Machine Learning Based Near-Real-Time Cell Outage Compensation

amdocs | make it amazing

Table of contents

Introduction

This paper introduces a comprehensive method for addressing cell outages in mobile networks with near-immediate compensation. By adjusting neighboring cells in near-real-time, the aim is to maintain network coverage and capacity in the face of sudden outages.

Cell outages significantly impact the Quality of Service (QoS), making it essential for network operators to respond quickly. The complexity of mobile networks and the interdependence of cells, due to mutual interference, make traditional Coverage and Capacity Optimization (CCO) methods inadequate for real-time outage compensation. This underscores the necessity of proactive planning for various outage and traffic scenarios, enabling an immediate and optimal response when an outage occurs.

Achieving this rapid response involves applying reinforcement learning (RL) to a digital twin of the network cell cluster. A sophisticated network simulation framework allows for the exploration of a wide range of realistic scenarios and to train agents to make optimal decisions under almost any condition. These agents are then implemented in the actual network to manage real-time issues.

As part of our evaluation, we investigated several RL strategies and algorithms, with a focus on the most significant ones and their distinctions. This paper explores the use of single-agent (SA) versus multi-agent (MA) systems, and the application of Deep Q-Network (DQN) and proximal policy optimization (PPO) algorithms. These agents are designed to either adjust a single cell or coordinate actions across multiple cells, based on generic or cell-specific observations.

The effectiveness of these strategies is assessed by analyzing Key Performance Indicators (KPI), such as coverage, capacity, and interference reduction. The impact of the agents' actions on the network is evaluated by observing changes in these metrics, particularly through adjustments in electrical tilt (e-tilt) in response to simulated cell outage scenarios.

Terminology

Uncontrolled cell outages

These outages are unexpected and occur due to various reasons, often related to hardware or software issues within the network's many components. Such outages lead to a drop in capacity within the affected cluster. In rural areas, a single cell outage can cause significant network coverage problems, directly affecting QoS. Anticipating these outages is challenging, and it is difficult to predict their duration with precision.

Controlled cell outages

With advancements in network automation and the increased flexibility offered by software-defined networks, such as the Open RAN (O-RAN) architecture illustrated in Figure 1 below, networks are evolving into more complex systems. The integration of various near-real-time xApps and non-real-time rApps from different vendors allows for multiple optimization goals to be pursued simultaneously. This complexity means networks must be agile in responding to and optimizing different network functions.

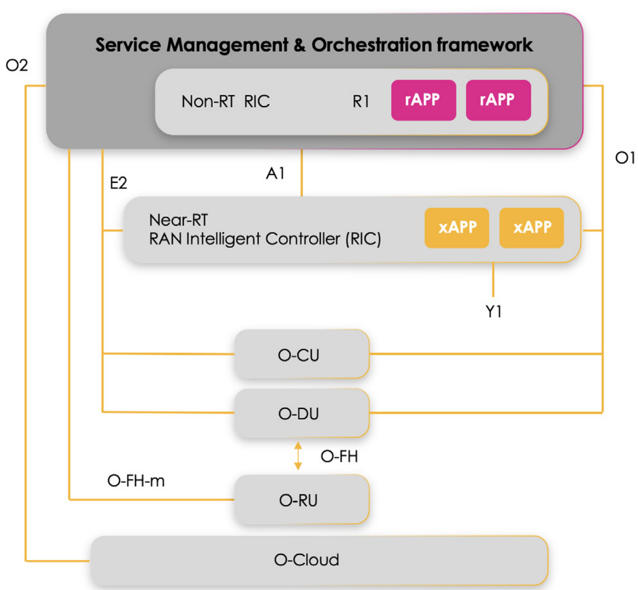

Figure 1: O-RAN Architecture Overview

One such function is the Energy Saving Management (ESM), which reduces capacity or turns off small cells not in use, particularly during low traffic periods like at night. ESM is becoming a key tool for network operators to save costs and reduce energy waste.

Planned network maintenance, sometimes conducted during peak hours, can also lead to controlled outages. These outages can affect large clusters and significantly impact network performance.

Cell outage compensation

Compensating for the negative effects of uncontrolled cell outages is critical to maintaining coverage, capacity, and most importantly, QoS. Operators need to ensure a seamless experience for their customers, but often they lack the tools for an automated and near-real-time response. Amdocs Cell Outage Compensation (COC) solution addresses this need by facilitating network self-repair through the automatic adjustment of the Remote Electrical Tilt (RET) of neighboring cells. Such adjustments are instrumental in mitigating coverage gaps and enhancing capacity in response to uncontrolled cell outages.

Amdocs Cell Outage Compensation ORAN App

The Amdocs Cell Outage Compensation xApp was tested within the Accelerating RAN Intelligence in 5G (ARI-5G) project, initiated in August 2022 and spanning 18 months. The ARI-5G Consortium is dedicated to exploring foundational use cases such as energy efficiency, massive MIMO, coverage and interference mitigation. This collaborative effort demonstrates how a united ecosystem can accelerate innovation, introduce new Open RAN products to mobile operators, and shorten development cycles.

Led by the Telecom Infra Project (TIP), the ARI-5G Consortium includes key players such as Accelleran, Amdocs, AttoCore, BT, and VIAVI Solutions. It's a significant initiative supported by the UK government's Department for Digital, Culture, Media and Sport (now DIST) to expedite the deployment of Open RAN products in UK mobile networks.

The Cell Outage Compensation xApp employs machine learning (ML) to identify the most effective strategies for cell outage compensation in near-real-time. It integrates seamlessly into the O-RAN architecture (ref. 1) and the ML workflow (ref. 2), reapplying 3GPP Self-Organizing Networks (SON) use cases like Cell Outage Detection and Compensation within the O-RAN framework, as discussed by the O-RAN alliance in the "Integrated SON Function within the O-RAN framework" use case (ref. 3).

In line with the O-RAN alliance ML workflow architecture (ref. 2), the functionality is divided between ML Training and ML Inference Host apps. Consistent with the initial discussion, ML Inference is implemented in the O-RAN near-real-time control loop as an xApp.

Both components are packaged in Docker containers and can be connected to commercial or open-source non-real-time and/or near-real-time RIC platforms via their open RIC App SDK and/or the R1 interface, depending on the ML workflow version being implemented. RAN telemetry data (CM, PM) are accessed through the Open RAN standardized O1 interface. ML actions from the ML Inference Host are conveyed through the O1 or the O-FH-m interface to the network's actors.

The initial version created for the ARI-5G project aligns with Scenario 1.4 of the O-RAN alliance AI/ML deployment scenarios, featuring both ML Training and ML Inference Host components as an xApp. These components are then connected to the project's near-real-time RIC (see Figure 2). In the ARI-5G project, a network simulator serves as a digital twin for ML training and reinforcement learning (see Figure 3).

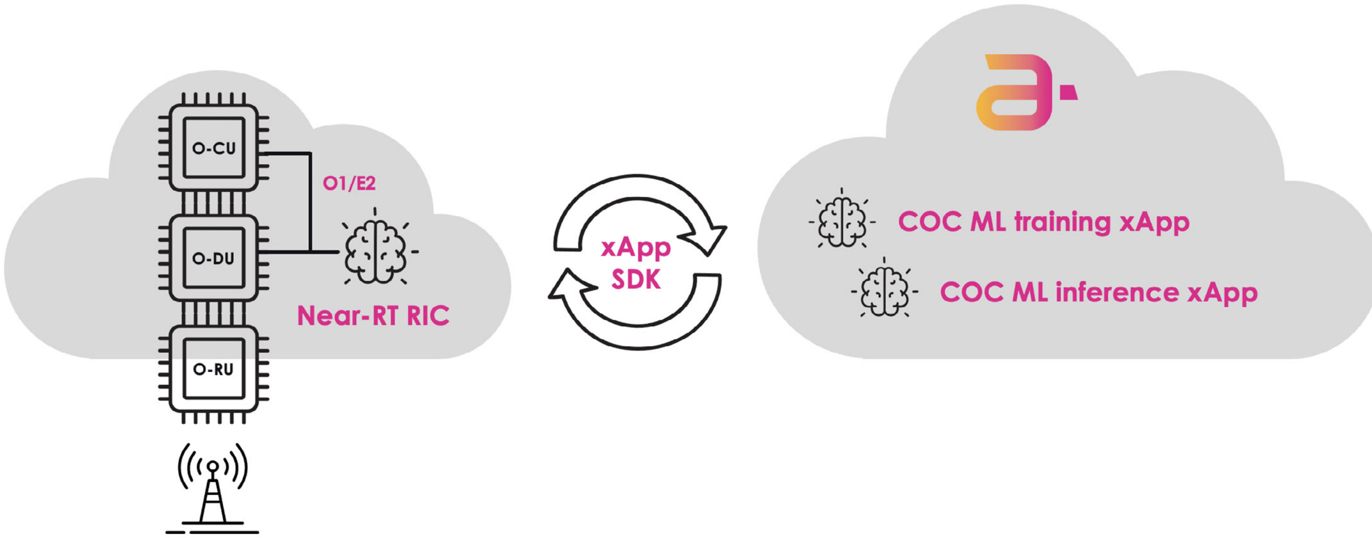

Figure 2: Amdocs Cell Outage Compensation implemented following O-RAN ML deployment Scenario 1.4 (ARI-5G version)

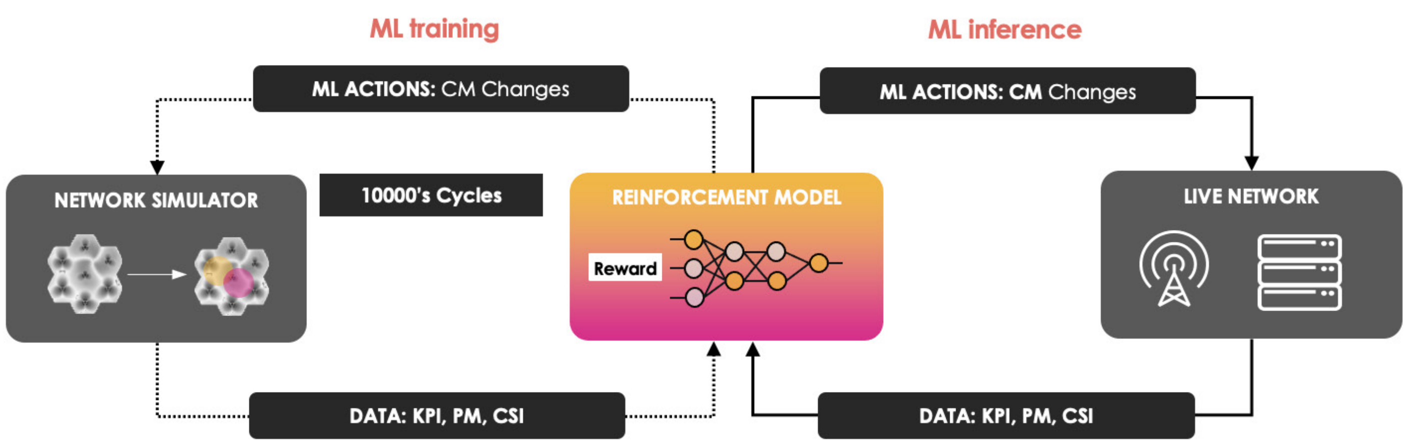

Figure 3: ML training supported by a network simulator as a digital network twin

An alternative approach for the Amdocs Cell Outage Compensation implementation might align with Scenario 1.2 outlined by the O-RAN alliance AI/ML deployment scenarios. This approach would transition the ML training process to the non-real-time control loop, utilizing the A1 interface to deploy the ML model into the near-real-time Inference Host xApp (refer to Figure 4).

Additionally, based on specific network operator requirements, the O-RAN ML deployment Scenario 1.3 could be an option. This scenario involves conducting ML training offline and managing the ML model externally from the non-real-time RIC, while keeping the ML Inference xApp within the near-real-time control loop.

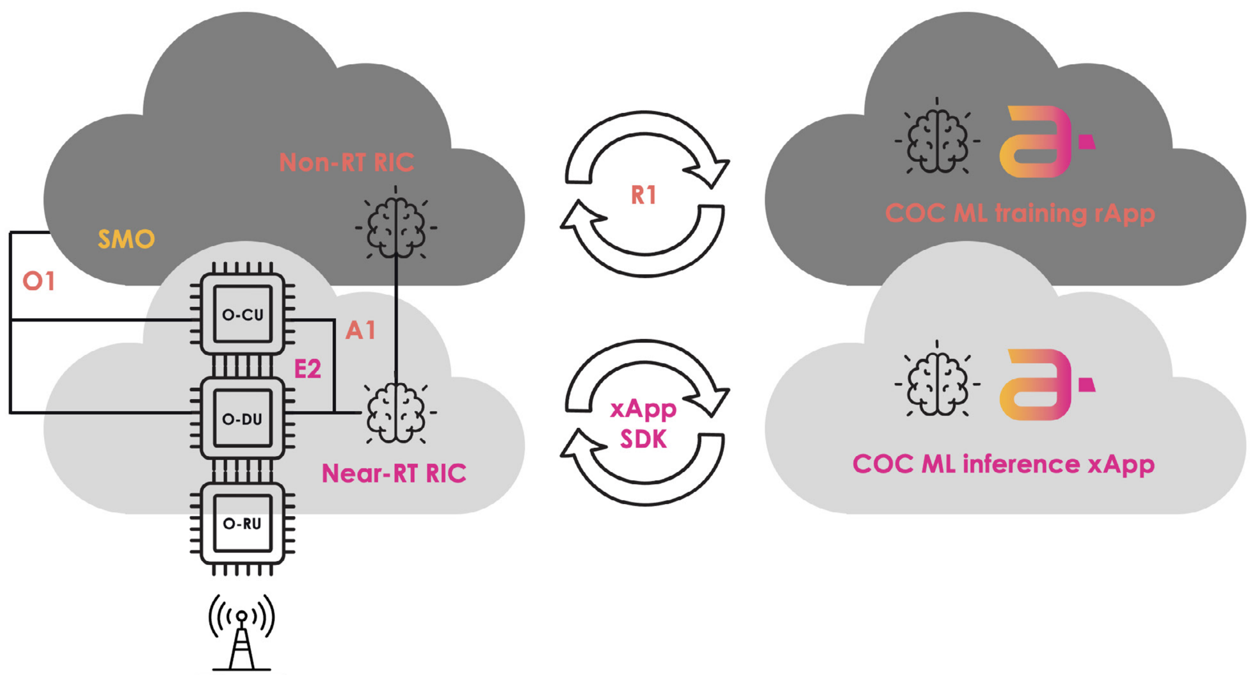

Figure 4: Amdocs Cell Outage Compensation implemented following O-RAN ML deployment scenario 1.2 (alternative version)

Machine Learning in Cell Outage Compensation

Addressing cell outages in near-real-time, particularly within a one-second window, necessitates a sophisticated approach. ML techniques are ideally suited to meet this challenge. The key is to train the system in advance using simulations that closely mimic real network conditions.

In scenarios like ours, specifically within the realm of network optimization, reinforcement learning proves to be the most effective method. This approach allows the agent to autonomously create a model that optimally maps observations to actions, relying solely on the rewards received – essentially, the network KPIs achieved during simulations.

This pre-training phase is efficiently managed by dividing it into smaller, manageable network clusters. These are then processed independently and in parallel, enabling us to scale the learning phase according to our needs. Each segment of this phase involves running a network simulator, providing the agents with realistic scenarios to learn from and adapt to different network states.

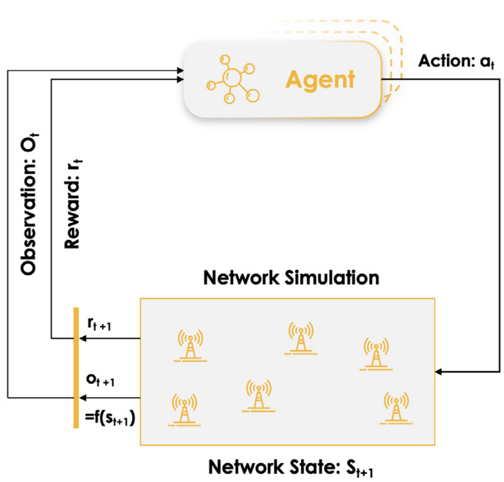

Figure 5: Reinforcement Learning Model

The training is driven by a reward function adjusted in accordance with our optimization goals. This function gives feedback to the agent, encouraging actions and state changes that align with our objectives. The ultimate aim for the agent during its training journey is to maximize the sum of these rewards over time, focusing on the most effective strategies for outage compensation.

Simulation environment

Creating a close-to-reality digital twin of an actual Radio Access Network (RAN) is crucial for our work. Our simulation framework is designed to precisely mimic the real network's topology, enabling us to cover a wide array of traffic scenarios. This includes everything from map-based situations and moving hotspots to time-based variations and their combinations.

At the core of our RAN simulation is the load-coupling model (ref 4,5,6,7), which accurately reflects the intricate interplay of interference and load dependencies among neighboring cells.

For Cell Outage Compensation (COC), our simulator is equipped to introduce cell outages randomly throughout the network, even simulating multiple outages simultaneously. This feature is invaluable during the training phase, as it allows the agents to adapt to various outage scenarios. Their learning focuses on maintaining the highest possible QoS despite these challenges.

Learning agent configurations

Exploring various configurations for RET optimization, we've tested multiple agents across different reinforcement learning (RL) setups, examining their effectiveness under various conditions:

A. Observation types

1. Single Cell Observation (SCO)
   - Pros: Operates independently without requiring data on neighboring cells
   - Cons: Limited to inputs from the targeted cell only, overlooking neighboring cell dynamics
2. Multiple Cell Observation (MCO)
   - Pros: Enhances decision-making by incorporating data from neighboring cells
   - Cons: Requires identifying relevant neighboring cells, achievable with neighborhood algorithms

B. RL algorithms

1. Deep Q Network (DQN)
   - Pros: Effectively generalizes across unseen network states
   - Cons: May not fully learn the upper end of the state range
2. Proximal Policy Optimization (PPO)
   - Pros: Achieves faster learning with fewer iterations
   - Cons: Risk of overfitting to simulated scenarios

C. Agent setup

1. Multi-Agent (MA)
   - Pros: Identifies and adapts to the distinct characteristics of each cell, enabling customized adaptation
   - Cons: Requires a dedicated model for each cell, increasing complexity
2. Single Agent (SA)
   - Pros: Simplifies model management by using one model for multiple cells
   - Cons: Faces challenges in learning individual cell peculiarities due to a generalized approach

D. Action execution

1. Multiple Actions (MAct)
   - Pros: Facilitates understanding of cell interdependencies through simultaneous adjustments
   - Cons: May limit exploration of the network state space
2. Single Action (SAct)
   - Pros: Enables thorough exploration of network states with targeted adjustments
   - Cons: Makes it challenging to grasp the interplay between cells

E-Tilt action range for agents

The range of actions available to an agent consists of potential e-tilt adjustments it can make. These ranges are subsets of the physical e-tilt options available for each cell, which vary based on the cell's specific constraints. The most effective method for determining an agent's action range involves using delta values relative to the current e-tilt of each cell. Our findings indicate that a maximum delta of ±2 degrees is generally adequate to address nearly all cell outages in urban and suburban environments.

Agent observations in the network

Agents are trained using simulations, acquiring their "knowledge" from these simulated environments. Once deployed in the actual network, they require comparable inputs from the real-world network to operate effectively. To this end, the observation space is designed to ensure that inputs in live networks – derived from cell performance measurement (PM) counters – are analogous to those in the simulation environment.

Results

To assess the effectiveness of our solution and compare various approaches and configurations, let's examine the results from a specific network scenario, as shown in Figure 6, below.

In this scenario, the agents are active in the three central cells. These cells are represented without hatching. One of these cells experiences an outage, indicated by a red flash. The neighboring cells attempt to compensate for this outage, but as the coverage plot reveals, there are coverage gaps along the cell boundaries (marked in red) where service cannot be provided because the neighboring cells are already at their capacity limits. Yellow areas highlight where traffic is redirected from overloaded cells to those with available capacity. The active agents gather data not only from their own cells but also from the neighboring cells, which are depicted with white hatching.

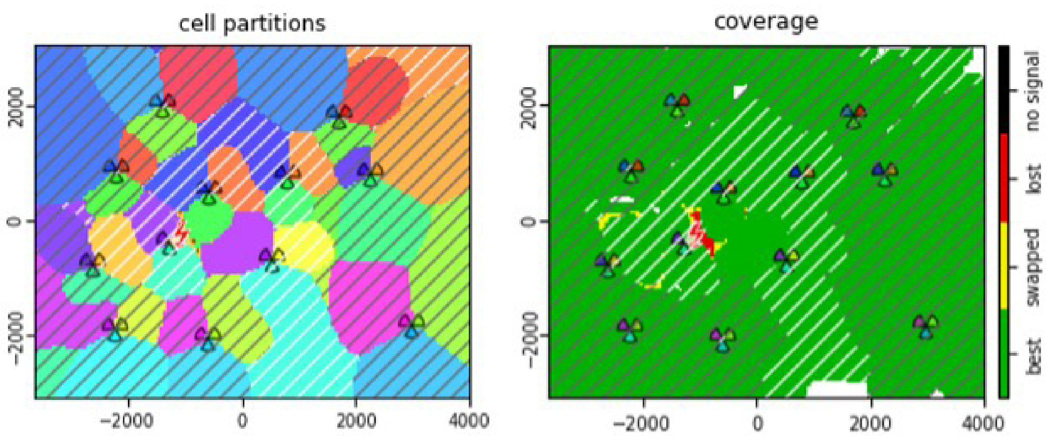

Figure 6: Cell Outage Example

Figure 7 reveals how the cell outage was effectively mitigated. The agents, guided by the measurements they observed, made tilt adjustments that successfully covered the previously exposed areas, demonstrating the system's ability to adapt and maintain coverage.

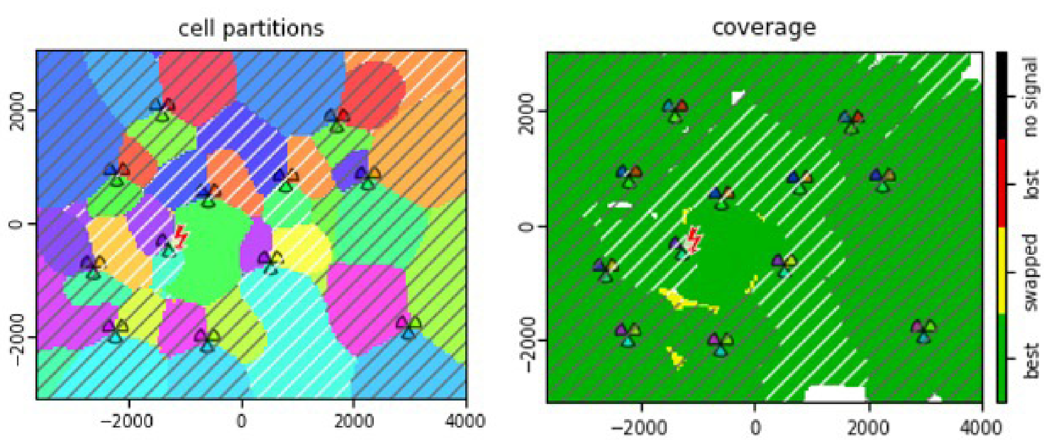

Figure 7: COC Example

Figure 8 tracks the optimization journey of the RAN cluster depicted in Figures 6 and 7. The x-axis represents each step in the process, showing how the agents adjusted e-tilts in response to evolving network conditions. When the network state remained constant, the agents retained the existing e-tilts, indicating a stable optimization. The right-most graph demonstrates the target Key Performance Indicators (KPIs) achieved with each action. Violin plots offer a comprehensive view of the potential KPIs across all network states within the cluster:

- The **black series** shows the KPIs realized by the active agents.
- The **brown series** indicates the network's KPIs had there been no e-tilt adjustments.
- The **green series** highlights the optimal KPIs theoretically attainable at each step.

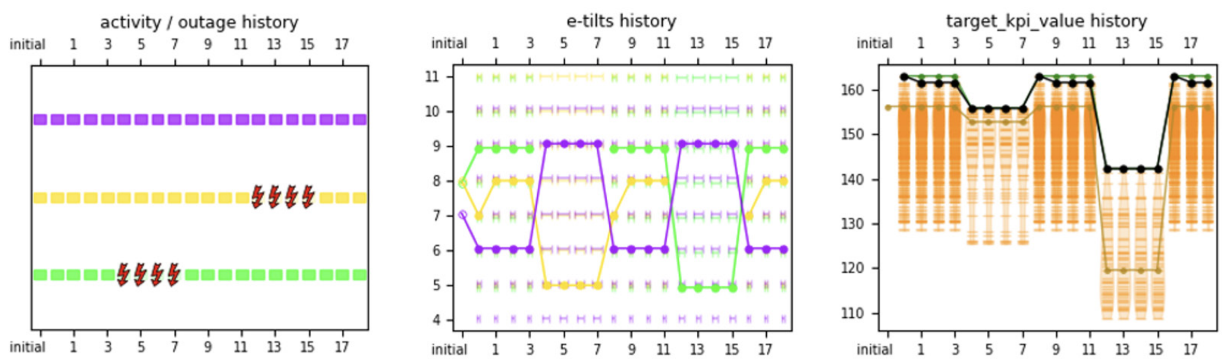

Figure 8: COC Trajectory

The KPI presented is a synthesis of fundamental metrics such as RSRP, SINR, THROUGHPUT, and LOAD. The action range for e-tilt adjustments was set to {-2, -1, 0, 1, 2} degrees, enabling precise adjustments to the cell's orientation to either improve or sustain coverage.

Figure 9 compares the RET optimization performance across various agent configurations, offering a clear visual representation of the effectiveness of each setup in enhancing network resilience and service quality.

SCO/MCO	DQN/PPO	MA/SA	MAct/SAct	KPI
SCO	DQN	MA	MAct	156
			SAct	158
		SA	SAct	155
	PPO	MA	MAct	151
			SAct	153
		SA	SAct	155
MCO	DQN	MA	MAct	159
			SAct	158
		SA	SAct	150
	PPO	MA	MAct	160
			SAct	160
		SA	SAct	157
Max achievable KPI				160
NO COC				136

Figure 9: COC Evaluation Results by optimizing RET

Overall, all agent configurations demonstrated significant improvement in network performance and effectively compensated for cell outages (COs). This improvement is particularly notable when compared to scenarios where no COC actions were taken.

However, it's clear that some configurations outperformed others. Notably, the Multi-Agent (MA) setup generally yielded slightly better results than the SA approach. This can be attributed to the MA configuration's ability to consider the specific conditions of neighboring cells (Multiple Cells Observation, or MCO), which proved more effective for COC tasks than relying solely on the data from the agent's own cell (Single Cell Observation, or SCO).

Furthermore, agents that were trained to act simultaneously (Multiple Actions, or MAct) demonstrated a better understanding and application of the interdependence of their actions, enhancing their overall performance in maintaining network quality.

To validate our COC solution we successfully demonstrated a closed-loop optimization process back into a network simulator, hereby replacing the cellular network on the right side in Figure 3 by another, independent network simulation to produce independent training and live data sets. This process further confirms the robustness and effectiveness of our approach.

Deployment and model management

The deployment and ongoing management of models are overseen by a dedicated machine learning model management (ref. 2), responsible for organizing the deployment and monitoring the status of each model. Upon deployment in the network, an agent model not only begins its optimization tasks but also continues its training, using real network data as online input. This continuous learning process allows the model to fine-tune its strategies and adjust to local network conditions that were not fully anticipated in the simulation phase.

However, if there's an external change to a cell's configuration that falls outside the model's predefined action range, the model is considered outdated. When this occurs, the model management is tasked with initiating a retraining process for the affected model to align it with the updated network environment. This ensures that the models remain effective and relevant, capable of responding to the dynamic nature of network conditions and configurations.

Conclusion

Addressing cell outages is key to preserving Quality of Service (QoS) in mobile networks, playing a crucial role in ensuring network coverage and enhancing capacity under the dynamic conditions of modern network operations. This paper underscores the effectiveness of Amdocs Cell Outage Compensation within the O-RAN near-real-time control loop and framework, demonstrating its potential to adapt and respond to the challenges presented by today's complex network environments.

Our findings affirm that reinforcement learning is a fitting method for compensating cell outages quickly and efficiently, allowing the network to adapt seamlessly when a cell becomes unavailable or is restored.

When considering deployment in actual network clusters, the robustness of the solution is paramount. To this end, our training approach encompasses a wide array of network conditions, including variations in tilts and other parameters, alongside diverse traffic scenarios, to promote agent generalization. The robustness of trained agents is further validated against scenarios not previously seen during their training, ensuring their preparedness for real-world deployments. Once deployed, these agents continue to evolve by continuous online learning from real network data, enabling them to fine-tune their responses to specific local conditions that simulations cannot fully replicate.

Finally, our approach aligns with various ML deployment and workflow scenarios outlined by the O-RAN alliance, demonstrating its versatility and applicability in current and future network architectures.

References

[1]

O-RAN Alliance: O-RAN Architecture Description, O-RAN.WG1.OAD-R003-v10.00, October 2023.

[2]

O-RAN Alliance: AI/ML workflow description and requirements, O-RAN.WG2.AIML-v01.03, October 2021.

[3]

O-RAN Alliance: Use Cases Detailed Specification, O-RAN.WG1.Use-Cases-Detailed-Specification-R003-v12.00, October 2023.

[4]

H. Klessig, A. Fehske, J. Voigt, and G. Fettweis: Improving Coverage and Load Conditions Through Joint Adaptation of Antenna Tilts and Cell Selection Rules in Mobile Networks, IEEE 2nd International Workshop on Self-Organizing Networks (IWSON) at the 9th International Symposium on Wireless Communication Systems (ISWCS), Paris, France, August 2012.

[5]

A. Fehske, H. Klessig, J. Voigt, and G. Fettweis: Concurrent Load-Aware Adjustment of Cell Selection and Antenna Tilts in Self-Organizing Radio Networks, IEEE Transactions on Vehicular Technology, Special Section on Self-Organizing Radio Networks, Vol. 62, Issue 5, ISSN 0018-9545, pages 1974 – 1988, June 2013.

[6]

H. Klessig, A. Fehske, J. Voigt, and G. Fettweis: Cell Load-Aware Energy Saving Management in Self-Organizing Networks, IEEE Vehicular Technology Conference (VTC '13 Fall), Las Vegas, NV, USA, September 2013.

[7]

A. Fehske, H. Klessig, J. Voigt, and G. Fettweis: Flow-Level Models for Capacity Planning and Management in Interference-Coupled Wireless Data Networks, IEEE Communications Magazine, vol. 52, no. 2, pages 164-171, February 2014.

Abbreviations

AI	Artificial Intelligence
CCO	Coverage and Capacity Optimization
COC	Cell Outage Compensation
DQN	Deep Q Network
ESM	Energy Savings Management
KPI	Key Performance Indicator
MA	Multi Agent
MAct	Multiple Actions
MCO	Multiple Cell Observation
ML	Machine Learning
PM	Performance Management
PPO	Proximal Policy Optimization
QoS	Quality of Service
RAN	Radio Access Network
rApp	Application on a non-real-time RIC
RL	Reinforcement Learning
SA	Single Agent
SAct	Single Action
SCO	Single Cell Observation
SON	Self-Organizing Networks
xApp	Application on a near-real-time RIC

Amdocs helps those who build the future to make it amazing. With our market-leading portfolio of software products and services, we unlock our customers' innovative potential, empowering them to provide next-generation communication and media experiences for both the individual end user and large enterprise customers. Our approximately 30,000 employees around the globe are here to accelerate service providers' migration to the cloud, enable them to differentiate in the 5G era, and digitalize and automate their operations.

Listed on the NASDAQ Global Select Market, Amdocs had revenue of $4.89 billion in fiscal 2023.

For more information, visit Amdocs at www.amdocs.com.

amdocs

make it amazing

www.amdocs.com