

Simulations of Ion Beams for NDCX-II[☆]

D. P. Grote^{a,*}

LBNL MS47-112, 1 Cyclotron Rd, Berkeley, CA 94720, 1-510-495-2961

A. Friedman^a, W. M. Sharp^a

^aLawrence Livermore National Lab, Livermore, CA 94550

Abstract

NDCX-II, the second neutralized drift compression experiment, is a moderate energy, high current accelerator designed to drive targets for warm dense matter and IFE-relevant energy coupling studies, and to serve as a testbed for high current accelerator physics. As part of the design process, studies were carried out to assess the sensitivities of the accelerator to errors, and to further optimize the design in concert with the evolving pulsed power engineering. The Warp code was used to carry out detailed simulations in both axisymmetric and full 3-D geometry. Ensembles of simulations were carried out to characterize the effects of errors, such as timing jitter and noise on the accelerator waveforms, noise on the source waveform, and solenoid and source offsets. In some cases, the ensemble studies resulted in better designs, revealing operating points with improved performance and showing possible means for further improvement. These studies also revealed a new non-paraxial effect of the final focus solenoid on the beam, which must be taken into account in designing an optimal final focusing system.

Keywords: particle accelerator, space-charge, plasma physics

1. Introduction

Ion beams are an attractive approach for heating targets to study warm dense matter[1] and high energy density physics, and ultimately for driving targets for fusion energy[2]. The second neutralized drift compression experiment, NDCX-II, is a moderate energy, high current accelerator that is designed to heat thin-foil targets and enable beam dynamics studies applicable to a driver for heavy ion driven inertial fusion energy. The facility accelerates and compresses an ion beam, providing a high current short pulse that can be focused down to a small spot on the target. It uses induction acceleration, with carefully designed waveforms that accelerate and shape the beam as it propagates. A head-to-tail velocity gradient (“tilt”) is applied to the beam so that it compresses as it approaches the target, increasing the power. A plasma is used to neutralize the space-charge of the beam as it compresses, allowing a higher degree of current amplification and removing the need for transverse confinement during compression. The initial configuration uses 12 induction cells to accelerate a Li⁺ beam up to 1.2 MeV over a 10 m long accelerator.

The beam is confined transversely by 27 solenoids, whose strength varies between 1 and 2 T to maintain a radius of approximately 2 cm. The beam is compressed from an initial 500 ns duration down to less than 1 ns duration on the target, putting 30 nC of charge onto the target with a peak current of typically around 36 A. An 8 T final focus solenoid focuses the beam down to less than a 1 mm RMS spot size. See [3] for more details on how the required beam manipulations are accomplished.

An important part of the design process was characterization of errors and how they affect the beam. Tolerances were required, based on an assessment of the levels of errors acceptable without significant degradation of the beam. A variety of errors were examined, including static errors, such as hardware misalignment and field offsets and non-uniformities, and dynamic errors such as field strength, timing jitter, and waveform noise. This is not an exhaustive list of errors, but covers what were expected to be the largest errors, and those with the most significant effect on the beam.

Designing NDCX-II was a complex process, with a large number of variables - it was not expected that the design, as currently laid out, would necessarily be optimal. Because of this, the error characterization process could be used as a sort of additional optimization, to find improved cases that result from the presence of “errors”. While the errors were being characterized, the hardware layout of the machine was held fixed - this included the source geometry and voltage waveforms, the number and arrangement

[☆]This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

*Corresponding author

Email addresses: dpgrote@lbl.gov (D. P. Grote),
afriedman@lbl.gov (A. Friedman), afriedman@lbl.gov
(W. M. Sharp)

of induction cells, and the accelerating waveform shapes (though they could be varied in amplitude and timing). It was found that some variation from the nominal of the dynamic quantities could lead to improved performance. For example, shifts in the timing of multiple induction waveforms could compound to produce a beam with a more linear velocity tilt, or errors in the solenoid strengths could modify the beam radius such that the changed longitudinal self-fields improved the beam behavior.

To help in characterizing the errors, it is useful to have a figure of merit. The main goal of NDCX-II is driving targets, so some measure of how well the beam performs that task was needed. There are several beam measures that are useful, such as the FWHM (full width half maximum) of the beam and the peak fluence on target, but these by themselves are insufficient. We define the peak fluence as the energy deposited within a 0.1 mm diameter hot spot, integrated over the beam duration. A diameter of 0.1 mm is the expected smallest diagnosable area on the target. One possible measure is the peak power density on target. A figure of merit was chosen that directly reflects how strongly the beam drives the target. Using a range of typical beam parameters as input, example target simulations were carried out and a result extracted that measures the effect on the target. The following fit to the results was developed.

$$\tau_0 = (0.42 - 0.004f)(E/2.8) \quad (1)$$

$$P = 0.02f \left(\frac{2.8}{E} \right) \left(\frac{\tau_0}{\tau} \right) \left(1 - \exp \left[\left(\frac{\tau}{\tau_0} \right)^3 \right] \right)^{\frac{1}{3}} \quad (2)$$

Here, f is the peak fluence in units of J/cm^2 , τ is the FWHM in ns, E is the ion energy in MeV, and P is the figure of merit, or performance metric. This is shown in Figure 1.

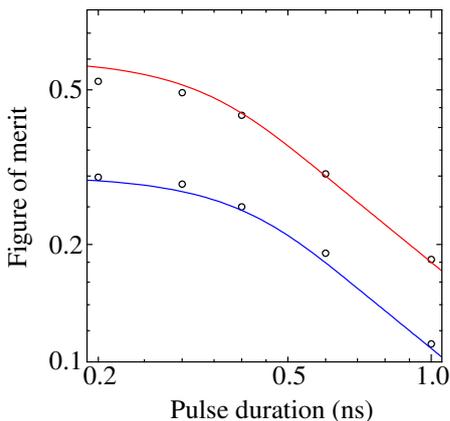


Figure 1: The figure of merit is a fit to target simulation results given a range of ion beam inputs. It is used as a measure of how hard the beam drives the target. The red and blue curves are fits to cases with $30 \text{ J}/\text{cm}^2$ and $15 \text{ J}/\text{cm}^2$ of beam fluence on target.

The simulations were carried out using the Warp code[4], which uses the particle-in-cell method to include the self-

field of the beam, combined with the applied field from the accelerator lattice elements. The simulations were launched at the source, a hot-plate alumina-silicate emitter, assuming space-charge limited emission and including the full geometry. Where appropriate, depending on the errors, axisymmetry was assumed, otherwise the simulations were 3-D. The beam was propagated through the accelerator including all applied and self-fields. For these simulations, the neutralizing plasma was assumed to be perfect - the space-charge was assumed to be completely neutralized. This was accomplished in the simulation by first zeroing out the charge of the part of the beam that is inside the plasma, solving Poisson's equation, and then zeroing out the potential inside the plasma (this method is only important as the beam enters the plasma). The beam was followed through the final focusing solenoid and onto the target.

For each type of error, ensembles of simulations were run. For each ensemble, a maximum value of the error was specified. For each case in the ensemble, the actual errors were chosen randomly, typically from a uniform distribution, with the specified maximum value - each case used a different random number seed. The results from the cases were combined to give an estimate of the effect of the error. In the plots of the ensembles that follow, each case is shown as a circle, showing the spread of the results, and the averages are shown with a red line. Note that each case also used a different seed for the generation of the particles. This results in a small spread in the results independent of the error; this can be seen in the plots, where there is finite spread even with zero errors.

2. Solenoid field errors

Despite careful construction, some level of field error is unavoidable. In the solenoids, there are two major sources of errors: field errors due to the construction of the solenoids, and errors due to misalignments. In the simplest terms, the solenoids consist of wires wrapped around the beam tube. The quality of the field is determined by the layout of the wires. Errors can be introduced by the layout itself, due to fabrication constraints and the need for leads, for example. Errors can also arise from inaccuracies in the construction. For example, wires bunched on one side and spread out on the other would produce an unwanted dipole field. Alignment errors can also produce dipole fields, since the solenoid magnetic axis would not be aligned with the beam axis. Both of these sources of errors were included in the simulations.

Prototype solenoids were built and measured - these had dipoles due to construction errors of order 25 G when the peak B_z was 1 T. This measured error, along with the profile of B_z , was imported into the simulations and replicated for each solenoid. Ensembles were carried out scaling the size of the dipoles, by factors of between 0 and 5. The results are shown in Figure 2. At the expected

error size, a scale factor of 1, the effect on the beam performance is small. Note that here, the errors were fixed - the ensembles only show the spread in the simulation results due to particle noise.

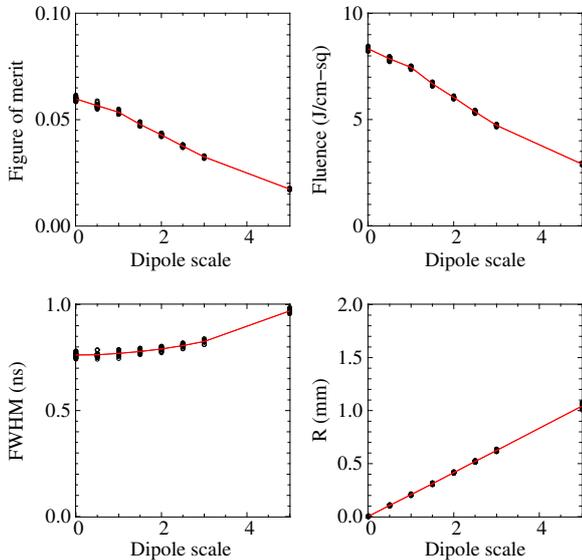


Figure 2: Ensembles with solenoid construction errors. The size of the error is scaled from the measured values from a prototype solenoid.

For solenoid misalignments, an absolute maximum on the tolerance is the requirement that the beam does not wander far enough off axis to scrape on the wall. Beyond this, it is desirable that the beam should be relatively well centered on the target with minimal degradation of the spot size and shape. Because of the longitudinal velocity tilt, transverse kicks on the beam from the misalignments will vary along the beam, giving rise to a helical or “corkscrew” deformation. On the target, this will spread out the focal spot. In the simulations, the ends of the solenoid were displaced randomly with a uniform distribution out to a specified maximum value. This resulted in both offsets and tilt errors. These simulations included the estimated dipole from the construction errors, with a scale factor of 1. Results are shown in Figure 3. Acceptable degradation is seen with offsets up to 0.5 mm, which is within the expected achievable construction tolerance. In all of the cases, up to 2 mm solenoid offsets, no particle loss was seen.

These errors are fixed in the machine once it has been constructed. While the beam may not perform as well as it would have without the errors, these parameters do not vary shot-to-shot. Compensation can counteract the errors and bring the performance closer to the ideal. NDCX-II has controllable dipoles in a number of places along the machine, designed for beam steering. (The ensemble calculations did not include steering.) It is calculated that the machine can perform rather well without steering - this avoids complication during machine commissioning. Steering will be employed later to further improve the beam

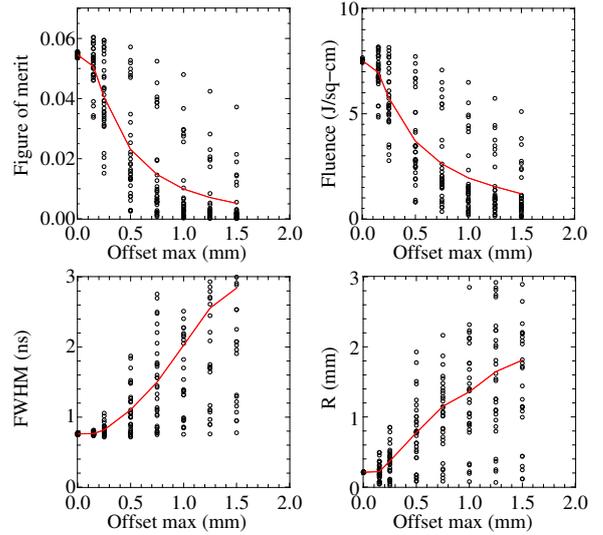


Figure 3: Ensembles with solenoid misalignments.

performance.

Independent of alignment and construction errors, there can be errors in the solenoid strengths as well. It is not expected that random, but repeatable, errors will be a problem. When tuning the system, the solenoid strengths will be adjusted to keep the beam well matched - the actual values of the solenoid fields will not matter very much. There can, however, be systematic errors that need to be taken into account. One possible issue that needs to be taken into account is slow, shot-to-shot shifts in the solenoid strengths over a period of operation, due to heating of the solenoid coils during operation (the heating increases resistance causing decreased current and field strength). Without any other errors, a scan was carried out to examine the effect of decreasing solenoid strength. Figure 4 shows the resulting figure of merit and FWHM. Without other errors, the degradation is small.

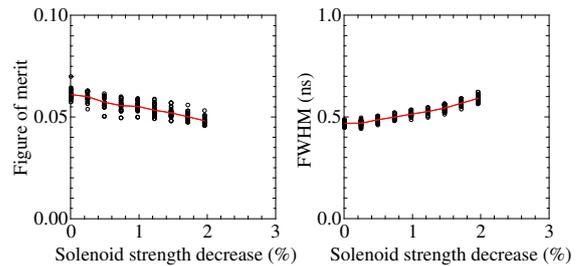


Figure 4: Ensembles with solenoid strength errors.

However, when combined with other errors, a shift in the dipole strengths can be more significant. Cases with combined solenoid offsets and strength decreases were examined. For each simulation, a particular set of randomly chosen solenoid offsets, with a maximum of 0.5 mm, was used scanning the decrease in the solenoid strength. Figure 5 shows some results from a typical parameter scan. As the solenoid strength varies, the size of the dipoles (re-

sulting from misalignments) varies as well. This changes the offset of the beam and moves the location of the focal spot. In the particular case shown, with a decrease of just over 2%, the focal spot would shift 0.1 mm, the size of the nominal diagnosable spot. This problem would need to be handled during operation and can be compensated for by one of several possible methods, for example adjustment of the solenoid drive to maintain constant current output and temperature regulation of the solenoid.

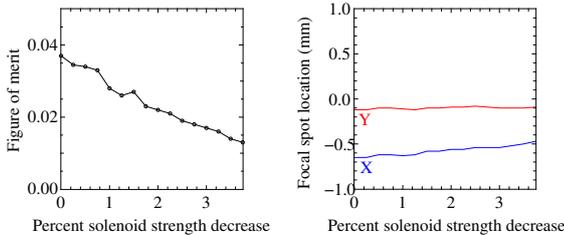


Figure 5: Effects of solenoid strength errors.

3. Accelerator waveform errors

Much of the complexity of designing the NDCX-II accelerator was in the accelerating waveforms. There is a very large parameter space over which optimization were carried out - the waveform shapes, amplitudes, durations, and timings for each of the twelve induction cells (or more, for longer versions of NDCX-II). The basic design was done using the ASP code[5], which is a 1-D particle code that was developed for this purpose. The code is fast, but has only approximate models. For example, since the code does not model the transverse size of the beam, it can only have an approximate model for the effect of the surrounding beam pipe, which reduces the longitudinal self-fields (because of image charges) - it can use one of several simplified models. The waveforms so derived are imported into Warp, which can do a more complete simulation.

Since the models in ASP are approximate, a set of waveforms optimized in ASP will not necessarily be optimal in the Warp calculation. Because of this, when examining errors in the waveforms, i.e. variations from the ASP generated waveforms, improved cases may be found. Because of the large parameter space, and the relative slowness of the Warp calculations limiting the number of simulations that can be done, our experience is that the error ensembles have proven to be a better approach to further optimization than standard multivariate optimization algorithms.

The waveform shapes and durations are constrained by the ATA Blumleins, but also by the custom lumped element pulsers fabricated for NDCX-II - change requires component replacement. The timing and amplitude, however, can be readily varied. It is not expected that the shot-to-shot variability of the amplitudes will be significant.

The most significant errors in the induction cells are timing errors. The first seven waveforms are nonlinear ramps of with individual profiles that were carefully designed to shape and compress the beam, as well as accelerate it. In the presence of a timing error, the beam will be accelerated by the wrong part of the ramp, affecting the shaping that is done. Furthermore, the overall acceleration of the beam will be in error, causing it to arrive at the next gap at the wrong time, etc. Similarly, the last three waveforms are ramps that build up a velocity tilt on the beam, setting it up for the neutralized drift compression. If there are errors in the degree of velocity tilt or in the beam energy, the plane of peak compression will shift away from the target plane, degrading performance. Random timing errors are problematic since they affect the shot-to-shot repeatability.

As with other errors, ensembles of calculations were done in Warp with randomly generated timing errors in the waveforms. In the simulations, all other aspects were held fixed, the most important being the location of the target plane. During optimization, the target plane can be adjusted to put it at the plane of peak compression. However, timing errors vary from shot-to-shot and so such fine tuning of the target plane is not possible.

Results are shown in Figure ???. It is estimated that the timing jitter on the waveforms will be about 2-3 ns, based on test stand measurements. The results for the baseline case show that this degree of jitter is not expected to be a problem. Note that for the figure of merit, as expected, there are cases which do better than the base cases with no jitter. An optimization path is to start with the best case and begin the ensembles about that point. The next iteration is shown in Figure 7a. This turned up cases that are significantly better. The best of those was used as the starting point for a new ensemble, shown in Figure 7b. The process can be continued, but note that the sensitivity of the figure of merit to jitter increases with each successive iteration. As the system is more finely tuned, it becomes more sensitive to errors. Except in case when the highest possible performance is needed, this increased sensitivity to errors should be avoided since it makes experiments less reliable.

4. Other errors

In addition to the errors discussed above, other errors were examined. Offsets of various components, such as the source and extractor plates in the injector, were considered. For offsets within what is reasonably achievable, up to about 1 mm, little to no degradation was seen in the beam performance.

Errors and noise on the waveforms were examined, both on the injector voltage waveform and on the waveforms in the induction cells. Over a wide range of frequencies of added noise, little to no performance loss was seen for noise up to several percent of the signal. For noise up to 10%, the figure of merit averaged over the ensembles

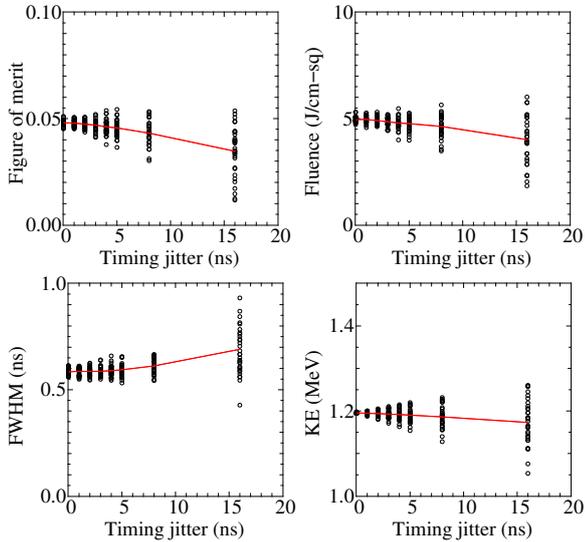


Figure 6: Ensembles with waveform timing errors.

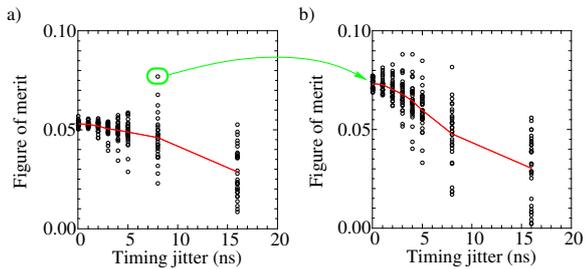


Figure 7: Optimization using ensembles with waveform timing errors. The case with the highest figure of merit in (a) is used as the base case in (b).

showed little degradation, but the spread in the results increased, with low frequency noise creating the largest spread. It is estimated that the shot-to-shot noise in the waveforms will be of order a percent.

5. Final focus solenoid

Another optimization carried out was fine tuning the final focus solenoid. This optimization cannot be done using the lower dimensional ASP code - it requires at least 2-D simulation and so was done with Warp. While the solenoid field strength is constrained by the physical design of the magnet, there is some flexibility. It might be expected that increasing the field to its maximum would always be better since it should give a smaller focal spot size, increasing the power density. This was found not to be the case.

Sample results are shown in Figure 8. Note that as expected, with increasing solenoid strength, the beam becomes more focused and the fluence on target steadily increases. However, the pulse duration (the FWHM) shows a marked increase with increasing solenoid strength. This increase is due to non-paraxial pulse broadening[6] - the action of the solenoid and focusing increases the transverse velocity component while decreasing the longitudinal velocity of particles off axis, causing them to be delayed, and therefore spreading out the arrival time at the target. This spreading out of the pulse dilutes the power density, eventually causing a decrease in the figure of merit. The z kinetic energy as shown in the figure (which includes only the longitudinal velocity component) drops with increasing solenoid strength because of the increasing transverse convergence angle. In the case shown, the optimal final focus field strength is around 10 T.

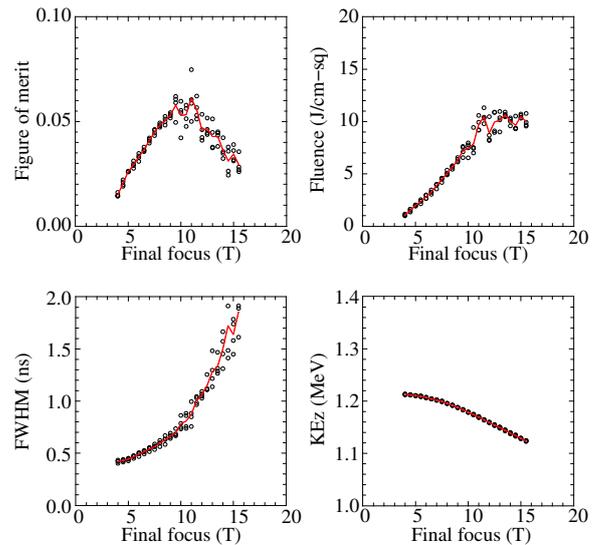


Figure 8: Optimization of the final focus strength. Because of the non-paraxial pulse broadening effect, the figure of merit has a maximum value.

The non-paraxial pulse broadening effect plays a role in the beam performance degradation visible in some of the cases shown earlier. In Figures 2 and 3, the FWHM can be seen to increase with increasing size of the error being varied. In both cases, the beam is being kicked further off axis increasing the non paraxial time delay. In Figure 4, the beam is becoming larger with decreasing solenoid strength, increasing the pulse broadening.

6. Conclusion

The NDCX-II accelerator was designed to produce moderate energy, high current beams that will be used to drive targets for warm dense matter studies and as a testbed for accelerator physics studies relevant to drivers for heavy ion driven inertial fusion energy. Because of the complex design and high space-charge of the beam, simulations were essential to designing the accelerator and will be essential to its effective use. Ensembles of simulations were carried out to examine various possible machine errors, both static and shot-to-shot. Overall, the simulations indicate that NDCX-II should operate without serious beam degradation if errors are within nominal tolerances. Solenoid offsets do show a potential to cause noticeable loss in performance, but it is expected that this can largely be compensated for by the use of beam steering, and should not affect the shot-to-shot reliability. In some cases, such as with timing errors on the waveforms, examining “errors” showed pathways to improved beam performance. Optimization of the final focus strength revealed a new effect, non paraxial pulse broadening, that needs to be taken into account.

7. Acknowledgements

We would like to thank Will Waldron and Diego Arbelez for useful discussions about constraints and tolerances on NDCX-II. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

8. References

- [1] J. Barnard, J. Armijo, R. More, A. Friedman, I. Kaganovich, B. Logan, M. Marinak, G. Penn, A. Sefkow, P. Santhanam, P. Stoltz, S. Veitzer, J. Wurtele, Theory and simulation of warm dense matter targets, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 577 (1-2) (2007) 275 – 283, proceedings of the 16th International Symposium on Heavy Ion Inertial Fusion, HIF06. doi:10.1016/j.nima.2007.02.062. URL <http://www.sciencedirect.com/science/article/pii/S0168900207003646>
- [2] W. M. Sharp, J. Barnard, R. Cohen, M. Dorf, A. Friedman, D. Grote, S. Lund, L. Perkins, M. Terry, F. Bieniosek, A. Faltens, E. Henestroza, J.-Y. Jung, A. Koniges, J. Kwan, E. P. Lee, S. Lidia, B. Logan, P. Ni, L. Reginato, P. Roy, P. Seidl, J. Takakuwa, J.-L. Vay, W. Waldron, R. Davidson, E. Gilson, I. Kaganovich, H. Qin, E. Startsev, I. Haber, R. Kishkek, Inertial fusion driven by intense heavy-ion beams, Proceedings of the 2011 Particle Accelerator Conference (2011) 1386 – 1393. URL <http://accelconf.web.cern.ch/AccelConf/PAC2011/papers/weoas1.pdf>
- [3] A. Friedman, J. J. Barnard, R. H. Cohen, D. P. Grote, S. M. Lund, W. M. Sharp, A. Faltens, E. Henestroza, J.-Y. Jung, J. W. Kwan, E. P. Lee, M. A. Leitner, B. G. Logan, J.-L. Vay, W. L. Waldron, R. C. Davidson, M. Dorf, E. P. Gilson, I. D. Kaganovich, Beam dynamics of the Neutralized Drift Compression Experiment-II, a novel pulse-compressing ion accelerator, Physics of Plasmas 17 (5) (2010) 056704. doi:10.1063/1.3292634. URL <http://link.aip.org/link/?PHP/17/056704/1>
- [4] D. P. Grote, A. Friedman, J.-L. Vay, I. Haber, The WARP code: Modeling high intensity ion beams, AIP Conference Proceedings 749 (1) (2005) 55–58. doi:10.1063/1.1893366. URL <http://link.aip.org/link/?APC/749/55/1>
- [5] W. Sharp, A. Friedman, D. Grote, E. Henestroza, M. Leitner, W. Waldron, Developing acceleration schedules for NDCX-II, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 606 (12) (2009) 97 – 101, proceedings of the 17th International Symposium on Heavy Ion Inertial Fusion, HIF09. doi:10.1016/j.nima.2009.03.229. URL <http://www.sciencedirect.com/science/article/pii/S0168900209005555>
- [6] D. P. Grote, A. Friedman, E. P. Lee, Non-paraxial pulse broadening in a solenoid focusing element, submitted to PRSTAB.