SLAC-TN-15-060

Detecting Energy Modulation in a Dielectric Laser Accelerator

Louis Lukaczyk
EE, Physics
University of Virginia
Class of 2016

SLAC National Accelerator Laboratory
Science Undergraduate Laboratory Internship

August 21st, 2015

Written in partial fulfilment of the requirements for the Department of Energy deliverables for the Science Undergraduate Laboratory Internship program at SLAC National Accelerator Laboratory

SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025
This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program, under Contract No. DE-AC02-76SF00515.

Abstract

The Dielectric Laser Acceleration group at SLAC uses micro-fabricated dielectric grating structures and conventional infrared lasers to accelerator electrons. These structures have been estimated to produce an accelerating gradient up to 2 orders of magnitude greater than that produced by conventional RF accelerators. The success of the experiment depends on both the laser damage threshold of the structure and the timing overlap of femtosecond duration laser pulses with the electron bunch. In recent dielectric laser acceleration experiments, the laser pulse was shorter both temporally and spatially than the electron bunch. As a result, the laser is theorized to have interacted with only a small portion of the electron bunch. The detection of this phenomenon, referred to as partial population modulation, required a new approach to the data analysis of the electron energy spectra. A fitting function was designed to separate the accelerated electron population from the un-accelerated electron population. The approach was unsuccessful in detecting acceleration in the partial population modulation data. However, the fitting functions provide an excellent figure of merit for previous data known to contain signatures of acceleration.

Acknowledgements

Thanks to my parents for supporting me through college. Thanks to Enrique Cuellar for being a great program director for SULI. Thanks to my mentor Joel England for giving me the opportunity to work on a fascinating project. Thanks to Kent Wootton, Ziran Wu, Igor Makasyuk, and Adi Hanuka for giving me advice on my research project and showing me how to be a good experimentalist. Thanks to the Department of energy, Stanford University, and SLAC National Accelerator Laboratory for making this summer research experience possible.

Outline

1. Introduction to the Dielectric Laser Accelerator
2. Experimental Setup for Dielectric Laser Accelerator at NLCTA
3. Data Analysis
4. Conclusion

Introduction to the Dielectric Laser Accelerator (DLA)

Dielectric laser acceleration is a novel concept which could make particle accelerators more compact and economical. Conventional accelerators based on radio frequency technology have reached a limit for the maximum accelerating gradient they can provide even with superconducting cavities. By using micro-fabricated dielectric structures and conventional high power lasers, dielectric laser acceleration is estimated to produce an accelerating gradient up to 2 orders of magnitude greater than gradients produced by conventional RF accelerators.

Particle accelerators have seen widespread use in modern science and technology. Important tools powered by accelerators such as the X-ray free electron laser are enabling groundbreaking research in physics, chemistry, and biology. Accelerators are also used in the medical physics industry as radiation sources for cancer therapy. The widespread use of particle accelerators is thwarted by the size, complexity, and cost of conventional radio-frequency structures. The development of the dielectric laser accelerator could herald a new age of scientific research enabled by compact accelerator technology.

Previously, the Dielectric Laser Acceleration group at SLAC observed high gradient (beyond 250 MeV/m) acceleration of relativistic electrons using a micro-fabricated diffraction grating structure [2]. The structure was powered by 800nm infrared light from a mode-locked Ti:Sapphire laser. The pictures below show the wafer-scale micro-fabricated structures.


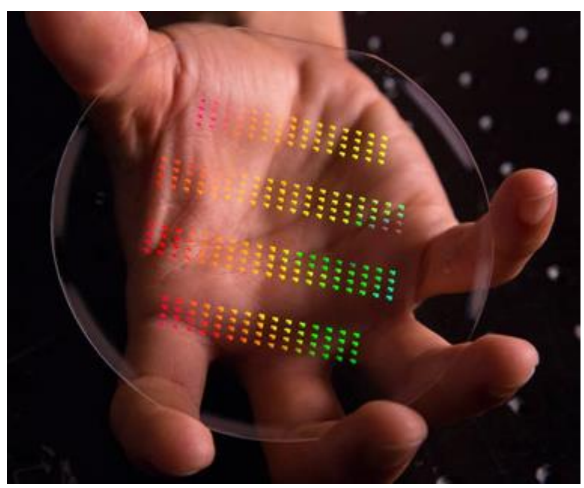

Fig 1a – A single micro-fabricated grating structure diced from a fused silica wafer (Left);
Fig 1b - A fused silica wafer containing multiple grating structures (Right) [1]

The diagram below shows the basic setup of the experiment.

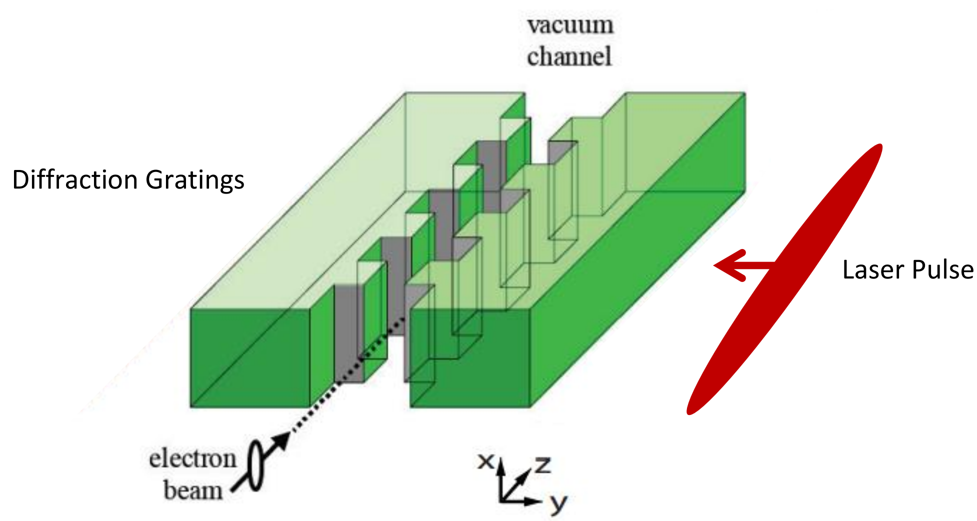

Fig 2 – Basic Setup for Accelerating Electrons with the DLA [1]

The structures consist of a sandwich of gratings etched in fused silica using conventional lithographic techniques. Electrons travel through a small vacuum channel between the two diffraction gratings. The purpose of the grating structure is to create longitudinal electric field modes which are phase matched to the velocity of the electrons [1]. The electrons are phased matched to the structure because the periodicity of the gratings is the same as the laser wavelength (800 nm).

The principle of operation of the structure can be easily understood by considering that the electric field is intensified in the pillar regions and diminished in the gap regions. This is because the electric field is stronger near the dielectric. The electrons perfectly matched to the laser phase experience a strong accelerating force in the pillar region, and a weak decelerating force in the gap region. The result is net energy gain after a single grating period. Current structures are a millimeter long and consist of hundreds of grating periods.

The structure is an example of a photonic crystal. Other photonic crystal structures, such as a Bragg fiber or a woodpile structure, have been proposed for use as dielectric laser accelerators [2]. These designs are difficult to fabricate since they require challenging dimensional tolerances. The diffraction grating structure circumvents this problem due to its simple planar structure. The structures used by the Dielectric Laser Acceleration group were produced at Stanford Nanofabrication Facility. The diagram below shows the fabrication process for the grating structure.

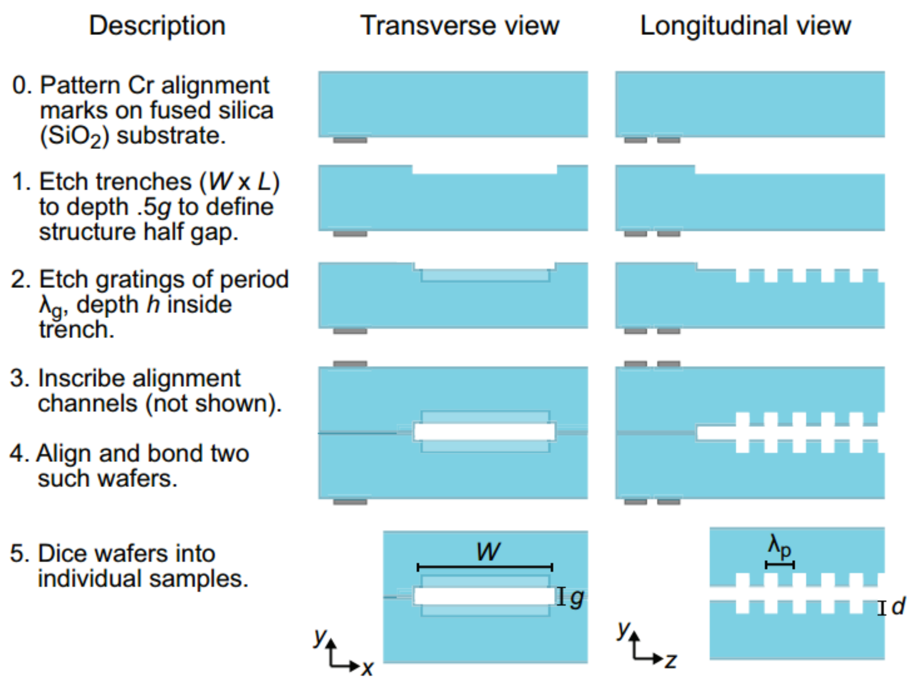

Fig 3 – Wafer-scale Bonding Fabrication Method [1]

The trenches that form the diffraction grating are etched using reactive ion etching. This enables control over the depth and quality of the etching. It is important to note that the fabrication process requires bonding two separate fused silica wafers together such that the gratings are precisely aligned on the top and bottom. The gap between the top and botton gratings is less than a micron.

Experimental Setup for the Dielectric Laser Accelerator at the NLCTA (Next Linear Collider Test Accelerator)

The Dielectric Laser Acceleration group uses relativistic electrons (60 MeV) from the first linac section of the Next Linear Collider Test Accelerator beamline at SLAC. This beamline is a conventional accelerator and uses metal resonant cavities pumped with X-band radiofrequency (12 GHz) electromagnetic waves from a klystron. A diagram of the facility is shown below.

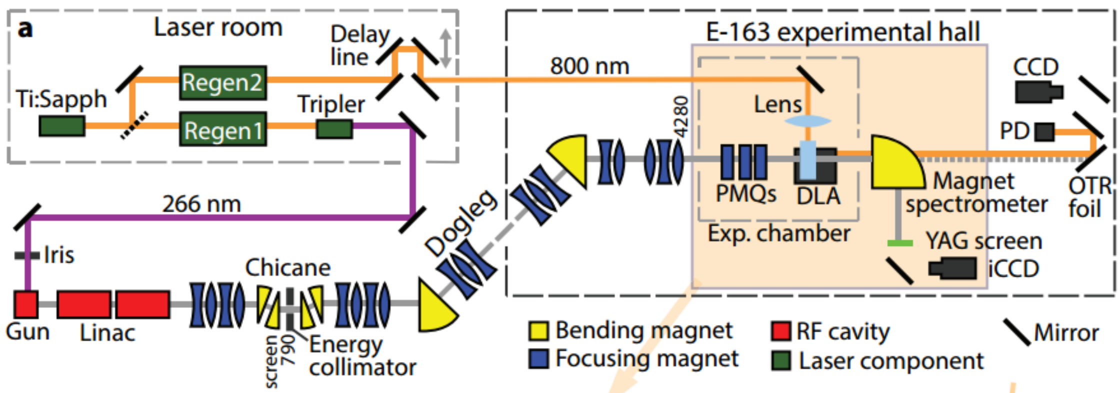

Fig 4 – Diagram of the NLCTA Facility, Including the E-163 Experimental Hall [1]

The process for experimenting with the dielectric laser accelerator begins in the laser room. Here, an 800nm master laser pulse from a Ti:Sapphire laser is generated and split into two laser beamlines. One beamline is used to generate electrons at the electron gun via the photoelectric effect. The laser pulse is amplified in a regenerative amplifier and passed through a Tripler to convert it to ultraviolet wavelength. The ultraviolet pulse is then directed onto the cathode of the electron gun. This creates a high quality beam which can be precisely timed with the laser pulse which drives the dielectric accelerator structure.

The timing of the drive laser pulse with the electron bunch is adjusted via a delay line. This consists of a translatable series of mirrors which vary the distance traveled by the laser pulse. One section of the delay line is for rough adjustment and is moved manually. This adjustment gets the timing correct to within a few picoseconds. The precise timing is performed using an automated voice coil delay stage. This is important because it allows the timing of the laser pulse to be scanned during the experimental run over a timespan of around 20 picoseconds with femtosecond resolution. The voice coil is scanned in a pseudo random fashion to eliminate systematic error. The laser pulse is turned on and off during the scan as a control. This is achieved by disabling the Pockels cell inside of the regen which prevents the laser pulse from exiting.

Electrons from the photocathode are accelerated to roughly 60 MeV after the first linac section. The electrons pass through focusing optics and a chicane which separates the electrons spatially by energy. Collimator jaws are then used to pass a certain energy band. The electron bunch is then sent through a dogleg into the experimental hall. The beam is focused by a series of electromagnet and permanent magnet quadrupoles before being sent through the dielectric laser accelerator. After traversing the dielectric laser accelerator, the electron energy distributions are analyzed using a spectrometer magnet. The electrons hit a YAG screen which is then imaged by a camera.

Much of the problems faced during an experimental run of the dielectric laser accelerator stems from poor electron transmission through the vacuum channel of the structure. The gap through which the electrons must traverse is on the order of the laser wavelength, 800nm. A large population of the electron bunch hits the structure and loses some energy. The diagram below depicts a typical single shot electron energy distribution.

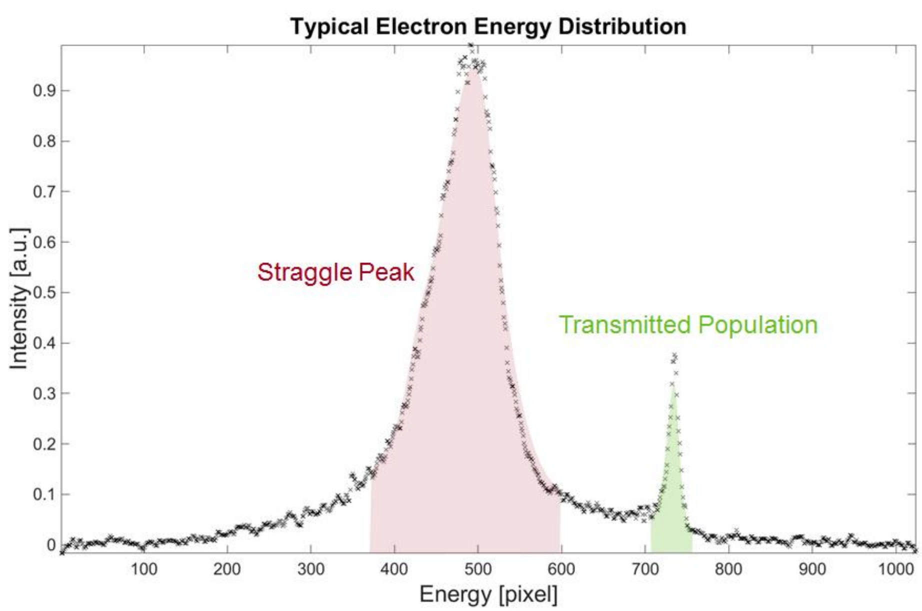

Fig 5 - Typical Single Shot Electron Energy Distribution

The electrons which hit the structure and lose energy show up on the energy distribution as a large straggle peak. The transmitted population shows up as a smaller and sharper peak to the right of the straggle peak.

Transmission through the structure is complicated by electron beam jitter. Jitter refers to instability in various beam parameters such as directionality, energy, or spot size. Jitter is caused by transient fluctuations in various beamline components. The most significant fluctuation is phase and amplitude jitter in the RF klystrons which power the NLCTA linac. Jitter can be caused by temperature fluctuations due to weather. Lastly, jitter can be caused by unstable current supplies for the optical components like quadrupole or bending electromagnets.

To minimize the effect of RF jitter, the phase of the linac RF with respect to the arrival of the electron gun laser pulse is adjusted to produce a chirped beam. This produces a long tail in the electron energy distribution. Electrons on the less jittery tail of the chirped energy distribution are then selectively passed by the collimator jaws in the chicane.

Data Analysis

In experiment, individual shots are filtered to eliminate extraneous data. The center of the transmitted population peak is one of the filter criteria. If the beam has significant energy jitter, the transmitted peak location will not be stable. Another filter criteria is the height of the transmitted peak. The height of the transmitted population depends on good transmission through the structure. Poor alignment of the electron bunch with the vacuum gap makes detection of the acceleration signature difficult because a poor transmission event might be mistaken for acceleration.

A consequence of powering an accelerator with optical frequency laser light is that the time and length scales required for correct phase matching of the electrons with the laser pulse are reduced to the order of an optical oscillation period (~1fs). The electron bunch produced by the NLCTA is on the order of 100 fs. This means that the electrons sample both the accelerating and decelerating phase of the laser pulse, creating an energy spectrum referred to as a double humped distribution. The plot below shows a typical example of a double humped distribution.

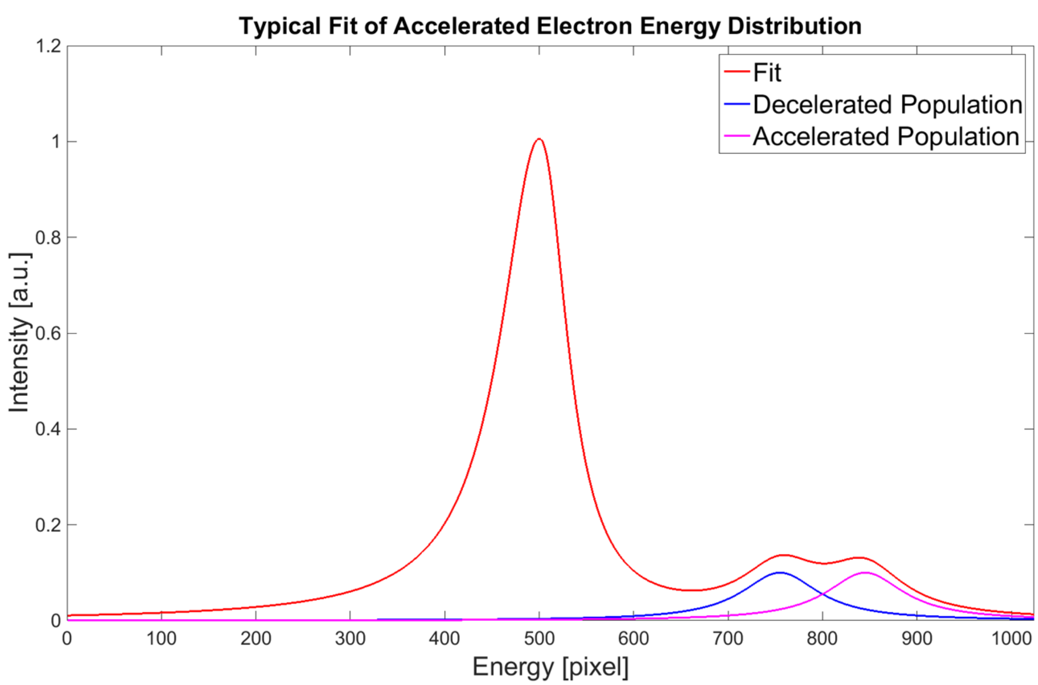

Fig 6 – Typical Electron Double Hump Distribution

Here the red trace is the overall energy distribution. The blue peak represents the decelerated electron bunch and the pink peak represents the accelerated electron bunch. Each peak is modeled by a Lorentzian distribution. During an experimental run, each acceleration shot is fitted with functions representing the expected distribution. Figures of merit based on the fit parameters for each shot are plotted for an entire run of events. A run of events consists of hundreds of energy spectra shots with similar experimental parameters. The following plot shows a figure of merit plot for a successful run of events.

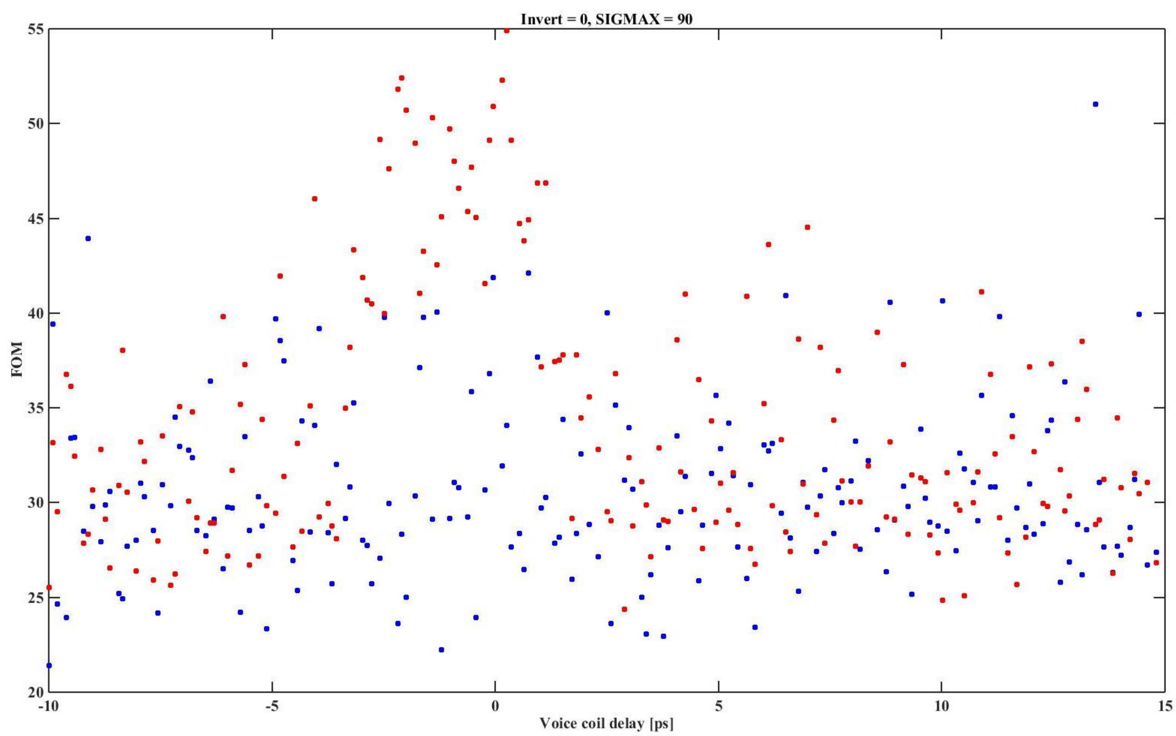

Fig 7 – Example of a Figure of Merit Plot for Successful Data using Full Modulation Fits

The plot shows laser on events as red dots and laser off events as blue dots. The y-axis is the figure of merit (FOM) for each event. The x-axis is the voice coil delay in picoseconds. Here, the figure of merit is the integral of the transmitted population over the height of the transmitted peak. This figure of merit represents how spread out the transmitted population is and has been used for most of the previous experiments because it doesn't depend on the specific form of the fitting function. A peak in the figure of merit plot for laser on events over the laser off events localized in a timespan of a few picoseconds indicates that acceleration was detected.

In the data plotted above, the electron bunch was fully covered by the laser pulse. In recent experiments, the laser pulse was shorted both spatially and temporally than the electron bunch. This makes detecting acceleration more difficult, since less of the electrons are accelerated. A solution to this problem is to change how the data is fitted. A ‘partial population modulation’ fitting function was implemented to try to detect the small amount of electrons that are accelerated. The diagram below shows how a short laser pulse can cause partial population modulation.

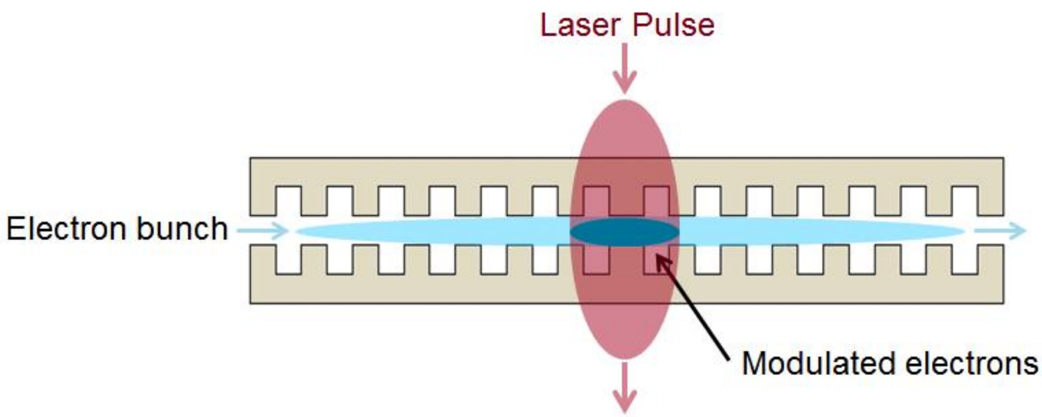

Fig 8 – Diagram of Partial Laser-electron Interaction

The partial population modulation fit consists of three transmitted peaks. One peak represents electrons transmitted without acceleration. The two remaining peaks represent the accelerated and decelerated electrons. The plot below shows an example of a partial population modulation fit.

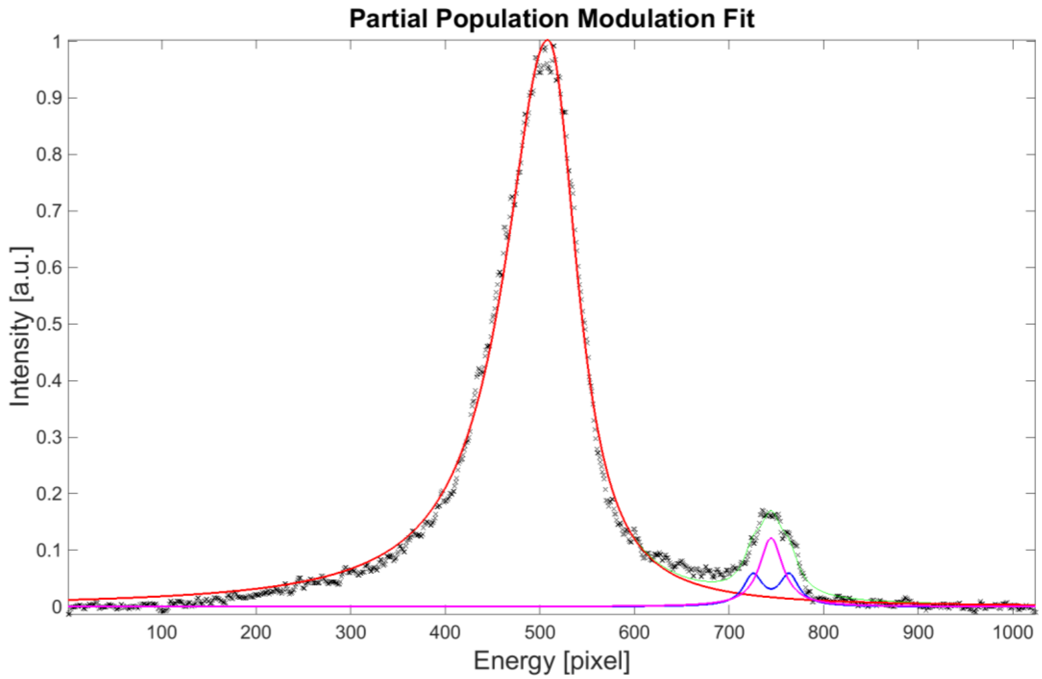

Fig 9 - Example of a partial population modulation fit

In the plot above, the blue trace represents the sum of the accelerated and decelerated peaks. The pink peak is the unmodulated peak. This fitting function is useful because the separation between unmodulated and modulated peaks is explicitly fit as a parameter. This may help quantify the energy gain of accelerated electrons. The other fit parameters include the height and width of the peaks. By adjusting the range of acceptable values for the parameters, a strong signal was found in the previous successful data. The population of accelerated electrons can also be explicitly integrated.

The plot below shows an example of a strong signal found using the partial population modulation fitting routine. The figure of merit used is the integral of the accelerated population.

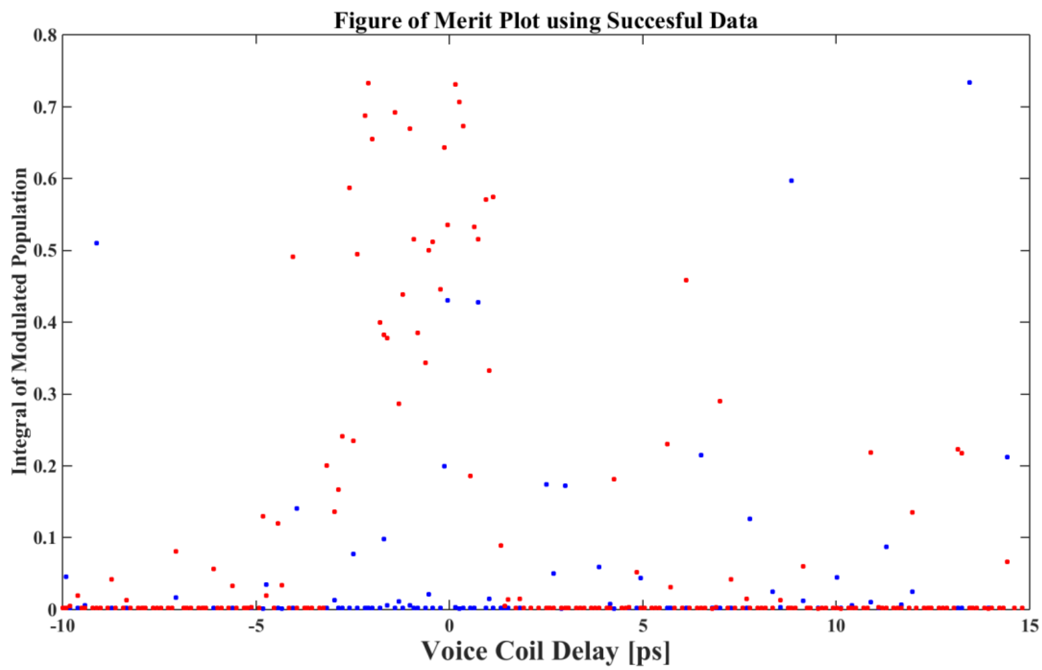

Fig 10 – Example of a Figure of Merit Plot on Successful Data Using Partial Population Modulation Fit

With the new fitting function, most events with no sign of acceleration have a null figure of merit. The peak of red dots above the blue dots clearly shows that acceleration occurred. The new partial population modulation fit seems to extract the signal very well from the previously successful data. However, it doesn't agree with the parameters of that run. The laser pulse parameters were such that the entire electron bunch should have been modulated. Thus, it is wrong to assume a partial population modulation distribution for that data.

The plot below shows an example of a typical figure of merit plot for the experimental runs expected to contain partial population modulation.

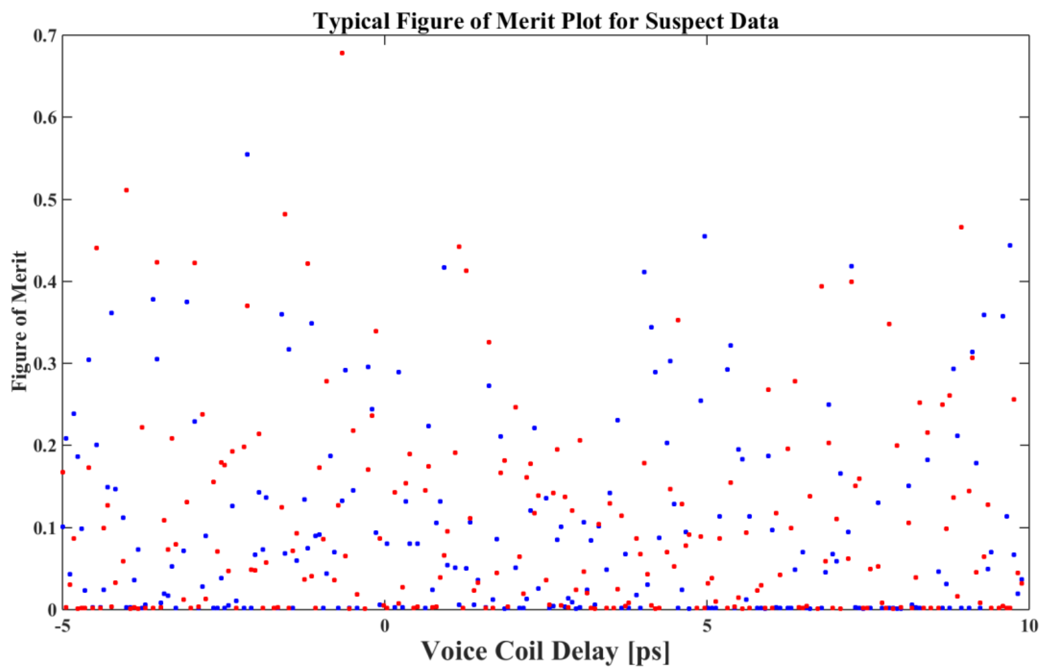

Fig 11 – Figure of Merit Plot of Partial Population Modulation Data

No signal is observable in the figure of merit plots for the partial population data. While the new fit is very sensitive to signs of acceleration, the fit parameter bounds require significant fine tuning to eliminate noise from the acceleration signal. This, combined with a small population of accelerated electrons, means that the fit was unsuccessful in detecting acceleration for that data.

Conclusion

The success of the dielectric laser accelerator (DLA) depends heavily on the interaction between the laser pulse and the electron bunch. Several factors can effect this. In recent experiments, detection of electron acceleration was jeopardized by inconsistent transmission of the electron bunch through the vacuum gap of the grating structure. Unfortunately, the partial population modulation fit was unsuccessful in detecting signal in the experiments with partial laser interaction. The fitting function may provide an effective means of detecting signal in future experimental runs. However, the fit is not suited for the final data analysis since it makes flawed assumptions about the amount of laser-electron interaction.

Bibliography

1. Peralta, Edgar A. *ACCELERATOR ON A CHIP: DESIGN, FABRICATION, AND DEMONSTRATION OF GRATING-BASED DIELECTRIC MICROSTRUCTURES FOR LASER-DRIVEN ACCELERATION OF ELECTRONS*. Diss. STANFORD UNIVERSITY, 2015. N.p.: n.p., n.d. Print.
2. Peralta, E. A., K. Soong, R. J. England, E. R. Colby, Z. Wu, B. Montazeri, C. Mcguinness, J. Mcneur, K. J. Leedle, D. Walz, E. B. Sozer, B. Cowan, B. Schwartz, G. Travish, and R. L. Byer. "Demonstration of Electron Acceleration in a Laser-driven Dielectric Microstructure." Nature 503.7474 (2013): 91-94. Web.