ADSP - Exercise 3: Model Solution and Teaching Notes

Teaching focus and timing

The central obstacle is not computing a DFT but understanding what the DFT-domain product represents. Require a prediction before each slider change. Reactivation: about 90 s individually, then plenary; pairs only in the oral examination. Keep circular convolution, linear convolution, and denser DTFT sampling distinct, and show that time-domain aliasing and too coarse frequency sampling are the same effect.

Notation aligned with Exercise 4 and the lecture notes, Part 4, slides 3-10 (the lecture usually leaves the zero-padding implicit): $v(n)$ and $h(n)$ with lengths M_1 and M_2, zero-padded sequences $\tilde{v}(n)$ and $\tilde{h}(n)$ of length M, their DFTs $\tilde{V}_M(\mu)$ and $\tilde{H}_M(\mu)$, $y_{\text{lin}}(n) = v(n) * h(n)$, and $y_{\text{cir}}(n) = \tilde{v}(n) \circledast \tilde{h}(n)$.

Suggested 45-minute teaching sequence:

Phase	Time	Teaching purpose
Reactivate DFT/IDFT and complete the processing chain	0-6 min	establish notation and periodic interpretation
Predict the finite convolution example	6-14 min	expose the circular/linear distinction before the demo
Explore the DFT-length slider	14-24 min	observe wrap-around, collisions, and the alias-free threshold
Index derivation of the zero-padding condition (core)	24-30 min	students derive $M \geq M_1 + M_2 - 1$ themselves
Spectral argument, guided (slide 9)	30-33 min	teacher-led extension; written out after class
Coherent DTFT example, short demo	33-35 min	predict $M = 16$, check with $M = 32$
Peer oral examination	35-43 min	practise explanation and follow-up questions
Take-away	43-45 min	secure the two central statements

The index derivation is the central joint result. The spectral argument is a guided extension: present the three steps, let students supply the conclusion of step 3, and leave the complete write-up to the follow-up. The comparison of $M = 32$ and $M = 128$ is exploration after class.

Planned and extra tasks. The worksheet is ordered by topic and does not mark this distinction, because all tasks are exam-relevant. The 45-minute plan covers the tasks in the middle column; the extra tasks are for self-study, and the answers below cover them. The comparison of $M = 4, 5, 6$ in Section 3 is part of the slider demo in class.

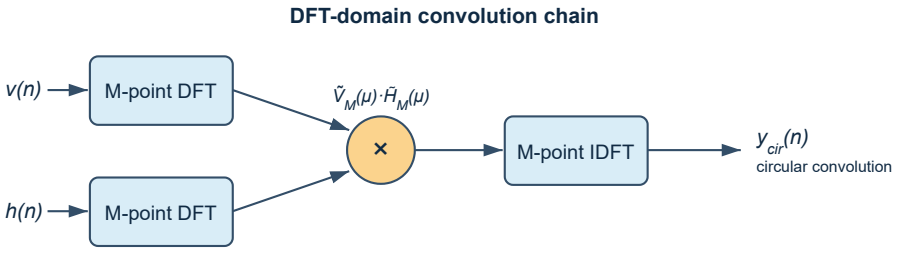
Section	Planned in class	Extra (self-study, exam-relevant)
1 Reactivation: read the DFT equations	1.1-1.4, processing chain	–
2 Predict circular convolution	2.1, 2.2, 2.3, 2.4	–
3 Demonstration: watch samples wrap around	steps 1-3, comparison of $M = 4, 5, 6$	–
4 Two derivations of the zero-padding condition	4.1, 4.2 (guided), 4.3.1, 4.3.2	comparison $M = 32/128$ (closing paragraph of 4.3)

1. Reading the DFT equations (0-6 min)

In the stated convention, n is the time index and μ is the frequency-grid index. An M-point DFT maps M time samples to M equally spaced samples around one 2π-period of frequency. The IDFT reconstructs the M-

sample sequence; $1/M$ supplies the inverse normalization.

The tilde denotes zero-padding to length $M \geq \max(M_1, M_2)$, never truncation. Circular convolution treats finite arrays as one period of M-periodic sequences (slide 5: $y(n) = y(n + \lambda M)$). This periodic interpretation causes wrap-around.



Linear convolution is obtained only if $M \geq M_1 + M_2 - 1$ (zero-padding to length M).

2. Circular and linear convolution (6-14 min)

For

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

the linear output length and result are

$$M_1 + M_2 - 1 = 6, \qquad y_{\text{lin}}(n) = \{1, 3, 6, 9, 7, 4\}.$$

This is y_0 of the Overlap-Add example in Exercise 4.

With $M = 4$, circular indexing maps $4 \mapsto 0$ and $5 \mapsto 1$. Therefore

$$y_{\text{cir}}(n) = \{1 + 7,\ 3 + 4,\ 6,\ 9\} = \{8, 7, 6, 9\}.$$

Nothing is numerically wrong with the DFT calculation. The theorem (slide 5) states

$$\text{IDFT}_M\{\tilde{V}_M(\mu)\tilde{H}_M(\mu)\} = \tilde{v}(n) \circledast \tilde{h}(n),$$

which is M-point circular convolution. It must not be identified with linear convolution without checking the length.

For $M = 5$, index 5 wraps to index 0, yielding

$$y_{\text{cir}}(n) = \{5, 3, 6, 9, 7\}.$$

For $M = 6$, all six samples have distinct destinations and the result equals the linear convolution. With $M = 8$, the same six values are followed by two zeros.

3. Expected demonstration observations (14-24 min)

M	Collision?	Wrapped indices	Result
4	yes	$4 \to 0,\ 5 \to 1$	circular result differs
5	yes	$5 \to 0$	circular result differs
6	no	none	equals linear convolution
8	no	none	equals linear convolution plus zeros

The notebook shows every linear-convolution sample before folding, its destination $n \bmod M$, and the sum at that destination. Samples with the same destination collide; this is time-domain aliasing.

4. Two derivations of the zero-padding condition (24-35 min)

4.1 Index argument (24-30 min)

The last possibly nonzero linear-convolution index is

$$n_{\max} = M_1 + M_2 - 2.$$

To represent indices $0, \ldots, n_{\max}$ without modulo reduction, require $M - 1 \geq n_{\max}$. Hence

$$\boxed{M \geq M_1 + M_2 - 1}.$$

4.2 Spectral argument, guided (slide 9; 30-33 min)

For $M \geq M_1$, all samples of $v(n)$ lie in $0, \ldots, M-1$, and the DFT sum equals the DTFT sum evaluated at $\Omega_\mu = 2\pi\mu/M$:

$$\widetilde{V}_M(\mu) = \sum_{n=0}^{M_1-1} v(n) e^{-j\frac{2\pi}{M}\mu n} = V(e^{j\Omega_\mu}).$$

The same holds for $h(n)$ when $M \geq M_2$. Hence

$$\widetilde{V}_M(\mu)\widetilde{H}_M(\mu) = V(e^{j\Omega_\mu})H(e^{j\Omega_\mu}) = Y(e^{j\Omega_\mu}),$$

i.e. the DFT product consists of samples of the spectrum of the **linear** convolution. For $M \geq \max(M_1, M_2)$, the product contains the correct M samples of $Y(e^{j\Omega})$; these may still be insufficient to recover the finite linear convolution without overlap. This shows up in the inverse step: the IDFT of M samples of $Y(e^{j\Omega})$ is

$$\frac{1}{M}\sum_{\mu=0}^{M-1} Y(e^{j\Omega_\mu}) e^{j\frac{2\pi}{M}\mu n} = \sum_k y_{\text{lin}}(k)\,\frac{1}{M}\sum_{\mu=0}^{M-1} e^{j\frac{2\pi}{M}\mu(n-k)} = \sum_{\lambda=-\infty}^{\infty} y_{\text{lin}}(n - \lambda M),$$

because the inner sum is one if $k = n - \lambda M$ and zero otherwise. Sampling the spectrum at M points periodizes the sequence with period M; this is exactly the folding relation of Section 2 and $y_{\text{p}}(n)$ on slide 10. The periodized copies do not overlap if and only if the period is at least the length of y_{lin}, i.e. $M \geq M_1 + M_2 - 1$. This is the lecture's formulation: $Y(e^{j\Omega})$ must be represented by at least $M_1 + M_2 - 1$ samples.

Teaching point: aliasing in time after too coarse sampling in frequency is the dual of the spectral replicas after too coarse sampling in time (Exercise 1).

4.3 What zero-padding changes (33-35 min)

Zero-padding supplies enough positions for the circular period to contain the complete linear-convolution support. It gives a denser sampling grid of the same DTFT, but adds no observations and does not improve the intrinsic resolution set by the record length.

Coherent example in the notebook. The sequence $v(n) = \sin(\pi n/4)$, $n = 0, \ldots, 15$, contains exactly two periods. Its DTFT consists of two main lobes at $\Omega = \pm\pi/4$ with zeros spaced by $2\pi/16$ and decaying side lobes; these are caused by the finite observation length. For $M = 16$ the grid spacing is also $2\pi/16$: the samples hit the two peaks ($|\widetilde{V}_{16}(2)| = |\widetilde{V}_{16}(14)| = 8$) and otherwise fall exactly on zeros of the DTFT. The DFT then *looks* like an ideal line spectrum, although the DTFT is not. $M = 32$ already places samples between the zeros and reveals main-lobe width and side lobes; $M = 128$ shows the curve almost continuously. The side lobes were present all along; only the grid made them visible.

Expected answer to 4.3.1: two clearly nonzero samples, but no ideal line spectrum; the grid can hide properties of the DTFT. Leakage and windowing follow in Exercise 6.

5. Common misconceptions

- **"DFT multiplication gives linear convolution."** It gives circular convolution of the selected length.
- **"The DFT product is wrong if M is too small."** For $M \geq \max(M_1, M_2)$, the product contains correct samples of $Y(e^{j\Omega})$, but possibly too few; their IDFT periodizes y_{lin}.
- **"The tail is discarded."** It wraps modulo M and is added to earlier samples.
- **"Few nonzero DFT samples mean a pure line spectrum."** The coherent example shows that the grid can fall on zeros of the DTFT.
- **"M must be a power of two."** That can be convenient for an FFT implementation but is not required by the convolution theorem.

6. Take-away (43-45 min)

1. Multiplication of two M-point DFTs followed by an IDFT produces their M-point circular convolution, i.e. the linear convolution periodized with period M.
2. It equals the complete linear convolution when $M \geq M_1 + M_2 - 1$, because only then do the periodized copies, created by sampling $Y(e^{j\Omega})$ at M points, not overlap.