

# Separable Beamforming For 3-D Medical Ultrasound Imaging

Ming Yang, *Student Member, IEEE*, Richard Sampson, *Student Member, IEEE*, Siyuan Wei, *Student Member, IEEE*, Thomas F. Wenisch, *Member, IEEE*, and Chaitali Chakrabarti, *Fellow, IEEE*

**Abstract**—Three-dimensional ultrasound imaging is a promising medical imaging technology because of its ease of use and improved accuracy in diagnosis. However, its high computational complexity and resulting high power consumption has precluded its use in hand-held applications. In this paper, we present a separable beamforming method that greatly reduces computational complexity. Our method is based on decomposing the delay term in a way that minimizes the root-mean-square error caused by the decomposition. We analyze tradeoffs between the approximation error caused by the decomposition and computational complexity. Then, we present enhancements to the Sonic Millip3De hardware accelerator for ultrasound beamforming to implement separable beamforming. Using hardware synthesis targeting standard cells in 45 nm, we show that the proposed method allows us to boost the Sonic Millip3De frame rate from 1–2 Hz to 32 Hz while maintaining power consumption at 15 W. We validate image quality of our method using cyst phantom simulations in Field II. Our evaluation demonstrates that the proposed separable beamforming method can produce 3-D images with high quality that are comparable to those generated by non-separable beamforming.

**Index Terms**—Beamforming, decomposition, hardware accelerator, separable, 3-D ultrasound.

## I. INTRODUCTION

FOR decades, ultrasound imaging has been one of the most important medical imaging modalities, as it is safe and less costly than CT and MRI. Moreover, ultrasound requires little energy, which makes ultrasound imaging a good candidate for hand-held applications. Until recently, most ultrasound applications relied on 2-D imaging and required a highly trained technician to search iteratively to position the scan probe for diagnosis. This trial-and-error approach is not necessarily ideal and can reduce diagnostic accuracy and throughput because of the 3-D nature of body anatomies. 3-D ultrasound imaging has been

Manuscript received December 31, 2013; revised April 27, 2014 and August 09, 2014; accepted October 27, 2014. Date of publication November 20, 2014; date of current version December 08, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. John McAllister. This work was supported by NSF CCF-1406739, CCF-1406810, CCF-0815457, CSR-0910699, and grants from ARM Inc.

M. Yang, S. Wei and C. Chakrabarti are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: m.yang@asu.edu; siyuan.wei@asu.edu; chaitali@asu.edu).

R. Sampson and T. F. Wenisch are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 (e-mail: rsamp@umich.edu; twenisch@umich.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2014.2371772

TABLE I  
SYSTEM PARAMETERS

Property	Value
Pitch, $\mu\text{m}$	192.5
Array size, element	$120 \times 88$
Subaperture size, element	$32 \times 32$
Number of scanlines	$48 \times 48$
View angle, square degree	$45^\circ \times 45^\circ$
Max depth, cm	10
Center frequency, MHz	4
6dB transducer bandwidth, MHz	2
A/D sampling rate, MHz	40

shown to improve system ease of use, facilitating higher diagnostic accuracy and patient throughput [1], [2].

However, the benefits of 3-D ultrasound imaging come with challenges that are exacerbated when trying to make the system hand-held. A 2-D transducer array for a typical 3-D ultrasound system contains over ten thousand transducers, compared to hundreds of transducers in a typical 2-D system. The number of scanlines required in a 3-D volume is also at least one order of magnitude higher than in a 2-D image. The resulting computational complexity of 3-D ultrasound image formation is thus hundreds of times larger than its 2-D counterpart. So, it is no surprise that high quality 3-D imaging is currently available only on a few bulky, ultrasound machines.

Over the past few years, several specialized hardware platforms have been proposed for 2-D and 3-D ultrasound imaging. 2-D imaging architectures include a specialized DSP architecture for color Doppler [3], an FPGA architecture for beamforming [4] and GPGPU implementations for sonoelastography and color Doppler [5], [6]. 3-D imaging systems have been implemented using DSPs [7], computing clusters [8] and multi-FPGA systems [9], [10]. All of these systems target conventional imaging platforms and have large power budgets; none is suitable for a hand-held system.

In our earlier work, we proposed the Sonic Millip3De architecture, a 3-D die-stacked accelerator for the digital front end of a 3-D synthetic aperture imaging system that can produce high fidelity 3-D images at up to 1–2 frames per second with a power budget of 15 W in 45 nm technology [11], [12]. Though promising, the existing design has a frame rate that is too low for some applications, and falls short of the 5 W target for safe contact of the probe head with human skin at 45 nm technology node. For a typical imaging configuration (shown in Table I), the number of delay-sum operations in beamforming is on the order of  $10^{12}$  per frame, consuming 60% of Sonic Millip3De's total power. Thus reducing the beamforming complexity would help

in both increasing the frame rate and/or decreasing the power consumption.

To support 3-D ultrasound imaging with reduced computational complexity, researchers have proposed 2-D sparse array designs [13]–[15]. Although sparse-array based-systems employ fewer transducer elements than traditional 2-D arrays, such systems typically have high sidelobe levels and suffer from low SNR. Another technique, preliminary beamforming for reconfigurable arrays [16], [17], reduces the number of A/D converters and the complexity of the digital front end. However this approach has not yet been extended to support 3-D synthetic aperture ultrasound imaging.

Alternatively, the computational complexity of the digital front end can be greatly reduced without compromising the image quality by a promising method called separable beamforming. Separable beamforming reduces computational complexity by decomposing 2-D array beamforming into a series of 1-D array operations. This method was first introduced for frequency domain beamforming in sonar applications [18], [19] and later to ultrasound imaging [8]. Although amenable to parallelization [8], prior delay decomposition methods incur large error for small f-number, resulting in poor resolution. A later separable beamforming proposal [20] improves delay decomposition error at the cost of only using vertical scanlines, limiting the approach to plane-wave systems with rectangular views that are too narrow to image large anatomies.

In an earlier, short version of this article [21], we proposed a generalized delay decomposition method that is not restricted to any transmit scheme or scan geometry. Our approach designs the beamforming decomposition to minimize root-mean-square (RMS) error. When applied to subaperture-based synthetic aperture ultrasound, as in Sonic Millip3De, this new decomposition reduces computational complexity by 19 $\times$  over the non-separable baseline. The reduction in computation can be leveraged either to reduce average system power or to increase frame rate within a fixed power budget. We pursue the latter in our hardware design and evaluation, in which we achieve 32 frames per second compared to 1–2 frames per second in [12] with 15 W power at 45 nm technology node, projected to scale to 5 W power at 16 nm technology node.

In this article, we present a more detailed study of separable beamforming. Our key contributions include:

- We investigate key factors that affect phase error due to delay decomposition, and propose a method to find the right decomposition form and ordering (row first versus column first) that minimizes phase error.
- We present the trade-offs between RMS phase error and computational complexity. We use this analysis to select subaperture size.
- We validate the image quality of our proposed separable beamforming system through simulated beamforming of several phantoms including upright and angled cyst phantoms.
- We present extensions to the Sonic Millip3De hardware accelerator to support the proposed separable beamforming method. Our analysis includes detailed power and synthesis results targeting a 45 nm standard cell library.

The rest of the paper is organized as follows. In Section II, we present our separable beamforming delay decomposition method. We present the hardware architecture in Section III. In Section IV, we present our image quality and power analysis. Section V concludes the paper.

## II. SEPARABLE BEAMFORMING

### A. Why Separable Beamforming?

Consider a 3-D ultrasound system described by the configuration shown in Table I. To generate a 3-D image of size 10 cm  $\times$  45 $^\circ$   $\times$  45 $^\circ$ ,  $1.1 \times 10^{12}$  delay-sum operations must be performed per frame. Such high computational complexity results in correspondingly high power consumption and thus limits hand-held applications of 3-D ultrasound.

Separable beamforming can significantly reduce computational complexity by decomposing 2-D array beamforming into two stages of 1-D beamforming. In the first stage, the beams are computed along one direction (e.g., horizontal), and in the second stage the beams are computed along the other direction (e.g., vertical). The image quality of separable beamforming depends on the accuracy of the beamforming delay decomposition. In this section, we describe a separable beamforming method that minimizes delay decomposition error.

### B. Separable Beamforming Process

We present our separable beamforming method in the context of a system that uses subaperture transmit and receive, meaning that only a subset of transducers (e.g., 32  $\times$  32, each referred to as a subaperture) are used in each firing. In each firing, the active subaperture shifts by 8 transducer elements, partially overlapping the preceding subaperture. This process repeats until the whole aperture (e.g., 120  $\times$  88 transducer array) is traversed.

Without loss of generality, we assume the 3-D coordinate system shown in Fig. 1. Let  $(R, \theta, \phi)$  be the coordinates of a focal point  $P$ . Here  $R$  is the radial distance from the origin  $O$  to point  $P$ . Point  $P'$  is the orthogonal projection of  $P$  in the  $yz$  plane.  $\phi$  is the elevation angle between line  $OP'$  and the  $z$  axis.  $\theta$  is the azimuth angle between  $OP$  and its orthogonal projection  $OP'$  in the  $yz$  plane. For a transducer array element at  $(x, y, 0)$ , the distance between the transducer element and the focal point  $P$  is given by

$$d_{tx} = \sqrt{R^2 + x^2 - 2Rx \sin(\theta) + y^2 - 2Ry \cos(\theta) \sin(\phi)} \quad (1)$$

Similarly the distance between the firing virtual source located at  $(x_v, y_v, z_v)$  and the focal point  $P$  is given by

$$d_{tx} = \left( R^2 + x_v^2 + y_v^2 + z_v^2 - 2x_v R \sin \theta - 2Ry_v \cos \theta \sin \phi - 2Rz_v \cos \theta \cos \phi \right)^{\frac{1}{2}} \quad (2)$$

Assuming that the ultrasound speed is  $c$ , and the round-trip delay between the origin and the focal point is  $2R/c$ , the round-trip delay at the transducer relative to that at the origin is given by

$$\tau(x, y, R, \theta, \phi) = \frac{(2R - d_{tx} - d_{rx})}{c} \quad (3)$$

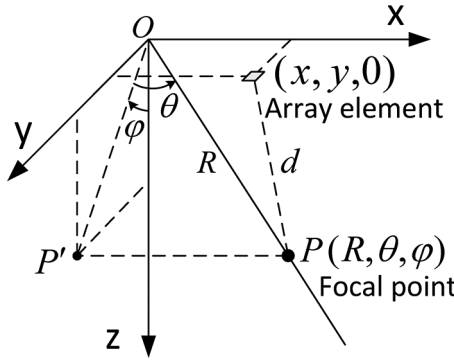


Fig. 1. 3-D coordinate system used in the beamforming formulation.

Let  $\tau(n_x, n_y, m_R, m_\theta, m_\phi)$  be the discrete form of  $\tau(x, y, R, \theta, \phi)$ , where  $n_x$  and  $n_y$  are variables associated with the coordinates of receive elements, and  $m_R, m_\theta$  and  $m_\phi$  are variables associated with the coordinates of focal points. Then the non-separable beamforming corresponding to subaperture  $l$  whose left corner indices are  $i_l$  and  $j_l$ , is described as

$$F_l(m_R, m_\theta, m_\phi; t) = \sum_{n_x=i_l}^{i_l+N_x-1} \sum_{n_y=j_l}^{j_l+N_y-1} A_l(n_x, n_y) \cdot S_l(n_x, n_y, t - \tau(n_x, n_y, m_R, m_\theta, m_\phi)) \quad (4)$$

where  $S_l(n_x, n_y, t)$  is the signal received by transducer element  $(n_x, n_y)$  at  $l$ th firing and  $A_l(n_x, n_y)$  is the corresponding apodization coefficient.  $F_l(m_R, m_\theta, m_\phi; t)$  is the low resolution 3-D image generated by subaperture  $l$ , which is sampled at  $t = 2R/c$  for dynamic focusing. For a synthetic aperture ultrasound system, the final high resolution image is obtained by summing all the low resolution images from all subapertures.

Now, if  $\tau(n_x, n_y, m_R, m_\theta, m_\phi)$  can be decomposed as

$$\tau(n_x, n_y, m_R, m_\theta, m_\phi) = \tau_1(n_x, n_y, m_R, m_\theta) + \tau_2(n_y, m_R, m_\theta, m_\phi) \quad (5)$$

then (4) can be represented by a two-stage separable beamforming process:

$$F_l^{(1)}(n_y, m_R, m_\theta; t) = \sum_{n_x=i_l}^{i_l+N_x-1} A_l(n_x, n_y) S_l(n_x, n_y, t - \tau_1(n_x, n_y, m_R, m_\theta)) \quad (6)$$

$$F_l^{(2)}(m_R, m_\theta, m_\phi; t) = \sum_{n_y=j_l}^{j_l+N_y-1} F_l^{(1)}(n_y, m_R, m_\theta; t - \tau_2(n_y, m_R, m_\theta, m_\phi)) \quad (7)$$

In the first stage, the beamforming is along the  $x$  axis, which functions as a spatial filter that steers the receive plane to azimuth angle  $\theta$ . The process repeats for all combinations of  $m_R, n_y$  and  $m_\theta$  and results in a partially beamformed intermediate signal  $F_l^{(1)}$ . In the second stage, 1-D beamforming is performed along the  $y$  axis, and corresponds to steering receive plane to elevation angle  $\phi$ . The second stage beamforming is repeated for




Fig. 2. The principle of separable beamforming.

all combinations of  $m_R, m_\theta$  and  $m_\phi$ . The principle of the proposed separable beamforming method is shown in Fig. 2.

The number of delay-sum operations of separable beamforming for one subaperture is  $N_x N_y M_R M_\theta + N_y M_R M_\theta M_\phi$  in contrast to  $N_x N_y M_R M_\theta M_\phi$  in conventional, non-separable beamforming. Thus, the computational complexity reduction is  $N_x M_\phi / (N_x + M_\phi)$ . For the configuration shown in Table I with a  $32 \times 32$  subaperture size and  $48 \times 48$  scanlines, our approach achieves about  $19\times$  complexity reduction.

The separable beamforming method is based on the assumption that the wave propagation delay  $\tau$  can be decomposed into  $\tau_1$  and  $\tau_2$ . However the decomposition is not exact and its effectiveness depends on the accuracy of the  $\tau_1$  and  $\tau_2$  approximations. Next, we describe the proposed delay decomposition method, which designs  $\tau_1$  and  $\tau_2$  to minimize RMS phase error.

### C. Delay Decomposition Method

The beamsum delay  $\tau$ , which is a function of five variables, cannot be strictly decomposed into a sum of two functions with fewer variables because the distance calculation involves a square root operation (as shown in (1) and (2)). The Taylor series of the square root includes functions involving multiple variables, also referred to as cross terms, which can not be easily decomposed. To make the delay equation separable, some of these cross terms must be dropped. Although the effect of the cross terms diminish with large  $R$ , for small depths, the cross terms in the delay calculation can be significant. There are several factors that affect the error incurred by the delay decomposition: the number of variables in each decomposed function, the combination of variables, and the sequence of computations in the two-stage beamforming.

To build foundation for our proposed decomposition, we first consider a simpler strawman wherein  $\tau_1$  and  $\tau_2$  are each functions of three variables and the delay decomposition is given by

$$\tau(x, y, R, \theta, \phi) \approx \tau_1(x, R, \theta) + \tau_2(y, R, \phi) \quad (8)$$

We justify this first decomposition as follows. For dynamic focusing, both  $\tau_1$  and  $\tau_2$  depend on  $R$ , thus  $R$  should be included in the variable list of both  $\tau_1$  and  $\tau_2$ .  $\tau_1$  is also a function of  $\theta$  and  $x$  because 1-D beamforming along the  $x$  direction allows the array system to filter out signals from all azimuth directions except those with azimuth angle  $\theta$ . Hence  $\theta$  and  $x$  should not be separated. Similarly, 1-D beamforming along the  $y$  direction allows the array system to filter out signals from all elevation directions except those whose elevation angle is  $\phi$ , so  $\phi$  and  $y$  should not be separated. In this decomposition, since  $(x, \theta)$  and  $(y, \phi)$  are interchangeable, either  $\tau_1$  or  $\tau_2$  can be used in the first stage of beamforming. Unfortunately, this simple decomposition approach leads to large errors primarily because  $\theta$  and  $\phi$  are separated and the cross terms involving  $\theta$  and  $\phi$  are lost.

To improve the approximation,  $\tau_1$  and  $\tau_2$  must capture additional important dependencies. To simplify our presentation, let us assume that we perform first-stage beamforming along the  $x$  axis while the second stage is along the  $y$  axis. For  $\tau_1$ , consider adding  $\phi$  or  $y$  to the variable list already consisting of  $R$ ,  $x$  and  $\theta$ . There is no benefit in adding  $\phi$  because 1-D beamforming in the first stage is along the  $x$  direction, and does not have enough resolution along  $\phi$ . However there is an advantage of adding  $y$  to the variable list of  $\tau_1$ , because  $\tau_1$  already includes  $x$  and thus adding  $y$  helps to retain cross terms between  $x$  and  $y$ . For  $\tau_2$ , there are two candidate variables that can be added to the variable list, namely  $x$  and  $\phi$ . There is no benefit of adding  $x$  to  $\tau_2$ 's variable list because neither the input signal of second-stage beamforming  $F^{(1)}$  nor output of second stage beamforming  $F^{(2)}$  relates to  $x$ . However, adding  $\theta$  to  $\tau_2$  has the advantage of preserving cross terms involving  $\theta$  and  $\phi$ . Thus, both  $\tau_1$  and  $\tau_2$  are best represented as functions of four variables and our final decomposition is of the form

$$\tau(x, y, R, \theta, \phi) \approx \tau_1(x, y, R, \theta) + \tau_2(y, R, \theta, \phi) \quad (9)$$

Alternatively, if the first beamforming stage is performed along the  $y$  axis, the candidate decomposition is instead

$$\tau(x, y, R, \theta, \phi) \approx \tau_1(x, y, R, \phi) + \tau_2(x, R, \theta, \phi) \quad (10)$$

Compared to the decomposition in (9), here  $\tau_1$  includes  $\phi$  because beamforming along the  $y$  axis allows the system to distinguish signals coming from different elevation angles  $\phi$ .

Fig. 3 compares the RMS phase errors for the three alternative decompositions defined in (8), (9) and (10). Each decomposition is obtained by setting up an RMSE minimization problem and is solved using the Euler-Lagrange equation. By increasing the number of variables of  $\tau_1$  and  $\tau_2$  from three to four, the approximation error is reduced by one decade for relatively large depth. For small depths, the error of three-variable decomposition increases rapidly, while the error of four-variable decomposition approaches zero. This behavior occurs because, as  $R$  becomes very small, the cross terms involving  $x$  and  $y$  become significant, and both the four-variable decomposition forms retain these cross terms in  $\tau_1$ .

Also note that beamforming along  $x$  first results in lower RMSE. We find this ordering is better because the RMSE of approximation (9) is primarily due to dropping cross terms involving  $x$  and  $\phi$ , which is less than the RMSE caused by drop-

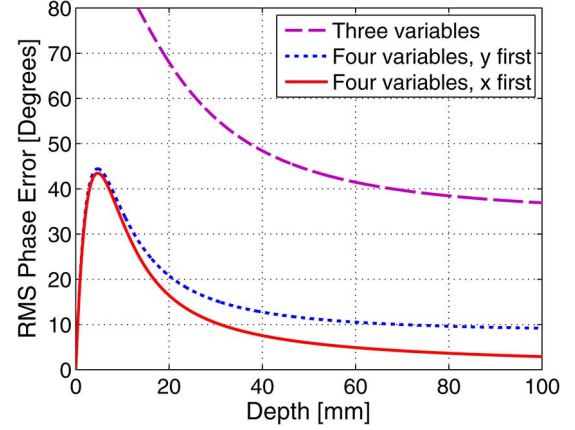


Fig. 3. RMS phase error as a function of depth for different types of decomposition corresponding to three-variable decomposition defined in (8), four-variable decomposition with first-stage beamforming in  $x$  dimension defined in (9), and four-variable decomposition with first-stage beamforming in  $y$  dimension defined in (10).

ping cross terms involving  $y$  and  $\theta$  in approximation (10). The cross term involving  $y$  and  $\theta$  is an artifact of the coordinate system; the distance calculation ((1)) contains a natural cross term involving  $y$  and  $\theta$ , namely  $2Ry \cos(\theta) \sin(\phi)$ . The proposed decomposition ((9)) results in delay error under  $9^\circ$  when depth is larger than 3 cm, and reduces rapidly as depth increases.

#### D. Generating Delay Functions $\tau_1$ and $\tau_2$

Given this decomposition, we must next generate  $\tau_1(x, y, R, \theta)$  and  $\tau_2(y, R, \theta, \phi)$  such that the error due to approximation is minimized. Minimizing RMSE is equivalent to minimizing

$$E = \int_{\phi_1}^{\phi_2} \int_{x_1}^{x_2} [\tau(x, y, R, \theta, \phi) - (\tau_1(x, y, R, \theta) + \tau_2(y, R, \theta, \phi))]^2 dx d\phi \quad (11)$$

where  $x_1$  and  $x_2$  are the lateral boundary of the corresponding subaperture, and  $\phi_1$  and  $\phi_2$  are the range of the elevation angle—a classic calculus of variation problem. By solving the corresponding Euler-Lagrange equations [22] and discretizing  $\tau_1$  and  $\tau_2$ , we get the following solution that minimizes RMSE:

$$\begin{aligned} \tau_1(n_x, n_y, m_R, m_\theta) &= \frac{1}{M_\phi} \sum_{m_\phi=1}^{M_\phi} \tau(n_x, n_y, m_R, m_\theta, m_\phi) - \rho(n_y, m_R, m_\theta) \end{aligned} \quad (12)$$

$$\begin{aligned} \tau_2(n_y, m_R, m_\theta, m_\phi) &= \frac{1}{N_x} \sum_{n_x=i_l}^{i_l+N_x-1} \tau(n_x, n_y, m_R, m_\theta, m_\phi) - \rho(n_y, m_R, m_\theta) \end{aligned} \quad (13)$$

$$\begin{aligned} \rho(n_y, m_R, m_\theta) &= \frac{1}{2N_x M_\phi} \sum_{n_x=i_l}^{i_l+N_x-1} \sum_{m_\phi=1}^{M_\phi} \tau(n_x, n_y, m_R, m_\theta, m_\phi) \end{aligned} \quad (14)$$

Note that (12) and (13) are not the only form that minimizes RMSE. Assuming we add an arbitrary term  $\xi(n_y, m_R, m_\theta)$  on the right-hand side of (12) while subtracting it from the right-hand side of (13), the summation of  $\tau_1$  and  $\tau_2$  and the RMSE both remain the same. In this work, we choose  $\xi(n_y, m_R, m_\theta) = 0$  so that the means of  $\tau_1$  and  $\tau_2$  are the same. From an architectural perspective, this formulation makes the delay line length or buffer depth roughly equal in the two beamforming stages.

### E. Online Iterative Separable Delay Calculation

Next we focus on efficient calculation of  $\tau_1$  and  $\tau_2$ . Equations (12) and (13) are straight forward, and can be used to generate look-up tables for  $\tau_1$  and  $\tau_2$ . However, storing  $\tau_1$  and  $\tau_2$  as look-up tables is not practical due to the required look-up table size. For our system configuration, look-up tables of  $\tau_1$  and  $\tau_2$  for 96 subapertures include at least (considering symmetry) 5.7 billion and 8.9 billion constants, respectively.

Fortunately, the delay values of consecutive samples on a scanline do not change much. Hence, it is possible to iteratively calculate the delay value for the  $i$ th focal point on a scanline from the delay value of the  $(i - 1)$ th focal point. We use piece-wise quadratic curves to approximate the delay difference between consecutive samples along a scanline. For example, let  $\tilde{\tau}(m_R)$  be the delay corresponding to the  $m_R$ th focal point for fixed  $n_x$  and  $n_y$ . Let  $\eta(m_R) = \tilde{\tau}(m_R + 1) - \tilde{\tau}(m_R)$ , then  $\eta(m_R)$  can be approximated by  $am_R^2 + bm_R + c$ . Instead of storing the delay look-up table directly, the coefficients  $a$ ,  $b$  and  $c$  and the initial delay are stored, and the delays are iteratively calculated using these coefficients. The iterative calculation method does not require multiplications, it can be implemented in a simple circuit using only three additions. We had employed a similar iterative delay calculation for non-separable beamforming in [11].

To get an accurate approximation, each scanline is divided into 2–4 sections and the delay in each section is approximated by a quadratic curve. For our system configuration, where the depth ranges from 2 cm to 10 cm, we cannot use a 2 section configuration since it results in significantly large approximation error. We choose a 3 section configuration over a 4 section configuration since it requires 23% lower storage with comparable approximation error.

The storage requirements of this method are as follows. Each section is characterized by three constants and an initial point, and each scanline requires an additional start index. Thus, each scanline requires 13 constants. A total of 38 M constants must be stored; 15 M constants are required for  $\tau_1$  and the remaining 23 M for  $\tau_2$ . The 15 M constants for  $\tau_1$  correspond to 13 constants/scanline  $\times$  48 scanlines  $\times$  1,024 transducers/subaperture  $\times$  96 subapertures, divided by 4 due to symmetry (the delay term is symmetric in both  $x$  dimension and  $y$  dimension and so it is sufficient to store only 1/4 of the constants). The number of constants for  $\tau_2$  is calculated in a similar way. Each constant requires 12 bits on average, resulting an overall storage requirement of 55 MB.

Fig. 4 shows simulation results of this iterative delay calculation method with double-precision floating-point and 12-bit

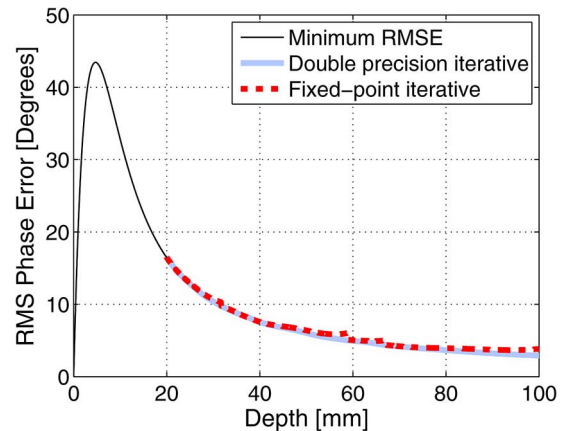


Fig. 4. RMS phase error as a function of depth when using iterative delay calculation method with both floating-point double-precision and 12-bit fixed-point precision.

fixed-point precision. The iterative method with double precision floating-point coincides with minimum RMSE curve, and fixed-point approximation only slightly increases RMSE.

## III. HARDWARE ACCELERATOR

### A. System Architecture Overview

We implement separable beamforming as an extension to our existing Sonic Millip3De beamforming accelerator. Sonic Millip3De [11] is a system architecture and accelerator for 3-D ultrasound that combines numerous hardware design techniques to minimize power while simultaneously generating high fidelity 3-D images. The full system comprises three distinct die layers that are stacked vertically using modern 3-D die-stacking techniques and are connected with through-silicon-vias (TSVs) [23]. This unique hardware layout allows for a dense, highly parallel design that can be easily integrated directly into the ultrasound scanhead, performing all front-end computation locally in a hand-held wand without the need for a large external system. We focus our design on the computation and data-intensive front-end of ultrasound image formation; presentation of a projection/view of the image or further post-processing can then be done in a separate GPU-based system integrated into a tablet computer.

The Sonic Millip3De system is split across its three distinct die layers as shown in Fig. 5. The first of these layers contains a  $120 \times 88$  grid of capacitive-micromachined ultrasonic transducers (CMUTs) and associated analog support circuitry [24]. Using our sub-aperture scheme for transmit and receive, these transducers are multiplexed into 1,024 output channels, which are then fed into the second layer. This process allows us to reduce the necessary hardware of the logic layers to only support a single sub-aperture as only 1,024 transducers are active at a time.

Sonic Millip3De's second layer is made up of ADCs and SRAM storage. For each of the 1,024 processing channels, there is a 12-bit ADC as well as a 6 kB SRAM array to store the digital signal during the first stage of beamforming. Additionally, this layer has a secondary set of 1,536 6 kB SRAM arrays which are

used to store partially beamformed data for the second stage of separable-beamforming with data being fed back to this storage from the accelerator for the second pass. This second set of arrays is necessary to prevent the original echo data from being overwritten during beamforming as it is reused over a series of scanlines.

The final layer is the beamforming accelerator, which reads echo data from the SRAM arrays and generates beamformed output. To perform separable beamforming, data must pass through this layer twice, once for each partial beamforming operation. In the following section, we provide a more complete description of this layer and its operation.

### B. Beamforming Accelerator

The beamforming accelerator is the central component of Sonic Millip3De, combining massive parallelism with a hardware-efficient implementation of the piecewise quadratic approach to delay estimation. The accelerator comprises 1,024 parallel processing channels, which each read data from separate input channels and process 16 scanlines at a time. Each of these channels is further broken into a three-unit pipeline, which translates raw echo data stored in the SRAM layer into the beamformed data for the image (Fig. 5). During the first stage of separable beamforming, partial beamforming is performed within 32-channel clusters that perform a summation within the cluster and write partially beamformed data back to secondary SRAM storage in the second (memory) layer. The partially beamformed data is then fed through the accelerator a second time, where it is again delay-aligned and summed across all 1,024 channels to generate the final image. The image is then written to external memory via a ARM Cortex M-3 control processor.

As noted, each beamforming channel comprises three units. The first unit (interpolation unit) reads echo data from the SRAM storage and applies a pre-loaded channel-specific constant apodization to the signal. The apodization weights the channel's impact on the final image based on the corresponding transducer's position in the sub-aperture. After apodization, this unit then performs a  $4\times$  linear interpolation to up-sample the signal from 40 MHz to 160 MHz, a common optimization in existing commercial designs to reduce the ADC sampling frequency.

Next, the expanded data is streamed into the next unit for the beamformation process to begin. The interpolated signal is transferred from the interpolation unit to the select unit. The select unit iteratively calculates the delays between consecutive focal points along a scanline and identifies the interpolated sample that most closely corresponds to the focal point (i.e., it selects the sample from its channel nearest to each focal point). The select unit operates in parallel on 16 scanlines. 16 sub-units iterate over the interpolated data in a block-synchronized fashion each aligning the input signal to its assigned scanline. As described previously, the iterative delay calculation algorithm determines how many samples to advance an input channel to arrive at the sample nearest a focal point using our piecewise quadratic delay estimation formula. The hardware is easily able to estimate the delta between selected samples using

three adders and the pre-computed quadratic constants, thereby iteratively solving the quadratic equation and producing each estimated delta as needed. Using these estimates, the sub-units know how far along the data stream to iterate before selecting their next output value. The delay-adjusted scanline data for the 16 neighboring scanlines is then fed forward to the summation network.

The final unit of each channel sums partially beamformed data across the channels though the use of adders connected via a reconfigurable mesh network. The network is reconfigured between beamforming operations to connect adders into a pipeline appropriate to the necessary summation operation: within clusters of 32 channels in the first beamforming stage, and across clusters in the second beamforming stage. The reconfigurability of the summation network is one of the key changes required over the baseline Sonic Millip3De design to enable separable beamforming.

The output of the summation network is written either to secondary SRAM arrays on the memory layer (for the first beamforming stage), or are passed to an ARM Cortex M-3 control processor to write final image data to external memory. (Fig. 6) illustrates the reconfigurable network and the data flow in the two beamforming stages.

## IV. RESULTS

### A. Methodology

We evaluate image quality through simulated beamforming of cyst phantoms using Field II [25], [26] and MATLAB. The simulation parameters are listed in Table I. The system employs a 2-D transducer array comprising  $120 \times 88$  transducer elements with a central frequency of 4 MHz and 50% fractional bandwidth. The scan view is  $45^\circ$  in both elevation and azimuth angles. The maximum depth of view is 10 cm.

In [12], we proposed a firing method based on overlapping subapertures with each virtual source corresponding to a sub-aperture, a generalization of the 1-D sub-aperture scheme in [27]. There are 96 overlapping subapertures each of size  $32 \times 32$  to cover the whole array of size  $120 \times 88$ . The use of virtual sources helps improve the SNR. When a virtual source fires, the corresponding subaperture receives and a 3-D sub-image is formed. After 96 firings, a complete 3-D image is created. The apodization coefficients for each subaperture are generated according to the equivalent aperture method [28]. We decompose the 2-D apodization windows into products of 1-D apodization windows and optimize those 1-D windows with respect to SNR and sidelobe levels.

### B. Analysis of Delay Decomposition Error

The delay decomposition error arises due to dropping cross terms in our delay decomposition approximation. We simulate various configurations using MATLAB, and analyze the RMS phase errors for different subaperture sizes and wide view angles.

As shown in Fig. 8, RMS phase error is approximately proportional to subaperture size, but the error decreases rapidly as

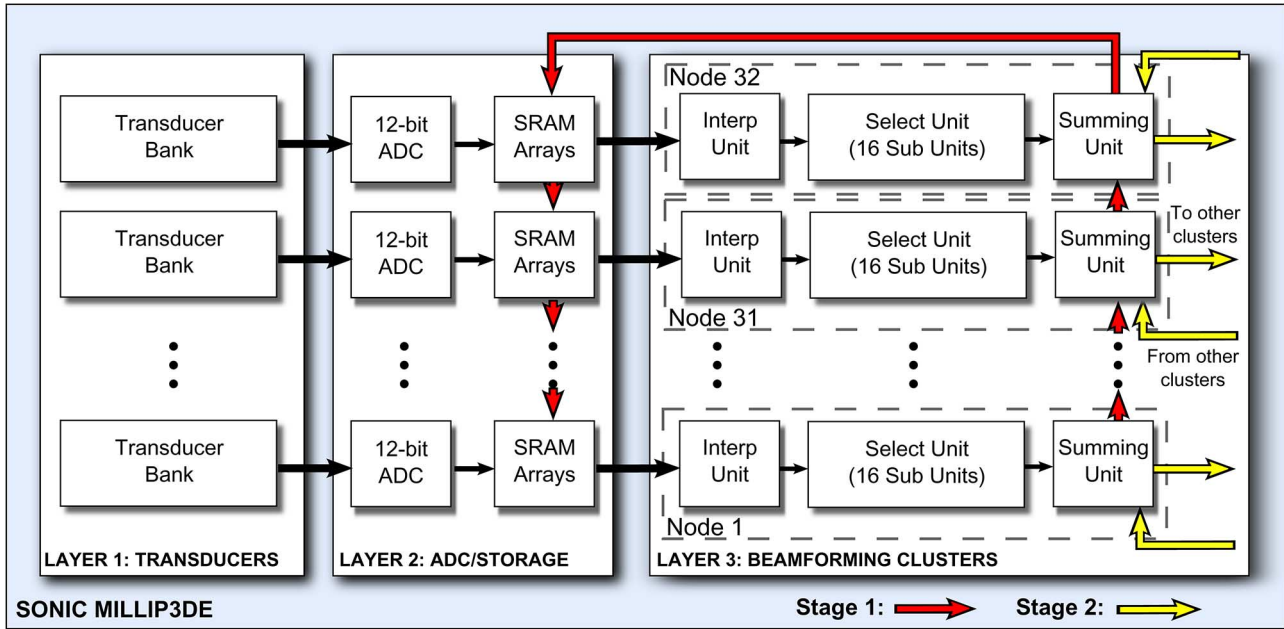


Fig. 5. **Sonic Millip3De Hardware Overview.** The full hardware design is laid out over three distinct die layers connected vertically via TSVs. Layer 1 ( $24 \times 18$  mm) comprises  $120 \times 88$  transducers, with the analog transducer outputs multiplexed for each sub-aperture and routed over TSVs to Layer 2, comprising 1024 12-bit ADC units operating at 40 MHz and SRAM arrays to store incoming samples. Data buffered in the SRAMs are transferred via face-to-face links to Layer 3 for processing in one of the 1024 3-unit pipelines of the beamsun accelerator. The interpolation unit upsamples the signal to 160 MHz and performs apodization. The select unit maps signal data from the receive time domain to the image space domain. The summing unit combines the data across the 32 channels belonging to a particular cluster to construct the partial beamsum. The partially beam-formed data is transferred back to the SRAM layer to store until the second beamforming stage. In the second stage, the data are again sent through beamforming accelerator; however, the summing units are reconfigured to sum across all 1024 pipelines, arriving at the final 3-D image, which is then written to external memory.

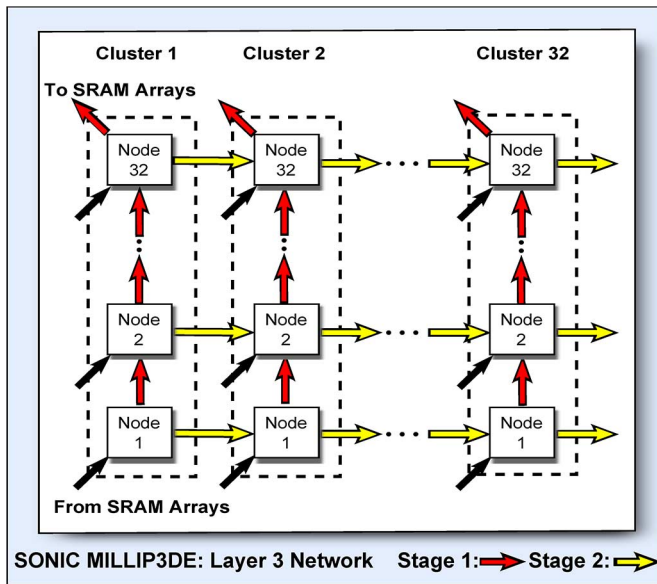


Fig. 6. Network data flow in stages 1 & 2 of separable beamforming: The beamforming accelerator units in a dashed box form a cluster with the black arrows corresponding to data flow from SRAM arrays to nodes in a cluster. In the 1st beamforming stage the data is summed from bottom to top and is written back to the secondary SRAM arrays. During the 2nd stage, the network is re-configured so that summation occurs from left to right, after which the fully beamformed data is stored in the DRAM.




Fig. 7. Overlapped subaperture processing for SAU system: 96 overlapping subapertures of size  $32 \times 32$  to cover an array of size  $120 \times 88$ .

depth increases. For depths larger than 2 cm, even for a subaperture size of  $64 \times 64$ , the error drops below  $28^\circ$ .

Although smaller subapertures result in less delay error, a smaller subaperture implies more firings to traverse the same

transducer array, and correspondingly more computations, when compared to a larger subaperture. For instance a  $(16 \times 16)$  configuration requires  $4 \times$  more firings, resulting in a doubling of computational complexity and cutting the peak frame rate (due to transmit limits) by  $4 \times$ . On the other hand, the larger  $(64 \times 64)$  configuration quadruples on-chip storage requirements and increases the number of parallel pipelines in the accelerator, increasing area requirements. The trade-off




Fig. 8. RMS phase error as a function of depth for different subaperture sizes; the errors are proportional to subaperture size.

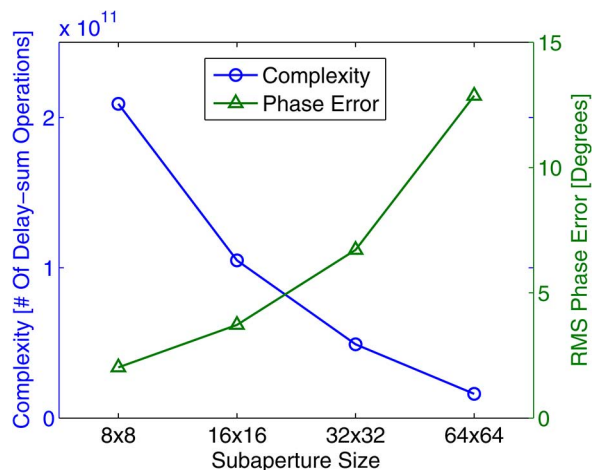


Fig. 9. Complexity and RMS phase error as a function of subaperture size. As subaperture size increases, the RMS phase error increases, while the computational complexity decreases.

between delay error and beamforming complexity is shown in Fig. 9. The beamforming complexity in this figure is reported as the number of delay-sum operations required to produce a complete frame with multiple firings. The RMSE is averaged over depths from 2 cm to 10 cm, for a  $45^\circ \times 45^\circ$  angle view. We see that the  $(32 \times 32)$  configuration has both low RMSE ( $< 6^\circ$ ) and low complexity, so we choose this configuration for our system.

Next, we present the RMS phase error map across the  $120 \times 88$  transducer array for subaperture size of  $32 \times 32$ . The RMS errors are averaged across different depths (2–10 cm), different angles ( $45^\circ \times 45^\circ$  field of view) and 96 subapertures. The result is shown in Fig. 10(a). The RMS phase errors are mostly below  $20^\circ$ ; the worst case is at the corners where RMS phase error reaches  $30^\circ$ .

We also present the RMS phase errors for different angles. Although our proposed system is for  $45^\circ \times 45^\circ$  field of view, in order to study the performance of the separable beamforming method for larger angles, we extend the RMS phase error map to  $90^\circ \times 90^\circ$ . The RMS phase errors are averaged across different depths (2–10 cm), different transducers and different subaper-




Fig. 10. RMS phase error maps for subaperture size of  $32 \times 32$ . The highest phase error arises in the four corners of the transducer array and scanlines with large azimuth angle and large elevation angle. (a) RMS phase error map as a function of transducer element position; (b) RMS phase error map as a function of scanline angle.

tures, and presented in Fig. 10(b). The highest error arises in the four corners where both  $\theta$  and  $\phi$  are large.

Ideally, the phase error of our configuration (4 MHz central frequency and 160 MHz sampling rate after  $4 \times$  interpolation), should be  $\pm 9$  degrees. In our system this can be achieved when the depth is larger than 3 cm or the elevation angle is within  $\pm 10$  degrees. Our simulation results for different depths and different angles show that the overall image quality of the separable system is good and comparable to that of the non-separable system, and are presented in Section IV-C.

### C. Separable Beamforming

We use Field II [25], [26] and MATLAB to simulate a 3-D imaging system with the parameters listed in Table I. We consider two simulation cases, as illustrated in Fig. 11. Both cases have twelve anechoic cysts located in a  $20 \text{ mm} \times 15 \text{ mm} \times 80 \text{ mm}$  volume of random scatterers. The diameters of the cysts range from 2 mm to 7 mm. In Case A (Fig. 11(a)), the volume containing cysts and scatterers is vertical, corresponding to  $\theta = \phi = 0^\circ$ . In Case B (Fig. 11(b)), the volume containing cysts and scatterers, with cysts located at  $\theta = \phi = 30^\circ$ . Consequently, the field of scan view is increased from  $45^\circ \times 45^\circ$  to  $90^\circ \times 90^\circ$ , and the number of scanlines is increased from  $48 \times 48$  to  $96 \times 96$ .



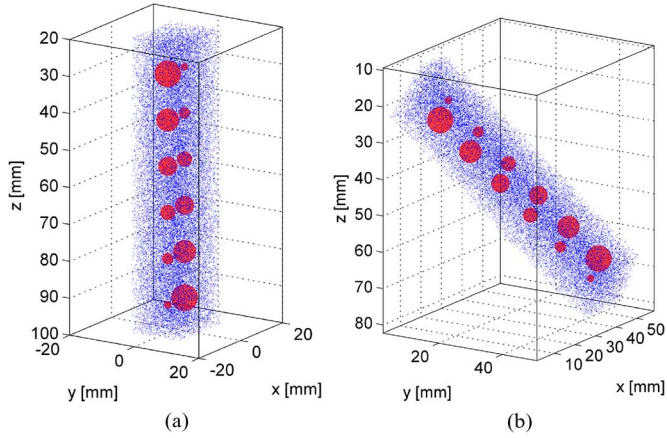


Fig. 11. The scatterer and cyst distribution for two Field II simulation cases; 12 anechoic cysts with diameter ranges from 2 mm to 7 mm. (a) Case A: Scatterers and cysts in upright position; (b) Case B: Scatterers and cysts swung to  $\theta = \phi = 30^\circ$ .

We quantify image quality via Contrast-to-Noise Ratio (CNR) and Contrast Ratio (CR). The CNR and CR are defined as follows

$$\text{CNR} = \frac{|\mu_{\text{cyst}} - \mu_{\text{bgnd}}|}{\sqrt{\sigma_{\text{cyst}}^2 + \sigma_{\text{bgnd}}^2}} \quad (15)$$

$$\text{CR} = \frac{\mu_{\text{bgnd}} - \mu_{\text{cyst}}}{\mu_{\text{bgnd}} + \mu_{\text{cyst}}} \quad (16)$$

where  $\mu_{\text{cyst}}$  and  $\mu_{\text{bgnd}}$  correspond to mean brightness of cyst and background, and  $\sigma_{\text{cyst}}$  and  $\sigma_{\text{bgnd}}$  are the standard deviation of cyst and background.

The image quality of the 2-D slice images obtained in Case A by the baseline non-separable beamforming (shown in Fig. 12(a)) and our proposed separable beamforming method (shown in Fig. 12(b)) are nearly indistinguishable; both achieve an average CNR of 2.0 and an average CR of 0.55.

The 2-D slices of 3-D images obtained in Case B by non-separable beamforming and separable beamforming method are shown in Fig. 13(a) and Fig. 13(b), respectively. We perform coordinate transformation and scan conversion in order to display the 2-D slices vertically. In these images, the vertical axis indicates depth  $R$  rather than the  $z$  axis coordinate. The images produced by non-separable method achieve an average CNR of 1.55 an average CR 0.55, while the images produced by separable method achieve an average CNR of 1.45 and an average CR of 0.55.

Finally, we confirm that the fixed-point performance of the proposed methods matches the quality of full double-precision floating point. We compare results of 12-bit and 14-bit separable beamforming in Fig. 14(a) and Fig. 14(b). Both the 14-bit and 12-bit implementation achieve the same average CNR of 2.0, as in the double-precision-floating-point separable beamforming, although the 12-bit implementation has a slightly lower average CR compared to the 14-bit implementation (0.54 vs 0.55). Compared to the 14-bit nonseparable beamforming suggested in [12], 12-bit is sufficient for separable beamforming, because truncations can be done on partial beamforming data  $F_l^{(1)}(n_y, m_R, m_\theta; t)$  to prevent overflow without affecting




Fig. 12. 2-D slices of 3-D simulation images, Case A:  $\theta = \phi = 0^\circ$ , 40 dB dynamic range. (a) Non-separable beamforming, average CNR of 2.0, average CR of 0.55; (b) separable beamforming, average CNR of 2.0, average CR of 0.55.

image quality. Hence we propose a 12-bit data path in our hardware implementation for separable beamforming.

#### D. Power Analysis

To analyze the power and performance of our separable beamforming system, we use power estimates for each component of the design. For the beamforming accelerator, we use RTL-level Verilog synthesis results of the accelerator hardware using an industrial 45 nm standard cell library. SRAM values are generated using an industrial SRAM compiler, and our network power is obtained using SPICE models of our wires

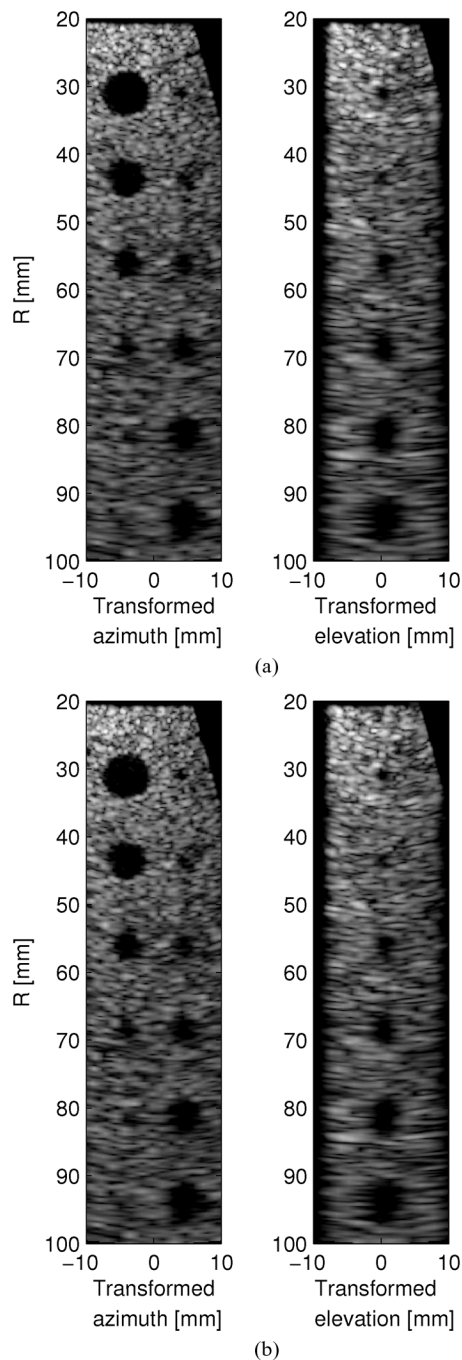


Fig. 13. 2-D slices of 3-D simulation images, Case B:  $\theta = \phi = 30^\circ$ , 40 dB dynamic range. (a) Non-separable beamforming, average CNR of 1.55, average CR of 0.55; (b) Separable beamforming, average CNR of 1.45, average CR of 0.55.

in 45 nm. Published state-of-the-art power numbers are used for ADC [31], DRAM [32], and memory interconnect (ARM Cortex M-3) [33].

In addition to our power analysis at 45 nm, we also project power requirements to 11 nm technology using published trends. ADC scaling uses values from [30], technology scaling is taken from [29], and we assume network wire power does not scale other than the shortening of the wires due to transistor area scaling.

Fig. 15 shows the complete power breakdown of each component and a total system power at 45 nm of just below 15 W




Fig. 14. 2-D slices of 3-D simulation images generated by fixed-point algorithm, Case A:  $\theta = \phi = 0^\circ$ , 40 dB dynamic range. (a) Separable beamforming, for 12-bit data path, average CNR of 2.0, average CR of 0.54; (b) separable beamforming, for 14-bit data path, average CNR of 2.0, average CR of 0.55.

for a frame rate of 32 Hz. Based on the scaling trends, Sonic Millip3De is just within our 5 W target at the 16 nm node and falls well below the target power by the 11 nm node. Thus, compared to our previous implementations in [11], [12], the separable beamforming method improves the frame rate from 1–2 Hz to 32 Hz without increasing the whole system power consumption.

## V. CONCLUSION

In this paper, we presented a separable beamforming method that reduces the computational complexity of 3-D ultrasound

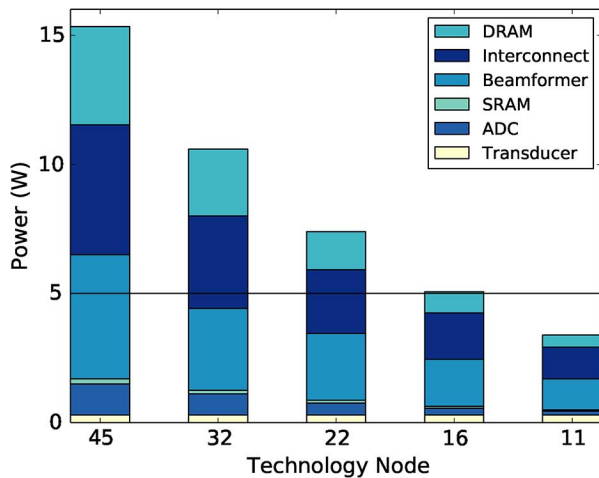


Fig. 15. Power Breakdown Across Technology Nodes. Scaling projections based on trends reported in [29], [30]. We project meeting the 5 W power budget at the 16 nm node.

imaging systems. The method is based on decomposing the delay term in a way that minimizes RMS phase error. We investigated several key factors for the separable delay functions, including the form of decomposition and order of beamforming and studied trade-offs between approximation error and computational complexity. We also proposed extensions to the Sonic Millip3De beamforming hardware accelerator to efficiently implement the separable beamforming method. Synthesis targeting an industrial 45 nm standard cell library indicated that the design can produce 3-D images at 32 Hz frame rate within a 15 W power budget, compared to 1–2 Hz frame rate in [12] without increasing the power consumption. We validate image quality via Field II cyst phantom simulations, which show that the proposed separable beamforming synthetic aperture ultrasound system can produce high quality images that are comparable to those produced by the non-separable baseline method.

ACKNOWLEDGMENT

The authors thank the reviewers for their insightful comments and J. Brian Fowlkes, Oliver Kripfgans and Matthew O’Donnell for many helpful discussions.

REFERENCES

[1] S. Campbell, C. Lees, G. Moscoso, and P. Hall, “Ultrasound antenatal diagnosis of cleft palate by a new technique: The 3D reverse face view,” *Ultrasound Obstetr. Gynecol.*, vol. 25, no. 1, pp. 12–18, 2005.  
 [2] S. Yagel, S. M. Cohen, I. Shapiro, and D. V. Valsky, “3D and 4D ultrasound in fetal cardiac scanning: A new look at the fetal heart,” *Ultrasound Obstetr. Gynecol.*, vol. 29, no. 1, pp. 81–95, 2007.  
 [3] T.-H. Yu, S.-Y. Sun, C.-L. Ding, P.-C. Li, and A.-Y. Wu, “Reconfigurable color Doppler DSP engine for high-frequency ultrasonic imaging systems,” in *Proc. IEEE Workshop Signal Process. Syst.*, 2007, pp. 187–192.

[4] A. B. Abche, A. Maalouf, R. Ayoubi, E. Karam, and A. M. Alameddine, “An FPGA implementation of a high resolution phase shift beamformer,” in *Proc. IEEE Int. Conf. Signal Process. Commun.*, 2007, pp. 1319–1322.  
 [5] E. Montagnon, S. Hissoiny, P. Despres, and G. Cloutier, “Real-time processing in dynamic ultrasound elastography: A GPU-based implementation using CUDA,” in *Proc. Int. Conf. Inf. Sci., Signal Process., Appl.*, 2012, pp. 472–477.  
 [6] I.-H. Lee, Y.-H. Chen, N.-S. Huang, and A.-Y. Wu, “Accelerating motion-compensated adaptive color Doppler engine on CUDA-based GPU platform,” in *Proc. IEEE Workshop Signal Process. Syst.*, 2013, pp. 1–6.  
 [7] C. R. Hazard and G. R. Lockwood, “Theoretical assessment of a synthetic aperture beamformer for real-time 3-D imaging,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 46, no. 4, pp. 972–980, 1999.  
 [8] A. C. Dhanantwari, S. Stergiopoulos, L. Song, C. Parodi, F. Bertor, P. Pellegritti, and A. Questa, “An efficient 3D beamformer implementation for real-time 4D ultrasound systems deploying planar array probes,” in *Proc. IEEE Ultrason. Symp.*, Aug. 2004, vol. 2, pp. 1421–1424.  
 [9] J. A. Jensen, H. Holten-Lund, R. T. Nilsson, M. Hansen, U. D. Larsen, R. P. Domsten, B. G. Tomov, M. B. Stuart, S. I. Nikolov, M. J. Pihl, Y. Du, J. H. Rasmussen, and M. F. Rasmussen, “SARUS: A synthetic aperture real-time ultrasound system,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 60, no. 9, pp. 1838–1852, 2013.  
 [10] Y.-T. Shen, A. Freibert, S. Stergiopoulos, and K. Plataniotis, “Computing architecture for the portable four-dimensional ultrasound diagnostic imaging system,” in *Proc. IEEE Ultrason. Symp.*, 2012, pp. 2059–2062.  
 [11] R. Sampson, M. Yang, S. Wei, C. Chakrabarti, and T. F. Wenisch, “Sonic Millip3De: Massively parallel 3D stacked accelerator for 3D ultrasound,” in *Proc. 19th IEEE Int. Symp. High Perform. Comput. Architect.*, Feb. 2013, pp. 318–329.  
 [12] R. Sampson, M. Yang, S. Wei, C. Chakrabarti, and T. F. Wenisch, “Sonic Millip3De with dynamic receive focusing and apodization optimization,” in *Proc. IEEE Int. Ultrason. Symp.*, Jul. 2013, pp. 557–560.  
 [13] M. Cetin, E. Bossy, R. Cleveland, and W. C. Karl, “Sparsity-driven sparse-aperture ultrasound imaging,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2006, vol. 2, pp. 89–92.  
 [14] M. Karaman, I. O. Wygant, O. Oralkan, and B. T. Khuri-Yakub, “Minimally redundant 2-D array designs for 3-D medical ultrasound imaging,” *IEEE Trans. Med. Imag.*, vol. 28, no. 7, pp. 1051–1061, Jul. 2009.  
 [15] B. Diarra, H. Liebgott, M. Robini, P. Tortoli, and C. Cachard, “Optimized 2D array design for ultrasound imaging,” in *Proc. 20th Eur. Signal Process. Conf.*, 2012, pp. 2718–2722.  
 [16] B. Savord and R. Solomon, “Fully sampled matrix transducer for real time 3D ultrasonic imaging,” in *Proc. IEEE Symp. Ultrason.*, Oct. 2003, vol. 1, pp. 945–953.  
 [17] F. M. Hooi, K. E. Thomenius, R. Fisher, and P. L. Carson, “Hybrid beamforming and steering with reconfigurable arrays,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 6, pp. 1311–1319, Jun. 2010.  
 [18] A. Tawfik, S. Stergiopoulos, and A. C. Dhanantwari, “A generic beamforming structure allowing implementation of adaptive processing schemes for 2-D and 3-D arrays of sensors,” in *Proc. MTS/IEEE OCEANS Conf.*, 1997, vol. 1, pp. 369–373.  
 [19] S. Stergiopoulos, *Advanced Signal Processing Handbook: Theory and Implementation for Radar, Sonar, and Medical Imaging Real Time Systems*. Boca Raton, FL, USA: CRC Press, 2010.  
 [20] K. Owen, M. I. Fuller, and J. A. Hossack, “Application of X-Y separable 2-D array beamforming for increased frame rate and energy efficiency in handheld devices,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 59, no. 7, pp. 1332–1343, Jul. 2012.  
 [21] M. Yang, R. Sampson, T. F. Wenisch, and C. Chakrabarti, “Separable beamforming for 3-D synthetic aperture ultrasound imaging,” in *Proc. IEEE Workshop Signal Process. Syst.*, Oct. 2013, pp. 207–212.  
 [22] B. Dacorogna, *Introduction to the Calculus of Variations*. Singapore: World Scientific, 2004.  
 [23] B. Black, M. Annavaram, N. Brekelbaum, J. DeVale, L. Jiang, G. H. Loh, D. McCaule, P. Morrow, D. W. Nelson, D. Pantuso, P. Reed, J. Rupley, S. Shankar, J. Shen, and C. Webb, “Die stacking (3D) microarchitecture,” in *Proc. 39th Annu. IEEE/ACM Int. Symp. Microarchitect.*, Dec. 2006, pp. 469–479.

- [24] O. Oralkan, A. S. Ergun, J. A. Johnson, M. Karaman, U. Demirci, K. Kaviani, T. H. Lee, and B. T. Khuri-Yakub, "Capacitive micro-machined ultrasonic transducers: Next-generation arrays for acoustic imaging?," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 49, no. 11, pp. 1596–1610, Nov. 2002.
- [25] J. A. Jensen, "FIELD: A program for simulating ultrasound systems," in *Proc. 10th Nordicbaltic Conf. Biomed. Imag., Vol. 4, Suppl. 1, Pt. 1*, 1996, vol. 4, pp. 351–353.
- [26] J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers," *IEEE Trans. Ultrason. Ferroelectr., Freq. Control*, vol. 39, no. 2, pp. 262–267, Mar. 1992.
- [27] M. Karaman, H. S. Bilge, and M. O'Donnell, "Adaptive multi-element synthetic aperture imaging with motion and phase aberration correction," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 45, no. 4, pp. 1077–1087, Jul. 1998.
- [28] G. R. Lockwood, J. R. Talman, and S. S. Brunke, "Real-time 3-D ultrasound imaging using sparse synthetic aperture beamforming," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 45, no. 4, pp. 980–988, Jul. 1998.
- [29] H. Esmailzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *Proc. 38th Int. Symp. on Comput. Architect.*, Jun. 2011, pp. 365–376.
- [30] B. Murmann, "ADC performance survey 1997–2013," [Online]. Available: <http://www.stanford.edu/murmann/adcsurvey.html>
- [31] B. Verbruggen, M. Iriguchi, and J. Craninckx, "A 1.7 mw 11b 250 ms/s 2-times interleaved fully dynamic pipelined SAR ADC in 40 nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 2880–2887, Dec. 2012.
- [32] K. Malladi, F. Nothaft, K. Periyathambi, B. Lee, C. Kozyrakis, and M. Horowitz, "Towards energy-proportional datacenter memory with mobile DRAM," in *Proc. 39th Int. Symp. Comput. Architect.*, June 2012, pp. 37–48.
- [33] ARM, Cortex-m3 40g specifications [Online]. Available: <http://www.arm.com/products/processors/cortex-m/cortex-m3.php>




**Ming Yang** (S'09) has an M.S. in electrical engineering from Beijing University of Posts and Telecommunications. He is a Ph.D. candidate in the School of Electrical, Computer and Energy Engineering at Arizona State University. His research focuses on the development of algorithms and low-power hardware for a handheld 3D ultrasound imaging device.



**Richard Sampson** (S'07) Sampson has a B.A. in Physics and a B.S. in Computer Engineering from Columbia University as well as an M.S. in Computer Science and Engineering from the University of Michigan. He is a Ph.D. candidate in the Department of Computer Science and Engineering at the University of Michigan. His research interests include hardware system and accelerator design for imaging and computer vision applications.



**Siyuan Wei** (S'14) has a B.E. in electrical engineering from Huazhong University of Technology and Science. He is a Ph.D. student in the School of Electrical, Computer and Energy Engineering at Arizona State University. His research focuses on algorithm-architecture codesign of ultrasound imaging systems, especially those based on Doppler imaging.



**Thomas F. Wenisch** (M'07) received his Ph.D. from Carnegie Mellon University. He is an Associate Professor of Electrical Engineering and Computer Science at the University of Michigan and a member of the Advanced Computer Architecture Lab (ACAL). His research interests center on computer architecture with particular emphasis on multiprocessor and multicore systems, smartphone architecture, data center architecture, 3D medical imaging applications, and performance evaluation methodology.



**Chaitali Chakrabarti** (M'89–SM'02–F'11) received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, in 1984, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1986 and 1990, respectively. She is a Professor with the School of Electrical Computer and Energy Engineering, Arizona State University (ASU), Tempe. Her research interests include low power embedded systems design, and VLSI algorithm-architecture codesign of signal processing and communication systems.