ADSP - Exercise 3: What Does the DFT Compute When We Convolve?

Learning objectives

After this exercise, you should be able to

- state the DFT and IDFT and interpret the index μ as a frequency-grid index,
- use the convolution theorem to explain multiplication in the DFT domain,
- distinguish circular from linear convolution,
- derive the condition for avoiding time-domain aliasing by zero-padding, both from the indices and from sampling the spectrum, and
- explain the DFT as samples of the DTFT of a finite sequence.

Guiding question: What does the multiplication of two DFTs followed by an IDFT compute, and when is the result the linear convolution?

Notation (lecture notes, Part 4, slides 3-10): $v(n)$ and $h(n)$ are finite sequences of lengths M_1 and M_2 that begin at $n = 0$; M is the DFT length; κ is the summation index of a convolution; $\circledast$ denotes circular convolution. For

$$M \geq \max(M_1, M_2),$$

a tilde marks the sequences extended with zeros to length M:

$$\tilde{v}(n) = \begin{cases} v(n), & 0 \leq n < M_1, \\ 0, & M_1 \leq n < M, \end{cases} \qquad \tilde{h}(n) = \begin{cases} h(n), & 0 \leq n < M_2, \\ 0, & M_2 \leq n < M. \end{cases}$$

Thus $\tilde{v}(n)$ and $\tilde{h}(n)$ contain the complete original sequences followed by zeros; neither sequence is truncated. The lecture usually leaves this zero-padding implicit and writes $V_M(\mu)$; the tilde makes it explicit. The M-point DFT and IDFT are

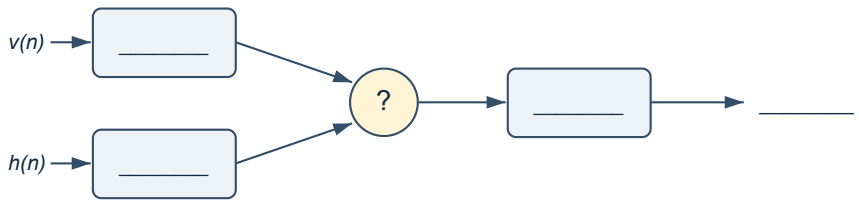
$$\tilde{V}_M(\mu) = \sum_{n=0}^{M-1} \tilde{v}(n) e^{-j\frac{2\pi}{M}\mu n}, \qquad \tilde{v}(n) = \frac{1}{M} \sum_{\mu=0}^{M-1} \tilde{V}_M(\mu) e^{j\frac{2\pi}{M}\mu n}, \qquad \mu, n = 0, \ldots, M-1.$$

1. Reactivation: read the DFT equations

1. Which index belongs to time and which to frequency?
2. Why does an M-point DFT contain exactly M frequency samples?
3. What does the factor $1/M$ in the stated convention do?
4. Is the DFT input interpreted as a finite sequence or as one period of a periodic sequence when circular convolution is formed?

Complete the processing chain:

Complete the DFT-domain convolution chain



Label the transforms, operation, output, and convolution type.

2. Predict circular convolution before calculating it

Let

$$v(n) = \{1, 2, 3, 4\}, \qquad h(n) = \{1, 1, 1\},$$

so $M_1 = 4$ and $M_2 = 3$. The linear convolution

$$y_{\text{lin}}(n) = v(n) * h(n) = \sum_{\kappa} v(\kappa)\, h(n - \kappa)$$

has length $M_1 + M_2 - 1$.

1. Determine this length without calculating all output values.
2. Predict what happens if both sequences are transformed with an $M = 4$ DFT, their DFTs are multiplied, and a 4-point IDFT is applied.
3. Mark which indices of the linear convolution cannot remain at their original position when only four output positions are available.
4. State the smallest DFT length that can reproduce the complete linear convolution without time-domain aliasing.

Use the folding relation of the lecture (slide 10)

$$y_{\text{cir}}(n) = \tilde{v}(n) \circledast \tilde{h}(n) = \sum_{\lambda=-\infty}^{\infty} y_{\text{lin}}(n - \lambda M), \qquad n = 0, \ldots, M-1.$$

For the finite example, only the terms inside the support of y_{lin} contribute.

3. Demonstration: watch samples wrap around

The demonstration increases the DFT length step by step, starting at $M = 4$. Before each step, predict:

- whether any linear-convolution samples collide,
- where the last nonzero sample will appear after reduction modulo M, and
- whether the IDFT result equals the complete linear convolution.

In the notebook, follow one color at a time:

1. its original index in the linear convolution,
2. its destination $n \bmod M$, and
3. the sum formed when several samples share one destination.

Explain why the frequency-domain multiplication is correct even when the result is not the desired linear convolution.

Compare $M = 4, 5, 6$, and one larger value. Identify the first alias-free length and explain what the additional output positions contain.

4. Two derivations of the zero-padding condition

4.1 Index argument

The last possibly nonzero index of the linear convolution is

$$n_{\max} = (M_1 - 1) + (M_2 - 1).$$

Use this expression to derive a condition on M that prevents two different linear-convolution indices from being mapped to the same circular index.

Complete:

$$\boxed{M \;\geq\; \underline{\qquad\qquad\qquad}}$$

4.2 Spectral argument (lecture, slide 9)

The DTFT of the finite sequence $v(n)$ is

$$V(e^{j\Omega}) = \sum_{n=0}^{M_1-1} v(n) e^{-j\Omega n}.$$

1. Show that for $M \geq M_1$

$$\widetilde{V}_M(\mu) = V(e^{j\Omega})\big|_{\Omega=\Omega_\mu}, \qquad \Omega_\mu = \frac{2\pi}{M}\mu.$$

2. Conclude that $\widetilde{V}_M(\mu)\widetilde{H}_M(\mu)$ consists of M samples of the spectrum $Y(e^{j\Omega}) = V(e^{j\Omega})H(e^{j\Omega})$ of the linear convolution.
3. The lecture states that $y_{\text{lin}}(n)$ is represented uniquely by samples of $Y(e^{j\Omega})$ only if their number is at least $M_1 + M_2 - 1$. Explain this statement with the folding relation of Section 2: what does too coarse sampling in frequency cause in time?

4.3 What zero-padding changes

The second notebook view shows the magnitudes $|V(e^{j\Omega})|$ and $|\widetilde{V}_M(\mu)|$ for a fixed sequence of $M_1 = 16$ samples, $v(n) = \sin(\pi n/4)$.

1. Predict for $M = 16$: how many DFT samples will be clearly nonzero? Does the sequence therefore have an ideal line spectrum? Check your prediction with $M = 32$.
2. Explain why zero-padding changes the DFT grid and the circular period, but does not append new information to the original sequence.

Compare $M = 32$ and $M = 128$. Which displayed properties belong to the sequence, and which belong only to the chosen grid?

5. Peer oral examination

Work in pairs using examination cards A and B.

- 3 min: Student A examines student B using card A.
- 3 min: Change roles; student B examines student A using card B.
- 2 min: Give feedback using the checklist.

6. Take-away

Complete and explain:

1. Multiplication of two M-point DFTs followed by an IDFT produces ...
2. It equals the complete linear convolution if ..., because ...