International Journal of Engineering Trends and Technology
ISSN: 2231-5381/ https://doi.org/10.14445/22315381/IJETT-V74I7P101

Volume 74 Issue 7, 1-8, July 2026

Short Communication

Simulation and Analysis of Multicast Traffic through IGMP: A Comparative Study with Physical Network Measurements

Adrian Shehu[1], Fatmir Basholli[2], Lea Likaj[3], Bexhet Kamo[4]

[1]TCN, Tirana, Albania.
[2]Department of Engineering, European University of Tirana, Tirana, Albania.
[3]Technical Department, DigitAlb, Tirana, Albania.
[4]Department of Electronics and Telecommunications, Polytechnic University of Tirana, Tirana, Albania.

[1]Corresponding Author : adshehu@tcn.al

Received: 27 October 2025 Revised: 11 May 2026 Accepted: 03 June 2026 Published: 28 July 2026

Abstract - *Multicast transmission is a fundamental mechanism used to deliver content to multiple receivers/users and is still an essential method in television broadcasting services and data distribution networks. Efficient management of multicast traffic has a direct impact on bandwidth consumption and helps to prevent network congestion, especially during high-demand periods that are common in the broadcasting industry. This article investigates the effectiveness of the Internet Group Management Protocol (IGMP) on optimizing the multicast traffic and the utilization of resources by creating a representative simulated network in Cisco Modeling Labs (CML), based on previous experiments and extended versions of them that are carried out on physical devices. The simulated model allows the evaluation and monitoring of two key performance metrics, including bit rate and packet rate, for two operating conditions: case one, with IGMP enabled, and case two, with IGMP disabled. Results show an average of 79.4% reduction in forwarded multicast traffic on active ports when IGMP is enabled, and a 4.85× increase in bandwidth usage when IGMP is disabled. These results confirm how much control IGMP provides in reducing unnecessary data replication and its direct impact on improving network efficiency. The alignment between the simulated and real-world results strengthens the reliability of the results. The combined evidence provides a solid foundation for understanding the role of IGMP in controlling and improving multicast traffic by offering practical guidance for designing and implementing multicast communication systems in networking and broadcasting domains.*

Keywords - *Bandwidth optimization, IGMP Protocol, Multicast communication, Network Simulation, TV Broadcasting.*

1. Introduction

In the digital broadcasting industry, where live television and video content must reach large audiences with thousands of viewers, even small inefficiencies, due to many reasons, can scale into significant risks of network congestion. Therefore, proper management of infrastructure resources is essential to ensure service sustainability.

Multicast is a communication method that sends a single data stream to multiple destinations simultaneously. It includes a method of delivery and discovery of senders and receivers of multicast data, which are transmitted using IP multicast addresses, called groups [1]. By sending a single stream to many endpoints/users, multicast reduces unnecessary load on the network and, for this reason, is widely adopted in IPTV and other media distribution systems. However, without proper control mechanisms, multicast traffic may be forwarded even to inappropriate downstream interfaces. As a result, bandwidth and device processing capacities that could support other applications become consumed and increase the operational expenses for service providers, and at the same time, potentially degrade the quality of service or raise the costs for customers. Such practices are not only limited to media and entertainment industries, but also appear in other large-scale or one-to-many data distribution services, like video conferencing and IoT networks, as mentioned in prior studies [14]. These challenges evident the need for measurable and reproducible evaluations of multicast control mechanisms, such as IGMP, which is our motivation for this study.

The Internet Group Management Protocol, known as IGMP, addresses these issues and ensures that the multicast traffic is forwarded only to interested parties. Routers and switches in the network, using port information, manage in a dynamic way the multicast group memberships through the exchange of control messages. Hosts send 'Membership Reports' to routers to join a group or 'Leave Group' messages when they no longer wish to receive data. The router then uses periodic 'Queries' to confirm group memberships and maintains a list of which devices are interested and in which multicast addresses. This dynamic control reduces the consumption of bandwidth in a multi-access LAN environment to avoid flooding the entire VLAN [2]. Although IGMP is widely implemented, its benefits are rarely analysed under controlled experimental conditions. IGMP is particularly useful in multimedia applications, such as LAN TV, desktop conferencing, and collaborative computing, that use multipoint communication (that might be from one-to-many hosts or originate from many hosts and with destinations to many other hosts, known as many-to-many) [1].

However, despite its widespread adoption, IGMP's impact is rarely validated by using controlled experimental conditions. Prior studies are based on either a theoretical overview of the protocol or observations taken from production networks, where conditions vary, and results cannot be easily isolated.

Can a virtualized simulation environment reproduce the same optimization patterns observed in real hardware, and can this environment be generalized or scaled for broader applications? This study fills that gap by combining both theory and controlled experimentation and pursuing the following objectives:

- Quantify the impact of IGMP.
- Provide a reference that future research can use when designing or optimizing multicast systems.

The aim of this article is to recreate a real physical setup inside a simulation environment and offer a stronger and measurable understanding of the protocol's functionality and its contribution to reduce the redundant traffic. It demonstrates that the optimization patterns seen in production can be reproduced in an idealized simulation model. This makes the results reproducible, generalizable, and suitable as a reference point for future multicast research and operational planning.

2. Materials and Methods

This article evaluates the impact of the IGMP protocol, based on the measurement methods and techniques from prior physical experiments. Two complementary analysis approaches were considered:

- Port-level traffic analysis, used to track traffic volume and link utilization across switch interfaces;
- Packet-level analysis, focused on captured streams, allowing a further review of packet forwarding behavior and protocol exchanges.

Both approaches were analyzed under two operational modes:

1. IGMP Enabled: the switch actively managed multicast group memberships;
2. IGMP Disabled: The switch functioned in flooding mode, sending multicast traffic throughout all ports in the VLAN. [2].

For each mode and analysis approach, identical preprocessing steps were applied to produce clean and comparable datasets. Quantitative evaluations were then carried out to examine how packet forwarding mechanisms affected delivery and bandwidth consumption.

2.1. Simulation Environment

All experiments were performed in a virtualized lab environment using Cisco Modelling Labs (CML). Figure 1 shows the network simulation topology, which consisted of three nodes linked, with the following roles:

- Video-Server: Multicast video source, running Ubuntu Server 24.04.2 LTS operating system, responsible for streaming three simultaneous video channels.
- Switch: identified as an IGMP Device, a Cisco Layer 2 switch handling multicast group management, IGMP protocol support, and querier functionality. The function of the switch as an IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to extract group membership information. (This function applies for IGMP on VLANs having an IP address, when there is no other device in the VLAN, such as a multicast router to act as a querier) [1].
- Receiver Nodes (User-1 and User-2): Clients subscribing to specific channels available in the network, both running Ubuntu Server, identical to the system used for the video source, and handling the role of IGMP Hosts.

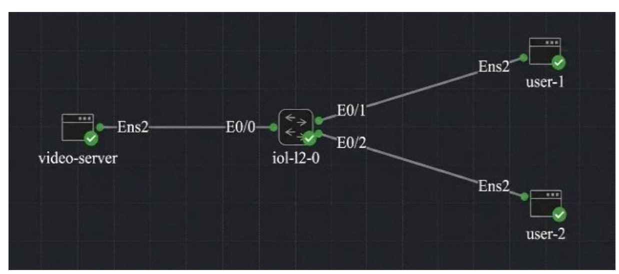

Fig. 1 Network simulation configuration in CML showing nodes, links and their respective roles

Configurations and service management tasks were executed exclusively via the command line, ensuring that measurements were purely analytical and not affected by user-experience factors such as displayed video quality or media player user interface. The CML simulation, operated within a virtual machine provisioned with 8 GB RAM, was distributed among all applied nodes.

2.2. Streaming Configuration

The streaming content consisted of three video files, representing three multicast channels. Each stream was encoded using H.264 video codec at 500kbps bitrate and AAC audio codec at 64 kbps, configured with preset "very-fast". Streams were encapsulated in MPEG-TS containers and encoded using low-profile parameters, adopted by broadcasters to balance content quality with resilience to network fluctuations. Using fixed encoding parameters, maintained an established load across all test executions.

During each experiment, the following channels were transmitted to their respective multicast UDP addresses:

- Channel 1: udp://239.1.1.1
- Channel 2: udp://239.1.1.2
- Channel 3: udp:// 239.1.1.3

User-1 subscribed to Channel 1, and User-2 subscribed to Channel 2, replicating a typical broadcasting scenario.

CML was selected as the experimental environment due to its support for virtual machines, running Ubuntu OS, which provides a reliable platform for executing of FF-mpeg. FF-mpeg is an open-source multimedia framework, widely used for processing video and audio streams [7]. It allows encoding, decoding, transcoding, streaming, and playback of multimedia content by providing a reliable tool for managing video streams in experimental and production environments. Integrating FF-mpeg was essential for handling encoding, streaming, and playback processes for all nodes. The video server used FF-mpeg to generate and transmit the streams, while the end users used FF-play to receive them. The simulated Cisco iol-l2 switch in CML provided four Gigabit Ethernet (10/100/1000 Base-TX) interfaces operating in full-duplex mode: one interface linked to the streamer, two interfaces to receiver nodes, and one left intentionally unconnected. Provided switch image, fully supported IGMP functionality, and running IGMP traffic control features.

2.3. Measurement Procedure

Two key performance metrics were recorded for port-level and packet-level analyses, each reflecting a different aspect of network efficiency:

1. Bit rate (bps or kbps)
2. Packet rate (bps or kbps)

Packet captures were collected from each interface within CML for detailed and further verification. Metrics used in port-level analysis represent traffic volume transmitted through each interface over time, and they are important because, through them, we can quantify how much unnecessary traffic is eliminated when IGMP is active. When combined with packet captures from every interface, they provide detailed verification of IGMP's effect on the network. In that way, we ensure that the results are not dependent on a single measurement but are supported by consistent patterns across all data points.

The following steps summarize the complete experimental procedure:

1. Initiate three multicast streams from the video server.
2. Configure receiver nodes to join their respective multicast groups.
3. Alternate IGMP functionality (enabled or disabled) on the switch, to compare operational modes.
4. Collect traffic metrics and packet captures for each operational mode.
5. Repeat experiments, to confirm consistent of the results, to evaluate the performance parameters, and their variability.

3. Results and Discussion

3.1. Port-Level Analysis

To evaluate in an accurate way the multicast streaming behavior and IGMP's traffic control features under both conditions (IGMP Enabled and IGMP Disabled), five identical tests were performed using the same configuration scripts and node parameters. Each test lasted approximately five minutes, allowing the network to reach a steady state and minimize transient traffic variations. Across all runs, measurements showed negligible variation, confirming a stable and appropriate simulation environment. Due to the deterministic nature of the simulation, repeated tests produced approximate values. Results, as in the tables below, represent mean values calculated across all tests performed, providing a reliable basis for analysis.

Two primary performance metrics were considered:

1. Bit rate (bits/sec): Represents transmitted/received data seen by the interfaces, measured as the average input/output rate over the last five minutes.
2. Packet rate (packets/sec): Represents the number of packets transmitted per Second across the corresponding interfaces.

Table 1. Traffic measurements on switch ports, bit rate, and packet rate, IGMP enabled operational mode

	Bit rate (kbits/sec)	Packet rate (packets/sec)
Video-Server	1401	166
User-1, Channel 1	469	64
User-2, Channel 2	462.3	64
Unused Port	0	0

With IGMP Enabled, as shown in Table 1, the video-server transmitted 1,401,000 bps (≈1,401 kbps), while the two ports connected to receivers recorded 469,000 bps (≈469 kbps) and 462,333 bps (≈462 kbps) respectively (average traffic for receivers ≈466 kbps). The unused switch port did not record any data (0 kbps). Combined traffic forwarded to receiver ports totaled 931,333 bps (931.3 kbps), representing the amount of traffic being managed efficiently through IGMP. Both metrics maintained uniform patterns across test repetitions, with the packet rate metric reflecting a similar trend.

Table 2. Traffic measurements on switch ports, bit rate, and packet rate, IGMP disabled operational mode

	Bit rate (kbits/sec)	Packet rate (packets/sec)
Video-Server	1633.7	185
User-1, Channel 1	1523.7	177
User-2, Channel 2	1523.3	178
Unused Port	1471.3	171

When IGMP was disabled, as shown in Table 2, traffic distribution changed significantly. The three observed edge ports recorded bitrates of 1,523,557 bps, 1,523,333 bps, and 1,471,333 bps, respectively, with the total forwarded traffic reaching ≈4,518.3 kbps. The port connected to the video server carried a slightly higher load, since it transmitted all three multicast streams simultaneously.

A comparison of these totals shows that with IGMP-enabled mode, the traffic delivered to receiver ports represents only 20.6% of the traffic observed without IGMP, indicating a *≈79.4% reduction* in unnecessary data delivery. Conversely, the network consumed *≈4.85x* more bandwidth when IGMP was inactive.

The absolute bandwidth, saved in this scenario, equals ≈3,587 kbps, that would otherwise be wasted by forwarding streams to all switch ports. These findings confirm IGMP's operating mechanism: When active, the switch forwards multicast streams only to interfaces with interested subscribers. In contrast, disabling IGMP causes the switch to flood the multicast group to all ports. Figure 2 graphically shows a comparison of the total port-level traffic for considered modes.

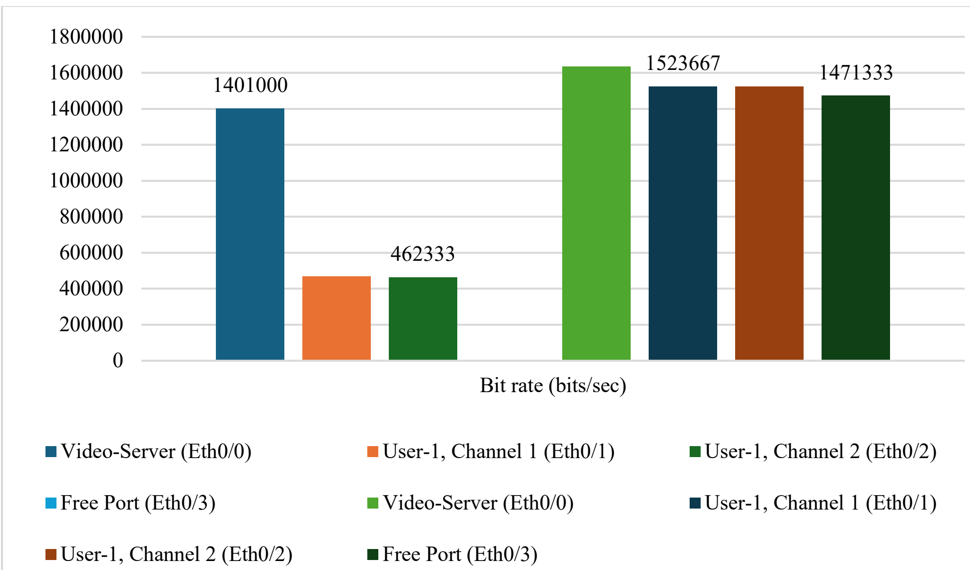

Fig. 2 Comparison of traffic measurements on switch ports, under both operational modes, IGMP enabled (left) and IGMP disabled (right)

3.1.1. Comparison with Real Network Measurements

When compared with the real network data, simulation results demonstrated a consistent pattern in interface occupancy. In each environment, enabling IGMP allowed selective forwarding to subscriber ports only, whereas by disabling it caused the switch to indiscriminately transmit data to all interfaces.

Although traffic values in the simulation were lower than those measured in the physical setup, the overall operational pattern remained the same.

Table 3, altogether with Figure 3, presents measurements and configuration from the physical network used as reference.

Table 3. Reference port-level traffic measurements from physical network

Port (IGMP Enabled)	Traffic (kbps)
HP-1, Port 1 (Source, Channel A, B, C)	1981
HP-1, Port 2 (Client-1, Channel A)	664
HP-2, Port 1 (Client-2, Channel B)	685
Port (IGMP Disabled)	
HP-1, Port 1 (Source, Channel A, B, C)	2125
HP-1, Port 2 (Client-1, Channel A)	2130
HP-2, Port 1 (Client-2, Channel B)	2117

Minor variations between two environments can be attributed to several factors:

1. In the simulation, the switch forwards traffic to all ports when IGMP is disabled (including unused ones), whereas in physical implementations, certain firmware or hardware may limit this practice.
2. The virtual network configured is more idealized, with fewer intermediate devices and more stable links. It isolates multicast traffic, while real networks carry multiple types of traffic simultaneously.
3. Hardware and software differences, such as device models, firmware, or IGMP implementation, can also affect the mentioned mechanisms.

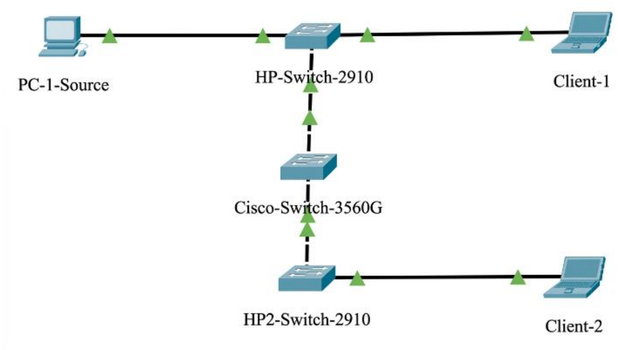

Fig. 3 Physical network configuration of the reference setup used for validating the simulation environment

Overall, the similarity between simulated and real network results highlights the reliability of the CML environment for reproducing switching mechanisms, correctly handling protocols, and evaluating multicast network performance.

3.2. Packet-Level Analysis

To complement port-level observations, a detailed packet-level analysis was considered by using packet captures from Cisco Modelling Labs (CML). This approach, focused on MPEG Transport Stream (MPEG-TS) packets, which carry encapsulated video and audio data. MPEG-TS plays a critical role in the integration of modern broadcasting technologies, particularly in IP-based networks and Digital Video Broadcasting (DVB) standards. [8]

Network configurations remained identical to the previous setup. Packet captures that were examined in Wireshark to observe traffic distribution, under both IGMP modes, and to verify IGMP's operating mechanism, was correctly followed.

Because capture durations varied slightly between links, packet counts were normalized over time to ensure accurate rate comparisons, independent of timing differences:

3.2.1. IGMP Enabled:

- Video Source: 9,409 / 282 s ≈ 33.4 pps
- Receiver 1: 4,290 / 382 s ≈ 11.2 pps
- Receiver 2: 3,960 / 360 s ≈ 11.0 pps

3.2.2. IGMP Disabled

- Video Source: 9,435 / 282 s ≈ 33.5 pps
- Receiver 1: 10,586 / 319 s ≈ 33.2 pps
- Receiver 2: 10,631 / 320 s ≈ 33.2 pps

Table 4. Packet-level overview during IGMP enabled operational mode

	Packets	**Duration (sec)**	**packets per second (pps)**
Video-Server	9409	282	33.4
User-1, Channel 1	4290	382	11.2
User-2, Channel 2	3960	360	11.0

As detailed in Table 4, when IGMP was enabled, User-1 recorded 4,290 MPEG-TS packets over 382 seconds, corresponding to its subscribed Channel 1 (239.1.1.1), while User-2 recorded 3,960 packets over 360 seconds for accessing Channel 2 (239.1.1.2).

The video source transmitted 9,409 packets over 282 seconds, representing the combined traffic of all multicast channels, being streamed over the network.

Table 5. Packet-level overview during IGMP disabled operational mode

	Packets	**Duration (sec)**	**packets per second (pps)**
Video-Server	9435	282	33.5
User-1, Channel 1	10586	319	33.2
User-2, Channel 2	10631	320	33.2

In contrast, with IGMP Disabled, User-1 recorded 10,586 packets, and User-2 recorded 10,631 packets during similar durations (≈320 seconds each), while the source transmitted 9,435 packets in total, showed in Table 5. Despite the source rate, remaining constant to its rate with IGMP active, receiver nodes processed nearly three times more packets. Packet counts indicate that in the absence of IGMP, the source continued to send all multicast traffic to every connected host.

Normalizing by time shows that each end user's packet rate is increased from approximately 11 pps (IGMP Enabled) to 33 pps (IGMP Disabled), which represents a *3x factor* increase per-receiver packet load. Collected results correspond to results from port-level findings. Enabling IGMP reduced packet load on each receiver to one-third (1/3) of the opposite flooding mode.

Figure 4 graphically presents the normalized packet-level rates comparison for the considered modes. When IGMP packets were filtered out, their control exchanges became visible by following the standard message flow, as described in the documentation of [1, 3]. The capture included 'Membership Query' packets from the multicast router, 'Membership Report' messages from the subscriber joining respective groups, and 'Leave Group' messages when the stream was terminated. By integrating packet-level and port-level analyses, demonstrate that IGMP optimizes multicast delivery by forwarding data exclusively from the source to interested receivers.

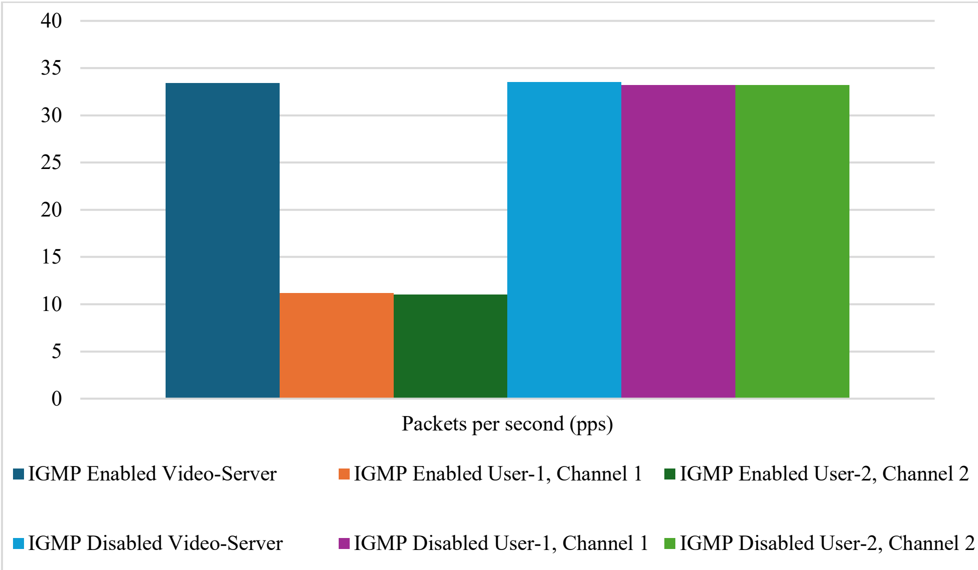

Fig. 4 Normalized Packet-level rates for each link, under both operational modes

3.3. System Resources Utilization

A critical factor in any industry network deployment is the efficient use of system resources. In large-scale environments, with thousands of concurrent subscribers, resource demand increases considerably.

Proper resource planning is crucial to ensure reliable service delivery, especially during peak broadcasting periods when backup capacity becomes vital to prevent service disruptions.

As mentioned previously, the simulation environment was executed within a Virtual Machine (VM) hosting Cisco Modeling Labs. The VM was provisioned with 8 GB of RAM in total, of which 2 GB was assigned to each virtual node, added to the network, and a smaller allocation to the simulated layer-2 switch.

When IGMP was active, as shown in Table 6, video-server transmitting multicast channels used 71% CPU and 65% RAM, while receiver nodes each revealed moderate usage, 6% CPU and 33% RAM. The switch showed negligible usage of 4% CPU and 19% RAM.

Table 6. Devices utilization, CPU and RAM, IGMP enabled operational mode

	CPU Usage	RAM Usage
Video-Server	72%	65%
Cisco L2 Switch	4%	19%
User-1	6%	33%
User-2	6%	33%

With IGMP disabled, Table 7 shows that the video server reached 102% CPU usage, indicating a fully utilized processor, a condition that can lead to its saturation, packet drops, encoding/decoding delays, and degraded Quality of Service.

Receiver nodes also experienced higher utilization levels (9% CPU, 48% RAM), while the switch showed a minor increase (6% CPU). These readings reflect additional processing workload, caused by flooding to all interfaces in the absence of IGMP control.

Table 7. Devices utilization, CPU and RAM, IGMP disabled operational mode

	CPU Usage	RAM Usage
Video-Server	102%	94%
Cisco L2 Switch	6%	16%
User-1	9%	48%
User-2	9%	48%

Figure 5 provides a graphical representation of device load utilization with IGMP Enabled and Disabled. It is important to note that these results are affected by the virtualized environment. In this setup, each node operated as a virtual instance with CPU and memory resources, limited by the host. To have a better understanding of the scalability is suggested to make tests on the dedicated hardware, by using scenarios with higher loads. An alternative way is to perform tests with larger simulated networks that may be configured as more complicated.

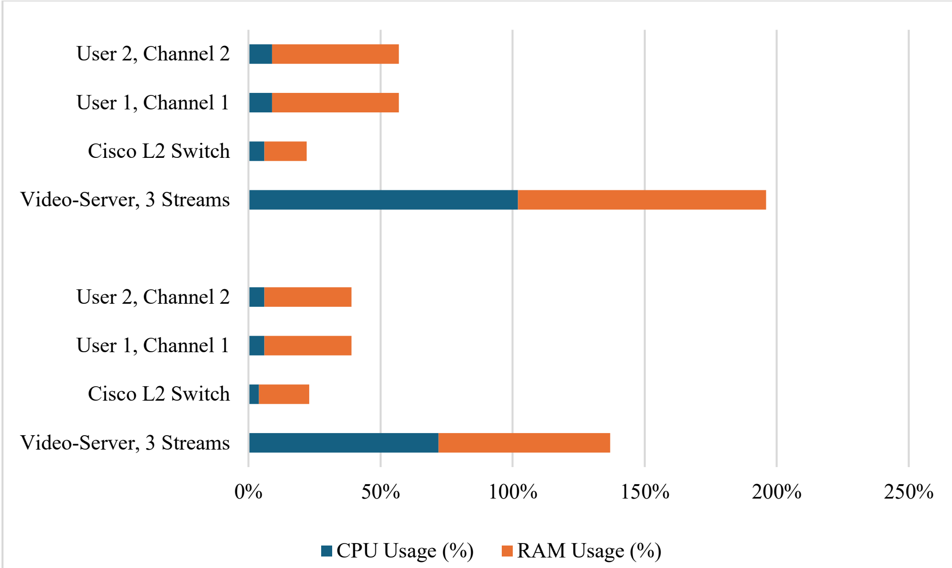

Fig. 5 Graphical comparison of devices load utilization with IGMP enabled and disabled.

4. Conclusions

By using controlled simulations in Cisco Modeling Labs, we quantified the IGMP protocol's impact. Port-level measurements revealed that enabling IGMP reduced forwarded traffic volume on active switch interfaces by nearly 79.4 %. In contrast, without IGMP, flooding in all connected interfaces led to a 4.85× increase in bandwidth usage. At the packet level, observations further confirmed IGMP's effectiveness and functionality: disabling the protocol led to a threefold increase in redundant traffic on links that are not active. An improper manner of multicast management will result in a service degradation, even at a small scale, and in serious scalability issues, in larger scales. Despite variations caused by hardware and firmware differences that may be between simulation and physical deployments, the overall results showed the consistency of experimental and theoretical expectations. By combining controlled simulation with reference physical measurements, is expected to provide both quantitative and reproducible insights.

Future work could explore larger and more complex networks, with additional switches and nodes, and may investigate how routing decisions will impact the multicast efficiency.

Ethical Considerations and Declarations

The authors declare that there is no conflicts of interest related to this work. Besides, no personal or sensitive data were collected in this article, and all results are derived from simulated network traffic and reference measurements.

References

[1] ProCurve Networking, Multicast and Routing Guide for the ProCurve 6200 yl, 5300 zl, & 3500 yl Series Switches, ProCurve Networking, 2006. [Online]. Available: https://ftp.hp.com/pub/networking/software/3500_5400_6200_MulticastRteGde-July2006-59914692.pdf

[2] Cisco Nexus 1000V Troubleshooting Guide, Release 4.2(1) SV2(2.1), Chapter 18 Multicast IGMP, Cisco Systems, Inc., 2015. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus1000/sw/4_2_1_s_v_2_2_1/troubleshooting/configuration/guide/n1000v_trouble.pdf?utm

[3] Catalyst 3560 Switch Software Configuration Guide, Cisco IOS Release 12.2(25) SEC, Cisco Systems, 2005. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3560/software/release/12-2_25_sec/configuration/guide/3560scg.pdf

[4] Consolidated Platform Configuration Guide, Cisco IOS XE 3.3SE (Catalyst 3850 Switches), Chapter Configuring SSM, Cisco Systems, 2015. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/3se/consolidated_guide/configuration_guide/b_consolidated_3850_3se_cg/b_consolidated_3850_3se_cg_chapter_010000.html

[5] ProCurve, Management and Configuration Guide, 2910al, 2009. [Online]. Available: https://data.nag.wiki/HP/2910_Switch_Management_and_Configuration_Guide.pdf

[6] Kevin Dooley, What is Multicast Networking and How Does it Work?, 2024. [Online]. Available: https://www.auvik.com/franklyit/blog/multicast-networking/

[7] FFmpeg, About FFmpeg, FFmpeg, 2026. [Online]. Available: https://www.ffmpeg.org/about.html

[8] SAMIM, Role of MPEG Transport Stream in Modern Broadcasting, SAMIM, 2025. [Online]. Available: https://www.samimgroup.com/blog/role-mpeg-ts-broadcasting/

[9] Emad Abd-Elrahman et al., "Optimization of TV Multicast Delivery," *Journal of Electronic Systems*, pp. 135-147, 2013. [Google Scholar] [Publisher Link]

[10] Shoyeb Al Mamun Shohag, and Mozammel Bin Motalab, "An Observation and Analysis of IPTV and Multicasting Traffic," *International Journal of Computer Applications*, vol. 30, no. 1, pp. 1-6, 2011. [CrossRef] [Google Scholar] [Publisher Link]

[11] Brad Cain et al., "Internet Group Management Protocol, Version 3," *RFC 3376*, 2002. [Google Scholar] [Publisher Link]

[12] CISCO DevNet, Introduction to Cisco Modeling Labs, version 2.9 x, CISCO DevNet, 2026. [Online]. Available: https://developer.cisco.com/docs/modeling-labs/2-9/introduction/#introduction-to-cisco-modeling-labs

[13] CISCO DevNet, VM Images for CML Labs, CISCO DevNet, 2026. [Online]. Available: https://developer.cisco.com/docs/modeling-labs/vm-images-for-cml-labs/#vm-images-for-cml-labs

[14] Ruiqi Ouyang et al., "A Scalable Video Multicast Scheme based on user Demand Perception and D2D Communication," *Sensors*, vol. 23, no. 17, pp. 1-19, 2023. [CrossRef] [Google Scholar] [Publisher Link]