IEPC-2002-030

Interactions within a Cluster of Low Power Hall Thrusters

W. A. Hargus, Jr
william.hargus@edwards.af.mil
AFRL/PRSS
1 Ara Road
Edwards AFB, CA
USA

Abstract

This work examines a long duration Hall thruster start transient caused by the vacuum chamber environment. During operation of a cluster of four Hall thrusters, large anode discharge fluctuations, visible as increased anode current and a more diffuse plume structure, occur in an apparently random manner. For single thrusters, the transient appears as a smoothly decaying elevated anode current with a diffuse plume which persists for less than 500 seconds. The start transient is characterized by severe 18 kHz oscillations which dominate the anode discharge. This contrasts with typical steady state behavior of a strong DC component overlaid with a low amplitude 25 kHz component. The main discharge chamber has been previously determined to be the source of this behavior. The work shows that transient appears to be a result of the release of water previously hydrated on to on the surface layer of the boron nitride acceleration channel insulator.

Introduction

During testing of a cluster of Hall thrusters at the Air Force Research Laboratory (AFRL) Electric Propulsion Laboratory, large amplitude oscillations of the anode currents occurred regularly during the first start after evacuation of the vacuum facility. These oscillations visibly manifested themselves as thrusters variously entered and exited *jet* and *diffuse plume modes* over periods of several seconds. These mode transitions often appeared to be synchronized between several thrusters. The behavior appears to be random with a thruster transitioning from the preferred jet mode to more diffuse plume mode with a higher (hence less efficient) discharge current. Often, this triggers change in one of the other thrusters within the cluster.

For example; during one specific instance of cluster operation, two thrusters would alternatively enter and exit jet and diffuse plume modes cyclically with a period of approximately 5 seconds. Once the thrusters were operated for greater than approximately 500 seconds, the oscillatory behavior disappeared and did not reappear in subsequent starts. However when the chamber was returned to ambient atmospheric conditions and subsequently evacuated, a similar behavior returned, although in this case the mode changes appeared more random.

Eventually, test procedure included provision for the conditioning of the thrusters following any exposure to the ambient laboratory conditions. The data during this conditioning time period (approximately 10 minutes) was generally ignored. It was soon realized that this *start transient* exhibits the strongest interaction between independent thrusters, yet observed.

It should be noted that in this work the term *start transient* will refer to an increase in the anode current evident for a period of approximately 500 seconds. This terminology, while grammatically correct, is not necessarily the definition of the commonly discussed Hall thruster *start transient* which describes the transient plasma processes which first initiate the plasma discharge and has a duration on the order of milliseconds.

In order to better understand this phenomenon, it was decided to first examine the behavior of a single thurster. This is also important on its own merit. Unlike ground testing, Hall thrusters on spacecraft do not use an independent mass flow controller. Rather, the anode current is used in a negative feedback loop to throttle the propellant flow. Therefore if the current is anomalously high for a particular flow rate, the closed loop system will regulate the anode current by reducing propellant flow. This could unexpectedly reduce

thrust and impulse from an initial short duration firing. For this reason, understanding the characteristics of this anode current start transient may also be important for missions with short duration initial firings.

Facilities

All measurements in this study were performed in Chamber 6 at the AFRL Electric Propulsion Facility at Edwards AFB, CA, USA. Chamber 6 is a stainless steel chamber with a 1.8 m diameter and 3 m length. It has a measured pumping speed of approximately 32,000 l/s on xenon. Pumping is provided by four single stage cryo-panels (APD single stage cold heads at ~25 K) and one 50 cm two stage APD cryo-pump (<12 K). The chamber is roughed by a oil free Stokes Stealth® mechanical and blower. The chamber is configured such that the 50 cm cryo-pump may be isolated from the chamber and cryo-panels. During the pump down of the chamber, the chamber is first roughed to approximately 4 Pa using the mechanical and blower pumps. At this time, cryo-panel cooling begins. When a single thruster is operated, the background pressure measured by a cold cathode vacuum gauge is approximately $7x10^{-4}$ Pa (corrected for Xe).

The thrusters used for this test are four BHT-200-X3 thrusters which are described in detail elsewhere [1,2]. This is a cluster of four 200 W thrusters being tested to determine the engineering aspects of clustering higher power Hall thrusters on a single spacecraft [3-4]. The cluster of four thrusters is shown in Fig. 1. The thrusters are placed in a 2x2 grid with a center-to-center separation of 115 mm.

Each BHT-200 in the cluster is independently connected to four power supplies. A Sorensen DHP-400-5 is used for the main discharge, a Sorensen DCS-600-1.7E is used to power the cathode keeper, and two Sorensen DLM-40-15's provide power to the magnetic circuit and cathode heater. During all testing the thruster is operated at the nominal conditions shown in Table 1. Following exposure of the thrusters to ambient atmospheric conditions, the cathodes of each thruster are conditioned by flowing 1 sccm (98 µg/s) of xenon and heating for approximately 90 minutes.

An inductive-capacitive (L= 250 µH, C = 13 µF) filter is placed between the anode and cathode external to the chamber. This filter approximately duplicates the impedance characteristics of a power processing unit (PPU) for the BHT-200. The aim of the circuit in this work is not to attempt to replicate PPU characteristics, but to provide isolation of the power supplies from the discharge oscillations of the plasma and to insure that the oscillations measured are not a product of feedback between the power supplies and the plasma.

Xenon propellant (99.995%) flow to the thrusters was metered by use of Unit Instruments model 7301 mass flow controllers (MFC) calibrated for xenon (± 1% of reading). Flow for each anode and cathode is individually metered through a separate MFC. Ten thruster electrical operating parameters are recorded during operation of the system using an Agilent 34970A data acquisition and switch unit. These parameters include the currents and potentials of the anode, cathode, heater, keeper, and magnet circuits. These data are taken at approximately 1 Hz.

Table 1: Thruster Operating Parameters

Anode flow	840 µg/s
Cathode flow	98 µg/s
Anode potential	250 V
Anode current	0.83 A
Keeper current	0.50 A
Magnet current	1.0 A
Heater current	3.0 A

Experimental Observations and Analysis

Figure 2 shows the cluster anode current time history of an unconditioned start. The four thrusters are started within 15 seconds. During a period of approximately 400 seconds, the thrusters all exhibit higher than nominal anode currents. Closer examination as provided by the insert of Fig. 2 shows that the behavior is complex. All thrusters immediately assume a higher than nominal anode current mode with a characteristic diffuse plume. Approximately 65 seconds later, three thrusters leave the diffuse mode, and simultaneously enter the lower current jet mode. Individually, two

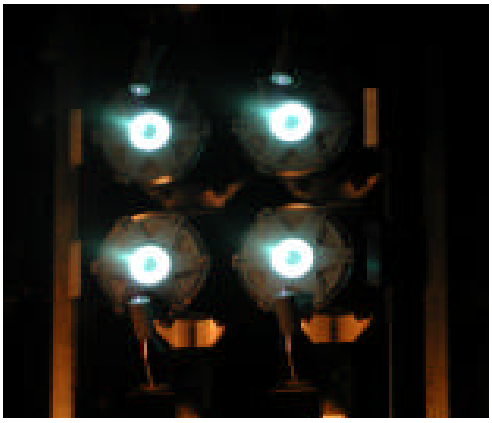

Fig 1. The cluster of BHT-200 Hall thrusters firing within AFRL Chamber 6. The thrusters are numbered 1 through 4 starting from the upper left corner and proceeding counter-clock-wise.

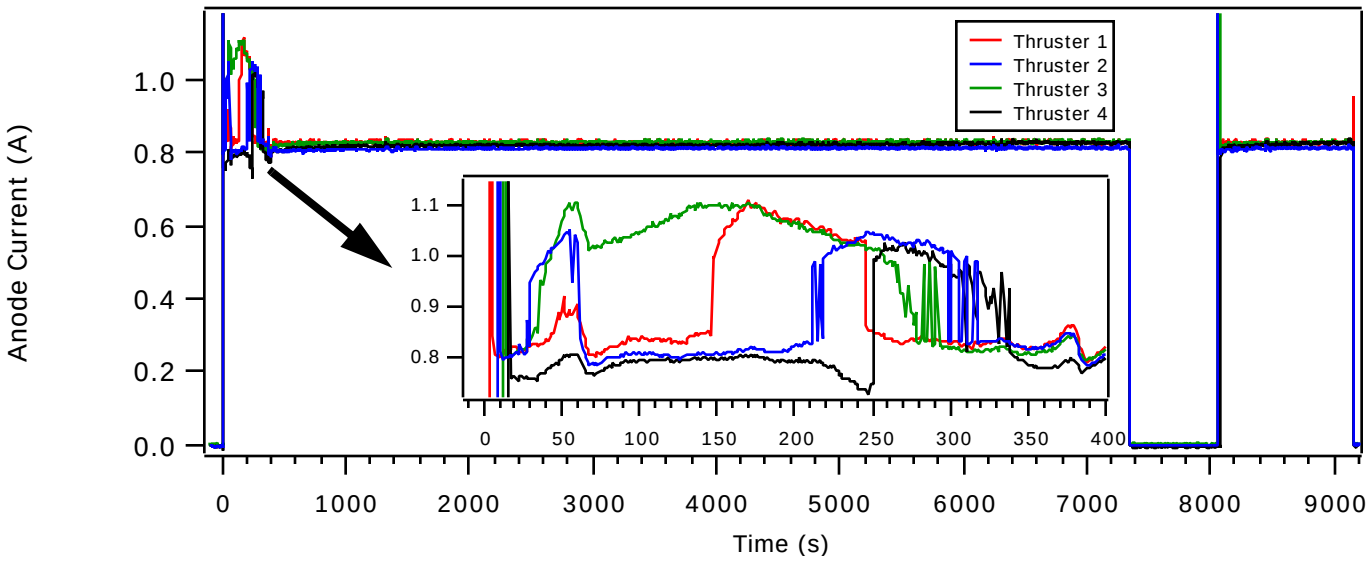

Fig. 2. Anode current transient of four nearly simultaneously started thrusters. Note the transient persists for ~400 seconds and does not reappear in a later restart of the cluster.

thrusters spontaneously leave the jet mode and reenter diffuse mode at later times. Then at approximately 250 seconds after start, two thrusters apparently exchange modes. Subsequent behavior becomes more complex as the thrusters enter and exit jet mode and appear to affect their neighbors. Eventually, the behavior decays and the thruster anode currents all approach their nominal value near 830 mA. Restarting the thrusters after conditioning does not reproduce the start transient. The cathode potential during this start is shown in Fig. 3. The variation of the cathode potential mirrors the anode potential, indicative of significant changes occurring within the Hall thruster anode discharge plasma during the start transient.

In order to better understand the behavior and source of this behavior, it was decided that studying the isolated behavior of a single thruster would prove more feasible

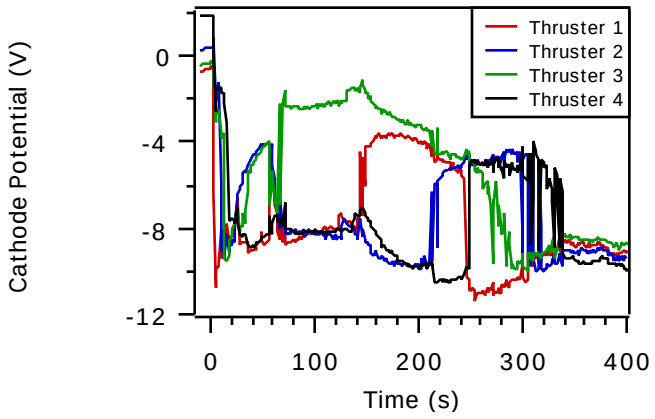

Fig. 3. Cathode potential during the transient of four clustered Hall thrusters.

and could better isolate the effect and its source. The behavior of a single isolated thruster is similar to that shown in Figs. 2 and 3. However, the transition out of the diffuse mode and into the lower current jet mode is consistently repeatable following exposure to ambient atmospheric conditions. Figure 4 shows a typical trace of the anode current during a start after the chamber has been opened to atmosphere. The initial anode current spike (1.5 A) is due to the start procedure where the anode is current limited and the magnet current is off. After the magnet current is switched to its nominal value, the transient consists of anode currents as much as 50% greater than nominal, lasting approximately 300 to 500 seconds. The anode current transient only occurs for the first 300-500 seconds and does not return on restart once the thruster has been conditioned. With subsequent atmospheric exposure, the transient returns. Figure 5 shows that this behavior persists even if the thruster anode discharge is cycled during the time period associated with the anomaly. Interestingly, the thrust level during the period of the anode current transient is unchanged [5]. The 50% increase in anode current simply reduces the efficiency by a like amount.

The start transient shown in Figs. 2-5 occurs under two distinct circumstances. First, it occurs when the chamber has been opened to atmosphere, exposing the thrusters to ambient air. Second, it occurs when the chamber pumps have been turned off, allowing the cryo-panels to regenerate, exposing the thrusters to a 70 Pa *dirty* vacuum. In the latter case, the cryo-pump is valved off from the chamber. Then, it and the cryo-panels are allowed to warm to room temperature. The chamber pressure rises to approximately 70 Pa. This atmosphere within the chamber consists of primarily

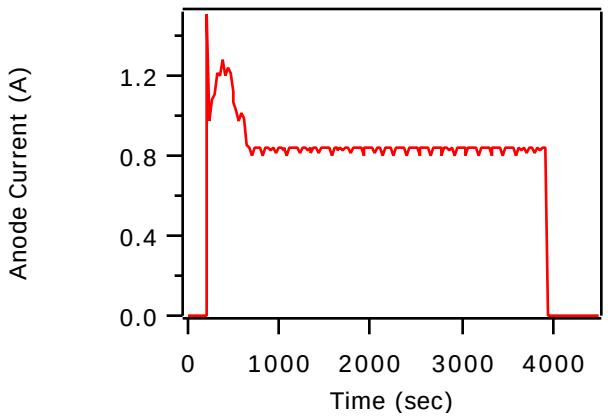

Fig. 4. Anode current trace characteristic of a start following exposure of the thruster to atmospheric conditions for several days. Duration of the anode current transient is 420 seconds. Note: initial current spike to 1.5 A is caused by start-up sequence.

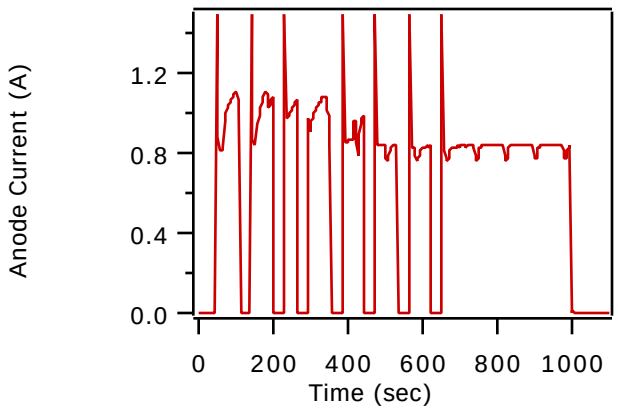

Fig. 5. Anode current trace of the transient with a number of nominally 60 second firings spaced 30 seconds apart. The anomaly persisted for approximately 4.25 min. Note that the thruster has been exposed to regenerated cryo-panel products for a period of 12 days and not exposed to atmosphere.

xenon gas sublimated from the cryo-panels and water which was trapped by the panels during the initial pump down. However, molecular oxygen can be pumped by the cryo-panels at 25 K, so there may be substantial oxygen and nitrogen in the chamber as well [6]. The atmosphere within the chamber during this conditions will be sampled in future tests. Subsequent use of the chamber does not require repressurization to atmosphere. The chamber may be roughed to approximately 4 Pa at which time the cryo-pump and cryo-panels are activated.

The duration of the anode current start transient appears to depend on the degree of exposure to the atmosphere, or gases sublimated from the pumping surfaces. If exposed to atmospheric conditions for several days, the anode current will return to the nominal 830 mA value within 500 seconds. Subsequent exposures to the regenerated vacuum appear to reduce the length of the anode current transient by approximately 20%. For example, thruster 1 was started after exposure to atmosphere and the initial transient persisted for 425 seconds. The chamber was cycled without being opened to atmosphere and the subsequent current transient lasted 340 seconds. This chamber was again cycled without being opened and the anode current transient persisted for only 255 seconds. This trend is typical of all four thrusters.

There are actually two effects shown in Figs. 4 and 5. The anode current transient is accompanied by a periodic (~60 sec) dip in the anode current. Figure 5 shows quite clearly that this period is not affected by the thruster on/off condition. This behavior has been correlated to the pressure fluctuations in the propellant lines of the thruster. The general effect has been documented previously in this facility and elsewhere [5,7]. Due to the low flow of propellant from the pressure bottle through the MFC, the bottle pressure regulator is opening causing a pressure increase in the lines between the regulator and the MFC. The valve on the regulator then closes until the pressure drops by some fraction of the set pressure and then the process cyclically repeats. The MFC responds to these periodic increases in upstream pressure and overcompensates, reducing the flow rate. This overcompensation is due to the flow sensing element being upstream of the throttling valve in the MFC. The MFC does not recognize that it has over compensated and for approximately 3 seconds the flow is reduced by approximately 10%, reducing the anode current by a similar fraction. This issue has been eliminated by replacing the single stage pressure regulator with a two stage regulator. Since much of the data was acquired prior to this action, caution should be taken not to confuse this much slower variation in the anode currents with collective behavior.

The frequency of the anode current oscillations during a start were measured using Tektronix TCP202 current probes and an TEK1103 probe power supply connected to an Stanford Research Systems SR785 dynamic signal analyzer (DC-104 kHz). A plot of the oscillatory behavior during the anode current transient is given in Fig. 6 and compared to the steady state behavior. During the duration of the start transient, there is a strong frequency response at approximately 18 kHz and at 4 harmonics. This behavior is radically different than the steady state behavior also shown in Fig. 6. In the steady state case, there is only a single broad peak at approximately 25 kHz. The steady state spectra is significantly less energetic with peak magnitudes approximately 20 dBA less than the transient case.

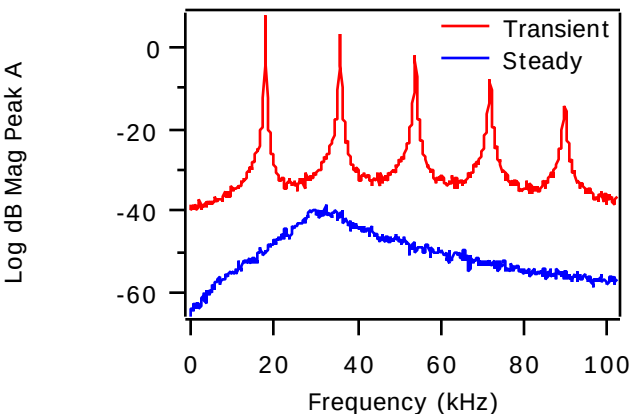

Fig. 6. Two spectra of current oscillations. The first during the transient with an average current of 1.14 A, and the second during steady state operation with an average current of 0.84 A.

Figure 7 shows the time evolution of the anode current frequency spectrum during the start transient. Although the frequencies of the major features are increasing slightly and the higher orders are decreasing in magnitude, these frequency components retain significant power until they suddenly collapse into the steady state single broad peak at 25 kHz. The periodic shifts in the peaks correspond to periods of lowered propellent flow during the pressure regulator approximately 60 second cycle.

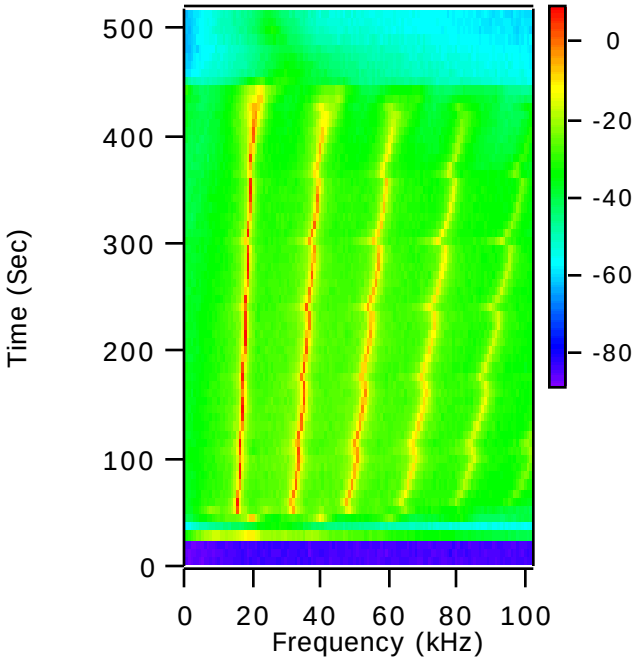

Fig. 7. Contour plot showing the time evolution of the anode Log Magnitude Peak Current (dBA) during a start where the anode current transient is present.

The behavior during the start transient and steady state are very different. Examination of the time domain behavior of the anode current using two Tektronix TCP202 current probes connected to a TDS3012 100 MHz bandwidth oscilloscope is shown in Fig. 8. It shows the dramatic difference between the steady state behavior and the behavior during the anode current transient. During steady state, there is a strong DC anode current component overlaid by a weak 25 kHz component. The behavior during the anode current transient is nearly the opposite. The anode current is primarily AC with peaks measured as high as 9 A. In this mode of operation, the thruster is literally turning itself on and off every 50-60 µs.

The effect of the LC filter is shown in Fig. 9 where the anode current measured at the thruster is compared to the anode current on the power supply side of the filter. The thruster anode sees 18 kHz on/off behavior, while the power supply does not see the oscillations. The LC circuit shields the power supplies' exposure to the oscillations by 30-50 dB. The use of the LC circuit is not strictly necessary since the inherent impedance of the power supplies produces a similar filtering. Here, the filter is used to provide an additional degree of isolation.

A previous effort has demonstrated that the insulator surface condition is responsible for the start transient [5]. An examination of the main discharge chamber of the Busek BHT-200-X3 Hall thruster (similar many other Hall thrusters) reveals that the portions of the thruster in contact with the plasma are limited to the alumina plasma sprayed coated front plate of the magnetic circuit, the anode, and the boron nitride (BN) acceleration channel insulator. Of these three surfaces, the most likely source of contamination are the BN insulator inserts due to their high porosity (~15%) and thermodynamic tendency to hydrate. Some BN grades are capable of absorbing up to 3.5% of their weight in water in

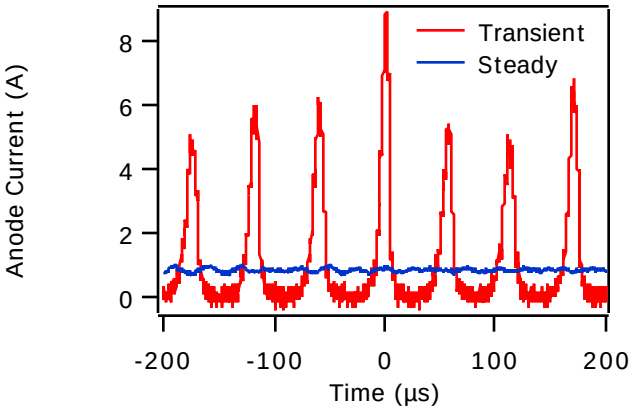

Fig. 8. Current oscillations of the transient and steady state cases seen in the time domain. Conditions are analogous to those in Fig. 4.

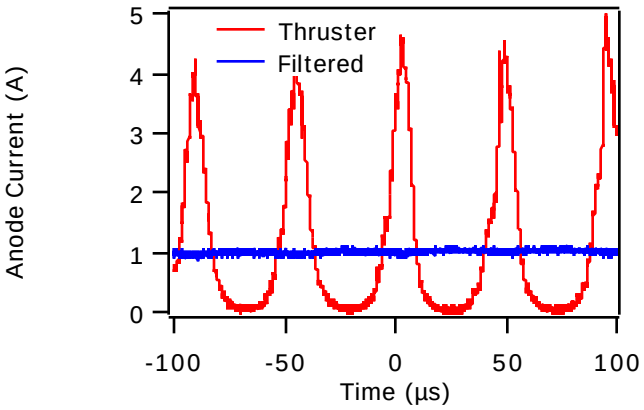

Fig. 9. Current oscillations during the anode current start transient as seen at the thruster and behind the filter at the power supply.

conditions of high relative humidity. Higher grade boron nitrides generally absorb less water (<1%). The hydration is believed to be primarily a surface phenomena, but the depth of penetration is unknown. No data on the depth of penetration is appears to be available. In general, the hydration of BN may be reversed by heating. For example, Saint-Gobain suggests heating their BN products to approximately 400°C for one hour to reverse hydration caused by exposure to ambient conditions.

The plasma within the acceleration channel is either etching away the hydrated layer of the insulator, or heating the hydrated surface and driving the water out of the BN matrix. The impact of relatively high energy xenon ions into the wall affects insulator wall properties. Some wall materials are certainly etched, but ion recombination energy (12.1 eV) as well as kinetic energy (0-200 eV) will be deposited on the wall surface by each ion impact, substantially heating the surface layer. Despite the high thermal conductivity of BN, any given area of the accelerator channel wall will be periodically heated to very high temperatures which result in water molecules being driven from the BN matrix. Considering the time scale of 300-500 seconds, this appears to be the most likely scenario.

The most surprising result is the difference in the acceleration channel current conduction between the anode current transient and steady state. These two cases represent distinct operating conditions. In steady state operation, the small amplitude approximately 25 kHz frequency component is generally attributed to the so-called *breathing mode oscillation* [8]. Here, the ionization occurs in a planar sheet which oscillates transversely within the acceleration channel. The frequency of which is related to the neutral residence time (25 kHz -> 40 µsec). During the start transient, an amplified breathing mode oscillation appears to dominate the main discharge.

In order to test the hypothesis that the anomalous start transient is due to water liberated from the insulator walls, a series of emission spectroscopy experiments were performed using an Ocean Optics USB-2000 fiber optic spectrometer. The light was collected from the Hall thruster anode discharge at approximately 45° to the plume center line through a quartz window by a lens which focused the incident light onto a fiber optic launcher which then routed the light to the spectrometer.

Spectra from the anode discharge between 400 and 1000 nm were gathered at 30 second intervals to determine the change in the spectra over time. Figure 10 shows two typical spectra. The first is of the discharge during the start transient soon after the discharge has been started (~30 s) and the anode current is near maximum (1.20 A). The second spectra is 720 seconds at which time the anode current has reached its nominal value (830 mA). Each emission spectrum represents 100 averaged samples with 5 ms acquisition times. Neither spectrum is intensity calibrated.

Upon examination of Fig. 10, it is immediately evident that the two spectra have some similarities, but also exhibit significant differences. The strongest neutral xenon transitions (823, 828, and 882 nm) are of nearly equal intensity in both cases. However, the ionic xenon lines which form the bulk of the spectrum below 600 nm are of much lower intensity for the steady state case than during the start transient.

There are several possible explanations for this behavior. First, the electron temperature of the plasma may be higher during the transient period. The 18 kHz enhanced breathing mode with high instantaneous anode currents is likely producing strong unsteady electric fields within the acceleration channel. This may preferentially raise the bulk electron temperature and increase the density of excited state ions and thereby the emission from the excited states seen in

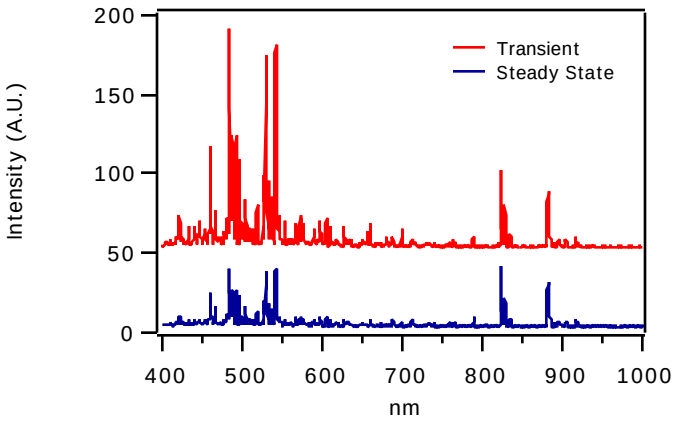

Fig. 10. Visible Spectrum of main discharge. Red-shortly after start with maximum discharge current during start transient. Blue-steady state operation.

Fig. 10. Since the anode electron current is up to 150% of the nominal value, the electron density may be higher within the acceleration channel. This could also raise the density of the ionic excited states. It may also be possible that ion density is increased due to less efficient ionization or longer ion residence time during the start transient.

Figure 11 shows the emission time history of three representative xenon transitions. The ionic transitions are well characterized by the $5d[3]_{7/2} - 6p[2]^0_{5/2}$ transition at 605 nm which is represented by the solid triangles. All ionic transitions sampled tend to follow the behavior illustrated by this trace. It is characterized by an initial low value when the anode discharge is in a glow discharge mode with the magnetic current at zero. (In this case, the time period with no magnetic field was deliberately extended to examine the glow discharge behavior.) As soon as the magnetic current is energized, the signal rises to its largest value and then decays to a value of approximately 40% of maximum. The behavior occurs in the period as the start transient.

The neutral line intensities appear to fall into two broad categories as shown in Fig. 11. The first is represented by the $6s[3/2]^0_1 - 6p[1/2]_0$ transition at 828 nm and the second by the $6s[3/2]^0_2 - 6p[5/2]_3$ transition at 882 nm. In both cases, the signal is at its peak value during the initial glow discharge. After the anode current has reached steady state, emission from both transitions eventually reach 50% of glow discharge mode. However, the behavior is very different during the intermediate period. The neutral transitions which are optically coupled to the ground state immediately drop to an intermediate value which remains near constant, except for a small rise just prior to the anode current reaching steady state. When the anode current reaches steady state, the signal drops suddenly to its final value. Alternatively, several transitions with metastable (MS) lower states (i.e. not optically coupled to lower states) exhibit different behavior. After the magnetic circuit is energized, the emission level drops further than the other neutral case. The emission level then rises to its steady state value.

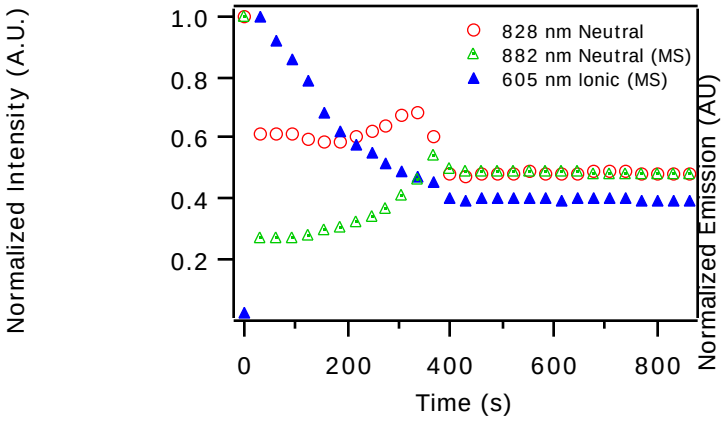

Fig. 11. Normalized representative line intensities for neutral xenon (828 and 882 nm) and ionic xenon (605 nm). Note that the neutral 882 nm and ionic 605 nm transitions have metastable lower states.

In Fig. 11, it is interesting to point out that the neutral relative intensity ratio during the initial glow period and steady state are the same, indicating the upper states of these transitions are similarly populated. This implies that the temperature is also similar in the two cases. That the intensity ratios differ during the transient implies that the temperature is different during this period. By only examining these highly excited states and without greater knowledge of the equilibrium state, any further discussion is unfortunately purely speculative.

A 0.75 m focal length vacuum UV monochrometer (Acton 704) was placed normal to the plume center line and imaged the exit plane of the Hall thruster. Figure 12 shows the atomic hydrogen Lyman alpha emission (121.6 nm) collected during the start transient relative to the *excess anode current* (defined as actual current minus nominal current). Also, included in the plot is the emission signal of the atomic hydrogen Balmer alpha transition acquired by the fiber optic spectrometer.

Figure 12 shows that both hydrogen emission profiles follow the excess anode current closely. Further reinforcing the hypothesis that the excess current is a product of water vapor released into the discharge channel. As expected, subsequent restarts did not produce measurable hydrogen emission. Figure 12 implies a strong correlation between the emission of hydrogen and the excess anode current. The near constant intensity ratio between the two hydrogen transitions indicates that relative populations are not changing and that the electron temperature within the plasma is also likely invariant during the transient time. This

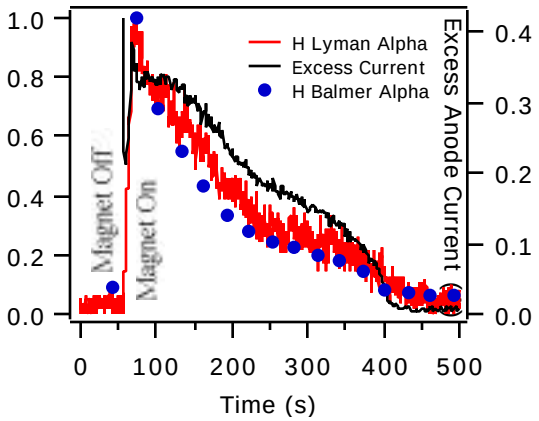

Fig. 12. Normalized hydrogen emission (Balmer alpha, and Lyman alpha) compared to the excess anode current (defined as actual current minus nominal current).

behavior is consistent to that indicated by neutral xenon emission in Fig. 11.

If the anode current is simply proportional to the volumetric flow rate, approximately an additional 5 sccm would be required to raise the anode current by 50%. Making the simple assumption that the excess anode current is solely due to the influx of water vapor liberated from the insulator walls, it is possible to determine the quantity of water required raise the current. A water vapor flow rate of 5 sccm of 500 s duration with a 50% duty cycle yields a total mass of 17 mg. This is a maximum value since each atom in the water molecule could conceivably be ionized and accelerated by the anode discharge. Assuming a 1% hydration (the maximum for higher grades of BN) and a typical density of 1.8 g/cm^3, approximately 1 cm^3 of BN is hydrated. The total BN surface area exposed to the plasma discharge is near 14 cm^2. This yields an average depth of penetration of approximately 700 μm. This is the representative depth of hydration when the BN insulator is exposed to atmospheric conditions (generally less than 20% relative humidity at AFRL). It may also be the depth to which the plasma is able to drive out water vapor. In either case, the depth is on the order of 1 mm and is significant.

This hydration depth calculation should be viewed only as an estimate and likely high by a factor of 3, or more. The assumption that the increase in current is only due to an additional volumetric flow rate is questionable. Low frequency oscillations are known to govern the electron diffusion across magnetic field lines [8-9]. It is likely that some of the additional anode current is due to increased electron conduction due to the high amplitude oscillations shown in Figs. 8 and 9.

Conclusions

Due to this effort, the start transient described in this work is better understood. Specifically, the effect on a single thruster. This start transient examined is only of interest during initial Hall thruster operation. Its effect is generally neglected in the laboratory for good reason. The transient only occurs within the first 500 seconds of operation after exposure to atmospheric water vapor which is hydrated into the surface of the BN acceleration channel insulator. After the thruster has been conditioned, subsequent restarts, without exposure to water vapor, do not exhibit the start transient.

For single thruster operations, this transient can generally be neglected. Station-keeping systems with lifetimes of thousands of hours may neglect the first 500 seconds of degraded performance. Systems that may not be able to neglect this behavior are smaller satellite systems which are launched en-masse. These systems may require short duration firings with known impulses to accurately place formations of satellites into precise orbits. The increasing interest in the space community in distributed architectures indicates that this issue may soon have to be addressed.

For a cluster of Hall thrusters, the difficulties associated with the start transient can be eliminated by operating each independently to sequentially condition each thruster. Once conditioned, the more chaotic start transient fluctuations will not return until again exposed to water vapor. This option works well in the laboratory, but may not in orbit due to spacecraft dynamics. The mechanisms governing the behavior of the current transient on a cluster of Hall thrusters firing in close proximity needs to be further examined. There is still no understanding as to the mechanism which results in thrusters variously entering and exiting jet mode during an unconditioned cluster start. Based on experience gathered from a single thruster this behavior appears counter intuitive. Future efforts will attempt to understand the mechanisms which drive this behavior in cluster operations.

Acknowledgements

The author would like to thank Professor Cappelli and Dr. Gascon of Stanford University for the loan of the two spectrometers used in this study, as well as Dr. Katz of JPL for his interest and thoughts concerning these measurements.

References

1. R. Spores, G. Spanjers, M. Birkan, and T. Lawrence, “Overview of the USAF Electric Propulsion Laboratory,” AIAA-2001-3225, *37th Joint Propulsion Conference and Exhibit*, 8-11 July, 2001, Salt Lake City, UT.

2. V. Hruby, J. Monhwiser, B. Pote, C Freedman, and W. Connally, “Low Power Hall Thruster Propulsion System,” IEPC-1999-092, *Proceedings of the 26th Electric Propulsion Conference*, 1999, Kitakyushu, Japan.

3. B. Beal, A. Galimore, and W. Hargus, “Preliminary Plume Characterization of a Low-Power Hall Thruster Cluster,” AIAA-2001-4251, *38th Joint Propulsion Conference and Exhibit*, 7-10 July, 2002, Indianapolis, IN.

4. W. Hargus and G. Reed, “The Air Force Clustered Hall Thruster Program,” AIAA-2001-3678, *38th AIAA Joint Propulsion Conference and Exhibit*, 7-10 July, 2002, Indianapolis, IN.

5. W. Hargus and B, Pote, “Examination of a Hall Thruster Start Transient,” AIAA-2001-3956, *38th AIAA Joint Propulsion Conference and Exhibit*, 7-10 July, 2002, Indianapolis, IN

6. D. Lide, Ed., *Handbook of Chemistry and Physics: 74th Edition*, CRC Press, Boca Raton, FL, 1993.

7. D. Morgan, “Ultrapure Gas Delivery: Devising Specifications for Optimizing Point-of-Use Pressure Regulators,” *MicroMagazine.com*, Los Angeles, Sept. 2000.

8. J. Boef and L. Garrigues, “Low Frequency Oscillations in a Stationary Plasma Thruster,” *Journal of Applied Physics*, Vol. 84, No. 7, Oct 1998.

9. S. Yoshikawa and D. Rose, “Anomalous Diffusion of a Plasma Across a Magnetic Field”, Physics of Fluids, Vol. 5, No. 3, March 1963.