



X-ray Instrumentation Associates
8450 Central Ave
Newark, CA 94560
Phone: (510) 494-9020
FAX: (510) 494-9040

PHASE I FINAL REPORT

Ultra-Low Background Alpha Activity Counter

DE-FG02-04ER84120

W.K. Warburton, Ph.D.
Principal Investigator

X-ray Instrumentation Associates
8450 Central Avenue
Newark, CA 94560-3430

Executive Summary

In certain important physics experiments that search for rare-events, such as neutrino or double beta decay detections, it is critical to minimize the number of background events that arise from alpha particle emitted by the natural radioactivity in the materials used to construct the experiment. Similarly, the natural radioactivity in materials used to connect and package silicon microcircuits must also be minimized in order to eliminate "soft errors" caused by alpha particles depositing charges within the microcircuits and thereby changing their logic states. For these, and related reasons in the areas of environmental cleanup and nuclear materials tracking, there is a need that is important from commercial, scientific, and national security perspectives to develop an ultra-low background alpha counter that would be capable of measuring materials' alpha particle emissivity at rates well below $0.00001 \alpha/\text{cm}^2/\text{hour}$.

This rate, which corresponds to 24 alpha particles per square meter per day, is essentially impossible to achieve with existing commercial instruments because the natural radioactivity of the materials used to construct even the best of these counters produces background rates at the $0.005 \alpha/\text{cm}^2/\text{hr}$ level. Our company (XIA) had previously developed an instrument that uses electronic background suppression to operate at the $0.0005 \alpha/\text{cm}^2/\text{hr}$ level. This patented technology operates as follows. First, it sets up an electric field, in a large parallel plate geometry, between a large planar sample and a large planar anode, and fills the gap between them with pure Nitrogen. When an alpha particle enters the chamber it ionizes the Nitrogen, producing a "track" of electrons, which are attracted to the anode by the electric field. Because alpha particles emitted from the anode produce tracks close to the anode, it takes less than 10 microseconds (μs) to the electrons to be collected, producing a signal with a $10 \mu\text{s}$ risetime from the instrument's preamplifier. Tracks produced by alphas emitted from the sample, on the other hand, have to drift across the full gap between sample and anode, which takes about $35 \mu\text{s}$ and produces a signal with a $35 \mu\text{s}$ risetime. By digitizing the preamplifier signals and analyzing them in a digital signal processor, we can easily distinguish between these two risetimes and thereby count only the alpha particles emitted by the sample. Because the sheet of sample absorbs any alpha particle emitted from the sample tray before they reach the counter's active volume, the tray's emissivity does not contribute to the counter's background either. In addition, extensions of the method to the counter's sidewalls similarly allow us to reject alpha particles emitted from the sidewalls. We are thus able to obtain background rates that are up to a factor of 1000 lower than in conventional instruments that do not employ this active rejection principle.

Attempting to extend this principle to count at the $0.00001 \alpha/\text{cm}^2/\text{hour}$, level encounters difficulties, however. In particular, at this emissivity, there will typically be only 2.4 alpha particles per square meter per day! Since approximately 6 true counts are required to measure activity at the 95% confidence level, large sample areas are required if measurements are to proceed in reasonable time periods. Unfortunately, increasing the counter's anode area to a square meter raises its capacitance so much that the preamplifier noise levels swamp the alpha particle signals and make counting impossible.

In this SBIR we proposed to solve this dilemma by segmenting the single large area electrode into several smaller independent electrodes whose individual capacitances would be small enough so that we could still detect the alpha particles reliably. Since each electrode would then have its own electronics, we would have to capture signals from all of them in coincidence (since an alpha track might well deposit charge on more than one electrode), but we had already developed simple techniques for doing so for XIA's gamma-ray spectrometers. Therefore, in Phase I we proposed to show proof of principle by subdividing our original $1,800 \text{ cm}^2$ electrode into 4 square segments, each 625 cm^2 and demonstrating that signal noise on individual channels has been reduced as expected. Because the Phase II counter with a 1 m^2 segmented anode would require 16 such segments plus a segmented guard as well, we also proposed to design low cost signal processing electronics to instrument it in Phase II.

Our Phase I effort met our major proof of principle goals by carrying out our stated research plan. In particular, we showed that reducing the anode size by a factor of 4 in area reduced electronic noise by 3. Noise reduction did not scale exactly with electrode area because a significant fraction of electrode capacitance is to its neighboring electrodes and this term does not drop as fast as the area term. We also

developed an analytical model of signal generation as the charges in the track drift across the counter and showed that the features observed in the real signals closely resembled those predicted by the model. In addition, we captured events where the track charge was indeed collected by two electrodes and showed that, by summing them, we could recover the preamplifier signal shape appropriate to collection by a single electrode. We thereby showed that we could achieve noise levels that would allow us to analyze our signals with even higher precision than in our current instrument while being able to increase the measured sample area to the values needed to attain the proposed sensitivity. We also examined the Phase II electronics issue and produced a high level design showing that we could construct digital signal processing electronics that would fully meet our segmented electrode needs for a cost of less than \$500 per channel. The only Phase I goal we did not meet was to produce a software viewer that would have simplified the collection and inspection of multiple signal traces simultaneously. The reason for this shortfall was that we encountered unexpected difficulties with microphonic noise and had to expend our limited time and money resources in removing this noise in order to provide the critical, proof of principle demonstration that signal noise reduced with electrode size as predicted.

Table of Contents

1. Significance, Opportunity, Technical Approach & Public Benefit	5
1.1. Identification and Significance of the Problem (adapted from our Phase I proposal)	5
1.2. Technical approach	5
1.3. Anticipated Public Benefits	9
2. Phase I demonstration of technical feasibility	11
2.1. Phase I Research objectives	11
2.1.1. The overall goal of the Phase I research – proof of principle	11
2.1.2. Specific Phase I technical objectives (abbreviated from the Phase I Proposal)	11
2.2. Tasks proposed to meet the Phase I Objectives (abbreviated from the Phase I Proposal)	11
2.2.1. Objective 1: Modify existing detector	11
2.2.2. Objective 2: Firmware and software modifications	11
2.2.3. Objective 3: Noise tests and measurements	12
2.2.4. Objective 4: Phase II design and proposal	12
2.3. Modification of existing detector	12
2.3.1. Task 1: Segmented anode design and fabrication	12
2.3.2. Anode pickoff PC board design and fabrication	12
2.3.3. Build up test detector	14
2.4. Firmware and software modifications	14
2.4.1. Data collection software changes	15
2.5. Noise tests and measurements	15
2.5.1. Quiescent noise measurements	15
2.5.1.1. Noise source identification and removal	16
2.5.1.2. ADC traces and Fourier transforms	16
2.5.1.3. Detector calibration using α -source	18
2.5.1.4. Noise performance comparison to original detector	19
2.5.1.5. Pulse fitting issues in the new detector	23
2.5.2. Clean chamber background rate	27
2.5.3. Estimation of achievable backgrounds	28
2.6. Electronics design for Phase II	30
2.7. Other work: Nitrogen enclosure for sample preparation	33
2.8. Prepare Phase I report	34
3. Other products developed under this award	35
3.1. Publications – none	35
3.2. Patents - none	35

1. Significance, Opportunity, Technical Approach & Public Benefit

1.1. Identification and Significance of the Problem (adapted from our Phase I proposal)

The search for rare-events, such as neutrino or double beta decay detections is primarily a fight to reduce the incidence of background events to sufficiently low levels that the events of interest can be seen. Two of the major sources of background are cosmic rays, commonly muons but including other particles as well, and natural radioactivity in materials in the vicinity of the experiment, including both the materials of the environment and the materials from which the experiment's components are constructed. Cosmic rays are reduced both by operating deep underground and by using coincidence shielding techniques. Natural radioactivity in the environment is reduced both by selecting experimental locations low in natural activity and by shielding the experiment with low activity materials (e.g. antique lead). Since the experiment cannot be shielded from its own components, the solution there is to construct all components from materials having as low a concentration of radioactive impurities as possible. Similar issues apply in the detection of rare events at radioactive beam facilities where the environment cannot be tailored and the need for ultra-low activity shielding becomes even more extreme.

The DOE, which is a major supporter of rare-event search activities, therefore has an active interest in the development of ultra-low activity materials. By extension, there is also an active interest in developing instrumentation for characterizing such ultra-low activity materials, since it is impossible to improve something that cannot be measured.

In this proposal we seek to develop an instrument for characterizing the α -particle emissivity of material surfaces, an instrument whose lower detection limit will be at least 1000 times lower than that of the best extant commercial instruments. α -particle characterization is important to rare-event instrumentation for two reasons. First, α -particles emitted by materials constituting the detectors or their nearby environment cannot be excluded from detection by shielding, since they occur within the shielding, inside the detectors themselves. Second, because α -particles are emitted at various steps along the Uranium and Thorium decay chains, which also emit γ -rays at other steps in the chains, they can therefore also be a sensitive marker for the presence of these other members of the decay chains. As a historical example of these issues, we note that the ^3He -filled proportional counters used in the Sudbury Neutrino Detector had to be specially manufactured from ultra-low α -activity Ni in order to reach the detection limits required for that experiment's success.

The present state-of-the-art in low background α -particle detection (the Model 1950 Alpha Counter from IICO/Spectrum Sciences of Saratoga, CA) employs a wire chamber set up as an ultra-thin window, gas-filled proportional counter. α -particles entering the detector through the window produce ionization tracks in the counter gas. The electrons in the tracks are collected to the thin wire anode which is operated in proportional mode. After 20 years of identifying and removing sources of α -particles within the chamber volume, the detection limit of this device is set by cosmic background radiation at about $0.005 \alpha/\text{cm}^2/\text{hr}$. This process, which we refer to as "passive background reduction" becomes exponentially more difficult as background levels drop and has apparently reached an asymptotic value in this design.

1.2. Technical approach

In the proposed approach we will develop an "active background reduction" approach based upon digital signal processing and pulse shape analysis techniques to reduce background levels still further. In this method we digitize our detector's output pulses and then examine each event individually to determine whether it arose from an α -particle emitted from the sample or from some part of the detector chamber. By discriminating against the chamber-source α -particles we expect to lower our detection limit for sample-source α -particles by over 1000. Several innovations enhance the power of this approach: 1) we specifically shape our counter chamber to maximize pulse shape differences between different types of events; 2) we operate in ionization mode so that pulse risetime is proportional to charge

drift time; and 3) we introduce a veto electrode to specifically identify sidewall emissions. Using this approach, we already have results which indicate our ability to reach detection limits of $0.0001 \alpha/\text{cm}^2/\text{hr}$.

Figure 1 shows a schematic representation of an existing XIA low α activity counting system. The active volume is a Plexiglas chamber which comprises an upper shell sitting on a sample tray **1** and filled with ultra-pure nitrogen boil-off gas. The height R of the active volume is about 20 cm, significantly larger than the maximum α -particle range in N_2 (about 4 cm at 5 MeV). The sample is placed in the chamber on sample tray **1**, essentially covering the chamber floor. An α -particle emitted from the sample creates an ionization track **2** in the N_2 which ends a distance d_s from the top of the chamber, which is fitted with a large central anode **3** surrounded by a peripheral guard ring **4**. An α -particle emitted from the anode leaves an ionization track **5** that ends a distance d_a from the top of the chamber. An α -particle emanating from the chamber sidewall creates ionization track **6**. Filtered high voltage **8** supply biases the anode and guard electrodes, applying voltage V to the chamber and also to field shaping electrodes **10** (which prevent the electric fields from distorting near the edges of the parallel plate chamber), so that the electrons in the ionization tracks are swept to the anode and/or guard electrodes, depending upon their location within the chamber. The preamplifiers **7** amplify the resultant induced charge signals and pass them to digital signal processors, which can exchange coincidence timing signals, and capture the signal waveforms in their onboard memory. Data is then downloaded to the host computer via a DB25 connector **9**.




Figure 1: Schematic of the XIA ultra-low α activity counting system

In order to understand the operation of this counter, it is necessary to recall (per the Shockley-Ramo theorem) that (induced) current flows in the external circuit only while a charge (e.g. an electron) is actually moving within the counter's active volume due to the applied potential V , and ceases when the charge is 'collected' at an electrode. If this concept seems counter-intuitive, it is helpful to recall that the electric field from the electron reaches out to infinity, so that the external circuit is reacting to the presence (and motion) of the electron long before it reaches the "collecting" electrode.

The ‘induced charge’ integrating preamplifiers **7** then integrate this induced current and thereby produce a signal whose risetime equals the total electron drift time and whose amplitude increases with the distance the moving electron travels within the detector. In our nominally parallel plate geometry, the induced current is the same as the “current” of the drifting electron q itself: $i_q = q \cdot v_q / R$. The electron’s drift velocity v_e in N_2 equals its mobility μ_e times the electric field V/R (i.e. $v_e = \mu_e V/R$), so it is straightforward to compute both its drift time and the total induced charge integrated onto the preamplifier’s feedback capacitor C_f due to its motion.

Alpha track from the sample:

For a uniform ionization track of N electrons emanating from the sample **2**, the situation is a bit more complex, since first the track drifts across the chamber and then starts to be collected as the electrons in its leading end reach the anode. At this point the entire charge in the track does not disappear instantaneously, but instead decreases linearly in time as its component electrons are collected, vanishing completely as the last electrons (those originating nearest to the sample) finally reach the anode. Thus the induced charge signal resulting from the track’s motion is initially constant, as the ionization track drifts across the chamber, and then decays linearly as the track is collected. The preamplifier integrates this current and thus produces an output signal $S_a(t)$ that increases linearly in time while the track is crossing the chamber and then increases parabolically to an asymptotic final value while the charge is collected. Writing out the equations explicitly shows that $S_a(t)$ is then given by:

$$S_s(t) = \begin{cases} \frac{N \cdot e \cdot \mu_e \cdot V}{C_f \cdot R^2} t, & 0 \leq t \leq t_s \\ \frac{N \cdot e \cdot \mu_e \cdot V}{C_f \cdot R^2} \left(t - \frac{(t - t_s)^2}{2(t_R - t_s)} \right), & t_s < t \leq t_R \end{cases}, \quad (1)$$

where N is the number of electrons in the track, e is the electron charge, μ_e is the electron mobility in the N_2 gas, and C_f is the value of the preamplifier’s feedback capacitor. Here t_s is the time it takes for the leading electrons in the track to drift the distance d_s across the chamber and marks the transition between linear and parabolic behavior. t_R is the risetime (i.e. the time it takes for that trailing electrons to drift the distance R across the chamber) and S_{aMAX} is the maximum amplitude. Their values are:

$$t_s = \frac{d_s R}{\mu_e V}, \quad t_R = \frac{R^2}{\mu_e V}, \quad \text{and } S_{sMAX} = \frac{N \cdot e}{2C_f} \left(1 + \frac{R - d_s}{R} \right). \quad (2)$$

Alpha track from the anode:

When an alpha track emanates from the anode, Eqn. 1 becomes much simpler because the time t_s is zero – the “leading electrons” are generated right at the anode and have zero distance to travel to reach it. Therefore, again assuming a uniform charge distribution along the track, all we see is the parabolic portion of Eqn. 1:

$$S_a(t) = \frac{N \cdot e \cdot \mu_e \cdot V}{C_f \cdot R^2} \left(t - \frac{t^2}{2t_a} \right), \quad (3)$$

where the risetime t_a is the time for an electron to drift the distance d_a and the output signal amplitude S_{aMAX} become:

$$t_a = \frac{d_a R}{\mu_e V} \quad \text{and} \quad S_{aMAX} = \frac{N \cdot e}{2C_f} \frac{d_a}{R}. \quad (4)$$

Thus, in the case of alpha track emanation from the anode, we only see the parabolic behavior.

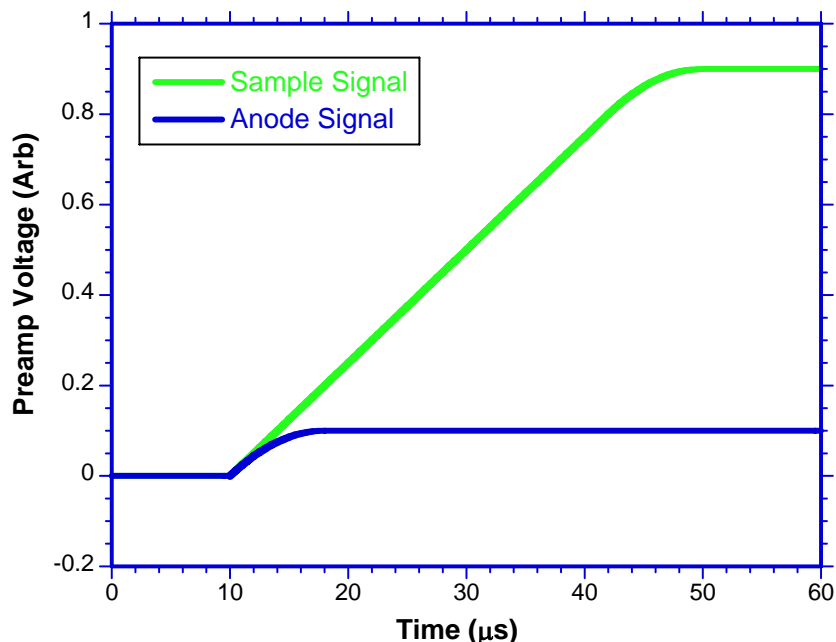


Figure 2: Analytic model results for signals from the sample and anode, assuming that the chamber height R was 20 cm, that the two track lengths were 4 cm, that the tracks were normal to the electrodes, and that the voltage V had been adjusted for a maximum transit time of 40 μs .

Figure 2 presents Eqns. (1) and (3) graphically, assuming that the voltage V was adjusted so that the maximum transit time was 40 μs , and assuming a chamber height of 20 cm and alpha track lengths of 4 cm. On this scale the anode signal risetime is only 8 μs (easily distinguished from 40 μs from a sample track) and its amplitude for the same energy is only $1/9^{\text{th}}$ as big. Clearly these two types of pulses will be easy to distinguish. We note that the parabolic nature of the traces is most prominent at the end of the signal risetimes. In the case of the sample signal, the transition between linear and parabolic behavior occurs at $16/20^{\text{th}}$ of the total drift time, or 32 μs , but the transition between the two forms is subtle and difficult to spot exactly. It is also important to stress that this is the infinite parallel plate solution and, in fact the traces from our original counter look very much like this, since the separation between the electrodes is 20 cm and their horizontal extent, including the guards, is 50. However, when we start to look at traces from our segmented anode, we can expect that there may be significant changes since they are only as wide (20 cm) as the separation to the sample (also 20 cm). The general shapes are preserved however, and this simple model is perfectly adequate to understand most of the physics of the counter.

In reality, the alpha particles are emitted at all possible angles from the electrode normal, so that their projected lengths on the normal can vary from 0 to their maximum track length. In the case of the sample tracks, the maximum drift length remains unaffected by this emission angle, since one end of the track always starts at the sample. For the anode tracks however, the drift time scales as the normally projected track length and so can vary from zero to the value given in Eqn. 4. Thus the case with the smallest difference between rise times is the normal emission case and if we design the chamber so that R is at least twice r , we can readily distinguish the two sources of emanation. Designing for distinguishability is an important part of our method. For the case of $R = 20$ cm and $V = 1000$ V, then t_R is of order 35 μs , t_{amax} is 10 μs or less and the two cases are readily distinguished.

An important design point is that, because the chamber is filled with N_2 gas at only 1 atm of pressure, it is extremely insensitive to both gamma-rays and cosmic ray muons – the former because their scattering cross section in N_2 is negligible, the latter because a muon ionization track contains charges equivalent to only about 50 keV, compared to the 5 MeV deposited by α -particles, and are difficult to see

above our noise levels, which are about 100 keV rms in our existing counter and will be about 50 keV rms in our proposed counter.

The technical problem:

At this current state of development, the interaction between two issues prevents us from extending the method an order of magnitude to reach detection limits of $0.00001 \alpha/\text{cm}^2/\text{hr}$. The first issue is that we need to be able to test much larger sample areas. Our current design can handle a sample of $1,800 \text{ cm}^2$ (0.18 m^2). At $0.0001 \alpha/\text{cm}^2/\text{hr}$, only 0.2 events/hour will occur, implying a need to count for approximately 30 hours to achieve measurements (6 counts) that exceed zero by 2.5 times their standard deviation. At $0.00001 \alpha/\text{cm}^2/\text{hr}$, there will typically be only a single event! While obviously terrific news from the point of the rare-event detector builder, the rate of screening materials for this level of activity will be intolerably slow. The obvious approach is to increase the sample area. At 1.0 m^2 , the count rate will become 0.1 event/hour, or 5 events every 2 days, which will allow several materials to be characterized per week.

The second issue, however, is that, because the anode and sample areas are approximately equal, and because signal noise scales linearly in detector capacitance – which is proportional to anode area, we cannot arbitrarily increase sample size while retaining our ability to distinguish between sample and anode source alpha particles. In fact, with the $1,800 \text{ cm}^2$ anode, noise levels are already too high to achieve the rejection ratios needed to reach detection limits of $0.00001 \alpha/\text{cm}^2/\text{hr}$. We are therefore in the unpleasant situation of desiring to both increase and decrease our anode area in order to meet our target detection limit.

In this SBIR we are proposing to solve this dilemma by segmenting our single large area electrode into several smaller independent electrodes, whose individual signals can, if necessary, be viewed in coincidence to detect α particle events. Our Phase I we have showed proof of principle by modifying our original $1,800 \text{ cm}^2$ electrode into 8 smaller segments and demonstrating that signal noise on individual channels has been reduced as expected. In Phase II we propose to proceed by constructing a counter with a 1 m^2 segmented anode and develop low cost signal processing electronics to instrument it.

1.3. Anticipated Public Benefits

The public benefits that will result from a successful conclusion to the proposed research fall into three areas. The first, which is of primary importance to the DOE, is that it will provide a major advance in the instrumentation available to develop ultra-low background materials required to support a variety of cutting edge physics experiments that seek to expand our understanding of the basic principles upon which the universe operates. The second area lies in the realm of environmental measurements of alpha radiation, which affects both the public at large and is of interest to the DOE in support of various environmental remediation efforts. Many of these measurements currently require samples to undergo complex concentrating procedures before they can be measured because existing instruments lack the sensitivity to measure them directly. These procedures are time consuming, expensive and, in many cases, difficult to calibrate because their process yield is uncertain. Direct measurement capability would offer a significant advantage in these cases.

The third area, which has both the largest social and economic impact, lies in the area of testing electronic packaging materials for alpha emissivity. Current roadmaps for high density digital circuitry call for a great expansion of the technique of directly mounting chips onto PC boards, using variants of bump bonding, which bypass the traditional enclosure of the chips in their own packages made of low α -particle emitting plastics. In the new designs the lead solder bumps are in direct contact with the chips, so that any α -particle emissions are extremely likely to cause “soft errors” in the chips by depositing ionization charge into memory storage cells or in the vicinity of sensitive transistor gates. Environmentally friendly “lead-free” solder alternatives are no better in this regard, being composed of equally contaminated metals such as Bi or Sn. Industry goals, therefore, are that all materials likely to come into contact with chips should have very low α -particle emissivities, $0.0001 \alpha/\text{cm}^2/\text{hr}$ or better. This includes not only solders, but also passivating coatings, underfill epoxies, and so forth. The economic necessity of meeting these goals is enormous, which means that all such materials will be

subject to rigorous screening requirements, by both material and chip manufacturers. As noted above, the best commercial instrument is less sensitive than this by a factor of 50, implying the existence of a substantial market for a new instrument sensitive enough to measure these values reliably. Sales of the current instrument are approximately 40/year with an installed base of about 500. Simply replacing this base with a new instrument would be approximately \$25 Million worth of business and there should be increased demand as well.

2. Phase I demonstration of technical feasibility

2.1. Phase I Research objectives

2.1.1. The overall goal of the Phase I research – proof of principle

The overall goal of the Phase I research was to demonstrate that our noise reduction scheme will work – that we can achieve improved signal quality levels by using 625 cm² anode electrodes compared to our current 1,600 cm² electrode. Achieving this goal would justify the Phase II effort of developing both the specialized electronics required to make the concept commercially feasible and the specialized sample handling procedures required to measure ultra-low background materials accurately.

2.1.2. Specific Phase I technical objectives (abbreviated from the Phase I Proposal)

Our Phase I effort had four specific technical objectives:

- 1) Modify an existing detector for test purposes: We proposed modifying one of our existing 1,800 cm² detectors for the proposed measurements – primarily by replacing the existing anode with a test design having four anode and four guard pixels and adding 8 preamplifier channels. However, for this comparison to be valid, because these are very low noise instruments we would require the same set of special low-noise design features as in the original and take the same care in implementing them.
- 2) Produce the firmware and software modifications required by the new configuration: Rather than design data collection electronics from scratch for this project, we proposed to borrow a pair of our general purpose, 4 channel, laboratory digital spectrometers and modify their firmware to work with signals with these very slow risetimes. With 75 MHz sampling rates, these modules normally detect 100 ns risetimes and need to be slowed down by almost 3 orders of magnitude.
- 3) Make noise tests and measurements: Once we have the hardware, firmware and software in working order, then we will make the necessary noise tests and measurements.
- 4) Create a detailed Phase II design and proposal: Finally, based on our Phase I experience, we will produce designs for a Phase II prototype with at least a 1 m² sample area and write the Phase I Final Report and Phase II proposal. A specific goal in Phase II would be to develop processing electronics that can be produced at a sales price of \$750/channel or less, so that implementing 20 electronics channels on the 1 m² detector (16 anode, 4 guard) would not increase its price excessively.

2.2. Tasks proposed to meet the Phase I Objectives (abbreviated from the Phase I Proposal)

To meet the above objectives, we proposed to carry out the following tasks, most of which are self explanatory in the overall context and will be discussed below in detail.

2.2.1. Objective 1: Modify existing detector:

- a) Task 1: New anode design and fabrication.
- b) Task 2: Anode pickoff PC board design and fabricate. This board collects signals from the anodes and delivers them to the preamplifiers.
- c) Task 3: Preamplifier mounting card and shielding: design and fabricate.
- d) Task 4: Build up test detector.

2.2.2. Objective 2: Firmware and software modifications

- a) Task 5: Adapt Pixie-4 gate array for alpha chamber timescale
- b) Task 6: Modify IGOR control software for Pixie-4.

2.2.3. Objective 3: Noise tests and measurements

- Task 7: Quiescent noise measurement – system shakedown. This task includes identifying and removing all stray noise sources, which proved to be challenging, even with our experience in building the existing counter.
- Task 8: Capture data sets at using an alpha source.
- Task 9: Analyze data sets to quantify noise levels.
- Task 10: Measure clean chamber background data set.
- Task 11: Develop rejection ratios and estimate achievable background rates.

2.2.4. Objective 4: Phase II design and proposal

- Task 12: Design 1 m² alpha particle counter
- Task 13: Phase I report (this section of Phase II proposal).

2.3. Modification of existing detector

2.3.1. Task 1: Segmented anode design and fabrication

Because recently completed material characterization measurements conducted in our original detector revealed that mirror-polished, chemically cleaned stainless steel has a low background rate of only 0.0014 α /cm²/hr, we decided to manufacture our electrode segments out of this material as it is affordable, readily available from a local distributor, and we can complete the cleaning process in house.

The anode electrode structure (see Figure 3) consists of 4 electrically isolated central pixels 400 cm² in size, with 4 electrically isolated L-shaped outer pixels of approximately 190 cm². The outer pixels are actually constructed of 8 individual plates of stainless steel which are electrically connected in four corner locations to create 4 effective L-shaped pixels. We had the pieces sheared at a local sheet metal house and holes drilled in their corners so they could be attached to a Plexiglas backing after we cleaned them with dilute HClO₃, rinsed several times in distilled water, and finally alcohol rinsed. They were assembled onto the backing in a glove bag filled with boil-off N₂ so that they would not pick up Radon contamination from the atmosphere, and then installed in the detector, which was then N₂ purged.

This work completes Task 1.



Figure 3: Segmented electrode structure (with signal collection cards in green and blue, at right). The screws that bring anode signals through the Plexiglas chamber ceiling are shown.

2.3.2. Anode pickoff PC board design and fabrication

In order to connect all eight electrode channels to the charge integrating preamplifiers, several new PC boards were designed (see Figure 3 and Figure 4). These were: a) Signal Collection card; b) Signal

Routing Card; c) Preamp Motherboard card; and d) HV Filter Card. The horizontal (green in Figure 3) signal Collection cards collect the signals from the anode pixels and route them to a central card edge connector. Mating with this connector is the long, vertically mounted Signal Routing card (blue in Figure 3), which delivers the signals up through a shielding wall (see Figure 4), to the Preamplifier Mother Board. The Preamplifier Mother Board was also redesigned to accommodate the additional preamplifiers required for each new segments and to separate the preamplifier signal routing functions from the HV filtering functions, which were placed on a robustly mounted HV Filter Card. Filtering the DC supply voltages for the preamplifiers is carried out on the Preamplifier Mother Board. This set of PC boards was designed to be capable of handling 4 segments. Two sets of the cards (see Figure 3) are thus required to instrument the new anode design.




Figure 4: Segmented Anode Electronics

An alternate version of the Signal Routing card containing microstrip technology has also been fabricated. We believe that by additionally shielding the signal lines with microstrips we can further reduce system noise pickup. However, since this technique was not used in our original detector, we have postponed studying its effectiveness in order to make the comparison between new and old designs as unbiased as possible.

Only a simple Al shielding box was required to house the end of the Preamplifier Mother Board that protrudes above the shielding wall. The preamplifier output signals were lead out through holes in the box walls to the digital spectrometer system using shielded cables terminated in SMA connectors.

This work completes Task 2 and Task 3.

2.3.3. Build up test detector

By designing the new electrodes and PC boards to have the same overall dimensions as the original detector, very few changes were required to integrate the modifications listed above into the prototype detector. An additional hole in the shielding wall was necessary to accommodate signals from the new segments on the far side of the electrode. Both sets of preamplifiers, preamp motherboards, and HV filter cards are shielded in aluminum enclosures to minimize additional noise pickup from external sources. We have further minimized any additional system noise pickup by removing all electronics from the immediate proximity of the sensitive front end preamplifiers. Power supplies, HV bias supply, and digital processing electronics are now all external to the detector, and power lines and signals are delivered in shielded coaxial cables. The issues involved in making these changes will be discussed below.

This work completes Task 4.

2.4. Firmware and software modifications

In the period between the writing of the Phase I proposal and its execution, XIA introduced a PXI version of its original CAMAC-based digital spectrometer, the DGF-4C that was proposed for this work. Since all the engineers were thus extremely current with its firmware design, we chose to use this newer module, the Pixie-4, for our work, reasoning that it would be significantly more efficient than having the engineers go back and try to reconstruct the details of the DGF-4C's firmware.

The Pixie-4 cards are PXI modules that are designed for instrumenting multi-element detectors in nuclear physics experiments. These four channel data acquisition modules are extremely flexible, but are optimally suited for detecting fast signals, such as the 20-200 ns risetimes typically generated in nuclear radiation detectors. As discussed above, the alpha particle tracks generated in our counter create signals with 5-50 μ s risetimes. Thus it was necessary to modify the Pixie-4 firmware in order to adapt the modules to our present – much slower – timescale.

Each DGF Pixie-4 card has four channels for processing signals from four preamplifiers. The signal chain in each channel comprises analog signal conditioning, a 75 MHz ADC, a field programmable gate array (FPGA) wherein the primary digital signal processing is accomplished, and an attached first-in-first-out (FIFO) memory, where captured signal traces can be stored. The four channels are serviced by a single, 75 MHz digital signal processor (DSP) whose job it is to handle such housekeeping chores as adjusting signal gains and performing any desired calculations. In particular, the DSP is responsible for capturing values from the FPGA's digital filter from which it reconstructs the energies of pulses detected by the FPGAs. A second FPGA and a Cypress PXI chip interface the module to the PXI backplane.

1.3.2.1 FPGA Design

Each FPGA contains the following digitally implemented components: a 200 ns trapezoidal trigger filter and a threshold comparator used to detect pulses in the preamplifier signal, pileup inspection and rejection, a “decimate-by-N” that sums consecutive blocks of 2^N samples into single values at $1/2^N$ th the clock rate in order to implement longer timescale energy filters, a trapezoidal energy filter whose peaking and flat-top times can be adjusted from 0.25 μ s to about 16 μ s, and circuitry to gate the digitized preamplifier signal into the FIFO. By timing the gating circuit it is possible to capture a few μ s of signal approximately centered on a detectable pulse. These captured signals are typically used for performing pulse shape analysis on solid state detector signals.

While analyzing captured pulses for risetime information is central to our approach for distinguishing the point of emanation of α -particle tracks, we need a much longer data trace, approximately 100 μ s so that we can establish the value of the baseline before the pulse and check it for pileup after it occurs. As originally implemented, then, the FPGA firmware was not optimally suited for our application. Specifically, even the fastest rising pulses in our system (6 μ s) are too slow to be detected by the standard 200 ns trigger filter and the maximum limit of 10 μ s storage time in the FIFO is too short to capture not only the risetime of a typical pulse in our system, but also the also 20 – 30 μ s of signal on each side of the risetime region so that the fitting algorithm will have enough data to achieve reliable fits and not be deceived by noise fluctuations. The solution to this problem was to insert a second decimation block

upstream of the entire previous design, allowing its entire time operation to be scaled by a factor of 2^M . In practice, using a decimation value of $M = 6$ worked well, slowing everything by a factor of 64. Our trigger filter is now adjustable over a risetime range of $0.64 \mu\text{s} - 6.4 \text{ms}$, and the FIFO can capture traces up to $300 \mu\text{s}$ in length.

This work completes Task 5.

2.4.1. Data collection software changes

As it turned out, only minor software modifications were required since, between submitting this proposal and its execution, our software engineers had developed a new application for another new XIA product, a very high performance 16 channel PXI card intended for instrumenting high speed cross strip detector arrays, the Pixie-16. This software already had adequate display capabilities and only needed to be modified to allow correct parameter values to be input and passed to our modified FPGA.

One issue that will have to be addressed in Phase II is that the Pixie-4 cards do not produce data in the same format as the modified Polaris spectrometers that we used to build our original alpha counter. Therefore all of the automated analysis software that we have developed for use on the original counter (in particular the extremely useful ability to click on a point in an "Risetime – Energy" plot and immediately be able to see its trace, the fit to its trace, and the resulting fit parameters and classification. One of our immediate Phase II tasks will be to harmonize the data formats so that we can not only recover this ability but also view the traces on an anode element's nearest neighbors.

This work completes Task 6.

2.5. Noise tests and measurements

2.5.1. Quiescent noise measurements

Once the detector system had been modified and assembled we had to identify and eliminate external noise sources. We spent an inordinate amount of time on this task, far more than we had budgeted for, because of two unforeseen issues. In the first place, the new detector seems to be far more sensitive to mechanical vibrations (microphonics) than the old design, even though they differ substantially only in the anode design. After much trial and error, we found, for a first example, that the fan noise from the Pixie-4's PXI crate added significant noise; for a second example, while we initially used a linear power supply for our preamplifiers to eliminate switching noise, we eventually discovered that their transformer "hum" (magnetostriction in the core) showed up in the detector noise spectrum as an enormous set of low frequency peaks. In the second place, because the processing electronics were set outside the counter body, we introduced a complex set of ground loops that picked up all sorts of random frequencies and radio stations and that took a long time to track down and eliminate. We eventually solved this problem by actually eliminating the solid ground strap between the counter chassis and the PXI crate/DC supply and connecting them only through the grounds of the shielded preamplifier output signal cables.

While this work consumed several months of frustrating effort, we had little choice, since it would be pointless to try to compare the intrinsic noise of the old and new systems if the new system noise were entirely dominated by external effects. As a result, we were forced to cut back on some of our later data set capturing tasks. However we have obtained enough data to clearly demonstrate that we have indeed reduced the intrinsic noise by the predicted amount due to the reduction in the area of the anode pixels.

In reducing system noise, we initially conducted a system shakedown and made quiescent noise measurements by capturing noise traces and examining both the traces themselves as well as their Fourier transforms. We then proceeded to calibrate our system using an alpha particle source of known energy. Determining the characteristic gain of our new system allowed us to make direct comparisons to our existing detector. Noise comparisons were finally made using captured pulses, since it is here that the noise has its most critical impact, both in setting trigger levels and in disturbing our pulse shape analysis algorithm. We therefore captured a large number of pulses from an alpha source and then computed the rms noise on the baseline section of each. These rms values were averaged over the full set of pulses to remove effects of pulse-to-pulse fluctuations and the resultant "mean rms" noise compared between the

old and new systems. As we will show, we obtained an overall reduction mean rms noise corresponding to our reduction in electrode area and capacitance.

Before direct comparisons could be drawn from measurements between our current and original system, the new detector had to reach a state of stable noise performance. We determined the noise performance of our system by using an oscilloscope to probe the outputs of the signal preamplifiers. Fourier analysis of these signals helped to identify polluting frequencies characteristic of various noise sources located somewhere upstream of our outputs. By systematically identifying and removing these sources we were eventually able to continue with our system's characterization and gain calibration.

2.5.1.1. Noise source identification and removal

We increased the amount of front end electronics in our detector by a factor of about four in order to instrument the new channels present with the segmented anode. This correspondingly increased the power consumption of our system and necessitated the addition of a large power transformer to supply this new demand. Additionally, the PXI crate - which houses and powers the DGF Pixie-4 cards - contains a new power supply. It has been our experience that linear and switching power supplies can be a source of low frequency noise at and around the line voltage frequency (60, 120, 180 Hz). Identifying pickup noise from these components is relatively straightforward as the offending frequencies readily stand out in a Fourier transform of a captured background noise trace. Eliminating these noise sources, however, is often a matter of trial-and-error. Through careful positioning of new components, additional shielding, rerouting power supply lines, and reducing the resistance of critical current return paths, we were eventually able to reduce noise pickup from our system to very low levels that were comparable to or less than those in the old system.

2.5.1.2. ADC traces and Fourier transforms

After minimizing the noise effects from new components in our system, we can compare the spectral noise density measured on one of our new pixel electrodes to that measured on our original detector. In the plots shown in Figure 5, Fast Fourier Transforms (FFTs) from the original detector are depicted in blue, while those from the new segmented detector are in red. In Figure 5a, the spectral response of the two detectors is compared in the low frequency regime (0-3.3kHz); in Figure 5b the range is 0 – 10 kHz; and in Figure 5c the range is 0 – 100 kHz.

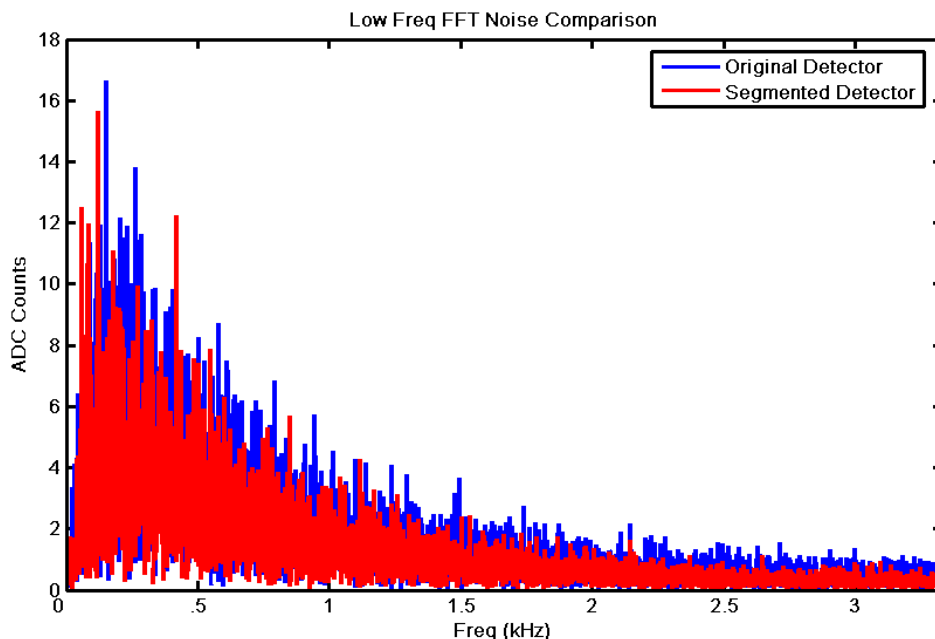


Figure 5a: Noise power spectrum (FFT) comparison between original and segmented anode counters, frequency range 0 – 3.5 KHz.

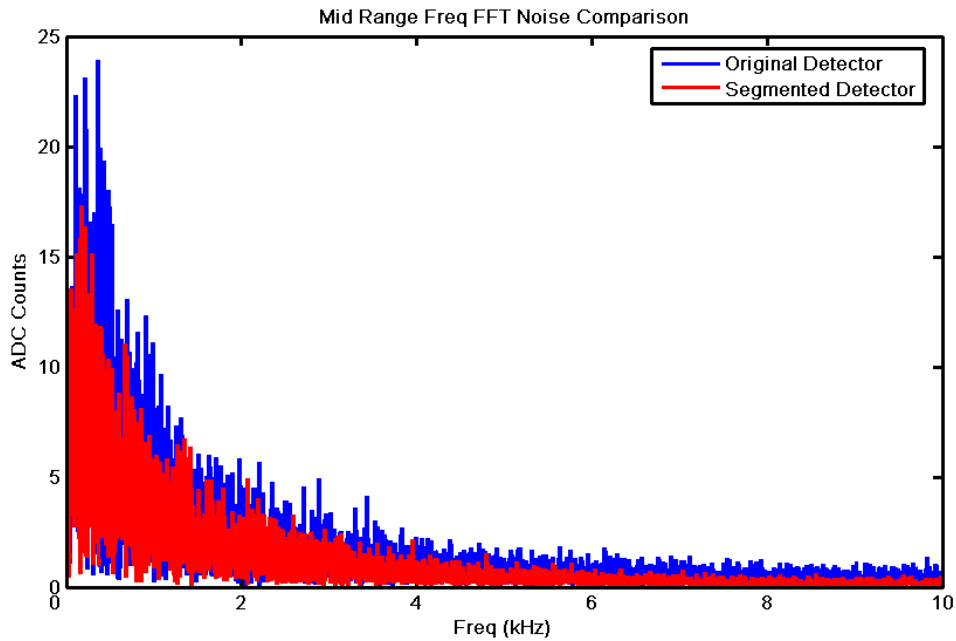


Figure 5b: Noise power spectrum (FFT) comparison between original and segmented anode counters, frequency range 0 – 10 KHz.

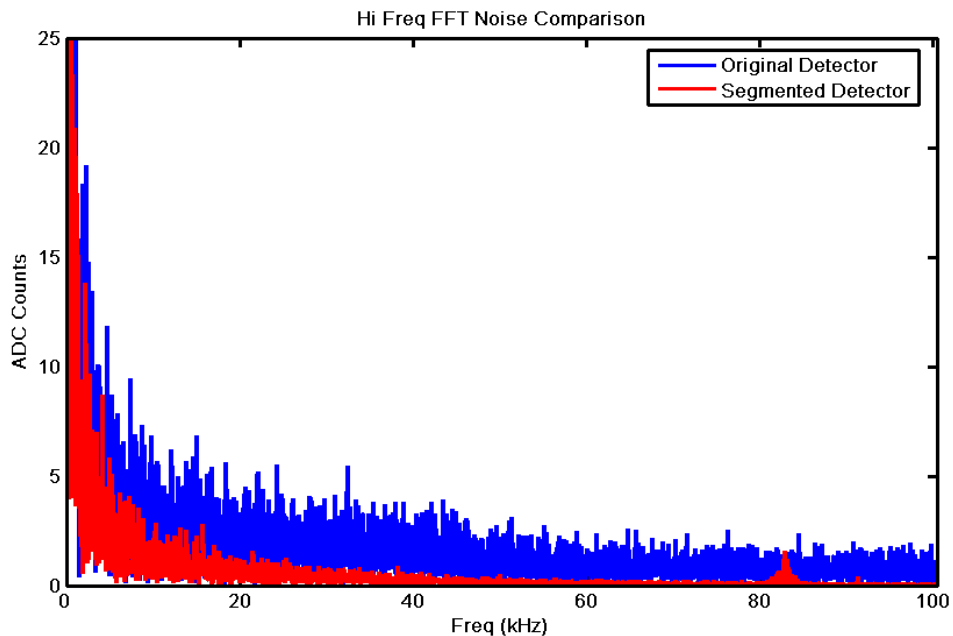


Figure 5c: Noise power spectrum (FFT) comparison between original and segmented anode counters, frequency range 0 – 100 KHz.

Looking at the FFTs in Figure 5a and Figure 5b we see that the new and old detector have a similar noise power spectra in low to mid range frequency ranges, with the new counter being consistently about 20% lower, a minor improvement. In the high frequency range, Figure 5c, above 10 kHz, the improvement is far more dramatic, with the new system noise being only about 20% of the old system noise. We believe that the majority of low frequency noise is due to power supply regulation and microphonic pickup, which is nearly identical between the two detectors. The effects disappear above 5-10 kHz, however, so the observed reduction of noise power at higher frequencies may be directly attributed to the lower capacitance of the segmented anodes and is a first demonstration of our thesis.

This work completes the requirements of Task 7.

2.5.1.3. Detector calibration using α -source

Following the system shakedown and the measurement of the improved noise power results of Sect. 2.5.1.2, we used the new counter to capture traces from our known alpha source so that we could calibrate the detector's gain. In particular, we placed our 250 count/sec ^{230}Th α -particle source (4.77 MeV) on the sample tray, directly underneath the center of one of the anode pixels. We collected 4000 events, and ran our standard analysis algorithm which measures risetimes and energies from individual signals. **Figure 6** depicts a scatter plot made from the 4000 event dataset.




Figure 6: Scatter plot of risetime versus energy of the 4,000 α -particle events captured in new segmented anode counter.

The dense red cluster of points on the right side of the figure results from essentially full energy alpha track events. The mean energy in this cluster is 2860 ADC steps, which corresponds to the full collection of charge generated by a 4.77 MeV alpha emitted from the sample. This value therefore characterizes our new counters gain at about 1.6 keV/ADC step. Events in the tail to the left of this cluster come from α -particles that lose varying amounts of energy as they depart the source at shallow emission angles. Points in the "halo" about the cluster are superimposed on electronic noise that distorts either their risetime or energy somewhat. The small cluster of events with risetimes less than 50 and energies between 0 and 250 are a combination of alpha particles emitted from the anode electrode itself, and noise triggers and are not readily mistaken for alpha particle events.

It is particularly important to note here the complete absence of points between 10 μs risetime and 25 μs risetime at full alpha particle energy. This means that we are not going to misinterpret a true alpha particle as something else and that, therefore, our counting efficiency for true alpha particles will be essentially 100%.

A similar gain measurement was previously conducted on our original detector using the same α -particle source. We present those results again here in Figure 7 in order to demonstrate the difference in gain (i.e. the voltage gain in the spectrometer's analog front end) between the two systems.

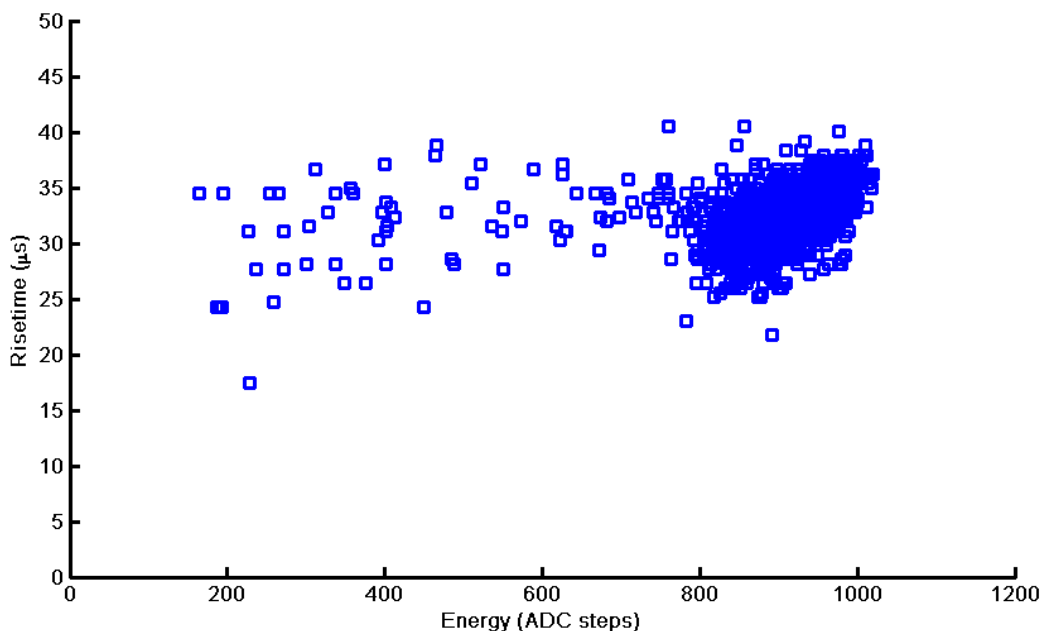


Figure 7: Scatter plot of α -particle events captured in our original counter

The mean energy in this distribution is centered about 915 ADC steps, which corresponds to a gain of 5.2 keV/ADC step. We note that in this figure there are no low energy, low risetime events detected. This is because, with the lower system noise in the new detector, we set our threshold lower and so were able to detect pulses that we had not been able to see in the old detector. For a true, one-to-one comparison between the two systems, we will need to repeat the new measurement with the gain adjusted threshold set to match the old system. We also notice that the correlation between risetime and energy that was present in the old data set has disappeared in the new data set. This correlation was presumably introduced by the response of our fitting algorithm to noise on the leading and trailing edges of the pulse's risetime and reducing this noise has removed it.

Comparing Figure 6 to Figure 7 we also note that the percentage width of the alpha pulse distribution has not noticeably changed between the two, whereas we might have expected the new plot to have a tighter distribution due to the lowered trace noise. Our current explanation of this is that, just because we have lowered the noise, we are now seeing the trace shapes more accurately and our simple three-linear-segment fitting function is no longer adequate to the task. We will comment on this further as we examine some of the traces in the following sections.

This work completes the requirements of Task 8.

2.5.1.4. Noise performance comparison to original detector

Noise in our system manifests itself as random fluctuations appearing on our signal traces. This is problematic because our analysis algorithm deals directly with these traces. In particular, our algorithm extracts both risetime and amplitude information via a least-squares method by fitting each trace (see Figure 8) with a trial function comprised of three straight lines: two horizontal lines of variable length τ (shown below at 30 μ s), and a variable length line joining the end of the first to the start of the second. The four adjustable parameters are the ending time and voltage of the first segment and the starting time and voltage of the third segment. The fitting algorithm varies all four parameters and then seeks out the set of values that has the minimum $t^2 = \chi^2/N$ value, where χ is the normal variance for the fit and N is the number of data points in the fit (which varies as the distance between the first and third segments is adjusted). Looking at the fit in the vicinity of the joints at both ends of the middle segment it is easy to see how noise fluctuations might easily cause risetime/amplitude correlations. It is also clear, however, that this fitting function is a pretty good approximation to the shape of the pulse's rising edge.




Figure 8: Signal trace collected from original detector with 3-segment fit overlaid

As evident from the figure above, the noise on the trace immediately before and after the pulse (i.e. in the first and third segment fitting regions) significantly contributes to error in the overall fits. By lowering this noise we hope to be able to set our fitting thresholds lower and produce more accurate results. Our thesis is that the majority of this noise is due to detector capacitance. Since capacitor noise scales linearly with capacitance and we have reduced the anode segment areas by a factor of four, we therefore expect to see trace noise reduced by a factor of about four. The factor will not be exactly four, since, in addition to parallel plate capacitance across the detector, there is also edge capacitance between the anode segments and also between them and the guard strip and this term is larger in the new design than in the old.

Our next step, therefore, is to make a direct comparison of trace noise between the two detectors. Figure 9 shows comparable, full alpha particle energy signal traces captured from the two detectors. By visual inspection, it is apparent that we have made a substantial improvement in noise performance.




Figure 9: Signal trace captured from the prototype detector with 3-segment fit overlaid.

To quantify this improvement, we will explicitly measure the amount of noise on the signals' baselines. Here we define the 'baseline' section as the $t_b = 30 \mu\text{s}$ of the signal trace immediately preceding the rising edge of the pulse, where t_b is the length of the first fitting segment. For each individual trace i we: first we find the baseline's average value v_i ; second, compute the standard deviation σ_i of the points from v_i over t_b , third, histogram the values σ_i for all 4,000 events in **Figure 6**, (giving **Figure 10**), and, fourth, fit with a Gaussian to compute the mean standard deviation $\langle\sigma\rangle_N$, which is the best estimate of the average noise for the new counter. Then we repeat this analysis for the dataset from the original detector (**Figure 7**), to produce **Figure 11** below and the value $\langle\sigma\rangle_O$ from the old data set.

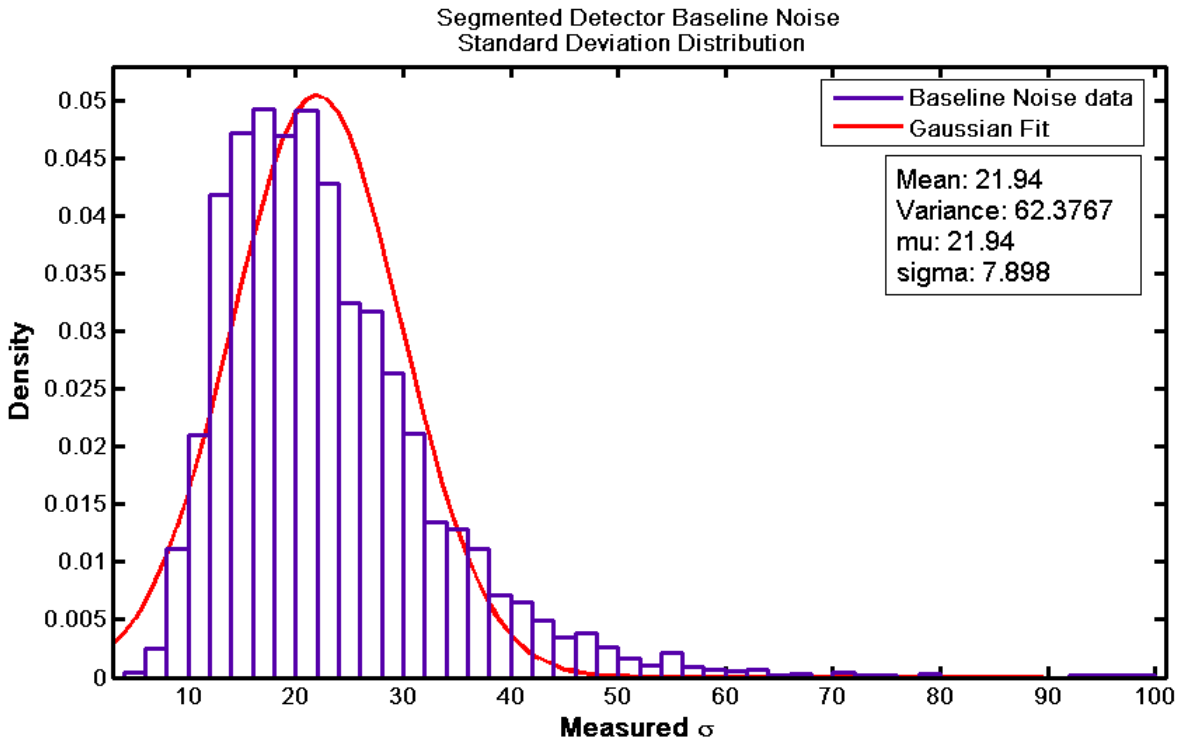


Figure 10: Distribution of baseline standard deviations σ_i from the segmented anode detector.

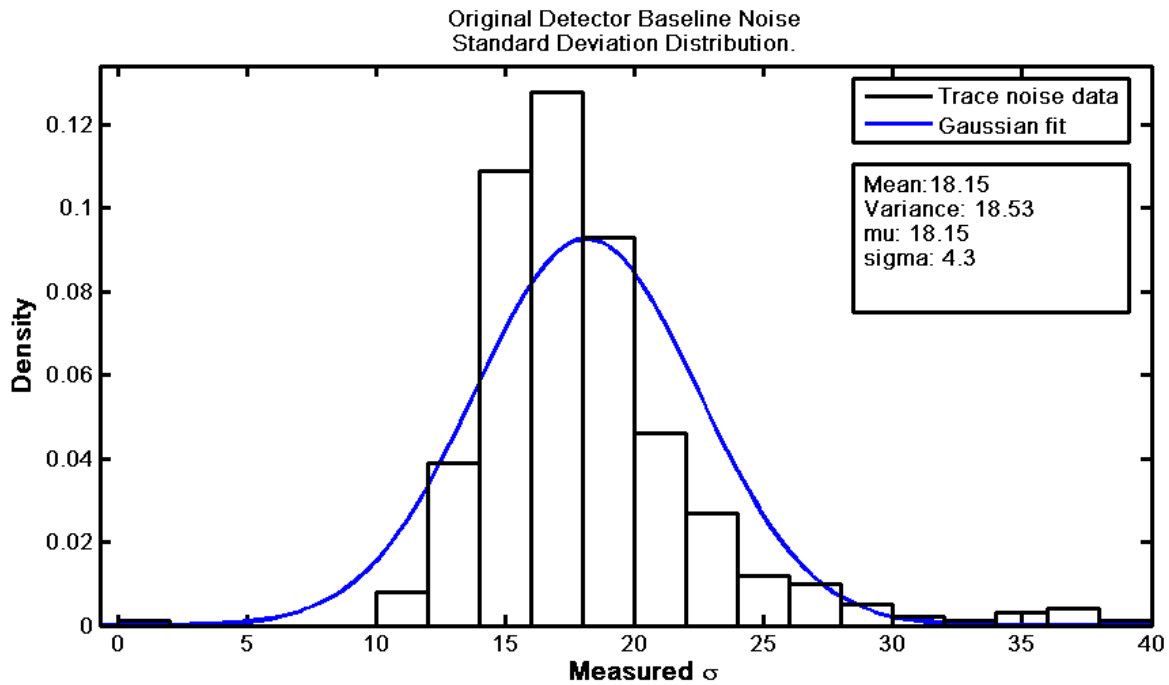


Figure 11: Distribution of baseline standard deviations σ_i from the original detector.

Case	Mean noise value (curve centroid)	Detector Gain	Scaled mean noise value
Original Detector	18.2 ADC	5.2 keV/ADC	94.4 KeV
Segmented Detector	21.9 ADC	1.6 keV/ADC	35.0 KeV

Table 1: Scaling measured mean noise values in ADC units into absolute KeV units.

The mean values of the two distributions $\langle\sigma\rangle_N$ and $\langle\sigma\rangle_O$ are the values "mu" shown in the figures. In order to complete the comparison, we finally need to scale the results into absolute (KeV) units by multiplying by the two systems' different gains. This is done in **Table 1**, where, after adjusting for the much larger gain in the original detector, we see that we have reduced average trace noise from 94 KeV to 35 KeV, a factor of nearly 3, which is approaching the factor of four predicted from the change in anode capacitance due to change in anode area (1,600 cm² goes to 400 cm²). The result is not expected to be exactly four because, although we have reduced the area by a factor of four, we have only reduced the segment's edge capacitance by a factor of 2 (i.e. from 160 cm to 80 cm of ground plane 0.6 cm away).

This result definitively proves our thesis that we can reduce noise levels by going to a segmented anode structure and that, by using segmented anodes, we can thereby extend our pulse shape analysis method of rejecting alpha particle tracks anodes that do not come from the sample being measured to measure samples of any size desired. In particular, we will be able to build a 1 m² alpha particle counter in Phase II that will be able to achieve detection limits of 0.00001 α /cm²/hour.

2.5.1.5. Pulse fitting issues in the new detector

As we stated in the Phase I proposal, we expect pulses in the new detector to differ from those in the old one in at least the following ways:

- lower electronic noise, as discussed above;
- different signal shapes since the electrodes no longer act as infinite parallel plates;
- signals split across neighboring electrodes.

The traces shown **Figure 12** illustrate the issues. These were captured from a pair of neighboring electrodes using our standard alpha source placed on the sample tray at a point under one edge of the electrode pair. The figures show a sequence of cases, starting with (**Figure 12a**) an alpha track that deposits all of its charge on electrode segment 1 but comes very close to reaching electrode segment 2. The next case (**Figure 12b**) is an alpha track that reaches a bit under segment 2 and deposits about 15% of its charge there. The last case (**Figure 12c**) is a track that deposits charge nearly equally (60/40) between the two segments.




Figure 12a: traces from a single alpha track that was collected on a single electrode segment 1 (red) but came very close to electrode segment 2 (blue).




Figure 12b: traces from a single alpha track that was about 85% collected on a single electrode segment 1 (red) but deposited about 15% of its charge on electrode segment 2 (blue).




Figure 12c: traces from a single alpha track that crosses between electrode segment 1 (red) and electrode segment 2 (blue).

In order to be sure that we understood these traces, we modeled them using a LabView code developed to model induced charge effects in cross strip semiconductor detectors. This code uses the method of image charges to compute the induced charge on an array of rectangular pixels and, by monitoring those charges in time, computes the signals that integrating preamplifiers attached to those pixels would see. Details are in XIA's US Patent 6,169,287. In this case we set the pixel size to 20 cm square with a drift length of 20 cm as well, and adjusted the electrons mobility to give a 40 μ s crossing time under 1000 V applied bias. Ion motilities were set to zero. We investigated four cases for a compact charge (as opposed to a track). It is important to note that, just as in the real case, when charge is created in the detector it initially consists of an equal number of holes and ions and appears to be electrically neutral when viewed from outside. Only after charge separation occurs, as the electrons and ions drift apart under the applied field, do they induce different charges on the electrodes. The investigated cases were:

- Charge created under the pixel center and drifting 20 cm to the anode (sample event).
- Charge created under the pixel center and drifting 4 cm to the anode (anode event).
- Charge created in pixel corner (5 cm from both edges) and drifting 20 cm to anode.
- Charge created in pixel corner (5 cm from both edges) and drifting 4 cm to anode

The results are shown in Figure 13a and Figure 13b. We see that, in the centered case, the charge generated close to the anode produces a signal that looks very like our parallel plane theory, with a

linearly rising signal until the charge collects. The signal induced on the neighboring segment is small and becomes negative as the moving electrodes leave the positively charged ion behind. The charge moving across the full detector now has a curved profile, becoming straighter as it approaches the anode and the anode becomes relatively more "infinite" in extent. Similarly, the charge induced in the neighbor first rises as the moving charge gets closer to the anode plane and then dropping as it gets close to the anode and all its field lines collapse onto the first electrode. In this case the remaining ions do not contribute since they are immediately collected at the sample plane and not left hanging in the counter gas.




Figure 13a: Traces from charge generated under anode segment center and drifting 20 cm (red/violet) and 4 cm (blue/green) to the anode.




Figure 13b: traces from charge generated under point inset from anode segment corner by 5 cm both ways and drifting 20 cm (red/violet) and 4 cm (blue/green) to the anode.

When we look at the case where the charge is generated closer to the corner of the segment, all the effects seen before exaggerate (and will continue to do so as the generation point gets even closer to the segment edges). In the case of the sample generated charge, the signal seen by the neighbor is now

starting to get quite large (it can be as much as 30% of the main signal at the very edge of the segment). Similarly, the curvature of central segment signal is much more pronounced. We would have a very hard time generating an accurate fit to this curve with our three-linear-segment fitting function. Meanwhile, the anode generated charge continues to behave fairly classically. The signal generated on the neighbor is more negative now because it sees more of the ion's field after the electrons' departure than it did when the ions were farther away under the electrode center.

Comparing Figure 13a to Figure 13b shows a particularly interesting feature of "anode generated" charge collection. Notice that in Figure 13a the anode signal maximum is about 0.35, while in Figure 13b it reaches 0.45. Meanwhile the neighboring signal went from -0.05 to -0.10. Thus we see that the amount of residual ion charge can affect the amplitude of the anode signals. Noting that the central segment actually has 2 nearest neighbors, we see that the real charge from the electrons only is $0.35 - 2 \cdot 0.05 = 0.25$, while in the corner case it is $0.45 - 2 \cdot 0.10 = 0.25$, which is the same in the two cases. This means that if we want to correctly estimate collected charges from anode events we will have to make some correction for the distance they travel to be collected. A Phase II task will therefore be to model real charge tracks, rather than localized charges, to fully understand the issues.

This simple analysis suggests that summing neighboring segments' signals will then make them behave more like the infinite parallel plate analysis. We test this idea using the traces first presented in **Figure 12a-c**. These traces are repeated, with their sums, in Figure 14a – c. The shape improvement is striking! All three curves, which had previously shown significant degrees of curvature, are now quite straight, and may be expected to perform well with our three-linear-segment fitting algorithm. We also notice that, because the noise of the two segments added in quadrature, it does not appear to have significantly worsened, particularly when compared to the noise shown from the original detector in Figure 8. Further, after the summation, all the pulses have essentially identical amplitudes, which was not the case at all prior to summation. It is therefore clear that, in order to achieve optimum results from our new detector design, we will have to pay particular attention to signal summation issues in Phase II.




Figure 14a: Figure 12a repeated with the two signals added.




Figure 14b: Figure 12b repeated with the two signals added.




Figure 14c: Figure 12c repeated with the two signals added.

These summation results do conclusively demonstrate that signals from neighboring segments can be successfully summed to reconstruct complete signal traces whose behavior is quite close to that predicted by the infinite parallel plate model and which will therefore be easy to analyze, once the issues of automating the addback process are dealt with.

This work completes Task 9.

2.5.2. Clean chamber background rate

Because of the time required to eliminate spurious noise sources and obtain the excellent noise reduction results shown in the previous section, we were not able to complete the software required to operate our new counter in full chamber counting mode and therefore were unable to use it to take a "clean chamber" (i.e. using the lowest background sample that we currently have available) background run. Instead, we present a clean background run taken on our old detector, which, since we have shown it to have higher noise levels should provide an upper bound on what we can expect to achieve. This




Figure 15: Scatter plot of events capture during an 8 hour run using the previous detector and a polished stainless steel sample pan.

was only an 8 hour run, and we measured 21 blue events that passed both the risetime and energy cut tests. The large numbers of red points that failed the tests are primarily noise events (since we set our thresholds low to be sure not to miss any real events), but there are a few odd points that appear to have risetimes that correspond to points of emanation neither on the sample nor on the anode, but from the chamber gas itself. In Phase II we will investigate the source of these pulses, to determine whether they are real (and if so what their genesis is) or some sort of rare noise artifact.

Twenty one good alpha events over 1800 cm^2 in 8 hours corresponds to a sample emissivity of $0.0015 \alpha/\text{cm}^2/\text{hour}$ from our stainless steel sample tray. Since this is a property of the tray, we expect that we would measure the same value using the new counting chamber. It is also clear that there are only three "dubious" events in this scatter plot, those at amplitude about 400 ADC steps and risetimes of $22 - 26 \mu\text{s}$. These points appear possibly to be part of the red noise distribution of points just on the other side of the energy cut line and should be eliminated in the improved detector. Treating these points at this detector's "error rate", we divide by 8 hours and the $1,800 \text{ cm}^2$ sample area to get a limiting background rate of $0.0002 \alpha/\text{cm}^2/\text{hour}$, which is encouraging and strongly suggests that we will be able to successfully reach our target of $0.00001 \alpha/\text{cm}^2/\text{hour}$ with a larger sample area and the better pulse shape algorithms that we will be able to develop to take advantage of the lower noise levels that come with our new segmented anode design.

This work completes Task 10.

2.5.3. Estimation of achievable backgrounds

We attempted to estimate the achievable backgrounds with the new, segmented electrode, by placing our 250 cps alpha source in the center of one of the anode electrodes and capturing 10,000 events. Our plan was to estimate how many of these would be mistaken for alpha events coming from the sample. The results are shown in Figure 17 and turned out not to be useful for the intended purpose.




Figure 16: Scatter plot of risetime versus energy of the 4,000 α -particle events captured in new segmented anode counter with an alpha source placed in the center of an anode segment.

There are two problems with this data. First, the fitting algorithm is not working correctly, as is evidenced by the presence of bands in the risetime values. We will discuss this issue in more detail in the Phase II proposal, but believe that it is essentially due to the improved energy resolution that makes the fitting function less smoothly varying as the segment endpoints are adjusted.

The more serious issue is that, at even with a low source rate of only 250 cps, we are seeing such a large number of pulse pileup events that they are swamping the fitting errors or other pulse identification artifacts that might cause anode events to be mistaken for sample events at the very low data rates we will actually encounter. This is because, when two events pileup, they can sum their energies, risetimes, or any combination in between. Thus, for example, two anode pulses with risetimes of 10 μ s can add to produce a "single" pulse with a risetime of 25 μ s, which our fitting algorithm classifies as a sample event.

Figure 17 shows a pileup pulse that demonstrates how longer risetimes can occur than the sum of the individual risetimes. One pulse followed another by about 15 μ s, with the result that the fitting algorithm assigned the total pulse a risetime of 35 μ s and an amplitude of 1,450 ADC steps. Looking back at Figure 6, we see that this could easily be classified as a "partial energy alpha pulse.




Figure 17: Trace of a piled up pulse that was found in **Figure 6**.

We can estimate the number of pileups that we would expect in our data set. Assuming that we are looking in a $50 \mu\text{s}$ region for good risetime values, then the chance of a second pulse coming in that interval, given the presence of a first pulse, is just $\text{ICR} \times 50 \mu\text{s}$. The alpha source strength is 250 cps, so the probability is 0.0125. Multiplying this by 10,000 events gives an expected 125 pileup events. In a real experiment, of course, the ICR will be more than 10,000 times lower (less than counts/minute) and so pileup will not be an issue. In order for this measurement to be meaningful, we will need to repeat it using a much weaker source (say 2 cps). In the final analysis, we cannot use pulse shape analysis to remove the piled up pulses just because the normal test for pileup is an exaggerated risetime. In the present work there are also real mechanisms that lead to exaggerated risetimes, namely sample alphas.

Therefore, for the present, we will have to make do with an estimate derived from our previous detector, where we showed that we could reject anode events with at least a 1000 to 1 ratio. If the anode electrodes are then made of our low emissivity stainless steel, which give off $0.0015 \alpha/\text{cm}^2/\text{hour}$, we would expect to see 15 α particles/hour, of which we would fail to recognize 0.015. Therefore, in a 48 hour measurement period (the time required to make measurements at the $0.00001 \alpha/\text{cm}^2/\text{hour}$ level) we will then see 0.72 falsely identified anode events compared to the 5 events we would see from a true $0.00001 \alpha/\text{cm}^2/\text{hour}$ sample. Thus the error rate is non-negligible, but still small compared to the sample rate. To remove even this error component, in Phase II we will investigate making our anodes out of materials with much lower alpha emissivity – either electroformed Copper or CVD Nickel. These materials have emissivities more than 2 orders of magnitude lower than our stainless material and, coupled with our active background rejection, will completely eliminate this source of error.

This work completes Task 11.

2.6. Electronics design for Phase II

Task 12: Produce a preliminary electronics design for the Phase II detector at \$500/channel.

The goal of this objective was to produce a digital signal processing design that can process a relatively large number of anode/guard channels at a modest price. Our target goal was be \$500/channel, compared to the \$3,000/channel cost of the Pixie-4 cards we used in Phase I. This cost would result in a total instrument electronics cost (of final sales price) of between \$10,000 and \$15,000, depending upon whether the optimum number of anode electrodes is 16 or 25. This price is nicely in line with a target instrument price of \$40K to \$50K.

We were able to accomplish this by taking advantage of the fact that the signals we are looking at have quite slow risetimes and therefore need digitization rates only in the 1-2 MHz range, whereas low-cost ADCs can easily run at rates of 40 MHz or more. This suggests that considerable costs savings could arise if the signals were multiplexed so that a single ADC and gate array processor could service a large number of channels. This forms the basis of our conceptual electronics card design.

A conceptual sketch of the processing electronics is shown in Figure 18. Here we are contemplating putting 32 channels onto a single 3U card, the same size that currently carries 4 channels on the Pixie-4 card. By using this format, we will be easily able to adapt the card to the PXI format in the future as a commercial signal processing product completely apart from alpha counter applications.

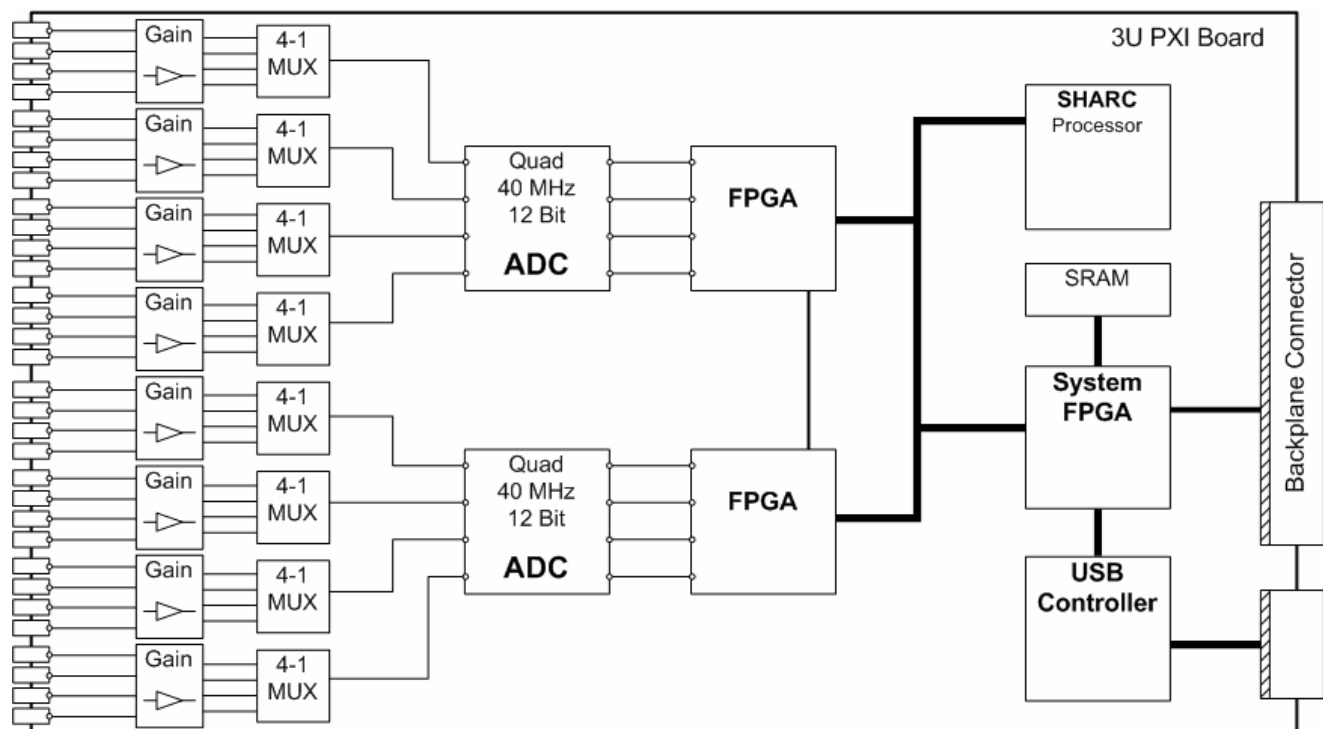


Figure 18: Conceptual diagram of 32 channel digital signal processing card for alpha counter.

The design works as follows. Each channel passes through an analog input buffer stage with gain. This allows us to match the preamplifier output signals to the input range of the ADC. This stage also applies Nyquist filtering to the input signals. According to Nyquist's sampling theorem, there cannot be any frequency in the sampled signal that is greater than $\frac{1}{2}$ of the sampling frequency. Otherwise that frequency is aliased into the signal at a lower frequency as noise and cannot be removed. As we will show below, we will probably be running the ADC at 48 MHz, so that we sample each of the four multiplexed signals at 12 MHz, with the samples between the different signals being interleaved as shown in Figure 19. In practice the same sampling pattern will be duplicated on each of the Quad ADC's four inputs. As a result, our Nyquist filter will be designed with a 3 dB rolloff point of about 3-4 MHz so that it will be well down at the 6 MHz Nyquist frequency.




Figure 19: Sketch showing four signals with multiplexed sampling by a single ADC input. The first few samples are numbered consecutively in the order collected by the multiplexer.

Following the buffer/Nyquist filter stage is an analog multiplexer chip (e.g. the Analog Devices ADG759, which settles in 14 ns and costs \$1.25) and then into the Analog Devices AD9229 ADC, which is a Quad-50-65 MHz part and costs about \$50. This is an interesting part in that it accepts differential analog inputs and supplies digital information as differential serial output. This is important both because differential signals are inherently low noise and also because this means that we only need to make a small number of input connections to the FPGA that follows the ADC. This means in turn that we can use a much smaller FPGA package, which is both cheaper and uses less board space, reducing board costs. In this design the ADC will be run at 48 MHz, which means that the multiplexer will have 20.8 ns to settle, and each of the 4 multiplexed input signals will be sampled at 12 MHz. This is already far faster than we will need. Thus, in the FPGA, our first step will be to "decimate" the incoming signals by summing consecutive sets of six signal samples, producing a single, averaged sample value at 2 MHz, or once every 500 ns. This will be adequate to produce 20 samples on the rising edge of a 10 μ s risetime pulse, which will be more than adequate.

The FPGA will be a medium size device in the Spartan-III family, which includes onboard multiplication (which we will not use here) and large blocks of memory (which we will). In this design we see that each FPGA will be processing 16 channels worth of information that arrives in four data streams. All pulse detection and trace capture will be implemented in the FPGAs, using variants of the same FPGA firmware that was developed for this project. The only development issue will be whether to process the pulses using interleaved filters or to break the signals out into 16 completely parallel processing chains inside the FPGA. This is an engineering decision that will depend upon an analysis of the amount of FPGA real estate and FIFO memory that is needed for each channel (assuming they are done separately) compared to the added complexity of doing them in an interleaved fashion. We will apply our usual design methods to this issue – develop conceptual circuits, estimate required resources, develop simulation models to evaluate reliability and timing constraints – and then select an approach. The interface to the onboard SHARC processor, system FPGA, and USB controller will borrow heavily from existing XIA designs.

The card is being designed as a 32 channel device for two reasons. First is that with Quad multiplexers and Quad ADCs, the number 16 is a rather natural number to pipeline into a processing FPGA. Our initial space studies indicate that this will not consume a great deal of space and so we plan to put two of them on the same card. A simple digital line then allows the two FPGAs to share common trigger information so that either can know if the other discovers a signal pulse. Since our Phase I electrode design uses 20 channels (16 anode and 4 guard) we need more than 16 channels in any case. With 32 channels we can also consider the possibility of a 25 anode design at small additional cost if the gain in signal to noise appears to be promising. If it turns out that we can only fit 16 channels on a single board, then we will implement connections between the two boards through the backplane so that they can share the required common trigger information.

Initially the SHARC processor will not be used. However, in the second year of the work we will evaluate whether it is feasible (and/or useful) to transfer our pulse fitting algorithms to the SHARC so that only fit data and no traces need to be transferred off the card. If feasible, this would be a significant step toward producing a completely stand-alone alpha counter.

We note that the primary function of the System FPGA is to handle initial bootup of the processing FPGAs and the SHARC, and after that to handle communications between the card and the external world through the USB interface chip. These are all functions that have already been implemented on other XIA products and do not carry any significant risk. In fact, even the multiplexing design does not really carry much risk either since all parts of the circuit chain will be modeled, either in SPICE or with Xilinx's digital modeling tools, before we ever start layout.

The per channel costs of this design are quite low. The expensive parts, the \$100 FPGAs and \$200 SHARC, are shared by 16 and 32 channels, respectively, so that the total parts cost per channel is only about \$25 - \$30, meaning that the complete card can be sold and supported commercially for less than \$500/channel, sales price being 8 to 10 times parts cost as a good rule of thumb. This design therefore completely satisfies our Phase I circuit design goals.

2.7. Other work: Nitrogen enclosure for sample preparation

Since one of our Phase II goals would be the development of sample handling methods that could prevent sample exposure to air, we invested a few thousand dollars in a custom built N₂ housing for our old detector, plus a pair of movable glove boxes for sample preparation with the intention of placing our counter inside this housing and delivering samples to it without having to expose them to air. A photo of the units is shown in Figure 20. It is sitting on a table that is isolated from the floor by pneumatic isolators that we have found to work very well. The white HDPE plastic structure is the housing for the detector, the clear Polycarbonate box is a sample glove-box. The large holes are where the gloves are to be attached. The sample box seals to the housing and a pair of doors (with gas purge between them) can be opened to pass the sample into the detector. The extension on the housing is to accommodate the sample drawer that is used to insert samples into the counter.



Figure 20: N₂ housing for the old alpha detector, including a sample glove box.

Our experience with this structure was not good overall, which meant that we learned a lot about how the problem should really be addressed, as we will see in Phase II. The major problem was that it was very difficult to operate our counter inside the N₂ housing since we couldn't touch anything without completely disassembling the housing. Specifically, this excluded access to the counting chamber for maintenance or cleaning. It also meant that all the controls had to be rerouted. Overall, we have concluded that the counter should be designed so that it provides its own N₂ operating environment to assure that the sample is never exposed to air.

The sample handling boxes, on the other hand worked fine. They could be purged separately, sample manipulations carried out with the gloves, and then easily attached to the counter housing when it was time to transfer the sample to the counter. Having them as separate units seems like a very good idea. For example, when a sample arrives (protected by some wrapping), it can be placed into a box and purged for a few days to remove both any Radon and let any daughter products decay. Then the packaging can be unwrapped and the sample moved to the counter. After counting, the sample can be returned to its protective wrappings and then returned to its owner.

2.8. Prepare Phase I report

Task 13: Phase I report.

This report completes and satisfies this Task.

3. Other products developed under this award

As noted below, no other products were produced under this award than the work documented above.

3.1. Publications – none

3.2. Patents - none