

Title of Grant / Cooperative Agreement:	Continuous Electrode Inertial Electrostatic Confinement Fusion
Type of Report:	Summary of Research
Name of Principal Investigator:	Raymond J. Sedwick
Period Covered by Report:	05/01/2017 - 05/31/2018
Name and Address of recipient's institution:	University of Maryland, 3112 Lee Bldg, College Park, MD 20742
NASA Grant / Cooperative Agreement Number:	NNX17AJ72G

Reference 2 CFR § 1800.908 or 14 CFR § 1260.28 Patent Rights as applicable (abbreviated below)

The recipient may use whatever format is convenient to disclose subject invention required in subparagraph (c)(1). NASA prefers that the recipient use either the electronic or paper version of NASA Form 1679, Disclosure of Invention and New Technology (Including Software), to disclose subject inventions. Both the electronic and paper version of the NASA Form 1679 may be accessed at the electronic New Technology Reporting Web site <https://invention.nasa.gov>.

A final new technology summary report listing all subject inventions (or a statement certifying there were none) for the entire award period; which report shall be submitted within 90 days after the end date for the period of performance within the designated system noted within the award document."

Have any Subject Inventions / New Technology Items resulted from work performed under this Grant / Cooperative Agreement?	No ☐	Yes ☒
If yes a complete listing should be provided here: Details can be provided in the body of the Summary of Research report.	Patent Application No. 20180033496. Systems Methods and Devices for Inertial Electrostatic Confinement	

Reference 2 CFR § 1800.907 or 14 CFR § 1260.27 Equipment and Other Property as applicable (abbreviated below)

A Final Inventory Report of Federally Owned Property, including equipment where title was taken by the Government, will be submitted by the Recipient no later than 60 days after the expiration date of the grant. Negative responses for Final Inventory Reports are required.

Is there any Federally Owned Property, either Government Furnished or Grantee Acquired, in the custody of the Recipient?	No ☒	Yes ☐
If yes please attach a complete listing including information as set forth at § 1260.134(f)(1).		

Attach the Summary of Research text behind this cover sheet.

Reference 2 CFR § 1800.902 or 14 CFR § 1260.22 Technical publications and reports as applicable (abbreviated below)

Reports shall be in the English language, informal in nature, and ordinarily not exceed three pages (not counting bibliographies, abstracts, and lists of other media).

A Summary of Research (or Educational Activity Report in the case of Education Grants) is due within 90 days after the expiration date of the grant, regardless of whether or not support is continued under another grant. This report shall be a comprehensive summary of significant accomplishments during the duration of the grant.

SUMMARY OF RESEARCH

Ion Confinement Physics

The NIAC Phase I project on Inertial Electrostatic Confinement was a continuation of early stage research that was funded by an NSTRF. The student on the project, Andrew Chap, was funded by the NSTRF from Fall 2013 through the Summer of 2017, and then was funded on the NIAC through the completion of his PhD. A significant amount of work targeting the plasma confinement physics was the focus of his NSTRF, and over the course of that effort he developed a number of analyses and computational tools that leveraged GPU parallelization. A detailed discussion of these models can be found in his dissertation, which has been included as Appendix D in this report. As a requirement for the NSTRF, Andrew's full dissertation was submitted at the end of the program.

Having developed the computational tools, a substantial amount of simulation and analyses leveraging those tools were conducted during the Fall of 2017, under the auspices of the NIAC funded research. Much of this work targeted optimization of the confinement fields, investigating their structure and the possible advantages of having them be time-varying. The results of these simulations can also be found in Appendix D.

One of the main results from this research is that the density of ions electrostatically confined within the system can indeed be increased by several orders of magnitude by optimizing the radial potential distribution, and by dynamically varying these fields to maintain compressed ion bunches. An electron population can also be confined within the core by a static radial cusped magnetic field, which helps to support a greater ion density within the core. The issue with the confinement mechanism is that as the ion densities are increased toward fusion-relevant levels, the electrostatic forces generated by the confined electron population become so great that the ions are no longer energetic enough to leave the device core. As their excursions into the outer channels are diminished, the mechanism that is used to maintain their non-thermal velocity distributions becomes ineffective, and eventually the ions become fully confined within the core, where they thermalize.

A possible fix to the problem comes by discarding the active ion control (a main pillar of the concept) but retaining the structure of the permanent magnet confinement of the electron population. Such cusped field confinement has been the focus of other IEC approaches (e.g. Polywell), but the high transparency of the permanent magnet structure lends itself to better ion extraction and power conversion (a second pillar of the concept). The question then becomes whether any influence on the ion evolution within the core can be achieved to slow the thermalization of the ions. Such approaches have been studied in highly idealized analytic models, but face major criticisms within the literature. While this is a possible path forward, the uncertainty in the approach did not warrant committing NIAC Phase II resources to investigating the concept at this time.

System Modeling

Another main focus of the NIAC Phase I was to look at the technology from a systems perspective, more broadly than just the confinement physics. The results of this study are codified in a paper that was presented at the 2018 American Astronautical Society Guidance & Control Conference in Breckenridge, CO. This paper has been included in Appendix B of this report. It provides a higher-level overview of the potential system performance under the assumption that the ion confinement issues can be resolved. If the thermalization can be resolved, the system performance becomes very

attractive, with a high conversion efficiency and low specific mass. Detailed performance parameters are discussed in the paper.

A computer animation that provides an overview of the concept was commissioned as part of the Phase I effort. The movie file can be accessed here:

https://www.dropbox.com/s/9aznyvuy2ftl61m/Continuous_Grid_IEC.mov?dl=0

Patent Pending

A patent application for the full system was submitted on July 26, 2017. This application is included in Appendix A of this report.

Core Confinement

A first step to investigating the evolution of the ion population confined to the core is to create a better model for how the electron population in the core dynamically changes the confinement magnetic field. The magnetic field induces surface currents in the electrons which in turn modify the magnetic field. The impact is to create larger field gradients that can actually aid in confinement. This research was started in the Fall of 2017, but continued through May, 2018, which was the main motivation for requesting a no-cost extension. The results of the research were presented as a student paper at the AIAA Region I student conference, and the paper is included as Appendix C in this report.

The modeling of the electron confinement proved challenging, and while three different models were developed, none of them were able to produce satisfyingly consistent results. There is some indication that the issues may simply have arisen from instability in the numerical algorithms that were being employed. While the current results are inconclusive, there is research in the literature that indicates enhanced confinement in high- β plasmas can be achieved. More research on the topic is planned outside of NIAC support, and if positive results can be obtained a follow-on proposal will be submitted.

Appendix A:

Patent Application No. US 2018/0033496 A1
Systems, Methods, and Devices for Inertial Electrostatic Confinement

Raymond Sedwick
Andrew Chap


Patent application filed July 26, 2017



US 20180033496A1

(19) **United States**(12) **Patent Application Publication** (10) **Pub. No.: US 2018/0033496 A1****SEDWICK et al.**(43) **Pub. Date: Feb. 1, 2018**(54) **SYSTEMS, METHODS, AND DEVICES FOR
INERTIAL ELECTROSTATIC
CONFINEMENT**(52) **U.S. Cl.**
CPC **G21B 1/03** (2013.01); **B64G 1/408**
(2013.01); **B64G 1/405** (2013.01)(71) Applicant: **University of Maryland, College Park,**
College Park, MD (US)(72) Inventors: **Raymond J. SEDWICK**, University
Park, MD (US); **Andrew CHAP**,
Chevy Chase, MD (US)(21) Appl. No.: **15/659,962**(22) Filed: **Jul. 26, 2017****Related U.S. Application Data**(60) Provisional application No. 62/367,410, filed on Jul.
27, 2016.**Publication Classification**(51) **Int. Cl.**
G21B 1/03 (2006.01)
B64G 1/40 (2006.01)(57) **ABSTRACT**

A continuous electrode (CE) inertial electrostatic confinement (IEC) device has particle paths radially extending from a central core region. Each particle path has a corresponding particle path aligned on an opposite side of the central core region. Sidewalls bounding the particle paths provide continuous surfaces radially extending from a cathode region proximal to the central core region to an anode region remote from the central core region. Electrodes are coupled to the sidewalls to provide an electric field that varies along each particle path from the cathode region to the anode region. The CE-IEC device can be used for particle fusion by directing ions along the particle paths to the central core region, for example, to generate power or to propel a spacecraft.






FIG. 1A




FIG. 1B




FIG. 1C




FIG. 1D




FIG. 1E




FIG. 1F




FIG. 1G




FIG. 2A




FIG. 2B




FIG. 3A




FIG. 3B




FIG. 3C




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 5




FIG. 6




FIG. 7A




FIG. 7B




FIG. 7C




FIG. 7D




FIG. 8




FIG. 9




FIG. 10




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 11D




FIG. 11E




FIG. 12A




FIG. 12B

SYSTEMS, METHODS, AND DEVICES FOR INERTIAL ELECTROSTATIC CONFINEMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of U.S. Provisional Application No. 62/367,410, filed Jul. 27, 2016, which is hereby incorporated by reference herein in its entirety.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] This invention was made with government support under NNX13AL44H awarded by the National Aeronautics and Space Administration (NASA). The government has certain rights in the invention.

FIELD

[0003] The present disclosure generally relates to nuclear fusion, and, more particularly, to fusion by inertial electrostatic confinement using a device with continuous walls.

SUMMARY

[0004] In embodiments, a continuous electrode (CE) inertial electrostatic confinement (IEC) device employs sidewalls with substantially continuous surfaces radially extending from a central core to define radial particle paths. Electrodes coupled to the sidewalls provide an electric field that varies along each particle path to accelerate ions within the particle paths toward the core. Interaction of the ions within the core can result in nuclear fusion, which may be used for electricity generation or for spacecraft propulsion. The CE-IEC device can include one or more features designed to decrease distances between ions, for example, by compacting ion bunches as they travel along the particle paths and/or neutralizing space charge of ion bunches within the core using a population of electrons captured therein.

[0005] In one or more embodiments, a device comprises a central core region, particle paths, sidewalls, electrodes, and a control module. Each particle path can radially extend from the central core region and can have a corresponding particle path aligned therewith on an opposite side of the central core region. The sidewalls can extend in a radial direction. Each particle path can be bounded by a corresponding set of the sidewalls. The electrodes can be coupled to the sidewalls so as to provide an electric field that varies along each particle path from a cathode region proximal to the central core region to an anode region remote from the central core region. The control module can control the electrodes to provide the electric field. Each sidewall can provide a continuous surface radially extending from the cathode region to the anode region.

[0006] In one or more embodiments, a fusion method comprises directing ion bunches along particle paths that radially extend from a central core region. Each particle path can be bounded by a corresponding set of radially extending sidewalls and can have a corresponding particle path aligned therewith on an opposite side of the central core region. Each sidewall can provide a continuous surface radially extending from a cathode region proximal to the central core region to an anode region remote from the central core region. The method can further comprise generating an

electric field that varies along each particle path from the cathode region to the anode region such that the ion bunches are accelerated toward the central core region, and fusing ions from the ion bunches within the central core region. The method can also comprise allowing fusion products to travel from the central core region to beyond the anode region via the particle paths.

[0007] Objects and advantages of embodiments of the disclosed subject matter will become apparent from the following description when considered in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Embodiments will hereinafter be described with reference to the accompanying drawings, which have not necessarily been drawn to scale. Where applicable, some features may not be illustrated or otherwise simplified to assist in the illustration and description of underlying features. Throughout the figures, like reference numerals denote like elements.

[0009] FIG. 1A is an isometric view of a CE-IEC device employing a square configuration, according to one or more embodiments of the disclosed subject matter.

[0010] FIG. 1B is a view of FIG. 1A along one of the particle paths, according to one or more embodiments of the disclosed subject matter.

[0011] FIG. 1C is a cross-sectional view of a CE-IEC device, according to one or more embodiments of the disclosed subject matter.

[0012] FIG. 1D is a graph of an exemplary electrical potential profile along one of the particle paths of FIG. 1C, according to one or more embodiments of the disclosed subject matter.

[0013] FIG. 1E-1G are cross-sectional views of a CE-IEC device with ion bunches traveling and interacting to produce fusion products, according to one or more embodiments of the disclosed subject matter.

[0014] FIG. 2A is a cross-sectional view of a CE-IEC device employing sidewalls with variable resistivity, according to one or more embodiments of the disclosed subject matter.

[0015] FIG. 2B is a cross-sectional view of a CE-IEC device employing sidewalls with segments having different electric potentials applied thereto, according to one or more embodiments of the disclosed subject matter.

[0016] FIG. 3A is a cross-sectional view showing construction of a sidewall in a CE-IEC device that has a permanent magnet therein, according to one or more embodiments of the disclosed subject matter.

[0017] FIG. 3B is a cross-sectional view showing construction of a sidewall in a CE-IEC device that is formed of a permanent magnet, according to one or more embodiments of the disclosed subject matter.

[0018] FIG. 3C shows exemplary magnetic field lines for a CE-IEC device employing the sidewall construction of either FIG. 3A or FIG. 3B.

[0019] FIG. 4A is a cross-sectional view showing construction of a sidewall in a CE-IEC device that has multiple permanent magnets therein, according to one or more embodiments of the disclosed subject matter.

[0020] FIG. 4B is a cross-sectional view showing construction of a sidewall in a CE-IEC device that is formed of multiple permanent magnets, according to one or more embodiments of the disclosed subject matter.

[0021] FIG. 4C shows exemplary magnetic field lines for a CE-IEC device employing the sidewall construction of either FIG. 4A or FIG. 4B.

[0022] FIG. 5 is a magnified cross-sectional view of a CE-IEC device employing protective standoffs, according to one or more embodiments of the disclosed subject matter.

[0023] FIG. 6 is an isometric view of a CE-IEC device employing a square configuration with conduits at vertices between adjacent sidewalls, according to one or more embodiments of the disclosed subject matter.

[0024] FIG. 7A is a cross-sectional view showing an exemplary configuration for electrical coupling for a variable resistivity sidewall via a conduit, according to one or more embodiments of the disclosed subject matter.

[0025] FIG. 7B is a cross-sectional view showing an exemplary configuration for electrical coupling for a variable resistivity sidewall via a conduit with a permanent magnet, according to one or more embodiments of the disclosed subject matter.

[0026] FIG. 7C is a cross-sectional view showing an exemplary configuration for electrical coupling for a segment of a sidewall via a conduit, according to one or more embodiments of the disclosed subject matter.

[0027] FIG. 7D is a cross-sectional view showing an exemplary configuration for electrical coupling for multiple segments of sidewalls via a conduit, according to one or more embodiments of the disclosed subject matter.

[0028] FIG. 8 is a cross-sectional view showing an exemplary configuration for ion supply to the particle pathways via a conduit, according to one or more embodiments of the disclosed subject matter.

[0029] FIG. 9 is a cross-sectional view showing an exemplary configuration of sensors and signal communication via a conduit, according to one or more embodiments of the disclosed subject matter.

[0030] FIG. 10 is a simplified schematic of a CE-IEC system for particle fusion, according to one or more embodiments of the disclosed subject matter.

[0031] FIGS. 11A-11C show stages in converting fusion products to electricity using standing wave direct energy conversion (SW-DEC), according to one or more embodiments of the disclosed subject matter.

[0032] FIG. 11D is a simplified schematic of an RLC circuit employing SW-DEC for direct energy usage, according to one or more embodiments of the disclosed subject matter.

[0033] FIG. 11E is a cross-sectional view of a CE-IEC device with SW-DEC rings in an outer power conversion region, according to one or more embodiments of the disclosed subject matter.

[0034] FIG. 12A is an isometric view of a CE-IEC device employing a special irregular truncated icosahedron (SITI) configuration, according to one or more embodiments of the disclosed subject matter.

[0035] FIG. 12B is a cutaway view of FIG. 12A showing exemplary magnetic field lines, magnets embedded in walls, and an exemplary electric potential profile.

DETAILED DESCRIPTION

[0036] In embodiments, an inertial electrostatic confinement (IEC) device has radially extending sidewalls that define radial particle paths emanating from a central core. Electrodes can be coupled to these sidewalls in order to provide an electric field that varies along the particle paths,

for example, from a radially-outer anode region (remote from the central core) to a radially-inner cathode region (proximal to the central core). As such, embodiments employing continuous sidewalls with coupled electrodes are referred to herein as continuous electrode (CE). Although the sidewalls are referred to as continuous, this does not require that the sidewalls be monolithic or isotropic. Rather, the continuous sidewalls provide a substantially continuous surface from the anode region to the cathode region and can have different properties at different radial or azimuthal locations (e.g., formed of different material segments at different radial locations and/or having a resistivity that varies radially).

[0037] The electric field can accelerate ions toward the core, where the ions interact to yield nuclear fusion. Non-reacting ions can pass through the core to an opposite (aligned) particle path, where the electric field therein slows the ions to reverse direction and accelerate back toward the core for further interactions. Unlike conventional IEC devices that employ multiple independent grids of different radii centered on the core to provide an electric field, embodiments of the present disclosure employing continuous sidewalls allow the electric potential to be imposed continuously over the particle path and/or to dynamically adjust the electric potential as ions travel within the particle paths (e.g., to compact traveling ion bunches prior to introduction to the core). The continuous sidewalls can also provide a very high grid transparency (i.e., greater than 70%, for example, ~85%) as seen from the core without otherwise sacrificing structural rigidity.

[0038] The sidewalls also provide real estate for conduits from external to the device toward the central core, for example, to provide electrical connections for power or signal transmission, to provide a magnetic field using permanent magnets, and/or to feed ions into the particle paths. While the sidewalls provide an additional surface area that ions could strike (representing a loss to the system) that multiple independent grids would otherwise lack, such ions would be on non-radial trajectories and thus would likely not contribute to nuclear fusion anyway.

[0039] Since the ions interact in the core, fusion products are only generated in the core and leave along predominantly radial paths, e.g., along the radially extending particle paths. In general, the IEC device should be as transparent as possible to these energetic particles. In other words, the construction of the IEC device (e.g., sidewalls) should subtend as little solid angle (as seen by the core) as possible. To that end, any radially outer structure that falls within the "shadow" of radially inner structures, as viewed from the device center, will not diminish the transparency of the system. Thus, in embodiments, structures are arranged to fall within this "shadow" of the innermost structure, which is designed to subtend as little solid angle as possible. In other words, the sidewalls could be considered a radially outward extrusion of the innermost edge (adjacent to the core).

[0040] FIG. 1A shows an isometric view of a simplified CE-IEC device 100. FIG. 1B shows a side view of the CE-IEC device 100. The simplified CE-IEC device 100 employs a cubic configuration, where the innermost edge 104 formed by radially extending sidewalls 106 is a cube surrounding a central core region 112. The outermost edge 102 formed by the sidewalls 106 can also be a cube. The sidewalls 106 can be continuous between the innermost edge 104 and outermost edge 102 and define three radial particle

paths 108a-108c. Each sidewall 106 can abut an adjacent sidewall 106 along a radially extending vertex 114. Each particle path is bounded on four sides (defining a square face perpendicular to the radial direction) by the sidewalls 106, such that each half of the particle path outside the core region 112 takes the form of a truncated pyramid. Note that the particle paths 108a-108c can be considered a single path that extends through the central core region 112, or as separate but aligned paths on opposite sides of the central core region.

[0041] Referring to FIGS. 1C and 1E-1G, cross-sectional views of CE-IEC device 100 are shown in order to describe various features thereof. In particular, the CE-IEC device 100 can be considered to have four main regions along the radial direction 116: (1) an innermost core region 122, (2) an outer core region 112 surrounding the innermost core region 122, (3) a focusing region 110 between sidewalls 106 and along particle paths 108a-108c, and (4) a power conversion region 150 beyond the sidewalls 106.

[0042] The focusing region 110 is where the CE structure (i.e., sidewalls 106) is disposed. The open channels between facing sidewalls 106 in the focusing region 110 form the particle paths 108a-108c, along which ion bunches recirculate as they pass into and out of the core. A radially outer region of the sidewalls 106 can be biased at a relatively higher voltage to form anode region 126, while a radially inner region of the sidewalls 106 can be biased at a relatively lower voltage to form cathode region 124. The resulting electric field in the focusing region 110 causes the ions to accelerate toward the core 112. Since the potential profile is generated solely in the focusing region 110, the ions will drift through the core at a constant speed.

[0043] Over time, the traveling ions self-assemble into bunches 120 due to two-stream instability. Moreover, due to cross-talk, the bunches 120 synchronize between different particle paths 108a-108c, as shown in FIG. 1E, such that the bunches 120 arrive at the inner core region 122 at substantially the same time, as shown in FIG. 1F. Alternatively or additionally, ions could be introduced into the particle paths 108a-108c in a pulsed manner (i.e., at different times), thereby forming preliminary ion bunches. The formation of bunches and the synchronization results in a beneficial situation since the ion bunches 120 only cross in the inner core 122.

[0044] Low angle scattering between counter-streaming ions, which would normally result in the rapid global thermalization of the ion population, is suppressed by local thermalization within each of the bunches 120 near the anode region 126. Moreover, the low-angle collisions among opposing bunches within the core merely reshuffle the specific radial trajectories among the particle paths 108a-108c that the individual ions will follow to exit the inner core 122. Upon refocusing (in focusing region 110), the velocity distribution in the azimuthal direction 118 within each bunch 120 should be indistinguishable from the start of the previous pass, cancelling any azimuthal momentum growth. Low angle scatters among non-opposing ion bunches can introduce both azimuthal and radial (energy) scattering. However, on average these scattering events will both up-scatter and down-scatter the ions equally. Upon refocusing (in focusing region 110), the net ion energy within each bunch 120 due to low angle ion-ion collisions should remain the same.

[0045] High angle scatters are confined to the core where the resulting ion trajectories will still be approximately radial. In other words, if the scattering angle of the ion does not otherwise cause it to collide with the inner edge of a sidewall 106, the ion should simply end up in a different channel (proceeding along a different particle path 108a-108c) rather than being lost. Upon refocusing (in focusing region 110), the ion will be merged into the traveling ion bunch 120. Thus, the electric potential within the particle pathways can help keep scattered ions from being lost.

[0046] To increase the density of ion bunches as they pass into the core 112, where the bulk coulomb repulsion of the ions would tend to push them apart, a population of electrons is confined to the core of the device in order to neutralize the space charge of the ions. For example, an electron population can be generated within the core by completely ionizing the fusion fuel (e.g., boron), which may be only singly ionized initially prior to injection. For example, any remaining electrons of the fusion fuel can be stripped, either through collisions with other ions passing through the core or as a result of a nuclear fusion event. Alternatively or additionally, electrons can be injected directly into the core, for example, to replenish those that might be lost over time and that would otherwise not be replenished by further ionizing the fusion fuel.

[0047] To help confine the electrons to the core, the sidewalls 106 may further provide another anode region 128 adjacent to the core (i.e., radially between the cathode region 124 and the outer core 112). The resulting electric field can create a reversed potential well for electrons that keeps them from escaping the core along particle pathways. For example, FIG. 1D shows a graph of an exemplary electric potential 130 along one of the particle paths 108a-108c. Region 132 represents a confinement field for the ions (i.e., between the sidewalls 106 along the particle pathways) while region 134 represents the confinement field for the electrons (i.e., within the core). Alternatively or additionally, a magnetic field may be provided to keep electrons from escaping along the particle paths 108a-108c, as described in further detail elsewhere herein.

[0048] The ions thus travel from the anode region 126 to the cathode region 124 (see FIG. 1E) and on to the outer core region 112 (see FIG. 1F), which is a transition region where the ions interact with the confined electron population. As the ions pass into the outer core 112, the electron population responds by being attracted to the ions, thereby neutralizing their space charge and allowing the ions to further compress along their radial trajectories as they move toward the inner core region 122. Within the inner core region 122, the ions traveling in crossing or opposite directions along the particle paths collide, exchanging energy and momentum and in some cases undergoing nuclear fusion. In general, the central core region may be substantially empty except for the electrons confined therein, ions traveling therethrough between the particle paths, and products resulting from interaction of the traveling ions.

[0049] Unreacted ions 120 and/or fusion products 140 travel along substantially radial paths out of the inner core region 122 to the outer core region 112, where they leave behind the confined electrons before passing back into the focusing region 110 (see FIG. 1G). The outer edge of the focusing region 110 (i.e., outermost edge of the anode region 126) represents the maximum radial extent to which the ion

bunches 120 are allowed to reach, with only fusion products 140 being energetic enough to proceed into power conversion region 150.

[0050] The fusion products 140 can thus continue through focusing region 110 and escape to power conversion region 150 beyond sidewalls 106, where the fusion products 140 can be converted to electricity or otherwise used to generate work (e.g., to propel a spacecraft). The fusion products 140 can enter the power conversion region 150 with a nearly isotropic angular distribution and can arrive in pulses a few nanoseconds long separated by several microseconds. The pulsed output of the fusion products 140 can be converted to electricity using a direct conversion process, for example, Traveling Wave Direct Energy Conversion (TW-DEC) or Standing Wave Direct Energy Conversion (SW-DEC), as described in further detail elsewhere herein.

[0051] When the ions leave the outer core 112, their self-charge will still tend to cause the bunches of ions to expand. Thus, as they travel along the particle paths 108a-108c in the focusing region 110, the ions 120 can also be continuously refocused and compressed (due to the electric potential provided by the sidewalls 106 and/or magnetic fields from permanent magnets of the sidewalls 106) to combat this natural spreading due to space charge and/or low angle collisions.

[0052] As noted above, embodiments of the CE-IEC device can provide a potential that varies continuously in the focusing region 110. FIG. 2A illustrates a configuration where the continuous sidewalls 106 can provide such an electric potential. In particular, the sidewall 106 can be formed of a material 200 that has a resistivity that varies along the radial direction 116. Portions of the sidewall 106 at different radii can be connected to a potential difference. For example, a voltage source 204 can be connected between a radially outer portion 202 (e.g., an anode region 126) and a radially inner portion 206 (e.g., a cathode region 124). For example, the outer portion 202 can be set to ground while the inner portion 206 can be set to -50 kV.

[0053] At each radial location on the sidewall, the resistivity normal to the radial direction can be very low, such that locations on the same sidewall at the same radial distance from the core can be held at substantially the same potential. In other words, portions of the sidewall 106 at the same radii would act as an isopotential conductor. Moreover, portions of different sidewalls 106 (i.e., falling along isopotential line 208 in FIG. 2A) can also be at substantially the same potential. Note that the isopotential line 208 connects locations on the sidewalls that are at the same potential, not between the sidewalls in the open regions of the particle paths where the potential would necessarily be different.

[0054] The variable resistivity material 200 may be accomplished, for example, by engineering a composite material having strips of different material layers at different radii extending in a direction perpendicular to the radial direction 116. Adjacent strips are electrically coupled to each other along contacting faces perpendicular to the radial direction 116 to provide the radially varying resistivity, while otherwise acting as isopotential conductors in a direction perpendicular to the radial direction. The composition of the material 200 can be customized to achieve any desired radial potential profile. For example, the potential profile can monotonically decrease from the anode region to the cathode region, or can be a complex profile that does not

necessarily monotonically decrease). For example, the sidewall material can be formed via 3-D printing.

[0055] Alternatively or additionally, a customized potential profile can be achieved using a segmented continuous sidewall structure, as illustrated in FIG. 2B. Each sidewall 106 (or a selection of sidewalls) can include a plurality of segments 252 at different radii from the core 112. Each segment 252 can be separated from an adjacent segment 252 in the radial direction by an insulating spacer 254, such that each segment 252 is electrically isolated from the other segments 252 of the sidewall 106. The segments 252 can thus be independently controlled (for example, by providing independent voltage sources 256 connected to the segments at 258) to provide a custom electric potential along the particle paths 108a-108c. Each segment 252 can be composed of the same material or of different materials, and can act as isopotential conductors in a direction perpendicular to the radial direction.

[0056] Although shown in FIG. 2B as being the same size, the segments 252 can also be of different sizes. For example, one or more intermediate segments 252 between anode region 126 and cathode region 124 can be made smaller than the other segments 252. These smaller segments 252 can be used to dynamically vary the potential as ion bunches pass, for example, to compact the ion bunches by providing a perturbation that slows ions at the front of the bunch and/or accelerates ions at the rear of the bunch. In addition, the number of spacers 254 and/or segments 252 illustrated in FIG. 2B are exemplary only, and other numbers are also possible according to one or more contemplated embodiments.

[0057] Moreover, the features of FIG. 2A and the features of FIG. 2B are not intended to be mutually exclusive. Rather, in some embodiments, the features of FIGS. 2A-2B can be combined to particular advantage. For example, one or more of the segments 252 of FIG. 2B can be formed of a material 200 that has a resistivity that varies along the radial direction 116. Other combinations and variations should be readily apparent to one of ordinary skill in the art.

[0058] As noted above, embodiments of the CE-IEC device can provide a magnetic field to help confine electrons to the core region 112. Since only enough electrons are needed to neutralize the traveling ions, the magnetic field requirement is relatively low and can be satisfied using permanent magnets (e.g., formed of a rare-earth material, such as neodymium, or other permanent magnetic material). For example, radially polarized permanent magnets may be incorporated into sidewalls 106 with same poles (either north or south) facing the core 112. The resulting magnetic field has field lines 312 extending through the channels formed by the sidewalls 106, e.g., substantially following the particle paths 108a-108c, as shown in FIG. 3C. At the core 112, the permanent magnets form a cusped magnetic field 314, which acts as a magnetic mirror that repels electrons from impacting ends of the sidewalls 106 facing the core 112. Moreover, the cusped magnetic field at the intersection of the particle paths 108a-108c and the outer boundary of the core 112 keep electrons from escaping along the particle paths 108a-108c.

[0059] The provision of permanent magnets in the sidewalls 106 allows more material to be used, thereby resulting in a stronger magnetic field, without otherwise compromising transparency to particles exiting the core 112. For example, the permanent magnets 306 can be incorporated

into each sidewall 106 between conductive panels 302 thereof, as shown in FIG. 3A. Alternatively or additionally, the permanent magnets 310 can be conductive panels of the sidewall 106 itself and used to provide both the magnetic and potential profiles, as shown in FIG. 3B. Depending on the desired profile of the magnetic field, the amount of magnetic material may be uniform or vary along the radial direction, and/or can be the same or different between different sidewalls at the same radius. The magnets can optionally be separated from the conductive panels 302 or other magnets 310 by an insulating spacer 304.

[0060] Each magnet would have a polar orientation 308 extending radially (i.e., with one pole adjacent to the core 112, and the opposite pole spaced at a radially outer location). Although a particular polar orientation is illustrated in FIGS. 3A-3B, embodiments of the disclosed subject matter are not limited thereto. Indeed, a polar orientation for each magnet opposite that illustrated may be adopted with similar effect.

[0061] Although FIGS. 3A-3C illustrate a configuration with the permanent magnets extending along the length of the sidewalls, embodiments of the disclosed subject matter are not limited thereto. Indeed, in some cases, electrons that approach the cusped magnetic field with a sufficiently low pitch angle (i.e., more parallel to the field lines) may not be turned around before the point of maximum field strength. As a result, some amount of electron leakage from the core may occur. Electrons that escape the core would be accelerated by the focusing region 110 to leave the CE-IEC device, which loss would be undesirable.

[0062] To avoid such electron losses, additional cusps 414 can be provided within the focusing region 106 along the particle paths 108a-108c, as shown in FIG. 4C. Electrons escaping the core 112 that encounter the cusp 414 can be guided to the sidewall 106 by the magnetic field 312, and thus absorbed at a potential that is closer to that of the core 112. The cusps 414 may be repeated along the radial direction—due to randomization of electron velocities between null regions of the magnetic fields, an electron that makes it through one cusp 414 may not necessarily make it through a subsequent cusp 414. As a result, power loss by escaping electrons can be reduced. In addition, these magnetic field lines can help ions from spreading too far in the azimuthal direction and potentially impacting the sidewalls 106. Moreover, the periodic magnetic field (experienced by the ion bunches as they recirculate along the particle paths 108a-108c) can help to compress the ion bunches in the azimuthal direction.

[0063] Such cusped magnetic fields can be generated by incorporating multiple separate magnets in the sidewalls along the radial direction. For example, radially polarized permanent magnets 402A-402C may be incorporated into sidewalls 106, with polar orientations 404A-404C alternating along the radial direction, as shown in FIG. 4A. Alternatively or additionally, the permanent magnets 406A-406C can be conductive panels of the sidewall 106 itself and used to provide both the magnetic and potential profiles, as shown in FIG. 4B. The magnets can optionally be separated from the conductive panels 302 or other magnets 406A by insulating spacers 304 and/or 410. The innermost magnet 402C of the sidewalls 106 can have the same pole facing the core.

[0064] Each magnet would have a polar orientation 404A-404C extending radially (i.e., with one pole closer to the core 112, and the opposite pole spaced at a radially outer

location). The resulting magnetic field lines 312 extend through the channels formed by the sidewalls, but with cusped regions 414 directed toward the sidewalls. Although a particular polar orientation is illustrated in FIGS. 4A-4B, embodiments of the disclosed subject matter are not limited thereto. Indeed, a polar orientation for each magnet opposite that illustrated may be adopted with similar effect. Although an alternating polar orientation is illustrated in FIGS. 4A-4B, in certain embodiments, the polar orientation of each magnet 402A-402C of the sidewall 106 can be the same (i.e., non-alternating, with a north pole facing the south pole of the adjacent magnet).

[0065] Moreover, the features of FIG. 3A-4C are not intended to be mutually exclusive. Rather, in some embodiments, the features of FIGS. 3A, 3B, 4A, and/or 4B can be combined to particular advantage. For example, one or more sidewalls can be formed of a permanent magnet, while other sidewalls can be formed of permanent magnet segments. Other combinations and variations should be readily apparent to one of ordinary skill in the art.

[0066] Although the magnetic field configurations of FIGS. 3A-4C have been separately illustrated from the electric field configurations of FIGS. 2A-2B, embodiments of the disclosed subject matter can employ both techniques, for example, to cooperatively contain electrons within the core. For example, the electrons can be prevented from hitting ends of the sidewalls 106 facing the core 112 by the cusped magnetic field 314 provided by the permanent magnets, despite the field generated by inner anode region 128 of sidewalls 106. Meanwhile, electrons that are able to escape the core and enter the particle paths 108a-108c between sidewalls 106 can be turned back by the electric field generated by the cathode region 124 of the sidewalls 106.

[0067] Despite the provision of a cusped magnetic field, high-angle scattered ions and/or fusion products can impact the edges of the sidewalls 106 facing the core 112. These impacts can cause undesirable heating of the sidewall 106 and its components (e.g., permanent magnets, embedded electrodes, sensors, etc.). Heating of the permanent magnets is especially undesirable as it may lead to de-magnetization. To avoid damaging the sidewalls 106, protective standoff 502 (i.e., shields) can be provided at the innermost edge of the sidewalls 106, as shown in FIG. 5, to absorb impacts from the ions and fusion products. Since energetic particles impacting the standoff 502 will deposit both energy and momentum, which will cause both heating and sputtering, the standoffs 502 can be formed of a material having a very high melting point (e.g., greater than 2000K) and be resistant to sputtering. For example, the standoff 502 can be formed of tungsten or carbon.

[0068] Instead of forming the standoff 502 from the heat-resistant material, the standoff 502 can be coated with a layer of the heat-resistant material. For example, the heat-resistant material could be flowed through a pipe extending along sidewall 106, for example, via expanded channel 604 of FIG. 6 (or temporarily along the particle paths 108a-108c between the sidewalls 106), to coat the radially inner surface of the standoff 502. This material could be allowed to sputter away as a result of particle impacts, which can provide a mechanism for heat removal in addition to protecting the standoff 502 and the sidewalls 106 from erosion.

[0069] Any heat absorbed by the standoff 502 may be passively radiated away or actively cooled by a separate

mechanism (e.g., a heat transfer fluid circulating through the standoff 502). The standoff 502 can also be formed of sufficient thickness (or coated with sufficient thickness) such that any sputtering that does occur would still allow for a sufficient lifetime of operation before failure. For example, if sputtering resulted in loss of approximately 0.013 monolayers/second, a standoff 502 having a 1 cm thick layer of carbon could have a lifetime of several years.

[0070] In embodiments, transparency of the CE-IEC device can be maintained by taking advantage of otherwise unused real estate of the sidewall structures for various functions, such as, but not limited to, electrical connections for voltage or signals, supporting permanent magnets, feeding fuel (e.g., ions) for fusion, and coating standoffs. For example, one or more of the vertices 602 between adjacent sidewalls 106 can be expanded into channel 604 to accommodate an electrical feed line 606, as shown in FIG. 6. The electrical feed line 606 can be coupled to one or more sidewalls 106 at an internal attachment point 608 to provide a bias to the sidewall 106 to generate the radially varying potential field. Outside of the coupling 608, the feed line 606 within the channel 604 can be electrically isolated so as to avoid impacting the potential profile.

[0071] The channel 604 may accommodate a single electrical feed line 606, for example, to set a potential at a single radial location. For example, FIG. 7A illustrates a configuration of the vertex channel 604 where a single feed line 606 is coupled to adjacent sidewalls 106 at internal attachment points 608. The feed line 606 is insulated along its length outside of attachment point 608 by insulating wall 702. Such a configuration may be employed, for example, when the sidewall is formed of a material that has a resistivity that varies along 704, i.e., the radial direction, although the configuration can also be applied to a sidewall formed of multiple segments 252, as illustrated in FIG. 7C.

[0072] Alternatively or additionally, the channel 604 may accommodate multiple electric feed lines 606 to set potentials on different sidewalls, while also having a permanent magnet therein. For example, FIG. 7B illustrates a configuration of the vertex channel 604 where multiple feed lines 606 are disposed around an embedded permanent magnet 306 and separated therefrom by insulating material 702. Each feed line 606 can be coupled to a respective sidewall 106 at a corresponding attachment point 608. Again, the feed line 606 may be insulated along its entire length outside of attachment point 608 by insulating material 702.

[0073] Alternatively or additionally, the channel 604 may accommodate multiple electric feed lines 606, for example, to set potentials at more than one radial location. For example, FIG. 7D illustrates a configuration of the vertex channel 604 where multiple feed lines 606 extend to different radial depths and are separated from each other and the sidewalls 106 by insulating material 702.

[0074] The features of FIGS. 7A-7D are not intended to be mutually exclusive. Rather, in some embodiments, the features of FIGS. 7A-7D can be combined to particular advantage. For example, some vertex channels 604 may have the configuration of FIG. 7C, while other feed channels may have the configuration of FIG. 7D, especially if the number of available vertices may be otherwise limited. Other combinations and variations should be readily apparent to one of ordinary skill in the art.

[0075] The vertex channels 604 could also be used to provide fuel to the CE-IEC device for fusion. For example,

protons and boron ions can be generated at appropriate radial locations within a particle path so that the relative energy matches the fusion cross-section resonance, and so that the center mass of the reaction is stationary at the device core (i.e., zero net momentum). To achieve this, neutral atoms can be fed into the device to an appropriate radius prior to ionization. If the feed tubes were placed within the particle paths, they would be subject to fusion product bombardment. To avoid such bombardment, the feed tubes 804 are placed between the sidewalls 106, for example, at the vertices 602 between adjacent sidewalls, for example, as shown in FIG. 8. Non-ionized fuel 810 is introduced at inlet 802 of the fuel feed tube 804, where it is conveyed down to an appropriate radius before being ionized by ionizer 806. The resulting ions can then be injected transversely to the particle path via outlets 808 extending through sidewalls 106. The injected ions 812 mix with existing ions and join the ion bunches traveling along the particle paths. Injection slightly above the required energy level can allow for energy losses and for the ion to approach the resonance peak from above.

[0076] The vertex channels 604 could also be used to convey signals to/from locations within the sidewalls 106, for example, to convey sensor signals. As discussed above, the potential may be dynamically controlled to compact ion bunches as they travel along the particle paths. In some embodiments, the potential may be controlled without feedback (i.e., open loop), for example, by establishing a time-varying profile and allowing the traveling ion bunches to synchronize to the profile.

[0077] Alternatively, one or more sensors 902, as illustrated in FIG. 9, may be disposed along the sidewalls 106 to monitor the ion bunches as they travel along the particle path. The profile can then be dynamically adjusted in real time to compact the ion bunches or for any other purpose (e.g., to transfer energy to the ion bunch). The sensors 902 may be disposed on a surface of the sidewalls 106 (as shown), within the sidewalls 106, or behind the sidewalls 106 (i.e., within channel 604). Other sensor types and configurations are also possible according to one or more contemplated embodiments, for example, to monitor variable indicative of operation of CE-IEC device, such as temperature, electron population in core, standoff thickness, etc. Signals from the sensors 902 can be communicated via signal wires 904 for subsequent use, for example, by controller 1006. Although shown as extending along vertex channel 604, it is also possible for the signal wires 904 to be disposed between adjacent panels of sidewalls 106 away from the vertex 602.

[0078] Referring to FIG. 10, an overview of a system 1000 including a CE-IEC device 100 is shown. When the CE-IEC device 100 is used in a space environment, the CE-IEC device 100 may be housed in an enclosure 1002 that may be open to the environment (i.e., vacuum). Otherwise, the enclosure 1002 of the CE-IEC device 100 is a chamber that maintains a vacuum environment.

[0079] The CE-IEC device 100 can be coupled to a controller 1006 that controls operation thereof. Such control by the controller 1006 can include providing a static potential and/or a dynamic potential (e.g., to compact ion bunches) to the sidewalls 106 using voltage source 1008. Although shown as a single component, voltage source 1008 can include multiple voltage sources and/or be capable of generating multiple independent voltages (for example, as

needed for the multiple sidewall segments of FIG. 2B). The controller 1006 can also control fuel supply 1004 to supply non-ionized fuels to the CE-IEC device 100, for example, to be ionized in situ as in FIG. 8. The control by the controller 1006 may be responsive to feedback from one or more sensor signals 1010, for example, the sensors of FIG. 9.

[0080] The fusion products 1012 can be directed from the focusing region 110 of the CE-IEC device 100 for subsequent use, for example, directly utilized 1014 (e.g., propulsion of a spacecraft) and/or converted for use 1016 (e.g., converted to electricity using electrostatic deceleration and/or dynamically oscillating potentials). In the latter utility, the kinetic energy of the fusion products can be directly converted into electrical energy. Alternatively or additionally, the fusion products can be used to charge a conducting plate, which resulting charge can be used to drive a high impedance load.

[0081] For example, FIGS. 11A-11C show aspects of a 1-D version of a standing wave direct energy conversion (SW-DEC) process 1100 for converting the fusion products to electricity. The fusion products 140 pass along the particle paths 108a-108c and are energetic enough to escape the focusing region 110 and reach the power conversion region 150. In the power conversion region 150, a plurality of ring-shaped electrodes 1102-1108 are sequentially disposed along a path collinear with one of the particle paths 108a-108c. Alternating ring electrodes are coupled together, such that odd numbered electrodes are in-phase with each other, and the even numbered electrode are 180° out of phase with the odd numbered electrodes.

[0082] As the fusion products 140 approach the first electrode 1102 in FIG. 11A, the first and third electrodes 1102, 1106 are rising toward their peak potential, while the second and fourth electrodes 1104, 1108 are falling toward their minimum potential. The potential 1110 along the axis seen by the fusion products 140 has a positive gradient, which decelerates the fusion products. As the fusion products 140 pass through the first electrode 1102 in FIG. 11B, the potential is approximately zero as the potentials reverse. Once past the first electrode 1102, the first and third electrodes 1102, 1106 are falling toward their minimum potential, while the second and fourth electrodes 1104, 1108 are rising toward their peak potential, as shown in FIG. 11C. Thus, the fusion products continue to experience a positive gradient that further causes deceleration. The process of FIGS. 11A-11C continues for each subsequent ring, extracting additional energy with the passing of each subsequent electrode. Although illustrated as equally spaced in FIGS. 11A-11C, it is also possible to place consecutive rings closer together to compensate for reduced velocity of the fusions products.

[0083] For example, the electrodes 1102-1108 can form the capacitive element of a tuned resistor-inductor-capacitor (RLC) circuit 1120, as shown in FIG. 11D. The first and third electrodes 1102, 1106 can be connected together as one plate of the capacitor while the second and fourth electrodes 1104, 1108 can be connected together as the opposing plate of the capacitor. A resistor 1122 and an inductor 1124 can be connected in series between the plates of the capacitor. As the fusion products 140 lose kinetic energy passing through electrodes 1102-1108, a corresponding amount of energy is pumped into the circuit 1120, thereby increasing its oscil-

lation amplitude. The resistive element 1122 of the circuit 1120 can be a load, which could be operating equipment or an energy storage device.

[0084] The description of FIGS. 11A-11D is for a simplified 1-D SW-DEC configuration. An SW-DEC configuration applied to practical embodiments of the CE-IEC device would necessarily be more complex, with ring or wire mesh electrodes provided at each end of particle paths 108a-108c to capture energy of fusion products emanating therefrom. For example, to implement the SW-DEC configuration in the power conversion region 150 of the CE-IEC, each electrode ring can be at a separate radius from the central core. However, instead of individual rings, each radial layer would be thin walled 3-D honeycomb structure, similar to the construction of the sidewalls forming the particle path channels. Such a configuration would maintain the high transparency of the CE-IEC, while allowing the fusion products to strongly couple to the electrodes. At the outermost radius of the SW-DEC, the fusion products can be driven into an electrode (not shown) where they would be neutralized and allowed to pass out into space. The charging of this electrode due to electron loss could also potentially be used to convert any remaining percentage of the kinetic energy of the fusion products into electricity.

[0085] An exemplary configuration of a 3-D SW-DEC is illustrated in the cross-sectional view of FIG. 11E, where ring electrodes 1102-1108 are disposed along common radial lines in the power conversion region 150. The ring electrodes 1102-1108 may be supported in position with respect to each other by supports 1132, which may also connect the ring electrodes 1102-1108 to the sidewalls 106 of the CE-IEC device 100. The supports 1132 may also provide electrical connectivity between the various ring electrodes and/or other components. Other configurations are also possible according to one or more contemplated embodiments. For example, although four electrodes (rings in FIGS. 11A-11D, cubes in FIG. 11E) have been illustrated, fewer or greater number of electrodes may be provided for the SW-DEC system.

[0086] Although a cubic configuration for the CE-IEC device 100 is illustrated in FIGS. 1-11, this is merely the simplest configuration for explanation of the features of the present disclosure and embodiments are not limited to such geometries. For example, the innermost edge 104 and/or the outermost edge 102 can lie on a sphere, such that each half of the particle path outside the core region 112 takes the form of a truncated cone (defining a circular face perpendicular to the radial direction).

[0087] Indeed, practical embodiments of the disclosed subject matter may employ other configurations with different particle path geometries and/or number of particle paths. To this end, the sidewall structure can be formed by radial extrusion of the edges of any polyhedron, so long as the faces of the polyhedron come in diametrically opposed pairs in order to create the aligned particle paths on opposite sides of the core. For example, the number of particle pathways can be increased and the geometry of the innermost edge 104 and/or outermost edge 102 can be more complex, such as a truncated icosahedron (e.g., soccer ball geometry), as described in further detail below with respect to FIGS. 12A-12B. Accordingly, geometries other than those specifically illustrated are also possible according to one or more contemplated embodiments.

[0088] An entire family of highly symmetric CE options is provided by the geometry of fullerenes—carbon molecules that form closed polyhedral cages. Fullerenes are labeled as C_N , where N is the number of carbon atoms in the molecule. Of particular interest are those of icosahedral ($-I_h$) symmetry, such as C_{20} , C_{60} , C_{80} , C_{240} , etc., where a necessary condition is that N must be a multiple of 20. For each of these, 12 faces are always pentagons and the rest are always hexagons. A regular truncated icosahedron (RTI) has edges of equal length, but the hexagonal faces have an area that is about 60% larger than the pentagons. This is the geometry of the C_{60} molecule (and the soccer ball). However, the truncation of the icosahedron can be done in such a way that instead of ending up with the edges all the same length, one can achieve a geometry where the two types of faces can be inscribed by circles having substantially the same area. This makes the resulting particle paths more equivalent. FIG. 12A shows an exemplary CE-IEC device 1200 employing such geometry, which is referred to as the special irregular truncated icosahedron (SITI). FIG. 12B shows a cutaway version of the CE-IEC device 1200 and illustrates exemplary magnetic field lines and an exemplary electric potential as well as other hidden components (e.g., internal permanent magnets).

[0089] Embodiments of the CE-IEC device can employ various nuclear fusion fuels. For example, the CE-IEC device can employ the deuterium-tritium (D-T) reaction or the deuterium-deuterium (D-D) reaction. For single species fuel, such as D-D, only two diametrically opposed ion bunches 120 will be present along each particle path at any given time, one bunch on each side of the core. While the burning of D-T has the highest cross-section, it may suffer from the production of highly energetic neutrons, which can be absorbed into the nuclei of other materials and create unstable radioactive isotopes. In addition, scattering of these neutrons can dislocate atoms from their lattices resulting in structural degradation over time.

[0090] In another example, the CE-IEC device uses an aneutronic fuel such as $p\text{-}^{11}\text{B}$, which is the fusion of a hydrogen atom with the most common isotope of boron. The result of the fusion process is three helium nuclei (alpha particles) with a total energy of about 8.7 MeV. For two-species fuel such as $p\text{-}^{11}\text{B}$, four diametrically opposed bunches 120—one pair for each of the two species—will be present along each particle path at any given time, two bunches on each side of the core, separated radially.

[0091] Although features of the various figures have been separately illustrated, embodiments of the disclosed subject matter can combine one or more of the separately illustrated features. For example, embodiments can include the side-wall geometry features of FIGS. 1A-1C or 12A-12B, the continuous electrode features of FIGS. 2A-2B, the permanent magnet features of FIGS. 3A-4C, the protective stand-off features of FIG. 5, the electric feed line features of FIGS. 6-7D, the fuel feed features of FIG. 8, the sensor features of FIG. 9, the control features of FIG. 10, and/or the power generation features of FIGS. 11A-11C. Accordingly, embodiments of the disclosed subject matter are not limited to the configurations specifically illustrated.

[0092] Moreover, although the above description has focused on the use of the CE-IEC device for nuclear fusion, embodiments of the disclosed subject matter are not neces-

sarily limited thereto. Indeed, aspects of the disclosed subject matter may be employed in other applications that have traveling ions.

[0093] It will be appreciated that the aspects of the disclosed subject matter can be implemented, fully or partially, in hardware, hardware programmed by software, software instruction stored on a computer readable medium (e.g., a nontransitory computer readable medium) or a combination of the above.

[0094] For example, components of the disclosed subject matter, including components such as a controller, processor, or any other feature, can include, but are not limited to, a personal computer or workstation or other such computing system that includes a processor, microprocessor, microcontroller device, or is composed of control logic including integrated circuits such as, for example, an application specific integrated circuit (ASIC).

[0095] Features discussed herein can be performed on a single or distributed processor (single and/or multi-core), by components distributed across multiple computers or systems, or by components co-located in a single processor or system. For example, aspects of the disclosed subject matter can be implemented via a programmed general purpose computer, an integrated circuit device (e.g., ASIC), a digital signal processor (DSP), an electronic device programmed with microcode (e.g., a microprocessor or microcontroller), a hard-wired electronic or logic circuit, a programmable logic circuit (e.g., programmable logic device (PLD), programmable logic array (PLA), field-programmable gate array (FPGA), programmable array logic (PAL)), software stored on a computer-readable medium or signal, an optical computing device, a networked system of electronic and/or optical devices, a special purpose computing device, a semiconductor or superconductor chip, a quantum computing chip or device, a software module or object stored on a computer-readable medium or signal.

[0096] When implemented in software, functions may be stored on or transmitted over as one or more instructions or code on a computer-readable medium. The steps of a method or algorithm disclosed herein may be embodied in a processor-executable software module, which may reside on a computer-readable medium. Instructions can be compiled from source code instructions provided in accordance with a programming language. The sequence of programmed instructions and data associated therewith can be stored in a computer-readable medium (e.g., a nontransitory computer readable medium), such as a computer memory or storage device, which can be any suitable memory apparatus, such as, but not limited to quantum-based memory, read-only memory (ROM), programmable read-only memory (PROM), electrically erasable programmable read-only memory (EEPROM), random-access memory (RAM), flash memory, disk drive, etc.

[0097] As used herein, computer-readable media includes both computer storage media and communication media, including any medium that facilitates transfer of a computer program from one place to another. Thus, a storage media may be any available media that may be accessed by a computer. By way of example, and not limitation, such computer-readable media may comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, quantum-based storage, or any other medium that may be used to

carry or store desired program code in the form of instructions or data structures and that may be accessed by a computer.

[0098] Also, any connection is properly termed a computer-readable medium. For example, if the software is transmitted from a website, server, or other remote source using a transmission medium (e.g., coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared, radio, and microwave), then the transmission medium is included in the definition of computer-readable medium. Moreover, the operations of a method or algorithm may reside as one of (or any combination of) or a set of codes and/or instructions on a machine readable medium and/or computer-readable medium, which may be incorporated into a computer program product.

[0099] One of ordinary skill in the art will readily appreciate that the above description is not exhaustive, and that aspects of the disclosed subject matter may be implemented other than as specifically disclosed above. Indeed, embodiments of the disclosed subject matter can be implemented in hardware and/or software using any known or later developed systems, structures, devices, and/or software by those of ordinary skill in the applicable art from the functional description provided herein.

[0100] In this application, unless specifically stated otherwise, the use of the singular includes the plural, and the separate use of "or" and "and" includes the other, i.e., "and/or." Furthermore, use of the terms "including" or "having," as well as other forms such as "includes," "included," "has," or "had," are intended to have the same effect as "comprising" and thus should not be understood as limiting.

[0101] Any range described herein will be understood to include the endpoints and all values between the endpoints. Whenever "substantially," "approximately," "essentially," "near," or similar language is used in combination with a specific value, variations up to and including 10% of that value are intended, unless explicitly stated otherwise.

[0102] It is thus apparent that there is provided in accordance with the present disclosure, systems, methods, and devices for inertial electrostatic confinement. Many alternatives, modifications, and variations are enabled by the present disclosure. While specific examples have been shown and described in detail to illustrate the application of the principles of the present invention, it will be understood that the invention may be embodied otherwise without departing from such principles. For example, disclosed features may be combined, rearranged, omitted, etc. to produce additional embodiments, while certain disclosed features may sometimes be used to advantage without a corresponding use of other features. Accordingly, Applicant intends to embrace all such alternatives, modifications, equivalents, and variations that are within the spirit and scope of the present invention.

1. A device comprising:

a central core region;

particle paths radially extending from the central core region, each particle path having a corresponding particle path aligned therewith on an opposite side of the central core region;

sidewalls that extend in a radial direction, each particle path being bounded by a corresponding set of the sidewalls;

electrodes coupled to the sidewalls so as to provide an electric field that varies along each particle path from a

cathode region proximal to the central core region to an anode region remote from the central core region; and a control module that controls the electrodes to provide said electric field,

wherein each sidewall provides a continuous surface radially extending from the cathode region to the anode region.

2. The device of claim 1, wherein at least one of the sidewalls is a continuous planar piece of material from the cathode region to the anode region.

3. The device of claim 1, wherein at least one of the sidewalls is composed of segments separated from each other in the radial direction by insulating spacers.

4. The device of claim 3, wherein the electrodes are coupled to respective ones of the segments such that an electric potential of each segment can be independently controlled.

5. The device of claim 1, wherein the controller is configured to control the electrodes to provide an electric field that accelerates ions along each particle path to cause fusion of said ions within the central core region.

6. The device of claim 1, wherein the controller is configured to control the electrodes to provide time varying electric fields that compact ion bunches traveling along the particle paths.

7. The device of claim 6, further comprising:

sensors that monitor the traveling ion bunches and generate signals responsively thereto, wherein the controller controls the time varying electric fields based on the signals from the sensors.

8. The device of claim 1, wherein electrical connections to the electrodes are routed through insulated conduits at respective intersections of adjacent sidewalls.

9. The device of claim 1,

wherein permanent magnets are disposed between adjacent sidewalls, and along each radial direction, no more than a single magnet is disposed.

10. The device of claim 9, wherein the permanent magnets are arranged to form a cusped magnetic field adjacent to the central core region, with field lines within each particle path following the radial direction.

11. The device of claim 1,

wherein permanent magnets are disposed between adjacent sidewalls, and along each radial direction, more than a single magnet is disposed.

12. The device of claim 11, wherein the permanent magnets are arranged to form a cusped magnetic field adjacent to the central core region, with at least some field lines crossing the radial direction within each particle path.

13. The device of claim 1, wherein at least a portion of each sidewall is formed of a permanent magnet.

14. The device of claim 1,

wherein the cathode region is separated from the central core region by a second anode region, and the second anode region has an electric potential higher than that of the cathode region so as to confine electrons to the central core region while allowing fusion products and ions to escape from the central core region.

15. The device of claim 1,

wherein each sidewall is formed of a material having a resistivity that varies along the radial direction, and

each sidewall acts as an isopotential conductor in an azimuthal direction.

16. The device of claim 1, further comprising shields disposed between the central core region and ends of the sidewalls facing the central core region, the shields being thermally isolated from the sidewalls and protecting the sidewalls from heat and/or impact from particles deviating from the particle paths.

17. The device of claim 1, wherein, when viewed along the respective radial direction, each particle path bounded by the corresponding set of sidewalls has a shape of a polygon with at least three sides.

18. The device of claim 17, wherein areas of said shapes at a same radial distance are substantially equal.

19. The device of claim 1, wherein the central core region is substantially empty except for electrons confined therein, ions traveling therethrough between particle paths, and products resulting from interaction of the traveling ions.

20. The device of claim 1, wherein electrodes are provided at different radial distances for each set of the sidewalls so as to provide different potentials along the corresponding particle path.

21. The device of claim 1, wherein fuel feed paths for injecting ions to the particle paths are provided at respective intersections of adjacent sidewalls.

22. A fusion method comprising:

directing ion bunches along particle paths that radially extend from a central core region, each particle path being bounded by a corresponding set of radially extending sidewalls and having a corresponding particle path aligned therewith on an opposite side of the central core region, each sidewall providing a continuous surface radially extending from a cathode region proximal to the central core region to an anode region remote from the central core region;

generating an electric field that varies along each particle path from the cathode region to the anode region such that the ion bunches are accelerated toward the central core region;

fusing ions from the ion bunches within the central core region; and

allowing fusion products to travel from the central core region to beyond the anode region via said particle paths.

23. The method of claim 22, wherein the generating an electric field comprises varying the electric field with respect to time such that the ion bunches traveling along the particle paths are compacted.

24. The method of claim 23, detecting the ion bunches along the particle paths, wherein the varying of the electric field is responsive to the detecting.

25. The method of claim 22, further comprising confining a population of electrons to the central core region so as to neutralize the ion bunches passing into the central core region.

26. The method of claim 25, wherein the confining comprises generating a cusped magnetic field adjacent to the central core region using a plurality of permanent magnets as said radially extending sidewalls or between adjacent ones of the radially extending sidewalls.

27. The method of claim 26,

wherein more than one permanent magnet is disposed along the radial direction and the poles of the disposed magnets alternate along the corresponding radial direction, and

further comprising using the magnetic fields of the permanent magnets to direct electrons escaping the central core region to the sidewalls.

28. The method of claim 25, wherein the confining comprises controlling the electric field such that a region having a higher potential than that of the cathode region is formed between the central core region and the cathode region along the radial direction.

29. The method of claim 22, further comprising, protecting ends of the sidewalls facing the central core region from heat and/or particle impact using a plurality of shields that are thermally isolated from the sidewalls.

30. The method of claim 22, wherein the central core region is substantially empty except for electrons confined therein, ions traveling therethrough between particle paths, and products resulting from interaction of the traveling ions.

31. The method of claim 22, further comprising, directing the fusion products out of a spacecraft so as to propel the spacecraft.

* * * * *

Appendix B:

AAS 18-037: Continuous Electrode Inertial Electrostatic Confinement Fusion

Raymond Sedwick

Andrew Chap

Nathan Schilling

Paper presented at the American Astronautical Society Guidance & Control Conference
Feb 1-7, 2018, Breckenridge, CO

CONTINUOUS ELECTRODE INERTIAL ELECTROSTATIC CONFINEMENT FUSION

R. J. Sedwick,* A. M. Chap and N. M. Schilling

One of the greatest impediments to space exploration is a lack of abundant power, in particular as missions extend farther from the sun. Nuclear fusion, while not technically a renewable like solar photovoltaics, offers such a high energy density that the distinction becomes meaningless over foreseeable mission lifetimes. An ideal implementation of fusion power for space would be aneutronic, removing mass requirements for reactor shielding and eliminating material activation and damage. It would ideally also leverage direct energy conversion, eliminating the need for massive radiators required to support thermodynamic power conversion. Continuous Electrode Inertial Electrostatic Confinement Fusion is a concept currently under development that may lend itself to utilizing the proton-boron reaction with power conversion provided by a standing wave direct energy conversion concept. This paper presents an overview of the technology and provides a high-level, top-down system design for a 1 MW reactor.

INTRODUCTION

Inertial Electrostatic Confinement (IEC) as a concept for fusion power generation has been around since the mid-1960's. Since that time, the concept has been researched and refined as a fringe approach to fusion power, with the general consensus being that fundamental issues such as ion thermalization will remain insurmountable to ever achieving – let alone exceeding – breakeven operation¹. A particular development path, ongoing since 2000, has sought to address this issue through advancement of the ion confinement mechanism. A key development was the realization that the single cathode accelerating mechanism is naturally defocusing to the ion stream, causing a mean ion lifetime of only a dozen or so passes through the system before impacting the cathode. Since the mean free path for fusion is on the order of a million passes, the yield for such devices is quite low². It was also seen that the electrode feed stalks created large field asymmetries that contributed to rapid ion loss to the electrodes.

The introduction of additional electrodes at different radial locations was found to stabilize the ion paths, even against feed stalk asymmetries, increasing their lifetime in the device 1000-fold. As the ion lifetimes grew, a coupling between the two-stream instability and the natural trap kinematics was shown to cause the ions to form long-lived diametrically opposed bunches along the channels formed by the electrode grid openings. These bunches became synchronized among the channels, all reaching the center of the device simultaneously³. It was theorized that by limiting the counter-streaming ion bunch interaction to the core of the device that the low-angle collisional processes that lead to rapid ion thermalization could be kept in check.

* Department of Aerospace Engineering, University of Maryland, College Park, Maryland 20740, USA.

For such a device to work, the ions would need to be well localized, both radially and azimuthally, perhaps using active control along the channels. They would also need to be shielded against their own space charge while in the core. A means of providing highly controllable electric and magnetic fields along the channels and confining an electron population in the core would be required, while still offering a highly transparent and unobstructed path for fusion products to escape so that direct energy conversion can take place. Thus, was born the concept of the continuous electrode (CE-) IEC.

CONCEPT OVERVIEW

The notional configuration of the CE-IEC can be divided into four main radial regions, as shown in Figure 1. In radially increasing order, these are the 1) inner core, 2) outer core, 3) focusing region, and 4) power conversion region. The focusing region is where the continuous electrode structure resides. The bulk electrode structure can be thought of as a radial extrusion of the edges of a polyhedron. The cavities formed where the polyhedral faces would normally reside then become open channels through which the ion bunches recirculate as they pass in and out of the core. This region is where the ion packets are continuously refocused and compressed to maintain them in a compact form against the natural spreading that results from their own space charge and low angle collision events. Although the ions are in localized bunches, the paths along the centers of these cavities from one end of the device to the other will often be referred to as beamlines. For a single species fuel (D-D for example) only two diametrically opposed ion bunches will be present along each beamline at any given time, so only one bunch per channel. For a two-species fuel (p-¹¹B for example), four diametrically opposed packets will be present – one pair for each of the two species – and therefore there will be two bunches per channel, separated radially.




Figure 1. Four distinct radial regions of the device

The recirculating ion bunches start near the outermost edge of the CE and fall through the potential drop from the focusing region radially inward, first into the outer core. The outer core is the transition region where the ion bunches pass initially into and then out of the confined electron population on the opposite side. As they pass into the outer core, the electron population responds by being attracted to the ion bunches to neutralize their space charge. This allows the ions to then further compress along their primarily radial trajectories as they fall into the inner core. The inner core is then the central region within which the ion packets interact collisionally. The size of this region is ideally small, and is determined by the bunch sizes and trajectories. The collisional processes within the inner core include energy and momentum exchange, and of course fusion.

Slowing/Halting Thermalization

Low-angle collisions among opposing bunches, which would normally result in azimuthal thermalization, instead simply reshuffle the specific radial trajectories along the beamline that the individual ions will follow as they exit the core. Upon refocusing, the azimuthal velocity distribution within each bunch should be indistinguishable from the start of the previous pass, cancelling the azimuthal momentum growth.

Likewise, low angle scatters among the non-opposing packets can introduce both azimuthal and radial (energy) scattering. However, on average these scattering events will both up-scatter and down-scatter the ions equally, so again upon refocusing, the net ion energy within each packet due to low angle ion-ion collisions should remain the same.

Ion Particle and Energy Loss Mechanisms

A separate effect is the drag experienced by the ions due to the electrons, transferring energy out of the ion packets. This represents a net energy loss, however continuous injection of new energetic ions into the packets may be sufficient to address this loss. This balance is to be investigated.

High angle scatters present the possibility of scattering the fuel ions into the inner edge of the continuous electrode structure, therefore it is paramount to construct the electrode to be as transparent as possible as viewed from the center of the system. If the scattering angle of the ion does not cause a collision with the electrode, then the ion should simply end up within a different channel, where the refocusing process will merge the ion into the local packet. If the electrode transparency as seen from the core were 85%, a close encounter between fuel ions will only result in a possible wall collision 15% of the time. This allows more opportunity for fusion to occur.

Similarly, high transparency is necessary to allow the fusion products to escape the core through the focusing region to reach the power conversion region. An 85% transparent electrode will subtend 15% of the fusion product paths, meaning that 15% of the fusion power output could end up as heat on the electrode. This will ultimately establish the operational limits on the system and will factor into system sizing.

Energy Conversion

The outer edge of the focusing region represents the maximum radial extent to which the fusion fuel ion bunches must reach, with only fusion products being energetic enough to extend out into the power conversion region. The fusion products enter this outermost region with a nearly isotropic angular distribution, arriving in pulses a few nanoseconds long separated by several microseconds. Such a pulsed output may be ideal for using a technology such as Traveling or Standing Wave Direct Energy Conversion (T/SWDEC)⁴.

TOP-DOWN REFERENCE DESIGN

The following presents a first cut at a self-consistent set of design and operating parameters for a full power generation system based on the continuous electrode design. The choices of parameters are not meant to be optimal, but allow for establishing a baseline from which a better point design can evolve. Detailed modeling, analysis and optimization of each of the subsystems is ongoing and will be discussed in future publications.

Continuous Electrode Geometry

As mentioned, the continuous electrode can be thought of as a radial extrusion of the edges of a polyhedron. The main requirement for selection of the polyhedron is that its faces come in diametrically opposed pairs, creating the opposing channels of a beamline. So, while something as simple as a cube could work, a pyramid would not. If the walls are assumed to be thin, then the total amount of volume that all of the ion bunches have to expand into is generally the same

regardless of the total number channels (faces). The thing that more channels do provide is a greater level of control over the ion bunches. The ions divided among smaller volume bunches will be closer to the channel walls on average, and therefore more affected by the electric and magnetic fields that they generate.

However, as the number of channels grows, so does the number of edges, and at the inner electrode edge, a minimum achievable wall thickness will mean an ever-decreasing system transparency. For the remainder of the analysis it will be assumed for reference that the inner radius of the continuous electrode is at $r_1 = 0.5 \text{ m}$, the outer radius is at $r_2 = 1.0 \text{ m}$, and the wall is uniformly $t_w = 1.0 \text{ cm}$ thick.

An entire family of highly symmetric electrode options is provided by the geometry of the fullerenes – carbon molecules that form closed polyhedral cages⁵. Fullerenes are labeled as C_N , where N is the number of carbon atoms in the molecule. Of particular interest are those of icosahedral ($-I_h$) symmetry, such as C_{20} , C_{60} , C_{80} , C_{240} , etc., where a necessary condition is that N must be a multiple of 20. For each of these, 12 faces are always pentagons and the rest are always hexagons.

A regular truncated icosahedron (RTI) has edges of equal length, but the hexagonal faces have an area that is about 60% larger than the pentagons. This is the geometry of the C_{60} molecule, also called Buckminster Fullerene. However, the truncation of the icosahedron can be done in such a way that instead of ending up with the edges all the same length, one can achieve a geometry where instead the two types of faces can be inscribed by circles having the same area. This makes the resulting channels more equivalent. An image of this geometry is shown in Figure 2, where it can be seen that the hexagons are now only 3-fold symmetric. This geometry will be referred to as the special irregular truncated icosahedron (SITI) and will be used as the baseline for design purposes.

Transparency, Volume and Mass

Adequate geometric relationships are achieved by considering an RTI and approximating the total surface area of the polyhedron as being equivalent to the total curved surface area of a circumscribed sphere. This surface, at a given radius (r), is equally divided among F flat faces, 12 of which are pentagonal ($A_F = 1.7 l_E^2$), and $F-12$ of which are hexagonal ($A_F = 2.6 l_E^2$), with A_F the area of the face, and l_E the length of an edge. The total edge area can be mapped by attributing the nearest half of the edges surrounding each face to that face, and then summing over all faces. The transparency is then approximated by taking the ratio of the total edge area to the spherical surface area, and subtracting from unity, as given by Equation (2)

$$T \approx 1 - \frac{A_E}{A_S} = 1 - \frac{3l_E t_w (F-2)}{4\pi r_1^2} \approx 1 - 1.3 \frac{t_w}{r_1} \sqrt{F} \quad (1)$$

For the RTI (and approximately for the SITI) the transparency using the baseline design values is ~85%, consistent with the example value assumed above. For comparison, the C_{240} geometry with the same dimensions would have an inner edge transparency of only ~71%.

However, while Equation (1) is appropriate for the recirculating fuel ions, this actually represents an overestimate for the fusion products because it assumes that all of the fusion takes place at a single point at the center of the core. As the spatial extent of the fusion volume (inner core) grows, some fusion products will have trajectories that can hit the broad side of the wall surfaces with grazing incidence, reducing the effective system transparency. As the details depend on better knowledge of the achievable core volume compression, this effect will not be addressed in the current analysis.




Figure 2. Continuous electrode from an SITI

The volume, and therefore the mass of material making up the electrode can be estimated by treating each wall segment as the difference between the sector of a circle with outer radius r_2 and inner radius r_1 . Summing over all E edges of the polyhedron, the material volume can then be estimated as

$$V = 3.3 t_w \sqrt{F} (r_2^2 - r_1^2) \quad (2)$$

and for the assumed parameters, this gives a total electrode volume of 0.14 m^3 . The structure of the electrode will be complex, with radial segmentation, embedded feedthroughs, permanent magnets and internal structure to maintain the potential profiles. As a gross over-estimate, if the electrode were to be constructed out of solid neodymium (to provide the magnetic field – discussed later), the present electrode would have a mass of $\sim 1000 \text{ kg}$ (1 metric ton). Assuming that this represents a large fraction of the total system mass (say $\sim 50\%$), then to achieve a target system specific mass of 2 kg/kW , the power output of the system must be at least 1 MW .

Fusion Fuel Parameters

As a conservative estimate, an overall power conversion efficiency of 75% will be assumed for now. Given the large energy density of fusion fuels, the impact of the conversion efficiency is less concerned with the efficient use of the fuel, and more about disposition of the wasted power. This will be addressed later in the paper. The fusion fuel that will be assumed for the baseline design is $\text{p-}^{11}\text{B}$ (see Figure 3.). Advanced fuels are notoriously challenging to burn as a result of lower fusion cross sections and higher energy requirements, however the aneutronic nature of $\text{p-}^{11}\text{B}$, make it an appealing candidate for use in space. One issue with $\text{p-}^{11}\text{B}$ is that to burn the fuel in a thermal plasma near its peak fusion cross section ($1.2 \text{ barns @ } 550 \text{ keV}$) actually produces greater energy loss due to Bremsstrahlung within the electron population than what is generated by the fusion process¹. It is expected that the temperature of the neutralizing electron population in the current system will be such that this loss is not significant.

Ion Energy

Even targeting the peak cross-section, it may be possible that the average electron temperature in the core can be maintained below this level in the current system, since as will be shown the ions only pass through the electron core with a low duty cycle. However, for now the baseline system will avoid the issue altogether by targeting the narrow resonance in the $p\text{-}^{11}\text{B}$ fusion cross section that is located very near to 148 keV. This resonance offers a fusion cross section of 0.1 barns (10^{-25} cm^2), which is only 8% of the $p\text{-}^{11}\text{B}$ fusion peak (1.2 barns) and about 2% of the D-T fusion peak (5 barns @ 64 keV)⁶. However, it also reduces the cost of losing an ion by 400 keV per ion. The reason that leveraging this resonance is a possibility is that the ion focusing may be able to maintain a somewhat narrow spread in energy that can be placed on top of the resonance. Although the required ion densities will be higher, other system losses are substantially reduced.

Duty Cycle and Power Output

The radial potential profile across the device will not be parabolic, and in fact it is desirable from an ion bunch stability standpoint to have the period of oscillation increase with ion energy⁷. It is also necessary that the proton and boron bunches oscillate at the same frequency, requiring a non-parabolic profile. Nevertheless, for the purpose of estimating the period and duty cycle of the fusion output, a parabolic profile with a drift region across the core will be assumed. To achieve the necessary 148 keV center of mass energy with zero net momentum in the device frame, the speeds of the proton and boron bunches as they pass through the core of the device must be $5.1(10^6) \text{ m/s}$ and $4.6(10^5) \text{ m/s}$ respectively. The peak potential to confine the protons is then 136 kV. For the boron ions, the confining potential when they are fully ionized is only 2.4 kV, but when they are first introduced as singly ionized particles they must be injected at 12 kV. Injecting the boron atoms so low into the potential well is one of the many challenges.

Since the potential profile is generated solely in the focusing region, the ions will drift through the core at constant speed, which will add time onto the oscillation period. Under a parabolic profile with drift, the transit time of the protons across the device is $0.7 \mu\text{sec}$, and that of the boron atoms is $4.8 \mu\text{sec}$, demonstrating that the profile will need to be modified to match the transit times across the system. Using the boron period as the limiter, the fusion pulses would occur at just over 208 kHz.

The fusion occurs only within the inner core, where the fuel atoms collide. Ideally, the inner core volume would be highly focused, confined to a region that is perhaps no more than a few centimeters in diameter. A more conservative value might be an inner core diameter of 10 cm in the current example, exhibiting a 10:1 ion compression (in bunch radius) from the outer to the inner core. At their peak velocity, the proton bunches will pass through this inner volume in only 2 nsec, so that the fusion output occurs at a duty cycle of only $DC = 2 \text{ nsec} / 4.8 \mu\text{sec} = 4.2(10^{-4})$. This is a possible disadvantage of the current approach, as the fusion output during the pulse must be 2600




Figure 3. Fusion Cross-sections for various fuels

times greater to achieve the same average power output as if it occurred continuously. In general, the average power output is given by

$$P \approx \eta \frac{16\pi}{3\tau} e E_F n_p n_B \sigma_F r_{ic}^4 \quad (3)$$

where τ is the transit time of the fuel ions and the rest is the energy output per pulse. At constant inner core density, it can be seen that large gains can be made by increasing the inner core radius (less focusing, but larger fusion volume). However, this also means more recirculating current through the device and a lower transparency to fusion products. Resolving this trade-off will be the subject of future analyses.

Core Density and Recirculating Current

For the assumed performance parameters, the density of both proton and boron atoms in the inner core must be $n_p = n_B = 6.32(10^{22}) m^{-3}$, assuming $\eta = 75\%$ and $E_F = 8.7 MeV$. As they expand through the outer core, the ions separate into 32 individual proton bunches and 32 individual boron bunches, all moving toward their respective channels and spreading out, at least azimuthally but possibly also radially. As each bunch passes out of the outer core, the bunch density relative to the inner core is therefore reduced by a factor of up to 64,000 (recall the 10x radial compression) to a value of $n_p = n_B = 9.9(10^{17}) m^{-3}$, with the protons continuing to expand to the outer edge of the focusing region, where the density is again reduced by up to a factor of eight. The proton current through a given channel as the bunch leaves the outer core would then be 79 kA, and the boron current at the same location (1 μ sec later) would be 36 kA.




Figure 4. Cut-away of CE showing B-field lines, magnets embedded in walls and potential profile.

Core Electron Population and Confinement

When the ions are outside of the core, the electrons will spread out due to their own space charge. The electron population that is needed to fully neutralize the ions passing through the core is equal to one electron for each proton and five for each boron, assuming full ionization. When the electrons spread out, they will occupy a volume that is 1000 times larger (0.5 m outer core diameter), and if uniform the electron density is then $n_e = 3.8(10^{20}) m^{-3}$. The primary means of electron confinement in the core is by a cusped magnetic field, although a reversed electric field at the inner edge of the focusing region might also be established. The effectiveness of the electron confinement will depend on the equilibrium temperature of the electrons.

The magnetic field is achieved through the placement of permanent rare Earth magnets within the continuous electrode walls. The configuration will have either all North or all South poles

aligned with the inner edge of the electrode. The field lines must then thread up through the channel openings to connect with the opposite magnetic poles, as shown in Figure 4. This will create a cusped field on the inner edge of the walls, and down the middle of each channel. Recent work has shown that electron confinement of this type in Polywell™ devices improves with increasing β , the ratio of plasma pressure to magnetic pressure⁸. Surface currents in the electron population resulting from steep pressure gradients compress the magnetic field and constrict the electron escape paths. At $\beta = 1$, the electrons are held in check by the full magnetic pressure, provided such a condition can be achieved.

To estimate the field strength in the channels, consider that the entire magnetic flux leaving the inner edge of the electrode must pass through the channel opening at the inner radius and continue up the channel. Assuming a uniform field across the edge of the wall, and likewise into the channel opening, flux conservation requires that the ratio of the magnetic field strength through the channel to that on the magnet surface must be inversely proportional to the respective flux areas. This ratio is known from the transparency. At the assumed 85% transparency, the ratio of the strength of the field in the channel to that at the face of the electrode is then 15/85, or about 18% of the field strength at the wall edge. The surface field strength for an N52 grade neodymium magnet can be as high as 7400 G, which would mean a 1300 G peak field in the channel⁹.

The maximum electron temperature that can be confined solely using a magnetic field is found by equating the electron pressure with the magnetic field pressure

$$T_{ev} = \frac{p}{n_e e} = \frac{\beta B^2}{2\mu_0 n_e e} \quad (4)$$

which for $\beta = 1$ and the parameters given above will support an electron temperature of only 110 eV. While it is expected that the electron temperature will be less than the ion energies, this value seems problematic. However, it is also possible to raise the potential within the core relative to ground to provide additional confinement. Assuming the electrons to be Maxwellian distributed, the fraction of the particles that can escape a well of depth ϕ is given by

$$\frac{n(\phi)}{n_0} = \operatorname{erfc}(\bar{x}) + \frac{2}{\sqrt{\pi}} \bar{x} e^{-\bar{x}^2} \quad (5)$$

where $\bar{x} = \sqrt{\frac{\phi}{T_{ev}}}$. As the well potential increases relative to the electron temperature (increasing $\bar{x}$), a smaller fraction of the particle density can escape the well. If it is assumed that the electron temperature is on the order of 1 keV, the trap potential that allows Equation 4 to be satisfied is 6 kV, whereas if the electron temperature reaches 5 keV, the trap potential must be 24 kV.

Even with the pressure balance satisfied, electrons will still be lost through the cusps. Electrons that approach the cusp with a sufficiently low pitch angle (more parallel to the field line) will not be turned around before the point of maximum field strength. Therefore, some amount of electron leakage current will exist. Figure 4 showed the magnetic field lines extending from the innermost edge of the CE to the outermost edge, however, a better configuration would be to have the outer termination point of the field lines intersect the walls farther down in the potential well. It is then possible that electrons passing through the cusp will be guided to the wall by the magnetic field and be absorbed at a potential that is closer to that of the core, reducing the power loss by escaping electrons. This would be far preferable to the electrons being accelerated to the top of the CE, where each electron lost would represent a cost of 148 keV.

Additional magnetic field loops can be placed farther up the wall by alternating the North and South poles of the magnets. In the null region between each magnetic field loop, the electron velocities will tend to randomize, so making it through the first cusp will not guarantee that an electron would make it through consecutive ones, further increasing the chances that the electron paths will terminate on the wall prior to reaching the anode potential.

Direct Energy Conversion

Fusion products (alphas) that end up passing through the open channels will be energetic enough to escape the confining potential of the focusing region and reach the energy conversion region. Here their kinetic energies will be converted via a radio frequency direct energy conversion (DEC) system. One such system is the Standing Wave DEC (SWDEC)¹⁰. The operation of the SWDEC is illustrated frame-by-frame for a simplified 1-D implementation in Figure 5. In 1-D, consider a series of ring-shaped electrodes, through which passes a localized bunch of ions. The ring electrodes are biased such that all odd numbered electrodes share the same phase, and all even numbered electrodes are 180 degrees out of phase. In frame #1, the ion bunch approaches the first ring, which is rising toward its peak potential. The bunch therefore sees a positive potential gradient and is decelerated. As the bunch passes through the center of the electrode, the phase of the electrode biases reverse, so again the ion bunch sees a positive gradient and continues to slow. By placing consecutive rings closer together to compensate for the reduced bunch velocity, this mechanism continues until most of the energy has been removed.




Figure 5. Frame-by-frame illustration of SWDEC deceleration mechanism using ring electrodes¹⁰.

If the electrodes are part of the capacitive element of a tuned RLC circuit, the ion bunches act as a driving force on the circuit, increasing the oscillation amplitude by pumping energy into it. This energy corresponds to the loss of kinetic energy of the bunch. The resistive element of the circuit is the load, which could be operating equipment or an energy storage device. An analysis of such a device in Ref. 10 demonstrated conversion efficiencies in excess of 90%. For implementation in the outer region of the CE-IEC, each ring in Figure 5 would be at a specific radius. However, instead of individual rings, each radial layer would consist of a thin walled honeycomb, offering a very high transparency, but allowing the ions to strongly couple to the electrodes. At the outermost radius of the SWDEC, the alphas would be driven into an electrode where they would be neutralized and allowed to pass out into space. The charging of this electrode due to electron loss could also potentially be used to convert the last remaining 10% of the ion kinetic energy into electricity.

One challenge with the SWDEC implementation is that ideally all fusion products would be produced at the same energy, arriving at the same time and with the same velocity. In addition to

some random variation, it is well-known that the three alpha particles of the $p\text{-}^{11}\text{B}$ reaction are not formed simultaneously with the same energy. As far back as 1936 it was determined that two of the (secondary) alpha particles are formed with equal energies ranging from 3.8 MeV to 4.4 MeV each from the decay of an excited state of ^8Be , with the primary alpha particle receiving typically less than 1.0 MeV¹¹. While it would be nice to extract the energy from all three alpha particles, the two higher energy alpha particles typically contain at least 88% of the available energy of the reaction. Of greater importance is designing the SWDEC so that it can convert as much of the secondary alpha energy as possible, despite the 0.6 MeV spectrum width. In addition, it is desirable that the primary alphas are able to leave the system, rather than being captured in the fuel recirculation region.

Power Dissipation and Sputtering

The innermost edge of the CE will see heating due to high angle scattering of fuel ions, as well as intercepting a fraction of the fusion products. Heating of the permanent magnets is undesirable, as it may lead to de-magnetization. To mitigate this, a standoff will be constructed conformally to the inner edge of the CE, separated from it by an insulating layer. The material of this stand-off would ideally have a very high melting point and be resistant to sputtering. These requirements are related, since energetic particles impacting this surface will deposit both energy and momentum, which will cause heating and sputtering.

Regarding the heating, the majority of the power deposited into the material will come from the assumed 15% of the fusion products that it intercepts on their way out of the core. At 1 MW output, 1.14 MW of ion power must reach the SWDEC, if it is assumed that only 88% of the power (100% of secondary alphas only) is extracted. This 1.14 MW must then represent the 85% of ions that pass through the channels, leaving 200 kW (15%) of power deposited onto the standoff. The area of the standoff is 0.47 m², so to radiate this much power (assuming an emissivity of 0.8) requires a surface temperature of 1750 K. For reference, the melting temperatures of tungsten and carbon are 3700 K and 4500 K respectively, so there are materials that can likely withstand the temperatures. The radiated energy (peaking near 1 μm) will experience some absorption into other surfaces as it leaves the system, which is an effect that will have to be evaluated.

Sputtering is a potential problem due to the energetic alphas, carrying energies up to 4.4 MeV. Sputtering data for alpha particles at these energies are fairly elusive, but to get an idea of expected sputtering levels, a comparison will be made to work conducted at Sandia National Labs in 2005¹². In this work, the focus was on the evolution of surface morphology of diamond by gallium ions with 20 KeV incident energy. The comparison is admittedly apples to oranges (the gallium atoms are singly charged initially, but can potentially reach a fully charged state of +31) but at a molecular weight of nearly 70, the momentum carried by 20 KeV Ga ions is nearly 40% of the momentum carried by 4.4 MeV alphas. The normal incidence sputtering due to Ga^+ was found to be ~ 2.5 atoms per incident ion, so it will be assumed that the alphas will sputter material at roughly 5 atoms per incident ion. The flux of alphas to the surface ($\sim 400 \text{ kW/m}^2$) is then $\sim 5.6(10^{17})$ particles/m²/sec. The surface density of carbon atoms in a monolayer (assuming 154 pm bond length) is $4.2(10^{19})$ particles/m², so the depth erosion rate is estimated to be 0.013 monolayers/second, or 63 $\mu\text{m}/\text{year}$. Even at 10 times this erosion rate, a 1 cm thick standoff could last for several years before needing replaced.

CONCLUSIONS

This paper has presented an introduction to the concept of the Continuous Electrode Inertial Electrostatic Confinement approach to fusion power generation in space. The top-down design parameters were presented assuming a nominal power output of 1 MW. Sizing of the system was arbitrary to establish a baseline, and current efforts are underway to determine the minimum system size that is required to produce a given power output. Such an analysis will then lead to an understanding of the specific mass (kg/kW) of the device as it scales to different power levels, with a desired target of less than 2 kg/kW. Successful implementation of the concept will rely on comprehensive modeling of the electron confinement, in particular as the ion bunches traverse in and out of the core region, as well as the evolution of the ion population and the effectiveness of active control. Both of these efforts are also currently underway. There are many other implementation details of the technology, such as injection of the fuel into the potential well, many of which have preliminary solutions that are yet to be fleshed out.

ACKNOWLEDGMENTS

This work has been funded in part by the NASA Space Technology Research Fellowship (NSTRF) Program (Chap – Grant #NNX13AL44H), the University Maryland Department of Aerospace Engineering AEROS Scholarship Program (Schilling), and most recently the NASA Innovative Advanced Concepts (NIAC) Program (Grant #NNX17AJ72G). The authors wish to acknowledge collaboration with the NASA Johnson Spaceflight Center Propulsion and Power Division on the topic of RF Direct Energy Conversion.

NOTATION

A_F	=	Area of the face of a polyhedron (m^2)
A_S	=	Surface Area of a circumscribed sphere (m^2)
E	=	Number of polyhedral edges (-)
E_F	=	Energy release during a fusion event (eV)
e	=	fundamental unit of charge (C)
F	=	Number of polyhedral faces (-)
l_E	=	Length of a polyhedral edge (m)
N	=	Number of polyhedral nodes (-)
n_P	=	Density of protons during fusion (m^{-3})
n_B	=	Density of boron atoms during fusion (m^{-3})
n_e	=	Density of electrons in the core (m^{-3})
P	=	Fusion power generated (W)
$\mathcal{P}$	=	Period between fusion events (s)
r_1	=	Inner (core) radius of electrode (m)
r_2	=	Outer radius of electrode (m)
r_{ic}	=	Radius of the inner core (m)
T_{eV}	=	Electron temperature (eV)
t_w	=	Wall thickness of electrode (m)
ϕ	=	Raised core trap potential (V)
σ_F	=	Fusion cross section (m^2)

1. References

- ¹ Rider, T.H., "Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium", *Physics of Plasmas* **4**, 1039 (1997); doi: <http://dx.doi.org/10.1063/1.872556>
- ² McGuire, T.J. and R.J. Sedwick, "Improved Confinement in Inertial Electrostatic Confinement for Fusion Space Power Reactors", *AIAA J. of Propulsion and Power*, 2005, Vol.21: 697-706, 10.2514/1.8554
- ³ Dietrich, C., Eurice, L. and R.J. Sedwick, "Experimental Verification of Enhanced Confinement in a Multi-Grid IEC Device", 44th AIAA/ASME/SAE/ASEE JPC, 2008.
- ⁴ Chap, A.M. and R.J. Sedwick, "One-dimensional semi-analytical model for optimizing the standing-wave direct energy converter", *J. of Propulsion and Power*, September, Vol. 31, No. 5, pp. 1350-1361, 2015.
- ⁵ Schwerdtger, P., L.N. Wirz and J. Avery, "The topology of fullerenes", *WIREs Comput Mol Sci* 2015, 5:96-145. doi: 10.1002/wcms.1207
- ⁶ <http://www.fisicanucleare.it/documents/0-19-856264-0.pdf>
- ⁷ Zajfman, D., O. Heber, M.L. Rappaport, H.B. Pedersen, D. Strasser, and S. Goldberg, "Self-bunching Effect in an Ion-trap Resonator," *J. Opt. Soc. Am. B*, **20** (5), pp. 1028-1032 (May 2005).
- ⁸ Park, J., N.A. Krall, P.E. Sieck, D.T. Offermann, M. Skillicorn, A. Sanchez, K. Davis, E. Alderson and G. Lapenta, "High-Energy Electron Confinement in a Magnetic Cusp Configuration", *Phys. Rev. X*, **5**(2), 021024 (2015); doi: 10.1103/PhysRevX.5.021024
- ⁹ <https://www.kjmagnetics.com/calculator.asp?calcType=block>
- ¹⁰ Chap, A.M. and R.J. Sedwick, "One-Dimensional Semi-analytical Model for Optimizing the Standing-Wave Direct Energy Converter", *AIAA Journal of Propulsion and Power*, Vol. 31, No. 5 (2015), pp. 1350-1361. doi: 10.2514/1.855439.
- ¹¹ P.I. Dee, C.W. Gilbert, *Proc. R. Soc. Lond. A* **154**, 279 (1936)
- ¹² Mayer, T.M., D.P. Adams, M.J. Vasile, and K.M Archuleta, "Morphology evolution on diamond surfaces during ion sputtering", *J. Vac. Sci. Technol. A: Vacuum, Surfaces, and Films* **23**(6), 1579 (2005) doi:10.1116/1.2110386

Appendix C:

Electron Confinement in an IEC Fusion Reactor Core Using Permanent Magnets

Nathan Schilling
Raymond Sedwick

Paper Presented at the 2018 Region I Student Conference
April 13-14, 2018, Clarkson University, Potsdam, NY

Electron Confinement in an IEC Fusion Reactor Core Using Permanent Magnets

Nathan M. Schilling,¹ and Raymond J. Sedwick.²

University of Maryland, College Park, Maryland, 20742, United States of America

In order for safe, quick, and routine manned space exploration to interplanetary destinations within the solar system, a high-thrust, high-specific impulse propulsion system is needed. The precondition for any such system is high-specific power, potentially achievable through Inertial Electrostatic Confinement Fusion (IECF). The primary issue facing such systems currently is that they are not net-power positive. Previous research has suggested that adding permanent magnets along the beamline of such devices may improve confinement and produce a net-power positive fusor. This project attempts to determine the optimal configuration of permanent magnets through building several models of the IEC core; the first model was built in COMSOL, and the next three in MATLAB. The COMSOL model, while interesting, necessitated the development of the MATLAB models. All MATLAB models are 2D and 1D electron fluid models. Research with the 2D model uncovered results regarding the limitations of electron fluid models to resolve interactions in the IEC core. The process of validating a similar 1D model has demonstrated the unexpected complexity in developing a fluid electron model with self-consistent magnetic field coupling.

I. Nomenclature

$\vec{A}$	= Magnetic vector potential, T*m
$\vec{B}$	= Magnetic field, T
$\vec{\Gamma}$	= Electron flux, s ⁻¹ m ⁻²
i	= Grid Point, x-dimension
j	= Grid Point, y-dimension
k	= Iteration number
n_x	= Number of nodes in the x-direction
n_y	= Number of nodes in the y-direction
IEC	= Inertial Electrostatic Confinement
IECF	= Inertial Electrostatic Confinement Fusor
p	= Pressure, N
$\vec{E}$	= Electric field, V/m
ϕ	= Electric Potential, V
T_{ev}	= Temperature, eV
n	= Particle density, m ⁻² or m ⁻¹
R	= Simulation length, m
λ_0	= Debye Length at $x=0$, m
σ	= Dimensionless parameter that determines the steepness of the increase of the Gaussian Error function
u_0	= Dimensionless parameter that determines at which x location along the domain the increase of the Gaussian Error function occurs
Δx	= Distance between nodes in x -axis, m

Fundamental Constants

c	= Speed of light in a vacuum, 3.0×10^8 m/s
e	= Fundamental charge, 1.6×10^{-19} C

¹ Undergraduate Researcher, Department of Aerospace Engineering, 3181 Glen L. Martin Hall, Student Member.

² Associate Professor, Department of Aerospace Engineering, 3181 Glen L. Martin Hall, Associate Fellow.

ϵ_0 = Permittivity of free space, 8.85×10^{-12} F/m
 μ_0 = Permeability of free space, $4\pi \times 10^{-7}$ H/m

II. Introduction

In its simplest configuration, IECF is composed of a negatively charged spherical wire cage in the center, called the cathode grid, surrounded by a positively charged anode on the outside. As ions are created near the anode, the electric field transfers energy to them until they get sufficient energy for fusion in the center core. However, the chance for fusion in the center is actually quite low, so ions must recirculate throughout the device millions of times without impacting any grids. With only a single cathode grid, this proves to be difficult, however research by Dietrich, et al [4] has shown that adding grids inside and outside the cathode grid will ‘focus’ the ion beams. This is referred to as the Multi-grid Configuration. This configuration allows ion trajectories to be better controlled and stimulates ion bunching [4]. Ion bunching is further stimulated in the logical extension of the Multi-grid Configuration – the Continuous-grid Configuration [2] – in which the electrodes are extended radially in the form of a truncated icosahedron (Fig. 1).




Figure 1. IEC geometry [2]. Ions travel in the either pentagonal or hexagonal channels where they can fuse with other ions.

Ions travel through the channels in the device, and potentials can be set continuously along the side walls, allowing for extremely fine granularity in shaping the potential well, and therefore the ion bunch shaping along the beamline [2]. Compact ion bunching would potentially allow the device to operate in a non-thermal ion regime that could allow for the use of aneutronic fuels like $p\text{-}^{11}\text{B}$ [2]. Use of this fuel reduces the need for heavy radiation shielding, and as the products are charged alpha particles, their kinetic energy can potentially be converted directly into electricity [2].

However, to improve confinement, and fusion power, as suggested Chap et al. [1] the ion packets must be sufficiently neutralized with electrons as they pass through the device core. Electrons must be kept confined to the core to facilitate this, primarily with magnetic fields in a magnetic mirror configuration similar to the Polywell [3]. However, because the electrons are moving at average speeds that are orders of magnitude faster than the ions, in order to simultaneously resolve electron and ion motion it is best to model the electrons as a fluid [5]. A fluid model was developed in [1], but it does not capture important physics. Namely, the model neglects the effect of the magnetic field the moving electrons create as they stream out from the center of the simulation, which greatly changes the magnetic field. Adding in this physics provides a numerically tractable model that allows for optimization of the confining magnetic field requirements and investigation into the use of permanent magnets to provide the field. However, adding in the physics in a way the couples with the magnetic field in a self-consistent manner has proven unexpectedly challenging.

III. Fluid Model Context

Development of the fluid model was initiated after testing permanent magnet configurations with COMSOL Multiphysics. Two primary configurations were considered: a configuration with eight bar magnets in a circle, and a configuration with ten wedge magnets also in a circle. Both configurations were considered to facilitate confinement in the core, hence the circular arrangement to create the magnetic mirror effect. Parameters considered included magnet size and magnetic strength along the beamline. To validate the software, the magnetic field of a 0.6m by 0.3m 2T bar magnet was calculated, plotted, and visualized. Having established validity and proficiency with the software,

the magnetic field of eight such bar magnets were visualized. These data were compared with a single bar magnet with symmetric boundary conditions applied, and no significant differences were found.

However, to further validate the symmetric model, the strength of the magnetic flux density was calculated and plotted along two rays as shown in Fig. 2. The blue ray is meant to track the magnetic field along the beamline and the green ray is meant for the surface of the magnet. The green ray is to ensure the model is working properly by validating that the strength of the magnet. The magnets had the same dimensions as before, but they were increased to 7T. As shown in Fig. 3 and Fig. 4 there were no significant differences between the symmetric and asymmetric models.




Figure 2. Ray path traces for symmetric and asymmetric models




Figure 3. Magnetization vs. Distance from origin for asymmetric model. The magnet surface line was included to ensure the model is working correctly and the surface of the magnet is about 7T. Additionally, this plot is not

markedly different from Fig. 4, so the symmetric and asymmetric models are qualitatively similar enough that the symmetric model can be used.






Figure 4. Magnetization vs. Distance from origin for symmetric model. The magnet surface line was included to ensure the model is working correctly and the surface of the magnet is about 7T. Additionally, this plot is not markedly different from Fig. 3 so the symmetric and asymmetric models are qualitatively similar enough that the symmetric model can be used.

Because the symmetric approach took half the computation time, symmetric approaches were used going forward. A symmetric approach was used to compare the bar magnet configuration with that of ten wedge magnets (with dimensions stipulated in Fig. 5a), modeled in symmetric halves. The magnetic field along rays similar to above (along beamline and magnet) were plotted to facilitate this comparison. The magnetic field strength of each magnet was reduced by ten a more realistic value of 0.7T, which complicates comparison. However, this experiment was meant to determine the tradeoff between field strength along the beam line and magnetic density (magnet size) with high strength and high density (low size) being optimal.



a)



b)

Figure 5. Wedge magnet geometrical configuration (right) and ray diagram (left).




Figure 6. Magnetic Field strength vs. Distance from center (r).

As shown in Fig. 6, while a 0.7T wedge magnet is more magnetically dense, it has much lower field strength along the beamline. Increasing the angle subtended increased the strength along the beamline, but decreased the magnetic density. However, a lower field strength along the beamline might still be acceptable if electron confinement can still be maintained, and this necessitated the development of a model that incorporated the permanent magnetic field and electron motion.

IV. Fluid Electron Model Development

Two fluid electron models have been developed, and one is still in development, but none so far have given numerical results consistent with analytical results for test cases. All have had difficulty coupling electron motion to magnetic field generation, which has manifest itself in different ways for each simulation. All simulations were developed in MATLAB.

A. Hybrid Particle in Cell Direct Extension

Development of a fluid-electron model that self-consistently couples magnetic field strength and fluid electron motion began with the extension of the hybrid particle-in cell model developed in [1]. The model in [1] features static magnetic fields generated by 5A current-carrying wires that affect the paths of magnetized electrons streaming out from a Gaussian source in the center. However, the model developed in [1] does not include how the electron motion changes the magnetic field. To remedy this, Eq. (1) (Ampere's law) was added to the simulation. In order to capture the effect, the electrons have on the in-plane magnetic field, the simulation had to be extended to 2.5D. This was done by assuming electrons could move freely in the z-direction but variations in the z-direction were assumed to be zero. Any electrons that moved up and out of the plane (+z-direction) were assumed to be replaced by electron from below, and electrons moving down and out of the plane (-z-direction) were assumed to be replaced by electrons from above. Both of these assumptions allow for the derivation of the error functions Eq. (2)-(4) from Eq. (1).

$$\nabla \times \vec{B} = \mu_0 e \vec{\Gamma} \quad (1)$$

$$g_3 = \frac{\partial B_z}{\partial y} - \mu_0 e \Gamma_x \quad (2)$$

$$g_4 = \frac{\partial B_z}{\partial x} + \mu_0 e \Gamma_y \quad (3)$$

$$g_5 = \frac{\partial B_y}{\partial x} - \frac{\partial B_x}{\partial y} - \mu_0 e \Gamma_z \quad (4)$$

In extending the previous model, a direct method was used to iteratively drive the five error functions (two previous and three new) to zero, described as the steady-state model in Part C of Section 3 of [1] using the Jacobian. External boundaries were set to zero like in [1], but for this case internal boundaries had to be set as well, since defining the magnetic field inside the cross section of the wires was outside the scope of the model. Two methods for setting internal boundaries were used. The first method modeled the wires as circles with a radius equal to 20% of the simulation width. In Fig. 7, the center of the wire is at the red dot, so any node inside the white circle would be considered a wire node, and the magnetic field would be set to zero at those nodes. The pink nodes are unaffected.




Figure 7. Visualization of first internal boundary setting method.

This method did not result in convergence, so a second method was attempted. The second method modeled the wires as rectangular, and took the four closest nodes that fully encapsulated the wire (white nodes in Fig. 8). Additionally, the nodes one node away from each of these nodes (red nodes in Fig. 8) were also considered part of the wire and also set to zero.




Figure 8. Visualization of second internal boundary setting method.

Using this method for setting internal boundaries also did not result in convergence, so an entirely new model was developed that did not require resolution of internal boundaries.

B. Hybrid Particle in Cell with Magnetic Vector Potential Extension

The next model was composed of two parts. The first part was the original model developed in [1]; it determined modeled electron motion based on an applied magnetic field. The second part used Eq. (5), a reformation of Ampere's law based on the magnetic vector potential, to determine in-plane magnetic fields from applied electron fluxes.

$$\nabla^2 \vec{A} = -\mu_0 e \vec{\Gamma} \quad (5)$$

After convergence of the first part (electron motion), the output of the first part would be used as input to the second part, and then the output of the second part would be used as input to the first part (the second part is linear and does not need to converge). This process would continue as the model would slowly approach the steady-state solution, with each part alternately running. Splitting the method in two allowed each part to be validated independently before integration, potentially reducing the amount of error. The second part was validated first by inputting the wire fluxes from [1] and comparing the magnetic field to that of the one analytically derived from the Bio-Savart Law in [1].




Figure 9. Comparison of analytical magnetic field (left) with numerically derived magnetic field (right). White circles denote wire locations

As shown in Fig. 9, both magnetic fields are nearly identical. The numerically derived magnetic field has the locations of the wires as slightly off because the wires are modeled as being at the nearest node, but this is consistent with numerical approximation. The magnetic field from the second (field-deriving) part was fed into the first part and compared with prior results from [1] and was found to agree well. After properly validated the model, it was run for 1000 iterations (Fig. 11) for an 80 by 80 node grid to see if convergence was achieved. Convergence was determined based on the RMS error (Eq. (6)) of the length of the delta vector (Fig 10) at each node.




Figure 10. Definition of the delta vector

$$RMS = \sqrt{\frac{\sum_{i=1}^{nx} \sum_{j=1}^{ny} \|\vec{\Delta}_{i,j}\|}{nx*ny}} \quad (6)$$




Figure 11. RMS error over 1000 iterations of part two code.

As shown in Fig. 13, the error does not seem to decrease, nor behave in a regular fashion; it seems random and noisy. It was surmised this model did not converge because of the complexity and scope in developing a fluid electron model with self-consistent magnetic field coupling. Accordingly, the complexity considered in the model was reduced by scaling down to a 1D, Cartesian simulation and assuming a magnetic field profile.

C. 1D assumed magnetic field model

The last fluid model was developed using a simple balance of forces at each point along the domain. The electromagnetic force on the charged particles at every point along the domain is given by the Lorentz Force in Eq. (7)

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}) \quad (7)$$

Using the fluid approximation for the electron population, evaluating the cross-product in 1D, and adding the force due to the pressure term yields Eq. (8).

$$\sum \vec{F} = 0 = -E(x)en(x) - \frac{B(x)}{\mu_0} \frac{dB}{dx}(x) - \frac{dp}{dx}(x) \quad (8)$$

Substituting in Poisson's equation (Eq. (9)) and using the definition of pressure results in Eq. (10).

$$\frac{dE}{dx}(x) = -\frac{en(x)}{\epsilon_0} \quad (9)$$

$$0 = \epsilon_0 E(x) \frac{dE}{dx}(x) - \frac{B(x)}{\mu_0} \frac{dB}{dx}(x) + \epsilon_0 T_{ev} \frac{d^2 E}{dx^2}(x) \quad (10)$$

Eq. (10) can be used directly, or integrated to yield the Integrated Constraint Equation, Eq. (11).

$$K = \epsilon_0 \frac{E(x)^2}{2} - \frac{B(x)^2}{2\mu_0} + \epsilon_0 T_{ev} \frac{dE}{dx}(x) \quad (11)$$

Using the fact that E is dimensionally similar to $\frac{T_{ev}}{R}$, and taking $B = \bar{B}B_{\max}$ results in the Non-Dimensional Integrated Constraint Equation, Eq. (12).

$$K = \frac{\bar{E}(x)^2}{2} - M^2 \frac{\bar{B}(x)^2}{2} + \frac{d\bar{E}}{dx}(x) \quad (12)$$

M and K in the Eq. (12) are given by Eq. (13) and Eq. (14) below.

$$M = \frac{RB_{\max} c}{T_{ev}} \quad (13)$$

$$K = \frac{\bar{E}(0)^2}{2} - M^2 \frac{\bar{B}(0)^2}{2} + \frac{d\bar{E}}{dx}(0) = \frac{d\bar{E}}{dx}(0) = -\frac{R^2}{\lambda_0^2} \quad (14)$$

This model assumes a function for the magnetic field at steady-state, given by Eq. (15) below

$$\bar{B}(x) = 0.5 \left(1 + \operatorname{erfc}(\sigma(x - u_0)) \right) \quad (15)$$

$$u_0 = 0.5 \frac{R}{B_{\max}} \quad (16)$$

where erfc in Eq. (15) denotes the Complementary Error Function. This specific distribution was chosen because the cusped magnetic field profile the researchers want to properly contain the electrons along the beamline is qualitatively similar – zero near the center and then rising rapidly to a peak value.

Eq. (16) is a constraint function that needs to be added to the model to account for the fact that as the particles press up against the magnetic field, the energy of the magnetic field increases. Eq. (16) allows for proper modeling of the non-dimensionalized specific energy of the system in Eq. (18). If the solution does not converge properly, it could be because this assumption is invalid (which serves as another check). Because the magnetic field is defined numerically, the non-dimensionalized Integrated Constraint Equation must be numerically integrated along the domain.

$$\bar{\epsilon} = \sum_{i=1}^{nx} \left\{ \frac{\bar{E}_i^2}{2} + M^2 \frac{\bar{B}_i^2}{2} - \frac{3}{2} \left(\frac{\bar{E}_{i+1} - \bar{E}_{i-1}}{\Delta x} \right) \right\} \quad (18)$$

Using the central difference scheme for numerally approximating the first derivative results in Eq. (19), which can be rearranged to isolate $\bar{E}_i$ and progressively solved across the entire domain.

$$K = \frac{\bar{E}_{i-1}^2}{2} - M^2 \frac{\bar{B}_{i-1}^2}{2} + \frac{\bar{E}_i - \bar{E}_{i-2}}{\Delta x} \quad (19)$$

Once this model gives an answer consistent with an analytical solution for a test case, σ can be varied and the resulting $\bar{\epsilon}$ found until the σ that gives the smallest $\bar{\epsilon}$ is found. This minimum energy case will be nature's solution that properly accounts for the magnetic field due to electron motion.

V. Results

At this stage, the model does not give answers consistent with analytic solutions. The model was properly validated against the Boltzmann distribution (with $B_{\max} = 0$) and was found to agree (Fig. 12), however, for $B_{\max} \neq 0$ it currently gives results inconsistent with analytical results.




Figure 12. Comparison of potential distributions for the Boltzmann distribution (left) and model (right).

For example, the solver currently gives a negative density for sufficiently high values of $B_{\max}$. For lower values of $B_{\max}$ it gives distributions that vary significantly and do not seem to make sense with the given magnetic field. Work is still ongoing, but at this point the aforementioned numerical fluid electron models do not give results consistent with analytic models. Developing an electron fluid model to capture the electron motion coupled magnetic field, and other essential physics, at the center of an IECF has proven more difficult than anticipated.

VI. Conclusion

Significantly more work is needed to properly develop and validate a fluid-electron model for the core of an IECF. Three models have been tried, and none have produced analytically consistent results so far. Once a model is developed, it will allow for proper analysis and tuning of the permanent magnetic fields around an IECF, which could result in a net-power positive system. If the IECF can be demonstrated to produce sufficient yield, this would lead to a revolution in energy both here on Earth and in space. IECF would allow for significantly cheaper green energy production, potentially slowing climate change and giving millions more cheaper access to power. In space, it would allow for significantly reduced interplanetary mission times, reducing the risk of radiation for manned missions and improving safety. All in all, further development of these IEC core fluid-electron models could lead to a significant change in the way humans live and work in the solar system.

Acknowledgments

The authors would like to acknowledge the Aerospace Department of the A. James Clarke School of Engineering at the University of Maryland, the Maryland Technology Enterprise Institute (Mtech), and the Aerospace Engineering Research Opportunity Scholars (AEROS) Program for their support. General support received from the graduate students of the UMD Space Power and Propulsion Lab (SPPL) was helpful in completing this project. The authors would especially like to thank Andrew Chap for his assistance with this project. This project is funded in part by the AEROS Program and the "A Scholar's Program for Industry-Oriented Research in Engineering" (ASPIRE) program.

References

- [1] Chap, Andrew, and Raymond Sedwick. "A Hybrid Particle-in-cell Simulation for a Multiple Grid Magnetic Core Inertial Electrostatic Confinement Device." In 50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, p. 3516. 2014.
- [2] Chap, Andrew, and Raymond Sedwick. "Simulation of an Inertial Electrostatic Confinement Device Using a Hermite N-body Individual Time-step Scheme." In 51st AIAA/SAE/ASEE Joint Propulsion Conference, p. 3860. 2015.
- [3] Bussard, Robert W. "Some physics considerations of magnetic inertial-electrostatic confinement: a new concept for spherical converging-flow fusion." *Fusion Technology* 19, no. 2 (1991): 273-293.
- [4] Dietrich, Carl C., Leslie J. Eurice, and Raymond J. Sedwick. "Experimental verification of enhanced confinement in a multi-grid IEC device." In 44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, p. 4760. 2008.
- [5] Hagelaar, G. J. M. "Modelling electron transport in magnetized low-temperature discharge plasmas." *Plasma Sources Science and Technology* 16, no. 1 (2007): S57.

- [6] Cheng, C. Z., and Jay R. Johnson. "A kinetic-fluid model." *Journal of Geophysical Research: Space Physics* 104, no. A1 (1999): 413-427.

Appendix D:

Simulation and Optimization of The Continuous Electrode Inertial Electrostatic
Confinement Fusor

Andrew Chap

PhD Dissertation defended November, 2017

ABSTRACT

Title of dissertation: SIMULATION AND OPTIMIZATION OF
THE CONTINUOUS ELECTRODE
INERTIAL ELECTROSTATIC
CONFINEMENT FUSOR

Andrew M. Chap, Doctor of Philosophy, 2017

Dissertation directed by: Professor Raymond J. Sedwick
Department of Aerospace Engineering

A concept for generating nuclear fusion power and converting the kinetic energy of aneutronic fusion products into electric energy is proposed, and simulations are developed to design and evaluate this concept. The presented concept is a spherical fusor consisting of linear ion acceleration channels that intersect in the sphere center, where the converging ions form a high-energy, high-density fusion core. The geometry is that of a truncated icosahedron, with each face corresponding to one end of an ion beam channel. Walls between the channels span radially from the outer fusion fuel ionization source to an inner radius delimiting the fusion core region. Voltage control is imposed along these walls to accelerate and focus the recirculating ions. The net acceleration on each side of the channel is in the direction of the center, so that the ions recirculate along the channel paths. Permanent magnets with radial polarization inside the walls help to further constrain the ion beams while also magnetizing electrons for the purpose of neutralizing the fusion core region. The natural modulation of the ion beams along with a proposed phase-locked

active voltage control results in the coalescence of the ions into “bunches”, and thus the device operates in a pulsed mode. The use of proton-boron-11 ($p\text{-}^{11}\text{B}$) fuel is studied due to its terrestrial abundance and the high portion of its energy output that is in the form of charged particles.

The direct energy converter section envelopes the entire fusion device, so that each fusion fuel channel extends outward into a fusion product deceleration region. Because the fusion device operates in a pulsed mode, the fusion products will enter the energy conversion region in a pulsed manner, which is ideal for deceleration using a standing-wave direct energy converter. The charged fusion products pass through a series of mostly-transparent electrodes that are connected to one another in an oscillating circuit, timed so that the charged fusion products continuously experience an electric field opposite to the direction of their velocity. In this way the kinetic energy of the fusion products is transferred into the resonant circuit, which may then be connected to a resistive load to provide alternating-current energy at the frequency of the pulsed ion beams.

Preliminary calculations show that a one-meter fusor of the proposed design would not be able to achieve the density required for a competitive power output due to limits imposed by Coulomb collisions and space charge. Scaling laws suggest that a smaller fusor could circumvent these limitations and achieve a reasonable power output per unit volume. However, ion loss mechanisms, though mitigated by fusor design, scale unfavorably with decreasing size. Therefore, highly effective methods for mitigation of ion losses are necessary. This research seeks to evaluate the effectiveness of the proposed methods through simulation and optimization.

A two-dimensional axisymmetric particle-in-cell ion-only simulation was developed and parallelized for execution on a graphics processing unit. With fast computation times, this simulation serves as a test bed for investigating long-timescale thermalization effects as well as providing a performance output as a cost function for optimization of the electrode positions and voltage control.

An N -body ion-only simulation was developed for a fully 3D investigation of the ion dynamics in an purely electrostatic device. This simulation uses the individual time-step method, borrowed from astrophysical simulations, to accurately model close encounters between particles by slowing down the time-step only for those particles undergoing sudden high acceleration.

A two-dimensional hybrid simulation that treats electrons as a fluid and ions as particles was developed to investigate the effect of ions on an electrostatically and magnetically confined electron population. Electrons are solved for at each time-step using a steady-state iterative solver.

A one-dimensional semi-analytic simulation of the direct energy conversion section was developed to optimize electrode spacing to maximize energy conversion efficiency.

A two-dimensional axisymmetric particle-in-cell simulation coupled with a resonant circuit simulation was developed for modeling the direct energy conversion of fusion products into electric energy.

In addition to the aforementioned simulations, a significant contribution of this thesis is the creation of a new model for simulating Coulomb collisions in a non-thermal plasma that is necessary to account for both the low-angle scattering

that leads to thermalization as well as high-angle scattering that leads to ion departure from beam paths, and includes the continuous transition between these two scattering modes.

The current implementation has proven problematic with regard to achieving sufficiently high core densities for fusion power generation. Major modifications of the current approach to address the space charge issues, both with regard to the electron core population and the ion population outside of the core would be necessary.

Simulation and Optimization of the Continuous Electrode Inertial
Electrostatic Confinement Fusor

by

Andrew Mark Chap

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2017

Advisory Committee:

Professor Raymond J. Sedwick, Chair/Advisor
Professor Christine M. Hartzell
Professor James D. Baeder
Professor Adil Hassam, Dean's Representative
Professor William Dorland

© Copyright by
Andrew Chap
2017

Dedication

For my wife, Dasha.

Acknowledgments

This thesis was supported by a NASA Space Technology Research Fellowship (NSTRF), grant number NNX13AL44H. I feel very lucky to have been supported by this fellowship, and I thank the NSTRF team for their patience with my sometimes late reports and/or renewal applications.

Thanks to John H. Scott, Sonny White, and everyone at NASA Johnson Space Center for making my time in Houston worthwhile and giving me the the real-world work experiences necessary for me to get the most out of graduate school.

My thesis committee deserves recognition for some very beneficial late-stage input on this work, especially Professor Christine Hartzell for reminding us that the contributions need to be the forefront of this work, and Professors Adil Hassam and Bill Dorland for giving some much needed outside criticism of the CE-IEC design. A special thanks is in order for Professor Dorland who helped direct a lot of the calculations in Chap. 3 and made sure that I would make sure that readers of this thesis can fully understand the pitfalls of the CE-IEC fusor design.

Thank you to my advisor, Dr. Raymond Sedwick, for putting me on this project when I didn't have much of an idea of what I wanted to research, and for giving me plenty of freedom and room for creativity over the course of this research. The skills I learned while working on this project (GPU programming, optimization, stochastic processes) will be valuable throughout the rest of my career.

Thanks to my labmates in the UMD Space Power and Propulsion Laboratory, who were great workmates and even better friends.

Thank you to my parents and family, for their unconditional support over my years in and out of school.

Lastly, and most importantly, I thank my wonderful wife Dasha, for brightening my world and making every day a good day, and of course for giving me the encouragement I needed to finish my disseratation in a relatively timely manner.

I'm forever grateful for how much you believe in me an for your continually strong support. I really don't know what I would do without you.

Table of Contents

1	Introduction	1
1.1	Outline of material to be presented	2
1.2	A conceptual introduction to the CE-IEC	3
1.2.1	Geometry	4
1.2.2	Electrostatic focusing	6
1.2.3	Pulsed operation	6
1.2.4	Active voltage control	8
1.2.5	Permanent magnet geometry	9
1.2.6	Core electron confinement	9
1.2.7	Direct energy conversion	10
1.3	Summary of contributions	13
2	Background and Previous Research	17
2.1	Fusion for energy production	17
2.1.1	Calculation of fusion power	20
2.1.2	Thermal plasma vs. non-thermal plasma for fusion	21
2.2	The two-grid inertial electrostatic confinement fusor	22
2.3	The multi-grid inertial electrostatic confinement fusor	24
2.4	From the multi-grid to the continuous electrode	25
3	Preliminary Calculations for a CE-IEC	27
3.1	Required fuel density for a useful fusor	27
3.2	Bremsstrahlung radiation loss analysis	28
3.3	Space-charge limitation of ion bunch density in the non-neutralized regions	32
3.3.1	Limitation of density due to bunch expansion parallel to the beamline	32
3.3.2	Limitation of density due to bunch expansion transverse to the beamline	33
3.4	Limitation on core density due to the two-stream instability	33
3.4.1	Derivation of the two-stream instability dispersion relation	34
3.4.2	Application of the dispersion relation to a density constraint	37
3.5	Coulomb collisions	38
3.5.1	High probability, low-angle Coulomb collisions	38

3.5.2	Low probability, high-angle Coulomb collisions	39
3.6	Power deposited on the electrodes and thermal management	41
3.7	Power balance between protons, boron ions, and electrons	44
3.8	Limits on electron confinement	46
3.8.1	Space charge limitation on confined electrons	46
3.8.2	Number of electrons required for complete neutralization	47
3.8.3	Electron line cusp loss frequency	47
3.8.4	High- β loss rate along beamline cusps	49
3.9	Scaling laws of the CE-IEC	50
3.9.1	Overcoming space-charge limitations by scaling down	51
3.9.2	Scaling of energy input	53
3.9.3	Scaling of surface erosion	53
3.9.4	Size of a small CE-IEC with significant power density	54
3.9.5	Structural limitations of a small CE-IEC	55
3.9.6	Lawson criterion estimation	56
4	Particle-in-cell Modeling	58
4.1	Domain	58
4.1.1	Axial cell spacing	59
4.1.2	Radial cell spacing	60
4.1.3	Cell volumes	61
4.2	Particle-in-cell algorithm and parallelization	62
4.2.1	Particle-to-cell interpolation to find charge density	63
4.2.2	Calculation of electric potential from charge density	64
4.2.3	Calculation of electric field from electric potential	67
4.2.4	Cell-to-particle interpolation of electric and magnetic field	68
4.2.5	Particle position and velocity updates	68
4.2.6	Particle-particle collision modeling	70
4.2.7	Particle-boundary interactions	70
4.3	Fusion calculation	71
4.4	Fuel species	74
4.5	Optimization routine	75
4.5.1	Algorithm for the optimization wrapper	76
4.6	Optimization results	77
4.6.1	Without magnetic field	77
4.6.2	With magnetic field	80
4.7	Conclusions of the particle-in-cell optimizer	83
5	N-body Simulation	84
5.1	Calculation of the electric field due to electrode voltages	85
5.2	Calculation of the magnetic field due to permanent magnets	88
5.3	The N -body individual time-step method with Hermite integrator	90
5.4	Overestimation of Coulomb scattering due to macroparticle weighting	94
5.5	Testing on two particles with a known scattering angle	95
5.6	Ion simulation results	96

5.7	Electron simulation results	100
5.8	Conclusions of the N -body simulation	101
6	A Fluid Treatment of IEC Electrons	105
6.1	Governing equations	106
6.2	The numerical model	108
6.3	The time-stepping and steady-state models	109
6.4	Test problem and results	111
6.5	Comparison of the fluid model to a particle model	113
7	A Coulomb Collision Model for Nonthermal Plasma Simulation	115
7.1	An overview of Coulomb collisions in plasma simulations	115
7.1.1	A cumulative Coulomb collision model	116
7.2	Relevant previous research on Coulomb collision models	119
7.3	Improvements of this model over previous models	120
7.4	The cumulative binary collision approximation	121
7.4.1	The limit for small a	125
7.5	The validity of the cumulative binary collision approximation	126
7.5.1	Shortcomings of the cumulative binary collision approximation	129
7.6	Heuristic formulae for the cumulative scattering angle	132
7.6.1	Functional fits for numerical data	132
7.6.2	Scattering angle as a function of a random seed	137
7.6.3	A comparison of function fits with numerical data	137
7.6.4	Trends for σ , κ , U_{low} , and U_{high}	139
7.7	Comparison to previous methods	144
7.8	Implementation and comparison to an N -body simulation	146
7.9	Discussion of small impact parameters	148
7.9.1	Fusion event	151
7.9.2	de Broglie wavelength	152
7.9.3	Potential energy equal to kinetic energy	153
7.10	Concluding remarks on the Coulomb collision model	154
8	The Standing Wave Direct Energy Converter	156
8.1	SWDEC overview	157
8.1.1	Past research	158
8.1.2	SWDEC vs. TWDEC	159
8.2	SWDEC simulation overview	159
8.3	A 1D1V semi-analytical simulation of the SWDEC	161
8.3.1	Point-charge description of the ion bunches	161
8.3.2	Comparison between the particle-in-cell simulation of the ion bunches and the point-charge approximation	163
8.3.3	Effect of velocity modulation on ion bunch lifetime	164
8.3.4	Effect of space-charge expansion on ion bunch lifetime	165
8.3.5	1D1V simulation overview	169
8.3.6	Determination of electrode charge distribution	171

8.3.7	The circuit equation	176
8.3.8	Ion bunch deceleration due to charged electrodes	178
8.3.9	Partial validation of the model through demonstration of conservation of energy	179
8.3.10	1D1V electrode spacing optimization	180
8.3.11	Circuit resistance calculation for steady-state operation	183
8.3.12	Demonstration of a self-consistent steady-state simulation	184
8.3.13	Analytical efficiency optimization accounting for ion bunch expansion	184
8.3.14	1D1V optimization results	187
8.4	A 2D3V particle-in-cell simulation of the SWDEC	189
8.4.1	Modeling of floating electrodes	190
8.4.2	Implementation of the circuit equation	191
8.4.3	Calculation of the magnetic field due to a solenoid	192
8.4.4	2D3V simulation results	193
9	Conclusion	195
9.1	Summary of contributions	196
9.2	Problems that still need solutions	198
9.3	Recommendations for future work	199
9.3.1	3D simulation	199
9.3.2	Introduction of optimization degrees-of-freedom	200
9.3.3	Possible fast optimization by finding an unchanging initial particle distribution	201
9.3.4	Global simplex method	202
9.4	Summary on the difficulties of achieving net-power fusion in a CE-IEC	202
	Appendix A Derivation of ion bunch expansion	204
	Bibliography	208

List of Figures

1.1	Conceptual diagram of the continuous electrode inertial electrostatic confinement fusor (CE-IEC): Feed-throughs inside the walls must be supplied for the (a) cathode and the (b) inner anode. The voltage at other points along the wall can be controlled by (c) radially varying resistance along the walls. Along the (d) center of the beamline the (e) electric potential has a “W”-shape	4
1.2	Modified truncated icosahedron with a wall thickness of 0.08 radians.	5
1.3	Possible geometries	5
1.4	Illustration of the “bunching effect” when the kinematic criterion ($dT/dE > 0$) is satisfied.	7
1.5	Frame-by-frame diagram of acceleration of charged particles using time-varying voltages rather than a potential well. The reverse process must be used to decelerate the particles so that particles don’t escape the potential well.	8
1.6	The radially polarized permanent magnet (maroon) shown in a cut-away of the IEC	9
1.7	Conceptual diagram of electron confinement in the CE-IEC. (a) An electron in the fusion core region is (b) prevented from escaping along the beamline by the negative potential of the cathode and the magnetic mirror effect and (c) prevented from striking the inner anode by the magnetic mirror effect.	10
1.8	Cutaway of the CE-IEC with electric potential plotted in the $x-y$ plane and 3D magnetic field lines drawn.	11
1.9	Schematic of a static direct energy convertor	12
1.10	Schematic of an SWDEC array surrounding the CE-IEC	13
2.1	A hierarchy of fusion plasma confinement methods.	19
2.2	A diagram of the IEC ion acceleration mechanism.	22
2.3	<i>Images from Ref. [8]</i> (a) The Multi-grid electrodes in the vacuum chamber. (b) low-vacuum operation of the Multi-grid IEC	24
2.4	Schematic of the multi-grid IEC. Additional grids biased positively relative to the cathode to counteract the defocusing nature of the cathode.	25

3.1	The ratio of fusion power to bremsstrahlung radiation power for various mixture ratios of monoenergetic ions and thermally equilibrated electrons. n_1 is the first species of each fuel as written in the legend, and n_2 is the second. The dependence of the power ratio on electron density only occurs in the Coulomb logarithm, and so changing the electron density has little effect (increases in the electron density moves the power ratios to slightly more favorable values.) An electron density of 10^{22} m^{-3} was chosen for this plot. Maxima occur at the following points: (a) p- ^{11}B (150 keV), $M = 7.4$, $T_e = 54.4 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 0.3$; (b) p- ^{11}B (600 keV), $M = 5.4$, $T_e = 120 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 3.1$; (c) D- ^3He , $M = 2.2$, $T_e = 56.2 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 24.7$; (d) D-T, $M = 0.9$, $T_e = 19.6 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 430$	31
3.2	Dispersion relation $\omega(k)$	37
3.3	With starting energies of $E_{p+} = 550 \text{ keV}$ and $E_{B^{5+}} = 50 \text{ keV}$ and $T_e = 0$, the ion temperatures equilibrate with one another on a faster time-scale than with the electrons. As $tn_e \rightarrow \infty$ the energies are depleted to Bremsstrahlung radiation.	45
4.1	Cell locations (dots) and cell boundaries (lines) for the particle-in-cell domain. A low number of cells is used for this figure for the purpose of clear illustration. The resolution used in the simulation is about four times greater, for a factor of 16 increase in the number of cells as compared to this figure.	62
4.2	Modification of the discrete Poisson equation on a skewed grid	64
4.3	Fusion cross section as a function of center-of-mass velocity for p- ^{11}B fuel. The three sections of Eq. (4.20) are delineated by vertical dashed lines.	74
4.4	Frame of the output of the optimization routine of the 2D3V CE-IEC optimizer.	78
4.5	The cost function output as a function of periods completed, with red circles denoting the iterations where the simulated annealing algorithm found a new optimum.	78
4.6	Frame of a long-timescale simulation of the CE-IEC beamline without a magnetic field.	79
4.7	Frame from an optimization of the CE-IEC with a magnetic field. . . .	81
4.8	The cost function output as a function of periods completed, with red circles denoting the iterations where the simulated annealing algorithm found a new optimum.	81
4.9	The optimal voltage output of the hybrid optimizer moving from the 5 th to the 6 th period.	82
4.10	Frame from the long-timescale simulation of the optimization results with a magnetic field.	82

5.1	Point charge values of the discretized electrodes for electrode voltages (from inner radius to outer radius) of -50 kV, -75 kV, -10 kV, and +10 kV. The electric potential in the x - y plane due to these point charges is shown as well.	87
5.2	Visualization of the discretization of permanent magnets in the calculation of the CE-IEC magnetic field. The volume of the sphere representing each dipole is the same as the Vol_p term in Eq. 5.5. . . .	89
5.3	Testing of the Hermite integrator individual time-step method on a known 90 degree scatter for different values of η . <i>Left</i> : Simulation of a 90° scatter with equal scaling of the x and y axes. <i>Right</i> : Same simulation with the x and y axes of different scaling to the illustrate differences between trajectories.	96
5.4	<i>Left</i> : Comparison of final scattering angle vs. computation time for different values of η . <i>Right</i> : Comparison of the percentage change in total energy vs. computation time for different values of η	97
5.5	A frame from simulation of ions in a truncated icosahedron IEC. . . .	97
5.6	Frame-by-frame plots of data from an ion simulation. <i>Top</i> : The phase space of all particles projected onto one beam line. <i>Middle</i> : The ion density in the x - y plane. <i>Bottom</i> : The beam current along one beam line through the center of the device.	98
5.7	Velocity distribution in the x -dimension of ions in the core region, with one beamline aligned with x	98
5.8	Impact points of ions onto the surface of the CE-IEC over the course of a simulation.	99
5.9	Impact points of ions onto the surface of the CE-IEC over the course of a simulation.	100
5.10	Electrons simulated under the influence of electric and magnetic fields in the CE-IEC showing the relation between power input, electron density, and electron mean lifetime.	102
5.11	Impact points of electrons onto the surface of the CE-IEC over the course of a simulation.	103
6.1	Test problem for the 2D hybrid PIC simulation. Six wires, three of which have positive current perpendicular to the plane and three of which have negative current create a confining magnetic field. A electron source function replenishes electrons in the center of the domain.	111
6.2	Comparison between the time-stepping method (left) and steady-state method (right) solutions of the electron density in the test problem.	112
6.3	Comparison between computation times for the time-stepping model and steady-state model. “ Δt ” is the length of the time step used as determined by the CFL number, the grid spacing, and the characteristic velocity of either the electrons (time-stepping model) or the ions (steady-state model).	112

6.4	Test problem for the 2D hybrid PIC simulation. <i>Top row, l-r:</i> The electron source term, steady-state state density solution, electric potential created by the electrons. <i>Bottom row, l-r:</i> The drift term ($\mu n_e \nabla \Phi$), the diffusion term ($\mu \nabla (n_e T_e)$), positions of the ion macroparticles.	113
6.5	Side-by-side comparison of the electron fluid simulation with a particle-in-cell simulation of electrons using equivalent conditions.	114
7.1	A 2-dimensional cross-sectional schematic of the N -body simulation for testing the cumulative binary collision approximation. The test particle travels a distance of $v\tau = 2$ mm through a sphere of field particles but only experiences a force from field particles within a distance of $b_{\max} = 1$ mm.	127
7.2	Probability distribution functions for varying values of $b_{\max}$ with $v_\alpha = 10^3$ m/s, $m = 1$ AMU, $n = 10^{11}$ m $^{-3}$ and $\tau = \mu$ s. <i>Top:</i> Results of the N -body simulation with fixed field particles. <i>Bottom:</i> Results of the cumulative binary collision approximation.	130
7.3	Comparison of scattering angles produced by the cumulative binary collision approximation with scattering angles produced by the three pieces of Eq. (7.28).	139
7.4	Trends for σ , $\tilde{\sigma}$, κ , u_{low} , and u_{high} for cases in which the scattering angle is very small and the results depend only on N	141
7.5	Plots of σ , $\tilde{\sigma}$, κ , u_{low} , and u_{high} along with best-fit functions for a range of a and N . Selected contours of constant value are plotted to aid in comparison.	143
7.6	A comparison of the probability distribution functions for the scattering angle between the cumulative binary collision approximation (Sec. 7.4), the N -body simulation (Sec. 7.5), the Nanbu method, the Takizuka-Abe method, and the present method (Sec. 7.6)	145
7.7	The relative discrepancy of the mean scattering angle for the Nanbu method and the present method as compared to the results of the cumulative binary collision approximation.	146
7.8	A frame from the counter-streaming N -body simulation used for testing the collision model.	147
7.9	Time-averaged density for four different simulations of counterstreaming ion beams. The plots are axisymmetric about the z -axis and plane-symmetric about the r axis. The envelope of the beam sourced at $z = 5$ mm is visible as a dark shade and the envelope of the beam sourced at $z = -5$ mm is visible as a light shade in all plots. The density resulting from high-angle scatters permeates the remainder of the domain and is displayed using contour lines of constant value. Densities down to 10^6 m $^{-3}$ are resolved by time-averaging the density over 0.5 ms. Densities below 10^6 m $^{-3}$ are not resolved.	149

7.10	A comparison of scattering angle probabilities with the probabilities of u_{fusion} (a fusion event), $u_{\text{de Broglie}}$ (significant interaction of matter waves), and $u_{\text{potential}}$ (potential energy exceeding kinetic energy) occurring.	154
8.1	Schematic for TWDEC test article at NASA Johnson Space Center.	157
8.2	Frame-by-frame illustration of the SWDEC deceleration mechanism using four ring-shaped electrodes. Each electrode has an alternating electric charge, creating a standing wave along the axis. A correctly timed ion will consistently experience a positive potential gradient, resulting in the deceleration of the ion.	161
8.3	A frame-by-frame comparison of the point-charge description of the modulated ion beam with the 2D axisymmetric particle-in-cell simulation of the modulation process. Particles are moving from left to right. The two methods are simulated separately and then superimposed upon one another for comparison. The modulator electrodes do not have any effect on the point-charge bunches. Axial and radial axes are of different scales for clarity.	164
8.4	A comparison of the effect of modulation voltage on ion bunch formation. A higher voltage (top) results in quick formation of bunches, while a lower voltage (bottom) leads to longer bunch lifetimes. The simulation uses a low beam current (1 ampere) so that the expansion of the bunches due to space charge is low.	166
8.5	2D axisymmetric particle-in-cell simulation developed in [51] of the expansion of an initially spherical ion bunch. Left: symmetric expansion in the absence of a magnetic field. Middle: radial expansion limited by an axial magnetic field increases the rate of axial expansion. Right: comparison with the theoretical ion bunch radius.	168
8.6	The circuitry schematic for an SWDEC with eight electrodes. The system is an RLC circuit with the odd/even electrodes acting as a capacitor. Converted energy from the decelerating ions is stored in the inductor and capacitive electrodes, and dissipated in the resistor.	170
8.7	An infinitesimal increase in potential $d\Phi$ on ring j due to a charge dq on an infinitesimal segment of electrode ring i . Due to axial symmetry and the assumption that each electrode ring is equipotential, the position of $d\Phi$ can be chosen for convenience.	172
8.8	A potential increase $d\Phi$ at location z due to a charge dq on an infinitesimal segment of electrode ring i	178

8.9	Demonstration of the conservation of energy: The ion bunch enters a region of eight equally spaced decelerator electrodes shortly after 4 microseconds into the simulation, and excites the RLC circuit, where the bunch energy is transferred into an oscillation alternating between the inductor and capacitor. The bunch leaves the decelerator region shortly before 6 microseconds into the simulation, and the oscillating circuit energy is dampened and dissipated by the resistor. The total energy remains unchanged throughout the simulation.	180
8.10	All units arbitrary. A demonstration of the electrode spacing optimization. Over each iteration the circuit amplitude is increased, and the electrode spacing is modified to correspond with the deceleration of the test particle. By the 84th iteration, the conversion efficiency has achieved approximately 90%.	182
8.11	A frame-by-frame demonstration of the steady-state operation of the simulation, with asterisks denoting the axial positions of the ion bunches. The operation of the SWDEC in this simulation is self-sustaining, in that the only power input is the incoming ion bunches. The energy gained by the circuit from the decelerating bunches is offset by the energy dissipated in the resistor, and so the amplitude stays constant. This simulation demonstrates what is illustrated in Fig. 8.2: the ion bunches only experience “uphill” potentials while in the decelerator.	185
8.12	Optimization of efficiency as a function of beam current. Efficiency is capped at 90% to allow accurate calculation of the decelerator region length and other parameters.	188
8.13	A magnetic field resulting from two solenoids, discretized according to the black dots plotted in the domain	194
8.14	A frame from a particle-in-cell simulation of the decelerator electrodes of the SWDEC. This simulation served as a test of the electrode deceleration optimization using the particle-in-cell method. Particles are moving from left to right. The axial and radial axes are of different scales for clarity.	194
9.1	The truncated icosahedron can be split into 120 symmetric slices. One symmetric slice (raised area) contains part of a hexagon and part of a pentagon.	199
9.2	A single symmetric slice of the truncated icosahedron IEC with the wall sections shown.	200

Chapter 1

Introduction

This thesis introduces and evaluates a new concept for powering spacecraft via nuclear fusion named the Continuous Electrode Inertial Electrostatic Confinement (CE-IEC) fusor. Through a lightweight design and efficient direct conversion of fusion energy into electrical power, this fusor is studied as a possible breakthrough alternative to existing space power technology. The fusor consists of intersecting beam channels, each of which confines a population of recirculating ionized fusion fuel. The beam channels intersect at a common open center point where each ion has kinetic energy suitable for a fusion event with an ion traveling in the opposite direction.

1.1: Outline of material to be presented

- The remainder of this chapter gives a top-level introduction of the CE-IEC Fusor design and summarizes the contributions of this thesis.
- In Chap. 2 the classification of the CE-IEC is put in context within the field of nuclear fusion. Then, more specifically, the lineage of the CE-IEC concept is presented.
- In Chap. 3 some preliminary order-of-magnitude calculations are made to roughly define the operating conditions necessary for a useful space-based CE-IEC Fusor. Limitations on operating conditions due to the relevant plasma physics are then evaluated, and scaling laws are defined.
- In Chap. 4 a parallelized 2D3V axisymmetric particle-in-cell (PIC) simulation is presented, and the use of the simulation as an optimization tool is presented. The results of optimization and long-timescale simulation are presented and discussed.
- In Chap. 5 an individual time-step (ITS) N -body simulation is presented and used for observing bunching synchronization among beamlines, detecting ion transfer between beamlines, simulating electron confinement, and profiling ion and electron surface impact points.
- In Chap. 6 a Scharfetter-Gummel electron fluid simulation is presented and evaluated against a particle-in-cell simulation of an identical scenario.

- In Chap. 7 the Coulomb collision model that was developed for use in the PIC simulation in Chap. 7 is presented.
- In Chap. 8 the Standing Wave Direct Energy Converter (SWDEC) is presented, and simulations are used to optimize electrode spacing and evaluate energy conversion performance.

1.2: A conceptual introduction to the CE-IEC

The CE-IEC design provides a means to confine both ions and electrons using a carefully structured electric potential well geometry and permanent magnets. The following features are designed to maximize confinement time and minimize energy losses. The features can be roughly summarized as follows:

- Electrostatic focusing is employed to minimize ion collisions with electrodes and other surfaces.
- Permanent magnets assist ion confinement along the beam channels, and magnetic cusps in the center help confine electrons.
- Operation of the ions in a “pulsed”, or “bunched” manner limits ion counterstreaming to only the fusion area, minimizing the thermalization process of the ions.
- A mostly transparent (as seen from the center of the device) structure and the use of direct energy conversion result in a lightweight, high efficiency device.




Figure 1.1: Conceptual diagram of the continuous electrode inertial electrostatic confinement fusor (CE-IEC): Feed-throughs inside the walls must be supplied for the (a) cathode and the (b) inner anode. The voltage at other points along the wall can be controlled by (c) radially varying resistance along the walls. Along the (d) center of the beamline the (e) electric potential has a “W”-shape

1.2.1 Geometry

The geometry chosen for this study is that of a truncated icosahedron, modified to increase the area of the pentagonal faces at the expense of the hexagonal faces, so that all faces have near-equal area, which is best for symmetry between beamlines. The truncated icosahedron shown in Fig. 1.2 has a transparency (as seen from the center of the device) of about 80%. Other symmetric geometries are possible: removing the hexagonal faces reveals a dodecahedron, while further subdivision of an icosahedron adds additional rings of hexagons around each pentagon (Fig. 1.3). With fewer faces, a higher transparency (as seen from the center of the device) is possible. A high transparency is desirable for two reasons: ions undergoing a high-angle scatter may scatter onto a different beamline rather than striking the inner edge, and fusion products have a higher chance of making it to the direct




Figure 1.2: Modified truncated icosahedron with a wall thickness of 0.08 radians.




Figure 1.3: Possible geometries

energy conversion unimpeded. On the other hand, larger numbers of faces decrease the transparency but allow closer control of the beamline potential (because of the narrower beamlines) and also provide more beamlines to contribute to fusion.

1.2.2 Electrostatic focusing

A static-voltage linear ion accelerator is an inherently defocusing device, because the low-voltage electrode (cathode) necessary for accelerating the ions will accelerate these ions towards the surface and away from the central beamline. This is most easily remedied by placing electrodes before and/or after the cathode, biased positively relative to the cathode, to direct ions back towards the beamline.

1.2.3 Pulsed operation

Pulsed operation, also referred to as the “bunching” of the ions, occurs naturally, and has been observed in electrostatic ion traps [1] and in the multi-grid IEC experiments [2]. The bunching arises when the kinematic criterion is satisfied:

$$\frac{dT}{dE} > 0 \quad (1.1)$$

where E is the energy of a particle and T is the the oscillation period of that particle if it were the only particle in the system, absent from space-charge or collisional effects from other particles. When Eq. 1.1 is satisfied, the geometry of the potential well is such that increasing the energy of an ion will increase the oscillation period of the ion. In such a potential well geometry the bunching will naturally form




Figure 1.4: Illustration of the “bunching effect” when the kinematic criterion ($dT/dE > 0$) is satisfied.

from a continuous beam by the following process: If there is a perturbation or non-uniformity in the ion beam that results in a slight increase in ion density, the ions near the front of this perturbation (where the “front” is in the direction the ions are moving) will be accelerated by the space charge of the perturbation, thereby gaining energy. With the extra energy, these ions will travel farther up the well in the turnaround region, and their period will increase (Eq. (1.1)), causing them to take more time to traverse the beam line, thereby moving them towards the rear of the perturbation. The ions near the rear of the perturbation undergo the opposite process: they are decelerated by the space charge of the bunch and lose energy, they don’t travel as far up the potential well as the ions in front of them so that their period is shortened, causing them to traverse the beam path in less time, moving them to towards the front of the bunch. In this way, the so-called “trap kinematics” are causing ions to move towards regions of higher density, so that any small perturbation will grow.




Figure 1.5: Frame-by-frame diagram of acceleration of charged particles using time-varying voltages rather than a potential well. The reverse process must be used to decelerate the particles so that particles don't escape the potential well.

1.2.4 Active voltage control

Another dimension of control over the ion dynamics is achieved with time-varying electrode voltages. Voltages could be modulated as the bunches pass by electrodes for added bunch compression to increase peak density before entering the core. In an extreme version of active bunch control, the electrodes are operated in the fashion of a particle accelerator, with ions accelerated towards the core and decelerated when traveling away from the core. In this case, a potential well may not even be necessary if all ion acceleration is performed by oscillating voltages, but may require a separate voltage feed to each electrode, as illustrated in Fig. 1.5.




Figure 1.6: The radially polarized permanent magnet (maroon) shown in a cut-away of the IEC

1.2.5 Permanent magnet geometry

The walls of the CE-IEC are proposed to contain (or be constructed of) radially polarized permanent magnets, which can be fabricated with a magnetic field strength of approximately $M \approx 1$ T. If it is assumed that the magnets occupy half of the walls (see Fig. 1.6) then the transparency of the bare magnets in this case will be approximately $t_m \approx 90\%$. Due to the conservation of magnetic flux, the magnetic field strength along the beamlines will be approximately $B \approx \frac{1-t_m}{t_m} M = 0.11$ T.

1.2.6 Core electron confinement

The “W”-shape of the beamline potential of the CE-IEC, along with the cusped magnetic structure of the permanent magnets, lends itself to an electron confinement region. Electrons that are in the fusion core region are prevented from escaping along the beam paths by the strongly negative cathode grid and are pre-




Figure 1.7: Conceptual diagram of electron confinement in the CE-IEC. (a) An electron in the fusion core region is (b) prevented from escaping along the beamline by the negative potential of the cathode and the magnetic mirror effect and (c) prevented from striking the inner anode by the magnetic mirror effect.

vented from directly hitting the inner anode grid by the magnetic mirror effect of the field cusps (see Fig. 1.7). The complete CE-IEC prototype with the electric potential and magnetic field lines is shown in Fig. 1.8, plotted using methods described in Chap. 5.

1.2.7 Direct energy conversion

The primary draw of using $p\text{-}^{11}\text{B}$ fuel over D-T (deuterium-tritium) fuel is that $p\text{-}^{11}\text{B}$ fusion produces only charged α -particles rather than neutrons. Not only are high-energy α -particles much more easily stopped by matter than neutrons are, but their charge also allows for a more efficient conversion of energy than is possible




Figure 1.8: Cutaway of the CE-IEC with electric potential plotted in the x - y plane and 3D magnetic field lines drawn.




Figure 1.9: Schematic of a static direct energy convertor

with thermodynamic energy conversion.

1.2.7.1 Static direct energy conversion

Energy of the α -particles may be captured by biasing the entire fusor to a negative potential, so that escaping α -particles are decelerated by the potential difference (Fig. 1.9), and in doing so raise the potential difference by a small amount which can then be used to power an electric load. However, such an approach requires a very high (≈ 4 MV) potential to fully decelerate all fusion products.

1.2.7.2 Standing-wave direct energy conversion

Direct energy conversion for the CE-IEC is proposed to be achieved through the standing-wave direct energy converter (SWDEC) [3]. A series of mostly-transparent electrodes surround the CE-IEC (either ring electrodes that extend the beamlines,




Figure 1.10: Schematic of an SWDEC array surrounding the CE-IEC

or gridded electrodes) and an oscillating potential is induced between the alternating even and odd electrodes, which are connected via an inductor and resistor. The oscillation is timed so that the passing α -particles are decelerated by the electrodes, thereby driving the oscillation, which must be damped by a resistor to maintain steady-state. The resistor in this case is the power load of the spacecraft and electric propulsion system, to which the SWDEC provides alternating-current electricity.

1.3: Summary of contributions

The contributions of this thesis to the area of inertial electrostatic confinement research may be summarized as follows:

- The CE-IEC concept was developed as a natural evolution of the previous generation device, the Multi-grid IEC.

- A 2D3V axisymmetric particle-in-cell (PIC) simulation of a single channel of the CE-IEC was developed using MATLAB, C, and CUDA for execution on a general purpose graphics processing unit (GPU) and may be used by other researchers studying this or related concepts. In this research, it was used for two purposes:

- Optimization of the ion channel voltage profile was performed using a cost function to maximize the bunching behavior of the ions and minimize ion losses. Successful optimization demonstrated long ion lifetimes on the order of 3000 oscillation periods when the ion density was under the space-charge limit.
- Long time-scale simulation of the IEC to reach an oscillatory steady state to evaluate the effect of thermalization on ion behavior. These simulations concluded that thermalization of the ion bunches appears to continue despite kinematic constraints on the system that formed the bunches initially. Active control of thermalization was an intended path of investigation, but it is not included in the current work.

- A fully 3D N -body simulation was developed using methods borrowed from the field of astrophysical systems, allowing for analysis of the interaction between beamlines and effects of a non-uniform cylindrical beamline potential profile. This simulation reached the following conclusions:

- Ions can transfer between beamlines due to high-angle scattering, though newly transferred ions are often lost shortly thereafter due to their trajec-

tories being far off-axis.

- Electron simulation demonstrated a steady state density close to that of the core ion density, though neither density was high enough for a significant fusion power output.
- Ion impacts are mostly limited to the inner edge of the device, and electron impacts are exclusively limited to the inner edge of the device.
- A Scharfetter-Gummel simulation was developed to simulate electrons as a thermal fluid under the influence of a static external magnetic field. This may be used in future research for investigating the effect of the electron pressure on the magnetic field in the CE-IEC. However, comparisons between an electron fluid simulation and an electron particle simulation for a test problem did not agree well enough to continue along this path.
- A Coulomb collision model was developed to account for both low-angle Coulomb scatters that lead to thermalization as well as high-angle scatters that throw ions off of beamlines.
- A 1D1V semi-analytic simulation of the Standing Wave Direct Energy Converter (SWDEC) was developed to optimize electrode spacing for optimal direct energy conversion efficiency.
- A 2D3V PIC code was developed for the SWDEC to test the optimized results and demonstrated the direct conversion of fusion products into electricity at a 50% conversion efficiency, based on a design that was optimized via the 1D1V

model to operate at a 65% efficiency.

Chapter 2

Background and Previous Research

2.1: Fusion for energy production

Nuclear fusion is the process by which two atomic nuclei unite, and the energy gained or released in this process is related by the difference in mass of the product(s) with that of the mass of the reactants by

$$\mathcal{E} = (m_{\text{products}} - m_{\text{reactants}}) c^2 \quad (2.1)$$

where c is the speed of light. This process is exothermic ($\mathcal{E} > 0$) for light nuclei (when the product is lighter than Iron-56). Fusion is contrasted with nuclear fission which is the splitting of a nucleus. Fission is exothermic for heavy nuclei, and has been utilized with success for terrestrial energy production.

Fusion reactions occur naturally in stars, where gravity confines and heats

matter (mostly hydrogen and other light atoms). Nuclei are repelled from one another by the electrostatic force and attracted to one another by the strong nuclear force. The electrostatic force is dominant at long distances, down to within a few femtometers of the nucleus, requiring a significant relative kinetic energy for nuclei to overcome this repulsive barrier. At sufficiently high temperatures, a significant fraction of the fusion fuel ions will be energetic enough to reach this close proximity, aided by the process of quantum tunneling.¹ The energy produced by a fusion event in a star provides heat to the surrounding matter and in this way the process is self-sustaining. The Sun produces power via nuclear fusion at a rate of approximately one watt per cubic meter. In any fusion scenario, the energies are high enough that the fusion fuel atoms are completely ionized, and the fuel ions and electrons together form a quasineutral plasma.

Efforts to generate energy through terrestrial fusion reactions require both a method of energizing the fusion fuel plasma so that a significant fusion reaction rate is present, and a method of confining the plasma to prevent energy loss of the fuel. The ionizing and energizing of fuel is generally easily achievable in the laboratory, but the simultaneous confinement of such a fusion fuel plasma at a sufficient density remains elusive. Thus, terrestrial methods of fusion for energy production are often classified according to their confinement schemes.

The largest share of investment in fusion power research for electricity generation resides in magnetic confinement. The charged particles of a plasma are

¹This thesis is not concerned by the physics of this process, but will use instead the simplified notion of a fusion cross-section for the calculation of fusion events




Figure 2.1: A hierarchy of fusion plasma confinement methods.

inhibited from moving perpendicular to a magnetic field by the Lorentz force, and so a magnetic confinement fusion reactor consists of magnetic fields either parallel to the walls of the chamber or in a cusped configuration to reflect particles back into the plasma. The magnetic fields are typically generated by strong electric currents, either applied externally or appearing internally within the plasma. The fusion reactions heat the plasma, and excess heat from the reactor is converted thermodynamically to electricity in the same manner as nuclear fission and hydrocarbon power plants.

Inertial confinement involves a multi-directional transfer of momentum to the fusion fuel to direct nuclei onto collision paths with one another, resulting directly in fusion reactions and/or thermalized heating that produces fusion reactions. What

can be considered the only successful macro-scale fusion reaction for local energy production, the hydrogen bomb, belongs in this category, but notably requires a fission reaction to bring the hydrogen and/or lithium to the required fusion energy. Laser fusion involves a similar implosion of a fuel pellet, using short, high-energy focused laser pulses to heat a shell surrounding a fusion fuel core.

Inertial electrostatic confinement is one of few confinement schemes that attempts a completely non-thermal approach to confining ionized fuel. High electric fields accelerate ions moving in opposite directions to fusion energies. Since the fusor of this work falls into this category, a more detailed look at the history, methods, and challenges of this confinement scheme follows.

2.1.1 Calculation of fusion power

The fusion power per unit volume produced by a plasma consisting of ions of species 1 and 2 is

$$\frac{P}{\text{Vol}} = \mathcal{E} \int_{\mathbf{v}_1} \int_{\mathbf{v}_2} f_1(\mathbf{x}, \mathbf{v}) f_2(\mathbf{x}, \mathbf{v}) |\mathbf{v}_1 - \mathbf{v}_2| \sigma(|\mathbf{v}_1 - \mathbf{v}_2|) d^3\mathbf{v}_1 d^3\mathbf{v}_2 \quad (2.2)$$

where $\mathcal{E}$ is the energy released when an ion of species 1 fuses with an ion of species 2, and $\sigma(v)$ is the velocity-dependent fusion cross section. The fusion cross section $\sigma(v)$ is unique for each fuel species pair, and is determined primarily experimentally. For species pairs of interest to the field of laboratory fusion, $\sigma(v)$ typically peaks at center-of-mass energies of 50 to 3000 keV, and so laboratory devices must produce voltages on this scale to achieve fusion. The cross section for proton-boron-11 fusion

(p-¹¹B) has two primary peaks, at 150 keV and 600 keV, the latter of which is considered for the majority of this work.

Eq. 2.2 may be simplified for thermal plasma through the use of an integrated mean velocity and cross section product, $\langle \sigma v \rangle$

$$\frac{P}{\text{Vol}} = n_1 n_2 \langle \sigma v \rangle_{1-2} \mathcal{E} \quad (2.3)$$

and approximated for a non-thermal plasma with species 1 and 2 both monoenergetic and moving at a relative speed v_{1-2} from one another as

$$\frac{P}{\text{Vol}} = n_1 n_2 v_{1-2} \sigma(v_{1-2}) \mathcal{E} \quad (2.4)$$

which is the method that will be used for parameter estimation in this chapter.

2.1.2 Thermal plasma vs. non-thermal plasma for fusion

A thermal plasma is one in which the species all have the same mean energy and have energies and velocities that follow Maxwellian distributions. Any non-thermal plasma will “thermalize” over time unless there is some process to actively keep it in a non-thermal state. The primary driver of plasma thermalization is Coulomb collisions. A non-thermal plasma fusor must produce more fusion energy than the energy required to maintain the non-thermal state.




Figure 2.2: A diagram of the IEC ion acceleration mechanism.

2.2: The two-grid inertial electrostatic confinement fusor

The first inertial electrostatic confinement (IEC) experimental fusor was built and tested by Hirsch and Farnsworth in 1967 [4]. This device, along with many following ones, consisted of two spherically concentric, mostly transparent electrodes, with the inner cathode biased to a negative voltage relative to the outer anode (Fig. 2.2.) At the center, the ions have enough energy to overcome their mutual electrostatic repulsion and fuse. The probability of fusion at each pass is low. Ions that do not fuse are decelerated on the other side of the potential well, turning around just prior to reaching the anode, to be accelerated towards the cathode grid once again.

In the two-grid IEC, the following phenomena preclude net power generation:

- The cathode grid defocuses the ion beams, causing ions to stray off of the beam

paths onto trajectories that are not likely to result in fusion, often striking the cathode grid wires instead.

- Coulomb collisions tend to scatter ions off of beam paths as well. These collisions happen along the entirety of each beam path, while fusion events may only happen in the high-energy core. The result is that scatter collisions greatly outnumber fusion events and plasma thermalization becomes problematic.
- If ion lifetimes are not limited by collisions with the cathode, they are typically limited by collisions with neutrals within the vacuum chamber.
- The voltage feed stalks of the cathode create an asymmetry in the potential well, so beam paths tend to be curved in the direction of these feeds.

Despite these barriers to net power generation, the Hirsch Farnsworth two-grid IEC remains the canonical fusor, and has been built in many research universities such as the University of Wisconsin-Madison, USA, University of Sydney, Australia, Kyoto University, Japan, and Tokyo Technical Institute, Japan; the private company Phoenix Nuclear Labs in Madison, Wisconsin; and even in the garages and basements of hobbyists. The end-goal of this type of fusor is typically for the safe and compact generation of high-energy neutrons from fusion reactions for purposes such as medical isotope production and neutron imaging. The majority of the neutrons produced in these fusors are due to ion-neutral (“beam-background”) rather than ion-ion (“beam-beam”) fusion events.




Figure 2.3: *Images from Ref. [8]* (a) The Multi-grid electrodes in the vacuum chamber. (b) low-vacuum operation of the Multi-grid IEC

Researchers of the two-grid IEC have concluded that the only pathway to net-power generation in an IEC fusor is through sustained beam-beam fusion [5]. This requires a high vacuum for long mean free paths to reduce collisions with neutrals, and a way to limit ion collisions with the cathode grid. One method of reducing ion-grid collisions is to replace the physical cathode with a “virtual cathode” of confined electrons (e.g. the Polywell [6] or the Penning Trap [7]). Another method is the multi-grid IEC, described in the next section.

2.3: The multi-grid inertial electrostatic confinement fusor

To overcome the defocusing nature of the accelerating cathode grid, the “multi-grid” approach of Sedwick, Dietrich, McGuire, and Eurice [2,8,9], shown in Fig. 2.3 was to introduce additional electrode grids inside and outside of the cathode grid, biased positively relative to the cathode grid, to push ions back towards the beamline axis after being accelerated and pulled away from the axis by the cathode grid.

The multi-grid research demonstrated an improvement in ion confinement times of up to thousands of passes before loss. The increase of the average ion lifetime due to




Figure 2.4: Schematic of the multi-grid IEC. Additional grids biased positively relative to the cathode to counteract the defocusing nature of the cathode.

electrostatic focusing revealed another phenomenon: the tendency of ions to coalesce into bunches so as to operate in a pulsed manner rather than as time-invariant recirculating beams. This arises when the “kinematic criterion” is satisfied [1], that is, when an increase in ion energy results in a lengthening of the ion’s oscillation period, which was discussed in Sec. 1.2.3.

2.4: From the multi-grid to the continuous electrode

If an additional two electrodes inside and outside the cathode were found to be beneficial for ion confinement, a next logical step is to continuously experiment with higher numbers of electrodes to obtain fine-grain control of the potential well structure. At this extreme the grids become so close together that connecting them radially becomes beneficial. This is how the concept of the CE-IEC was born, by replacing radially-spaced grids with continuous electrode walls. While it may seem disadvantageous to introduce more surfaces in an environment that is sensitive to

ion loss, ions that strike the walls separating the channels were most likely already on non-radial paths that would not result in fusion. Following from the multi-grid findings, voltage feeds inside the walls to both the cathode and inner anode would be necessary, while radially varying resistance between grids can provide additional control over the radial potential profile, as illustrated in Fig. 1.1. This would also avoid the feed-stalk asymmetry problem, and would likely require additive manufacturing of the electrodes, magnets, resistors and insulators, though it is possible that the electrodes and their feeds could be manufactured from Neodymium-based ferromagnetic material.

Chapter 3

Preliminary Calculations for a CE-IEC

3.1: Required fuel density for a useful fusor

The purpose of this research is the advancement of a space-based fusor for low- α propulsion systems, where α is the mass of the power systems divided by the energy produced, usually expressed in kg/kW. NASA's technology roadmap cites the need for a specific mass "well under 3 kg/kW" for enabling sustained trips to and from Mars at a cost comparable to NASA's budget [10]. As a baseline, a 1 meter radius IEC producing 1 MW of energy will be considered. For this, the 600 keV peak cross section of p- ^{11}B is chosen. Some useful constants for this fuel and peak are as follows:

Properties of p-¹¹B at peak fusion cross section

Property	Symbol	Value
Fusion cross section	σ	$1.2 \times 10^{-28} \text{ [m}^2\text{]}$
Fusion output energy	$\mathcal{E}$	8700 [keV]
Center-of-mass energy	E_{COM}	600 [keV]
Center-of-mass velocity	v_{COM}	$1.12 \times 10^7 \text{ [m/s]}$
Proton energy	E_p	550 [keV]
Boron energy	E_B	50 [keV]

Using Eq. 2.4, the average plasma density n required for a given fusion power P is

$$n = \sqrt{\frac{P}{\frac{4}{3}\pi R_c^3 v_{\text{COM}} \sigma \mathcal{E}}}. \quad (3.1)$$

Using $P = 10^6 \text{ W}$ and $R_c = 5 \text{ cm}$, the required density is on the order of $n_p, n_B \approx 10^{21} \text{ m}^{-3}$. The required density may be lowered by an order of magnitude due to the multiple beamlines all contributing to fusion. The density in the acceleration and turnaround regions may be lowered by perhaps two orders of magnitude due to the converging nature of the fusion core, so these regions may only need to accommodate a density of 10^{19} m^{-3} . The electron confinement region is two orders of magnitude larger than the fusion core, so unless the electrons are regenerated anew at each cycle, the electron density must be on the order of $n_e \approx 10^{19} \text{ m}^{-3}$.

3.2: Bremsstrahlung radiation loss analysis

Energy losses due to Bremsstrahlung radiation, if the radiation is not converted into usable energy, renders impossible the use of p-¹¹B for thermonuclear fusion, as the Bremsstrahlung radiation power density exceeds the fusion power density for

any plasma temperature and plasma density. In this section, the ratio of fusion power density to Bremsstrahlung power density for electrons in thermal equilibrium with monoenergetic ions (rather than thermalized ions) is considered.

When the electrons are in thermal equilibrium with monoenergetic ions, the energy radiated by the electrons through Bremsstrahlung is equal to the energy transfer from ions to electrons:

$$P_{\text{Brem}} = P_{i \rightarrow e}. \quad (3.2)$$

Bremsstrahlung radiation power density for relativistic electrons, normalized by the square of the electron density, is

$$\frac{P_{\text{Brem}}}{n_e^2} = 1.69 \times 10^{-38} \sqrt{T_e} \left\{ \frac{3}{\sqrt{2}} \frac{T_e}{m_e c^2} + \sum_i \frac{Z_i^2 n_i}{n_e} \left[1 + 0.7936 \frac{T_e}{m_e c^2} + 1.874 \left(\frac{T_e}{m_e c^2} \right)^2 \right] \right\} \quad (3.3)$$

where T_e and $m_e c^2$ are in eV. The ion to electron collisional power transfer density, normalized by the square of the electron density, is given by Ref. [11].

$$\frac{P_{i \rightarrow e}}{n_e^2} = 7.61 \times 10^{-34} \frac{\log \Lambda}{T_e^{\frac{3}{2}}} \left(1 + \frac{0.3 T_e}{m_e c^2} \right) \sum_i \frac{Z_i^2 g_i}{\bar{m}_i} \left(1 + \frac{m_e \frac{2}{3} E_i}{m_i T_e} \right)^{-\frac{3}{2}} \left(\frac{2}{3} E_i - T_e \right) \quad (3.4)$$

where E_i , T_e and $m_e c^2$ are in eV, $\log \Lambda = 24 - \log \frac{\sqrt{n_e}}{T_e}$, $\bar{m}_i$ is the ion mass in AMU, and g_i is the ratio of ion density to electron density. The fusion power produced by

two species of monoenergetic ions is

$$\frac{P_{\text{fusion}}}{n_e^2} = g_1 g_2 \sigma v_{1-2} \mathcal{E}. \quad (3.5)$$

The fuel mixture ratio is defined as $M \equiv \frac{n_1}{n_2}$. For a fully neutralized plasma, the condition $n_e = \sum_i Z_i n_i$ must be met, resulting in the ratios $g_1 = \frac{1}{Z_1 + Z_2/M}$ and $g_2 = \frac{1}{M Z_1 + Z_2}$. It is assumed for the following analysis that all fuels will be fully ionized. The ratio of fusion power to Bremsstrahlung power is calculated by finding the electron temperature that satisfies Eq. 3.2, and then calculating $P_{\text{fusion}}/P_{\text{Brem}}$. This power ratio is given for different mixture ratios of different fuels in Fig. 3.1. At the 150 keV resonance, Bremsstrahlung radiation power exceeds fusion power, however, at the 600 keV resonance, fusion power is approximately three times that of Bremsstrahlung power at a mixture ratio of 5.4 parts hydrogen to 1 part boron and an electron temperature of $T_e = 120$ keV.

For the CE-IEC, Bremsstrahlung not only happens during the fusion-producing counterstreaming inside the core, but also when the ion bunches are within the electron neutralization area and approaching the core. The distance over which they would travel through electrons would need to be limited to the size of the fusion core. Alternatively, conversion of Bremsstrahlung radiation energy would loosen these limitations.




Figure 3.1: The ratio of fusion power to bremsstrahlung radiation power for various mixture ratios of monoenergetic ions and thermally equilibrated electrons. n_1 is the first species of each fuel as written in the legend, and n_2 is the second. The dependence of the power ratio on electron density only occurs in the Coulomb logarithm, and so changing the electron density has little effect (increases in the electron density moves the power ratios to slightly more favorable values.) An electron density of 10^{22} m^{-3} was chosen for this plot. Maxima occur at the following points: (a) $p\text{-}^{11}\text{B}$ (150 keV), $M = 7.4$, $T_e = 54.4 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 0.3$; (b) $p\text{-}^{11}\text{B}$ (600 keV), $M = 5.4$, $T_e = 120 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 3.1$; (c) $\text{D-}^3\text{He}$, $M = 2.2$, $T_e = 56.2 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 24.7$; (d) D-T , $M = 0.9$, $T_e = 19.6 \text{ keV}$, $\frac{P_{\text{fusion}}}{P_{\text{Brem}}} = 430$.

3.3: Space-charge limitation of ion bunch density in the non-neutralized regions

Outside of the fusion core, there is no electron neutralization, so the density of the ions in these regions is space-charge limited. Space-charge causes the ions to warp the acceleration electric field along the beamline and also causes expansion of the ion bunches towards the electrode walls.

3.3.1 Limitation of density due to bunch expansion parallel to the beamline

The maximum bunch density limited by expansion parallel to the direction of acceleration of the fuel ions is analyzed using the Child-Langmuir current law:

$$nv = \frac{4}{9}\epsilon_0 \sqrt{\frac{2}{qm}} \frac{V^{3/2}}{d^2} \quad (3.6)$$

For protons, using $V \approx 550$ kV and $d \approx 1$ m, the limitation is $nv \approx 10^{21} \text{ m}^{-2}\text{s}^{-1}$. At a velocity of 10^6 m/s, the maximum density that can be accelerated by the IEC is approximately 10^{15} .

3.3.2 Limitation of density due to bunch expansion transverse to the beamline

The maximum bunch density is also limited by the transverse expansion in the acceleration and turnaround regions. Eq. (A.16) is the time for expansion of the bunch from an initial to final radius:

$$\tau = \sqrt{\frac{3 m_i \epsilon_0}{2 q_i^2 n}} \left\{ \frac{r_\tau}{r_0} \sqrt{1 - \frac{r_0}{r_\tau}} + \frac{1}{2} \ln \left[2 \frac{r_\tau}{r_0} \left(1 + \sqrt{1 - \frac{r_0}{r_\tau}} \right) - 1 \right] \right\}. \quad (3.7)$$

The turn-around time for an ion bunch in a 1 meter IEC is on the order of 10^{-7} s. For a bunch size on the order of 5 cm radius, an expansion of up to 10 cm may be acceptable, which limits the bunch density to approximately 10^{14} . This limit could be increased by the axial magnetic field, but likely not to more than 10^{15} .

The estimated limit of 10^{15} is for the acceleration and turnaround regions. Methods of increasing the density limit above 10^{15} include a decrease in acceleration distance, and a neutralization of the acceleration region (similar to the Multiple Ambipolar Recirculating Beam Line Experiment [12]).

3.4: Limitation on core density due to the two-stream instability

While counter-streaming ion bunches are passing through one another in the fusion core, they can be analyzed as uniform counter-streaming ion beams over

a time scale of $t = v_0/R_c$ where v_0 is the peak cross-section velocity and R_c is the radius of the fusion core. A two-stream instability grows exponentially, and so a situation in which the argument of the exponential is above unity should be considered problematic, while if it is below unity the instability should not be so significant that the perturbations are not smoothed out during the transit of the bunch out to the turnaround point and back inwards once again.

3.4.1 Derivation of the two-stream instability dispersion relation

Consider two streams of identical ions with uniform densities (n_0) moving with opposite velocities ($\pm v_0$) in one dimension. A perturbation is introduced in the density, velocity, and electric field with a temporal frequency ω [rad/s] and spatial frequency k [rad/m]. The functions for the densities, velocities, and electric field for the $-v_0$ and $+v_0$ populations are

$$n_- = n_0 + \tilde{n}_- e^{i(kx - \omega t)} \quad (3.8a)$$

$$n_+ = n_0 + \tilde{n}_+ e^{i(kx - \omega t)} \quad (3.8b)$$

$$v_- = -v_0 + \tilde{v}_- e^{i(kx - \omega t)} \quad (3.8c)$$

$$v_+ = v_0 + \tilde{v}_+ e^{i(kx - \omega t)} \quad (3.8d)$$

$$E = \tilde{E} e^{i(kx - \omega t)} \quad (3.8e)$$

where $(\tilde{n}_-, \tilde{n}_+) \ll n_0$ and $(\tilde{v}_-, \tilde{v}_+) \ll v_0$. The equations of mass conservation, momentum conservation, and electric field are

$$\frac{\partial}{\partial t} n_- + \frac{\partial}{\partial x} (n_- v_-) = 0 \quad (3.9a)$$

$$\frac{\partial}{\partial t} n_+ + \frac{\partial}{\partial x} (n_+ v_+) = 0 \quad (3.9b)$$

$$\frac{\partial}{\partial t} v_- + v_- \frac{\partial}{\partial x} v_- = \frac{q}{m} E \quad (3.9c)$$

$$\frac{\partial}{\partial t} v_+ + v_+ \frac{\partial}{\partial x} v_+ = \frac{q}{m} E \quad (3.9d)$$

$$E = \frac{q}{\epsilon_0} (n_- + n_+). \quad (3.9e)$$

Inserting the values of Eqs. 3.8 into Eqs. 3.9, discarding second-order small terms, and simplifying, results in

$$\tilde{n}_- = \frac{kn_0}{\omega + kv_0} \tilde{v}_- \quad (3.10a)$$

$$\tilde{n}_+ = \frac{kn_0}{\omega - kv_0} \tilde{v}_+ \quad (3.10b)$$

$$\tilde{v}_- = \frac{iq/m}{\omega + kv_0} \tilde{E} \quad (3.10c)$$

$$\tilde{v}_+ = \frac{iq/m}{\omega - kv_0} \tilde{E} \quad (3.10d)$$

$$\tilde{E} = \frac{q/\epsilon_0}{ik} (n_- + n_+). \quad (3.10e)$$

Inserting Eqs. (3.10c) and (3.10d) into Eqs. (3.10a) and (3.10b) respectively and then inserting the resulting forms of Eqs. (3.10a) and (3.10b) into (3.10e) results in

the dispersion relation of the counter-streaming instability:

$$1 = \frac{q^2 n_0}{\epsilon_0 m} \left(\frac{1}{(\omega + kv_0)^2} + \frac{1}{(\omega - kv_0)^2} \right) \quad (3.11)$$

where the coefficient is the square of ion plasma frequency $\omega_i^2 \equiv \frac{q^2 n_0}{\epsilon_0 m}$. Eq. 3.11

produces four solutions for $\omega(k)$:

$$\omega(k) = \pm \omega_i \sqrt{\left(\frac{kv_0}{\omega_i} \right)^2 + 1} \pm \sqrt{4 \left(\frac{kv_0}{\omega_i} \right)^2 + 1} \quad (3.12)$$

The signs of ω and k only correspond to a phase difference (thus the symmetry among the four quadrants of the ω - k plane), and so there are two unique solutions of interest, a purely real solution, ω_r and a complex solution, ω_c :

$$\omega_r = \omega_i \sqrt{\left(\frac{kv_0}{\omega_i} \right)^2 + 1} + \sqrt{4 \left(\frac{kv_0}{\omega_i} \right)^2 + 1} \quad (3.13a)$$

$$\omega_c = \omega_i \sqrt{\left(\frac{kv_0}{\omega_i} \right)^2 + 1} - \sqrt{4 \left(\frac{kv_0}{\omega_i} \right)^2 + 1} \quad (3.13b)$$

A plot of ω_r and the real and imaginary components of ω_c are shown in Fig. 3.2.

The maximum of the imaginary component of ω is $\frac{i}{2}\omega_i$ at $k = \frac{\sqrt{3}\omega_i}{2v_0}$. This defines the maximum instability growth rate, which from Eqs. 3.8 goes as $e^{(ikx - i\omega t)}$.

Discarding phase information, the growth rate is $e^{(\frac{\omega_i}{2}t)}$, and making the substitution for ω_i , the maximum instability occurs at

$$\omega_{\max} = \frac{1}{2} \sqrt{\frac{q^2 n_0}{\epsilon_0 m}} \quad \text{at} \quad k_{\max} = \frac{\sqrt{3}}{2v_0} \sqrt{\frac{q^2 n_0}{\epsilon_0 m}} \quad (3.14)$$




Figure 3.2: Dispersion relation $\omega(k)$

3.4.2 Application of the dispersion relation to a density constraint

For timescales of $t < \omega_{\max}^{-1}$ the growth rate of an instability is considered to be negligible, and for $t > \omega_{\max}^{-1}$ a small perturbation of spatial frequency near $k_m a x$ will cause a significant disruption of the counter-streaming state. Thus the time constraint on the counter-streaming state is

$$t < 2\sqrt{\frac{\epsilon_0 m}{q^2 n_0}}. \quad (3.15)$$

Using the relation $t = R_c/v_0$, the upper limit on the allowed density of the ion bunches:

$$n_0 < 4 \frac{\epsilon_0 m v_0^2}{R_c^2 q^2}. \quad (3.16)$$

For a fusion core radius of $R_c = 5$ cm, the counter-streaming time is about 5 ns.

This puts a limit on bunch density of approximately $n < 10^{16}$ with the strongest

instability occurring at $\omega \approx 10^8$ rad/s and $k \approx 10^4$ rad/m, corresponding to an instability wavelength of about 0.5 mm.

3.5: Coulomb collisions

Coulomb collisional effects can be split into two types: the continuously occurring low-angle collisions that are principally responsible for thermalization, and the rare-event high-angle collisions between two ions that can suddenly and radically change the trajectory of a particle. Both kinds of collisions push ions off of the desired trajectories by some amount. It is assumed that the ions only collide while streaming through the device center. The opening angle of the beamline from the device center as measured from the beam axis is approximately 0.3 rad (17°), so scatters in this range are considered problematic.

3.5.1 High probability, low-angle Coulomb collisions

Literature on low-angle collisional processes is vast and well-established [13]. The rate of change of momentum in the x -dimension for an ion encountering a counter-streaming beam is

$$m \frac{dv_x}{dt} = n \left(\frac{q_1 q_2}{4\pi\epsilon_0} \right)^2 \frac{4\pi \log \Lambda}{\mu v^2} \quad (3.17)$$

It is assumed that the center-of-mass frame of the collisions is that of the device, and the scope is limited to investigating only small cumulative changes in angle over each pass through the fusion core so that energy exchange is negligible. In

this way, $\frac{dv_x}{dt} \rightarrow \frac{\Delta v_x}{\Delta t}$ where $\Delta t = \frac{R_c}{v}$ is the amount of time the ion bunches spend passing through the fusion core. The change in x -velocity can be expressed as $\Delta v_x = v(1 - \cos \theta)$ where θ is the average cumulative scattering angle of an ion over an amount of time Δt , and so Eq. (3.17) can be expressed as

$$\frac{mv^2(1 - \cos \theta)}{R_c} = n \left(\frac{q_1 q_2}{4\pi\epsilon_0} \right)^2 \frac{4\pi \log \Lambda}{\mu v^2} \quad (3.18)$$

which can be rearranged into an expression for the maximum allowable density:

$$n < \left(\frac{4\pi\epsilon_0}{q_1 q_2} \right)^2 \frac{\mu m v^4 (1 - \cos \theta)}{R_c 4\pi \log \Lambda} \quad (3.19)$$

Using a maximum tolerable scattering angle of $\theta = 0.03$ rad (1% of the wall angle) and a value of $\log \Lambda = 22$, the maximum allowable density is $n = 6 \times 10^{21} \text{ m}^{-3}$.

3.5.2 Low probability, high-angle Coulomb collisions

High-angle scattering events are similar to fusion events in that the interactions are binary, have a very low probability of happening to any single ion during a single pass, but nonetheless have a significant impact on device operation. Thus the ratio of the frequency of high-angle scatters to the frequency of fusion for a single ion is investigated. The frequency of a fusion event for an ion is

$$\nu_{\text{fusion}} = nv\sigma. \quad (3.20)$$

The frequency of low probability, high-angle collisions follows from the Rutherford scattering probability (Eq. 3 of Ref. [14])

$$\nu_{\theta} = \left(\frac{q_1 q_2}{4\pi\epsilon_0\mu} \right)^2 \frac{\pi n}{v^3 \tan^2(\theta/2)}. \quad (3.21)$$

The ratio of high-angle scatters to fusion events is

$$\frac{\nu_{\theta}}{\nu_{\text{fusion}}} = \left(\frac{q_1 q_2}{4\pi\epsilon_0\mu} \right)^2 \frac{\pi}{v^4 \tan^2(\theta/2) \sigma}. \quad (3.22)$$

Since these are low probability events, the maximum tolerable angle is chosen to be the wall angle of $\theta = 0.3$ rad. The 600 keV peak of $\text{p-}^{11}\text{B}$ results in a ratio of $\frac{\nu_{\theta}}{\nu_{\text{fusion}}} = 80$, i.e. for every fusion reaction there are 80 high-angle scatter events in which an ion certainly leaves the beam path. Note that this ratio is density-independent. The expected value of energy output of a fusion event is $t\mathcal{E}_{\text{fusion}}$ which accounts for the portion t of alpha particles that strike the inner surface. The expected value of energy loss due to a high angle scatter of $\theta > 0.3$ rad is $(1-t)E_{\text{COM}}$ which accounts for the portion $1-t$ of ions that may scatter into a different channel. If there were no other energy losses present, and it is assumed that the ratio of ions that scatter onto different channels rather than striking the wall is equal to t , then Q (ratio of power output to power input) would be

$$Q = \frac{\nu_{\text{fusion}}}{\nu_{\theta}} \frac{\mathcal{E}_{\text{fusion}}}{E_{\text{COM}}} \frac{1}{1-t}. \quad (3.23)$$

Setting $Q = 1$ and rearranging for t results in the minimum allowable transparency of a truncated icosahedron CE-IEC

$$t > \frac{1}{1 + \frac{\mathcal{E}_{\text{fusion}}}{E_{\text{COM}}} \frac{\nu_{\text{fusion}}}{\nu_{\theta}}} \quad (3.24)$$

which, for $\frac{\nu_{\text{fusion}}}{\nu_{\theta}} = \frac{1}{80}$ and $\theta = 0.3$ rad results in $t > 0.84$, which is independent of device size and fuel density.

3.6: Power deposited on the electrodes and thermal management

The maximum power deposited on the electrodes will happen on the inside-facing surface, where the walls meet the fusion core area. If all the fusion energy is produced at a single point at the device center, the power deposited on the walls per unit area will be

$$\frac{P_f}{A} = \frac{P}{4\pi R_i^2} \quad (3.25)$$

where R_i is the inner radius of the device. With a transparency as seen from the device center of $t \approx 0.8$, the total power radiated on the inner surface is simply $P_f = (1 - t)P$. The inner surface will radiate with power according to the Stephan-Boltzmann law

$$\frac{P_i}{A} = e\sigma T_i^4. \quad (3.26)$$

where T_i is the temperature of the inner surface, e is the emissivity of the surface, and $\sigma = 5.67 \times 10^{-8} \frac{\text{W}}{\text{m}^2\text{K}^4}$ is the Stefan-Boltzmann constant. However, some of the inner surface radiates onto other parts of the inner surface, and so a fraction approximately equal to t escapes, so that the effective emissivity of the inner surface is te , so that the actual power radiated per unit area will more accurately be

$$\frac{P_i}{A} = te\sigma T_i^4. \quad (3.27)$$

Assuming that the inner surface is of a high emissivity material and is thermally insulated from the innermost electrode, then $\frac{P_f}{A} = \frac{P_i}{A}$ and the equilibrium temperature of the inner surface will be

$$T_i = \left(\frac{P}{te\sigma 4\pi R_i^2} \right)^{\frac{1}{4}}. \quad (3.28)$$

Some fraction f_1 of the power radiated by the inner surface will impinge on the channel walls, so the power per unit area received by the walls is

$$P_w = f_1(1 - t)P. \quad (3.29)$$

Of this power absorbed by the walls, some fraction f_2 will be radiated into other parts of the walls, and so the effective emissivity of the channel walls will be $(1 - f_2)e$.

The equilibrium temperature of the walls then will be

$$T_w = \left(\frac{f_1(1 - t)P_f}{(1 - f_2)e\sigma A_w} \right)^{\frac{1}{4}}. \quad (3.30)$$

where A_w is the area of the channel walls. The edge length of a unit regular truncated icosahedron is 0.4, and so the approximate area of one wall is $(0.4)(R_0 - R_i)\frac{R_0 + R_i}{2}$. There are 90 edges, and two walls per edge, so $A_w \approx (180)(0.4)(R_0 - R_i)\frac{R_0 + R_i}{2} = 36(R_0 - R_i)(R_0 + R_i)$. For $P_f = 10^6$ W, $R_i = 0.25$ m, $R_0 = 1$ m, $e = 1$, $t = 80\%$, $f_1 = 0.5$, and $f_2 = 0.9$, the equilibrium temperatures of the inner surface and channel walls are $T_i = 2300$ K and $T_w = 850$ K. To limit thermal conduction from the inner surface to the walls, insulation could separate the two. The thermal conduction power per unit area through the insulation of thickness a and thermal conductivity k between the inner surface and the wall is

$$\frac{P_t}{A} = \frac{k(T_i - T_w)}{a}. \quad (3.31)$$

In order to limit the power transfer through the walls to a fraction f_3 of the power radiated by the inner surface requires that the length a be

$$a = \frac{k(T_i - T_w)}{f_3} \frac{4\pi R_i^2}{P} \quad (3.32)$$

which, for $f_3 = 1\%$ and $k = 0.05 \frac{\text{W}}{\text{mK}}$ requires an insulation thickness of 5 mm.

3.7: Power balance between protons, boron ions, and electrons

The power transfer between ions of species 1 and 2 at different energies is given

by [15]

$$\frac{P_{1 \rightarrow 2}}{n_e^2} = 4.208 \times 10^{-44} \frac{\sqrt{m_1 m_2} Z_1^2 Z_2^2 g_1 g_2 \log \Lambda (E_1 - E_2)}{(m_1 E_1 + m_2 E_2)^{3/2}}. \quad (3.33)$$

where E is in eV and $g_i \equiv n_i/n_e$. This is likely an overestimation of power transfer for purely counterstreaming beams, since the center-of-mass velocity in the lab frame is ideally zero, but may be a good estimator for cross-streaming beams in the CE-IEC, and so is used as an order-of-magnitude estimator here. The power balance is found by numerical integration of the following equations

$$\frac{\partial E_{p^+}}{\partial t} = \frac{-P_{p^+ \rightarrow e} - P_{p^+ \rightarrow B^{5+}}}{n_e g_{p^+}} \quad (3.34a)$$

$$\frac{\partial E_{B^{5+}}}{\partial t} = \frac{-P_{B^{5+} \rightarrow e} + P_{p^+ \rightarrow B^{5+}}}{n_e g_{B^{5+}}} \quad (3.34b)$$

$$\frac{\partial T_e}{\partial t} = \frac{P_{p^+ \rightarrow e} + P_{B^{5+} \rightarrow e} - P_{\text{Brem}}}{n_e} \quad (3.34c)$$

Using a fuel mixture ratio of $M = 5.4$, the numerical results are plotted in Fig. 3.3.

For the nominal IEC system, the time over which the ion bunches are in transit is on the order of $t = R_c/v_{\text{fusion}} \approx 5 \times 10^{-9}$ s. At an electron density of 10^{22} m^{-3} the product $tn_e = 5 \times 10^{13}$ over which the equilibration time-scale is negligible. After hundreds to thousands of passes, however, the energy transfer becomes significant, but this would be naturally mitigated by introducing new fuel ions to replaced lost




Figure 3.3: With starting energies of $E_{p^+} = 550$ keV and $E_{B^{5+}} = 50$ keV and $T_e = 0$, the ion temperatures equilibrate with one another on a faster time-scale than with the electrons. As $t n_e \rightarrow \infty$ the energies are depleted to Bremsstrahlung radiation.

or fused ions. The probability of fusion at each pass through the fusion core is

$$\mathcal{P}_1 = g_2 n_e \sigma v \quad (3.35)$$

which for $n_e = 10^{22} \text{ m}^{-3}$, $\mathcal{P}$ is on the order of 10^{-8} , meaning that on average, 10^8 passes are required for a fusion event to occur, and for break-even energy production approximately 10^7 passes of an ion are required to happen before ion loss or ion fusion. Over this many passes through the core, the product $t n_e = 5 \times 10^{20}$ now appears prohibitively large in Fig. 3.3, and so some active method of draining energy from the boron ions and transferring that energy to the proton ions is necessary, which is conceivably possible through active voltage control.

3.8: Limits on electron confinement

Effective electron confinement in the core region faces three challenges: space charge limitation when ions are not present in the core, thermal leakage of electrons along beamline point cusps, and loss of electrons to the inner surface line cusps. Lost electrons must be replaced, and the power required for replenishment is proportional to the energy of the lost electrons. The energy of the lost electrons comes from both the electron source and the energy transfer from ions to electrons, the latter of which is an issue for long electrons lifetimes (see Fig. 3.3.)

3.8.1 Space charge limitation on confined electrons

The density of the electrons in the fusion core is limited by the space charge of electrons relative to the potential difference between the confining cathode and the inner anode V_c . The potential of a sphere of electrons of radius R_i is

$$\Phi_e = \frac{R_i^2 n e}{3\epsilon_0}. \quad (3.36)$$

and so the density is limited to

$$n_e < V_c \frac{3\epsilon_0}{R_i^2 e} \quad (3.37)$$

which, for an electron confinement potential on the order of $V_c = 25$ kV, limits the electron core density to $n_e = 10^{14} \text{ m}^{-3}$ for cold electrons. Higher temperature

electrons will eject from the core more frequently.

3.8.2 Number of electrons required for complete neutralization

As opposed to all electrons staying in the core region while the ions are not present, some electrons could be newly generated in pulses synchronized with the passing ion bunches. As previously estimated, the electron density required for core neutralization for a useful fusor is $n_e \approx 10^{22} \text{ m}^{-3}$. In a fusion core of radius $R_c = 5 \text{ cm}$, the number of electrons needed is $N_e = 5 \times 10^{18}$. If each electron must be created anew each time the bunches pass through the fusion core, the power required for electron neutralization is

$$P_e = T_e N_e f \quad (3.38)$$

where T_e is in joules, and $f \approx 10^6 \text{ Hz}$ is the oscillation frequency of the device. To limit the electron neutralization power to 1 MW, the temperature of the electrons must be held to approximately 1 eV. To maintain this power loss limit with electrons of energies up to 10 keV, then 10^{14} (0.01%) of the neutralizing electrons may be lost at each pass.

3.8.3 Electron line cusp loss frequency

Assuming the electrons are prevented from escaping along the beamlines by the cathode potential, the electron loss frequency is calculated from the mirror ratio of the cusps. Theoretically the mirror ratio is infinite, since the magnetic field is zero at the device center. An effective mirror ratio can be calculated from the point

where the electrons become magnetized, which is taken to be the point at which the Larmor radius is equal to the fusion core radius, where the magnetic field will be $B = \frac{mv}{eR_c}$ and so the mirror ratio $\mathcal{R}$ is estimated as

$$\mathcal{R} = M \frac{eR_c}{m_e v_{th}} \quad (3.39)$$

And an electron is lost to a cusp if its velocity vector has an angle less than θ measured from the magnetic field, where θ is defined by

$$\sin \theta = \frac{1}{\sqrt{\mathcal{R}}} \quad (3.40)$$

The cusps in this configuration are line cusps (rather than the more commonly analyzed point cusps) so rather than a loss cone, there is the two-dimensional loss sector (or loss arc). For a random electron velocity direction, the probability $\mathcal{P}$ that the electron is in a loss cone is the area of a spherical sector of height $2 \sin \theta$ divided by the surface of the sphere, which simplifies to $\mathcal{P} = \sin \theta$. At an electron density of 10^{14} m^{-3} and at a temperature on the order of a keV or above, the electrons are collisionless, and so the frequency at which an electron has a “chance” to escape is based on its transit time across the core region, which is v_{th}/R_c . This frequency multiplied by the loss probability gives an estimation of the loss frequency:

$$\nu_{\text{loss}} = \frac{v_{th}}{R_c} \sqrt{\frac{m_e v_{th}}{e M R_c}} \quad (3.41)$$

At a temperature of $T_e = 120$ keV and magnetization of $M = 1$ T, the loss probability will be 7% and the loss frequency will be $\nu_{\text{loss}} = 6 \times 10^7 \text{ s}^{-1}$ which is slightly higher than the oscillation frequency of the ions. Lower temperatures will lower this loss rate, but 120 keV is the expected temperature for long-lifetime electrons (see the analysis in Sec. 3.2). Note that this result does not take into account space charge effects, and so is only applicable when the density of electrons is low ($n_e < 10^{14} \text{ m}^{-3}$).

3.8.4 High- β loss rate along beamline cusps

If the temperature and space charge of the electrons is high enough such that some electrons would overcome the potential barrier of the cathode, leakage of these electrons is limited by the point-cusp nature of the magnetic fields along the beamlines. The portion of ions that have enough energy to exit along the cusps is given by the Boltzmann Factor, and accounts for the space-charge potential of the electrons:

$$\mathcal{B} = \exp\left(\frac{V_c - \Phi_e}{T_e}\right) \quad (3.42)$$

where V_c is the voltage of the cathode relative to the fusion core and Φ_e is the potential spike of the electrons. In the best-case scenario, the electrons will be in a high- β state ($\beta = \frac{n_e T_e}{B^2/2\mu_0} \approx 1$) and the cusp-rate ion loss for a high- β plasma is given by [6]

$$\frac{I}{e} = \frac{\pi^2}{9} n_e v_{th} r_l^2 \quad (3.43)$$

where $r_l = \frac{m_e v_{th}}{eB}$ is the Larmor radius of electrons. In the point cusps, the strength of the magnetic field is approximately $\frac{1-t_m}{t_m}M$ and so the loss rate, accounting for the limiting potential of the cathode, is

$$\frac{I}{e} = n_e v_{th}^3 \left(\frac{m_e}{eM} \frac{t_m}{1-t_m} \right)^2 \exp \left(\frac{V_c - \Phi_e}{T_e} \right) \quad (3.44)$$

where M is the magnetization of the permanent magnets and t_m is the transparency of the device if only the permanent magnets were present (i.e. t_m is the fraction of the spherical surface area of the CE-IEC that is not magnetized and $t_m > t$.)

For $t_m = 0.9$ and $M = 1$ T, and at a temperature of $T_e = 10$ keV and $V_c - \Phi_e = 25$ kV, and a density of $n_e = 10^{22} \text{ m}^{-3}$, the electron loss rate at a point cusp is 10^{34} electrons per second. For $N_e = 5 \times 10^{18}$ the loss rate for a single electron over the 32 point cusps of the CE-IEC is $\nu = 32 \times 10^{34} / N_e = 5 \times 10^{16} \text{ s}^{-1}$. To lower this number to the ion oscillation frequency of approximately 10^7 s^{-1} , the ratio $\frac{V_c - \Phi_e}{T_e}$ must be lowered by a factor of 10.

3.9: Scaling laws of the CE-IEC

In this section, a change in the device length scale L denotes a change in each part of the device by the same ratio. The nominal scale for the preceding sections was based on $L \approx 1$.

3.9.1 Overcoming space-charge limitations by scaling down

The limitation on a human-sized IEC or larger is due to the space-charge of recirculating ions in non-neutralized regions. If the figure-of-merit for an IEC device is net power output *per unit volume* $\mathcal{P} \equiv \frac{P}{\text{Vol.}}$ rather than net power output, an IEC device can overcome space charge limitations by a reduction in size. $\text{Vol.} = \frac{4}{3}\pi R_0^3$ is the volume of the fusor. The electric potential of an unneutralized ion bunch goes as

$$\Phi_{\text{bunch}} \propto nL^2 \quad (3.45)$$

(this can be seen by taking the potential difference between the center and edge of a uniformly charged sphere: $\Phi = \frac{Q}{8\pi\epsilon_0 R_c}$ and expressing the charge as density times volume $Q = qn\frac{4}{3}\pi R_c^3$.) A space charge limitation implies that a device has a maximum ratio of bunch potential to device voltage

$$\frac{\Phi}{V} < \text{constant} \quad (3.46)$$

where the constant is likely on the order of $\frac{1}{100}$. Since V is determined by the maximum fusion cross section, it is constant, and so the potential of the bunch must also be limited by a constant, and so Eq. 3.45 can be re-written as

$$n \propto \frac{1}{L^2} \quad (3.47)$$

which is the same trend found in Eqs. (3.6) and (3.16). The fusion power per unit volume in the device core is

$$\mathcal{P}_{\text{fusion}} = n^2 v \sigma \mathcal{E} \quad (3.48)$$

where v is the center-of-mass velocity between the counter-streaming proton and boron ions, which is constant at the peak value of σ , and so Eq. (3.48) can be written as

$$\mathcal{P}_{\text{fusion}} \propto n^2 \quad (3.49)$$

and combining Eqs. (3.47) and (3.49) results in

$$\mathcal{P}_{\text{fusion}} \propto \frac{1}{L^4} \quad (3.50)$$

which can be rephrased as

$$\mathcal{P}_{\text{fusion}} L^4 = \text{constant}. \quad (3.51)$$

To reiterate, this trend holds only when space charge is the limiting factor, and works on the principal that reducing the size of an ion bunch allows for an increase in ion bunch density without an increase in space charge, such that the potential from the ion bunch remains at the same ratio to the accelerating potential at various scales. To think of it another way, decreasing the size of the CE-IEC while maintaining the same voltage increases the electric field, and decreasing the size of the ion bunch decreases the electric field of the space charge, allowing the density of the bunch more room for increase.

It should also be noted that at small sizes, the effect of the magnetic field

becomes negligible. The magnitude of the magnetic field is independent of scale, which means the Larmor radius is also independent of scale. At large scales, the ions and electrons are effectively very tightly bound to magnetic field lines, but at small scales the Larmor radius can become large relative to the scale length.

3.9.2 Scaling of energy input

Energy input, is estimated here to scale as the rate of ion loss, commonly referred to as conduction loss, or P_{cond} . The ion loss frequency per unit volume scales as the density of the ions (n) multiplied by the collision frequency of a single ion (n , Eq. 3.17) multiplied by the frequency of oscillation f (because each pass through the system is another “chance” to hit an electrode). Since $f \propto \frac{1}{L}$ and again using Eq. 3.47 the power input scales as:

$$\mathcal{P}_{\text{cond}} \propto \frac{1}{L^5} \quad (3.52)$$

which is not favorable to small scaling unless the electrostatic focusing can be improved such that ion collisions with the electrodes are an extremely rare occurrence over each pass through the system.

3.9.3 Scaling of surface erosion

The operational lifetime of a CE-IEC fusor will be limited by erosion of the surfaces due to impacts from both ionized fuel straying from beampaths as well as fusion products. The lifetime will scale as the inverse of the erosion rate relative to

the scale length of the device. The erosion rate r will scale as the sum of the fusion power and the input power per unit area at the location of the inner surface with area $A = 4\pi R_i^2$

$$r \propto \frac{P_{\text{fusion}} + P_{\text{cond}}}{A}. \quad (3.53)$$

The lifetime of the fusor is related to the erosion rate by

$$T_{\text{life}} \propto \frac{L}{r} \quad (3.54)$$

and so, making the substitution $P \equiv \mathcal{P} \text{Vol.} \propto \mathcal{P}AL$ the fusor lifetime is given by

$$T_{\text{life}} \propto \frac{1}{\mathcal{P}_{\text{fusion}} + \mathcal{P}_{\text{cond}}}. \quad (3.55)$$

No estimate is made here on the actual erosion rate due to the complexity of the process.

3.9.4 Size of a small CE-IEC with significant power density

From the simulations of Chap. 4 it was found that the maximum achievable density in the core for long-lifetime ions was on the order of $n = 10^{14} \text{ m}^{-3}$. This results in a fusion power density of approximately $\mathcal{P}_{\text{fusion}} = 10^{-6} \frac{\text{W}}{\text{m}^3}$. Using Eq. 3.51, the CE-IEC size $\tilde{L}$ required for a power density of $\tilde{\mathcal{P}}_{\text{fusion}} = 10^6 \frac{\text{W}}{\text{m}^3}$ (one megawatt per cubic meter) is given by the relation

$$\tilde{L} = L \left(\frac{\mathcal{P}_{\text{fusion}}}{\tilde{\mathcal{P}}_{\text{fusion}}} \right)^{\frac{1}{4}} \quad (3.56)$$

which results in $\tilde{L} = 1$ mm.

3.9.5 Structural limitations of a small CE-IEC

Practical reduction of the size of a fusor is limited chiefly by three possible factors. To analyze these factors, diamond is proposed as an inter-electrode insulator within the CE-IEC walls due to its high compressive strength and high dielectric strength.

The first possible limitation is that the electric force between electrodes will cause structural failure of the fusor at small scales. The force per unit area between two electrodes within an CE-IEC wall is approximated as the force between two parallel electrodes

$$F = 2\epsilon_r\epsilon_0 \frac{V^2}{d^2} \quad (3.57)$$

where ϵ_r is the dimensionless relative permittivity of the inter-electrode material and $d \propto L$ is the space between the electrodes. For a wall thickness of 0.04 radians, $d \approx 0.04L/4$, and for $L = 1$ mm, and using an inter-electrode medium of diamond ($\epsilon_r \approx 7$) the attractive force between the electrodes is $F = 300$ GPa (gigapascals) whereas the maximum pressure of diamond is 600 GPa, so it appears that inter-electrode pressure does not immediately make a 1 mm fusor impossible.

The second limitation is the dielectric breakdown of the interelectrode medium.

This could be theoretically limited by operating only in the vacuum of space, as well as using a very high dielectric strength material (e.g. diamond) as an insulator. The

electric field between electrodes in a CE-IEC wall will be

$$E = \frac{V}{d} \quad (3.58)$$

which, again using diamond, is $5 \times 10^{10} \frac{\text{V}}{\text{m}}$, whereas the dielectric strength of diamond is $2 \times 10^9 \frac{\text{V}}{\text{m}}$, suggesting that a fusor of this size would cause dielectric breakdown of the diamond spacing.

The third limitation is that the manufacturing of a very small fusor could be limited by the precision of the manufacturing process.

3.9.6 Lawson criterion estimation

The Lawson criterion is met when the electric power generated by the fusor and energy converter exceeds the power required to operate the fusor. The Lawson criterion can be estimated as

$$\mathcal{P}_{\text{net}} = (t \eta_{\text{DEC}} \mathcal{P}_{\text{fusion}} - \mathcal{P}_{\text{cond}} - \mathcal{P}_{\text{brem}}) \quad (3.59)$$

where once again $\mathcal{P}$ is power density, t is the transparency of the device as viewed from the center point, and η_{DEC} is the efficiency of the direct energy converter.

The criterion is met when $\mathcal{P}_{\text{net}} > 0$. From section 3.2 it was found that the best case scenario for bremsstrahlung radiation is $\mathcal{P}_{\text{brem}} = \frac{\mathcal{P}_{\text{fusion}}}{3}$. Eq. 3.59 can then be expressed as

$$\mathcal{P}_{\text{net}} = \left[\left(t \eta_{\text{DEC}} - \frac{1}{3} \right) \mathcal{P}_{\text{fusion}} - \mathcal{P}_{\text{cond}} \right]. \quad (3.60)$$

In order for $t\eta_{\text{DEC}} - \frac{1}{3} > 0$, for $t = 0.8$ the energy conversion efficiency η_{DEC} must be at least 42%. A suboptimal Standing Wave Direct Energy Converter was shown in Chap. 8 to have a conversion efficiency of 50% for mono-energetic α -particles.

At 50% conversion efficiency, the raw fusion power output of the core would only need to exceed the input power by a factor of 10.

Unfortunately, from simulations that will be discussed in Chap. 4, space charge limits the fusion power output of the current approach on the order of a microwatt. These issues would need to be addressed if any variants on the current approach are to yield success.

Chapter 4

Particle-in-cell Modeling

A 2D3V (two spatial dimensions and three velocity dimensions) axisymmetric particle-in-cell (PIC) simulation was created for the Continuous Electrode Inertial Electrostatic Confinement (CE-IEC) Fusor. The simulation models one half of a single beamline, approximated as axisymmetric, and is run in parallel on a general purpose graphics processing unit (GPU) for fast execution, enabling both high-resolution simulation as well as optimization.

4.1: Domain

The simulation domain exists in two-dimensional axial-radial cylindrical coordinates with azimuthal symmetry assumed. The axis of symmetry extends along the center of a single IEC beamline. Ions move primarily along the axial dimension (x). The radial dimension (r) is transverse to the beamline center, and is not to be

confused with the spherically radial dimension of a three-dimensional IEC. Planar symmetry exists at $x = 0$ where the ion bunches pass through the device center.

A single channel of the continuous electrode IEC has greatest width (radial extent) at the outer radius (axial extent) and tapers down to a minimum width at the inner radius, where it then opens up in the central fusion region. The angle of the domain boundary wall is calculated as the angle of the wall of a pentagonal channel aligned with the x -axis where it intersects with the x - y plane when the wall thickness is 0.08 radians. This geometry is represented on a structured grid by increasing the grid spacing both axially and radially with increased x . In the IEC, the ions tend to be more spread out in the turnaround region near the outer radius, and so a lower grid resolution is needed in this area. This is contrasted with the fusion core region (near the axis origin) in which the grid resolution is greatest. An extra region of cells is added in the radial direction to emulate the open region in the center of the IEC.

For the cell spacing formulae below, indices i and j refer to the axial index of the axial and radial address of the cell respectively, and the coordinates x and r are the axial and radial locations of the cell nodes respectively. Index values start at zero, with the origin point at $i = 0$ and $j = 0$.

4.1.1 Axial cell spacing

A non-constant spacing in the axial dimension is used to avoid the overuse of computational cells in the turnaround region where the inter-particle spacing is

generally larger. The cell spacing algorithm is required to be simple to calculate, and the inverse calculation (finding the non-integer cell location of a particle) should not be computationally intensive and should not require a lookup table. The cell spacing formula chosen is:

$$x_i = x_b + k ((ci + 1)^2 - 1) \quad (4.1)$$

where x is the location of cell i , x_b is the value at which the cell spacing becomes non-constant, and k and c are constants that determine the scale and the rate of change of cell spacing respectively. The user inputs the desired cell spacing for the beginning as well as the end of this region, along with the beginning and end points, and k and c are found using MATLAB's *lsqnonlin* function to match the beginning and end cell spacings as close to those specified by the user as possible, while maintaining the exact endpoints specified by the user.

4.1.2 Radial cell spacing

To generate the angled wall of the IEC beampath without losing the structured nature of the grid, the radial cell spacing is a function of axial position. The function for radial cell spacing is

$$r(i, j) = \begin{cases} x_i \frac{j}{N_r - 1} \tan(\theta) & x_i > x_b \\ x_b \frac{j}{N_r - 1} \tan(\theta) & x_i \leq x_b \end{cases} \quad (4.2)$$

where N_r is the number of radial cell locations and θ is the angle of the wall with respect to the axis of symmetry, which is set to 17.7° . For a given location x_i , the radial cell spacing is constant, and will be denoted as Δr_i . The formula for Δr_i follows:

$$\Delta r_i = \begin{cases} x_i \frac{1}{N_r - 1} \tan(\theta) & x_i > x_b \\ x_b \frac{1}{N_r - 1} \tan(\theta) & x_i \leq x_b \end{cases} \quad (4.3)$$

4.1.3 Cell volumes

The three-dimensional volume of each cell is found by extruding each cell around the axis of symmetry. Eqs. (4.1) and (4.2) specify the cell locations. The cell locations are typically the center of each cell, except for the boundary cells in which case the cell location is on the domain boundary. Calculation of cell volumes requires the positions of cell boundaries as well. Evaluating Eqs. (4.1) and (4.2) at the half-index values results in the cell boundary locations. These boundaries are used to determine if a particle is inside of a particular cell. The volume of each cell is a sum of the four or fewer “sub-cells” that make up each cell. A sub-cell is only used for finding cell volumes, and is made by extending sub-cell boundaries from the cell location point to the cell boundary lines. The volume of the sub-cell is given by the four nodes that make up the corners of the sub-cell, where the four nodes, clockwise from the bottom-left corner, are (x_1, r_{1a}) , (x_1, r_{1b}) , (x_2, r_{2b}) , and (x_2, r_{2a}) :

$$\text{Vol}_{\text{sub}} = \frac{\pi (x_2 - x_1)}{3} (r_{2b}^2 + r_{1b}^2 - r_{2a}^2 - r_{1a}^2 + r_{2b}r_{1b} - r_{2a}r_{1a}) \quad (4.4)$$




Figure 4.1: Cell locations (dots) and cell boundaries (lines) for the particle-in-cell domain. A low number of cells is used for this figure for the purpose of clear illustration. The resolution used in the simulation is about four times greater, for a factor of 16 increase in the number of cells as compared to this figure.

and the volume of a cell is $\text{Vol} = \sum \text{Vol}_{\text{sub}}$. The cell volumes are used for calculating charge density in the particle-in-cell simulation.

4.2: Particle-in-cell algorithm and parallelization

Setup of the simulation domain and initial parameters is performed in the MATLAB language and environment. The time-stepping portion of the simulation is written in C and is compiled and executed by MATLAB using the MEX (MATLAB executable) interface. The C routine contains all memory allocation on the GPU and all transfer of memory between the CPU and GPU. During each time-step, the CPU manages calls to CUDA kernels, which execute functions on the particles and cells using the GPU's processors.

4.2.1 Particle-to-cell interpolation to find charge density

Particle charges are deposited at the cell centers using linear interpolation. Interpolation in the x -dimension is straightforward. Interpolation in the r -dimension is accomplished using the “cylindrical cloud-in-cell” linear interpolation from Ruyten [16].

For particle p located at (x_p, y_p) between nodes i and $i + 1$ in the x -dimension and between nodes j and $j + 1$ in the r -dimension, the weighting in each dimension determines the portion of the particle that is scattered towards the i side or j side respectively, and are calculated as follows

$$w_x = \frac{x_{i+1} - x_p}{x_{i+1} - x_i} \quad (4.5a)$$

$$w_r = \left(\frac{r_{i,j+1} - r_p}{\Delta r_i} \right) \left(\frac{3 + r_{i,j}/r_p}{4} \right) \quad (4.5b)$$

and so the contribution of particle p to the charge density at the four nearest cells is

$$\begin{bmatrix} \rho_{i,j} & \rho_{i,j+1} \\ \rho_{i+i,j} & \rho_{i+i,j+1} \end{bmatrix}_p = q_p \begin{bmatrix} \frac{w_x w_r}{\text{Vol}_{i,j}} & \frac{w_x(1-w_r)}{\text{Vol}_{i,j+1}} \\ \frac{(1-w_x)w_r}{\text{Vol}_{i+1,j}} & \frac{(1-w_x)(1-w_r)}{\text{Vol}_{i+1,j+1}} \end{bmatrix} \quad (4.6)$$

where q_p is the charge of the particle (accounting for both the macroparticle weighting and the ionization level). This task is parallelized by particle, which could result in a “race condition” whereby two or more processes read the same value and attempt to increment that value one after another, but the second process reads the original value before the first process incremented it, and then writes a new value that does not contain the incrementation applied by the first process. The net effect




Figure 4.2: Modification of the discrete Poisson equation on a skewed grid

is that the algorithm will sum the contributions of processes to an erroneously low value. To overcome this, the *atomicAdd* CUDA function is used for each evaluation of Eq. 4.6 so that the contribution of each particle p to the density ρ is correctly summed without errors due to the race condition.

4.2.2 Calculation of electric potential from charge density

The electric potential is found through discretization and Jacobian iteration of the axisymmetric form of Poisson's equation:

$$\frac{\partial^2 \Phi}{\partial x^2} + \frac{1}{r} \frac{\partial \Phi}{\partial r} + \frac{\partial^2 \Phi}{\partial r^2} = -\frac{\rho}{\epsilon_0} \quad (4.7)$$

The skewed nature of the grid requires modification of the calculation of the x -derivative, since the cell centers are no longer aligned in the x -dimension. The x -derivative requires the two closest cells in the x -direction. As shown in Fig. 4.2,

the second derivative of x can be approximated as

$$\left[\frac{\partial^2 \Phi}{\partial x^2} \right]_{i,j} = \frac{1}{\Delta x^2} \left(-2\Phi_{i,j} + \frac{a\Phi_{i-1,j} + b\Phi_{i-1,j+1}}{a+b} + \frac{c\Phi_{i+1,j} + d\Phi_{i+1,j-1}}{c+d} \right) \quad (4.8)$$

where a , b , c , and d are the distances illustrated in Fig. 4.2, and it is assumed (for this expression only) that $\Delta x = x_{i+1} - x_i = x_i - x_{i-1}$. Further modification of Poisson's equation is warranted by the non-uniform spacing in the x -dimension of the cells. The modification of the second derivative with non-uniform grid spacing is given by Sfakianakis [17]:

$$\begin{aligned} \left[\frac{\partial^2 \Phi}{\partial x^2} \right]_i &= -\frac{2}{(x_{i+1} - x_i)(x_i - x_{i-1})} \Phi_i + \frac{2}{(x_i - x_{i-1})(x_{i+1} - x_{i-1})} \Phi_{i-1} \\ &\quad + \frac{2}{(x_{i+1} - x_i)(x_{i+1} - x_{i-1})} \Phi_{i+1} \end{aligned} \quad (4.9)$$

By combining the effects of Eq. 4.8 and Eq. 4.9, the discrete form of Eq. 4.7 becomes

$$\begin{aligned} &\frac{2}{(x_{i+1} - x_i)(x_i - x_{i-1})} \Phi_i + \frac{2}{(x_i - x_{i-1})(x_{i+1} - x_{i-1})} \frac{a\Phi_{i-1,j} + b\Phi_{i-1,j+1}}{a+b} \\ &+ \frac{2}{(x_{i+1} - x_i)(x_{i+1} - x_{i-1})} \frac{c\Phi_{i+1,j} + d\Phi_{i+1,j-1}}{c+d} + \frac{1}{r_{i,j}} \frac{\Phi_{i,j+1} - \Phi_{i,j-1}}{2\Delta r_i} \\ &+ \frac{\Phi_{i,j} - \Phi_{i,j-1} - \Phi_{i,j+1}}{\Delta r_i^2} = -\frac{\rho_{i,j}}{\epsilon_0} \end{aligned} \quad (4.10)$$

On the $x = 0$ and $r = 0$ boundaries, the derivative of the potential perpendicular to the boundary vanishes due to symmetry (Neumann boundary conditions). This is enforced by replacing $\Phi_{-1,j}$ with $\Phi_{1,j}$, and $\Phi_{i,-1}$ with $\Phi_{i,1}$ in Eq. 4.10 on nodes where

$i = 0$ and/or $j = 0$. On the other (Dirichlet) boundaries, the boundary potentials are known values and do not need to be solved for. The coefficients in Eq. 4.10 can be rewritten as

$$A\Phi_{i-1,j} + B\Phi_{i-1,j+1} + C\Phi_{i,j-1} + D\Phi_{i,j} + E\Phi_{i,j+1} + F\Phi_{i+1,j} + G\Phi_{i+1,j+1} = -\frac{\rho_{i,j}}{\epsilon_0} \quad (4.11)$$

At each location (i, j) , the coefficients of Eq. 4.11 makes of a row of the linear system

$$\mathbf{A}\vec{\Phi} = \vec{b} \quad (4.12)$$

where $\vec{\Phi}$ is a column vector of the unknown potentials of the cells in the simulation domain, $\mathbf{A}$ is the matrix of coefficients, and $\vec{b}$ is the source term $-\rho/\epsilon_0$ added to any known boundary (Dirichlet) potentials. Solution of Eq. 4.12 is accomplished using a Jacobian iteration method, which is chosen due to its straightforward parallelization. An initial guess of the potential $\vec{\Phi}^0$ is chosen, and the potential is solved for iteratively:

$$\vec{\Phi}^{k+1} = \mathbf{D}^{-1}\vec{b} - (\mathbf{D}^{-1}\mathbf{N}) \vec{\Phi}^k \quad (4.13)$$

where k is the iteration index, $\mathbf{D}$ is the diagonal of $\mathbf{A}$ and $\mathbf{N}$ is the non-diagonal part of $\mathbf{A}$ so that $\mathbf{A} = \mathbf{D} + \mathbf{N}$. Both $\mathbf{D}^{-1}$ and $\mathbf{D}^{-1}\mathbf{N}$ can be precalculated. Each row of $\mathbf{D}^{-1}\mathbf{N}$ has at most 7 entries, so the maximum amount of memory needed is on the order of 7 multiplied by the number of cells, which is easily achievable on a GPU.

Each row of Eq. 4.13 is evaluated in parallel. After each iteration, the GPU must be synchronized (all processes allowed to complete) so that at the next iteration

each process has access to the updated data. Since many iterations of Eq. 4.13 must be performed, this part of the PIC algorithm takes a significant portion of the computation time.

4.2.3 Calculation of electric field from electric potential

The electric field components are calculated in a similar manner to the electric potential, but the method is explicit rather than implicit, and therefore faster and more simple. The electric field is the negative gradient of the potential

$$\begin{bmatrix} E_x \\ E_r \end{bmatrix} = - \begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial r} \end{bmatrix} \Phi \quad (4.14)$$

Once again, because of the non-uniform cell spacing in x and the skewed nature of the grid, the numerical derivatives must be modified using the numerical first derivative of Sfakianakis [17] and the same skewed grid modification as Eq. 4.8

$$\begin{aligned} [E_x]_{i,j} = & -\frac{a\Phi_{i-1,j} + b\Phi_{i-1,j+1}}{a+b} \frac{x_{i+1} - x_i}{(x_i - x_{i-1})(x_{i+1} - x_{i-1})} \\ & + \Phi_{i,j} \left(\frac{1}{x_i - x_{i-1}} - \frac{1}{x_{i+1} - x_i} \right) \\ & + \frac{c\Phi_{i+1,j} + d\Phi_{i+1,j-1}}{c+d} \frac{x_i - x_{i-1}}{(x_{i+1} - x_i)(x_{i+1} - x_{i-1})} \end{aligned} \quad (4.15a)$$

$$[E_r]_{i,j} = \frac{\Phi_{i,j+1} - \Phi_{i,j-1}}{2\Delta r_i} \quad (4.15b)$$

Eqs. 4.15a are easily parallelized by cell on the GPU and so this part of the PIC algorithm takes only a small portion of the simulation time.

4.2.4 Cell-to-particle interpolation of electric and magnetic field

The acceleration of particles due to the electric and magnetic fields are interpolated from the cell values first by recalling the particle weights from Eq. 4.5, then weighting the fields to the particles:

$$\begin{bmatrix} a_{E_x} \\ a_{E_r} \\ a_{B_x} \\ a_{B_r} \end{bmatrix}_p = \frac{q_p}{m_p} \left(w_x w_r \begin{bmatrix} E_x \\ E_r \\ B_x \\ B_r \end{bmatrix}_{i,j} + (1 - w_x) w_r \begin{bmatrix} E_x \\ E_r \\ B_x \\ B_r \end{bmatrix}_{i+1,j} + w_x (1 - w_r) \begin{bmatrix} E_x \\ E_r \\ B_x \\ B_r \end{bmatrix}_{i,j+1} + (1 - w_x)(1 - w_r) \begin{bmatrix} E_x \\ E_r \\ B_x \\ B_r \end{bmatrix}_{i+1,j+1} \right) \quad (4.16)$$

Like in Sec. 4.2.1, there is an issue of multiple processes attempting to access the same data on GPU memory, but in this case it is read access. No explicit coding is necessary to resolve this conflict, and the GPU performs this part of the PIC algorithm quite quickly.

4.2.5 Particle position and velocity updates

Particle positions and velocities are updated at each timestep, in a method equivalent to the leapfrog method with constant value time-steps. Particle velocities

are first updated using the Boris method [18] for particle movement in a magnetic field and are then updated using the Birdsall method [19] for moving particles in curvilinear coordinates. Together, these methods are as follows, where a superscript k refers to the time-step.

$$\mathbf{v}_m = \mathbf{v}^k + \frac{\Delta t}{2} \mathbf{a}_E \quad (4.17a)$$

$$\mathbf{v}_t = \frac{\Delta t}{2} \mathbf{a}_B \quad (4.17b)$$

$$\mathbf{v}_c = \mathbf{v}_m - (\mathbf{v}_m \times \mathbf{v}_t) \quad (4.17c)$$

$$\mathbf{v}_s = \frac{2\mathbf{v}_t}{1 + |\mathbf{v}_t|^2} \quad (4.17d)$$

$$\mathbf{v}_l = \mathbf{v}_m + (\mathbf{v}_c \times \mathbf{v}_s) \quad (4.17e)$$

$$\mathbf{v}'^{k+1} = \mathbf{v}_l + \frac{\Delta t}{2} \mathbf{a}_E \quad (4.17f)$$

$$x' = x_r^k + v_r'^{k+1} \Delta t \quad (4.17g)$$

$$y' = v_\theta'^{k+1} \Delta t \quad (4.17h)$$

$$r' = \sqrt{x'^2 + y'^2} \quad (4.17i)$$

$$\theta = \sin^{-1} \frac{y'}{r'} \quad (4.17j)$$

$$\mathbf{v}^{k+1} = \begin{bmatrix} v_x'^{k+1} \\ \cos \theta v_r^k + \sin \theta v_\theta^k \\ -\sin \theta v_r^k + \cos \theta v_\theta^k \end{bmatrix} \quad (4.17k)$$

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \mathbf{v}^{k+1} \Delta t \quad (4.17l)$$

Eqs. 4.17a are independent for each particle, and so are easily parallelized, and this part of the PIC algorithm takes a small portion of overall computation time.

4.2.6 Particle-particle collision modeling

Collisions are modeled in the pair-matching Monte-Carlo scheme of Takizuka and Abe [20]. Particles are sorted into cells (performed in the density scattering of Sec. 4.2.1) and each cell is assigned a thread for GPU execution. First the particle list in each cell is shuffled into a random order using the Fisher-Yates algorithm. A loop through each neighboring pair of particles in the list is performed, and the scattering angle θ is calculated using the method outlined in Chap. 7. The relative velocity v_{rel} for calculating θ is the velocity difference of the particle pair, and the density n used for the calculation of θ is the density of the cell, so that in this way each particle gets a random sampling of the velocity space of the cell and over many time steps the collisional effects are approximately integrated over the entire velocity space. Once θ is calculated, a random azimuthal angle ϕ is generated uniformly between 0 and 2π . In the center-of-mass frame of the particle pair, one particle has its velocity changed by these two angles and the change in velocity of the other particle is calculated such that the post-collision momentum of the pair is unchanged. The changes in velocity are then applied back to the laboratory frame.

4.2.7 Particle-boundary interactions

Particles that cross the domain boundary on either the axis symmetry at $r = 0$ or the plane of symmetry at $x = 0$ are reflected (as if “bouncing” off of these boundaries) by checking each particle for a negative position value in each dimension, and in the case of a negative value, changing it to positive value, as well

as switching the sign of the velocity in that dimension.

Particles that cross any other boundaries are removed from the domain, and the kinetic energy of the removed particles is summed for the calculation of power input of the fusor. Lost particles are replaced at the beginning of the next period, in the fusion core with fusion velocity.

Calculation of which particles need to be removed from the domain is accomplished by parallel process on the GPU. The algorithm for removing particles from the simulation requires a GPU-to-CPU memory transfer of boolean values. The CPU then loops through the array of particles, and upon encountering a particle in need of removal, replaces that particle's data on the GPU with the data of the last active particle in the array via the *cudaMemcpyDeviceToDevice* option in the *cudaMemcpy* function.

4.3: Fusion calculation

The pair-matching algorithm used for collision modeling doubles as a fusion calculation tool. The contribution of each particle pair to the fusion rate is calculated using the relative velocity of the particles, the ion densities in the cell, and the number of macroparticles in the cell using the fusion rate equation for fusion between species i and j

$$R_{\text{fusion}} = \frac{n_i n_j \sigma(v_{\text{rel}}) v_{\text{rel}}}{N_p} \left[\frac{\#}{\text{m}^3 \text{s}} \right] \quad (4.18)$$

where N_p is the number of particle pairs in the cell contributing to fusion so that the fusion rate is averaged over the each pair. The contribution from each pair in the cell

is summed and then multiplied by the volume of the cell, then the contribution from each cell is summed and multiplied by the energy per fusion reaction and divided by the simulation time-step to get the overall fusion power:

$$P_{\text{fusion}} = \frac{\mathcal{E}}{\Delta t} \sum_k^{N_k} \frac{\text{Vol}_k}{N_p} \sum_p^{N_p} n_i n_j \sigma(v_{\text{rel}}) v_{\text{rel}} \text{ [W]} \quad (4.19)$$

Fit equations for the fusion cross section of $\text{p-}^{11}\text{B}$ as a function of the center-of-mass energy are given by Nevins and Swain [21]. More useful for simulation is the cross section as a function of the relative velocities of the particles, which, using the Nevins and Swain equations, are produced below. Because a pair of ions from the fictional species s (Eq. 4.23) has the same relationship between relative velocity and center of mass as a proton and boron nucleus pair, the cross section as a function of velocity is the same for $\text{p-}^{11}\text{B}$ as it is for species s .

$$\sigma(v) = \frac{1}{v^2} \exp\left(-\frac{v_{\text{Gamow}}}{v}\right) \times \begin{cases} a_0 + a_1 v^2 + a_2 v^4 + \frac{a_3}{(v^2 - u_{148}^2)^2 + w_{2.35}^4} & u_{20} < v < u_{400} \\ b_0 + b_1 (v^2 - u_{400}^2) - b_2 (v^2 - u_{400}^2)^2 - b_3 (v^2 - u_{400}^2)^5 & u_{400} < v < u_{642} \\ c_0 + \frac{c_1}{(v^2 - u_{581.3}^2)^2 + w_{85.7}^4} + \frac{c_2}{(v^2 - u_{1083}^2)^2 + w_{234}^4} & u_{642} < v < u_{3500} \\ + \frac{c_3}{(v^2 - u_{2405}^2)^2 + w_{138}^4} + \frac{c_4}{(v^2 - u_{3344}^2)^2 + w_{309}^4} & \end{cases} \quad (4.20)$$

v_{Gamow} is the Gamow velocity

$$v_{\text{Gamow}} = 2c\pi\alpha Z_1 Z_2 \quad (4.21)$$

a_0	$4.118 \times 10^{-26} \left[\frac{\text{m}^4}{(10^{-7} \text{ s})^2} \right]$
a_1	$2.400 \times 10^{-26} [\text{m}^2]$
a_2	$1.105 \times 10^{-26} [(10^{-7} \text{ s})^2]$
a_3	$1.662 \times 10^{-29} \left[\frac{\text{m}^8}{(10^{-7} \text{ s})^6} \right]$
b_0	$6.898 \times 10^{-26} \left[\frac{\text{m}^4}{(10^{-7} \text{ s})^2} \right]$
b_1	$6.610 \times 10^{-26} [\text{m}^2]$
b_2	$9.711 \times 10^{-26} [(10^{-7} \text{ s})^2]$
b_3	$8.275 \times 10^{-25} \left[\frac{(10^{-7} \text{ s})^8}{\text{m}^6} \right]$
c_0	$9.156 \times 10^{-28} \left[\frac{\text{m}^4}{(10^{-7} \text{ s})^2} \right]$
c_1	$2.347 \times 10^{-27} \left[\frac{\text{m}^8}{(10^{-7} \text{ s})^6} \right]$
c_2	$5.179 \times 10^{-28} \left[\frac{\text{m}^8}{(10^{-7} \text{ s})^6} \right]$
c_3	$1.224 \times 10^{-28} \left[\frac{\text{m}^8}{(10^{-7} \text{ s})^6} \right]$
c_4	$5.188 \times 10^{-28} \left[\frac{\text{m}^8}{(10^{-7} \text{ s})^6} \right]$

Table 4.1: Coefficients for Eq. (4.20).

where c is the speed of light and α is the dimensionless fine structure constant ($\alpha \approx 0.007297$). For $\text{p-}^{11}\text{B}$, $v_{\text{Gamow}} = 6.87 \times 10^7 \text{ m/s}$. The constants u_E and w_E are relative velocities at center-of-mass energy E expressed in keV, i.e. $u_E, w_E = \sqrt{\frac{2E[\text{keV}]}{\mu} \frac{1}{1000e}}$. Some of the coefficients in SI units are too small to be represented in single-precision floating-point format. Instead, all units of seconds (s) are converted to units of 10^{-7} s . The coefficients have approximate values as shown in Table 4.3. In the simulation, use of this conversion only requires an additional step of multiplying the relative velocity v by 10^{-7} , while the output σ (Eq.(4.20)) remains in units of m^2 . The cross section as a function of center-of-mass velocity (valid for both $\text{p-}^{11}\text{B}$ and $s\text{-s}$) is shown in Fig. 4.3.




Figure 4.3: Fusion cross section as a function of center-of-mass velocity for p-¹¹B fuel. The three sections of Eq. (4.20) are delineated by vertical dashed lines.

4.4: Fuel species

The present PIC simulation uses a single species of ions as a stand in for p-¹¹B fuel. As a stand in, the single ion species s is defined such that two such ions have the same argument of the Rutherford scattering formula at a given energy, and the same fusion rate production for counter-streaming beams (assuming that the proton and boron ions are moving exclusively in opposite directions).

$$\frac{Z_p Z_B}{\mu_{pB}(v_p + v_B)^2} = \frac{Z_s^2}{\mu_{s-s}(v_s + v_s)^2} \quad (4.22a)$$

$$\sigma_{pB} (\mu_{pB}(v_p + v_B)^2) (v_p + v_B) = \sigma_{pB} (\mu_{s-s}(v_s + v_s)^2) (v_s + v_s) \quad (4.22b)$$

Because of the complex nature of the fusion cross section as a function of center-of-mass energy $\sigma_{\text{PB}}(E_{\text{COM}})$, it follows that the only solution is $Z_s = \sqrt{Z_p Z_B}$, and $m_s = 2\mu_{\text{PB}}$ so that $v_s = (v_p + v_B)/2$ for a given energy. For the values $Z_p = 1$, $Z_B = 5$, $m_p = 1 \text{ AMU}$, and $m_B = 11 \text{ AMU}$, the following values satisfy Eqs. (4.22):

$$Z_s = \sqrt{5} \quad (4.23a)$$

$$m_s = \frac{11}{6} [\text{AMU}]. \quad (4.23b)$$

The voltage required to accelerate to the velocity $v_s = 5.56 \times 10^6$ (so that the relative velocity between two ions is the peak fusion cross sectional relative velocity $2v_s$) is

$$V_s = \frac{1}{2} \frac{m_s}{e Z_s} v_s^2 = 132 [\text{keV}]. \quad (4.24)$$

4.5: Optimization routine

The goal of the optimizer is to choose the voltage profile along the IEC wall that results in the best bunching of the ions. The cost function is evaluated when the particles are passing through the fusion core and is defined as

$$C = \frac{1}{N} \left[\sqrt{\sum_i^N \left| \frac{\mathbf{x}_i}{L} \right|^2} + \sqrt{\sum_i^N \left| \frac{\mathbf{v}_i - \mathbf{v}_{\text{fusion}}}{v_{\text{fusion}}} \right|^2} + N_{\text{loss}} \right]. \quad (4.25)$$

Particles are born into the system “pre-bunched”, that is, they are generated in the fusion core with a random normal distribution of offsets in both position and velocity. Over one period of oscillation within the simulation the particles travel

from the fusion core to the turnaround region and back into the core, returning to near their original positions. The pre-bunched particles, however, are not inherently in a steady-state structure, and the phase space distribution of the bunch changes quite drastically over its first few passes through the core. Therefore it is not useful to only optimize over a single period because after the first period the ions will almost surely behave sub-optimally. It is also not feasible to optimize over a large number of periods since the final state of the ions is extremely sensitive to the electrode voltages. The routine for optimizing the CE-IEC instead starts with optimization over one period and then builds up to larger numbers of periods, to better mimic steady-state operation.

4.5.1 Algorithm for the optimization wrapper

Starting with a value of $P = 1$, and initializing the wall voltage V_{best} and the particle positions and velocities $f_{\text{best}}(\mathbf{x}, \mathbf{v})$ to initial-guess values do the following:

1. Initialize the best-cost for P periods $C_P = 1$. Run the optimization routine with initial guess V_{best} , and initialize the particle positions with $f_{\text{best}}(\mathbf{x}, \mathbf{v})$, evaluating the cost function C after P periods of oscillation at each run. For each evaluation of C if $C < C_P$, set $C_P \leftarrow C$ and record V_{best} and $f_{\text{best}}(\mathbf{x}, \mathbf{v})$.
2. If steady-state of the cost function has been reached (i.e. $C_P \approx C_{P-1}$), the optimization is complete. Otherwise, set $P \leftarrow P + 1$, and go to step (1).

The reasoning for storing $f_{\text{best}}(\mathbf{x}, \mathbf{v})$ for each optimization period is to more quickly approach steady-state operation than if the bunches were re-initialized to a purely

monoenergetic state at each run. The optimization routine is a hybrid global-local optimizer. Each optimization starts using MATLAB's bounded simulated annealing function *simulannealbnd*. The optimal point from the simulated annealing routine is then used as the starting point for a MATLAB's Nelder-Mead simplex local optimizer *fminsearch* which was modified to include bounds and to specify an initial simplex size (so that smaller simplexes can be initialized when P is large).

4.6: Optimization results

The optimization method was tested on a CE-IEC with inner radius 0.25 m, outer radius 1 m, wall angle 17.7° and a total number of confined ions of species s per half-beamline of 2×10^9 which are grouped into 5000 appropriately weighted macroparticles. The computational grid consists of 3722 computational cells. The optimizer was then tested both without a magnetic field and with a magnetic field.

4.6.1 Without magnetic field

A frame of the optimization without a magnetic field is shown in Fig. 4.4. The optimization visualization is designed to be intuitive for the user by supplying important information on both the state of the simulation as well as the state of the optimization algorithm. The plot of the cost function output at each iteration is shown in Fig. 4.5, plotted against the number of periods over which the optimization was performed. Long term simulation can be used to study thermalization of the CE-IEC. The components of the temperature of the bunch,




Figure 4.4: Frame of the output of the optimization routine of the 2D3V CE-IEC optimizer.




Figure 4.5: The cost function output as a function of periods completed, with red circles denoting the iterations where the simulated annealing algorithm found a new optimum.




Figure 4.6: Frame of a long-timescale simulation of the CE-IEC beamline without a magnetic field.

normalized by the bunch energy, are calculated as follows:

$$T_x = \text{mean} \left\{ \frac{[|v_x| - \text{mean}(|v_x|)]^2}{v_{\text{fusion}}^2} \right\} \quad (4.26a)$$

$$T_r = \text{mean} \left\{ \frac{v_r^2}{v_{\text{fusion}}^2} \right\} \quad (4.26b)$$

$$T_\theta = \text{mean} \left\{ \frac{v_\theta^2}{v_{\text{fusion}}^2} \right\} \quad (4.26c)$$

$$T = T_x + T_r + T_\theta \quad (4.26d)$$

Fig. 4.6 shows the increase in temperature over 527 oscillations of the IEC to where it approaches steady-state operation. The T_θ temperature in this case is the most

“pure” metric of thermalization, as this simulation contains no magnetic field and so the only way a particle can obtain a θ component of velocity is through collisions. Both T_x and T_r are subject to the shape of the electric potential in the space that the ions occupy. The long steady increase in T_x in Fig. 4.6 is due to the spreading of the ion bunch along the beamline.

4.6.2 With magnetic field

With a magnetic field, the optimizer has less difficulty in maintaining a lower cost function. Fig. 4.7 shows a frame from an optimization with a magnetic field, which also results in better ion bunching behavior than the optimization without the magnetic field. The simulated annealer also has better success in finding optima. This is likely due to the role of the magnetic field in lessening the transverse expansion of the ion bunches. The cost function is most heavily penalized by lost particles, so when particle loss is mitigated by the magnetic field, the simulated annealer can “focus” more on reducing the cost associated with position and velocity spread, as shown in Fig. 4.8. The simulated annealer also demonstrated the ability to leave a local optimum, which was the purpose of including simulated annealing in the hybrid optimizer. Fig. 4.9 shows that the optimal voltage from the 5th to the 6th period of optimization made a drastic change, especially in the sign of the voltage difference between the first two electrodes.

The results of the optimization are tested up to a time of 210 microseconds, or 655 periods of oscillation, and a frame from this simulation is shown in Fig. 4.10.




Figure 4.7: Frame from an optimization of the CE-IEC with a magnetic field.




Figure 4.8: The cost function output as a function of periods completed, with red circles denoting the iterations where the simulated annealing algorithm found a new optimum.




Figure 4.9: The optimal voltage output of the hybrid optimizer moving from the 5th to the 6th period.

This simulation used 50,000 macroparticles each having the mass and charge of 40,000 real particles of species s . The computational grid is 3722 cells. With a time-step of 50 picoseconds, and a oscillation period of 0.32 microseconds, the entire simulation consisted of four million time-steps, with pair-matching collisions at each time-step, over an execution time of 4 days on a Nvidia c2070 GPU. Even at this time-scale, the simulation has not reached an oscillatory steady state, as evidenced by the changing temperature components up until the simulation end time.

4.7: Conclusions of the particle-in-cell optimizer

The particle-in-cell optimizer was successful insofar as the optimal voltages demonstrated long confinement times such that the limitation on ion lifetimes was due to thermalization rather than space-charge. This work shows that ion loss due to thermalization can not be contained by static voltage control, due to the fact that thermalization is cumulative over many passes through the system. Active control of CE-IEC voltages may provide greater control over the thermalization process.




Figure 4.10: Frame from the long-timescale simulation of the optimization results with a magnetic field.

Chapter 5

N-body Simulation

In Chap. 4, the PIC simulation domain consisted of one half of one beamline. However, an important aspect of the CE-IEC is that the beamlines intersect at various angles at a center point, and so the interaction between beamlines must be studied in a 3D model. A fully 3D PIC simulation would be infeasible for two reasons. First, solving Poisson's equation on a 3D grid at high-enough resolution to accurately simulate the IEC would be quite computationally intensive. Second, the geometry of the channel walls would require significant modification to the computational grid geometry, or likely the use of an unstructured grid. Instead, an *N*-body simulation is used, with electric and magnetic fields calculated via point charge and dipole discretizations respectively. Inter-particle forces are calculated directly using Coulomb's law, avoiding the need for a Poisson solution at each time-step.

5.1: Calculation of the electric field due to electrode voltages

The electrodes are modeled as conductive surfaces with a radial position but no radial thickness (like a spherical shell with holes for the beam channels). The voltage on the edges of one of these electrodes (on the surfaces of the beam channel walls) is of interest, so these edges are discretized into point charges. The vertices of the channel wall edges are

$$\begin{bmatrix} \begin{bmatrix} \Phi_{1,1} \\ \Phi_{1,2} \\ \vdots \\ \Phi_{1,N_p} \end{bmatrix} \\ \begin{bmatrix} \Phi_{2,1} \\ \Phi_{2,2} \\ \vdots \\ \Phi_{2,N_p} \end{bmatrix} \\ \vdots \\ \begin{bmatrix} \Phi_{N_e,1} \\ \Phi_{N_e,2} \\ \vdots \\ \Phi_{N_e,N_p} \end{bmatrix} \end{bmatrix} = \frac{1}{4\pi\epsilon_0} \mathbf{S} \begin{bmatrix} \begin{bmatrix} q_{1,1} \\ q_{1,2} \\ \vdots \\ q_{1,N_p} \end{bmatrix} \\ \begin{bmatrix} q_{2,1} \\ q_{2,2} \\ \vdots \\ q_{2,N_p} \end{bmatrix} \\ \vdots \\ \begin{bmatrix} q_{N_e,1} \\ q_{N_e,2} \\ \vdots \\ q_{N_e,N_p} \end{bmatrix} \end{bmatrix} \quad (5.1)$$

where $\mathbf{S}$ is the elastance matrix (the inverse of the capacitance matrix). The entries of the elastance matrix are

$$S_{(e_1,p_1),(e_2,p_2)} = \begin{cases} \frac{1}{|\mathbf{x}_{e_1,p_1} - \mathbf{x}_{e_2,p_2}|}, & e_1 \neq e_2 \vee p_1 \neq p_2 \\ \frac{1}{\left[\frac{1}{2}\Delta r_{\min,e_1,p_1}\right]}, & e_1 = e_2 \wedge p_1 = p_2 \end{cases} \quad (5.2)$$

where the second case of Eq. 5.2 (diagonal entries of $\mathbf{S}$) represent the self-capacitance of each point, modeled as a conducting sphere of radius $\frac{1}{2}\Delta r_{\min,e_1,p_1}$ where $\Delta r_{\min,e,p_i} = \min_{p_i \neq p_j} |\mathbf{x}_{e,p_i} - \mathbf{x}_{e,p_j}|$ is the distance to the closest neighboring point charge on the same electrode.

Eq. 5.1 for the unknown charge vector is performed using MATLAB's backslash operator. From the charge vector, the electric potential at each point is calculated from

$$\Phi(\mathbf{x}) = \frac{1}{4\pi\epsilon_0} \sum_e^{N_e} \sum_p^{N_p} q_{e,p} \frac{1}{|\mathbf{x} - \mathbf{x}_{e,p}|}. \quad (5.3)$$

and the electric field at each point is calculated from

$$\mathbf{E}(\mathbf{x}) = \frac{1}{4\pi\epsilon_0} \sum_e^{N_e} \sum_p^{N_p} q_{e,p} \frac{\mathbf{x} - \mathbf{x}_{e,p}}{|\mathbf{x} - \mathbf{x}_{e,p}|^3}. \quad (5.4)$$

The electric potential in the x - y plane of the IEC is plotted in Fig. 5.1, with half the CE-IEC represented as structure and the other half represented by the discretized point charges, drawn as spheres with the radius $\frac{1}{2}\Delta r_{\min,e_1,p_1}$.




Figure 5.1: Point charge values of the discretized electrodes for electrode voltages (from inner radius to outer radius) of -50 kV, -75 kV, -10 kV, and +10 kV. The electric potential in the $x-y$ plane due to these point charges is shown as well.

5.2: Calculation of the magnetic field due to permanent magnets

The permanent magnet structure is discretized into N_p dipole points, where the strength of each dipole is the magnetization M of the magnet multiplied by the volume Vol_p of the part of the magnet that the dipole is responsible for. Since the magnets are radially polarized, the dipole vectors always point in the radial direction. Choosing the origin to be at the center of the device, the dipole m of the discretized point p is

$$\mathbf{m}_p = M \text{Vol}_p \frac{\mathbf{x}_p}{|\mathbf{x}_p|}. \quad (5.5)$$

The magnetic field at any other point in the domain is given by

$$\mathbf{B}(\mathbf{x}) = \frac{1}{4\pi} \sum_p^{N_p} \left(\frac{3\mathbf{x}(\mathbf{m}_p \cdot \mathbf{x})}{|\mathbf{x}|^5} - \frac{\mathbf{m}_p}{|\mathbf{x}|^3} \right). \quad (5.6)$$

The magnetic field in the x - y plane is shown in Fig. 5.2, where on half of the CE-IEC is displayed as structure while on the other side the dipoles are represented as spheres. The electric and magnetic fields are calculated over a 3D grid of points encompassing the entire IEC domain (typically of size $300 \times 300 \times 300$) and the values are interpolated to the particles as necessary using the 3D version of Eq. 4.16 with linear weighting from the 8 nearest points. Though the evaluation of fields in this way is quite computationally intensive, it only must be calculated once at the beginning of the simulation (or loaded from saved values on the disk) and only




Figure 5.2: Visualization of the discretization of permanent magnets in the calculation of the CE-IEC magnetic field. The volume of the sphere representing each dipole is the same as the Vol_p term in Eq. 5.5.

needs to be referenced throughout the simulation. This makes the approximation that the fuel ions have no effect on the point charge distribution in Eq. 5.1. If this approximation could not be made, the contribution of ion charges to the potential on the electrodes would need be accounted for by means of another matrix-vector product in Eq. 5.1, and this equation would need to be solved and the electric field re-solved at each time-step, rendering the simulation computationally intractable.

5.3: The N -body individual time-step method with Hermite integrator

N -body methods are widely used for astrophysical gravitational simulations [22] but are applied here to charged particles in a plasma. In a global time-step method, all particle trajectories over the time-step are calculated simultaneously, and the global time-step must remain small enough to accurately capture the motion for all particles. Any particles undergoing close encounters with other particles (Coulomb collisions) will require the global time-step to be reduced accordingly, which can become computationally burdensome. To remedy this, the individual time-step method [23] evaluates particles in a queue. When it is time for a particle to be updated, the simulation calculates the trajectory of that particle over that particle's time-step, updates the particle's time, and calculates a new time-step for that particle. The trajectories are calculated using a high-order predictor-corrector method. [24]

The remainder of this section outlines the procedure by which each particle is

updated. It is assumed that the following parameters in the simulation are known for each particle $j \in \{1, \dots, N\}$: position ($\vec{x}_j$), velocity ($\vec{v}_j$), acceleration due to inter-particle forces ($\vec{a}_j$), jerk due to inter-particle forces ($\vec{k}_j \equiv d\vec{a}_j/dt$), electric field at the particle position ($\vec{E}_j$), magnetic field at the particle position ($\vec{B}_j$), the particle time-step (Δt_j), and the time (in simulation time) at which all these parameters are known for each particle (t_j). In the first step, the particle with the lowest value of $t_j + \Delta t_j$ (referred to now as particle i) is chosen.

$$i = \min_j (t_j + \Delta t_j) \quad (5.7)$$

The global simulation time is then updated

$$t = t_i + \Delta t_i \quad (5.8)$$

The time difference between the current time and the time at which the position of each particle is known is defined as

$$\delta t_j := t - t_j. \quad (5.9)$$

Note that δt_j will always be positive for all j and that $\delta t_i = \Delta t_i$. The position of each particle is predicted at the particle's current time using the leapfrog method and the Boris method [18] as used in standard particle-in-cell methods [19], modified slightly to account for the inter-particle force terms. First, the positions of all the particles are predicted at a time halfway between their last known time and the

current time.

$$\mathbf{x}'_j = \mathbf{x}_j + \mathbf{v}_j \left(\frac{\delta t_j}{2} \right) \quad (5.10)$$

Then the 3-D magnetic and electric field values that are known at discrete nodes over the domain are linearly interpolated to the position of particle i . The velocity is updated using the Boris method (Eqs. 5.11, 5.12, 5.13 and 5.14) including the acceleration contribution from the inter-particle forces as well as those from the externally applied E and B fields.

$$\mathbf{v}'_j = \mathbf{v}_j + \left(\mathbf{a}_j + \frac{q_j}{m_j} \mathbf{E}_j \right) \left(\frac{\delta t_j}{2} \right) \quad (5.11)$$

$$\mathbf{e}_j = \frac{q_j}{m_j} \mathbf{B}_j \delta t_j \quad (5.12)$$

$$\mathbf{v}''_j = (\mathbf{v}'_j + \mathbf{v}'_j \times \mathbf{e}_j) \times \left(\frac{2\mathbf{e}_j}{1 + |\mathbf{e}_j|^2} \right) \quad (5.13)$$

$$\mathbf{v}'''_j = \mathbf{v}''_j + \left(\mathbf{a}_j + \frac{q_j}{m_j} \mathbf{E}_j \right) \left(\frac{\delta t_j}{2} \right). \quad (5.14)$$

The predicted positions from the Boris method are updated

$$\mathbf{x}''_j = \mathbf{x}'_j + \mathbf{v}'''_j \left(\frac{\delta t_j}{2} \right) \quad (5.15)$$

and finally, the contributions to the position and velocity due to the jerk are added, resulting in the predicted position and velocity of all particles at the time at which particle i is to be updated.

$$\mathbf{x}_j = \mathbf{x}''_j + \frac{1}{6} \mathbf{k}_j (\delta t_j)^3 \quad (5.16)$$

$$\mathbf{v}_j = \mathbf{v}_j''' + \frac{1}{2} \mathbf{k}_j (\delta t_j)^2 \quad (5.17)$$

The next step is to calculate the acceleration and jerk on particle i based on the predicted positions of all other particles. The relative position of particle i with respect to all other particles is $\mathbf{r}_j := \mathbf{x}_i - \mathbf{x}_j$ and the relative velocity of particle i with respect to all other particles is $\mathbf{u}_j := \mathbf{v}_i - \mathbf{v}_j$. The acceleration of particle i due to the force from all other particles is

$$\mathbf{a}_i = \frac{q_i}{4\pi\epsilon_0} \sum_{j \neq i}^N \frac{q_j \mathbf{r}_j}{|\mathbf{r}_j|^3} \quad (5.18)$$

and the jerk of particle i , $\mathbf{k}_i = d\mathbf{a}_i/dt$ is

$$\mathbf{k}_i = \frac{q_i}{4\pi\epsilon_0} \sum_{j \neq i}^N \left(\frac{q_j \mathbf{u}_j}{|\mathbf{r}_j|^3} - \frac{(\mathbf{u}_j \cdot \mathbf{r}_j) \mathbf{r}_j}{|\mathbf{r}_j|^5} \right). \quad (5.19)$$

The higher order derivatives of acceleration are estimated from the jerk and acceleration

$$\ddot{\mathbf{a}}_i = \frac{6(\mathbf{a}_i - \mathbf{a}_o) - \Delta t_i (2\mathbf{k}_i + 4\mathbf{k}_o)}{\Delta t_i^2} \quad (5.20)$$

$$\dot{\mathbf{k}}_i = \frac{12(\mathbf{a}_o - \mathbf{a}_i) + 6\Delta t_i (\mathbf{k}_i + \mathbf{k}_o)}{\Delta t_i^3} \quad (5.21)$$

where a_o and k_o are the acceleration and jerk of particle i that were previously known before the values that were calculated in Eqs. 5.18 and 5.19 respectively. Finally, the new position and new velocity of particle i is updated from the predicted values

that were found in Eqs. 5.15 and 5.14 respectively.

$$\mathbf{x}_i = \mathbf{x}_i + \frac{1}{24}\ddot{\mathbf{a}}_i(\Delta t_i)^4 + \frac{1}{120}\ddot{\mathbf{k}}_i(\Delta t_i)^5 \quad (5.22)$$

$$\mathbf{v}_i = \mathbf{v}_i + \frac{1}{6}\ddot{\mathbf{a}}_i(\Delta t_i)^3 + \frac{1}{24}\ddot{\mathbf{k}}_i(\Delta t_i)^4 \quad (5.23)$$

The next time-step for particle i is updated according to the formula

$$\Delta t_i = \sqrt{\eta \frac{|\mathbf{a}_i| |\ddot{\mathbf{a}}_i| + |\mathbf{k}_i|^2}{|\mathbf{k}_i| |\ddot{\mathbf{k}}_i| + |\ddot{\mathbf{a}}_i|^2}} \quad (5.24)$$

where η is a chosen dimensionless parameter. The process then repeats, returning to Eq. 5.7 to select the next particle to be updated.

5.4: Overestimation of Coulomb scattering due to macroparticle weighting

The Coulomb scattering angle θ of a particle in a plasma of density n scales as

$$\theta \propto \sqrt{n} \frac{q^2}{m} \quad (5.25)$$

If particles in the plasma are replaced by macroparticles of weight w such that the charge density and mass density stay the same, then the new number density $\tilde{n}$ is related to the old number density by $\tilde{n} = n/w$, and the new charge and mass of each particle are related to the unweighted values by $\tilde{q} = wq$ and $\tilde{m} = wm$ respectively.

Substituting these values into Eq. 5.25 results in

$$\tilde{\theta} \propto \sqrt{wn} \frac{q^2}{m}. \quad (5.26)$$

And so Coulomb scattering angles are over-calculated by a factor of $\sqrt{w}$. A correction to this is not straightforward, since the space-charge effect is well captured by weighted particles in an N -body simulation. In the simulation results that follow, the macroparticle weighting is on the order of one million, and so the Coulomb scatters are overestimated by a factor of one thousand. This means that high-angle scatters that transfer particles between beamlines happen one thousand times more often, and that thermalization happens one thousand times faster. However, this is not completely detrimental to the research, since the CE-IEC is chiefly space-charge limited, the overestimation of Coulomb scatter makes the observation of Coulomb scattering more feasible on shorter time-scales.

5.5: Testing on two particles with a known scattering angle

To find an appropriate value for η , two equally charged particles are simulated undergoing a binary Coulomb collision. The solution to this collision is known analytically, and the results from simulations over a range of values of η can then be compared in both scattering angle and conservation of energy. The results of this test for a 90° scatter are shown in Fig. 5.3. For most values of η tested a scattering angle of close to 90° is calculated. However, for $\eta = 0.6$ the simulation “misses” this scatter by using time-steps that are too large. The computational




Figure 5.3: Testing of the Hermite integrator individual time-step method on a known 90 degree scatter for different values of η . *Left:* Simulation of a 90° scatter with equal scaling of the x and y axes. *Right:* Same simulation with the x and y axes of different scaling to illustrate differences between trajectories.

performance of this method is also compared to the more basic leapfrog method with individual time-steps, where scattering accuracy and conservation of energy are plotted vs computational time (Fig. 5.4).

5.6: Ion simulation results

Rather than creating ions in a pre-bunched configuration as was done in Chap. 4, ions are instead created continuously in time at points near the end of the channel, and are removed from the simulation when striking a wall. The bunching behavior is shown by this simulation not only to arise naturally, but also to be synchronized between beamlines. A frame of this simulation is shown in Fig. 5.5.

A frame-by-frame of the particle phase space (projected onto one beam line), core beam current, and core density is shown in Fig. 5.6. The velocity distribution function of the ions in the fusion core region is shown in Fig. 5.7. The impact points of




Figure 5.4: *Left:* Comparison of final scattering angle vs. computation time for different values of η . *Right:* Comparison of the percentage change in total energy vs. computation time for different values of η .




Figure 5.5: A frame from simulation of ions in a truncated icosahedron IEC.




Figure 5.6: Frame-by-frame plots of data from an ion simulation. *Top*: The phase space of all particles projected onto one beam line. *Middle*: The ion density in the x - y plane. *Bottom*: The beam current along one beam line through the center of the device.




Figure 5.7: Velocity distribution in the x -dimension of ions in the core region, with one beamline aligned with x .




Figure 5.8: Impact points of ions onto the surface of the CE-IEC over the course of a simulation.

ions on the CE-IEC surfaces can be mapped by saving the last position of a particle before it is deleted from the simulation due to being found inside the walls of the device. The ion impact points are shown in Fig. 5.8. The primary region of impact is clearly the inner edge of the device, with some impacts occurring on the wall surfaces near the inner radius and very few impacts occurring near the outer radius.

Finally, the simulation demonstrates that ions are transferred between beam-lines due to high angle collisions in the core. Transfers were detected qualitatively




Figure 5.9: Impact points of ions onto the surface of the CE-IEC over the course of a simulation.

by coloring all the particles of a particular beamline red, so that any red particles that show up in a different beamlines and any non-red particles that show up in the beamline of red particles are known to have arrived there via high-angle scatter.

Fig. 5.9 shows an ion scattered onto a different beamline, but quite far off of the beamline axis, resulting in its impact with the surface soon after. In fact, all the ions that were observed to transfer onto a different beamline were observed to be lost soon after, typically not even lasting another oscillation period, due to not being scattered into the “bulk” of the on-axis particle beam.

5.7: Electron simulation results

Ions and electrons move over drastically different time-scales and so the only barrier to simulating ions and electrons simultaneously is the constraint of computation speed, i.e. the electron evolution is easily captured but the ions cannot

be evolved to steady-state over reasonable computation times when electrons are present in the simulation.

Confined electrons are simulated and a frame of this simulation is shown in Fig. 5.10. For an electron input of approximately 8 amperes, the electron density in this simulation is approximately 10^{12} m^{-3} over a radius of 0.25 m and produces a potential drop of 400 V in the center. The electron density displays the expected spherical shell-like distribution due to the space charge of the electrons and the mirror effect of the magnetic line cusps. The electron impacts primarily happen in the line cusps and in this simulation no electrons were observed to have exited the simulation along the beamlines. In this simulation the electrons are generated at source points along each beamline, and are deposited at a higher voltage at the surfaces so that the power input is quite high (10 kW). To lower this power requirement, a better path may rely on thermionic emission of electrons from the inner edge so that the emitted voltage and the absorbed voltage of the electrons are identical. The electrons impact points are shown in Fig. 5.11.

5.8: Conclusions of the N -body simulation

The N -body simulation was used to investigate aspects of the CE-IEC that were not able to be investigated by the 2D simulation. The conclusions drawn from the N -body simulation are:

- High-angle collisions that transfer ions between beamlines do occur, but typically the newly transferred ions do not last more than half an oscillation




Figure 5.10: Electrons simulated under the influence of electric and magnetic fields in the CE-IEC showing the relation between power input, electron density, and electron mean lifetime.




Figure 5.11: Impact points of electrons onto the surface of the CE-IEC over the course of a simulation.

thereafter. The short lifetimes of newly transferred ions is theorized to be due to the trajectories being close to the wall that separates the old beamline and the new beamline, rather than being close to the axis of the new beamline.

- Most ion-surface collisions occur on the inner edge of the CE-IEC. It is also theorized that the majority of the α -particles would strike the inner edges. An effective sputter shield stand-off would need to be implemented to maximize the lifetime of the CE-IEC, and thermal insulation between the shield and the rest of device would be required to more effectively radiate waste heat directly from the inner edge.
- Electron losses are primarily to the inner edge rather than along beamlines, which means that magnetic mirror effect along the inner line cusps is the limiting factor on electron confinement.

Chapter 6

A Fluid Treatment of IEC Electrons

Simulating both ions and electrons as particles simultaneously in the CE-IEC is impractical because of the exceedingly small time-step ($\Delta t \approx 10^{-10}$ s) required for electron simulation. An alternative is to assume the electrons are thermalized and magnetized (Larmor radius much smaller than the scale length of the simulation) and to simulate them as a fluid via the Sharfetter-Gummel method [25]. Not only could the time-step for electron simulation be increased, but a steady-state solution may also be calculated at each ion time-step, such that the electrons are continuously in a steady-state that slowly evolves with the movement of the ions.

6.1: Governing equations

The electron conservation equation [26] is

$$\frac{\partial n_e}{\partial t} + \nabla \cdot \vec{\Gamma} = S. \quad (6.1)$$

where S is the electron source term, and the electron flux $\vec{\Gamma}$ which arises due to the drift (due to the electric field $\nabla\Phi$) and thermal diffusion ($\nabla n_e T_e$) of the electron population

$$\vec{\Gamma} = \bar{\mu} [n_e \nabla\Phi - \nabla(n_e T_e)] \quad (6.2)$$

where $\bar{\mu}$ is the electron magnetic mobility tensor such that the electron mobility parallel to the magnetic field is $\mu_0 = \frac{e}{m_e \nu}$ and the electron mobility perpendicular to the magnetic field is $\frac{\mu_0}{1+\Omega^2}$ where $\vec{\Omega} = \frac{q\vec{B}}{m\nu}$ is the vectorized Hall parameter. To simplify the derivation of $\bar{\mu}$, the flux term is written as

$$\vec{\Gamma} = \bar{\mu} \nabla \mathcal{E} \quad (6.3)$$

where $\mathcal{E}$ is the effective energy-per-unit-volume that is the source of electron flux. $\vec{\Gamma}$ can be broken up into components parallel and perpendicular to the magnetic field: $\vec{\Gamma}_{\parallel}$ and $\vec{\Gamma}_{\perp}$ respectively. The flux parallel to the magnetic field is the mobility parallel to the magnetic field multiplied by the directional derivative of $\mathcal{E}$ in the direction of magnetic field

$$\vec{\Gamma}_{\parallel} = \mu_0 \hat{\Omega} \left(\hat{\Omega} \cdot \nabla \right) \mathcal{E}. \quad (6.4)$$

The flux perpendicular to the magnetic field is the mobility perpendicular to the magnetic field multiplied by the gradient of $\mathcal{E}$ with the component parallel to the magnetic field subtracted out

$$\vec{\Gamma}_{\perp} = \frac{\mu_0}{1 + \Omega^2} \left[\nabla \mathcal{E} - \hat{\Omega} \left(\hat{\Omega} \cdot \nabla \right) \mathcal{E} \right]. \quad (6.5)$$

Therefore the electron flux is

$$\vec{\Gamma} = \vec{\Gamma}_{\parallel} + \vec{\Gamma}_{\perp} = \mu_0 \left[\frac{\nabla \mathcal{E}}{1 + \Omega^2} + \left(1 - \frac{1}{1 + \Omega^2} \right) \hat{\Omega} \left(\hat{\Omega} \cdot \nabla \right) \mathcal{E} \right] \quad (6.6)$$

and simplifying, becomes

$$\vec{\Gamma} = \frac{\mu_0}{1 + \Omega^2} \left[\nabla \mathcal{E} + \vec{\Omega} \left(\vec{\Omega} \cdot \nabla \right) \mathcal{E} \right]. \quad (6.7)$$

where the equivalence can be made that

$$\nabla \mathcal{E} + \vec{\Omega} \left(\vec{\Omega} \cdot \nabla \right) \mathcal{E} \equiv \left[\vec{I} + \vec{\Omega} \otimes \vec{\Omega} \right] \nabla \mathcal{E} \quad (6.8)$$

where $\vec{I}$ is the identity tensor and $\otimes$ is the vector outer product. It can then be deduced that

$$\vec{\mu} \equiv \frac{\mu_0}{1 + \Omega^2} \left(\vec{I} + \vec{\Omega} \otimes \vec{\Omega} \right) \quad (6.9)$$

and so

$$\vec{\Gamma} = \frac{\mu_0}{1 + \Omega^2} \left(\vec{I} + \vec{\Omega} \otimes \vec{\Omega} \right) [n_e \nabla \Phi - \mu \nabla (n_e T_e)] \quad (6.10)$$

which, in Cartesian coordinates, is

$$\vec{\Gamma} = \frac{\mu_0}{1 + \Omega^2} \begin{bmatrix} 1 + \Omega_x^2 & \Omega_x \Omega_y & \Omega_x \Omega_z \\ \Omega_x \Omega_y & 1 + \Omega_y^2 & \Omega_y \Omega_z \\ \Omega_x \Omega_z & \Omega_y \Omega_z & 1 + \Omega_z^2 \end{bmatrix} [n_e \nabla \Phi - \mu \nabla (n_e T_e)]. \quad (6.11)$$

In this work, the electron temperature is considered constant over the domain, and the simulation is limited to two dimensions, with $d/dz = 0$ and $B_z = 0$. Poisson's equation for the electric potential due to the electron and ion densities is

$$\nabla^2 \Phi = \frac{e}{\epsilon_0} (n_e - n_i). \quad (6.12)$$

which, for a static magnetic field, closes the system and makes a solution possible.

6.2: The numerical model

The numerical model presented here is in two dimensions ($\partial/\partial z = 0$) with no z -component of the magnetic field ($B_z = 0$ and $\Omega_z = 0$). The derivatives in Eq. 6.1 are discretized through the Scharfetter-Gummel scheme [25]. The domain is thus limited to the x - y plane, and is discretized into equally spaced nodes, with $x_i = i\Delta x$ and $y_j = j\Delta y$. The discretization of the second term in Eq. 6.1 at point $[i, j]$ is

$$\nabla \cdot \vec{\Gamma}_{i,j} = \frac{\Gamma_{x;i+\frac{1}{2},j} - \Gamma_{x;i-\frac{1}{2},j}}{\Delta x} + \frac{\Gamma_{y;i,j+\frac{1}{2}} - \Gamma_{y;i,j-\frac{1}{2}}}{\Delta y} \quad (6.13)$$

with

$$\Gamma_{x;i+\frac{1}{2},j} = -\frac{eT_e}{m_e\nu} \frac{1}{\Delta x} \left((1 + \Omega_x^2) \Pi_{[i+1,j];[i,j]} + \Omega_x \Omega_y \Pi_{[i+\frac{1}{2},j+\frac{1}{2}];[i+\frac{1}{2},j-\frac{1}{2}]} \right) \quad (6.14a)$$

$$\Gamma_{x;i-\frac{1}{2},j} = -\frac{eT_e}{m_e\nu} \frac{1}{\Delta x} \left((1 + \Omega_x^2) \Pi_{[i,j];[i-1,j]} + \Omega_x \Omega_y \Pi_{[i-\frac{1}{2},j+\frac{1}{2}];[i-\frac{1}{2},j-\frac{1}{2}]} \right) \quad (6.14b)$$

$$\Gamma_{y;i,j+\frac{1}{2}} = -\frac{eT_e}{m_e\nu} \frac{1}{\Delta y} \left((1 + \Omega_y^2) \Pi_{[i,j+1];[i,j]} + \Omega_x \Omega_y \Pi_{[i+\frac{1}{2},j+\frac{1}{2}];[i-\frac{1}{2},j+\frac{1}{2}]} \right) \quad (6.14c)$$

$$\Gamma_{y;i,j-\frac{1}{2}} = -\frac{eT_e}{m_e\nu} \frac{1}{\Delta y} \left((1 + \Omega_y^2) \Pi_{[i,j];[i,j-1]} + \Omega_x \Omega_y \Pi_{[i+\frac{1}{2},j-\frac{1}{2}];[i-\frac{1}{2},j-\frac{1}{2}]} \right) \quad (6.14d)$$

and defining

$$\Pi_{[A],[B]} \equiv n_{e;A} \frac{\frac{\Phi_A - \Phi_B}{T_e}}{\exp\left(\frac{\Phi_A - \Phi_B}{T_e}\right) - 1} - n_{e;B} \frac{\frac{\Phi_B - \Phi_A}{T_e}}{\exp\left(\frac{\Phi_B - \Phi_A}{T_e}\right) - 1}. \quad (6.15)$$

Poisson's equation is discretized in the usual way

$$\frac{-2\Phi_{i,j} + \Phi_{i+1,j} + \Phi_{i-1,j}}{(\Delta x)^2} + \frac{-2\Phi_{i,j} + \Phi_{i,j+1} + \Phi_{i,j-1}}{(\Delta y)^2} = \frac{e}{\epsilon_0} (n_{e;i,j} - n_{i,i,j}). \quad (6.16)$$

6.3: The time-stepping and steady-state models

The method by which Eq. 6.1 is advanced over time-steps defined by the electron movement time-scale is referred to as the **time-stepping** model. In this method, the electron conservation term is discretized in time as

$$n_e^{k+1} + (\Delta t) \nabla \cdot (\mu n_e^{k+1} \nabla \Phi^k - \mu \nabla (n_e^{k+1} T_e)) = n_e^k + (\Delta t) S. \quad (6.17)$$

with n_e^{k+1} solved for implicitly. In the **steady-state** model, it is assumed that the electrons are in a steady-state that slowly evolves with the changing ion positions. In this case, the time-dependence is considered negligible ($dn_e/dt = 0$). Stability in this model most easily achievable when the system is solved implicitly and simultaneously for both n_e and Φ . The equations that describe the system are not all linear, so the system is solved via iteration using a method [27] that starts by taking the Jacobian of the system. First, two variables are defined:

$$g_1 = \nabla^2 \Phi - \frac{e}{\epsilon_0} (n_e - n_i) \quad (6.18a)$$

$$g_2 = \nabla \cdot \bar{\mu} (n_e \nabla \Phi - \nabla (n_e T_e)) - S \quad (6.18b)$$

The solutions to n_e and Φ are found when $g_1 \rightarrow 0$ and $g_2 \rightarrow 0$. g_1 and g_2 are defined at all points on the computational mesh and Eqs. 6.18 are discretized in an identical manner to Eqs. 6.13 and 6.16. The Jacobian of the system is

$$-\begin{bmatrix} g_1 \\ g_2 \end{bmatrix} = \begin{bmatrix} \frac{\partial g_1}{\partial \Phi} & \frac{\partial g_1}{\partial n_e} \\ \frac{\partial g_2}{\partial \Phi} & \frac{\partial g_2}{\partial n_e} \end{bmatrix} \begin{bmatrix} \delta \Phi \\ \delta n_e \end{bmatrix}. \quad (6.19)$$

Solving for $\delta \Phi$ and δn_e , a new iteration is found by

$$\Phi^{k+1} = \Phi^k + \delta \Phi \quad (6.20a)$$

$$n_e^{k+1} = n_e^k + \delta n_e \quad (6.20b)$$

Eq. 6.19 is then redefined using the values from Eqs. 6.20 and the process is repeated




Figure 6.1: Test problem for the 2D hybrid PIC simulation. Six wires, three of which have positive current perpendicular to the plane and three of which have negative current create a confining magnetic field. A electron source function replenishes electrons in the center of the domain.

until convergence is reached.

6.4: Test problem and results

To observe a 2D implementation of this model in a pseudo-IEC setting, a test problem was developed. In the test problem, electrons are produced at a constant rate in the center of the domain, and six current-carrying wires create a magnetic field to limit the movement of the electrons from the source to the boundaries (see Fig. 6.1). Dirichlet conditions are imposed at the boundaries, with $n_e = 0$ and $\Phi = 0$.

The results from both the time-stepping model and the steady-state model using the same initial conditions are shown in figure 6.2. Despite using an implicit Scharfetter-Gummel scheme, the time-stepping simulation produces spurious oscillations near steep gradients and thus produces negative electron densities in some locations. While this problem does decrease with increased grid resolution, the steady-state solution avoids these spurious oscillations, even at low grid resolu-




Figure 6.2: Comparison between the time-stepping method (left) and steady-state method (right) solutions of the electron density in the test problem.

Grid size	100 x 100		200 x 200	
Method	Time-stepping	Steady-state	Time-stepping	Steady-state
Δt	1.94 ns	103 ns	0.48 ns	51.3 ns
$= \frac{\text{computation time}}{\text{time step}}$	1.6 s	4.1 s	3.2 s	13.9 s
$= \frac{\text{simulation time}}{\text{computation time}}$	1.2×10^{-9}	2.5×10^{-8}	1.5×10^{-10}	3.7×10^{-9}

Figure 6.3: Comparison between computation times for the time-stepping model and steady-state model. “ Δt ” is the length of the time step used as determined by the CFL number, the grid spacing, and the characteristic velocity of either the electrons (time-stepping model) or the ions (steady-state model).

tions. Additionally, for the parameters used, there appeared to be little difference between the time-stepping and steady-state behaviour when the fluid model was implemented into the PIC model.

The computation time for both the time-stepping and steady-state hybrid PIC models is shown in Fig. 6.3, for two different grid sizes. In both cases, the computation time for the steady-state model is approximately one fifth of the computation time of the time-stepping model. Due to the lack of spurious oscillations, as well

while the PIC simulation models them with the correct mass. Additionally, the fluid treatment does not allow for a non-thermal velocity distribution, while the PIC simulation does. Future work on the electron fluid model should continually




Figure 6.5: Side-by-side comparison of the electron fluid simulation with a particle-in-cell simulation of electrons using equivalent conditions.

verify results through comparison to a particle-in-cell model, and if the results do not agree, one or both simulations should be modified until agreement is reached so that the limitations and approximations that each simulation makes are well understood.

Chapter 7

A Coulomb Collision Model for Nonthermal Plasma Simulation

7.1: An overview of Coulomb collisions in plasma simulations

The velocity of a single charged particle in a population of other charged particles is affected by the Coulomb electric force between that particle and all other charged particles. In the simulation of charged particle plasmas, well-established methods for accounting for the Coulomb force include the following:

- The plasma fluid approximation [25], outlined in Chap. 6 is suited for plasmas in which the particle velocities follow a Maxwell-Boltzmann distribution, the velocity of any one particle changes quickly relative to the time-scale of the

plasma, and the spacing between particles is small in comparison to the length scale of the plasma. In this way, the bulk velocity and thermal velocity of each plasma species are well-distinguished.

- Poisson’s equation, typically as part of a particle-in-cell approach [19], as discussed in Chap. 4, is effective at calculating the long-range force between particles by weighting these particles to a spatial grid, but the resolution of short-range forces is limited by both the magnitude of particle weighting as well as the resolution of the spatial grid.
- The N -body simulation method [22], used in Chap. 5 is the truest method of calculating both short-range and long-range forces between particles, however the resolution is severely limited by particle weighting for systems in which the real particle count is high, and the treatment of boundary conditions in N -body simulations typically requires a separate approach.

7.1.1.1 A cumulative Coulomb collision model

This chapter is dedicated to the study of short-timescale changes in the velocities of charged particles in a non-thermal plasma. To this end, a single non-weighted charged particle (hereafter referred to as the “test particle”) moving through a uniform population of non-weighted charged particles (hereafter referred to as the “field particles”) is examined in order to develop an approximation for Coulomb scattering that can be applied to kinetic plasma simulations. The change in velocity angle of the test particle is referred to as a “scattering,” and the probability distribution

of such a scattering is dependent on the field particle density, the relative velocity between the test particle and field particle, and the amount of time over which the scattering occurs. In the development of the present method, it is assumed that there is no change in the density or the velocity distribution function of the field particles over the scattering time. It is also assumed that the center-of-mass frame stays constant over the scattering time, so that there is no energy exchange between the test particle and field particles. The energy exchange between the test particle and field particles is realized through the conversion from the center-of-mass frame to the laboratory frame.

For application to the PIC simulation of Chap. 4, the model presented here is implemented by randomly pairing macroparticles at each time step. In the center-of-mass frame of a pair, the first macroparticle is represented by the test particle and the second macroparticle by the field particles. The field particles are assumed to all have velocity equal to that of the second macroparticle, and density equal to the local density of the field particle species. After applying the present model to the first macroparticle (the test particle), the second macroparticle receives the reverse of the same collision and in this way momentum and energy are conserved. If the simulation time-step is small compared to the time-scale of the plasma evolution, then collisions implemented this way will collectively model the collision-driven thermalization of the plasma.

This chapter is organized as follows:

- In Sec. 7.2 other collision models used for non-thermal plasma simulations are

reviewed.

- In Sec. 7.4 the “cumulative binary collision approximation” is presented and a method for efficiently calculating a cumulative scattering angle from a large number of binary collisions without energy transfer is outlined. These calculations serve as the basis for which the heuristic model is later derived.
- In Sec. 7.5 the validity of the cumulative binary collision approximation is evaluated by comparing its results to the results of N -body simulations of identical scenarios.
- In Sec. 7.6 heuristic formulae are presented for recreating the effect seen in Sec. 7.4 for a plasma simulation. This section contains the complete collision model that is the focus of this chapter.
- In Sec. 7.7 results obtained from the present collision model are compared to those obtained by other collision models.
- In Sec. 7.8 the collision model is implemented in a particle-in-cell simulation of a highly non-thermal, weakly collisional plasma and the results are compared to a true N -body simulation of an identical scenario.
- In Sec. 7.9 a discussion on low impact parameters is presented in the context of commonly used formulae for calculating a minimum impact parameter.

7.2: Relevant previous research on Coulomb collision models

A method for simulating Coulomb collisions of macroparticles was first proposed by Takizuka and Abe [20] and included details on a pair-matching Monte Carlo implementation, but no comparison to direct calculation of binary collisions was performed.

The effect of a series of binary collisions on a charged particle was first addressed by Nanbu [28] who used direct calculations of binary collisions to find the scattering angle distribution functions and created a collision model to replicate it. This work included an analytical derivation for the scattering angle to approximate the effect of low-angle collisions.

Dimits *et al.* [29] argued that Nanbu's binary collision method was identical to the Lorentz collision operator and assessed Nanbu's analytical model as such. However, both Nanbu and Dimits failed to identify the heavy tail of the probability distribution of the scattering angle that is clearly present from the results of Nanbu's data from simulating a series of binary collisions. Additionally, none of the referenced works offer an analysis of the validity of simulating a cumulative Coulomb scatter as a series of binary collisions.

Rutherford's famous discovery of the nucleus [14] involved a derivation of the probability distribution for high-angle scattering of light ions off of gold nuclei. Conte [30] applied this formula to counter-streaming charged particle beams and

used it to calculate beam particle loss due to high-angle Coulomb collisions but did not apply it to cumulative low-angle scatters.

7.3: Improvements of this model over previous models

The model presented in this chapter seeks to identify both the cumulative effect of many small-angle scatters as well as the effect of a single high-angle scatter and to recover both in a piecewise continuous heuristic model. This model is the first to identify that the probability distribution of a cumulative Coulomb scattering angle Θ transitions from an exponential form $f_{\Theta}(\theta) \sim \exp(-\theta^2)$ to a power-law form $f_{\Theta}(\theta) \sim \theta^{-3}$ as θ increases. Additionally, the present model differs from previous models in that it is based entirely on the results of numerical experiments, rather than relying on the Coulomb logarithm which is not well defined for highly non-thermal and non-neutral plasmas. Like previous models, this model uses the assumption that when the distance between two particles is large, they can be considered to have no interaction at all. The cut-off distance at which this assumption is applied is denoted as $b_{\max}$ and physically symbolizes either the distance at which space charge is accounted for via another calculation such as Poisson's equation [19], or the distance at which Debye shielding [13] is significant. The present work also benefits from the general advancements in computing that have taken place in the twenty years since the publication of Nanbu's work. At the time of Nanbu's publication, the computational resources required to calculate the number binary collision calculations used in the present work were simply not available. Despite this,

Nambu's model is still used in contemporary charged particle simulation [31] though it is the aim of this work to present a more accurate model.

7.4: The cumulative binary collision approximation

A test particle of species α traveling through a field of N randomly positioned charged particles of species β will have its velocity vector changed by some angle Θ after an amount of time τ . The interactions that cause this change in angle may be approximated as the cumulative effect of independent binary collisions between the test particle and each field particle. The angle of scatter for a Coulomb collision between the test particle and a single field particle in the center-of-mass frame is [14]:

$$\theta = 2 \tan^{-1} \left(\frac{q_\alpha q_\beta}{4\pi\epsilon_0 \mu_{\alpha\beta} v_{\alpha\beta}^2 b} \right) \quad (7.1)$$

where q_α and q_β are the particle charges, $\mu_{\alpha\beta} \equiv (m_\alpha^{-1} + m_\beta^{-1})^{-1}$ is the reduced mass, $v_{\alpha\beta} \equiv |\mathbf{v}_\alpha - \mathbf{v}_\beta|$ is the relative speed between the particles, and b is the impact parameter (the perpendicular distance between the initial paths of the two particles in the center-of-mass frame). Because a collision model is typically only applied over a local region, only field particles with impact parameters $b < b_{\max}$ are considered. Over an amount of time τ of a particle simulation (usually equal to the simulation timestep), a particle of species α moving at a velocity $v_{\alpha\beta}$ relative to a population of particles of density n_β , will undergo a number of binary Coulomb

collisions approximately equal to

$$N = n_{\beta} v_{\alpha\beta} \tau \pi b_{\max}^2 \quad (7.2)$$

which is the field particle density n_{β} multiplied by the volume of a cylinder with radius $b_{\max}$ and length equal to the relative distance the test particle travels over time τ .

Let the initial velocity of a test particle be aligned with the z -axis, and let the axis rest in the center-of-mass frame of a single test particle/field-particle pair. The final velocity after N binary collisions will have a final scattering angle of Θ with respect to the z -axis. Because of the azimuthal symmetry of the problem, the final azimuthal angle is uniformly distributed between 0 and 2π . Let θ_i be the angle of the velocity vector before the i^{th} collision, $[\Delta\theta]_i$ be the change in the angle of the velocity vector due to the i^{th} collision given by Eq. (7.1), and $[\Delta\phi]_i$ be the azimuthal angle of this change, randomly selected between 0 and 2π . The azimuthal angle before the i^{th} collision, ϕ_i , has no effect on the final probability distribution function and so may be chosen to equal zero for the purpose of this derivation. The velocity vector after the i^{th} collision is found by rotating $\hat{\mathbf{z}}$ about the y -axis by $[\Delta\theta]_i$, then rotating the resultant vector about the z -axis by $[\Delta\phi]_i$ and lastly rotating that result about the y -axis by θ_i to effectively give $\hat{\mathbf{z}}$ the correct “starting position”. In

summary, the new velocity vector after the i^{th} collision is

$$\begin{aligned} \hat{\mathbf{v}}_{i+1} = & \begin{bmatrix} \cos(\theta_i) & 0 & \sin(\theta_i) \\ 0 & 1 & 0 \\ -\sin(\theta_i) & 0 & \cos(\theta_i) \end{bmatrix} \\ & \times \begin{bmatrix} \cos([\Delta\phi]_i) & -\sin([\Delta\phi]_i) & 0 \\ \sin([\Delta\phi]_i) & \cos([\Delta\phi]_i) & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ & \times \begin{bmatrix} \cos([\Delta\theta]_i) & 0 & \sin([\Delta\theta]_i) \\ 0 & 1 & 0 \\ -\sin([\Delta\theta]_i) & 0 & \cos([\Delta\theta]_i) \end{bmatrix} \hat{\mathbf{z}}. \end{aligned} \quad (7.3)$$

The z -component of $\hat{\mathbf{v}}_{i+1}$ is equal to $\cos(\theta_{i+1})$, and so the new angle is found in a simple manner by evaluation of the z -component of Eq. (7.3):

$$\begin{aligned} \cos(\theta_{i+1}) = & \cos(\theta_i) \cos([\Delta\theta]_i) \\ & + \sin(\theta_i) \sin([\Delta\theta]_i) \cos([\Delta\phi]_i). \end{aligned} \quad (7.4)$$

To randomly distribute the field particles uniformly in a cylinder of radius b_{max} , the impact parameter of each particle is calculated as $b_i = b_{\text{max}}\sqrt{U_i}$ where each U_i is independently and uniformly distributed in $(0, 1)$ and so the angle of scatter from Eq. (7.1) becomes

$$[\Delta\theta]_i = 2 \tan^{-1} \left(\frac{a}{\sqrt{U_i}} \right) \quad (7.5)$$

with the dimensionless parameter a introduced as

$$a \equiv \frac{q_\alpha q_\beta}{4\pi\epsilon_0\mu_{\alpha\beta}v_{\alpha\beta}^2 b_{\max}}. \quad (7.6)$$

The azimuthal angle is equally likely to take any value between 0 and 2π and so is calculated as

$$[\Delta\phi]_i = 2\pi V_i \quad (7.7)$$

where each V_i is independently and uniformly distributed in $(0,1)$. Combining Eqs. (7.4), (7.5) and (7.7), and making the definition $C_i \equiv \cos(\theta_i)$, the recursive relation is

$$C_{i+1} = \frac{U_i - a^2}{U_i + a^2} C_i + \frac{2a\sqrt{U_i}}{U_i + a^2} \sqrt{1 - C_i^2} \sin(2\pi V_i) \quad (7.8)$$

where $C_0 = 1$ and the final cumulative scattering angle is $\Theta \equiv \cos^{-1}(C_N)$. In this formulation, the probability distribution of Θ is dependent only on the dimensionless variables a and N (defined in Eqs. (7.6) and (7.2) respectively). Eq. (7.8) is used for generating numerical data for cases in which a is large enough that evaluation of $U_i + a^2$ is not limited by machine precision.

7.4.1 The limit for small a

For small values of a , the evaluation of $U_i + a^2$ in floating point arithmetic may result in significant error. It is found that $a \lesssim 10^{-6}$ generates noticeable error in the evaluation of Eq. (7.8) in double-precision floating-point format. Taking the limit as $a \rightarrow 0$, Eq. (7.5) becomes

$$\lim_{a \rightarrow 0} [\Delta\theta]_i = \frac{2a}{\sqrt{U_i}}. \quad (7.9)$$

With Eq. (7.7) unchanged by this limit, the scattering is now equivalent to a random walk in a 2D plane with step length $2a/\sqrt{U_i}$. By separating this 2D walk into the x and y components of the now flat θ -plane, the final scattering angle can be expressed as the magnitude of the summation of each component:

$$\lim_{a \rightarrow 0} \Theta = 2a \sqrt{\left(\sum_{i=1}^N \frac{\cos(2\pi V_i)}{\sqrt{U_i}} \right)^2 + \left(\sum_{i=1}^N \frac{\sin(2\pi V_i)}{\sqrt{U_i}} \right)^2}. \quad (7.10)$$

The scattering angle in the $a \rightarrow 0$ regime now scales linearly with a , though the dependence on N remains non-trivial. To avoid calculating scattering angles greater than π , Eq. (7.10) can be replaced with

$$\lim_{a \rightarrow 0} \Theta = 2 \tan^{-1} \left[a \sqrt{\left(\sum_{i=1}^N \frac{\cos(2\pi V_i)}{\sqrt{U_i}} \right)^2 + \left(\sum_{i=1}^N \frac{\sin(2\pi V_i)}{\sqrt{U_i}} \right)^2} \right] \quad (7.11)$$

which reduces to Eq. (7.5) for $N = 1$ but avoids the machine precision limitation inherent in Eq. (7.8) for small values of a .

7.5: The validity of the cumulative binary collision approximation

The validity of equation Eq. (7.8) in calculating the angle of the change in velocity of a particle over a time-step is examined by comparing it to an N -body simulation using identical parameters. For this validation to remain numerically tractable, the field particles are held in fixed locations ($m_\beta = \infty$, $v_\beta = 0$, $v_{\alpha\beta} = v_\alpha$). The field particles are randomly and uniformly distributed throughout a sphere of radius R at a density of n_β and the test particle starts at the sphere center moving with an initial velocity of $\mathbf{v}_\alpha$ parallel to the z -axis. At each time-step the test particle is accelerated only by those field particles that lie within a distance $b_{\max}$ of the test particle. To ensure that the simulated domain is large enough to keep the $b_{\max}$ sphere fully populated at all times, the radius of the simulation domain is $R = v_\alpha \tau + b_{\max}$ so that $\tilde{N} = n_\beta \frac{4}{3} \pi R^3$ field particles must be generated. A diagram of this method is shown in Fig. 7.1.

The N -body method used here is similar to that used in previous research [32] which is in turn based on the work of Aarseth [22]. The test particle trajectory is calculated using the following steps starting with $t_0 = 0$ and repeating until $t_k = \tau$ (where $t_k \equiv \sum_{k'=0}^k [\Delta t]_{k'}$):

1. Advance the position of the test particle over the first half of the time-step:

$$\mathbf{x}_\alpha(t_{k+1/2}) = \mathbf{x}_\alpha(t) + \mathbf{v}_\alpha(t)[\Delta t]_k/2. \quad (7.12)$$




Figure 7.1: A 2-dimensional cross-sectional schematic of the N -body simulation for testing the cumulative binary collision approximation. The test particle travels a distance of $v\tau = 2$ mm through a sphere of field particles but only experiences a force from field particles within a distance of $b_{\max} = 1$ mm.

2. For each i of $\tilde{N}$ field particles, find if it lies within a sphere of radius $b_{\max}$ centered on the test particle:

$$\mathbf{1}_{\alpha\beta,i} = [|\mathbf{x}_{\alpha\beta,i}| < b_{\max}] \quad (7.13)$$

where $\mathbf{x}_{\alpha\beta,i} \equiv \mathbf{x}_{\alpha}(t_{k+1/2}) - \mathbf{x}_{\beta,i}$.

3. Calculate the acceleration of the test particle due to the force from all field particles within the sphere of radius $b_{\max}$:

$$\mathbf{a}_{\alpha}(t_{k+1/2}) = \frac{q_{\alpha}q_{\beta}}{4\pi\epsilon_0 m_{\alpha}} \sum_{i=1}^{\tilde{N}} \mathbf{1}_{\alpha\beta,i} \frac{\mathbf{x}_{\alpha\beta,i}}{|\mathbf{x}_{\alpha\beta,i}|^3}. \quad (7.14)$$

4. Advance the velocity of the test particle over the full time-step:

$$\mathbf{v}_{\alpha}(t_{k+1}) = \mathbf{v}_{\alpha}(t) + \mathbf{a}_{\alpha}(t_{k+1/2})[\Delta t]_k. \quad (7.15)$$

5. Advance the position of the test particle over the second half of the time-step:

$$\mathbf{x}_{\alpha}(t_{k+1}) = \mathbf{x}_{\alpha}(t_{k+1/2}) + \mathbf{v}_{\alpha}(t_{k+1})[\Delta t]_k/2. \quad (7.16)$$

6. Calculate the value of the next time-step using the minimum of a method of Aarseth [22] or a maximum timestep:

$$[\Delta t]_{k+1} = \min \left([\Delta t]_{\max}, \sqrt{\eta_1} \frac{|\mathbf{a}_{\alpha}(t_{k+1/2})|}{|\dot{\mathbf{a}}_{\alpha}(t_k)|} \right) \quad (7.17)$$

where $\dot{\mathbf{a}}_\alpha(t_k) = (\mathbf{a}_\alpha(t_{k+1/2}) - \mathbf{a}_\alpha(t_{k-1/2})) / [\Delta t]_k$ and the maximum allowed timestep is $[\Delta t]_{\max} = \eta_2 / (n_\beta^{1/3} v_\alpha)$. η_1 and η_2 are chosen such that further decreasing either value does not significantly change the results of the simulation.

When $t_k = \tau$ the simulation stops and the cumulative scattering angle Θ_i is recorded as the angle between the initial velocity and the final velocity of the test particle. This process is repeated M times for a set of input parameters, where M is chosen such that the probability distribution function $f_\Theta(\theta)$ is smooth enough for confident comparison with other probability distribution functions.

Typically in a particle simulation, the time-step will be held to a value such that $\tau < Cd/v$, where C is the Courant number [33], and the distance d is either the distance between grid points or the Debye length. For these cases, the distance a particle travels in a given time step τ will almost always be less than the value $b_{\max}$, so tests of this method need not explore the parameter space where $b_{\max} \ll v_{\alpha\beta}\tau$. The probability distributions of the scattering angles for different values of $b_{\max}$ (holding constant $v\tau = 1$ mm) are shown in Fig. 7.2.

7.5.1 Shortcomings of the cumulative binary collision approximation

The cumulative binary collision approximation tends to overestimate scattering angles because it assumes a complete collision between the test particle and all field particles. But if the assumption is that particle interactions should be neglected at distances greater than $b_{\max}$, field particles with an impact parameter b close to




Figure 7.2: Probability distribution functions for varying values of $b_{\max}$ with $v_{\alpha} = 10^3$ m/s, $m = 1$ AMU, $n = 10^{11}$ m $^{-3}$ and $\tau = \mu$ s. *Top*: Results of the N -body simulation with fixed field particles. *Bottom*: Results of the cumulative binary collision approximation.

the value of $b_{\max}$ will only impart a partial collision to the test particle and so the effect of long-range Coulomb collisions becomes lessened.

Another shortcoming of the cumulative binary collision approximation is that it assumes a random distribution of field particles, but in an actual plasma, particles are not randomly distributed. Rather, the randomness of particle positions is a function of temperature. At a high temperature, where the trajectories of particles are relatively straight compared to the inter-particle distance, the instantaneous positions of particles can be close to truly random. At a lower temperature the particles of the system must stay organized in a low-energy state and so the particle positions are distinguished from a random distribution. To test the effect of the randomness of field particle positions on the scattering angle, the same N -body test was performed with field particles positioned using MATLAB's *haltonset* function [34] to uniformly fill the test volume in a quasi-random distribution (a lower-energy state than a random distribution). It was found that this uniformity had a significant effect on shifting the peak of the probability distribution to a lower angle, while the high-angle portion of the distribution remained unchanged. The probability distribution results of this test are also shown in Fig. 7.2.

Finally, it can be noted that the cumulative binary collision approximation has only two degrees of freedom, a and N , while the N -body fixed field particle simulation has three: a , N , and a third quantity: $n_{\beta} b_{\max}^3$, which scales as the number of particles within a sphere of radius $b_{\max}$. Because of the higher computational cost of the N -body fixed particle simulation as well as the additional degree of freedom it requires, the remainder of this article uses the cumulative binary collision

approximation, in spite of its shortcomings, as a baseline for which to compare the formulated heuristics of the collision model that follows.

7.6: Heuristic formulae for the cumulative scattering angle

With the assumption that the cumulative binary collision approximation can be made, calculations are feasible enough such that the probability distribution function of Θ can be found over a range of a and N (from Eqs. (7.6) and (7.2)). The collision model outlined in this section takes a single random number input U , uniformly distributed on $(0, 1)$, and produces a scattering angle output with a probability distribution function that approximates that of the cumulative binary collision approximation. The convention has been chosen so that decreasing (increasing) the random number input results in an increasing (decreasing) of the output scattering angle. Though at times counterintuitive, this convention is preferable both for plotting purposes and because the randomly generated numbers have finer resolution when closer to zero [35].

7.6.1 Functional fits for numerical data

Three regions of behavior based on the scattering angle after a large number of Coulomb collisions have been identified:

- The high-probability **low-angle region** is the collective effect of all scattering events over the time-step. It contains the angle of highest probability and is described by an exponential function.

- The low-probability, **high-angle region** is the result of the effect of one high-angle collision that is large in magnitude compared to all other scatters in that time-step. This region is well described by the analytically-determined probability distribution function for the closest expected Coulomb collision. In other words, it is the result of a single collision so large that all other collisions over the time-step are negligible.
- The mid-range **transition region** bridges the low-angle region with the high-angle region. It is best described by a linear fit of the logarithms of the variables involved, resulting in a power law.

For a single binary Coulomb collision, the cumulative distribution function of the scattering angle, or the probability that the resulting angle Θ will be greater than or equal to θ , ($F_{\Theta}(\theta) \equiv P(\Theta \geq \theta)$) is found in a straightforward manner from Eq. (7.5) by recognizing that U_1 is identical to $1 - F_{\Theta, N=1}$:

$$F_{\Theta, N=1}(\theta) = \begin{cases} 0 & \theta < 2 \tan^{-1}(a) \\ 1 - \frac{a^2}{\tan^2(\frac{\theta}{2})} & \theta \geq 2 \tan^{-1}(a) \end{cases} . \quad (7.18)$$

Eq. (7.18) is suitable for the $N = 1$ case, but for large N there is no analytical solution, and so a heuristic model is formulated instead.

7.6.1.1 High-angle region

The high-angle region is found to be well described by choosing a dummy value of $b_{\max}$ such that $N = 1$ in Eq. (7.2), i.e. $\tilde{b}_{\max} \equiv (nv\tau\pi)^{-1/2} = b_{\max}/\sqrt{N}$. A dummy version of a is defined using Eq. (7.6) with $\tilde{b}_{\max}$ in place of $b_{\max}$: $\tilde{a} \equiv a\sqrt{N}$.

The high-angle region of the cumulative distribution function then follows from Eq. (7.18) and results in

$$F_{\Theta,\text{high}}(\theta) = 1 - \frac{a^2 N}{\tan^2\left(\frac{\theta}{2}\right)}. \quad (7.19)$$

Note that $b_{\max}$ is not present in the term $a^2 N$. The probability distribution function $f_{\Theta}(\theta) \equiv \frac{d}{d\theta} F_{\Theta}(\theta)$ for the high angle region is

$$f_{\Theta,\text{high}}(\theta) = \frac{a^2 N}{\sin^2\left(\frac{\theta}{2}\right) \tan\left(\frac{\theta}{2}\right)} \quad (7.20)$$

which is equivalent to Eq. (2) in Ref. [14] (known as Rutherford Scattering). It is important to note that this equation demonstrates that the probability distribution of the scattering angle is a heavy-tailed distribution, and that any collision model that produces only an exponential probability distribution of scattering angles will tend to drastically underestimate the frequency of high-angle collisions. Inclusion of Eq. (7.20) in a collision model ensures that the collision model accurately produces high-angle scatters with the correct probability.

The continuous independent variable $u \in (0,1)$ is introduced as the domain

of possible values of the discrete random number input U . With this convention it follows that $u \equiv 1 - F_{\Theta}(\theta)$ and so $\theta(u)$ for the high-angle region is found from Eq. (7.19) and is quite similar to Eq. (7.5):

$$\theta_{\text{high}}(u) = 2 \tan^{-1} \left(\frac{a\sqrt{N}}{\sqrt{u}} \right). \quad (7.21)$$

7.6.1.2 Low-angle region

The low angle region is described by the work of Nanbu [28] and modified here to include a newly defined constant κ (dependent on a and N) which is less than unity to account for the fact that this region is not independently normalized. Additionally, a constant σ is used which corresponds to the most probable scattering angle, i.e. the maximum value of $f_{\Theta}(\theta)$, and scales generally as $a/\sqrt{N}$, but asymptotes to a value of $\pi/2$ when the effects of collisions approach isotropy (high a and/or high N). Using the formulation of Nanbu as a starting point, the probability distribution function of the low-angle region is found to be well-described by

$$f_{\Theta, \text{low}}(\theta) = \kappa \frac{\varsigma \sin(\theta) \exp(\varsigma \cos(\theta))}{2 \sinh(\varsigma)} \quad (7.22)$$

where ς is defined as

$$\varsigma \equiv \cos(\sigma) / \sin^2(\sigma). \quad (7.23)$$

The cumulative distribution function can be found by integrating the probability distribution function of Eq. (7.22), i.e. $F_{\Theta}(\theta) \equiv \int_0^{\theta} f_{\Theta}(\theta') d\theta'$:

$$F_{\Theta, \text{low}}(\theta) = \kappa \left[1 - \frac{\exp(\varsigma \cos(\theta)) - \exp(-\varsigma)}{2 \sinh(\varsigma)} \right] \quad (7.24)$$

and the scattering angle as a function of u is

$$\theta_{\text{low}}(u) = \cos^{-1} \left\{ \frac{1}{\varsigma} \log \left[\exp(-\varsigma) + 2 \sinh(\varsigma) \left(\frac{u-1}{\kappa} + 1 \right) \right] \right\}. \quad (7.25)$$

For values of $\varsigma \gtrsim 100$ the evaluation of Eq. (7.25) results in exponential overflow, so the following can be used for these cases:

$$\theta_{\text{low}, \varsigma > 100}(u) = \cos^{-1} \left\{ 1 + \frac{1}{\varsigma} \log \left(\frac{u-1}{\kappa} + 1 \right) \right\}. \quad (7.26)$$

7.6.1.3 Transition region

In between the low-angle and high-angle regions is a transition region that is not easily defined but is continuously monotonic. The transition region is chosen to be a linear fit in logarithmic space that minimizes the error when compared to the cumulative binary collision approximation. The bounds of the transition region are defined as u_{low} and u_{high} . The transition region is chosen to be a linear fit of the logarithms of θ and u , i.e. a power law. The chosen fit is

$$\theta_{\text{transition}}(u) = \theta_{\text{low}}(u_{\text{low}}) \left[\frac{u}{u_{\text{low}}} \right] \wedge \left[\frac{\log \left\{ \frac{\theta_{\text{low}}(u_{\text{low}})}{\theta_{\text{high}}(u_{\text{high}})} \right\}}{\log \left(\frac{u_{\text{low}}}{u_{\text{high}}} \right)} \right]. \quad (7.27)$$

7.6.2 Scattering angle as a function of a random seed

A piecewise function is created from Eqs. (7.21), (7.25) (or (7.26)), and (7.27):

$$\theta_{\text{Chap}}(u) = \begin{cases} \theta_{\text{high}}(u) & u < u_{\text{high}} \\ \theta_{\text{transition}}(u) & u_{\text{high}} < u < u_{\text{low}} \\ \theta_{\text{low}}(u) & u > u_{\text{low}} \end{cases} \quad (7.28)$$

A single scattering angle is calculated from a single random number input U in the this model as $\Theta = \theta_{\text{Chap}}(U)$.

7.6.3 A comparison of function fits with numerical data

Numerical data is produced using Eq. (7.8). All calculations are performed in MATLAB and executed in parallel on an NVIDIA Tesla c2070 in double-precision floating-point format, which performs at an effective rate of approximately 1 nanosecond per binary collision evaluation including random number generation. M trials are performed, and in each trial, Eq. (7.8) is evaluated N times. The result is a collection of independently produced values Θ_i , $i = 1 \dots M$. The cumulative distribution function from these trials is

$$F_{\Theta}(\theta) = \frac{1}{M} \sum_{i=1}^M \mathbf{1}_{\Theta_i \leq \theta}. \quad (7.29)$$

Defining $\tilde{\Theta}$ as an ordering of Θ such that $\tilde{\Theta}_1 \geq \tilde{\Theta}_2 \geq \dots \geq \tilde{\Theta}_M$, a function that relates a random number input to a scattering angle in a manner that replicates the numerical cumulative distribution function Eq. (7.29) is

$$\theta_{\text{binary}}(u) = \tilde{\Theta}_{\lceil Mu \rceil} \quad (7.30)$$

where $\lceil \cdot \rceil$ is the ceiling function. These functions imply a probability of $1/M$ for each Θ_i . By choosing the constants σ , κ , u_{low} , and u_{high} such that the error is minimized between Eq. (7.30) and Eq. (7.28), then the scattering angles produced by Eq. (7.28) will have similar probability distributions to those produced by N evaluations of Eq. (7.4). The cost function for the optimization of the fit function is

$$C = \sum_{i=1}^M \left[\tilde{\Theta}_i - \theta_{\text{Chap}}\left(\frac{i-1/2}{M}\right) \right]^2. \quad (7.31)$$

By adjusting the values of σ , κ , u_{low} , and u_{high} so that the cost function is minimized, Eq. (7.28) becomes a good approximation for Eq. (7.30).

For the values of $a = 10^{-3}$ and $N = 10^3$, and using $M = 10^7$, the resulting plot of $\theta_{\text{binary}}(u)$ (smoothed for clarity) is shown in Fig. 7.3. The cost function was minimized using MATLAB's nonlinear least squares solver *lsqnonlin* [34] and resulted in values of $\sigma = 0.132$ rad, $\kappa = 0.912$, $u_{\text{low}} = 0.194$, and $u_{\text{high}} = 0.00481$ which are used to plot the three pieces of Eq. (7.28).




Figure 7.3: Comparison of scattering angles produced by the cumulative binary collision approximation with scattering angles produced by the three pieces of Eq. (7.28).

7.6.4 Trends for σ , κ , U_{low} , and U_{high}

To be useful for a plasma simulation, the values of σ , κ , u_{low} , and u_{high} need to be easily approximated for a given pair of a and N . These parameters can be found by repeating the process outlined in Sec. 7.6.3 for a range over both a and N and then finding fit functions that closely follow the values found.

7.6.4.1 Low-angle regime

For low values of a and/or N the peak scattering angle is low ($\sigma \rightarrow 0$, $\varsigma \rightarrow \infty$) and the values of $\tilde{\sigma} \equiv \frac{\sigma}{a\sqrt{N}}$, κ , u_{low} , and u_{high} have logarithmic dependence on N , and no dependence on a . The functional fits chosen for these four parameters, using

the shorthand $x_{a \rightarrow 0} \equiv \lim_{a \rightarrow 0} [x(a, N)]$ are

$$\tilde{\sigma}_{a \rightarrow 0} = -K_{\sigma}^{(1)} \exp \left(-K_{\sigma}^{(2)} N^{K_{\sigma}^{(3)}} \right) + K_{\sigma}^{(4)} \quad (7.32a)$$

$$\kappa_{a \rightarrow 0} = -K_{\kappa}^{(1)} \exp \left(-K_{\kappa}^{(2)} N^{K_{\kappa}^{(3)}} \right) + 1 \quad (7.32b)$$

$$u_{\text{low}}_{a \rightarrow 0} = K_{u_{\text{low}}}^{(1)} \exp \left(-K_{u_{\text{low}}}^{(2)} N^{K_{u_{\text{low}}}^{(3)}} \right) \quad (7.32c)$$

$$u_{\text{high}}_{a \rightarrow 0} = K_{u_{\text{high}}}^{(1)} \exp \left(-K_{u_{\text{high}}}^{(2)} N^{K_{u_{\text{high}}}^{(3)}} \right) \quad (7.32d)$$

where all values of K are positive. Only values of $N \geq 1000$ are used for finding the best-fit parameters, and so the equations are only to be considered valid in this range. In cases where N is large, M must be small so that computation time (which is approximately $MN \times 10^{-9}$ s per data point) remains reasonable. Numerical data for different values of M are shown alongside plots of Eqs. (7.32) in Fig. 7.4. Best-fit values for K were found using again MATLAB's nonlinear least squares solver *lsqnonlin*.

7.6.4.2 High-angle regime

When a and/or N are not low, $\tilde{\sigma}$, σ , κ , u_{low} , and u_{high} have dependence on both a and N . The functional fits chosen for extending Eqs. (7.32) into the high-angle




Figure 7.4: Trends for σ , $\tilde{\sigma}$, κ , u_{low} , and u_{high} for cases in which the scattering angle is very small and the results depend only on N .

regime are

$$\sigma(a, N) = a\sqrt{N} \tilde{\sigma}_{a \rightarrow 0} \left\{ 1 + \left(\frac{2}{\pi} a\sqrt{N} \tilde{\sigma}_{a \rightarrow 0} \right)^{K_{\sigma}^{(5)}} \right\}^{\frac{1}{K_{\sigma}^{(5)}}} \quad (7.33a)$$

$$\kappa(a, N) = \min \left[1, \frac{\kappa}{a \rightarrow 0} \exp \left(K_{\kappa}^{(4)} \sigma^{K_{\kappa}^{(5)}} \right) \right] \quad (7.33b)$$

$$u_{\text{low}}(a, N) = \min \left[1, \frac{u_{\text{low}}}{a \rightarrow 0} \exp \left(-K_{u_{\text{low}}}^{(4)} \sigma^{K_{u_{\text{low}}}^{(5)}} \right) \right] \quad (7.33c)$$

$$u_{\text{high}}(a, N) = \min \left[1, \frac{u_{\text{high}}}{a \rightarrow 0} \exp \left(-K_{u_{\text{high}}}^{(4)} \sigma^{K_{u_{\text{high}}}^{(5)}} \right) \right]. \quad (7.33d)$$

Note that Eqns (7.33b), (7.33c), and (7.33d) are dependent on (7.33a). The best-fit values for the functional fits are:

$$\begin{aligned} K_{\sigma} &= \begin{bmatrix} 1.040 \times 10^6 \\ 11.76 \\ 3.289 \times 10^{-3} \\ 10.41 \\ 4.17 \end{bmatrix}, K_{\kappa} = \begin{bmatrix} 6.776 \times 10^7 \\ 19.92 \\ 3.803 \times 10^{-3} \\ 0.4890 \\ 2.576 \end{bmatrix}, \\ K_{u_{\text{low}}} &= \begin{bmatrix} 1.926 \times 10^9 \\ 20.72 \\ 3.164 \times 10^{-4} \\ 166.5 \\ 6.193 \end{bmatrix}, K_{u_{\text{high}}} = \begin{bmatrix} 5.307 \times 10^7 \\ 22.61 \\ 1.720 \times 10^{-3} \\ 6.248 \\ 1.618 \end{bmatrix}. \end{aligned} \quad (7.34)$$

The fits for these equations are plotted in Fig. 7.5. The discrepancy in the case of u_{low} is likely due to differences in the implementation of the optimizer between the




Figure 7.5: Plots of σ , $\tilde{\sigma}$, κ , u_{low} , and u_{high} along with best-fit functions for a range of a and N . Selected contours of constant value are plotted to aid in comparison.

low-angle and high-angle regimes, but this difference is not large enough to cause a significant change in scattering angles generated by the this model.

7.7: Comparison to previous methods

Previous methods include the work of Takizuka and Abe [20] and Nanbu [28].

Takizuka and Abe define a scattering angle variance which can be rewritten in terms of the parameters a and N as

$$\langle \delta^2 \rangle = 2a^2 N \log \left(\frac{1}{2a} \right). \quad (7.35)$$

A normally distributed random number, δ , is produced with variance $\langle \delta^2 \rangle$ and the scattering angle is calculated as

$$\theta_{\text{Takizuka-Abe}}(\delta) = 2 \tan^{-1}(\delta). \quad (7.36)$$

Nanbu defines an isotropy parameter, s , which may be written in terms of the parameters a and N as

$$s = 4a^2 N \log \left(\frac{1}{2a} \right) \quad (7.37)$$

and the parameter A is defined in terms of s as

$$\coth A - A^{-1} = \exp(-s). \quad (7.38)$$




Figure 7.6: A comparison of the probability distribution functions for the scattering angle between the cumulative binary collision approximation (Sec. 7.4), the N -body simulation (Sec. 7.5), the Nanbu method, the Takizuka-Abe method, and the present method (Sec. 7.6)

From this value of A , the scattering angle as a function of u is

$$\theta_{\text{Nanbu}}(u) = \cos^{-1} \{ A^{-1} \log [\exp(-A) + 2u \sinh(A)] \}. \quad (7.39)$$

It is worthwhile to note that Nanbu's parameter A has a similar role to the parameter ς defined in Eq. 7.23. Results from Nanbu's formulation can be compared to the present model as well as with the results of the cumulative binary collision approximation. The results of this comparison are shown in Fig. 7.6. The results of the Nanbu method are slightly upshifted from the results of the cumulative binary collision approximation, which in turn was shown to be upshifted from the N -body simulation of Sec. 7.5. Most glaringly, however, neither the Takizuka-Abe method nor the Nanbu method recreate the low-probability, high-angle scattering above 0.5




Figure 7.7: The relative discrepancy of the mean scattering angle for the Nanbu method and the present method as compared to the results of the cumulative binary collision approximation.

radians seen in both the binary and N -body collision data, as well as the present model. For comparison over a range of a and N , the errors of the average scattering angle relative to the results of the cumulative binary collision approximation for both the Nanbu method as well as the present method are shown in Fig. 7.7.

7.8: Implementation and comparison to an N -body simulation

The present collision model is tested by implementation into a 2D3V axisymmetric particle-in-cell (PIC) similar to that used in Chap. 4. This simulation was then compared to the results of an N -body simulation of an identical scenario, as shown in Fig. 7.8. The scenario chosen is that of counterstreaming ion beams, to




Figure 7.8: A frame from the counter-streaming N -body simulation used for testing the collision model.

demonstrate the effect of high-angle scatters in a situation that is illustrative of the conditions encountered in inertial electrostatic confinement fusion [4]. Two beam sources are placed facing one another at a distance of 10 mm apart, each producing monoenergetic protons with an initial axial velocity of 10^4 m/s, at a density of 10^{13} m $^{-3}$ in an initial beam radius of 0.1 mm. Such a small-scale scenario is chosen so that the N -body simulation can simulate real particles rather than macroparticles. After the simulations reach steady-state, the densities are time-averaged over a long enough duration ($t \approx 0.5$ ms) so that the density plot is smooth.

The collision model is implemented into the PIC simulation using the Monte-carlo approach described by Takizuka and Abe [20], in which particles are randomly matched pairwise with other particles in the same simulation cell. For comparison, this PIC simulation was run using the present collision model, Nanbu's collision model, as well as a baseline case of no collision implementation at all. The particles in the PIC simulation are oversampled ($w = 0.1$) to ensure that simulation cells

within the beam envelope are well populated.

The N -body simulation for this scenario uses the method of Ref. [32], which in turn is based on Aarseth [22] and is similar to the N -body method described in Sec. 7.5. It uses a particle weighting of unity so that the macroparticle approximation is avoided. The results from these simulations are compared in Fig. 7.9.

Some inherent differences in these simulations preclude exact agreement. Due to limitations in computational power, the chosen beam radius is quite small compared to the beam density, such that the mean inter-particle spacing (≈ 0.03 mm) is not small compared to the beam radius, meaning that the beam is not as axially symmetric as the initial conditions may suggest, and may also be the reason that the beam envelope is less sharply defined in the density profile of the N -body simulation as compared to the PIC simulation. Another inherent difference is that the N -body simulation has completely open boundary conditions, which is not feasible within a PIC simulation. To reduce unwanted boundary effects, Dirichlet boundary conditions in the PIC simulation were placed at twice the axial extent ($x = 10$ mm) and twice the radial extent ($x = 2$ mm) so that the boundaries would not have a significant effect on the beam envelope.

7.9: Discussion of small impact parameters

Collision models generally make use of a minimum impact parameter, below which collisions are not considered. In the development of this model, charged




Figure 7.9: Time-averaged density for four different simulations of counterstreaming ion beams. The plots are axisymmetric about the z -axis and plane-symmetric about the r axis. The envelope of the beam sourced at $z = 5$ mm is visible as a dark shade and the envelope of the beam sourced at $z = -5$ mm is visible as a light shade in all plots. The density resulting from high-angle scatters permeates the remainder of the domain and is displayed using contour lines of constant value. Densities down to 10^6 m^{-3} are resolved by time-averaging the density over 0.5 ms. Densities below 10^6 m^{-3} are not resolved.

particles are treated as points with no spatial extent, and no minimum impact parameter is assumed. In actuality, the scattering angle is limited by the size and nature of the participating particles. As an example, collisions between protons and boron-11 are considered under IEC fusion conditions. The following considerations are present in relevant literature concerning the lower limit of impact parameters:

- A fusion event occurs if the impact parameter between any two particles is below the experimentally determined maximum fusion impact parameter.
- The particles may come within a de Broglie wavelength of each other, suggesting that their matter waves have overlapped to such an extent that the point-charge Coulomb force is no longer an accurate representation of the interaction between them. The distance at which this occurs is used in some models as the minimum impact parameter [37].
- The potential energy of a particle pair may exceed the kinetic energy of the particle pair in the center-of-mass frame. The distance between particles at this limit is used as a minimum impact parameter for some collision models [38] though it serves only as a relevant scale and has no immediately obvious physical significance. Many other models use similar scales pertaining to the potential energy of the particle pair [39, 40].

In assessing these conditions, it is assumed that a high cumulative scattering angle results from a single high-angle scatter that makes all low-angle scatters negligible over the time-step, i.e. $U \ll u_{\text{high}}$. The scattering angle for this region is given by Eq. (7.21). In this limit, the minimum impact parameter experienced by the test

particle is the impact parameter that would result in this scattering angle from a binary collision with one field particle:

$$b_{\min}(U \ll u_{\text{high}}) = \sqrt{\frac{U}{n_{\beta} v_{\alpha\beta} \tau \pi}}. \quad (7.40)$$

The distance of closest approach r_0 during a single binary collision as a function of the impact parameter and scattering angle is

$$r_0(\theta, b) = b \frac{\cos\left(\frac{\theta}{2}\right)}{1 - \sin\left(\frac{\theta}{2}\right)}. \quad (7.41)$$

Combining Eqs. (7.21), (7.40), and (7.41) reveals the minimum of r_0 among all binary collisions, i.e. the minimum of the closest approaches between the test particle and all field particles:

$$r_{\min}(U \ll u_{\text{high}}) = ab_{\max} \left(1 + \sqrt{1 + \frac{U}{a^2 N}} \right). \quad (7.42)$$

7.9.1 Fusion event

The maximum impact parameter for a fusion event is $b_{\text{fusion}} = \sqrt{\frac{\sigma_{\text{fusion}}}{\pi}}$, so a randomly generated scattering angle that suggests a lower impact parameter than b_{fusion} can be assumed to have resulted in the fusion of the test particle with a field particle. The range of u for which U results in a fusion event is defined as

$$u_{\text{fusion}} \leq n_{\beta} v_{\alpha\beta} \tau \sigma_{\text{fusion}}. \quad (7.43)$$

u_{fusion} is equivalent to the probability of the test particle fusing with a field particle during an amount of time τ .

7.9.2 de Broglie wavelength

The de Broglie wavelength of a particle is

$$\lambda_{\text{de Broglie}} = \frac{h}{p} \quad (7.44)$$

where h is the Planck constant and p is the particle momentum. The criterion of interest is if at any time the distance between the test particle and any field particle becomes less than the sum of their de Broglie wavelengths. This criterion is satisfied if and only if the minimum of the distances of closest approach given by Eq. (7.42) is less than or equal to the sum of the de Broglie wavelengths of the particles:

$$r_{\min}(U \ll u_{\text{high}}) \leq \lambda_{\text{de Broglie}}. \quad (7.45)$$

From the difference of the initial kinetic energy and the potential energy at closest approach, the momentum of the particle pair at closest approach can be found:

$$p = 2\mu_{\alpha\beta} \sqrt{v_{\alpha\beta}^2 - \frac{e^2}{4\pi\epsilon_0\mu_{\alpha\beta}r_{\min}}}. \quad (7.46)$$

Combining Eqs. (7.42), (7.45) and (7.46), the range of values for which U results in the test particle coming within a distance of any field particle less than or equal to

the sum of their de Broglie wavelengths is defined:

$$u_{\text{de Broglie}} \leq a^2 N \left\{ \frac{1}{4} \left[\sqrt{1 + \left(\frac{4\pi\epsilon_0 v_{\alpha\beta} h}{q_{\alpha} q_{\beta}} \right)^2} - 1 \right]^2 - 1 \right\}. \quad (7.47)$$

7.9.3 Potential energy equal to kinetic energy

The potential energy of the particle pair exceeds its kinetic energy when its potential energy at closest approach exceeds half its initial kinetic energy:

$$\frac{q_{\alpha} q_{\beta}}{4\pi\epsilon_0 r_{\min}} \geq \frac{1}{4} \mu_{\alpha\beta} v_{\alpha\beta}^2. \quad (7.48)$$

Combining Eqs. (7.42) and (7.48) results in the range of values for which U results in a test particle field particle pair having a higher potential energy than kinetic energy at closest approach:

$$u_{\text{potential}} \leq 8a^2 N. \quad (7.49)$$

For comparison of u_{fusion} , $u_{\text{de Broglie}}$, and $u_{\text{potential}}$ under p-¹¹B fusion conditions, either species can be assigned as the α species and the other assigned to the β species. The velocities are chosen such that the center-of-mass energy is equal to the resonant center-of-mass peak fusion cross-section that occurs at approximately 148.3 keV [41] where the fusion cross section is approximately $\sigma_{\text{fusion}} = 10^{-29} \text{ m}^2$.

The densities are chosen to be $n_p = n_B = 10^{16} \text{ m}^{-3}$ with $\tau = 10^{-8} \text{ s}$. To avoid electron shell effects, boron nuclei are simulated ($q_B = 5e$.) The values of u_{fusion} , $u_{\text{de Broglie}}$, and $u_{\text{potential}}$ are plotted in Fig. 7.10. It is clear that above these limits,




Figure 7.10: A comparison of scattering angle probabilities with the probabilities of u_{fusion} (a fusion event), $u_{\text{de Broglie}}$ (significant interaction of matter waves), and $u_{\text{potential}}$ (potential energy exceeding kinetic energy) occurring.

high angle scattering beyond that which is predicted by Nanbu's model is present in this particular IEC fusion scenario.

7.10: Concluding remarks on the Coulomb collision model

A collision model for non-thermal plasma simulation has been formulated based on data obtained by numerical experimentation on the effect of repeated binary collisions on a test particle. The work presented in this chapter expands on previous efforts by accounting for low-probability, high-angle scatters and by limiting the model input to two parameters: a and N (Eqs. (7.6) and (7.2) respectively). From these two parameters, the values σ , κ , u_{low} , and u_{high} are calculated from Eqs. (7.32) and (7.33). Finally, the scattering angle is calculated from a random number input using Eq. (7.28). Numerical experiments show that this model recov-

ers high-angle scatters not seen in previous models, and conforms well to numerical data produced by the cumulative binary collision approximation. Lastly, a significant range of high-angle scatters was shown to be present at impact parameters above commonly defined forms of a minimum impact parameter $b_{\min}$ in a highly non-thermal plasma.

Chapter 8

The Standing Wave Direct Energy Converter

Direct energy conversion, or the conversion of the kinetic energy of charged fusion products directly into electricity, is necessary for keeping the specific mass (mass per unit power) of a space power system low enough provide a game-changing alternative to current space-based power systems.

The Traveling Wave Direct Energy Converter (TWDEC) [42] was conceived as a way of direct energy conversion that produced alternating current power and did not require megavolt voltage levels. This chapter introduces the Standing Wave Direct Energy Converter (SWDEC) as a simplified version of and possible milestone towards the TWDEC and to facilitate a general understanding of the physics of the TWDEC as well as to simplify the modeling and results. The SWDEC may also stand alone as an alternative to the TWDEC.

8.1: SWDEC overview

An SWDEC or a TWDEC is a linear particle decelerator that may consist of two sections of electrodes, a modulator section and a decelerator section. An experimental TWDEC setup is shown in Fig. 8.1, with an ion beam source as a stand-in for charged fusion products. The modulator section is only necessary for




Figure 8.1: Schematic for TWDEC test article at NASA Johnson Space Center.

a continuous beam input. In the modulator section the α -particles pass through a series of electrodes with time-varying voltages, accelerating some of the α -particles and decelerating others, so that the density of the beam becomes modulated and the ions become bunched.

Downstream from the modulator section, the decelerator electrodes use the kinetic energy of the ion bunches to excite an oscillating circuit, from which power can be drawn. The TWDEC differentiates itself from other direct energy conver-

sion methods [43] in that it provides its electric power in alternating current — advantageous for the direct drive of proposed radio frequency (RF) space propulsion systems [44–46] — and can operate at a lower voltage relative to the ion energy than a direct charge capturing system [47]. The mechanism of energy conversion in a TWDEC is analogous to a linear particle accelerator operating in reverse. Rather than imparting electric field energy to a particle in order to accelerate a group of ions, the ions are decelerated while exciting an oscillating resistor-inductor-capacitor (RLC) circuit, thereby providing an alternating current power source. The impending particle bunches and the oscillation of the circuit are synchronized so that the particle bunches consistently experience a positive potential gradient, as shown in Fig. 8.2 for the case of a standing wave. Though the time-varying electric field due to the passing ion bunches is what causes the oscillation of the RLC circuit, the decelerator electrodes at peak oscillation are the dominant source of electric field, able to impart a significant deceleration on the ions.

8.1.1 Past research

Past work on the TWDEC includes a study on a concept for a $D-^3He$ fusion reactor incorporating the TWDEC [48] and a system level study on the effect of TWDEC implementation on the specific mass of a variety of theoretical fission and fusion powered spacecraft [49]. The physics of particle deceleration in the TWDEC has been studied numerically in [50] which found significant differences between the 1D and 2D models. The approach used for these models was that of

an externally imposed voltage on the decelerator electrodes. These studies did not directly model the conversion of kinetic energy into electric energy. The goal of the research presented in this chapter is to directly model the conversion.

8.1.2 SWDEC vs. TWDEC

The SWDEC differentiates itself from the TWDEC in that the SWDEC operates with an electrode spacing equal to one half the wavelength, while the TWDEC has an increased number of electrodes per wavelength so that the waveform imposed by the voltages of the decelerator electrodes can travel with the moving particles in order to increase the deceleration efficiency. The electrode spacing in both systems must be adapted (tapered) to match the changing particle velocity. While most basic principles are common to each, the SWDEC has a more simple circuitry and is chosen for study to facilitate understanding, computation, and analysis.

8.2: SWDEC simulation overview

Two simulations were created for optimizing and studying the SWDEC. The first simulation is a 1D1V semi-analytical method, that takes advantage of fast simulation times for optimizing the electrode spacing. It is semi-analytical in the sense that it does not use a computational grid but instead calculates the electric interaction between on-axis point particles and ring electrodes using analytical expressions, and then advances the simulation time-step numerically. This model simulates the bunches of α -particles as point charges.

The second simulation is a 2D3V axisymmetric particle-in-cell simulation from which the code in Chap. 4 was developed. This model assumes that α -particles generated from the CE-IEC have been collimated into a beam and are near-monoenergetic. While the CE-IEC would produce pulses of fusion products, this chapter investigates both a pulsed beam of α -particles as well as a continuous beam, so that the results are applicable to a variety of fusors. In the case of a continuous beam, a modulator electrode section is needed in front of the decelerator electrode section in order to first change it into a pulsed beam. The 2D3V PIC model is used to study the physics of the conversion of the kinetic energy of the ions into electrical power, study the beam modulation process, validate the simplifications made in the 1D1V model, and test the optimization results.

Both the 2D3V and 1D1V model simulate the electrodes at floating potentials connected through a simple resistive circuit, allowing the direct measurement of converted power while maintaining conservation of energy. Particle-in-cell methods are inherently computationally intensive and not suited to optimization schemes in which the simulation must iterate over a parametric sweep. For the purposes of parametric studies and optimizations, a fast method of SWDEC simulation with direct applicability to TWDEC simulation is presented in the following section.




Figure 8.2: Frame-by-frame illustration of the SWDEC deceleration mechanism using four ring-shaped electrodes. Each electrode has an alternating electric charge, creating a standing wave along the axis. A correctly timed ion will consistently experience a positive potential gradient, resulting in the deceleration of the ion.

8.3: A 1D1V semi-analytical simulation of the SWDEC

8.3.1 Point-charge description of the ion bunches

To expedite simulation and make quick optimization schemes possible, the ion bunches that result from a modulated ion beam are approximated as point charges

traveling along the axis of the device. These point charges have the same charge and mass as the number of ions they represent and have a spacing equal to the wavelength of the modulated beam. Some of the physics of the ion bunches is lost to this approximation, and some considerations need to be made to account for discrepancies between the point-charge description and a full particle-in-cell simulation:

1. There are differences in potential due to a finite-sized ion bunch and an equally-charged point charge placed at the center of the bunch. These differences are investigated in section 8.3.2.
2. The velocity modulation of the beam that takes place in the modulator section limits the lifetime of the ion bunches. It is shown in section 8.3.3 that this limitation can be overcome by lowering the modulator voltage and increasing the distance between modulator and decelerator sections.
3. The space charge expansion of the ion bunches limits their lifetime. This is accounted for in the model by an analytical approximation to the bunch expansion in section 8.3.4.
4. In the decelerator section, the finite size of the bunches causes non-uniform deceleration of the ions, which generally leads to the point-charge approximation overestimating the energy conversion. This is investigated by a particle-in-cell simulation of the deceleration process in section 8.3.14.
5. The kinetic energy of the ions from radial and azimuthal velocities that arises due to the bunch expansion, the axial magnetic field, and other irreversibilities,

cannot be converted by the decelerator and leads to a decrease in converted energy when these velocities are present.

8.3.2 Comparison between the particle-in-cell simulation of the ion bunches and the point-charge approximation

The focus of this study is energy conversion, and so the chief concern of this section is the discrepancy of the electric potential at the electrode radius between that due to the particle-in-cell ions and that predicted by the point-charge approximations of those bunches. The discrepancy is due to two factors. First, the modulation process is not perfect and not all of the ions are moved into the bunched regions. Some ions will stay in the space between bunches and will not be decelerated and their energy will not be converted. This will cause the simulation to *overestimate* the amplitude of alternating voltage induced on the decelerator electrodes. Second, the point-charge ion bunches exist only on the axis of the device, and do not pass as close to the electrode rings as off-axis ions. This causes the simulation to *underestimate* the potential induced on the decelerator electrodes. How these two opposing errors offset one another is chiefly a relation between the scale length of the device as compared to the radius of the electrodes. These effects are shown in Fig. 8.3. The results obtained by the model will therefore have increased error for very large or very small ratios of electrode spacing to electrode radius. Coincidentally, this is the same regime within which the space-charge expansion model presented in section 8.3.4 is valid. This does not mean that designs outside of this regime are

necessarily suboptimal but it does imply that there is some parameter space which this study does not enter.




Figure 8.3: A frame-by-frame comparison of the point-charge description of the modulated ion beam with the 2D axisymmetric particle-in-cell simulation of the modulation process. Particles are moving from left to right. The two methods are simulated separately and then superimposed upon one another for comparison. The modulator electrodes do not have any effect on the point-charge bunches. Axial and radial axes are of different scales for clarity.

8.3.3 Effect of velocity modulation on ion bunch lifetime

The modulator section of a TWDEC or SWDEC imparts a velocity broadening to the beam, which transforms the continuous input beam into a series of

ion bunches. The velocity modulation that creates the bunches also limits their lifetime. Simulations using a low beam density were performed to investigate this effect. Using a low beam density ensures that the bunch lifetime is limited only by the velocity spread due to the modulation process rather than by the space charge expansion of the ion bunch. Fig. 8.4 shows two simulations using different modulation voltages. The higher modulation voltage results in a quick formation of bunches upon exit of the deceleration region, though their lifetime is short due to their high velocity spread. A lower modulation voltage results in a greater bunch lifetime, though there is also a longer distance required for the bunches to form before they are useful for energy conversion. While staying in the realm of tractable domain sizes, the simulations showed that for low beam densities there is no limit on increasing ion bunch lifetime by reducing modulation voltage and increasing the device length, though sometimes increasing the number of modulator electrodes was also necessary to maintain good bunch formation at low voltage amplitudes. For this reason, the model does not account for bunch lifetime limitation due to the modulation process.

8.3.4 Effect of space-charge expansion on ion bunch lifetime

The previous section used low beam currents to isolate the velocity spread effect. When the beam density is raised the limiting factor on bunch lifetime is space-charge expansion. This is accounted for in the model by an analytic approximation to bunch expansion. In the decelerator section, the radial expansion will be limited




Figure 8.4: A comparison of the effect of modulation voltage on ion bunch formation. A higher voltage (top) results in quick formation of bunches, while a lower voltage (bottom) leads to longer bunch lifetimes. The simulation uses a low beam current (1 ampere) so that the expansion of the bunches due to space charge is low.

by the axial magnetic field, preventing ions from striking the ring electrodes. The axial expansion will be affected (and likely limited) by the decelerator electrodes but the full extent of this effect is not known. Thus, in this model the axial expansion of the bunch is assumed to have no limitation. Because of this, as the ion bunches decelerate, the spacing between the bunches decreases. Once the bunches overlap, the effectiveness of the electrodes in decelerating the ion packets will diminish. A limitation on the maximum ion bunch size is set to the instantaneous wavelength, $\lambda = \frac{v}{f}$. The wavelength correlates to the electrode spacing; there are two electrodes

per wavelength in the SWDEC, greater than two electrodes per wavelength in the TWDEC, and the wavelength decreases with the decreasing velocity of the particles. The analytical approximation determines the lifetime of the bunches as a function of the beam current and other parameters; a higher beam current leads to increased space-charge expansion. To formulate the approximation for expansion, a spherical ion bunch is considered. Acceleration of a single ion on the edge of an ion bunch due to the bulk charge of the bunch is calculated. The acceleration will be dependent on the ion charge, the bunch charge, the ion mass, and the time-dependent radius of the bunch:

$$\ddot{r} = \frac{F}{m_i} = \frac{1}{m_i} \frac{1}{4\pi\epsilon_0} \frac{Q_i Q_b}{r^2} \quad (8.1)$$

where Q_i is the charge of a single ion and Q_b is the total charge of the bunch. Assuming all ions have an ionization level of Z the differential equation can be written in terms of the number of ions in the bunch N_b as

$$\ddot{r} = \frac{1}{m_i} \frac{1}{4\pi\epsilon_0} \frac{N_b (Ze)^2}{r^2}. \quad (8.2)$$

The solution to this differential equation in terms of the initial radius r_0 and radius r_τ at time τ is

$$\sqrt{2 \frac{1}{m_i} \frac{1}{4\pi\epsilon_0} N_b (Ze)^2} \tau = r_0^{\frac{3}{2}} \left\{ \frac{r_\tau}{r_0} \sqrt{1 - \frac{r_0}{r_\tau}} + \frac{1}{2} \ln \left[2 \frac{r_\tau}{r_0} \left(1 + \sqrt{1 - \frac{r_0}{r_\tau}} \right) - 1 \right] \right\} \quad (8.3)$$

and shows the trade-off between the number of ions in each bunch and the time over which the bunch may be decelerated subject to the initial and final bunch

sizes which are correlated to the electrode spacing. The derivation of Eq. (8.3) is detailed in Appendix A. The approximation assumes a spherically expanding bunch.




Figure 8.5: 2D axisymmetric particle-in-cell simulation developed in [51] of the expansion of an initially spherical ion bunch. Left: symmetric expansion in the absence of a magnetic field. Middle: radial expansion limited by an axial magnetic field increases the rate of axial expansion. Right: comparison with the theoretical ion bunch radius.

However, the radial expansion of the bunch is limited by the axial magnetic field. A comparison between the expansion as modeled by Eq. (8.3) with particle-in-cell simulations of the expansion with and without an axial magnetic field are shown in Fig. 8.5. Early on in the expansion (up until the bunch radius has grown by a factor of about 2.5) the spherically symmetric and radially limited cases agree to within 10%. Typically the bunch spacing is on the order of 4 times that of the bunch length, so though the analytical model for bunch expansion becomes inaccurate

(off by approximately 20%) near the end of the expansion process, it provides an estimate on the bunch lifetime for the purpose of quick simulation. Additionally, the decelerator electrodes can have a compressing effect on the bunches (evident in simulation of the deceleration process in Fig. 8.14) thereby increasing ion bunch lifetimes. Considering the time at which the ion bunch enters the decelerator region as $t = 0$ and the time at which it leaves the region as $t = \tau$, τ is related to number of electrodes and the oscillation frequency as $\tau = \frac{N_e}{2f}$ for two electrodes per wavelength, where N_e is the number of decelerator electrodes. The number of ions per bunch is related to the beam current I_b by $I_b = N_b Z e f$, where f is the both the frequency of oscillation as well as the frequency of the passing ion bunches. Finally, r_τ is replaced with r_f , the final radius of the bunch upon leaving the decelerator region. Making these substitutions into Eq. (8.3) results in

$$\sqrt{2 \frac{1}{m_i} \frac{1}{4\pi\epsilon_0} I_b Z e} \frac{N_e}{f^{\frac{3}{2}}} = r_0^{\frac{3}{2}} \left\{ \frac{r_f}{r_0} \sqrt{1 - \frac{r_0}{r_f}} + \frac{1}{2} \ln \left[2 \frac{r_f}{r_0} \left(1 + \sqrt{1 - \frac{r_0}{r_f}} \right) - 1 \right] \right\}. \quad (8.4)$$

Eq. (8.4) relates the beam current to the oscillation frequency, the initial and final radii of the expanding ion bunches, and the number of decelerator electrodes. This expression is used in the optimization scheme to ensure the model stays in a region of realistic ion bunch lifetimes in section 8.3.13.

8.3.5 1D1V simulation overview

In the SWDEC, electrodes are connected in an “even/odd” fashion (Fig. 8.6). The simulation is performed via a discrete time-stepping scheme. Prior to the




Figure 8.6: The circuitry schematic for an SWDEC with eight electrodes. The system is an RLC circuit with the odd/even electrodes acting as a capacitor. Converted energy from the decelerating ions is stored in the inductor and capacitive electrodes, and dissipated in the resistor.

simulation start, the capacitance matrix is calculated, as outlined in section 8.3.6.1.

From the capacitance matrix, the aggregate capacitance between the odd and even electrodes is calculated, also outlined in section 8.3.6.1. After the simulation start, the following steps are performed during each discrete time-step:

1. The capacitance between the ion bunch and each electrode is calculated, as in outlined section 8.3.6.2.
2. The capacitive voltage (i.e. the voltage difference between the odd and even electrodes) is calculated from the ion bunch–electrode capacitance as well as any capacitive charge, also in section 8.3.6.2.
3. The current and charge of the RLC circuit are advanced over the time-step using the circuit equation as outlined in section 8.3.7.
4. The charge on each individual electrode is calculated using known voltages

and the capacitance matrix equation (found in section 8.3.6.1).

5. The electric potential along the axis is calculated from the charge on each electrode as outlined in section 8.3.8.
6. The acceleration of an ion bunch from the electric potential is calculated, also outlined in section 8.3.8. The velocity and position are advanced over the time-step accordingly.

8.3.6 Determination of electrode charge distribution

8.3.6.1 Charge on each electrode from a voltage difference

To determine the quantity of charge manifested on each electrode when the capacitor system is charged to a voltage, the capacitance of every possible electrode pair must be calculated. These capacitances form a capacitance matrix $\mathbf{C}$ which corresponds to the capacitance equation $\mathbf{q} = \mathbf{C}\mathbf{V}$, where $\mathbf{q}$ and $\mathbf{V}$ are vectors containing the respective charges and voltages of each electrode, C_{ij} is the capacitance between the i^{th} and j^{th} electrode, and C_{ii} is the self-capacitance of electrode i . The non-diagonal values of $\mathbf{C}^{-1}$ (also known as the elastance matrix) can be found by assuming a charge on ring i and finding the resulting potential on ring j . The thickness of each ring is assumed to be negligible in comparison to the distance between rings. This is not just a simplifying assumption: decreasing the thickness of the rings increases their coupling with the ion bunches while decreasing their coupling with each other. The potential on any point of ring j resulting from a charge dq on

an infinitesimal segment of ring i is:

$$d\Phi_j = \frac{1}{4\pi\epsilon_0 r_{ij}} dq_i \quad (8.5)$$

where r_{ij} is the distance from the charge on ring i to an arbitrary point on ring j as shown in Fig. 8.7, and is defined by

$$r_{ij}^2 = (z_i - z_j)^2 + 2R^2(1 - \cos \theta). \quad (8.6)$$

Plugging this value for r into Eq. (8.5) and integrating around ring i results in




Figure 8.7: An infinitesimal increase in potential $d\Phi$ on ring j due to a charge dq on an infinitesimal segment of electrode ring i . Due to axial symmetry and the assumption that each electrode ring is equipotential, the position of $d\Phi$ can be chosen for convenience.

$$\Phi_j = \frac{q_j}{4\pi\epsilon_0} \int_0^{2\pi} [(z_i - z_j)^2 + 2R^2(1 - \cos\theta)]^{-\frac{1}{2}} d\theta \quad (8.7)$$

and the non-diagonal elements of the elastance matrix are then

$$S_{ij} = \frac{1}{4\pi\epsilon_0} \int_0^{2\pi} [(z_i - z_j)^2 + 2R^2(1 - \cos\theta)]^{-\frac{1}{2}} d\theta \quad (8.8)$$

with $S_{ij} = S_{ji}$ for all i and j . These integrals are best calculated numerically. The self-elastance of each electrode can be found using the analytical expression [52]

$$S_{ii} = \frac{\ln(\frac{8R}{a})}{(4\pi\epsilon_0)(\pi R)} \quad (8.9)$$

where a is the radius of thickness of the electrode. This expression is valid for $R \gg a$. The odd and even electrodes form a capacitor. The following steps outline the calculation of the total capacitance from the capacitance matrix of the two sets of equipotential electrodes in the SWDEC. The matrix system $\mathbf{q} = \mathbf{C}\mathbf{V}$ can be ordered such that

$$\begin{bmatrix} \mathbf{q}_o \\ \mathbf{q}_e \end{bmatrix} = \begin{bmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Y} & \mathbf{Z} \end{bmatrix} \begin{bmatrix} \mathbf{V}_o \\ \mathbf{V}_e \end{bmatrix} \quad (8.10)$$

where a subscript o denotes odd electrodes and a subscript e denotes even electrodes:

$\mathbf{q}_o = [q_1, q_3 \dots]^T$ and $\mathbf{q}_e = [q_2, q_4 \dots]^T$; $\mathbf{X}$, $\mathbf{Y}$, and $\mathbf{Z}$ correspond to the quadrants of the symmetric capacitance matrix; and $\mathbf{V}_o = [V_o, V_o \dots]^T$ and $\mathbf{V}_e = [V_e, V_e \dots]^T$.

Summing equipotential rows results in

$$\begin{aligned}\sum \mathbf{q}_o &= xV_o + yV_e \\ \sum \mathbf{q}_e &= yV_o + zV_e\end{aligned}\tag{8.11}$$

where $x \equiv \sum \mathbf{X}$, $y \equiv \sum \mathbf{Y}$, and $z \equiv \sum \mathbf{Z}$. Because $\sum \mathbf{q}_o = -\sum \mathbf{q}_e \equiv q$, a capacitance matrix for the odd and even electrodes is revealed:

$$\begin{bmatrix} q \\ -q \end{bmatrix} = \begin{bmatrix} x & y \\ y & z \end{bmatrix} \begin{bmatrix} V_o \\ V_e \end{bmatrix}\tag{8.12}$$

This determines the voltage difference between the odd and even electrodes as

$$V_o - V_e = \frac{x + 2y + z}{xz - y^2}q\tag{8.13}$$

which defines the capacitance as

$$C = \frac{xz - y^2}{x + 2y + z}.\tag{8.14}$$

8.3.6.2 Voltage difference between the electrodes induced by a nearby ion bunch

The introduction of a positively charged ion bunch, approximated as a point charge on the axis, will raise the voltage of nearby electrodes. This effect can be

represented as

$$\mathbf{V} = \mathbf{C}^{-1}\mathbf{q} + \mathbf{C}_Q^{-1}Q \quad (8.15)$$

where $\mathbf{C}_Q = [C_{1Q}, C_{2Q}, \dots]^T$ is a column vector containing the capacitance between the ion bunch and the electrodes and Q is the charge of the ion bunch. The ion bunches are always on-axis, thus all parts of the electrode ring are equidistant from the ion bunch. The voltage induced on any electrode by an ion bunch of charge Q is therefore given by

$$\Phi_i = \frac{Q}{4\pi\epsilon_0 r_{iQ}} \quad (8.16)$$

Where r_{iQ} , the distance between the ion bunch and a point on the electrode ring, is defined by

$$r_{iQ}^2 = (z_i - z_Q)^2 + R^2 \quad (8.17)$$

so the elastance between a ring electrode and the ion bunch becomes

$$S_{iQ} = \frac{1}{4\pi\epsilon_0} [(z_i - z_Q)^2 + R^2]^{-\frac{1}{2}} \quad (8.18)$$

and the capacitance is found by taking an element-wise inverse of the elastance vector (i.e. $\mathbf{C}_Q = [S_{1Q}^{-1}, S_{2Q}^{-1}, \dots]$). For reasons that will become apparent later, Eq. (8.15) is rearranged as

$$\mathbf{q} = \mathbf{C}\mathbf{V} - \mathbf{C}\mathbf{C}_Q^{-1}Q. \quad (8.19)$$

The rows are rearranged in the same manner as Eq. (8.10)

$$\begin{bmatrix} \mathbf{q}_o \\ \mathbf{q}_e \end{bmatrix} = \begin{bmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Y} & \mathbf{Z} \end{bmatrix} \begin{bmatrix} \mathbf{V}_o \\ \mathbf{V}_e \end{bmatrix} - \begin{bmatrix} \mathbf{\Gamma}_o \\ \mathbf{\Gamma}_e \end{bmatrix} Q \quad (8.20)$$

defining $\mathbf{\Gamma} \equiv \mathbf{C}\mathbf{C}_Q^{-1}$, with $\mathbf{\Gamma}_o$ designating the odd rows of $\mathbf{\Gamma}$ and $\mathbf{\Gamma}_e$ designating the even rows of $\mathbf{\Gamma}$. A summing of equipotential rows and a rearrangement results in

$$\begin{bmatrix} V_o \\ V_e \end{bmatrix} = \begin{bmatrix} x & y \\ y & z \end{bmatrix}^{-1} \left(\begin{bmatrix} q \\ -q \end{bmatrix} + \begin{bmatrix} \gamma_o \\ \gamma_e \end{bmatrix} Q \right) \quad (8.21)$$

with $\gamma_o \equiv \sum \mathbf{\Gamma}_o$ and $\gamma_e \equiv \sum \mathbf{\Gamma}_e$. The voltage difference between the odd and even electrodes can then be expressed as

$$\Delta V = \frac{q}{C} + f(\mathbf{C}_Q) \quad (8.22)$$

with C defined in Eq. (8.14) and

$$f(\mathbf{C}_Q) \equiv \begin{bmatrix} x & y \\ y & z \end{bmatrix}^{-1} \begin{bmatrix} \gamma_o \\ \gamma_e \end{bmatrix} Q. \quad (8.23)$$

8.3.7 The circuit equation

The differential equation describing an RLC circuit is

$$L\ddot{q} + R\dot{q} + \Delta V = 0 \quad (8.24)$$

where ΔV is the potential difference between the two terminals of the capacitor. In the SWDEC, the terminals of the capacitor are the even and odd electrodes, and the potential difference is a result of both the total electrode charges ($+q$ and $-q$) as well as any potential difference induced by an ion bunch:

$$L\ddot{q} + R\dot{q} + \frac{q}{C} + f(\mathbf{C}_Q) = 0 \quad (8.25)$$

where L and R are the inductance and resistance of the circuit, with $\dot{q} \equiv dq/dt$ and $\ddot{q} \equiv d^2q/dt^2$. For multiple ion bunches the last term will be replaced by a term for each bunch. This modification is trivial, so for simplicity this derivation will continue to assume the presence of only one ion bunch. A vector that determines the state of the circuit and its time derivative are

$$\mathbf{x} = \begin{bmatrix} q \\ \dot{q} \end{bmatrix}, \quad \dot{\mathbf{x}} = \begin{bmatrix} \dot{q} \\ -\frac{R}{L}\dot{q} - \frac{1}{LC}q + \frac{1}{L}f(\mathbf{C}_Q) \end{bmatrix}. \quad (8.26)$$

In a time-stepping simulation the state vector can be updated at each time-step by

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \frac{\dot{\mathbf{x}}^{n+1} + \dot{\mathbf{x}}^n}{2} \Delta t \quad (8.27)$$

and from this the solution to the state vector at each new time-step is

$$\begin{bmatrix} q \\ \dot{q} \end{bmatrix}^{n+1} = \begin{bmatrix} 1 & -\frac{\Delta t}{2} \\ \frac{\Delta t}{2LC} & 1 + \frac{\Delta t R}{2L} \end{bmatrix}^{-1} \left\{ \begin{bmatrix} 1 & \frac{\Delta t}{2} \\ -\frac{\Delta t}{2LC} & 1 - \frac{\Delta t R}{2L} \end{bmatrix} \begin{bmatrix} q \\ \dot{q} \end{bmatrix}^n + \begin{bmatrix} 0 \\ f(\mathbf{C}_Q)^n \end{bmatrix} \Delta t \right\} \quad (8.28)$$

where the assumption $f(\mathbf{C}_Q)^{n+1} \approx f(\mathbf{C}_Q)^n$, has been made for simplicity, which is a valid approximation for small values of Δt .

8.3.8 Ion bunch deceleration due to charged electrodes

When the voltage difference between the electrodes is known, the charge on each electrode can be calculated as described in section 8.3.6.1. Because the system $\mathbf{q} = \mathbf{C}\mathbf{V}$ is linear, the charge distribution over the electrodes can be calculated just once for a given electrode setup and a test charge and then scaled for later use. The calculation of the electric potential at a point along the axis due to the charged electrodes is similar to the steps taken in Eq. (8.16), Eq. (8.17), and Eq. (8.18). The potential at an axial point z from a charge dq on an infinitesimal segment of electrode i is (see Fig. 8.8)

$$d\Phi(z) = \frac{dq_i}{4\pi\epsilon_0} [(z - z_i)^2 + R^2]^{-\frac{1}{2}}. \quad (8.29)$$

Integrating each side of the equation is trivial due to the symmetry of the problem,




Figure 8.8: A potential increase $d\Phi$ at location z due to a charge dq on an infinitesimal segment of electrode ring i .

and so the potential from all electrodes is

$$\Phi(z) = \sum_{i=1}^{N_e} \frac{q_i}{4\pi\epsilon_0} [(z - z_i)^2 + R^2]^{-\frac{1}{2}} \quad (8.30)$$

where N_e is the total number of electrodes. The electric field at point z is related to the derivative of the electric potential, and assuming there is an ion bunch at this location the acceleration of the ion bunch is

$$a = -\frac{Q}{m_i} \frac{d}{dz} \Phi(z_{\text{bunch}}) \quad (8.31)$$

where Q is the charge of the ion bunch and $\Phi(z_{\text{bunch}})$ is the potential due to all electrode rings at the location of the ion bunch as found in Eq. (8.30). The velocity and displacement of the bunch are then updated accordingly.

8.3.9 Partial validation of the model through demonstration of conservation of energy

Sections 8.3.6 through 8.3.8 have outlined the method by which the SWDEC is simulated. As a partial validation of these methods, the conservation of energy in the system is demonstrated. In a test case (Fig. 8.9) a single bunch traverses the decelerator electrode region. The decelerator electrodes are initially uncharged but the capacitance between the moving ion bunch and electrodes induces a charge in the electrodes which starts an oscillation in the RLC circuit. The oscillation is dampened by the resistor, and the total energy dissipated in the resistor is equal to

the loss in kinetic energy of the ion bunch. Though the conversion efficiency in this test case is low due to the lack of an initial oscillation of the circuit, accurate energy conservation is demonstrated.




Figure 8.9: Demonstration of the conservation of energy: The ion bunch enters a region of eight equally spaced decelerator electrodes shortly after 4 microseconds into the simulation, and excites the RLC circuit, where the bunch energy is transferred into an oscillation alternating between the inductor and capacitor. The bunch leaves the decelerator region shortly before 6 microseconds into the simulation, and the oscillating circuit energy is dampened and dissipated by the resistor. The total energy remains unchanged throughout the simulation.

8.3.10 1D1V electrode spacing optimization

During the optimization of the electrode positions, the effect of the charged bunch on the RLC circuit is disregarded and circuit amplitude damping due to the resistor is removed. The physics that remains is the deceleration of a single test ion, with a prescribed charge-to-mass ratio, by a constant amplitude LC circuit. This test ion has negligible charge relative to the oscillating charge on the electrodes, so the effect of the ion on the circuit is negligible. However, the effect of the circuit and electrodes on the ion is non-negligible, and any energy lost by the ion must be

gained by the circuit (since there is no other mechanism for loss in this closed, idealized system). This single-ion simplification is an imitation of the electric field that ions will experience when the SWDEC is operating in steady-state. The expected outcome is the need for the downstream electrodes to be more closely spaced so that the circuit oscillation remains synchronized with the transit of the decelerating ions.

The goal is to reach a situation for which the ions only experience an “uphill” potential while in the decelerator section. Starting with a nominal electrode spacing (a constant electrode spacing based on initial ion velocity and operating frequency) an ion is decelerated slightly by a small amplitude LC circuit. The position of the test charge each time the polarity of the capacitor charge switches (from positive to negative or vice versa) is noted, and on the next iteration the electrodes are moved to these positions. The new positions will be only slightly shifted from the old positions if the process is stable. Through this process, the circuit amplitude required to achieve this deceleration is specified and can be chosen freely. The determination of optimal resistance requires no further simulation, only calculation. During the optimization of the electrode spacing, the circuit equation, Eq. (8.25) with the resistor removed and neglecting any effect of the ion bunches on the circuit, reduces to

$$L\ddot{q} + \frac{q}{C} = 0. \quad (8.32)$$

A solution to Eq. (8.32) is

$$\dot{q} = \dot{q}_0 \sin(2\pi ft) \quad (8.33)$$

where $\dot{q}_0$ is the amplitude of the circuit (in amperes) at which the electrode spacing




Figure 8.10: All units arbitrary. A demonstration of the electrode spacing optimization. Over each iteration the circuit amplitude is increased, and the electrode spacing is modified to correspond with the deceleration of the test particle. By the 84th iteration, the conversion efficiency has achieved approximately 90%.

was optimized in section 8.3.10 and f is the frequency of both the incoming bunches and the frequency of oscillation of the circuit, $2\pi f = (LC)^{-1/2}$. This is the solution to the circuit oscillation with no resistance and no influence from passing ion bunches.

8.3.11 Circuit resistance calculation for steady-state operation

Now a system with the resistor in place as well as the influence of passing ion bunches is considered. In steady-state the energy lost by the decelerating ions is transferred into the circuit and dissipated by the resistor. This means that in Eq. (8.25) $R\dot{q} = -f(\mathbf{C}_q)$, so that Eq. (8.25) simplifies to Eq. (8.32) and therefore Eq. (8.33) is a valid solution for the full system as well. The next step is to find the resistance which results in $R\dot{q} = -f(\mathbf{C}_q)$, so that the system will operate at the amplitude specified in section 8.3.10. To find this resistance, the relation between resistance, power, and current is used:

$$P = \dot{q}^2 R. \quad (8.34)$$

Plugging in the solution to the LC differential equation (Eq. (8.33)) which is also desired as the solution to the steady-state RLC system, results in

$$P = \dot{q}_0^2 \sin^2(2\pi ft) R \quad (8.35)$$

and time-averaging over one period simplifies to

$$\langle P \rangle = \frac{\dot{q}_0^2}{2} R. \quad (8.36)$$

Here the time-averaged circuit power is (neglecting other energy losses) beam power multiplied by the efficiency

$$\langle P \rangle = \eta P_{\text{beam}} \quad (8.37)$$

and so the resistance required for the energy conversion efficiency specified by the optimization process is

$$R = \frac{2\eta P_{\text{beam}}}{\dot{q}^2} \quad (8.38)$$

8.3.12 Demonstration of a self-consistent steady-state simulation

To demonstrate a working system in steady-state, a series of ion bunches representing the modulated ion beam is sent through the decelerator electrodes. To reach steady-state, the oscillations of the circuit must initially be externally established, at a one-time start-up energy cost. Once the first of the series of ion bunches passes through the decelerator electrodes, the power source can be turned off and the simulation will become self-sustaining with only the ion beam power as the input. In steady-state the energy lost by the decelerating ions is transferred into the RLC circuit and balances the energy dissipated in the resistor. In reality, the converted energy only comes from the decrease in kinetic energy associated with a decrease in the axial velocity of the ion bunches. Any kinetic energy associated with radial or azimuthal velocity of the ions (which can arise due to an axial magnetic field) as well as any thermal energy, will not be converted by this process.




Figure 8.11: A frame-by-frame demonstration of the steady-state operation of the simulation, with asterisks denoting the axial positions of the ion bunches. The operation of the SWDEC in this simulation is self-sustaining, in that the only power input is the incoming ion bunches. The energy gained by the circuit from the decelerating bunches is offset by the energy dissipated in the resistor, and so the amplitude stays constant. This simulation demonstrates what is illustrated in Fig. 8.2: the ion bunches only experience “uphill” potentials while in the decelerator.

8.3.13 Analytical efficiency optimization accounting for ion bunch expansion

The model may be deemed valid insofar as the bunches do not expand to the extent that the leading edge of one ion bunch overtakes the trailing edge of the preceding ion bunch. In other words, the final distance between bunch centers, (equivalent to wavelength) λ_f , may not be less than twice the final bunch radius, r_f (from Eq. (8.4)), or

$$\lambda_f \geq 2r_f. \quad (8.39)$$

Assuming that optimal cases are reached when Eq. (8.39) is an equality, the final bunch velocity is related to the oscillation frequency and final bunch radius by

$$v_f = \lambda_f f = 2r_f f. \quad (8.40)$$

The energy conversion efficiency is then

$$\eta = \frac{E_i - E_f}{E_i} = 1 - \frac{v_f^2}{v_i^2} \quad (8.41)$$

where E_i and E_f are the initial and final ion energies respectively and v_i is the initial ion velocity, so then

$$\eta = 1 - \frac{4r_f^2 f^2}{v_i^2}. \quad (8.42)$$

For a given ion energy, beam radius, and beam current there exists an optimal operating frequency which can be found from the preceding equations. The steps for analytically finding the optimal operating frequency are briefly summarized below:

1. For a chosen beam current, beam energy and beam radius, choose a range of frequencies over which the optimal frequency is expected to be found. It is assumed that the beam radius corresponds to the initial bunch radius and is maintained by an axial magnetic field.
2. Sweep over all frequencies and solve Eq. (8.4) implicitly for the final bunch radius at each frequency.
3. Use Eq. (8.42) to find the efficiency at each frequency. Choose the frequency

that results in the highest efficiency.

The maximum efficiency for a given beam current, beam energy, and beam radius is now known. The electrode spacing can then be optimized to achieve this efficiency using the analytical-numerical method outlined in section [8.3.10](#).

8.3.14 1D1V optimization results

Three important approximations that may have an effect on the validity of the model were made:

1. In the interaction between the ion bunches and electrodes, the bunches are treated as point charges when in reality the bunches will occupy a finite volume. The reduced radial extent of the bunch that results from this approximation tends to underestimate the bunch-electrode capacitance, while the reduced axial extent results in an overestimate.
2. The ion bunches are assumed to be initially spherical with a radius equal to the beam radius, but as previously discussed the ion bunches will have shapes dependent on the ratio of the device length to the beam radius in any particular SWDEC/TWDEC. This may invalidate the model in regimes where the wavelength to beam radius ratio is far from unity.
3. In the analytical model of ion bunch expansion, the bunches were assumed to expand spherically but will actually expand primarily axially due to the axial magnetic field. This results in the expansion possibly being underestimated.

However, this expansion also does not take into account the electric field from the decelerator electrodes which may tend to compress the ion bunch, possibly resulting in the expansion being overestimated.

Fig. 8.12 shows the optimization results for a beam energy of 3 MeV and a beam radius of 10 cm. For each beam current, the oscillation frequency is chosen so as to optimize the efficiency. To ensure stability of the model, the maximum allowable efficiency is set to 90%. In the plot for converted power, a maximum for each number of electrodes exists. This maximum corresponds to the ideal beam current at which the SWDEC should operate for maximum power output. The




Figure 8.12: Optimization of efficiency as a function of beam current. Efficiency is capped at 90% to allow accurate calculation of the decelerator region length and other parameters.

fact that increasing the beam current above a certain value decreases the power output warrants explanation. A larger beam current corresponds to higher ion bunch densities, which increases the rate of ion bunch expansion. The increased ion bunch

expansion rate requires the length of the decelerator region to be decreased so that the ion bunches spend less time decelerator region. This shorter decelerator region length corresponds to a smaller difference between the initial electrode spacing and the final electrode spacing (because there is a lower limit on final electrode spacing) which corresponds to a smaller decrease in the velocity of the ion bunches. A smaller decrease in velocity is equivalent to a lower deceleration efficiency. Above a certain value for beam current this decreasing efficiency overcomes the benefit of increasing the beam power input (beam current) such that the power output decreases.

The electrode spacing, while variable, is on the order of the decelerator region length divided by the number of decelerator electrodes. As long as the electrode spacing is on the order of the electrode radius, the model may be deemed valid. The case shown (Fig. 8.12) remains in this regime. To accommodate larger beam currents (and larger output powers) the electrode radius may be increased.

In Fig. 8.12 a trade-off between output power and circuit amplitude (operating voltage) becomes apparent. While a lower number of electrodes corresponds to a larger output power, this also results in an increase in circuit amplitude. Additionally, increased output power can also mean increased oscillation frequency. Trends for other beam energies and radii can be calculated as well, and while they are similar to what is shown, the electrode spacing does not always remain comparable to the bunch radius.

8.4: A 2D3V particle-in-cell simulation of the SWDEC

The details of the 2D3V PIC method covered in Chap. 4 will not be repeated here. The components of the PIC method unique to this chapter are the addition of an RLC circuit equation to the solution of Poisson's equation, and the calculation of the magnetic field due to a solenoid.

8.4.1 Modeling of floating electrodes

The discrete Poisson equation for a particle in cell simulation is

$$\mathbf{A}\vec{\Phi} = \frac{-\vec{\rho}}{\epsilon_0}. \quad (8.43)$$

The specifics of this equation in a 2D3V domain are discussed in Sec. 4.2.2, and this section introduces the concept and modeling of floating electrodes. In the SWDEC, the decelerator electrodes are not at a known potential, nor can they be solved for as a standalone potential point like the rest of the unknown potentials on the simulation grid. Instead, each grid point on a single floating electrode is constrained to a single voltage, and the charge at each point varies to maintain this voltage, though the net charge stays the same. For the case of the SWDEC, there are two unknown potentials V_e and V_o corresponding to the even and odd electrodes respectively, and there are also unknown charges at each point on each electrode, which can be represented as column vectors $\vec{q}_e$ and $\vec{q}_o$. Without considering an

electrical connection between V_e and V_o , Eq. 8.43 takes the following form:

$$\begin{bmatrix} [\mathbf{A}]_{N_\Phi \times N_\Phi} & [\mathbf{S}]_{N_\Phi \times (N_e + N_o)} \\ [\mathbf{E}_e]_{(N_e - 1) \times N_\Phi} & [\mathbf{0}]_{(N_e - 1) \times (N_e + N_o)} \\ [\mathbf{E}_o]_{(N_o - 1) \times N_\Phi} & [\mathbf{0}]_{(N_o - 1) \times (N_e + N_o)} \\ [\mathbf{0}]_{2 \times N_\Phi} & [\mathbf{C}]_{2 \times (N_e + N_o)} \end{bmatrix} \begin{bmatrix} [\vec{\Phi}]_{N_\Phi \times 1} \\ [\vec{q}_e]_{(N_e) \times 1} \\ [\vec{q}_o]_{(N_o) \times 1} \end{bmatrix} = \frac{-1}{\epsilon_0} \begin{bmatrix} [\vec{\rho}]_{N_\Phi \times 1} \\ [\vec{0}]_{N_e \times 1} \\ [\vec{0}]_{N_o \times 1} \end{bmatrix} \quad (8.44)$$

where the sizes of sub-matrices and sub-vectors are denoted for convenience, $\mathbf{A}$ is taken from Eq. 8.43, $\mathbf{S}$ has the form $(\epsilon_0 \text{Vol})^{-1}$ and accounts for the effect of the $\vec{q}$'s on $\vec{\Phi}$, $\mathbf{E}_e$ and $\mathbf{E}_o$ ensure the equipotential condition on the parts of $\vec{\Phi}$ that are equal to V_e and V_o respectively, and $\mathbf{C}$ ensures conservation of charge on the electrodes: $\sum \vec{q}_e = 0$.

8.4.2 Implementation of the circuit equation

When the even and odd electrodes are connected with an inductor and resistor (neither of which is modeled physically within the domain) the circuit equation is given once again by Eq. 8.24. This equation is implemented directly into the Poisson matrix system (Eq. 8.44) by the addition of two more degrees of freedom: the charge difference between the electrodes $q \equiv \sum \vec{q}_e \equiv \sum \vec{q}_o$ and the current between the electrodes $\dot{q}$ where the voltage difference is $\Delta V \equiv V_e - V_o$. Once again using the form Eq. 8.26 and Eq. 8.27, the expressions for q and $\dot{q}$ at the new time-step are

labelqqdot

$$\dot{q}^{n+1} - \frac{\Delta t}{2} \ddot{q}^{n+1} = \dot{q}^n + \frac{\Delta t}{2} \ddot{q}^n \quad (8.45a)$$

$$\dot{q}^{n+1} - \frac{\Delta t}{2} \left(-\frac{R}{L} \dot{q}^{n+1} - \Delta V^{n+1} \right) = \dot{q}^n + \frac{\Delta t}{2} \left(-\frac{R}{L} \dot{q}^n - \Delta V^n \right) \quad (8.45b)$$

which simplifies to the matrix form

$$\begin{bmatrix} & & & \\ & & & \\ 0 & 0 & 1 & -\frac{\Delta t}{2} \\ \frac{\Delta t}{2} & -\frac{\Delta t}{2} & 0 & 1 + \frac{\Delta t R}{2L} \end{bmatrix} \begin{bmatrix} V_e \\ V_o \\ q \\ \dot{q} \end{bmatrix}^{n+1} = \begin{bmatrix} & & & \\ & & & \\ 0 & 0 & 1 & \frac{\Delta t}{2} \\ -\frac{\Delta t}{2} & \frac{\Delta t}{2} & 0 & 1 - \frac{\Delta t R}{2L} \end{bmatrix} \begin{bmatrix} V_e \\ V_o \\ q \\ \dot{q} \end{bmatrix}^n \quad (8.46)$$

which is then implemented directly into Eq. 8.44 where V_e and V_o may refer to any two nodes on an even or odd electrode respectively. The top two rows of each matrix are blank because the coefficients for determining V_e and V_o are already held in Eq. 8.44. The left hand side of the bottom two rows of Eq. 8.46 are inserted into the bottom two rows of A and $\vec{\Phi}$ with correct references to V_e and V_o nodes, and the right hand side is evaluated and inserted into the bottom two entries of b of Eq. 8.43 using information from the previous time-step.

8.4.3 Calculation of the magnetic field due to a solenoid

The SWDEC/TWDEC may require the use of a magnetic field to limit the expansion of bunches into the ring-shaped electrodes. The magnetic field is gen-

erated by a solenoidal current wire outside of the electrodes, and is calculated by discretizing the current wire into discrete current elements and using the Bio-Savart law to calculate the magnetic field. The solenoid is defined by the number of loops N_l , the radius R , the axial extent d and the current I . The solenoid is first discretized axially into N_r individual current rings, where the current of each ring is $I \frac{N_l}{N_r}$. Each current ring is then discretized azimuthally into N_d current elements, where the strength of each current element is

$$\mathbf{i}_{i,j} = I \frac{2\pi R N_l}{N_d N_r} (\hat{x} \sin \theta_i + \hat{y} \cos \theta_i) \quad (8.47)$$

where i ranges from 1 to N_d and j ranges from 1 to N_r and $\theta_i = \frac{2\pi i}{N_d}$. The locations of the elements are

$$\mathbf{x}_{i,j} = \hat{z} \left(z_0 + j \frac{d}{N_r} \right) + R (\hat{x} \cos \theta_i + \hat{y} \sin \theta_i). \quad (8.48)$$

The magnetic field at any point $\mathbf{x}$ is found by summing the contributions of each current element using the Bio-Savart law:

$$\mathbf{B}(\mathbf{x}) = \frac{\mu_0}{4\pi} \sum_j \sum_i \frac{\mathbf{i}_{i,j} \times (\mathbf{x} - \mathbf{x}_{i,j})}{|\mathbf{x} - \mathbf{x}_{i,j}|^3} \quad (8.49)$$

and in this way is used to calculate the magnetic field over the SWDEC domain, and example of which is shown in Fig. 8.13




Figure 8.13: A magnetic field resulting from two solenoids, discretized according to the black dots plotted in the domain

8.4.4 2D3V simulation results

The electrode spacing optimization that resulted from the 1D1V simulation of Sec. 8.3 was tested in the 2D3V axisymmetric SWDEC particle-in-cell code. A frame from this simulation is shown in Fig. 8.14. The electrodes were optimized for a conversion efficiency of 75%. In the particle-in-cell code, the actual conversion efficiency is 66%. The conversion efficiency is lower because the ions in the center of the beam experience a lower decelerating electric field than the ions closer to the electrodes. Compression of the ion bunches near the exit of the decelerator is evident. In this case the effect of the decelerating electric field is stronger than the force of space-charge expansion from the ion bunches, resulting in some axial compression of the bunches.




Figure 8.14: A frame from a particle-in-cell simulation of the decelerator electrodes of the SWDEC. This simulation served as a test of the electrode deceleration optimization using the particle-in-cell method. Particles are moving from left to right. The axial and radial axes are of different scales for clarity.

Chapter 9

Conclusion

This thesis contributes to the field of inertial electrostatic confinement fusion by introducing the Continuous Electrode IEC fusor as an advancement upon previous fusor concepts. Particle-in-cell, N -body, and fluid simulations were all developed specifically for investigating the CE-IEC and related fusors. The most successful simulation was the 2D3V axisymmetric skewed-grid PIC simulation that was parallelized for execution on a GPU and used for electrode voltage optimization. The optimizer “re-discovered” a voltage profile similar to that of the Multi-grid IEC that is the predecessor to the CE-IEC, and output a wall voltage profile tuned to maximize ion bunching behavior while minimizing ion losses.

The calculations and simulations of this thesis showed that the ion density in the acceleration region was space charge limited to the extent that the theoretical maximum fusion power output for a one-meter radius CE-IEC fusor was on the order

of one microwatt. Though the ion lifetimes averaged thousands of passes through the device, thermalization was found to continue, despite the kinematic constraints imposed by the optimal potential profile. It would appear that a passive mechanism for halting thermalization has not been achieved. The possibility of actively halting this process may still exist, however for the current design, this is not the limiting factor for fusion power generation.

9.1: Summary of contributions

The work of this thesis contributes to the fields of IEC fusion, non-thermal plasma simulation, and ion acceleration optimization. The conclusions drawn from each contribution may be summarized as follows:

- The calculations in Chap. 3 show that a megawatt CE-IEC is impossible due primarily to space charge limitations. However a low power CE-IEC aimed at neutron production for non-power source application (neutron imaging or medical isotope production) could mark an improvement in the rate of neutrons produced per unit power input. Further, scaling laws show that the space charge effects might be mitigated by scaling the CE-IEC to a small size to reach the desired power-per-unit volume. However, an acceptable power output is not reached until the device radius is on the order of 1 millimeter, which presents difficulties in manufacturing due to the high voltages involved.
- The 2D3V axisymmetric particle-in-cell simulation features:
 - a fully parallel execution on a GPU with minimal memory transfer be-

tween the CPU and GPU including pair-wise collision implementation

- an original Coulomb collision model that takes into account high-angle scatters

which enabled use of the simulation for:

- optimization of the voltages along the CE-IEC acceleration channels for minimizing ion loss
 - simulation of the long timescale thermalization of the ion bunches using optimized voltages.
- The N -body simulation, which successfully demonstrated the full 3D simulation of the CE-IEC with 16 beamlines, is a tool for:
 - electron confinement simulation in a spherical cusped magnetic field, demonstrating a spherical-shell electron distribution
 - a method of investigating the interaction between CE-IEC beamlines, which showed that ions transferred between beamlines have short lifetimes.
 - a method for visualizing the points of ion and electron impact on the CE-IEC surfaces, showing that the majority of impacts happen on the inner surface, necessitating a sputter shield in this location thermally insulated from the rest of the device.
 - The Scharfetter-Gummel electron simulation method described in Chap. 6 provides a capability for calculating the steady-state electron distribution at each

ion time-step, and is a first step in the simulation high- β electron confinement in the CE-IEC core.

- The standing-wave direct energy converter (SWDEC) was presented, and two simulations served a dual purpose:
 - A 1D1V semianalytical simulation was developed for optimizing the electrode spacing in the SWDEC, which is necessary for an effective energy conversion efficiency
 - A 2D3V PIC simulation used the optimized electrode spacing to demonstrate the direct conversion of kinetic energy into electricity at an efficiency of more than 50% for mono-energetic fusion products.

9.2: Problems that still need solutions

Some inherent limitations were discussed in Chap. 4, mostly due to the space charge limitation and the difficulties of ion beam neutralization via electrons. The most prominent problem is the Child-Langmuir limitation on the acceleration of ions in the acceleration region. The limitation on density is also seen in the simulations, where the maximum density reached by a one-meter radius CE-IEC is on the order of 10^{14} m^{-3} .

Concerning the simulation methods, an obvious improvement that must be made is the simulation of two fuel species instead of one. While it is not difficult to add a new species, the optimization becomes more complicated due to the different reactions to the electrostatic potential geometry of each species, and so optimizations

might become significantly more computationally intensive.

9.3: Recommendations for future work

Since every thesis must come to an end, there remain many ideas that never made it to fruition. These ideas are discussed below:

9.3.1 3D simulation

A 3D simulation of the IEC could be performed by taking advantage of the symmetry of the truncated icosahedron. The truncated icosahedron, even when modified so that the pentagonal and hexagonal faces have close to equal area, may be split into 120 symmetric pieces, as shown in Fig. 9.1. Simulation of a single




Figure 9.1: The truncated icosahedron can be split into 120 symmetric slices. One symmetric slice (raised area) contains part of a hexagon and part of a pentagon.

symmetric slice, then, would effectively simulate the entire domain requiring only $1/120^{\text{th}}$ of the grid size and particle count of a full simulation. The geometry of the domain does not easily permit a structured grid, so the best method would likely be an unstructured particle-in-cell simulation using the finite element method for field value solutions. A diagram of the 3D domain of this proposed simulation is shown in Fig. 9.2.




Figure 9.2: A single symmetric slice of the truncated icosahedron IEC with the wall sections shown.

9.3.2 Introduction of optimization degrees-of-freedom

Full exploration of optimization options was forgone in the interest of finishing this thesis, but the degrees of freedom in the CE-IEC are vast and a full optimization

over all of them is a formidable task. Possible optimization degrees-of-freedom include:

- *Ion density, or total number of confined ions*
- *Optimization of the overall size of the device* — The effects of changing the size scale were investigated in Sec. 3.9 but other effects could be present. A limitation would need be placed from below perhaps by maximum allowable electric field, and from above by maximum structural size.
- *Ratio of the inner radius to the outer radius*
- *Wall angle* — Though considering this as a continuously-free parameter is not valid, it may help decide which geometry (see Fig. 6.18) is best
- *Fuel ratio* — Two-species fuel has not been implemented in this simulation but if it is in the future, the fuel ratio will be another degree of freedom. The simulation will also need to take into account Bremsstrahlung radiation power loss if applicable.

9.3.3 Possible fast optimization by finding an unchanging initial particle distribution

The optimizer solved the problem of the effect of the initial particle positions on the cost function output by optimizing over a large number of oscillation periods until it was reasonable that the initial particle positions were no longer having any significant effect on the final particle positions, that is, the simulation had reached an

oscillatory steady state. As an alternative to this workaround, the optimizer could also optimize the initial particle phase-space distribution, with the cost function being the change in the particle distribution over one period. Though this would require many more degrees of freedom due to the complexity of the 5D $(x, r, v_x, v_r, v_\theta)$ phase space distribution in addition to the electrode voltages, the optimizer would only be required to execute one oscillation period for each iteration, making many more iterations possible. As another alternative, some combination between the method just described and the existing method could be concocted.

9.3.4 Global simplex method

Currently, the PIC code optimizes by alternating between the simulated annealing global optimizer and the Nelder-Mead simplex local optimizer, using MATLAB's optimization tools. As an alternative to this, a custom-built optimizer may be able to find a global minimum more quickly and more reliably. A starting point may be investigation into a multi-path global simplex optimizer, such as that in Ref. 53.

9.4: Summary on the difficulties of achieving net-power fusion in a CE-IEC

The primary issues preventing net-power generation in the IEC are the space-charge and thermalization issues. The CE-IEC space-charge issue is particularly difficult due to the pulsed nature of the device, which was introduced to mitigate the

thermalization issue. However in order to achieve the density required for fusion, either significant neutralization is required, or a shrinkage of the non-neutralized regions must occur. The long-term simulation of the device showed that though the beam-beam thermalization is slow due to the short amount of time that ions spend in the counter-streaming state, the effects of thermalization nonetheless accumulate over time and contribute to significant long-timescale ion loss even at low densities. Thermalization and fusion power both scale as the square of ion density, and so these trends can expect to continue even if the space charge problem is solved.

Chapter A

Derivation of ion bunch expansion

Consider a sphere of radius r_0 randomly and uniformly populated with ions of density n , with all ions initially at rest $T_i = 0$. The radial electric field at the initial state ($t = 0$) and at any radius $r \leq r_0$ is

$$E_r(r \leq r_0) = \frac{q_i n}{3\epsilon_0} r \quad (\text{A.1})$$

which causes the sphere to expand. Since the electric field (and therefore the acceleration of each ion in the sphere) inside the sphere depends on r linearly, the distance that each ion travels will have a linear dependence on that ion's initial distance r from the center of the sphere, so in this way the density of the expanding sphere will remain spatially uniform. For the rest of the derivation, r will be the radial position of an ion on the edge of the sphere, and is equivalent to the time-dependent radius

of the sphere. The electric field at the edge of the sphere is dependent on the total charge of the sphere Q

$$E_r(r) = \frac{Q}{4\pi\epsilon_0 r^2} \quad (\text{A.2})$$

and substituting in $Q = q_i n \frac{4}{3}\pi r_0^3$ results in

$$E_r(r) = \frac{q_i^2 r_0^3 n}{3\epsilon_0 m_i} \frac{1}{r^2} \quad (\text{A.3})$$

The acceleration of an ion at the edge of the sphere (and therefore the acceleration of the radius of the sphere) is

$$\ddot{r} = \frac{q_i}{m_i} E_r(r) \quad (\text{A.4})$$

and combining Eqs. (A.3) and (A.4) results in

$$\ddot{r} = \frac{q_i^2 r_0^3 n}{3\epsilon_0 m_i} \frac{1}{r^2}. \quad (\text{A.5})$$

Note that n is the initial density of the sphere (not the time-dependent density) and so by definition of the constant $k \equiv \frac{q_i^2 r_0^3 n}{3\epsilon_0 m_i}$ the second order non-linear ordinary differential equation to be solved is

$$\ddot{r} = \frac{k}{r^2}. \quad (\text{A.6})$$

Because it was assumed that there was no radial bunch velocity at time $t = 0$

($\dot{r}(0) = 0$) $\ddot{r}$ can be expressed as

$$\ddot{r} = \frac{d\dot{r}}{dt} \frac{dr}{dr} = \dot{r} \frac{d\dot{r}}{dr} \quad (\text{A.7})$$

and so the differential equation becomes

$$\dot{r} d\dot{r} = k \frac{1}{r^2} dr. \quad (\text{A.8})$$

Integrating each side

$$\int_0^{\dot{r}} \dot{r} d\dot{r} = k \int_{r_0}^r \frac{1}{r^2} dr \quad (\text{A.9})$$

results in a first order differential equation

$$\frac{\dot{r}^2}{2} = k \left[\frac{1}{r_0} - \frac{1}{r} \right] \quad (\text{A.10})$$

which, when rearranged, yields

$$dr = \sqrt{2k \left[\frac{1}{r_0} - \frac{1}{r} \right]} dt. \quad (\text{A.11})$$

Integrating again,

$$\int_{r_0}^{r_f} \frac{1}{\sqrt{2k \left[\frac{1}{r_0} - \frac{1}{r} \right]}} dr = \int_0^\tau dt \quad (\text{A.12})$$

and rearranging the integral results in

$$\sqrt{2k} \int_0^\tau dt = \int_{r_0}^{r_f} \left(\frac{1}{r_0} - \frac{1}{r} \right)^{-\frac{1}{2}} dr \quad (\text{A.13})$$

where τ is the amount of time the bunch takes to traverse the decelerator electrode region. Performing the integration on each side of Eq. (A.13) results in

$$\sqrt{2k} \tau = r_0^{\frac{3}{2}} \left\{ \frac{r}{r_0} \sqrt{1 - \frac{r_0}{r}} + \frac{1}{2} \ln \left[2r \left(1 + \sqrt{1 - \frac{r_0}{r}} \right) - r_0 \right] \right\} \Big|_{r_0}^{r_\tau} \quad (\text{A.14})$$

which, evaluated at the limits of integration, becomes

$$\sqrt{2k} \tau = r_0^{\frac{3}{2}} \left\{ \frac{r_\tau}{r_0} \sqrt{1 - \frac{r_0}{r_\tau}} + \frac{1}{2} \ln \left[2 \frac{r_\tau}{r_0} \left(1 + \sqrt{1 - \frac{r_0}{r_\tau}} \right) - 1 \right] \right\}. \quad (\text{A.15})$$

Substituting in the value of k and rearranging results in the time for an expansion of the sphere from radius r_0 to r_τ :

$$\tau = \sqrt{\frac{3 m_i \epsilon_0}{2 q_i^2 n}} \left\{ \frac{r_\tau}{r_0} \sqrt{1 - \frac{r_0}{r_\tau}} + \frac{1}{2} \ln \left[2 \frac{r_\tau}{r_0} \left(1 + \sqrt{1 - \frac{r_0}{r_\tau}} \right) - 1 \right] \right\}. \quad (\text{A.16})$$

Bibliography

- [1] D. Zajfman, O. Heber, L. Vejby-Christensen, I. Ben-Itzhak, M. Rappaport, R. Fishman, and M. Dahan. Electrostatic bottle for long-time storage of fast ion beams. *Phys. Rev. A*, 55(3):R1577–R1580, mar 1997.
- [2] T. J. McGuire and R. J. Sedwick. Improved confinement in inertial electrostatic confinement for fusion space power reactors. *Journal of Propulsion and Power*, 21(4):697–706, jul 2005.
- [3] Andrew M. Chap and Raymond J. Sedwick. One-dimensional semianalytical model for optimizing the standing-wave direct energy converter. *Journal of Propulsion and Power*, 31(5):1350–1361, sep 2015.
- [4] Robert L. Hirsch. Inertial Electrostatic Confinement of Ionized Fusion Gases. *Journal of Applied Physics*, 38(11):4522–4534, Oct 1967.
- [5] George H. Miley and S. Krupakar Murali. *Inertial Electrostatic Confinement (IEC) Fusion: Fundamentals and Applications*. Springer, 2013.
- [6] Jaeyoung Park, Nicholas A. Krall, Paul E. Sieck, Dustin T. Offermann, Michael Skillicorn, Andrew Sanchez, Kevin Davis, Eric Alderson, and Giovanni Lapenta. High-energy electron confinement in a magnetic cusp configuration. *Phys. Rev. X*, 5:021024, Jun 2015.
- [7] D. C. Barnes, T. B. Mitchell, and M. M. Schauer. Beyond the brillouin limit with the penning fusion experiment. *Physics of Plasmas*, 4(5):1745–1751, 1997.
- [8] Carl Dietrich, Raymond Sedwick, and Leslie Eurice. Experimental Verification of Enhanced Confinement in a Multi-Grid IEC Device. In *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, Hartford, CT, 2008.

- [9] Thomas J. McGuire and Raymond J. Sedwick. Improving iec particle confinement times with multiple grids. In *US-Japan IEC Workshop on Small Plasma and Accelerator Neutron Sources*, Argonne National Laboratory, 2007.
- [10] NASA technology roadmaps TA 3: Space power and energy storage. Technical report, 2015.
- [11] Todd H. Rider and Peter J. Catto. Modification of classical spitzer ion electron energy transfer rate for large ratios of ion to electron temperatures. *Physics of Plasmas*, 2(6):1873–1885, 1995.
- [12] Alex Klein. Marble: Multiple ambipolar recirculating beam line experiment. 13th US-Japan Workshop on Inertial Electrostatic Confinement Fusion, 2011.
- [13] Lyman Spitzer. *Physics of Fully Ionized Gases: Second Revised Edition*. Dover Publications, 2013.
- [14] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philosophical Magazine Series 6*, 21(125):669–688, may 1911.
- [15] Todd H. Rider. *Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [16] Wilhelmus M. Ruyten. Density-conserving shape factors for particle simulations in cylindrical and spherical coordinates. *Journal of Computational Physics*, 105(2):224 – 232, 1993.
- [17] Nikolaos Sfakianakis. *Finite Difference schemes on non-Uniform meshes for Hyperbolic Conservation Laws*. PhD thesis, University of Crete, 2009.
- [18] Jay P. Boris. Relativistic plasma simulation optimization of a hybrid code. In *Proceedings of the Conference on the Numerical Simulation of Plasmas (4th)*, November 1970.
- [19] C. K. Birdsall and A. B. Langdon. *Plasma physics via computer simulation*. New York: Taylor and Francis, 2005.
- [20] Tomonori Takizuka and Hirotada Abe. A binary collision model for plasma simulation with a particle code. *Journal of Computational Physics*, 25(3):205–219, nov 1977.
- [21] W.M. Nevins and R. Swain. The thermonuclear fusion rate coefficient for p- 11 b reactions. *Nuclear Fusion*, 40(4):865, 2000.
- [22] S.J. Aarseth. *Gravitational N-Body Simulations: Tools and Algorithms*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2003.

- [23] J. Makino. Optimal order and time-step criterion for Aarseth-type N-body integrators. *Astrophysical Journal, Part 1*, 369:200–212, March 1991.
- [24] J. Makino and S. J. Aarseth. On a Hermite integrator with Ahmad-Cohen scheme for gravitational many-body problems. *Publications of the Astronomical Society of Japan*, 44:141–151, apr 1992.
- [25] D.L. Scharfetter and H.K. Gummel. Large-signal analysis of a silicon read diode oscillator. *Electron Devices, IEEE Transactions on*, 16(1):64–77, Jan 1969.
- [26] G J M Hagelaar. Modelling electron transport in magnetized low-temperature discharge plasmas. *Plasma Sources Science and Technology*, 16(1):S57, 2007.
- [27] L. Andor, H.P. Baltes, A Nathan, and H.G. Schmidt-Weinmar. Numerical modeling of magnetic-field-sensitive semiconductor devices. *Electron Devices, IEEE Transactions on*, 32(7):1224–1230, Jul 1985.
- [28] K. Nanbu. Theory of cumulative small-angle collisions in plasmas. *Physical Review E*, 55(4):4642–4652, apr 1997.
- [29] Andris M. Dimits, Chiaming Wang, Russel Caflisch, Bruce I. Cohen, and Yanghong Huang. Understanding the accuracy of nanbus numerical coulomb collision operator. *Journal of Computational Physics*, 228(13):4881 – 4892, 2009.
- [30] M. Conte. Beam loss by single coulomb scattering in collider rings. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 228(2-3):236–239, jan 1985.
- [31] D. Wu, X. T. He, W. Yu, and S. Fritzsche. Monte carlo approach to calculate proton stopping in warm dense matter within particle-in-cell simulations. *Physical Review E*, 95(2), feb 2017.
- [32] Andrew M. Chap and Raymond J. Sedwick. Simulation of an inertial electrostatic confinement device using a hermite n-body individual time-step scheme. In *51st AIAA/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics (AIAA), jul 2015.
- [33] R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen differenzengleichungen der mathematischen physik. *Mathematische Annalen*, 100(1):32–74, dec 1928.
- [34] MATLAB. *version 9.1 (R2016b)*. The MathWorks Inc., Natick, Massachusetts, 2016.
- [35] Cleve B. Moler. *Numerical Computing with Matlab*. Society for Industrial & Applied Mathematics (SIAM), jan 2004.

- [36] Andrew M. Chap, Alfonso G. Tarditi, and John H. Scott. Numerical and experimental investigation on the traveling wave direct energy converter concept. In *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics (AIAA), jul 2014.
- [37] A.P.L. Robinson, D.J. Strozzi, J.R. Davies, L. Gremillet, J.J. Honrubia, T. Johzaki, R.J. Kingham, M. Sherlock, and A.A. Solodov. Theory of fast electron transport for fast ignition. *Nuclear Fusion*, 54(5):054003, apr 2014.
- [38] James D. Callen. Draft material for fundamentals of plasma physics book. <http://homepages.cae.wisc.edu/callen/book.html>. Accessed: 2017-03-28.
- [39] Michael Keidar and Isak Beilis. *Plasma Engineering: Applications from Aerospace to Bio and Nanotechnology*. Academic Press, 2013.
- [40] R. Alexandre and C Villani. On the landau approximation in plasma physics. *Annales de l'Institut Henri Poincaré (C) Non Linear Analysis*, 21(1):61–95, jan 2004.
- [41] H. W. Becker, C. Rolfs, and H. P. Trautvetter. Low-energy cross sections for $^{11}\text{B}(p, 3\alpha)$. *Zeitschrift für Physik A Atomic Nuclei*, 327(3):341–355, sep 1987.
- [42] Hiromu Momota. Direct energy conversion of 15 mev fusion protons. Technical report, Los Alamos Natl. Lab, 1990. LA-11808-C.
- [43] Robert W. Bussard. Some physics considerations of magnetic inertial-electrostatic confinement; A new concept for spherical converging-flow fusion. *Fusion Technology*, 19:273–293, 1990.
- [44] Vito Lancellotti, Riccardo Maggiora, Giuseppe Vecchi, Daniele Pavarin, Simone Rocca, and Cristina Bramanti. Radiofrequency plasma thrusters: Modelling of ion cyclotron resonance heating and system performance. In *43rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics, jul 2007.
- [45] Daniele Pavarin, Andrea Lucca Fabris, Fabio Trezzolani, Marco Manente, Martina Faenza, Francesca Ferri, Antonio Selmo, Konstantinos Katsonis, and Chloé Berenguer. Low power RF plasma thruster experimental characterization. In *48th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*. American Institute of Aeronautics and Astronautics, jul 2012.
- [46] David Brady, Harold White, Paul March, James Lawrence, and Frank Davies. Anomalous thrust production from an RF test device measured on a low-thrust torsion pendulum. In *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*. American Institute of Aeronautics and Astronautics, jul 2014.
- [47] R.W. Moir and W.L. Barr. “venetian-blind” direct energy converter for fusion reactors. *Nucl. Fusion*, 13(1):35–45, jan 1973.

- [48] H. Momota, A. Ishida, Y. Kohzaki, G. H. Miley, S. Ohi, M. Ohnishi, K. Yoshikawa, K. Sato, L. C. Steinhauer, and Y. Tomita. Conceptual design of D-He-3 FRC reactor ARTEMIS. *Fusion Science and Technology*, 21(4):2307–2323, July 1991.
- [49] John H. Scott, Jeffery A. George, and Alfonso G. Tarditi. Direct energy conversion for low specific mass in-space power and propulsion. Albuquerque, NM, February 2013. Nuclear and Emerging Technologies for Space, American Nuclear Society.
- [50] Ryoh Kawana, Motoo Ishikawa, Hiromasa Takeno, Takayoshi Yamamoto, and Yasuyoshi Yasaka. Performance analysis of small-scale experimental facility of TWDEC. *Energy Conversion and Management*, 49(10):2522–2529, oct 2008.
- [51] Andrew M. Chap, John Scott, and Alfonso Tarditi. A particle-in-cell simulation for the traveling wave direct energy converter (TWDEC) for fusion propulsion. In *49th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, San Jose, CA, July 2013. American Institute of Aeronautics and Astronautics.
- [52] TSE Thomas. The capacitance of an anchor ring. *Aust. J. Phys.*, 7(2):347, 1954.
- [53] Akbar Karimi and Patrick Siarry. Global simplex optimization a simple and efficient metaheuristic for continuous optimization. *Engineering Applications of Artificial Intelligence*, 25(1):48 – 55, 2012.

About the Author

Andrew M. Chap grew up in Chevy Chase, Maryland and received his B.A. in physics from the College of the Holy Cross in Worcester, Massachusetts. After spending some time trying to eek out a music career first in Maryland and later in New York City (some of his music can still be found on Spotify and around the internet under the moniker “Rearview Minor”) Andrew entered graduate school at the University of Maryland in the aerospace engineering department in 2011. He asked his advisor Raymond Sedwick for a research project and the result is this dissertation. With some luck Andrew became a NASA Space Technology Research Fellow for the years 2013-2017 and spent a total of about 15 months working at NASA Johnson Space Center in Houston, Texas on the Traveling Wave Direct Energy Converter project and the Q-thruster (also known as EM Drive) project.

August 19th, 2017 was the happiest day (so far) of Andrew’s life, when he married the love of his life and his best friend, Dasha. The publication date of this thesis would likely be in 2018 (or later) if not for Dasha’s encouragement. While writing this sentence, Andrew and Dasha are on an airplane one-way to Colorado, where Andrew will soon be starting a job at Tech-X Corporation. We couldn’t be more excited.