

WIRELESS SYNCHRONIZATION AND TIMING

Synchronization Procedures for User Devices and Base Stations

David E. Newman and R. Kemp Massengill

Motivation

Wireless messages in 5G-Advanced, and especially 6G, face new, critical timing requirements. Legacy synchronization procedures were developed for early, low-performance systems, and are not ideal for the fast-cadence, high-throughput demands of the next generation in wireless networking. Leaner, simpler, faster synchronization procedures are needed.

Executive Summary

Every modulated symbol of every uplink message must arrive at the base station within strict symbol-times and subcarrier frequencies. Any deviation results in message faulting, with inevitable latency and reliability issues. On downlink, user devices must adjust their local clocks to maintain coherence on reception, but this is not enough. Each user device must also adjust its timing advance to compensate for the critical uplink propagation delay, and mobile users must correct for the Doppler frequency shift, to match the base station's resource grid. At 6G frequencies, even a slight time/frequency deviation will cause unacceptable inter-symbol interference, subcarrier crosstalk, and phase noise.

Legacy procedures for clock alignment, propagation compensation, and frequency correction were developed back when demand was light and frequencies were modest. Message-heavy and error-prone procedures were satisfactory then, but are poorly matched to the high expectations of 6G. Improved lean means are needed to enable user devices to adjust both timing and frequency, in real-time, for both uplink and downlink, with minimal consumption of resources and energy.

Therefore, disclosed below are new timing and synchronization procedures suitable for next-generation networking. Sharp timing is provided by timestamp waveforms that include a readily detectable transition point. Timestamp signals enable nearly instantaneous synchronization of all users in a network, yet cost just a few resource elements at certain scheduled times. Versions can also provide time marks in the guard-space of each OFDM symbol, for further resource economy. In addition, simplified versions provide satisfactory alignment based on amplitude measurements alone, which even the most basic receivers can readily perform. New ultra-lean timestamp exchange procedures provide full synchronization based on just four or five brief signal pulses, with minimal resource and power consumption. Further procedures accommodate mobile users, including high-Doppler applications with unlimited uplink-downlink propagation asymmetry. These procedures enable each user device to automatically ensure compliance with tight base station requirements in time and frequency, while optimizing reception reliability at the user device, with very minimal energy consumption.

Public acceptance of next-generation wireless communications will depend on how well the timing and synchronization problems are solved. All synchronization procedures must be transparent to the user, accessible to IoT and reduced-capability devices, low in power and resource cost, while remaining stable even in the most demanding urban network environments. The timing and synchronization procedures described herein are designed to provide those solutions.

Mid-Symbol Timestamp Point

The base station can initiate synchronization by transmitting, at a pre-scheduled time and frequency, a "mid-symbol timestamp point" signal. The transmission can be unicast to a particular user device, or broadcast to the entire network simultaneously. Although extremely brief, the timestamp point conveys the specific timing information needed by user devices.

Figure 1 shows an example of a mid-symbol timestamp point based on a sudden phase change (central arrow), for an "idealized" waveform. For a real transmitter, the timestamp point is the moment of maximum rate of change of phase. Most base station transmitters can handle this. Also plotted is the transmitted phase versus time, with oscillations suppressed. The receiver can determine the time of phase reversal, relative to the receiver's symbol boundaries, thereby revealing any timing offset relative to the base station's clock. The user device can then correct its local clock according to the offset detected.




Fig. 1: A phase-shift timestamp point.

The time resolution achievable in a single measurement depends on the bandwidth allocated. For enhanced precision, the timestamp point can be surrounded by blank subcarriers, up to a complete OFDM symbol if traffic permits. The receiver can also apply techniques to enhance the measurement precision. For example, the phase change can be measured both forward and backward in the data and averaged, canceling many errors. Algorithmic corrections, based on the allocated bandwidth, can be applied to the data before or after analysis. In addition, the receiver can slide a "fitting template" along the time axis and determine where the best fit is obtained. The template shown in the inset is shaped like the timestamp point itself. The fitting template therefore has a very high quality-of-fit only at the timestamp point, and a very low quality-of-fit everywhere else. This is indicated by the "fit" peak.

Figure 2 shows an alternative waveform that may be easier for the transmitter to produce. Here the transmitted wave reverses phase by imposing a slight hold (1/2 wavelength) at a waveform crest, instead of the sharp reversal of the previous example. This version results in a wider transition, as shown in the phase plot, and lower bandwidth, but is easier for both transmitter and receiver to process. An appropriate fitting template is also shown, along with the (slightly) wider quality-of-fit peak. The achievable timing precision still corresponds to a small fraction of the symbol-time, sufficient for next-generation networking.




Fig. 2: Alternative phase shift at mid-symbol.

To demodulate QAM signals, the receiver separates a received waveform into orthogonal I and Q "branch" components. Figure 3 shows a timestamp point configured for QAM modulation, in which the branch signals are suddenly reversed in the middle of a symbol-time. Such a branch reversal requires a little extra transmitter power, because the transition is a 90-degree phase change, as opposed to the 180-degree phase changes in Figs. 1 and 2. The receiver then separates the two branches and detects the time of the branch reversal. The figure also shows the branch amplitudes versus time, with the oscillations suppressed. Due to various bandwidth and hysteresis effects, the receiver will take a certain amount of time to recognize that the I branch has dropped to zero, and for the Q branch to reach full amplitude. The receiver can then optimize the time resolution according to an algorithm "Transition" based on the I-Q sum and difference, as shown.

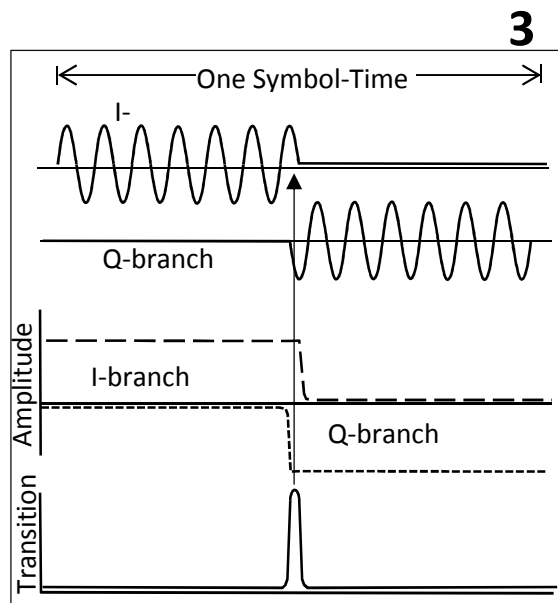


Fig. 3: QAM branch-alternation timestamp.

Figure 4 shows several phase-change timestamp points at various positions in a resource grid. The phase reversal is indicated by light and dark shading. Additional precision can be obtained by providing one or more blank subcarriers (diagonal hash) above and below the subcarrier of the timestamp point, thereby multiplying the perceived bandwidth and accommodating much faster transitions. Alternatively, multiple timestamp points, spaced apart in frequency as shown, can emulate a larger bandwidth, and therefore sharper time resolution at negligible cost. In addition, the timestamp point may be made longer by occupying multiple symbol-times (3 shown), and surrounded by blank subcarriers (2 blanks shown on each side) thereby providing enhanced SNR. For even greater bandwidth, the timestamp point can occupy an otherwise blank OFDM symbol at a scheduled time.




Fig. 4: Resource grid with timestamps.

For further SNR, the transmitted power can be boosted during the timestamp signal, compensating for the lack of power in the blank resource elements. The timing of the timestamp point, relative to the receiver's symbol boundaries, indicates the timing offset of the receiver clock relative to the base station clock. The transmitter transmits the timestamp point precisely in the center of the transmitter's symbol-time. If the receiver's clock is aligned with the base station, the timestamp point will appear centered in the receiver's symbol-time as well. If the receiver's clock is mis-aligned, the timestamp position, as received, is displaced from the receiver's symbol center. The timing offset equals that displacement. The receiver can then adjust its clock setting according to the timing offset, thereby resynchronizing with the base station. In addition, by setting the receiver's clock equal to the scheduled transmission time of the timestamp point, the receiver can automatically compensate the downlink propagation time. This enables the receiver to properly receive downlink messages from the base station.

Guard-Space Phase-Tracking

Each modulated message element of a message is transmitted in one symbol-time on one subcarrier of a resource grid. Each resource element is divided into a guard-space interval followed by the message data. The guard-space prevents inter-symbol interference and accommodates users at different distances from the base station. Conventionally, the guard-space contains a cyclic prefix, which is a copy of the final portion of the message data. The cyclic prefix simplifies certain signal processing procedures, but is not essential.

Figure 5 shows a timestamp point inserted into the guard-space of a single resource element. The first half of the guard-space is a copy of the first half of the final portion data, and the second half is the rest of the final portion but inverted in phase. The receiver can determine the exact moment when the guard-space signal suddenly differs from the final portion data, revealing any timing offset. To measure the timestamp time, the receiver can simply add the guard-space values to the final portion data. This yields zero during the phase-inverted half, and non-zero values in the other half of the guard-space, thereby indicating the timestamp point, and the timing offset, precisely.

Figure 6 shows a phase-inversion timestamp inserted into the guard-space of an OFDM symbol, which has a much larger bandwidth. An OFDM symbol is a superposition of multiple subcarrier signals within the same symbol-time. The modified cyclic prefix is a copy of the final portion of the message data, with a sudden phase reversal indicating the exact center of the guard-space, as transmitted. The receiver determines the timing offset by the position of the phase reversal, as received in the OFDM symbol guard-space.

Figure 7 shows how a receiver can determine the exact timestamp position in the receiver's guard-space. The receiver first digitizes the received signal, and then adds or subtracts the final portion of the message data to/from the guard-space signal. Half of the resulting points will be zero due to the phase inversion, and the other half will be random numbers but not zero. The timestamp is the moment when the sum and difference data switch from zero to non-zero amplitude.

If desired, the receiver can reconstruct the original cyclic prefix by re-inverting the inverted half of the guard-space data.




Fig. 5: Timestamp point in the guard-space of a single subcarrier.

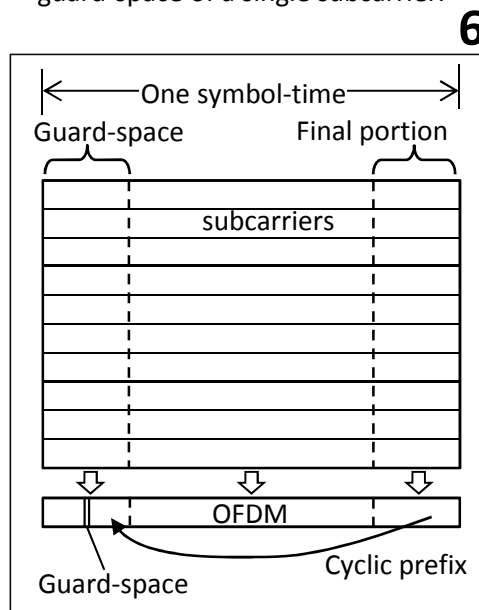


Fig. 6: Timestamp point in guard-space of an OFDM symbol.




Fig. 7: Sum/Difference reveals exact position of timestamp point within the guard-space.

Amplitude-Only Synchronization

Some receivers, such as low-cost IoT receivers, may be unable to do the timing measurements or signal processing involved in the previous examples. For them, a different synchronization procedure is available, involving only amplitude measurements that all 3GPP-compliant receivers can do natively. The transmitter generates a signal, as uniform as possible, filling one symbol-time at a pre-scheduled time, and also provides 1 or 2 blank symbol-times before and after the signal. The signal may be in a single subcarrier or a broadband OFDM block, depending on the resolution required. The receiver measures the received energy (or time-averaged amplitude) in all three symbol-times, as-received. If the user clock is mis-aligned, the user's symbol boundaries will be displaced relative to the base station's boundaries, causing some of the transmitted signal to be received in one of the two outer symbol-times. The receiver calculates the time offset according to a ratio of energy appearing in the outer symbol-times divided by the central symbol-time. The user device thereby determines a precise time offset without actually making any timing measurements, just amplitude measurements.

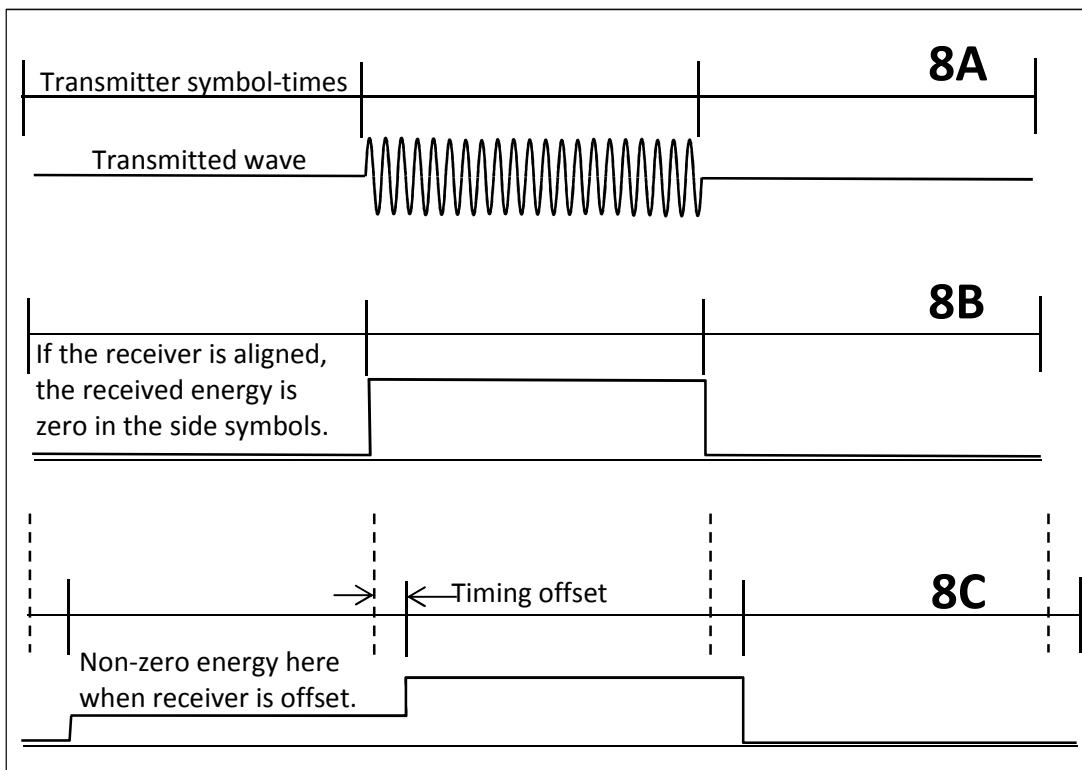


Fig. 8A: Transmitter produces signal filling one symbol-time according to transmitter's clock. 8B: Here the receiver is perfectly aligned, so none of the energy appears in the side symbols. 8C: Here the receiver is offset, causing spill-over energy into one side symbol.

Figure 8A shows a signal transmitted by a base station within one symbol-time, as determined by the base station clock, with 1 or 2 blank symbol-times before and after. Figure 8B shows the same signal received by a receiver that is perfectly aligned with the base station. All the energy is received within the central symbol-time, and nothing is received in the adjacent two symbol-times. Figure 8C shows what happens when the receiver clock is offset relative to the base station. Some of the energy overlaps one of the side intervals. The time offset is determined by the ratio of amplitudes received in the symbols.

For extra sharpness, the transmitted pulse can be shaped (i.e. edge-emphasized, with extra power at the edges) to provide fast turn-on and turn-off, compensating for the finite bandwidth.

Full Synchronization for Non-Mobile Users

Figure 9 shows how a base station can provide full synchronization service to all of its user devices in a compact, resource-conserving procedure. The base station is on the left and the user device on the right, with signals shown delayed by the uplink and downlink propagation times. With just four brief signals containing timestamp points, the user device determines its correct timing settings for both downlink and uplink, plus the frequency, all compliant with the base station's subcarriers and symbol boundaries.

At a previously scheduled time T_{30} , the base station transmits a timestamp signal, which the user device receives at T_{31} , delayed by the downlink propagation time ΔT_{down} . The user device sets its clock equal to the scheduled transmission time, T_{30} . This ensures that the downlink propagation time is automatically compensated in subsequent downlink receptions, synchronous with the user device's symbol boundaries.

After a scheduled delay D_{freq} , the base station sends a second timestamp signal at T_{32} , which the user device receives at T_{33} . The user device calculates a frequency error proportional to D_{freq} minus the interval $T_{33}-T_{31}$, and corrects its clock rate (frequency) to agree with the base station's clock rate.

Then at T_{34} , the user device transmits an uplink timestamp point, which the base station receives at T_{35} , delayed by the uplink propagation time ΔT_{up} . After a predetermined delay D_{reply} , the base station transmits another downlink timestamp point at T_{36} , which the user receives at T_{37} . The user then calculates the timing advance equal to the interval $T_{37}-T_{34}$ minus D_{reply} , which equals the round-trip propagation time $\Delta T_{up}+\Delta T_{down}$. The user device then applies this timing advance to all uplink messages (relative to the user's reception symbol boundaries). As mentioned, the user's clock was set behind the base station's clock by ΔT_{down} in the first step, so the timing advance actually causes the uplink signals to be transmitted early by ΔT_{up} , the uplink propagation delay. The uplink message then arrives at the base station exactly centered in the base station's symbol-times, as required. The round-trip timing advance automatically compensates both propagation delays for the user and the base station.

In summary, the user device sets its clock so that downlink messages arrive at the receiver aligned within the user's (adjusted) symbol boundaries despite the downlink propagation delay, and then sets its clock rate or frequency to agree with the base station's frequency. The user device sets the timing advance so that uplink signals arrive at the base station synchronized with the base station's symbol boundaries automatically, and the uplink signals remain centered in the base station's subcarriers, thereby providing optimal communication both ways. This procedure is far simpler and less resource-hungry than legacy synchronization procedures that require back-and-forth messaging and complex measurements. The new auto-synchronization procedure is brief and economical, compatible with reduced-capability devices, independent of any propagation asymmetries, not dependent on satellite signals or other external standards, and not dependent on costly precision clocks.

9

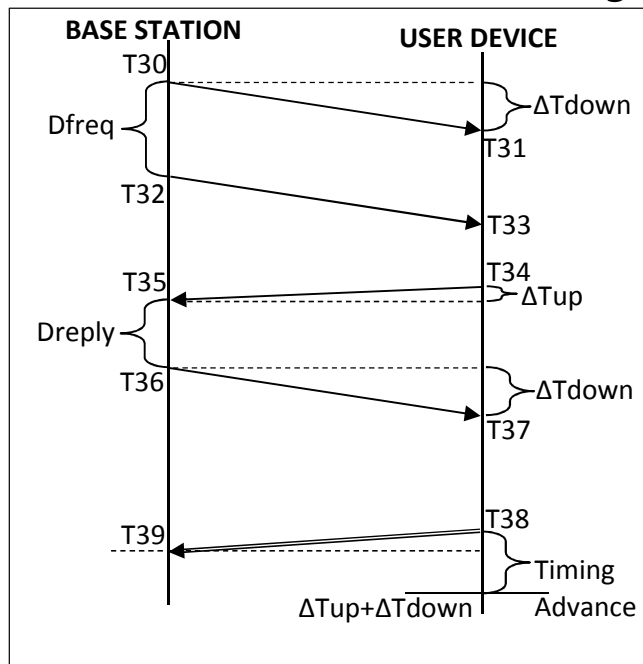


Fig. 9: Base station aligns user device using only timestamp points at scheduled times.

Full Synchronization for Mobile Users

Figure 10A shows a synchronization procedure for a mobile user (tilted trajectory line). As in the previous example, the base station transmits two scheduled timestamp signals at T40 and T42, at frequency f_0 , spaced apart by the known interval D_{freq} . They are received by the moving user at T41 and T43, at frequency $f_0 + DS$. (DS is the Doppler shift.) The mobile user sets its clock time according to the scheduled T40 time, and initially matches the received frequency of $f_0 + DS$. Then, at T44, the user transmits an uplink timing signal using the initial frequency $f_0 + DS$. At T45, the base station receives the signal, double-Doppler-shifted to $f_0 + 2DS$. After a known delay D_{reply} , the base station transmits a downlink timestamp point at T46, but using a special frequency $f_0 - 2DS$ (for this signal only). The user receives it at T47, but now Doppler shifted to $f_0 - DS$. The user device then calculates the correct Doppler shift, equal to half the difference between its transmitted frequency at T44 and the received frequency at T47. The user also applies the timing advance TA, equal to the time interval $T47 - T44 - D_{reply}$, (relative to its own symbol boundaries), to match the base station's symbol boundaries. The user then sends all further uplink messages using the Doppler-corrected frequency $f_0 - DS$ and the time advance TA, arriving perfectly aligned with the base station's symbol boundaries and subcarrier frequency f_0 , upon reception.

Figure 10B shows an alternative procedure that does not require the base station to transmit at a custom frequency. Instead, the base station informs the user device of the Doppler shift based on timing. The user transmits the uplink signal at T54 using the frequency it has received, $f_0 + DS$. The base station transmits the Dreply signal at T56, then waits a variable delay time D_{shift} which is related to DS by a formula, then sends another timestamp signal to the user. The user calculates DS from the interval $T59 - T57$ using the formula, and applies the Doppler-corrected frequency $f_0 - DS$ for all uplink messages thereafter.

10A




Fig. 10A: Base station aligns a mobile user device using only timestamp points at scheduled times.

10B

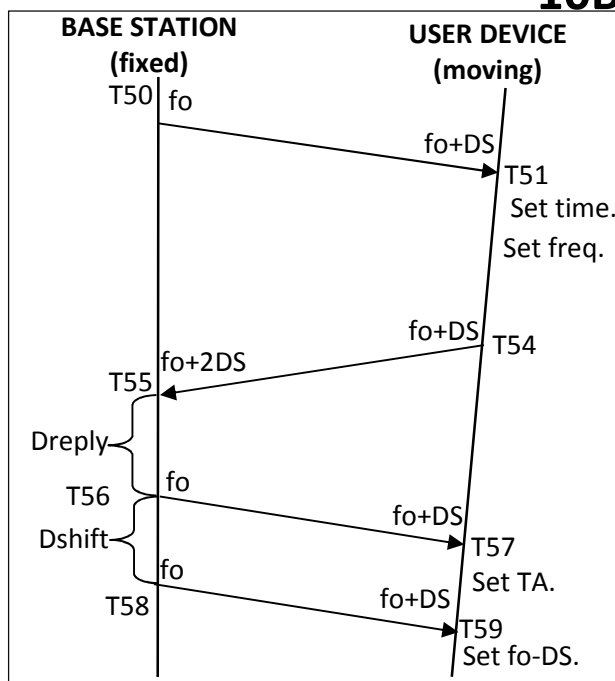


Fig. 10B: Alternative method to inform user of Doppler shift using timing instead of frequency.

User-Centric Grid Timing: Non-Mobile User

The base station of a wireless network must communicate simultaneously with multiple user devices. This is feasible only if all of the user devices precisely adjust their timing so that their uplink messages arrive at the base station within the base station's symbol boundaries, and with frequencies centered in the base station's subcarriers. The user device, on the other hand, receives messages from only one entity, the base station. Therefore the user has the flexibility to accommodate time-retarded and Doppler-shifted messages on the downlink, as long as it arranges its uplink messages to fit the base station's resource grid.

Figures 11A-B-C show how the user device can accommodate the base station's requirements of uplink synchronization and syntonization, while still receiving downlink messages coherently. In Figure 11A, the base station's resource grid is shown, with symbol-times horizontally and subcarriers vertically. One block-slot is depicted.

Figure 11B shows the user device's resource grid for receiving downlink messages. Since downlink messages are delayed by the downlink propagation time, the user device delays its reception grid by the same amount, ΔT_{down} , the timing retard for reception. The user device does this by setting its clock according to the scheduled time of the base station's timestamp points, as described in Fig. 9.

Figure 11C shows the user device's grid for transmitting uplink messages to the base station. Uplink messages must arrive coincident with the base station's symbol boundaries. Therefore the user applies a timing advance of $\Delta T_{up} + \Delta T_{down}$ (the round-trip travel time), relative to the user's reception grid. Since the user's reception grid is already delayed by ΔT_{down} , this timing advance causes the uplink messages to be transmitted early by ΔT_{up} . After propagating to the base station, the message therefore arrives at the base station exactly aligned with the base station's grid, as required.

When configured as shown, the user device can receive the downlink messages coherently, and the base station can receive the user's uplink messages coherently, without further timing complexity or corrections. The synchronization procedure of Figure 9 provides the user device with reception and transmission grids (or the corresponding symbol boundaries) properly timed for both uplink and downlink communications, compliant with the base station's resource grid.




Fig. 11ABC: User device retards its reception grid to account for the downlink propagation time, and advances its transmission grid so that the uplink messages will arrive at the base station coincident with the base station's grid.

Figure 12 shows how a mobile user device can arrange its reception and transmission grids in frequency, to cancel the Doppler shift when communicating with the base station. (The timing advance is assumed to be already taken care of by the procedure of Fig. 11.) The user device initially sets its frequency to that received from the base station, which is Doppler shifted $+1 DS$ for downlink messages. Accordingly, the user device's reception grid is frequency adjusted by $+1DS$ to compensate. (In this example, motion toward the base station is assumed.) The user device can then receive messages from the base station without further frequency adjustment, since the Doppler shift compensation is already built-in.




Fig. 12: User sets own frequency $+1 DS$ Doppler shifted relative to base station, for ease of reception. For uplink transmission, user applies a $-2DS$ reverse Doppler shift relative to own clock, so that after uplink $+1DS$ Doppler shift, message will arrive at the base station's frequency.

For uplink, on the other hand, the user device must arrange for its transmissions to arrive at the base station without a Doppler shift, in order to match the base station's subcarrier frequencies. The user device first measures DS by running the synchronization procedure shown in Fig. 10, and adjusts the transmission grid down by $-2DS$ (relative to the reception grid). Since the reception grid already has $+1DS$ frequency shift, the $-2DS$ downward adjustment of the transmission grid causes the uplink messages to be transmitted at a frequency shift of $-1DS$, relative to the base station. Then, after propagating to the base station, it is again Doppler shifted up to f_0 , which is the subcarrier frequency that the base station expects. When configured as directed, the user device's reception grid is arranged for downlink message reception transparently, and the transmission grid is arranged for uplink communications compliant with base station subcarrier requirements.

It is important to note that the entire synchronization procedure involved only four or five brief timestamp points, transmitted at pre-scheduled times which the user device knows. This is far simpler and more resource-efficient than legacy procedures involving bulky messages back and forth between the user device and the base station. Another nice feature is that the procedure is completely independent of the uplink/downlink asymmetry. Most of the old-fashioned synchronization procedures depend on perfect propagation symmetry, which is questionable at FR2 and higher. These lean and efficient procedures avoid the waste and complexity of prior-art timing procedures, while fully matching the base station requirements, while serving both fixed and mobile users transparently, and enabling all of the users in a network - including low-performance IoT devices - to precisely synchronize with the base station, all at very low cost.

Conclusion

Synchronization between user devices and base stations will become increasingly crucial as the tight tolerances of 6G begin to take effect. Legacy synchronization procedures are not ideal for next-generation networking, due to resource consumption and unnecessary complexity. The procedures described herein provide resource-efficient, rapid, precise time and frequency alignment, using measurements that even basic IoT and low-cost sensor/actuator devices perform natively.

Wireless planning committees and standards organizations, such as 3GPP, should consider the proposed synchronization procedures as advantageous alternatives to the bulky timing procedures of previous releases. The lean procedures outlined herein are well-suited to next-generation networking goals. With prompt action, the new synchronization procedures may be available to all network users before 5G-Advanced is fully rolled out. Network users will appreciate the improved speed and reliability.

Glossary

3GPP (Third Generation Partnership Program): the primary wireless standards organization.
OFDM (Orthogonal Frequency-Division Multiplexing): transmit multiple frequencies at the same time.
IoT (Internet of Things): low-cost wireless sensors and actuators.
SNR (Signal-to-Noise Ratio): as used herein, includes interference and clock drift.
FR1 and FR2 are frequency ranges. FR1 is 7.125 GHz and below. FR2 is 24.25 GHz and up.
QAM (Quadrature Amplitude Modulation): data encoded in two orthogonal signal components.
Resource grid: a schedule of symbol-times and subcarrier frequencies for message transmission.
Message element: a modulated resource element of a wireless message.
Guard-space: a brief interval separating the symbol-times of a wireless message.
Cyclic prefix: a copy of the final portion of a message element into the guard-space.
Synchronization: setting a clock equal to another clock.
Syntonization: setting a clock rate equal to another clock rate.

References

[1] The following synchronization and timing patents can be found at www.UltraLogic6G.com.

<u>US Patent/Publication</u>	<u>Title</u>
11,722,980	Guard-Space Timestamp Point for Precision Synchronization in 5G and 6G
11,737,044	Mid-Symbol Timestamp Point for Precision Synchronization in 5G and 6G
11,800,480	Ultra-Lean Timing Signal for Precision Synchronization in 5G and 6G
11,805,491	Compact Timing Signal for Low-Complexity 5G/6G Synchronization
11,832,204	Ultra-Lean Synchronization Procedure for 5G and 6G Networking
11,950,197	Precision Timing for Low-Complexity User Devices in 5G/6G Networks
11,956,746	Precision Synchronization Using Amplitude Measurements in 5G and 6G
2024/0031968	Rapid Low-Complexity Synchronization and Doppler Correction in 5G/6G
2024/0196354	Phase-Shift Guard-Space Timestamp Point for 5G/6G Synchronization
2024/0205857	Fast, Resource-Efficient Timestamp Generation and Measurement in 5G/6G
2024/0214961	Simultaneous Timing Synchronization of User Devices in a 5G/6H Wireless Network