Ultrasound goes GPU: real-time simulation using CUDA

Tobias Reichl[1,2], Josh Passenger[2], Oscar Acosta[2], Olivier Salvado[2]

[1] Computer-Aided Medical Procedures (CAMP), TUM, Munich, Germany
[2] CSIRO, ICTC, The Australian e-Health Research Centre – Biomedical Imaging
Royal Brisbane and Women's Hospital, Herston, Australia

ABSTRACT

Despite the increasing adoption of other imaging modalities, ultrasound guidance is widely used for surgical procedures and clinical imaging due to its low cost, non-invasiveness, and real-time visual feedback. Many ultrasound-guided procedures require extensive training and where possible training on simulations should be preferred over patients. Computational resources for existing approaches to ultrasound simulation are usually limited by real-time requirements. Unlike previous approaches we simulate freehand ultrasound images from CT data on the Graphics Processing Unit (GPU). We build upon the method proposed by Wein et al. for estimating ultrasound reflection properties of tissue and modify it to a computationally more efficient form. In addition to previous approaches, we also estimate ultrasound absorption properties from CT data. Using NVIDIA's "Compute Unified Device Architecture" (CUDA), we provide a physically plausible simulation of ultrasound reflection, shadowing artifacts, speckle noise and radial blurring. The same algorithm can be used for simulating either linear or radial imaging, and all parameters of the simulated probe are interactively configurable at runtime, including ultrasound frequency and intensity as well as field geometry. With current hardware we are able to achieve an image width of up to 1023 pixels from raw CT data in real-time, without any pre-processing and without any loss of information from the CT image other than from interpolation of the input data. Visual comparison to real ultrasound images indicates satisfactory results.

Keywords: Surgical Simulation, Ultrasound Guidance, Image-Guided Therapy, Visualisation

1. INTRODUCTION

Despite the increasing adoption of Computed Tomography (CT) and Magnetic Resonance Imaging, medical ultrasound (US) remains one of the most used imaging modalities, due to its low cost, widespread availability and high patient acceptance. However, procedures like US-guided needle insertions require a high degree of hand-eye coordination and spatial reasoning, so extensive training is needed to gain sufficient experience, and for ethical and patient safety reasons it is desirable to train surgeons on simulations rather than on patients.[1]

1.1. Related Work

Real-time simulation of US images has been claimed for at least a decade now, with advances mainly being made in the realism of the simulation and the amount of manual annotation required for preparation.

Different approximating methods have been proposed since Jensen's landmark paper[2] about exhaustive simulation on the scatterer level and Aiger and Cohen-Or's paper[3] about real-time simulation from 3-D US. One drawback of Jensen's approach is the long time needed for simulation, 11 hours for one image as stated in the original publication. Slicing 2-D images from a 3-D US volume, like proposed by Aiger and Cohen-Or, is strongly dependent on the acquisition parameters and the quality of the prerecorded 3-D data, including artifacts like shadows.

Many recent approaches, as proposed e.g. by Hostettler et al.,[4] Zhu et al.,[5] Wein et al.,[6] and Shams et al.,[7] show promising results, but detailed information about the achievable image resolution is commonly omitted. Possible approaches include an estimation of acoustic properties from CT data[4,6,7] or slicing from a 3-D model,

T. Reichl (corresponding author): reichl@in.tum.de, +49 89 289 19412 | J. Passenger: Josh.Passenger@csiro.au, +61 7 3253 3622 | O. Acosta: Oscar.Acosta@csiro.au, +61 7 3253 1620 | O. Salvado: Olivier.Salvado@csiro.au, +61 7 3253 3658

based on a segmentation of CT data, using synthetic 3-D textures, based on real 2-D ultrasound image samples.[5] While the latter approach produces convincly textured organs, it can only be as accurate or detailed as the underlying segmentation. Recently, an approach has been presented by Shams et al.[7] to combine both an estimation of acoustic properties from CT data, similar to Wein et al.,[6] and a scatterer image texture computed using the FIELD II algorithm developed by Jensen and Nikolov.[8]

Usually, there exists a trade-off between realism of the simulation and the computation effort needed. Harnessing the highly specialised hardware on the Graphics Processing Unit (GPU) of modern graphics cards offers a possible solution to simulating freehand US with high resolution in real-time. Taking advantage of GPUs requires additional effort for the parallel implementation of algorithms. However, the increasing number of publications around this topic shows that there is a significant interest and that the design of an implementation on GPU will be useful to the scientific community. It is clearly necessary to accelerate computation in order to achieve realism and real-time performance at the same time, and it is especially promising to use the parallel processing capabilities on the GPU.

2. DESCRIPTION OF PURPOSE

The purpose is to develop a physically accurate simulation of freeehand ultrasound in real time, that incorporates acquisition parameters such as ultrasound frequency, intensity, and gain, as well as tissue properties. To achieve this we accelerate the simulation by using the GPU for parallel processing. All three approaches[4–6] mentioned above involve some form of ray-casting, and as a consequence not all computations can be done on each pixel separately and to some extent synchronisation is needed on the GPU.

Vidal et al.[9] already published a description of an implementation on GPU. However, their simulation still has to be computed partly on the CPU, and for reproducing US-like reflections they rely on a rather vague "enhancement" of horizontal edges without further consideration of tissue types or probe geometry. Zhu et al.[5] use the GPU only to simulate radial blurring.

The well-known OpenGL Shader Language (GLSL) does not offer synchronisation between threads on the GPU, and thus we use NVIDIA's "Compute Unified Device Architecture" (CUDA). CUDA 1.1 does not yet include support for 3-D textures (those are scheduled for the upcoming release of CUDA 2.0), so we present a modified algorithm that is feasible in 2-D. From this, not only implementations in CUDA will benefit, but this in general reduces the computational complexity. For slicing 3-D data established techniques in GLSL can be used.

3. METHODS

The developed simulation framework consists of several stages as depicted in figure 1. First, given the position and orientation of the probe, a 2-D reformatted image is extracted from the 3-D volume. Then, the physical phenomena involved in US image generation are simulated from the data. Thus, reflection, absorption, and transmission are computed taking into account the properties of the structures and their boundaries (acoustic impedance, orientation) in interaction with US waves. Finally, speckle noise and blurring are added in a post-processing step, resulting in an image (see figure 9) that can be updated in real time according to the user-defined parameters.

3.1. Acoustic Impedance

It is possible to assume an approximately proportional relation of X-ray attenuation to tissue density. However, the assumption of a constant speed of sound is invalid, as e.g. the speed of sound in bones reaches up to 4330 m/s,[10] i.e. almost three times the speed in soft tissue, 1540 m/s. For a better estimation of the acoustic impedance we use use quadratic extrapolation from the known values of air, water, and an estimated value for bone: $Z = 7$ Mrayl and $\mu = 1000$ HU. For a screen shot of original CT data and estimated acoustic impedance cf. figure 3.

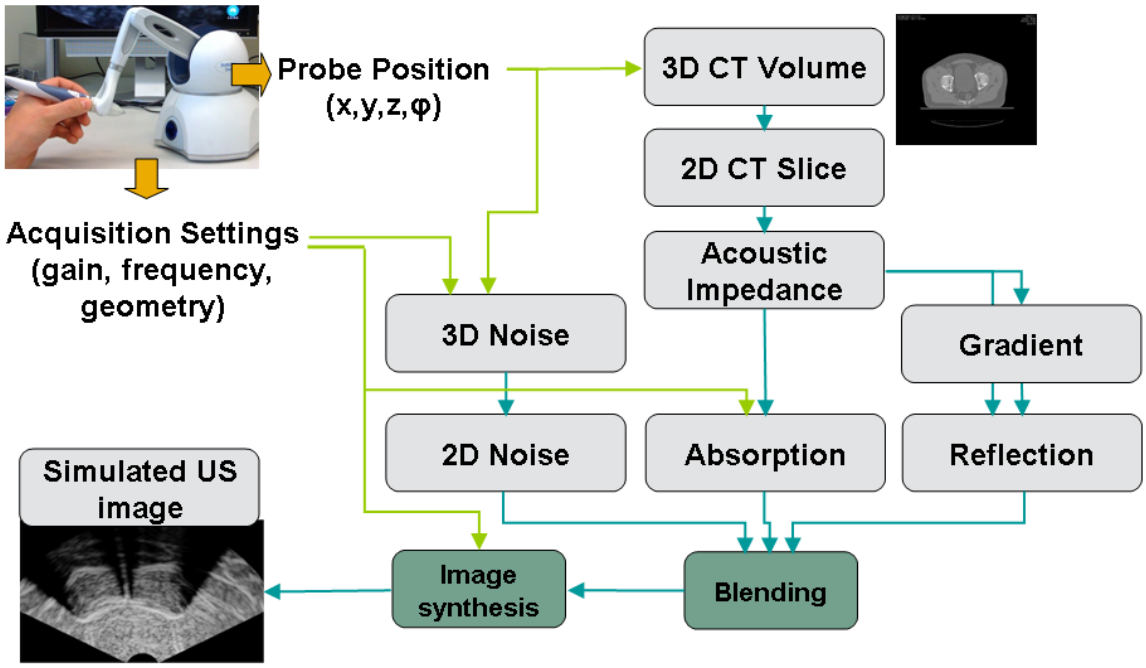

Figure 1. From 3-D CT data, 3-D noise, probe position, and acquisition settings an US image is simulated.

3.2. Reflection

It can be shown[12] that as in optics, Snell's law applies to ultrasound, and the intensity reflected at a specular interface can be computed as

$$\frac{I_r}{I_i} = \left(\frac{Z_2 \cdot \cos\theta_i - Z_1 \cdot \cos\theta_t}{Z_2 \cdot \cos\theta_i + Z_1 \cdot \cos\theta_t}\right)^2 \overset{\theta\approx 0}{\approx} \left(\frac{Z2 - Z1}{Z2 + Z1}\right)^2 \tag{1}$$

where Z_1 and Z_2 are the acoustic impedances of the two types of tissue and θ_i and θ_t are the angles of incidence and transmission, respectively. This equation is usually simplified to the case of perpendicular incidence, where no or unreliable information about the local normal exists.

However, we're only interested in intensity that is reflected back to the transducer and in the general case we consider diffuse rather than specular reflection. Assuming ideal diffuse reflection, the same intensity is reflected in all directions across a hemisphere in 3-D (respectively a semicircle in 2-D). A general model of diffuse reflection is Lambert's cosine law

$$\frac{I_r}{I_i} = \cos(\theta) \tag{2}$$

where θ is the angle between the incident radiation and the surface normal. Usually, a combination of diffuse and specular reflection is approximated by using

$$\frac{I_r}{I_i} = \cos(\theta)^n \tag{3}$$

with $n = 1$ for a perfect diffuse reflection and $n > 1$ for a combination of specular and diffuse reflection. Specular reflection does occur in US images, e.g. at tissue-bone interface, but the scale of inhomogeneities relevant

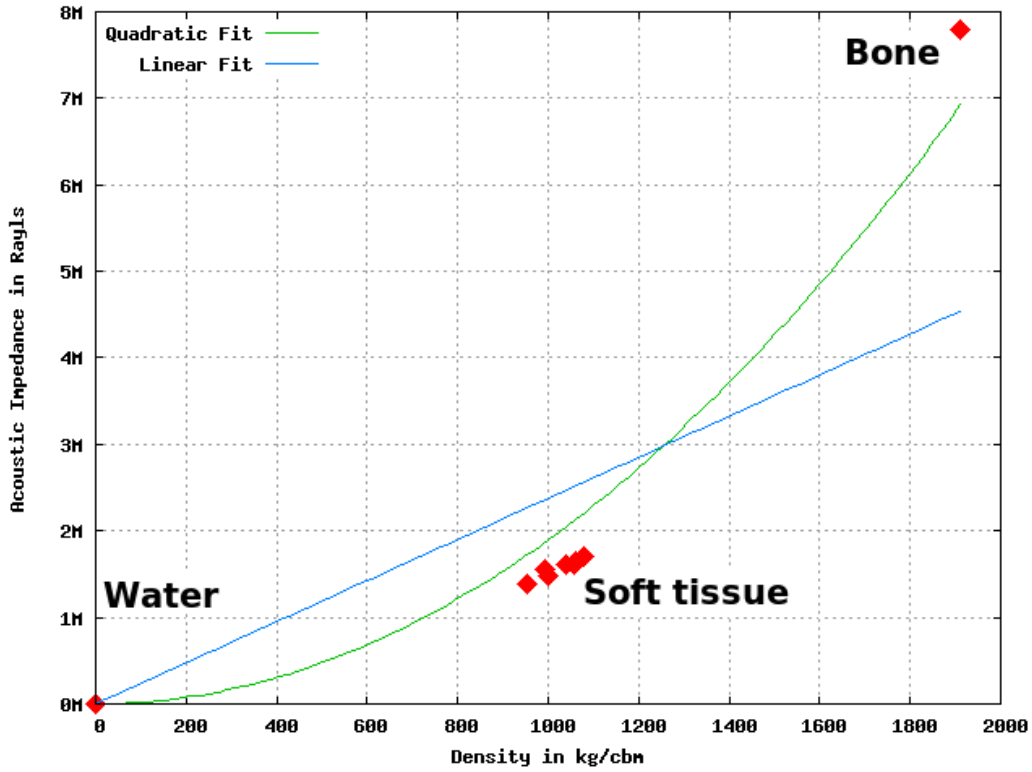

Figure 2. Known values for tissue density and acoustic impedance of different tissue types (water, , with linear and quadratic fit functions, values from Schneider et al.[11]

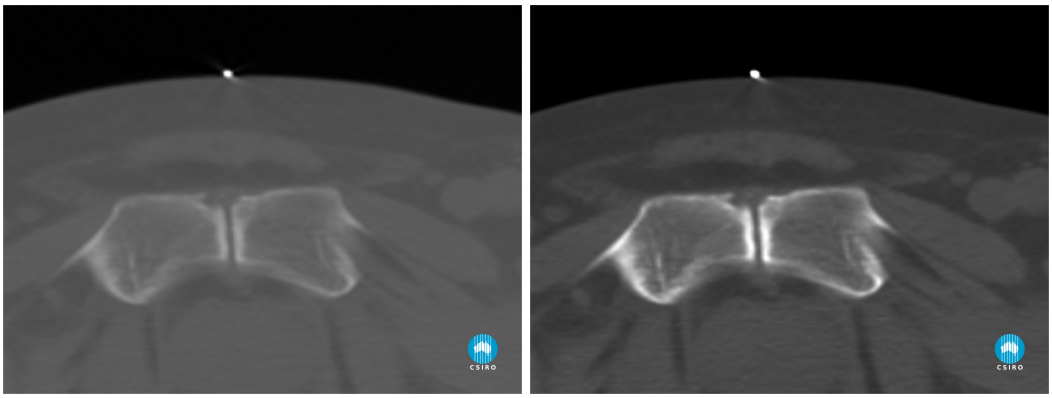

Figure 3. Screen shot of our original CT data (left) and estimated acoustic impedance (right).

to US reflection (wavelength approx. 0.6 to 0.15 mm) is beyond the usual resolution of abdominal CT images (e.g. 0.88 mm), so we may choose $n = 2$, as it further simplifies the equations.

We combine $\cos(\theta)^2$ with equation 1, as proposed by Wein et al.[6] Substituting the cosine with the dot product of a unit vector d in ray direction and the normalised gradient vector, we derive

$$\frac{I_r}{I_i}(x) \overset{est.}{=} \cos(\theta)^2 \left(\frac{\|\nabla Z(x)\|}{2 \cdot Z(x)}\right)^2 = \left(d^T \cdot \frac{\nabla Z(x)}{\|\nabla Z(x)\|}\right)^2 \left(\frac{\|\nabla Z(x)\|}{2 \cdot Z(x)}\right)^2 = \left(\frac{d^T \cdot \nabla Z(x)}{2 \cdot Z(x)}\right)^2 \tag{4}$$

where $d^T \cdot \nabla Z(x)$ equals the gradient magnitude along the ray, which can be easily computed in the simulated 2-D US scan plane. Note that there is no longer information needed about the original 3-D gradient, so all subsequent computation can be done in 2-D.

3.3. Absorption

Absorption, i.e. energy transfer into a localised heating of the tissue, accounts for nearly all US attenuation in soft tissue, at least 95% according to Lyons et al.[13] Ultrasound absorption can be characterised by an exponential law, similar to X-ray attenuation.

$$\frac{I}{I_0} = e^{-\beta x} \tag{5}$$

Often this is normalised to unit length as $\alpha = -10 \cdot \log_{10}(I/I_0)$, where α is the attenuation in Decibels per unit length. The relationship between α and US frequency is approximately linear over the range from 200 kHz to 100 MHz.[14,15] Attenuation for different types of soft tissue usually varies between 0.6 and 1.2 dB/cm at 1 MHz, with bone ranging up to 20 dB/cm. Typical absorption coefficients for various media are known, so we estimate absorption coefficients by an interpolation from the values of air, water, and bone, and thus the fraction of intensity absorbed in tissue along a known distance d can be estimated for every pixel as

$$\frac{I}{I_0} \overset{est.}{=} 10^{-\alpha \cdot d \cdot f/10} \tag{6}$$

where f is the currently used US frequency.

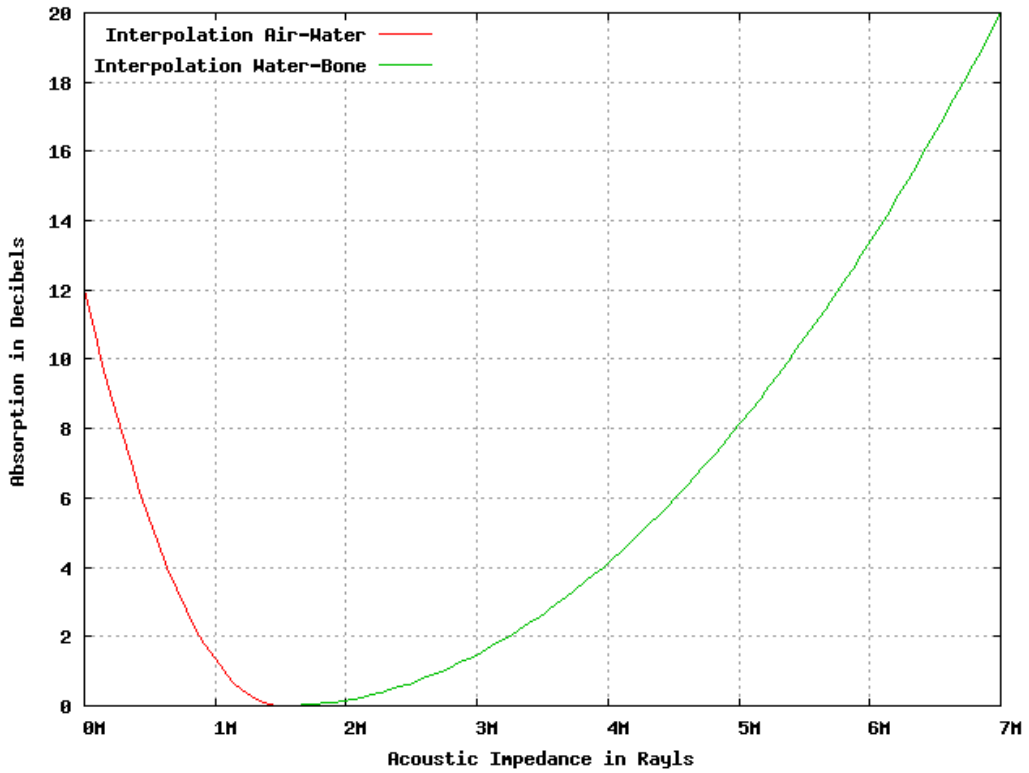

Figure 4. Mapping from acoustic impedance to ultrasound absorption, separately interpolated between air and water resp. water and bone, using a quadratic interpolation each.

Also the transducer will be able to detect more reflected intensity, if it is closer to the point of reflection. In addition to our tissue-dependent absorption we also model via the solid angle the distance-dependent ratio of energy that escapes the transducer.

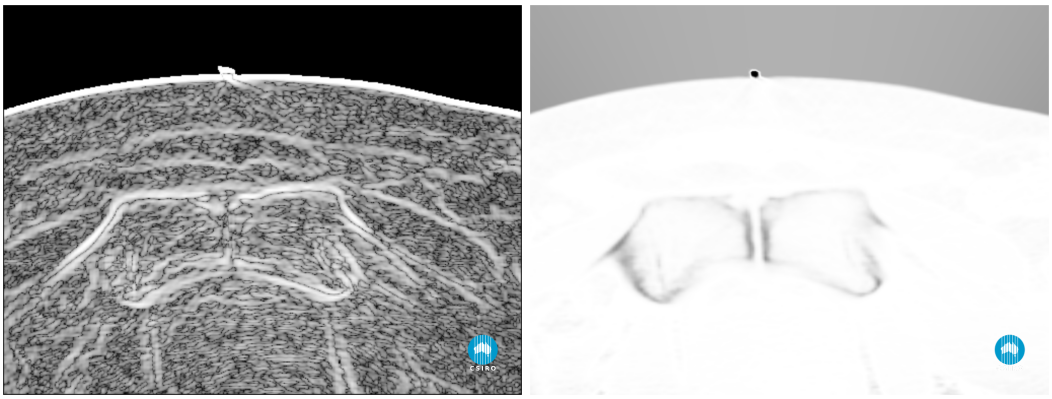

Figure 5. Screen shot of estimated reflection ratio (left) and estimated absorption ratio (right). For example, bones provide strong reflections, visible as bright borders in the left image, but also absorb intensity, visible as dark areas in the right image.

For a screen shot of estimated reflection and absorption ratios cf. figure 5.

3.4. Transmission

To simulate the image, we track US propagation through the tissue, this being is the only step that can not be done for each pixel separately. Initially, the US instensity at the location of the probe is set. We track the amount of intensity transmitted along each column of the image, and for every pixel in the column we compute the amount of intensity that is reflected or absorbed (based on the propagation characteristics explained above). After that we subtract reflection and absorption from the incident intensity at that pixel and step to the next. The same attenuation is assumed for the reflected signal on its way back to the transducer.

Finally, the estimations for reflection, transmission and absorption can thus be combined to estimate the reflected intensity for any depth x along the scanline in the ultrasound image:

$$I(x) \overset{est.}{=} I_0 \cdot \left(\prod_{0 \le k < x} \underbrace{\left[\frac{I_t}{I_i}(k) \right]}_{\text{transmission}} \underbrace{\left[1 - \frac{I_a}{I_i}(k, d) \right]}_{\text{absorption}} \right)^2 \cdot \underbrace{R(x)}_{\text{reflection}} \tag{7}$$

For a screen shot of transmitted and reflected intensity cf. figure 6. The "blending" step is illustrated in figure 7. Special care has to be taken for the absorption, because it depends on the length of tissue d between the evaluation points. Equation 7 assumes that the evaluation points on one ray are regularly spaced – this should be the case in all implementations based on any rectangular (pixel) grid.

In our implementation we use one CUDA thread per column, the currently transmitted intensity is maintained as one row in shared memory, and all threads are synchronised before and after updating those values, because in the general case intensity is interpolated from two pixels in the row above.

So far, we've synthesised an ultrasound reflection image, and in post-processing steps we add artifacts characteristic for ultrasound images, namely speckle noise and blurring.

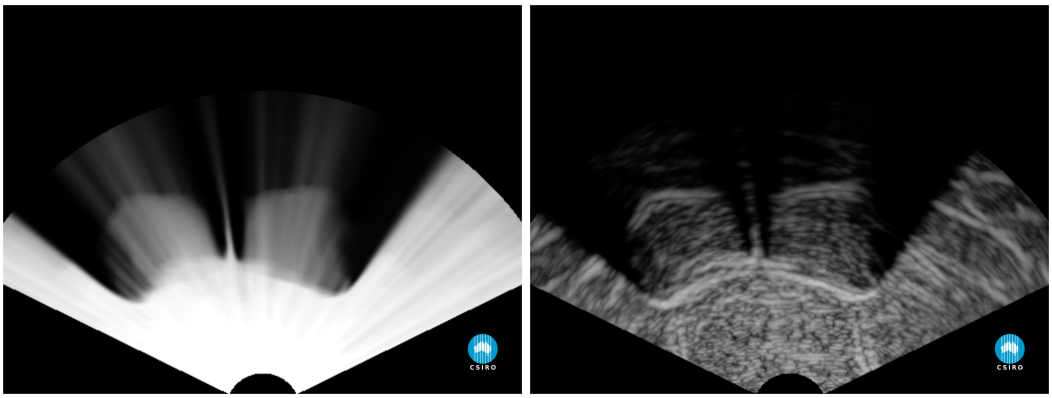

Figure 6. Screen shot of estimated transmitted intensity (left) and estimated reflected intensity (right).

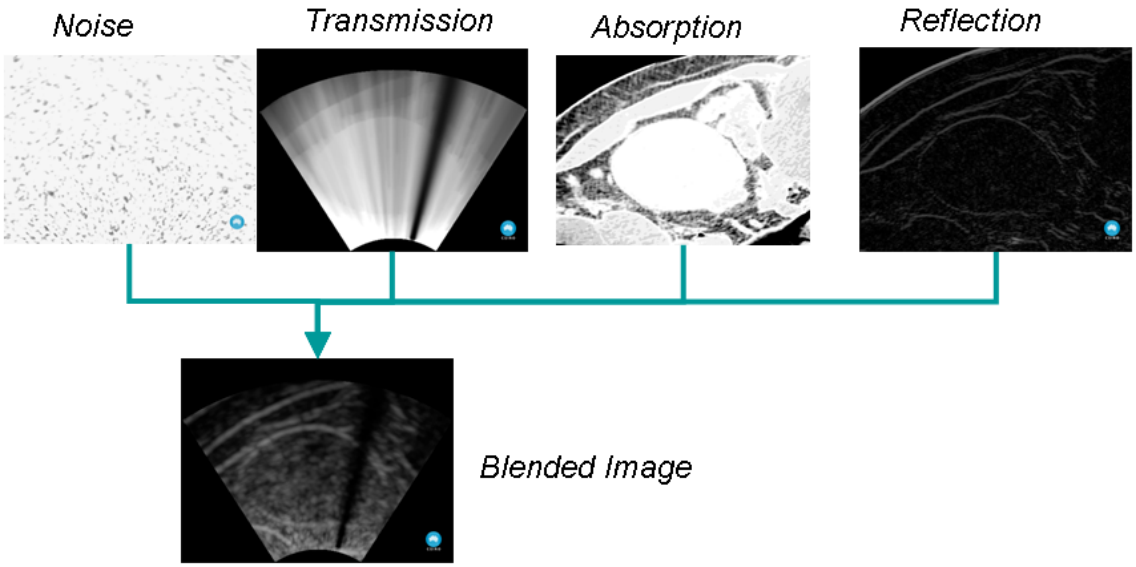

Figure 7. Integration of noise, transmission, absorption, and reflection in "blended image".

3.5. Speckle noise and blurring

A "fully developed speckle pattern", i.e. pure diffuse scattering with many (i.e. > 10) scatterers per resolution cell, can be well characterised by a Rayleigh distribution.[16,17] The additive component is generally accepted to be negligible compared to the multiplicative component. And as commonly done in work dealing with the filtering of speckle noise, our definition above implies that the noise is uncorrelated, i.e. we each pixel is distorted independently. This allows us to pre-compute the noise in the form of a 3-D "noise volume". Using a sufficiently large size, e.g. $128 \times 128 \times 128$ pixels, this texture can be used for the whole image without any noticeable repetition.

Speckle noise is added in a post-processing step: we use GLSL to extract a 2-D slice from the 3-D Rayleigh noise texture and CUDA for blending. Radial blurring (in lateral and axial direction) is added using CUDA again.

3.6. Refraction

Refraction is not taken into account so far. However, the framework can be extended to also model refraction as soon as hardware support for atomic functions (CUDA compute capability 1.1) will be available on NVIDIA's more advanced graphics processors.

3.7. Log-compression

To amplify small values a logarithmic compression is be applied, according to the following function:[6]

$$r_a(x) = \log(1 + a \cdot x) * \log(1 + a) \tag{8}$$

In order to provide a more meaningful parameter value, a is computed from an input value b in Decibels as $a = 10^{b/10}$. Usually input values in the range of 50 to 70 dB yield appropriate results. Like with commonly used US machines the amount of compression is variable for different imaging depths via the so-called "time gain compensation".

4. RESULTS

We benchmarked our implementation by simulating an ultrasound image of size 512×384 pixels (aspect ratio 3:4) from a raw CT volume of size $512 \times 512 \times 148$ voxels on a work station equipped with a GeForce 8800 GTX, Dual Intel Xeon 3 GHz, and 2 GB RAM. For detailed timing information cf. figure 8 and tables 1 and 2.

Note that transfers between CPU and GPU amount to 4.08 ms, because as of CUDA 1.1 transfers between OpenGL and CUDA use host memory, but in total we met the requirement of usually 30 frames per second for an interactive simulation.

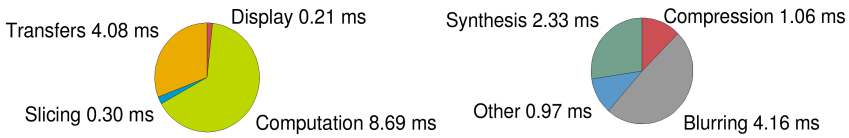

Figure 8. Timing overview for all simulation steps (left) and a breakdown of the "Computation" step on GPU (right).

Our ultrasound simulation has been integrated with a haptic input device and force feedback. One possible application is simulating a transrectal ultrasound examination, as needed for e.g. ultrasound-guided needle biopsy of the prostate. For a visual comparison of our simulation with a typical image from a real examination cf. figure 9. As visible in the image, we provide a realistic simulation of US reflection and artifacts like shadows, attenuation, speckle noise, and radial blurring. A more general abdominal ultrasound examination has also been implemented.

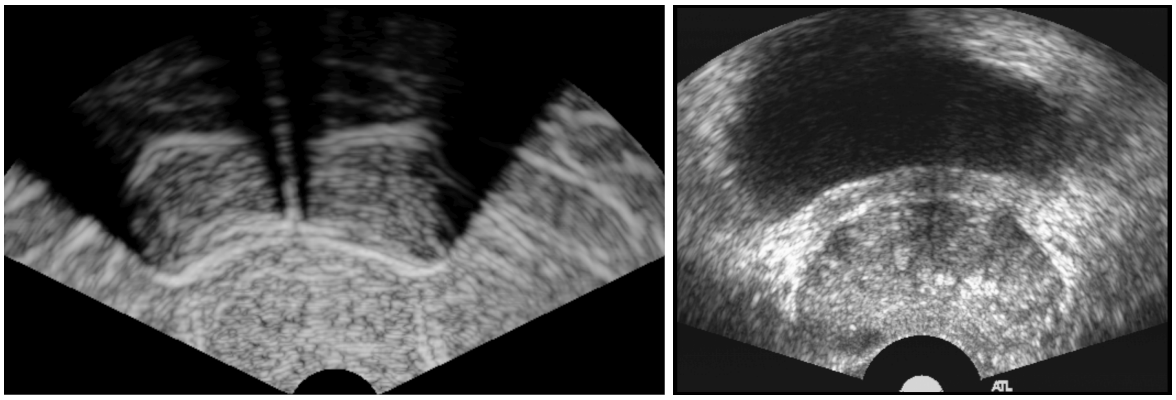

Figure 9. Screen shot of our ultrasound simulation (left), configured like a transrectal ultrasound probe, with a similar image from a real transrectal ultrasound examination of a different patient (right, image courtesy of Robarts Research Institute). For display purposes, both images have been manually and identically brightened.

Slicing	0.300 ms
Transfer OpenGL-host	1.864 ms
Transfer host-CUDA	1.114 ms
Computation on GPU	8.690 ms
Transfer CUDA-host	0.302 ms
Transfer host-OpenGL	0.796 ms
3-D Rendering	0.208 ms

Table 1. Timing information for our ultrasound simulation.

Impedance Estimation	0.067 ms
Gradient Computation	0.617 ms
Propagation	0.152 ms
Synthesis	2.325 ms
Blurring	4.157 ms
Stenciling	0.072 ms
Log-compression	1.060 ms
Data Scaling	0.060 ms

Table 2. Timing information for the image computation on GPU.

5. DISCUSSION

The achievable resolution for commonly used US frequencies is limited by the wavelength $\lambda = c/f$, and for clinically relevant frequencies on the order of 0.2 mm. The achievable resolution of the US simulation is strongly dependent on the resolution of the underlying CT data, and thus in order to increase the resolution of the simulation CT with a higher resolution will be needed.

The generation of US images from a geometric model, which can be a segmentation of CT data, and synthetic 3-D textures, which can be based on real US images, yields convincingly textured organs, as shown by Zhu et al.[5] However, this simulation can only be as detailed as the underlying model or segmentation, respectively, so smaller structures are not included in the simulated, if they have not been segmented.

Also, we spend a significant amount of time for blurring the image, cf. figure 8, so there might be potential for further optimisation.

6. CONCLUSIONS

We modified an existing physically plausible approach to simulate US images so that computation is entirely feasible within the virtual US scan plane and with only one slice from the original 3-D CT image. We present the design of an implementation in CUDA, which provides an acceleration compared to previous CPU-based approaches. No manual annotation or adjustment of the input data is necessary, and for a simulation of US images with the same resolution (512 pixels) as commonly used CT images 8.7 ms are needed for the image computation step on GPU.

Ultrasound reflection, attenuation, shadowing artifacts, speckle noise, and radial blurring are reproduced. To our knowledge no method of simulating ultrasound absorption from CT data has been published before beyond "counting bone pixels" as implemented by Zhu et al.[5] or Vidal et al.[9] Like a real probe, all acquisition parameters can be interactively changed during the simulation, including US frequency, US intensity, time-gain compensation, and field geometry, as well as speckle size and radial blurring. A wide range of different probe types can be simulated and our simulation framework is generally applicable to training of different US examinations or US-guided procedures.

7. ACKNOWLEDGMENTS

This work has been partially supported by the Deutsche Forschungsgemeinschaft grant NA 620/2-1.

REFERENCES

1. A. Ziv, P. R. Wolpe, S. D. Small, and S. Glick, "Simulation-based medical education: an ethical imperative.," *Academic Medicine* **78**, pp. 783–788, August 2003.
2. J. A. Jensen, "FIELD: A program for simulating ultrasound systems," *Medical & Biological Engineering & Computing* **34**, pp. 351–353, 1996.
3. D. Aiger and D. Cohen-Or, "Real-time ultrasound imaging simulation," *Real-Time Imaging* **4**, pp. 263–274, 1998.
4. A. Hostettler, C. Forest, A. Forgione, L. Soler, and J. Marescaux, "Real-time ultrasonography simulator based on 3D CT-scan images," *Medicine Meets Virtual Reality 13* **111**, pp. 191–193, 2005.
5. Y. Zhu, D. Magee, R. Ratnalingam, and D. Kessel, "A virtual ultrasound imaging system for the simulation of ultrasound-guided needle insertion procedures," in *Proceedings of Medical Image Understanding and Analysis*, 2006.
6. W. Wein, S. Brunke, A. Khamene, M. Callstrom, and N. Navab, "Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention," *Medical Image Analysis* **12**, September 2008. In Press.
7. R. Shams, R. Hartley, and N. Navab, "Real-time simulation of medical ultrasound from CT images," in *MICCAI*, *Lecture Notes In Computer Science* **5242**, pp. 734–741, Springer Berlin / Heidelberg, (New York, USA), September 2008.
8. J. A. Jensen and S. I. Nikolov, "Fast simulation of ultrasound images," in *Proceedings of IEEE International Ultrasonics Symposium*, 2000.
9. F. P. Vidal, N. W. John, A. E. Healey, and D. A. Gould, "Simulation of ultrasound guided needle puncture using patient specific data with 3D textures and volume haptics," *Computer Animation and Virtual Worlds* **19**, pp. 111–127, May 2008.
10. K. Raum, J. Reißhauer, and J. Brandt, "Frequency and resolution dependence of the anisotropic impedance estimation in cortical bone using time-resolved scanning acoustic microscopy.," *Journal of Biomedical Materials Research Part A* **71**, pp. 430–438, Dec 2004.
11. U. Schneider, E. Pedroni, and A. Lomax, "The calibration of CT Hounsfield units for radiotherapy treatment planning.," *Physics in Medicine and Biology* **41**, pp. 111–124, January 1996.
12. K. Shung and G. Thieme, *Ultrasonic Scattering in Biological Tissues*, CRC Press, 1993.
13. M. Lyons, R. Chivers, and K. Parker, "Absorption dominates attenuation in soft tissues," in *Proc. IEEE 1986 Ultrasonics Symposium*, R. Chivers, ed., pp. 871–874, 1986.
14. R. Kuc, "Clinical application of an ultrasound attenuation coefficient estimation technique for liver pathology characterization," **BME-27**(6), pp. 312–319, 1980.
15. D. Dowsett, P. Kenny, and R. Johnston, *The physics of diagnostic imaging*, Chapman & Hall Medical, 1998.
16. R. Wagner, S. Smith, J. Sandrik, and H. Lopez, "Statistics of speckle in ultrasound B-scans," *IEEE Transactions on Sonics and Ultrasonics* **30**(3), pp. 156–163, 1983.
17. D. Kaplan and Q. Ma, "On the statistical characteristics of log-compressed Rayleigh signals: theoretical formulation and experimental results," in *Proc. Ultrasonics Symposium IEEE 1993*, Q. Ma, ed., pp. 961–964 vol.2, 1993.