Ericsson White Paper
GFTL-24:001335
November 2024

ERICSSON

Energy Performance of 6G Radio Access Networks: A once in a decade opportunity

Content

Introduction

High and growing energy consumption receives more and more attention in the telecommunication industry. This also applies to discussions about 6G where delivering continued service growth combined with reduced network energy consumption is a key consideration.

The term "energy performance" in the title of this white paper denotes minimizing the network energy consumption for a given set of performance requirements, such as user throughput, capacity, and latency. Energy performance is a system-wide aspect that can be easily compromised by mistakes in standards, implementations, or network management. Every component of the system matters, and it only takes one bad solution in one part for the system to end up with poor overall energy performance. A good standard is essential for achieving excellent energy performance, but it is not enough on its own.

It takes time to see the global benefits of improved design practices and infrastructure technology due to the long lifespans of network equipment and telecom generations. This makes it even more important to double down and ensure that cellular technology continues to advance in the energy performance dimension.

There have already been important advancements in the energy performance area, particularly with the introduction of lean design with 5G NR, which has demonstrated significant energy savings compared to 4G LTE. The lean design reduces signaling and enables prolonged sleep time between the mandatory transmissions and receptions required for providing access coverage.

In the cellular industry, we are now preparing to take the next step to design the new 6G standard to enable even greater energy savings than 5G NR.

Some solutions, such as those related to UE idle-mode functions like system-information broadcast, random-access, and paging can only be changed when a new generation is introduced. At the same time, the energy cost of providing these fundamental idle-mode related functions is generally much higher than the additional energy used for serving user data traffic in a nationwide network. A new generation like 6G is therefore a once-in-a-decade opportunity to address the dominating energy cost of a cellular network, which is the cost of providing area coverage, and enabling idle-mode UEs to become active. This white paper presents key areas for standardization, that would enable increased energy-saving potential in 6G compared to 5G NR, for further discussion.

Networks today and in the future – traffic and deployment

When aiming to reduce network energy consumption, it is necessary to understand how cellular networks are used, and where most energy is consumed today to target the areas with the largest potential. The lion's share of energy consumed in cellular networks is associated with the radio access network (RAN) (76 percent according to [1]). RAN sites are distributed and typically support several generations of radio-access technologies with a substantial part of legacy equipment that can be 10 years or older.

Understanding traffic characteristics is critical when designing an energy-efficient network. Based on measurements taken in live networks, we observe that almost all data sessions are small, with only around one percent being larger than 20 MB. However, these larger sessions make up the majority of the data volume in the network.

The average traffic load in different cellular networks is depicted in Figure 1. The top sub-figure in Figure 1 depicts the average resource utilization in different LTE networks (measured in PRB utilization) while the bottom sub-figure depicts the average resource utilization in different NR networks (measured in RBSymb utilization). The averaging is done over 24 hours and includes all cells in each of the networks. The weighted average numbers (dashed red line) consider the number of cells in each of the studied networks. The resulting bars are sorted from high to low utilization per customer identity of the networks (not shown in the figure). We note that in a nationwide network, the average physical layer resource utilization in LTE is around 20 percent and in NR it is around 6 percent today. One reason for the significantly lower resource utilization in 5G NR compared to LTE is the widespread deployment on non-stand-alone (NSA) 5G where many packets are still transmitted over the 4G connection.

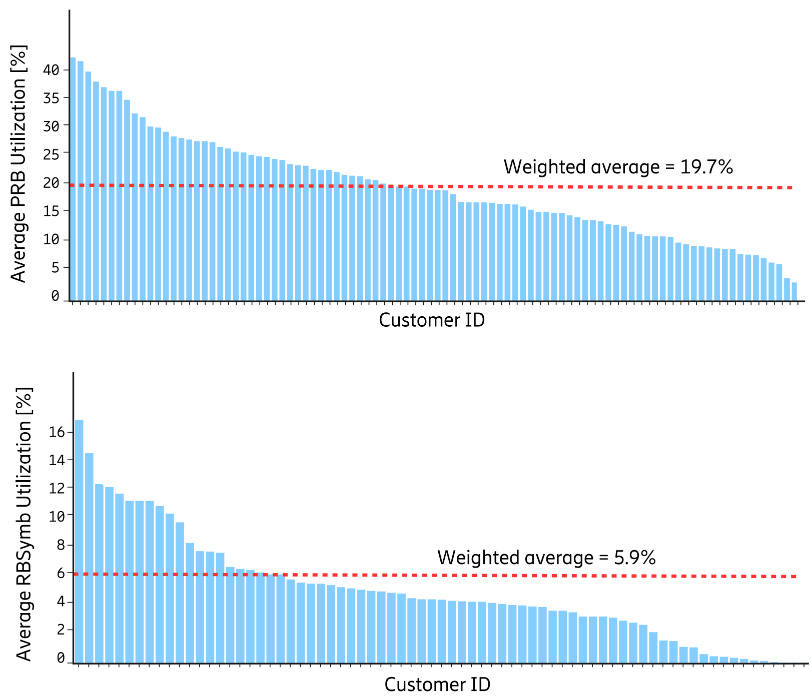

Figure 1: Twenty-four hour and network-wide average resource utilization in different LTE networks (top) and NR networks (bottom), respectively.

Hence, a typical LTE or NR cell is most likely to have zero traffic and serve zero users in any given millisecond. When a cell is not empty, the typical data transmissions are bursty and interleaved with periods of low or no activity. Even in a network that gets congested during peak hours and in hot-spots, the traffic is low most of the time in most locations. Networks need to be dimensioned for peak traffic demands, leading to low average utilization most of the time in most locations, creating substantial potential for network energy savings. Most cells are not hot-spots, and most hours are not peak hours. Even during peak hours, there are several short durations on a timescale of milliseconds with zero traffic that can be effectively utilized for saving energy.

Increased energy consumption during a data burst is unavoidable and can still result in high energy efficiency, in terms of bits/Joule. In the context of minimizing network energy consumption, we also seek a system that enables low energy consumption for the most common case when there is no or low data traffic. Moreover, we want the network power usage to be proportional to the traffic load, and when there is zero traffic, the power usage should be as close to zero Watt as possible.

For 6G, we need to ensure that we can benefit, in terms of reduced network energy consumption, from deployment architectures where RAN processing is more centralized (see Figure 2). In radio products, there is also a trend toward lower-layer split becoming the norm. This development can enable larger gains with hardware pooling, virtualization, adaptation, and coordination of RAN processing resources. The notion of a "site" might become less important in case processing related to an "area" is performed in one central location. In scenarios with many frequency bands and transmission points, there are additional opportunities for reducing network energy consumption. For example, not every carrier and node needs to support system-information broadcast and connected-mode functionality.

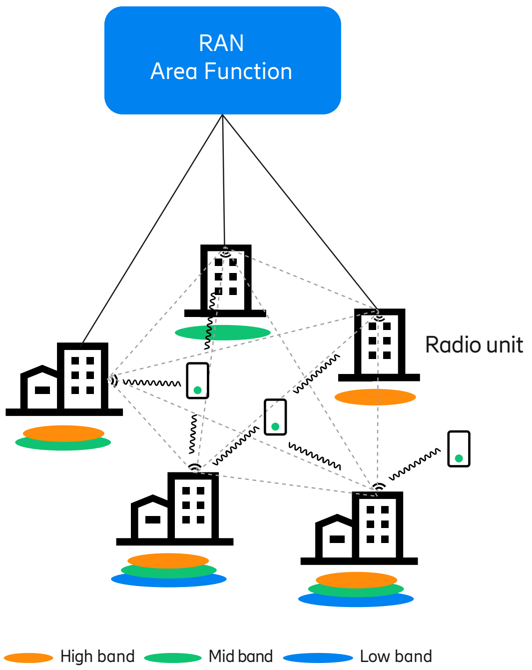

Figure 2: 6G needs to support multi-carrier and multi-point transmission and reception from day one in an energy efficient manner. The relevance of the site is diminishing when processing related to an area is increasingly centralized.

Lean design in 6G — Separation of concerns

The introduction of lean design in 5G NR, which focuses on minimizing transmissions not related to data transfer, has been a tremendous success enabling large network energy savings due to micro-sleep between transmissions. For 6G, we should continue to build on the lean design success story and do more of what has proven to work well in 5G. A 6G lean design concept can be summarized as follows:

Enhanced lean design in the time domain: This implies further reducing the time-domain footprint of idle-mode signals, related to system-information broadcast, paging, and random-access. The discontinuous reception of multiple UEs can sometimes be aligned to better support discontinuous transmission on the NW side, and vice versa. This will further enhance opportunities for micro-sleep reception and micro-sleep transmission in the network equipment.

Extending lean design to the spatial domain: System-information transmission can occur from a subset of the transmission points. The system-information broadcast function can also use single-frequency-network (SFN) transmission formats to further enhance the coverage of the system-information broadcast.

Extending lean design to the frequency domain: We should further improve solutions that enable carriers to operate without transmitting system-information. Also, not all carriers need periodic downlink mobility reference signal (RS) transmissions. In general, certain requirements and functions may only be supported on a subset of the carriers. This will enable to dynamically adapt the capabilities of transmission points on selected carriers to match the current traffic requirements.

Note that lean design in the spatial and frequency domain has been gradually introduced in later 5G-NR specification releases. However, carriers without Synchronization Signal Block (SSB) transmissions, although supported in 5G NR, have not yet been deployed in live networks. This is primarily due to a lack of support in UEs that are compliant only with older 5G NR release specifications. In a new 6G system, we can further enhance these lean design enhancements and ensure that all UEs support them in the first release.

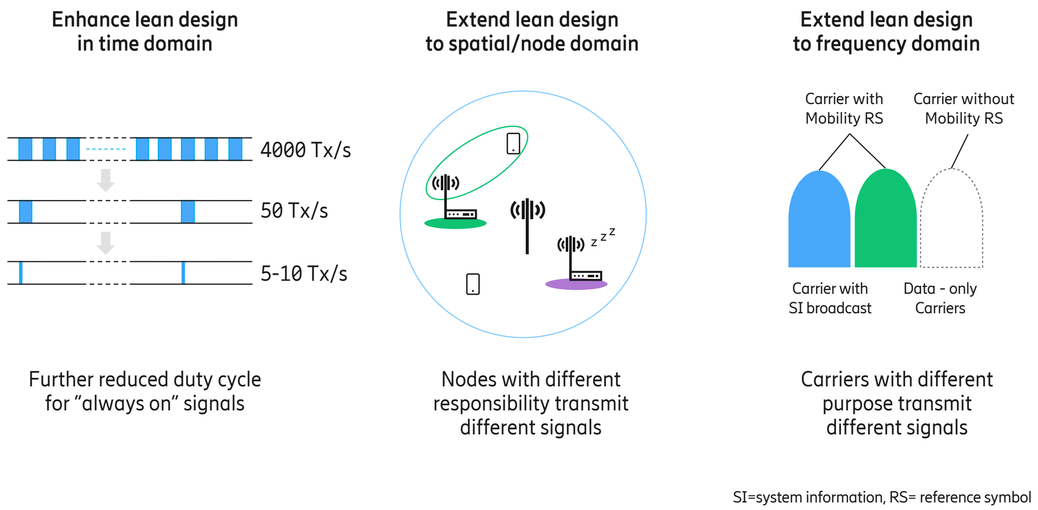

Figure 3: 6G should deliver enhanced lean design in the time domain (left) and extend the lean design to spatial (middle) and frequency domains (right).

For 6G, we want the basic period for idle-mode cell search to be increased from 20 ms (which is the maximum SSB periodicity for stand-alone operation in 5G NR) to something significantly longer. Initial evaluations of the impacts of signaling periodicity and signaling size on energy consumption are shown in Figure 4. The first (left-most) bar is the 5G NR reference case with 20 ms SSB periodicity, system-information transmission, and random-access reception. The middle bars show the energy consumption when using 40, 80, 160, 320, 840, and 1280 ms periodicity for SSB transmission and random-access reception, respectively. The final, right-most configuration is the energy consumption of a base station in deep sleep [2].

We note that the energy saving for the 160 ms case is significant (77 percent) compared to the left-most case representing 5G NR. We consider this configuration to be a good achievable target for normal 6G operation. A longer SSB periodicity provides greater energy-saving benefits. However, at some point, the UE energy consumption and time spent for the initial cell search will become unacceptably high. These configurations can however be useful as an alternative to turning off a cell completely, but we cannot expect idle-mode UEs to detect these extended periodicity SSB transmissions. For example, an extended SSB periodicity of 1280 ms uses only twice the energy, compared with putting a cell in a deep sleep. With side information from the network, the connected-mode UEs can still measure on these cells.

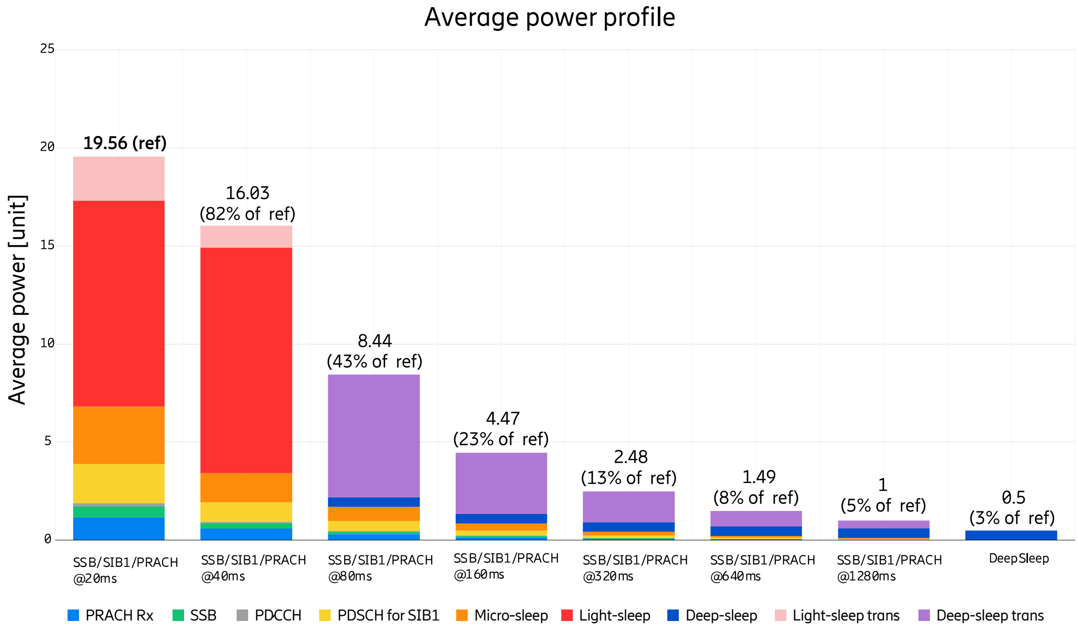

Figure 4: Breakdown of network energy consumption for SSB/SIB1/PRACH periodicities equal to or larger than 20 ms.

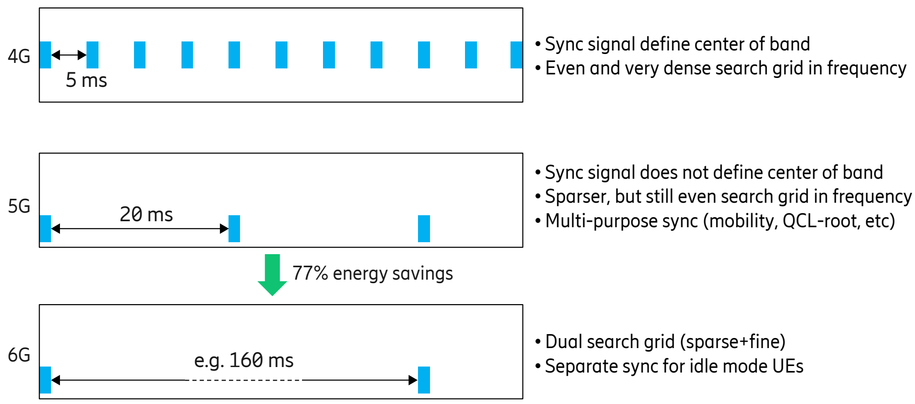

Figure 5: The idle-mode synchronization signal periodicity increased from 5 ms in 4G to 20 ms in 5G. For 6G, we should target a significantly larger value, such as 160 ms.

One problem in the design of the first 5G NR release is that several functions with very different requirements are coupled together through the SSB signal. This coupling often hampers the implementation of energy-saving optimizations. Although this has been partly addressed in the subsequent 5G NR releases, these enhancements have not been widely deployed in real networks, primarily due to a lack of support in legacy UEs. It may appear efficient to reuse signals for multiple purposes instead of defining and transmitting separate reference signals for distinct functions, but what often happens is that the function with the toughest requirements defines the design. Some functions require fine spatial resolution (like connected-mode mobility) while others (like system-information broadcast and paging), do not. When the same reference signal is reused, it needs to fulfill requirements in every dimension, and the resulting design can quickly become very dense in the time, space, and/or frequency domains, resulting in high energy consumption.

Different functions in a cellular system have different requirements. Especially requirements on idle-mode functions (such as initial cell search, system-information broadcast, random-access, and paging) and connected-mode functions. The connected-mode functions can be further divided into serving-cell related functions (such as data transmission, reception, and measurements), and target-cell related functions (such as mobility measurements). These functions have very different requirements in the frequency, time, spatial, overhead, and adaptability domains:

Frequency: Idle-mode functions are not always needed on all frequency bands. For the initial access, it may not be necessary for a specific spectral resource (such as carrier frequency band, etc.) to be visible to idle-mode UEs. However, for connected-mode mobility, it might be necessary for UEs to be able to measure on that frequency band.

Time: For idle-mode operation, we want to extend the maximum transmission periodicity to at least 160 ms for the default 6G operational mode (similar to the 20 ms SSB periodicity supported for NR stand-alone mode). To achieve success in this area, we should design the idle-mode functions while only considering idle-mode requirements. This approach can involve the introduction of additional signals to support the needs of connected-mode functions.

Spatial: The required beam shape of an idle-mode SSB and a connected-mode SSB can be very different. The idle-mode SSB transmissions can be made more efficient if we can transmit them with wider beams and from more than one transmission point. For connected-mode SSB signals, we typically want to have narrow beams to provide high signal-to-interference-plus-noise ratio (SINR) on the synchronization signals or to enable mobility directly to a good beam in a target node.

Overhead: The additional transmissions needed for an idle-mode and a connected-mode SSB are quite different. The idle-mode SSBs need to be associated with system-information broadcast transmissions, while the connected-mode SSBs do not.

Adaptability: It is often preferred that all idle-mode signals used for UE initial cell search can be static and fixed. There are potential negative consequences of changing an idle-mode signal in a way that changes the area coverage of the system. For most operators, it is critically important that once coverage is provided in an area, it should always be there. Reducing the network operational cost is often not a good enough argument to jeopardize reliable area coverage. For connected-mode signals, there are no such issues.

Increasing the idle-mode SSB periodicity in the time domain from 20 ms (used in NR) to something larger, such as 160 ms, must be done without degrading the UE idle-mode cell search time and the UE battery life, or increasing UE complexity.

In 5G NR, the idle-mode sync-signal periodicity was extended from 5 ms to 20 ms by eliminating the requirement for the sync signal (that is, 5G SSB) to be positioned at the center of the band. As a result, we made the search frequency-domain space sparser, without negatively impacting the total UE search time, see Figure 5. For 6G, we can extend this principle further only by requiring that every operator should have at least one frequency location in every frequency range (such as, low, mid-low, mid-high, and high) where a sparse SSB periodicity of 160 ms can be used. To support idle-mode functions in special deployments, for instance, small stand-alone private indoor systems, some additional frequency points can be placed on a fine search grid supporting only a shorter SSB periodicity (like 20 ms). This is an acceptable tradeoff since (1) such special cases are rare overall, (2) idle-mode functions are only needed on a small number of bands, and (3) small indoor networks already have small total energy consumption (such as 10-20 W per node and a small number of nodes in total) due to the low output power and the small coverage area.

When connected to a network, the UE might receive additional information about SSB periods, available NW services, and band exclusion for raster point subsets for the current geographical area to expedite future initial cell searches. This can significantly decrease the average time (and increase battery life) for finding a cell in idle-mode.

Scalable by design – Fast adaptation

As discussed above, when designing complex systems (like 6G) it is often good to follow a separation of concerns design principle. This ensures that each function can be independently optimized without considering requirements from other functions in the system (see Figure 6).

If the idle-mode synchronization signal in 6G, here denoted as I-SSB to separate it from other connected-mode related signals, is tailored exclusively for idle-mode functions, it may not have to be transmitted from every transmission point or potential beam in the network. In that case, connected-mode mobility measurements on target cells may need to be based on a different signal than I-SSB, for example, on a new downlink mobility reference signal. In connected-mode, the UE shall not assume that an idle-mode synchronization signal can also be used for connected-mode measurements. Connected-mode mobility decisions can also be based on UL transmissions, secondary-carrier predictions, or derived from UE positioning, in which case additional mobility reference signals may not be needed.

Here we can see why a separation of signals for idle-mode and connected-mode functions is required in 6G. This separation also enables a much more dynamic usage of synchronization signals such as allowing them to be turned on or off on a fast time scale, dynamically change the transmission power or beam shape, or move synchronization signals from a source node to a target node as the UE moves in the network. In 5G NR, there was always a risk of accidentally creating a coverage hole if a cell-defining SSB transmission (CD-SSB) is adapted in any way. In contrast, connected-mode functions in 6G should primarily rely on aperiodic and on-demand transmission of signals rather than periodic transmissions.

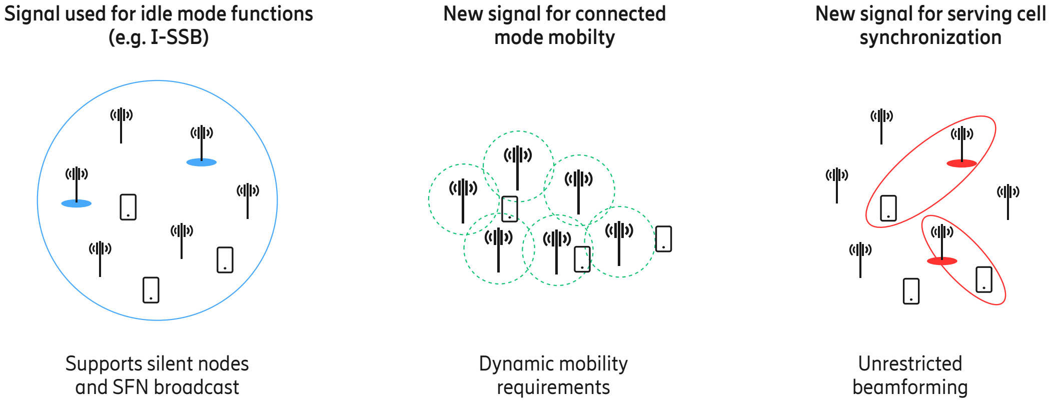

Figure 6: In 6G, distinct functions such as "idle-mode SI broadcast" (left), "measurements related to target cells" (middle), and "additional signals related to the serving cells" (right) should by default rely on different signals, thereby enabling independent optimization of each function.

Requirement management – Service requests and activation

Enhanced observability of both network energy consumption and user experience will be required in 6G networks. Lack of sufficient observability typically results in network designs with static overprovisioning as the main solution for managing demanding requirements.

In 6G, requirements can be expected to be much more demanding across several dimensions. The energy cost of providing these requirements everywhere all the time has the potential to hinder the rollout of 6G in several areas. To ensure the commercial success of 6G networks, we therefore need to develop the capabilities required to operate the network such that the delivered capacity and performance are not too much, not too little, but just right.

Demanding 6G requirements, not constrained in time and space can otherwise force the entire network into a state of high alert with exceedingly high energy consumption. This is because today the main driver of network energy consumption is increased requirements rather than increased traffic volume. Due to the traffic characteristics and lack of fast network adaptability discussed above, the energy consumption in a country-wide network has a very low correlation with traffic volume. However, the network's energy consumption greatly depends on its design requirements, such as cell edge user throughput, latency, reliability, and capacity.

A solution for RAN hardware-requirement management that can dynamically adapt the configurations of hardware units in correspondence with current quality-of-service requirements is needed. This involves mapping service requirements to the relevant hardware units and converting service requirements (rate, latency, priority, etcetera) to requirements pertinent to reducing the energy consumption of radio units (including linearity, distortion, transmitted RF power, bandwidth, number of antenna ports) and RAN compute units (such as required compute power, power gating, clock gating, dynamic voltage and frequency scaling (DVFS) configuration). The overarching goal of such a requirement management function is to ensure that energy-demanding requirements are limited in time and space, and thereby enable energy-saving features to be safely applied without the fear of any negative impact on relevant key performance indicators.

Summary and conclusion

The lean design of 5G NR was a remarkable success and for 6G there is still more to gain by further improving a future 6G system based on these principles. For example, by extending the periodicity of the default synchronization signal from 20 ms to 160 ms, the idle-mode energy consumption of 6G can be reduced by a factor of 4 compared to 5G NR. In addition, we expect further energy savings by extending the lean design to the node and frequency domain, resulting in a total idle-mode network-energy savings of more than 4 times for 6G compared to 5G NR. To achieve this, the design should follow a separation of concerns principle such that connected-mode requirements (bitrate, reliability, latency, etcetera) do not impact the design of idle-mode functions.

An energy efficient network also needs to be carefully tuned based on real-time traffic and performance requirements. Scalability and fast adaptability need to be considered in the design to allow the network to always operate with just the right amount of activated hardware components. This requires enhanced observability of end-user experience as well as real-time network energy usage to enable fast and fine granular hardware configuration and management.

References

1. GSMA Intelligence, Going green: measuring the energy efficiency of mobile networks, February 2024.
2. 3GPP, Study on network energy savings for NR (Release 18), TR 38.864.

Authors

Pål Frenger is an esteemed expert at Ericsson with over 25 years of experience. He is currently working at the radio access department at Ericsson Research in Linköping, Sweden, where he specializes in RAN energy performance. Pål has been involved in standardization and development efforts for both HSPA, LTE, and NR and is now focused on defining the scope of future cellular networks including 6G. With over a decade of experience in radio network energy performance, he also has a rich background in link-level research covering areas such as massive MIMO, OFDM, and Turbo codes. Pål has filed more than 300 patent applications and was honored with the Ericsson Inventor of the Year Award in 2017. He holds a Ph.D. in electrical engineering from Chalmers University of Technology, Gothenburg, Sweden.

Ylva Jading is the senior specialist in network energy performance at Ericsson Research in Stockholm, where she has been working since 1999. Throughout her tenure at Ericsson, she has held several roles in product development, standardization, and research. Currently, Ylva is the technical manager of the networks energy performance program, focusing on researching and developing low-energy concept for future 6G networks She also leads the technical early-phase energy-performance activities in the business unit networks. Ylva holds a Ph.D. in physics from the Johannes Gutenberg University Mainz in Germany.

Ali Nader joined Ericsson in 1999 and has held various positions within R&D since then. He has extensive experience in the systemization and development of a wide range of 3GPP-based (2G, 3G, and 4G) network access functionalities for mobile phones, platforms, and chipsets. In 2014, he shifted his focus to radio base stations, initially focusing on concept development, standardization, systemization, and the development of the initialreleases of the massive IoT technology NB-IoT. His work later expanded to include systemization of various mobile broadband 5G-related functionalities. Since 2018, Ali has also been involved in the concept development, systemization, and standardization of energy performance features for both devices and networks.