



BE GREEN

Energy-Efficient and Intelligent RAN Evolution for B5G and 6G

BeGREEN Project Final White Paper

March 2026



Co-funded by
the European Union

6GSNS

Editor(s):

Jesús Gutiérrez (IHP)

Author(s)/Contributor(s):

Mir Ghoraishi (GIGASYS)

Miguel Catalán-Cid (I2CAT)

Simon Pryor (ACC)

*This White Paper presents the results, conclusions, and perspectives derived from the research and innovation activities carried out by the **BeGREEN consortium** throughout the project. The findings reflect the collective work of the project partners and are based on the design, implementation, and validation of the proposed solutions.*

*This work is supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101097083, **BeGREEN project**. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS-JU. Neither the European Union nor the granting authority can be held responsible for them.*

1 Introduction

The European Union (EU) is committed to ensuring that the development and deployment of future 6G networks align with the objectives of the European Green Deal and broader sustainability ambitions. This requires a combination of research and innovation (R&I) in energy-efficient technologies, regulatory and policy measures to reduce the carbon footprint of digital infrastructures and services, and close public-private cooperation to shape a green digital future. In this context, the EU, the 6G Infrastructure Association (6G-IA), and the Smart Networks and Services Joint Undertaking (SNS JU) are strongly supporting R&I efforts to develop technologies that contribute to a climate-neutral and sustainable Europe. Communication networks are a key enabler of both the digital and green transitions, and their evolution must ensure that these transitions remain mutually reinforcing [1]. As the mobile ecosystem progresses towards 6G, energy efficiency has emerged as one of the defining challenges of next-generation network design. Standardisation discussions on 6G started in early 2025, initiating a broader process of studies, workshops, and industry alignment around the future mobile generation. This process is expected to mature further during 2026, while Standard Development Organisations (SDOs) continue to discuss the architectural and technological building blocks of Radio Access Networks (RANs) and Core Networks in line with the International Telecommunication Union (ITU) IMT-2030 framework [2].

Future 6G systems are expected to support higher performance, denser deployments, greater flexibility, and a much broader service portfolio. Under these conditions, energy efficiency becomes a critical design requirement rather than a secondary optimisation objective. Intelligent RAN concepts, based on Artificial Intelligence (AI) and Machine Learning (ML), provide new opportunities to manage resources dynamically and reduce energy consumption while preserving performance. However, the introduction of AI-driven control and orchestration also adds compute and processing overheads, which must themselves be managed in an energy-aware manner. Within mobile networks, the RAN has long been recognised as the main contributor to total energy consumption, with the radio domain representing the most energy-intensive part of the network. In particular, RF chains, power amplifiers (PAs), and transceivers account for a substantial share of base station energy consumption. At the same time, the growing adoption of virtualisation, cloud-native architectures, and AI-assisted control is shifting part of the energy burden towards compute and orchestration layers. As a result, future sustainable mobile networks must optimise energy consumption holistically across radio, compute, and control domains.

During the evolution from 5G towards 6G, an important intermediate step has been the enhancement of 5G building blocks beyond current specifications, particularly to support improved energy efficiency, advanced programmability, and evaluation in dense and heterogeneous environments. In this context, the Open Radio Access Network (O-RAN) architecture has emerged as an important enabler, thanks to its open interfaces, modular design, and support for AI-driven programmability.

Despite this progress, several major challenges still delay the realisation of truly energy-efficient B5G and future 6G networks:

1. **Architectural transformation:** making informed decisions on the adoption of new architectures and deployment models that reduce overall RAN energy consumption.
2. **Resource heterogeneity:** managing and allocating resources efficiently in increasingly heterogeneous network environments.
3. **Energy-aware control of RUs:** enabling intelligent and dynamic operation of Radio Units (RUs), which remain the most energy-consuming elements of the RAN.
4. **Energy cost of AI/ML:** ensuring that the introduction of AI/ML techniques for optimisation does not offset the intended gains through excessive compute, memory, and training requirements.

5. **Integration of enabling technologies:** combining baseline 3GPP and O-RAN technologies with new capabilities, such as RIS, relays, ISAC, and advanced edge/cloud orchestration, to achieve sustainable network operation.

Against this background, one of the central research questions for 6G is how future wireless systems can simultaneously deliver very high performance and strong sustainability. This is precisely where BeGREEN makes its contribution: by treating energy consumption as a core system design metric and by demonstrating how multiple architectural, algorithmic, and hardware innovations can be combined to support sustainable, intelligent, and scalable future mobile networks.

1.1 BeGREEN proposition

BeGREEN, together with its peer SNS Phase 1 projects, was conceived as a response to the need for early-stage 6G research and for the development of sustainable and trustworthy technologies. The project focuses particularly on the RAN and addresses how it can be redesigned so that different strategies across multiple RAN elements combine to form an energy-smart network.

To address the challenges outlined above, BeGREEN adopts the following core design principles:

1. Energy efficiency as a primary design principle and KPI

The BeGREEN architecture is designed with energy efficiency as a first-class system objective. This enables the definition of quantitative indicators, such as energy scores and energy ratings, to benchmark solutions and to assess the energy footprint of AI/ML-driven control and orchestration mechanisms.

2. Architectural innovation and novel algorithmic solutions for sustainability

BeGREEN investigates innovative architectures and advanced optimisation methods to reduce the environmental impact of cellular networks. Examples include the evaluation of cell-free (CF) networks, as well as the integration of enabling technologies such as Reconfigurable Intelligent Surfaces (RISs), relays, and Integrated Sensing and Communication (ISAC).

3. Energy-efficient implementation of RAN elements

The project explores how hardware acceleration and advanced processing architectures can reduce the energy consumption of computationally intensive RAN functions.

4. Advanced AI/ML built on the baseline O-RAN architecture

At the heart of BeGREEN lies the Intelligence Plane, which provides a programmable and extensible framework for integrating AI models and novel rApps/xApps to support energy-efficient RAN operation without compromising performance.

5. Integration of enabling technologies and new interfaces

BeGREEN extends the baseline O-RAN architecture to integrate Edge and Core domains and to support technologies such as RIS and relays, which are currently beyond the scope of traditional O-RAN implementations.

Taken together, these principles position BeGREEN as a holistic approach to energy-efficient network evolution, combining architectural innovation, intelligent control, and practical validation to support sustainable B5G and 6G systems.

This White Paper presents the BeGREEN vision for energy-efficient and intelligent RAN evolution, the main architectural and technological innovations developed in the project, the key validation outcomes, and the lessons and opportunities that emerge for future research, standardisation, and industrial uptake.

2 BeGREEN System Design

BeGREEN was a pioneering project in enhancing 5G and O-RAN building blocks and interfaces while also incorporating new capabilities for joint evaluation with standardised solutions. Its overall goal is to improve the reliability, adaptability, and sustainability of future RANs.

Throughout the project, BeGREEN has developed and validated a range of methods, tools, and solutions to maximise network performance while limiting energy consumption. These include mechanisms to reduce transmit power, solutions integrated into compatible O-RAN RUs, and enabling technologies such as RIS, relays, and ISAC. Together, these approaches support context-aware and adaptive power optimisation under varying traffic and environmental conditions.

Figure 1 depicts the final BeGREEN architecture, which consists of the following main elements:

- **BeGREEN Intelligence Plane**, the central element of the overall architecture, acting as a cross-domain management and optimisation entity based on advanced AI/ML technologies;
- **Enabling RAN technologies**, including RISs, relays, and ISAC-capable RUs, which extend the capabilities of the RAN;
- **Enhanced RAN units**, namely the O-CU, O-DU, and O-RUs, incorporating energy-aware mechanisms and improved implementations;
- **BeGREEN new and extended interfaces and control components**, introduced to support the management and monitoring of Edge and Core domains and to integrate RAN technologies that are not currently addressed by baseline O-RAN specifications.

Taking O-RAN as the baseline architecture, BeGREEN has introduced several architectural extensions under the umbrella of the Intelligence Plane [3][4][5]. Once incorporated into the O-RAN framework, these enhancements can be exploited by novel AI/ML-driven control mechanisms to reduce the energy consumption of both the RAN and Edge domains. In the RAN, the proposed energy-saving optimisations focus on virtual RAN (vRAN) resource allocation, dynamic 5G carrier deactivation and traffic offloading, and the use of relay and RIS technologies to improve coverage and efficiency. In the Edge and Core domains, resource allocation strategies are applied to User Plane Function (UPF) and Edge AI services, while data dimensionality reduction is investigated to lower the energy cost of AI/ML processing without compromising accuracy.

2.1 Intelligence Plane

The BeGREEN Intelligence Plane, central to the BeGREEN overarching architectural concept and depicted in the upper part of Figure 1, is a cross-domain management entity that integrates advanced AI/ML technologies to monitor, analyse, and optimize network operations [3] [4]. By leveraging the Intelligence Plane's AI Engine, Service Management and Orchestration (SMO), and RICs, BeGREEN addresses the multifaceted challenges of energy efficiency. These main components are described as follows:

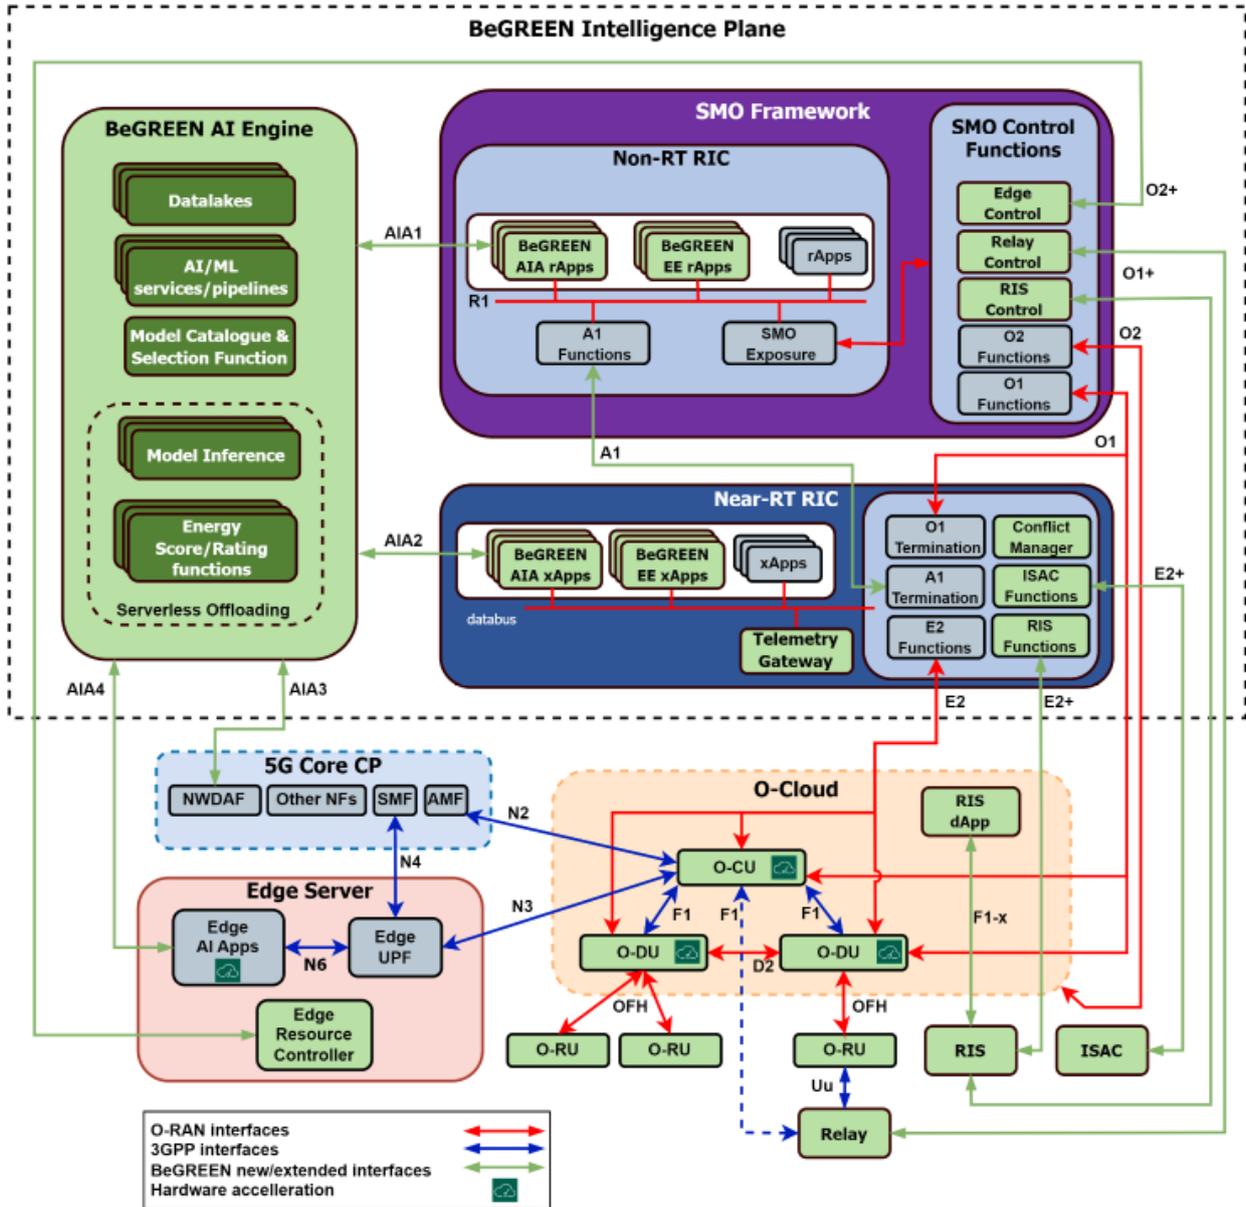


Figure 1: Final BeGREEN Architecture

- AI Engine:** The key element of the Intelligence Plane, providing the framework to implement and expose AI/ML services in a loosely coupled manner. This means that ML models are trained and hosted within the AI Engine, with their inference results exposed to rApps and xApps rather than being embedded directly in them. This approach promotes model reuse and allows ML and RAN control developers to focus on their respective domains of expertise, leveraging AI/ML and RAN management as independent yet interoperable components. Furthermore, the AI Engine framework supports serverless inference, enabling model offloading via serverless computing and hardware (HW) acceleration. This allows the benefits of the AI Engine to be utilized not only in the Non-RT and -Near-RT RIC domains but also in other domains, such as the 5G Core or Edge services, as illustrated in Figure 1 through the AI Engine Assist (AIA)-x interfaces.

Additionally, the AI Engine provides Energy Score and Energy Rating functions to identify network areas or components requiring energy-saving policies and to assess the benefits of performed optimizations. The Energy Score reports an absolute measure of energy efficiency (bits/Joule) based on data volume and energy consumption, while the Energy Rating monitors relative performance by comparing historical data from equivalent network entities, such as cells of the same vendor within

the same area. Based on these insights, control rApps determine appropriate A1 energy-saving policies and communicate them to the Near-RT RIC over the A1 interface.

- **SMO and non-RT RIC:** The Non-Real-Time RAN Intelligent Controller, hosts rApps for RAN optimization, utilizing AI/ML models trained on data collected from the emulated E2 nodes. These models are deployed in the AI Engine, and exposed through the AIA rApps. AIA rApps decouple ML model management from control loops, acting as a proxy between ML models and control rApps. During inference, AIA rApps function as data producers, exposing model outputs via the Data Management and Exposure (DME) service of the R1 interface.
- **Near-RT RIC:** The Near-RT RIC handles fast telemetry and control operations for the RAN. It uses an Energy Savings xApp to manage RAN configurations and implement energy-saving policies received from the Non-RT RIC. Integrated into dRAX¹, it features a Telemetry Collector that gathers data from both O-RAN and non-O-RAN compliant interfaces. This telemetry is processed into 3GPP-compliant messages, distributed via the dRAX data bus, exposed through the O1 interface, and pre-processed to be visualized in a Grafana dashboard. Additionally, the Near-RT RIC enforces E2-based RAN control, manages conflicts, and monitors RAN KPIs by interfacing with the emulated RAN environment.

AI/ML-assisted procedures have been developed in BeGREEN, including the Intelligence Plane itself, as summarized in the following points:

- **Dimensionality reduction:** This method proposes a solution to systematically reduce the input data required to train and retrain models such as predictors. The objective is to enhance the energy efficiency of the models without impacting the required model accuracy. Initial results are based on a real dataset provided by an MNO.
- **Compute resource allocation in vRAN:** Addresses the problem of compute resource allocation in virtualized RAN under shared computing infrastructure. According to an initial experimental characterization, proposes Reinforcement Learning (RL) based solution which adapts the compute resources allocated to virtual Base Stations (vBSs) according to network demands, avoiding over- and under-provisioning issues. The method is evaluated experimentally in a testbed.
- **AI/ML and data-driven strategies for energy-efficient 5G carrier on/off switching:** This solution considers scenarios with capacity and coverage cells, as is the case of current 5G Non-Stand Alone (NSA) deployments, in order to propose on/off switching and traffic offloading strategies. According to the real data from a MNO, available energy saving opportunities are studied, evaluating how different heuristics and ML-driven strategies could be applied to match them while ensuring a specific QoS.
- **AI/ML-based algorithmic solutions for relay-enhanced RAN control:** Proposes the utilisation of Relays and RUEs to enhance energy efficiency of areas with high traffic demands and poor propagation conditions, avoiding the installation of new cells or the increase of cell transmission power. Several AI/ML-based methods are defined to provide a complete solution, including CHD, fixed relay placement decision, identification of candidate RUs, and relay activation and deactivation. Initial evaluation is based on simulations according to a realistic scenario in a University Campus.
- **Traffic-aware compute resource management to enhance UPF energy efficiency:** This method deals with the dynamic management of the compute resources allocated to a UPF according to the traffic demand. According to the experimental characterization of a high-performance open-source UPF implementation, energy efficient strategies are proposed and evaluated. Additionally, proactive

¹ ACCELLERAN'S 5G DRAX™ CLOUD-NATIVE OPEN RAN SOFTWARE, <https://acceleran.com/accelerans-5g-drax-cloud-native-open-ran-software-now-available/>

management of the strategies is provided by ML models forecasting the traffic demand.

- Joint orchestration of vRANs and Edge AI services: Addresses the problem of optimizing the allocation of resources to vRANs and Edge AI services, considering the intertwined relationships and trade-offs between RAN and Edge configurations and their impact on performance. To solve it, and according to the results from an experimental characterization, an online learning algorithm formulated as a contextual bandit is proposed.
- Intelligence Plane validation: Presents the evaluation of the Intelligence Plane, focusing on the integration of the AI Engine and Non-RT RIC components through AIA rApps. The emulated RAN is implemented using Viavi's TeraVM AI RSG6. This tool enables the scalable and realistic evaluation of rApps and xApps by emulating O-RAN compliant and real-world deployments. It also generates 3GPP-compliant KPMs, enabling the collection of data for AI model training and testing. We developed and evaluated different rApps and xApps using this framework, such as a KPM producer rApp, dataset producer rApp, energy predictor rApp, load predictor rApp, throughput predictor rApp, energy saving xApp and smart handover xApp [13][14]².

2.2 BeGREEN new capabilities

From 6G use case definition and IMT-2030 framework, it is obvious the need to incorporate enabling technologies that offer new capabilities. These were already anticipated in 2022, when the project content and objectives were initially drafted. The following subsections outlines the technologies that BeGREEN identifies as key to improve energy efficiency in the RAN.

2.2.1 RISs

RISs have significant potential to improve energy efficiency. They can extend RAN coverage in urban areas, particularly in locations where propagation is blocked by buildings. RIS operate at extremely low power and, in some cases, can substitute for a RU to provide coverage extension.

A novel solution, namely Metasurface Absorption and Reflection for Intelligent Surface Applications (MARISA), leverages Hybrid Reconfigurable Intelligent Surfaces (HRISs), as plug-and-play devices with reflection and power-sensing capabilities [10]. The development represents an autonomous RIS that can optimize wireless communication signals without needing a dedicated control channel or continuous external power supply. MARISA optimizes the HRISs based on an appropriately designed codebook, which allows for the estimation of the BS-RIS and RIS-UEs AoAs in a distributed manner.

Also, awareness of UE location allows for more effective deployment of energy-saving strategies, such as reducing beam measurements and Channel State Information (CSI) reporting, which can conserve significant power. Current RIS-based indoor positioning approaches rely on probing reflection configurations from precomputed codebooks to maximize signal gain, but face challenges including limited resolution, low-gain changes due to strong direct paths, and phase misalignment from multipath interference. To overcome these limitations, BeGREEN introduces the "Echoes" solution, which integrates RIS with the BeGREEN Intelligence Plane to provide indoor user localization. Echoes infers user positions from power measurements of uplink reference signals, leveraging the non-singular property of RIS steering configurations to combine angle-of-arrival estimation with phase offset adjustments.

Echoes operates in an O-RAN-enabled scenario, where a small cell gNB interacts with the RIS through a locally deployed dApp, while an xApp in the near-RT RIC handles measurements, reconfiguration, and positioning computations. The workflow involves four phases: receiving location requests, collecting uplink power

² https://www.youtube.com/watch?v=a_4MaX7a1aY

measurements, tuning the RIS according to a codebook with phase offsets, and estimating user positions via an optimization procedure. Experimental validation with a 10×10 unit-cell RIS prototype operating at 5.3 GHz demonstrated the ability to accurately estimate user locations, with a Mean Absolute Error (MAE) of 2.15° for azimuth angles, 14.15° for elevation angles, and a distance MAE of 0.92 m [5]. These results confirm that Echoes can effectively exploit RIS for indoor localization while supporting energy-efficient RAN operation, highlighting the potential of RIS-enabled ISAC solutions within the BeGREEN architecture.

2.2.2 Relays

Regarding the studies about B5G RAN enhanced through relay nodes, we have mainly focused on energy efficiency improvements through relays nodes, where several conclusions were derived from the study performed in [7]:

- The insertion of a relay at an indoor position with high enough spectral efficiency can bring important energy savings, which increase with the required bit rate and depend on the power consumption model parameters. For a required bit rate of 50 Mbps, energy savings range between 50% and 80% depending on the model parameters, reaching up to 90% for larger bit rates.
- The usage of relays brings notable improvements in energy efficiency. Specifically, a properly placed relay can increase the number of positions on a floor achieving high efficiency: from over 85% of positions below 1×10^5 bit/J to nearly the entire floor reaching 2.6×10^5 bit/J.

AI/ML-based algorithmic solutions for relay-enhanced RAN control targeting enhanced energy efficiency have been also addressed [5], focusing on detecting coverage holes (CHs) and traffic hotspots to decide the placement of fixed relays and the activation/deactivation of fixed relays or Relay UE (RUE). Results, published in peer-reviewed conferences and journals, showed that relays can significantly boost spectral efficiency in certain scenarios, while reducing energy consumption (up to 40%) and Capital Expenditure (CAPEX) by lowering the number of required gNB deployments. Additionally, in some scenarios, The RUEs to mitigate some CHs provides performance and power consumption results close to the case of fixed relays and further reduces the deployment costs. An architecture to integrate relays and the required AI/ML models within O-RAN and the Intelligence Plane was also defined.

Finally, a comparison between the use of a RIS or a relay has been presented focused on assessing their achievable energy savings concerning the case in which none of these elements is used to enhance the coverage in certain areas [4]. It has been found that in general, the RIS provides higher energy savings when its pathloss with the base station is lower than 90 dB and for UEs located in a relatively narrow angular region for the RIS pointing direction, while for larger pathlosses or outside this region the relay outperforms the RIS. Also, an analysis in a realistic university campus scenario in which most of the poor coverage areas are found indoors has led to the observation that in most of the studied situations, the relay provides higher energy savings than the RIS, while the RIS only outperforms the relay in very specific conditions where it can be placed outdoors.

2.2.3 ISAC

Two approaches concerning ISAC were followed in BeGREEN focused on the improvement of the energy efficiency of the network. The first one is to use ISAC to assist the beam training process to make it more attractive for use in different RUs and wireless technologies. Our analysis showed that this approach is superior to the State-of-the-Art (SotA) approaches, which will solve the issues due to which electronic beam steering technologies are not widely and fully used. The second approach is to use ISAC to detect potential users, therefore enabling better network resource allocation and will reduce the idling time of the devices in the RAN. The obtained results show excellent results in terms of range and angular resolution, allowing for easy estimation of the number of potential users.

In the context of BeGREEN it was proven and demonstrated that ISAC can be used to detect potential users and, based on their distribution, to optimally assign network resources. Making wireless access technologies able to sense the environment and detecting moving objects and people as well as their features (size, speed and direction), type (i.e. human/object), etc. In BeGREEN, different algorithms have been developed to implement these sensing functionalities. These algorithms are used to create different heat maps representing the detected obstacles, the moving objects or persons, etc.

A Sub-6 ISAC solution has been developed and experiments both in a controlled environment and in indoor scenarios have been performed [10]. With these experiments, it was shown that the sensing functionality of the radio interface has the potential to be used for improving the beam search and beam tracking functionalities of the PHY layer, making them more attractive for use in the next generation radio networks. This should enable using higher antenna gains which will indirectly improve the power efficiency of the network.

In addition, we leverage sensing information to increase the efficiency of the beamforming. We considered a Multiple-Input Multiple-Output (MIMO) radar system with multiple users and estimate the Directions of Arrival (DoAs) of the users using the multiple signal classification (MUSIC) algorithm at the receiver side of the MIMO radar. We then use this information to design the beam codebook, namely, we locate more beams on the estimated user areas to decrease the beamwidth and increase energy efficiency, while covering the idle areas by only beam. Numerical examples have demonstrated the effectiveness of the proposed approach; the success rate and an improvement in energy efficiency of the beam search compared to the case where no sensing parameter estimation is performed.

The RIS-assisted ISAC solution at Sub-6 GHz achieved **0.7 m range resolution** and **>90% person-count precision**.

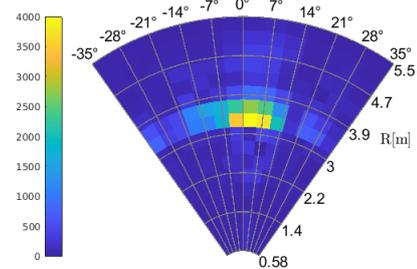
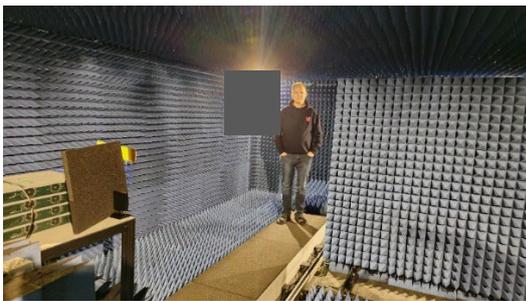


Figure 2: ISAC setup with a person in front of the sensing system

2.3 Energy-Efficient RAN Components

This section presents power optimization techniques that target the RAN components of the O-RAN architecture. On the radio side, we introduce novel energy-saving mechanisms integrated into the O-RU, aiming to reduce the power consumption of the most energy-demanding components of the RAN. On the computing side, we demonstrate how HW acceleration and the adoption of advanced processing architectures can significantly improve the energy efficiency of computationally intensive functions in the O-DU and O-CU, as well as in the RIC. Overall, this section highlights that achieving substantial energy savings in future RANs requires a joint optimization of radio HW, baseband processing, and computing platforms, rather than focusing on any single component in isolation. BeGREEN O-RU

In the RU domain, which is the major power consumer in cellular network, several power saving techniques were explored and simulated with promising results such as AI-based Digital-Pre-Distortion (DPD) and Envelope Tracking (ET), which can lead to more than 60% RF Power Amplifier (PA) energy savings. DPD is a common technique to cope with signal Peak-to-Average Power Ratio (PAPR) it is mainly used to increase the RF PA efficiency. PA ET is a method in which the power supply voltage applied to the PA is continuously

adjusted ensuring the PA is operating at peak power efficiency [10].

In addition, a novel technique called PA Blanking was developed. PA blanking stops the PA energy consumption for symbols without active data allocations. When resources of a downlink symbol have no allocated data, the related RF PA will be blanked (momentarily shut down) for the symbol duration. To perform the blanking, a PA blanking module needs to be integrated in the RU [12]. By concentrating the allocated Resource Element Blocks (REBs) into a subset of symbols, the DU scheduler maximizes the number of non-allocated symbols, thus increasing the opportunities for blanking. Measurements demonstrate that more than 90% of the PA energy can be saved during blanked symbols. At the system level, this translates into energy savings exceeding 80% under low traffic load, over 50% at average traffic conditions, and above 30% during busy-hour operation. The PA switching time is below one microsecond, enabling effective blanking for more than 90% of the symbol duration even with short slot lengths. These results confirm that PA blanking is a practical and highly effective mechanism to significantly reduce RU power consumption under realistic traffic patterns, complementing AI-based DPD and ET techniques for holistic energy-efficient RU operation.

2.3.1 BeGREEN O-DU

The Baseband processing of the received waveforms in the DU is highly complex and requires a large amount of energy. The highest energy consumption is in the Low-Density Parity Check (LDPC) decoder. This means that optimizing the LDPC decoder has the highest potential for energy efficiency improvement. Within the duration of the project we have extended the implementation and testing of the LDPC decoder beyond the implementation carried out on x86 and ARM, the LDPC decoder is implemented on a GPU. Several optimization iterations are carried out to be able to efficiently utilize the GPU and use all its available threads. The results showed that the LDPC wireless performance of the GPU was the same as the x86 and the ARM as expected. The power consumption of the GPU was higher when compared to ARM CPU using multiple cores. However, when comparing to using only a single ARM core the GPU was much more energy efficient.

In the DU acceleration demo, we compare the power consumption of a sphere decoder running on CPU and on GPU. The sphere decoder advantage in wireless performance for the scenario being tested in the PoC is shown in [11], where we can see a 2-dB improvement over the basic linear receiver for the case where we have correlation of 0.03 for the Tx antennas and the Rx antennas. This applies to both the CPU and the GPU implementations. As for **power consumption**, when running the 120 Resource Blocks (RBs) sphere decoder scenario on CPU the power consumption is 26W, and when running on GPU the power consumption is 24W, which is similar. However, the **running times** are dramatically different. It takes the CPU **5.5 ms** to complete the task where the GPU completes it in **0.5 ms**. Hence, the GPU energy score is **100 kbit/J** and the CPU energy score is **9 kbit/J**, resulting in a 10x improvement. To summarize, this demo shows that for complex L1 tasks, such as the sphere decoder, GPUs have the potential to significantly reduce power consumption.

2.3.2 BeGREEN O-CU and RIC

The CU acceleration shows the transition from x86 to ARM architectures to improve energy efficiency and performance in O-RAN networks. The CU is crucial for managing control functions and user data transmission. Given the significantly higher traffic volume handled by the CU-UP, the porting process prioritizes the Arm architecture due to its superior power efficiency and processing capabilities. This transition involves adapting C++ code to ARM using a cross-compilation toolchain and deploying the CU components as Docker images for both x86 and Arm platforms. The integration with Kubernetes further facilitates seamless multi-architecture support and optimized resource utilization. Initial test results reveal that while x86 shows better energy efficiency in scenarios with a limited number of simultaneous users, the Arm architecture excels in large-scale, high-throughput environments, achieving greater overall power savings. Specifically, as the data throughput increases, the energy savings with Arm become more pronounced compared to x86, as this is

attributed to Arm's efficient handling of parallelized tasks. A collaborative approach between project partners was pivotal in these advancements. Joint efforts in testing and validation have shown that the aggregated power reduction when using Arm covers a wide range of scenarios, making it a highly suitable choice for future scalable and energy-efficient network deployments. For some of the methods initial power measurement results are shared, which show significant power reduction compared to legacy network for some of the scenarios.

Initial measurements confirmed that the **CU** could operate effectively on both x86 and ARM servers under realistic RAN conditions. In single UE scenarios, x86 showed approximately 22% better energy efficiency, making it suitable for light-load environments. However, when scaling to higher UE densities, the ARM platform outperformed x86, achieving up to 10% energy savings and demonstrating better scalability under sustained traffic. This validated ARM's suitability for dense deployments where cumulative energy consumption becomes critical.

Performance testing of the **Near-RT RIC** on both architectures revealed distinct behavioural differences. ARM servers exhibited significantly lower average power consumption (up to 33% less) during computational tasks such as matrix multiplication and hashing, with consistent scaling across cores, as discussed in D3.3 section 2.3 [12], and as shown in Figure 3. Conversely, x86 systems delivered lower latency and better performance for small, time-sensitive operations. These trade-offs highlight a design space where ARM can support energy-sensitive functions, while x86 remains advantageous for ultra-low-latency workloads.

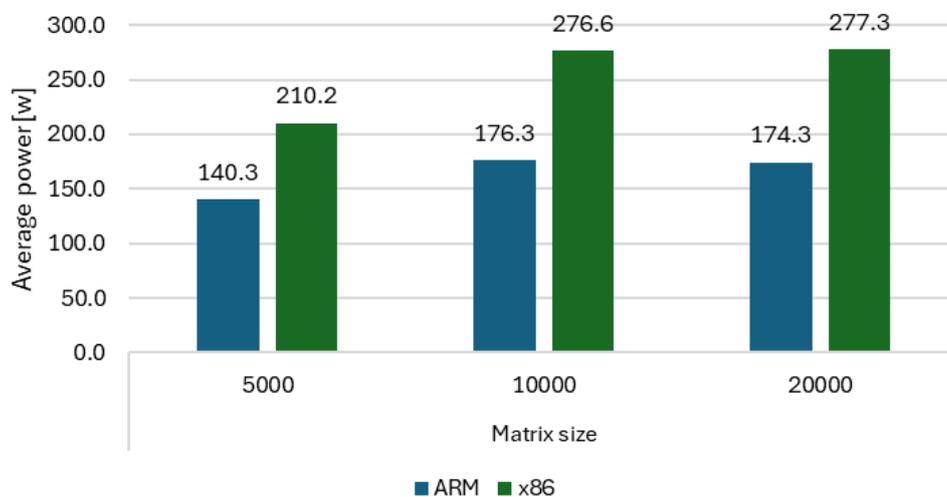


Figure 3: Average power consumption comparison for ARM and x86

3 Validation of BeGREEN Propositions

The baseline configuration for the assessment on the achieved improvement in terms of energy savings consider, as depicted in Figure 4, a traditional cell-like RAN deployment with no RIS and relays available, with O-RUs that do not feature any enhancement of those mentioned in sections 2.2.3 and 0, no acceleration or implementation of O-DU, RIC and O-CU functionalities in a different implementation rather than x86, existing O-RAN management and orchestration platforms and no energy efficient mechanisms or implementations at the UPF, edge and core.

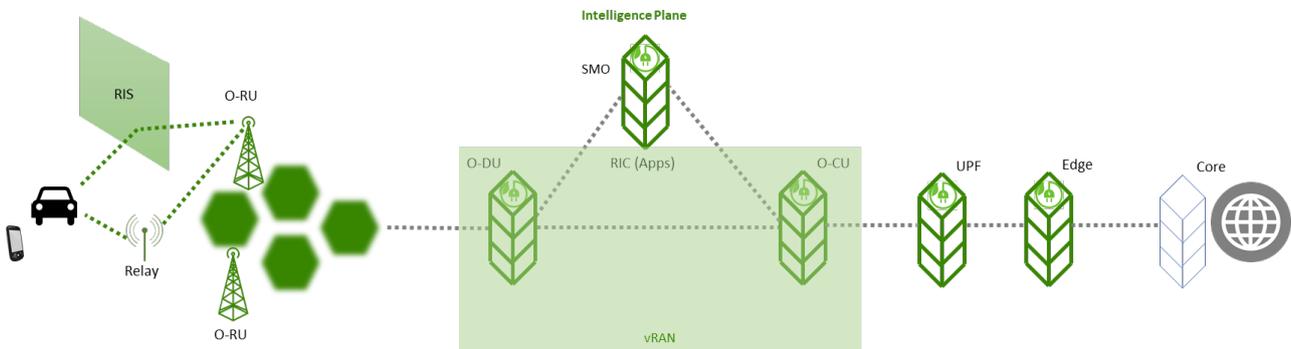


Figure 4: Sketch of all involved components for an energy-efficient RAN

BeGREEN propositions to make this RAN more energy efficient do provide the following advantages and yield the following energy savings:

- RIS energy-efficiency improvements can be up to **80%**, value which is dependent on the scenario, system design, and optimization strategy.
- Different algorithmic solutions with the purpose to identify the presence of CHs in cellular networks and mitigate them by means of relays have been proposed. The **CHs** can be addressed with a power consumption reduction in the range between **35%-70%** depending on the BS and the relay power consumption model. The possibility of ensuring **Relay Control** brings an energy reduction of **100Wh/day**.
- The proposition of CF architectures for future 6G networks has been addressed in the context of BeGREEN with the goal of reducing energy consumption. In particular, a CF system was compared to a traditional Centralized MIMO (C-MIMO) deployment in a realistic urban-dense scenario. It was concluded that a CF architecture reduces the transmit power by **30 dB** compared with a traditional MIMO deployment. In addition, the CF system provides functional Bit Error Rate (BER) for 8 streams, making this technology suitable for high-data-rate applications.
- Novel RU-related power optimizations, such as power amplifier (PA) blanking, envelope tracking (ET), and AI-based digital pre-distortion (DPD) algorithms have been developed to improve the energy efficiency at the RU. PA-blanking confirmed **64 % PA power savings**, dropping RU PA draw from 5.41 W to 1.94 W using scheduler-driven blanking. AI-based DPD and ET, which can lead to more than **60% RF PA energy savings**. The conclusion of the research was that DPD can be implemented with a Deep Neural Network (DNN) with only about 16 neurons and 2 layers, which may be considered small. However, when approaching the saturation point, the random carrier phase effect is more pronounced and reduces the Digital Pre-Distortion (DPD) performance. Changing the size of the DNN does not bring any improvements. It is also noted that a DPD amplifies the signal and, therefore, the DNN implementation will require more digital resources, i.e., larger multipliers, wider buffers, etc. We concluded that ET is a promising technology to reduce PA power consumption by a factor of

10dB. However, a good performance is expected only with a DPD that is also aware of the power supply input.

- **BeGREEN** ISAC activities at Sub-6 have achieved **0.7 m range resolution** and **>90% person-count precision**. The angular resolution, however, heavily depends on number of antennas used, meaning that an individual can be seen under a certain angle depending on the distance to the sensing solution. To really reap the benefit of perceiving a faithful representation of the environment ahead of the ISAC solution, an adaptive background modelling is required. Despite the non-3GPP nature of the system and its (in principle) incompatibility with an O-RAN system, opportunities to gather sensing information stemming from it have been devised and an initial implementation is available. This allows providing sensing information to the Near Real-Time RIC for visualization/post-processing.
- The effectiveness of the **BeGREEN** Intelligence Plane has been assessed, achieving **energy savings of up to 40%** without compromising network performance. The **BeGREEN** Intelligence Plane is set to continue evolving, with future developments aimed at expanding its capabilities and real-world applicability. Planned advancements include enhancing AI models to address more complex traffic scenarios, leveraging VIAVI TeraVM AI RSG to generate realistic data for even more accurate simulations. Evaluating the **BeGREEN** solution in scenarios that closely mimic real deployments, using the ability of TeraVM AI RSG to simulate thousands of UEs and cells, is a critical next step. Integration into live networks will validate the system's performance in operational environments, ensuring its reliability and adaptability. Additionally, collaboration with industry partners will drive the adoption of energy-efficient O-RAN solutions across a broader ecosystem, amplifying its impact. This journey underscores how AI-driven automation can revolutionize mobile network management.
- Running the LDPC algorithm on an ARM-based platform leads to **~20% power reduction** in average compared to legacy x86 platforms.
- HW-accelerated DU Graphics Processing Unit (GPU)-based sphere decoder **cut execution time by 11×** and **improved energy efficiency from 9 kbit/J (CPU) to 100 kbit/J**—evidence that off-the-shelf accelerators are pivotal for green Layer-1 processing.
- Energy-aware **CU & O-RAN RIC** was ported to ARM servers yielded up to **10% savings** under high traffic, with 33% lower idle power versus x86, validating edge-grade, low-carbon cloud-native RAN.
- Allocation of virtualized resources in vRANs, analysing the impact of shared O-Cloud resources on the performance of virtual Base Stations, and providing an AI-driven optimisation which aims to minimize the operating cost of the vRAN infrastructure. The second use case focus on the utilisation of CPU resources by virtualised UPFs hosted in Commercial Off-The-Shelf (COTS) servers, characterising the relationship between CPU states, UPF performance and energy consumption, and developing AI-driven solutions to optimise energy efficiency. In both cases, we successfully designed and evaluated AI/ML-driven optimisations that reduced energy consumption by **~20-30%** without degrading network performance. These results have been validated through several high-ranked publications.

Collectively, the PoCs demonstrate an end-to-end 6G blueprint capable of achieving **20-35 %** system-level energy reduction without degrading user experience, supporting the project's vision of a carbon-aware mobile network stack. This highlights the commitment of **BeGREEN** to sustainable and intelligent network design, paving the way for scalable, high-performance B5G and 6G networks that align with global environmental goals.

4 Key Insights and Forward-Looking Perspectives

This chapter highlights some of the lessons learnt through the integration and demonstration of the different technologies considered in the project, as well as some of the open points that to be addressed in the future.

4.1 Technical potential

From the start of the BeGREEN project, the claim was that the RF part was typically reported to be the dominant component within RAN energy use. Recent hardware advances in that part – such as ET, DPD, adaptive PA power control, and beam-level optimization – help reducing RF energy consumption, especially under variable traffic. BeGREEN’s own measurements highlight such techniques that adapt PA linearity and channel usage to reduce power draw while maintaining RF performance. Currently, industry vendors report that energy savings of ~39% or more in newer 5G radio generations due to RF and system optimization. At the same time, modern massive MIMO radios can achieve much higher energy efficiency per delivered bit, lowering the RF share even as overall capacity increases.

Given these improvements in the RF part, conversely, in AI-RAN deployments (O-RAN, vRAN, Cloud-RAN), virtualization and AI workloads shift energy consumption toward DSP/compute layers, especially:

- Baseband processing (vDU/vCU) running on general-purpose CPUs, NPUs, GPUs.
- AI inference and orchestration functions continually adapting radio parameters.
- Edge AI workloads co-hosted with RAN compute.

Recent analyses point out that AI-based optimization can reduce RAN energy by ~10-12% or more through smarter power control and sleep mode management. Moreover, studies focusing on joint radio plus compute optimization show that compute energy can grow relative to RF, especially as more intelligence and dynamic traffic shaping are deployed.

This updated breakdown reflects that RF is still a major consumer, but its relative share is reduced due to faster switching, sleep modes, adaptive PAs, energy-efficient massive MIMO, together with virtualization and AI tasks increasing compute energy needs. These shifts are consistent with both BeGREEN’s own technological focus and broader literature highlighting rising compute influence and AI’s role in RAN energy. The original BeGREEN proposition that a significant part of RAN energy is consumed by RF units remains directionally correct, but needs refinement in the era of Intelligent RAN:

- RF energy share is still large but likely reduced to ~40–55% of RAN energy due to smarter RF technologies and dynamic operation.
- Compute and AI workloads increasingly constitute ~25–40%, reshaping the energy paradigm.
- A full sustainable strategy must target cross-domain optimization across RF, compute, cooling, and AI logic.

The main technical lessons emerging from BeGREEN can be summarised as follows:

- **Cell-Free MIMO** offers significant energy-efficiency advantages over traditional cellular systems, particularly by reducing transmission power at the cell edge and lowering energy per transmitted bit. However, practical deployment challenges remain, especially with regard to fronthaul, synchronisation, infrastructure cost, and the overall energy cost of distributed access-point deployment.
- **O-RAN interoperability still presents practical challenges**, especially when integrating heterogeneous RUs and DUs and when supporting advanced or non-conventional MIMO schemes. Open interfaces alone do not guarantee straightforward integration.

- **Hybrid compute strategies are essential.** ARM platforms offer superior energy efficiency under sustained traffic loads, while x86 remains more appropriate for latency-sensitive functions. This supports a hybrid deployment vision in which performance and energy efficiency are jointly optimised.
- **RIS-assisted sensing and localisation are feasible and valuable,** even though current hardware constraints still limit some capabilities. These functions can support more energy-aware coverage and resource-management strategies.
- **Traffic regularity and strong load–energy correlation support predictive control,** particularly for carrier on/off switching. However, the interaction between 4G and 5G layers must be managed carefully to avoid QoS degradation during traffic offloading.
- **AI/ML-based energy optimisation shows strong promise,** but future work must explicitly account for the energy overhead of AI itself. This is especially important as AI-RAN concepts evolve and larger models, accelerators, and third-party AI workloads are introduced into the network.

Overall, and in line with BeGREEN’s assessments, future 6G energy-saving ambitions target approximately 40–50% reduction in RAN energy consumption compared with current 5G baselines through a combination of hardware innovation, architectural change, and intelligent control.

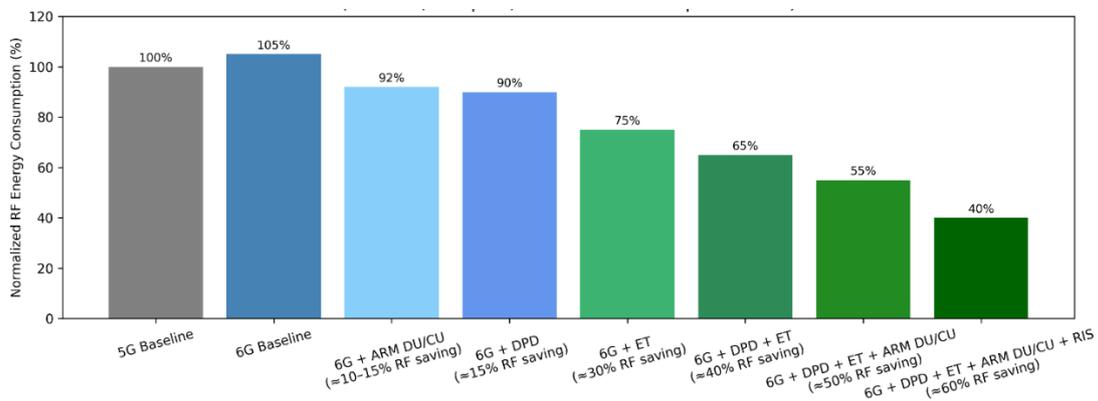


Figure 5 Normalized RAN energy consumption in 5G and intelligent 6G radio access networks. Energy savings are achieved through BeGREEN techniques

From a sustainability perspective, the project highlights several important lessons:

- the historical emphasis on RF power must be updated, as compute and orchestration layers are becoming increasingly relevant;
- energy optimisation must be holistic and cross-domain;
- AI itself must be energy-aware, otherwise gains in RAN efficiency may be offset by AI processing overheads;
- future energy monitoring and optimisation frameworks must capture dynamic and workload-dependent behaviour rather than focusing only on static RF consumption.

4.2 Standardisation potential

BeGREEN has generated several results and architectural insights with clear potential to inform future standardisation and pre-standardisation activities.

A first important contribution is the concept of the **BeGREEN Intelligence Plane**, which provides a unified AI/ML framework across RAN, Edge, and Core domains. By decoupling model hosting and inference from control logic, it enables the reuse of AI outputs across multiple optimisation functions and domains. This

architectural principle could help inform future O-RAN and 3GPP discussions on AI-enabled network management, energy-aware orchestration, and shared exposure mechanisms. It also supports the idea of cross-domain service and data exposure to verticals, in line with broader industry initiatives such as GSMA CAMARA.

Second, BeGREEN's work on **energy-aware control and telemetry** points to the need for richer information and service models to support energy-related monitoring, policy enforcement, and closed-loop optimisation. This applies in particular to the exposure of telemetry, control actions, and energy KPIs across the O-RAN stack.

Third, the project's work on **RIS, relays, and ISAC integration** provides useful input on how emerging non-conventional RAN components can be interfaced with the O-RAN framework. In particular, the extension of sensing capabilities towards the DU and RIC domains, and the exposure of sensing telemetry via extended E2-based mechanisms, could provide valuable input to ongoing discussions in ETSI ISG ISAC and related pre-standardisation fora.

Fourth, BeGREEN's work on **energy-efficient O-RU implementation**, including DPD, ET, and PA blanking, highlights the need for implementation-aware flexibility within standardised RU frameworks. While DPD itself is not standardised as a specific algorithm in O-RAN, the performance requirements imposed on O-RUs often imply the need for such techniques. BeGREEN's results provide useful evidence for implementation and deployment trade-offs, especially for higher-power RF modules.

Several relevant standardisation and industry venues can be considered for future exploitation of BeGREEN outcomes, including:

- **O-RAN Alliance**, particularly with regard to AI-enabled control, telemetry exposure, energy-aware rApps/xApps, and the integration of non-conventional RAN functions;
- **3GPP**, especially in relation to energy-aware RAN evolution, management, and future 6G study items;
- **ETSI ISG ISAC**, for contributions related to integrated sensing and communication interfaces and service models;
- **Telecom Infra Project (TIP) OpenRAN Project Group**, both in component-level and integrated deployment workstreams.

Overall, while BeGREEN was not a pure standardisation-driven project, it has produced a set of technically relevant results, design directions, and architectural concepts that position it well for future standardisation contributions and ecosystem influence.

4.3 Business potential

BeGREEN demonstrates significant business potential by transforming energy efficiency from a cost-control measure into a strategic capability for future Intelligent RAN and 6G systems. Through AI-driven optimisation, cross-domain orchestration, advanced power-saving mechanisms, and carbon-aware operation, BeGREEN enables energy savings in the range of 20–40% in RAN operation, depending on deployment conditions and maturity. This can improve profitability, accelerate return on investment for 5G evolution and 6G roll-out, and reduce exposure to volatile energy costs.

Beyond cost reduction, BeGREEN opens new revenue opportunities. Energy-intelligent RAN sites, particularly those evolving towards AI-RAN and integrated edge computing, can participate in smart-grid programmes such as demand response and peak shaving. Base stations equipped with batteries, renewable integration, and intelligent workload scheduling can move from being passive energy consumers to active participants in the energy ecosystem. This creates potential new revenue streams and strengthens operators' position within broader energy-transition strategies.

BeGREEN also supports Environmental, Social and Governance (ESG) goals, which are becoming increasingly important for investors, regulators, and enterprise customers. Carbon-aware scheduling, energy transparency KPIs, and green service-level agreements can enable operators to offer more sustainable connectivity services, support sustainability reporting, and improve access to green financing mechanisms.

For equipment vendors, BeGREEN creates differentiation opportunities through energy-aware controllers, carbon-intelligent orchestration platforms, and hardware–software co-design for low-energy operation. In this context, energy performance becomes a competitive parameter alongside throughput, latency, and coverage. Software-based energy optimisation services and AI-enabled orchestration functions also create opportunities for recurring revenue models and longer-term operator partnerships.

Cloud and AI infrastructure providers may also benefit, as Intelligent RAN architectures increasingly rely on energy-efficient accelerators and optimised AI pipelines. BeGREEN's focus on efficient AI models, workload consolidation, and hardware-aware deployment supports the market for low-power accelerators, sustainable edge platforms, and tighter telecom–cloud convergence.

Private 6G and industrial network deployments represent another promising business domain. Energy-aware RAN solutions can reduce total cost of ownership and align with sustainability objectives in vertical sectors such as manufacturing, logistics, utilities, and smart cities. In such settings, renewable-aware operation and guaranteed energy-performance metrics can become a differentiating factor.

In the longer term, BeGREEN also points towards new business models such as energy-aware service-level agreements, carbon-based differentiation of connectivity services, and monetisation of shared AI-RAN infrastructure. As the RAN evolves into a programmable compute–network platform, the value proposition shifts from maximising traffic capacity alone to maximising intelligent service delivery per unit of energy.

Overall, the business potential of BeGREEN lies in enabling a structural change: from energy being treated as a fixed operating cost to energy becoming a managed, optimised, and potentially monetisable resource in future 6G systems.

4.3.1 Cell-Free MIMO

Cell-Free massive MIMO has strong long-term strategic potential, especially for private networks and future 6G evolution. In the short term, however, its most promising commercial pathway lies in high-value vertical environments rather than direct competition with established macro cellular infrastructure.

Its main technical strength is the ability to provide more uniform performance without classical cell-edge degradation. At the same time, fronthaul bandwidth, latency constraints, synchronisation complexity, and centralised processing scalability all affect economic viability. Cell-Free MIMO is therefore best understood today as a capacity-enhancing and coordination-oriented architecture that can coexist with macro deployments under favourable infrastructure conditions.

4.3.2 RF optimisation strategies

The business case for RF optimisation in O-RAN RUs is based on using advanced AI-driven techniques to improve the performance of DPD and ET modules in RF PAs. These solutions are largely software-defined and may support recurring licensing models.

The current TRL of the AI-based DPD and ET solutions is around TRL 3, corresponding to experimental proof of concept. Their future integration into commercial O-RU products will depend on whether the achieved power savings justify the cost increase associated with additional modules and control functions. Initial analysis suggests that these techniques are more likely to be cost-effective in higher-power RU products than in current small-cell products such as the RunEL Sparq-2025-SRU.

By contrast, PA blanking has reached around TRL 4, corresponding to technology validation in the laboratory. It is currently being incorporated into the RunEL Sparq-2025-SRU commercial RU product and is expected to become available to future customers as a differentiated energy-efficiency feature.

4.3.3 RIS

The BeGREEN results on RIS can support the evolution of current prototype solutions and help strengthen data-driven product portfolios for next-generation virtualised and O-RAN-compliant RAN systems. In particular, the project's work on joint AI/ML-based control of RAN and RIS can support the development of xApps/rApps and motivate the definition of new service and information models for O1, E1, and related interfaces where needed.

4.3.4 Energy efficiency optimisation in O-RAN

The proposed Intelligence Plane enables a clear separation between ML model development and control application development. This facilitates reuse, modularity, and easier integration into operational workflows. O-RAN assets developed within the project, including the Non-RT RIC and related rApps for energy optimisation, have clear exploitation potential through licensing and technology-transfer agreements.

The work on CU and RIC hardware acceleration also creates business potential by enabling more energy-efficient execution environments for intelligent xApps and rApps, including those targeting RU power management and broader energy-aware RAN control.

4.3.5 ISAC

ISAC also offers long-term exploitation potential, although its commercial maturity is currently lower than that of more established RAN energy-saving techniques. In the shorter term, one realistic pathway is to combine radar-like sensing capabilities with evolving 3GPP sensing functions in 4G/5G and pre-6G systems. These directions are already reflected in 3GPP Release-19 6G use cases and are expected to become increasingly relevant in Release-20 study items, as well as in industry initiatives such as ETSI ISG ISAC.

5 Conclusions and Future Work

The results achieved in BeGREEN demonstrate that introducing an entity such as the Intelligence Plane into a baseline 3GPP/O-RAN system, together with energy-efficient techniques applied across different RAN elements and enabling technologies, can bring substantial value to B5G and future 6G networks seeking flexibility, adaptability, and controlled energy consumption.

The Intelligence Plane was tightly integrated with O-RAN and enabled the offloading of AI/ML complexity to a dedicated AI Engine, while supporting the integration of enabling technologies such as RIS, relays, and ISAC. This confirms the value of a cross-domain and extensible framework for energy-aware network operation.

Taken together, the BeGREEN propositions provide a credible end-to-end blueprint for a carbon-aware mobile network stack capable of achieving system-level energy reductions in the order of 20–35% without degrading user experience. In addition, several of the project results have relevance for ongoing discussions in O-RAN, 3GPP, and related ecosystem activities, while a number of the developed technologies are already being extended in follow-on SNS and partner-led initiatives.

Looking ahead, several important directions for future work emerge from BeGREEN:

- **Validation in larger-scale and more realistic environments**

Future work should extend the current validation activities towards digital twins and large-scale test environments capable of emulating realistic 5G and pre-6G deployments, including real-time traffic variability, mobility, and multi-domain interactions.

- **Conflict management and multi-loop orchestration**

As more AI/ML-based optimisation loops are introduced into the RAN, future work should address conflict detection and mitigation, model retraining strategies, and parameter tuning across multiple interacting control functions.

- **Broader integration of compute-aware and energy-aware AI mechanisms**

Future developments should explicitly account for the energy cost of AI itself, including training, inference, model orchestration, and hardware acceleration, so that intelligent control remains net-positive from an energy perspective.

- **Further exploration of advanced compute optimisation techniques**

This includes the use of C-states, deeper workload-aware resource scheduling, more advanced reinforcement learning approaches, and additional hardware–software co-design strategies for DU, CU, RIC, and Edge execution environments.

- **Progressive integration into real-world deployments**

Future work should focus on integrating the developed solutions into operational environments and broader use cases, enabling more realistic assessment of performance, robustness, scalability, and standardisation relevance.

In conclusion, BeGREEN has shown that energy efficiency can and should become a core design principle of future intelligent mobile networks. By combining architectural innovation, AI-driven optimisation, hardware-aware implementation, and practical validation, the project lays a strong foundation for the next generation of sustainable, high-performance, and intelligent B5G and 6G systems.

6 Bibliography

- [1] Sustainability in SNS JU Projects: Targets, Methodologies, Trade-offs and Implementation Considerations Towards 6G Systems, https://smart-networks.europa.eu/wp-content/uploads/2025/05/sns_ju_sustainabilitytf_whp_june2025_v1.0-1.pdf
- [2] IMT towards 2030 and beyond (IMT-2030) <https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2030/Pages/default.aspx>
- [3] BeGREEN D4.1, "State-of-the-Art Review and Initial Definition of BeGREEN O-RAN Intelligence Plane", December 2023. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [4] BeGREEN D4.2, "Initial evaluation of BeGREEN O-RAN Intelligence Plane, and AI/ML algorithms for NFV user-plane and edge service control energy efficiency optimization", July 2024. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [5] BeGREEN D4.3, "Final Architecture and Evaluation of BeGREEN O-RAN Intelligence Plane, and AI/ML Algorithms for NFV User-Plane and Edge Service Control Energy Efficiency Optimisation, September 2025. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [6] ITU-R, "Future technology trends of terrestrial International Mobile Telecommunications systems towards 2030 and beyond", Report ITU-R M.2516-0, November 2022.
- [7] BeGREEN D2.2, "Evolved Architecture and Power Enhancement Mechanisms", July 2024. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [8] BeGREEN D2.3, "Energy Efficient RAN Architecture and Strategies", March 2025. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [9] BeGREEN D3.1, "State-of-the-Art on PHY Mechanisms Energy Consumption and Specification of Efficiency Enhancement Solutions", January 2024. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [10] BeGREEN D3.2, "Initial Developments and Evaluation of the Proposed Enhancements and Optimisation Strategies", August 2024. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [11] BeGREEN D5.3, "BeGREEN Use Case Demonstration and KPI Assessment", September 2025. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [12] BeGREEN D3.3, "Final Evaluation and Benchmarking of the Implemented PHY Layer Solutions", March 2025. [Online] Available: <https://www.sns-begreen.com/deliverables>
- [13] BeGREEN White Paper on AI-Driven Energy Efficiency for O-RAN. SNS BeGREEN. May 2025. Available at: <https://www.sns-begreen.com/news/begreen-white-paper-on-ai-driven-energy-efficiency-for-o-ran>
- [14] M. Catalan-Cid, D. Reiss, G. Castellanos and J. Armstrong, "Demo: BeGREEN Intelligence Plane for AI-driven Energy Efficient O-RAN management," IEEE INFOCOM 2025 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), London, United Kingdom, 2025, pp. 1-2, doi: 10.1109/INFOCOMWKSHPS65812.2025.11152833.
- [15] O-RAN Alliance Working Group 1 (Use Cases and Overall Architecture), "O-RAN Network Energy Saving Use Cases Technical Report 2.0", O-RAN.WG1.Network-Energy-Savings-Technical-Report-R003-v02.00, June 2023.