Intelligent Traffic Steering xApp for Load Balancing in O-RAN: A Deep Reinforcement Learning Approach

Myoungjin Oh, Sangmin Kim*, Inkyu Lee*

School of Computer Science and Engineering, Chungang Univ.,
*School of Electrical Engineering, Korea Univ.

omjin7n@cau.ac.kr, *smgeem@korea.ac.kr, *inkyu@korea.ac.kr

O-RAN 환경에서 부하 분산을 위한 심층 강화학습 기반의 지능형 트래픽 스티어링 xApp

오명진, 김상민*, 이인규*
중앙대학교 소프트웨어학부, *고려대학교 전자전기공학부

Abstract

Towards 6G O-RAN, conventional RSRP-based handover mechanisms present a significant limitation, failing to prevent network-wide outages under dynamic traffic surges. To address this, we propose an intelligent traffic steering xApp for load balancing, which can be deployed on the Near-Real Time RAN Intelligent Controller (Near-RT RIC). Specifically, we employ a Deep Reinforcement Learning (DRL) approach utilizing the Proximal Policy Optimization (PPO) algorithm. Simulation results demonstrate that the proposed xApp achieves strict load fairness among base stations, effectively preventing single-node overloads and ensuring robust network stability.

Ⅰ. Introduction

The evolution of Beyond 5G technologies requires mobile networks to handle unprecedented dynamic traffic conditions without service outages. In this context, the Open Radio Access Network (O-RAN) architecture enables data-driven control via RAN Intelligent Controller (RIC). Specifically, the Near-Real Time RIC (Near-RT RIC) hosts third-party applications, known as xApps, to execute programmable network control logic.

However, conventional RSRP-based handover mechanisms defined in 3GPP specifications [1] present a significant limitation as they are cell-centric, often failing to prevent network-wide outages under dynamic traffic surges.

To address this limitation, we propose a system-centric Traffic Steering (TS) xApp utilizing the Proximal Policy Optimization (PPO) algorithm. Unlike legacy methods, our approach dynamically balances traffic loads among base stations by optimizing fairness-aware reward function to minimize load imbalance. We formulate the load balancing problem as a Markov Decision Process (MDP) and design a lightweight PPO agent. Simulation results demonstrate that the proposed xApp effectively prevents service outages and ensures robust network stability.

Ⅱ. System Model and Problem Formulation

We consider a downlink network environment consisting of multiple base stations (BSs) and mobile users, managed by the O-RAN architecture. The objective is to optimize traffic steering policies to prevent base station overloading by dynamically balancing traffic among adjacent BSs.

1. O-RAN Architecture

As illustrated in Fig. 1, the proposed Traffic Steering (TS) xApp operates on the Near-RT RIC to optimize network load distribution. The xApp interacts with E2 Nodes via the standard E2 interface, utilizing the E2 Service Model for Key Performance Measurement (E2SM-KPM) [2] to collect real-time UE RSRP and cell load statistics. Based on these observations, the xApp executes handover control decisions to redistribute UEs from traffic hotspots to adjacent cells, ensuring that load balancing is performed only when signal coverage constraints are met.

2. MDP Formulation

We formulate the traffic steering problem as an MDP defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \gamma \rangle$. First, the state $s_t = [L_{1,t}, L_{2,t}, L_{3,t}]$ represents the normalized traffic demand of each BS at time t, collected via E2SM-KPM. Second, the action a_t determines the handover (HO) direction. The discrete action space $\mathcal{A}$ consists of a

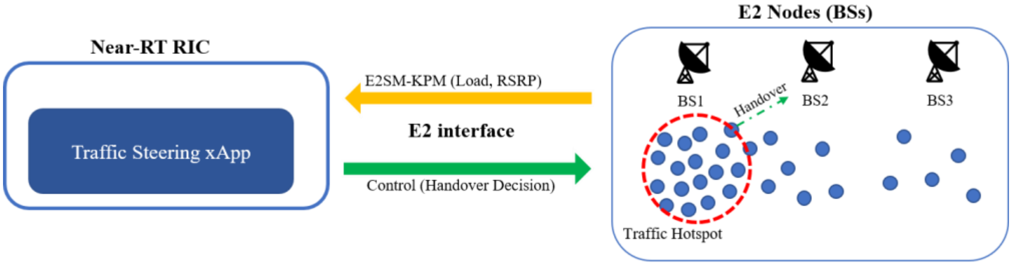

Fig. 1. System architecture of the proposed traffic steering xApp in O-RAN. The xApp collects network stats via E2SM-KPM and executes load balancing controls.

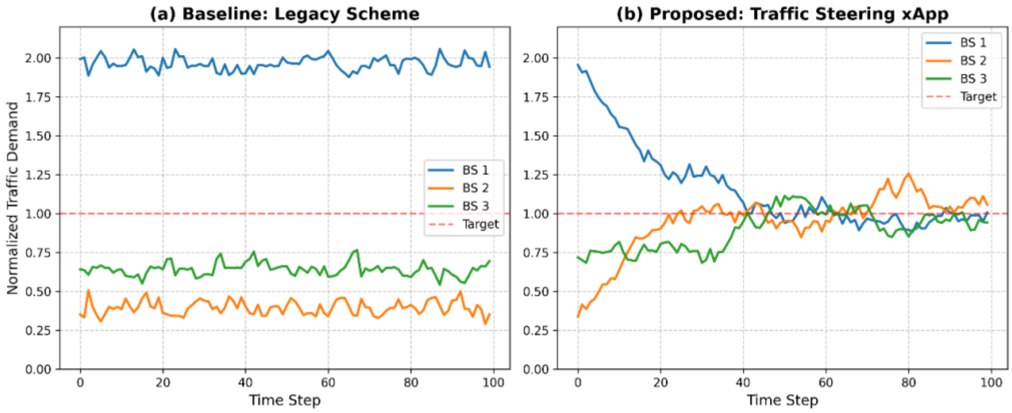

Fig. 2. Comparison of load balancing performance in a traffic hotspot scenario: (a) Baseline (Max-RSRP) scheme; (b) Proposed Traffic Steering xApp.

"No-Op" and directional handovers between adjacent BSs (i.e., $a_t \in \{No - Op\} \cup \{HO_{i \to j}\}$). Finally, to minimize load imbalance while strictly preventing service outages, we define the reward function $\mathcal{R}$ at time t as:

$$r_t = -\alpha\sigma(s_t) - \beta \sum_{i=1}^{3} I(L_{i,t} > 1.0)$$

where $\sigma(s_t)$ is the standard deviation of loads, and $I(\cdot)$ is an indicator function that applies α penalty β when any BS is overloaded (> 1.0).

Ⅲ. Simulation Results

We evaluate the scheme in an O-RAN environment with 3 BSs and 60 UEs, following the 3GPP Urban Micro (UMi) model [3]. To simulate congestion, 70% of UEs are concentrated at BS 1, while the rest are uniformly distributed. The proposed method is compared against a standard Max-RSRP baseline (3GPP Event A3 [1]). Fig. 2 presents the load balancing performance. In the baseline case (Fig. 2(a)), BS 1 remains heavily overloaded (Normalized Demand ≈ 2.0) because UEs do not detach from the strongest cell. This confirms the inefficiency of load-agnostic handovers. On the contrary, the proposed xApp (Fig. 2(b)) successfully redistributes traffic. The agent redirects UEs from the hotspot to adjacent cells (BS 2, BS 3) that meet the minimum RSRP threshold (−110 dBm). As a result, the load across all BSs converges to the target level (≈ 1.0), ensuring service continuity and efficient resource usage.

Ⅳ. Conclusion

We proposed a traffic steering xApp utilizing the O-RAN E2 interface to resolve load imbalance in 5G networks. By formulating the problem as an MDP and applying the PPO algorithm, our approach effectively distributes user traffic from congested cells to underutilized neighbors while maintaining signal quality constraints. Simulation results demonstrated that the proposed scheme prevents the "sticky client" problem inherent in legacy Max-RSRP handovers, achieving load convergence across base stations. In future work, we plan to extend this framework to dynamic mobility scenarios and investigate the impact of diverse traffic patterns on reinforcement learning performance.

ACKNOWLEDGMENT

This work was supported in part by the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT (MSIT), Korea Government under Grant RS-2022-NR070834.

REFERENCES

[1] 3rd Generation Partnership Project (3GPP), "NR; Radio Resource Control (RRC); Protocol specification," TS 38.331, V 18.7.0, Oct. 2025.

[2] O-RAN Alliance, "O-RAN Near-Real-time RAN Intelligent Controller E2 Service Model (E2SM), KPM 1.0," O-RAN.WG3.E2SM-KPM-v01.00, Feb. 2020.

[3] 3rd Generation Partnership Project (3GPP), "Study on channel model for frequencies from 0.5 to 100 GHz," TR 38.901.