

ENERGY EFFICIENCY IN WIRELESS COMMUNICATIONS

Higher Network Performance with Lower Energy Consumption

David E. Newman and R. Kemp Massengill

Executive Summary

Energy efficiency is a critical requirement for next-generation wireless. Energy costs for transmission are rising exponentially with demand growth, while mobile and IoT users face declining battery lifetimes due to increased computation and retransmission burdens. The success of our ambitious 5G-Advanced and 6G goals will depend on developing better wireless for less energy.

Disclosed below are options for reducing energy consumption while maintaining or improving network performance. Immediate energy-saving options include: (a) lean beam control without beam scanning, (b) simplified initial access without blind searching, (c) rapid detection of downlink messages without excessive computation, (d) efficient phase noise mitigation at high frequencies, and (e) fault correction without a retransmission. All of these energy-saving options are readily available [1] for implementation as software updates, with no changes required in the hardware. Each option provides substantial energy reductions and, in certain versions, improved network performance overall.

It is essential that designers be made aware of these energy-saving options before the next-generation standards are finalized.

A. Energy Savings by Beam Alignment - Without Beam Scanning

In 6G, every communication is carried on a narrow transmission beam focused on the recipient, and every recipient uses a narrow reception beam focused on the transmitter. Unfortunately, the old-fashioned "beam scanning" alignment procedures involve numerous transmissions at different angles, followed by multiple heavy feedback messages, followed by multiple further transmissions with various beam widths and transmission powers, followed by further time-consuming steps. More energy is consumed in every step. The entire beam scanning procedure must be repeated for each user in the network. If anything changes, such as locations or atmospheric conditions, it must be done all over again.

Therefore, a quick and efficient beam alignment procedure has been developed, which enables all of the user devices in a network to align their beams simultaneously, at a total cost of just a few resource elements. Figure 1 (below) shows how it works. At a scheduled time, the base station transmits a special "angle-dependent alignment pulse" in which the modulation depends on the angle. Each user device detects the angle-dependent signal and measures the modulation parameter at its location. The user device then calculates the alignment angle directly from the measurement, using a predetermined algorithm. The user device can then aim its uplink transmission beam and its downlink reception beam toward the base station antenna, for optimal communications at minimal power. The user device can then inform the base station of the alignment angle upon the next uplink communication. The base station and the user device can then use directed beams thereafter, with resulting energy savings for both.

Figure 1 is a graph of an angle-dependent signal. The base station transmits a brief (single-resource-element) signal at a prescribed time, extending all around the base station antennas, modulated so that the phase is different for each angle. In the depicted case, the phase equals the angle, so 360 degrees in phase corresponds to 360 degrees in angle. The user device simply measures the phase of the signal at its location, and instantly knows its direction relative to the base station. All of the user devices in the network can measure their directions simultaneously from a single angle-dependent signal. There is no beam-scanning and no feedback, other than a brief reply message informing the base station of each user device's angle (preferably appended to another message, such as an acknowledgement, and only if the angle has changed).

Antennas capable of beamforming are generally able to produce such a signal distribution, with no changes other than the beamforming instructions. All of the energy waste of the legacy beam alignment procedures has been eliminated.

Optionally, the base station can transmit a second "calibration" signal isotropically, with a uniform phase (such as zero degrees of phase) all around the base station. The user devices can compare the two signals to get a better measure of their directional angle. As a further option, the base station can send out a "vernier" signal with greater phase variation, such as 360 degrees of phase in each 90 degrees of angle, thereby making four complete phase cycles around the antenna. The user device can then determine its angle more precisely by comparing the vernier phase with the original angle-dependent signal phase.

As an alternative, the angle-dependent signal may be modulated in amplitude instead of phase. The user device determines its angle by measuring the received amplitude, relative to a second unmodulated, isotropic, calibration pulse. For disambiguation, the base station may transmit two or three angle-dependent signals with different amplitude-vs-angle distributions. The user can then determine its direction based on the amplitude ratios, using an algorithm. Absolute amplitude calibration is not required.

At very minimal cost, the base station can transmit the angle-dependent signals periodically, so that the user devices (especially mobile) can check for any changes in alignment. If the alignment angle has changed, the user device can inform the base station. For example, the user can append feedback to the next uplink message, or reply at an assigned time after making the measurement.

The angle-dependent signal provides instantaneous beam alignment for all the user devices in the network simultaneously. It eliminates time and energy costs of legacy beam-scanning procedures, and enables faster responses to changing conditions. It also avoids generating backgrounds and interference. Since the new procedure is lean and efficient, the base station can broadcast an angle-dependent signal periodically, such as once per frame or subframe, with negligible impact on throughput. The base station can also transmit an angle-dependent signal on the broadcast channel, for example after each SSB message, thereby enabling new users to align their beams before attempting to communicate with the base station. This also enables all users to check their alignment whenever desired, at zero cost.

In summary: Using only one or a few modulated resource elements, all of the user devices in the network can update their beam alignment, thereby enabling precise control of the transmission and




Fig. 1: Base station transmits a single modulated pulse modulated to have a different phase (or amplitude) at each angle.

reception beams, for greatly enhanced signal quality and range. Energy savings are due to the reduced transmitted power required for good reception, avoided retransmissions due to higher reliability, and the greatly reduced energy cost involved in the beam-alignment procedure itself. These improvements can be implemented in software, with no changes required of the hardware, including the base station antennas.

Value-chain Analysis - Beam Alignment

For users:

Efficient beam alignment enables users to obtain longer battery life by avoiding message-heavy beam-scanning. They also obtain improved signal quality due to more frequent beam adjustments. The improved uplink signal results in fewer interruptions and fewer retransmission requests, which results in lower latency and greatly reduced power. In addition, mobile users can update their alignment every 20 ms at zero cost, thus keeping aligned even as they move and turn. Rapid real-time beam alignment also enables narrower beams without fear of mobile users drifting out of the beam.

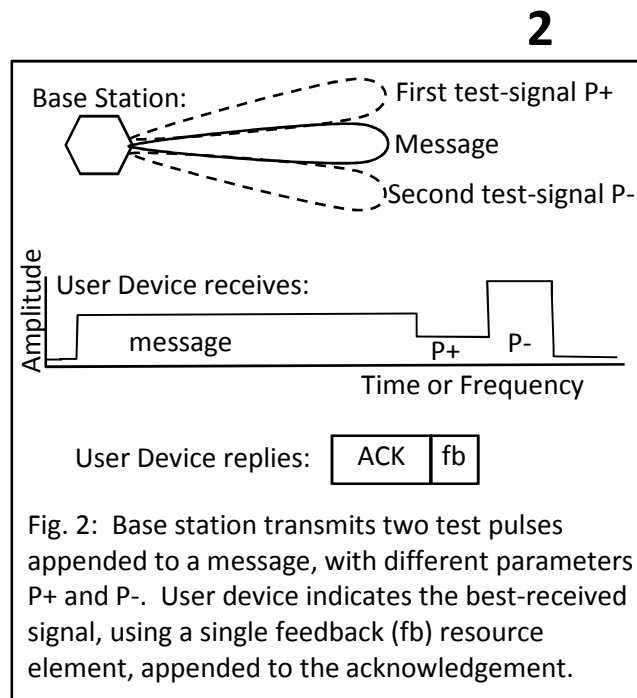
For networks:

Networks can save substantial power by avoidance of beam scanning for each individual user; the angle-dependent signal aligns them all simultaneously. With improved beamforming, the base station will experience fewer dropped calls, due to better signal quality in both uplink and downlink. Customers will be happier, due to the superior communication experience, and will select such networks again whenever there is a choice. In addition, the cost of transmitting the alignment pulses is negligible.

B. Energy Savings by Lean Real-Time Feedback - Without SRS/CSI

Besides the beam direction, there are several other beam parameters that must be controlled to optimize energy efficiency, such as beam width, transmission power, polarization, use of repetitions and/or FEC bits, among others. Frequent feedback is also needed to keep the beams optimized in changing atmospheric and background conditions, and of course for mobility. If current feedback protocols were used for real-time fine-tuning, the energy and resource costs would be enormous. Therefore, a terse, energy-efficient, real-time feedback code is disclosed, which the base station and every user device can apply, at near-zero cost.

Figure 2 shows the new procedure. The base station appends two short, single-resource-element "test signals" P+ and P- at the end of each downlink message that requires a reply. The two test signals are transmitted differently, such as at slightly higher and lower beam widths (or power or polarization or direction or elevation or frequency, etc.) relative to the message itself. The recipient can then send a reply or acknowledgement (ACK), plus a single feedback resource element (fb) modulated to indicate whether the first or second test signal, or the message itself, was best received. By this



method, the base station can apply small incremental fine-tuning adjustments frequently, thereby keeping the downlink beam optimized at very minimal cost in power and resource elements.

The user device can keep its uplink transmission beam optimized the same way, by appending two test signals to each uplink message that requires a response, and receiving the base station's very brief feedback indicator appended to that response.

The feedback modulation can also indicate whether the transmission power level is too high, too low, or just right. Based on changing absorption or interference conditions, transmitters adjust the power so that the received signal provides a sufficiently reliable reception, but no more. The beam alignment selection and the power adjustment can be multiplexed in the single-element feedback signal. Proper adjustment of the transmission power saves energy, and also minimizes interference for the other users, all at a cost of just three resource elements.

The feedback resource element can include multiplexed information about the best-received test signal, and also the need for any power adjustment, such as incremental increases or decreases in transmitted power by a predetermined increment. For example, the feedback message element can include at least eight modulation states, corresponding to selecting the first or second test signal, or the message itself, as best-received, multiplexed with a request to increase or decrease or leave unchanged the power. If no changes are needed (that is, the message was best-received and the power is correctly set) then the feedback message element can be left blank, and the transmitting entity will understand that blank means no changes are needed.

For mobile users, the Doppler shift changes each time the user accelerates or changes direction. Frequent Doppler adjustment is essential to avoid subcarrier crosstalk at the base station and the user device. Instead of message-heavy legacy feedback protocols, the user can append two test signals with slightly different frequencies to each uplink message. The user device can maintain tight frequency control incrementally, thereby compensating changes in speed and direction.

In summary: Test signals, appended to messages that require a reply, and a single message element feedback, provide an extremely low-cost incremental beam control option for base stations and user devices. The real-time incremental adjustments save energy by avoiding complex feedback procedures, while enabling multiparameter beam optimization with very high granularity, thereby greatly reducing retransmission costs and energy. The lean feedback procedure can be implemented by software updates alone, without hardware changes.

Value-chain Analysis - Lean Feedback

For users:

Frequent incremental beam adjustments result in optimal reception, especially in places where the backgrounds vary faster than the network can perform beam scanning. Lean feedback also avoids excessive transmission power, thus saving battery time. The short single-resource-element feedback, appended to an acknowledgement, greatly reduces transmission energy relative to legacy processes - in part because it replaces an entire cycle (or two) of message exchanges. For even greater efficiency, the user can simply skip the feedback addendum when no changes are needed; transmitting nothing costs nothing. Advantageously, no hardware changes are needed, just a software update.

For networks:

The short test signals, appended to another message, use far less resource area and transmission power than legacy beam-scanning since no request/response messaging is needed. In addition, optimizing the beam direction and width and power ensures satisfactory reception, while minimizing base station

energy costs, and consuming less resource area. The procedure allows networks to separately optimize a wide range of transmission parameters. Keeping the signal reception high for each mobile user, through responsive beam adjustments, generates superior QoE and happy customers. No hardware changes are needed, just a software update.

C. Energy Saving by Automatic Receptivity Compensation - Without Dead Zones

Mobile users and battery-operated IoT devices are especially sensitive to energy efficiency. The base station can help mobile users to conserve energy by avoiding unnecessary positional updating messages, and by proactively compensating for receptivity changes based on the user's location.

The base station can keep an internal map (Figure 3) to track its mobile users in real-time. Each mobile user initially indicates its location, speed, and direction of travel on entry. The base station locates the user on the internal map, and then follows the mobile user's subsequent trajectory, adjusting the downlink beam direction and power level accordingly. Users can inform the base station whenever they change direction, and otherwise remain silent, thereby saving further transmission energy.

For even greater power control, the base station can measure the wireless receptivity versus location within its coverage area. The base station can then automatically correct the downlink power to maintain adequate signal quality at each user device, thereby compensating receptivity changes. This eliminates "dead zones" due to hills, tall buildings, and the like. For example, the base station can predict when a mobile user will enter and exit a dead zone, and can increase or decrease the power accordingly to maintain adequate signal quality. Users will appreciate not losing their connections.




Fig. 3: The base station calculates each vehicle's location based on data in the vehicle's entry message, and adjusts downlink power using an internal map. Base station also boosts power when the user is in a predetermined dead zone.

The base station can also tell the mobile user device to increase or decrease its uplink power at the same time as the base station adjusts the downlink power. However, most mobile users have very limited uplink communication needs, and therefore such real-time adjustment may be unnecessary in most cases. Instead, the mobile user device can keep a copy of the receptivity map for its own use, thereby adjusting its uplink transmission power autonomously when needed to compensate for changing propagation losses, without exchanging unnecessary position messages with the base station.

In summary: Base stations and mobile users can save energy by adjusting their transmission power for adequate signal quality, despite changing distances and receptivity. Excess transmission power wastes energy and generates troublesome backgrounds for other users, whereas insufficient transmission power leads to low communication reliability, frequent retransmissions, and dropped calls, all of which

cost additional energy. Instead, at zero cost, the base station can trace the predicted location of each user device in real-time, automatically adjusting the downlink beam direction and power to keep the mobile user adequately served, without wasting excess power. Energy-intensive legacy positional feedback messaging can be avoided while improving signal quality. These improvements can be implemented by software updates, with no change in hardware. The resulting energy savings, and improved reliability, and reduced latency, will provide a substantially better communication experience for mobile users.

Value-chain Analysis - Receptivity Compensation

For users:

Avoidance of legacy positional messaging saves transmission energy, and hence battery life for mobile users. Continuous proactive optimization of beam parameters by the base station provides uninterrupted communications for mobile users, despite changing angles and distances. Users will notice, and appreciate, that contact is maintained even when passing through areas previously considered dead zones.

For networks:

Tracking mobile users according to current traffic patterns is free. This avoids frequent positional messaging, thereby saving energy and avoiding background generation. Automatically adjusting downlink power, according to the receptivity at the user's location, also maintains signal quality and avoids dropped calls, and their associated costs and delays. Advantageously, no hardware changes are needed, just software updates, plus a one-time charge for preparing the regional receptivity map.

D. Energy Savings by Network Database - Without Blind Searching

Before a mobile user device can join a network, it must perform a complex and energy-intensive series of steps, with unknown chance of success. First, the user must perform a blind search of numerous frequencies that might, or might not, be used by a base station. The scan necessarily involves isotropic reception because the alignment direction is unknown. If the user device is lucky, it discovers an SSB message, and finally the user device can begin attempting to contact the base station, all without knowing whether the base station is the most appropriate one for the user device. In many cases, it is not, and the user device must start over at the beginning. To actually contact the base station, the user must wait for a random access interval, then hail the base station repeatedly, using successively higher power until finally getting a reply. After some number of failed entry requests, the hapless user device is required to abandon the process and resume blind-searching yet again. Needless to say, this procedure costs precious time and battery energy, with uncertain results.

Figure 4 shows a simpler and much more energy-efficient access procedure. A user device




Fig. 4: A mobile device uses its internal Network Database to find the nearest base station and the broadcast frequency. User then aligns its reception beam toward the antenna, receives system info messages, adjusts its power according to the distance, and transmits an initial access request.

simply checks a "network database" file, which lists the locations of all base stations in each region. The network database also lists the broadcast frequencies, so that the user can bypass all the arbitrary and uncertain steps, proceeding immediately to acquire the system information messages using a directional reception beam. Optionally, the network database can provide all of the information in the SSB and SIB1 messages, enabling the user device to skip that step too. The user can then adjust its power according to the distance, and send the entry request (random-access preamble) without further delay. If the user is in motion, it can also compensate the Doppler shift according to its direction. The user device can also inform the base station of the user's location shortly after making initial contact, so that both the user device and the base station can employ directed beams and appropriate power, right from the start. The resulting improvement in signal quality, for both uplink and downlink messages, thereby avoids many potential faults while saving time and energy on initial access and thereafter.

The user device can possess the network database in any number of ways. The network database file may be installed by default in each mobile wireless device, or it may be downloaded at a convenient time before the trip. It may be updated periodically, such as after registering with each base station.

In summary: The network database saves energy by (a) indicating the closest base station for best reception, (b) avoiding the tedious blind search and several other peculiar steps not mentioned, (c) enabling the user to set the proper beam direction without a beam scan, (d) setting the uplink transmission power based on the distance, without a frustrating power scan, and then (e) begin communicating in an energy-efficient directed beam right from the start. The network database can be prepared, downloaded, and employed by mobile users at zero energy cost, other than simply receiving the database. The network database, and the other procedures described, can all be implemented by software, requiring no changes to current and planned future hardware.

Value-chain Analysis - Network Database

For users:

The Network Database enables new users to avoid the costly blind searches of legacy procedures, and immediately select the closest public access point. Users can then begin listening on the right frequency, as indicated in the database. Users can also avoid the RACH power scan by calculating the distance. Mobile users can also correct for the Doppler shift, without feedback, based on the velocity and position relative to the base station.

For networks:

For a small subscription fee, each base station and access point in the network can be listed in a widely-distributed database that all users can freely obtain, thereby enabling immediate access into the subscriber networks. Since users can calculate the alignment direction before receiving the SSB message, networks no longer need to broadcast the SSB multiple times in different directions, thus saving transmission power, resource area, and background generation.

E. Energy Savings with SSB Location Symbol - Without Beam Scanning

Another way that the base station can help prospective user devices is to indicate the location of the base station in the system information messages (SSB and/or SIB1) which are broadcast by base stations periodically. For example, the latitude and longitude of the base station antenna can be appended to the SSB message as shown in Figure 5. New user devices can align their transmission and reception beams toward that location, and adjust their uplink power according to the calculated distance, and can also calculate the Doppler shift if mobile, all at zero cost, before transmitting the random-access preamble to the base station. This saves energy by enabling the user device to use a directed reception beam for

receiving the SIB1 message, and avoids subcarrier crosstalk by adjusting the Doppler shift, thereby enhancing reception - without a beam scan and without a power search. The user device can then use a directed transmission beam for the random-access preamble, at the estimated correct power level, and with the Doppler shift negated, thereby providing excellent communication upon first call, and all at a cost of just one extra symbol in the SSB message.

The figure also shows an additional feature that assists prospective user devices, or any other user device in its network, to optimize its downlink reception beam empirically. The base station appends a series of uniform tuning signals to each SSB transmission. Tuning signals are identical transmissions, broadcast isotropically and sequentially in time. The user device can then receive the uniform tuning signals while varying its reception beam direction and width, thereby optimizing reception before attempting to receive the SIB1 message. The user can also inform the base station of the user location, appended to the Msg3 or MsgA, or in a subsequent uplink message, or at another convenient time, thereby enabling the base station to use energy-efficient downlink beams as well.

In summary: The base station indicates its (antenna) coordinates, and optionally a series of uniform tuning signals, appended to each SSB message. A new user device can then calculate the distance, direction, and Doppler shift based on the location and motion of the user device. The user can also fine-tune its reception beam using the uniform tuning signals. The user device can then receive the SIB1 message using a directed reception beam for high reliability, and can transmit the initial access request using a directed transmission beam, at the correct power and frequency for optimal reception by the base station. In addition, the user device can provide the location of the user device in, or appended to, a follow-on message, thus enabling the base station to use directed beams and appropriate power for communicating with the user device thereafter. Importantly, these benefits are obtained at zero energy cost by the user device and negligible extra transmission power by the base station. Transmission energy is then saved for both user and base station, by enabling the user device and the base station to use directed beams during the initial access procedure instead of relying on isotropic signals that are intrinsically wasteful and interference-prone, and for all communications thereafter. All of these improvements can be implemented using software updates alone, requiring no hardware changes of any kind.

Value-chain Analysis - SSB Location Symbol

For users:

The modified SSB enables each user to align its beams right after detecting the SSB, thus obtaining better signal quality for receiving the other system information messages. Based on the antenna location, users can transmit the RACH and other entry messages on aligned transmission beams, thus reducing faults and failed initial-access attempts. Users can also set their power level according to the distance, thereby avoiding a time-consuming RACH power scan.

5




Fig. 5: The first system information message (SSB), followed by the latitude and longitude of the base station's antenna, and a series of uniform tuning signals to help user devices align their reception beams.

For networks:

For the negligible cost of adding one OFDM symbol to the SSB, networks can provide valuable location data to all users, thereby enabling new users to obtain greatly improved signal quality from the start, and avoiding costly faults during reception of the other system information messages. Users that have a choice of networks will preferentially select networks that make access easy.

F. Energy Savings in Downlink Control Message Discovery - Without Search Spaces

A big problem for user devices is that they never know when they have a downlink control (DCI) message, and therefore must constantly scan the downlink control channel. The task is greatly complicated by the user's identification code being scrambled with the error-detection code, thus forcing the user to decrypt, unscramble, and demodulate all possible combinations of starting time, starting subcarrier, and message length - a monumental task. The user device must do it all within one symbol-time of each downlink control interval to avoid missing its messages. Unfortunately, it is impossible to correct message faults when the ID is scrambled with the error-detection code, and this forces automatic retransmissions when a fault is seen, with further energy cost. Unfortunately, the size of the error-detection code had to be increased from 16 to 24 bits, due to the number of false positives, which further burdens the receiver. As a patch-up, the network can restrict users to certain "search spaces", but this only marginally reduces the energy cost while increasing the average latency for those users.

Figure 6 shows a much more economical solution. Each message on the downlink control channel is marked by an easily-recognized demarcation, at the start and end of the message. In this case, the demarcation is a blank (no transmission) resource element B at the start and end of the message M. The user device readily determines where each downlink control message starts and ends, greatly reducing the amount of computation required. The error-detection code can again be returned to the 16-bit size, saving further power and resources. Blank demarcations are easy for the base station to insert; they can simply transmit no power in the subcarrier corresponding to the demarcations. The resulting energy savings will enable many battery-constrained applications such as IoT.

The figure shows further energy-saving options. The base station can include the user's identification code (ID) in the message, but NOT scrambled with the error-detection code. This enables the user device to determine whether the message is intended for it, and also enables the user to recover faulted messages, as detailed below. The message can also include a demodulation reference signal (RS) for immediate amplitude and phase correction. The fourth message is time-spanning and indicates the length (LL) up front. As a further option, the user device may request that its downlink data messages (on PDSCH) be demarked in the same way. The user can find its downlink data messages clearly demarked, in which case its DCI messages are no longer needed. Omitting the redundant DCI messages saves further energy.




Fig. 6: DCI messages (M) can be demarked by a blank (B) resource element at start and end. They can also include the recipient's identification (ID), a demodulation reference signal (RS), and the message overall length (LL), to assist the user device.

In summary, demarcations enable the beleaguered user device to retrieve its messages easily, saving huge amounts of battery-draining energy for the user, at negligible cost to the network. The error-detection code could be safely returned to 16 bits, further saving energy. The byzantine search-space limitations could be eliminated, reducing latency. Identifying the recipient separately from the error-detection code, enables the user to recover its messages even when faulted, (using strategies detailed below), avoiding a costly retransmission and saving further energy. All of these improvements can be implemented in software, without requiring hardware changes in the base station or the user device.

Value-chain Analysis - DCI Discovery

For users:

Users can save a huge amount of compute energy by simply scanning the raw branch data of the PDCCH signal for a unique pattern: zero Q and I amplitudes. Users can avoid inefficient custom search spaces, thereby regaining low latency. In addition, by not scrambling the recipient ID with the CRC, users can recognize their messages even when faulted, and then can attempt to correct them without a retransmission.

For networks:

With DCI demarcations, the base station is relieved of the burden of user-specific search spaces, greatly simplifying downlink scheduling. In addition, by removing the primary barrier to low-cost IoT devices, DCI demarcations enable numerous new use cases that were previously impractical.

G. Energy Savings by Advanced Modulation - Without Faulting

Most wireless messages are modulated in BPSK, QPSK, or one of the QAM orders such as 16QAM. Most receivers demodulate messages by separating the received waveform into orthogonal amplitude-modulated I and Q branches. However, additional information can be provided by adding modulation states that are easily differentiated, resulting in shorter messages and less energy consumption.

Figure 7 shows how BPSK can be enhanced with a third state of zero power. The receiver can easily differentiate zero power from the regular BPSK states. Messages can therefore be made shorter due to the increased information density (1.58 bits per symbol instead of 1) at zero energy cost.

Figure 8 shows a similar improvement in QPSK, with an additional zero-power state. The energy cost to transmit a particular message is substantially reduced with five modulation states instead of four.

Figure 9 shows how a similar approach can increase the information density in 16QAM by providing eight new states with zero power in one of the I and Q branches, plus a central zero-power state in both I and Q. The branch amplitude levels have been altered slightly




Fig. 7: BPSK has two modulation states of the same amplitude and differing 180 degrees in phase. A new zero-power state provides additional information density at no cost.




Fig. 8: QPSK has four modulation states of the same amplitude and differing 90 degrees in phase. The central zero-power state provides additional information density.

for equal spacing. Due to the extra states, the information density has been raised from 4 to 4.64 bits per symbol, with concomitant energy savings, and all at zero cost. Modulation errors are not expected to increase because receivers can easily differentiate the added states from the regular states.

Some may wonder whether the transmitter amplifier would suffer efficiency degradation due to large swings in power output when transmitting the zero-power states. In practice, this is not a concern because OFDM symbols comprise hundreds of subcarrier signals added together, thereby smoothing out power fluctuations in the actual transmission.

In summary: Zero-power modulation states provide higher information density, thereby saving energy due to shorter messages and use of states with no transmitted power. No hardware changes are needed. Therefore, it seems reasonable to add zero-power states to BPSK, QPSK, and QAM modulation schemes in future releases of the standards.




Fig. 9: 16QAM has 16 states with orthogonal amplitude-modulated branches I and Q. Here 9 additional states are added, having zero power in one or both branches. Messages can be shorter, saving both energy and time.

Value-chain Analysis - Advanced Modulation

For users:

Additional zero-power modulation states result in shorter messages, and hence less transmission energy consumption. The incremental cost of a zero-power state is zero.


For networks:

The zero-power modulation states also save energy for base stations, for the same reasons as for users. In addition, zero-power options provide networks with impactful parameters that they can adjust for optimal performance, based on current demand, background conditions, and fault rates.

H. Energy Savings by Phase Noise Mitigation - Without Loss of Throughput

Phase noise is a limiting factor in high-frequency communications. It is mainly caused by clock variations in the transmitter and receiver, and also by signal distortions due to interference. Demodulation errors from phase noise are proportionally worse in the projected 8.4 GHz band, and especially in the highly desirable FR2 frequency bands. The current solution is to limit all messages to low-order modulation, include bulky FEC bits in each message, and automatically retransmit the message upon any fault, all of which greatly increase the latency and energy consumption.

Figure 10 shows a "constellation chart" of 16QAM. The receiver separates the waveform into orthogonal I and Q branches. Each branch is amplitude-modulated according to



four levels, as mentioned. However, phase noise does not affect the branches directly; it alters the phase of the received waveform. The waveform amplitude (AMP) of each modulation state is the radius, and its phase is the angle relative to the horizontal axis. The gray regions show the effect of moderate phase noise, which spreads the received modulation states in phase only. As the figure makes clear, there are actually only three amplitude levels in 16QAM, since the middle eight states all have the same waveform amplitude. Even more crucial, the phase separation between adjacent pairs of states is only 36.9 degrees. This is an insufficient phase margin, and is responsible for most phase faulting in 16QAM.

In contrast, Figure 11 shows the modulation chart of amplitude-phase waveform modulation with 16 states. In amplitude-phase modulation, the transmitter modulates the amplitude and phase of the overall waveform according to the message bits. "Polar" and APSK modulation are types of waveform amplitude-phase modulation. The receiver still separates the incoming signal into I and Q branches as usual, but then calculates the waveform amplitude and phase before demodulating. Substantially increased phase margins can be obtained, at a negligible cost if just a couple lines of code to convert the branch amplitudes to waveform amplitude and phase before demodulating. No hardware changes or signal processing changes are needed.






Figure 11 demonstrates enormous advantages of demodulating in waveform amplitude and phase. The phase levels are equally separated by a full 90 degrees (unlike QAM) thereby providing much larger, and uniform, phase margins. There are four waveform amplitude levels (unlike QAM), and they are equally spaced (unlike QAM). The noise distribution is again shown in gray, using the same scale as in Figure 10. The phase spread looks smaller here because the phase margins are so much larger. The close phase overlaps of 16QAM are eliminated with waveform amplitude-phase modulation due to the much larger phase margins, even assuming the same amount of phase noise as in the QAM chart of Fig. 10. In addition, the four waveform amplitude levels can provide higher throughput than 16QAM, especially in congested networks, by providing additional, and equally spaced, amplitude modulation levels.

Amplitude-phase modulation also allows much greater versatility in modulation than QAM. Figure 12 shows a non-square amplitude-phase modulation table that provides a full 180-degree phase margin on every modulation state. There are eight amplitude levels in this case, 16 states in all. Note: the amount of phase noise plotted here is much higher than in Fig. 10, yet the states are safely separated due to the 180-degree phase margins. Thus the asymmetric modulation scheme easily accommodates even the worst-case phase noise, which QAM could never survive. This modulation scheme (2x8) would be ideal in situations where phase faulting is a problem and amplitude faulting is not, such as at high frequencies. Eight amplitude levels are feasible so long as a demodulation reference is close to the message (such as prepended to the message itself). Other non-square modulation examples have 3 or 5 equally-spaced phase levels, and as many amplitude levels as the noise environment allows. Such asymmetric modulation is impossible in QAM and the other planned 6G



modulation schemes, but is easy with waveform amplitude-phase modulation. As mentioned, the receiver processes the received signal by separating the I and Q branches as usual, and then calculates the waveform amplitude and phase before demodulating. Faults are avoided, even in extreme phase noise.

A further advantage of amplitude-phase modulation is that it enables fault diagnosis by type. As shown in Figure 13, adjacent-amplitude faults are distorted by one amplitude level, adjacent-phase faults by one phase level, and non-adjacent faults by multiple levels. Fault-type diagnosis is impossible with QAM because noise affects every modulation state differently. Network operators can use amplitude-phase modulation to diagnose message faults by type, and respond by changing modulation parameters accordingly. For example, if most faults are adjacent-phase, then switch to a modulation scheme with wider phase margins. If the faults are mainly adjacent-amplitude, use fewer amplitude levels and possibly more phase levels. If non-adjacent faults predominate, the problem is most likely bursty interference, which case small changes are unlikely to help. Instead, the network can move traffic into quieter time windows and frequency bands.



In summary: Amplitude-phase modulation saves energy by providing larger phase margins, and therefore fewer phase faults, especially at high frequencies. It also saves energy by increasing message throughput, especially in crowded networks, due to the four waveform amplitude levels instead of three. It also saves energy by making asymmetric modulation feasible, thereby optimizing performance in real environments. In addition, amplitude-phase modulation provides the ability to categorize faults according to cause, enabling a rational response, and thereby saving energy in the form of avoided retransmissions and reduced power per message. Amplitude-phase modulation can be implemented in software, with no changes in hardware. Messages are easily demodulated in the receiver after the usual I-Q branch signal processing. For these reasons and others, the option of amplitude-phase modulation should be made available by networks right now in 5G, and certainly for 6G where phase margins will be critical.

Value-chain Analysis - Phase Noise Mitigation

For users:

Waveform amplitude-phase modulation provides substantially larger phase margins than QAM of the same order, greatly reducing phase faulting. Avoidance of the interruptions and retransmission costs associated with phase faulting represents a major cost reduction. In addition, the large phase margins of waveform modulation enable much higher frequencies to be used on both uplink and downlink, which is essential for high-volume 6G traffic.


For networks:

In addition to all the value items listed for users, the base station also benefits by being able to use asymmetric modulation to precisely target fault rates and other problems. Diagnosing faults by fault-type allows the base station to specifically mitigate the observed fault types, instead of blindly reducing the throughput, as is currently done. Networks can select a modulation scheme, often an asymmetric modulation scheme, to optimize network performance in ways that QAM cannot.

I. Energy Savings by Fault Correction - Without a Retransmission

Currently, the automatic response to a corrupted message is another transmission - either a retransmission of the message itself, or a transmission of the FEC bits which may or may not be helpful.

Figure 14 shows a far more efficient solution, enabling the receiver to correct the fault autonomously, without a retransmission. In this example, the receiver determines the likely faulted message elements according to the "modulation deviation" - the distance between the received signal and the closest proper modulation state. Good message elements tend to be modulated close to one of the predetermined modulation levels, while faulted message elements tend to have random modulation values, far from the predetermined modulation levels. The receiver can identify faulted message elements (x) according to the deviation between the received modulation value and the closest proper modulation level (black dot).



In addition, the receiver can also use certain properties of the signal waveform. Faulted message elements tend to have a wider variation ("fluctuation") in phase and amplitude, than the good message elements. Faulted message elements tend to have a different x-y polarization angle than the rest of the message. Faulted message elements tend to have a larger frequency deviation Δf relative to the central subcarrier frequency. These fault signatures are already available to most receivers, buried in the digitized branch data, and could be extracted with software. The receiver can then calculate an overall quality factor based on the modulation deviation, the waveform signatures, and other factors, thereby accurately revealing the faulted message element(s) at zero energy cost, and without a retransmission.

After identifying the faults, the receiver can then repair the message in several ways. The receiver can alter the faulted message elements to each proper modulation state, using the embedded error-detection code to find the corrected version. If the message has only one fault, the correct value of that message element can be back-calculated from the error-detection code. If the fault is in the error-detection code itself, and the message otherwise appears fault-free, then the message may be accepted as-is. For even greater effectiveness, the receiver can use an AI model to discern the likely meaning or intent of the message, based on prior unfaulted messages, determining the most likely corrected version using the waveform data and the unfaulted message elements as input, at zero cost.

Because fault mitigation can be performed within the receiver, not dependent on other network systems, automatic fault mitigation is outside the purview of 3GPP. Instead, automatic fault mitigation is an excellent opportunity for an equipment producer to develop proprietary fault mitigation procedures. The producer can thus provide better message reliability and better energy efficiency than competitors, which naturally confers a competitive advantage.

In summary, the receiver can save energy by correcting faults internally, thereby avoiding a retransmission. The receiver can localize the likely faulted message elements by its modulation deviation and numerous waveform characteristics. The receiver can then correct the message by an AI model to determine the meaning or intent of the message, at zero cost by either method. Autonomous fault correction is a business opportunity with enormous profit potential.

Value-chain Analysis - Fault Mitigation in Receiver

For users:

The user can avoid the delays and energy costs involved in retransmissions by correcting faults autonomously, in real-time, transparently to the user. The incremental cost of identifying faults according to waveform signatures, and correcting them by CRC or other means, are negligible.

For networks:

The base station's receivers obtain the same benefits as listed the users. In addition, avoidance of retransmissions greatly simplifies the task of downlink scheduling, a further benefit for the base station.

Conclusions

Wasted energy is not an inevitable consequence of wireless communications, even at the highest frequencies. We propose low-cost options for improving energy efficiency, options that enable simpler, more efficient procedures. More specifically:

1. For beam alignment, a single modulated pulse indicates the alignment direction instantly.
2. Incremental feedback maintains beam parameters economically, even in changing conditions.
3. Base stations can adjust transmission power automatically, using a receptivity map.
4. Users can select the best base station using a network database, without blind searching.
5. Users can also set their communication parameters automatically using the network database.
6. Base stations can include their location in the SSB message, for immediate beamforming.
7. Base stations can provide a series of uniform signals for reception beam optimization.
8. DCI messages can be identified by special demarcations, saving computation energy.
9. Each PDSCH message may also be demarked and labeled, eliminating the need for DCI.
10. Modulation states, blank in one or both branches, enable shorter messages and save power.
11. Amplitude-phase modulation greatly increases phase margins, reducing retransmission costs.
12. Amplitude-phase modulation also enables asymmetric modulation, enabling optimization.
13. Amplitude-phase modulation also reveals faults by type, enabling corrective changes.
14. Faulted message elements can be identified according to the modulation deviation.
15. Faulted message elements can also be identified by their waveform characteristics.
16. Faulted messages can be corrected using AI analysis, based on prior message data.

These energy-saving options reduce unnecessary messaging, avoid costly retransmissions, and enhance beam control. They can be implemented in software, without hardware changes, at zero or near-zero cost. Users will appreciate the improved communication quality, and especially the reduced energy consumption for battery-operated applications, resulting from these energy-saving upgrades.

6G development is proceeding at a rapid pace. Let us not commit to the old energy-wasting legacy procedures, since better options are clearly available. It is important that developers be aware of the energy-saving concepts proposed here, and to include these options in the next-generation plan. This must happen before the 6G specifications proceed further - or we will again be stuck with outdated limitations and inefficiencies. Therefore, it is proposed that all of the energy-saving methods disclosed in this paper should be included in the next 3GPP release.

For more information, please contact:

Kemp Massengill, President
709 Via Del Monte
Palos Verdes Estates, CA 90274 USA
760.390.1410 (pacific time)
kemp.massengill@UltraLogic6G.com

For technical questions, please contact:

David E. Newman, PhD
VP and Chief Scientist
UltraLogic 6G, LLC
760-809-7601 (pacific time)
david.newman@Ultralogic6G.com

Glossary

"Base station", as used herein, includes all network assets communicating with users, including access points, access relay stations, roadside monitors, satellite relays, and the like. The term also includes the core network, backhaul, and other internal systems of the network assets, unless otherwise called out.

"User device", as used herein, refers to the radio portion of user equipment, specifically the transmitter, receiver, antenna, signal processing electronics, and demodulation processor. The term also includes AI models for fault mitigation and message interpretation and the like, when present.

3GPP (Third Generation Partnership Program) is the primary organization for wireless technical specifications, and with seven "Partner" organizations, promulgates universal wireless standards.

OFDM (Orthogonal Frequency-Division Multiplexing) means transmitting message data in multiple frequencies (subcarriers) at the same time. The receiver then measures the subcarrier signals to separate and demodulate the message elements.

IoT (Internet of Things) are low-cost, reduced-capability wireless sensors and actuators.

SNR (Signal-to-Noise Ratio), as used herein, includes interference, stochastic noise, clock drift, and all other effects causing message faults, unless specifically indicated.

FR1 and FR2 are frequency ranges. FR1 is 7.125 GHz and below (and up to 8.4 GHz in 6G). FR2 is 24.25 GHz and up. FR2 is often called mmWave, although a wavelength of 1 mm actually corresponds to a frequency of 300 GHz.

BPSK (binary phase-shift keying) is phase modulation at constant amplitude with 2 states separated by 180 degrees, carrying 1 bit per symbol.

QPSK (quadrature phase-shift keying) is phase modulation at constant amplitude with 4 states separated by 90 degrees, carrying 2 bits per symbol

QAM (Quadrature Amplitude Modulation) is a modulation scheme in which the message data is encoded in the amplitudes of two orthogonal signal components, termed I and Q.

A resource grid is an array of resource elements, arranged by symbol-times in time and subcarriers in frequency.

A message element is a single modulated resource element of a wireless message.

A message is "time-spanning" if the message elements are sequential in time on the same subcarrier, and "frequency-spanning" if the message elements are sequential in frequency at the same symbol-time.

PDSCH and PDCCH represent the downlink shared and control channels by which the base station communicates with each user device.

RACH, Msg3, MsgA and the like refer to various steps of the initial access procedure.

References

[1] The following energy-saving patents and publications can be found at www.UltraLogic6G.com.

<u>US Patent</u>	<u>Title</u>
12,185,297	Direct AI Management of 5G/6G Network Operations
12,166,619	Procedures for Efficiently Defaulting QAM Messages in 5G and 6G
12,155,471	Message Fault Localization and Correction in 5G and 6G
12,149,985	Artificial Intelligence for Optimizing 5G/6G Wireless Network Performance
12,149,397	Multi-Parameter Fault Detection and Mitigation in 5G and 6G Messaging
12,143,189	Low-Complexity Procedure for 5G/6G Beam Alignment
12,126,439	Rapid Fault Localization Based on Waveform and CRC in 5G/6G
12,095,573	Fault Determination by AI Waveform Analysis in 5G and 6G
12,074,741	Identifying Faulted Message Elements by Modulation Consistency in 5G/6G
12,047,220	AI-Based Correction of Corrupted 5G/6G Messages
11,387,935	Error Detection and Correction by Modulation Quality in 5G/6G
11,387,960	Downlink Demarcations for Rapid, Reliable 5G/6G Messaging
11,398,876	Error Detection and Correction in 5G/6G Pulse-Amplitude Modulation
11,405,131	AI-Based Error Detection and Correction in 5G/6G Messaging
11,411,612	Location-Based Beamforming for Rapid 5G and 6G Directional Messaging
11,411,795	Artificial-Intelligence Error Mitigation in 5G/6G Messaging
11,418,372	Low-Complexity Demodulation of 5G and 6G Messages
11,424,787	AI-Based Power Allocation for Efficient 5G/6G Communications
11,438,033	Location-Based Power for High Reliability and Low Latency in 5G/6G
11,438,834	Searchable Database of 5G/6G Network Access Information
11,451,429	Modulation Including Zero-Power States in 5G and 6G
11,509,381	Resource-Efficient Beam Selection in 5G and 6G
11,516,065	Identifying Specific Faults in 5G/6G Messages by Modulation Quality
11,522,636	Modulation Quality and Fault Mitigation in 5G/6G
11,522,637	Selection of Faulted Message Elements by Modulation Quality in 5G/6G
11,522,638	Artificial Intelligence Fault Localization in 5G and 6G Messages
11,522,745	Identification and Mitigation of Message Faults in 5G and 6G Communications
11,523,334	Network Database for Rapid, Low-Complexity 5G/6G Network Access
11,528,178	Zero-Power Modulation for Resource-Efficient 5G/6G Messaging
11,533,084	Automatic Adjustment of Transmission Power for 5G/6G Messaging
11,546,111	Demarking the Start and End of 5G/6G Downlink Messages
11,546,201	Selection of Message Elements based on Modulation Quality in 5G and 6G
11,558,236	Single-Branch Reference for High-Frequency Phase Tracking in 5G and 6G
11,563,515	Fault Recovery by Selection based on Modulation Quality in 5G/6G
11,581,919	Transmission Power Compensation by Attenuation Mapping in 5G and 6G
11,601,150	Demodulation for Phase-Noise Mitigation in 5G and 6G
11,611,375	Location-Based System Information and Doppler Correction in 5G/6G
11,616,679	Detection and Mitigation of 5G/6G Message Faults
11,637,649	Phase-Noise Mitigation at High Frequencies in 5G and 6G
11,644,522	Triangular Beam Configurations for Rapid Beam Alignment in 5G and 6G
11,652,533	Rapid Alignment of User Directional Beams in 5G/6G Networks
11,671,305	Extremely Compact Phase-Tracking 5G/6G Reference Signal
11,689,249	Geographical Localization of 5G/6G Network Users and Base Stations
11,695,612	Method to Locate Faulted Message Elements Using AI in 5G and 6G
11,736,320	Multiplexed Amplitude-Phase Modulation for 5G/6G Noise Mitigation
11,736,332	Enhanced Fault Correction and Noise Avoidance in 5G/6G Networking
11,736,333	Information Content in Zero-Power Modulation States in 5G and 6G

11,770,207	Method for Mitigating Branch-Amplitude Faults in 5G and 6G Messages
11,770,209	Signal Quality Input for Error-Detection Codes in 5G and 6G
11,770,815	Lean Deterministic Beam/Power Feedback During 5G/6G Initial Access
11,777,547	Phase-Tracking Demodulation Reference and Procedure for 5G and 6G
11,777,639	How to Maximize Phase-Noise Margins in 5G and 6G
11,782,119	Phased Beam-Alignment Pulse for Rapid Localization in 5G and 6G
11,784,764	Artificial Intelligence for Fault Localization and Mitigation in 5G/6G
11,785,591	Multiplexed Code for ACK/SR/Power/Beam Feedback in 5G and 6G
11,799,585	Error Correction in 5G and 6G using AI-Based Analog-Digital Correlations
11,799,608	Low-Complexity Method for Identifying Downlink Messages in 5G and 6G
11,805,506	Concise Feedback for Downlink Beam and Power Adjustment in 5G and 6G
11,811,565	Demodulation Using Two Modulation Schemes in 5G and 6G
11,811,579	Recovery of Corrupted 5G/6G Messages by Modulation Quality
11,812,421	AI-Managed Channel Quality Feedback in 5G/6G
11,817,950	AI Means for Mitigating Faulted Message Elements in 5G/6G
11,824,667	Waveform Indicators for Fault Localization in 5G and 6G Messages
11,832,118	Beam Adjustment by Incremental Feedback in 5G and 6G
11,832,128	Fault Detection and Mitigation Based on Fault Types in 5G/6G
11,843,468	Fault Detection, Localization, and Correction by 5G/6G Signal Quality
11,848,774	AI-Based Analog-Digital Fault Detection and Localization in 5G/6G
11,848,788	AI-Based Waveform Analysis for Fault Localization in 5G and 6G
11,849,349	Fault Mitigation Using Signal Quality and Error-Detection Codes in 5G/6G
11,871,254	Incremental Realtime Signal-Quality Feedback in 5G/6G
11,996,971	Enhanced Throughput and Reliability with Zero-Power States in 5G and 6G
11,996,973	Scheduling Single-Branch Phase-Tracking References in 5G and 6G
12,003,323	AI Model with Error-Detection Code for Fault Correction in 5G/6G
12,021,614	5G/6G Network Operations with AI-Based Message Fault Correction
12,028,158	Low-Complexity Fault Mitigation in 5G/6G Messages
12,034,571	Modulation and Demodulation for Enhanced Noise Margins in 5G and 6G
12,038,519	Low-Complexity Beam Alignment by Directional Phase in 5G and 6G
12,040,891	How to Maximize Throughput and Phase Margin in 5G/6G Communications
12,047,219	Fault Detection and Correction by Sum-Signal Modulation in 5G or 6G
12,052,129	Ultra-Compact Phase-Tracking Demodulation Reference for 5G/6G
12,057,936	Fault Correction Based on Meaning or Intent of 5G/6G Messages
12,069,635	Resource-Efficient Low-Complexity Beamforming Feedback in 5G/6G
2023/0089694	Message Fault Recovery Without Retransmission in 5G and 6G
2023/0100826	Throughput Enhancement by Location-Based Power Adjustment in 5G and 6G
2023/0179261	Automatic Condition-Based Adjustment of Transmission Power in 5G and 6G
2023/0262595	Automatic Base Station Discovery, Selection, and Registration in 5G/6G
2023/0283322	Beamforming Communication Method for Wireless Devices in 5G and 6G
2023/0387952	Selecting a Modulation Scheme to Mitigate Specific Fault Types in 5G and 6G
2024/0015721	Resource-Efficient Low-Complexity Beamforming Feedback in 5G/6G
2024/0045013	Deterministic Low-Complexity Beam Alignment for 5G and 6G Users
2024/0080144	Enhanced Reliability by Waveform Analysis in 5G/6G Communications
2024/0097852	Eliminating Unnecessary Downlink Control Messages in 5G and 6G
2024/0154679	Test Signals and Feedback for Beam Optimization in 5G and 6G