

Experimenting with Bit Index Explicit Replication

Louis Navarre*

UCLouvain

Belgium

louis.navarre@uclouvain.be

Nicolas Rybowski

UCLouvain

Belgium

nicolas.rybowski@uclouvain.be

Olivier Bonaventure

UCLouvain

Belgium

olivier.bonaventure@uclouvain.be

ABSTRACT

Bit Index Explicit Replication (BIER) is a recent multicast architecture that solves several problems with deployed IP multicast protocols. BIER embeds an implicit multicast tree representation inside each transmitted packet. With this new mechanism, it becomes possible to reconsider multicast applications. However, no open-source implementation of BIER can be found. This paper presents our open-source implementation of the BIER forwarding mechanism and a companion socket-like API. Additionally, we show simulations of the implementation with ns-3 DCE.

CCS CONCEPTS

• **Networks** → **Routing protocols; Programming interfaces; Logical / virtual topologies.**

KEYWORDS

Bit Index Explicit Replication, Multicast

ACM Reference Format:

Louis Navarre, Nicolas Rybowski, and Olivier Bonaventure. 2022. Experimenting with Bit Index Explicit Replication. In *CoNEXT Student Workshop 2022 (CoNEXT-SW '22)*, December 9, 2022, Roma, Italy. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3565477.3569151>

1 INTRODUCTION

Multicast protocols enable efficient one-to-many communication as packets are duplicated in the network only when needed, thus ensuring that at most a single copy of a same packet will reach each router. IP Multicast, the current deployment of this architecture, suffers from scalability issues [2]. One of them concerns the states that each node of a multicast tree must keep as the tree is explicitly built. This memory consumption linearly grows with the number of multicast groups in a network, thus becoming a problem for large-scale multicast deployments.

The IETF standardized the Bit Index Explicit Replication (BIER) mechanism [12], a new source-routed multicast transmission mechanism. Each BIER packet embeds a *bitstring*, where each bit uniquely represents an egress router in the network. A bit set indicates that the corresponding router should receive the packet. Based on the *bitstring*, forwarding nodes rely on their Bit Index Forwarding Table

(BIFT) to choose the output interface(s) for the packet (and its potential copies), thus implicitly building the multicast tree. The table is constructed statically or following the Interior Gateway Protocol (IGP). In opposition to IP multicast, the intermediate routers need constant memory state, independently of the number of multicast groups in the network.

BIER opens-up possibilities for easier deployment of multicast applications. New transport protocols can be built atop this mechanism and benefit from its source-routed architecture. Other research problems could also be reconsidered, such as reliable multicast and multicast congestion control algorithms using the traffic-engineering extensions for BIER (BIER-TE) [3]. Recent works already explore that direction. However, they either rely on simulation models [1], P4 designs [6], or non-open-source implementation [5].

Open-source implementations of protocols help to drive research and evaluate new designs. In this paper, we present our open-source implementation of the BIER data-plane. Its objective is to provide to the research community a building block to design and experiment protocols on top of this new multicast architecture. It leverages the C standard library to represent a BIER data-plane daemon. As no formal API has been defined by the IETF to communicate with BIER, we propose and implement a socket-like API to interact with our daemon. The main challenge consisted at designing a simple API still exposing the full functionalities of BIER to the application. The BIER packets are forwarded using IPv4/IPv6 tunnels [14] to accommodate existing networks that do not natively support the architecture. We also provide an implementation of the BIER-TE extension [3]. To demonstrate it, we simulate a small network topology with ns-3 DCE and our daemon. The implementation essentially focuses on source-specific multicast, with (S,G) representing the multicast Source and Group addresses pair. Section 2 presents the BIER data-plane implementation and the API to communicate with the daemon. The ns-3 DCE simulation is detailed in Section 3. Section 4 concludes this paper. The data-plane uses ~1500 SLOC and the API uses 400 SLOC. The open-source implementation of BIER and our socket-like API are available on GitHub [7].

2 IMPLEMENTATION

We design our daemon to work on existing Linux network stacks. As such, BIER packets are encapsulated in IPv4/IPv6 tunnels [14] between two BIER Forwarding Routers (BFR). Routers not supporting BIER can still co-exist in the network and simply forward them as pure IP packets. The destination of the tunnel header is the address of the *next BFR* that will process the packet. The *next BFR* is found by inspecting the packet *bitstring*. For each bit set in the *bitstring*, the daemon queries its BIFT, and erases the potential other bits that are mapped to the same outgoing interface. Finally, it sends to the *next BFR* a copy of the modified packet.

*Louis Navarre is an F.R.S.-FNRS Research Fellow.

Each router running the daemon listens to incoming BIER packets using a raw IP socket. This socket is also used to forward the packet copies after processing.

The implementation supports the BIER-TE extension [3] for traffic engineering purposes. With BIER-TE, each link (or IP tunnel) between two BFRs is also assigned a bit in the *bitstring*. To forward a packet, a BFR now inspects the bits corresponding to its outgoing interfaces. BIER-TE offers strong traffic-engineering capabilities, at the cost that the source must explicitly encode the entire path towards the BIER egress routers of the multicast flow.

The two forwarding mechanisms are distinct and have different outcomes. We leverage the BIFT-id of the BIER header [13] to let the data-plane daemon understand which procedure to use when it processes a BIER packet.

The data-plane daemon runs in its own process to forward and create BIER packets. However, a BIER Forwarding Router may also act as an egress router, meaning that it could receive packets destined to its connected hosts. These packets should leave the BIER data-plane and be forwarded to an upper layer. Additionally, a multicast source must be able to send packets towards its receivers through the BIER data-plane.

To this end, we designed a socket-like API to let upper-layer applications, e.g., multicast transport protocols, communicate with the BIER daemon. As the daemon runs in its own process, the channel uses datagram UNIX sockets in both directions. The API (Listing 1) mimics the socket API.

```

1 ssize_t sendto_bier(int socket, const void *buf,
  size_t len, const struct sockaddr *dest_addr,
  socklen_t addrlen, uint16_t proto, bier_info_t
  *bier_info);
2 ssize_t recvfrom_bier(int socket, void *buf,
  size_t len, struct sockaddr *src_addr,
  socklen_t *addrlen, bier_info_t *bier_info);
3 int bind_bier(int socket, const struct sockaddr_un
  *bier_sock_path, bier_bind_t *bind_to);
4 int unbind_bier(int socket, const struct
  sockaddr_un *bier_sock_path, bier_bind_t *
  unbind_from);

```

Listing 1: Socket-like API.

`sendto_bier` forwards to the daemon the payload `buf` of length `len`. The UNIX path to the data-plane daemon is represented by `dest_addr`, of length `addrlen`. The `proto` value specifies the protocol of the payload, as defined in [12]. The last parameter, `bier_info`, is a new structure containing (i) the BIFT-id that a BFR uses to forward the packet, (ii) the *bitstring* of the packet and (iii) the *bitstring* length [13].

`recvfrom_bier` reads packets received from the data-plane daemon in `buf`, of capacity `len`. The `src_addr` is the address of the upstream BFR. `bier_info` carries the position in the *bitstring* of the upstream BIER router.

`bind_bier` informs BIER of the interest of an application in a multicast flow. `bier_sock_path` is again the UNIX path to communicate with the daemon. `bier_to` is yet another structure that specifies the (S,G) pair.




Figure 1: Small network topology with 6 nodes.

Similarly, `unbind_bier` communicates to the daemon the intent to leave the multicast group. Upon reception of these messages, the BIER data-plane forwards to the BIER ingress router a BIER packet containing the multicast join/leave message.

3 SIMULATING BIER

We demonstrate our implementation using ns-3 [9] and its Direct Code Execution (DCE) extension [11]. In a nutshell, DCE enables the execution of userspace code within ns-3 without source code changes. We wrote a wrapper to simulate our daemon in ns-3 DCE [10]. For illustration, we simulate a small network of 6 routers. The topology is represented in Figure 1. All links have a cost of 1, except 1-3 and 2-5 with a cost of 10. For simplification, we simulate hosts directly on the nodes and compute BIFTs beforehand using a shortest-path algorithm. Our host receivers and senders leverage our socket-like API. Node 1 simulates a multicast sender, waiting for all other nodes (the receivers) to send a multicast join request. We run the simulation with BIER and BIER-TE, and collect packet captures from each interface of all routers. The multicast trees constructed by BIER are represented in bold orange by analyzing these packet captures.

With BIER, the multicast tree is implicitly built using the shortest path towards its leaves. As expected, the flow uses the shortest path to reach each receiver (Figure 1a) from Node 1. BIER-TE [3] enables the source to specify the entire multicast tree for traffic engineering purposes. Once the receivers join the multicast flow, we set the *bitstring* so that the multicast tree uses the two links with a cost of 10. Figure 1b shows that the path is correctly enforced despite the higher link costs.

4 CONCLUSION

This paper presented our open-source implementation of the Bit Index Explicit Replication (BIER) forwarding mechanism [12]. We provide our BIER implementation and propose a simple socket-like API to communicate with our daemon, as no such formal API has been standardized. This implementation can be deployed in networks that do not support BIER by leveraging IP tunnels to carry the BIER packets [14]. With this mechanism, several research problems related to multicast transport protocols can be tackled more easily, such as reliable communication, multicast congestion control and scalable deployment. We hope that this work will stimulate the research community into new efforts in developing multicast protocols above BIER. As our implementation follows the BIER specifications [12], it should be interoperable with P4 designs [6]. Future work will consist at integrating a control-plane above open-source implementations of routing algorithms, such as IS-IS [4] or OSPF [8].

REFERENCES

- [1] Yoann Desmouceaux, Thomas Heide Clausen, Juan Antonio Cordero Fuertes, and W Mark Townsley. 2018. Reliable multicast with BIER. *Journal of Communications and Networks* 20, 2 (2018), 182–197.
- [2] Christophe Diot, Brian Neil Levine, Bryan Lyles, Hassan Kassem, and Doug Balensiefen. 2000. Deployment issues for the IP multicast service and architecture. *IEEE network* 14, 1 (2000), 78–88.
- [3] Toerless Eckert, Michael Menth, and Gregory Cauchie. 2022. Tree Engineering for Bit Index Explicit Replication (BIER-TE). RFC 9262. <https://doi.org/10.17487/RFC9262>
- [4] Les Ginsberg, Tony Przygienda, Sam Aldrin, and Zhaohui (Jeffrey) Zhang. 2018. Bit Index Explicit Replication (BIER) Support via IS-IS. RFC 8401. <https://doi.org/10.17487/RFC8401>
- [5] Alessio Giorgetti, Andrea Sgambelluri, Francesco Paolucci, Nicola Sambo, Piero Castoldi, and Filippo Cugini. 2017. Bit Index Explicit Replication (BIER) multicasting in transport networks. In *2017 International Conference on Optical Network Design and Modeling (ONDM)*. IEEE, 1–5.
- [6] Daniel Merling, Steffen Lindner, and Michael Menth. 2021. Hardware-based evaluation of scalable and resilient multicast with bier in p4. *IEEE Access* 9 (2021), 34500–34514.
- [7] Louis Navarre. 2022. Bit Index Explicit Replication data-plane and socket-like API implementation. <https://github.com/louisna/bier-socket-api.git>. Accessed: 2022-10-21.
- [8] Peter Psenak, Nagendra Kumar Nainar, Ijsbrand Wijnands, Andrew Dolganow, Tony Przygienda, Zhaohui (Jeffrey) Zhang, and Sam Aldrin. 2018. OSPFv2 Extensions for Bit Index Explicit Replication (BIER). RFC 8444. <https://doi.org/10.17487/RFC8444>
- [9] George F Riley and Thomas R Henderson. 2010. The ns-3 network simulator. In *Modeling and tools for network simulation*. Springer, 15–34.
- [10] Nicolas Rybowski and Olivier Bonaventure. 2022. Evaluating OSPF Convergence with ns-3 DCE. In *Workshop on ns-3 (WNS3)*.
- [11] Hajime Tazaki, Frédéric Uarbani, Emilio Mancini, Mathieu Lacage, Daniel Camara, Thierry Turletti, and Walid Dabbous. 2013. Direct code execution: Revisiting library os architecture for reproducible network experiments. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*. 217–228.
- [12] Ijsbrand Wijnands, Eric C. Rosen, Andrew Dolganow, Tony Przygienda, and Sam Aldrin. 2017. Multicast Using Bit Index Explicit Replication (BIER). RFC 8279. <https://doi.org/10.17487/RFC8279>
- [13] Ijsbrand Wijnands, Eric C. Rosen, Andrew Dolganow, Jeff Tantsura, Sam Aldrin, and Israel Meilik. 2018. Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks. RFC 8296. <https://doi.org/10.17487/RFC8296>
- [14] Zheng Zhang, Zhaohui (Jeffrey) Zhang, Ijsbrand Wijnands, Mankamana Prasad Mishra, Hooman Bidgoli, and Gyan Mishra. 2022. *Supporting BIER in IPv6 Networks (BIERin6)*. Internet-Draft draft-ietf-bier-bierin6-05. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-bier-bierin6/05/> Work in Progress.