Reducing energy use with 5G-Advanced

The essential guide to RAN energy savings in 3GPP Release 18

White paper

Daniela Laselva and Mads Lauridsen

Contents

Executive summary

Communications service providers strive to provide coverage, capacity and superior service quality while facing the challenges of ever-increasing traffic volumes, carbon neutrality targets, and higher costs of energy. Reducing the energy use of mobile networks to address these challenges is an urgent imperative. The efforts from the industry primarily revolve around the radio access network (RAN), which has the greatest impact as it consumes over 80% of mobile network energy. Nokia is making a continuous effort to decrease RAN energy consumption and improve the energy efficiency using the many available levers and measures: from network modernization and renewable energy sourcing solutions to advanced energy-saving features across the entire product portfolio.

As the radio technology strongly influences the ability to minimize energy use, Nokia is also actively engaged in the 3GPP standardization of new enablers in 5G-Advanced. These enablers are designed to facilitate dynamic energy-saving techniques for 5G base stations (gNBs). The objective is to reduce gNB energy use by operating the radios more efficiently than today without compromising service quality. The key techniques of 5G-Advanced allow the gNB to adapt active antennas, transmission power and time resources to changes in traffic load and end users' QoS needs. Specifically, these techniques allow adaptations to be more dynamic and at a more granular level than in today's networks. Further, they can use enhanced feedback from the device to determine the optimal adaptation to be applied by the network. Base station hardware can then be deactivated when the load decreases and quickly reactivated when the load increases again, in turn reducing energy use especially during the often occurring low-to-medium load scenarios.

Nokia has been a key contributor to define and assess the network energy-saving (NES) techniques of 5G-Advanced for their potential energy-saving gains and user throughput performance impact. As an essential tool for the evaluation, within the 3GPP, we have defined a common evaluation methodology and, for the first time, a gNB power consumption model.

In this white paper, we examine the 5G RAN energy-saving techniques introduced in 3GPP Release 18, describe how these can strengthen the broad energy-saving toolbox offered by Nokia, and provide recommendations on their use. Dynamic adaptations of transmission power and antennas improve the energy efficiency of the base station transmissions the most and obtain 15–30% energy savings under low-to-medium cell load levels. When looking at a daily average load, these techniques when jointly used also provide the best tradeoffs between energy savings and throughput impact. The summary of the network energy-saving gain and throughput impact obtained by the four key techniques of Release 18 standalone as well as the joint use of antenna and power adaptation (see results in the light blue) is provided in Figure 1 for different cell load levels. In general, these techniques must be applied with a careful design to mitigate the throughput loss they may cause.

This white paper concludes with our view on the anticipated additional energy-saving techniques that may be expected beyond Release 18.

Figure 1. Performance of the key Release 18 NES techniques standalone, and joint antenna and power adaptation (light blue background) at different load levels in terms of NES gain and user perceived throughput loss[1]

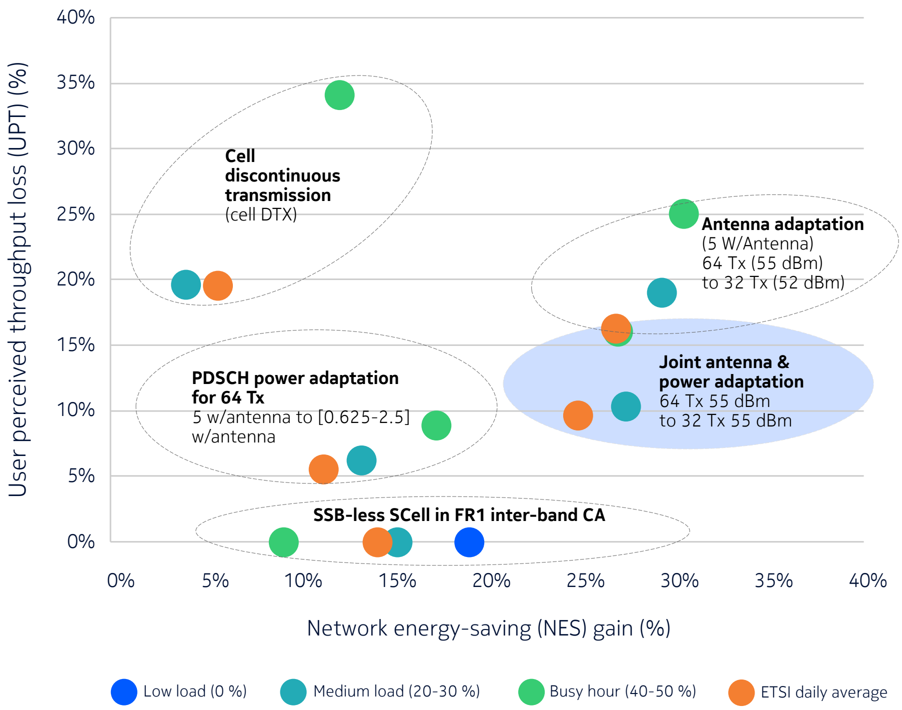

1 The results are based on the 3GPP-defined base station power consumption model and performance evaluation methodology [14].

Introduction

Energy consumption is one of the key metrics closely monitored by companies of every industry. It impacts costs and, in turn, profitability and is also critical in meeting commitments to reduce greenhouse gas emissions. Market expectations for most publicly traded companies are to meet ESG (Environmental, Social, Governance) criteria. Communications service providers (CSPs) and enterprises using wireless networks are no exception. In the journey towards net-zero emissions, many of them have set targets to be 'carbon neutral' or 'climate neutral' in the next decades for both Scope 1 and 2 (covering respectively direct and indirect CO_2 emissions from purchased electricity, heating, and cooling) [1]. At the same time, energy prices have been following an increasing trend [2], which is putting pressure on CSPs' profit margins, energy being one of the main cost components of the CSPs' OPEX (operational expenses) representing 20–25% share [3].

Since about 95% of mobile radio network product life cycle emissions occur while in use, minimizing energy consumption during product operation is the key to achieving reduction of both emissions and OPEX. In current mobile networks, most of the energy is consumed by the radio access network (RAN). According to the latest GSMA report, 87% of the energy of the operators surveyed is consumed by their RANs. The core network and data centers consume 12% and other operations account for the remaining 1%. Thus, minimizing energy use in the RAN is the key to improving profitability and achieving climate targets.

Minimizing energy use is critical also for 5G base stations. Despite being able to improve energy efficiency (bits/Joule) by up to 20 times as compared to 4G [4], 5G base stations contribute to electricity bills and actions towards mitigating their energy consumption are required. Particular attention should be paid to making the operations of power-hungry massive MIMO (mMIMO) radios more energy efficient during all load scenarios.

How 5G-Advanced enhancements reduce RAN energy use

On RAN energy use

Currently, base stations have the ability to deactivate unnecessary radio resources for energy savings while maintaining the requested network performance. The base station can enter a sleep state or mute resources when not needed by deactivating certain hardware components in the radio unit(s) and/or baseband unit (see Figure 2).

Figure 2. Simplified schematic of a traditional base station

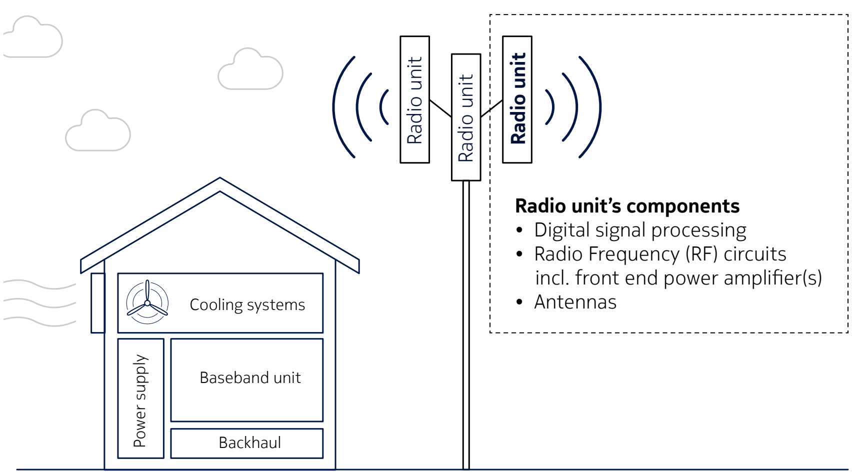

Sleep states range from micro discontinuities in the power amplifier at the symbol level — sometimes referred to as micro discontinuous transmission (DTX) — to longer deactivations at the cell level when a capacity cell can stop provisioning service during low-load periods.

Clustering the transmissions of different user equipments (UEs) in time can further enhance micro DTX and baseband resources can also be scaled. Further, mMIMO muting is a well-known technique for antenna resource muting (transparent to the UE) to enable energy-efficient transmissions. These techniques, shown in Figure 3, are part of the strong Nokia portfolio of energy-saving RAN software features.

Figure 3. Illustration of current key energy-saving RAN software features

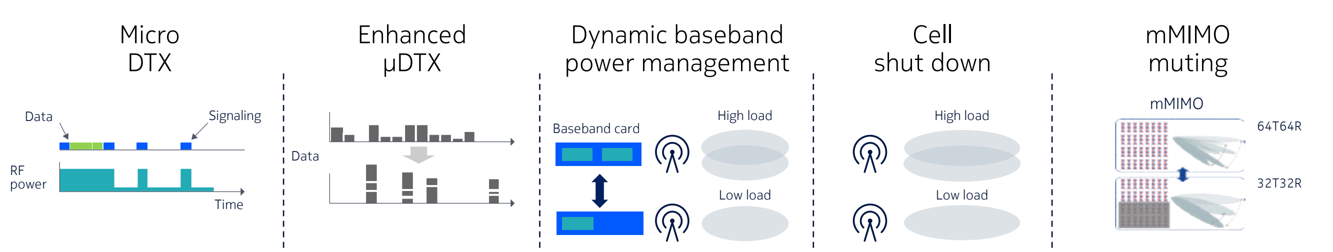

In recent years, several innovations have further contributed to the reduction of base station's energy consumption, particularly, by reducing the energy consumption for cooling. Moving the radio unit's components closer to the antenna is an example of innovation that removes the RF coaxial cables connected to the antenna and the associated RF signal losses. RAN site components that can be used outdoors are another example of innovation that decrease the need for air-conditioning.

Radio technologies, it is important to note, strongly influence the ability of the base station to use sleep states and resource muting. For example, 5G New Radio (NR) considerably increases cell sleep periods compared to LTE. With NR's lean carrier design, the regular Synchronization Signal Block (SSB) transmission is made (typically) every 20 ms [13]. This allows frequent use of micro sleep between SSB transmissions under low load. In contrast, LTE's always-on signalling (i.e., cell specific reference signals) are sent every 1 ms, continuously interrupting cell sleep. NR also increases energy efficiency (bits/Joule) compared to LTE thanks to mMIMO and beamforming technologies that bring a large boost in spectral efficiency.

The absolute energy consumption of the RAN site, nonetheless, may increase when introducing 5G [5]. A gNB employing mMIMO operations with 64 transceivers (TRXs) and 100 MHz system bandwidth could provide 10 to 20 times higher throughput than an LTE base station with 8 TRXs and 20 MHz, but it also consumes more power. The power consumption significantly depends on the hardware of the Radio Frequency (RF) radio module [6], particularly the power amplifier(s) and RF chipset, and on the capability to switch off RF hardware when idle or lowly loaded. A severe pain point with the majority of the first 5G products was that when high throughput was not needed, they lacked the capability to switch off mMIMO components, which resulted in unnecessarily high-power usage. Recently, promising advancements in the RF chipsets enable RF sub-components like mMIMO antenna sub-arrays, TX chains, front end components and circuitry, and baseband cards, to be turned off using extremely low current draw, such as in deep sleep, allowing to significantly optimize the power usage under low load.

RAN energy saving techniques in 3GPP Release 18

In Release 18 of 5G NR, 3GPP has been investigating techniques to reduce the energy consumed by 5G mMIMO base stations [14]. These techniques can turn off hardware components that are not needed in a timely and flexible way. The techniques use dynamic adaptation by the radio and corresponding hardware resources to the actual traffic load. Their goal is to manage the tradeoff between end-user performance (throughput and latency) and network energy consumption. Importantly, as per Release 18, the gNB can leverage the knowledge of the radio channel quality for different adaptation levels based on enhanced feedback from the device. This differs from similar vendor-specific implementation solutions that are transparent to the device.

In the following, we discuss the complete set of techniques and related new 3GPP standard enablers that are being defined in 3GPP Release 18 [7]. Notably, their selection was based on extensive investigations during the study item phase [14] that particularly targeted low-to-medium load scenarios. Although these scenarios consume lower power than high load scenarios (see Figure 9), they have the largest energy-saving potential. The reasons are twofold. Firstly, energy-saving gain can be achieved in low-to-medium load scenarios without compromising traditional performance. Secondly, the low-to-medium scenarios occur the most frequently in today's networks. In fact, the average network radio utilization is only ~30% during busy hours. Only a fraction of the cells is needed to provide coverage, so most cells are lightly loaded most of the time. The remaining cells are dimensioned for peak hour traffic demand, which represents a small fraction of the 24-hour period, and consequently these cells are also low-loaded most of the time (for example, during early morning). Figure 4 shows how unequal the traffic distribution is over the mobile network. 80% of the traffic is carried by 30% of the base stations, whereas the remaining 70% of base stations only carry 20% of the traffic; therefore, their utilization is low.

Figure 4. Typical traffic distributions in the mobile network

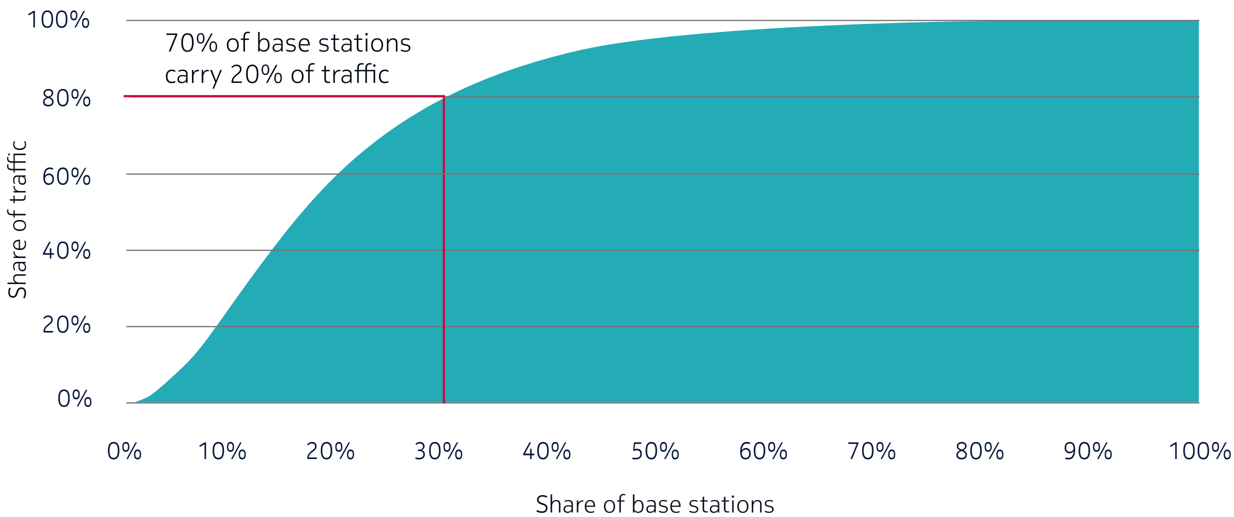

Energy-savings can be obtained at the base station (BS) by leveraging two main gain mechanisms:

- Enable more cell sleep opportunities: this mechanism utilizes low-power sleep states when no Tx or Rx operations must be made. It can be achieved (a) by removing downlink transmissions and/or uplink receptions of common channels (e.g., SSB, system information, paging transmissions) especially at low load, and (b) by clustering user-plane data in fewer time instances (slots) at low-to-medium load so as to reduce the number of slots with transmissions or receptions.
- Enable more energy-efficient transmissions or receptions: this mechanism adapts radio resources by switching off unneeded hardware components. The adaptation can be in time (occupied slots), frequency (occupied bandwidth), transmission power, and/or space domain (active TX and RX antenna ports and elements) and be made according to the actual needs (e.g., traffic amount, QoS demands, and device locations in the cell). It can be employed particularly at low-to-medium load when the peak capacity may not be entirely needed.

For each gain mechanism, several energy-saving techniques are being defined in 3GPP Release 18 according to the Work Item 'Network Energy Savings for NR' [11]. The techniques in scope are listed in Table 1 categorized according to the corresponding gain mechanism. Further, the table also provides an overview of the key 3GPP features for network energy savings available prior to Release 18.

Table 1. List of 3GPP techniques for Network Energy Savings in 5G NR

Gain mechanism	Technique	3GPP Release
Enable more cell sleep opportunities	SSB/SIB1 transmission periodicity of (typically) 20 ms in standalone 5G, and up to 160 ms in non-standalone 5G	Release 15
	SSB-less SCell operations in intra-band Carrier Aggregation	Release 15
	User Equipment DRX	Release 15
	SSB -less SCell in inter-band Carrier Aggregation for FR1 and co-located cells	Release 18
	Cell DTX/DRX	Release 18
	Conditional handover (CHO) procedure enhancements for Network Energy Savings	Release 18
	Restricting paging in a limited area of the cell	Release 18
	Xn inter-node signalling for requesting SSB-beam reactivation	Release 18
Enable more energy-efficient transmissions or receptions	Bandwidth part (BWP) adaptations	Release 15
	Fast SCell activation and deactivation	Release 15-17
	Small data transmissions	Release 17-18
	Antenna adaptation	Release 18
	PDSCH power adaptation	Release 18
	Preventing legacy UEs camping on cells adopting Release 18 NES techniques	Release 18

Techniques to enable more cell sleep opportunities

SSB-less SCell in inter-band carrier aggregation for FR1 and co-located cells

In the carrier aggregation (CA) scenario, the user equipment (UE) can be served by a primary cell (PCell) and one or more secondary cells (SCell). In the legacy operations, all SCells must transmit SSBs and system information (SIB1) regularly. On the contrary, this technique allows the SCells co-located with the PCell to omit the regular SSB (and SIB1) transmissions altogether. Compared to Release 15 where the intra-band SSB-less scenario is supported, in Release 18 the support is being extended to SCells operating in different bands than the PCell. As shown in Figure 5, the UE should be able to acquire coarse physical synchronization in time and frequency with the SSB-less SCell, for example, based on the SSBs provided by the PCell (as shown in Figure 5). Optionally, the UE can acquire synchronization from fine time tracking based on UE-specific reference signals provided by the SCell. Under low load, the energy-saving potential may be like switching off the SCell. However, if the SCell operates in SSB-less mode, it can exit its sleep state more dynamically to temporarily boost the network capacity when needed, resulting in increased throughput.

Figure 5. Illustration of SSB-less SCell operations

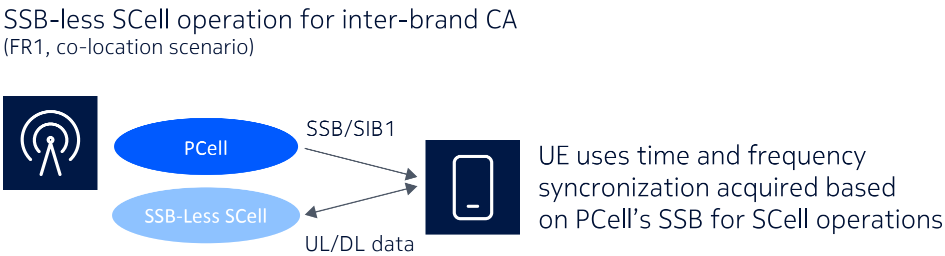

Cell discontinuous transmission and reception

Currently, a cell is allowed to sleep using micro-DTX when no transmission must be made. Further, Nokia's packet scheduler is enhanced for energy-saving attempts by enabling the clustering of transmissions of different UEs in the same scheduling interval to reduce the total number of transmissions (see enhanced micro DTX in Figure 3). The cell is still required, however, to wake up to perform certain transmissions and receptions. For example, the cell should transmit the Channel State Information Reference Signal (CSI-RS) reference signals when the UEs expect them and may need to wake up for that purpose. With cell DTX/DRX, the cell is instead allowed to omit, partially or completely, certain transmissions and receptions during non-active times (i.e., OFF periods for the cell).

The affected transmissions and receptions (that can be omitted) are related both to data traffic (e.g., semi-persistent scheduling) and reference signals. The UEs are configured with the cell DTX/DRX pattern, and may be provided with a dynamic (i.e., L1/L2) indication of (de)activation of the pattern to enhance network flexibility. In turn, the UEs can also avoid certain transmissions/receptions when the network is non-active and sleep instead, resulting in increased UE energy savings. To avoid impact to legacy UEs, cell DTX/DRX operations do not impact SSB transmissions. An illustration of a periodic cell DTX pattern is provided in Figure 6.

Figure 6. Illustration of a periodic cell DTX pattern

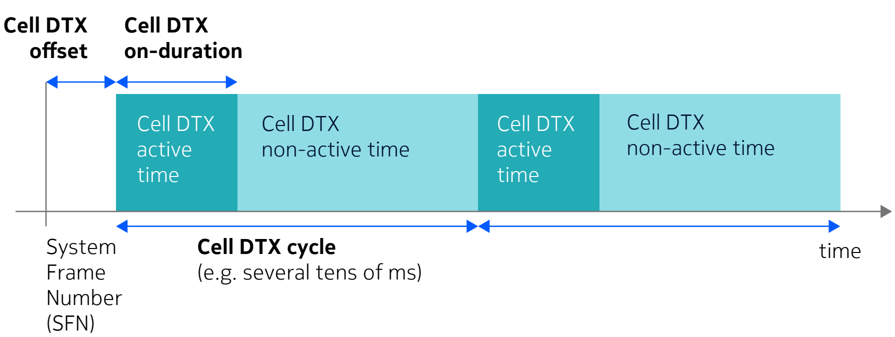

This technique may increase the network flexibility to enter and remain in a sleep state to save energy under low load, whenever no data traffic is present and expected. The challenge, however, is to make accurate AI/ML predictions of traffic variations to avoid latency impact. For this use case, new RAN data collection would be beneficial [17].

Conditional handover procedure enhancements for network energy savings

Currently, when an NES mode is being activated (e.g., the cell is being shut down), the cell needs to hand over to another cell each active UE individually. This causes unnecessary delay and may increase the likelihood of radio failures, especially if the NES modes are switched frequently. To avoid this, the active UEs can leverage the Release 16 conditional handover (CHO) framework to offload to a neighbor cell. With Release 16 CHO [8], the handover is executed by the UE whenever the configured conditions are met, without requiring a measurement report or handover command by the network. For NES, the Release 16 CHO procedure is being extended to account for the NES mode of the source and/or target cell. For example, the UE may initiate a conditional handover whenever the source cell enables an NES mode and may avoid performing the conditional handover to a target cell in NES mode.

Restricting paging in a limited area of the cell

Paging transmissions towards the UE are typically done initially by the last serving cell in a beam-swept manner — i.e., via all the SSB beams — because the network lacks awareness of the beam in which the UE performs the paging monitoring, which is up to the UE implementation. To limit paging transmissions to fewer spatial directions, the UE may be paged via a subset of the SSB beams within the cell. This may be beneficial for stationary UEs (e.g., smartphones during nighttime), which may be paged successfully via their last serving beam. However, the energy-saving potential may be limited as the cell has to wake up in any case to perform the paging transmission(s), albeit in a reduced number of beams.

This technique requires the network to know the last serving beam(s) of the UE, for which new in-network signalling is being defined in Release 18. Importantly, it requires the network to predict whether a UE is semi-stationary, thus eligible for this paging restriction (using, for example, AI/ML means introduced in Release 18 by [16]). Paging escalation will be triggered, however, if the UE is unreachable, for example, if the UE has moved away from the last serving beam. In turn, this would increase paging latency and overhead.

Xn inter-node signalling for requesting SSB-beam reactivation

Currently, during low load periods, cell shutdown is used as an effective means to save power, however, the deactivated cell(s) cannot provide any service at all. To increase flexibility, it can be beneficial to use deactivation at a finer granularity than cell level — i.e., on a per SSB beam level. For example, if low load occurs in a confined geographical area of the cell, the corresponding SSB beam(s) that illuminate(s) this area can be deactivated. Currently, SSB beam deactivation is supported in 5G NR via semi-static reconfiguration of the SSB beam grid, primarily for coverage purposes. To also leverage this functionality for energy savings, new (Xn) signalling between neighboring nodes is being introduced in Release 18. It will enable a node to request a neighbor node to reactivate an SSB beam when load increases.

Techniques to enable more energy-efficient data transfer

Antenna adaptation

The gNB employs mMIMO technology and is equipped typically with a large number of Tx antennas, and hence power amplifiers (PAs), for capacity boosting. One way to reduce the power consumption of these PAs is through antenna adaptation, i.e., by a dynamic and partial muting of certain antenna elements and the corresponding PAs. An example of this antenna adaptation is shown in Figure 7, where the gNB dynamically switches from 64 Tx to 32 Tx or 16 Tx based on the actual needs. This is also referred to as mMIMO muting. Depending on implementation, either all or only a subset of antenna elements associated with a logical antenna port can be disabled or enabled.

This is different from its implementation-based feature counterpart, which is transparent to the UE (see mMIMO muting in). The Release 18 antenna adaptation uses enhanced feedback from the UE to assist the network adaptation. Specifically, the enhanced channel state information (CSI) feedback will indicate the performance of different candidate antenna patterns for more optimal muting decisions by the network.

Figure 7. Illustration of antenna domain adaptations

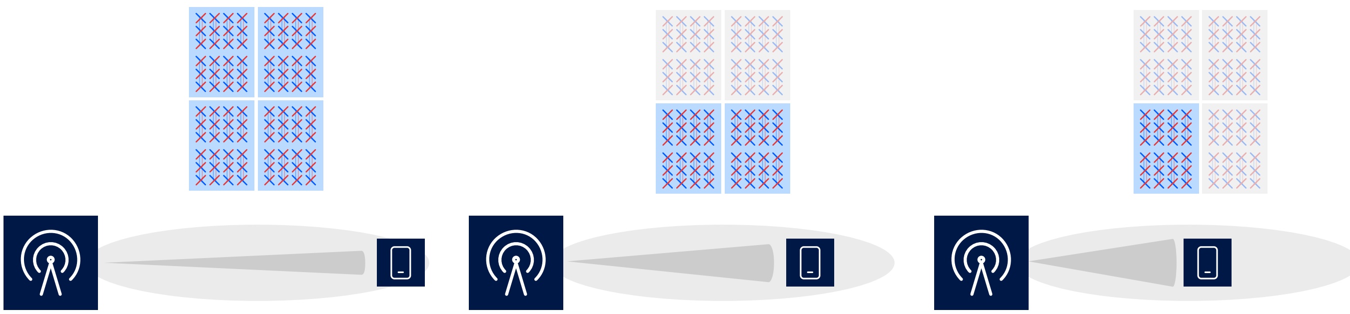

Blue and red colors indicate dual-polarized antennas that are active, grey color indicates inactive antennas.

PDSCH power domain adaptation

This technique dynamically adjusts the Tx power of the physical downlink shared channel (PDSCH). For example, higher power levels can be used for transmissions to cell edge UEs while lower levels are used for cell center UEs, as shown in Figure 8. In legacy NR operations, adaptation is done semi-statically and transparently to the UE. The UE derives the channel state information (CSI) for aiding the PDSCH transmission, accounting for a semi-statically configured power level (i.e., the power offset between the PDSCH and CSI-RS). In Release 18, a more dynamic power adaptation is envisioned based on enhanced CSI feedback from the UE. This enables the network to select the optimal power offset based on the radio performance observed by the device.

Both the reduction of active antennas and transmit power may be associated with a reduction in throughput, which in turn may lead to a longer time to transmit a payload. From an energy-saving perspective, fortunately, the additional activity does not outweigh the significant power reduction obtained by these techniques, as will be shown in the section on performance evaluation below.

Figure 8. Illustration of power domain adaptations

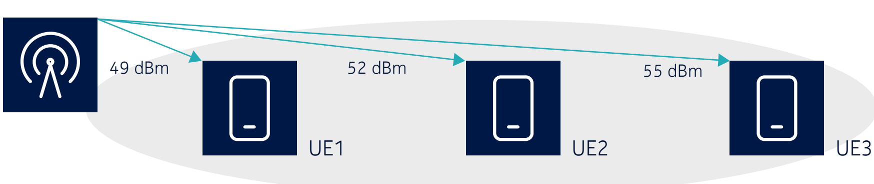

Preventing legacy UEs camping on Release 18 cells

Although legacy 5G UEs (Releases 15–17) can be served by a cell that is using Release 18 NES techniques (such as cell DTX/DRX and antenna adaptation), in certain cases it may be beneficial to prevent the presence of legacy UEs in the cell to enable higher energy-saving gains. For example, this may be used in certain capacity cells depending on the penetration of Release 18 devices. Therefore, new mechanisms to control (prevent) the camping of legacy 5G UEs on cells using NES are being defined in Release 18.

3GPP-defined 5G base station power consumption model

The ability to evaluate accurately the energy-saving potential of different techniques is essential for the optimization of design and use of energy-saving solutions. 3GPP has therefore defined in [14], for the first time, a 5G base station power consumption model to evaluate and compare in a fair and realistic manner the energy-saving gain of the techniques studied in 3GPP Release 18. The model is based on the consolidated input provided by the major telecommunication companies and network equipment vendors, including Nokia. The configuration of radio units of the BS largely influences its power consumption, thus the model is defined for three typical radio configurations in terms of frequency range, duplex type, number of transmit antennas, and transmit power, as shown in Table 2. As per [14], the power consumed by the base station is given for the defined power states as follows:

- The active states: these are defined separately for downlink transmissions (active downlink state) and uplink receptions (active uplink state). Furthermore, a scaling of the power consumed in the active downlink and uplink states is defined as a function of the occupied symbols and bandwidth within a slot, number of active component carriers, number of active RF chains, and transmit power. This allows us to study energy-saving adaptations in all four radio domains: time, frequency, antenna and power.
- The sleep states: these include deep, light and micro sleep states, during which no transmission nor reception can take place, thus the base station can sleep and save energy. Whereas micro sleep can be achieved instantaneously, for the advanced sleep states, i.e., light and deep sleep, we assume that more hardware sub-components will be turned off to reach a lower power consumption. As a consequence, a longer transition time is needed to enter and exit the advanced sleep states as shown in Table 3. We note that the transition of the BS from any active state to a given sleep state is characterized by a total transition time, during which the BS enters and leaves the sleep state and during which it consumes a certain energy (see Table 3).

Table 2. 3GPP-defined Reference Radio Configurations Set 1 – 3

	Set 1, FR 1	Set 2, FR 1	Set 3, FR2
BS type	Macro	Macro	Macro
Duplex	TDD	TDD	TDD
System bandwidth	100 MHz	20 MHz	100 MHZ
Subcarrier spacing	30 kHz	15 kHz	120 kHz
Number of transmission points	1	1	1
Number of DL TX	64	32	2
Total DL power (Tx Power)	55 dBm	49 dBm	33 dBm
Total number of UL Rx	64	32	2

The power consumed in each state (on a per slot basis) is given relative to the power consumed in the most efficient state (i.e., deep sleep). In Table 3, the power values and transition time for BS Category 2 are shown. Although based only on a model describing a complex reality, it may be considered realistic considering the state-of-art BS hardware components and anticipated technological developments [14]. In addition, 3GPP also defines a BS Category 1 model, which assumes more advanced hardware components, which are not expected to be feasible in the coming years because both extremely fast transitions and very low-power sleep states would be required.

Table 3. 3GPP-defined 5G relative power consumption values of BS active and sleep states and total transition time per sleep state for BS Category 2

Power state	Relative power			Transition time
	Set 1	Set 2	Set 3	
Deep sleep	1	1	1	10 s
Light sleep	2.1	2.1	2.1	640 ms
Micro sleep	5.5	5	3	0 ms
Active UL	6.5	5.8	4.2	
Active DL	32	26	17.6	

To exemplify the model, Figure 9 shows 3GPP-defined 5G BS power consumption for the reference Radio Configuration Set 1 (Macro, mMIMO, 64 Tx, TDD, FR1, with 55 dBm total transmit power) as a function of the RF load per slot. The RF load of 100% means that all 14 OFDM symbols in a slot are occupied. The power reduction values in the figure are shown compared to the reference (64 Tx, 55 dBm and 100% RF load). Note that the largest power reduction, according to the 3GPP 5G BS power model, is attained by using the micro sleep and deep sleep states, which enable power reductions of 82% and 96% respectively. The power reduction obtained by reducing the RF load per slot is 41%, 58%, and 70-80% for RF loads of 50%, 30% and 2–10% respectively. Another effective means to reduce the BS power consumption is by reducing the number of Tx chains. Using 32 Tx rather than 64 Tx results in up to 52% power reduction for RF loads of 100% (see the light green curve in Figure 9).

Figure 9. Illustration of 3GPP-defined 5G BS power consumption vs. RF load for Radio Configuration Set 1 (Macro, mMIMO, 64 Tx, TDD, FR1) with scaling with antennas (64TX and 32TX) and transmit power (52 dBm and 55 dBm).

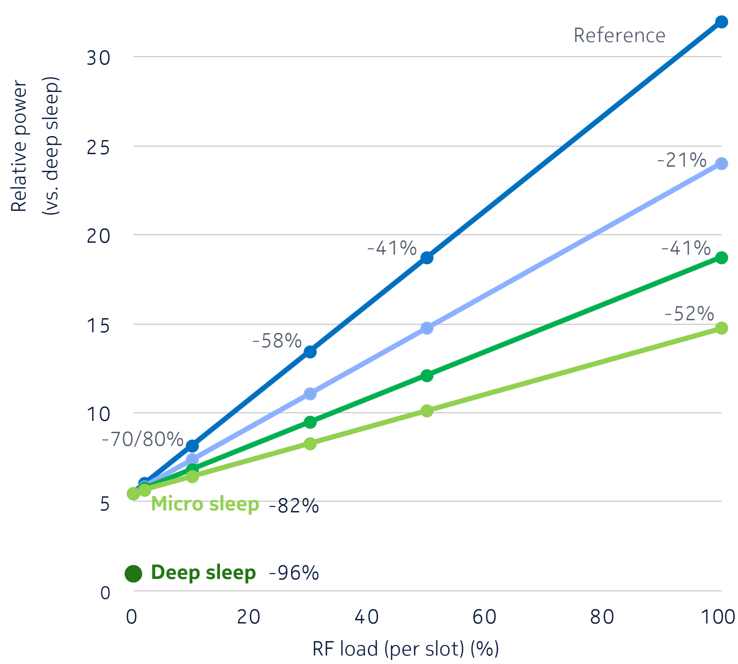

Performance evaluation

The 3GPP Release 18 techniques aim at achieving a good tradeoff between BS energy savings and UE performance impact, targeting the highest energy savings with no or limited impact to the user-perceived throughput. An extensive performance evaluation was performed by Nokia and other companies during the 3GPP study item phase [14]. In the following, we present the evaluation of antenna and PDSCH Tx power adaptations, which are among the most promising techniques in terms of the tradeoff between gain and impact [9].

System-level simulations are carried out according to the 3GPP-defined evaluation methodology in a typical dense urban macro cell deployment for Radio Configuration Set 1 (Macro, mMIMO, 64 Tx, TDD, FR1, 55 dBm total transmit, see Table 2) and using the BS Category 2 power model (see Table 3). For further details see Annex B in [14]. The potential NES gain and the corresponding cost of degraded user throughput scale with the cell load; thus, the evaluation is carried out at low-to-medium load in terms of physical resource block (PRB) usage.

As seen in Figure 10, the energy-saving gain attainable with the investigated antenna and power adaptations increases as a function of the cell load. The gain for all curves is given compared to the baseline BS configuration using 64 Tx and the maximum transmit power of 55 dBm. Performing Tx power adaptation when using 64 Tx (blue curves) can achieve up to a 39% energy-saving gain. Joint antenna and Tx power adaptations using 32 Tx (green curves) can achieve a higher energy-saving gain compared to the baseline (up to 56%). The energy savings increase with the load because, although the slot occupancy also increases with the load (due to the lower throughput), the lower power consumption per transmission (52% lower, see the light green curve in Figure 9) outweighs the extra power consumed by the additionally occupied slots.

It may also result, however, in a loss of user-perceived throughput due to reduced beam-forming gain when using fewer antenna elements, as shown in Figure 11. At the lowest investigated load, the loss is 12% for the configuration with 3 dB power reduction. At 40% load, the loss of the same configuration is only 2%. In most scenarios, the loss remains constrained or decreases when the load increases to very high levels (i.e., 80%), because the throughput tends to be more influenced by the resource sharing occurring at higher loads. The loss in general, however, may be mitigated by algorithm design, for example, by targeting only cell center UEs rather than cell edge UEs, using additional frequency allocation, and power boosting to compensate for the loss. Those aspects are not accounted for in the presented results.

Figure 10. Network energy-saving gain vs. PRB usage for power adaptations (blue curves) and joint antenna and power adaptations (green curves)

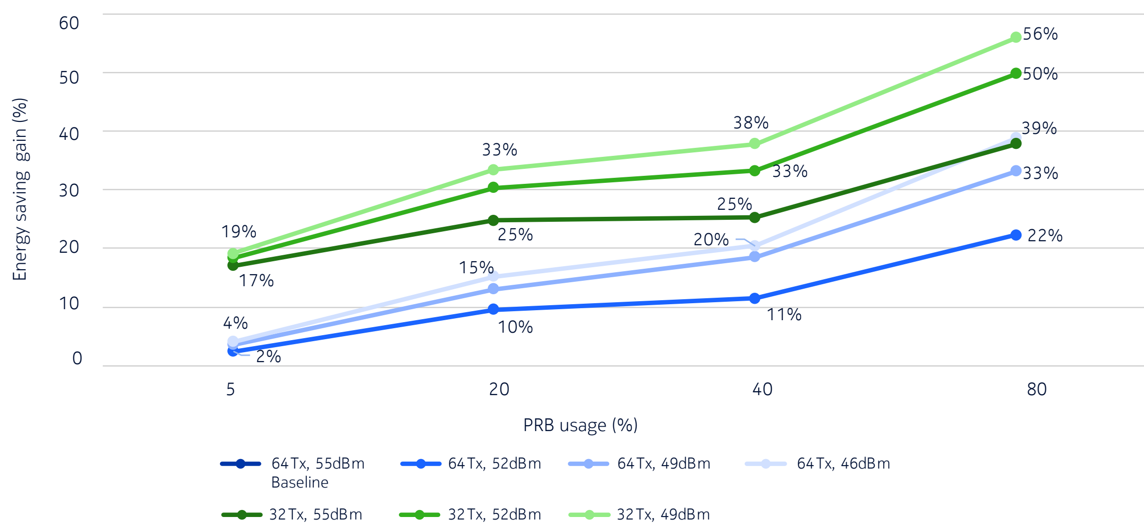

Figure 11. Packet throughput vs. PRB usage for power adaptations (blue bars) and joint antenna and power adaptations (green bars)

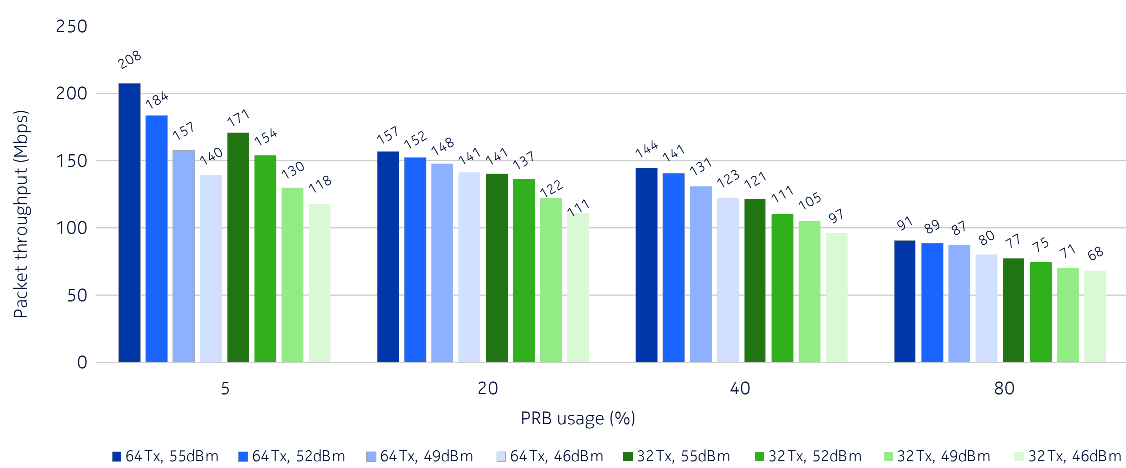

Recommendations

For reducing BS energy consumption, what really matters is the NES gain provided by a technique throughout the entire day. NES gain will vary according to the load variability during the 24-hour period, and it depends on the performance of the given technique at different load levels, which is highly variable, as observed in the previous section.

Figure 12 summarizes the NES gain and impact of four key standalone Release 18 techniques as well as the joint use of antenna and power adaptation (see results in the grey background). The results are shown for different cell loads including for the daily average cell load [19]. The load levels in Figure 12 are in terms of PDSCH resource usage for user-plane data. The results are cell-wide and do not consider potential enhancements related to, for example, cell edge and cell center users, etc. Notably, the results assume 100% penetration of Release 18 devices supporting the investigated techniques. In reality, actual penetration of supportive devices will likely be less than this, meaning lower performance levels than those recorded here, at least for some time.

At low load (i.e., signalling-only periods), omitting SSB transmissions on a (capacity) SCell is the only technique, which can provide gains because the cell DTX, power and antenna adaptation schemes only impact the user data transmissions, which do not occur at low load. The SSB-less functionality can generally provide ~35 % energy savings on the SCell but is only applicable in CA scenarios. Note that SIB1 is also omitted by the SCell (besides SSBs), and this is the largest contributor to the gain.

At medium load and busy hour, the antenna adaptation brings the largest NES gain (>30%), but at significant impact to the user perceived throughput. At these loads, the power adaptation may provide a more balanced tradeoff between NES gain and throughput loss, being in the order of 2:1. In the investigated 3GPP scenario, the cell DTX suffers from throughput loss because user scheduling may be delayed due to their clustering in the time domain.

When looking at the daily average load, the techniques that improve the energy efficiency of the transmissions (power and antenna adaptations) provide the best tradeoffs between NES gain and throughput loss especially when used jointly (see results in the grey background in Figure 12). The SSB-less SCell can be deployed without performance impact and brings significant gains, but it is only applicable in the CA cases. Cell DTX is generally not a favorable technique in the evaluated 3GPP scenarios, where the large packet size makes large frequency allocations necessary for single UEs thus prohibiting advantageous frequency domain multiplexing across different UEs during cell DTX active periods.

Besides joint use of antenna and power adaptation shown in Figure 12, it is worth noting that other techniques can be combined too. For example, one can apply both power and antenna adaptations together with cell DTX, while the SSB-less SCell can be applied simultaneously on the capacity carrier if available. However, the gains when combining different techniques are not additive.

Concluding, at non-zero load, we recommend a careful use of power and antenna adaptations jointly to achieve the desired tradeoff between energy savings and performance impact. In general, one should account for various factors that might strongly influence the most beneficial technique(s) to be enabled such as:

- The actual scenario including cell loading levels, UE locations in the cell, traffic characteristics (e.g., packet size)
- The acceptable tradeoff between NES gain and performance impact considering the QoS targets
- The deployed base station hardware and software

Notably, AI/ML algorithms could be employed to optimally operate the Release 18 energy-saving techniques leveraging their ability to predict, for example, UE mobility and cell load variations, which are essential to determine the optimal power state of the base station at a given time. These algorithms can leverage RAN data collection being defined in Release 18 [16] and its evolution, which is expected in Release 19 [17].

Figure 12. Performance of the key Release 18 NES techniques standalone, and joint antenna and power adaptation (light blue background) at different load levels in terms of NES gain and user perceived throughput loss.[2]

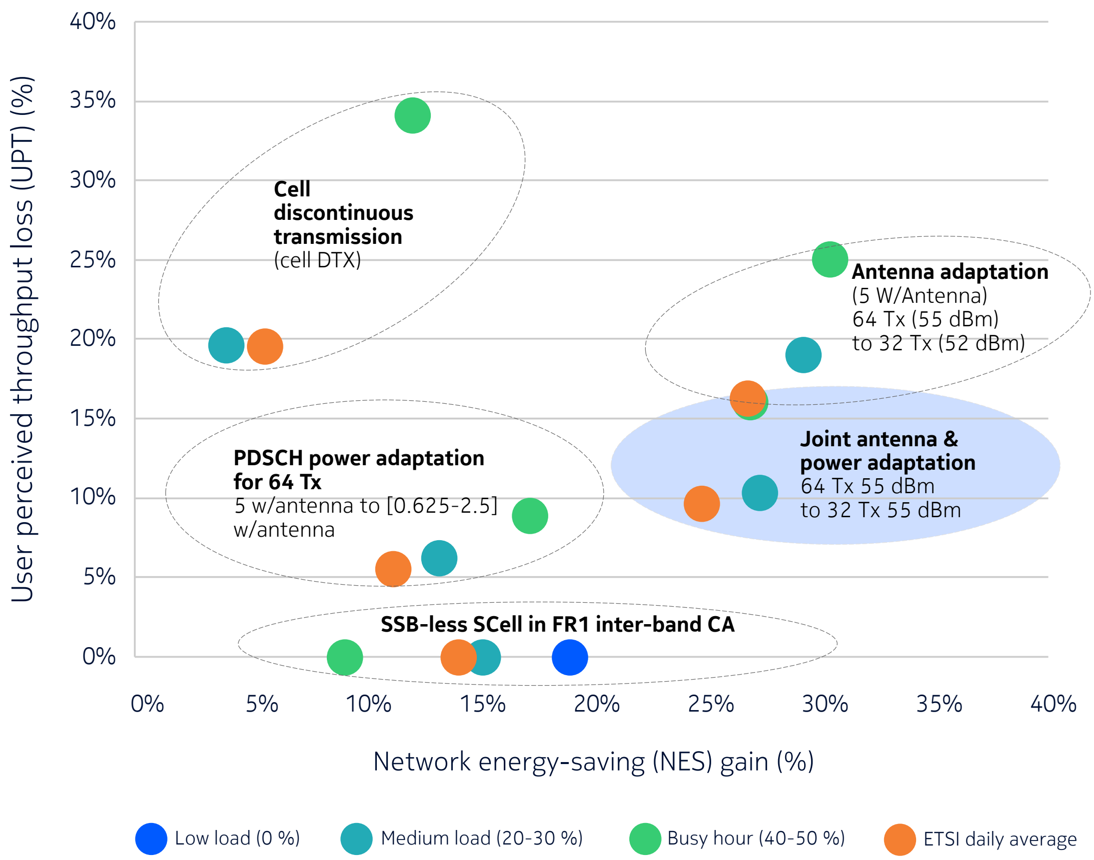

2 The gains achieved by each feature independently cannot be combined when considering the case that more than one feature is enabled at the time.

Outlook for RAN energy savings in 3GPP Release 19

The NES enhancements of 5G NR Release 18, described in the previous section, are only a subset of the many identified as promising in 3GPP TR38.864 [14]. However, it is beneficial to continue the normative work in 3GPP Release 19 aiming at defining relevant techniques to enable additional energy savings, especially if the standardization efforts are reasonable. Particularly, as the Release 18 techniques apply primarily to cells serving radio resource control (RRC) connected UEs, in Release 19 the focus could be on techniques that also provide energy savings when RRC-idle/inactive UEs are present, which may be the majority of UEs in low load scenarios.

The exact scope of the RAN energy-saving evolution in Release 19 will be defined only after completion of the Release 18 work, with the final definition of the overall Release 19 work and study items expected in December 2023. Nevertheless, the key technical areas that should be considered for further energy-saving enhancements in Release 19 are discussed in the following. (See further details in [15].)

Further SSB/SIB1 transmission adaptations for low load scenarios: any active 5G cell as per Release 15–18 must transmit SSB/SIB1 to accommodate legacy UEs (except for SSB-less SCells discussed in the previous section). This is irrespective of devices being present that benefit from the signalling. Thus, SSB/SIB1 transmissions can be further optimized by:

- Enabling on-demand SSB transmissions for RRC-connected UEs based on the request by the device
- Enabling on-demand SIB1 transmissions for RRC-idle/inactive UEs based on the request by the device
- Enabling SIB1-less cell operations for RRC-idle/inactive UEs in co-located scenarios where another carrier can provide the SIB1 for the UE

Further enhancements that may be considered for Release 19 and beyond are discussed below.

Paging enhancements for low load scenarios: according to 5G Release 15–18, the paging occasions of different UEs in a cell are by design distributed in time to increase paging capacity and distribute resource usage over time. But this negatively impacts the cell sleeping opportunities at low load because the cell must wake up for paging transmissions according to the paging occasions. Paging transmissions should instead be clustered in time to obtain energy savings.

Enhancements for cell shutdown procedures: the cell shutdown procedure is one of the most effective means used to save energy, but it is relatively slow, and limited by the signalling mechanisms available in 5G Release 15–18 to gracefully power down the cell. This may prevent it from being used frequently. In Release 18, conditional handover enhancements for saving energy, introduced in the previous section, are considered primarily for robustness, for example, to ensure traffic can be offloaded at cell shutdown in a failure-free fashion. Further enhancements would be beneficial to speed up the powering down of the cell shutdown procedure. In addition, enhancements related to gNB's components shutdown would be also beneficial in the gNB disaggregated architecture, where the gNB is split in Centralized Unit and Distributed Units. This would enable such split gNB deployments to obtain the same energy-saving potentials as in traditional gNB deployments.

Moreover, standards for collecting RAN data to enable AI/ML algorithms for energy savings is expected to continue after the closing of the related Release 18 work [16]. This may evolve in Release 19 towards enabling collection of RAN data to operate some of the Release 18 NES techniques. Examples that can be considered include the prediction of preferred cell DTX/DRX patterns and the prediction of the opportunities for beam-level switch-on/off [17].

Summary and outlook

Nokia engages in a continuous effort to help optimize network energy consumption. It helps CSPs and enterprises using wireless networks to achieve their sustainability targets and reduce their energy bills. Nokia, for example, has the ambition to achieve substantial power consumption reduction for new generations of base station hardware and strengthen the offering of energy-saving features (especially SON) and AI/ML-based energy-saving automations.

As part of the 3GPP Release 18 work item on network energy savings [11], Nokia is engaged in defining several promising RAN energy-saving techniques that enable a more dynamic and granular control of energy use of the gNB under low-to-medium loading scenarios. The aim is to obtain the highest potential energy savings while ensuring no or limited impact to user throughput. As expected, the performance of these techniques is highly load dependent.

Among the Release 18 techniques, transmit power and antenna adaptations provide the best tradeoffs between energy-saving gain and throughput impact in the 3GPP evaluation scenarios. They obtain a 15–30% energy-saving gain under low-to-medium load levels and slightly less gain under daily average loads, assuming 100% penetration of Release 18 devices. Furthermore, these techniques provide above 30% energy savings at higher load levels, however, an intelligent control of when to enable these techniques is necessary for throughput loss mitigation.

In contrast to their implementation-based counterparts, the Release 18 power and antenna adaptations leverage enhanced feedback from the UE to assist the network adaptation. In general, the technique(s) to be deployed in each network at a given time should be determined in a holistic manner accounting for the actual cell load level, UE locations in the cell, traffic characteristics, QoS targets, the penetration of the devices supporting the techniques, and deployed RAN hardware and software.

Furthermore, several other 3GPP activities on network energy savings are in full swing at present. Among others, 3GPP is currently introducing AI/ML-assisted energy-saving means for cell shutdown in Release 18 [16]. Using new RAN signalling (Xn/F1/Uu), the AI/ML algorithms can use RAN data to optimize the decisions related to cell shutdown and traffic offloading, for example, by predicting UE mobility, cell loads and base station power states. Further, the 3GPP is studying energy efficiency as service criteria considering, for example, energy consumption exposure to verticals and energy credit limits as enhanced means to monitor and control the actual energy use in the network [10].

Looking ahead, we see the benefit of introducing additional energy-saving features both in the 5G RAN and core network in 5G-Advanced Release 19. The development of energy-saving features in 5G-Advanced will also benefit 6G, paving the way to stronger energy efficiency foundations, which are industry priorities. In 6G, the notion of a leaner carrier and dynamic resource muting for increased energy efficiency should be supported from the get-go, by the entire device population. This would allow us to control the energy usage when considering the anticipated increase in TRX chains (e.g., 128 of 6G vs 64 of 5G), bandwidth size (e.g., 400 MHz of 6G vs. 100 MHz of 5G) and, consequently, transmit power [18].

Abbreviations

3GPP	Third-generation partnership project
BS	Base station
BWP	Bandwidth part
CA	Carrier aggregation
CHO	Conditional handover
CSI	Channel state information
CSI-RS	Channel state information reference signal
CSP	Communications service provider
DRX	Discontinuous reception
DTX	Discontinuous transmission
ESG	Environmental, social, governance
FR1	Frequency range 1
GSMA	Groupe Spéciale Mobile Association
gNB	Next generation node B
LTE	Long-term evolution (often interchangeable with 4G)
MIMO	Multiple input, multiple output
mMIMO	Massive MIMO
NES	Network energy-saving
NG-RAN	Next-generation radio access network (5G)
NR	New radio (5G)
QoS	Quality of service
OFDM	Orthogonal frequency-division multiplexing
OPEX	Operating expenses
PA	Power amplifier
PCell	Primary cell
PDSCH	Physical downlink shared channel (NR)
PRB	Physical resource block
RAN	Radio access network
RF	Radio frequency
RRC	Radio resource control
RX	Receiver
SCell	Secondary cell
SI	System information
SIB1	System information block type 1
SSB	Synchronization signal block
TRX	Transceiver
TX	Transmitter
UE	User equipment
Xn	Network interface between NG-RAN nodes

NOKIA

References

1. Telefonica Climate Action Plan, June 2022, and AT&T Climate Strategy & Transition, 2019.
2. International Energy Agency, “Electricity Market Report 2023”, Feb 2023.
3. GSMA, “5G energy efficiencies: Green is the new black,” Nov 2020.
4. Orange, “Sustainable networks: aiming for net zero carbon,” Jul 2022.
5. Wang, F., “China Mobile Reduces the Power Consumption of 5G Base Station,” Equal Ocean, Jul 2021. https://equalocean.com/news/2021070616439
6. SK Telekom and NTT DoCoMo, “Green Mobile Network: Energy Saving Efforts,” White Paper, Feb 2023. https://www.docomo.ne.jp/english/binary/pdf/corporate/technology/rd/docomo6g/GreenMobileNetworksWhitePaper_22February2023.pdf
7. 3GPP Work Item RP-223540, “Network energy savings for NR”, Dec 2022. https://portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workitemId=981037.
8. Nokia, “Rock solid mobility innovations from 5G to 5G-Advanced”, white paper, 2023. https://onestore.nokia.com/asset/212564
9. M. Oikonomakou, A. Khlass, D. Laselva, M. Lauridsen, M. Deghel, and G. Bhatti “A power consumption model and energy saving techniques for 5G-Advanced base stations,” IEEE ICC, Rome, Italy, 2023.
10. 3GPP Study Item S1-221232 “Study on Energy Efficiency as service criteria”, May 2022. http://www.3gpp.org/ftp/tsg_sa/WG1_Serv/TSGS1_98e_EM_May2022/Docs/S1-221232.zip
11. 3GPP Work Item RP-223540, “Network energy savings for NR”, Dec 2022. https://www.3gpp.org/ftp/Information/WI_Sheet/RP-223540.zip
12. Nokia, “How network adaptations for 5G devices will lead to superior battery life,” 2021. https://onestore.nokia.com/asset/210981
13. 3GPP TR 37.910 “Study on self-evaluation towards IMT-2020 submission”, 2022. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3190
14. 3GPP TR 38.864 “Study of Network Energy Savings for NR”, Dec 2022. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3987
15. Nokia, “Network Energy Saving Enhancements”, Sep 2023. https://www.3gpp.org/ftp/meetings_3gpp_sync/ran/Docs/RP-231891.zip
16. 3GPP Work Item RP-213602, “AI/ML for NG-RAN”, Jun 2022. https://www.3gpp.org/ftp/tsg_ran/TSG_RAN/TSGR_94e/Docs/RP-213602.zip
17. Nokia “AI/ML for NG-RAN”, Jun 2023. https://www.3gpp.org/ftp/tsg_ran/TSG_RAN/TSGR_AHs/2023_06_RAN_Rel19_WS/Docs/RWS-230023.zip
18. Nokia Bell Labs, “Energy efficiency in next-generation mobile networks”, Nov 2022. https://www.bell-labs.com/institute/white-papers/energy-efficiency-in-next-generation-mobile-networks/
19. ETSI 202 706-1, “Environmental Engineering (EE) - Metrics and measurement method for energy efficiency of wireless access network equipment - Part 1: Power consumption - static measurement method,” V1.7.1, Sep 2022. https://www.etsi.org/deliver/etsi_es/202700_202799/20270601/01.07.01_60/es_20270601v010701p.pdf

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2023 Nokia

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: CID213599 (November)