

Ultrafast Medical Ultrasound Imaging Using Diverging Waves: Optimization Strategies for Enhanced Image Quality



**Università
di Genova**

Zahraa Alzein

Supervisor: Prof. Daniele Caviglia

Dr. Marco Crocco

Department of Electrical, Electronic, Telecommunications Engineering and
Naval Architecture (DITEN)
University of Genoa

This dissertation is submitted for the degree of
Doctor of Philosophy

Joint Doctorate in Interactive and
Cognitive Environments - Cycle 38

December 2025

I would like to dedicate this thesis to my loving Family ...

Declaration

I hereby declare that, except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 65,000 words including bibliography, footnotes, tables and equations and has fewer than 150 figures.

Zahraa Alzein
December 2025

Acknowledgements

This thesis marks the culmination of three years of hard work and dedication, and I acknowledge with deep gratitude the many people who have supported me along the way.

First, I would like to express my sincere gratitude to my supervisor, Prof. Daniele Caviglia, for his guidance, helpful feedback, and unwavering support. His knowledge and commitment have greatly influenced my research. I also want to thank Dr. Marco Crocco, my co-supervisor at Esaote S.P.A., invaluable supervision, support, and continuous help throughout my PhD journey. I am also especially grateful to Prof. Herve Libgott from the University of Lyon 1 for his support and supervision during my doctoral mobility collaboration.

I also appreciate the University of Genova in Italy and Esaote S.P.A. for creating a positive work environment with their dedicated staff.

Finally, I would like to thank my parents, friends for their support, encouragement, and appreciation. This work is dedicated to all of you.

Zahraa Alzein

Abstract

Ultrafast imaging has rapidly advanced the field of medical ultrasound by enabling data acquisition and image reconstruction at extremely high frame rates, approximately 1 to 2 orders of magnitude higher than conventional methods. Several advanced applications, such as shear-wave elastography, high-sensitivity Doppler imaging, contrast-enhanced imaging, and functional ultrasound have been made effective by the availability of ultrafast acquisition methods. Ultrafast imaging is typically achieved using either plane waves (PWs) or diverging waves (DWs). Both methods employ all transducer elements for transmission, PWs using varying steering angles, and DWs by spherical waves emitted from virtual sources (VS) placed behind the probe. The use of DWs in ultrafast imaging has two-fold advantages: high frame rate and wide field of view; however, despite these advancements, ultrafast imaging faces fundamental trade-offs—most notably, a compromise in image resolution, contrast, and frame rate.

This thesis aims to enhance image quality for ultrafast ultrasound imaging using diverging waves, maintaining high frame rate and real-time performance. In particular, I first derived a closed-form approach that maps, under suitable hypotheses, the transmit apodization weights used in synthetic aperture imaging into the compound mask applied to diverging wave imaging. The approach draws inspiration from a successful technique developed for plane wave imaging, leveraging synthetic aperture imaging as a reference due to its superior image quality. Unlike the previous work on plane waves, the proposed approach is not limited to linear probe geometries but also works seamlessly with convex ones, thus expanding the scope of applicability of diverging wave imaging. Moreover, it handles arbitrary spatial arrangements of virtual sources generating divergent waves. The approach has been validated through simulated data using both linear and convex probes. After that, the weights are integrated into the beamforming pipeline of the Verasonics scanner, validating its efficacy across three deterministic VS configurations (linear, curvilinear, and tilted distributions). Experimental results demonstrate that the compound mask improves the quality of B-mode images with all distributions of virtual sources for linear and convex arrays, all without compromising real-time performance.

Since the number and spatial distribution of these VSs affect both image quality and frame rate, their optimization is of high interest. I propose a multi-objective genetic algorithm to optimize VS spatial distributions with a compound mask weighting strategy to enhance beam coherence and reduce artifacts during optimization to further improve the image quality. The framework was evaluated across different numbers of VSs to quantify performance trade-offs under fewer transmission events. In both simulations and experimental trials, the proposed approach achieved improved image quality metrics against deterministic methods, preserving these gains even with reduced transmission events.

Next, I developed a computationally efficient receive beamformer and coherent compounder, suited to work with both radio frequency band and base band data related to plane-wave/diverging-wave insonifications under the constraints of a commercial ultrasound scanner, using GPU acceleration. The beamformer, implemented in GPU and taking into account the constraints of typical commercial ultrasound scanners, integrates multiple optimizations (precomputed delays, transmit/receive decoupling, symmetry exploitation, and carrier phase reconstruction) to drastically reduce computation and data transfer overhead. Moreover, it integrates the optimization strategies previously developed (compound mask, optimized VS). The GPU beamformer is validated using in-vitro datasets collected from a tissue-mimicking phantoms, using a Verasonics Vantage system with a convex probe. Tests are performed with deterministic and optimized virtual source distributions and multiple VS configurations, including 20/10/6/4 transmissions. The proposed GPU-based receive beamformer achieves real-time performance while maintaining image quality even when the number of virtual sources is significantly reduced.

Finally, the compound-mask apodization scheme generalizes from 2-D to fully volumetric ultrafast diverging-wave imaging with 2-D matrix probes, providing a practical basis for next-generation 3-D ultrasound. I devise a geometric, voxel-wise mapping that converts the apodization derived in 3-D synthetic transmit aperture imaging (STAI) into 3-D compound weights: for each virtual source and voxel, the source–voxel ray is traced to its intersection on the array face, yielding spatially consistent weights across the entire volume. The approach was assessed with FIELD II simulations, using matrix probes with full and randomly sparse elements and consistently reduce the sidelobe levels across depths.

Keywords: Ultrafast Ultrasound Imaging, Diverging wave, Image Quality, transmit apodization, virtual source distribution, multi-objective optimization, 3D ultrafast imaging, receive Beamformer, GPU- implementation, in-vitro data

Contents

List of Figures	xvii
List of Tables	xxv
1 Introduction	1
1.1 Motivation	1
1.2 Objectives and Contributions	3
1.2.1 Optimize compound weights for each virtual source transmission with real-time implementation on verasonics scanner	3
1.2.2 Multi-objective optimization of virtual source spatial distribution for ultrafast diverging wave imaging	4
1.2.3 Optimized GPU-based receive beamformer implementation for effi- cient and high-quality ultrafast diverging-wave imaging	4
1.2.4 Extension of compound Mask approach to 3D ultrafast imaging using 2D-matrix probe:	5
1.3 Thesis Outline	5
1.4 Research Outputs	6
1.4.1 International Journals:	7
1.4.2 Conference Papers:	7
1.4.3 Book Chapter:	8
1.4.4 Attendance at Conferences	8
2 Literature Review	9
2.1 Fundamentals of Ultrasound	9
2.1.1 Principle of ultrasound imaging	9
2.1.2 Conventional Ultrasound (Focused Imaging)	11
2.1.3 Alternative receive beamforming	15
2.1.4 Image quality in Ultrasound	17

2.2	Advanced Focusing Strategies	20
2.3	Unfocalized transmission methods for Ultrafast Ultrasound Imaging	22
2.3.1	Synthetic Transmit Aperture (STA)	23
2.3.2	Plane Wave Imaging (PWI)	23
2.3.3	Diverging Wave Imaging (DWI)	24
2.3.4	Coherent Compounding	25
2.4	Discussion and Problem Statement	27
2.5	Global Optimization Algorithms	28
2.5.1	Deterministic and Stochastic Optimization	28
2.5.2	Multi-Objective Optimization	32
2.5.3	Multi-objective optimization algorithms based on non-dominated solutions	33
2.5.4	Discussion	34
3	Compound Mask for Diverging Wave Imaging: Simulation Study	37
3.1	Introduction	37
3.2	Methodology	39
3.2.1	Synthetic Aperture Imaging (SA) using Linear and Convex Probes	39
3.2.2	Diverging Wave Imaging using Linear and Convex probe	40
3.2.3	Relation between transmit apodization in STAI and Compound Mask in DWI	41
3.2.4	Compound Mask Computation	44
3.2.5	Algorithm for compound mask computation	45
3.3	Results and Discussion	47
3.3.1	Geometrical warping	47
3.3.2	Validation of compound Mask using simulated data	48
3.4	Conclusion	51
4	Experimental Validation of Compound Mask in Real Time	53
4.1	Introduction	53
4.2	Methodology	53
4.2.1	Integration of the compound-mask apodization in the Verasonics pipeline	54
4.2.2	Virtual Sources Spatial Arrangement for Linear and Convex Arrays in DWI	54
4.2.3	Experimental Setup	55
4.2.4	Evaluation Metrics	56

4.3	Results and Discussion	57
4.3.1	B-mode Images using Linear probe L7-4:	57
4.3.2	B-mode Images using Convex probe C5-2	61
4.3.3	Comparative Analysis of Computational Complexity	62
4.3.4	Statistical significance testing for paired comparisons (one-sided)	63
4.4	Discussion and Conclusion	64
5	Multi-Objective Optimization of Virtual Source Distributions for Ultrafast Diverging Wave Imaging	67
5.1	Introduction	67
5.2	Methodology	70
5.2.1	Multi-Objective optimization Framework	70
5.2.2	Evaluation Metrics	74
5.2.3	Simulation Setup	76
5.2.4	Deterministic Distributions	76
5.2.5	Phantom with Anechoic Cysts and Reflectors	77
5.3	Results and Discussions	78
5.3.1	Optimization Results with and without Compound Mask	78
5.3.2	Optimization Results with Different Number of Virtual Sources	79
5.3.3	Selection of Representative Pareto-Optimal Solutions	79
5.3.4	Validation of the optimized distributions Using Silico Data	81
5.3.5	Validation of the optimized distributions Using Multi-purpose Phantom	83
5.4	Conclusion	85
6	Optimization-Driven GPU-Based receive Beamformer for Efficient and High-Quality Ultrafast Diverging Wave Imaging	87
6.1	Introduction	87
6.2	Methodology	88
6.2.1	RF-Band and Baseband Receive Beamformer	88
6.2.2	GPU-Implementation using CUDA kernels	91
6.2.3	Datasets for Experimental Validation	92
6.3	Results	93
6.3.1	B-mode Images for Multi-Purpose Phantom	93
6.3.2	Computational Complexity	94
6.4	Conclusion	96

7	Extending the compound mask apodization for 3D Ultrafast Imaging using 2D-matrix probe	99
7.1	Introduction	99
7.2	Methodology	102
7.2.1	Compound Mask Computation for 3D Divergent Wave Imaging . .	102
7.2.2	Computation of Weights based on virtual Aperture	106
7.2.3	Static Hanning window apodization for each virtual source for Comparison	107
7.2.4	Simulation Setup	108
7.3	Results	109
7.3.1	3D Contouring of the Compound Mask Results for a fixed image point	109
7.3.2	PSF Simulation using 32x32 2D matrix probe	111
7.3.3	PSF Simulation using 32x32 2D matrix probe with 25% random sparsity factor	111
7.4	Discussion	113
7.4.1	Analysis of Results	113
7.4.2	Limitations and Future Directions	116
7.5	Conclusion	116
8	Conclusion and Future Work	119
8.1	Conclusion	119
8.2	Future Directions	121
	Bibliography	125

List of Figures

2.1	Schematic of backscattering. (1) Outgoing wave emitted by the source; (2) incident wave at the target (scatterer); (3) incident wave together with the backscattered echo following the target's response.	10
2.2	Schematic of a linear transducer array showing element width, height, pitch, and kerf, with the (x, y, z) axes indicated [1]	11
2.3	Principle of conventional ultrasound: transmission of focused waves, receive focusing to form an image column (RF column), repeated to build the entire image [2].	12
2.4	Example of measurement of lateral resolution using a PSF	18
2.5	Contrast measurement on an image of a cyst phantom [3].	19
2.6	Conventional focused and ultrafast ultrasound imaging sequences in a typical medical setup: (a) conventional focused imaging (≈ 25 fps), (b) single plane-wave imaging (≈ 18 kfps), (c) coherent plane-wave compounding with 17 angles (≈ 1 kfps), and (d) compounding with 40 angles (≈ 350 fps). Image quality improves as more sub-images are coherently compounded, at the expense of frame rate. Adapted from [4].	24
2.7	Schematic of a diverging (spherical) transmit formed by a virtual point source (black dot) positioned behind a phased-array transducer. Element-wise transmit delays are chosen to synthesize a wavefront that spans the shaded region of interest, characterized by angular aperture β and steering angle θ , given the probe aperture[5].	25
2.8	Steps of image formation by compounding using plane waves([2]).	26
2.9	Example of a convex function (a) and a nonconvex function (b).	29

2.10	Dominance of solutions in the objective space (bi-objective minimization). The blue step curve indicates the non-dominated (Pareto) front. Points marked with circles are non-dominated; crosses are dominated. Region A illustrates a heavily dominated area (worse in both objectives), while region B marks the ideal improvement zone.	34
3.1	Comparison of SA Linear and Convex Probe Configurations with distances between image point (x, z) and probe elements coordinates in transmission (x_t, z_t) and reception (x_r, z_r)	40
3.2	Comparison of DW Linear and Convex Probe Configurations with distances between virtual source (x_i, z_i) and image point (x, z) and probe elements coordinates in transmission (x_t, z_t) ; moreover also the distance between image point and probe elements coordinates in reception (x_r, z_r) is displayed.	42
3.3	Examples of values of $H(i, t)$ matrix in Eq. 7. Row and column indexes are related to virtual sources and probe elements indexes, respectively	43
3.4	A schematic illustrating the generation of the compound mask for Diverging Wave Imaging (DWI). Starting from each image point, a continuous transmit apodization curve is first defined over the probe surface, following the principles of Synthetic Aperture Imaging (SAI). Then, a set of virtual source (VS) positions is specified behind the probe, which can follow arbitrary spatial distributions. For each VS, a straight-line trajectory is constructed toward the image point, and the corresponding intersection with the probe surface is determined. At these intersection points, the apodization curve is sampled, and the resulting values are assigned as weights to the respective virtual sources. The same methodology applies to convex arrays by replacing the linear probe intersection with angular (θ_E) intersections and is valid for every image point regardless of the virtual source distribution.	45
3.5	Comparison of transmit apodization in Synthetic Transmit Aperture Imaging (STAI) and virtual source-based apodization in Diverging Wave Imaging (DWI). (a) Transmit apodization $v(x_T)$ in STAI for Hanning, Blackman, and Tukey (0.5 roll-off) windows. (b) The corresponding virtual source-based apodization $w(x_i, z_i)$ in DWI, computed using the proposed method, for a fixed point located at $(0, 20)$ mm. The virtual sources are placed linearly behind the linear probe at a distance of 20 mm.	48

3.6	Two-dimensional representation of (a) transmit apodization $v_t(x, z)$ in Synthetic Transmit Aperture Imaging (STAI) using Hanning, Blackman-Harris, and Tukey (0.5 roll-off) windows as baseline apodization, and (b) the corresponding virtual source-based apodization $w_t(x, z)$ in DWI Linear Array with the virtual sources are placed linearly behind the linear probe at 20 mm. Both maps are displayed for $x = 0$	49
3.7	PSF of SALA, DWLA, SACA, and DWCA for a point at 20 mm depth (first row), and 30mm depth (second row) with F-number=0.7	50
4.1	Spatial arrangements of virtual sources (VSs) used in the experimental validation. (a) Linear distribution for the linear array, with the virtual sources placed at a depth $z_d = 20$ mm above the probe. For the convex array, three configurations are shown: (b) Linear arrangement, where the virtual sources are distributed linearly; (c) Curvilinear arrangement, where the virtual sources follow the curvature of the probe, with angular placement confined within a maximum aperture angle of 33° ; and (d) Tilted arrangement, where the virtual sources are placed along a radial line at a depth equal to R_{convex} . In all configurations, all probe elements are used in both transmission and reception, with $N = 20$ virtual sources used for each arrangement.	56
4.2	Example of regions of interest (ROIs) used for contrast ratio (CR) computation. Green circles indicate the regions inside the cysts used to estimate μ_{in} , and the orange circle indicates the background region used to estimate μ_{out}	57
4.3	B-mode images generated using the L7-4 linear array probe and Verasonics scanner for Diverging Wave Imaging (DWI), and DWI with the proposed compound mask. The images correspond to the sequences described in Figure 4.1(a). The first row displays the resolution measurements, while the second row shows the imaging results for three greyscale cysts (anechoic, -6 dB, +6 dB) located at a depth of 45 mm.	58
4.4	a)Boxplots of the Full Width at Half Maximum (FWHM) measured on the five reflectors embedded in the multi-purpose phantom, shown in Figure 4.3. The FWHM values are compared for Diverging Wave Imaging (DWI) and DWI with the proposed compound mask using linear array(L7-4).b)Zoomed-in Figure for the Cysts presented in 4.3. From left to right anechoic cyst, -6dB cyst, +6dB cyst. The number under each cyst is the measured contrast ratio.	59

4.5	B-mode images generated using the C5-2 convex array probe and Verasonics scanner for Diverging Wave Imaging (DWI), and DWI with the proposed compound mask. The DWI images correspond to the three transmission strategies described in Figure 4.1: (a) linear distribution, (b) curvilinear distribution, and (c) tilted distribution.	60
4.6	Boxplots of the Full Width at Half Maximum (FWHM) measured on the five reflectors embedded in the multi-purpose phantom, corresponding to the images in Figure 4.5 using a convex array. The FWHM values are compared for Diverging Wave Imaging (DWI) and DWI with the proposed compound mask. The DWI results are further categorized by the three transmission strategies: linear, curvilinear, and tilted.	61
5.1	Overview of the Multi-Objective Genetic Algorithm (MOGA) for optimizing virtual source (VS) distributions. The algorithm initializes a population of VSs distributions within a predefined search space and evaluates them using two objective functions. Evolutionary operations include tournament selection ($k=4$), crossover (60%), and adaptive mutation to refine the solutions iteratively. The process continues until convergence criteria are met, yielding an optimized set of VS configurations	71
5.2	Study configuration from top to bottom: the arrangement of VSs, where the orange VSs indicate those selected for optimization; a convex probe with N elements; scatterer positions, with grey ones used in the optimization process; and the imaging grid defining the region of interest.	74
5.3	Impact of the compound mask on the optimization of Virtual Source (VS) distributions using a Genetic Algorithm (GA):(a) Pareto front showing the average values of Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL) obtained from GA optimization without the compound mask.(b) Pareto front obtained from GA optimization with the compound mask . . .	77
5.4	Pareto-optimal front for different numbers of Virtual Sources (VSs) in the optimization process: 20 VSs (blue), 10 VSs (orange), 6 VSs (purple), and 4 VSs (green).	80

5.5	The Pareto front for 20 VSs (a) includes three selected solutions: optFWHM (red), midopt (dark blue), and optPSL (light pink) alongside deterministic configurations: 20VS-linear (green square) and 20VS-curvilinear (orange square). (b) represents the Pareto fronts for reduced VS configurations (10, 6, and 4 VSs) feature two selected solutions— midopt and optPSL —for each case. The selected solutions for 10 VSs (orange and red), 6 VSs (purple and blue), and 4 VSs (green and yellow) respectively	81
5.6	Spatial distribution of VSs for both deterministic and optimized distributions chosen for validation: (a) 20VS-Lin, (b) 20VS-Curv, (c) 20VS-optFWHM, (d) 20VS-midopt, (e) 20VS-optPSL, (f) 10VS-midopt, (g) 10VS-optPSL, (h) 6VS-midopt, (i) 6VS-optPSL, (j) 4VS-midopt, (k) 4VS-optPSL.	81
5.7	B-mode images of the simulated phantom with anechoic cysts and reflectors for (a) 20VS-Lin, (b) 20VS-Curv, (c) 20VS-optFWHM, (d) 20VS-midopt, (e) 20VS-optPSL, (f) 10VS-midopt, (g) 10VS-optPSL, (h) 6VS-midopt, (i) 6VS-optPSL, (j) 4VS-midopt, and (k) 4VS-optPSL.	82
5.8	Boxplots for (a) FWHM measured on the 9 reflectors (b) CR measured on 9 anechoic cysts from the phantom images (Fig. 5.7) for all selected VSs.	82
5.9	B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for left: 20VS-Lin, right: 20VS-Curv.	84
5.10	B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for 20VS-optFWHM, 20VS-midopt, 20VS-optPSL	84
5.11	B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for 10VS-midopt, 4VS-midopt	85
5.12	Boxplots for FWHM measured on the 6 reflectors from the phantom images (Fig. 5.9, 5.10 and 5.11) for selected VSs	86

6.1	Block diagram of the GPU-based receive beamformer for ultrafast diverging-wave imaging. The workflow illustrates the integration of offline pre-computations (compound mask, receive aperture ranges, apodization profiles, transmit and receive delays) with the online beamforming pipeline. Pre-computed data are efficiently translated along channels and combined to generate focalization terms. Depending on the beamforming mode, the input data are either shifted to the RF band (RF beamforming) or processed directly in baseband (I/Q signals) with carrier delay multiplication. Local linear interpolation ensures accurate delay alignment, followed by element-wise apodization weighting and summation across channels. Compounded images are obtained by combining multiple transmissions, ultimately producing the final high-quality B-mode image.	89
6.2	Virtual source distributions used for experimental validation of diverging-wave imaging. Five transmit configurations are compared: 20 virtual sources (VSs) arranged in a conventional curvilinear pattern, 20 VSs optimized, 10 optimized VSs, 6 optimized VSs, and 4 optimized VSs.	93
6.3	Beamformed B-mode images for the multi-purpose tissue-mimicking phantom using the GPU-based receive beamformer under different virtual source (VS) distributions. First row (20,10,6 and 4)VSs arranged in a curvilinear deterministic distribution, second row the same number of VS with the optimized distributions presented in figure 6.2.	94
6.4	Image quality assessment of GPU-based receive beamforming under varying numbers of virtual source (VS) transmissions for deterministic and optimized distributions. (a) Lateral resolution, expressed as full-width-at-half-maximum (FWHM), was measured along the central axis at depths of 30, 40, and 60 mm. (b) Contrast ratio (CR) measured for an anechoic cyst located at 55 mm depth, expressed in decibels(dB).	95
6.5	Memory requirements for GPU-based receive beamforming as a function of the number of transmissions. (a) Raw data storage, which scales linearly with the number of transmissions, highlights the increasing data transfer demand from acquisition hardware to GPU memory. (b) Breakdown of focalization data memory usage, including received delays, apodization coefficients, compounded mask, and transmitted delays.	96
7.1	Geometry of 3D DWI showing virtual sources (blue dots), probe elements (black grid), and image voxel (red point) with relevant distances D_i , D_{it} , D_t and D_r annotated.	103

- 7.2 Illustration of the virtual aperture concept in 3D Diverging Wave Imaging. **Left:** rays from virtual sources (blue) intersect the physical aperture (black) for voxels within the central field of view, defining valid compound mask weights. However, voxels near the periphery may yield intersection points outside the probe footprint. **Right:** introduction of a virtual aperture (red) surrounding the physical aperture ensures that all voxel–source rays intersect a valid region, allowing weights to be computed and stored. 106
- 7.3 Static 2D Hanning apodization profile applied to virtual sources. **(a)** 2D Hanning surface defined on an extended 7×7 grid, ensuring that apodization weights taper smoothly without vanishing at the edges. **(b)** Sampling of the weights at the positions of the actual 5×5 virtual source constellation. . . . 108
- 7.4 Comparison between transmit apodization in STAI and the mapped compound mask in DWI for two selected points in the 3D imaging volume. **(a)** Apodization profiles obtained from STAI at points $(0, 0, 20)$ mm and $(-1, 2.3, 21.4)$ mm are shown as continuous contour maps($\{(x_e, y_e)\}$ represents the probe elements). **(b)** Corresponding compound mask profiles in DWI, derived using the proposed compound mask mapping, demonstrate strong agreement with the STAI reference. The local structure of the apodization, including its smooth tapering and symmetry, is preserved in the mapped DWI profiles. 110
- 7.5 **PSF images (xz slice at $y = 0$; full 32×32 array).** Rows indicate VS layout; columns indicate VS weighting strategy. **Row A:** *Grid distribution* (5×5) located at $z_{VS} = -10$ mm. **Row B:** *Concatenated rings + central VS* with $r_{outer} = 5$ mm, $r_{inner} = 2.5$ mm, $z_{VS} = -10$ mm. **Columns:** (1) no weights, (2) static apodization, (3) compound-mask apodization. Images are envelope-detected and log-compressed to a 60 dB display range. 112
- 7.6 **Lateral PSF profiles (full 32×32 array).** **(a)** 5×5 uniform grid over the full array aperture. **(b)** *Concatenated rings + central VS*:, all VS located at $z_{VS} = -10$ mm behind the probe. Four on-axis point scatterers at 5, 20, 40, and 65 mm (colors match the depth colorbar). Curves are *peak-normalized* to 0 dB, Line styles denote the VS weighting: **dashed** = no weights, **dotted** = static apodization, **solid** = compound-mask apodization. 112

- 7.7 **PSF images (xz slice at $y = 0$; 32×32 with 25% random active elements).** Rows indicate VS layout; columns indicate VS weighting. **Row A:** *Uniform grid* (5×5) located at $z_{VS} = -10$ mm. **Row B:** *Concatenated rings + central VS* with $r_{outer} = 5$ mm, $r_{inner} = 2.5$ mm, $z_{VS} = -10$ mm. **Columns:** (1) no weights, (2) static apodization, (3) compound-mask apodization. Images are envelope-detected and log-compressed to a 60 dB display range. . . . 114
- 7.8 **Lateral PSF profiles (32×32 with 25% random active elements).** (a) *Grid distribution* of virtual sources (VS), 5×5 uniformly over the full aperture. (b) *Concatenated rings + central VS:*, all VS located at $z_{VS} = -10$ mm behind the probe. Four on-axis point scatterers at 5, 20, 40, and 65 mm (colors match the depth colorbar). Curves are *peak-normalized* to 0 dB, Line styles denote the VS weighting: **dashed** = no weights, **dotted** = static apodization, **solid** = compound-mask apodization. 115

List of Tables

2.1	Speed of sound and attenuation in biological media	10
2.2	Focused HFR strategies: qualitative comparison.	21
3.1	Comparison of FWHM values at 20 mm and 30 mm for different techniques with F-number 0.7	50
4.1	Contrast ratio (CR) values (in decibels) for three nominal target contrasts (anechoic, -6 dB, +6 dB) DWI, and DWI with the compound mask using linear array (L7-4).	60
4.2	Contrast ratio (CR) values (in decibels) for greyscale targets (anechoic, -6 dB, +6 dB) using various virtual source distributions using a convex array (C5-2).	62
4.3	Computational complexity (per frame) for common apodization strategies in diverging/plane-wave compounding. Here K is the number of image pixels, N the number of active receive channels, and M the number of transmissions.	63
4.4	Paired one-sided sign test for FWHM (masked vs. unmasked) combining all 20 paired measurements from linear and convex probe experiments. Negative differences indicate improvement (masked < unmasked).	64
5.1	CR values measured on anechoic cyst located at 45mm from the phantom images(Fig.5.9,5.10 and 5.11)for selected VSs.	85
6.1	Processing time for RF-band and Baseband beamforming with different numbers of transmissions.	96

Chapter 1

Introduction

1.1 Motivation

Ultrasonography is one of the most widespread medical imaging modalities due to its portability, low cost, and safety because of the absence of ionizing radiation[6]. In conventional 2D ultrasound imaging (US), images are formed by transmitting sequential focused beams and reconstructing the image line-by-line to produce high-quality images, but at the expense of frame rate that hardly exceeds the order of tens of frames per second[7][8]. To reduce the number of transmissions and improve the frame rate, ultrafast imaging using plane waves (PWs)[9] or diverging waves (DWs)[10][11] has been introduced. In PWs imaging, transducer elements are excited in order to produce a locally plane wavefront that insonifies all the area of interest. Since one image frame can be in principle obtained with just one transmission, this approach allows an increase of roughly two order of magnitudes in frame rate with respect to line-by-line insonification (up to thousands of frames per second). The loss of image quality, due to the absence of transmit focalization, can be recovered by sequentially transmitting a set of plane waves with different steering angles and coherently compounding the corresponding reconstructed frames. DWs imaging is based on a similar concept but plane waves are replaced by diverging waves, typically characterized by approximately circular wavefronts in the image plane. Diverging waves are generated by setting transmit delays such that the resulting wavefront seems to originate from a virtual source typically placed behind the probe surface. The concept of steering angle related to PWs is replaced in DWs by the change in the coordinates of the virtual source. In both cases the underlying principle is to differentiate the local angle of the wavefront at every image point, as a function of transmission index, in order to recover, by coherent compounding in reception, an equivalent of the transmit focalization. The high frame rate capabilities, yielded by both PWs and DWs

imaging have paved the way to innovative applications such as shear wave elastography[12], ultrafast doppler imaging [13], and functional ultrasound imaging[14].

While PWs imaging is a well established research topic, DWs imaging has received comparatively less attention. Despite this, DWs bear some advantages with respect to PWs, like the ability to cover a wider Field of View (FOV) with less transmissions, and a better fitting with convex probe geometries since, on average the wavefront is locally more aligned with the direction of the peak of the radiation lobe of the probe elements, thus granting an higher sensitivity and a reduction of artifacts due to grating lobes. On the contrary, the main disadvantage of DWs is the faster power drop with propagation depth since the total energy is spread over an increasingly wider surface.

In DWs imaging, the improvement in image quality is generally achieved increasing the number of transmissions, i.e., the number of virtual sources (N_{VS}); however, this comes at the expense of frame rate. Therefore, achieving an high image quality while maintaining an high frame rate is a relevant research goal. Moreover, the position of the virtual sources with respect to the probe elements affects the overall divergent transmit field. Even if the number of virtual sources is fixed, there are countless ways to position them[15]. In the literature on virtual sources, they are generally distributed periodically above the probe, and it has been shown that this periodicity causes artifacts that degrade image quality [16]. Therefore, it is interesting to investigate non-periodic distribution of virtual sources that simultaneously offers better contrast and better spatial resolution than a regular distribution.

Another important aspect in unfocalized transmission, both with PWs and DWs, is the demanding requirement in terms of computational power and data transfer bandwidth. This is due to the combination of high transmission rate, easily reaching thousands of Hertz, and massive amount of beamforming operations in reception in order to reconstruct a whole frame per transmission event. Therefore, efficient implementations respecting the hardware/firmware constraints of real medical ultrasound equipment should be taken into consideration to guarantee real-time applications. While the above challenges are critical in 2D imaging, they become even more pronounced in 3D volumetric ultrafast imaging.

This thesis aims to improve image quality in DWI without affecting real-time performance and under hardware constraints. To this end, it focuses on optimizing the transmit apodization weights for each virtual source transmission and optimizing the distribution of these virtual sources to maximize image quality metrics, including both spatial resolution and contrast, even with a low number of transmissions. Furthermore, GPU-based implementations are developed to accelerate the beamforming and reconstruction processes, ensuring real-time execution with optimized transmission strategies. The methodologies proposed in this work can be extended to 3D volumetric imaging, enabling high-quality and computationally

efficient ultrafast ultrasound imaging systems suitable for advanced clinical and research applications.

1.2 Objectives and Contributions

The main contributions of this thesis are summarized as follows :

1.2.1 Optimize compound weights for each virtual source transmission with real-time implementation on verasonics scanner

In diverging wave imaging (DWI) the overall image quality can be optimized by acting on the multiplicative weights applied to post beamforming data during the coherent compounding phase. These weights are in general function of the image point coordinate and of the virtual source position. To our knowledge no closed form methods is available in the literature to compute optimal values for this weights. On the contrary, in Synthetic Aperture Imaging (STAI), a method based on consecutive transmissions from single probe elements and coherent compounding in reception, these weights can be easily derived from the standard transmit apodization windows currently adopted in line-by-line imaging [17]. Though STAI is considered a reference method in terms of spatial and contrast resolution it suffers from poor signal to noise ratio due low transmit energy emitted by each element. On the contrary DWI overcome this limitation by using all or a great part of probe elements in transmission to shape the divergent wavefront. Therefore, this study aims to optimize the weights in DWI and consists of two main parts:

- *Theory and Simulation:* we derive a closed-form mapping from STAI compound weights to DWI compound weights, and we refer to those weights as the compound mask. The mapping is obtained through a geometric warping process that preserves the local apodization values applied in STAI, however, associating them with virtual source positions rather than probe elements. The method is applicable to both linear and convex arrays. Comprehensive Field II simulations demonstrate that the proposed compound mask accurately reproduces the image quality characteristics of STAI
- *Experimental Validation using Verasonics* The proposed compound mask was integrated into the verasonics beamforming pipeline and experimentally validated using a tissue-mimicking phantom and multiple virtual source (VS) spatial distributions with both linear and convex probes. Compared to standard DWI without weighting,

the proposed method consistently improves image quality and introduces negligible computational overhead.

1.2.2 Multi-objective optimization of virtual source spatial distribution for ultrafast diverging wave imaging

This study focuses on enhancing the image quality of DWI by optimizing the spatial distribution of VSs using a multi-objective genetic algorithm (MOGA) and the previously developed compound mask approach. The work is divided into two main parts:

- *Optimization Framework and Simulation:* A multi-objective optimization framework based on genetic algorithms was developed to optimize the lateral and axial positions of virtual sources using convex arrays. The optimization simultaneously minimizes the Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL) to improve lateral resolution and contrast. Compound mask weights are integrated into the optimization process to adaptively adjust the contribution of each VS during transmission. Simulation results demonstrate significant performance gains, including a reduction in FWHM and improvement in contrast ratio (CR) compared to deterministic VS distributions, even with a low number of transmission events.
- *Experimental Validation:* The optimized VS distributions were validated experimentally using a Verasonics scanner and a convex array probe (C5-2). Results obtained from tissue-mimicking phantoms confirm that the optimized configurations outperform standard deterministic arrangements, enhancing both resolution and contrast while maintaining high frame rates.

1.2.3 Optimized GPU-based receive beamformer implementation for efficient and high-quality ultrafast diverging-wave imaging

This study develops a real-time receive beamforming pipeline on a GPU that preserves image quality while significantly reducing computational and memory costs. The work is organized into two parts:

- *GPU architecture and algorithmic optimizations:* A real-time beamformer is designed for RF-band and baseband (I/Q) data processing under commercial scanner constraints. The pipeline integrates: (i) offline precomputation of transmit/receive delays and apodizations; (ii) transmit–receive decoupling to convert 4D focalization data matrices into more compact 3D ones; (iii) symmetry-based translation of receive profiles to

further reduce receive focalization data from 3D to 2D size; (iv) transmission-wise parallelization on GPU (v) local linear interpolation with carrier-delay phase rotation (baseband) in order to avoid raw data oversampling; (vi) seamless integration of compound masks and optimized virtual-source (VS) layouts from previous optimization approaches.

- *Experimental validation and performance–quality trade-offs:* The GPU beamformer is validated using in-vitro datasets collected with deterministic and optimized virtual source distributions using Verasonics vantage system with a convex probe for tissue-mimicking phantoms and multiple VS configurations at 20/10/6/4 transmissions. The proposed GPU-based receive beamformer achieves real-time performance while maintaining image quality even when the number of virtual sources is significantly reduced.

1.2.4 Extension of compound Mask approach to 3D ultrafast imaging using 2D-matrix probe:

The final contribution extends the compound mask apodization framework from 2D imaging to fully 3D ultrafast diverging wave imaging using 2D matrix probes, establishing a key foundation for future volumetric ultrasound systems. A voxel-geometric mapping approach is introduced to compute 3D compound weights from the corresponding compound weights in 3D STA, enabling accurate and spatially consistent apodization across the 3D field of view. The method was validated through comprehensive FIELD II simulations using a matrix probe of varying numbers of active elements. This work lays the ground for future research in 3D ultrafast imaging, including sparse-array design and large-aperture optimization, paving the way toward real-time high-quality volumetric imaging with reduced hardware and processing requirements.

1.3 Thesis Outline

The structure of this thesis is as follows:

- **Chapter 2** presents a comprehensive review of ultrasound principles, existing ultrafast imaging modalities, problem statement and overview of optimization algorithms.
- **Chapter 3** Introduces the theoretical formulation of the compound mask concept and its derivation from the transmit apodization used in synthetic aperture imaging. The chapter details the methodology for computing the compound mask for both linear

and convex probes, describes the simulation setup, and validates the approach through simulation, demonstrating improved image quality.

- **Chapter 4** Describes the experimental setup for validating the compound mask approach using a Verasonics ultrasound system. It presents the results obtained with both linear and convex probes, demonstrating that the proposed apodization technique can be implemented in real time and achieves enhanced image quality compared to conventional DWI configurations.
- **Chapter 5** Focuses on the optimization of virtual-source distributions using a multi-objective genetic algorithm framework. It describes the optimization criteria, simulation and experimental setups, and performance metrics such as Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL). The results demonstrate that optimized virtual-source configurations yield superior image quality compared to deterministic distributions.
- **Chapter 6** Presents the development and implementation of a GPU-based receive beamformer capable of processing both RF-band and baseband data. The chapter details the CUDA-based parallelization strategies, evaluates computational performance, and provides quantitative comparisons in terms of memory occupancy, processing time using the optimization strategies developed in chapters 3,4, and 5.
- **Chapter 7** Extends the compound mask concept to three-dimensional Diverging Wave Imaging. It discusses the adaptation of the apodization strategy to volumetric imaging using 2D matrix probes and evaluates image quality through simulated 3D Point Spread Functions (PSFs). Comprehensive simulations on fully populated and sparse matrix arrays demonstrate that the proposed approach reduces sidelobe levels, establishing a solid foundation for future research on real-time volumetric ultrafast imaging and optimized sparse-aperture designs.
- **Chapter 8** Summarizes the main findings of the thesis, emphasizing the contributions to the field of ultrafast ultrasound imaging. It also outlines possible directions for future research.

1.4 Research Outputs

This dissertation has resulted in several research outputs, including international journal publications, conference papers, and a book chapter, as well as participation in international conferences listed as follows:

1.4.1 International Journals:

1. **Z. Alzein**, M. Crocco, and D. D. Caviglia, "Compound Mask for Divergent Wave Imaging in Medical Ultrasound," *IEEE Access*, vol. 13, pp. 59675–59686, 2025, <http://dx.doi.org/10.1109/ACCESS.2025.3557085>.
2. **Z. Alzein**, M. Crocco, D. D. Caviglia, and H. Liebgott, "Multi-Objective Optimization of Virtual Source Distributions for Ultrafast Diverging Wave Imaging," *Ultrasonics*, <http://dx.doi.org/10.2139/ssrn.5378898>.
3. **Z. Alzein**, H. Liebgott, M. Crocco, and D. D. Caviglia, "Virtual Source-Based Apodization for Divergent Wave Imaging: An Experimental Study" (submitted to *Ultrasonic Imaging*, under review).
4. **Z. Alzein**, M. Crocco, D. D. Caviglia, and H. Liebgott, "Transmit Apodization for 3D Diverging Wave Ultrafast Imaging" (preparing the first draft to be submitted to a Q1 journal).

1.4.2 Conference Papers:

1. **Z. Alzein**, H. Liebgott, M. Crocco and D. D. Caviglia, "Experimental Validation of Compound Mask for Enhanced Diverging Wave Ultrasound Imaging," 2025 IEEE International Ultrasonics Symposium (IUS), Utrecht, Netherlands, 2025, pp. 1-3, <http://dx.doi.org/10.1109/IUS62464.2025.11201552>.
2. **Z. Alzein** and D. D. Caviglia, "Optimization of Compound Weights in Ultrafast Ultrasound Imaging: An Experimental Study," 2025 IEEE International Ultrasonics Symposium (IUS), Utrecht, Netherlands, 2025, pp. 1-4, <http://dx.doi.org/10.1109/IUS62464.2025.11201824>.
3. **Z. Alzein** and D. D. Caviglia, "A Multi-Objective Optimization Framework for Compound Weights in Divergent Wave Ultrasound Imaging," 2025 *IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS)*, Madrid, Spain, 2025, pp. 987–990, <http://dx.doi.org/10.1109/CBMS65348.2025.00197>.
4. **Z. Alzein** and D. D. Caviglia, "Genetic Algorithm-Optimized Apodization for Ultrafast Plane-Wave Compounding," 2025 *IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS)*, Madrid, Spain, 2025, pp. 983–986, <http://dx.doi.org/10.1109/CBMS65348.2025.00196>.

5. **Z. Alzein**, A. Ibrahim, M. Crocco, M. Merlanti, and D. D. Caviglia, “Comparative Analysis of Complexity–Image Quality Trade-offs in Ultrasound Systems for Efficient Hardware Implementation.”<http://dx.doi.org/10.11159/icbes24.153>.
6. **Z. Alzein**, M. Crocco, and D. D. Caviglia, “Computationally Efficient RF Band and Baseband Beam-former for Coherent Plane Wave Imaging,” *2024 19th Conference on Ph.D Research in Microelectronics and Electronics (PRIME)*, Larnaca, Cyprus, 2024, pp. 1–4, <http://dx.doi.org/10.1109/PRIME61930.2024.10559667>.
7. A. Bassal and **Z. Alzein**, “Ultrafast Divergent Wave Imaging in 2D Echography: A Parametric Study,” *11th International Conference on Biomedical Engineering and Systems (ICBES)*, <http://dx.doi.org/10.11159/icbes25.197>.

1.4.3 Book Chapter:

1. **Z. Alzein**, A. Ibrahim, A. Wehbe, and D. D. Caviglia (2025). “Plane-Wave Ultrasound Imaging: Implementation and Evaluation of Different Interpolation Schemes.” In: Ruo Roch, M., Bellotti, F., Berta, R., Martina, M., Motto Ros, P. (eds) *Applications in Electronics Pervading Industry, Environment and Society. ApplePies 2024. Lecture Notes in Electrical Engineering, vol 1369*. Springer, Cham. https://doi.org/10.1007/978-3-031-84100-2_12.

1.4.4 Attendance at Conferences

1. Attended the *62nd IEEE International Ultrasonics Symposium (IUS)*, Utrecht, Netherlands, September 15–18, 2025. **(Two Poster Presentations)**.
2. Attended the *2025 Artimino Conference on Medical Ultrasound Technology*, Lyon, France. **(Oral Presentation)**.
3. Attended the *2025 IEEE 38th International Symposium on Computer-Based Medical Systems (CBMS)*, Madrid, Spain. **(Two Oral Presentations)**.
4. Attended the *IEEE 19th International Conference on Ph.D Research in Microelectronics and Electronics (PRIME)*, Larnaca, Cyprus. **(Oral Presentation)**.
5. Attended the *12th International Conference on Biomedical Engineering and Systems (ICBES)*. **(Virtual Presentation)**.

Chapter 2

Literature Review

This chapter provides an overview of the basic concepts of ultrasound in Section 2.1 and of the developments it has undergone over the past decades in Section 2.2 and 2.3, followed by a discussion and problem statement in Section 2.4. Since part of the work involves the use of stochastic optimization algorithms, a state-of-the-art review on this topic is provided in Section 2.5.

2.1 Fundamentals of Ultrasound

In this section, the fundamentals of ultrasound imaging are detailed, including principles of ultrasound imaging, image reconstruction, beamforming, and image quality metrics.

2.1.1 Principle of ultrasound imaging

Ultrasound imaging relies on the emission and reflection of high-frequency acoustic waves within a medium. Biological tissues are generally modeled as collections of microscopic scatterers, each capable of reflecting part of the incident ultrasound energy. A commonly used approximation considers these scatterers to re-radiate energy uniformly in all directions. To illustrate the wave propagation principle, consider a point source that emits an ultrasound pulse at time t_{emission} , radiating spherically throughout the medium. The pulse reaches a scatterer located at depth z_0 after a time delay, while part of the wavefront is reflected back toward the source. The backscattered echo is received at a later time $t_{\text{reception}}$ (see Figure 2.1).

Assuming a constant sound speed c within the medium, the depth of the scatterer can be estimated from the round-trip delay $\Delta T = t_{\text{reception}} - t_{\text{emission}}$ using the following equation:

$$z_0 = \frac{c \cdot \Delta T}{2} \quad (2.1)$$

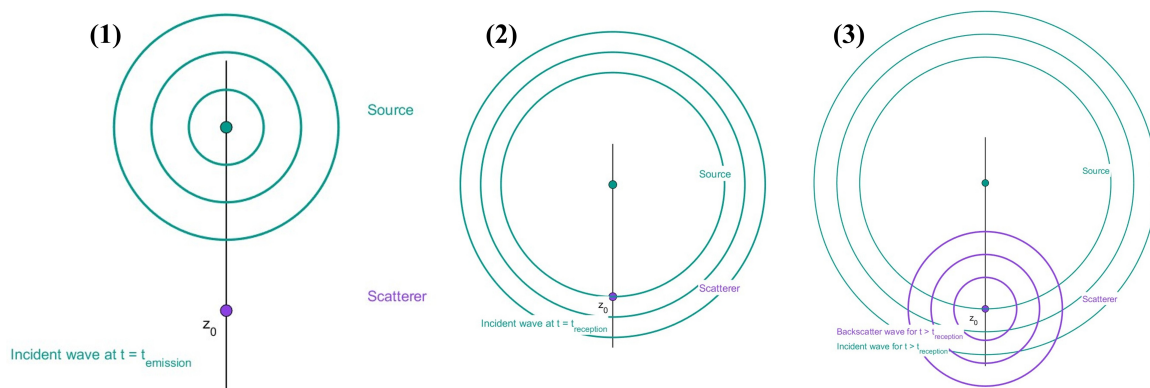


Fig. 2.1 Schematic of backscattering. (1) Outgoing wave emitted by the source; (2) incident wave at the target (scatterer); (3) incident wave together with the backscattered echo following the target's response.

This fundamental relationship is analogous to methods used in sonar systems for underwater ranging. In clinical practice, biological tissues consist of layers with varying acoustic properties, notably sound speed and attenuation. Table 2.1 presents typical acoustic parameters for common biological media [18].

Table 2.1 Speed of sound and attenuation in biological media

Medium	Speed of Sound (m/s)	Attenuation Coefficient (dB/MHz/cm)
Water	1480	0.002
Blood	1575	0.15
Fat	1450	0.6
Soft Tissue	1540	0.75
Muscle	1585	1.5
Bone	4080	10.0

Although the actual speed of sound varies with tissue type, a nominal value of $c = 1540$ m/s is often used for soft tissue imaging. Additionally, acoustic waves are attenuated as they travel through tissue, and the degree of attenuation increases with frequency. To compensate for this depth-dependent signal loss, clinical ultrasound systems implement *Time Gain Compensation (TGC)* to amplify echoes originating from deeper structures, assuming the signal-to-noise ratio (SNR) remains adequate. A characteristic feature of ultrasound images is the presence of *speckle*, a granular texture resulting from the interference of echoes from randomly distributed scatterers. Each scatterer emits a spherical wavefront upon insonification, and the overlap of these waves leads to constructive and destructive




Fig. 2.2 Schematic of a linear transducer array showing element width, height, pitch, and kerf, with the (x, y, z) axes indicated [1] .

interference patterns. While speckle can obscure anatomical detail, it also provides useful information about tissue texture and microstructure.

Modern ultrasound systems use transducers composed of multiple piezoelectric elements, collectively known as array probes. Each element can both transmit ultrasound pulses (by converting electrical energy into mechanical vibrations) and receive returning echoes (by converting acoustic signals back into electrical form). Ideally, elements would be infinitesimally small in order to emit a spherical wavefront. However, practical designs involve elements of finite width, separated by small gaps called *kerf*. The *pitch* is defined as the distance between the centers of two adjacent elements (element width + kerf), as illustrated in Figure 2.2. The finite size of each element determines a radiation pattern, function of the transmit/receive wavelength, with a peak in the direction orthogonal to the element surface and gradually decreasing as a function of the angle from the normal to the surface. The radiation pattern causes the element sensitivity to be variable as a function of the transmit/receive angle. A variety of probe geometries are used in clinical settings depending on the imaging application. For instance, linear arrays are commonly used for vascular and superficial imaging, while phased arrays are preferred for cardiac applications due to their small footprint and wide angular coverage where convex arrays are suited for abdominal application requiring wide fields of view.

2.1.2 Conventional Ultrasound (Focused Imaging)

In conventional ultrasound, the image is reconstructed line by line, where a “line” is defined by a starting point from the probe surface and a depth d along a given direction position




Fig. 2.3 Principle of conventional ultrasound: transmission of focused waves, receive focusing to form an image column (RF column), repeated to build the entire image [2].

x. The organization of the lines depends on the probe geometry: typical arrangements are on a cartesian grid for linear geometries or polar grid for convex geometries, but other arrangements are also exploited. Each image line is computed using a subset of probe elements that form an *aperture*. The steps of transmitting a focused wave, receiving echoes, and reconstructing an RF column are illustrated in Figure 2.3.

The transmission steps involves the activation of a subset of the probe elements, i.e. the transmit aperture, centered around the image lines starting point, and the application of a set of transmit delays that produce a transmit field converging on a point of the line at a given depth, named focal depth. The whole process is defined *transmit focalization* or transmit beamforming, and the resulting transmit field intensity has a typical hourglass-like shape ensuring that all the points along the line are insonified and the field is concentrated in the nearby of the line. In the reception phase, the echoes originating from the insonified scatterers arrive at each element probe following a path of different length that determines a propagation delay. Consequently, it is necessary to apply a set of receive delays in order to compensate the propagation delays and properly align signals coming from the region of interest, before summing them coherently to produce the signal related to a given position. The process takes the name of *receive focalization* or receive beamforming, more specifically delay-and-sum beamforming. In the first implementations of ultrasound scanners the set of receive delay was equal to set of transmit delays for a given line. This can be regarded as *static focusing*, to emphasize that the same delay is applied to all pixels in an i -th RF line (denoted RF_i^{post}), regardless of depth d . Equation 2.2 expresses this:

$$\text{RF}_i^{\text{post}}(d) = \sum_{n=n_{\text{start}}(i)}^{n_{\text{stop}}(i)} w_{ni} \text{RF}_{ni}^{\text{raw}}(t(d) - \tau_{ni}), \quad (2.2)$$

where $n_{\text{start}}(i)$ and $n_{\text{stop}}(i)$ are the first and last active elements for the receive aperture related to image line i , $\text{RF}_{ni}^{\text{raw}}$ is the signal received by element n and corresponding to transmission i , $w_n(i)$ is the weight multiplied to the raw RF signal of element n for the line i , $t(d)$ is the time of arrival of signal to the probe corresponding to the depth d , and τ_{ni} is the delay applied to element n for the line i . The weighting w_{ni} implements a spatial apodization window whose aim is typically to reduce the level of side lobes in the receive directivity pattern. These side lobes are due to abrupt edges of finite-size receive apertures and may affect the image quality, generating artifacts. A similar apodization can be applied on the transmit side, to prevent the formation of side lobes in the transmit pattern. The RF image—denoted $\text{IRF}(i, d)$ —is then reconstructed from the collection of N_l RF lines:

$$\text{IRF}(i, d) = \{\text{RF}_1^{\text{post}}, \dots, \text{RF}_i^{\text{post}}, \dots, \text{RF}_{N_l}^{\text{post}}\}, \quad (2.3)$$

The static receive focusing is not optimal since it guarantees a perfect signal alignment only at one focal depth (notice that both τ_{ni} and w_{ni} are function of column index i both not of depth d). Therefore, instead of summing the raw RF signals with static delays, each pixel of an RF column can be computed using the round-trip time-of-flight $T_{ni}(d)$ of the focused wave:

$$T_{ni}(d) = T_{\text{forward}} + T_{\text{return}}, \quad (2.4)$$

$$T_{\text{forward}} = \frac{|d|}{c}, \quad T_{\text{return}} = \frac{\|\vec{r}_i - \vec{r}_n^{EL}\|}{c}, \quad (2.5)$$

with $\vec{r} = (x, z)$ denoting the 2D coordinate of the image pixel and $\vec{r}^{EL} = (x^{EL}, z^{EL})$ denoting the coordinate of the probe element. The Time-of-flight is then used to reconstruct the RF column:

$$\text{RF}_i^{\text{post}}(d) = \sum_{n=n_{\text{start}}(i,d)}^{n_{\text{stop}}(i,d)} w_{ni}(d) \text{RF}_{ni}^{\text{raw}}(T_{ni}(d)), \quad (2.6)$$

This RF-column reconstruction is called *dynamic focusing*. since the delays applied in reception are function of the signal time of arrival, that in turn is function of the echo depth. It is worth noting that also the receive aperture and the apodization window can be made dynamic, in order to optimize the trade-off between lateral resolution and signal to noise

ratio at every depth. This process typically defines an aperture law that makes the aperture to increase as a function of depth.

The raw RF signals used to form the image have a frequency band centered around a carrier, therefore they appear to oscillate around zero value. The resulting $\text{IRF}(i, d)$ image therefore also contains zero-centered oscillations arising from the echoes. To make the image easier to read, *envelope detection* is applied to each RF column. In early scanners this was performed in analog hardware; today it is commonly done by applying the Hilbert transform to each image column and then taking the magnitude of the resulting analytic signal. This envelope detection is summarized in 2.7:

$$I_{\text{env}} = |\text{IRF} + i \mathcal{H}(\text{IRF})|, \quad (2.7)$$

where \mathcal{H} denotes the Hilbert transform performed along d , and i the imaginary unit. To emphasize weak echoes for display and avoid saturation of strong echoes a compression curve is applied. In the simplest case, the compression curve is a logarithm:

$$I_{\log} = \alpha \log_{10} \left(\frac{I_{\text{env}}}{I_0} \right), \quad (2.8)$$

where I_0 is the lowest value to be represented, corresponding to an output value of 0, and α is a scaling factor mapping the highest representable intensity to a predefined value, typically 255 for an 8-bit image. Since the image is computed along lines with arbitrary geometric organization it is needed a final step to map the values of I_{\log} function of line index i and depth d into a Cartesian image, a function of coordinates x and z . This geometric warping takes the name of *scan conversion* [19], and is achieved by local 2D interpolation.

Two key parameters define the timing constraints of conventional imaging:

- **Pulse Repetition Frequency (PRF):** The PRF sets the maximum number of transmission-reception cycles per second. It is limited by the maximum imaging depth z_{max} to ensure that returning echoes are received before the next pulse is emitted:

$$\text{PRF} = \frac{c}{2z_{\text{max}}} \quad (2.9)$$

- **Frame Rate (Frames Per Second, FPS):** The imaging frame rate depends on the PRF and the number of scan lines N_{cycle} required to form one image frame:

$$\text{FPS} = \frac{\text{PRF}}{N_{\text{cycle}}} \quad (2.10)$$

In clinical practice, conventional focused imaging achieves frame rates of up to approximately 100 FPS. While this is sufficient for most routine diagnostics, it may fall short in capturing fast transient phenomena such as rapid cardiac motion or high-frequency tissue dynamics.

2.1.3 Alternative receive beamforming

The receive delay-and-sum (DAS) beamformer with dynamic delays and aperture/apodization described in the previous section (see Eq. 2.6) is the prevailing baseline in diagnostic ultrasound imaging [20]. It is to note that all the sets of delays, apertures and apodization are data independent, i.e. they can be computed once, before the start of the beamforming process since they do not depend on the received raw data at the probe elements.

These features make DAS simple and computationally efficient. The drawback of DAS is its sensitivity to off-axis and sidelobe contributions that can degrade contrast, particularly under challenging apertures or sparse sampling[21],[22]. To address these limitations, the literature proposes a broad class of *adaptive*, *coherence-based*, and *non-linear* beamformers. For brevity, we will refer to these categories collectively as “adaptive beamformers” in the remainder of this chapter, noting that they differ in their underlying principles and computational cost.

Minimum-Variance (MV) and variants. The minimum-variance (MV) beamformer adaptively selects apodization weights to minimize the output power while preserving unit gain in the look direction, thereby sharpening the mainlobe and improving lateral resolution [23][24]. Practical implementations estimate the data covariance from delayed RF channels and require matrix inversion; both steps increase complexity and can impact robustness when sample support is limited. Moreover, despite its resolution gains, MV can exhibit reduced contrast in some scenarios—similarly to DAS—and is sensitive to sparse apertures and model mismatch[21],[22]. Several enhancements have been proposed: eigenspace-based MV (ESBMV) exploits the covariance eigenstructure to attenuate off-axis components and improve contrast [25], while forward–backward averaging (and related spatial averaging schemes) improves covariance estimation stability and MV robustness.

Coherence-based weighting. Coherence-based methods compute pixel-wise weights that modulate the DAS output according to the spatial coherence of delayed channels. The classical *coherence factor* (CF)[26][27] compares the coherent sum energy to the incoherent total energy across the aperture:

$$\text{CF}(i, d) = \frac{\left| \sum_{n=n_{\text{start}}(i, d)}^{n_{\text{stop}}(i, d)} \text{RF}_{ni}^{\text{raw}}(T_{ni}(d)) \right|^2}{N_{i, d} \sum_{n=n_{\text{start}}(i, d)}^{n_{\text{stop}}(i, d)} \left| \text{RF}_{ni}^{\text{raw}}(T_{ni}(d)) \right|^2}, \quad (2.11)$$

with $N_{i, d} = n_{\text{stop}}(i, d) - n_{\text{start}}(i, d) + 1$. After proper focusing, on-axis echoes tend to be phase-aligned (high CF), whereas off-axis/sidelobe contributions appear incoherent (low CF); multiplying the DAS image by CF thus suppresses sidelobes and improves contrast. The *generalized coherence factor* (GCF) extends CF by accounting for low spatial-frequency content, which can better preserve speckle statistics in some cases[28]. Beyond amplitude-only measures, the *phase coherence factor* (PCF) leverages instantaneous phase consistency across channels: mainlobe signals exhibit tight phase dispersion (PCF near one), while off-axis energy shows larger phase variability[29][30].

Non-linear compounding: DMAS and derivatives. Delay-Multiply-and-Sum (DMAS) is a non-linear approach that enhances contrast by emphasizing coherent channel pairs [31]. Instead of summing delayed channels directly, DMAS forms pairwise products before summation. A numerically stable formulation uses sign-preserving square roots to keep physical units and polarity:

$$y_{\text{DMAS}}(t) = \sum_{i=1}^{N-1} \sum_{j=i+1}^N \hat{s}_i(t - \tau_i) \hat{s}_j(t - \tau_j), \quad \hat{s}_i(u) = \text{sign}(s_i(u)) \sqrt{|s_i(u)|} \quad (2.12)$$

where $s_i(u)$ represents the pre-beamformed receive signal at channel i , τ_i is the propagation/focus delay for that channel, $\hat{s}_i(u)$ is the sign-preserving square-root mapped version of $s_i(u)$, t is the beamforming (focusing) time for the current image point, i and j index the N active receive channels (with $1 \leq i < j \leq N$), and $y_{\text{DMAS}}(t)$ is the DMAS beamformed sample. Because the multiplication introduces a DC term and second harmonics, band-pass filtering is typically applied, yielding filtered DMAS (F-DMAS)[31][32]. The naive double sum has $\mathcal{O}(N^2)$ complexity; an algebraic reformulation reduces it to $\mathcal{O}(N)$ operations[33]:

$$y_{\text{DMAS}}(t) = \frac{1}{2} \left[\left(\sum_{i=1}^N \hat{s}_i(t - \tau_i) \right)^2 - \sum_{i=1}^N |\hat{s}_i(t - \tau_i)|^2 \right] \quad (2.13)$$

Several variants tailor DMAS to specific modalities, e.g., double-stage DMAS for photoacoustics [34] or DwMAS, which integrates weighting windows to further suppress sidelobes [35].

In summary, DAS remains the reference due to its simplicity and speed, but off-axis sensitivity motivates more advanced strategies. MV-type methods improve resolution at a higher cost and with contrast sensitivities; coherence-based weights and non-linear compounding offer robust contrast gains.

2.1.4 Image quality in Ultrasound

The images obtained are more or less able to help clinicians to make a diagnosis. Talking about image quality directly relates to this ability to help the clinician identify abnormalities, such as breast lesions or calcifications in the kidneys. Studies have been conducted to determine the criteria that reflect good image quality as perceived by clinicians [36]. The study indicates that this notion is highly subjective: it depends on the individual and their specialty. In parallel, quality criteria have been established to evaluate ultrasound performance objectively. However, some studies show that certain objective criteria are not very representative of the subjective criteria used by healthcare professionals [37][38]. Despite these disparities, objective metrics such as spatial resolution and contrast still make it possible to benchmark ultrasound performance.

2.1.4.1 Spatial Resolution

Spatial resolution can be described as the ability to delineate a fine structure in an image, or even to separate fine structures from one another when they are close together. In optical imaging, resolution is introduced with the Rayleigh criterion: it is the ability to distinguish two luminous objects in an image. This criterion is also used in ultrasound imaging. Another measure often used to quantify the spatial resolution of an imaging system is the full width at half maximum (FWHM)[39] of the image of an isolated point-like scatterer. This type of image is representative of the point spread function (PSF) of the system. For log-compressed images, the lobe width is often measured at 6 dB (Figure 2.4)[3].

The smaller this width, the finer the spatial resolution according to this criterion. This measurement can be made axially (along the probe axis, i.e., vertically in an image) or laterally (horizontally in an image).

2.1.4.2 Contrast

In general, *contrast* refers to the difference in brightness between a bright region and a dark region in an image. In ultrasound, a *hyperechoic* region (which produces many echoes when traversed by an ultrasonic wave) appears with higher gray-level intensity than a *hypoechoic* region (which produces few echoes). Contrast can be assessed on tissue-mimicking *phantoms*

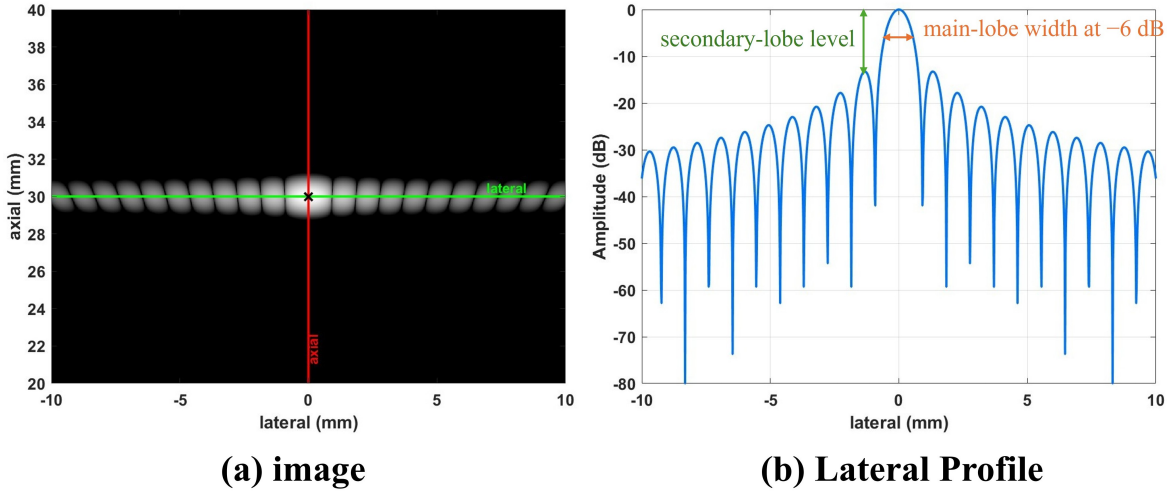


Fig. 2.4 Example of measurement of lateral resolution using a PSF

that reproduce the texture of cysts. A “phantom” is a test object that mimics the mechanical properties of various biological tissues and is used to evaluate imaging system performance. Figure 2.5 shows an ultrasound image of a cyst phantom [3].

The medium contains several *anechoic* cysts (which do not return echoes), outlined by yellow circles. These cysts appear dark compared with the surrounding medium when the latter is hyperechoic. In the ultrasound literature, the contrast between an anechoic cyst and its hyperechoic background is commonly quantified using the *Contrast Ratio* (CR, Eq. (2.14)) and the *Contrast-to-Noise Ratio* (CNR, Eq. (2.15)) ([39]; [3]):

$$CR = \frac{|\mu_{\text{bck}} - \mu_{\text{cyst}}|}{(\mu_{\text{bck}} + \mu_{\text{cyst}})/2}, \quad (2.14)$$

$$CNR = \frac{|\mu_{\text{bck}} - \mu_{\text{cyst}}|}{\sqrt{(\sigma_{\text{bck}}^2 + \sigma_{\text{cyst}}^2)/2}}, \quad (2.15)$$

where μ and σ denote, respectively, the mean and standard deviation of pixel values measured over the hypoechoic and hyperechoic regions. The abbreviation “bck” (*background*) refers to the cyst background, i.e., the hyperechoic area (delimited by green circles in Fig.2.5), while “cyst” refers to the interior of the cyst considered for measurement (delimited by red circles). CR (Eq. 2.14) measures contrast, whereas CNR (Eq. 2.15) captures both contrast and the influence of noise.

Sidelobes of an ultrasound array, in both transmit and receive, contribute to contrast degradation: a high sidelobe level brightens otherwise dark hypoechoic areas ([?], pp. 193–194). Here, “sidelobe” denotes any lobe other than the mainlobe, including *grating lobes* (caused




Fig. 2.5 Contrast measurement on an image of a cyst phantom [3].

by the spatial periodicity of array elements) and *side lobes* arising from unapodized apertures. The link between sidelobe level and contrast is sometimes exploited to *estimate* contrast. For example, [40] proposed including the sidelobe level in a cost function to be minimized by an optimization algorithm with the goal of improving image contrast. A PSF can also be used to measure sidelobe levels, as illustrated in Figure 2.4. While CR and CNR remain closer to the general definition of contrast (difference in average brightness), it has been shown that certain image reconstruction algorithms modify the image dynamic range, which can lead to overestimation of contrast measured with Eqs. 2.14 and 2.15 ([41]).

To address this sensitivity, a contrast metric called the *generalized Contrast-to-Noise Ratio* (gCNR) was introduced ([42]), designed to be robust to dynamic-range alterations. In its baseline form:

$$\text{gCNR}_0 = \frac{C_0^- - C_1^-}{C_0 - 1}, \quad (2.16)$$

with

$$C_0 = \frac{3}{2N_{\text{El}} \text{SNR}_1 + 3}, \quad (2.17)$$

where N_{EI} is the number of elements used in transmit/receive, and SNR_1 is the single-channel electrical signal-to-noise ratio:

$$\text{SNR}_1 = \frac{v_S}{v_N}, \quad (2.18)$$

with v_S the average speckle power on a single channel and v_N the average uncorrelated noise power on that channel. Here, “speckle” denotes the received signal produced by echoes from a sufficiently dense collection of scatterers such that individual reflections are not fully separated in the channel data. According to [42], gCNR provides a faithful measure of lesion detectability in ultrasound images.

2.2 Advanced Focusing Strategies

The standard beamforming approach described till now follows a serial, line-by-line acquisition: each transmit–receive cycle forms (at most) one scan-line, and this process is repeated until a full frame (often $\approx 10^2$ lines) is completed [43]. This approach has been the reference method in commercial equipments for several years. However, this approach suffers from two fundamental limitations: image non-uniformity and limited frame rate. The fixed focus in transmission results in a variable lateral resolution along depth, with a maximum around the focal depth and gradually worsening far from it. The line-by-line insonification and single line reconstruction severely limits the frame rate: given a maximum depth D and a number of TX/RX lines N_ℓ the maximum frame rate is given by $FR = c/(2DN_\ell)$. Considering that some nonlinear techniques, like Pulse Inversion or Contrast Pulse Sequence (CPS) require two or three subsequent transmissions to produce a single line, the frame rate can be further reduced, especially for anatomical districts, like abdomen, in which the maximum depth is particularly high. Moreover the limited frame rate is particularly critical in cardiological applications in which it is necessary to monitor dynamic phenomena. Finally, line-by-line approach renders real-time volumetric (3D/4D) imaging impractical [44] due to the sheer number of lines to be reconstructed. To address these issues, in the last decade research has enabled the development of more advanced techniques, aimed at decoupling frame rate from the number of sequential shots while preserving diagnostic image quality [45]. We will first review advanced methods based on focalized transmission and subsequently method based on unfocalized transmission specifically devised for ultrafast imaging.

Multi Line Acquisition (MLA) A straightforward method to increase frame rate is to enlarge the width of the transmit beam, by reducing the transmit aperture, to cover a wider area with one transmission and reconstruct multiple parallel line from the same receive data.

Table 2.2 Focused HFR strategies: qualitative comparison.

Strategy	Frame-rate gain	Image quality gain	Dominant limitation / baseline
MLA [48]	Multiple line reconstructed with a single transmission	No particular improvement	Potential artifacts at boundaries between receive line groups
RTBF [43]	Multiple lines reconstructed with a single transmission	Improved image uniformity along depth;	tight trade-off between improvement in image quality vs frame rate.
SAMF / SAMI [50, 51]	Improvement wrt sequential transmissions at different focal depths. No improvement wrt standard beamforming	Improved image uniformity along depth;	Depth-uniformity trade-off at constant frequency; addressed by FM-SAMI [51]
MLT [52]	Multiple simultaneous transmit focused lines. Up to $\sim 8 \times$ FR gain	No particular improvement	Inter-beam crosstalk; mitigated by F-DMAS and harmonic imaging [53, 52]

The so-called multi-line acquisition [46–48] allows to increase the frame rate by a factor given by the number of parallel receive lines but this comes at cost of reduction in image quality, especially due to artifacts visible in the image when moving across different groups of parallel receive lines.

Retrospective transmit beamformation (RTBF). RTBF performs transmit focusing in post-processing: information encoded by a standard fixed-focus transmission field is re-focused retrospectively by coherently summing contributions from different transmissions with overlapping beams, after proper phase corrections in order to align the signals coming from the same spatial location. The method can achieve near two-way dynamic focusing across the field, improving resolution and sidelobes, especially away from the transmit focal depth [43, 49]. Variants of RTBF underlie commercial coherent volume formation pipelines [43].

Simultaneous Axial Multifoci (SAMF) SAMF (a.k.a. SAMI) forms multiple axial foci from a *single* transmission by superposing appropriately delayed waveforms on each element, enabling focused imaging over a wide depth range with standard DRF on receive [50]. A key limitation is depth non-uniformity when a single center frequency is used: higher frequencies

provide resolution but attenuate at depth, while lower frequencies penetrate but blur laterally [51]. Frequency-multiplexed SAMI (FM-SAMI) assigns higher frequency to shallow foci and lower frequency to deeper foci, flattening resolution across depth while preserving penetration, and has been demonstrated in real time [51, 54].

Multi-Line Transmission (MLT) MLT transmits N_{MLT} focused beams in parallel to achieve an N_{MLT} -fold frame-rate gain [44]. Its central challenge is crosstalk among simultaneously transmitted beams, especially in the near field and in richly scattering tissue. For example, in [53], the method yielded perturbations around ~ -40 dB (up to ~ -27 dB in vivo) at the fundamental frequency. Nonlinear beamforming—Filtered Delay-Multiply-and-Sum (F-DMAS)—suppresses crosstalk by exploiting spatial coherence, improving contrast ratio (e.g., ~ 8 dB over DAS in 4-MLT) while narrowing main-lobe and lowering sidelobes [52, 55, 56]. Second-harmonic imaging further reduces crosstalk to ~ -60 dB (or ~ -45 dB in vivo) [53].

Summary Advanced focused HFR strategies address the limitations in image uniformity and frame rate of traditional beamforming, trying to improve one of the two features or finding a better trade-off between them. MLA and MLT offer pragmatic speedups with their own artifact and engineering trade-offs. SAMF (with FM-SAMI) improves image uniformity but keep unchanged the frame rate. RTBF kilely provides the best trade-off between frame-rate and image quality. In practice, the combination of CE, coherence-aware/nonlinear beamforming, frequency-domain reconstruction, and GPU/FPGA acceleration with optimized data movement has closed much of the historical gap, enabling real-time focused HFR with preserved diagnostic quality [45, 57, 58]. Critically, the principal limitation of all the approaches based on focused transmission is that the field related to each transmission event covers only a limited portion of the region of interest and thus cannot approach the ultrahigh (kHz-scale) rates. Theses frame rates are maybe not necessary for standard B-mode maging but are of great benefit and sometimes mandatory for advanced applications like high-sensitivity Doppler imaging [59], transient elastography [60], and vector flow imaging [61].

2.3 Unfocalized transmission methods for Ultrafast Ultrasound Imaging

To overcome the frame–rate limitations inherent to focalized transmission, *ultrafast* ultrasound replaces many focused transmits, each one covering a limited portion of the image

area, with fewer unfocalized transmissions covering all or great part of the image area. In the following sections, we detail *Synthetic Transmit Aperture (STA)*, *Plane Wave Imaging (PWI)* and *Diverging Wave Imaging (DWI)*. Moreover we discuss in more detail the concept of coherent compounding since it is a transversal approach to recover image quality while preserving high frame rate acquisition.

2.3.1 Synthetic Transmit Aperture (STA)

STA is considered a reference method providing, in ideal conditions, the best image quality achievable within the scope of linear data-independent approaches. This is due to the fact that a theoretically optimal transmit focusing is achieved at every image point. In STA every transmission is performed by a single element, or a very small group of them, generating a quasi-spherical wavefront that covers great part of the field of view (the limitation being the radiation pattern of the single element). The optimal focalization in transmission for every point is achieved virtually in reception by coherently compounding the contributions of every transmission with proper phase corrections, analogously to RTBF, but extending the recombination to much more different contributions [62, 63]. Coherent combination across N_e shots yields range-independent lateral resolution and two-way dynamic focusing over depth [62, 63]. Unfortunately STA suffer from poor penetration and low SNR due to the low energy emitted by each single element in every transmission. Practical reports show measurable SNR losses (e.g., several dB) and potential partial resolution degradation if energy is insufficient [45, 64]. Moreover STA is not per se an ultrafast method (if every element transmit separately the number of transmissions is the same order magnitude of line-by-line transmission) unless the number of transmissions is drastically reduced by sparsely sampling the transmit elements: unfortunately this further reduces the penetration and worsen the general image quality. Moreover the computational burden scales with the number of emissions, and explodes in 3D imaging, historically making image formation slower than data acquisition. [65]. Frequency-domain (wavenumber) reconstructions using 2D FFTs can be an order-of-magnitude faster than classical time-domain DAS while preserving resolution [58].

2.3.2 Plane Wave Imaging (PWI)

Plane Wave Imaging (PWI) enables ultrahigh frame rates by replacing sequential, line-by-line focused transmissions with one (or a few) unfocused plane-wave insonifications of the entire field of view, followed by parallel reception and reconstruction across all channels [66, 67]. Because the frame rate scales inversely with the number of transmit events, using

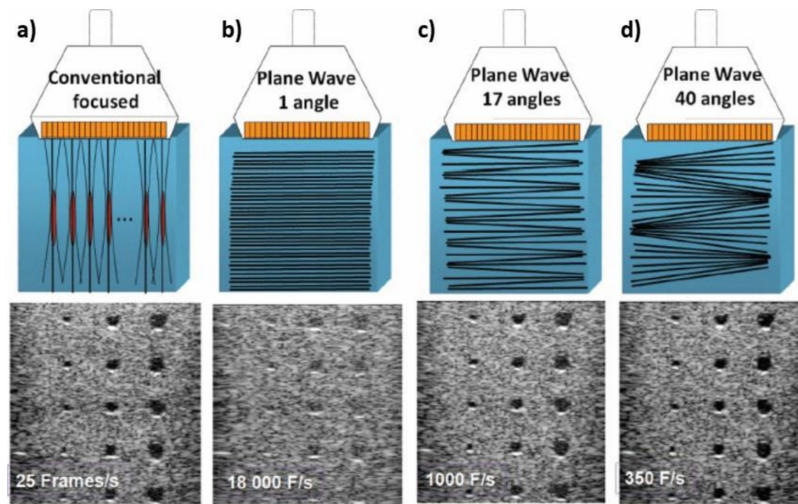


Fig. 2.6 Conventional focused and ultrafast ultrasound imaging sequences in a typical medical setup: (a) conventional focused imaging (≈ 25 fps), (b) single plane-wave imaging (≈ 18 kfps), (c) coherent plane-wave compounding with 17 angles (≈ 1 kfps), and (d) compounding with 40 angles (≈ 350 fps). Image quality improves as more sub-images are coherently compounded, at the expense of frame rate. Adapted from [4].

few plane waves allows frame rates in the kHz regime and underpins applications such as ultrafast vascular/functional imaging and quantitative ultrasound (QUS) [4, 67, 68]. A single plane-wave transmit lacks transmit focusing, producing a broad PSF with stronger sidelobes and reduced contrast. Coherent plane-wave compounding (CPWC) mitigates this loss by steering multiple plane waves over distinct angles and coherently summing the corresponding reconstructions, which narrows the mainlobe and suppresses clutter/grating lobes (Figure 2.6)[69, 70]. However, CPWC introduces a fundamental trade-off: more angles improve image quality but reduce temporal resolution because total acquisition time increases with the number of transmits [71, 72]. Differently from STA, all the probe elements are involved in every transmission, partially compensating the lack of transmit focalization in terms of maximum penetration.

2.3.3 Diverging Wave Imaging (DWI)

Diverging-wave imaging (DWI) emits wide spherical waves synthesized via element delays equivalent to originating from a virtual point behind the array (Figure 2.7); reception/compounding then reconstructs the image at ultrafast rates [73, 74]. The formulation is closely linked to STA imaging; however, like in PWI all or great part of elements are involved




Fig. 2.7 Schematic of a diverging (spherical) transmit formed by a virtual point source (black dot) positioned behind a phased-array transducer. Element-wise transmit delays are chosen to synthesize a wavefront that spans the shaded region of interest, characterized by angular aperture β and steering angle θ , given the probe aperture [5].

in every transmission. Moving from single-element (STA) to multi-element virtual-source transmissions (DWI) increases emitted energy while retaining ultrafast coverage [66, 75].

For wide/deep fields—especially with convex arrays—DWI can extend the usable aperture and the field of view (FOV) relative to PWI, improving clutter suppression in challenging anatomies. For example, microvascular imaging of renal cortex reported a higher signal-to-clutter ratio (SCR) with DWI than with PWI under identical conditions [73]. As with CPWC, coherent compounding over virtual-source locations sharpens the PSF; encoded/distributed schemes (e.g., HDMSA or distributed compounding across push–detect events) preserve kHz-rate tracking for elastography while recovering image quality [75].

2.3.4 Coherent Compounding

Compounding refers to reconstructing a final image from multiple *sub-images*. In medical ultrasound, this concept dates back to the 1980s ([76]; [77]), where sub-images were first formed using conventional focused imaging. Later, images reconstructed from plane-wave or diverging-wave methods were used as sub-images, improving image quality compared with a single transmission while also increasing frame rate relative to conventional ultrasound [78]. In that case, the compounding can be performed on *envelope-detected* images.




Fig. 2.8 Steps of image formation by compounding using plane waves([2]).

The term *coherent compounding*, introduced in [79], designates summation at the *RF* level, i.e., before envelope detection. Let $\text{IRF}_m(x, z)$ denote the beamformed RF image obtained from the m -th transmission (plane wave, diverging wave, STA shot, etc.), and let $w_m(x, z)$ be the (possibly pixel-dependent) compound weight. The coherently compounded RF image is:

$$\text{IRF}_{\text{cmp}}(x, z) = \sum_{m=1}^{N_{\text{TX}}} w_m(x, z) \text{IRF}_m(x, z), \quad (2.19)$$

the principle of plane-wave compounding is illustrated in Fig. 2.8(adapted from [2]). This approach achieves image quality comparable to conventional focused imaging while using roughly ten times fewer transmit events. Consequently, the frame rate can be increased by about a factor of 10 over conventional imaging. It is also flexible, allowing a tunable trade-off between image quality and frame rate by choosing the number of compounded transmissions.

Limitations: Frame-Rate Trade-offs, Complexity, and Artifacts CPWC/DW compounding improves image quality but reduces effective frame rate as the number of transmits grows [71]. Advanced adaptive beamformers (e.g., full MV) can be computationally prohibitive due to per-pixel covariance processing (cubic in aperture size), motivating low-complexity iterative MV and nonlinear options nearer to quadratic cost for real-time feasibility [80]. Fi-

nally, overly aggressive coherence/nonlinear weighting can degrade speckle statistics (“black” artifacts or SNR drop); modern designs therefore balance contrast gains with texture fidelity [81, 82].

2.4 Discussion and Problem Statement

This chapter has reviewed two alternative paths for achieving a better trade-off between image quality and high frame rate (i) *focused* strategies that seek to keep two-way focusing while reducing sequential shots and/or improve image uniformity (e.g., RTBF, SAMF, MLT, MLA), and (ii) *ultrafast* strategies—particularly Plane Wave Imaging (PWI) and Diverging Wave Imaging (DWI)—that replace many focused transmissions with a few wide insonifications and recover image quality through coherent compounding [20, 43, 69, 73]. The review also covered beamforming families beyond DAS (minimum-variance, coherence-based, and non-linear DMAS/F-DMAS) and objective image-quality metrics (FWHM/PSL, CR/CNR, gCNR) [23, 28, 31, 42].

PWI/DWI deliver kilohertz frame rates because frame time scales with the number of transmit events, but a single unfocused transmit yields broader PSF and lower contrast; coherent compounding mitigates this at the cost of additional emissions and thus lower effective frame rate [69]. Focused high-frame-rate (HFR) methods (RTBF, SAMF, MLT, MLA) retain transmit focusing but yields only moderate improvement in frame-rate.[43, 51, 52]. DWI is particularly attractive for wide/deep fields and convex probes, yet its final image quality depends strongly on the *virtual-source* geometry and the *compound weights* used during coherent summation [73, 74].

Gaps and open issues distilled from the state of the art.

1. **High frame-rate vs. image-quality trade-off in CPWC/DWC.** Image quality improves with the number of compounded emissions, but each added wave reduces the frame rate; principled ways to *minimize* the number of emissions for a *target* quality level are still underexplored [69, 71].
2. **Virtual-source (VS) design for DWI is suboptimal in practice.** Regular or heuristic VS layouts can introduce structured clutter and non-uniform PSFs across the FOV, especially with convex arrays. A systematic, *constraint-aware* optimization of VS number and placement remains an open problem [73, 74].
3. **Computational efficiency and deployability.** As discussed in Section 2.1.3, adaptive and non-linear beamformers markedly improve resolution and contrast but can impair

real-time performance; sustaining ultrafast rates on clinical systems therefore requires carefully engineered implementations with per-sample complexity with GPU/FPGA acceleration [23, 31, 58].

Problem statement. *How can we design and validate a DWI transmit–receive strategy—based on optimized virtual-source placement and principled compound weights—that (i) delivers spatially **uniform** high image quality across the field (low FWHM/PSL and low variance), (ii) preserves **ultrafast** acquisition with a **minimal** number of emissions, and (iii) is realizable on a reference ultrasound equipment like ULA-OP[83]/Verasonics Vantage System [84] with quantified computation and memory footprints?*

Thesis objectives and scope. Guided by these questions, this thesis will improve the image quality in DWI by :

1. Propose a computationally efficient method to optimize the weights applied during coherent compounding for DWI.
2. Develop a *constraint-aware*, multi-objective optimization of virtual-source positions (and, when relevant, weights), targeting low mean/variance of FWHM and PSL over representative scatterer fields while keeping the number of emissions small.
3. Build an *end-to-end* beamforming pipeline (RF- and baseband-paths) for ULA-OP/Verasonics data that operates without proprietary toolboxes and reports PSF/contrast metrics (FWHM/CR) alongside compute, memory, and throughput.

2.5 Global Optimization Algorithms

One of the objectives of this thesis is to identify one or more optimal distributions of virtual sources, according to imaging performance criteria. An optimization algorithm can be employed to place the virtual sources to improve these imaging performance metrics. If such an algorithm is to be used, it must be carefully selected and configured with respect to several considerations; these considerations are presented in the following paragraphs.

2.5.1 Deterministic and Stochastic Optimization

The objective (or *cost*) function is the central element guiding the choice of an optimization algorithm. When the objective is *convex*, it has a single local minimum, which is also




Fig. 2.9 Example of a convex function (a) and a nonconvex function (b).

the global minimum. As a simple example, consider

$$f(x) = (x + 0.4)^2 + 0.2,$$

which is strictly convex and attains its unique global minimum at $x^* = -0.4$ (Figure 2.9a). In such cases, deterministic algorithms (e.g., gradient descent, Newton-type methods) reliably converge to the optimal solution. In many signal-processing and imaging problems, however, the objective is *nonconvex* and may exhibit multiple local minima. For instance, the tilted double-well function

$$g(x) = (x^2 - 1)^2 + 0.25x$$

has multiple stationary points and at least two distinct local minima (Figure 2.9b). Here, deterministic methods are strongly initialization-dependent: if the starting point lies in an unfavorable “basin,” the method may converge to a local minimum rather than the global one. To improve exploration of the search space and increase the likelihood of reaching the global minimum, one often adopts *stochastic* optimization strategies that inject randomness to escape local basins. In the literature on sparse ultrasonic array design and related antenna optimization, two prominent families are *simulated annealing* [85] and *genetic algorithms* [86].

2.5.1.1 Simulated annealing

Simulated annealing takes its name from a metallurgical process used to strengthen an alloy: annealing. The goal is to eliminate imperfections in the material (heterogeneities, micro-cracks, delaminations, etc.) by reheating it—i.e., melting it—and then cooling it.

The cooling schedule is chosen so that the alloy cools sufficiently slowly, thereby limiting the formation of defects. This idea was adapted to computing by S. Kirkpatrick [85]. In simulated annealing, a candidate solution $\tilde{\chi}$ is generated from the last accepted solution χ (modification). The two are then compared (evaluation), and the candidate is accepted under certain conditions (acceptance).

Modification: Communication mechanism The generation of a candidate solution $\tilde{\chi}$ from χ is carried out through a *communication mechanism* Θ that explores the solution space Ω . The kernel $\Theta(\chi, \tilde{\chi})$ specifies the probability of proposing a move from χ to $\tilde{\chi}$ via a *neighborhood system*. This neighborhood typically alters only a localized part of the configuration based on its surroundings—for example, flipping a single element within an array of transducers [87]. For simulated annealing to converge properly, certain conditions must be met [88]. In brief: (i) the proposal should satisfy a symmetry (or detailed-balance-compatible) property so that the probability of proposing a move from χ to $\tilde{\chi}$ equals that of proposing the reverse move, $\Theta(\chi, \tilde{\chi}) = \Theta(\tilde{\chi}, \chi)$; and (ii) the proposal must be *irreducible* (ergodic), meaning any solution in Ω can be reached from any other in a finite number of neighborhood moves.

Conditional acceptance When the candidate solution has a lower cost than the current one (i.e., $\Delta E \leq 0$), it is accepted unconditionally and replaces the incumbent. If the candidate is worse ($\Delta E > 0$), it can still be accepted with probability P that depends on the current “temperature”:

$$P = \exp\left(-\frac{\Delta E}{T}\right) \quad \text{equivalently} \quad P = \exp(-\beta \Delta E), \quad (2.20)$$

where

$$\Delta E = E_{n+1} - E_n, \quad (2.21)$$

E_n and E_{n+1} denote the energies (costs) of the incumbent and candidate solutions, T is the temperature, and $\beta = 1/T$ is the inverse temperature. The larger the cost increase ΔE , the smaller the acceptance probability; likewise, as temperature decreases, uphill moves become less likely.

Cooling schedules. The temperature is decreased during the optimization according to a prescribed *cooling schedule*. In the simulated annealing literature, several schedules have been proposed and are known to critically affect performance. From a theoretical standpoint, a logarithmic schedule $T_n \propto 1/\ln(n+1)$ guarantees convergence to a global minimum under

suitable conditions [89], but is typically far too slow in practice [90]. Geometric/exponential schedules are widely used to accelerate convergence [91].

Remark (stochastic continuation). The proposal (communication) mechanism can be adapted over iterations (e.g., changing the neighborhood scale) to further accelerate convergence; such annealers are often termed *stochastic continuation* schemes.

2.5.1.2 Genetic Algorithms

Genetic algorithms (GAs) have been employed to tackle optimization problems in both ultrasound and acoustic imaging [92–95]. Inspired by Charles Darwin’s theory of evolution, they emerged in the optimization literature in the 1990s [86, 96]. The core idea is to generate a population of candidate solutions, evaluate their quality, and retain the best; new candidates are then produced by recombining (*crossover*) and perturbing (*mutation*) the selected ones. These steps are iterated until a stopping or stability criterion is met. The three main stages—evaluation, selection, and crossover & mutation—are outlined below.

Evaluation: fitness function Individuals are evaluated via the objective (cost) function, called the *fitness* in genetic algorithms (GA). This procedure scores each individual so they can be compared against one another.

Selection A selection step is required to retain the best individuals as prospective parents. Several selection schemes exist in the literature, including rank selection, tournament selection, and uniform (roulette-wheel) selection, among others.

Crossover and mutation Once parents have been selected, offspring (i.e., new candidate solutions) are created by recombining parental “genes”—typically by taking part of one solution and part of another. This procedure is called *crossover*. Offspring can also be altered by *mutation*, which modifies a small portion of the genes (for example, changing the position of a single element within an array). Mutation helps the algorithm escape local minima of the objective function.

The stopping condition typically refers to population stability (i.e., convergence). If all individuals in the population have converged to the same type, the algorithm terminates. In practice, this loop is often replaced by a fixed iteration budget.

2.5.1.3 Particle Swarm Optimization

Particle Swarm Optimization (PSO) is an algorithm inspired by collective behaviors such as fish schools, bird flocks, or insect swarms. The key idea is to share information among

particles (analogous to individuals in GA) to conduct a group search. Each particle χ (i.e., a solution) moves in the decision space from position \vec{r}_χ to a new position $\vec{r}_{\tilde{\chi}}$ with velocity $\vec{v}_{\tilde{\chi}}$ (i.e., $\vec{v}_{\tilde{\chi}} \equiv \vec{v}_\chi^{k+1}$). The velocity is composed of three terms:

- an inertia term, i.e., the previous velocity \vec{v}_χ ,
- a *cognitive* term toward the particle's personal best position \vec{r}_{χ_p} ,
- a *social* term toward the global best position \vec{r}_{χ_g} found by the swarm.

The updated velocity is computed by weighting these terms. Let k denote the current iteration; then \vec{v}_χ^{k+1} is the next-iteration velocity, while \vec{v}_χ , \vec{r}_χ , \vec{r}_{χ_p} , and \vec{r}_{χ_g} refer to their values at iteration k :

$$\vec{v}_\chi^{k+1} = w_i \vec{v}_\chi + w_p (\vec{r}_{\chi_p} - \vec{r}_\chi) + w_g (\vec{r}_{\chi_g} - \vec{r}_\chi), \quad (2.22)$$

where the coefficients w_i (inertia), w_p (cognitive), and w_g (social) control the exploration–exploitation (divergence–convergence) trade-off. The optimization steps are summarized.

Here, `NI τ` is the fixed number of iterations (analogous to the number of generations in GA). Convergence speed depends on the chosen weights. If the algorithm converges quickly, particles cluster around what is deemed the optimal location visited by one particle. Conversely, if the influence of personal and global best terms is reduced, convergence slows and exploration increases.

This section has outlined four stochastic optimization methods: simulated annealing, stochastic continuation, genetic algorithms, and particle swarm optimization. The latter two, which operate on *populations* of candidates, belong to the broader family of evolutionary algorithms [97].

2.5.2 Multi-Objective Optimization

The algorithms cited above can be applied when the goal is to minimize a single objective. When multiple objectives must be minimized simultaneously, these algorithms need to be adapted. Two common strategies for handling multiple objectives are outlined below.

Objective aggregation Each objective is a scalar quantity. A first approach consists of combining several objectives into a single scalar. For example, suppose Objective 1 and Objective 2 are to be minimized. They can be aggregated as

$$\text{objective} = \text{objective}_1 + w \times \text{objective}_2, \quad (2.23)$$

where w is a coefficient that balances the relative importance of Objectives 1 and 2. This weighting must be chosen carefully so that one objective does not overwhelm the other. The a priori risk associated with such weighting can be avoided by, secondly, adopting an optimization strategy based on non-dominated solutions.

Non-dominated solutions (NDS) in the objective space The *objective space* (also referred to as the solution space Ω in some texts) is the space covering all attainable objective vectors; it is defined on $\mathbb{R}^{N_{\text{obj}}}$ with $N_{\text{obj}} \in \mathbb{N}^*$ the number of objectives. Any solution χ can be represented as a point in this space. Solutions can also be represented in the *variable space*, defined on $\mathbb{R}^{N_{\text{var}}}$ with $N_{\text{var}} \in \mathbb{N}^*$ the number of scalar decision variables.

Whether a solution is “better” depends on whether the objectives are minimized or maximized. In the case of minimization, objective₁(χ_1) is considered better than objective₁(χ_2) if objective₁(χ_1) < objective₁(χ_2). Conversely, for maximization, objective₁(χ_1) is considered better than objective₁(χ_2) if objective₁(χ_1) > objective₁(χ_2).

Definition of Dominance[98] A solution χ_1 is said to *dominate* another solution χ_2 if and only if the following two conditions hold:

1. χ_1 is no worse than χ_2 for *all* objectives (i.e., when compared directly via their objective values, or equivalently, by their locations in objective space); and
2. χ_1 is *strictly* better than χ_2 in at least one objective.

Figure 2.10 illustrates this notion of dominance for a bi-objective minimization. The set of non-dominated solutions (highlighted in orange) forms a *front* of non-dominated solutions (often called the Pareto front; see [99]). Any solution outside this front and lying within one of the blue regions is dominated by at least one member of the front. A solution located in region A (the darkest blue area) is dominated by *all* non-dominated solutions. Conversely, if a solution were found in region B (green box), it would become the *only* non-dominated solution, and all current non-dominated solutions would then be dominated. This notion of dominance enables handling multiple objectives simultaneously without requiring any pre-optimization balancing (weighting) of the objectives.

2.5.3 Multi-objective optimization algorithms based on non-dominated solutions

In the multi-objective optimization literature, at least two algorithms based on the selection of non-dominated solutions are widely cited: *MOPSO* (Multi-Objective Particle Swarm Optimization, [100]) and *NNSGA-II* (Non-Dominated Sorting Genetic Algorithm II, [99]).

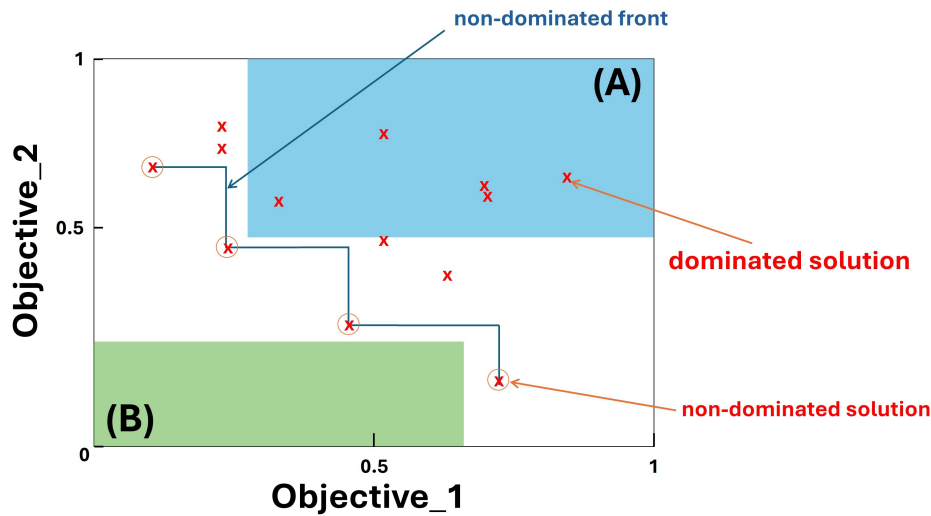


Fig. 2.10 Dominance of solutions in the objective space (bi-objective minimization). The blue step curve indicates the non-dominated (Pareto) front. Points marked with circles are non-dominated; crosses are dominated. Region A illustrates a heavily dominated area (worse in both objectives), while region B marks the ideal improvement zone.

These two methods are, respectively, adaptations of PSO and GA described previously. PSO and GA are population-based algorithms: at each iteration, an entire new population is generated and evaluated, in contrast to simulated annealing, where only a single candidate is handled per step. This population property makes it straightforward to extend single-objective variants to multi-objective formulations grounded in non-dominated solutions, since the non-dominated sorting or selection can be applied directly to each newly formed population without the need to explicitly maintain a separate archive of prior solutions.

2.5.4 Discussion

Different types of optimization coexist, and a multitude of algorithms have been developed over the past decades. A stochastic optimization strategy is essential for searching for optimal virtual-source distributions: the objective function(s) that relate virtual-source positions to image quality metrics are complex. These functions are rarely convex or linear; this is why stochastic algorithms are used in the literature on probe (or other antenna) design by optimization. The goals of the dissertation, defined in Section 2.4, require considering two quality metrics to optimize virtual-source distributions. A multi-objective optimization strategy based on the search for non-dominated solutions is the chosen option: it provides a safe way to balance the different objectives and, at the end of the optimization routine, offers different trade-offs among the objectives. In this study, we adopt a genetic-algorithm

(GA)-based multi-objective optimizer based on preliminary tests with a similar optimization problem [101] that yielded more stable convergence and superior Pareto fronts than simulated annealing or particle swarm approaches.

Chapter 3

Compound Mask for Diverging Wave Imaging: Simulation Study

This chapter introduces the derivation of the compound mask weights for diverging wave imaging from synthetic aperture transmit apodization, detailing the methodology, theoretical background, mathematical formulation, and computational steps involved in generating the compound mask. It also presents simulation-based validation to demonstrate the effectiveness of the proposed approach for both linear and convex probe geometries.

3.1 Introduction

In conventional ultrasound with focused transmission, apodization is applied by varying the amplitude of transmitted signals across the transducer elements, shaping the transmitted beam profile to minimize acoustic energy in off-axis directions. First studies investigated the relationship between predefined apodization windows, like Hamming, Hann, Blackman, Tukey, etc., with the resulting beam pattern shape, extending the well-known findings of the narrowband case to the wideband one employed in most of the medical ultrasound modalities[102]. Similar approaches have been applied in the reception phase, dynamically varying the apodization with the time of arrival of the echoes[103][104]. Besides predefined windows, apodization weights have been synthesized through cost function minimization, aimed at optimizing the trade-off between side lobe level and main lobe width. Among them, approaches based on least squares[105] or stochastic methods [106]. Another approach derives sub-optimal apodization for each focal depth through nonlinear transformation of a baseline window related to a reference depth[107]. Recent advancements have introduced adaptive apodization techniques, such as minimum variance beamforming, which dynami-

cally adjust weights based on received data to further enhance image quality[108]. However, these methods often require iterative optimization, increasing computational complexity and limiting real-time applicability. To cope with this limitations other adaptive approaches select or combine a limited set of predefined apodization windows based on the real-time analysis of received data [109][110], thus reducing the number of variables involved and the computation time.

Since focused transmission achieves optimal focusing just at one fixed depth, unfocused transmit methods have been explored to overcome this limitation. A reference unfocused method is the so called Synthetic Transmit Aperture Imaging (STAI) [111], in which every probe element transmits in sequence and the received signals related to each transmission are aligned and compounded together to beamform every image point. Transmit apodization in STAI is applied in the compounding step of the reception phase: in this way it can be dynamically tailored to each focal depth. STAI transmit apodization can be set borrowing from methods developed for focused transmission, just setting independently the apodization window for each focal depth. Though STAI achieves a theoretically optimal beam pattern it suffers from poor SNR due to limited amount of energy transmit by each element. Moreover the frame rate remains comparable to that one of focused transmissions.

To overcome these limitations other unfocused transmission methods, based on the generation of coherent plane waves (PWs) [9]and divergent waves (DWs) [10][11], have been developed, making it possible to achieve ultrafast imaging with high frame rates and wide fields of view. As for STAI, receive signals related to different transmissions are compounded together. However, differently from STAI, the computation of optimal compound weights has not a straightforward relation with transmit anodization, and it is subject of ongoing research. For PWs some studies have drawn parallels between coherent plane-wave compounding (CPWC) and STAI, suggesting that STAI apodization techniques can be adapted for angular weights in CPWC [17]. Other approaches aim to dynamically estimate the compound weights from the received data, for example adopting Capon's minimum variance beamforming [112]. Recent advancements have incorporated Independent Component Analysis (ICA) into CPWC for angular apodization. ICA decomposes the received signals into statistically independent components, which helps estimate the angular weights more effectively [113]. This method assumes that each plane-wave transmission provides a non-independent observation of the target field, thus enabling the reconstruction of high-quality images through optimal weight estimation.

These adaptive approaches have not yet been investigated for DWI. Various researchers have proposed different compounding strategies for DWI [114]. Two predominant transmit sequences for coherent compounding have been described thus far. One method involves

generating a divergent wave using the full aperture, whereby the point spread function (PSF) is rotated between transmits by moving the virtual focus point along an arc at a constant distance from the probe's center [10] [11]. The other method uses only part of the aperture to generate the divergent wave, tilting the PSF by sliding the active aperture across the array [16]. It remains uncertain which of these methods yields superior results. According to [114], spatial compounding by shifting the active aperture of a diverging wave source can enhance image quality more effectively than rotating a diverging wave generated with the full aperture. However, this study does not address how to apply transmit apodization in DWI or what the optimal weights to apply during coherent compounding are.

While prior research has advanced the understanding of apodization in synthetic aperture imaging and coherent compounding in plane wave imaging, the optimization of compound weights for divergent wave imaging remains an open challenge. Existing methods either rely on empirical weight assignment or computationally intensive adaptive algorithms, limiting their practicality in real-time applications. This work addresses this gap by deriving a relationship between an arbitrary apodization law in STAI and the compound mask weights in DWI, thus leveraging established optimization criteria in STAI transmit apodization. In particular, under suitable hypotheses, a closed-form solution is computed, allowing straightforward mapping of STAI apodization coefficients into DWI compound mask coefficients. The method works for both linear and convex probe geometries and arbitrary arrangements of virtual source positions, thus allowing the approach to be widely applicable.

3.2 Methodology

3.2.1 Synthetic Aperture Imaging (SA) using Linear and Convex Probes

Synthetic Aperture Imaging (SA) is used in ultrasound imaging to combine signals from multiple transmit-receive events to form a high-resolution image. SA allows for dynamic focusing in transmission and reception and improves image quality and resolution. The signal reconstructed at any point (x, z) in the imaging field can be mathematically represented by the following equation:

$$S_{SA}(x, z) = \sum_{r=1}^M u_r(x, z) \sum_{t=1}^N v_t(x, z) h_{tr} \left(\frac{D_t}{c_0} + \frac{D_r}{c_0} \right) \quad (3.1)$$

where $u_r(x, z)$ and $v_t(x, z)$ are the receive and transmit apodization windows function element index r and t and image point coordinates x, z . The transmit apodization $v_t(x, z)$ is the weight

applied to the signal transmitted by element t ¹. The receive apodization $u_r(x, z)$ is the weight applied to the signal received by the element r . The indices M and N are the number of active transmit and receive elements, respectively. h_{tr} is the two-way impulse response from transmit element t to receive element r ; D_t and D_r are the distances between the point (x, z) and elements t and r , whose coordinates are x_t, z_t and x_r, z_r respectively, and c_0 is the medium's sound speed. The representations for both synthetic aperture with linear array (SALA) and convex array (SACA) are shown in Figure 3.1.

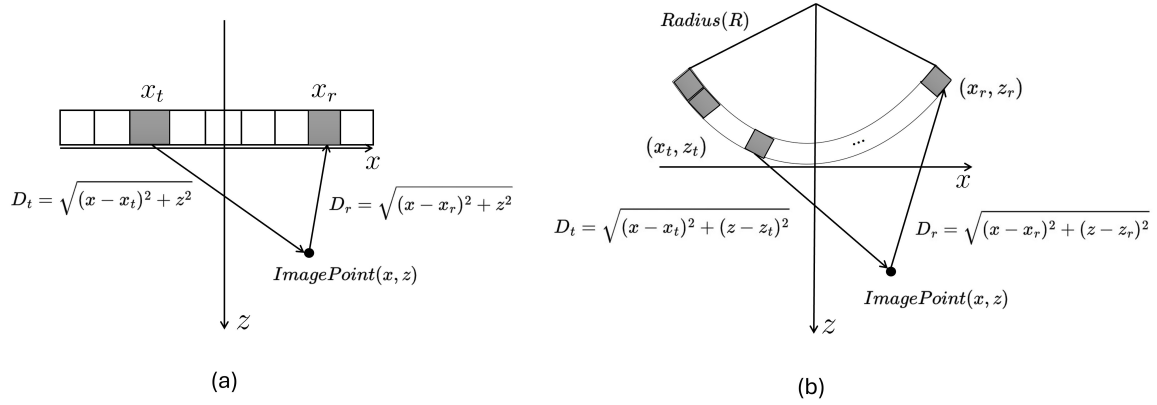


Fig. 3.1 Comparison of SA Linear and Convex Probe Configurations with distances between image point (x, z) and probe elements coordinates in transmission (x_t, z_t) and reception (x_r, z_r)

3.2.2 Diverging Wave Imaging using Linear and Convex probe

A standard approach to generate diverging waves is to consider a virtual source point i , of coordinates (x_i, z_i) , and compute the relative arrival times of a spherical wave propagating from the virtual source to each active probe element location. These relative times of arrival are then applied as a transmission delay curve to the probe elements to generate a divergent field approximating the spherical one in the image plane. The effect of angular compounding, achieved in plane wave transmission by steering the plane wave angle, is achieved here by moving the virtual source in different locations. The optimal placement of virtual sources will be detailed in chapter 5. Given a set of virtual source positions, the signal reconstructed at every point (x, z) can be expressed as follows:

$$S_{DW}(x, z) = \sum_{r=1}^M u_r(x, z) \sum_{i=1}^V w_i(x, z) \hat{h}_{ir} \left(\frac{D_i}{c_0} + \frac{D_r}{c_0} \right) \quad (3.2)$$

¹Notice that in SA imaging, transmit apodization is actually applied to the received signal. This is mandatory to make it a function of the imaging point (x, z)

where $w_i(x, z)$ denotes the compound mask function of virtual source index i and image point coordinates. The compound mask is defined as the set of weights applied to the beamformed data related to a specific virtual source in transmission, at each image point (x, z) [115]. \hat{h}_{ir} is the two-way impulse response of the diverging wave from the virtual source i to the image point (x, y) and back to the receive element r , and D_i is the distance between the virtual source of coordinates (x_i, z_i) and the image point. The impulse response of the diverging wave field can be described as a linear combination of impulse responses related to each active probe element in transmission with proper relative delays:

$$\hat{h}(t)_{ir} = \sum_{t=1}^N h_{tr}(t - D_{it}/c_0) \quad (3.3)$$

where D_{it} is the distance between the virtual source i and the transmit element t . Substituting the above expression into the equation for the diverging wave imaging, one obtains:

$$S_{DW}(x, z) = \sum_{r=1}^M u_r(x, z) \sum_{i=1}^V \sum_{t=1}^N w_i(x, z) h_{tr} \left(\frac{D_t}{c_0} + \frac{D_r}{c_0} + \tau_{it} \right) \quad (3.4)$$

where τ_{it} is given by:

$$\tau_{it} = (D_i - D_t - D_{it})/c_0 \quad (3.5)$$

In Figure 3.2 the respective distances D_i , D_r , D_{it} are displayed with the linear and convex probe geometries. From the above expression, it can be noticed that both $S_{SA}(x, z)$ and $S_{DW}(x, z)$ can be thought as a linear combination of two-way probe elements impulse responses $h_{tr}(t)$. This allows, under opportune hypotheses, to derive a simple relationship between the compound mask of S_{DW} and apodization weights of S_{DW} , as described in the next Section.

3.2.3 Relation between transmit apodization in STAI and Compound Mask in DWI

In the work of Rodriguez-Molares et al [17], the authors derived a relationship between apodization in STAI and CPWC, showing that under certain conditions, the angular compounding weights can be derived in closed form from the transmit apodization window of STAI. Here, we adapt their approach to diverging waves generated by virtual sources and arbitrary probe geometries. First of all, under the same assumptions of [17] a linear relationship between the compound mask for diverging waves and transmit apodization can be derived assuming that the last expression of $S_{DW}(x, z)$ and the expression for $S_{SA}(x, z)$ are




Fig. 3.2 Comparison of DW Linear and Convex Probe Configurations with distances between virtual source (x_i, z_i) and image point (x, z) and probe elements coordinates in transmission (x_t, z_t) ; moreover also the distance between image point and probe elements coordinates in reception (x_r, z_r) is displayed.

equivalent. The following expression yields:

$$\sum_{i=1}^V w_i(x, z) H(i, t) = v_t(x, z) \quad (3.6)$$

where

$$H(i, t) = \frac{E(\tau_{it})}{E(0)} \quad (3.7)$$

where $E(t)$ is the convolution of the transmit pulse with the two-way electroacoustic transfer function of the probe elements (see Eq. 14 in [17]). Collecting the above equations for each apodization weight $v_t(x, z)$, a linear system is derived from which the compound mask can be computed, as described in [17]. If we display the matrix $H(i, t)$ in Figure 3.3, we can see that only a few elements, corresponding to values of τ_{it} close to zero, are different from zero. This means that under the additional hypothesis of short-time transmitted pulse, already introduced in [17], it is possible to approximate the matrix $H(i, t)$ in Eq. 6 as diagonal with non-zero values equal to 1. As a consequence, the only non-zero contribution of the sum of Eq. 6 is the one for which $\tau_{it} = 0$. In other words Eq. 3.6 becomes $w_i(x, z) = v_t(x, z)$ for the couple of (i, t) fulfilling the condition $\tau_{it} = 0$. Looking at the definition of τ_{it} , we see that it represents the difference between the distance from the virtual source to the image point and the sum of the distances between the virtual source and the probe element and between the probe element and the image point, all converted in time. By the triangular inequality, the difference can be zero only if the probe element is aligned along the line connecting the virtual source and the image point. Therefore, we computed the intersection point between

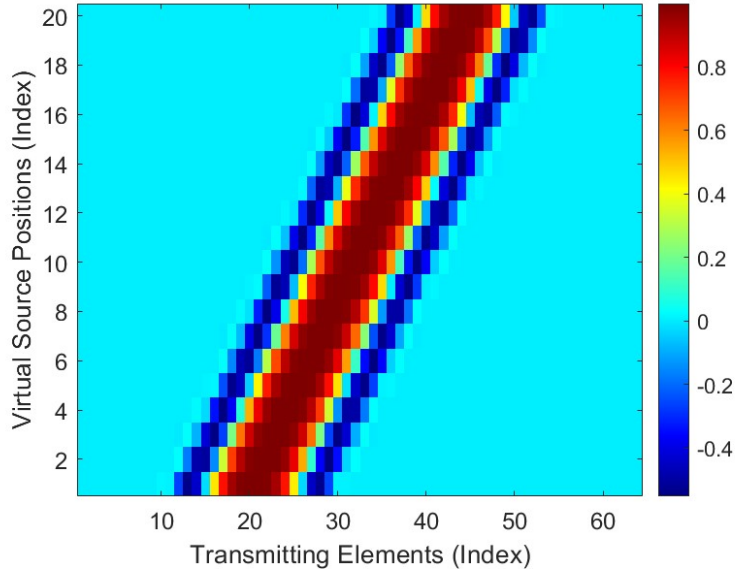


Fig. 3.3 Examples of values of $H(i, t)$ matrix in Eq. 7. Row and column indexes are related to virtual sources and probe elements indexes, respectively

the above line and the probe surface, for each virtual source position, thus obtaining the associated probe element and the corresponding apodization weight. For a convex probe of radius r_{probe} , the equation yielding the intersection points (x_e, z_e) is given by:

$$z_e = \frac{b}{m^2} + \frac{\sqrt{\left(\frac{b}{m^2}\right)^2 - \left(1 + \frac{1}{m^2}\right) \left(\left(\frac{b}{m}\right)^2 - r_{probe}^2\right)}}{1 + \frac{1}{m^2}} \quad (3.8)$$

$$x_e = \frac{z_e - b}{m} \quad (3.9)$$

where:

$$m = \frac{z - z_i}{x - x_i} \quad \text{and} \quad b = z - \left(\frac{z - z_i}{x - x_i}\right) \cdot x \quad (3.10)$$

Based on these intersection coordinates, the angle of intersection θ for a convex probe is calculated as follows:

$$\theta_e = \arctan\left(\frac{x_e}{z_e}\right) \quad (3.11)$$

For a linear probe, the x coordinate of the intersection point can be expressed as:

$$x_e = \left(-\frac{b}{m} \right) \quad (3.12)$$

The resulting intersection points are then used to map the transmit aperture onto the compound mask. Since the intersection point generally does not correspond to the center of a probe element, an interpolation across adjacent probe elements on apodization weights is carried out to compute the corresponding compound weight. In particular, for a linear geometry, defining $[x_t]$ the set of x coordinates of the transmit elements and $[v_t(x, z)]$ the corresponding set of transmit apodizations in SA for image point (x, y) , the compound weight is computed as:

$$w_i(x, z) = \text{interp}([x_t], x_e, [v_t(x, z)]) \quad (3.13)$$

where $\text{interp}(a, b, c)$ is an interpolating function computing an off-grid value at coordinate b starting from a vector of values c a vector of on-grid coordinates a . A similar approach holds for convex geometries, where compound weights are computed as:

$$w_i(x, z) = \text{interp}([\theta_t], \theta_e, [v_t(x, z)]) \quad (3.14)$$

with $[\theta_t]$ being the set of angles related to transmit elements computed as in Eq. (3.11).

3.2.4 Compound Mask Computation

As mentioned before, the compound mask computation starts from an arbitrary set of transmit apodizations $v_t(x, z)$ defined for SA imaging. A common choice, here adopted, for $v_t(x, z)$ to ensure image uniformity, is to start from a baseline apodization function \hat{v}_t defined for a reference image point $(0, z_{ref})$ and compute the overall set $v_t(x, z)$ by interpolation of scaled and translated versions of \hat{v}_t . For linear geometries, this can be done by:

$$v_t(x, z) = \text{interp}([x_t], \frac{(x_t - x)z_{ref}}{z}, [\hat{v}_t]) \quad (3.15)$$

For convex geometries the corresponding equation is:

$$v_t(x, z) = \text{interp}([\theta_t], \frac{(\theta_t - \theta)\rho_{ref}}{\rho}, [\hat{v}_t]) \quad (3.16)$$

where θ and ρ are the polar coordinates of image point (x, z) and ρ_{ref} is the radius polar coordinate corresponding to z_{ref} when $x = 0$. The baseline SA apodization \hat{v}_t can be selected

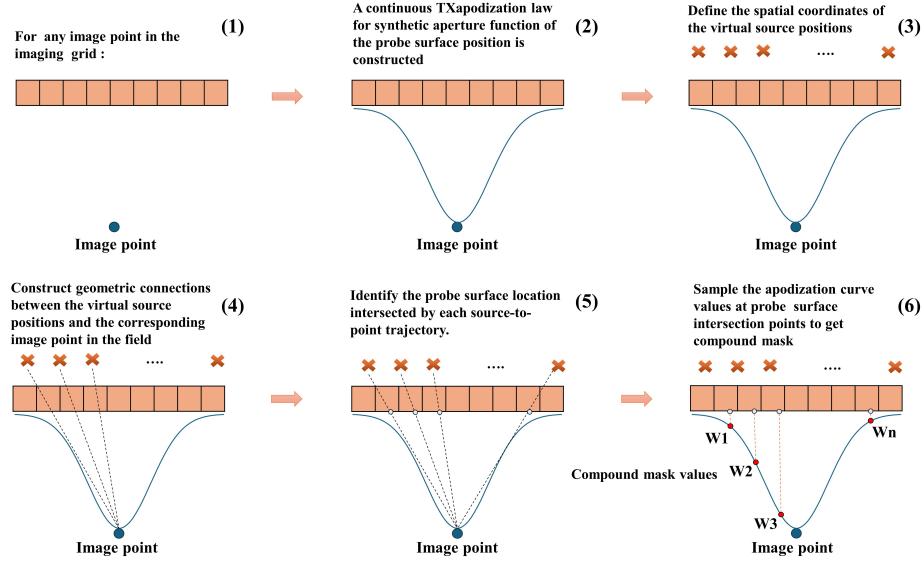


Fig. 3.4 A schematic illustrating the generation of the compound mask for Diverging Wave Imaging (DWI). Starting from each image point, a continuous transmit apodization curve is first defined over the probe surface, following the principles of Synthetic Aperture Imaging (SAI). Then, a set of virtual source (VS) positions is specified behind the probe, which can follow arbitrary spatial distributions. For each VS, a straight-line trajectory is constructed toward the image point, and the corresponding intersection with the probe surface is determined. At these intersection points, the apodization curve is sampled, and the resulting values are assigned as weights to the respective virtual sources. The same methodology applies to convex arrays by replacing the linear probe intersection with angular (θ_E) intersections and is valid for every image point regardless of the virtual source distribution.

according to standard apodization windows like Hamming, Hann, Tukey, truncated Gaussian, each one with well-known trade-off between main lobe width and side lobe height.

3.2.5 Algorithm for compound mask computation

This section provides a structured overview and a pseudo-code for the algorithm used to derive the compound mask for DWI. The algorithm inputs are given by the probe geometry, the position of virtual sources, and the imaging grid. There are two preliminary stages and a third stage that is the actual core of the algorithm:

Baseline Transmit Apodization Definition: The methodology begins by defining a baseline transmit apodization profile across the active probe elements \hat{v}_t . Standard apodization windows like Hamming, Hann, etc. can be employed.

Apodization Mapping to Imaging Points: Since the apodization window must adapt to varying depths and lateral positions, it undergoes a transformation based on the image

Algorithm 1 Compound Mask Computation for Diverging Wave Imaging for linear and convex geometries

Input : Sets of probe elements coordinates $[(x_t, z_t)]$, virtual sources coordinates $[(x_i, z_i)]$, imaging grid coordinates $[(x, z)]$, probe radius r_{probe} (for convex probes)

Output : Set of compound mask weights $[w_i(x, z)]$

```

1 Define a baseline SA apodization  $\hat{v}_t$  and a reference depth  $z_{ref}$ 
2 for each imaging point  $(x, z)$  do
3   for each transmit element  $t$  do
4     // Compute SA transmit apodization
5     if probe is convex then
6        $v_t(x, z) = \text{interp}([\theta_t], \frac{(\theta_t - \theta)\rho_{ref}}{\rho}, [\hat{v}_t])$ , Eq. 16
7     if probe is linear then
8        $v_t(x, z) = \text{interp}([x_t], \frac{(x_t - x)z_{ref}}{z}, [\hat{v}_t])$ , Eq. 15
9   for each virtual source  $i$  do
10    // compute intersection between probe ceramic and line connecting
11    // image point  $(x, z)$  and virtual source  $i$ 
12    if probe is convex then
13       $\theta_e$  Compute intersection angle  $\theta_e$  by Eqs. 8-11
14    if probe is linear then
15       $x_e$  compute intersection coordinate  $x_e$  by Eqs 10,12
16    // Compute compound mask from interpolation of SA transmit
17    // apodization
18    if probe is linear then
19       $w_i(x, z) = \text{interp}([x_t], x_e, [v_t(x, z)])$  (Eq. 13)
20    if probe is convex then
21       $w_i(x, z) = \text{interp}([\theta_t], \theta_e, [v_t(x, z)])$  (Eq. 14)
22 return  $[w_i(x, z)]$ 

```

point coordinates. This is achieved by interpolation of a translated and scaled version of the baseline apodization profile. The results is a set of apodization profiles $v_l(x, z)$, one for each image grid point (Eqs. 15-16).

Compound Mask Computation: In the final stage, for every combination of virtual source i and image point (x, y) , a compound mask value $w_i(x, z)$ is computed from the apodization profile $v_l(x, z)$. The process involves two sub-steps. In the first one, the intersection between probe surface and line connecting image point (x, z) and virtual source point (x_i, z_i) is computed (Eqs. 8-12). In the second one, the value of apodization, associated to the probe surface location previously computed, is assigned to the compound mask value. Since the surface location is in general outside of the probe elements' grid, an interpolation function is adopted (Eqs. 13,14). The algorithm holds for both linear and convex geometries. What changes are the specific expressions for computing SA apodizations, intersections on the probe surface, and final interpolation to get the compound mask value. ²

3.3 Results and Discussion

3.3.1 Geometrical warping

To illustrate the mapping process, Figure 3.5 and Figure 3.6 show the geometric warping of three tapering windows (Hanning, Blackman-Harris, and Tukey with 0.5 roll-off) for an F-number of 0.7. In Figure 3.5 (a), the transmit apodization is shown for a fixed point located at (0,20) mm in SAI, while Figure 3.5 (b) displays the corresponding mapped weights $w_i(x, z)$ for the same point in DWI, the same for Figure 3.6 in 2D representation. Both figures demonstrate how the numerical values are locally preserved for a fixed image point., highlighting the effectiveness of the geometric transformation. As can be seen the transformation consists of a geometric warping that locally preserves the numerical values, the apodization function demonstrates a smooth distribution of weights, peaking at the center and tapering off towards the edges, ensuring optimal energy distribution during transmission.

²The content of this chapter has been published in the IEEE access journal: **Z. Alzein**, M. Crocco, and D. D. Caviglia, "Compound Mask for Divergent Wave Imaging in Medical Ultrasound," *IEEE Access*, vol. 13, pp. 59675–59686, 2025, <http://dx.doi.org/10.1109/ACCESS.2025.3557085>.




Fig. 3.5 Comparison of transmit apodization in Synthetic Transmit Aperture Imaging (STAI) and virtual source-based apodization in Diverging Wave Imaging (DWI). (a) Transmit apodization $v(x_T)$ in STAI for Hanning, Blackman, and Tukey (0.5 roll-off) windows. (b) The corresponding virtual source-based apodization $w(x_i, z_i)$ in DWI, computed using the proposed method, for a fixed point located at (0, 20) mm. The virtual sources are placed linearly behind the linear probe at a distance of 20 mm.

3.3.2 Validation of compound Mask using simulated data

3.3.2.1 Simulation set-up

The effectiveness of the method was validated through simulated data generated from Field II software[116][117] for both linear and convex probes. A 192-element array with a 5 MHz central frequency and 0.25 mm pitch was used for the linear probe, while a 3 MHz central frequency and 0.3 mm pitch were used for the convex probe with a radius of 50mm. A Hanning window was used in transmission and reception, with an F-number of 0.7 and 1.4 with a 2-cycle sinusoidal excitation signal. Pixel beamforming was performed on a grid of $Z=164$ and $X=156$. The image point spread function (PSF) for a point scatterer located at 20 mm and 30 mm depths was evaluated for the four cases, i.e. Synthetic Aperture Linear Array (SALA), Synthetic Aperture Convex Array (SACA), Diverging Wave Linear Array (DWLA), and Diverging Wave Convex Array (DWCA). For DWLA, a linear distribution of virtual sources placed at a fixed distance behind the probe surface has been adopted. This arrangement is depicted in Figure 3.2(a), where the virtual sources are evenly spaced along the x-axis directly behind the probe elements. For DWCA, the virtual sources were positioned along a circle behind the probe surface (Figure 3.2(b)). Research on optimized virtual source distributions for linear probes and convex arrays in 2D imaging is limited. Therefore, a deterministic distribution of virtual sources is employed for both probe types according to the proposal detailed in [118]. For the linear probe, 20 virtual sources are uniformly distributed along the lateral axis, at a distance of 50 mm from the surface of the probe. These sources are positioned at regular intervals across a range of -25 to 25 times pitch. For the convex probe,




Fig. 3.6 Two-dimensional representation of (a) transmit apodization $v_t(x, z)$ in Synthetic Transmit Aperture Imaging (STAI) using Hanning, Blackman-Harris, and Tukey (0.5 roll-off) windows as baseline apodization, and (b) the corresponding virtual source-based apodization $w_i(x, z)$ in DWI Linear Array with the virtual sources are placed linearly behind the linear probe at 20 mm. Both maps are displayed for $x = 0$.

the distribution of 20 virtual sources follows a curvilinear pattern to match the geometry of the probe elements. The angular range for these sources is defined by the maximum aperture angle set to 33° , covering a symmetric span on both sides of the probe's central axis. The curvilinear pattern of virtual sources is placed behind the probe, intersecting the probe center

3.3.2.2 PSF Results

The images of point-like scatterers are displayed in Figure 3.7 using F-number 0.7, and their evaluation in terms of full width at half maximum (FWHM) is reported in Table 3.1. For linear probes, the Full Width at Half-Maximum (FWHM) with DWLA increased by 7.5% at 20 mm and by 9% at 30 mm compared to SALA. In contrast, DWCA showed a




Fig. 3.7 PSF of SALA, DWLA, SACA, and DWCA for a point at 20 mm depth (first row), and 30mm depth (second row) with F-number=0.7

much smaller increase of 1.64% at 20 mm and a substantial increase of 26.56% at 30 mm compared to SACA, with a slight increase in sidelobes. Overall, the proposed approach for the computation of the compound mask seems to reproduce quite accurately the effect of transmit apodization in Synthetic Aperture in terms of spatial resolution, except for the case of DWCA at 30 mm. This result accounts for the fact that the spatial distribution of virtual sources also plays a crucial role in image optimization, and it may limit the accuracy of mapping between STAI apodization and DWCA compounding. These aspects will be investigated in future studies. Concerning the side-lobe level, the worsening is likely due to the non-ideal nature of the divergent wavefield, given by the boundary effect introduced by the finite aperture.

Technique	FWHM at 20 mm	FWHM at 30 mm
SALA	0.40	0.43
DWLA	0.43	0.47
SACA	0.61	0.64
DWCA	0.62	0.81

Table 3.1 Comparison of FWHM values at 20 mm and 30 mm for different techniques with F-number 0.7

In addition to the results reported here, the proposed methodology was also evaluated with different F-number to investigate its robustness and adaptability. The results of this extended

analysis, which further confirm the effectiveness of the compound mask formulation, are presented in our recent publication [119].

3.4 Conclusion

In this chapter, we derived the compound mask weights for DWI directly from the transmit apodization used in SAI. The proposed approach, based on a closed-form geometric mapping, enables a consistent transfer of apodization characteristics to virtual source transmissions for both linear and convex array geometries. Through theoretical formulation, computational implementation, and simulation-based validation, we demonstrated that the method preserves key image quality metrics, including main-lobe width, while extending the benefits of SAI apodization to unfocused wave imaging.

The proposed compound mask framework serves as a foundational tool for the rest of this thesis. Its experimental validation and real-time implementation are presented in the next chapter 4, where it is integrated into a Verasonics beamforming system and evaluated on tissue-mimicking phantoms. Furthermore, the same methodology is leveraged in Chapter 5 to enhance the optimization algorithm for virtual source placement, enabling more efficient image-quality enhancement. The compound mask concept is also integrated into the GPU-based beamformer design presented later in the thesis, where it contributes to achieving high-quality imaging with significantly reduced computational cost in chapter 6. Finally, this methodology is extended to three-dimensional volumetric DWI in chapter 7.

Chapter 4

Experimental Validation of Compound Mask in Real Time

This chapter validates the compound mask approach detailed in 3 experimentally. To do so, real-time acquisition of a multi-purpose phantom data has been performed with a Verasonics scanner using different spatial arrangements of virtual sources (VS) for both probe geometries linear and convex probe. The goal was to assess the improvement in image quality in DWI imaging obtained by applying the compound mask, with respect to a standard DWI imaging in which all contributions from all the transmissions are equally weighted in reception.

4.1 Introduction

To evaluate the compound mask effectiveness in more sophisticated scenarios, such as in-vitro data, and not relying only on the simulated data, this chapter presents the experimental validation of the proposed method in enhancing the image quality without adding computational cost overhead in the beamforming pipeline for Verasonics. The proposed mapping approach is applied on a Verasonics platform with negligible computational burden during online processing, as the weights are seamlessly integrated into the beamforming pipeline. The method was evaluated across multiple virtual source configurations with linear and convex probes while maintaining real-time imaging capability.

4.2 Methodology

This section provides a concise overview of the key concepts and equations of the proposed method to calculate compound mask weights in DWI based on the theoretical framework

established in 3, followed by a detailed description of the experimental validation setup and integration to the Verasonics scanner.

4.2.1 Integration of the compound-mask apodization in the Verasonics pipeline

The compound-mask apodization was fully integrated into the Verasonics real-time beamforming pipeline using only standard data structures and buffers, ensuring compatibility with both linear and curvilinear probes. For each virtual-source (VS) transmission event $m = 1, \dots, M$, the transmit and receive apodization coefficients were assigned through the built-in fields `Tx(m).Apod` and `Receive(m).Apod`. These apodizations define the element-level taper specific to each VS configuration and are automatically interpreted by the Verasonics sequence controller, without any modification of the underlying delay-and-sum (DAS) beamforming kernel. The pixel-wise compound-mask weights $w_m(x, z)$, precomputed once on the imaging grid for all VS indices m , were applied during the coherent compounding stage. Each event-specific complex image $p_m(x, z)$ was first reconstructed by the Verasonics reconstructor and stored in a dedicated `InterBuffer` frame. This was configured by setting `ReconInfo(m).mode = 'replaceIQ'` and mapping each event to a separate frame index. After all M VS events were reconstructed, a lightweight custom `Process` function was executed to perform masked coherent compounding. In this step, each per-event image $p_m(x, z)$ was multiplied by its corresponding weight map $w_m(x, z)$ and accumulated as:

$$p(x, z) = \sum_{m=1}^M w_m(x, z) p_m(x, z),$$

where $p(x, z)$ is the final compounded image. The result was written to the `ImageBuffer` using `ReconInfo(end).mode = 'accumIQ_replaceIntensity'` for subsequent envelope detection and log compression. This integration preserves real-time operation, adding negligible computational overhead (one scalar multiplication per pixel per VS).

4.2.2 Virtual Sources Spatial Arrangement for Linear and Convex Arrays in DWI

The proposed approach was evaluated using different spatial distributions of virtual sources (VSs) for linear and convex arrays. For linear arrays, 20 virtual sources were distributed linearly along the x -axis, with a spacing of $\pm 32 \times \text{pitch}$ (where $\text{pitch} = 298 \mu\text{m}$), at a fixed depth $z_v = 20\text{mm}$ behind the probe. For convex arrays, three well-established transmission

strategies for wide-focus imaging were used, as previously described in [118]. The first strategy, referred to as the linear distribution, follows the same configuration used for the linear array, with virtual sources arranged linearly. The second strategy, known as the curvilinear distribution, aligns the virtual sources with the geometry of the convex probe. In this configuration, the angular placement of the sources is confined within a maximum aperture angle of 33° , distributed symmetrically around the central axis of the probe. In the third strategy, termed the tilted distribution, the virtual sources are placed along a radial line with a constant radius $R_{\text{convex}} = 41.42$ mm. The center of this curve is positioned at the origin, with the same angular placements used in the curvilinear strategy. All distributions are visually presented in Figure 4.1 for both linear and convex probes.

4.2.3 Experimental Setup

To rigorously evaluate the efficacy of the proposed compound mask mapping method, a comprehensive experimental validation was performed using a Verasonics Vantage 256 research ultrasound system (Verasonics, Redmond, WA, USA). The experimental setup included a 128-element linear array probe (L7-4) with a central frequency of $f_0 = 5$ MHz and a 128-element convex array probe (C5-2) with a central frequency of $f_0 = 3.1$ MHz. Both transducers were excited with a 2-cycle sinusoidal signal, and RF-data were collected with a sampling frequency of 12 MHz, allowing the evaluation of the method across linear and curvilinear geometries. A tissue-mimicking phantom (Model 040GSE, CIRS Inc., USA) was used to emulate realistic imaging conditions, allowing for the quantitative assessment of key image quality metrics such as spatial resolution and contrast. Data acquisition was performed in Diverging Wave Imaging (DWI) and DWI with a compound mask integrated into the Verasonics Beamforming pipeline. In DWI, diverging waves were generated from 20 virtual sources positioned behind the probe surface, with the spatial distribution described in 4.2.2, arranged linearly for the linear probe (Figure 4.1a) and in three distinct configurations (linear, tilted, and curvilinear) for the convex probe (Figure 4.1 b, c, d).

B-mode images were acquired in real time for diverging wave imaging (DWI), and DWI with the proposed compound mask using a Hanning window as the chosen apodization law for both linear and convex array probes. Under the experimental configuration used in this study, including the selected image grid and 20 virtual sources (VSs), the compound mask integration introduced no computational overhead during real-time operation. However, the computational complexity of generating compound mask weights, $w_i(x, z)$, scales with both the imaging grid ($N_x \times N_z$) and the number of virtual sources (N_{VS}). For computationally constrained systems, these weights can be precomputed offline and stored in memory, following established apodization methods [115]. This enables seamless integration into the




Fig. 4.1 Spatial arrangements of virtual sources (VSs) used in the experimental validation. (a) Linear distribution for the linear array, with the virtual sources placed at a depth $z_d = 20$ mm above the probe. For the convex array, three configurations are shown: (b) Linear arrangement, where the virtual sources are distributed linearly; (c) Curvilinear arrangement, where the virtual sources follow the curvature of the probe, with angular placement confined within a maximum aperture angle of 33° ; and (d) Tilted arrangement, where the virtual sources are placed along a radial line at a depth equal to R_{convex} . In all configurations, all probe elements are used in both transmission and reception, with $N = 20$ virtual sources used for each arrangement.

Verasonics beamforming pipeline while maintaining $\leq 5\%$ additional latency compared to standard DWI.

4.2.4 Evaluation Metrics

The performance of the proposed method was evaluated by comparing the image quality metrics obtained from DWI with and without the compound mask. The lateral resolution was measured using the Full Width at Half Maximum (FWHM) [3] for the point scatterer embedded in the multi-purpose phantom (Model 040GSE, CIRS Inc., USA), and the contrast was measured using anechoic cysts on envelope images. The contrast ratio (CR) was computed as:

$$CR = 20 \log_{10} \left(\frac{\mu_{in}}{\mu_{out}} \right), \quad (4.1)$$

where μ_{in} and μ_{out} denote the mean pixel intensities inside and outside the cyst region, respectively. The regions extracted to compute μ_{in} and μ_{out} are visualized in Fig. 4.2, with all selected regions having the same size and are located at the same depth.

1

¹The content of this chapter has been partially published in the IEEE IUS conference, and the full experiment is submitted to Ultrasonic Imaging journal (under revision): **Z. Alzein**, H. Liebgott, M. Crocco, and D. D. Caviglia, "Virtual Source-Based Apodization for Divergent Wave Imaging: An Experimental Study" // **Z. Alzein**, H. Liebgott, M. Crocco, and D. D. Caviglia, "Experimental Validation of a Real-Time Compound Mask for Enhanced Diverging Wave Imaging"




Fig. 4.2 Example of regions of interest (ROIs) used for contrast ratio (CR) computation. Green circles indicate the regions inside the cysts used to estimate μ_{in} , and the orange circle indicates the background region used to estimate μ_{out} .

4.3 Results and Discussion

4.3.1 B-mode Images using Linear probe L7-4:

Figure 4.3 shows the B-mode images obtained using the Verasonics machine with a linear array probe (L7-4) for Diverging Wave Imaging (DWI), and DWI after applying the proposed compound mask. The virtual sources (VSs) in DWI were distributed linearly above the probe, as illustrated in Figure 4.1(a). The images in the first row first row of Figure 4.3 display a set of point-like scatterers useful to measure the spatial resolution, with the corresponding Full Width at Half Maximum (FWHM) values reported in the boxplots in Figure 4.4(a). The images in the second row show three cysts located at a depth of 45 mm, useful to measure the contrast ratio (CR) defined in Equation (2), with respect to the surrounding uniform speckle. The B-mode images visually demonstrate that the resolution of DWI after applying the compound mask is higher than DWI without the mask. This observation is quantitatively confirmed by the boxplots in Figure 4.4(a). DWI with the compound mask shows a 30% improvement in resolution compared to DWI alone. Similarly, the contrast measurements reported in Table 4.1 and in the zoomed-in figure 4.4(b) reveal that the application of the compound mask improves the CR for all three cysts, significantly outperforming DWI for the predefined VS distribution.




Fig. 4.3 B-mode images generated using the L7-4 linear array probe and Verasonics scanner for Diverging Wave Imaging (DWI), and DWI with the proposed compound mask. The images correspond to the sequences described in Figure 4.1(a). The first row displays the resolution measurements, while the second row shows the imaging results for three greyscale cysts (anechoic, -6 dB, $+6$ dB) located at a depth of 45 mm.

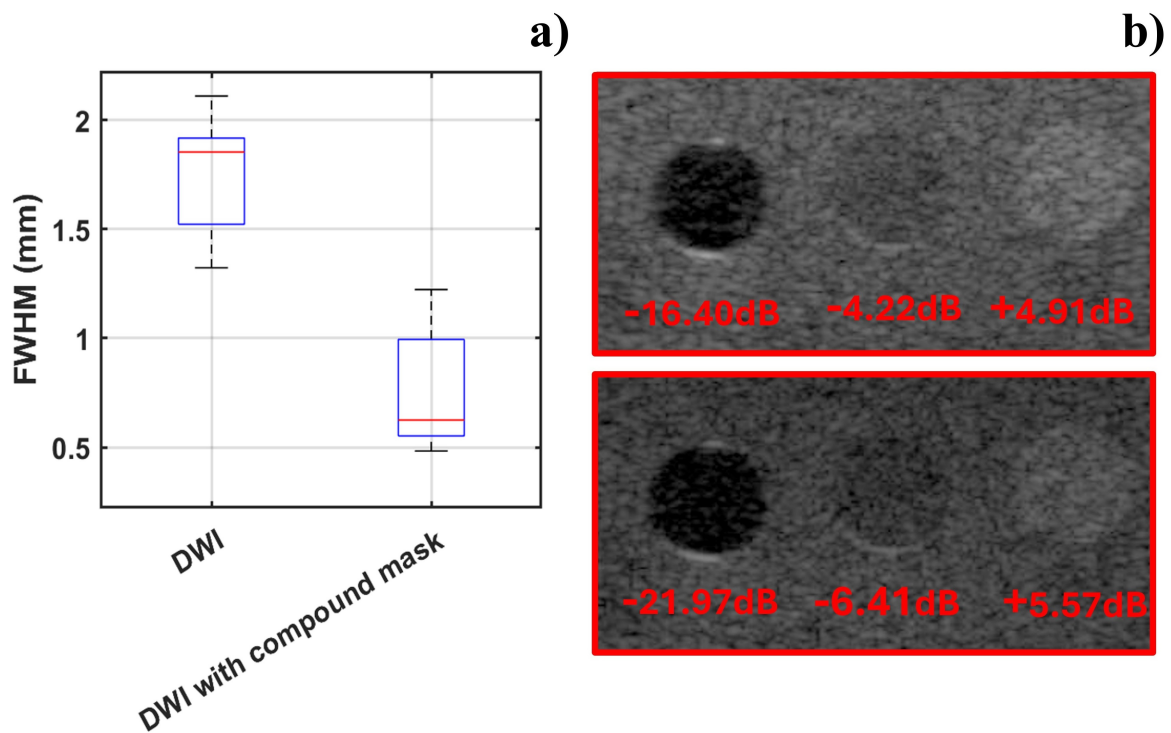


Fig. 4.4 a) Boxplots of the Full Width at Half Maximum (FWHM) measured on the five reflectors embedded in the multi-purpose phantom, shown in Figure 4.3. The FWHM values are compared for Diverging Wave Imaging (DWI) and DWI with the proposed compound mask using linear array (L7-4). b) Zoomed-in Figure for the Cysts presented in 4.3. From left to right anechoic cyst, -6dB cyst, +6dB cyst. The number under each cyst is the measured contrast ratio.

Method	anechoic	-6 dB	+6 dB
DWI	-16.40	-4.22	4.91
DWI with compound mask	-21.97	-6.41	5.57

Table 4.1 Contrast ratio (CR) values (in decibels) for three nominal target contrasts (anechoic, -6 dB, +6 dB) DWI, and DWI with the compound mask using linear array (L7-4).




Fig. 4.5 B-mode images generated using the C5-2 convex array probe and Verasonics scanner for Diverging Wave Imaging (DWI), and DWI with the proposed compound mask. The DWI images correspond to the three transmission strategies described in Figure 4.1: (a) linear distribution, (b) curvilinear distribution, and (c) tilted distribution.

4.3.2 B-mode Images using Convex probe C5-2

The B-mode images obtained using the convex array probe C5-2 are presented in Figure 4.5, showing DWI with and without the compound mask for the three different VS distributions described in Section 4.2.2 and illustrated in Figure 4.1(b-c-d). The corresponding FWHM and CR values are reported in the boxplots in Figure 4.6 and Table 4.2, respectively. As evident from the figures, each VS distribution yields different resolution and contrast measurements, reflecting a trade-off between these metrics. However, the application of the compound mask consistently improves both resolution and contrast, regardless of the VS distribution. This improvement is further confirmed by the FWHM values, which demonstrate that the performance of DWI with the compound mask outperforms that of standard DWI across all three distributions. Additionally, the improvements extend beyond resolution to include contrast, as evidenced by the better-defined cyst boundaries observed after applying the compound mask, with a contrast improvement of up to **6 dB** for anechoic cyst and up to **1.8 dB** for -6 dB and $+6$ dB cysts (Table 4.2). For enhanced visualization, all B-mode images were post-processed with linear intensity adjustments ($+20\%$ brightness gain, -40% contrast reduction), improving display clarity on the Verasonics interface.




Fig. 4.6 Boxplots of the Full Width at Half Maximum (FWHM) measured on the five reflectors embedded in the multi-purpose phantom, corresponding to the images in Figure 4.5 using a convex array. The FWHM values are compared for Diverging Wave Imaging (DWI) and DWI with the proposed compound mask. The DWI results are further categorized by the three transmission strategies: linear, curvilinear, and tilted.

Method	anechoic	-6 dB	+6 dB
Linear Distribution	-16.27	-3.42	3.2
Linear with mask	-21.08	-4.46	4.28
Curvilinear Distribution of VS	-18.04	-5.06	3.51
Curvilinear with mask	-23.74	-6.3	5.18
Tilted Distribution of VS	-18.68	-4.91	4.0
Tilted with mask	-24.65	-5.8	5.84

Table 4.2 Contrast ratio (CR) values (in decibels) for greyscale targets (anechoic, -6 dB, +6 dB) using various virtual source distributions using a convex array (C5-2).

4.3.3 Comparative Analysis of Computational Complexity

The proposed compound-mask apodization is *data-independent*: the spatial weights are precomputed offline from the probe geometry and virtual-source (VS) layout and then applied as scalar multipliers to the beamformed data, during coherent compounding. Consequently, the additional on-scanner computational load per frame is

$$\mathcal{O}(KM), \quad (4.2)$$

where K is the number of reconstructed pixels and M is the number of diverging/plane-wave transmissions used for compounding. This computational overhead is negligible with respect to the standard delay-and-sum beamforming, involving $\mathcal{O}(KNM)$ operations, where N is the number of active probe elements in reception.[120] [121][122].By contrast, *data-adaptive* beamformers estimate local signal statistics for each pixel (or subaperture) For example, the coherence-factor (CF) family computes a ratio of coherent to incoherent energy across channels,

$$\text{CF} = \frac{|\sum_{n=1}^N u_n|^2}{N \sum_{n=1}^N |u_n|^2},$$

and applies it as a weight to the DAS output[123] [124].In its classic 1-D form, CF adds $\mathcal{O}(N)$ operations per pixel (sums of u_n and $|u_n|^2$), yielding an overall order of $\mathcal{O}(KNM)$; however, extended coherence designs that explicitly form spatial-lag correlations can incur even higher per-pixel cost (often $\mathcal{O}(N^2)$)[125] [126].

Minimum-variance (MV, Capon) beamforming solves a per-pixel (or subaperture) quadratic program that reduces to computing and inverting a sample covariance matrix $\mathbf{R} \in \mathbb{C}^{N \times N}$, followed by a weight normalization [127] [128] [129]. Even with diagonal loading and snapshot reuse, the dominant costs are covariance formation $\mathcal{O}(N^2)$ and matrix inversion $\mathcal{O}(N^3)$ per update, yielding an overall complexity between $\mathcal{O}(KN^2M)$ and $\mathcal{O}(KN^3M)$ depending on

implementation and update rate. This gap explains the common need for GPU acceleration or reduced subapertures for real-time MV in imaging.

Table 4.3 summarizes these results. The proposed compound mask adds negligible overload with respect to standard DAS, while avoiding per-pixel statistical estimation and matrix operations across channels, which is advantageous for real-time diverging-wave imaging.

Table 4.3 Computational complexity (per frame) for common apodization strategies in diverging/plane-wave compounding. Here K is the number of image pixels, N the number of active receive channels, and M the number of transmissions.

Method	Main operation	Complexity
CF (classic)	Coherent/incoherent energy ratio	$\mathcal{O}(KNM)$
MV (Capon)	Covariance \mathbf{R} and \mathbf{R}^{-1} per pixel/subaperture	$\mathcal{O}(KN^2M)$ to $\mathcal{O}(KN^3M)$
Compound mask (proposed)	Precomputed mask applied	$\mathcal{O}(KM)$

4.3.4 Statistical significance testing for paired comparisons (one-sided)

To quantitatively assess whether the compound mask produces a statistically significant improvement in image resolution, we performed a paired, one-sided nonparametric hypothesis test on the full dataset of FWHM measurements. Each FWHM value was measured on the same scatterers (pins) before and after applying the compound mask, resulting in naturally paired observations. Because the compound mask is expected to reduce FWHM (indicating improved lateral resolution), the alternative hypothesis is directional (one-sided). Moreover, because the total number of measurements is relatively small ($n = 20$ pairs), and the underlying distribution of differences is unknown, a nonparametric approach is most appropriate. For these reasons, we use the paired sign test, which is exact, distribution-free, and robust for small sample sizes[130] [131][132].

Test formulation and p-value calculation. Let x_i and y_i denote the FWHM for the i -th scatterer without and with the compound mask, respectively, and define the paired difference:

$$d_i = y_i - x_i, \quad i = 1, \dots, n.$$

For this metric, an ‘‘improvement’’ occurs when $y_i < x_i$, i.e., $d_i < 0$. The sign test considers only the sign of each paired difference. Let S_- be the number of negative differences among the n_0 nonzero pairs ($n_0 \leq n$, after excluding any ties with $d_i = 0$). Under the null hypothesis H_0 (no systematic effect), positive and negative signs are equally likely and independent,

such that:

$$S_- \sim \text{Binomial}(n_0, 0.5).$$

The one-sided p-value for testing $H_1: \text{median}(d) < 0$ is then given by the upper tail of the binomial distribution:

$$p_{\text{sign}} = \mathbb{P}(S_- \geq s_-^{\text{obs}} | H_0) = \sum_{k=s_-^{\text{obs}}}^{n_0} \binom{n_0}{k} 2^{-n_0},$$

where s_-^{obs} is the observed number of negative differences.

The resulting one-sided p-value is:

$$p_{\text{sign}} = \sum_{k=15}^{20} \binom{20}{k} 2^{-20} = 0.041.$$

This value is below the conventional significance level $\alpha = 0.05$, indicating that the improvement in resolution is statistically significant. In addition to the p-value, we report the mean paired difference \bar{d} and the *Hodges–Lehmann* estimator (HL), defined as the median of $\{d_i\}$, both of which are negative, further supporting a significant reduction in FWHM.

Table 4.4 Paired one-sided sign test for FWHM (masked vs. unmasked) combining all 20 paired measurements from linear and convex probe experiments. Negative differences indicate improvement (masked < unmasked).

Metric	n_0	# negatives	Mean Δ (mm)	HL (mm)
FWHM (all scatterers)	20	15	−0.642	−0.597
One-sided sign-test p-value ($\alpha = 0.05$): $p_{\text{sign}} = 0.041$ (significant)				

The results are summarized in 4.4, which indicates that the compound mask consistently improves spatial resolution across different probe geometries and virtual source configurations. The significance of the result ($p_{\text{sign}} = 0.041 < 0.05$) confirms that the observed improvement is unlikely to be due to random variation.

4.4 Discussion and Conclusion

This work provides a detailed experimental validation of the compound mask-based apodization technique in DWI. The results highlight significant improvements in spatial resolution and contrast, maintaining real-time processing capabilities. Notably, the compound mask

method results in a 30% reduction in Full Width at Half Maximum (FWHM) measurements using a linear array incomparable with standard DWI. It also consistently improves resolution and contrast ratios using a convex array across various virtual source (VS) configurations, including linear, curvilinear, and tilted distributions.

The observed variations in image resolution and contrast across the different virtual source (VS) distributions indicate that the spatial arrangement of the virtual sources plays a critical role in determining image quality. Specifically, the linear distribution resulted in the best resolution, while the tilted configuration provided the highest contrast (see Figure 4.6 and Table 4.2). This further highlights the significant influence of VS geometry on image performance, suggesting that thoughtful optimization of the VS distribution can lead to notable enhancements in both resolution and contrast.

In contrast to data-driven iterative methods such as Minimum Variance Beamforming[133] or Independent Component Analysis (ICA)[134], the proposed compound mask approach eliminates the need for iterative updates or optimization techniques, which are computationally intensive. This significantly reduces the computational load and memory overhead, making it suitable for real-time applications in high-frame-rate applications. The ability to precompute the weights offline further aligns with hardware implementations on platforms such as Verasonics or ULA-OP [83], ensuring that this method can be deployed in practical scenarios without compromising processing time. These advantages approach is particularly well suited for clinical applications that require fast imaging, such as cardiac, abdominal, and vascular imaging [135][136] [137]. It is also important to acknowledge that the present study was conducted under ideal imaging conditions, without explicitly accounting for non-ideal propagation effects such as phase aberrations, attenuation, or inhomogeneities typically observed in biological tissues. The proposed compound mask apodization scheme is expected to exhibit a relatively high degree of robustness in realistic scenarios. This can be attributed to its data-independent nature: the apodization coefficients are derived from the interpolation of transmit apodization windows and remain fixed for a given imaging configuration. Consequently, the approach does not rely on the inversion of data-dependent covariance matrices, as occurs in adaptive beamforming techniques such as the Minimum Variance (MV) method. In the latter case, small phase distortions or noise fluctuations may lead to ill-conditioned matrices and unstable weighting, potentially resulting in image quality degradation under aberrated conditions. The next chapter will utilize these mapped weights in the optimization process to enhance the optimization approach to finalize the optimal virtual source distribution to further enhance the image quality metrics.

Chapter 5

Multi-Objective Optimization of Virtual Source Distributions for Ultrafast Diverging Wave Imaging

The studies in Chapters 3 and 4 focused on introducing the compound mask approach, validating it with in-silico data and in-vitro data using deterministic virtual sources. To further enhance image quality, in this chapter, we introduce a multi-objective optimization approach using genetic algorithms to find the optimal virtual source spatial arrangement. The chapter introduces the proposed multi-objective optimization methodology and details its integration with compound mask weighting to enhance image quality metrics such as resolution and contrast, the simulation setup and evaluation strategy used to assess the effectiveness of the optimized configurations, followed by extensive experimental validation using a Verasonics ultrasound scanner.

5.1 Introduction

DWs exploiting spherical wavefronts, emitted from virtual sources (VSs) placed behind the transducer, have been explored for convex array configurations. Ghigo et al.[138] investigated how the configuration of the VSs impacts image quality using a convex array. The observed performance was nearly equivalent, likely due to using only three compounding waves. An alternative method for synthesizing a wide, unfocused wavefront across the full aperture was introduced in [139][140]. This method employs a straightforward transmission strategy, a linear delay profile for convex arrays. The resulting wavefront was shown to follow an Archimedean spiral, varying with the steering angle. Another Fourier-domain beamforming

approach was proposed [141], using a tilted diverging wave transmission. While this approach demonstrated enhanced image quality, its effectiveness was restricted to regions with a narrow azimuthal angle. In [?], three transmission strategies for generating diverging waves were analyzed: lateral virtual-source-based Diverging Wave Imaging (latDWI), tilted virtual-source-based Diverging Wave Imaging (tiltDWI), and Archimedean-spiral-based Imaging (AMI). latDWI achieved superior resolution compared to the other methods. However, this came at the cost of stronger axial lobe levels, leading to the poorest image contrast. In contrast, both tiltDWI and AMI offered improved image contrast, albeit with a slight reduction in resolution.

Existing studies on DWs using convex arrays have primarily focused on deterministic distributions or specific transmission strategies. While previous approaches have demonstrated improvements in resolution and contrast, they are limited by their reliance on fixed configurations. For instance, deterministic distributions, such as uniformly spaced VSs, fail to account for the spatial variations in beam coherence and sidelobe levels across the imaging field. Similarly, transmission strategies like tilted diverging waves or Archimedean spirals, while effective in specific regions, are constrained by azimuthal dependencies and may not provide uniform performance across the entire field of view. In DWs, the spatial distribution and number of VSs critically influence fundamental performance trade-offs between frame rate and image quality. As each VS corresponds to a distinct transmission event, the achievable frame rate exhibits an inverse proportionality to the number of VSs. Increasing the number of VSs enhances imaging quality by improving spatial resolution and contrast through coherent compounding of multiple diverging wavefronts, however, at the cost of framerate. Therefore, achieving a high image quality while maintaining a high frame rate of original unfocused wave imaging is of high interest.

While prior studies have explored deterministic VS distributions, this study proposes a multi-objective evolutionary optimization framework to optimize VS spatial distributions, aiming to minimize the Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL). The optimization framework leverages insights from prior work in related fields, where multi-objective methods have been successfully applied to solve complex trade-offs. In sparse array configurations, optimization algorithms such as simulated annealing[85] and genetic algorithms[86] have been employed to enhance lateral resolution and suppress sidelobes while minimizing the number of active elements[87]. These methods have demonstrated significant improvements in imaging quality, particularly in applications requiring high spatial resolution and low artifact levels. Similarly, in acoustic source localization[95], and radar systems[142], evolutionary optimization techniques have been used to optimize array

geometries, ensuring accurate source detection even under challenging conditions, such as high noise levels or complex propagation environments .

The multi-objective genetic algorithm MOGA, from the Global Optimization Toolbox in MATLAB[143], has been chosen for optimization following better preliminary results compared to other approaches in [101]. The optimization framework was enhanced by integrating compound mask weights introduced in chapters 3 and 4, which were previously derived from synthetic aperture imaging (SAI) transmit apodization. The optimization was performed using varying numbers of virtual sources (20, 10, 6, and 4) chosen to explore the trade-off between image quality and frame rate systematically and to evaluate the method's adaptability in enhancing frame rate while preserving diagnostic image quality. The optimized VSs distributions were subsequently validated using both in silico data and experimental data. The proposed multi-objective optimization framework for virtual source (VS) distribution is designed as an offline calibration procedure. The optimization algorithm is executed prior to imaging using representative simulated data to identify optimal VS spatial configurations that balance image quality and frame rate. Once these optimized configurations are obtained, they are fixed and employed during real-time ultra-fast imaging acquisitions without any further computational overhead. Therefore, the approach is fully compatible with high frame rate imaging, as the computationally intensive optimization does not interfere with live data acquisition or processing. In summary, in this chapter :

- We propose a multi-objective optimization framework to optimize the spatial distribution of VSs in DWI, maximizing image quality metrics.
- We enhance the optimization algorithm by integrating compound mask weights, adjusting VSs' contributions to the image points during the optimization process.
- We demonstrate that the optimized VSs distributions achieve up to a **50% reduction in FWHM** and a **60% improvement in CR**, significantly enhancing both lateral resolution and contrast over deterministic distributions.
- We further show that the proposed method could work at a high frame rate with low transmission events.
- We validate the framework through extensive simulations and experimental trials using a Verasonics scanner.

5.2 Methodology

This section details the methodology developed for optimizing the spatial distribution of virtual sources (VSs) in DWI using a multi-objective genetic algorithm framework. It explains the optimization setup, including the search space definition, initialization, selection, crossover, mutation, and termination criteria, as well as the Pareto-based approach used to balance image quality metrics. The evaluation metrics employed, namely Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL), are described along with their formulation and role in guiding the optimization process. The section also presents the simulation setup used to generate synthetic data, the deterministic distributions employed for comparison, and the design of the synthetic phantom for quantitative assessment. Finally, it outlines the experimental setup for in-vitro validation using a Verasonics ultrasound scanner, demonstrating how the proposed optimization strategy was integrated and evaluated under realistic imaging conditions.

5.2.1 Multi-Objective optimization Framework

The genetic algorithm (GA) is a well-established optimization technique inspired by biological evolution processes[144]. In this study, GA was employed to optimize the distribution of VSs for convex arrays in DWI. The optimization process employs two fitness functions: one for the Full Width at Half Maximum (FWHM) and one for the Peak Sidelobe Level (PSL), which are critical metrics for assessing image quality in ultrasound imaging. To implement the optimization, the GA was executed in MATLAB using the `gamultiobj` function[145], and the overall framework is illustrated in Figure 5.1. The following subsections detail the key components of the optimization framework, including initialization, selection, crossover, mutation, termination criteria, and post-processing steps.

5.2.1.1 Initialization and Search Space

The algorithm starts by initializing a population of 1500 individuals, with each individual \mathbf{x}_i randomly sampled from a predefined search space \mathcal{X} . In our framework, the search space is defined by the lateral (x) and axial (z) positions of the N virtual sources (VSs) within predefined physical limits. The search space, or decision space[146], was restricted to a rectangular region to reduce the complexity of the objective functions and accelerate the optimization process[88]. The population was defined as:

$$P(0) = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{1500}\}, \quad \mathbf{x}_i \sim \mathcal{U}(\mathcal{X}), \quad (5.1)$$

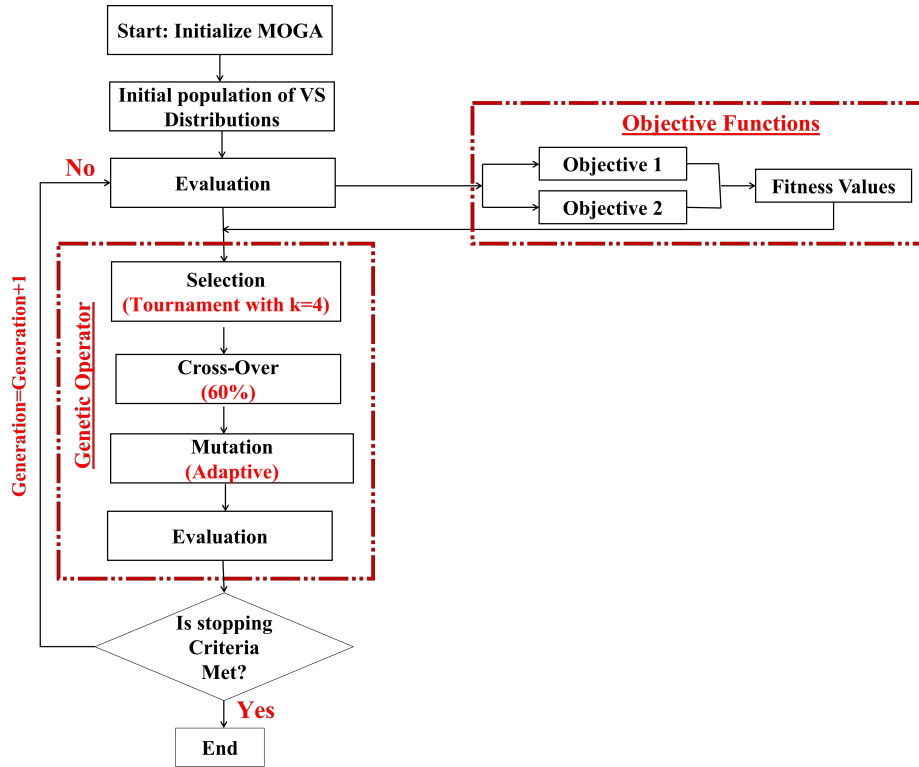


Fig. 5.1 Overview of the Multi-Objective Genetic Algorithm (MOGA) for optimizing virtual source (VS) distributions. The algorithm initializes a population of VSs distributions within a predefined search space and evaluates them using two objective functions. Evolutionary operations include tournament selection ($k=4$), crossover (60%), and adaptive mutation to refine the solutions iteratively. The process continues until convergence criteria are met, yielding an optimized set of VS configurations

where \mathcal{U} denotes a uniform distribution over the search space, and \mathbf{x}_i represents the i -th individual in the population. This random initialization promotes sufficient population diversity to support a thorough exploration of the decision space. Subsequent sections detail the fitness functions, the specific limits of the search space, and the imaging parameters adopted in this study.

5.2.1.2 Selection, Crossover, and Mutation

In multi-objective optimization, Pareto dominance is a fundamental concept where multiple conflicting objectives must be optimized simultaneously. Consider two candidate solutions, \mathbf{x}_i and \mathbf{x}_j , each evaluated by a vector of objective functions:

$$\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_m(\mathbf{x})],$$

where m represents the number of objectives. A solution \mathbf{x}_i is said to *dominate* another solution \mathbf{x}_j (denoted $\mathbf{x}_i \prec \mathbf{x}_j$) if and only if:

$$f_k(\mathbf{x}_i) \leq f_k(\mathbf{x}_j) \quad \forall k \in \{1, \dots, m\},$$

and $\exists k$ such that $f_k(\mathbf{x}_i) < f_k(\mathbf{x}_j)$. (5.2)

where f_k denotes the k -th objective function, \mathbf{x}_i is at least as good as \mathbf{x}_j in all objectives and strictly better in at least one. A solution that is not dominated by any other solution in the population is called a *non-dominated solution*. The collection of all such solutions forms the *first Pareto front*, also referred to as rank 1. After removing the first Pareto front, the next set of non-dominated solutions constitutes the *second front* (rank 2), and this process continues iteratively. Solutions with a higher Pareto dominance rank are generally closer to the optimal trade-off frontier and are less dominated by other solutions.

In the context of this study, there are two objectives:

$$\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x})], \tag{5.3}$$

where $f_1(\mathbf{x})$ represents FWHM and $f_2(\mathbf{x})$ represents the PSL that will be explained in detail in section II-B. The `gamultiobj` function automatically performs multi-objective selection based on Pareto dominance. This means that instead of ranking solutions by a single fitness score, the algorithm identifies non-dominated solutions—those that cannot be improved in one objective without degrading another. The selection process retains diverse solutions to explore trade-offs between resolution and contrast. To select individuals for reproduction, a tournament selection of size 4 was applied based on analytical studies in [147][148] that indicate tournament sizes in the range of 2–5 provide balanced selection pressure while preserving population diversity. In this process, subsets of 4 individuals were randomly selected, and the individual with the highest Pareto dominance rank within each subset was chosen as:

$$\mathbf{x}_{\text{selected}} = \arg \max_{\mathbf{x}_j} \{ \mathbf{f}(\mathbf{x}_j) \mid \mathbf{x}_j \in \text{Tournament Subset} \}, \tag{5.4}$$

where $\mathbf{f}(\mathbf{x}_j)$ represents the multi-objective fitness evaluation of individual \mathbf{x}_j . For crossover, 60% of the population was combined to produce offspring based on preliminary experiments with commonly observed parameter ranges (typically 60–90%) in genetic algorithm literature [149][150] that effectively balances exploration and exploitation for most applications. In our genetic algorithm, each individual \mathbf{x} represents a candidate solution encoding the

spatial coordinates of all N virtual sources:

$$\mathbf{x} = [x_1, z_1, x_2, z_2, \dots, x_N, z_N],$$

where x_n and z_n denote the lateral and axial positions of the n -th virtual source, respectively. The crossover operation combines two parent individuals, \mathbf{x}_A and \mathbf{x}_B , to generate an offspring $\mathbf{x}_{\text{child}}$ as follows:

$$\mathbf{x}_{\text{child}} = \alpha \mathbf{x}_A + (1 - \alpha) \mathbf{x}_B, \quad \alpha \in [0, 1], \quad (5.5)$$

where α is a random weighting factor that introduces variability while preserving key traits from the parent solutions. The offspring inherit beneficial characteristics from both parents, promoting convergence toward optimal solutions. Then, a mutation was applied to the individuals. The mutation operation for a *gene* x_j within a chromosome \mathbf{x} that corresponds to a single coordinate value (either lateral or axial) of one virtual source. The mutation operation perturbs this gene by adding a small random offset δ :

$$x_j^{\text{mutated}} = x_j + \delta, \quad \delta \sim \mathcal{U}([- \Delta, \Delta]), \quad (5.6)$$

where $\mathcal{U}([- \Delta, \Delta])$ denotes a uniform distribution over the interval $[- \Delta, \Delta]$. It represents the feasible mutation range dynamically adjusted during optimization by `mutationadaptfeasible` function, which modifies the size of the mutation step based on the boundaries of the variables and the current search space to ensure that the algorithm explores the search space effectively without premature convergence.

5.2.1.3 Termination Criteria and non-dominated solutions (NDSs)

The algorithm terminated under the following criteria:

- A maximum of 500 generations was reached.
- No significant improvement in the NDSs was observed over 50 consecutive generations, with a tolerance of 5%.

The chosen criteria balanced computational efficiency and thorough exploration of the search space, ensuring that the algorithm converged to high-quality solutions without excessive computational overhead. Fitness evaluation was performed for each individual to compute the objectives of FWHM and PSL, defined in the next subsection. The optimization sought non-dominated solutions, representing trade-offs between FWHM and PSL, providing valuable insights into the trade-offs between resolution and contrast, enabling the selection of optimal VS configurations for specific imaging scenarios.

5.2.1.4 Parallel Computing and Clustering

To accelerate the evaluation process, MATLAB’s Parallel Computing Toolbox[151] was used, significantly reducing runtime by enabling parallel fitness calculations. Once the non-dominated solutions (NDSs) were identified, k-means clustering was applied using MATLAB’s Statistics and Machine Learning Toolbox[152] based on similarity in objective space to reduce redundancy among the final NDSs solutions.

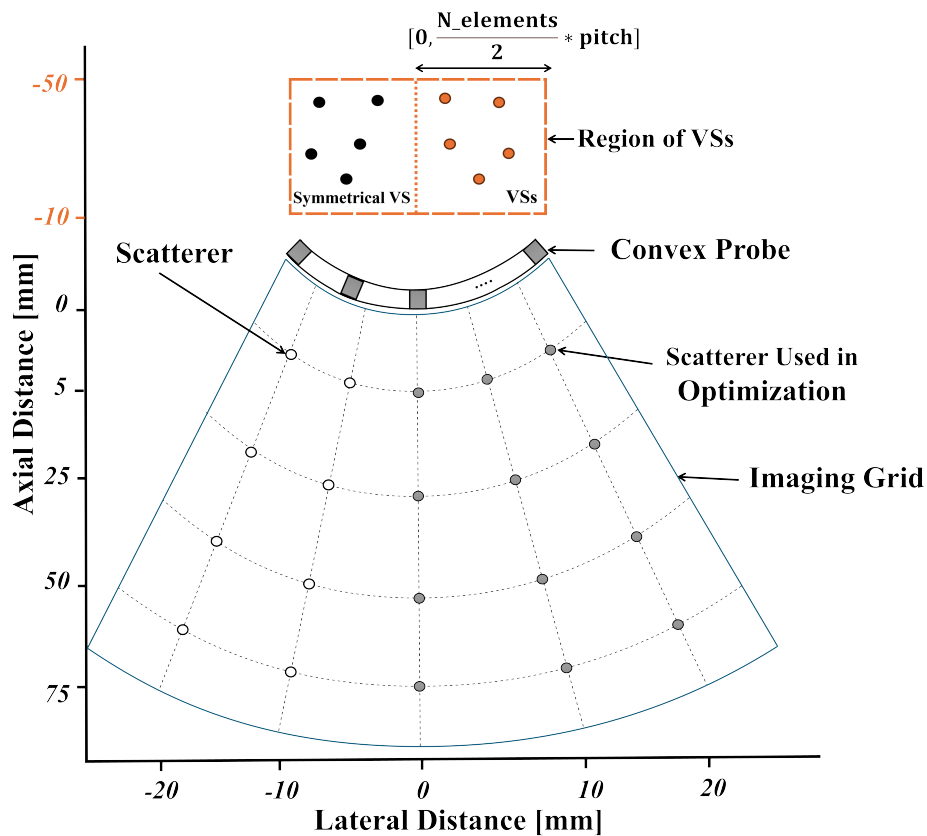


Fig. 5.2 Study configuration from top to bottom: the arrangement of VSs, where the orange VSs indicate those selected for optimization; a convex probe with N elements; scatterer positions, with grey ones used in the optimization process; and the imaging grid defining the region of interest.

5.2.2 Evaluation Metrics

The imaging performance of the virtual source (VS) distributions was quantitatively assessed using two key metrics: *Full Width at Half Maximum (FWHM)* and *Peak Sidelobe Level (PSL)*. These metrics are widely recognized in ultrasound imaging for evaluating lateral resolution and contrast, respectively, and were employed to ensure a comprehensive assessment of

image quality across the imaging volume. The FWHM quantifies the lateral resolution of the imaging system by measuring the width of the main lobe of the Point Spread Function (PSF) at -6 dB relative to its peak intensity[153]. The PSL measures the contrast degradation caused by sidelobes [?], which is critical to distinguish low-contrast structures such as cysts or lesions. The PSL was computed as the ratio of the amplitude of the highest sidelobe to the amplitude of the main lobe, expressed in decibels (dB). The sidelobe peaks were identified using a peak detection function, `findpeaks`[154], ensuring precise localization of the highest sidelobe.

In the optimization process, to evaluate the imaging performance across the field of view, FWHM and PSL values were computed for a set of scatterers located at different depths and lateral positions within the imaging volume to ensure a robust assessment of imaging quality across the entire volume, capturing variations in resolution and contrast at different spatial locations. The values of FWHM and PSL were concatenated into two datasets, denoted as $\mathcal{D}_{\text{FWHM}}$ and \mathcal{D}_{PSL} , respectively:

$$\mathcal{D}_{\text{FWHM}} = [\text{FWHM}_1, \text{FWHM}_2, \dots, \text{FWHM}_N](mm) \quad (5.7)$$

$$\mathcal{D}_{\text{PSL}} = [\text{PSL}_1, \text{PSL}_2, \dots, \text{PSL}_N](dB) \quad (5.8)$$

where N represents the total number of scatterers used for evaluation. The optimization process aimed to minimize two objective functions derived from the FWHM and PSL datasets. The first objective function, $f_1(\mathbf{x})$, was defined as the mean value of the FWHM data set, $\mathcal{D}_{\text{FWHM}}$, which represents the average lateral resolution across the image volume. The second objective function, $f_2(\mathbf{x})$, was defined as the mean value of the PSL dataset, \mathcal{D}_{PSL} , representing the average sidelobe level. These objectives were formulated as:

$$f_1(\mathbf{x}) = \text{avg}(\mathcal{D}_{\text{FWHM}}), \quad (5.9)$$

$$f_2(\mathbf{x}) = \text{avg}(\mathcal{D}_{\text{PSL}}), \quad (5.10)$$

where $\text{avg}(\cdot)$ denotes the arithmetic mean. By minimizing Objective_1 , the optimization process sought to enhance lateral resolution, while minimizing Objective_2 aimed to reduce sidelobe levels, thereby improving contrast and reducing image artifacts. The use of average values for the objective functions prioritizes consistent imaging performance across the field of view, avoiding biases that could arise from extreme values in the datasets.

5.2.3 Simulation Setup

The received raw data were generated using the Field II software[116][117], using a 128-element convex array transducer with a center frequency of $f_0 = 3$ MHz. The transducer was excited with a 2-cycle sinusoidal signal and with a sampling frequency of 12 MHz. A Hanning window was applied for both transmit (TX) and receive (RX) apodization, with the window scaled according to the aperture size in receive, the latter being set according to the focal number equal to 0.7. Compound mask weights, defined as functions of the virtual source (VS) positions, were mapped directly from synthetic aperture imaging transmit apodization to ensure that only contributions from properly insonified regions were retained. Subsequently, the simulated RF signals were processed using a standard time-domain delay-and-sum beamforming algorithm. The beamforming process was followed by coherent recombination of data related to each VS position, weighted by the above-defined compound mask, envelope extraction, and logarithmic compression to produce B-mode images. The optimization setup included 20 point scatterers distributed at various depths of [5, 25, 50, 75]mm along the lateral positions of the axial direction $[-20, -10, 0, 10, 20]$ mm. All scatterers were assigned equal amplitudes to ensure uniform reflectivity. Figure 5.2 illustrates the organization of the imaging grid and the distribution of the virtual sources (VS) utilized in the multi-objective optimization process. The lateral positions (x) of the VSs were constrained between 0 mm and 27.2 mm, corresponding to half the total lateral width of the convex array. Axial positions (z) were restricted to a range between -50 mm and -10 mm. These constraints ensured that each VS remained within the practical imaging boundaries. To reduce computational complexity, symmetry about the transducer's central axis was exploited: only half of the total number of VSs was explicitly optimized, and the remaining VSs were mirrored. Consequently, configurations with 20, 10, 6, and 4 VSs were optimized by considering only 10, 5, 3, and 2 VSs, respectively. This approach effectively halves the dimensionality of the decision space; rather than optimizing $2N$ parameters (the x and z positions for N VSs), only N are considered. Similarly, the symmetry reduces the number of scatterers needed, from an initial set of 20 scatterers, only 12 were required for the optimization since symmetrical scatterer positions yield identical image-quality metrics.

5.2.4 Deterministic Distributions

For direct comparison with the optimized genetic algorithm solutions, two deterministic configurations of 20 virtual sources (VSs) were evaluated: a linear arrangement and a curvilinear arrangement selected based on configurations investigated in prior studies [118], serving as established references for comparison. In both cases, all VSs were positioned at

the same axial distance of $R_{\text{convex}} = 50$ mm from the probe surface. In the linear arrangement, the 20 VSs were uniformly spaced across a lateral range spanning -32 to $+32$ times the transducer pitch (pitch = $425.56 \mu\text{m}$). For the curvilinear arrangement, the VSs follows a curvilinear pattern that aligns with the geometry of the probe, with angular placement of the sources confined within a maximum aperture angle of 33° , distributed symmetrically about the central axis of the probe. Figures 6(a) and 6(b) shows the spatial distribution of linear and curvilinear VS distributions, respectively.

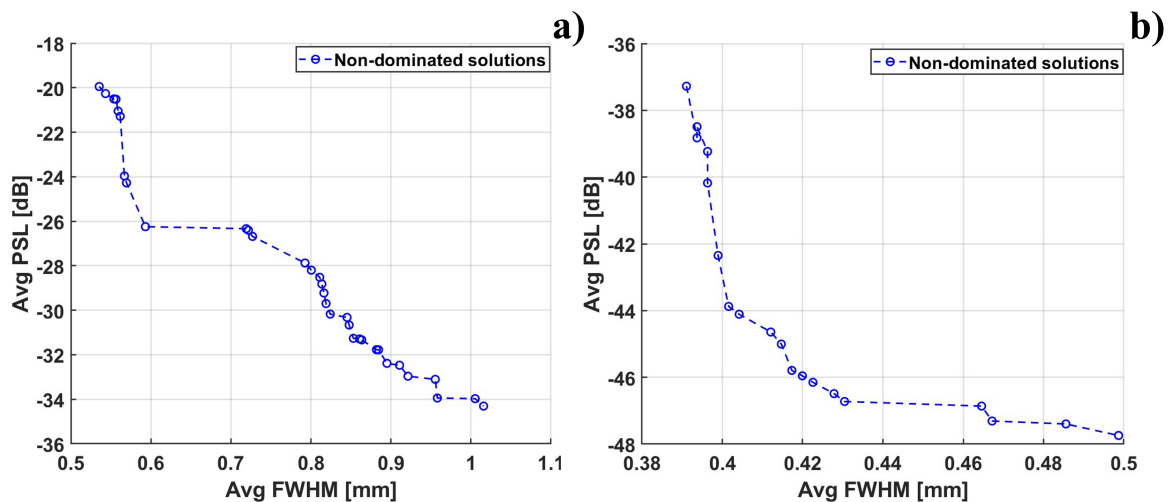


Fig. 5.3 Impact of the compound mask on the optimization of Virtual Source (VS) distributions using a Genetic Algorithm (GA):(a) Pareto front showing the average values of Full Width at Half Maximum (FWHM) and Peak Sidelobe Level (PSL) obtained from GA optimization without the compound mask.(b) Pareto front obtained from GA optimization with the compound mask

1

5.2.5 Phantom with Anechoic Cysts and Reflectors

To evaluate the imaging performance of the optimized virtual source (VS) distributions, a synthetic phantom was designed to simulate a tissue-mimicking environment for 2D ultrasound imaging. The phantom dimensions were defined to match the convex transducer footprint, with its lateral extent determined by the number of elements and their respective pitch. The axial depth extended to 70 mm, starting from 5 mm below the transducer surface. Within this domain, scatterers were randomly distributed with lateral and axial

¹The content of this chapter has been submitted to *Ultrasonics*: Z. Alzein, M. Crocco, D. D. Caviglia, and H. Liebgott, "Multi-Objective Optimization of Virtual Source Distributions for Ultrafast Diverging Wave Imaging," *Ultrasonics*, <http://dx.doi.org/10.2139/ssrn.5378898>

positions sampled from a uniform distribution. Their amplitudes followed a Gaussian distribution to mimic the heterogeneous nature of biological tissues. The synthetic phantom includes anechoic cysts and high-amplitude reflectors to assess contrast and spatial resolution. Specifically, 9 anechoic cysts—scatterer-free regions—were positioned at axial depths of 20 mm, 40 mm, and 60 mm. The cysts of radii of 2 mm, 3 mm, and 4 mm, respectively, with lateral positions spanning from -20 mm to 20 mm. The regions ensure the creation of distinct low-echo zones, enabling a quantitative assessment of contrast on envelope images. The contrast ratio (CR) was computed as:

$$CR = 20 \log_{10} \left(\frac{\mu_{\text{in}}}{\mu_{\text{out}}} \right), \quad (5.11)$$

where μ_{in} and μ_{out} denote the mean pixel intensities inside and outside the cyst region, respectively. The high-amplitude point reflectors were placed at axial depths of 10 mm, 30 mm, and 50 mm, with lateral positions ranging from -20 mm to 20 mm. Their amplitude was scaled by a factor of 10 relative to the background scatterers, ensuring strong reflectivity. All deterministic and optimized VS distributions obtained through the multi-objective optimization framework were used to reconstruct B-mode images of the synthetic phantom. The simulation process was conducted under the same imaging configuration described previously in section II.C.

5.3 Results and Discussions

5.3.1 Optimization Results with and without Compound Mask

The set of non-dominated solutions (NDSs) commonly referred to as the Pareto-optimal front[155][156], resulting from the optimization of virtual source (VS) distributions, with and without the application of a compound mask, are illustrated in Figure 5.3(a) and Figure 5.3(b), respectively, using 20 VSs. These plots demonstrate the trade-offs between average FWHM and average PSL, where a decrease in average PSL generally accompanies an increase in average FWHM, and vice versa. The set of (NDSs) obtained using the compound mask consistently outperforms those without the mask across both metrics. The Pareto fronts of NDSs, exhibit discontinuities. Such discontinuities are commonly observed in multi-objective optimization problems and arise due to the inherent structure of the non-dominated solution set (NDS)[98]. The use of the compound mask significantly enhanced imaging quality, as reflected in the improved trade-offs between FWHM and PSL. Without the compound mask, the average FWHM ranged from 0.6 mm to 1.0 mm, while the average PSL spanned -18 dB

to -34 dB. In contrast, applying the compound mask improved values of 0.38 mm to 0.5 mm for FWHM and -36 dB to -48 dB for PSL, corresponding to improvements of 37% and 41%, respectively. These enhancements are attributed to the weighting mechanism introduced by the compound mask, which prioritizes contributions from regions exhibiting higher beam coherence while effectively suppressing sidelobe artifacts. Unlike conventional coherent compounding methods, which fail to adapt to varying insonification conditions, the compound mask adjusts the transmission weights for each virtual source position, addressing the limitations of traditional approaches, particularly in convex arrays, where the insonification field is inherently non-uniform.

5.3.2 Optimization Results with Different Number of Virtual Sources

Figure 5.4 presents the optimal Pareto fronts obtained for configurations with 20, 10, 6, and 4 VSs, illustrating the trade-off between the Avg FWHM and the Avg PSL. As expected, the configuration with 20 VSs offers the most favorable balance between resolution and contrast, achieving Avg FWHM values between 0.38mm and 0.55mm and Avg PSL levels approaching -50dB. However, reducing the number of VSs does not severely compromise image quality, thanks to the robustness of the proposed optimization framework. For instance, with 10 VSs, the Pareto front shifts to achieve an Avg FWHM below 0.2mm, while Avg PSL degrades by roughly 3–5dB relative to the 20-VS configuration. A similar trend is observed with 6 VSs, where the Avg FWHM remains in the range of less than 0.45mm and the Avg PSL hovers around -45dB. Even with just 4 VSs, the method continues to preserve image quality, maintaining Avg FWHM values typically below 0.6mm and Avg PSL near -45dB.

5.3.3 Selection of Representative Pareto-Optimal Solutions

To assess the impact of different optimized VS configurations on imaging performance, three specific solutions were selected from the Pareto front for further validation:

- **Optimized for FWHM(optFWHM):** Prioritizing minimum beamwidth, offering higher lateral resolution at the expense of increased PSL.
- **Optimized for PSL(optPSL):** Prioritizing lower sidelobe levels, improving contrast while allowing for a slightly broader main lobe.
- **Mid-range trade-off(midopt):** A balanced solution, offering an intermediate compromise between resolution and contrast.




Fig. 5.4 Pareto-optimal front for different numbers of Virtual Sources (VSs) in the optimization process: 20 VSs (blue), 10 VSs (orange), 6 VSs (purple), and 4 VSs (green).

Figure 5.5 illustrates the chosen solutions based on visual inspection of the objective space and corresponding image quality metrics from the Pareto front for the 20 VS configuration, plotted alongside deterministic VS distributions (linear and curvilinear) previously described in section II-D. For lower VS configurations (10, 6, and 4 VSs)(Figure 5.5(b)), the selection strategy was adapted based on the trade-offs observed in the Pareto front. The optimized FWHM solutions for these configurations exhibited significantly higher PSL levels, which would lead to degraded image quality. Therefore, only solutions corresponding to the mid-range trade-off and optimized PSL were selected for validation. These choices ensure that the imaging performance remains comparable to the 20 VS configuration. The spatial distributions of the selected VS configurations across all tested cases (20, 10, 6, and 4 VSs) are visualized alongside the deterministic ones in Figure 6. The solutions chosen for the midoptFWHM configuration in Figure 6—panels (d), (f), (h), and (j) corresponding to 20, 10, 6, and 4 VS transmissions, respectively—illustrate spatial distributions where one of the VS is placed along the central axis. As a result, its symmetry VS will be aligned with it. This central alignment was previously demonstrated in [16], which showed that positioning one VS at the center is advantageous. However, in our study, such centering was not necessarily in optimizing PSL levels, as evidenced by the optimized configurations in Figure 5.6 panels (e), (g), (i), and (k), where the sources are distributed asymmetrically around the axis. All the selected solutions were evaluated through both simulation-based and experimental validation and analyzed quantitatively using metrics such as *FWHM*, and *contrast ratio (CR)* to validate performance.

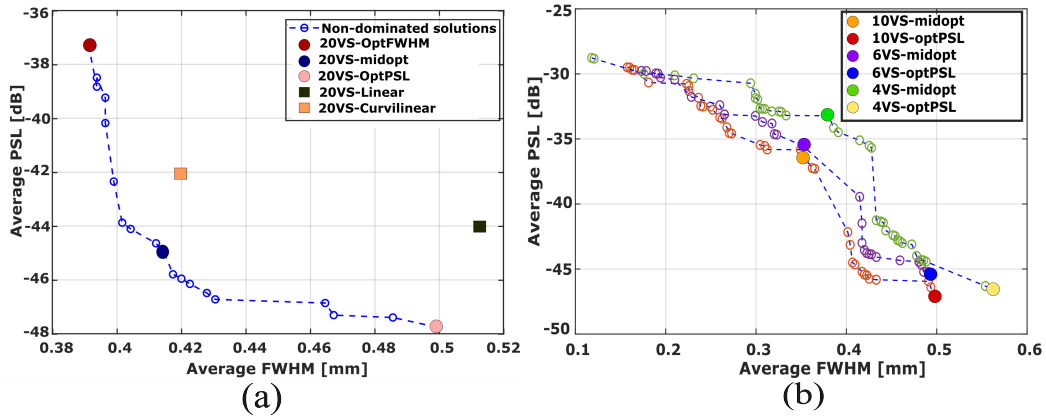


Fig. 5.5 The Pareto front for 20 VSs (a) includes three selected solutions: **optFWHM** (red), **midopt** (dark blue), and **optPSL** (light pink) alongside deterministic configurations: **20VS-linear** (green square) and **20VS-curvilinear** (orange square). (b) represents the Pareto fronts for reduced VS configurations (10, 6, and 4 VSs) feature two selected solutions—**midopt** and **optPSL**—for each case. The selected solutions for 10 VSs (orange and red), 6 VSs (purple and blue), and 4 VSs (green and yellow) respectively

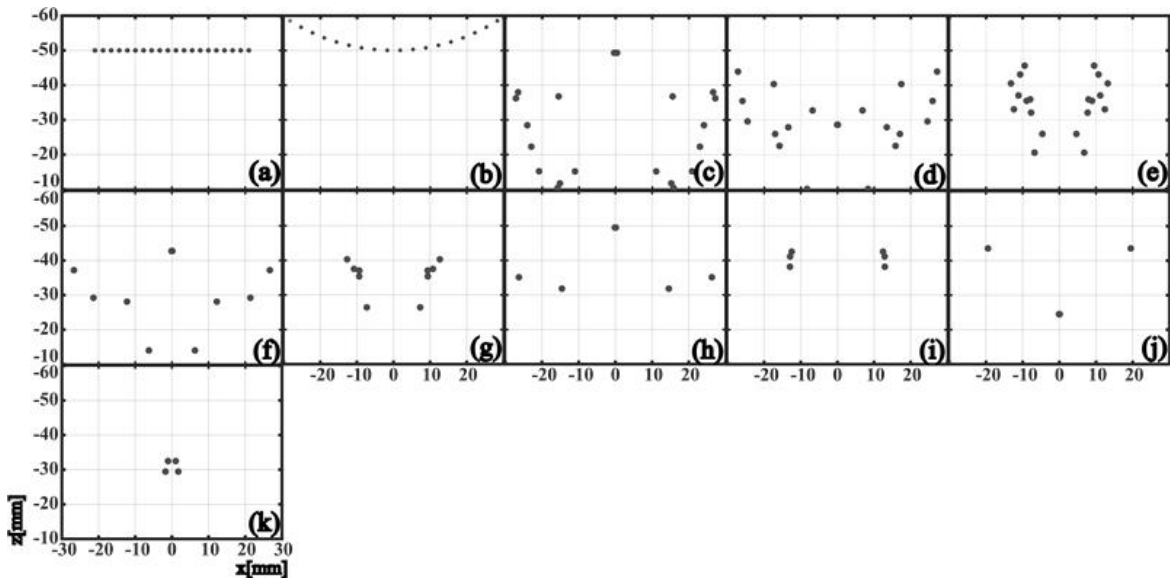


Fig. 5.6 Spatial distribution of VSs for both deterministic and optimized distributions chosen for validation: (a) 20VS-Lin, (b) 20VS-Curv, (c) 20VS-optFWHM, (d) 20VS-midopt, (e) 20VS-optPSL, (f) 10VS-midopt, (g) 10VS-optPSL, (h) 6VS-midopt, (i) 6VS-optPSL, (j) 4VS-midopt, (k) 4VS-optPSL.

5.3.4 Validation of the optimized distributions Using Silico Data

The reconstructed B-mode images for the synthetic phantom, described in Section II-E, are presented in Figure 5.7. These images illustrate the performance of both optimized chosen

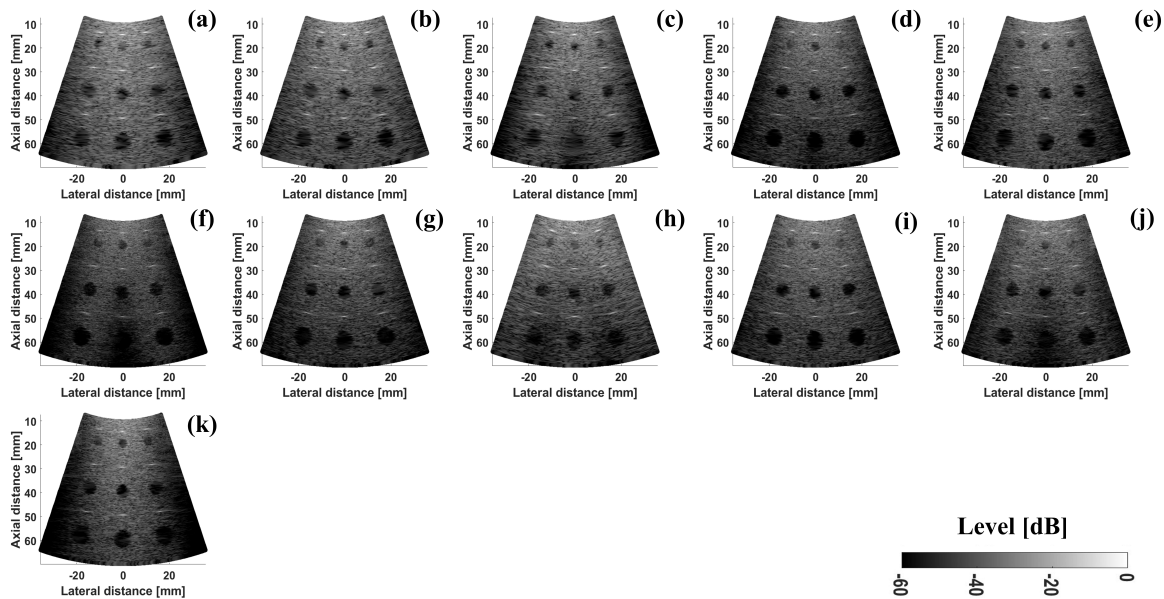


Fig. 5.7 B-mode images of the simulated phantom with anechoic cysts and reflectors for (a) 20VS-Lin, (b) 20VS-Curv, (c) 20VS-optFWHM, (d) 20VS-midopt, (e) 20VS-optPSL, (f) 10VS-midopt, (g) 10VS-optPSL, (h) 6VS-midopt, (i) 6VS-optPSL, (j) 4VS-midopt, and (k) 4VS-optPSL.

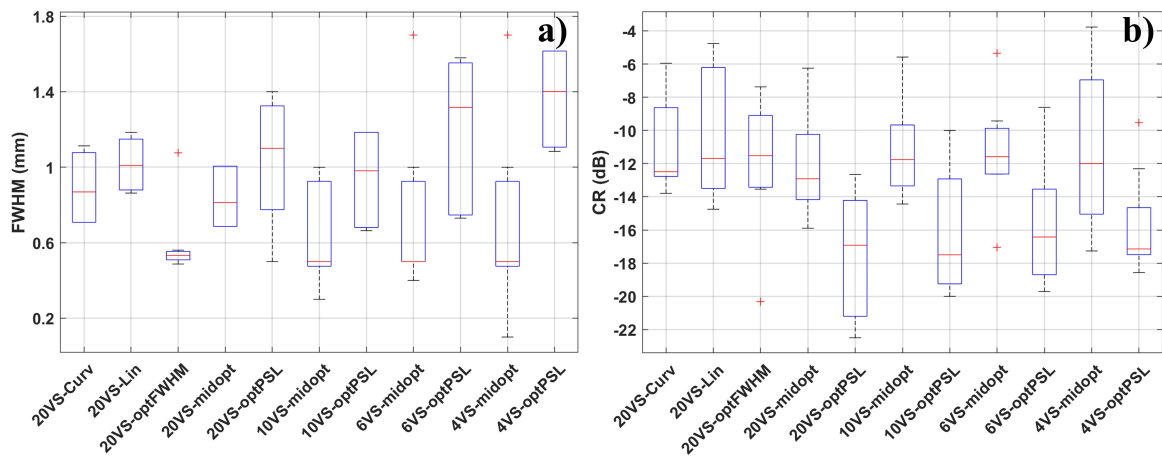


Fig. 5.8 Boxplots for (a) FWHM measured on the 9 reflectors (b) CR measured on 9 anechoic cysts from the phantom images (Fig. 5.7) for all selected VSs.

solutions and deterministic VS configurations. The FWHM was measured on high-intensity reflectors, while the CR was computed for anechoic cysts, as defined by Equation (10). The quantitative results are summarized using boxplots in Figure 5.8, showing the values of FWHM(a) and CR(b) across different imaging regions. The optimized solution for FWHM (**optFWHM**) with 20 VSs achieved a median FWHM of 0.5 mm, a 50% reduction compared to curvilinear and linear deterministic distributions (1.0 mm and 1.1 mm, respectively). Similar performance was observed for mid-optimal (**midopt**) solutions with fewer VSs: 10 VSs achieved a median FWHM of 0.55 mm (45–50% reduction), while 6 VSs and 4 VSs maintained FWHM values of 0.6 mm and 0.65 mm, respectively. The optimized configurations with fewer VSs outperform deterministic configurations with 20 VSs. For CR, the optimized solution for PSL (**optPSL**) with 20 VSs achieved a median CR of -16 dB, a 25–30% improvement over deterministic curvilinear and linear distributions (-12 dB and -10 dB, respectively). Similar trends were observed for reduced (**optPSL**) VS configurations: 10 VSs maintained a CR of -15 dB, while 6 VSs and 4 VSs achieved -15.5 dB and -14 dB, respectively. Optimized solutions consistently outperformed deterministic distributions, even with fewer transmission events.

5.3.5 Validation of the optimized distributions Using Multi-purpose Phantom

5.3.5.1 Experimental-SetUp

To further test the robustness of the optimized VS, the sequences are then implemented a Vantage 256 ultrasound research scanner (Verasonics, Redmond, WA, USA) and we used CA5-2 curved array with 128 elements ($Radius = 41.2$ mm, $pitch = 424.56$ m, $f_0 = 3.1$ MHz). The excitation parameters are set the same as in section 5.2.3. The sampling rate was set to 12 MHz, and imaging was performed on the 410SCG LE Multi-Purpose Accreditation Phantom. The acquired RF data were subsequently beamformed and processed to evaluate imaging performance metrics contrast, and resolution. The full aperture is used for transmission.

5.3.5.2 B-mode Image acquired using Verasonics scanner

All the optimized distribution validated by simulation are implemented on the Verasonics scanner and validated in [157]. Here, we are presenting the most outperforming ones with comparison with the deterministic (linear and curvilinear). The B-mode images of the tissue-mimicking phantom obtained with the deterministic distribution are shown in Figure 5.9,




Fig. 5.9 B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for left: 20VS-Lin, right: 20VS-Curv.




Fig. 5.10 B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for 20VS-optFWHM, 20VS-midopt, 20VS-optPSL

and with 20 VS optimized in Figure 5.10, and with a reduced number of VSs in Figure 5.11, illustrating the performance of both deterministic and chosen optimized(VSs) distributions.

The FWHM was measured on the point scatterers, and the CR ratio was measured on the anechoic cyst located at 45mm. The numerical results for FWHM and CR are summarized in a box plot in Figure 5.12 and Table 5.1, respectively. The **optFWHM** solution with 20 VSs achieved a median FWHM of 0.9 mm, representing reductions of 35.7% and 28.6% from the linear (**20VS-Lin**, 1.4 mm) and curvilinear (**20VS-Curv**, 1.26 mm) configurations, respectively. The mid-optimal solution (**20VS-midopt**) maintained a median FWHM of 1.0 mm, corresponding to reductions of 28.6% and 20.6%. Similar trends were observed with fewer VSs: the **10VS-midopt** achieved an FWHM of approximately 0.6 mm (52 % reduction), while the 4 VS midopt solutions demonstrated approximately FWHM values of 0.8 mm (36% reduction) and 0.7 mm (44% improvement over **20VS-Curv**), respectively. CR measurements (Table 5.1) indicate enhanced contrast performance for optimized distributions. The **optPSL** solution with 20 VSs achieved a CR of -26 dB, a 60% improvement over the curvilinear deterministic distribution (-8.1 dB). The mid-optimal solution (**20VS-midopt**)




Fig. 5.11 B-mode images of the multi-purpose phantom acquired from the Verasonic scanner for 10VS-midopt, 4VS-midopt

reached -27 dB, a 62.5% improvement. For reduced VS configurations, the **optPSL** solution with 10 VSs achieved -20.3 dB (40% improvement), while 4 VS configurations yielded CR values of -17.9 dB (32.4% improvement), respectively. The experimental results confirm that the proposed method demonstrates robust performance on the Verasonics Vantage 256 scanner, closely aligning with the trends observed in the simulations.

Table 5.1 CR values measured on anechoic cyst located at 45mm from the phantom images(Fig.5.9,5.10 and 5.11)for selected VSs.

Virtual Source Distribution	CR [dB]
20VS-Lin	-16.66
20VS-Curv	-8.1027
20VS-optFWHM	-15.21
20VS-midopt	-27
20VS-optPSL	-26
10VS-midopt	-20.3
4VS-midopt	-17.35

5.4 Conclusion

This chapter presented a multi-objective optimization framework for designing virtual-source (VS) distributions in ultrafast diverging-wave imaging. By coupling a genetic algorithm with compound-mask weights, the approach explicitly balances lateral resolution (FWHM) and sidelobe suppression (PSL) and yields Pareto-optimal VS layouts that outperform deterministic arrangements. Across simulations and phantom experiments on a Verasonics




Fig. 5.12 Boxplots for FWHM measured on the 6 reflectors from the phantom images(Fig.5.9,5.10 and 5.11)for selected VSs

system, the optimized configurations delivered substantial gains—up in image quality metrics while preserving robustness at reduced numbers of transmissions, thus supporting high-frame-rate operation. The optimization is executed offline, and the resulting VS distributions are fixed at run time, making the method practical for real-time imaging. The optimized VS sets, together with the compound mask, are subsequently integrated into the GPU-based receive beamformer described in the next chapter to enable efficient, high-quality ultrafast imaging.

Chapter 6

Optimization-Driven GPU-Based receive Beamformer for Efficient and High-Quality Ultrafast Diverging Wave Imaging

This chapter presents the implementation of a computationally efficient GPU receiver while preserving image quality by using image quality enhancement optimization strategies introduced in 4,5, and 6.

6.1 Introduction

In ultrafast imaging, the ultrafast unfocalized transmission by all channels comes at the cost of an increased computational load, memory occupancy, and required transfer bandwidth, implying significant challenges for real-time implementation. Traditional RF beamforming for both focalized and unfocalized transmission involves applying time delays to RF signals from the transducer's channels for image reconstruction. In the late 1970s, baseband I/Q beamformers were introduced, which demodulate RF signals to the baseband and apply time delays and phase rotations. This approach allows the downsampling of baseband signals, reducing the cost of the delay-and-sum process [158]. Various hardware architectures, such as bandwidth (BW) sampling, pipelined sampling, and direct sampling, have been proposed to implement the I/Q beamformer. These approaches have successfully reduced computational complexity and power consumption, making them particularly useful for portable or high-frequency systems [159][160]. Moreover, the I/Q beamformer can decrease the data

transfer rate required for transmitting channel data from hardware to software in software-based beamforming systems [161][162]. This reduction in channel data is advantageous for implementing parallel beamformers, such as plane/diverging wave beamforming, which requires the storage of large volumes of channel data [163]. Recent advancements in graphics processing unit (GPU) technology have enabled the implementation of various ultrasound algorithms in software, including plane wave imaging [164]. GPUs, with their high memory bandwidth and massively parallel multiprocessing cores, offer outstanding computational power, making them suitable for fully PC-based programmable ultrasound systems. This chapter objectives can be summarized as follows :

- It presents a real-time implementation of coherent plane-wave/diverging-wave ultrasound imaging receive beamformer under the constraints of a commercial ultrasound scanner, using GPU acceleration.
- An efficient beamforming pipeline is introduced, integrating multiple optimizations (precomputed delays, transmit/receive decoupling, symmetry exploitation, and base-band processing) to drastically reduce computation and data transfer overhead.
- Comprehensive experimental validation is provided, including phantom studies with a Verasonics Vantage system and C5-2 probe, demonstrating that our approach achieves high image quality (resolution and contrast) comparable to standard imaging while operating at real-time speeds.

6.2 Methodology

6.2.1 RF-Band and Baseband Receive Beamformer

In ultrasound diverging wave imaging (DWI), a sequence of spherical waves is transmitted, each one with a virtual source (VS) located behind the probe. The raw data received for each transmission are converted into analytic signals by Hilbert filtering, dynamically delayed and multiplied by a proper apodization function, and summed over the channels for each point on the image grid to be reconstructed; next, a further sum is performed over the transmission index to coherently combine the images related to each transmission. For B-mode imaging, the processing chain continues with envelope extraction and logarithmic compression on each image point. The process can be resumed in the following equation:

$$I(x, z) = \sum_{n=1}^N \sum_{h=H_{\min}}^{H_{\max}} S_n(h, \mathcal{T}(x, z, h, n)) \cdot a(x, z, h, n) \quad (6.1)$$

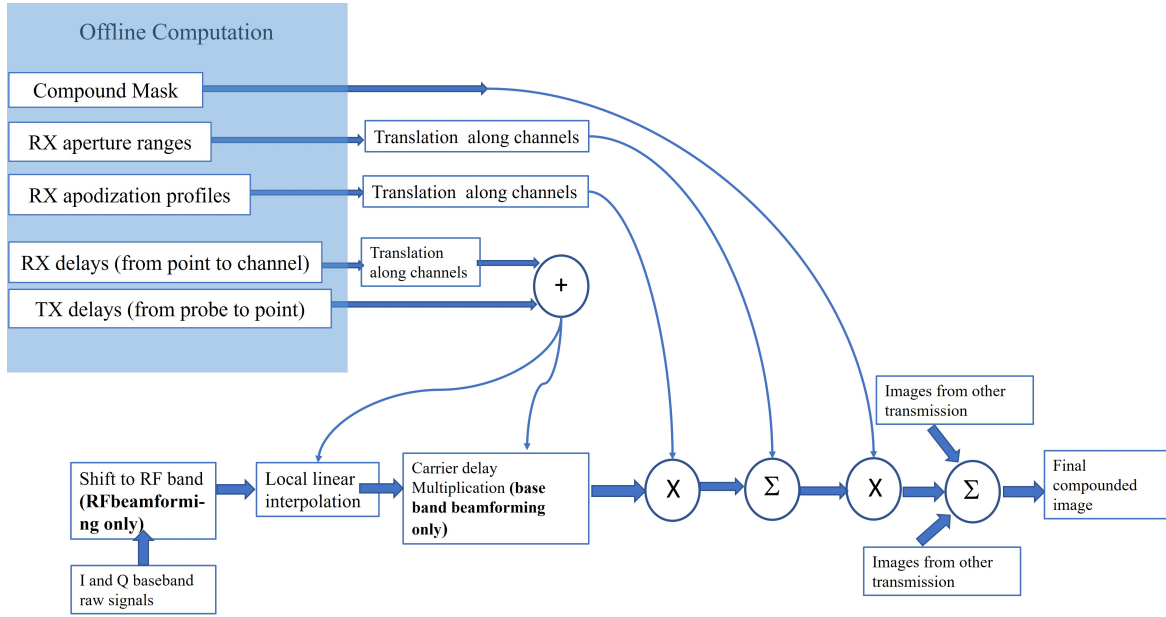


Fig. 6.1 Block diagram of the GPU-based receive beamformer for ultrafast diverging-wave imaging. The workflow illustrates the integration of offline pre-computations (compound mask, receive aperture ranges, apodization profiles, transmit and receive delays) with the online beamforming pipeline. Pre-computed data are efficiently translated along channels and combined to generate focalization terms. Depending on the beamforming mode, the input data are either shifted to the RF band (RF beamforming) or processed directly in baseband (I/Q signals) with carrier delay multiplication. Local linear interpolation ensures accurate delay alignment, followed by element-wise apodization weighting and summation across channels. Compounded images are obtained by combining multiple transmissions, ultimately producing the final high-quality B-mode image.

where :

- $I(x, z)$: final compounded image
- N : number of transmissions
- H_{\min}, H_{\max} : indexes of minimum and maximum channel in receive aperture
- $S_n(h, t)$: analytic (complex) temporal signal related to n -th TX received at channel h . Depending on the receive hardware, the signal can be in RF band or downmixed to baseband.
- $\mathcal{T}(x, z, h, n)$: two-way delays, function of imaging grid, channel index, and VS index
- $a(x, z, h, n)$: apodization profile, function of imaging grid, channel index, and VS index.

Notice that a pixel-based beamforming is adopted, i.e. the delay and sum is performed only on a predefined spatial grid that is independent from the sampling frequency. Delays and apodization can be computed online at each iteration, but their computation can be quite cumbersome if a high precision is required. For example, the computation of delays should take into account the refraction effects into the probe lens, while proper apodization should take into account the shape of the transmitted field that has to be simulated. For these reasons and considering that they are data-independent, it is more convenient to compute them offline just once and store them in memory. However, since both apodization and delays are four-dimensional matrices, their size could exceed the hardware capabilities or slow down their access, especially considering GPU-based architectures. A first memory reduction can be achieved by decoupling delays and apodization into a transmit part and a receive part, as shown in the following equation:

$$I(x, z) = \sum_{n=1}^N \sum_{h=H_{\min}}^{H_{\max}} S_n(h, \mathcal{T}_{tx}(x, z, n) + \mathcal{T}_{rx}(x, z, h)) \cdot a_{cp}(x, z, n) \cdot a_{rx}(x, z, h) \quad (6.2)$$

where $T_{tx}(x, z, n)$ is the travel time of TX wavefront from the probe to grid point (x, z) , $\mathcal{T}_{rx}(x, z, n)$ is the travel time of the echo backscattered from grid point (x, z) to probe channel h , $a_{rx}(x, z, h)$ is the RX apodization applied to each channel for beamforming the point (x, z) and $a_{cp}(x, z, n)$ is the compound mask that weights for each grid point the contribution of transmission n . A method for the computation of the compound mask was previously introduced in 3. Notice that the former 4D matrices boiled down to a couple of 3D matrices with a notable computational saving and only an additional sum (for delays) and multiplication (for apodization). Also, the compound mask is multiplied outside the inner sum over channels. In this implementation, the higher memory demand is expected to be due to the receive part since the channel index may span even more than 200 channels, while the VS index span usually spans a few tens of transmissions. However, the receive focalization data can be further reduced with a specific organization of the grid by leveraging translation symmetries. Let us assume that the ratio between the probe inter-element spacing and the distance between two adjacent x coordinates is an integer number K . Notice that this approach works with both cartesian grids with linear probe geometries and polar grids with convex probe geometries (in the second case, distances are substituted with angles). With this assumption it is easy to demonstrate that only a subset of K RX delays and apodization profiles have to be computed, corresponding to K different x_0 coordinates, while all the other profiles for the remaining x coordinates can be computed on the fly by simple translation over the channel indexes as

displayed in the following equation:

$$I(x, z) = \sum_{n=1}^N \sum_{h=H_{\min}}^{H_{\max}} S_n(h, \mathcal{T}_{tx}(x, z, n) + \mathcal{T}_{rx}(x_0, z, h - k)) \cdot a_{cp}(x, z, n) \cdot a_{rx}(x_0, z, h - k) \quad (6.3)$$

where x_0 belongs to the subset of x coordinates of size K and k in an integer index such that $x_0 = x - kP$ where, P is the probe pitch. With this approach, the first dimension of RX matrices is greatly reduced from 100 – 200 to 1 – 2, i.e. the typical number of beamformed lines per pitch.

Since delays \mathcal{T}_{tx} and \mathcal{T}_{rx} are real numbers, their application to sampled temporal data is performed by selecting a couple of samples closest to the current delay and applying local linear interpolation. For RF band signals, this is the whole procedure, while for base band signals also the phase of the carrier must also be properly delayed. This is done by multiplying the IQ data by a complex exponential function of the delay, the carrier frequency, and the sampling frequency. This additional operation involves the computation of a cosine and a sine, five real multiplications, and two real sums for each delayed sample. The overall proposed beamforming implementation, with the two variants of RFB and BBB, is summarized in the diagram in Figure 6.1.

6.2.2 GPU-Implementation using CUDA kernels

The beamforming algorithm is implemented using PyTorch on an NVIDIA GeForce RTX 4070 GPU with CUDA. The raw data is processed in parallel, with the GPU processing multiple pixels simultaneously along the z -axis of the imaging grid for each transmission. The beamformed signals are then transmitted to the PC for post-processing, which includes, in the case of B-mode imaging, envelope extraction, followed by log compression and spatial interpolation to construct high-resolution images at video. We employed a depth-first parallelization, where each GPU thread iteratively computes beamformed sample values along the z -axis (depth) for a given lateral position (x). This approach is in line with prior strategies that exploit independence along depth for GPU acceleration—for example, pixel-driven schemes where each pixel is computed in its own thread to allow massive concurrency. In addition, computation is distributed across virtual source (VS) transmissions: for each transmit event, the depth-stack per (x) is processed in parallel, and the compounding across multiple VS is done sequentially. Combining depth-wise threading with transmission-wise sequencing ensures that we can exploit both within-transmission as well as across-transmission parallelism, maximizing GPU occupancy while keeping memory usage in check. Specifically, the raw RF or baseband data for each transmit event is stored in GPU

global memory. For a given lateral location x , a block of threads (one thread per depth sample) reads the channel data, applies the appropriate delay (precomputed and passed into the kernel), samples via linear interpolation, and accumulates contributions from all receive channels. This constitutes one fully formed depth-line for that transmission. The value of our implementation lies in its practical simplicity and efficiency: using only depth-wise and transmit-wise parallelism, we achieve real-time beamforming by building on proven GPU methodologies. The integration of the optimized compound mask and VS distributions, described respectively in previous chapters within this minimal implementation, makes our solution both efficient and straightforward to reproduce on commercially viable hardware, while preserving image quality at reduced computational cost.

6.2.3 Datasets for Experimental Validation

The experimental validation of the proposed GPU-based receive beamformer was carried out using the same datasets previously acquired and described in our earlier optimization study 5. Specifically, we adopted the midrange virtual source (VS) distributions identified in the Pareto-optimal front of the multi-objective optimization framework, which offer a balanced compromise between lateral resolution and sidelobe suppression. These distributions were tested for different transmissions of 20, 10, 6, and 4 VSs, providing a consistent basis for assessing both image quality preservation and computational efficiency. In addition, we used a deterministic curvilinear distribution with 20, 10, 6, and 4 VSs as a baseline for comparison with the optimized configurations. The chosen virtual source distributions are shown in Figure 6.2. Two processing configurations were considered: RF-band beamforming and Baseband beamforming. The datasets were reconstructed using the proposed GPU beamformer, integrating the previously developed compound mask strategy and optimized VS distributions. This ensured that the evaluation captured the full impact of the optimization framework on both image quality (in terms of full width at half maximum and contrast ratio) and computational performance (processing time, memory occupancy).

1

¹The content of this chapter was partially published in IEEE Prime conference: **Z. Alzein**, M. Crocco, and D. D. Caviglia, "Computationally Efficient RF Band and Baseband Beam-former for Coherent Plane Wave Imaging," *2024 19th Conference on Ph.D Research in Microelectronics and Electronics (PRIME)*, Larnaca, Cyprus, 2024, pp. 1–4, <http://dx.doi.org/10.1109/PRIME61930.2024.10559667>.




Fig. 6.2 Virtual source distributions used for experimental validation of diverging-wave imaging. Five transmit configurations are compared: 20 virtual sources (VSs) arranged in a conventional curvilinear pattern, 20 VSs optimized, 10 optimized VSs, 6 optimized VSs, and 4 optimized VSs.

6.3 Results

6.3.1 B-mode Images for Multi-Purpose Phantom

The Beamformer was validated using in-silico data with plane wave transmission in [115]. In this chapter only experimental data is presented, figure 6.3 illustrates the reconstructed B-mode images obtained from the multi-purpose phantom using different virtual source (VS) distributions. The reference case with 20 VSs arranged in a curvilinear geometry [Fig. 6.2(a)] provides a baseline image quality with acceptable lateral resolution but suffers from reduced contrast, particularly in the depiction of anechoic structures. When the same number of transmissions is optimized [Fig. 6.2(b)], a substantial improvement in contrast and boundary definition is observed, indicating the benefit of optimized source placement even at higher transmission counts.

A progressive reduction in the number of VSs to 10, 6, and 4 demonstrates the expected trade-off between frame rate and image quality. While fewer transmissions yield higher frame rates and reduced memory requirements, the image quality is increasingly degraded in terms of contrast resolution. Nevertheless, when optimized distributions are employed, the degradation is less severe, as can be seen in Figure 6.3, particularly in preserving cyst detectability and maintaining lateral resolution across depth.

This qualitative observation is further supported by the quantitative analysis of image quality metrics (Figure 6.4). The lateral resolution, expressed as the full-width at half maximum (FWHM) along the central axis, shows a monotonic degradation with depth for all configurations. The optimized 20-VS case achieves the lowest FWHM values across depths (30, 40, and 60 mm), outperforming both the deterministic 20-VS curvilinear distribution and the reduced transmission schemes. Notably, the 10-VS optimized distribution maintains

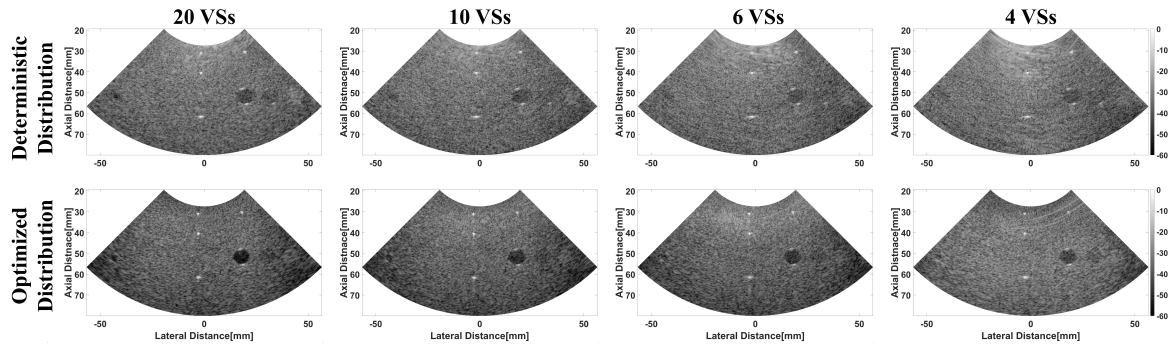


Fig. 6.3 Beamformed B-mode images for the multi-purpose tissue-mimicking phantom using the GPU-based receive beamformer under different virtual source (VS) distributions. First row (20,10,6 and 4)VSs arranged in a curvilinear deterministic distribution, second row the same number of VS with the optimized distributions presented in figure 6.2.

resolution values close to the 20-VS optimized case, showing that optimized configurations can effectively mitigate the loss of resolution when reducing the number of transmissions.

In terms of contrast resolution, quantified as the contrast ratio (CR) of three cysts located at 55 mm depth (Figure 6.4(b)), the optimized distributions consistently outperform the deterministic baseline. The 20-VS optimized configuration achieves the best CR (27 dB), significantly surpassing the curvilinear case (17 dB). As the number of transmissions decreases to 10, 6, and 4, the CR values of optimized distributions degrade (20.3 dB, 19.1 dB, and 17.9 dB, respectively), but remain superior to the deterministic 20-VS configuration. This highlights the effectiveness of optimization in preserving cyst contrast even under transmission reduction.

6.3.2 Computational Complexity

The computational complexity of the proposed beamforming strategies was evaluated in terms of processing time, achievable frame rate, and memory storage requirements for both raw data and focalization parameters. These metrics were analyzed as a function of the number of transmissions, considering 20, 10, 6, and 4 VS configurations.

Table 6.1 shows the processing time per frame for both RF-band and baseband implementations. As expected, the processing time increases approximately linearly with the number of transmissions. For the full 20-VS configuration, the processing time reaches 54 ms (RF-band) and 67 ms (baseband), while reducing the number of VSs to 10, 6, and 4 results in substantial gains, with times dropping below 30 ms and approaching 10 ms for the lowest transmission count. This reduction directly translates to improved frame rates.

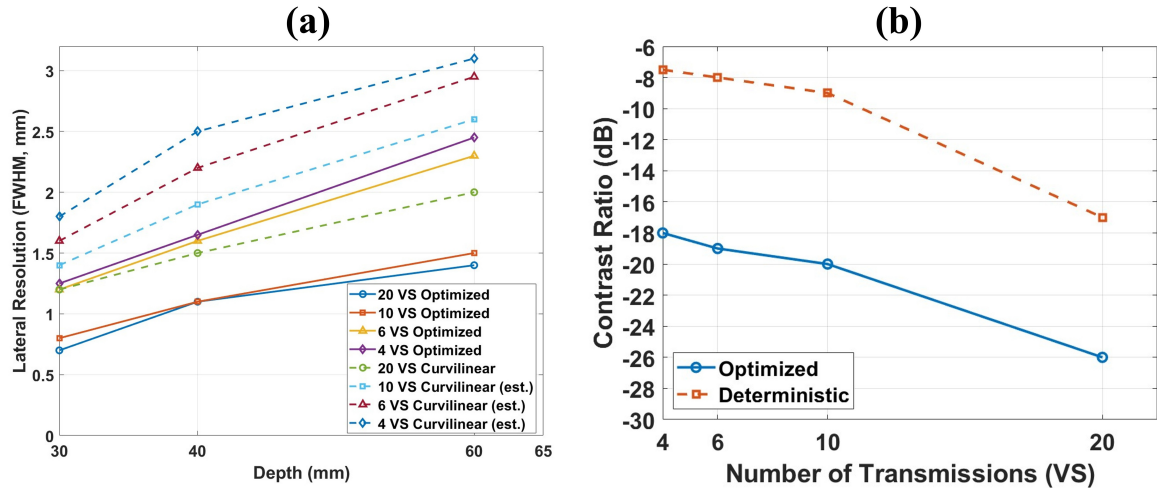


Fig. 6.4 Image quality assessment of GPU-based receive beamforming under varying numbers of virtual source (VS) transmissions for deterministic and optimized distributions. (a) Lateral resolution, expressed as full-width-at-half-maximum (FWHM), was measured along the central axis at depths of 30, 40, and 60 mm. (b) Contrast ratio (CR) measured for an anechoic cyst located at 55 mm depth, expressed in decibels(dB).

The memory requirements for one frame were assessed separately for raw data, i.e. data coming from the probe, and focalization data (Fig. 6.5(a)). The raw data storage per frame scales directly with the number of transmissions, ranging from 31 MB at 20 VSs to 6 MB at 4 VSs. It is worth noticing that it is not mandatory to store together raw data for all the different transmissions at the same time but in any case, the number of transmissions per second would impact the transfer bandwidth requirement from the hardware to the GPU. This scaling highlights the strong dependency of acquisition bandwidth on transmission count and underlines the potential of optimized low-VS schemes in reducing data transfer and storage demands without proportional degradation in image quality.

Focalization data, comprising received delays, apodization profiles, compounded masks, and transmitted delays, follow a similar trend [Fig. 6.5(b)]. While the receive delays and apodization remain constant across transmission numbers due to their dependence on the image grid rather than transmission count, the compounded mask and transmitted delays scale linearly with the number of VSs. At 20 transmissions, the total focalization storage requirement reaches 34 MB, which reduces by nearly a factor of five at 4 transmissions. Importantly, these reductions in storage also decrease the memory bandwidth pressure during GPU processing, thereby contributing to faster execution times.

A key observation is that the optimized VS distributions preserve computational efficiency by maintaining favorable image quality even at low transmission counts. In particular, the 10-VS optimized configuration achieves image quality metrics close to the 20-VS optimized case,

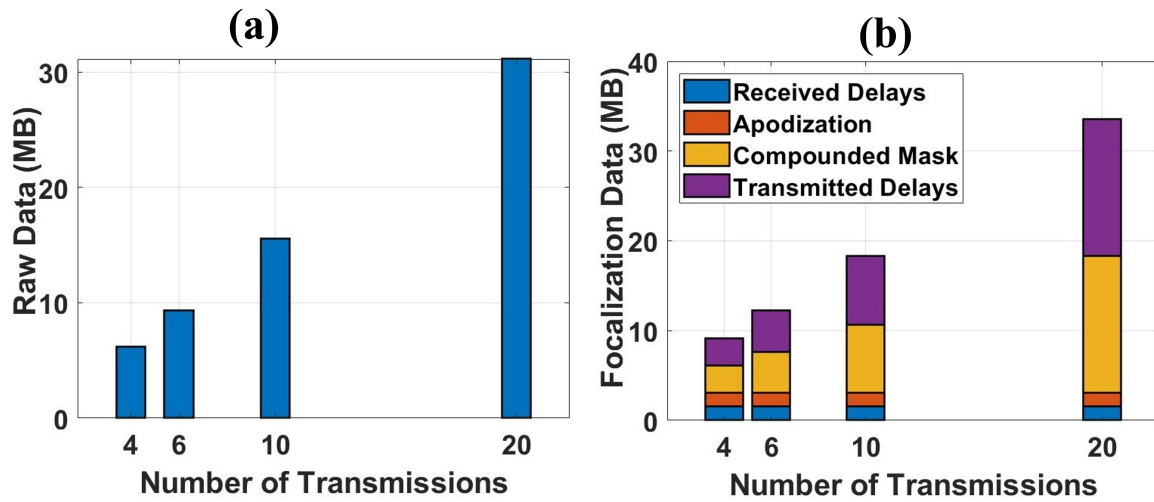


Fig. 6.5 Memory requirements for GPU-based receive beamforming as a function of the number of transmissions. (a) Raw data storage, which scales linearly with the number of transmissions, highlights the increasing data transfer demand from acquisition hardware to GPU memory. (b) Breakdown of focalization data memory usage, including received delays, apodization coefficients, compounded mask, and transmitted delays.

while halving the processing time and memory load. Similarly, reducing to 6 VSs further cuts resource requirements by two-thirds while still yielding diagnostically acceptable images, particularly in terms of cyst detectability. This demonstrates the ability of optimization to shift the trade-off frontier between complexity and quality, enabling significant savings in processing and memory without proportionally sacrificing imaging performance.

Table 6.1 Processing time for RF-band and Baseband beamforming with different numbers of transmissions.

Number of Transmissions	RF-band (ms)	Baseband (ms)
20	54	66
10	28	33
6	16	20
4	11	13

6.4 Conclusion

This work presents an optimization-driven GPU-based receive beamformer tailored for coherent diverging wave imaging under the constraints of commercial ultrasound systems. The proposed framework integrates multiple algorithmic optimizations, including precomputed

apodization masks, transmit-receive decoupling, and spatial symmetry exploitation, to minimize computational and memory overhead. Both RF-band and baseband beamformers were implemented and validated using experimental data acquired from a Verasonics scanner with a C5-2 probe.

Chapter 7

Extending the compound mask apodization for 3D Ultrafast Imaging using 2D-matrix probe

This chapter finalizes the thesis by extending the compound-mask transmit apodization framework from 2D to fully 3D DWI with a 2D matrix probe. We derive a voxel-adaptive mapping from STAI apodization to DWI virtual-source weights, and validate the method via FIELD II simulations on dense and sparse arrays. The developed methodology opens avenues for future work, including sparse/optimized apertures for 3D ultrafast volumetric imaging.

7.1 Introduction

Most diagnostic ultrasound examinations were performed in two dimensions (2D), which provides limited anatomical context and heavily relies on the operator's experience to mentally reconstruct the three-dimensional (3D) anatomy [165]. The importance of moving to 3D volumetric ultrasound is well recognized: a volume scan captures the entire region of interest, allowing the clinician to visualize arbitrary imaging planes and the true spatial relationships of structures. This reduces the uncertainty and operator dependence inherent in interpreting multiple 2D slices. Combining the benefits of volumetric imaging with the high frame rates of PWs/ DWs, holds promise for real-time 3D visualization of dynamic phenomena (cardiac function, blood flow, tissue elasticity, etc.) that were previously inaccessible.

Achieving real-time 3D ultrasound, however, poses significant technical challenges. Early approaches to 3D scanning involved mechanical movement of a 1D array (or freehand sweeps

with a position sensor) to acquire a volume, which was slow and produced static volumes [166]-[167]. Modern systems increasingly rely on 2D matrix array transducers that have transducer elements arranged in a matrix (e.g., 32×32 elements) to electronically scan a volume. Fully sampling a matrix array requires a very large number of channels (on the order of N^2 elements; for example, a 32×32 array has 1024 elements), which leads to enormous data throughput and hardware complexity.

Matching the lateral and elevational resolution of a conventional 1D array with a 2D matrix requires many elements in *both* directions. As a simple estimate, achieving the performance of a 128-element linear array would need about $128 \times 128 = 16,384$ independently driven elements. Because free beam steering needs element-level control, this implies the same number of channels. Most current systems expose roughly 256 channels, so a one-channel-per-element design would be equivalent to using ~ 64 systems in parallel, and the wiring/electronics for thousands of elements would make the probe and cable impractically large. A second constraint comes from element scaling: packing many elements into the same footprint makes each element small and therefore increases its electrical impedance [168, 169]. In standard linear arrays, impedances are typically 100Ω – $1 \text{ k}\Omega$ depending on center frequency and aperture [170]; in 2D matrices, the much smaller elements can be up to $\sim 64 \times$ higher, which reduces power transfer and receive sensitivity, lowering SNR [171].

Data and compute are also limiting. With thousands of channels, raw throughput and storage become very large. For example, a 40×40 (1,600-element) matrix sampled at 80 MHz with a 12-bit ADC produces on the order of 1,500 Gb/s, i.e., about 2 TB per second of acquisition [40]. The computational cost to reconstruct volumetric data from these streams is similarly heavy. To address these issues, prior work reduces either the number of independently driven elements or the volume of data sent to the system. Common strategies include *micro-beamforming* [172]-[173], *multiplexing* [174]-[175], *row-column addressing* [176], and *2D sparse arrays* [177]-[178]. All these techniques have specific tradeoffs between frame rate, contrast, signal-to-noise ratio (SNR), and system complexity. In general, 3D ultrasound tends to suffer from lower SNR and resolution than 2D, because the energy is distributed over a volume and the aperture in each dimension is limited. Improving image quality for 3D ultrafast ultrasound, while managing the data and hardware constraints, is an active area of research.

One key technique for improving image quality in beamformed ultrasound is transmit apodization – the application of amplitude weighting across the transmit aperture to shape the emitted beam. In conventional focused array imaging, applying a tapering window (e.g., Hanning or Hamming) to the transmit elements can suppress sidelobes and thus improve contrast resolution, at the cost of a slightly widened main lobe [102][103]. Transmit

apodization effectively weights down the contributions of outer elements, reducing the intensity of sidelobe beams. However, in unfocused transmission schemes like PWs or DWs, all elements typically transmit together without individual focusing delays, so one cannot directly apply different voltages per element based on each transmission to achieve apodization in the same way. Instead, it is possible to apply weights to the image grid related to each transmission before coherent compounding. To date, most 3D ultrafast systems either use no transmit apodization (i.e., all elements firing uniformly per plane-wave or divergent-wave) or rely on simple fixed tapers on the aperture, due to the difficulty of designing location-dependent weights.

For DWs in 2D (using a linear or curved array), the problem of transmit apodization has only recently been addressed in our previous work (chapter 3). We proposed a closed-form method to map any chosen STAI transmit apodization window onto a set of virtual source weights for coherent diverging-wave compounding. This compound mask mapping was derived by evaluating the local angle of impact of the diverging wavefront at every image point, and retrieving the apodization weight of STAI associated with the probe element seeing the image point under the same angle. The method, inspired by the plane-wave angular apodization approach [17], provides a deterministic way to calculate compounding weights for each diverging wave emission without trial-and-error weighting. It was validated in simulations in chapter 3 and in vitro data in chapter 4 for 2D imaging with both linear and convex arrays, demonstrating that the mapped weights preserve the point-spread function (PSF) characteristics of the reference synthetic aperture method (with only a few percent difference in lateral resolution). A key advantage is that the weights can be computed offline in advance and applied during imaging with negligible computational overhead. In practice, this means that the apodization mapping can be integrated into real-time imaging pipelines (e.g., on the Verasonics research scanner, ULAOP[83]) without impacting frame rate, since only a lookup of pre-computed weight factors is needed during beamforming. Once the compound mask was applied, significant improvements in image quality metrics (higher contrast ratio and improved lateral resolution) were observed compared to unweighted diverging wave compounding, validating the effectiveness of compound weight based on equivalent transmit apodization in ultrafast imaging.

Volumetric (3D) ultrafast imaging using DWs, however, still lacks an established transmit apodization strategy. To our knowledge, there have been almost no studies addressing how to design optimal weighting for multiple diverging wave transmissions in 3D imaging. In 3D DWI, one must determine appropriate weights for each virtual source position (or each transmit event) in a diverging wave set that spans the volume. This is more challenging than the 2D case since the aperture of a matrix array is inherently smaller. Simply applying a fixed

2D window (e.g., a radial Tukey or Hanning) to the virtual sources is unlikely to be optimal for all points in the 3D volume, and an unweighted average of diverging wave images may suffer from high sidelobes.

In this chapter, we address the problem of transmit apodization for 3D ultrafast ultrasound with diverging waves. Building upon our previous 2D solution, we derive a method to compute weighting coefficients for each diverging wave emission in a volumetric scan by mapping a desired synthetic-aperture apodization function onto the set of virtual source positions. This effectively shapes the transmit beam in 3D space in a similar way that conventional apodization shapes a focused beam. The proposed apodization mapping were compared with a conventional approach of using a static transmit window (uniform weights or a basic taper applied equally to all transmits). Through simulations using Field II software[116][117], we demonstrate that the mapped apodization yields notable improvements in image quality, including lower sidelobe levels and enhanced contrast in 3D diverging wave compounding (DWI).

7.2 Methodology

This section details the methodology adopted for extending the compound mask framework to 3D DWI. It presents the mathematical formulation linking Synthetic Aperture and Diverging Wave imaging through apodization mapping, the derivation of compound mask weights via a geometrical approach, the implementation of virtual aperture and static apodization strategies. The final part describes the simulation setup used to validate the proposed approach under various aperture sizes and apodization schemes.

7.2.1 Compound Mask Computation for 3D Divergent Wave Imaging

The coherent reconstruction of volumetric ultrasound images using Synthetic Transmit Aperture Imaging (STAI) and Diverging Wave Imaging (DWI) can be unified within a common geometrical framework that relates the transmit apodization weights of STAI to the compound mask weights required for DWI. Let's consider an imaging voxel located at $\mathbf{r} = (x, y, z)$, a planar probe with N transmit elements and M receive elements. In STAI imaging, the reconstructed signal at voxel \mathbf{r} is expressed as a coherent summation over transmit and receive elements:

$$S_{\text{STAI}}(x, y, z) = \sum_{r=1}^M u_r(x, y, z) \sum_{t=1}^N v_t(x, y, z) h_{tr} \left(\frac{D_t + D_r}{c_0} \right), \quad (7.1)$$

where $u_r(x, y, z)$ and $v_t(x, y, z)$ denote the receive and transmit apodization functions, respectively, and $h_{tr}(\cdot)$ represents the transmit–receive impulse response. The delay distances are defined as

$$D_t = \sqrt{(x - x_m)^2 + (y - y_m)^2 + z^2}, \quad (7.2)$$

$$D_r = \sqrt{(x - x_m)^2 + (y - y_m)^2 + z^2}, \quad (7.3)$$

where (x_m, y_m) denote the coordinates of the transmit and receive elements on the probe surface ($z = 0$). In DWI, the transmitted field is generated by a set of V virtual sources




Fig. 7.1 Geometry of 3D DWI showing virtual sources (blue dots), probe elements (black grid), and image voxel (red point) with relevant distances D_i , D_{it} , D_t and D_r annotated.

$\mathbf{S}_i = (S_x, S_y, S_z)$, placed behind the probe surface as shown in Fig. 7.1, producing a diverging wavefront that insonifies the imaging region. The reconstructed signal at \mathbf{r} can thus be written as

$$S_{DW}(x, y, z) = \sum_{r=1}^M u_r(x, y, z) \sum_{i=1}^V w_i(x, y, z) \hat{h}_{ir} \left(\frac{D_i + D_r}{c_0} \right), \quad (7.4)$$

where $w_i(x, y, z)$ are the compound mask weights corresponding to each virtual source, and $\hat{h}_{ir}(\cdot)$ is the diverging-wave impulse response expressed as

$$\hat{h}_{ir}(t) = \sum_{t=1}^N h_{tr} \left(t - \frac{D_{it}}{c_0} \right). \quad (7.5)$$

Here, the distances are

$$D_i = \sqrt{(x - S_x)^2 + (y - S_y)^2 + (z - S_z)^2}, \quad (7.6)$$

$$D_{it} = \sqrt{(x_m - S_x)^2 + (y_m - S_y)^2 + S_z^2}, \quad (7.7)$$

and D_r is as defined in (7.3), all represented in Figure 7.1 . As demonstrated in chapter 3, STAI transmit apodization and DWI compound mask can be linked by the linear system:

$$\sum_{i=1}^V w_i(x, y, z) H(i, t) = v_t(x, y, z), \quad (7.8)$$

with

$$H(i, t) = \frac{E(\tau_{it})}{E(0)}, \quad \tau_{it} = \frac{D_i - D_t - D_{it}}{c_0}, \quad (7.9)$$

where $E(\cdot)$ denotes the convolution of the transmit pulse with the two-way electroacoustic transfer function of the probe elements. Equation (7.9) defines the relative delay τ_{it} entirely from 3D path lengths D_i , D_t , and D_{it} . Under the short-pulse approximation, the coupling matrix \mathbf{H} becomes approximately diagonal, as shown in Chapter 3 for the 2D case. Therefore, the only significant contribution in (7.8) is given by the pair (i, t) satisfying $\tau_{it} = 0$, which yields

$$w_i(x, y, z) \approx v_t(x, y, z) \quad \text{for the pair } (i, t) \text{ such that } \tau_{it} = 0. \quad (7.10)$$

Geometrically, $\tau_{it} = 0$ holds when the ray from \mathbf{S}_i to \mathbf{r} aligns with the transmit element t . This alignment is enforced by computing the intersection of the line segment $\mathbf{S}_i \rightarrow \mathbf{r}$ with the probe plane $z = 0$ and assigning to $w_i(x, y, z)$ the SA apodization sampled at the corresponding probe location (with interpolation over the discrete element grid). The following subsections will detail the definition of the transmit apodization function in 3D SA imaging and the derivation of the compound mask weights via intersection mapping and interpolation.

7.2.1.1 Transmit Apodization in Synthetic Aperture Imaging

The transmit apodization weights $v_t(x, y, z)$ in SA are typically derived from a baseline apodization window (Hanning Window), dynamically adjusted for voxel depth and lateral offsets. For a static Hanning window defined over normalized aperture coordinates $u \in [-0.5, 0.5]$, the per-element transmit weight is expressed as

$$v_s(u) = \begin{cases} \frac{1}{2} [1 - \cos(2\pi(u + \frac{1}{2}))], & |u| \leq \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases} \quad (7.11)$$

For the m -th transmit element located at (x_m, y_m) and voxel $\mathbf{r} = (x, y, z)$, the normalized lateral offset is

$$u_m(x, y, z) = F_{\#} \frac{\sqrt{(x - x_m)^2 + (y - y_m)^2}}{z}, \quad (7.12)$$

where $F_{\#}$ is the transmit F-number. The dynamic apodization weight for element m is then $a_{tx,m}(x, y, z) = v_s(u_m(x, y, z))$ subsequently used to form the apodization map $a_{tx}(\cdot, x, y, z)$ over all transmit elements, which is then interpolated in the compound mask computation.

7.2.1.2 Derivation of Compound Mask Weights via Intersection Mapping and Interpolation

To map the SA transmit apodization weights to the compound mask weights in 3D DWI, we leverage the geometric intersection of the line connecting the virtual source \mathbf{S}_i and voxel \mathbf{r} with the probe surface (assumed planar at $z = 0$). The parametric equations of this line are:

$$x(t) = S_x + t(x - S_x), \quad (7.13)$$

$$y(t) = S_y + t(y - S_y), \quad (7.14)$$

$$z(t) = S_z + t(z - S_z). \quad (7.15)$$

Solving $z(t_0) = 0$ gives

$$t_0 = -\frac{S_z}{z - S_z}, \quad (7.16)$$

and thus the intersection coordinates on the probe plane:

$$x_e = S_x + t_0(x - S_x) = S_x - \frac{S_z}{z - S_z}(x - S_x), \quad (7.17)$$

$$y_e = S_y + t_0(y - S_y) = S_y - \frac{S_z}{z - S_z}(y - S_y). \quad (7.18)$$

The compound mask weight $w_i(x, y, z)$ for virtual source i and voxel \mathbf{r} is computed by interpolating the per-element SA transmit apodization weights $a_{tx,m}(x, y, z)$ from the discrete element grid $\{(x_m, y_m)\}$ to the off-grid intersection point (x_e, y_e) :

$$w_i(x, y, z) = \mathcal{I}_2 \{a_{tx}(\cdot, x, y, z); \{(x_m, y_m)\}\} (x_e, y_e), \quad (7.19)$$




Fig. 7.2 Illustration of the virtual aperture concept in 3D Diverging Wave Imaging. **Left:** rays from virtual sources (blue) intersect the physical aperture (black) for voxels within the central field of view, defining valid compound mask weights. However, voxels near the periphery may yield intersection points outside the probe footprint. **Right:** introduction of a virtual aperture (red) surrounding the physical aperture ensures that all voxel–source rays intersect a valid region, allowing weights to be computed and stored.

where \mathcal{I}_2 denotes a 2D interpolation operator such as cubic spline or bilinear interpolation. Explicitly, this can be written as a kernel-weighted sum:

$$w_i(x, y, z) = \sum_{m=1}^{N_e} a_{tx,m}(x, y, z) \Phi \left(\frac{x_e - x_m}{\Delta x}, \frac{y_e - y_m}{\Delta y} \right), \quad (7.20)$$

where Φ is the interpolation kernel, N_e is the number of elements locally involved in the interpolation, and Δx , Δy are the uniform pitch spacings between elements in the lateral and elevational directions, respectively.

7.2.2 Computation of Weights based on virtual Aperture

Since the typical 2D matrix probes have usually a limited footprint, especially for cardiovascular applications, some voxel–source rays intersect the probe plane outside the physical aperture. In such cases, directly interpolating the SA apodization map $a_{tx}(\cdot, x, y, z)$ on the discrete element grid $\{(x_m, y_m)\}$ is not possible and would leave the compound weights undefined. To guarantee well-defined transmit weights over the whole 3D field of view

the following procedure can be adopted. First of all we define an extended probe element grid (\hat{x}_m, \hat{y}_m) centered on the probe surface center, in which the grid steps $(\Delta x, \Delta y)$ on the two spatial direction correspond to the physical probe elements spacing, and covering a rectangular region:

$$\mathcal{V} = [x_{\min}^V, x_{\max}^V] \times [y_{\min}^V, y_{\max}^V]. \quad (7.21)$$

The extents of \mathcal{V} are chosen once to *guarantee* that the intersection of the lines connecting every virtual source and every image point for the imaging box $\mathcal{G} = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}] \times [z_{\min}, z_{\max}]$ onto the plane $z = 0$ fall within the boundaries of the region \mathcal{V} itself. In practice we created an extended probe aperture in which the inner elements correspond to the physical probe elements and the outer ones are virtual elements. As illustrated in Fig. 7.2, the schematic shows the physical aperture and the virtual aperture surrounding it, together with the image-grid regions that remain covered by the proposed virtual aperture computation. Then we proceed to compute an extended SA transmit apodization map $\hat{a}_{tx}(\cdot, x, y, z)$ following the procedure described in subsection 7.2.1.1 but substituting the grid (x_m, y_m) with the extended grid (\hat{x}_m, \hat{y}_m) . Finally, we compute the compound mask for the DWI method for each virtual source as an interpolation of the values of the extended STAI apodization map, following the procedure in Eqs 7.19 and 7.20:

$$w_i(x, y, z) = \mathcal{I}_2\{\hat{a}_{tx}(\cdot, x, y, z); \{(\hat{x}_m, \hat{y}_m)\}\}(x_e, y_e), \quad (7.22)$$

This approach guarantees the existence of valid STAI apodization values in the nearby of every interpolation point.

All compound weights $w_i(x, y, z)$ for all voxels and virtual sources are computed offline and saved in memory; during delay-and-sum beamforming the only additional operation is the multiplication of per-virtual-source beamformed data by these precomputed weights before coherent compounding (chapter 4).

7.2.3 Static Hanning window apodization for each virtual source for Comparison

In addition to the compound mask formulation, we consider a baseline apodization strategy that applies a *static two-dimensional (2D) Hanning window* directly on the virtual source (VS) constellation, where the 2-D Hanning apodization is computed as the outer product of two 1-D Hanning windows. The goal is to reproduce the well-established apodization technique widely used in array signal processing and provide a consistent reference against the proposed mapping. Notice that with this baseline apodization, the compound mask is

function only of the 2D position of virtual sources on x and y coordinates and not dependent on the voxel position.

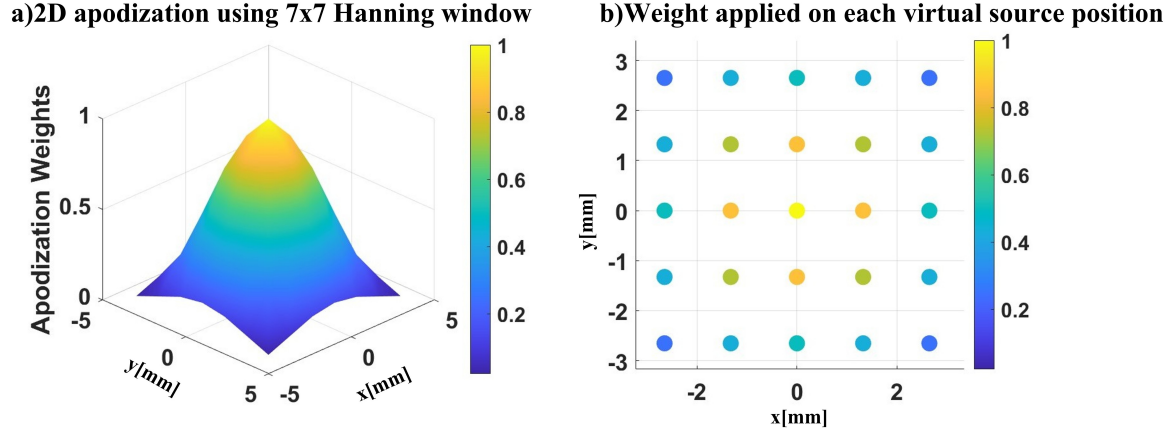


Fig. 7.3 Static 2D Hanning apodization profile applied to virtual sources. **(a)** 2D Hanning surface defined on an extended 7×7 grid, ensuring that apodization weights taper smoothly without vanishing at the edges. **(b)** Sampling of the weights at the positions of the actual 5×5 virtual source constellation.

Since the values of the Hanning window tend to zero at the boundaries of the window, sources located at the periphery receive near-zero weights, effectively removing them from contributing to the compounded image. To address this issue, we define the 2D Hanning window over an extended area covering a 7×7 rectangular grid. Then we pick up the values over the inner subgrid of size 5×5 and assign them as weights to the corresponding virtual sources. The latter are therefore arranged on a rectangular grid of size 5×5 on a (x, y) plane at a fixed negative value of z . In this way, the outermost virtual sources of the effective grid no longer fall into the near-zero region of the Hanning function, thereby retaining non-zero weights while maintaining the desired tapering profile. An illustration of the continuous 2D Hanning profile and its sampled weights for the 5×5 VS constellation is reported in Fig. 7.3.

7.2.4 Simulation Setup

The simulations were performed using MATLAB FIELD II toolbox [116][117], considering a fully populated 2D matrix probe comprising 32×32 elements, for a total of 1024 transmit/receive channels. Each element exhibits a fractional bandwidth of approximately 40% in pulse-echo mode. The interelement pitch was set to $\lambda/2 \approx 257 \mu\text{m}$, where $\lambda = c_p/f_0$ is the wavelength corresponding to a propagation speed of $c_p = 1540 \text{ m/s}$ and a center frequency of $f_0 = 3 \text{ MHz}$. A kerf of $\lambda/20$ was included between adjacent elements. With these parameters, the effective footprint of the probe results in an aperture size of $8.4 \times 8.4 \text{ mm}^2$. To generate

diverging wave transmissions, a total of 25 virtual sources (VSs) were employed arranged on a regular grid and as concatenated rings and positioned at a distance of 10 mm along the negative z -axis with respect to the probe surface. The proposed methodology was validated using simulated point-like scatterers positioned along the central axis of the imaging volume. Four scatterers were considered at axial depths of $z = 5, 20, 40,$ and 65 mm, with lateral coordinates fixed at $x = 0$ and $y = 0$. Three different apodization strategies were compared to assess their impact on image quality:

- **No apodization:** virtual sources were assigned uniform weights across the grid.
- **Static apodization:** a two-dimensional Hanning window, defined on an extended 7×7 lattice and sampled at the 5×5 virtual source positions, was applied as described in Section 7.2.3.
- **Compound mask apodization:** voxel weights were computed based on the proposed mapping framework described in section Section 7.2.1 and Section 7.2.2

To assess the dependence of the compound mask method on channel number, we repeated all experiments with a 32×32 matrix probe with a *25% random sparsity factor*, i.e., a uniformly random subset of 256 active elements out of 1024. The virtual-source layout (25 VSs at $z_s = -10$ mm) and all other parameters were kept identical across configurations, and the three apodization strategies (*no*, *static 2-D Hanning*, and *compound mask*) were evaluated in each setting.¹

7.3 Results

7.3.1 3D Contouring of the Compound Mask Results for a fixed image point

To demonstrate the validity of the compound mask formulation in 3D DWI, we first analyze the transmit apodization profiles obtained for selected voxels in the imaging grid. Figure 7.4 reports a comparison between the reference STAI apodization and the corresponding mapped profiles generated by the compound mask in DWI. Two representative voxels were considered: one located on the central axis at $(0, 0, 20)$ mm and another off-axis point at $(-1, 2.3, 21.4)$ mm. For both cases, the STAI apodization is shown as a smooth and radially

¹The content of this chapter will be submitted to IEEE Open Journal of Ultrasonics, Ferroelectrics, and Frequency Control: **Z. Alzein**, M. Crocco, D. D. Caviglia, and H. Liebgott, ‘‘Transmit Apodization for 3D Diverging Wave Ultrafast Imaging’’

symmetric function on the probe aperture. The corresponding compound mask derived for DWI reproduces the same contour distribution extended to virtual source positions instead of probe elements. Importantly, the numerical values of the transmit weights are locally preserved within the DWI formulation, as evidenced by the near-identical range of amplitudes (0.88–0.98). In both panels, the contours of the STAI map is obtained by sampling the dynamic transmit apodization $a_{tx,m}(x,y,z)$ on the full 32×32 elements, whereas the DWI map is computed from the compound-mask weights $\{w_i(x,y,z)\}$ defined on a much sparser 5×5 virtual-source grid. This difference in sampling density explains why the DWI contours appear more polygonal/segmented than the smoother STAI contours.

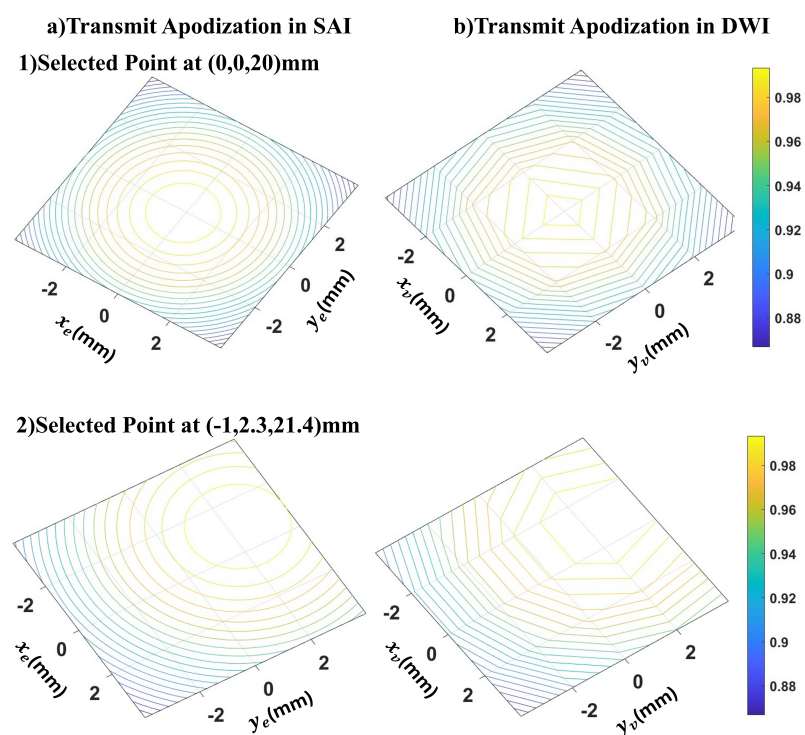


Fig. 7.4 Comparison between transmit apodization in STAI and the mapped compound mask in DWI for two selected points in the 3D imaging volume. **(a)** Apodization profiles obtained from STAI at points $(0, 0, 20)$ mm and $(-1, 2.3, 21.4)$ mm are shown as continuous contour maps ($\{(x_e, y_e)\}$ represents the probe elements). **(b)** Corresponding compound mask profiles in DWI, derived using the proposed compound mask mapping, demonstrate strong agreement with the STAI reference. The local structure of the apodization, including its smooth tapering and symmetry, is preserved in the mapped DWI profiles.

7.3.2 PSF Simulation using 32x32 2D matrix probe

We first analyze the point-spread function (PSF) obtained with a fully populated 32×32 2-D matrix probe for four on-axis point scatterers positioned at depths of 5, 20, 40, and 65 mm, and considering two virtual-source (VS) layouts: (i) *Grid distribution* (5×5 uniformly spanning the aperture) and (ii) *Concatenated rings + central VS*, consisting of **12 sources on an outer ring** (10 mm diameter), **12 sources on an inner ring** (5 mm diameter), and **one central source**—for a total of **25 VSs**; all VSs lie at $z_{VS} = -10$ mm and all VSs lie at $z_{VS} = -10$ mm. Figure 7.5 shows xz slices ($y = 0$) for the two VS layouts (rows) and the three weighting strategies (columns). With both layouts, the *compound-mask* reduces the sidelobe levels relative to the *no-weights* case. At larger depths (40–65 mm), qualitative differences among weightings become less evident in the images. Figure 7.6 reports lateral PSF cuts at x for the four depths (color-coded). For both VS layouts, the compound mask consistently lowers the first and subsequent sidelobes compared with no weights, with the largest difference at shallow depths (5–20 mm). The *concatenated rings* layout further smooths the lateral response and slightly reduces residual asymmetries compared with the grid distribution. Across depths, the impact of *static apodization* was negligible when compared with *no weights*: in the peak-normalized lateral profiles of Fig. 7.6 the two curves are almost coincident, and the small differences become visible only under strong zoom. Consistently, these tiny profile deviations do not translate into clear visual changes in the PSF images of Fig. 7.5. In contrast, the *compound-mask apodization* yields a systematic reduction of sidelobe levels, most prominently with the **Grid distribution (layout A)** and still noticeable with the **Concatenated rings** layout. Using the standard definition of peak-sidelobe level (PSL)—i.e., the difference between the mainlobe peak (0 dB after normalization) and the highest lateral sidelobe—measurements on the profiles of Fig. 7.6 indicate first-sidelobe reductions of approximately ~ 8 –10 dB at 5 mm, ~ 5 –7 dB at 20 mm, and ~ 3 –5 dB at 40–65 mm for layout A, with slightly smaller but consistent improvements (~ 3 –6 dB at shallow and ~ 1 –3 dB at deep targets) for the rings layout, which aligns with the image-level behavior in Fig. 7.5, where differences are most evident in the near field and progressively diminish at larger depth.

7.3.3 PSF Simulation using 32x32 2D matrix probe with 25% random sparsity factor

In order to consider the limited number of channels of most commercial systems (e.g., 256 Tx/Rx channels) and to reduce hardware and data-transfer burden, we repeated the experiment using the same 32×32 matrix probe but randomly enabling only 25% of the elements (256




Fig. 7.5 PSF images (xz slice at $y = 0$; full 32×32 array). Rows indicate VS layout; columns indicate VS weighting strategy. **Row A:** Grid distribution (5×5) located at $z_{VS} = -10$ mm. **Row B:** Concatenated rings + central VS with $r_{outer} = 5$ mm, $r_{inner} = 2.5$ mm, $z_{VS} = -10$ mm. **Columns:** (1) no weights, (2) static apodization, (3) compound-mask apodization. Images are envelope-detected and log-compressed to a 60 dB display range.

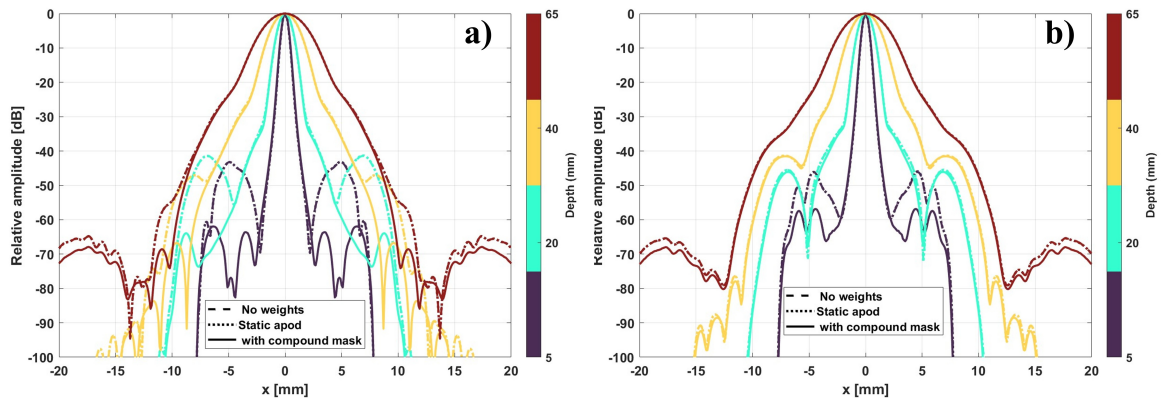


Fig. 7.6 Lateral PSF profiles (full 32×32 array). (a) 5×5 uniform grid over the full array aperture. (b) Concatenated rings + central VS: all VS located at $z_{VS} = -10$ mm behind the probe. Four on-axis point scatterers at 5, 20, 40, and 65 mm (colors match the depth colorbar). Curves are peak-normalized to 0 dB, Line styles denote the VS weighting: **dashed** = no weights, **dotted** = static apodization, **solid** = compound-mask apodization.

active out of 1024). As in the full-array study, we considered two virtual-source (VS) layouts: (A) *Uniform grid* (5×5 across the aperture) and (B) *Concatenated rings with a central VS*. Because the static apodization produced negligible improvement in the fully populated case (Sec. 7.3.2), we omit it here and focus on *no weights* vs. *compound-mask* apodization. Figure 7.7 shows xz slices of PSFs for layouts A and B with and without compound-mask apodization. As expected, random sampling increases the sidelobe levels and introduces more pronounced off-axis energy relative to the full array. Nevertheless, the compound mask visibly suppresses sidelobes for both layouts. At deeper targets (40–65 mm), qualitative differences remain but they are limited in the displayed FOV, consistently with off-axis energy moving outside the lateral span. Figure 7.8 quantifies these effects with peak-normalized lateral cuts at each depth. Across both layouts, compound-mask apodization lowers the first and subsequent sidelobes in comparison to the no-weights case. Using the standard peak-sidelobe level (PSL) metric on the curves in Fig. 7.8, we observe typical first-sidelobe reductions of approximately $\sim 6\text{--}8$ dB at 5 mm and $\sim 4\text{--}6$ dB at 20 mm, tapering to $\sim 2\text{--}4$ dB at 40–65 mm. Layout B (rings + center) tends to produce a slightly smoother, more isotropic lateral response, while layout A (grid) can yield a narrower mainlobe in the near field at the expense of stronger sidelobes without masking. Reducing the active aperture to 256 channels degrades PSF quality, due to the rise of the side lobes intensity, but *compound-mask apodization* consistently mitigates these losses, improving PSL, thereby improving visual contrast in the near field for both VS layouts.

7.4 Discussion

7.4.1 Analysis of Results

This work extends the compound-mask apodization introduced in chapter 3 for 2D DWI to volumetric (3D) DWI with a matrix probe. We evaluated two virtual-source (VS) layouts—(A) *Uniform grid* and (B) *Concentric rings with a central VS*—on a fully populated 32×32 array and on a sparse configuration with only 25% active elements (≈ 256 channels).

Static transmit apodization had *negligible effect* in the matrix-probe setting: in the peak-normalized lateral profiles (Fig. 7.6) the curves with and without static apodization are almost coincident; the small differences become visible only under strong zoom and do not yield perceptible changes in the corresponding PSF images (Fig. 7.5). In contrast, the **compound-mask** systematically reduced sidelobe levels (PSL), with the largest gains in the near field where sidelobes lie within the field of view (FOV). Using the standard PSL metric (difference between the mainlobe peak and the highest lateral sidelobe on the normalized cuts), we

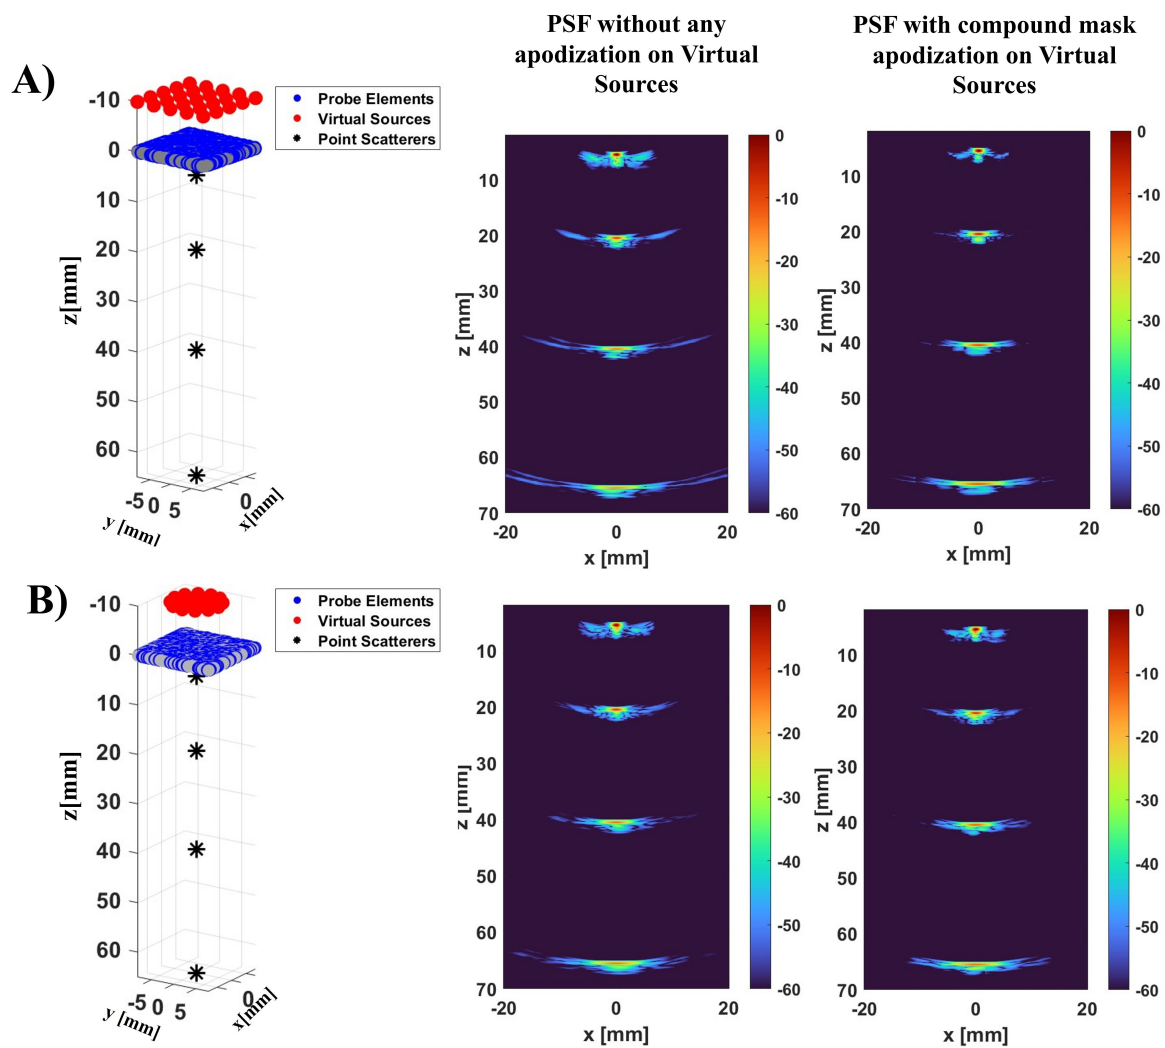


Fig. 7.7 **PSF images** (xz slice at $y = 0$; 32×32 with 25% random active elements). Rows indicate VS layout; columns indicate VS weighting. **Row A:** Uniform grid (5×5) located at $z_{VS} = -10$ mm. **Row B:** Concatenated rings + central VS with $r_{\text{outer}} = 5$ mm, $r_{\text{inner}} = 2.5$ mm, $z_{VS} = -10$ mm. **Columns:** (1) no weights, (2) static apodization, (3) compound-mask apodization. Images are envelope-detected and log-compressed to a 60 dB display range.

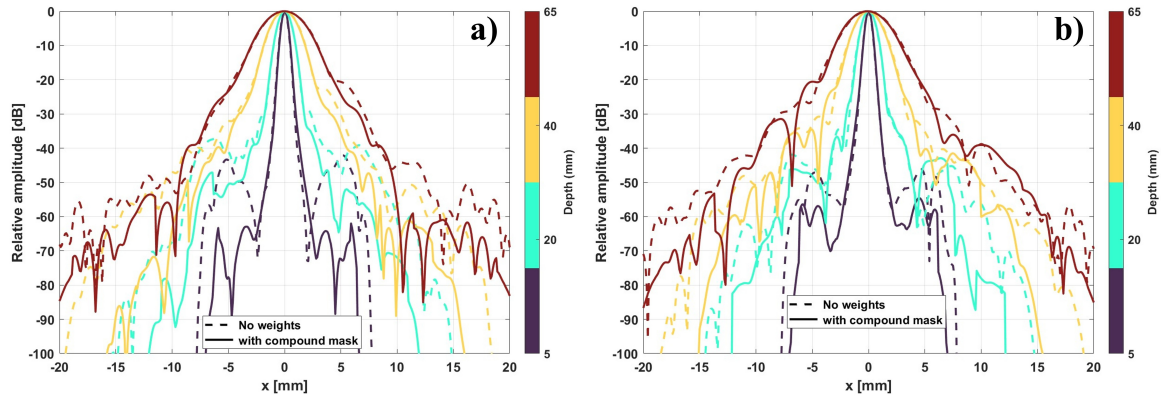


Fig. 7.8 Lateral PSF profiles (32×32 with 25% random active elements). (a) *Grid distribution* of virtual sources (VS), 5×5 uniformly over the full aperture. (b) *Concatenated rings + central VS*; all VS located at $z_{VS} = -10$ mm behind the probe. Four on-axis point scatterers at 5, 20, 40, and 65 mm (colors match the depth colorbar). Curves are *peak-normalized* to 0 dB, Line styles denote the VS weighting: **dashed** = no weights, **dotted** = static apodization, **solid** = compound-mask apodization.

observed typical first–sidelobe reductions of ~ 8 – 10 dB at 5 mm, ~ 5 – 7 dB at 20 mm, and ~ 3 – 5 dB at 40–65 mm for layout A; the rings+center layout B yielded slightly smaller yet consistent improvements. These trends match the image-level behavior: differences are most evident at shallow depths and gradually diminish as off-axis energy moves outside the displayed lateral extent at larger depths.

Subsampling the aperture raised the overall sidelobe levels and made off-axis energy more apparent, as expected. Even when reducing the number of channels, the compound mask mitigated the loss of image quality: lateral PSL gains of about ~ 6 – 8 dB at 5 mm and ~ 4 – 6 dB at 20 mm were typical, tapering to ~ 2 – 4 dB at 40–65 mm. The rings+center layout tended to produce smoother, more isotropic response, whereas the grid layout preserved a slightly narrower mainlobe in the near field. Taken together, the results indicate that precomputed compound masks are effective in 3D for sidelobes and preserving contrast, and that this persists under severe channel count reduction.

The compound mask is *data independent* and can be precomputed from a virtual (enlarged) aperture’s static apodization and mapped onto the physical VS layout. This keeps the online cost identical to standard DAS (no covariance inversions or iterative updates) and is suitable for real-time use and hardware deployment, as we demonstrated in chapter 4 and 6. In 3D, the dominant cost remains the DAS backprojection over all VS events and receivers; the mask adds only pointwise multiplications. Although this study is done only on simulation, simulation-based results often correlate with in vivo findings for similar imaging methods[179][180]. The compound mask’s data independence and reliance on geometric

optics (rays/angles) rather than on high-order adaptive statistics further support its robustness when moving from ideal media to experiments.

7.4.2 Limitations and Future Directions

In the PSF analysis we restricted evaluation to a single imaging plane ($y=0$) and to on-axis point scatterers centered with respect to the probe; where future studies should extend the analysis to off-axis targets (lateral positions) and to multiple planes (e.g., x - y slices at different depths and oblique cuts) to fully characterize full 3D volume. Also, in this study, we assumed a homogeneous speed of sound and neglected aberration and electronic noise. While the compound mask is relatively robust to such effects (no ill-conditioned weight estimation), quantitative gains in PSL or resolution may differ in heterogeneous tissue. Future work should include *aberrator phantoms* and *in-vivo* validation, and assess contrast metrics such as CNR and gCNR alongside PSL and FWHM. Moreover, we tested two *deterministic* VS configurations and a fixed number of VSs. In chapter 5, we prove that VS geometry strongly governs the trade-off between resolution, sidelobes, and uniformity. An optimized VS positions and mask weights via multi-objective optimization (e.g., minimizing PSL and FWHM variance across depth while constraining mainlobe width and frame rate) will further enhance image quality. Also, the 25% active-element case used a random subset of elements. Optimizing sparse layouts [181][182] would better control spatial aliasing and reduce grating lobes. Combining *optimized sparse arrays* with the *compound mask* and optimized VSs should further mitigate sidelobe levels while retaining low channel count. Compared with long 1D apertures (linear/convex arrays), a square matrix probe has a smaller effective aperture in each lateral dimension. This inherently limits lateral resolution and the extent to which apodization alone can reduce sidelobes, therefore using *larger physical apertures* [183] together with compound masks is expected to improve both FWHM and PSL further.

7.5 Conclusion

In this work, we have extended the concept of compound mask apodization from two-dimensional imaging to fully three-dimensional Diverging Wave Imaging (DWI) using a matrix probe. The proposed framework introduces a voxel-adaptive mapping strategy that transfers transmit apodization laws from Synthetic Transmit Aperture Imaging (STAI) into the DWI domain. By combining source-voxel-aperture intersection analysis with interpolation, the method ensures that the local characteristics of STAI apodization are preserved in DWI, while also providing a consistent treatment of out-of-footprint voxels

through the use of virtual apertures. The performance of the approach was validated through FIELD II simulations with a 32×32 element matrix probe. Both qualitative contour analysis and quantitative point spread function (PSF) measurements confirmed that the compound mask preserves the benefits of adaptive apodization, maintaining narrow main lobes similar to those of uniform excitation while significantly suppressing sidelobes, outperforming static apodization schemes such as the 2D Hanning window on virtual sources.

Chapter 8

Conclusion and Future Work

8.1 Conclusion

Ultrafast ultrasound can be achieved by transmitting unfocused wavefronts—either plane waves (PWs) or spherical/diverging waves (DWs). While PW imaging is now a mature line of research, DW imaging has been explored less extensively. DW transmissions offer two key benefits: they naturally cover a large field of view and support very high frame rates. Their main drawback is a faster decay of acoustic power with depth, since the emitted energy spreads over an expanding surface. In practice, DW image quality is often improved by increasing the number of emissions (i.e., virtual sources), but this reduces the frame rate. Balancing high image quality and a high frame rate remains an active area of research.

This dissertation aims to improve image quality in ultrafast imaging using DWs without sacrificing frame rate and real-time performance. In chapter 3, we introduced a closed-form compound-mask apodization for DWI that transfers any STAI transmit apodization law (e.g., Hanning, Blackman–Harris, Tukey) into *compound weights* for diverging-wave compounding, valid for linear and convex arrays. Under a short-transmit-pulse condition, the matrix linking the vector of STAI apodization and the vector of compound mask values can be approximated as diagonal, yielding a simple geometric construction: for each image point and VS, we compute the intersection of the VS–pixel(or voxel) ray with the probe surface and *sample* the continuous STAI apodization at that location. This produces per-transmit weights that (i) emulate the desired transmit taper without iterative optimization and (ii) can be precomputed offline and applied at negligible runtime cost. Field II simulations confirmed that the mapped DWI preserves STAI-like mainlobe widths with only minor deviations (e.g., SALA 0.40 mm vs. DWLA 0.43 mm at 20 mm depth). Next, in chapter 4, we present real-time experimental validation on Verasonics scanner, where we integrated the compound mask into the Verasonics pipeline using standard buffers ($\text{Tx}(m)$.Apod, $\text{Receive}(m)$.Apod,

per-event InterBuffer frames, and a lightweight masked compounding step). This design preserved real-time performance adding *one scalar multiplication per pixel per VS*. On a multipurpose phantom, DWI with the mask (L7-4 linear probe) improved lateral resolution by $\approx 30\%$ and enhanced contrast ratio (CR) across anechoic and grayscale cysts (e.g., anechoic CR from -16.40 to -21.97 dB; -6 dB target from -4.22 to -6.41 dB; $+6$ dB target from 4.91 to 5.57 dB). With a C5-2 convex probe and three VS layouts (linear, curvilinear, tilted), the mask consistently improved both resolution and contrast, with anechoic CR gains up to ~ 6 dB and ~ 1.8 dB on -6 and $+6$ dB targets. A paired one-sided sign test on FWHM (20 paired measurements) yielded $p = 0.041$, confirming a statistically significant resolution improvement (mean paired difference -0.642 mm; Hodges–Lehmann -0.597 mm). Complexity-wise, the mask is $\mathcal{O}(KM)$ per frame (pixels \times transmits), substantially lighter than adaptive apodization approaches like coherence-factor ($\mathcal{O}(KNM)$) or minimum-variance beamforming ($\mathcal{O}(KN^2M)$ to $\mathcal{O}(KN^3M)$).

To further enhance the image quality and improve frame rate capabilities, chapter 5 introduced a multi-objective optimization of VS distributions with compound masking for convex arrays. In this chapter, we formulated a Pareto optimization of VS positions using a genetic algorithm (`gamultiobj`) with two objectives averaged over multiple scatterers: lateral resolution (FWHM) and peak sidelobe level (PSL). The compound mask was *embedded* in the optimization algorithm. Compared with deterministic layouts (uniform linear/curvilinear VS grids), the optimized fronts dominated in both objectives. For 20 VSs, the *optFWHM* solution, Pareto-optimal solution that minimizes the average FWHM, halved median FWHM (~ 0.5 mm vs. 1.0 – 1.1 mm, a $\sim 50\%$ reduction), and the *optPSL*, Pareto-optimal solution that minimizes the average PSL, improved median CR to ~ -16 dB versus -12 and -10 dB for curvilinear and linear baselines. Crucially, reducing transmissions preserved quality: 10-, 6-, and even 4-VS *optimized* layouts maintained sub-millimetric FWHM and CR in the -14 to -16 dB range in simulation, outperforming deterministic 20-VS references while enabling higher frame rates.

To translate these methods into a practical, high-throughput pipeline, in chapter 6 we designed a GPU-based delay-and-sum beamformer that (i) decouples transmit/receive delays and apodization, (ii) exploits lateral translation symmetries to shrink receive look-up tables, and (iii) supports both RF-band and baseband (I/Q) processing with carrier-delay correction. On a single NVIDIA RTX-class GPU, per-frame processing times scaled approximately with the number of transmissions. Memory usage was characterized end-to-end: raw-data and focalization stores both decreased markedly with fewer VSs (e.g., raw-data ~ 31 MB $\rightarrow \sim 6$ MB from 20 to 4 VSs). Importantly, optimized VS sets preserved image quality at reduced transmissions: e.g., anechoic CR ~ -26 dB for optimized 20-VS versus ~ -10 dB

for a deterministic curvilinear reference, while the optimized 10-VS and 6-VS case remained close in both FWHM and CR. Finally, we generalized the compound-mask mapping to 3D volumes and 2D apertures in chapter 7. The ray–aperture intersection and interpolation were extended to two lateral dimensions, and a *virtual aperture* was introduced to handle intersections falling just outside the physical footprint, ensuring well-defined weights over the entire volume of interest. FIELD II studies with 32×32 and 32×32 matrices with 25% sparse random elements showed that the 3D compound mask preserves local STAI apodization contours at representative voxels and provides a deterministic, low-overhead path to transmit apodization in volumetric diverging-wave imaging.

Summary. Taken together, this dissertation brings the following contributes: (i) a closed-form, geometry-based apodization mapping from STAI to DWI that is data-independent, portable across probes, and efficient at runtime; (ii) the first real-time experimental validation of compound-mask DWI on a research scanner, with statistically significant resolution gains and consistent contrast improvements across VS strategies; (iii) a Pareto-optimization of VS distributions that, coupled with the mask, delivers superior FWHM/PSL trade-offs and maintains quality with far fewer transmissions; (iv) an efficient GPU receive beamformer that realizes these ideas under realistic computation and bandwidth constraints; and (v) a principled extension of the mapping to 3D with matrix arrays. Collectively, these advances bridge theory, algorithms, and systems, enabling high-quality ultrafast DWI at practical frame rates.

8.2 Future Directions

To advance the proposed framework for ultrafast diverging–wave (DW) imaging and broaden its clinical impact, we outline the following research directions. Each item targets a specific limitation observed in this work or in the broader literature and proposes concrete experimental protocols and metrics for rigorous evaluation.

- In-vivo abdominal and Doppler validation with convex arrays: a key next step is to validate the compound mask and the optimized VS distributions *in vivo* on abdominal organs (e.g., liver, kidney) and vascular beds using convex arrays. Ultrafast plane/diverging-wave imaging has already shown clinical promise for abdominal and vascular applications when paired with clutter-robust Doppler processing and coherent compounding [184]. Ultrafast Doppler methods (power/vector Doppler with SVD/low-rank clutter suppression) resolve slow microvascular flows while maintaining wide fields of view and high frame rates [185, 186]. Given the resolution/contrast gains we

observed with compound weighting and VS optimization on phantoms, a structured *in vivo* study should:

- compare DWI with and without compound weights under identical convex-probe sequences for liver and renal parenchyma, quantifying FWHM, CR, and Doppler SNR (including wall-filter/SVD operating points);
- assess flow sensitivity (slow-flow detectability, velocity bias/noise) in parenchymal and vascular targets with ultrafast Doppler pipelines [185, 186];
- report frame-rate/latency trade-offs versus number of VS transmissions (e.g., 20/10/6/4), using our optimized distributions to retain diagnostic image quality at reduced transmission counts.

These experiments will test whether the compound mask and optimized VSs translate to clinically relevant abdominal imaging, where the field of view, penetration, and Doppler sensitivity are critical.

- To move toward product-grade deployment, the optimization-driven GPU beamformer should be embedded inside a scanner pipeline, co-designed with acquisition firmware and host software. Several engineering refinements can further reduce latency and memory traffic while preserving accuracy:
 - **Phase rotation via LUT/CORDIC in BBB:** replace per-sample trigonometric calls with table-lookup (LUT) or CORDIC micro-kernels for carrier phase rotation in baseband beamforming [187]. This eliminates expensive transcendentals, reduces jitter, and improves runtime determinism.
 - **Subsampled focalization with on-the-fly interpolation:** store receive delay/apodization tables on coarser lateral/elevational grids and reconstruct per-pixel values with linear (or cubic) interpolation on GPU, cutting memory footprint while preserving alignment accuracy (building on the digital interpolation principles underlying delay-and-sum [188]).
 - **Streaming/tiling of raw data per transmit:** stream channel data one transmit (TX) at a time and tile depth along the z axis to minimize peak memory; this is compatible with the depth-first kernels used in our implementation and with established GPU beamforming practice [189, 190].

A system-level evaluation should report frame rate, latency, and bandwidth headroom against clinical presets (e.g., abdominal general, vascular, obstetric), benchmarking both RF-band and baseband beamforming paths [189, 190].

- From 2D to resource-efficient 3D: sparse apertures + optimized VS + compound weights:

Our 3D extension in chapter 7 showed that voxel-based compound weights can be mapped in volumetric DWI and that the approach remains beneficial. The next phase is to *co-design* (i) sparse or row–column-addressed matrix apertures, (ii) optimized VS constellations in 3D, and (iii) compound weights, targeting high image quality at minimal hardware resources. The next important step would be to apply the above described approach to volumetric high-frame-rate (HFR) Doppler for *in vivo* assessment of valve dynamics, intracardiac flows, and fast transients where spatio-temporal resolution is important; leveraging the maturing body of work in HFR/ultrafast echocardiography and volumetric Doppler [184, 191]. This line of work aims to make volumetric ultrafast imaging practical with reduced hardware while enabling applications that benefit from 3D high frame rate (e.g., valve regurgitation timing, microvascular perfusion mapping, and fast flow transients [192] [193]).

Bibliography

- [1] Matthieu Toulemonde. *New beamforming strategy for improved ultrasound imaging: application to biological tissues nonlinear imaging*. Ph.d. dissertation, 2014. Available via ResearchGate. pages xvii, 11
- [2] M. Y. Benane. *Ultrafast, broadband and multi-pulse transmissions for ultrasonic imaging*. Ph.d. dissertation, Université de Lyon, Lyon, France, 2018. pages xvii, 12, 26
- [3] Hervé Liebgott, Alfonso Rodriguez-Molares, Lucilio Cordero-Grande, Jørgen Arendt Jensen, and Olivier Bernard. Plane-wave imaging challenge in medical ultrasound. In *Proceedings of the IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, September 2016. pages xvii, 17, 18, 19, 56
- [4] M. Tanter and M. Fink. Ultrafast imaging in biomedical ultrasound. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2014. pages xvii, 24
- [5] Daniel Posada, Jonathan Porée, Arnaud Pellissier, Boris Chayer, François Tournoux, Guy Cloutier, and Damien Garcia. Staggered multiple-PRF ultrafast color doppler. *IEEE Transactions on Medical Imaging*, 35(6):1510–1521, June 2016. pages xvii, 25
- [6] S. Wang, J. A. Hossack, and A. L. Klibanov. From anatomy to functional and molecular biomarker imaging and therapy: Ultrasound is safe, ultrafast, portable, and inexpensive. *Investigative Radiology*, 55(9):559–572, September 2020. pages 1
- [7] Thomas L. Szabo. *Diagnostic Ultrasound Imaging: Inside Out*. Academic Press Series in Biomedical Engineering. Elsevier Academic Press, Amsterdam; Boston, 2004. OCLC: ocm57744110. pages 1
- [8] P. R. Hoskins, K. Martin, and A. Thrush. *Diagnostic Ultrasound: Physics and Equipment*. Cambridge University Press, Cambridge, U.K., 2010. pages 1
- [9] L. Sandrin, S. Catheline, M. Tanter, X. Hennequin, and M. Fink. Time-resolved pulsed elastography with ultrafast ultrasonic imaging. *Ultrasonic Imaging*, 21(4):259–272, October 1999. pages 1, 38
- [10] H. Hasegawa and H. Kanai. High-frame-rate echocardiography using diverging transmit beams and parallel receive beamforming. *Journal of Medical Ultrasonics*, 38(3):129–140, July 2011. pages 1, 38, 39
- [11] J. Poree et al. High-frame-rate echocardiography using coherent compounding with doppler-based motion-compensation. *IEEE Transactions on Medical Imaging*, 35(7):1647–1657, July 2016. pages 1, 38, 39

- [12] M. Dighe and M. Bruce. Elastography of diffuse liver diseases. *Seminars in Roentgenology*, 51(4):358–366, 2016. pages 2
- [13] M. Bruce, A. Hannah, C. Tremblay-Darveau, and P. Burns. High frame-rate visualization of blood flow with ultrasound contrast agents. *Journal of the Acoustical Society of America*, 140(4):3028, 2016. pages 2
- [14] E. Macé et al. Functional ultrasound imaging of the brain. *Nature Methods*, 8:662–664, May 2011. pages 2
- [15] J. Provost, C. Papadacci, J. E. Arango, M. Imbault, M. Fink, J.-L. Gennisson, M. Tanter, and M. Pernot. 3d ultrafast ultrasound imaging in vivo. *Physics in Medicine and Biology*, 59(19):L1–L13, oct 2014. pages 2
- [16] C. Papadacci, M. Pernot, M. Couade, M. Fink, and M. Tanter. High contrast ultrafast imaging of the heart. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(2):288–301, 2014. pages 2, 39, 80
- [17] A. Rodriguez-Molares, H. Torp, B. Denarie, and L. Løvstakken. The angular apodization in coherent plane-wave compounding [correspondence]. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(11):2018–2023, 2015. pages 3, 38, 41, 42, 101
- [18] M. O. Culjat, D. Goldenberg, P. Tewari, and R. S. Singh. A review of tissue substitutes for ultrasound imaging. *Ultrasound in Medicine & Biology*, 36(6):861–873, 2010. pages 10
- [19] Frederick W. Kremkau. *Sonography: Principles and Instruments*. Elsevier, Philadelphia, 10 edition, 2020. pages 14
- [20] V. Perrot, M. Polichetti, F. Varray, and D. Garcia. So you think you can DAS? a viewpoint on delay-and-sum beamforming. *Ultrasonics*, 111:106309, 2021. pages 15, 27
- [21] Y. Wang, Y. Wang, M. Liu, Z. Lan, C. Zheng, and H. Peng. Minimum variance beamforming combined with covariance matrix-based adaptive weighting for medical ultrasound imaging. *BioMedical Engineering OnLine*, 21(1):40, 2022. pages 15
- [22] B. M. Asl and A. Mahloojifar. Contrast enhancement and robustness improvement of adaptive ultrasound imaging using forward-backward minimum variance beamforming. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 58(4):858–867, 2011. pages 15
- [23] J. F. Synnevåg, A. Austeng, and S. Holm. Adaptive beamforming applied to medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(8):1606–1613, 2007. pages 15, 27, 28
- [24] J.-F. Synnevåg, A. Austeng, and S. Holm. Benefits of minimum-variance beamforming in medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(9):1868–1879, 2009. pages 15

- [25] B. M. Asl and A. Mahloojifar. Eigenspace-based minimum variance beamforming applied to medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(11):2381–2390, 2010. pages 15
- [26] R. Mallart and M. Fink. Adaptive focusing in scattering media through sound-speed inhomogeneities: The van cittert–zernike approach and focusing criterion. *The Journal of the Acoustical Society of America*, 96(6):3721–3732, 1994. pages 15
- [27] K. Hollman, K. Rigby, and M. O’Donnell. Coherence factor of speckle from a multi-row probe. In *Proceedings of the 1999 IEEE Ultrasonics Symposium*, volume 2, pages 1257–1260. IEEE, 1999. pages 15
- [28] P.-C. Li and M.-L. Li. Adaptive imaging using the generalized coherence factor. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 50(2):128–141, 2003. pages 16, 27
- [29] J. Camacho, M. Parrilla, and C. Fritsch. Phase coherence imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(5):958–974, 2009. pages 16
- [30] C. Fritsch, J. Camacho, and M. Parrilla. New ultrasound imaging techniques with phase coherence processing. *Ultrasonics*, 50(2):122–126, 2010. pages 16
- [31] G. Matrone, A. S. Savoia, G. Caliano, and G. Magenes. The delay multiply and sum beamforming algorithm in ultrasound b-mode medical imaging. *IEEE Transactions on Medical Imaging*, 34(4):940–949, 2014. pages 16, 27, 28
- [32] F. Prieur, O. M. H. Rindal, and A. Austeng. Signal coherence and image amplitude with the filtered delay multiply and sum beamformer. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 65(7):1133–1140, 2018. pages 16
- [33] M. Ramalli, G. Scaringella, G. Matrone, A. Dallai, E. Boni, A. S. Savoia, L. Bassi, G. E. Hine, and P. Tortoli. High dynamic range ultrasound imaging with real-time filtered delay multiply and sum beamforming. In *Proc. IEEE Int. Ultrasonics Symp. (IUS)*, pages 1–4. IEEE, 2017. pages 16
- [34] M. Mozaffarzadeh, A. Mahloojifar, M. Orooji, S. Adabi, and M. Nasirivanaki. Double stage delay multiply and sum beamforming algorithm: Application to linear-array photoacoustic imaging. *IEEE Trans. Biomed. Eng.*, 65(1):31–42, 2017. pages 16
- [35] A. Vayyeti and A. K. Thittai. A filtered delay weight multiply and sum (f-dwmas) beamforming for ultrasound imaging: Preliminary results. In *Proc. IEEE Int. Symp. Biomed. Imaging (ISBI)*, pages 312–315. IEEE, 2020. pages 16
- [36] C. Keeble, S. Wolstenhulme, A. Davies, and J. Evans. Is there agreement on what makes a good ultrasound image? *Ultrasound*, 21(3):118–123, 2013. pages 17
- [37] S. C. Metcalfe and J. A. Evans. A study of the relationship between routine ultrasound quality assurance parameters and subjective operator image assessment. *The British Journal of Radiology*, 65(775):570–575, 1992. pages 17

- [38] S. Wolstenhulme, A. G. Davies, C. Keeble, S. Moore, and J. Evans. Agreement between objective and subjective assessment of image quality in ultrasound abdominal aortic aneurysm screening. *The British Journal of Radiology*, 88(1046):20140482, 2015. pages 17
- [39] M. van Wijk and J. Thijssen. Performance testing of medical ultrasound equipment: Fundamental vs. harmonic mode. *Ultrasonics*, 40(1–8):585–591, may 2002. pages 17, 18
- [40] E. Roux. *2D Sparse Array Optimization and Operating Strategy for Real-Time 3D Ultrasound Imaging*. PhD thesis, Université de Lyon; Università degli Studi di Firenze, 2016. pages 19, 100
- [41] O. M. H. Rindal, A. Austeng, A. Fatemi, and A. Rodriguez-Molares. The effect of dynamic range alterations in the estimation of contrast. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 66(7):1198–1208, July 2019. pages 19
- [42] A. Rodriguez-Molares, O. M. H. Rindal, J. D’hooge, S.-E. Masoy, A. Austeng, M. A. Lediju Bell, and H. Torp. The generalized contrast-to-noise ratio: A formal definition for lesion detectability. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(4):745–759, April 2020. pages 19, 20, 27
- [43] Retrospective transmit beamformation. https://cdn0.scrvt.com/39b415fb07de4d9656c7b516d8e2d907/1800000000064734/cc581ee74214/Whitepaper_Bradley_1800000000064734.pdf. Accessed: 2025-11-10. pages 20, 21, 27
- [44] Accelerated focused ultrasound imaging. <https://pmc.ncbi.nlm.nih.gov/articles/PMC2813841/>. Accessed: 2025-11-10. pages 20, 22
- [45] A review of ultrasound imaging methods and techniques to enhance their frame rate. <https://ijeee.iust.ac.ir/article-1-3632-en.pdf>. Accessed: 2025-11-10. pages 20, 22, 23
- [46] D. P. Shattuck, M. D. Weinschenker, S. W. Smith, and O. T. von Ramm. Explososcan: A parallel processing technique for high speed ultrasound imaging with linear phased arrays. *The Journal of the Acoustical Society of America*, 75:1273–1282, 1984. pages 21
- [47] O. T. von Ramm, S. W. Smith, and H. G. Pavy. High-speed ultrasound volumetric imaging system. ii. parallel processing and image display. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 38(2):109–115, 1991. pages
- [48] A. Rabinovich, Z. Friedman, and A. Feuer. Multi-line acquisition with minimum variance beamforming. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(12):2521–2534, December 2013. pages 21
- [49] Imaging with unfocused regions of focused ultrasound beams. *Ultrason. Imaging*. Accessed: 2025-11-10. pages 21

- [50] Simultaneous axial multifocal imaging using a single acoustical transmission: a practical implementation. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*. Accessed: 2025-11-10. pages 21
- [51] Optimized simultaneous axial multifocal imaging via frequency multiplexed focusing. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*. Accessed: 2025-11-10. pages 21, 22, 27
- [52] High frame-rate, high resolution ultrasound imaging with multi-line transmission and filtered-delay multiply and sum beamforming. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2017. Accessed: 2025-11-10. pages 21, 22, 27
- [53] Multi-line transmission in medical imaging using the second-harmonic signal. *Ultrasonics*, 2014. Accessed: 2025-11-10. pages 21, 22
- [54] Optimized simultaneous axial multifocal imaging via frequency multiplexed focusing. <https://cris.tau.ac.il/en/publications/optimized-simultaneous-axial-multifocal-imaging-via-frequency-mul>. Accessed: 2025-11-10. pages 22
- [55] Performance of f-dmas beamforming with adjustable maximum spatial lag in multi-line transmission ultrasound imaging. *Applied Sciences*, 8(4):486, 2018. Accessed: 2025-11-10. pages 22
- [56] Spatial coherence of backscattered signals in multi-line transmit ultrasound imaging and its effect on short-lag filtered-delay multiply and sum beamforming. *Applied Sciences*, 8(4):486, 2018. Accessed: 2025-11-10. pages 22
- [57] Requirements and hardware limitations of high-frame-rate 3-d ultrasound imaging systems. *Applied Sciences*, 12(13):6562, 2022. Accessed: 2025-11-10. pages 22
- [58] Computational complexity reduction of synthetic-aperture focus in ultrasound imaging using frequency-domain reconstruction. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2015. Accessed: 2025-11-10. pages 22, 23, 28
- [59] D. H. Evans and W. N. McDicken. *Doppler Ultrasound: Physics, Instrumentation, and Signal Processing*. Wiley, Chichester; New York, 2000. pages 22
- [60] L. Sandrin, S. Catheline, M. Tanter, X. Hennequin, and M. Fink. Time resolved pulsed elastography with ultrafast ultrasonic imaging. *Ultrasonic Imaging*, 21(4):259–272, 1999. pages 22
- [61] J. A. Jensen, S. Nikolov, A. C. Yu, and D. Garcia. Ultrasound vector flow imaging: I: Sequential systems. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, pages 1–1, 2016. pages 22
- [62] Ultrasonic synthetic apertures: Review. *Archives of Acoustics*. Accessed: 2025-11-10. pages 23
- [63] Synthetic transmit aperture in ultrasound imaging. <https://oldwww.ippt.pan.pl/Repository/o1220.pdf>. Accessed: 2025-11-10. pages 23

- [64] Snr analysis of multi-aperture ultrasound and photoacoustic imaging systems. *The Journal of the Acoustical Society of America*, 157(2):1228–1240. Accessed: 2025-11-10. pages 23
- [65] Henrik Andresen. *Synthetic Aperture Beamforming in Ultrasound using Moving Arrays*. PhD thesis, Technical University of Denmark, 2007. Accessed: 2025-11-10. pages 23
- [66] Ultrafast ultrasound imaging. MDPI Books, 2021. Online book. pages 23, 25
- [67] Plane wave imaging through interfaces. *Sensors*, 2021. pages 23, 24
- [68] Experimental evaluation of spectral-based quantitative ultrasound imaging using plane wave compounding. *Ultrasound Med. Biol.*, 2014. pages 24
- [69] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink. Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(3):489–506, 2009. pages 24, 27
- [70] Coherent plane wave compounding combined with tensor completion applied for ultrafast imaging. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2021. pages 24
- [71] A two-dimensional angular interpolation based on radial basis functions for high frame rate ultrafast imaging. *Ultrasonics*, 2023. pages 24, 26, 27
- [72] Grating lobe reduction in plane-wave imaging with angular compounding using subtraction of coherent signals. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2022. Open-access summary at PMC. pages 24
- [73] Wide field-of-view ultrafast curved array imaging using diverging waves. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2020. pages 24, 25, 27
- [74] Compound mask for divergent wave imaging in medical ultrasound. arXiv preprint, 2025. pages 24, 27
- [75] Distributing synthetic focusing over multiple push-detect events enhances shear wave elasticity imaging performance. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control / PMC*, 2019. pages 25
- [76] M. Berson, A. Roncin, and L. Pourcelot. Compound scanning with an electrically steered beam. *Ultrasonic Imaging*, 3(3):303–308, 1981. pages 25
- [77] S. K. Jespersen, J. E. Wilhjelm, and H. Sillesen. Multi-angle compound imaging. *Ultrasonic Imaging*, 20(2):81–102, 1998. pages 25
- [78] M. Tanter, J. Bercoff, L. Sandrin, and M. Fink. Ultrafast compound imaging for 2-d motion vector estimation: Application to transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 49(10):1363–1374, oct 2002. pages 25

- [79] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink. Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(3):489–506, mar 2009. pages 26
- [80] Iterative minimum variance beamformer with low complexity for medical ultrasound imaging. *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2018. pages 26
- [81] Dynamic coherence factor based on the standard deviation for coherent plane-wave compounding. *Ultrasonics*, 2019. pages 27
- [82] An improved spatio-temporally smoothed coherence factor combined with delay multiply and sum beamformer. *Electronics*, 2023. pages 27
- [83] Enrico Boni, Luca Bassi, Andrea Dallai, Francesco Guidi, Alessandro Ramalli, Stefano Ricci, and Piero Tortoli. Ula-op 256: A 256-channel open scanner for development and real-time implementation of new ultrasound methods. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(10):1488–1495, October 2016. pages 28, 65, 101
- [84] Verasonics, Inc., Kirkland, WA, USA. *Vantage Research Ultrasound Systems: Sales Brochure*, 2021. Accessed Nov. 10, 2025. pages 28
- [85] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983. pages 29, 30, 68
- [86] J. H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Oxford, England, 1975. pages 29, 31, 68
- [87] E. Roux, A. Ramalli, P. Tortoli, C. Cachard, M. C. Robini, and H. Liebgott. 2-d ultrasound sparse arrays multidepth radiation optimization using simulated annealing and spiral-array inspired energy functions. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(12):2138–2149, December 2016. pages 30, 68
- [88] M. C. Robini. Theoretically grounded acceleration techniques for simulated annealing. In I. Zelinka, V. Snáßsel, and A. Abraham, editors, *Handbook of Optimization*, volume 38, pages 311–335. Springer, Berlin, Germany, 2013. pages 30, 70
- [89] Bruce Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13(2):311–329, 1988. pages 31
- [90] Marc-Christophe Robini and Pierre-Jean Reissman. From simulated annealing to stochastic continuation: A new trend in combinatorial optimization. *Journal of Global Optimization*, 56(1):185–215, May 2013. pages 31
- [91] Olivier Catoni. Rough large deviation estimates for simulated annealing: Application to exponential schedules. *The Annals of Probability*, 20(3):1109–1146, 1992. pages 31

- [92] P. K. Weber, A. Austeng, S. Holm, and N. Aakvaag. 1d- and 2d sparse-array-optimization. *Instrumentation Science & Technology*, 27(4):235–246, September 1999. pages 31
- [93] S. Caorsi, A. Lommi, A. Massa, S. Piffer, and A. Trucco. Planar antenna array design with a multi-purpose ga-based procedure. *Microwave and Optical Technology Letters*, 35(6):428–430, December 2002. pages
- [94] P. Yang, B. Chen, and K.-R. Shi. A novel method to design sparse linear arrays for ultrasonic phased array. *Ultrasonics*, 44:e717–e721, December 2006. pages
- [95] F. Le Courtois, J.-H. Thomas, F. Poisson, and J.-C. Pascal. Genetic optimization of a plane array geometry for beamforming. application to source localization in a high-speed train. *Journal of Sound and Vibration*, 371:78–93, June 2016. pages 31, 68
- [96] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, 1996. pages 31
- [97] Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Genetic and Evolutionary Computation. Springer, New York, 2nd edition, 2007. pages 32
- [98] K. Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. Springer, London, U.K., 2001. [Online]. Available: https://link.springer.com/chapter/10.1007/978-0-85729-652-8_1. pages 33, 78
- [99] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *Parallel Problem Solving from Nature — PPSN VI*, volume 1917 of *Lecture Notes in Computer Science*, pages 849–858, Berlin, Heidelberg, 2000. Springer. pages 33
- [100] Carlos A. Coello Coello, Gregorio T. Pulido, and M. Salazar Lechuga. Handling multiple objectives with particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, 8(3):256–279, 2004. pages 33
- [101] G. Le Moign, P. Masson, O. Basset, and H. Liebgott. Optimized virtual sources distributions for 3-d ultrafast diverging wave compounding imaging: A simulation study. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 70(10):1319–1328, October 2023. pages 35, 69
- [102] P. J. Thoen. Aperture apodization to reduce the off-axis intensity of the pulsed mode directivity function of linear arrays. *Ultrasonics*, 20(5):231–236, 1982. pages 37, 100
- [103] C. M. W. Daft and W. E. Engeler. Windowing of wide-band ultrasound transducers. In *Proc. IEEE Ultrasonics Symposium*, pages 1541–1544, 1996. pages 37, 100
- [104] K. E. Thomenius. Evolution of ultrasound beamformers. In *Proc. IEEE Ultrasonics Symposium*, pages 1615–1622, 1996. pages 37
- [105] D. A. Guenther and W. F. Walker. Optimal apodization design for medical ultrasound using constrained least squares—part ii: simulation results. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(2):343–358, 2007. pages 37

- [106] S. Repetto and A. Trucco. A stochastic approach for the apodization of very short arrays. *Ultrasonics*, 42:425–429, 2004. pages 37
- [107] S. M. Sakhaei, A. Mahloojifar, and H. Ghassemian. A transformation-based method to design ultrasound array. *Ultrasonics*, 49(2):179–184, 2009. pages 37
- [108] J. F. Synnevåg, A. Austeng, and S. Holm. Adaptive beamforming applied to medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 54(8):1606–1613, 2007. pages 38
- [109] C. H. Seo and J. T. Yen. Sidelobe suppression in ultrasound imaging using dual apodization with cross-correlation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 55(10):2198–2210, 2008. pages 38
- [110] J. F. Synnevåg, A. Austeng, and S. Holm. A low-complexity data-dependent beam-former. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 58(2):281–289, 2011. pages 38
- [111] J. A. Jensen, S. I. Nikolov, K. L. Gammelmark, and M. H. Pedersen. Synthetic aperture ultrasound imaging. *Ultrasonics*, 44:e5–e15, 2006. pages 38
- [112] O. M. H. Rindal and A. Austeng. Double adaptive plane-wave imaging. In *Proc. IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, Tours, France, 2016. pages 38
- [113] S. Goudarzi, A. Asif, and H. Rivaz. Angular apodization estimation using independent component analysis in coherent plane-wave compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(5):989–998, 2020. pages 38
- [114] V. Komini, P. Santos, and J. D’hooge. Diverging wave compounding: Direct comparison of two popular approaches. In *Proc. IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, Washington, DC, USA, 2017. pages 38, 39
- [115] Z. Alzein, M. Crocco, and D. D. Caviglia. Computationally efficient rf band and baseband beam-former for coherent plane wave imaging. In *2024 19th Conference on Ph.D Research in Microelectronics and Electronics (PRIME)*, pages 1–4, Larnaca, Cyprus, 2024. pages 41, 55, 93
- [116] J. A. Jensen. Field: A program for simulating ultrasound systems. In *10th Nordic-Baltic Conference on Biomedical Imaging, Medical & Biological Engineering & Computing*, volume 34, pages 351–353, 1996. pages 48, 76, 102, 108
- [117] J. A. Jensen and N. B. Svendsen. Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 39(2):262–267, 1992. pages 48, 76, 102, 108
- [118] S. Liang and L. Wang. A study of wide unfocused wavefront for convex-array ultrasound imaging. *Ultrasonics*, 134:107080, Sep. 2023. pages 48, 55, 76
- [119] Z. Alzein, M. Crocco, and D. D. Caviglia. Compound mask for divergent wave imaging in medical ultrasound. *IEEE Access*, 13:59675–59686, 2025. pages 51

- [120] Vincent Perrot, Adrien Besson, Olivier Bernard, and Hervé Liebgott. So you think you can das? a viewpoint on delay-and-sum beamforming. *Ultrasonics*, 116:106524, 2021. pages 62
- [121] Thomas L. Szabo. *Diagnostic Ultrasound Imaging: Inside Out*. Elsevier, 2 edition, 2014. pages 62
- [122] Jørgen Arendt Jensen. Linear description of ultrasound imaging systems. Technical report, Technical University of Denmark, 1999. Technical notes. pages 62
- [123] M. Hisatsu, S. Misumi, and M. Yoshizawa. Generalized coherence factor estimated from real signals. *Journal of Medical Ultrasonics*, 2020. pages 62
- [124] Y. Wang, X. Li, Q. Zhang, and Y. Zhou. Dynamic coherence factor based on the standard deviation. *Computers in Biology and Medicine*, 107:174–183, 2019. pages 62
- [125] K. Ozgun, P. Li, A. Basarab, D. Friboulet, and H. Liebgott. A spatial coherence beamformer design for power doppler imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 66(12):1936–1946, 2019. pages 62
- [126] K. E. Song, J. H. Lee, and C. Kim. A comparison of coherence factor and sign coherence. *Journal of Medical Biomechanics*, 2023. pages 62
- [127] J. Capon. High-resolution frequency–wavenumber spectrum analysis. *Proceedings of the IEEE*, 57(8):1408–1418, 1969. pages 62
- [128] J.-F. Synnevåg, A. Austeng, and S. Holm. Benefits of minimum-variance beamforming in medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(9):1868–1879, 2009. pages 62
- [129] I. K. Holfort and J. A. Jensen. Broadband minimum variance beamforming for ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(2):314–325, 2009. pages 62
- [130] Babak M. Asl and Ali Mahloojifar. Minimum variance beamforming combined with adaptive coherence weighting applied to medical ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 56(9):1923–1931, 2009. pages 63
- [131] W. J. Conover. *Practical Nonparametric Statistics*. Wiley, 3 edition, 1999. pages 63
- [132] Philip Good. *Permutation, Parametric, and Bootstrap Tests of Hypotheses*. Springer, 3 edition, 2005. pages 63
- [133] Qian You, Jan D’hooge, Alfonso Rodriguez-Molares, Jørgen Arendt Jensen, and Andreas Austeng. Pixel-oriented adaptive apodization for plane wave imaging based on recovery of the complete dataset. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 69(2):512–522, February 2022. pages 65

- [134] S. Goudarzi, A. Asif, and H. Rivaz. Angular apodization estimation using independent component analysis in coherent plane-wave compounding. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(5):989–998, May 2020. pages 65
- [135] Vincent Perrot, Sara Goudarzi, Angelo Deho, Shiqi Niu, Hassan Rivaz, and Hervé Liebgott. Translation of simultaneous vessel wall motion and vectorial blood flow imaging in healthy and diseased carotids to the clinic: A pilot study. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(3):558–569, March 2021. pages 65
- [136] Thomas L. Szabo and Peter A. Lewin. Ultrasound transducer selection in clinical imaging practice. *Journal of Ultrasound in Medicine*, 32(4):573–582, April 2013. pages 65
- [137] Maja Cikes, Lucy Tong, George R. Sutherland, and Jan D’hooge. Ultrafast cardiac ultrasound imaging: Technical principles, applications, and clinical benefits. *JACC: Cardiovascular Imaging*, 7(8):812–823, August 2014. pages 65
- [138] N. Ghigo, A. Ramalli, S. Ricci, P. Tortoli, and H. Liebgott. Sequence optimization for high frame rate imaging with a convex array. In *Proc. IEEE Int. Ultrason. Symp. (IUS)*, pages 1–4, Las Vegas, NV, USA, 2020. pages 67
- [139] A. Besson, F. Wintzenrieth, B. Hériard-Dubreuil, and C. Cohen-Bacrie. On archimedean-spiral-based imaging. In *Proc. IEEE Int. Ultrason. Symp. (IUS)*, pages 1–4, Las Vegas, NV, USA, 2020. pages 67
- [140] A. Besson, F. Wintzenrieth, and C. Cohen-Bacrie. Vector-flow imaging in convex-array configurations. In *Proc. IEEE Int. Ultrason. Symp. (IUS)*, pages 1–4, Las Vegas, NV, USA, 2020. pages 67
- [141] S. Liang and L. Wang. Fourier beamformation for convex-array diverging wave imaging using virtual sources. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 69(5):1625–1637, May 2022. pages 68
- [142] R. L. Haupt. Thinned arrays using genetic algorithms. *IEEE Transactions on Antennas and Propagation*, 42(7):993–999, July 1994. pages 68
- [143] MathWorks. *Global Optimization Toolbox*. Natick, MA, USA, 2021. pages 69
- [144] Z. Michalewicz and M. Schoenauer. Evolutionary algorithms for constrained parameter optimization problems. *Evolutionary Computation*, 4(1):1–32, 1996. pages 70
- [145] MathWorks, Natick, MA, USA. *Global Optimization Toolbox*, 2021. pages 70
- [146] C. A. Coello, G. B. Lamont, and D. A. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer, Cham, Switzerland, 2nd edition, 2007. pages 70
- [147] D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. In *Foundations of Genetic Algorithms*, University of Illinois at Urbana-Champaign, 1991. pages 72

- [148] A. Brindle. *Genetic Algorithms for Function Optimization*. PhD thesis, Department of Computer Science, University of Alberta, 1981. pages 72
- [149] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2nd edition, 2003. pages 72
- [150] D. Whitley. A genetic algorithm tutorial. *Statistical Computing*, 4(2):65–85, 1994. pages 72
- [151] The MathWorks, Inc., Natick, MA, USA. *Parallel Computing Toolbox*, 2024. pages 74
- [152] MathWorks, Natick, MA, USA. *MATLAB*, 2024. [Online]. Available: <https://www.mathworks.com>. pages 74
- [153] H. Liebgott, A. Rodriguez-Molares, F. Cervenansky, J. A. Jensen, and O. Bernard. Plane-wave imaging challenge in medical ultrasound. In *Proc. IEEE Int. Ultrason. Symp. (IUS)*, pages 1–4, September 2016. pages 75
- [154] The MathWorks, Inc., Natick, MA, USA. *Signal Processing Toolbox User's Guide*, 2024. [Online]. Available: <https://www.mathworks.com/help/signal/ref/findpeaks.html>. pages 75
- [155] C. A. C. Coello, G. T. Pulido, and M. S. Lechuga. Handling multiple objectives with particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, 8(3):256–279, June 2004. pages 78
- [156] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *Parallel Problem Solving from Nature PPSN VI*, volume 1917. 2002. pages 78
- [157] Zahraa Alzein, Marco Crocco, D. D. Caviglia, and Hervé Liebgott. Multi-objective optimization of virtual source distributions for ultrafast diverging wave imaging. *Ultrasonics*, 2025. Submitted. pages 83
- [158] R. G. Pridham and R. A. Mucci. Digital interpolation beamforming for low-pass and bandpass signals. *Proceedings of the IEEE*, 67(6):904–919, June 1979. pages 87
- [159] K. Ranganathan, B. C. Byram, J. J. Dahl, and G. E. Trahey. Direct sampled i/q beamforming for compact and very low-cost ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 51(9):1082–1094, September 2004. pages 87
- [160] M. Kim, J. Park, T. Song, and C. Kim. Computationally efficient architecture for ultrasound phase rotation beamforming. *Electronics Letters*, 52(1):18–19, 2016. pages 87
- [161] J. Ma, J. Luo, C. Yu, H. Liu, and J. Liu. Ultrasound phase rotation beamforming on multi-core dsp. *Ultrasonics*, 54(1):99–105, 2014. pages 88

- [162] D. M. Becker, C. A. Tafoya, S. L. Becker, G. H. Kruger, and M. J. Tafoya. The use of portable ultrasound devices in low- and middle-income countries: A systematic review of the literature. *Tropical Medicine & International Health*, 21(3):294–311, 2016. pages 88
- [163] B. Y. S. Yiu, I. K. H. Tsang, and A. C. H. Yu. Gpu-based beamformer: Fast realization of plane wave compounding and synthetic aperture imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 58(8):1698–1705, August 2011. pages 88
- [164] O. C. Eidheim, J. Skjermo, and L. Aurdal. Real-time analysis of ultrasound images using gpu. In *International Congress Series*, volume 1281, pages 284–289, May 2005. pages 88
- [165] J. Liu, Y. Zhang, J. Liu, and X. Cheng. A review on real-time 3d ultrasound imaging technology. *Journal of Healthcare Engineering*, 2018:1–13, 2018. Article ID 3734568. pages 99
- [166] S. W. Hughes, D. Jones, C. M. Gutteridge, M. A. Cox, D. J. Griffiths, and E. Berry. Volume estimation from multiplanar 2d ultrasound images using a remote electromagnetic position and orientation sensor. *Ultrasound in Medicine & Biology*, 22:561–572, 1996. pages 100
- [167] W. E. Moritz, M. Berger, J. G. Miller, and B. E. Sobel. An ultrasonic technique for imaging the ventricle in three dimensions and calculating its volume. *IEEE Transactions on Biomedical Engineering*, BME-30:482–492, 1983. pages 100
- [168] V. T. Rathod. A review of electric impedance matching techniques for piezoelectric sensors, actuators and transducers. *Electronics*, 8(2):169, February 2019. pages 100
- [169] R. Wodnicki, H. Kang, D. Li, D. N. Stephens, H. Jung, Y. Sun, R. Chen, L.-M. Jiang, N. E. Cabrera-Munoz, J. Foiret, Q. Zhou, and K. W. Ferrara. Highly integrated multiplexing and buffering electronics for large aperture ultrasonic arrays. *BME Frontiers*, 2022:9870386, January 2022. pages 100
- [170] T. R. Gururaja. Composite piezoelectric transducer arrays with improved acoustical and electrical impedance. US Patent 6,225,728, May 2001. pages 100
- [171] K. Hynynen and J. Song. Ultrasound transducer and method for making the same. US Patent 9,327,317, May 2016. pages 100
- [172] J. D. Larson III. 2-d phased array ultrasound imaging system with distributed phasing. US Patent 5,229,933, July 1993. pages 100
- [173] G. Matrone, A. S. Savoia, M. Terenzi, G. Caliano, F. Quaglia, and G. Magenes. A volumetric CMUT-based ultrasound imaging system simulator with integrated reception and μ -beamforming electronics models. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(5):792–804, 2014. pages 100

- [174] N. Wang, Y. Qiang, C. Qiu, Y. Chen, X. Wang, Y. Pan, R. Liu, W. Wu, H. Zheng, W. Qiu, and Z. Zhang. A multiplexed 32×32 2D matrix array transducer for flexible sub-aperture volumetric ultrasound imaging. *IEEE Transactions on Biomedical Engineering*, 71(3):831–840, 2024. pages 100
- [175] J. Yu, H. Yoon, Y. M. Khalifa, and S. Y. Emelianov. Design of a volumetric imaging sequence using a vantage-256 ultrasound research platform multiplexed with a 1024-element fully sampled matrix array. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 67(2):248–257, 2020. pages 100
- [176] J. A. Jensen, M. Schou, L. T. Jørgensen, B. G. Tomov, M. B. Stuart, M. S. Traberg, I. Taghavi, S. H. Øygaard, M. L. Ommen, K. Steenberg, et al. Anatomic and functional imaging using row–column arrays. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 69(10):2722–2738, 2022. pages 100
- [177] D. Turnbull and F. Foster. Beam steering with pulsed two-dimensional transducer arrays. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 38(4):320–333, 1991. pages 100
- [178] X. Li. Design of 2d sparse array transducers for anomaly detection associated with a transcranial ultrasound system. Unpublished work / technical report, 2020. pages 100
- [179] C. Papadacci, M. Pernot, M. Couade, M. Fink, and M. Tanter. High contrast ultrafast imaging of the heart. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(2):288–301, February 2014. pages 115
- [180] Y. Chen, Z. Zhuang, J. Luo, and X. Luo. Doppler and pair-wise optical flow constrained 3d motion compensation for 3d ultrasound imaging. *IEEE Transactions on Image Processing*, 32:4501–4516, 2023. pages 115
- [181] E. Roux, A. Ramalli, H. Liebgott, C. Cachard, M. C. Robini, and P. Tortoli. Wideband 2-D array design optimization with fabrication constraints for 3-D us imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 64(1):108–125, 2016. pages 116
- [182] B. Diarra, M. Robini, P. Tortoli, C. Cachard, and H. Liebgott. Design of optimal 2-d non-grid sparse arrays for medical ultrasound. *IEEE Transactions on Biomedical Engineering*, 60(11):3093–3102, 2013. pages 116
- [183] L. Ratsimandresy, P. Mauchamp, D. Dinet, N. Felix, and R. Dufait. A 3 MHz two dimensional array based on piezocomposite for medical imaging. In *2002 IEEE Ultrasonics Symposium, Proceedings*, volume 2, pages 1265–1268, Munich, Germany, 2002. IEEE. pages 116
- [184] Mickaël Tanter and Mathias Fink. Ultrafast imaging in biomedical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(1):102–119, 2014. pages 121, 123
- [185] C. Demené, M. Pernot, V. Biran, M. Alison, H. Peperstraete, A. Tissier, J.-M. Correas, M. Fink, and M. Tanter. Spatiotemporal clutter filtering of ultrafast ultrasound data highly increases doppler and fultrasound sensitivity. *IEEE Transactions on Medical Imaging*, 34(11):2271–2285, 2015. pages 121, 122

-
- [186] J. Baranger, C. Demené, E. Mace, and M. Tanter. Adaptive spatiotemporal svd clutter filtering for ultrafast doppler. *Nature Communications*, 12:158, 2021. (Functional ultrasound context; robust ultrafast Doppler processing). pages 121, 122
- [187] J. E. Volder. The cordic trigonometric computing technique. *IRE Transactions on Electronic Computers*, EC-8(3):330–334, 1959. pages 122
- [188] R. G. Pridham and J. C. Mucci. Digital interpolation beamforming for narrowband and broadband arrays. *Proceedings of the IEEE*, 67(6):904–919, 1979. pages 122
- [189] Billy Yiu and Alfred C. H. Yu. Gpu-based beamformer for real-time high-frame-rate ultrasound imaging. In *IEEE International Ultrasonics Symposium (IUS)*, pages 691–694, 2011. pages 122
- [190] B. Gasse, C. Fei, C. Bo Henriksen, et al. Supra: Open-source software defined ultrasound processing for real-time applications. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(1):61–75, 2021. pages 122
- [191] J. Provost et al. High frame-rate echocardiography: A review, 2023. arXiv preprint. pages 123
- [192] P. Santos, G. U. Haugen, L. Løvstakken, E. Samset, and J. D’hooge. Diverging wave volumetric imaging using subaperture beamforming. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 63(12):2114–2124, December 2016. pages 123
- [193] L. Castrignano, P. Tortoli, G. Matrone, M. Crocco, A. S. Savoia, and A. Ramalli. On the impact of microbeamformers in 3-d high frame rate ultrasound imaging: A simulation study. *IEEE Transactions on Biomedical Engineering*, 72(6):1941–1950, June 2025. pages 123

