Simulation of Vacuum Chamber Pressure Distribution with Surrogate Modeling and Uncertainty Quantification

Charles P. Lipscomb*, Iain D. Boyd†
University of Colorado Boulder, Boulder, CO, 80305

Kaelan B. Hansson ‡
AMA Inc at NASA Ames Research Center, Mountain View, CA, 94035, USA

Joshua D. Eckels§, Alex A. Gorodetsky¶
University of Michigan, Ann Arbor, MI, 48109

A major challenge in understanding differences in electric propulsion performance in ground tests and in space operations concerns the pressure distribution within the test vacuum chamber. The chamber backpressure is much higher than experienced in space, modifying thruster performance and plume dynamics. Numerical simulation is a key element to determining the background conditions in non-ideal vacuum chamber environments. An important parameter for the accurate simulation of chamber backpressure is the probability that an atom will stick to a cryogenic panel used to pump away the plume gases. This quantity can be used to model vacuum pumps in particle-based kinetic numerical methods. In this work, a three-dimensional direct simulation Monte Carlo code is used to model neutral xenon atoms flowing from the anode of the H9 Hall Effect Thruster within the University of Michigan's Large Vacuum Test Facility. Simulated pressures are compared with ion gauge pressure measurements to infer the effective sticking coefficient of the chamber's vacuum pumps. A pressure predicting surrogate model is developed for inference of pump sticking coefficients and for uncertainty quantification. This information enables accurate and useful kinetic simulations of electric propulsion thruster plasma plumes in vacuum chambers.

I. Nomenclature

Kn	=	Knudsen number
λ	=	Mean free path (m)
L_C	=	Characteristic length of the system (m)
$\vec{F}$	=	Force (N)
$\vec{v}$	=	Particle velocity (m/s)
N_{pairs}	=	Total particle pairs to be assessed for collision
N_P	=	Number of macroparticles
n	=	Number density (m^{-3})
d	=	Collision diameter (m)
σ	=	Collisional cross section (m^2)
g	=	Relative speed (m/s)
Δt	=	Simulation timestep (s)
k_B	=	Boltzmann constant (J/K)
m	=	Mass (kg)
ν	=	Collision frequency (s^{-1})

---

[*] Ph.D. Student, Ann and H.J. Smead Department of Aerospace Engineering Sciences, charles.lipscomb@colorado.edu

[†] H. T. Sears Memorial Professor, Ann and H.J. Smead Department of Aerospace Engineering Sciences, AIAA Fellow, iain.boyd@colorado.edu

[‡] Research Scientist, Aerothermodynamics Branch.

[§] Ph.D. Candidate, Department of Aerospace Engineering, Student Member AIAA

[¶] Assistant Professor, Department of Aerospace Engineering, Member AIAA

ref = Reference value
$\mathcal{U}$ = Uniform distribution
$\mathcal{N}$ = Normal distribution

II. Introduction

Electric propulsion (EP) thrusters have become widely used for station-keeping and orbit-raising due to their high specific impulse, longevity, and reliability [1]. Hall effect thrusters (HETs) are particularly attractive due to the relative simplicity of their design and their high thrust-to-power ratios [2]. The higher specific impulse of electric propulsion over chemical propulsion allows mission designers to allocate less mass for on-board propellant, reducing launch costs and/or increasing the payload mass allowance. The aerospace community is extending the application of EP systems into the domain of orbital transfers. Opting to use electric propulsion for a transfer decreases the required mass of propellant, but increases the time to complete the transfer. For some transfers, minimizing the duration is of critical importance. During the common transfer from LEO to GEO, satellites enter the Van-Allen radiation belts. Prolonged exposure to the radiation belts can cause mission threatening damage to solar arrays and payload electronics [3]. Short transfer times are also essential for enabling human spaceflight.

High-power EP (~100 kW) will enable expeditious and fuel efficient space travel. For the proliferation and flight of high-power EP, thrusters must be tested in vacuum chambers on the surface of the Earth because testing in space is prohibitively expensive. As the power of an EP system being tested in a vacuum chamber is increased, the mass flow rate of the thruster also increases, and the ability of the vacuum chamber to emulate the space environment is decreased [4] [5]. Since facility pumping capacity is inherently limited by chamber surface area constraints, a greater mass flow rate will lead to an increase in facility backpressure. Elevated background pressures alter the amount of gas ingested by the thruster, the production of charge-exchange ions, and the divergence of the plume. It is imperative that ground-based tests be accompanied by predictive models to confidently extrapolate thruster performance and lifetime to in-space operation.

An accurate model of facility vacuum pumps is needed to calculate the spatially varying chamber backpressure. Vacuum pumps operate at extremely cold temperatures and remove gas by freezing it onto pump surfaces. Pumps can be effectively modeled by assigning a sticking coefficient to pump surfaces. The sticking coefficient sets the fraction of particle-pump interactions that result in sticking versus reflection. This work adopts an approach (similar to that developed in [6]) for inferring the pump sticking coefficient by comparing simulated pressure distributions with pressure measurements.

Since the employed gas dynamics model relies on uncertain inputs, the output of the model is also uncertain. Uncertainty quantification applied to quantities like chamber backpressure is essential for confidently extrapolating thruster performance and lifetime from vacuum chamber tests to in-space operation. The remainder of this manuscript presents the experimental setup in Section III, the numerical methods in Section IV, the results in Section V, and the conclusions in Section VI.

III. Experimental Overview

A. Vacuum Facility

All experiments are performed in the University of Michigan's Large Vacuum Test Facility (LVTF). The cylindrical chamber is 9 meters long and 6 meters in diameter. The LVTF is equipped with two types of vacuum pumps. The first type of pump, referred to as a cryopump, is the PHPK-TM1200i re-entrant vacuum pump which has a reported pumping speed of 35,000 l/s for xenon. The LVTF is populated by 13 cryopumps each of which is housed in a baffled, liquid nitrogen cooled casing. The second type of pump, referred to as cryosails, were developed by the Plasmadynamics and Electric Propulsion Laboratory (PEPL) at the University of Michigan. They are liquid nitrogen free cryogenic pumps operating with a theoretical pumping speed of 39,600 l/s for xenon [7]. At the time of these experiments, there were a total of five cryosails within the LVTF. Figure 1 shows a picture of the interior of the LVTF. The baffled structures in Fig. 1 are the cryopumps and the octagonal surfaces are the cryosails.

Two sets of experimental pressure measurements are collected in this work. One set comes from experiments using a high pressure configuration and the other set uses a low pressure configuration. The high pressure configuration has three cryopumps active along the top of the chamber as well as two inactive pumps with active liquid nitrogen shrouds. The low pressure configuration uses 13 cryopumps and four cryosails. Figure 2 shows a rendering of each configuration.

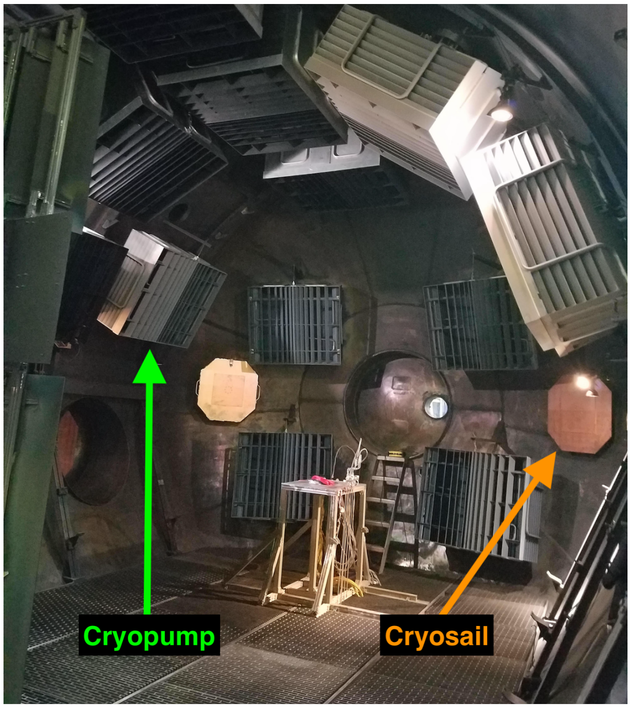

Figure 1. The interior of the LVTF equipped with octagonal, PEPL designed cryosails and baffled PHPK-TM1200i cryopumps.

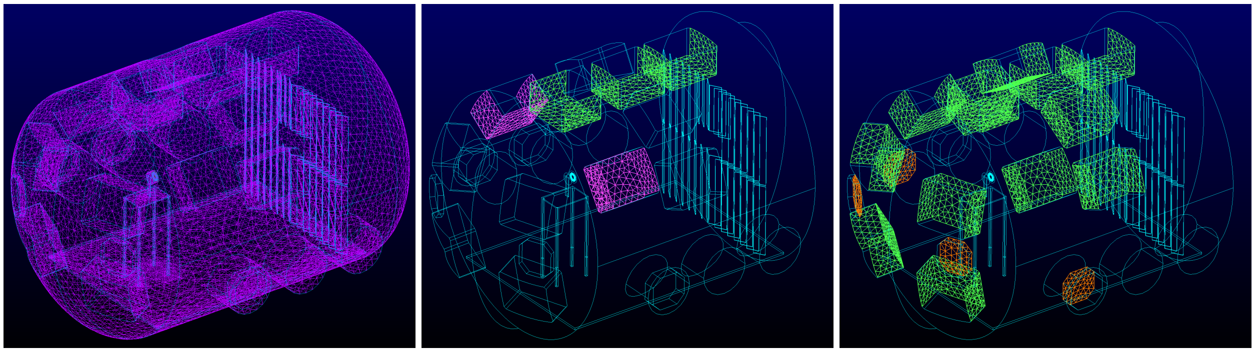

Figure 2. A computational representation of the LVTF (left), the high pressure configuration (center) with cryopumps highlighted in green and inactive pumps with liquid nitrogen shrouds in pink, and the low pressure configuration (right) with cryopumps highlighted in green and cryosails in orange.

B. Ionization Gauges

Two Granville-Phillips 370 hot-cathode Bayard-Alpert ionization gauges running on a 370 series vacuum gauge controller are used for measuring chamber pressure. These gauges are capable of measuring pressures in the range of 10^{-10} to 10^{3} Torr. The accuracy of the gauges is taken to be ±10% based on heritage data.

One ionization gauge, referred to as Gauge 1, faces downstream (the direction in which the thruster expels gas) and is located 1 m away from the thruster in the thruster exit plane (as is recommended in [8]). The other ionization gauge, referred to as Gauge 2, faces away from the thruster and is located 1 m behind the thruster (the front of the thruster being defined as the side with the discharge channel which points towards the beam trap). Figure 3 shows the approximate locations of the ionization gauges in the LVTF. Pressure measurements from both gauges are calibrated for nitrogen gas by the manufacturer and are corrected to xenon.

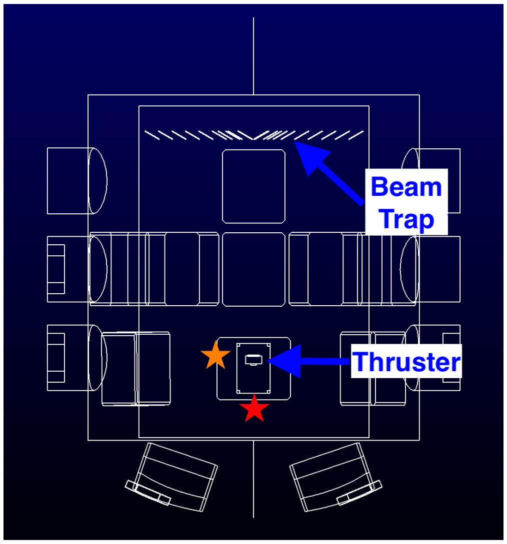

Figure 3. Top-down view of the LVTF showing ionization gauge locations with Gauge 1's location being indicated by the orange star and Gauge 2's by the red star. Thruster and beam trap locations are pointed to with arrows to provide orienting information.

C. Hall Thruster

The thruster used in all experiments is the magnetically shielded 9-kW H9 HET. The centrally mounted lanthanum hexaboride cathode, the anode/gas distributor, and the discharge chamber geometry of the H9 were all inherited from the unshielded 6-kW H6 HET. More information about the H9 may be found in [9]. In all experiments, neutral xenon atoms flow through the anode of the H9 without a plasma discharge. Mass flow rates and the corresponding pressure measurements from both ionization gauges are reported in Tables 1 and 2 for the high and low pressure configurations, respectively.

Table 1 Data from high pressure experiments of cold flow through the H9 HET.

Mass Flow Rate (sccm)	Gauge 1 Pressure (µTorr)	Gauge 2 Pressure (µTorr)
100	11.6	10.6
200	19.9	18.9
300	25.9	26.5
400	31.6	33.9

Table 2 Data from low pressure experiments of cold flow through the H9 HET.

Mass Flow Rate (sccm)	Gauge 1 Pressure (µTorr)	Gauge 2 Pressure (µTorr)
250	5.2	4.0
300	6.1	4.6
350	6.9	5.3
400	7.6	6.0

IV. Numerical Methods

A. Physics Based Modeling

The modeling approach to simulate dilute gas flow ejected from the thruster within the vacuum chamber is to use a direct simulation Monte Carlo (DSMC) code. The code, known as MONACO, is a parallelized 3D DSMC code that accommodates unstructured meshes. MONACO is provided with static boundary conditions and returns time-averaged steady state results. Kinetic descriptions of gases, the DSMC method, and the boundary conditions used in this model are discussed in the following three subsections.

1. Kinetic Modeling of Gas Dynamics

The extent of rarefaction of a gas flow may be quantified by the following non-dimensional parameter known as the Knudsen number:

$$Kn = \frac{\lambda}{L_C} \tag{1}$$

where λ is the mean free path of the gas and L_C is a characteristic length of the system. When $Kn < 0.01$, the continuum assumption is valid and fluid conservation equations accurately describe the system. The case when $Kn > 1$ is known as free molecular flow. In this regime, a kinetic description of the flow that accounts for the molecular behavior of the gas is needed. A kinetic description is necessary for large Knusden number, nonequilibrium flows in which the lack of collisions allows for the presence of non-Maxwellian velocity distribution functions. The transitional regime, where typical EP plumes reside [10], is characterized by $0.01 < Kn < 1$. In this regime, collisions are important, but insufficient for bringing the system into equilibrium. When modeling transitional gases, it is necessary to use a kinetic approach.

To accurately resolve the distribution functions within a rarefied gas, kinetic methods that capture the physics of the Boltzmann equation are required [11]. This work employs a particle-based kinetic approach in which the motion of individual macroparticles which represent a much larger number of real particles are tracked through the computational domain. The transport of neutrals and the collision dynamics thereof are simulated using the DSMC method [12].

2. Direct Simulation Monte Carlo

The DSMC method takes advantage of three physical properties of a dilute gas [11]:

1) Molecules move in free flight without interaction for time scales on the order of the local mean collision time.
2) The impact parameters and initial orientations of colliding molecules are random.
3) There are an enormous number of real molecules per cubic mean free path but only a small fraction of model particles need to be simulated to obtain an accurate molecular description of the flow.

The first property allows for the separation of particle translation and particle-particle collisions which reduces computational cost and simplifies the modeling of both processes. Assuming random orientations of colliding molecules reaps significant savings over deterministic collisional algorithms such as those used in molecular dynamics simulations. This is an excellent assumption in most circumstances as there is no physical mechanism to cause an inherent bias in the parameters that define the initial conditions of a molecular collision. The assumption that only a small fraction of molecules need to be simulated to accurately describe a dilute gas is statistically justified. Highly accurate distribution functions can be obtained with relatively low computational expense with this method.

Leveraging the properties listed above is what allows the DSMC method to model dilute gases using feasible computational resources. To further decrease the wall clock simulation duration, MONACO is highly parallelized, allowing it to be run in a distributed sense across many CPUs on a supercomputer. MONACO is equipped to simulate flows on unstructured computational grids, enabling the simulation of gas flows within complicated geometries.

Collisions in the DSMC method are performed stochastically. The No-Time-Counter scheme [12] is used to calculate the total number of potential collision pairs within each computational cell:

$$N_{pairs} = \frac{1}{2} N_P n (\sigma g)_{max} \Delta t \tag{2}$$

where N_p is the total number of macroparticles, n is the number density, $(\sigma g)_{max}$ is an estimate of the maximum value of the product of the collisional cross section with the relative speed of the colliding particles, and Δt is the timestep. Whether or not a given pair collides is decided by comparing the collision probability with a random number. The collision probability is calculated as the ratio of σg to $(\sigma g)_{max}$. Various collision models yield different values of σ.

MONACO utilizes the variable hard-sphere model [13] to obtain momentum exchange collision cross sections. In this model, the cross section is a function of relative speed:

$$\sigma = \sigma_{ref} \left(\frac{g_{ref}}{g} \right)^{-2\omega} \tag{3}$$

where g is the relative speed and the power law exponent, ω, is a fitting parameter related to the gas viscosity.

3. Boundary Conditions

At the thruster exit plane, the number density, velocity, and temperature of xenon atoms are prescribed. The atoms are assumed to be moving at the speed of sound with a temperature of 300K. The number density is obtained from the following mass flow rate equation:

$$\dot{m} = Amnv \tag{4}$$

where $\dot{m}$ is the mass flow rate, m is the propellant molecular mass, A is the area of the thruster exit plane, n is the number density, and v is the velocity. For the number density calculation, the plume is assumed to be unidirectional. Particle-surface interactions are modeled with an accommodation coefficient that sets the probability that a particle interacting with a solid wall will reflect diffusely rather than specularly. An accommodation coefficient of 0.9 means that 90% of particle reflections will be diffuse and 10% will be specular. Walls are set to 300K in the simulations. The vacuum pumps are modeled using a sticking coefficient that sets the fraction of pump interactions that result in removal of an incident atom. Cryopumps are nominally set to 85K and cryosails are set to 40K.

B. Uncertainty quantification

Uncertainty arises in the DSMC evaluation of vacuum chamber pressure due to imprecise knowledge of the experimental conditions, the model fit coefficients, and natural variability in the system (i.e. measurement uncertainty, numerical tolerance, etc.). Uncertainty in the model inputs confounds understanding of the model outputs, and so this work seeks to quantify output uncertainty by propagating input uncertainty through the DSMC code. Input uncertainty is represented by assigning a probability distribution function (PDF) to each input that encodes a prior belief in their value and degree of variability (i.e. wider distributions encode a greater degree of uncertainty in the value of the inputs).

Table 3 summarizes the model input uncertainties for the DSMC pressure predictions. Each input is assigned a uniform prior distribution between conservative upper and lower bounds. The uncertainty in the collision diameter comes from low temperature, low pressure xenon viscosity data [14]. The mass flow rate is controlled by a mass flow controller with ±1% accuracy. The wall accommodation coefficient uncertainty is reported in [15]. Appropriate bounds for the pump sticking coefficients are informed by [6]. The baffled cryopumps are enshrouded by a liquid nitrogen (LN2) casing at 77K. While the inner, helium gas-cooled sticking surface is at about 15K, the outer casing will be much closer to the temperature of the LN2. Particles that pass within the baffle system are extremely likely to stick to the inner surface. Therefore, any particles that contact the pump and return to the chamber will likely be thermalized to the temperature of the LN2 casing. A range of 75–95K is chosen in the absence of temperature data on the outer casing. Two types of pumps are involved in this analysis. DT-670 temperature diodes measure the surface temperature of the cryosails which have no casing. Since these pumps account for much less surface area than the cryopumps and their temperatures are known within 1K, the uncertainty in their surface temperature is ignored in this analysis.

Table 3 Summary of input uncertainty for the DSMC pressure model.

Parameter	Symbol	Units	Distribution	Domain
Collision diameter	d	nm	$\mathcal{U}(0.564, 0.584)$	$[0.564, 0.584]$
Mass flow rate	$\dot{m}$	sccm	$\mathcal{U}(\pm 1\%)$	$[100, 400]$
Wall accommodation coefficient	α	-	$\mathcal{U}(0.9, 1)$	$[0.9, 1]$
Cryopump surface temperature	T_p	K	$\mathcal{U}(75, 95)$	$[75, 95]$
Cryopump sticking coefficient	θ_{cp}	-	$\mathcal{U}(0.2, 0.6)$	$[0.2, 0.6]$
Cryosail sticking coefficient	θ_{cs}	-	$\mathcal{U}(0.2, 1)$	$[0.2, 1]$

The goal of this work is two-fold: to infer the value of the pump sticking coefficients from experimental pressure data and to understand the effects of the input uncertainty on the model outputs. These tasks are respectively referred to as Bayesian inference and forward uncertainty quantification (UQ). Since both require many forward evaluations of the model (on the order of thousands to millions), it is infeasible to use the full DSMC code directly. Each DSMC run takes approximately 5 hours on 72 processors. Instead, a surrogate model is constructed to learn the input-output behavior of the DSMC pressure model. Once constructed during an offline training phase (using a total of 312 DSMC simulations evenly distributed over the input domain in Table 3), the surrogate model can then be used in place of the full DSMC model for UQ at comparatively negligible computational cost. A simple linear surrogate model with 3rd order polynomial features is used in this work to map the inputs in Table 3 to the DSMC pressure predictions at two ion gauge locations in the chamber. The coefficients of the linear model are learned by linear least squares regression with L_2 regularization.

Bayesian inference: This work seeks to estimate the value of a pump sticking coefficient θ given experimental ion gauge measurements at two locations in the chamber. In light of the uncertainties present in the system, the experimental data is used to update the prior belief in the distribution of the sticking coefficient. In terms of Bayesian inference, we determine a posterior distribution $p(\theta|y)$ given the experimental data y. The posterior distribution is given by Baye's rule:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}, \tag{5}$$

where $p(y|\theta)$ is the likelihood of observing the experimental data for a given sticking coefficient, $p(\theta)$ is the prior distribution on the sticking coefficient, and $p(y)$ is the evidence term. All other uncertainties present in the model (Table 3) are grouped into a single term $\phi = [d, \dot{m}, \alpha, T_p]$, and $p(\phi)$ is the associated PDF. The data can be treated as the result of a forward model prediction $f(\theta, \phi)$ and additive Gaussian noise ξ:

$$y = f(\theta, \phi) + \xi, \quad \text{where } \xi \sim \mathcal{N}(0, \psi^2), \tag{6}$$

for which the likelihood is available as

$$p(y|\theta) = \int p(y|\theta, \phi)p(\phi)\mathrm{d}\phi, \tag{7}$$

and the conditional likelihood can be evaluated directly as

$$p(y|\theta, \phi) = \frac{1}{\psi\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{y - f(\theta, \phi)}{\psi}\right)^2\right]. \tag{8}$$

Eq. (8) is recognized as the Gaussian PDF centered at the model evaluation $f(\theta, \phi)$ and scaled by the experimental noise ψ (this form results from the typically reasonable noise model in Eq. (6)). This work assumes a 10% experimental noise for the ion gauge measurements, that is $2\psi = 0.1y$. The presence of uncertainty in ϕ necessitates the integration in Eq. (7), since we are only interested in inferring the value of θ. The integration is performed by the Monte Carlo estimate with M samples:

$$p(y|\theta) \approx \frac{1}{M}\sum_{i=1}^{M} p(y|\theta, \phi_i), \quad \phi_i \sim p(\phi), \tag{9}$$

where each ϕ_i is sampled from the distributions in Table 3. The evidence term $p(y)$ is similarly evaluated as a marginalization over θ:

$$p(y) = \int p(y|\theta)p(\theta)\mathrm{d}\theta. \tag{10}$$

The sticking coefficient θ is a single variable of integration (whereas ϕ was a total of four), so this integration is easily computed numerically over the bounds of $p(\theta)$ with Simpson's rule. It is evident from Eq. (9) that many evaluations of the model $f(\theta, \phi)$ are necessary to obtain the posterior $p(\theta|y)$, so it is infeasible to use the DSMC code directly; instead the surrogate linear model is used in place.

Forward UQ: After inferring the posterior distribution of the sticking coefficient $p(\theta|y)$, this work then studies the impact of the input uncertainties on the pressure predictions using the sampling-based Monte Carlo method. For a given mass flow rate, a set of N input samples are drawn: $\{\theta_i\}_{i=1}^N \sim p(\theta|y)$ and $\{\phi_i\}_{i=1}^N \sim p(\phi)$, and then propagated through the surrogate model to obtain the pressure predictions $y_i = f(\theta_i, \phi_i)$. Statistics of the output distribution $p(y)$ can then be estimated using Monte Carlo estimators, e.g. the mean via $\bar{y} \approx \frac{1}{N}\sum_{i=1}^N y_i$. This method allows plotting uncertainty bounds when comparing model predictions to experimental data.

This work also employs the global, variance-based Sobol' method for sensitivity analysis to quantify how uncertainty in the output arises from uncertainty in the inputs [16, 17]. The Sobol' method decomposes the output variance into unique contributions from individual inputs, as well as contributions from interactions between inputs. The quantities of interest in the Sobol' method are the first-order Sobol' indices: $S_i = V_i/V(y)$, where V_i is the partial variance due to the i-th input x_i, and $V(y)$ is the total observed variance in the output. The partial variances are computed by

$$V_i = V_{x_i}\left(\mathbb{E}_{x_{\sim i}}[y|x_i]\right), \tag{11}$$

where the inner expectation is taken of the output y for a fixed input x_i over all possible values of the other inputs $x_{\sim i}$. The outer variance is then taken over all possible values of x_i. This work also considers higher-order and total-order indices using numerical estimators found in [18, 19]. The results of this global sensitivity analysis not only indicate inputs with greater impact on output uncertainty, but also the relative magnitude of their importance compared to other inputs.

V. Results and Discussion

A. Overview

A picture of the 3D unstructured grid used by MONACO is shown in the left panel of Fig. 2. The mesh is comprised of approximately 130,000 cells and contains detailed information regarding the location and size of the thruster, the chamber walls, the vacuum pumps, the beam trap, and the floor. Each DSMC simulation takes approximately 5 hours to run on 72 processors. The DSMC simulations contain between 3 and 4 million particles at steady state. Simulations take about 40,000 iterations to reach steady state with a timestep of 10^{-4} seconds. The mean collision time per particle is on the order of 10^{-2} seconds. Once simulations reach steady state, samples are taken every timestep for 100,000 iterations to obtain meaningful statistical results [20]. The mean free path of the gas is on the order of 1 m and the edges of the computational cells don't exceed 0.3 m. The high pressure case is analyzed first in order to obtain the sticking coefficient of the cryopumps which is subsequently applied to the low pressure case for inference of the sticking coefficient of the cryosails.

B. High Pressure Pump Configuration

In this configuration only the cryopumps are operating, the cryosails are inactive. Representative simulated pressure distributions for cold flow through the H9 HET operating in the LVTF using the high pressure configuration are presented in Fig. 4. These results show how chamber pressure is reduced as the pump sticking coefficient is increased.

Figure 5 shows how pressure at the two gauge locations changes with mass flow rate. The plots present measurements and simulation results at each gauge location. The simulations are run using sticking coefficients of 0.2, 0.4, and 0.6. From the two plots, the sticking coefficient of the cryopumps is inferred to lie between 0.2 and 0.4. Experimental uncertainties are shown in Fig. 5. To quantify the uncertainty in simulation results, a surrogate model is developed. The surrogate model will also narrow down a best fit value for the sticking coefficient using a thorough Bayesian calibration.

The high pressure surrogate model is used to numerically obtain the posterior distribution $p(\theta_{cp}|y)$ as given by Baye's rule in Eq. (5). A Gaussian distribution is fit to the numerical PDF and is shown in Fig. 6. The mean and standard deviation of the Gaussian are 0.258 and 0.008, respectively (that is, $p(\theta_{cp}|y) \approx \mathcal{N}(0.258, 0.008)$). The most likely value of $\theta_{cp} = 0.258$ matches the intuition from Fig. 5 that the best-fit sticking coefficient is between 0.2 and 0.4. The relatively small standard deviation indicates high confidence in this value.

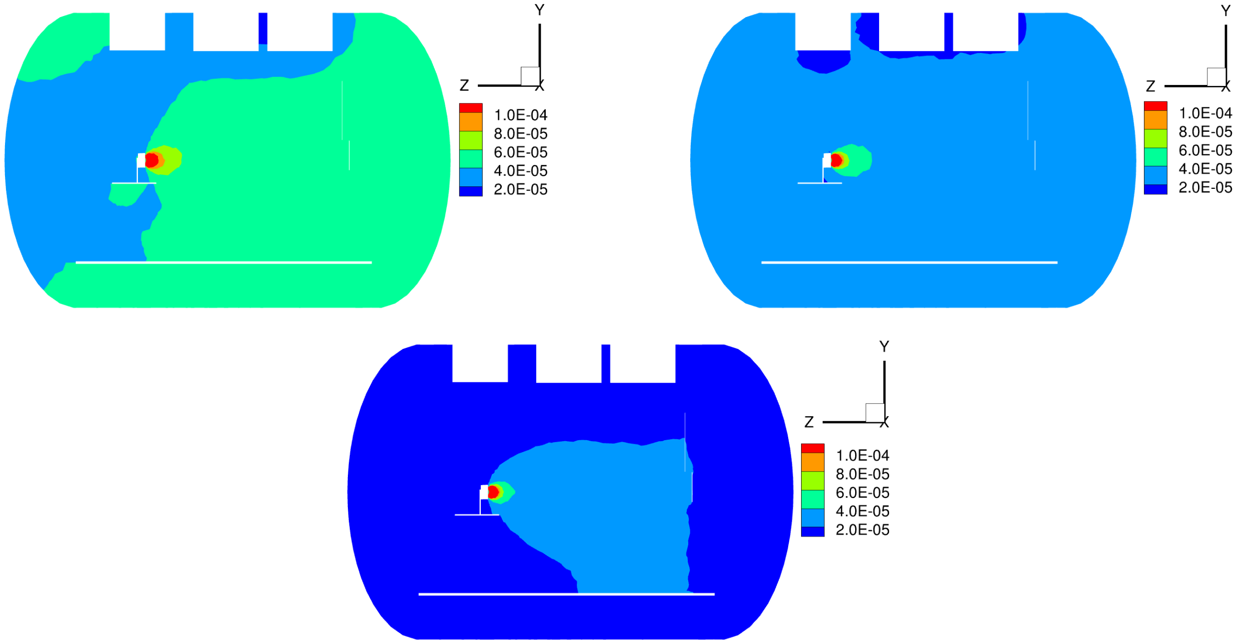

Figure 4. Pressure (Torr) contours throughout a vertical slice of the 3D solution in the LVTF from cold flow simulations using a pump sticking coefficient of 0.2 (upper left), 0.4 (upper right), and 0.6 (bottom middle).

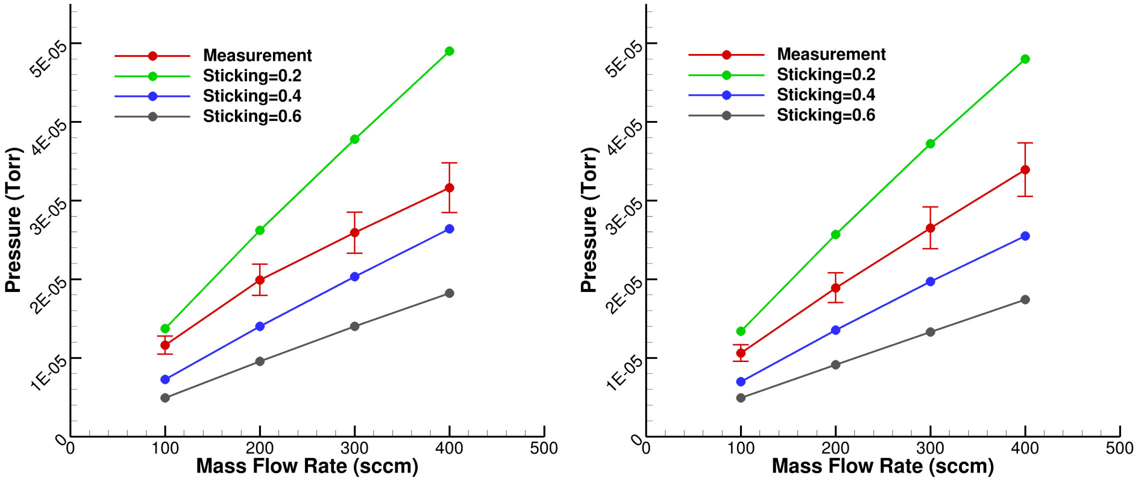

Figure 5. Pressure versus mass flow rate including simulation and experimental results at the location of Gauge 1 (left) and Gauge 2 (right).

Figure 6 also shows a comparison of the model predictions against experimental data for various mass flow rates (for both ion gauges). The shaded regions indicate 5th and 95th percentiles of model predictions over $N = 1000$ Monte Carlo samples of the uncertain inputs. The solid line is the median (50th percentile) model prediction. The increase in model uncertainty for increasing mass flow rates is consistent with the increasing experimental uncertainty (as indicated by the errorbars on the experimental data points). Overall, the model predictions quantitatively agree well with the experimental data given all uncertainties present in the system.

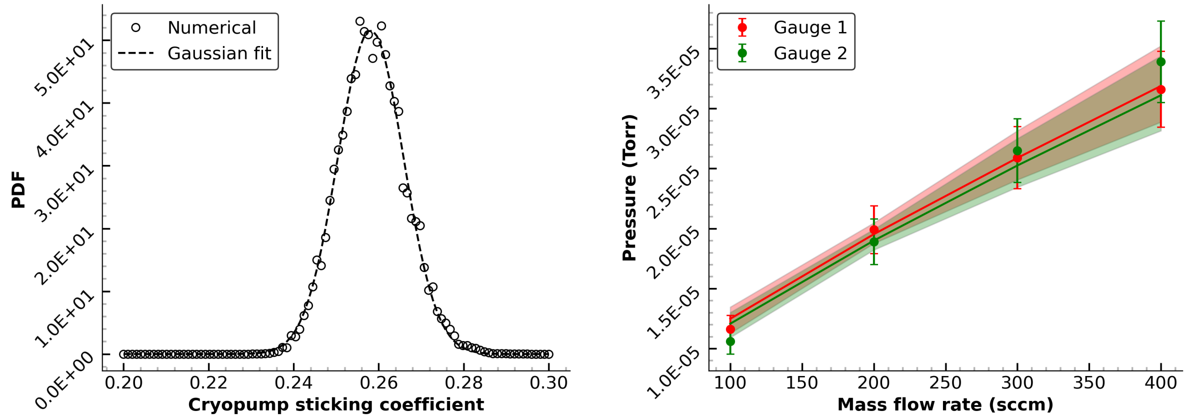

Figure 6. Posterior distribution of the cryopump sticking coefficient (left) and model pressure predictions against experimental data for both ion gauges (right). The markers with error bars indicate the experimental data and noise. The solid lines indicate 50th percentile model predictions over $N = 1000$ samples, with the shaded regions indicating 5th and 95th percentiles.

C. Low Pressure Pump Configuration

For the low pressure configuration, the sticking coefficient of the cryopumps is set using the high pressure inference described above; this allows the sticking coefficient of the cryosails to be inferred separately using the low pressure data. The posterior distribution $p(\theta_{cs}|y)$ and the resulting model predictions are shown in Fig. 7. A skew-normal (SN) distribution is fit to the numerical PDF and has a parameterization of $p(\theta_{cs}|y) \approx SN(\xi = 0.228, \omega = 0.169, \alpha = 3.29)$, with a most likely value occurring at $\theta_{cs} = 0.305$.

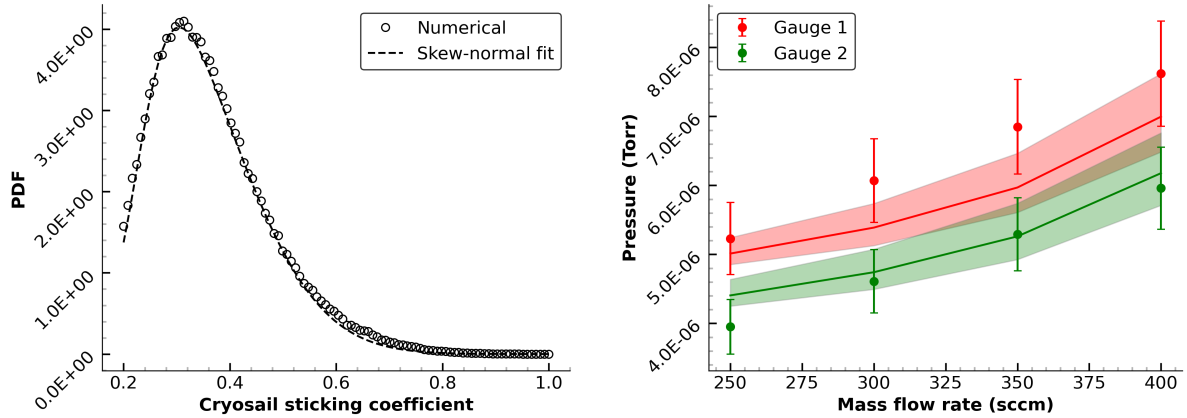

Figure 7. Posterior distribution of the cryosail sticking coefficient (left) and model pressure predictions against experimental data for both ion gauges (right). The markers with error bars indicate the experimental data and noise. The solid lines indicate 50th percentile model predictions over $N = 1000$ samples, with the shaded regions indicating 5th and 95th percentiles.

This PDF reveals a much higher degree of uncertainty in θ_{cs} than is observed in θ_{cp}, given the data at hand. When viewing the comparison of model predictions to experimental data in Fig. 7, it can be seen that a better fit to Gauge 1 data could be obtained with a lower sticking coefficient, while the opposite is true for Gauge 2 data. The PDF intuitively balances this interaction by choosing the middle ground and spreading the uncertainty over a wider area (while it is still skewed towards lower sticking coefficients to obtain predictions more in line with Gauge 1 data). Overall, there is quantitative agreement with the data within model and experimental uncertainties.

This work also seeks to quantify the relative impact of each source of uncertainty on the model predictions of pressure. For this, the first order (S_1) and total order (S_T) Sobol' indices are computed for the low pressure case over varying mass flow rates, as shown in Fig. 8. The quantity of interest for this sensitivity analysis is the pressure at Gauge 1's location. The sensitivity results are very similar when using the pressure at Gauge 2's location. The first order indices exist on the domain $[0, 1]$, with higher values indicating a greater contribution of a given input's *unique* contribution to the output uncertainty; the total order indices also account for interactions between parameters. The fact that S_1 and S_T are nearly identical for each parameter indicates that the model is fairly well decoupled, i.e. the model responds mostly to independent changes in each parameter.

An intuitive result from Fig. 8 is that the two pump sticking coefficients (θ_{cp} and θcs) dominate the output uncertainty; this is indeed why their calibration to experimental data is targeted in this study. Together, they account for over half of the output uncertainty at most flow rates. On the other hand, the 1% relative uncertainty in the mass flow rate is shown to have negligible effect on the model predictions, although its importance increases slightly for higher mass flow rates. This increase is the result of assigning a relative uncertainty to mass flow rate and plotting the indices over increasing flow rates.

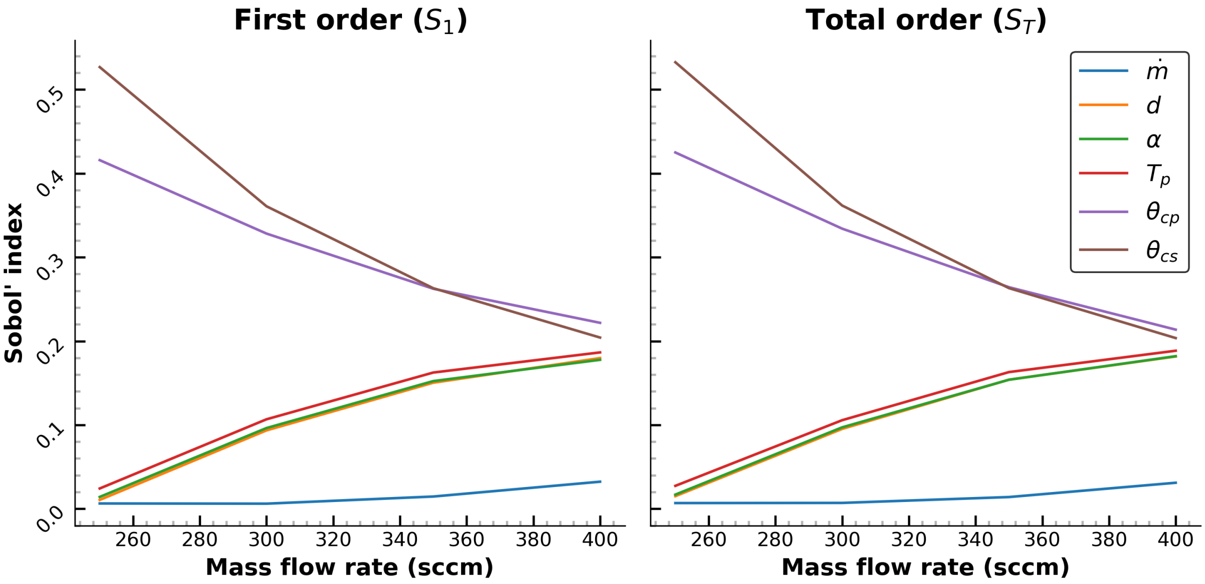

Figure 8. Sobol' first order (left) and total order (right) indices versus mass flow rate for the low pressure configuration.

The increase in the importance of the collision diameter, d, is expected because more collisions take place in gases of greater density. A surprising result is the increasing importance of α and T_p for increasing flow rates. One possible interpretation of this result is that higher flow rates magnify the importance of the accommodation coefficient and pump temperature because higher flow rates cause higher incident fluxes to walls and pumps. This interpretation is in agreement with the corresponding increase in total model uncertainty for increasing flow rates as seen in Fig. 7.

VI. Conclusions

Ground-based electric propulsion testing facilities interfere with thruster operation and with the dynamics of their emitted plasma plumes. Elevated background pressures caused by the inherently limited pumping capabilities of these facilities alter the amount of gas ingested by the thruster, the production of charge-exchange ions, and the divergence of the plume. These effects and others undermine confidence in the ability of performance and lifetime test results conducted in ground-based facilities to accurately reflect how thrusters will behave in space environments. High-power electric propulsion thrusters will exacerbate these facility effects. Thruster and plume models that account for facility effects must accompany ground-based tests to extrapolate from non-ideal chambers to space environments in predictive fashion. This work is an initial step towards equipping thruster and plume models with a means of accurately recreating the elevated background pressure experienced by thrusters in ground-based facilities.

Posterior distributions of sticking coefficients have been obtained for the two types of pumps that populate the University of Michigan's Large Vacuum Test Facility. Surrogate models, trained on hundreds of high fidelity DSMC simulations, enabled detailed Bayesian inference of pump sticking coefficients and thorough uncertainty quantification. The Sobol' method was employed for sensitivity analysis and showed that pump sticking coefficients are the largest contributor to uncertainty in model pressure predictions. Results from simulations using the inferred sticking coefficients contain ionization gauge measurements within error bounds. The agreement between backpressure measurements and simulation predictions indicates the utility of this approach to modeling chamber vacuum pumps. Accurately modeling facility backpressure within simulations of plasma flow experiments is a crucial capability for investigating the role of this facility effect on electric propulsion thrusters and their plumes. Future work will involve the application of these sticking coefficients within simulations of plasma flow experiments of the H9 HET in the LVTF to validate a hybrid particle-fluid plasma plume model. The plasma plume modeling effort will focus on uncertainty quantification in plasma properties and the detailed sticking coefficient uncertainty information obtained here will be invaluable.

VII. Acknowledgements

The authors gratefully acknowledge support from the Joint Advanced Propulsion NASA Space Technology Research Institute. This work utilized resources from the University of Colorado Boulder Research Computing Group, which is supported by the National Science Foundation, the University of Colorado Boulder, and Colorado State University. Resources supporting this work were also provided by the NASA High-End Computing Program through the NASA Advanced Supercomputing Division at Ames Research Center.

References

[1] Saleh, J. H., Geng, F., Ku, M., and Walker II, M. L. R., "Electric propulsion reliability: Statistical analysis of on-orbit anomalies and comparative analysis of electric versus chemical propulsion failure rates," *Space Weather*, Vol. 16, No. 10, 2018, pp. 1561, 1569. https://doi.org/10.1029/2018SW001910.

[2] Goebel, D. M., and Katz, I., *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*, JPL Space Science and Technology Series, John Wiley and Sons, New Jersey, 2008.

[3] Saleh, J. H., Geng, F., Ku, M., and Walker II, M. L. R., "Five-Year Results From the Engineering Radiation Monitor and Solar Cell Monitor on the Van Allen Probes Mission," *Journal of Propulsion and Power*, Vol. 15, No. 2, 1999, pp. 345, 357. https://doi.org/10.2514/2.5432.

[4] Hofer, R. R., Peterson, A. Y., and Gallimore, A. D., "Characterizing Vacuum Facility Backpressure Effects on the Performance of a Hall Thruster," *27th International Electric Propulsion Conference*, 2001.

[5] Walker, J. A., Langendorf, S. J., Walker II, M. L. R., Khayms, V., King, D., and Pertson, P., "Electrical Facility Effects on Hall Current Thrusters:Electron Termination Pathway Manipulation," *Journal of Propulsion and Power*, Vol. 32, No. 6, 2016, pp. 1365, 1377. https://doi.org/10.2514/1.B35904.

[6] Walker II, M. L. R., Gallimore, A. D., Boyd, I. D., and Cai, C., "Vacuum Chamber Pressure Maps of a Hall Thruster Cold-Flow Expansion," *Journal of Propulsion and Power*, Vol. 20, No. 6, 2004, pp. 1127, 1131.

[7] Viges, E. A., Jorns, B. A., Gallimore, A. D., and Sheehan, J. P., "University of Michigan's Upgraded Large Vacuum Test Facility," *36th International Electric Propulsion Conference*, 2019.

[8] Dankanich, J. W., Walker, L. R. W., Swiatek, M. W., and Yim, J. T., "Recommended Practice for Pressure Measurement and Calculation of Effective Pumping Speed in Electric Propulsion Testing," *Journal of Propulsion and Power*, Vol. 33, No. 3, 2017, pp. 668–680.

[9] Hofer, R. R., Cusson, S. E., Lobbia, R., and Gallimore, A. D., "The H9 Magnetically Shielded Hall Thruster," *35th International Electric Propulsion Conference*, 2017.

[10] Choi, M., *Improved Hall Thruster Plume Simulation by Including Magnetic Field Effects*, PhD Thesis, University of Michigan, 2016.

[11] Boyd, I. D., and Schwartzentruber, T. E., *Nonequilibrium Gas Dynamics and Molecular Simulation*, Cambridge Aerospace Series, Cambridge University Press, 2017.

[12] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Clarendon Press, 1994.

[13] Bird, G. A., "Monte-carlo simulation in an engineering context," *Progress in Astronautics and Aeronautics*, Vol. 74, 1981, pp. 239, 255.

[14] Grisnik, S. P., "Measurement of Xenon Viscosity as a Function of Low Temperature and Pressure," *34th Joint Propulsion Conference*, AIAA, 1998.

[15] Agrawal, A., and Prabhu, S. V., "Survey on measurement of tangential momentum accommodation coefficient," *Journal of Vacuum Science and Technology A*, Vol. 26, 2008, pp. 634–645.

[16] Sobol, I. M., "Sensitivity analysis for non-linear mathematical models," *Mathematical modelling and computational experiment*, Vol. 1, 1993, pp. 407–414.

[17] Sobol, I. M., "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates," *Mathematics and computers in simulation*, Vol. 55, No. 1-3, 2001, pp. 271–280.

[18] Saltelli, A., "Making best use of model evaluations to compute sensitivity indices," *Computer physics communications*, Vol. 145, No. 2, 2002, pp. 280–297.

[19] Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., and Tarantola, S., "Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index," *Computer physics communications*, Vol. 181, No. 2, 2010, pp. 259–270.

[20] Chen, G., and Boyd, I. D., "Statistical Error Analysis for the Direct Simulation Monte Carlo Technique," *Journal of Computational Physics*, Vol. 33, No. 3, 2017, pp. 668–680.