

Mobility-Aware Edge Rebinding in 5G Networks using SR4MEC

Louis Royer
IRIT, Université de Toulouse,
CNRS, Toulouse INP
Toulouse, France
louis.royer@irit.fr

Emmanuel Lavinal
IRIT, Université de Toulouse,
CNRS, Toulouse INP
Toulouse, France
emmanuel.lavinal@irit.fr

Emmanuel Chaput
IRIT, Université de Toulouse,
CNRS, Toulouse INP
Toulouse, France
emmanuel.chaput@irit.fr

Abstract—The integration of edge computing in 5G networks requires the ability to dynamically select and access service instances located close to mobile users. While existing 5G mechanisms preserve PDU session continuity during handover, they do not directly address the problem of updating the edge instance used by an ongoing session. This paper extends SR4MEC, an architecture based on Segment Routing over IPv6 for the Mobile User Plane, to support mobility-aware edge rebinding. SR4MEC exposes the SRv6 domain as a single logical UPF to the 5G control plane and relies on an SRv6 controller to translate PFCP updates into data-plane rules. We define procedures that integrate edge rebinding with 5G handover, allowing the network either to preserve the current edge instance or to immediately switch to an instance in the target area. The proposed approach remains transparent to the UE and applications, while avoiding the reconfiguration of chains of GTP-U tunnels.

Index Terms—5G/6G, Edge Computing, Segment Routing.

I. INTRODUCTION

Edge computing is a key enabler for latency-sensitive services in 5G and beyond mobile networks. By deploying computing and storage resources close to users, Mobile Network Operators (MNOs) can reduce end-to-end latency and improve quality of experience. In such an environment, a service may be deployed through several instances distributed across central and edge data networks, and the network must determine which instance should serve a given User Equipment (UE).

This decision involves both binding and access. Binding associates a UE, or more precisely a PDU session, with a service instance according to an operator policy. Access ensures that packets are effectively forwarded to the selected instance. These operations should remain transparent to the UE and to the application, while allowing the MNO to enforce policies depending on the slice, the service, the user location, or the current state of the edge infrastructure.

Mobility makes this problem harder. Standard 5G handover procedures preserve the PDU session by updating the path between the UE and the mobile core, but they do not address the specific issue of edge-instance selection. After a handover, the previously selected instance may no longer be optimal, as the UE may have moved closer to another edge. Conversely, changing the instance must be carefully coordinated with the handover procedure to preserve service continuity and avoid packet loss. Mobility-aware edge computing therefore requires

the network to jointly handle handover, session continuity, and edge-instance rebinding.

In previous work, we introduced SR4MEC [1], an architecture that relies on Segment Routing over IPv6 for the Mobile User Plane (SRv6 MUP) to provide flexible and scalable access to edge-hosted service instances. SR4MEC exposes the SRv6 domain as a single logical User Plane Function (UPF) to the Session Management Function (SMF), following a “one big UPF” abstraction. An SRv6 controller receives PFCP rules from the SMF, applies operator-defined edge policies, and configures only the ingress and egress nodes of the SRv6 domain. SR4MEC therefore adds an edge-policy abstraction and a scalable control-plane integration, but mobility-aware edge rebinding was left open.

This paper extends SR4MEC with mobility-aware edge rebinding. The proposed architecture uses handover events as triggers to update the SRv6 rules that bind an ongoing PDU session to an edge service instance. Depending on the operator policy and application requirements, the current instance can either be preserved during handover or immediately replaced by an instance located in the target area. These updates are performed without modifying the UE, without exposing instance identifiers to the application, and while remaining compatible with the 5G control plane.

The main contributions of this paper are the following:

- We extend SR4MEC with mobility-aware edge rebinding procedures that support both instance preservation and immediate instance switching during handover.
- We show how 5G mobility events and PFCP updates can be translated into SRv6 data-plane rule updates, while keeping the number of SMF interactions independent of the number of intermediate user-plane nodes.
- We implement and evaluate the proposed mechanisms in an open, container-based experimental platform with a mobility-capable RAN and 5G control-plane emulator, demonstrating seamless edge switching during handover.

The remainder of this paper is organized as follows. Section II discusses related work on edge-service access in 5G networks. Section III provides background on SRv6 MUP. Sections IV and V recall the SR4MEC architecture and introduce its mobility-aware extensions. Section VI presents the implementation and experimental evaluation. Finally, Section VII concludes the paper.

II. RELATED WORK

Existing approaches to accessing edge-service instances can be broadly classified according to where the binding decision is made. In application-level approaches, the UE is explicitly directed to a particular instance of the service. ETSI describes, for example, application-assisted mechanisms in which the application developer makes the instance identifier available to the UE, as well as DNS-based mechanisms in which the UE resolves a service name to an instance [2]. These approaches can be deployed without modifying the 5G user plane, but they leave the binding decision partly outside the control of the MNO. With application-based solutions, such as URL redirection, runtime and mobility-aware rebinding depend on application-specific logic. With DNS-based solutions, cached answers make fast rebinding difficult after a handover. As a result, these approaches provide limited support for transparent, operator-controlled edge switching.

The closest 5G-native alternative is traffic steering for MEC, where the MNO controls the user-plane path toward local data networks. In the standard 3GPP approach, the SMF can insert an Uplink Classifier (UL-CL) in the PDU-session path to steer selected traffic toward a local UPF, while the remaining traffic is forwarded to a central data network [3]. This approach preserves integration with the 5G control plane and avoids exposing instance identifiers to the UE. However, it relies on GTP-U tunnel chaining and PFCP rules distributed across the UPFs composing the path. Runtime updates, especially after inter-area handovers, may therefore require modifications at several UPFs, making the signaling cost depend on the number of nodes in the user-plane path.

III. BACKGROUND: SRV6 FOR THE MOBILE USER PLANE

SRv6 MUP defines how SRv6 can be used to steer traffic in mobile networks while preserving the 3GPP user-plane model [4]. SR4MEC relies on the *Enhanced Mode with Unchanged gNB GTP-U Behavior*, which allows existing gNBs to keep using GTP-U on the N3 interface while introducing an SRv6 domain behind an SR gateway (SRGW). On the uplink, the SRGW receives GTP-U packets from the gNB, decapsulates them, and maps the traffic into an SR policy by pushing an IPv6 header and a Segment Routing Header (SRH) containing the segment list to be followed inside the operator network. On the downlink, an SRv6 endpoint encapsulates packets toward the SRGW, which reconstructs the appropriate IPv4/UDP/GTP-U headers and forwards the traffic to the gNB. This mode is particularly suitable for edge computing because the gNB remains unchanged, while the operator can still use multiple SRv6 segments to steer packets through selected anchors or edge functions, and update these paths without installing per-session state on intermediate nodes.

IV. SR4MEC ARCHITECTURE

SR4MEC integrates an SRv6-based mobile user plane with the existing 5G control plane to provide flexible access to distributed edge-service instances [1]. Its main principle is to decouple the service identifier used by a UE from the

identifier and location of the concrete instance serving its traffic. Instance selection therefore remains under the control of the mobile network operator and can be changed without requiring the UE to discover a new address or establish a new PDU session. The architecture combines the *One Big UPF* abstraction, operator-defined *edge policies*, and an SRv6 controller that translates 5G session rules into forwarding rules at the boundaries of the SRv6 domain.

A. One Big UPF Abstraction

From the SMF perspective, the complete SRv6 domain is exposed as a single logical UPF through the standard PFCP interface. This logical UPF is implemented by an SRv6 controller and the forwarding nodes under its control. The controller terminates the PFCP messages sent by the SMF and configures the underlying SRv6 domain through a southbound interface. As shown in Fig. 1, the data plane comprises SRv6 gateways (SRGW) at the boundary with the RANs and SRv6 endpoints, or anchors, providing access to data networks.

Consequently, the SMF does not need to discover the internal topology or maintain PFCP state for every node crossed by a user-plane path. In contrast to UPF chaining, adding an intermediate router or network function does not require an additional PFCP session or a new pair of GTP-U tunnels. The internal path is instead represented by a segment list carried in the packet. This separation preserves compatibility with the 5G control plane while delegating fine-grained path computation and installation to the SRv6 controller.

B. Edge Policies

An edge policy expresses the operator's binding decision between a service and its execution location:

$$edgePolicy : Slice \times Area \times Service \rightarrow Edge \quad (1)$$

The *slice* is identified by its S-NSSAI, while an *area* represents a portion of the access network, typically defined as a set of gNBs associated with an SRGW. A *service* is identified by a stable name or network address known independently of its deployed instances. In the baseline model, at most one instance of a given service is deployed on each edge; selecting an edge therefore implicitly selects the corresponding instance. The policy can be extended to return an $(Edge, Instance)$ pair when several instances coexist at the same edge.

Policies are stored in an operator-controlled repository and may incorporate placement decisions, resource availability, application requirements, or load-balancing objectives. A default policy directs traffic that does not match an edge-specific rule toward a central data network.

C. Binding and SR Rule Installation

During PDU-session establishment or modification, the SMF sends to the logical UPF the usual uplink and downlink PFCP rules. The controller extracts the session context required to build the data paths, including the UE address, GTP-U tunnel identifiers, and the addresses of the relevant gNB and SRGW. It also determines the slice and current area




Fig. 1. Overview of SR4MEC architecture in a mobility context.

of the UE. For each edge-enabled service associated with the slice, the controller queries the edge-policy repository and resolves the selected edge to an SRv6 anchor.

The resulting binding is implemented using two match-action rules. An uplink rule is installed at the SRGW. It identifies traffic using the GTP-U context and service destination, decapsulates the incoming GTP-U packet, and applies an SR policy leading to the selected anchor. A downlink rule is installed at that anchor. It matches traffic destined for the UE and creates an SRv6 encapsulation containing the information required by the SRGW to reconstruct the appropriate GTP-U tunnel toward the gNB. Intermediate SRv6 nodes process the segment list without maintaining per-PDU-session state.

Following a handover, the controller can re-evaluate the edge policy and update the affected SR rules while preserving both the UE's service identifier and its PDU session. The mobility-specific mechanisms used to coordinate these updates are presented in the following sections.

V. MOBILITY MANAGEMENT

A. Edge Rebinding

An edge rebinding changes the service instance associated with a UE and redirects its traffic from a source edge to a target edge. It may be triggered by UE mobility, for example after a handover changes the UE's serving area, or by an independent event such as load balancing or an update of the operator's edge policy. The operation is performed by the SRv6 controller and does not require the UE to establish a new PDU session or learn the address of the target instance.

To preserve connectivity during the transition, the controller applies a make-before-break procedure involving three SR rules (illustrated in Fig. 2). First, it installs a new downlink rule at the SRv6 endpoint of the target edge. This rule establishes the return path from the target instance toward the UE through the appropriate SRGW and GTP-U tunnel. Once the target downlink path is available, the controller updates the uplink rule at the SRGW so that subsequent packets sent by the UE are steered toward the target edge. At this point, new requests reach the target instance, while the downlink rule at the source edge remains temporarily active.

The source edge may still hold packets in transit or responses generated for requests received before the uplink path was changed. The controller therefore removes the old downlink rule only after a short guard timer expires. This temporary overlap ensures that a valid return path exists




Fig. 2. SR rule updates during edge rebinding (make-before-break).

throughout the transition and prevents packets from being dropped because of premature state removal.

B. Inter-Area Handover with Post-Handover Edge Rebinding

In a first approach, mobility is handled by preserving the edge instance currently associated with the UE during the N2 handover procedure [5]. The objective is to remain fully aligned with the standard 5G handover sequence: the handover first updates the access-side data path so that the UE can move from the source gNB to the target gNB without changing its PDU session, and the edge rebinding decision is performed only after the handover has completed.

As shown in Fig. 1, we focus on inter-area mobility, where the source and target gNBs are attached to different SRGWs. From the point of view of the 5G control plane, each SRGW is exposed as a distinct GTP interface of the logical *One Big UPF*. Therefore, when the UE moves to a new area, the SMF updates the PFCP rules of the *One Big UPF* as it would do with a conventional access UPF. The SRv6 controller then translates these PFCP updates into SR rules deployed at the boundary of the SR domain. The main control-plane interactions and SR rule updates are illustrated in Fig. 3.




Fig. 3. Handover procedure with edge preserved.

During the preparation phase, the SMF creates a new PFCP context for the target area (P1) and configures the uplink path associated with the target gNB. The controller installs the corresponding uplink SR rule on the target SRGW (R1), while keeping the selected edge instance unchanged. In parallel, the standard handover procedure prepares the target gNB and allocates the GTP-U information required for the downlink

tunnel. Since the UE is not yet synchronized with the target gNB, downlink packets must still be forwarded toward the target area through the handover forwarding path. The SMF therefore configures temporary forwarding rules, which the controller maps to SR forwarding rules between the source and target SRGWs (R2).

Once the UE has synchronized with the target gNB and the target gNB sends the Handover Notify message, the handover enters its final stage. The SMF updates the downlink PFCP rule with the F-TEID of the target gNB (P4). The controller uses this information to update the downlink SR rule on the endpoint of the currently selected edge, so that packets sent by the preserved instance are now encapsulated toward the target SRGW and the target gNB. Finally, the PFCP state associated with the source area is removed (P5), and the controller deletes the obsolete SR rules on the source SRGW (R4).

At this point, the N2 handover is complete and the UE still accesses the same edge instance as before the mobility event. The controller can then re-evaluate the edge policy with the new area as input. If the policy now selects a different edge, the controller triggers the edge rebinding procedure described previously (cf. section V-A). This post-handover strategy is simple and robust, but it introduces a transient period during which the UE remains attached to an instance that may no longer be optimal for its new location.

C. Inter-Area Handover with Immediate Edge Rebinding

The post-handover approach separates mobility management and edge rebinding into two successive procedures. However, when the target area is known during the preparation phase, the controller can already evaluate the edge policy for the UE's future location. If the policy selects a different edge, and if the target instance is available, SR4MEC can combine the N2 handover and the edge rebinding into a single procedure, illustrated in Fig. 4.

The key idea is to prepare the path toward the target instance before the UE completes the handover. As in the previous case, the SMF first creates the PFCP context associated with the target area, and the controller installs the uplink SR rule on the target SRGW (R1). After the target gNB has allocated the F-TEID for the new downlink tunnel, the SMF configures the PFCP rules required by the handover forwarding mechanism (P2, P3), which are mapped to SR forwarding rules (R2). At this stage, the controller has enough information to create an SR rule (R3) on the endpoint of the target edge. This rule builds the SRv6 encapsulation required to send packets from the target instance toward the target SRGW and, ultimately, toward the GTP-U tunnel of the target gNB. As a result, the target instance can be used as soon as the UE is attached to the target gNB, without waiting for a separate edge rebinding procedure after the handover.

During the transition, the forwarding path associated with the handover still guarantees continuity for packets generated before the switch. The previous downlink rule on the source edge is therefore kept temporarily, so that in-flight packets can still reach the UE. After the handover has completed and




Fig. 4. Handover procedure with immediate edge rebinding.

a guard timer has expired, the controller removes the obsolete downlink rule on the source edge (R4), and the temporary forwarding rules associated with the source area (R5).

This immediate rebinding procedure has two advantages. First, it shortens the time during which the UE uses a non-optimal edge instance. Second, it reduces the number of SR rule updates compared with performing an ordinary handover followed by a separate edge rebinding.

D. Control-Plane Overhead

We now evaluate the control-plane signaling cost of the inter-area mobility scenarios described above. Table I compares a standard 3GPP UL-CL deployment with SR4MEC. The first scenario corresponds to the standard N2 handover case, where the edge binding is preserved. The two other scenarios, namely post-handover edge rebinding and immediate edge rebinding, are not specified by 3GPP. They are introduced here as UL-CL baselines that follow the same high-level logic as SR4MEC, in order to estimate the signaling cost that would be required to support equivalent edge-rebinding operations with a tunnel-based user plane.

For the UL-CL model, we distinguish the UPFs located at the extremities of the paths from the intermediate UPFs in the transport network. The extremities are the access UPFs, directly connected to the gNBs, and the PSA UPFs, directly connected to the data networks. These UPFs account for the constant terms in Table I. In addition, we denote by UPF_s , UPF_t , and UPF_i the intermediate UPFs on the source, target, and transition paths, respectively.

For UL-CL, the constant terms correspond to the mandatory PFCP operations of the N2 handover. In S_1 , five PFCP messages are required: creating the uplink and downlink rules on the target access UPF, installing a forwarding rule on the source access UPF, updating the downlink rule on the source PSA, and deleting the obsolete source access UPF. The remaining terms depend on the number of intermediate UPFs: removing the old path requires one PFCP deletion per UPF on the source path, while creating the new uplink and downlink paths requires two PFCP updates per UPF on the

TABLE I
CONTROL MESSAGES DURING INTER-AREA HANDOVER

Scenario	UL-CL PFCP	SR4MEC PFCP	SR
S ₁ : Edge binding preserved	$5 + 2 UPF_t + UPF_s $	5	4
S ₂ : Post-handover edge rebinding	$11 + 2 UPF_t + UPF_s + 3 UPF_i $	5	7
S ₃ : Immediate edge rebinding	$7 + 2 UPF_t + UPF_s $	5	5

target path. In S₂, an edge rebinding is performed after the handover. Compared with S₁, the UL-CL baseline requires six additional PFCP messages, and the transition path must also be configured and then removed, which introduces the additional dependency on $|UPF_i|$. In S₃, the edge rebinding is integrated directly into the handover. The UL-CL baseline thus avoids the full post-handover rebinding procedure, but still requires seven PFCP messages at the extremities. The variable part is limited to creating the target path and removing the source path.

With SR4MEC, the SMF always interacts with a single logical UPF. Therefore, the PFCP cost remains limited to the five messages of the inter-area N2 handover and does not depend on the number of intermediate SRv6 nodes. The additional operations are performed inside the SR domain: S₁ requires four SR rule updates to preserve the instance, S₂ requires seven updates because the edge is changed after the handover, and S₃ requires only five updates by directly configuring the target edge during the handover.

VI. EVALUATION

A. Implementation

We implemented SR4MEC in an open and reproducible container-based experimental platform built around open-source 5G components from multiple providers. This platform allows us to validate the integration of SR4MEC with existing 5G implementations, notably by testing its interoperability with free5GC core network [6] and UERANSIM [7] for the emulation of UEs and gNBs. Since N2 handover procedures are not currently supported by UERANSIM, we also developed and integrated NextMN-Lite [8]. NextMN-Lite is our lightweight emulator for the RAN and the 5G control plane, designed to reproduce mobility procedures while remaining interoperable with 5G user-plane components through standard N3 and N4 interfaces. This enables us to evaluate SR4MEC both in standard PDU session establishment scenarios and in controlled mobility scenarios involving inter-area handovers.

B. Experiments

We evaluate the impact of mobility-aware edge rebinding on the user-plane latency with stateless applications. All experiments are conducted on our container-based open experimental platform deployed on the Grid'5000 testbed [9].

As shown in Fig. 5, the evaluated service, denoted Δ , is deployed as three instances: δ_0 in the central data network, δ_1 and δ_2 in edge networks located respectively in Areas A and B. In our setup, a UE connected to a gNB in Area A has an




Fig. 5. Experimental topology for the mobility evaluation.

RTT of approximately 30–32 ms when accessing δ_1 , and a UE connected to a gNB in Area B observes a similar RTT when accessing δ_2 . When the UE is connected to Area B but continues to access the edge located in Area A, the RTT increases to approximately 35–36 ms. Accessing the central data network from either area results in an RTT of approximately 50 ms.

The experiment consists in continuously measuring the RTT from a UE initially connected to a gNB in Area A while it accesses service Δ . At $t = 10$ s, an inter-area N2 handover is triggered toward a gNB in Area B. We evaluate the two SR4MEC mobility procedures: post-handover edge rebinding and immediate edge rebinding. For each procedure, we consider two initial bindings: one where the UE is initially bound to the central instance δ_0 , and one where it is initially bound to the Area A edge instance δ_1 . We compare SR4MEC with a UL-CL baseline in which the handover preserves the initial binding and no edge rebinding is performed.

1) *Post-Handover Edge Rebinding*: In the post-handover case, the SR4MEC controller performs an edge rebinding after the handover has completed. In our experiment, this rebinding is triggered at $t = 13.5$ s, i.e., 3.5 s after the handover. The results are shown in Figs. 6a and 6b.

When the UE is initially bound to δ_0 , the RTT remains close to 50 ms before and immediately after the handover, since the UE continues to access the central data network. At $t = 13.5$ s, SR4MEC rebinds the service to δ_2 , the instance located in the target area. The RTT then drops to approximately 31 ms, confirming that traffic is redirected to the local edge in Area B. In contrast, the UL-CL baseline remains close to 50 ms because no edge rebinding is performed.

When the UE is initially bound to δ_1 , the RTT starts at approximately 32 ms while the UE is in Area A. After the handover, the binding is preserved and the UE continues to access the Area A edge from Area B. As a result, the RTT increases to 35–36 ms. At $t = 13.5$ s, SR4MEC rebinds the UE to δ_2 , and the RTT decreases back to approximately 32 ms. The UL-CL baseline remains around 36 ms after the handover, as it keeps using the initial edge instance.

2) *Immediate Edge Rebinding*: We then evaluate the immediate edge rebinding procedure, in which the target edge path is prepared during the N2 handover. In this case, the UE can use the new instance as soon as it is synchronized with the target gNB. The results are shown in Figs. 6c and 6d.




Fig. 6. RTT during inter-area handover (at $t = 10$ s) with post-handover and immediate edge rebinding.

When the UE is initially bound to δ_0 , SR4MEC switches directly from the central instance to the target edge instance during the handover. The RTT therefore drops at $t = 10$ s from approximately 50–51 ms to approximately 32–33 ms. This confirms that the edge rebinding is effective immediately after the UE attaches to the target gNB, without waiting for a separate post-handover procedure.

When the UE is initially bound to δ_1 , the RTT starts at approximately 31–32 ms in Area A. Around $t = 10$ s, a short latency spike is observed, which is caused by the N2 handover procedure and the transient indirect forwarding between the source and target gNBs. After this transient phase, the RTT remains close to 31–32 ms, showing that the UE is immediately served by δ_2 in Area B. In contrast, the UL-CL baseline remains around 36 ms after the handover because it preserves access to δ_1 .

Overall, these measurements confirm the benefit of integrating edge rebinding into mobility management. Post-handover rebinding eventually restores the optimal edge path, but only after an additional delay. Immediate edge rebinding removes this delay by combining the handover and the instance update in a single procedure. The UE therefore reaches the target-area edge as soon as the handover completes, while preserving session continuity.

VII. CONCLUSION

This paper presented a mobility-aware extension of SR4MEC for edge rebinding in 5G networks. Building on the *One Big UPF* abstraction, the SR4MEC controller trans-

lates PFCP updates generated during mobility procedures into SRv6 rule updates at the boundaries of the SRv6 domain. This enables the network to either preserve the current edge instance during handover, or immediately rebind the UE to an instance located in the target area, without modifying the UE, the application, or the 5G control plane.

Compared with a UL-CL-based architecture, SR4MEC avoids reconfiguring chains of GTP-U tunnels and keeps the SMF interacting with a single logical UPF. Our implementation and experiments show that post-handover rebinding restores access to the local edge after mobility, while immediate rebinding allows the UE to reach the target-area edge as soon as the handover completes. Future work includes coupling these procedures with application-state migration and evaluating larger mobility scenarios.

REFERENCES

- [1] L. Royer, E. Lavinal, and E. Chaput, “SR4MEC: Enabling Flexible and Scalable Edge Computing in 5G Networks,” in *IEEE ICC*, June 2025.
- [2] D. Sabella *et al.*, *Developing Software for Multi-Access Edge Computing*, ser. ETSI W. Papers. ETSI: MEC ISG, Feb. 2019, no. 20, 2nd ed.
- [3] S. Kekki *et al.*, *MEC in 5G networks*, ser. ETSI White Papers. ETSI: MEC ISG, Jun. 2018, no. 28, 1st ed.
- [4] S. Matsushima *et al.*, “Segment Routing over IPv6 for the Mobile User Plane,” RFC 9433, Jul. 2023.
- [5] 3GPP, “Procedures for the 5G System,” TS 23.502, Dec. 2025, v19.6.0.
- [6] free5GC Project. free5GC. [Online]. Available: <https://free5gc.org>
- [7] A. Güngör *et al.* (2024, Oct.) UERANSIM. [Online]. Available: <https://github.com/aligungr/UERANSIM>
- [8] L. Royer. Next Mobile Network open platform. NextMN and NextMN-Lite. [Online]. Available: <https://github.com/NextMN>
- [9] Grid’5000. (2024, Oct.) Grid’5000 testbed for experiment-driven research. [Online]. Available: <https://www.grid5000.fr/>