SAND2013-10110C

Bubble Masks for Time-Encoded Imaging of Fast Neutrons

James Brennan, Erik Brubaker, Aaron Nowack, John Steele, Melinda Sweany, and Daniel Throckmorton

***Abstract*—Time-encoded imaging is an approach to directional radiation detection that is being developed at SNL with a focus on fast neutron directional detection. In this technique, a time modulation of a detected neutron signal is induced—typically, a moving mask that attenuates neutrons with a time structure that depends on the source position. An important challenge in time-encoded imaging is to develop high-resolution two-dimensional imaging capabilities; building a mechanically moving high-resolution mask presents challenges both theoretical and technical. We have investigated an alternative to mechanical masks that replaces the solid mask with a liquid such as mineral oil. Instead of fixed blocks of solid material that move in pre-defined patterns, the oil is contained in tubing structures, and carefully introduced air gaps—bubbles—propagate through the tubing, generating moving patterns of oil mask elements and air apertures. Compared to current moving-mask techniques, the bubble mask is simple, since mechanical motion is replaced by gravity-driven bubble propagation; it is flexible, since arbitrary bubble patterns can be generated by a software-controlled valve actuator; and it is potentially high performance, since the tubing and bubble size can be tuned for high-resolution imaging requirements. We have built and tested various single-tube mask elements, and will present results on bubble introduction and propagation for different tube sizes and cross-sectional shapes; real-time bubble position tracking; neutron source imaging tests; and reconstruction techniques demonstrated on simple test data as well as a simulated full detector system.**

I. Introduction

We are investigating the possibility of extending the technique of fast neutron time-encoded imaging [1] from the one-dimensional case of large standoff detection of special nuclear material (SNM) to the more demanding application of high-resolution two-dimensional imaging, such as for arms control treaty verification. Briefly, in time-encoded imaging, the detected neutron signal is modulated in time in a defined fashion, typically using a moving mask that attenuates neutrons with a time structure that depends on the source position. The technique is characterized by simple and robust detector elements, low channel counts, and low cost, in contrast with other imaging approaches. However, time-encoded imaging of fast neutrons has never been demonstrated in a high-resolution mode neede for these applications.

For high-resolution imaging in two dimensions, we need a time-dependent mask that produces unique, and preferably orthogonal, attenuation patterns on a neutron-sensitive detector; the patterns produced must vary over relevant distance scales in the source space; and the system must be physically buildable and operable. In order to sufficiently attenuate fast neutrons, mask elements are typically made of ~ 4 in thick HDPE or other hydrogenous material.

We explored a potential technique that replaces the solid mask with a liquid such as mineral oil. Instead of fixed blocks of solid material that move in pre-defined patterns, the oil is contained in tubing structures, and carefully introduced air gaps—bubbles—propagate through the tubing, generating moving patterns of oil mask elements and air apertures. A conceptual model of such a system is shown in Fig. 1. Compared to current moving-mask techniques, the bubble mask is simple, since mechanical motion is replaced by gravity-driven bubble propagation; it is flexible, since arbitrary bubble patterns can be generated by a software-controlled valve actuator; and it is potentially high performance, since the tubing and bubble size can be tuned for high-resolution imaging requirements. Bubble masks could be a key enabling technology for high-resolution time-encoded imaging, with applications in arms control treaty verification and radiological emergency response. Bubbles can be introduced in such a way that a URA pattern moves through the tube(s) cyclically, satisfying the requirement for artifact-free image reconstruction in the ideal limit. In reality, for fast neutron imaging we cannot achieve the ideal conditions of full mask opacity, negligible mask thickness, far-field sources, and so on. With further study and simulations, the optimal mask pattern may turn out to be something other than a cyclical URA; the greater flexibility of a bubble mask relative to a mechanical mask allows the exploration of a larger space of time-encoded masks. In fact, with a bubble mask, data can be acquired with different mask patterns in the same measurement, possibly reducing systematic reconstruction effects relative to the use of a single fixed mask pattern.

II. Bubble Characteristics

To test this concept, we decided a simple experimental path would yield the best results. Our initial design was inspired by an art exhibit by artist Bruce Shapiro. Mr. Shapiro designed multiple bubble displays that use small tubes to align bubbles in patterns, sometimes depicting faces of historical figures [2]. Our requirements push us in slightly different directions, but the basic concepts remain the same as the art exhibit: select a tube material and diameter, fill the tube with mineral oil or another viscous liquid, and inject compressed air into the tubes.

An illustration of the test system is depicted in Fig. 2. Our initial design consisted of a 1.5 in ID acrylic tube that is 2 ft long. We used a simple 24 V power supply to deliver DC

Manuscript received November 22, 2013.
All authors are with Sandia National Laboratories, Livermore, CA.
e-mail: ebrubak@sandia.gov

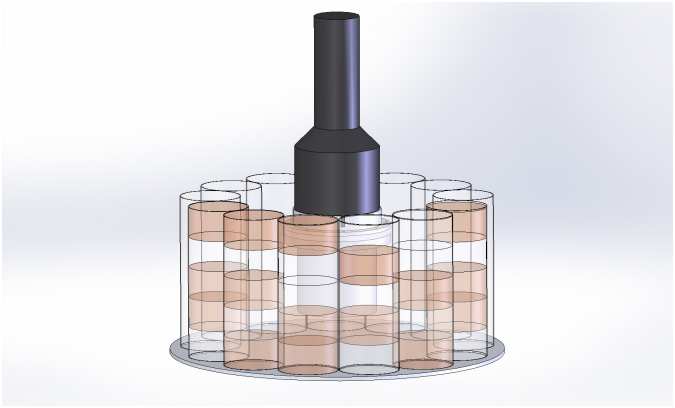

Fig. 1. A conceptual model of a time-encoded imager based on a bubble mask. Bubbles are introduced into an array of individually controlled vertical tubes of mineral oil or similar material.

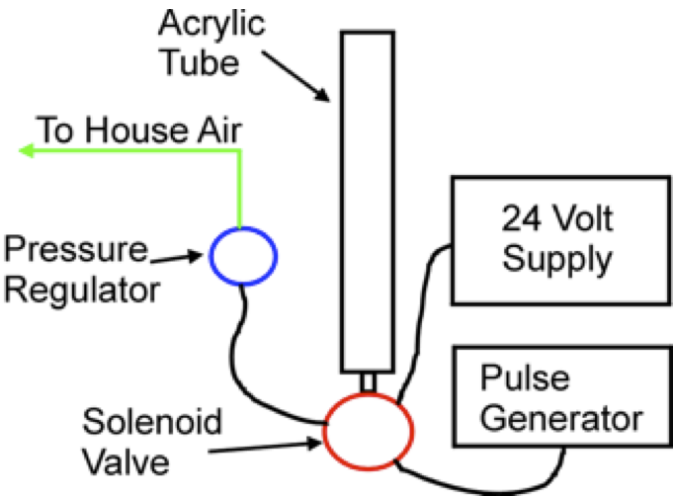

Fig. 2. Line drawing of initial test configuration

power to the Asco solenoid valve mounted to a check valve on the bottom of the tube. A DG535 delivered the pulse signal to the valve, allowing us to control repetition rate and pulse duration. House compressed air was regulated through a Parker air pressure regulator to allow fine control of the air delivered to the valve. Mineral oil was poured in to the open top of the acrylic tube to a level of about 18 in.

As seen in Fig. 3, bubbles emerge from the bottom of the tube. As the solenoid valve closes, small bubbles that remain in the valve to tube connection slowly exit and trail the main bubble up the tube. These bubbles contribute additional small bubbles that remain in the mineral oil. Additionally, as the bubble size is increases the main bubble separates from the injection point as the buoyant force increases. This leads to multiple large bubbles being introduced into the tube instead of the desired single bubble. This problem keeps us from making larger bubbles and reduces the possible bubble configurations. While large bubbles are difficult to produce, smaller bubble sizes are very reproducible and transit up the tube at a very consistent velocity. If bubbles are introduced into the tube where the previous bubble is at minimum greater than one bubble length away from the subsequent bubble, the bubbles

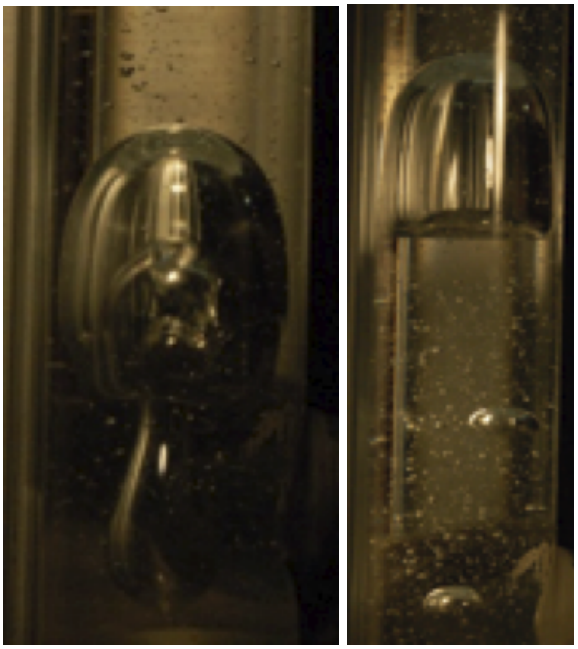

Fig. 3. First image shows bubble as it emerges from the solenoid valve. Second image shows small after bubbles following the main bubble up the tube.

remain well separated. As an example, if the bubble introduced is 2 in long, as long as the next bubble isn't closer than 2 in from the previous they will remain separate. As the gap between the bubbles is decreased, there is a probability that the lower bubble will catch and combine with the previous bubble. This doesn't happen every time, but is undesirable because of the unpredictable nature of propagation.

After various modifications to address the problems encountered in initial operations, consistent and reproducible bubble production and propagation was achieved. Some of the changes included injecting bubbles from the top of a side inlet tube rather than from the bottom of the main tube, using a hose pinch solenoid valve to reduce abruptness in the opening and closing action, and replacing mineral oil with glycerine. An example of a bubble propagating through glycerine is shown in Fig. 4. In this case, the bubble bottom is convex. Additionally, the bubble is less prone to disturbance as the next bubble is introduced which means the bubble shape remains consistent as it travels up the tube. Finally, to add contrast for the video bubble tracking described below, blue dye was mixed into the glycerine solution.

In models of a full detector system such as that shown in Fig. 1, cylindrical mask tubes are not ideal due to the reduced contrast at the edges of the tube where the particle path length through the fluid is small or zero. We therefore investigated tube cross-sections that would provide more consistent contrast. We started with a simple 1 in x 1 in square tube with an internal wall dimension of 0.83 in per side. Initial tests were performed using mineral oil. Initial results showed similar performance of both bubble injection and propagation. Bubble shapes were very consistent from bubble to bubble. However, it was noted that the bubble had the same shape as those in the cylindrical tube, while the hope was for bubbles that

Fig. 4. A bubble transiting up the tube filled with glycerine.

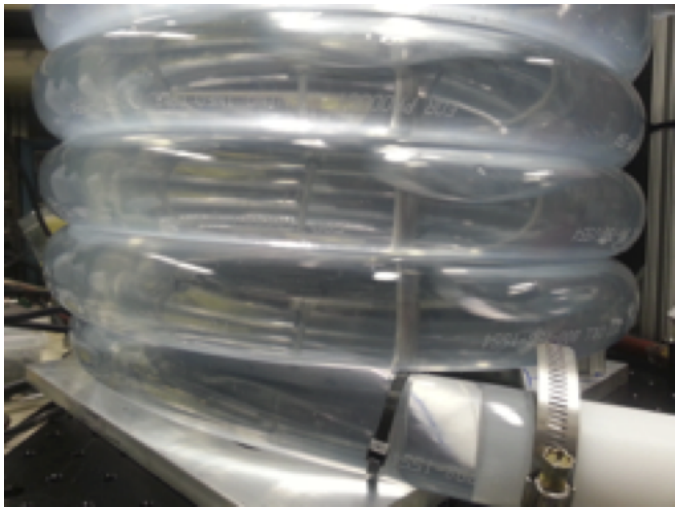

Fig. 5. Helical mask with mineral oil and air injected.

filled more of the square tube cross-section. With larger bubble sizes, an interesting effect was observed. As the bubble size increases, the bubble expands towards the wall of the tube until it reaches about 75% of the tube cross-section, after which it expands only along the length of the tube. Additionally, as the bubble became very large (several cross-section lengths long), the fluid flow in the corners of the tube would begin to shear the sides of the bubble, resulting in inconsistent bubble shapes and a large number of very small stationary bubbles produced in the tube. Further testing using 1 in x 4 in and 2 in x 4 in interior cross-section tubes yielded similar observations.

Another possible mask design could be a single tube of flexible material wrapped in a helix. This design could have advantages as there would only need to be one solenoid valve and the associated electronics would be simpler. To test this design we coiled a length of flexible tubing around a structural frame. As hypothesized, however, the injected bubbles did not fill the now almost horizontal tube cross-section, but rather extends along the longitudinal direction in a small region near the top of the tube, occupying only about 5% of the cross-sectional area. (Fig. 5). Any helical design would need to use a fluid with significantly different viscosity and cohesion/adhesion properties, or a substructure within the mask tube to permit the desired propagation.

III. Bubble position tracking

In order to reconstruct radiation sources the shape of the mask must be known at each time. This is done by tracking the position and shape of each bubble over a data run. While we have the information of when a bubble was injected into the tube and with how much air, it is not straightforward to extend that information to bubble size and position for a variety of tube shapes, liquids, and other conditions. Instead we acquire a video of the run and use techniques from the field of computer vision for measuring bubble position and size. Using the freely available OpenCV library [3], we developed a generic position tracking algorithm and a more specific and simple method. For this method, the OpenCV Python bindings were used to access its functionality.

The generic algorithm is based on measuring significant changes in time on a per-pixel basis. A running average of the past several frames is calculated and subtracted from the current frame. The resulting image only displays areas of motion within a frame, displayed as a green overlay in Fig. 6, left. These areas of motion are then surrounded by a bounding box. Any overlapping bounding boxes are expanded and then combined to group neighboring regions for a bubble. This results in the real-time object position tracking in Fig. 6, right. While this method works in a variety of cases and is able to track most bubbles it cannot tightly constrain the bubble due to the churning motion of the liquid that trails it. Also, the bounding boxes constantly adjust their size making it difficult to measure the center position of a bubble.

The next approach is simpler and is based on comparing the values of pixels within a region before bubbles are injected into the tube and afterwards. A box indicating the region of interest is made and calculations are restricted to this region. The pixels are averaged horizontally, resulting in an array of pixel values for each vertical position. This is done for the red, green, and blue (RGB) channels separately. A baseline array for each RGB channel is made by time-averaging this array over the period before bubbles are injected into the tube. Once bubbles are injected into the tube the baseline array is subtracted from the calculated array in the current frame. The bubbles are clearly visible as deviations above this value and a simple threshold on this difference tightly tracks each bubble as seen in Fig. 7, left. While this method is simpler it turned out to be more useful as the edges of the bubbles are clearly defined. This allows us to input an accurate mask into the source reconstruction algorithm.

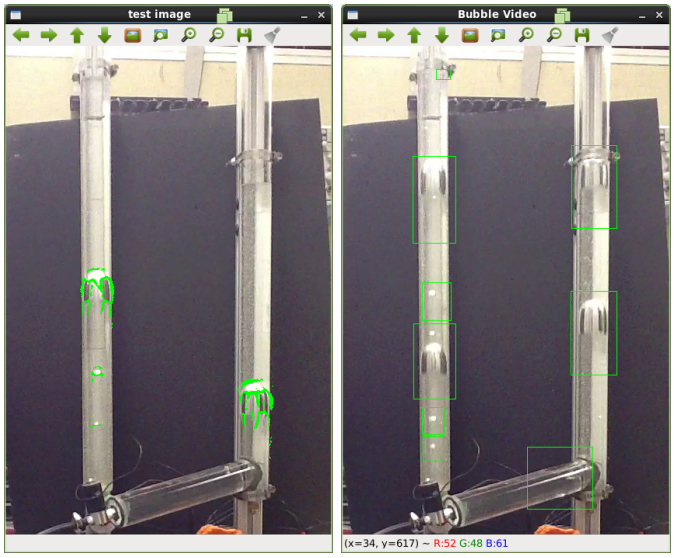

Fig. 6. Left: the differences of the current frame compared to previous frames is overlaid in green. Right: the resulting bounding boxes for areas of motion.

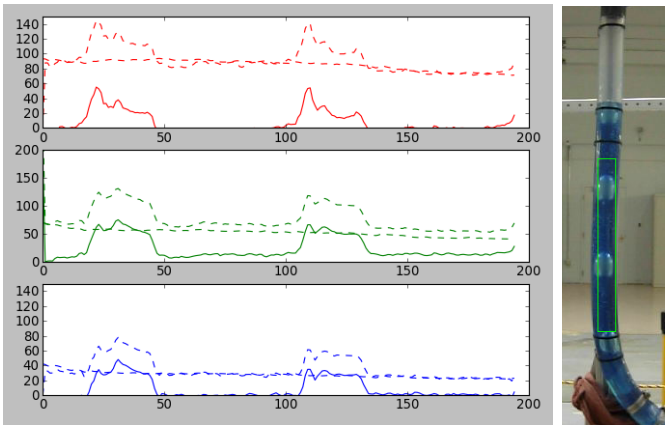

Fig. 7. On the left the baseline and current vertical pixel values are shown (dashed) and the deviation from the baseline is show (solid). Right: The region of interest is overlaid on the image

IV. Image Reconstruction

An important aspect of using the bubble mask approach to time-encoded imaging is that the mask pattern is not fully known in advance of the measurement, but is detected in real time as described in Sec. III. Any image reconstruction technique relies on a firm grasp of the system response function. The challenge then is not only to perform image reconstruction using e.g. maximum likelihood expectation maximization (MLEM), but in fact to build the system response function itself in real time. Previous techniques to map out the system response due to neutron attenuation in the mask used hours to days of CPU time, and were therefore not feasible for mask patterns that change in every measurement.

To resolve this difficulty, we separate the system response calculation into two parts. Additionally, we consider the mask as being composed of many thin "slices" of liquid, each of which may be present or not present at any given moment in time. A bubble, then, is simply the absence of a contiguous set of liquid slices in the tube.

The first part of the response calculation is the most computationally intensive, and may take many hours to complete, but does not need to be repeated for every mask pattern. In this first stage, we calculate the attenuation for each combination of source pixel, detector, and liquid slice based on the length of the particle's path through that slice. Since the physical extent of those three objects is of the same order, the attenuation must in fact be an average over all possible trajectories connecting the source pixel and detector. We use a Monte Carlo averaging technique by selecting a series of random locations in the source pixel and detector (respecting the expected distribution of interaction depths) for the endpoints of the trajectory. The average number of interaction lengths is stored for each such combination.

The second part of the response calculation requires as an input the specific mask pattern as a function of time. For each snapshot in time, the overall detector response to a given source pixel is calculated from the attenuation length summed over all liquid slices present in the tube at that time. This part of the calculation requires simple multiplication and exponentiation operations, and can be performed as the data and the bubble pattern are acquired, allowing image reconstruction to be performed in real time.

For these studies, a single interaction length of 5 cm is used as a typical value for fission-energy neutrons. Note that we assume any interaction results in non-detection of the neutron, and scattering *into* the detector is not included in the final system response function.

Finally, given the system response function, the images shown in this report are reconstructed using a standard MLEM iterative algorithm [4].

V. One-Dimensional Imaging Experiment

The experimental setup, shown in Fig. 8, consists of a tube of colored glycerol with a source stand and a neutron detector on opposite ends. The bubbles are released through an air-pressure solenoid valve which is controlled through an Arduino UNO over USB to the data acquisition computer. The source stand can be moved up and down (defined as the z axis), always with the mask tube between the source and the detector. The detector comprises a PMT coupled to a 1 in dia. x 1 in cylinder of liquid scintillator; full pulse waveforms are read out and saved for offline analysis. The neutron detector was calibrated using standard procedures, both to determine the (electron-equivalent) energy scale and the pulse shape discrimination performance. The latter information was used to select for imaging only those events that were very likely to be neutrons based on the observed pulse shape.

An early test was made to measure the effects on the neutron rate caused by propagating bubbles. A constant bubble shape is injected into the tube at a constant rate of 1 Hz for one hour. By plotting the number of detected neutrons as a fraction of the bubble period, the average rate of neutrons for many bubbles can be calculated. The resulting plot in Fig. 9 shows that an air bubble allows at most 30% more neutrons to pass through to the detector for this tube and detector size.

In order to combine the video and detector data, a new acquisition scheme was devised. This system used a Logitech

Fig. 8. Experimental setup of the bubble acquisition system: a PMT is placed opposite of a source stand with the glycerol tube and bubbles placed in between

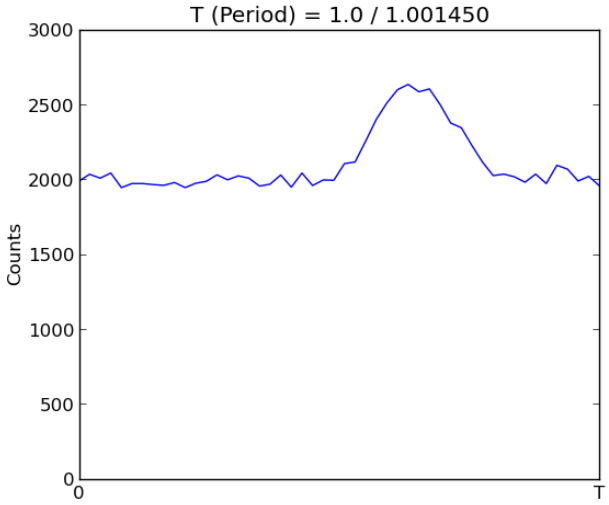

Fig. 9. The number of neutrons for an hour-long run as a function of the bubble period. A contrast of 30% is measured

C210 USB web cam to capture the bubble position images and a VME crate containing a Struck SIS33350 digitizer and an Iseg VHS3040n high voltage supply to capture the neutron detector response. The choice of these alternate data input sources allowed one program on one PC to synchronously receive data from each and store that data appropriately time-stamped in a single data file. The data acquisition software was written entirely in C++, using an existing framework developed in-house for other similar applications. The control and collection program code linked directly to and depended on OpenCV to manipulate the video frames. The frame rate achieved varied from program run to program run, typically just under the standard NTSC TV rate of 30 frames per second but occasionally dipping to as low as 10 frames per second.

The arduino program controlling the bubble introduction was modified to allow random (within defined limits) bubble size and spacing, and this mode was used for the test described here. An example of a reconstructed source distribution from

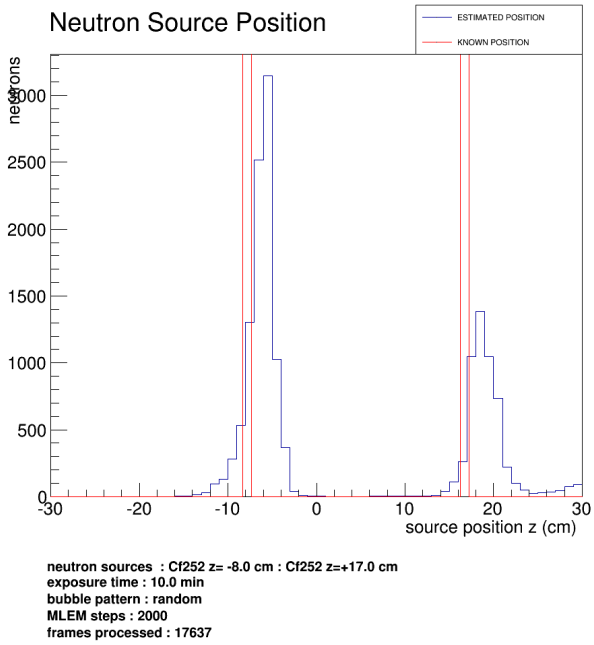

Fig. 10. Two ^{252}Cf sources, 10 min.

one such data run is shown in Fig. 10, which used two equal-strength ^{252}Cf sources at different z positions: $z = -8$ cm and $z = +17$ cm. The one-dimensional image correctly identifies the two sources, although there is an offset in the reconstructed z positions. This may be due to a delay between the camera frame acquisition and readout, resulting in a systematic offset between the true bubble pattern over time and that assumed in the reconstruction.

VI. Modeling

We modeled both a simple imaging configuration, similar to the one used in Sec. V, and a complete two-dimensional imager with many mask tubes. The detector geometry is described using the ROOT data analysis package, and a source space in (ϕ, z) is defined at a given radius away from the tube.

Figure 11 shows the reconstructed source position of a 10 min simulated data run, for two point sources as in the experimental case above, albeit in different locations. The red line indicates the true source positions. Each source has a detected neutron rate of 5 Hz. There is no background field present. Although the configuration is not exactly that of the experimental data, the general features of the result are similar (note the log y scale). This gives us confidence in the simulated results.

Figure 12 is a rendering of a full detector geometry, in which the mask is a series of tubes arranged in a circle around one detector: the tubes are 4.48 cm in diameter, chosen to be the neutron attenuation length through water.

Although the tube geometry is the most likely two-dimensional extension of our current experimental setup, it is not ideal because the area between tubes will have poor masking of neutron sources. The reconstruction of a ring source (20 min, 100 Hz detected neutron rate, no background) is shown in Fig. 13 for an alternative annulus geometry, in

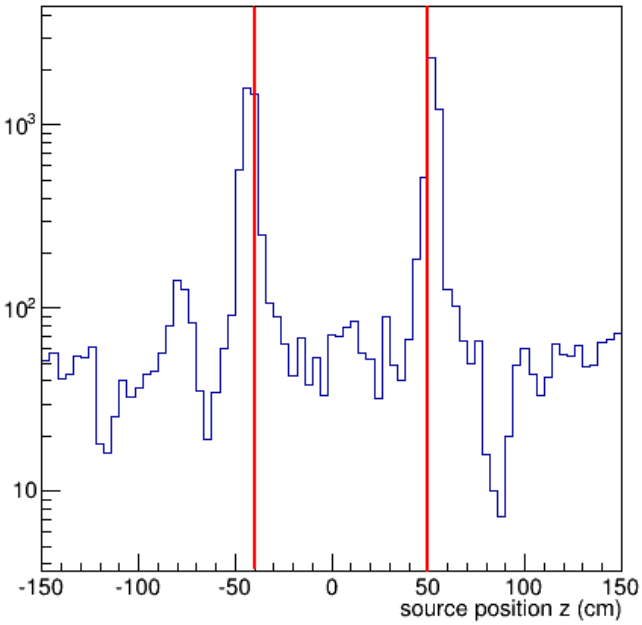

Fig. 11. The reconstructed source positions for a one-dimensional bubble mask, where a 10 min run is simulated for two point sources at -40 cm and 50 cm. The source strength is such that the rate of detected neutrons is 10 Hz, and there is no background present.

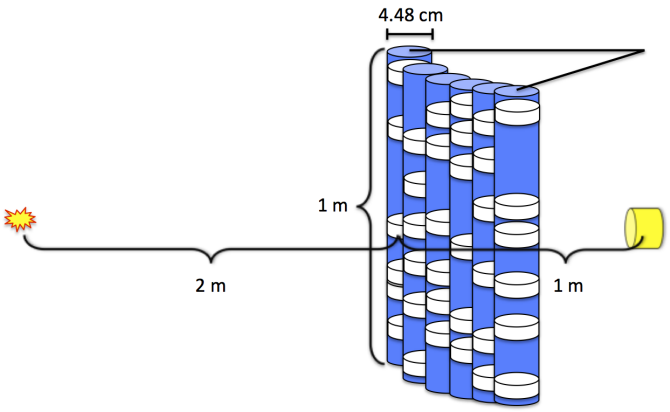

Fig. 12. The simulated two dimensional tube geometry: the detector is at the center, the tube array is at a radius of one meter, and the source space is at a radius of 3 meters in reference to the detector. Each tube is 4.48 cm and one meter tall.

which the bubbles are assumed to fill each wedge-shaped azimuthal segment. This eliminates the effects of reduced attenuation and contrast on the sides of the tubes, and significantly improves the imaging performance; however, as seen in Sec. II, bubbles filling tubes with non-circular cross-sections are difficult to produce. This is therefore an optimistic result, but demonstrates the potential of time-encoded imaging based on bubble masks under practical conditions.

VII. Conclusion

The goal of this work was to determine the feasibility of building bubble masks for neutron time-encoded imaging, and identify any limitations on their design or performance. We have successfully demonstrated controlled introduction and

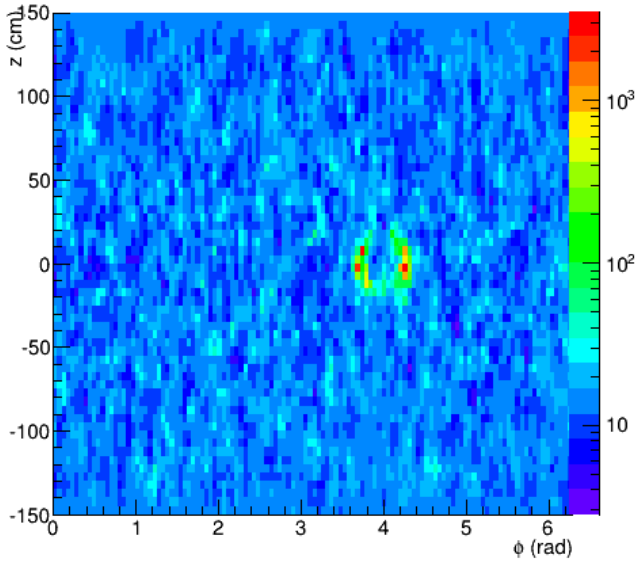

Fig. 13. Two-dimensional reconstruction of a circular source with the annulus geometry. The simulated data run was 20 min long and had no background.

propagation of bubbles through a tube of viscous hydrogenous liquid in a pattern relevant for time-encoded imaging of fast neutrons. We developed a technique for tracking bubble positions in real time and methods to reconstruct a source distribution from acquired neutron rate data. Results from a simple experimental setup provide a demonstration of the concept, and give us confidence in simulation results. The simulation was then used to extrapolate beyond the experimental results to explore the potential performance of a large-scale imaging system based on bubble masks.

Two major limitations of the bubble mask technique were identified. The first is the challenge of achieving a high mask contrast, which would require tubes (and therefore bubbles) with large cross-sectional area, and ideally a rectangular shape. The second is the inherent correlations among the response vectors of nearby source positions, due to the motion of the bubbles through the tubes. Low mask contrast and non-ideal orthogonality both contribute to slower resolving times, or the need for more statistics for a given detection threshold or image resolution.

Advantages of the bubble mask technique include the length of the encoding pattern, which is limited only by the length of the acquisition, in contrast to fixed moving masks, which must be cyclical with some finite length. In addition, the ability to respond to observations by changing the mask pattern (adaptive encoding) has potential to improve resolving times and open up a new class of algorithms for source detection and imaging.

Acknowledgment

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S.

Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND 2013-XXXX.

REFERENCES

[1] P. Marleau, J. Brennan, E. Brubaker, M. Gerling, A. Nowack, P. Schuster, and J. Steele, "Time encoded fast neutron/gamma imager for large standoff SNM detection," in *2011 IEEE Nuclear Science Symposium Conference Record*. IEEE, Oct. 2011, pp. 591–595. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6154118

[2] B. Shapiro, "The art of motion control: Pipedream series," http://taomc.com/art/permanent_installations/pipedream_series.html, accessed: 2013-11-21.

[3] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[4] L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography." *IEEE transactions on medical imaging*, vol. 1, no. 2, pp. 113–22, 1982. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/18238264