

# Comparative Study of Routing Protocols

Ammar Rangwala  
Computer Engineering Department  
Government Polytechnic Dahod  
Dahod, India

Bhavesh Patil  
Computer Engineering Department  
Government Polytechnic Dahod  
Dahod, India

Paresh Patel  
Computer Engineering Department  
Government Polytechnic Dahod  
Dahod, India

**Abstract** - Routing protocols are fundamental to the efficient and reliable transfer of data across networks, ensuring optimal path selection for packet delivery. This paper presents a comparative analysis of three prominent routing protocols: Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP). RIP, an early distance-vector protocol, is well-suited for small networks but struggles with scalability and convergence. OSPF, a link-state protocol, offers improved scalability and rapid convergence, making it ideal for large and complex networks. BGP, essential for interdomain routing, enables the exchange of routing information across different Autonomous Systems (AS), supporting the vast scale of the internet. By examining the underlying algorithms, operational features, and use cases of these protocols, this paper provides insights into their strengths, limitations, and optimal applications in modern networking environments.

**Keywords** – Autonomous System (AS), Interior Gateway Protocol (IGP), Exterior Gateway Protocol (EGP), Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Border Gateway Protocol (BGP)

## I. INTRODUCTION

Routing protocols are the backbone of modern network infrastructures, enabling the efficient and accurate transfer of data across complex and dynamic systems. These protocols determine the best possible paths for data packets to travel, ensuring not only speed and reliability but also optimized use of network resources. As networks grow in size and complexity, selecting the right routing protocol becomes increasingly critical to maintaining high performance, scalability, and security.

A key component of networking is the concept of an Autonomous System (AS), which refers to a collection of IP networks managed by a single organization with a unified routing policy. Within an AS, routing is handled by Interior Gateway Protocols (IGPs), while routing between different ASes is managed by Exterior Gateway Protocols (EGPs). IGPs focus on optimizing routes within a localized network environment,

whereas EGPs ensure smooth data transfer between networks on a global scale [1].

This paper focuses on three widely-used routing protocols: Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) [2].

To understand how routing protocols address the challenges of modern networking, we begin by exploring the earliest and simplest protocol, RIP. Despite its simplicity, RIP laid the foundation for more advanced protocols and provides key insights into the evolution of routing technologies.

## II. ROUTING INFORMATION PROTOCOL

It is one of the earliest and simplest distance-vector routing protocols designed for small to medium-sized networks. It was first standardized as part of the IGP suite and is widely used in local and enterprise networks to enable routers to dynamically exchange routing information. RIP is suitable for networks with a relatively small number of routers and moderate traffic, as it relies on the hop count as the primary metric for determining the best path to a destination [6]. RIP operates at Layer 3 (Network Layer) of the OSI model and supports both IPv4 (RIPv1, RIPv2) and IPv6 (RIPng). Despite its simplicity, RIP has significant limitations, which include slow convergence and a maximum hop count of 15, which makes it unsuitable for larger, complex networks [3]. At the heart of RIP lies the Distance Vector Algorithm, also known as the Bellman-Ford algorithm [3]. The core idea is that each router maintains a table (routing table) that holds the best-known routes to all possible destinations.

This table includes:

- Destination IP address
- Next hop router
- Metric (hop count)

Each router periodically shares its routing table with its immediate neighbours, allowing them to learn about new routes or update existing ones.

#### A. Steps of the Distance Vector Algorithm

1. Each router initializes its routing table with a route to itself, with a hop count of 0.
2. Routers periodically (every 30 seconds by default) exchange their routing tables with neighbouring routers. This information exchange is known as RIP updates.
3. Upon receiving updates from neighbours, a router checks if a better route to a destination exists. A better route is defined as a route with a lower metric (hop count). If a better route is found, the router updates its table with this new route.
4. Over time, as routers share and update information, the network converges, meaning that all routers have consistent and accurate routing information.

RIP operates by broadcasting or multicasting its routing table updates to all neighbours at regular intervals. It follows a hop-count metric to determine the shortest path to a destination, where each hop between routers is assigned a cost of 1. However, the maximum allowable hop count is 15, and any destination that requires 16 or more hops is considered unreachable. This limitation makes RIP suitable only for small to medium-sized networks.

RIP uses a set of timers to manage routing updates and maintain the freshness of routing information [6]:

- Update Timer: Sends routing table updates every 30 seconds.
- Invalid Timer: Marks a route as invalid if no updates for the route are received within 180 seconds.
- Flush Timer: Removes a route from the routing table if no updates are received for 240 seconds.

To prevent routing loops, RIP implements a technique known as Split Horizon, which ensures that a router does not advertise a route back to the neighbour from which it learned that route.

Another mechanism for loop prevention is Route Poisoning, where a router advertises an unreachable route (hop count of 16) to inform other routers that the route is no longer valid.

In the event of a network change, such as a route becoming unavailable, RIP sends a triggered update instead of waiting for the regular update interval, speeding up convergence.

While RIP is easy to implement in small networks, its limitations in scalability and convergence time make it less suitable for larger, dynamic networks. This led to the development of more sophisticated protocols like OSPF, which addresses many of RIP's shortcomings through a link-state approach.

### 3. OPEN SHORTEST PATH FIRST

OSPF is a link-state routing protocol designed to overcome the limitations of distance-vector protocols like RIP. OSPF is widely used in large enterprise networks due to its scalability, efficiency, and fast convergence [4]. Unlike RIP, OSPF does not use hop count as a metric but instead calculates the cost of each link based on bandwidth. This allows OSPF to find the most efficient path between routers.

OSPF operates within a single AS and is classified as an IGP. It is designed to be scalable and highly efficient in handling complex topologies with multiple routers and links. It uses a hierarchical structure with areas to optimize traffic and minimize resource consumption.

OSPF uses the Link-State Algorithm, which is based on Dijkstra's Shortest Path First (SPF) algorithm [1]. This algorithm enables OSPF to build a complete map of the network by exchanging link-state information with neighbouring routers. Each OSPF router maintains a Link-State Database (LSDB) that contains the network's topology, allowing it to calculate the shortest path to each destination.

#### A. Steps of the Link-State Algorithm

1. Each router creates a Link-State Advertisement (LSA) that contains information about its links, including the state of the link (up or down) and the cost associated with it.
2. OSPF routers exchange LSAs with their neighbours. The LSAs are flooded throughout the network so that each router has an identical view of the network's topology.
3. Each router builds its own Link-State Database based on the received LSAs. This database represents a complete map of the network's links and nodes.
4. Using the Dijkstra algorithm, each router calculates the shortest path to each destination in the network. The result is stored in the router's Routing Information Base (RIB), which is used to forward packets.
5. Once all routers have a synchronized view of the network and have calculated the shortest paths, the network is considered to have converged.

OSPF operates with a hierarchical structure [7] that helps improve scalability and efficiency. The network is divided into multiple areas, which reduces the size of the link-state databases and limits the scope of LSA flooding.

OSPF divides large networks into multiple areas as shown in Fig 1. Each area is connected to a central backbone area (Area 0). This hierarchical design minimizes LSA flooding and reduces processing overhead.




Figure 1 OSPF Hierarchical Design

In networks with multiple routers, OSPF elects a Designated Router (DR) and a Backup Designated Router (BDR) to reduce the number of adjacencies and the amount of LSA flooding. The DR manages communication between routers on a shared network segment, and the BDR takes over if the DR fails.

OSPF uses a cost metric based on link bandwidth. The cost is inversely proportional to the bandwidth of the link [7], meaning higher-bandwidth links have lower costs. This allows OSPF to choose the most efficient path for routing traffic.

OSPF quickly detects network changes, such as link failures, and recalculates the shortest paths using Dijkstra's algorithm. This fast convergence is critical in large, dynamic networks.

OSPF routers use the "Hello Protocol" to discover and maintain adjacencies with their neighbours. Hello packets are sent periodically to detect link failures, and routers that do not receive Hello packets within a certain interval declare the link down.

Although OSPF excels within an Autonomous System, managing routing between different Autonomous Systems across the internet requires a protocol designed for interdomain routing. This is where the BGP plays a critical role, enabling communication between networks with distinct policies and ensuring the stability of global internet routing.

#### 4. BORDER GATEWAY PROTOCOL

The Border Gateway Protocol (BGP) is the protocol that governs how packets are routed across the internet. It is classified as an Exterior Gateway Protocol (EGP) and is responsible for exchanging routing information between different ASes. An AS is a collection of IP networks and routers under the control of a single

organization that presents a common routing policy to the internet.

Unlike interior protocols such as RIP or OSPF, which operate within a single AS, BGP manages routing between multiple ASes, making it essential for interdomain routing across the global internet [5]. BGP is designed to handle vast, dynamic networks with millions of routes and offers fine control over routing decisions based on policies rather than purely metric-based calculations like hop count or bandwidth.

BGP uses a Path-Vector Algorithm, which is an extension of the distance-vector algorithm [5]. Instead of tracking just the distance to each destination, BGP tracks the full path (the sequence of ASes) that packets take to reach their destination. This enables BGP to avoid routing loops and enforce routing policies [5].

##### A. Steps of the Path-Vector Algorithm

1. Each BGP router advertises the network prefixes it can reach, along with the full AS path to those destinations. This path information allows BGP routers to make more informed routing decisions.
2. When a router receives multiple paths to a destination, it evaluates them based on a variety of attributes, such as AS path length, next hop IP address, and local routing policies. The router then selects the best path and advertises this to its neighbours.
3. By tracking the complete AS path, BGP can detect and prevent routing loops. If a router sees its own AS in the AS path of a received route, it discards the route to avoid a loop.
4. BGP routers continually exchange path information and update their routing tables, allowing the network to converge on stable paths to all destinations.

BGP operates over Transmission Control Protocol (TCP), which ensures reliable delivery of routing updates between BGP peers. BGP forms connections between routers in different Autonomous Systems, called External BGP (eBGP), as well as between routers within the same AS, known as Internal BGP (iBGP).

Each route advertisement includes an AS Path attribute [9], which lists all the ASes that the route has traversed. This attribute is used to detect routing loops and to prefer shorter paths.

BGP also uses the Next Hop attribute, which indicates the IP address of the next router along the path to the destination. This allows routers to correctly forward packets.

BGP does not rely on simple metrics like hop count or bandwidth. Instead, it uses a series of attributes to determine the best path.

One of BGP's key strengths is its ability to implement complex routing policies. Network administrators can define policies to control which routes are advertised, which routes are accepted, and how traffic should be routed.

BGP establishes peering sessions between routers. When two routers form a BGP connection, they exchange full routing tables at the start of the session and subsequently send incremental updates to reflect network changes.

BGP is designed to scale to very large networks, such as the internet. It is capable of handling millions of routes and is highly efficient in terms of memory and CPU usage [8].

Each of the three protocols—RIP, OSPF, and BGP—has its own strengths and limitations, tailored to different network environments and requirements. To better understand when and where each protocol is most effective, we now present a direct comparison of their key features and operational characteristics.

### 5. COMPARISON

Feature	RIP	OSPF	BGP
Algorithm	Bellman-Ford (Distance Vector)	Dijkstra's Algorithm (Shortest Path First)	Path-Vector Algorithm
Convergence Time	Slow	Fast	Moderate
Hierarchical Design	No	Yes	No
Loop Prevention Mechanism	Split Horizon, Route Poisoning	Link-state routing and SPF calculation	AS Path information prevents loops
Maximum Hop Count	15	No Limit	No Limit
Metric	Hop Count	Cost (based on link bandwidth)	Policy based
Network Size	Small to Medium	Large Enterprise	Internet Scale
Protocol Type	Distance Vector	Link State	Path Vector
Scalability	Low	High	Very High
Security	Basic	Moderate	High
Update Mechanism	Periodic Broadcasts	Triggered Updates	Incremental Updates
Use Case	Small networks, educational purposes	Large enterprise networks, data centers	Internet-wide routing, ISPs

With an understanding of the specific capabilities and applications of RIP, OSPF, and BGP, it is clear that choosing the right routing protocol depends on the scale and needs of the network. In the following conclusion, we summarize the critical takeaways from our analysis and highlight the significance of selecting appropriate routing protocols for different networking contexts

### CONCLUSION

In this paper, we examined three prominent routing protocols: RIP, OSPF, and BGP, and highlighted their distinct characteristics, strengths, and limitations.

RIP, while straightforward and easy to implement, is best suited for small networks due to its limited scalability and slow convergence. OSPF, on the other hand, offers significant improvements in scalability, faster convergence, and network efficiency, making it ideal for large enterprise environments. BGP, as the primary protocol for interdomain routing, is crucial for

managing the vast and complex infrastructure of the global internet. Its ability to handle millions of routes and enforce policy-driven routing decisions makes it indispensable for large-scale networks and Internet Service Providers (ISPs).

The choice of a routing protocol is not a one-size-fits-all decision. Each protocol is optimized for specific environments—RIP for small networks, OSPF for large enterprises, and BGP for interdomain routing across the internet.

In conclusion, the evolution of routing protocols from RIP to OSPF and BGP reflects the growing complexity and scale of modern networking. As network demands continue to expand, selecting the appropriate routing protocol will remain a critical factor in ensuring network performance and reliability.

### REFERENCES

- [1] Doyle, J., & Carroll, J. (2005). Routing TCP/IP, Volume I (2nd ed.).
- [2] Kurose, J. F., & Ross, K. W. (2016). Computer Networking: A Top-Down Approach (6th ed.).
- [3] Malkin, G. (1998). RFC 2453: RIP Version 2. Internet Engineering Task Force (IETF).
- [4] Moy, J. (1998). RFC 2328: OSPF Version 2. Internet Engineering Task Force (IETF).
- [5] Rekhter, Y., & Li, T. (2006). RFC 4271: A Border Gateway Protocol 4 (BGP-4). Internet Engineering Task Force (IETF).
- [6] Cisco Systems. (2006). Introduction to RIP.
- [7] Cisco Systems. (2016). OSPF Design Guide.
- [8] White, R. (2005). Optimal Routing Design.
- [9] Halabi, B. (2000). Internet Routing Architectures (2nd ed.).
- [10] Tanenbaum, A. S., & Wetherall, D. J. (2011). Computer Networks (5th ed.).