5G for the Greenhouse Horticulture

Considerations for network architecture, cloud federation, radio network coverage and energy consumption

TNO Public ⟩ TNO 2025 R13094

December 2025

ICT, Strategy & Policy
www.tno.nl
+31 88 866 00 00
info@tno.nl

TNO 2025 R13094 – December 2025

5G for the Greenhouse Horticulture

Considerations for network architecture, cloud federation, radio network coverage and energy consumption

Author(s) Sarah Lim Choi Keung, Pascal Heijnen, Stan van Nieuwamerongen, Ljupco Jorguseski, Yohan Toh, Aditya Ganesh, Anthony Pages

Classification TNO Public
Number of pages 56 (excl. front and back cover)
Number of appendices 0

Contents

List of Figures

List of Tables

Abbreviations and Terminology

5G Access Network. The part through which the user (wirelessly) communicates with the network, e.g., NG-RAN: Next Generation Radio Access Network (which is used for 5G). The interface technology that is used for 5G radio is known as NR: New Radio.

5G: 5th generation wireless technology

5GC: 5G Core Network. The part of the network that provides data routing from and to mobile subscribers, application services, and other networks.

5G RedCap: 5G Reduced Capability

5GS: 5G System

AMF: Access and Mobility Function. The function which receives requests to connect to the network, then checks with authentication network functions if the user is allowed to communicate with the network and which then maintains a signalling connection with the UE for other 5G Core network functions.

Container: An isolated environment running on a host Operating System, typically provided a share of resources with the intention of running a specific application, complete with its specific dependencies.

DNN: Data Network Name

FSPL: Free Space Path Loss

gNB: gNodeB. The wireless base stations/radio equipment that sends and receives radio signals to and from the UEs.

IaC: Infrastructure as Code. A paradigm for computer infrastructure configuration, deployment and operation using software development concepts.

K8s: Kubernetes. An open source system for the orchestration of container based applications

Liqo: An open source tool for resource and workload sharing (federation) between two or more independent Kubernetes clusters.

MetalLB: A load balancer designed to operate with standard network equipment.

MNO: Mobile Network Operator

NF: Network Function. A modular specification of a specific function that the network is required in order to engage in network traffic and communication servicing, e.g. AMF, UPF, SMF.

Open5Gs: An open source 5GS implementation

OpenStack: An open standard for cloud infrastructure deployment provided in an “as a Service” model

PDU: Packet Data Unit

RF: Radio Frequency

Slice: a virtually independent network reserved for a subscriber. Because 5G has virtualized network functions and is unbound with regards to the physical hardware, it is possible to create independent networks not just by setting up a new physical infrastructure, but also by creating multiple network functions of the same type, separating them, and reserving these for different clients or purposes. These independent virtual networks are called slices

SME: Small and Medium Enterprises

SMF: Session Management Function. The part of the 5G Core network that sets up and manages communication sessions with UEs.

S-NSSAI: Single Network Slice Selection Assistance Information. a unique identifier used in 5G networks to specify the characteristics and requirements of a single network slice, enabling service differentiation, dynamic resource allocation, and end-to-end quality of service enforcement

TAC: Tracking Area Code

Terraform: A standard for cloud infrastructure and application deployment in a programmatic manner

UE: User Equipment. Devices that use 5G to access network services.

UPF: User Plane Function. The function which routes data to and from user devices.

VM : Virtual Machine

Management Summary

The ***Do IoT voor de Glastuinbouw*** project is a collaboration between leading field labs in South Holland: the DoIoT Fieldlab with partners TUDelft TNO and MCS, the RoboValley Foundation, and the Tomatoworld Foundation. The project aims to accelerate digital innovation in the horticultural sector by enabling companies to develop, test, and demonstrate advanced sensor systems, robotics, and autonomous vehicles that rely on high-quality mobile data connections, particularly 5G. Reliable, high-capacity, and secure wireless communication is key for this transformation.

To link the work to real impact for the greenhouse sector, TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. That way the greenhouse could be managed at the location where the produce is meant to be sold.

Therefore, TNO focused on 4 areas of research to explore the options for network architecture, cloud federation, network coverage as well as investigate the impact of energy consumption.

The first research topic is on a hybrid 5G network architecture, where TNO investigated a setup that maximizes the control over the (growth) data in the greenhouse, while minimizing the effort to control the remote network. This provides insight into how an exported greenhouse could be managed remotely by a local Mobile Network Operator (MNO), while the data stays within the greenhouse.

A Non-Public Network (NPN) architecture featuring a Shared Radio Access Network (RAN) and Control Plane provides the best option, with two separate core networks deployed: a public network core and a private network core. This option creates the best of both worlds as the greenhouse provider owns the parts that hold the data of the greenhouse and at the same time the control of the network is handled by an MNO or service provider to operate the network.

The second research topic is on cloud federation, where TNO developed and tested a solution where the local cloud (data) in the greenhouse is federated with a cloud at the local MNO and a cloud within the owner company in the Netherlands and/or a public cloud. This federation provides freedom in where to store the data and where to operate on that data. Cloud federation in particular brings a demonstrable benefit of increased resilience of core network functions.

A Proof of Concept was built that shows that sensitive operational data can remain within the local greenhouse cloud, while investment in the local greenhouse cloud can be minimized by offloading less sensitive processing or backup to external clouds, while ensuring compliance and privacy.

The third research topic is on measuring the effect of plant growth on the radio transmission in a tomato greenhouse. Here, TNO opted for a setup with different paths through the rows of tomato plants and use a complete frequency range. That way the results also cover Wi-Fi signals and 4G signals.

Our measurements in the Tomatoworld greenhouse clearly demonstrate that vegetation significantly impacts wireless signal propagation. The measurements showed a 5G coverage

reduction down to 40 meters when the tomato plants are fully grown. Translated into real-life impact: in an empty greenhouse, one 5G base station can provide full coverage for a commercial-scale greenhouse (≈115 × 115 m). However, with fully grown crops, four base stations are required to maintain reliable connectivity.

The fourth research topic is related to a new simplified user device in 5G systems, called 5G RedCap, which stands for 5G Reduced Capability. Such a device is simpler with a limited capability, but at the same time uses less energy.

TNO measured the energy consumption of 5G RedCap modems to evaluate efficiency compared to traditional 5G modems. The RedCap modem consistently consumed less energy than the legacy 5G modem across all tested conditions. The RedCap modem used less than half the normalized energy compared to the 5G modem. Only a modest increase in energy consumption was observed as signal quality decreased, which aligns with expectations.

As technology matures and enhanced RedCap (eRedCap) modules become commercially available, further improvements in power consumption are expected, potentially making RedCap a strong candidate for energy-constrained IoT applications.

1 Introduction

The ***Do IoT voor de Glastuinbouw*** project is a collaboration between leading field labs in South Holland—Do IoT Fieldlab, RoboHouse, and Tomatoworld. In implementing the project, TU Delft (as founder and manager of the Do IoT Fieldlab) is collaborating with partners TNO and MCS. The project aims to accelerate digital innovation in the horticultural sector by enabling companies to develop, test, and demonstrate advanced sensor systems, robotics, and autonomous vehicles that rely on high-quality mobile data connections, particularly 5G.

The Dutch agriculture and food sector plays a vital economic and societal role, contributing to almost 10% of the country's national income and employment[1]. However, the sector faces increasing challenges, including international competition, labor shortages, rising energy and resource costs, and growing sustainability demands. Digital transformation is essential to maintain and strengthen the sector's global leadership.

Reliable, high-capacity, and secure wireless communication is key for this transformation. Yet, greenhouses present a challenging environment for wireless technologies due to the presence of biomass, steel structures, and glass surfaces. 5G technology is as a promising solution to overcome these limitations and support data-driven cultivation, quality control, and autonomous operations.

To link the work to real impact for the greenhouse sector, TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. That way the greenhouse could be managed at the location where the produce is meant to be sold.

1.1 Objectives

The main objective of the project is to support the innovative small and medium enterprises in the horticulture sector in developing products and services enabled by high-quality connectivity, and to raise awareness among growers and other horticultural end-users and supply chain partners about the opportunities and best practices in applying these innovations.

To achieve this, the project focused on three main activities:

1. Further development of testing and experimentation environments for 5G in the fieldlabs.
2. Perform innovation projects with SMEs in greenhouse horticulture to develop and test new innovations.
3. Activating and connecting the greenhouse horticulture sector to increase the regional impact of 5G technology.

In the following sections, we outline the activities and outcomes achieved by TNO in these three focus areas. The remainder of the document then goes into detail on the research and experimentation for 5G in Chapters 2-5. Finally, in Chapter 6, we provide some conclusions and main takeaways.

[1] https://dutchindustry.org/5/

1.2 Research into further development of testing and experimentation environments for 5G in the fieldlabs

To link the work to real impact for the greenhouse sector, TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. That way the greenhouse could be managed at the location where the produce is meant to be sold. An exportable greenhouse not only consists of the building structure, but next to power, watering system, heating/cooling, it also includes connectivity and control over the data.

Therefore, TNO focused on 4 areas of research to explore the options for network architecture, cloud-federation, network coverage and also investigate the impact of energy consumption.

The premise for the horticulture sector was that the digital transformation is needed to enable more robots in the greenhouses and to go to a growth as a service, specifically remotely. The latter means that a greenhouse can be exported fully equipped to the countries that buy the produce like vegetables and flowers. That means that next to watering, heating lightning, etc. there is connectivity. And then instead of a several Wi-Fi systems, one for each robotics provider, a single 5G network for the connectivity is provided in the greenhouse. In 5G, provisions are featured to ensure Quality of Service for the individual applications.

In the first research topic on a hybrid 5G network architecture, TNO investigated a setup that would maximize the control over the (growth) data in the greenhouse, while minimizing the effort to control the remote network. This provides insight into how an exported greenhouse could be managed remotely by a local MNO, while the data stays within the greenhouse.

The second research topic was on cloud federation, where TNO developed and tested a solution where the local cloud (data) in the greenhouse is federated with a cloud at the local MNO and a cloud within the owner company in the Netherlands and/or a public cloud. This federation provides freedom in where to store the data and where to operate on that data.

The third research topic was on measuring the effect of plant growth on the radio transmission in a tomato greenhouse. Here TNO opted for a setup with different paths through the rows of tomato plants and use a complete frequency range. That way the results also cover Wi-Fi signals and 4G signals.

The final research topic is related to a new simplified user device in 5G systems. It is called 5G RedCap and stands for 5G Reduced Capability. Such device can be simpler and with a limited capability, but at the same time less energy use. Especially for IoT devices this is needed for long term battery operations. TNO created a setup to measure the energy consumption of such 5G RedCap device. Although only a pre-production device was available, we could see a halving of the energy usage compared to normal operations.

1.3 Innovation and demonstration projects with SMEs

In the project several deep dives were done as starting point for SMEs to figure out what their needs are in relation to 5G and what could be analysed and tested in the Tomatoworld greenhouse. TNO participated in these deep dive sessions for all use cases:

- **Plense Technologies** in April 2024 – to empower breeders and growers with sensor data-driven insights on their crops.
- **PATS** in September 2024 – to investigate 5G-related aspects relevant to the PATS-C system, which captures infrared images of flying insects to track the pest population. PATS focuses on utilizing camera and drone technology to monitor and eliminate flying pests in greenhouses.
- **SenseNL** in September 2024 – to investigate 5G module in place of LoRaWAN for wireless communication for their CARA MET sensors, which provides highly accurate insight into the plant's root zone for precise irrigation.
- **The Hague University of Applied Sciences** and **Koppert** in October 2024 – to investigate the effects of 5G (radiofrequency electromagnetic field exposure) on bumblebees and their behaviour in the greenhouse.
- **Mizusense** in September 2025 - the pilot validates soil-moisture sensors in a hybrid IoT connectivity context (BLE Mesh, LoRa, 5G). Mizusense will compare reliability, minute-scale latency, power consumption, and operational feasibility in a greenhouse environment (including Tomatoworld). It also investigates AI use in the Edge of the 5G network.
- **OpusWatch** in November 2025 - to investigate what connectivity with 5G would mean for the watch. OpusWatch can track harvesting of flowers in a greenhouse. To build a more scalable and future-proof communication foundation, OpusWatch aims to validate the use of a private network (first 4G then 5G).

1.4 Activating and connecting the greenhouse horticulture sector

TNO participated in several outreach and dissemination activities, aiming to activate and connect the greenhouse horticulture sector, in addition to communication about the research work performed:

- Participation of Pascal Heijnen in a panel discussion on 9 November 2023 at the opening of a new 5G test facility at Tomatoworld.
- Radio interview[2] on Hortibiz news radio in February 2024 with Aart van de Bos (Tomatoworld) and Pascal Heijnen on the project.
- Participation of team members in the opening of a new 5G test facility at RoboHouse in June 2024.
- Pascal Heijnen gave a presentation “5G Research in DoIoT in de Glastuinbouw – Project at TomatoWorld” at the High Tech in Horticulture & Agriculture event on 24 March 2025, Delft, in the Session “5G in the Greenhouse: The Future of Wireless Communications”. At the event, Stan van Nieuwamerongen showcased the 5G RedCap setup in the exhibition area of the event.
- Stan van Nieuwamerongen gave a presentation “TomatoWorld – 5G mogelijkheden en onderzoeksresultaten” at the Smart Robots in the Greenhouse – Knowledge Sharing Event on 30 October 2025, in Honselersdijk.
- Pascal Heijnen attended the Tomatoworld event for growers on 4 December 2025 together with the other partners in the project.

[2] https://goedemorgen.podbean.com/e/tno-en-tomatoworld-over-het-nieuwe-5g-ketenwatchen-special-hortibiz-newsradio/

2 5G Network Architecture Options

2.1 Introduction

As an introduction to the topic of Non-Public Network architectures, we provide in this section some excerpts from the TNO report TNO-2023-10276 "5G Non-Public Network Architectures : An overview of deployment options and implementation considerations for enterprises"[3]:

The 5G network is designed for three key service types:

- **Ultra-Reliable Low-Latency Communication**: important for mission critical applications such as critical infrastructure, factory automation, robotics, or self-driving cars.
- **Massive Machine-Type Communication**: important for Internet of Things applications such as sensor networks, self-driving cars, and monitoring and control services.
- **Enhanced Mobile Broad Band**: important for mobile user experience-focused applications such as streaming and virtual reality, and real-time monitoring services.

To accomplish this, the 5G system is designed in various parts, namely Radio (e.g. antenna spectrum) and Core (network functions e.g. authentication, billing, and sending data). An important part of the Core design is that these functions are virtualized and can be run on most generic processing devices, such as (cloud) servers.

2.1.1 The basic 5G architecture

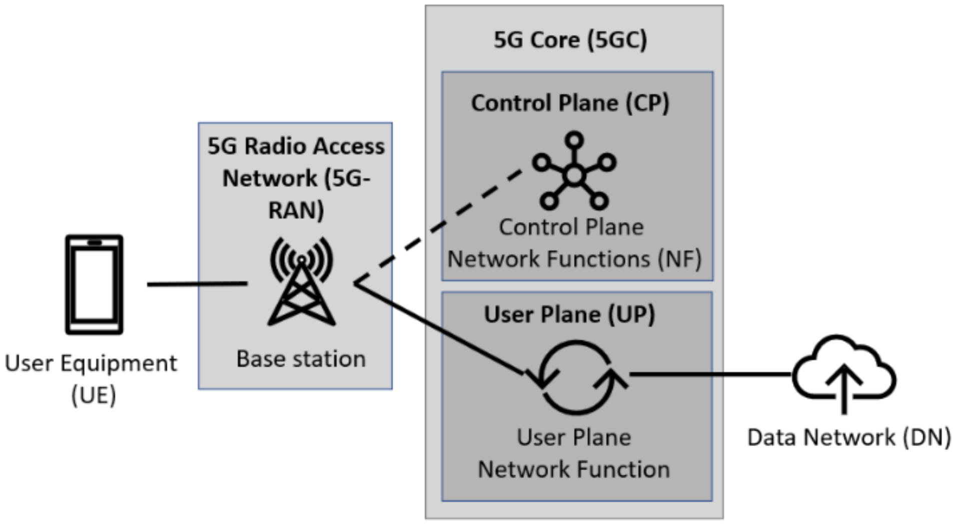

Figure 2.1: The basic 5G architecture

[3] Bergman, W.A.L. *et al.* 5G Non-Public Network Architectures: An overview of deployment options and implementation considerations for enterprises, 2023, https://publications.tno.nl/publication/34640592/a6yjJe/TNO-2023-R10276.pdf

The basic architecture of a 5G network consists of 3 sections: UE, Core and Radio (**Figure 2.1**).

UE: you will need *User Equipment able to connect* to the network, i.e., equipment that has a *5G capability* and a *SIM with a subscription to the network*.

5GC: you will need access to a running *5G-Core*. The Core network is made up of Network Functions, which each fulfil a role in the network's operation. These can be separated into the Control Plan (for network-internal tasks like authentication) and User Plane (for transmitting user data). There are various implementations of 5G-Core, some of which are open source. This core will need to be hosted on a *server*, usually running on a physical machine or a virtual machine on a cloud server.

RAN: access to an operational *5G-Access Network such as a radio base station*, and *usage plans to the electromagnetic spectrum bands* which it will *use (and, if applicable, paid-for usage licenses)*.

2.1.2 Four possible non-public network architectures

To provide a private network, four basic levels of non-public 5G network architectures (private deployment vs. various public integrations) or are possible:

1. **Standalone**: fully independent private network managed by the enterprise.
2. **Shared RAN**: independent Core but sharing the Radio network with a service provider.
3. **Shared RAN and CP**: independent User Plane so enterprise data traffic stays within the enterprise but sharing the Radio network and Control functions with a service provider.
4. **Slicing**: fully dependent on a public network, in which both Core and RAN are shared, and a *virtually* independent 'network slice' is set up for secure, private use by the enterprise.

For a visual diagram of these architectures, see **Figure 2.2** below.

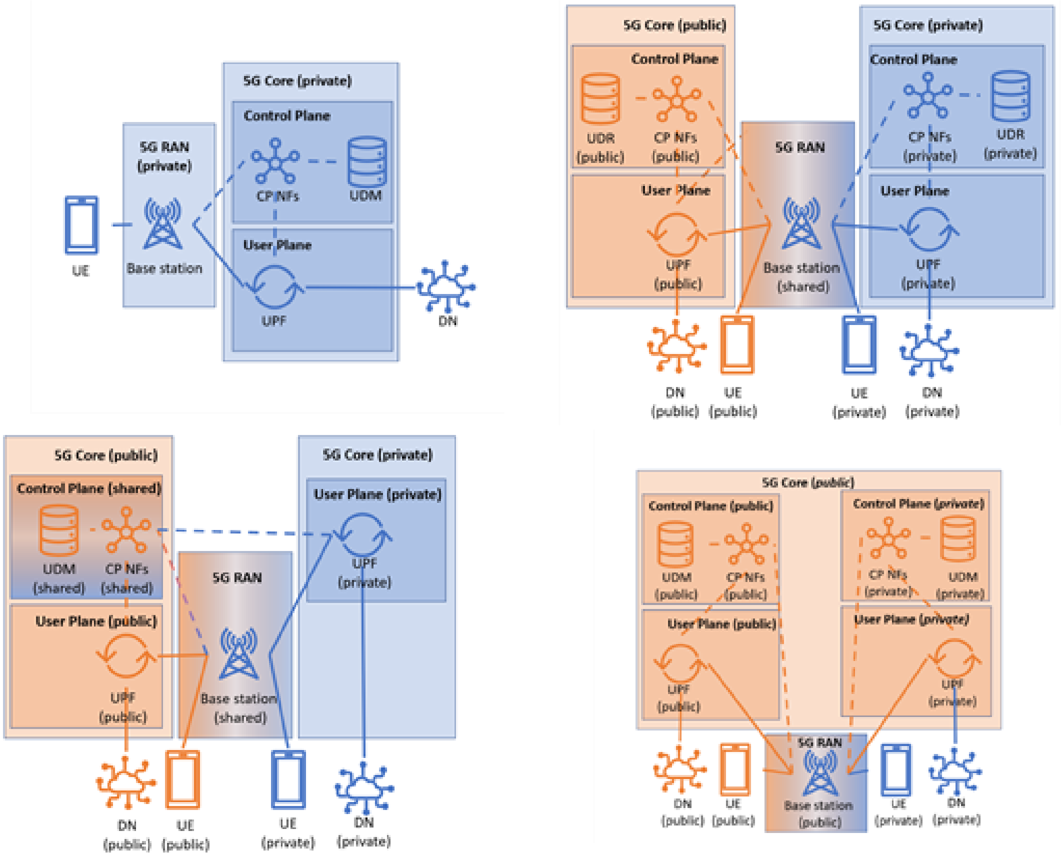

Figure 2.2: The four non-public network options, from Top-Left (TL) to Bottom-Right (BR): 1. Standalone (TL), 2. Shared Radio Access Network (TR), 3. Shared Radio Access Network and Control Plane (BL), 4. Network Slice (BR)

2.2 Network setup overview

For this project, TNO implemented a **Non-Public Network (NPN) architecture featuring a Shared Radio Access Network (RAN) and Control Plane**, corresponding to option 3 described previously. In this setup (Figure 2.3), two separate core networks are deployed: a public network core with UPF1 and a private network core with UPF2. The public network is responsible for providing the Control Plane, including essential network functions such as the AMF and the SMF, while the 5G RAN is shared between both the public and private networks.

The AMF manages the registration and authentication of 5G UE, ensuring that only authorized devices can access network services. The SMF is responsible for establishing, maintaining, and releasing user sessions, as well as coordinating the allocation, modification, and release of IP addresses and other session parameters for the UEs. Importantly, the SMF also selects the appropriate UPF to handle each user's data session, applying the necessary QoS parameters to meet application requirements.

This architecture enables flexible and secure connectivity by allowing both public and private data traffic to coexist on the same radio infrastructure, while keeping sensitive operational data within the private network.

This option creates the best of both worlds as the greenhouse provider owns the parts that hold the data of the greenhouse and at the same time the control of the network is handled by an MNO or service provider to operate the network.

Several UPF selection methods and their combinations were tested within this setup, providing valuable insights into how different approaches can optimize network performance, security, and resource utilization for the greenhouse environment.

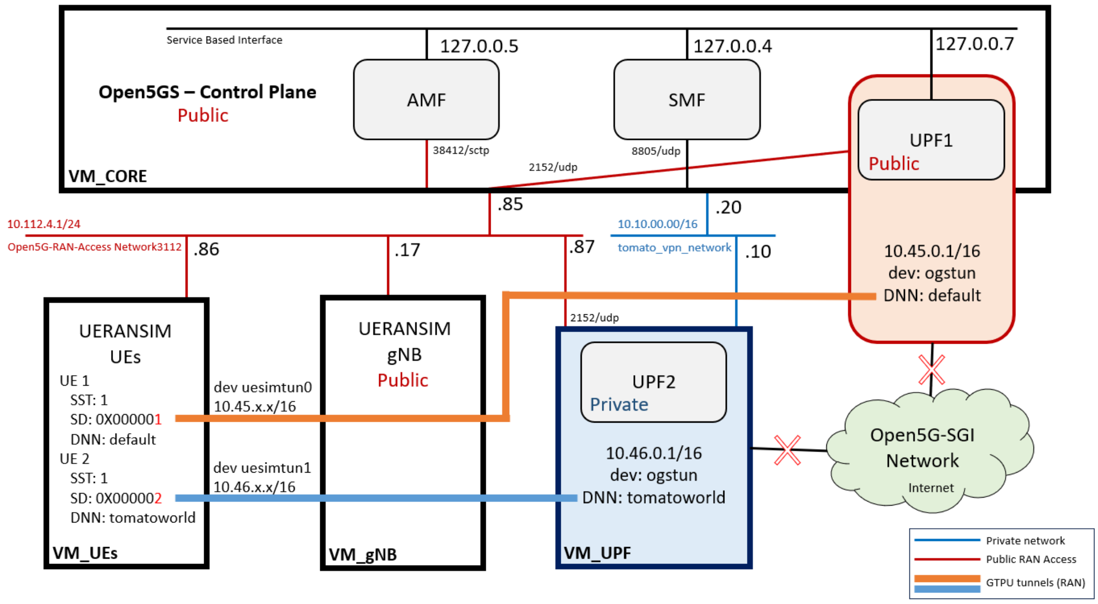

Figure 2.3: Architecture deployment: with two separate core networks: a public network core with UPF1 and a private network core with UPF2

2.3 Possible setup options

The UPF is central to establishing the user plane in a 5G network, which begins when the network activates a PDU session. A PDU session is a logical connection between UE and a data network, such as the internet or a private greenhouse network. The SMF, working in coordination with the AMF and the gNB, is responsible for assigning IP addresses, creating tunnels, and initiating data flow. The way in which the UPF is selected and configured directly impacts how efficiently and securely different devices and applications in the greenhouse can access network resources, making it a key consideration for supporting diverse operational needs.

2.3.1 One SMF, multiple UPFs

Single PDU session

In this configuration (**Figure 2.4**), a single PDU session is established, allowing the UE to connect to either the public or the private UPF based on the subscription profile or user identity. This setup is particularly valuable for scenarios where certain devices, such as sensors or automation systems, require exclusive access to the private network for reasons of security, data integrity, or operational control. At the same time, it enables straightforward segregation, so that UEs belonging to visitors, contractors, or temporary workers can be restricted to the public UPF, ensuring that their traffic remains isolated from sensitive greenhouse operations. This clear separation of network domains supports compliance, enhances security, and simplifies network management for greenhouse operators.

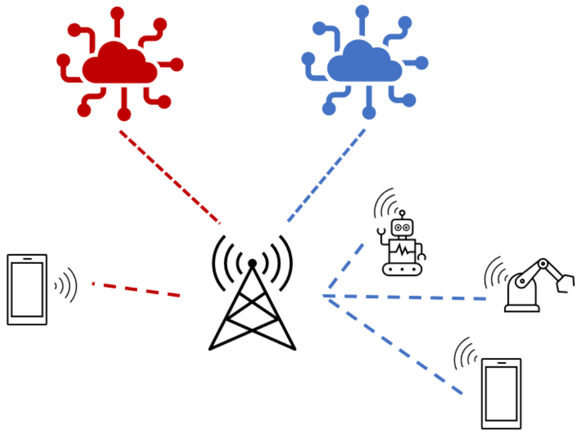

Figure 2.4: Single PDU session, allowing the UE to connect to either the public or the private UPF based on the subscription profile or user identity

Simultaneous PDU sessions

Alternatively, the UE can establish concurrent PDU sessions with both the public and private UPFs, enabling simultaneous data transmission and reception across multiple network domains (**Figure 2.5**). This dual connectivity is highly advantageous for administrative or supervisory UEs that need to oversee and coordinate activities spanning both public and private networks. By maintaining parallel connections, administrators can efficiently monitor and manage network resources, perform diagnostics, and ensure seamless communication and data flow between different segments of the greenhouse infrastructure. This capability not only enhances operational flexibility and responsiveness but also empowers administrators to optimize network performance, address issues proactively, and support evolving digital applications within the greenhouse environment.

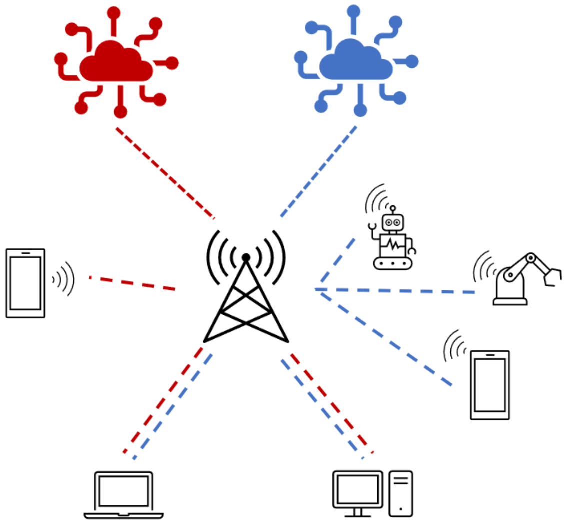

Figure 2.5: Concurrent PDU sessions with both the public and private UPFs, enabling simultaneous data transmission and reception across multiple network domains

2.3.2 UPF selection

The selection of the UPF, whether public UPF1 or private UPF2, can be achieved in multiple ways, each offering different levels of flexibility, security, and performance to address the diverse connectivity needs within the greenhouse environment.

Select UPF based on slicing (S-NSSAI)
In this method, the selection of the UPF is determined by the Single Network Slice Selection Assistance Information (S-NSSAI), which uniquely identifies a specific network slice within the 5G system. Each slice is configured to meet particular service requirements, such as low latency for robotics, high throughput for video monitoring, or enhanced reliability for critical control systems. By associating each S-NSSAI with a dedicated UPF, the network can direct user traffic to the most appropriate data path, ensuring that the performance characteristics of each slice are maintained.

This approach enables precise traffic separation and resource allocation, allowing the greenhouse network to simultaneously support diverse applications with different connectivity needs. For example, sensitive operational data from autonomous vehicles or environmental sensors can be routed through a private, high-security slice, while guest or visitor devices may be assigned to a separate, less critical slice. As a result, the system can optimize network performance, enhance security, and ensure that critical horticultural processes receive the necessary quality of service, even as the number and type of connected devices evolve.

Select UPF based on DNN
In the DNN-based approach, the UE can access multiple UPFs depending on the DNN it requests. Each DNN represents a specific data network or service domain, such as a private greenhouse management network, an external analytics platform, or a public internet connection. By associating each DNN with a dedicated UPF, the network can dynamically route user traffic according to the intended data network context.

This mechanism allows the gNB and the 5G core to flexibly select the most appropriate UPF for each session, based on the DNN specified by the UE. As a result, different types of devices and applications, such as environmental sensors, autonomous robots, or guest devices, can be directed to the correct data network, each with its own security policies, quality of service, and access controls. For example, sensitive operational data can be kept within a private UPF, while less critical or external communications are routed through a public UPF.

This approach not only accommodates diverse service requirements and network configurations efficiently, but also enhances security and operational flexibility. It enables greenhouse operators to segment network traffic according to business needs, support multiple service providers or partners, and adapt to evolving digital applications without major changes to the underlying infrastructure. As new services or data networks are introduced, additional DNNs and corresponding UPFs can be provisioned, ensuring the network remains scalable and future-ready.

Select nearby UPF according to the connected gNB based on TAC
In this method, the UPF is selected based on the Tracking Area Code (TAC) associated with the connected gNB. Each gNB is linked to a specific UPF, and the system uses TAC information to dynamically assign the nearest UPF to the user equipment. This approach helps minimize latency by routing data through the closest available UPF, but is less flexible and scalable compared to slicing or DNN-based selection, making it less suitable for complex or evolving greenhouse network environments.

2.3.3 Conclusion

In 5G networks, it is possible to combine different types of UPF selection methods, such as S-NSSAI slicing and DNN based selection. This combined approach allows for a more granular and flexible allocation of resources, catering to specific service requirements while leveraging network context information for efficient UPF assignment.

The adoption of a 5G network architecture with advanced options for UPF selection offers greenhouse owners and growers a robust, flexible, and future-proof connectivity solution. This architecture enables them to securely separate and manage different types of data traffic, ensuring that sensitive cultivation and business data remains protected within the greenhouse while still allowing for remote management and support. It also improves operational efficiency and reliability, as critical applications like robotics, sensors, and autonomous vehicles can be prioritized and guaranteed high-quality service. By consolidating connectivity needs into a single 5G infrastructure, growers can reduce costs and simplify network management, replacing multiple vendor-specific Wi-Fi systems.

Furthermore, the architecture is scalable and adaptable, making it easier to expand operations, integrate new technologies, and support ongoing innovation. Overall, these choices empower greenhouse operators to accelerate digital transformation, enhance productivity, and maintain a competitive edge in a rapidly evolving horticultural sector.

2.4 Experimentation

The experimentation phase of this project focused on validating the functionality of a shared public 5G network with private network capabilities, leveraging a setup with a 5G core network, multiple UPFs, and various access points. Our goal was to evaluate the network's ability to handle connections from multiple commercial devices, as well as to examine the effectiveness of different selection criteria for access point allocation, such as network slicing.

2.4.1 Network setup and device configuration

In the laboratory, we established a 5G core network connected to a gNB, enabling direct connection testing. The network configuration included multiple UPFs, which allowed us to simulate multiple access points within the network. A commercial smartphone was used to validate network connectivity, demonstrating the network's ability to seamlessly connect to various UPFs based on the selection criteria.

2.4.2 Testing access point selection criteria

Our testing covered multiple aspects:

1. Slicing Capability:
   - We tested different types of network slicing configurations, ensuring compatibility with various devices. This included evaluating how Quectel and Fibocom modules handled these slices, as both are known for supporting slicing through URSP.
   - We found that Amarisoft offers support for URSP, which enabled flexible traffic management by allowing devices to be routed to specific network slices based on URSP policies.

2. Device Support for URSP:
   - Our testing included both older and newer models of Quectel modules, as well as devices from Fibocom. While newer Quectel devices showed robust support for URSP, enabling effective slice selection, older models demonstrated limited functionality in this regard.
   - Both the Quectel and Fibocom modules showed satisfactory support for slicing within the 5G network, particularly when configured with URSP.

3. Performance Testing with Commercial Smartphones:
   - Using a smartphone, we validated the network's capability to dynamically switch between different UPFs based on predefined criteria. This demonstrated the network's flexibility in handling diverse UE and aligning with its connectivity requirements.

2.4.3 Observations and limitations

During the testing process, we observed that while most devices were able to support slicing and URSP configurations, the extent of this support varied between manufacturers and device models:

- Quectel devices: The latest Quectel modules showed improved URSP support, which facilitated seamless network selection. However, some older Quectel models lacked full compatibility with the advanced slicing features, which could limit their use in more complex network environments.
- Fibocom devices: These also demonstrated adequate slicing support, though their performance with URSP varied based on firmware and configuration settings.

2.4.4 Conclusion

The experimentation phase confirmed the practical viability of using multiple UPFs and network slicing to provide a shared public 5G network with private network functionalities in greenhouse environments. By incorporating advanced features such as URSP and device-specific slicing capabilities, the network demonstrated flexible resource allocation and optimized service delivery for different user demands. For growers and greenhouse owners, this means they can benefit from a connectivity solution that is both robust and adaptable: critical operations and sensitive data can be securely isolated within private network slices, while guest or external traffic is kept separate, all on the same shared infrastructure. This approach not only enhances security and operational efficiency but also simplifies network management and supports the integration of new digital applications as the sector evolves. Ultimately, these capabilities empower greenhouse operators to accelerate digital transformation, improve productivity, and maintain a competitive edge in a rapidly changing horticultural landscape.

2.5 Implementation requirements

Implementing a shared public 5G network with private capabilities involves several key requirements:

Functional Requirements

- *Provider Collaboration*: Network providers must establish agreements to share infrastructure, frequency bands, and resources, ensuring smooth operation and service continuity.
- *Infrastructure Sharing*: Efficient resource allocation and network slicing are necessary to manage shared infrastructure, enabling flexible service delivery across various user devices.

Performance Requirements

- *Low Latency and High Throughput*: The network must support high-performance requirements for real-time applications, achieved through load balancing and efficient resource allocation.
- *Quality of Service (QoS):* Network slicing should prioritize critical services, ensuring consistent quality across different types of traffic.

Interoperability Requirements

- *Device Compatibility*: Adherence to 3GPP standards ensures compatibility with diverse commercial devices and facilitates easy integration with legacy networks.
- *Cross-Network Integration*: The network must work seamlessly with other networks, supporting a smooth transition between 5G and private networks.

Security and Privacy Requirements

- *Data Protection*: Robust encryption and secure authentication are essential to safeguard user data and comply with privacy regulations.
- *Network Security*: Security measures such as firewalls and intrusion detection systems protect the network from threats, with policies to isolate traffic for enhanced security.

Scalability Requirements

- *Network Expansion*: The network should easily scale to accommodate growth, allowing for new UPFs, gNBs, and additional providers as needed.

- *Flexible Resource Allocation*: Dynamic network slicing enables efficient resource use, allowing the network to scale with demand and integrate new technologies.

2.6 Conclusion

To link the research to real impact for the greenhouse sector, TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. For the grower to stay in control it is imperative that the data stays within control of the grower operating the greenhouse.

A Non-Public Network (NPN) architecture featuring a Shared Radio Access Network (RAN) and Control Plane, provides the best option, with two separate core networks deployed: a public network core with UPF1 and a private network core with UPF2. This option creates the best of both worlds as the greenhouse provider owns the parts that hold the data of the greenhouse and at the same time the control of the network is handled by an MNO or service provider to operate the network.

In this setup with two UPFs, it is possible to combine different types of UPF selection methods, such as S-NSSAI slicing and DNN based selection. The combined approach allows for a more granular and flexible allocation of resources. The experimentation phase confirmed the practical viability of using multiple UPFs and network slicing to provide a shared public 5G network with private network functionalities in greenhouse environments. It demonstrated flexible resource allocation and optimized service delivery for different user demands.

This architecture offers greenhouse owners and growers a robust, flexible, and future-proof connectivity solution. It enables them to securely separate and manage different types of data traffic, ensuring that sensitive cultivation and business data remains protected within the greenhouse while still allowing for remote management and support. It also improves operational efficiency and reliability, as critical applications like robotics, sensors, and autonomous vehicles can be prioritized and guaranteed high-quality service, while guest or external traffic is kept separate, all on the same shared infrastructure. By consolidating connectivity needs into a single 5G infrastructure, growers can also reduce costs and simplify network management, replacing multiple vendor-specific Wi-Fi systems.

This approach not only enhances security and operational efficiency but also simplifies network management and supports the integration of new digital applications as the sector evolves. Ultimately, these capabilities empower greenhouse operators to accelerate digital transformation, improve productivity, and maintain a competitive edge in a rapidly changing horticultural landscape.

3 Cloud Federation for Private 5G Network

3.1 Introduction

The feasibility of deploying a private 5G network using cloud infrastructure techniques, specifically in a federated cloud environment, was explored through a Proof of Concept. The resultant implementation consisted of 2 distributed 5G cores, one on-premises, the other remotely situated at a TNO 5G testbed, federated using Liqo, a State of the Art container workload federation technology built on Kubernetes, an industry standard for container orchestration. Here, in Section 3.2, we present in brief the motivation behind the adoption of cloud federation and associated techniques from traditional cloud computing for private 5G. In Section 3.3, the implementation process followed for the Proof of Concept is explained in detail.

3.2 Why cloud federation for 5G Non-Public Networks?

While 5G non-public (private) networks represent a potentially viable solution for connectivity in challenging environments for connectivity such as greenhouses, deployment of the supporting mobile core infrastructure has traditionally been prohibitively expensive due to the need for specialized hardware and resultant vendor-lock-in.

Open-source projects such as Open5Gs have made significant contributions to overcoming this hurdle, by implementing 5G compliant core functionality, deployable as Virtual Network Functions on Commercial Off the Shelf (COTS) hardware, or even in public cloud infrastructure. Tools and techniques from traditional cloud computing can further lower the barrier to entry further, leveraging containerization, virtualization, and Infrastructure as Code to standardize and automate the processes of deployment and orchestration for this otherwise complex system. Cloud federation, a recent development in traditional cloud infrastructure, refers to the interconnection and collaboration of multiple independent cloud environments. This can be across different organizations, to share resources, workloads, and services in a seamless, secure, and scalable manner.

In a federated cloud infrastructure, each participating cloud (e.g. private, public, or edge cloud) retains its autonomy but agrees to common protocols and interfaces for interoperability. This enables organizations to:

- Dynamically allocate workloads across multiple clouds for resource availability, cost, or performance requirements.
- Achieve greater resilience and flexibility by avoiding vendor lock-in and enabling failover.
- Optimize data locality, which is especially important for applications requiring low latency or compliance with data sovereignty wishes or regulations.

For 5G private networks, specifically, cloud federation is a key enabler for scalability and operational flexibility, enabling the following:

- *Data Sovereignty and Security*: Sensitive operational data can remain within the local greenhouse cloud, while less sensitive processing or backup can be offloaded to external clouds, ensuring compliance and privacy.
- *Distributed 5G Core Functions*: 5G core network functions can be deployed across federated clouds, allowing critical workloads (like the User Plane Function, UPF) to run close to the edge (e.g., within the greenhouse), while control functions remain centralized or in a trusted public cloud.
- *Resource Pooling*: Federation allows seamless sharing of compute, storage, and network resources between the greenhouse's local cloud, the mobile network operator's cloud, and even public clouds. This supports dynamic scaling and efficient use of infrastructure.
- *Resilience and Flexibility*: If one cloud environment experiences issues, workloads can be migrated or replicated to another federated cloud, minimizing downtime.

3.3 Proof of concept implementation

In line with cloud deployment principles, a modular approach was considered, whereby an initial deployment was made, first on the server infrastructure contained within the greenhouse, then cloned in the remote infrastructure.

For the initial deployment, a simple distributed use case was considered with a core network and single edge network residing within the greenhouse infrastructure. Each was provisioned through OpenStack, consisting of a Kubernetes controller node and worker node, deployed as separate virtual machines. For the 5G core, a container based repackaging of Open5Gs, developed by TNO specifically for deployment on Kubernetes clusters, was used. Liqo was then used to federate the two clusters. In order to ensure reproducibility, this process was performed declaratively using Ansible.

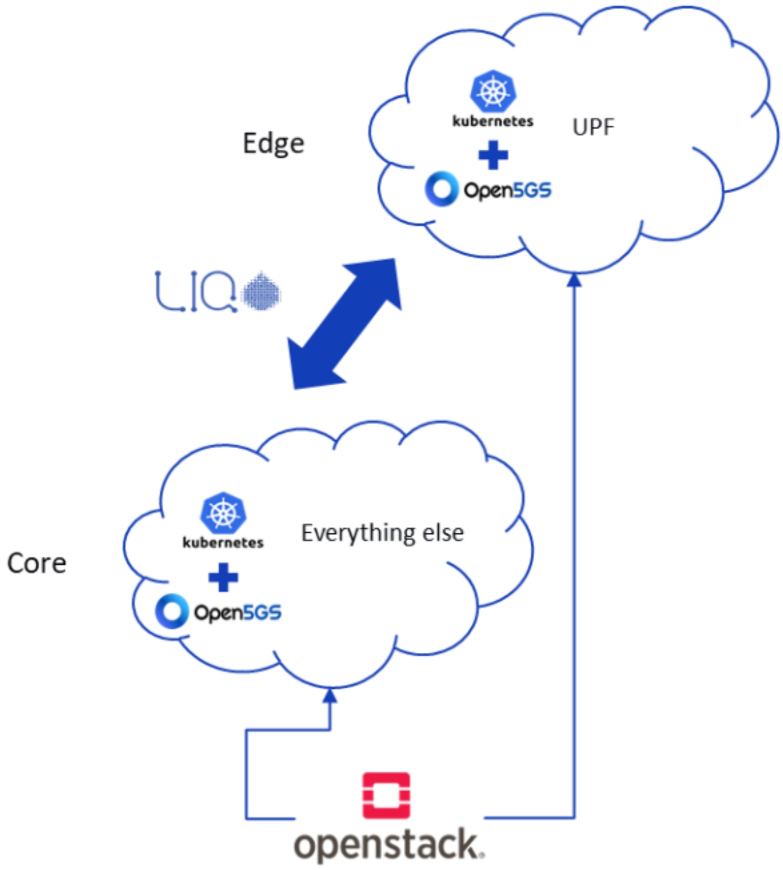

Figure 3.1: Single-site distributed Open5Gs deployment on Openstack using Liqo for cluster peering

As shown in **Figure 3.1**, the peering process (i.e. resource sharing setup) between the two clusters was performed using an out-of-band approach, whereby Ansible, relying on the Secure Shell protocol was used to share the necessary context information between the clusters.. Once peering was established, resource sharing and re-deployment was enabled unidirectionally, from the core to the edge, (i.e. it would be possible for workloads that are deployed on the core to be redeployed on the edge, but not vice versa).

The User Plane Function, responsible for routing subscriber internet data packets, was designated to run on the edge cluster, a deployment pattern proposed with 5G so as to minimize packet latency, and maintain adequate Quality of Service. All other 5G workloads (e.g. Session Management, which establishes connectivity between user devices and data networks) were designated to run on the core.

Correct functionality of the 5G core was verified using UERANSIM, an open source software emulation of 5G radio functionality, with the ability to simulate both 5G Base Station Radio Units (gNodeBs) and User Equipment.

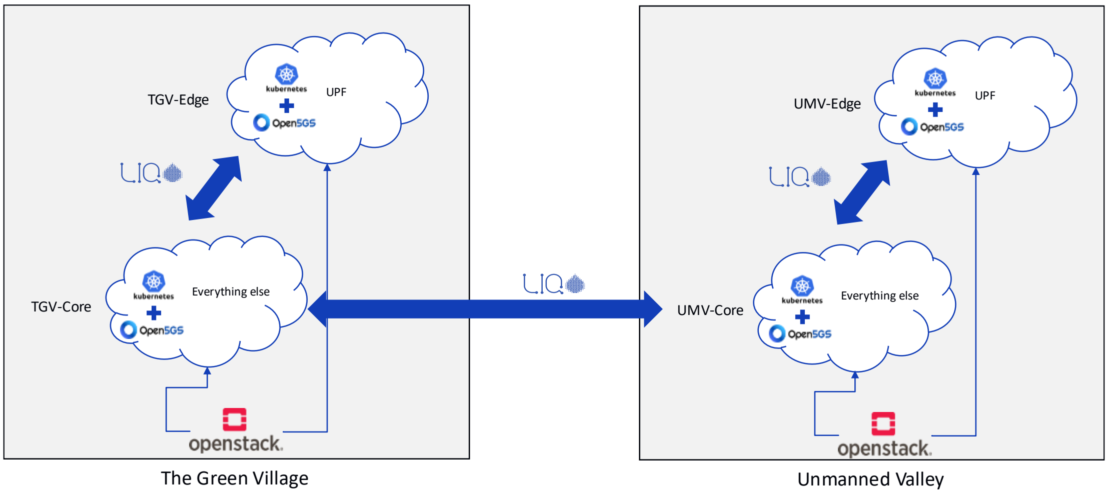

Figure 3.2: Multi site distributed Open5Gs deployment, with bidirectional offloading between sites

The greenhouse infrastructure (Openstack Virtual Machines, Kubernetes clusters, and liqo peering) and workload (Open5Gs) were then redeployed in the remote infrastructure. After instantiation, liqo was used to peer the core networks of the two sites in a bidirectional manner, allowing for either core to offload functions to the other (as shown in **Figure 3.2**). For both offloading setups (core to edge and between the cores), the requisite behaviour was verified by applying “taints” (a Kubernetes labelling system to control workload deployment) to a given cluster, preventing a specific function from being deployed on it, and then observing the failover process.

3.4 Conclusion

The distributed 5G network so created, demonstrated the feasibility of using a combination of Open Source technologies for networking and cloud computing, along with widely used, commercially available infrastructure to deploy a private 5G network. Cloud federation in particular brings a demonstrable benefit of increased resilience of core network functions. The caveat for such an implementation is that network connectivity between the two clusters is guaranteed, and that the network fabric used to provide key 5G functionality (such as radio access) must be mutually reachable between the two sites.

While the implementation created relies exclusively on Kubernetes for load balancing and routing, production 5G environments utilize the Border Gateway Protocol for communication between clusters. Kubernetes may be operated in conjunction with existing Border Gateway Protocol routers using MetalLB, a load balancer designed to operate with standard network equipment. However, this was not explored in the context of the project. Adaptation of the implementation undertaken here using MetalLB would thus constitute the most fruitful avenue for future work.

Although this was a very technical proof of concept (PoC), for the greenhouse use case the outcome is relevant. This PoC shows that sensitive operational data can remain within the local greenhouse cloud, while investment in the local greenhouse cloud can be minimized by offloading less sensitive processing or backup to external clouds, while ensuring compliance and privacy.

4 Radio Network Coverage

4.1 Introduction

To explore the feasibility of private 5G in greenhouses, a series of radio frequency (RF) path loss measurements were conducted between January and April 2024 in the Tomatoworld tomato greenhouse located in Honselersdijk, The Netherlands. The greenhouse measures 38 by 38 meters and is enclosed by glass walls.

The goal of the measurement campaign was to assess the impact of tomato plant growth on signal attenuation across frequencies ranging from 680 MHz to 6000 MHz. Measurements were performed under the following conditions:

- **Four distinct propagation paths** were selected within the greenhouse (36.6 m, 39.6 m, 48 m, and 31.2 m), each representing different distances and levels of signal obstruction.
- **Three receiver heights** were tested (1 m, 2 m, and 3 m) to reflect typical placements of wireless devices and sensors.
- **Both vertical and horizontal polarizations** were measured to capture the full range of signal behaviors.
- **Measurements were taken at three key plant growth stages:** when the greenhouse was empty, when plants were approximately 50% grown, and when plants were fully mature.

We present the methodology, results, and implications of these measurements for wireless resource planning in greenhouse environments. The central research question is:

How does plant growth in a greenhouse affect wireless signal propagation, and what are the implications for planning a private 5G network?

The remainder of this chapter is structured as follows: Section 4.2 introduces key concepts in wireless communication resource planning and the fundamentals of path loss and free space loss. Section 4.3 introduces the measurement setup and how the measurements are conducted, followed by Section 4.4, where the results of the measurements with the setup are discussed. Further, Section 4.5 describes what the results mean for the coverage of wireless networks within a greenhouse. Lastly, the overall conclusion can be found in Section 4.6.

4.2 Wireless Communication Resource Planning

4.2.1 Technical background

Wireless communication relies on the transmission of radio frequency (RF) signals between a transmitter and a receiver. As these signals propagate through space, they are subject to various physical phenomena that can degrade their strength and quality. A few concepts should be known to understand the remaining of this chapter.

Free space path loss

One of the most fundamental concepts is Free Space Path Loss (FSPL), which describes how signal strength decreases with distance and frequency, even in ideal conditions without obstacles. The formula for free space path-loss (FSPL) is:

$$FSPL = \left(\frac{(4\pi df)}{c}\right)^2$$

where d is the distance in meters, f the frequency in Hertz and c the speed of light in meter per second. This means that as the distance (or frequency) increases, the free space path loss grows proportionally to the square of that distance (frequency).

Decibel notation (dB)

Because of this quadratic relation in free space loss, the difference between the transmit and the receive signal strength value can be huge. For example, when transmitting with 10 watts of power at a frequency of 1 GHz, the receive power at 75 meters is roughly $1x10^{-6}$ watts. To make calculating with these quadratic relations, the decibel notation p_{dB} is often used where:

$$p_{dB} = \log_{10} p$$

Where p is the power in Watt. This means that the earlier described power value $1x10^{-6}$ can be written a -60dB watt (or dBW for short). So this means every 10 dB increase results in 10 times the power and visa versa. Lastly sometimes the term dBm is used, which is short for dB milli-watt.

$$p_{dBm} = p_{dBW} + 30$$

Environmental Effects

In real-world environments, additional factors come into play:

- **Multipath Fading** occurs when signals reflect off surfaces such as walls, glass, or metal, arriving at the receiver at slightly different times. This can cause constructive or destructive summation of the reflected signals, leading to fluctuations in signal strength.
- **Scattering** happens when signals encounter small objects or irregular surfaces—such as leaves or stems—which redirect the signal in multiple directions, often weakening it.
- **Absorption** refers to the loss of signal energy as it passes through materials that partially absorb RF energy, such as water-rich plant tissue.

Receiver Sensitivity

Another important concept is receiver sensitivity, which defines the minimum signal strength required for a user equipment (UE) to maintain a reliable connection. For typical 5G devices, this threshold is often around -120 dBm, though it varies depending on the device and network configuration.

Understanding these principles is essential for interpreting RF measurements and planning wireless infrastructure—especially in complex environments like greenhouses, where signal behaviour can be highly variable.

4.2.2 Wireless signals in greenhouses

Greenhouses present a complex and dynamic environment for wireless communication. Unlike typical indoor spaces, they contain large amounts of biomass, reflective surfaces, and structural materials that significantly affect radio signal behaviour.

Even when the greenhouse is empty, due to the reflective nature of greenhouse surfaces, the signal can be severely degraded by multipath fading. With multipath fading, signals may bounce off glass panels, metal structures, and even the moist surfaces of plants, arriving at the receiver with varying delays and phases. This causes fluctuations in signal strength and reduce reliability. This can even result in signal degradation at short distances compared to large open environments where the FSPL is more dominant.

Scattering becomes increasingly relevant as plants grow. Leaves, stems, and other plant components scatter RF signals, especially at higher frequencies. This leads to unpredictable signal paths and additional attenuation, which can vary depending on plant density and growth stage.

Absorption also plays a role, as water-rich plant tissue can absorb part of the signal energy. This effect becomes more significant as the greenhouse fills with mature plants, further reducing signal strength.

These factors combined make it challenging to maintain consistent wireless coverage throughout a greenhouse. The dynamic nature of plant growth means that signal conditions can change over time, requiring accurate planning and careful placement of network infrastructure.

Although theoretical models exist to estimate path loss, they often fall short in accurately predicting signal behavior in greenhouse environments. This is due to the large number of variables that influence wireless signal propagation—such as plant density, moisture levels, structural materials, and even minor changes in layout or growth stage. These factors can significantly alter the signal path and result in varying levels of attenuation.

Therefore, this report presents a series of measurements to provide an initial indication of the path loss that can be expected in a tomato greenhouse. It is important to note that these results are indicative rather than definitive. The number of measurements conducted is limited, and a more robust statistical path loss model would require a significantly larger dataset across varying conditions.

4.3 RF Path Loss Measurement Methodology

4.3.1 Greenhouse layout

The measurements were conducted in a tomato greenhouse located in Honselersdijk, the Netherlands. The greenhouse has a square layout measuring approximately 38 × 38 meters, bordered by glass walls and structured with rows of tomato plants spaced at regular intervals. The layout is illustrated in **Figure 4.1**.

To assess signal propagation under realistic conditions, four distinct receiver locations (RX1–RX4) were selected, each representing a different propagation path through the greenhouse:

RX1: A path of 36.6 m
RX2: A path of 39.6 m
RX3: The longest diagonal path across the greenhouse of 48 m
RX4: The path closest to the wall of the greenhouse of 31.2 meters. This path crosses the plant rows at a perpendicular angles through 17 rows of tomato plants.

Each path was chosen to represent a different level of obstruction and signal complexity, allowing for a comparative analysis of RF behaviour under varying conditions. Additionally the paths have different distance from the transmitter which could be used to estimate a pathloss model based on the distance between the transmitter and receiver.

The transmitter (TX1) was placed at a fixed location and a fixed height of 3 meter near one corner of the greenhouse, with the receiver antennas RX1-RX4 positioned at three different heights: 1 meter, 2 meters, and 3 meters. The directional log-periodic TX1 and RX1-RX4 antennas were fixed in either vertical or horizontal polarization mode. These heights were selected to reflect typical wireless device placements, including 5G access point, ground-level sensors and elevated equipment such as cameras.

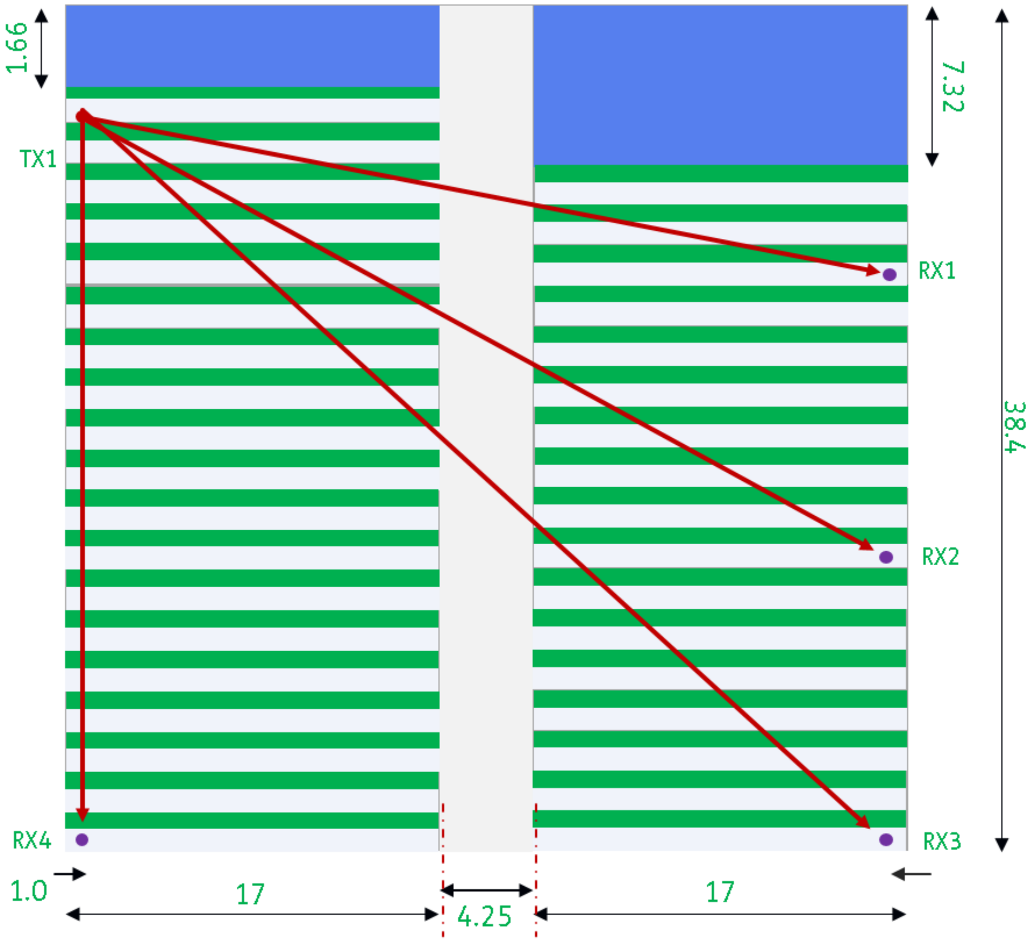

Figure 4.1: Greenhouse layout including measurement paths

This layout forms the basis for the measurement scenarios described in the following sections, where signal behaviour is analysed across different paths, heights, and polarizations.

4.3.2 Measurement setup

The setup was designed to capture signal attenuation across a wide frequency range and under varying spatial conditions.

Equipment configuration

The measurement system consisted of a Rohde & Schwarz ZVL-6 network analyzer, connected to two directional log-periodic antennas. These antennas provided a gain of approximately 5 dBi across the frequency range of 680 MHz to 6000 MHz, and were selected to minimize reflections from the rear side of the antenna. A low-noise amplifier (LNA) was used to improve signal sensitivity, and the setup was normalized to account for cable losses and amplifier gain, ensuring that only environmental effects influenced the measured signal strength. Figure 4.2 shows the measurement setup in more details.

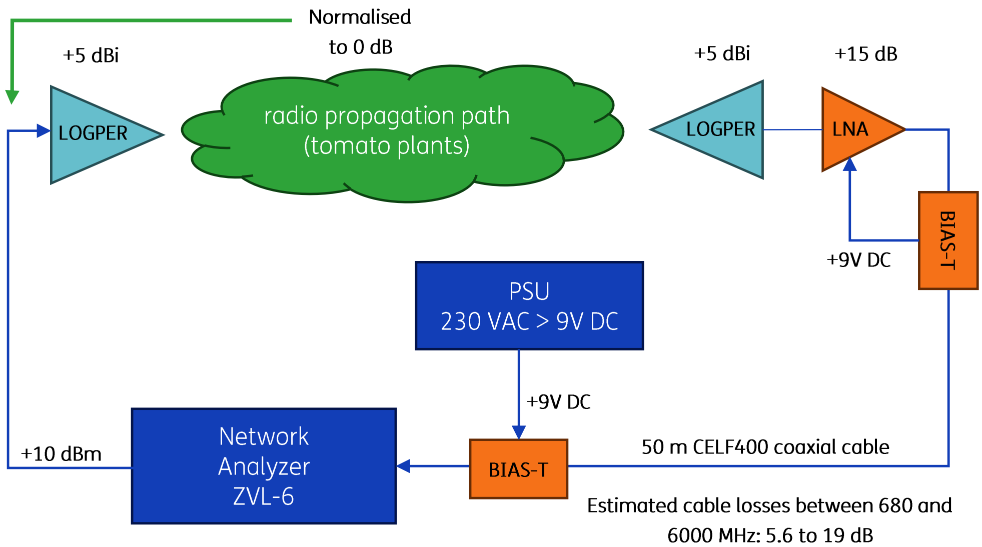

Figure 4.2: Measurement setup for measuring signal strength using a network analyzer

Measurement parameters

- Frequency sweep: 680 MHz to 6000 MHz
- Resolution bandwidth: 100 Hz
- Polarization: Both vertical and horizontal polarization where measured
- Sweep time of approximately 2 seconds
- TX at a fixed height of 3 meters
- Variable RX height at 1 meter, 2 meters and 3 meters.
- Variable TX – RX distance at 36.6, 39.6, 48 and 31.2 meters

4.3.3 Measurement conditions

The measurement were conducted under different conditions:

- **30 January 2024 – 0% grown**: Plants had just been planted. This session served as a baseline reference with minimal vegetation.
- **20 February 2024 – ~50% grown**: Plants had reached approximately half their full height, introducing moderate biomass into the signal paths.
- **18 March 2024 – Fully grown**: Plants were at full height and density. In this scenario, the vegetation exceeded the height of the transmit antenna (3 meters), resulting in a fully obstructed path.
- **10 April 2024 – 0% grown (again)**: A second measurement under minimal vegetation conditions was performed to validate the consistency and reliability of the setup.

These sessions allowed for a comparative analysis of signal attenuation across different growth stages. Notably, only during the fully grown phase did the tomato plants exceed the height of the transmit antenna. The measurement at 0% was repeated for the second time to validate the setup and the results obtained from the first measurement.

Note here that the greenhouse contained a variety of tomato plants, ranging from small cherry tomatoes to large beef tomatoes. This diversity in size of tomatoes introduces additional variability parameter, which might deviate from the conditions in production greenhouse situations where only one type of tomato plants are grown. This is an additional reason why the measurement results in this report should be considered only as an indication of additional vegetation loss that could be induced on the radio signal.

Figure 4.3 to **Figure 4.7** show the setup in the greenhouse during the various measurement days.

Figure 4.3: TX antenna setup and greenhouse overview 30 January 2024 at 0% plant growth

Figure 4.4: TX antenna setup 20 February 2024 at 50% plant growth

Figure 4.5: Greenhouse overview 20 February 2024 at 50% plant growth

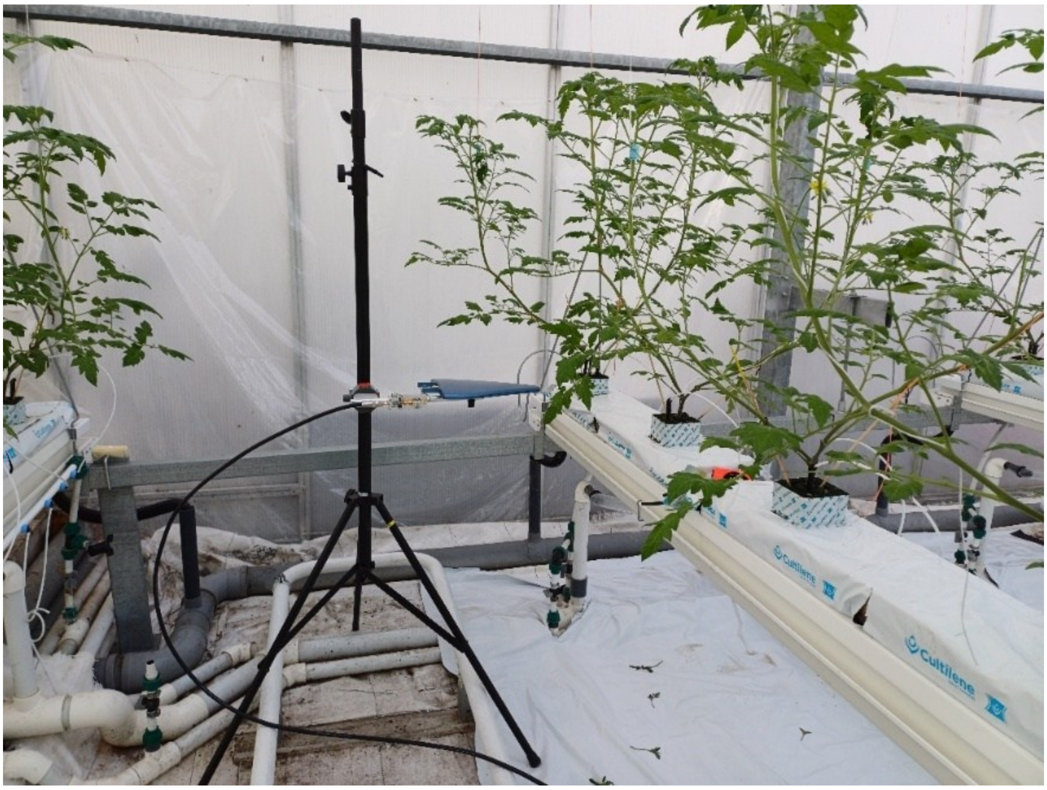

Figure 4.6: RX antenna setup 20 February 2024 at 50% plant growth

Figure 4.7: TX antenna setup and greenhouse overview 18 March 2024 at 100% plant growth

4.4 Results

4.4.1 Observations

Figure 4.8 shows measurement results of the three growth stages of the tomato plants with RX3 receiver location at 2 m height. The first measurements at 20 February 2024 show that when the plants are just potted, the propagation loss follows the free space loss with a pathloss variability for slight frequency changes that can be explained by multipath fading.

The increase of the path losses during the subsequent months after February, very distinctly followed the growth of the tomato plants. The attenuation effects occur across the entire swept frequency spectrum of 680 – 6000 MHz, although high frequencies suffer slightly more than low frequencies due to the additional vegetation loss.

For the furthest RX3 measurement location and fully grown tomato plants, when the frequency exceeded 3800 MHz, the received signal dropped below the receiver's sensitivity of the network analyzer. Therefore, for the frequency range above 3800 MHz in **Figure 4.8**, the path loss at RX3 location could not have been measured and only the fitted curve is visible.

There is a slight horizontal polarization preference over vertical polarization for frequencies below 3 GHz, in other words there is ≈ 3-4 dB less loss at horizontal polarization when compared to vertical polarization, see **Figure 4.9**.

The dotted line in the **Figure 4.8** represents a curve fitted to the measurement data and the curve fitting formulas are provided in the legend of **Figure 4.8**. To create this fit, the free space path loss curve is shifted downwards by adding a fixed dB value, until the average error with the measured values is minimized. The vertical distance by which the curve is shifted (in dB) quantifies the additional signal attenuation caused by vegetation. This value is referred to as the *vegetation attenuation*. This *vegetation attenuation* when the plants are fully grown, the four receiver locations RX1 to RX4 and horizontal/vertical polarisation is listed in **Table 4.1**.

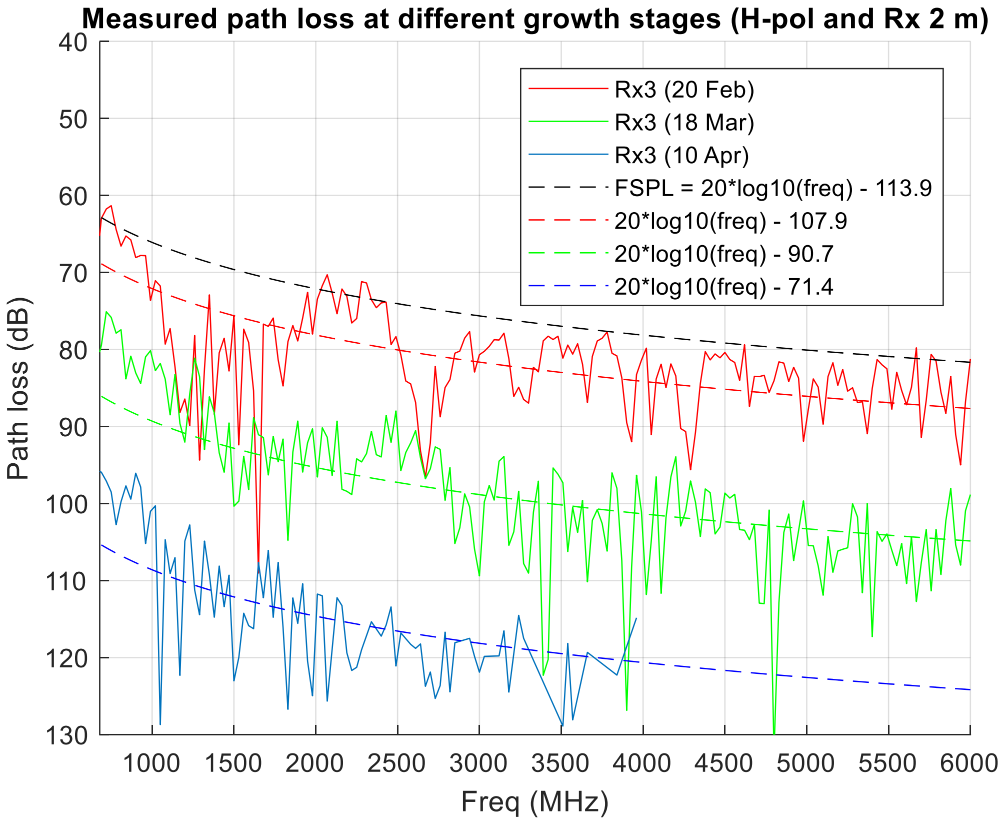

Figure 4.8: Measurement results for the three stages of plant growth. 20 February 2024 at 0%, 18 March 2024 at 50% and 10 April 2024 at 100%. Additional curve fitting (dotted lines) is added to estimate the losses for each growth state. The legend show the formula used for these curve fittings.

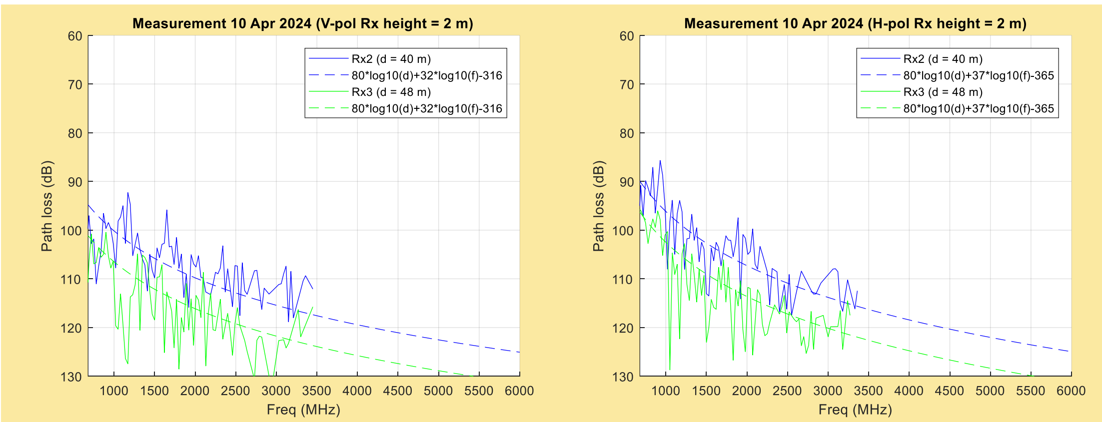

a) Vertical polarization   b) Horizontal polarization

Figure 4.9: Measurements and curve fitting showing vegetation loss at different polarizations

Table 4.1: Vegetation loss at 100% plant growth with RX height of 2 meters

	LOSS WITH VERTICAL POLARIZATION (DB)	LOSS WITH HORIZONTAL POLARIZATION (DB)
RX1 (37 m)	28	25
RX2 (40 m)	37	35
RX3 (48 m)	45	43
RX4 (31 m)	28	24

4.4.2 Additional vegetation attenuation estimation

In order to estimate the additional *vegetation loss* the measurements from the four different receiver locations RX1 to RX4 were combined and the additional loss is depicted in **Figure 4.10**. Here the additional loss is determined as the difference in the fitted FPSL values for zero grown and full grown tomato plants. Further, as the additional loss value was similar for the pairs of receiver locations e.g. (RX1, RX 4) and (RX2, RX3) the averaged additional loss value from the pair is plotted in **Figure 4.10**. It is important to note the following:

1. The additional vegetation loss is strongly influenced by the frequency and the relative position between the transmitter and receiver.
2. The higher the frequency the higher the additional vegetation attenuation. For example, if we compare the additional attenuation at 680 MHz till 6000 MHz we have ca 10-13 dB attenuation increase:
    a. Additional vegetation attenuation of 25 to 35 dB, respectively, for RX1 and RX4 positions
    b. Additional vegetation attenuation of 37 to 50 dB, respectively, for RX2 and RX3 positions.
3. Apparently, how many rows of tomato plants are crossed as well as how the signal propagates through the tomato plants rows (e.g. perpendicularly, in parallel, or diagonally) also has significant effect to the additional vegetation loss.
    a. We have additional attenuation loss that is ca. 12 dB to 15 dB lower for RX1 (signal propagating mostly parallel to tomato plants rows) and RX4 (signal propagating perpendicular to tomato plants rows)[4], when compared to RX2 and RX3 (signal propagates diagonal to tomato plants rows).

[4] Note that for RX4 we might also have effects from the glass wall between the TX and RX4.

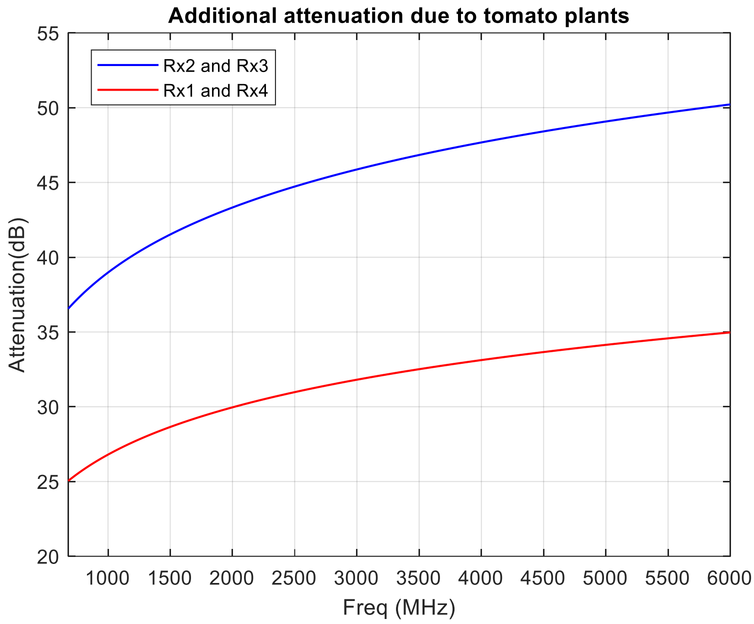

Figure 4.10: Overview of the additional vegetation attenuation for the measured frequency range and different RX positions

4.5 Impact of the tomato plant growth on wireless signals

This section translates the technical findings from our measurement campaign about the vegetation loss into practical insights. The vegetation loss in decibels (dB), which can be abstract for non-specialists, will be used to give insights on the impact on the coverage range for example Wi-Fi and 5G system configurations.

We focus on three frequencies commonly used in wireless technologies:

- 2.4 GHz (used in Wi-Fi as used in many consumer and industrial devices),
- 5 GHz (used in Wi-Fi with higher speed, shorter range),
- 3.5 GHz (used in private 5G for robust, high-capacity networks).

It is important to note that the conducted measurement campaign can be used as an insight, not an accurate estimation of the Wi-Fi or 5G coverage. The number of measurements was not sufficient to develop an accurate model for signal loss over distance within a tomato greenhouse. Secondly, the variety of tomato plants in each row makes the measurements less representative, also because the measured paths were traversing the tomato plan rows at different angles, causing each propagation path to pass through different types of tomato plants.

4.5.1 Effects on coverage range at different vegetation attenuation

From the trends illustrated in Section 4.4.1 and Section 4.4.2 about the additional vegetation attenuation in the frequency range of typical Wi-Fi systems (e.g. at 2.4 GHz or 5 GHz) or private 5G systems (e.g. at 3.5 GHz), we could examine the effect on the coverage range from additional vegetation attenuation of 45 dB. Here, 45 dB was calculated as additional vegetation loss at 48 meters on top of the free space loss fitted formula to the measurements at RX3 location and full grown tomato plants.

For maintaining coverage (i.e. feasible communication link), the received signal power P_{rx} must be above the **receiver sensitivity threshold** $P_{rx,sensitivity\_threshold}$.

Consequently, for a given distance *d*, we can determine if there is a coverage if:

$$P_{rx} = P_{tx} - FSPL(d) - L_{att} \geq P_{rx,sensitivity\_threshold}$$

Here, $FSPL(d)$ is the free space path loss at distance *(d)*, P_{Tx} the transmit power, and L_{att} the vegetation attenuation.

For an example Wi-Fi system with 40 MHz bandwidth a typical receive sensitivity reported in literature is $P_{rx,sensitivity\_threshold}$ = -90 dBm while the maximal allowed transmit power of a Wi-Fi terminal is P_{tx} = 20 dBm. The expected signal strength at 48 meters with a frequency of 2.4GHz based on a L_{att} = 45 dB vegetation attenuation is then:

$$P_{rx} = P_{tx} - FSPL(48) - L_{att}$$

$$P_{rx} = 20 - 73.7 - 45 = -98.7\ dBm$$

Based on the calculation for Wi-Fi at 2.4GHz it can be seen that the signal strength is too low at 48 meters to get a Wi-Fi connection i.e. -98.7 dBm is lower than the receiver sensitivity threshold of -90 dBm. These type of calculations can also be done for Wi-Fi at 5 GHz and 5G at 3.5 GHz. **Table 4.2** uses the same calculation method as indicated above to estimate at what distances a connection can still be established when the vegetation attenuation as in **Table 4.1** is considered.

Table 4.2: coverage of Wi-Fi and 5G at measured distances. Yes/No indicates if an connection can be established for that technologies at that distance

	Wi-Fi 2.4 GHz	Wi-Fi 5 GHz	5G 3.5 GHz at 40 MHz bandwidth	5G 3.5 GHz at 100 MHz bandwidth
Receive sensitivity (dB)	-90	-90	-93	-89
31 meter	yes	yes	yes	yes
37 meter	yes	yes	yes	yes
40 meter	yes	no	yes	no
48 meter	no	no	no	no

4.5.2 Effects on resource planning for 5G

Effective resource planning is essential when deploying private 5G networks in greenhouse environments. It involves determining the optimal number and placement of base stations (gNBs) to ensure reliable coverage and performance across the entire greenhouse area. While free space path loss models can provide a baseline, the presence of vegetation—as shown in the result in Section 4.5.1 —introduces significant additional attenuation that must be accounted for.

To illustrate this, we compare two scenarios in a greenhouse that is nine times larger than the one used in our measurement campaign. The scaled-up greenhouse measures 114.75 × 115.2 meters, representing a realistic commercial deployment size. See **Figure 4.11** and **Figure 4.12**.

Scenario 1: No Tomato Plants (0% Grown):
In an empty greenhouse, signal propagation is relatively unobstructed. Based on our measurements we experienced an additional initial vegetation attenuation of roughly 5 dB on top of the free space loss. With this small attenuation, a single 5G base station operating at 3.5 GHz is sufficient to cover the entire greenhouse. This single base station is illustrated as a blue dot in the middle of the green house in **Figure 4.11**. The signal can reach all corners with acceptable strength, assuming the method used for the calculations in **Table 4.2**.

Scenario 2: Fully Grown Tomato Plants (Vegetation Attenuation ≈ 45 dB)
When the greenhouse is filled with fully grown tomato plants, signal attenuation increases dramatically. Our measurements indicate an additional vegetation attenuation of up to 45 dB, which significantly reduces the effective 5G coverage range.

We can see from the measurements that at 48 meters there is no 5G coverage anymore therefore to cover the full scaled-up greenhouse multiple base stations are needed to get full coverage. To maintain reliable connectivity throughout the greenhouse under these full grown tomato plant conditions, four 5G base stations are required. This ensures that each section of the greenhouse receives sufficient signal strength despite the heavy attenuation caused by the plants. The four base stations are shown as purple dots in **Figure 4.12**.

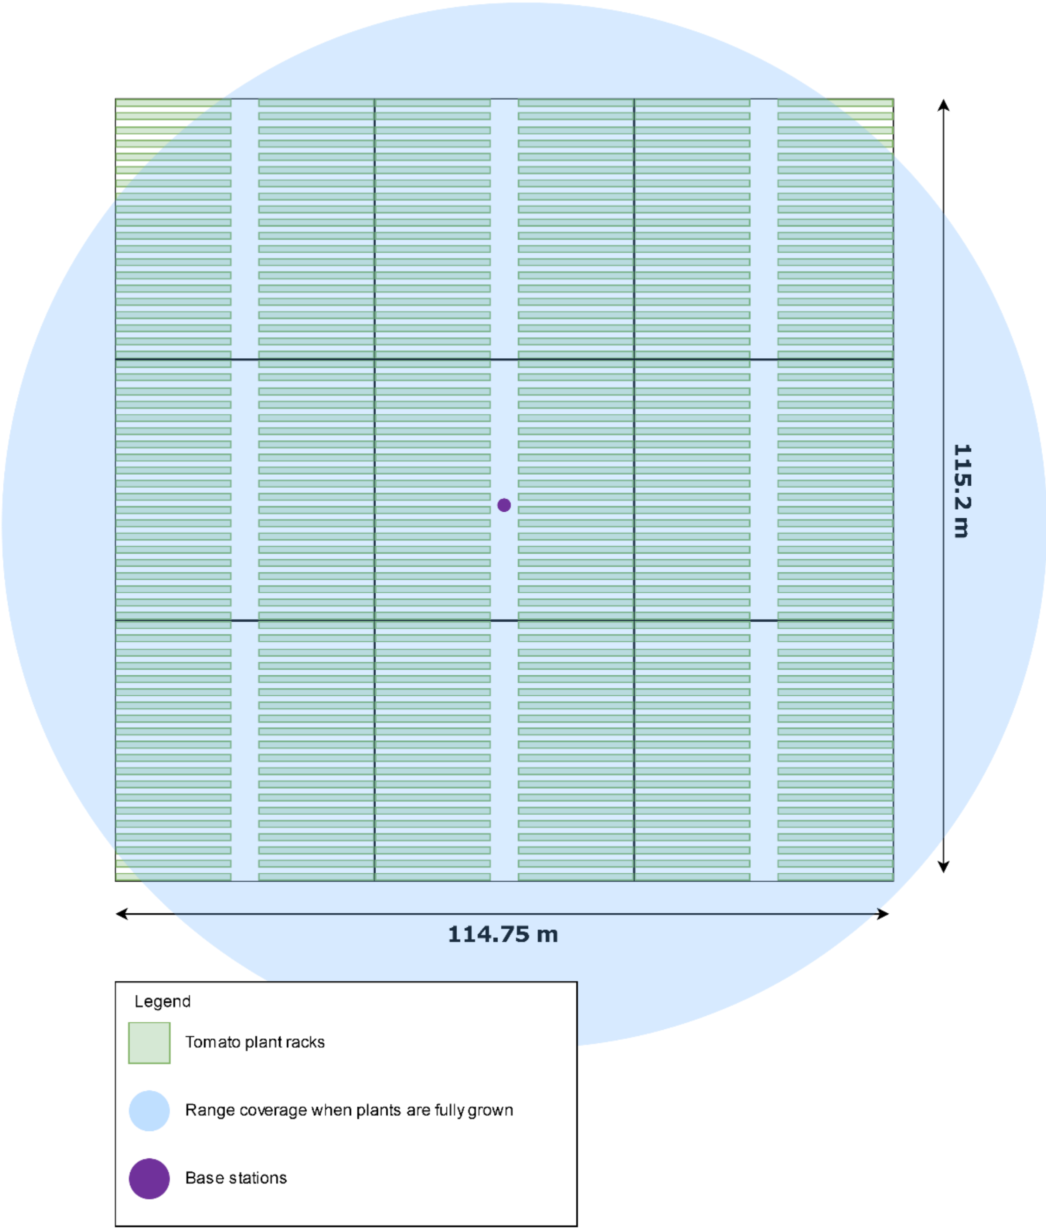

Figure 4.11: One central antenna can cover the whole greenhouse without any plants

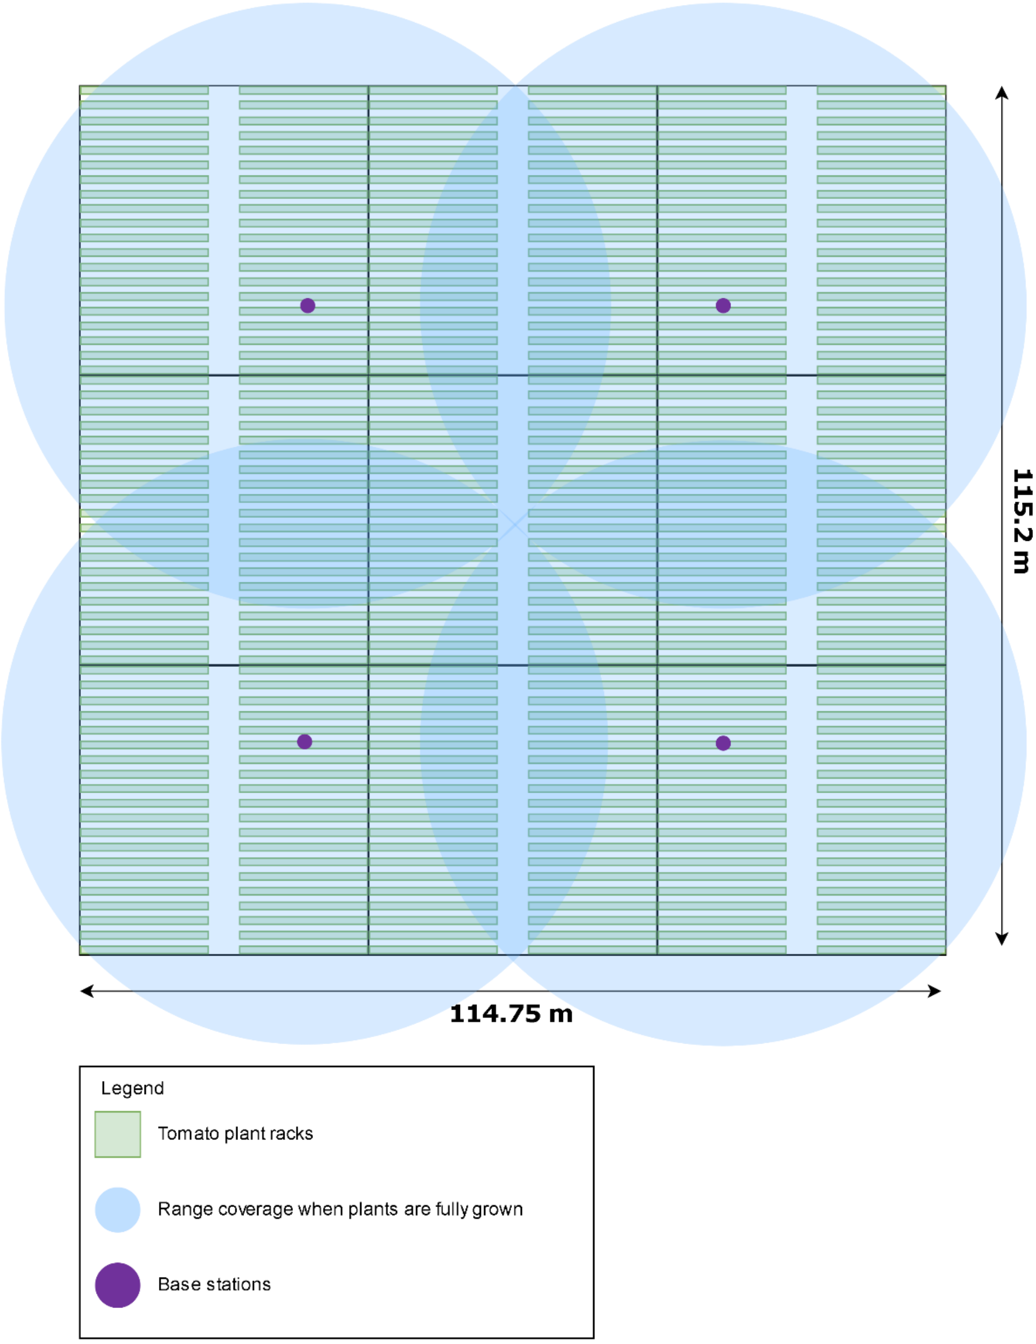

Figure 4.12: Due to the plants, the coverage per antenna is decreased, a total of 4 antennae are needed

This comparison highlights the critical importance of accounting for plant growth in wireless network design. Without proper planning, coverage gaps and performance issues may arise, especially in high-frequency bands like 3.5 GHz.

4.6 Conclusion

In this work, we tried to answer the research question:
How does plant growth in a greenhouse affect wireless signal propagation, and what are the implications for planning a private 5G network?

Our measurements in the Tomatoworld greenhouse clearly demonstrate that vegetation significantly impacts wireless signal propagation. When tomato plants are fully grown, additional vegetation attenuation of **up to 45 dB at 48 meters** was observed across the measured frequency range (680 MHz–6000 MHz), compared to an empty greenhouse. A translation of this vegetation attenuation via the free-space path loss formula and typical receiver sensitivity threshold values indicates different maximum coverage, depending on the wireless system and the used frequency bands:

- At **2.4 GHz** (as used in Wi-Fi), this simple translation indicates coverage up to 40 meters under full vegetation.
- At **5 GHz** (as used in Wi-Fi), this simple translation indicates coverage up to 37 meters under full vegetation.
- At **3.5 GHz** (as used in private 5G), this simple translation indicates coverage up to 37 meters for high data rates (100 MHz bandwidth) and 40 meters for medium data rate (40 MHz bandwidth) under full vegetation.

The measurement results also highlight that vegetation attenuation is not only frequency-dependent but also influenced by propagation paths between the two communication end-points. For example, diagonal paths through plant rows showed **12–15 dB higher losses** than parallel or perpendicular paths.

The practical implication of the additional vegetation attenuation for greenhouse wireless network design is substantial. Given the example calculation in Section 4.5.2 in an empty greenhouse, **one 5G base station** (or Wi-Fi access point) can provide full coverage for a commercial-scale greenhouse (≈115 × 115 m). However, with fully grown crops, **four 5G base stations** (or Wi-Fi access points) are required to maintain reliable connectivity[5]. This represents a **four-fold increase in infrastructure** and underscores the importance of accounting for additional vegetation attenuation for plant growth in private 5G base station or Wi-Fi deployment planning.

[5] For Wi-Fi deployments on higher frequency bands (e.g. 5GHz, 6 GHz) or private 5G deployments at 3.5 GHz with wider bandwidth of 100 MHz more than 4 antennas could be needed.

5 Energy Consumption Measurements

5.1 Introduction

This work investigated the energy efficiency of 5G Reduced Capability (RedCap) modems compared to traditional 5G modems. The motivation stems from the growing need for battery powered IoT-devices in greenhouses. Where extensive measurements of technologies like LTE-M, NB-IoT and LoRa exist, New Radio (NR) technologies like RedCap are not yet fully developed and still fall behind. RedCap, introduced in 3GPP Release 17, aims to reduce modem complexity and power consumption while maintaining compatibility with standalone 5G networks.

To evaluate RedCap's performance, a series of experiments were conducted simulating a scenario where 128-byte messages are transmitted every 300 milliseconds. While this frequency exceeds typical IoT transmission rates, it accelerates the measurement process and highlights energy consumption patterns. The report outlines the RedCap technology, describes the measurement setup and methodology, presents the experimental setups for both regular 5G and RedCap modem, and concludes with key insights.

5.2 RedCap overview

5G Reduced Capability (RedCap) was introduced in 3GPP Release 17[6]. As the name suggests, RedCap reduces the complexity of the modem to lower the cost and energy consumption. In comparison with LTE-M and NB-IoT technologies, RedCap does not require a 5G system with backward 4G (LTE) support. This makes it possible to support RedCap on any 5G standalone network without the need for additional (LTE) hardware.

Although 5G RedCap was introduced in 3GPP Release 17 (completed in 2022), the first development modules were introduced in mid-2024. This means that RedCap is still in the early phase of development. Nevertheless, it shows great potential for IoT devices. Release 17 RedCap can be placed between traditional 5G NR devices and LTE-M devices in terms of peak data rate. This means that RedCap is capable of higher data rates than LTE-M or NB-IoT, which presumably will come at the cost of higher power consumption. Two of the suggested use cases for RedCap are wearables (like smartwatches) and battery-powered (security) cameras. They send more data than, for example, simple sensors, while still being battery-powered. **Figure 5.1** illustrates how 5G NR, 5G RedCap Release 17, and 5G RedCap Release 18 relate to their LTE counterparts in terms of peak data rate.

[6] ETSI, "3GPP TR 21.917 version 17.0.1 Release 17," 2023.

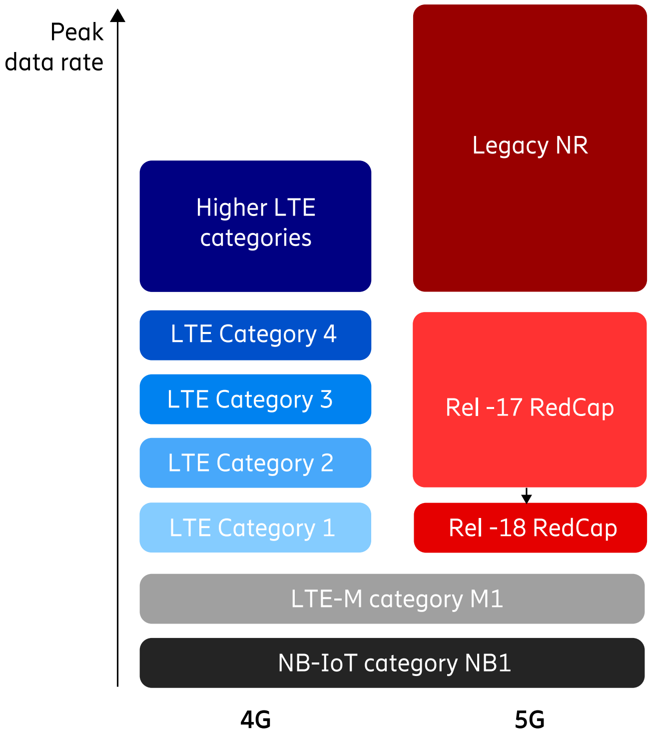

Figure 5.1: Position of release 17 RedCap in terms of data rate compared to other cellular technologies[7]

The peak data rate of release 17 RedCap depends on the configuration of the channel. The simplest RedCap device can have a peak data rate of around 50 Mbps (downlink) where the most complex version can achieve data rate up to 645 Mbps (downlink). See **Figure 5.2**.

3GPP release 18[8] specifies additional RedCap enhancements under the name eRedCap (enhanced Reduced Capability). The complexity and data rate of eRedCap devices is even further reduced to make it more power efficient. eRedCap utilizes a maximum bandwidth of 5 MHz with a corresponding maximal data rate of 10 Mbps for downlink traffic. This makes eRedCap similar to the needs fulfilled by LTE cat 1 (see **Figure 5.2**).

The current RedCap modules on the market only support Release 17 RedCap, which means they are expected to be more closely related to legacy 5G modems in terms of energy consumption than to NB-IoT and LTE-M. Secondly, the energy consumption may still not be comparable with NB-IoT and LTE-M because the RedCap modem is still in active development and thus not yet optimized for power-efficient applications. For that reason, this report compares the energy consumption of the RedCap modem with a legacy 5G modem to investigate if it is more power-efficient, even though the modem will not be fully optimized in terms of energy consumption.

[7] Ericsson, "RedCap - expanding the 5G device ecosystem for consumers and industries," 2023. [Online]. Available: https://www.ericsson.com/493d70/assets/local/reports-papers/white-papers/redcap-5g-iot-for-wearables-and-industries.pdf.
[8] ETSI, "3GPP TS 38.300 version 18.1.0 Release 18," 2024.

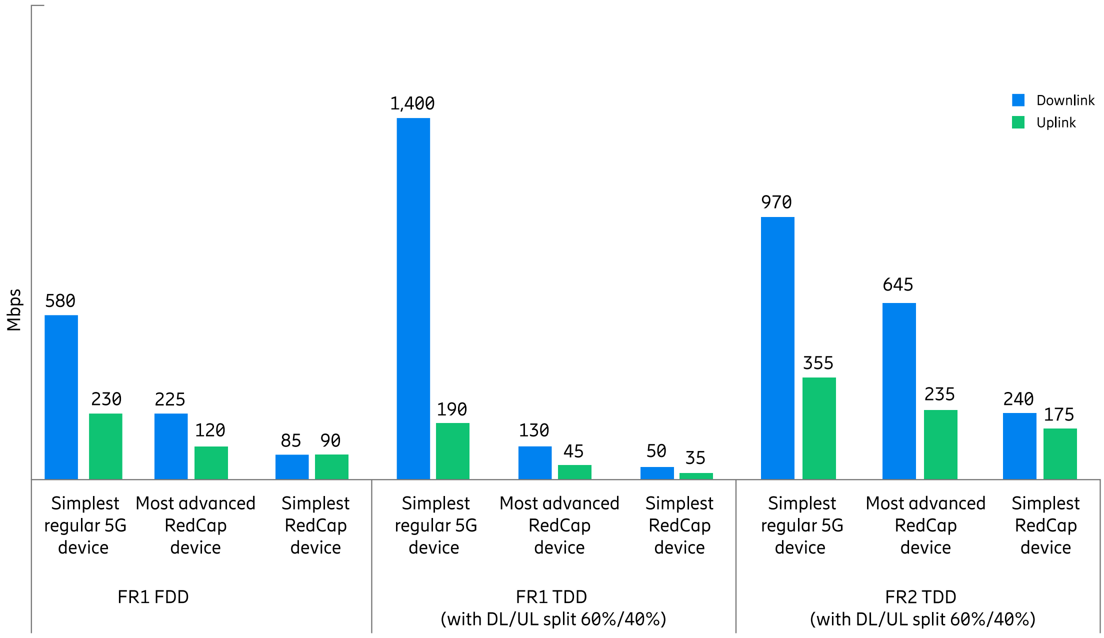

Figure 5.2: Release 17 RedCap data rates compared to regular 5G devices. FR1 (Frequency Range 1) covers sub-6 GHz bands, while FR2 (Frequency Range 2) refers to mmWave frequencies above 24 GHz. TDD (Time Division Duplex) uses the same frequency for uplink and downlink, separated in time, whereas FDD (Frequency Division Duplex) uses separate frequencies for simultaneous uplink and downlink[9] [10].

5.3 Measurement Setup and Methodology

5.3.1 Hardware Configuration

The measurement setup comprises four components

- **5G Standalone Network:** Amarisoft Callbox Classic[11].
- **Modems:** a RedCap modem (Quectel RG255C[12]) and a 5G modem (Fibocom FG150[13])
- **Power measurement:** the Otii Arc Pro[14]
- **Setup controller:** Ubuntu LTE 20.04.6

Due to restrictions on high-power transmission in the 3.5GHz band, the modems are positioned in close proximity, approximately one meter from the Amarisoft 5G box. In both setups, the modems are powered by Otii, which can measure the energy consumption in detail. Although the modems are located on top of a development board, only the modems themselves are powered and measured by the Otii system. The entire setup is controlled by a laptop.

[9] Ericsson, "RedCap - expanding the 5G device ecosystem for consumers and industries," 2023. [Online]. Available: https://www.ericsson.com/493d70/assets/local/reports-papers/white-papers/redcap-5g-iot-for-wearables-and-industries.pdf.
[10] Wikipedia, "5G NR frequency bands," [Online]. Available: https://en.wikipedia.org/wiki/5G_NR_frequency_bands.
[11] Amarisoft, [Online]. Available: https://www.amarisoft.com/test-and-measurement/device-testing/device-products/amari-callbox-classic.
[12] Quectel RG255C, [Online]. Available: https://www.quectel.com/product/5g-redcap-rg255c-series/.
[13] Fibocom FG150, [Online]. Available: https://www.fibocom.com/en/Products/5G-FG150-AE.html
[14] Otii Arc Pro, [Online]. Available: https://www.qoitech.com/otii-arc-pro/

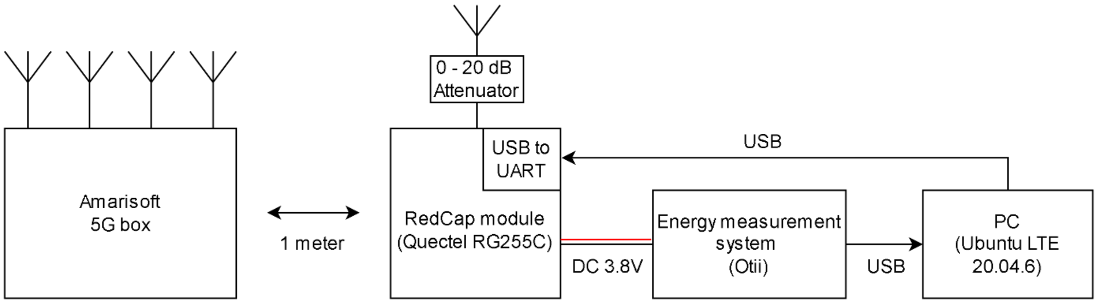

Figure 5.3: Measurement setup for 5G RedCap

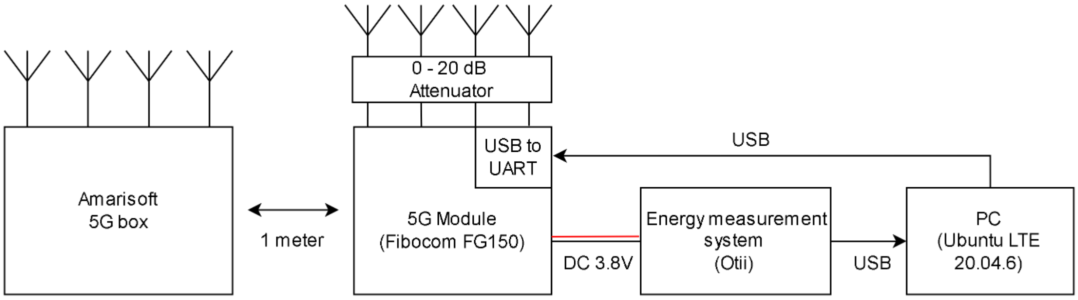

Figure 5.4: Measurement setup for a legacy 5G modem

To simulate different signal strengths, an attenuator is placed between the modem and the antenna(s). Varying the attenuation value allows us to observe the effect of channel quality on energy consumption. It is expected that the energy consumption will increase when the channel quality decreases, as more power is needed to transmit the same amount of data.

5.3.2 Measurement method

For the energy consumption measurements, a fully automated software solution was developed and run on the PC. This software enables precise control over the measurement parameters, allowing the user to configure the packet size, the total number of packets to be sent, and the interval between packets. All measurements were performed using UDP packets only, to simulate the data traffic typical of low data rate IoT devices.

During each experiment, the software transmitted a predefined number of UDP packets of 128 bytes each at a set interval (for these measurements, every 300 milliseconds). The energy consumption for each transmitted packet was automatically recorded. In addition to the transmission energy, the system also measured the energy consumed during idle periods. By subtracting the idle energy (i.e. when the modem is not transmitting or receiving) from the total, the additional energy required for each transmission could be calculated. This approach allows for a direct comparison of the relative energy consumption between different modem types, independent of the energy required to power the microcontroller.

For both the 5G and RedCap modems, three experiments were conducted, each with a different antenna attenuation (0 dB, 10 dB, and 20 dB) to simulate varying channel conditions. All measurements for each modem were performed on the same days to ensure consistency.

This automated methodology was used to perform 2000 consecutive measurements for each experiment, enabling an accurate and reproducible estimation of energy consumption by eliminating manual intervention and reducing the potential for measurement errors.

5.4 Results

5.4.1 Measurement samples

Figure 5.5 and **Figure 5.6** show two snapshots of the energy consumption measurements for a RedCap and 5G modem respectively. The idea of these snapshots is to get a general idea on how the energy consumption looks like. The top part of the snapshots show the measured current and the bottom part shows the total measured power.

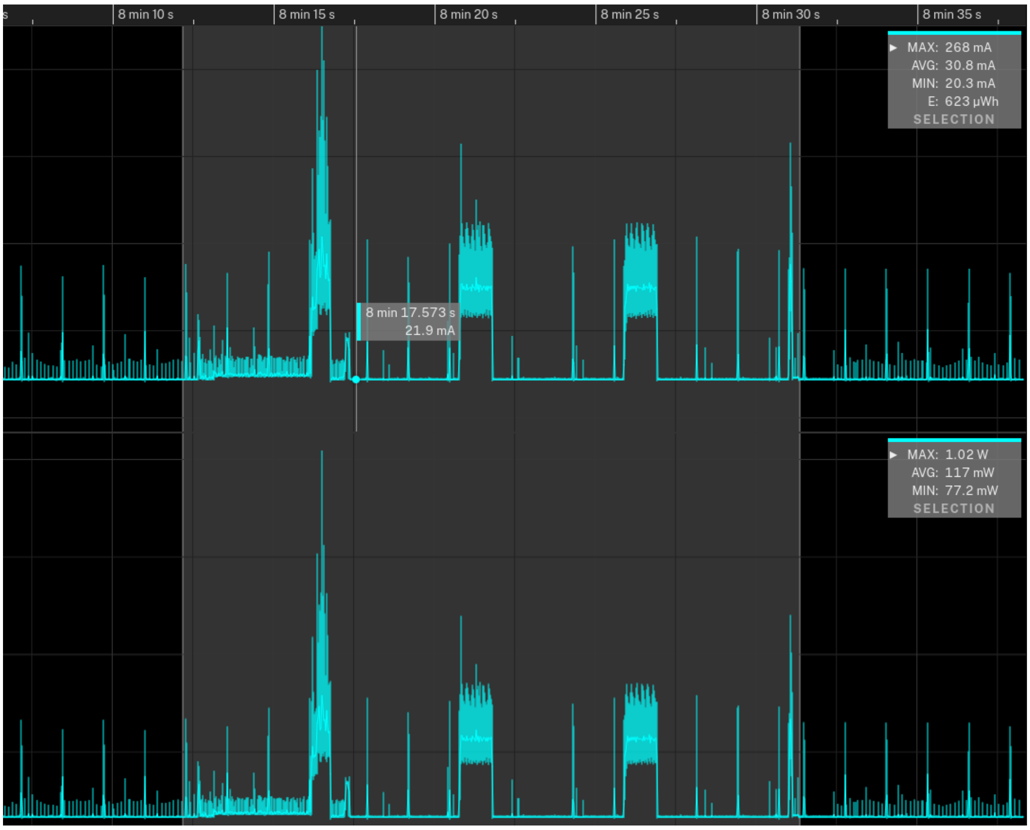

Figure 5.5: Sample of the energy consumption of the RedCap modem

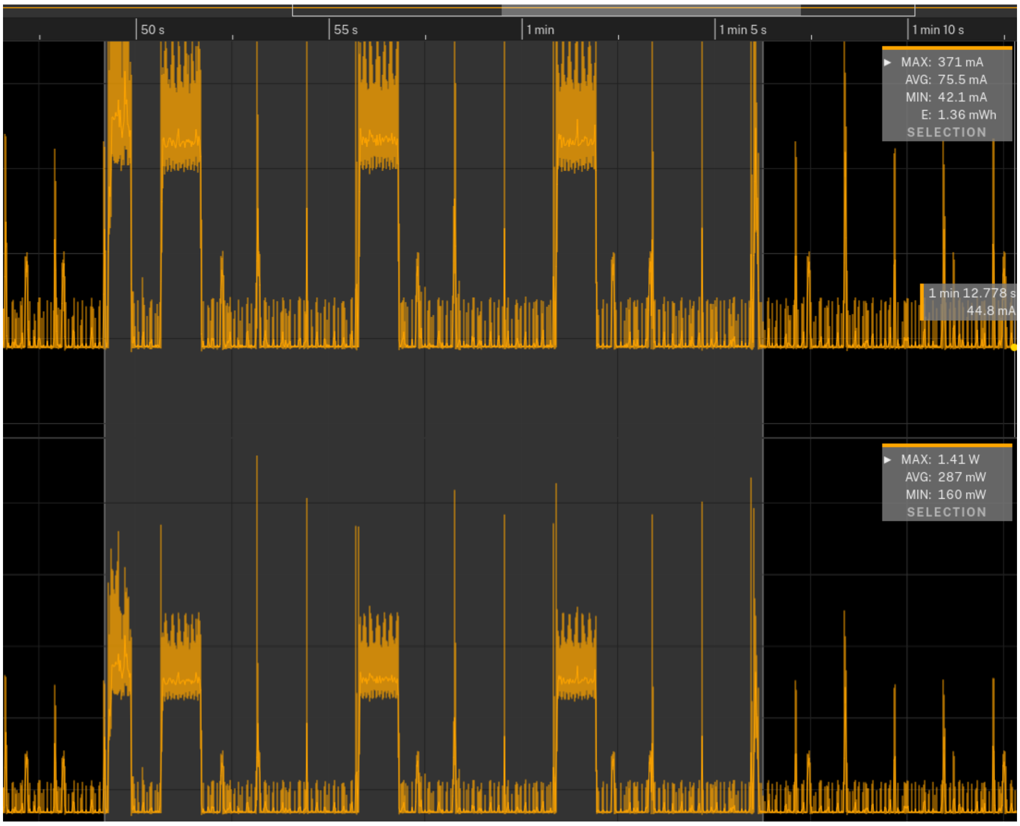

Figure 5.6: Sample of the energy consumption of the 5G modem

In **Figure 5.5** and **Figure 5.6**, multiple spikes in energy consumption are measured. The repeating smallest spikes are the paging messages of the modem. These messages basically ask the 5G station if there are new messages ready for reception. Within the selected parts of the measurements are some wider spikes visible. The first one is the part where the modem is transmitting a message. The other wider spikes occur when the modem is opening a receive window. Notice that the RedCap modem opens only two reception window and the 5G modem three. At least that is the case in this specific snapshot. It sometimes happens that the RedCap modem also opens three reception windows, it is not clear yet why this is not a constant.

5.4.2 Energy consumption measurements

To assess the energy efficiency of the RedCap modem compared to a legacy 5G modem, detailed measurements were conducted under varying signal conditions. The results, summarized in **Table 5.1** and **Table 5.2**, reveal consistent differences in power usage between the two technologies.

The analysis focuses on three attenuation levels—0 dB, 10 dB, and 20 dB—representing different channel qualities. For each condition, the modems transmitted 2,000 messages at fixed intervals, and both average and peak power consumption were recorded. The normalized energy values, which exclude idle power, offer a clear comparison of transmission efficiency.

Table 5.1: Measurement result of the RedCap modem transmitting 2000 messages with an interval of 300 ms

	Attenuation: 0dB	Attenuation: 10dB	Attenuation: 20dB
RSRP (dB)	-89	-100	-120
Average Power (mW)	182.3	187.06	187.29
Maximal Power (mW)	589.32	545.49	644.92
Energy (mWh)	30.91	31.74	32.07
Normalized energy (mWh)	27.10	27.92	28.22
Power per message (uWh)	13.55	13.96	14.11

Table 5.2: Measurement result of the legacy 5G modem transmitting 2000 messages with an interval of 300 ms

	Attenuation: 0dB	Attenuation: 10dB	Attenuation: 20dB
RSRP (dB)	-52	-61	-72
Average Power (mW)	431.28	432,37	460.55
Maximal Power (mW)	749.05	714.36	710.34
Energy (mWh)	83.99	83.98	90.01
Normalized energy (mWh)	64.08	64.12	70.03
Power per message (uWh)	32.04	32.06	35.02

The RedCap modem consistently consumes less energy than the legacy 5G modem across all attenuation levels. At 0 dB attenuation, the RedCap modem used 27.10 mWh of normalized energy, compared to 64.08 mWh for the 5G modem—less than half. This trend continues at higher attenuation levels: at 10 dB, the RedCap modem consumed 27.92 mWh, while the 5G modem used 64.12 mWh; and at 20 dB, RedCap reached 28.22 mWh, whereas the 5G modem rose to 70.03 mWh.

Additionally, the data also shows a modest increase in energy consumption as signal quality decreases (i.e., higher attenuation), which aligns with expectations. Lower RSRP values require the modem to use more power to maintain connectivity and transmit data reliably.

Notice that the RSRP (Reference Signal Received Power) of the 5G modem is in general much higher than that from the RedCap modem. Although the used bandwidth within both experiments is set to 20MHz, the 5G modem uses 2x2 MIMO for the downlink communication. This results in a higher effective bandwidth which directly relates to the total amount of power received by the modem. However, the difference is much higher than expected. In the best case, the bandwidth is doubled which means the maximum received power can maximal two times higher when using MIMO. In terms of decibels (dB) this would add an maximal of 3 dB to the received signal power. This is much lower than the 40dB difference measured during the experiments. Where these extremely high differences in RSRP are coming from is still unclear, however we can still draw conclusions based on the measured energy consumptions.

5.5 Conclusion

In this chapter, we presented multiple energy consumption measurements to evaluate the efficiency of 5G RedCap modems compared to traditional 5G modems. Since currently available RedCap modules are still in the development phase and not yet optimized for low power consumption, a direct comparison with LTE-M or NB-IoT is not appropriate at this stage. Instead, the focus has been on comparing RedCap with legacy 5G modems.

3GPP Release 18 introduced a more enhanced version of RedCap called eRedCap. This technology reduces the complexity of RedCap even further, which should lead to lower power consumption, comparable to LTE Cat1 technology. This would make the comparison between eRedCap and LTE-M/NB-IoT much fairer. However, unfortunately, there are no eRedCap modems available on the market now.

For both the RedCap and legacy 5G modems, three experiments were conducted, each with increasing antenna attenuation (0 dB, 10 dB, and 20 dB) to simulate different channel conditions. During the measurements each modem transmitted 2,000 messages at 300 ms intervals, and the results clearly demonstrate the energy-saving potential of RedCap technology.

The key findings of the measurements are:

- The RedCap modem consistently consumed less energy than the legacy 5G modem across all tested conditions.
- At 0 dB attenuation, the RedCap modem used less than half the normalized energy compared to the 5G modem.
- This trend continued when increasing the attenuation.
- Energy per message was significantly lower for RedCap, ranging from 13.55 µWh to 14.11 µWh, compared to 32.04 µWh to 35.02 µWh for the 5G modem.
- A modest increase in energy consumption was observed as signal quality decreased, which aligns with expectations.

Overall, while RedCap is not yet fully optimized, it already demonstrates a clear advantage in energy efficiency over traditional 5G modems. As the technology matures and eRedCap modules become commercially available, further improvements in power consumption are expected, potentially making RedCap a strong candidate for energy-constrained IoT applications.

6 Conclusions and Key Takeaways

Architecture

To link the research to real impact for the greenhouse sector, TNO envisioned a greenhouse that could be exported but still be managed from the grower in the Netherlands. For the grower to stay in control it is imperative that the data stays within control of the grower operating the greenhouse.

A Non-Public Network (NPN) architecture, featuring a Shared Radio Access Network (RAN) and Control Plane, provides the best option, with two separate core networks deployed: a public network core with UPF1 and a private network core with UPF2. This option creates the best of both worlds as the greenhouse provider owns the parts that hold the data of the greenhouse and at the same time the control of the network is handled by an MNO or service provider to operate the network. In this setup with two UPFs, it is possible to combine different types of UPF selection methods, such as S-NSSAI slicing and DNN based selection. The combined approach allows for a more granular and flexible allocation of resources.

The experimentation phase confirmed the practical viability of using multiple UPFs and network slicing to provide a shared public 5G network with private network functionalities in greenhouse environments. It demonstrated flexible resource allocation and optimized service delivery for different user demands.

This architecture offers greenhouse owners and growers a robust, flexible, and future-proof connectivity solution. It enables them to securely separate and manage different types of data traffic, ensuring that sensitive cultivation and business data remains protected within the greenhouse while still allowing for remote management and support. It also improves operational efficiency and reliability, as critical applications like robotics, sensors, and autonomous vehicles can be prioritized and guaranteed high-quality service, while guest or external traffic is kept separate, all on the same shared infrastructure. By consolidating connectivity needs into a single 5G infrastructure, growers can also reduce costs and simplify network management, replacing multiple vendor-specific Wi-Fi systems.

This approach not only enhances security and operational efficiency but also simplifies network management and supports the integration of new digital applications as the sector evolves. Ultimately, these capabilities empower greenhouse operators to accelerate digital transformation, improve productivity, and maintain a competitive edge in a rapidly changing horticultural landscape.

Cloud Federation

The distributed 5G network so created, demonstrated the feasibility of using a combination of Open Source technologies for networking and cloud computing, along with widely used, commercially available infrastructure to deploy a private 5G network. Cloud federation in particular brings a demonstrable benefit of increased resilience of core network functions. The caveat for such an implementation is that network connectivity between the two clusters is guaranteed, and that the network fabric used to provide key 5G functionality (such as radio access) must be mutually reachable between the two sites.

While the implementation created relies exclusively on Kubernetes for load balancing and routing, production 5G environments utilize the Border Gateway Protocol for communication between clusters. Kubernetes may be operated in conjunction with existing Border Gateway Protocol routers using MetalLB, a load balancer designed to operate with standard network equipment. However, this was not explored in the context of the project. Adaptation of the implementation undertaken here using MetalLB would thus constitute the most fruitful avenue for future work.

Although this was a very technical proof of concept (PoC), for the greenhouse use case the outcome is relevant. This PoC shows that sensitive operational data can remain within the local greenhouse cloud, while investment in the local greenhouse cloud can be minimized by offloading less sensitive processing or backup to external clouds, while ensuring compliance and privacy.

Radio Network Coverage

Our work tried to answer the research question:
How does plant growth in a greenhouse affect wireless signal propagation, and what are the implications for planning a private 5G network?

Our measurements in the Tomatoworld greenhouse clearly demonstrate that vegetation significantly impacts wireless signal propagation. When tomato plants are fully grown, additional vegetation attenuation of **up to 45 dB at 48 meters** was observed across the measured frequency range (680 MHz–6000 MHz), compared to an empty greenhouse. A translation of this vegetation attenuation via the free-space path loss formula and typical receiver sensitivity threshold values indicates different maximum coverage, depending on the wireless system and the used frequency bands:

- At **2.4 GHz** (as used in Wi-Fi), this simple translation indicates coverage up to 40 meters under full vegetation.
- At **5 GHz** (as used in Wi-Fi), this simple translation indicates coverage up to 37 meters under full vegetation.
- At **3.5 GHz** (as used in private 5G), this simple translation indicates coverage up to 37 meters for high data rates (100 MHz bandwidth) and 40 meters for medium data rate (40 MHz bandwidth) under full vegetation.

The measurement results also highlight that vegetation attenuation is not only frequency-dependent but also influenced by propagation paths between the two communication endpoints. For example, diagonal paths through plant rows showed **12–15 dB higher losses** than parallel or perpendicular paths.

The practical implication of the additional vegetation attenuation for greenhouse wireless network design is substantial. Given the example calculation in Section 4.5.2 in an empty greenhouse, **one 5G base station** (or Wi-Fi access point) can provide full coverage for a commercial-scale greenhouse (≈115 × 115 m). However, with fully grown crops, **four 5G base**

stations (or Wi-Fi access points) are required to maintain reliable connectivity[15]. This represents a **four-fold increase in infrastructure** and underscores the importance of accounting for additional vegetation attenuation for plant growth in private 5G base station or Wi-Fi deployment planning.

Energy Consumption Measurements

In this work, we presented multiple energy consumption measurements to evaluate the efficiency of 5G RedCap modems compared to traditional 5G modems. Since currently available RedCap modules are still in the development phase and not yet optimized for low power consumption, a direct comparison with LTE-M or NB-IoT is not appropriate at this stage. Instead, the focus has been on comparing RedCap with legacy 5G modems.

3GPP Release 18 introduced a more enhanced version of RedCap called eRedCap. This technology reduces the complexity of RedCap even further, which should lead to lower power consumption, comparable to LTE Cat1 technology. This would make the comparison between eRedCap and LTE-M/NB-IoT much fairer. However, unfortunately, there are no eRedCap modems available on the market now.

For both the RedCap and legacy 5G modems, three experiments were conducted, each with increasing antenna attenuation (0 dB, 10 dB, and 20 dB) to simulate different channel conditions. During the measurements each modem transmitted 2,000 messages at 300 ms intervals, and the results clearly demonstrate the energy-saving potential of RedCap technology.

The key findings of the measurements are:

- The RedCap modem consistently consumed less energy than the legacy 5G modem across all tested conditions.
- At 0 dB attenuation, the RedCap modem used less than half the normalized energy compared to the 5G modem.
- This trend continued when increasing the attenuation.
- Energy per message was significantly lower for RedCap, ranging from 13.55 µWh to 14.11 µWh, compared to 32.04 µWh to 35.02 µWh for the 5G modem.
- A modest increase in energy consumption was observed as signal quality decreased, which aligns with expectations.

Overall, while RedCap is not yet fully optimized, it already demonstrates a clear advantage in energy efficiency over traditional 5G modems. As the technology matures and eRedCap modules become commercially available, further improvements in power consumption are expected, potentially making RedCap a strong candidate for energy-constrained IoT applications.

[15] For Wi-Fi deployments on higher frequency bands (e.g. 5GHz, 6 GHz) or private 5G deployments at 3.5 GHz with wider bandwidth of 100 MHz more than 4 antennas could be needed.

Acknowledgements

This work was performed as part of the “Do IoT voor de Glastuinbouw” project, co-funded by Metropoolregio Rotterdam Den Haag (MRDH), Provincie Zuid-Holland and Gemeente Westland. TNO worked in collaboration with project partners: Do IoT Fieldlab, MCS, Tomatoworld, and RoboHouse.

We thank the following TNO colleagues for their assistance during this project: Peter Hoefsloot, Pawel Mackowiak, Borgert van der Kluit, Ewout Brandsma, Tim Bergman, Floris Drijver and Meerten Versluis.

ICT, Strategy & Policy

Anna van Buerenplein 1
2595 DA Den Haag
www.tno.nl