MyRiAD: Towards Transport Protocol Agnostic In-Network Resilience

Daniel Helmig and Martina Zitterbart
Institute of Telematics and KASTEL Security Research Labs
Karlsruhe Institute of Technology
Karlsruhe, Germany
daniel.helmig@kit.edu, zitterbart@kit.edu

***Abstract*—The shift of power systems towards renewable energy and decentralized energy generation requires additional efforts of operators to keep the grid stable. To accomplish this task, grid operators use Intelligent Electronic Devices (IEDs). These IEDs need to communicate with each other, both inside the LAN of a substation, and between distinct substations and control centers over a wide area network (WAN). Such grid control traffic requires high reliability and low latency. While the communication architecture in substation LANs is well established, communication across WANs is still a field of active research. Therefore, we propose MyRiAD, a duplication-based in-network resilience mechanism implemented in the network layer. MyRiAD leverages existing ring structures within network topologies to boost resilience. Our evaluation results demonstrate that MyRiAD reduces end-to-end packet loss by up to 93% compared with related techniques and achieves over 96% loss reduction compared to conventional shortest-path routing.**

***Index Terms*—network resilience, redundancy, grid control**

I. Introduction

Power grids shift from centralized electricity generation towards increasingly decentralized ways of providing electricity, making grid control more difficult. To keep their grids working seamlessly, grid operators need ways to orchestrate the interplay of different grid parts. This is realized using Intelligent Electronic Devices (IEDs) that report measurements and take necessary protective actions. IEDs need to exchange data among each other inside the LAN of a substation and also inside the grid operator's WAN to communicate with different substations and control centers. Grid control applications such as grid protection have stringent latency requirements (<10 ms) and demand high reliability (>99.99%) [1]. While the communication architecture and protocols inside a substation LAN are well established, the communication inside the operator WAN is still actively researched.

In WANs with longer propagation delays, packet retransmissions become non-viable to ensure reliable packet delivery under low latency requirements. In addition to packet loss, link failures provide a significant challenge towards low latency and high reliability. To mitigate link failures, failover-based mechanisms (e.g., MPLS Fast Reroute [2], Topology Independent Loop Free Alternate [3]) predefine backup paths. Upon detection of a link failure, traffic is switched onto the backup path, bypassing the failed link. However, such mechanisms are inherently reactive and need time to detect the link failure and react accordingly. Until the backup path is active, packets continue to get lost.

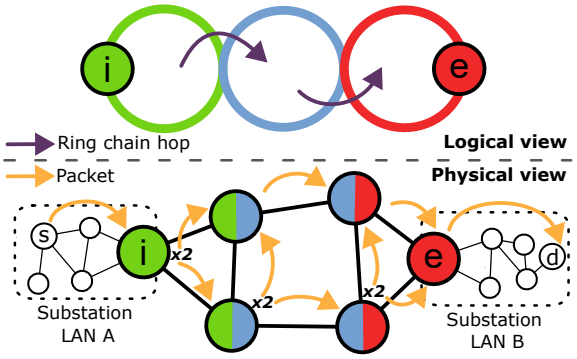

Fig. 1: Ring structure of MyRiAD with sender s, ingress node i, egress node e and destination d.

In this paper, we present MyRiAD[1], a duplication-based in-network resilience mechanism for WAN communication in power grid networks. MyRiAD can cope with packet losses and link failures while providing low latency. It duplicates and de-duplicates packets at the network layer, and thus is agnostic to the transport protocol being used by grid control applications. MyRiAD provides resilience even when faced with multiple packet losses and multiple link failures along an end-to-end path of a packet traveling from the sender s to its destination d. MyRiAD uses rings to construct ring chains in order to achieve end-to-end resilience. A ring chain is able to provide protection against at least one link failure or packet loss on each ring. Fig. 1 shows an example of such a ring chain. Packets are duplicated when entering a ring chain at the so-called ingress node i. The duplicates are sent clockwise and counter-clockwise along the ring, respectively. When transitioning from one ring to the next, only the first instance of the packet is duplicated again into the next ring. Further copies of the same packet get filtered out, i.e., de-

This work was supported by funding from the topic 46.23.02 Engineering Security for Energy Systems of the Helmholtz Association (HGF) and by KASTEL Security Research Labs.

ISBN 978-3-903176-82-9

[1]**M**any **Ri**ngs to **A**rrive at **D**estination

duplicated. The same is true for the egress node e of the ring chain, the first packet is sent into the substation LAN towards the destination d, while the duplicate packet arriving later is discarded.

Our contributions are the following:

- A novel duplication-based in-network mechanism to provide resilience for grid control applications against packet loss and multiple link failures in smart grid WAN topologies.
- A novel de-duplication scheme using random numbers in the IPv6 Hop-by-Hop options header.
- Evaluation in large networks. MyRiAD achieves a worst-case delivery ratio of 95.7% under adverse conditions i.e., packet loss and multiple link failures, among 1000 different sender-receiver pairs and reduces the end-to-end loss by 93% compared to related approaches.

II. Resilient communication in Wide Area Networks

In this section, we discuss the underlying communication infrastructure and assumptions we made. Then we present related work in the field of resilient wide area grid communication.

A. Wide area communication topologies in power grids

Grid operators, in addition to the power grid, also operate WANs [4] (e.g., as optical ground wires in transmission lines) to meet their communication needs. Fiber optic links connect the routers of different substations and control centers. Related work [5], [6] suggests that grid communication topologies are hierarchical and consist (at least in part) of rings. MyRiAD leverages these existing rings to improve resilience. Inside the network of a substation reside Remote Terminal Units (RTUs) and Intelligent Electronic Devices (IEDs), which are responsible for taking measurements and protective actions [6]. Teleprotection requires 10 ms latency [6], [7] and high reliability, thus these requirements also need to be met by the communication infrastructure.

The scope of MyRiAD is resilient communication in a smart grid between different substations that are interconnected by a WAN. We target communication between substations and control centers, the communication inside a substation LAN is protected by already established resilience approaches such as the Parallel Redundancy Protocol (PRP) or High-availability Seamless Redundancy (HSR) [8]. Each substation has at least one router connecting it to other substations over the aforementioned WAN.

Highly reliable and low latency communication requires resilient networks, as otherwise reliability targets can not be met in the presence of failures. Our goal is to establish resilient low-latency communication between routers of different substations.

B. Related Work

Existing approaches to resilient grid communication can roughly be separated into duplication-based and failover-based approaches. Duplication-based approaches send the same packet multiple times to increase the probability of successful packet delivery. Failover-based approaches switch traffic on pre-defined backup paths to quickly mitigate link failures.

Examples for duplication-based approaches are PRP and HSR. PRP employs two redundant LANs to transport a duplicate over each LAN, while HSR uses a ring to transport a duplicate in each direction along the ring. PRP and HSR operate at the link layer and are not suited for the use in WANs. Protocols working at the transport layer are Multipath TCP (MPTCP) [9], IP Parallel Redundancy Protocol (iPRP) [10] and Transparent Redundancy for UDP STreams (TRUST) [11]. MPTCP can be leveraged to send redundant packets along multiple paths, thus increasing resilience [6]. TRUST and iPRP operate by the same principle and send packets on two redundant paths, but they operate on top of UDP. Specifying a transport protocol like UDP in order to improve resilience restricts the use of higher layer protocols that use TCP (e.g., MMS [12]) and vice versa (e.g., GOOSE [12]). In order to be transport protocol agnostic, MyRiAD targets the network layer. OD^{3}R is another approach targeting the network layer. It increases resilience by routing packets via disjoint paths [13]. Using two redundant paths, the mitigation of a single link failure is possible, yet MyRiAD also covers multiple link failure scenarios.

Representatives of failover-based approaches are, e.g., the Media Redundancy Protocol (MRP) [14] and the Rapid Spanning Tree Protocol (RSTP) [15]. Like HSR and PRP, MRP and RSTP are unsuited due to being link layer protocols. MPLS Fast Reroute [2] is an extension to MPLS where precomputed backup-paths are used to quickly divert traffic in case of a link failure. A more specialized version for ring topologies exists in the form of MPLS-TP Shared-Ring Protection [5]. Packets are sent along a ring and are switched onto the backup path in the reverse direction upon link failure. Likewise, Topology Independent Loop-Free Alternate [3] is an approach using segment routing to divert traffic from failed links. However, all failover-based approaches are reactive and therefore cannot guarantee the 10 ms delay requirement due to their inherent delay < 50 ms [5]. Further, packet loss is not covered by failover-based approaches. Additionally, Fast Reroute has configuration overhead: Each link needs a backup path to provide resilience against any link failing. Fast Reroute is designed to protect against single pre-planned failures, thus a backup path for every link can not be expected [3]. In order to overcome the issues of multiple link failures and the inherent delay of reactive approaches we propose MyRiAD.

III. MyRiAD

The basic network building blocks of MyRiAD are rings. They are used to provide resilience against multiple packet losses and multiple link failures at the same time. MyRiAD uses packet duplication and de-duplication when entering and leaving a ring, respectively. We combine each substation's LAN and its router into a single node of a graph $G = (V, E)$ with nodes V and edges E, however some nodes can be

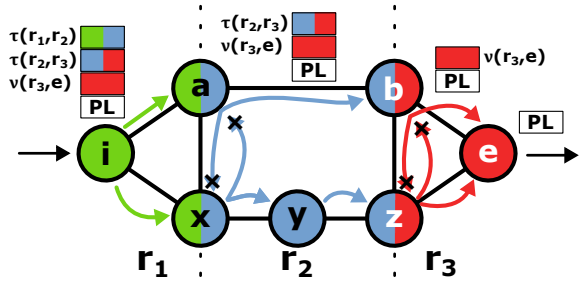

Fig. 2: Forwarding in MyRiAD with different label stacks and payloads along the path. Arrows are colored to indicate the current ring for forwarding and an x denotes de-duplication.

routers without an associated substation. For simplicity, we omit physical redundancy i.e., redundant routers or multiple links between routers.

Starting from the physical topology, MyRiAD first needs to determine rings $r_i, i \in 1, ..., n$ in G. Rings can either be obtained using cycle algorithms for graphs [16], [17] or defined manually e.g., because the topology is made up of rings to begin with [5], [6]. These rings are reused for packet delivery between different senders and receivers.

MyRiAD uses MPLS labels to forward packets. Transitions between adjacent rings consist of either one or two *transition nodes*. Every transition is assigned a *transition label* $\tau(r_i, r_j)$. Each node in a ring $r_i = \{n_j \mid j \in 1, .., k\}$ is assigned a unique MPLS label $\nu(r_i, n_j)$ called *node label*. The label $\nu(r_i, n_j)$ allows for identification of node n_j in ring r_i as egress node.

The ring chain $c = r_1, r_2, r_3$ with the *egress node* e on ring r_3, consists of two transition labels followed by a node label marking the egress node, i.e., $\tau(r_1, r_2), \tau(r_2, r_3), \nu(r_3, e)$. To achieve the correct forwarding behavior, every node needs flow rules (e.g., OpenFlow or P4) according to the assigned labels. MyRiAD needs four types of flow-rules: ingress rules, transition rules, egress rules and ring forwarding rules.

Starting with the ingress rules, packets arriving at the ingress node are duplicated and sent out on both *ring ports* i.e., clockwise and counter-clockwise. Transition nodes between ring r_i and r_j handle packets with the label $\tau(r_i, r_j)$. For a packet with a matching transition label, the label is popped and the packet duplicated again onto both ring ports. In addition to duplicating on the next ring, the transition nodes perform de-duplication on incoming packets, which will be explained later. If a packet with $\nu(r_i, n_j)$ arrives at the egress node n_j on ring r_i the packet is forwarded to the local substation. Packets not containing a matching node or transition label are simply forwarded along the ring.

Fig. 2 shows an example ring chain consisting of three rings r_1, r_2, r_3. Packets arriving at i are duplicated towards a and x. The transition nodes between r_1 and r_2, a and x, duplicate packets onto r_2 while filtering each others duplicate packets. Because y is not a transition node, packets arriving from x are simply forwarded along r_2 to z. Nodes b and z match on $\tau(r_2, r_3)$ and duplicate onto r_3, again receiving duplicated from each other an filtering them accordingly. Finally the packet arrives at the egress node e, where the top label matches $\nu(r_3, e)$, gets popped and the payload delivered.

A. Identifying duplicates

To distinguish duplicates from non-duplicates, MyRiAD uses a unique random number RN per packet. Two packets are considered duplicates if their RN matches. MyRiAD introduces a new IPv6 Hop-by-Hop option [18] in order to mark packets unambiguously. This option is inserted in the Hop-by-hop Options header of the IPv6 datagram at the ingress node, before duplication takes place. The structure of the option is depicted in Fig. 3. It consists of the necessary type and length fields and a 64 bit random number (RN). MyRiAD's option has type 0x1f.

Packets can have matching RNs by chance and be wrongfully discarded as duplicates. In this case, the effect is the same as for an isolated packet loss. In order to avoid false positives i.e., two packets carrying the same RN by chance, we used the approximation formula of the birthday paradox $1 - e^{-\frac{n*(n-1)}{2k}}$ [19]. The total number of packets in a time interval is denoted by n, while k is the number of possible RNs. To determine an appropriate size of RN, we evaluated different link rates and RN sizes. We assumed a minimum frame size of $82+l$ Bytes (l being the length of RN), consisting of an IPv6 header, MyRiAD's Hop-By-Hop Options header, no higher layer protocols or payload and assuming Ethernet at the link-layer. This gives us the maximum amount of packets possible for a fixed link rate and time interval. Further, we assumed a time interval of 20 ms, which is double the latency requirement for control applications. Such an extended time interval gives a conservative estimate for false positives, due to more packets potentially carrying matching RNs. Using different link rates and RN sizes we determined the probability for at least one false positive occurring.

The results are depicted in Fig. 4, with different RN sizes on the x-axis and different data rates on the y-axis. Consequently, we opted for 64 bit RNs where the probability for a false positive occurring in 20 ms is around $10^{-5}\%$ even at the full line rate of 100 Gbit/s.

B. De-duplication

MyRiAD targets high-reliability low-latency grid control traffic. We do not expect large data flows incurred by control traffic. Payload sizes vary between 4 and 157 Bytes, as suggested in [20]. Assuming a packet every 100 ms as in [6], [20], a single flow is responsible for about 18 kbit/s of data. However de-duplication is still necessary to prevent exponential duplication along a ring chain.

To de-duplicate packets, each de-duplication instance maintains an array to cache already seen RNs of previous packets. De-duplication based on an array results in constant memory requirements and $O(1)$ lookups. The de-duplication instance parses the Hop-by-Hop option of an arriving packet and reads the RN contained in the option. The next steps are shown in Fig. 5.

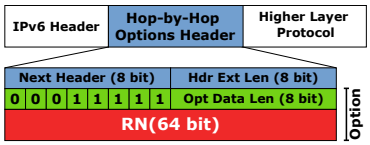

Fig. 3: MyRiAD's option with the random number RN inside the Hop-by-Hop Options header.

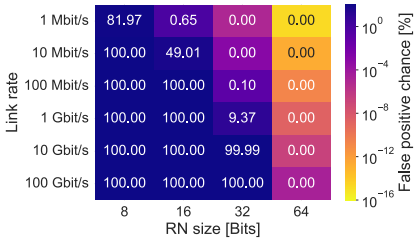

Fig. 4: Chance for two packets to have the same RN in a time interval of 20 ms for various RN sizes and data rates.

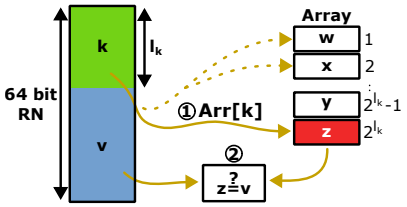

Fig. 5: De-duplication. RN gets split into k and v depending on the key size. k is used as index into the array.

TABLE I: Different simulation modes. FRR represents the failover-based mechanisms, RP represents other duplication based mechanisms.

Name	Description	Covers	
		Packet loss	Link failure
SP	Shortest path routing	No	No
FRR	Reroute around failure	No	Partially
RP	Two redundant paths	Yes	Single
MyRiAD	Presented in section III	Yes	Multiple

① The RN gets split into a key k and value v, consisting of l_k and l_v bits ($l_k + l_v = 64$). We call l_k the *key size*. The de-duplication array has 2^{l_k} entries of length l_v to use k as the index for a lookup.

② In order to determine if a packet is a duplicate, the value at the index is compared to v of the received RN. If v equals the saved value, then the packet is considered a duplicate and dropped. If the values differ, then the packet is considered new and v is saved into the array at this index for future lookups. Using this de-duplication mechanism the computational complexity and the memory requirements stays the same regardless of ring sizes and the time between duplicate packets. However the de-duplication performance degrades as more and more entries in the de-duplication array get overwritten. Therefore the de-duplication array must be sized accordingly. For our evaluation MyRiAD uses a key size of 18 bits, which results in roughly 1.4 MB of required memory.

IV. Evaluation Methodology

To evaluate on large topologies, we implemented MyRiAD in a simulation using the ns3 network simulator [21]. To compare MyRiAD with baselines, the simulation has four different modes which are explained in table I.

We construct a topology called *ring hierarchy* consisting of 50 nodes arranged in six rings to act as a baseline. 10 nodes are arranged in a central ring and the remaining 40 nodes are attached to this central ring to form five additional rings. Further we used the DFN topology from the topology zoo [22] as a general topology. For these topologies every link has a propagation delay of 1 ms and link speed of 100 Gbit/s.

We additionally used the topology generator POBTOG [23] to generate two large hierarchical topologies consisting of four layers. POBTOG generates a topology for a given country by partitioning the country based on a set population threshold. We left the population threshold at the default value and generated two topologies, one topology using Switzerland with about 1700 nodes and 2200 edges and one topology using Germany with about 13000 nodes and edges. For the Switzerland topology we configure the core layer to be fully connected and the other three layers to be arranged in rings. For the Germany topology we followed the example of [23] and configured the core layer to be fully connected, the layer below to be arranged in a star and the lower two layers to be rings. The links in both the Switzerland and Germany topology have link capacities, depending on the layer, of 400 Gbit/s, 100 Gbit/s, 10 Gbit/s or 1 Gbit/s, with links in the core having the most capacity. All links have a propagation delay according to their geographical length times a factor of $\frac{2}{3}$ the speed of light.

Before simulation start, ring chains from sender to receiver are calculated for the topologies. Ring chains run along the rings of different hierarchical layers in the topology, with transitions between layers. If no consecutive ring chain is possible i.e., due to the star layer in the Germany topology, multiple ring chains are interconnected using regular shortest path routing. Every simulation has a duration of 30 seconds. However, SP does not recalculate the shortest path, as we used simulation times of 30 seconds to enhance the effects of link failures and packet loss.

We determined two edge redundant paths for RP i.e., two paths that have no common edge. If full edge redundancy was not possible due to a necessary common edge, we calculated redundant path segments instead of fully redundant paths and used the common edges only where needed. For FRR, we calculated backup paths, i.e., shortest alternative paths, for the links in the core ring. The switchover to the backup path occurs after a delay of 50 ms, as suggested in [5]. In the DFN topology we calculated basic cycles using the algorithm in [16]. Like SP, FRR, RP and MyRiAD do no recalculation of paths or rings during simulation.

For all experiments, we assume no queuing delay as packets for grid control applications must experience minimal queuing delay to meet their stringent delay requirements and thus need to be subjected to high priority treatment, e.g., by using priority queues. Delay therefore is only determined by propagation delay. To simulate IEDs in a substation LAN sending and

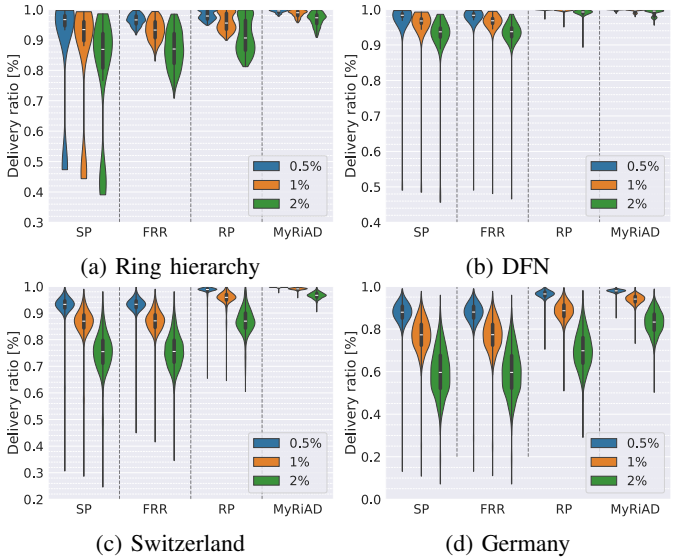

(a) Ring hierarchy (b) DFN

(c) Switzerland (d) Germany

Fig. 6: Packet delivery ratios for the evaluated topologies

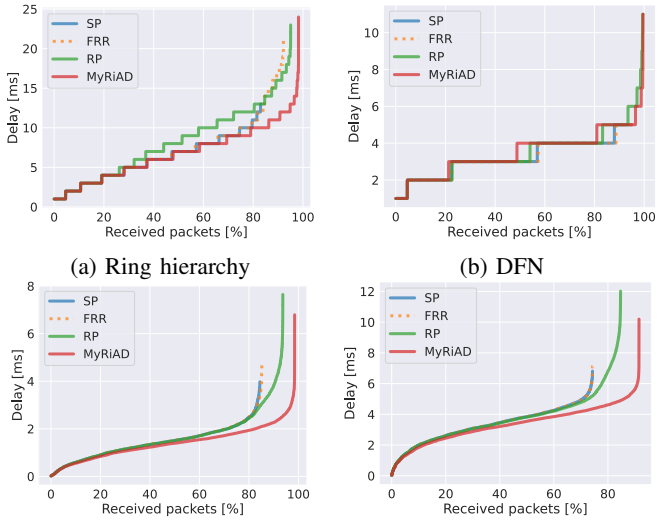

(a) Ring hierarchy (b) DFN

(c) Switzerland (d) Germany

Fig. 7: Experienced packet delay of all packets for the evaluated topologies

receiving packets, every node has one host attached, acting as sender or receiver. Each sender sends a UDP packet every 10 ms with a payload of 200 Bytes for a total of 3000 packets per sender. A packet every 10 ms roughly corresponds to 10 IEDs in the substation sending data.

V. Evaluation

The goal of MyRiAD is to achieve a high packet delivery ratio while maintaining low latency, even when faced with packet loss and link failures. Therefore, we evaluate the delivery ratio in section V-A and then take a look at one-way delay in V-B.

A. Delivery ratio

Small topologies: Hierarchical rings and DFN: For the experiment 100 sender-receiver pairs are chosen at random, these are the same for all modes. We repeat each experiment three times with 0.5%, 1% and 2% loss probability per link. For each sender-receiver pair, we determine the ratio of successfully delivered packets. We simulate a link failure by disabling one random link after half the simulation time.

Fig. 6a shows the resulting distribution of delivery ratios (y-axis) of all sender-receiver pairs for every mode (x-axis) for all runs for the hierarchical ring topology. The pairs that achieved $< 50\%$ delivery ratio for SP contain the failed link in the shortest path and thus no packets can be delivered after the link failure. FRR achieves significantly better results because the link failure can be mitigated by the backup path. Compared to RP, FRR performs worse because the random packet loss has higher impact, as a lost packet is lost for good, whereas with RP a duplicate is still underway. MyRiAD achieves the best overall delivery ratio in the hierarchical rings topology, achieving $> 97\%$ delivery ratio between all sender-receiver pairs for 0.5% packet loss.

Fig. 6b shows that FRR is not able to recover from every link failure, as there are Sender-Receiver pairs where Fast Reroute protects another link along the path instead of the failed link. Looking at the worst case delivery ratios for 0.5% loss ratio, MyRiAD achieves 98.6% compared to 97.2% for RP, and only 49.1% for SP and FRR.

Large topologies: Switzerland and Germany: We chose 1000 random sender-receiver pairs and again used packet loss ratios of 0.5%, 1% and 2% per link. Again, we repeated each experiment three times. To also evaluate multiple link failure scenarios, every 5 seconds a random link in either the core layer or a ring fails, for a total of 5 link failures. We excluded link failures in other parts of the topologies, e.g., star layers, as they would separate the topology in two non-connected parts. For FRR, we calculated shortest backup paths for core links. We did not calculate backup paths for every possible link and did not define backup paths for links on backup paths, as the number of needed backup paths does not scale well in this case, which is one drawback of relying on Fast Reroute.

The results for the Switzerland and Germany topology are shown in Fig. 6c and Fig. 6d respectively. For the Switzerland topology and 0.5% loss, SP and FRR perform very similar with a median delivery ratio of 93.3% and a worst case of 30.6% and 44.9% respectively. FRR is not able to increase the delivery ratio compared to SP for all sender-receiver pairs. This is because only core links are protected by a backup path and link failures on other links can not be mitigated. Again looking at 0.5% loss, RP achieves a median of 98.8% and a worst case delivery ratio of 65.4%. MyRiAD achieves a median of 99.8% and a worst case of 97.6%, corresponding to 93% less end-to-end loss compared to RP. MyRiAD performs especially well on the Switzerland topology as it consists of rings on each layer (except the fully connected core), playing into the strengths of MyRiAD.

Results for higher packet loss ratios and the Germany topology are similar, albeit worse in general due to the larger topology size and therefore longer paths, resulting in more packets being lost due to random loss on links.

B. One-way delay

To evaluate if MyRiAD is able to achieve high packet delivery ratios while adhering to the stringent latency requirements imposed by grid control applications we measured the one-way delay of each packet sent during our evaluation.

The results are shown in Fig. 7. The x-axis shows the total amount of delivered packets between all simulated sender-receiver pairs and the y-axis the cumulative distribution of one-way delay. SP delivers the least amount of packets overall, but does so over the shortest path. The curve of FRR matches that of SP but reaches a higher number of delivered packets as packets lost to a link failure can be mitigated. This effect is less pronounced on the larger topologies in Fig. 7c and Fig. 7d as not all links are protected by the reroute mechanism. RP achieves a higher number of delivered packets than SP and FRR, but Fig. 7a shows packets experiencing larger delays than the other approaches. This is caused by the calculation of redundant paths. It is possible that neither of the two redundant paths is the shortest path, hence a packet sent via redundant paths might take longer.

MyRiAD delivers the most packets overall while providing comparable latencies as SP. This is because MyRiAD implicitly takes the shortest path along rings. If the shortest path coincides with the ring chain i.e., runs along the rings, MyRiAD also uses the shortest path. However, in topologies not consisting of rings it is possible that packets delivered via MyRiAD take longer than packets sent along the shortest path if that path does not follow the rings used by MyRiAD. This effect can be seen in Fig. 7b, where some packets sent by MyRiAD take longer than packets sent via the shortest path.

VI. Conclusion

In this paper, we presented MyRiAD, an approach to achieve duplication-based transport protocol-agnostic in-network resilience. By using rings in a topology, MyRiAD is able to be resilient against multiple packet losses and multiple link failures along an end-to-end path. We showed a way to implement MyRiAD and presented a novel de-duplication scheme based on random numbers carried in a custom Hop-by-Hop option. We evaluated MyRiAD with simulations on large scale topologies. We achieved a worst case delivery ratio of 97% under adverse conditions with constant packet loss and multiple link failures, which improved upon the result of redundant packets by 93%. Additionally, we showed that MyRiAD has the same one-way delays compared to related approaches. In future work, we plan to evaluate the performance of MyRiAD on real hardware and improve the implementation.

References

[1] IEC, "Communication networks and systems for power utility automation – Part 5: Communication requirements for functions and device models." IEC 61850-5:2013 + A1:2022, 2020.

[2] A. Atlas, G. Swallow, and P. Pan, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels." RFC 4090, May 2005.

[3] A. Bashandy, S. Litkowski, C. Filsfils, P. Francois, B. Decraene, and D. Voyer, "Topology Independent Fast Reroute Using Segment Routing." RFC 9855, Oct. 2025.

[4] M. Adamiak, A. Apostolov, M. Begovic, C. Henville, K. Martin, G. Michel, A. Phadke, and J. Thorp, "Wide area protection—technology and infrastructures," *IEEE Transactions on Power Delivery*, vol. 21, no. 2, pp. 601–609, 2006.

[5] W. Cheng, L. Wang, H. Li, H. van Helvoort, and J. Dong, "MPLS-TP Shared-Ring Protection (MSRP) Mechanism for Ring Topology." RFC 8227, Aug. 2017.

[6] I. B. Fink, L. Ferlemann, M. Dahlmanns, C. Thimm, and K. Wehrle, "Emulating and Evaluating Transport Layer Protocols for Resilient Communication in Smart Grids," in *Proceedings of the 2025 IEEE/IFIP Network Operations and Management Symposium*, pp. 1–10, IEEE, 2025.

[7] R. Bächli, M. Häusler, and M. Kranich, "Teleprotection solutions with guaranteed performance using packet switched wide area communication networks," in *2017 70th Annual Conference for Protective Relay Engineers (CPRE)*, pp. 1–6, 2017.

[8] IEC, "Industrial communication networks – High availability automation networks – Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)." EN IEC 62439-3:2022 + AC:2023, 2023.

[9] A. Ford, C. Raiciu, M. J. Handley, O. Bonaventure, and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses." RFC 8684, Mar. 2020.

[10] M. Popovic, M. Mohiuddin, D.-C. Tomozei, and J.-Y. Le Boudec, "iprp: Parallel redundancy protocol for ip networks," in *2015 IEEE World Conference on Factory Communication Systems (WFCS)*, pp. 1–4, 2015.

[11] F. Neumeister, M. Göckel, and M. Zitterbart, "Trust: Transparent redundancy for udp streams," in *2023 IEEE 21st International Conference on Industrial Informatics (INDIN)*, pp. 1–7, 2023.

[12] IEC, "Communication networks and systems for power utility automation – Part 8-1: Specific communication service mapping (SCSM) Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3." IEC 61850-8-1:2011 + A1:2020, 2020.

[13] K. Altenhofen, J. Zobel, and B. Scheuermann, "Increasing network resilience through dynamic routing with disjoint paths," in *2024 IFIP Networking Conference (IFIP Networking)*, pp. 23–31, IEEE, 2024.

[14] IEC, "Industrial communication networks – High availability automation networks – Part 2: Media Redundancy Protocol (MRP)." EN IEC 62432-3:2021 + COR1:2023, 2023.

[15] "Ieee standard for local and metropolitan area networks–bridges and bridged networks," *IEEE Std 802.1Q-2022 (Revision of IEEE Std 802.1Q-2018)*, pp. 1–2163, 2022.

[16] K. Paton, "An algorithm for finding a fundamental set of cycles of a graph," *Commun. ACM*, vol. 12, p. 514–518, Sept. 1969.

[17] N. Deo, G. Prabhu, and M. S. Krishnamoorthy, "Algorithms for generating fundamental cycles in a graph," *ACM Trans. Math. Softw.*, vol. 8, p. 26–42, Mar. 1982.

[18] D. S. E. Deering and B. Hinden, "Internet Protocol, Version 6 (IPv6) Specification." RFC 8200, July 2017.

[19] F. H. Mathis, "A generalized birthday problem," *SIAM Review*, vol. 33, no. 2, pp. 265–270, 1991.

[20] M. Kuzlu, M. Pipattanasomporn, and S. Rahman, "Communication network requirements for major smart grid applications in han, nan and wan," *Computer Networks*, vol. 67, pp. 74–88, 2014.

[21] "ns-3 Network Simulator." https://www.nsnam.org/.

[22] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, "The internet topology zoo," *IEEE Journal on Selected Areas in Communications*, vol. 29, pp. 1765–1775, October 2011.

[23] F. Poignée, F. Loh, S. Schardt, F. Lipp, D. Hock, T. Hoßfeld, and International Conference on Networked Systems (Ilmenau, 01.-04.09-2025), "Pobtog: a population-based topology generator for country-wide communication networks," *Proceedings of the International Conference on Networked Systems 2025 (NetSys 2025): Technische Universität Ilmenau, 1 – 4 September 2025*, pp. 47–50, Aug 2025.