Electromagnetic Extraction and Annihilation of Antiprotons for Spacecraft Propulsion

by

Daniel Antonio Zayas

S.B. Aerospace Engineering, S.B. Physics
Massachusetts Institute of Technology, 2006

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2008

Author .....
Department of Aeronautics and Astronautics
May 23, 2008

Certified by .....
Prof. Raymond J. Sedwick
Visiting Assistant Professor of Aeronautics and Astronautics
Thesis Advisor

Certified by .....
James A. Bickford
Charles Stark Draper Laboratory
Thesis Supervisor

Accepted by .....
Prof. David L. Darmofal
Associate Department Head
Chair, Committee on Graduate Students

Electromagnetic Extraction and Annihilation of Antiprotons for Spacecraft Propulsion

by

Daniel Antonio Zayas

Submitted to the Department of Aeronautics and Astronautics
on May 23, 2008, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

A novel concept for the extraction and long-term trapping of antiprotons from planetary magnetospheres is developed. The excitation via rotating magnetic field of an electron population within a distributed neutral plasma is shown to produce a large-scale magnetic field approximating a dipole with r^{-2} or slower decay. These shallow field gradients are, in turn, found to suppress the natural tendency of the magnetic dipole field to shield against incident particle fluxes. Particle transport and plasma interaction models are developed and used to compare collection performance against various electrostatic and magnetostatic collection systems. Baseline system architectures are presented, and antiproton collection rates on the order of nanograms per day are estimated, far exceeding current Earth-based production rates.

To demonstrate the performance potential of antiproton fuels for spacecraft propulsion, a modified rocket equation with relativistic corrections and variable mass-energy conversion efficiency is derived. Antimatter engines operating via the catalysis of nuclear fission in a fissile material by antiprotons are found to be three to four $(3-4)$ times more efficient than traditional proton-antiproton annihilation or "beam-core" engines, and vastly more efficient than conventional chemical and electric propulsion systems. Baseline architectures for various missions are presented.

Thesis Advisor: Prof. Raymond J. Sedwick
Title: Visiting Assistant Professor of Aeronautics and Astronautics

Thesis Supervisor: James A. Bickford
Affiliation: Charles Stark Draper Laboratory

[This Page Intentionally Left Blank]

Acknowledgments

I would like to begin by thanking my advisor and friend Ray Sedwick, a constant source of encouragement and guidance. His willingness to serve as my advisor made my last three and a half years at MIT infinitely more tolerable. Though his escape to greener pastures came as a bit of a shock - an email from your advisor titled "Important Life-Changing Question" is almost certainly Bad News! - things seem to have worked out just fine. The Terrapins are certainly fortunate to have him.

I would also like to thank my supervisor at the Charles Stark Draper Laboratory, Jim Bickford, whose keen insight and bottomless enthusiasm were the backbone of this research. Though he is a mechanical engineer by trade, I like to think he is a physicist at heart. For the past two years, he has been a patient teacher and a great boss. Working with him has been a privilege. I owe a special debt of gratitude to two other Draper employees. Bill Schmitt was an integral part of this research effort, and a kindred spirit; an alumnus of the Department of Physics at MIT, he and I commiserated often over the nature of life at the Institute. Joe Bauer, my office-mate of two years, could always be counted on for engaging conversation and a cheerful disposition; the perfect start to a long day's work.

I would especially like to thank a few friends, without whom the past six years would have been very dull indeed. Tom Walker, currently pursuing his own Master's Degree, kept me entertained when I needed distractions, encouraged me when I needed support, and managed (more often than not) to reach base safely in the last inning so I could sneak in another at-bat. Alberto Gayon, also slated to receive his Master's Degree in the next few months, is about as great a friend and roommate as anyone could ask for – though I'll never quite understand his predilection for Japanese video games with huge robots. Lastly, Angel Solis remains to this day my closest friend and most trusted confidante. The past two years simply haven't been the same without the midnight runs to MacGregor - Gizzle? - but I am now pleased to have an extra degree to lord over him mercilessly.

Finally, I would like to thank my parents, Carlos and Gilda; my sister, Amy; y mis abuelas, Mima y Tatia; for their unconditional love and support. Their encouragement kept me motivated, and their praise gave my accomplishments meaning.

This thesis was prepared at The Charles Stark Draper Laboratory, Inc., under Contract No. 07605-003-069, sponsored by the National Aeronautics and Space Administration's

Institute for Advanced Concepts (NIAC). Publication of this thesis does not constitute approval by Draper or the sponsoring agency of the findings or conclusions contained herein. It is published for the exchange and stimulation of ideas.

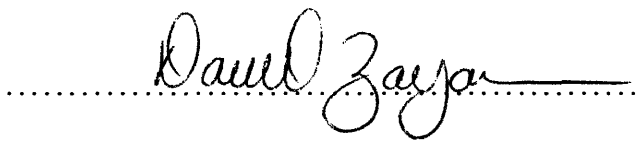

Daniel Antonio Zayas

Cambridge, MA

May 23, 2008

[This Page Intentionally Left Blank]

[This Page Intentionally Left Blank]

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Generally, the term "antimatter" refers to the various species of subatomic particles possessing equal mass and opposite color or charge relative to their "ordinary" counterparts. The antiproton, for instance, possesses the same rest mass as the proton, yet carries negative electric charge. This definition becomes less satisfactory when extended to neutral particles but it will suffice for the purposes of this introduction to simply note that, just as quarks may join together to form neutral particles (e.g., neutrons), so too may antiquarks combine to form antiparticles (e.g., antineutrons). The existence of particle-antiparticle pairs is a direct mathematical consequence of the theory of special relativity, and was first predicted mathematically by Dirac in 1928[1, 2]. All elementary particles can be associated with a corresponding antiparticle, with only a few exceptions (i.e., the photon and certain species of neutrino which are, in fact, their own antiparticle).

Although there is no known mechanism by which nature might favor matter over antimatter, or vice versa, the observable universe seems to be composed almost exclusively of traditional protons, neutrons and electrons. This fact remains a great mystery and a topic of popular study among physicists. Despite their apparent scarcity, however, antiparticle populations have been measured in both laboratory and natural environments, respectively, as byproducts of high-energy collisions and constituents of the cosmic ray flux.

When a particle-antiparticle pair come into contact, in appropriate quantum states, they annihilate to release all of their rest mass as energy. In the simplest case of a low-energy collision involving one electron and one positron, two gamma rays are produced such that

$$e^- + e^+ \rightarrow \gamma + \gamma \tag{1.1}$$

and the combined energy of the resulting gamma rays (provided the annihilation occurs at rest, as it often does) is precisely the rest energy of the electron-positron pair: $E = 2m_e c^2$. More generally, higher mass-energy collisions can release enough energy to drive the production of more exotic, heavy particles – in the case of proton-antiproton annihilation, high-energy pions are produced in relative abundance. These processes are two to three (2-3) orders of magnitude more mass-efficient than nuclear reactions and ten (10) orders of magnitude more mass-efficient than typical chemical reactions, and thus have enormous potential in a variety of applications.

Recently, medical applications have been envisioned wherein antimatter beams might be used to target and annihilate cancerous cells with minimal impact on surrounding healthy tissue. The Antiproton Cell Experiment (ACE) at the European Organization for Nuclear Research (CERN) near Geneva, Switzerland, demonstrated a four-fold increase in the destructive power of targeted antiproton beams relative to standard proton-beam therapy[3]. Applications involving improved-resolution imaging of solids, the analysis of annihilation products as a means of spectroscopy and materials analysis, and even cosmology and dark matter physics have also been envisioned. Perhaps the most technically challenging potential application is the use of antimatter as a fuel for spacecraft propulsion. The energy densities ($E/m \sim c^2$) associated with matter-antimatter annihilation have the potential to revolutionize our ability to explore our solar system and beyond.

In all cases, however, viability remains limited by the difficulties associated with producing and storing large quantities of antimatter. It is useful, then, to examine these difficulties in turn, their consequences for the development of large-scale antiparticle production, and how they might be mitigated in both laboratory and natural environments.

1.1 Laboratory Antimatter Production

The primary mechanism by which antiprotons are produced in a laboratory environment is the collision of high-energy proton beams with a solid (typically copper or nickel) target. The resulting debris, a small fraction (roughly ten antiprotons per million protons incident upon the target) of which is composed of antiprotons, may then be collected and separated using electromagnetic fields into its constituent particles. At present, only two facilities meet the requirements for sustained antimatter production (sufficiently high-energy particle

accelerators to support production, and dedicated electromagnetic extraction systems to separate and contain the collision products): CERN and the Fermi National Accelerator Laboratory (Fermilab or FNAL) in Batavia, IL, US.

Once the antiprotons are separated from their collision coproducts, they are decelerated and placed in confinement rings for use in subsequent experiments. The rates of production, however, are prohibitively low; at Fermilab, approximately 10^7 antiprotons per pulse can be confined for short periods, with production rates of approximately 10^{11} antiprotons per hour[4]. If the FNAL facilities were used to exclusively create and accumulate antiprotons a total of almost one nanogram (1 ng) would be produced over the course of a year, though a number of improvements have been discussed which could enhance production by a factor of ten (10) or more by increasing the efficiency of the collection process. A dedicated facility with an estimated cost of $5-$17 billion USD has also been proposed, which could increase production rates still further[5].

Schmidt et al.[4] discuss some of the fundamental energy cost constraints of producing antiprotons in particle accelerators. Due to the energy requirements for accelerating the proton beam, even with a wall plug efficiency of 50%, $0.10 USD per kilowatt-hour yields a net antiproton production cost of roughly $62.5 trillion USD per gram collected due to electricity costs alone. Other estimates have placed this figure as high as $160 trillion USD per gram[6].

A number of alternative production techniques have been proposed. Hora[7] suggests using high intensity lasers to produce antimatter. However, efficiently generating laser pulses with sufficient energy remains an obstacle. Chapline[8] proposes using heavy ion beams instead of proton beams to increase accelerator production, though the antiprotons are still generated isotropically making it difficult to collect the antiprotons from the ejecta debris. Cassenti[9] suggests redirecting pions generated during collisions, though this approach also suffers from the difficulty of containing and redirecting the debris. LaPointe[6] proposes using the Casimir force to suppress local vacuum fields as a means of generating the steep gradients required for proton-antiproton pair production at a potential boundary. The concept calls for holding two conducting metal plates nearly a meter on a side to within one nanometer (1 nm) of each other. This approach, too, has yet to be demonstrated in practice and remains outside the realm of manufacturability, at least in the near term. However, the basic physics of production can be validated with positrons, which the author suggests

could be accomplished with half-meter plates placed about 100 nanometers apart. Though this would be challenging, the production of positrons would validate the relevant physics.

1.2 Natural Antimatter Production

In general, pair production is possible whenever sufficiently energetic collisions occur. Under natural conditions in our solar system, only galactic cosmic rays (GCR) possess the high energies (on the order of a few GeV) required to create proton-antiproton pairs. There is, however, no shortage of target bodies capable of providing sufficiently large collision cross-sections.

The primary mechanism for the production of antiprotons in the Earth's magnetosphere, for example, is collision processes between the incident GCR flux and air molecules in the upper atmosphere. These collisions generate (among other species) free antineutrons, a fraction of which travel back into space before decaying into an antiproton, a positron and a neutrino. Those particles which possess a non-zero electric charge then become trapped in the Earth's magnetic field, and undergo periodic motion as described by the Lorentz force. That is, trapped charged particles in the Earth's magnetic field will spiral along their local magnetic field lines, while mirroring between the Northern and Southern hemispheres and drifting slowly around the planet. A fraction of the trapped particles, specifically those with pitch angles lying within the loss cone, are lost to annihilation as they penetrate the atmosphere near the magnetic poles, but the remaining population becomes stably magnetized and trapped in the Earth's magnetic field.

As particles are lost to diffusion or annihilation, new particles are generated to maintain a steady-state supply. Bickford et al.[10] developed a comprehensive phase-space model for the Earth's antiproton belt, and evaluated all other major objects in the Solar System for their ability to generate and sustain antiproton populations. Saturn's magnetosphere was found to be a particularly favorable environment for antiproton production, due to the relatively high local magnetic field strengths, and also to the favorable collision cross-sections produced by dust and debris in the planet's ring system. Figure 1-1, reproduced from [10], depicts the antiproton flux in Saturn's magnetosphere, and highlights the absorptive effects of the planet's rings and moons.

It is important to note that while larger planets tend to produce and trap larger quan-

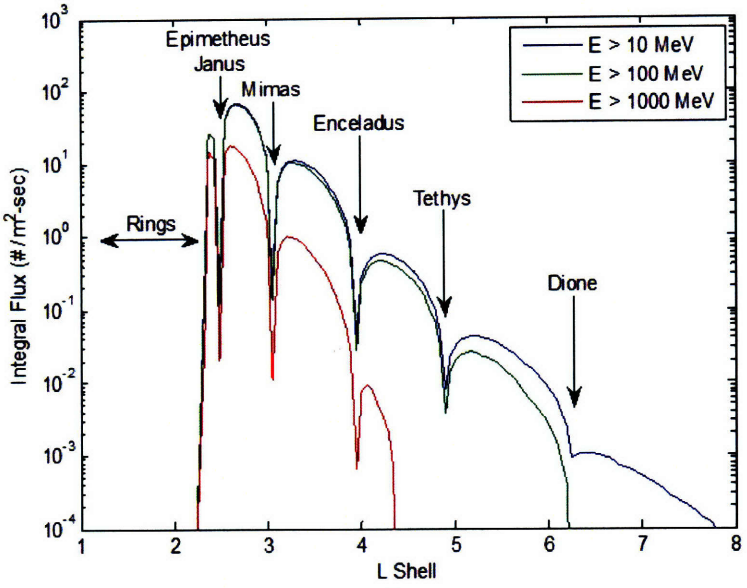

Figure 1-1: Saturn Integral Antiproton Flux

tities of antiparticles, the larger volumes result in reduced equatorial fluxes. In fact, the highest antiproton fluxes are conveniently found in the Earth's radiation belts, where radial transport processes are comparatively slow and result in increased residence times. Preliminary estimates suggest a quasi-static supply of roughly 160 nanograms, with steady replenishment.

1.3 Antimatter Storage

Long-term storage of antimatter is limited by several factors. In general, antiparticles must be stored in high-vacuum environments and suspended in electromagnetic fields to avoid annihilation with either background gases or, if applicable, the container walls. The first such trap, capable of storing 10^{10} antiprotons (roughly ten (10) femtograms), was 100 cm tall by 30 cm across and weighed 55 kg fully loaded[11]. The ratio of trap mass to stored antiproton mass was approximately 10^9 kg/μg though it is unclear how, or even if, this ratio would scale to the nanogram to microgram class storage levels needed for space applications.

Currently, the most advanced portable trap is the High Performance Antiproton Trap (HiPAT). HiPAT can, in principle, store approximately 10^{12} antiprotons (roughly 1 pg) for

days or more at a time by maintaining the trap at a temperature of 4 Kelvin. The storage density could, in theory, be increased by forming electrically neutral anti-hydrogen atoms to address space charge and Brillouin trapping limits, though this has not been demonstrated at a relevant scale.

A variety of other condensed-matter concepts have been discussed including leveraging photonic band-gap structures, quantum reflection, paraelectricity and other techniques to improve antiproton trapping, though none of these concepts have proven mature enough for practical applications[12].

1.4 Antimatter Propulsion

A variety of conceptual spacecraft propulsion systems have been envisioned which make use of antimatter beams in some form or another to achieve relativistic speeds. For instance, simple "beam-core" systems use electromagnetic "nozzles" to direct the high-energy products of annihilation and produce thrust[13]. More advanced concepts call for the use of low-energy antiparticle beams to catalyze nuclear fission and/or fusion reactions and produce thrust either directly (say, via directed ejecta[14]) or indirectly (say, by momentum transfer via the absorption of shockwaves from the detonation of small pellets of fissile materials[15]). Nuclear Electric Propulsion (NEP) systems have also been envisioned, where the heat of annihilation is converted to electrical energy and used to power other onboard electrical and propulsion systems, akin to modern radioisotope thermal generators (RTGs)[14].

The versatility of antimatter propulsion systems is readily apparent. For interstellar missions, where the duration of the mission is the leading constraint, the high energy densities associated with antimatter propulsion systems allow for very high ΔVs – typical interstellar cruise velocities range from $0.25c$ to $0.5c$ – such that scientifically interesting targets outside our solar system might be reached and studied within reasonable time scales. By contrast, minute quantities of antimatter – as little as a few nanograms, as will be shown in later chapters – would be sufficient to vastly expand our exploratory capabilities within our solar system, which have historically been limited by the low payload mass fractions associated with chemical propellants.

A significant detractor for most matter-antimatter propulsion systems is that they are inherently acceleration-limited by virtue of their very high specific impulses (I_{sp}). It is, of

course, the capability of the propulsion system to accelerate the spacecraft which drives performance, and not merely the velocity attained; however, higher accelerations typically result in dramatic mass and power penalties so a balance must be struck. Frisbee suggests that for missions up to a maximum of 40 LY, accelerations on the order of $0.01g_0$ to $0.03g_0$ would be preferred[13].

There are, clearly, several difficulties with antimatter propulsion. A large production infrastructure must be developed and maintained, and new technologies will be required to efficiently and effectively produce, isolate, and store appreciable quantities of antimatter fuel - the last of which has yet to be demonstrated on any significant scale. Further, gamma radiation produced in most antimatter annihilation reactions would pose a significant risk for onboard systems, and exposure would undoubtedly be lethal for any human passengers without proper shielding; indeed, proton-antiproton annihilation is in many cases the preferred process for space propulsion applications, as the products consist predominantly of both neutral and charged pions (π^0, π^+, π^-). This is not the case for electron-positron annihilation, for which the products are composed exclusively of high-energy gamma rays. Despite these challenges, however, nuclear and antimatter propulsion systems are widely regarded as the only practical choices for future interstellar missions.

1.5 Thesis Overview

In the remainder of this thesis, a novel concept for the extraction and annihilation of cosmic ray and radiation belt antiprotons will be explored. Chapter 2 will review the relevant physics governing field-particle interactions, classify the various possibilities for collection mechanisms, and discuss the computational approach used to analyze collection efficiency. Chapter 3 will detail the results of various simulations, and establish a baseline collection system that is both mass- and cost-efficient. Chapter 4 will assess the propulsive system performance achievable with localized concentrations of directed antiprotons in two competing antimatter propulsion concepts; "beam-core" propulsion, which calls for the direction of annihilation products via electromagnetic nozzle, and "catalyzed-fission" propulsion, where both annihilation products and fragments from catalyzed nuclear fission reactions serve as the reaction mass. Finally, chapter 5 closes with a summary of major findings and conclusions on the feasibility of antiproton propulsion systems, including re-

marks on a conceptual design for an antiproton-catalyzed fission engine, and sample mission architectures. Recommendations for future research are provided for the intrepid reader.

Chapter 2

Collection System Modeling

Recent work[10] at the Charles Stark Draper Laboratory in Cambridge, MA, suggests that localized, naturally occurring populations of antimatter can be exploited for a variety of applications including, among others, space propulsion. However, at best these populations occur in concentrations that are orders of magnitude below the storage densities required for operational systems. If we are to take advantage of natural sources of antimatter, we must first develop a means to concentrate these tenuous populations and trap them for extended periods. This chapter will introduce some of the proposed mechanisms for concentrating or focusing incident particle fluxes, as well as the analytical and computational models used to study them.

2.1 Field-Particle Interactions

2.1.1 The Magnetic Dipole Field

For the purposes of this thesis, we will restrict our analysis primarily to the motion of particles in dipole magnetic fields. Such fields are easy to model both mathematically and computationally, and closely approximate the magnetospheres of the Earth and Jovian planets. Magnetic dipoles arise from a closed circulation of electric current. The simplest case is a circular loop of wire having radius r and carrying current I, as depicted in Figure 2-1. Dipole fields can be characterized mathematically by a vector quantity known as the dipole moment, which in the above example would point through the wire loop according to the right-hand rule, and have a magnitude given by the product of the current and the loop area.

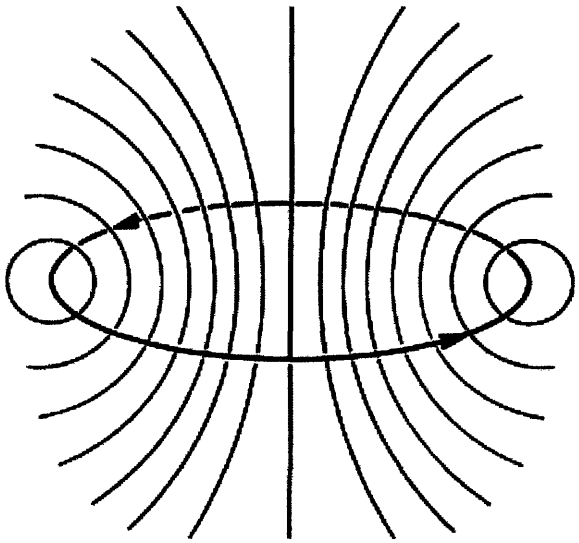

Figure 2-1: Magnetic Field Lines of a Circular Current Loop

The definition of the magnetic dipole moment is mathematically useful in that it allows for a basis of comparison between magnetic fields of differing strength and physical scale. A useful relation involving the magnetic dipole moment gives the strength B of the magnetic field at a point $(\vec{r}, \lambda)$ as

$$\vec{B}(\vec{r}, \lambda) = \frac{\mu_0}{4\pi} \frac{\vec{M}}{r^3} \sqrt{1 + 3sin^2\lambda}, \tag{2.1}$$

where λ is the magnetic latitude in radians, $\vec{M}$ is the magnetic dipole moment in ampere-square-meters, and μ_0 is the permeability of free space in Henrys per meter. Another useful relation defines the torque induced by an applied external magnetic field $\vec{B}$ on a magnetic dipole field. Specifically, the torque is given by

$$\vec{\tau} = \vec{M} \times \vec{B}. \tag{2.2}$$

2.1.1.1 Störmer Forbidden Regions

One particular aspect of the magnetic dipole field merits special attention, namely the existence of so-called "forbidden regions" which are inaccessible to particles below a specified energy. Störmer[16] showed that for magnetic dipole fields, there exist potential barriers which are impenetrable to incident particles defined by the relation

$$r = \sqrt{\frac{||\vec{M}||q}{pc}} \frac{cos^2\lambda}{1 + \sqrt{1 + cos^3\lambda}}, \tag{2.3}$$

where r is the radial distance from the center of the magnetic dipole field, $||\vec{M}||$ is the magnitude of the magnetic dipole moment, λ is the magnetic latitude, q and p are, respectively, the electric charge and momentum of the incident particle, and c is the speed of light in vacuum.

2.1.2 The Lorentz Force

The motion of any charged particle through an electromagnetic field is governed by the Lorentz force law,

$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B}). \tag{2.4}$$

A positively charged particle exhibits acceleration in the direction of an applied electric field, and curves perpendicularly around an applied magnetic field according to the right-hand rule.

In many cases, it is simpler to express a particle's trajectory as a superposition of a relatively fast circular motion around a particular point referred to as the "guiding center" and a (typically) slower drift of same. In the simplest case of a uniform static magnetic field, absent any other forces or fields, the Lorentz force is perpendicular to the magnetic field and the particle trajectory, resulting in helical trajectories at a constant speed. The gyro- or cyclotron frequency and the gyro- or Larmor radius are, respectively,

$$\omega_c = \frac{qB}{m}; \quad r_L = \frac{v_\perp}{\omega_c}. \tag{2.5}$$

The effects of additional fields and forces can be quantified by a general force drift,

$$\vec{v}_f = \frac{1}{q} \frac{\vec{F} \times \vec{B}}{B^2}. \tag{2.6}$$

The simplest cases are electric and gravitational fields, for which drift velocities are easily and intuitively derived. For other drifts, such as the so-called ∇B (pronounced "grad-

B") drift, the responsible force is more complex; say, the force on a magnetic dipole in a gradient. And in many cases, the particle's own acceleration gives rise to "fictitious" forces that account for specific drifts such as the "curvature" and "polarization" drifts. Despite these quirks, an accurate accounting of all forces and their respective drifts results in a complete description of a particle's motion.

2.1.3 Plasmas

If we wanted simply to influence the motion of a single particle in a vacuum, we could proceed with only the above equations of motion and design any number of relatively simple systems. In practice, most applications will require operation amidst some background population – a planetary atmosphere, a comet tail, or even the solar wind. Our framework is incomplete unless we also include certain bulk effects[17].

2.1.3.1 Debye Shielding

The screening of electric fields, a phenomenon commonly referred to as Debye shielding, places a fundamental limit on the influence of electric fields in plasmas. The presence of free charge carriers allows the plasma to shield local charge imbalances over distances on the order of a few Debye lengths,

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{n e^2}}, \tag{2.7}$$

where ϵ_0 is the permittivity of free space, k_B is the Stefan-Boltzmann constant, T and n are, respectively, the plasma temperature and density, and e is the fundamental electric charge. This places an upper bound on system size.

2.1.3.2 The Plasma Parameter

To supplement the idea of Debye shielding we may define a parameter Λ, referred to as the plasma parameter, given by

$$\Lambda = \frac{4}{3}\pi n {\lambda_D}^3, \tag{2.8}$$

where n and λ_D are as defined above. The plasma parameter may be interpreted physically as the number of particles within a given "Debye sphere" – that is, a sphere of radius λ_D.

An "ideal plasma" is one for which the relation $\Lambda \gg 1$ holds, such that collisions between plasma particles may be neglected, and the plasma bulk may be modeled as interacting uniformly with any external forces.

2.1.3.3 Plasma Waves

The propagation of electromagnetic waves in plasmas differs from that in a traditional gas, in that the presence of two distinct, charged populations results in a deviation from the vacuum dispersion relation $\omega = ck$. The simplest example is the class of waves referred to as "plasma oscillations." Consider a population of electrons displaced slightly from their equilibrium position within the plasma. Due to charge imbalance, these electrons are subject to a local electric field and begin to accelerate toward their equilibrium position. This relaxation invariably overshoots, however, and gives rise to a simple oscillation with frequency

$$\omega_p = \sqrt{\frac{ne^2}{\epsilon_0 m_e}}, \tag{2.9}$$

where n, e and ϵ_0 are as defined above, and m_e is the mass of the electron. We refer to this frequency as the "electron plasma frequency," and indeed a similar – though much lower, as $m_i \gg m_e$ – frequency exists for ions.

To see how the propagation of electromagnetic waves through a plasma is affected by local conditions within said plasma, consider the case of the "ordinary wave," a mode in which an oscillating electric field ($\vec{E} \parallel \vec{B}_0$) propagates in a direction perpendicular to the local magnetic field (that is, $\vec{k} \perp \vec{B}_0$). The dispersion relation for the ordinary wave is given by[17]

$$\omega^2 = {\omega_p}^2 + c^2k^2, \tag{2.10}$$

where c is the speed of light in vacuum, and k is the wavenumber. We may express the square of the phase velocity of the wave $v_\phi = \omega/k$ as

$${v_\phi}^2 = \frac{\omega^2}{k^2} = \frac{c^2}{N^2}, \tag{2.11}$$

or

$$N^2 = \frac{c^2 k^2}{\omega^2}. \tag{2.12}$$

The quantity N defined above is analogous to an index of refraction, such that propagation occurs only if $N^2 > 0$. In the case of the ordinary wave, this condition is satisfied by

$$N^2 = \frac{c^2 k^2}{(\omega_p{}^2 + c^2 k^2)^2} > 0, \tag{2.13}$$

which in turn implies that $k \in \Re$ or $\omega > \omega_p$. Indeed, for most wave modes pertinent to this study, the propagation of electromagnetic waves in a plasma occurs only at frequencies greater than the plasma frequency.

2.2 Antimatter Collection System Classification

Having described the various mechanisms by which electromagnetic fields and charged particles interact, let us now turn our attention to a few notable concepts for practical antimatter collection. For comparison's sake, we will consider both electrostatic and magnetostatic systems.

2.2.1 Electrostatic Systems

The ability of a plasma to screen electric fields places a fundamental limit on the efficiency of systems that rely solely on electric potential gradients to focus incident particles. The electron densities and temperatures under solar wind ($n \sim 10^3$ m^{-3} - 10^{11} m^{-3}, $kT_e \sim 1$ eV - 10 eV) or magnetospheric ($n \sim 10^6$ m^{-3} - 10^{11} m^{-3}, $kT_e \sim 1$ eV - 1 keV) conditions result in Debye lengths between 1 m and 100 m[18], such that any electrostatic fields generated by a collector would be capable of only localized penetration into the background plasma – far short of the length scales required to collect a significant flux. Further, increasing plasma densities resulting from the concentration of incident particles would amplify the shielding effect up to a critical density. Despite these challenges, however, several concepts merit consideration.

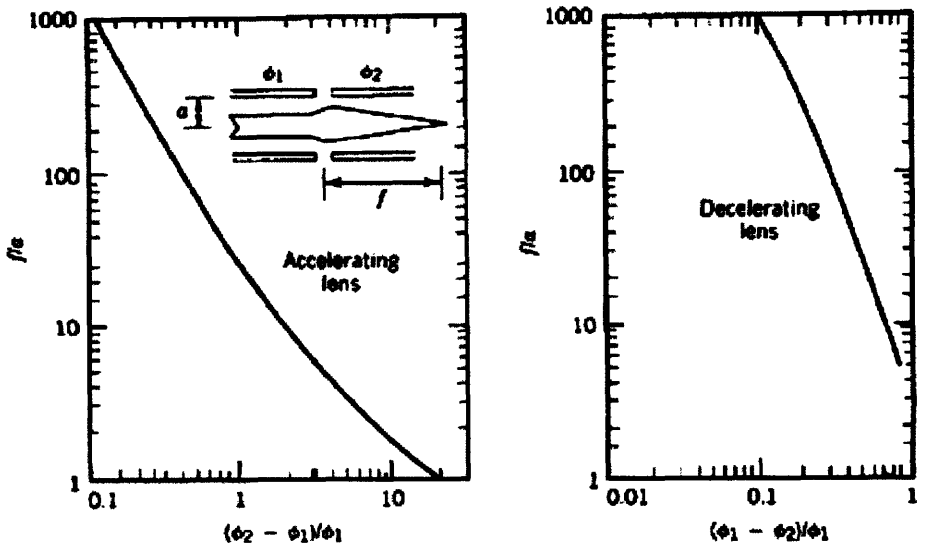

Figure 2-2: Focal Length Relations for Accelerating and Decelerating Einzel Lenses

2.2.1.1 Electrostatic Lensing

One such concept draws on the principles of ion and electron optics. Consider a series of discrete, circular structures, sharing a common axis of symmetry and biased to independent voltages. This configuration is known as an Einzel lens, and has been shown to focus incident beams of charged particles with well-defined focal lengths governed by stage separation and collection area. Figure 2-2, reproduced from [19], depicts the focal length relations for standard accelerating (left) and decelerating (right) Einzel lenses.

This "step-ladder" concept employing discrete, graduated electric fields is unfortunately hindered by its unidirectional nature. That is, the system has a preferred axis of acceptance leading to severe collimation problems when subjected to an incident flux of particles exhibiting primarily helical motion (as is the case for charged particles trapped in a planetary magnetosphere). The performance of electrostatic lens collection systems is discussed in more detail in Section 3.1.1.

2.2.1.2 The Jackson Sphere

In a recent study, Jackson[20] suggests the use of charged concentric spherical grids to focus $\sim$ 1 GeV antiprotons from the GCR flux. The proposed system consists of an outer grid biased at 10 MeV to reflect solar wind protons and positrons; an intermediate grid

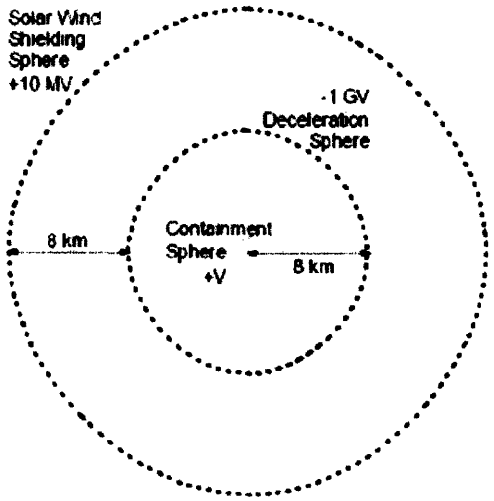

Figure 2-3: The Jackson Sphere

biased at 1 GeV to decelerate the incident antiprotons; and finally an inner trapping grid biased to some sufficiently positive voltage for long-term storage. A schematic is shown in Figure 2-3. The performance of the Jackson Sphere collection system is discussed in more detail in Section 3.1.2.

2.2.2 Magnetostatic Systems

Collection systems that rely solely on the use of static magnetic fields to concentrate an incoming flux of particles are essentially "drift-limited," in that their performance is limited by the drift velocities that they induce. Consider the case of a magnetic field created by a current-carrying loop, as depicted in Figure 2-1. As a particle traces its helical orbit along a field line it will experience a stronger magnetic field as it approaches the center of the loop, resulting in a decreasing gyroradius, until the particle loses all momentum in the direction of its travel and reverses direction; the ability of a particle to penetrate near to the center of the system is limited, as described in Section 2.1.1.1, by its kinetic energy, and its orientation (pitch angle) with respect to the field line it follows.

2.2.2.1 Current Loop Systems

The effect of the ∇B drift on low-energy particles can be a significant detriment to performance, as most particles are deflected away from the desired focal point to surrounding regions, where the magnetic field strength is comparatively low. The simplest (though least-

effective) magnetostatic collection systems, then, are composed of a single loop of current, and performance is governed by the physical size of the system and the magnitude of the operating current. Further, in light of the propensity of current loop systems to deflect any incident flux, variations on the single-loop collector scheme are also worth closer study – in particular, circular (as shown in Figure 2-4) and concentric arrangements of current loops arranged so as to exploit the deflection of particles to preferred regions, not necessarily centered on a given loop. Sections 3.2.1 and 3.2.2 describe in more detail the performance of current loop collection systems.

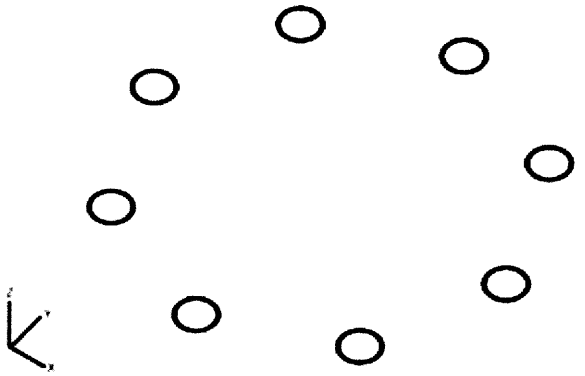

Figure 2-4: Three-Quarter View of an N-Loop System, for $N = 8$

2.2.2.2 The Plasma Magnet

Slough[21] has suggested the possibility of establishing large-scale magnetic fields with minimal infrastructure by using a rotating magnetic field (RMF) to magnetize electrons in a background plasma. The "plasma magnet" concept consists of a relatively small ($10 - 100$ m) two- or four-loop (2 or 4) antenna (see Figure 2-5) operated within a neutral plasma. By supplying current to opposing loops or loop pairs with a relative phase difference, a rotating magnetic field (RMF) is established near the antenna. The RMF is operated at a sufficiently high frequency $\omega \gg \omega_{ci} = eB/m_i$ so as to elicit a response in only the electron population; the motion of the ion population may be ignored, due to its relatively high inertia. The resulting current due to the electron motion sustains a large-scale magnetic field (hereinafter referred to as the LSMF) within the plasma, perpendicular to the plane of rotation of the electron population. Figure 2-6 depicts an artist's conception of the RMF antenna assembly.

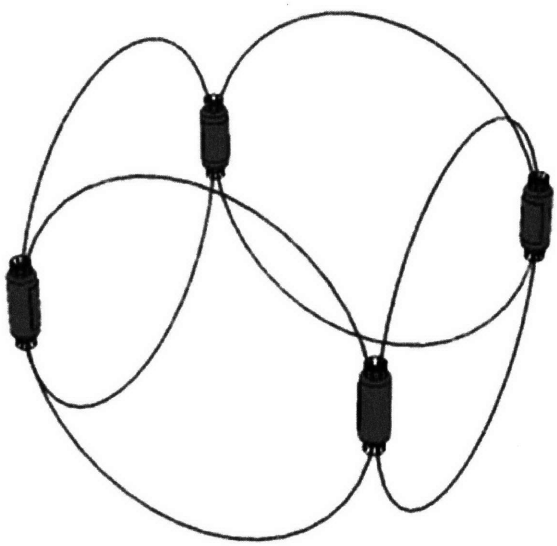

Figure 2-5: Artist's Rendition of an RMF Antenna Assembly

While the plasma magnet concept was originally developed for use as a novel form of "solar sail," the idea being to use the pressure forces supplied by an incident solar wind flux against a magnetized plasma to impart momentum upon the antenna assembly and any attached payload, the LSMF exhibits variable power-law behavior in radial decay as a function of the bulk plasma density and can be graduated to reduce gradients and mitigate the effects of collection-inhibiting drifts; requiring significantly less power and exerting influence over large volumes, the LSMF is an excellent candidate for antimatter collection as is shown in Section 3.2.3.

2.3 Computational Modeling

In practice, any conceivable collection system will generally require three modes of operation: (i) startup, where the creation of large-scale electromagnetic disturbances will require correspondingly large quantities of energy. The power system must be capable of driving the system to steady-state operation within a practical time scale, likely one year at worst; (ii) focusing, during which overall system performance is quantified by the rate at which antiparticles are collected. The ability of the system to concentrate and localize incident particle fluxes represents the fundamental limit of performance; and finally (iii) trapping, where some mechanism must be included to directly influence the localized antiparticle concentrations and ultimately transition them into long-term storage. We define

a "trapped" particle as one whose trajectory is bounded within some reasonable distance from the device center, for a reasonably long duration. All matter of electrostatic and electromagnetic traps are viable options.

For the remainder of this chapter, we will focus primarily on analytical and computational techniques used in the simulation of particle focusing. In particular, we wish to study the motion of a charged particle – in this case, an antiproton – in the presence of electromagnetic fields. These fields may be uniform or non-uniform, static or dynamic, so long as they are well defined in both space and time. Given a set of initial coordinates in phase space $(x0, y0, z0, u0, v0, w0)$, the path of the particle through the simulation volume must be determined. For an adequate statistical population, we apply large-scale Monte Carlo analyses over a range of physical parameters spanning the design space.

2.3.1 Collector Field Modeling

In order to accurately simulate particle trajectories, we must first develop high-fidelity models of any electromagnetic fields, ambient or artificial, present in the system. In most cases, such as the equatorial regions of planetary magnetospheres and interplanetary space, the ambient magnetic fields may be treated as uniform, and we will neglect ambient electric fields. Thus, it remains only to calculate the electromagnetic fields generated by the collector.

2.3.1.1 The Magnetic Field of a Loop of Current

In the case of a single current loop with current I and radius R, the approach is relatively straightforward. By making use of elliptic integrals, we may obtain a closed-form solution for the magnetic field at an arbitrary point[22]. In particular, we have that,

$$B_\xi = B_0 \left(E(k) \frac{1 - \alpha^2 - \beta^2}{Q - 4\alpha} + K(k) \right), \tag{2.14}$$

$$B_\rho = B_0 \frac{\gamma}{\pi\sqrt{Q}} \left(E(k) \frac{1 + \alpha^2 + \beta^2}{Q - 4\alpha} - K(k) \right), \tag{2.15}$$

where B_ξ and B_ρ are the axial and radial components of the magnetic field, respectively; $B_0 = \mu_0 I/2R$ is the magnetic field intensity at the center of the loop; $\alpha = \rho/R$, $\beta = \xi/R$, $\gamma = \xi/\rho$, $Q = (1+\alpha)^2 + \beta^2$ are non-dimensional spatial coordinates; and $K(k)$, $E(k)$ are

the complete elliptic integrals of the first and second kind, respectively. Their argument, $k = \sqrt{4\alpha/Q}$, is sometimes defined alternately as $m = k^2$, or $\alpha = sin^{-1}k$. Note that for generality we have introduced a fixed cylindrical coordinate system (ρ, ξ, θ) centered on the axis of symmetry.

2.3.1.2 The Plasma Magnet Field

For plasma magnet configurations the process is slightly more involved. The principle driver of performance is the RMF antenna's ability to influence the electron population; the more electrons driven by the RMF, the stronger the resultant LSMF. We can quantify this by establishing a magnetization criterion, and determining a penetration depth for the RMF into the bulk plasma that will depend on (i) the antenna coil radius; (ii) the antenna operating current; and (iii) the plasma temperature.

Consider an electron in motion at a point A near the equatorial plane of a dipole magnetic field. Following the methodology of Pugacheva[23], we impose the criterion $r_L \leq \frac{1}{10}\rho_{B,A}$ for magnetization; that is, we require that the electron gyroradius be at most one-tenth (1/10) the radius of curvature of the magnetic field line passing through the point A, evaluated at A. For a magnetic dipole field, the equatorial radius of curvature is given by $\rho_B = 3r$, where r is the radial distance from the center of the dipole. Thus we arrive at the condition for electron magnetization,

$$r_L = \frac{m_e v_\perp}{qB} \leq \frac{3}{10} r \tag{2.16}$$

or, equivalently,

$$B_{min}(r) = \frac{10 m_e v_\perp}{3qr} \tag{2.17}$$

where $B_{min}(r)$ is the minimum magnetic field strength required to magnetize an electron with velocity $v_\perp$ perpendicular to the local magnetic field at a distance r from the center of the dipole, along the equatorial plane.

Assuming a geometry for the RMF antenna, we may calculate explicitly the time-varying RMF. As described in the previous section, we model the RMF antenna as consisting of four current loops (i) oriented around a common center and perpendicular to each other in space; and (ii) driven by oscillating currents separated by 90° in phase. In this configuration, the

plane containing the centers of each of the four antennae is also the plane of rotation for the resulting RMF. Figure 2-6 depicts the RMF antenna configuration; the magnetic field lines are drawn in red, and the direction of rotation is labeled.

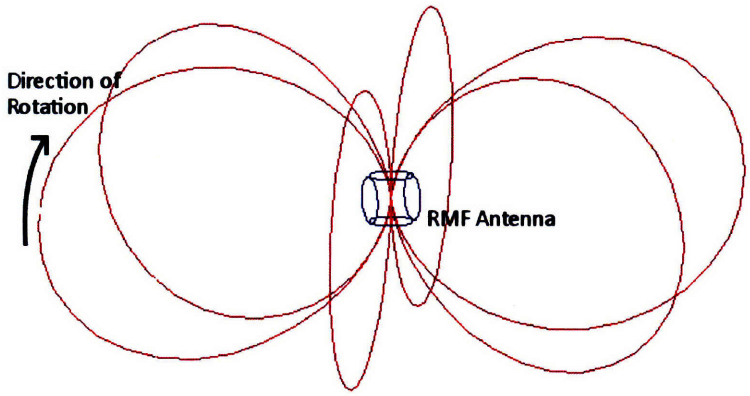

Figure 2-6: RMF Configuration for Plasma Magnet Systems

Next, we assume bulk characteristics for a background, neutral plasma, and use the aforementioned trapping criterion to determine the extent to which the rotating magnetic field penetrates the plasma. In particular, we expect to find a characteristic radius r_{max} beyond which the RMF is too weak to fully magnetize free electrons; that is, $B_{RMF}(r > r_{max}) < B_{min}(r)$. Figure 2-7 depicts the penetration of typical RMF antennas into a 10 eV background plasma.

Given a density profile for the background plasma, we may calculate a current density induced by the RMF and integrate to yield the overall plasma magnet field. In particular, recall from (2.14) and (2.15) that the magnetic field is the sum of a superposition of differential current loop elements, so that

$$\vec{B} = \iint B_0 \Gamma(\vec{r'}). \tag{2.18}$$

That is to say, each differential current loop contributes a central field intensity, which is then scaled by a geometric term $\Gamma(\vec{r'})$ (the parenthetical terms in 2.14 and 2.15) where $\vec{r'}$ denotes the relative separation of the differential current element from the point of calculation.

From the definition of B_0, and expanding, we have that

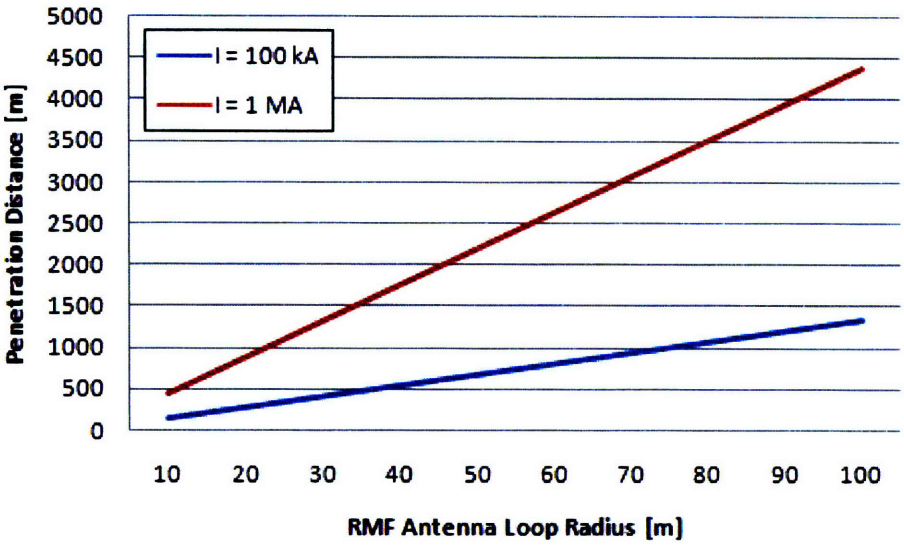

Figure 2-7: RMF Penetration for a 10 eV Plasma

$$\vec{B} = \frac{\mu_0}{2} \iint n(r) q \omega \Gamma(\vec{r'}) dA, \tag{2.19}$$

but we must take care to properly define the limits of integration. Recall that the RMF is only strong enough to magnetize electrons up to a radial distance r_{max} from the center of the antenna. Since magnetized electrons are confined to particular field lines, we include only those field lines that lie entirely within a distance r_{max} from the center of the antenna. Approximating again the field as a dipole, we thus discretize the space bounded by the minimally trapped field line into concentric cylindrical current shells, each with width dr and height

$$h(r) = r_{max}^{1/3} r^{2/3} \sqrt{1 - \frac{r^{2/3}}{r_{max}^{2/3}}}, \tag{2.20}$$

such that we arrive at a closed-form expression for the magnetic field intensity,

$$\vec{B} = 2\pi\mu_0 \int_0^{r_{max}} \int_0^{h(r)} n(r) q \omega h(r) \Gamma(\vec{r'}) dz dr. \tag{2.21}$$

2.3.2 Plasma Modeling

Having described the various assumptions and constraints pertinent to the design and simulation of the plasma magnet concept, let us take a few moments to describe the model used to simulate the bulk plasma. To first order, we assume a plasma parameter $\Lambda \gg 1$ and apply a variant of the two-fluid model in simulations pertaining to interactions between the RMF and the bulk plasma; one species of singly-charged ions (protons, as in the case of the solar wind) and a corresponding population of electrons. Neutrality is assumed, and viscous effects are neglected.

For sufficiently high ($\omega \gg \omega_{ci}$) RMF frequencies, the response of the ion population is negligible, and we may effectively reduce our model to an ideal gas of electrons moving against a non-interacting background of protons so as to satisfy neutrality. The inclusion of viscous effects alters the characteristics of the LSMF, however these effects are qualified in the next chapter.

2.3.3 Flux Input Models

As described in previous sections, there are two principal sources of antiprotons in our solar system. The first is the relatively dense population of antiprotons trapped in planetary magnetic fields, for which we can readily calculate differential flux models. In general, these antiprotons exhibit highly predictable bounce and precession motions as they respond to the planetary fields.

For a device operating in an equatorial orbit, we can reasonably model the input flux as a uniform shower of particles whose guiding centers travel with velocities parallel to the symmetry axis of the device. The distributions of energies and pitch angles are governed primarily by the production processes that replenish the population, as well as atmospheric interaction; the former shapes the distribution in energy, while the latter results in an upper limit (or "loss cone") on equatorial pitch angles. Around Earth, for example, antiprotons trapped in the radiation belts typically exhibit energies in the $10 - 250$ MeV range, with a loss cone of approximately $20°$. Figure 2-8, reproduced from [10], depicts the differential flux of antiprotons in the Earth's magnetosphere bound to field lines crossing the Earth's magnetic equator at a distance $L = 1.38$ from the center of the Earth, where $L = R/R_{Earth}$.

The relatively dense populations supported by planetary fields yield encouraging collec-

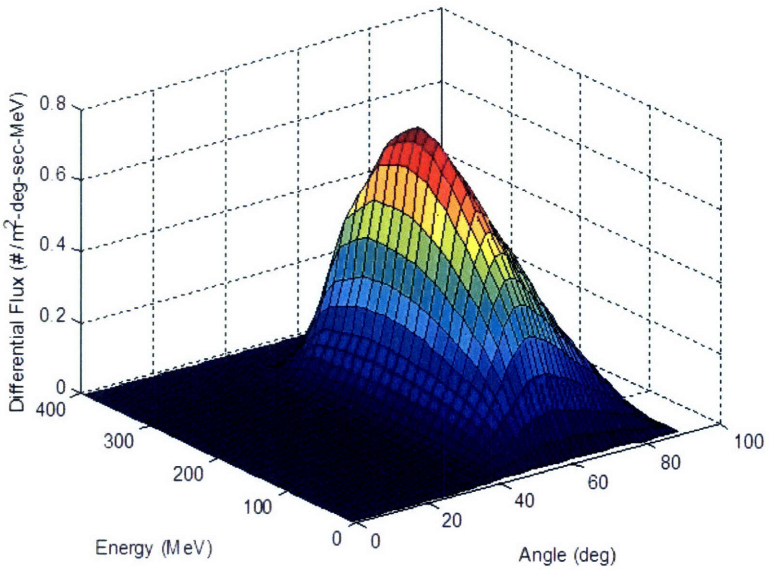

Figure 2-8: Differential Antiproton Flux for Earth, $L = 1.38$

tion rates, simply by virtue of a naturally larger incident flux. They are, however, fundamentally limited by loss processes; the steady-state supply of antiprotons trapped in the Earths radiation belts is estimated at roughly 160 ng[10]. Any significant collection from planetary sources would quickly exhaust the available supply, and the estimated timescales for replenishment from pair production processes are typically much longer ($\sim$ 2 ng/yr) than any timescales involved in collection. In that respect, the Jovian planets and GCR flux are promising alternatives.

As discussed in the previous chapter, Saturn's ring system – composed primarily of chunks of ice, rocks, and dust – yields highly favorable cross-sections for the production of antiparticles due to collisions with the incident GCR flux. Further, antiparticles generated directly and via the decay of antineutrons generated in the ring system need not be backscattered before becoming stably trapped in Saturn's magnetosphere. This gives rise to fantastic production rates and steady-state supplies of antiprotons around Saturn; as much as 240 μg/yr are generated in Saturn's magnetosphere, with a steady-state supply of roughly 10 μg! However, while the supply is vastly superior to the quantities available in the Earth's magnetosphere, the differential flux – and, by extension, the achievable collection rates – are significantly lower, due to the proportionally larger volume over which the antiproton population is trapped.

Collection directly from the GCR flux is also not without its difficulties. First, GCR

antiprotons exhibit energies much greater than those found in planetary magnetic fields; the peak in the GCR spectrum occurs at approximately 2 GeV, with non-negligible differential fluxes approaching as high as 10 GeV. This makes GCR antiprotons considerably more difficult to focus and trap, for they spend far less time under the influence of any electromagnetic fields. Further, while practically infinite in supply, the GCR antiproton flux is exceedingly tenuous. Despite these shortcomings, however, it remains true that the GCR flux represents an essentially inexhaustible supply, with no preferential direction or pitch angle to limit collection efficiency.

2.3.4 Path Integration

As described above, the motion of a charged particle through an electromagnetic field is governed by the Lorentz force law. We can express the particles trajectory using a system of six first order differential equations for position and relativistic momentum,

$$\frac{d\vec{s}}{dt} = \vec{V}, \tag{2.22}$$

$$\frac{d\vec{p}}{dt} = q(\vec{E} + \vec{V} \times \vec{B}), \tag{2.23}$$

which can be readily solved numerically. To that end, a simulation environment was developed in MATLAB using the well-known *ODE45* package. *ODE45* is an adaptive time-scale, fourth- and fifth-order pair Runge-Kutta-Fehlberg integration algorithm, and provides high accuracy for reasonable computation cost with minimal additional development. Other integration methods may be adapted to obtain similar results.

Each particle is assigned an initial position and velocity on the simulation boundary, and traced as it passes through the simulation space under the influence of electromagnetic forces induced by the collectors generated fields. To minimize computational cost, we consider a particles trajectory "complete" once it either (i) strikes the surface of the device; or (ii) attains a distance greater than 1% in excess of its initial distance from the device center (i.e., the particle passes outside the simulation boundary). Figure 2-9 shows an example trajectory for a GCR antiproton incident on the Earth's magnetic environment.

Rather than simulating single particles and applying the relevant Monte Carlo analyses "ex post facto," a novel approach was developed to vectorize the integration process. Ex-

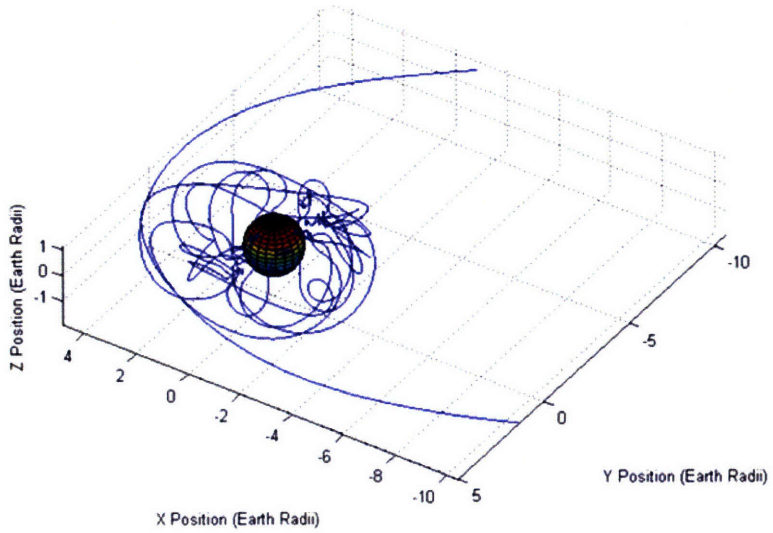

Figure 2-9: Example GCR Antiproton Trajectory, $E = 2$ GeV

panding the original system of six first-order differential equations into a matrix of six-by-N differential equations, we effectively solve all N particle trajectories simultaneously, eliminating a significant degree of overhead computation and reducing computational cost by as much as 15- to 100-fold, allowing for the simulation of a significant $(10^5 - 10^6)$ number of particle trajectories over the course of a few hours on a reasonably fast desktop computer (2.0 GHz, dual-core CPU, 1 GB RAM).

2.3.5 Quantifying Performance

2.3.5.1 Antiproton Collection

In order to quantify performance, we must establish a suitable metric to characterize a given system's ability to focus incident fluxes. The simplest and most practical metric is simply local density; for a given trapping efficiency, the greater the number of particles present in a given region the more particles may be trapped by the system for long-term storage. Taking advantage of axial symmetry, we discretize the simulation space using cylindrical coordinates r, z, θ. The resultant "bins" take the shape of rings, with square cross-sections, centered on the axis of symmetry of the device.

For a given particle, mapping the trajectory onto such a discretized space is equivalent to projecting the particles trajectory onto the $r - z$ plane. Applying Monte Carlo analyses, we can readily sum over the contributions from all particles, thereby "counting" the total

number of particles passing through each bin. Together with the total flux incident on the system, and given a total number of simulated particles, the raw particle counts are converted to a local flow rate; that is, the number of particles (alternatively, the mass) passing through each bin per unit time. It is reasonable to assume that any successful trapping mechanism will have a relatively localized sphere of influence, centered on the focal point of the collector. A useful metric for performance is then the local flow rate within some sufficiently practical distance from the center of the device.

Chapter 3

Collection System Performance

Having developed analytical field models for simple one- and two-loop systems, as well as RMF-driven systems, we may quantitatively assess the various antiproton collection concepts. In particular, we wish to estimate collection rates and efficiencies under various operating conditions (i.e., Earth's magnetosphere, the GCR flux). For the purposes of this thesis, we will focus primarily on magnetostatic systems, which at present seem to show more promise than their electrostatic counterparts. For the sake of comparison, however, top-level estimates of performance are included for the two electrostatic systems discussed in Chapter 2.

By calculating raw collection rates, we may intuitively match prototype systems to baseline missions – i.e., LEO, Jovian, Interstellar - requiring certain ΔV. It should be noted that, due to radical differences in scale and concept, it is not always reasonable to compare systems by their gross collection performance. In these cases alternative metrics are desired which incorporate such factors as mass, power, and/or cost, among others. In this chapter, the various collector concepts are analyzed. Where practical, we attempt to assess performance across a range of metrics; total collection rate (TCR), measured in units of mass per year; mass-specific collection rate (MSCR), measured in units of mass per year per unit device mass; and power-specific collection rate (PSCR), measured in units of mass per year per unit power.

For each of the various system concepts described below, we make several key assumptions. First, where applicable, the current state of the art in superconductor technology is assumed, limiting current densities to $1.7 \cdot 10^8$ A/m^2 and material densities to $9 \cdot 10^3$ kg/m^3[24]. Second, a period of one year is allotted for any start-up and/or transient

phenomena. Third, for all systems, launch costs are determined according to the present day Atlas V Programs incremental launch costs of approximately 8,800 USD/kg[25]. Finally, where relevant, power systems have been scaled according to current state-of-the-art production limits of approximately 40 W/kg (nuclear), though missions carried out within the confines of our solar system could potentially achieve 25 to 250 W/kg (photovoltaic)[26].

3.1 Electrostatic Systems

3.1.1 Electrostatic Lensing

As described in the previous chapter, a series of discrete, charged toroidal surfaces may be arranged in series and biased to graduated electric potentials to form an Einzel lens. These systems focus incident beams of charged particles to well-defined focal lengths. The principle drivers of performance are the collection aperture, and the potential distribution across successive stages.

The focal length of the system may be adjusted by biasing the various stages in the system so as to achieve characteristic changes in beam kinetic energy. For example, consider a two-stage configuration with an inlet radius of 100 km. Let ϕ_1 and ϕ_2 be the change in beam kinetic energy across the first and second stages, respectively. For $\phi_1 \sim -\phi_2$, the focal length is roughly 100 km downrange of the point lying between the two stages. Conservative estimates suggest collection rates ranging from 10^{-5} to 10^{-2} ng/yr are achievable, assuming the device can be kept dynamically stable.

In addition to the uninspiring collection rates, this concept suffers from several key flaws. First, operation within any substantial background population (i.e., the solar wind or a planetary magnetosphere) will render the system susceptible to shielding effects. Subsequent stages would require sufficient proximity so as to maintain the continuity in the electric potential distribution, however for decelerating lens configurations – where the successive potentials are of opposing polarity – this would introduce severe instabilities. Further, the unidirectionality of the system makes it a poor choice for harvesting from an omnidirectional source such as the GCR flux.

One attractive feature is the relatively low mass of the stages themselves; however, excessive background densities could result in prohibitively large electron currents to the surfaces of the stages, thus requiring prohibitively massive power systems to maintain the

desired potentials.

3.1.2 The Jackson Sphere

Operating in the GCR flux, the Jackson Sphere concept has the potential to harvest significant quantities of incident antiprotons, as well as confine them for extended periods, up to a space-charge-imposed limit depending on the volume of the collector. However, as with the electrostatic lens concept, the Jackson sphere is severely handicapped by several key factors.

First, the arrangement of the system is dynamically unstable and thus gives rise to significant control issues. Further, Jackson suggests positron cooling of antiprotons passing through the deceleration sphere; however, the mass of positrons required for complete elimination of momentum for antiprotons with residual energies of a few MeV to a few GeV is roughly 10^9 kg! Clearly, the problems of obtaining 10^9 kg of positrons for space-charge considerations and preventing electron/positron annihilation within the deceleration sphere represent enormous technical challenges, far surpassing those of the collection system itself.

3.2 Magnetostatic Systems

3.2.1 Single-Loop Systems

The performance of various single-loop collection systems has been previously assessed, however under faulty assumptions. Bickford's[27] approach is purely analytical, and assumes the conservation of the first adiabatic invariant for charged particles moving through an appropriate region of influence. Specifically, the relationship between the pitch angle (α), momentum (p) and magnetic field strength (B) for a particle in motion may be written as

$$\frac{p_1^2 \sin^2(\alpha_1)}{B_1} = \frac{p_2^2 \sin^2(\alpha_2)}{B_2} \tag{3.1}$$

for any two points (1) and (2) along the particle's trajectory.

While this methodology is valid for describing the motion of particles in relatively homogeneous magnetic fields, such as protons trapped in the Earth's radiation belts, it fails for regions in which the gradient of the field is large. For most reasonably sized single-loop systems, the high-field region is localized to volumes considerably smaller than a typical Lar-

mor orbit. Estimates based on this erroneous assumption yield collection rates approaching hundreds of micrograms (100 μg) per year; as will be shown below, subsequent simulations of explicit particle trajectories through concept collectors have tempered expectations.

For single-loop systems, we may optimize over two principal independent variables: (i) the radius of the superconducting coil; and (ii) the operating current. Figures 3-1 and 3-2 depict the TCR in nanograms per year for various single-loop configurations in both the Earth's Magnetospheric and GCR fluxes.

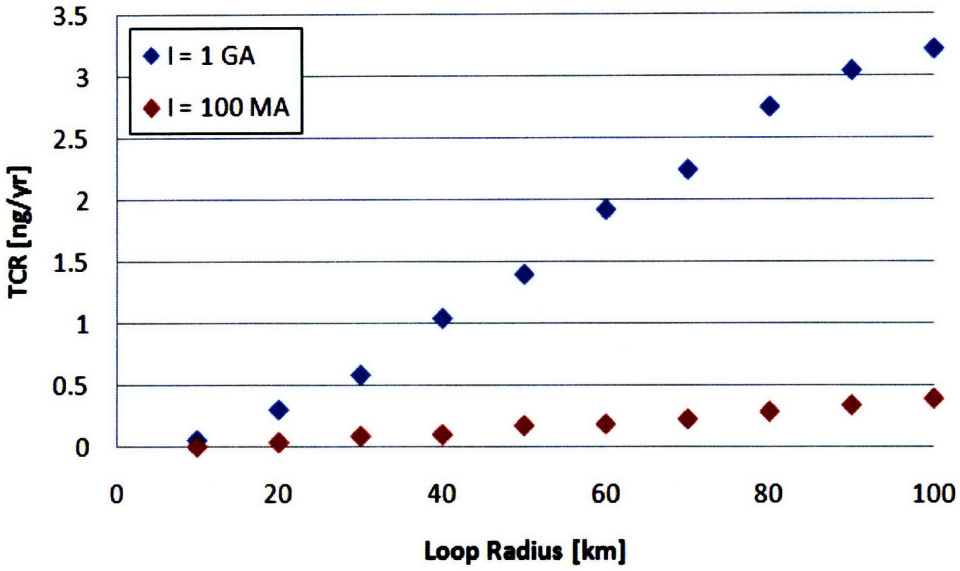

Figure 3-1: TCR for Single-Loop Systems in Earth's Magnetosphere

Immediately we observe that the total collection rate scales directly with the physical dimensions of the device. That is, for a given operating current, larger systems focus incident particles at a greater rate. Further, system performance is also found to scale with the operating current. Indeed, for single loop systems, maximizing TCR is as simple as building the largest coil possible, and operating at the highest attainable current. For devices measuring hundreds of kilometers and operating at billions of amperes, collection rates rival those of earth-based production. In practice, however, these systems are beyond the realm of practicality. Current superconductor technology sets a prohibitive lower limit on system mass, and typical systems draw power at rates exceeding the GW range, as shown in Table 3.1.

In light of the staggering costs and power draws associated with these systems, there is significant motivation to optimize not over raw collection rates, but rather over mass- and

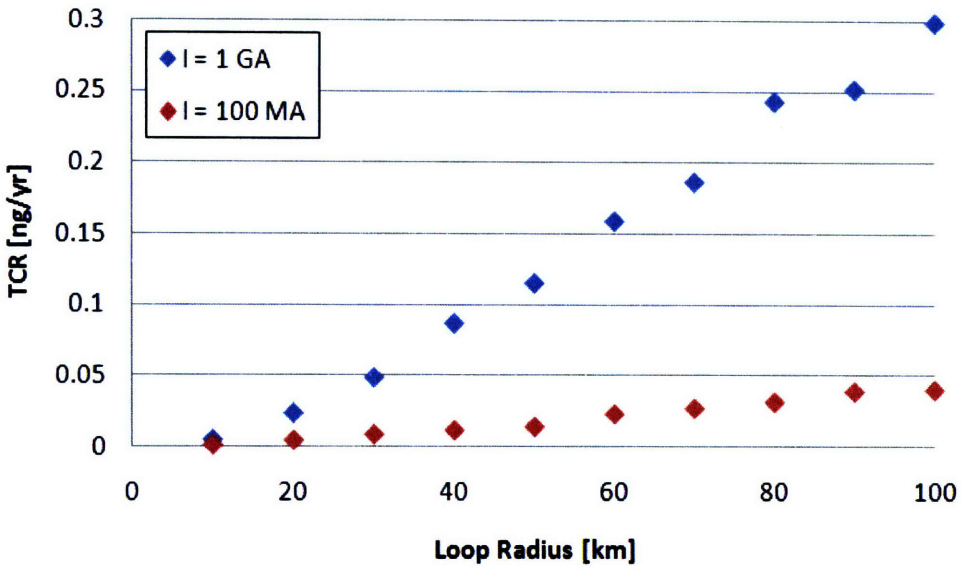

Figure 3-2: TCR for Single-Loop Systems in the GCR Flux

Radius \ Current	100 MA			1 GA		
	Mass [kg]	Power [MW]	Launch Cost [USD]	Mass [kg]	Power [MW]	Launch Cost [USD]
10 km	$3 \cdot 10^8$	28.2	$\$2.8 \cdot 10^{12}$	$3 \cdot 10^9$	$2.8 \cdot 10^3$	$\$2.8 \cdot 10^{13}$
40 km	$1 \cdot 10^9$	124	$\$1.1 \cdot 10^{13}$	$3 \cdot 10^{10}$	$1.24 \cdot 10^4$	$\$1.1 \cdot 10^{14}$
70 km	$2 \cdot 10^9$	225	$\$2.0 \cdot 10^{13}$	$3 \cdot 10^{10}$	$2.25 \cdot 10^4$	$\$2.0 \cdot 10^{14}$
100 km	$3 \cdot 10^9$	328	$\$2.8 \cdot 10^{13}$	$3 \cdot 10^{10}$	$3.3 \cdot 10^4$	$\$2.8 \cdot 10^{14}$

Table 3.1: Mass, Power and Cost Figures for Various Single-Loop Systems

power-specific collection rates as shown in Figures 3-3 through 3-6.

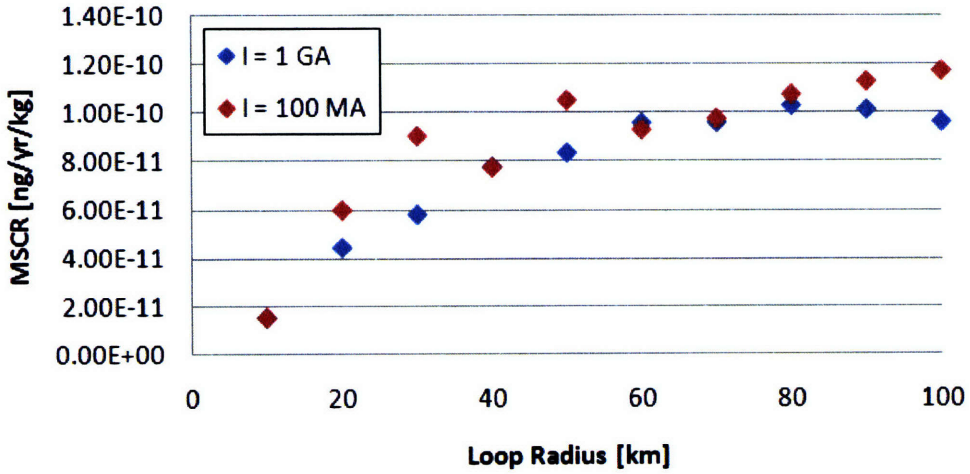

Figure 3-3: MSCR for Single-Loop Systems in Earth's Magnetosphere

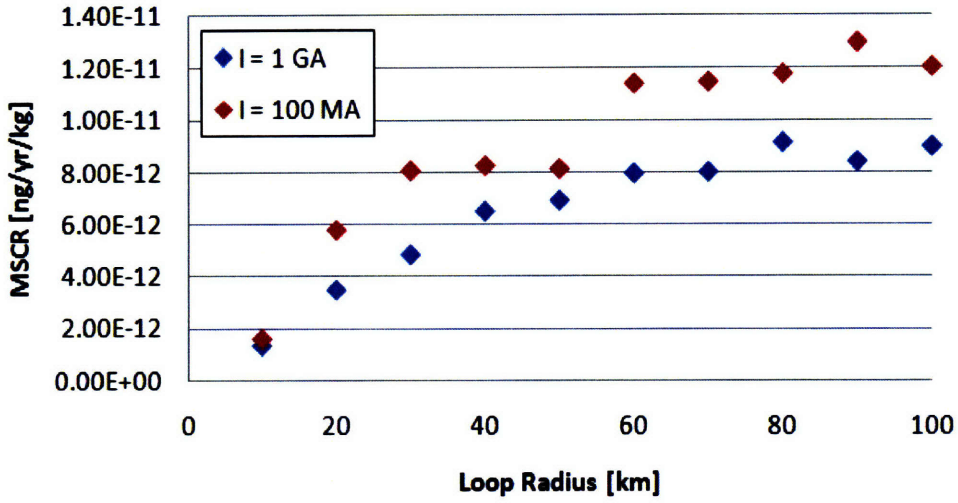

Figure 3-4: MSCR for Single-Loop Systems in the GCR Flux

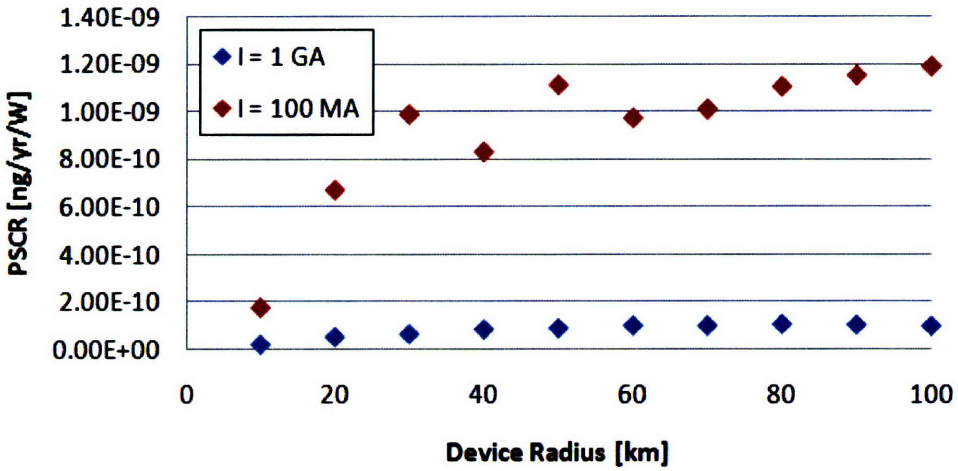

Figure 3-5: PSCR for Single-Loop Systems in Earth's Magnetosphere

Based on the above data, we find that lower-current systems are vastly more efficient from both a mass and power perspective. Armed with this knowledge, and given practical constraints such as a maximum launch mass or maximum power consumption, we can make objective comparisons between various points in the design space. In various single-loop simulations, strong trends emerge to suggest that optimal mass- and power-efficiency is achieved by maximizing the coil radius and minimizing operating current. However, because the limiting case of an infinitely large device operating at zero current is physically meaningless, we must bound the design space by specifying a minimum TCR. For single

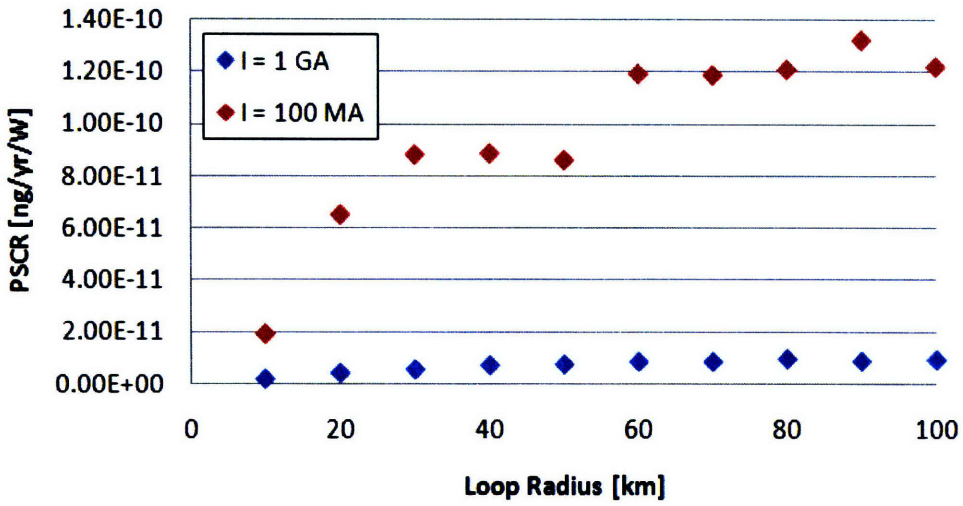

Figure 3-6: PSCR for Single-Loop Systems in the GCR Flux

loop systems, operating at currents below 100 MA fails to produce practical collection rates, and such systems are thus not considered. Having established a minimum current, and thus bounded the design space, we may proceed with objective comparisons between systems using only mass, power, and TCR as performance metrics; a 100 km, 100 MA collector, for example, will outperform a 10 km, 1 GA collector of equal mass with an 88% decrease in power consumption (see Table 3.1).

As described above, typical single-loop systems generate magnetic fields that are simply ill-suited to efficient collection. The vast majority of the incident flux is deflected away from the device, so collection must occur at either very low rates, or over impractically large volumes. The use of multiple single-loop systems in tandem, as described in the previous chapter, allows us to exploit this effect to improve performance. However, simulations suggest only a marginal increase (approximately a five-fold improvement in collection rate) in performance for N-loop systems, for $N = 10$. While efficiency improves with increasing N, so too does total power and mass, thus rendering N-loop systems equally impractical. The limiting case of $N = \infty$, however, is of particular interest.

3.2.2 Concentric-Loop Systems

Concentric-loop systems represent the mathematical and practical limit of an N-loop configuration with $N = \infty$. Figure 3-7 illustrates the superposition of current in each loop (shown in red), resulting in two net currents (shown in blue), each of which gives rise to its

own magnetic field. The opposing polarities of these fields, due to the opposing directions of current flow, give rise to localized null-points in the magnetic field (regions where the net magnetic field is negligible) which are conducive to improved focusing. These null points may be placed anywhere along the axis of symmetry by varying the respective magnetic dipole moments of the two current loops. The advantages of concentric-loop systems include a reduction in overall field gradients, and the localization of high-field regions nearer the coils, away from the center of the device.

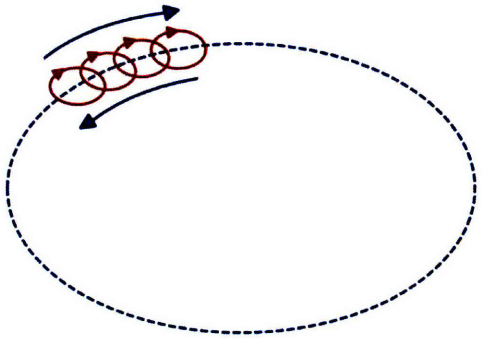

Figure 3-7: Superposition of Current for N-Loop Systems in the Limit $N \rightarrow \infty$

Relative to single-loop systems, the addition of a second coil introduces two new variables which relate directly to performance: (i) the radius of the inner coil; and (ii) the operating current of the inner coil. For simplicity, we replace these two variables with a set of non-dimensional counterparts; namely, (i) the ratio of inner loop radius R_i to outer loop radius R_o; and (ii) the ratio of inner loop current I_i to outer loop current I_o. The former is naturally limited to values between 0 and 1, while the latter is practically limited by the fact that a single power system must operate both current loops. As a conservative limit, we bound the ratio of currents to within one decade of unity, i.e. $0.1 \leq I_i/I_o \leq 10$. The first step in optimization is to eliminate one of these new variables. Specifically, we seek to determine a favorable operating point in either R_i/R_o or I_i/I_o. As shown in Figures 3-8 through 3-10, overall performance is largely insensitive to variations in the current ratio, except for the case when the ratio of loop radii approaches unity. In these cases, we observe a distinct linear trend in performance with increasing inner loop current. This is not altogether surprising – the limiting case of equal radii is equivalent to a single-loop system, and we have already seen that increasing the operating current leads to an increase in TCR for such systems. By operating at high ratios of loop radii and high inner loop currents, and

thus localizing the high-field regions near the coils and away from the focusing regions, we achieve significantly increased performance at minimal mass and power penalties relative to single-loop systems.

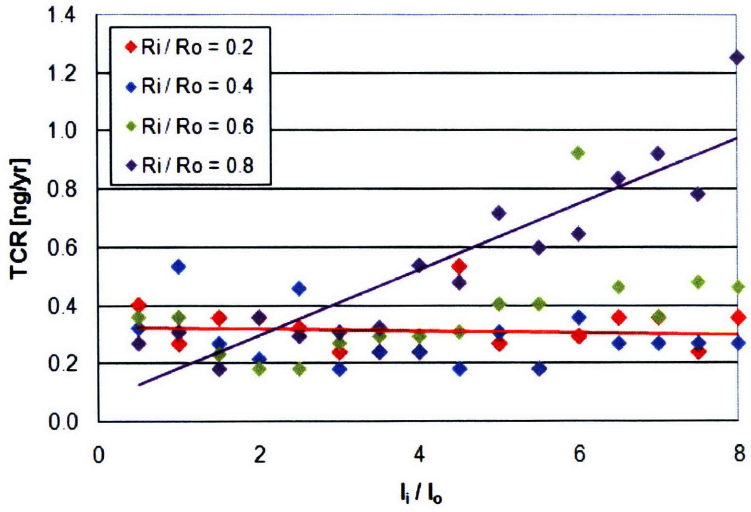

Figure 3-8: TCR for Concentric-Loop Systems in Earth's Magnetosphere, R_o = 10 km, I_o = 1 GA

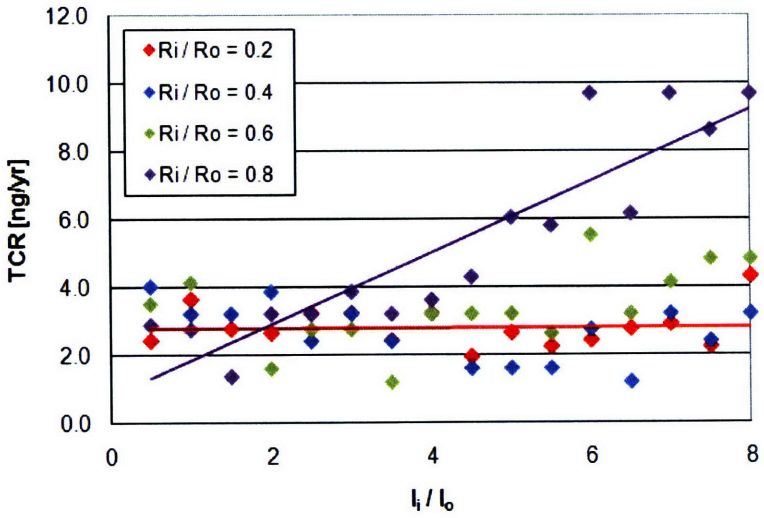

Figure 3-9: TCR for Concentric-Loop Systems in Earth's Magnetosphere, R_o = 30 km, I_o = 1 GA

Table 3.2 lists total collection rates for "practical" concentric-loop systems. The outer coil current is fixed at 100 MA, and the inner coil radius and current are 0.8 and 8 times, respectively, those of the outer coil. For large systems, collection rates approach tens of

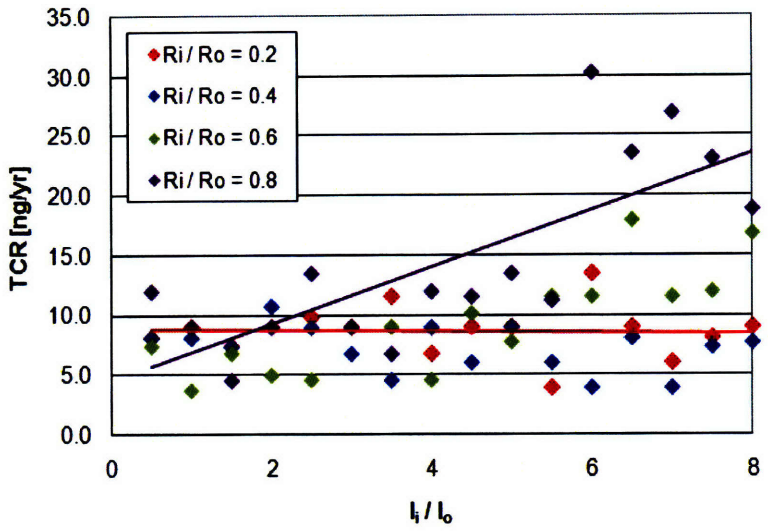

Figure 3-10: TCR for Concentric-Loop Systems in Earth's Magnetosphere, R_o = 50 km, I_o = 1 GA

Outer Coil Radius	$I_o = 100MA$, $R_i/R_o = 0.8$, $I_i/I_o = 8$			
	Mass [kg]	Power [GW]	Launch Cost [USD]	Collection Rate [ng/yr]
10 km	$3 \cdot 10^9$	3	$\$2.8 \cdot 10^{13}$	0.49
40 km	$1 \cdot 10^{10}$	10	$\$1.1 \cdot 10^{14}$	3.1
70 km	$2 \cdot 10^{10}$	20	$\$2.0 \cdot 10^{14}$	6.7
100 km	$3 \cdot 10^{10}$	30	$\$2.8 \cdot 10^{14}$	9.6

Table 3.2: Mass, Power and Cost Figures for Various Concentric-Loop Systems

nanograms per year (10 ng/yr) - a significant improvement over single- and N-loop systems.

3.2.3 Plasma Magnet Systems

So far we have seen encouraging collection rates from simple magnetostatic systems - despite enormous masses and power requirements, the best concentric-loop systems achieve collection rates exceeding current Earth-based production capabilities. However, as discussed in Section 2.2.2.2, there is just cause to expect considerably higher performance from the plasma magnet concept.

Recall from the previous chapter that RMF penetration (and, by extension, collection performance) was found to vary linearly with RMF antenna radius for a given operating current. Figures 3-11 and 3-12 depict TCRs for various plasma magnet configurations in the Earth Flux at 10^{18} m^{-3} and 10^{16} m^{-3} plasma densities, respectively, and an operating current of 100 kA. As expected, the "larger is better" trend observed in single-loop systems

remains prevalent.

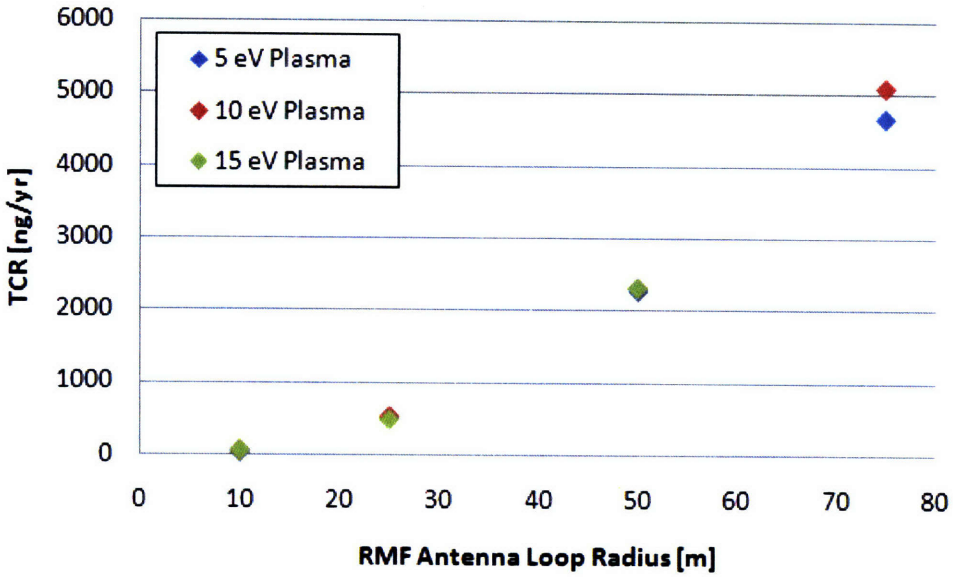

Figure 3-11: TCR for Plasma Magnet Systems with $n = 10^{18}$ m^{-3}, $I_{RMF} = 100$ kA

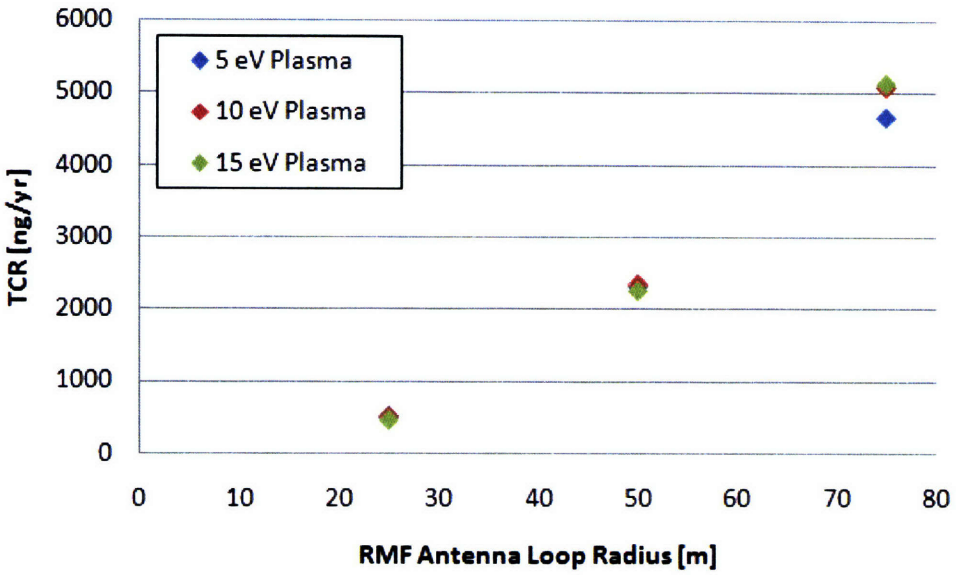

Figure 3-12: TCR for Plasma Magnet Systems with $n = 10^{16}$ m^{-3}, $I_{RMF} = 100$ kA

For most configurations studied, the intensities and volumes of influence of the resultant fields were comparable to those generated using single- or concentric-loop systems. However, the principal advantage of the plasma magnet lies in the use of a distributed, low-density plasma as the charge-carrying species; this leads to dramatically reduced gradients and

reduced system mass with a corresponding improvement in overall performance. Antiproton collection rates exceed micrograms per year for typical systems which are just 100 m in size and which draw just a few hundred kilowatts of power. Complete RMF systems could conceivably be inserted into their operational orbits with a single launch.

A surprising result is the relative insensitivity of performance to plasma density. Higher plasma densities should result in increased magnetic field intensity, and thus improved collection performance. However, the corresponding increase in the magnetic field gradient retards this effect, in much the same way it does for single- and concentric-loop systems.

In fact, closer analysis reveals favorable performance at lower plasma densities due to the reduction of dissipative losses. Assuming a uniform resistivity[17]

$$\eta = 5.2 \cdot 10^{-4} (kT_e)^{-\frac{3}{2}} \text{ohm·m} \tag{3.2}$$

throughout the magnetization region and integrating, we may express the power required to maintain the driven current as

$$P = \iiint_V \eta j_\theta^2 dV, \tag{3.3}$$

where $j_\theta(r) = ne\omega_{RMF}r$ is the induced azimuthal electron current density. We see then that the RMF power varies as the square of the plasma density, and inversely with plasma temperature. Figures 3-13 and 3-14 depict the mass-specific collection rates for various plasma magnet configurations in the Earth Flux at 10^{18} m^{-3} and 10^{16} m^{-3} plasma densities, respectively, and an operating current of 100 kA. By operating at relatively low plasma densities and high temperatures we reduce power consumption and thus system mass with virtually no penalty to performance.

While MSCR was found to scale favorably with decreasing plasma density and increasing plasma temperature, of interest is the reversal in behavior with respect to device size. At high densities the mass of the power system dominates and thus smaller devices are desired to mitigate the high power consumptions, despite the reduction in TCR. For lower densities, however, the mass of the power system becomes less significant, allowing for larger devices to maximize collection efficiency.

As with the single- and concentric-loop systems described in previous sections, we observe for plasma magnet systems similar monotonic trends in performance. Short of oper-

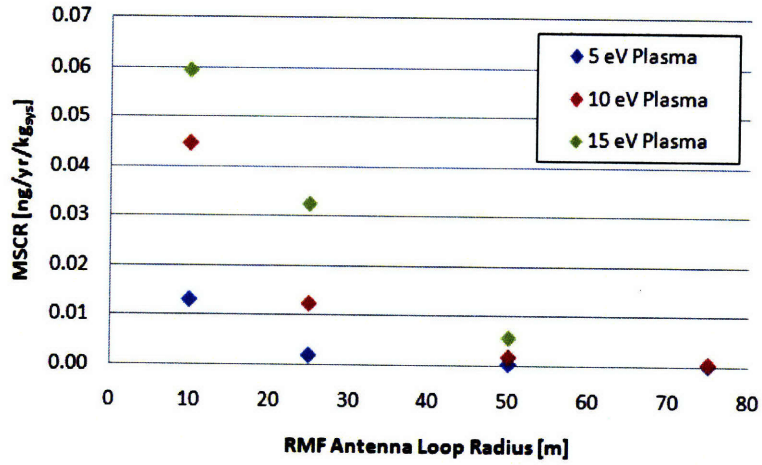

Figure 3-13: MSCR for Plasma Magnet Systems with $n = 10^{18}$ m^{-3}, $I_{RMF} = 100$ kA

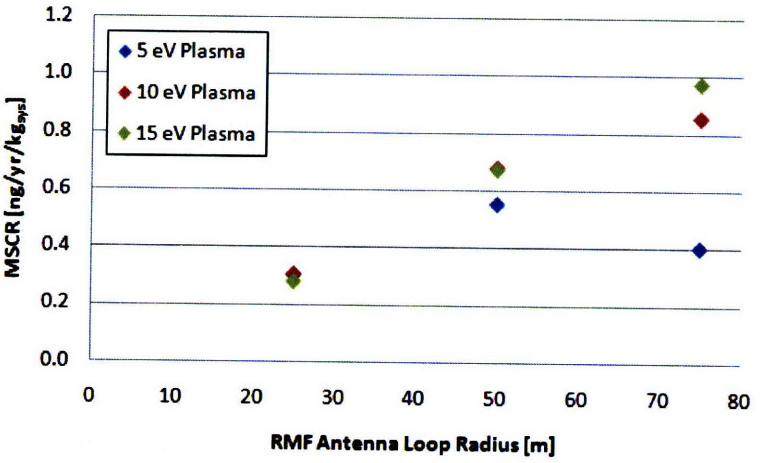

Figure 3-14: MSCR for Plasma Magnet Systems with $n = 10^{16}$ m^{-3}, $I_{RMF} = 100$ kA

ating in a zero-density, infinite-temperature plasma, we choose total capacity as a useful boundary condition. Specifically, we know that the confinement of likely charged particles in a magnetic field is limited[28] to a maximum achievable density of

$$n_B = \frac{B^2}{8\pi mc^2}, \tag{3.4}$$

above which the repulsive electrostatic and centrifugal forces become powerful enough to overcome the Lorentz force induced by the RMF. For antiprotons, plasma densities below 10^{16} m^{-3} are generally insufficient to drive plasma magnet fields powerful enough to trap meaningful quantities of antiprotons, so we choose as a lower bound a plasma density of

$n = 10^{16}$ m^{-3}. The following tables summarize the performance of various plasma magnet systems operating in the Earth flux, with varying plasma temperatures.

RMF Antenna Radius	Collection Rate [ng/yr]	Antenna Mass [kg]	RMF Power [kW]	Power System Mass [kg]	Total Mass [kg]
25	511.29	1,662.35	1.1	27.38	1,689.73
50	2,328.80	3,324.71	35	876.00	4,200.71
75	4,681.57	4,987.06	266	6,652.50	11,639.56

Table 3.3: Plasma Magnet Collector Performance for $n = 10^{16}$ m^{-3}, $kT_e = 5$ eV

RMF Antenna Radius	Collection Rate [ng/yr]	Antenna Mass [kg]	RMF Power [kW]	Power System Mass [kg]	Total Mass [kg]
25	508.45	1,662.35	0.158	3.94	1,666.29
50	2,345.50	3,324.71	5.04	126.03	3,450.73
75	5,085.84	4,987.06	38.3	957.00	5,944.06

Table 3.4: Plasma Magnet Collector Performance for $n = 10^{16}$ m^{-3}, $kT_e = 10$ eV

RMF Antenna Radius	Collection Rate [ng/yr]	Antenna Mass [kg]	RMF Power [kW]	Power System Mass [kg]	Total Mass [kg]
25	465.26	1,662.35	0.051	1.26	1,663.62
50	2,262.56	3,324.71	1.62	40.40	3,365.11
75	5,149.64	4,987.06	12.3	306.75	5,293.81

Table 3.5: Plasma Magnet Collector Performance for $n = 10^{16}$ m^{-3}, $kT_e = 15$ eV

3.3 Discussion

Having analyzed the various electrostatic and magnetostatic collection concepts, let us close the chapter with a few remarks on antiproton trapping and model limitations, as well as a summary of baseline collection system architectures.

3.3.1 Antiproton Trapping

In practice, we must assume that any incident flux should require a degradation in energy before it can be stably trapped. One means by which to accomplish this would be to simply place a slab of material near the focal point of the system. Particles passing through the slab would lose a fraction of their kinetic energy with each pass, ultimately achieving a suitably low momentum for capture. This mechanism is inefficient, however, in

that for reasonable momentum loss, the slab would need to be prohibitively thick and cover a large area, resulting in increased system mass and increased annihilation cross-sections. An alternative to the slab concept is to use polarized RF waves to slow incident antiprotons. This method, too, is inefficient; the propagation of RF waves would be shielded by the bulk plasma, thereby limiting the region of influence over incident particles. Table 3.6 highlights several possibilities for antiproton energy degradation as a means of diverting the incident flux onto closed magnetic field lines.

Concept	Description	Mass Penalty	Power Draw
Slab Interaction[1]	A slab of matter is placed in the path of incident antiprotons. Scattering processes during transit serve to degrade particle kinetic energy, by an amount proportional to the slab thickness and cross-section.	$\sim 10^7$ kg	N/A
Polarized RF Cooling[2]	Circularly polarized RF waves are used to selectively degrade particle kinetic energy, and promote the trapping of a particular species based on the electric charge (and thus polarization). In this manner, the trap may be biased to favor anti-protons over protons.	$\sim 10^2$ kg	$\sim 10^3$ W
Potential Barrier[3]	Charged bodies are used to exert a retarding force on incident particles by making them traverse a potential well. Here, again, the opposing polarities of protons and anti-protons can be used to bias the collection towards the desired species.	$\sim 10^3$ kg	$\sim$ 150 mW

Table 3.6: Summary of Antiproton Energy Degradation Mechanisms

Irrespective of the nature of the degradation mechanism used, the storage capacity of a given system will in principle be limited solely by (i) the volume of influence of the collection system; and (ii) the nature of the background particle population. The former condition is obvious, but the latter merits careful consideration.

For a neutral plasma, the Brillouin density limit – described in Section 3.2.3, but not specific to the plasma magnet configuration – places an upper bound on storage capacity, and is indeed the limit used in calculations of collector storage capacity. This limit can in theory be circumvented by artificially seeding the background population with oppo-

[1]Based on a collection aperture of radius 100 m, and a 4 cm uranium slab or an 18 cm aluminum slab for 250 MeV antiprotons.

[2]Based on an incident flux of approximately 8 /*mu*g/yr (plasma magnet collection rate) and 10% efficiency.

[3]Based on an incident flux of approximately 8 /*mu*g/yr (plasma magnet collection rate) and a transparency of 99%.

sitely charged species, thereby relaxing the space-charge limit, though this has yet to be demonstrated in simulations.

3.3.2 Comparative Performance

The plasma magnet concept is found to be the clear winner in direct comparison with all others, achieving collection rates in the microgram-per-year (μg/yr) range. As discussed in previous sections, the use of a distributed, low-density plasma for the expansion of the LSMF yields significantly shallower field gradients and thus suppresses the natural tendency of the magnetic field to shield against incident fluxes. Power consumption and total system mass are also significantly reduced, as the need for large, dense, superconducting coils is eliminated. By contrast, single- and concentric-loop systems operate at much lower collection rates, and draw considerably more power, yet still rival Earth-based production. Table 3.7 summarizes the relevant figures of merit for various baseline magnetostatic systems.

Parameter	Single-Loop	Concentric Loop	Plasma Magnet
Coil Radius	100 km	100 km	100 m
Operating Current	10^9 A	10^9 A	10^5 A
Plasma Density	N/A	N/A	$2 \cdot 10^{16}$ m^{-3}
Plasma Temperature	N/A	N/A	15 eV
Coil Mass	$3.3 \cdot 10^{10}$ kg	$6.6 \cdot 10^{10}$ kg	6600 kg
Total Power	32.8 GW	65.6 GW	200 kW
Power System Mass	$8.2 \cdot 10^8$ kg	$1.6 \cdot 10^9$ kg	5,200 kg
Total Mass	$3.4 \cdot 10^{10}$ kg	$6.8 \cdot 10^{10}$ kg	12,000 kg
Collection Rate	0.2 ng/yr	4 ng/yr	8.6 μg/yr

Table 3.7: Comparison of Baseline Collection Systems

Given the above figures for mass, power, and cost, it is readily apparent that a baseline plasma magnet collection system is well within reach given current or near-future technology. Such a system would revolutionize our ability to generate and sustain significant antimatter stores, as highlighted by Table 3.8.

Parameter	Earth-Based Colliders	Plasma Magnet Collector
Production Rate	2 ng/yr	25 ng/day
Storage Capacity	$\sim$ 4 pg	$\sim$ 110 ng
Storage Time	$> 10^5$ s	$> 10^7$ s
Production Cost	$\$4.8 \cdot 10^5$ USD/ng	N/A
Launch Cost	$\sim\$10^9$ USD/ng	$\sim\$9.0 \cdot 10^5$ USD/ng[1]

Table 3.8: Comparison of Baseline Plasma Magnet Collection System with Earth-Based Production Capabilities

3.3.3 Model Limitations

Recall from Section 2.3.2 that we have treated our plasma as an inviscid ideal gas of electrons responding to the RMF field against a backdrop of stationary protons. For the remainder of this chapter, let us discuss these assumptions, in particular the neglect of viscosity, and attempt to qualify their effects. For the plasma magnet baseline operating conditions (moderate plasma density, high plasma temperature) described in Section 3.3.2, we may readily calculate a plasma parameter $\Lambda \simeq 10^7 \gg 1$. Immediately, this confirms our expectations of collective behavior within the plasma.

The largest remaining source of potential error is the assumption of an inviscid, collisionless plasma. Recall from Section 2.3.1.2 that we have assumed complete magnetization of electrons within the plasma up to a distance r_{max}, the maximum penetration distance of the RMF into the bulk plasma. This assumption was used earlier in this chapter to derive a driven current density $j_\theta(r) = ne\omega_{RMF}r$, for $r < r_{max}$. The inclusion of collisional effects would modify the current profile, resulting in contributions to the driven LSMF at radii above r_{max}. Figure 3-15 depicts qualitatively the comparative shapes of the inviscid (red) and viscid (blue) current profiles.

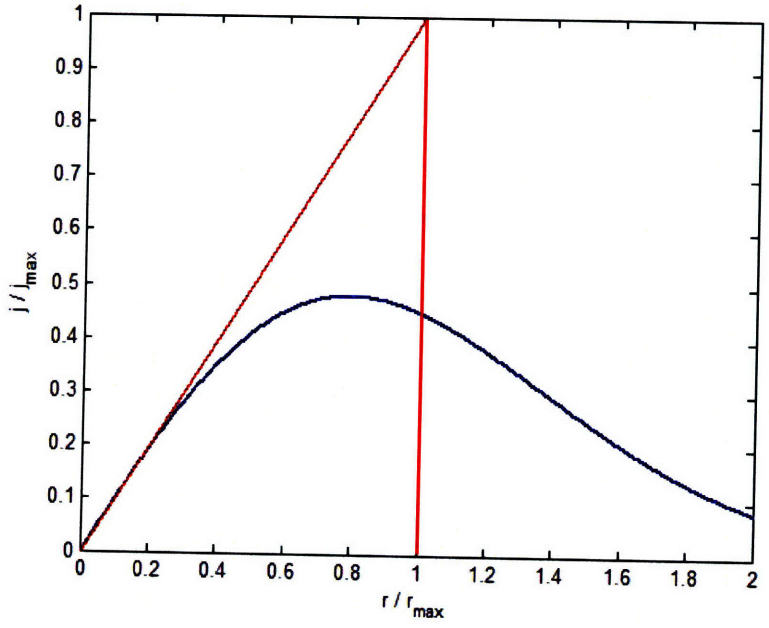

Figure 3-15: Qualitative Current Profiles for Inviscid (Red) and Viscid (Blue) Plasma Models

[1]This figure is based on a one-time collection of the stated (Brillouin-limited) storage capacity.

The inclusion of viscous effects, as shown above, should give rise to current densities beyond the inviscid RMF penetration distance, and thus a larger LSMF with shallower gradients and increased collection performance. Several attempts were made to model this phenomenon in MATLAB and COMSOL Multiphysics; however, no successful model was obtained to accurately describe the viscous effects on the behavior of the bulk plasma. A full, equilibrium MHD model of the RMF and its interactions with the bulk plasma would serve to validate the qualitative descriptions above.

The lack of collisional effects in the plasma model also results in an absence of diffusive losses. In actuality, any system will experience continual depletion as collisions within the plasma give rise to a mass flow out of the system. The rate at which a bulk plasma would require replenishment would, of course, depend on the plasma density and temperature. In the case of a fully ionized plasma, the flux associated with diffusion is proportional to[17]

$$\vec{\Gamma}_{\perp} \propto \frac{n}{B^2 \Sigma \sqrt{kT}} \nabla n. \tag{3.5}$$

With total bulk masses on the order of only a few hundred grams ($m_{bulk} \simeq 140.1$ g for the baseline plasma magnet collector operating point), however, replenishment is not likely to be a burden.

Chapter 4

Propulsive System Performance

If we are to leave the confines of our solar system and engage seriously in the exploration of nearby stars, it will be necessary to develop and implement radically new and vastly superior propulsion systems. Traditional chemical rockets are limited by their relatively low energy-density propellants; the Voyager spacecraft, currently traveling at just shy of 17 km/s, would require nearly 75,000 years to reach the closest star to our Sun, Proxima Centauri. While advanced electric and hybrid propulsion systems have advanced considerably in recent years, "fast" interstellar missions requiring $\Delta V \sim 0.5c$ remain outside the realm of practicality for modern propulsion systems.

Various electric and nuclear propulsion systems have been proposed and analyzed[14] for interstellar applications, ranging from advanced electric propulsion (EP) systems, to beamed energy systems such as solar and laser sails, to matter-antimatter propulsion systems. Of these concepts, only the third presents difficulties in analysis and prototyping, by virtue of the difficulty in producing and handling large quantities of antimatter and, consequently, has been developed only cursorily. In the following sections, we will refine estimates of propulsive performance for two distinct matter-antimatter propulsion concepts.

4.1 The Relativistic Rocket Equation

A useful performance metric in evaluating rocket propulsion systems is the dry mass fraction; that is, the ratio of "dry" mass (typically the payload and structural casings) to "wet" or total mass (i.e., the total mass of the rocket at launch). The use of this metric is convenient for several reasons, the most immediate of which are a relative ease

of calculation and a basic intuitiveness that applies generally to other propulsion concepts. Before we begin, let us first examine the fundamental equations which will govern our analysis. Consider a rocket of mass m traveling at a speed v and expelling propellant at an exit speed $u_e = g_0 I_{sp}$ relative to the body of the rocket. Conservation of momentum requires that the change in momentum of the rocket be equal to the change in momentum of the expelled propellant. Specifically,

$$\frac{d}{dt}(mv) = g_0 I_{sp} \frac{dm_p}{dt}, \tag{4.1}$$

where m_p is the mass of the expelled propellant. Expanding the derivative on the left, we have that

$$mdv + vdm = g_0 I_{sp} dm_p, \tag{4.2}$$

which we may simplify by noting that $dm = -dm_p$ by conservation of mass. In the reference frame of the moving rocket ($v = 0$), we thus have that

$$dv = -g_0 I_{sp} \frac{dm}{m}, \tag{4.3}$$

or

$$\Delta V = -g_0 I_{sp} \ln \frac{m}{m_0}. \tag{4.4}$$

Written in its classical form, the rocket equation relates explicitly a rocket's "wet" and "dry" masses to its propulsive performance. Specifically,

$$\frac{m}{m_0} = e^{-\frac{\Delta V}{g_0 I_{sp}}}, \tag{4.5}$$

where m is the mass of the rocket at any point in time, m_0 is the initial or "wet" mass, ΔV is the change in velocity of the rocket, and g_0 and I_{sp} are the acceleration due to gravity at the Earth's surface and the specific impulse, respectively. The specific impulse $I_{sp} = u_e/g_0$ of a rocket is given in seconds, and is a measure of the change in momentum imparted per unit of propellant; the higher the I_{sp}, the less fuel required to attain a given ΔV.

As with chemical rockets, antimatter propulsion systems are – to an extent – governed

by the rocket equation. However, if we wish to apply Equation 4.5 at relativistic velocities, we must rederive the rocket equation to account for relativistic effects. Recall that for an object moving with speed v,

$$m = \frac{m_0}{\sqrt{1 - v^2/c^2}}, \tag{4.6}$$

where m_0 is the rest mass of the object, and $c = 299,792,458$ m/s is the speed of light in vacuum. The sum w of two relativistic speeds u and v is given by

$$w = \frac{u + v}{1 + uv/c^2}. \tag{4.7}$$

Following the methodology of Frisbee and Leifer[14], consider a rocket of mass M expelling propellant opposite its direction of motion at a speed u_e relative to its own reference frame. In the center-of-mass frame, the resulting forward velocity of the rocket is v, and the reaction mass dm attains a velocity u.

Conservation of mass-energy yields

$$\frac{dM}{\sqrt{1 - v^2/c^2}}c^2 = -\frac{dm}{\sqrt{1 - u^2/c^2}}c^2, \tag{4.8}$$

while conservation of momentum yields

$$d\left(\frac{M}{\sqrt{1 - v^2/c^2}} \cdot v\right) = u \cdot d\left(\frac{m}{\sqrt{1 - u^2/c^2}}\right). \tag{4.9}$$

Applying the relativistic expression for additive velocities, we have that

$$u = \frac{u_e - v}{1 - u_e v/c^2}. \tag{4.10}$$

Finally, expanding the above derivatives and substituting for u, we have that

$$\frac{dM}{M} = -\frac{c^2}{u_e}\frac{dv}{c^2 - v^2}, \tag{4.11}$$

or, equivalently,

$$\ln M = -\frac{c}{2u_e}\ln\frac{c + v}{c - v}. \tag{4.12}$$

Here we have made the approximation $v \sim u_e$ so as to arrive at an analytic solution. Applying Equation 4.12 at the initial ($M = m_0$, $v = 0$) and final ($M = m$, $v = \Delta V$) boundary conditions, we arrive at the relativistic rocket equation,

$$\frac{m}{m_0} = \left(\frac{1 + \Delta V/c}{1 - \Delta V/c}\right)^{-\frac{c}{2g_0 I_{sp}}}. \tag{4.13}$$

Let us now turn our attention to the details of annihilation and reaction mass for the two concepts in question. We will find that for antimatter rockets, an additional correction must be made to the rocket equation to account for non-unity efficiency in the conversion of mass-energy at annihilation.

4.2 Beam-Core Propulsion

The simplest means of achieving thrust from controlled matter-antimatter annihilation is to direct, via magnetic nozzle, the electrically charged ejecta. This concept is typically referred to as "beam-core" antimatter propulsion, and was analyzed extensively by Frisbee[13]. The reactants of choice are proton-antiproton pairs which annihilate to produce pions, roughly two-thirds (2/3) of which are electrically charged and may, therefore, be directed magnetically. Because the pions possess a nonzero rest mass, not all of the proton-antiproton mass is converted into kinetic energy of products. In practice, the process is roughly 64% efficient – the rest mass of charged pions accounts for roughly 22% of the initial proton-antiproton mass, and the neutral pions account for another 14%. Table 4.1, reproduced from [13], summarizes the distribution of mass-energy and decay processes for proton-antiproton annihilation.

Because the conversion from proton-antiproton rest mass to charged pion rest mass is incomplete, we must further adjust the rocket equation to account for the "loss" of reaction mass. Again following the methodology of Frisbee and Leifer[14], we introduce the mass-energy conversion efficiency χ to denote the ratio of charged reaction mass to propellant mass; i.e., for proton-antiproton annihilation, $\chi = 0.223$. After considerable manipulation[1] Equation 4.11 then becomes

[1] A full derivation of the relativistic rocket equation for both unity and non-unity mass-energy conversion efficiency can be found in the Appendices of [14]

Species	Rest Mass [MeV]	Fraction [%]	Kinetic Energy [MeV]	Fraction [%]	Mass-Energy [MeV]	Fraction [%]
Initial Reactants						
p	938.3	49.97	0	0	938.3	49.97
e^-	0.5	0.03	0	0	0.5	0.03
$\bar{p}$	938.3	49.97	0	0	938.3	49.97
e^+	0.5	0.03	0	0	0.5	0.03
Initial Products						
$2.0\pi^\circ$	269.9	14.38	439.1	23.39	709.1	37.77
$1.5\pi^+$	209.4	11.15	374.3	19.94	583.7	31.09
$1.5\pi^-$	209.4	11.15	374.3	19.94	583.7	31.09
2γ			1.0	0.05	1.0	0.05
Decay Products						
$2.0\pi^\circ \rightarrow 4\gamma$			709.1	39.77	709.1	37.77
$1.5\pi^+ \rightarrow 1.5\mu^+$	158.5	8.44	288.5	15.36	446.9	23.80
$1.5\pi^+ \rightarrow 1.5\nu_m$			136.8	7.28	136.8	7.28
$1.5\pi^- \rightarrow 1.5\mu^-$	158.5	8.44	288.5	15.36	446.9	23.80
$1.5\pi^- \rightarrow 1.5\bar{\nu}_m$			136.8	7.28	136.8	7.28

Table 4.1: Mass-Energy Distribution of Matter-Antimatter Annihilation Products

$$\frac{dM}{M} = -\frac{c^2}{u_e}\frac{dv}{c^2\left(\chi + (1-\chi)\frac{v}{u_e}\right) - v^2} \tag{4.14}$$

or, equivalently,

$$\ln M = \frac{1}{\sqrt{(1-\chi)^2 + 4\chi u_e^2/c^2}} \ln \frac{-2u_e v/c^2 + (1-\chi) - \sqrt{(1-\chi)^2 + 4\chi u_e^2/c^2}}{-2u_e v/c^2 + (1-\chi) + \sqrt{(1-\chi)^2 + 4\chi u_e^2/c^2}}. \tag{4.15}$$

Note that in the limit $\chi \rightarrow 1$, Equations 4.14 and 4.15 reduce to Equations 4.11 and 4.12, respectively. Applying Equation 4.15 at the initial ($M = m_0$, $v = 0$) and final ($M = m$, $v = \Delta V$) boundary conditions, we have that

$$\frac{m}{m_0} = \frac{\left(-2g_0 I_{sp}\Delta V/c^2 + (1-\chi) - \chi_0\right)\left(1 - \chi + \chi_0\right)}{\left(-2g_0 I_{sp}\Delta V/c^2 + (1-\chi) + \chi_0\right)\left(1 - \chi - \chi_0\right)}^{-\frac{1}{\chi_0}}, \tag{4.16}$$

where $\chi_0 = \sqrt{(1-\chi)^2 + 4\chi {g_0}^2 I_{sp}^2/c^2}$.

Figure 4-1 depicts the dry mass fractions for relativistic rockets with various mass-energy conversion efficiencies. For the sake of comparison with Frisbee's results we have assumed a specific impulse $g_0 I_{sp} = \frac{1}{3}c$ (roughly 10^7 s); already we see a fantastic improvement over conventional chemical rockets, due largely to the relativistic velocities of the ejecta. The direction of decreasing efficiency is indicated, and the trend is just as we expect; decreasing

conversion efficiency implies a larger fraction of wasted propellant, and thus a lower dry mass fraction. The red curve denotes the particular case ($\chi = 0.223$) of proton-antiproton annihilation, where non-unity conversion efficiency reduces the attainable dry mass fraction by a factor of two to three $(2-3)$ for a given ΔV.

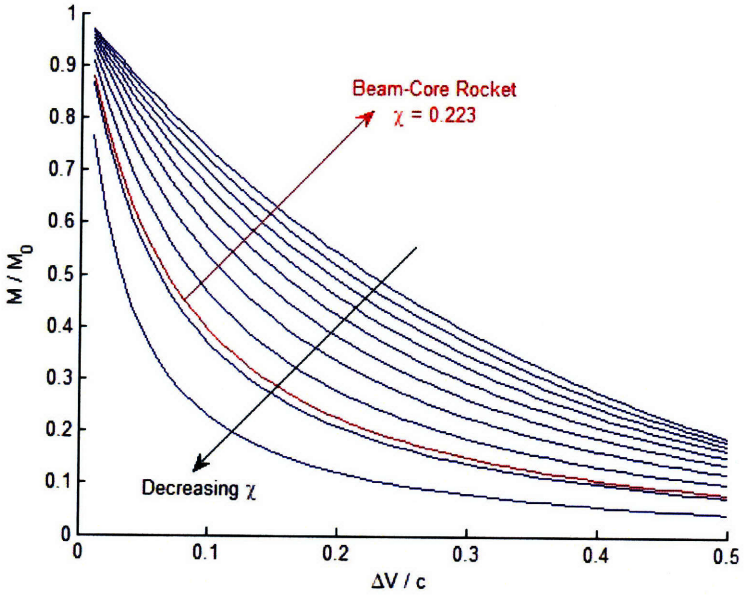

Figure 4-1: Dry Mass Fractions for Antimatter Rockets ($g_0 I_{sp} = \frac{1}{3}c$) with Various Mass-Energy Conversion Efficiencies

It is readily apparent that there is considerable room for improvement upon the beam-core concept. However, in order to truly appreciate the potential of antimatter propulsion in general, let us compare this – the least efficient of all antimatter engines – to a few present-day propulsion systems. Table 4.2 compares total propellant masses at various ΔVs for a beam-core rocket, typical Hall and Ion thrusters, and the Space Shuttle Main Engine. A 100 kg payload is assumed, and we find that there is indeed no contest; for relativistic ΔVs, the required mass of chemical propellants exceeds the mass of the observable universe!

ΔV	**Beam-Core** ($g_0 I_{sp} \simeq \frac{1}{3}c$)	**Ion Engine** ($I_{sp} \simeq 3000$ s)	**Hall Thruster** ($I_{sp} \simeq 1500$ s)	**SSME** ($I_{sp} \simeq 452$ s)
1 km/s	4.5 g	3.5 kg	7.0 kg	25.3 kg
10 km/s	44.9 g	40.5 kg	97.3 kg	853.8 kg
100 km/s	449.0 g	$2.9 \cdot 10^3$ kg	$8.9 \cdot 10^4$ kg	$6.2 \cdot 10^{11}$ kg
1000 km/s	4.5 kg	$5.7 \cdot 10^{16}$ kg	$3.3 \cdot 10^{31}$ kg	$8.8 \cdot 10^{99}$ kg

Table 4.2: Total Propellant Masses for 100 kg Payload Using Various Modern Propulsion Concepts

4.3 Nuclear Fission Catalysis

As shown above, the significant fraction ($\sim$ 77.7%) of "lost" mass-energy associated with beam-core propulsion places a fundamental limit on propulsive performance. One means of improving efficiency is the catalysis of nuclear fission within a fissionable material by a stream of incident antiprotons. Just as the annihilation of proton-antiproton pairs to produce a reaction mass of charged pions, nuclear fission produces a slew of energetic particles, many of which possess significant mass-energy and charge and are thus viable sources of thrust.

We will begin with a review of the basic physics of fission processes. Unless otherwise noted, the data presented in the following sections was obtained via various simulations in Geant4.[2]

4.3.1 Fission Cross-Sections

The mechanisms by which protons and antiprotons split atomic nuclei upon interaction are quite dissimilar. In the case of protons, fission occurs when a collision imparts sufficient energy upon a fissile nucleus, causing it to split into two intermediate-mass fragments and anywhere from zero to six (0-6) or more free neutrons. For the sake of validation we can compare the Geant4 data for proton-induced fission cross-sections in a uranium bulk, depicted in Figure 4-2, to published values depicted in Figure 4-3 (reproduced from [29]). Peak interaction occurs for incident protons of kinetic energy $T \geq 100$ MeV, with little to no variation at higher energies, and we find excellent agreement between predicted and observed cross-sections, with divergence only for very-high-energy ($E \geq 100$ MeV) protons. This divergence may be due to inconsistencies or a lack of high-energy data in the Geant4 libraries, but in any event it is irrelevant for our purposes, which will be limited to incident particles whose energies have been degraded for trapping and storage.

Clearly, the cross-section for fission is very low for protons; at best, only one to two in twenty ($1 - 2 : 20$) incident protons per cubic centimeter bulk material will yield a fission event, with significantly lower rates for lower-energy protons. By contrast antiprotons have unity probability of inducing fission once they come to a stop. The most useful metric is thus the range of incident antiprotons (or protons, as the process is electrostatic and thus

[2]Geant4 is a C++ toolkit developed and maintained by CERN for the simulation of particle and radiation transport through matter.

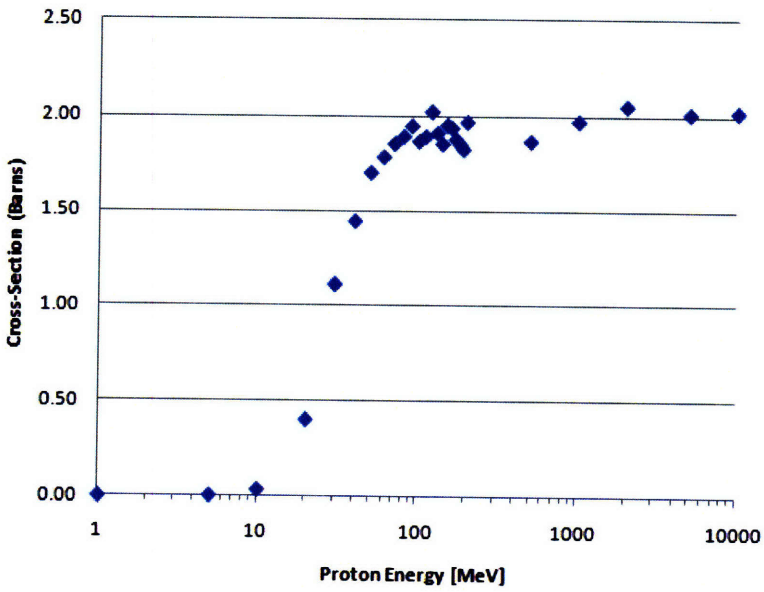

Figure 4-2: Geant4 Cross-Sections for Proton-Induced Fission in ^{238}U

independent of charge or color) in the bulk. Figure 4-4 depicts the mean stopping distance for antiprotons through a bulk of ^{238}U.

We expect that performance will ultimately scale with the rate of fission, and thus the mean stopping distance of antiprotons in the bulk serves as a characteristic thickness for the uranium slab. However, as will be discussed in the following chapter, other factors will place aggressive upper limits on slab thickness and thus a more thorough accounting is in order.

4.3.2 Fission Fragments

The interaction between incident antiprotons and the uranium bulk is perhaps more accurately described as a spallation process, as opposed to traditional fission. Contrary to proton- and neutron-driven fissions, which split the nucleus into two fragments of comparable mass, antiproton annihilation occurs on the surface of the nucleus resulting in the release of numerous light and heavy fragments. This process has been studied in-depth[30], and was simulated in Geant4 to arrive at the following results.

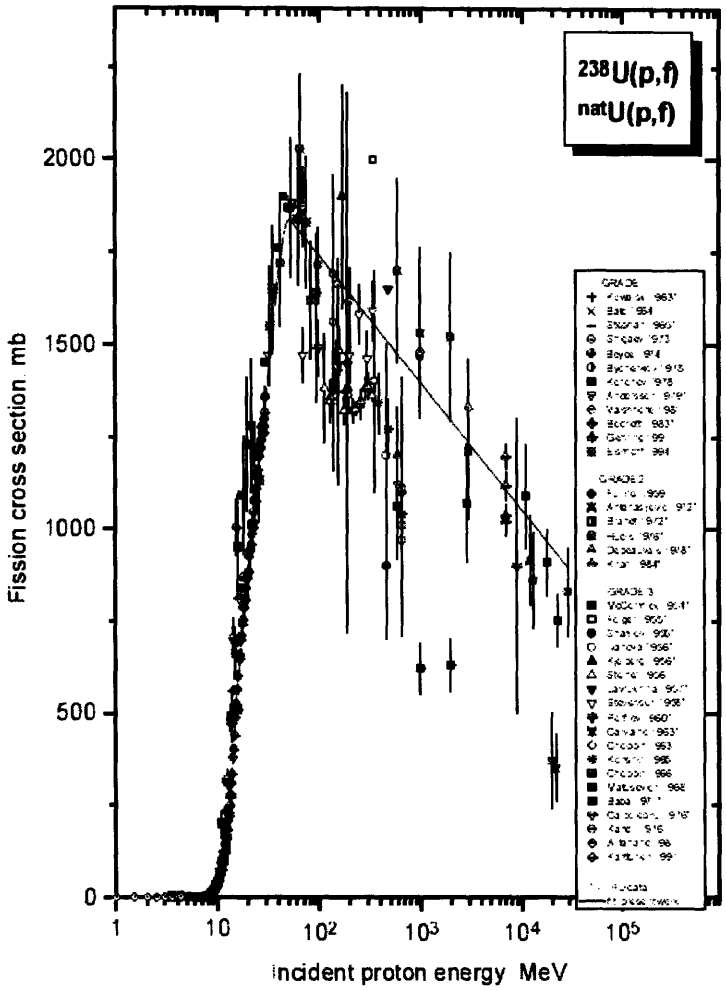

Figure 4-3: Published Cross-Sections for Proton-Induced Fission in ^{238}U

4.3.2.1 Light Fragments

Light fragments, defined for our purposes as those possessing an atomic mass $A < 6$ amu, were found to comprise roughly 84.8% of all antiproton-induced fission ejecta and consist (in order of decreasing abundance) of neutrons, charged pions, protons, deuterons, alphas, and tritons. Trace quantities of kaons, Helium-3 nuclei, electrons, and positrons were also observed. Table 4.3 highlights the average rates of production for the various light fragments per antiproton fission event.

It is interesting to note that the production rate of neutrons for antiproton-induced fission is considerably higher – by a factor of two to three $(2-3)$ – than the corresponding rate for traditional proton- or neutron-induced fission. This will be shown to be advantageous

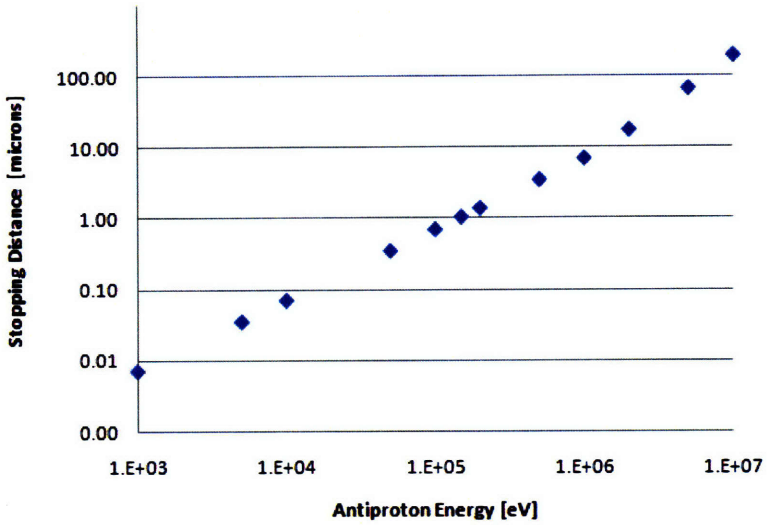

Figure 4-4: Mean Stopping Distance for Antiprotons in ^{238}U

Species	Average # Produced Per Antiproton	Rest Mass [amu]	Electric Charge
n	9.6456	1.0087	0
p	1.0273	1.0073	+1
π^+	0.8652	0.1498	+1
π^-	0.7450	0.1498	+1
^2H	0.2357	2.0136	+1
α	0.1765	4.0015	+2
^3H	0.0992	3.0160	+1
k^+	0.0414	0.5300	+1
^3He	0.0261	3.0160	+2
k^-	0.0208	0.5300	−1
e^-	0.0015	0.0005	−1
e^+	0.0012	0.0005	+1

Table 4.3: Light Fragment Spectrum for Antiproton-Induced Fission in ^{238}U

in later analysis, as a source for subsequent generations of chained fission reactions.

4.3.2.2 Heavy Fragments

Figure 4-5 depicts the production rates of heavy fission fragments for various values of the atomic mass A. Despite comprising only 15.2% of all ejecta, heavy fragments account for roughly 98.2% of the reaction mass. Clearly, then, any practical catalyzed fission propulsion system should be optimized around the extraction and acceleration of heavy fragments. We will explore this design optimization qualitatively in the next chapter.

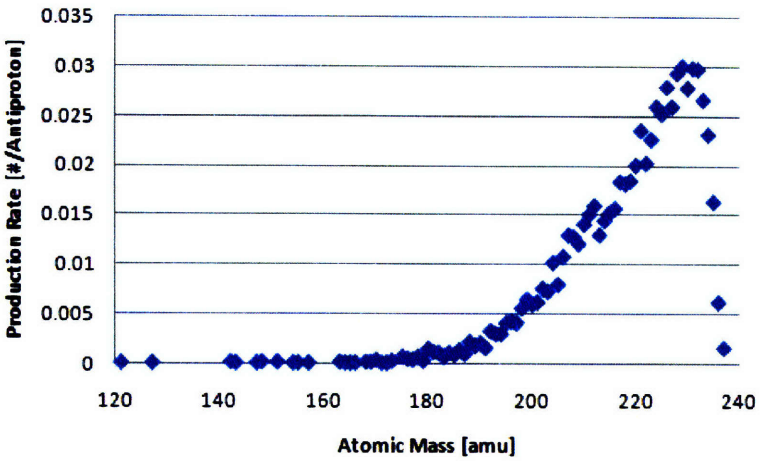

Figure 4-5: Heavy Fragment Spectrum for Antiproton-Induced Fission in ^{238}U

4.3.3 Secondary Fissions

As shown in the previous section, free neutrons represent a significant fraction of the reaction mass generated by the fission process. However, because the neutron flux is isotropic and carries no electric charge, it represents a significant loss in performance. One means of mitigating said loss is to secondary fission reactions. Let us begin by examining the neutron flux generated by antiproton-induced fission. Figure 4-6 depicts a histogram of the energy spectrum for free neutrons generated by antiproton-induced fissions in a bulk of ^{235}U. This data was obtained using the MCNP Neutron Cross Section Library. Most of the spectrum lies between one (1) and 100 MeV, such that the mean free path for interactions is roughly three centimeters (3 cm) (see Figure 4-7).

The fission of enriched ^{235}U by thermal neutrons is a well-documented phenomenon, and typically yields a relatively symmetric fission of the uranium nucleus. Figure 4-8 depicts the relative abundance of various atomic species as fission fragments for thermal neutron fission of ^{235}U. As previously discussed, we note marked dissimilarities relative to antiproton-induced fission; specifically, fission by thermal neutrons yields two distinct, heavy fragments, with peaks in the distribution at roughly 95 amu and 135 amu. Light fragments are atypical, with the exception of neutrons.

For each secondary fission driven by the neutron population, we effectively trade m_n worth of otherwise "lost" reaction mass for two charged, massive fragments, which may then be directed via electromagnetic nozzle for additional thrust. Secondary neutrons ($0 - 6$ per

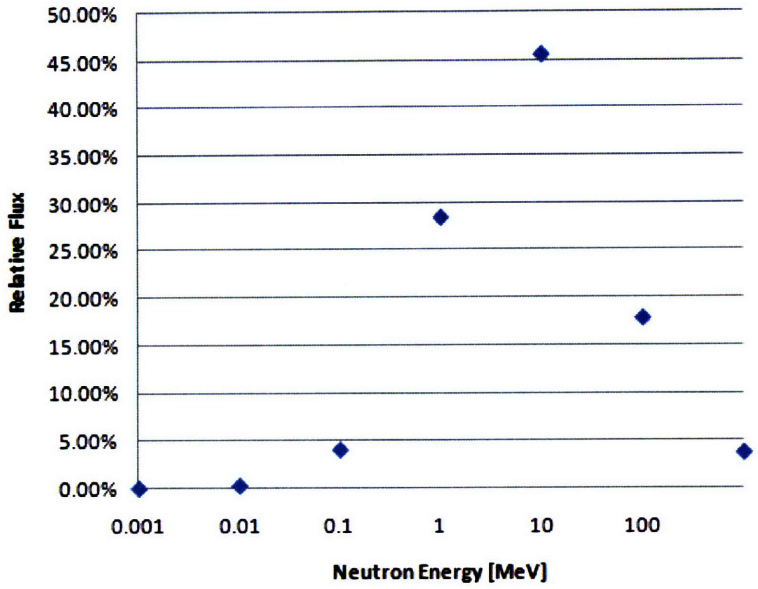

Figure 4-6: Neutron Energy Histogram for Antiproton-Induced Fission in ^{235}U

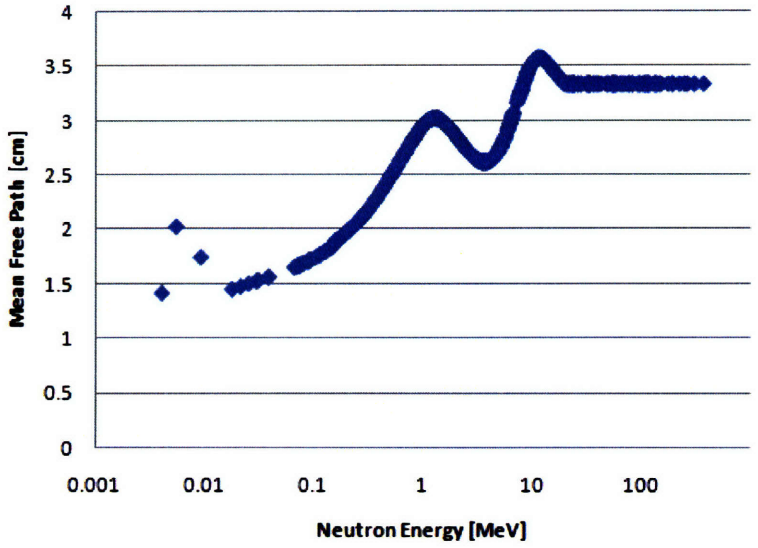

Figure 4-7: Neutron Mean Free Paths in ^{235}U

fission, on average) are also released, which in turn drive further generations of induced fission reactions given sufficient bulk mass and an appropriate geometry. The following section will discuss in further detail the effects of secondary neutron fission, and its impacts on propulsive performance.

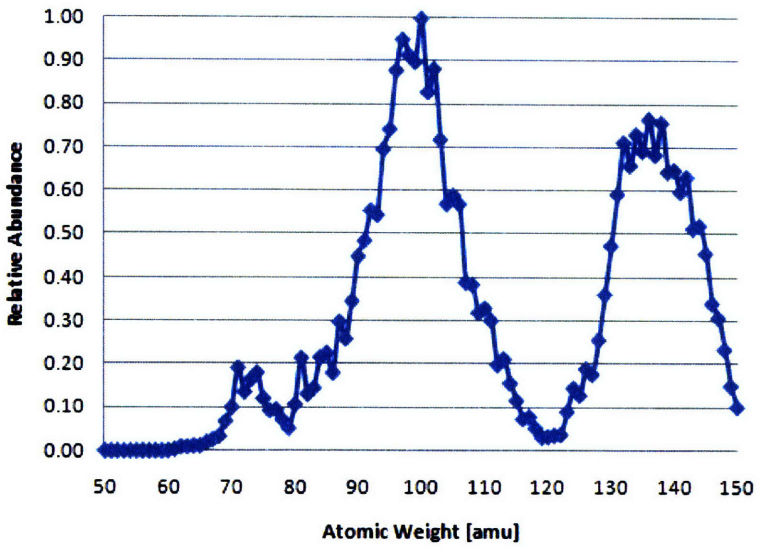

Figure 4-8: Fragment Distribution for Thermal Neutron Fission of ^{235}U

4.4 Propulsive Performance

Having characterized the spectra for both antiproton- and neutron-induced fission ejecta, we may now proceed to assess propulsive performance. Recall from Sections 4.1 and 4.2 that the performance of antimatter rockets is largely driven by the mass-energy conversion efficiency χ. For beam-core propulsion, the annihilation of proton-antiproton pairs results in a reaction mass of charged pions and is characterized by a conversion efficiency of $\chi \simeq 0.223$. For catalyzed fission, the presence of heavy fission fragments in the reaction mass together with the driving of secondary fissions via thermal neutrons results in a marked increase in mass-energy conversion efficiency.

4.4.1 Catalyzed Fission

Figure 4-9 depicts the efficiency χ for various values of the neutron fission fraction - that is, the fraction of secondary neutrons that trigger further fission reactions. Here we assume, as described in Section 4.3.2.1, an average of ~ 9.65 free neutrons per primary (antiproton) fission event, as well as an average of three (3) neutrons per secondary (neutron) fission event. The curves shown represent increasingly more aggressive estimates of performance, due to an increasing number N_g of neutron generations[3] triggering fissions.

[3]For the purposes of this analysis, we can approximate crudely the limiting case of a critical reaction as the limit of successive "generations" of neutrons which induce a fission event. For example, the "first"

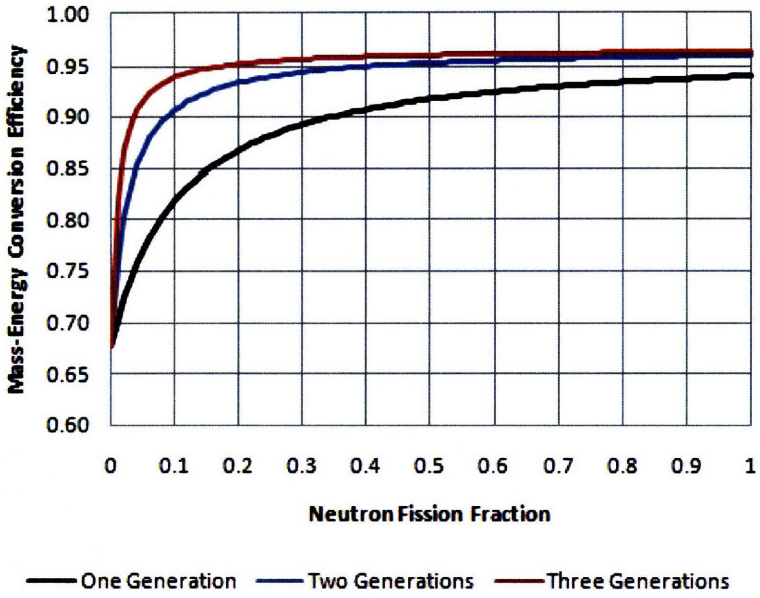

Figure 4-9: Mass-Energy Conversion Efficiency vs. Neutron Fission Fraction for Catalyzed Fission of ^{235}U

As expected, higher neutron fission fractions diminish the fraction of wasted neutron mass thereby improving efficiency. Indeed, we can liken the neutron fission fraction to the degree of enrichment in the uranium sample - thermal neutrons are typically absorbed by the natural isotope ^{238}U, whereas enriched ^{235}U is highly fissile for thermal neutrons. Further, efficiency is found to increase as the neutron-induced fission reaction approaches criticality (i.e., the limit $N_g \rightarrow \infty$). Assuming a reasonably enriched sample (say, 50% ^{235}U) we can observe more closely the relationship between χ and the sustainability of the neutron fission process. Figure 4-10 depicts the increase in χ towards a maximum value at criticality.

4.4.2 Comparative Performance

We may now compare directly the beam-core and catalyzed fission concepts. Assuming a sample of enriched (50% ^{235}U) uranium, Table 4.4 lists mass-energy conversion efficiencies for various operating assumptions.

Assuming the complete extraction of all fission fragments, even absent the effects of sec-

generation of neutrons are those liberated via the primary (antiproton) fission event; the "second" generation of neutrons are those liberated via fission events induced by "first" generation neutrons; and so on. Assuming a reasonable neutron fission fraction $f_n \simeq 0.5$ yields reasonably fast convergence, such that only four to five ($N_g \simeq 4-5$) generations are required to simulate criticality.

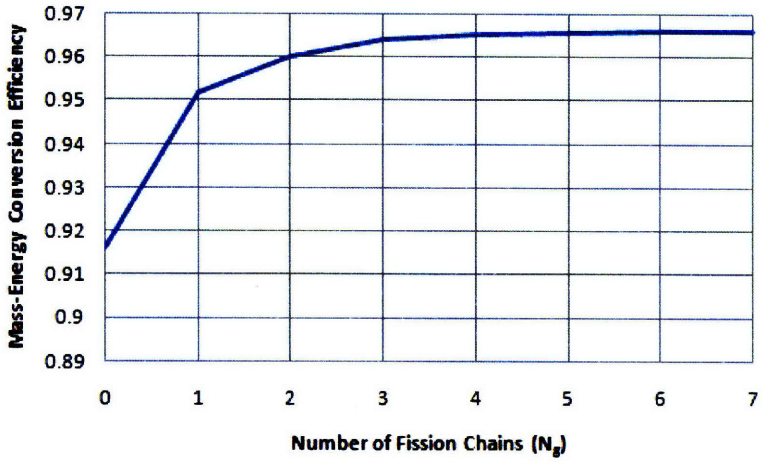

Figure 4-10: Mass-Energy Conversion Efficiency vs. Criticality for Enriched (50% ^{235}U) Uranium

Concept	Assumption	χ
Beam-Core	N/A	0.223
Catalyzed Fission	No Secondary Neutron Fission	0.677
	Critical Secondary Neutron Fission	0.967

Table 4.4: Mass-Energy Conversion Efficiencies for Beam-Core and Catalyzed Enriched (50% ^{235}U) Uranium Fission

ondary fission chains the catalyzed fission concept is roughly three (3) times as mass-efficient as traditional beam-core propulsion. If the free neutron population is allowed to induce a critical fission reaction, the efficiency soars to over 96.7%, with the remaining losses due primarily to the left-over neutron population. This represents a fantastic improvement, and can be attributed directly to the abundance of heavy fission fragments in the reaction mass. In the next chapter we will explore the validity of our assumptions regarding extraction, but for now it suffices to say that very generous extraction efficiencies could, in theory, be realized given adequate design and engineering.

We may relate these mass-energy conversion efficiencies to dry, or payload mass fractions for various ΔV as in Section 4.2. Figure 4-11 depicts the mass fractions for both beam-core and catalyzed fission, again assuming a specific impulse $g_0 I_{sp} = \frac{1}{3}c$. The latter concept is shown for both no secondary fission, and critical secondary fission.

For a given total mass, catalyzed fission rockets may accelerate roughly *twice* as much payload mass to appreciable fractions of the speed of light. The addition of a critical neutron fission chain further increases this ratio to roughly two and one-half (2.5).

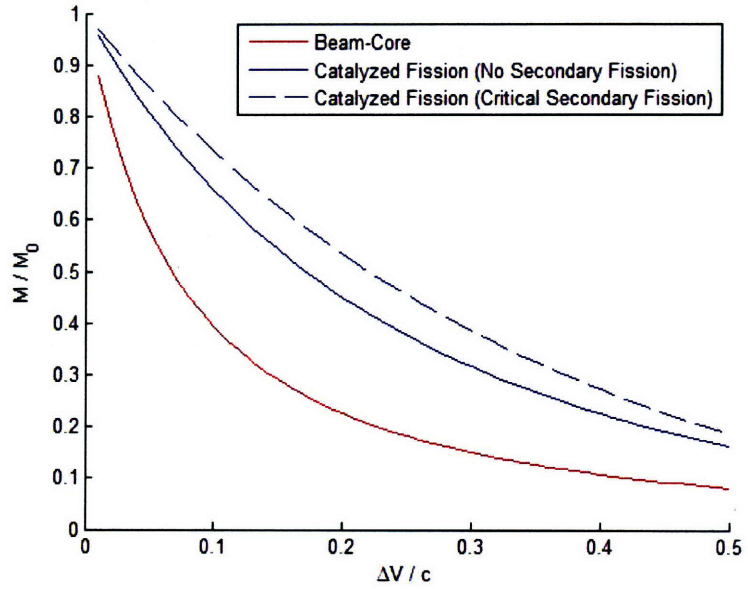

Figure 4-11: Dry Mass Fractions for Beam-Core and Catalyzed Fission Rockets

4.4.3 Antimatter Fuel Mass Fraction

Secondary fission of the uranium bulk by the neutron population also provides the added benefit of reducing the overall antiproton fuel mass fraction; that is, a greater fraction of the reaction-mass conversion process is driven by neutrons born of the fission process, as opposed to the expensive and tenuous antiproton population. Figure 4-12 depicts the fraction of fuel mass apportioned to antiprotons as the secondary fission process approaches criticality.

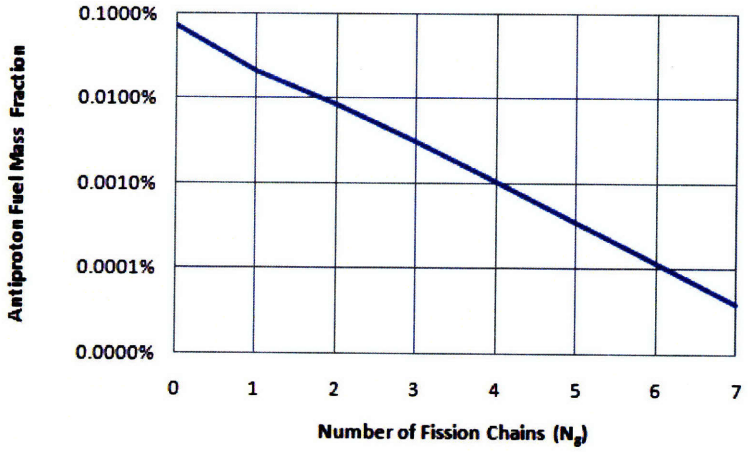

Figure 4-12: Antiproton Fuel Mass Fractions for Various Degrees of Criticality

As expected, the antiproton fuel mass fraction varies exponentially with criticality; mass fractions of one part in 10^6-10^8 ($0.0001\% - 0.000001\%$) are not unreasonable, and suggest that one kilogram (1 kg) of uranium could be fissioned by as little as 10 micrograms to 1 milligram (10-1000 μg) of antiprotons. However, given sufficient engineering and an appropriate geometry, the efficiency is theoretically limitless – there is nothing to suggest that a single antiproton could not trigger a critical fission chain. In any event, these antiproton mass-efficiencies represent a monumental improvement over the beam-core concept, where fully one-half of the fuel mass is composed of antiprotons.

Chapter 5

Summary and Conclusions

The purpose of this research was to assess the viability and performance potential of naturally occurring antiprotons as a fuel for spacecraft propulsion. Current particle accelerator technology limits laboratory-based production of antiprotons, and reliable long-term storage of antimatter has yet to be demonstrated at scales sufficient for practical propulsion systems. To that end, a novel electromagnetic collection system was described and analyzed for the harvesting of cosmic ray and planetary radiation belt antiprotons, and the performance of antimatter rockets was assessed for both beam-core annihilation and antiproton-catalyzed nuclear fission.

5.1 General Findings

5.1.1 Antiproton Collection

Chapter 3 analyzed various electromagnetic collection systems, to assess the feasibility of harvesting antiprotons from planetary magnetospheres. Among the concepts studied, a novel design derived from Slough's[21] magnetic sail concept (the "plasma magnet") was found to achieve high rates of collection, with mass- and power-specific performance vastly superior to the present state of the art in Earth-based production. Shallower gradients in the induced magnetic field, together with the use of a distributed plasma as the charge carrying species, combined to account for the significant reduction in mass and power relative to single- and concentric-loop systems. Table 5.1 outlines a baseline plasma magnet collection system.

The relatively low mass (12,000 kg) suggests that complete systems could conceivably

Parameter	Value
Antenna Coil Radius	100 m
Operating Current	10^5 A
Plasma Density	$2 \cdot 10^{16}$ m^{-3}
Plasma Temperature	15 eV
Coil Mass	6600 kg
Total Power	200 kW
Power System Mass	5,200 kg
Total Mass	12,000 kg
Collection Rate	8.6 μg/yr

Table 5.1: Baseline Plasma Magnet Antiproton Collection System

be deployed with a single launch, and the stated power draw is in line with near-future nuclear power system concepts, such as NASA's Prometheus program.

5.1.2 Antiproton Propulsion Systems

Chapter 4 characterized the propulsive performance of various antiproton propulsion systems. Specifically, catalyzed fission rockets (with and without secondary neutron fission) were compared to simple beam-core rockets, and were found to attain significantly greater – in most cases, by a factor of two to three $(2-3)$ – dry mass fractions. This, in turn, led to significant reductions in the antiproton fuel mass fraction. To conclude our analysis, let us translate our results into sample mission architectures and attempt to arrive at some qualitative conclusions on engine design.

5.1.2.1 Sample Mission Architectures

Given the propulsive performance estimates derived in the previous chapter, we may compare directly the various antiproton propulsion systems in terms of a sample mission architecture. Starting small, we choose as a baseline mission a one-way transfer of a 100 kg payload to Mars, requiring $\Delta V \simeq 5$ km/s. Table 5.2 lists the various mass fractions and propellant masses for several different spacecraft configurations; for comparison, figures are included for a mission utilizing standard chemical propellants ($I_{sp} \simeq 300$ s), a typical Hall thruster ($I_{sp} \simeq 1500$ s) and a catalyzed fission antimatter engine.

The use of antiproton fuels results in staggering reductions in propellant mass – for the catalyzed fission engine, a total of 5.174 g of propellant is consumed, only ~ 51.74 ng of

Parameter \ Concept	Chemical Rocket	Hall Thruster	Catalyzed Fission
Mass-Energy Conversion Efficiency χ	1	1	0.967
Payload Mass	100 kg		
Dry Mass Fraction	0.1829	0.7119	$1 - 5.174 \cdot 10^{-5}$
Propellant Mass	446.747 kg	40.469 kg	5.174 g
Annihilated Antiproton Mass	N/A	N/A	~ 51.74 ng

Table 5.2: Sample Mars Mission Architectures for Various Propulsion Concepts

which is in the form of antiprotons. Here, again, it is worth examining in more detail the precise value of the antiproton fuel mass fraction. Recall from Section 4.4.3 that the number of antiprotons required to initiate a critical reaction can in theory be made arbitrarily small given sufficient engineering and an appropriate geometry. The antiproton fuel mass fraction was also found to vary exponentially with the assumed number of neutron generations N_g triggering fissions. For all mission architectures described in this section, we have assumed a value $N_g = 10$ such that $m_{\bar{p}}/m_{fuel} = 10^{-8}$.

Consider now a mission to the Sun's gravity focus. At a distance of roughly 550 AU ($\sim 8.2 \cdot 10^{10}$ km), the gravity focus is a prime example of a scientifically interesting target that is inaccessible via modern propulsion technologies. We choose again a 100 kg payload, and select a cruise speed $v = 0.005c$ and acceleration $a = 0.01g_0$, where $g_0 = 9.81$ m/s^2 is the acceleration due to gravity at the Earth's surface. Table 5.3 lists the various mass fractions, propellant and launch masses, and flight times associated with the various antiproton propulsion concepts, so as to demonstrate relative performance. A relatively fast (~ 2 year) mission can be accomplished with microgram quantities of antiprotons!

Parameter \ Concept	Beam-Core	Catalyzed Fission (No Secondary)	Catalyzed Fission (Critical Secondary)
Mass-Energy Conversion Efficiency χ	0.223	0.677	0.967
Payload Mass	100 kg		
Dry Mass Fraction	0.8799	0.9570	0.9695
Propellant Mass	13.65 kg	4.49 kg	3.15 kg
Annihilated Antiproton Mass	6.83 **kg**	3.24 **g**	< 31.5 **μg**
Acceleration Time	0.4842 yr		
Cruise Time	1.2552 yr		
Deceleration Time	0.4842 yr		
Total Flight Time	2.2236 **yr**		

Table 5.3: Sample Mission Architectures for a Mission to the Solar Gravity Focus Using Various Antiproton Propulsion Concepts

Looking ahead to the possibility of interstellar exploration, we choose a one-way transfer

of the same 100 kg payload to the Alpha Centauri system, at a distance of roughly 4.3 light-years ($\sim 4.1 \cdot 10^{13}$ km) from Earth. A cruise speed $v = 0.2c$ and an acceleration $a = 0.01g_0$ are assumed. Table 5.4 lists the various mass fractions, propellant and launch masses, and flight times associated with the various antiproton propulsion concepts, so as to demonstrate relative performance.

Parameter \ Concept	Beam-Core	Catalyzed Fission (No Secondary)	Catalyzed Fission (Critical Secondary)
Mass-Energy Conversion Efficiency χ	0.223	0.677	0.967
Payload Mass	100 kg		
Dry Mass Fraction	0.1081	0.2254	0.2724
Propellant Mass	825.07 kg	343.66 kg	267.11 kg
Annihilated Antiproton Mass	412.53 **kg**	247.78 **g**	< 2.67 **mg**
Acceleration Time	19.3681 yr		
Cruise Time	2.1309 yr		
Deceleration Time	19.3681 yr		
Total Flight Time	40.8670 **yr**		

Table 5.4: Sample Interstellar Mission Architectures for Various Antiproton Propulsion Concepts

As expected, the presence of heavy fragments in the reaction mass – and thus higher conversion efficiency χ – for catalyzed fission rockets leads to a significant decrease in fuel mass, and thus overall launch mass. The introduction of the uranium bulk also reduces the antiproton fuel mass fraction, an effect which is amplified in the presence of secondary neutron fissions within the bulk. For the critical catalyzed fission case, a total required antiproton mass of 2.67 mg is not altogether discouraging. Recall from Chapter 4 that the antiproton fuel mass fraction is directly related to the criticality of the fission reaction; it is not unreasonable to assume that the fraction could be made arbitrarily low given sufficient bulk mass and a clever arrangement of bulk slabs and neutron-reflecting surfaces. Figure 5-1 depicts the relationship between the assumed antiproton fuel mass fraction and the attainable ΔV for given quantities of antiprotons. The potential for enormous gains in performance at low antiproton fuel mass fractions is readily apparent.

5.1.2.2 Qualitative Remarks on Catalyzed Fission Engines

As described in previous sections, the success of the catalyzed-fission concept is predicated in part on the ability of the free neutron population to induce fissions in the uranium bulk. The characteristic length scales for these interactions are well-defined, and may be

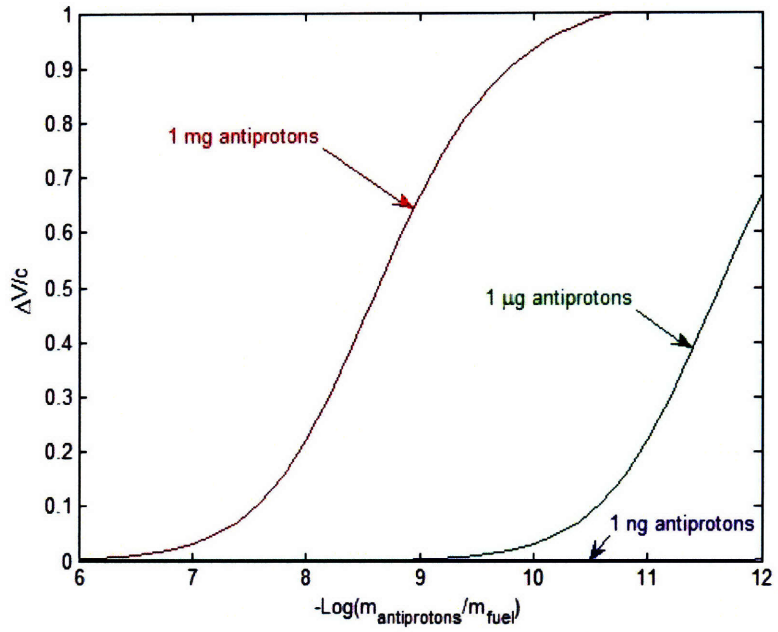

Figure 5-1: Catalyzed Fission Rocket Performance for Various Assumed Values of the Antiproton Fuel Mass Fraction

considered a good approximation of the effective traverse through the bulk required by a given neutron before it will induce a fission. These mean free paths for interaction were characterized as a function of neutron kinetic energy in Figure 4-7.

A more limiting factor, however, is the extraction of heavy fission fragments from the bulk. Contrary to light fragments – whose cross-sections for interaction were sufficiently low so as to allow them to escape arbitrarily thick slabs – heavy fragments were typically found to possess very short ranges within the bulk. Specifically, slab thicknesses exceeding one-tenth of a micron (0.1 μm) were sufficient to fully ionize most heavy fragments and thus prevent them from leaving the bulk and contributing to the reaction mass.

In light of these limitations, let us make a few qualitative observations about the design of a catalyzed fission engine. In particular, we will require two principal components: (i) a uranium target, arranged in such a way as to maximize the extraction of heavy fission fragments; and (ii) an electromagnetic nozzle (a simple current loop or permanent magnet would suffice, though an optional electrostatic accelerator grid may be added for increased performance) to deflect and accelerate the charged products of annihilation and in turn provide thrust.

In order to maximize the extraction of heavy fragments from the uranium bulk, recall

that we are limited to sub-micron slab thicknesses. For thermal neutrons with mean free paths for interaction of several centimeters, the number of slabs required would be prohibitively high (in the limiting case of a fission occurring at one end of the chamber, and a corresponding neutron fission occurring at the opposite end, the number of required slabs exceeds 10^5, assuming a slab thickness $t_s = 0.1$ μm. One possibility for improving performance is the use of neutron-reflecting materials on the chamber walls. This would increase the residence time of the neutron population within the chamber, thereby reducing the number of slabs required to sustain secondary fissions. Most neutron reflectors also have the fortunate effect of degrading the kinetic energy of reflected neutrons, which would in turn decrease their mean free paths and increase the rate of fission in the bulk.

5.2 Recommendations for Future Research

Several topics would stand to benefit from further consideration, particularly with respect to collection system modeling. A higher-fidelity model of the plasma magnet, in particular the RMF assembly and its interaction with the bulk plasma, would help to validate the results stated herein; specifically, additional transport modeling should be carried out to account for collisional and thus diffusive effects in the electron population. A full MHD equilibrium model would be ideal for the characterization of the driven electron current, though two-dimensional PIC modeling would be an adequate start.

Further research could also serve to refine estimates of storage capacity and trapping stability for the various collection systems studied. One possible avenue is the modeling of non-neutral plasmas for the circumvention of space-charge saturation in the Brillouin limit. A thorough treatment of this problem would have immediate real-world applications in the medium- to long-term storage of antimatter. The specifics regarding the transfer of incident particles onto closed field lines following energy degradation would also stand to benefit from further attention.

Finally, any practical propulsion or collection system must be designed to operate in tandem with other spacecraft systems. The use of superconducting materials and the presence of high-energy gamma radiation associated with the various antimatter engine concepts only serve to compound the difficulties associated with design. Further research should be conducted at a systems level, particularly in the areas of thermal and dynamic control, and

radiation shielding.

[This Page Intentionally Left Blank]

References

[1] P. A. M. Dirac. The Quantum Theory of the Electron. In *Proceedings of the Royal Society (London) A*, volume 117, pages 610–612, 1928.

[2] P. A. M. Dirac. The Quantum Theory of the Electron: Part II. In *Proceedings of the Royal Society (London) A*, volume 118, pages 351–361, 1928.

[3] European Organization for Nuclear Research (CERN). "Antiprotons Four Times More Effective than Protons for Cell Irradiation". Press Release, 31 October 2006.

[4] G. R. Schmidt, H. P. Gerrish, J. J. Martin, G. A. Smith, and K. J. Meyer. Antimatter Production for Near-term Propulsion Applications. In 35^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, 1999.

[5] B. W. Augenstein. "RAND Workshop on Antiproton Science and Technology, October 6-9, 1987: Annotated Executive Summary". Note N-2763-AF, RAND Corporation, October 1988.

[6] M. R. LaPointe. Antimatter Production at a Potential Boundary. In 37^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-2001-3361, 2001.

[7] H. Hora. Estimates of the Efficient Production of Antihydrogen by Lasers of Very High Intensities. *OptoElectronics*, 5:491–501, 1973.

[8] G. Chapline. Antimatter Breeders. *Journal of the British Interplanetary Society*, 35: 423–424, 1982.

[9] B. N. Cassenti. Concepts for the Efficient Production and Storage of Antimatter. In 29^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-93-2031, 1993.

[10] J. A. Bickford. "Extraction of Antiparticles Concentrated in Planetary Magnetic Fields". NIAC Phase II Final Report, The Charles Stark Draper Laboratory, Inc., August 2007.

[11] M. H. Holzscheiter, R. A. Lewis, E. Mitchell, J. Rochet, and G. A. Smith. Production and Trapping of Antimatter for Space Propulsion Applications. In *American Institute of Physics Conference*, 1997.

[12] S. D. Howe and G. A. Smith. Enabling Exploration of Deep Space: High Density Storage of Antimatter. In 35^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-99-2692, 1999.

[13] R. H. Frisbee. How to Build an Antimatter Rocket for Interstellar Missions. In 39^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-2003-4676, 2003.

[14] R. H. Frisbee and S. D. Leifer. Evaluation of Propulsion Options for Interstellar Missions. In 34^{th} *Annual AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, number AIAA-98-3476, 1998.

[15] General Atomic. "Nuclear Pulse Space Vehicle Study: Vol. I – Summary". Technical report, performed under NASA Contract NAS 8-11053, George C. Marshall Space Flight Center, September 1954.

[16] C. Störmer. *The Polar Aurora.* Oxford Press, 1955.

[17] F. F. Chen. *Introduction to Plasma Physics and Controlled Fusion.* Plenum Press, second edition, 1984.

[18] M.-B. Kallenrode. *Space Physics: An Introduction to Plasmas and Particles in the Heliosphere and Magnetospheres.* Springer-Verlag, third edition, 2004.

[19] S. Humphries, Jr. *Principles of Charged Particle Acceleration.* John Wiley and Sons, 1986.

[20] G. Jackson. "Antimatter Harvesting in Space". NIAC Phase I Final Report, Hbar Technologies, LLC, March 2006.

[21] J. Slough. "The Plasma Magnet". NIAC Phase II Final Report, The University of Washington, Department of Aeronautics and Astronautics, September 2006.

[22] D. B. Montgomery and J. Terrell. "Some Useful Information for the Design of Air-Core Solenoids". Technical report, Francis Bitter National Magnet Laboratory, 1961.

[23] G. Pugacheva, A. A. Gusev, U. B. Jayanthi, N. J. Schuch, W. N. Spjeldvik, and C. H. Choque. Antiprotons Confined in the Earth's Inner Magnetosphere. *Astroparticle Physics*, 20:257–265, 2003.

[24] R. Zubrin. "The Magnetic Sail". NIAC Phase II Final Report, Pioneer Astronautics, January 2000.

[25] S. J. Isakowitz, J. B. Hopkins, and J. P. Hopkins, Jr. *International Reference Guide to Space Launch Systems.* American Institute of Aeronautics and Astronautics, 2004.

[26] W. J. Larson and L. K. Pranke. *Human Spaceflight: Mission Analysis and Design.* The McGraw-Hill Companies, Inc., 2000.

[27] J. A. Bickford. "Extraction of Antiparticles Concentrated in Planetary Magnetic Fields". NIAC Phase I Final Report, The Charles Stark Draper Laboratory, Inc., April 2006.

[28] L. Brillouin. A Theorem of Larmor and Its Importance for Electrons in Magnetic Fields. *Physical Review*, 67:260–266, 1945.

[29] M. I. Baznat, K. K. Gudima, and S. G. Mashnik. Proton-Induced Fission Cross Section Calculation with the LANL Codes CEM2K+GEM2 and LAQGSM+GEM2. arXiv:nucl-th/0307014v1, July 2003.

[30] H. S. Plendl, H. Daniel, T. von Egidy, T. Haninger, F. S. Hartmann, P. Hoffman, Y. S. Kim, H. Machner, G. Riepe, J. Jastrzebski, A. Grabowska, W. Kurcewicz, P. Lubinski, A. Stolarz, A. S. Botvina, Ye. S. Golubeva, A. S. Iljinov, V. G. Nedorezov, A. S. Sudov, and K. Ziock. Antiproton-Nucleus Annihilation at Rest. *Physica Scripta*, 48:160–163, 1993.