# A Modern Edge-based Design for Cellular Roaming

Tenzin Samten Ukyab[1], Shaddi Hasan[2], Sylvia Ratnasamy[1], Scott Shenker[1,3]
[1]UC Berkeley, [2]Virginia Tech, [3]ICSI

## Abstract

Support for roaming in today's cellular architecture involves operating a private network that sits right next to the mainstream Internet. This network is called the IP Exchange Network (IPX) and it provides Mobile Network Operators (MNOs) with a mechanism to form roaming partnerships, resolve billing and QoS, and set up tunnels as necessary. We propose a design which we call the CPX (Consolidated IPX/IXP) where the roaming network converges with the Internet by leveraging existing edge providers to provide all the benefits that the IPX network provides. In addition to convergence, the CPX architecture provides lower latency for roaming users; in our simulation of a global CPX deployment, we demonstrate an average of 318% improvement in roaming connection latency.

## CCS Concepts

• **Networks** → **Network design principles**; **Mobile networks**.

## Keywords

Roaming, IPX Network, Cellular Architecture

**ACM Reference Format:**
Tenzin Samten Ukyab, Shaddi Hasan, Sylvia Ratnasamy, Scott Shenker. 2025. A Modern Edge-based Design for Cellular Roaming. In *The 24th ACM Workshop on Hot Topics in Networks (HotNets '25), November 17–18, 2025, College Park, MD, USA*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3772356.3772381

## 1 Introduction

Both the 1G cellular network and the Internet were conceived around the same time. However, they were launched separately and with drastically different architectures and design principles. ARPANET was built to allow diverse networks to interconnect without need for centralized control and focuses only on enabling best effort delivery. The 1G cellular network on the other hand was built as an extension of the existing telephone network, which used centralized control to ensure accurate billing, high availability, and high-quality connections.

With the success of the Internet, the cellular network has been slowly converging with the Internet's network model by moving to IP in 4G LTE and adopting a more decentralized system design in 5G. Convergence is important since it enables benefits such as consolidated physical infrastructure and fewer technologies for operators to master and manage. However, the *federation* of mobile networks, which is required to support roaming connections between Mobile Network Operators (MNOs), is the last place the Internet and the cellular network have failed to converge.

Today, cellular networks use an entirely separate "internetwork" for roaming, called the IP Exchange (IPX) network, that runs in parallel to the mainstream Internet; this network is private and not reachable by the public Internet. This made sense when the IPX network was originally built, since this separation enabled MNOs to: (i) provide quality-of-service (QoS) guarantees, which essentially is dedicated bandwidth, that the first iteration of the Internet would not have been able to achieve, and (ii) safeguard the legacy protocols that lacked security mechanisms. Now that elements like the 5G Security Edge Protection Proxy (SEPP) [20] address many legacy security issues and the Internet service quality has improved, there is little reason to keep these networks separate. And the rise of what we call edge providers – *i.e.,* entities that provide services like CDNs – has made flexible compute and other services easily available at the network edge. This paper describes how one can leverage these developments to consolidate the roaming network into the Internet, which we call the CPX (Consolidated IPX/IXP) network.

The IPX network consists of several providers, which are called IPX-Ps. Examples of IPX-Ps include: Telefónica, Syniverse, and BICs [24]. To explore how we might build the CPX, we must start by describing the functions that IPX-Ps currently provide. IPX-Ps facilitate roaming agreements between MNOs (i.e., by having a standard agreement that if you sign up with an IPX-P, you agree to allow roaming to/from other MNOs connected to the IPX network) and handles the billing for roaming charges between MNOs. When roaming occurs, we differentiate between the home MNO (the one the user has a contract with) and the visited MNO (the MNO where the user needs to gain access). IPX-Ps build tunnels between the home and roaming MNOs so that the home MNO can maintain control of their user's connections and QoS can be guaranteed on those connections. By "maintain

This work is licensed under a Creative Commons Attribution 4.0 International License.
*HotNets '25, College Park, MD, USA*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2280-6/25/11
https://doi.org/10.1145/3772356.3772381

control", we mean that the home network retains control over whether the user is eligible to receive service, can track usage (e.g., for billing purposes), can dictate the quality of service the user receives, and more broadly can define and modify the network policies that apply to their own users' traffic.

Note that the functions of an IPX span both the *infrastructure* and *contractual* arrangements between MNOs, and our CPX design will have to provide the same. Focusing first on the infrastructure, we believe that CPX can provide both the control and the QoS functions by leveraging edge providers (e.g., CDNs, Cloud or telecom providers with edge services) as follows. The Home MNOs run a virtualized core, which we call a "micro-core", at an edge server from an edge provider. The visited MNO uses locality-aware DNS provided by edge providers to locate a nearby micro-core belonging to the home MNO and then tunnels traffic to it as it does today. This means the home MNO is still in control of their user's connections (since the micro-core belongs to the Home MNO). Additionally, edge providers already provide mechanisms for bandwidth management [5, 7, 16] that can be used to establish tunnels (from the visited MNO to the micro-core) with the necessary QoS.

Our notion of micro-cores builds on recent efforts that virtualize core components [4]. There have been proposals to run these virtualized components in the cloud and even to run them close to the user (within the visited MNO or at the network edge) [35] to achieve much lower latency. Our work leverages these ideas, but extends them to realize a new goal: the elimination of the IPX infrastructure.

Having eliminated the need for a separate IPX infrastructure, we turn to the contractual arrangements facilitated by IPX-Ps. We note that in terms of contractual arrangements, IPX-Ps are essentially a third-party that arranges agreements and handles payments between MNOs. In CPX, these functions can be provided by any third-party through the same mechanisms and agreements that the IPX-Ps use today. We will call these third-parties *brokers*, and they can be any new or existing entity, such as IXPs, CDNs, or the IPX-Ps themselves (acting as brokers without using their network infrastructure). Indeed, the viability of such brokers is supported by recent CDN offerings in which they broker transit connectivity between small ISPs [10].

The convergence of cellular networks and the Internet is an important goal that offers immediate practical benefits (*e.g.,* consolidated infrastructure, streamlined feature sets) and longer-term architectural coherence (*e.g.,* with common design principles and assumptions). Yet, despite its importance, major gaps in convergence remain, with the cellular roaming network as a glaring example. In this paper, we sketch how roaming can be supported without standalone cellular infrastructure by instead leveraging common capabilities of the modern Internet — DNS, MPLS-based bandwidth management and edge compute. More generally, we hope that our paper stimulates discussion on the challenges that remain for full convergence and strategies to overcome them.

The remainder of this paper is organized as follows: we provide the necessary background in §2 for the proposed CPX design that follows in §3. Then we provide an analysis of the latency gains that CPX and micro-cores can provide globally in §4 and compare it to related work in §5.

# 2 Background

## 2.1 Cellular Architecture

Cellular networks are built out of a Radio Access Network (RAN) and a mobile core. For the remainder of this paper we will focus on the 5G network architecture, but our proposed design can be applied to the LTE architecture. The RAN is a set of radio towers – or gNodeBs (gNBs) for 5G – which connect to user equipment (UE) such as mobile phones. The gNBs manage the efficient use of the radio spectrum, ensuring it meets the QoS requirements of every UE. The mobile core connects the RAN to the Internet and manages the databases storing user information (*e.g.,* service plans, usage, *etc.*). The 5G core is assembled as a Service Based Architecture (SBA) or a collection of network functions that can be deployed in a cloud-based micro-service mesh. Instead of describing every part of the cellular core in detail, we focus on the components that are key to our design.

The User Plane Function (UPF) and Session Management Function (SMF) make up the data and control plane, respectively, of the gateway to the rest of the Internet. The UPF handles routing and QoS for user connections. The SMF handles session management, IP allocation, and traffic steering configuration for the UPF. Other important functions include: the Access and Mobility Management Function (AMF), Authentication Server Function (AUSF), Policy Control Function (PCF), and Unified Data Management (UDM), all of which are responsible for managing UE mobility and authentication. The NF Repository Function (NRF), Network Slice Selection Function (NSSF), and Short Message Service Function (SMSF) support discovery, slicing, and SMS, respectively.

## 2.2 Roaming Architecture

For two MNOs to be roaming partners, they need to connect their networks to achieve two things: data routing and user authentication. To achieve data routing, the SMFs and UPFs of each MNO must connect. For user authentication, the home MNO's AMF must connect to the visited MNO's AUSF, PCF, and UDM. Figure 1 depicts how the network functions connect during roaming. In order to resolve other relevant but less critical features in roaming (e.g., slicing, function discovery, SMS), other functions connect as well.

There are currently two options for routing roaming traffic: home routing and local breakout (LBO) [1]. Home routing – depicted in Figure 1 – reroutes user traffic through the gateway/UPF of the home MNO using the N9 interface connecting the home and visited UPFs [17]. In contrast, with LBO, user traffic is routed locally through the visited MNO's gateway, *i.e.,* traffic is routed directly from the Visited MNO's UPF to the Internet, incurring lower latency and inter-MNO traffic. Prior work has measured large latency penalties incurred from home roaming compared to local breakout [13–15, 25, 27, 28]. Despite these performance benefits, LBO is not widely deployed because it requires MNOs to delegate control of user traffic to visited networks. This delegation is both commercially and technically complex: home MNOs have limited means to validate that policies are applied as required, and visited MNOs would need to have mechanisms in place to support the diverse network policies all of their roaming partners might need applied to their users. As a result, the simpler home roaming approach is deployed ubiquitously today [12].

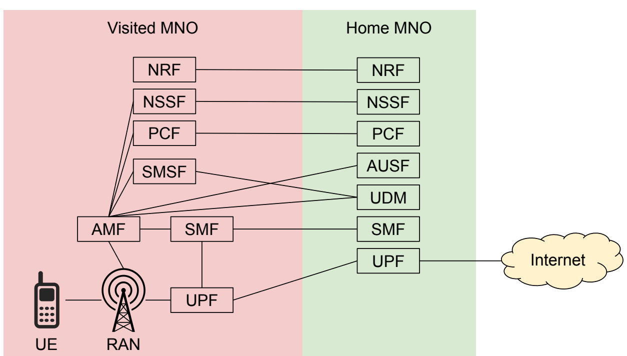


**Figure 1: Roaming architecture with home routing.**

## 2.3 IPX Architecture

The IPX network [18] consists of MNOs and IPX-Ps as shown in Figure 2. MNOs are allowed to contractually connect to a maximum number of two IPX-Ps, but can physically connect to multiple Points of Presence (PoP) of each contracted IPX-P for redundancy. In order to minimize traceability difficulty, the GSMA requires all IPX-P's peer with each other [18] (akin to IP peering between Tier-1 ASes on the Internet). The IPX-P to IPX-P peering can be done at Internet Exchange Points (IXPs) or through private connections.

As mentioned previously, IPX-Ps provide the following: (i) setting up contracts, (ii) tunneling, (iii) ensuring QoS, and (iv) settling billing between MNOs [19]. We discuss each of these in turn.

**Contracts:** Two MNOs can set up a bilateral roaming contract and just use the IPX network for tunneling. Alternatively, IPX-Ps can act as contract "hubs", where customers of the IPX-P have automatic roaming contracts with any of the other IPX-P's customers and the customers of their peers. So MNOs can connect to one IPX-P and get world-wide roaming coverage. This is one of the most useful functions on IPX.

**Tunneling:** IPX-Ps use proxies in their networks to handle the GPRS Tunneling Protocol (GTP) tunneling between the MNOs cores as described in §2.2. The proxies also implement a number of other functions that will be explained in the following sections.

By default, the IPX-Ps implement home roaming for all the customers. Additionally, they provide IPX Hub breakout, where the user data connection can break out to the Internet inside the IPX network in an attempt to reduce roaming latency. Unfortunately, IPX Hub breakout suffers from the same problem as LBO because H-MNO's lose control of their user's connections and thus is not commonly used.

**QoS:** Each MNO may provide different QoS guarantees for their users and have a different billing structure for each QoS. The proxies mentioned above support QoS negotiation between the roaming partners. Thus, one of the key benefits of the IPX Network is that it ensures the QoS guarantees (as determined by their contract with their home MNO) for roaming mobile connections.

**Billing:** Like contracts, billing can be resolved bilaterally. However, the proxies in the IPX networks are centralized places to record network usage for billing. To resolve billing between two MNOs, the IPX architecture provides the option of *cascaded billing* where the cost of the network usage is paid at each network connection point, rather than only at the ends. As an example, a UE roams onto the visited MNO (V-MNO) network and uses IPX-P1 to reach the home MNO (H-MNO), who is a customer of IPX-P2. The H-MNO pays IPX-P1 a fee for transiting the connection and the termination fee to be paid to the V-MNO. The termination fee is passed to the V-MNO through IPX-P1 then IPX-P2. IPX-P1 takes a cut of the transit fee and passes the rest to IPX-P2.

# 3 CPX Design

## 3.1 Overview

We will describe two approaches to building CPX, first a basic design and then an alternative design called CPX-direct (described in §3.6). Using either CPX design involves four steps: (1) setting up contracts, (2) discovering the home micro-core, (3) tunneling to the home micro-core, and (4) handling billing. Table 1 summarizes them and describes who implements them. Recall that the CPX network must provide both

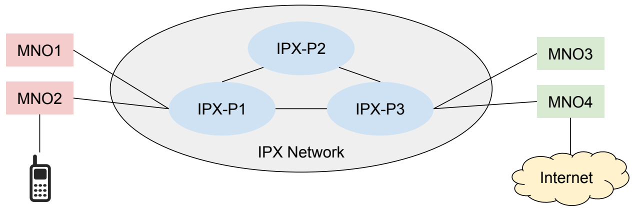


**Figure 2: The IPX network, where an UE's home MNO is MNO 4 and is roaming on MNO2.**

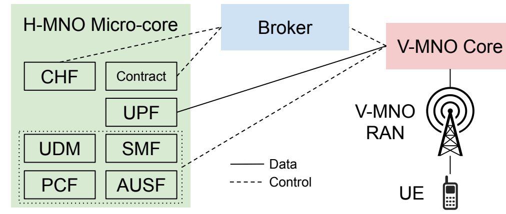


**Figure 3: Micro-core overview.**

| | IPX | CPX | CPX-direct |
|---|---|---|---|
| Contracts | IPX-P | Broker | On-demand |
| Discovery | N/A | DNS | DNS |
| Tunneling | IPX-P | V-MNO and H-MNO micro-core | V-MNO and H-MNO micro-core |
| Billing | IPX-P | Broker | V-MNO and H-MNO micro-core |

**Table 1: Function implementers for IPX and CPX.**

the contractual and the infrastructural functions of the IPX network. The first (1) and last (4) steps address the contractual functions and the middle two steps (2) and (3) address the infrastructural function.

The following sections will describe each of these steps in order, but before turning to these steps we provide some additional background. CPX involves encapsulating the key roaming functions of the mobile core in a micro-core, which is enabled by the change in the implementation of the cellular core from vendor appliances to software-based implementations (termed SBAs), which can be virtualized. As 5G infrastructure matures more cellular core functions are being implemented as virtualized containers [30]. To support CPX, an MNO must deploy micro-cores (*i.e.,* a subset of cellular core functions, discussed next) at existing edge providers such as Akamai, Cloudflare, *etc.*. Leveraging existing edge providers allows MNOs to easily achieve widespread deployment, while edge providers can host and support micro-core deployment as a new product offering for MNOs.

To handle the infrastructural functions, we equip each micro-core with the 5G functions UPF, SMF, UDM, PCF, and AUSF, as portrayed in Figure 3. These functions provide the tunneling and authentication necessary to route user data between MNOs, which we will describe in further detail in §3.4. For simplicity, we chose to not include other relevant roaming functions in the diagram, such as NRF, SMSF, and NSSF, and include the charging function (CHF) that we will further describe in §3.6.

To handle the contractual functions, CPX allows MNOs to employ a trusted third-party to accomplish the four steps, which we call a *broker*. Edge network providers where the micro-cores are already placed would be ideal brokers.

## 3.2 Contracts

For MNOs to enable roaming for their users, they must enter into a contractual agreement with other MNOs. Achieving a global set of pairwise contractual agreements is difficult and time-consuming. The IPX-P hubs provided a system to enable MNOs to plug into them and automatically establish roaming agreements with the rest of their customers or through peering with other IPX-Ps to reach the rest of the roaming MNOs. The previously mentioned brokers can provide the same functionality by serving as a "contract hub"; signing on with a broker would automatically create roaming contracts with all other customers who have signed up with a broker.

## 3.3 Discovery

For the visited MNO to route a roaming UE's data to the micro-core of the home MNO, the visited MNO must know where an appropriate micro-core is. For this, we can utilize DNS as a discovery mechanism by requiring each MNO to have a domain name that the visited MNO can use for the DNS query. An existing top level domain (TLD) such as ".com" or a new TLD such as ".mno" can be created by ICANN for the use of MNOs. The Public Land Mobile Network (PLMN) is the identifier for the specific MNO and is available in the SIM card of the UE. PLMNs are made up of the 3 digit Mobile Country Code (MCC) and 2 or 3 digit Mobile Network Code (MNC). Since the name field in a DNS record cannot be all numbers, we can append the strings "mnc" and "mcc" to each PLMN, resulting in a record such as "mnc01mcc001.mno" for a PLMN 001-01.

What IP address should this resolve to? CPX will leverage the managed DNS services (*e.g.,* [33]) of the edge providers they are already using to deploy their micro-cores for the following two tasks: (i) provide IP addresses for these micro-cores and (ii) manage updating DNS servers properly to return the IP address of the closest micro-core to the user.

## 3.4 Tunneling

We first describe how tunneling is implemented in the current IPX network. IPX proxies are used to set up GTP tunnels between the visited MNO and the home MNO. The visited MNO would set up a tunnel to the IPX-P proxy and the proxy would set up the rest of the tunnel to the home MNO. In acting as the middleman (like a VPN proxy), the IPX can help resolve protocol inconsistencies and handle QoS negotiation.

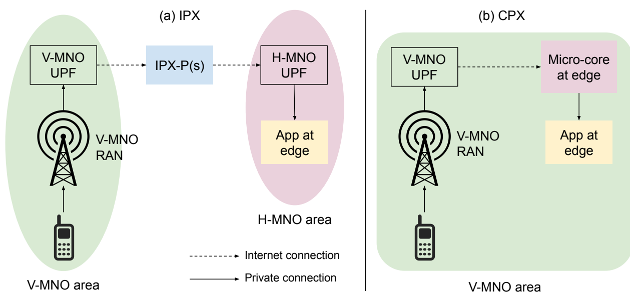


**Figure 4: Latency path of IPX vs CPX.**

As seen in Figure 4, the path for home roaming with IPXs starts at the visited MNO's area, or geographical location, at the UE and visited MNO's RAN. It then travels to the visited MNO's mobile core, where it exits the visited MNO's network into the IPX network. It travels through one or more IPX-Ps and arrives at the home MNO's area at the home MNO's core, where it finally breaks out to the Internet and reaches the nearest edge instance of the application it is trying to reach; note that nearest here would be with respect to the home MNO's core. The path in CPX is shortened: the path travels from the visited MNO directly over the Internet to the closest edge provider location (e.g., the CDN PoP that runs the Home MNO's micro-core). Since edge providers have a large global footprint, it is very likely that a micro-core will be deployed close to the visited MNO and hence UE; we evaluate these latency savings in §4. Finally, from the micro-core, traffic proceeds over the Internet to its final destination. In the modern Internet, this final destination might even be hosted within the same edge provider, thereby offering additional latency savings. While CPX does not reduce the number of logical (*i.e.,* provider) hops, in the CPX case each hop is, by design, shorter (i.e., lower latency).

To provide QoS guarantees, CPX can take advantage of another capability that edge providers already have. The micro-core can utilize managed bandwidth services provided by all major edge providers [5, 7, 16]. These services – which internally use MPLS circuits [9, 11] or other forms of bandwidth reservations – provide QoS along the path from the exchange point (IXP, not IPX) where the visited MNO and the edge provider meet, to the edge facilities where the micro-core runs. Hence we can leverage existing bandwidth management services to meet the QoS goals for CPX. Moreover, we speculate that explicit bandwidth management may be entirely unnecessary because (relative to the early days of the Internet) QoS on the Internet today is vastly improved even for best-effort traffic. However this conjecture is non-trivial to verify and hence we leave exploring it to future work.

## 3.5 Billing

The brokers already set up the contracts, so they can also resolve billing between the MNOs. Both the visited and home MNOs (the latter via its micro-core) can track usage (since user traffic flows through them) and report the same to the broker which then facilitates billing just as IPX-Ps do today.

## 3.6 Direct Approach

The approach presented above uses brokers for setting up contracts and resolving billing. However, brokers are not necessary to enable CPX. One can imagine the MNOs setting up on-demand roaming partnerships. These on-demand partnerships can be enabled using a protocol that facilitates contract negotiation and establishment between the visited and home MNO in a standardized way. In brief, when a UE attempts to attach to a new visited MNO that the home MNO does not have an established roaming partnership with, the visited and home MNO engage in protocol negotiation, where each MNO has a prior established negotiation policy. The MNO can use the discovery method described previously to find the home MNO. This does, however, require each MNO to build another component in their core that takes care of this contract protocol. For billing, the MNOs already have 3GPP standardized billing mechanisms in place for their users that can be used to resolve roaming billing, *e.g.,* the CHF [2], that work bilaterally between roaming MNOs.

## 3.7 Discussion

**Security:** When moving the roaming network from private to public, as in CPX, the surface of attacks opens up. Current roaming protocols in 2G through LTE already make users vulnerable to attacks, *e.g.,* surveillance attacks [29], protocol regression [20], and SMS hijacking [8]. With the roaming network moving to the Internet, MNOs are now more vulnerable to common network vulnerabilities, such as man-in-the-middle and denial of service attacks, on top of being more vulnerable to the cellular-specific attacks. Fortunately, the 3GPP standard has already enforced the Security Edge Protection Proxy (SEPP) mechanism in 5G [20] for roaming messages. SEPP acts as a secure gateway that protects inter-operator signaling, encrypts/authenticates messages, hides internal topology details, and enforces security policies.

**Incremental deployment:** Adoption of CPX should be relatively straightforward since an MNO can use the existing IPX and new CPX infrastructures in parallel with the expectation that CPX will eventually dominate by offering better performance for users, compounded by being able to seamlessly benefit from improvements in edge provider footprints and services. Since IPX-Ps already manage the commercial roaming relationships for MNOs, they can act as brokers in the new CPX network, though new entrants (e.g., CDN providers, cloud providers, Internet exchange providers) can also emerge as brokers. To get visited MNOs to tunnel to mini-core without DNS support, home MNOs can manually configure their existing bilateral roaming tunnels or the IPX-Ps they are using today to do the tunneling. This method is also an opportunity for aforementioned new entrants to provide the gains of CPX today, with no change required to visited MNOs.

**Adoption:** *Why would any of the cellular players chose to adopt CPX?* Besides the latency gains for users and the importance of the convergence of the Internet and cellular networks, roaming traffic is growing [21]. Niche use cases of roaming, such as IoT, put pressure on the roaming network [26] and continue to grow [32]. New small MNOs and existing entities that want to serve these emerging markets can lead the adoption of CPX.

## 4 Latency Gains from Micro-cores

In this section we evaluate the potential latency gains of using micro-cores at an edge server compared to home roaming with IPX-Ps, by simulating the round trip time experienced by a UE. To do this analysis, we use publicly available peering data from PeeringDB [31], AWS Cloudfront edge locations [6], and 2020 world population data from WorldPop [36]. We use a greedy approach to place UPFs, or mobile cores, across the globe at locations with the highest population to cover about 55% of the global population (assuming each UPF can reach 500 km). To simplify our simulation, we assume that UPFs are placed close to peering locations and that the latency between two points is proportional to the straight-line distance between them. We also assume the UE is accessing an application that uses an edge provider to distribute itself closer to the UE.

As a note, we leave a detailed evaluation of the rarely used LBO and IPX Hub breakout to future work, and for now note that LBO is a lower bound to CPX latency and IPX Hub break out is an upper bound. However, the important distinction is that neither of those are widely deployed today since they do not provide the home MNO an easy approach to maintaining control of roaming traffic as previously described, while CPX achieves this by leveraging edge providers.

Figure 5[1] depicts the latency difference using mini-cores with CPX instead of home roaming with IPX across the globe when the home location is in X-city in Northern California, USA. Locations closer to the UPFs placed and AWS Cloudfront edge location have higher latency gains. Whereas locations further from UPFs and edge locations experience the same latency using CPX or IPX. Figure 6 presents the CDF of the latencies, which show that CPX is clearly superior and, in particular, mean for CPX is 83.8 ms and for IPX is 346.09 ms, and median for CPX is 79.37 ms and for IPX is 301.85 ms. To show these results generalize to the entire globe, we performed the same simulation for 100 random locations across the globe. The average latency difference between CPX and IPX is 262.3 ms with a small standard deviation of 1.8 ms, which is a 318% improvement in latency.

## 5 Related Work

**Roaming measurement studies:** There have been many prior works that have measured roaming connections [13–15, 25, 27, 28]. All found there are large penalties in latency when UEs are subject to home roaming. CPX is able to provide lower latencies on top simplifying the network.

**IPX measurement studies:** Lutu et. al. [24, 26] obtained access to the private IPX-P network and was able to provide interesting and surprising insights into this previously opaque piece of critical infrastructure. For example, they find that much of the roaming traffic is from IoT devices (from embedded SIMs shipped in IoT devices), and most of the traffic is 2G/3G traffic. Vomhoff et. al. [35] provide a partial peering map of the IPX network from the private BGP routing table of an MNO. They also propose using regional breakout via the IPX networks to improve latency, since they find IPX Hub breakout is quite random in its breakout location. In contrast, we propose an overhaul of the IPX network that achieves the same end.

**Mobile Network Aggregators (MNAs):** Mobile Network Aggregators (MNA) aggregate services of multiple MNOs to provide optimized global coverage and sustained QoE. MNA companies, such as Twilio [23], Airalo [3], and Truphone [34], provide a "global SIM" product that achieves this. Prior work [4, 22] has looked into how these MNAs operate. MNA's can be "thick" in that they operate their own core and use roaming network's for their RAN or they can be "light" in that they do not operate a core and use SIMs from MNO to roam onto other MNO's networks. MNA's are a solution that solves an orthogonal problem than CPX and can even utilize CPX for "thick" deployments. Additionally, the prior work has found that MNAs suffer in performance due to home roaming and found evidence of poor breakout locations of

[1]The blank spaces over land areas are due to gaps in the input data.

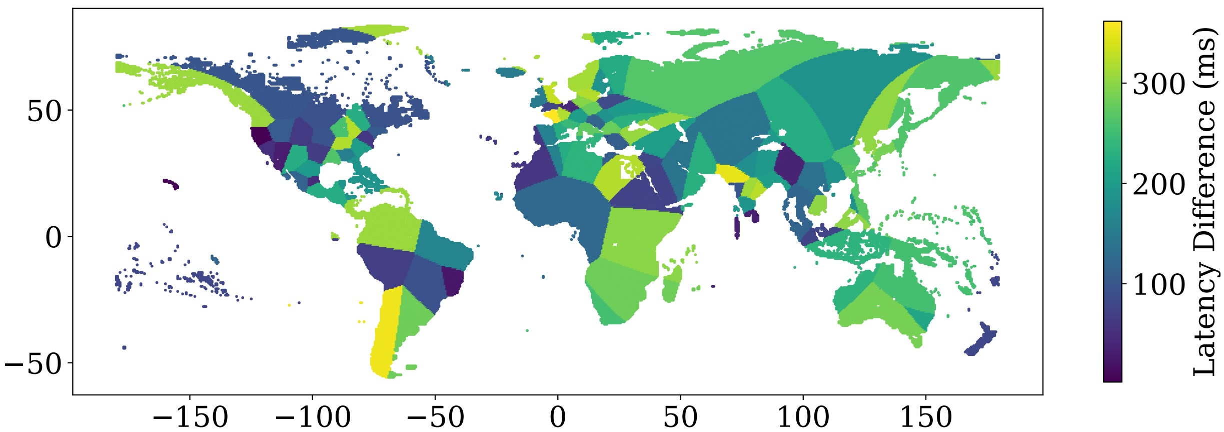


**Figure 5: Map of latency difference between IPX and CPX with X-city in Northern California, USA as the home location.**

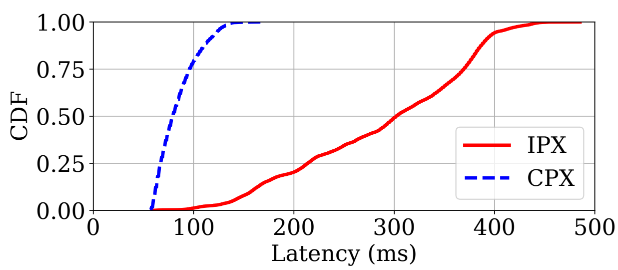


**Figure 6: CDF of the latency using IPX and CPX with X-city in Northern California, USA as the home location.**

IPX Hub breakout [35]. By using CPX, MNAs can cut down the latencies that their users experience.

**Local breakout:** Alcalá-Marín et. al. [4] suggest a LBO approach that brings the UPF of the home MNO closer to the roaming user through the use of edge providers' services or even inside the visited MNO's network. The first approach is a similar to ours; however, the approach to deploy the visited MNO's UPF inside the visited MNO's network is limited in that requires tight coupling between MNOs, which is also a weakness of all other LBO solutions where the roaming users use the visited MNO's UPF to break out. In addition, our goal was not just to improve performance but to merge roaming into the general Internet context.

**Systems built with edge providers:** Other work takes advantage of the various services edge networks provide to build useful systems. For example, Cloudflare allows small ISPs to use them as transit providers wherever they have PoPs rather than having to build longer backbones in [10].

## 6 Conclusion

It has been about 40-50 years since the birth of the Internet and cellular networks and we are still working towards convergence. The contribution of this paper lies in sketching what we believe is a very practical path to convergence for one of the last remaining hold outs.

## Acknowledgments

We thank the anonymous reviewers and our shepherd Dr. Ignacio De Castro Arribas for their insightful comments and helpful feedback. This research was funded in part by the National Science Foundation, Grant No. 2148230.

## References

[1] 3GPP. [n. d.]. TS 23.501: System Architecture for the 5G System.
[2] 3GPP. [n. d.]. TS 32.255: Telecommunication management; Charging management; 5G data connectivity domain charging; Stage 2.
[3] Airalo. [n. d.]. Global eSIMs. https://www.airalo.com/global-esim?srsltid=AfmBOopUnqKt88NbYJNGx2WKtqEeQfkSfAQoNZSc1ZzPHLr3bN1QXREK.
[4] Sergi Alcalá-Marín, Aravindh Raman, Weili Wu, Andra Lutu, Marcelo Bagnulo, Ozgu Alay, and Fabián Bustamante. 2022. Global mobile network aggregators: taxonomy, roaming performance and optimization. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services* (Portland, Oregon) *(MobiSys '22)*. Association for Computing Machinery, New York, NY, USA, 183–195. doi:10.1145/3498361.3538942
[5] Amazon. [n. d.]. Dedicated Interconnect overview. https://aws.amazon.com/directconnect/.

[6] Amazon. 2025. Amazon CloudFront Key Features. https://aws.amazon.com/cloudfront/features/?whats-new-cloudfront.sort-by=item.additionalFields.postDateTime&whats-new-cloudfront.sort-order=desc#Global_Edge_Network.

[7] Azure. [n. d.]. Azure ExpressRoute. https://azure.microsoft.com/en-us/products/expressroute.

[8] Crofton Black, Max Hoppenstedt, Omer Benjakob, Dominik Balmer, Jerome Hourdeaux, Marcel Rosenbach, Oliver Zihlmann, Hakan Tanriverdi, Gabriel Geiger, Melissa del Bosque, Tomas Statius, Beatriz Ramalho da Silva, Bashar Deeb, Nalinee Maleeyakul, Maud Jullien, Avi Scharf, and Daniel Howden. 2023. Ghost in the network. https://www.lighthousereports.com/investigation/ghost-in-the-network/.

[9] Cisco. 2009. MPLS Quality of Service (QoS). https://www.cisco.com/c/en/us/td/docs/ios/ios_xe/mpls/configuration/guide/mp_qos_xe.html

[10] Cloudflare. [n. d.]. Project Pangea. https://www.cloudflare.com/pangea/.

[11] Ross Callon Eric C. Rosen, Arun Viswanathan. 2001. Multiprotocol Label Switching Architecture. RFC 3031. doi:10.17487/RFC9405

[12] Ericsson. 2021. 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.

[13] Rostand A. K. Fezeu, Claudio Fiandrino, Eman Ramadan, Jason Carpenter, Daqing Chen, Yiling Tan, Feng Qian, Joerg Widmer, and Zhi-Li Zhang. 2024. Roaming across the European Union in the 5G Era: Performance, Challenges, and Opportunities. In *IEEE INFOCOM 2024 - IEEE Conference on Computer Communications*. IEEE, 2378–2387. doi:10.1109/INFOCOM52122.2024.10621234 ISSN: 2641-9874.

[14] Gabriel K. Gegenhuber, Wilfried Mayer, and Edgar Weippl. 2022. Zero-Rating, One Big Mess: Analyzing Differential Pricing Practices of European MNOs. In *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*. IEEE, Rio de Janeiro, Brazil, 203–208. doi:10.1109/GLOBECOM48099.2022.10001701

[15] Gabriel K. Gegenhuber, Wilfried Mayer, Edgar Weippl, and Adrian Dabrowski. 2023. MobileAtlas: Geographically Decoupled Measurements in Cellular Networks for Security and Privacy Research. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 3493–3510. https://www.usenix.org/conference/usenixsecurity23/presentation/gegenhuber

[16] Google Cloud. [n. d.]. Dedicated Interconnect overview. https://cloud.google.com/network-connectivity/docs/interconnect/concepts/dedicated-overview.

[17] GSMA. [n. d.]. NG.113: 5GS Roaming Guidelines.

[18] GSMA. 2007. IPX White Paper. https://www.gsma.com/newsroom/wp-content/uploads//IR.34-v17.0.pdf.

[19] GSMA. 2021. Guidelines for IPX Provider networks (Previously Inter-Service Provider IP Backbone Guidelines).

[20] GSMA. 2024. 5G Security Guide.

[21] GSMA Intelligence. 2021. Radar: Roaming in a post-Covid world.

[22] Hyunseok Daniel Jang, Matteo Varvello, Andra Lutu, and Yasir Zaki. 2024. Unraveling the Airalo Ecosystem. http://arxiv.org/abs/2408.14923 arXiv:2408.14923 [cs].

[23] KORE. [n. d.]. Super SIM. https://docs.korewireless.com/en-us/supersim.

[24] Andra Lutu, Byungjin Jun, Fabián E. Bustamante, Diego Perino, Marcelo Bagnulo, and Carlos Gamboa Bontje. 2020. A first look at the IP eXchange ecosystem. *ACM SIGCOMM Computer Communication Review* 50, 4 (Oct. 2020), 25–34. doi:10.1145/3431832.3431836

[25] Andra Lutu, Byungjin Jun, Alessandro Finamore, Fabián E. Bustamante, and Diego Perino. 2020. Where Things Roam: Uncovering Cellular IoT/M2M Connectivity. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Virtual Event USA, 147–161. doi:10.1145/3419394.3423661

[26] Andra Lutu, Diego Perino, Marcelo Bagnulo, and Fabián E. Bustamante. 2021. Insights from operating an IP exchange provider. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. ACM, Virtual Event USA, 718–730. doi:10.1145/3452296.3472930

[27] Anna Maria Mandalari, Andra Lutu, Ana Custura, Ali Safari Khatouni, Özgü Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jörg Ott, Martino Trevisan, Marco Mellia, and Gorry Fairhurst. 2022. Measuring Roaming in Europe: Infrastructure and Implications on Users' QoE. *IEEE Transactions on Mobile Computing* 21, 10 (Oct. 2022), 3687–3699. doi:10.1109/TMC.2021.3058787 Conference Name: IEEE Transactions on Mobile Computing.

[28] Anna Maria Mandalari, Andra Lutu, Ana Custura, Ali Safari Khatouni, Özgü Alay, Marcelo Bagnulo, Vaibhav Bajpai, Anna Brunstrom, Jörg Ott, Marco Mellia, and Gorry Fairhurst. 2018. Experience: Implications of Roaming in Europe. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking* (New Delhi, India) *(MobiCom '18)*. Association for Computing Machinery, New York, NY, USA, 179–189. doi:10.1145/3241539.3241577

[29] Gary Miller and Christopher Parsons. 2023. Finding You: The Network Effect of Telecommunications Vulnerabilities for Location Disclosure. https://citizenlab.ca/2023/10/finding-you-teleco-vulnerabilities-for-location-disclosure/.

[30] Iain Morris. 2025. Ericsson goes 'native' with Google in first 5G public cloud offer. https://www.lightreading.com/mobile-core/ericsson-goes-native-with-google-in-first-5g-public-cloud-offer.

[31] PeeringDB. 2025. PeeringDB. https://www.peeringdb.com.

[32] Juniper Research. [n. d.]. Global IoT Roaming Market 2024-2029.

[33] Kyle Schomp, Onkar Bhardwaj, Eymen Kurdoglu, Mashooq Muhaimen, and Ramesh K. Sitaraman. 2020. Akamai DNS: Providing Authoritative Answers to the World's Queries. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication* (Virtual Event, USA) *(SIGCOMM '20)*. Association for Computing Machinery, New York, NY, USA, 465–478. doi:10.1145/3387514.3405881

[34] Truphone. [n. d.]. Global IoT SIM Card. https://www.truphone.com/things/sim/.

[35] Viktoria Vomhoff, Marleen Sichermann, Stefan Geißler, Andra Lutu, Martin Giess, and Tobias Hoßfeld. 2024. A Shortcut Through the IPX: Measuring Latencies in Global Mobile Roaming with Regional Breakouts. In *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, Dresden, Germany, 1–10. doi:10.23919/TMA62044.2024.10559001

[36] WorlPop. 2018. The spatial distribution of population in 2020. https://hub.worldpop.org/geodata/summary?id=24777. doi:10.5258/SOTON/WP00647