

Fixed-Point CPWC Ultrasound Image Reconstruction

Ji Shi

Electrical and Computer Engineering Department
University of Victoria
Victoria, Canada
jishi@uvic.ca

Daler Rakhmatov

Electrical and Computer Engineering Department
University of Victoria
Victoria, Canada
daler@ece.uvic.ca

Abstract—Coherent plane-wave compounding (CPWC) ultrasonography is an important imaging modality that allows for very high frame rates. During CPWC image reconstruction, computationally expensive delay-and-sum beamforming can be replaced by faster Fourier-domain remapping. We show that the cost of such spectral reconstruction can be reduced further, by employing scaled fixed-point arithmetic amenable to efficient beamforming hardware realization.

Index Terms—Plane-wave ultrasound imaging, Fourier-domain image reconstruction, fixed-point computations

I. INTRODUCTION

Ultrafast plane-wave (PW) ultrasound imaging features the data acquisition rates reaching thousands of frames per second, which gives rise to new possibilities in echocardiography and elastography applications [1]. Typically, the PW image reconstruction process involves coherent compounding of multiple beamformed frames, each corresponding to a particular tilted PW emission (over a set of multiple angles). Such coherent plane-wave compounding (CPWC) improves the resulting image quality but reduces the overall frame rate.

There are numerous methods we can use to obtain compounded PW images in two spatial dimensions (z, x) , where z and x refer to the axial and lateral coordinates, respectively. For example, one common technique is the standard delay-and-sum (DAS) beamformer that operates in the (t, x) domain, where t represents the temporal axis (sampling time instances). Alternatively, image reconstruction can be done in the *spatio-temporal frequency domain*: the (f, k_x) -domain dataset is remapped into the (k_z, k_x) -domain dataset, where f denotes the temporal frequencies, while k_x and k_z denote the spatial frequencies (e.g., see [2]–[6]).

This paper focuses on one of the Fourier-domain remapping methods proposed in [6], offering substantially lower computational latency compared to conventional DAS beamforming. We demonstrate that it is possible to obtain high-quality compounded images using hardware-oriented *scaled fixed-point arithmetic* operations during spectral reconstruction, as opposed to using more accurate but costly software-oriented floating-point calculations.

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

II. BACKGROUND

This section briefly summarizes our Fourier-domain migration method for reconstructing coherently compounded PW images, borrowing from our previous work [6] with a slight change in notation to streamline the presentation.

Let θ represent a PW emission angle, and let $P(t, z, x, \theta)$ denote the resulting acoustic wavefield. Given the wavefield $P(t, 0, x, \theta)$ recorded over time at the surface (i.e., at depth $z = 0$), we want to reconstruct the subsurface image dataset $P(0, z, x, \theta)$ at time $t = 0$.¹ This goal can be accomplished using Fourier-domain interpolation as follows.

Let $\Psi(f, 0, k_x, \theta)$ and $\tilde{\Psi}(0, k_z, k_x, \theta)$ denote the Fourier transforms of known $P(t, 0, x, \theta)$ and unknown $P(0, z, x, \theta)$:

$$\Psi(f, 0, k_x, \theta) = \iint P(t, 0, x, \theta) e^{-j2\pi(k_x x + ft)} dx dt, \quad (1)$$

$$P(0, z, x, \theta) = \iint \tilde{\Psi}(0, k_z, k_x, \theta) e^{j2\pi(k_x x + k_z z)} dk_x dk_z. \quad (2)$$

We have $\Psi(f, 0, k_x, \theta)$ as an input, generated by the Fourier transform of $P(t, 0, x, \theta)$. We need to obtain $\tilde{\Psi}(0, k_z, k_x, \theta)$ from $\Psi(f, 0, k_x, \theta)$, so that sought $P(0, z, x, \theta)$ can be computed via the inverse Fourier transform of $\tilde{\Psi}$. In our previous work [6] (extending classic Stolt's migration method [7]), we proposed the remapping formula shown below, producing the *intermediate spectrum*

$$\tilde{\Psi}^*(0, k_z, k_x, \theta) = A(k_z, k_x, \theta) \cdot \Psi(f_{\text{mig}}(k_z, k_x, \theta), 0, k_x, \theta), \quad (3)$$

with the values of f_{mig} and A determined by

$$f_{\text{mig}}(k_z, k_x, \theta) = \frac{ck_z}{1 + \cos(\theta)} [1 + (k_x/k_z)^2], \quad (4)$$

$$A(k_z, k_x, \theta) = \frac{c}{1 + \cos(\theta)} [1 - (k_x/k_z)^2], \quad (5)$$

where c is the speed of sound. Applying the inverse Fourier transform to $\tilde{\Psi}^*(0, k_z, k_x, \theta)$ will yield a preliminary image dataset $P(0, z^*, x, \theta)$, from which we can obtain $P(0, z, x, \theta)$ by repositioning data points at locations (z^*, x) to their new coordinates $(z = z^* + x \tan(\theta)/2, x)$ [6].

¹This approach is based on the so-called *exploding reflector model*, assuming that the backscattered echoes are produced by reflectors “exploding” at $t = 0$ (e.g., see [7]).

To compound multiple angle-specific $P(0, z, x, \theta_n)$ over a given angular set $\{\theta_n \mid n = 1, 2, \dots, N_a\}$, we perform the following summation [6]:

$$C(k_z, x) = \sum_{n=1}^{N_a} e^{j\pi k_z x \tan(\theta_n)} \cdot \bar{\Psi}(k_z, x, \theta_n), \quad (6)$$

$$\bar{\Psi}(k_z, x, \theta_n) = \int \tilde{\Psi}^*(0, k_z, k_x, \theta_n) e^{j2\pi k_x x} dk_x, \quad (7)$$

where $\bar{\Psi}(k_z, x, \theta_n)$ represents the 1D inverse Fourier transform of $\tilde{\Psi}^*(0, k_z, k_x, \theta_n)$ along the k_x -axis, and $C(k_z, x)$ is the compounded (k_z, x) -domain dataset. We get the final image dataset, denoted by $D(z, x)$, via the 1D inverse Fourier transform of $C(k_z, x)$ along the k_z -axis:

$$D(z, x) = \int C(k_z, x) e^{j2\pi k_z z} dk_z. \quad (8)$$

Fig. 1 outlines the corresponding computational procedure. The 3D input $P[\cdot]$ is a raw RF channel dataset recorded over N_t time instances $\{t_l = 0, \Delta_t, \dots, (N_t - 1)\Delta_t\}$, N_x sensor locations $\{x_m = -\frac{N_x}{2}\Delta_x, \dots, 0, (\frac{N_x}{2} - 1)\Delta_x\}$, and N_a PW emission angles $\{\theta_n \mid n = 1, 2, \dots, N_a\}$, i.e., $P[\cdot]$ represents $P(t, 0, x, \theta)$. The other 3D inputs $M[\cdot]$, $A[\cdot]$, and $E[\cdot]$ represent $f_{\text{mig}}(k_z, k_x, \theta)$, $A(k_z, k_x, \theta)$, and $\exp(j\pi k_z x \tan(\theta))$, respectively. The 2D output $H[\cdot]$ is obtained via the Hilbert transform (along the z -axis) of the final dataset $D(z, x)$ represented by $D[\cdot]$ (lines 12 and 13 in Fig. 1). This output is useful for subsequent data processing (e.g., detecting the envelope), and it is computed by the HILBERT function for each x_m value.

Input: Raw dataset $P[\cdot]$, map $M[\cdot]$, scaler $A[\cdot]$, phase shift $E[\cdot]$
Output: Hilbert-transformed compounded image dataset $H[\cdot]$

```

1.  $C \leftarrow 0$ ;
2. for  $n = 1 : N_a$  do {
3.   for  $m = 1 : N_x$  do {  $F[:, m] \leftarrow \text{FFT}(P[:, m, n], N_t^{\text{FFT}});$  }
4.   for  $l = 1 : N_t^{\text{FFT}}$  do {  $F[l, :] \leftarrow \text{FFT}(F[l, :], N_x^{\text{FFT}});$  }
5.   for  $m = 1 : N_x^{\text{FFT}}$  do {
6.      $K[:, m] \leftarrow \text{REMAP}(F[:, m], M[:, m, n]);$ 
7.      $K[:, m] \leftarrow A[:, m, n] \times K[:, m];$  }
8.   for  $l = 1 : N_t^{\text{FFT}}$  do {  $K[l, :] \leftarrow \text{IFFT}(K[l, :], N_x^{\text{FFT}});$  }
9.   for  $m = 1 : N_x$  do {  $K[:, m] \leftarrow E[:, m, n] \times K[:, m];$  }
10.   $C \leftarrow C + K$ ; }
11. for  $m = 1 : N_x$  do {
12.   $D[:, m] \leftarrow \text{IFFT}(C[:, m], N_t^{\text{FFT}});$ 
13.   $H[:, m] \leftarrow \text{HILBERT}(D[:, m]);$  }
```

Fig. 1. Spectral migration algorithm, based on [6].

The 1D temporal and spatial Fourier transforms and their inverses are computed by the FFT and IFFT functions, using the power-of-2 transform lengths denoted by N_t^{FFT} and N_x^{FFT} . Upon execution of lines 3 and 4 in Fig. 1, we obtain the (f, k_x) -domain spectrum $F[\cdot]$ of size $N_t^{\text{FFT}} \times N_x^{\text{FFT}}$. Then, for each k_x bin indexed by $m = 1, 2, \dots, N_x^{\text{FFT}}$, we remap the f -axis points to the k_z -axis points according to $M[:, m, n]$ and multiply the resulting data pointwise by $A[:, m, n]$ (lines 6 and 7 in Fig. 1). In other words, we interpolate $F[\cdot]$ using $f_{\text{mig}}(k_z, k_x, \theta_n)$ and then scale it by $A(k_z, k_x, \theta_n)$ to get the (k_z, k_x) -domain spectrum $K[\cdot]$ of size $N_t^{\text{FFT}} \times N_x^{\text{FFT}}$. Next,

for each k_z bin indexed by $l = 1, 2, \dots, N_t^{\text{FFT}}$, we transform $K[\cdot]$ back to the (k_z, x) domain (line 8 in Fig. 1) and apply appropriate phase shifts $\exp(j\pi k_z x_m \tan(\theta_n))$ specified by $E[:, m, n]$. Finally, each angle-specific $K[\cdot]$ is added to $C[\cdot]$ that represents the compounded 2D dataset $C(k_z, x)$ of size $N_t^{\text{FFT}} \times N_x$. After processing all angles, the inverse Fourier transform of $C[\cdot]$ along the k_z -axis yields the (z, x) -domain image dataset $D[\cdot]$, which is then Hilbert-transformed along the z -axis to produce the desired output $H[\cdot]$ of size $N_z \times N_x$.

III. FIXED-POINT RECONSTRUCTION

Fig. 2 depicts our scaled fixed-point computational flow implementing the Fourier-domain reconstruction algorithm from the previous section (see Fig. 1). In Fig. 2, the Fourier transform calculations are performed by the 1D split-radix FFT blocks [8] (see lines 3 and 4 in Fig. 1); however, the number of iterations over both indices $m = 1, 2, \dots, N_x$ and $l = 1, 2, \dots, N_t^{\text{FFT}}$ is cut in half. Given that the input $P[\cdot]$ is real-valued, its (f, x) spectrum will be symmetric, i.e., we only need its positive- f portion (hence $N_t^{\text{FFT}}/2$ iterations) to compute the (f, k_x) half-spectrum for subsequent remapping. To get the positive- f portion of the (f, x) spectrum, we process the t -axis data for x_m and x_{m+1} in pairs (hence $N_x/2$ iterations): the former acts at the real part of the FFT input, while the latter acts as the corresponding imaginary part [8].

During fixed-point computations, the real/imaginary data values are restricted to the interval $[-1, +1]$. Whenever these values fall outside the allowed limits, they undergo binary scaling (i.e., division by some fitting power of 2) to enforce our range restriction. After completing the temporal FFTs over all x bins, we have the scaling factors $S(x)$ associated with the f -axis data vectors. Before starting the spatial FFTs, we equalize the x -axis data vectors using the maximum of $S(x)$. Similarly, after completing the spatial FFTs over all positive- f bins, we use the maximum of the scaling factors $S(f)$, associated with the k_x -axis data vectors, to equalize the f -axis data vectors prior to frequency remapping.

The REMAP/MULTIPLY block in Fig. 2 implements lines 6 and 7 in Fig. 1. Given a particular k_x bin, we let the k_z -axis data values equal those found (via linear interpolation) at $f_{\text{mig}}(k_z, k_x, \theta)$ and scaled by $A(k_z, k_x, \theta)$. In Fig. 2, the values of $f_{\text{mig}}(k_z, k_x, \theta)$ and $A(k_z, k_x, \theta)$ are provided by the inputs $M[\cdot]$ and $A[\cdot]$, respectively. Next, the resulting (k_z, k_x) half-spectrum is fed into the spatial IFFTs that yield the (k_z, x) half-spectrum and the scaling factors $S(k_z)$. To apply the phase shifts $\exp(j\phi(k_z, x, \theta))$, where $\phi(k_z, x, \theta) = \pi k_z x \tan(\theta)$, we first equalize the x -axis data vectors using the maximum of $S(k_z)$, and then perform CORDIC-based phase rotations [9] specified by $\phi(k_z, x, \theta)$. In Fig. 2, the values of $\phi(k_z, x, \theta)$ are provided by the input $R[\cdot]$ (as opposed to the input $E[\cdot]$ in Fig. 1).

After the phase-rotation step, each θ -specific half-spectrum in the (k_z, x) domain is compounded with the others (line 10 in Fig. 1) to produce half-sized $C[\cdot]$ over positive- k_z bins. Instead of expanding such $C[\cdot]$ into the full-sized symmetric (k_z, x) spectrum to undergo the N_t^{FFT} -point IFFTs, followed

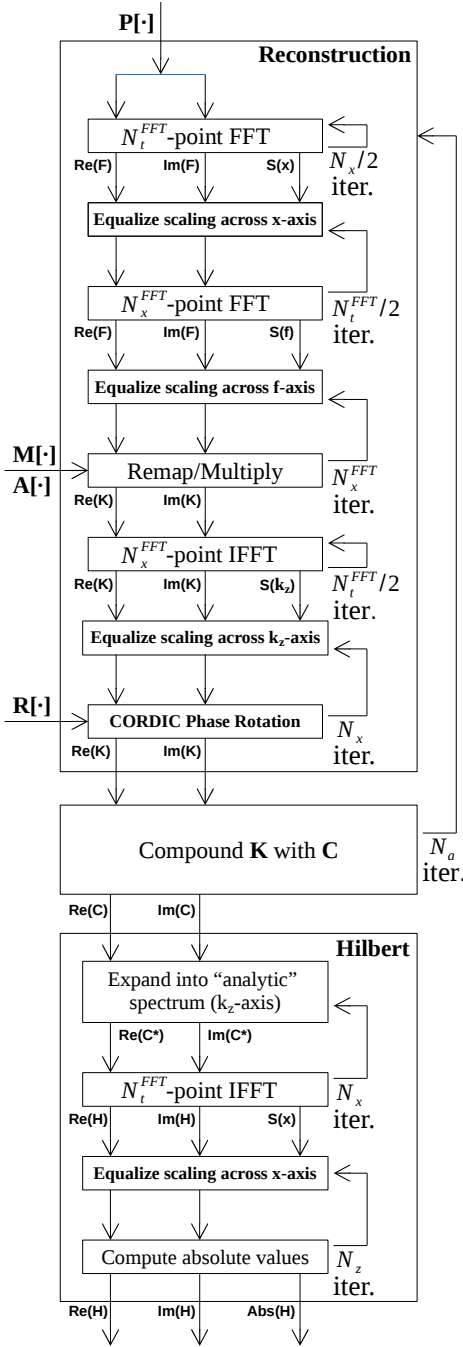


Fig. 2. Fixed-point scaled reconstruction.

by the Hilbert transforms (lines 12 and 13 in Fig. 1), we obtain the desired output $H[\cdot]$ directly from half-sized $C[\cdot]$. For each $m = 1, 2, \dots, N_x$, we first expand $C[\cdot]$ into its full-sized “analytic” version $C^*[\cdot]$ in accordance with [10]. Then, we compute the N_t^{FFT} -point IFFTs of the individual column vectors of $C^*[\cdot]$ to get the Hilbert-transformed compounded image dataset:

$$H[:, m] \leftarrow \text{IFFT}(C^*[:, m], N_t^{\text{FFT}}). \quad (9)$$

In other words, we have effectively eliminated line 13 in Fig.

1 by replacing symmetric $C[\cdot]$ with $C^*[\cdot]$ based on [10]. After computing the IFFTs along the k_z -axis for each x bin, we get the (z, x) -domain output $H[\cdot]$ and the scaling factors $S(x)$. The final computational block in Fig. 2 equalizes the z -axis data vectors using the maximum of $S(x)$, and it also outputs the absolute values of $H[\cdot]$ giving the envelope.

IV. EVALUATION RESULTS

For fixed-point arithmetic testing, we have used several experimental datasets from PICMUS-2016 [11] that utilize $N_a = 11$ plane waves emitted at angles $\pm 16^\circ$, $\pm 13^\circ$, $\pm 9.5^\circ$, $\pm 6.5^\circ$, $\pm 3.0^\circ$, and 0° . Specifically, we evaluate the following imaging cases:

- A) Two anechoic cylinder targets (cyst phantoms), Fig. 3;
- B) Seven wire targets (point phantoms), Fig. 4;
- C) Carotid artery – longitudinal section, Fig. 5;
- D) Carotid artery – cross section, Fig. 6.

For any given angle θ , the N_t -by- N_x size of raw RF channel data frames was 3328×128 in cases A and B, and 1536×128 in cases C and D. We generated the compounded B-mode images by log-compressing their respective normalized envelope sections of size 1216×128 , covering the imaging depth from 5 to 50 mm as shown in Fig. 3-6 using the 60-dB dynamic range. In all four cases, we let $N_t^{\text{FFT}} = 4096$ and $N_x^{\text{FFT}} = 256$, giving $M_t^{\text{FFT}} = \log_2 N_t^{\text{FFT}} = 12$ and $M_x^{\text{FFT}} = \log_2 N_x^{\text{FFT}} = 8$.

Our fixed-point parameter settings are summarized in Table I, with the wordlengths limited to 16 or 24 bits only. Since the values of $P[\cdot]$, $F[\cdot]$, and $K[\cdot]$ are restricted to the $[-1, +1]$ range, the integer part of their signed fixed-point representation is only 1 bit long. The number of fractional-part bits has been set to 14, which is equal to $\max\{M_t^{\text{FFT}}, M_x^{\text{FFT}}\}$ plus 2 extra bits, in order to match the fractional-part length of the FFT twiddle factors having sufficient resolution. When computing $C[\cdot]$ and $H[\cdot]$, we have increased their fixed-point wordlength by additional 8 bits for the benefit of CPWC.

TABLE I
FIXED-POINT PARAMETER SETTINGS.

| Parameter | P | M | A | R | F / K | C / H |
|------------|-----|----|-----|-----|-------|-------|
| Signed | Yes | No | Yes | Yes | Yes | Yes |
| Int. Part | 1 | 12 | 1 | 3 | 1 | 1 |
| Frac. Part | 14 | 12 | 14 | 12 | 14 | 22 |
| Wordlength | 16 | 24 | 16 | 16 | 16 | 24 |

To accommodate the permissible range of the f_{mig} values and to allow for the adequate interpolation accuracy, we let the integer and fractional parts of unsigned $M[\cdot]$ have 12 bits each, thus keeping its wordlength at 24-bit limit. As for $A[\cdot]$, its fixed-point representation has been made compatible with the data format of $K[\cdot]$ (via binary prescaling and redundant signedness of the A values). Since the phase rotation block in Fig. 2 may take any ϕ value between -2π and 2π , signed $R[\cdot]$ has the integer part of 3 bits. We have allocated 12 bits to its fractional part to maintain the target 16-bit wordlength.

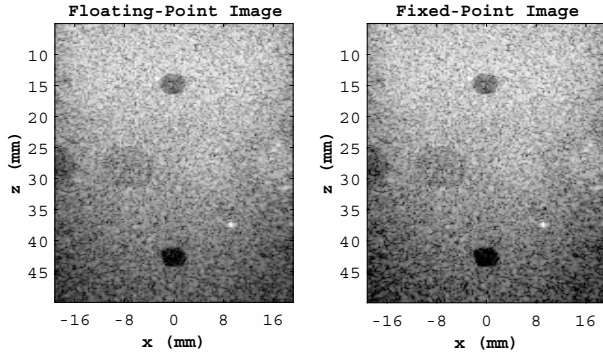


Fig. 3. Case A: Compounded images of the cyst phantoms (11 PWs).




Fig. 4. Case B: Compounded images of the point phantoms (11 PWs).

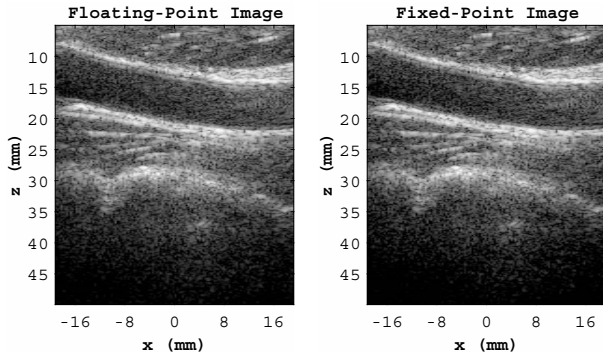


Fig. 5. Case C: Compounded images of the carotid artery longitudinal section (11 PWs).




Fig. 6. Case D: Compounded images of the carotid artery cross section (11 PWs).

The resulting B-mode images (reconstructed using our fixed-point settings from Table I) are displayed in Fig. 3-6, next to the reference images obtained with more precise but more expensive floating-point calculations. As one can see, the fixed-point and floating-point versions are almost indistinguishable visually. To evaluate reconstruction differences quantitatively, we have compared the fixed-point versions of a 2D normalized envelope prior to log-compression (i.e., 1216×128 datasets of values ranging from 0 to 1) to their respective floating-point references. Our comparisons rely on three measures of similarity listed in Table II: the structural similarity index (SSIM), peak-signal-to-noise ratio (PSNR), and mean absolute error (MAE). Table II also includes the average data value of each fixed-point dataset, which provides a baseline for MAE interpretation. Since the error values are at least an order of magnitude smaller than the corresponding averages, one can view the numerical distance between the fixed-point and floating-point datasets as insignificant. The PSNR values ranging from 48 to 65 dB and the SSIM values exceeding 0.99 confirm that the fixed-point reconstruction results are indeed very close to their respective floating-point references.

TABLE II
NORMALIZED ENVELOPE SIMILARITY BETWEEN FIXED-POINT AND
FLOATING-POINT COMPOUNDED DATA (11 PLANE WAVES).

| Case | Mean | MAE | PSNR | SSIM |
|------|------------------------|------------------------|----------|--------|
| A | 1.090×10^{-1} | 2.893×10^{-3} | 47.81 dB | 0.9965 |
| B | 1.288×10^{-2} | 4.436×10^{-4} | 65.09 dB | 0.9993 |
| C | 3.074×10^{-2} | 1.088×10^{-3} | 54.65 dB | 0.9981 |
| D | 1.701×10^{-2} | 9.604×10^{-4} | 54.21 dB | 0.9987 |

REFERENCES

- [1] M. Tanter and M. Fink, "Ultrafast imaging in biomedical ultrasound," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 61, no. 1, pp. 102–119, Jan 2014.
- [2] J. Cheng and J. Lu, "Extended high-frame rate imaging method with limited-diffraction beams," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 53, no. 5, pp. 880–899, May 2006.
- [3] D. Liu and T. Ji, "Plane wave image formation in spatial-temporal frequency domain," in *Proc. IEEE International Ultrasonics Symposium*, Sep 2016, pp. 1–5.
- [4] D. Garcia, L. Tarnec, S. Muth, E. Montagnon, J. Poree, and G. Cloutier, "Stolt's f - k migration for plane wave ultrasound imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 60, no. 9, pp. 1853–1867, Sep 2013.
- [5] A. Besson, M. Zhang, F. Varray, H. Liebgott, D. Friboulet, Y. Wiaux, J. Thiran, R. Carrillo, and O. Bernard, "A sparse reconstruction framework for Fourier-based plane-wave imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 63, no. 12, pp. 2092–2106, Dec 2016.
- [6] M. Albulayli and D. Rakhmatov, "Fourier-domain depth migration for plane wave ultrasound imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 65, no. 8, pp. 1321–1333, Aug 2018.
- [7] O. Yilmaz, *Seismic Data Analysis*, 2nd ed. OK: SEG, 2001, vol. 1.
- [8] H. Sorensen and C. Burrus, "Fast DFT and convolution algorithms," in *Handbook for Digital Signal Processing*, S. Mitra and J. Kaiser, Eds. NY: Wiley, 1993, ch. 8, pp. 491–610.
- [9] J. Muller, *Elementary Functions*, 3rd ed. NY: Springer, 2016.
- [10] S. Marple, "Computing the discrete-time 'analytic' signal via FFT," *IEEE Trans. Signal Processing*, vol. 47, no. 9, pp. 2600–2603, Sep 1999.
- [11] H. Liebgott, A. Rodriguez-Molares, F. Cervensky, J. Jensen, and O. Bernard, "Plane-wave imaging challenge in medical ultrasound," in *Proc. IEEE International Ultrasonics Symposium*, Sep 2016, pp. 1–4.