

Making Cellular Networks More Efficient By Roaming-in-Place

Tenzin Samten Ukyab¹, Lisa Suzuki², Demetrius Davis³, Zhihong Luo¹, Silvery Fu¹, Shaddi Hasan³,
Sylvia Ratnasamy¹, Scott Shenker^{1,4}

¹UC Berkeley, ²Waseda University, ³Virginia Tech, ⁴ICSI

Abstract

We propose Roaming-in-Place (RinP), a technique for dynamically sharing capacity across mobile network operators. RinP is a new form of infrastructure sharing that expands the traditional notion of roaming in cellular networks such that users may roam between operators with overlapping coverage areas based on load and performance conditions. Using simulation and small-scale experiments, we show that deploying RinP would allow operators to run their networks at higher utilization and provide users with higher availability and performance, while achieving 30-40% infrastructure savings in our typical evaluation scenarios. We present a design for RinP that can be incrementally deployed with modest changes to existing cellular infrastructure. We build a prototype RinP testbed, and show that our proposed design can be realized feasibly with modest changes to existing cellular infrastructure, requires no change to current protocol standards, and adds minimal latency overheads.

CCS Concepts

• **Networks** → **Network protocol design**; **Wireless access points, base stations and infrastructure**; *Network simulations*; *Network measurement*; *Network mobility*; **Mobile networks**.

Keywords

Roaming, Handover, Cellular Architecture

ACM Reference Format:

Tenzin Samten Ukyab¹, Lisa Suzuki², Demetrius Davis³, Zhihong Luo¹, Silvery Fu¹, Shaddi Hasan³, Sylvia Ratnasamy¹, Scott Shenker^{1,4}. 2025. Making Cellular Networks More Efficient By Roaming-in-Place. In *ACM SIGCOMM 2025 Conference (SIGCOMM '25)*, September 8–11, 2025, Coimbra, Portugal. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3718958.3750538>

1 Introduction

Roaming – in which an operator A directs its users to connect to the infrastructure of a different operator B – has long been a central feature of cellular networks, allowing A to ensure that its subscribers enjoy coverage even in areas where A lacks cellular infrastructure. In this paper, we explore the feasibility and benefits of generalizing roaming as follows: consider a scenario in which operators A and B both have towers (denoted T_a and T_b respectively) in the *same* vicinity but perhaps T_a is overloaded while T_b has spare capacity, or perhaps certain users in that vicinity enjoy better performance to T_b than to T_a . At such times, when a user U , who




Figure 1: Overview of RinP. HO: home operator. VO: visited operator.

is a subscriber of A , attempts to connect to T_a , operator A may (re)direct U to instead attach to T_b .

We call this generalization Roaming-in-Place (RinP) because the mechanism is essentially identical to roaming, only that it is not limited to scenarios where A has no deployments. In this paper, we present the design, implementation and evaluation of the RinP idea, demonstrating that RinP is practical: our design requires no modification to existing cellular standards, only modest changes to existing cellular core implementations, and can be incrementally deployed with immediate benefit to early adopters.

We argue that RinP is a valuable feature because it unlocks two important benefits for operators: (1) efficient capacity scaling, particularly in dense high-demand areas such as large cities, and (2) the ability to provide users with higher availability and performance service level objectives (SLOs). In what follows, we briefly expand on these benefits.

More efficient capacity scaling. The current cellular infrastructure has over 8B subscribers and carries over 190EB per month, yet bandwidth demand continues to grow exponentially with no end in sight [30]. The traditional approach to meeting this demand is via *densification*, where operators deploy more base stations per unit area, enabling better spatial reuse of spectrum. To date, this has been an effective – though quite expensive – approach with users enjoying data rates that typically double every two years, a scaling trend referred to as Cooper’s law [23].

Wireless experts have long warned us that Cooper’s law will eventually slacken because, beyond a certain point, denser deployments lead to increased interference and hence sub-linear capacity scaling. Unfortunately, these expert analyses reveal that we are now approaching this point of diminishing returns, a shift they liken to the end of Moore’s law and its impact on computing efficiency [14, 24, 46]. This raises the question: *Can we enable more efficient capacity scaling, particularly in dense high-demand areas such as large cities?*

To date, efforts to increase capacity through more spectrum and spectral efficiency have contributed a fairly small fraction of total growth [14]. Thus, we need other ways to scale capacity efficiently. A commonly suggested approach for building cost-effective infrastructure is sharing infrastructure across operators and various forms of passive and active sharing have been explored over the



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGCOMM '25, Coimbra, Portugal*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1524-2/2025/09

<https://doi.org/10.1145/3718958.3750538>

years. However, as we discuss in §2, these approaches often focus on challenges other than densification, or are less effective in dense high-demand areas where operators are unlikely to have excess resources to allocate to others on a long-term basis.

Taking a step back, we observe that previous efforts to share infrastructure were often intended to address the need for cost-efficient *coverage* (e.g., two operators agreeing to share infrastructure in remote areas where neither can individually provide coverage in a cost-effective manner). Another example of this are Mobile Virtual Network Operators (MVNOs) entities that pool multiple operators together to improve their user’s coverage. However, our focus is on sharing in high-demand areas where deployments are dense, and this places a very different set of requirements on sharing. In particular, to be effective in such contexts we need resource sharing that is *fine-grained*: i.e., we must assume that neither *A* nor *B* has significant unused resources to allocate over long timescales and instead we must be able to exploit smaller temporal variations so operators can avoid provisioning for peak (or close) demand. In short, we need statistical multiplexing at a finer granularity than is typically achieved by active sharing today than is typically achieved by MVNOs or active sharing today. For example, as discussed in §2.4, MVNOs typically leverage a small number (1-2) of MNOs with a single primary MNO and others as backup during lack of coverage conditions.

In this paper, we argue that (RinP) can help us deal with the problem of densification. We should make clear that RinP does not “solve” the problem of scaling cellular capacity, which remains a huge undertaking to cope with the growing needs. All we are claiming is that RinP helps offset some of the problems that arise with densification and the loss of linear scaling in capacity as more infrastructure is deployed.

Improved availability and performance. Even with well provisioned mobile networks, emerging applications are pushing our limits on both performance and availability. Applications, e.g., AR/VR, autonomous vehicles, and AI serving, demand ever-higher bandwidth and lower latency [11, 49, 56, 70]. Simultaneously, the growing adoption of cellular technology in the digitalization of traditional infrastructures (e.g., grid [71], water [33]), health, and critical services means that any downtime comes at a high cost, as evidenced by recent outages that affected 911 services [65, 66].

In datacenters, operators meet such demands by deploying redundant infrastructure (e.g., with dual top-of-rack switches for failover, overprovisioning to absorb demand spikes, and etc.), but doing the same for cellular infrastructure would be impractically expensive. This leads to the question: *Is there an efficient approach to improving the performance and availability SLOs that cellular operators offer their users?*

We propose RinP as one approach to meeting the above challenge. Using RinP, an operator can address outlier events - e.g., poor performance due to localized interference, a tower failure, load spikes - by opportunistically leveraging the infrastructure of a *different* operator. Prior work, such as [57], has shown that user driven selection of operators can improve performance. With RinP’s network-driven approach, however, an operator can also use insights from inside their network (e.g., tower load), as seen in Figure 1, to improve performance and availability for their users. RinP gives the operator detailed control over when to redirect a

subscriber to another operator, thus providing a fine-grained knob by which operators can manage performance anomalies and failure.

As with capacity scaling, we recognize that RinP does not improve the average case of performance and availability. It merely offers an efficient approach to improving the tail in performance or availability metrics. Similarly, RinP’s benefits are reduced if operators experience highly correlated conditions, though our results in §5.2 show that RinP achieves valuable benefits up to a significant degree of correlation.

RinP is an extreme form of the general notion of user sharing that has been proposed before [18, 48, 74]. However, to date, user sharing has been primarily of theoretical interest and typically studied from the viewpoint of improving coverage rather than as a solution to the challenges increasing user demand (§7). In addition, RinP is largely orthogonal to MVNO architectures and traditional forms of infrastructure sharing and is designed to address technical challenges of a more efficiently and plausibly deployable system. These challenges are:

- (1) **To be effective, it must enable fine-grained and flexible sharing:** RinP’s load-balancing is implemented per-user, per-tower and can vary in time. The decision to redirect a user is up to the operators involved and can be based on diverse factors such as performance, load, costs, etc.
- (2) **To be deployable, it must be minimally disruptive to the existing cellular infrastructure:** RinP reuses existing roaming mechanisms and as such is largely a shift in policy vs. a modification of cellular architecture or implementation. RinP could be used with existing phones with minor changes and network infrastructure with only over-the-top changes.
- (3) **To be deployable, it must be minimally disruptive to operational practice:** RinP can be realized as a simple extension to existing roaming agreements and does not require operators to cooperate in operating their infrastructure. Moreover, policy is not only fine-grained but need not be decided a priori, so operators can decouple building/operating their network from policies on when and where to offer RinP.
- (4) **To be future-proof, it must continue to provide benefits even at greater densities of deployment:** RinP’s efficiency gains *improve* with densification. This is because the demands from each operator at a given location are increasingly decorrelated as density increases, which in turn creates more opportunities for fine-grained sharing (§5.2).

In this paper, we review relevant background (§2) and then present the detailed design (§3) and implementation (§4) of RinP. We benchmark the overheads of our implementation (§5.1) and use simulation and small-scale experiments to evaluate the potential benefits of RinP (§5.1-§5.3).

2 Background

We briefly provide background on current cellular networks.

2.1 Cellular architecture

Cellular networks today are divided into the Radio Access Network (RAN) and the mobile core. The RAN consists of radio towers (called gNodeBs for 5G and eNodeBs for LTE) that connect to user equipment (UE), such as mobile phones. It manages the radio spectrum,

making sure it is used efficiently and meets the quality-of-service requirements of every user. The mobile core connects the RAN to the rest of the Internet and manages the databases storing user information (*e.g.*, service plans, usage, and *etc.*). The functions of the RAN, each component in the core network, and the communication between components and UEs, are defined in the 3rd Generation Partnership Project (3GPP) specifications. However, the implementation of each component is left to the operator's discretion. Instead of describing every part of the cellular core in detail, we focus on two components that are key to RinP: the Mobility Management Entity (MME) in LTE and the Access and Mobility Management Function (AMF) in 5G. The MME and AMF are responsible for managing UE mobility and authentication.

When a UE is turned on and wants to connect to a mobile network, it performs the *cell selection* procedure, which consists of two phases: *camping to the cell*, which involves tuning to the appropriate frequency and synchronizing with the cell's broadcast control channel and *attaching to the network*. Initially camping to a network requires scanning multiple frequency bands to identify nearby eNBs and available networks. The UE then chooses a tower to connect to using a 3GPP-specified procedure based on the home operator's Public Land Mobile Network (PLMN), or identifier, as specified in its Subscriber Identity Module (SIM) card and signal strength [7]. After camping, the UE "attaches" to the network [5]. This procedure includes authentication of the user and setting up of user plane and control plane connections, so the UE's traffic can be processed in the network. When a UE moves around in the physical world, it may need to "handover" to a cell that provides better signal quality. To do this, each cell shares information about the frequencies of adjacent "neighbor" cells to its attached UEs; a UE camped to a cell continually measures signal quality to these neighbor cells and reports this information back to the cell to which the UE is currently camped, which in turn uses these measurements to determine when to trigger a handover.

The RAN and core coordinate to handover the UE with minimal disruption to the connection. An eNodeB or gNodeB triggers a handover for a user when the UE measures a better signal quality from a neighboring cell. The adjacent cells can exchange user state via the appropriate interfaces (X2 for LTE or Xn for 5G), so the user can switch immediately to using the neighboring tower without having to go through the attachment process again. If those interfaces are not supported in the implementation being used, the MME/AMF must be involved to pass the user state through the connections between the tower and the core (S1 for LTE and NG for 5G) [3]. If the target neighboring cell is not connected to the same MME/AMF, the two MME/AMFs must communicate the user state through an additional interface (S10 interface between MMEs and the N14 reference point between AMFs).

2.2 Roaming today

When a UE enters an area where its home operator does not have any towers available and thus has no coverage, the UE attempts to "roam" by attaching onto an available tower. This roaming attachment only succeeds, however, if the visited operator has a roaming agreement with the UE's home operator. This roaming agreement

may come in the form of a bilateral agreement or they may be established through an IP Packet Exchange Provider (IPX-P). IPX-P's provide the infrastructure and functionality that enables roaming, such as building the data tunneling (GTP-U) between operators, and settling the billing between operators [61, 62].

Traditional roaming agreements, however, only occur in non-overlapping areas of coverage, such as different operators having coverage in different countries, and hence does not help operators combat the costs of scaling and densification in the areas they both cover. Domestic/national roaming agreements exist that enable roaming between operators within the same country, but they too are only enabled in areas of non-overlapping coverage [39]. RinP on the other hand assumes that the operators have RAN deployments in the same geographical region and both own their own separate spectrum and mobile core deployments.

The attachment procedure for roaming is largely the same as an attachment to a home operator tower, with three key changes. One change is during the tower selection procedure, the UE will preferentially select a tower to attempt to attach to based on an ordered list of roaming PLMNs listed in the SIM called the "Operator controlled PLMN Selector with Access Technology" [2]. As the name suggests, a UE's home operator configures this ordered list based on its roaming agreements and policy. Another change is that, in order to authenticate the UE, the visited operator core will communicate with the home operator core (*e.g.*, through, in LTE, the S6a interface [4] and, in 5G, the N32 reference point [41]). Finally, in order to enable billing of roaming traffic between operators the UE's data plane traffic may be routed differently. There are currently two options for routing roaming traffic: home routing and local breakout (LBO) [4]. Home routing is more widely deployed [29] and reroutes user traffic through the gateway of the home operator (through, in LTE, the S6 interface and in 5G, the N9 reference point [41]). In contrast, with LBO, users' traffic is routed locally through the visited operator's gateway, incurring lower latency and inter-operator traffic. This gain in efficiency comes at the cost of transparency: in home routing inter-operator migrations do not change the user's IP address, whereas LBO migrations do. RinP reuses the above procedures with no change.

2.3 Infrastructure Sharing

Various forms of infrastructure sharing have been explored in the literature, including:

(i) *Passive sharing*: in which operators share non-networking infrastructure, *e.g.*, buildings, towers, and power. This amortizes cost but does not fundamentally change the network capacity an operator must deploy to meet growing demands, and hence does not address the challenges of densification.

(ii) *Active sharing*: in which an operator *A* allocates a portion of its network infrastructure for use by an operator *B*. Active sharing can take many forms depending on the resources being shared (spectrum, radio, baseband, backhaul), who operates the allocated resources (*A*, *B*, or jointly), who owns the corresponding spectrum (*A* or *B*), etc. Regardless of the specifics, in practice the allocations are typically coarse-grained and static (*e.g.*, one operator allows another to use a fixed portion of their spectrum for a year or more).

This makes traditional active sharing less effective in dense high-demand areas where operators are unlikely to have excess resources to allocate to others on a long-term basis.

One might ask: can we extend active sharing to make fine-grained allocations? This is possible in principle, but in practice we believe this would be challenging and disruptive at both the technical and operational levels. On the technical front, it would require implementing fast and fine-grained resource allocation/revocation which in turn involves new mechanisms for virtualizing resources – such as the aforementioned radio, baseband, and backhaul – then providing mechanisms for scheduling those resources, enforcing isolation between those resources, and accounting for the use of those resources. It is worth noting here that even basic network/RAN slicing, which is a strict subset of what would be needed for fine-grained active sharing, has been many years in development and is still far from widely deployed [32, 63]. Perhaps more importantly, fine-grained active sharing would pose significant operational challenges, since it would require that operators cooperatively implement dynamic resource scheduling. To our knowledge, the only networks that have undertaken such cross-operator cooperative management involve subsidiaries that are under a common administrative entity [17, 69]. Finally, we note that, as currently implemented, neither active nor passive sharing helps with improving users' performance or availability SLO, which is one of the benefits we target with RinP.

2.4 Mobile virtual network operators (MVNOs)

In MVNO architectures – e.g., Google Fi [37] – a *virtual* operator (MVNO) acts as a broker between users and one or more infrastructure operators (MNOs). Most MVNOs today use one MNO; when an MVNO has multiple MNOs, it is for coverage in regions where the MNOs do not have overlapping coverage, rather than load-balancing or performance optimization. e.g., users with Fi can switch between networks automatically triggered by the loss of service¹. However this is not a part of the 3GPP standards and is only enabled for phones that have the proprietary Fi UE baseband that enables this [36], which restricts the adoption and efficacy of this architecture. In principle, an MVNO could load-balance users between MNOs in the dynamic and fine-grained manner that RinP advocates although it is unclear that MVNOs have as much incentive to do so since infrastructure efficiency is a more immediate concern for MNOs.

Architecturally, MVNOs may operate their own mobile core infrastructure (at least the control plane functions) but they do not operate their own RAN infrastructure. Rather they rely on MNO infrastructure to serve all their customers by leasing a portion of the MNO's RAN and leasing MNO spectrum. Conversely, RinP assumes that the MNOs using RinP own their own spectrum and have RAN deployments in the same geographical region. This difference is portrayed in Figure 2.

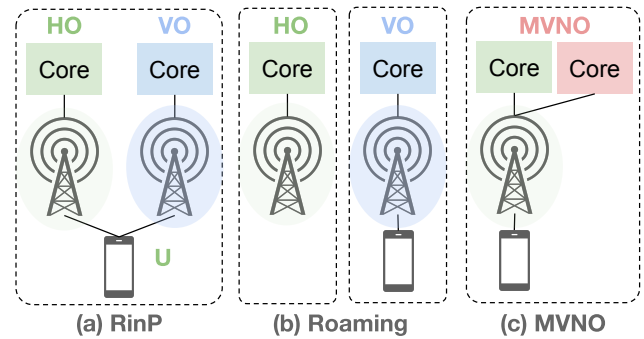


Figure 2: Comparing (a) RinP to (b) roaming today and (c) MVNO.

3 Design of Roaming-in-Place

3.1 Design Goals

We expect that operators will deploy RinP incrementally, and want the effectiveness of user sharing to improve with the degree of adoption. This requires our design to be easily adoptable with minimal changes to the operational practice (*i.e.*, resource management) and technical architecture (*i.e.*, current protocols and implementations) of operators.

In terms of operational practice, operators will be able to reuse their traditional roaming practices, such as international and national roaming agreements through IPX-Ps and roaming user prioritization, with RinP. Additionally, operators will not have to reinvent their architecture to enable active sharing, more specifically they do not have to implement dynamic and cross-operators/cooperative resource sharing.

Even within the scope of minimal changes, there is a distinction worth drawing between changes to (i) the internals of an operator's network (e.g., policy or software component) vs. (ii) the internal implementation of UEs vs. (iii) protocols (e.g., headers, signaling messages). The first only requires implementation by participating operators, the second by mobile device vendors, while the last requires a change to standards and hence broader adoption. While all three options are feasible, changing 3GPP standards is most difficult.

Therefore, we focus on approaches that avoid changes in 3GPP's specification and can be deployed with minimal coordination among participating operators. Put another way, we want to ensure (i) that from an operator perspective, RinP can be deployed with only internal configuration or policy changes to standards-based components or via the addition of over-the-top components² and (ii) that from a UE and network equipment provider perspective, RinP can be implemented with modest software changes, without changing the hardware and without violating existing specifications.

3.2 Overview

Based on the above goals, we design RinP such that operators can choose to offload users from a specific tower to another operator's

¹Fi no longer uses multiple networks and no longer supports this technology in new devices [59].

²By "over-the-top", we mean that the necessary components such as servers operate at the application level (above the cellular control and data plane), with message exchanges that are out-of-band (independent of any cellular protocol signaling).

network based on their own policies, such as high load on the tower, roaming pricing, etc. Operators communicate their willingness to accept roaming users to other operators through an over-the-top **readiness service**, described in §3.3. In addition, we need a procedure that takes a UE attached to one operator's network to roaming on another operator's network. We call this procedure a **roamover**, which we describe in more detail in §3.4, and provide an optimization for this procedure in §3.5.

Roaming agreements: Before using roamovers, participating operators must establish RinP roaming agreements. We assume that an operator *A* will only offload its users to (and accept RinP users from) operators with whom it has pre-established contractual agreements, either through a bilateral agreement or an IPX-P, similar to today's roaming agreements. We refer to these operators as *A*'s RinP "roaming partners". This decision allows *A* to (re)use standard authentication and settlement processes at the cost of limiting the roaming partners whose capacity *A*'s users can exploit.

3.3 Readiness Service

Before an operator *A* can redirect a user to another operator *B*'s tower, it must discover which (if any) of *B*'s towers are available to accept roamover users. To achieve this, a participating RinP operator deploys a *readiness service* that is made up of some number of readiness servers. The readiness server for an operator *A* gathers status information from (some subset of) towers in *A*'s infrastructure and, based on its internal policies, exposes the availability of its towers to other operators with which it has roaming agreements. The server also queries the readiness servers at *other* operators based on which it selects the target operator to which it can redirect its users.

Similar to how operators deploy a MME or AMF regionally for scalability and performance reasons, a readiness server can be deployed per geographic region, assigning towers in that region to it. This allows for the readiness server to have reasonable request latency and balances request load across readiness servers. Each tower periodically updates its load information to its associated server. If a tower is on the edge of two regions it can register to both servers. As trends toward virtualization of the RAN and core are actualized, it is not unreasonable to expect an operator to deploy another service alongside their existing regional services.

In each region, partnering operators exchange readiness information via pairwise exchanges. There are many options for what form this readiness information takes: e.g., a simple and most conservative YES/NO indicator for the operator to the most descriptive. The example of most descriptive is described in Appendix A, includes the cell information for the UE to optimize the radio frequency switch and QoS parameters to expose the available bandwidth on the cell. QoS implementations can vary by operator, so minimum expectation guidelines exist for supporting and enforcing QoS between operators during roaming [40, 41]. Roaming agreements generally include QoS support clauses, which operators would adhere to with RinP.

Each RinP participating operator exposes an API server to their roaming partners that their partners can query to discover the neighboring towers in the region and the readiness information of

those towers. Operators can find these API servers using standard techniques such as the DNS or other service discovery mechanisms.

The specifics of the implementation of the readiness server above is just one feasible option. Ultimately it is left up to the operators' discernment and agreements between partnering operators, enabling them to opt in and out of participating in RinP according to their business practices.

3.4 Roamover

We now step through the roamover procedure, depicted in Figure 3, in detail. Initially, the UE is attached to the tower of operator *A* (step 1). When the tower experiences a scenario where it wants to roamover a UE (e.g., during high load or a report of poor performance from the UE), it triggers a readiness request (step 2) and uses the readiness service to find a target operator *B* to roam to (steps 3-14). Having selected *B*, we still need a solution for how *A* will actually "roamover" the UE to *B*'s network. Our challenge is one of finding the right insertion strategy: *i.e.*, can we identify hooks/APIs in the existing standard that will allow us to achieve the behavior we desire. Our study of the 5G specification leads us to the following solution. Per the 5G specification, a UE that is not attached to its home network will look for another network to attach to by consulting a list of alternate networks, called the "equivalent PLMN list" [2], that it has stored locally. Thus, to roamover a UE, the operator *A* will first update the UE's equivalent PLMN list to contain operator *B*'s PLMN; this can be done with existing configuration APIs. Next, *A*'s core can use existing control commands to force the UE to detach from *A*'s network. The UE then detaches and by simply following the regular (re)attachment procedures will automatically select *B* as the new network it must attach to. Thus, to perform a roamover, the PLMN of the selected operator *B* is pushed to the UE (steps 14-17) and then the UE detaches from tower *A* (step 18, 19) and attaches to tower *B* (step 20). We now elaborate on these steps.

Selection: A tower might consider a roamover because it is experiencing high load and/or because its UE(s) are reporting poor performance. In such cases, the tower sends a request to its readiness server to get roaming partners it can offload UEs to. The readiness server queries the readiness servers of its partners to compile a list of towers that are near the UE for the tower. This can even be periodically gathered from partners to reduce request latency.

The tower will schedule a measurement report from the UE, passing on information about the frequencies at which the target towers operate (step 8); the UE's radio completes the requested measurements when there is no data to send or receive. Once the tower acquires the measurement report from the UE, it passes it to the readiness server (steps 9, 10). With the signal quality information (from the measurement reports) and the readiness information, the readiness server can now select an operator for the UE to roam to. The exact selection policy is up to each operator and can factor in additional information, such as roaming data prices. To avoid overloading a tower with available capacity, the roaming UEs' readiness server first requests the use of the cell from the readiness server of the target tower before forwarding the roamover decision to its UEs (steps 11, 12). The reserved cell from the chosen operator

must be the cell that provided the best signal quality from the measurement report of the UE, since the UE will pick that cell to camp on during the cell reselection process. As mentioned earlier, all communication between readiness servers is over-the-top.

The next steps required for roamer requires procedures (*i.e.*, detach) that must be initiated by the core according to the specification. Therefore, after the roamer decision is made, the readiness server communicates which UE and the PLMN to which it must roamer through the core (step 13).

EPLMN reordering: Next, the PLMN code preference list of the UE must be reordered, so that it will automatically pick the desired target operator when reattaching. This can be accomplished by having A's tower overwrite the equivalent PLMN (EPLMN) list [2] on the UE. Per the 3GPP UE specification, the EPLMN list is considered for attachment after the home PLMN and before any of the PLMNs stored on the SIM, and before attempting attachment to PLMNs detected in a scan.

In the 5G protocol, the PLMN reorder can be triggered through the UE reconfiguration update procedure [5]. The AMF initiates this procedure and indicates to the UE that it needs to register again (steps 15, 16). In the final confirmation message of the procedure, the EPLMN list can be passed to the UE and stored for future use (step 17).³

Detach: To get the UE to switch to another operator, the UE must detach first. The detach (for LTE [3] and deregistration for 5G [5]) procedure is initiated by the core (steps 18, 19); specifically, by the MME (for LTE) or the AMF (for 5G). To get the UE to attach after the detach, the detach message must include a reject code (*e.g.*, Cause #15: no suitable cells in tracking area [6]) that triggers a PLMN search.

Detaching and reattaching a user will lead to a disconnection of the user from any network for some time. Prior to attachment, the UE will usually scan a range of frequencies to discover nearby towers and measure their signal quality to make an attachment decision. This scanning process is usually done when the phone is first switched on or there is a total loss of coverage and can be slow. However, vendors can implement optimizations in the scanning process that can reduce this time to a reasonable number. For example, triggering measurement reports for cells belonging to the other operator allows the UE to learn which cells and frequencies are in use by that operator, avoiding a full scan and drastically reducing the time that the UE is disconnected [7]. With RinP, this information is certain to be in cache because of the measurement report requests made in the selection process (steps 8, 9). Such optimizations are not required by the 3GPP specification, but optional for vendor implementation [9].

The above roamer process incurs additional overheads relative to traditional handovers, but is also performed far less often. We evaluate the overheads of roamers in §5.1.

Attach: The target operator attachment procedure in RinP is the same as the roaming attachment procedure. When setting up the data plane connections, the home operator can reuse the IP

³The EPLMN list is soft state on the UE, so it gets refreshed after every registration procedure. Therefore, after attaching to the selected roaming tower, the most preferred PLMN will return to being the home PLMN. Therefore, when the UE performs the PLMN selection process after detaching from the visited tower, it will return to its home tower.

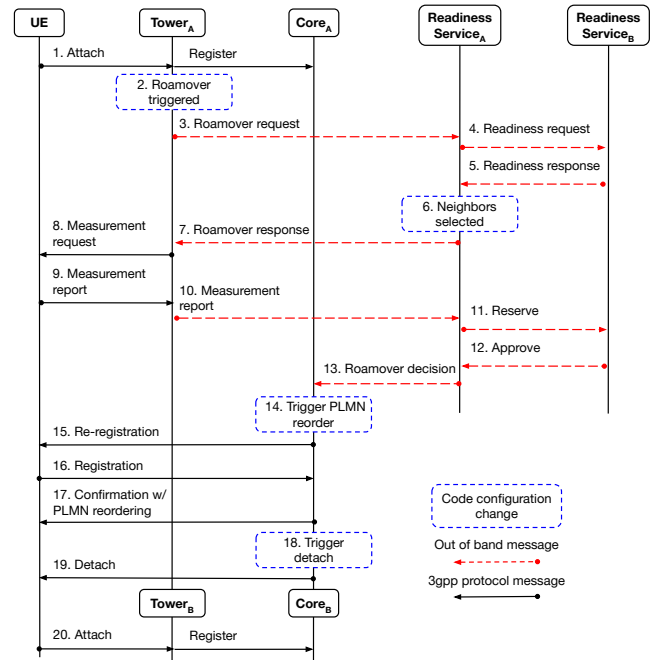


Figure 3: An overview of the roamer procedure where an operator A queries operator B for its readiness and reserves resources from operator B.

and Internet facing gateway of the previous connection to preserve data plane seamlessness. This is assuming the norm HR is used, as described in §2. With HR, switching between operators does not change the user's IP address and hence RinP will not disrupt the user's flows. However, with LBO, opposite true, disrupting connections based on traditional transport protocols like TCP. LBO provides a shorter path for the UE to an Internet gateway, which is valuable in international roaming where the home operator's gateway is far. For RinP, HR is less of a concern as both operators will have gateways in the same area.

Roamover when visiting: When a UE is attached to a visited tower, it can remain in the visited network and handover within it. In the 3GPP specification today, the UE checks the availability of its home operator at some configurable period t [2]. In the RinP scenario where the home operator has coverage when the UE is roaming, the UE will be pushed back to the home network which then can make decide to roam based on load again. If the visited tower wants to offload a visiting UE, it can perform the full RinP procedure to roamover the UE to its home or other operator.

Minor Optimizations: The protocol above describes a scenario where all necessary steps are run *on demand*. Instead, certain steps can be performed in advance to speed up the roamer. Such as pushing the PLMN reordering in the initial registration of the UE or periodically perform the measurement reports of neighboring towers to have them cached. Section 5.1 will evaluate the on-demand version of the roamer protocol which is the pessimistic case.

3.5 S10/N14 Optimization

As described in §2, today, towers coordinate with other towers, so handovers are "seamless" in that the disruption when the UE migrates between towers is minimal. Ideally, we would like RinP roamers to achieve low disruption migrations akin to handovers. Interestingly, examination of the 3GPP standards reveals that handover mechanisms to migrate a user across *different* mobile core components do exist in the standard, and hence RinP could exploit these to achieve a more seamless roamer compared to the cell re-selection approach described in §3.4. For example, the S10 interface in LTE facilitates handovers between eNodeBs connected to two distinct MMEs, potentially running in two separate core networks. The N14 reference point in 5G does the same by connecting AMFs. By implementing the S10/N14 interface in addition to the standard roaming interfaces (e.g., S6a and S8), roaming partners could facilitate RinP via inter-PLMN handovers without having to modify the standards.

Unfortunately, the S10/N14 interface is not implemented in the open-source systems available to us and, based on our discussions with vendors and operators, does not appear to be consistently implemented or used in practice, which is why we view this feature as an optional optimization in our design. However, since it is defined in 3GPP standards and has been implemented between operators before [22], it is feasible for operators to broadly adopt it to support RinP and reap its benefits. Operators can build up their own infrastructure to support the S10/N14 interface, or IPX-Ps can offer them in addition to the other services they offer.

If available, the roamer procedure would change in a few ways (see Figure 10 for a detailed view). The selection procedure would remain the same (steps 1-12 in Figure 3). The primary change to our design would be that, instead of asking the UE to detach and then re-attach, we directly trigger a handover via the S10/N14 interface to the target operator. For the target core operator to accept the source operator's handover request, the readiness service pre-loads information used to identify the UE (such as IMSI) in the target operator's core. The source operator passes all the information for the UE to be able to attach, however it does not pass along the UE security context, to protect its authentication keys. Therefore, the target operator's core needs to be modified to accept the UE transfer request. The request format remains protocol-compliant but simply holds temporary UE context fields. After the generic handover protocol specified in the standard, the target operator re-registers the UE to set up the UE context securely. As described in §5.1, we added support for the above to existing open source cellular systems and evaluated its performance.

3.6 Discussion

More user-centric designs: Today, handover target selection is network-driven; i.e., made by the operator of the user's current network. Our RinP design retains this model of network-driven selection and the corresponding control it gives operators. An alternative might have been to allow operator selection to be user-driven, so the UE makes this selection decision. This requires the UE to query the readiness information of roaming partners itself, which can be done over-the-top. But it will also need to be able to initiate

frequency measurements and migration from one operator to another, which is a requirement of UE baseband that is not common or specified in the standards. For these reasons, we opted to retain today's operator-centric approach in designing RinP.

Security and Privacy: Using a cellular network today, comes with users exposing themselves to privacy and security vulnerabilities, such as location tracking. RinP may be viewed as exposing a user to location tracking by additional operators. However, RinP does employ the existing roaming security and privacy mechanisms between operators, such as the use of Internet Protocol Security (IPSec) with Internet Key Exchange (IKE) in LTE and Security Edge Protection Protocol (SEPP) in 5G at the IPX-Ps [40, 41], which protect the operator networks and inter-operator connections from attack. Since RinP targets domestic roaming agreements, security and privacy guidelines, such as firewalling, that are not fully adhered to in the international sphere can be enforced by regulation [73].

Roamer policy: Since roamer, as regular handover, is network-driven, the operators have control over when the roamers occur. Additionally, operators participating in a roamer between them, preapprove the roamer before it proceeds, as described above. This luckily avoids the problems that naturally arise in a roamer mechanism that is user-driven or initiated, such as ping-pong roamers or high roamer frequencies, that can reduce the efficiency of the network. The burden of devising and implementing a roamer policy that avoids such problems and optimizes efficiency is put on the operators and left to future work.

4 Implementation

We implemented a prototype of RinP as depicted in Figure 4. It includes four components: the UE, the tower, the cellular core, and the RinP readiness server. The UE is a dual-SIM 5G capable "Oneplus Nord CE 2 Lite 5G" Android device [45] with Gailer programmable SIMs [34]. An x86 laptop running Ubuntu is connected to two USRP B210 [31] SDR devices that act as our gNBs. The laptop runs the open-source RAN program srsRAN [72] and the open-source 5G and roaming capable mobile core Open5GS [52] locally. We extend both srsRAN and Open5GS to implement RinP.

As discussed in §3.2 operators can choose to implement their readiness servers to their discretion. For simplicity, we choose to implement the service as a single centralized server. The server is implemented using gRPC [38] and can be run locally and as an AWS instance. The policies for advertising readiness information are flexible and up to operators. For our prototype, we implement a simplified policy as a proof of concept. A tower queries the server providing its own PLMN and the server serves the tower the PLMN and operating frequency of the roaming partner. Additionally, the readiness server forwards the roamer decision provided by the tower to the appropriate mobile core. Since there is only one tower for each operator in our prototype, we do not mockup the location filtering needed to differentiate neighboring towers.

4.1 Code Modifications

Implementing RinP does not require modifications to standard 3GPP protocols. However, it does require a few changes to the internal implementation of the towers and core, which is the srsRAN and



Figure 4: Our testbed uses an x86 laptop with two b210 SDRs acting as gNodeBs for two separate PLMNs and an Android phone with a programmable SIM.

Code	Original Line Count	Lines Added
Open5GS	1,366,453	4,824
srsRAN_Project	931,885	499
Total	2,298,338	5,323

Table 1: Lines of code added to implement RinP.

Open5GS code bases in our prototype, respectively. Our implementation of RinP only required 5,602 LoC changes as shown in Table 1. Our code bases are research-style implementations that are reduced in size and complexity from vendor implementations, but are 3GPP standards compliant and work with unmodified UEs. We also added the S10 interface implementation into the MME implementation Open5GS, extending the S1 handover procedure to send a *Forward Relocation Request* to a second MME instead of sending a *Handover Command* back to the eNB.

5 Evaluation

We address three questions: (i) Does RinP introduce significant new system overheads? We evaluate this using our prototype testbed in §5.1; (ii) What efficiency benefits can RinP provide? We evaluate this question via simulation in §5.2 and, (iii) What performance improvements might users enjoy with RinP? We evaluate this via small-scale experiments on real-world cellular networks in §5.3. We acknowledge that these questions are best answered via large-scale experiments and traces from production networks; our methodology reflects our best attempt given that we lack access to such networks.

Our work raises no ethical concerns.

5.1 RinP Prototype Evaluation

We evaluate the latency overhead during roamovers and how this overhead compares to traditional S1 handovers.

5.1.1 Baselines. We compare RinP to a baseline of S1/NG handover (described in 3.6). The standard handover represents a lower bound for the overhead/disruption that a UE experiences when switching between towers, as handovers (unlike roamovers) do not require the UE to deregister and register on the new tower.

RinP set up: We emulate two operators, each consisting of a single srsRAN [72] gNB and an Open5GS [52] core network running on a single x86 laptop. Each gNB is configured to run on a different NR ARFCN and each operator has unique PLMNs (001-01 and 999-70). An Android UE is installed with a SIM configured for the 001-01 PLMN. The two Open5GS cores are set up with a 5G roaming connection using a Security Edge Protection Proxy (SEPP) connection.

Handover set up: To evaluate handover, we configure the PLMN 001-01 operator with one Open5GS LTE core and two srsLTE eNBs.⁴ Each eNB operates on a different EARFCN and the UE is installed with a SIM registered to the PLMN 001-01. We trigger S1 handovers via the eNB’s console.




Figure 5: Breakdown of roamover latency in our testbed. Cell Reselection time measured as 40–60 seconds but can be optimized to 0.08 seconds (§5.1.2).

5.1.2 Overhead Breakdown. Figure 5 breaks down the roamover steps from switching from operator 001-01 to operator 999-70. There are 3 steps: (i) detach from operator 001-01, (ii) cell reselection, (iii) attach to operator 999-70. Each of those steps are further broken down in Figure 5 from when the connection is interrupted from Radio Resource Control (RRC) Release to the connection being re-established in Protocol (PDU) Establishment.

Cell reselection requires the UE to scan its supported bands to perform the PLMN search and cell selection logic. In our testbed implementation, we measured the time required for cell reselection to be 40–60 seconds of frequency scanning when we used the One-Plus UE. Clearly, many applications would not tolerate connection disruptions of this magnitude. Unfortunately, the closed nature of UE baseband hardware prevents us from diagnosing this overhead directly but, based on discussions with UE vendors and as discussed below, we believe this overhead is not fundamental and can be reduced with changes to UE baseband implementations.

⁴Open5GS includes both a 5G core and LTE core implementation; however, roaming is only implemented in the project’s 5G core, and handover is only implemented in their LTE core, hence we use LTE for handover tests.

5.1.3 UE Changes. One conjecture for the high latency is that our UE’s baseband implementation scans the entire list of roaming PLMNs, rather than exiting the scanning process after finding a high priority PLMN with good enough signal quality.⁵ Modifying this behavior to exit early would be a simple change to UE hardware. Alternatively (or in addition), reselection time could be shortened if the baseband implementation used cached measurements of cells supporting the PLMN of roaming partners. More generally, we can use the time required for a handover to approximate a lower bound for reselection time. This is because, in principle, the procedures performed in the access stratum for a handover and a cell reselection in which a UE had advance knowledge of target gNB radio parameters, are similar. Thus, the difference in time required for a detach-attach style roamover should be the latency required for a handover plus the time to perform the detach and attach procedures.

We measured the time it takes for a UE to transition from one eNB to another to be ~ 0.08 seconds, which is the value used as a placeholder in Fig. 5. This time was inferred from the S1 handover procedure by measuring the time between the *Handover Command* and *Handover Confirm* messages. The overall UE cell reselection time may take longer in roamover, since the phone must do the PLMN search and cell search procedure. However, we can use this as a lower bound and find that roamover time can be under 1 second on our testbed. Prior work [57, 60] has shown that overheads of this magnitude have minimal impact on application-level performance. This method, however, requires all UE device vendors to implement their basebands with the cached measurements, if the RinP were enabled.

5.1.4 S10/N14 Optimization. Another approach is to reuse the existing handover standards alongside RinP, as described in §3.5, to achieve the lower bound above. The difference between the S1 handover procedure and S10 handover procedure is only apparent in the preparation phase before the transition occurs to migrate the UE state between the source and target core components, such as the MME and SGW (or AMF and UPF in 5G), and after to clean up leftover UE state. Therefore, the connection interrupt time is ~ 0.08 seconds, which is the same as measured for the S1 handover above.

5.1.5 Signaling Cost. We find that roamovers, even optimized for interruption time, incur a bit over 2x the signaling load of handovers (see Appendix C). However, we note that roamovers are performed far more infrequently than handovers (for a small number of users, only under load spikes or performance anomalies), and hence we believe this tradeoff can be made without overloading the internals of the operator network.

5.2 The Efficiency Benefits of RinP

As mentioned, RinP is motivated in part by the challenge of densification. We briefly elaborate on this challenge and then present simulation results showing that RinP can provide network efficiency savings. We study (i) what savings RinP can achieve and (ii) the key factors that impact these savings.

5.2.1 Densification. Densification, where more towers are deployed per unit area, has driven the exponential growth of cellular capacity

⁵This is further supported by the fact that 40-60 seconds is a fairly typical time for a “cold scan”.

Parameter	Default Value
Correlation coefficient	1.0
Coefficient of variation	0.6 [75, 78]
User traffic distribution model	Log-normal [75, 78]
Path-loss exponent	4
Interference sensitivity	1.0
Tower placement	u.a.r, non-overlapping
User placement	u.a.r.

Table 2: Default simulation parameters.

for decades [14, 76]. When towers are far apart, cellular capacity scales *linearly* with the density of towers [13], which makes densification a cost-effective means for operators to meet growing demand. Densification has limits, however, and as the density increases, capacity enters a *sublinear* scaling regime with respect to infrastructure density [12, 14, 15, 26, 27, 58, 77] due to the complexity of managing interference from adjacent cells. Because of this, there is a need for new approaches that allow operators to meet future demand in a cost-effective manner. We believe RinP through “user sharing” can be one such approach.

5.2.2 Simulation Model. Our simulation models two operators that have users and cellular towers within the same area. Each operator participates in RinP but prioritizes serving its own users’ demand over that of the other operator [18].

The savings from RinP depend on a few key parameters that we explain below and summarize in Table 2. For our analysis, we pick default values for these parameters to match those from prior work. Where there is a range, we pick default values that are pessimistic in estimating RinP’s benefits.

User placement and demand. Each providers’ aggregate demand is modeled as a log-normal distribution with a coefficient of variation of 0.6 (as per [75, 78]). The sampled aggregate demands then determine each provider’s number of users who all have equal demand and are placed uniformly at random within a 2D grid area of size 1000m by 1000m.

User traffic demands in a local area may be correlated due to similar behaviors across the user populations of each operator. Intuitively, we expect that the more correlated the user demands, the lower the savings due to RinP since there will be fewer opportunities for offloading excess demand across operators. We thus correlate the operators’ demand distributions with a Pearson correlation coefficient x (in keeping with [78]) and experiment with x in the range [0.3, 1.0] where higher values represent higher correlation. We pessimistically set $x = 1.0$ as our default value, so both operators have an equal number of users and aggregate demand.

Tower placement. The main goal of operator provisioning is to deploy the minimum level of capacity C needed to meet a target Service Level Objective (SLO), where the SLO is typically defined in terms of the degree to which user demands are satisfied. (We define our SLO targets more precisely later in this section.) Given a user demand, our simulation adds towers into the grid until the operator’s target SLO is met. We assume each tower has a fixed capacity and hence the total capacity C deployed by an operator

is the number of towers times the per-tower capacity. We use C as our measure of the operator’s infrastructure costs.

The above process requires us to model how users are assigned to towers which in turn depends on how much of a tower’s capacity is available to meet a user’s demand given the impact of interference and noise. For a given user u and tower t , we model the downlink Signal to Interference and Noise Ratio (SINR) as in prior work: $\text{SINR}(u, t) = \frac{P}{I + sN}$ where P is the power of the signal from t , I is the sum of the interfering signals of other towers, N is such that the SNR at 100m is 20dB at all locations, and lastly we included a parameter s to dial the sensitivity the simulation has to interference which ranges from 0.0 to 1.0. P is modeled by the path-loss function $l(x) = x^{-\alpha}$ where x is the distance between u and the t and α is the decaying factor. Thus the capacity that a tower t has available for the user u is: $c(u, t) = B * \log_2(1 + \text{SINR}(u, t))$ where B is the bandwidth in hertz available at t . We set the various interference parameters to match prior work, as summarized in Table 2.

Given the above, we then assign users to towers as follows. For a user u , we first consider any tower that belongs to u ’s home operator and that has sufficient capacity to meet u ’s demand. When many such towers are available, we assign u to the tower where it consumes the least bandwidth (*i.e.*, to which it has the highest SINR). We do this for all users, reflecting our policy that operators prioritize their own users over those of their RinP peers. Up to this point, our model captures how users are serviced today in current networks. To model RinP, we include an additional step in which we check whether users that were not satisfied in the previous step (*i.e.*, by their home network) can be serviced by a tower belonging to its RinP peer using the same algorithm as above, but applied to towers in the other carrier.

Metrics. To measure the *savings* due to RinP, we simulate an instance of time reflecting current operation – in which each user’s demand can only be met by towers belonging to its home operator – and measure the total capacity needed as C_{today} . We then repeat the simulation under RinP in which users can be serviced by any tower and derive C_{RinP} . The savings due to RinP is thus measured as: $\frac{C_{today} - C_{RinP}}{C_{today}}$. We expect this to directly correlate to the CapEx and OpEx savings that operators will achieve with RinP.




Figure 6: Savings vs SLO target (default setup).

5.2.3 *Results.* We present results using the above.

(1) Baseline savings and the impact of SLO thresholds. We start by considering the savings under our default setup. As mentioned

above, the potential savings with RinP depends greatly on how much capacity the operator provisions which in turn depends on their target SLO.

Generally, we expect RinP to offer greater savings with more demanding SLOs; *i.e.*, SLOs that require the operator to provision closer to peak demand, or over-provision for the average demand, without RinP. This over-provisioning of the average demands creates more opportunities for fine-grained sharing to exploit since there is a higher chance that there is leftover capacity at any given time. While operators do not formally publish their SLOs, the information we encountered in discussions with industry experts was that operators typically provision to “meet over D% of user demand, at least T% of the time” with both D and T in the “high nineties”.

Our results in Figure 6 explore this regime. We plot the savings (y-axis) for values of D ranging from 0.8 to 0.98 (x-axis) under two scenarios for T. In the first (labeled avg), D% of user demand must be satisfied on average; in the second it must be satisfied at least 90% of the time (90%ile).

We see that: (1) RinP offers valuable cost savings – *e.g.*, at $D = 0.975$ we see between 30-40% savings depending on our SLO definition, and (2) as expected, savings increase with D and are greater with the more demanding 90%ile SLO target. Going forward, we use the avg SLO definition, reflecting a more conservative estimate of RinP savings.

One may wonder how these savings compare to the “roaming” fees that an operator must pay its peer operators. In all of our simulations, the actual volume of traffic that an operator offloads is a small fraction of its total demand – *e.g.*, in our simulations above the operator offloads a *maximum* of only 4% of its total traffic demand on average. Hence, for any reasonable pricing model, we expect that any loss in revenue to an operator (due to roaming settlements) is dwarfed by the savings in infrastructure costs. Moreover, depending on their relative provisioning levels, there may be no revenue loss since the operator reciprocally receives additional roaming revenue when it services the users of its RinP peers.

(2) Impact of user correlation. Next, we repeat the above simulation while varying the degree of correlation. Figure 7 shows savings for three Pearson correlation coefficient values 0.3, 0.6, and 1.0 reflecting (respectively) weak, medium and highly correlated demands across operators. As expected, savings improve as correlation decreases: *e.g.*, at $D = 0.975$, savings increase from 30 to 40% (a third increase in savings) as the correlation coefficient drops from 1.0 to 0.3. With lower correlation the chances that the towers of one operator are experiencing high load and the towers of the other are experiencing low load are higher, which means more opportunities for RinP to exploit and higher savings with RinP. Note that we see gains even with a 1.0 correlation coefficient due to randomness in individual user and tower placements.

(3) Sensitivity analysis. Here we present a sensitivity analysis on other parameters, further is in Appendix D.

(ii) *Demand distributions.* We experimented with distributions such as Exponential, Pareto, and Weibull and found no significant differences in the conclusions drawn – Figure 8.

(iii) *Tower placement.* An interesting effect we noticed was the impact of tower placement across the two operators; savings decrease with tower placements that are more correlated, with the baseline assumption that the user demand on each operator is correlated,




Figure 7: The impact of user demands on RinP's savings.




Figure 8: The impact of user demand distribution on RinP's savings.




Figure 9: The impact of tower overlap on RinP's savings.

Scenario	RSRQ gain %	Bandwidth gain %
Mobile in Tokyo (KDDI / SB)	10.98 / 9.83	25.62 / 55.01
Mobile in Berkeley (FN / VZ)	10.37 / 8.50	0.12 / 722.32
Stationary in Tokyo (KDDI / SB)	7.55 / 5.94	3.53 / 85.51
Stationary in Berkeley (FN / VZ)	25.96 / 0.96	0.03 / 254.28

Table 3: Potential perf. gains in real-world tests

as seen in Figure 9. At the extreme, where both operators place towers in identical locations, savings are reduced by as much as 50%. This is because the diversity of tower options decreases with increasingly correlated placement. When the towers of two operators are in overlapping locations, offloaded users from one operator will flood to the tower of the other in the same location, which may not be able to service all the offloaded users. With a higher diversity of tower location, the offloaded users can be divided into different towers, which have a higher chance of being able to service the users. We believe this has interesting implications for how operators select their RinP peers for maximum benefit; e.g., that an operator might gain more by peering with operators that they do not share infrastructure with.

5.3 Real-World Measurement

To evaluate RinP's potential to improve user performance today, we conducted measurement studies of the existing cellular networks in two cities: Tokyo (a densely populated city in Asia) and Berkeley (a mid-sized city in the US). In both cities, we used a pair of Google Pixel 8a phones connected to two major mobile operators in each country – denoted *KDDI* and *SoftBank (SB)* in Tokyo and *FirstNet (FN)*, *Verizon (VZ)* in Berkeley. We consider mobile and stationary scenarios each of 2 hours duration. The mobile scenario in Tokyo includes riding on multiple trains while the Berkeley it includes walking.

Each device collected *passive* and *active* observations of network performance. For the former, we use Android API calls (*Telephony-Manager*) to collect signal measurements every second. We focus our analysis on RSRQ since this is the metric typically used to trigger handovers [8]. For the latter, we use the Network Diagnostic Tool NDT7 (which uses TCP BBR, with Cubic as fallback) to measure network bandwidth between a client and an NDT server.

Table 3 summarizes the gains in RSRQ and bandwidth for each city and each operator. Gains are calculated by averaging the improvement in RSRQ or bandwidth that a user belonging to (for example) *K* would see if they were allowed to roamover to *S*. (Detailed timeseries of our measurements are shown in Appendix E.)

Overall, we see that there are non-trivial to significant opportunities for performance improvement that RinP can exploit. In some cases, such as for the mobile user in Tokyo, we see a larger number of opportunities for roamovers (because of the higher speed of the moving train and the high density of tower deployments along railway lines in Tokyo), resulting in significant gains (25-55% higher bandwidth). By contrast, the significant improvement for customers of *V* in Berkeley stem from the seemingly better provisioning and coverage levels of *F* (compared to *V*), rather than dynamically varying conditions. Regardless of their underlying cause, RinP's techniques can exploit these opportunities for improved performance, subject, of course, to operator policies.

On one hand, this analysis represents an upper bound for RinP, with real world gains being a function in large part of operator policy around roamover triggers. At the same time, this estimates potential gains for a UE in "normal" circumstances, as we have no reason to believe we observed poor signal quality from either operator during our mobile test; we imagine RinP would provide larger benefits to *tail* performance for UEs experiencing interference or congestion. Nevertheless, this limited measurement study demonstrates that current commercial mobile networks present opportunities for improved user performance through roamovers.

6 Summary of Findings

We evaluated both the feasibility and benefits of RinP. We showed that implementing RinP required the addition of less than 0.2% LoC to existing open source cellular systems. Our prototype evaluation in §5.1 identified two modifications that vendors should implement to achieve roamovers with latency comparable to existing handovers. The first involves a simple modification to the UE baseband to avoid unnecessary scanning of alternate PLMNs and/or to cache the results of previous scans. The second requires implementing the 3GPP S10/N14 interface (a 50 LoC change in our system). These changes are orthogonal: they can be pursued independently (the former by UE baseband manufacturers and the latter by cellular core vendors) and either would yield the performance improvements we seek. For example, our evaluation showed that with simple UE changes, the overhead of roamovers can be under 1 second. Importantly, neither modification involves a change to 3GPP standards.

In §5.2, we showed via simulation that RinP can improve network efficiency by 30-40%. Finally, in §5.3, we demonstrated that RinP can exploit performance variation that exist in real-world networks to yield valuable improvements in signal quality and channel bandwidth. We view our results as an encouraging indicator that RinP is both practical and beneficial, yet also recognize that industry adoption and deployment would be required to validate our results at scale.

7 Related Work

Infrastructure sharing in cellular networks: As discussed in §2.3, many forms of infrastructure sharing have been explored for cellular networks [10, 20, 21, 25, 28, 39, 51] that involve sharing physical network equipment (*e.g.*, MOCN). In contrast, RinP offers a qualitatively different approach: load-balancing users across independently managed infrastructures (*vs.* having operators jointly implement elastic resource management). Further, RinP does so *post-hoc* such that users can be load-balanced in a fine-grained and dynamic manner when and where operator incentives align and service quality improves *vs.* requiring advance planning and (critically) provisioning. Ultimately, we view RinP and active sharing as complementary approaches with tradeoffs, and we leave a study on combining them to future work.

MVNO architectures. As discussed in §2.4, MVNO architectures today are not built with network efficiency in mind, let alone providing better user performance in a fine-grained manner. Thus we view RinP as primarily relevant to the traditional MNO model which remains, by far, the most common cellular subscription model for users [35].

A recent clean-slate proposal [60] advocates UE-driven mobility in which UEs switch between any available operator to maximize performance. However, again they don't consider infrastructure efficiency as a motivation. More importantly, their approach is limited to an MVNO operator model, while RinP addresses the more common MNO model. Similarly, Magma [47] explores the design of a core network with minimal mobility support, finding many cellular use cases do not appear to require it: RinP can be applied in such a Magma network.

Steering of Roaming. Steering of Roaming (SoR) [42] enables home operators to influence UE behavior while roaming using the S6a interface and PLMN ordering. However, SoR only applies if the home network is unavailable [1, 2]; thus, SoR does not enable user sharing when there is *overlapping* coverage as RinP does.

Multi-access: Trends such as dual-SIM and dual-stack phones [43, 44] may enable a user-driven approach to RinP without operator involvement. This is a complementary approach that could realize some of the performance benefits of RinP though not its infrastructure efficiencies. Prior work, iCellular [57] develops a predictor that uses metrics available to the UE to dynamically switch to an operator that should provide better network performance. They rely on indirect measures, whereas RinP's operator-driven approach allows decisions that are based on direct knowledge of tower load and UE performance. Focusing on self-driving cars, CellFusion [67] takes advantage of multi-access by using each operator simultaneously. We note that dual-stack phones with multiple radios are not common on the market today.

User sharing: Within research, several more theoretical papers do consider offloading user traffic across operator boundaries, largely motivated by reducing base station energy consumption [16, 18, 19, 53–55, 64, 68]. Closer to our motivation is Berry *et al.* [18] which uses game theory to study whether the ability to offload users would disincentive operators from investing in infrastructure, finding no such disincentive exists provided operators prioritize their own subscribers before admitting those from other operators; we adopt this policy in §5.2.2. None of the above consider the systems and protocol issues involved in dynamically offloading users, nor the impact of densification costs.

802.11 Wireless Roaming: "Roaming" in the Wi-Fi context refers to when UEs move between cells. The protocols in the 802.11r/k/v [50] wireless standard outline roaming for the same Service Set Identifier (SSID), but do not touch on roaming between different SSIDs. The latter would be equivalent to the problem RinP is trying to solve within the cellular architecture, since it requires the detachment of the UE from one SSID and attachment/authentication to another.

8 Conclusion

RinP extends roaming so that users can dynamically roam between networks in overlapping coverage zones, based on load and service quality. Thus, RinP achieves efficient capacity scaling for operators and user performance improvements. RinP can be adopted by operators with minimal changes to the existing cellular infrastructure, by avoiding changing the standard protocols or requiring cooperative resource sharing.

Acknowledgments

We thank the anonymous reviewers and our shepherd Zhi-Li Zhang for their insightful comments and helpful feedback. We also thank Thomas Kampa, Hannaneh Barahouei Pasandi, and others for their discussions and contributions to the development of RinP. This research was funded in part by the National Science Foundation, Grant No. 2148230.

References

- [1] 3GPP. TS 22.011: Service accessibility.
- [2] 3GPP. TS 23.122: Access-Stratum (NAS) to Mobile Station (MS) in idle mode.
- [3] 3GPP. TS 23.401: General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access.
- [4] 3GPP. TS 23.501: System Architecture for the 5G System.
- [5] 3GPP. TS 23.502: Procedures for the 5G System.
- [6] 3GPP. TS 24.501: Non-Access-Stratum (NAS) Protocol for 5GS.
- [7] 3GPP. TS 35.304: User Equipment (UE) procedures in idle mode.
- [8] 3GPP. TS 36.331 LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification.
- [9] 3GPP. TS 38.304: User Equipment (UE) Procedures in idle mode and in RRC inactive state).
- [10] 3GPP. 3GPP RAN Sharing. <https://www.3gpp.org/news-events/1592-gush>, 2021.
- [11] AKAMAI. What is ai on edge networks? <https://www.akamai.com/glossary/what-is-ai-on-edge-networks>, 2024.
- [12] ALAMMOURI, A., ANDREWS, J. G., AND BACCELLI, F. Sinr and throughput of dense cellular networks with stretched exponential path loss. *IEEE Transactions on Wireless Communications* 17, 2 (2017), 1147–1160.
- [13] ANDREWS, J. G., BACCELLI, F., AND GANTI, R. K. A tractable approach to coverage and rate in cellular networks. *IEEE Transactions on communications* 59, 11 (2011), 3122–3134.
- [14] ANDREWS, J. G., ZHANG, X., DURGIN, G. D., AND GUPTA, A. K. Are we approaching the fundamental limits of wireless network densification? *IEEE Communications Magazine* 54, 10 (Oct 2016), 184–190.

- [15] ATZENI, I., ARNAU, J., AND KOUNTOURIS, M. Downlink cellular network analysis with los/nlos propagation and elevated base stations. *IEEE Transactions on Wireless Communications* 17, 1 (2017), 142–156.
- [16] BAO, Y., WU, J., ZHOU, S., AND NIU, Z. Bayesian mechanism based inter-operator base station sharing for energy saving. In *2015 IEEE International Conference on Communications (ICC)* (2015), IEEE, pp. 49–54.
- [17] BENDER, R., AND RAICE, S. European telecom companies race to merge. *Wall Street Journal* (Jun 2015).
- [18] BERRY, R., HONIG, M., NGUYEN, T., SUBRAMANIAN, V., ZHOU, H., AND VOHRA, R. On the nature of revenue-sharing contracts to incentivize spectrum-sharing. In *2013 Proceedings IEEE INFOCOM* (2013), IEEE, pp. 845–853.
- [19] BOUSIA, A., KARTSAKLI, E., ANTONOPOULOS, A., ALONSO, L., AND VERIKOUKIS, C. Game-theoretic infrastructure sharing in multioperator cellular networks. *IEEE Transactions on Vehicular Technology* 65, 5 (2015), 3326–3341.
- [20] CANO, L., CAPONE, A., CARELLO, G., CESANA, M., AND PASSACANTANDO, M. Cooperative infrastructure and spectrum sharing in heterogeneous mobile networks. *IEEE Journal on Selected Areas in Communications* 34, 10 (2016), 2617–2629.
- [21] CANO, L., CAPONE, A., CARELLO, G., CESANA, M., AND PASSACANTANDO, M. On optimal infrastructure sharing strategies in mobile radio networks. *IEEE Transactions on Wireless Communications* 16, 5 (2017), 3003–3016.
- [22] CISCO THE NEWSROOM. Cisco powers rakuten mobile's inter-network roaming service launch. <https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2020/m04/cisco-powers-rakuten-mobile-s-inter-network-roaming-service-launch-to-help-customers-stay-seamlessly-connected-from-one-mobile-network-to-another.html>.
- [23] COOPER, M. Antennas get smart. *Scientific American* 289, 1 (2003), 48–55.
- [24] DEB, P., MUKHERJEE, A., AND DE, D. A study of densification management using energy efficient femto-cloud based 5g mobile network. *Wireless Personal Communications* 101, 4 (Aug 2018), 2173–2191.
- [25] DENG, X., WANG, J., AND WANG, J. How to design a common telecom infrastructure for competitors to be individually rational and collectively optimal. *IEEE Journal on Selected Areas in Communications* 35, 3 (2017), 736–750.
- [26] DING, M., LÓPEZ-PÉREZ, D., MAO, G., WANG, P., AND LIN, Z. Will the area spectral efficiency monotonically grow as small cells go dense? In *2015 IEEE Global Communications Conference (GLOBECOM)* (2015), IEEE, pp. 1–7.
- [27] DING, M., WANG, P., LÓPEZ-PÉREZ, D., MAO, G., AND LIN, Z. Performance impact of los and nlos transmissions in dense cellular networks. *IEEE Transactions on Wireless Communications* 15, 3 (2015), 2365–2380.
- [28] DOYLE, L., KIBLDA, J., FORDE, T. K., AND DASILVA, L. Spectrum without bounds, networks without borders. *Proceedings of the IEEE* 102, 3 (2014), 351–365.
- [29] ERICSSON. Roaming in the 5g system: the 5gs roaming architecture. <https://www.ericsson.com/en/reports-and-papers/ericsson-technology-review/articles/roaming-in-the-5g-system>, 2021.
- [30] ERICSSON. Mobile data traffic forecast – mobility report. <https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast>, 2023.
- [31] ETTUS RESEARCH. Usp b210 usb software defined radio (sdr). <https://www.ettus.com/all-products/ub210-kit/>.
- [32] FITCHWIRE. EU network sharing decision may spur 5g deployment cooperation. <https://www.fitchratings.com/research/corporate-finance/eu-network-sharing-decision-may-spur-5g-deployment-cooperation-20-07-2022>, Jul 2022.
- [33] FOR WATER INNOVATION, N. A. <https://www.nawihub.org>.
- [34] GAILER. Gailer SIM Card. <https://www.gialer.com/collections/writable-sim-card/products/gialer-sim-card-program-kit-sim-card-tools-accessories-include-1-multi-sim-card-reader-5pcs-programmable-usim-cards-3-in-1-card-adaptor-kit-newest-grsim-software-programmer-tool>, 2024.
- [35] GLOBENEWSWIRE. United states mobile virtual network operator (mvno) market - growth, trends, and forecasts (2020 - 2025). <https://www.globenewswire.com/news-release/2021/06/15/2247282/0/en/United-States-Mobile-Virtual-Network-Operator-MVNO-Market-Growth-Trends-and-Forecasts-2020-2025.html>, 2022.
- [36] GOOGLE. Phones on Fi Guidebook: Use your own phone. <https://guidebooks.google.com/fi/phonesonfi/useyourownphone?hl=en-us>.
- [37] GOOGLE-FI. <https://fi.google.com/about/>, 2021.
- [38] gRPC. gRPC. <https://grpc.io/>.
- [39] GSM ASSOCIATION. Infrastructure sharing: An overview. <https://www.gsm.a.com/futurenetworks/wiki/infrastructure-sharing-an-overview/>.
- [40] GSMA. IR.188: LTE and EPC Roaming Guidelines.
- [41] GSMA. NG.113: 5GS Roaming Guidelines.
- [42] GSMA. Steering of Roaming Implementation Guidelines. Tech. Rep. 73, GSM Association (GSMA), 5 2022. Version 9.0.
- [43] GSMARENA. Asus Zenfone 2. https://www.gsmarena.com/asus_zenfone_2_ze551ml-6917.php, 2022.
- [44] GSMARENA. Huawei Mate 7. https://www.gsmarena.com/huawei_ascend_mate7-6543.php, 2022.
- [45] GSMARENA. OnePlus Nord CE 2 Lite 5G. https://www.gsmarena.com/oneplus_nord_ce_2_lite_5g-11344.php, 2022.
- [46] GUPTA, A. K., SABU, N. V., AND DHILLON, H. S. *Fundamentals of Network Densification*. Springer International Publishing, Cham, 2021, p. 129–163.
- [47] HASAN, S., PADMANABHAN, A., DAVIE, B., REXFORD, J., KOZAT, U., GATEWOOD, H., SANADHYA, S., YURCHENKO, N., AL-KHASIB, T., BATALLA, O., ET AL. Building flexible, low-cost wireless access networks with magma. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)* (2023), pp. 1667–1681.
- [48] HUA, S., LIU, P., AND PANWAR, S. The urge to merge: When cellular service providers pool capacity. *arXiv* (Nov 2011).
- [49] HUAWAI. Cloud vr network solution white paper. https://www-file.huawei.com/-/media/corporate/pdf/ilab/2018/cloud_vr_network_solution_white_paper_2018_en_v1.pdf, 2018.
- [50] IEEE. IEEE 802.11: Wireless LANs.
- [51] KIBLDA, J., DI FRANCESCO, P., MALANDRINO, F., AND DASILVA, L. A. Infrastructure and spectrum sharing trade-offs in mobile networks. In *2015 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)* (2015), IEEE, pp. 348–357.
- [52] LEE, S. Open5gs. <https://open5gs.org/open5gs/>, Dec. 2023.
- [53] LEE, S., MOON, S., AND YI, Y. On greening cellular networks by sharing base stations: A game-theoretic approach. In *Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools* (2016), pp. 87–94.
- [54] LEE, S., AND YI, Y. Distributed sharing of base stations for greening: A population game approach. In *International Conference on Game Theory for Networks* (2016), Springer, pp. 79–89.
- [55] LENG, B., MANSOURIFARD, P., AND KRISHNAMACHARI, B. Microeconomic analysis of base-station sharing in green cellular networks. In *IEEE INFOCOM 2014-IEEE Conference on Computer Communications* (2014), IEEE, pp. 1132–1140.
- [56] LI, E., ZENG, L., ZHOU, Z., AND CHEN, X. Edge ai: On-demand accelerating deep neural network inference via edge computing. *IEEE Transactions on Wireless Communications* 19, 1 (2020), 447–457.
- [57] LI, Y., DENG, H., PENG, C., YUAN, Z., TU, G.-H., LI, J., AND LU, S. iCellular: Device-Customized cellular network access on commodity smartphones. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)* (Santa Clara, CA, Mar. 2016), USENIX Association, pp. 643–656.
- [58] LIU, J., SHENG, M., LIU, L., AND LI, J. Network densification in 5g: From the short-range communications perspective. *IEEE Communications Magazine* 55, 12 (2017), 96–102.
- [59] LUMB, D. Google Fi Ditches Another Mobile Network, Leaving Only T-Mobile. <https://www.cnet.com/tech/mobile/google-fi-reportedly-drops-us-cellular-leaving-t-mobile-as-last-network/>.
- [60] LUO, Z., FU, S., THEIS, M., HASAN, S., RATNASAMY, S., AND SHENKER, S. Democratizing cellular access with cellbricks. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (2021), pp. 626–640.
- [61] LUTU, A., JUN, B., BUSTAMANTE, F. E., PERINO, D., BAGNULO, M., AND BONTJE, C. G. A first look at the ip exchange ecosystem. *ACM SIGCOMM Computer Communication Review* 50, 4 (Oct. 2020), 25–34.
- [62] LUTU, A., PERINO, D., BAGNULO, M., AND BUSTAMANTE, F. E. Insights from operating an ip exchange provider. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference (Virtual Event USA, Aug. 2021)*, ACM, p. 718–730.
- [63] MAREK, S. Uscellular ceo says network sharing will be a necessity. <https://www.fiercewireless.com/wireless/uscellular-ceo-says-network-sharing-will-be-necessity>, Aug 2022.
- [64] MOON, S., KIM, H., AND YI, Y. Brute: Energy-efficient user association in cellular networks from population game perspective. *IEEE Transactions on Wireless Communications* 15, 1 (2015), 663–675.
- [65] NEWS, A. AT&T interoperability issue. <https://abcnews.go.com/US/att-customers-facing-nationwide-issue-impacting-ability-call/story?id=110829533>, 2024.
- [66] NEWS, A. AT&T outage caused by software update, company says. <https://abcnews.go.com/US/att-outage-impacting-us-customers-company/story?id=107440297>, 2024.
- [67] NI, Y., ZHENG, Z., LIN, X., GAO, F., ZENG, X., LIU, Y., XU, T., WANG, H., ZHANG, Z., DU, S., YANG, G., SU, Y., CAI, D., LIU, H. H., XU, C., ZHAI, E., AND MA, Y. Cellfusion: Multipath vehicle-to-cloud video streaming with network coding in the wild. In *Proceedings of the 2023 ACM SIGCOMM 2023 Conference* (2023), pp. 668–683.
- [68] OH, E., KRISHNAMACHARI, B., LIU, X., AND NIU, Z. Toward dynamic energy-efficient operation of cellular network infrastructure. *IEEE Communications Magazine* 49, 6 (2011), 56–61.
- [69] SAERANG, S. The story of indosat, the most successful merger that empowers indonesia. <https://www.businesswire.com/news/home/20230228005739/en/The-Story-of-Indosat-the-Most-Successful-Merger-That-Empowers-Indonesia>, Feb 2023.
- [70] SCHAFFHALTER, P., KALRA, S., XU, L., GONZALEZ, J. E., AND STOICA, I. Leveraging cloud computing to make autonomous vehicles safer. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2023), IEEE, pp. 5559–5566.
- [71] SIEMENS. Siemens Energy Automation. <https://www.siemens.com/global/en/products/energy/energy-automation-and-smart-grid.html>.

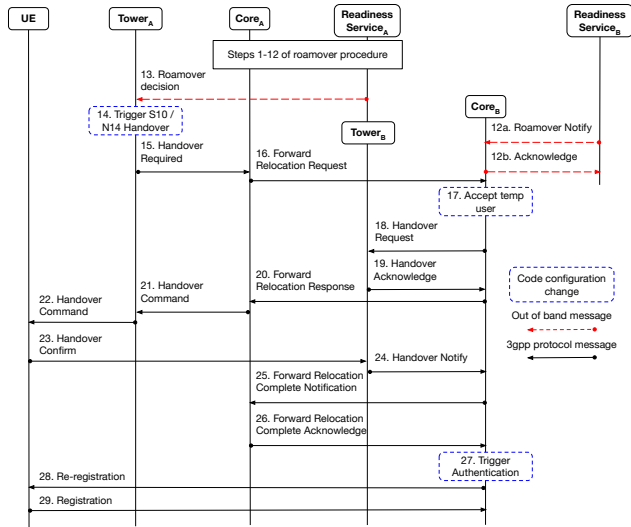


Figure 10: The roamover procedure optimized from Figure 3 using the S10/N14 interface. Messages 16, 20, 25-26 are for the MME transfer LTE procedure, the equivalent 5G messages are *Namf_Communication_CreateUEContextRequest/Response* and *Namf_Communication_N2InfoNotify* and *Ack*.

- [72] SOFTWARE RADIO SYSTEMS. srsRAN. <https://www.srsran.com/>.
- [73] THOMPSON, M. Cellular roaming’s inadequate security. *Criminal Legal News* (May 2024), 14.
- [74] VENKITESH, B., AND ROSENBERG, C. Smart roaming: How operator cooperation can increase spectrum usage efficiency at practically no cost. *IEEE Transactions on Network and Service Management* 16, 2 (Jun 2019), 690–700.
- [75] WANG, S., ZHANG, X., ZHANG, J., FENG, J., WANG, W., AND XIN, K. An approach for spatial-temporal traffic modeling in mobile cellular networks. *2015 27th International Teletraffic Congress* (Sep 2015), 203–209. arXiv: 1703.10804.
- [76] ZHANG, H., DONG, Y., CHENG, J., HOSSAIN, M. J., AND LEUNG, V. C. M. Fronthauling for 5g lte-u ultra dense cloud small cell networks. *IEEE Wireless Communications* 23, 6 (2016), 48–53.
- [77] ZHANG, X., AND ANDREWS, J. G. Downlink cellular network analysis with multi-slope path loss models. *IEEE Transactions on Communications* 63, 5 (2015), 1881–1894.
- [78] ZHAO, Z., LI, M., LI, R., AND ZHOU, Y. Temporal-spatial distribution nature of traffic and base stations in cellular networks. *IET Communications* 11, 16 (2017), 2410–2416.

Appendices are supporting material that has not been peer-reviewed.

A Readiness server details

Please refer to Table 4.

B S10/N14 Optimization

The full RinP protocol with the S10/N14 Optimization described in Section 3.5 can be found in Figure 10.

C Signaling cost of RinP

Compared to regular handovers, a roamover requires more signaling since it requires registration, deregistration, and the out-of-band messages to the readiness server. We measure the signaling cost as the number of messages exchanged for a single invocation of a procedure. Figure 11 shows this cost for a handover vs. a roamover in 5G, where registration and deregistration procedures make up




Figure 11: Assessing signaling costs in RinP in 5G NR.

~55% and ~38% of the roamover cost, respectively. For completeness, we also show the cost that the S10/N14 roamover optimization would incur.

D Further simulation parameter sensitivity analysis

Here we expand upon the sensitivity analysis in Section 5.2.3. (i) *Interference parameters.* We experimented with different values for our parameters that model interference sensitivity (s) as described in §5.2.2 in Figure 12 and path-loss exponent (α) in Figure 13. We see that, as expected, greater interference *increases* the savings due to RinP. For example, halving α from its default value of 4 to 2 leads to increased interference which improves savings from 22% to 40%. Similarly, dialing the interference sensitivity s from our default of 1 down to 0 reduced our savings from 25% down to 10%. (iv) *Increasing demand.* As shown in Figure 14, we experimented with increasing the total demand by increasing the per-user demands or the number of users and found that RinP’s savings hold in both.

E Real-World Measurement Data

Here we provide a detailed description of the real world measurement data summarized in Section 5.3.

Figure 15 - 22 depicts observed variation in RSRQ and bandwidth during the mobile and stationary tests with operators KDDI and SoftBank in Tokyo and FirstNet and Verizon in Berkeley. In RinP operators can choose metrics to trigger roamovers to craft their own policy. Here we employ RSRQ and bandwidth allocated to the UE (estimated as bandwidth observed by NDT), both of which are metrics available to operators. To project potential benefits from RinP, we consider two scenarios in which operators trigger roamovers: when the mean RSRQ and bandwidth on the other network exceeds that which they can provide on their own for at least $T = 30$ seconds. We choose this arbitrarily to represent the minimum time a UE would remain on a single operator’s network and to provide enough time to trigger a subsequent roamover and collect measurement reports. The purple vertical lines in the figures represent roamover opportunities in which the operator offering the best RSRQ over 30s window changes.

Table 5 expands on the Table 5 in Section 5.3. Table 5 depicts the percent of time roamover opportunities are available based on different RSRQ thresholds, Th .

API Call	Source/Destination	Content
RoamoverRequest()	$T_{source} \rightarrow RS_{source}$	T_{source} location and CellID, UE ID
ReadinessRequest()	$RS_{source} \rightarrow RS_{target}$	T_{source} location and CellID
ReadinessResponse()	$RS_{target} \rightarrow RS_{source}$	List of [CellInfo, QoS Param., Max UE Count, Time Limit]
RoamoverResponse()	$RS_{source} \rightarrow T_{source}$	List of [CellInfo]
MeasurementReport()	$T_{source} \rightarrow RS_{source}$	List of [CellInfo, Measurement Report]
Reserve()	$RS_{source} \rightarrow RS_{target}$	CellInfo, UE Count
Approve()	$RS_{target} \rightarrow RS_{source}$	CellInfo, UE Count, Time Limit
RoamoverDecision()	$RS_{source} \rightarrow Core_{source}$	CellInfo, UE ID

Table 4: Readiness API, for roamover of multiple UEs in a single procedure. T_{source} = Source Tower, RS_{source} = Source Readiness Service, RS_{target} = Target Readiness Service. CellInfo includes frequency, PLMN, CellID, and Location.




Figure 12: The impact of interference sensitivity (s) on RinP's savings.




Figure 13: The impact of the path-loss exponent (α) on RinP's savings.




Figure 14: The impact of increasing user demands by doubling per-user demand vs the number of users on RinP's savings.




Figure 15: RSRQ signal of riding on trains in Tokyo.




Figure 16: RSRQ signal of the stationary devices in Tokyo.




Figure 17: RSRQ signal of walking around in Berkeley.




Figure 18: RSRQ signal of the stationary devices in Berkeley.




Figure 19: Bandwidth observed riding trains in Tokyo.




Figure 20: Bandwidth observed of the stationary devices in Tokyo.




Figure 21: Bandwidth observed walking around in Berkeley.




Figure 22: Bandwidth observed of the stationary devices in Berkeley.

Scenario	RSRQ gain %	% roamover time		Bandwidth gain %	% roamover time	
		by RSRQ ($Th = 3dB$)	by bandwidth ($Th = 1Mbps$)		by bandwidth ($Th = 2Mbps$)	
Mobile in Tokyo (KDDI / SB)	10.98 / 9.83	18.37 / 14.99	30.21 / 47.09	25.62 / 55.01	22.40 / 35.76	
Mobile in Berkeley (FN / VZ)	10.37 / 8.50	20.17 / 15.47	1.51 / 97.27	0.12 / 722.32	0.85 / 97.12	
Stationary in Tokyo (KDDI / SB)	7.55 / 5.94	5.09 / 7.80	15.37 / 72.61	3.53 / 85.51	10.10 / 68.26	
Stationary in Berkeley (FN / VZ)	25.96 / 0.96	59.11 / 0.66	0.20 / 99.49	0.03 / 254.28	0.20 / 99.25	

Table 5: RinP performance comparison across different scenarios