10.1109/ULTSYM.2013.0300

# Breast imaging using ultrasound tomography: From clinical requirements to system design

Olivier Roy*†, Steven Schmidt*†, Cuiping Li*†, Veerendra Allada*‡, Erik West*, David Kunz*, and Neb Duric*†

*Delphinus Medical Technologies, 46701 Commerce Center Drive, Plymouth, MI 48170, USA
†Karmanos Cancer Institute, 4100 John R, Detroit, MI 48201, USA
‡Iowa State University, 2229 Lincoln Way, Ames, IA 50011, USA

*Abstract*—**Ultrasound tomography (UST) is a breast imaging modality that is radiation free, operator independent, and does not involve breast compression. In the UST system under consideration, the breast is surrounded by a transducer ring that moves along the coronal axis from the chest wall to the nipple region. The deployment of UST in a clinical setting is technically challenging from three major standpoints: acquisition speed, storage capability, and computational requirements. Data acquisition must be fast to maximize patient throughput and minimize image registration artifacts. Unlike traditional ultrasound, hundreds of gigabytes of data must be acquired, buffered, and processed to image various characteristics of breast tissues such as sound speed, attenuation, and reflectivity. The tomographic image reconstruction methods are non-linear, iterative algorithms with significant computational complexity. Moreover, the scanner hosting the acquisition and reconstruction components must satisfy stringent cost, power, and size requirements. For decades, the above constraints have hindered the practicality of UST in a clinical scenario. We describe the design of a UST system that addresses the relevant clinical requirements as a means to demonstrate the feasability of UST deployment in a clinical setting.**

## I. INTRODUCTION

Ultrasound tomography (UST) is a technique that uses computed tomography methods to solve an inverse problem involving ultrasound signals. In the scenario of interest, the breast to be imaged is surrounded by an ultrasound transducer ring. Each transducer element emits, in turn, an ultrasound pulse that propagates through the medium and gets recorded by all the elements of the ring aperture. Acoustic properties of the medium are inferred from the recorded ultrasound data using image reconstruction algorithms. Unlike conventional ultrasound, both the transmitted and reflected components of the waveforms are used to characterize breast tissue. The amount of data acquired by a UST scanner, and the computational complexity of the associated tomographic reconstruction techniques, make it difficult to apply this technology clinically.

A number of investigators have developed scanners based on UST principles. Examples include the work by Marmarelis *et al.* [1], Johnson *et al.* [2], and Ruiter *et al.* [3]. Our group at the Karmanos Cancer Institute (KCI) has also focused on the development of a UST research prototype for breast cancer detection [4]. The clinical feedback obtained from the studies performed over the past decade at KCI's breast center has guided the continuous development of that prototype, as well as its upgrade to SoftVue, the commercial system developed by our start-up company, Delphinus Medical Technologies[1]. The upgrade is challenging from many standpoints. The scanner must house a water conditioning system, a powerful computing platform, and a tomographic acquisition device, while satisfying stringent clinical requirement ranging from size, and weight, to acquisition speed and computational power. We explain the rationale behind these requirements and describe a system architecture that addresses them. Our discussion focuses on the acquisition and computing aspects of the scanner, although many of the requirements we touch upon have broader implications.

The outline of the paper is as follows. In Section II we review the clinical requirements that motivate the design of SoftVue. Section III describes the design in details, in particular data acquisition, data pre-processing, networking, and image reconstruction. Section IV concludes the discussion.

## II. CLINICAL REQUIREMENTS

The design of a medical device for breast cancer detection is driven by clinical requirements that touch upon a great deal of aspects including cost, safety, patient throughput, and imaging capability. We discuss a few of these aspects and describe some of their implications pertaining to system design.

### A. Cost

One of the key objectives of UST is to combine the cost effectiveness of mammography and conventional ultrasound, with the superior imaging performance of magnetic resonance imaging (MRI). Installation and maintenance costs can be kept to a minimum by designing a system with a small footprint which does not require a dedicated examination room that is expensive to build and operate. This leads to size, weight, and power requirements that further impose limitations on the computing platform that can be hosted in the scanner enclosure.

### B. Safety

National cancer screening statistics indicate that only 51% of eligible women undergo annual mammograms in the United States [5]. That rate is even lower for African American women and/or those of lower socioeconomic groups. Fear of radiation and discomfort are some of the factors that contribute to the

---

[1] SoftVue has not yet received FDA or CE mark approval. The authors have financial interest in Delphinus Medical Technologies.

low participation rate. UST, by virtue of its compression and radiation free imaging capability, addresses some of these concerns, and offers the potential to increase participation rate which would lead to cancer detection at an earlier stage leading to a greater survival rate.

### *C. Patient Throughput*

The ability to sustain a high patient throughput is important to reduce the operating cost of the device in order to make it a viable alternative to existing technologies (e.g., mammography) for routine use. In a screening scenario, it is beneficial to be able to quickly schedule a follow up exam in the case of abnormal findings. In other words, we should minimize the duration between the time the patient enters the examination room, and the time a decision can be made by the radiologist. Our system has been designed with the goal of sustaining a throughput of four patients per hour (bilateral exams).

Sustaining such a throughput requires the scan of a single breast to be done in less than a couple of minutes on average. It also enforces stringent requirements on the time taken to reconstruct and store the images for review. In this context, the design of a scalable computing platform is of utmost importance to be able to ride Moore's law and benefit from the latest advances in computing technology. Ease of use (e.g., during patient positioning), comfort, and operator independence allow for faster exam procedures, hence promoting higher patient throughput. The image review process must also be streamlined, for example using computer aided diagnosis (CAD) tools, to help the radiologist with the analysis and evaluation of the medical images. The development of CAD for UST is a work in progress and is not addressed in this paper.

### *D. Imaging Capability*

Imaging plays an important role in the early detection of breast cancer. Mammography, while having been shown to greatly reduce the mortality rate of woman over the age of 50, is known to have poor imaging capability in dense breast tissue. At the same time women with dense breast are at the highest risk for developing breast cancer [6]. Consequently many cancers are missed at their earliest stages when they are the most treatable. MRI can significantly improve on these limitations at the expense of long exam times and the use of contrast agents.

UST has the potential to rival the imaging performance of MRI at a fraction of the cost. To this end, the data acquisition system and image reconstruction algorithms must work in concert to produce images that are clinically relevant for early beast cancer detection. Reconstruction algorithms must be accurate enough for tissue characterization and lesion detection, while being viable from a computational standpoint. The UST system must also be capable of imaging the entire breast from the chest wall to the nipple region.

## III. SYSTEM DESIGN

### *A. Overview*

The SoftVue scanner is shown in Figure 1. During a scan, the patient is placed in the prone position on a flat table with the breast suspended through a hole in the table in a water bath lying just below the table surface. The water is a requirement that ensures minimal distortion of the breast while allowing strong coupling of acoustic waves to the tissue. SoftVue hardware platform comprises three main building blocks: acquisition, reconstruction, and control.

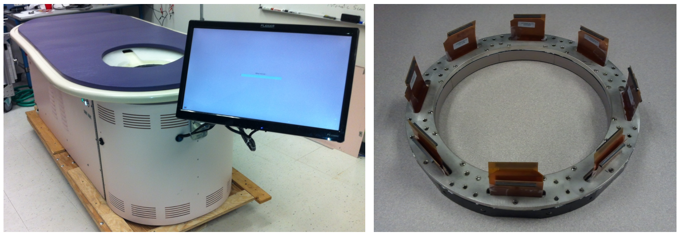
(a) (b)

Fig. 1. The SoftVue scanner. (a) The SoftVue table. The breast is suspended through a hole in the table in a water bath lying just below the table surface. The system is operated through a touchscreen monitor. (b) The 2048 element transducer ring that moves in the water bath along the coronal axis using a motorized gantry.

The acquisition engine acquires the raw ultrasound data using a set of pre-defined acquisition parameters. At the end of each slice acquisition, the raw data is transferred to the reconstruction engine for image reconstruction.

The reconstruction engine is a blade server architecture that can host up to ten blades. It is currently operated with one master blade and four compute blades. The master blade is directly connected to the acquisition engine, and is equipped with two quad-core Intel Xeon E5620 CPUs, and 192 gigabytes (GB) of random access memory (RAM). It preprocesses each slice of raw data and transfers it to a compute blade selected according to a minimum load scheduling strategy. A compute blade has two quad-core Intel Xeon E5620 CPUs, two Nvidia Tesla M2070 GPU devices, and 96 GB of RAM. Each processor is capable of executing 8 threads by using the Intel hyper-threading technology. Each compute blade reconstructs the images using the raw data it receives. It then sends the results to the master blade for image aggregation, stacking, and storage in DICOM files.

The goal of SoftVue's control engine is to coordinate the data acquisition and image reconstruction process with the rest of the operations of the scanner (e.g., water conditioning, power switching, patient positioning). SoftVue is operated through a graphical user interface (GUI) displayed on a touchscreen monitor. The images produced by SoftVue are transferred to an external PACS server. For research purposes, the raw data acquired by SoftVue can be saved and transferred to an external FTP server.

### *B. Data Acquisition*

During data acquisition, the breast is surrounded by a solid state ring transducer that can move along the coronal axis with a 16 cm span using a motorized gantry. The diameter of the ring is 22 cm. These numbers have been chosen to accommodate about $4\sigma$ of the US population, according to the demographics reported in [7] as well as statistics computed on patients of our own research study [4]. The

TABLE I. SIZE IN GIGABYTES OF A SLICE OF RAW DATA AS A FUNCTION OF THE NUMBER OF ARRAY ELEMENTS AND THE SAMPLING FREQUENCY.

| Sampling frequency (MHz) | Number of elements | | |
|---|---|---|---|
| | 256 | 512 | 1024 |
| 10 | 0.21 | 0.86 | 3.44 |
| 12 | 0.26 | 1.03 | 4.13 |
| 14 | 0.30 | 1.20 | 4.81 |

transducer ring comprises 2048 elements multiplexed on 512 acquisition channels to reduce the cost and complexity of the acquisition hardware. Each acquisition channel is equipped with variable time gain control, anti-aliasing filtering, and A/D conversion (ADC) capability. The transducers operate at a central frequency of 3 MHz with significant energy in the 1 to 5 MHz band. The sampling rate is 12 MHz. The recording length is set to 176 $\mu s$ to ensure that reflections at the center of the ring are captured in scenarios where the speed of sound is as low as 1.25 km/s. A fixed gain is applied to all the waveforms to make use of the entire ADC range while avoiding saturation. Samples are quantized on 14 bits, and zero padded to 16 bits. Each slice is made of a number of shots. In each shot, a transducer fires a broadband ultrasound pulse which is recorded by the multiplexed receive elements. The inter-shot time interval is chosen such that the residual energy from the previous shot falls below the noise level. When the acquisition is done, the data is transferred to the master blade of the reconstruction engine. While the data are being transferred, the transducer ring is moved by a configurable distance, referred to as the inter-slice spacing.

The above parameters enable trade-offs between signal coverage, signal-to-noise ratio (SNR), scan rate, size of the acquired dataset, image reconstruction time, and image quality. The size of a single slice of raw data for different acquisition parameters is listed in Table I. Extensive simulation and experimental tests have been performed to select operational parameters that address the clinical requirements stated in Section II (see [8] for details). In SoftVue, the slice is acquired using 1024 emitters and receivers, a 800 $\mu s$ inter-shot time interval, and a 3 mm inter-slice interval. With these parameters, the size of a single slice of data is 4.125 GB. A slice is acquired in 1.64 seconds, and transferred to the reconstruction engine using a 16 lane PCI express link in about 0.66 second. It should be noted that, while these parameters have been fixed, SoftVue has been designed to be a fully scalable computing platform. Operational parameters can be adjusted with only minimal modifications.

### C. Data Preprocessing

The raw data are acquired at a much higher rate than it can be processed for image reconstruction. Typically, it takes about 2.3 seconds to acquire and transfer a slice of data, and a couple of minutes to process it using a single CPU-GPU configuration on a single compute blade. The data thus need to be buffered. Note that with 4 compute blades, 8 such tasks can be run in parallel to output a set of images every 20 seconds on average. To maximize the number of slices that can be stored as well as to put it in a form that is suitable for image reconstruction, pre-processing is applied on the master blade before the data can be sent to a compute blade for reconstruction.

Firstly, the data are transposed such that the temporal samples of a waveform are stored contiguously in memory. The data samples are then converted from their original 14 bit 2's complement format to the 16 bit 2's complement format. Saturated waveforms are detected and possibly discarded. The waveforms corresponding to reciprocal emitter-receiver pairs are averaged. Finally, the data pre-processing has the capability to discard portions of the waveform that are not used by the reconstruction algorithms; e.g., the portion of the received waveform that is recorded before the first arrival (time-of-flight). Reciprocal averaging and data cutting allows us to reduce the size of a slice by a factor 4. Pre-processing is a time critical operation since it may limit the sustainable scan rate in the case of a large breast where hundreds of gigabytes of data need to be acquired in a short time frame. SoftVue's pre-processing core is multi-threaded and heavily relies on SSE instructions. It can sustain a processing rate of 3.17 GB/s.

### D. Networking

In order to sustain a reasonable scan rate, the pre-processed data needs to be quickly distributed from the master blade to the compute blades. To this end, we have developed a dedicated multi-threaded run-time communication service over Infiniband (IB) network hardware and its remote direct memory access (RDMA) protocols. IB supports different message transfer protocols of which the send/receive based channel semantics and RDMA based read/write memory semantics are relevant for this paper. The lowest level interface to programming the IB is via a set of standard functions referred to as verbs; it has low software overhead and offers the best performance. One major advantage of RDMA over TCP/IP protocols is that the operating system is not involved in read and write operations, allowing high throughput and low-latency networking. The blades of the reconstruction engine are connected with a high speed IB network switch (Mellanox InifiniScale IV). Each blade has a single port QDR Infiniband adapter (Mellanox ConnectX) supporting a theoretical bandwidth of 40 Gb/s per blade.

The raw data are disseminated among the compute blades via the RDMA write/read protocols. The service also exposes a well-defined application programming interface. The service runs as a dedicated thread on the operating system, and takes the role of a client (master blade) or server (compute blade) based on the arguments given at start up. For the data a push/pull protocol is implemented using the RDMA write/read calls, respectively. With the verbs implementation, an average throughput of 2.69 GB/s has been achieved on the raw data transfers. Pre-processing is applied on the master node to reduce the amount of data by a factor of 4 for further reduction of transfer time.

With the chosen operating parameters, the scan rate that can be sustained by SoftVue is limited by the acquisition chain to 0.43 slice per second. An average breast of length 7.19 cm (in water) can thus be scanned in less than a minute using a 3 mm inter-slice spacing. Further software and hardware optimization is likely to improve upon these results.

### E. Image Reconstruction and Storage

Image reconstruction is a key component of the processing chain. Image reconstruction algorithms turn raw data into

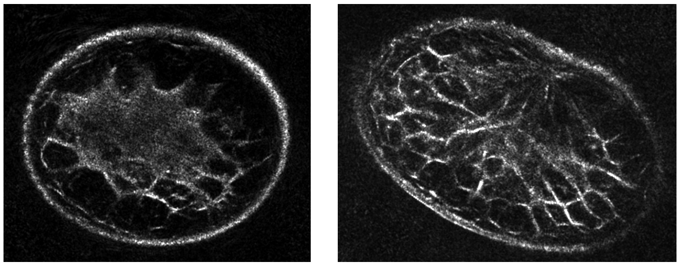

(a) (b)

Fig. 2. Examples of reflection images produced by SoftVue (the images have been cropped). The sound speed and attenuation images are used during the synthetic aperture beam-forming process for delay and gain compensation.

images by solving an inverse problem. On each compute blade, two reconstruction tasks run in parallel, each using the resources of a single GPU card. The accuracy, robustness, and computational complexity of the inverse problem are essentially driven by the forward model under consideration. The forward model is the model chosen to simulate the propagation of ultrasound waveforms in the medium whose characteristics need to be estimated. SoftVue can image three characteristics of breast tissues: sound speed, attenuation, and reflectivity. Examples of reflection images produced by SoftVue are shown in Figure 2.

Sound speed images are computed using a non-linear iterative reconstruction technique that alternates between estimating the propagation paths (or rays) between a set of emitters and receivers, and estimating the sound speed along these rays. For each emitter, we select an aperture of receivers centered around the receiver facing the emitter on the ring. The size of the aperture is set to a fraction 3/4 of the ring. The ray trajectories at a given iteration are computed by solving the Eikonal equation using the current sound speed model, and then finding the paths of minimum travel time between every emitter-receiver pair. At each iteration, the sound speed model is updated using a conjugate gradient method that attempts to minimize residuals between observed and measured travel times at the receiver locations. Travel time information is extracted from the raw data using a travel time picker based on a statistical information criterion. The sound speed is updated from a coarse 4 mm pixel size grid to a finer 1 mm pixel size grid using a fixed number of iterations. This multi-resolution approach increases both the robustness of the method to noise and the reconstruction speed. To further increase reconstruction speed, many components of the reconstruction chain (e.g., travel time picking, ray tracing, conjugate gradient optimization) have been ported to GPU using the CUDA programming framework introduced by Nvidia. The use of a forward propagation model that does not assume straight propagation paths increases the computational complexity of the algorithm but provides better reconstruction accuracy.

Using the rays computed by the sound speed reconstruction algorithm described above, attenuation reconstruction is a linear inverse problem that minimizes residuals between observed and measured power loss ratios. These power ratios are computed by integrating the analytic magnitude of the transmitted portion of the received waveforms in both a reference water propagation medium (obtained during a calibration phase), and in the inhomogeneous medium of interest (breast). Most of these computations are done on GPU.

The reflection images are produced using a synthetic aperture beam-forming algorithm. For each pixel, the contribution of a set of emitter-receiver pairs are summed up with appropriate delay and gain compensation. The delay and gain are computed by integrating the slowness (inverse of sound speed) and attenuation, respectively, along the reflection path. For each emitter, we select an aperture of receivers centered around the emitter. The size of this aperture is set to a fraction 2/3 of the ring. A local spatial averaging filter is applied on the absolute values of the pixels to obtain the final image. Most of these computations are done on GPU.

The images reconstructed by the compute blades are sent to the master blade for aggregation and stacking. The master blade stores these stacks in DICOM files on the reconstruction engine internal storage. Since no UST DICOM format is currently available, the MRI DICOM object is used instead. A number of private tags that describe the units of the images produced by SofVue are also included. SoftVue's control engine is responsible for pushing the studies to an external PACS server for review by the radiologist.

## IV. CONCLUSION

We have described SoftVue, a UST scanner that allows for fast, operator independent, compression-free, non-invasive, and radiation-free imaging of the breast. The clinical requirements taken into account for its design have been reviewed, and the technical specifications of the modality have been detailed. Further research and development efforts concentrate on optimizing the overall acquisition and reconstruction chain, as well as developing CAD tools to assist with the image review process.

## REFERENCES

[1] V. Z. Marmarelis, J. Jeong, D. C. Shin, and S. Do, "High-resolution 3-D imaging and tissue differentiation with transmission tomography," in *Acoustical Imaging*. Springer Netherlands, 2007, vol. 28, pp. 195–206.

[2] S. A. Johnson, D. T. Borup, J. W. Wiskin, F. Natterer, F. Wuebbling, Y. Zhang, and S. C. Olsen, "Apparatus and method for imaging with wavefields using inverse scattering techniques," US Patent 6,005,916, 1999.

[3] N. V. Ruiter, G. Göbel, L. Berger, M. Zapf, and H. Gemmeke, "Realization of an optimized 3D USCT," in *SPIE*, vol. 7968, Mar. 2011.

[4] N. Duric, P. Littrup, L. Poulo, and A. Babkin, "Detection of breast cancer with ultrasound tomography: First results with the Computed Ultrasound Risk Evaluation (CURE) prototype," *Medical Physics*, vol. 34, no. 2, pp. 773–785, Feb. 2007.

[5] American Cancer Society, "Cancer prevention & Early detection facts & Figures 2009," *Atlanta, GA: American Cancer Society*, pp. 34–37, 2009.

[6] G. Ursin, L. Hovanessian-Larsen, Y. R. Parisky, M. C. Pike, and A. H. Wu, "Greatly increased occurrence of breast cancers in areas of mammographically dense tissue," *Breast Cancer Research*, vol. 7, no. 5, pp. R605–R608, 2005.

[7] J. M. Boone, N. Shah, and T. R. Nelson, "A comprehensive analysis of dgn ct coefficients for pendant-geometry cone-beam breast computed tomography," *Medical Physics*, vol. 31, no. 2, pp. 226–235, 2004.

[8] N. Duric, P. Littrup, C. Li, O. Roy, S. Schmidt, O. Rama, and L. Bey-Knight, "Breast ultrasound tomography: Bridging the gap to clinical practice," in *SPIE*, vol. 8320, Feb. 2012.