This is the authors' accepted version of the article. The final version published by IEEE is M.A. Rahman, A. Duque, N. Chukhno, B. Niżnik, M. Gramaglia, and A. Garcia-Saavedra, "Conflict Mitigation of xApps and Interoperability of O-RAN Components: The ORIGAMI Approach," 2026 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit), doi: TBD.

Conflict Mitigation of xApps and Interoperability of O-RAN Components: The ORIGAMI Approach

Md Arifur Rahman*, Alexis Duque†, Nadezhda Chukhno‡, Bartosz Niżnik*, Marco Gramaglia§, and Andres Garcia-Saavedra¶
*IS-Wireless, Poland †Net AI Tech Ltd, UK ‡IMDEA Networks, Spain
§Univ. Carlos III of Madrid, Spain ¶NEC Laboratories Europe, Germany

***Abstract*—As 6G architectures evolve toward pervasive Network Intelligence (NI), managing concurrent and potentially conflicting control actions from diverse Artificial Intelligence (AI)-driven xApps becomes a critical challenge. In the disaggregated Open Radio Access Network (O-RAN) ecosystem, independent xApps optimizing for different metrics—such as energy efficiency versus throughput—can lead to network instability and degraded performance. This paper presents a conflict management solution within the ORIGAMI framework, specifically addressing the contention between the Interoperable Machine Learning models for RAN Energy efficiency (IMLE) xApp (energy optimization) and a competing Throughput xApp (throughput maximization). We demonstrate how ORIGAMI's Global Service-Based Architecture (GSBA) facilitates the real-time discovery and mediation of these conflicting objectives. Our experimental results in a distributed O-RAN testbed show that by leveraging GSBA-exposed services, the proposed conflict management logic effectively prioritizes critical network requirements, ensuring stability and resource efficiency without compromising the performance targets of high-priority services.**

***Index Terms*—O-RAN, Conflict Mitigation, Network Intelligence, xApp, Service-Based Architecture.**

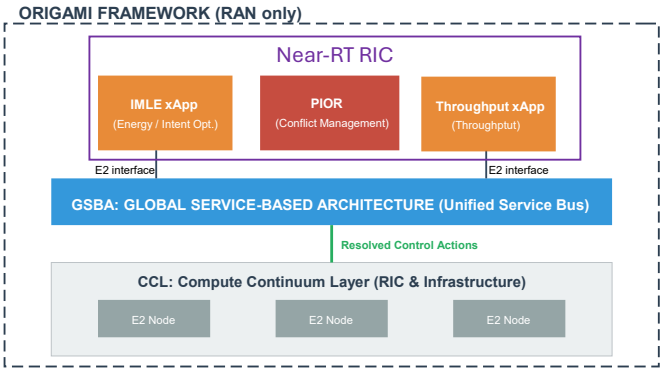

Fig. 1: The ORIGAMI's GSBA acts as the interaction layer for specialized xApps. In this setup, PIOR acts as a mediator to resolve conflicts between the IMLE xApp (energy optimization) and a testing xApp (a PRB quota optimization xApp for throughput optimization).

I. Introduction

The transition from 5G to 6G is characterized by an unprecedented integration of Artificial Intelligence (AI) and Machine Learning (ML) across all network layers. To manage architectural and operational complexity, the industry has turned to disaggregated architectures, such as Open Radio Access Network (O-RAN), that enable the deployment of specialized applications (xApps) to optimize specific network functions. By decoupling the control logic from the radio hardware, the Near-Real-Time RAN Intelligent Controller (Near-RT RIC) allows operators to onboard "best-of-breed" solutions from multiple vendors. However, this flexibility introduces significant operational risks. As independent xApps operate simultaneously on the same underlying infrastructure, the probability of conflicting control actions increases drastically. This is particularly critical when one application's optimization objective (*e.g.*, energy efficiency) directly counteracts another's (*e.g.*, throughput maximization), leading to "parameter flipping", network instability, and degraded Quality of Service (QoS).

The ORIGAMI project[1] addresses these challenges by proposing a unified architecture designed to transcend current domain silos and ensure coherent network orchestration. The ORIGAMI framework is built upon three fundamental pillars: the Global Service-Based Architecture (GSBA), the Zero Trust Exposure Layer (ZTL), and the Compute Continuum Layer (CCL) [1]–[3]. While the ZTL ensures secure service exposure and authentication across multi-domain environments, and the CCL provides a distributed execution environment for network functions, the GSBA serves as the central orchestration layer.

As illustrated in Fig. 1, the GSBA extends the service-based principles of the 5G Core into the Radio Access Network (RAN). Unlike legacy O-RAN implementations that often rely on point-to-point interfaces, the GSBA provides a "unified service bus" and a set of standardized Global Service APIs. This architecture enables a degree of cross-app awareness and coordination previously unattainable, allowing the platform to intercept and mediate conflicting control intents before they reach the radio hardware.

This paper focuses on the Platform for Interoperable O-RAN (PIOR) use case as a primary validation of the conflict mitigation capabilities enabled by ORIGAMI's GSBA. Defined formally in the project specifications [4] as the *Conflict Mitigation of xApps and Interoperability of O-RAN components*, PIOR is not merely another xApp, but a platform-level functionality designed to detect and resolve conflicts between various Network Intelligence (NI) solutions. To evaluate this

[1]https://sns-origami.eu/

capability, we establish a conflict scenario involving the Interoperable Machine Learning models for RAN Energy efficiency (IMLE) use case, also defined in [4]. IMLE utilizes interoperable ML models to optimize RAN energy efficiency by forecasting traffic demand and dynamically adjusting Physical Resource Block (PRB) quotas.

To generate a contention for resources against IMLE, we introduce a distinct PRB optimization xApp designed to maximize throughput (hereafter referred to as the "Throughput" xApp). In this context, PIOR acts as the architectural mediator, leveraging the GSBA to monitor the conflicting control requests from IMLE and the Throughput xApp. By analyzing the intent of both applications, PIOR enforces global policies that ensure infrastructure stability and compliance, effectively preventing the Throughput xApp from destabilizing the energy-saving profile of IMLE.

The main contributions of this work are as follows:

- We describe the integration of conflict management logic within ORIGAMI's GSBA, demonstrating how service exposure enables the real-time detection of overlapping control loops between multi-vendor xApps.
- We present the design and experimental validation of the PIOR conflict mitigation mechanism in a distributed O-RAN testbed.
- We demonstrate how the architecture successfully mediates between the conflicting objectives of IMLE (energy) and the Throughput xApp, ensuring that energy-saving intents do not compromise critical infrastructure Key Performance Indicators (KPIs) or Service-Level Aggreement (SLA) requirements for high-priority slices.

The remainder of this paper is structured as follows. Section II provides the necessary background on O-RAN conflict types and the specific xApps involved. Section III details the design of the Conflict Mitigation framework and the operational procedures of the PIOR mediator. Section IV describes the implementation details and the experimental testbed setup. Finally, Section V presents the performance evaluation and quantitative results, followed by the conclusion in Section VI.

II. Background and Problem Statement

The evolution toward O-RAN introduces functional disaggregation and programmable intelligence into the RAN, enabling flexible deployment of control applications from multiple vendors. While this increases innovation and adaptability, it also creates challenges related to coordination and conflict management among independently operating controllers. This section reviews O-RAN disaggregation and intelligence, defines key types of multi-vendor conflicts, and presents ORIGAMI's GSBA-based approach for conflict management.

A. O-RAN Disaggregation and Intelligence

The O-RAN architecture represents a fundamental shift in RAN design, transitioning from monolithic, proprietary hardware to a disaggregated stack composed of centralized, distributed, and radio units. A cornerstone of this architecture is the Near-RT RIC, which introduces a programmable abstraction layer between the radio hardware and the control logic. The Near-RT RIC hosts xApps—microservice-based applications—that optimize radio resource management through the E2 Interface (E2). These xApps utilize the Shared Data Layer (SDL) for persistent data storage and state management, decoupling logic from state. Control operations are standardized via Service Models. For instance, the Key Performance Metric (KPM) provides high-granularity measurement reporting, while the E2 Service Model for RAN Control (E2SM-RC) enables granular control over cell configurations, scheduling policies, and User Equipment (UE)-level associations.

B. The Multi-Vendor Conflict Problem

A primary commercial and technical advantage of O-RAN is the ability to onboard xApps from multiple vendors, creating a "best-of-breed" ecosystem. However, this flexibility introduces significant operational risks that did not exist in single-vendor, legacy networks.

In legacy RAN, the Radio Resource Management (RRM) functions (*e.g.*, scheduler, handover control, power control) are tightly integrated and validated by a single vendor. In O-RAN, xApps are designed as independent, closed-loop controllers targeting specific metrics (*e.g.*, IMLE focuses on reducing power consumption, while a traffic steering xApp focuses on load balancing). In a multi-vendor environment, two xApps may simultaneously attempt to modify the same parameters without awareness of each other's actions.

Without a centralized authority to ensure data consistency, these overlapping loops can lead to unstable network states, known as "parameter flipping". In this state, the system oscillates between conflicting configurations (*e.g.*, toggling a handover parameter back and forth), leading to signaling storms and degraded user experience. According to O-RAN technical specifications [5], [6], these conflicts are categorized as follows:

- **Direct Conflict:** Multiple xApps request different values for the same parameter (*e.g.*, PRB quota) simultaneously. This is the most severe form of conflict as it leads to immediate race conditions in the database.
- **Indirect Conflict:** xApps target different parameters that influence the same metric. For example, one xApp increases transmit power to improve throughput, while another tilts the antenna down to reduce interference; both actions fundamentally alter the Signal-to-Interference-plus-Noise Ratio (SINR) landscape, potentially leading to over-correction or "overshooting" the optimization target.
- **Implicit Conflict:** Optimization of one metric (*e.g.*, coverage) unintentionally harms another metric (*e.g.*, throughput in a neighboring cell). These are often difficult to detect without holistic network visibility.

In this paper, we specifically address **Direct Conflicts**, where both the Throughput and IMLE xApps attempt to modify the same RAN control parameter (slice-level PRB quota) to optimize throughput or energy consumption, respectively.

C. ORIGAMI and the GSBA Solution

The ORIGAMI framework addresses these issues by mitigating the reliance on point-to-point xApp-to-E2 communication. In standard O-RAN deployments, conflict mitigation is often left to the internal logic of a specific vendor's RIC. ORIGAMI introduces a standardized architectural approach:

1) **GSBA**, a unified service bus where control intentions are exposed as services, serves as the central orchestration layer, extending the service-based principles of the 5G Core into the RAN.
2) **ZTL** ensures that xApps are authenticated and authorized before accessing sensitive radio data.
3) **CCL** provides the distributed execution environment.

By intercepting service-based control requests via the GSBA, the PIOR use case acts as a global conflict management function. It verifies whether a request violates existing policies or conflicts with higher-priority tasks before any signaling is sent over the E2 interface.

III. Design of the Conflict Mitigation Framework

The optimization of RAN resources through xApps running in a Near-RT RIC introduces the risk of conflicting operations on shared RAN parameters. To enable seamless coordination between independent control loops, we propose a Direct Conflict Mitigation (DCM) framework. This framework serves as the reference implementation of the PIOR NI solution defined in [4], embedded within the ORIGAMI's GSBA. It leverages the unified service bus to detect and resolve conflicts in real-time by monitoring the targeted control parameters and KPIs exposed by active xApps.

A. Problem Formulation of the PIOR Logic

We define the set of active xApps registered in the GSBA as $\mathcal{A} = \{a_1, a_2, ..., a_n\}$. Each xApp a_i generates a control request R_i at time t, targeting a specific set of RAN parameters $P_i \subset \mathcal{P}$, where $\mathcal{P}$ is the universal set of controllable E2SM-RC parameters (*e.g.*, *RRMPolicyRatio*, *CellTransmissionPower*).

The PIOR conflict detection logic identifies a **Direct Conflict** C_{dir} at time t if the intersection of the target parameter sets of any two active xApps is non-empty:

$$C_{dir}(t) = \begin{cases} 1, \exists\{a_i, a_j\} \subseteq \mathcal{A} : (p \in \mathcal{P}_i \cap \mathcal{P}_j) \wedge (V_i(p) \neq V_j(p)), \\ 0, \text{otherwise.} \end{cases} \tag{1}$$

where $V_i(p)$ and $V_j(p)$ represent the target values requested by xApps a_i and a_j, respectively. For simplicity, the time index t is omitted on the right-hand side here and thereafter.

In our specific scenario, we consider two agents:

- a_{IMLE}: the IMLE xApp (energy efficiency).
- a_{Thp}: the Throughput xApp (throughput maximization).

Both agents target the parameter $p_{quota} \in \mathcal{P}$, representing the slice-level PRB quota. Since $P_{IMLE} \cap P_{Thp} = \{p_{quota}\}$, a direct conflict is guaranteed whenever the requested values for p_{quota} differ.

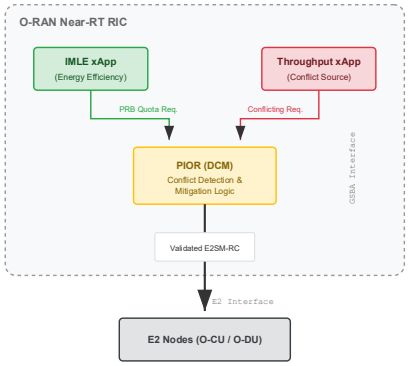

Fig. 2: High-level schematic of the conflict mitigation use case. The PIOR mediator intercepts requests from the IMLE (energy) and Throughput xApps to prevent conflicting write operations.

B. Priority-Based Mitigation Strategy

To mitigate detected conflicts, the PIOR framework enforces a priority-based resolution strategy (see Fig. 2). A priority hierarchy is established where one xApp's decision takes precedence based on the current network state. For instance, a policy can be enforced dictating that during periods where PRB demand exceeds availability, specific slices maintain precedence, while lower-priority allocations are reduced or rejected. More specifically, a priority function $\Phi(a_k)$ is assigned to each xApp or, equally, to the specific slice intent (*e.g.*, precise requirements that a network must meet). Thus, the resolution logic ensures that:

$$R_{final} = R_{k^*}, \text{where } k^* = \arg\max_{k \in \{i,j\}} \Phi(a_k). \tag{2}$$

C. Operation Procedures of the CM Framework

The proposed DCM logic operates by continuously monitoring the control parameters targeted by each xApp. By utilizing the GSBA as the interaction layer, the framework gains visibility into the intents of all registered services.

1) Conflict Detection: At the stage of conflict detection, the framework leverages the central database at the Near-RT RIC. By querying the intent definitions of active xApps registered in the GSBA, the system identifies conflicts by matching the specific Object IDs of the target parameters (*e.g.*, the specific *RRMPolicyRatio* associated with a Single Network Slice Selection Assistance Information (S-NSSAI)).

2) Resolution and Actuation: Once a conflict is identified, the mitigation strategy triggers the prioritization logic before the control decision is sent to the underlying E2 nodes. This prevents the "ping-pong" effect of conflicting writes. To illustrate this, Fig. 3 provides a detailed sequence comparison: without mitigation (Fig. 3a), conflicting requests from the Throughput xApp and IMLE cyclically overwrite each other at the E2 node level. Conversely, with mitigation enabled (Fig. 3b), the PIOR logic intercepts the request within the GSBA, enforces the priority policy, and blocks the conflicting command before it reaches the RAN.

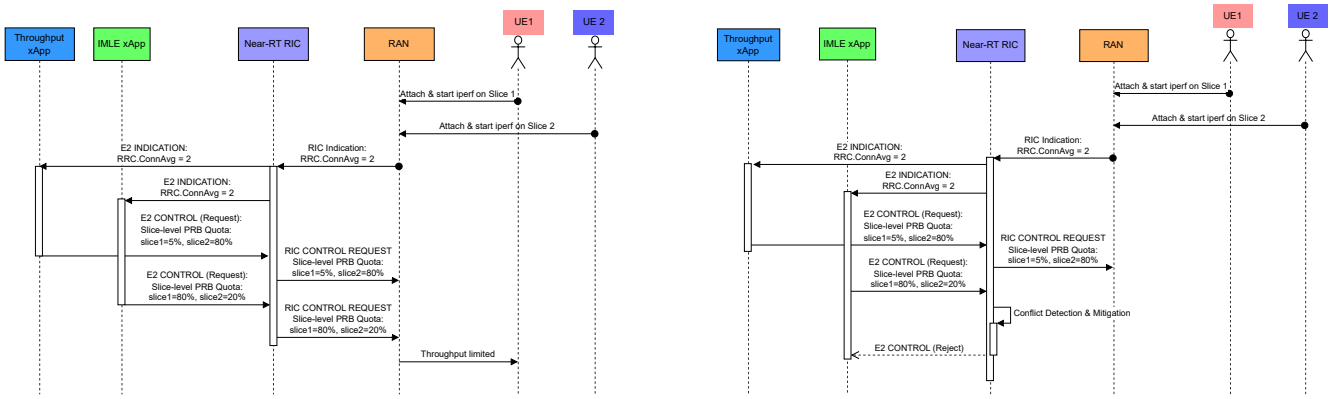

(a) Without conflict mitigation logic.

(b) With conflict detection and mitigation.

Fig. 3: Sequence diagram showing conflict detection and resolution logic (PIOR) between the Throughput and IMLE xApps.

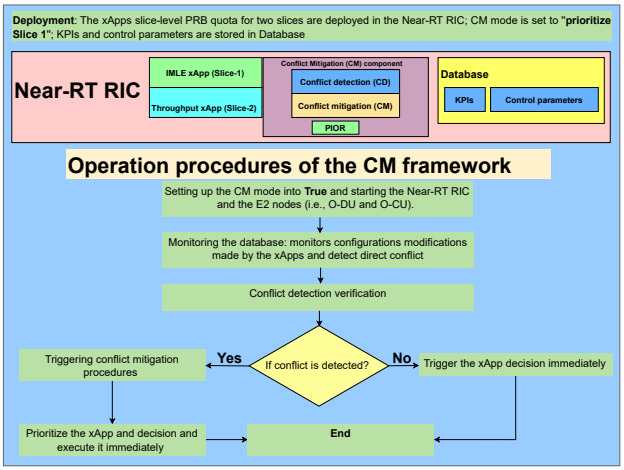

Fig. 4: The operational procedures of the conflict mitigation framework.

The operational procedures, as depicted in Fig. 4, follow a strict sequence:

1) The Near-RT RIC and E2 nodes are initialized with Conflict Mitigation (CM) mode enabled.
2) The DCM logic monitors the GSBA service registry and shared database for incoming configuration modifications.
3) Upon detecting overlapping requests (*e.g.*, IMLE and Throughput xApps are targeting the same slice), the prioritization logic is executed.
4) Only the validated decision is forwarded to the RAN.

This scheme improves interoperability and ensures seamless operation among multiple vendors. We demonstrate this in our implementation by enabling seamless coordination between xApp decisions, avoiding conflicting write operations, and maintaining stable throughput even when the IMLE and Throughput xApps generate competing requests on the same platform.

D. Interoperable ML model for PRB Quota determination

The IMLE xApp serves as the primary test subject for this interoperability framework. Its objective is to optimize RAN energy consumption through highly accurate forecasting of PRB utilization, thereby enabling more efficient and scalable network resource management.

IMLE subscribes to KPMs—specifically, PRB usage in both Uplink (UL) and Downlink (DL)—exposed via the GSBA from the E2 nodes. Furthermore, the xApp consumes data from the RAN Intelligent Controller (RIC) SDL, which acts as a central repository for information contributed by various third-party xApps. By subscribing to this data, ORIGAMI xApps gain access to a rich pool of information essential for fine-tuning underlying ML models.

However, relying on third-party data introduces inherent risks related to data quality. To mitigate this challenge, IMLE leverages the resolution capabilities built into the GSBA-enabled platform. These mechanisms serve as a governance layer, ensuring data consistency in the SDL and arbitrating conflicting control requests. This is particularly crucial given the potential for unreliable forecasts stemming from opaque third-party xApp behaviors.

E. Actuation and Control Logic via O-RAN Interfaces

Upon generating traffic forecasts, the IMLE optimization logic triggers reconfiguration commands to adapt the network capacity. The actuation mechanism relies strictly on O-RAN E2 interface standards, specifically utilizing the E2SM-RC specifications.

The xApp transmits an E2 Control Request structured to comply with standards. The payload utilizes the *E2SM-RC Control Header Format 1* and *E2SM-RC Control Message Format 2*. Within this message, the logic targets the RRM policy parameters, specifically those associated with *Slice-specific Quota Control*. The control action is directed at the *RRMPolicyRatio* attribute, enabling the dynamic modification of PRB quotas—specifically adjusting the **min**, **max**, and **dedicated PRB ratios** for specific *S-NSSAIs*. This granular control allows the system to selectively scale radio resources for each network slice based on real-time traffic demand and SLA requirements.

Fig. 5: Performance evaluation of the PIOR CM framework. The telemetry data (captured via Grafana) demonstrates the stability of Slice 2 PRB allocation (top center) and Throughput (bottom center) during the mitigation period (approx. 10:47:00 to 10:51:00), despite concurrent conflicting requests from the Throughput xApp.

IV. Implementation

This section details the experimental platform, system configuration, and validation methodology used to evaluate the proposed DCM framework.

A. Experimental Setup

To validate the proposed DCM framework, we deployed a distributed O-RAN evaluation testbed configured for a realistic 5G New Radio (NR) deployment scenario, operating in TDD mode on band N78 with a 40 MHz channel bandwidth. The network slicing configuration is partitioned into two distinct slices:

- **Slice 1 (SST1):** Configured as a lower-priority service.
- **Slice 2 (SST2):** Configured as a high-priority service with strict QoS requirements.

The Near-RT RIC hosts the critical components required for the validation scenarios. To isolate the performance of the conflict mitigation logic from the stochastic nature of ML training, we utilize an xApp emulating the IMLE behavior. This emulator generates E2SM-RC control messages representing an optimized, energy-efficient resource allocation (specifically, a stable 80% quota for the high-priority slice). This is contrasted against the Throughput xApp, which acts as the conflict generator.

B. Test Methodology

We defined a test protocol to isolate the impact of the conflict mitigation logic over a 55-second window.

1) Phase 1: Baseline (CM Mode False): In the initial phase, the Near-RT RIC and RAN are started with the DCM mode set to "False". Downlink traffic is initiated using *iperf3* on both slices. We establish the baseline behavior where the IMLE emulator attempts to enforce its policy (70%-80% allocation), while the Throughput xApp operates blindly.

2) Phase 2: Conflict Generation: To test resilience, the Throughput xApp triggers aggressive, conflicting E2SM-RC decisions every 10 seconds, attempting to overwrite the IMLE emulator's configuration. Without mitigation, this results in a "flipping behavior" where PRB allocations and throughput oscillate rapidly.

3) Phase 3: Mitigation (CM Mode True): In the final phase, the conflict mitigation mode is enabled. The DCM logic utilizes the higher priority assigned to the IMLE intent to reject the PRB adjustment requests from the Throughput xApp. We validate this success by observing stable throughput and unchanged PRB allocations over the entire test period.

V. Performance Evaluation

In the following, we evaluate the proposed conflict mitigation framework.

A. Experimental Metrics and Analysis

To evaluate the effectiveness of the PIOR conflict mitigation logic, we conducted a series of tests comparing network performance with the mitigation logic enabled and disabled. The primary metrics for evaluation were the stability of the PRB quota allocation and the resulting UE throughput for the high-priority slice (*i.e.*, Slice 2).

1) Visual Analysis of Mitigation: Fig. 5 presents the real-time telemetry captured during the execution of the conflict scenario with CM mode enabled. The dashboard visualizes the behavior of two network slices under contention:

- **Slice 1 (Low Priority):** Allocated a minimal operational quota (≈5%).
- **Slice 2 (High Priority):** Targeted for 80% resource allocation to support high-throughput services.

As observed in the "Slice 2 PRB quota (DL)" panel (top center), the allocation remains highly stable at 80% for the duration of the test (from 10:47:00 to 10:51:00). This stability persists despite the concurrent execution of the conflicting Throughput xApp, which actively attempts to request resource reallocation. The corresponding throughput for Slice 2 (bottom

TABLE I: Experimental comparison of throughput stability.

Evaluation Stage	Test	Slice 1 TP (5% PRB)	Slice 2 TP (Target 80%)
Conflict Mitigation OFF	1	10 Mbps	148 Mbps (80%)
Conflict Mitigation ON	1	10 Mbps	130 Mbps (80%)
Conflict Mitigation OFF	2	10 Mbps	**28.7 Mbps (20%)**
Conflict Mitigation ON	2	10 Mbps	**130 Mbps (80%)**

center) mirrors this stability, maintaining a steady rate of approximately 130 Mbps.

This visual evidence confirms that the DCM logic successfully intercepts and neutralizes the conflicting control messages from the Throughput xApp, thereby enforcing the energy-efficient policy defined by the IMLE emulator.

2) Throughput Stability Comparison: To provide a more granular analysis, we aggregated the throughput statistics for two distinct test runs: Test 1 (Baseline) and Test 2 (Conflict Scenario). Table I summarizes the results.

In Test 1, with the conflict mitigation disabled but no active conflict, the system performs as expected, delivering 148 Mbps to Slice 2. However, the critical observation occurs in Test 2. In the scenarios where **Conflict Mitigation is OFF**, the system is vulnerable to the "flipping behavior" observed every 10 seconds. As seen in Table I, the aggressive Throughput xApp successfully overwrites the configuration intended by the IMLE emulator, causing Slice 2's PRB allocation to plummet to 20%. Consequently, the throughput for the high-priority service degrades significantly, dropping to 28.7 Mbps.

In contrast, when mitigation is enabled in the same conflict scenario (Test 2, **Conflict Mitigation is ON**), the DCM framework enforces priority, ensuring a consistent throughput of 130 Mbps. This effectively isolates the high-priority service from multi-xApp instability.

B. Discussion and Operational Implications

The experimental results highlight the critical role of the DCM framework in maintaining network stability.

1) Stability vs. Agility: Without the DCM framework, the network enters a state of instability driven by the frequent update intervals (every 10 seconds) of the conflicting Throughput xApp. This "flipping" behavior not only degrades the QoS for high-priority slices but also indicates a lack of coordination in the control plane. By enforcing a strict priority rule, the DCM logic acts as a necessary stabilizer, preventing rapid oscillations that would otherwise disrupt user sessions.

2) Policy Enforcement in O-RAN: The results demonstrate that conflict resolution can be effectively managed at the platform level. By exposing xApp intents and monitoring the database for modifications, the ORIGAMI framework enables the enforcement of global policies that supersede individual xApp logic. This capability ensures that critical infrastructure targets, such as the 80% resource allocation for Slice 2, are preserved even when experimental or lower-priority applications attempt to modify the shared configuration.

VI. Conclusion

In this paper, we addressed the critical challenge of interoperability and conflict management in multi-vendor O-RAN environments. As AI-driven xApps become more prevalent, the risk of "parameter flipping" and service degradation due to overlapping control loops increases significantly.

We proposed a Direct Conflict Mitigation (DCM) framework integrated within the ORIGAMI's GSBA. By shifting from point-to-point interfaces to a unified service bus, our architecture enables real-time visibility into the intents of active xApps. We validated this approach through a specific use case involving the IMLE xApp, which seeks to minimize energy consumption, and a "Throughput" xApp, which aggressively optimizes for throughput. Experimental results from our distributed O-RAN testbed demonstrated that without mitigation, these competing xApps lead to severe instability, with PRB allocations oscillating rapidly and throughput dropping to as low as 28 Mbps. However, upon activating the DCM logic, the system successfully enforced priority policies, stabilizing the high-priority slice at 130 Mbps and 80% PRB utilization despite concurrent conflicting requests. These findings confirm that the ORIGAMI's GSBA provides a robust foundation for 6G network orchestration, ensuring that the pursuit of energy efficiency does not compromise critical QoS standards. Future work will extend this framework to handle indirect conflicts and integrate more complex, ML-driven arbitration policies.

Acknowledgments

This work has been supported by the European Commission through Grant No. 101139270 (ORIGAMI). N. Chukhno's work is funded by the European Union, Grant Agreement No. 101206327 (6G-AI-TANGO). This work is also funded and supported by project FENG.01.01-IP.02-4543/23.

References

[1] L. E. Chatzieleftheriou, M. Gramaglia, A. Garcia-Saavedra, S. Gebert, G. Garcia-Aviles, S. Geissler, M. Fiore, P. Patras, A. Lutu, D. Tsolkas *et al.*, "Towards 6G: Architectural Innovations and Challenges in the ORIGAMI Framework," in *2024 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, 2024, pp. 1139–1144.

[2] L. E. Chatzieleftheriou, D. D. A. Hernandez, S. Bizzarri, M. Fiore, M. Fodrini, A. Garcia-Saavedra, M. Gramaglia, E. Municio, and D. Tsolkas, "6G Standardization Potential of the ORIGAMI Novel Architectures and Use Cases," in *2025 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, 2025, pp. 387–392.

[3] ORIGAMI Consortium, "Initial Architecture of ORIGAMI, New Business Models and Refinement of Requirements and KPIs," ORIGAMI Project (Horizon Europe GA 101139270), Tech. Rep. Deliverable D2.2, Apr. 2025. [Online]. Available: https://sns-origami.eu

[4] ——, "Description of NI Solutions and Services Exploiting Zero-Trust, Exposure and Compute Continuum Layers," ORIGAMI Project (Horizon Europe GA 101139270), Tech. Rep. Deliverable D3.1, Oct. 2024. [Online]. Available: https://sns-origami.eu

[5] O-RAN Working Group 3 (Near-Real-time RAN Intelligent Controller and E2 Interface), "Conflict Mitigation," O-RAN Alliance, Technical Report O-RAN.WG3.TR.ConMit-R004-v01.00, October 2024.

[6] C. Adamczyk and A. Kliks, "Conflict Mitigation Framework and Conflict Detection in O-RAN Near-RT RIC," *Comm. Mag.*, vol. 61, no. 12, p. 199–205, Dec. 2023. [Online]. Available: https://doi.org/10.1109/MCOM.018.2200752