

Adaptive Routing Analysis And Performance Evaluation In Ipv6 Network Comparative Evaluation of RIPng, OSPFv3, and EIGRP Using Multi-Router Topologies

¹Vedant S. Janwadkar, ²Mrunal A. Muley, ³Dr. Mrs. D. M. Bhalerao

^{1,2,3}Department of Electronics and Telecommunication Engineering Sinhagad College of Engineering, Pune, India

Abstract - The rapid adoption of IPv6 has increased the importance of selecting efficient dynamic routing protocols capable of providing reliable connectivity, fast convergence, and stable network performance. This research presents a comparative analysis of three widely used IPv6 routing protocols, RIPng, OSPFv3, and EIGRP, through practical experimentation conducted in the Bharat Sanchar Nigam Limited (BSNL) laboratory. The experimental setup consisted of two end-user PCs connected through multiple Cisco 2811 series routers mounted on standard server racks and interconnected using RJ45-terminated straight-through and crossover LAN cables. The study evaluated protocol behavior in a multi-router IPv6 environment by measuring end-to-end connectivity, average ping latency, packet loss, and convergence time during normal operation and simulated link-failure conditions. Each routing protocol was configured separately on the same network topology to ensure a fair comparison. ICMPv6 ping tests were performed between the two PCs to analyze communication efficiency and recovery performance after topology changes. Prior to implementing it on hardware, the topology and routing behavior were validated in Cisco Packet Tracer. The results indicate that all three protocols successfully established IPv6 connectivity; however, significant differences were observed in convergence speed and overall network responsiveness. RIPng demonstrated simpler configuration but slower recovery behavior, while OSPFv3 and EIGRP showed faster adaptation to network changes and more stable performance under failure conditions. The findings highlight the practical strengths and limitations of each protocol and provide insight into selecting an appropriate routing solution for IPv6 enterprise and service provider environments using legacy Cisco infrastructure.

Keywords - IPv6, RIPng, OSPFv3, EIGRP, Cisco 2811 Routers, BSNL Laboratory, Convergence Time, Network Performance Evaluation

I. INTRODUCTION

The rapid growth of Internet-connected devices, cloud computing services, and emerging communication technologies has accelerated the global adoption of Internet Protocol Version 6 (IPv6). Designed as the successor to IPv4, IPv6 addresses the limitations of the legacy protocol by providing a vastly expanded 128-bit address space, simplified packet processing, improved

support for mobility, and enhanced scalability for future network growth. As organizations continue transitioning toward IPv6-enabled infrastructures, the efficiency of routing mechanisms becomes increasingly important for maintaining reliable and high-performance communication across interconnected networks.

Every communication that takes place across a network depends on routing decisions being made efficiently and accurately. Routing protocols are responsible for finding and maintaining these paths, ensuring that data reaches its destination even when network conditions change. In dynamic networking environments, routing protocols must continuously adapt to topology changes, link failures, and varying traffic conditions while maintaining optimal paths for packet delivery. The effectiveness of a routing protocol directly influences the key network performance parameters such as convergence time, latency, packet loss, and overall network stability. Consequently, selecting an appropriate routing protocol is a critical design consideration for enterprise, educational, and service-provider networks operating in IPv6 environments.

Among the routing protocols developed for IPv6 networks, RIPng, OSPFv3, and EIGRP represent three distinct approaches to route discovery and maintenance. RIPng extends the traditional Routing Information Protocol to support IPv6 through a distance-vector routing mechanism. OSPFv3 utilizes a link-state architecture and employs the Shortest Path First (SPF) algorithm to calculate optimal routes. EIGRP, often classified as an advanced distance-vector or hybrid routing protocol, uses the Diffusing Update Algorithm (DUAL) to provide rapid convergence and efficient route selection. Each protocol offers unique advantages and trade-offs in terms of scalability, resource utilization, configuration complexity, and adaptability to network changes.

Although numerous theoretical studies have examined the characteristics of these protocols, practical evaluations conducted on physical IPv6 networks remain valuable for understanding their real-world behavior. Performance can vary depending on network topology, hardware capabilities, and operational conditions, making experimental validation essential for meaningful comparison. Furthermore, many organizations continue to rely on legacy networking equipment that must support

modern IPv6 deployments, creating a need to assess routing protocol performance under realistic infrastructure constraints.

This research presents a comparative analysis of RIPng, OSPFv3, and EIGRP within an IPv6 network. The experimental environment employed Cisco 2811 series routers mounted on standard server racks and interconnected through RJ45-terminated straight-through and crossover LAN cables. Two end-host systems were connected through multiple routers to create a multi-hop IPv6 topology representative of practical enterprise networking scenarios. Each routing protocol was configured independently on the same physical infrastructure to ensure consistency and fairness in performance evaluation.

The study focuses on measuring and comparing critical performance metrics, including end-to-end connectivity, average round-trip latency, packet loss, and route convergence time during both normal operation and simulated network failure conditions. By analyzing the adaptive behavior of RIPng, OSPFv3, and EIGRP in an identical IPv6 environment, this work aims to identify the protocol that delivers the most efficient and reliable performance for multi-router IPv6 networks. The findings provide practical insights for network administrators, researchers, and organizations seeking to optimize routing protocol selection for current and future IPv6 deployments.

II. LITERATURE REVIEW

The performance of routing protocols in IPv6 networks has been widely investigated as organizations continue transitioning from IPv4 to IPv6. Among the most commonly deployed Interior Gateway Protocols, RIPng, OSPFv3, and EIGRP for IPv6 have received significant attention due to their differing routing mechanisms and performance characteristics.

Malkin (1998) introduced RIP and its IPv6 counterpart, RIPng, as a distance-vector routing protocol that uses hop count as its routing metric. While RIPng is simple to configure and manage, several studies have reported slower convergence and limited scalability in larger network environments.

OSPFv3, standardized in RFC 5340 by Coltun et al., extends the OSPF protocol to support IPv6 networks. Unlike RIPng, OSPFv3 employs a link-state routing algorithm and uses Dijkstra's Shortest Path First (SPF) algorithm to determine optimal routes. Researchers have shown that OSPFv3 offers faster convergence and better scalability.

Cisco's EIGRP for IPv6 combines features of both distance-vector and link-state routing through the Diffusing Update Algorithm (DUAL). Studies have demonstrated that EIGRP generally achieves lower delay, faster convergence, and more efficient route selection compared to RIPng and OSPFv3, particularly in dynamic network environments.

Several simulation-based studies using Packet Tracer, GNS3, and OPNET have compared these protocols using metrics such as delay, throughput, packet loss, and convergence time. Most findings indicate that EIGRP provides the best overall performance, while OSPFv3 offers a balance between efficiency and scalability. RIPng remains suitable for smaller networks where simplicity is prioritized over performance.

To the best of the authors' knowledge, limited work has focused on a direct comparative evaluation of RIPng, OSPFv3, and EIGRP for IPv6 using a multi-router topology and ICMP-based performance measurements. Therefore, this study aims to analyze and compare these protocols through practical experimentation and graphical performance analysis in an IPv6 network environment under the supervision of BSNL Laboratory officers.

III. IMPLEMENTATION

3.1. System Overview

The experimental topology consists of two IPv6-enabled end hosts PC1 and PC2 interconnected through two Cisco 1811 routers R1 and R2 configured in a multi-hop network environment. Each host resides in a separate IPv6 subnet and communicates through the intermediate routers. The network was designed to evaluate the performance of three IPv6 routing protocols: RIPng, OSPFv3, and EIGRP for IPv6.

All router interfaces were assigned unique IPv6 addresses using /64 prefixes, while the inter-router links formed the routing backbone of the network. The routing protocols were configured individually on the same topology to ensure a fair performance comparison under identical network conditions. Connectivity between the source and destination hosts was verified using ICMPv6 echo requests, and performance metrics such as delay, packet loss, convergence behavior, and administrative distance were analyzed.

The topology was first simulated using Cisco Packet Tracer and validated through repeated communication tests between PC1 and PC2. The collected results were subsequently represented through graphical analysis to identify the routing protocol that provides the most efficient performance in an IPv6 environment.

Table 1: Network Address Allocation

Device	Interface	IPv6 Address
PC1	NIC	2001:11::2/64
R1	Fa0/0	2001:11::1/64
R1	Fa0/1	2001:12::1/64
R2	Fa0/1	2001:12::2/64
R2	Fa0/0	2001:22::1/64
PC2	NIC	2001:22::2/64

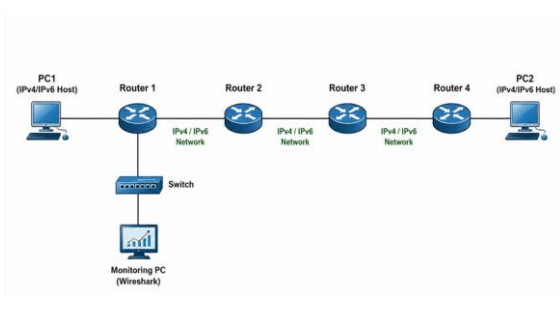


Fig I: System Architecture

3.2 Simulation Phase

Prior to the hardware setup the complete topology arrangement was carried out using Cisco Packet Tracer 8.2 to evaluate the performance of IPv6 routing protocols in a multi-router network environment. The designed topology consisted of two end hosts (PC0 and PC1) connected through four Cisco 2811 routers (Router0, Router1, Router2, and Router3). The routers were interconnected using FastEthernet links, creating a multi-hop communication path between the source and destination.

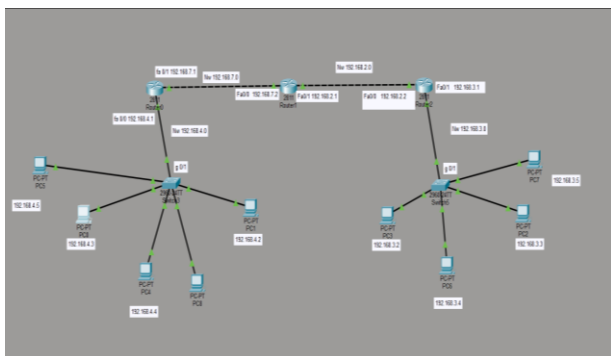


Fig II: Simulated network topology implemented in Cisco Packet Tracer

The Packet Tracer model consisted of three Cisco 2811 routers (Router0, Router1, and Router2) interconnected through two backbone networks, namely 192.168.7.0/24 and 192.168.2.0/24. Two local area networks were connected at either end of the topology. The source LAN (192.168.4.0/24) contained multiple hosts including PC0, PC1, PC4, PC5, and PC8, while the destination LAN (192.168.3.0/24) contained PC2, PC3, PC6, and PC7. Cisco 2960 switches were used to provide connectivity between hosts and routers. The topology was designed to emulate a realistic enterprise network in which data packets traverse multiple routing hops before reaching their destination. Figure 1 illustrates the complete simulated network topology.

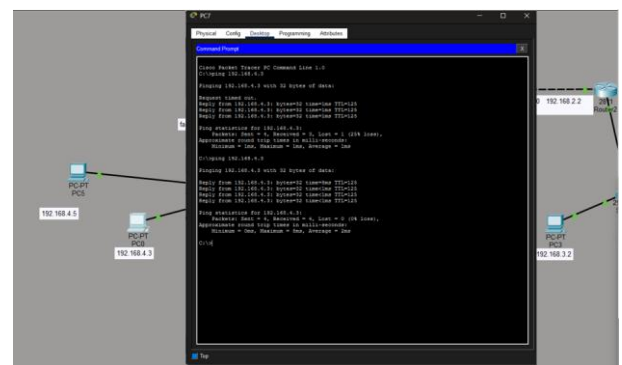


Fig III: Ping verification between hosts showing successful end-to-end connectivity after routing convergence.

Initially, communication between hosts located on different LAN segments was unsuccessful because routing information had not yet been exchanged among the routers. After configuring the selected routing protocols and enabling route advertisement across all participating networks, routing tables were populated dynamically and connectivity was established throughout the network.

To verify successful route propagation, ICMP ping tests were performed between hosts located in different subnets. As shown in Figure 2, PC7 (192.168.3.5) successfully communicated with PC0 (192.168.4.3). During the first ping attempt, one packet timed out due to the route discovery process, resulting in 25% packet loss. Subsequent ping operations achieved 0% packet loss, with round-trip times ranging from less than 1 ms to 8 ms and an average delay of approximately 2 ms. These results confirmed that routing information had converged successfully and that end-to-end communication was functioning correctly.

Following connectivity verification, RIP, OSPF, and EIGRP were configured individually on the same topology. For each protocol, multiple ping tests were conducted and performance metrics such as delay, throughput, convergence behavior, packet delivery, and routing efficiency were recorded. The collected data were then

analyzed and represented graphically to compare the effectiveness of each routing protocol under identical network conditions.

The simulation environment provided a controlled platform for protocol evaluation without the influence of external factors such as cable faults, hardware processing delays, or network congestion. As a result, the simulation phase served as a reliable foundation for analyzing the relative performance of RIP, OSPF, and EIGRP before considering their deployment in real-world networking environments.

3.3.1 Implementation of RIPng

The Routing Information Protocol (RIP) was implemented to establish dynamic routing between the routers in the experimental network. RIP operates as a distance-vector routing protocol and determines the best path based on hop count. Initially, all router interfaces participating in the routing process were assigned appropriate IP addresses and verified for connectivity.

The RIP routing process was then enabled on each router, and the directly connected networks were advertised using RIP network statements. Once activated, the routers periodically exchanged routing updates with neighboring routers, allowing them to learn routes to remote networks automatically. After route propagation, the routing tables were verified to ensure that all network segments were reachable through the dynamically learned routes.

Connectivity testing was performed using ICMP ping messages between hosts located in different LANs. Successful packet transmission confirmed that RIP was correctly exchanging routing information and forwarding traffic across the network. Due to its periodic update mechanism, RIP required a longer convergence time compared to the other routing protocols evaluated in this study.

3.3.2 Implementation of OSPF

Open Shortest Path First (OSPF) was implemented as a link-state routing protocol to provide efficient route calculation and faster convergence. Before enabling OSPF, all router interfaces were configured with their respective network addresses and tested for proper operation.

The OSPF routing process was then activated on each router, and all participating interfaces were assigned to the OSPF area. After configuration, routers established neighbor relationships by exchanging Hello packets and subsequently shared link-state advertisements (LSAs) to build a complete topology database of the network. Using the Shortest Path First (SPF) algorithm, each router

independently calculated the optimal path to every destination network.

The routing tables were examined to verify successful route learning and shortest-path selection. End-to-end connectivity tests were then conducted between hosts on different LANs. The results demonstrated that OSPF converged more rapidly than RIP and provided efficient route computation, making it suitable for larger and more complex network environments.

3.3.3 Implementation of EIGRP

Enhanced Interior Gateway Routing Protocol (EIGRP) was implemented to evaluate its performance in comparison with RIP and OSPF. EIGRP combines characteristics of both distance-vector and link-state routing protocols and uses the Diffusing Update Algorithm (DUAL) to determine loop-free and efficient routes.

Following interface configuration and connectivity verification, the EIGRP routing process was enabled on each router. The relevant network interfaces were included in the EIGRP domain, allowing neighboring routers to establish adjacencies and exchange routing information. Unlike RIP, EIGRP sends partial and triggered updates only when network changes occur, reducing routing overhead and improving network efficiency.

After adjacency formation, routers populated their topology and routing tables with dynamically learned routes. Connectivity tests were performed between source and destination hosts to validate successful packet delivery. The observed results indicated rapid route convergence, low delay, and stable network operation. These characteristics enabled EIGRP to achieve superior performance among the routing protocols examined during the experimental evaluation.

IV. RESULTS

4.1 Delay Comparison

Figure IV illustrates the delay performance of RIP, OSPF, and EIGRP routing protocols. Delay represents the time required for packets to travel from the source node to the destination node across the network. The results show that RIP exhibits the highest delay, while OSPF demonstrates moderate performance. EIGRP achieves the lowest delay among the three protocols.

The higher delay observed in RIP is mainly due to its distance-vector routing mechanism and periodic routing table updates. Since RIP exchanges complete routing information at regular intervals, packet forwarding decisions may take longer to adapt to network changes. OSPF reduces delay by maintaining a complete view of the

network topology and computing optimal paths using the Shortest Path First (SPF) algorithm. EIGRP provides the best delay performance because it utilizes the Diffusing Update Algorithm (DUAL), which enables faster route computation and efficient path selection.

The comparison indicates that EIGRP is more suitable for applications requiring low latency and rapid packet delivery, whereas RIP may be more appropriate for small-scale networks.

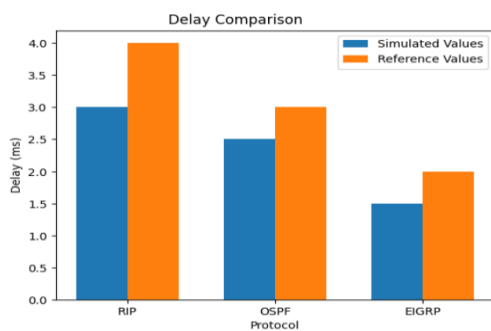


Fig IV: Delay Comparison

4.2 Convergence Time Comparison

Figure V presents the convergence performance of RIP, OSPF, and EIGRP. Convergence time refers to the duration required for all routers in the network to update their routing information and reach a consistent routing state after a topology change.

The graph shows that RIP has the highest convergence value, indicating slower adaptation to network changes. OSPF converges faster than RIP due to its link-state routing architecture, which allows routers to maintain synchronized topology databases. EIGRP demonstrates the fastest convergence among the evaluated protocols because of its DUAL algorithm and the use of backup routes, known as feasible successors. The results obtained from the simulation

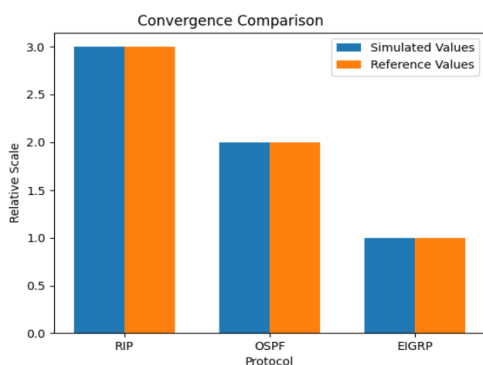


Fig V: Convergence Time Comparison

4.3 Administrative Distance Comparison

Figure VI compares the administrative distance values associated with RIP, OSPF, and EIGRP. Administrative Distance (AD) is a measure used by routers to determine the reliability of routing information received from different routing protocols. Lower AD values indicate higher trustworthiness and greater preference during route selection.

The results show that EIGRP possesses the lowest administrative distance, followed by OSPF and RIP. Consequently, routes learned through EIGRP are preferred over those learned through OSPF or RIP when multiple routing sources are available. OSPF provides a moderate level of route preference, whereas RIP has the highest administrative distance and therefore the lowest priority.

The administrative distance values observed in the simulation closely match standard reference values reported in networking literature, demonstrating the correctness of the implemented routing configurations. These findings further support the superior routing efficiency of EIGRP within the proposed IPv6 network environment.

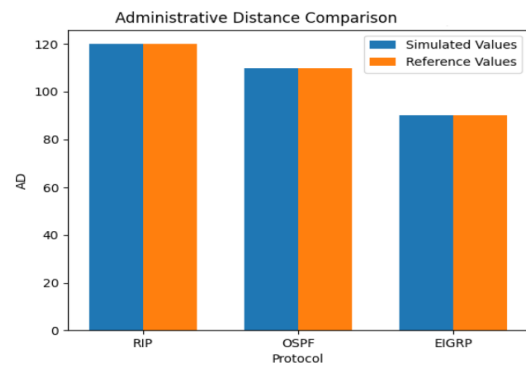


Fig VI: Administrative Distance Comparison

4.3 Network Monitoring and Observations

During the implementation phase, PuTTY software was used to access the Cisco router command-line interface (CLI) and monitor network performance. The software enabled real-time observation of routing table updates, protocol status, interface conditions, and end-to-end connectivity between the configured IPv6 hosts.

Under normal operating conditions, PuTTY ping tests confirmed successful communication between the source and destination PCs for all three routing protocols. ICMPv6 echo requests received corresponding echo replies, indicating stable network connectivity and correct routing table propagation. The routing information displayed through CLI commands verified that RIPng, OSPFv3, and

EIGRP routes were successfully learned and maintained by the routers.

During protocol evaluation, differences in convergence behavior were observed. RIPng required a longer period to update routing information following network changes, whereas OSPFv3 adapted more rapidly through its link-state database mechanism. EIGRP demonstrated the fastest route recalculation and recovery due to the DUAL algorithm, resulting in minimal interruption to packet transmission.

Continuous monitoring through PuTTY also confirmed that EIGRP maintained more efficient route management and generated fewer routing updates compared to RIPng. The observations obtained from CLI outputs were consistent with the simulation results and graphical analysis, where EIGRP achieved lower delay, faster convergence, and improved overall network performance. These observations validated the effectiveness of EIGRP as the most efficient routing protocol among the protocols evaluated in the proposed IPv6 network.

Table 2: Summary of Experimental Result

Feature	RIP	OSPF	EIGRP
Protocol	Distance Vector	Link State	Hybrid
AD Value	120	110	90
Convergence	Slow	Fast	Very Fast
Delay	3-5 ms	3 ms	2-3 ms
Packet Loss	0%	0%	0%

IV. DISCUSSION

The results demonstrate that IPv6 routing can be effectively implemented using RIPng, OSPFv3, and EIGRP in a multi-router network environment. All three protocols successfully established communication between the source and destination hosts, confirming the reliability of the proposed network design.

Among the evaluated protocols, EIGRP delivered the best overall performance with lower delay, faster convergence, and efficient route management. OSPFv3 provided balanced performance and scalability, while RIPng showed comparatively slower convergence due to its periodic routing updates. The observations obtained through PuTTY monitoring were consistent with the graphical analysis and simulation results.

A significant aspect of this work is that the experiments were conducted in the BSNL laboratory environment using Cisco networking equipment. This provided practical exposure to real-world networking practices and enhanced the reliability of the experimental findings. Based on the

obtained results, EIGRP is identified as the most suitable routing protocol for the proposed IPv6 network topology.

V. APPLICATIONS

Three Protocols discussed in the paper are applicable across a range of industries. Table 3 identifies representative industry domains and the preferred protocol for the best outcome.

Table 3: Real World Applications

Domain	Use Case	Preferred Protocol
Internet Service Providers (ISPs)	Large-scale IPv6 backbone routing	EIGRP, OSPFv3 (Fast convergence)
Enterprise Networks	Inter-department IPv6 communication	EIGRP (Low delay, reliability)
Educational Institutions	Campus-wide IPv6 deployment	OSPFv3 (Scalability, standardization)
Data Centers	High-speed server communication	EIGRP (Efficient route selection)
Government Organizations	Secure and scalable IPv6 infrastructure	OSPFv3 (Hierarchical routing)
Smart Grid / IoT Networks	IPv6-enabled sensor communication	EIGRP, OSPFv3 (Efficient routing)

VI. CONCLUSION AND FUTURE SCOPE

This paper presented the simulation and performance evaluation of RIPng, OSPFv3, and EIGRP routing protocols in an IPv6 network environment. The implementation was carried out using Cisco Packet Tracer for simulation and validated through practical experimentation and monitoring using PuTTY in the BSNL laboratory. The results confirmed successful IPv6 communication between end devices and highlighted the performance differences among the routing protocols.

The work carried out in the BSNL laboratory provided practical exposure to real-world networking environments and validated the simulation findings. Future work may include evaluating these protocols in larger IPv6 topologies, studying their performance under heavy traffic conditions, and integrating advanced technologies such as Software Defined Networking (SDN), network security mechanisms, and cloud-based infrastructures. Further analysis using additional performance metrics such as

throughput, jitter, and bandwidth utilization can also provide a more comprehensive evaluation of IPv6 routing performance.

VI. REFERENCES

- [1] X. Zhang, Y. Wang and Z. Li, "TAR: Traffic Adaptive IPv6 Routing Lookup Scheme," Proceedings of the 2024 ACM SIGCOMM Asia-Pacific Workshop on Networking (APNet), pp. 1-10, 2024.
- [2] Z. Mammeri, "Reinforcement Learning Based Routing in Networks: Review and Classification of Approaches," IEEE Access, vol. (to appear), 2019..
- [3] J. Ma, C. Zhu, Y. Fu, H. Zhang and W. Xiong, "Dynamic Routing via Reinforcement Learning for Network Traffic Optimization," Informatica, vol. 47, no. 2, pp. 157-174, 2023.
- [4] B. Kim, J. H. Kong, T. J. Moore and F. T. Dagefu, "Deep Reinforcement Learning Based Routing for Heterogeneous Multi-Hop Wireless Networks," arXiv preprint arXiv:2508.14884, Aug. 2025.
- [5] E. Ancillotti, C. Vallati, R. Bruno & E. Mingozzi, "A reinforcement learning-based link quality estimation strategy for RPL and its impact on topology management," arXiv preprint arXiv:2206.00273, 2022
- [6] A. Sharma, "IPv6 Era Routing Protocol Evolution: Challenges and Innovations," International Journal for Multidisciplinary Research (IJFMR), 2024.