure no longer present
 - Change of the primary BFR-NBR through reconvergence

- ▶ Forwarding operation for BFER
 - If BEA set, apply backup forwarding entry
 - Apply backup F-BM and send to backup BFR-NBR via the backup forwarding action
 - Otherwise, apply primary F-BM and forward packet

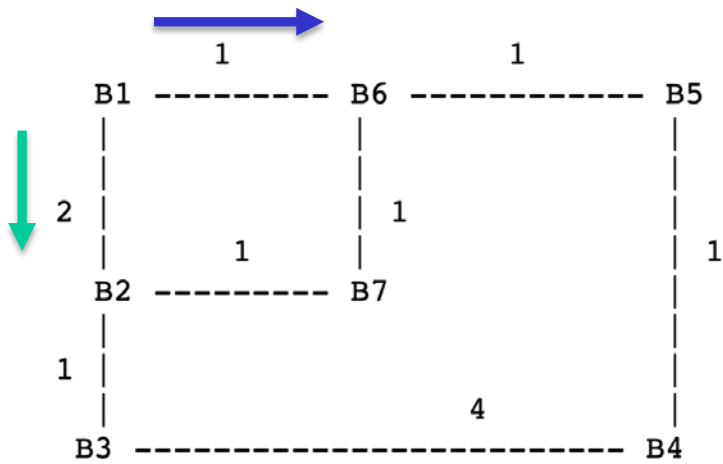
- ▶ Important: apply all active backup forwarding entries first!



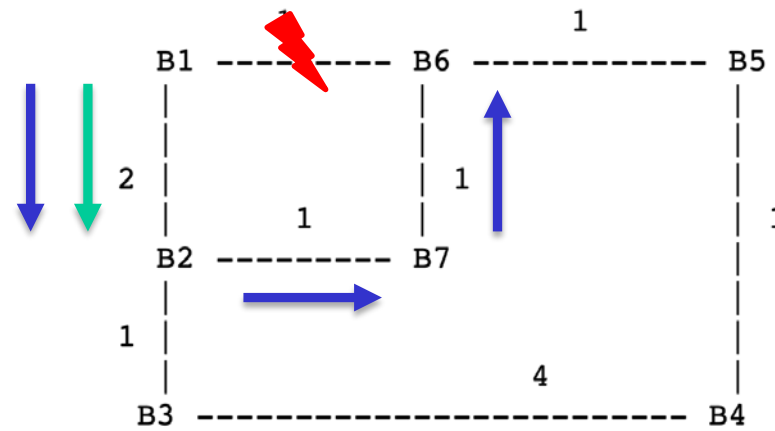
Prioritization of Backup Forwarding Entries (1)

► BIER forwarding from B1 to {B2, B6}

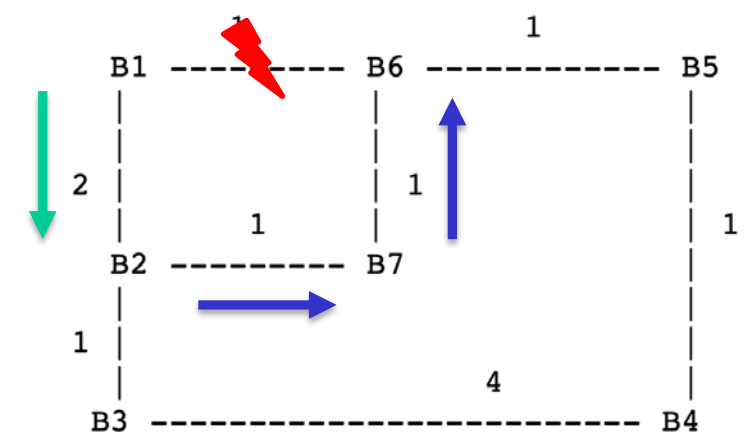
failure-free case



failure case w/o prioritization



failure case w/ prioritization





► Two (easier?) variants

- Primary BIFT and single backup BIFT
- Primary BIFT and failure-specific backup BIFTs (for each BFR-NBR)
(Preferred variant may be platform-specific)

► Operation

- First, process backup BIFT entries with BEA activated
- Then, process primary (regular) BIFT

BFR-id	F-BM	BFR-NBR
2	0000110	B2
3	0000110	B2
4	1111000	B6
5	1111000	B6
6	1111000	B6
7	1111000	B6

B1's primary (regular) BIFT.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects failure of
2	1111110	B6	Plain		Link B1->B2
3	1111110	B6	Plain		Link B1->B2
4	1111110	B2	Plain		Link B1->B6
5	1111110	B2	Plain		Link B1->B6
6	1111110	B2	Plain		Link B1->B6
7	1111110	B2	Plain		Link B1->B6

B1's (single) backup BIFT.



▶ Two (easier?) variants

- Primary BIFT and single backup BIFT
- **Primary BIFT and failure-specific backup BIFTs (for each BFR-NBR)**

(Preferred variant may be platform-specific)

▶ Operation

- Activate backup BIFT instead of single backup entry; process activated backup BIFT
- If no backup BIFT activated, process primary (regular) BIFT

BFR-id	F-BM	BFR-NBR
2	0000110	B2
3	0000110	B2
4	1111000	B6
5	1111000	B6
6	1111000	B6
7	1111000	B6

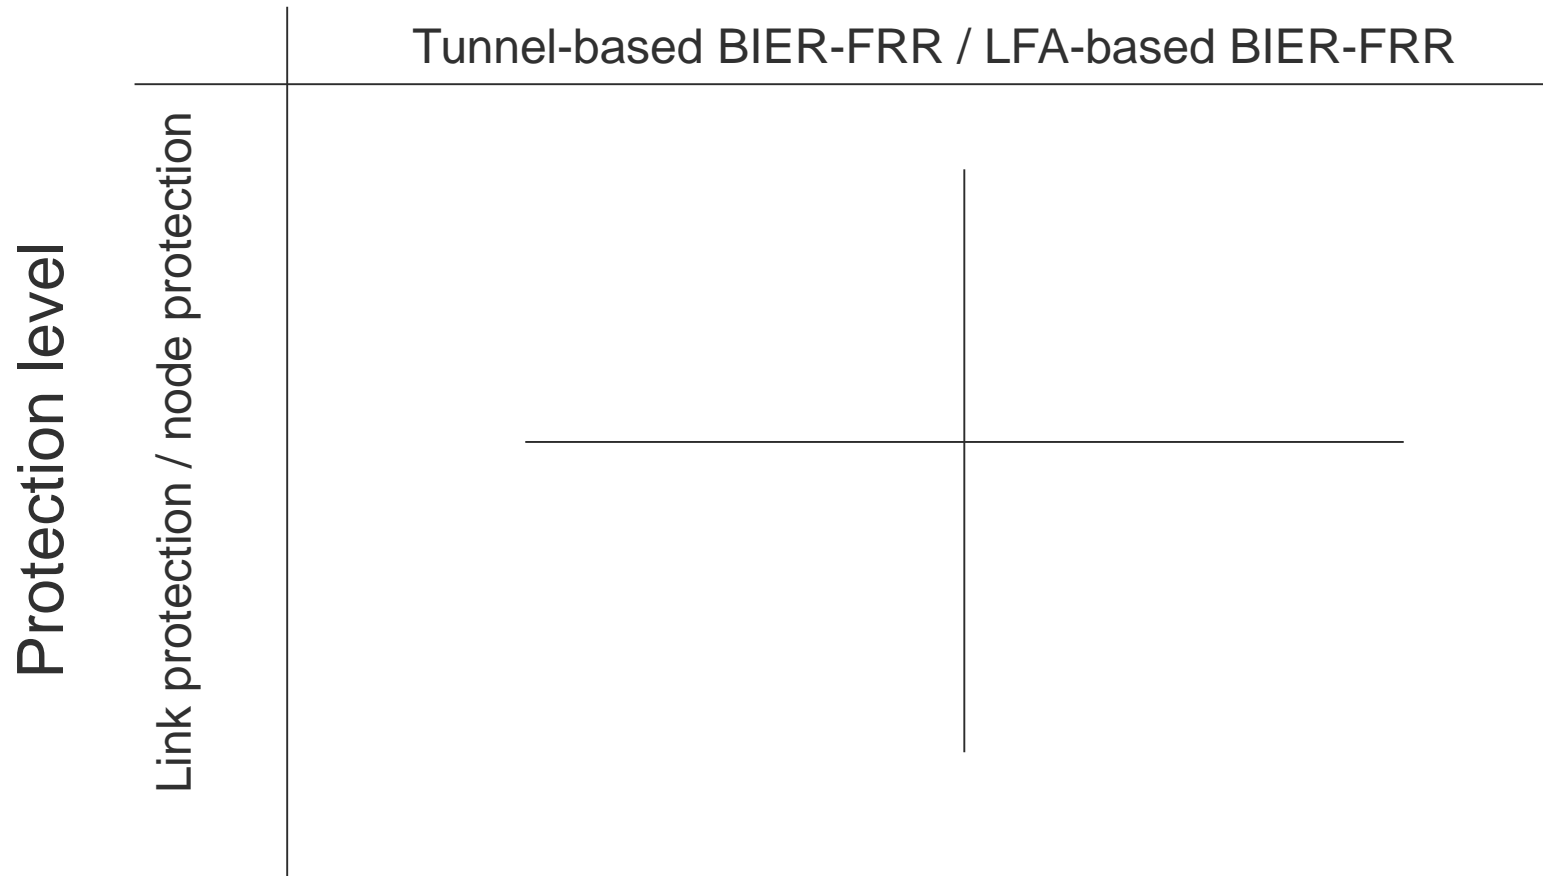
B1's primary (regular) BIFT.

BFR-id	F-BM	BFR-NBR	BFA
2	1111110	B2	Plain
3	1111110	B2	Plain
4	1111110	B2	Plain
5	1111110	B2	Plain
6	1111110	B2	Plain
7	1111110	B2	Plain

B1's failure-specific backup BIFT for failure of B6.



Backup strategy





▶ Tunnel-based BIER-FRR

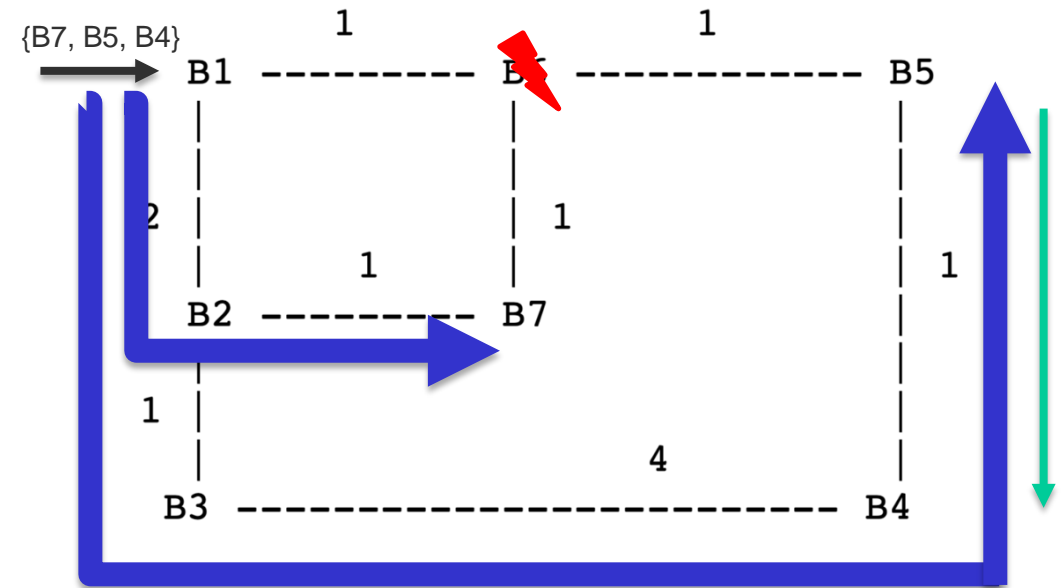
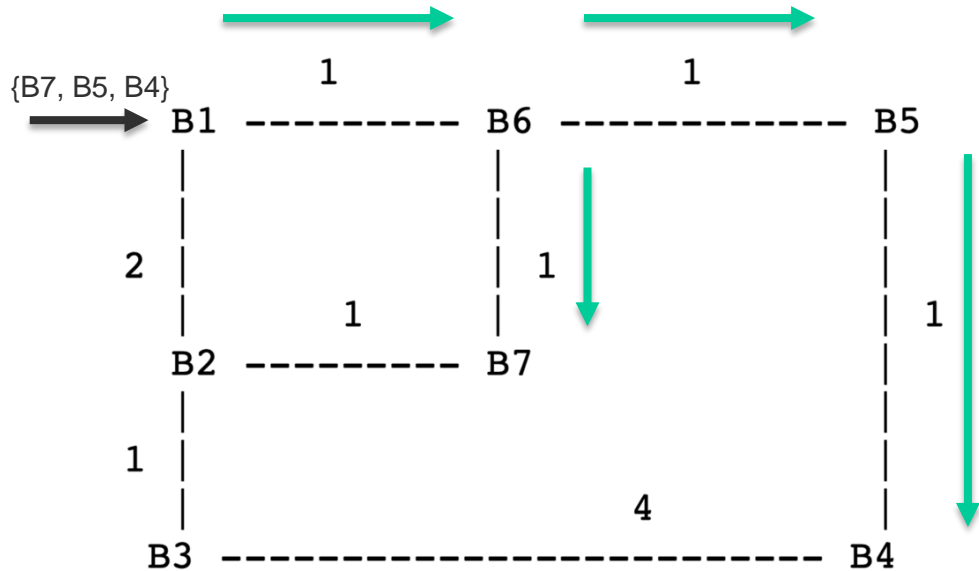
- Leverages FRR capabilities of the routing underlay
- Reaches BFR-NBRs through a tunnel via routing underlay
 - Link protection: use next hop BFR
 - Node protection: use next hop BFR and next-next hop BFRs

▶ LFA-based BIER-FRR

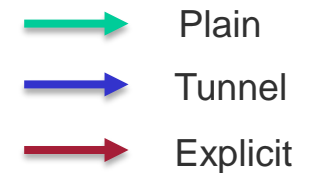
- Leverages LFA principle
 - Well-known from IP
 - But IP-LFA \neq BIER-LFA: LFA node must be a BFR!
- Backup Forwarding Actions for BIER-LFAs
 - Normal BIER-LFA: Plain
 - Remote BIER-LFA: Tunnel
 - Topology-independent (TI) BIER-LFA: Explicit



Example for Tunnel-based BIER-FRR

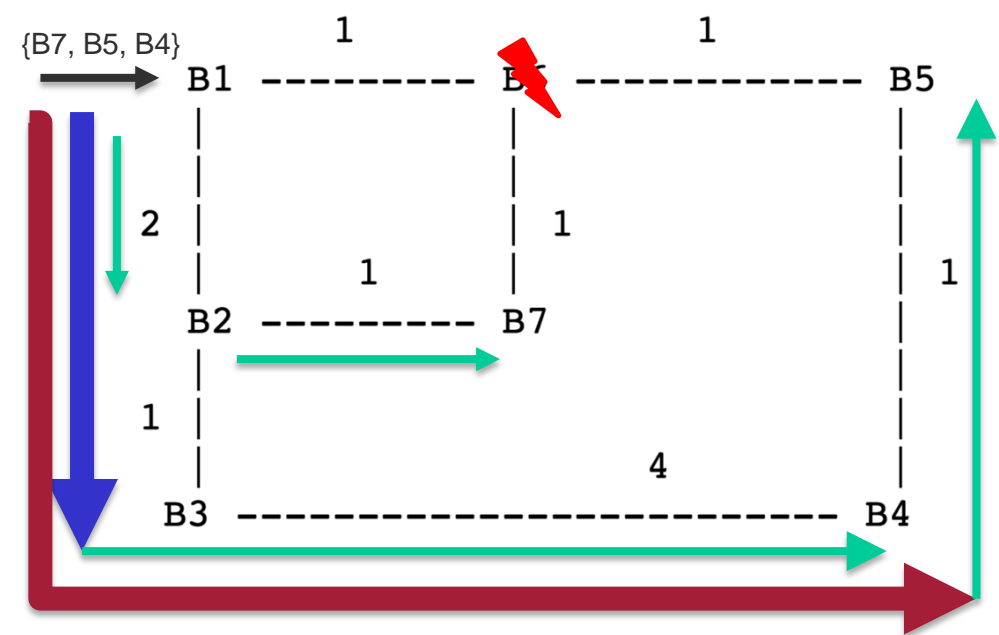
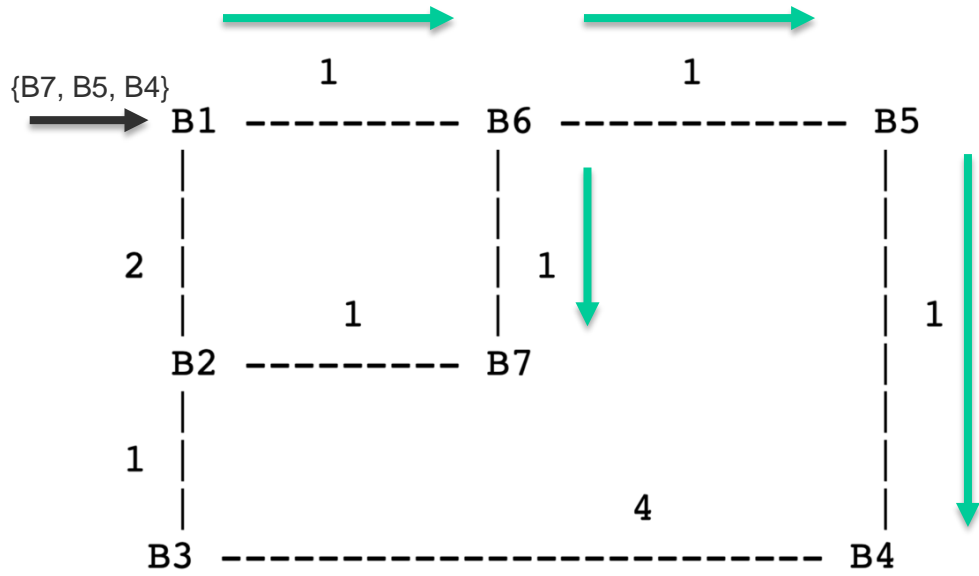


- ▶ Packet delivered to {B7, B4, B4}
- ▶ One redundant packet on link B1-B2
- ▶ Longer backup path for B4

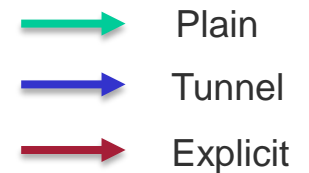




Example for LFA-based BIER-FRR



- ▶ Packet delivered to {B7, B4, B4}
- ▶ Two redundant packets on link B1-B2



► Advantages

- Re-uses existing FRR of routing underlay
- Simple computation of backup forwarding entries
 - Derived from primary BIFTs of a PLR and its BFR-NBRs
- Forwarding action Explicit not needed

► Disadvantages

- Requires FRR mechanism w/ desired protection level on the underlay
- Header overhead due to tunnel header
- Potential path stretch due to triangular routing via next hop BFR or next-next hop BFR
- May cause redundant packets on backup link in spite of prioritization

► Advantages

- Does not rely on FRR in the underlay
- Can provide better protection on the BIER layer than on the IP layer
 - Example: node protection on BIER layer while no protection or FRR w/ link protection on IP layer
- Avoids header overhead for normal BIER-LFAs

► Disadvantages

- Computation more complex
 - Requires routing information from the underlay
 - But reuse of IP-LFAs if all IP nodes are BFRs
- Remote and TI-BIER-LFAs require header overhead
- May cause redundant packets on backup link in spite of prioritization
 - But fewer than tunnel-based BIER-FRR
- Explicit paths needed for TI-LFAs (complex)
 - More remote and TI-LFAs needed than for IP-FRR if some nodes are not BFRs



▶ Comments welcome

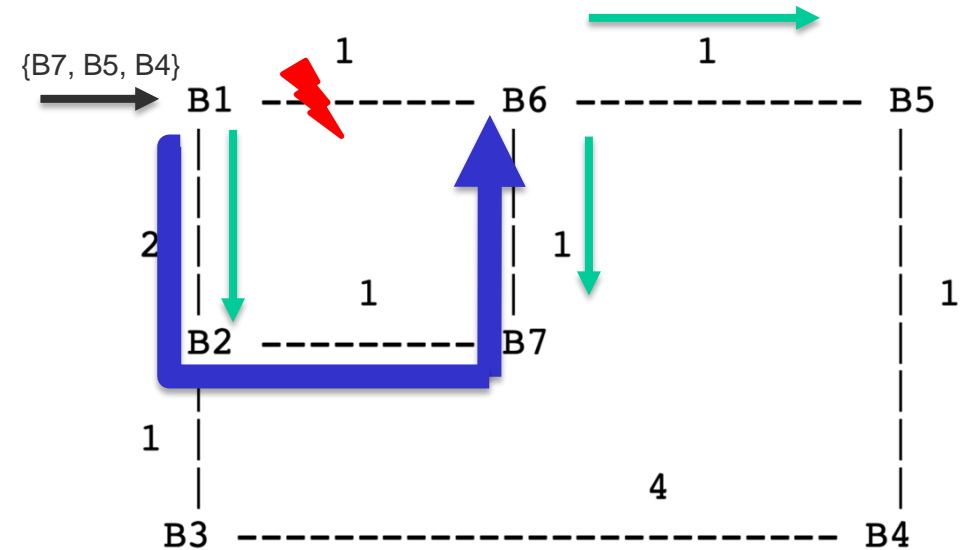
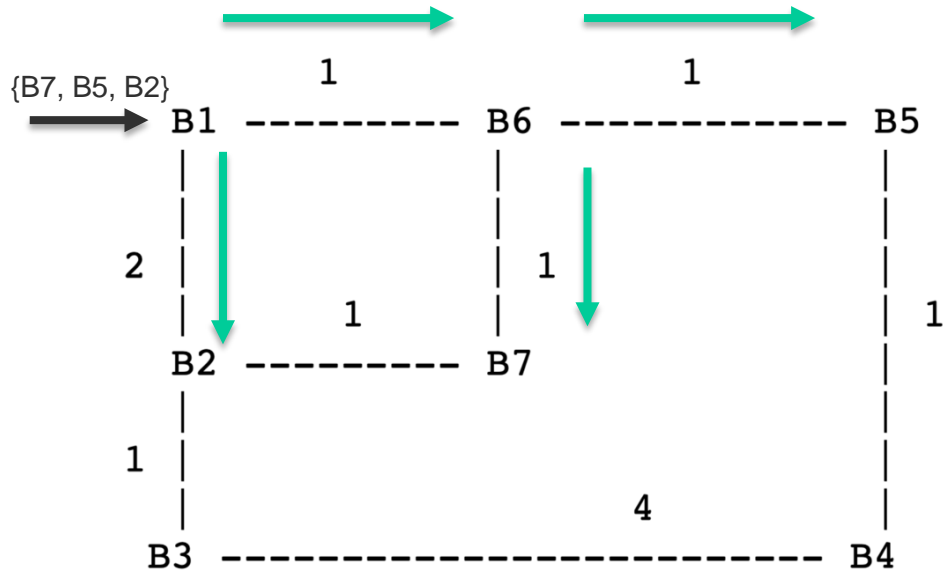
▶ Request for Adoption



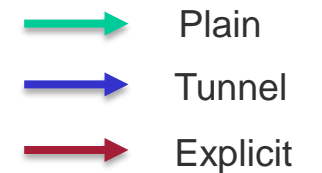
Backup



Example for Tunnel-based BIER-FRR (2)

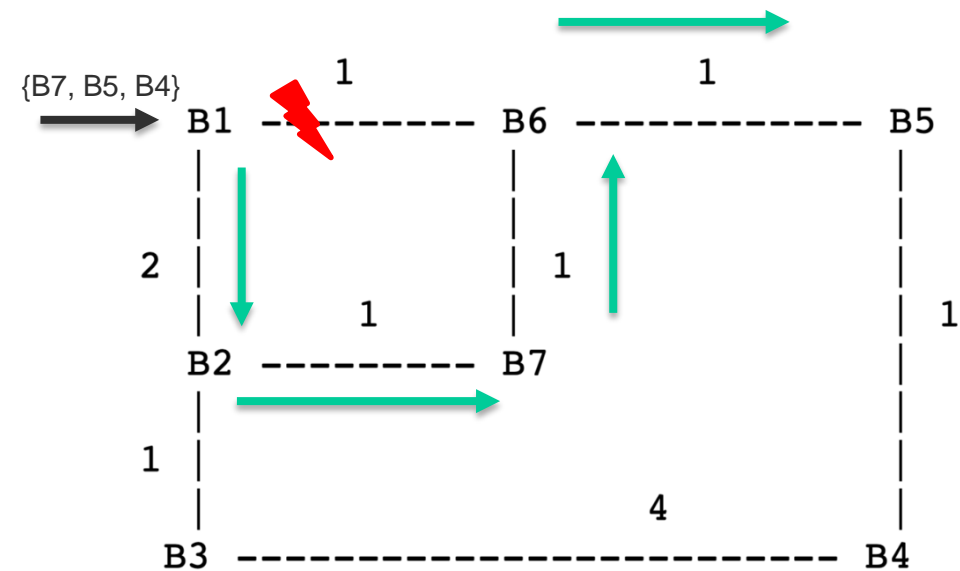
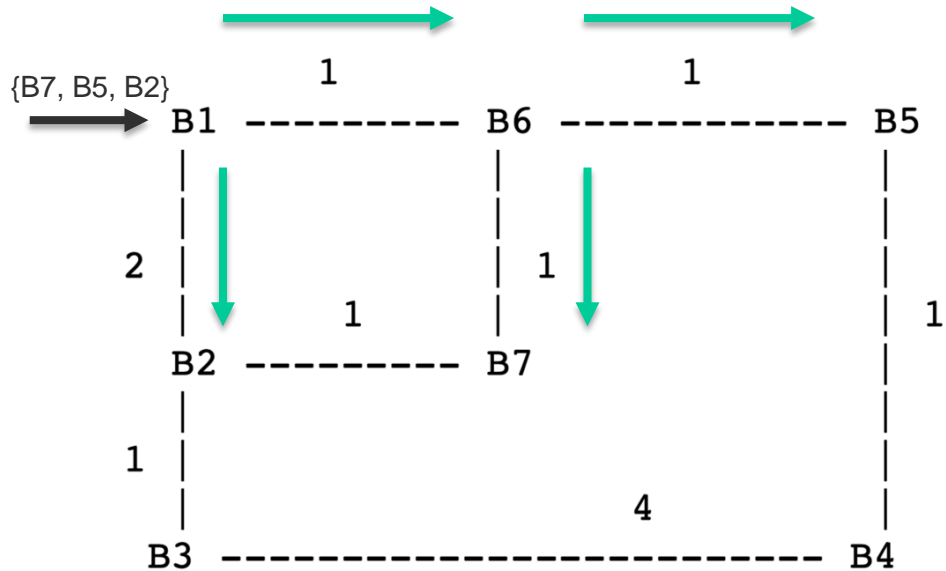


- ▶ Packet delivered to {B7, B5, B2}
- ▶ One redundant packet on link B1-B2
- ▶ Longer backup path for B4





Example for LFA-based BIER-FRR (2)



- ▶ Packet delivered to {B7, B5, B2}
- ▶ No redundant packet on link B1-B2

