

Island-Ready 5G Deployments: Decentralized Core Networks Enabling Crisis Connectivity at the Edge

Leon Janzen,* Bastian Bloessl,* Matthias Hollick*†

*Secure Mobile Networking Lab, Technical University of Darmstadt, Darmstadt, Germany

{ljanzen,bbloessl,mhollick}@seemoo.tu-darmstadt.de

†Secure Mobile Systems Group, IMDEA Networks Institute, Madrid, Spain

matthias.hollick@networks.imdea.org

Abstract—The concept of *island readiness* envisions that communication networks can fall back to local connectivity islands when global connectivity breaks, e.g., after natural disasters or targeted attacks. Island connectivity would enable users to use crisis-relevant applications hosted at the local edge during crisis response. However, realizing island readiness involves many stakeholders, and this paper is the first to consider the perspective of mobile network operators (MNOs). This paper explains why today’s 5G-Advanced deployments are not island-ready and outlines the options MNOs have to support island connectivity. We introduce a design for island-ready 5G and beyond core networks and demonstrate its functionality in an end-to-end lab testbed with commercial smartphones.

Index Terms—5G/6G, Island Readiness, Resilient Networks.

I. INTRODUCTION

The availability of the 5G system is crucial to everyday life, as billions of mobile devices rely on 5G connectivity to access the Internet. Cellular networks are considered critical infrastructure, which has brought increased attention to their resilience [1], [2]. Most mobile network operators (MNOs) keep the topology of their 5G systems private, but today’s deployments are expected to feature a few ($N \approx 5$) 5G cores per country [3]. These *high-availability* core networks feature synchronized replicas across different regions, minimizing the risk of all 5G cores being unavailable simultaneously. However, as traffic is backhauled across the country to one of the 5G cores, backhaul failures can still result in the disconnection of thousands of users from the 5G core, and, consequently, from the Internet. Recent crises have shown that such scenarios can lead to areas forming *connectivity islands*, where no connections to the outside world are possible [4], [5]. This is particularly devastating during crises, since smartphone applications and cellular services play a crucial role during crisis response [6], and most smartphone apps are rendered useless without Internet connectivity.

The concept of *island connectivity* envisions 5G-based communication networks that support island operation, enabling users on the same island to communicate with one another

This work was supported by the Federal Ministry of Education and Research of Germany in the project Open6GHub (grant number: 16KISK014), by the Federal Ministry of Research, Technology and Space of Germany in the project Open6GHub+ (grant number: 16KIS2407), and by the LOEWE initiative (Hesse, Germany) within the emergenCITY center [LOEWE/1/12/519/03/05.001(0016)/72].

ISBN 978-3-903176-82-9 © 2026 IFIP

and connect to locally available application servers during crises [7]. However, realizing island readiness involves many stakeholders [8], and this paper considers the perspective of MNOs. In principle, the 5G specifications include the necessary functionality to support island connectivity, such as cloud-native network functions (NFs) and the ability to connect multiple access and data networks to the 5G core [9]. 5G-Advanced (5G-A) introduced local breakouts of the user plane based on distributed user plane functions (UPFs) [10]. However, the control plane of today’s 5G-A systems typically resides in only a few data center locations. This georedundancy might be robust against attacks targeting individual 5G cores, but it takes a much more fine-grained distribution of 5G cores to achieve island readiness. As such, MNOs need to address non-trivial challenges to realize island-ready 5G and beyond systems. We contribute towards the vision of island readiness structured by two research questions:

RQ1: How can MNOs deploy 5G and beyond systems that achieve island readiness?

Island readiness promises local 5G connectivity in areas isolated from the global Internet. By design, this requires 5G core functionalities available in each island-ready area. However, deploying a highly distributed core network of full-scale 5G cores is unlikely to be feasible due to operational costs. In Sec. IV, we discuss options MNOs have to achieve island readiness. **We propose two components that MNOs can add to 5G-A cores to make them island-ready** while maintaining compliance with the 3rd Generation Partnership Program (3GPP) standards. These components synchronize stateful 5G NFs across all replicas and monitor the connectivity within the core network. We also propose the concept of failover 5G cores that operate only the synchronization and connectivity components during normal operation and transition into fully functional 5G cores in island scenarios. Sec. V describes the experiments we conduct to demonstrate the functionality of our solution, i.e., the correct synchronization of stateful NFs, the activation of failover cores, and the seamless transition from normal to island operation. To the best of our knowledge, **we are the first to develop and integrate a fully functional island-ready 5G testbed**

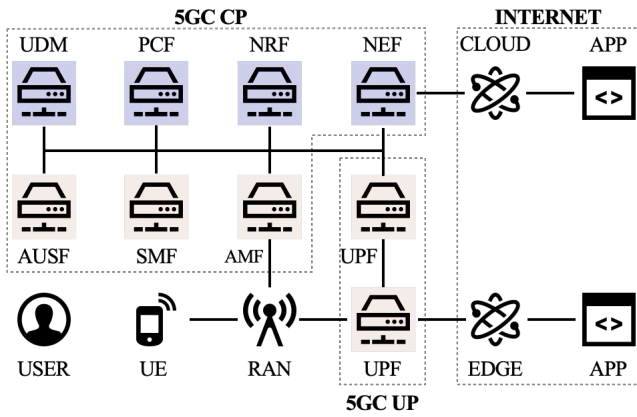


Fig. 1: *The 5G-Advanced System*

that supports end-to-end communication with commercial smartphones.

RQ2: How does island readiness affect the availability of the 5G system? Which overhead does it incur?

Island readiness promises resilience and availability advantages over today’s typical 5G deployments. In Sec. VI, **we evaluate the availability of country-wide 5G systems**, considering the topology and replication factor of 5G core networks as well as the availability of core instances and backhaul links. However, increasing the availability of 5G deployments will likely cause higher costs for MNOs. **We discuss the resource overhead of island readiness** and elaborate on options for MNOs to reduce costs when realizing island readiness. We publish our source code and the evaluation scripts for reproducibility and to facilitate future work.¹

II. 5G AND BEYOND COMMUNICATION NETWORKS

In 5G communication networks, the MNO operates a radio access network of base stations (BSs) that connect user equipment (UE) to the 5G core. From there, the user’s IP traffic is forwarded to the Internet, where applications are hosted on servers. Cloud and edge computing use virtualization to enable the flexible deployment of applications in appropriate locations. Relative to the user’s geographic location, the Internet can be divided into local parts (available at the *edge*) and remote parts (available on the *global Internet*). Today’s typical edge applications are characterized by the demand for low latency or high throughput, which outweighs the costs of edge deployments. However, deploying application servers close to the user can also be beneficial for crisis-relevant applications [7], [11].

A. The 5G and 5G-Advanced System

With 5G, the cellular core (Fig. 1) was redesigned to natively support user access from multiple access networks,

such as 5G New Radio (5G-NR), Long-Term Evolution (LTE), or Wi-Fi, and interfaces to multiple data networks, such as local and remote parts of the Internet, i.e., edge and cloud [9]. The NFs of the 5G core are separated into the user plane and the control plane.

The *control plane* implements a cloud-native, service-based architecture, in which each function provides a dedicated set of services. For instance, when a user requests a connection to the Internet, the access and mobility management function (AMF) receives the request, uses the authentication server function (AUSF) to authenticate the user, and instructs the session management function (SMF) to open a session [12]. Most NFs are stateless and rely on state management provided by a small set of stateful NFs: The unified data management (UDM) for subscriber data, the policy control function (PCF) for policy rules, and the network exposure function (NEF) that enables trusted applications to securely access core network services. Additionally, the network repository function (NRF) enables NFs to locate each other. The control plane is typically hosted remotely in a centralized core network of a few 5G cores per country. 5G-A supports local deployment of the *user plane*, i.e., the UPF, to give users a shortcut to the local network edge [9]. This shortcut is available after an initial connection to the control plane for authentication and session management. Therefore, the local deployment of the 5G user plane can improve the latency and performance of 5G connectivity to the edge, but not its resilience, because the control plane remains available only in a few central locations.

B. Why Communication Networks Fail During Crises

Internet outages often accompany natural disasters. Crises can cause disruptions, for example, due to broken connections, or as a cascading effect resulting from failures of other critical infrastructures, such as power outages that leave infrastructure without electricity. After the 2011 Great East Japan earthquake, the cellular network remained operational but was overloaded as survivors tried to communicate with loved ones [13]. During the 2022 Russian invasion of Ukraine, communication networks in Ukraine exhibited degraded performance levels as they adapted to the altered usage patterns [14]. The infrastructure of the Internet, cellular networks, and applications heavily depends on electricity [15]. For instance, the 2025 blackout on the Iberian Peninsula disrupted cellular systems, the local Internet, and connected applications, such as payment systems [16]. Similarly, attacks on Ukraine’s critical infrastructure have led to intermittent Internet outages since 2022 [14]. While backup power solutions for BSs [17] and data centers [18] exist, they are not designed to power the entire end-to-end chain of components. Rather, their purpose is to allow the graceful shutdown of individual components. During the 2021 Western European flooding, citizens of the Ahr valley were left without Internet access for several days after the collapse of streets and bridges holding fiber backhaul links [4]. Undersea Internet cables are publicly mapped [19] and have recently been targets of attacks [20]. The global Internet heavily relies on undersea cables, and previous incidents have led to the

¹<https://github.com/seemoo-lab/island-ready-5g-deployments>

isolation of areas, such as the 2024 West Africa Submarine Cable Outage [5].

III. RELATED WORK ON CRISIS CONNECTIVITY

Different categories of equipment exist to *restore* 5G connectivity, such as directional antennas, mobile BSs, and mobile 5G cores that can be deployed to crisis-struck areas. However, such equipment must be available in the affected area, and it takes time to make it operational. Instead, this section surveys mechanisms that aim to *maintain* connectivity in crisis scenarios with 5G network deployments, as these approaches align better with the concept of island readiness.

A. Maintaining 5G Connectivity During Crises

The connection of BSs to the 5G core is referred to as *backhaul*. When the backhaul fails, BSs are disconnected from the 5G core, and users in affected areas are unable to use 5G connectivity. Over the past 30 years, most copper-based backhaul cables have been replaced by fiber or point-to-point microwave links [21]. Dense areas often require high-capacity fiber backhauling, while microwave has advantages in rural settings [22]. By 2030, the industry expects an even mix of microwave and fiber backhaul [23]. There is little public information on how MNOs design backhaul networks, but related work suggests that transport networks are often co-located with other critical infrastructure, such as roads and gas or electricity networks [24]. Repairing or replacing broken fiber cables is a challenging and time-consuming process that requires human intervention. Alternatively, establishing point-to-point *microwave links* to replace the broken fiber cable can recover the backhaul faster, as practiced in the Ahr valley during the 2021 Western European flooding [25]. *Integrated access and backhaul (IAB)* enables wireless backhaul between BSs, where one BS connected to the 5G core acts as an IAB-donor for another IAB node [9, Clause 5.35].

5G-A studied access networks based on *non-terrestrial networks (NTNs)* for the 5G core, considering both spaceborne and airborne access networks [26]. In 2025, Eutelsat conducted the first successful test of a 5G connection with a satellite-based access network [27], but NTN-based backhauling is generally not in a mature state. More applications are expected for 6G [28], [29], e.g., there are proposals to use NTN-based access networks for public safety networks [30].

While global connectivity gives users access to the global Internet, *island-ready* 5G systems are equipped to maintain connectivity to the locally available subset of the Internet (the island) when global connectivity breaks. This can be a game-changer during crisis response, e.g., island readiness enables users to contact other users on the same connectivity island [7] and use crisis-relevant applications hosted at the local edge [11]. Notably, *private* or *campus* 5G networks and island connectivity are related in that both provide connectivity limited to a region. However, the distinct difference is that private 5G networks do not serve the general public, which is an important goal of island connectivity.

B. Non-5G Crisis Connectivity

Most 5G UEs, such as smartphones, support *Wi-Fi*. While Wi-Fi enables users to access the Internet when 5G connectivity is unavailable, its usable coverage is limited compared to 5G connectivity. There are ideas to merge private Wi-Fi networks into a city-scale network for crisis communication [31], but the practicability of this approach remains to be verified. Related work has studied the deployment of *satellite*-based crisis communication infrastructure [32], which bridges the 5G system to directly connect users to the Internet. Similarly, commercial solutions offering non-5G Internet connectivity via satellite are available. For example, Starlink was deployed during crisis responses in Germany [33] and Ukraine [34]. However, today's satellite communication solutions require the deployment of complex technology, and they are not ready to provide the capacity required to seamlessly support crisis connectivity for an entire area. *Device-to-device (D2D)* communication is a well-understood research area, and there have been many proposals for D2D crisis communication [35], often utilizing protocols natively supported by commercial smartphones (Bluetooth, Wi-Fi, and 5G baseband).

Research Gap: Today's 5G communication networks are susceptible to the isolation of geographical areas as a result of backhaul failures. The considered scenarios highlight that crises can disconnect areas from the global Internet, and users cannot rely on 5G connectivity during crisis response. Related work does not adequately address this issue, as recovering backhaul links is time-intensive, and NTN-based solutions are not yet mature.

IV. DESIGN: MAKING 5G-A CORES ISLAND-READY

RQ1 asks how MNOs can deploy island-ready 5G systems. In fact, island readiness poses non-trivial challenges to MNOs:

- *Synchronization.* Island-ready deployments need to provide distributed 5G core functionality, which requires the state replication of stateful NFs. When islands reconnect to the core network, their local state might have diverged from the global state. In that case, MNOs need to solve synchronization issues.
- *Connectivity.* Island-ready cores need to monitor the connectivity within the core network to transition between global and island connectivity if necessary.
- *Scalability.* Only areas equipped with core functionality can be island-ready, but operating many cores is hardly feasible considering the operational costs of full-scale 5G cores. MNOs need to find solutions that achieve island readiness but also scale for normal operation, considering operational costs and communication overhead.

As the first major contribution of this paper, **we propose an island-ready 5G core design** that allows MNOs to make state-of-the-art 5G-A cores island-ready by adding synchronization and connectivity functionality. This section shows how our design solves the aforementioned synchronization, connectivity,

and scalability challenges. We demonstrate the feasibility of our solution by integrating it into a state-of-the-art Open5GS core implementation. Our implementation is based on the `docker_open5gs` project² that shares configuration files for the joint setup of Open5GS and srsRAN. Our repository contains the source code to reproduce our results.

A. Synchronizing Distributed 5G Control Planes

Decentralizing *stateless* components amounts to deploying them to all desired areas and redirecting workloads to them. In contrast, decentralizing *stateful* components requires synchronizing state, which causes communication overhead. There are different interpretations of NFs's statefulness. Our interpretation aligns with that of Open5GS, as the NFs we consider stateful are those that, in Open5GS, use the unified data repository (UDR) to maintain state for them. Fig. 1 highlights stateless NFs in beige and stateful NFs in blue. The stateful NFs are the NEF, NRF, PCF, and UDM.

NFs register with the *NRF*, so it can support NFs in addressing each other. Assuming that decentralized cores are full instances in the sense that they feature a replica of each NF, the state the NRF maintains is only relevant within that core. As a result, the NRF state does not need to be synchronized to other cores. The *PCF* oversees adherence to policy rules. NFs, such as the AMF or SMF coordinate with the PCF to apply prevailing policy rules. As such, the state maintained by the PCF is the set of active policy rules, and treating all users equally requires synchronization of PCF state across all cores. The *NEF* grants third parties secure access to the 5G core. Therefore, the access level of a third party determines the state maintained by the NEF. Access to collected data requires synchronizing and potentially aggregating data across all cores. The *UDM* stores subscriber data, which is accessed by multiple NFs, such as the AUSF, to authenticate users. Therefore, the UDM state needs to be synchronized.

We solve the synchronization of stateful NFs by storing their state in a database that we replicate across all 5G cores. This process is 3GPP-compliant and aligns well with the cloud-native nature of NFs since the release of 5G. In Open5GS, the state is stored in a MongoDB container shared between the stateful NFs. We integrate synchronization functionality by implementing and adding a `sync` container that replaces the `open5gs-mongodb` container. Our `sync` container features a self-managed MongoDB replica set [36] that replicates the state across all Open5GS cores, so that the NFs of all cores are in sync at all times during normal operation. When an area forms an island, the local core can continue to provide 5G connectivity using the most recent synchronized state.

B. Monitoring Connectivity in the Distributed Core Network

The idea of island readiness is that areas that are disconnected from the outside Internet form an island and transition to island connectivity. Transitioning between normal operation and island operation requires the 5G cores to monitor their

connection to the remaining 5G core network. This allows them to activate island connectivity, which, for instance, involves pausing the state synchronization and providing connectivity based on the most recent state.

We integrate connectivity functionality by implementing a `connect` container that is added to each core. The `connect` container monitors the connectivity to the core network through Internet Control Message Protocol (ICMP) packets. When the `connect` container signals the transition into island operation, the `sync` component pauses participation in the replica set and continues storing the state in a local database. Once a connection to the core network is reestablished, the `connect` container signals the transition back into normal operation. In that case, the `sync` component rejoins the replica set, fixes any synchronization issues, and resumes participation in the state replication.

C. The Scalability of Island-Ready Solutions

It is fair to assume that MNOs have already solved the state synchronization and connectivity challenges for small core networks, since this functionality is required to operate today's typical high-availability core deployments. However, island readiness demands many more cores per country than state-of-the-art deployments, making scalability a challenge we must address. Increasing the 5G availability increases operational costs and communication overhead for MNOs.

1) *Operational Costs*: In areas where the permanent operation of an *active core* is not resource-efficient, we propose the provisioning of *failover cores*. In normal operation, failover cores route user traffic to the closest active core. This behaviour is equivalent to areas without a local core in traditional 5G deployments. The `sync` container of failover cores participates in state replication to keep in sync with the active cores. When the connectivity component signals that the active core is unavailable, the failover core activates and starts serving users with 5G connectivity. Failover cores enhance the 5G system's availability while reducing operational costs for MNOs compared to deploying full cores at scale.

We implement failover cores by deploying only the `sync` and `connect` components in an area. Each failover core connects to an active core. In normal operation, the failover core's `connect` component monitors the active core's availability. If the active core becomes unavailable, the `connect` component triggers the local deployment of all 5G NFs, transitioning into island mode and enabling island connectivity.

2) *Synchronization Complexity*: The state synchronization among the cores causes overhead that puts load on the MNOs' network. Without knowledge of the real-life database architectures of MNOs, we are bound to argue about the complexity of MongoDB replica sets [37]. Each replica set has up to 50 members, of whom one is elected the *primary*. The primary handles all write operations and records all database changes in an operation log. The secondaries replicate the primary's operation log and apply write operations to their own data sets, keeping their databases in sync with the primary. This scheme allows efficient scaling to up to 50 active cores, which

²https://github.com/herlesupreeth/docker_open5gs

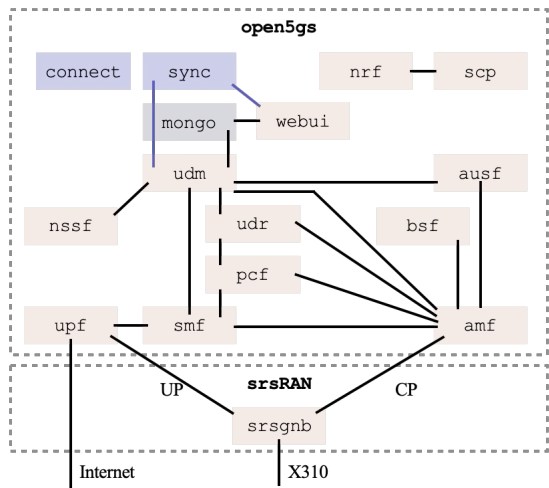


Fig. 2: Node Setup

is sufficient for country-wide 5G deployments. If desired, failover cores do not need to be members of the replica set because they only read the state. Instead, active cores can periodically replicate the current state to their failover cores.

V. EXPERIMENTS: ISLAND-READY END-TO-END TESTBED

We run experiments on three lab computers (*nodes*) equipped with an additional 10 Gbit Ethernet card. A USRP type X310 from Ettus Research (*radio*) with a board-mounted GPSDO is connected to each node. We use a Motorola Edge 30 Pro and a Google Pixel 6 as *UE* with programmable SIM cards from Sysmocom. Fig. 2 shows the setup deployed on each node. Each node hosts an Open5GS core and an srsRAN access network. We use Docker Compose to orchestrate deployments. The UEs connect to the radio via a 20MHz 5G-NR cell in the n78 band, and the radios connect to the nodes via an SFP+ cable to the 10Gbit Ethernet cards. All nodes are connected to the same local network, ensuring IP connectivity among them.

A. Experiment 1: Synchronizing Stateful Network Functions

This experiment (Fig. 3) aims to show that the `sync` component correctly synchronizes the subscriber information between UDMs on different cores. We run Open5GS cores on *node1* and *node2*. Open5GS features a web interface to manage the subscriber data in the UDM. We add the subscriber information of the SIM card of *ue1* through the `webui` of *core1*. We confirm that *ue1* can access the website via route `radio1`, `ran1`, `core1`. Next, we observe that the subscriber information is also available in the `webui` of *node2*. We shut down `radio1` and start `radio2`. We confirm that *ue1* can access the website via route `radio2`, `ran2`, `core2`. The experiment shows that the subscriber information is successfully synchronized between the UDMs in *core1* and *core2*. This demonstrates that **the `sync` component correctly synchronizes the state of stateful NFs**, which solves the challenge of distributing the 5G control plane.

B. Experiment 2: Activation of Failover Replicas

This experiment (Fig. 4) aims to show that failover cores are ready to take over when the active core becomes unavailable. We run srsRAN on *node1* and *node2*, but deploy Open5GS only on *node1*. We connect *ue1* to `radio2` and configure `ran2` to connect to `core1`, simulating *node1* as an area with an active core and *node2* as an area with a failover core. We confirm that we can access the website from *ue1* via route `radio2`, `ran2`, `core1`. We terminate `core1` on *node1* to simulate the active core being unavailable. We observe that `connect2` on *node2* detects the unavailability of `core1` and activates `core2`. We confirm that we can access the website from *ue1* via route `radio2`, `ran2`, `core2`. This demonstrates that **the `connect` component activates failover core replicas** to seamlessly take over connectivity when the active core fails.

C. Experiment 3: Maintaining Island Connectivity

This experiment (Fig. 5) aims to show that island-ready cores can transition to island connectivity during outages. We begin with the same setup as in the second experiment, where *node1* hosts an active core and *node2* has a failover core. At the local network edge of *node1* and *node2*, we deploy web applications `webapp1` and `webapp2`, respectively. We connect *ue1* to `radio1` and *ue2* to `radio2`. We confirm that both UEs can access both web applications and the website, which shows that *global* connectivity is intact. Then, we disconnect *node2* from the local lab network to simulate the area's isolation. As a result, the `connect` component of *node2* activates `core2`, which we confirm by observing that *ue2* indicates 5G connectivity. From *ue1*, we successfully connect to `webapp1` and the website, confirming that *area node1* was unaffected by the disconnection and is still in normal operation. From *ue2*, connecting to `webapp1` and the website fails because *global* connectivity is broken. However, *ue2* can still connect to `webapp2` hosted at the local edge of *node2*. The experiment shows that **island-ready 5G systems based on our island-ready core design can transition to island connectivity**.

Regarding RQ1: We present a two-component solution that MNOs can add to 3GPP-compliant 5G-A cores to make them island-ready. Our `sync` component synchronizes the state of stateful NFs, and our `connect` component monitors the connectivity within the core network. We integrated our solution into Open5GS and demonstrated its functionality in an island-ready end-to-end testbed with commercial smartphones.

VI. EVALUATION

RQ2 asks how island readiness affects the availability of 5G deployments and what overhead it incurs. This section introduces a system model for country-wide 5G deployments and evaluates the availability of 5G connectivity, considering parameters related to island readiness.




Fig. 3: Experiment 1 (Sec. V-A)




Fig. 4: Experiment 2 (Sec. V-B)






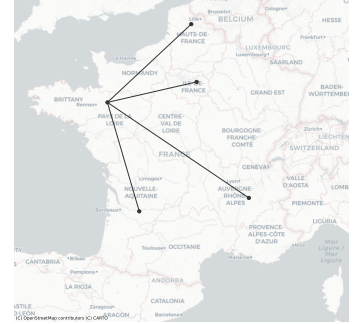
Fig. 5: Experiment 3 (Sec. V-C)



(a) Access (L_3)



(b) Transport (L_1+L_2)



(c) Core (L_0) with $N_0 = 5$ und $T_0 = \text{star}$

Fig. 6: Cellular Network Topology for France. We configured POBTOG with a population quota of $P = 10\,000$ citizens, $L = 4$ layers, $N = [5, 50, 400]$, and $T = [\text{star}, \text{star}, \text{ring}, \text{ring}]$. Our evaluation considers $T_0 \in \{\text{ring}, \text{star}\}$.

A. System Model: A Country-Wide 5G Deployment

We use the population-based topology generator (POBTOG) [38] to generate topologies for the access, transport, and core networks. POBTOG generates country-wide communication networks using the population as a proxy for network endpoints, such as smartphones and BSs, recursively partitioning the country into Voronoi cells until each cell holds P or fewer inhabitants. It expects a configuration stating the country, the population threshold per Voronoi cell P , the number of hierarchical layers L , the number of nodes N_l per layer $l \in [0, L-2]$, and the topology T_l per layer $l \in [0, L-1]$.

Fig. 6 shows the generated topology for France with a configuration informed by industry partners resembling real-world cellular and transport networks [38]. Layer L_3 contains a node at the centroid of each Voronoi cell ($N_3 = 9751$). They are divided into $N_2 = 400$ clusters using population-weighted k -means clustering. The centroid of each cluster becomes an L_3 node, and the topology $T_3 = \text{ring}$ connects the L_3 cluster and its L_2 centroid. This step is repeated for the remaining layers. Specifically, the nodes in the current layer are clustered, and the topology connects each cluster to its centroid. The centroid then becomes a node of the next layer. We consider L_3 to be the access network, L_2 and L_1 to be the transport network, and L_0 to be the core network. We fix the topologies to $T_3 = T_2 = \text{ring}$, $T_1 = \text{star}$, and evaluate $T_0 \in \{\text{star}, \text{ring}\}$.

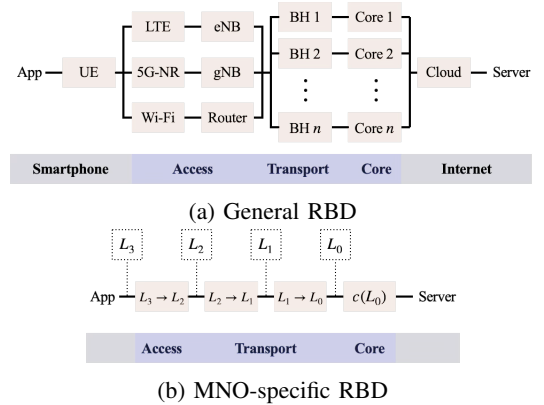


Fig. 7: Cellular Connectivity Availability. The RBDs illustrate the components required for end-to-end 5G connectivity.

B. Availability


The *availability* of a system is defined as “the probability that a system is operational and functioning correctly at a given point in time” [39]. Users who connect to the Internet via 5G connectivity rely on the availability of all components between their smartphone and the closest 5G core. Therefore, the end-to-end 5G availability A depends on the entire stack of Fig. 1. Fig. 7 depicts these components in reliability block diagrams (RBDs). We start with the RBD depicted in Fig. 7a and refine it to match the MNOs’ perspective. To this end, we consider which components are in the responsibility of MNOs and

analyze how relevant sub-availabilities can be refined to align with our system model. The MNO has no control over the UE or the Internet. On the other hand, the components from BSs to the 5G core are within the scope of the MNO. Accordingly, for our analysis, we set $A_{UE} = A_{Internet} = A_{App} = 1$. We abstract from the radio access technology (RAT) of the access network (such as 5G-NR, LTE, or Wi-Fi), assuming that the availability of 4G, 5G, and Wi-Fi BSs is similar. We set $A_{RAT} = A_{BS} = 1$ because failures of BSs or the radio interface are outside the scope of this work. This leaves $A_{AN} = A_{L_3 \rightarrow L_2}$, i.e., the access network availability equals the availability of the link connecting the L_3 node to the L_2 node in our system model (Fig. 6a). We model the availability of backhaul links as a function of their distance, and introduce the **link availability per kilometer** A_b as a parameter for our evaluation. For any nodes u, v with a direct link, we define $A_{u \rightarrow v} := A_b^{\text{dist}(u,v)}$.


The link lengths through the access and transport networks depend on the positions of the closest 5G core. This distance decreases as the number of cores per country increases (assuming a reasonable geographic distribution of the deployed cores). Therefore, we introduce the **core replication factor** N_0 as a parameter for our evaluation. The topologies used in the access, transport, and core networks impact the availability. While star topologies only feature a single path connecting source and target, there are two disjunct paths in ring topologies. Our system model fixes the topologies of access and transport networks to $T_3 = T_2 = \text{ring}$, and $T_1 = \text{star}$, so A_{AN} needs to consider two paths for $A_{L_3 \rightarrow L_2}$, and A_{TN} needs to consider two paths for $A_{L_2 \rightarrow L_1}$ and one path for $A_{L_1 \rightarrow L_0}$. Function $a(u, v)$ considers these topologies and computes the combined availability of all paths connecting source u and target v . This allows us to refine the availabilities of the access and transport networks to $A_{AN} = A_{L_3 \rightarrow L_2} = a(L_3, L_2)$ and $A_{TN} = A_{L_2 \rightarrow L_1} \cdot A_{L_1 \rightarrow L_0} = a(L_2, L_1) \cdot a(L_1, L_0)$.

The core network is of central interest in our work, so we introduce two core-related parameters for our evaluation: **core network topology** T_0 and **core replica availability** A_c . Considering T_0 and A_c in our evaluation allows detailed insights into achievable availabilities based on the core network configuration. The core availability A_{CN} depends on core replication factor N_0 , core network topology T_0 , and core availability A_c . Therefore, we define the function c to compute the core availability $A_{CN} = c(C_i)$ for any core replica $C_i \in L_0$. Our analysis of sub-availabilities allows us to adjust the end-to-end availability equation to $A = a(L_3, L_2) \cdot a(L_2, L_1) \cdot a(L_1, L_0) \cdot c(L_0)$, which is based on components under MNO control. Our repository contains the evaluation scripts to reproduce our findings.


1) *Backhaul Link Availability*: We fix $A_c = .99999$ and $T_0 = \text{ring}$ to evaluate the effect of A_b and N_0 on the 5G availability A . For better readability, we only report the availability levels, e.g., *two nines* for $A \geq .99$ and *five nines* for $A \geq .99999$. Fig. 8a plots A depending on N_0 and A_b , illustrating which core replication factor is necessary to achieve an availability level for a given link availability. With $N_0 = 5$, the assumed default in today's 5G systems, it takes



(a) Evaluation of link availability A_b (with $A_c = .99999, T_0 = \text{star}$)



(b) Evaluation of core availability A_c (with $A_b = .99999, T_0 = \text{star}$)



(c) Evaluation of core topology T_0 (with $A_b = A_c = .99999$)

Fig. 8: *Evaluation*. We compare the effect of backhaul link availability A_b , core replica availability A_c , replication factor N_0 , and core network topology T_0 on the 5G availability A .

$A_b = .999$ to achieve *one nine*, and even with very high backhaul availability, traditional 5G deployments fall short of *four nines*. *Five nines* can be achieved with $A_b = .999999$ and $N_0 \geq 10$, or with $A_b = 0.99999$ and $N_0 = 45$. These results indicate that higher core replication factors allow achieving availability with lower link availability. At the top of the spectrum, increasing $N_0 \geq 40$ facilitates availability levels of *five, six, and seven nines* with $A_b = .999999$.

2) *Core Replica Availability*: We fix $A_b = .99999$ and $T_0 = \text{ring}$ to evaluate the effect of A_c and N_0 on A . Fig. 8b shows that A_c has little impact on A compared to N_0 . The graphs of all reasonable core availabilities ($A_c \geq .9$) are practically identical. All graphs for $A_c < .9$ converge to the graph of $A_c \geq .9$ with growing N_0 . This means that the

impact of A_c on 5G availability is negligible for higher core replication factors. Also, high 5G availability is only possible with high core replication factors.

3) *Core Network Topology*: We fix $A_b = A_c = .99999$ to evaluate the effect of T_0 and N_0 on A . Fig. 8c depicts the mean and interquartile range (IQR) of star and ring topologies. It shows that T_0 has a negligible impact on A compared to N_0 , considering that the topology generation can result in slightly shifted coordinates, which explains the gaps between both graphs, such as for $N_0 \in [30, 40]$.

Our evaluation indicates that A_b **has the highest impact on A , but higher core replication can mitigate the effect of a lower backhaul link availability**. As such, MNOs can achieve high availability levels by increasing N_0 when increasing A_b is not possible. In particular, the results indicate that **high core replication factors ($N_0 \geq 45$ for country-wide 5G deployments) are required to achieve five nines**. Deploying that many full-scale cores would cause high operational costs for MNOs. Failover cores can increase availability while reducing operational costs during normal operation.

Regarding RQ2: Our evaluation shows that the effects of the core availability A_c and core network topology T_0 are negligible compared to the core replication factor N_0 . In particular, we found $N_0 \geq 10$ necessary for 5G availability of *four nines* and $N_0 \geq 45$ for *five nines*. MNOs can alleviate overhead and costs by augmenting the 5G core network with failover cores.

VII. DISCUSSION

A. The MNO's Perspective on Island Readiness

1) *Distributing 5G Core Functionality*: High core replication factors are required to achieve a 5G availability of *five nines*. However, decentralizing 5G cores is hard because they consist of many stateful NFs. We solved this challenge by introducing a `sync` component that synchronizes the state of stateful NFs across distributed 5G cores. The synchronization functionality builds the basis for making 5G and beyond communication networks island-ready.

2) *Operating Highly-Distributed 5G Deployments*: Deploying, connecting, and operating many 5G cores is resource-intensive, and running cores in areas with a few users may not be cost-effective. We presented *failover* 5G cores that remain passive in normal operation but periodically poll the current state from active cores. Failover cores are relatively resource-efficient in normal operation and can be activated to provide island connectivity when the area forms an island. The option to provision failover cores where active cores are not cost-effective allows MNOs to achieve high 5G availabilities with moderate operational costs.

3) *The Size of Island-Ready Areas*: Since each area needs an available 5G core (active or failover) to be island-ready, the absolute number of 5G cores determines the maximum number of islands the 5G system can support. In typical 5G deployments today, countries can fall into only a few ($N_0 \approx 5$)

islands. As N_0 increases, the 5G system can fall into more islands and is more resilient to backhaul failures. At the same time, the required resources per 5G core decrease because each 5G core has to handle fewer users.

B. Limitations and Future Work

The goal of our work is to maintain connectivity for users in crisis-struck areas as far as possible. The solutions presented in this paper provide MNOs with an option to make 3GPP-compliant 5G-A cores island-ready. Our experiments prove the feasibility of island-ready 5G deployments and demonstrate our proposed solution in an end-to-end testbed with commercial smartphones. We encourage further experiments to complement our work, such as large-scale testbeds that consider a broader range of users. One process that requires additional attention is the reverse transition from island to normal operation, which likely poses additional challenges. In that process, islands need to resynchronize their state with the rest of the network, potentially resolving divergent states. Given the proposed high number of active and failover cores in island-ready 5G deployments, reaching consensus about divergent states is a non-trivial challenge.

For our system model, we generated a country-wide topology using POBTOG, as MNOs seldom share official public information on 5G deployments. The resulting network topology is informed by industry partners resembling real-world cellular and transport networks [38], but only MNOs could verify its accuracy. To get a suitable system model, we had to abstract details and make assumptions, such as abstracting from the availability of different RAT and assuming a backhaul link availability of *five nines* per kilometer. Reproducing our results under different assumptions can enhance the generalizability of our findings. Future work may investigate available data sources on 5G network deployments and combine them into topologies that closely resemble the real world, while respecting the necessary discretion for cellular networks as critical infrastructure. Ideally, MNOs and researchers could collaborate in this endeavor. The focus of this work was on the end-to-end availability of 5G communication networks. Future work could analyze other goals, such as sovereignty.

The vision of island-ready communication networks involves many stakeholders [8]. This paper considered the MNOs' perspective to supplement previous work on the user perspective [7]. MNOs play a crucial role in achieving island readiness, as the 5G system serves as the backbone of the communication network, connecting users with each other and with applications. Designing 5G and beyond systems that support the concept of island readiness is a milestone on the road to island-ready communication networks. Still, there are many more stakeholders involved than MNOs, and numerous additional challenges to address.

VIII. CONCLUSION

This paper demonstrates how MNOs can make 5G and beyond *resilient by design* against backhaul failures by applying the concept of island readiness. We present two components

that add synchronization and connectivity functionality to 3GPP-compliant 5G-A cores to address the challenges of operating highly distributed 5G core networks (RQ1). We demonstrate the feasibility of the island connectivity concept by conducting experiments in our end-to-end communication testbed using commercial smartphones. Our evaluation shows that high core replication factors ($N \geq 10$) achieve higher availability than today's typical 5G deployments ($N_0 \approx 5$). Augmenting the core network with failover cores reduces the synchronization overhead and operational costs for MNO (RQ2). Realizing the concept of island readiness can be a game-changer during crisis response in isolated areas and, on top of that, increase the resilience of 5G connectivity in everyday life. The ongoing standardization of 6G should prioritize island connectivity to ensure future communication networks are resilient by design against backhaul failures.

REFERENCES

- [1] CISA, "CISA 5G strategy: Ensuring the security and resilience of 5G infrastructure in our nation;" 2020, Accessed: Oct 10, 2025. [Online]. Available: <https://www.cisa.gov/resources-tools/resources/5g-strategy>
- [2] BNetzA, "Resilience of telecommunications networks," 2022, Accessed: Oct 10, 2025. [Online]. Available: <https://www.bundesnetzagentur.de/DE/Fachthemen/Telekommunikation/Resilienz/start.html>
- [3] J. P. Rula, "Adopting a gateway centric view for cellular network content delivery," Ph.D. dissertation, Northwestern University, 2017.
- [4] A. Fekete and S. Sandholz, "Here comes the flood, but not failure? Lessons to learn after the heavy rain and pluvial floods in Germany 2021," *Water*, vol. 13, no. 21, p. 3016, 2021.
- [5] Internet Society, "2024 West Africa submarine cable outage report," 2024, Accessed: Nov 05, 2025. [Online]. Available: <https://www.internetsociety.org/resources/doc/2024/2024-west-africa-submarine-cable-outage-report/>
- [6] K. Eismann, O. Posegga, and K. Fischbach, "Collective behaviour, social media, and disasters: A systematic literature review," in *Twenty-Fourth European Conference on Information Systems*, 2016.
- [7] L. Janzen, F. Putz, M.-A. Kaufhold, K. Straub, and M. Hollick, "The user perspective on island-ready 6G communication: A survey of future smartphone usage in crisis-struck areas with local cellular connectivity," in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 2025, pp. 1–22.
- [8] L. Janzen and M. Hollick, "Making cellular networks crisis-proof: Towards island-ready, resilient-by-design 6G communication networks," *Computer Communications*, vol. 247, p. 108407, 2026.
- [9] 3GPP, "System architecture for the 5G system (5GS)," 3GPP, Technical Specification (TS) 23.501, 2025, version 18.10.0.
- [10] —, "5G system enhancements for edge computing," 3GPP, Technical Specification (TS) 23.548, 2025, version 18.10.0.
- [11] E. Lukau, J. Schiller, and U. Meissen, "Public safety campus networks: Towards a reference architecture for post-disaster edge-computing and crisis-communication," in *2024 IEEE WFPST*, 2024, pp. 74–80.
- [12] 3GPP, "Procedures for the 5G system (5GS)," 3GPP, Technical Specification (TS) 23.502, 2025, version 18.10.0.
- [13] H. Yamamura, K. Kaneda, and Y. Mizobata, "Communication problems after the Great East Japan Earthquake of 2011," *Disaster medicine and public health preparedness*, vol. 8, no. 4, pp. 293–296, 2014.
- [14] A. Jain, D. Patra, P. Xu, J. Sherry, and P. Gill, "The Ukrainian Internet under attack: An NDT perspective," in *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022, pp. 166–178.
- [15] L. Khaloopour, Y. Su, F. Raskob, T. Meuser, R. Bless, L. Janzen, K. Abedi, M. Andjelkovic, H. Chaari, P. Chakraborty *et al.*, "Resilience-by-design in 6G networks: Literature review and novel enabling concepts," *IEEE Access*, vol. 12, pp. 155 666–155 695, 2024.
- [16] R. González-Pozo, "Social profiles and response patterns during the 2025 Iberian peninsula power outage. The case of Spain," *International Journal of Disaster Risk Reduction*, vol. 130, p. 105813, 2025.
- [17] F. Wang, X. Fan, F. Wang, and J. Liu, "Backup battery analysis and allocation against power outage for cellular base stations," *IEEE Transactions on Mobile Computing*, vol. 18, no. 3, pp. 520–533, 2018.
- [18] Y. Chen, K. Shi, M. Chen, and D. Xu, "Data center power supply systems: From grid edge to point-of-load," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 11, pp. 2441–2456, 2022.
- [19] J. Franken and C. Reuter, "The subsea data cable security map – fusing public information for enhanced critical maritime infrastructure security," in *4th Maritime Security Conference Proceedings*, 2024.
- [20] D. Guilfoyle, T. P. Paige, and R. McLaughlin, "The final frontier of cyberspace: The seabed beyond national jurisdiction and the protection of submarine cables," *International & Comparative Law Quarterly*, vol. 71, no. 3, pp. 657–696, 2022.
- [21] O. Tipmongkolsilp, S. Zaghoul, and A. Jukan, "The evolution of cellular backhaul technologies: Current issues and future trends," *IEEE Communications Surveys & Tutorials*, vol. 13, pp. 97–113, 2011.
- [22] R. Lombardi, "Microwave and millimetre-wave for 5G transport," *ETSI White Paper*, vol. 25, 2018.
- [23] Ericsson, "Backhaul media for 5G and beyond," 2023, Accessed: Oct 14, 2025. [Online]. Available: <https://www.ericsson.com/en/reports-and-papers/microwave-outlook/reports/2023>
- [24] S. Anderson, L. Salamatian, Z. S. Bischof, A. Dainotti, and P. Barford, "iGDB: Connecting the physical and logical layers of the Internet," in *Proceedings of the 22nd ACM Internet Measurement Conference*, 2022.
- [25] KOWID, "Richtfunk als Alternative zur Glasfaser," *Behörden Spiegel*, 2022, Accessed: Oct 14, 2025. [Online]. Available: https://kowid.de/wp-content/uploads/2022/04/Beh%C3%B6rden SpiegelPTL_OR926.pdf
- [26] 3GPP, "Non-terrestrial networks," 2025, Accessed: Oct 12, 2025. [Online]. Available: <https://www.3gpp.org/technologies/ntn-overview>
- [27] MediaTek, "Eutelsat, MediaTek and Airbus announce world's first 5G non-terrestrial network connection, leveraging OneWeb LEO satellites," 2025, Accessed: Oct 13, 2025. [Online]. Available: <https://www.mediatek.com/press-room/eutelsat-mediatek-and-airbus-announce-worlds-first-5g-non-terrestrial-network-connection-leveraging-oneweb-leo-satellites>
- [28] M. M. Azari, S. Solanki, S. Chatzinotas, O. Kodheli, H. Sallouha, A. Colpaert, J. F. Mendoza Montoya, S. Pollin, A. Haqiqatnejad, A. Mostaani, E. Lagunas, and B. Ottersten, "Evolution of non-terrestrial networks from 5G to 6G: A survey," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 4, pp. 2633–2672, 2022.
- [29] F. Rinaldi, H.-L. Maattanen, J. Torsner, S. Pizzi, S. Andreev, A. Iera, Y. Koucheryavy, and G. Araniti, "Non-terrestrial networks in 5G & beyond: A survey," *IEEE Access*, vol. 8, pp. 165 178–165 200, 2020.
- [30] F. Völk, R. T. Schwarz, M. Lorenz, and A. Knopp, "Emergency 5G communication on-the-move: Concept and field trial of a mobile satellite backhaul for public protection and disaster relief," *International Journal of Satellite Communications and Networking*, vol. 39, no. 4, pp. 417–430, 2021.
- [31] Z. Liu, O. Chabra, J. Lynch, A. Martin, C. Li, and H. Balakrishnan, "Scalable routing in a city-scale Wi-Fi network for disaster recovery," in *1st New Ideas in Networked Systems (NINeS 2026)*, vol. 139, 2026.
- [32] F. Kagai, P. Branch, J. But, R. Allen, and M. Rice, "Rapidly deployable satellite-based emergency communications infrastructure," *IEEE Access*, vol. 12, pp. 139 368–139 410, 2024.
- [33] Heise, "Nach Hochwasser-Katastrophe: Rheinland-Pfalz stellt Starlink-Antennen auf," 2025, Accessed: Oct 13, 2025. [Online]. Available: <https://www.heise.de/news/Nach-Hochwasser-Katastrophe-Rheinland-Pfalz-stellt-Starlink-Antennen-auf-6143652.html>
- [34] The Independent, "Elon Musk helps Ukraine with SpaceX's Starlink satellites," 2022, Accessed: Oct 13, 2025. [Online]. Available: <https://www.independent.co.uk/news/world/europe/elon-musk-helps-ukraine-satellites-b2024893.html>
- [35] H. Nishiyama, M. Ito, and N. Kato, "Relay-by-smartphone: Realizing multihop device-to-device communications," *IEEE Communications Magazine*, vol. 52, no. 4, pp. 56–65, 2014.
- [36] MongoDB, Inc., "Deploy a self-managed replica set," Online, 2025, Accessed: Dec 03, 2025. [Online]. Available: <https://www.mongodb.com/docs/manual/tutorial/deploy-replica-set/>
- [37] —, "Replication," Online, 2025, Accessed: Feb 06, 2026. [Online]. Available: <https://www.mongodb.com/docs/manual/replication/>
- [38] F. Poignée, F. Loh, S. Schardt, F. Lipp, D. Hock, and T. Hoßfeld, "POBTog: A population-based topology generator for country-wide communication networks," *Proceedings of the International Conference on Networked Systems 2025*, pp. 47–50, 2025.
- [39] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, pp. 11–33, 2004.