FERMILAB-FN-1014-LDRD

2014 Fermilab | Laboratory Directed Research & Development Annual Report

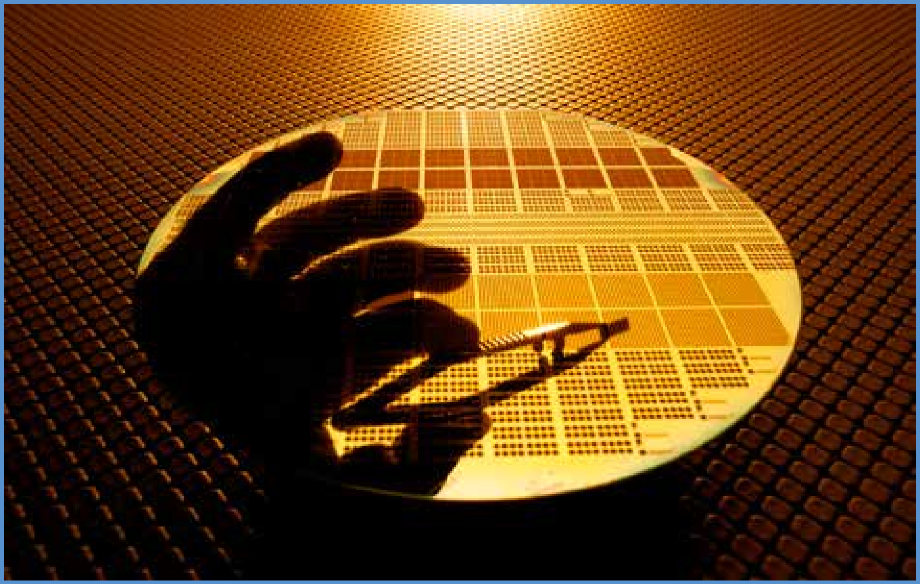

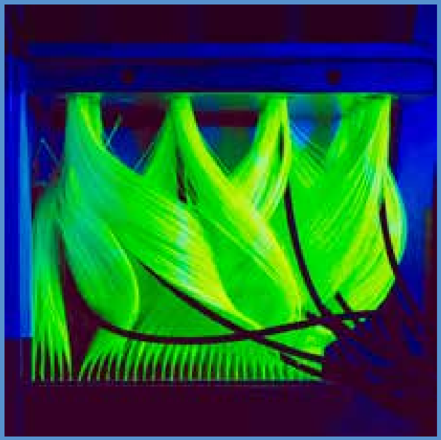

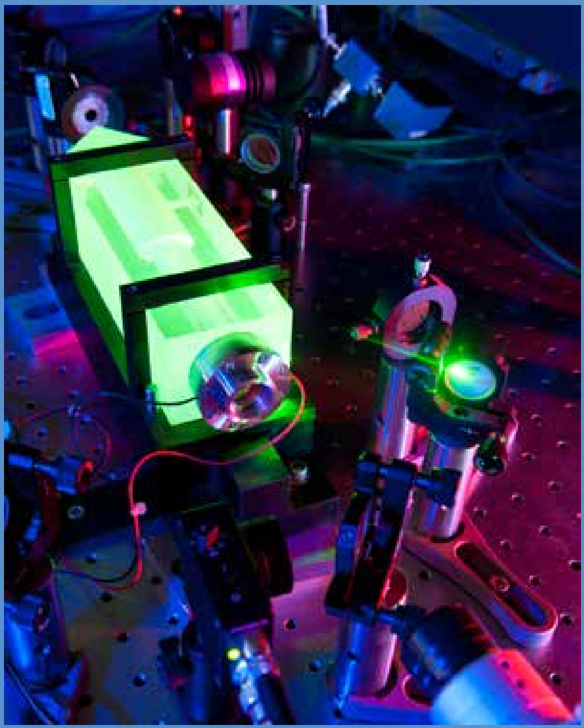

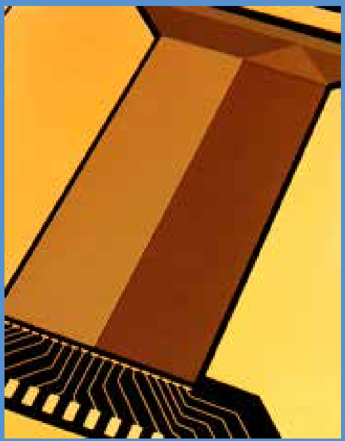

Fermilab

U.S. DEPARTMENT OF ENERGY | Office of Science

Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy.

Table of Contents

I. Executive Summary

After initiation by the Fermilab Laboratory Director, a team from the senior Laboratory leadership and a Laboratory Directed Research and Development (LDRD) Advisory Committee developed an implementation plan for LDRD at Fermilab for the first time. This implementation was captured in the approved Fermilab 2014 LDRD Program Plan and followed directions and guidance from the Department of Energy (DOE) order, DOE O 413.2B, a "Roles, Responsibilities, and Guidelines, …" document, and examples of best practices at other DOE Office of Science Laboratories. At Fermilab, a FY14 mid-year Call for Proposals was issued. A LDRD Selection Committee evaluated those proposals that were received and provided a recommendation to the Laboratory Director who approved seven LDRD projects. This Annual Report focuses on the status of those seven projects and provides an overview of the current status of LDRD at Fermilab. The seven FY14 LDRD approved projects had a date of initiation late in FY14 such that this report reflects approximately six months of effort approximately through January 2015. The progress of these seven projects, the subsequent award of six additional new projects beginning in FY15, and preparations for the issuance of the FY16 Call for Proposals indicates that LDRD is now integrated into the overall annual program at Fermilab. All indications are that LDRD is improving the scientific and technical vitality of the Laboratory and providing new, novel, or cutting edge projects carried out at the forefront of science and technology and aligned with the mission and strategic visions of Fermilab and the Department of Energy.

Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy.

II. Program Overview

Beginning in FY 2014, Fermilab has initiated a LDRD program as authorized by a DOE order, DOE O 413.2B, to enhance and realize the mission of the laboratory in a manner that also supports the laboratory's strategic objectives and the mission of the Department of Energy. LDRD funds enable scientific creativity, allow for exploration of "high risk, high payoff" research, and allow for the demonstration of new ideas, technical concepts, and devices. LDRD also has an objective of maintaining and enhancing the scientific and technical vitality of Fermilab.

LDRD is able to fund employee-initiated proposals that address the current strategic objectives and that better position Fermilab for future mission needs. The request for such funds is made in consideration of the investment needs, affordability, and directives from DOE and Congress.

The FY 2014 Fermilab Annual LDRD Program Plan was approved and used to implement the FY 2014 LDRD Program. The laboratory sought and was granted approval for an LDRD expenditure comprising up to 0.4% ($1.5M) of the laboratory's total operating / capital budget.

Following an approved Annual LDRD Program Plan, a Call for Proposals was issued and 50 Preliminary Proposals were received. These Preliminary Proposals were typically one page in length and required supervisory and divisional concurrence as to the proposed scope of effort and materials requested by the Principal Investigator (PI). An LDRD Selection Committee reviewed these Preliminary Proposals and PI's were advised of their approximate competitive standing along with feedback should they wish to submit a Full Proposal.

In response to the Preliminary Proposal stage, a total of 29 Full Proposals were prepared, submitted, and fully evaluated by the LDRD Selection Committee. Each PI made a brief presentation to the Selection Committee. The LDRD Selection Committee evaluated the Full Proposals on a 5-point rating scale across 10 scoring criteria. The scoring criteria included an evaluation of the scientific or technical significance, innovativeness / novelty, the qualifications of the PI, the overall quality of the proposal, the likelihood of success, mission relevance, and relevance to the initiative as spelled out in the Call for Proposals, the strategic fit, enduring capability, and the likelihood to enhance the laboratory's reputation. In short, the scoring criteria encapsulated the key objectives and aspects that LDRD has as its purpose.

In consideration of the scope of proposals received, the LDRD Selection Committee made a recommendation to the Laboratory Director who approved the funding of the initial 7 LDRD projects at Fermilab. Table 1 shows the flowdown for proposals and funded projects for FY14 and FY15 (although FY15 projects are not the subject of this report).

Table 1: Number of LDRD Proposals received in response to the annual Call for Proposals, the number of Full Proposals prepared, and the number of awarded and funded LDRD Projects

Fiscal Year	Preliminary Proposals	Full Proposals	Funded LDRD Projects
FY14 (this report)	50	29	7
FY15	34	12	6

The seven FY14 LDRD projects, once approved, had internal laboratory project and task numbers assigned with budgetary information recorded such that financial tracking of effort and spending could be monitored. On roughly a monthly basis, project financial information was compiled, shared with the PI and the LDRD Coordinator. Each PI has also been asked to provide a short progress report to the LDRD Coordinator on approximately a monthly basis. In December 2014, the LDRD Selection Committee conducted a mid-year review of each of the seven projects and recommended project continuation that was approved by the Laboratory Director. Each PI provided the Project Summary contained within this document. Currently, all seven projects are being executed and no project has yet reached the stage of completion.

III. LDRD and Laboratory and Agency Mission

Department of Energy Mission

The mission of the Energy Department is to ensure America's security and prosperity by addressing its energy, environmental and nuclear challenges through transformative science and technology solutions.

Fermilab Mission

Our mission is to drive discovery in particle physics by

- *building and operating world-leading accelerator and detector facilities*
- *performing pioneering research with global partners*
- *transforming technologies for science and industry*

The Mission statement of Fermilab reflects the pursuit of excellence in scientific research in the area of particle physics. Particle physics addresses scientific mysteries in matter, energy, space and time through the Intensity, Energy, and Cosmic frontiers as enabled by theoretical, computational, accelerator, instrumentation and other technical capabilities. Fermilab's Mission statement reflects Fermilab's role in support of the overall mission of the Department of Energy. In particular, "transformative science and technology solutions" in the area of particle physics will be furthered through the use of LDRD. LDRD provides flexibility and efficiency that enables investigators to carry out creative new projects in forefront areas that enrich the current Fermilab program and strategically put Fermilab in a better position to deliver the mission objectives of DOE and Fermilab for the future. The Project Summaries describe the relevance of each project to the missions of Fermilab and DOE.

IV. Summary of Fermilab LDRD Costs

The costs associated with the Fermilab LDRD program are reported as part of the annual CFO database upload required at the beginning of each fiscal year. Costs associated with the administration of LDRD are absorbed into the Laboratory's overhead. Table 2 shows a list of approved projects and the FY14 spending for each project. Also shown is the spending for each project from the project inception through Jan 15, 2015 (the date most representative of the preparation of the Project Summaries).

Table 2: List of FY14 Fermilab LDRD Projects and the associated spending with each project. Shown is the actual spending in FY14 and uploaded to the CFO database. An updated number from Laboratory accounting shows spending for each project from the beginning of the project through Jan 15, 2015.

LDRD Project Number	Project Name	FY2014 Spending	Spending through 1/15/2015
FNAL-LDRD-2014-010	Cosmic Microwave Background Detector Development at Fermilab	$18,568	$90,278
FNAL-LDRD-2014-012	Development of HTS Based Rapid-Cycling Accelerator Magnets	$12,677	$39,669
FNAL-LDRD-2014-016	High Frequency Gallium Nitride Driver	$10,080	$40,111
FNAL-LDRD-2014-025	The Sinuous Target	$7,014	$15,319
FNAL-LDRD-2014-027	From Magic to Method: Characterizing High Voltage in Liquid Argon Time Projection Chambers with the Breakdown in liquid argon cryostat for high voltage experiments	$20,158	$104,066
FNAL-LDRD-2014-028	Deployment and operation of prototype CCD array at Reactor Site for detection of Coherent Neutrino-Nucleus Interaction	$70,518	$120,713
FNAL-LDRD-2014-038	Application-Oriented Network Traffic Analysis based on Graphical Processing Units	$48,693	$140,936
Totals		**$187,708**	**$551,093**

The small amount of spending at a level of $187,708 relative to the approved $1.5M in FY14 reflects the fact the projects were approved near the end of the fiscal year and many projects were just beginning by the end of FY14.

V. Project Summaries

Each of the seven approved and funded LDRD projects are described in a following Project Summaries. The proposal number and title, key authors including the PI are listed. A short project description is provided along with a statement addressing the relevance of the project to the Laboratory. There is a description of initial results and accomplishments of each project along with a list of publications (if any) that have been produced.

Proposal Number and Title: FNAL-LDRD-2014-010
Cosmic Microwave Background Detector Development at Fermilab

Authors: (PI) Bradford A. Benson (Fermilab); (Co-Is) John Carlstrom (University of Chicago), Clarence Chang (Argonne National Lab), Hogan Nguyen (Fermilab)

Project Description:
This project proposes to carry out R&D necessary to establish a high-throughput, high-quality characterization, packaging, and testing of prototype transition edge sensors (TES) and arrays of sensors at a scale never before performed for demonstration towards meeting the requirements of the next generation cosmic microwave background (CMB) experiments that will have greater than 10 times the scale and sensitivity compared with current experiments.

Relevance:
Fully realizing the scientific potential of CMB polarization will require new detector technologies that enables instruments with >10 times more sensitivity than current experiments. The technological challenges for future CMB experiments are well matched to, and solely filled by national labs like Fermilab. This proposal explicitly follows the P5 report's recommendation to: "***Increase particle physics funding of CMB research and projects in the context of continued multiagency partnerships***." The instrumentation developed will lead to CMB experiments that aim to answer some of the most exciting questions in cosmology and that are at the heart of the high-energy physics mission: to constrain inflationary physics at grand-unified theory energy scales (~ 10^{16} GeV), to measure the sum of the neutrino masses at a level below the minimum mass expected from neutrino oscillations (< 0.06 eV), and to precisely constrain the relativistic energy density of the universe and any "dark radiation" component.

Results and Accomplishments:
There have been five key accomplishments in the first 8 months of this proposal:

1) The first of two high-throughput detector test cryostats was received and successfully passed initial cryogenic tests. The design and development of cryogenic mechanical supports and wiring is ongoing, with full cryogenic validation expected by March 2015.
2) Lab A at Sidet has been equipped with the necessary infrastructure to run three helium compressors, for future cryogenic test-beds and camera integration. This

has included the purchase of a new water chiller, with the associated piping, and high-current (60 A) electrical wiring.

3) Several superconducting resonator prototypes, designed in collaboration with Argonne National Lab (ANL), have been characterized in an existing sub-Kelvin fridge at Sidet.
4) Fermilab scientists and technicians have begun to simulate and design prototype antenna-coupled TES detectors, work that has already fed-back into the design of TES detectors for the SPT-3G camera.
5) A process for the packaging and wire-bonding of prototype TES detector arrays has been developed, which has been recently used for the packaging of the first TES detector array fabricated by ANL.

Publications:

Benson, B. A., Ade, P. A. R., Ahmed, Z., et al. 2014, "SPT-3G: A Next-Generation Cosmic Microwave Background Polarization Experiment on the South Pole Telescope", Proceedings of the SPIE, 9153, 91531P

Project ID: FNAL-LDRD-2014-012
Project title: Development of HTS Based Rapid-Cycling Accelerator Magnets

Authors: (PI) Henryk Piekarz, James Blowers and Steven Hays

Project description:

The goal is design, fabricate, and test a short-sample rapid-cycling superconducting accelerator magnet using a combined novel core and power cable designs. The power cables are based on a high temperature superconductor (HTS) to expand the working temperature margin. These unique core and cable designs will minimize limitations with the normal conducting (copper cable) rapid-cycling magnets due to power loss induced by the eddy currents and magnetic hysteresis. They will potentially pave the way for extended capabilities in ramp rate, lower power, and long-term operation of such magnets. The rapid-cycling test magnet will be coupled to a 500 Hz ringing mode power supply and to a 20 Hz white circuit power supply to achieve 2000 T/s and 20 T/s ramp capabilities, respectively.

Relevance:

Rapid-cycling magnets constitute a critical component of many particle accelerator systems including fast-cycling synchrotrons such as Booster, Main Injector, and FAIR. The key limitations in achieving high beam intensities in these accelerators are due to power losses induced by a high ramp rate of the magnetic field. Successful R&D in this area would provide extended capabilities for possible new magnet and new synchrotron designs for Fermilab future accelerator complex and other facilities within DOE, as well as HEP facilities elsewhere, e.g. CERN in Europe, HEP in China.

Results and accomplishments:

The design of a reference window-frame magnet to determine the strength of the energizing current and to examine the B-field quality is completed. A novel design of a dual C-shaped magnetic core assembled in both horizontal and vertical orientation, is also completed. The engineering design of two types of HTS cables for a short-sample test-magnet: (1) multiple wide YBCO tape strands stacked vertically, and (2) multiple narrow YBCO strands wounded helically over a former pipe, are completed. The two designs differ in a way they minimize power losses due to strands self-field coupling and screen currents. However, due to rather high complexity of the HTS stack structures the cable application in a test magnet is required to determine with some certainty which cable type generates lower power losses. In both cable types the HTS conductor assemblies are placed inside the elliptical liquid helium channels which provide the strands cooling. The engineering designs of the splice connections between the strands of two connecting sub-cables, as well as the sub-cables cryogenic support system, are in progress.

Publications:
A Milanese [1], H. Piekarz [2] and L. Rossi [1], "Concept of a Hybrid (NC + SC) Bending Magnet for (80-100) km Hadron/Lepton Colliders, TUOCB01, Proceedings of IPAC-2014, Dresden, Germany ([1] CERN, [2] FNAL)

Proposal Number and Title: FNAL-LDRD-2014-038
HF GaN Driver

Authors: (PI) Gregory Warren Saewert

Project Description:
The purpose of this project is to develop an electronic switch employing Gallium nitride (GaN) field effect transistors (FET) having capabilities that are not available commercially. The capabilities to be achieved are: (1) switching transition times in the 1 to 2 nanosecond (ns) range; (2) repetition rates in the tens of megahertz, whether in bursts or continuous; (3) operating voltage in the 500 volt range; and (4) deliver direct current (DC) voltage, pulsed waveforms to a load. The switch needs to be constructed by connecting a number of identical GaN FET switch stages in a series connection for a number of reasons. One is to keep power dissipation of each GaN FET manageable, and the other is due to the limited voltage rating of currently available GaN FETs. Silicon based metal oxide field effect transistors (MOSFET) are typically used to build this type of switch but are incapable of achieving the above requirements for a couple of reasons. GaN FETs, however, are becoming commercially available that overcome the limitations of MOSFETs—both in speed and capacitance. Several challenges exist. One is to develop a suitable circuit using discrete parts to drive the GaN FET, because commercial FET driver ICs are too slow. Another is to develop an isolation gate triggering scheme having very high transient immunity, since the requirements for transient immunity and timing precision are beyond what is commercially available. Lastly, the voltage sharing ability of the GaN FET stages must be demonstrated.

Relevance:
There are a variety of applications in the accelerator community for fast transition, moderately high voltage, high duty factor switches. One significant interest is to kick out particle beam on a bunch-by-bunch basis. New linear accelerators being constructed at laboratories at Fermilab and elsewhere are looking to selectively kick out individual bunches of beams that are spaced 6.2 ns apart. To do this, irregular waveforms having high duty factors are required on the part of the driver to deliver voltage pulses to the beam-kicking structure. A DC coupled switch based driver has, in principle, some advantages over the only known competing scheme, namely, an AC coupled, very wide band, high voltage amplifier driver. This expensive scheme has not been demonstrated either because of its own set of challenges including potential baseline settling issues, which a DC coupled switch driver does not have. Other possible applications for very fast switches include those for driving Pockel Cell laser Q switches or electrodes of electron devices. Other techniques and technologies employing multiple transistors to deliver fast transitioning voltage pulses in the 500 Volt range have been used. These include using a Marx generator topology, voltage summing techniques or proprietary FID pulsers. However, all these are suitable for only low duty factor applications. On the other hand, GaN FETs are ten times faster than silicon MOSFETs. GaN FET technology enables building switches for accelerator applications that were not possible before. Side benefits of this LDRD effort useful to pulsed power systems for accelerators will be to develop: (1) a very fast and efficient drive technique for GaN FETs to utilize their speed, and (2) a very fast triggering scheme having high transient immunity.

Results and Accomplishments:
The project has progressed with a number of things: (1) identified available GaN FETs; (2) developed a driver circuit design using vendor-provided SPICE models; (3) begun the printed circuit board (PCB) layout of the driver circuit; (4) identified the isolation, gate triggering scheme to implement. Because the load current requirement is not high, 600V "low" current rated GaN FETs are not yet available from vendors. The need is for low capacitance that only accompanies lower current rated parts. Therefore, the plan is to use 200V GaN FETs from a commercial vendor until higher voltage parts become available. The vendor of the devices provides models of all their GaN and DMOS FETs that have been used to pursue driver development using SPICE. The PI has arrived at a driver design optimized for the following features: (1) adequate GaN FET drive to assure the 2 ns maximum transition time; (2) provision to adjust and fix both the turn-on and turn-off rates such that all are matched; (3) provision to adjust and fix the delay such that all stages transition simultaneously; (4) inherent insensitivity to temperature affects on the transition timing; (5) a GaN driver stage that is reasonably efficient; and (6) an interface stage that anticipates an eventual isolated triggering scheme. SPICE modeling proved highly beneficial for the time spent. The PI has chosen a design and is in the process of designing its printed circuit board (PCB) layout. This prototype is isolated from ground and will allow me to test even a multi-FET switch – should it work well enough. Lastly, very stable voltages are required on each stage; and power needs to be delivered by a means having very low common mode impedance (this is again an isolation issue). The PI has incorporated a scheme to do this in this prototype layout.

Proposal Number and Title: FNAL-LDRD-2014-025
The Sinuous Target

Authors: (PI) Robert Zwaska

Project Description:
The project is to generate a new, engineered material for use in high-power accelerator targets. The sinuous material will be composed of a multitude of interlaced rills: wires of small dimensions. The sinuous bulk material will have improved resistance to thermal shock due to the very low effective modulus of elasticity. Furthermore, the interlaced nature of the wires makes it resilient to individual wire failures. The material will enable targets to accommodate higher incident beam power with more efficient secondary beam production. This project will develop production techniques of several matrixing approaches to the material, and test the mechanical and thermal properties of the engineered bulk.

Relevance:
A high-power target is an integral part of a neutrino beam, muon beam, other intensity frontier beams for high energy physics as well as neutron and rare isotope beams outside of high energy physics. Thermal, mechanical, and radiation effects limit the degree to which targets can be subject to high incident beam power. If successful, the material developed has applications at accelerators at Fermilab and within the DOE complex (including beam facilities for NP, BES, FES, and other areas). The most proximal application of this technology will be for the production of multi-megawatt neutrino beams at Fermilab.

Results and Accomplishments:
A number of analogue, manufactured materials have been identified to act as reference materials: metal foams, flexible graphite, and reticulated metal foams and glassy carbon. A number of prototypes are being procured and will be subjected to mechanical experiments and inspection. A suite of tests is being designed for these materials.

Publications:
The Sinuous Target IPAC 2015, May, 2015. Accepted for contribution.

Proposal Number and Title: FNAL-LDRD-2014-027
From Magic to Method: Characterizing High Voltage in Liquid Argon TPCs with Breakdown in liquid argon cryostat for high voltage experiments

Authors: (PI) Sarah E. Lockwitz, Brian Rebel, Hans Jostlein

Project Description:
Liquid argon is a popular detector medium for neutrino and dark matter experiments due to its ionization detection ability and its relative affordability compared to other noble liquids. Producing a strong signal in such experiments requires a high electric field, yet

achieving and maintaining voltages for such fields continues to challenge experimentalists. The goal of this project is to better understand the conditions related to dielectric breakdown in liquid argon by focusing on two key areas: the dielectric strength of liquid argon, and the performance of insulating materials under electrical stress in liquid argon. The experiments are being carried out in two stages: first, in an open-to-air cryostat where preliminary tests are being performed, and then in a closed cryostat where refined testing can take place in ultra-pure argon similar to the conditions of a liquid argon experiment. A better understanding of these parameters would lead to a more efficient design process for high voltage components in liquid argon experiments in effect taking the approach from magic to method.

Relevance:
This project is relevant to the DOE's Office of Science in that it helps to support the Fermi National Accelerator Laboratory's (FNAL) mission to build and operate world-leading detector facilities and perform pioneering research. Successfully defining the design and operation parameters leading to stable high voltage operation in liquid argon would enable FNAL to efficiently design and use robust, cutting-edge liquid argon detectors in particle physics experiments with minimal downtime, and reduced risk of damaging components due to high voltage failures. In the near term, LAr1ND, a liquid argon time-projection chamber planned for one of FNAL's neutrino beams, is being designed and could immediately benefit from the results of this project. In the future, the Long-Baseline Neutrino Facility (LBNF) would also directly benefit from the research and could further benefit in the design process by understanding the limits of high voltage stability. More generally, liquid argon detectors have the potential to unlock the mysteries related to the ubiquitous yet elusive neutrinos and dark matter thereby increasing our understanding of the fundamental components of the universe.

Results and Accomplishments:
Since July 2014, we have procured a cryostat and a condenser for our future high voltage testing in pure argon. We have been working with an engineering team on the new cryostat and condenser assembly to ensure a safe apparatus, and their first drafts of the engineering design documents have been prepared. In January 2015, we will begin installation with the help of technicians to expedite the process. During the Summer of 2014, we performed the first round of preliminary insulator tests in an open-to-air cryostat. Here, the basic mechanics of the test were evaluated; we learned how to make test pieces of insulators, how to vary the length, and how to ensure dielectric breakdown over the piece under study. This information was then used for a second round of insulator testing in the open cryostat near the end of 2014. For these tests, we increased the number of breakdowns, tested longer lengths, and evaluated additional samples bringing the total number of tested materials to 11. We found that some materials suffer mechanical failure under thermal stress when introduced to liquid argon, while others fail only after numerous exposures to high electrical stress. We have identified a subset of the test materials warranting further evaluation in the closed cryostat in 2015.

Proposal Number and Title: FNAL-LDRD-2014-028
Deployment and operation of a prototype CCD array at Reactor Site for detection of Coherent Neutrino-Nucleus Interaction

Authors: (PI) Juan Estrada, Gustavo Cancelo, Javier Tiffenberg

Project Description:
Neutrinos are one type of the elementary particles that make up our universe. These weakly interacting subatomic particles have no electric charge. The detection of neutrinos is an experimental challenge because of their low interaction probability. The state of the art to address this challenge is based on very massive particle detectors with the capability to identify signals from secondary particles generated in a neutrino interaction. This project has the goal of developing a new way of detecting neutrinos, based on the direct observation of their elastic collisions with nuclei. The very small signal produced by these low energy nuclear recoils is typically below the detection threshold. For this project we establish a much lower detection threshold by using extremely sensitive charge couple devices (CCDs). We use a 4 gigawatt nuclear power plant as a source of neutrinos to test this detector concept. The cosmic and environmental radiation is suppressed using a polyethylene and lead radiation shield. The array of CCDs is located 30 meters away from the core of the reactor, inside the radiation shield.

Relevance:
The elastic collisions of neutrinos and nuclei is theoretically well established, but has escaped detection until now. At low energies the probability of this interaction is expected to increase due to the large number of elementary particles inside a nucleus (coherent interaction). The successful detection of this process will validate for the first time this important aspect of our understanding of elementary particles. This neutrino detection technique could also open a new window to physics beyond our standard model of particles and their interactions, looking at processes that are enhanced at low interaction energies in the neutrino sector. A compact detector capable of identifying neutrinos produced in a reactor, could also be used as a tool for monitoring the operations in nuclear power plants.

Results and Accomplishments:
The experimental setup was assembled at Fermi National Acceleration Laboratory (Batavia, IL) for testing in July 2014. The equipment was then shipped to the nuclear power plant in Angras do Reis, Brazil in October 2014. The detector has been collecting data with a partially assembled radiation shield since November 2014. Stable operation, with low threshold has been demonstrated at the reactor site. We are currently finalizing the radiation shield assembly, evaluating the environmental radiation background at the site, and establishing a baseline for the detection of neutrinos generated at the reactor.

Publications:
"Charge Coupled Devices for the Detection of Coherent Neutrino-Nucleus Scattering", G.F. Moroni, Juan Estrada, Gustavo Cancelo, Eduardo Paolini, Javier Tiffenberg, Jorge Molina. Submitted to Physical Review D, arXiv:1405.5762 (2014).

Proposal Number and Title: FNAL-LDRD-2014-038
Application-Oriented Network Traffic Analysis Based on GPUs

Authors: Philip DeMar (PI); Dr. Wenji Wu (Co-Investigator)

Project Description:
The project seeks to develop a network traffic analysis platform that provides a generalized framework for real time traffic pattern identification. The platform is intended to be capable of supporting these services within advanced, high-performance networks that utilize 40GE and 100GE network technologies. Existing packet capture and analysis tools do not perform adequately in such advanced network environments. The project platform will consist of a lossless packet capture engine, and a set of basic GPU libraries for manipulating and analyzing the captured data packets with GPUs. The development use case will be the identification of high impact, large-scale science data flows. A customized set of GPU libraries will be developed to perform the specific traffic characterization services that detect and identify high-impact data flows of interest. While this project is specifically targeted at identification of high-impact data flows, the platform design is intended to be adaptable a wide range of traffic detection and identification capabilities. In the final stage of the project, an interface for the platform to an OpenFlow network controller will be developed to demonstrate proof-of-concept reconfiguration of network infrastructure to accommodate specifically identified traffic.

Relevance:
The Department of Energy's (DOE) Office of High Energy Physics (HEP) has heavily invested in US participation in the Large Hadron Collider (LHC) experiments. The LHC experiments have distributed computing models of global scale, and move extreme volumes of data among collaborating sites, including DOE-funded National Laboratories and research university HEP programs. This project is targeted at helping facilitate use of emerging network technologies, such as Software-Defined Networking (SDN), that will be needed to effectively support this extreme data movement. The focus of our project is on identification of high-impact traffic flows for network traffic management purposes. We target those traffic management purposes to be: (1) provision of adequate network resources for the specified flows; (2) isolation of the specified data flows from general Internet traffic; and (3) facilitation of custom computer security risk profiles for the specified traffic. We envision interfacing our product to SDN controllers, in order to enable efficient use of high performance network infrastructure, such as 40GE and 100GE, for LHC data movement.

Results and Accomplishments:
The project has four components: (1) a high-performance packet capture engine; (2) a set of general GPU libraries that manipulate the captured packets for GPU analysis; (3) a set of custom GPU libraries that perform the traffic analysis and classification; and (4) an output engine that interfaces to SDN controllers for manipulating network services based on identified traffic. The initial project plan specified a 24-month development road map, with items (1) and (2) completed by the end of year one, with items (3) and (4)

developed in year two. The project is currently six months into its development plan. Item (1) has been completed with 10GE network interface cards (NICs) and is being expanded to include 40GE NICs. Preliminary design work on item (2) is currently under way, with an expectation that development will be completed by the end of year one. A stand-alone software package based on item (1) has been presented at the ACM/IMC 2014 conference (see publication listing below). The stand-alone software package is named WireCap. A number of researchers and commercial entities have expressed interest in the WireCap software, and the technology that it uses.

Publications:
ACM/IMC 2014; Vancouver, British Columbia: "WireCAP: a Novel Packet Capture Engine for Commodity NICs in High-speed Networks" - Dr. Wenji Wu and Philip DeMar (Fermilab)

VI. Conclusions

Fermilab is successfully conducting a LDRD program in accordance with the terms of its authorization. The initial seven approved projects are addressing R&D in several areas of scientific and technical expertise that exist at the Laboratory. All projects are aligned with the mission of the Laboratory and DOE and have begun to make progress in the typical six month period that this report reflects. Already, there have been several publications that are related to the LDRD work along with initial discussions of intellectual property that have been generated. The number of proposals already submitted and subsequent discussion with those who have submitted the proposals also indicate that the LDRD program is strengthening the scientific and technical vitality of Fermilab.