Enhancing Local 5G Resilience through Satellite-Terrestrial Integrated Backhaul and Adaptive Resource Control

Akihiro Nakao
Professor
School of Engineering
The University of Tokyo

1. Introduction

As Beyond 5G evolves, interest in non-terrestrial networks (NTNs) grows because they extend coverage and strengthen communication systems against terrestrial infrastructure failures. Within NTNs, satellite communication provides wide-area coverage and an independent backhaul path for Local 5G systems.

In this article, Local 5G resilience means the ability to maintain or rapidly restore essential connectivity and service quality when terrestrial backhaul fails, access demand rises sharply, or available network capacity becomes constrained. Satellite backhaul improves path diversity, but a backup path alone does not guarantee service continuity. Satellite links generally provide less bandwidth and higher latency than terrestrial links, so traffic that moves to the satellite path can create congestion and degrade critical applications.

This study therefore combines satellite-terrestrial integrated backhaul with adaptive control in the 5G Core (5GC), the user plane, and the network-slicing management layer. The approach addresses three complementary aspects of resilience: connection-establishment resilience during simultaneous access, priority-data resilience under user-plane contention, and application-service resilience under constrained backhaul capacity.

The University of Tokyo conducts this work under the NICT-commissioned project “Research and Development of Satellite-Terrestrial Integration Technology for Beyond 5G” (No. 21901). The project develops Local 5G backhaul technologies that use satellite communication as an alternative path when terrestrial connectivity becomes unavailable or insufficient.

2. Local 5G Resilience Framework

Figure 1 illustrates the basic resilience concept. A Local 5G system normally uses a terrestrial backhaul path and switches essential traffic to a satellite path when the terrestrial path fails. This path diversity preserves connectivity, while network slicing and adaptive resource control determine which services continue to receive resources when the backup path has limited capacity.

Figure 2 shows the integrated testbed and the resilience-control sequence. The system monitors available links, detects a fault, switches the backhaul route, allocates routes according to application requirements, and adjusts QoS according to the available backhaul capacity. Together, these functions form a resilience loop that detects a change, adapts the network, and

■ **Figure 1: Satellite backup path for resilient Local 5G**

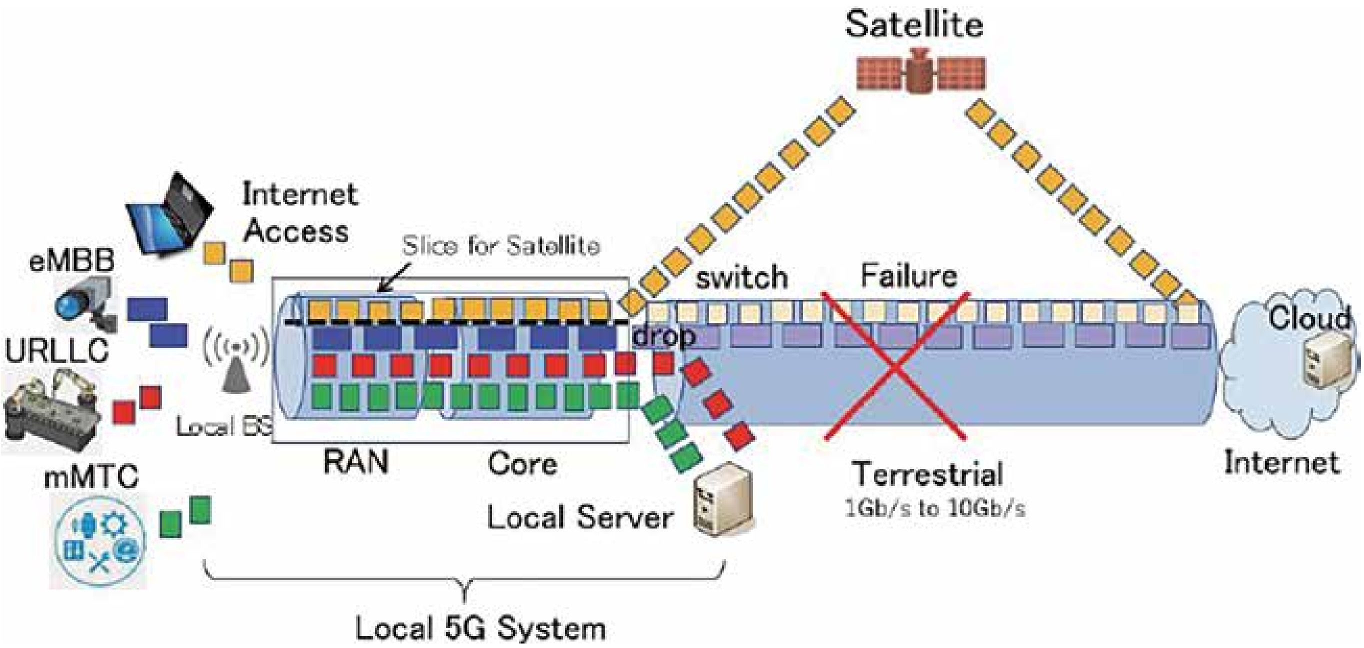

■ Figure 2 : Resilience architecture for satellite-terrestrial integrated Local 5G

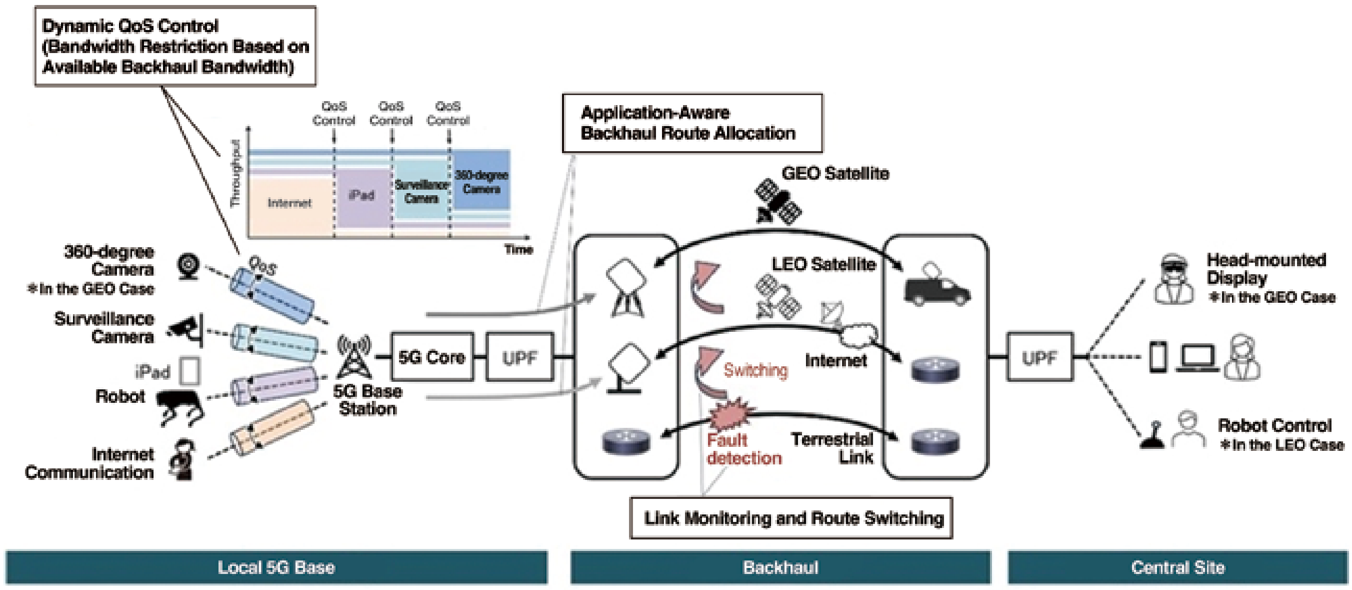

protects priority services.

To make the purpose and evaluation directly correspond, this study defines the following three resilience objectives:

Objective	Resilience objective	Control mechanism	Evaluation focus
Objective 1	Connection-establishment resilience	Prioritize 5GC control-plane network functions during simultaneous UE access.	Average number of UEs that establish a connection.
Objective 2	Priority-data resilience	Allocate user-plane processing resources to critical IoT slices under competing traffic.	Number of MQTT messages transmitted by each sensor class.
Objective 3	Application-service resilience	Apply dynamic 5QI, GFBR, and MFBR control when backhaul capacity is constrained.	Continuity and quality of the priority camera service.

The integrated testbed includes terrestrial and real satellite links. The results below focus on the control and resource-management mechanisms that sustain essential services after network conditions change. End-to-end failover metrics, such as switchover time and service-interruption duration, remain outside the present evaluation scope.

3. Resilience Technologies and Evaluation

3.1 Control-Plane Resilience during Simultaneous Access

A disaster or backhaul outage can cause many devices to reconnect at nearly the same time. The resulting signaling surge can congest the 5GC control plane before user equipment (UE) establishes a Packet Data Unit (PDU) session. This congestion prevents devices from regaining network access precisely when communication becomes most important.

To improve connection-establishment resilience, we apply a control-plane resource-prioritization method, referred to here as control-plane slicing, to selected 5GC network functions (NFs). The method assigns different CPU scheduling priorities to the Network Repository Function (NRF) and the Access and Mobility Management Function (AMF), where signaling processing can become a bottleneck. Figure 3 shows the evaluation environment based on free5GC and UERANSIM.

The evaluation activates 70 UEs simultaneously and compares the average number of successful connections under different NF priority settings. The Linux nice command controls CPU scheduling priority. Nice values range from -20 to 19, and a lower value gives a process a higher scheduling priority.

Without priority allocation, an average of 12.4 UEs establishes a connection. The average rises to 29.75 when the NRF receives higher priority, 20.9 when the AMF receives higher priority, and 24.45 when the NRF and AMF receive equal priority. NRF prioritization produces the best result and reaches approximately 2.4 times the baseline.

These results show that targeted control-plane resource allocation improves the ability of Local 5G to absorb a simultaneous-access surge and reconnect more devices. The best configuration does not connect all 70 UEs within the measurement interval, so the result represents a component-level resilience improvement rather than complete recovery. Further evaluation examines scaling behavior, connection-completion time, and the interaction among additional control-plane NFs.

■ Figure 3: Evaluation environment for control-plane resilience using free5GC and UERANSIM

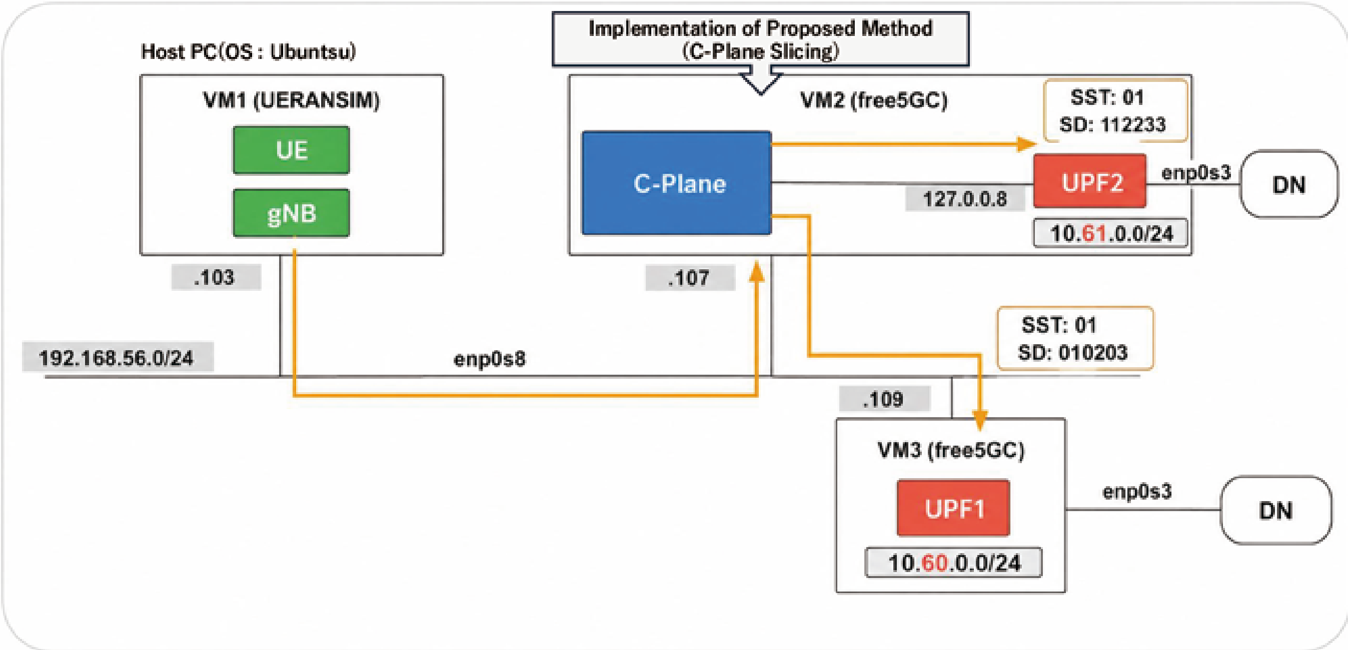

■ Figure 4: Average number of connected UEs under different NF CPU-priority configurations

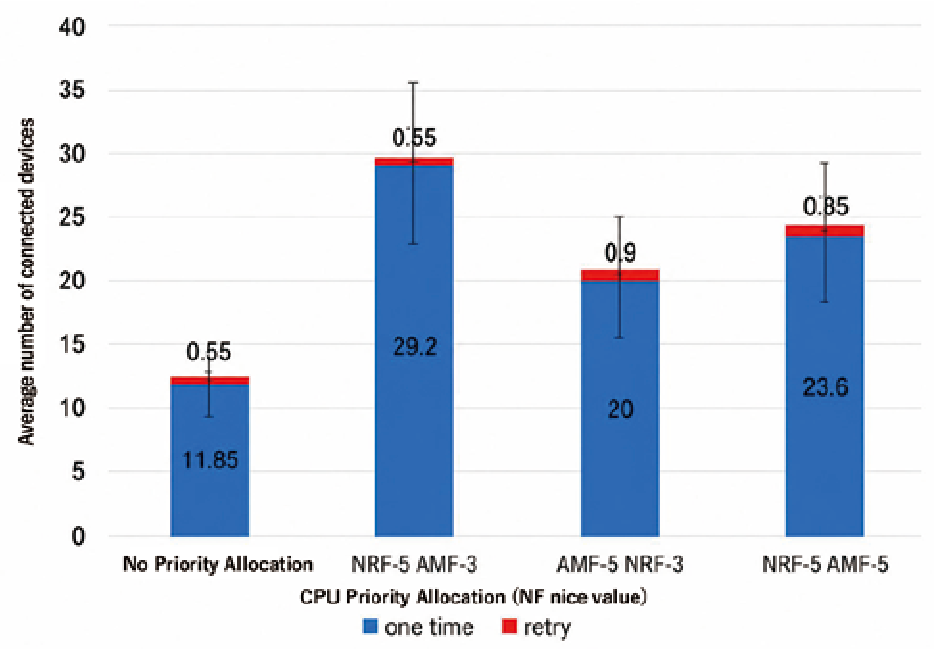

3.2 User-Plane Resilience for Priority IoT Traffic

After devices establish a connection, Local 5G resilience also depends on the continued delivery of critical data. Disaster-monitoring systems commonly use Message Queuing Telemetry Transport (MQTT), a lightweight messaging protocol for IoT applications. When multiple sensor classes share the same user-plane resources, high-volume or non-critical traffic can reduce the processing capacity available to urgent monitoring data.

The proposed network-slicing method assigns separate user-plane processing resources to pressure-sensor and seismometer

■ Figure 5: Slice-aware UPF resource allocation for resilient MQTT communication

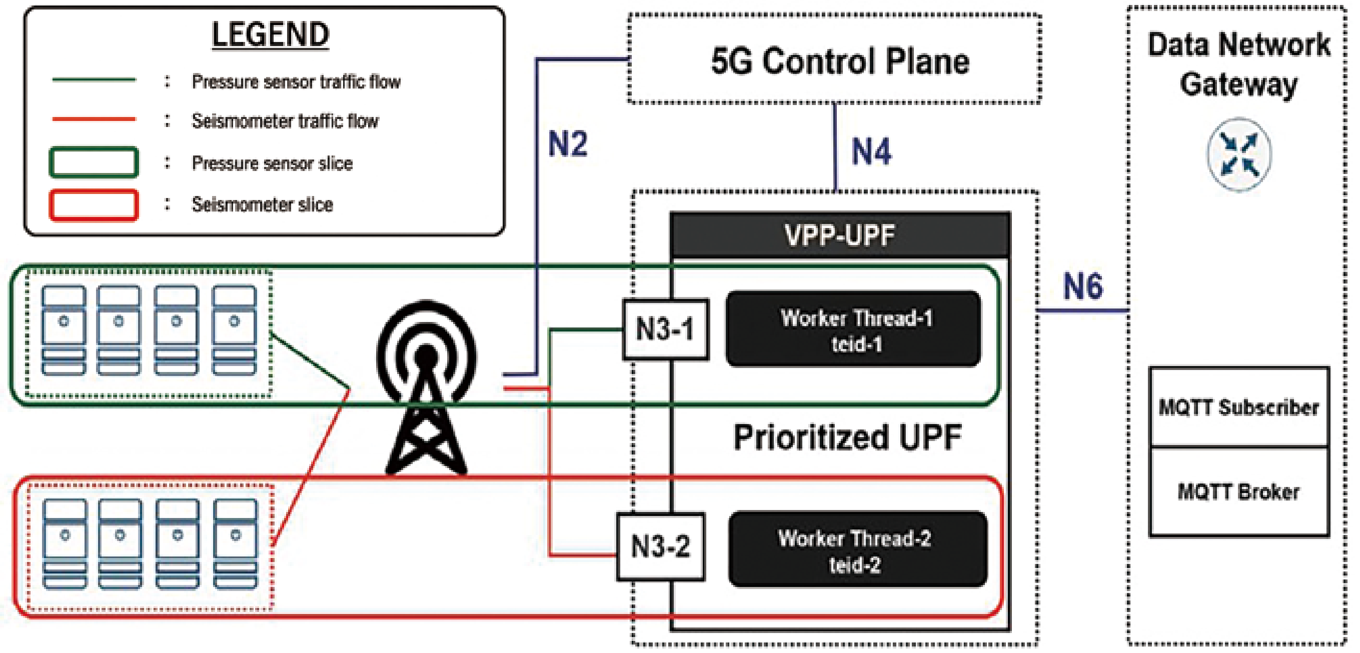

traffic and gives the critical seismometer slice higher processing priority in the User Plane Function (UPF). Figure 5 illustrates the slice-aware UPF configuration, and Figure 6 compares the baseline and proposed methods at MQTT QoS levels 0, 1, and 2.

In the demonstration, seismometers transmit data at a high rate. At every MQTT QoS level, the proposed method increases the number of transmitted seismometer messages by more than 1,000 relative to the baseline while preserving the pressure-sensor service. This result indicates that slice-aware user-plane resource allocation helps sustain priority-data throughput under contention.

The present evaluation focuses on throughput. Mission-critical resilience also depends on end-to-end delivery ratio, latency, jitter, and duplicate-message behavior, particularly at MQTT QoS levels 1 and 2. The next evaluation stage measures these indicators and

■ Figure 6: Total number of MQTT messages transmitted with the baseline and proposed methods

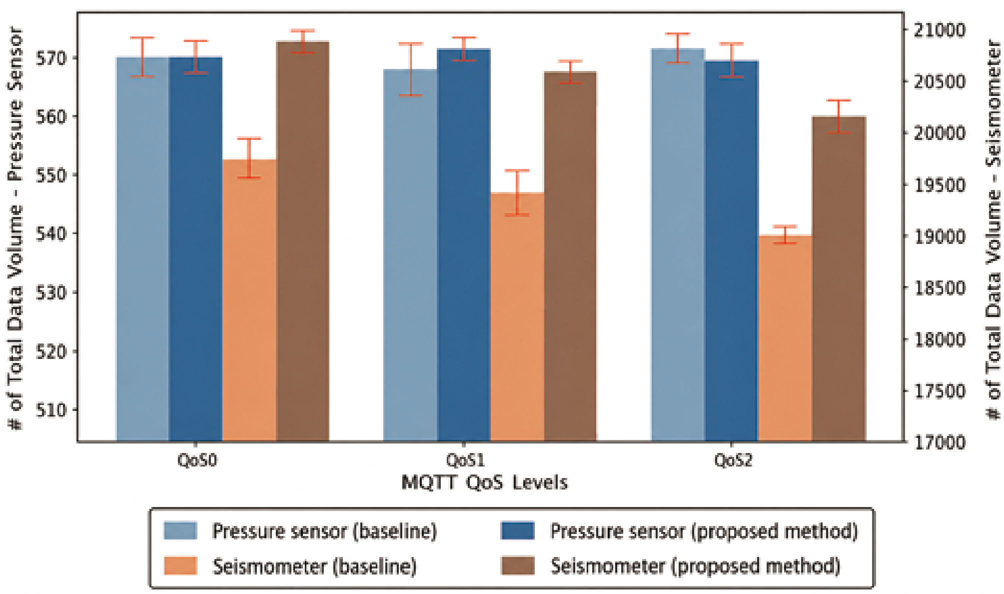

■ Figure 7: Dynamic QoS control for priority-service continuity under constrained backhaul

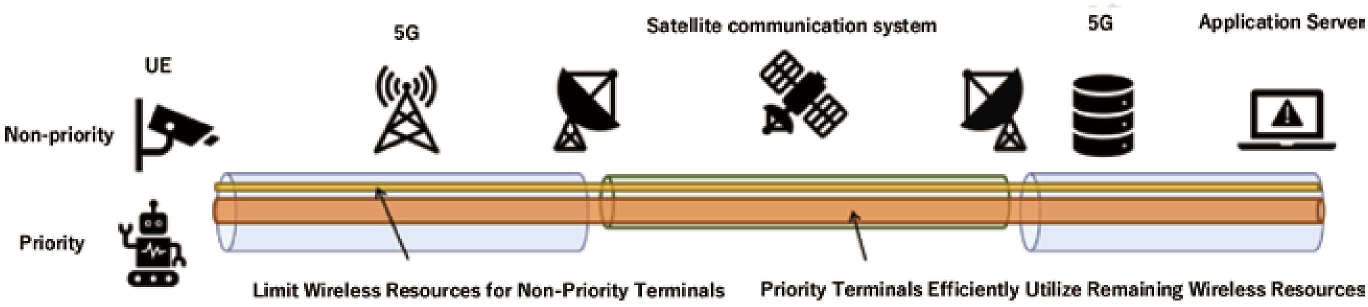

examines whether the method maintains stable performance when satellite-link delay and loss vary.

3.3 Application-Service Resilience under Constrained Backhaul

When traffic moves to a bandwidth-constrained backhaul path, priority applications can still experience video degradation or freezing even though basic connectivity remains available. Application-service resilience therefore requires the network to adapt resource allocation according to the importance of each service and the currently available capacity.

The proposed management method uses network slicing and 5G QoS control. Each application receives a Single Network Slice Selection Assistance Information (S-NSSAI) value and a 5G QoS Identifier (5QI). QoS flows use either a Guaranteed Bit Rate (GBR) or non-GBR resource type. GBR flows include a Guaranteed Flow Bit Rate (GFBR) and a Maximum Flow Bit Rate (MFBR), which allow the controller to set explicit bandwidth parameters.

When a user designates a camera as a priority application, the controller applies a GBR-type QoS profile with configured GFBR and MFBR limits to non-priority traffic. This control restricts the bandwidth used by non-priority slices and preserves capacity for the priority camera. Figure 7 shows the resulting resource-control concept across the Local 5G and satellite-communication path.

In the remote disaster-monitoring demonstration, the controller responds to an on-demand request for a camera with guaranteed quality. The evaluation observes less image degradation and shorter freezing intervals after dynamic bandwidth control is applied. Because the present result is qualitative, further evaluation quantifies freeze duration, frame loss, video bitrate, latency, and objective image-quality indicators.

Operational resilience also requires timely and repeatable control. The next implementation stage automatically selects and applies network-slice policies according to application state, user priority, and measured backhaul capacity. This automation reduces dependence on manual operation and lowers the risk of human error during an emergency.

4. Conclusion

This article frames satellite-terrestrial integration as a Local 5G resilience architecture rather than as a backup-link function alone. Satellite backhaul provides path diversity, while control-plane prioritization, slice-aware user-plane processing, and dynamic QoS management help the system use limited backup capacity effectively.

The component evaluations address three directly corresponding resilience objectives. Control-plane prioritization raises the average number of connected UEs from 12.4 to 29.75 in the best configuration in our testbed. This result shows our proof of concept for control-plane slicing, but we plan to extend our idea to a larger scale experiment. Slice-aware UPF resource allocation increases transmitted seismometer messages by more than 1,000 at each MQTT QoS level. Dynamic QoS control preserves capacity for a priority camera and reduces observed image degradation and freezing intervals.

Taken together, these results support the component-level feasibility of a Local 5G system that absorbs access surges, sustains critical IoT traffic, and protects priority services when backhaul capacity becomes constrained. The current evaluation does not yet provide a complete end-to-end resilience assessment because it does not quantify route-switchover time, service-interruption duration, satellite-link latency, jitter, or packet loss.

Further work integrates automatic policy control with link monitoring and route switching and evaluates the complete resilience sequence under repeatable failure and congestion conditions. The resulting architecture provides a practical basis for robust Local 5G deployments in disaster monitoring and other environments where terrestrial connectivity is unavailable, damaged, or insufficient.