



QUICPro: Integrating Deep Reinforcement Learning to Defend against QUIC Handshake Flooding Attacks

Y A Joarder* & Carol Fung

Concordia Institute for Information Systems Engineering
Concordia Institute for Information Systems Engineering (CIISE),
Concordia University
Montreal, Quebec, Canada


Outline

- ✓ Introduction of QUICPro
- ✓ QUIC Overview
- ✓ Previous Research Works
- ✓ Motivations
- ✓ Problem Statements
- ✓ QUICPro Framework Details
- ✓ Planned Implementation Details
- ✓ Expected Outcomes and Benefits
- ✓ Conclusion & Future Work

Introduction of QUICPro



It utilizes Deep Reinforcement Learning (DRL) with the Proximal Policy Optimization (PPO) algorithm for dynamic security optimization against handshake flooding attacks



It employs real-time rate limiting, connection prioritization, and traffic shaping





It implements the Isolation Forest for anomaly detection and Support Vector Machine (SVM) for pattern recognition



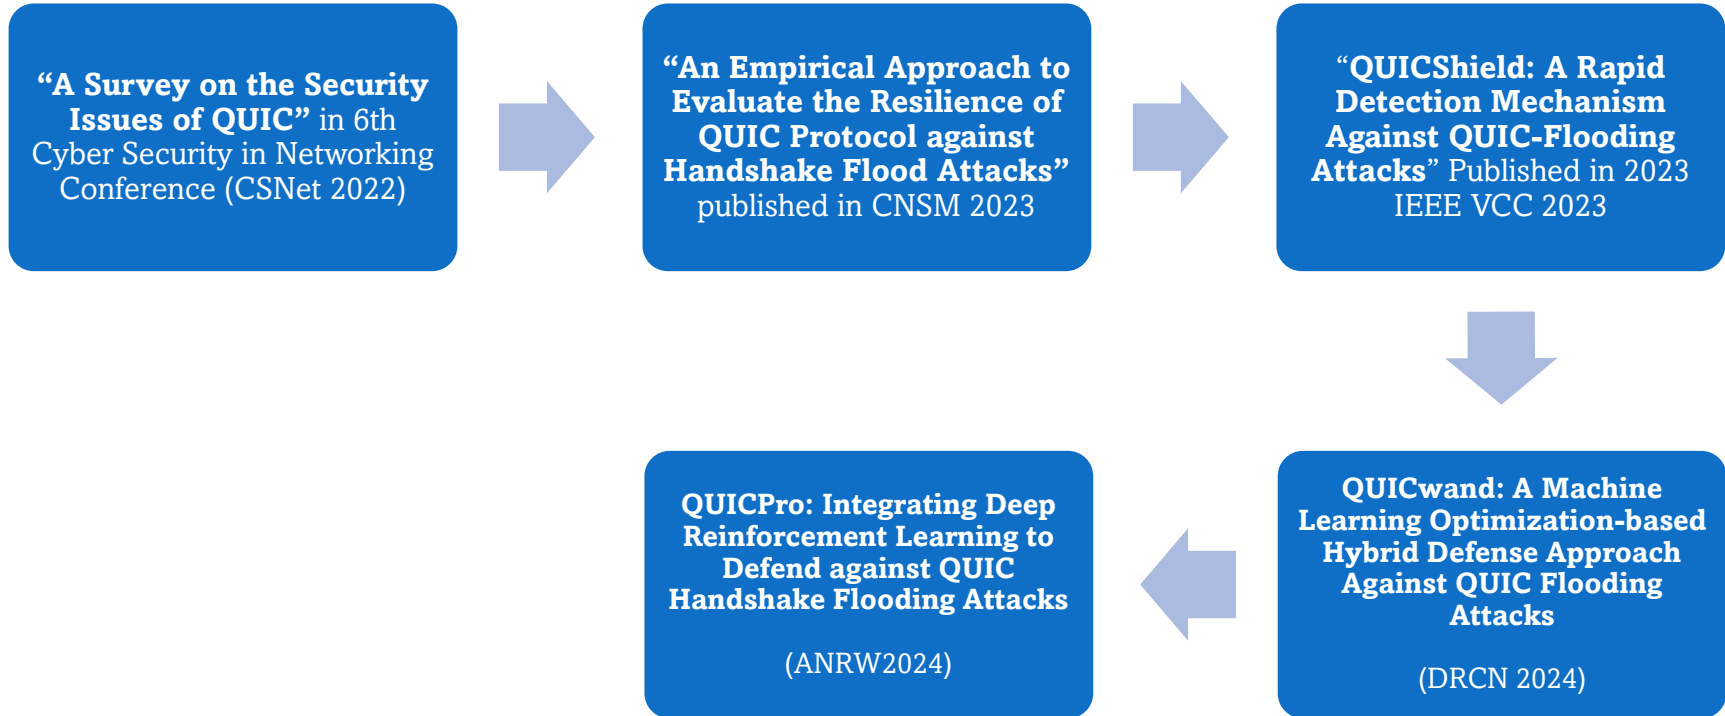
It features a feedback loop for ongoing improvement of security policies

QUIC Overview

- A UDP-Based Multiplexed and Secure Transport Protocol
- It Provides flow-controlled streams for structured and efficient communication
- It gives Low-latency connection establishment
- It supports network path migration
- It offers 0-RTT and connection migration feature
- It enhances security with measures for confidentiality, integrity, and availability



Previous Research Works



Motivations

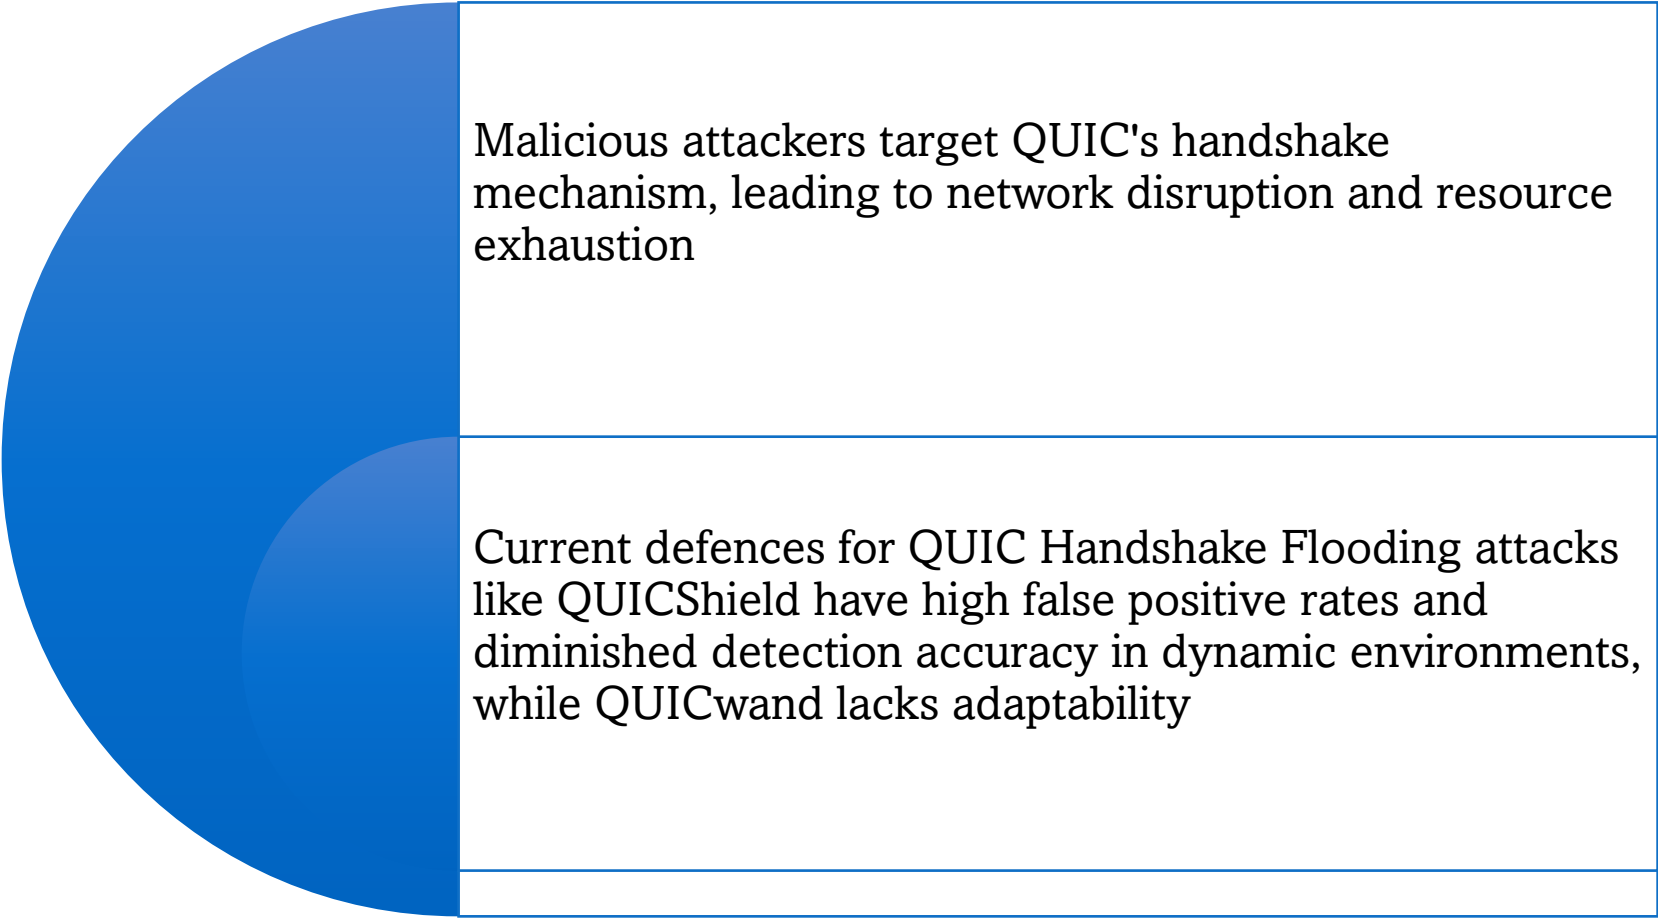
QUIC's handshake process is vulnerable to flooding attacks with a 4.6x CPU amplification factor

Traditional TCP's SYN-flooding detection mechanisms are ineffective due to QUIC's distinct handshake design

QUICShield, a prior defense solution, has a higher false positive rate, potentially blocking legitimate traffic

QUICwand lacks adaptability to rapidly evolving network conditions and sophisticated attacks

Problem Statements



Malicious attackers target QUIC's handshake mechanism, leading to network disruption and resource exhaustion

Current defences for QUIC Handshake Flooding attacks like QUICShield have high false positive rates and diminished detection accuracy in dynamic environments, while QUICwand lacks adaptability


QUICPro Framework



Core Technology: Utilizes Deep Reinforcement Learning (DRL) with Proximal Policy Optimization (PPO) to dynamically optimize security measures against handshake flooding attacks



Functionality: Detects and mitigates QUIC flooding attacks by continuously monitoring network traffic and adjusting defense mechanisms in real-time

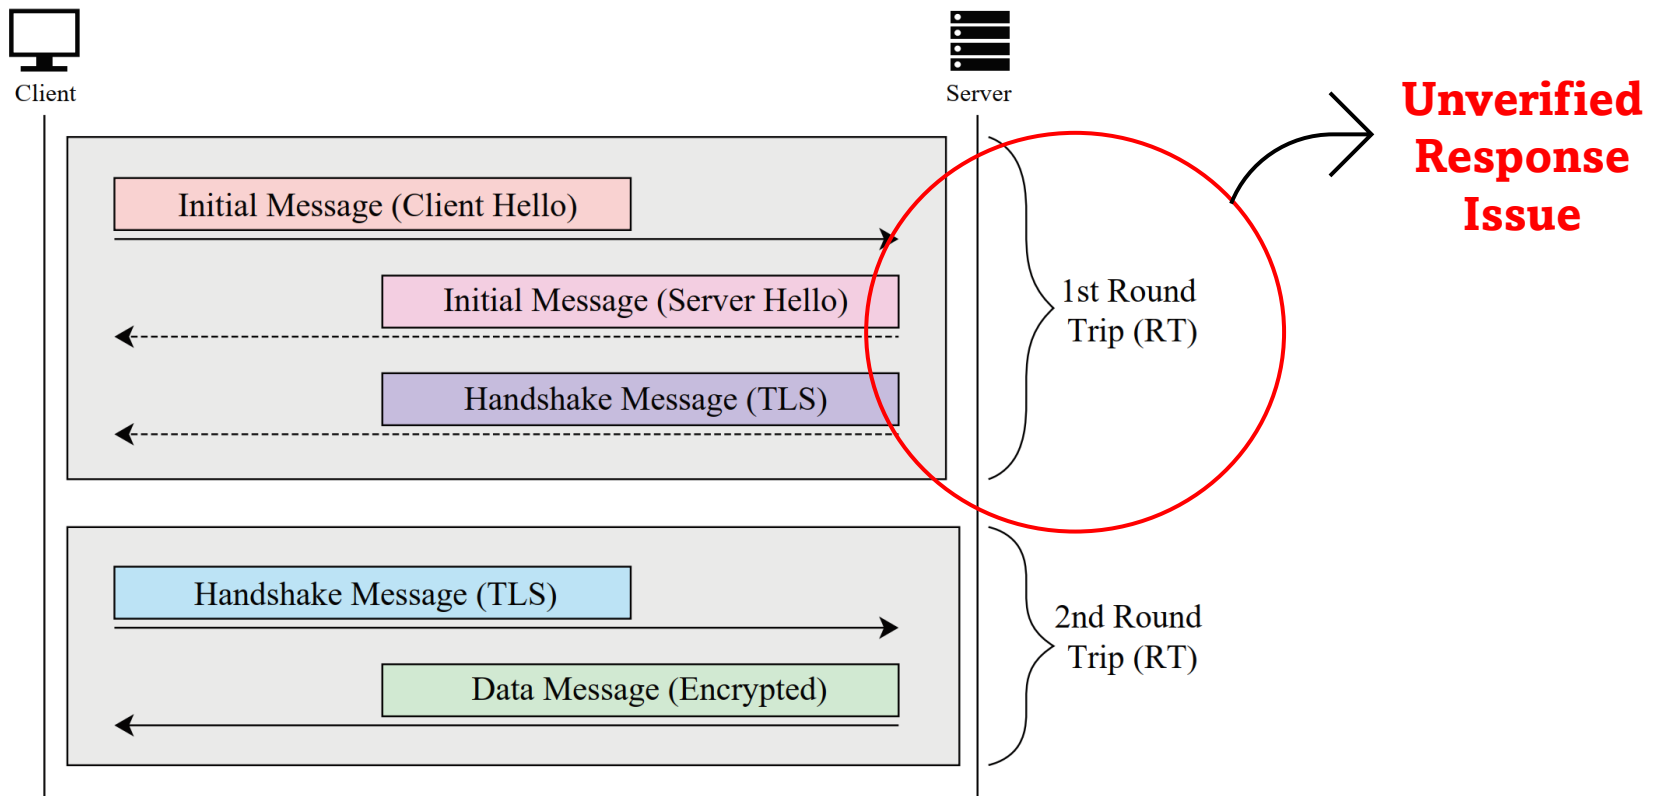


Dynamic Optimization: Leverages machine learning to adapt detection and mitigation strategies based on current network conditions and emerging threats

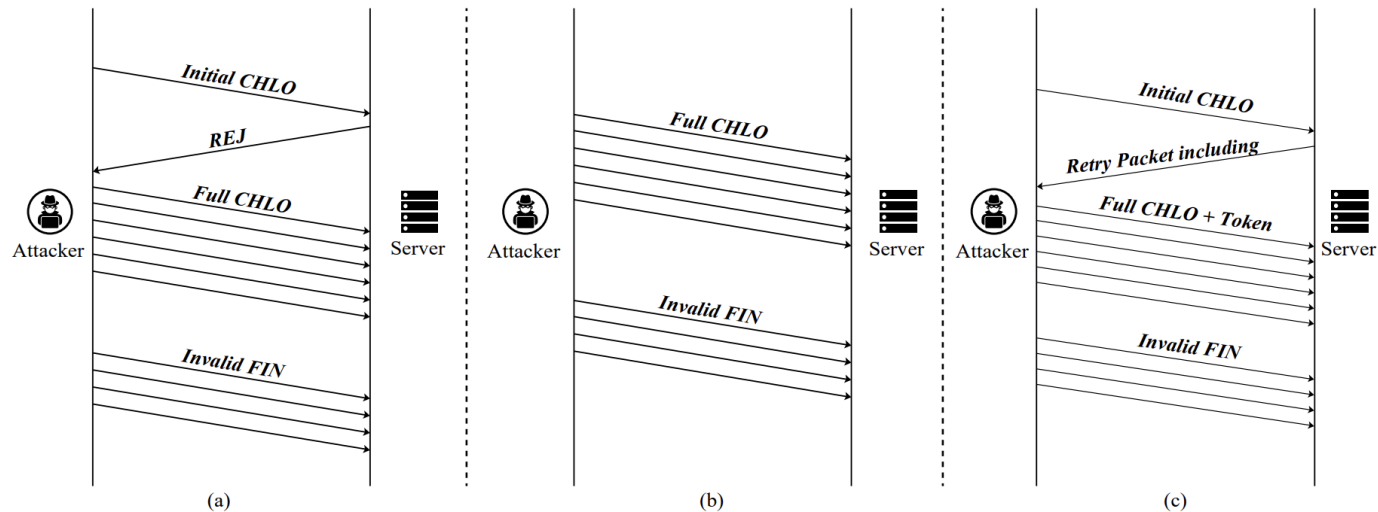


Improved Accuracy: Reduces false positives and increases detection accuracy compared to previous methods, such as QUICShield and QUICwand

QUIC Simplified Handshaking



QUIC Handshake Flooding Attacks



QUIC Flooding Attack: (a) Version negotiation; (b) Without address validation; (c) With address validation, showcasing security layers

Core Components of QUICPro

1

- Deep Reinforcement Learning (DRL) Agents

2

- Network Traffic Monitoring Module

3

- Adaptive Defense Mechanisms

Deep Reinforcement Learning (DRL) Agents



Use Proximal Policy Optimization (PPO)
for learning and optimizing security
policies based on environmental
feedback




Neural networks with Convolutional
Neural Networks (CNNs) for feature
extraction followed by fully connected
layers for decision-making

Network Traffic Monitoring Module



Continuously monitors incoming traffic to QUIC servers



Uses machine learning algorithms like Isolation Forest for anomaly detection and Support Vector Machines (SVM) for pattern recognition

Adaptive Defense Mechanisms



Use Employs dynamic rate limiting, connection prioritization, and traffic shaping techniques



Adjusts in real-time based on feedback from DRL agents and the traffic monitoring module

Planned Implementation Details of QUICPro

Integration with Existing QUIC Implementations

- Incorporates with popular QUIC libraries (e.g., *aiquic*, *quicly*, *mvfst*) for seamless operation

Deep Reinforcement Learning (DRL) Setup

- **Algorithm:** Uses Proximal Policy Optimization (PPO)
- **Neural Networks:** Employs CNNs for feature extraction, followed by fully connected layers for decision-making

Network Traffic Monitoring

- **Continuous Monitoring:** Analyzes packet headers and payload data
- **Machine Learning:** Utilizes Isolation Forest for anomaly detection and SVM for pattern recognition

Expected Outcomes and Benefits

High Detection Accuracy

- Achieves high detection accuracy against handshake flooding attacks by continuously adapting to new patterns

Reduced False Positives

- Minimizes legitimate traffic blocks by reducing false positives

Dynamic Adaptation

- Adjusts security parameters in real-time for effective mitigation of attacks

Enhanced Resilience

- Improves the resilience of QUIC servers, ensuring robust and uninterrupted protection against evolving handshake flooding attacks

Conclusion & Future Work



Effective Detection and Mitigation



Low False Positive Rates



Enhanced Resilience



- Implementation
- Testing Against Diverse Attacks
- Developing a Comprehensive QUIC Attacks Dataset

Thank You so much!



References

- [1] Y A Joarder and Carol Fung. 2024. QUICwand: A Machine Learning Optimization-Based Hybrid Defense Approach Against QUIC Flooding Attacks. In 2024 20th International Conference on the Design of Reliable Communication Networks (DRCN). IEEE, Montreal, QC, Canada, 92–99. <https://doi.org/10.1109/DRCN60692.2024.10539170>
- [2] B. Teyssier, Y A Joarder, and C. Fung, “QUICShield: A Rapid Detection Mechanism Against QUIC-Flooding Attacks,” in 2023 IEEE Virtual Conference on Communications (VCC), pp. 43–48, Nov. 2023.
- [3] B. Teyssier, Y. A. Joarder, and C. Fung, “An Empirical Approach to Evaluate the Resilience of QUIC Protocol Against Handshake Flood Attacks,” in 2023 19th International Conference on Network and Service Management (CNSM), pp. 1–9, Oct. 2023. ISSN: 2165-963X.
- [4] Y A Joarder and C. Fung, “A Survey on the Security Issues of QUIC,” in 2022 6th Cyber Security in Networking Conference (CSNet), pp. 1–8, Oct. 2022. ISSN: 2768-0029.
- [5] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. 2017. Deep Reinforcement Learning: A Brief Survey. IEEE Signal Processing Magazine 34, 6 (Nov. 2017), 26–38. <https://doi.org/10.1109/MSP.2017.2743240>

*****The remaining references are in the main QUICPro Paper.***