

# OFDM, the Cyclic Prefix & 5G Numerology

*The algebra behind "one division per subcarrier," and the numbers behind numerology.*

---

## 1. The IDFT is the modulator

Place one data symbol  $X[k]$  (e.g. a QAM point) on each of  $N$  subcarriers and take the inverse DFT:

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} X[k] e^{j2\pi kn/N}, \quad n = 0, \dots, N-1.$$

The receiver takes the forward DFT,  $X[k] = \sum_n x[n] e^{-j2\pi kn/N}$ , and reads one symbol per bin. In continuous time the subcarriers are  $e^{j2\pi k \Delta f t}$  with spacing  $\Delta f = 1/T_u$ , where  $T_u$  is the useful (FFT) symbol length. They are orthogonal over one symbol:

$$\frac{1}{T_u} \int_0^{T_u} e^{j2\pi kt/T_u} e^{-j2\pi mt/T_u} dt = \delta_{km}.$$

Each subcarrier completes an integer number of cycles in  $T_u$ , so every other subcarrier integrates to zero. Their spectra overlap (sinc-shaped), yet each spectrum is zero at every other subcarrier's center.

## 2. Why the prefix must be cyclic: diagonalization

Let the channel be an FIR filter  $h[0], \dots, h[L_h - 1]$ . Prepend a cyclic prefix of length  $L$ , copying the last  $L$  samples:

$$\tilde{x} = [x[N-L], \dots, x[N-1], x[0], \dots, x[N-1]].$$

If  $L \geq L_h - 1$ , then after the receiver discards the first  $L$  samples, the remaining  $N$  samples are

$$y[n] = \sum_{l=0}^{L_h-1} h[l] x[(n-l) \bmod N] + w[n] = (h \circledast x)[n] + w[n],$$

a **circular** convolution. Wherever  $n-l$  would reach into the previous block, the prefix has supplied exactly the wrapped-around sample  $x[(n-l) \bmod N]$ . By the circular-convolution theorem, the DFT turns this into a product:

$$\boxed{Y[k] = H[k] X[k] + W[k]}, \quad H[k] = \sum_{l=0}^{L_h-1} h[l] e^{-j2\pi kl/N}.$$

## Matrix view

After CP removal,  $\mathbf{y} = \mathbf{H}_c \mathbf{x} + \mathbf{w}$  with  $\mathbf{H}_c$  **circulant**. Every circulant matrix is diagonalized by the DFT:  $\mathbf{H}_c = \mathbf{F}^H \mathbf{\Lambda} \mathbf{F}$ , where  $\mathbf{F}$  is the unitary DFT matrix and  $\mathbf{\Lambda} = \text{diag}(H[0], \dots, H[N-1])$ . A zero-padded (blank) guard also removes ISI, but leaves a *Toeplitz* matrix, which the DFT does not diagonalize. The cyclic content is what makes the channel diagonal.

**One-tap equalizers.** Zero-forcing:  $\hat{X}[k] = Y[k]/H[k]$ . MMSE, which avoids amplifying noise in spectral nulls:

$$\hat{X}[k] = \frac{H^*[k]}{|H[k]|^2 + \sigma_w^2/E_s} Y[k].$$

## 3. What the prefix costs

Time efficiency:  $\eta = T_u/(T_u + T_{cp}) = N/(N + L)$ . Energy is lost in the same ratio, since the prefix carries power but no new data.

Peak-to-average power: with many independent subcarriers the time signal is approximately complex Gaussian, so rare large peaks occur. The worst case is  $N$  (i.e.  $10 \log_{10} N$  dB). In practice, PAPR exceeded with probability  $10^{-4}$  is roughly 10–12 dB for large  $N$ , which forces power-amplifier back-off. DFT-spread OFDM (an extra DFT before the subcarrier mapping) restores single-carrier-like peaks on the 5G uplink.

## 4. 5G NR numerology

Subcarrier spacing  $\Delta f = 15 \cdot 2^\mu$  kHz. A slot is 14 symbols (normal CP), and there are  $2^\mu$  slots per 1 ms subframe. Normal-CP length scales as  $2^{-\mu}$ . The figures below are for most symbols; the first symbol of each 0.5 ms half-subframe gets a slightly longer CP.

| $\mu$ | $\Delta f$ | $T_u = 1/\Delta f$ | Normal CP    | Slot         | Typical use                          |
|-------|------------|--------------------|--------------|--------------|--------------------------------------|
| 0     | 15 kHz     | 66.7 $\mu$ s       | 4.69 $\mu$ s | 1 ms         | FR1, wide-area, LTE-like             |
| 1     | 30 kHz     | 33.3 $\mu$ s       | 2.34 $\mu$ s | 0.5 ms       | FR1 mid-band (e.g. 3.5 GHz)          |
| 2     | 60 kHz     | 16.7 $\mu$ s       | 1.17 $\mu$ s | 0.25 ms      | FR1 or FR2; extended CP 4.17 $\mu$ s |
| 3     | 120 kHz    | 8.33 $\mu$ s       | 0.59 $\mu$ s | 0.125 ms     | FR2 mmWave data                      |
| 4     | 240 kHz    | 4.17 $\mu$ s       | 0.29 $\mu$ s | 62.5 $\mu$ s | FR2 synchronization (SSB)            |

Release 17 adds 480 and 960 kHz for 52.6–71 GHz (FR2-2). Normal-CP overhead is about 6.6% at every  $\mu$ : the CP and the symbol scale together, so the fraction stays constant while the absolute delay-spread tolerance shrinks.

## 5. The trade-off numerology exposes

- **Delay spread** must fit inside the CP, which favors small  $\Delta f$  (long symbols).

- **Frequency error, Doppler, and phase noise** cause inter-carrier interference that grows with the normalized offset  $\varepsilon = f_{\text{off}}/\Delta f$ . For a small carrier-frequency offset, ICI power is approximately  $(\pi\varepsilon)^2/3$  relative to the signal. This favors large  $\Delta f$ .
- **Latency** scales with slot length, which favors large  $\Delta f$ .

### Illustration: Doppler at mmWave

At  $f_c = 28$  GHz and  $v = 30$  m/s (108 km/h):  $f_D = v f_c / c = 2.8$  kHz. Treating that shift like a frequency offset gives a rough guide: at 15 kHz,  $\varepsilon = 0.19$  and  $\text{ICI} \approx -9$  dB, which is unusable. At 120 kHz,  $\varepsilon = 0.023$  and  $\text{ICI} \approx -28$  dB, which is fine. That is why FR2 uses wide spacing.

## 6. Self-test, worked

### Exercise 1: pick a numerology from a delay spread

An outdoor cell has a *maximum excess delay* of about 2  $\mu\text{s}$ . The CP must cover the longest significant echo, not the RMS delay spread, which is usually several times smaller. From the table:

- 15 kHz: CP 4.69  $\mu\text{s} \geq 2$   $\mu\text{s}$ , comfortable margin.
- 30 kHz: CP 2.34  $\mu\text{s} \geq 2$   $\mu\text{s}$ , works with little margin.
- 60 kHz (normal CP) 1.17  $\mu\text{s}$  and 120 kHz 0.59  $\mu\text{s}$ : too short, so ISI returns.

Overhead in every case:  $4.69 / (66.7 + 4.69) \approx 6.6\%$ . Rather than paying more overhead, NR changes the absolute CP by changing  $\mu$ . The one exception is the extended CP at 60 kHz (4.17  $\mu\text{s}$ , about 20% overhead), which buys long-delay tolerance at a wide spacing.

### Exercise 2: why cyclic, not blank

A blank guard of length  $\geq L_h - 1$  also stops the previous symbol's echoes from landing in the FFT window, so ISI is gone. But inside the window the received block is then a *linear* convolution truncated to  $N$  samples. Its DFT is not  $H[k]X[k]$ : energy smears between subcarriers (the matrix is Toeplitz, not circulant). The cyclic copy makes the block look periodic to the channel, so the result is a circular convolution, which the DFT diagonalizes into one complex gain per subcarrier and one-tap equalization.