Optical and single element transducers for the generation of arbitrary acoustic fields

Michael David Brown

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Medical Physics and Biomedical Engineering

University College London

April 12, 2018

I, Michael David Brown, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

Precise control over the temporal and spatial properties of acoustic fields in 2 or 3-D is essential for nearly all modern, biomedical applications of ultrasound. At present, piezoelectric arrays dominate, however, despite their ubiquity they have a number of drawbacks that compromise the fidelity with which the output field can be manipulated, particularly at high frequencies and in three dimensions. The development of new novel alternatives for manipulating acoustic fields in 3-D is therefore essential. This thesis presents several new techniques through which this can be achieved using both the optical generation of ultrasound and single element piezoelectric transducers. First, the use of multiple Q-switch laser sources in combination with binary amplitude holograms is investigated for the generation of single and multi-focal acoustic fields. The conditions required for the generation of a focus are established numerically and the method is validated experimentally. Next, two approaches are developed for the generation of arbitrary spatial distributions of pressure using a single optical pulse. The first employs multi-layer optical absorbers: structures composed of several absorbing layers each individually patterned such that the field constructively interferes at a set of target points. The second uses tailored optically absorbing surface profiles: arbitrary surface shapes, fabricated through 3-D printing, designed to geometrically focus over a continuous pattern. Finally, the last chapter of the thesis investigates the use of multi-frequency kinoforms for mapping the field of single element piezoelectric transducers onto multiple complex target distributions. The properties of these kinoforms are explored in depth numerically and experimentally it is shown that multiple complex distributions can be generated in a target plane using this approach.

Impact statement

This thesis is comprised of a set of numerical and experimental studies exploring different techniques for generating arbitrary spatial distributions of pressure in 3-D. This could provide a direct benefit to researchers working across a range of areas in biomedical and physical acoustics. It is first shown that multiple Q-switch laser pulses can be combined with binary amplitude holograms to create complex steady state acoustic fields. This could be useful for researchers looking to create a system for dynamically manipulating biological particles to investigate cell properties. Two different approaches are then demonstrated for generating complex acoustic fields using a single optical pulse. The first uses multi-layer patterned absorbers, the second uses tailored optically absorbing surface profiles. Both methods could be used to create customised transducers for different biomedical and industrial applications. For example, focused transducers correcting for wave-front aberrations introduced by soft-tissues could be fabricated for targeted drug delivery. Alternatively, transducers conforming to particular mechanical components could be designed for industrial testing. Investigations into the design of multi-frequency kinoforms for mapping the output of single element transducers onto multiple target distributions are then described. This could be used to correct for skull-aberrations in transcranial neuromodulation or for the assembly or propulsion of synthetic objects. The thesis then concludes with an appendix investigating damage mechanisms that can occur in the fibre coupling of high-energy high peak-power optical pulses. The results of this study would be useful to anyone looking to stably couple high energy Q-switch lasers.

Acknowledgements

First and foremost, I would like to thank my supervisors Bradley Treeby and Ben Cox for their continual support throughout my PhD. Without their enthusiasm, patience, and sensible (and occasionally not so sensible) ideas the completion of this thesis would not have been possible.

I would also like to thank the rest of the Biomedical Ultrasound group for providing a pleasant working environment, helpful discussions, and excellent home cooking. Particular thanks to Louis Robertson who was undoubtedly a worse influence on me than I was on him.

Within the wider Medical Physics department I am grateful to a number of people for their assistance during the project. Daniel Nikitichev and Eve Hatton who helped with the design and fabrication of the 3-D printed components, James Guggenheim, Edward Zhang, Robert Ellwood, Olumide Ogunlade, and Thomas Allen who helped with the operation of the Fabry-Perot sensor and the Q-switch sources, and Elly Martin who assisted with several field characterisations.

Thanks also to my friends for their moral support. Particularly Lu An and Martina Fonseca for their unwavering positivity, Daniel Elson (and Lu An again) for being excellent housemates, and, Tom Walden, Holly Tetlow, Adam Bailey, and Paul Judge for various food and drink based excursions.

Finally, thanks to my Dad for providing a sensible perspective throughout the 4 years.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Since the initial demonstrations for diagnosis [1, 2] and therapy [3] in the 1940's and 50's, the use of ultrasound in biomedicine has undergone continual growth. Today, the most significant application is for diagnostic imaging [4], however, ultrasound is also used to heat or ablate tissues [5], apply mechanical forces for the manipulation of particles [6], and more recently for the stimulation of neurological structures [7]. For each of these applications, the ability to precisely control the spatial and spectral properties of the acoustic fields in two and three dimensions is essential. However, achieving the requisite control with existing technology can be a challenge. If a transducer with the ability to arbitrarily vary the spatial distribution of acoustic pressure generated in 3-D could be developed, it would enable the realisation of new biomedical applications for ultrasound.

In therapeutic ultrasound, fields with multiple foci can enable more precise heat deposition, improving treatment outcomes [8, 9] and reducing therapy time [10, 11]. The field can also be tailored to avoid certain highly absorbing areas such as the ribs or plaques to eliminate unwanted tissue damage [12, 13]. In ultrasound stimulation, patterned distributions of pressure have been proposed for stimulating the retina [14] and multi-focal fields would enable the simultaneous targeting of different brain structures [15]. Similarly, for ultrasound mediated drug delivery, the ability to precisely control the pressure distribution and correct for tissue aberrations would enable far more accurate targeting of biological structures [16]. For particle manipulation, it has already been shown that with full control over both the

phase and amplitude of the pressure in 3-D, particles can be forced into complex arrangements [17] as well as stably levitated, translated [18], and rotated [19]. This can be used to sort and mix cells [20], or alternatively to measure their mechanical properties [21]. Outside of biomedicine, arbitrary control over distributions of pressure has also shown potential for haptics for both 2-D displays [22] and 3-D virtual reality systems [23].

1.1 Current systems

At present, nearly all clinical and industrial ultrasound systems are constructed from piezoelectric materials. These are materials which respond to the application of a potential difference with the generation of mechanical stress and vice versa enabling both the transmission and reception of ultrasound. For most modern systems, these materials are diced into separate elements and arranged in 1-D, 2-D or annular configurations to form arrays [24]. Despite their ubiquity these arrays do have several drawbacks.

In order to manipulate the properties of the acoustic field in 3-D, a transducer is required for which the generated pressure can be spatially and temporally controlled across a 2-D aperture. This requires dense 2-D arrays of separately addressable piezoelectric elements. One major constraint for these arrays is that the element pitch in each direction should be less than $\frac{\lambda}{2}$ (where λ is the shortest operational wavelength of the transducer) to avoid the generation of grating lobes [25], intense side beams generated by constructive interference in periodic arrays. In practice the spacing used is usually slightly larger than $\frac{\lambda}{2}$, however, the element size and separation in each dimension is still limited to under 100 μm for frequencies of 10 MHz and above [26]. Consequently, arrays with even a moderate aperture (1.5×1.5 cm) operating at this frequency require extremely high numbers (22500) of separate elements. Fabricating functional arrays possessing these numbers of elements poses several challenges. Principally, the number of channels or electronic driving signals available to most commercial systems is limited to 256, so driving each element conventionally isn't possible. Additionally, assuming enough channels were avail-

able creating electrical contacts to each element would lead to excessively bulky probes and heating problems [27].

Despite these challenges, several large fully addressed 2-D systems ($\sim$9000 elements) have been developed commercially for volumetric imaging [28]. These work by performing part of the signal processing using application specific integrated circuits (ASIC) attached directly to the array elements. This enables the number of independent channels to be kept low ($\sim$128) while still driving each element [29]. The flexibility of driving signals for each element depends on the integrated circuit. For one approach [30], the transducer aperture is divided into a number of sub-apertures each associated with a set number of system channels. The elements in each sub-aperture can be connected to any of the associated channels, which allows each element to be driven from a small set of different signals. These systems do possess certain drawbacks: they are extremely expensive to purchase and develop [31], are subject to discretisation errors introduced by sub-aperture processing [32], and at present the operating frequencies are still limited to less than 10 MHz.

Sparse or under-sampled 2-D arrays utilising random or irregular arrangements of elements are often considered as an alternative [33]. The appearance of grating lobes is related to the periodicity of the array, so their amplitude can be minimised by either randomising, or directly optimising the element distribution [25]. However, the fabrication of these arrays can be more difficult, and it is still possible for grating lobes or clutter to be generated [34]. In addition to these issues, piezoelectric materials have certain inherent limitations [35]. Elements are resonant, constraining the operating bandwidths, and fabricating arrays for high frequencies is difficult, even in 1-D. This is due to the challenges involved in dicing piezoelectric materials to micron-sized elements, creating contacts with such small elements, and eliminating crosstalk [36].

Capacitive micro-machined ultrasound transducers (CMUT's) are one alternative to piezoelectric arrays. These are easier to fabricate, have wider bandwidths, and the integration of electronics is significantly simpler. The manufacture of large

(128×128) 2-D arrays has already been demonstrated [37]. However, at present CMUT arrays suffer from issues with acoustic cross-talk, and dielectric breakdown [38], and at present the largest fully-addressed arrays are smaller than the piezoelectric equivalents [39].

1.2 Project goal

Given the expense associated with traditional piezoelectric arrays and their inherent limitations, the development of new novel alternatives for the manipulation of acoustic fields in 3-D is necessary. Ideally, these methods should scale easily either to large apertures or to higher frequencies in order to overcome the drawbacks of the existing technologies. Two promising modalities for achieving this are (1) the optical generation of ultrasound via the photoacoustic effect, which occurs when a time-varying optical source is incident on an optically absorbing medium, and (2) the use of large single element piezoelectric transducers combined with structures that modulate the output field. The goal of this thesis is to investigate different approaches within these two paradigms for generating arbitrary 3-D distributions of acoustic pressure.

Chapter 2

Background

In this Chapter, the relevant literature for this thesis is summarised. In Sec. 2.1, different methods for focusing broadband and continuous wave acoustic fields in 3-D are reviewed. In Sec. 2.2, the optical generation of ultrasound is introduced. This is followed in Sec. 2.3 by a review of the previous experimental approaches for temporal and spatial modulation of optically generated acoustic fields. Section 2.4 then briefly summarises the previous methods for modulating single element piezoelectric transducer fields, and the Chapter concludes with a description of the thesis structure in Sec. 2.5.

2.1 Focusing of acoustic fields

The generation of arbitrary acoustic fields in 3-D can be divided into 2 problems. The first is how to control the temporal and spatial properties of the pressure $p(x,y,z,t)|_{source}$ generated across two or in some cases three dimensions. The second is with the ability to manipulate the time varying pressure, how can a desired distribution of acoustic pressure $p(x,y,z,t)$ be realised. This is solved in different ways depending on the properties of the source and the target properties of the field.

2.1.1 Geometric focusing and lenses

The two simplest methods for generating a focus is either to focus geometrically using a concave transducer or to use an acoustic lens. For geometric focusing, the transducer surface is shaped such that the distance from each point on the surface to a target focal point is identical, which results in strong constructive interference

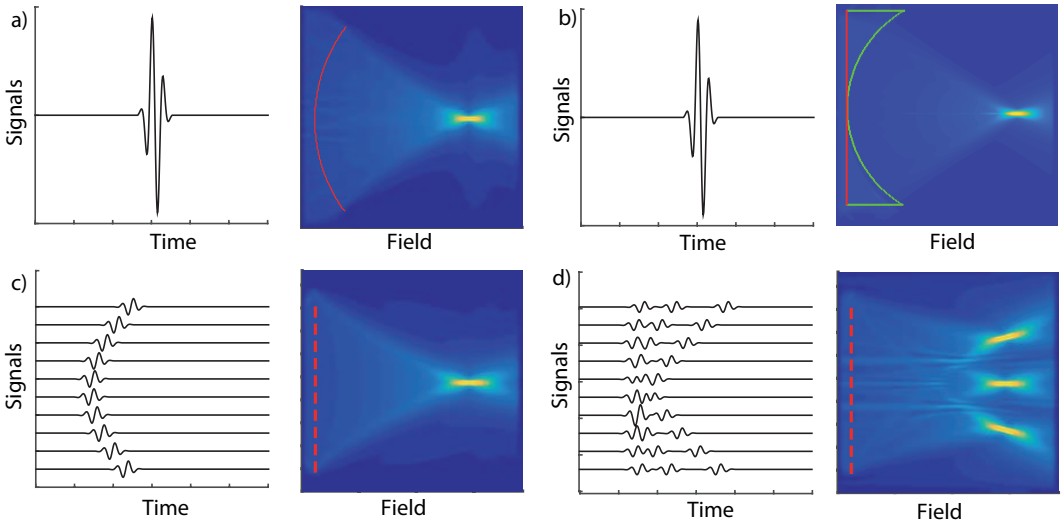

Figure 2.1: Illustrative figure of different methods for focusing: (a) geometric focusing, (b) acoustic lens, (c) electronic focusing, and (d) electronic focusing for multiple foci. For the acoustic lens, the green line indicates the physical structure of the lens.

at that point when the transducer is driven (Fig. 2.1 (a)) [40]. With a lens, a shaped profile with a different sound speed to the coupled medium is attached to a flat transducer surface. The thickness of the lens across the surface is varied such that the sound is focused to a single point [41]. Both methods have the advantage that only a single driving signal is necessary. However, the focal position is fixed and generalising to multiple foci or arbitrary fields isn't possible.

2.1.2 Electronic focusing

A more flexible approach is to use electronic focusing with a transducer array. This is similar to geometric focusing except rather than physically shaping the source or using a lens, the driving signal to each element of the array is delayed so that the signal generated by each arrives simultaneously at a target focal point. This results in constructive interference and the generation of a focus (Fig. 2.1 (c)) [42]. For a homogeneous medium, to generate a focus at a position (x,y,z) using a set of elements positioned at (x_i,y_i,z_i), where i is the element index, the appropriate delay Δt_i for each element can be calculated using

$$\Delta t_i = -\frac{r_i}{c_0} + \max_i\left(\frac{r_i}{c_0}\right). \tag{2.1}$$

Here c_0 is the sound speed in the medium, r_i is the distance between each transducer element and the focus, and $\max_i(r_i/c_0)$ is the longest travel time to the focus from any element.

Electronic focusing has the advantage that the focal position can be freely steered both laterally and axially by altering the delays. It can also be generalised to the simultaneous generation of several foci by summing the driving signals for individual foci (Fig. 2.1 (d)). However, for pulsed signals, this results in different time-varying signals $p_i(t)$ for each element across the array. As discussed in Sec. 1.1, existing 2-D systems rely on a large degree of element multiplexing so the aperture for which arbitrary focusing of broadband signals can be achieved (conventionally) is extremely limited.

2.1.3 Time reversal

Applying either of the previous methods to focus in inhomogeneous media results in aberrations. These aberrations occur as portions of the emitted sound encounter regions with differing sound speed which distort the shape of the wavefront. When focusing through the skull or ribs, these can be particularly severe due to the large contrast in sound speed between bone and tissue. To correct for aberrations, time reversal focusing can be used. Time reversal was first demonstrated by Fink in 1992 [43]. It is based on the observation that in a non-dissipative medium the acoustic wave equation is invariant under time reversal. This means that for a diverging burst of sound $\psi(r,t)$ emitted by a source and reflected and refracted by a inhomogeneous medium, there exists a sets of waves $\psi(r,-t)$ that retrace the original paths and converge on the original source [44].

To take advantage of this property, an acoustic source is placed at a desired focal point, driven, and the resulting field $p_i(t)$ is sampled across a set of transducers. The recorded signals are then reversed and retransmitted from each of the transducers. This creates a wave that converges on the original source, creating a focus and compensating for the inhomogeneities in the medium. If the properties of the medium are known in advance then the first step of this process can be done numerically by simulating propagation through the medium with a full-wave model

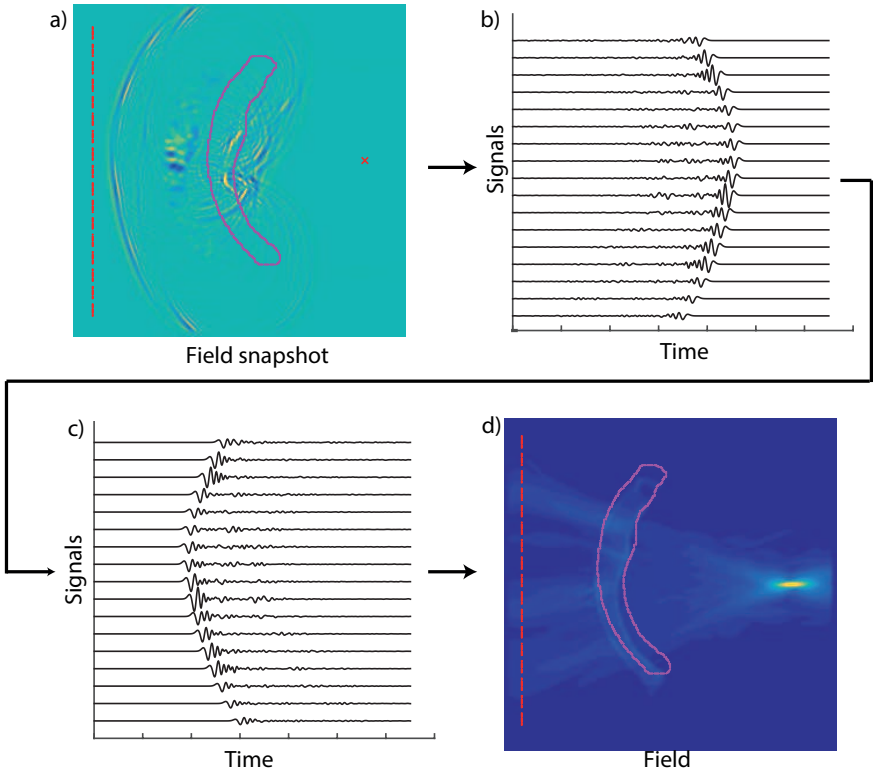

Figure 2.2: Illustrative figure of time reversal focusing. (a) Propagation through the medium from the target point is simulated or measured. The red cross indicates the target point. The purple outline indicates the structure of a region with significantly higher sound speed. Dashed points represent the transducer positions. (b) Signals are recorded at a set of transducer locations. (c) Recorded signals are reversed in time. (d) Signals are retransmitted creating a wavefront that refocuses on the original target point.

[45]. This can be used to derive driving signals for a physical transducer. This is illustrated by Fig. 2.2.

2.1.4 Holography

Each of the approaches described in Sections 2.1.1-2.1.3 can be used for either pulsed or continuous wave sources. In the case of continuous wave, the calculated time delays for electronic focusing can be converted to phase offsets $\Delta\phi_i$. Similarly, for time reversal, phase and amplitude values that correct for the aberrations can be extracted from the steady state signals [45]. However, for solving the general problem of how to generate a desired field using an acoustic source, of frequency f, with different (amplitude or phase) constraints the best approach is holography.

Holography originates in the field of optics with the work of Dennis Gabor

in 1948 [46]. He established that it is possible to obtain three dimensional reconstructions of a optical wave-field if information on both the phase and amplitude distribution can be captured across a surface normal to the wave propagation [47]. In optics where only the intensity can be directly recorded, this is achieved by interfering the wave-field scattered by an object with a coherent reference source and storing the resulting intensity as a transparency (the hologram). When this transparency is illuminated by a beam matching the original reference beam, a real and virtual reconstruction of the original object is generated.

While the initial demonstration of holography was in optics, holography has also been adopted in acoustics. Acoustic frequencies are significantly lower, so direct measurement of both the phase and amplitude of the wave-field is possible and a reference source is not required. If the time varying pressure generated by a source can be recorded across a 2-D plane then these measurements can be used to numerically reconstruct the full 3-D wave-field [48] for both steady state and broadband signals [49].

2.1.5 Computer generated holography

For manipulating acoustic fields, the inverse problem is of interest. That is, starting from a desired 3-D distribution of pressure $p(x,y,z)$ and possessing the ability to control properties of the source such as the amplitude $A(x,y)$ or phase $\phi(x,y)$ across a 2-D plane, how to derive a complex distribution $A(x,y)e^{i\phi(x,y)}$ (or hologram) that will best generate the target field. In optics, the development of algorithms that address this inverse problem is known computer generated holography (CGH).

2.1.5.1 Forward models

To solve the holography inverse problem, models that can calculate the 3-D distribution of acoustic pressure generated by a given hologram distribution are necessary. In acoustics, due to the length scales involved, two models are most commonly used: the angular spectrum approach (ASA) and the Rayleigh-Sommerfeld integral.

The ASA was originally developed in optics [50] for the rapid evaluation of the field generated by a known planar source distribution across other parallel planes.

Assuming the complex distribution of the pressure $p(x,y,z_0)$ is known across some plane z_0, the pressure at a depth z is evaluated using the ASA by first taking the spatial Fourier transform of the pressure using

$$P(k_x,k_y,z_0) = \int\int p(x,y,z_0)e^{jk_x x + jk_y y}dxdy. \tag{2.2}$$

Here k_x and k_y are the spatial frequencies in x and y respectively. This decomposes the field into a set of plane waves propagating at different angles defined by the spatial frequencies. Each of these plane waves is then multiplied by a propagator term $H(k_x,k_y,\Delta z)$ which accounts for the phase change introduced by propagation to the new target depth. This is given by

$$H(k_x,k_y,\Delta z) = \begin{cases} e^{(-i\Delta z\sqrt{(k^2-k_x^2-k_y^2)})}, & \text{if } k_x^2+k_y^2 < k^2. \\ e^{(-\Delta z\sqrt{(k_x^2+k_y^2-k^2)})}, & \text{if } k_x^2+k_y^2 > k^2. \end{cases} \tag{2.3}$$

Here k is the wavenumber, and $\Delta z = (z - z_0)$. Where $k_x^2 + k_y^2 < k^2$ the waves are real and propagate, and where $k_x^2 + k_y^2 > k^2$, the waves are evanescent so decay exponentially. After applying this propagator, the field across z is reconstructed using the inverse spatial Fourier transform. In order to ensure numerical accuracy, there are additional requirements. When the field is back-propagated (i.e., $\Delta z < 0$), evanescent waves are amplified exponentially so need to be cutoff. Additionally, the ASA is subject to undersampling for higher spatial frequencies which introduces aliasing. To eliminate this, it is necessary either to significantly extend the source plane or to restrict the spectrum that is propagated [51].

The ASA is useful for problems where the target distribution of pressure and source occupy 2 parallel planes, as the field can be rapidly propagated between the two using Fourier transforms. However, if the source is non-planar or the target distribution consists of a few discrete points, it becomes less efficient to propagate the whole field. Where this is the case, it is more convenient to use the Rayleigh-Sommerfeld integral. For a planar source $p(x',y',0)$ mounted in an infinite baffle,

this gives the pressure at a point elsewhere in the field as [47]

$$p(x,y,z) = \int_S p(x',y') \frac{z}{2\pi r} \left(\frac{ik}{r} + \frac{1}{r^2} \right) e^{ikr} dS. \tag{2.4}$$

Here dS is a area element, S defines the area over which the source is non-zero, and r is the distance between the $(x',y',0)$ and (x,y,z) coordinates given by $r = \sqrt{(x-x')^2 + (y-y')^2 + z^2}$. When numerically accurate, the Rayleigh-Sommerfeld and ASA are equivalent [50, 52], so the appropriate forward model depends on the problem geometry.

2.1.5.2 Complex modulation

The difficulty of calculating a computer generated hologram depends on the properties of the transducer. In the ideal case of a dense 2-D array for which both the amplitude and phase can be freely modulated, the problem can be simple [53, 54]. For example, for a planar target distribution $A_t(x,y)$ at a depth z, the field can be back-projected to the hologram plane using the ASA. Using the phase and amplitude values that result from this back-projection as driving signals generates a good realisation of the target distribution. This is illustrated in Fig. 2.3. For non-planar distributions, such as a set of arbitrarily placed focal points r_j, the phase conjugation approach described by Ibbini can be used [55]. This uses Eq. 2.4 to calculate the complex amplitude $A_{ij}e^{\phi_{ij}}$ generated at each target point r_j by each transducer position r_i. The sum of the complex values is then taken for each transducer element to give a single phase and amplitude $p_i = A_i e^{\phi_i}$, and the complex conjugate of this value is used as a driving signal.

For many practical systems, however, it is only possible to control either the amplitude or the phase. Moreover, the values these can occupy are usually quantised, in the extreme case to binary values. Where the source has these constraints, the simple back-propagation approach no longer yields acceptable results. This is illustrated by Fig. 2.4, which shows the pressure distribution generated for the same problem as Fig. 2.3 with different constraints applied to the source after back-propagation. From the top row to the bottom it shows: phase and amplitude modula-

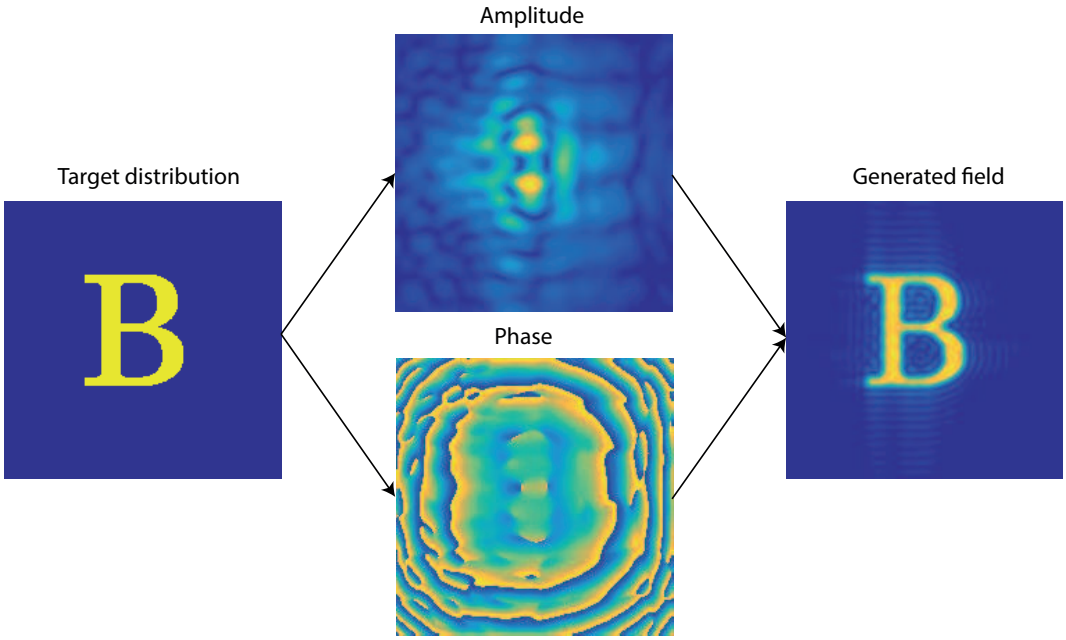

Figure 2.3: Generation of an arbitrary pressure distribution by using back-propagation. First column is the target distribution which resembles a B. The ASA approach is used to back-propagate to the target plane giving the phase and amplitude values shown in the middle column. When these values are forward propagated with the ASA approach the distribution on the right is generated at the target depth. This closely matches the target distribution

tion, only phase modulation, binary phase modulation, only amplitude modulation, and binary amplitude modulation. In each case, the generated field is significantly distorted compared to the ideal case shown in Fig. 2.3. To design holograms for sources possessing these limitations, a wide variety of algorithms have been developed. Two popular approaches are iterative Fourier transform algorithms, and direct binary search algorithms.

2.1.5.3 Iterative Fourier transform algorithms

Iterative Fourier transform algorithm (IFTA) is a general term for a group of algorithms that originated with the work of Gerchberg and Saxton in 1972 [56]. They addressed the problem of retrieving the phase distribution across two planes, the fields of which are related by a Fourier transform, from amplitude measurements on both planes. This is a relevant problem for several fields such as astronomy and electron microscopy. These methods have also been applied for the calculation of

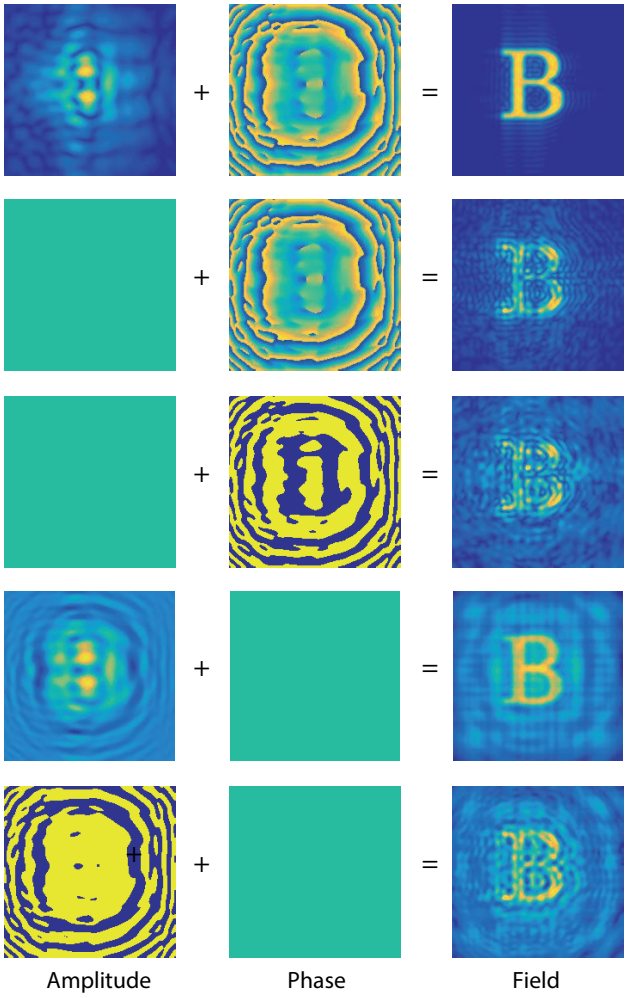

Figure 2.4: Attempting to generate the same arbitrary pressure distribution as Fig. 2.3 with different constraints on the ability to manipulate the phase and amplitude. From top to bottom the constraints applied are: None, only phase modulation, binary phase modulation thresholded to 0 and π, only amplitude modulation calculated taking the real component of the complex modulation, binary amplitude modulation calculated by thresholding the real component about 0.

phase only holograms $e^{\phi_h(x,y)}$ that generate target planar intensity (or pressure) distributions $A_t(x,y)$ [57]. The basic algorithm is illustrated in Fig. 2.5. It starts by imposing a starting phase distribution (usually random) onto the target amplitude distribution $A_t(x,y)$. This field is then propagated to the hologram plane where the phase $\phi_h(x,y)$ is kept and the hologram amplitude $A_h(x,y)$ is replaced with a known distribution (i.e., the spatial profile of the beam). The new field is then propagated back to the target plane where the phase is retained and the amplitude is replaced by $A_t(x,y)$. The process is iterated until the algorithm stagnates or the error, eval-

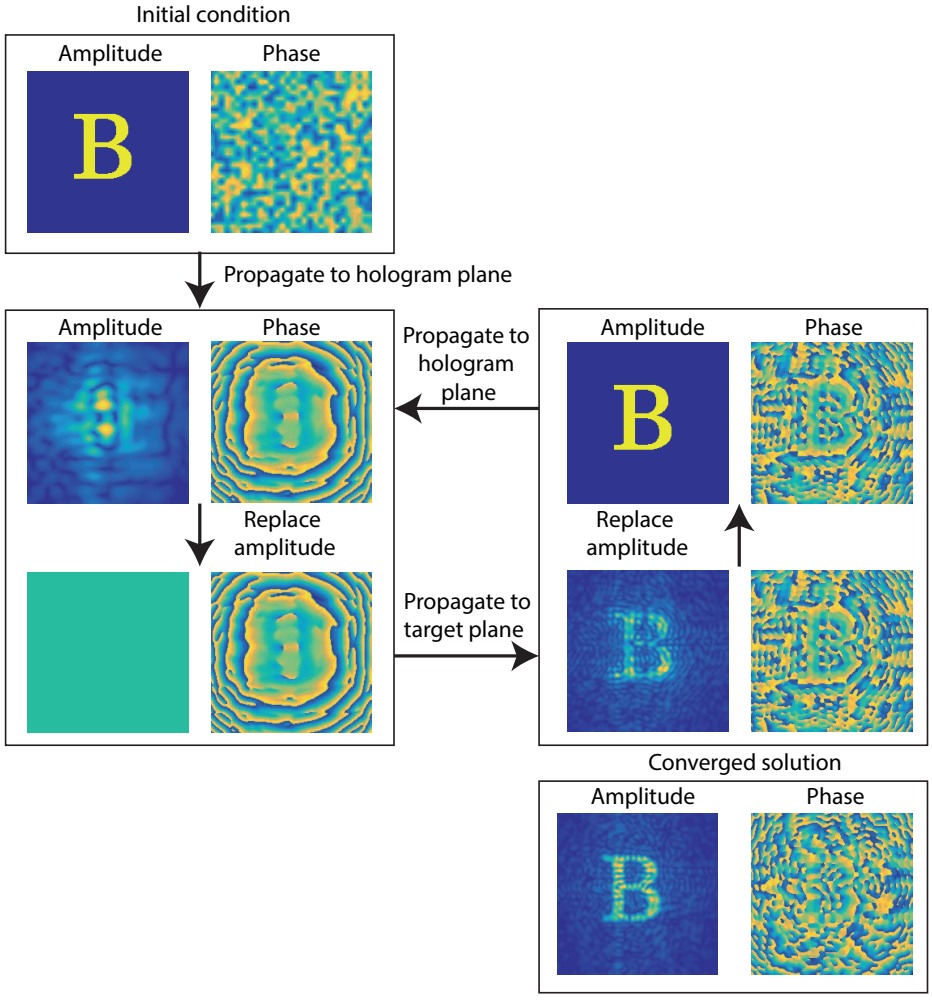

Figure 2.5: Illustration of iterative Fourier transform algorithm for the same problem as Fig. 2.3. The converged solution is shown in the bottom right. The improvement compared to the solution from the initial iteration (directly above) can clearly be seen.

uated by comparison with the target amplitude, falls below a threshold. For planes related by a Fourier transform, the amplitude error in the reconstruction decreases or remains constant with each iteration due to the unitary property of the Fourier transform [56].

The model relating the two planes depends on their positioning. In the original implementations of the IFTA, holograms were designed to generate target distributions in the far-field or in the focal plane of a lens [58], where the fields are related by a Fourier transform. However, they have also been designed to operate in the Fresnel region where planes are related by the Fresnel transform [59], and in the near-field where they are related by the ASA [60].

One drawback of IFTA is the difficulty of calculating quantised phase holograms. If the phase distribution $\phi_h(x,y)$ is quantised, in each iteration errors are introduced and the algorithm is prone to getting stuck in poor local minima [61]. To avoid this, the phase can be quantised gradually, either by defining limits around the discrete levels for which the values are quantised that gradually expand, or by slowly decreasing the number of available phase levels as the iteration number increases [62].

2.1.5.4 Direct search algorithms

An alternative to IFTA are a group of algorithms that directly perturb the hologram amplitude $A_h(x,y)$ or phase $\phi_h(x,y)$ in order to find a solution that best approximates a target amplitude. These are used most often to calculate holograms for which the available modulation is highly quantised (in most cases binary). The most basic version of this approach has been dubbed direct binary search [63], and is illustrated in Fig. 2.6 for the design of a binary amplitude hologram.

The hologram is initialised in an arbitrary binary state $A_h(x,y)$ and a target amplitude distribution $A_t(x,y,z)$ is defined. The amplitude $A(x,y,z)$ generated by the current hologram state is calculated using a forward model. The quality of this solution is then evaluated by comparing the current amplitude to the target distribution using a cost function, for example

$$C = \sum_x \sum_y \sum_z |A_t(x,y,z) - A(x,y,z)|. \tag{2.5}$$

Here the value of C is the 'cost' or current solution quality, where lower is better. From this initial condition, individual pixels on the hologram aperture are selected, the amplitude is perturbed, and the change in C is evaluated. If C is decreased the change is kept, otherwise it is reversed. This process continues until no further changes are accepted or a stopping criterion is fulfilled. One disadvantage of the direct binary search approach is that the algorithm is greedy. Since only positive changes are accepted at each step, it is inevitably trapped in the first minima it encounters. This is defined by the initial hologram state and the pixel selection

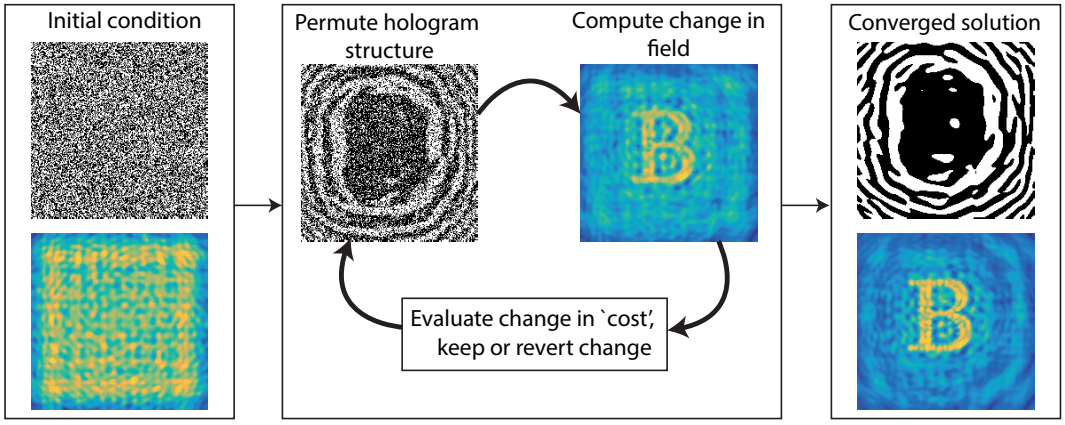

Figure 2.6: Illustration of direct search algorithm for designing a binary amplitude hologram for the same design problem as Fig. 2.4. The converged solution in this case can be compared to the final row of Fig. 2.4 the improvement can again be clearly seen.

criteria. If the solution space contains large numbers of bad solutions, then the algorithm will perform poorly.

To overcome this drawback several works have proposed simulated annealing [64, 65, 66]. This introduces a probability of keeping bad changes which increase the cost into the optimisation. This is done by defining a temperature T associated with the optimisation. Each time a pixel is changed, as before if C decreases it is kept, however, if C increases then the change can still be kept with a probability calculated based on the current value of T and the magnitude of the change ΔC. This probability P is usually calculated using

$$P = \exp\left(-\frac{\Delta C}{T}\right). \tag{2.6}$$

As the optimisation proceeds, the temperature is gradually decreased. This reduces the probability of accepting bad changes until the algorithm effectively becomes direct binary search. Adopting this method gives the optimisation a chance of escaping very poor solutions reached early in the optimisation. The actual convergence will be determined by the choice of initial temperature and the cooling rate.

Another proposed alternative are genetic algorithms [67, 68], which can in principle converge to the global minimum. These work by starting with a population of initial solutions $H_n(x,y)$. The cost C_n of each is then evaluated and the solu-

tions are 'bred' and 'mutated'. Breeding involves creating new solutions $H_{new}(x,y)$ by combining 2 or more existing solutions. The solutions that are bred are selected with a probability based on the cost. Mutation involves perturbing a random set of pixels on each solution. After these steps the cost of each solution is re-evaluated and a percentage with the greatest cost are eliminated. The process is repeated through hundreds of generations until a stopping criterion, based either on stagnation or a target solution quality, is reached. Having a large family of solutions and introducing random mutations allows the optimisation to avoid stagnating in local minima. It is, however, extremely sensitive to tuning the various parameters. Additionally, genetic algorithms are computationally intensive as the acoustic field of a large family of holograms needs to be repeatedly evaluated.

In comparison to IFTA algorithms, direct search approaches have several advantages. For the design of binary amplitude or binary phase holograms they can perform better as the quantisation is directly included in the model. Furthermore, the choice of forward model is arbitrary, so either the target distribution or the hologram distribution can be three dimensional without additional complexity. The disadvantage is that these involve significantly more computational effort. A typical IFTA algorithms requires $\sim$50-100 iterations to converge, while calculating a 128×128 pixel hologram with a direct search method requires hundreds of thousands of operations due to the nature of the algorithm.

2.1.6 Summary

There are a wide range of approaches that can be used to generate arbitrary focused distributions of pressure in 3-D. The ideal method to use for a given transducer will depend on the desired properties of the field (i.e., steady state or broadband), the properties of the transducer, and the trade off between solution quality and computational effort.

2.2 Optical generation of ultrasound

The optical generation of ultrasound was discovered in 1880 by Alexander Graham Bell, who observed that rapidly chopped sunlight incident on a thin disk produced

an audible sound [69]. However, it wasn't until the development of the laser in the 1960's that the modality began to be investigated, initially for industrial applications [70]. There are a range of mechanisms through which incident light can be converted into elastic waves, including dielectric breakdown, ablation, thermally-induced pressure rise, electrostriction, and radiation pressure. However, the first two mechanisms result in damage to the absorbing medium while the last two mechanisms in most cases result in low acoustic pressures [71, 72]. Therefore, for this thesis, only the optical generation of ultrasound through a thermally-induced pressure rise will be considered and this mechanism will subsequently be referred to as the photoacoustic effect.

2.2.1 Photoacoustic wave equation

The photoacoustic effect occurs when a time-varying optical source is incident on an optically absorbing medium. The incident photons are absorbed, promoting electrons in the medium to a higher energy state. From there, the electrons decay on a very short time scale through variety of pathways. A percentage decay non-radiatively converting the absorbed energy to heat. This causes a localised temperature increase, and, if the heating is sufficiently rapid, a simultaneous pressure increase. This pressure increase subsequently relaxes generating an acoustic pulse [73].

The acoustic field $p(r,t)$ generated by an optical pulse incident on an optically absorbing medium can be modelled using two equations

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right) p(r,t) = -\frac{\beta}{\kappa c^2}\frac{\partial^2 T(r,t)}{\partial t^2}, \tag{2.7}$$

and

$$\rho C_v \frac{\partial T(r,t)}{\partial t} = H(r,t) + \kappa \nabla^2 T(r,t). \tag{2.8}$$

Here $T(r,t)$ is the temperature rise above equilibrium, β is the thermal expansion coefficient, κ is the isothermal compressibility, c is the sound speed, ρ is the density, C_v is the constant volume specific heat capacity, and $H(r,t)$ is the spatial and temporal heat deposition. The first equation is the acoustic wave equation with a

source term based on the temperature variation, and the second equation is the temperature diffusion equation with a heating source term [74]. If the heating is driven solely by optical absorption then it can be related to the incident fluence using

$$H(r,t) = \mu_\alpha(r)\phi(r,t). \tag{2.9}$$

Here μ_α is the optical absorption coefficient which can vary spatially in the medium, and $\phi(r,t)$ is the spatial and temporal variation in the fluence (light) incident at each point in the medium.

There are two conditions that, if satisfied, allow modelling the photoacoustic generation of sound to be significantly simplified. The first is thermal confinement, which requires that the duration of the optical heating (i.e., the optical pulse length) τ_p is much less than the time taken τ_{th} for heat to diffuse across a characteristic length D_α of the heated medium, which might be a blood vessel or thin absorbing layer. The value of τ_{th} is determined by the thermal diffusivity χ and can be calculated using

$$\tau_{th} = \frac{D_\alpha^2}{4\chi}. \tag{2.10}$$

Satisfying thermal confinement implies that the temperature increases arise solely from optical heating. This allows the second term in Eq. 2.8 to be neglected [75] leading to

$$\frac{\partial T(r,t)}{\partial t} = \frac{1}{\rho C_v} H(r,t). \tag{2.11}$$

This can be substituted into Eq. 2.7, giving the following differential equation for the acoustic pressure

$$\left(c^2\nabla^2 - \frac{\partial^2}{\partial t^2}\right) p(r,t) = \Gamma \frac{\partial H(r,t)}{\partial t} \tag{2.12}$$

Here Γ is the Gruneisen parameter given by $\Gamma = \frac{\beta}{C_v \kappa \rho} = \frac{\beta c^2}{C_p}$, and C_p is the specific heat capacity at constant pressure. Thermal confinement also allows the heating source to be split into separate spatial and temporal components, assuming the

medium is stationary [76]. In this case, the heating is given by

$$H(r,t) = H(r)f(t), \tag{2.13}$$

where $f(t)$ is the normalised temporal shape of the pulse. The second condition is stress confinement. This is satisfied if the time taken τ_s for an acoustic wave to propagate across the width of the absorbing region is much longer than the length of the illuminating pulse. This can be calculated using

$$\tau_s = \frac{D_\alpha}{c}. \tag{2.14}$$

Satisfying stress confinement implies that all the optical energy is deposited before any energy is able to propagate away from the heated region as an acoustic pulse. Where this is the case, generation is most efficient, the optical pulse can be modelled as instantaneous (i.e., a Dirac delta function), and Eq. 2.12 reduces to

$$\left(c^2\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right) p(r,t) = \Gamma H(r)\frac{\partial \delta(t)}{\partial t}. \tag{2.15}$$

This can be equivalently expressed as an initial value problem where the photoacoustic pressure can be modelled using the wave equation

$$\left(c^2\nabla^2 - \frac{\partial^2}{\partial t^2}\right) p(r,t) = 0. \tag{2.16}$$

Subject to the initial conditions that

$$p|_{t=0} = p_0(x) = \Gamma(r)\mu_\alpha(r)\phi(r) \tag{2.17}$$

and

$$\left.\frac{\partial p}{\partial t}\right|_{t=0} = 0. \tag{2.18}$$

Here $p_0(x)$ is the initial distribution of pressure after the application of the optical pulse. When the incident optical source is longer, such that stress confinement is not

satisfied (but thermal confinement still is), the resulting field is obtained by solving the above set of equations to obtain $p_\delta(r,t)$, then convolving the solution with the temporal shape $I(t)$ of the optical signal. That is

$$p(r,t) = \int_{-\infty}^{\infty} p_\delta(r,\tau) I(t-\tau) dt. \tag{2.19}$$

2.2.1.1 Solutions

Analytic solutions to the above set of equations (Eq. 2.16-2.19) have been obtained for certain simple source distributions in 1, 2 and 3-D [77, 78, 79]. In particular, Diebold showed that stress confined heating of a uniform sphere results in an 'N' shaped acoustic signal with a temporal width determined by the radius, while for a planar source, the acoustic signal is directly mapped from the heating profile. These are shown in Fig. 2.7 (a-b). Diebold also analysed the case where the incident optical signal is significantly longer than the dimensions of the absorber. He found that for a planar source, the temporal profile of the incident pulse is directly mapped onto the acoustic signal, while for a spherical source, the signal is proportional to the time derivative of the incident illumination. This is illustrated in Fig. 2.7 (c-d) which show the time series generated by a Gaussian pulse for a sphere and a plane.

In general, however, the shape of the heat deposition won't match one of these simple distributions and an analytical solution can't be found. Where this is the case, the resulting field has to be evaluated numerically. Several approaches have been developed to achieve this, for example, k-space methods [80, 81].

2.2.2 Photoacoustic transducers

The most common biomedical application for the optical generation of ultrasound is photoacoustic tomography (PAT) [82]. In PAT, nanosecond pulses of light are used to illuminate tissue. The absorption of the light results in the generation of acoustic transients via the photoacoustic effect which are detected at the tissue surface by a 2-D sensor array. These measurements are used to reconstruct a 3-D image of the initial acoustic pressure distribution, which is closely related to the inherent tissue absorption. This allows for high-resolution images with molecular contrast between

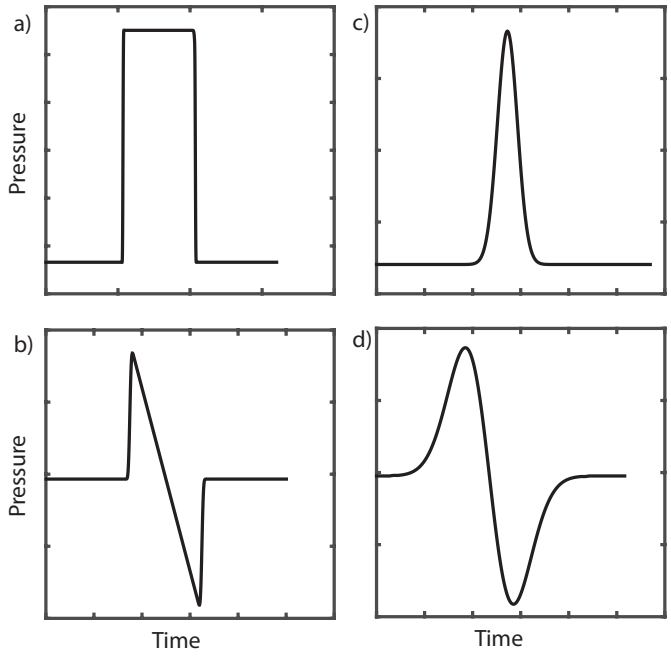

Figure 2.7: Acoustic signals generated by different photoacoustic sources. (a) Uniform impulsive heating of a plane. (b) Uniform impulsive heating of a sphere. (c) Time signal generated for the source in (a) excited by a long Gaussian pulse. (d) Time signal generated for the source in (b) excited by a long Gaussian pulse.

different tissue types to be generated, at greater depths than other optical imaging modalities.

In recent years, there has also been growing interest in using the photoacoustic effect as a source for traditional biomedical applications of ultrasound (i.e., as a transducer) [83, 84, 85]. A typical configuration for this, illustrated in Fig. 2.14, uses a thin layer of a highly optically absorbing material deposited on a transparent substrate. A short high-energy optical pulse incident on the layer generates an acoustic pulse which propagates into the medium. This has several potential advantages over piezoelectric generation: for thin layers and short pulses, extremely broadband signals can be produced [86], no electrical contacts are required, and the element size can be varied by altering the illumination.

Despite these advantages, application of the photoacoustic effect to ultrasonics was held back until the early 2000's by the poor photoacoustic conversion efficiencies and therefore pressures that are generated from most materials [87, 88, 89]. From Eq. 2.18 it can be seen that the magnitude of the optically generated pressure is determined by the fluence, the optical absorption, and the Gruneisen coefficient,

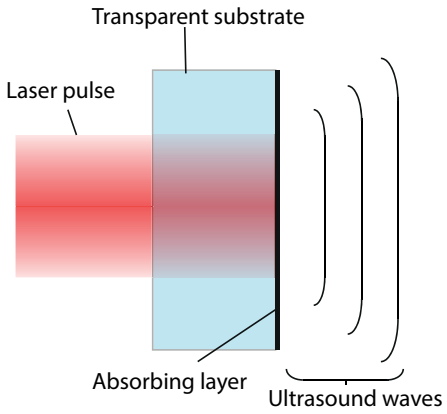

Figure 2.8: Configuration for the optical generation of ultrasound.

which is determined by thermal expansivity, heat capacity, and sound speed. The growing interest in the modality has been driven by research into materials that improve the conversion efficiency by optimising these parameters. These materials can largely be divided into two categories: carbon nano-composites, and metallic nano-particles.

2.2.2.1 Carbon nano-composites

The carbon-based absorbers consist of composites formed of different carbon allotropes mixed with a polymer. This approach was first reported by Biagi et al [90] and Buma et al [91]. The allotropes of carbon investigated include carbon black [90], multi-walled carbon nanotubes [92, 84], carbon nanofibres [93], soot derived carbon nano-particles [94], and carbon film [95]. These are dispersed into or overcoated with a polymer, most commonly PDMS, although polyurethane and epoxy have also been tested [76]. To form the absorber into a thin layer of the composite it is spin-coated onto a transparent substrate [94], or applied using a blade-based applicator [76], or dip-coated onto an optical fibre [84].

The generation of ultrasound using these composites is particularly effective, as the carbon allotropes are able to efficiently absorb incident light to become thermalised. This thermalisation is then rapidly transferred to the polymers that are directly interfaced with the nanoparticles. Since the polymer has a much higher expansivity, high acoustic pressures are generated [96].

2.2.2.2 Gold nano-particles

Metallic composites consisting of gold nano-particles, either arranged into arrays or deposited onto the tips of optical fibres, have also shown promise as transmitters. An early demonstration of this approach was by Hou et al [35]. He fabricated a 2-D array of gold nano-particles overcoated with a thin PDMS layer. The absorption spectrum of gold nano-particles is variable, depending on the size, shape, and local environment. In principle, this allows the absorption to be tailored to particular wavelengths. Similar to the carbon allotropes, after preparation the nano-particles are overcoated with PDMS or dispersed within the polymer to increase the generated pressures [97]. In general, the amplitudes generated by these metallic composites are lower than the carbon based ones, as the optical absorption and damage threshold are both lower [35].

2.2.2.3 Pressure

The signals and pressures generated by these photoacoustic transmitters are highly sensitive to certain parameters. The acoustic attenuation of the composite materials has been found to be extremely high, so the layer thickness needs to be kept to a minimum. Similarly, maximising the optical absorption (and minimising thickness) should in principle maximise the bandwidth, given the direct mapping of the absorption profile onto the signal for impulsive heating (Fig. 2.7). However, for too thin a layer, the optical pulse width can become a limiting factor [76]. Overall the generation of plane waves with peak pressures of over 5 MPa have been reported by multiple works [76, 94] and focused pressures of 50 MPa have also been achieved [92]. These are comparable to existing piezoelectric sources, so photoacoustic transmitters represent a viable method for ultrasound transduction for biomedical applications.

2.3 Controlling optically generated acoustic fields

The previous section introduced the optical generation of ultrasound and briefly summarised work that has been done to optimise the signal amplitudes. The problem addressed in this thesis is how to manipulate this transduction mechanism to

form arbitrary spatial distributions of pressure. From Eq. 2.16-2.19, it's clear that the time-varying acoustic pressure generated at a point is dependent on two parameters, the spatial distribution of the initial pressure across the planar source, and the temporal variation of the incident illumination. These can be manipulated experimentally using a variety of different techniques.

2.3.1 Spatial distribution

Equation 2.17 shows that the spatial distribution of initial pressure depends on the optical fluence $\phi(r)$ incident on each point, as well as the optical absorption $\mu(r)$ and Gruneisen parameter $\Gamma(r)$ at that location. It isn't possible to arbitrarily control the fluence incident at each point across a volume, as the light reaching a point is non-linearly dependent on the optical absorption elsewhere in the medium. However, it is possible to control the fluence across a planar absorber.

Two spatial distributions are most commonly reported in the literature: annular patterns, and diffraction gratings. Exciting an annular (or ring) shaped distribution creates an inward and outward propagating wave on the surface of a material. This inward propagating wave constructively interferes along the axis through the centre of the ring. This results in an increased signal amplitude and allows the optical energy to be distributed across the surface to avoid ablation [70]. The field generated by a ring shaped initial pressure is shown in Fig. 2.9. In comparison, exciting a diffraction grating (or series of lines) allows the generated acoustic spectra to be controlled [98] and the directivity to be enhanced [99]. For a series of lines generated on a planar surface, the acoustic signal measured, in the same plane, perpendicular to the grating consists of a set of acoustic pulses arriving with an even spacing Δt determined by the sound speed and the line separation. The acoustic spectra therefore has peaks at $f = \frac{1}{\Delta t}$ and its harmonics. Out of plane, the peak frequency varies, as the temporal spacing Δt changes as a function of the detection angle. This variation can be described using the grating equation $d \sin\theta = n\lambda$ [100], where λ is the wavelength of the peak frequency. This is demonstrated by Fig. 2.10, which shows the peak frequency recorded in a simulation of a diffraction grating pressure distribution as a function of detection angle.

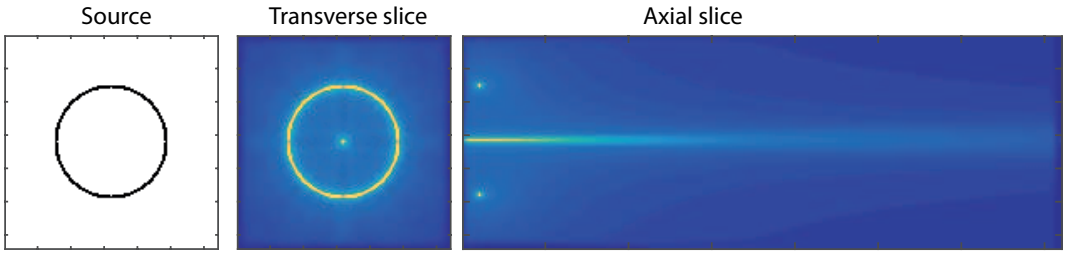

Figure 2.9: Illustrative figure of using an annular source pattern. Constructive interference is generated at the ring's centre and along the axis passing through it.

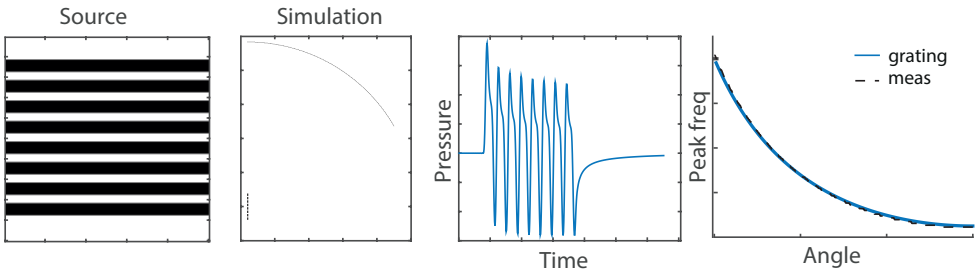

Figure 2.10: Illustrative figure of using a diffraction grating source pattern. The time series perpendicular to the array consists of a train of pulses. The repetition frequency of the pulses is determined by the grating spacing and the angle of the detection point to the array. These can be related using the grating equation.

In order to generate an annular pattern, a positive axicon lens with a conical surface can be used. Dixon et al [101] and Cielo et al [102] both used this approach then measured the peak amplitude of the resulting surface waves as a function of radial distance from the centre. Increases in the peak pressure by factors of 7.5 and 20, respectively, were measured at the focus. An alternative approach was reported by Wang et al [103, 104] who altered a laser cavity to generate optical pulses with a ring shaped profile. These pulses were made incident on an aluminium plate and the generated signal was measured on the far side. The presence of defects within the plate was found to result in conversion between different acoustic modes which could be identified in the generated signal.

To generate diffractions gratings, the most common method is the transient grating technique (TG). A single laser pulse is split in two, then made to interfere on the surface or interior of a sample. The interference of the two pulses creates a fringe pattern with a spacing determined by the angle between the two. This has been reported by both Cachier [105] and Maznev et al [106], and has been used to study transient changes in materials on femtosecond timescales [107]. An extension

to this method is to frequency shift one of the two laser pulses. In that case, a grating is generated which sweeps across the surface of the sample. This sweeping speed can be matched to the phase velocity of the excited frequency to further amplify the signal [108]. Diffraction gratings can also generated with a variety of optical elements. For example, using a cylindrical lens combined an optical diffraction grating [109], lenticular arrays (an array of cylindrical lenses) [100], and optical masks [110].

Arbitrary 2-D spatial distributions of pressure can be generated by using optical holograms to control the incident fluence. As discussed in Sec. 2.1.4, holography can be used to manipulate patterns of optical intensity in 3-D by modulating the phase or amplitude of coherent light across a 2-D aperture. Holograms that generate a desired distribution of light can be calculated using the CGH algorithms described in Sec. 2.1.5. These can then either be physically fabricated using etching [111] or lithography [112] (among other techniques), or more commonly displayed on a spatial light modulator (SLM). SLM is a general term for devices able to dynamically modulate the phase, amplitude, or polarisation of an incident optical field. They come in various forms: electro and acoustic-optic modulators, digital micro-mirror devices (DMD), and liquid crystal (LC). Of these, the latter two are the most common for this application. Both consist of arrays of microscopic pixels, in the former, mechanical mirrors that can switch 'on' or 'off' to modulate the amplitude, and in the latter, liquid crystal cells with an electronically variable refractive index which can modulate the phase. These have been employed in a number of works. For example, Zamiri et al utilised a LC SLM to generate an annular pattern for the detection of flaws [113], while Kalms et al used one to tailor the shape of the incident optical intensity to match the physical structure of a composite being tested [114].

Rather than modulate the incident fluence, it is also possible to pattern the absorbing layer itself. In this case, the pressure distribution is controlled by differences in absorption $\mu(r)$ and photoacoustic efficiency $\Gamma_\alpha(r)$ over this layer. Stratoudaki et al have developed cheap optical transducers (CHOTs) which utilise this approach [115]. These are patterned absorbers that can be deposited onto samples for re-

mote inspection with a laser. They constrain the spectrum of the resulting acoustic waves allowing for narrowband detection, improving signal to noise (SNR). Both diffraction grating and Fresnel zone plate patterns have been manufactured using this approach, for application to surface [116, 117] and longitudinal acoustic waves [118], respectively. This was also reported much earlier by Von Gutfeld, who also used a zone plate pattern to focus 10 MHz waves off axis [119]. Compared to modulation of the incident fluence, this approach is significantly easier to implement, and the cost is lower. The disadvantage is that dynamic variation of the generated pressure is not possible. The relative efficiency of the two methods depends on the light utilization efficiency of the SLM.

2.3.2 Temporal modulation

Modulating the temporal shape of the incident fluence $f(t)$ allows for the shape of the resulting acoustic signal to be controlled and the spectra to be manipulated. The exact shape of the signal generated will also depend on the spatial distribution of pressure. However, for a small spherical 'point' source, as discussed in Sec. 2.2.1.1, the resulting signal will follow the temporal derivative of the incident intensity.

The most common source for the optical generation of ultrasound are Q-switch lasers, which generate short high peak-power optical pulses by modulating the loss of the laser cavity. The pulse lengths and energies of these sources vary drastically depending on the properties of the laser. Lengths ranging from a few nanoseconds to 100's of microseconds and energies from micro joules to potentially several joules. Using a Q-switch laser pulse generates a bipolar acoustic pulse with a broadband spectra, which is illustrated in Fig. 2.11. This is useful for PAT and conventional laser generated ultrasound due to the high energies and short lengths, however, the holographic approaches described in Sec. 2.1.5 require a narrowband signal which requires alternate sources.

An alternative signal that is often used is a train of discrete short optical pulses with an even temporal spacing Δt. Similar to the diffraction grating described in Sec. 2.3.1, this generates signals with a narrowband spectra peaked at the pulsing frequency $\frac{1}{\Delta t}$ and any harmonics present in the spectra of the individual pulses

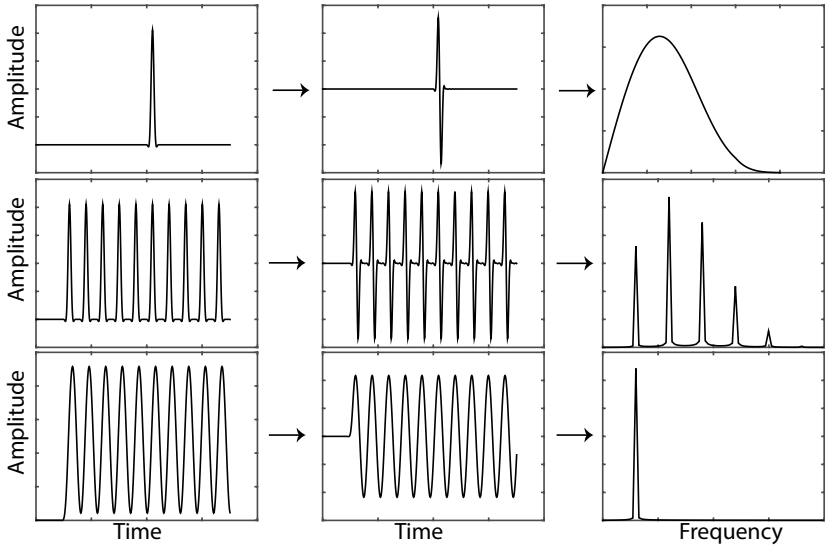

Figure 2.11: Three examples of the effect of different optical pulse shapes on the properties of the resulting acoustic pulse. Left column: incident optical energy. Middle column: measured acoustic signal. Right column: spectra of acoustic signal. Top row: single pulse. Middle row: multiple pulses. Bottom row: sinusoidal modulation.

as illustrated in Fig. 2.11. To generate these pulse trains, mode-locked lasers are commonly used [120, 121, 122, 123, 124]. These are lasers for which the optical cavity is modified to generate a short pulse on each round trip. This can be combined with Q-switching to create a train of pulses with a repetition rate defined by the cavity length. Conventionally this results in repetition rates too high for most biomedical applications (>80 MHz). However, Deaton et al developed a long cavity mode-locked laser using a white cell for which the cavity length could be varied in discrete steps from 11-60 m. This enabled the generation of pulse trains with repetition rates varying between 2.5-13.4 MHz [123, 124]. The same group also created pulse trains by repetitively Q-switching a laser during a single flashlamp pulse. However, the frequency with this method was limited to 56 kHz. Other reported techniques include the use of a fibre laser [125] and a Fabry-Perot cavity with an adjustable spacing [126]. The most flexible approach reported to date is to use an array Q-switch sources that can be independently triggered [127, 128]. In this case, any arbitrary pulsing frequency can be created by adjusting the electronic trigger, and the peak energies are sufficient to generate high acoustic pressures.

Rather than discrete pulses, continuous wave or sinusoidal signals have also

been used. In this case, the generated spectra has a single peak at the modulation frequency (Fig. 2.11). To generate these sinusoidal signals, diode or fibre lasers are the most flexible approach. Murray et al [129, 130, 131] used an electroabsorption modulated distributed feedback diode laser (EML) coupled to an Erbium doped fibre amplifier (EDFA), giving a 1 W output. The EML could be amplitude modulated at frequencies of up to 2.5 GHz using a signal generator. The disadvantage of diode and fibre lasers is that, at present, the output powers are low resulting in low acoustic pressure. For the work by Murray, the signals could only be detected via a lock-in detection scheme. To generate narrowband signals with greater amplitudes, long Q-switch pulses can be used in combination with external amplitude modulation. This was reported by Pierce et al [132] who used an acousto-optic modulator (AOM) to modulate the intensity of a 100 μs pulse generated by a Nd:YAG Q-switch laser sinusoidally at a frequency of 1 MHz. This same approach was applied by several later groups (at varying frequencies) [133, 134, 135]. An electro-optic modulator (EOM) has also been used to achieve the same outcome. With this method, the polarisation of a long Q-switch pulse is made to rotate at the desired frequency, which is then transformed into an intensity modulation by use of a polariser [136].

2.3.3 Spatial and temporal modulation

The two previous sections described different methods that can be used to modulate the spatial and temporal properties of ultrasound generated via the photoacoustic effect. In the following sections different works that have employed these techniques to steer or focus acoustic fields in 3-D are discussed.

2.3.3.1 Optical ultrasound phased arrays

Most early efforts for steering optically generated ultrasound fields attempted to replicate piezoelectric phased arrays and implement electronic focusing. Multiple groups created an array by splitting a single Q-switch pulse and coupling it into several optical fibres with differing lengths. The output of each fibre was then positioned on the surface of a sample with a set spacing. The fibre lengths and spacing could be combined in different configurations to control the relative delay and steer

or focus the field in particular directions [137, 138, 139, 140]. A similar method was used by Noroy et al who, rather than splitting a source, used multiple Q-switch sources that could be separately triggered. These different sources were again made incident on a sample with a set separation and the electronic delays tuned to steer or focus the field [127, 141]. This method was also used by Karpinnen et al [142] and Murray et al [128]. A continuous approximation of a phased array was reported by Ing et al who used a long optical pulse that was deflected across the surface of a sample at a set velocity using an AOM to create a directed wavefront [143]. Other works used the same method but instead used a rotating mirror to deflect the pulse [144, 145].

This method of focusing using spatially separated time delayed optical pulses has significant drawbacks. The number of optical sources available to each work was limited from 2-10, so the number of 'elements' is significantly constrained. Using optical fibres to introduce delays also fixes the temporal spacing, so steering or moving the focus requires impractical reorientating of the fibres. Using separate Q-switch sources is more flexible in that the trigger can be varied. However, extending this method to the element numbers available to modern piezoelectric systems ($\sim$9000) is, at present, infeasible.

2.3.3.2 Optoacoustic holography

Later works have used holographic approaches to generate a desired field. The first report of a system combining spatial control over the source with a narrowband temporal signal was by Linanne et al [146]. They combined a Q-switch mode-locked pulse train with an 80 MHz repetition rate with binary phase holograms fabricated in quartz. These holograms were designed, using a direct search approach, to generate series of arcs the spacing of which was tailored to match the acoustic travel time to the optical pulse spacing. The result was strong constructive interference at the arc focus and a narrowband signal. Later works employed the same system, replacing the quartz hologram with an SLM, for acoustic microscopy of material surfaces [147]. Alternate optical patterns were also investigated including arrays of arcs with differing orientations to create multiple foci in 2-D [148] and 1-D diffrac-

tive acoustic phase elements that achieved the same effect [149].

Grunsteidl et al developed a similar system [150, 151, 152]. In place of a Q-switch mode-locked laser, a more flexible laser diode source was used combined with a fibre amplifier. Spatial modulation was achieved using a LC SLM able to modulate the phase of the incident light. The spatial patterns consisted of a series of evenly spaced lines or rings. The driving frequency of the diode was then scanned through a wide bandwidth. When the wavelength of the signal matched the feature spacing, constructive interference occurred. By varying the spacing of the grating and the frequency, the dispersion relations of different materials could be accurately measured.

Both those systems only considered control over the acoustic field in 2-D. Full control in 3-D using holographic patterns to manipulate the optically generated pressure has been reported by a group headed by Monika Ritsch-Marte, who developed two separate systems. In an initial work, a Ti:Sapphire laser with a wavelength of 780 nm, pulse energies of up to 550 mJ, and pulse lengths of 10 μs was used in combination with an acousto-optic modulator (AOM), and a liquid crystal SLM [135]. The AOM was used to modulate the intensity of the 10 μs laser pulse to vary sinusoidally at a desired frequency. This modulated pulse was then made incident on the LCD SLM which rotated the polarization of the incoming light by either 0 or 90 degrees. A polariser then converted this into an intensity modulation. This enabled binary amplitude hologram source patterns to be generated, combined with a narrowband driving frequency. The system was tested using several fundamental holographic patterns including diffraction gratings and Fresnel zone plates. The acoustic field in each case was found to match the expected optical result.

Later, a system for phase holography was developed [136][153]. The set-up used the Ti:Sapphire laser, LCD SLM and an electro-optic modulator (EOM). The EOM was used to induce rotation of the polarization of the laser pulse at the desired ultrasound frequency. The SLM then further rotated the polarisation within each pixel by either 0 or 90, introducing a phase offset to the rotation of different pixels. The polarization modulation was then transformed into an intensity modulation

using a polarizer. The result was a pulse in which all the pixels were sinusoidally oscillating in intensity, but with a phase shift. This allowed for binary phase holograms to be used to control the field in 3-D. The system was again tested on several fundamental holograms in this case a diffraction grating, a Fresnel zone plate, a cylindrical lens, and a Bessel beam. Acoustic fields with multiple foci were also demonstrated by displaying a superposition of two Fresnel zone plates. The general idea of this approach is illustrated in Fig. 2.12.

The two systems serve as an excellent demonstration that arbitrary acoustic fields can be created with the optical generation of ultrasound. However, they did have drawbacks. Mainly, the acoustic pressures achieved were limited to 10 kPa. This was the result of two major factors: the absorber used wasn't optimised, limiting the conversion efficiency, and the pulse energy was limited by the damage threshold of the EOM. There were also trade-offs in the absorber thickness. If stress confinement was satisfied then the pressure generated from successive modulation cycles overlapped, dropping the modulation depth. However, if stress confinement wasn't satisfied, then the conversion efficiency dropped. This led to a limited bandwidth in which the system worked most efficiently.

2.3.4 Absorbing surface profiles

An alternate approach that can be used to control optically generated acoustic fields is to manipulate the shape of the surface that makes up the absorbing layer. Since no electrical connections are required, transducers with effectively arbitrary shape can in principle be generated, providing the intended profile can be fabricated from a transparent substrate, and an optical absorber can be deposited on the surface. An example of this approach is illustrated in Fig. 2.13.

This can be exploited in several ways. The simplest example is to deposit an absorber on a concave lens. This creates a geometrically focused transducer which results in a strongly focused field. Baac et al demonstrated that focused pressures of 50 MPa can be created using this approach [92, 154]. The principle has also been demonstrated using other structures. Gspan et al fabricated a corkscrew surface that generated a rotating field with a phase singularity able to impart angular momentum

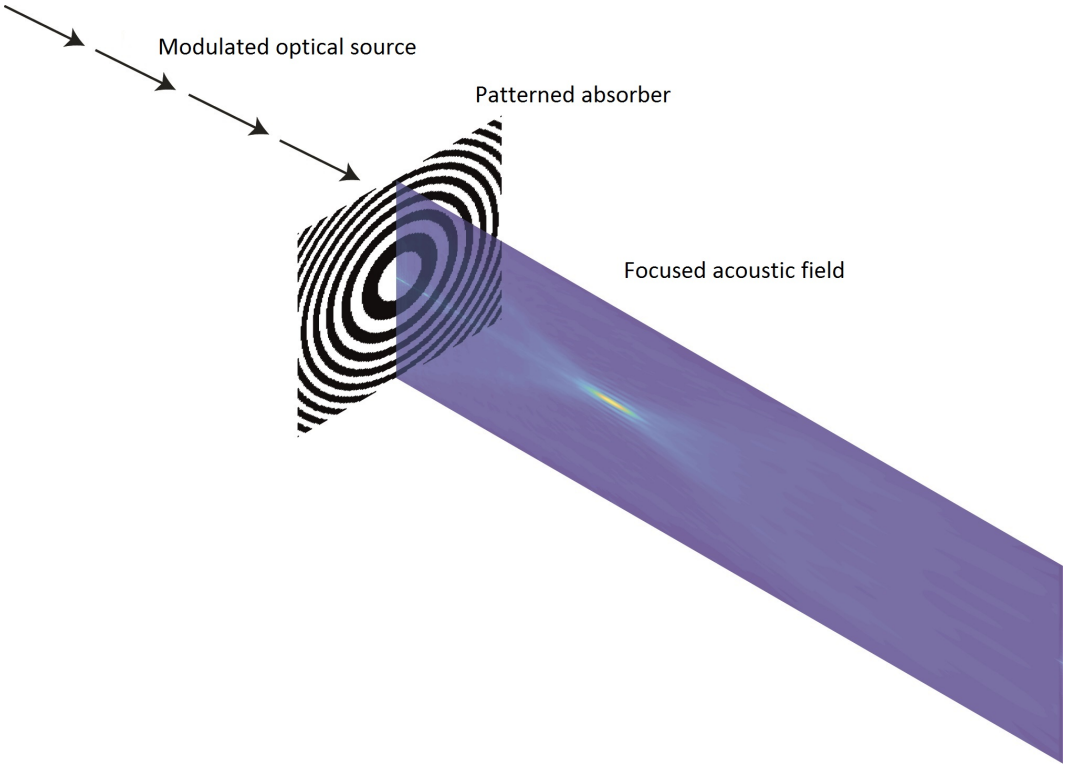

Figure 2.12: Controlling optically generated acoustic fields in 3-D. Using a modulated optical source to generate single frequency optoacoustic pulses, and spatially patterning the initial pressure distribution using holograms allows the field to be focused at arbitrary points in 3-D.

[155]. Passler et al used a conical surface to create a non-diffracting Bessel beam which had an elongated focus [156]. More recently, Chan et al created a stepped surface profile that divided the aperture into two radially. The outer half consisted of a concave surface that generated a focused wave, while the inner half was flat and generated a plane wave with a set temporal delay compared to the spherical wave [157]. Importantly, 3-D printing using a transparent substrate was used to create this surface, offering an easy method to create more arbitrary structures.

2.4 Single element transducers

As discussed in Sec. 1.2, another promising alternative to traditional arrays is to use a large single element transducer and design a structure that modulates the phase and/or amplitude of its output field to manipulate the resulting pressure. The basic idea has been used for decades in the form of acoustic lenses that focus waves in the elevation direction or at a particular depth [41]. However, recently, due in part

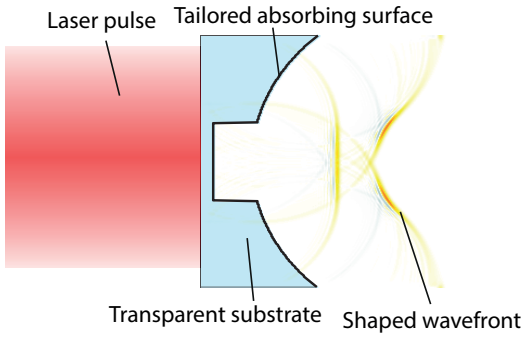

Figure 2.13: Illustrative figure of a tailored optically absorbing surface profile. In this case designed to generate a converging spherical wave followed by a planar wave.

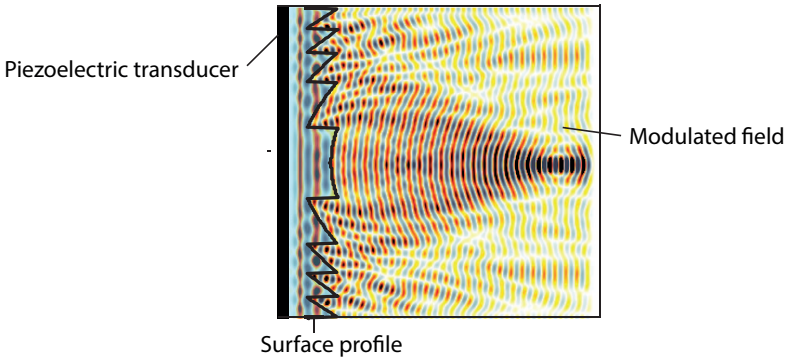

Figure 2.14: Illustrative figure of the use of a profile to modulate the output of the piezoelectric transducer.

to the increasing availability of high resolution 3-D printers, this area has seen rapid growth with a variety of complex structures able to modulate the field in different ways being reported. This is illustrated by Fig. 2.14.

One method that is particularly popular are physical arrays composed of sub-wavelength 'bricks' with a complex internal structure creating an acoustic metamaterial [158]. The properties of these metamaterial bricks can be tuned to provide different phase and amplitude modulations to an incident wave. By assembling them into different 2-D arrangements, holographic phase and amplitude distributions can be generated from an incident wave [159]. This enables the formation of arbitrary distributions of acoustic pressure. A variety of metamaterial structures have been reported: blocks with labyrinthine internal structures [158], Helmholtz resonators [160], and coiled up structures [161]. However, at present the frequencies to which this can be applied are limited to $<$ 100kHz due to the fabrication scales required for the internal structures of the bricks.

A method that scales better to higher frequencies was recently demonstrated by Melde et al [17] who developed 3-D printable structures called acoustic kinoforms that were able to slot onto single element transducers and map the output field onto pre-determined phase distributions. This mapping was achieved by varying the thickness of a 3-D printed structure across the transducer surface. Since the sound speed in the 3-D printed material differs from the surrounding medium, a phase offset was introduced at each point that was dependent on the thickness of the material at that point. This allowed for the phase to be controlled across the transducer aperture with a resolution limited by the 3-D printer. As pixel sizes on the order of 100 μm are presently achievable, this approach can be applied for frequencies up to $\sim$ 10 MHz. In the original work, phase holograms were designed using a IFTA, acoustic kinoforms were fabricated for generating these phase holograms, and the output field was characterised experimentally. The generated fields were found to match well with predictions made using the ASA. The kinoforms were then used both to create complex arrangements of particles and to physically propel objects using pre-defined phase tracks. The concept of mapping the output of a single element transducer onto a phase hologram was demonstrated earlier for focusing at a single point, though it was not identified as such [162, 163].

2.5 Thesis content

This thesis investigates different approaches for the generation of arbitrary acoustic fields using the optical generation of ultrasound and single element transducers, some of which expand on approaches reviewed within this Chapter. With the exception of Chapters 3, and 4 the thesis is divided into Chapters that each address a different method.

Chapter 3 investigates binary amplitude holograms excited by trains of optical pulses for the formation of a single foci. The properties of the acoustic fields and how they're influenced by different experimental parameters are explored numerically. Experimental validation using laser cut patterned absorbers is then reported.

Chapter 4 extends the work of Chapter 3 to multi-focal binary amplitude holo-

grams excited by the same optical source. An algorithm based on the direct search approach discussed in Sec. 2.1.5.4 is developed for the calculation of these holograms. The performance of this algorithm is optimised and the properties of the fields are again investigated numerically. Experimental validation is then presented using the method described in Chapter 3.

Chapter 5 investigates the generation of multi-focal acoustic fields using a single optical pulse. This is achieved by using multi-layer patterned absorbers. A ray-tracing algorithm is developed to rapidly approximate the impulse response of pixels in each absorbing layer to a pulsed excitation at a target point. An optimisation approach based on modified direct search is then developed to design non-overlapping absorbing patterns for each layer. Experimental measurements from a 3 layer structure demonstrate this method can be used to create multi-focal acoustic fields using a single optical pulse.

Chapter 6 also investigates the generation of patterned acoustic fields using single optical pulses. In this case, the approach discussed in Sec. 2.3.4 of shaping the surface profile is generalised to the creation of multiple foci distributed arbitrarily in three dimensions. An optimisation is again adapted to address the calculation of surface profiles for this new problem. Experimental validation using a transparent 3-D printed sample onto which a thin absorbing layer is deposited is then reported.

Chapter 7 then develops the concept of multi-frequency acoustic kinoforms. These extend the approach developed by Melde et al by enabling different target distributions to be encoded onto different frequencies. The direct search approach reported in Chapter 6 is adjusted for the design of these multi-frequency kinoforms and the approach is again demonstrated experimentally using two test cases. These demonstrate that both continuous patterns of pressure and discrete foci can be generated using this approach. A detailed numerical investigation is then carried out into how different design parameters influence the performance of these kinoforms.

The overall conclusions of the work are then presented in Chapter 8.

Chapter 3

Single focus optoacoustic holograms

3.1 Introduction

This Chapter investigates the generation of acoustic fields with a single focus using binary amplitude holograms excited by trains of short high-power optical pulses. This is similar to the approach adopted by Meyer et al [135], who used an optical source that was sinusoidally modulated temporally then spatially modulated to create a system for binary optoacoustic holography. However, the use of a train of optical pulses introduces harmonics to the spectrum, the effect of which requires characterisation. In Sec. 3.2, the basic properties of single focus binary amplitude holograms are introduced. This is followed in Sec. 3.3 by an introduction to the numerical model used in this and in later Chapters. Sections 3.4-3.8 describe a series of numerical investigations into the properties of the acoustic fields generated using this approach and how they are influenced by different properties of the source. Finally, in Sec. 3.9 the model is experimentally validated. This uses patterned optical absorbers, fabricated by laser cutting thin optically absorbing layers formed by aerosol deposition, excited by a Q-switched laser source.

The journal article in [164] has been modified and adapted to form parts of this chapter, with reprint permission under CC BY 4.0 License.

3.2 Single focus holograms

A binary amplitude hologram designed to focus at a single point is known as a Fresnel zone plate (FZP). These have several well known properties that have been

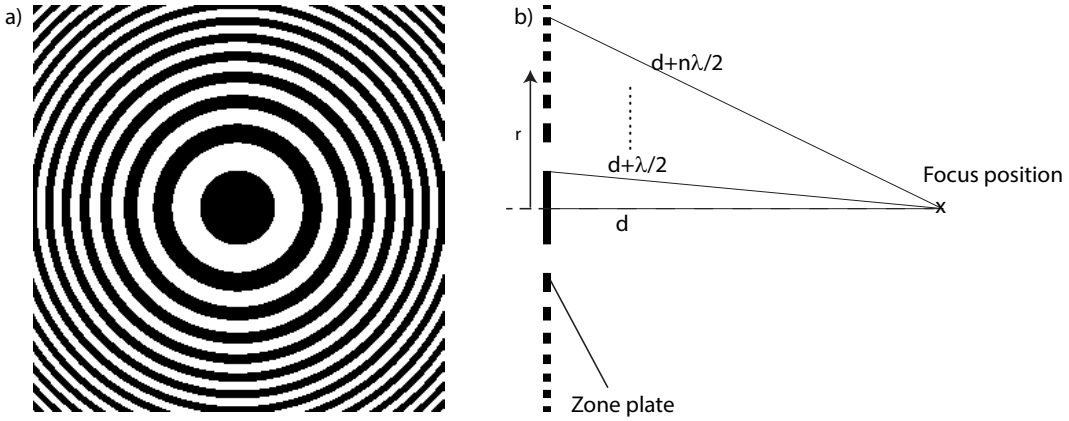

Figure 3.1: (a) Example FZP with a 3 cm aperture designed to focus a frequency of 3 MHz at a depth of 3 cm. In this image and for the rest of this Chapter black represents a transmitting point and white is a blocked portion. (b) Schematic illustrating the geometric relations underlying the zone plate structure.

studied both for optics and for x-rays [165, 166]. A FZP consists of a series of rings called zones, that alternate between opaque and transparent, alternately blocking and transmitting light. The radii for which the FZP transitions between opaque and transparent zones is constructed such that the path length difference to the focus across each zone varies by half a wavelength of the design frequency. This ensures that for an incident plane wave, all points on the aperture that constructively interfere at the design focus are transmitted while all those that destructively interfere are blocked. This results in strong constructive interference at the target depth. Figure 3.1 illustrates this showing an example FZP along with its underlying geometric relations.

Based on the requirement for the zone widths, it is simple to derive an expression for one set of radii r_n for a FZP with a design wavelength λ with a focus at depth d. If the first zone is considered to start at $r = 0$, then the first transition will occur when the distance to the focus is $d + \frac{\lambda}{2}$. Similarly the transition between the $(n-1)^{th}$ and n^{th} zones will occur when the distance to the focus is $d + \frac{n\lambda}{2}$. The radius for this is given by

$$\sqrt{r_n^2 + d^2} = d + \frac{n\lambda}{2}, \tag{3.1}$$

which can be rearranged to yield

$$r_n = \sqrt{n\lambda d + \frac{n^2\lambda^2}{4}}. \tag{3.2}$$

For most FZP $\lambda \ll d$, so for small n Eq. 3.2 can be approximated by

$$r_n \approx \sqrt{n\lambda d}. \tag{3.3}$$

An alternative approach to define a FZP starts by placing a spherical point source at the target depth d. The pressure generated elsewhere in the field by this point source is given by

$$p(l) = \frac{1}{4\pi l} e^{i\left(\frac{2\pi l}{\lambda} + \Phi\right)}. \tag{3.4}$$

Here Φ is an arbitrary phase factor and l is the distance from the source. Taking the imaginary component of the pressure across the hologram aperture and thresholding about 0 creates a FZP. To show this, the imaginary component of the pressure across the aperture (from Euler's formula) is given by

$$p(r) = \frac{1}{4\pi l} \sin\left(\frac{2\pi\sqrt{d^2 + r^2}}{\lambda} + \Phi\right). \tag{3.5}$$

Here r is the radial distance on the aperture from the focal position. If the pressure $p(r)$ is thresholded about 0, the transition from opaque to transparent will occur when the argument of the sin term is a multiple of π. From Eq. 3.5 it can be seen that this occurs when

$$\frac{2\pi\sqrt{d^2 + r^2}}{\lambda} + \Phi = n\pi. \tag{3.6}$$

Expanding and rearranging this expression yields

$$r = \sqrt{\frac{\lambda^2 n^2}{4} + \frac{\lambda^2\Phi^2}{4\pi^2} - \frac{n\lambda^2\Phi}{2\pi} - d^2}. \tag{3.7}$$

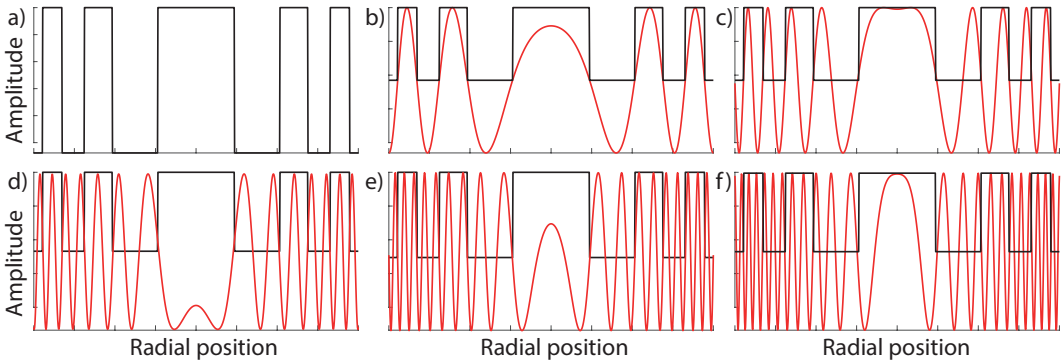

Figure 3.2: (a) Example zone plate amplitude as a function of radius designed to focus at a depth d. (b-f) Real component of pressure generated across zone plate aperture for point sources at different depths. (b) Depth - d. (c) Depth - $d/2$. (d) Depth - $d/3$. (e) Depth - $d/4$. (f) Depth - $d/5$. For (c) and (e) the pressure across each zone integrates to 0, for (b), (d), and (f) the pressure across each zone doesn't cancel out.

If Φ is set to $-\frac{2\pi d}{\lambda}$ then Eq. 3.7 becomes

$$r = \sqrt{n\lambda d + \frac{n^2\lambda^2}{4}}. \tag{3.8}$$

This is identical to Eq. 3.2, showing a thresholded spherical source yields a FZP. This was the approach used to calculate the zone plates used in this Chapter.

A key property of FZPs is that in addition to the main focus they also generate several additional foci called high order foci. These occur along the focal axis at distances of roughly d/m, (m = 1, 3, 5,...). The higher order foci are generated at these depths as for distances of d/m, the difference in path length across each zone is approximately $m\lambda/2$ and, similarly, the path length difference from equivalent areas on neighbouring transmitting zones is $m\lambda$. So for odd m, constructive interference will occur resulting in a focus. For even m, the waves transmitted from each zone should (ideally) cancel out. This is illustrated by Fig. 3.2, which shows the pressure as a function of radius across the zone plate aperture generated by point sources placed at depths of d/m. For even m (Fig. 3.2 (c) and (e)) the pressure across each zone approximately integrates to zero, for odd m (Fig. 3.2 (b), (d), and (f)) the pressure doesn't cancel out.

3.3 k-Wave toolbox

In order to characterise the acoustic fields generated by zone plates excited by a train of optical pulses, a set of simulations varying different parameters was carried out. These were run using the k-Wave toolbox, an open-source software for time-domain simulations of acoustic fields [81]. This toolbox is used extensively throughout this thesis for the numerical evaluation of the acoustic fields generated by different sources, in different media.

The k-Wave toolbox models the time-domain propagation of acoustic waves in a heterogeneous and absorbing medium by solving the first-order differential equations that arise from mass and momentum conservation, and the pressure density relation. In the case of a homogeneous medium with no losses, these are [167]

$$\frac{\partial u}{\partial t} = -\frac{1}{\rho_0}\nabla \cdot p, \tag{3.9}$$

$$\frac{\partial p}{\partial t} = -\rho_0 \nabla \cdot u, \tag{3.10}$$

$$p = c_0^2 \rho, \tag{3.11}$$

where u is the acoustic particle velocity, ρ is the acoustic density, and ρ_0 and c_0 are the ambient density and sound speed respectively. In k-Wave, these equations are solved using a k-space pseudospectral method. The medium is discretised into a regularly spaced cartesian grid on which the sound speed and density (along with other acoustic parameters as required) are defined. The spatial gradients of the pressure and velocity are calculated using a Fourier collocation method that fits a Fourier series to the entire domain. The temporal derivatives are calculated using a k-space corrected first-order finite difference. To eliminate wave wrapping introduced by the Fourier collocation method, a perfectly matched layer (PML) is used to attenuate waves as they leave the medium [168].

Absorption and non-linearity can be modelled using additional terms in the system of coupled equations, however, unless otherwise stated, for this thesis they are not included. Pressure can be introduced into the domain in several ways. Time

varying force or mass sources can be added using additional terms in Eq. 3.9 or Eq. 3.10, respectively. Alternatively, for pulsed photoacoustic problems, initial pressure distributions can be explicitly defined with the properties outlined in Sec. 2.2.1. In this thesis, only mass sources and initial pressure distributions are used. For each simulation, the relevant properties of the grid, medium, and source are defined in the text.

3.3.1 Simulation parameters

For all the simulations in this Chapter, the medium was homogeneous with a density of 1000 kg m^{-3} and a sound speed of 1500 m s^{-1}. Unless otherwise stated, the simulations were carried out on a regular $512 \times 512 \times 512$ cartesian grid with a grid point spacing of 25 μm, and time steps of 5 ns. This gave a maximum supported frequency of 30 MHz. The grid spacing determines the points per wavelength (PPW) for each frequency. For each experiment, the design frequency was set such that it had at least 3 PPW to minimise numerical errors [45]. The thickness of the PML was set to 10 grid points. Holograms were inserted along the edge of the simulation domain as a binary source mask. The train of optical pulses incident on the hologram was represented temporally as impulses at discrete simulation time points. These were inserted at the nearest time point to the analytic pulsing frequency. This quantisation of the time step introduced a small error into the simulation spectra, however, this was found to have no noticeable effect on the simulation results. The temporal input was low-pass filtered at 30 MHz to remove frequencies not supported by the spatial grid. The simulations were primarily run on a compute server consisting of 2 $\times$ 8-core Xeon E5-2650 2.0 GHz CPUs, and 256 GB of 1600 MHz DDR3 memory.

In the next two Chapters, a wide range of different values are used for the design frequency (1-10 MHz), aperture (1-7.5 cm), and focal depth (1-7.5 cm) of the single focus and later multi-focal holograms investigated. This is used to demonstrate that the approach is robust across a range of different bandwidths and spatial scales.

3.4 Pulsed sources vs monochromatic

The construction of a FZP discussed in Sec. 3.2 assumes a monochromatic input spectrum. However, as discussed in Sec. 2.3.2, the spectrum of the acoustic wave generated by a short optical pulse is broadband or, if generated at a set repetition rate, has peaks at the pulse repetition frequency (PRF) and harmonics of the PRF contained in the broadband spectrum of the single pulse.

To establish the effect the difference between the assumed spectra and the actual spectra has on the resulting field, the monochromatic and pulsed inputs were compared in simulation. A hologram was calculated for a design frequency of 5 MHz, with a 0.03 mm pixel spacing, a 1.5×1.5 cm aperture, and a focus at 1.5 cm. The fields generated from this hologram were simulated on a $540\times540\times1536$ grid, with a grid spacing of 0.03 mm. The maximum supported frequency was 25 MHz. The hologram was inserted as a source distribution in the first plane. The two optical inputs were represented in the simulation using a sinusoid (for the monochromatic input), and using a continuous train of delta functions spaced at 5 MHz (to represent a pulsed laser input). These are shown in Fig. 3.3 (a) and (b). Figure 3.3 (c) shows the maximum amplitude projection of the peak pressure in each voxel of the hologram field generated with a pulsed laser input. Figure 3.3 (d) is the same field generated by a monochromatic input.

There are several significant differences introduced by the harmonics in the acoustic spectrum. Firstly, the volume of the design focus, labelled f_0 in both figures, measured by thresholding the maximum pressure data at 50% of the overall maximum, is reduced by 92% compared to the monochromatic input. Secondly, the signal to noise of the focus relative to the background is decreased. Thirdly, there are several additional foci present in the pulsed field not present in the monochromatic field. These are generated by the harmonics in the source spectrum. From Eq. 3.3 it can be seen that halving the wavelength λ and keeping the radii constant approximately doubles the focal length d. Therefore, the hologram focuses each harmonic at a point further away from the hologram aperture. As with the fundamental, each harmonic also generates higher-order foci at points where the path

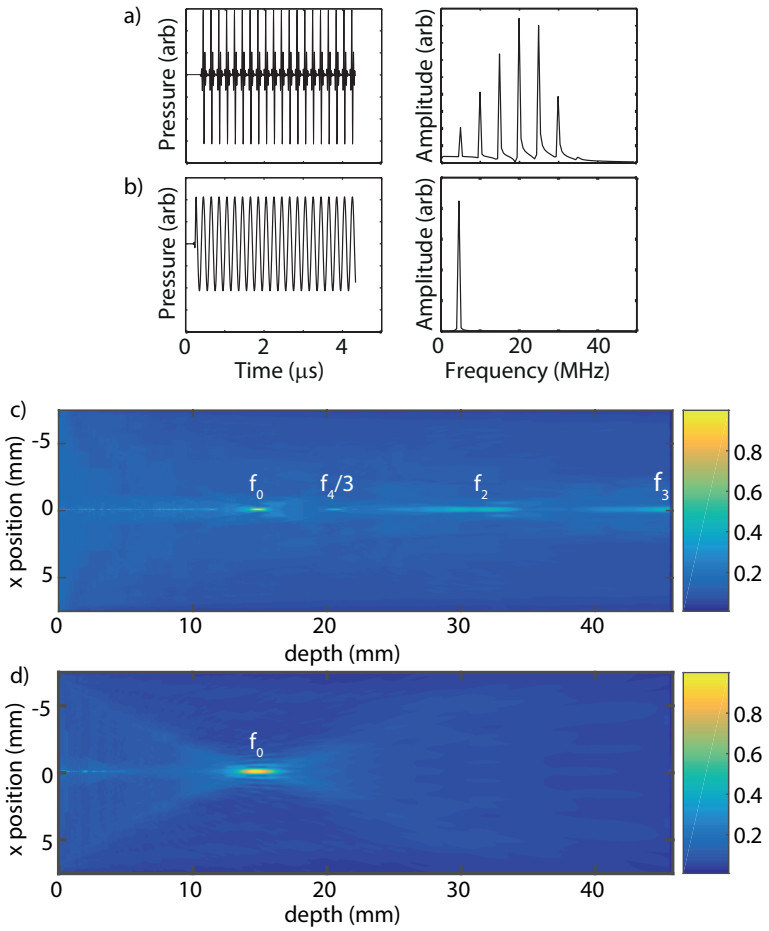

Figure 3.3: (a) Acoustic signal and spectra generated for the pulsed temporal input. (b) Acoustic signal and spectra generated for the sinusoidal temporal input (c) Normalised maximum amplitude projection of maximum pressure generated by the hologram with a pulsed input. (d) Normalised maximum amplitude projection of maximum pressure generated by the hologram with a monochromatic input. The pulsed input generates a smaller focus volume but also generates several additional foci and has lower signal to noise. Figure reprinted from [164] under CC BY 4.0 License.

length from subsequent rings (or zones) differs by an odd number of wavelengths of that harmonic. An example is labelled in Fig. 3.3 (c) as $\frac{f_4}{3}$, which is the third order focus for the fourth harmonic. These observations were verified in simulation by recording the spectrum at each focal point. The peak frequency in the spectrum was found for each focal point to match the predicted harmonic.

The higher-order foci of certain harmonics also spatially coincide with the design focus, for example $\frac{f_3}{3}$. This spatial overlapping is responsible for the reduction in focus volume observed in the pulsed input relative to the monochromatic input.

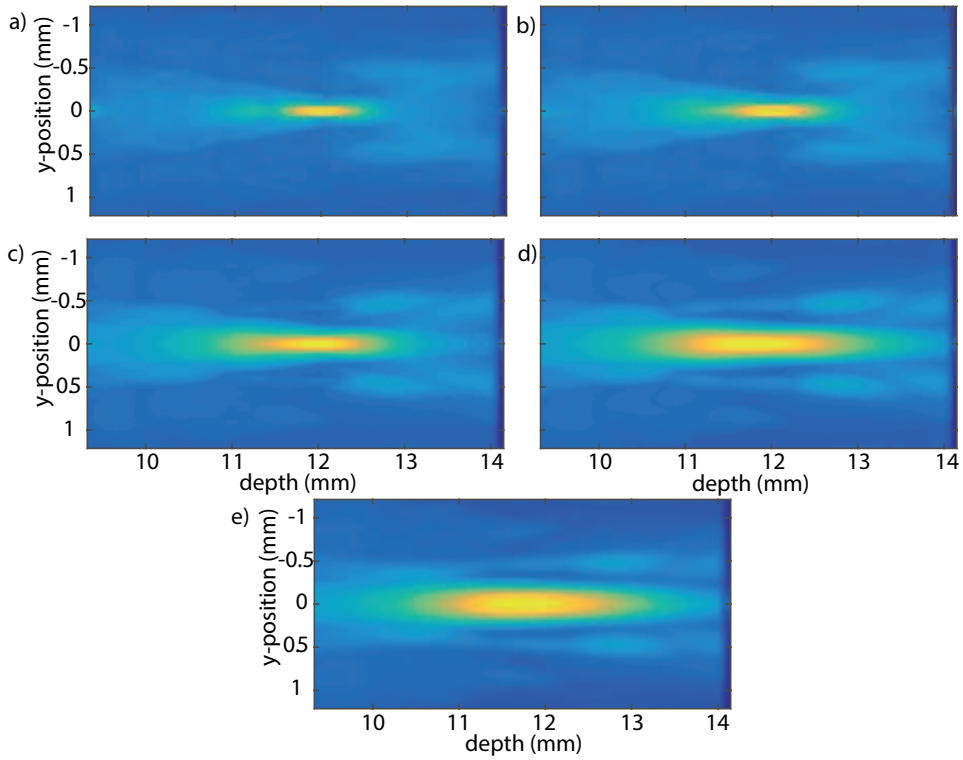

Figure 3.4: Maximum amplitude projections of the maximum pressure generated by a 5 MHz hologram possessing a 1.2×1.2 cm aperture, and a focus at 1.2 cm. Simulations were carried out on a 540×540×600 grid with a spacing of 0.024mm which supported frequencies up to 31.25 MHz. The source was low-pass filtered with different cut-off frequencies. A small area surrounding the focus was selected for visualisation. (a) 27.5 MHz cutoff (4 harmonics) . (b) 22.5 MHz (3 harmonics). (c) 17.5 MHz (2 harmonics). (d) 12.5 MHz (1 harmonic). (e) 7.5 MHz (0 harmonics). The volume of the focus increases with decreasing filter cutoff.

To observe the effect of different harmonics on the structure of the resulting acoustic field directly, the fields generated by a 5 MHz hologram were simulated. For these simulations, the input source was low pass filtered starting at 27.5 MHz, the cutoff was then decreased down to 7.5 MHz in intervals of 5 MHz. Maximum amplitude projections of the acoustic field around the focus from each of these simulations are shown in Fig. 3.4.

It can be seen that as the cutoff frequency decreases, the volume of the focus increases. Based on the observation that the volume reduction is caused by higher-order foci for odd harmonics, the volume of the focus should only increase between simulations that eliminate odd harmonics. However, the filter isn't perfectly flat so this is not observed. Apart from the change in focal volume, changes in the pressure

generated in the near field of the hologram before the focus were also found to occur. As the cutoff decreases, the average pressure generated in the near field decreases, and the number of foci also lowers. This is a result of the other higher-order foci that don't coincide with the focus being eliminated.

3.5 Harmonic zone plates

The reduction in focal volume caused by harmonics in the hologram field is useful. For example, it allows for small foci to be generated without the need to pulse or design the hologram for a high frequency. However, the generation of additional foci is undesirable. For example, in therapeutic ultrasound it could result in heating at targets away from the desired focus. For other applications, the additional foci are less of a concern (i.e., in imaging they could be windowed out). However, adjusting the design of the zone plate to account for the presence of these harmonics is still worthwhile. This can be done by altering the thickness of each ring on the zone plate to change the distribution of pressure generated by each harmonic at its design focus and at each of its higher-order foci. As the zone plate is designed only to optimise focusing for the fundamental, when harmonics are present, this is not necessarily optimal.

For a zone plate designed to focus a frequency f at a distance d away from the zone plate, the relative maximum pressure $p_{\max}(f)$ generated by the fundamental and the first two harmonics at the focus d can be approximated by the following equations

$$p_{\max}(f) \propto \int_0^{\pi} sin(\theta) d\theta, \tag{3.12}$$

$$p_{\max}(f_2) \propto 2\int_0^{2\pi} sin(\theta) d\theta, \tag{3.13}$$

$$p_{\max}(f_3) \propto 3\int_0^{3\pi} sin(\theta) d\theta. \tag{3.14}$$

Here $p_{\max}(f_2)$ and $p_{max}(f_3)$ are the maximum pressures of the second and third harmonic respectively, and the factors of 2 and 3 in Eq. 3.13 and 3.14 are added

to account for the scaling of focal gain with frequency. These equations derive from the path length difference between acoustic waves arriving from the inner and outer radii of a ring being $\lambda/2$, which corresponds to a phase difference of π for the fundamental. It can be seen from Eq. 3.13 that the second harmonic almost entirely cancels out at the focus, and from Eq. 3.12 & 3.14 that the third harmonic and fundamental should generate similar pressures.

Alternatively, consider a zone plate calculated for a frequency of f_3 to focus at a distance d away from the zone plate from which 2 out of 3 rings have been removed. This has an equivalent effect to adjusting the thickness of each of the rings on the fundamental zone plate by a specific amount. An example is shown in Fig. 3.5 (a). For this zone plate the maximum pressure $p_{\max}(f)$ generated by the fundamental and the first two harmonics relative to each other can be approximated by

$$p_{\max}(f) \propto \int_{\frac{\pi}{3}}^{\frac{2\pi}{3}} sin(\theta) d\theta, \tag{3.15}$$

$$p_{\max}(f_2) \propto 2 \int_{\frac{\pi}{6}}^{\frac{5\pi}{6}} sin(\theta) d\theta, \tag{3.16}$$

$$p_{\max}(f_3) \propto 3 \int_{0}^{\pi} sin(\theta) d\theta. \tag{3.17}$$

This was derived from the observation that the path length difference between acoustic waves arriving from the inner and outer radii of each ring is $\frac{0.5\lambda}{3}$. This is known because the harmonic zone plate was calculated initially from a frequency of f_3 to focus at an identical point to the original zone plate. So geometrically, the path length difference for each ring remaining on the zone plate must be $\frac{0.5\lambda}{3}$. This corresponds to a phase difference of $\frac{\pi}{3}$ for f, $\frac{\pi}{2}$ for f_2 and π for f_3. Comparing Eq. 3.12-3.14 and 3.15-3.17, the relative maximum pressure predicted by the two is 1:1.19. This is because, the second harmonic now contributes to the pressure generated at the target focus, and the fundamental and third harmonic have not significantly reduced. This suggests that despite having lower number of transmitting

pixels, the harmonic zone plate should generate higher peak pressures. Extending this to 5 harmonics, the predicted ratio becomes 1:1.24.

A simulation was carried out to validate these approximations. The field from Fig. 3.3 (c), which was generated by a 5 MHz zone plate with a 1.5×1.5 cm aperture and a focus at 1.5 cm, was compared against the field generated by a harmonic zone plate simulated on the same grid for the same input (i.e., same number of harmonics). The harmonic zone plate used is shown in Fig. 3.5 (a). This was generated by first calculating a hologram for a 15 MHz design frequency with a 1.5×1.5 cm aperture and a focus at 1.5 cm, then removing two out of three rings. Maximum amplitude projections of both acoustic fields are shown in Fig. 3.5 (b)-(c).

It can be seen that the maximum pressure generated by the harmonic zone plate is greater than the conventional zone plate. The ratio between the two differs slightly from the prediction due to the maxima generated for different harmonics not temporally coinciding due to phase differences, and because the input spectra in the simulation wasn't perfectly flat. Other changes between the two fields are that the volume of the focus generated by the harmonic zone plate is decreased by a further 33% relative to the conventional zone plate, the signal to background ratio is improved, and the pressure of the additional foci is decreased. The spectra at the design focus was measured in both simulations, these are shown in Fig. 3.5 (d) and (e). These show that, as predicted, the second and fourth harmonics are almost completely absent for the conventional zone plate and for the harmonic zone plate they have a larger amplitude.

These results suggests that harmonic zone plates could outperform conventional zone plates for broadband pulsed input sources. The lower number of transmitting pixels on the harmonic zone plate also means that for a hologram generated by spatial modulation of an incident laser pulse, the optical fluence for each pixel could be greatly increased.

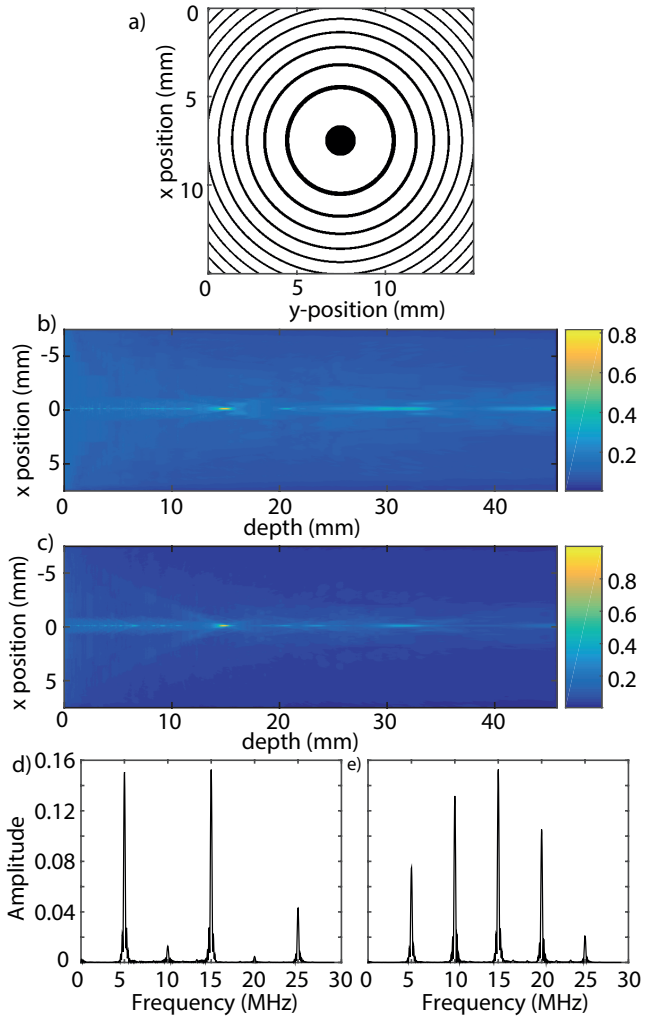

Figure 3.5: (a) 15 MHz harmonic zone plate, designed for a 3×3 cm aperture with a focus at 3 cm. (b) Acoustic field from Fig. 3.3(a), pressure normalised relative to maximum in (c). (c) Normalised maximum amplitude projection of the maximum pressure generated by the hologram from (a) with a pulsed input. The harmonic zone plate generates greater pressure and has lower background compared to the conventional zone plate. (d) Spectra generated at target focus for conventional zone plate. (e) Spectra generated at target focus for harmonic zone plate.

3.6 Zone plate resolution and pulse number

After establishing the basic properties of the fields generated by a train of optical pulses, the parameters required to represent a hologram and generate an axial focus were investigated. This was done by measuring the effect of resolution (pixel size) and the number of laser pulses on the focus volume and peak pressure. The results are plotted in Fig. 3.6. Simulations were carried out using four equivalent holograms calculated with resolutions of $\frac{\lambda}{16}$ $\frac{\lambda}{8}$, $\frac{\lambda}{4}$ and $\frac{\lambda}{2}$ relative to the design frequency of

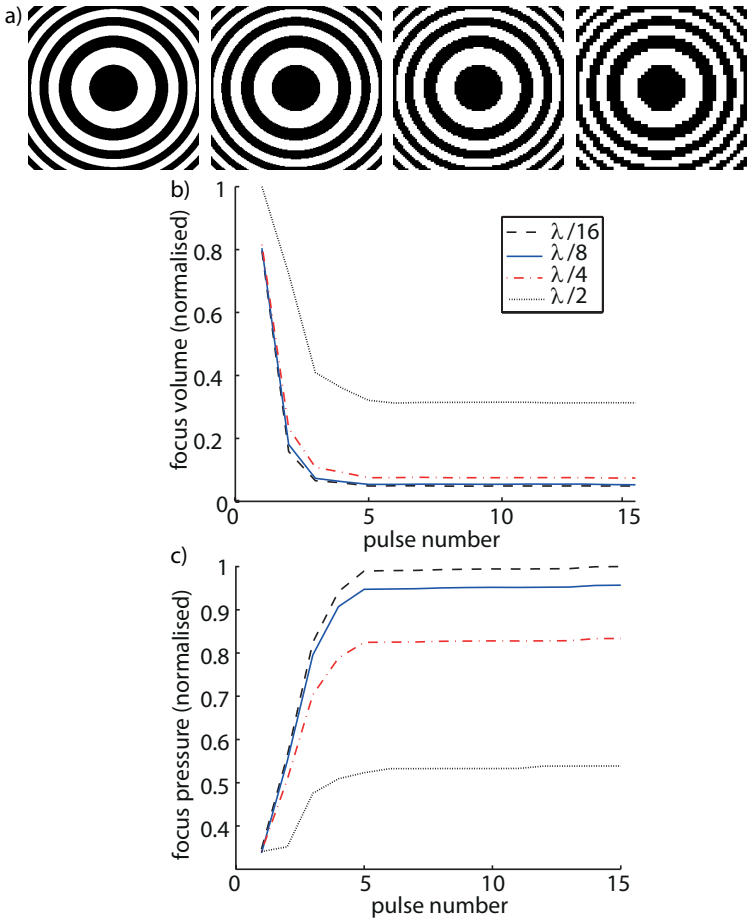

Figure 3.6: (a) 4 holograms with differing resolutions used in simulations. The resolutions from left to right are $\frac{\lambda}{16}$, $\frac{\lambda}{8}$, $\frac{\lambda}{4}$ and $\frac{\lambda}{2}$. (b) Variation in focus volume with the number of pulses applied in the simulation for each hologram. (c) Variation in maximum pressure at the focus with the number of pulses for each hologram. The maximum pressure increases and focus volume decreases with increasing hologram resolution and increasing pulse numbers. Figure reprinted from [164] under CC BY 4.0 License.

3.75 MHz. Each hologram had a 1×1 cm aperture and a focus at a depth of 1 cm. For each hologram, 15 simulations were carried out increasing the number of pulses applied from 1 to 15. The volume of the focus was evaluated in each simulation by thresholding the maximum pressure at each simulation voxel at 50% of the maximum pressure at the focus, then calculating the number of voxels in the cluster containing the target focus point.

As shown in Fig. 3.6 (b), the volume of the focus decreases with each additional pulse, reaching a steady state for each hologram at $\sim$6 pulses, which matches the number of rings on the hologram. It is also the point where the temporal spacing

between the first and last pulses in the input signal matches the travel time difference between waves from the edge and centre of the hologram to the focus. The largest decrease in focal volume is seen between the 1 and 2 pulse inputs. This is because, for a single pulse, the acoustic field isn't axially focused as there is no temporal overlap between the acoustic signal from different zones. For two or more pulses, the acoustic signal from neighbouring rings can constructively interfere. For example, the second ring excited in the first pulse temporally overlaps with the first ring excited in the second pulse, and this generates a focus.

The volume of the focus at any given pulse number decreases with increasing hologram resolution. However, the fractional improvement reduces with each subsequent increase in resolution, e.g., a gain of only $\sim 8\%$ is seen between the $\frac{\lambda}{8}$ and $\frac{\lambda}{16}$ holograms. This is likely because the decrease in volume arises due to imperfect representation of the hologram at lower resolutions (staircasing). As the hologram resolution increases, the size of the discrepancy also reduces and the result starts to converge.

The change in field with increasing pulse numbers is illustrated more clearly by Fig 3.7. This shows maximum amplitude projections for 4 simulations carried out using a 10 MHz zone plate applying increasing pulse numbers. Figures 3.7 (a)-(b) clearly show the generation of a foci that occurs after the application of 2 pulses to the hologram. Figures 3.7 (c)-(d) illustrate the further decrease in the volume of the focus that occurs as a greater number of pulses are applied and the field approaches a steady state.

Similar trends to the change in volume are observed for maximum pressure as shown in Fig. 3.6 (c). The pressure at the focus increases with increasing resolution and pulse numbers, and this also reaches a steady state for each resolution at approximately 6 pulses. A large gain $\sim 85\%$ is measured in the steady state pressure at the focus as the resolution is increased from $\frac{\lambda}{2}$ to $\frac{\lambda}{16}$. As with the focus volume, the fractional gain in pressure reduces with each increase in resolution. For example, an increase of only 4.5% is measured between the $\frac{\lambda}{8}$ and $\frac{\lambda}{16}$ holograms.

The spacing between rings on a zone plate, and their width, decreases with in-

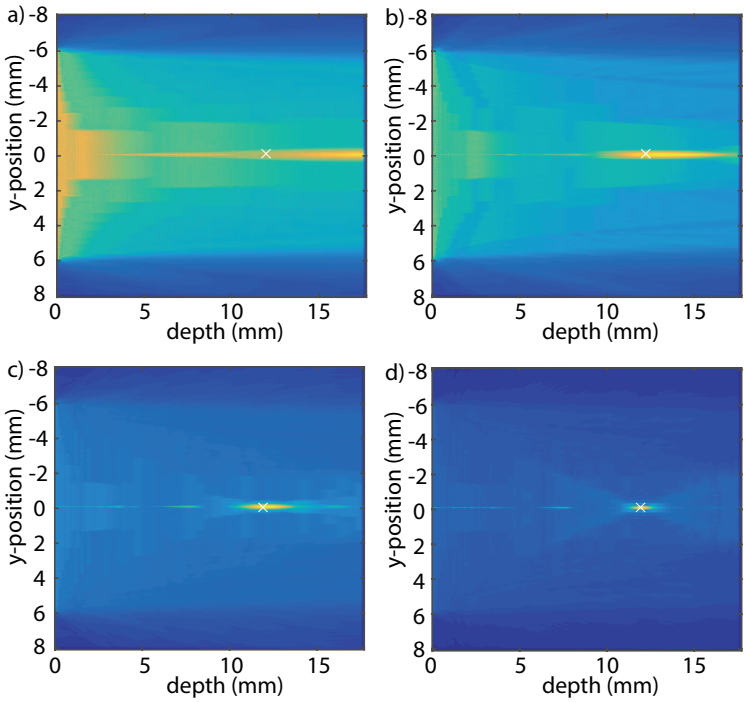

Figure 3.7: Maximum amplitude projections from 4 simulations carried out using a 10 MHz zone plate. The zone plate had a 1.2×1.2 cm aperture and a focus at 1.2 cm. Simulations were carried out on a 540×540×540 grid with a spacing of 0.024 mm. The number of pulses applied to hologram was varied between simulations. (a) 1 pulse. (b) 2 pulses. (c) 5 pulses. (d) 40 pulses. The white cross indicates target focus.

creasing distance from the centre (this can be seen from Fig. 3.1). This means that both the resolution required to represent the zone plate and the number of pulses needed for the field to converge will increase with increasing aperture size. Both parameters will also vary depending on the depth of focus relative to the aperture and if the focus is moved off the hologram centre. However, the simulations demonstrate that a focus can be generated with a relatively low hologram resolution $\sim \frac{\lambda}{2}$ (corresponding to 75 μm for 10 MHz), and with as few as two optical pulses. The limit on the aperture size can also be bypassed using similar observations to those applied for harmonic zone plates in Sec. 3.5. Composite zone plates can be created for which the inner portion is designed to conventionally focus at a depth d and the outer portion is designed to focus at $3d$ [169], such that the higher order focus generated from this area of the aperture coincides with the target point. The efficiency of this surrounding section is lower, however, the zone thickness is also increased enabling fabrication for cases where the resolution is limited. A composite zone

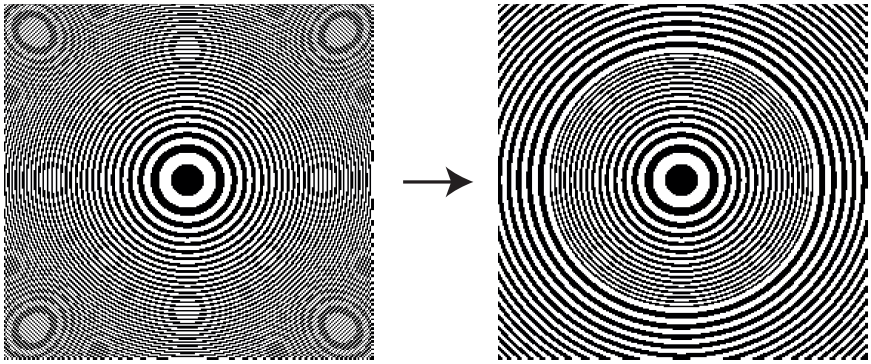

Figure 3.8: Illustrative figure of the structure of a composite zone plate designed such that the outer half of the aperture generates a higher order foci at the main focus.

plate is shown in Fig. 3.8.

3.7 Optical pulsing frequency

After establishing the conditions required for convergence of the field from a particular hologram, the degree to which the acoustic field of individual holograms could be varied was considered. It has already been found that the focal length of a hologram varies depending on the input frequency. As the spectrum generated by a train of optical pulses is determined by the pulse repetition frequency (PRF), it should be possible to vary the axial focus position by altering the PRF.

A prediction of the focal depth of each ring on the hologram can be calculated by inverting Eq. 3.2

$$d_{\text{pred}} = \frac{1}{N}\sum_{n=1}^{N}\left(\frac{r_n^2}{n\lambda(PRF)} - \frac{n\lambda(PRF)}{4}\right). \tag{3.18}$$

Here d_{pred} is the prediction for the focal depth, $\lambda(PRF)$ is the wavelength of the pulsing frequency, and N is the number of rings on the zone plate. Equation 3.18 averages the focal depth predicted by each ring radii r_{n} for a particular PRF.

To validate Eq. 3.18, predictions made using this formula were compared against simulations. A hologram was calculated for a design frequency of 9 MHz, with a 0.75×0.75 cm aperture, and a focus at 0.75 cm. The fields generated by this hologram were simulated, altering the PRF between 4.5 MHz and 13.5 MHz. For each simulation, the input source was filtered at 2.2 points per wavelength (PPW)

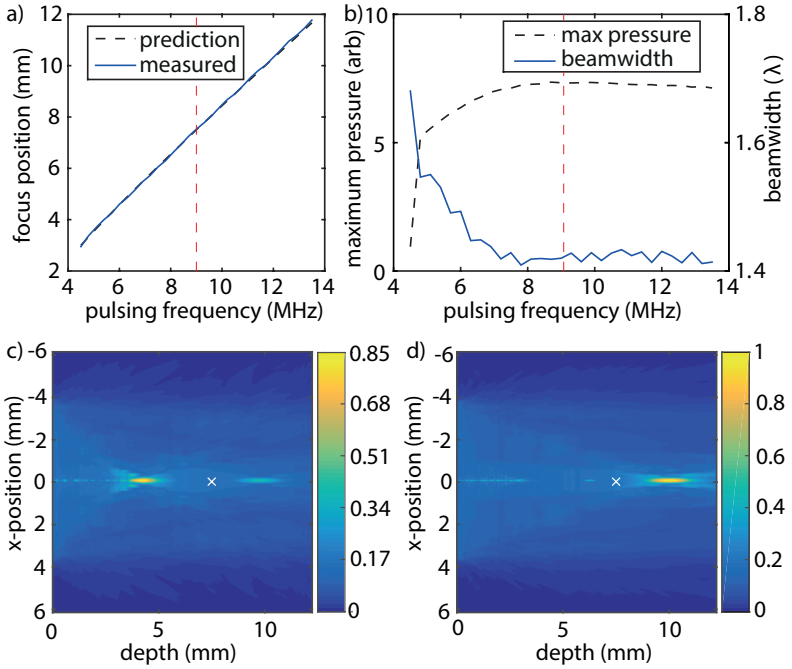

Figure 3.9: (a) Variation in focus depth with pulsing frequency, as measured in simulation and predicted by Eq. 3.18. The red dotted line indicates the design frequency. (b) Variation in peak pressure and focus width with pulsing frequency. Red dotted line indicates design frequency. (c) Maximum amplitude projection of the maximum pressure field for a pulsing frequency of 5.7 MHz. Pressure is normalised relative to maximum in (d). (d) Maximum amplitude projection of the maximum pressure field for a pulsing frequency of 11.7 MHz. The simulated data and the prediction show good agreement, and demonstrate it is possible to vary the focus depth over a wide range. Figure reprinted from [164] under CC BY 4.0 License.

relative to the pulsing frequency, and the magnitude was normalised. The focus position was recorded in each simulation as the depth from the hologram at which the peak pressure occurred. The peak pressure and lateral FWHM of the focus (beamwidth) were also evaluated.

Figure 3.9 (a) shows the measured variation in focus depth with pulsing frequency compared to the prediction from Eq. 3.18. The two are in close agreement. Figure 3.9 (b) shows the variation in peak pressure and beamwidth at the focus as a function of pulsing frequency. With increasing pulsing frequency compared to the design, the pressure at the focus and the beamwidth remain approximately constant. Conversely, as the pulsing frequency drops, the peak pressure lowers and the beamwidth increases. One possible cause is that the variation in the focus depth predicted by different rings becomes larger as the pulsing frequency decreases. At

much lower pulsing frequencies, the focus is observed to split into a series of smaller foci at different depths. This is accompanied by a large drop in peak pressure. Maximum amplitude projections of the maximum pressure recorded in each voxel of the simulation are shown in Fig. 3.9 (c) and (d) for pulsing frequencies of 5.7 and 11.7 MHz. These demonstrate the change in focus position.

3.8 Zone plate phase offset

A series of zone plates with different radii can be generated by varying the phase offset Φ in Eq. 3.4. It has been found for optical zone plates that the choice of this parameter can influence both the magnitude generated at the focus and the amplitude of the side lobes [170]. To test the impact of this parameter for acoustic holograms, a set of 60 equivalent holograms were generated by varying Φ between 0 and 2π. Each hologram had a design frequency of 9 MHz, a 1×1 cm aperture and a focus at 1 cm. The acoustic field generated by each hologram was simulated with the application of 2, 4 and 40 pulses. The measured variation in beamwidth and maximum pressure are shown in Fig. 3.10.

For two pulses, the variation in beamwidth between the best and worst phase offset is 66% (Fig. 3.10 (a)). As the number of pulses used increases, a similar trend is observed, however, the fractional variation between the best and worst case decreases. For 40 pulses, where the field has reached steady state, a variation of 9% occurs. From Fig. 3.10 (c)-(d) it can be seen that the beamwidth is maximised when the area of the hologram's central region is large and active, and is minimised for a small active central region. This is also the case, with small variations, for lower pulse numbers, before the field has reached steady state. The peak simulation pressure is also seen to change with phase offset, however, the measured variation is lower than that seen in the beamwidth (Fig. 3.10 (b)).

Slices through the field in the plane containing the focus for the simulations that generated the highest and lowest beamwidths at 2 and 40 pulses are shown in Fig. 3.11 (a)-(d). These show that there are significant differences in the field structure beyond just the beamwidth. For example, the pressure generated at the

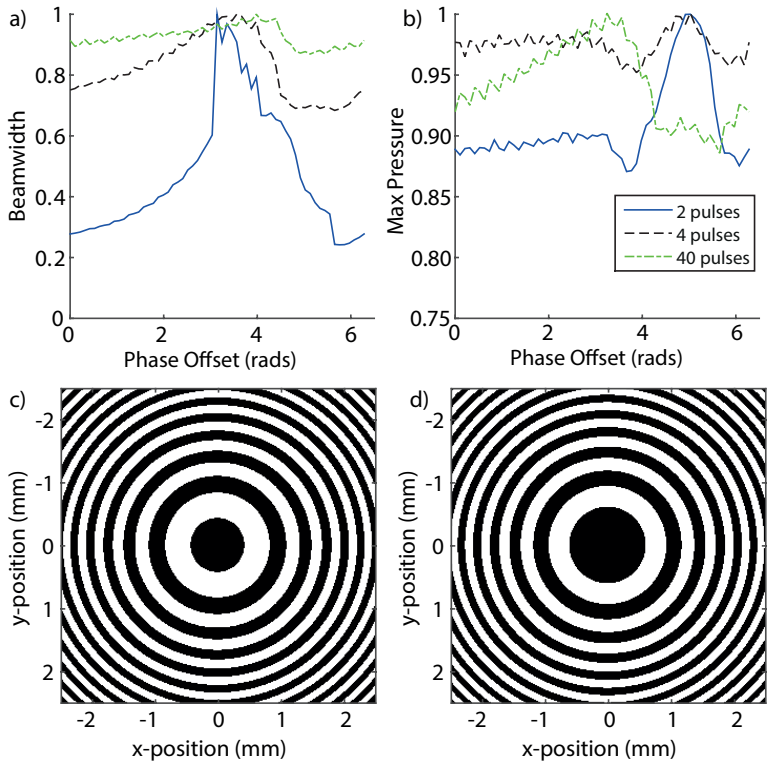

Figure 3.10: (a) Variation in focus diameter with choice of phase offset Φ for simulations using 2,4 and 40 pulses. Data is normalised against the maximum beamwidth. (b) Variation in maximum pressure at the focus with choice of phase offset for 2, 4 and 40 pulses. Data is normalised against overall maximum pressure. (c) Hologram that generated smallest beamwidth with 40 pulses. (d) Hologram that generated maximum beamwidth for 40 pulses. Choice of phase offset has significant effect on beamwidth, particularly for lower pulse numbers. Figure reprinted from [164] under CC BY 4.0 License.

second order focus in Fig. 3.11 (c) is significantly higher than that generated at the second order focus in (a). Conversely, the large region of pressure overlapping the focus visible in (a) is absent in (c). Even in the converged field, structural differences are apparent. For example, the pressure between the second order and design focus in (d) is significantly higher than the same region in (b), despite the pressure at each foci being approximately the same in both fields. This suggests that depending on the prioritised features of the acoustic field and the experimental parameters, the choice of phase offset is significant. Alternate metrics could be used to capture these features, for example, one that prioritised minimising the pressure in the near-field before the focus.

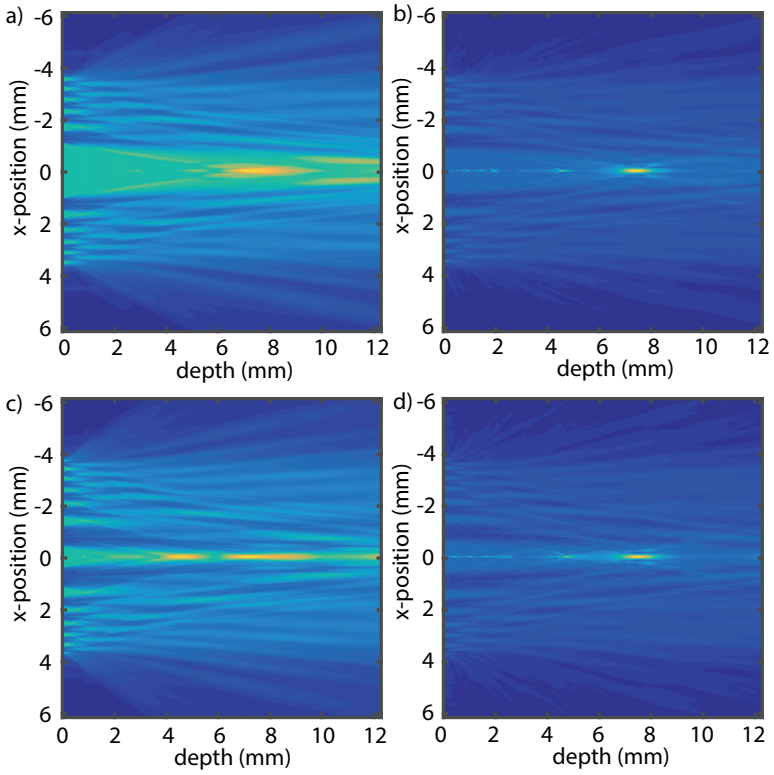

Figure 3.11: Maximum pressure measured in simulation on slice through central hologram axis for four different simulations. (a) Field of hologram that gave the highest beam width with two pulses. (b) Field of hologram that gave the highest beam width with 40 pulses. (c) Field of hologram that gave the lowest beam width with 2 pulses. (d) Field of hologram that gave lowest beam width with 40 pulses. Fields show significant feature differences in addition to the measured variation in beam width.

3.9 Experimental validation

3.9.1 Overview

Following the numerical tests, experimental measurements were carried out to validate the accuracy of the numerical model and demonstrate that acoustic foci can be generated by binary amplitude holograms. To do this, an optical absorber patterned with a calculated FZP was fabricated, and the acoustic field photoacoustically excited from this absorber by a pulsed laser was measured.

The absorber was fabricated on a $50\times30\times5$ mm transparent PMMA substrate. This was covered with a thin absorbing layer via aerosol deposition of an optically absorbing polymer composite (Super Satin, Plasti Kote, Valspar, US). A laser cutter (VLS4.60, Universal Laser Systems, US) was used to etch the calculated hologram

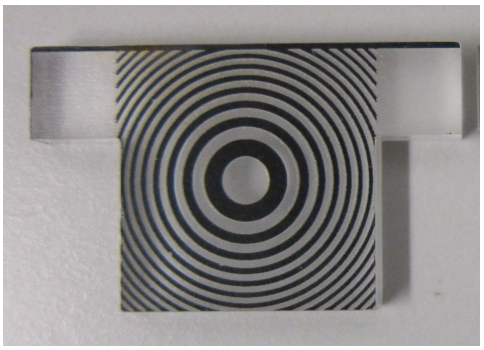

Figure 3.12: Single focus hologram fabricated using laser cutting on spray painted absorber. The hologram has dimensions of 3×3 cm, the focus is 3 cm from the hologram, and the design frequency is 3 MHz.

from the painted slide. The resolution of the laser cutter was first estimated by etching sets of lines with thicknesses and spacings of 50, 100, 150, 200, and 300 μm. Line spacings over 200 μm were accurately resolved, therefore 250 μm was used as the pixel size for the hologram. It was observed in this test that the laser cutter consistently removed more material than intended by the design pattern. This didn't significantly effect the fabrication of the zone plate, however, will be relevant in subsequent Chapters. The fabricated zone plate can be in Fig. 3.12. This had an aperture size of 30×30 mm, and was designed for a PRF of 3 MHz using a pixel spacing of 0.25 mm. The focus was 3 cm from the hologram aperture and placed on the central hologram axis.

A Litron Nano L-200 laser (Litron, Rugby, England) was used to generate ultrasound from the fabricated absorber. This consisted of two Q-switched lasers with spatially coincident beams that could be separately triggered. Both lasers had a wavelength of 1064 nm, a pulse length of 6 ns, and a repetition rate of 15 Hz.

3.9.2 Two-pulse measurements

An assumption made in the numerical experiments is that there is no variation in the magnitude or shape of the acoustic signal generated by a material when two or more optical pulses are incident on it in with a very small timing gap ($< \mu$s). The validity of this assumption for the absorbers in this experiment was tested using the Litron Nano L-200.

The experiments were carried out in a 22×40×28 cm tank filled with deionised

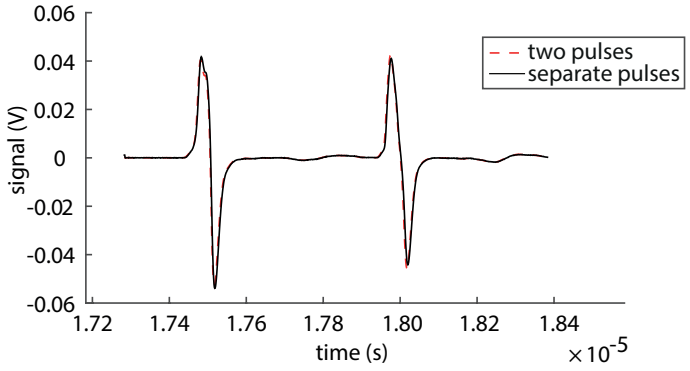

Figure 3.13: Voltage recorded on hydrophone from two pulse experiment with delay between pulses set to 0.5 microseconds. The red dotted line is the data with both pulses incident on the absorber. The black solid line is the combined data from both lasers used individually with the appropriate delay added back in. Figure reprinted from [164] under CC BY 4.0 License.

water. An unpatterned absorber fabricated as described in Sec. 3.9.1 was attached to the side. The laser was aligned with the absorber and the pulse energy of both lasers was set to 7 mJ. This output energy is significantly lower than that used for later experiments (Sec. 3.9.2), however, the spatial profile was narrower so the incident fluence was comparable. The generated acoustic signals were recorded using a focused PVDF transducer (Precision Acoustics, Dorchester, UK), which had a centre frequency of 25 MHz, and a focal length of approximately 24 mm. The hydrophone was suspended approximately 2 cm behind the absorber. Both lasers were externally triggered, and the delay between the two was varied between 0.02 and 2 μs. For each delay, three measurements were taken, one where both lasers were triggered together and two where each laser was used independently. Signals were recorded using a digital oscilloscope with a sampling rate of 400 MHz and 256 averages. The DC offset was removed from each dataset and the single laser measurements were summed accounting for the trigger delay. Each summed dataset was then overlaid with the two pulse measurement for comparison.

Figure 3.13 shows the recorded voltage from the two pulse and combined single pulse measurements where the trigger delay was 0.5 μs. It can be seen that there is no significant variation in magnitude in the second pulse between the two measurements (the correlation coefficient between the two signals was 0.989). This was found to be the case for separations down to 20 ns, on the order of the pulse

width, so the assumptions of the numerical model are valid up to frequencies of at least 50 MHz using conventional Q-switched lasers.

3.9.3 Field measurements

Next, measurements were made of the acoustic field generated from the fabricated zone plate. These were performed in a 40×40×60 cm tank with a two axis computer controlled positioning system (Precision Acoustics, Dorchester, UK). A calibrated PVDF needle hydrophone with a 0.2 mm active element was used to record the pressure (Precision Acoustics, Dorchester, UK). One set up for the experiments is shown in Fig. 3.14 (note this photo was taken from a separate experiment with a slightly altered set-up using a fibre coupled Q-switched laser and membrane hydrophone). The laser was positioned 50 cm from the tank edge. A concave lens with a 40 mm focal length was placed 30 cm from the laser to expand the beam. The radius of the beam at the edge of the tank was ∼30 mm. The absorber was suspended 30 mm inside the edge of the tank, and was aligned with the laser beam. The needle hydrophone was suspended ∼7 mm away from the absorber and aligned with its centre.

Since the Litron was limited to generating two pulses at a high repetition rate, two sets of measurements were made. For the first, one of the two lasers within the Litron Nano L-200 was used to excite the absorbers and additional pulses were added to the experimental data in post-processing. This was used to evaluate the steady state field. For the second, both lasers were triggered using the approach described in Sec. 3.9.1. The temporal spacing of the pulses was matched to the design frequency of the zone plate. This was used to validate the results of Sec. 3.6 that showed that a focus could be created using a minimum of two optical pulses.

3.9.3.1 Single pulse

For the single pulse measurements, signals were recorded over a 30×30 mm plane parallel to the absorber using a step size of 0.3 mm. The pulse energy of the Litron was set to ∼300 mJ. At each position, time-domain signals were recorded using a digital oscilloscope with a sampling rate of 400 MHz, taking 6 averages. The

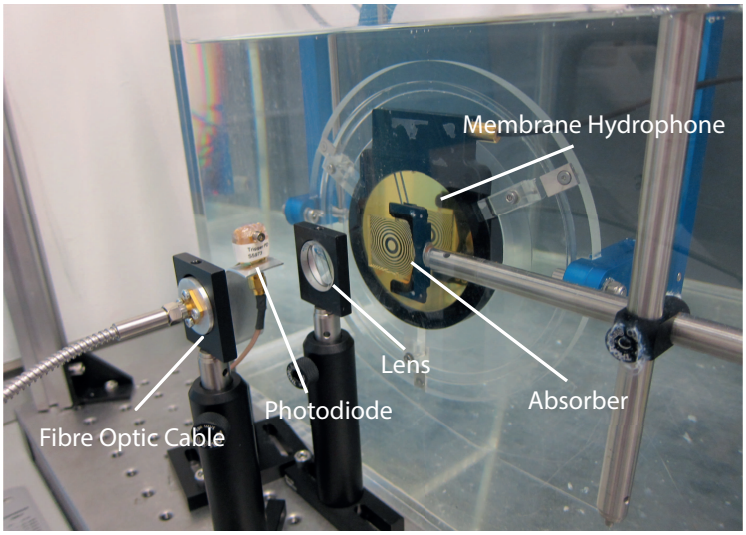

Figure 3.14: Photograph of experimental set-up. Figure reprinted from [171] © IEEE 2014.

dataset was then spatially up-sampled to a grid spacing of 0.0375 mm to support frequencies up to 20 MHz in the simulation and low pass filtered at 15 MHz to give a minimum PPW of 2.666. The filtering was used to limit the grid size needed for the projection. The data was also temporally down-sampled to a step size of 7 ns. The measurements were then additionally high pass filtered at 100 kHz to eliminate slowly varying background noise.

After processing the data, 15 additional pulses were created in the time domain signals. This was done by adding the data recorded at each point to itself translated by multiples of 0.33 μs. This temporal spacing corresponds to a 3 MHz pulsing frequency. After this, the data was forward and back-propagated by 100 and 20 mm, respectively, using the k-Wave toolbox to recover the 3-D field.

The resulting field was compared to a numerical simulation run using the same hologram, grid, pulse number, and cutoff frequency. The maximum amplitude projection of the 3-D field from both the experimental data and the simulation data is shown in Fig. 3.15 (a)-(b). The two show good agreement. Both acoustic fields are focused at 4 distinct points along the central hologram axis, and each foci occurs at approximately the same depth in both sets of data.

The first foci in the field is the design focus, the second is the 3rd order focus of the fourth harmonic, and the third and fourth are generated by the second and third harmonics, respectively. This is confirmed by the spectra measured at each

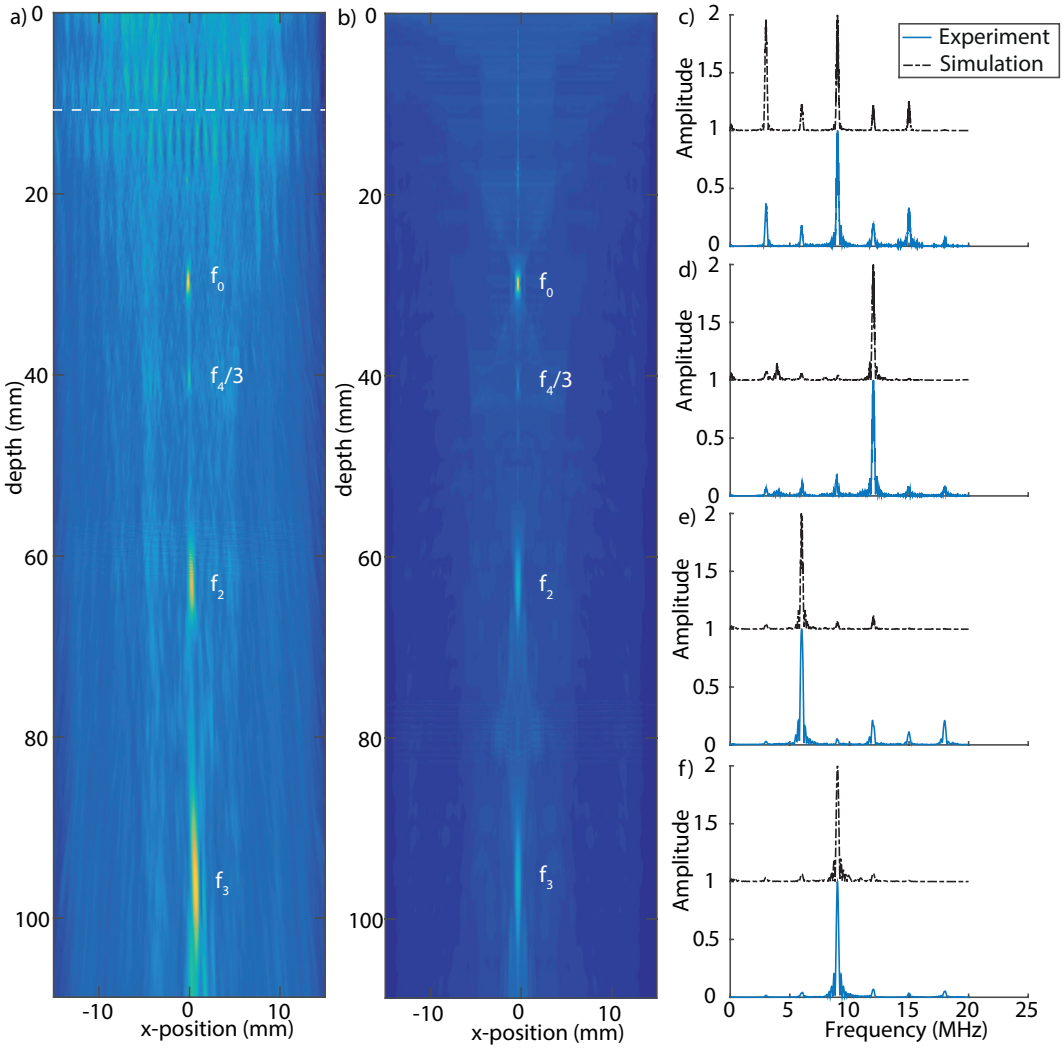

Figure 3.15: Comparison between the maximum amplitude projections of the maximum pressure fields produced from the single focus hologram shown in Fig. 3.12 in (a) experiment and (b) simulation. In both projections, the hologram is located along the top image axis. Positions of different harmonics foci are labelled. (c)-(f) Spectra measured at each of the focal points labelled in (a) and (b). Each spectra is normalised to its maximum. The simulation spectrum is offset by 1 in each case relative to the experiment. The dashed line in (a) indicates the measurement plane. Figure reprinted from [164] under CC BY 4.0 License.

focus (Fig. 3.15 (c)-(f)). The frequency with the maximum amplitude in the spectra matches the predicted harmonic for the 2nd, 3rd and 4th foci in both the simulation and experimental data (Fig. 3.15 (d)-(f)). At the design focus, the third harmonic has a higher amplitude than the fundamental in both spectra (Fig. 3.15 (c)). This is because, as discussed in Sec. 3.4, the third order focus for the 3rd harmonic also occurs at this point.

The peak pressure measured at the design focus in the experiment was 187.3 kPa. The maximum pressure in the data that was propagated was 79.8 kPa giving a focal gain of 2.347. However, this represents only the spectral content up to 15 MHz. The peak pressure in the experimental data up to the 30 MHz to which the needle hydrophone was calibrated was 326 kPa. Assuming equivalent focal gain this gives a pressure of 764.17 kPa at the focus. This could be further improved by the use of an optimised absorber instead of black spray paint [92]. The pressure achieved is over an order of magnitude higher than pressures reported in previous works utilising longer modulated pulses as the optical source [135]. Pulsed lasers could thus be useful in the future for the generation multi-focal or patterned fields in applications that require higher pressures.

3.9.3.2 Two-pulses

For the second set of measurements, both lasers were used and their energies were maximised. This corresponded to a pulse energy of $\sim$300 mJ for the first laser, and $\sim$220 mJ for the second. The time domain signals were recorded over a 32×32 mm area using a larger step size of 0.4 mm. The measured signals were then spatially up-sampled to a grid spacing of 0.0667 mm to support frequencies of up to 12 MHz, temporally down sampled to 10 ns and bandpass filtered between 0.1-10 MHz. After this, the data was forward and back-propagated by 40 and 20 mm with k-Wave assuming linear acoustic propagation. The propagated data was again compared against a simulation run on the same grid, for the same hologram, applying 2 pulses. The maximum pressure generated in a plane of the field containing the focus in both the experimental and simulated data is shown in Fig. 3.16.

Figure 3.16 (a) shows that a focus has clearly been generated along the main axis, at the target depth, in the experimental data. This confirms the observation of Sec. 3.6 that a focus can be generated with a minimum of two optical pulses. The simulated data also shows the generation of a focus, and there are notable similarities between both fields, particularly close to the source plane. However, the background is significantly higher in the experimental data and the focus is elongated. This discrepancy is the result of apodisation of the hologram introduced by

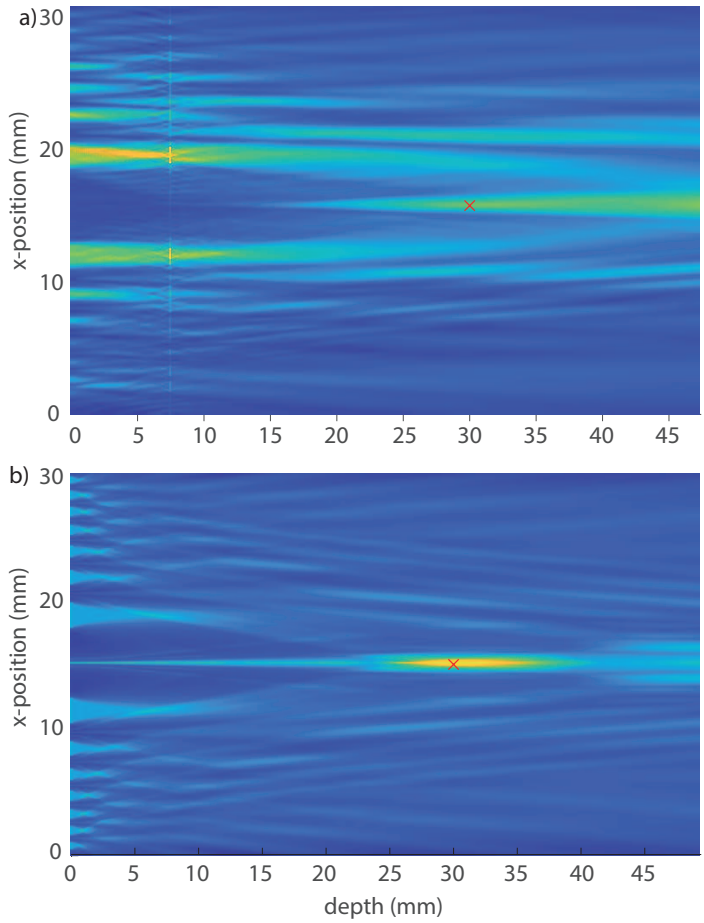

Figure 3.16: Slice through plane containing focus in (x-z) plane for two pulse measurement in (a) experimental data and (b) simulated data. For both datasets the red-cross marks the target focal position.

the Gaussian profile of the optical pulse. The difference in the pulse energies of the two lasers also impacted the data. Investigating numerically, the principal effect was found to be decrease in the ratio of the focal pressure to the background pressure. Matching the pulse energies would therefore improve the SNR, at the expense of focal gain.

In practice, for two pulses, temporal coincidence at the focus will only occur for neighbouring rings, making the large number of zones included on the fabricated hologram unnecessary. This observation suggests that the structure of the zone plate could potentially be optimised for different pulse numbers (i.e., identifying the optimal placement and width of two zones when using two pulses).

3.10 Conclusion

This Chapter has demonstrated, for the first time, that optically generated acoustic fields can be focused at a single point in 3-D by using optical absorbers spatially patterned to form Fresnel zone plates and excited by a train of narrowly spaced optical pulses. This expands the work of Meyer et al [135], and demonstrates that a rapidly pulsed optical source can also be used for the holographic control of optoacoustic fields. This has potential applications across a range of areas. Most immediately for imaging (using a small number of pulses) by steering a foci in 3-D, however, also, potentially, for particle manipulation or neuro-stimulation.

A detailed numerical investigation was carried out to establish the impact of different properties of the source on the characteristics of the resulting acoustic fields. This found that a focus can be generated using a minimum of two optical pulses, the focal depth can be controlled by varying the PRF of the pulse train, and that the harmonics introduced to the spectrum by the use of short pulses generate a number of additional foci to the field. A fraction of these foci were found to spatially coincide with the design focus, reducing the focal volume compared to a monochromatic source. The effect of the zone plate phase was also investigated. It was found that for a small number of pulses, having a large central transmitting zone greatly increases the focal width. A modified design method taking advantage of the harmonics present in the spectrum was also introduced. This showed that by altering the width of rings on the hologram, the focal gain, SNR, and focal volume can all be improved when harmonics are present. Finally, the technique was validated experimentally with two sets of measurements from a spatially patterned optical absorber. Both sets of measurements agreed well with k-Wave simulations of the equivalent field, supporting the numerical experiments.

There are several future lines of inquiry along which the work in this Chapter could be developed. First, the approach for generating the hologram by patterning an optical absorber is inflexible. The field is largely fixed and can only be varied by changing the pulsing frequency. As discussed in Chapter 2, the same effect can be achieved by modulating the incident fluence with a SLM. Employing an SLM

and a uniform absorber would enable the generated hologram to be dynamically varied. Similarly, the post-processing addition of pulses in the experimental data is also problematic, as the focus wasn't actually generated in the acoustic field. Again as discussed in Chapter 2, this can be overcome by using a greater number of Q-switch lasers, or, by splitting a single pulse into optical fibres of varying lengths. Implementing both changes would enable a full, dynamic system for optoacoustic holography to be realised. The design of the single focus holograms could also be further optimised. As discussed in Sec. 3.9.2.2, for a low number of pulses there is only temporal coincidence at the focus for a small fraction of the aperture, which could be incorporated into the hologram calculation. Similarly, Sec. 3.5 showed that the SNR can be substantially improved by modifying the zone structure accounting for harmonics. This concept could be applied more generally with an optimisation approach that directly perturbed the zone structure with the aim of altering certain properties of the field, for example, reducing the peak or average near field pressure.

Chapter 4

Multi-focal optoacoustic holography

4.1 Introduction

The previous Chapter investigated the generation of a single acoustic focus using optoacoustic holography, that is, binary amplitude holograms combined with trains of regularly spaced optical pulses. This Chapter investigates the more general problem of creating arbitrary distributions of pressure with the same approach. Firstly, in Sec. 4.2, a direct search algorithm is described for the calculation of binary amplitude holograms to generate arbitrary fields. In Sec. 4.3, the ability of this approach to generate multi-focal and continuous patterns is investigated numerically. In Sec. 4.4, the approach is experimentally validated. Section 4.5 then introduces an adjusted forward model able to account for harmonics in the spectrum, and in Sec. 4.6, this new model is experimentally validated.

The journal article in [164] has been modified and adapted to form parts of this chapter, with reprint permission under CC BY 4.0 License.

4.2 Direct search algorithm

4.2.1 Introduction

In Chapter 2, several algorithms were discussed for the calculation of holograms with different constraints. As the experimental approach in this Chapter is limited to binary amplitude modulation, the direct binary search approach outlined in Sec. 2.1.5.4 was chosen for this work. In this case, quantisation is directly included in

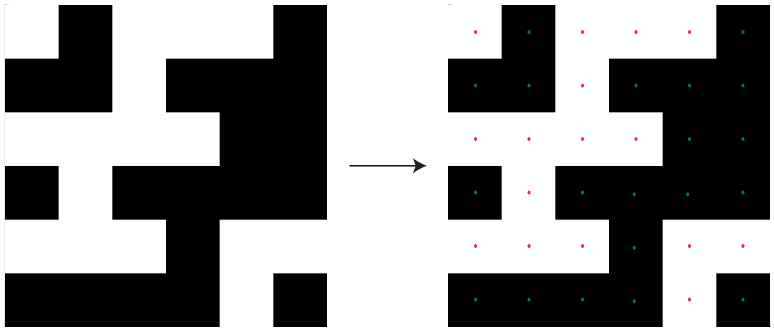

Figure 4.1: Approximation used to discretise holograms for calculation. Green is a transmitting point, red is not.

the forward model. Additionally, the major disadvantage of the approach (the computation time) is not an issue, as dynamic variation of the patterns is not considered.

The goal of the direct search algorithm is to find a 2-D binary pattern $H(x,y)$ that, when excited at a frequency f, generates a target distribution of pressure $p_{\text{tar}}(x,y,z)$. To solve this, the binary pattern was first discretised such that a square L×L hologram $H(x,y)$ is approximated by a N×N set of point sources with spacing $\Delta = \frac{L}{N}$, centred on the coordinates (x_i, y_j). Each point source (x_i, y_j) can be switched on or off, and the current state of the pattern is denoted H_{ij}. This approximation is shown in Fig. 4.1.

As discussed in Sec. 2.1.5.4, to solve this optimisation, the hologram is initialised in an arbitrary state, which is then perturbed. The change in solution quality or cost C caused by this perturbation is evaluated and the change kept if the solution is improved. There are a variety of functions that can be used to evaluate this cost. Many works use cost functions that compare the pressure over the whole field to a target pressure distribution, for example

$$C = \iiint |p(x,y,z) - p_{tar}(x,y,z)| dxdydz. \tag{4.1}$$

However, these are computationally inefficient when only a certain area of the field is of interest. They also require either normalisation or an estimate of the pressure achievable in the final solution to calculate the target pressure. For this work, to overcome these limitations a state-based cost function was used. This was first introduced by Clarke for the design of binary optical phase holograms [172]. Rather

than looking to minimise the difference across the whole field $p(x,y,z)$, the target field is decomposed into a set of points r_k. The state-based cost function aims simply to maximise the pressure $p(r_k)$ at this set of target points. To do this, the cost C is given by

$$C = -\overline{|p(r_k)|} + \alpha \text{std}(|p(r_k)|), \tag{4.2}$$

where $\overline{|p(r_k)|}$ is the average magnitude of the complex pressure at the target points, $\text{std}(|p(r_k)|)$ is the standard deviation of the pressure magnitude over the target points, and α is a factor weighting the two terms. The first term in Eq. 4.2 rewards increases in the peak pressure across the target points, and the second penalises variations in pressure. The advantage of this cost function is that it requires no estimate of the pressure prior to starting the optimisation, and it limits the points at which the pressure has to be evaluated.

To calculate the cost, a forward model was used to relate the hologram state to the complex pressure at the target points. This pressure was calculated assuming the hologram H_{ij} was placed in the $z = 0$ plane, where

$$p(r_k) = \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{H_{ij}}{((x_k - x_i)^2 + (y_k - y_j)^2 + z_k^2)^{0.5}} e^{ik((x_k-x_i)^2+(y_k-y_j)^2+z_k^2)^{0.5}}. \tag{4.3}$$

This approximates the hologram as a collection of point sources, as described above.

4.2.2 Optimisation algorithm parameters

4.2.2.1 Introduction

The performance of the optimisation and the result to which it converges is dependent on a variety of parameters. The major ones are: the starting condition (i.e., the state in which the hologram is initialised), the search procedure (i.e., the selection process for pixels), the stopping criteria, and the weighting of the two cost function terms. Different strategies were investigated for each of these parameters.

A standard problem was used for the investigations where the target pattern was a 'G'. This was chosen as it consists of both linear and curved components, and doesn't have any symmetry. It is also reflective of the types of patterns that are of

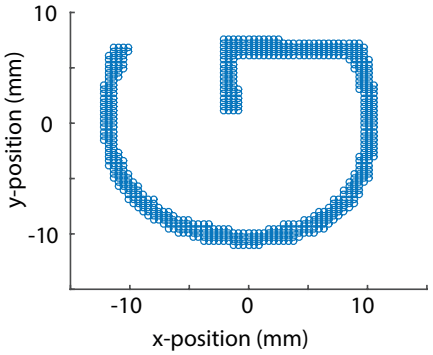

Figure 4.2: Distribution of points used for representing 'G' in standard test of hologram search parameters.

interest for this thesis (i.e., continuous patterns of pressure). The 'G' was discretised into a set of points and scaled as shown in Fig. 4.2. Each hologram was calculated for a 7.5×7.5 cm aperture, using a pixel spacing of 0.5 mm and a design frequency of 1 MHz. The target plane for the 'G' was 7.5 cm from the hologram aperture. The acoustic field generated by each hologram was simulated using a $192\times192\times192$ grid, with a grid spacing of 0.5 mm and a maximum supported frequency of 1.5 MHz. In each simulation, the maximum pressure generated at each target point and the variation in maximum pressure over the target points was recorded. These were input into Eq. 4.2 to evaluate the hologram cost. The time taken for the optimisation to run was also recorded for each implementation of the algorithm.

4.2.2.2 Starting condition

The starting condition for the hologram can impact the rate at which the optimisation converges and whether it gets trapped in bad, but locally optimal, solutions. Four conditions were considered: a random binary pattern, all pixels on, all pixels off, and an 'informed' hologram. Each condition is illustrated in Fig. 4.3. The informed hologram was calculated by placing a point source at each of the target points r_k and thresholding the sum of the pressure generated by each source across the hologram aperture. The optimisation was run 40 times for each of these starting conditions recording the time taken for the optimisation to converge and the final value of the cost function. The results are shown in Fig. 4.4.

It can be seen that a starting condition of all-off produces the worst holograms

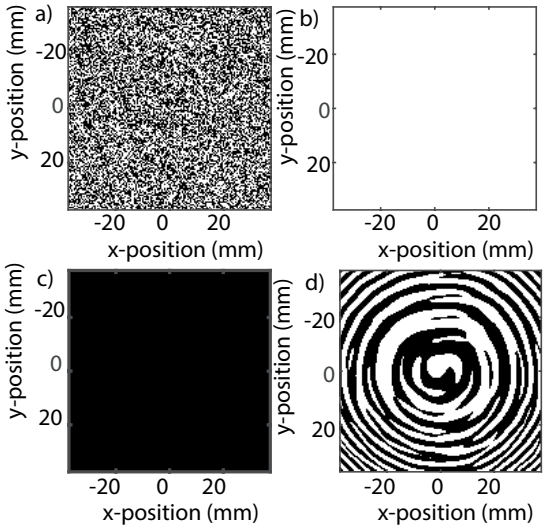

Figure 4.3: Four starting conditions compared for the optimisation algorithm. (a) Random binary pattern. (b) All pixels off. (c) All pixels on. (d) Output of ray-based algorithm for the same set of target points.

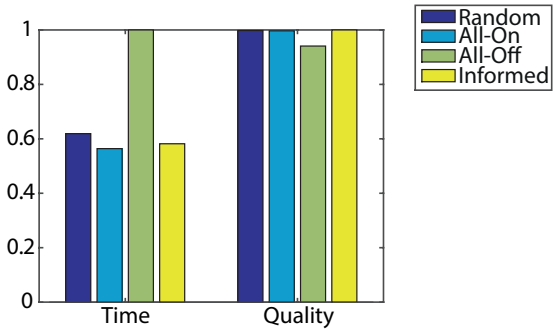

Figure 4.4: Bar chart showing measured variation in hologram quality, and time taken for optimisation to converge for each of the starting conditions outlined in Fig. 4.3.

on average, but that there are minimal differences in the quality of the solution for the other three starting conditions. The time taken for each to converge differs, with all-on being slightly better. However, this difference, evaluated using a t-test, is not statistically significant so random seeding was chosen for this Chapter to provide more variation in the starting conditions for the algorithm.

4.2.2.3 Search parameters

The method used to select pixels to test can also impact the rate at which the optimisation converges and the result to which it converges. There are a range of search methods that could be used, and four were considered.

1. Random exhaustive - each pixel on the hologram is tested before repeating,

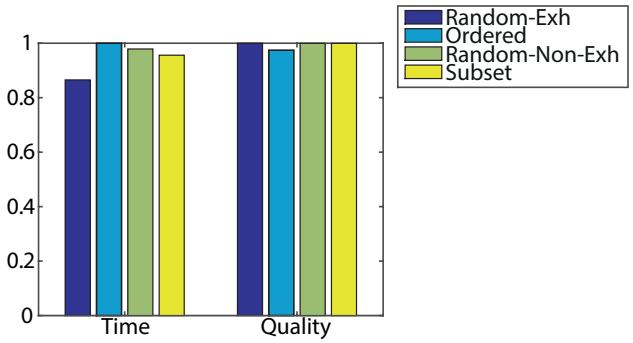

Figure 4.5: Bar chart showing variation in average time taken for algorithm to converge and average quality for four different search types. Quality is measured by the value of the cost function given in Eq. 4.3

but the ordering is random and changed between iterations.

2. Random non-exhaustive - pixels are tested randomly with no consideration to previous choice.

3. Ordered - each pixel tested in order and each tested once before repeating.

4. Subset - a subset of pixels are flipped simultaneously, starting the size of the subset at 20 and decreasing it by 1 after each iteration.

Each of these was tested running the optimisation 40 times. It can be seen from Fig. 4.5, that again there is minimal difference in solution quality between three of the search methods (random-exhaustive, random, and subset). However, the random-exhaustive converges most rapidly so was selected as the search method for the remainder of the work.

4.2.2.4 Cost weighting

The weighting term α in the cost function in Eq. 4.2 determines the degree to which variations in pressure at the target points are penalised by the optimisation versus increases in the pressure generated. An appropriate value to select for this parameter was investigated. To do this, the value of α was varied between 0-4 in increments of 0.1, a hologram was generated for each value, and the acoustic field generated by each hologram simulated. The change in the average peak pressure and in the variation in peak pressure over the target points with cost weighting is shown in

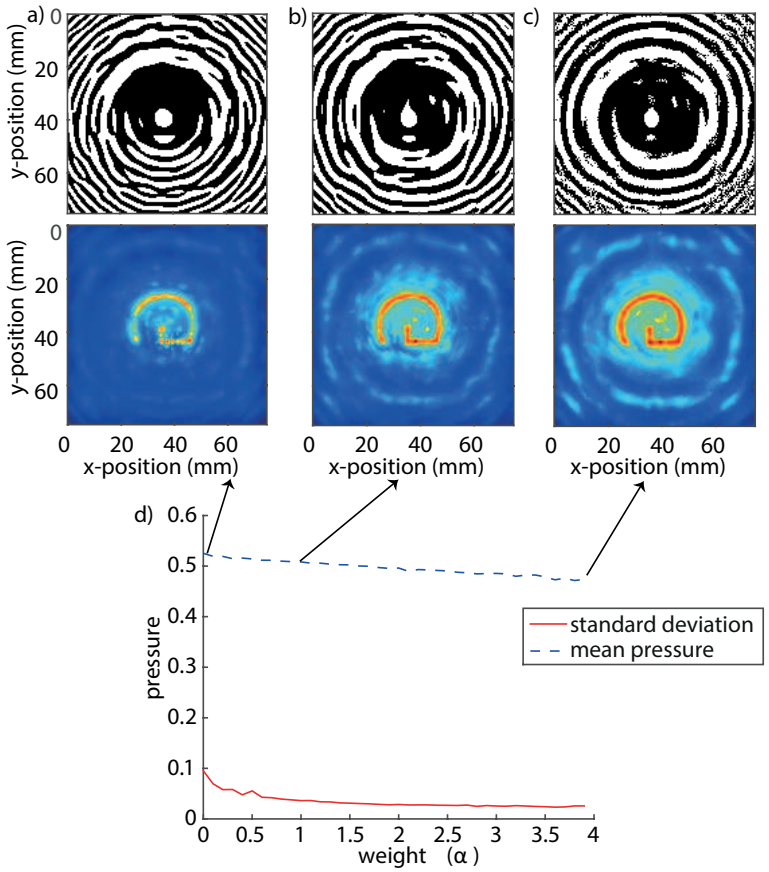

Figure 4.6: (a)-(c) Holograms generated by optimisation and the corresponding acoustic field in the target plane generated by these holograms for 3 different target weightings. (a) $\alpha = 0$. (b) $\alpha = 1$. (c) $\alpha = 3$. (d) Variation in average maximum pressure at target points, and variation in maximum pressure at target points, with α.

Fig. 4.6 (d). The hologram and acoustic field in the target plane of the pattern for 3 different cost weightings are shown in Fig. 4.6 (a)-(c).

For cost weightings between 0-0.5, the average pressure over the target points is largest. However, the variation in pressure is large. The effect of this is illustrated by Fig. 4.6 (a), in which it is difficult to resolve the outline of the G. For a cost weighting of 4, the G can be clearly resolved (Fig. 4.6 (c)). However, the background pressure is greater and the width of the target distribution appears wider than for lower cost weightings (Fig. 4.6 (b)). Additionally, the hologram generated using this cost weighting shows regions with a speckly structure, which are harder to accurately fabricate due to the tendency of the laser cutter to remove too much material. A cost weighting of 1 shows no speckly structure in the hologram and the

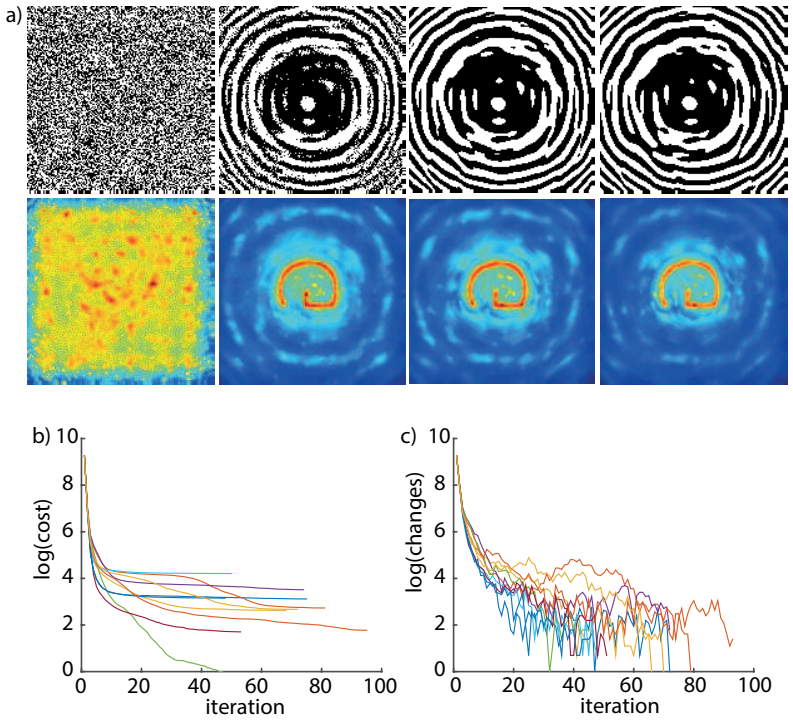

Figure 4.7: (a) Hologram state and acoustic field generated by the current hologram state for 4 different points in the optimisation. First: starting point, Second: 1 iteration, Third: 15 iterations, Fourth: 72 iterations. (b) Log of the cost at the end of each iteration for 10 different implementations of the optimisation. Each cost was scaled such that the lowest over all implementations was 1. (c) Log of the number of changes made in the previous iteration until convergence for 10 different implementations of the optimisation.

target pattern can be clearly resolved (Fig. 4.6 (b)). Values of α between 1-2 were therefore used in the remainder of this Chapter.

4.2.2.5 Stopping criteria

The stopping criteria also has an impact on the performance of the optimisation. If the optimisation is stopped prematurely then the performance of the holograms will be poor. If it is too strict then computational effort is wasted. The convergence of the hologram over time was investigated. To do this, the optimisation was run tracking the change in cost between iterations, the number of changes made in an iteration, and the state of the hologram itself. The results of 10 optimisation runs are shown in Fig. 4.7.

Figure 4.7 (b)-(c) show that the bulk of the changes and the majority of the decrease in cost occur within a single iteration (i.e., in the first test for each pixel).

The second hologram and plane displayed in Fig. 4.7 (a) were taken after a single iteration and show that the hologram is already recognisably generating the target pattern. The optimisation then spends a large number of iterations making small changes ($\sim$50) to the hologram, which gradually reduce the cost and improve the target pressure.

The criteria imposed for terminating the optimisation for this work was set to when the number of changes in an iteration fell below 0.5% of the pixels on the hologram. This corresponds to 112 pixels for a 150×150 hologram. Less strict criteria could have been imposed. However, the optimisation can be evaluated relatively quickly for the problems in this work, and computational efficiency was not essential.

4.2.2.6 Converged solutions

The number of states explored in an instance of the optimisation is $\sim N^2$ while the number of possible states for the hologram is 2^{N^2}. Given the direct search method only accepts locally optimal changes, it is clear that each 'run' of the algorithm converges to a local minimum (the fraction of states explored is $1/10^{26}$ for $N = 10$). The variation in the properties of these local minima was investigated by running the optimisation 100 times with the same starting condition and comparing the resulting holograms.

Figure 4.8 (a) shows a histogram of the fractional variation in the structure of each calculated hologram. This shows that in each case the optimisation converged to a different result, with the smallest difference between two solutions $\sim 7\%$. Given the size of solution space and the random approach used for pixel selection, this is expected. However, the fractional variation is on average $\sim 15\%$, so there are significant similarities in the generated holograms. This is confirmed by Fig. 4.8 (c) which shows a sum over the structure of all 100 holograms (i.e., the number of times each pixel is 'on' in a converged solution). This shows that there is significant cross-over between the solutions, where some pixels are 'on' in every converged hologram.

Overlaying the lateral position of the target distribution (Fig. 4.8 (d)) shows

that the target distribution is surrounded in each case by a broad central 'on' region. The reason for this is the cost function aims to maximise the peak pressure across the target points. Due to spherical spreading, points on the aperture directly coincident with the target points contribute a greater pressure, and the optimisation is therefore biased to these locations. The variation between holograms arises principally from small perturbations in the phase across the target distribution, which alter the shape of the fringes.

Figure 4.8 (b) shows the normalised variation in average peak pressure across the target points, and the normalised variation in the standard deviation in peak pressure across the target points. This shows there is no significant different between the different solutions ($\sim 2.5\%$ at most), and thus the algorithm consistently converges to good local minima.

4.2.3 Alternate cost functions

A major advantage of the direct search approach is that it offers significant flexibility over the resulting holograms. Additional terms can be added to the cost function to control the acoustic field in different ways. For example, the pressure at different points can be constrained to different relative values by introducing a weighting vector 'w' specifying the desired pressure at different points into the cost function. This is done using the following equation

$$C = -\left(\frac{1}{K}\sum_{k=1}^{K}\frac{1}{w_k}|p(r_k)|\right) + \alpha \mathrm{std}(\frac{1}{w_k}|p(r_k)|) \tag{4.4}$$

Here w_k is the weighting of the k^{th} point, $p(r_k)$ is the pressure at the k^{th} point, K is the overall number of points, and std is the standard deviation of the weighted pressures. Two different examples of using this cost function to generate varying peak pressures are shown in Fig. 4.9. The two tests used were: a line focus where the pressure linearly increased from the edge to double the value at the centre, and a line focus where the pressure from the edge linearly decreased to half the value at the centre. It can be seen from Fig. 4.9 (b) and (d) that the generated pressure accurately replicates the targeted pressure.

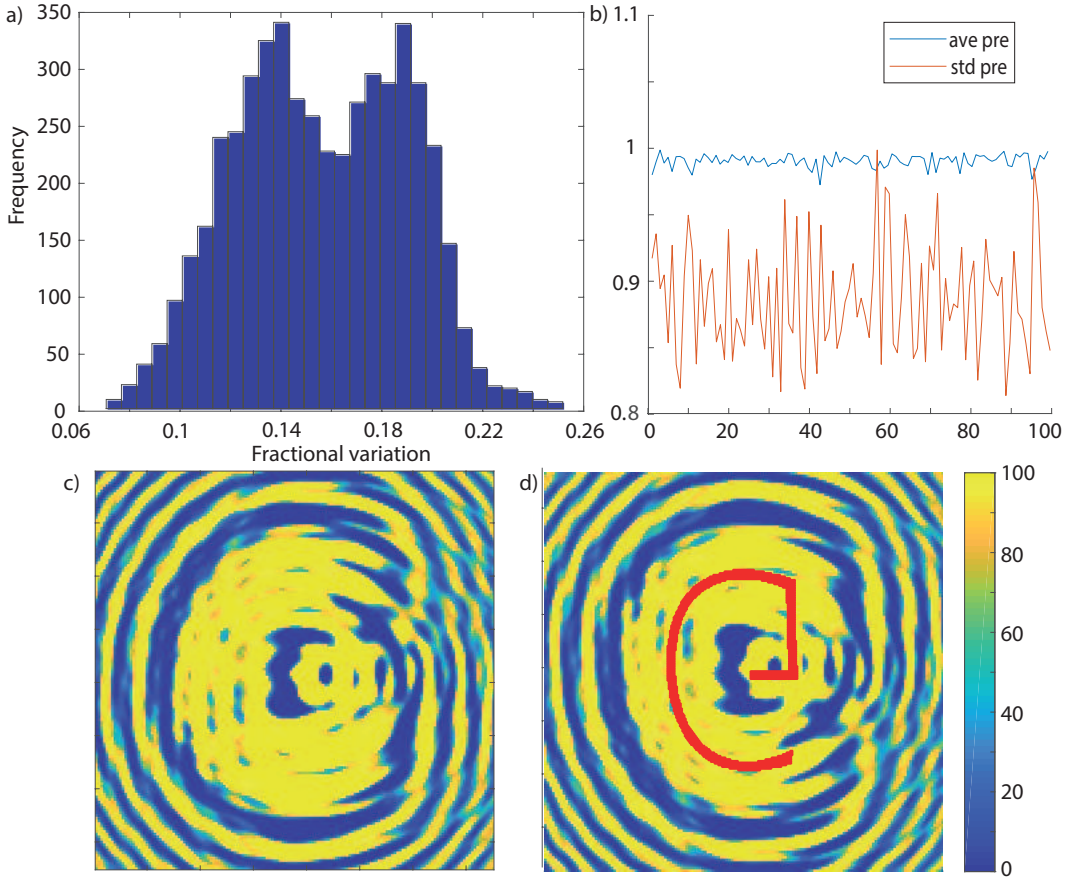

Figure 4.8: (a) Absolute fractional difference between converged outputs of optimisation. (b) Variation in average peak pressure across converged solutions normalised (blue), and variation in standard deviation of peak pressure across converged solutions normalised (red). (c) Sum of the converged solutions. (d) Sum of the converged solutions with lateral position of points overlaid on hologram.

Another useful adjustment that can be made is the addition of an L1 norm. This is based on the observation that the field generated by the hologram can be separated into two components: a plane wave with a magnitude proportional to the average value of all pixels on the hologram, and a portion that generates pressure at the desired target points. The addition of an L1 norm allows the optimisation to reduce the first component without significantly reducing the second component. This can be introduced to the cost function as follows

$$C = -|\bar{p}| + \alpha\sigma + \beta\bar{H}. \tag{4.5}$$

Here β is a weighting for the third term, and $\bar{H}$ is the average value for a pixel in

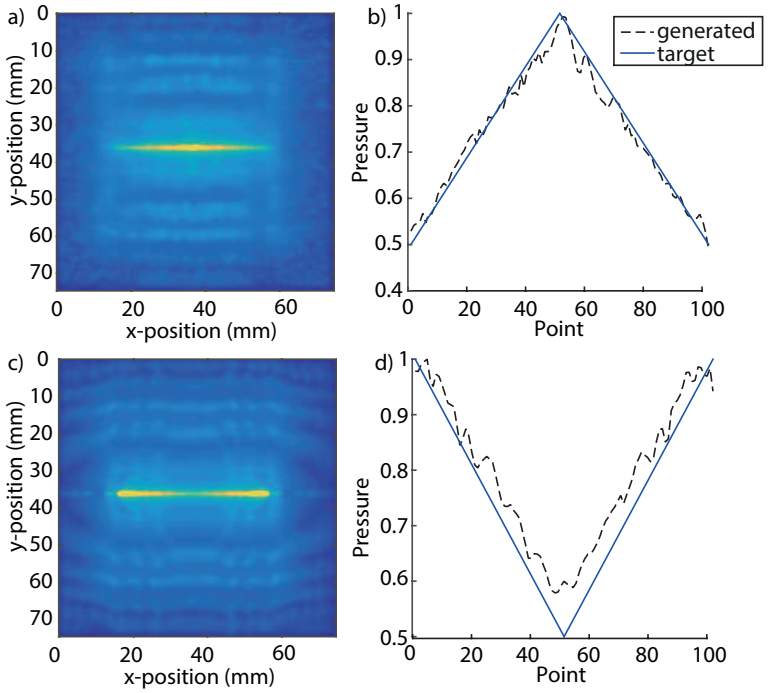

Figure 4.9: (a) Acoustic field in target plane of line-focus hologram where target pressure at the centre of the line is double that at the edge. (b) Measured peak pressure at each target point, and target peak pressure. (c) Acoustic field in target plane of line-focus hologram where target pressure at the centre of the line is half that at the edge. (d) Measured peak pressure at each target point, and target peak pressure. Holograms were calculated for a 7.5×7.5 cm aperture, the target plane was 7.5 cm from the hologram, and the design frequency was 1 MHz. Simulations were carried out on a 192×192×192 grid with a spacing of 0.5 mm.

the hologram. This cost function was tested on the problem outlined in Sec. 4.2.2.1 Holograms were generated varying the value of β between 0 and 16000, keeping α constant at 1. The average maximum pressure over the target plane (i.e., the background), the average pressure at the target points, and the variation in pressure at each of the target points, were then calculated for each hologram.

Increasing β causes the pressure at both the target points and the background to reduce. However, the pressure over the background decreases by a greater amount so the signal to noise is improved (Fig. 4.10 (e)). This is clear from Fig. 4.10 (c) and (d), where the background is visibly less intense in the L1 norm case. Observing the structure of the holograms (Fig. 4.10 (a) and (b)), it can be seen that the L1 term effectively acts as a threshold on the hologram pixels, such that only pixels that contribute significantly to the target points are included on the pattern. Depending

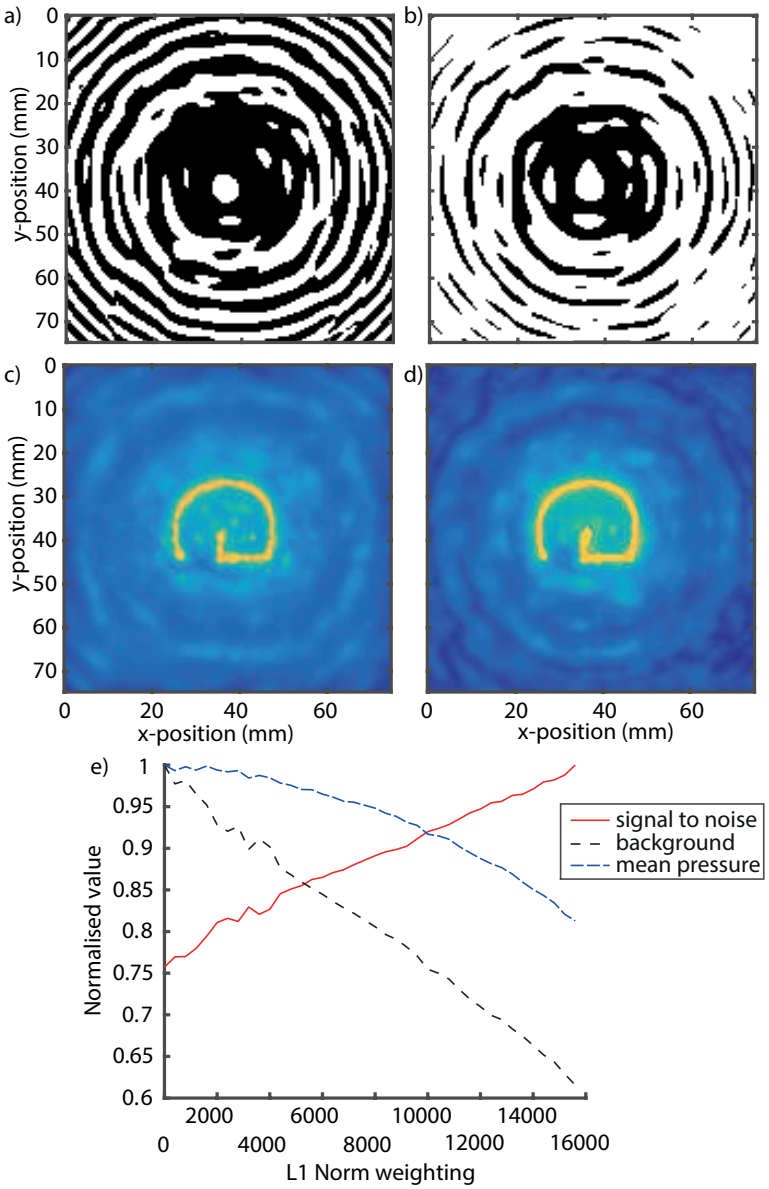

Figure 4.10: (a) Hologram generated for $\beta = 0$. (b) Hologram generated for $\beta = 16000$. (c) Acoustic field generated by hologram in (a) in target plane of 'G'. (d) Acoustic field generated by hologram in (b) in target plane of 'G'. (e) Variation in average pressure at target points (mean pressure), average pressure in target plane (background), and ratio of both values with value of β (signal to noise)

on the goal, the L1 norm could be a useful addition. Additionally, as with the harmonic zone plates, the lower number of active pixels on the hologram could allow for greater energy density when spatial modulation of the optical source is used to produce the pattern.

4.2.4 Simulated annealing

As discussed in 4.2.2.5, the direct search approach converges to the first local minima it encounters. The small variation in the quality of different converged solutions suggests that this is not a concern, however, a simulated annealing approach was tested to validate that better results couldn't be achieved using a more thorough search method.

As outlined in Sec. 2.1.5.4 simulated annealing is similar to direct binary search except a finite chance of accepting 'bad' changes is introduced. The probability of accepting the bad change is calculated using

$$P = \exp\left(\frac{-\Delta C_{\mathrm{i}}}{T}\right), \tag{4.6}$$

where ΔC_{i} is the change in cost associated with flipping a pixel i, and T is the simulated annealing temperature. The temperature determines the chance of accepting negative changes and is gradually lowered as the optimisation continues.

For this work, the initial temperature was selected based on a desired initial acceptance probability P using an approach implemented by Meister et al [173]. This works by initially sampling ΔC_{i} and setting T from its distribution. For reducing the temperature, the method used was to decrease it by a scaling factor between each iteration using

$$T_{\mathrm{new}} = \gamma T_{\mathrm{old}}, \tag{4.7}$$

where T_{new} is the temperature for the next iteration, γ is a parameter determining the rate of cooling, and T_{old} is the old temperature.

The simulated annealing approach was compared to direct binary search for several different starting and cooling conditions. Three different choices for the initial acceptance probability and cooling were used: a 'fast' one with $P = 0.6$ and $\gamma = 0.75$, a 'medium' one with $P = 0.85$ and $\gamma = 0.90$, and a 'slow' one with $P = 0.9$ and $\gamma = 0.95$. The optimisation was run 40 times for each schedule and the resulting cost and computational time evaluated in each. The results are shown in Fig. 4.11.

It can be seen that there is effectively no benefit to using a simulated annealing

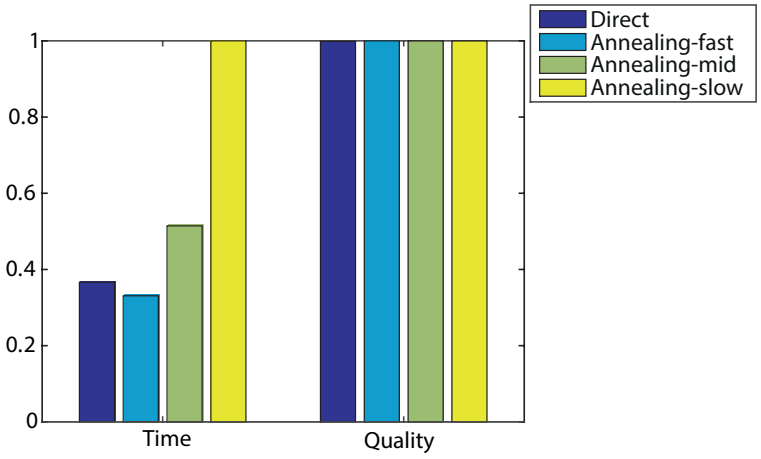

Figure 4.11: Bar chart showing average time to converge and quality of solution for simulated annealing using different cooling schedules vs direct binary search. Quality is measured by the cost function given in Eq. 4.2.

approach over direct binary search. The quality of solution is the same on average in each case and computationally more effort is expended. While the fastest cooling case did converge faster than direct binary search in these tests, this was not by a statistically significant amount and is due to the direct binary search and simulated annealing being effectively the same at low T.

4.2.5 Final algorithm

A flow chart summarising the default final algorithm is shown in Fig. 4.12. The hologram is initialised in a random binary state. Individual pixels are then selected in a random exhaustive manner, flipped, and the change is either kept or reversed depending on whether the cost decreased. The iterations are terminated when the number of changes falls below 0.5% of the pixels on the surface. Variations on this algorithm structure are utilised in later Chapters (alterations to the structure will be given in the text).

4.3 Multi-foci holograms simulations

Two cases were investigated to test the properties of multi-foci holograms: an acoustic field with a line focus, and a field with a square array of foci. These were used to investigate the sampling rate needed to represent a continuous shape, and the change in pressure and foci size as a field is made to focus at several points. Simu-

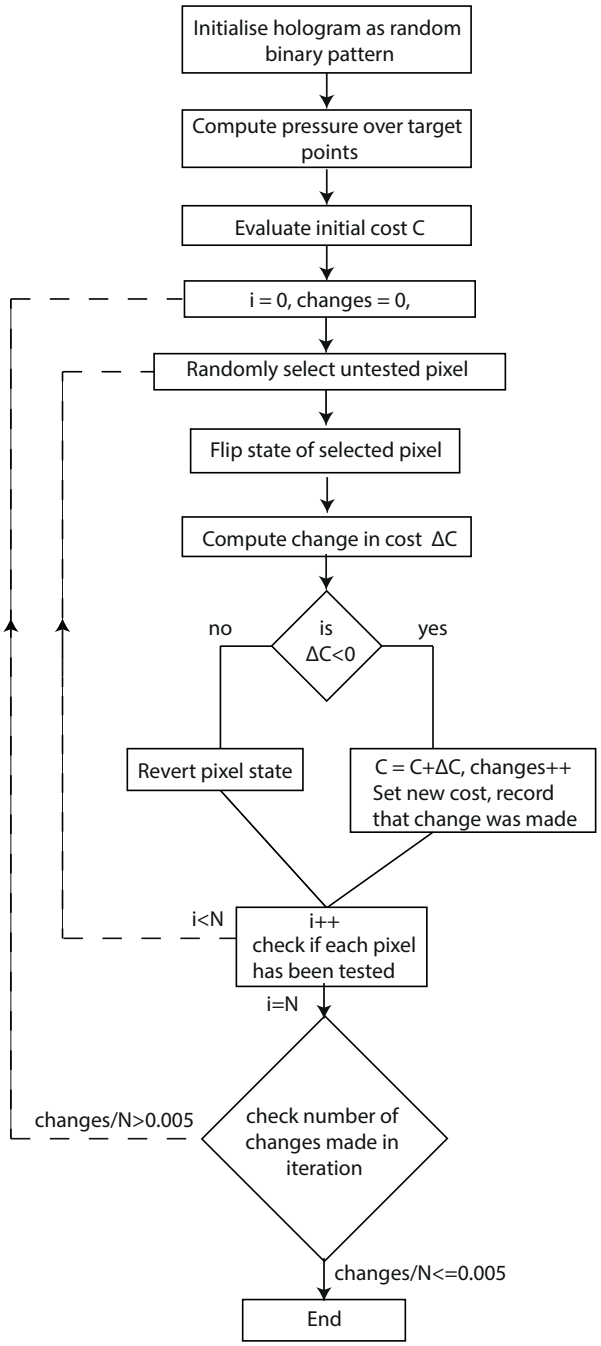

Figure 4.12: Flow chart summarising the optimisation algorithm. Figure reprinted from [164] under CC BY 4.0 License.

lations were carried out on a 256×256×256 grid, with a grid spacing of 0.167 mm, time steps of 33 *ns*, and a maximum supported frequency of 4.5 MHz. Holograms were calculated for a design frequency of 3 MHz and an aperture of 3.33×3.33 cm.

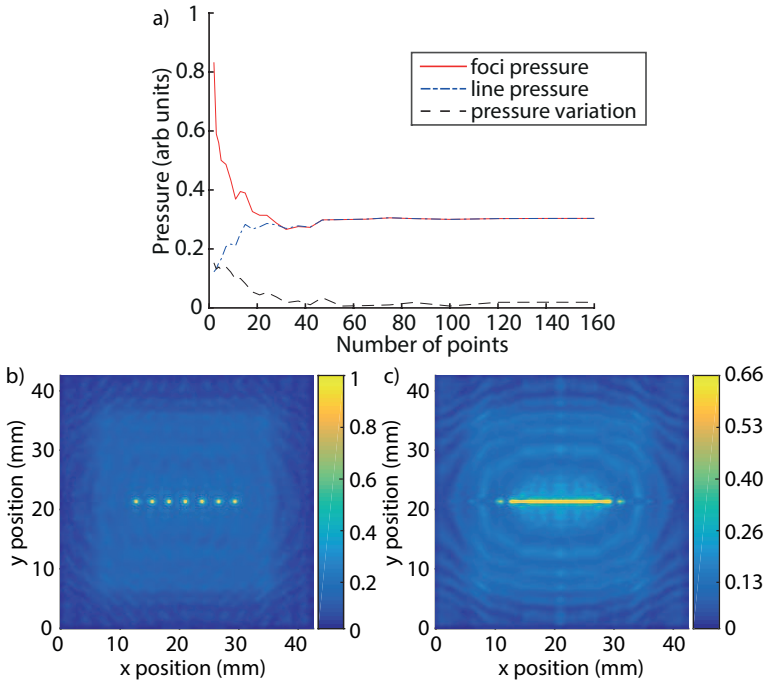

Figure 4.13: (a) Measured variation in mean line pressure, mean focal pressure, and standard deviation in pressure over the line with the spacing of points used in discretising the line. (b) Maximum pressure in target plane of line for a hologram generated using a point spacing of 5.55λ. (c) Maximum pressure in target plane of line for a hologram generated with a point spacing of 0.1675λ. Pressure normalised relative to maximum in (b). Figure adapted from [164] under CC BY 4.0 License.

4.3.1 Line focus

The generation of a continuous pattern of pressure was investigated first. The target used was a line focus, 1.66 cm in length in a plane 3.33 cm from and parallel to the hologram aperture. As an input to the optimisation, the line had to be discretised into a set of points. A set of holograms were calculated varying the number of points used to discretise the line between 2-160 points. The acoustic field generated by each hologram was simulated recording the average peak pressure in each voxel of the line, the variation in peak pressure over the line, and the peak pressure at each target point. These are each plotted in Fig. 4.13.

From Fig. 4.13 (c) it can be seen that it's possible to generate an acoustic field focused over a line. The average maximum pressure over this line increases with increasing point number up to 45, which corresponds to a point spacing of $\sim 0.75\lambda$. At this point it remains stable showing only small variations between

different point densities (Fig. 4.13 (a)). The maximum pressure at each individual foci decreases with decreasing point spacing until approximately 0.75λ. At this point it remains stable and matches the value of the average line pressure. This spacing also coincides with a large drop in the variation of maximum pressure over the line, which demonstrates that 0.75λ is the maximum point spacing that can produce a continuous pattern without significant variations in pressure. At lower spacings, a continuous line is also generated but there is a larger computational cost associated with calculating the hologram. For more sparse distributions, the line is realised as a set of discrete foci (Fig. 4.13 (b)). This result will be impacted by the choice of focal depth and aperture, which change the point spread function (PSF).

4.3.2 Multiple foci

The generation of multiple discrete foci was investigated next. In this case, the target was a square array of evenly spaced points. The points were distributed over a 1×1 cm area centred on the hologram in a plane 3.33 cm away. The number of target points in the array was varied from 2×2 to 11×11. The variation in the average area occupied by each foci, and the contrast between the target focal points and the background are shown in Fig. 4.14. To evaluate the average area occupied by each focus, the maximum pressure data was linearly interpolated onto a denser grid to improve differentiation, then thresholded at 50%. The area was then given by the size of the cluster centred on each target point in the 2-D target plane.

Figure 4.14 (b)-(c) show that arrays of points with evenly distributed intensities can be generated. However, there are several limitations. As the number of points in the array increases, the ratio between the pressure at each target point and the background pressure decreases (Fig. 4.14 (a)). This is due to both the background pressure increasing and the pressure at each focal point decreasing. For the 11×11 array there is only a 16% difference between the two.

As the number of points in the array increases, the average area occupied by each focus increases gradually until the 8×8 array (Fig. 4.14 (a)). A 46% increase is seen between the 2×2 and 8×8 arrays. From the 9×9 array onward the pressure at background exceeds 50% of the pressure at the foci. This results in the effective

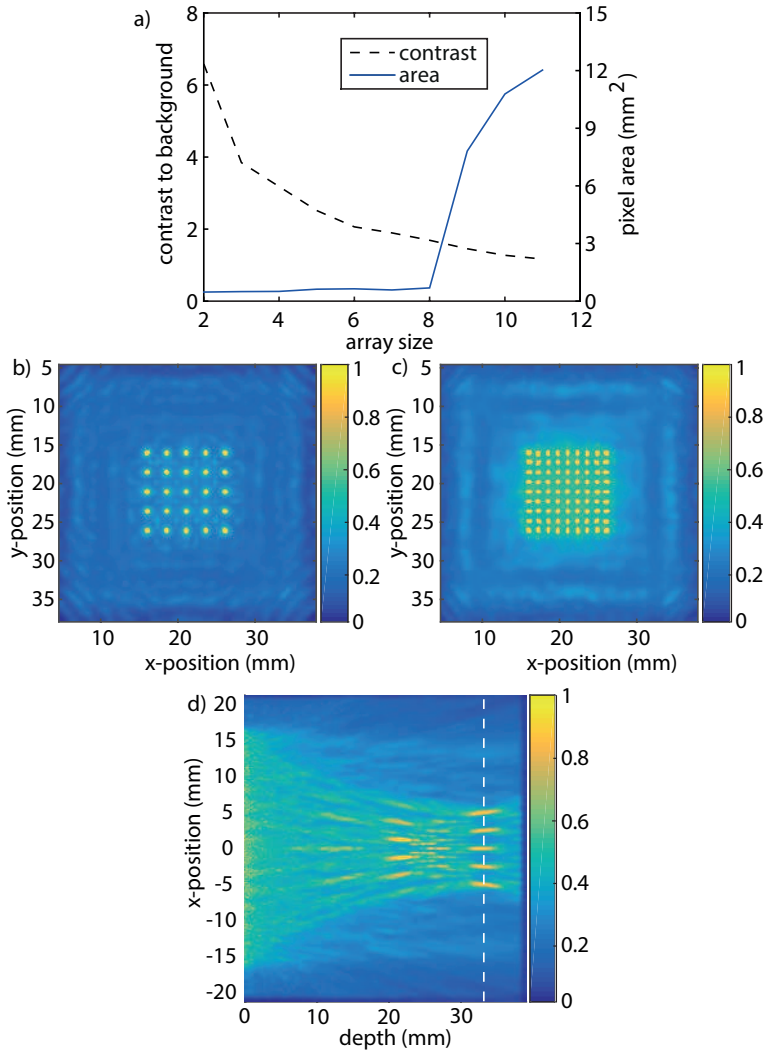

Figure 4.14: (a) Variation in the ratio of the average of the maximum pressure at each focus to the average of the maximum pressure over the background pixels as the size of the array of points increases. Also shown is variation in average area occupied by each focus as the size of the array increases. (b) Maximum pressure in target plane of hologram for a 5×5 array of points (c) Maximum pressure in target plane for a 9×9 array of points. (d) Maximum amplitude projection of maximum pressure generated through $(x, depth)$. The dashed line indicates target plane for array of foci. Figure reprinted from [164] under CC BY 4.0 License.

'area' occupied by each focus significantly increasing as different focal points can't be separately resolved by the metric. The closest spacing of discrete points that can be achieved is therefore limited to 2.2λ. Note this is distinct from the 0.75λ found in Sec. 4.3.1, which represents the maximum discrete point spacing which generates a continuous acoustic pressure.

The maximum amplitude projection of the hologram field perpendicular to the

target and hologram planes for the 5×5 array is shown in Fig. 4.14 (d), where the target plane is indicated by the dotted line. It can be seen that high pressure foci are generated in planes other than the target plane. For applications related to neurostimulation or ultrasound therapy, this is a potential drawback, as heating or stimulation could occur at unintended targets. However, these focal points occur in the near-field for the hologram aperture, so are not dissimilar to those that occur for conventional piezoelectric arrays.

4.4 Experimental validation

The accuracy of the direct search algorithm and the forward model, which approximated the hologram as a series of monochromatic point sources, was experimentally validated. This was done using a hologram designed to generate a line focus. The hologram had an aperture of 3×3 cm, a pixel size of 250 μm, and was designed for a frequency of 3 MHz. The target line had a length of 10 mm and was centred on and parallel to the hologram aperture at a depth of 3 cm. To calculate the hologram, the line was discretised into 400 points. This oversampled the line by a large factor (∼7) compared to the minimum necessary which was done to ensure sufficient sampling. The hologram was fabricated using the approach described in Sec. 3.9.1. A photo of the resulting sample can be seen in Fig. 4.15 (a).

Measurements of the acoustic field were then made using the set-up described in Sec. 3.9.3 using the calibrated 0.2 mm PVDF needle hydrophone. To generate sound, the first of the two Litron lasers was used to excite the sample with a pulse energy of ∼300 mJ. The acoustic field was then measured over a 34×34 mm area using a step size of 0.4 mm. Time domain signals were recorded at each position taking 6 averages.

The data was up-sampled to a spacing of 0.133 mm and low pass filtered at 4 MHz to eliminate harmonics. Additional pulses were then created in the data by adding the time series recorded at each point to itself translated by multiples of 0.33 μs. The data was then forward and back projected by 30 and 20 mm respectively. The hologram plane was located in the back-projected data and the target

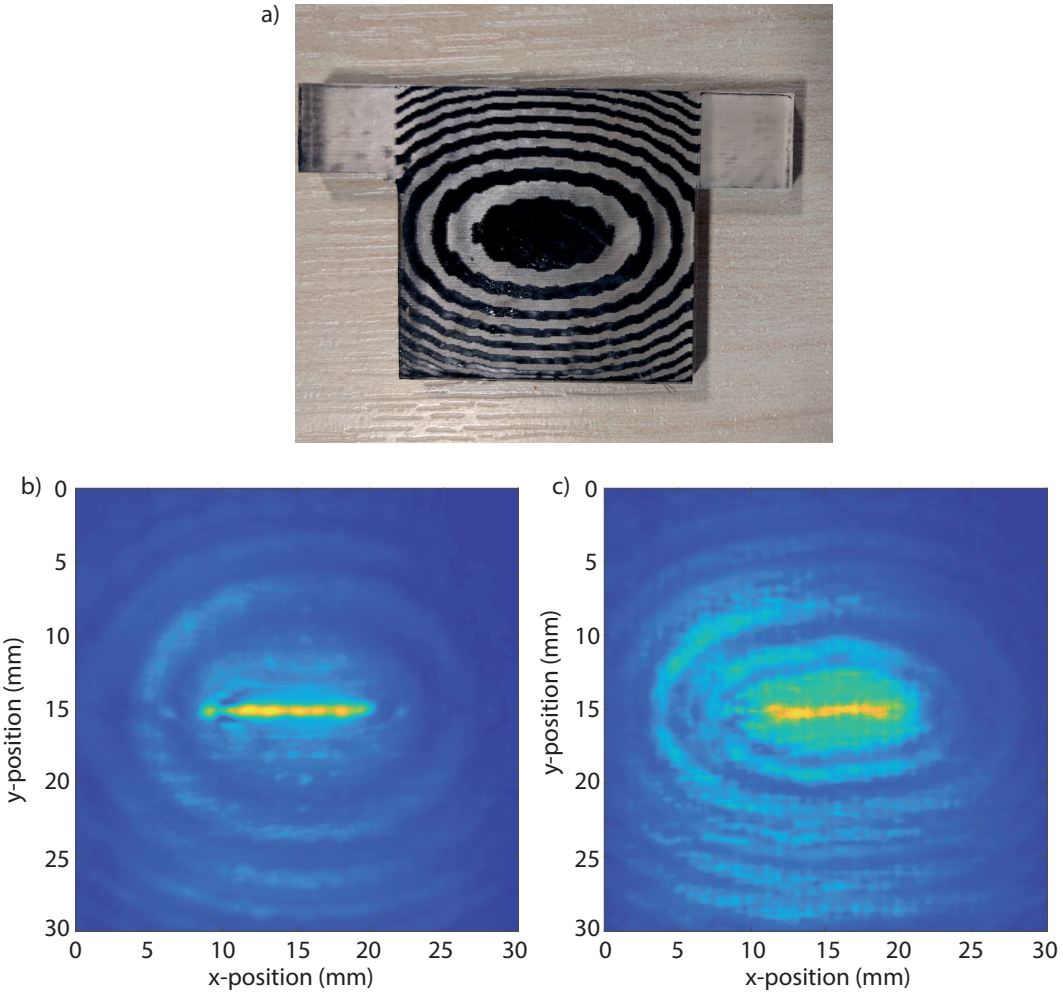

Figure 4.15: (a) Fabricated hologram designed to focus a frequency of 3 MHz over a 1 cm line in a plane 3 cm from the hologram. (b) Maximum pressure generated in a plane 3 cm from the hologram low-pass filtering the data at 4 MHz to eliminate harmonics. (c) Maximum pressure generated in a plane 3 cm from the hologram low-pass filtering the data at 10 MHz to include 2 harmonics.

line identified in a plane 30 mm from the hologram aperture. The maximum pressure generated across this plane can be seen in Fig. 4.15 (b). This shows that a line focus has been successfully generated in the experimental data, confirming that continuous patterns of pressure can be generated using binary amplitude holograms.

4.5 Harmonic correction

The initial analysis in Sec. 4.4 and the simulations in Sec. 4.3 were carried out for a single frequency, pulsed, but filtering out the contribution from harmonics. This can be a valid approximation, for example, for high frequency holograms in

highly attenuating media or for thick absorbing layers with low optical absorption. However, in most cases harmonics will be present in the generated spectra and, as seen in Chapter 3, they can potentially have a significant influence on the field.

To test the influence of the harmonics, the experimental dataset described in Sec. 4.4 was re-analysed. The data was up-sampled by a factor of 6 to a spacing of 0.0667 mm and low-pass filtered at 10 MHz such that two harmonics were included in the propagated data. Additional pulses were then added to the time series and the measurements forward and back-propagated. The field generated in the same plane as Fig. 4.15 (b) when two harmonics are present in the spectra is shown in Fig. 4.15 (c). The line now shows significant aberration, demonstrating that having just two harmonics in the field significantly distorts the generated pattern. This suggests that for applications at lower frequencies where a large number of harmonics will be present, a different forward model is required.

To account for the harmonics, a time-domain approach using k-Wave was adopted. First, a forward simulation was run. The input to this simulation was a single point source with a pulsed temporal input containing the desired number of harmonics. The time varying pressure $p(t)$ was recorded at a distance d from the source. The pressure $p_{ij}(r_k,t)$ generated by each hologram coordinate (x_i,y_j) at each target point r_k was then calculated from this signal. This was done by shifting the time signal $p(t)$ based on the distance between the target point and hologram coordinate using

$$p_{ij}(r_k,t) = p\left(t - \frac{d_{ijk}}{c_0}\right). \tag{4.8}$$

Here d_{ijk} is the distance between the target point and the hologram coordinate and c_0 is the sound speed. This is similar to delay and sum beamforming described in Sec. 2.1.1. Since the signal $p(t)$ is periodic, a window $T = \frac{1}{f}$ about some t_0 was extracted from the calculated time series to reduce memory consumption.

The pressure generated by a given hologram at each target point was then evaluated using

$$p(r_k,t) = -\sum_{j=1}^{N}\sum_{i=1}^{N}\frac{d}{d_{ijk}}H_{ij}p_{ij}(r_k,t). \tag{4.9}$$

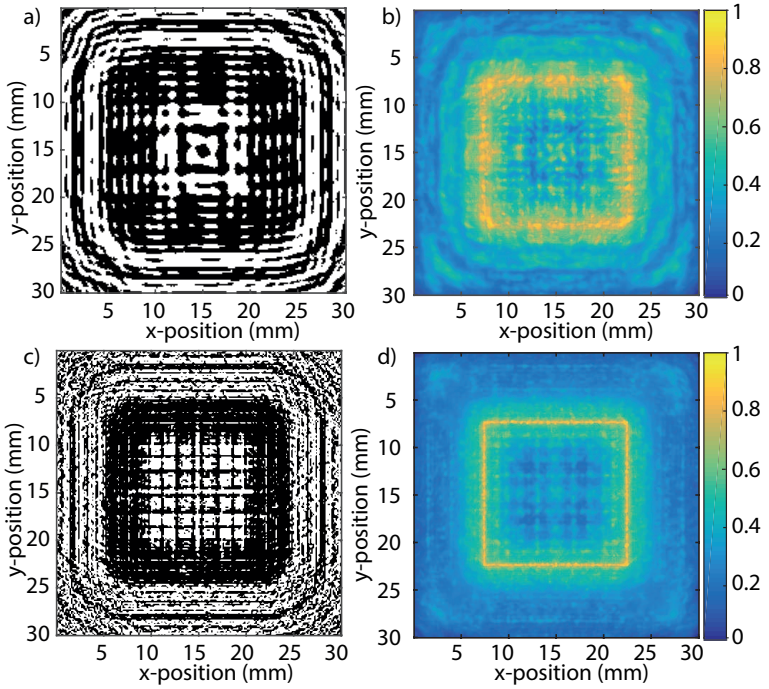

Figure 4.16: (a) Binary hologram to generate square outline. (b) Maximum pressure generated by hologram in (a) in target plane of square, simulated including first 2 harmonics in the input spectra. The outline of the square is very uneven. (c) Hologram generated using the harmonic corrected forward model from Eq. (8). Target and parameters are identical to (a). (d) Maximum pressure generated by hologram in (c) in target plane of square. The square outline is significantly more even than the field in (b). Figure reprinted from [164] under CC BY 4.0 License.

Here the term $\frac{d}{d_{ijk}}$ is included to account for spherical spreading. Since the complex pressure was no longer available, the cost function was altered to

$$C = -\left(\frac{1}{K}\sum_{k=1}^{K}\max_t(p(r_k,t))\right) + \text{std}(\max_t(p(r_k,t)). \tag{4.10}$$

Here $\max_t(p(r_k,t))$ is the maximum pressure generated at r_k at any time t and K is the number of target points. As with Eq. 4.2, the first term aims to maximise the peak pressure at each point, and the second aims to minimise the variation in peak pressure.

This model was tested using a hologram designed to generate the outline of a square. This had an aperture of 30×30 mm, was discretised using a spacing of 150 μm, and was designed for a frequency of 3 MHz. The square had a width of 1.5 cm and was placed in a plane 3 cm from the hologram. The resulting hologram

and maximum pressure are plotted in Fig. 4.16 (c)-(d). For comparison, a hologram calculated using the non-corrected algorithm and its corresponding field are shown in Fig. 4.16 (a)-(b). These were both simulated on a 512×512×512 grid with a spacing of 75 μm. Both holograms were up-sampled by a factor of 2, and two harmonics were included in the simulation spectra.

It can be seen that the outline of the square is successfully recovered using the new model. The spectral difference introduced by utilising a series of short pulses can therefore be accounted for in this way. However, the hologram generated by this algorithm (Fig. 4.16 (c)) shows significantly greater speckle or smaller structures than the corresponding hologram in Fig. 4.16 (a). This makes fabrication a greater challenge.

4.6 Experimental validation: harmonics

The model accounting for harmonics was then also validated experimentally. The hologram used as a test case had identical parameters to that of Sec. 4.4. However, it was calculated using the new forward model accounting for frequencies up to 15 MHz, and the target pattern was changed to a 2 cm line located in a plane 3 cm from the hologram. Prior to fabrication, clusters of pixels on the calculated hologram smaller than 15 were removed, because, as discussed in Sec. 3.9.1, these were not well realised by the laser cutter. The fabricated sample is shown in Fig. 4.17.

The field was measured using the set-up described in Sec. 3.9.3. A single pulse was again used with an energy of 300 mJ. The field was recorded over a 3×3 cm aperture using a 0.3 mm spacing, and 6 averages were taken at each position. The dataset was spatially up-sampled to a grid spacing of 0.0375 mm to support frequencies up to 20 MHz in the simulation, and low pass filtered at 15 MHz to match the number of harmonics corrected for by the model. The data was additionally high pass filtered at 100 kHz and 15 additional pulses added to the data. The measurements were then forward and back propagated by 30 and 20 mm, respectively. The plane of the hologram was found in the back propagated data, and the target

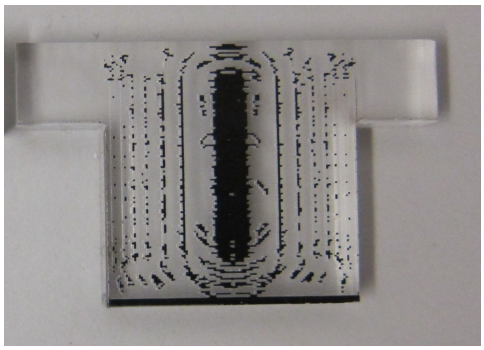

Figure 4.17: Line focus hologram fabricated using laser cutting on spray painted absorber. The hologram has dimensions of 3×3 cm, the focus is 3 cm from the hologram, and the design frequency is 3 MHz.

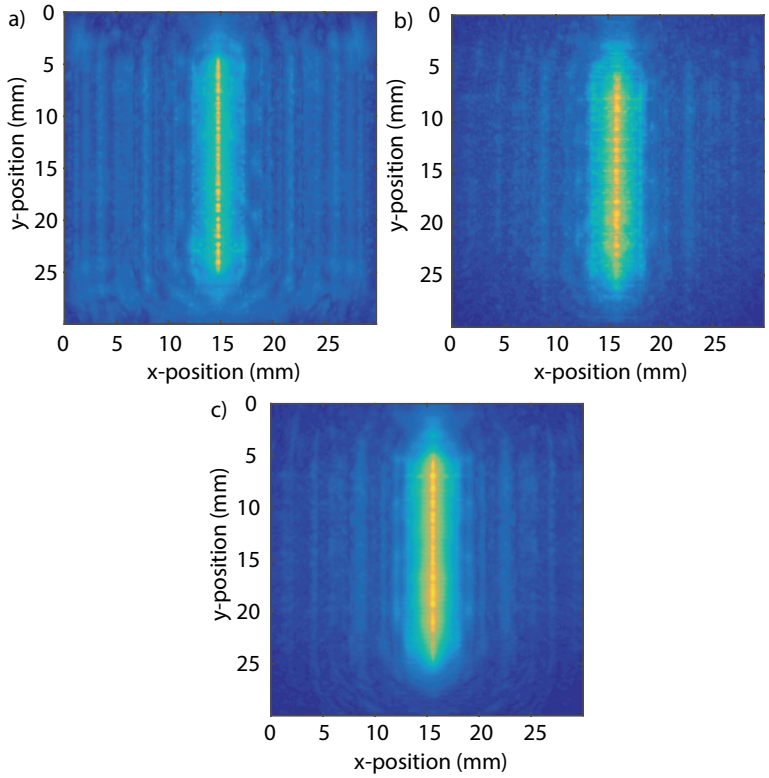

Figure 4.18: (a) Maximum pressure in the target plane of line for the line focus hologram shown in Fig. 4.17 simulated using k-Wave. (b) Maximum pressure in the target plane of line for the line focus hologram shown in Fig. 4.17 recorded in experiment. (c) Maximum pressure in target plane of the line for the line focus hologram shown in Fig. 4.17 simulated using k-Wave correcting for the spatial profile of laser using experimental data. Figure adapted from [164] under CC BY 4.0 License.

line was identified 30 mm from this hologram plane. This was compared against a simulation of the acoustic field generated by the same hologram. The experimental and simulated fields are plotted in Fig. 4.18 (a)-(b).

It can be seen from Fig. 4.18 (b) that a line focus is generated by the fabricated hologram, validating the model. However, while qualitatively similar, there are deviations between the experimental data and the simulation (Fig. 4.18 (a)). This is principally due to the spatial profile of the excitation laser. To account for this, a second simulation was run correcting for this spatial variation by registering the maximum pressure data from the plane of the hologram in the back-propagated data to the binary hologram. The results of this simulation are shown in Fig. 4.18 (c). This shows that when the spatial profile of the laser is accounted for, good agreement is achieved between the simulated and experimental data.

4.7 Conclusions

This Chapter has shown that acoustic fields focused over patterns or at multiple discrete points in three dimensions can be generated using binary amplitude holograms combined with trains of optical pulses. This represents the first demonstration that arbitrary arrangements of foci, or continuous arrangements of pressure can be generated using this technique. This technique has potential applications across a range of fields for example, for the dynamic manipulation of arrangements of small particles [17], or, for the stimulation of structures in the retina [12].

A direct search algorithm has been adapted from optics for the design of these binary amplitude holograms. This searches for a binary pattern that best approximates the target field by sequentially perturbing individual pixels and evaluating the change in solution quality using a cost function. The starting conditions and search method for this algorithm were optimised and the convergence rate and the properties of the converged results established. Two modifications for the algorithm were then introduced. The first allowed the relative pressure across the target pattern to be controlled by adding a weighting term to the cost function, the second allowed for the SNR to be improved by introducing a penalty term on the number of transmitting pixels. Two different forward models were described for the algorithm. The first approximated the hologram as a collection of monochromatic point sources, and was accurate when the spectrum contained no (or weak) harmonics.

The second used a time-domain delay-and-sum approach, and was accurate when harmonics were present. Both forward models were validated experimentally using measurements from patterned optical absorbers fabricated using laser cutting.

There are several future directions for the work presented in this Chapter. First, as in Chapter 3, employing spatial light modulation in conjunction with a temporally modulated optical source would enable the generated hologram to be dynamically changed, and therefore the acoustic field to be varied. Further extensions to the optimisation can also be considered. For example, simultaneous control over the phase and amplitude at the target points could be achieved by modifying the cost function to include a phase term. Non-binary modulation could also be explored. If a spatial light modulator is used to control the incident fluence then the amplitude values for each pixel aren't limited to 'on' or 'off'. It has been shown in optics that for zone plates, the higher order foci can be eliminated with gray-scale modulation [174]. This could therefore be investigated as a technique for further controlling the relative influence of harmonics on the generated field.

Chapter 5

Multi-layer absorbers

5.1 Introduction

In Chapter 3 and 4, the generation of single and multi-focal acoustic fields using binary amplitude optoacoustic holography was investigated. This technique can, in principle, be used to dynamically control high-frequency acoustic fields. However, the requirement for either a modulated optical source or a minimum of 2 high peak-power pulses to generate a focused field significantly increases both the cost and complexity of implementing the method experimentally. Ideally, the generation of arbitrary focused acoustic fields would be possible using a single optical pulse.

One method through which this could be achieved is to, rather than use a single absorbing layer, use an absorber made up of several layers each separated by a transparent substrate of a known thickness. By introducing a spatial offset between the different layers, temporal overlap between the signals from different parts of the aperture can be achieved. Each layer can then be individually patterned such that the acoustic signal generated on each arrives simultaneously at a set of target points and constructively interferes. This concept is illustrated in Fig. 5.1.

This Chapter validates that this proposed method can be used to generate multi-focal acoustic fields. In Sec. 5.2, the design problem is outlined and a ray-tracing model is described for the rapid approximation of impulse responses across several layers to a set of arbitrarily placed target points. An optimisation approach is then developed to design the multi-layer absorbers from the calculated impulse re-

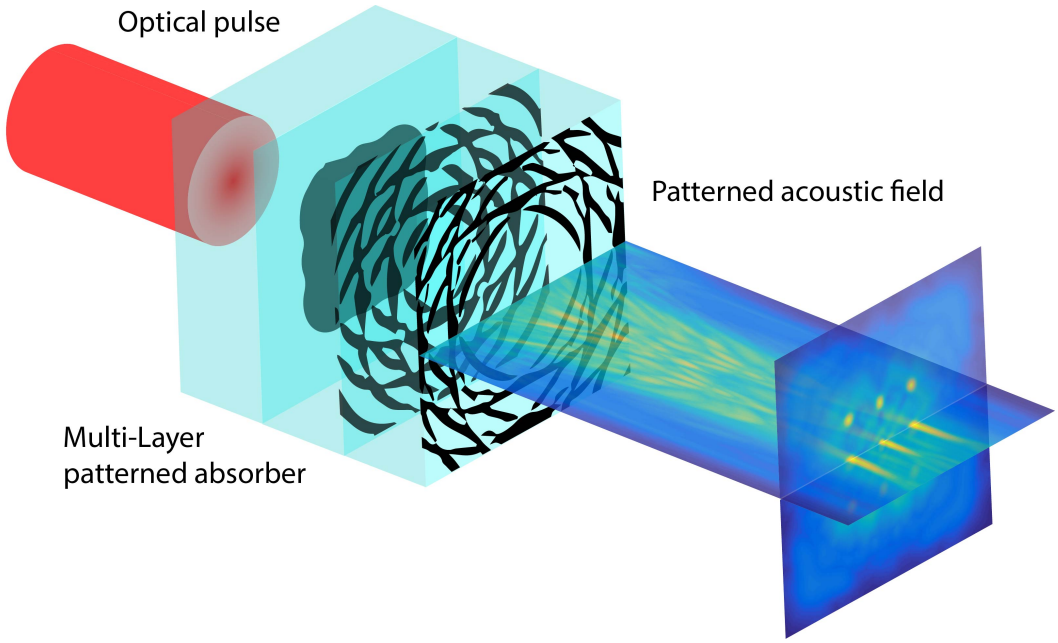

Figure 5.1: Illustration of a multi-layered patterned absorber for the generation of patterned optically generated acoustic fields. The pattern on each layer and their spacing are tailored such that constructive interference is achieved at the desired focal points. Figure reprinted from [175] © IEEE 2016.

sponses. In Sec. 5.3, the approach is validated experimentally using a 3-layer test case designed to generate a 3×3 array of foci. In Sec. 5.4, an alternative fabrication method using 2 material 3-D printing is introduced. Finally, in Sec. 5.5, the conclusions of the work are presented.

The conference proceeding in [175] © IEEE 2016 has been adapted to form parts of this chapter.

5.2 Multi-layer absorber design

5.2.1 Design problem

The goal in designing a multi-layer absorber is to find a set of binary patterns that, when excited optically by a short pulse, will generate high peak pressures at a set of target points. As a first step in calculating these patterns it is necessary to know the acoustic signal $p(r_k,t;x_i,y_j,d_l)$ generated at a target point r_k at a time t, by each pixel (x_i,y_j) in each layer d_l, where d_l is depth of the l^{th} absorbing layer, so that the signal generated by a each set of patterns can be predicted.

Since the absorbing layers are embedded in a transparent substrate with dif-

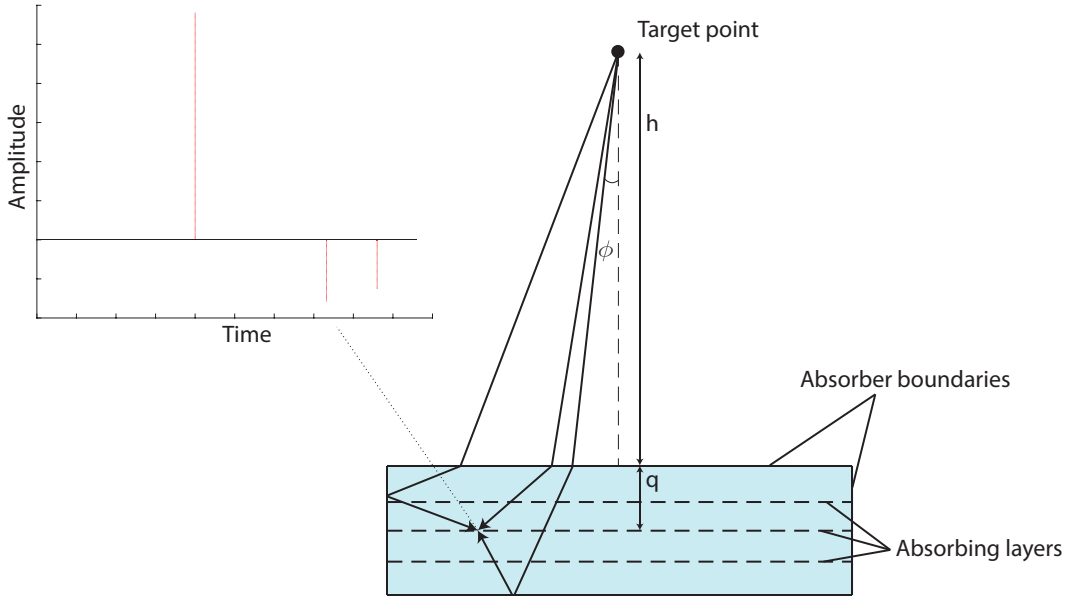

Figure 5.2: Illustrative figure of the ray-tracing approach. Rays are traced from the target point into the absorbing structure.

ferent acoustic properties to the surrounding medium, calculating these impulse responses is less trivial than in free-space. Refraction on leaving the absorber, reflections off the internal structure, and transmission losses all have to be considered. These can be modelled using a full-wave model such as the k-Wave toolbox. However, this requires running a separate simulation for every target point to separate the contributions of different points. This makes the approach computationally inefficient. As a result, a ray-tracing model was developed to rapidly approximate the impulse response for each pixel in each absorbing layer to arbitrarily placed target points. This is illustrated by Fig. 5.2.

5.2.2 Ray-tracing model

The inputs to the ray-tracing model consisted of: a set of targeted focal points r_k, the dimensions and shape of the absorber and its position relative to the focal points, the set of coordinates making up each absorbing layer (x_i, y_j, d_l), and the sound speed and density of the absorber (c_h, ρ_h) and the coupled medium (c_m, ρ_m).

To select the ray-trajectories, one approach would be to select a set of angles $\phi \in [0, \pi]$ and $\theta \in [0, 2\pi]$ that are evenly distributed on the surface of a sphere, then restrict these to those that intersect with the absorber. However, due to wave-

spreading, it was found that a large number of trajectories were required to sufficiently sample and approximate the pressure in pixels at higher transmitted angles. Therefore, the ray-trajectories were selected by evenly sampling between 0-$\frac{\pi}{2}$ in the transmitted angle (ϕ_t). These angles were then transformed into a set of incident angles (ϕ_i) using Snell's law:

$$\phi_i = \sin^{-1}\left(\frac{c_m}{c_h}\sin\phi_t\right). \tag{5.1}$$

For simplicity, the problem was orientated such that normal incidence to the absorber corresponded to $\phi_i = 0$ (Fig. 5.2) and the absorbing layers were placed in the (x,y) plane. For each ϕ_i, the number N of radial angles θ used was weighted based on the interaction radius with the absorber surface. This was calculated using

$$N = Ah\tan\phi_i + B. \tag{5.2}$$

Here h is the distance from the plane of the target point to the top layer of the absorber indicated in Fig. 5.2, and A and B are constants. The constant B was used to ensure that $\phi_i = 0$ was sampled, and to ensure a minimum sampling for small values of ϕ_i. The constant A was used to scale the number of rays. The set of values for θ for each ϕ_i were chosen using

$$\theta_n = \frac{2\pi n}{N} + \frac{2\pi C}{N}. \tag{5.3}$$

Here C is a random number between $0-1$ and n is an integer between $1-N$. This ensures the whole range $0-2\pi$ is sampled. The random factor was introduced for practical reasons to prevent pixels at certain radial angles from being systematically missed.

The propagation of each ray was modelled using Snell's law to transmit into the absorber and the Fresnel equations to account for subsequent reflection losses. Since the set of ray-trajectories used was uneven, the amplitude was calculated using an analytic solution introduced by Young [176]. This gives the pressure amplitude p at a particular depth for a point source above an infinite half-space as

$$p = \frac{T \cos \phi_i}{(h + q\frac{c_h}{c_m} \cos \phi_i \sec \phi_t)^{\frac{1}{2}} (h + q\frac{c_h}{c_m} \cos^3 \phi_i \sec^3 \phi_t)^{\frac{1}{2}}}. \tag{5.4}$$

Here q is the distance of the point below the interface shown in Fig. 5.2, and T is the pressure transmission coefficient for the incident angle calculated using

$$T = \frac{2\rho_h c_h \cos \phi_i}{\rho_h c_h \cos \phi_i + \rho_m c_m \cos \phi_t}. \tag{5.5}$$

For each intersection of a ray with a pixel (x_i, y_j, d_l), the time of arrival t, the amplitude, and the original trajectory (ϕ_i, θ_i) of the ray were stored. The amplitude was calculated using Eq. 5.4, accounting separately for losses caused by any subsequent reflections within the absorber. These were calculated using

$$R = \frac{\rho_m c_m \cos \psi_i - \rho_h c_h \cos \psi_t}{\rho_m c_m \cos \psi_i + \rho_h c_h \cos \psi_t}, \tag{5.6}$$

where ψ is the angle between the boundary normal and the incident ray. The ray was judged to have intersected a pixel if it passed through the square defined by $\{(x_i - \frac{\Delta}{2}, y_j - \frac{\Delta}{2}, d_l), (x_i - \frac{\Delta}{2}, y_j + \frac{\Delta}{2}, d_l), (x_i + \frac{\Delta}{2}, y_j - \frac{\Delta}{2}, d_l), (x_i + \frac{\Delta}{2}, y_j + \frac{\Delta}{2}), d_l)\}$. To avoid over or underestimating the pressure in a given pixel, the properties of rays with similar initial trajectories incident on the same pixel were averaged. For two rays (ϕ_1, θ_1) and (ϕ_2, θ_2) the properties were averaged if

$$|\phi_1 - \phi_2| < 1.5^o \text{ or } |\theta_1 - \theta_2| < 1.5^o \tag{5.7}$$

The output of the ray-tracing was a set of impulse responses $I(r_k, t; x_i, y_j, d_l)$. To get the acoustic signal $p(r_k, t; x_i, y_j, d_l)$ generated by each pixel, these were convolved with the signal generated optically from a pixel in the absorber $p(t)$. This was approximated by using k-Wave to simulate a point source on a 3-D grid, with a spacing Δ, driven by a low-pass filtered impulse.

The ray-tracing model was validated by comparison with k-Wave. This comparison used a cylindrical absorber with a radius of 11 mm, a height of 4 mm, a sound speed of 2450 $\mathrm{m\,s^{-1}}$, and a density of 1180 $\mathrm{kg\,m^{-3}}$. A layer 2 mm below the

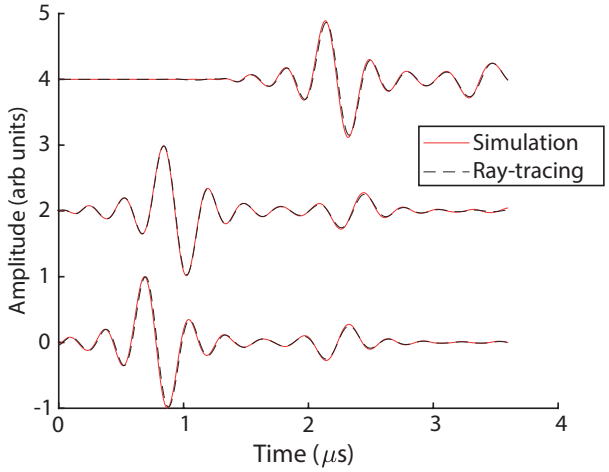

Figure 5.3: Comparison of the acoustic response calculated at three different points using the ray-tracing approach (dashed line) and simulation (continuous line). The time series for the ray-tracing and simulation are each normalised to the maximum pressure across all three points for the respective method. At all 3 points good agreement is found over the first 4 μs. Figure reprinted from [175] © IEEE 2016.

top surface was used for the comparison. The target point was placed 1.8 cm above the cylinder displaced 2 mm from its centre in both the x and y directions. The surrounding medium was assumed to be water with a sound speed of 1500 $\mathrm{m\,s^{-1}}$ and a density of 1000 $\mathrm{kg\,m^{-3}}$. The ray-tracing was run for these inputs and the output compared against a k-Wave simulation.

The k-Wave simulation was carried out on a $256\times256\times288$ domain with a grid spacing of 0.1 mm. The absorber was inserted at one end of the domain as a cylinder of matching dimensions, with acoustic properties matching those of the absorber. The rest of the medium was set to a sound speed of 1500 $\mathrm{m\,s^{-1}}$ and density 1000 $\mathrm{kg\,m^{-3}}$. The target point was inserted as a source point 180 grid points above the cylinder, displaced by 20 grid points from its centre in both x and y. The pressure was recorded across a plane 20 grid points into the cylinder. A bipolar pulse low-pass filtered at 6 MHz was used as the acoustic signal $p(t)$ for both the simulation and the ray-tracing. A comparison of the ray-tracing and k-Wave outputs from 3 separate positions is shown in Fig. 5.3. The two models show good agreement over the first 4 μs or, equivalently, the first 2 reflections.

5.2.3 Multi-layer optimisation

5.2.3.1 Introduction

After evaluating the impulse responses, they were used to determine the binary pattern for each layer using a direct search approach. As discussed in Sec. 5.2.1, the goal in designing the multi-layer absorber is to find a set of binary patterns H_{ijl} that will maximise the temporal peak pressure $\max_t(p(r_k,t))$ at each target point r_k. This is subject to the constraint that at each (x_i,y_j) position, only a single pixel can be active across each of the layers d_l (i.e., the binary patterns must have no overlap). This arises because, assuming the absorber is illuminated from below, the incident light will be almost entirely absorbed by the first pixel with which it interacts at each (x_i,y_j) position. This is due to the need for high optical absorption to maximise the acoustic pressure. Therefore, pixels at the same position (x_i,y_j) on subsequent layers will not contribute a signal.

It is, in principle, possible to tailor the optical absorption such that signals from the same position on multiple layers could be generated or, alternatively, materials that absorb over different bandwidths could be used for different layers. However, this would require a separate optical model and the fabrication would become significantly more complex, so this was not considered as part of this thesis.

5.2.3.2 Forward model and cost function

To begin with, a forward model and cost function for the optimisation were defined. The acoustic signal $p(r_k,t)$ generated at each target point by an arbitrary set of binary patterns H_{ijl} can be calculated from the derived impulse responses using

$$p(r_k,t) = \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{l=1}^{L} H_{ijl}\,p(r_k,t;x_i,y_j,d_l). \tag{5.8}$$

Here L is the total number of absorbing layers. The cost function used was the same as for the harmonic corrected direct binary search given in Eq. 4.10, that is

$$C = -\left(\frac{1}{K}\sum_{k=1}^{K}\max_t(p(r_k,t))\right) + \mathrm{std}(\max_t(p(r_k,t))). \tag{5.9}$$

This again aims to balance maximising the temporal peak pressure while minimising the variation between the target points.

5.2.3.3 Initialisation

The initialisation and search approach used in Chapter 4 were found to be inappropriate for the design of multi-layer holograms. Applying them naively, the optimisation converged to patterns that generated low peak pressures across the target points, and contained a large amount of speckle or small pixel clusters. As discussed in Chapter 4, accurately generating these small pixel clusters using laser cutting is a challenge. Additionally, assuming they could be generated with an alternate method, such patterns would be more sensitive to registration errors between the different layers, model errors, or misalignment of the optical source. Both the initialisation and search approach were therefore designed to avoid their generation.

To initialise the binary patterns H_{ijl}, the sum of the pressure generated by each target point r_k, across each plane d_l, at a time t_{thres}, was thresholded about 0. This was evaluated using

$$H_{ijl} = \begin{cases} 1 \text{ if } p_{ijl}(t_{thres}) > 0, \\ 0 \text{ if } p_{ijl}(t_{thres}) \leq 0. \end{cases} \tag{5.10}$$

Here the term $p_{ijl}(t_{thres})$ is the sum of the pressure generated by each target point at each pixel at a time t_{thres}. This was evaluated using

$$p_{ijl}(t_{thres}) = \sum_{k=1}^{K} p(r_k, t_{thres}; x_i, y_j, d_l), \tag{5.11}$$

where K is the number of target points. Using this method, the initial set of patterns H_{ijl} generates a high peak pressure at each target point at a pre-determined time t_{thres}. However, it also, usually, results in patterns possessing substantial overlap between different layers which has to be removed. This is illustrated by the third column in Fig. 5.4.

Due to the fabrication constraints discussed above, the removal of this overlap was done by considering overlapping clusters rather than individual pixels. First, the overlap between two layers d_n and d_m was evaluated. All independent, simply

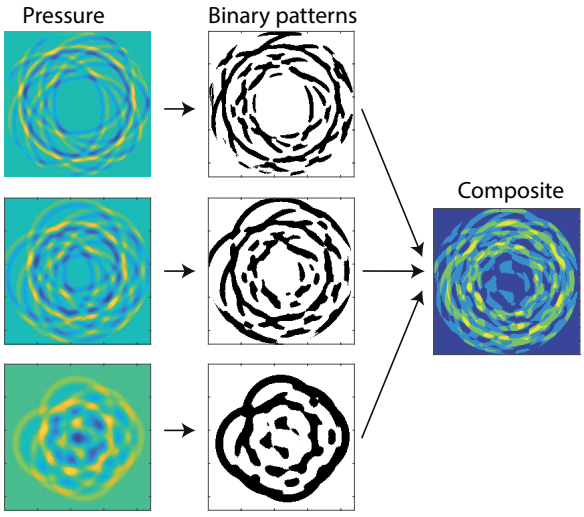

Figure 5.4: Thresholding approach used to initialise optimisation. First column shows the pressure generated by each target point across each absorbing layer d_l at a time t_{thres}. This is thresholded about 0 to create the initial binary patterns, this is shown in the second column. Using this approach, however, generates patterns with substantial overlap that has to be removed, this is shown in the third column.

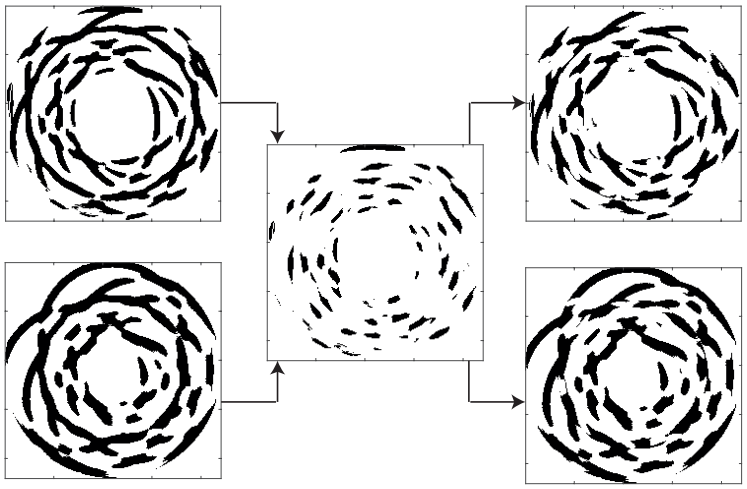

Figure 5.5: Figure illustrating the process used to remove the overlap between layers. The overlap between two layers is calculated. Simply connected clusters on this overlap are extracted, these are sequentially assigned to the layer on which they minimise the cost function in Eq. 5.9. This is repeated for every layer pair until all overlaps are removed.

connected, groups of pixels (clusters) on this overlap were then identified. The cost of the pattern H_{ijl} was analysed for each cluster sequentially when it was 'on' on d_n or d_m. Each was assigned to the layer on which it generated the lowest cost. The testing order for the clusters for each layer was selected randomly. This process, illustrated for a pair of layers by Fig. 5.5, was repeated for every pair of layers.

The choice of the time point t_{thres} was also optimised. For too early a time, waves from the target points won't have had time to propagate to each of the layers. For too late a time, the contribution of the pixels will be weaker due to repeated reflections. To optimise this parameter, an exhaustive search was done. The pattern was initialised separately for each t_{thres} and the time which generated the lowest initial cost was used. In an ideal case, the time t_{thres} could be varied for each individual target point. However, this is impractical to evaluate exhaustively where the pattern consists of more than a few points.

5.2.3.4 Direct search

From the initialised state H_{ijl}, an optimisation step was used to find further iterative improvements. As in Chapter 4, the state of individual pixels on H_{ijl} were perturbed, the change in cost ΔC was evaluated, and the change kept if the solution improved ($\Delta C < 0$). However, again due to the requirement discussed in Sec. 5.2.3.3, rather than test every pixel only a subset of pixels (x_i, y_j) were examined on each layer d_l. This subset was determined, separately for each iteration, as any pixel for which

$$H_{(i+1)jl} + H_{(i-1)jl} + H_{i(j+1)l} + H_{i(j-1)l} - 4H_{ijl} \neq 0. \tag{5.12}$$

This selected any pixel, on or off, on the edge of an existing cluster. Where a pixel (x_i, y_j) was on the edge of clusters on two or more layers, the cost was evaluated for each possible state (i.e., on for each layer and off), and the pixel was assigned the state that minimised Eq. 5.9. The iterations continued until the number of pixel positions changed in an iteration fell below 0.5% of the pixels of the aperture. The pixel selection was, again, random.

5.3 Experimental validation

5.3.1 Absorber properties

To validate that the proposed multi-layer absorbers could be used to generate patterned acoustic fields, a test case was fabricated and its acoustic field characterised. The dimensions of the absorber for this test case were 50×30×6 mm with a pat-

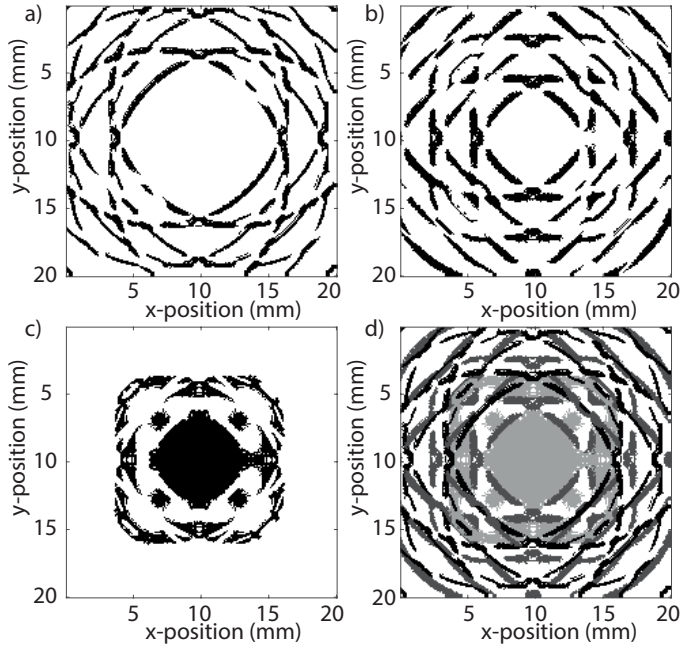

Figure 5.6: Figure showing the 3 patterns output by the optimisation for the experimental test case. (a) First layer of calculated absorber. (b) Second layer of calculated absorber. (c) Third layer of calculated absorber. (d) Composite of calculated absorber, with different layers represented by different shades. The lack of overlap can be seen. Figure reprinted from [175] © IEEE 2016.

tern aperture of 20×20 mm centred on the middle of the absorber. It had 3 distinct absorbing layers. These were spaced every 2 mm with the first layer occurring directly above the substrate. The absorber was composed of PMMA, which was taken to have a sound speed of 2750 $\mathrm{m\,s^{-1}}$ and a density of 1190 $\mathrm{kg\,m^{-3}}$.

The target pattern was a 3×3 array of evenly spaced points with a separation of 3 mm. This array was located in a plane parallel to the absorber, centred on the aperture, 18 mm from the top layer. The signal generated from the absorber $p(t)$ was taken to be a broadband bipolar pulse filtered at 3.5 MHz. The algorithm described in Sec. 5.2 was used to calculate the pattern for each layer. The resulting set of patterns are shown in Fig. 5.6.

5.3.2 Absorber fabrication

The absorber was fabricated from a set of three 50×30×2 mm transparent PMMA slides. These were each coated with a thin absorbing layer via aerosol deposition (Super Satin, Plasti Kote, Valspar, US). A laser cutter (VLS4.60, Universal Laser Systems, US) was used to etch the appropriate pattern onto each slide. These slides

were then glued (Serious Glue, Bostik, Wisconsin, US) together in the appropriate order to create the sample. A photograph of the fabricated sample is shown in Fig. 5.7 (a).

5.3.3 Field measurements

A fibre coupled Q-switch Nd:YAG laser (Ultra, Big Sky Laser Technologies, Bozeman MT) was used to excite ultrasound from the fabricated absorber. This had a wavelength of 1064 nm, a pulse length of 8 ns, a repetition rate of 20 Hz, and a pulse energy of 50 mJ. Measurements of the acoustic field were made using a previously described optical scanner [177]. This uses a planar Fabry-Perot thin film interferometer for acoustic detection over a 2-D plane. Acoustic waves incident on the sensor induce changes in its optical thickness, which modulates the reflectivity of the interferometer. This modulation is detected by a continuous wave interrogation laser focused on the plane of the sensor. The laser is scanned point by point over the sensor to map a slice of the acoustic field.

A photo of the experimental set up can be seen in Fig. 5.7 (b). The absorber was suspended in water 10 mm above the surface of the sensor and parallel to it. The laser fibre tip was then placed 20 cm above the absorber and manually aligned with its centre. The diameter of the beam at the absorber surface was approximately 20 mm. Signals were recorded over a 2.46×1.99 cm area centred on the absorber using a step size of 0.15 mm. Four averages were taken at each position. To calculate the 3-D wave-field from the planar measurements, k-Wave was used to both forward and back propagate the planar measurements by 25 mm assuming linear propagation. Prior to this, the dataset was low pass filtered at 3.5 MHz to match the design bandwidth, and high pass filtered at 300 kHz to eliminate slowly varying background noise.

5.3.4 Numerical experiments

The experimental results were compared against a simulation of the acoustic field generated by the same set of patterns. This simulation was carried out using k-Wave on a $534 \times 324 \times 324$ grid, with a grid spacing of 0.1 mm. The medium for

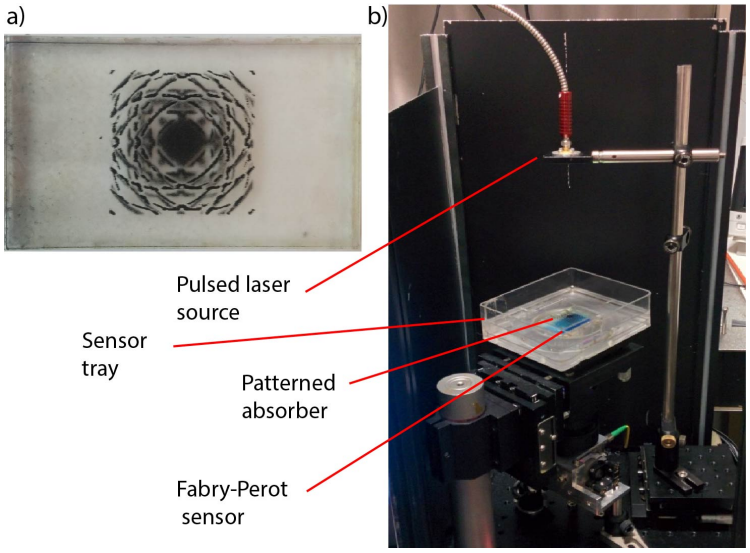

Figure 5.7: (a) Photo of fabricated multi-layer absorber. (b) Annotated photo of experimental set-up. Figure reprinted from [175] © IEEE 2016.

this simulation was inhomogeneous. The absorber was inserted at one end of the simulation domain as a $500\times300\times60$ region with sound speed of 2750 $\mathrm{m\,s^{-1}}$ and density of 1190 $\mathrm{kg\,m^{-3}}$. The rest of the medium had a sound speed of 1500 $\mathrm{m\,s^{-1}}$ and a density of 1000 $\mathrm{kg\,m^{-3}}$ to represent water.

The three patterned layers were inserted as source masks over 3 simulations planes. The first was placed 1 grid point above the absorber, and the second and third were placed at 20 and 40 grid points within the absorber. The pressure $p(t)$ generated from the absorber was represented by a broadband bipolar pulse high-pass filtered at 3.5 MHz. As in Sec. 4.5, a correction was applied for the spatial profile of the laser. To do this, the experimental data was back-propagated to $t = 0$ to recover the initial pressure distribution. Each of the absorber planes were then identified in this initial pressure distribution and registered to the binary patterns displayed in Fig. 5.6 (a)-(c). This registration step is shown in Fig. 5.8. The registered spatial profiles were used to weight the source mask in the forward simulation for a better comparison with the experimental data.

5.3.5 Results

The maximum pressure generated in the target plane of the multi-layer absorber in both the experimental and simulated data is shown in Fig. 5.9. It both cases the

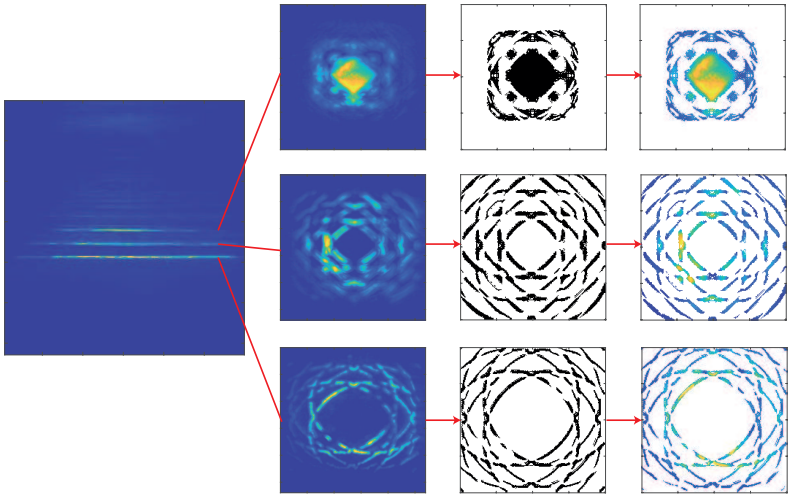

Figure 5.8: Process used for the simulation to correct for the spatial profile of the laser. Time series are back-propagated to t = 0, and the absorber planes are identified in the initial pressure distribution. The planes are registered using an affine transform to the calculated binary patterns. The resulting spatial profiles are then used to weight the source in simulation.

targeted 9 foci are clearly realised and there is good qualitative agreement between the simulation and experiment. This confirms that multi-layer absorbers can be used to generated multi-focal acoustic fields using a single optical pulse.

In both datasets, however, the pressure at the central foci is significantly higher than at the surrounding target points. This is due to the Gaussian spatial profile of the laser, which introduced a weighting for pixels contributing to the central focus. A high background can also be seen surrounding the central focus. This is also, in part, generated by the laser profile. However, it is principally due to the low focal gain at the target points. This is a result of the large layer spacing and low number of absorbing layers within the sample. This limits both the number of pixels on the aperture able to contribute a signal and the amplitude contributed by each point.

The layer spacing was constrained by the thicknesses available for the PMMA substrate and could be improved by employing a thinner transparent substrate. However, the fabrication method has other drawbacks. Namely, registration errors are introduced between the different layers, and the use of glue introduces weak acoustic reflections between each absorbing layer. The reflections from between the layers can be eliminated by using a solvent to chemically weld the layers. However, this was found to warp the pattern of the absorbing layer as the underlying

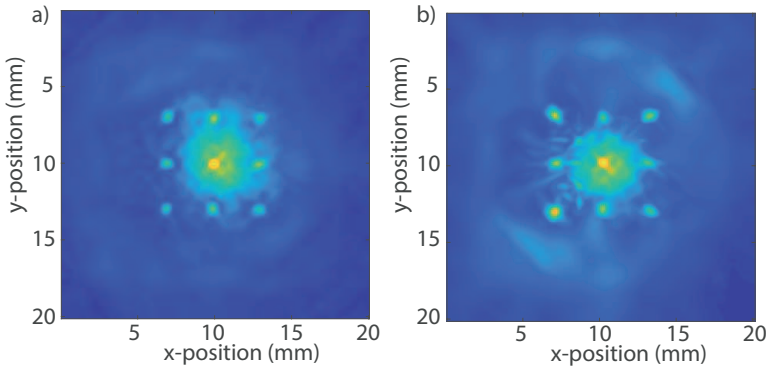

Figure 5.9: (a) 9 foci generated in the target plane of the absorber in the experiment. (b) 9 foci generated in target plane of absorber in simulation correcting for apodisation. Figure reprinted from [175] © IEEE 2016.

substrate was dissolved. These drawbacks make identifying other methods for fabricating these multi-layer absorbers a necessity.

5.4 3-D printing

One alternative is to directly 3-D print an absorber using two materials, one optically absorbing to form the patterned layers and one transparent to form the substrate. With this approach, a multi-layer structure can be constructed, cheaply, for which the layer spacing is limited only by the resolution of the 3-D printing process ($\sim$100 μm for current systems). The layers would also be automatically registered with low error. As a proof of concept for this approach, a test case was fabricated. The absorber in this case was a cylinder with a radius of 12 mm, and a height of 5 mm. This was discretised using a pixel spacing 100 μm and was approximated as having 50 absorbing layers evenly spaced every 100 μm within the sample. The target pattern was a square 2$\times$2 array of foci with a width of 7.1 mm. This was centred on the absorber, in a plane parallel to it, 25 mm from the top layer of the sample. The acoustic signal $p(t)$ was a bipolar pulse high-pass filtered at 6 MHz.

In order to fabricate the sample, the calculated H_{ijl} had to be converted into STereoLithography (STL) file format. This was done using an open source MATLAB function `CONVERT_voxels_to_stl` [178] able to directly convert binary 3-D matrices into STL files with dimensions based on supplied voxel widths. Initially, two files were prepared, one for the optically absorbing layers and one for

the transparent substrate. This was formed by subtracting the optically absorbing structure from the absorber volume. However, structures in STL format have to be simply connected, so printing H_{ijl} in this way requires a separate file for each cluster (> 1000). To circumvent this limitation, the approximation discussed in Sec. 5.2.3 can be used. If the optical absorption in each layer is sufficiently high then the incident light is almost entirely absorbed within the first absorbing voxel it interacts with at each position. This implies that the composition of the structure after the initial voxel doesn't effect photoacoustic signal, allowing H_{ijl} to be extruded after this depth to form a simply connected volume A_{ijl} using

$$A_{ijl} = \begin{cases} 1 \text{ if } l \geq M_{ij}, \\ 0 \text{ if } l < M_{ij}. \end{cases} \tag{5.13}$$

Here M_{ij} is a map of the layer l across the aperture for which H_{ijl} is 1, where the bottom layer has an index of 1. The corresponding volume for the transparent material T_{ijl} is then given by

$$T_{ijl} = \begin{cases} 1 \text{ if } A_{ijl} = 0, \\ 0 \text{ if } A_{ijl} = 1. \end{cases} \tag{5.14}$$

After calculating the volumes T_{ijl} and A_{ijl}, the sample was fabricated using a high-resolution polyjet printer (Objet250 Connex, Statasys, Eden Prairie, MN, USA), which has simultaneous two-material printing capability. The two materials used were veroClear, a rigid transparent material, and tangoBlack, an optically absorbing rubber-like material. The resulting multi-layer absorber along with both calculated volumes and the depth map M_{ij} can be seen in Fig. 5.10 (a-d). The acoustic field generated by this sample was then measured using the set-up described in Sec. 5.3.4. The field was recorded over a 22.92×22.56 mm area using a step size of 0.12 mm taking four averages at each position. This data was spatially up-sampled by a factor of 2, bandpass filtered between 0.3-6 MHz, and forward and back projected by 20 mm.

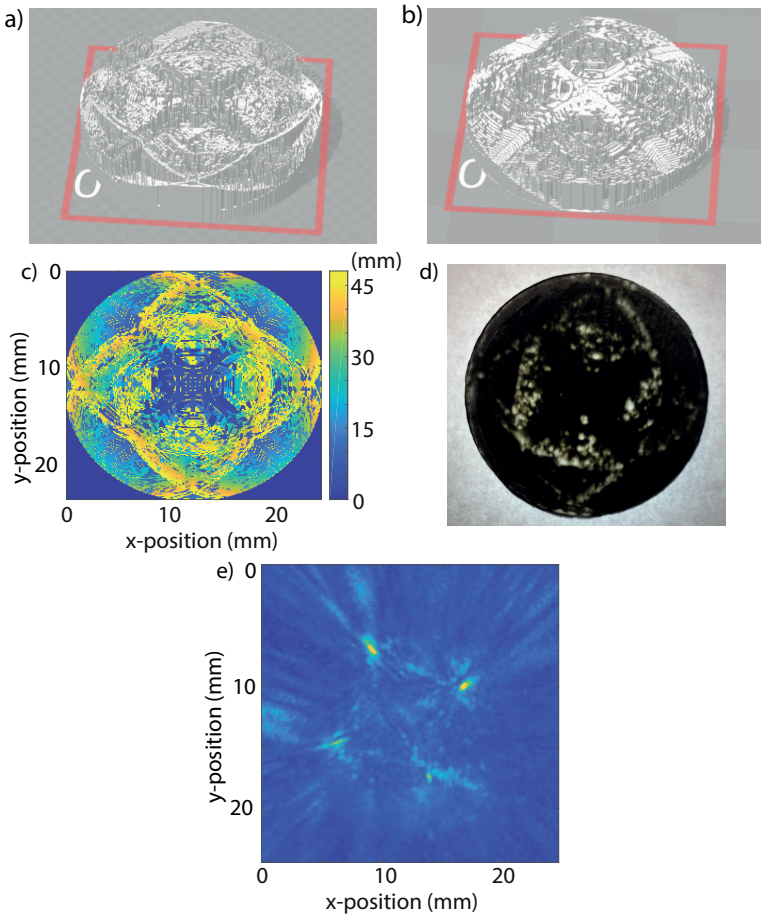

Figure 5.10: (a) Rendering of the structure of the transparent volume of the 3-D printed multi-layer absorber. (b) Rendering of the structure of the optically absorbing volume of the 3-D printed multi-layer absorber. (c) Depth map across aperture of layer for which H_{ijl} is 'on'. (d) Photo of the fabricated 3-D printed sample. (e) Field generated in target plane of sample when excited by a single optical pulse. Four foci are visible, however, the pressure generated varies significantly.

The maximum pressure generated in the target plane of the absorber in the projected data is shown in Fig. 5.10 (e). The four target foci have been successfully generated by the 3-D printed sample, confirming that the 3-D printing process has replicated the target structure. However, the discrepancy in pressure between the different foci is high, and the background surrounding each foci is also significant. These aberrations were the result of two factors. First, the sound speed of tangoBlack varies by 20% compared to veroClear. This introduces reflections off the internal structure and refraction that were not accounted for in the evaluation of the impulse responses. Second, the acoustic attenuation of the rubber-like tangoB-

lack is high ($\sim$11 dB/MHz/cm). This will have significantly attenuated the higher frequency content of the signal from the lower layers in the sample.

While these drawbacks disrupted the field of this test sample, both were introduced by the choice of optically absorbing material tangoBlack. This was determined by the optical properties of the materials available on the polyjet printer used to fabricate the sample. Other 3-D printable materials are commercially available, with identical acoustic properties to veroClear, that are optically absorbing. In the future these could be used to fabricate a homogeneous multi-layer sample with much lower acoustic absorption.

5.5 Conclusion

This Chapter has demonstrated a new technique for optically generating multi-focal acoustic fields using a single optical pulse. This uses optical absorbers made up of several spatially separate absorbing layers with each layer individually patterned such that the acoustic signal excited across it constructively interferes at a set of target points. This circumvents the requirement of most previous techniques for multiple costly optical sources or external amplitude modulation and has the potential to make the application of the optical generation of focused fields much simpler.

A ray-tracing approximation has been described for the rapid approximation of the impulse response of pixels in each layer to arbitrarily placed target points. This model avoids the need to run a costly full-wave model and has been validated by comparison with k-Wave. The direct search approach described in Chapter 4 has also been adjusted for designing the patterns within these multi-layer absorbers. This uses a new forward model based on the calculated impulse responses and a new initialisation that thresholds the sum of the pressure generated by each target point across each absorbing layer at a particular time point. The initialisation is followed by the removal of any overlap between layers based on overlapping clusters and an optimisation step that finds further refinements for each pattern through minimisation of a cost function.

The technique and optimisation approach has been experimentally validated

using two sets of measurements. The first used a 3-layer sample fabricated by gluing together several patterned absorbers fabricated through laser cutting. The second used a 50-layer sample fabricated by two-material 3-D printing. The targeted field was generated in both cases, however, the pattern generated from the 3-D printed sample was distorted by aberrations introduced by sound speed variations between the two materials.

One drawback of this technique is that, at present, each pattern is fixed. Dynamically varying the pattern as considered in Chapters 3 and 4 is also more complex due to the multi-planar nature of the source. However, it could, in principle, be achieved by using multiple optical wavelengths combined with a multi-wavelength phase hologram displayed on an SLM [179]. This would need to be combined with materials for each layer that absorb over a specified bandwidths. Another interesting possibility is to design multi-layer absorbers with arbitrary substrate shapes. Since the photoacoustic pressure is generated inside a homogeneous substrate, the shape of the substrate and, importantly, the surface can be varied arbitrarily. This could in theory be used to create focused transducers that conformed to particular objects for use in industrial evaluation. In addition to these experimental developments, the tolerance of the method to different fabrication errors should also be investigated numerically. For example, the effect of systematic errors in layer thickness, the assumed sound speed, or in the registration of the layers should be characterised.

Chapter 6

Tailored optoacoustic surface profiles

6.1 Introduction

The previous Chapter demonstrated that multi-focal acoustic fields could be generated from a single optical pulse by using optical absorbers composed of multiple individually patterned layers. This ability to generate focused fields with a single pulse has a number of potential applications, however, the multi-layer approach has drawbacks. Most significantly, the need to account for refraction and reflections within the substrate requires running a costly forward model, and, at present, the performance is limited by the large layer spacing that must be used to realise the absorbers experimentally.

An alternative approach, discussed in Sec. 2.3.4, is to instead manipulate the shape of the surface onto which the optical absorber is deposited. This allows the shape of the optically generated wavefront to be controlled, which can, in principle, be tailored such that almost any desired wavefield is generated. This concept is illustrated by Fig. 6.1. Several works have demonstrated the generation of specific fields such as single foci or non-diffracting x-waves using this approach. However, a general method for designing a surface for any arbitrary input target field has yet to be reported. In this Chapter, an optimisation approach based on direct search is developed to solve this problem.

In the next section, the design problem is outlined, and the forward model and cost function are introduced along with an initialisation tailored to the problem. An

extension to the optimisation is then presented that allows the temporal variation in peak pressure across the target pattern to be controlled. Then, in Sec. 6.3, the design algorithm is validated using a surface profile designed to generate the numeral '7'. This is fabricated by 3-D printing the substrate using a transparent plastic followed by aerosol deposition of an optical absorber. Finally, in Sec. 6.4 the conclusions of the work are presented along with several future research directions.

The journal article in [180] has been modified and adapted to form parts of this chapter, with reprint permission under CC BY 4.0 License.

6.2 Design algorithm

6.2.1 Introduction

The goal in designing a tailored optoacoustic surface profile is similar to that explored in Chapter 5 for the design of multi-layer absorbers. Briefly, an optically absorbing surface profile $h(x,y,z)$ is sought that, when illuminated by a single optical pulse, will generate a high peak pressure across a set of K target points $r_k = \{(x_k, y_k, z_k), k = 1, .., K\}$. Compared to the multi-layer absorbers however, the problem is simplified as the source is generated on a substrate directly coupled to

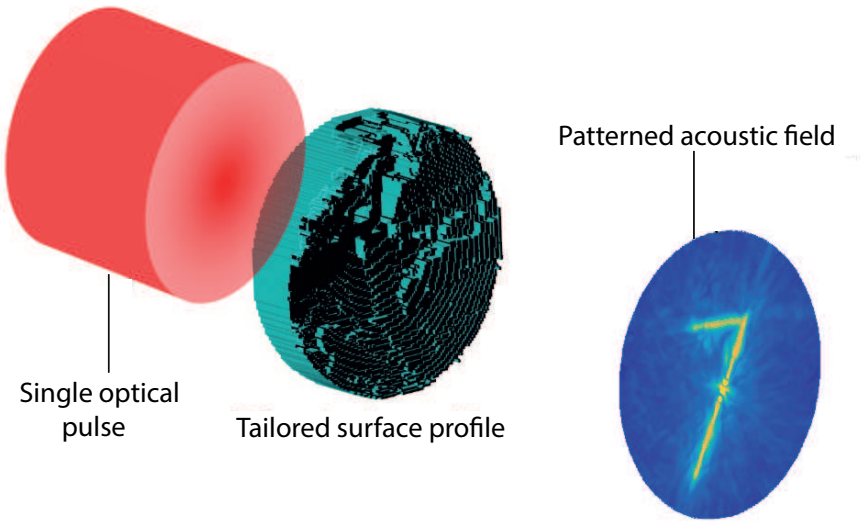

Figure 6.1: Illustration of the use of tailored surface profiles for the generation of patterned optically generated acoustic fields in 3-D. Figure reprinted from [180] under CC BY 4.0 License.

the surrounding medium. This allows the generation and subsequent propagation of sound to be modelled as occurring in free-space. This approximation, however, is only accurate if the height of the surface varies smoothly and slowly. For 'rough' surfaces with rapid oscillations in height, the generated sound will reflect off the substrate causing the approximation to breakdown. This requires the design algorithm to be constrained to avoid convergence to surfaces possessing these rapid oscillations.

Rather than make the free-space approximation, an alternative method would be to use k-Wave to model interactions with the substrate. However, the direct search approach requires too many calls of the forward model ($\sim 10^6$) for this to be feasible. Additionally, there are other practical reasons for requiring a smooth surface. Namely, depositing an absorber evenly on a surface with rapid jumps in height would be difficult, particularly with an aerosol based method, and, the physical structure of the substrate could potentially become unstable depending on the material properties.

6.2.2 Forward model and cost function

In free space, the acoustic signal $p(\mathbf{r},t)$ at a point $\mathbf{r} = (x,y,z) \in \mathbb{R}^3$ and time t due to a photoacoustic source distribution $p_0(\mathbf{r})\delta(t)$, can be written as [80]

$$p(\mathbf{r},t) = \frac{1}{c_0^2}\int_\Omega p_0(\mathbf{r}')\frac{\partial G}{\partial t}(\mathbf{r}-\mathbf{r}',t)d\Omega(\mathbf{r}'), \tag{6.1}$$

where c_0 is the sound speed, p_0 is compactly supported in $\Omega \subset \mathbb{R}^3$, and G is the free-space Green's function for the wave equation given by

$$G(\mathbf{r}-\mathbf{r}',t) = \frac{\delta(|\mathbf{r}-\mathbf{r}'|-ct)}{4\pi|\mathbf{r}-\mathbf{r}'|}. \tag{6.2}$$

In this Chapter, Ω is restricted to a thin region of thickness δ around a surface defined as follows. Consider a square region of the xy plane, $[-L/2,L/2]\times[-L/2,L/2]$, tiled by N^2 square pixels of side $\Delta = L/N$ centred on the coordinates (x_i,y_i). Each square (i,j) is displaced in the z-direction to a height h_{ij}. The surface

around which p_0 is defined is the discontinuous surface formed by the set of squares centred at (x_i, y_i, h_{ij}). To define other apertures, the set of squares is restricted to those that lie within the desired shape. The surface heights are also discretised such that $h_{ij} = n\Delta_z$, $n \in \mathbb{N}$.

The pressure at each target point was evaluated for each p_0 by numerically solving Eq. 6.1. The surface was approximated by a set of point sources of amplitude A located at (x_i, y_j, h_{ij}), where $p_0 \approx \sum_{i,j} A_{ij}\delta(x - x_i, y - y_i, z - h_{ij})$. Here the amplitudes A_{ij} were introduced to account for the spatial profile of the optical illumination. The signal $p(t)$ generated by each point source was calculated by using k-Wave to simulate a photoacoustic point source on a grid with spacing Δ. This signal was then shifted in time and modulated in amplitude for each target point r_k and source point (x_i, y_j, h_{ij}), based on the distance, using the free-space Green's function (Eq. 6.2).

The cost function for the optimisation was then defined as

$$C = -\left(\frac{1}{K}\sum_{k=1}^{K} \max_t(p(r_k,t))\right) + \alpha \mathrm{std}(\max_t(p(r_k,t)) + \beta S. \tag{6.3}$$

The first two terms are identical to those in Eq. 5.8, aiming to maximise the average temporal peak pressure and minimise the variation in temporal peak pressure across the target points, respectively. The variables α and β are weighting parameters for the different terms, and S is a smoothness term designed to penalise large variations in height between neighbouring pixels. This was introduced due to the requirement, discussed in Sec. 6.2.1, that the surface avoid large local variations in height. This smoothness term S was given by a second-order-accurate central difference approximation of the Laplacian.

$$S = \sum_{j=1}^{N}\sum_{i=1}^{N} \nabla^2 h_{ij} = \frac{h_{(i+1)j} + h_{(i-1)j} + h_{i(j+1)} + h_{i(j-1)} - 4h_{ij}}{4\Delta^2}. \tag{6.4}$$

The two weights α and β in Eq. 6.3 were assigned heuristically by calculating profiles for a range of values of α and β and simulating the fields they generated,

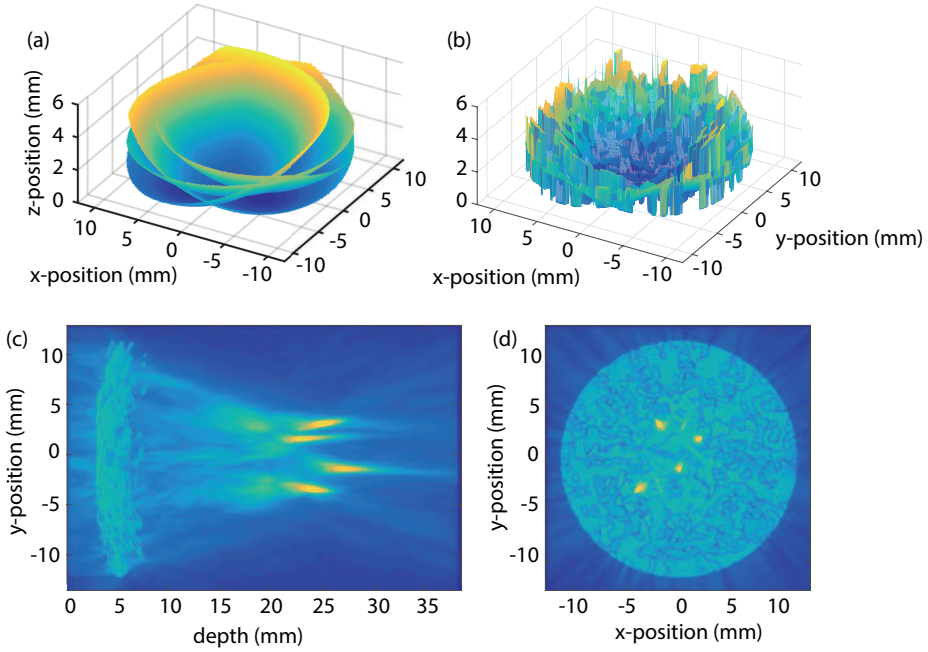

Figure 6.2: (a) Initial state of a pattern consisting of a set of overlapping bowls centred on each of the target foci. (b) Output of the optimisation. (c-d) Maximum amplitude projections (MAP) through the y-z plane (c) and the x-y plane (d) of the acoustic field generated by the profile in (b) when it is illuminated by a single optical pulse. Simulated using the k-Wave toolbox. Figure reprinted from [180] under CC BY 4.0 License.

with backing material, using k-Wave. The target distribution for this test profile consisted of four foci distributed across multiple planes and is shown in Fig. 6.2 (c) and (d). The test found that acceptable values for alpha ranged from $0.4-3$, while acceptable values for β ranged from $(0.1-3)\eta$, where $\eta = \frac{1}{SK}\sum_{k=1}^{K} \max_t(p(r_k,t))$. The variable η was evaluated after initialising the surface as described below. The profiles in the remainder of this Chapter were calculated using a value for α of 1 and a value for β of η.

6.2.3 Initialisation

The optimisation was initialised by choosing h_{ij}^0 in the following way. First, a set of spherical sections (bowls) centered on the focal points r_k were defined as the surfaces $b_k = \{(x-x_k)^2+(y-y_k)^2+(z-z_k)^2 = z_k^2 \,|\, x^2+y^2 \leq L^2/4\}$. Second, the union of these surfaces, $\cup\{b_k\}$, was sampled at the source points (x_i, y_j) and rounded to the nearest $n\Delta_z$ to give h_{ij}^{multi}, which is multivalued with up to K possible

surface heights for each source pixel (i, j). In the third step of the initialisation, one of these values must be chosen for each source pixel, as the absorbing surface must be single-valued.

The selection process involved taking each pixel (i, j) in turn, at random but not repeating any pixel, and considering a 5×5 pixel region around it. If any pixels in this region had already been assigned a single-valued height, the value of h_{ij}^{multi} chosen was that which minimised the average difference to these surrounding heights, with one condition. It was found, in simulation, that large planar regions with the same value of h_{ij} generated a high peak pressure at locations away from the targeted focal points. To ameliorate this, only height differences with the neighbouring pixels of $> 3\Delta$ were considered. When multiple values of h_{ij}^{multi} had equal difference to the surrounding heights, or when no pixels in the surrounding region had yet been assigned single-valued heights, the value of h_{ij}^{multi} that minimised the sum of the first two terms in Eq. 6.3 was chosen. The result was a single-valued initial surface profile, h_{ij}^{0} .

In defining the multi-valued surface, if the foci were distributed over multiple planes an additional step was incorporated. The area occupied by each spherical section was reduced, starting from the edge, by a fraction determined from the ratio of the target point's depth to the deepest target focus. This was to ensure that the peak pressure generated at each target point by h_{ij}^{multi} was approximately equal. Otherwise, it was found when selecting h_{ij}^{0} from h_{ij}^{multi} that points with lower peak pressure were prioritised when no neighbouring points had been assigned, due to the second term in Eq. 6.3. The effect of this was to trap the surface later in the initialisation, when the heights were principally constrained by the heights of the surrounding region, resulting in surfaces that maximised the pressure at deeper focal points.

6.2.4 Direct search

Following the initialisation, iterative refinements to the surface heights h_{ij} were obtained using the direct search approach. Equation 6.3 was minimised, within a specified range of heights, for each pixel sequentially in random order. The range

evaluated for each position was the set of heights for which $p(r_k,t_k) \neq 0$ for at least one k, where t_k is the time at which the peak occurs. This process gave the surface profile $h_{ij}^{(1)}$. This was then iterated to give the converging series $h_{ij}^{(m)}$. The iterations were terminated when the number of pixels changed from the previous iteration fell below a threshold. This, as in previous Chapters, was set at 0.5% of the total pixels on the surface. Figure 6.2 (a-b) illustrates the initial (Fig. 6.2 (a)) and final (Fig. 6.2 (b)) state of a surface profile designed to generate 4 foci at different depths. The acoustic field generated by this profile when illuminated by a single optical pulse may be seen in Fig. 6.2 (c-d). This was simulated using the k-Wave toolbox and clearly demonstrates the generation of four discrete foci.

6.2.5 Discussion

An alternative way to choose the surface profile would be to form an array of K bowl transmitters each focussed on one focal point r_k. However, such an approach would suffer from considerably worse lateral focussing due to the solid angle subtended by each bowl at its focal point being smaller, by a factor of about K, than the solid angle available in the optimisation approach. In the optimisation approach, any pixel can contribute to the focus at any point. The optimisation approach also offers more flexibility over the resulting profiles. The relative pressure at the different target points can be constrained, as in Sec. 4.2.3, by adding a weighting vector to the cost function. It is also possible to control the relative time t_k at which the peak pressure is generated at each target point.

This control over the relative time at which the peak pressure is generated is achieved by modifying the optimisation in two ways. First, in initialising the surface, the radii of the bowl centred on each target point is adjusted such that $b_k = \{(x-x_k)^2 + (y-y_k)^2 + (z-z_k)^2 = ct_k^2 \,|\, x^2+y^2 \leq L^2/4\}$. This sets the radii of each bowl such that the time at which the focus is generated matches the targeted temporal variation. Second, the cost function is modified to:

$$C = -\left(\frac{1}{K}\sum_{k=1}^{K} p(r_k,t_k)\right) + \alpha \mathrm{std}(p(r_k,t_k))) + \beta S. \tag{6.5}$$

Rather than maximise the average temporal peak pressure, this cost function aims to maximise the pressure at each target point r_k at the specified time t_k.

Figure 6.3 (a) shows an example surface profile calculated with a constrained temporal profile. This was designed to generate two 7.5 mm perpendicular lines intersecting at one end in a plane 25 mm from the surface. The temporal profile for the peak pressure was designed to vary linearly across each line by 1.5 μs starting from the intersecting point. The acoustic field generated by this profile was simulated using k-Wave and the time varying pressure $p(r_k, t)$ generated at each target point was measured.

The peak pressure generated in the target plane of the calculated profile can be seen in Fig. 6.3 (b), where the two lines have been generated with high contrast. The time t_k at which the peak pressure was generated at each target point can be seen in Fig. 6.3 (c) alongside the targeted variation. This shows that the intended temporal profile has been accurately generated by the optimisation; the average deviation at each target point is just 7.3 ns. These small deviations arise, principally, because the cost function aims to maximise the pressure at the specified time rather than actively penalising deviations. Figure 6.3 (d) also shows the temporal variation in peak pressure across the target points, in this case plotted at the (x, y) positions from which they were measured. These results demonstrate that the temporal variation in peak pressure across the target points can be controlled by appropriately adjusting the optimisation.

6.3 Experimental validation

6.3.1 Fabrication

To demonstrate that the optimisation approach is effective, samples were fabricated via additive manufacturing. This was done using a high-resolution polyjet printer (Objet350 Connex, Stratasys, Eden Prairie, MN, USA). The substrate was printed using a transparent material (VeroClear-RGD810) that is rigid and nearly colourless. The absorbing surface was formed through aerosol deposition of an optically absorbing polymer composite (Gloss Super, Plastikote, Valspar, US).

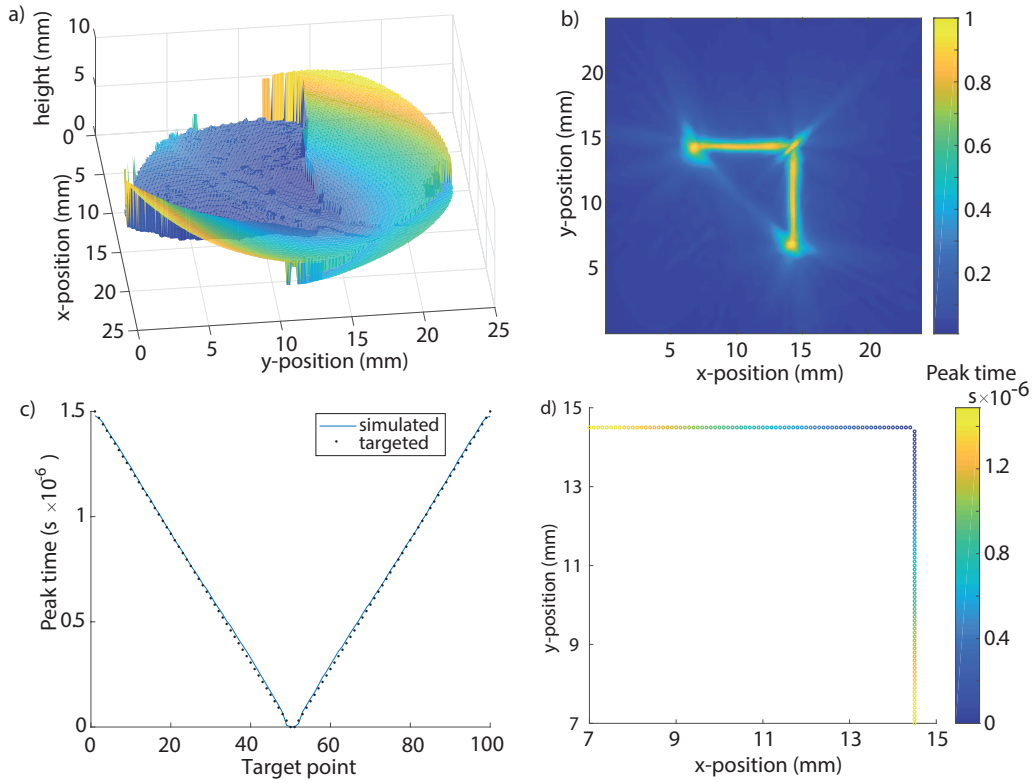

Figure 6.3: (a) Surface profile designed to generate two intersecting perpendicular lines with a pre-determined temporal variation in peak pressure. (b) Peak pressure generated by the profile in the target plane in simulation when illuminated by a single optical pulse. (c) Variation in the time t_k at which the peak pressure occurred at each target point r_k in simulation (continuous line) plotted against the targeted variation (dotted line). Both lines have had the minimum t_k subtracted to set the baseline to 0. (d) Variation in t_k plotted at the spatial coordinates at which it was measured.

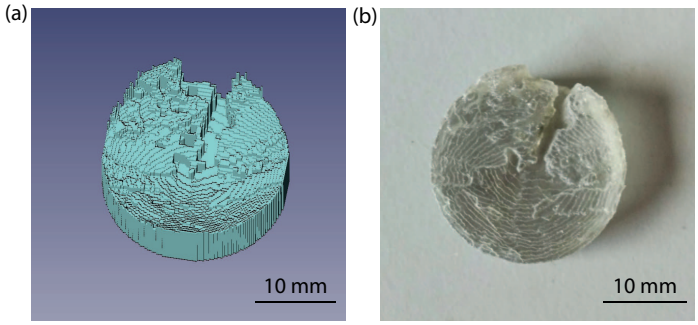

Figure 6.4: (a) Prepared surface mesh visualised using the freely-available open-source FreeCad software (Juergen Riegel, Werner Mayer, Yorik van Havre, Open-Source, freecad.com). (b) Photo of fabricated sample before absorber deposition. Figure reprinted from [180] under CC BY 4.0 License.

An example surface profile is shown in Fig. 6.4 (b). This was calculated for a circular aperture of diameter 23 mm discretised using a spacing Δ and Δ_z of 200

μm. The target pattern was a continuous distribution of high pressure resembling the numeral '7'. This was chosen as it is asymmetric and continuous and therefore represents a challenging target distribution. This makes it a good test of the limits of the inversion [181, 182]. This '7' had a height of 12 mm, a width of 6 mm, and was placed in a plane 25 mm from the bottom of the surface profile. The '7' was discretised using 61 evenly spaced points for input into the optimisation. It was assumed that the laser pulse incident on the surface had a top-hat profile. The forward model in Eq. 6.1 was solved accounting for frequencies up to 6 MHz. To prepare the sample from the calculated surface profile, backing material was added to each position up to the calculated height. A further 2 mm was added across the bottom of the surface to provide rigidity. The backing material is optically transparent and acoustically absorbing so this layer can be increased in thickness arbitrarily to eliminate the backward travelling wave. The surface mesh and the 3-D printed transparent sample can be seen in Fig. 6.4 (a-b).

6.3.2 Measurement

The field generated by the fabricated surface profile was characterised experimentally. A fibre coupled Q-switched Nd:YAG laser (Minilite, Continuum, San Jose, CA, USA) was used to excite ultrasound from the fabricated sample. The process used to couple this laser is described in Appendix 1. The optical pulses had a wavelength of 1064 nm, a pulse length of 6 ns, pulse energies of $\sim$14 mJ, and an approximately Gaussian profile. Measurements of the acoustic field were made using the optical scanner described in Chapter 5.

A schematic of the set-up used is shown in Fig. 6.5. The fibre was attached to a circular tube with a length of 10 cm and a diameter of 2.5 cm. The natural divergence of the laser beam expanded the beam to 2.5 cm by the end of the tube. The surface profile was attached to a retaining ring which was connected to the other end of the tube. The front face of the surface profile was then immersed in water approximately 1 cm above the optical scanner. Signals were recorded over a 22.20$\times$22.05 mm area using a step size of 0.15 mm. Four averages were taken at each position. The 3-D wavefield was calculated from the experimental measurements using linear

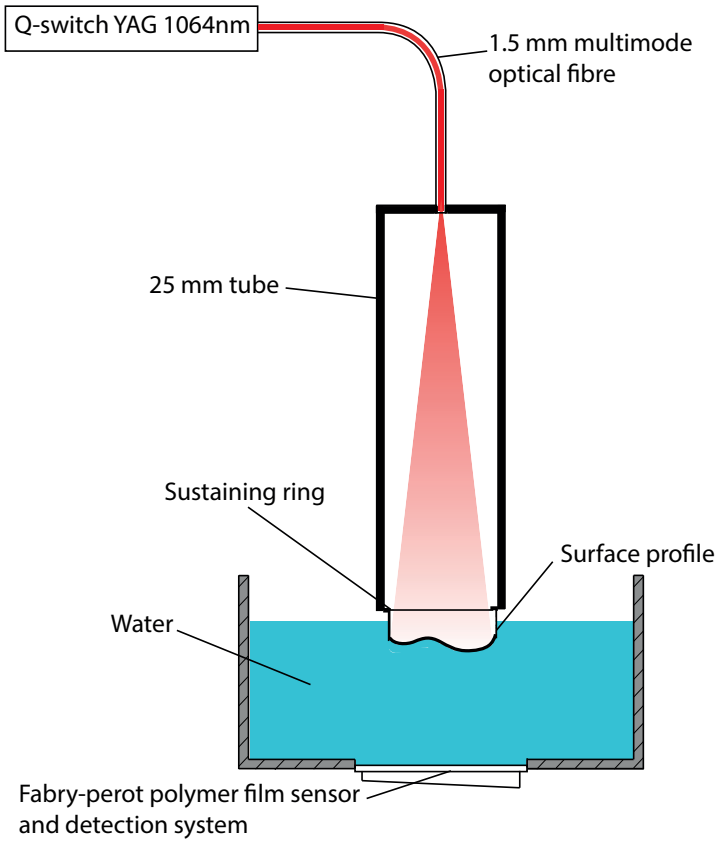

Figure 6.5: Schematic of experimental set-up used for the characterisation of the acoustic fields generated from the tailored optoacoustic surface profiles. Figure reprinted from [180] under CC BY 4.0 License.

acoustic holography. Prior to this the data was band-pass filtered between 30 kHz and 6 MHz and was spatially up-sampled by a factor of 2. The filter was applied to remove low frequency noise introduced by the detection system and to limit the measurement data to the design bandwidth.

The experimental measurements were compared against a simulation carried out using the k-Wave toolbox. The simulation grid was $288 \times 288 \times 384$ grid points with a spacing of 0.1 mm. The calculated h_{ij} was up-sampled by a factor of 2 in (x, y) such that each position (i, j) was represented by 4 pixels, and then inserted as a source mask in the simulation. The medium for the simulation was inhomogeneous. The backing material VeroClear was inserted as a region with a sound speed of 2495 m s^{-1} and a density of 1190 kg m^{-3}. The rest of the medium was set to a sound speed of 1500 m s^{-1} and a density of 1000 kg m^{-3}. A bipolar pulse filtered at 6 MHz was used to represent the acoustic signal generated from the sample. The spatial profile

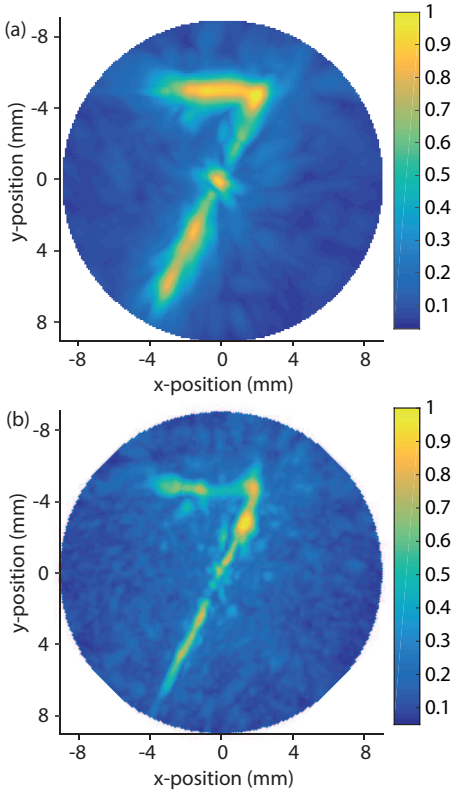

Figure 6.6: (a) Normalised peak pressure in the target plane of the surface in the simulation. (b) Normalised peak pressure in the target plane of the surface in the experimental data. Figure reprinted from [180] under CC BY 4.0 License.

of the laser was approximated by adding a 2-D Gaussian apodisation to the source. This had a standard deviation of 5 mm.

6.3.3 Results

The peak pressure in the target plane of the surface for both the experimental data and the simulation may be seen in Fig. 6.6. In both cases, the numeral 7 has been clearly realised and the two show excellent agreement in the overall shape of the pattern, and in the low background pressure. The focal gain of the surface profile in the experimental data is 2. This gave pressures of approximately 60 kPa across the target pattern. The signal to noise ratio (SNR) in the target plane of the experimental data is 4.1. This was quantified by creating a mask encompassing the target pattern and calculating the ratio of the average peak pressure within the mask to outside it. The time window between the first and last occurrence of a peak across the target pattern was 2.6 μs. There are however noticeable differences in the amplitude

generated across the 7 in both datasets. As in previous Chapters, these are generated by variations in the apodisation of the sample in the experimental set-up compared to the simulation. However, in this case, this can't be corrected via registration of the back-propagated pressure due to the complex shape of the source. In the future, pre-measurement of the apodisation A_{ij} in the experimental set-up for input into both the optimisation and the simulation would eliminate this variation.

6.4 Conclusion

This Chapter has demonstrated that complex patterned distributions of pressure can be generated with a single optical pulse by using tailored optoacoustic surface profiles. An optimisation approach has been developed for the design of these profiles from an arbitrary set of input foci. This builds on the direct search approach reported in earlier Chapters, incorporating a new initialisation that ensures the surface is initialised in a smooth state that generates a high peak pressure across the target points. A modified version of this optimisation that allows for temporal control over the peak pressure has also been reported and numerically validated. This adjusts the cost function to reward increases in the pressure at the specified time points. This could also be applied to the algorithms described in Chapters 4 and 5. Finally, the design method has been experimentally validated and it has been shown that the calculated profiles can be fabricated using widely available, cheap, techniques such as 3-D printing.

Several factors determine the complexity of the acoustic fields that can be generated and the focal gain and SNR that can be achieved with this method. The surface profile consists of a set of partial bowl transducers each centred on a target point and taking up a fraction of the aperture. The focal gain is therefore fundamentally limited by the aperture size, the centre frequency, and the total area occupied by the target pattern. However, the optimisation allows certain positions on the aperture to contribute to the peak pressure at multiple points, so the precise scaling varies with the distribution of target points and the bandwidth of the acoustic signal. In addition to these factors, the performance also improves with smaller pixel size

and the sizes of features that can be realised in the pattern are diffraction limited.

There are several directions along which the work in this Chapter could be developed. First, the profiles could be combined with a modulated optical source for the generation of a narrowband spectra. This could enable the creation of optoacoustic phase holograms by adapting the approach reported by Melde et al for single element piezoelectric transducers [17]. Second, the use of optimised nanocomposite absorbers in-conjunction with the profiles could be investigated for increasing the achievable pressures. Chan et al [157] recently demonstrated that these composites can be spin-coated onto simple 3-D printed surfaces, however, the viability of this method for more complex shapes needs to be established. Finally, different practical applications for the profiles could be explored. Focused laser generated transmitters are at present being investigated for use in cell cleaving and micro-bubble mediated drug delivery [92]. The method proposed in this Chapter would enable more precise shaping of the focus, or, for correction of medium-induced wavefront aberrations, improving the efficacy of these techniques.

Chapter 7

Multi-frequency acoustic kinoforms

7.1 Introduction

The work in this thesis thus far has addressed the generation of arbitrary acoustic fields using the optical generation of ultrasound. This modality has advantages namely, wide bandwidths, and ease of spatial modulation. However, at present, limitations on existing optical sources constrain the applications to which these methods can be applied. In this Chapter, the generation of arbitrary distributions of pressure using single element piezoelectric transducers is explored instead. This builds directly on work by Melde et al discussed in Sec. 2.4. Briefly, Melde developed a new approach for generating complex, diffraction limited, acoustic fields from planar single element piezoelectric transducers [17]. This used 3-D printed surface profiles, dubbed acoustic kinoforms, that, when attached to the front face of a transducer, introduced a thickness dependent phase offset at each position due to variations in sound speed. These were used to map the output field of the transducer onto different computer generated phase holograms, allowing the field to be diffracted onto different pre-determined acoustic fields.

This approach is extremely cheap to implement, foregoes the complicated driving electronics required for 2-D arrays, and scales easily to MHz frequencies. However, one drawback is that the acoustic field is fixed. Each kinoform is designed for a single driving frequency and target distribution (which can be multi-planar). To overcome this limitation, multi-frequency kinoforms for which different target

distributions are encoded onto different frequencies could be used. The field at a particular depth could then be varied simply by altering the driving frequency of the transducer. This concept is illustrated by Fig. 7.1.

This Chapter investigates the design and properties of these multi-frequency kinoforms. In the next section, a modified direct search approach developed for the design of these kinoforms is introduced. This is then validated in Sec. 7.3 using two test cases, each designed to generate 3 target distributions at 3 distinct driving frequencies. Section 7.4 then reports a numerical investigation into the effect of different design parameters on the performance of these kinoforms. This is followed in Sec. 7.5 by the introduction of two extensions to the design algorithm improving its accuracy and flexibility. Finally, in Sec. 7.6 the conclusions of the work are presented along with future research directions.

7.2 Design algorithm

7.2.1 Background

Within optical holography, multi-frequency kinoforms are of interest for a variety of applications, particularly colour displays. This has led to the development of a variety of algorithms for their design. Two simple approaches are spatial and depth division [185]. In the former, the aperture of the kinoform is sub-divided to separate regions for each frequency [186]. For the latter, a kinoform is calculated for a single frequency with target images at different depths using a multi-plane iterative Fourier transform algorithm (IFTA). The focal distance of kinoforms in the nearfield scales inversely with wavelength, so different patterns are shifted into the target plane as the frequency is changed [187]. These algorithms are fast, however, their performance is limited compared to other approaches. For spatial multiplexing this is because segmenting the hologram limits the aperture that contributes to each pattern. For depth division this is because the phase offset generated by the

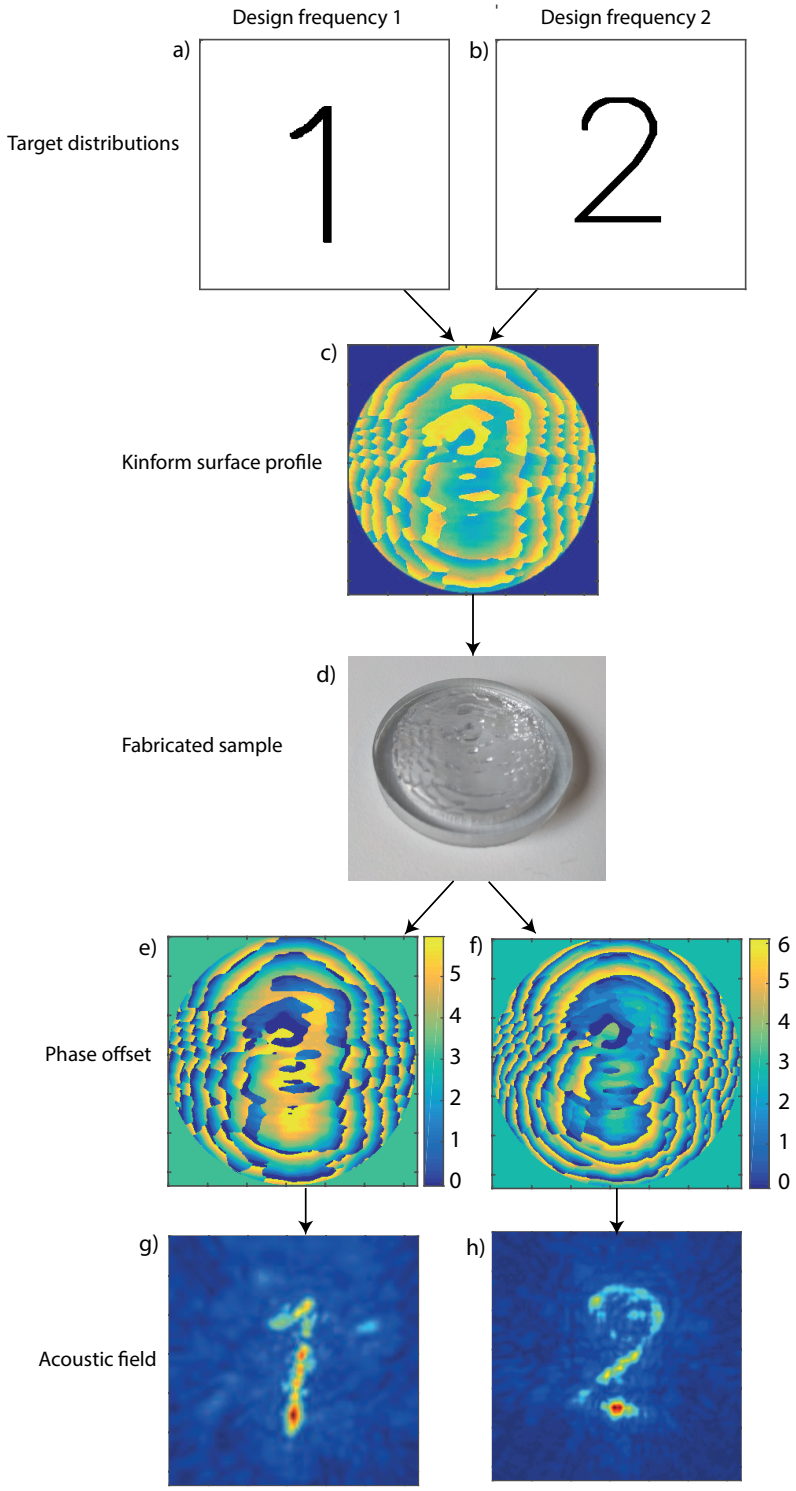

Figure 7.1: Illustrative figure demonstrating the concept of multi-frequency kinoforms. (a-b) Target distributions are defined for each driving frequency. (c) A direct search approach is used to calculate the kinoform. (d) Kinoform is fabricated via 3-D printing. (e-f) Kinoform is attached to a transducer where it generates different phase offsets for each frequency. (g-h) Output field is characterised and compared to the target distributions. Figure adapted from [183] © IEEE 2017.

kinoform is only correct (2π) for one of the design frequencies (Eq. 7.1), decreasing the efficiency. This results in variable efficiency for the different wavelengths, and a limited frequency spacing for which the approximation is accurate.

Other works have instead used methods that search directly for a thickness at each position that maximises a quality-metric based on the target field for each frequency. Several approaches have been used to solve this optimisation problem including direct search [188, 189], optimal-rotation-angle (ORA) [190, 191], and the Yang-Gu algorithm for phase retrieval [192]. To design the kinoforms in this Chapter, the direct search approach was again adopted. The implementation was adapted directly from Chapter 6 with alterations to the forward model and cost function that are outlined below.

7.2.2 Direct search for multi-frequency kinoforms

The inverse problem when designing a multi-frequency kinoform is to find a kinoform surface profile $h(x,y)$ that for a set of design frequencies f_n, $n \in \mathbb{N}$, maps the field of the transducer onto a set of frequency-specific target pressure distributions $p_n(x,y,z_d)$. Here z_d is the depth of the target plane and n is the index of the driving frequency. As in previous Chapters, rather than define continuous target distributions $p_n(x,y,z_d)$, sets of target points $p_n = \{(x_k,y_k,z_d), k = 1,\ldots,K\}_n$ were instead defined for each frequency.

For a kinoform with a spatially varying thickness described by $h(x,y)$ attached directly to a planar transducer driven at a frequency f, the acoustic field in a plane $p(x,y)$ directly after propagation through the element can be approximated by $p(x,y) = A(x,y)e^{i\Delta\phi(x,y)}$. Here $A(x,y)$ is the surface pressure of the transducer and $\Delta\phi(x,y)$ is the phase offset introduced by propagation through the kinoform. This phase offset is related to the thickness by

$$\Delta\phi(x,y) = 2\pi f \left(\frac{1}{c_{med}} - \frac{1}{c_{kino}}\right)(h_m - h(x,y)), \tag{7.1}$$

where c_{kino} and c_{med} are the sound speed in the kinoform and coupled medium, respectively, and h_m is the maximum thickness of the kinoform. This model approx-

imates the kinoform as an infinitely thin phase screen and neglects transmission losses.

Assuming the transducer is placed at $z = 0$, the acoustic field $p(x,y,z)$ for $z > 0$ can then be calculated using the Rayleigh-Sommerfeld integral [47]

$$p(x,y,z) = \int_{\Omega} A(x',y') e^{i\Delta\phi(x',y')} K(R) d\Omega,$$
$$K(R) = \frac{1}{2\pi}\left(-\frac{ik_c}{R} + \frac{1}{R^2}\right) e^{ik_c R}. \tag{7.2}$$

Here Ω defines the area over which the transducer output is non-zero, $d\Omega$ is a planar area element (in the $z = 0$ plane), and R is the distance between the coordinates (x,y,z) and $(x',y',0)$.

To solve the inverse problem, first the surface coordinates of the kinoform are discretised as in Chapter 6. Briefly, a square transducer occupying an L×L region of the (x,y) plane is represented by a set of N×N point sources with a spacing $\Delta = \frac{L}{N}$. The surface heights are also discretised such that the height at each coordinate (x_i, y_i) is given by $h_{ij} = m\Delta_z$, $m \in \mathbb{N}$ and $m\Delta_z < h_m$.

The optimisation searched for a map h_{ij} that maximised the pressure across the different sets of target points by minimising the following cost function

$$C = -\mathrm{E}(|p_n(f_n)|) + \alpha\,(\mathrm{std}(|p_n(f_n)|)) + \beta S. \tag{7.3}$$

Here $p_n(f_n)$ is the value of the pressure at a set of target points p_n, evaluated for a driving frequency f_n. The first term $E(|p_n(f_n)|)$ is the average pressure generated over the target points for each design frequency when the transducer is driven at that design frequency. The second term is the standard deviation of the same parameter. The three terms in Eq. 7.3 each serve the same purpose as those in Eq. 6.3. Namely, maximising the average amplitude across the target points, minimising the variations, and penalising sharp variations in height. The thin-phase element approximation neglects interaction with the kinoform structure so the smoothness term was again included.

The kinoform was initialised with a uniform height $h^0_{ij} = h_m/2$. Iterative re-

finements to the surface to form h_{ij}^{m} were then obtained using the approach described in Sec. 6.2.4, with one modification. For a continuous wave source, each height on the kinoform contributes a signal to each of the target points. Therefore, at each position the full range of heights were tested for each update.

If the smoothness term S was weighted inappropriately then the optimisation converged to kinoforms that were either flat or had large height differences between adjacent pixels. To find a suitable weighting, the optimisation was run once to convergence with $\beta = 0$. The output was then used to initialise a second instance of the optimisation with β scaled such that $E(|p_n(f_n)|) = 0.5S$ for the initial kinoform state. The value of α was assigned to 0.3.

7.3 Experimental validation

7.3.1 Fabrication

Two test cases were fabricated to demonstrate the use of multi-frequency kinoforms in practice. The first was designed to generate the numerals '1', '2' and '3' at three distinct frequencies. The second was designed to generate 3 sets of 8 discrete foci distributed evenly on circles of radii 10.5, 7, and 3.5 mm. These were selected to demonstrate that both patterns of pressure and discrete foci can be created with this approach.

Both kinoforms were designed for a 3.1 cm PZT piston transducer (Olympus, Japan) with a centre frequency of 2.25 MHz and a -6 dB bandwidth between 1.9 and 3.1 MHz. The image depth z_d in all cases was 2.5 cm and both kinoforms were discretised using a spacing Δ and Δ_z of 100 μm. The kinoform sound speed was assumed to be to 2495 $\mathrm{m\,s^{-1}}$ [180]. The maximum thickness h_m was set to 2 mm to provide a 2π modulation for the lowest design frequency. This was calculated by inverting Eq. 7.1 and using 1500 $\mathrm{m\,s^{-1}}$ as the sound speed for the coupled medium. The design frequencies for the three patterns were set to 1.9, 2.5 and 3.1 MHz to maximise the frequency spacing within the transducer bandwidth. The effect of the choice of both h_m and the frequency spacing is discussed in Sec. 7.4.

The calculated kinoforms were printed using a high-resolution PolyJet printer

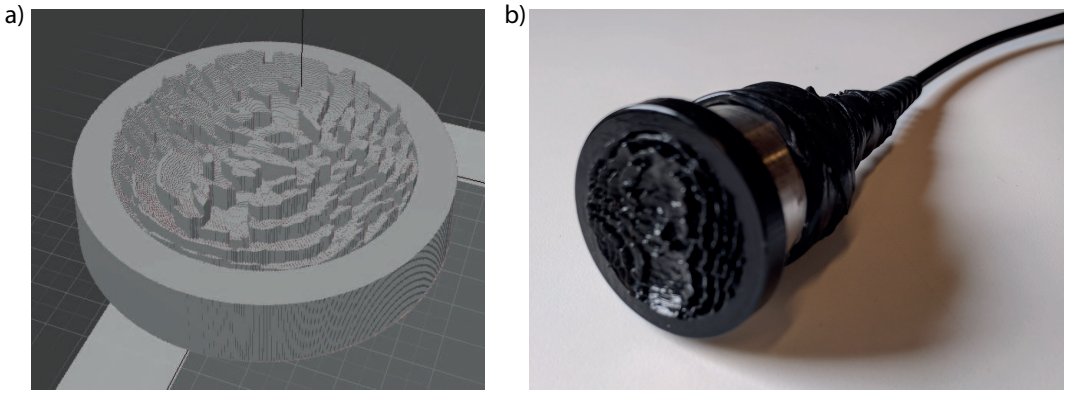

Figure 7.2: (a) 3-D rendering of the design of the three-pattern kinoform. (b) Photograph of fabricated kinoform attached to front face of piston transducer. Figure adapted from [184] under CC BY 4.0 License.

(Objet350 Connex, Stratasys, Eden Prairie, MN, USA) using VeroBlack as the substrate. A 1 mm thick backing was added to both kinoform designs for structural rigidity. Additionally, a ring was added to the outside to enable the kinoform to slot onto the front face of the transducer. A rendering of the kinoform design, and the fabricated kinoform for the 3 pattern sample can be seen in Fig. 7.2.

7.3.2 Measurement

The acoustic fields generated by the kinoforms were measured in a $40\times40\times60$ cm test tank with a two axis computer controlled positioning system (Precision Acoustics, Dorchester, UK). Measurements were made using a calibrated 0.2 mm needle hydrophone (Precision Acoustics, Dorchester, UK). The transducer was suspended inside the tank and aligned such that its front face was parallel to the measurement plane. The needle hydrophone was then suspended $\sim$5 cm from the transducer aligned with its beam axis. A photo of the experimental set-up can be seen in Fig. 7.3.

The transducer was driven using a signal generator (33522A, Agilent Technologies, Santa Clara, CA, USA) connected via a 75 W power amplifier (A075, E&I, Rochester, NY, USA). The driving signal in each case was a sinusoid at one of the design frequencies. The number of cycles was varied between 35 and 56 to ensure the signal covered the minimum and maximum travel times from the kinoform to the measurement plane. The driving voltage varied from 40-60 V peak-to-peak

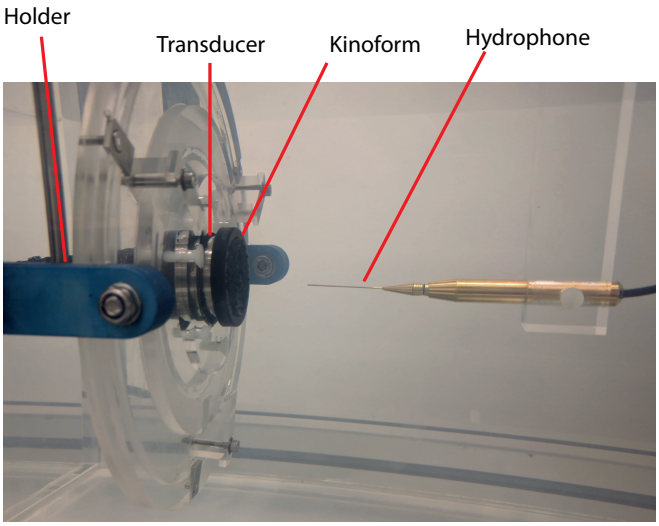

Figure 7.3: Photo of experimental set-up used to characterise the acoustic fields generated from the multi-frequency kinoforms.

for different scans. The field was recorded over a 6×6 cm area centred on the middle of the kinoform. This was done for both kinoforms and for each design frequency. The step size was varied from 0.24-0.3 mm to ensure a $\lambda/2$ sampling criteria was met for each driving frequency. Time domain signals were recorded at each position using a digital oscilloscope with a sampling rate of 400 MHz and 32 averages. To calculate the 3-D wavefield from each measurement set, the k-Wave toolbox was used to back-propagate them. Prior to projection, each dataset was upsampled by a factor of two and bandpass filtered between 100 kHz and 4 MHz to reduce noise.

The output field of each kinoform for each driving frequency was also simulated to compare with the experimental measurements. These simulations were carried out using the k-Wave toolbox on a $360\times360\times360$ grid with a 100 μm point spacing. The kinoform designs were inserted into the simulation as volumes with a sound speed of 2495 $\mathrm{m\,s^{-1}}$ and a density of 1190 $\mathrm{kg\,m^{-3}}$ [180]. The transducer was modelled by a circular source mask of diameter 3.1 cm driven at a single design frequency f_n. This was embedded in the kinoform volume. In the background medium the acoustic properties were set to those of water.

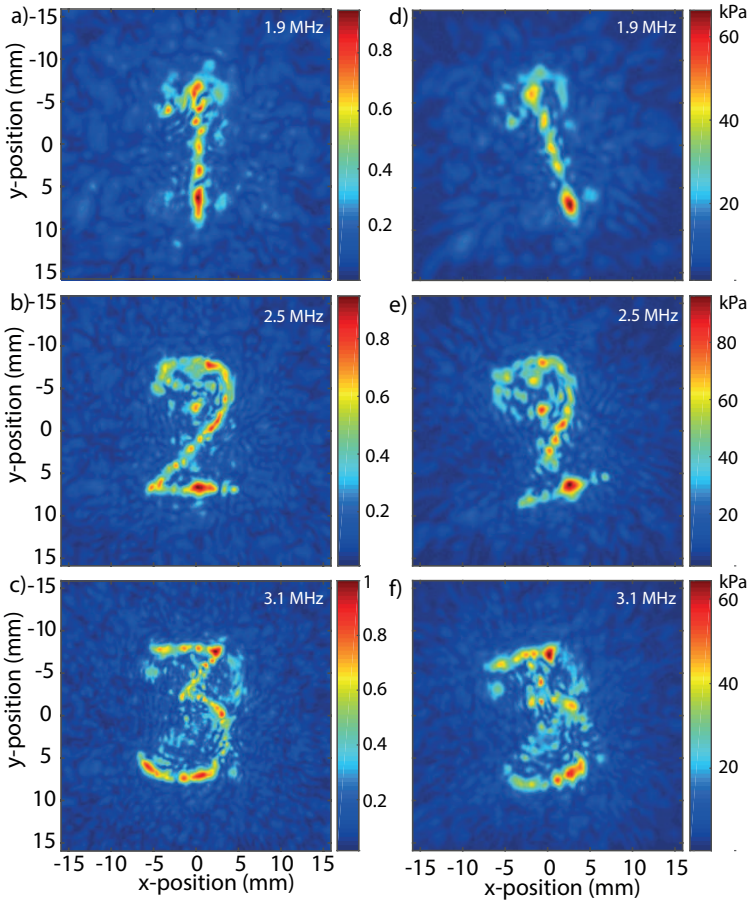

Figure 7.4: (a-c) Maximum pressure generated in the target plane of the continuous pattern kinoform in simulation for driving frequencies of (a) 1.9 MHz (b) 2.5 MHz and (c) 3.1 MHz. (d-f) Maximum pressure experimentally generated in the target plane of the continuous pattern kinoform for driving frequencies of (d) 1.9 MHz (e) 2.5 MHz and (f) 3.1 MHz. The rotation is due to the alignment of the kinoform on the transducer. Figure reprinted from [184] under CC BY 4.0 License.

7.3.3 Results

The maximum pressure generated in the target plane z_d for each design frequency in both simulation and experiment can be seen in Fig. 7.4 (a-f) for the continuous pattern kinoform and in Fig. 7.5 (a-f) for the discrete foci kinoform. For both kinoforms, the target distributions are clearly realised in the experimental data. This demonstrates that using multi-frequency kinoforms, different complex pressure distributions can be generated by switching the driving frequency. Additionally, the variation in pressure across the pattern and in the background is extremely similar between the simulations and experiments. This confirms that the k-Wave toolbox

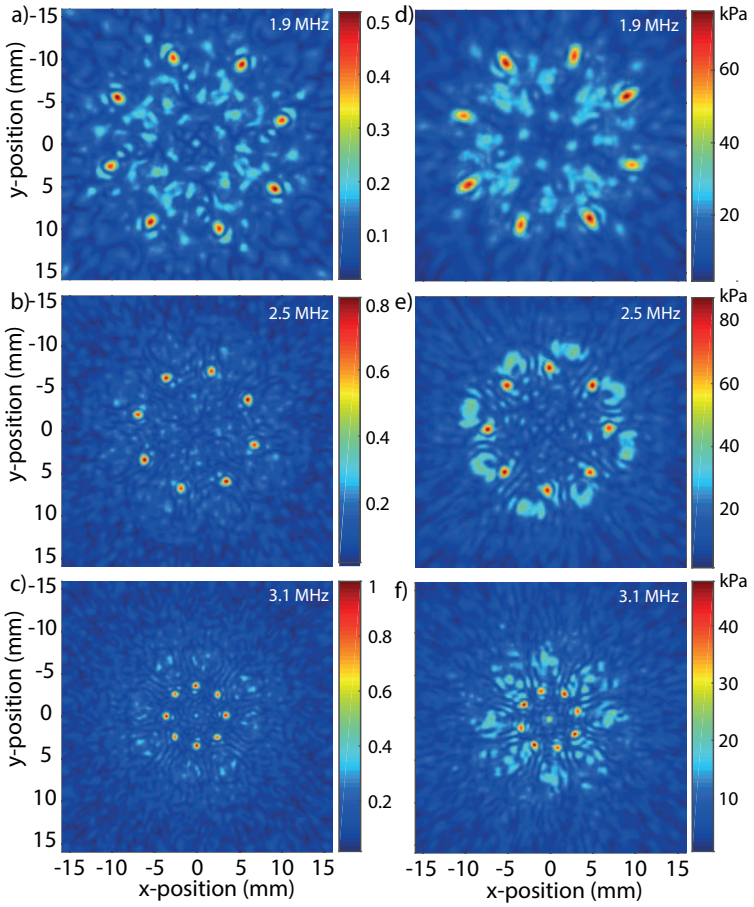

Figure 7.5: (a-c) Maximum pressure generated in the target plane of the discrete foci kinoform in simulation for driving frequencies of (a) 1.9 MHz (b) 2.5 MHz and (c) 3.1 MHz. (d-f) Maximum pressure experimentally generated in the target plane of the discrete foci kinoform for driving frequencies of (d) 1.9 MHz (e) 2.5 MHz and (f) 3.1 MHz. Figure reprinted from [184] under CC BY 4.0 License.

can be used to accurately predict the fields prior to fabrication.

Despite the clear generation of each targeted distribution, for both samples there are notable variations in the pressure across each pattern and the surrounding background pressure is high. This is in part due to constraints on the feature size imposed by the diffraction limit for each design frequency. However, it is principally caused by 'cross-talk' between the target distributions encoded onto the different frequencies.

7.3.4 Cross-talk

Cross-talk refers to the generation in the field at one design frequency of distributions encoded onto other design frequencies. It can be seen from Eq. 7.1 that the

phase offset generated for each frequency is identical aside from a scaling factor. For low maximum modulation depths (on the order of 2π) and narrowly spaced frequencies, this means the phase offsets generated for different design frequencies vary only slightly. As a result, each frequency has a similar diffraction pattern and each distribution encoded on the kinoform is generated in the field of each design frequency.

Fortunately, near-field kinoforms have a chromatic variation in focal depth that occurs due to the wavelength dependence of the phase offsets generated by free-space propagation [193]. This was observed in the frequency variation in zone plate focal depth explored in Sec. 3.7. This means that while all encoded patterns will appear in the field simultaneously, the patterns for other design frequencies will be out-of-plane. For a target depth z_d, the depth z_c at which each pattern is realised can be approximated from the ratio of the design and driving frequencies using

$$z_c = z_d \frac{f_{dr}}{f_{de}}, \tag{7.4}$$

where f_{de} is the design frequency for the pattern and f_{dr} is the frequency driving the transducer. This is illustrated by Fig. 7.6 and is integral to the operation of the multi-frequency kinoforms in this Chapter.

The amplitude of the crosstalk can in principle be reduced by observing that the diffracted field is determined by the 2π modulus of the phase rather than the absolute value. Therefore, increasing the thickness and using a kinoform with a high maximum phase modulation depth $> 4\pi$ can allow more distinct phase patterns to be generated for each frequency, as a given value of ϕ_1 will correspond to multiple values of ϕ_2 [190]. However, the ability to increase the modulation depth is limited by another factor: the accuracy of the thin phase element approximation used in the forward model.

7.3.5 Forward model accuracy

In optics it has been observed that the scalar or thin phase element approximation for diffractive elements breaks down as the scale of features (or spatial frequen-

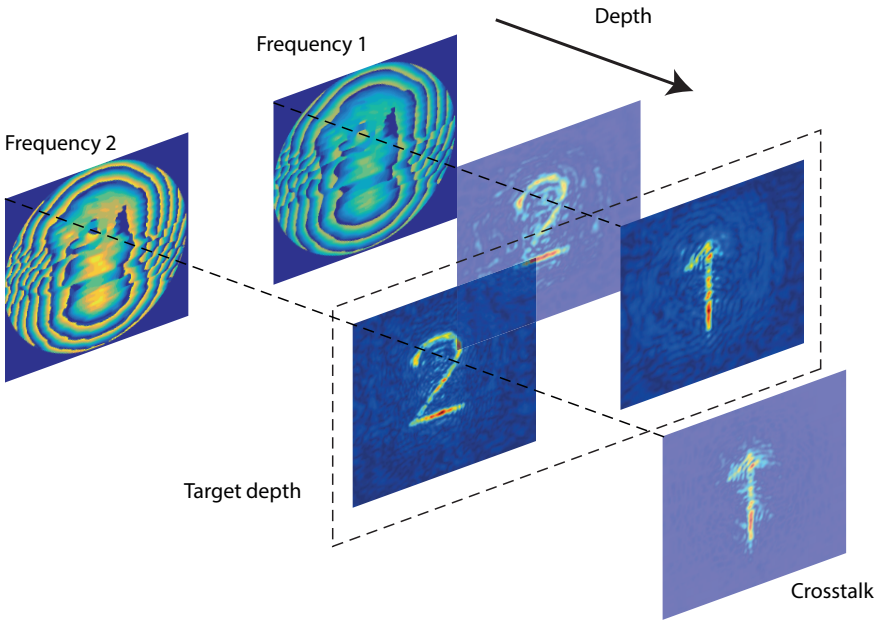

Figure 7.6: Illustrative figure of the cross-talk generated by the multi-frequency kinoforms.

cies) across the kinoform approach the size of the wavelength, and, as the thickness increases [194, 195]. In practice, this means that if h_m is increased arbitrarily, the forward model will break down and the optimisation will converge to structures that don't generate the target fields.

To illustrate this effect, a phase profile $\phi(x,y)$ was calculated, designed to focus at a single point for a single design frequency. The phase profile was transformed into two kinoforms $h_1(x,y)$ and $h_2(x,y)$ with differing thickness by inverting Eq. 7.1 to give

$$h_m - h(x,y) = \frac{\Delta\phi(x,y)}{2\pi f(\frac{1}{c_{med}} - \frac{1}{c_{kino}})}, \tag{7.5}$$

and using values of 2495 m s^{-1} and 1700 m s^{-1} for c_{kino}. The maximum thickness h_m was calculated for each using

$$h_m = \frac{1}{f\left(\frac{1}{c_{med}} - \frac{1}{c_{kino}}\right)}. \tag{7.6}$$

The output of both kinoforms was then simulated, with the appropriate sound speed, using k-Wave. The maximum pressure in a slice of the field containing the focus can be seen for both kinoforms in Fig. 7.7 (b-c). Just by altering the thickness,

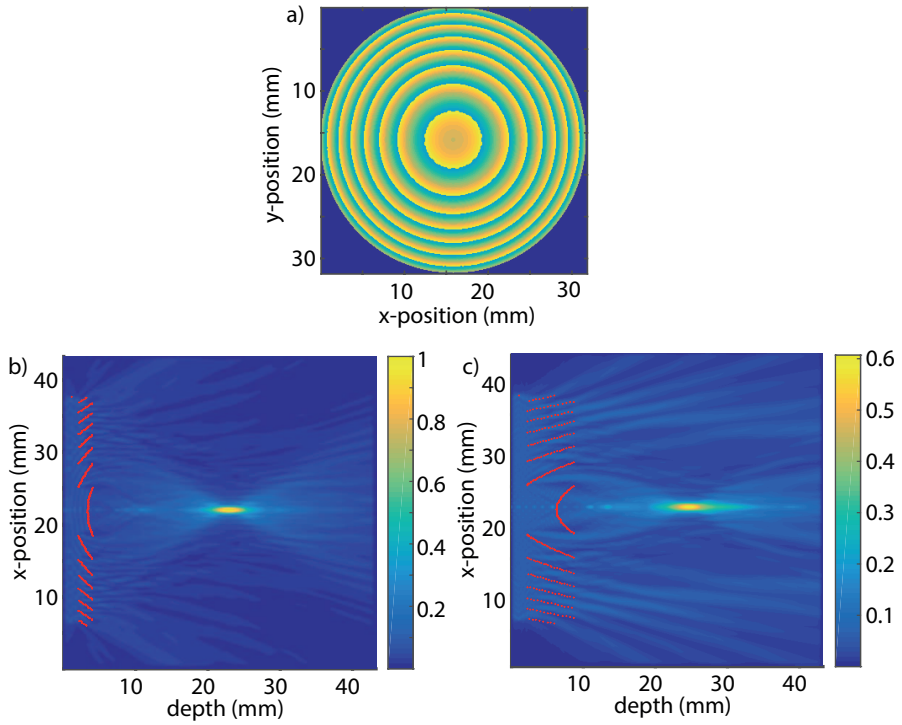

Figure 7.7: (a) Example single focus kinoform used to illustrate the breakdown of the thin phase approximation. (b) Maximium pressure in slice of acoustic field containing the focus for the kinoform calculated for a sound speed of 2495 m s^{-1}. (c) Maximium pressure in slice of acoustic field containing the focus for the kinoform calculated for a sound speed of 1700 m s^{-1}. In both cases the dotted line indicates the surface height of the kinoform at each point.

the output field has been changed substantially. The pressure at the focus of the thick kinoform has dropped to 60% of the thin kinoform and the background is significantly higher.

7.4 Optimisation parameters

From Secs. 7.3.4 and 7.3.5 it is clear that the performance of a multi-frequency kinoforms for a given problem will be determined by several different design parameters. Most significant are: the maximum thickness h_m, the ratio of the design frequencies f_n, and the target depth z_d. The value of these parameters will determine the trade-off between cross-talk amplitude and the forward model accuracy. The impact of these parameters was investigated numerically.

7.4.1 Numerical experiment design

All numerical experiments were carried out for a kinoform with a diameter of 3.1 cm discretised with a spacing Δ of 100 μm. Two binary patterns were encoded in each case: the numerals '1' and '2' seen in Fig. 7.1 (a-b). The sound speed and density, unless otherwise stated, were set to 2495 $\mathrm{m\,s^{-1}}$ and 1190 $\mathrm{kg\,m^{-3}}$.

Two approaches were used to calculate the field from each kinoform. For the first method, the kinoforms were approximated as infinitely thin phase elements that imparted phase offsets given by Eq. 7.1. A uniform incident pressure was assumed and the ASA was used to propagate the field to the target depth. For the second method, propagation through each kinoform was simulated using the k-Wave toolbox. This was done to assess the severity of the inaccuracies of the thin phase approximation. The simulations were carried out on a $360\times360\times128$ grid with a spacing of 100 μm. The kinoform volumes h_{ij} were inserted into the simulation as a region with a sound speed and density of 2495 $\mathrm{m\,s^{-1}}$ and 1190 $\mathrm{kg\,m^{-3}}$. The transducer was modelled by a circular source mask of diameter 3.1 cm driven at a design frequency f_n placed within a 1 mm buffer that was added to each kinoform structure. This buffer was extended into the perfectly matched layer eliminating reflections off the back from the model. The sound speed and density in the rest of the medium were set to 1500 $\mathrm{m\,s^{-1}}$ and 1000 $\mathrm{kg\,m^{-3}}$. In each simulation, the time varying pressure $p(x,y,t)$ was recorded across an (x,y) plane 0.2 mm above the kinoform. The phase $\theta(x,y)$ and amplitude $A(x,y)$ at the driving frequency f_n in steady state were extracted from these signals. The ASA was then used to propagate the field to other depths in steps of 100 μm to calculate the field across a volume.

Two key metrics were of interest. The SNR achieved in the target plane for each encoded pattern and the crosstalk between the patterns. To calculate the metrics, the images of '1' and '2' were used to define a foreground and a background for both patterns. The SNR was defined as the ratio of the average pressure across the foreground and the background. This was evaluated at the target depth for both models and each design frequency. The crosstalk was defined using the same

method, except the SNR of the '2' was evaluated in the field of the '1' and vice-versa. This was measured for only the full-wave model. As the crosstalk image depth varies from the target depth (Eq. 7.4), the metric was evaluated for all depths and the plane with the highest value was selected for each frequency. For each test, the SNR and cross-talk were averaged across both patterns to give a single value.

7.4.2 Maximum thickness and modulation depth

As discussed in Sec. 7.3.4, the maximum thickness h_m or the maximum phase modulation depth $\Delta\phi$ has a significant impact on the resulting field. If the different frequencies can be considered to have ideal phase distributions (ϕ_1,...,ϕ_n) then these can be approximated with much lower error for high modulation depths [196]. To test the effect of h_m, a set of kinoforms were calculated in which it was varied between 1-6 mm. In each case, the two design frequencies were set to 2 and 3 MHz and the image depth was set to 2.5 cm.

Figure 7.8 (a) shows the variation in SNR and cross-talk as a function of maximum thickness. As expected due to the reduction in phase error, the SNR for the ASA model improves with increasing h_m and the full-wave crosstalk decreases. Both converge for $h_m > 4$ mm or a 4π modulation for 2 MHz, suggesting the optimisation can find few further improvements. However, the full-wave model SNR behaves very differently. For h_m between 1 and $\sim$1.8 mm it increases. This is followed by a slow drop between 1.8-4 mm, then a rapid fall-off between 4-6 mm. This occurs because, as discussed in Sec. 7.3.5, the thin element approximation becomes increasingly inaccurate with increasing h_m, as not just the thickness but also the spatial frequencies on the kinoform profile increase. After 4 mm, extremely rapid jumps appear in the kinoform height causing the optimisation forward model to completely breakdown. This can be seen in Fig. 7.8 (c).

Next, the value of h_m was fixed at 2.5 mm and a set of kinoforms were calculated scaling the modulation depth from π to 4.4π by varying the kinoform sound speed c_{kino}. This was done to separate errors introduced by the increasing thickness from errors introduced by the changing structure (i.e., the spatial frequencies). For each kinoform, the two design frequencies were again set to 2 and 3 MHz and z_d to

2.5 cm.

Figure 7.8 (d) shows the measured variation in SNR and cross-talk as a function of modulation depth for fixed thickness. In this case, the peak SNR now occurs for a modulation depth of $\sim 3.3\pi$ (as opposed to $\sim 1.8\pi$). This is because, with the thickness fixed, the improvements in solution quality resulting from the increased modulation depth compensate for the decrease in model accuracy. This can be seen from Fig. 7.8 (e) which shows the ratio of ASA SNR to full-wave SNR as a function of modulation depth, for both experiments. In both cases, the accuracy (mostly) decreases with increasing modulation depth, however, when the thickness is fixed the decrease is slower allowing the improvement in solution quality to compensate. These experiments demonstrate that it is necessary to constrain h_m to maximise the performance, although the optimal value will depend on both the medium and kinoform properties.

7.4.3 Frequency spacing

The spacing of the design frequencies is also an important parameter. In practice, it is constrained by the transducer bandwidth, however, the choice of frequencies within that range can have a significant impact on performance. To investigate, kinoforms were calculated with the design frequency for the '1' fixed at 1.5 MHz, and the design frequency for the '2' increased in increments of 100 kHz starting at 1.6 MHz. This was done for h_m set to 2.5 or 5 mm, corresponding to modulation depths of 2π and 4π at 1.5 MHz.

Figure 7.9 (a) shows the variation in SNR and crosstalk for a maximum thickness of 2.5 mm. This shows that as the ratio of frequencies increases, the overall SNR increases and the crosstalk decreases for both the full-wave and ASA model. This is because the phase modulation depth increases for the higher frequency (due to its shorter wavelength) reducing the average phase error. The higher frequency also has a smaller point spread function (PSF) improving the fidelity of the target pattern. The decrease in the crosstalk is again due to less similarity between the phase distributions as wavelength difference and modulation depth increase.

Figure 7.9 (b) shows the variation for a maximum thickness of 5 mm, which is

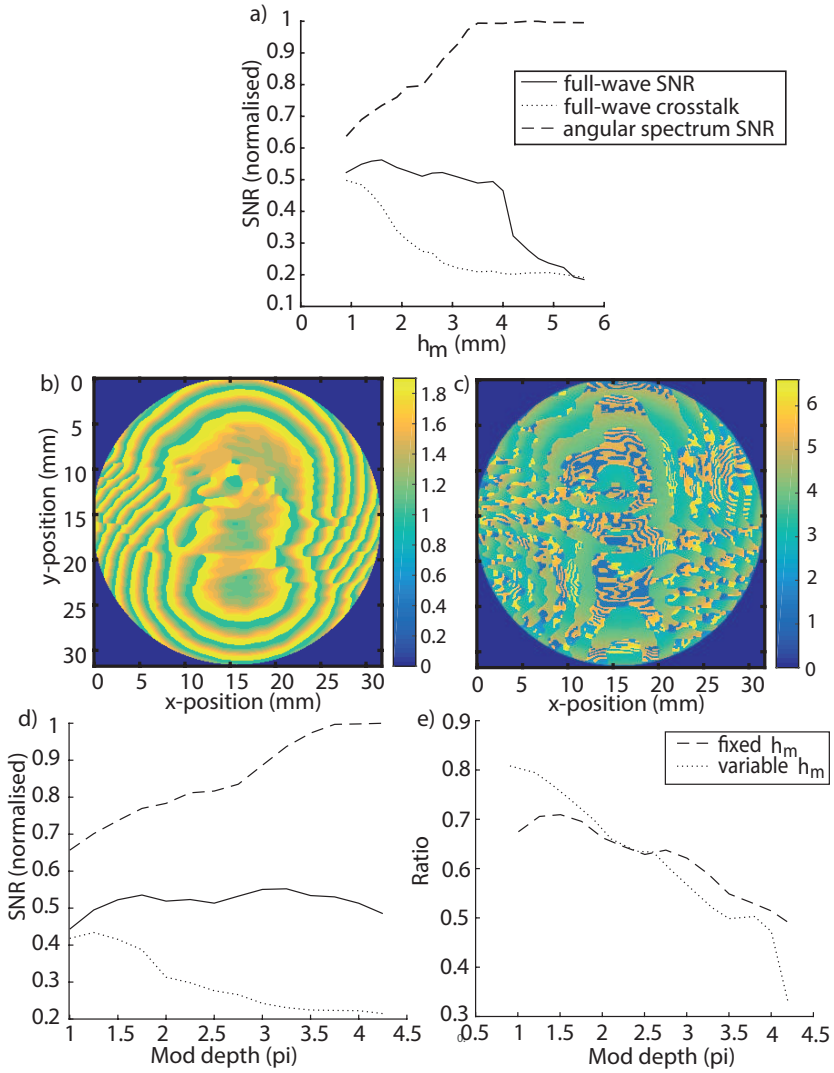

Figure 7.8: (a) Variation in SNR and cross-talk for the two pattern kinoform as a function of maximum thickness h_m. Normalised to the peak value of SNR. (b) Kinoform calculated for h_m of 1 mm. (c) Kinoform calculated for h_m of 6 mm. The rapid jumps in height as the modulation depth increases are apparent. (d) Variation in SNR and cross-talk for the two pattern kinoform, fixing the maximum thickness, altering the modulation depth by changing the sound speed. Normalised to peak value of SNR. (e) Variation in ratio of ASA SNR to full-wave SNR for experiments in (a) and (d). Figure adapted from [183] © IEEE 2017.

notably different. Initially, as the ratio of frequencies increases from 1 to ∼1.6, both the ASA and full-wave SNR increase. However, as the higher frequency approaches the second harmonic, both start to decrease. This decrease is due to how the phase offsets for each frequency are distributed relative to each other and can be explained by considering two cases: $f_2 = 2f_1$ and $f_2 = 1.5f_1$. For $f_2 = 2f_1$, for a given phase value ϕ_2 there are only 2 possible phase values for ϕ_1, despite the 4π modulation

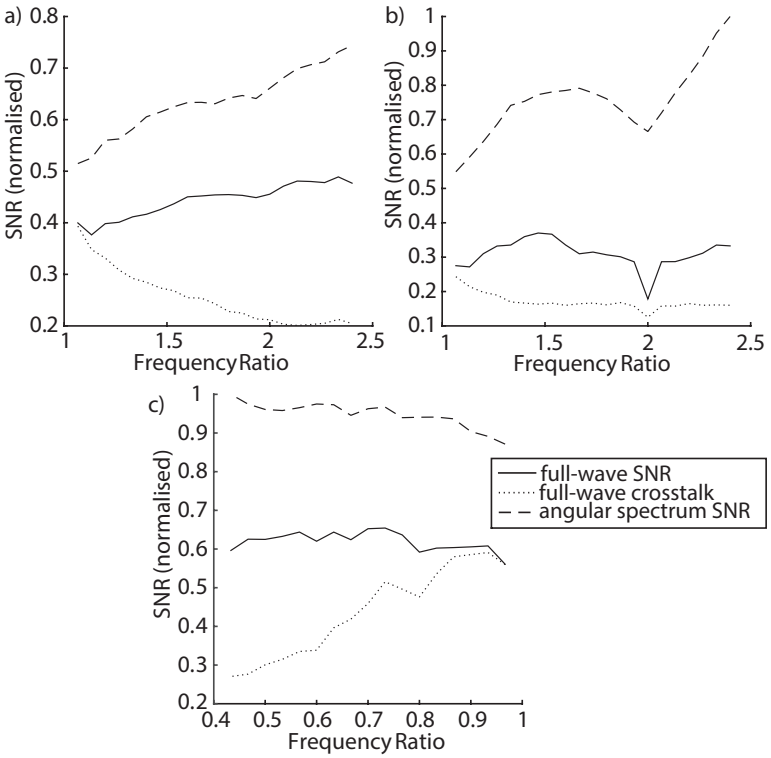

Figure 7.9: Variation in average SNR and cross-talk for the two-pattern kinoform as a function of the frequency ratio. (a) Variation for h_m = 2.5 cm for a constrained lower frequency, normalised to peak across (a) and (b). (b) Variation for h_m = 5 cm for a constrained lower frequency, normalised to peak across (a) and (b). (c) Variation for constrained upper frequency normalised to peak in (c). Figure reprinted from [183] © IEEE 2017.

depth. For $f_2 = 1.5f_1$, there are 3 possible ϕ_1 values for each ϕ_2. Thus despite the lower modulation depth, the latter will have a lower average phase error [197].

This hypothesis, that the error increases due to an increase in the phase error with frequency ratio, was confirmed using a simple simulation. A set of 10^6 random phase pairs $(\phi_n(1), \phi_n(2), n = 1, ..., 10^6)$ were generated. The thickness $h = n\Delta_z$ which minimised the combined error in approximating each pair was then calculated, fixing the lower design frequency at 1.5 MHz and varying the upper frequency between 1.6-4.5 MHz. The average phase error as a function of frequency ratio, with h_m set to 2.5 and 5 mm, can be seen in Fig. 7.10. This confirms that, for a 4π modulation depth, the error is smaller for a ratio of $\sim$1.5 than for a ratio of 2. This demonstrates that for a low modulation depth and fixed lower frequency, the spacing should be maximised, while for higher modulation depths harmonics

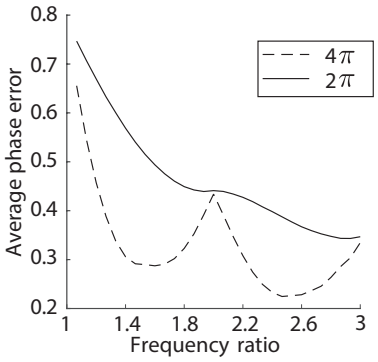

Figure 7.10: Average minimum phase error for randomly distributed phase pairs for the two-pattern kinoform as a function of frequency ratio.

should be avoided.

Next, the case of fixed upper frequency was investigated. To test this, the frequency of the '2' was fixed at 3 MHz and the frequency of the '1' was decreased in increments of 100 kHz starting at 2.9 MHz. The maximum thickness h_m was varied for each kinoform so that a 2π modulation depth was provided for the '1' in each case. The results are shown in Fig. 7.9 (c). In this case, the SNR for the ASA model still increases with increasing spacing, however, this happens more slowly. The initial improvement is principally due to the removal of in-plane crosstalk. It can be seen from Eq. 7.4 that for a close spacing, the patterns overlap at the target depth. The subsequent improvement is again principally due to increases in the modulation depth of the higher frequency. These improvements are offset by increases in the PSF for the lower frequency which decrease the fidelity of the '1'. Slightly different results can be seen for the full-wave SNR, where a small initial increase is followed by stability. The SNR remains stable rather than increasing, as the accuracy of the ASA model decreases due to the greater element thickness.

7.4.4 Target depth

The variation in SNR with target depth z_d is also interesting. Intuitively, the solution quality would largely be expected to increase as the target depth is reduced due to the larger numerical aperture. However, the fringe spacing (or spatial frequency) also increases simultaneously, increasing the aberrations. To test the effect of image depth, kinoforms were calculated varying z_d between 2-5 cm in increments of 3 mm.

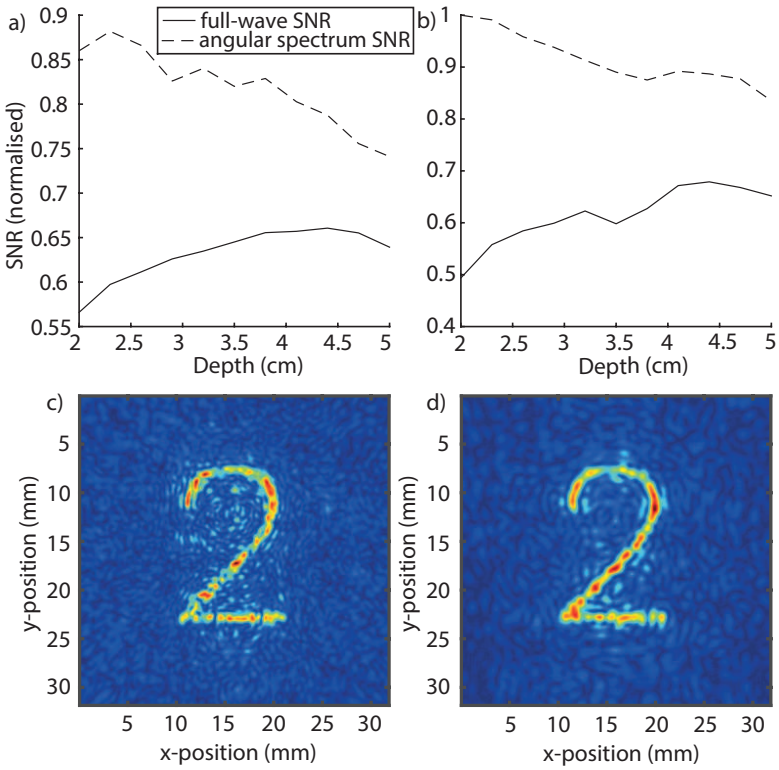

Figure 7.11: Variation in average SNR for the two pattern kinoform as a function of the target depth for (a) h_m of 2 cm (b) h_m of 3 cm. All lines normalised to peak across (a) and (b). (c) Peak pressure in the target plane for a depth of 2 cm. (d) Peak pressure in the target plane for a depth of 4.7 cm. Figure adapted from [183] © IEEE 2017.

This was done for h_m of 2 and 3 mm and fixed design frequencies of 2 and 3 MHz.

Figure 7.11 shows the resulting variation. As expected, in both cases the ASA model SNR decreases with increasing target depth, due to the increasingly broad PSF. However, the full-wave model SNR increases up to a target depth of 4 cm, as the reduction in aberrations compensates for the decrease in the fidelity of the target patterns. Similar trends are seen for both values of h_m tested. In this case, the overall peak in SNR occurs for a h_m of 3 mm suggesting that the optimal modulation depth will vary for different problems.

7.4.5 Additional parameters

Other design parameters also have an influence on the kinoform performance. As stated in Sec. 7.3.3, and encountered elsewhere, the target feature size is diffraction limited for each design frequency. Therefore, increasing the design frequencies and

aperture size (where possible) will improve the image fidelity. However, for too high a value the structures on the kinoform could become too fine to realise with 3-D printing, the spacing of phase levels will be greater, and acoustic attenuation will limit the deliverable pressures. The choice of sound speed (where possible) also influences the performance. As seen in Sec. 7.4.2, increasing the sound speed enables a higher modulation depth to be achieved in the same thickness, reducing the aberrations. However, the impedance mismatch between the coupled medium, the kinoform, and the front face of the transducer also increases, resulting in more reflected energy. Figure 7.12 illustrates the effect that changing the design frequencies, and their spacing, has on the acoustic field more clearly. It shows the field of the continuous pattern kinoform used in the experiment compared to that of a kinoform designed to generate the same patterns using design frequencies of 2.5, 4, and 6 MHz. The reduction in crosstalk and significant improvement in the fidelity are apparent.

7.5 Multi-frequency kinoform extensions

There are several possible extensions to the basic design algorithm introduced in Sec. 7.2 that can be used either to improve the flexibility of the generated kinoforms or to improve their performance. In the following sections, two such extensions are described.

7.5.1 Phase and amplitude control

The original work by Melde et al showed that by controlling the phase across the target distribution in addition to the amplitude, objects could be propelled along predefined paths. Replicating this across multiple frequencies would allow for different, pre-defined, mechanical operations on particles to be encoded onto the same kinoform, greatly improving the flexibility. As in Chapter 6, this combined phase and amplitude control can be achieved by making several modifications to the optimisation.

First, a desired phase $\theta_n = \{\phi_k, k = 1, \ldots, K\}_n$ is defined for each target point $p_n = \{(x_k, y_k, z_d), k = 1, \ldots, K\}_n$ at each frequency. Next, the initialisation is ad-

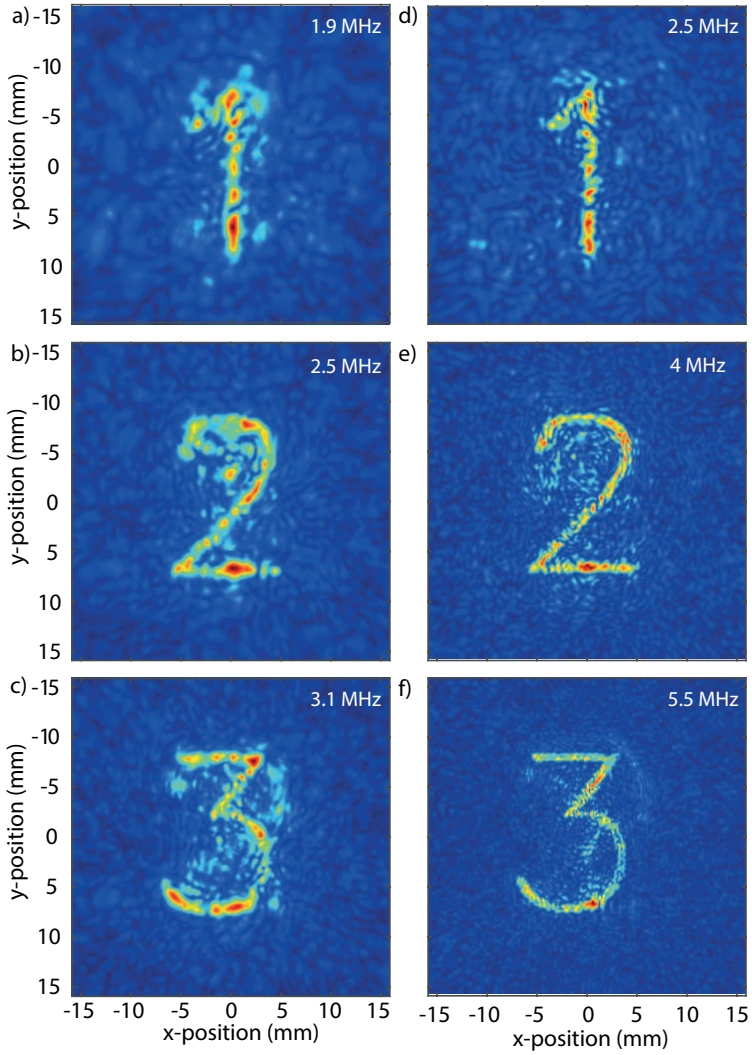

Figure 7.12: (a)-(c) Maximum pressure generated in the target plane of the continuous pattern kinoform in simulation for driving frequencies of (a) 1.9 MHz, (b) 2.5 MHz, and (c) 3.1 MHz. (d)-(f) Maximum pressure generated in the target plane of a kinoform designed to generate the same patterns as (a)-(c) but for driving frequencies of 2.5, 4, and 5.5 MHz. Figure reprinted from [184] under CC BY 4.0 License.

justed. Rather than start from a flat kinoform, at each point the height h_{ij} is assigned the value that minimises the L2 norm between the phase of the wave generated from the source point at each target point and the target phase θ_n. This is evaluated, for each possible height, using

$$D = \sum_n ((\Delta\phi_{ij} + \arg(e^{ikR_{ijn}})) - \theta_n)^2. \tag{7.7}$$

Here, R_{ijn} is the distance between $(x_i, y_j, 0)$ and each target point p_n, and $\Delta\phi_{ij}$ is

the phase offset introduced by a given height h_{ij} evaluated using Eq. 7.1. The cost function for the optimisation is then altered to

$$C = -\mathrm{E}(|p_n(f_n)|) + \alpha\,(\mathrm{std}(|p_n(f_n)|)) + \beta S + \gamma(\arg(p_n(f_n)) - \theta_n)^2, \qquad (7.8)$$

where γ is an additional weighting factor, and $\arg(p_n(f_n))$ is the current phase generated by the kinoform across each target point. This penalises variations in the kinoform pattern phase from the target phase.

To validate the model, a kinoform encoding two patterns was calculated. The target distributions for this kinoform were two 2 cm lines orientated perpendicular to each other, with design frequencies of 1.9 and 3.1 MHz. The phase for both lines was designed to vary linearly by 4π, while the pressure across both lines was designed to be identical. The kinoform diameter was 3.1 cm, the spacings Δ and Δ_z were both set to 100 μm, and the target depth was 2.5 cm. The weighting term γ was assigned a value, after initialising the surface, such that $\mathrm{E}(|p_n(f_n)|) = (\arg(p_n(f_n)) - \theta_n)^2$. This was to equally weight the initial accuracy of the phase term and the amplitude term. The calculated kinoform, along with both target distributions can be seen in Fig. 7.12 (a-c).

Figure 7.12 (d-e) shows the simulated peak pressure generated in the target plane for both design frequencies by the kinoform. In both cases, the lines have been clearly generated. The variation in phase across the target pattern for both frequencies is shown in Fig. 7.12 (f-g). The targeted phase distributions have been accurately replicated by the optimisation confirming that simultaneous control over the amplitude and phase can be achieved. However, the deviation from the target pressure (uniform) is greater than the deviation from the target phase. This suggests that the weighting of the phase term was too high. Additionally, the developed method does have drawbacks. Most significantly, the absolute phase is constrained while, in practice, only the relative phase across the target points is of interest. However, this only significantly effects the initialisation. Once initialised, variation of the entire phase profile is difficult due to the small perturbation to the pressure applied by changing a single pixel.

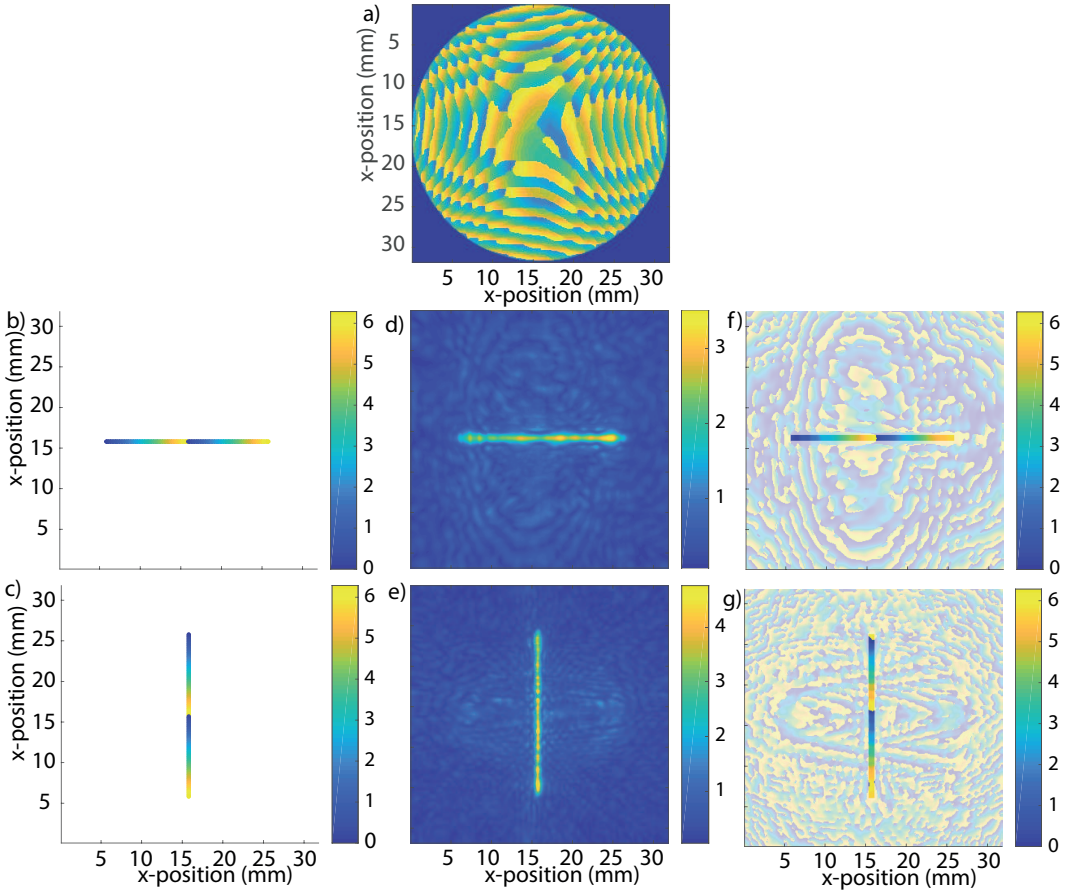

Figure 7.13: (a) Multi-frequency kinoform calculated for simultaneous control over phase and amplitude. (b-c) Target phase distributions for the two design frequencies of the kinoform in (a). (d-e) Maximum pressure generated at the target depth for both design frequencies. (f-g) Variation in phase across the target depth for both design frequencies.

7.5.2 Full-wave model optimisation

A major obstacle encountered in this Chapter are the limitations imposed by the inaccuracies of the thin-phase approximation. Ideally, this could be avoided by using k-Wave to calculate updates to the field. However, as in Chapter 6, this is impractical due to the number calls of the forward model required to calculate a single kinoform $\sim 10^7$. This arises due to the number of pixels on the kinoform surface, which is typically of the order 250×250. For certain problems, however, this pixel number can be substantially reduced enabling the forward model to be replaced. The most obvious example of such a problem is a set of axisymmetric target distributions for which the target pressure varies only radially. In this case, the kinoform pixels can be clustered into groups g_l defined, for a kinoform of radius

G, using $g_l = \{((l-1)\Delta < x_i^2 + y_j^2 <= l\Delta), l = 1, \ldots, \frac{G}{\Delta}\}$.

To establish whether the kinoform performance can be improved with a more accurate forward model, a test case was designed. The target patterns were two rings of radii 4 and 7 mm encoded onto the frequencies 2 and 3 MHz. The aperture for the kinoform was a circle of diameter 3.1 cm, the spacings Δ and Δ_z were set to 150 μm, and the target depth z_d was set to 2 cm. To reduce the computational effort further, the optimisation was first run to convergence using the thin-phase element approximation, clustering the pixels as described above.

The kinoform output by the thin phase model was used to initialise the full-wave step. First, the steady state pressure generated across both target patterns was evaluated for the input kinoform using the method described in Sec. 7.4.1. Iterative updates the surface were then sought. A group of pixels g_l on the surface was selected at random, the height h_l of this group was perturbed to both $h_l + \Delta$ and $h_l - \Delta$, and the field was evaluated, with k-Wave, for both updates. The position was then assigned the height from $(h_l - \Delta, h_l, h_l + \Delta)$ that maximised the cost function and further groups of pixels were updated, chosen in a random-exhaustive manner. After 10 iterations the optimisation was terminated. The cost function for this full-wave optimisation was altered from the usual state-based equation. The field across the entire target plane was available for no additional computational cost so the metric maximised was the SNR between the foreground and background pressure. This was evaluated as defined in Sec. 7.4.1.

The input and output kinoforms from this optimisation are shown in Fig. 7.14 (a-b). These show that the optimisation has made a number of perturbations to the kinoform structure. The peak pressure generated in the target plane by the input and output kinoform, for each design frequency, can be seen in Fig. 7.14 (c-f). The generation of the target patterns has been significantly improved in the field of the optimised kinoform, particularly for the lower design frequency. Overall the ratio between the foreground and background pressure was increased by 117% for the 2 MHz pattern and by 17% for the 3 MHz pattern. This suggests that using a full-wave model for axisymmetric problems could be valuable. However, further steps

are required to confirm this.

First, the accuracy of k-Wave for this problem needs to be validated experimentally. If fabrication errors are substantial or if the model errors are too large, then the perturbations made by the full-wave step may not improve the output field. Second, initialising using the output from the thin-phase approximation may also have introduced errors. The converged kinoform for the thin-phase model may lie close to a local minimum that is not optimal when employing a more accurate model. For the test case this step was included as the forward model is still expensive ($\sim$30s) to evaluate. Finally, it is unclear whether it will be possible to find substantial improvements to the SNR in all cases, or, whether the performance is inherently limited due to the physical structure. This needs to be established.

7.6 Conclusion

This Chapter has shown several complex, arbitrary distributions of pressure can be generated from a single element piezoelectric transducer using multi-frequency kinoforms. This represents a significant extension to the work by Melde et al, greatly improving the flexibility of the technique. A direct search approach has been developed for the design of these kinoforms, and the accuracy of this model has been experimentally validated using two 3-D printed test cases. These experiments have shown that both continuous distributions of pressure and arrangements of discrete foci can be created using these kinoforms.

A detailed numerical investigation has also been carried out quantifying the effect of different design parameters on the performance of the kinoforms. It was found that the frequency spacing should be maximised within the transducer bandwidth, the target depth should be determined from the diffraction limit of each design frequency and the structure of the target patterns, and the maximum thickness should be constrained due to limitations on the forward model accuracy. Finally, two extensions to the design algorithm have been introduced. The first showed that by altering the optimisation cost function and initialisation, simultaneous control over the phase and amplitude distribution across multiple frequencies can be

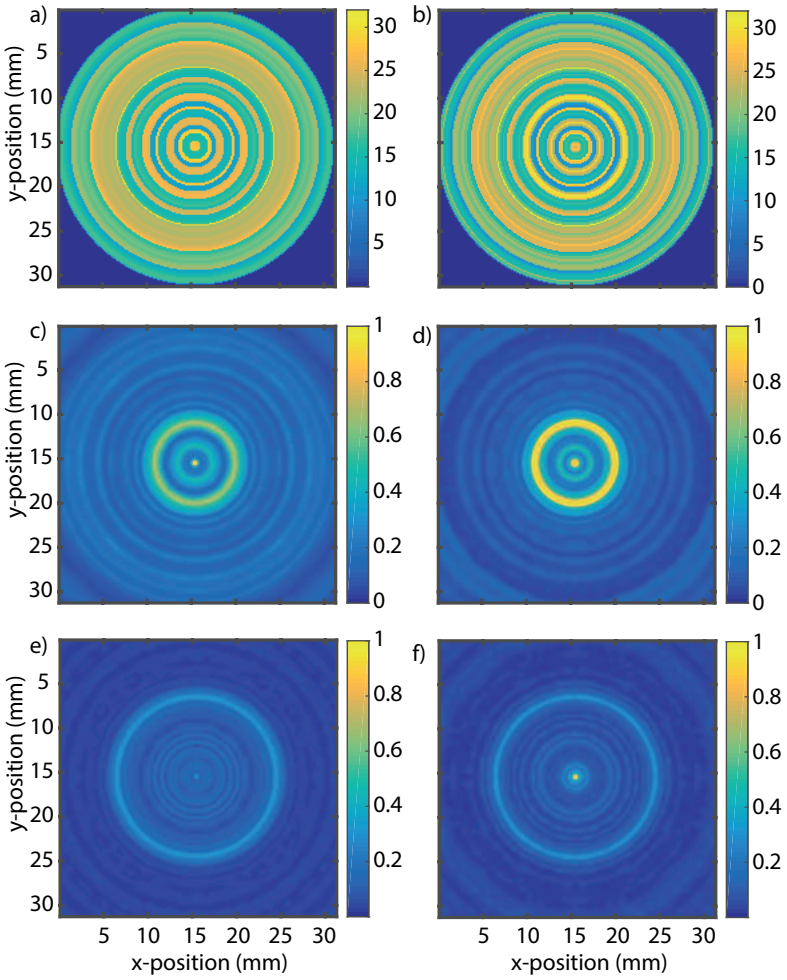

Figure 7.14: (a-b) Initial (a) and final (b) state for full-wave model optimised kinoform. (c-d) Peak pressure generated in the target plane by the initial (c) and final (d) state of the kinoform for a driving frequency of 2 MHz. (e-f) Peak pressure generated in the target plane by the initial (e) and final (f) state of the kinoform for a driving frequency of 3 MHz.

achieved. The second showed that for axisymmetric problems, the optimisation forward model can be replaced with a full-wave model to more accurately model interactions with the kinoform structure.

There are several further directions along which the work in this Chapter can be developed. First, as in Chapter 6, the use of these kinoforms for correcting wavefront aberrations should be investigated. A recent work showed that this can be done for a single point. However, it can potentially be generalised to multiple points for a single frequency, or, to multiple points for different driving frequencies. Second, Chapter 5 demonstrated that using two material 3-D printing high

resolution structures with different optical or acoustic properties can be fabricated. If materials with appropriate parameters could be identified then this could be used to simultaneously modulate both the phase and amplitude of the transducer by varying the attenuation. This would enable precise realisation of highly complex target distributions as discussed in Sec. 2.1.5.2. Finally, in place of a planar transducer, kinoforms could instead be designed for a bowl transducer. When placed across a bowl, the complex field across the kinoform surface is related to the field across the focal plane via a Fourier transform. The advantage of this property is that, rather than change the focal depth, varying the driving frequency instead alters the size of the target distribution. This could be used to steer arrays of foci, in a plane, by varying the driving frequency.

Chapter 8

General conclusions

The aim of this thesis was to investigate a range of new techniques for the generation of arbitrary spatial distributions of acoustic pressure in three dimensions. This was motivated by a number of emerging applications in biomedical ultrasound for which precise control over the spatial and temporal properties of the acoustic fields is of paramount importance. A requirement for each of the techniques explored was that they avoid the primary limitations of the contemporary approaches. Namely, poor scaling to larger apertures and higher frequencies, and significant cost. Based on this, the optical generation of ultrasound and single element piezoelectric transducers were identified as promising avenues to explore. The thesis began with a summary of a range of different techniques that can be used to manipulate acoustic fields in 3-D. Following this, previous experimental techniques used to implement these methods, with the modalities of interest, were summarised.

The first technique, investigated in Chapters 3 and 4, was the use of a regularly spaced train of optical pulses to excite a spatially modulated optical absorber. It was shown that combining these techniques allows for the use of binary amplitude holography to control the resulting pressure distribution in 3-D. Experimentally, it was demonstrated that both individual foci and continuous patterns of steady state pressure can be generated with this approach. Numerically, it was found that a minimum of two optical pulses are required for the generation of a focus, the position of this focus can be shifted by varying the frequency content of the incident signal, and that the choice of phase in the design of single focus holograms has a signif-

icant effect on the distribution of pressure in the resulting fields. Additionally, an optimisation approach, adapted from optics, was introduced for the design of these holograms. Different initialisations and search procedures were explored for this algorithm to maximise its performance and efficiency. The work in these chapters serves as the first demonstration that focusing in 3-D can be achieved using a train of discrete optical pulses in addition to a sinusoidal source.

The second technique, investigated in Chapter 5, was the use of a single optical pulse to excite a multi-layer patterned absorber. It was shown, experimentally, that by appropriately tailoring the spatial pattern of optical absorption across multiple layers, multiple acoustic foci could be generated with a single optical pulse. Two methods were introduced for fabricating these samples, both of which were validated. The first involved attaching together multiple discrete absorbers fabricated through laser cutting, the second 3-D printing with two materials possessing different optical properties. An algorithm was developed for designing these multi-layer absorbers. This built on the direct binary search approach, modifying the forward model, initialisation, and search method. For the forward model, a ray-tracing approach was developed for the rapid approximation of the impulse response of pixels embedded in a homogeneous substrate to a target point. The work of this chapter is the first demonstration of multi-layer absorbers for the generation of multi-focal acoustic fields.

The third technique, investigated in Chapter 6, was the use of a single optical pulse to excite a tailored optically absorbing surface profile. It was demonstrated experimentally, that by appropriately tailoring the shape of the absorbing surface, acoustic fields focused over a continuous pattern of pressure could be generated. A new algorithm, also building on direct binary search, was developed for designing these surface profiles from an arbitrary set of input foci. Temporal control over the distribution of peak pressure was demonstrated, and, it was shown that the calculated profiles can be cheaply generated through 3-D printing. The work of this chapter is a generalisation of an approach employed by a number of earlier investigators. It extends it to arbitrary spatial distributions of acoustic pressure, greatly

improving the flexibility of the technique.

The final technique, investigated in Chapter 7, was the use of multi-frequency acoustic kinoforms to map the output field of a planar single element transducer onto multiple diffraction-limited distributions of pressure that could be switched by altering the driving frequency. Experimentally it was shown that the field of a single element transducer could be mapped onto a minimum of 3 complex target distributions with good SNR. These patterns could consist of both arrangements of discrete foci or continuous patterns of pressure. Numerically it was established that to maximise the kinoform performance, the frequency spacing should be maximised and the maximum thickness should be constrained. An optimisation approach was again developed for the design of these kinoforms and several extensions were introduced for improving the flexibility and performance of the base algorithm. The work of this chapter builds on a technique introduced by Melde et al [17], greatly improving its flexibility and therefore its possible applications. It also contributes the first investigation of the impact of the forward model accuracy on the performance of these acoustic kinoforms.

Overall, this thesis constitutes a highly varied body of work into the spatial manipulation of acoustic fields in 3-D. The potential applications to which this can be applied were not the focus of the work and, in practice, will vary for each of the techniques. However, different future research directions have been proposed for each technique, and, as biomedical ultrasound develops and technology improves, the contributions of this thesis could form the basis for substantial future research.

Appendix A

Appendix A: Fibre coupling

A.1 Introduction

The experimental set-ups that were reported in Chapter 5 and 6 both used high peak power Q-switched lasers for which the output pulses were coupled into 1 and 1.5 mm multi-mode optical fibres. The fibre coupling of these sources is challenging, and can result in optical damage to the fibres under a variety of conditions.

As part of this thesis, a study was carried out investigating the coupling of these high-energy high-peak power optical pulses into 1.5 mm multimode optical fibres. The aims of the study were to establish:

1. The maximum energy that can be coupled safely into 1.5 mm multimode optical fibres.

2. The conditions under which damage occurs to these fibres and the different damage mechanisms.

3. Whether either the maximum energy that can be coupled or the conditions for damage vary when lenses with a different numerical aperture (NA) or focal length are used.

A.2 Equipment and methods

Two lasers were used in this study: a Litron Nano L 200-15, and a Minilite. The key properties of each are summarised in Tab. A.1. These were used in conjunction

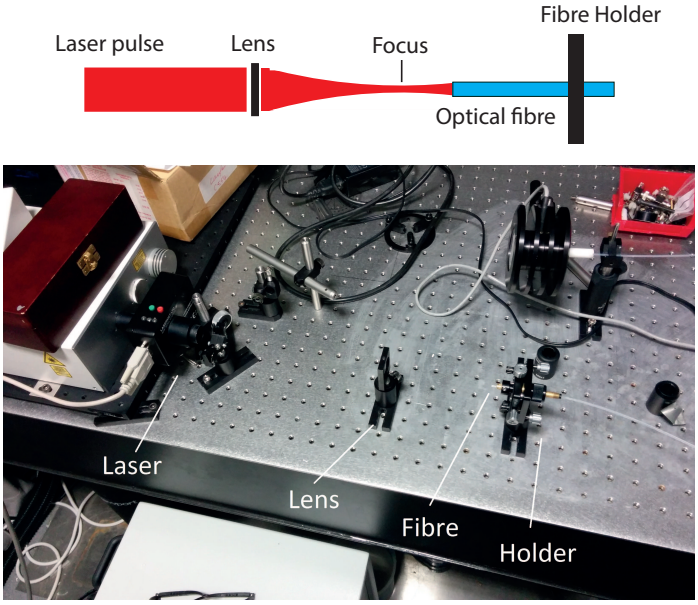

Figure A.1: Photo and schematic of the standard set-up for fibre coupling.

with 4 lenses with identical apertures and focal lengths of 8 cm, 6.3 cm, 4 cm and 2.5 cm.

Laser	Mode structure	Max pulse energy
Litron Nano 200-L	Gaussian TEM_{00} Multimode	14 mJ 300 mJ
Minilite	Gaussian TEM_{00}	37 mJ

Table A.1: Summary of the properties of the different lasers used.

A standard set-up was used for coupling the lasers; as shown in Fig. A.1. The optical fibre was placed in a multi-axis fibre positioner able to vary the position (x,y,z) of the optical fibre tip and the angle (θ,ϕ) of the fibre axis relative to the laser. A lens of varying focal length was used to focus the laser. The tip of the optical fibre was positioned behind the focal point of the lens. The laser pulse was then coupled into the fibre by attenuating the laser output using an ND filter and manually aligning the fibre tip with the beam using a viewing card. To test whether the coupling was on axis, the position of the fibre within the holder was varied by a few centimetres and the alignment tested at this new depth.

A.3 Litron Nano L 200-15 - single mode

A.3.1 Experiments

An initial set of experiments were carried out to establish the conditions under which fibre damage occurs. These involved coupling the pulses generated by the Litron Nano L 200-15. For each experiment, a lens was selected and the output of the laser coupled using the approach described in Sec. A.2. The coupling efficiency was started out low ($\sim 50\%$) by positioning the fibre at a large distance from the focal point of the lens. The energy of the laser was then set to a minimum and gradually increased to a maximum recording the input and output energy of the fibre using a power meter. The fibre was then repositioned closer to the lens to increase the coupling efficiency and the measurements repeated. This was continued until the coupling efficiency reached $\sim 80\%$ or the fibre was damaged. At this point the experiments were repeated for a different lens.

A.3.2 Results and discussion

Table A.2 summarises the conditions under which damage to the fibre occurred and the nature of the damage observed. From this a number of observations can be made. Firstly, there are 2 distinct modes of damage: damage to the front face of the fibre, and damage several centimetres down its length.

Damage to the front of the fibre is caused by proximity of the fibre tip to the focal point of the lens. It was not observed further than 1.43f (where f is the focal length of the lens) for any lens and was generally observed to occur at high pulse energy. The chance of damage appears inversely proportional to lens focal length occurring for greater relative depths (1.43f vs 1.29f) for the shortest focal length. However, the sample size is too small to confirm this.

Damage down the length of the fibre was observed under more variable conditions. On average, it happened at greater depths than sparking on the front of the fibre, occurring at distances ranging from 1.4-2.16f. The chance of breaking appears proportional to both the input energy of the laser and the lens focal length, as the instances of damage all occurred at close to the maximum laser output and for

Lens (focal length)	Distance (mm) (distance/focal length (d/f))	Coupling efficiency (%)	Laser output (mJ)	Damage description
75 mm	131 (1.75)	77	14	Ablation 8 cm into the fibre when coupling at max output
75 mm	105 (1.4)	80	13.3	Small drop in coupled energy observed, subsequently fibre snapped when moving it. Damage occurred at 7.6 cm into the fibre
75 mm	162 (2.16)	70	$\sim$18	Ablation 8 cm into the fibre, damage occurred when immediately coupling the output of the laser. This coincided with a brief peak in the laser output estimated at 18 mJ
63 mm	98 (1.56)	76-78	12.2	Ablation 8 cm into fibre
63 mm	75 (1.19)	80	10	Sparking damaging the front face of the fibre
75 mm	97 (1.29)	76	11.7	Sparking damaging the front face of the fibre
40 mm	57 (1.43)	75	14	Sparking damaging the front face of the fibre

Table A.2: Summary of conditions under which damage was recorded to optical fibres and the nature of the damage that occurred for experiments using the Litron L 200-15 with TEM00 operation.

lenses with a focal length over 63 mm. However, between the 5 instances reported in Tab. A.2 and further measurements, the pitch and yaw (θ,ϕ) of the fibre holder was adjusted. After this adjustment, no further instances of damage down the length of the fibre were observed. This suggests that the damage mechanism is angularly dependent. To validate these observations, higher pulse energy was required to provide a greater range over which the damage could occur to identify any focal length dependence.

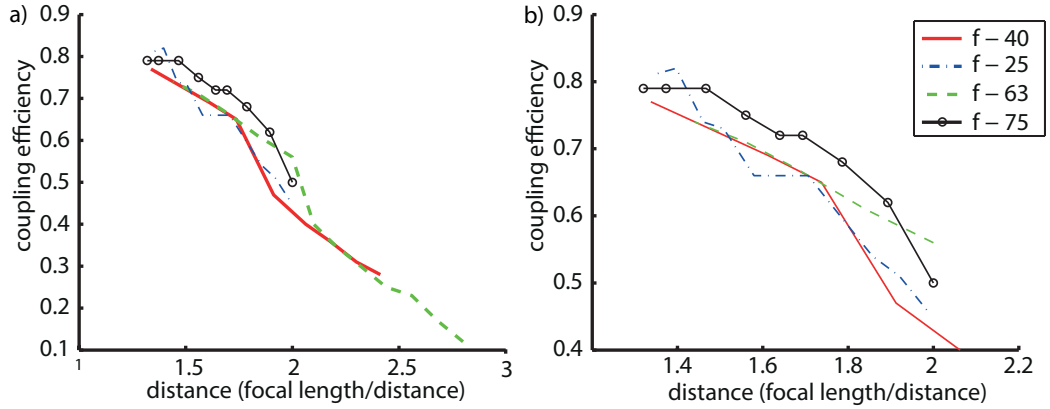

Figure A.2: (a) Variation in coupling efficiency with distance from lens for lens with focal lengths of 25, 40, 63, and 75 mm. (b) Subset of data from (a) for clarity.

A.4 Minilite

A.4.1 Experiments

The experiments from Sec. A.3 were repeated using the Minilite. This had a higher maximum pulse energy between 33-37 mJ, however, the output couldn't be varied. Therefore, the coupled energy was varied by changing the depth of the fibre relative to the lens in short steps. As in Sec. A.3, the coupling efficiency was gradually increased for each lens until it either reached $\sim 80\%$ or damage to the fibre occurred.

A.4.2 Results and discussion

Figure A.2 shows the measured variation in coupling efficiency with fibre depth for the four lenses used in this study. It can be seen that it decreases approximately linearly with distance from the lens and is mostly independent of focal length. Slightly higher coupling efficiencies (2-4%) are found for lenses possessing longer focal lengths. The distance at which the coupling efficiency stops increasing is $\sim$1.35f for each lens.

Table A.3 summaries the conditions under which the fibres were recorded breaking or sparks generated on the front face for the Minilite. It can be seen that the only form of damage observed for the Minilite was sparking and damage to the front face of the optical fibre, explicable by proximity to the lens. The shorter focal lengths were again found to generate sparks at a greater distance (1.36f vs 1.29f).

Lens (focal length)	Distance (mm) (distance/focal length (d/f))	Coupling efficiency (%)	Laser output (mJ)	Damage description
75 mm	97 (1.29)	78	33	Sparking on fibre, when fibre was removed from holder it snapped at approximately 9 mm down the fibre length
40 mm	49 (1.23)	N/A	33	Sparking on fibre, no change in coupling efficiency or damage observed
25 mm	34 (1.36)	N/A	33	Sparking on fibre, no change in coupling efficiency or damage observed

Table A.3: Summary of conditions under which damage was recorded to optical fibres and the nature of the damage that occurred for experiments using the Minilite.

A.4.3 Angular coupling measurement

Section A.4.2 provided more evidence that damage down the length of the fibre is a consequence of angular misalignment between the fibre and laser axes, as it wasn't observed with any lens despite a 3-fold increase in pulse energy. A set of measurements were carried out intentionally misaligning the optical fibre and laser to test this hypothesis.

As an initial test, the fibre holder was skewed to a maximum angle. The laser was then coupled into this skewed fibre using the 75 mm lens. A low coupling efficiency was observed (60%) followed by damage a short way into the fibre (0.5 cm). The damage happened geometrically where the laser first hit the fibre edge and was localised to that side. Figure A.3 illustrates this experiment. This demonstrated that damage within the fibre can be generated by angular misalignment.

As a second test again using the 75 mm lens, the fibre was aligned directly with the lens minimising the skew. The angle of the fibre holder was then adjusted gradually (2 turns between each measurement), realigned after each adjustment and the coupling efficiency recorded. For the first 4 adjustments of the fibre angle, the coupling efficiency remained constant. After the 5th adjustment, the fibre broke

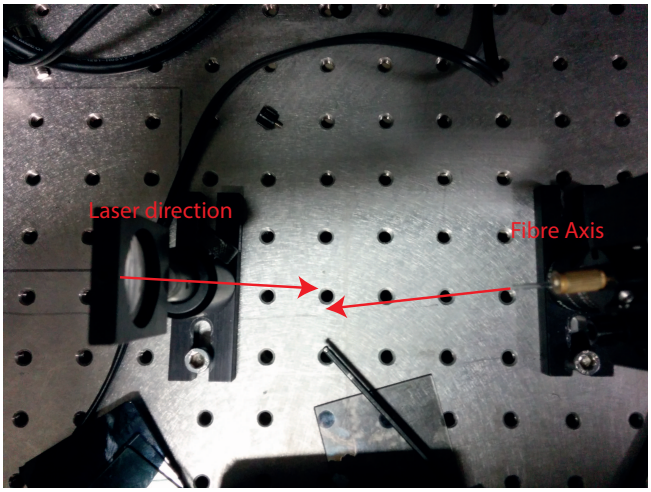

Figure A.3: Photograph of the first experiment used to test the effect of intentional misalignment.

ablating 7.8 cm down the length of the fibre. This depth is consistent with the results in Sec. A.3.2 and suggests that damage down the length of the fibre is a consequence of skew between the fibre axis and the laser.

As a third test, the 40 mm lens was used to check whether the chance of ablation changes using a more divergent beam. As before, the fibre was initially coupled minimising the skew. The skew was then gradually increased with the fibre being realigned at each step, and the coupling efficiency re-recorded. For the 40 mm, lens the fibre was successfully coupled without a drop in efficiency up to 9 turns. At 10 turns, the coupling efficiency reduced by 5%, and at 11 turns sparking occurred on the front face of the fibre damaging it. This is different from the measurements using the 75 mm lens and indicates that the chance of damage by ablation down the length of the fibre can be reduced by using a lens with a shorter focal length. This is supported by the lack of any occasion in Sec. A.3.2 of damage down the length of the fibre using a shorter focal length. However, the sparking damage did occur at a similar point for misalignment (11 turns) as the ablation using the 75 mm lens, so this needs further measurements and a reliable technique for measurement of the coupling angle between the fibre and laser.

A.4.4 Time dependence

Within this study, relatively little work was done to look at the stability of the fibre coupling (i.e., does fibre damage occur after a lengthy period of time). This was due to the need to run the laser for $\geq$10 hours to simulate an actual experiment from this thesis and the large parameter space of damage mechanisms. The damage that has been reported was all observed either instantaneously or within $\sim$1 minute of switching on the laser.

A set of 4 tests were done coupling the Minilite at high efficiency and running the laser for 30-60 minutes to try to explore this. For 3 of these tests the laser was stably coupled. In 1 case, the fibre broke at approximately 8.5 cm down the length of the fibre after 15 minutes of operation. The nature of the damage matched the angular dependent damage, but, occurred at a greater depth (by 0.5 cm) than was observed in any other instance. This suggests that the angularly dependent damage has a temporal property, and that where a slight skew is present, damage could be caused after prolonged running of the laser.

A.5 Multi-mode

For a final set of tests, the aperture was removed from the Litron L 200-15 laser cavity. This had the effect of changing the structure of the laser pulse to a multi-mode while increasing the maximum pulse energy to 300 mJ. A set of measurements were made coupling this increased output, the results of which are shown in Fig. A.4.

Figure A.4 shows that high energies can be safely coupled into the fibre when the beam output is multi-mode, with the peak output achieved being $\sim$130 mJ. No damage was observed, although it remains to be tested whether there was any temporal instability in the coupling. The input energy was stopped at 180 mJ as at this energy using a lens to focus the beam resulted in loud sparking and air breakdown, even using a lens with a longer (75 mm) focal length.

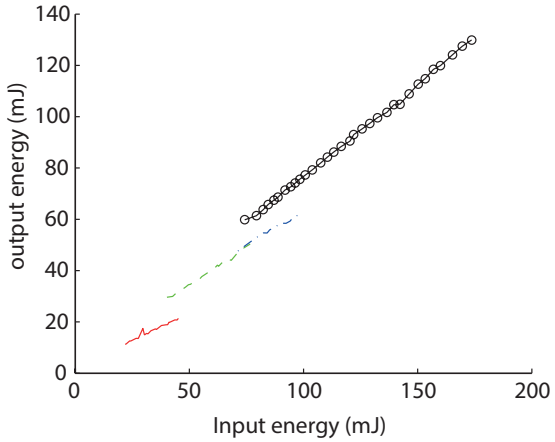

Figure A.4: Recorded input and output energy of the fibre for measurements using the Litron L 200-15 with multi-mode pulse structure. The different curves were acquired from different positions of the fibre

A.6 Conclusions

Fibre coupling of high energy high peak power optical pulses into multi-mode optical fibres has been investigated. Two modes of damage have been identified:

- Sparking and front face damage. This is a consequence of proximity of the fibre to the lens focus and was generally found to occur at greater relative depths for shorter focal lengths.
- Damage down the length of the fibre. This has an angular dependence and potentially a direct dependence on lens NA.

To minimise the chances of damage when coupling the optical fibres it is recommended to

- Position the fibre at least 1.4 (f/D) away from the lens, regardless of the focal length being used. If attempting to maximise the energy then 1.35 (f/D) is largely safe, though damage has been observed at this point.
- Check the beam is on axis over a large range $\sim$5-10 cm to minimise any angular misalignment.

For choice of lens

- Longer focal length lenses are easier to couple with high efficiency from a practical standpoint.

- Shorter focal length lenses can be used, however, it was found generally the maximum coupling efficiencies were a couple of percent lower (Fig. A.2) and the distance between high coupling efficiency and damage is shorter so it poses a higher risk. Additionally, sparking and air-breakdown were both seen using the 25 and 40 mm lenses in free-space. However, the angularly dependent damage was never observed for a shorter focal length lens, so these potentially represent a more stable option.

Appendix B

Appendix B: Publications, conferences, and awards

B.1 Journal publications

1. M. D. Brown, J. Jaros, B. T. Cox, B. E. Treeby, Control of broadband optically generated ultrasound pulses using binary amplitude holograms, *Journal of the Acoustical Society of America*, 139 (4), pp. 1637-1647, 2016.

2. M. D. Brown, D. I. Nikitichev, B. E. Treeby, B. T. Cox, Generating arbitrary ultrasound fields with tailored optoacoustic surface profiles, *Applied Physics Letters*, 110 (9), p. 094102, 2017.

3. M. D. Brown, B. T. Cox, B. E. Treeby, Design of multi-frequency acoustic kinoforms, *Applied Physics Letters*, 111(24), p. 244101, 2017.

B.2 Conference publications

1. M. D. Brown, T. J. Allen, B. T. Cox, B. E. Treeby, Control of optically generated ultrasound fields using binary amplitude holograms, In: IEEE International Ultrasonics Symposium, Chicago, IL, USA, 2014.

2. M. D. Brown, E. Martin, B. T. Cox, and B. E. Treeby, Single pulse illumination of multi-layer photoacoustic holograms for patterned ultrasound field

generation, In: IEEE International Ultrasonics Symposium, Tours, France, 2016.

3. M. D. Brown, B. T. Cox, and B. E. Treeby, Investigating the effect of thickness and frequency spacing on multi-frequency acoustic kinoforms, In: IEEE International Ultrasonics Symposium, Washington D.C., USA, 2017.

B.3 Awards

1. **Annual PhD prize for best paper** - Department of Medical Physics and Biomedical Engineering, 2017.

2. **Student paper competition finalist** - IEEE International Ultrasonics Symposium, 2017.

3. **Student travel award** - $700, IEEE International Ultrasonics Symposium, 2017.

4. **Student travel award** - $380, IEEE International Ultrasonics Symposium, 2016.

Bibliography

[1] Karl Theo Dussik. Uber die Moglichkeit , hochfrequente mechanische Sehwingungen als diagnostisehes Hilfsmittel zu verwertenl . *Zeitschrift für die gesamte Neurologie und Psychiatrie*, 174(1):153–168, 1941.

[2] I Edler and CH Hertz. Use of ultrasound reflectoscope for the continuous recording of movements of heart walls. *Kungl Fysiogr Sallsk Lund Forh*, 24:40, 1954.

[3] W. J. Fry. Use of intense ultrasound in neurological research. *American Journal of Physical Medicine*, 37(3):143–147, 1958.

[4] Thomas L Szabo. *Diagnostic Ultrasound Imaging*. 2013.

[5] Jean-Michel Escoffre and Ayache Bouakaz. *Therapeutic Ultrasound*. 2015.

[6] J. Nilsson, M. Evander, B. Hammarström, and T. Laurell. Review of cell and particle trapping in microfluidic systems. *Analytica Chimica Acta*, 649(2):141–157, 2009.

[7] Omer Naor, Steve Krupa, and Shy Shoham. Ultrasonic neuromodulation. *Journal of Neural Engineering*, 13(3):031003, 2016.

[8] Emad S. Ebbini and Charles A. Cain. Multiple-Focus Ultrasound Phased-Array Pattern Synthesis: Optimal Driving-Signal Distributions for Hyperthermia. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 36(5):540–548, 1989.

[9] Emad S. Ebbini, Shin Ichiro Umemura, Mohammed Ibbini, and Charles A. Cain. A Cylindrical-Section Ultrasound Phased-Array Applicator for Hyperthermia Cancer Therapy. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 35(5):561–572, 1988.

[10] Douglas R. Daum and Kullervo Hynynen. A 256-element ultrasonic phased array system for the treatment of large volumes of deep seated tissue. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 46(5):1254–1268, 1999.

[11] H Wan, P VanBaren, E S Ebbini, and C A Cain. Ultrasound surgery: Comparison of strategies using phased array systems. *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 43(6):1085–1098, 1996.

[12] Y. Hertzberg and G. Navon. Bypassing absorbing objects in focused ultrasound using computer generated holographic technique. *Medical Physics*, 38(12):6407–6415, 2011.

[13] Svetlana Bobkova, Leonid Gavrilov, Vera Khokhlova, Adam Shaw, and Jeffrey Hand. Focusing of high-intensity ultrasound through the rib cage using a therapeutic random phased array. *Ultrasound in Medicine and Biology*, 36(6):888–906, 2010.

[14] Yoni Hertzberg, Omer Naor, Alexander Volovick, and Shy Shoham. Towards multifocal ultrasonic neural stimulation: pattern generation algorithms. *Journal of Neural Engineering*, 7(5):056002, 2010.

[15] Michael D Menz, Omer Oralkan, Pierre T Khuri-Yakub, and Stephen a Baccus. Precise neural stimulation in the retina using focused ultrasound. *The Journal of Neuroscience*, 33(10):4550–60, mar 2013.

[16] Dana Gourevich, Yoni Hertzberg, Alexander Volovick, Yaron Shafran, Gil Navon, Sandy Cochran, and Andreas Melzer. Ultrasound-Mediated Targeted Drug Delivery Generated by Multifocal Beam Patterns: An In vitro Study. *Ultrasound in Medicine and Biology*, 39(3):507–514, 2013.

[17] Kai Melde, Andrew G. Mark, Tian Qiu, and Peer Fischer. Holograms for acoustics. *Nature*, 537(7621):518–522, 2016.

[18] Asier Marzo, Sue Ann Seah, Bruce W. Drinkwater, Deepak Ranjan Sahoo, Benjamin Long, and Sriram Subramanian. Holographic acoustic elements for manipulation of levitated objects. *Nature Communications*, 6(May):8661, 2015.

[19] Bt Hefner and Pl Marston. An acoustical helicoidal wave transducer with applications for the alignment of ultrasonic and underwater systems. *Journal of the Acoustical Society of America*, 106:3313–3316, 1999.

[20] Thomas Laurell, Filip Petersson, and Andreas Nilsson. Chip integrated strategies for acoustic separation and manipulation of cells and particles. *Chem. Soc. Rev.*, 36(3):492–506, 2007.

[21] Xiaoyun Ding, Sz-Chin Steven Lin, Brian Kiraly, Hongjun Yue, Sixing Li, I-Kao Chiang, Jinjie Shi, Stephen J Benkovic, and Tony Jun Huang. On-chip manipulation of single microparticles, cells, and organisms using surface acoustic waves. *Proceedings of the National Academy of Sciences of the United States of America*, 109(28):11105–11109, jul 2012.

[22] Charles Hudin, José Lozada, and Vincent Hayward. Localized tactile feedback on a transparent surface through time-reversal wave focusing. *IEEE Transactions on Haptics*, 8(2):188–198, 2015.

[23] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. Rendering volumetric haptic shapes in mid-air using ultrasound. *ACM Transactions on Graphics*, 33(6):1–10, 2014.

[24] Bruce W. Drinkwater and Paul D. Wilcox. Ultrasonic arrays for non-destructive evaluation: A review. *NDT and E International*, 39(7):525–541, 2006.

[25] Bakary Diarra, Marc Robini, Piero Tortoli, Christian Cachard, and Herve Liebgott. Design of optimal 2-d nongrid sparse arrays for medical ultrasound. *IEEE Transactions on Biomedical Engineering*, 60(11):3093–3102, 2013.

[26] David E. Dausch, John B. Castellucci, Derrick R. Chou, and Olaf T. Von Ramm. Theory and operation of 2-D array piezoelectric micromachined ultrasound transducers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 55(11):2484–2492, 2008.

[27] R W Prager, U Z Ijaz, A H Gee, and G M Treece. Three-dimensional ultrasound imaging. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 224(2):193–223, 2010.

[28] Philippe Acar, Laia Battle, Yves Dulac, Marianne Peyre, Hélène Dubourdieu, Sébastien Hascoet, Marion Groussolles, and Christophe Vayssière. Real-time three-dimensional foetal echocardiography using a new transabdominal xMATRIX array transducer. *Archives of Cardiovascular Diseases*, 107(1):4–9, 2014.

[29] B. Savord and R. Solomon. Fully sampled matrix transducer for real time 3D ultrasonic imaging. *IEEE Symposium on Ultrasonics*, pages 945–953, 2003.

[30] Kjell Kristoffersen, Bruno Haider, Thomas Halvorsroed, and Steinar Bjaerum. Transmit beamforming in 3-dimensional ultrasound, 2012.

[31] Y. Lu, H. Tang, Q. Wang, S. Fung, J. M. Tsai, M. Daneman, B. E. Boser, and D. A. Horsley. Waveguide piezoelectric micromachined ultrasonic transducer array for short-range pulse-echo imaging. *Applied Physics Letters*, 106(19), 2015.

[32] Kang Qiao Zhao, Tore Gruner Bjastad, and Kjell Kristoffersen. Error analysis of subaperture processing in 1-D ultrasound arrays. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 62(4):663–672, 2015.

[33] O. Martínez, C. J. Martín, G. Godoy, and L. G. Ullate. 2D array based on fermat spiral. *Physics Procedia*, 3(1):399–406, 2010.

[34] F. Reverdy, G. Ithurralde, and N. Dominguez. Advanced Ultrasonic 2D Phased-Array Probes. *18th World Conference on Nondestructive Testing*, pages 16–20, 2012.

[35] Yang Hou, Jin Sung Kim, Shai Ashkenazi, Sheng Wen Huang, L. Jay Guo, and Matthew O'Donnell. Broadband all-optical ultrasound transducers. *Applied Physics Letters*, 91(7):89–92, 2007.

[36] C. Sheaff and S. Ashkenazi. Characterization of an improved polyimide-etalon all-optical transducer for high-resolution ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 61(7):1223–1232, 2014.

[37] Butrus T Khuri-Yakub and Omer Oralkan. Capacitive micromachined ultrasonic transducers for medical imaging and therapy. *Journal of micromechanics and microengineering : structures, devices, and systems*, 21(5):54004–54014, 2011.

[38] A. I. H Chen, Lawrence Wong, and John Yeow. Recent advances in Capacitive Micromachined ultrasonic transducer imaging systems. In *Medical Imaging*, pages 253–273. 2013.

[39] Anshuman Bhuyan, Jung Woo Choe, Byung Chul Lee, Ira O Wygant, Amin Nikoozadeh, Ömer Oralkan, and Butrus T Khuri-Yakub. Integrated circuits for volumetric ultrasound imaging with 2-D CMUT arrays. *IEEE transactions on biomedical circuits and systems*, 7(6):796–804, dec 2013.

[40] Armen Sarvazyan, Laurent Fillinger, and Leonid R. Gavrilov. Time-reversal acoustic focusing system as a virtual random phased array. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 57(4):812–817, 2010.

[41] Ching-hua Chou, S Kino, and I Introduction. Lens Design for Acoustic Microscopy. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 35(4):464–469, 1988.

[42] Hongxia Yao. *Synthetic aperture methods for medical ultrasonic imaging*. PhD thesis, 1997.

[43] M Fink. Time reversal of ultrasonic fields. I. Basic principles. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 39(5):555–566, 1992.

[44] Mathias Fink. Time-reversal acoustics. *Journal of Physics: Conference Series*, 118:012001, 2008.

[45] James L. B. Robertson, Ben T. Cox, J. Jaros, and Bradley E. Treeby. Accurate simulation of transcranial ultrasound propagation for ultrasonic neuromodulation and stimulation. *The Journal of the Acoustical Society of America*, 141(3):1726–1738, 2017.

[46] Dennis Gabor. A New Microscopic Principle. *Nature*, 171:439 – 440, 1948.

[47] Oleg A. Sapozhnikov, Sergey A. Tsysar, Vera A. Khokhlova, and Wayne Kreider. Acoustic holography as a metrological tool for characterizing medical ultrasound sources and fields. *The Journal of the Acoustical Society of America*, 138(3):1515–1532, 2015.

[48] J. D. Maynard, E. G. Williams, and Y. Lee. Nearfield acoustic holography: I. Theory of generalized holography and the development of NAH. *The Journal of the Acoustical Society of America*, 78(4):1395–1413, 1985.

[49] Efren Grande. *Near-field acoustic holography with sound pressure and particle velocity measurements*. PhD thesis, 2012.

[50] Joseph W. Goodman. *Introduction to Fourier Optics, Third Edition*. 2004.

[51] Xiaozheng Zeng and Robert J. McGough. Evaluation of the angular spectrum approach for simulations of near-field pressures. *Journal Acoustical Society of America*, 123(1):68–76, 2008.

[52] Robert L Lucke. RayleighSommerfeld diffraction and Poisson's spot. *European Journal of Physics*, 27(2):193–204, 2006.

[53] L G Neto, D Roberge, and Y Sheng. Full-range, continuous, complex modulation by the use of two coupled-mode liquid-crystal televisions. *Applied Optics*, 35(23):4567–76, 1996.

[54] Alexander Jesacher, Christian Maurer, Andreas Schwaighofer, Stefan Bernet, and Monika Ritsch-Marte. Near-perfect hologram reconstruction with a spatial light modulator. *Optics Express*, 16(4):2597–2603, 2008.

[55] Mohammed S. Ibbini and Charles A. Cain. A Field Conjugation Method for Direct Synthesis of Hyperthermia Phased-Array Heating Patterns. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 36(1):3–9, 1989.

[56] R W Gerchberg and W O Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35(2):237–246, 1972.

[57] J R Fienup. Phase retrieval algorithms: a comparison. *Appl. Optics*, 21(15):2758–2769, 1982.

[58] N C Gallagher and B Liu. Method for Computing Kinoforms that Reduces Image Reconstruction Error. *Applied Optics*, 12(10):2328, 1973.

[59] Rainer G Dorsch, Adolf W Lohmann, and Stefan Sinzinger. Fresnel ping-pong algorithm for two-plane computer-generated hologram display. *Applied Optics*, 33(5):869–875, 1994.

[60] Stephen Mellin and Gregory Nordin. Limits of scalar diffraction theory and an iterative angular spectrum algorithm for finite aperture diffractive optical element design. *Optics Express*, 8(13):705, 2001.

[61] Olivier Ripoli, Ville Kettunen, and Peter Herzig Hans. Review of iterative Fourier-transform algorithms for beam shaping applications. *Optical Engineering*, 43(11):2549–2556, 2004.

[62] Frank Wyrowski. Diffractive optical elements: iterative calculation of quantized, blazed phase structures. *Journal of the Optical Society of America A*, 7(6):961, 1990.

[63] Michael A. Seldowitz, Jan P. Allebach, and Donald W. Sweeney. Synthesis of digital holograms by direct binary search. *Applied Optics*, 26(14):2788, 1987.

[64] Nobukazu Yoshikawa and Toyohiko Yatagai. Phase optimization by simulated annealing. *Applied Optics*, 33(5):863–868, 1994.

[65] Andrew G. Kirk and Trevor J. Hall. Design of binary computer generated holograms by simulated annealing observation of metastable states. *Journal of Modern Optics*, 39(12):2531–2539, 1992.

[66] Yen-wei Chen, Shinichiro Yamauchi, and Ning Wang. A Fast Kinoform Optimization Algorithm Based on Simulated Annealing. *IEICE Transactions Fundamentals of Electronics,Communications and Computer Sciences*, 1(4):774–776, 2000.

[67] N. Yoshikawa, M. Itoh, and T. Yatagai. Use of genetic algorithm for computer-generated holograms. *Proceedings of SPIE*, 2577:150–157, 1995.

[68] Jin Seon Yoon and Nam Kim. Optimization of Diffrative Optical Elements by Genetic Algorithm. *Journal of the Optical Society of Korea*, 4(1):30–36, 2000.

[69] Alexander Graham Bell. On the production and reproduction of sound by light. *Journal of the Society of Telegraph Engineers*, 9(34):404–426, 1880.

[70] S J Davies, C Edwards, G S Taylor, and S B Palmer. Laser-generated ultrasound: its properties, mechanisms and multifarious applications. *Journal of Physics D: Applied Physics*, 26(3):329–348, 1993.

[71] M. W. Sigrist. Laser generation of acoustic waves in liquids and gases. *Journal of Applied Physics*, 60(7):R83–R121, 1986.

[72] D. A. Hutchins. Ultrasonic generation by pulsed lasers. *Physical Acoustics*, 18:21–123, 1988.

[73] Ben Cox, Jan G. Laufer, Simon R. Arridge, and Paul C. Beard. Quantitative spectroscopic photoacoustic imaging: a review. *Journal of Biomedical Optics*, 17(6):061202, 2012.

[74] Olivier Simandoux. *Photoacoustic imaging : contributions to optical-resolution photoacoustic endoscopy and experimental investigation of thermal nonlinearity*. PhD thesis, 2015.

[75] Yong Zhou, Junjie Yao, and Lihong V. Wang. Tutorial on photoacoustic tomography. *Journal of Biomedical Optics*, 21(6):061007, 2016.

[76] Srinath Rajagopal, Toby Sainsbury, Bradley E Treeby, Ben T Cox, and A Carbon-polymer Nanocomposites. Laser Generated Ultrasound Sources Using Polymer Nanocomposites for High Frequency Metrology. *IEEE International Ultrasonics Symposium*, pages 3–6, 2017.

[77] G J Diebold and T Sun. Properties of potoacoustic waves in one, two, and three dimensions. *Physical Review Letters*, 67(24):3384–3387, 1991.

[78] M Khan. The photoacoustic effect generated by an isotropic solid sphere. *Ultrasonics*, 33(4):265–269, jul 1995.

[79] C G A Hoelen and F F M De Mul. A new theoretical approach to photoacoustic signal generation. *Journal Acoustical Society of America*, 106(2):695–706, 1999.

[80] B. T. Cox and P. C. Beard. Fast calculation of pulsed photoacoustic fields in fluids using k-space methods. *The Journal of the Acoustical Society of America*, 117(6):3616, 2005.

[81] Bradley E. Treeby and B. T. Cox. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of Biomedical Optics*, 15(2):021314, 2010.

[82] P. Beard. Biomedical photoacoustic imaging. *Interface Focus*, 1(4):602–631, 2011.

[83] Sung-Liang Chen. Review of Laser-Generated Ultrasound Transmitters and Their Applications to All-Optical Ultrasound Transducers and Imaging. *Applied Sciences*, 7(1):25, 2016.

[84] Richard J. Colchester, Edward Z. Zhang, Charles A. Mosse, Paul C. Beard, Ioannis Papakonstantinou, and Adrien E. Desjardins. Broadband miniature optical ultrasound probe for high resolution vascular tissue imaging. *Biomedical Optics Express*, 6(4):1502, 2015.

[85] Pieter Kruizinga, Ben Cox, Nico De Jong, Paul Beard, Anton F W Van Der Steen, and Gijs Van Soest. Plane wave ultrasound imaging with a broadband photoacoustic source. *IEEE International Ultrasonics Symposium, IUS*, pages 1414–1416, 2012.

[86] Bradley E Treeby, Benjamin T Cox, Edward Z Zhang, Sarah K Patch, and Paul C Beard. Measurement of broadband temperature-dependent ultrasonic attenuation and dispersion using photoacoustics. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 56(8):1666–1676, aug 2009.

[87] R. M. White. Generation of elastic waves by transient surface heating. *Journal of Applied Physics*, 34(12):3559–3567, 1963.

[88] R. J. von Gutfeld. Thermoelastic generation of elastic waves for non-destructive testing and medical diagnostics. *Ultrasonics*, 18(4):175–181, 1980.

[89] Elena Biagi, Massimo Brenci, Stefano Fontani, Leonardo Masotti, Massimiliano Pieraccini, and Massimiliano Pieraccinil. Photoacoustic generation: optical fiber ultrasonic sources for non destructive evaluation and clinical diagnosis. *Optical Review*, 4(4):481–483, 1997.

[90] Elena Biagi, Fabrizio Margheri, and David Menichelli. Efficient laser-ultrasound generation by using heavily absorbing films as targets. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 48(6):1669–1680, 2001.

[91] Takashi Buma, Monica Spisar, and Matthew O'Donnell. A high-frequency, 2-D array element using thermoelastic expansion in PDMS. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 50(9):1161–1176, 2003.

[92] Hyoung Won Baac, Jong G Ok, Adam Maxwell, Kyu-Tae Lee, Yu-Chih Chen, a John Hart, Zhen Xu, Euisik Yoon, and L Jay Guo. Carbon-nanotube optoacoustic lens for focused ultrasound generation and high-precision targeted therapy. *Scientific reports*, 2:989, jan 2012.

[93] Bao Yu Hsieh, Jinwook Kim, Jiadeng Zhu, Sibo Li, Xiangwu Zhang, and Xiaoning Jiang. A laser ultrasound transducer using carbon nanofibers-polydimethylsiloxane composite thin film. *Applied Physics Letters*, 106(2):2–7, 2015.

[94] Wei Yi Chang, Wenbin Huang, Jinwook Kim, Sibo Li, and Xiaoning Jiang. Candle soot nanoparticles-polydimethylsiloxane composites for laser ultrasound transducers. *Applied Physics Letters*, 107(16):1–5, 2015.

[95] E Vannacci, L Belsito, F Mancarella, M Ferri, G P Veronese, A Roncaglia, and E Biagi. Miniaturized fiber-optic ultrasound probes for endoscopic tissue analysis by micro-opto-mechanical technology. *Biomedical microdevices*, 16(3):415–426, jun 2014.

[96] Srinath Rajagopal. Ultrasound sources for metrology. Technical report, 2017.

[97] Nan Wu, Ye Tian, Xiaotian Zou, Vinicius Silva, Armand Chery, and Xingwei Wang. High-efficiency optical ultrasound generation using one-pot synthesized polydimethylsiloxane-gold nanoparticle nanocomposite. *Journal of the Optical Society of America B*, 29(8):2016–2020, 2012.

[98] Victoria Ageeva. *Endoscopic Cheap Optical Transducers (CHOTs) for On-Wing Ultrasonic Inspection*. PhD thesis, 2016.

[99] A. M. Aindow, R. J. Dewhurst, and S. B. Palmer. Laser-generation of directional surface acoustic wave pulses in metals. *Optics Communications*, 42(2):116–120, 1982.

[100] A. D. W. Mckie, J. W. Wagner, J. B. Spicer, and C. M. Penney. Laser generation ultrasound of narrow-band and directed ultrasound. *Ultrasonics*, 27:323–330, 1989.

[101] S Dixon, T Harrison, Y Fan, and P A Petcher. Thermoelastic laser generated ultrasound using a ring source. *Journal of Physics D: Applied Physics*, 45(17):175103, 2012.

[102] P Cielo, F Nadeau, and M Lamontagne. Laser Generation of Convergent Acoustic-Waves for Materials Inspection. *Ultrasonics*, 23(2):55–62, 1985.

[103] X Wang and M G Littman. Laser cavity for generation of variable-radius rings of light. *Optics letters*, 18(10):767–8, 1993.

[104] Xiao Wang, Michael G. Littman, John B. McManus, Mohsen Tadi, Young Sik Kim, Attila Askar, and Herschel Rabitz. Focused bulk ultrasonic

waves generated by ring-shaped laser illumination and application to flaw detection. *Journal of Applied Physics*, 80(8):4274–4281, 1996.

[105] G. Cachier. Optical excitation of high-amplitude surface waves. *Applied Physics Letters*, 17(10):419–421, 1970.

[106] A. A. Maznev, K. A. Nelson, and J. A. Rogers. Optical heterodyne detection of laser-induced gratings. *Optics Letters*, 23(16):1319–1321, 1998.

[107] Tobias Brixner, Jens Stenger, Harsha M. Vaswani, Minhaeng Cho, Robert E. Blankenship, and Graham R. Fleming. Two-dimensional spectroscopy of electronic couplings in photosynthesis. *Nature*, 434(7033):625–628, 2005.

[108] Hideo Nishino, Yusuke Tsukahara, Yoshihiko Nagata, Toshio Koda, and Kazushi Yamanaka. Excitation of high frequency surface acoustic waves by phase velocity scanning of a laser interference fringe. *Applied Physics Letters*, 62(17):2036–2038, 1993.

[109] J Huang, S Krishnaswamy, and J D Achenbach. Laser generation of narrow-band surface waves. *Journal of the Acoustical Society of America*, 92(5):2527–2531, 1992.

[110] F. Laivza Di Scnlra, Francesco Lanza, Tobias P Berndt, James B Spicer, Associate Member, and B Boro Djordjevic. Remote laser generation of narrow-band surface waves through optical fibers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 46(6):1551–1557, 1999.

[111] H Andersson, M Ekberg, S Hård, S Jacobsson, M Larsson, and T Nilsson. Single photomask, multilevel kinoforms in quartz and photoresist: manufacture and evaluation. *Applied Optics*, 29(28):4259–4267, 1990.

[112] N Yoshikawa, M Itoh, and T Yatagai. Binary computer-generated holograms for security applications from a synthetic double-exposure method by electron-beam lithography. *Optics letters*, 23(18):1483–1485, 1998.

[113] S Zamiri, B Reitinger, S Bauer, and P Burgholzer. Converging Laser Generated Ultrasonic Waves using Annular Patterns Irradiation. *Journal of Physics: Conference Series*, 520(1):012001, 2014.

[114] Michael Kalms, Sandra Hellmers, Philipp Huke, and Ralf B. Bergmann. Beam shaping using liquid crystal-on-silicon spatial light modulators for laser ultrasound generation. *Optical Engineering*, 53(4):044110, 2014.

[115] T. Stratoudaki, J. A. Hernandez, M. Clark, and M. G. Somekh. Cheap optical transducers (CHOTs) for narrowband ultrasonic applications. *Measurement Science and Technology*, 18(3):843–851, 2007.

[116] Ahmet Arca, Jon Aylott, Leonel Marques, Matt Clark, Mike Somekh, Richard Smith, Steve Sharples, Teti Stratoudaki, and Xuesheng Chen. CHOTs optical transducers. *Nondestructive Testing and Evaluation*, 26(3-4):353–366, 2011.

[117] I. J. Collison, T. Stratoudaki, M. Clark, and M. G. Somekh. Measurement of elastic nonlinearity using remote laser ultrasonics and CHeap Optical Transducers and dual frequency surface acoustic waves. *Ultrasonics*, 48(6-7):471–477, 2008.

[118] Theodosia Stratoudaki, Matt Clark, Michael Somekh, and Ahmet Arca. CHeap Optical Transducers (CHOTs) for generation and detection of longitudinal waves. *IEEE International Ultrasonics Symposium, IUS*, pages 961–964, 2012.

[119] R. J. Von Gutfeld, D. R. Vigliotti, C. S. Ih, and W. R. Scott. Thermoelastic hologram for focused ultrasound. *Applied Physics Letters*, 42(12):1018–1020, 1983.

[120] M. Clark, F. Linnane, S. D. Sharples, and M. G. Somekh. Frequency control in laser ultrasound with computer generated holography. *Applied Physics Letters*, 72(16):1963–1965, 1998.

[121] Y Fukuda, T Kohmoto, M Kunitomo, K Ishikawa, M Tanigawa, Y Takahashi, and T Hashi. Optical generation of continuous-wave accumulated coherent phonons. *Journal of the Optical Society of America*, 9(11):2054–2058, 1992.

[122] P E Dyer. Generation and Detection of 160-MHz Acoustic Waves Using a TEA CO2 Laser Irradiated Polymer Film Combination. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 35(1):2–4, 1988.

[123] J. B. Deaton, A. D W McKie, J. B. Spicer, and J. W. Wagner. Generation of narrow-band ultrasound with a long cavity mode-locked Nd:YAG laser. *Applied Physics Letters*, 56(24):2390–2392, 1990.

[124] Jr John B Deaton and James W Wagner. Variable-cavity-length mode-locked Nd:YAG laser for noncontact generation and spectral control of narrow-band ultrasound. *Applied optics*, 33(6):1051–1058, 1994.

[125] Takahiro Hayashi and Ken Ishihara. Generation of narrowband elastic waves with a fiber laser and its application to the imaging of defects in a plate. *Ultrasonics*, 77:47–53, 2017.

[126] U Störkel, K L Vodopyanov, and W Grill. GHz ultrasound wave packets in water generated by an Er laser. *Journal of Physics D: Applied Physics*, 31(18):2258–2263, 1998.

[127] M. H. Noroy, D. Royer, and M. Fink. Transient elastic wave generation by an array of thermoelastic sources. *Applied Physics Letters*, 63(24):3276–3278, 1993.

[128] T.W. Murray, J.B. Deaton Jr, and J.W Wagnera. Experimental evaluation of enhanced generation of ultrasonic waves using an array of laser sources. *Ultrasonics*, 34(1):69–77, 1996.

[129] T. W. Murray and O. Balogun. High-sensitivity laser-based acoustic microscopy using a modulated excitation source. *Applied Physics Letters*, 85(14):2974–2976, 2004.

[130] O. Balogun and T. W. Murray. A frequency domain laser based ultrasonic system for time resolved measurement of broadband acoustic transients. *Journal of Applied Physics*, 100(3), 2006.

[131] Suraj Bramhavar, Bruno Pouet, and Todd W. Murray. Superheterodyne detection of laser generated acoustic waves. *Applied Physics Letters*, 94(11):10–13, 2009.

[132] R Pierce. Temporal modulation of a laser source for the generation of ultrasonic waves. *Ultrasonics*, 33(2):133–137, 1995.

[133] Terry Sanderson. Experimental and numerical results for intensity modulated laser ultrasonics. *The Journal of the Acoustical Society of America*, 104(4):2207–2212, 1998.

[134] Terry Sanderson, Charles Ume, and Jacek Jarzynski. A comparison of Q-switched and intensity modulated laser pulses for ultrasonic NDT. *Journal of Nondestructive Evaluation*, 17(4):199–208, 1998.

[135] A Meyer, S Gspan, S Bernet, and M. Ritsch-Marte. Tailoring ultrasonic beams with optoacoustic holography. In *Laser Resonators and Beam Control VI. Proc SPIE 4969*, volume 4969, pages 105–114, 2003.

[136] Stefan Gspan, Alex Meyer, Stefan Bernet, and Monika Ritsch-Marte. Synthetic acoustic holograms realized via optoacoustic methods. *SPIE Photons plus Ultrasound*, 5697:99–106, 2005.

[137] Y.H. Berthelot and Y Jarzynski, J. The use of optical fibers to enhance the laser generation of ultrasonic waves. *The Journal of the Acoustical Society of America*, 85(1):158–162, 1989.

[138] Yves H. Berthelot and Jacek Jarzynski. Directional laser generation and detection of ultrasound with arrays of optical fibers. *Journal of Nondestructive Evaluation*, 9(4):271–277, 1990.

[139] J. Yang, N. DeRidder, C. Ume, and J. Jarzynski. Non-contact optical fibre phased array generation of ultrasound for non-destructive evaluation of materials and processes. *Ultrasonics*, 31(6):387–394, 1993.

[140] Jr-syu Yang, Terry Sanderson, Charles Urne, and Jacek Jarzynski. Laser Phased Array Generated Ultrasound for Nondestructive Evaluation of Ceramic Materials. *Journal of Nondestructive Evaluation*, 16(1):1–9, 1997.

[141] M. H. Noroy, D. Royer, and M. A. Fink. Shear-Wave Focusing with a Laser-Ultrasound Phased-Array. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(6):981–988, 1995.

[142] Pasi Karppinen, Ari Salmi, Petro Moilanen, Timo Karppinen, Zuomin Zhao, and Jussi Timonen. Phase-delayed laser diode array allows ultrasonic guided wave mode selection and tuning. *Journal of Applied Physics*, 113:144904, 2013.

[143] R K Ing, F Gires, and M Fink. Focusing and Beamsteering of Laser Generated Ultrasound. *IEEE Ultrasonics Symposium*, pages 539–544, 1989.

[144] Yves H Berthelot and Ilene J Busch-vishniac. Thermoacoustic radiation of sound by a moving laser source. *The Journal of the Acoustical Society of America*, 81(2):317–327, 1987.

[145] Kazushi Yamanaka. Selective scanning excitation of single-mode of a laser beam acoustic waves by phase velocity. *Applied Physics Letters*, 58:1591–1593, 1991.

[146] F Linnane, M Clark, and Mik Somekh. Surface acoustic wave generation with customized optical beam distributions. *IEEE Ultrasonics Symposium*, pages 479–482, 1996.

[147] Steve D Sharples, Matthew Clark, and Mike G Somekh. Spatially resolved acoustic spectroscopy for fast noncontact imaging of material microstructure. *Optics Express*, 14(22):10435–10440, 2006.

[148] M. Sharples, S. D. Clark, M. Somekh. Efficient and flexible laser ultrasound generation using spatial light modulators. *Electronics Letters*, 1(18):1–2, 2001.

[149] M Clark, S Sharples, and M Somekh. Diffractive acoustic elements for laser ultrasonics. *The Journal of the Acoustical Society of America*, 107(6):3179–85, 2000.

[150] Istvan A. Veres, Clemens M. Grunsteidl, Jurgen Roither, Peter Burgholzer, Thomas Berer, and Todd W. Murray. Direct measurement of SAW dispersion relations in the k-?? domains; Numerical and experimental studies. *IEEE International Ultrasonics Symposium, IUS*, pages 687–690, 2013.

[151] C. M. Grunsteidl, I. A. Veres, T. Berer, P. Burgholzer, T. Berer, and P. Burgholzer. Application of SLM generated patterns for laser-ultrasound. *IEEE International Ultrasonics Symposium, IUS*, pages 1360–1363, 2014.

[152] Cl. Grunsteidl, I. A. Veres, J. Roither, P. Burgholzer, T.W. Murray, and T. Berer. Spatial and temporal frequency domain laser-ultrasound applied in the direct measurement of dispersion relations of surface acoustic waves. *Applied Physics Letters*, 102(1):011103, 2013.

[153] A. Meyer, S. Gspan, S. Bernet, and M. Ritsch-Marte. Binary optoacoustic holography with a spatial light modulator. *Journal of Applied Physics*, 96(10):5886–5891, 2004.

[154] Hyoung Won Baac, Taehwa Lee, and L. Jay Guo. Micro-ultrasonic cleaving of cell clusters by laser-generated focused ultrasound and its mechanisms. *Biomedical Optics Express*, 4(8):1442, 2013.

[155] Stefan Gspan, Alex Meyer, Stefan Bernet, and Monika Ritsch-Marte. Optoacoustic generation of a helicoidal ultrasonic beam. *The Journal of the Acoustical Society of America*, 115(3):1142–1146, 2004.

[156] K. Passler, R. Nuster, S. Gratt, P. Burgholzer, and G. Paltauf. Laser-generation of ultrasonic X-waves using axicon transducers. *Applied Physics Letters*, 94(6):1–4, 2009.

[157] Weiwei Chan, Thomas Hies, and Claus Dieter Ohl. Laser-generated focused ultrasound for arbitrary waveforms. *Applied Physics Letters*, 109(17):1–7, 2016.

[158] Gianluca Memoli, Mihai Caleap, Michihiro Asakawa, Deepak R. Sahoo, Bruce W. Drinkwater, and Sriram Subramanian. Metamaterial bricks and quantization of meta-surfaces. *Nature Communications*, 8:1–8, 2017.

[159] Ye Tian, Qi Wei, Ying Cheng, and Xiaojun Liu. Acoustic holography based on composite metasurface with decoupled modulation of phase and amplitude. *Applied Physics Letters*, 110(19):191901, 2017.

[160] Yong Li, Xue Jiang, Bin Liang, Jian-chun Cheng, and Likun Zhang. Metascreen-Based Acoustic Passive Phased Array. *Physical Review Applied*, 4(2):024003, aug 2015.

[161] Yong Li, Xue Jiang, Rui Qi Li, Bin Liang, Xin Ye Zou, Lei Lei Yin, and Jian Chun Cheng. Experimental realization of full control of reflected waves with subwavelength acoustic metasurfaces. *Physical Review Applied*, 2(6):1–11, 2014.

[162] Y Takeuchi. Ultrasonic low-loss, phase continuous Fresnel lens for focal plane array imaging. *IEEE International Ultrasonics Symposium, IUS*, 2:1139–1142, 2002.

[163] Kazuyoshi Mori, Ayano Miyazaki, Hanako Ogasawara, Toshiaki Nakamura, and Yasuhito Takeuchi. Numerical Analysis of Sound Pressure Fields Focused by Phase Continuous Fresnel Lens Using Finite Difference Time Domain Method. *Japanese Journal of Applied Physics*, 46(7B):4990–4997, 2007.

[164] Michael D. Brown, Jiri Jaros, Ben T. Cox, and Bradley E. Treeby. Control of broadband optically generated ultrasound pulses using binary amplitude holograms. *The Journal of the Acoustical Society of America*, 139(4):1637–1647, 2016.

[165] Melvin H. Horman and Henry H. M. Chau. Zone Plate Theory Based on Holography. *Applied Optics*, 6(2):317, 1967.

[166] David Attwood. *Soft X-Rays and Extreme Ultraviolet Radiation: Principles and Applications*. 2007.

[167] A.D Pierce. *Acoustics: an introduction to its physical principles and applications*. 1989.

[168] Makoto Tabei, T. Douglas Mast, and Robert C. Waag. A k -space method for coupled first-order acoustic propagation equations. *The Journal of the Acoustical Society of America*, 111(1):53–63, 2002.

[169] M. J. Simpson and A. G. Michette. Imaging properties of modified fresnel zone plates. *Optica Acta*, 31(4):403–413, 1984.

[170] G. W. Webb, I. V. Minin, and O. V. Minin. Variable reference phase in diffractive antennas: Review, applications, new results. *IEEE Antennas and Propagation Magazine*, 53(2):77–94, 2011.

[171] Michael D. Brown, Thomas J. Allen, Ben T. Cox, and Bradley E. Treeby. Control of optically generated ultrasound fields using binary amplitude holograms. *IEEE International Ultrasonics Symposium, IUS*, 139(4):1037–1040, 2014.

[172] Matt Clark. *Direct-search method for the computer design of holograms*. PhD thesis, Imperial College London, 1997.

[173] Martin Meister and R. J. Winfield. Novel approaches to direct search algorithms for the design of diffractive optical elements. *Optics communications*, 203(March):39–49, 2002.

[174] Vicente Moreno, Juan Félix Román, and José Ramón Salgueiro. High efficiency diffractive lenses: Deduction of kinoform profile. *American Journal of Physics*, 65(6):556–562, 1997.

[175] Michael D. Brown, Eleanor Martin, Ben T. Cox, and Bradley E. Treeby. Single pulse illumination of multi-layer photoacoustic holograms for patterned ultrasound field generation. *IEEE International Ultrasonics Symposium, IUS*, pages 3–6, 2016.

[176] R W Young. Sound pressure in water from a source in air and vice versa. *Journal of the Acoustical Society of America*, 53(6):1708–1716, 1973.

[177] Edward Zhang, Jan Laufer, and Paul Beard. Backward-mode multiwavelength photoacoustic scanner using a planar Fabry-Perot polymer film ultrasound sensor for high-resolution three-dimensional imaging of biological tissues. *Applied Optics*, 47(4):561, 2008.

[178] Adam Aitkenhead. CONVERT_voxels_to_stl (https://uk.mathworks.com/matlabcentral/fileexchange/27733-converting-a-3d-logical-array-into-an-stl-surface-mesh), 2010.

[179] A Jesacher. *Applications of spatial light modulators for optical trapping and image processing*. PhD thesis, 2007.

[180] M. D. Brown, D. I. Nikitichev, B. E. Treeby, and B. T. Cox. Generating arbitrary ultrasound fields with tailored optoacoustic surface profiles. *Applied Physics Letters*, 110(9):1–5, 2017.

[181] Yangbo Xie, Chen Shen, Wenqi Wang, Junfei Li, Dingjie Suo, Bogdan-Ioan Popa, Yun Jing, and Steven A. Cummer. Acoustic Holographic Rendering with Two-dimensional Metamaterial-based Passive Phased Array. *Scientific Reports*, 6(1):35437, 2016.

[182] Omer Naor, Yoni Hertzberg, Esther Zemel, Eitan Kimmel, and Shy Shoham. Towards multifocal ultrasonic neural stimulation II: design considerations for

an acoustic retinal prosthesis. *Journal of Neural Engineering*, 9(2):026006, 2012.

[183] Michael D. Brown, Ben T. Cox, and Bradley E. Treeby. Investigating the effect of thickness and frequency spacing on multi-frequency acoustic kinoforms. *IEEE International Ultrasonics Symposium, IUS*, 2017.

[184] Michael Brown, Ben Cox, and Bradley Treeby. Design of multi-frequency acoustic kinoforms. *Appl Phys Lett*, 244101:1–5, 2017.

[185] David Bowman, Philip Ireland, Graham D. Bruce, and Donatella Cassettari. Multi-wavelength holography with a single spatial light modulator for ultracold atom experiments. *Optics Express*, 23(7):8365–8372, 2015.

[186] Michal Makowski, Izabela Ducin, Karol Kakarenko, Jaroslaw Suszek, Maciej Sypek, and Andrzej Kolodziejczyk. Simple holographic projection in color. *Optics Express*, 20(22):25130–25136, 2012.

[187] Michal Makowski, Maciej Sypek, Izabela Ducin, Agnieszka Fajst, Andrzej Siemion, Jaroslaw Suszek, and Andrzej Kolodziejczyk. Experimental evaluation of a full-color compact lensless holographic display. *Optics Express*, 17(23):20840–20846, 2009.

[188] Ganghun Kim, José a. Domínguez-Caballero, and Rajesh Menon. Design and analysis of multi-wavelength diffractive optics. *Optics Express*, 20(3):2814–2823, 2012.

[189] T R Sales and D H Raguin. Multiwavelength operation with thin diffractive elements. *Applied Optics*, 38(14):3012–3018, 1999.

[190] J Bengtsson. Kinoforms designed to produce different fan-out patterns for two wavelengths. *Applied Optics*, 37(11):2011–2020, 1998.

[191] Jorgen Bengtsson. Design of fan-out kinoforms in the entire scalar diffraction regime with an optimal-rotation-angle method. *Applied Optics*, 36(32):8425–8444, 1997.

[192] B Z Dong, G Q Zhang, G Z Yang, B Y Gu, S H Zheng, D H Li, Y S Chen, X M Cui, M L Chen, and H D Liu. Design and fabrication of a diffractive phase element for wavelength demultiplexing and spatial focusing simultaneously. *Applied optics*, 35(35):6859–64, 1996.

[193] María S Millan, Joaquín Otón, and Elisabet Pérez-Cabré. Chromatic compensation of programmable Fresnel lenses. *Optics express*, 14(13):6226–42, 2006.

[194] Uriel Levy, Emanuel Marom, and David Mendlovic. Thin element approximation for the analysis of blazed gratings: Simplified model and validity limits. *Optics Communications*, 229(1-6):11–21, 2004.

[195] Drew A. Pommet, M. G. Moharam, and Eric B. Grann. Limits of scalar diffraction theory for diffractive phase elements. *Journal of the Optical Society of America A*, 11(6):1827, 1994.

[196] Alexander Jesacher, Stefan Bernet, and Monika Ritsch-Marte. Colored point spread function engineering for parallel confocal microscopy. *Optics Express*, 24(24):27395–27402, 2016.

[197] Adam J Caley and Mohammad R Taghizadeh. Analysis of the effects of bias phase and wavelength choice on the design of dual-wavelength diffractive optical elements. *Journal of the Optical Society of America*, 23(1):193–198, 2006.