Application of Bayesian Inference to Develop an Air-Core Magnetic Circuit for a Magnetically Shielded Hall Thruster

William Hurley*, Thomas Marks †, Alex Gorodetsky‡, and Benjamin Jorns§
University of Michigan, Ann Arbor, MI, 48109

The magnetic field profiles of two Hall thrusters, the H-9 and the SPT-100, are reproduced computationally using an air-core magnetic circuit. Air-core circuits avoid the saturation limits exhibited by ferromagnetic material that can limit achievable field strengths. A Bayesian framework was implemented to learn the current needed in each loop of wire to match a set of magnetic field data. The results are validated through COMSOL, a multi-physics tool. A trade space is explored that balances accuracy of the magnetic field (compared to measurement data), number of coils, and maximum current in any one coil. It is found that the current needed to reproduce and then increase the magnetic field strength may be prohibitively large for traditional copper wires. A novel design solution to this problem using bitter magnets is discussed.

I. Nomenclature

$\bar{A}$	=	Matrix of Geometric factors
a	=	current loop radius
α	=	Constant
B_r	=	Radial Magnetic field
B_z	=	Axial Magnetic field
B_{mod}	=	Model magnetic field
B_{msd}	=	Measured magnetic field
d	=	Magnetic Field data set
d'	=	Revised magnetic field data set
E	=	Electric Field
E_1	=	Elliptic integral of the first kind
E_2	=	Elliptic integral of the second kind
$f(\vec{\theta}$	=	Prior
Γ	=	Noise covariance matrix
g	=	geometric factor: Axial Magnetic Field
h	=	geometric factor: Radial Magnetic Field
I	=	Current
I_b	=	Ion beam current
I_e	=	Electron Current
J_e	=	Electron Current Density
J_i	=	Ion Current Density
L	=	Acceleration Zone length
$L(d\|\vec{\theta})$ Likelihood function μ	=	Prior mean
μ_{post}	=	Posterior Mean
μ_0	=	Permeability of free space
m_e	=	Electron mass
m_i	=	Ion mass

*PhD Candidate, Department of Aerospace Engineering
†PhD Candidate, Department of Aerospace Engineering, AIAA Student Member
‡Assistant Professor, Department of Aerospace Engineering
§Associate Professor, Department of Aerospace Engineering, AIAA Associate fellow

n	=	Number density
η_b	=	current utilization efficiency
P	=	Pressure Tensor
q	=	Fundamental Charge
σ_{en}	=	Electron neutral cross section
Σ	=	Prior covariance matrix
Σ_{post}	=	Posterior covariance matrix
T_e	=	Electron Temperature
$\vec{\theta}$	=	vector of parameters
u_n	=	neutral velocity
V	=	Discharge Voltage
$\vec{v}$	=	velocity
v_{te}	=	electron thermal velocity
ν_{ex}	=	electron collision frequency with species 'x'
y	=	current loop model
ζ	=	Gaussian Noise

II. Introduction

Hall thrusters, axisymmetric electric propulsion devices that utilize orthogonal electric and magnetic fields, are attractive candidates for scaling to the high power operation (>100 kW) needed to support crewed, interplanetary exploration [1]. This stems from their high efficiency (>60%), specific impulse (1500-2000s), and long operational lifetimes (>10,000 h) [2, 3]. With that said, state-of-the-art Hall thrusters operate at the 10 kW level, and only preliminary studies have been performed to demonstrate this technology at the targeted 100 kW for crewed exploration. These studies ultimately have revealed there are many technical challenges that remain to achieve this order of magnitude increase in capability [4, 5].

One of the major challenges is the specific mass of the thruster. Specific mass in this context refers to the ratio of thruster mass to input power. This metric is a crucial consideration for reducing footprint and also for enabling more rapid acceleration. [6] While state-of-the-art Hall thrusters have lower specific mass compared to technologies like gridded ion thrusters, their specific mass is higher than electromagnetic systems like magnetoplasmadynamic thrusters. The typical Hall thruster specific mass (estimated to be approximately 2.4 kg/kW) [7], stems largely from heuristic design rules that place bounds on the achievable specific power and thrust density. The net consequence of this scaling law is that as Hall thrusters are increased to a larger size to meet required power levels, they can become prohibitively large for crewed deep-space missions. With that said, if Hall thruster power density could be increased by an order of magnitude, this technology could rival magnetoplasmadynamic thrusters as a low specific mass option for high-power EP.

To this end, there have been several efforts to overcome this challenge of increasing power density for Hall thrusters. These generally have followed two strategies: increasing the discharge voltage for a given thruster current density or increasing the current density for a fixed voltage. Efforts in the 1960s adopted the former approach, operating at a discharge current of 10 A and a discharge voltage of 10,000 V (compared to the typical 300-600 V regime of modern Hall thrusters) [4]. While highly efficient (>75%), the high specific impulse (8000 s) led to low thrust-to-power ratios and thus the inability to generate the acceleration required to reach Mars in a one or two-year time horizon [6].

Alternatively, more recent efforts have focused on increasing power density by running more current through a given thruster area for a fixed discharge voltage (specific impulse). This has the benefit of maintaining higher thrust to power while decreasing the specific mass. The challenge with this approach, however, is that there are three factors that may limit the achievable current density. First, with higher flux through the thruster, erosion due to ion bombardment of the discharge chamber walls will increase, limiting thruster life. Second, there will be an increase in thermal flux to the walls, potentially beyond the thermal margins of the materials. Third, there is some evidence that electron confinement is reduced with higher power density [8]. In effect, these limits have resulted in an empirical design rule for current in thrusters, which is they generally should not exceed 100 mA/cm^2.

With these limitations in mind, one strategy that has been explored to date is not to increase current density in the channel but to better utilize the footprint of the thruster. This has resulted in the nested channel design in which multiple channels are stacked concentrically, each of which is still limited by the empirical upper bound in current density. For example, this technique was utilized in the X3, a Hall thruster developed by the University of Michigan, Air Force Research Laboratory, and NASA from 2011-2013 [5]. This system was demonstrated at 100 kW operation at 300 V, a record for Hall thrusters [5]. This program demonstrated a wide throttling capability and a reduced footprint, but the heavy electromagnetic circuits still resulted in a significant mass penalty. This program also revealed a number of potential limitations with the nesting approach. Foremost among these was the added complexity of powering, providing flow, and magnetizing multiple nested channels. Moreover, there are still unresolved questions about how the channels couple together and the overall thruster performance.

In light of these recent results, although channel nesting is a promising approach, ideally we would find a way to overcome this empirical limit in current density. To this end, there is reason to believe–in large part due to recent advances in thruster technology [9, 10]–that this limit may be reconsidered. Most notably, the advent of "magnetic shielding" has impacted both the erosion rates and thermal loading of the channel. This technique, which bends the field lines convexly toward the anode, maintains cool electrons and high electric potentials near the discharge walls, significantly reducing the high energy ion flux into the walls. This strategy may helps mitigate two of the major issues with increasing current density: erosion of the channel walls, and increased thermal flux.

With that said, shielding does not necessarily address the third major limitation related to reduced electron confinement at high current densities. Classically, electrons cross field lines as a result of collisions, with the rate of this process given by the collision frequency. As we increase in current density, the number density of ions and neutrals increase, thus increasing our collision frequency and cross field transport of electrons. In principle, confinement can be improved by increasing the radial magnetic field strength, but SOA thruster designs have material limits that prevent further increase.

For Hall thrusters, there is an upper bound in achievable magnetic field strength due to the onset of saturation in the ferromagnetic material that is employed for the magnetic circuit. Saturation occurs when no additional magnetic field domains within a ferromagnetic material can be aligned with the externally applied field [11]. In order to increase the achievable magnetic field strength, one strategy is to remove the ferromagnetic material and adopt an air-core circuit.

There are two major challenges with an air-core design. The first is that air-cored based circuits can require orders of magnitude more current to generate the same magnetic field strengths as ferromagnetic based circuits. This poses technical obstacles related to thermal loading and implementation. Assuming this thermal problem can be solved (and indeed, there are air core circuits capable of 1 T of field [12]), an equally important question is whether this type of architecture can achieve the types of unique field shapes necessary for magnetic shielding. In light of this challenge, the need is apparent for a detailed investigation into the feasibility of a shielded architecture with an air core circuit

The goal of this work is to design an air-core circuit (one free of ferromagnetic material) that can provide a shielded geometry. To this end, this paper is organized in the following way. In the first section, we formally present the need for an increased magnetic field strength, demonstrate the problem of magnetic saturation, and lay out a framework to develop an air core circuit. In the next section, we introduce the two Hall thrusters and their magnetic fields that we will use the air-core circuit to match. We then present the results of the simulations, followed by a discussion analyzing these results. Finally, we conclude with a summary of the work done and propose avenues for future work.

III. Theory

Here we formally present the need for increased magnetic field strength with higher current density, j_i, demonstrate how saturation limits this field for current Hall thruster designs, and present an overview of our model set up and relevant inference equations used to develop our air core circuit.

A. The need for higher magnetic field with increasing current density

Fig. 1a shows an idealized geometry for the channel in a Hall thruster. It is characterized by an anode, discharge chamber, inner and outer magnets, and ferromagnetic material. Fig. 1b shows the electron and ion currents in the channel, as well as the magnetic field that resists the electron current.

The transverse magnetic field in Hall thrusters impedes the electron current from the cathode back to the anode. Since $m_i >> m_e$, ions are not magnetized and are accelerated out of the channel by the axial electric field (creating thrust), while electrons spiral around field lines and drift azimuthally. These electrons ionize the propellant and increase the local resistivity which in turn drives an electric field that accelerates the ions. Therefore, maintaining confinement of electrons with a strong magnetic field is paramount to efficient Hall thruster operation.

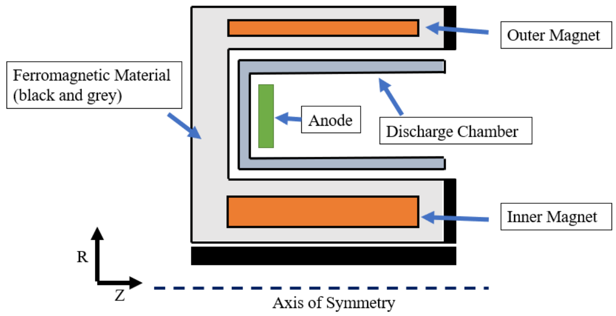

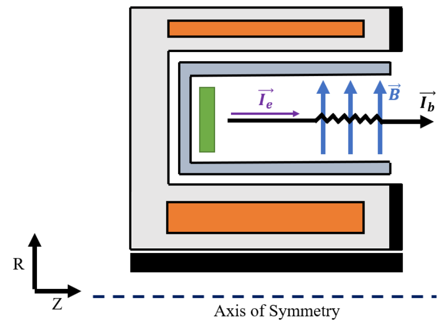

Fig. 1 Typical a) components and b) currents and magnetic field in a Hall thruster.

We can quantify the impact of electron confinement on thruster efficiency with the current utilization:

$$\eta_b = \frac{I_b}{I_d}, \tag{1}$$

where I_b is the ion beam current and I_d is the discharge current. This metric captures the fact that not all current from the supply is converted to accelerated exhaust. Indeed, assuming that the discharge current is approximately equal to the current of electrons to the anode, we can re-write this current utilization efficiency as

$$\eta_b = \frac{I_b}{I_b + I_e}, \tag{2}$$

where I_e is the average electron current from the cathode to the anode. This equation indicates that a major driver for lowering current utilization efficiency (and by extension mass utilization) efficiency is the degree of electron current across the confining field in the channel.

We can approximate the scaling for this cross-field electron current by invoking a generalized Ohm's law averaged over the channel geometry shown in Fig. 1:

$$I_e = A\frac{m_e n_e \nu_{ex}}{B_r^2}\left[E_z + \frac{1}{|q|n_e}\frac{\partial P}{\partial z}\right], \tag{3}$$

where A denotes the channel area, n_e is the electron number density, q is the fundamental charge, ν_{ex} is the electron collision frequency with species x (ions, neutrals), B_r is the radial magnetic field, E_z is the axial electric field, and P is the pressure tensor. Note that we have taken the limit that the electron Hall parameter is large, i.e. $\nu_{ex}/(qB_z/m_e) \gg 1$. This result shows that the electron current depends not only the forces directed along the channel center line, i.e. the electric field and pressure gradients, but the magnetic field and collision frequency as well. Physically, this result captures the intuitive trend that the magnetic field impedes electron motion. Collisions, on the other hand, work to demagnetize the plasma, thereby allowing more cross-field transport.

We in turn can substitute this result into Eq. 2 to determine the scaling

$$\eta_b \approx \frac{1}{1 + \frac{m_e \nu_{ex}}{q u_i B_r^2}\left[E_z + \frac{1}{|q|n_e}\frac{\partial P}{\partial z}\right]}, \tag{4}$$

where we have used the relation for ion beam current $I_b = qAu_i n_e$, where u_i is the ion velocity. To further simplify this relation, we make the approximation that the electric field dominates over the pressure gradient and can be approximated as $E_z = V/L$ where V is the discharge voltage and L, the characteristic length of the acceleration zone. We similarly invoke $u_i = \sqrt{2qV/m_i}$ to find

$$\eta_b \approx \frac{1}{1 + \frac{m_e \nu_{ex}\sqrt{V}}{q\sqrt{2q/m_i}B_r^2 L}}. \tag{5}$$

To determine the actual scaling of the beam utilization, we need to know the form of the collision frequency. The exact form of this collision frequency is still an active area of research [13], but for the purpose of this motivating section, we follow Dannanmayer and Mazouffre [8] in assuming that the collision frequency scales like the electron neutral collision frequency. Classically, this yields $\nu_{ex} = n_n \sigma_{en} v_{te}$, where n_n is the neutral number density, σ_{en} is the electron neutral collision cross section, and v_{te} is electron thermal speed. We thus find

$$\eta_b \approx \frac{1}{1 + \frac{m_e n_n \sigma_{en} v_{te} \sqrt{V}}{q\sqrt{2q/m_i} B_r^2 L}}. \tag{6}$$

As a next step, we assume approximately 100% mass utilization efficiency to write $q u_n n_n = j_i$, where u_n denotes the neutral velocity. We thus find

$$\eta_b \approx \frac{1}{1 + \frac{m_e j_i \sigma_{en} v_{te} \sqrt{V}}{u_n q^2 \sqrt{2q/m_i} B_r^2 L}}. \tag{7}$$

Finally, for Hall thrusters, we have approximately [2], $T_e \approx 0.1V$ such that

$$\eta_b \approx \frac{1}{1 + \alpha j_i \frac{\sqrt{m_e} \sigma_{en} V^{3/2}}{u_n q^2 \sqrt{q/m_i} B_r^2 L}}, \tag{8}$$

where α is a constant order unity.

From this relation, we see that for fixed discharge voltage and magnetic field, the beam utilization efficiency evidently decreases. Physically, this stems from the fact that higher current corresponds to more neutrals in the channel, which in turn promotes more collisions. To this point, it has been found in recent work by Su and Jorns that beam utilization does in fact appear to decrease with higher current densities [14]. This can be seen in Fig. 2

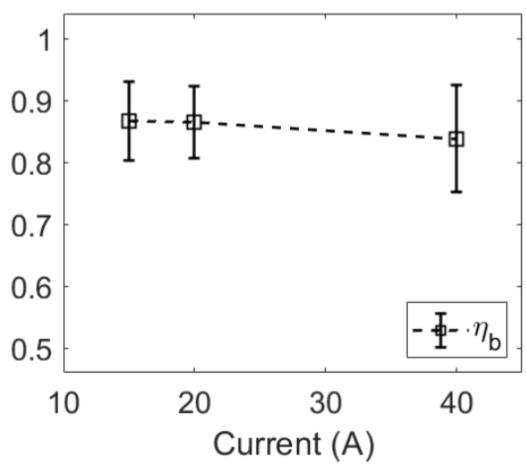

Fig. 2 Current utilization efficiency as a function of discharge current for the H9 Hall thruster [14].

While Eq. 8 shows the adverse impact of increasing current density, it also provides a possible mitigation strategy. Indeed, to counteract the enhanced cross field current, we can increase our perpendicular magnetic field, B_r. Based on this scaling law and assuming a typical value (based on SPT-100) of 150 G, if nominal is 100 mA/cm^2, then to achieve to 1000 mA/cm^2, we need a field strength of 475 G. With that said, there is a practical upper bound to achievable field strengths that stems from saturation in the ferromagnetic circuit. We discuss this in the following section.

B. The problem of magnetic saturation

Typical Hall thrusters consist of an inner and outer magnetic circuit that couple to produce a desired magnetic field (Fig. 1). The initial magnetic field generated by passing current through the inner and outer electromagnetic circuits is guided and amplified by ferromagnetic cores to produce a magnetic field. Careful placement of the ferromagnetic material, along with specific applied currents to the inner and outer circuit, allows one to tailor the magnetic field to a desired shape and strength.

Saturation occurs once all the magnetic field domains (dipole moments) in a ferromagnetic material are aligned with the externally applied field [11]. Saturation therefore limits the achievable field strengths in Hall thrusters. To illustrate this, we plot in Fig. 3 notionally the normalized peak radial field strength of a typical Hall thruster as a function of the current applied to its electronmagnetic coils.

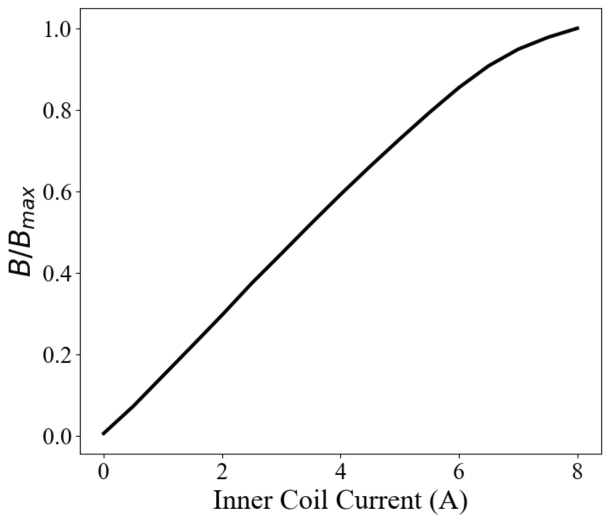

Fig. 3 Notional pot of peak radial magnetic field strength along channel center-line as a function of inner coil current.

We see that initially the relationship between the magnetic field and coil current is linear. As the coil current is increased further, the response of the magnetic field begins to asymptote. Physically, this is an indication of saturation in the ferromagnetic material.

Due to uneven magnetic flux through the ferromagnets, some areas will saturate before others, asymmetrically changing the field. To ensure a robust field design and thruster operation, we typically operate within the linear response regime for the entire operational envelope. This ultimately limits the achievable field strengths for Hall thrusters. Furthermore, a materials susceptibility to magnetic saturation increases with temperature [15]. Therefore, we may asymmetrically change the field by virtue of operating a higher powers and thus increased thermal loads.

It is clear that the ferromagnetic material is a limiting factor for magnetic field strength. Indeed, if alternate ways to produce a strong field for a shielded thruster are not found, it may severely limit efficient, high power density operation. To circumvent the limits of saturation, we explore a circuit without ferromagnets, i.e. an air-core circuit.

C. Model Definition

The ultimate goal of this approach is to design an air-core magnetic circuit (i.e. one without ferromagnetic material) that can match a given magnetic field topology in a Hall thruster channel (Fig. 1b). As our initial attempt at addressing this question, we propose a simple air core circuit design comprised of a series of idealized, infinitesimal loops of coil. This is shown notionally in Fig. 4. We then attempt to determine if there is a combination of currents in each of these coils that reproduces the same magnetic field as generated by a ferromagnetic circuit. We formulate this problem by defining a model for the radial and axial components of magnetic field in the channel produced by these series of coils at location, (r, z):

$$\vec{B}_{mod}(r, z) = \sum_{i=1}^{n} (B_{r(i)}, B_{z(i)}), \tag{9}$$

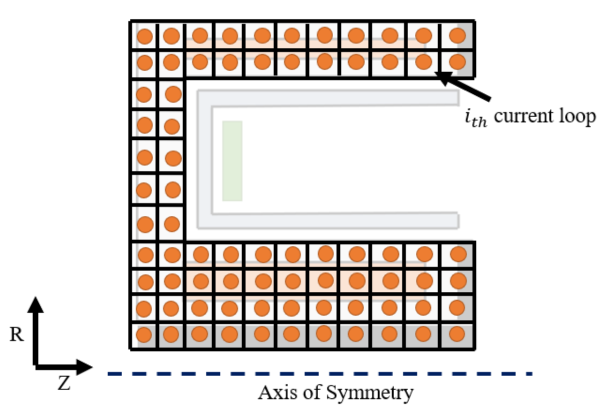

Fig. 4 Discretized thruster body with a current loop (orange) placed at the center of each cell.

where $B_{r(i)}, B_{z(i)}$ denote the radial and axial components of the magnetic field from the i^{th} coil produced at location (r, z). For an infinitesimal current loop, these are given by

$$B_{z(i)} = g(r_i, z_i)I_i, \tag{10}$$

$$B_{r(i)} = h(r_i, z_i)I_i, \tag{11}$$

where I_i is the current in the i^{th} loop and we have defined geometric factors

$$g(r_i, z_i) = \frac{\mu_0}{2\pi\sqrt{z_i^2 + (a + r_i)^2}}\left(\frac{a^2 - z_i^2 - r_i^2}{z_i^2 + (r_i - a)^2}E_2(k^2) + E_1(k^2)\right), \tag{12}$$

$$h(r_i, z_i) = \frac{\mu_0 z}{2\pi r_i\sqrt{z_i^2 + (a + r_i)^2}}\left(\frac{a^2 - z_i^2 - r_i^2}{z_i^2 + (r_i - a)^2}E_2(k^2) + E_1(k^2)\right), \tag{13}$$

where μ_0 is the permeability of free space, a_i is the radius of the current loop, r_i and z_i are the coordinates from the loop center to the point we want to compute the magnetic field, E_1 and E_2 are the elliptic integrals of the first and second kind, and k is a geometric factor defined as

$$k = \sqrt{4r_i a(z_i^2 + (a + r_i)^2)^{-1}} \tag{14}$$

[16]Note that if we fix the position of the loops and the places at which we wish to compute the magnetic field, the magnetic field is only a linear function of the current I_i. This property allows us to use linear Gaussian process regression to determine the currents needed in each loop to match a given magnetic field topography.

D. Bayesian Inference

We employ the method of Bayesian inference to find the set of values for the current loops, $\vec{\theta} = I_1, I_2, ..., I_n$, that yield the desired magnetic field configuration. To this end, we begin by discretizing this desired magnetic field in the channel as a known dataset $d = \{\vec{B}_{(msd)}(r_1, z_1), ..., \vec{B}_{(msd)}(r_j, z_j), ..., \vec{B}_{(msd)}(r_N, z_N)\}$ where B_{msd} is the measured magnetic field and B_{mod} is the magnetic field of our model. In the Bayesian approach, we represent our knowledge of the currents, θ, as probability distributions conditioned on the targeted magnetic field shape, d. This is formally stated with Bayes' rule:

$$\pi(\vec{\theta}|d) \propto L(d|\vec{\theta})f(\vec{\theta}). \tag{15}$$

Here $f(\theta)$ denotes our prior belief about the parameters; $\pi(\theta|d)$ is a posterior distribution conditioned on the measured data, d; and $L(d|\theta)$ is a likelihood function, which gives the probability of observing data, d, given our current parameters θ. Since the magnetic field model is linear in the parameters (current), we can define our model y as

$$y = \bar{A}\vec{\theta} + \zeta, \tag{16}$$

where $\bar{A}$ is a matrix of geometric factors, and ζ is our noise model which we define to be normally distributed with zero mean and variance Γ. Γ is then given by the accuracy of our Gaussmeter (1 G). Our likelihood function is then given by

$$L(d|\vec{\theta}) \propto \exp[-(y - \bar{A}\vec{\theta})^T \Gamma^{-1} (y - \bar{A}\vec{\theta})], \tag{17}$$

where the $\bar{A}$ matrix is formally defined as

$$\bar{A} = \begin{bmatrix} h_1(z_1, r_1) & h_2(z_1, r_1) & \dots & h_n(z_1, r_1) \\ h_1(z_2, r_2) & h_2(z_2, r_2) & \dots & h_n(z_2, r_2) \\ \vdots & \vdots & & \vdots \\ h_1(z_d, r_d) & h_2(z_d, r_d) & \dots & h_n(z_d, r_d) \\ \vdots & \vdots & & \vdots \\ g_1(z_1, r_1) & g_2(z_1, r_1) & \dots & g_n(z_1, r_1) \\ \vdots & \vdots & & \vdots \\ g_1(z_d, r_d) & g_2(z_d, r_d) & \dots & g_n(z_d, r_d) \end{bmatrix}, \tag{18}$$

where $h_1(z_1, r_1)$ and $g_1(z_1, r_1)$ is the radial and axial magnetic field of the first current loop at our first data point (z_1, r_1). A has dimensions $2DxN$ where D is the the length of our data and N is the number of parameters. Note that since the magnetic field has two components (B_z, B_r), we must stack them in our matrix which results in the first dimension being of length $2D$.

We assume a Gaussian prior of the form

$$f(\theta) \sim N(\mu, \Sigma), \tag{19}$$

where μ is a $N \times 1$ vector of prior means, and Σ is our $N \times N$ prior covariance. We make the assumption that the parameters are independent, which results in a diagonal prior covariance matrix with these values giving the prior variance in the parameters.

Since the prior and likelihood distributions are both Gaussian and we have a linear model, our resulting posterior distribution will also be Gaussian. When the resulting posterior is of the same form as the prior, the prior distribution is conjugate to the likelihood. Using this property, we can analytically update the prior distribution to resulting posterior distribution. These are known as update equations[17] and are defined as

$$\mu_{post} = \mu + \Sigma\bar{A}^T(\bar{A}\Sigma\bar{A}^T + \Gamma)^{-1}(\vec{d'} - \bar{A}\mu), \tag{20}$$

$$\Sigma_{post} = \Sigma - \Sigma\bar{A}^T(\bar{A}\Sigma A^T + \Gamma)^{-1}\bar{A}\Sigma, \tag{21}$$

where μ_{post} and Σ_{post} are the posterior mean and covariance, and we have defined a revised dataset $\vec{d'} = (B_{r(msd)}(r_1, z_1), ...B_{z(msd)}(r_1, z_1),$. Note that these equations provide some intuition on how data updates our belief in the parameters. Looking at Eq. 20, we add to our prior mean (μ) a factor that is scaled by the distance that our prior is from the data ($y - \bar{A}\mu$). In our case, if the prior currents in our loops closely match the magnetic field data, $y - \bar{A}\mu$ is small, and we do not change our prior mean significantly. Looking at Eq. 21 we see that we subtract a quantity from our prior covariance. This indicates that introducing data only increases our certainty in our parameters (decreases covariance). Similarly, if we are highly uncertain in the accuracy of our data (large Γ), our update will be small and our confidence in the parameters will remain largely unchanged.

IV. Case studies

Now that we have established the relevant background and framework for our analysis, we consider two case studies to demonstrate its capabilities. We use magnetic field data from two Hall thrusters, the SPT-100 and the H9 to inform our Bayesian analysis. The SPT-100 is one of the most widely flown and successful Hall thrusters to date, and the H9 is a new class of shielded Hall thruster with a particularly complex magnetic field topology.

A. H9

The H9 is a 9kW magnetically shielded Hall thruster developed jointly by the University of Michigan, JPL, and ARFL (Fig. 5a) [18]. It has been operated at discharge voltages up to 800 V, and discharge currents up to 40 A. [14]. It has a centrally mounted laB6 cathode, [19] and shares design heritage with Aerojet Rocketdyne's Advanced Electric Propulsion system (AEPS)[20]. The H9 has been operated in the laboratory setting since it was built in 2017. A Gaussmeter at the University of Michigan was used to measure the magnetic field of this thruster.

B. SPT-100

The SPT-100 is a 1.3 kW Hall thruster developed originally by FAKEL (Fig. 5b) [21]. It has been operated extensively since the 90's in both laboratory settings and in space. The magnetic field of this thruster was designed to be primarily radial, with peak center-line field strength of $\sim$ 140 G. The geometry of this thruster is summarized in table 1.

Table 1 Geometry of the SPT-100

Inner Radius	Outer Radius	Channel Depth	Width	Height
35 mm	50 mm	40 mm	85 mm	75 mm

(a) (b)

Fig. 5 The a) H9 Hall thruster and b) SPT-100 Hall thruster

V. Results

In this section, we present the results of our air-core magnetic field simulations for both the H9 and SPT-100. Optimally, for implementation purposes, we want an accurate simulation that uses the fewest amount of current loops with the smallest magnitude current in any one loop. In reality, these three ideas are highly correlated and form a trade space over which our solution must be tuned. Tuning the solution will include changing the prior covariance to limit currents or using a subset of coils rather than the full discretization. Before tuning the solution, we first show that we are able to match a set of magnetic field data by discretizing the entire thruster body with a loop of current at the center of each 5 mm $\times$ 5 mm cell.

To access the accuracy of the simulation we take the L2 norm between the measured magnetic field data (B_{msd}) and the magnetic field generated with current in each loop set to the posterior mean (B_{pred}). The L2 norm is defined as

$$L2 = (\frac{1}{2D}\sum_{i=1}^{2D}((B_{msd})_i - (B_{pred})_i)^2)^{\frac{1}{2}}. \tag{22}$$

We now compare the results of the air core simulation to the measured data for both the SPT-100 and the H9. We show the magnetic field lines in the channel, radial magnetic field along center-line, and the effect of the prior on the solution. In addition, for the SPT-100, we show a pixel map of the learned currents in each loop.

A. SPT-100

1. Magnetic Field Lines

The shape of the magnetic field in the discharge chamber is paramount to efficient thruster operation. The SPT-100 magnetic field is primarily radial and is compared to the air-core circuit in Fig. 6. Note that the anode is located at z = 0, and the exit plane is at 0.025 m. The channel walls are located at r = 0.035 m and 0.05 m. The L2 norm for this solution was 0.25 G.

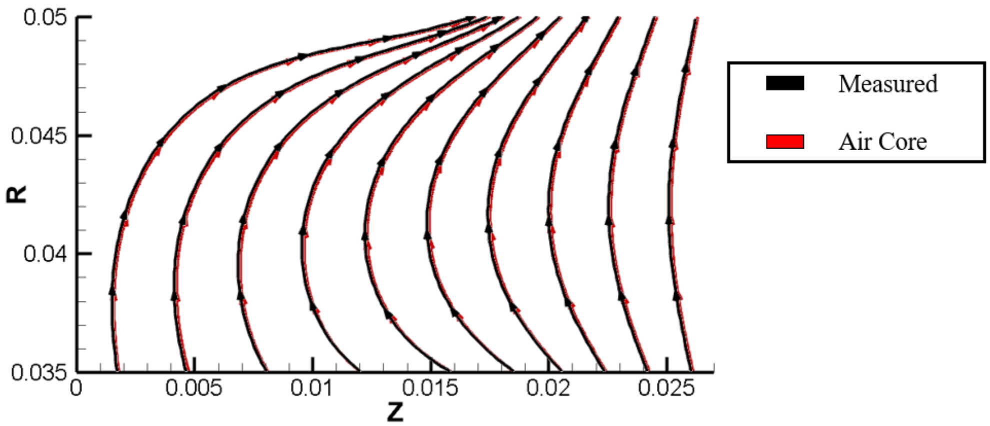

Fig. 6 Comparison between the the air core model and measurement data of magnetic field lines in the discharge chamber of the SPT-100.

Fig. 6 shows that the air-core circuit was able to match the measured magnetic field topology of the SPT-100. We see the primarily radial field lines in the channel intersect the upper and lower chamber walls. The field lines almost perfectly line up, with very minute differences. This is corroborated by the low L2 norm of our solution. While matching the shape of the field lines is important, we must also ensure that the simulation matches the required field strengths for efficient operation.

2. Radial profile

The peak radial center-line field must be sufficiently strong so that electrons remain locally confined to the acceleration region. Fig. 7 shows the center-line radial magnetic field strength for the H9 and the SPT-100 compared to that predicted by the air-core circuit.

The radial profile of the air core circuit shows very good agreement with the measured data. We see an increase in field strength from the anode, and a peak just past the exit plane. The peak field is just over 140 G as expected. Matching both the radial magnetic field strength and magnetic field topology (Fig. 6) indicates that the air core solution can fully replicate the magnetic field in the channel.

3. Maximum current

It is important to look at the maximum current in any one cell of the air-core solution. Traditional Hall thruster magnetic fields are amplified by ferromagnetic cores. Therefore an air-core design will necessarily need orders of magnitude more current to produce the same magnetic field strengths as its counterpart. We can control the magnitude of currents by changing the prior covariance. The L2 norm and max current in any one cell is plotted as a function of prior covariance for the SPT-100 in Fig. 8. Note that the solution used in the previous sections had a prior covariance of $10^8 A$.

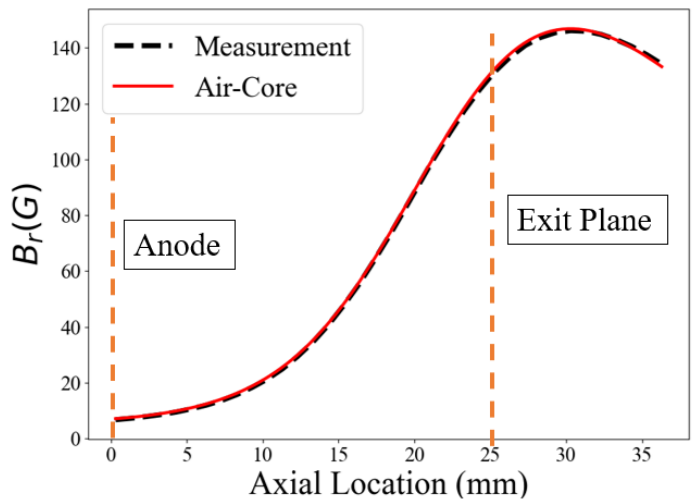

Fig. 7 Comparison between the the air core model and measurement data of the radial magnetic field strength along channel center-line of the SPT-100.

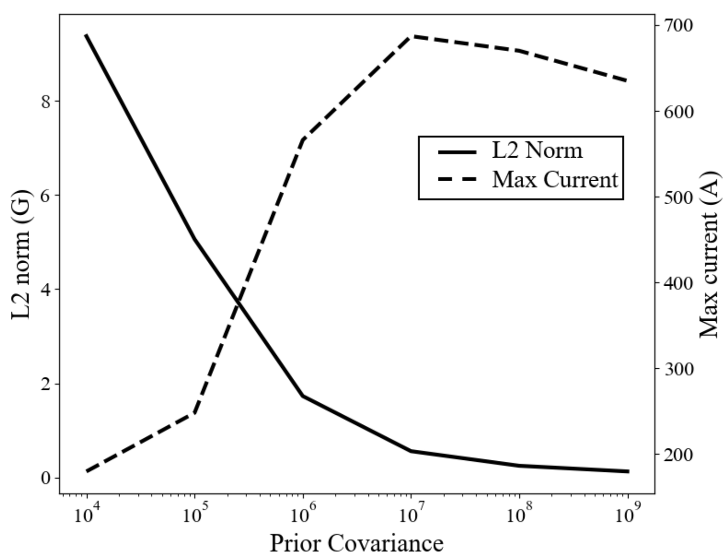

Fig. 8 L2 norm and maximum current for the air core simulation of the SPT-100 as a function of prior covariance.

As we change the prior covariance, we see an inverse relationship between the L2 norm and the maximum current. For the L2 = 0.25 solution used previously (prior covariance = 10^8), the maximum current was 670 A. As the prior covariance increases, our L2 norm decreases (accuracy increases), and the maximum current used tends to increase. This highlights a trade-off between solution accuracy, and potential machinability. Indeed, if we are unable build a circuit that can withstand the high currents, we may need to implement a less accurate solution that uses less current. This will be discussed further in Section VI.D.

The L2 norm asymptotes to a fixed solution accuracy. This is an indication that there is an inherent upper bound in achievable solution accuracy for a fixed number of current loops discretized in 5 mm by 5 mm blocks. This upper-bound corresponds to the MLE solution, which does not place bounds on the currents.

4. Current distribution

Ultimately, to make the solution more machinable, we want to reduce the amount of current loops used. To approach this problem, we need to first understand how the set of currents is distributed when the full body is discretized. A pixel map that shows the currents used in each loop to match the SPT-100 data is shown in fig. 9. The cell sizes for this mesh were 5 mm x 5 mm, which resulted in 231 current loops.

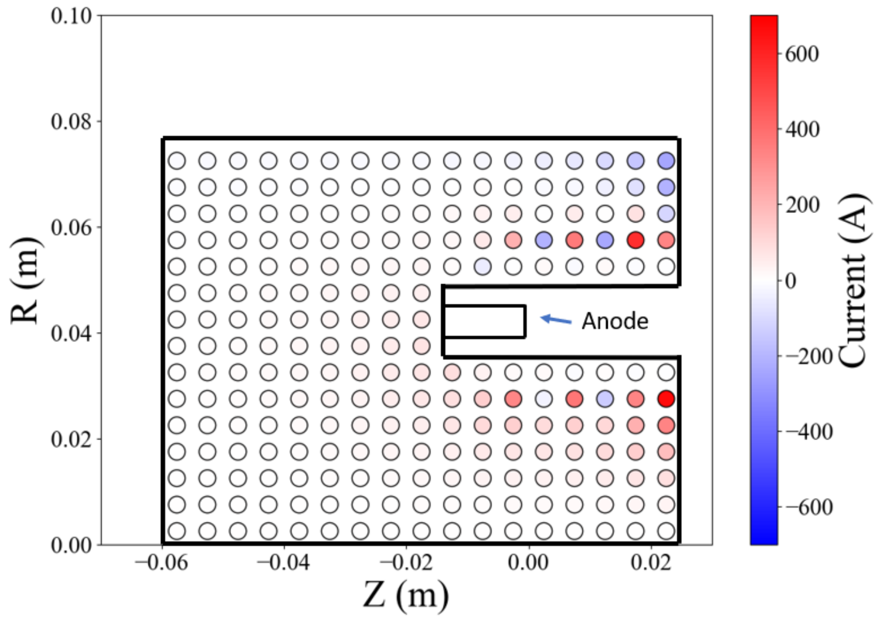

Fig. 9 A pixel map of currents used to reproduce the magnetic field topology for the SPT-100. The outline of the thruster body is shown in black, and the L2 norm was 0.25 G.

Fig. 9 highlights that the majority of current is located around the discharge chamber. Most of the current loops have less than |2| A of current in them, and therefore do not have a large impact on the magnetic field. This demonstrates that we may be able to reduce the amount of current loops we need, making our solution more practical to build. This will be discussed further in section VI.B. Fig. 9 also shows that 1-2 cells have far more current in them than the any of the other cells (~ 2x more). This information is critical for accessing practical limitations for actually constructing the circuit.

B. H9

The H9, in contrast to the SPT-100, is a magnetically shielded Hall thruster. It utilizes an inherently more complex magnetic field toplogy than the SPT-100 and may even be harder for our air-core circuit to match. The results of H9 simulations are given in the following sections.

1. Grazing Line

For a magnetically shielded Hall thruster, the grazing line is the most important magnetic field line. This is the magnetic field line that "grazes" the anode and runs parallel to the discharge chamber walls. It ensures that high potentials and low electron temperatures are maintained near the walls. The grazing line for the air-core circuit is compared to the measured grazing line of the H-9 in fig. 10. The L2 norm for this solution was 0.50 G.

The air-core circuit indeed matches the measured grazing line. The line runs parallel to the discharge chamber walls, effectively "shielding" them. This indicates that the air-core circuit in can replicate a shielded topography, thus reducing erosion and mitigating thermal loads.

2. Radial profile

The radial magnetic field in shielded Hall thrusters plays the same role as its un-shielded counterpart. It must be sufficiently strong to confine electrons, setting up a localized resistance and potential drop that accelerates ions. The

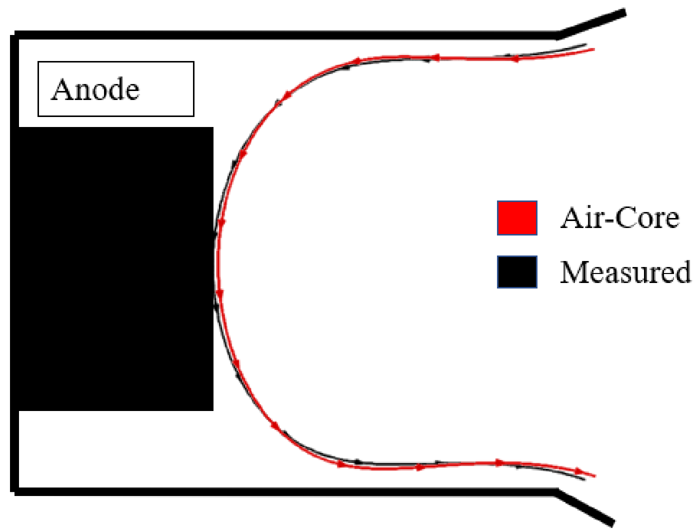

Fig. 10 H9 Grazing line for the measured magnetic field data (Black) compared to the air-core circuit (Red).

normalized radial magnetic field along center-line for the air-core circuit is compared to the measured data in Fig. 11. Note that this is the same air-core solution used for the grazing line in the previous section (L2 = 0.50)

Once again, the air-core circuit matches the measured radial profile well. We see the lowest values near the anode (z/z_{max} = 0), and a gradual increase throughout the discharge chamber. Since we matched the grazing line and the radial profile for the H9, we confirm that the air-core circuit can create a magnetically shielded topography.

3. Maximum current

For implementation purposes, we must understand what the maximum current is in any loop, and how the prior covariance controls this and our solution accuracy. A graph of the L2 norm and maximum current as a function of covariance for the H9 is given in Fig. 12

Similar to the SPT-100, we see an inverse relationship between the L2 norm and maximum current as the prior covariance increases. For the solution used in the previous sections (covariance = 10^8 A, and L2 = 0.50 G), the maximum current is 809 A. This is larger than the current used in the SPT-100 solution at the same prior covariance (670 A), even though the H9 solution is less accurate (L2 = 0.50 G vs. L2 = 0.25 G). While previously stated qualitatively, this is an indication that the magnetically shielded topography of the H9 is indeed harder to match accurately than the SPT-100. This is corroborated by the L2 norm that each solution asymptotes to (MLE solution). This corresponds to ≈ 0.07 G for the SPT-100, and ≈ 0.47 G for the H9.

Interestingly, the H9 air-core circuit was more accurate and had about the same maximum current as the SPT-100 for prior covariances between $10^4 - 10^6 A$. Since the H9 is larger than the SPT-100, there is inherently more total current loops than SPT-100 when the full body is discretized. While loops far away from the discharge chamber do not have a large impact, it is inherently easier to match a set of magnetic field data as you increase the number of current loops. It appears that, at lower prior covariances, the greater number of current loops in the H9 solution outweighed the increased difficulty in matching a magnetically shielded topography.

4. Current Distribution

Similar to the SPT-100, we found that the majority of current was distributed around the discharge chamber. 2-3 coils contained large currents (~ 800 A), with many loops have little or no current. Once again, this motivates a solution using fewer coils, the results of which will be discussed in section VI.B

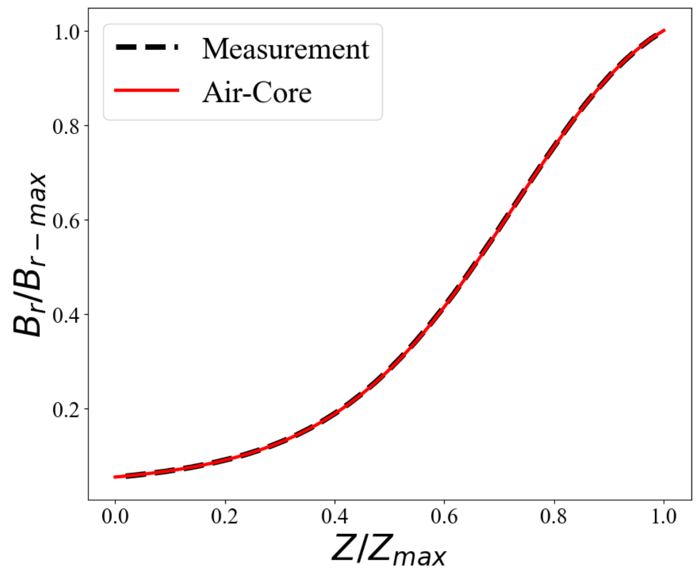

Fig. 11 Comparison between the air-core model and measurement data of the radial magnetic field along channel center-line of the H9.

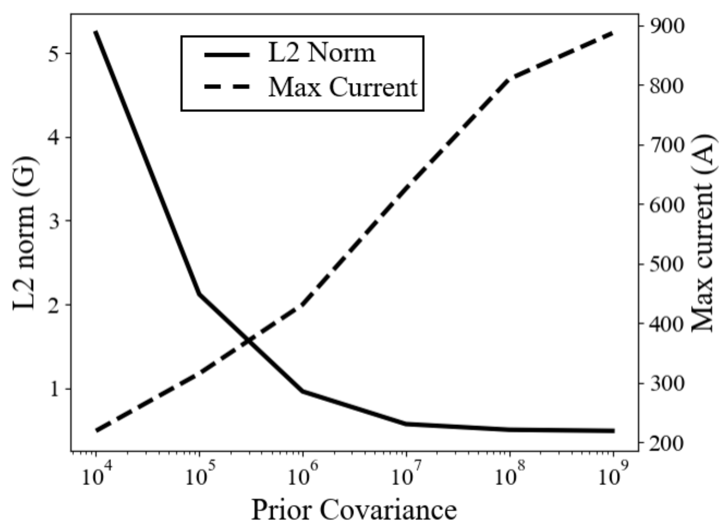

Fig. 12 L2 norm and maximum current for the air core simulation of the H9 as a function of prior covariance.

VI. Discussion

In the previous section, we showed that it is in principle possible to create a standard Hall thruster magnetic field configuration with air core circuits. In practice, however, there are a number of potential challenges with the implementation of this type of geometry. These stem from the fact that 1) the model was highly idealized and may not map to a real system with finite coil widths, 2) the air core design requires several more and independent current loops than a standard ferromagnetic based core and 3) from the fact that required currents may be prohibitively large. With this in mind, we turn in this section to a discussion of the feasibility of implementation–focusing on these issues.

A. Finite Coil Simulation

As a first concern with our solution, we note that in our model formulation, we assumed the diameter of the current loops was infinitesimally small. This invites the question, in accessing the validity of our air-core circuit, about how well loops of current approximate coils of finite size. To answer this question, we modeled our air-core circuit in COMSOL.

Unlike the air-core solution, we model each full 5 mm by 5 mm discretized block as a square current loop. While we may not machine the solution exactly in this fashion, ensuring that the solutions at least match for this case will help validate our model. We model each loop as copper, with the surrounding domain as vacuum. An adaptive mesh with 2 levels of refinement was used. We use the results of the air-core solution to determine the current needed in each finite coil. Rather than use the full discretization (Fig. 9), we use a reduced discretization to validate (Fig. 15). The reduced discretization is ultimately a more machinable solution and similarly easier to implement COMSOL. We discuss how this discretization was chosen in section VI.B.

Fig. 13 compares the magnetic field lines in the discharge chamber of the SPT-100 for both the COMSOl and air-core simulations. The center-line radial field of the SPT-100 is compared in Fig. 14.

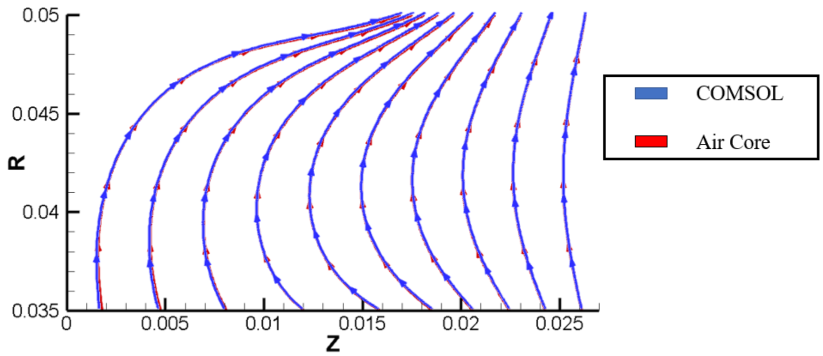

Fig. 13 Comparison of the magnetic field lines in the discharge chamber of the SPT-100 for the air core and COMSOL simulations.

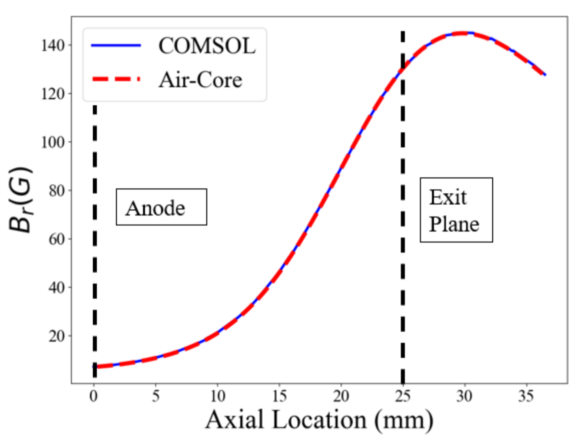

Fig. 14 Comparison of the radial magnetic field along center-line of the SPT-100 predicted by the air-core circuit and COMSOL simulation.

We see near perfect agreement for both the field line shape, and radial field strength between the COMSOL and air-core models. This validates that our model, and shows that it represents coils of finite size well. This gives us confidence that when the circuit is built the magnetic fields predicted by the air core circuit will match up with the measured field.

B. Tuning Solution for Machinability

Our ultimate goal is to accurately reproduce a shielded topography using the fewest amount of current loops with the lowest magnitude currents. We have already seen that adjusting the prior covariance can limit the maximum current used at the expense of the accuracy of the simulation. Now we aim to reduce the total amount of coils used.

The pixel map of currents used in the SPT-100 simulation shown in fig. 9 highlights that the majority of current is located around the discharge chamber. This is the result of the inverse relationship between magnetic field strength and distance from a given coil. Since the majority of currents away from the discharge channel have little current, it is likely these do not have a large impact on the final solution. We can use these results to motivate a solution which uses fewer total coils, making it easier to build.

We define a new discretization that only uses the coils in the vicinity of the discharge chamber. This geometry and the currents needed to reproduce the magnetic field of the SPT-100 is shown in fig. 15.

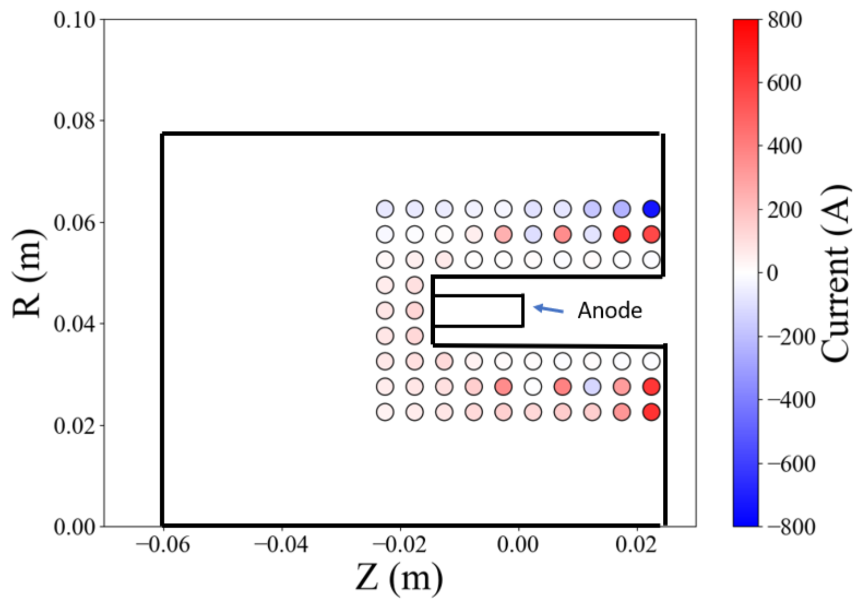

Fig. 15 Pixel map of the currents to match the SPT-100 using coils surrounding the discharge chamber.

The simulation used the same prior covariance (10^8 A) as the full discretization for comparison. Table 2 compares the results from the full discretization to the results from the reduced coil simulation.

Table 2 Comparison between full and reduced discretization for the SPT-100.

	Full Discretization	Reduced Discretization
Number of Coils	231	66
L2 Norm (G)	0.25	0.29
Max Current (A)	670	735

The reduced discretization had over 150 fewer coils than the full discretization, making it more feasible to machine and operate. The reduced coil configuration still accurately represents the magnetic field, evidenced by the small L2 norm, at the expense of an increase in current. With that said, we believe (see Sec. VI.C) that novel circuitry can be designed to handle these high currents. Therefore, the trade-off between number of current loops and maximum current used may be warranted.

We repeated this same study for the H9 where we show in Table 3 a comparison of important metrics between the full discretization and the reduced one for an initial covariance of 10^8 A.

The reduced discretization for the H9 used over 500 fewer coils than the full discretization. The L2 norm indicated that the solution was still accurate, and the maximum current actually decreased compared to the full discretization.

Table 3 Comparison between full and reduced discretization for the H9.

	Full Discretization	Reduced Discretization
Number of Coils	-	-
L2 Norm (G)	0.50	0.51
Max Current (A)	809	801

These results indicate that coils away from the discharge chamber have almost no impact on the fidelity of the simulation, and can be removed with little penalty.

Overall, we see that using current loops centered around the discharge chamber produced accurate (low L2 norm) solutions. Ultimately, due to physical constraints such as flow lines and support structures, we may not be able to place loops of current exactly in the configuration shown in Fig. 15. Therefore, our solution must be modified to accommodate these design constraints accordingly.

C. A Limit on Maximum Current

Ultimately, the currents used in the simulations may be too high for the material limits of traditional copper wires. Typical design rules place an upper bound in current density at 6 A/mm^2.[22] Based on this limit, the maximum current allowed in a 5mm $\times$ 5mm cell is 150 A, far lower than what the air-core circuit requires for an accurate solution. Similarly, the goal of our effort was to increase the magnetic field beyond its nominal operating range to increase electron confinement, requiring even more current than our simulations predict.

With this in mind, rather than use traditional copper wires to produce the magnetic field, one possible solution is to employ specially designed Bitter magnets that are capable of handling ultra-high currents. [12] Bitter magnets use sheets of copper surrounded by insulation and may be able to handle the high current loading required by an air-core circuit.

D. Relaxing requirements

If bitter magnets or other novel techniques cannot handle the current loading, we may need to relax our requirements on solution accuracy in favor of using less current. Returning to Figs. 8 and 12, we can limit the magnitude of current used by changing the prior covariance.

With that said, less accurate solutions must still uphold the key tenets of magnetic shielding to mitigate erosion and thermal loading. Magnetic shielding takes advantage of two well known properties that hold along magnetic field lines: electron isothermality and equipotentialization [2, 3]. To utilize these properties to shield the discharge chamber walls, the grazing line must extend deep in to the channel (fig. 10), where electrons are cold and potentials are high. The grazing line must also run adjacent to the channel walls in order to maintain this ideal equipotentialization and reduce the high energy ion flux and thermal loading. Therefore, when analyzing the accuracy and applicability of the different air-core configurations, we must ensure that the grazing line does not intersect channel walls.

As expected, in Fig. 10 the fully discretized accurate solution (L2 = 0.5 G) matches the measured grazing line of the H9. They both run parallel to the discharge chamber walls and therefore we expect our air-core circuit design for this solution to effectively shield the thruster. However, the maximum current in this solution is 809 A. To increase our magnetic field strength to the required value for 10$\times$ current density operation, the new maximum current for this solution will be greater than 2500 A. This ultimately may be too high for even novel circuitry to handle.

Using the prior covariance to limit the currents, we look at the air-core solution for the H9 corresponding to a prior of $10^4 A$, which has an L2 norm of 5.24 G, and a maximum current of 219 A. To analyze whether this solution achieves a shielded topography, we compare the grazing line to the measured values in Fig. 16.

Even using a less accurate solution, the grazing line of the air-core circuit matches well with the measured data. The grazing line runs parallel to the walls, and we expect this solution to effectively shield the thruster. Fig. 16 does not make it readily obvious as to why the L2 norm was higher for this solution. To see this, we plot the radial magnetic field along center line and compare to the measured data in Fig. 17.

Fig. 17 highlights that the inaccuracy of the L2 = 5.24 G solution is in large part due to an under prediction of magnetic field strength. Since this solution uses less current than our highly accurate solution, it is intuitive that the

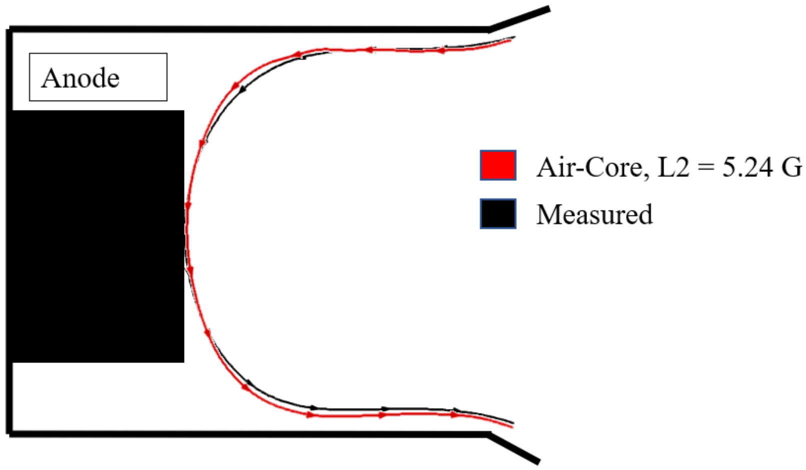

Fig. 16 H9 grazing line comparison between the air core solution (L2 = 5.24 G) and measurement data.

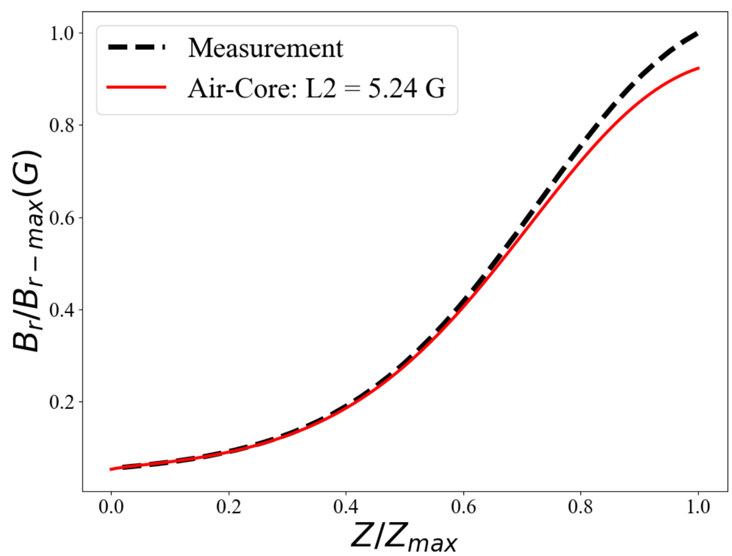

Fig. 17 H9 center line radial magnetic field comparison between the air core solution (L2 = 5.24 G) and measurement data.

magnetic field would be lower than the measured data. To remedy this, we simply scale the set of currents by a constant factor to match the desired field strength. This results in a maximum current of 243 A for this solution. Scaling the magnetic field and the set of currents by just over 3 × for high current density operation puts the new maximum current at ≈ 770 A. Comparing this to the scaled currents of the more accurate solution discussed previously (>2500 A), we see a difference of over 1750 A. This indicates that scaling a less accurate solution may be a feasible option if we cannot handle the current loading required by more accurate solutions.

VII. Conclusion

We demonstrated that an air-core circuit could replicate a given set of magnetic field data. Air-core circuits are not limited in field strengths by saturation in ferromagnetic material that hampers traditional Hall thruster circuits. Therefore, an air-core circuit can achieve the higher field strengths necessary to mitigate some performance losses at high current density operation. To design this circuit, we discretized a Hall Thruster body, and used Bayesian inference to learn the spacial distribution of current density needed to reproduce the magnetic field. Two Hall thrusters, the H9 and the SPT-100 were used as case studies to validate that this design would work for both shielded and un-shielded thrusters. Analytic calculations of the magnetic field due to a set of current loops showed that an air-core circuit can in fact reproduce the fields of these thrusters. Similarly, the solution was validated in COMSOL using coils of finite length. Ultimately, for implementation purposes, we wanted an accurate solution that minimized both the number of distinct current loops and the magnitude of current in each loop. We found that using a reduced discretization with coils centered around the discharge chamber could reproduce the fields with little loss in accuracy or gain in maximum current. However, to increase the field strengths by the necessary amount ($\approx 3.5\times$), requires an order of magnitude more current than traditional bounds of copper wire. In order to withstand this demanding current flow, advanced electromagnetic circuitry is required. One potential path forward is the use of helical resistive plate electromagnets known as Bitter magnets. Despite the large thermal and electrical challenges associated with recreating existing field topographies to high precision, our solution space demonstrates that relaxed geometries which still provide magnetically shielded operation can be achieved with more practical electrical loads.

References

[1] "Space Nuclear Propulsion for Human Mars Exploration," Tech. rep., National Academies of Sciences, Engineering, and Medicine, 2021.

[2] Goebel, D., and Katz, I., *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*, John Wiley and Sons, 2008.

[3] Mikellides, I., Katz, I., Hofer, R., and Goebel, D., "Magnetic Shielding of a Laboratory Hall thruster," *Journal of Applied Physics*, Vol. 115, No. 043303, 2014. https://doi.org/10.1063/1.4862313.

[4] Kim, V., and Popov, G., "Electric Propulsion Activity in Russia," *International Electric Propulsion Conference*, 2005.

[5] Hall, S. J., "Characterization of a 100-kW Class Nested-Channel Hall Thruster," Ph.D. thesis, University of Michigan, Ann Arbor, MI, 2018.

[6] Dankanich, J., Vondra, B., and Ilin, A. V., "Fast transits to Mars using electric propulsion," *46th AIAA/ASME/SAE/ASEE Joint Propulsion Confernce and Exhibit*, 2010.

[7] Hofer, R., and Randolph, T., "Mass and cost model for selecting thruster size in electric propulsion systems," *Journal of Propulsion and Power*, Vol. 29, No. 1, 2013, pp. 166–177.

[8] Dannenmayer, K., and Mazouffre, S., "Elementary Scaling Relations for Hall Effect Thrusters," *Journal of Propulsion and Power*, Vol. 27, No. 1, 2011, pp. 236–245. https://doi.org/10.1051/eucass/201102601.

[9] Mikellides, I. G., Katz, I., Hofer, R. R., Goebel, D. M., de Grys, K., and Mathers, A., "Magnetic Sheilding of the channel walls in a Hall plasma accelerator," *Physics of Plasmas*, Vol. 18, No. 033501, 2011. https://doi.org/10.1063/1.3551583.

[10] Grys, K. D., Mathers, A., Welander, B., and Khayms, V., "Demonstration of 10,400 Hours of Operation on a 4.5 kW Qualification Model Hall Thruster," *46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, 2010.

[11] Griffiths, D. J., *Introduction to Electrodynamics: Fourth Edition*, Pearson Education, Glennview, IL, 2013.

[12] Bird, M., and Toth, D., "Design of the next generation of Florida-Bitter magnets at the NHMFL," *IEEE Transactions on Applied Superconductivity*, 2004.

[13] Boeuf, J. P., "Tutorial: Physics and Modeling of Hall Thrusters," *Journal of Applied Physics*, Vol. 121, No. 1, 2017.

[14] Su, L. L., and Jorns, B. A., "Performance at High Current Densities of a Magnetically Shielded Hall Thruster," *AIAA Propulsion and Energy Forum*, 2021.

[15] Luborsky, F. E., Walter, J. L., Liebermann, H. H., and Wohlfarth, E. P., "The Effect of Temperature on Magnetic Saturation of Some Amorhous Alloys," *Journal of Magnetism and Materials*, Vol. 15-18, 1980, pp. 1351–1354.

[16] Simpson, J., Lane, J., Immer, C., and Youngquist, R., "Simple Analytic Expressions for the Magnetic Field of a Circular Current Loop," 2013.

[17] Bishop, C., *Pattern Recognition and Machine Learning*, Springer, 2006.

[18] Cusson, S., Hofer, R., and Lobbia, R., "Performance of the H9 Magnetically Sheilded Hall Thrusters," International Electric Propulsion Conference, Atlanta, Georgia, 2017.

[19] Hofer, R., Goebel, D., and Walker, M., "Compact Lab6 Cathode Hollow Cathode for a 6 kW Hall Thruster," *54th JANNAF Propulsion Meeting*, 2018.

[20] Jackson, J., Miller, S., Cassady, J., Soendker, E., Welander, B., and Barber, M., "13 kW Advanced Electric Propulsion Flight System Development and Qualification," 36th International Electric Propulsion Conference, Vienna, Austria, 2019.

[21] Sankovic, J., Hamley, J., and Haag, T., "Performance Evaluation of the Russian SPT-100 Thruster at NASA LeRC," *23rd International Electric Propulsion Conference*, 1993.

[22] "Current Carrying Capacity of Copper Conductors," https://www.multicable.com/resources/reference-data/current-carrying-capacity-of-copper-conductors/, 2017.