

Chapter 7

Optimizing Clinical Performance of a Time-Domain Breast Imaging System

The results of the system characterization presented in Chapter 6 are complimented here through the assessment of issues related to signal processing and limitations to measurements within the constraints of a clinical setting. The results of several phantom tests are explored and issues related to optimizing system performance are discussed. Next, preliminary image results are presented along with suggested improvements. The limitations to allowed signal power are defined with respect to government imposed limitations of skin exposure. Concerns related to localization of tissue boundaries in a clinical setting are explored. The trade-off space between the number and positions of sources and gate delays is motivated and discussed. Finally, the optimal measurement sequence and limitations on the total number of measurements as limited by patient tolerance are explored.

7.1 Background Considerations

The first step in processing the raw image file is to subtract off the background. A typical background image is shown in image (A) of Figure 7.1 for a gain voltage of 600 Volts and an integration time of 27 msec, indicating an offset on the order of 45 out of 4096 counts. The difference between two successive background images is shown in image (B) of Figure 7.1, indicating that the fluctuation in the background was substantially randomly distributed. A small degree of fixed-pattern noise was evidenced, however.

The method of background subtraction assumed by the acquisition program was found to be inadequate. The method inhibited the high voltage of the MCP, which essentially acquired a single image frame at the lowest gain of 260 Volts. In essence, this only gives the background for the CCD. There are three problems with this approach. First, at least 10 to 30 background frames should be averaged, in general, so that the shot noise in the photocathode thermionic emission and the CCD readout noise do not add




Figure 7.1 A typical background image is shown for a gain voltage of 600 Volts and an integration time of 27 msec on the left in image (A). The difference between two successive background images is shown in image (B), indicating that the fluctuation in the background is substantially randomly distributed.

significantly to reducing the SNR. Secondly, for MCP gain voltages above 600 Volts, the background due to thermionic emission is sufficiently high that it needs to be subtracted off to prevent nonlinear response, as described in Chapter 6, Section 6.1.9. Thirdly, there was generally some level of background incident on the photocathode from ambient light conditions not accounted for, which would also lead to nonlinear response. Figure 7.2 illustrates this point quite dramatically. The image in (A) of Figure 7.2 was acquired at an integration time of 100 msec and a MCP gain voltage of 800 Volts. The room lights were on during the measurement, but the probe was covered with several layers of black cloth.

It was also possible that some light could have coupled in through the optical fiber bundles between the probe and the camera. Aluminum foil had been placed over the most sensitive section of the fiber bundle, within 1 meter of the camera, to minimize ambient light coupling. This illustrates the sensitivity of the instrumentation to ambient light.

Image (B) of Figure 7.2 shows the background image acquired by the acquisition routine.




Figure 7.2 The raw image file in (A) was taken at 800 Volts MCP gain and 100 msec integration time with a cloth over the phantom stand. Image (B) was the background image acquired by the software, which inhibited the gain. Clearly the use of image (B) would result in nonlinear performance, as it only represents the background of the CCD, not the ICCD as a system.

Notice that a very significant error would result, if the image in (B) were used in place of the actual background of image (A). Again, background from ambient light could be dramatically reduced by incorporation of an appropriate interference filter into the camera system to reject out-of-band photons. Even with a filter, however, there would be background from the thermionic emission as described in Chapter 6 that should be accounted for. Thus, it would always a good idea to acquire a background image of the complete ICCD, not just the CCD.

Generally, there is not a significant penalty to be paid, with respect to measurement time, for acquiring a statistically relevant number of background frames for the case of a

constant ambient background. Changes in ambient light conditions over the period of a measurement can be problematic, as previously noted. Therefore, it may be prudent or necessary to utilize DC lighting near the measurement. One could use light shields or opaque cloths to reduce the background effects. Caution should be exercised, however, as the boundary conditions can easily be altered by differences in the position of the cloth with respect to the sides of the tissue or tissue phantom. This could lead to deviations from the assumptions of the forward model, and consequently, errors in the reconstructed optical properties.

Caution must also be exercised for the case of tissue or tissue phantoms with low absorption coefficients. In that case, the laser's pulse-to-pulse period of 12.5 nanoseconds may not be long enough to allow the scattered field to attenuate sufficiently. The tail of the TPSF could bleed over into the rising edge of the TPSF, distorting the signal for early time gates as illustrated in Figure 7.3. If left uncorrected for, this would result in errors to

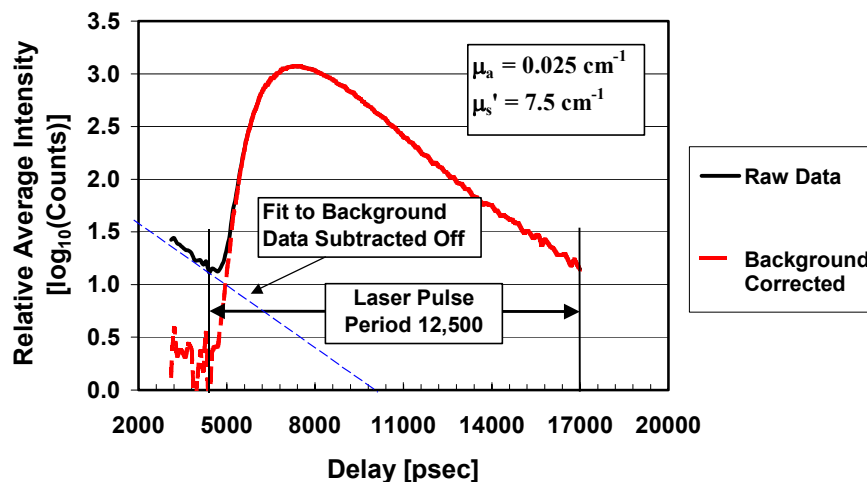


Figure 7.3 A background-corrected TPSF is shown for a low absorption phantom. Such correction procedures may be important in the case of breast imaging with compression plates, for which absorption coefficients approaching 0.025 cm^{-1} are not unexpected.

the fit to the forward model and thus errors in the reconstructed image of the optical properties. As shown in Figure 7.3, for a background μ_a of 0.025 cm^{-1} and μ_s' of 7.5 cm^{-1} , the slope of the asymptotic tail of the TPSF could be used to provide a background fit for correcting the data. For the case of breast imaging using compression plates, such low absorptions are not unexpected, thus appropriate care must be exercised.

7.2 Phantom Measurement Analyses

The heterogeneous phantom and the two homogeneous phantoms described in Chapter 4 were used extensively to acquire data sets to provide further insight into the performance and characterization of the Time-Domain Optical Breast Imaging System. The goal was to use this insight to develop recommendations for improving clinical performance.

The use of the phantoms allowed for further investigation into the character of the TPSF's under various conditions. For example, Figure 7.4 shows the effect of

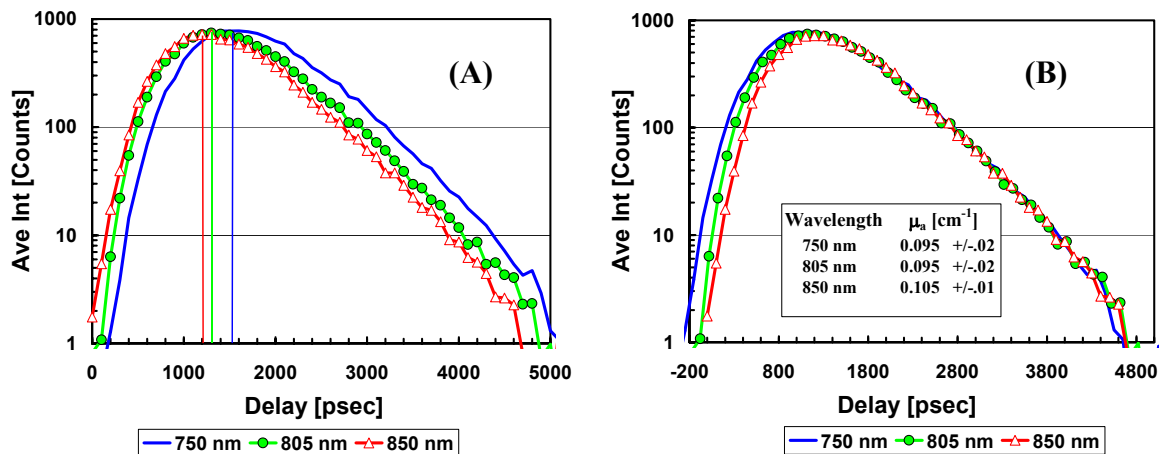


Figure 7.4 The TPSF's for a central source fiber on the high absorption phantom are shown in plot (A) as a function of wavelength. When the TPSF's are shifted, as in the plot (B), to superimpose for long delays, we see that the FWHM changes along with the delay to peak intensity, shown in (A), indicating the change in scatter with wavelength. The change in absorption with wavelength is minimal in comparison, as indicated by the similar slope of the asymptotic tails.

wavelength on the TPSF's for the high absorption homogeneous phantom. Plot (A) of Figure 7.4 shows the raw untouched TPSF's. It can be seen that the 805 and 750 nm TPSF's are disproportionately delayed in time with respect to the 850 nm TPSF, indicating that the scattering coefficient increases nonlinearly with decreasing wavelength, as would be expected from the use of 30 nm APS TiO₂ particles. In plot (B) of Figure 7.4, the 750 nm TPSF was shifted by 450 psec and the 805 nm TPSF by 180 psec to the left in order to superimpose their asymptotic tails. The similar asymptotic slopes show that the absorption coefficients were nearly the same for all three wavelengths, but the width of the TPSF's increased with decreasing wavelength. Thus, it can be seen that increased scattering coefficients lead to increased TPSF widths and time to maximum intensity as expected from the theory presented in Chapter 3.

Next, the effect of the boundaries of the tissue phantoms on the calculated values of the background optical properties was explored. The phantoms were made in a precise manner as described in Chapter 4, and were assumed homogeneous throughout. The most direct approach to determining the absorption coefficient μ_a was from calculation of the slope of the asymptotic tail. As described in Chapter 3, once the source term has had an opportunity to spread out over the entire slab geometry and is no longer localized or growing, the signal at any point of detection on the surface of the slab would decay at a rate proportional to $\exp(-\mu_a v)$, where again, v is the speed of the electromagnetic wave in the medium. As explained in Chapter 6, a MatLab routine was written to select the 148 source and detector pairs directly across from each other and calculate the absorption coefficient for each source location from the tail of the TPSF. A 2-dimensional absorption map was produced as shown in Figure 7.5. Any pixel at which no source was

located was set to the mean value of the full data set to improve visualization. The data was acquired at a MCP gain voltage of 500 Volts and 50 msec integration time. Again, as the photocathode and intensifier gain saturation effects were not known at the time of the data collection, the data was not corrected for saturation effects, as no calibration was

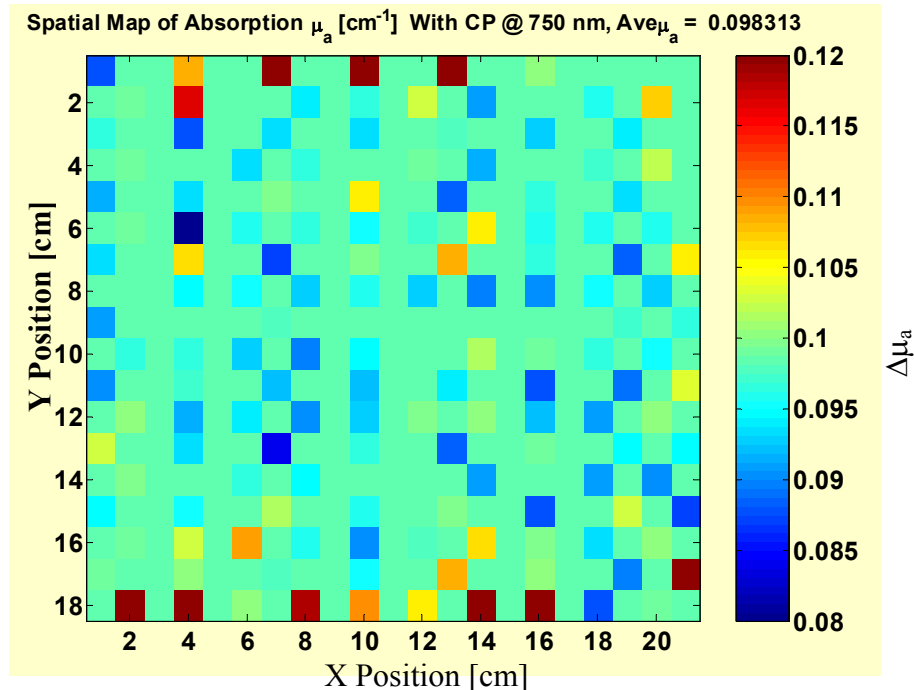


Figure 7.5 The absorption map for the high absorption homogeneous phantom at 750 nm is shown for all source positions and the detector position immediately opposite the source. The average absorption coefficient μ_a was shown to be 0.098. A slight increase in calculated absorption was observed between the center and the edge of the phantom.

acquired. The saturation effects would likely be small, however, as a raster scan through the sources was used, which as described in Chapter 6, minimizes saturation effects.

The most important point to note in the absorption map of Figure 7.5 is the general increase in calculated absorption in moving from the center to the edges (top and bottom) of the phantom. This weak, but general, trend was also observed for Figures 6.5 and 6.6 in Chapter 6. The edge source fiber positions were just 5 mm inside the edge of the phantom on both the top and bottom of the image. Light escaping from the boundaries,

the sides of the phantom, could account for the increased asymptotic slope near the edge. It is helpful at this point to take a look at the TPSF's themselves. Figures 7.6 and 7.7 show TPSF's of one source-detector pair near the center of the phantom and one near the central edge of one side for the same data as shown in the absorption map of Figure 7.5. The data also shows differences between the presence and absence of the compression plates. There is an observed difference in the absorption coefficient between the cases of with and without the compression plates. The use of compression plates resulted in a 12% decrease in measured absorption near the center and 5% near the edge relative to no compression plates. The shift of 50 psec was caused by the several millimeters of high refractive index polycarbonate that comprised the compression plates. The fall off in the absorption coefficient could be partially due to the use of the compression plates, as no

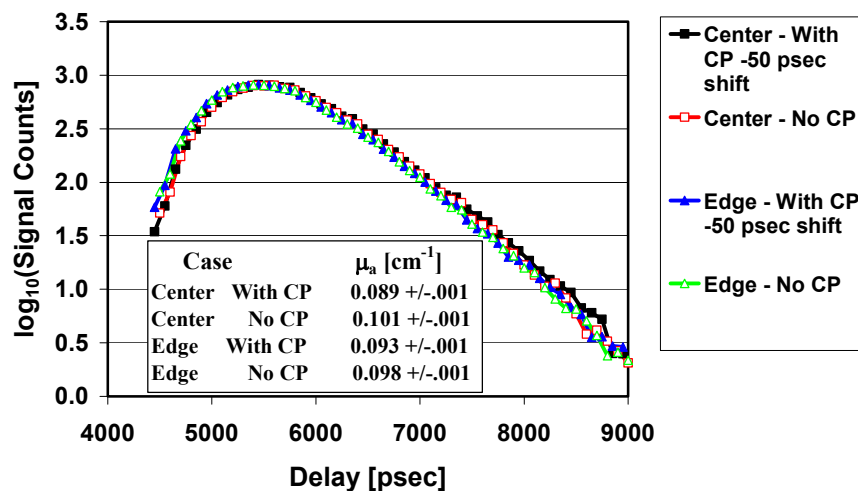


Figure 7.6 The TPSF's for a source-detector pair in the center of the phantom and a pair near the edge are shown for the cases of with and without the compression plates (CP).

significant difference was noted in the fall off toward the edge without compression plates for this high absorption phantom. There was also a shift in the delay observed between the center and edge for both with and without the compression plates. This

phantom had an increase in thickness of several millimeters toward the center, as there was no top on the mold during fabrication. This could account for the delay shift toward shorter times for the edge data. The phantom was machined down to have uniform thickness after these measurements were acquired. The increased signal for long delays for the case of with compression plates, may have been caused by light reflecting around the edge of the curved compression plates. The masking of the compression plates as discussed in Chapter 6, would be expected to reduce this cause to insignificance.

The high absorption phantom data shown above was taken with no black tape

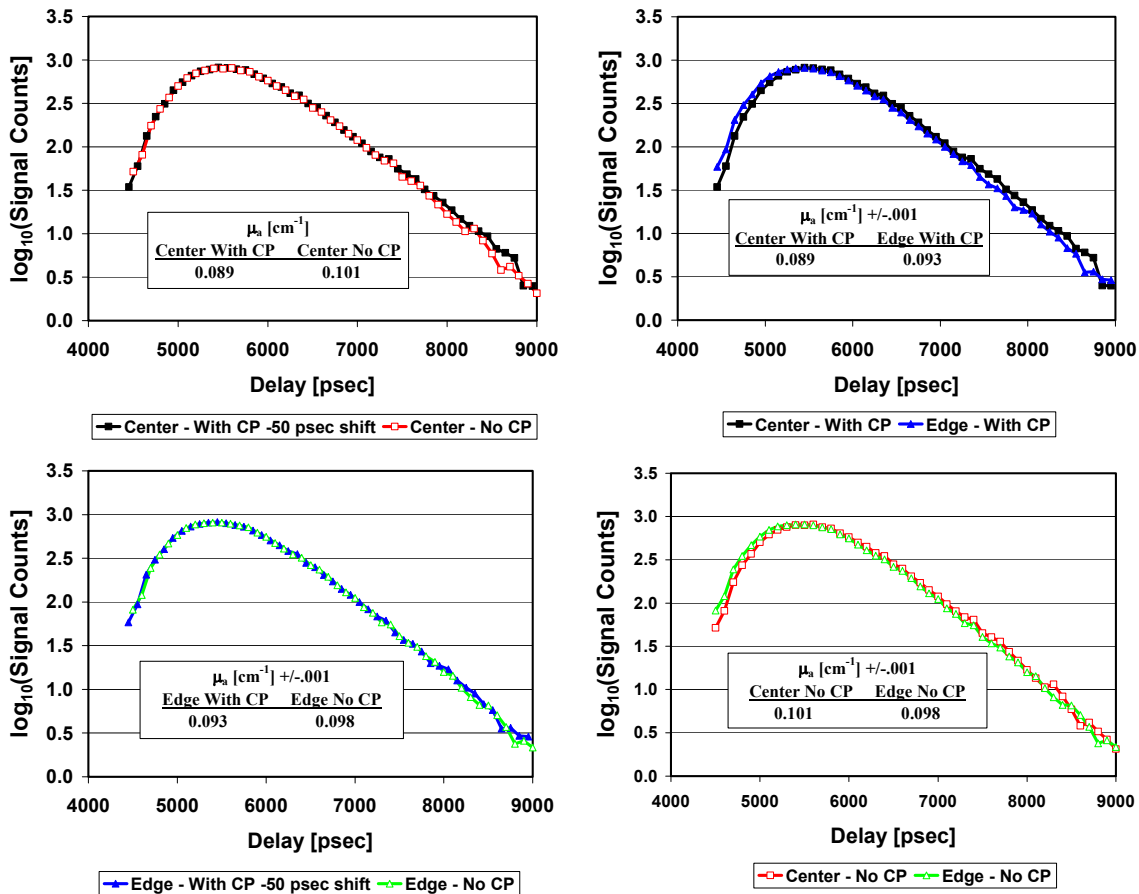


Figure 7.7 The four plots show the differences in TPSF's between the cases of with and without the use of the compression plates (CP). Also shown are the affects of moving from the center to the edge of the phantom.

surrounding the outside of the phantom. This resulted in sensitivity to reflections off the black cloth or sections of the phantom stand that could corrupt the data near the boundaries unless great care was exercised. For this reason, black tape was added to the outside of the phantoms. This reduced sensitivity to stray reflections shunting to the detection fibers near the outside, but did result in an increased change in the width of the TPSF between the center and edge fibers, which would be expected to lead to a decrease

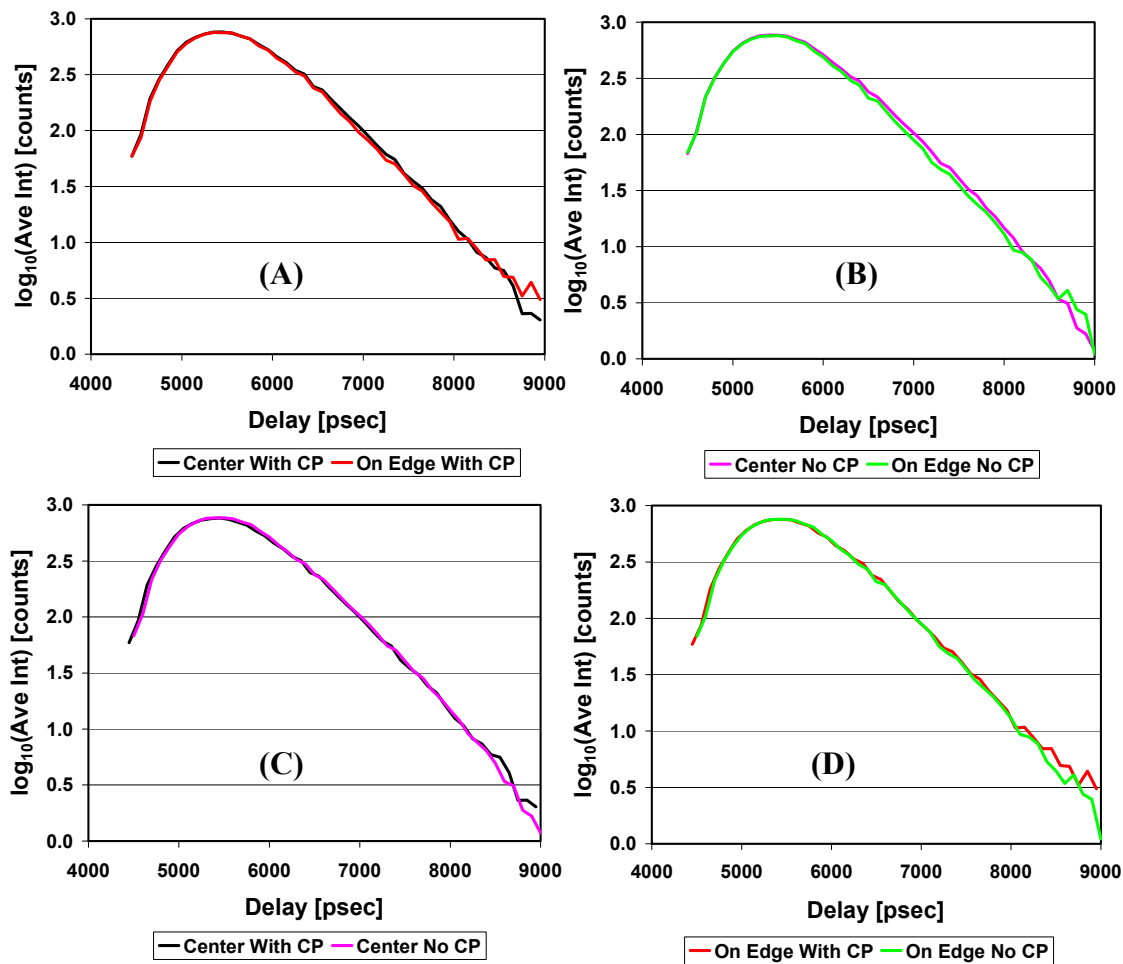


Figure 7.8 The TPSF's are shown for the high absorption homogeneous phantom, both with and without the compression plates (CP). The center measurements were comprised of an average of 5 centrally located source fiber positions. The edge measurements were comprised of an average of 4 source positions along the edge of the phantom. The edge of the phantom was covered by black tape.

in calculated scattering coefficient. These TPSF's are shown in Figure 7.8. The center fiber data comprised an average of 5 source fiber positions near the center of the phantom and the edge fiber data an average of 4 source fiber positions along the edge. This averaging was done to improve SNR. The data for the case with the compression plates was shifted by 40 psec to account for the path length difference due to the finite thickness of the compression plates. Plots (A) and (B) of Figure 7.8 illustrate the decrease in width of the TPSF's between the center and edge for the cases of with and without the compression plates, respectively. There was still an increased signal for long delays in the case of with compression plates that led to a decrease in the calculated absorption.

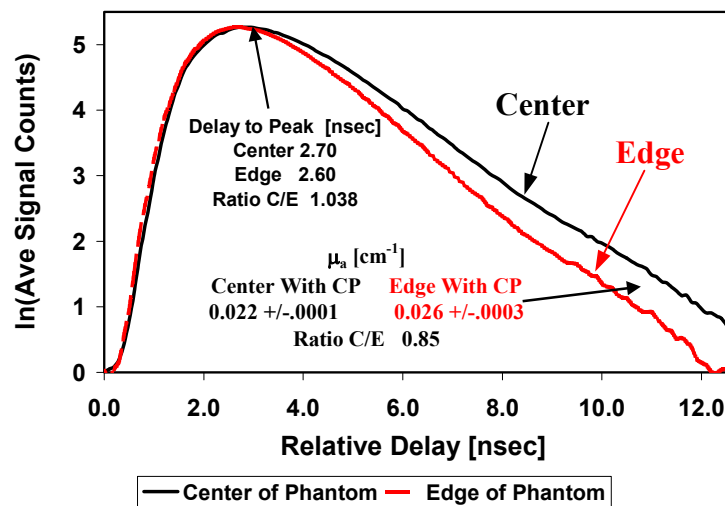


Figure 7.9 The average of 12 TPSF's each is shown for source-detector pairs near the center of the homogeneous phantom versus near the edge. Photons escaping out of the side of the phantom cause the signal to attenuate faster and thereby resulting in a higher absorption near the edge. The affect on the time to peak, which is primarily a function of the scattering coefficient, does not seem to be affected much,

The effect of the boundaries is much more pronounced for the low absorption homogeneous phantom TPSF's as shown in Figure 7.9. As stated above, this phantom also had black tape on its edges. The plot shows the natural log of the average intensity

versus the relative delay for the average of 10 fibers located near the center in comparison with an average of 10 located near the edge. The magnitude at the edge was scaled up to that of the center to make the comparison easier to interpret. The 4% shift in the delay to peak intensity would be expected to lead to reduced scattering coefficients toward the edge. We can understand the mechanism for the increased absorption and decreased delay by considering the boundary effect for a source detector pair near the edge of the phantom. Photons that strike the black tape on the boundary would be substantially absorbed. For the case of a semi-infinite slab, a portion of these photons would be scattered back and contribute to the signal at the detection fiber. This increased loss at the boundaries would account for the increased slope of the asymptotic tail. Apparently, this loss more than compensated for the slight shunting due to the compression plate that was observed near the edge for the case of the high absorption phantom. The effective decrease in the delay for the fiber source on the edge can be understood by considering that some of the photons that would have scattered into the region beyond the boundary would have contributed to the signal at the detection fiber. These photons would have had a long mean path and thus suppressing them leads to an apparent shift toward shorter delays when the peak magnitude of the edge source was scaled up to that of the center fiber. The 15% increase in the absorption at the edge was significant and would lead to an error in the reconstructed data unless accounted for in the forward model. This is one of the greatest shortcomings of our forward models, as they do not properly account for perturbations due to the boundaries.

The problem of incorporating boundary effects into any forward model is further complicated by the fact that the boundary conditions are so dependent on what lies

beyond the tissue boundary. It must be determined how to block stray light most effectively, but at the same time have the least impact on changing the boundary conditions from those assumed by the forward model. Not only should there be concern about the index mismatch at the boundaries, but there should also be concern about the effect of any materials placed beyond the boundary, which could either absorb or scatter the photon flux that passes through the boundary. Perhaps a bladder could be developed which contains a solution with background optical properties similar to breast tissue to be placed around the boundaries of the breast tissue. This bladder could have an opaque outer cover, sufficiently far from the boundary, to cut down on stray ambient light from entering the breast tissue.

Another boundary condition that may be problematic is that of the breast tissue sloping away from the compression plates near the boundary. This can lead to a gap between the tissue and the source and detector fibers. Generally, due to gravity, this problem would be most serious for the detector fibers that are located on the top. Two errors could be introduced. First, the field of view of the detection fiber could grow from a few millimeters to the order of a centimeter or more. Second, the tissue thickness would be different from the assumed inter-compression plate thickness. Both of these issues would lead to substantial errors in the reconstructed data. Again, a possible solution may be the conformable bladder, which could force the breast tissue up against the compression plates without any discomfort to the patient.

The absorption coefficient was quantified above for several different cases. The scattering coefficient was also characterized with respect to the effect due to the compression plates and the effect of moving from the center of the phantom toward the

edge. Figure 7.10 shows the geometry of groupings of sources and detectors that were used on the high absorption phantom to characterize the scattering coefficient. The groupings move from the center for (A) and increasingly approach the edge in moving toward group (E). The scattering coefficients, shown in Table 7.1, were determined by using the calculated absorption from the tails of the TPSF's, and a CW model fit to the integrated TPSF from the time-domain data. This approach led to a decrease in the scatter coefficient by about 25% in moving from center to edge. This method of calculating the scatter coefficient appears to have overestimated the change in scattering coefficient, in comparison with the TPSF's described above. The reason for this may be that the technique assumed a constant value for the absorption, which as was pointed out above, was not the case. The calculation showed an increase in the scattering coefficient of 7.4% with a standard deviation of 1.8% for use of the compression plates as compared to no compression plates. This also does not seem to be supported by the TPSF data of Figure 7.7. The shift in delay may not have been properly accounted for, which could lead to some of the difference between with and without compression plates.

In summary, the effect of the compression plates on the boundary for the case of determining the absorption coefficient is very small. The absorption effect could be further reduced by incorporation of a black absorption mask on both surfaces of the compression plates. The effect of the compression plates is also small when the TPSF's are compared directly. There was a difference noted in the scatter coefficients when the fitting routine was used, but it is likely that the model incorporated some systematic error that would account for the difference, as the TPSF's should speak for themselves.

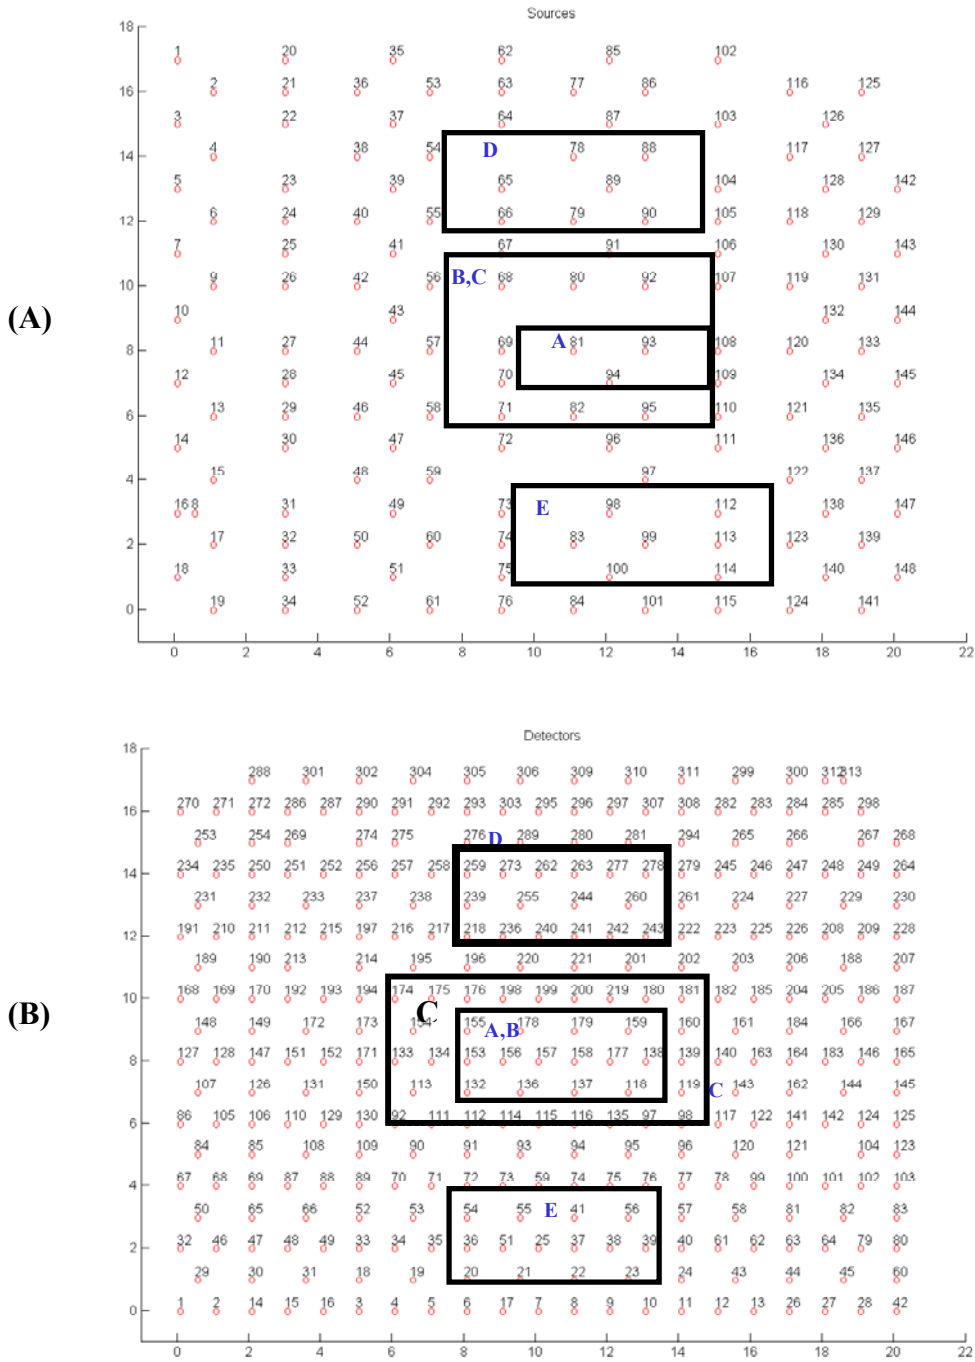


Figure 7.10 Plot (A) shows the groupings of sources and plot (B) shows the groupings of detectors used in mapping out the affect of source and detector position on the calculated homogeneous scattering coefficient.

Data Set	With Compression Plate	Without Compression Plate
A	7.90	7.32
B	7.76	7.07
C	6.52	6.03
D	6.18	5.88
E	5.99	5.64

Table 7.1 The scattering coefficients μ_s' [cm^{-1}] of a homogeneous phantom are shown as a function of the data set for both the cases of with and without the compression plates in place on the probe. The increasing letters represent data regions starting near the center with data set A, to near the edge for data set E. The case with the compression plate is increased from the case without the compression plate by an average of 7.4% with a standard deviation of 1.8%.

7.3 Preliminary Imaging Results

The heterogeneous phantom was used to collect a series of absorption data. A solution of India ink was mixed with water. Its absorption coefficient was determined by extinction measurements in a standard 1 cm path length spectroscopic cuvette. Next 500 ml of deionized water was added to a beaker and then combined with 17.3 ml of a 20% intralipid solution to achieve a scattering coefficient of 7.5 cm^{-1} at 800 nm. The beaker was then poured into the flow beaker of the dynamic heterogeneous phantom. Titrations of absorption were carried out from μ_a of 0.00 to 0.10 cm^{-1} in steps of 0.01 cm^{-1} . The ink absorption would be additive to that of the intrinsic absorption of the intralipid solution, so the actual absorption would have been slightly higher. The background optical properties of the heterogeneous phantom were calculated to be $\mu_a = 0.05 \text{ cm}^{-1}$ and $\mu_s' = 7.5 \text{ cm}^{-1}$ at 800 nm from TPSF data in a region of the phantom not containing any heterogeneities. Figure 7.11 shows a preliminary reconstruction of the data carried out by a colleague, Anand Kumar, indicating $\Delta\mu_a$ of the heterogeneity. At the time of this reconstruction, the saturation issues were not yet appreciated, so the data was not

corrected for saturation effects, but again due to the raster scanning, saturation effects were likely small. The data was collected at an MCP gain voltage of 500 Volts and an integration time of 50 msec. The reasons for using real background images for the background subtraction as discussed in Section 7.1, above, were not yet appreciated, so

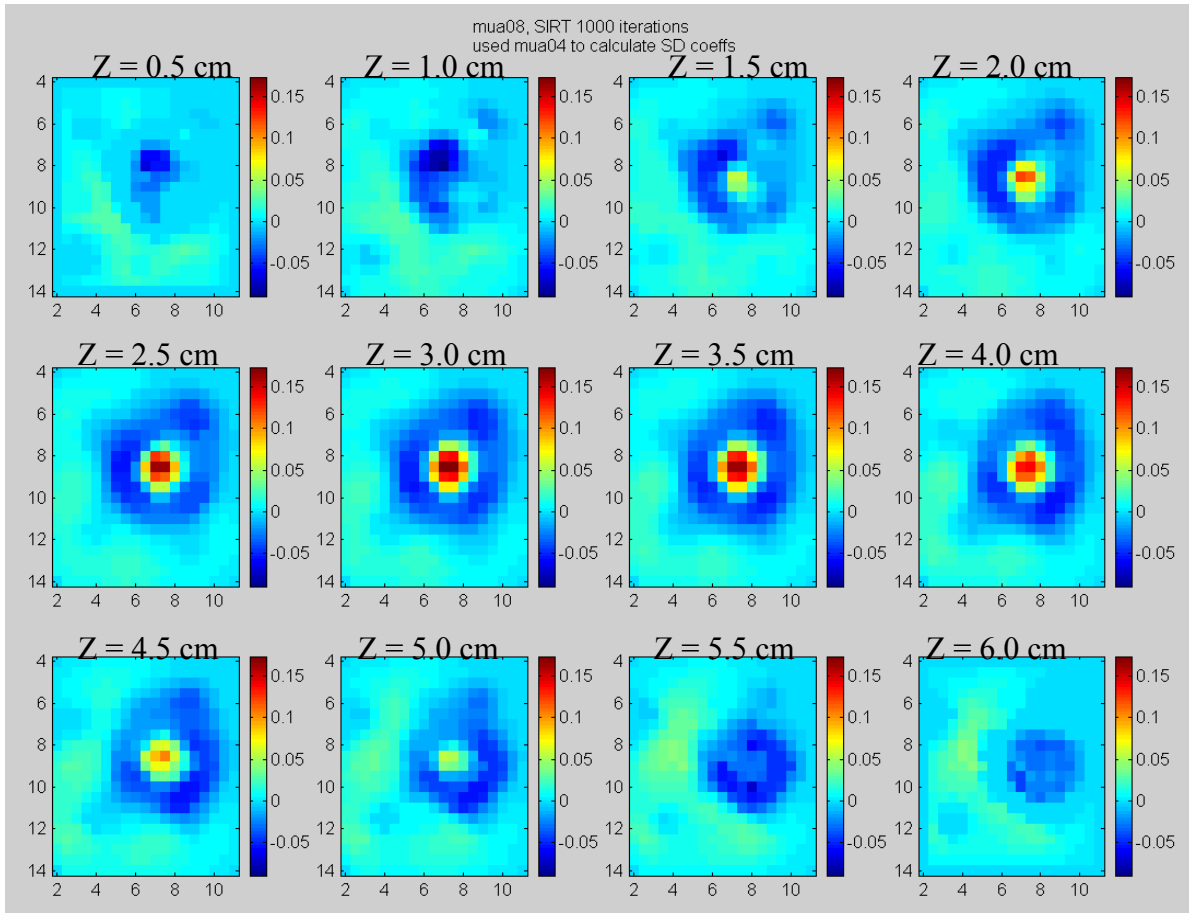


Figure 7.11 A preliminary image result is shown for an heterogeneous absorption phantom showing the image of a 1.5 cm diameter glass sphere at a depth of 3 cm. The background optical properties for the phantom were calculated to be $\mu_a = 0.05 \text{ cm}^{-1}$ and $\mu_s' = 7.5 \text{ cm}^{-1}$. The $\Delta\mu_a$ was 0.04 cm^{-1} , so it is apparent that it was overestimated by the 1000 SIRT iterations.

the standard background procedure assumed by the acquisition program was used. This would not be expected to have much affect at the relatively low MCP gain of 500 Volts. The SIRT algorithm described in Chapter 3 was used with 1000 iterations. Aside from the lack of calibration just described, it seems that the number of SIRT iterations may

have been too great, as this may be one of the causes of overestimation of the reconstructed $\Delta\mu_a$. One hint of this is that the FWHM appears to be somewhat smaller than would be expected, given that the system should have a finite point spread function (SPF), which would be expected to smear the reconstructed volume out over a greater number of voxels than the actual width of 1.5 cm (the voxel resolution was 0.5 cm). It is a well-known property of the SIRT algorithm that excessive iterations can lead to underestimated volumes and overestimated $\Delta\mu_a$'s. Even so, it is expected that the product of the volume and $\Delta\mu_a$ would be conserved, and that too seems to be overestimated by about a factor of 2. The failure to correct for saturation could have led to poor fits to the forward model and errors in estimation of the optical properties. Now that the groundwork has been laid to understanding methods for optimizing image performance, future investigators will have the tools they need to produce images of greater contrast and accuracy. One last point on this image data, suggested by David Boas, is that one could use a constraint on negative values of $\Delta\mu_a$ to prevent the negative ringing observed in the reconstruction of Figure 7.11. In this case, it is known a priori that there should be no negative values. All this being said, it is impressive that a preliminary image with such good contrast and localization of the actual heterogeneity was achieved. From these results, we can be very optimistic that the Time-Domain Breast Imaging System will prove to be an excellent clinical tool in the near future.

7.4 Tissue Boundary Localization

Localization of tissue boundaries is an important consideration for a clinical breast measurement. In most cases, not all of the source fiber positions would be covered with

tissue. Some source fibers would have an unobstructed view of some number of detector fibers. This would result in signal on the photocathode well above the saturation limit. This would cause short-term degradation in the photocathode response and could even result in permanent damage to the intensifier. Thus, a safe mechanism must be identified for localizing tissue boundaries. The best way to accomplish this is to provide a means for all source fibers to be filled uniformly at a very low average power level. The

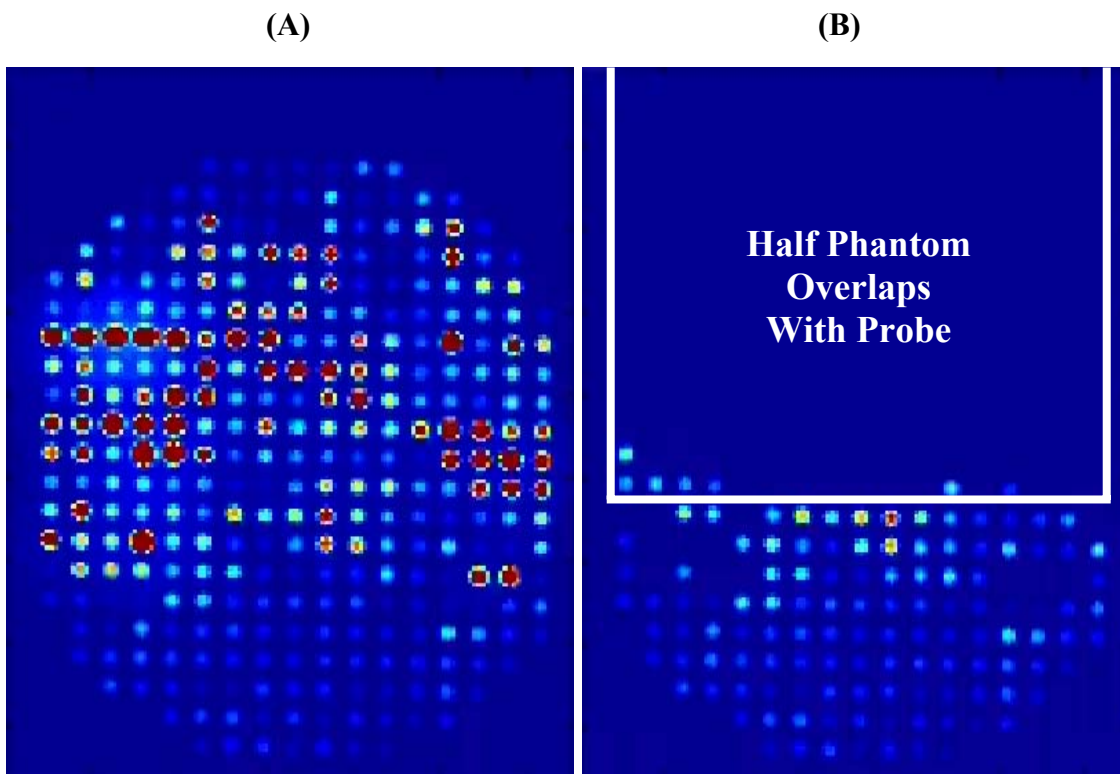


Figure 7.12 Raw data image files (A) and (B) represent a means of identifying tissue boundaries prior to scanning a breast. Image (A) represents the small signal that passes directly from source to detector fibers in the probe with no tissue or tissue phantom present. The source signal was effectively attenuated by several orders of magnitude by positioning the galvanometer mirrors such that the laser spot imaged between fibers, resulting in a small scattered signal filling all fibers within the Multiplexer. If tissue or a tissue phantom is present, as in image (B) for which the phantom filled the upper half of the probe's field of view, the outline of the tissue is readily discerned. Those source fibers outside the tissue boundary would not be scanned so as to prevent saturation of the camera's photocathode. A preferred implementation of this technique would be to fill the aperture stop of the Multiplexer with radiation from a beamsplitter-coupled LED. This would result in uniformly filling all fibers.

aperture stop of the telecentric scan lens system lies between the galvanometer mirrors. This aperture stop must be uniformly filled in angle space over the numerical aperture of the lens system, to insure uniform filling of all source fibers. This could be done by positioning a beamsplitter proximally to the galvanometer mirrors to inject signal from an LED illuminator. The LED system would be designed to re-image its aperture stop onto the aperture stop of the telecentric lens. The drive electronics for the LED source could be controlled by the computer to allow the current to be set to the appropriate level for a given gain and integration time. The LED-coupling beamsplitter could be comprised of a thin uncoated window, as very little LED illumination would be required.

Figure 7.12 represents an example of the tissue boundary localization approach. An LED illuminator has not yet been implemented, so the simple approach of using scattered laser power was demonstrated. The image of the laser was positioned by the galvanometer mirrors to fall midway between two of the source fibers within the Multiplexer's quick-connect fiber source array. Images (A) and (B) of Figure 7.12 are raw image files and have not been position mapped. The reasonably good 1-to-1 mapping between the probe and the camera fiber array gives a good idea of the concept. In image (A), there is no phantom present, a signal can be seen in all detector fiber positions. In image (B), the source-detector positions blocked by the tissue phantom are absent. This method of positioning the laser image between source fibers clearly does not result in good signal uniformity over all the fibers, but it does illustrate the point.

A tissue-localizing algorithm must also consider the possibility of saturation problems for source positions near the tissue boundary. Light would be expected to leak out of the side of the breast or breast phantom for source positions within, but near the

boundary. The LED source would not have enough power to assess this. The laser source could be used for this purpose. The source positions nearest the boundary could be scanned by the laser. It would not be wise, however, to use the full laser power. Three basic approaches could be used to reduce the laser power for this purpose. The polarizer could be motorized and rotated to change laser power. As mentioned in Chapter 6, it may be desirable to have the polarizer motorized for the additional purpose of compensating for the fall-off in camera response with increasing wavelength. Another method to modify the laser output power, which reaches the camera, would be to introduce a set of neutral density filters into the camera lens path. These filters could also be used to assess intensifier linearity. Another approach to modifying the laser output would be to scan the image of the laser over a given source fiber. This was shown to work with excellent repeatability, as described in the SNR measurement analysis of Chapter 6. The advantage of the last approach is that it is by far the fastest and would have the least impact on the total measurement time.

Much of the concern over saturation outside the tissue boundaries could be relieved by the use of the bladder described in Section 7.2. The exact tissue boundaries would not be required in that case, with respect to saturation concerns. There would still be some advantage to knowing the approximate boundary limits, however, as it would be desirable to minimize the number of required sources as to increase the number of delays, in order to optimize the SNR for a given patient.

7.5 Trade-Off Analysis for Number and Positions of Delays and Sources

Given that the total time available for acquiring data in a clinical setting is limited, the total number of measurements would be limited. The trade-off between the number of source positions and the number of delays must be understood to optimize the contrast to noise ratio (CNR) of the reconstructed optical properties. The number of detectors for our system was 313 and there was no penalty for using all of them for any given measurement, as they were all collected in parallel. The sources and delays were scanned serially and therefore limit the total measurement time. The relative positions of the source fibers and delays would also affect the CNR. Figure 7.13 shows a typical TPSF for which 12 different gate positions are indicated. Given a limited number of gate delays, the CNR for scattering would be optimized for delay positions located on the rising edge and near the peak, as the time to peak intensity of the TPSF is affected most

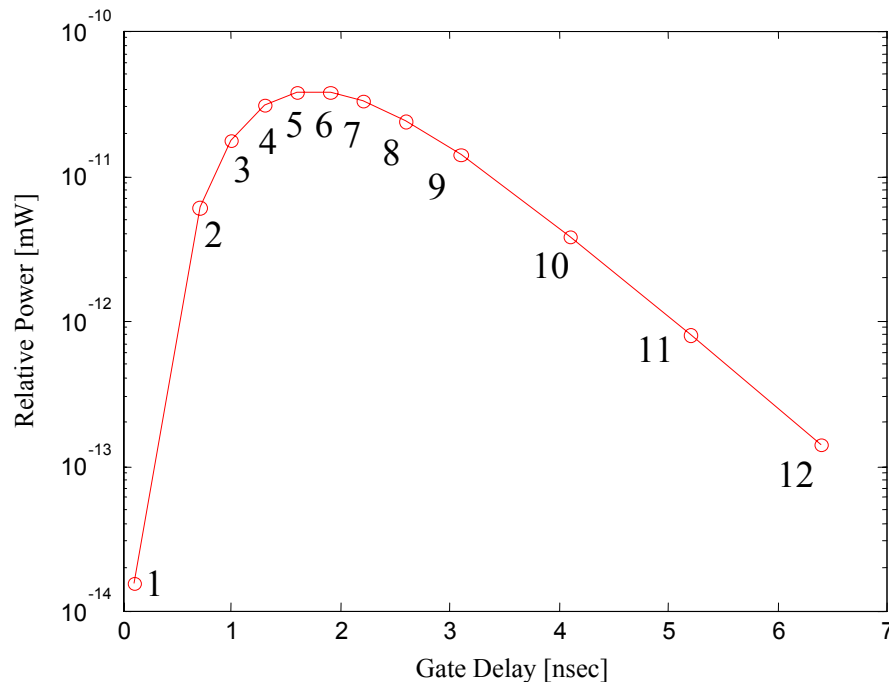


Figure 7.13 An example of a typical TPSF for a slab geometry is shown. The numbers indicate delay positions. For a subset of only 5 to 6 total delays, scattering contrast-to-noise ratio (CNR) would be optimized for a set biased toward early times and absorption CNR would be biased toward later delays.

by the scattering coefficient. Likewise, the CNR for the absorption would be optimized for delays near and after the peak intensity. Thus, one must decide prior to the measurement, whether there is greater interest in imaging scatter or absorption and bias the positions of the delays accordingly.

Simulations studies should be carried out to determine the optimal trade-off between the number of sources and the number of delays. The simulations should hold the product of the number of delays and the number of source positions constant. The magnitude of the noise would be expected to affect the outcome and should therefore be consistent with the system performance, which is near photon limited. For a given number of sources and delays, there would be an optimal configuration of relative positions for the sources and delays. Care must be exercised, however, in setting up the simulations. For example, if one were to place a heterogeneity midway between the source and detector plates, it may be concluded that the CNR is optimized for more delays and fewer source positions, as the large number of detector positions would result in reasonably good spatial resolution, even for a relatively small number of sources. If, however, the heterogeneity were biased in the direction of the sources, the result may be skewed in favor of more sources and fewer delays. The position of the heterogeneity could also be sensitive to its position in the plane parallel to the source and detector array (x - y plane). For any fixed number of delays, the positions of the delays with respect to the features of the TPSF would also affect the outcome and must be optimized. Thus, the trade-off analysis for sources and delays is a complex multi-dimensional simulation. Intuitively, it seems likely that something on the order of 6 delays would be sufficient to allow a good fit with 2 delays on the rising edge, 1 delay either side of the peak and 2 delays on the tail. The

number of source positions would affect the limiting resolution of the system. Perhaps the most significant limitation of DOT is its relatively poor spatial resolution, so understanding the limitations of spatial resolution as a function of the number of source positions is important. It would be useful to understand at what source density would no further improvement in spatial resolution result from increasing the number of sources.

7.6 Maximum Permissible Exposure (MPE) Limits

The American National Standards Institute's ANSI standard Z136.1-1993 was used to calculate the maximum permissible skin exposure (MPE) for the Time-Domain Breast Imaging System. In general, the document is convoluted and poorly written, so it is often difficult to interpret MPE for specific instrument configurations. Specific sections are referenced here for traceability. Fortunately, there are a number of examples in the appendix of the ANSI standard that are reasonably helpful.

A footnote for Section 8.2.2.2 of the ANSI standard states that for repetition rates that are so high that multiple pulses occur in a time frame less than 18 msec (0.4 to 1.05 μm), pulse energies delivered within those time frames are summed directly such that one can assume the energy acts as if it were delivered as a single pulse. Section 8.4 of the ANSI standard states that MPE values for skin exposure to laser beam are given in Table 7 and that those levels are worst-case conditions based on the best available information. Section 8.4.1 of the ANSI standard states that for repetitively pulsed lasers the MPE's for skin exposure are applied as follows: Exposure of the skin shall not exceed the MPE based upon a single-pulse exposure, and the average irradiance of the pulse train shall not exceed the MPE applicable for the total pulse train T . For the purposes of the MPE

calculations, Table 8 of the ANSI standard defines the limiting aperture for the case of skin in the spectral region between 0.400 and 1.400 μm to be 3.5 mm in diameter. The limiting aperture reflects the fact that due to the relatively low absorption of tissue and its scattering properties, there is little difference in the effect of a 100 micron source fiber incident on the tissue in comparison with a 3.5 mm fiber, both with the same total energy dose. That is to say that after a number of scattering events their respective effects are nearly indistinguishable. This is a point of confusion in the interpretation of the ANSI standards that causes some investigators to use lower than necessary average powers.

Example 3 of Appendix B3.1.2 is referred to, which describes the MPE for repetitively pulsed laser with a very high pulse repetition frequency (PRF). According to the example, since the PRF is above 55 kHz, the average radiant exposure limitation applies. According to Table 7 of the ANSI standard, the MPE in $[\text{J}/\text{cm}^2]$ for exposure durations t between 10^{-7} and 10 seconds is given by

$$MPE[\text{J} / \text{cm}^2] = 1.1C_A t^{1/4} \quad (7.1)$$

where C_A is the wavelength correction factor from Table 6 of the ANSI standard given by

$$C_A = 10^{2(\lambda - 0.700)} \quad (7.2)$$

where λ is the laser wavelength expressed in microns. Use of these equations results in the plot of Figure 7.14 (A), below. The MPE is shown for integration times, or equivalently pulse durations, of 25, 50, 100 msec and 1 and 10 sec. At 10 sec, the MPE is equal to the steady state, or CW limit, which is shown in Figure 2 for reference. For the phantom characterization, a maximum of about 40 mW was used at the tissue surface, thus according to Figure 7.14 (A), the average power for a 50 msec pulse duration could be increased by more than an order of magnitude for a single integration time exposure.




Figure 7.14 The maximum permissible exposure (MPE) is shown plot (A) as a function of wavelength for a single integration time ranging between 25 msec and 10 sec. The 10 sec values for plot (A) are equivalent to the steady state CW MPE. The MPE is in terms of the average power entering the tissue. Plot (B) explores the added complexity of the Breast Imaging System by assuming a worst case minimum number of source positions of 40. If it is assumed that the limiting rate is due to the integration time, we can scan a maximum of 10 delays (revisits to the same tissue area) within a 10 second period for a 25 msec integration time, thereby decreasing the limits in plot (A) by a factor of 10. The same restrictions applied to 50 and 100 msec, give a maximum number of revisits in a 10 second period of 5 and 2.5 times, respectively. Plot (B) indicates these more conservative MPE limits. For the 1 and 10 second integration times, the more restrictive limits in plot (A) apply, as the number of revisits would only be 0.25 and 0.025, respectively. Note that these are worst case assumptions in plot (B), as in general the time between integration times will be limited by more than the integration time itself.

It must be consider that the Time-Domain Optical Breast Imaging System's measurement scheme is somewhat more complex than addressed in the ANSI standard. The measurement as a whole must be considered. It is comprised of a number of sets of pulse trains. Consider the case of a small breast for which only 40 of the 150 available source fiber positions are used. If it is assumed that the limiting acquisition rate is due to the integration time in the worst case, a maximum of 10 delays could be scanned (revisits to the same tissue area) within a 10 second period for a 25 msec integration time, thereby decreasing the limits in plot (A) by a factor of 10. The same restrictions applied to 50 and 100 msec, give a maximum number of revisits in a 10 second period of 5 and 2.5 times, respectively. Plot (B) of Figure 7.14 indicates these more conservative MPE limits. For

the 1 and 10 second integration times, the more restrictive limits in plot (A) of Figure 7.14 apply, as the number of revisits would only be 0.25 and 0.025, respectively. Note that these are worst case assumptions in plot (B), as in general the time between integration times would be limited by more than the integration time. This more restrictive analysis limits us to about a factor of 5 times higher average laser power at 25 msec, than was used in the SNR analysis presented in Chapter 6. The real limit then, to the maximum average laser power for this breast imaging system, may be due to the saturation effects described in Chapters 5 and 6, and not the ANSI MPE limits.




Figure 7.15 The maximum permissible exposure (MPE) for skin is shown plotted against wavelength for the case of CW laser radiation, defined as exposure times in excess of 10 seconds.

7.7 Optimizing the Number and Order of Measurements for a Clinical Time-Domain Breast Imaging System

For a clinical Time-Domain Breast Imaging System integrated with X-ray mammography, the total number of measurements would be limited to the number that could be made in a period on the order of 120 seconds. This number is limited by patient tolerance, as the patient must remain in compression following the optical image for the co-registered X-ray image. The time to switch from optical to X-ray is also critical, and is

included here in the time limit for optical imaging. The quick release design features of the optical probe allow the source and detector plates to be removed within the order of 15 seconds.

The time required to change delay was on the order of 500 msec for the current version of the system. The combined switch and settle time for the Multiplexer's galvanometers to move the laser focus from one fiber to another in a raster scan was on the order of 300 μ sec. The time to switch wavelengths from 805 nm to 830 nm or from 830 nm to 760 nm, was on the order of 15 seconds. For three-wavelength spectroscopy, this would require two wavelength changeovers. Consider the case of taking measurements at 3 different wavelengths, 6 delays, and 120 source positions. The total number of source positions would be dependent on the boundaries of the breast tissue and would therefore vary from patient to patient. One hundred twenty source positions is a reasonable upper limit that is likely to be presented by any given patient.

Time Element	Number	Switch / Dwell Time [sec]	Time Multiple	Total Time [sec]
Tissue Boundary Scan	1	5	1	5
Background Image	30	0.025	30	0.75
Determine Delay Positions	1	5	1	5
Source Dwell Time	120	0.025	120 x 6 x 3	54
Source Switching		0.0003	120 x 6 x 3	0.648
Time Gate Delays	6	0.5	5	2.5
Wavelength Change	3	15	2	30
Remove Optical Probe	1	15	1	15
Total				112.90

Table 7.2 The time elements are shown for the case of a Clinical Breast Imaging measurement, indicating that the total time required to scan a large breast with 3 wavelengths, 6 delays, and 120 source positions is on the order of 120 seconds for an integration time of 25 msec.

The analysis of Table 7.2 shows that it would be feasible to acquire data from 120 source positions, 3 wavelengths, 6 delays, 30 background images, and a tissue boundary scan all within the order of 120 seconds. The CCD camera used in the ICCD did not have a TTL synchronization pulse output, so the time between fiber positions was not optimized. The actual time to scan the 2160 measurements was found to be 210 seconds. Future CCD cameras would not have this limitation. Even with the present camera, however, over 60 source fiber positions could be used. This system represents a significant improvement in measurement time relative to systems described in the literature.

7.8 Summary

Recommendations were presented with regard to improving signal processing toward the goal of optimizing clinical performance. Phantom data was discussed that revealed additional means for improved performance. Preliminary imaging results were evaluated and suggestions made for using the performance characteristics of the system to improve image contrast and noise. The limitations posed by ANSI skin MPE were found not to be overly restrictive or limiting for this system architecture. Issues related to tissue localization were explored and a bladder device was described to minimize requirements for tissue boundary localization and to minimize boundary effects on the forward model. The trade offs between the number and positions of the sources and gate delays was motivated and discussed. Finally, an optimal measurement scheme and measurement times for an optimized clinical measurement was described.