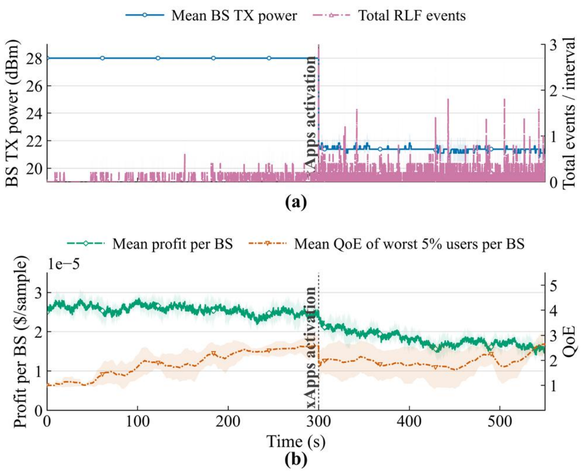
Exploring KPI Trade-offs in O-RAN Conflict Resolution with Deep Reinforcement Learning

Michele Simeone[a,b], Francesco Spinelli[c], Antonio Bazco[a],
Enrico Boffetti[d], Esteban Municio Hernandez[c], Vincenzo Mancuso[a,e],
[a] IMDEA Networks, Madrid, Spain; [b] Universidad Carlos III de Madrid, Spain; [c] Fundació i2CAT, Barcelona, Spain;
[d] Politecnico di Bari, Italy; [e] Università degli Studi di Palermo, Palermo, Italy.

Introduction

O-RAN makes the RAN open and programmable through standardized interfaces. The **Near-RT RIC** hosts **xApps** that use RAN measurements to optimize the network, and several xApps can operate **independently** on the same RAN.



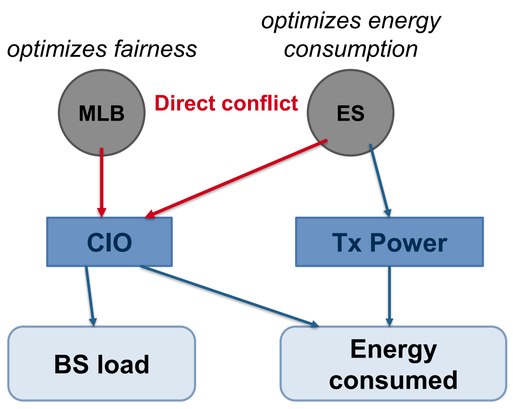
- Energy Saving xApp: **TX power decreases**
- Mobility xApp: **handover behaviour changes**
- After activation (dotted line) RLF events grow and profit per BS drops

Locally good decisions
≠
globally good configuration

Can a learning-based mediator find better compromises between independently optimized xApps?

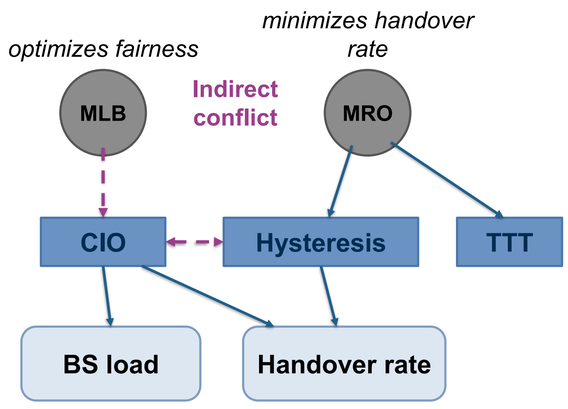
Direct vs. Indirect Conflicts

Direct conflict



Same parameter, competing intents

Indirect conflict



Different parameters, coupled effects

Detecting who conflicts is only the first step: resolving a conflict requires understanding its consequences on the whole network.

Problem Formulation

The mediator minimizes a discounted weighted sum of **QoE cost** (mapped from users' MOS) and **network operational cost**:

$$\min_{\{\boldsymbol{\delta_m}(t)\}_{t\geq 0}} \sum_{t=0}^{\infty} \gamma^t \mathbb{E}\left[w_{\mathrm{QoE}} L_{\mathrm{QoE}}(\boldsymbol{\theta}_t, \boldsymbol{\sigma}_t) + w_{\mathrm{cost}} C_{\mathrm{op}}(\boldsymbol{\theta}_t, \boldsymbol{\sigma}_t)\right]$$

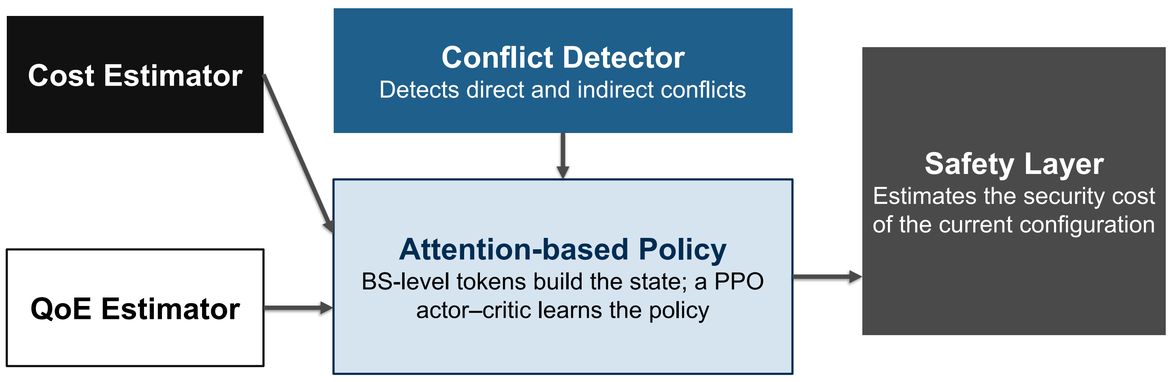
Unknown dynamics, long-term effects and a non-convex, time-varying control space make direct optimization hard.
We cast it as a **Markov Decision Process**:

- **State:** $s_t = (\theta_{t-1}, \Delta_t, m_t, F_t)$: previous configuration, xApp updates, conflict mask, BS context
- **Action:** $\delta_m(t)$, new values for the conflicting parameters only
- **Reward:** negative weighted cost minus a boundary-risk penalty g_t

$$r_t = -(w_{\mathrm{QoE}} L_{\mathrm{QoE}}(\theta_t, \sigma_t) + w_{\mathrm{cost}} C_{\mathrm{op}}(\theta_t, \sigma_t)) - g_t$$

SOCRATEX Architecture

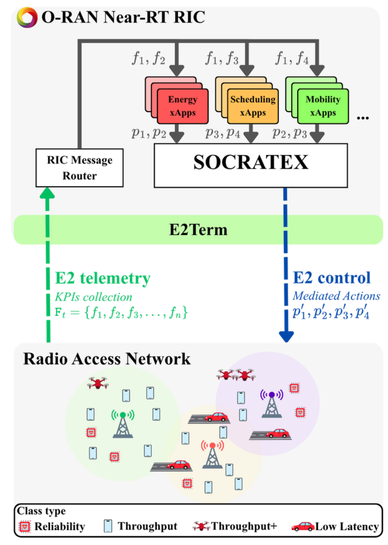
xApps propose their updates; **SOCRATEX** mediates the conflicting ones, weighing their joint effect on QoE and cost.



Simulation Setting

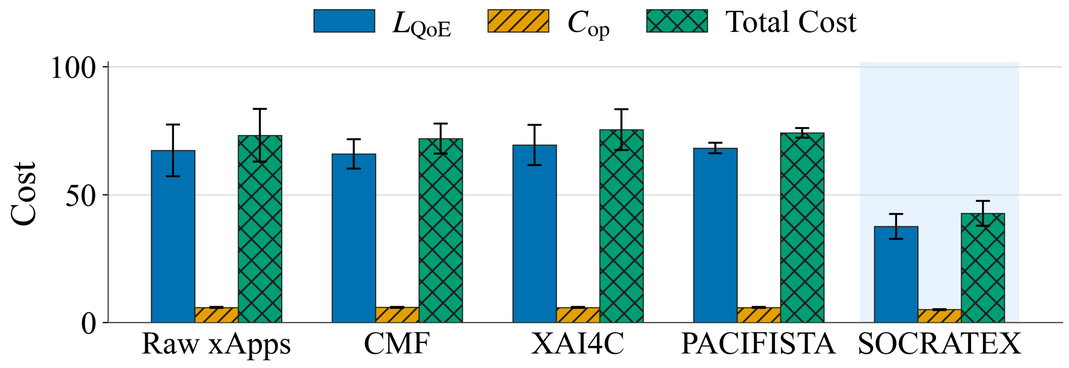
Integrated multi-xApp scenario in our own O-RAN simulator.

Topology	3 BSs, radius 1125 m
Users	225 (training), 450 (stress test)
xApps	MRO · MLB · Energy Saving · LL Sched. · TP+ Sched.
Controlled params	CIO · Hyst · TTT · TX power · ρLL · ρTP+
Mobility	80% pedestrian / 20% vehicular
Service classes	LL · TP · REL · TP+
Weights ω	0.5 QoE / 0.5 operational cost



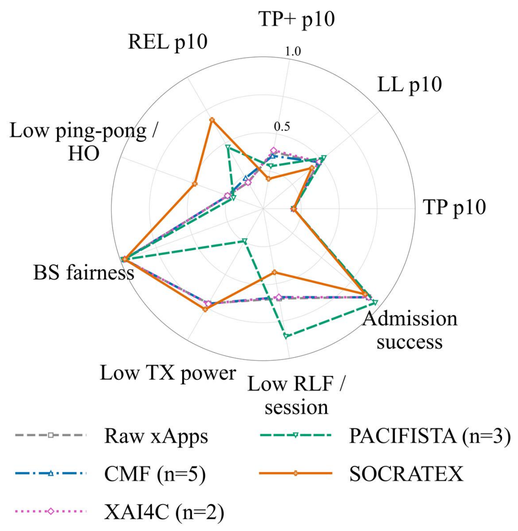
Results

SOCRATEX reduces total cost by 40%



- Stress test with **450 users** (trained with 225); lower is better.
- Main gain: **lower bad-tail QoE loss**, while network operational cost remains comparable.

Exploring the trade-off



A balanced network operating point

- No single KPI is optimized in isolation
- SOCRATEX balances QoE, fairness, mobility stability and energy efficiency
- Global gains may require controlled trade-offs on individual KPIs

Conclusion

- Independent xApps need **coordinated mediation**: their decisions interact through the same network.
- **SOCRATEX brings QoE into the control loop**, weighing it against operational cost.
- QoE-aware coordination improves the **global operating point**.

Future directions: ablation study, Non-RT RIC integration and implicit conflicts, hardware implementation.

Reference

[1] P. Brach del Prever et al., PACIFISTA: Conflict Evaluation and Management in Open RAN, IEEE TMC, 2025.
[2] N. Varshney et al., XAI4C: XAI for Conflict Detection and Mitigation in O-RAN Near-RT RIC, IEEE VT Magazine, 2025.
[3] J. Schulman et al., Proximal Policy Optimization Algorithms, 2017, 1707.06347.

Acknowledgement

Funded by the European Union through the FINALITY grant (101168816). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the granting authority can be held responsible for them.

Funded by the European Union

© Institute IMDEA Networks