Demo: Enhancing the Networking Performance of IPv6 IoT Devices Using Machine Learning and IVI

Mengying Zhou, Zhiyang Sun[*], Xin Wang, Yang Chen
School of Computer Science, Fudan University, China
{myzhou19,sunzy,xinw,chenyang}@fudan.edu.cn

ABSTRACT

IPv6, with its larger address space, is well-suited for IoT scenarios, allowing each device to have a unique address for efficient communication. However, its adoption in sensor networks is hindered by underdeveloped IPv6 infrastructure and compatibility issues with IPv4-based networks. To address these challenges, this paper proposes a dual-stack load balancing approach to enhance seamless and reliable connectivity between IPv6 and IPv4. By adaptively distributing traffic between IPv6 and IPv4 channels, our system alleviates performance bottlenecks and improves reliability. A real-world deployment in the campus network of Fudan University demonstrates a 1.9x increase in throughput and a 27% reduction in congestion-related failures with only CPU overhead of 2.44% and a translation delay of 0.58 ms. These results highlight the potential of intelligent load balancing and IPv4/IPv6 translation in supporting scalable and efficient sensor networks and facilitating the transition to IPv6.

CCS CONCEPTS

- **Networks** → **Network protocols**; *Sensor networks.*

KEYWORDS

Sensor Networks, IPv4/IPv6 Transition, Load Balance

ACM Reference Format:
Mengying Zhou, Zhiyang Sun[*], Xin Wang, Yang Chen. 2024. Demo: Enhancing the Networking Performance of IPv6 IoT Devices Using Machine Learning and IVI. In *ACM Conference on Embedded Networked Sensor Systems (SenSys '24), November 4–7, 2024, Hangzhou, China.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3666025.3699408

1 INTRODUCTION

Embedded sensor systems are a critical component of the Internet of Things (IoT), and their rapid growth highlights the necessity for efficient, scalable, and robust network protocols. IPv4, the prevalent protocol, faces the address exhaustion problem. IPv6, with its vastly larger address space and enhanced abilities, is widely adopted in IoT as a solution to these limitations [1, 2].

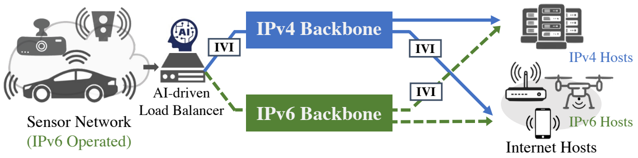

Figure 1: Framework of Our Proposal

However, the transition to IPv6 in sensor networks is hindered by two factors. First, the limited IPv6 infrastructure can lead to congestion or even single points of failure at strategic gateways. In high-traffic environments like the Internet of Vehicles, where each vehicle generates up to 300 Mbps [5] of data, a single IPv6 link lacks such capability due to insufficient infrastructure support. Second, communication compatibility is required between IPv6-based sensor networks and mainstream IPv4-based backbone networks. To accommodate the constrained resources and the large number of sensor nodes, a lightweight and scalable transition mechanism is essential.

To address these issues, we explores strategies for seamlessly integrating IPv6 and IPv4 in sensor networks. It employs dual-stack load balancing to enhance connectivity of IPv6-based sensor networks while ensuring compatibility with the IPv4 backbone network, as depicted in Figure 1. By adaptively distributing traffic across IPv6 and IPv4 routes, the load balancing mechanism mitigates performance bottlenecks and enhances communication reliability.

For traffic routed through the IPv4 network, we deploy IVI-equipped devices at gateways or points of presence (PoPs) to ensure compatibility between the two protocols when traffic is forwarded from IPv6 sensor networks to the IPv4 backbones. IVI [3] is a stateless IPv4/IPv6 translation technique with parallelization, making it ideal for handling high-concurrency traffic without excessive overhead. We develop and validate the effectiveness of load-balancing integrated IVI devices in a real-world campus network at Fudan University. The results demonstrate the feasibility of IPv4/IPv6 integration in sensor networks, achieving an improved transmission performance with minimal overhead.

2 SYSTEM DESIGN

Our proposal is equipped with an AI-driven load balancer and several IVI translators that ensures compatibility between IPv6 and IPv4.

2.1 AI-driven Load Balancer

The AI-driven load balancer dynamically distributes network traffic between IPv4 and IPv6 routes to optimize connectivity for IoT

---

[*] This work is supported by the National Natural Science Foundation of China No. 62472101. Zhiyang Sun is the Corresponding Author.

SenSys '24, November 4–7, 2024, Hangzhou, China

ACM ISBN 979-8-4007-0697-4/24/11.
https://doi.org/10.1145/3666025.3699408

Table 1: Header translation between IPv6 and IPv4

IPv6 Field	IPv4 Field
Version (0x6)	Version (0x4)
discarded	IHL
Traffic Class	Type of Service
Payload Length=Total Length - 20	Total Length
Hop Limit	TTL
Next Header	Protocol
IVI address mapping	Source Address
IVI address mapping	Destination Address
Flow Label	-
-	IHL = 5
-	Header Checksum recalculated
-	Options
-	Identification
-	Flags
-	Offset
-	Header Checksum

devices. By selecting the most suitable path for each packet, the load balancer ensures an efficient resource utilization, reduces congestion, and enhances network reliability.

The load balancer integrates a supervised machine learning model into the network layer, enabling it to make routing decisions based on real-time link information and historical feedback [6]. The model is trained using features derived from network characteristics, including latency, Time to First Byte, retransmission rates, packet size, traffic type, and historical transmission performance of IPv4 and IPv6 paths. For the machine learning model, we use XGBoost, a high-performance gradient boosting algorithm optimized for real-time decision-making with minimal overhead. We use grid search to identify the optimal model parameters for best performance. The model is trained on labeled data from [4], with each instance specifying the correct routing decision. Once trained, the model predicts the most efficient routing path for each packet, adapting to changing network conditions.

The load balancer's ability to adapt in real time is valuable in IoT environments where network conditions vary significantly. For example, the load balancer can respond to sudden increases in traffic by rerouting packets through less congested paths, preventing performance degradation. Additionally, using machine learning allows the system to learn from past network behaviors, improving decision-making capabilities over time. This proactive approach helps ensure consistent performance and reliability, which is critical for IoT applications requiring low latency and high availability.

2.2 IVI Translator

IVI [3] is a stateless IPv4/IPv6 translation technology that ensures compatibility between IPv4 and IPv6 by mapping addresses. The translation process involves:

1. **IPv4 to IPv6.** The IVI translator embeds the IPv4 address into an IPv6 address using a fixed p-bit prefix, the 32-bit IPv4 address, and a 96-p-bit zero suffix. The prefix's last byte is set to all ones to indicate a translated address.
2. **IPv6 to IPv4.** The translator extracts the IPv4 portion from the IPv6 address and converts it back to a standard IPv4 address. IVI adheres to preset rules listed in Table 1 for translating network layer headers, discarding or defaulting fields as necessary due to differences in header structures.

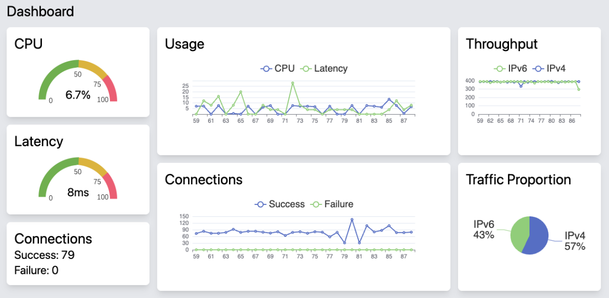

Figure 2: Dashboard of Our Proposal

Owing to its stateless design, IVI efficiently handles large volumes of concurrent data with minimal memory overhead and processing delay, ensuring scalability. IVI also can be easily implemented on commercial network cards, utilizing hardware optimizations to further decrease translation latency.

3 PRELIMINARY DEPLOYMENT

We conduct a preliminary evaluation of our proposal in a real-world environment. The network topology is shown in Figure 1. It involves a sensor network at the campus of Fudan University, a CDN service in Shanghai Alibaba Cloud datacenter, and a home-based wireless sensor network. These network communicate via both IPv4 and IPv6 links. The IPv4 link is provided by the first-generation China Education and Research Network (CERNET), while the IPv6 link is supported by a separate pure IPv6 CERNET2. Both links offer 1 Gbps bandwidth, 50 ms delay, and a 0.1% packet loss rate.

We also develop a dashboard to monitor the load balancer, as shown in Figure 2. Compared with the baseline without load balancing, where IPv6 traffic is transmitted solely over the IPv6 link, our adaptive method increases throughput from 393 Mbps to 764 Mbps, achieving 1.9x the baseline performance by offloading traffic to the IPv4 link. Additionally, our intelligent distribution reduces congestion-related request failures by 27%. We also monitor the overhead of IVI devices, including CPU usage and translation delays. Running on a server with a Quad Core Xeon CPU and 12 GB RAM, the IVI translator has an average CPU overhead of 2.44% and a minimal delay of 0.58 ms, which is negligible in the overall request time. Our current software-based prototype can be further optimized by implementing it using dedicated hardware.

REFERENCES

[1] Anwar Hussain et al. 2021. A Resource-Efficient Hybrid Proxy Mobile IPv6 Extension for Next-Generation IoT Networks. *IEEE Internet of Things Journal* 10, 3 (2021), 2095–2103.

[2] Antonio J Jara et al. 2013. The Internet of Everything through IPv6: An Analysis of Challenges, Solutions and Opportunities. *J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl.* 4, 3 (2013), 97–118.

[3] Xing Li et al. 2021. The China Education and Research Network (CERNET) IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition. RFC 6219. Available: https://datatracker.ietf.org/doc/html/rfc6219. Accessed: 2024-09-01.

[4] Jia Zhang et al. 2023. Reducing Mobile Web Latency Through Adaptively Selecting Transport Protocol. *IEEE/ACM Transactions on Networking* 31, 5 (2023), 2162–2177.

[5] Xumiao Zhang et al. 2021. EMP: Edge-assisted Multi-vehicle Perception Xumiao. In *Proc. of MobiCom.*

[6] Mengying Zhou et al. 2022. FlexHTTP: An Intelligent and Scalable HTTP Version Selection System. In *Proc. of EuroMLSys.*