A Fully Microfabricated Two-Dimensional Electrospray Array with Applications to Space Propulsion

by

Blaise Laurent Patrick Gassend

Diplôme d'Ingénieur, École Polytechnique, France, 2001
Master of Science, Massachusetts Institute of Technology, 2003

Submitted to the
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Author . . . .
Department of Electrical Engineering and Computer Science
May 23rd, 2007

Certified by . . . .
Manuel Martínez-Sánchez
Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by . . . .
Arthur C. Smith
Professor of Electrical Engineering and Computer Science
Chairman, Department Committee on Graduate Students

A Fully Microfabricated Two-Dimensional Electrospray Array with Applications to Space Propulsion

by

Blaise Laurent Patrick Gassend

Submitted to the Department of Electrical Engineering and Computer Science
on May 23$^{\text{rd}}$, 2007,
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

This thesis presents the design, fabrication and testing of a fully-integrated planar electrospray thruster array, which could lead to more efficient and precise thrusters for space propulsion applications. The same techniques could be used for making arrays to increase throughput in many other electrospray applications.

Electrospray thrusters work by electrostatically extracting and accelerating ions or charged droplets from a liquid surface to produce thrust. Emission occurs from sharp emitter tips, which enhance the electric field and constrain the emission location. The electrospray process limits the thrust from a single tip, so that achieving millinewton thrust levels requires an array with tens of thousands of emitters. Silicon batch microfabrication has been used, as it is well suited for making large arrays of emitters.

The thruster is made using Deep Reactive Ion Etching (DRIE) and wafer bonding techniques, in a six mask process, and comprises two components. The emitter die with up to 502 emitters in a 113 mm^2 area, is formed using DRIE and SF6 etching, and is plasma treated to transport liquid to the tips in a porous black-silicon surface layer. The extractor die incorporates the extractor electrode, a Pyrex layer for insulation, and springs which are used to reversibly assemble the emitter die. This versatile assembly method, with 10 μm RMS alignment accuracy and 1.3 μm RMSD repeatability, allows the extractor die to be reused with multiple emitter dies, and potentially with different emitter concepts than the one presented.

The thruster, weighing 5 g, was tested with the ionic liquids EMI-BF4 and EMI-Im. Time of flight measurements show that the thruster operates in the ion emission regime most efficient for propulsion, with a specific impulse around 3000 s at a 1 kV extractor voltage. Emission starts as low as 500 V. Currents of 370 nA per emitter have been recorded at 1500 V, for an estimated thrust of 26 nN per emitter or 13 μN total, and a 275 mW power consumption. The thrust efficiency is estimated around 85%. In good operating conditions, the current intercepted on the extractor electrode is well below 1 %, increasing to a few percent at the highest current levels. The beam divergence half width half maximum is between 10 and 15°.

Thesis Supervisor: Manuel Martínez-Sánchez
Title: Professor of Aeronautics and Astronautics

Acknowledgments

I like to compare graduate school to a long tunnel. The day you enter the tunnel, you naively think that it is just another tunnel. A few years later, when you realize your mistake, the exit of the tunnel is nowhere in sight, the entrance is a long forgotten memory, you are getting very tired of the scenery, and you realize that Jorge Cham's PhD comic is the true story of your life. In this section, I would like to thank all the people who helped me keep the faith that the tunnel has an exit. Without them, I would surely have given up along the way a long time ago.

First of all, I would like to thank my thesis committee. I am very grateful to Prof. Manuel Martínez-Sánchez for taking me into his group to do experimental work, despite my computer science background. Prof. Akintunde Ibitayo Akinwande wasn't really given a choice. I showed up in his office one day, and announced that he was my MTL faculty supervisor. He took it in stride, and took upon himself to educate me in fields as diverse as field emission, economics and philosophy. Dr. Luis Fernando Velásquez-García has saved me years if not months of frustration by advising me on all things MEMS related. Under his guidance and mentoring, I was able to make devices that worked the first time, avoiding most of the mistakes that beginners usually make. Finally, Prof. Srinivas Devadas has had an inestimable role in allowing me to follow my dreams at MIT. In my six years at MIT, I left his research group not once but twice. Most people would have given up on me after the first time, but not Srini, who continued to fund me during my first year working on electrospray propulsion, and ended up serving on my thesis committee. Over the years Srini has also been an excellent mentor to whom I owe a lot, and I doubt I would have had as much flexibility in setting my path without his constant support.

Next, I would like to thank Prof. Paulo Lozano, whose open office door I have often turned to when looking for background on electrosprays, to try out crazy ideas, or just for help in finding equipment in the lab. I am also very grateful for to apparatus left behind from his PhD work, which I used to test my thruster, saving myself weeks or months of effort in the process.

My devices were fabricated in the facilities of the Microsystems Technology Laboratories (MTL). MTL is a remarkable community of extremely competent people, with whom it has been a pleasure to work. I owe a great debt to the MTL specialists who devote their days to repairing the machines that students like myself specialize in taking out of service. Remarkably, they are very open to experimentation, even when it might mean more repair or cleaning work for them down the road. Their relentless fight to keep the lab running, allows the lab users to focus on their work and creativity. Much of the credit for this productive environment goes to Dr. Vicky Diadiuk, who has also always been available to provided me with process related guidance.

Even though their data did not arrive in time to be in this thesis. I am very grateful to Michael McGuirk and Seth McElhinney at Busek Co., who have been burning the midnight oil trying to get me thrust data in time for publication.

I would never have seen the light at the end of the tunnel had it not been for the support of the numerous friends and colleagues I have met over the years at CSAIL, RLE, SPL, MTL, Underwater Hockey, Rocket Team and MIT in general. I will abstain from listing you all, as there are too many of you to avoid making some grievous omissions. I would nevertheless like to mention Carl Dietrich and Andrew Heafitz, for the influential role they have had in my life as founders of the MIT Rocket Team. Without the Rocket Team, I would probably never have taken 16.512, met Prof. Martínez-Sánchez, or ended up working on electrospray propulsion. It is ironic how the most important decisions in life are usually guided more by luck than by intent. Carl and Andrew have helped me realize the importance of leadership. When they set out to build their rocket engine, I wonder if they realized the many ways, large and small, in which they would change the course of people's lives.

Graduate school is notorious for putting a strain on family life. This doesn't appear to be the case when you have a perfect wife like Valérie. I couldn't have picked a better person with whom to embark on the emotional roller-coaster of graduate school and of life. I didn't pick my parents, but if I had the opportunity to do so, I would pick the same ones. They have done a lot of things right, to get me to where I am today. I just wish I knew exactly what all the things they did right were in order to repeat them.

Last, but not least, I would like to acknowledge my funding sources. This work was funded by the Air Force Office of Scientific Research, Mitat Birkan monitor, the Space and Naval Warfare Systems Center award N66001-04-1-8925, manager Richard Nguyen, and the US Army Soldiers Systems Command award W911QY-05-1-0002, manager Henry Girolamo. I am also very grateful to Prof. Srinivas Devadas for financially support me during my first year of work on this project.

And for those of you still working on your PhDs, keep the faith! The light does appear at the end of the tunnel. But beware, when it does appear, the exit of the tunnel arrives faster than you could imagine, so be prepared.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Electrospray is a process which occurs when a conductive liquid surface is placed in a strong electric field. Under the influence of surface tension and the field, the liquid surface deforms into a static cone shape, first described mathematically by G.I. Taylor [1]. Near the cone tip, Taylor's model breaks down due to the high field, and charged particle emission is observed. To facilitate the generation of the strong electric field needed to start this process, electrosprays are usually produced from liquid located on a field enhancement structures, such as a needle or the end of a capillary tube (we shall call these *emitters*). Because the amount of emission that can be stably produced from a single emitter is limited, arrays of emitters may be needed to reach the levels of emission that are desired for a specific application.

This thesis describes the design, fabrication and characterization of a fully-integrated, two-dimensional electrospray array using silicon microfabrication technology. The main challenges that need to be addressed are: fabrication of emitters that are able to transport liquid to their tips, accurate location of the extraction electrode so that the emitted particles pass through apertures in the electrode, and electrical insulation of the electrode from the underlying emitters. The target application for our array is space propulsion, though many other applications of this technology exist. Some key features of our implementation are that the electrospray operates in the pure ion regime which is most efficient for space propulsion, and that the extractor electrode can be assembled and disassembled by hand, to ease experimentation.

We begin this chapter by introducing the reader to electrosprays in Section 1.1, and rocket propulsion in Section 1.2. Then we combine these two fields in Section 1.3, which gives background on electrospray propulsion. Finally, section 1.4 reviews other work on electrospray arrays, and Section 1.5 outlines the contents of this thesis.

1.1 Background on Electrospray Sources

Figure 1-1 shows a diagram of a simple electrospray source. Liquid is fed through the tip of a hollow needle, where a strong electric field creates a Taylor cone. The electric field is created by biassing the needle relative to an extractor electrode. The particles created by the electrospray process pass through an aperture in the extractor electrode.

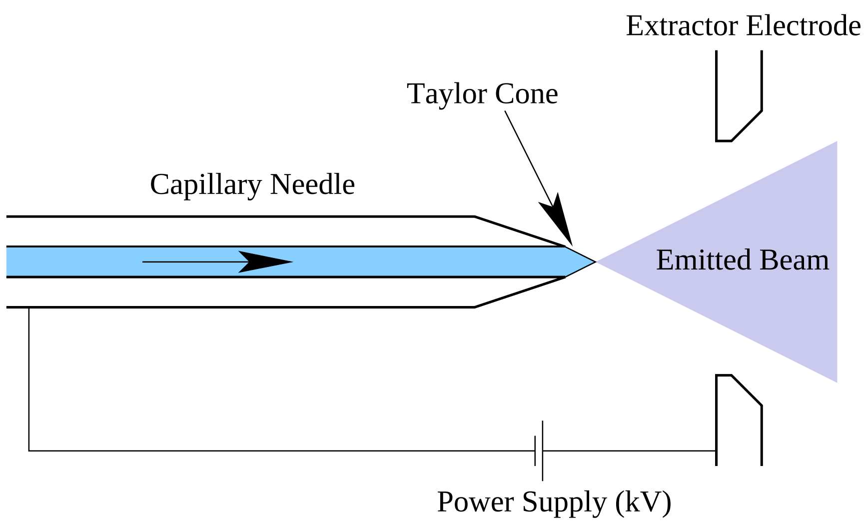

Figure 1-1: Basic elements of an electrospray source

The mechanism through which particles are emitted at the tip of the Taylor cone depends on the regime of operation. In the *cone-jet* regime, the tip of the cone deforms into a cylindrical jet, which breaks up into charged droplets [2, 3]. In the *ionic* regime, the field strength near the tip of the cone is strong enough for ions to be emitted directly from the liquid surface [4]. A mixed regime is also possible in which both ions and charged droplets are emitted [5].

1.1.1 Electrospray Applications

Numerous applications have been found for electrosprays. So far, the most successful application is as an ion source for mass spectrometry of large biomolecules, which resulted in a Nobel prize in Chemistry for J. Fenn in 2002 [6]. In electrospray mass spectrometry, the solution to be analyzed is electrosprayed in the cone-jet regime. The solvent in the emitted droplets evaporates, and the droplets fragment until only the ions to be studied remain. Other applications of the cone-jet regime include: fuel atomization in combustors [7, 8], nanoparticle formation [9], aerosol generation [10], fiber formation by electro-spinning [11] (in this case the jet does not break up into droplets), crop spraying [12], spray-painting [13], polymer coating [14], etching [15] and printing [16]. The ion regime has been used to make liquid metal ion sources, which can produce ion beams of Gallium, Indium or other metals [17, 18]. In this thesis, we will focus on the application of electrosprays to space propulsion, in which the acceleration of the emitted particles in the electric field is used to produce a net thrust [3, 19].

1.1.2 The Case for Electrospray Arrays

The physics of the electrospray process sets limits on the rate of emission occurring from a single Taylor cone in a particular stable operating regime. To increase the amount of emission, arrays of electrospray emitters can be used. By allowing batch production of hundreds to hundreds of thousands of emitters or more per wafer, microfabrication techniques seem to be good candidates for cheaply producing such arrays. Compared with traditional manufacturing methods, these techniques should allow higher emitter packing density, lower cost per emitter, and a higher manufacturing uniformity between emitters.

Many of the applications we have mentioned would benefit from electrospray emitter arrays. Nanoparticle production, aerosol generation, spray-painting, crop spraying, polymer coating, fuel injection, and space propulsion, can all benefit from simply multiplying the rate of material emission by operating many emitters in parallel. For other applications, more elaborate arrays are needed. In mass spectrometry, electrospray arrays are routinely used for high throughput experiments. In this case, though, each emitter in the array is connected to a separate liquid source to be analyzed, and each emitter is fired separately [20]. In more elaborate mass spectrometry setups, multiple electrospray sources with independent liquid supplies may be used simultaneously [21]. First, carefully selected subsets of the sources are fired in succession. Then, the measurements from each subset of sources are processed to determine the composition of each source. With this method, greater sensitivity is obtained for a given amount of mass spectrometer time than would be possible if each source had been fired alone. Printing applications require yet another type of electrospray array. This time, all the emitters in the array emit the same liquid, and can use a common liquid supply. However, each emitter must now be controlled independently so that a pattern can be printed. More complex printing applications could use more than one liquid supply if multiple materials (e.g., colors, RNA bases, etc.) need to be printed.

Clearly, a lot of variety is possible when considering electrospray arrays. This thesis focuses on the simplest case in which all the emitters are fired at once from a single liquid supply. This is the simplest case for extractor design, but raises some questions about undesirable interactions between emitters via the liquid feed system.

1.2 Background on Rocket Propulsion

Propulsion is an integral part of our daily lives. Every time we move from one location to another, be it by foot, car, boat, air, or even wheeled office chair, we are using propulsion. In each of these cases, we set ourselves in motion by pushing on objects in our environment: the floor, the road, the ocean, or the atmosphere. By Newton's third law, the object we push on pushes us in return, and by Newton's second law, this reaction sets us in motion. Likewise, when we are ready to stop, we push on our environment, and the reaction makes us stop. All the motion in our everyday lives is internal to the Earth-Atmosphere system. As we move about, the center of gravity of the system continues unperturbed in its endless dance with the Sun, Moon and planets.

As we turn to space exploration, a fundamental problem appears. For a spacecraft in outer space, there is no longer anything to push on to get in motion. Because of this, a fundamentally different approach to propulsion is needed. This approach is called *rocket propulsion*. With rocket propulsion, the spacecraft that wishes to accelerate pushes on a bit of its own mass, called *propellant*, to accelerate. When this happens, the center of gravity of the spacecraft+propellant system is unperturbed, but the propellant goes in one direction, and spacecraft begins its journey in the other.

Suppose that the spacecraft is initially at rest. It accelerates a small mass of propellant $\mathrm{d}m$ up to a velocity u by applying a constant force T over a brief period of time $\mathrm{d}t$. The acceleration of the propellant mass is $\frac{u}{\mathrm{d}t}$, so by Newton's second law,

$$T = \mathrm{d}m\frac{u}{\mathrm{d}t} = \dot{m}u. \tag{1.1}$$

In this expression, $\dot{m} = \frac{\mathrm{d}m}{\mathrm{d}t}$ is the average rate of propellant ejection during $\mathrm{d}t$. By Newton's third law, T is also the force that the spacecraft undergoes, i.e., the *thrust* of the propulsion system. Equation (1.1) immediately generalizes to a spacecraft equipped with a *thruster* continuously ejecting mass at a rate $\dot{m}$.

1.2.1 Specific Impulse

Once it is in motion, if the spacecraft again wishes to change course, the previously used propellant is no longer at hand, and so new propellant needs to be ejected. As the spacecraft maneuvers, more and more propellant is ejected, and the spacecraft mass slowly dwindles until there is no propellant left, and all that remains is structure, and its payload. At this point, the spacecraft can no longer alter its course.

How much maneuvering can the spacecraft do before it runs out of propellant? We measure the amount of maneuvering by the quantity Δv. This is a measure of how much the spacecraft could change its velocity if it used all its propellant to accelerate in a single direction. By integrating Newton's second law, and applying

Equation (1.1),

$$\Delta v = \int_{t_1}^{t_2} \frac{T\mathrm{d}t}{m} = \int_{m_0}^{m_0+m_p} \frac{u\,\mathrm{d}m}{m} = u \ln\left(\frac{m_0 + m_p}{m_0}\right), \tag{1.2}$$

where m_p is the propellant mass, and m_0 is the "dry" mass of the spacecraft. This relation is usually reversed to give the *Rocket Equation*

$$\frac{m_p}{m_0} = e^{\frac{\Delta v}{u}} - 1, \tag{1.3}$$

which states that the amount of propellant needed for a mission increases exponentially with the amount of Δv that is required.

The exhaust velocity u is determined by thruster technology, and Δv depends on the missions to be accomplished. Thus, Equation (1.3) says that the only reasonable missions are the ones that require a Δv comparable to or smaller than the exhaust velocities for existing technologies, and explains why thruster technologies allowing high exhaust velocity are an essential ingredient for extending the bounds of space exploration.

In practice, it is uncommon to specify the exhaust velocity of a thruster. Instead it is its *specific impulse*, i.e., the impulse (force integrated over time) $T\mathrm{d}t$ that can be produced per unit of propellant weight g$\mathrm{d}m$, which is specified. Taking the ratio of these two quantities, we find that specific impulse is given by

$$I_{\mathrm{sp}} = \frac{T}{\dot{m} g_0},$$

where g_0 is the acceleration of gravity on the Earth's surface. Specific impulse is measured in seconds, and can be seen as amount of time during which a quantity of propellant can support its own weight on the surface of the Earth. Specific impulse and exhaust velocity can be used interchangeably, in any case, the units indicate which one is being used.

1.2.2 Thrust-to-Weight Ratio

So far, we have seen that depending on how much maneuvering is needed (Δv), a thruster with a sufficiently high specific impulse (I_{sp}) is required. In practice, the time it takes to accomplish a maneuver is also important. The most striking example of this fact occurs when you try to get into space from the surface of the Earth. The Δv of getting into orbit is not a fundamental constant; it depends on how fast you get off the ground. Indeed, as a rocket lifts off the launch pad, a large proportion of its thrust is spent fighting gravity. Imagine a rocket that produces just enough thrust to fight its own weight. It can hover just above the ground spending Δv, but will never get into space. Thus, a thruster for a launch mission, needs to produce more thrust than the weight of the rocket it is trying to launch. Likewise, any space mission has minimum thrust requirements, and we need a metric for evaluating the effectiveness of a technology at producing high thrust.

Any given thruster technology can reach any thrust level simply by using multiple thrusters. However, multiplying the number of thrusters increases the spacecraft mass. A relevant measure of a technology's ability to produce high thrust must therefore take into account the mass of the thruster, and the natural measure is the thrust-to-mass ratio, or, with the influence of the imperial system of measurement, the thrust-to-weight ratio. The thrust-to-mass ratio has the dimension of an acceleration and expresses the acceleration that would be achieved if the thruster was only accelerating itself. The thrust-to-weight ratio is non-dimensional and expresses how many times its own weight the thruster can levitate at the surface of the Earth.

1.2.3 Energy

As the propellant is accelerated in a thruster, its kinetic energy increases. The power put into kinetic energy is

$$P_{\mathrm{k}} = \frac{1}{2}\dot{m}u^2. \tag{1.4}$$

The ratio of thrust to power (which is the same as the ratio of impulse to energy) is thus

$$\frac{T}{P_{\mathrm{k}}} = \frac{2}{u}.$$

Increasing specific impulse reduces propellant consumption but implies that more energy is needed.

Energy can be a limiting factor in two respects. There is a fixed mass for the energy handling systems which gets lumped into the thrust-to-weight ratio, and a consumable mass which decreases as the thruster is fired which gets lumped into the specific impulse of the thruster.

1.2.4 Rocket Engine Classification

The existence of specific impulse and thrust-to-weight ratio as two fundamental metrics for rocket propulsion introduces a lot of richness in the field. Indeed, there is no ideal rocket technology that is best for all missions. On one extreme, thrusters with thrust-to-weight ratios well above unity are needed to get off the Earth. On the other extreme, deep space missions, or long duration missions, can benefit from specific impulses in the tens of thousands of meters per second. Of course, many complementary metrics such as reliability and cost also exist, and further add to the richness of rocket engine technologies.

Chemical engines are no doubt the best known type of rocket engine. In these engines the propellant undergoes chemical reactions to power its own acceleration. During thruster operation, the propellant is combusted, producing a jet of hot gas, which is expanded in a nozzle to produce thrust. These engines have thrust-to-weight ratios as high as about 100, but are fundamentally limited in specific impulse by the energy density contained in the chemicals. The maximum specific impulse for chemical rockets is around 450 s for a hydrogen-oxygen rocket. Chemical engines are currently the only viable option for launch to space, but their low specific impulse makes them unsuitable for high Δv missions.

To get beyond the bounded energy density of chemicals reactions and reach higher specific impulses, the source of energy needs to be separated from the propellant. Some of the options which have been explored are nuclear energy, solar energy, and energy beamed from the ground.

In electric propulsion, the thruster is supplied with electrical energy from an arbitrary source, typically solar or nuclear. This electrical energy is used by the thruster to accelerate propellant to high velocities, for high specific impulses. There are three principal types of electric thrusters: thermal-electric, $j \times B$, and electrostatic. Thermal-electric thrusters are like chemical thrusters except that the gases are heated electrically. The $j \times B$ thrusters use the force exerted by a magnetic field on moving charges to accelerate ions. Finally, electrostatic thrusters use an electric field to accelerate charged particles.

1.2.5 Thrust Efficiency

In electric thrusters, the mass of the thrust system is usually dominated by the mass of the power supply. It is therefore important to make the best possible use of the available power. To quantify how well a thruster uses the available energy, we introduce the overall thrust efficiency

$$\eta_0 = \frac{T^2}{2\dot{m}P_{\text{in}}},$$

where P_{in} is the electric power supplied to the thruster. One reason for non-unity thrust efficiency is that the ejected propellant has a velocity distribution. To evaluate this effect, we can decompose η_0 into a kinetic and a non-kinetic component

$$\eta_{\text{nk}} = \frac{P_{\text{k}}}{P_{\text{in}}} \text{ and } \eta_{\text{k}} = \frac{T^2}{2\dot{m}P_{\text{k}}},$$

where P_{k} is the power delivered to the kinetic energy beam. η_{nk} measures how efficiently energy is delivered to the beam, and η_{k} measures how efficient the energy in the beam is for propulsive purposes. With the expressions for T and P_{k} from equations (1.1) and (1.4), the beam efficiency is unity. However, if particles have a distribution of velocities, the expressions for P_{in} and T have to be averaged over the distribution, and we get

$$\eta_{\text{k}} = \frac{\langle \boldsymbol{u} \rangle^2}{\langle \boldsymbol{u}^2 \rangle},$$

where $\langle \cdot \rangle$ denotes averaging weighted by mass flow rate, and $\boldsymbol{u}$ is the ejection velocity, which must now be represented as a vector quantity to account for dispersion of the beam. This expression can be rewritten as

$$\eta_{\text{k}} = 1 - \frac{\left\langle (\boldsymbol{u} - \langle \boldsymbol{u} \rangle)^2 \right\rangle}{\langle \boldsymbol{u}^2 \rangle}.$$

In this expression, $\left\langle (\boldsymbol{u} - \langle \boldsymbol{u} \rangle)^2 \right\rangle$ is just the variance of $\boldsymbol{u}$, so η_{k} is unity only when $\boldsymbol{u}$ has no spread. Thus, any angular spread of $\boldsymbol{u}$ or any spread in the speed at which particles are emitted will lead to reduced thrust efficiency. To maximize efficiency, we therefore need to minimize the angular divergence of the beam, and if all the particles are being accelerated by an identical electric field, we need to have as narrow as possible a distribution of charge-to-mass ratios for the emitted particles.

1.3 Electrospray Propulsion

By accelerating particles to high velocities, an electrospray can be used to produce thrust. Figure 1-2 shows an electrospray setup for thrust production. Compared with the basic electrospray source in Figure 1-1, two enhancements have been made. First, an accelerator electrode has been added so that the energy of the particles can be controlled independently of the voltage needed to extract them.[1] Second, two emitters are now present, one firing positive particles, the other firing negative particles, so that a neutral beam results [22].

[1] Our thruster will not incorporate an accelerator electrode, as there were other more critical hurdles to overcome to get a basic thruster working.

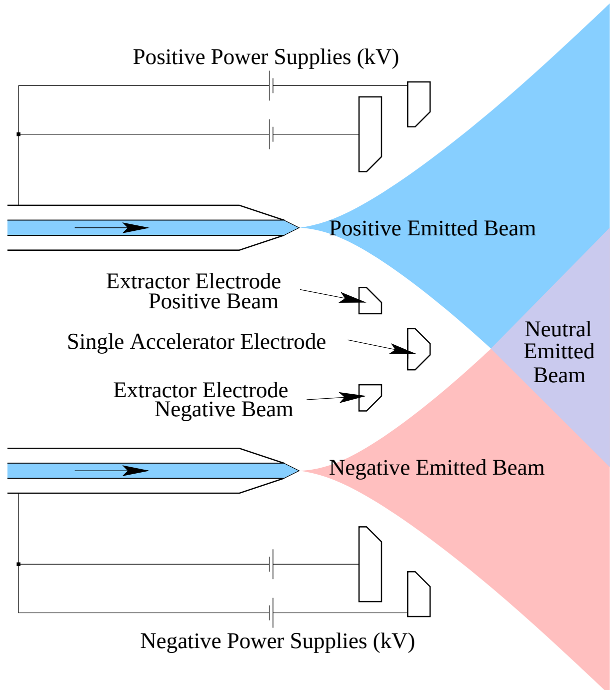

Figure 1-2: Basic elements of a bipolar electrospray thruster

1.3.1 Advantages of Electrospray Propulsion

It is difficult to make a general comparison of electrospray thrusters with other thruster types because of the variety of operating regimes that are available. Electrospray thrusters operating in the droplet regime have lower I_{sp}, but higher thrust per emitter than electrospray thrusters operating in the ionic regime. Thrusters that can tune between the ion and droplet regime via the intermediate mixed regime [23] are very promising for their flexibility. In this thesis we will focus on electrospray thrusters emitting in the ion regime with an ionic liquid as a propellant (see Section 2.1).

Figure 1-3 shows the main elements of an electrostatic thrust system. Propellant is stored in a propellant tank. When the thruster is in operation, propellant is brought to the thruster, ionized, accelerated in a region with an electric field and ejected from the thruster. Generally, only positive particles can be produced, and an electron-emitting cathode is used to neutralize the emitted beam.[2] Electrospray thrusters have

[2]If the beam were not neutralized, the spacecraft's potential promptly charge such that the beam

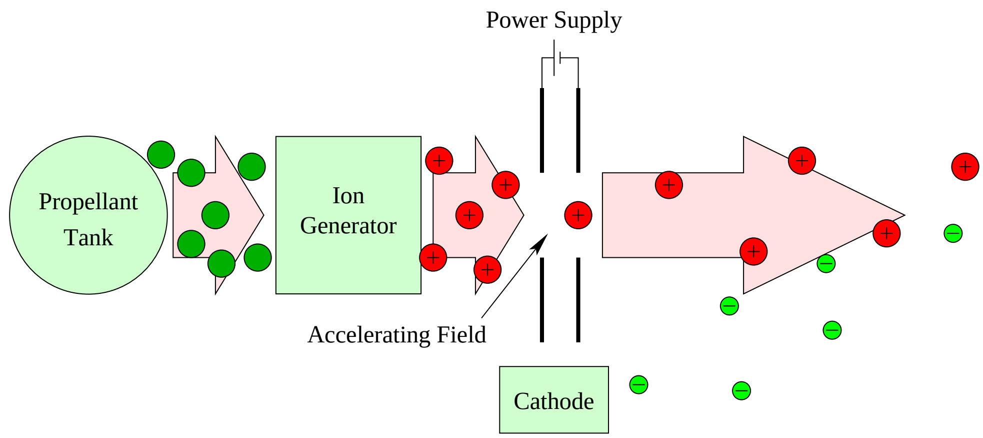

Figure 1-3: Block diagram of an electrostatic engine

excellent properties in all these steps compared with ion engines or hall thrusters.

Propellant Storage: Electrosprays use liquid propellant, leading to lighter and more compact propellant tanks than for pressurized gas storage.

Ionization: The cost of producing a charged particle in an electrospray thruster is very low, about 7 to 8 eV [24] for the propellant EMI-BF_4. All the ions are created near the tip of the Taylor cone and follow trajectories that bring them through the aperture in the extractor. In a typical ion engine, ions are created in a plasma discharge. Energy is lost when electrons collide with the anode with residual kinetic energy, and also when an ion which has been created collides with a thruster wall and recombines with an electron. The net result is that in an ion engine, the typical cost for one accelerated ion is 150 to 400 eV. The well controlled ionization in the electrospray thruster greatly reduces the energy lost in ion creation.

Acceleration: Since the ions are all created at the same location, they emerge from the extraction stage with a very narrow energy spread, around 7 to 8 eV [24]. Consequently, if power is scarce, an acceleration electrode can be used to reduce the specific impulse to the 500 to 1000 s range, down from the 3000 s to 4000 s

would be attracted back toward the thruster, negating the thrust that was generated when the beam was first accelerated.

due to the extraction voltage. The same electrospray thruster can just as easily be operated at a specific impulse of 10,000 s, just by adjusting the electrode voltages.

Neutralization: With ionic liquids or doped polar solvents, the electrospray process can emit both positive and negative particles. This eliminates the need for a cathode. All that is needed is two thrusters that are operated simultaneously, one in the positive polarity, and one in the negative polarity [22].

One major competitor of the electrospray thruster is the Field Effect Electric Propulsion (FEEP) thruster [25, 26]. It is, in effect, an electrospray thruster that operates using liquid metals, typically indium or cesium. The main advantages of the ionic-liquid electrospray thruster over FEEP technology are: reduced operating voltage because of the lower surface tension of the liquids used (1 kV instead of 5 kV or more), reduced power consumption because liquid metals need to be heated before they are liquid, possibility of operating in the bipolar ion regime, and increased flexibility because there are many ionic liquids to choose from.

1.4 Related Work

Zeleney [27] is the first to have photographed electrospray cone-jets formed from Glycerine in 1914. The mathematical description of this near perfect cone structure was given by Taylor in 1964 [1], predicting a cone half-angle of 49.3°. Also in the 1960s, the first wave of electrospray thruster development was underway. It was targeted at the main propulsion system of a spacecraft, and therefore needed operation in the highly-stressed multicone droplet regime with high voltages in the 10 to 15 kV range, to get thrusts in the millinewton range. The development of the Kaufman Ion Engine [28] which could achieve similar performance with greater simplicity led to the demise of this initial wave of development. Many key ideas still in use today were developed during this first wave: the time-of-flight apparatus for characterizing the emissions from the thruster was developed by Shelton [29] and Cohen [30], the notion of bipolar propulsion was introduced by Perel [22], the mixed ion-droplet regime was studied by Hunter [31].

The advent of mission profiles with small (100 kg) satellites, and the interest generated by electrospray in the mass-spectrometry community by Fenn [6] sparked a new wave of electrospray thruster development, starting in the late 1990s. The need for smaller thrust levels made electrosprays, which naturally operate in the nanonewton and micronewton ranges, more attractive. A review of the early stages of this revival was written by Martinez-Sanchez et al. [19]. One of the major new developments is the use of ionic liquids, whose extremely low (has not been successfully measured yet) vapor pressure avoids evaporation in vacuum. These ionic liquids can electrospray in a pure ion emission regime which is good for thrust efficiency [4]. The value of these ionic liquids was enhanced by the introduction of the externally wetted emitter by Lozano [32], inspired from work on liquid metal ion sources [17, 18], which facilitates stable operation at low flow rates. Without external wetting few ionic liquids were able to reach the pure ion emission regime [33], De la Mora *et al.* [34] report that with externally wetted emitters, a much broader class of ionic liquids are able to emit pure ions. Before the advent of the pure ion regime, the term colloid thrusters was most commonly used for these thrusters because of the fine droplets they emitted. In the ion regime, the term colloid loses its meaning, so the term "electrospray thruster", which we use, is progressively taking over.

The first mission that is being considered for electrospray propulsion is the LISA pathfinder mission, for which a thruster operating in the droplet regime is under development by Busek Co. (Natick, MA) [35]. Currently, a lot of development effort is turning toward microfabricated arrays. Velásquez-García tested a fully integrated linear array [36] operating in the droplet regime. The linear configuration allows long flow channels to be implemented, providing high hydraulic impedance. A system of clips is used to hold the extractor electrode in place in front of the emitters, with the extractor held perpendicular to the emitter substrate. A similar system of clips has been used for the present work. The extractor electrode is patterned in Tungsten on a thick silicon oxide film. Most efforts have now turned toward planar arrays. All the work except [37, 38] uses cylindrical internal capillary emitters operating in the droplet regime. Paine et al. [39,40] design an array that was not tested due to electrical breakdown problems. Heavy focus was placed on the ability to address clusters of emitters independently to throttle thrust. Recent work by Paine continues this push for throttleability [41], but given the performance of the thrusters presented in this thesis, it is not clear that independently addressable emitters or clusters of emitters are needed to achieve wide throttleability ranges. Stark *et al.* [42] have produced planar arrays of emitters which are pressure fed with EMI-BF_4. So far there are no reports of an integrated extractor electrode. Xiong et al. [43, 44] report on a single-emitter (apparently) planar thruster that should be scalable to more emitters, they report thrusts of a few μN, suggesting that they are operating in the droplet regime, with applied voltages from 1400 to 2800 V; no information is given about their fuel supply system or propellant. A glass spacer, bonded with epoxy is used to align their extractor with their emitters, no indication of how they align the two is given. Cardif et al. have reported on a planar array which uses a novel system of silicon oxide pillars to support a thin film nickel extractor [45]. They also report on preliminary time-of-flight measurements, but it is unclear if they are using their integrated electrodes for these tests [45]. Velásquez-García also pioneered a concept of externally wetted planar electrospray emitters operating in the pure ion regime [37]. These emitters are the basis for the present work.

Microfabricated electrospray arrays for FEEP thrusters have also been made by Lenard et al. [46]. They use cylindrical internal capillary emitters, like most of the microfabricated electrospray thruster concepts. A fully integrated design with individually addressable emitters is proposed, in which Pyrex with sand-blasted apertures is used for insulation.

Electrospray arrays are also being developed for applications outside space propulsion, mainly in mass spectrometry (see [47] for a review). Many linear designs exist, with emission occurring from the side of a planar substrate [48–51]. The more relevant designs for our work are planar with cylindrical internal capillary emitters (some improvements exist for sharpening the emitter tips [52]). After some early work which used non-integrated extractors [8, 10, 53, 54], Waits et al. have reported a design with integrated electrode [55]. Like Lenard et al., they use anodically bonded Pyrex for insulation, and, unlike ours, their array cannot be disassembled. Also, they suffer from large leakage currents through their insulator, which would hurt the thrust efficiency in propulsive applications.

Another area with relevant microfabrication work is the area neural implants, in which arrays of needles are used to probe neural activity in the brain. To access target cerebral areas, these probes need tall structures with sharp tips. In one particularly interesting pioneering case [56], an array of needles is produced by die-sawing a crisscross pattern into a wafer and then sharpening the resulting array of beams with a wet-etch. In some cases channels are introduced in these needles to transport neurotransmitters to the target site.

1.5 Outline

This thesis is about the fabrication of an planar electrospray thruster array. In Chapter 2, we will begin with a general discussion of electrospray design choices, before focusing on choices related to our specific thruster architecture. Our thruster presents challenges related to the externally wetted emitter architecture we are using, which will be covered in Chapter 3. We then turn to the problem of integrating an array of emitters with an extractor to make a complete thruster. Chapter 4 tackles the problem of aligning the emitter array with the apertures in the extractor electrode. We focus on methods in which the emitters and the extractor are fabricated separately, and then assembled use self-aligning hand assembly. These assembly methods all allow disassembly, which is very useful for laboratory experimentation where the emitter component may need to be inspected or replaced. Next, Chapter 5 develops one of the assembly methods into a complete electrospray thruster, focusing in particular on the problem of electrical insulation. The resulting thruster underwent testing to characterize its operation. The experimental data from the testing are presented in Chapter 6, followed by a discussion of what can be learned from these data in Chapter 7. The thesis concludes with a summary of the thesis contributions and suggested directions for future work in Chapter 8.

Chapter 2

Electrospray Thruster Design

The overall goal of this thesis is to develop technology which will allow a large array of electrospray emitters, integrated with an extractor electrode. This chapter reviews the high-level architectural choices for an electrospray thruster in Section 2.1. These choices include the mode of operation, the propellant selection, the type of propellant feed, the emitter architecture, and the dimensionality. Then, Section 2.2 focuses on detailed design questions: what substrate to use, what the important emitter geometry parameters are, what the hydraulics should be like, how the extractor is configured, and how the thruster can be electrically insulated.

2.1 Architectural Choices

2.1.1 Ion or Droplet Emission

The most general goal for an electrospray array is to produce a beam of charged particles from an electrically conductive liquid. Which particles are desired, and which properties are important depends on the application. Three stable operating regimes exist (see Figure 2-1): the *droplet regime*, in which a charged jet emerges from the tip of the Taylor cone and subsequently breaks up into droplets [2]; the *mixed regime*, in which some ion evaporation additionally takes place near the cone-jet transition [23] where high fields are present; and the *pure ion regime*, in which only ions are extracted from the liquid [4]. The ion regime occurs at the lowest flow rates, where the field is strongest, and the rate at which ions need to be evaporated is smallest.

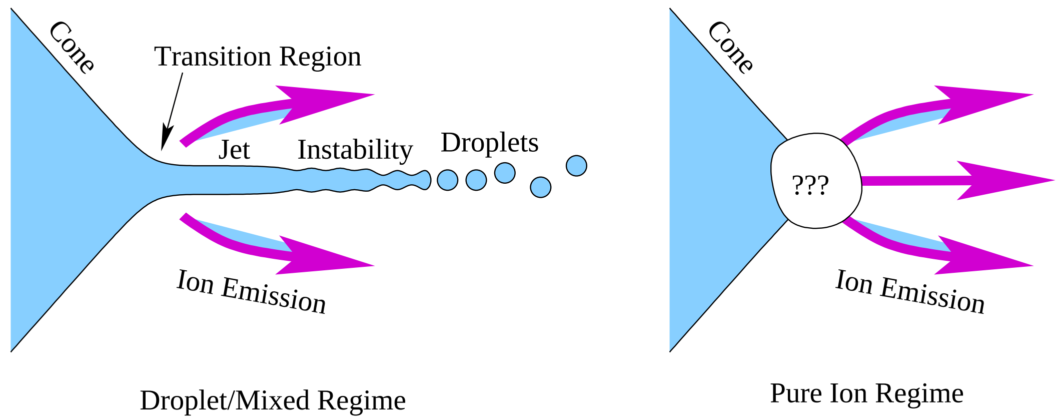

Figure 2-1: Electrospray operating regimes

We have selected the ion regime for our thruster. Indeed, the high charge to mass ratio of ions compared with droplets implies that the ion regime is the best candidate for high I_{sp} operation at low voltages. Moreover, in the ion regime, only a few comparable discrete particle mases coexist. This implies a small spread in particle velocity, and a good thrust efficiency. The drawback with the pure ion emission regime is that it is not well understood. SEMs of operating liquid metal ion sources show a jet forming with ion emission occurring in the highly enhanced fields at the tip of the jet [57]. But in Section 7.5, we will see that in our thruster, ion emission may be taking place without the formation of a Taylor cone. (This type of emission is also reported in [57].)

2.1.2 Propellant Selection

Only certain ionic liquids can achieve the pure ion emission regime we seek. EMI-BF_4 is commonly used because it is the easiest propellant to operate in the ionic regime. However, it has the drawback of containing fluorine. If electrochemical reactions take place at the emitter-liquid interface, the fluorine may etch the emitter (this is in particular the case for silicon which will be our material of choice). EMI-Im is a bit more difficult to operate in the ionic regime, but does not have any corrosive byproducts. In Chapter 6, we run tests with both EMI-BF_4 and EMI-Im, and find that our thruster is able to operate in the ionic regime with either one. These ionic liquids are also very suitable for externally wetted electrospray propulsion because their negligible vapor pressure precludes propellant evaporation when the thruster is in space.

2.1.3 Dimensionality

In moving from a single electrospray emitter to an array of emitters, a common choice is to make a linear array of emitters. Indeed, numerous technologies such as photolithography or laser cutting can make arbitrary two dimensional patterns in a planar substrate, so that complex hydraulic systems can easily be integrated with the electrospray source. A linear array can be made by cutting out needle shapes from a planar substrate and sharpening the needle tips. The same patterning method can define the hydraulics of the electrospray array. This approach has been extensively used [48–51] for mass spectrometry applications, and even for space propulsion [36].

The number of emitters can be further increased by stacking multiple linear arrays to make a two-dimensional array. However, this is a costly approach as the array cost is multiplied by the number of layers in the stack. A more promising approach is to make a two-dimensional array in which all the emitters are batch fabricated on a single substrate, and which we will call a *planar array*. This is a more challenging approach as it requires real three dimensional fabrication to make the emitters, and leaves little room for the thruster hydraulics. If these difficulties can be overcome, the advantage is that a whole array can be made from a single substrate, with a greatly reduced total effort and cost. One of the major goals of this thesis is to develop the technology needed for an integrated planar array.

2.1.4 Propellant Feed System

How liquid is fed to the emission site has a big impact on the complexity of an electrospray system. In this section we review the basic feed system architectures.

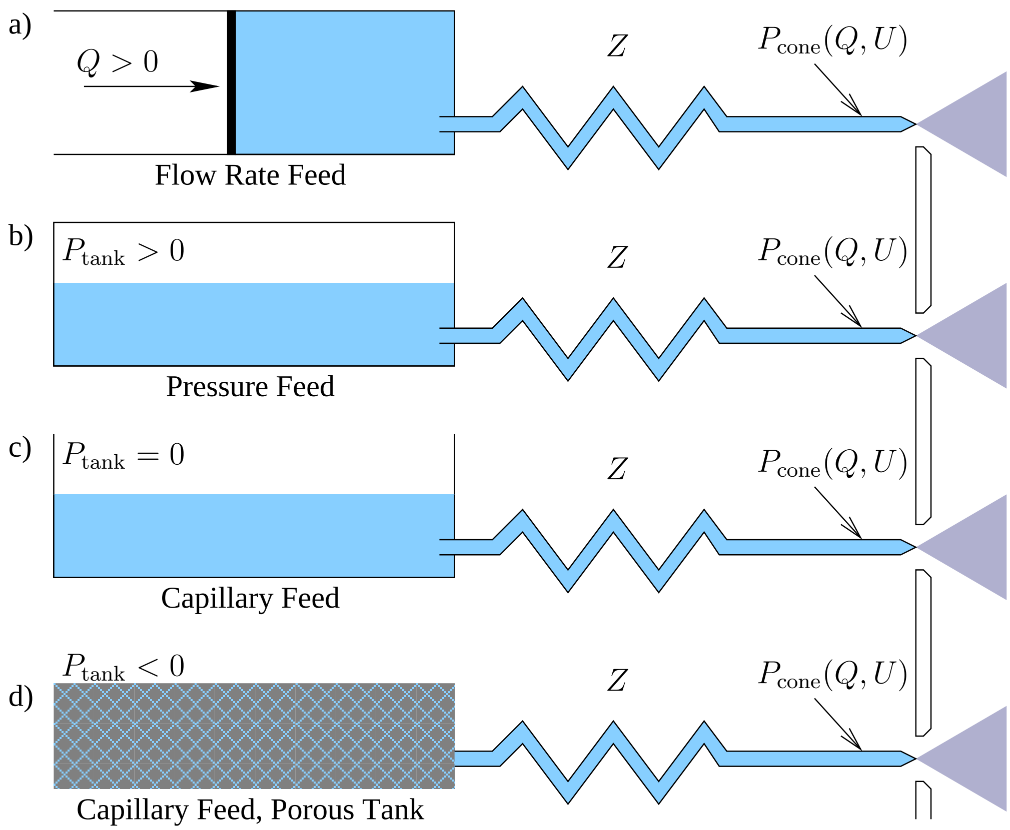

Figure 2-2: Architectures for supplying properlant to an electrospray emitter

The most direct way to operate an electrospray is to use a positive displacement pump or syringe pump to force liquid to the emission site at a desired rate Q, as in Figure 2-2.a. However, the flow rates are so low that to get steady flow it is easier to use a *pressure fed* system as shown in Figure 2-2.b [3, 23, 42]. In this configuration, the propellant tank is pressurized to a pressure P_{tank} which is large compared with the pressures $P_{\text{cone}}(Q,U)$ at the base of the Taylor cone. Thus the general formula for flow rate,[1]

$$Q = \frac{P_{\text{tank}} - P_{\text{cone}}(Q, U)}{Z} \tag{2.1}$$

can be simplified to

$$Q = \frac{P_{\text{tank}}}{Z}. \tag{2.2}$$

In this situation, the flow rate is set by the viscous flow in the capillary leading to the needle tip, and the details of the electrospray process which determine $P_{\text{cone}}(Q, U)$ can be ignored. In effect, the tank pressure directly sets the flow rate of the electrospray. The downside of this approach is that liquid can overflow from the tip of the needle. The risk is particularly high when the electrospray process is inactive, as all the liquid arriving at the tip can do is overflow. One author has suggested limiting the pressure in the tank to the pressure that can be supported by the meniscus at the needle tip [58], so that the thruster does not overflow when the electrospray is turned off.

When low flow rates are acceptable, *capillary feed* can be used (Figure 2-2.c) [59–61]. In this case, the tank is unpressurized, and it is the electric field at the emission site which creates a negative pressure to draw liquid from the tank. Because there is no pressure to make the meniscus bulge out, increasing the field on it, this configuration has higher starting voltages, for a given needle-tip diameter. Moreover, the flow rate is now set by a combination of the electrospray process and the hydraulic impedance Z of the feed system, as Equation (2.1) reduces to:

$$Q = \frac{-P_{\text{cone}}(Q, U)}{Z} \tag{2.3}$$

The reader should note that in general $P_{\text{cone}}(Q, U)$ depends on the flow rate Q so this equation does not imply that Z limits the electrospray emissions. Because liquid is no longer being pushed to the emission site, the risk of liquid overflowing is gone. The situation is analogous to a burning candle, where fuel is supplied to the flame at the same rate at which it is consumed. When working in vacuum, capillary feed greatly

[1]This formula assumes a constant impedance for the flow channel, which is a good assumption as far as we can tell, even for the externally wetted emitters. A justification for this assumption in the externally wetted case is given in Section 2.2.3.

simplifies the experiment by eliminating the pressurization system and feedthroughs. However, the dependence of $P_{\text{cone}}(Q, U)$ on Q and U is not well understood yet. Under the assumption that $P_{\text{cone}}(Q, U)$ increases with flow rate Q, capillary feed will never be able to rival pressure feed in terms of maximum flow rate, as it can only operate in the region where $P_{\text{cone}}(Q, U)$ is negative.

As we will see in Section 2.1.5, some emitter architectures transport liquid to the emission site via a porous material. This material can be seen as a small porous propellant reservoir. As depicted in Figure 2-2.d, this small reservoir can be naturally extended by a porous propellant tank leading to a system with very simple hydraulics. In this case the porous propellant tank simply needs to be contacted to the porous emitter surface, eliminating all piping and potential clogging points from the system. The pressure of the propellant supply P_{tank} is now negative and is determined by the scale of the porosity in the propellant tank. This leads to a reduced flow rate which is set by Equation (2.1). Hybrid systems are also possible. For example, a porous reservoir could be periodically resupplied from a pressurized reservoir.

2.1.5 Emitter Type

Figure 2-3 shows three common electrospray emitter architectures. The *internal capillary emitter* architecture is most common. In this configuration, emission occurs from the tip of a hollow needle. Liquid flows through the needle to the emission site. Internal capillaries are commonly combined with a pressure fed propellant supply. In addition to the risk of overflowing the thruster, this architecture has the risk that a single bubble can clog the capillary, shutting down the emitter, or causing a messy spurt when the bubble reaches the emission site.

Internal capillary emitters can also be used with capillary feed. This configuration leads to much smaller flow rates and has been known in the mass spectrometry community as nanoelectrospray [60, 61].

For vacuum applications such as space propulsion, capillary emitters have the benefit of preventing liquid evaporation, except for the small exposed area at the emitter tip. With the advent of ionic liquids, which for practical purposes do not

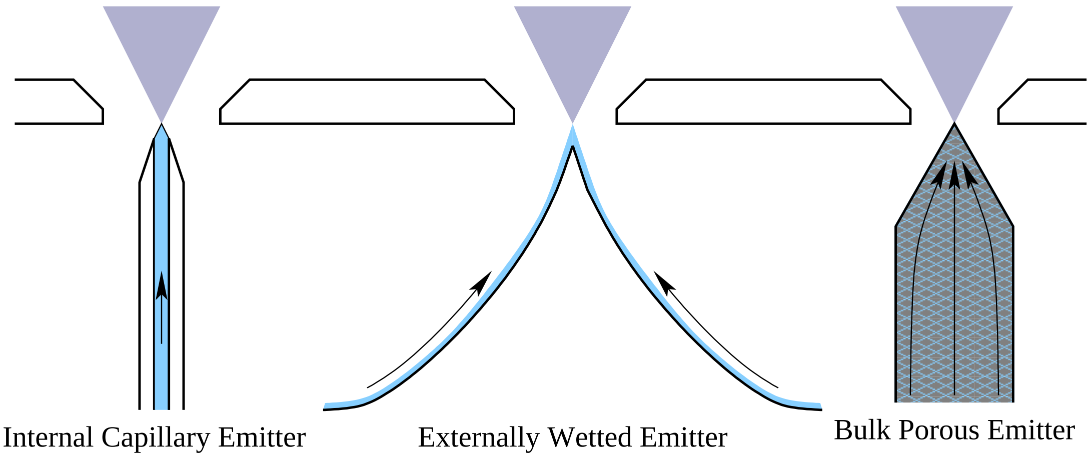

Figure 2-3: Electrospray emitter architectures

evaporate, shielding the liquid from vacuum is no longer necessary. This has enabled *externally wetted emitters*, as depicted in Figure 2-3, to be considered for space propulsion. In this architecture, liquid is transported to the emission site along the roughened surface of the emitter. This configuration was also used for liquid metal ion sources [17, 18] and FEEP thrusters [25] and was introduced to the electrospray thruster community by [32] as an easy method to obtain pure ion emission from ionic liquids. This configuration has some major advantages over capillary feed. It is impossible for liquid to overflow from the emitter, as its whole surface is inherently wet. Bubbles can leave via the exposed surface with less disturbance to the flow. There is a large number of redundant flow paths making clogging unlikely. The tip diameter is no longer determined by the diameter of the capillary, making lower starting voltages possible. Recent work suggests that it is easier to achieve the pure ion emission regime with externally wetted emitters [34], perhaps because the smaller tip allows stable operation at lower flow rates. The major drawbacks are that vacuum operation is only possible with low vapor pressure liquids, and that now, the emitter surface must be treated in a way that allows it to transport liquid. As we shall see, this surface treatment consists of roughness at the micron scale which is difficult to produce in a repeatable manner, though good control of this roughness is critical as it determines the hydraulic impedance of the feed system.

For externally wetted emitters, pressure feed is not applicable as pushing liquid to the emitter will cause it to flood. Therefore, capillary feed must be used.

A final architecture that has been considered is the *bulk porous emitter* architecture [62–64]. It has all the attributes of the externally wetted emitter except that the whole emitter is porous, offering a reduced hydraulic impedance to the flowing liquid. This could lead to higher emission per emitter.

In this thesis, we will be working with the externally wetted architecture. Indeed, we are targeting ion emission, and it is easier to produce ion emission with this type of emitter than with an internal capillary emitter. Moreover, the internal capillary emitter needs some flow impedance for each emitter to avoid hydraulic interactions between the firing emitters. With the internal capillary designs, providing this impedance means providing long channels for which there is little room in a planar array. The bulk porous emitter architecture looks like an excellent candidate for future work. We did not select it because unlike the external architecture, for which a proof of existence emitter fabrication scheme existed already [37], the fabrication of bulk porous emitters would have been a completely open problem. Instead, we preferred to focus the innovation in the extractor electrode packaging problems, and in better understanding and implementation of the existing external architecture.

For simplicity, we will use the surface of the emitter array as our only propellant tank. In Section 7.1.4 we briefly mention how a larger propellant supply could be integrated with our thruster.

2.2 Implementation Choices

2.2.1 Substrate

In this work we have chosen the silicon wafer as our building block. This choice has been dictated by the availability of microfabrication technology for silicon, inherited from the electronics industry. The emitter densities presented in this thesis could be achieved with other manufacturing technologies, but if we are to scale to smaller dimensions, nothing will rival the capabilities of silicon microfabrication.

One caveat with the choice of silicon is that it is susceptible to chemical etching, for example by fluorine which is present in EMI-BF_4. Indeed, during thruster operation, a current flows from the emitter to the liquid (and then to the plume). Since carriers cannot cross the emitter to liquid boundary, charges build up at the interface forming a charge double layer. With continued passage of current, the voltage across the double layer increases, and once a few volts are reached, electrochemical reactions can take place, possibly releasing atoms like fluorine which can damage the emitter, when they are present in the propellant. These electrochemical processes can be avoided by alternating the polarity of an emitter faster than the voltage necessary to start the reaction can build up [65]. If this precaution is not deemed sufficient, the silicon emitters can also be coated in a more chemical resistant materials such as platinum, once they have been fabricated.

2.2.2 Emitter Geometry

The externally-wetted electrospray emitter has two main functions: field enhancement and liquid transport. These two functions are usually antagonistic since field enhancement is best with sharp slender emitters, whereas, as we shall see in the next section, liquid transport is improved when there is a larger surface area through which to transport the liquid. One option to improve flow without hurting field enhancement, is to create groves on the side of an emitter to provide a low impedance channel for bulk liquid flow. In Section 3.2, we present some emitter geometries that try to provide this low impedance channel, but which we were unfortunately unable to test.

We base our design on [37]. In this design, emitters are pencil-shaped, with a slender 100 μm diameter shank, capped by a conical tip. The height of a pencil is about 400 μm tall. In Chapter 3 we review this design, and present the emitter geometry that we finally selected, which resembles a triangular pyramid. The main parameters for the emitter geometry are: tip radius of curvature, tip half-angle, shank diameter, and emitter height, and emitter depth.

2.2.3 Hydraulics

Liquid at rest on a surface

If a small amount of liquid is at rest on a smooth surface, it forms a spherical drop under the influence of surface tension [66].[2] Surface tension is a force that tries to reduce the surface of the interface between two phases. Taking any fragment of the surface, there is a force acting on the boundary of that surface that tends to reduce the area of the fragment. The force per unit boundary length is the surface tension γ. Equivalently, γ can be seen as the surface energy of the surface. Each unit area of the surface has an energy γ. The surface tension is acting to reduce the total surface, thus reducing the total surface energy.

For the drop we are considering, surface tension acts at two places. On the free surface of a phase, surface tension relates curvature and change in pressure across the interface

$$P = \gamma \left(\frac{1}{R_1} + \frac{1}{R_2} \right), \tag{2.4}$$

where γ is the surface tension of the interface (in our example this is the liquid-vapor interface), P is the pressure in the drop, and R_1 and R_2 are the principal radii of curvature of the drop. Because P is uniform in the static drop, we find that the curvature is uniform on the surface of the drop.

The second place where surface tension acts is at the contact line where three phases meet. Here, the three phases are the solid surface, the liquid, and the ambient gas (or vacuum). Along this line, three surface tensions are in competition: γ the liquid-vapor surface tension, γ_{SV} the solid-vapor surface tension, and γ_{SL} the solid-liquid surface tension (see Figure 2-4). The contact line can move along the surface, but not leave it since the solid is rigid. Thus, equilibrium is reached when the tangential components of the surface tensions cancel out. These three surface tensions are in equilibrium when the liquid meets the surface at an angle θ_{c} such that

$$\cos(\theta_{\mathrm{c}}) = \frac{\gamma_{\mathrm{SV}} - \gamma_{\mathrm{SL}}}{\gamma}. \tag{2.5}$$

[2] The assumption that the amount of liquid is small means that we can neglect gravity.

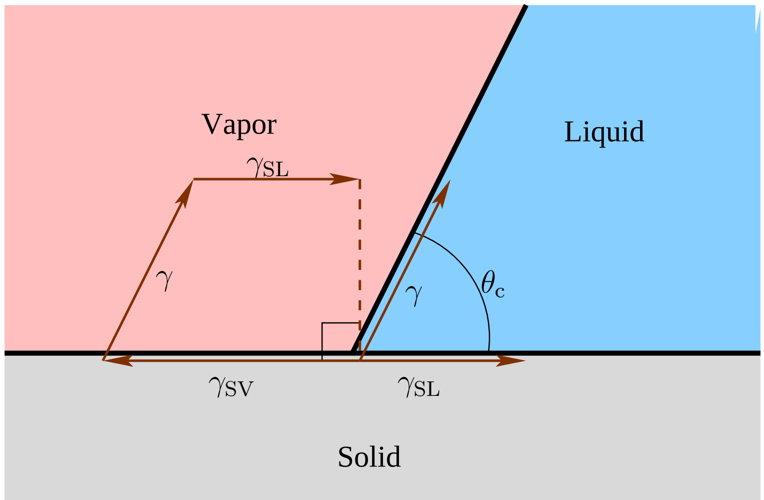

Figure 2-4: The contact angle at a contact line is set by balancing the tangential components of the three surface tensions

If $\theta_c < 90°$ then $\gamma_{SV} > \gamma_{SL}$, so the surface energy of the solid surface is reduced when the liquid covers it. We say that the solid is wetting.[3] In general, though, the liquid does not completely cover the solid because to do so it would have to create a large liquid-vapor surface, so the net energy would be increased. If, however, $\gamma_{SL} + \gamma < \gamma_{SV}$, then the advantage of wetting the solid is so great that the liquid tends to completely cover the solid surface. In that case the solid is said to be *completely wetting*. Conversely, when $\theta_c > 90°$, the solid surface prefers to remain dry, and it is said to be non-wetting. Nevertheless, a small amount of solid-liquid contact still does occur because this can reduce the liquid-vapor interface. If $\gamma_{SV} + \gamma < \gamma_{SL}$ then the surface is completely non-wetting. Silicon is wetting for the propellants we want to work with, so we will now focus on the wetting case.

On non-flat surfaces, the situation is more complex. For example, liquid will tend to penetrate cracks on a wetting surface because it can wet the solid without creating a liquid-vapor interface. Generally speaking, liquids on wetting surfaces will be drawn away from convex regions where it takes a lot of liquid surface to cover a solid surface, in favor of concavities where a lot of solid surface can be covered with minimal liquid

[3]Strictly speaking, the property of being wetting depends on which liquid is being considered, but this is the common terminology.

free surface. Moreover, all connected components of liquid will have uniform pressure and curvature at equilibrium.

The emitters we are considering have positive (convex) curvature over their whole area. The base, on the other hand, is flat, and the corner between the emitters and the base has negative curvature (as long as there isn't too big of a fillet between the needle and the base). Thus, it is unclear how the uniform film of liquid covering the externally-wetted emitter in Figure 2-3 is possible. It has to have uniform curvature by Equation (2.4), but has to have varying curvature to follow the surface. Figure 2-5 shows a more likely liquid configuration on these surfaces.

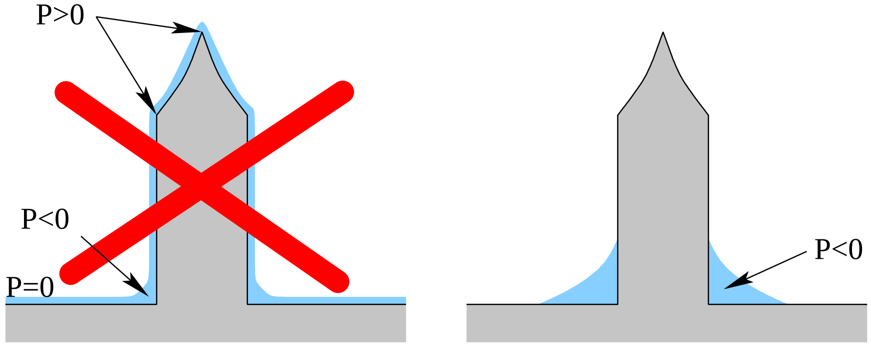

Figure 2-5: Liquid accumulates at the base of a needle rather than form a uniform film on the needle

The case of complete wetting is strange because it is energetically favorable to cover the whole surface, but, were the whole surface covered, the curvature of the liquid surface would not be uniform. This paradox is resolved by realizing that Equation (2.4) will force areas with high positive curvature to have a thin liquid film of molecular dimensions, which invalidates the continuum mechanics results we have been using.

Liquid on a spiked surface

How then can we generate the liquid film depicted in Figure 2-3? The key insight is that we do not need the film to cover the whole surface. It just needs to cover most of the surface. Indeed, consider a target surface which has been covered in a dense pattern of tall thin cylinders (see Figure 2-6). (By dense, we mean that between two cylinders, the orientation of the surface does not change much, and the distance between cylinders is much smaller than the height of a cylinder.) The forest of cylinders occupies some volume. Place a drop with a slightly smaller volume on the surface, and assume that the cylinders and the target surface are wettable with a low contact angle. There is enough liquid that the bases of all the cylinders will be covered, but there is not enough liquid for any of the cylinders to be completely covered. The radius of curvature of the liquid surface is comparable to the spacing between cylinders because its surface has to connect neighboring cylinders, and must contact those cylinders at the liquid's contact angle. This liquid surface is the desired approximation of the target surface. Each cylinder can be seen as a singularity where the liquid surface accumulates just the right amount of positive curvature to counterbalance the negative curvature of the liquid surface. The average curvatures of the target surface and the liquid surface end up being the same at a scale of a few cylinder spacings.

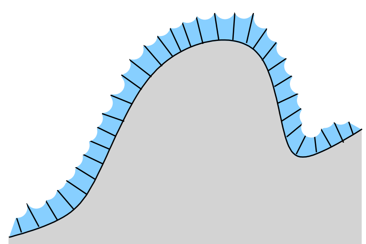

Figure 2-6: A forest of spikes on a surface allows a liquid film to form on an arbitrary surface

This construction suggests that if we are able to cover our emitters with little spikes, we should be able to cover them in a thin film of liquid that resides below the tops of the spikes. The spikes can be tapered a bit, but their sidewalls must be sufficiently steep that liquid contacting a spike at its contact angle is angled towards the base of the spike. Otherwise, surface tension is trying to push the liquid toward the base of the spike, and it is energetically favorable for the liquid to form a drop somewhere on the liquid surface.

This angle constraint can be seen by a simple energetic argument. Consider a surface covered in conical spikes with half-angle α. Seen from far away, the surface looks smooth. However, the area of the surface is greater than the area of a smooth surface a factor $\frac{1}{\sin(\alpha)}$. Therefore, from far away, the rough surface looks like a smooth surface with effective surface energies $\gamma'_{\mathrm{SV}} = \frac{\gamma_{\mathrm{SV}}}{\sin(\alpha)}$ and $\gamma'_{\mathrm{SL}} = \frac{\gamma_{\mathrm{SL}}}{\sin(\alpha)}$. By Equation (2.5), the contact angle θ'_{c} of the surface is given by

$$\cos(\theta'_{\mathrm{c}}) = \frac{\cos(\theta_{\mathrm{c}})}{\sin\alpha}.$$

When $\alpha < 90° - \theta_{\mathrm{c}}$, the effective contact angle θ'_{c} is undefined, and the surface is perfectly wetting. In the opposite case, drops with a finite contact angle will form on the surface.

In the case where the perfectly wetting condition is satisfied, what happens when a drop is placed on the surface? This situation is similar to the situation of a drop spreading placed on a porous surface which has been extensively studied [67, 68]. The drop spreads to some macroscopic contact angle, while in parallel liquid seeps into the porous surface. If the surface is large enough, the drop can be completely absorbed by the surface. Garza, also did some characterization and modeling of drop spreading speeds on a roughened silicon surface for our group [69, 70], and observed a halo spreading around the drop on the liquid surface (halos were also observed in [58]), which corresponds to the liquid propagating below the level of the surface roughness (see Figure 2-7). The existence of this halo is what is needed to create a liquid film suitable for transporting propellant to the tip of an electrospray emitter. However, determining how fast this halo actually spreads is not directly relevant for electrospray operation. Indeed, the spreading speed of the halo is determined by how liquid is able to creep from a wetted area into a neighboring area that is not yet wet, a process fraught with subtlety, as it involves motion of the liquid contact line [71,72]. However, when the electrospray is operated, the emitter surface is already covered in liquid and the problem reduces to a conceptually much simpler viscous flow problem.

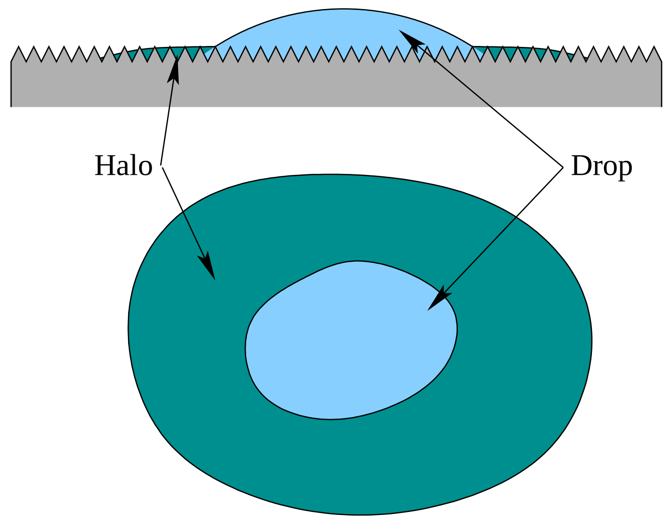

Figure 2-7: Halo surrounding a liquid droplet spreading on a porous surface

Scale of the surface roughness

We now turn to the problem of viscous flow along a spiked surface. Much insight on this problem can be gleaned from the much simpler study of flow in open capillaries [73], as this is just a two-dimensional variant of the same problem.

For a moment, let us assume that the flow of liquid along the spiked surface can be described by a two-dimensional version of Darcy's law

$$\boldsymbol{q_s} = -\frac{\kappa_{\mathrm{s}}}{\mu}\boldsymbol{\nabla} P, \tag{2.6}$$

where q_s is the surface flow rate density the fluid (volume crossing a unit length of surface per unit time), κ_{s} is the surface permeability of the surface, which has the dimensions of a volume, and μ is the viscosity of the liquid. For this formula to hold, we need to be at low Reynolds numbers which is usually the case at the small flow rates and scales we are considering. Moreover, the free surface of the liquid needs to be stationary, or the cross-section that is available for liquid to flow through will vary. The shape of the liquid surface is determined by the local pressure and so in a more general version of Equation (2.6), the surface permeability would be a function of P.

How much does P need to change before κ_s changes significantly? Consider a spiked surface on which spikes are separated by a distance ϵ, and which is full of liquid (i.e., the liquid surface is nearly flat). To significantly change the liquid surface, the radius of curvature of the liquid must be comparable to ϵ. From Equation (2.4), the pressure is thus $P \sim \frac{\gamma}{\epsilon}$. Assuming an unreasonably large spacing of $\epsilon = 100$ μm, and the ionic liquid EMI-BF_4 for which $\gamma = 0.052$ N/m at room temperature [74], we get $P \sim 520$ Pa. This pressure is large compared with pressures involved in the electrospray process which are typically in the tens of pascals range [75], and will get even larger as ϵ is decreased to more reasonable scales. We conclude that for a surface that is saturated in liquid the assumption that κ_s is constant is a good one, and Equation (2.6) is valid.

To design the spiked surface, we need to determine what values of the surface permeability κ_s are tolerable, and what spike separation ϵ will yield that value of κ_s. To estimate κ_s, we consider a $z_2 = 400$ μm high conical emitter with a $\theta = 30°$ half-angle shown in Figure 2-8, in which liquid is flowing to within $z_1 = 1$ μm of the tip, and will require that for a 1 μA emission of EMI-BF_4, the pressure drop across the emitter is at most 10 Pa, a pressure difference which is known to cause significant changes in emission [75].

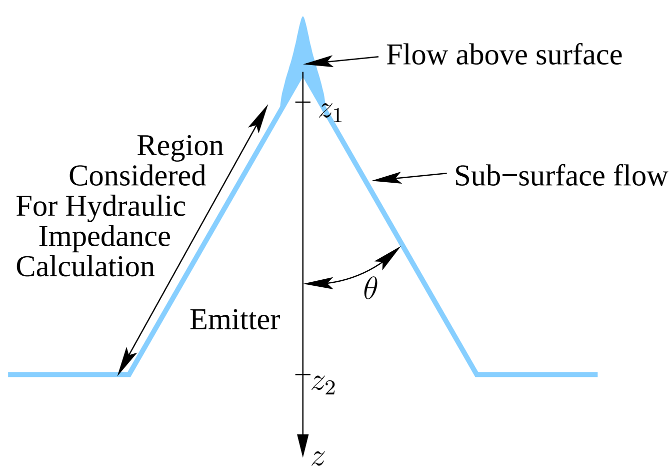

Figure 2-8: Geometry used for calculating hydraulic impedance of emitter

The flow rate can be deduced from the current by

$$Q = \frac{I \langle M \rangle}{\mathcal{N} e \rho}, \tag{2.7}$$

where $\langle M \rangle \sim 0.2$ kg/mol is the average molar mass of the emitted particles, $\mathcal{N} = 6.02 \cdot 10^{23} \mathrm{mol}^{-1}$ is Avogadro's number, $e = 1.6 \cdot 10^{-19}$ C is the elementary change and $\rho = 1300$ kg/m^3 is the density of EMI-BF$_4$. The flow rate we are considering is thus $Q \sim 1.6 \cdot 10^{-15}$ m^3/s. At 10 μm below the tip, assuming a spike height of 1 μm, this corresponds to an average velocity of 44 μm/s, and a Reynolds number of $1.1 \cdot 10^{-5}$, which confirms that we are well in the viscous flow regime.

Applying Equation (2.6) to the cone geometry, the pressure drop across the emitter is

$$\begin{aligned} \Delta P &= -\frac{\mu}{\kappa_{\mathrm{s}}} \int_{z_1}^{z_2} \frac{Q}{2\pi z \sin(\theta)} \mathrm{d}z \\ &= -\frac{\mu Q \ln(\frac{z_2}{z_1})}{2\pi \sin(\theta) \kappa_{\mathrm{s}}}. \end{aligned} \tag{2.8}$$

Using $\mu = 0.0052$ Pa $\cdot$ s, we get $\kappa_{\mathrm{s}} \gtrsim 1.5 \cdot 10^{-18}$ m^3.

To estimate κ_{s}, we consider that the flow is running through semi-circular channels with radius ϵ, and recall that for Poiseuille flow [66] in a cylindrical capillary with length L,

$$Q = \frac{\pi \epsilon^4 \Delta P}{8 \mu L}. \tag{2.9}$$

Comparing with Equation (2.6), by substituting $\frac{\Delta P}{L} = \boldsymbol{\nabla} P$ and $Q = 4\epsilon q_s$ (one cylindrical capillary is equivalent to two open half-capillaries, occupying a width of 4ϵ, as in Figure 2-9), we get

$$\kappa_{\mathrm{s}} \sim \frac{\pi \epsilon^3}{32}. \tag{2.10}$$

Taking the previously computed value of κ_{s}, we get $\epsilon \sim 2.5$ μm.

Thus our goal will be to make a surface with micrometer-scale spikiness. Because of the logarithmic dependence on emitter height in Equation (2.8), the height of the emitter will not change this result.

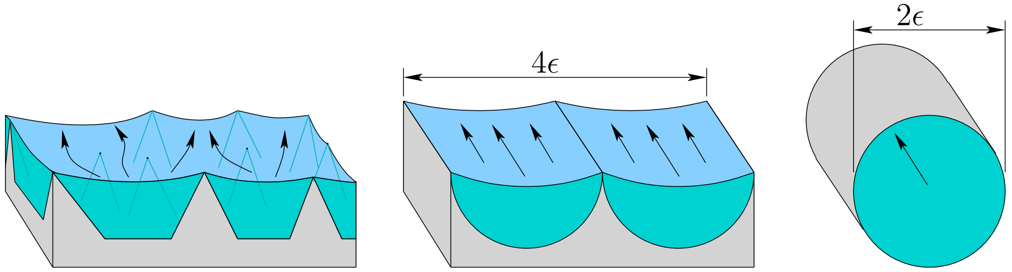

Figure 2-9: Approximate reduction of surface flow to Poiseuille flow

2.2.4 Extractor Configuration

In Figure 1-2, we presented a thruster with two extractor electrodes (one per polarity), and one accelerator electrode. The extractor electrode needs to be placed close to the emitters to provide the high fields necessary to start emission. On the other hand, the accelerator electrode does not need to be made to the same tight specification because it can surround clusters of emitters. Therefore, we will focus our efforts exclusively on the extractor electrode in this thesis.

The most natural extension from a single electrospray emitter to an array is to have an extractor with one circular hole per emitter. However, as we shall see, one of the limiting factors on emitter packing density is the size of the opening that is needed in the extractor to avoid beam impingement. By merging a row of holes into a slot, we no longer need to worry about beam impingement in the direction of the slot, and emitters can be packed much more densely in that direction (see Figure 2-10). The spacing between slots has the same constraints as the spacing between holes so this approach only improves density in one direction. This gain in density will certainly come at the cost of a reduction in electric field and hence an increase in startup voltage. By using extractor slots instead of holes we were able to nearly double the emitter density from 2.5 tips/mm^2 in [8] to 4.4 tips/^{2} in the present work.

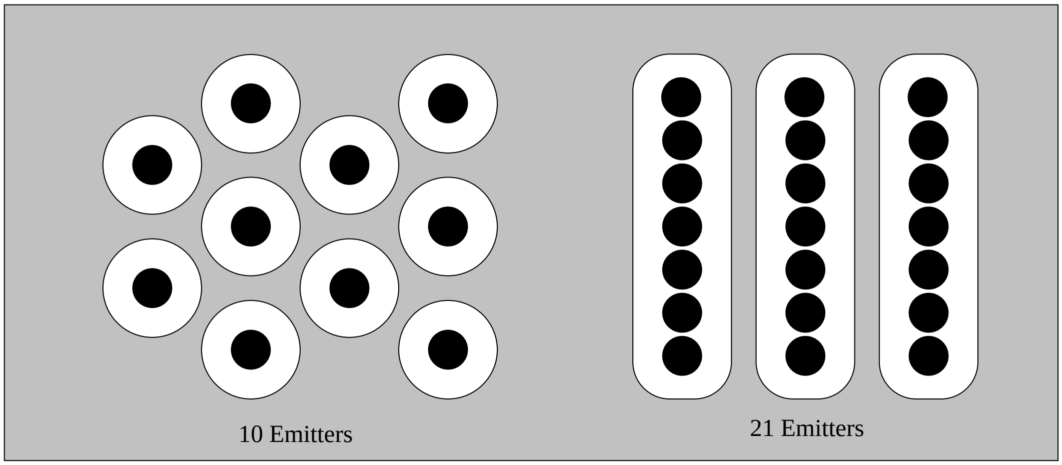

Figure 2-10: Extractor slots allow greater packing density

Extractor Slot Width

To maximize emitter density, the extractor slots should be as narrow as possible. They cannot be too narrow, however, or the emitted beam will impinge on the electrode, which will waste energy and propellant, and cause erosion of the extractor.

Figure 2-11 shows the situation with an extractor with a slot spacing of a, in which neighboring slots are separated by b, and for which the emitters are located at a level h below the top of the extractor. If the emitted beam has a half-angle θ_{beam}, the non-impingement condition is

$$\frac{a-b}{2h} > \tan(\theta_{\text{beam}}). \tag{2.11}$$

This equation means that the extractor needs to be as thin as possible, and that the emitters should be as close as possible to the height of the top of the extractor (though this will have an impact on the electric field).

For our design, we selected a=750 μm, and b=200 μm, which allows just over 325 μm for h for a conservative beam divergence θ_{beam}=40° [32]. As we shall see, this leaves just enough room for the extractor thickness and the depth of the emitter tips (see Section 3.4.2).

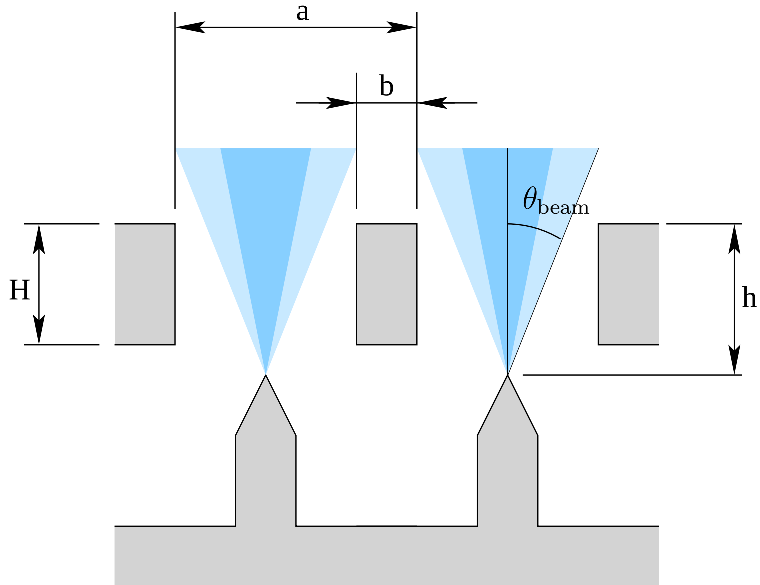

Figure 2-11: Parameters of the extractor slot

Extractor Thickness

During thruster operation, the extractor electrode experiences electrostatic forces that attract it to the emitter substrate. We estimate the extractor thickness that is required to limit the deflection that the extractor experiences.

First, we evaluate the electric pressure on the extractor. For this purpose, we consider the extractor and the emitter substrate to form a parallel plate capacitor. The Electric field between the plates is $\frac{V}{d}$ where V is the potential difference between the electrodes, and d is the distance between them. The electric pressure on the electrode is thus

$$\frac{P_{\text{el}} = \epsilon_0 V^2}{2d^2}. \tag{2.12}$$

Next we calculate the deflection of the beams that make up the extractor. We model it as a fixed-fixed beam with thickness H, length L and Young's modulus E, undergoing a uniform pressure P, to find the maximum beam deflection δ [76].

$$\delta = \frac{PL^4}{32EH^3} \tag{2.13}$$

The extractor is made up of beams with spacing a and width b. Assuming, conservatively, that the pressure P_{el} from Equation (2.12) gets concentrated into the area where a beam is actually present, we get

$$P = \frac{a}{b} P_{\text{el}}. \tag{2.14}$$

Combining equations (2.12), (2.13) and (2.14), we get a conservative formula for the beam displacement.

$$\delta = \frac{\epsilon_0 V^2 a L^4}{64 d^2 b E H^3} \tag{2.15}$$

How much displacement is too much? Too much displacement will change the electric field, so δ should be much smaller than the extractor to emitter distance. Moreover, δ should be less than $\frac{d}{3}$, which is the displacement at which the pull-in instability sets in, and the beam gets pulled down to the underlying substrate [77]. Generally, this pull-in condition is subsumed by the small-electric-field-change condition. Finally, δ should be small enough that the beam does not yield.

To make the extractor easy to fabricate, we decided to set its thickness no lower than 200 μm. This is to avoid excessively weakening of the wafer, and to allow nearly open-loop control of the etching. With the extractor sizes selected in the previous section, the L=1.6 cm thinned extractor region we will have in Chapter 5, an emitter base to extractor distance $H = 200$ μm, an extraction V voltage of 5000 V, a spacing d of 200 μm, and a Young's modulus E for silicon of 145 GPa [77], we compute from Equation (2.15) a conservative deflection of $\delta = 18$ μm. This deflection is acceptable, and probably much larger than what we will see in practice, given the conservative assumptions that led to it.

To make sure the extractor does not fail, we need to check that the sheer stress and moments in the beam are not excessive. The sheer stress is given by [76]

$$\tau = \frac{PL}{2H} = \frac{a\epsilon_0 V^2 L}{4bd^2 H} \simeq 0.41 \text{ MPa}, \tag{2.16}$$

which is negligible for silicon which fails around 1 GPa [78]. The maximum moment in a fixed-fixed beam is at its endpoint. The moment causes a stress [76]

$$\sigma = \frac{PL^2}{2H^2} = \frac{a\epsilon_0 V^2 L^2}{4bd^2 H^2} \simeq 33 \text{ MPa}, \tag{2.17}$$

which is also small for silicon (in fact, since $L > H$, we only ever need to check the moment condition as it is L/H times greater than the shear stress condition).

2.2.5 Electrical Insulation

With current designs, electrospray emission occurs at voltages from 1 to 3 kV. These voltages causes a formidable insulation problem. The small spacings, in the range of hundreds of micrometers, lead to a risk of flashover if the distance between electrodes along the insulator surface is too short [79]. Liquid present on the thruster migrates under the influence of the electric field to areas where the field is stronger. With poor design this effect can bring the conducting liquid to the insulators, leading to a risk of short-circuiting. In particular, if the insulators are porous (which is typically the case for sintered ceramics), the liquid permeates the insulator and promptly closes the circuit. Finally, the insulators are at risk from material emanating from the electrospray process, particularly when transients cause unsteady operation. Each one of these risks can be addressed in the thruster design.

Dry Insulation

For dry insulation, the solutions we have found empirically are:

- Always ensure a long length (millimeters) along the insulator surface between the two electrodes.
- Use ceramic or glass rather than plastics. Plastics like polyimide or parylene may have good dielectric strength, but they fail to prevent flashover.
- Laser-cut Pyrex side-walls seem very susceptible to breakdown.

A description of the experiments that were carried out to reach these conclusions appears in Section 5.1.3.

Wet Insulation

The thruster must be designed so that liquid does not spread to the insulator regions. During thruster operation, large electric fields act on the liquid that is present on the area surrounding the emitters. These forces can cause the liquid to migrate around the thruster in unexpected ways that are not revealed without the application of voltage. Two main effects can cause this migration.

Non-uniform electric field: An electric field applies an outwards pressure of $\frac{1}{2}\epsilon_0 E^2$ to a conductive surface. For a non-uniform electric field, this pressure is non-uniform, and can therefore cause a liquid that was formerly at equilibrium to reconfigure itself.

Electrowetting: Consider a conductive drop on a conductive surface. In the presence of an electric field, a surface charge forms on the free interface, and, to conserve charge, this causes a charge double-layer to form between the liquid and the conductive surface. The energy stored in that double layer reduces the liquid-solid surface energy and reduces the contact angle. This causes the liquid to spread more than it did without the field present. This phenomenon is called electrowetting [80, 81].

Some simple measures can avoid getting liquid on the insulators because of these effects.

- Provide non-wetting surfaces for the liquid to cross before it reaches insulators. Failing that, at the very least, surfaces leading to the insulators should not wick liquid.

- Force the liquid to pass through areas of low electric field before it can reach the insulators. The electric field will drive the liquid away from these regions.

- Provide traps where liquid can accumulate.

Deposition

To avoid deposition of stray emitted particles on the insulators, the strategy we have adopted is to use shadowing: if there is no line-of-sight path from the particle emission sites to the insulators, then the insulators will be protected. In the design stage, it is easy to provide features that will serve as shields for the insulators.

2.3 Conclusion

In this chapter, we have outlined the major design choices for an electrospray thruster. Our thruster will be a planar array of capillary fed, externally wetted emitters operating in the pure ion emission regime with EMI-BF_4 or EMI-Im as propellants.

The thruster will be microfabricated out of silicon and use the emitter geometry pioneered by Velásquez-García [37]. Calculations suggest that these emitters will need to be coated with a micron scale roughness if we do not want to have excessive hydraulic impedance. Chapter 3 will delve into the details of emitter formation.

We have decided not to provide an accelerator electrode in our thruster, greatly simplifying the design. Moreover, the accelerator electrode does not need to be microfabricated because it can surround clusters of emitters, and could therefore be added later. To maximize packing density while minimizing the risk of beam interception by the extractor electrode, a slotted configuration has been selected for the extractor. This configuration allows tight packing of emitters in the slot direction without taking any risks in selecting the dimensions of the slots. We have picked conservative dimensions for the slots, which will be 550 μm wide, with 200$\times$200 μm cross-section beams between them. In Chapter 4, we will design an assembly system which positions the emitters under the extractor slot with a precision in the ten micrometer range. The challenges of electrical insulation have been briefly described, and strategies for solving them proposed. We will apply these strategies in Chapter 5 when completing the thruster design.

Chapter 3

Emitters

The heart an electrospray array is its emitters. In fact, many groups focus all their effort on fabricating and characterizing emitters without moving to the stage of integration with the extractor electrodes [10, 42]. This chapter focuses on externally wetted silicon emitters, based on the work of Velásquez-García [37], extended with an improved understanding of the emitter formation process. The main goal is to have a viable electrospray source to integrate into the thruster that is presented in Chapter 5.

In this chapter, we will first consider how to make emitters with a desired geometry in Section 3.1. Then we will study the surface treatment which allows liquid transport along the surface in Section 3.3. Unfortunately, the surface treatment we obtained cannot treat arbitrary geometries, which forces us to use a compromise geometry as described in Section 3.4.

3.1 Emitter Geometry

3.1.1 Overview of SF_6 and DRIE Etching

Because of its repeatability, uniformity and versatility, plasma processing is one of the staples of the semiconductor industry [82]. In a plasma process, a plasma is generated in a chamber containing the sample to be processed. Depending on the gases that are present and the plasma conditions, material can be deposited on the sample or etched away. By subtle combinations of deposition and etching, features can be etched with a high degree of control over the sidewall profile (see [83] for an interesting review).

Two principal mechanisms of etching coexist. In chemical etching, reactive ions are generated in the plasma. These ions react with the atoms of the sample, and the gaseous byproducts are carried away. In physical etching, the impact of high energy ions from the plasma causes sample atoms to be sputtered away. Chemical etching is mainly isotropic, whereas physical etching is directional because of the highly directional velocity distribution of the ions that have been accelerated in the plasma sheath. Depending on the electrical bias between the substrate and the plasma, and the species present in the plasma, one or the other of these processes will dominate.

Sulfur hexafluoride (SF_6) is a commonly used gas for silicon etching. In the plasma, SF_6 produces F^- radicals that chemically etch the silicon by producing the volatile SiF_4^- ion. At low bias voltages, SF_6 etching is therefore roughly isotropic. As bias voltage is increased, physical etching increases and the etch becomes more directional.

For etching features with straight sidewalls, SF_6 can be mixed with other species, such as fluorocarbons. These fluorocarbons cause a fluoropolymer to form on the sample surface. With careful tuning, a net fluorocarbon buildup can form on the sidewalls, while the addition of physical etching at the base of the feature is sufficient to allow etching to occur. This strategy works well for shallow etches, but is inefficient because the etching and deposition processes interact, wasting a lot of reactants.

In Deep Reactive Ion Etching (DRIE) [84] this inefficiency is eliminated by temporally separating the etching and deposition steps. First, a fluoropolymer layer is deposited on the sample. Then some etching is performed. During the etch, the directionality of the etch implies that the polymer at the base of the etched feature is removed while the sidewalls are still coated. There is then a window of time during which the sample material at the base of the feature is etched, while the sidewalls are protected by fluoropolymer. Before that protective layer depletes, the etch is stopped, and the process repeats with a new fluoropolymer deposition stage. This process can etch silicon at rates of a few tens of micrometers per minute, and is highly selective to photoresist or silicon oxide masks. Trenches with aspect ratios as high as 25:1 can be created.

3.1.2 Basic Emitter Process

To make emitters, we have been using a combination of DRIE and pure SF_6 etching, as pioneered by Velásquez-García [37]. In that work, silicon needles were created by placing a star shaped layer of silicon oxide on a silicon surface, as shown if Figure 3-1. An SF_6 etch was used to partly undercut the oxide mask. Then a DRIE etch created a tall slender needle structure. The silicon under the oxide mask overhang was protected during the DRIE etch, so the tall sidewalls were further from the axis of the needle than the undercut region below the oxide mask. Finally, another SF_6 etch completely undercut the oxide mask until it fell off. Simultaneously, the final SF_6 etch etched the tall sidewalls inward, making the needle slenderer.

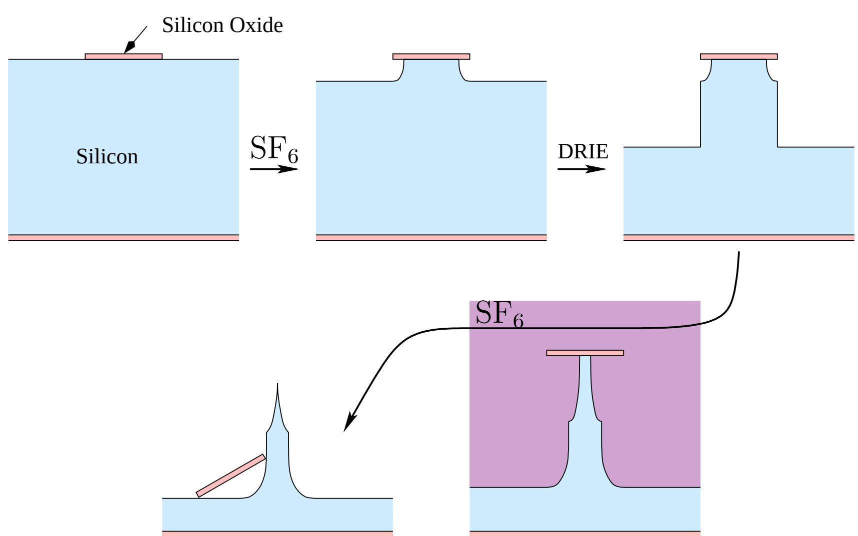

Figure 3-1: Creating the emitters with an alternation of SF_6 and DRIE etches

3.1.3 Emitter Geometry Model

Velásquez-García used simple geometrical shapes for the emitter etch mask, and a simple SF_6-DRIE-SF_6 etching sequence. A lot of variation is possible on this basic theme. What is the best mask shape to use? What is the best timing of the SF_6 and DRIE etches? More complex mask shapes are also possible, as are more complex etch sequences. What kind of improvements could be made to the emitters by introducing this complexity? To answer these question, we have developed a simple model that allows the geometry of an emitter to be deduced from the shape of the oxide mask and the etch sequence to be performed. Complete understanding of the DRIE and SF_6 etching process is a complex problem of plasma physics, and we have chosen a highly simplified model to work with. The simplicity of the model makes it useful to hone the intuition of the emitter designer, and fast to implement. Our hope is that it will be sufficiently accurate to help improve emitter geometry at the design stage.

In this model, SF_6 etching is modeled as an etch with a rate that is a function only of the normal of the surface being etched. If this function is uniform, we get an isotropic etch. However, we find better agreement with experiment if vertical features are etched faster than horizontal ones. A slight anisotropy related to the silicon crystal planes can further improve the model accuracy. For DRIE etching, the etch proceeds straight down at a uniform rate, except in areas directly below an oxide mask, which are not etched at all.

In the numerical implementation of the model, a partially etched wafer is represented by two 2D functions that vary along the surface of the wafer. The first function $M(x, y)$ defines the areas of the wafer that are protected by an oxide mask. It only takes on the values 1 or 0 indicating the presence or absence of oxide, respectively. The second function defines the shape of the etched silicon surface. Assuming that the shape of the silicon can be defined by how much silicon has been removed at a given (x, y) location along the wafer (i.e., there is no undercutting of the silicon), the shape can simply be represented by $z(x, y)$, which indicates how far down the silicon has been etched at (x, y).

After a DRIE operation of depth d, the function z becomes

$$z'(x, y) = z(x, y) + d \cdot M(x, y). \tag{3.1}$$

Stated in plain simpler terms, points below the oxide mask have remained unchanged, while points outside the mask have been etched down a distance d.

The SF_6 operation is more complicated. It is characterized by a concave shape function $s(\tilde{x}, \tilde{y})$ which can take on real values or $-\infty$. For normalization purposes, we take $s(0, 0) = 1$. Generally, the shape function is smooth wherever it is finite. After an SF_6 operation, the function z becomes

$$z'(x, y) = \max_{\substack{\tilde{x}, \tilde{y} \\ z(x+d\cdot\tilde{x}, y+d\cdot\tilde{y})>0}} \left(z(x, y),\ z(x + d \cdot \tilde{x}, y + d \cdot \tilde{y}) + d \cdot s(\tilde{x}, \tilde{y}) \right). \tag{3.2}$$

The maximization considers all the positions where the silicon is exposed (i.e., where it is not still directly in contact with the oxide mask) to see how deep the etch has progressed at a given point. In effect, the shape function defines the shape that an oxide mask with a point hole in it creates during SF_6 etching, once the silicon directly below the hole has receded by a unit distance. The model of the SF_6 etch simply states that during an SF_6 etch, etching proceeds from the exposed silicon surface in all directions, with a speed that is related to the tilt of the sidewall by the shape function s.

Typically, we find that s defines a roughly ellipsoidal shape, with a z-axis of length 1 and x and y axes of length around 0.5. Moreover, a small dependence in crystal plane orientation in the SF_6 etching seems to cause a slight bulge in the x and y axes compared with axes rotated by 45° in the xy-plane.

In our design efforts, we have taken

$$s(x, y) = \sqrt{1 - 4x^2 - 4y^2 - \alpha x^2 y^2}. \tag{3.3}$$

When the square root is not real, we set $s(x, y) = -\infty$. For $\alpha = 0$, this results in an ellipsoid with a z-axis of length 1 and x and y axes of length around 0.5. For non-zero α, there is a slight bulge in the x and y axes compared with the $x + y$ and $x - y$ axes.

In our simulations, we have selected the value of α which gives a 90% bulge, i.e.,

$$\alpha = \frac{64 \cdot (1 - 0.9^2)}{0.9^4} \approx 18.5$$

.

Near the end of emitter formation, the oxide mask falls off, and should no longer be considered when modeling DRIE etches. SF_6 etches are unaffected by the mask falling off, because the converging etching fronts that meet to produce the tip cause the tip height to reduce faster than downward etching at the tip does. Thus, the etching at the tip is unaffected by the loss of the mask. This fact will be revisited in Section 3.1.6, and illustrated in Figure 3-8(a).

With these models of the etch processes, it is easy to simulate an arbitrary sequence of DRIE and SF_6 etches. The result will be an alternation of vertical sidewalls and inclined surfaces. The vertical sidewalls are generated during the DRIE steps, and etched inward by subsequent SF_6 etches. Each SF_6 etch generates an inclined surface. The surface for a given SF_6 etch is generated by moving a scaled shape function along the contour of the mask, at the height of the emitter base at the beginning of the etch. The shape function is scaled proportionally to the amount of SF_6 etching that occurs after the current SF_6 step (inclusive). In some cases, one of the sidewalls can disappear when the adjacent inclined surfaces meet. With overetching, whole portions of the emitter can also disappear. Figure 3-2 shows this decomposition of the emitter geometry in a simple case.

3.1.4 Ridge Lines

When observing SEMs of emitters fabricated from angular masks with only SF_6 etching (see Figure 3-9(b)), one is struck by the existence of ridge lines in an otherwise smooth surface. As we shall see, these discontinuities in the surface normal are predicted by our model, and play an important part in producing high quality tips.

According to the model, these edges occur when control of the depth jumps discontinuously from one point to another on the edge of the mask (i.e., where etching

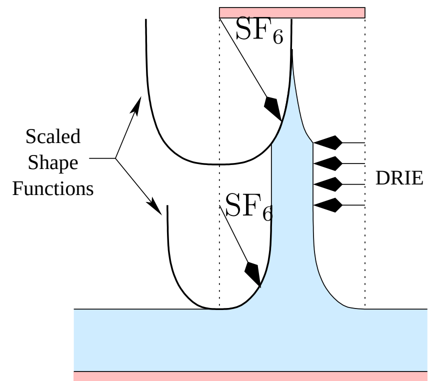

Figure 3-2: The emitter geometry is engendered by one scaled shape function per SF_6 etch, and one straight sidewall per DRIE etch

fronts meet). Indeed, as illustrated in Figure 3-3, the depth of most points below the mask is set by a single point on the edge of the mask. In the figure, each colored region corresponds to points whose depth is set by a given concave stretch of the mask boundary (a stretch in which the border only turns away from the inside of the mask). In each region, the maximum in Equation (3.2) is reached at exactly one point on the corresponding concave stretch of the boundary. When two regions meet the maximum is reached in more than one way. In this case, moving slightly away from the boundary causes different mask-edges to dominate. The regions dominated by two distant mask edges will not agree on the surface normal, and so a normal discontinuity is present where the two regions intersect. Likewise, when three regions intersect, a tip is formed.

If we assume that $s(x, y)$ has symmetry of revolution, then there is an interesting connection between ridge lines and the notion of Medial Axis Transform (MAT) in computational geometry [85]. The MAT of a region is the set of points for which there is more than one closest point on the boundary of the region. For a polygon, the MAT is made up of portions of angular bisectors of edges, perpendicular bisectors of vertices, and parabolas that have one edge as a directrix and one vertex as a focus. The MAT can easily sketched by hand, and is an excellent aid in designing mask shapes.

Another notion of computational geometry, the distance transform, is also relevant

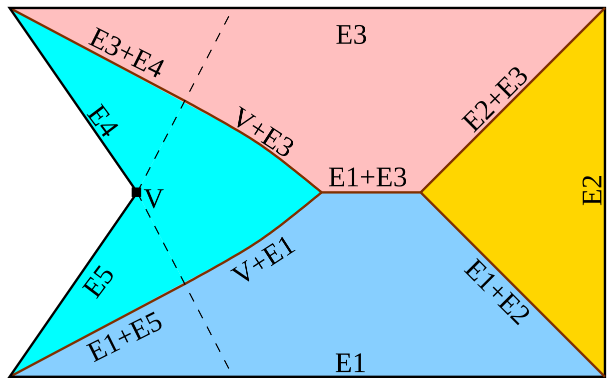

Figure 3-3: A hypothetical mask shape, the regions determined by concave stretches of boundary, and the ridge lines that will result. (The dotted lines delimit the area which is controlled by the vertex V, the boundaries of which are parabolic)

when $s(x, y)$ has symmetry of revolution. The distance transform of a mask shape is a function $\mathrm{DT}(x, y)$ that maps each point in the plane to the distance from that point to the closest point outside the mask. This function is zero in areas not covered by the mask, and steadily increases inside the mask [86]. Equation (3.2) is closely related to the definition of the distance transform, so that the final emitter shape can be expressed by

$$z(x, y) = f\left(\mathrm{DT}(x, y)\right),$$

where f is a function that depends on the sequence of etches that was performed. The ridges we have been describing arise because $\mathrm{DT}(x, y)$ is not smooth on the MAT.

So far we have been considering emitters formed from a single SF_6 etch. Another type of ridge appears when DRIE etching participates in the emitter formation, because the DRIE steps introduce discontinuities in f. These discontinuities lead to discontinuities in z(x,y), with an edge at the top of the corresponding vertical DRIE sidewalls.

Understanding where the ridges of an emitter are going to be located is very helpful in understanding what the shape of the emitter is going to be like, and, as we shall see, for producing high quality emitter tips. Connecting the emitter shape with existing notions in computational geometry helps the designer get an intuitive feel for the relation between the mask shape and the geometry of the final emitter.

This relationship doesn't hold when α is non-zero, but the departure from symmetry of revolution is small enough that it does not have a significant impact on the final emitter shape.

3.1.5 Key Emitter Properties and Uniformity

The model we have just presented allows the detailed geometry of an emitter to be estimated. However, early in the design process, it is often more useful to have a reduced-dimension analytical model to compute basic emitter parameters. In this section, we show how to estimate emitter tip half-angle α_t, tip height h_t and tip depth d_t from the mask radius R_m mask, the DRIE depth h_{DRIE}, and the amount of undercutting R_u.

Figure 3-4 shows what these different parameters are. The tip half-angle α_t is the angle between the vertical direction and the steepest surface arriving at the tip, and is an important parameter for field enhancement. The tip height h_t is the height of the tip above the lowest etched level, which is also important for field enhancement. The tip depth d_t is the distance between the tip and an unetched silicon surface, which is important because we will be using the unetched silicon surface as a reference height is Chapter 5. The mask radius R_m is the horizontal distance from the tip to the edge of the mask, which, for a mask yielding a single tip, is the radius of the largest inscribed circle for the mask. The DRIE depth h_{DRIE} is the total vertical distance etched by DRIE. The amount of undercutting R_u is the distance etched inward during the different SF_6 steps.

The formulas given here assume that the shape function is an ellipsoid with a vertical dimension f times greater than the horizontal dimensions. They also assume that a tip is actually formed ($R_u > R_m$). Moreover, they assume that etching begins with an SF_6 etch, and that it is this etch which engenders the tip (See Figure 3-2), and that no DRIE etching takes place after the mask has fallen off (in Section 3.1.3 we saw that the mask falling off affects DRIE etches but not SF_6 etches). The last assumption is relevant because, as long as the goal is to minimize h_t, any design that does not meet these assumptions could be improved by removing the first etching

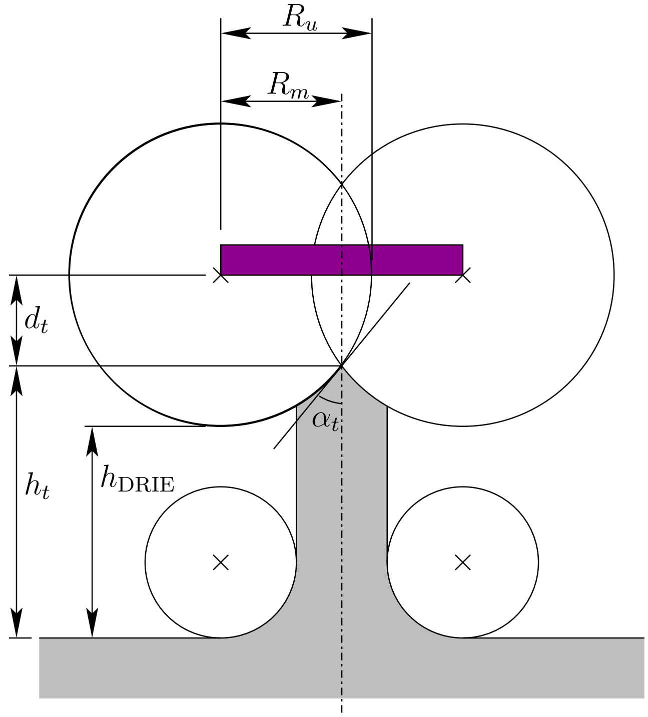

Figure 3-4: Key emitter properties. The loss of the mask does not affect SF_6 etching, as discussed in Section 3.1.3, so it is appropriate act as if the mask were still attached in this figure

step.[1]

The expressions for the different parameters follow directly from simple geometry in the case shown in Figure 3-4, where $f = 1$. To generalize to other values of f, we note that they can be reduced to the $f = 1$ case by scaling by a factor f in the vertical direction. This scaling operation leaves R_m and R_u unchanged, and divides r_t, h_t and h by a factor f. The tip angle is transformed into α_t' such that $\tan(\alpha_t') = f\tan(\alpha_t)$. The resulting relations are:

$$d_t = f\sqrt{R_u^2 - R_m^2} \tag{3.4}$$

$$h_t = h_{\mathrm{DRIE}} + f(R_u - \sqrt{R_u^2 - R_m^2}) \tag{3.5}$$

$$\tan(\alpha_t) = \frac{1}{f}\sqrt{\left(\frac{R_u}{R_m}\right)^2 - 1}. \tag{3.6}$$

[1]Despite this fact, the etch sequence used for the final emitters does not satisfy the second condition. Because we tried to reduce sensitivity to process variation at the cost of a slightly increased h_t.

These equations can be rearranged to give the undercut that is needed to achieve a particular tip half-angle, as well as the depth and the tip-to-total height ratio for that half angle:

$$
\begin{aligned}
d_t &= f^2 R_m \tan(\alpha_t) && (3.7)\\
R_u &= R_m \sqrt{f^2 \tan^2(\alpha_t) + 1} && (3.8)\\
\frac{h_t}{d_t + h_t} &= 1 - \frac{1}{\sqrt{1 + \frac{1}{f^2 \tan^2(\alpha_t)}}} \quad \text{when } h_{\mathrm{DRIE}} = 0. && (3.9)
\end{aligned}
$$

In practice, h_{DRIE} is uniform to within 10%, but R_u and f have been observed to vary by over 50%, mainly due to systematic radial etch non-uniformity. How do we avoid being too sensitive to these parameters? All three properties we have computed have infinite derivatives as R_u approaches R_m from above. Thus, to reduce sensitivity to process variation, it is preferable to overetch tips to some extent. (In any case, when R_u is just above R_m, the etch model is not very accurate, as the tip takes on the hourglass shape shown in Figure 3-5.[2]) In designing emitters, this sensitivity needs to be traded off with the fact that as R_u increases, the tip half-angle and depth d_t increase, while the height h_t decreases.

If tall emitters with good height uniformity are desired, the best option is to use a small mask. This way, the height of the emitters will mainly be set by h_{DRIE}, as the value of R_u can be kept small. The risk with thin tall emitters is that they may be undercut if there is any negative taper. Figure 3-6 shows a picture of a serrated blade emitter we made. Its Gothic arch appearance is caused by slight negative taper during the DRIE etch. The base of the blade was completely etched through, except at the thicker parts surrounding the tips. Other devices that were etched simultaneously were completely etched through and destroyed. To get around this problem, we lengthened the first SF_6 etch and decreased the height of the blade a bit. Subsequent serrated blade emitters no longer had this Gothic appearance. Another way to fix this problem would be to use a different DRIE recipe. The recipe we used

[2]This hourglass shape could be caused by charging of the oxide mask deflecting ions downward or by the limited supply of grazing incidence ions which can reach areas just below the mask.

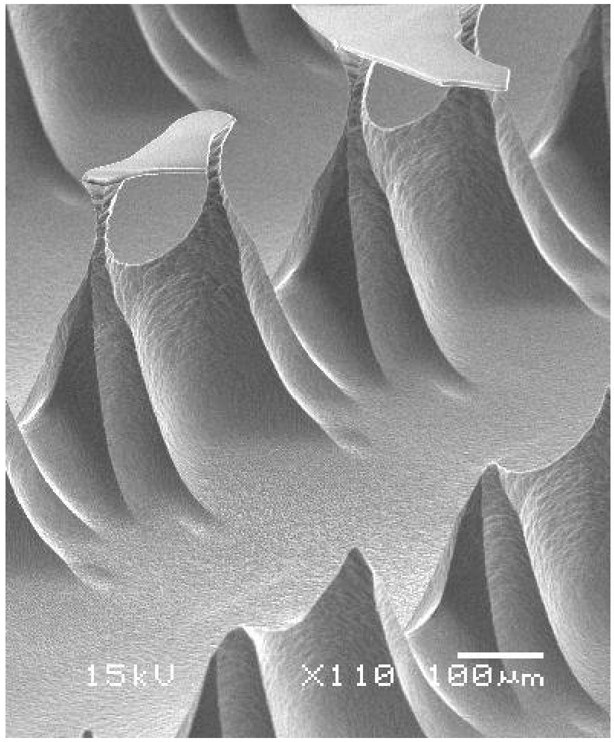

Figure 3-5: In reality, etching progresses more slowly just below the mask, leading to an hourglass shape

(JBETCH) is optimized for narrow trenches and has negative taper for wide features; recipes that do not have this drawback exist, but were not tried [55].

3.1.6 Making Tips Sharp

The model predicts that emitters will have infinitely sharp tips, which is unphysical. Figure 3-7 shows some SEMs of a number of tips. Some of the tips look very nice (Figures 3-7(a) and 3-7(b)), with radii of curvature of less than 1 μm. But there are many imperfections:

- Figure 3-7(c) shows a tip that looks like a blade. This blade apparently formed because of a slight etch non-uniformity. Normally, four ridges should meet to form the tip on this emitter. However, because of etch non-uniformity, the ridges were slightly shifted so that they met two by two, with a blade forming between the meeting point. An non-uniformity of only a few micrometers

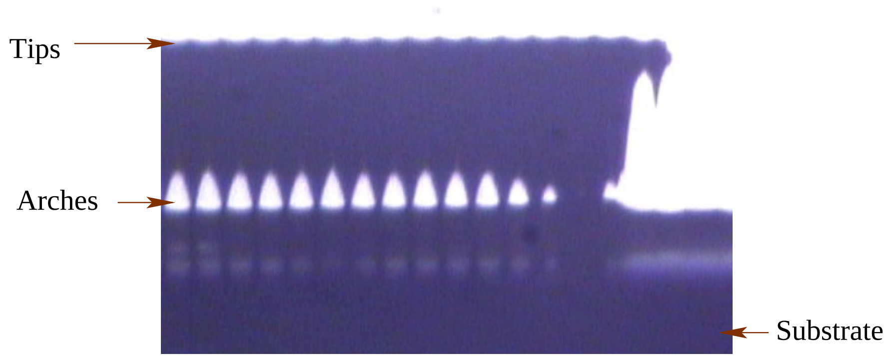

Figure 3-6: A serrated blade emitter for which negative DRIE taper etched through the base of the blade, except the slightly thicker portions near the tips

in the photolithography or etching (R_m=200 μm) is sufficient to cause this effect, which could have dramatic consequences on the emission properties of the emitter. Instead of having one Taylor cone firing upwards from this tip, we can expect two Taylor cones firing at an angle to form. Fortunately, avoiding this problem should be simple. Instead of the tip being formed at the meeting point of four ridges, we should design it to be at the meeting point of only three ridges. This way, if one of the ridges gets shifted because etching is progressing faster on one side than on the other, the position of the tip is moved slightly, but the tip does not become a blade. Moreover, the three ridges should ideally meet at 120° from each other so that the tip is symmetric. Otherwise the tip will be angled in the direction of the ridge deficit instead of pointing straight up.

- Figures 3-7(d), 3-7(e) and 3-7(f) show that in many cases the tip of the emitter isn't a clean point, but rather looks broken. We think this appearance may be due to the way in which the oxide mask falls off the emitter once it is completed. Up to now, we have considered that the oxide mask falls off when the emitter gets sharpened to a tip by the converging etching fronts. In practice, it is possible that the oxide mask causes the tip to break slightly before this time, resulting in a tip with a fractured appearance. This break could take place, for example, when the dies are removed from the etcher to examine the progress of

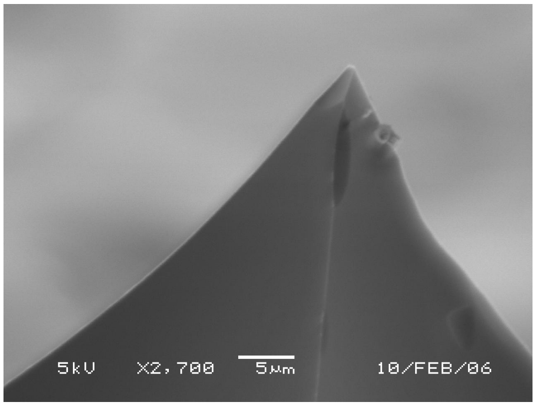

(a) A nice looking mountain tip

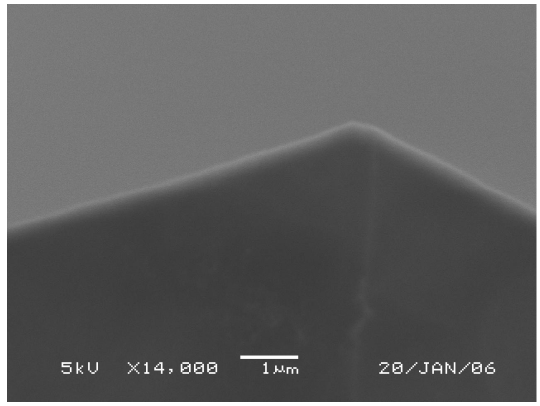

(b) A serrated blade tip seen from the side of the blade

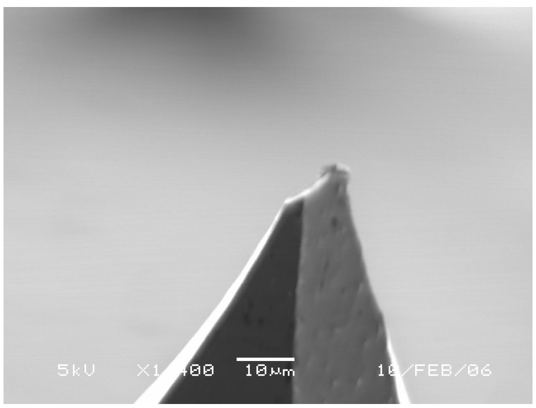

(c) A blade shaped mountain tip

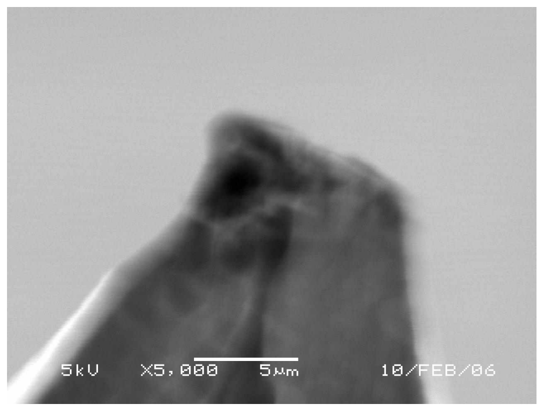

(d) A pencil tip

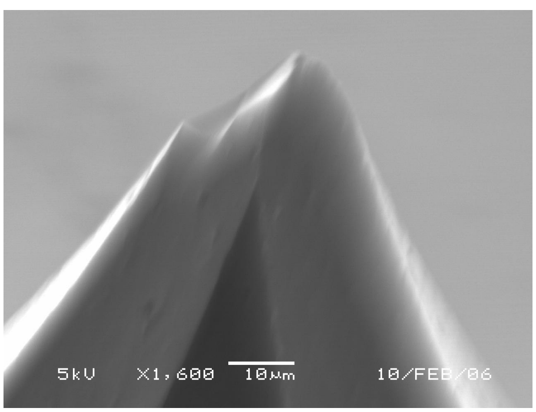

(e) A broken mountain tip

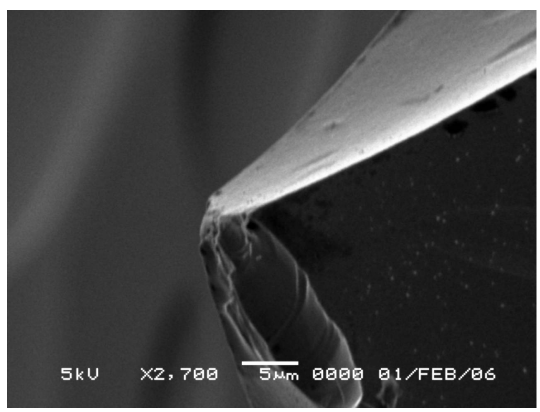

(f) A nearly repaired mountain tip

Figure 3-7: SEMs of emitter tips

the etching. Since the fractured tip is likely to have bad electrospray emission characteristics, we need to find ways to get nice sharp tips. Here, as with process uniformity, overetching may provide some relief. Indeed, by overetching, the sides of the emitter will continue to be etched, and this etching will erode away the fractured part of the tip until it disappears (Figure 3-8). At this point, only a clean tip should remain. Another possibility might be to remove the oxide mask using hydrofluoric acid just before it breaks off, and etching a bit more. Hopefully, removing the mask will make the tip less susceptible to air currents, and will prevent it from breaking before it comes to a sharp point. This method has the added advantage that it does not let oxide masks detach, littering the wafer and the plasma etcher, and causing undesired mesas to form wherever the masks happen to land.

3.1.7 Model Evaluation

We now turn to a qualitative evaluation of the model. To evaluate the model, we compare its predictions with actual emitters that were fabricated.

Figure 3-9 shows various emitter geometries we have made. These emitters are all 300 to 400 μm tall. The pencil emitter [37] requires 50 μm of underetching before its top detaches, as well as 300 μm of DRIE etching. The mountain emitter uses the same mask as the pencil emitter, but scaled up four times. It needs 200 μm of underetching, and no DRIE etching at all. The masks for the serrated blade emitters are rectangles with periodic triangular outcroppings. These triangular outcroppings define little tips along the blade that Taylor cones can anchor to for stable operation. These tips are spaced by about 100 μm. The top blade needs 100 μm of underetching, while the bottom one needs 50 μm of underetching. The model correctly predicts the position of the tips relative to the end of the blades. The model pictures show the tips rising above the rest of the blade more than in reality. We think this effect is due to the discretization in our implementation of the model, not the model itself.

Figure 3-10(a) was used to calibrate the amount of anisotropy in the model. This

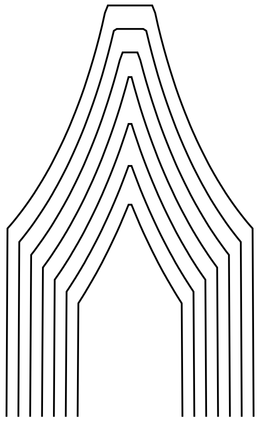

(a) Simulated

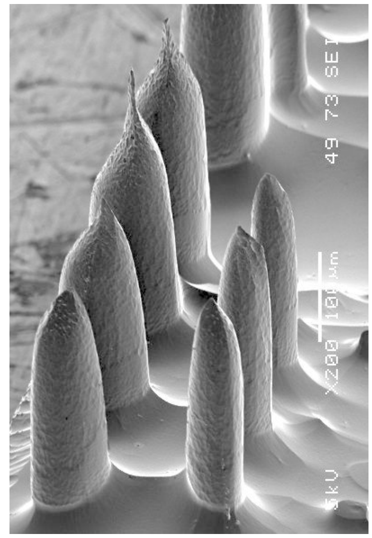

(b) Pillars with progressively increasing mask diameters

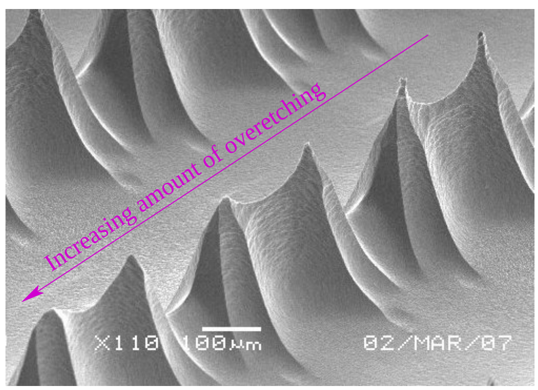

(c) Tips with progressively more overetching

Figure 3-8: Profiles of an emitter after the oxide mask has broken off, compared with reality. Overetching a broken finger will regenerate the tip and reduce sensitivity to etch duration, but will increase the half-angle of the tip and reduce the emitter base diameter

figure shows a top view of a mountain created by a star shaped mask as shown in Figure 3-9. The eight points of the star produce visible ridges going up the sides of the mountain. The edge anisotropy causes the ridges to deviate slightly so that the ridges meet two by two before reaching the tip of the mountain. Each pair of meeting ridges merge into a single ridge that continues straight up the to the tip of the mountain. Figure 3-10(b) shows the same behavior when anisotropy is present in the etch model. Unfortunately, the amount of anisotropy seems not to be very consistent for different

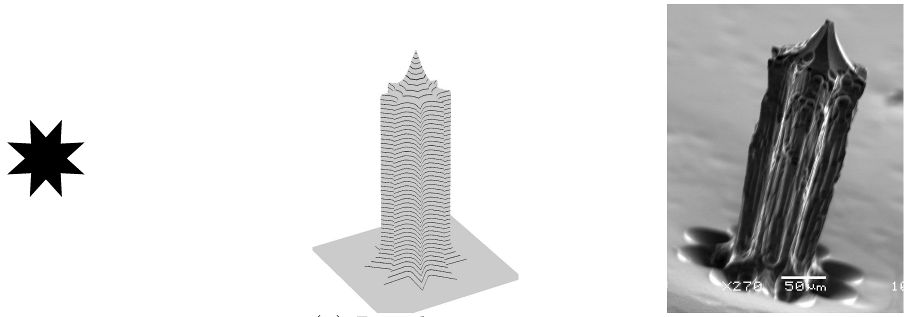

(a) Pencil emitter

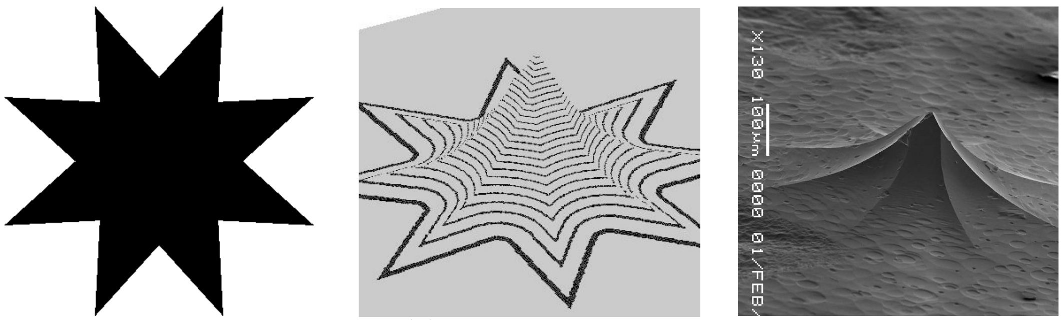

(b) Mountain emitter

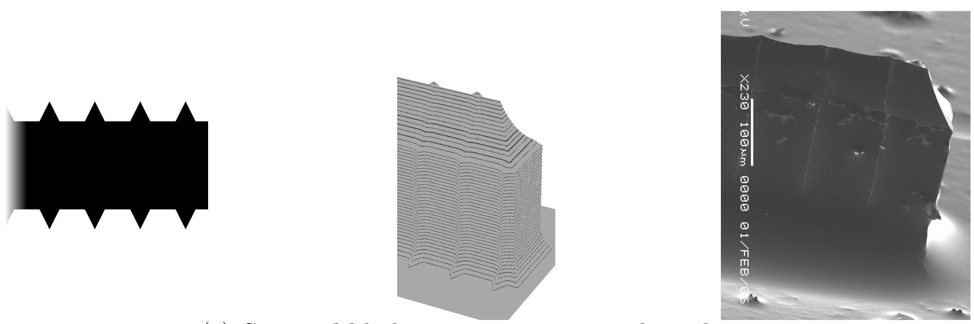

(c) Serrated blade emitter, 50 μm underetching

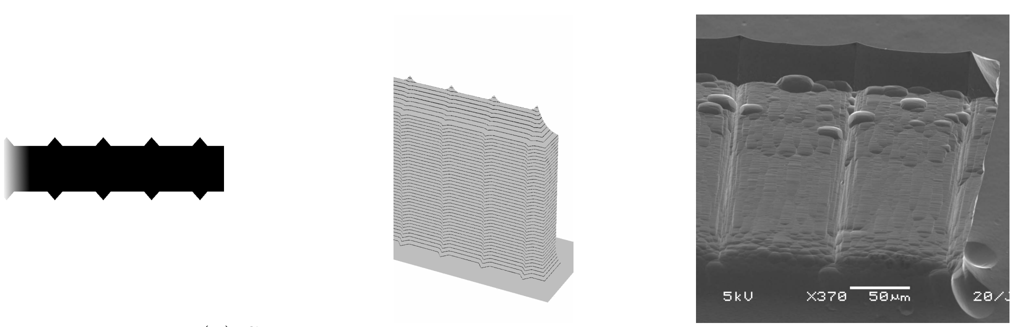

(d) Serrated blade emitter, 100 μm underetching

Figure 3-9: Comparison between modeled and observed structures for a pencil emitter, a mountain emitter, and two kinds of serrated blade emitters

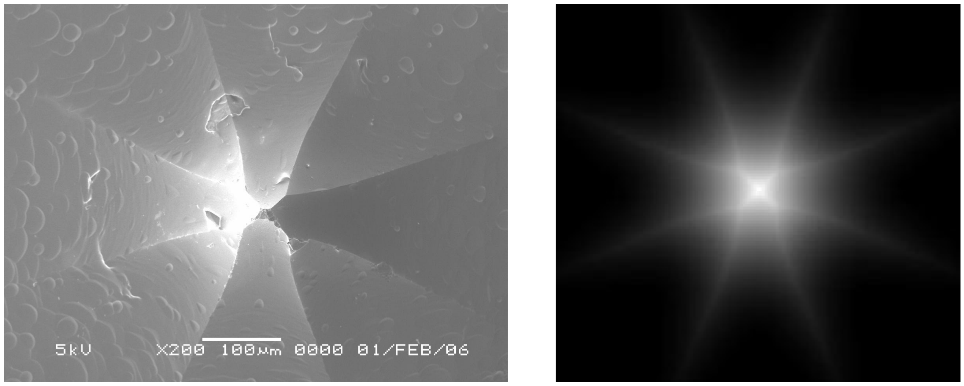

(a) SEM (b) Model

Figure 3-10: Top view of a mountain emitter showing etching anisotropy

emitter dies. Other mountains we have observed have less anisotropy than the one presented here. This inconsistency may be due to problems we had mounting the dies to a handler wafer [38] when shaping the emitters, the temperature of the die during processing could well have contributed to the variation in anisotropy from die to die.

3.2 Fancy Emitters

Our first generations of emitters were designed without the benefit of the modeling we have described above. We were able to make blades with regularly spaced tips, but these tips are only slightly higher than the rest of the blade, which limits the field enhancement they afford. Given the model, a number of new strategies come to mind for making better emitters.

First, it is now reasonable to perform more alternations between SF_6 and DRIE etching, as we will have an idea of what to expect from the extra alternations. With the SF_6-DRIE-SF_6 etching we have done so far, the shape of the emitter tip is very constrained. The top portion of the emitter has a profile which flares out below the tip. Then there is a sharp transition to the straight sidewall, and finally a curved transition to the horizontal plane of the substrate. By adding extra alternations, we can ease the transition from the top portion to the strait sidewalls, making a sharper tip with more field enhancement. We have limited ourselves to two DRIE steps interlaced with three SF_6 etches, but in the limit of a large number of transitions, it would be possible to produce near-conical emitters, or emitters with more complex profiles. For complex profiles, the main limit is the amount of non-uniformity in the etching.

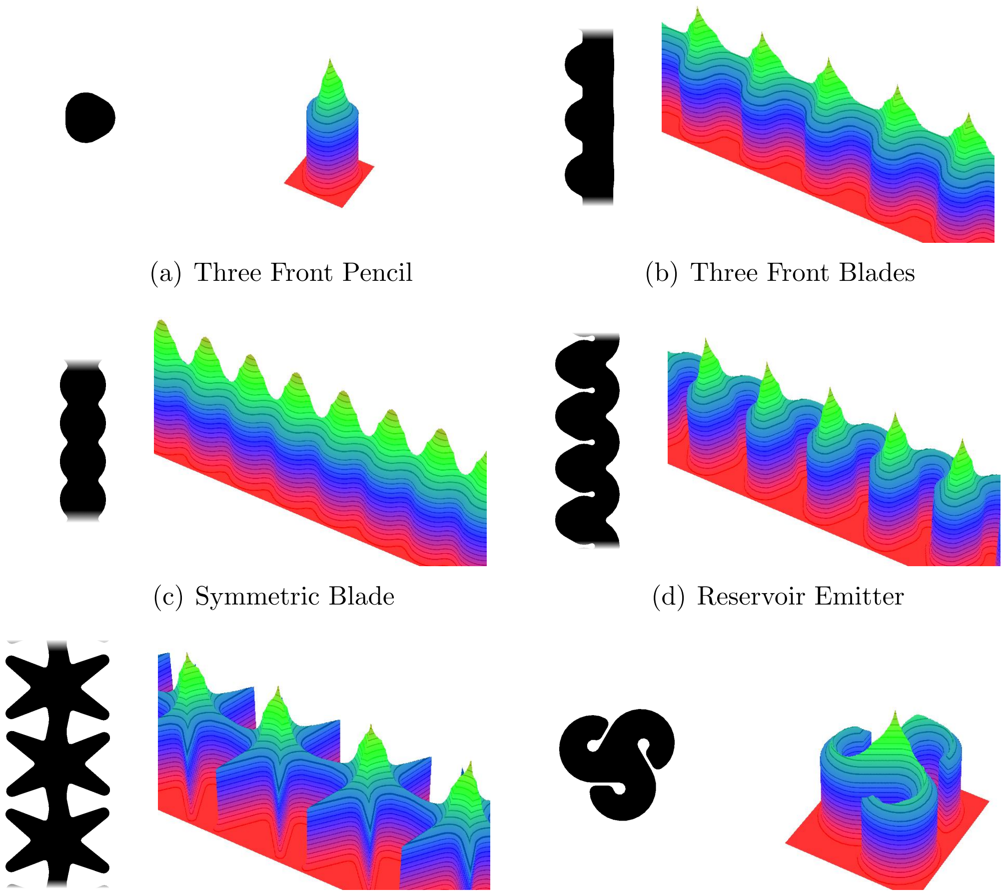

(a) Three Front Pencil
(b) Three Front Blades
(c) Symmetric Blade
(d) Reservoir Emitter
(e) Star Blade
(f) Swirl Reservoir Emitter

Figure 3-11: Improved emitter designs using multiple DRIE/SF_6 alternations

Another enhancement we can consider is the use of nested masks (i.e., multiple masks that are present simultaneously, but which can be stripped one at a time). With nested masks, it would be possible to use a large photoresist mask to define the large-scale shape of the emitter, with a large base to resist the negative taper of the DRIE etch. Then a smaller oxide mask, present only in the center of the emitter could be used to define the tips. The small size of the inner mask would reduce the effect of SF_6-etch non-uniformity on the overall geometry of the emitter.

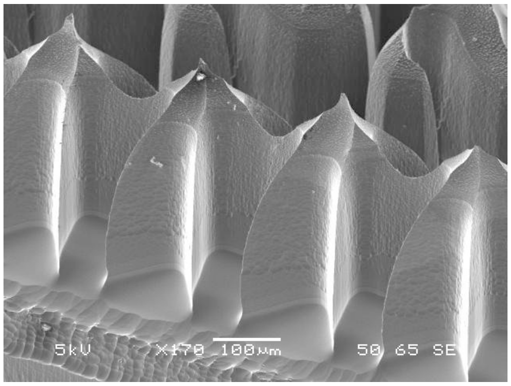

(a) Star Blade

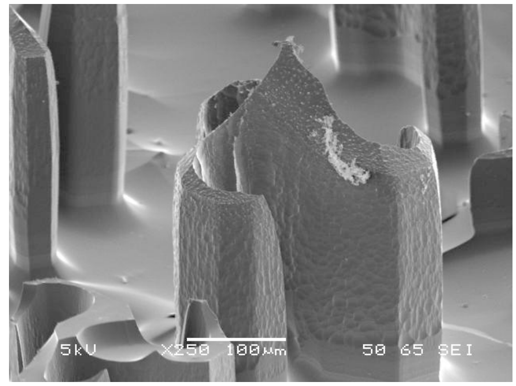

(b) Swirl Emitter

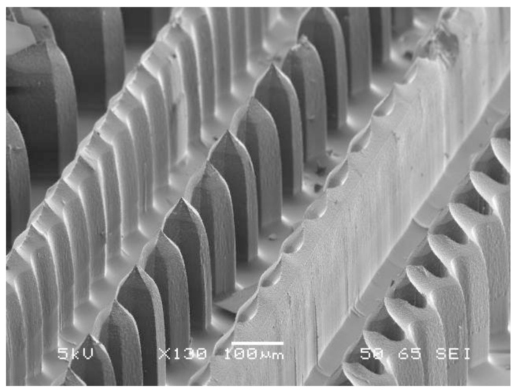

(c) Other new blades, a bit overetched

Figure 3-12: Scanning electron micrographs of the new emitter geometries

Figure 3-11 shows a number of emitter designs we have created using the model. These designs are all intended to be etched with the same three-SF_6, two-DRIE etch sequence. We begin by a 20 μm undercut SF_6 etch, a 90 μm down DRIE etch and a 40 μm undercut SF_6 etch. These etches create a sharp tip with minimum flaring. Then a 190 μm DRIE etch defines the tall sidewalls, and a 20 μm undercut in SF_6 etch finishes sharpening the tip. We designed the geometries to tolerate 20% variation in the SF_6 etching across the wafer. All but the *Symmetric Blade* emitters use the three converging fronts method to ensure nice sharp tips. The three *Reservoir* emitters are designed so that the liquid meniscus will rise up the channels in the emitter, possibly reducing the flow impedance to the emission site. The fancier geometries are also intended to stress-test the etch model.

These new geometries were fabricated as part of the hybrid assembly scheme (see

Section 4.5.2). During the fabrication, we found that the non-uniformity of the SF_6 etching was actually near 50%, possibly because of micro-loading effects. Thus, many emitters were over or underetched. Figure 3-12 show some promising results, and in some cases striking similarities with the predictions. Unfortunately, these emitters were never fired, because we were unable to reproduce a surface treatment that can coat vertical sidewalls.

3.3 Wicking Surface Treatment

3.3.1 Overview of Fabrication Methods

In Section 2.2.3, we discussed the surface treatment that is needed to allow a film of liquid to exist stably on the surface of an emitter. We now look at the fabrication techniques to make this surface treatment. Because the surface to be treated is non-planar, top-down methods, in which the detailed shape of the roughness is directly dictated by the designer, are not applicable. Instead, we must resort to bottom-up methods, in which the texture of the surface arises naturally from an appropriate physical process. A number of different physical processes could be used to make the desired texture. However, there is considerable process development work in going from the original idea of using a particular process, and actually getting a consistent treatment at the desired scale from that process.

Plasma Etching: Under certain conditions, plasma etching can lead to the formation of black patches on a silicon surfaces, called black silicon. The black color of the patches is caused by high aspect ratio surface roughness. Light arriving on this surface is reflected around multiple times by the surface roughness and is greatly attenuated before it makes its way out, leading to the dark appearance. According to [87], black silicon can occur when an etching recipe is tuned to create nearly vertical sidewalls. In such conditions, any micromasking (localized masking due to the presence of a particle, or a slightly thicker oxide layer) will create a pillar on the surface being etched. If there is a high density of these pillars, the etched surface appears to be covered in vertical blades of grass. Though they do not discuss what causes the micromasking, they do give a simple method to find conditions in which black silicon should form, which we were able to successfully apply. As detailed in Section 3.3.2, plasma etching will be our method of choice for the emitter surface treatment.

Porous Silicon: There exists a range of HF-containing environments in which silicon, when anodized, forms a porous structure. Depending on the silicon doping, the solution that is used, the lighting conditions, and the crystallographic orientation, the scale and structure of the pores can vary wildly. In some cases [88], these pores form interconnected networks that could be used as our surface treatment. We ran a few experiments with mixtures of HF, Nitric Acid and water, for which electrodeless anodization can take place in the stain-etch method [89–91]. In these experiments, we were able to make black-looking silicon surfaces. However, the samples that were observed in an SEM showed no structure above the 100 nm scale. Moreover, the appearance of the blackened area depended strongly on how fast the samples were removed from the etching solution, suggesting that good uniformity and repeatability would be difficult to obtain.

Sol-Gel Methods: In sol-gel methods, a colloidal suspension of particles solidifies into a solid gel, which can have a high degree of porosity [92]. Sol-gel methods can be amenable to dip-coating methods and may be useful to form a porous layer on a non-planar silicon surface. There is concern, however, that the emitter tip radius may be increased as a result of such a coating.

Electrodeposition: In electrodeposition, material in solution is deposited on an electrode used to flow current through the solution. Electrodeposited films often have a high degree of porosity [93, 94]. In some cases, the deposit takes on highly porous dendritic structures. In others, columnar grain structures are formed, which could, by selective grain boundary etching, be converted into arrays of pillars. Much research has gone into finding and documenting methods of avoiding the formation of porosity so that robust non-porous coatings can be deposited. Unfortunately, the literature is not as helpful when the goal is to design a porous layer with specific scale of porosity. Moreover, electrochemistry is sensitive to the shape of the object undergoing deposition, so the emitter tips are likely to be coated differently from the rest of the sample. An platinum treatment that appears to have such a caveat can be found in [95]. Electrodeless plating may avoid this problem.

Sputter deposition: In sputter deposition, material is sputtered from a target under bombardment by a plasma onto the sample to be coated. There is a lot of similarity between the material structures created by electrodeposition and sputtering, so sputtering is also a promising candidate to make porous surface layers [94, 96].

3.3.2 Plasma Treatments for Black Silicon Formation

The first use of plasma processing to create a rough surface for transport of liquid in an externally wetted electrospray array is by Velásquez-García [37, 58]. This work uses a Cl_2-He plasma in a LAM Research Corp. 490-B etcher to create the surface treatment. This treatment is reported to be conformal, and, using optical microscopy, liquid placed on the surface can been seen to spread along the surface and climb vertical emitter sidewalls.

Later, Garza set out to reproduce this Cl_2-He treatment in a Plasmaquest 145 RIE/PECVD system [69, 70]. The change in etcher was motivated by Microsystems Technology Laboratories (MTL) contamination rules, because it was believed that the wafers would need to undergo processing that would not allow the LAM 490-B to be used for the treatment. (This assumption turned out not to be true.) This study eventually led to a surface treatment that allowed emitters to be fired. However, the emitter operation was not very satisfactory. The surface treatment was also very fragile. For example, liquid would not wet samples that had been cleaned in acetone[3].

[3] Perhaps the cleaning in acetone slightly increased the microscopic contact angle of the surface

There was good evidence that liquid would spread effectively on a flat treated surface, but no proof that the same was true of the inclined surface of an emitter. Finally, SEM images of a treated surface reveal roughness on the 100 nm scale, which is much smaller than the target set in Section 2.2.3.

Since it was not clear whether Garza's treatment would provide adequate liquid transport, since in the end there was nothing preventing the use of the LAM 490-B, and since the LAM 490-B is much more convenient to use than the Plasmaquest, we decided to use the LAM 490-B for the present work. Unfortunately, the shower-head of the LAM 490-B had been changed since the work by Velásquez-García. The recipe that had previously been used yielded a surface that was black in areas, but that did not appear to transport liquid effectively. Since Velásquez-García's treatment had been found after a few exploratory attempts, we decided to try adjusting the parameters to find a treatment that would work.

3.3.3 Evaluation of a Treatment using Water

Up to now, evaluation of a surface treatment had always been carried out outside the cleanroom by applying drops of the target propellant EMI-BF_4 to the surface, and observing the spreading speed of the halo that forms around them. Given the discussion in Section 2.2.3, we expect that whether or not a liquid can wet the surface depends on its microscopic contact angle, but that there will be a fairly sharp cutoff angle, comparable with the inclination of the roughness. For a surface that can be impregnated, once the impregnation has taken place, the surface permeability should be independent of the liquid that is used, as long as there is sufficient liquid on the surface to fill the surface roughness.

Therefore, for this work, we used water to test the effectiveness of surface treatments. Water is a good liquid to use, as its high surface tension implies that if water is able to infiltrate a surface, most liquids will be able to. Moreover, this test could be carried out in the cleanroom as soon as the wafer had been processed. A drip of water

so that the angle α of the roughness defined in Section 2.2.3 was no longer sufficient to make the surface perfectly wetting.

would be placed on the treated surface using tweezers. If a large contact angle was found, the surface was clearly unsuitable. In other cases, the drop would spread out until it was nearly flat, and a small halo would extend from the drop, characteristic of liquid spreading in the porous surface layer. This halo would stabilize at a fixed diameter once a balance was reached between evaporation and liquid being drawn from the drop into the porous surface. The larger the halo, the lower the surface permeability of the treatment.

In addition to being easier to carry out, this method of evaluating the liquid transport properties of a surface may actually provide more relevant information. Indeed, once steady state is reached the complex phenomenon of motion of the halo wetting front along the rough surface has halted. Thus our measurement is no longer tainted by the details of the wetting mechanism, which are unimportant for the transport of liquid to the tip of an already wetted electrospray emitter tip.

In the work presented here, the water-halo-size test was used qualitatively as a pass/fail criterion for a surface treatment. For this method to be applicable quantitatively, two precautions should be taken. First, the ambient gas around the sample should be controlled so that the evaporation rate is consistent. Indeed, we observed that much larger halos could be obtained by placing the sample in a box, instead of leaving it in the well ventilated cleanroom air. Second, better precision will be obtained if the halo recedes before reaching steady state. This could be achieved by placing the sample in a confined environment, letting the drop spread, and returning the sample to a better ventilated environment. Indeed, the process by which the wetting front advances can be very subtle. By using a receding halo front for the measurement we will be sure that the halo did not get pinned, while advancing, in a configuration with low capillary pressure. Evaporation will prevent a receding halo from being pinned.

This method, as we have described it so far, only applies to flat surfaces. To evaluate a treatment on an inclined surface, a wafer with concentric rings of inclined surfaces can be treated. A drop is then placed in the center of the concentric rings, and the measurement is carried out as usual.

3.3.4 Summary of Surface Treatment Development

We now present the highlights of the experiments that were carried out on the LAM 490-B in search of a suitable surface treatment. We tried treating both whole wafers and individual 1×1 cm dies. When dies were treated, they were set on a handler wafer and run through the system. In a few cases, the handler wafer contained a recess a few hundred micrometers deep to prevent the dies from sliding around, this precaution proved unnecessary, and the recess apparently made no difference to the surface treatment.

The reason for working with dies was to to reduce the cost of performing tests (preparing each wafer with emitters on it takes a day of processing in the best of cases), and to allow different treatments to be applied to different emitter dies in the final thruster, to allow comparison of their effectiveness during electrospraying.

- Some exploration was attempted around the Cl_2-He treatment from Velásquez-García. Some surface treatments with good wicking properties on flat surfaces were found, but they only worked on the central region of the wafer. They did not produce wettable dies.

- The black silicon tuning method from [87] was applied for an SF_6-O_2 plasma.[4] Starting with an existing etch recipe, we significantly increased the O_2 flow rate, and progressively decreased it until the smooth surface turned dark. The corresponding parameters are listed in Table 3.1 as recipe R_0. This treatment wetted very well when applied to whole wafers. With pieces, a deep black surface was produced which did not wet at all. Under SEM observation, micron-scale features were visible for the whole-wafer treatment, while the treated die revealed texture at the hundred nanometer scale, barely resolved by the microscope that was used. We suspect that heat transfer may have been the determining difference between pieces and wafers. Indeed, the treatment worked well in whole wafers with oxide on the backside, excluding electrical effects as an explana-

[4] We also tried to transpose the method for Cl_2-He plasmas, and even for other settings of an SF_6-O_2 plasma without success. So there is definitely still some luck involved.

tion, and placing dies in a recessed handler to eliminate the height difference did nothing to improve the treatment.

Setting	Recipe R_0	Recipe R_1
Pressure (mtorr)	200	200
Power (W)	400	400
Electrode Gap Size (cm)	0.5	1
O_2 Flow (sccm)	50	50
SF_6 Flow (sccm)	40	40
Duration (min)	5	5

Table 3.1: Black silicon recipe for LAM 490-B

- After some exploration around R_0, we found that increasing the gap between the electrode and the wafer yielded dies with good wetting properties on flat surfaces. This is recipe R_1 in Table 3.1. Unfortunately, when this recipe was applied to dies with sloped features, the treatment only worked up to inclinations of about 45°, as shown in Figure 3-13. Despite much exploration around R_1, we never managed to increase this inclination. Making emitters with angles that do not exceed 45° without DRIE etching is not very compelling because of the 2:1 aspect ratio of the SF_6 etch. Indeed, Equation (3.9) indicates that the tip-to-total height ratio in that case is 11%, meaning that we will have a very short tip that is only 11% of the height of the recess that gets etched in the wafer. Thus this treatment would be very difficult to use to treat emitters.

- Giving up on dies, we tried R_0 on a die with inclined features, and found that the treatment is able to reach at least 70°, as shown in Figure 3-14. However, the treatment is not conformal, and does not work on vertical or inverted features, as the roughness is always oriented in a vertical direction. In all the SEM observations we made, we never saw any roughness that did not have an up-down orientation. Unfortunately, no SEMs of treated vertical sidewalls from Velásquez-García were available for comparison.

Treatment R_0 is not perfect as it does not work on vertical structures. Nevertheless, it has roughness at the right scale, and can approach verticality. At low

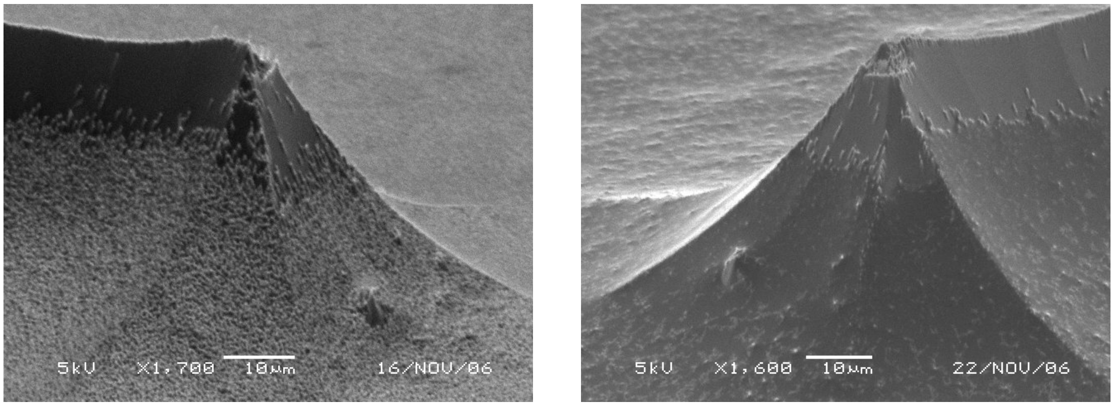

(a) Dry surface

(b) Wet surface

Figure 3-13: Same emitter with and without liquid. Little cones, seen as texture, form on the emitter surface up to an inclination of about 45°. Liquid rises only as high as the cones

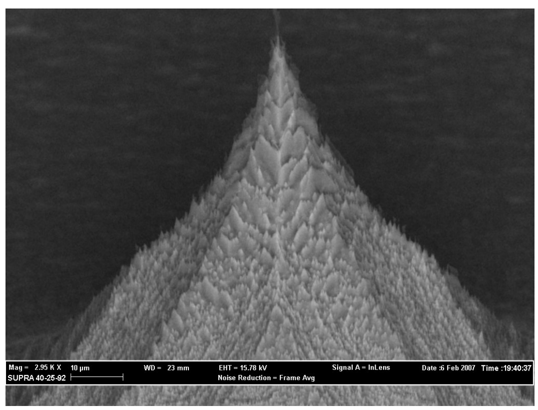

Figure 3-14: Surface treatment reaching at least 70°

inclinations, the treated surface appears to be covered with cones, almost exactly as in the theoretical discussion in Section 2.2.3. Figure 3-15 shows a high magnification SEM of the cones-like structure of the surface treatment on a horizontal surface. At surface inclinations, the cones seem to group into rows, with flow channels between rows.

This treatment did not work on wafers bearing a photoresist mask, but did work with oxide masks. It appears that the existence of a native oxide is important for the process. Indeed, when the treatment was applied to wafers straight out of a plasma etcher, the wafer remained smooth. If the wafer was then dipped in water, dried and

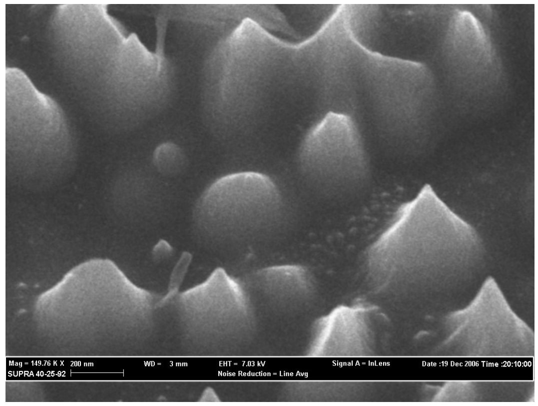

Figure 3-15: Closeup of the surface roughness showing cones with micron-scale spacing between them

retreated, the treatment worked perfectly.

When the first emitter dies were finally completed, taken from the cleanroom, and applied with EMI-BF_4, they did not wet at all. It was realized (after a few hours of profound worry) that the last processing step done on the emitters was an HF oxide strip. It is well known that smooth silicon emerging from an HF bath is hydrophobic (nonwetting). Placing the wafers in an oxygen plasma for a few seconds was sufficient to restore their hydrophilic nature, after which they wetted without difficulty.

3.4 Compromise Emitter Geometry

We have now seen how to make emitters with nice vertical sidewalls, and how to surface-treat wafers with no vertical sidewalls. Clearly some compromise is needed to make emitters that are treatable. This section describes the detailed design of the emitters that were finally used in the thruster.

The simplest solution to this problem is to use only SF_6 etching to make mountain shaped emitters. However, this method offers poor tradeoffs. We need an emitter that is tall for field enhancement. We need the tip depth to be as low as possible. We need to limit the inclination at the tip for the surface treatment to work, and for the sensitivity to process variation to be reduced. To be sure that the treatment which is known to work up to 70° will work, we take the tip half-angle to be 20°. Equation (3.9) with $f = 2$ gives a tip-to-total height ratio of 41%, which is still rather low. Increasing this angle is risky because we are unsure of the surface treatment, and besides, the hourglass effect from Figure 3-5 will be troublesome at higher inclinations.

We have already mentioned nested masks as a possible solution to this problem, however time constraints made us prefer an alternate approach which uses DRIE etching without leaving any vertical sidewalls.

3.4.1 Tapered Etch Recipe

The fact that vertical sidewalls must be avoided does not necessarily preclude the use of DRIE etching. Indeed, when a vertical sidewall has an inclined surface above it, the top of the wall gets etched away during SF_6 etching, whereas the base of the wall stays at a fixed height (see Figure 3-16). Therefore, a low enough vertical wall will disappear given sufficient subsequent SF_6 etching. This fact implies that we can introduce some DRIE etching into the emitter process.

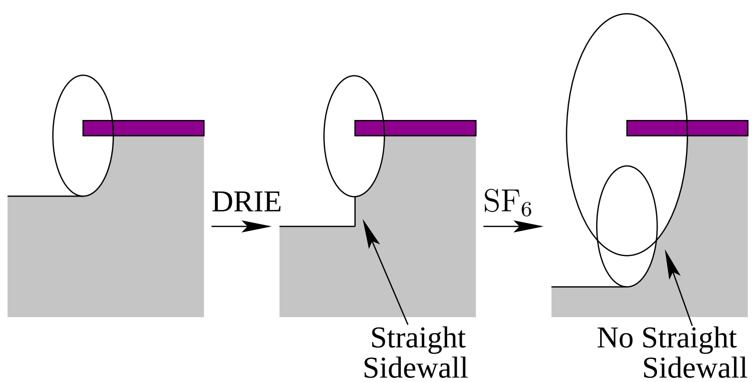

Figure 3-16: A short vertical sidewall can disappear with SF_6 etching

Now, consider the set of circles in Figure 3-17. The centers of the circles are aligned and equally spaced out. The radii of the circles decrease by a constant amount when moving from one circle to the next. The envelope of these circles is a straight line. Referring back to Section 3.1.3, we notice that if we stretch the circles into ellipsoids, we find that the envelope of the circles is exactly the sidewall configuration that would be obtained for a tight alternation of SF_6 and DRIE etches. Any sidewall angle can in theory be obtained by this method, and by varying the proportion of SF_6 and DRIE etching with depth, a wide range of non-straight profiles can also be obtained. The only caveat, is that the top of the profile is set by the largest ellipsoid, so if there is a maximum allowed inclination, the top of the profile will have to be etched away, increasing the tip depth. Equation (3.7) gives the tip depth, and, for a tapered emitter, the tip height satisfies $h_t \tan(\alpha_t) = R_m$. Thus, the tip-to-total height becomes

$$\frac{h_t}{d_t + h_t} = \frac{1}{1 + f^2 \tan^2(\alpha_t)} \quad \text{for a tapered emitter.} \tag{3.10}$$

To demonstrate the improvement due to the tapered etch, Figure 3-18 compares Equation (3.9) for the pure SF_6 etch with Equation (3.10). The tapered etch has a tip-to-total depth ratio of 65% at 20°, a marked improvement over SF_6 only (which was at 41%).

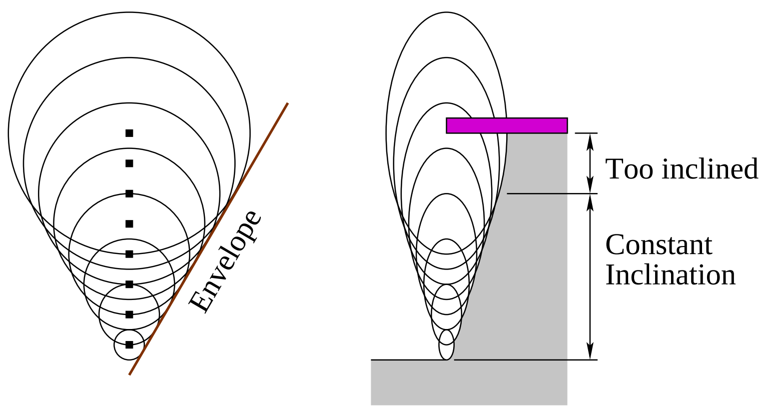

Figure 3-17: A constant-angle taper etch can be made by tightly alternating SF_6 and DRIE etching

Rather than implement an alternation of dozens of DRIE and SF_6 etches, we use the fact that DRIE is already an alternation of passivation and etching steps to create the same effect. We shortened the passivation step of the DRIE recipe, and increased the etching step in such a way that the sidewall passivation runs out before the end of the etch step. In this way we effectively get a single DRIE cycle followed by a short SF_6 etch, which is exactly the alternation we were looking for. Table 3.2 gives the

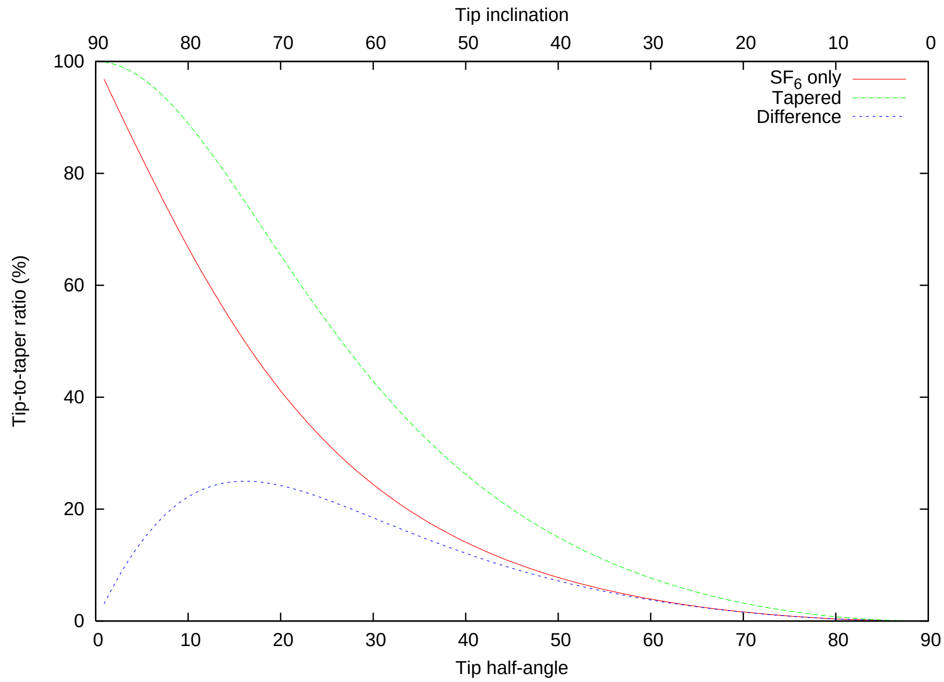

Figure 3-18: Comparison of the tip-to-total depth ratio for the plain SF_6 and the uniform taper emitters when $f = 2$

Parameter	BGJBMOD		JBETCH	
Pressure (mtorr)		28		28
	Etch	Passivate	Etch	Passivate
Duration (s)	30	1.9	13	11
C_4F_8 Flow (sccm)	0	35	0	50
SF_6 Flow (sccm)	105	0	105	0
Coil Power (W)	800	600	800	600
Platen Power (W)	5	6	15	6

Table 3.2: The tapered etch (BGJBMOD) compared with the straight etch (JBETCH) it was derived from

details of the recipe that was used.

An even more promising possibility that we did not have time to explore was to tune the DRIE etch to have positive taper without undercutting the oxide mask. Indeed, this would in principle allow an arbitrarily small mask to be used, and hence, an arbitrarily small amount of lowering of the tip below the level of the unetched silicon. This type of etch should be easy to obtain, at least for angles not too far from vertical. Parameters that that control the inclination of the ions bombardment should be adjusted, such as increased platen power to increase the vertical velocity component, or decreased pressure to reduce scattering of ions before they reach the substrate [55].

3.4.2 The Final Emitters

Two types of emitters were finally fabricated. The masks and emitters are shown in Figure 3-19. To generate the simulated emitter shapes, the tapered etch was modeled as a tight alternation of DRIE and SF_6 etches.

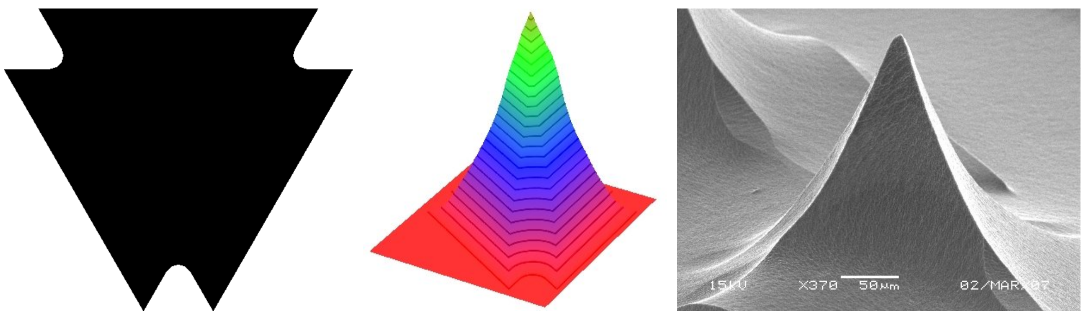

(a) The single tip, 500 μm/tip emitter

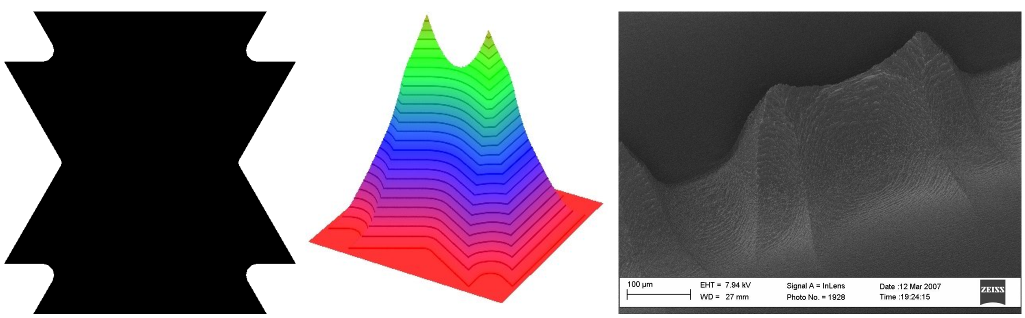

(b) The double tip, 275 μm/tip emitter

Figure 3-19: The final emitters, in theory and in practice

The first emitter type had a triangular mask from which the corners had been removed to improve packing density, producing an emitter shaped like a triangular pyramid. The designed amount of SF_6 etching to form a point was 175 μm. This relatively large value was selected because of uncertainty on the tapered etch at the time when the mask was designed. We wanted to guarantee that a point would not form before the emitter had reached a sufficient height. The triangular shape ensured that three ridge lines would meet at the tip, preventing blade formation. The etching proceeded in three steps: first an SF_6 etch that undercuts 50μm, then a tapered etch that was continued until the emitter was sufficiently tall, and finally an SF_6 etch that rapidly finished the tip formation at the target height.[5]

The first type of emitter had an emitter spacing of 500 μm, because of the size of the oxide mask. Indeed, neighboring masks must not touch, as this would lead to one of two problems: (1) an extra ridge line would be created connecting the two adjacent tips, or (2) the mask pattern would be asymmetric relative to the slot, and it would be difficult to ensure that the tips are actually along the slot centerline. However, it is possible to put two triangular masks together head-on to create a two-tip emitter with only three ridge lines per tip, which is what we did for the second emitter type.

Type (number of emitters)	1	216	502
Emitter type (tips)	1	1	2
Tip spacing (μm)	n/a	500	275
Rows	1	13	15
Row spacing (μm)	n/a	750	750
Array diameter (mm)	n/a	11	12
Array area (mm^2)	n/a	95	113
Emitter density (tips/mm^2)	n/a	2.3	4.4

Table 3.3: Three types of arrays were built

Table 3.3 describes the three types of emitter dies that were fabricated. The single-emitter dies were made so that data on an isolated emitter could be gathered. The 216 emitter dies are designed more conservatively than the 502 emitter dies; a smaller area was covered to guarantee that beam divergence would not cause beam interception by the Pyrex present in the final thruster (see Chapter 5). The 216 and 502 emitter arrays are shown in figures 3-20 and 6-38(a), respectively. Often, emitters at the edge of an array see a larger electric field because there is less shielding from neighboring emitters. To limit this non-uniformity, a slightly smaller single tip emitter

[5]The first 50μm etch is designed to reduce the sensitivity of the etch by reducing the angle at the tip compared with the tapered sidewall, at the cost of a slightly increased tip depth.

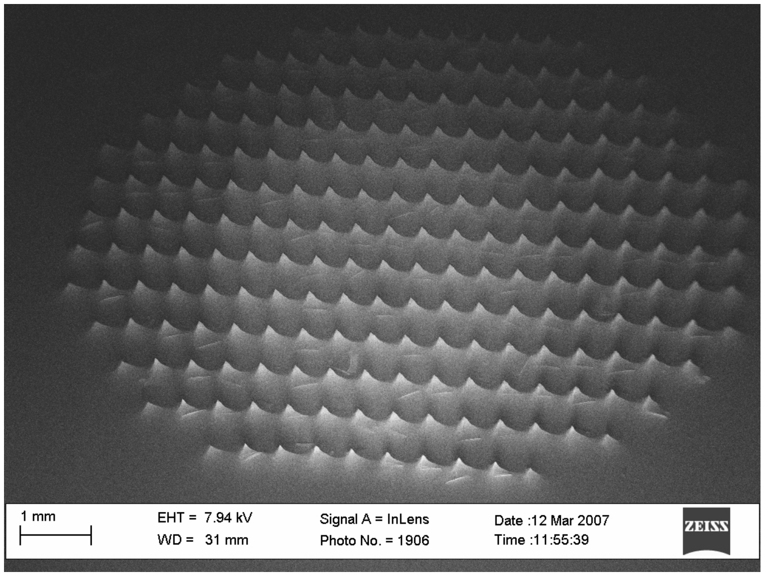

Figure 3-20: An array of 216 single emitters

was introduced at the end of each row.

The finished emitters have a height of about 250 μm, for a tip-to-total height around 60%. This value is 5% less than expected, which is excellent agreement given the processing non-uniformity, and the 50 μm undercut that was added to the start of the emitter process.

Looking at the fabricated dies in Figure 3-21(a), we found that the tips are not all aligned. Instead, they form a bimodal distribution with a shift of about 40 μm between the two maxima. One hypothesis to explain this bimodal distribution is that when the tip is formed, the oxide mask can fall to one side or the other of the emitter. The side on which it falls will be partially shielded from etching, causing the emitter to be lopsided. Figure 3-21(b) shows marks beside the emitters that could be caused by the mask after it has fallen off. These marks indeed appear on either side of the emitter. We still need to explain why, given that the mask has threefold symmetry, all the masks fall toward one of the two side faces of the emitter, and never toward the one which is lined up with the row of emitters. Here, the likely explanation is that the third face etches slower than the other two because of shielding from the mask of the neighboring emitter. Therefore, the tip is shifted toward the third face,

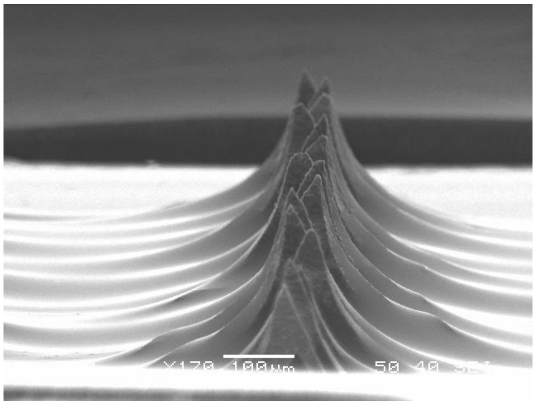

(a) Variation in tip position along an emitter row

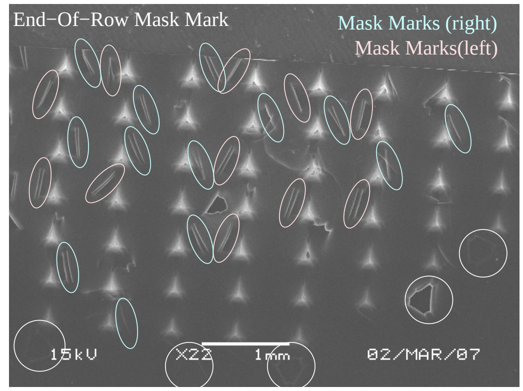

(b) Masks (marks left by them are circled) can fall to one side or the other when the tips (lighter areas) are formed

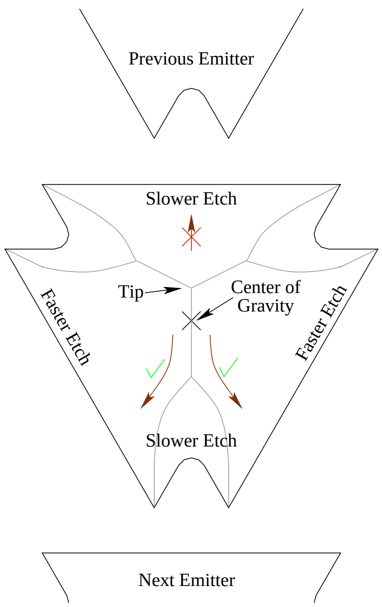

(c) By offsetting the tip, the neighboring emitter ensures that the mask does not fall toward it

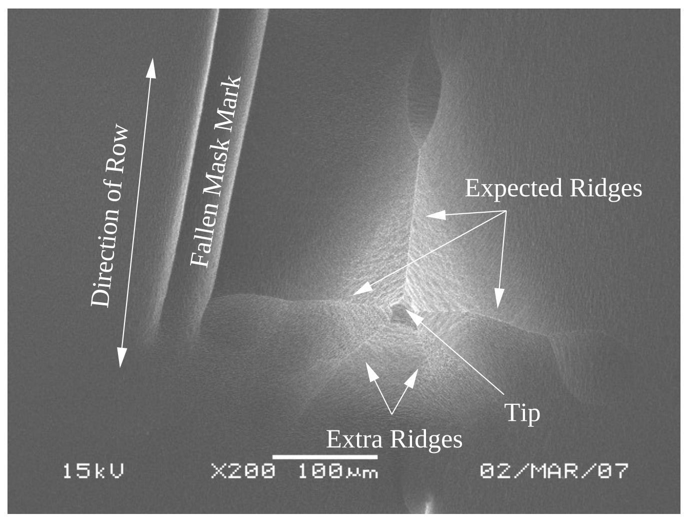

(d) The ridges on an emitter

Figure 3-21: Emitter tips do not all line up

and when the mask falls off, it falls away from the third face, and is then deflected by the opposite ridge to one side or the other, as shown in Figure 3-21(c). Observation of an emitter from above confirms that the neighboring mask has had an influence, as evidenced by some extra ridge lines, not predicted by our model, but which could be caused the sharp corners of the neighboring emitter's mask (see Figure 3-21(d)). Emitters at one end of a row, undergo the opposite effect. The neighbor that usually causes the tip shift is gone, and the remaining neighbor slightly pushes the tip the opposite way. Figure 3-21(b) shows that five out of eight emitters on the end of the row had the mask fall towards the missing neighbor position. These observations strengthen the case that it could be beneficial to strip the oxide masks just before they fall off if better tip centering is desired.

3.4.3 Wetting Results

When the emitted dies are wetted by placing a drop of EMI-BF_4 on the flat area beside the emitters, liquid spreads across the die. However, despite the surface treatment extending all the way to the top of the emitters, Figure 3-22(a) shows that liquid does not make it to the top of most emitters, though some emitters do appear to have some traces of liquid on them (see Figure 3-23). When the emitter is fired for the first time, it starts emitting at a higher voltage than is later needed to start operation. Looking at the emitter after it has been fired, reveals that liquid now extends all the way to the tips (Figure 3-22(b)). It appears that the capillary pressure driving the wetting front up the emitter is small, but that once the difficult initial wetting occurs, liquid flows well along the emitter surface.

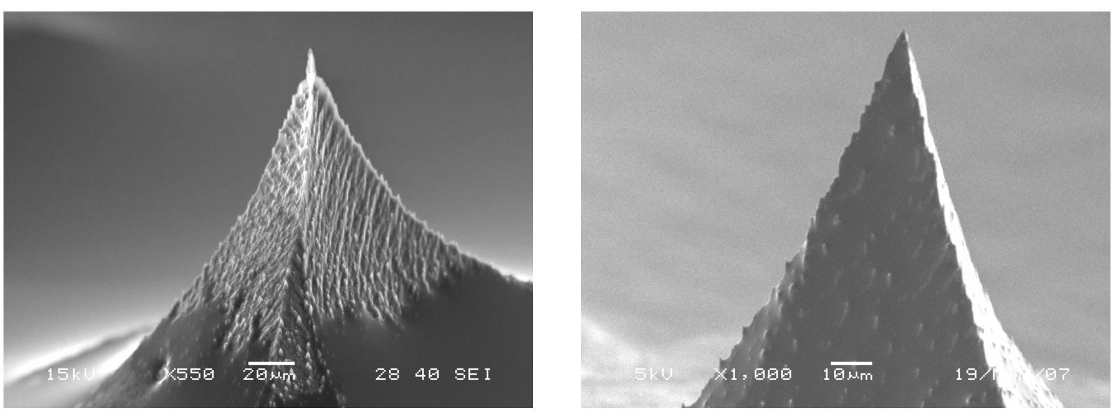

(a) Before firing (b) After firing

Figure 3-22: Emitters which have been treated still do not wet up to the tip when supplied with liquid. However, after firing the emitter is completely wetted, suggesting that the electric field was able to draw liquid to the tip of the emitter

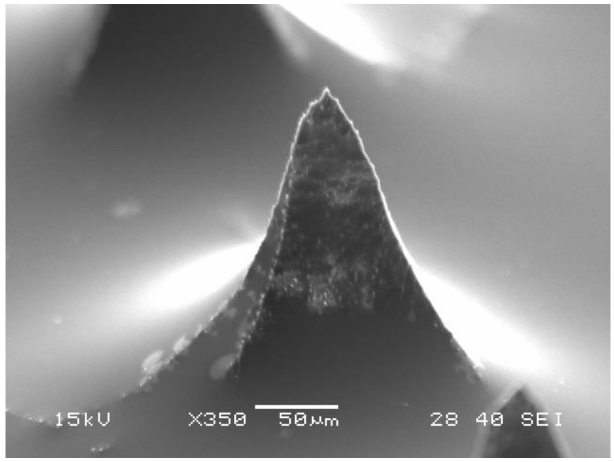

Figure 3-23: An emitter that partially wetted without electric field application

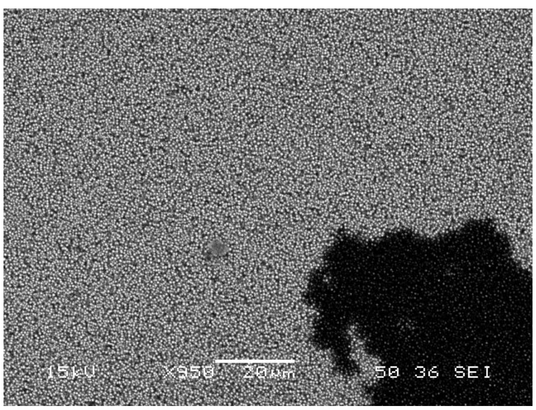

(a)

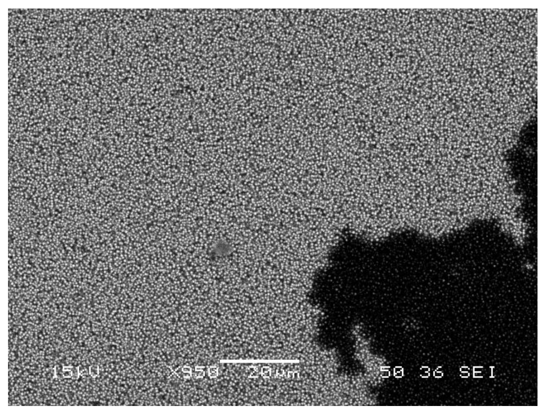

(b)

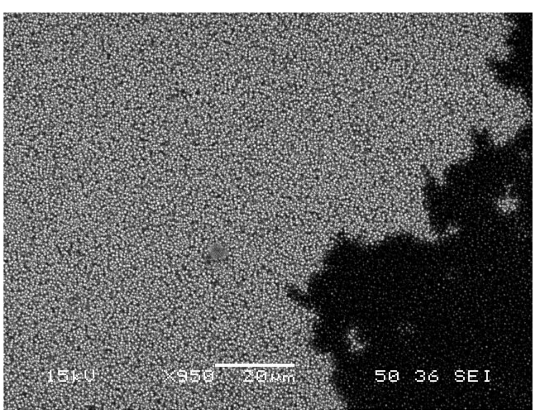

(c)

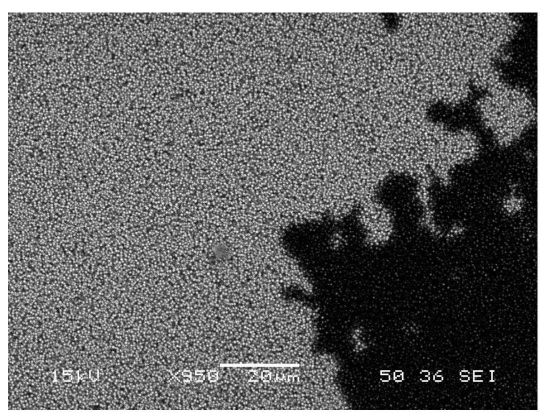

(d)

Figure 3-24: Liquid progresses irregularly across a previously unwetted surface when there is a very small liquid supply

We have been able to take SEMs of EMI-BF_4 progressing along a flat surface (see Figure 3-24). These SEMs show that the motion of the liquid front can be complicated and irregular, because the liquid has to find a path with sharp corners through the valleys of the surface treatment. Once this path has been found, the liquid level increases rapidly and formulations such as Darcy's equation (Equation (2.6)) become appropriate.

3.5 Concluding Remarks

In this chapter, we have extended the basic emitter geometry work from [37] in a number of promising directions. Increasing the number of alternations between SF_6 and DRIE leads to needles with smoother profiles. With hundreds of alternations, tapered sidewalls can be created, which enables tall emitters without vertical sidewalls. With nested masks, even more versatile structures should be possible.

All these extra degrees of freedom increase the difficulty of designing an emitter. The proposed emitter geometry model allows the designer to visualize the effect of varying the oxide mask shape and processing times. Despite its crude geometric assumptions, the model gives good qualitatively predictions of the emitter shape. A more detailed model would be needed to capture the interaction between neighboring emitters and the depletion of reactants when there is too much exposed silicon area, and to correctly predict the geometry of partially enclosed areas of the emitter geometry. However, because of the large amount of process variation across a wafer, emitter designs that rely on fine understanding of processing conditions are unlikely to be worthwhile. Designs that demonstrate robustness to process variation in the simplified model we have proposed will be easy to fabricate in practice, and in this respect, our model is sufficiently accurate.

The improved understanding of the emitter fabrication process has allowed improvements in tip geometry by avoiding fractured tips and blade shaped tips. It has also led to methods of mitigating the effects of process variation.

Unfortunately, most of the improvements in emitter geometry control have failed to lead to improvements in electrospray emission because of our inability to roughen vertical surfaces. Indeed, making a surface that can wick and transport liquid has proven to be one of the greatest challenges in this project. In the end we are able to treat surfaces up to at least 70° inclination, but vertical surfaces, or, worse, surfaces with slightly negative taper cannot be treated. All our electrospray firing has therefore been with tapered emitters, and the benefits of more complex geometries are unknown. In particular we do not know if emitters could benefit from liquid transport in channels with scales in the 10 μm range.

The most important fact is that emitters with good liquid transport properties were made so that the integration technologies that will be discussed in the rest of this thesis could be tested. These same integration technologies would apply just as well to almost any planar emitter geometry.

Chapter 4

Assembly Methods

In MicroElectroMechanical Systems (MEMS), it is common to find centimeter scale devices with functional features on the micrometer scale. In many cases, these devices can be made on a single substrate with features defined using photo-lithographic techniques. For some applications, however, the single substrate approach is not sufficient. This could happen when incursions into the third dimension of more than a few millimeters are necessary, when multiple substrate materials are needed, or when different parts of the device require incompatible process flows. In these cases, a number of components can be manufactured on independent substrates and then assembled. Often, this assembly must be made with micrometer-scale accuracy in order for the finished device to be functional.

We consider this assembly problem in the context of an electrospray array. The extractor electrode, which is electrically biased relative to the needles, poses a number of challenges. For each needle in the array, an aperture is needed to let the particles emitted from the needle leave the thruster. The accuracy with which these apertures are aligned determines how tightly needles can be packed onto the thruster, and hence the thrust density. The needles must be well aligned with the extractor aperture to limit the off-axis velocity component of the emitted beam. Therefore, accurate alignment of the electrode is critical. Moreover, large potential differences are involved, and to prevent short circuits, we need to limit the paths through which liquid can flow from the needles to the extractor. Thus, we have decided that the electrode and the needles should be separate components, connected only near the periphery of the thruster. The question we address in this chapter is how to assemble the extractor above the emitters.

4.1 Finger-based Assembly

We propose to design the extractor and emitter components so that they can be manufactured separately and then clipped together by hand. This approach is common with macro-scale plastic components, and we shall show that it extends well to the accuracy requirements of MEMS devices. In our approach (patent pending), the emitter die is held in place by an holder system comprising a number of springs, which we will call “fingers”. The person who assembles the device first roughly positions the emitter die into a recess on the top surface of the thruster (Figure 4-1.a), and then rotates it (clockwise in the figure). As the die rotates, the shape of the die holder forces the die to align itself to within 50 μm of the center of the thruster (Figure 4-1.b). Then it begins to flex the fingers. Finally, each fingertip falls into a notch in the holder, preventing further rotation (Figure 4-1.c).

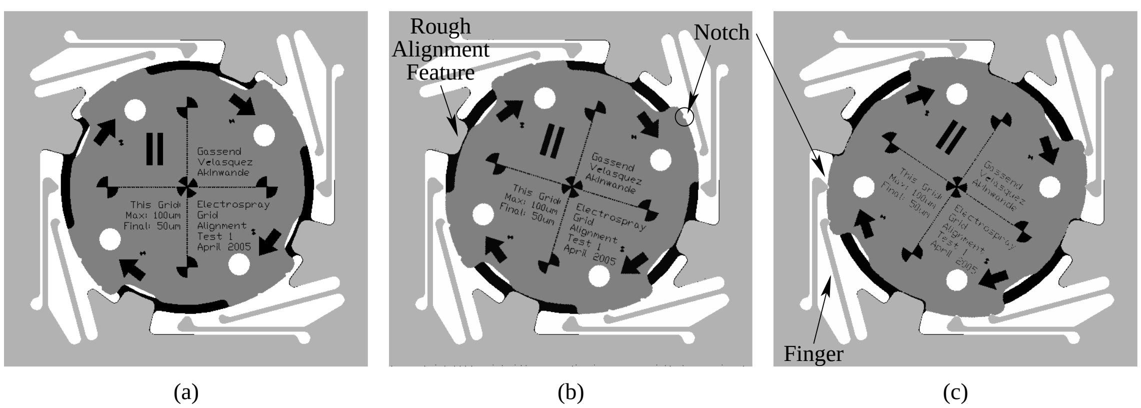

Figure 4-1: Hand assembly of the emitter die. First, the die is placed in its recess (a). As it is rotated, the die is forced into rough alignment (b). Finally, the fingers precisely clamp the die into place (c)

The idea of micro-machined clips is not new. They have been used for surface micro-machined components to build 3D structures from planar components [97, 98], but with components smaller than a millimeter, whereas we work with centimeter scale components. Planar snap fasteners have also been made [99], but they are also at a much smaller scale, and allow some play between parts once they have been assembled. The classic Micro-mechanical Velcro paper [100] shows how to fabricate micro-hooks on a silicon surface that mate in a way similar to Velcro. The original Micro-mechanical Velcro does not include positioning, but some enhancements provide for mating features that force the parts to be mated to assemble in a specific location [101]. Also, unlike conventional Velcro, the micromechanical variant cannot be disassembled without destroying the micro-hooks. Snap systems have also been used many times to attach optical fibers to substrates [102].

In our group, clips have already been used to assemble the extractor of an electrospray thruster with a linear array of emitters [36]. This clip method has been advanced as a flexible, precise and modular way to assemble MEMS devices [103]. The planar implementation improves over the linear implementation by a number of features. Notches for the fingertips prevent accidental disassembly of the emitter die. Once the die is assembled, it is firmly held in place by the high axial stiffness of the fingers, but the lateral flexibility of the fingers allows the device to accommodate for differential thermal expansion, or consistent over/under-etching of the side-walls during fabrication. In the linear thruster case, the use of clips was justified by the three dimensional nature of the assembly, where the plane of the extractor electrode was perpendicular to the substrate on which the needles were located. In the planar case, wafer bonding techniques could be used to bond the extractor and emitters together. We believe that clip based methods are superior to wafer bonding for a number of reasons. They are non-permanent, allowing different electrode designs to be tested. They do not depend on the nature of the materials to be bonded, leaving more flexibility in the design. They do not involve high temperatures, which allows temperature sensitive coatings to be used. Finally, the hand assembly step requires no special equipment, which improves manufacturing flexibility.

Electrospray thrusters are not the only application of this assembly method. It could be useful to assemble the extractor in any kind of electrospray device. Also, if the positional accuracy can be slightly improved, clips could be used to assemble the extractor in a field ionization device. In the latter case, the open architecture, which leaves room for gases at the periphery of the device to penetrate between the electrode and the emitters could improve the flow of material in the device. Electrostatic lenses could be assembled with our method, in which case they could be easily replaced if a different device configuration is necessary.

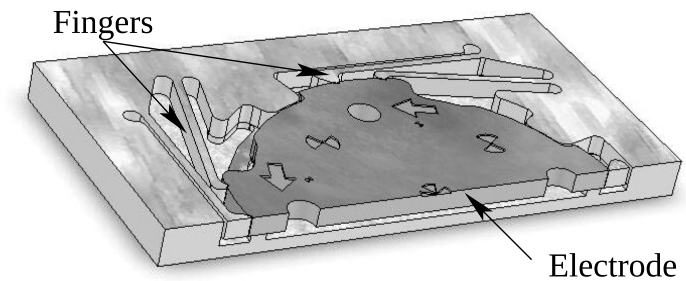

Figure 4-2: Section view of the device

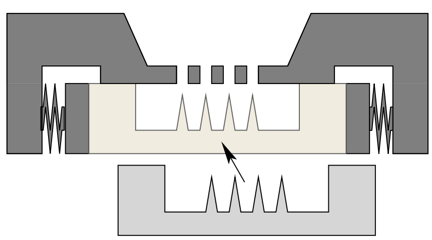

(a) Artist's view of complete thruster

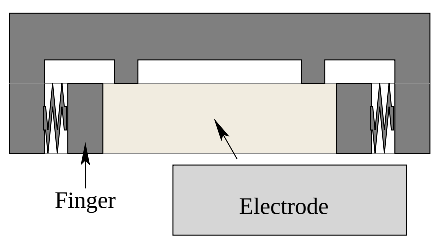

(b) Side diagram of the test device

Figure 4-3: Diagram of an electrospray thruster, compared with the simplified test device we first built to evaluate the alignment method

Figure 4-3(a) shows what an electrospray thruster array using this alignment method could look like. In the rest of this chapter we will discuss the design, fabrication and characterization of a simple test device (shown in Figure 4-3(b)) comprising only an electrode, an electrode holder, a base, and alignment marks. This test device was used to evaluate the assembly method we are proposing. In chapter 5 we will show how this assembly method was integrated into a complete thruster.

We will begin this chapter by describing the design of the test device, in Section 4.2. Then, in Section 4.3 we describe the process that was used to fabricate it.

Section 4.4 covers the characterization that was performed on the device once it was fabricated. Finally, in Section 4.5, we will look at two other assembly methods which were evaluated as part of this thesis.

4.2 Device Design and Modeling

In this section we describe the design of our assembly method. The driving features in the design of this method were:

1. Ensure stiff and accurate positioning of the assembled parts.
2. Be flexible: allow component materials to be changed, avoid restrictive high temperature steps during assembly.
3. Be robust, i.e., difficult to break, tolerant to process variation, easy to assemble.

4.2.1 Finger and Notch Design

The initial choice of a finger-based mechanical assembly method was guided by the desire for flexibility. As long as the electrode has the correct shape and sufficient mechanical properties, it will be possible to assemble it.

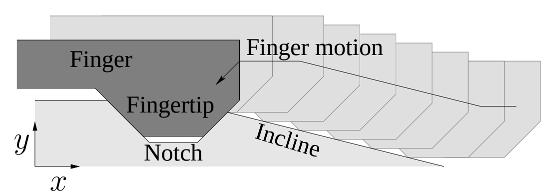

Figure 4-4: The motion a finger follows during assembly

Figure 4-4 shows the motion a finger follows during assembly. The parts to be mated are translated relative to one another in the x direction. Initially the finger is free, then it comes into contact with the part to be mated to. It is slowly deflected by a gentle incline, and then finally clicks into its notch. In the notch, the finger is still deflected, so there is always a force clamping the fingertip in its notch.

We place the finger so its direction of deflection is the y direction, perpendicular to the direction of relative motion of the two parts to be assembled. This prevents the clamping force from creating a systematic shift in the finger's final resting position. Moreover, the finger has low stiffness in the y direction (beam in flexion) and high stiffness in the x direction (beam in tension/compression). Thus, by placing the stiff direction of the finger parallel to the direction of assembly, we provide a high stiffness in the direction of assembly.

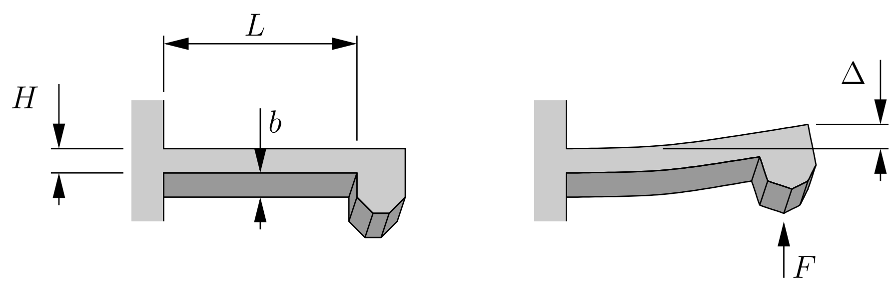

Figure 4-5: Notation used for describing finger geometry and flexion

The interface between the fingertip and the notch consists of two inclined planes. The angle of inclination is the same on both sides. This way, any systematic over or under etching will change the final deflection of the fingertip, but will leave the relative x positions of the parts unchanged. For steep angles of the interface planes, a small x error translates to a large deflection in the fingers, so there is a large force trying to center the fingertip in its notch. Also, a smaller coefficient of friction is needed at the interface to prevent the finger from slipping out of its notch under the influence of a force in the x direction. However, steep walls lead to higher sensitivity to process variation since a small over or under etch corresponds to a larger change in deflection for the assembled fingertip. We chose 45° as a compromise angle for the contact plane inclinations.

The amount of motion the finger undergoes during assembly is another tradeoff. As we shall see, for a given clamping force, the finger size increases with its maximum allowable deflection. This favors small deflections. But the deflections have to be large enough to guarantee proper assembly even when process variations are considered. We have decided to be robust to over/under-etching of up to 10 μm. This means that the assembled deflection should be at least 25 μm to accommodate the worst case under-etching with 5 μm to spare. For comparison, in our fabricated devices, we have experimented with different deflections, as shown in Table 4.1. We have also targeted a clamping force in the hundreds of millinewtons range.

Finger geometry, base width (μm)	Maximum / assembled deflections (μm)	Post geometry
I. Tapered, 700	A. 50 / 25	8 circular posts
II. Straight, 550	B. 100 / 50	Complete ring
III. Tapered, 1000	C. 150 / 75	

Table 4.1: Finger geometry, deflection and post geometry were varied independently across dies

To size the fingers, we use the analysis from [103]. The notation is introduced in figure 4-5. We model the finger as a slender uniform cross-section beam, fixed at one end, and with a point force applied at the other end. If E is the Young's modulus of the beam, its stiffness and maximum stress $\sigma_{\max}$ are given by

$$k_\perp = \frac{F}{\Delta} = \frac{E\,b}{4}\left(\frac{H}{L}\right)^3 \qquad \text{and} \qquad \sigma_{\max} = \frac{3EH\Delta}{2L^2}. \tag{4.1}$$

Eliminating H in Equations (4.1) yields an equation that clearly shows the tradeoff between high force, large deflection and small beam length/footprint:

$$\frac{F\Delta^2}{L^3} = \frac{2\,b\,\sigma_{\max}^3}{27E^2}, \tag{4.2}$$

where the right hand side is only determined by the substrate we use to make the fingers. For a 650 μm thick silicon wafer, $b = 650$ μm and $E = 145$ GPa [77]. Limiting the maximum stress to a conservative value of $\sigma_{\max} < 100$ MPa (this is ten times less than can be expected from [78]), we find that $F\Delta^2/L^3 < 2.3$ mN/m.

In our design, we wanted to fit 6 devices on a 6" wafer, which limits the finger length to about a centimeter. We opted for a maximum displacement of 100 μm. Equation (4.2) then limits the force to about 0.2 N, and Equations (4.1) give the corresponding slenderness ratio $L/H \approx 20$. The actual final dimensions we settled on for our uniform cross section fingers were $L = 1.09$ cm, and $H = 550$ μm.

The limitations imposed by Equation (4.2) only apply for uniform cross-section beams. Improved tradeoffs are possible for tapered beams, which are thick near their base and skinny further out where the moment is reduced. In the optimal beam geometry, the maximum stress is reached all along the beam, and the beam width H goes like the square root of the distance to the point where the force is applied. For simplicity, we have considered linearly tapered fingers (Figure 4-6). We found that the best tradeoff was when H decreases by 71% across the finger. We fabricated two different tapered finger sizes (Table 4.1). The type I finger is 700 μm wide at its base, and can provide the same force as the linear beam (type II) with 50% extra deflection. The type III finger is limited to 100 μm deflections, but nearly three times as stiff as the linear beam.

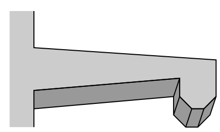

Figure 4-6: A linearly tapered beam

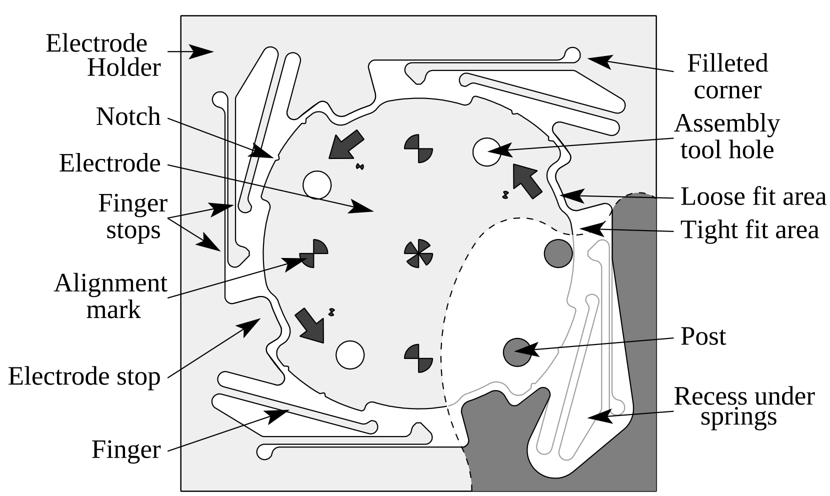

Figure 4-7: Major features of the assembly system. Part of the top layer has been cut away

4.2.2 Maximizing Robustness

Now that we understand how to design fingers, we can use them to build a robust assembly system. First, we have chosen to assemble the electrode by a circular motion, with a number of benefits. If there are 3 or more fingers, then as long as the fingertips remain seated in their notches, the stiffness of the electrode will be determined by the axial stiffness of the fingers, as opposed to the much lower flexural stiffness. Moreover, in the circular configuration, the electrode is held only by fingers (as opposed to the linear case where it could be supported by a hard contact on one side as in [103]). Thus, any systematic over/under etching, or any differential thermal expansion, simply causes the electrode to appear slightly larger or smaller than expected. This change in size can be accommodated by flexing the fingers with no first order error in electrode position. In fact, differences in etching between two fingers or notches still does not introduce any position error, as long as each fingertip or notch is consistently etched.

The most fragile part of this device is the long slender fingers. They are flexible, but can only tolerate a few hundred micrometers of deflection. Whenever possible we have tried to protect them from accidental damage. The reader can use Figure 4-7 to identify the features we shall now describe, and Figure 4-1 to follow the assembly steps. The first line of defense is to place the fingers on the electrode holder, where they are better sheltered than they would be on the periphery of the electrode. This keeps them protected until assembly time. When the electrode is first placed within the electrode holder, the person who is assembling only needs to place it to within about half a millimeter of its intended location for it to fall into the same plane as the electrode holder. From this position, the electrode can be wiggled around without risk to the fingers, as electrode stops prevent it from coming into contact with the fingers. As the electrode is rotated into place, the parts of the electrode near the electrode stops get wider, forcing the electrode into alignment to within 50 μm, making it impossible to wiggle the electrode beyond the breaking point of the fingers. The electrode stops also prevent it from being rotated backward, or more than 50 μm past the normal assembled position. Moreover, the fingers are placed so that if they are excessively deflected outward, they will hit a hard stop (finger stop), preventing them from flexing farther and breaking. We have also tried to maximize the electrode strength by avoiding sharp corners throughout the design. Finally, in case some fingers do break, we have decided to have 8 fingers to provide redundancy. More fingers could be used, but they would occupy valuable surface area, and contribute little to reliability, since we never lost more than two fingers in our experiments.

4.2.3 Expected Performance

We have mentioned that our clamping method offers a high clamping stiffness. Indeed, consider that the electrode undergoes a small translation (t_x, t_y), and assume that the fingertips remain in their notches without buckling or snapping. The force that is exerted by a finger with an axis along the x direction is determined by the axial and flexural stiffnesses of the beam $k_\parallel$ and $k_\perp$.

$$\begin{pmatrix} F_x^1 \\ F_y^1 \end{pmatrix} = -\begin{pmatrix} k_\parallel & 0 \\ 0 & k_\perp \end{pmatrix} \begin{pmatrix} t_x \\ t_y \end{pmatrix} \qquad \text{where} \qquad k_\parallel = \frac{Eb\bar{H}}{L} \tag{4.3}$$

For a set of n fingers evenly placed around the device, the total force is then:

$$\begin{pmatrix} F_x \\ F_y \end{pmatrix} = -\sum_{i=0}^{n-1} \begin{pmatrix} \cos\frac{2\pi}{n} & \sin\frac{2\pi}{n} \\ -\sin\frac{2\pi}{n} & \cos\frac{2\pi}{n} \end{pmatrix} \begin{pmatrix} k_\parallel & 0 \\ 0 & k_\perp \end{pmatrix} \begin{pmatrix} \cos\frac{2\pi}{n} & -\sin\frac{2\pi}{n} \\ \sin\frac{2\pi}{n} & \cos\frac{2\pi}{n} \end{pmatrix} \begin{pmatrix} t_x \\ t_y \end{pmatrix} \tag{4.4}$$

If $n > 3$, the sum simplifies to an isotropic stiffness of $k = n(k_\parallel + k_\perp)/2$. If the fingers are not evenly spaced around the device, but can be partitioned into groups of symmetrical fingers, then this result still applies. Since $k_\parallel \gg k_\perp$ for a slender beam, a finger with average width $\bar{H}$ has a clamping stiffness of about:

$$k = \frac{nEb\bar{H}}{2L} \tag{4.5}$$

For our 8-finger devices and type II fingers, we get $k \approx 21$ N/μm. If we assume that the contact between the fingertip and its notch is frictionless, then a finger will pop out of its notch when the force it is transferring is equal to its clamping force. For our device, this limit occurs for a total force around 1 N, so the displacement of the electrode when the fingers start sliding out of their notches is around 50 nm. Thus, as far as we are concerned, the electrode is held in place rigidly. To put into perspective the 1 N maximum force on the electrode, consider that the electrode's mass is around 0.5 g, so the electrode should be held in place for accelerations up to 200 g. Similarly, we can compute the rotational stiffness of the assembly:

$$k_\theta = \frac{nEb\bar{H}r^2}{L}, \tag{4.6}$$

where r is the distance from the center to the fingertip. It differs from the translational stiffness by a factor $2r^2$. The factor 2 occurs because in a rotation all the fingers are identically loaded along their axial direction, and the r^2 factor arises from converting forces and linear displacements into torques and angular displacements.

As a final check, we must verify that the fingers are not buckling or shearing. Indeed, we find that these failures occur for 100 N (from the Euler condition) and 35 N respectively, well after the fingertips have popped out of their notches.

4.3 Fabrication of the Assembly Test Device

We now describe the fabrication of the devices, which took place at the Microsystems Technology Laboratories (MTL) at MIT. Figure 4-8 shows the main steps in fabricating the device. Six devices are fabricated from two 6" diameter, double side polished, 650 μm thick silicon wafers. First, an oxide is grown to protect the wafer surface for bonding, and alignment marks are patterned on both sides of each wafer (step 1). The *top* wafer is etched through with Deep Reactive Ion Etching (DRIE) to make the electrodes and electrode holders of the device (step 2a). A recess is DRIE-etched in the *bottom* wafer (step 2b), to be closer to an actual device, and to make recesses below the fingers that prevent the fingers from bonding to the bottom wafer. The bottom wafer is bonded to the top wafer, simulating the area where the emitters are located, and providing alignment marks to evaluate the device (step 3). Finally, the completed devices are ready for hand assembly of the electrode (step 4).

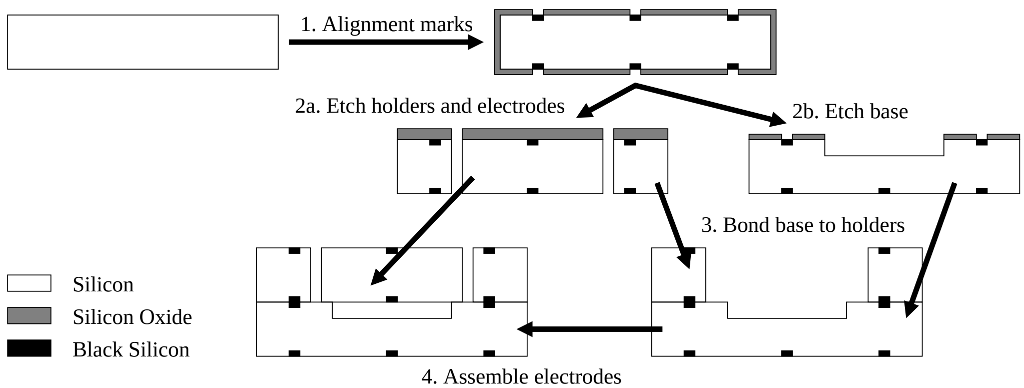

Figure 4-8: Fabricating the test device

This process requires 3 photo-masks. One for all the alignment marks that appear on the device, one to describe the recess to be etched in the bottom wafer, and one to delineate the electrode and electrode holder on the top wafer. Indeed, since the electrodes fit in the empty space inside the electrode holders, we were able to combine the electrodes and electrode holders on a single mask, removing a third of the processing steps. Six devices are placed on each 6" wafer, equally spaced around the outside of the wafer. The central region of the wafer is left empty to facilitate the onset of wafer bonding.

To increase our chances of getting working devices we introduced small variations between the devices (see Table 4.1 in Section 4.2.1 for details). Finger geometry was varied, trading off clamping force and maximum stress in the finger, and the electrodes were sized to cause varying amounts of deflection during assembly. At assembly time, any combination of finger geometry and deflection can be chosen because one is determined by the electrode, and the other by the electrode holder. Finally, two different geometries were explored for the posts that support the electrode in the vertical direction: small circular posts, and a ring support that goes all around the device (in case the small posts get caught on electrode features during assembly). In the end both post designs worked flawlessly so our worries were unjustified.

4.3.1 Process Details

Initially, all the wafers are processed identically. A protective oxide is grown on all the wafers to prevent surface damage or contamination which could hinder wafer fusion bonding. Then, alignment marks are successively patterned on the front and back of each wafer. To pattern alignment marks on a side of the wafer, thin photo-resist is spun on both sides of the wafer, and one side is exposed to the alignment mark mask. After development, buffered oxide etchant (BOE) is used to etch through the oxide where the alignment marks should be, then a chlorine-based plasma treatment is used to blacken the underlying silicon. Finally, the photo-resist is stripped using an oxygen plasma and a piranha (H_2O_2/H_2SO_4 1:3) bath. After this initial preparation step, the wafers are separated into tops and bottoms, which undergo separate processing.

The bottom wafers are coated in 10-μm-thick photo-resist which is patterned with the desired recess shape. Then, the exposed oxide is etched with BOE, and a recess about 100 μm deep is created using DRIE. At the end of the DRIE etch, a 30 second isotropic SF_6 etch is performed to smoothen the sidewall roughness and reduce the risk of crack formation [78]. Then, the photo-resist is stripped with an oxygen plasma and piranha.

In designing the process, we were particularly careful to obtain the straightest possible side-walls for the electrode and fingers, as it is the contact between the two that determines the precision of assembly. For this reason, our first iteration of the process used an aluminum etch mask to mask the DRIE etch that defines the top wafer. Unfortunately, we had photo-resist peeling problems when wet etching the aluminum and underlying oxide. Therefore, for the second iteration of the process, we use a more conventional oxide mask, which gives us excellent results. First, we deposit 5 μm of oxide using Plasma Enhanced Chemical Vapor Deposition (PECVD) followed by a one hour anneal at 950°C in a nitrogen atmosphere. Then we spin on a layer of 10-μm thick photo-resist and pattern it with the shape of the electrodes and electrode holders. The exposed oxide is etched in an Ar, CHF_3, CF_4 plasma. The wafer is mounted using photo-resist onto a quartz handler wafer. Then DRIE etching is used to etch through the wafer. Once the pieces have been cut out, and slightly smoothed by an isotropic SF_6 plasma etch, they are dismounted from the quartz wafer with acetone, and cleaned with an oxygen plasma and piranha.

Now we bond the top and bottom wafers together (except the electrodes, which are left aside). First, we clean the wafers using an oxygen plasma and piranha. Then we strip the protective oxide with hydrofluoric acid (HF), do an RCA clean (without HF). The wafers are then aligned, fusion bonded together, and annealed at 1050°C in a nitrogen atmosphere. Finally, the die saw is used to separate the 6 finished bases, which are then ready for assembly.

4.3.2 Alternate Device Configurations

To study electrical insulation between the electrode and the base, and to demonstrate the versatility of our assembly method, we produced a number of variants of the device we just described. They will be described in more detail in Section 5.1.3.

4.4 Experimental Results and Discussion

Two device batches were built, assembled and measured (Figures 4-9(a) and 4-9(b)). The first one uses an aluminum etch mask to define the spring wafer, the second one uses an oxide mask. After some general observations on our experience assembling and disassembling devices, we will look in detail at how accurately the electrode is located relative to the base, and present methods to insulate the electrode from the base.

4.4.1 Assembly, Disassembly, Robustness

We assemble the electrode by manually placing it into its recess. Then we used a small laser-cut plastic tool to assemble the devices (Figure 4-10). It engages the electrode's assembly tool holes (Figure 4-7), making it easier to apply torque. The torque needed to rotate the electrode into place depends perceptibly on the roughness

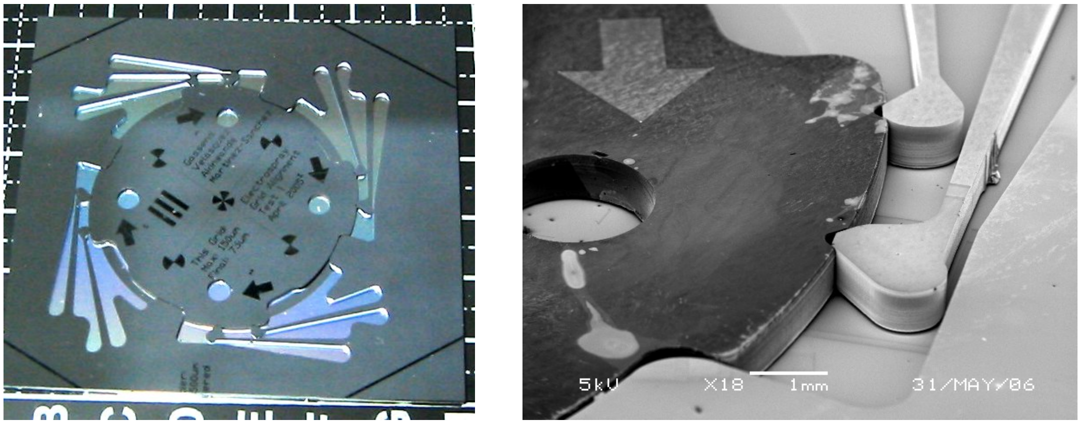

(a) Picture of completed test device (b) Scanning Electron Micrograph of an assembled finger

Figure 4-9: The completed test device

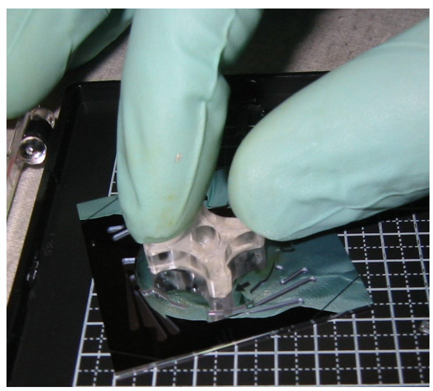

Figure 4-10: A device being assembled by hand with the help of a small plastic assembly tool

of the side-walls, the stiffness of the springs and how much they are deflected. When assembly is complete, the fingers fall into place, usually with a satisfying click. The whole assembly process can easily be done in under 30 seconds. Once assembled, the electrode can be disassembled by prying it up using a pair of tweezers. For some of the looser fitting electrodes, we have also been able to detach the electrode by applying sufficient torque in the disassembly direction.

An electrode that has been partially pried up can be pushed back down as long as it is still engaged by all the fingers. This is an important observation as it will allow the electrode to be rotated into place without interfering with the electrospray needles and then pushed down so the electrode surrounds the needles as in Figure 4-11. This configuration is advantageous as it increases the electric field at the needle tip, thus lowering the startup voltage. It also allows a smaller aperture size to be used for a given beam divergence, which in turn allows a tighter needle spacing and higher thrust densities.

During assembly, one must be careful to fully rotate the electrode into place. In our initial measurements, we found that a few of the devices had large misalignments (50 μm). The optical microscope revealed that the fingertips weren't completely in their notches. Wiggling the electrode with the assembly tool, allowed the fingertips to fall into their notches and reduced the misalignment to acceptable levels. This problem mainly occurred in the first batch of devices, perhaps because their rougher side-walls prevented the fingers from fully falling into their notches.

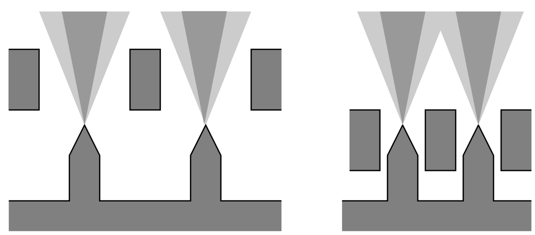

Figure 4-11: A lower electrode permits denser needle packing for a fixed beam divergence

Overall, the robustness of the devices has been quite satisfactory. We have fabricated a total of 36 devices: 18 in each batch. Only three fingers broke during fabrication: two while removing die-saw tape, and one because a speck of photoresist prevented it from properly separating from the scrap silicon surrounding it. We have only assembled seven devices from the first batch, breaking a total of four fingers. Two of those were when assembling type C electrodes (the ones with the largest deflection), the other two broke when we assembled a polyimide electrode that we believe was oversized (see Section 5.1.3). All 18 devices from the second batch were assembled without any breakage. Moreover, no fingers have been lost during disassembly, or other manipulation of the devices. If we only consider the 25 devices which have been assembled, our finger yield is thus 96%. But since the devices with broken fingers did not exhibit larger misalignments than the other devices, we can claim a device yield of 100%.

4.4.2 Assembly Alignment Measurements

We now look at how accurately the electrodes are aligned with the base. In doing so, we will only consider the second batch of devices, as its improved process seems to have dramatically enhanced alignment repeatability, and slightly improved alignment accuracy.

We used an Electronic Vision Group TBM8 front-to-back alignment checker to measure the misalignment between alignment marks on the electrode and corresponding marks on the back of the device. Marks were placed in the center and at four corners of the device so that both linear and angular alignment errors could be measured. (On most devices we only measured the center misalignment and one of the edge misalignments, after checking for a few cases that the five measurable misalignments were consistent with an isometry – i.e., a transformation which preserves distances.) We made misalignment measurements on all 18 devices. Moreover, two of the devices were assembled and disassembled 10 times each to evaluate the consistency of the assembly. Finally, we looked at the misalignment of the electrode in the vertical direction for a few devices using an interferometric microscope, and found that the top of the electrode is aligned with the top of the electrode holder to within a few micrometers.

The TBM8 has two roughly aligned cameras that look at the front and back of the device. By subtracting two measurements between which the device has been rotated 180°, any misalignment of the cameras is canceled out, and we get a precise measure of the front-to-back misalignment on the device. A one pixel error when determining the alignment mark location corresponds to a 0.21 μm measurement error. When taking multiple measurements of the same misalignment, we have always found the same measurement to within 0.42 μm, which is a two pixel error. This error is negligible because, as we shall see, it is lower than the accuracy and repeatability of our alignment method. We did note, however, that when the same measurement was performed on a device that had been rotated by 90°, 180° or 270° was made, we obtained measurements differing by several micrometers. The difference between the measurements was consistent with a constant offset being added to each measurement, perhaps due to some tilt in the TBM8 stage. We subtracted out this constant offset for the measurements we present here.

Table 4.2 shows the measurements that were made on the 18 devices in the second batch. The reported values (illustrated in Figure 4-12) are x and y, the misalignment at the center of the electrode, θ, the angular misalignment, d_0, the misalignment distance at the center of the electrode, and $d_{\mathrm{M}} = d_0 + r\,|\theta|$, the maximum misalignment across the central region of the electrode which has a radius $r = 7.5$ mm. Because the electrode is held in position by fingers at its periphery, we expect that if its size were increased, θ would decrease, while d_{M}, x and y would remain unchanged. The device name indicates the wafer number, electrode type, post geometry and finger type (Table 4.1 in Section 4.2.1). For example, 2CrIII, indicates a device obtained by inserting a type C electrode from wafer stack 2 into a base from wafer stack 2 with type III posts and a ring post geometry.

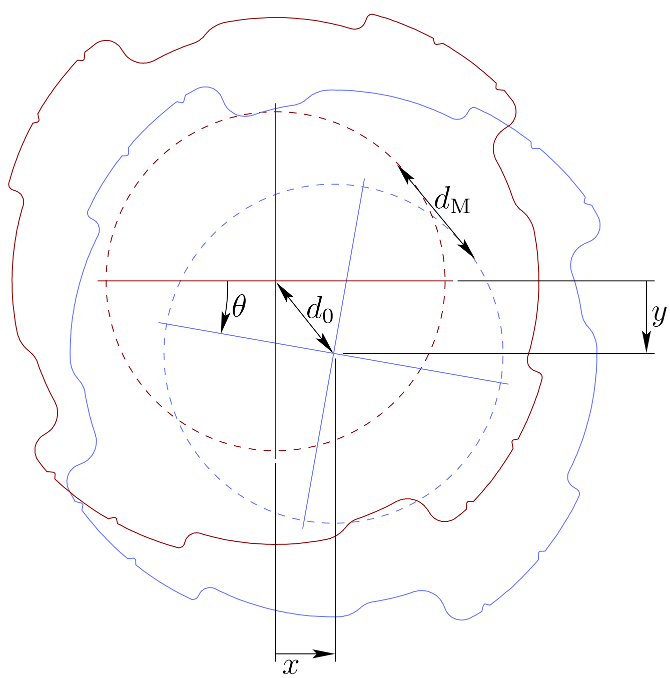

Figure 4-12: Variables used to measure misalignment

Device	x (μm)	y (μm)	θ (')	d_0 (μm)	d_{M} (μm)	Comment
1ApI	-6.3	-1.4	-2.0	6.5	10.9	
1BpII	-5.7	1.4	-0.3	5.9	6.5	
1CpIII	-9.8	-1.4	-2.6	9.9	15.6	
1ArII	4.1	4.0	0.9	5.7	7.7	
1BrIII	7.6	1.2	2.4	7.6	12.9	
1CrI	10.6	3.6	2.8	11.2	17.2	
2ApIII	-7.7	-2.3	-0.1	8.0	8.2	
2BpI	-7.4	1.0	-0.8	7.5	9.2	
2CpII	-7.2	-1.4	-1.4	7.4	10.4	
2ArI	-0.3	-3.2	0.5	3.2	4.3	
2BrII	4.7	-2.3	1.7	5.3	9.0	
2CrIII	11.7	-4.7	2.2	12.6	17.4	one broken finger
3ApII	-7.9	-3.4	1.4	8.6	11.6	
3BpIII	-12.0	-6.4	-0.3	13.6	14.3	
3CpI	-10.7	-3.0	2.1	11.1	15.7	
3ArIII	7.6	-1.9	0.6	7.8	9.1	
3BrI	-4.2	11.4	3.8	12.1	20.4	
3CrII	10.6	0.3	3.0	10.6	17.1	had to wiggle it
Minimum	-12.0	-6.4	-2.6	3.2	4.3	
Maximum	11.7	11.4	3.8	13.6	20.4	
Average	-1.3	-0.5	0.8	8.6	12.1	
RMS	8.4	4.0	2.0	9.3	13.2	
RMS isometry corrected	3.8	5.0	1.7	6.2	9.6	
Std. Dev. 10 reassemblies	1.0	1.2	0.5	0.6	1.3	

Table 4.2: Misalignment measurement results for the x and y offsets at the center of the electrode, the angular offset θ, the offset distance at the center of the electrode d_0, and the worst offset across the electrode d_{M}

In a few cases we also looked at the alignment of the electrode in the out of plane direction. These observations suggest that the tops of the pillars are directly resting on the bottom face of the electrode, as one would expect.

4.4.3 Analysis of Observed Misalignment

Looking at the measurement results, our first observation is that the precision of alignment is sufficient for the intended electrospray electrode alignment application – beam divergence forces the electrode apertures to have a diameter comparable with the wafer thickness (recall Figure 4-11), i.e., hundreds of micrometers in diameter, compared with 20 μm misalignment in the worst measured case. Moreover, in the data from the first batch of devices, broken fingers do not appear to reduce the alignment precision, even for two side-by-side broken fingers.

We believe that much of the observed misalignment is introduced during mask alignment and when aligning wafers for bonding, and is as such not directly related to our assembly method. Such misalignments could be avoided in a process which combines on the same mask the fingers or notches, and the features that require precise alignment. We will therefore try to estimate the error introduced by mask and bonding alignment, and take it out of our measurements to better appreciate the misalignment that is due to our assembly method.

In our devices, these misalignments can be modeled by using two isometries[1] (i.e., 6 degrees of freedom) for each wafer stack. One isometry models the misalignment between the springs and the alignment mark on the base, and the other models the misalignment between the notches and the alignment mark on the electrode. Any misalignment introduced when aligning masks and bonding should be described by these isometries.

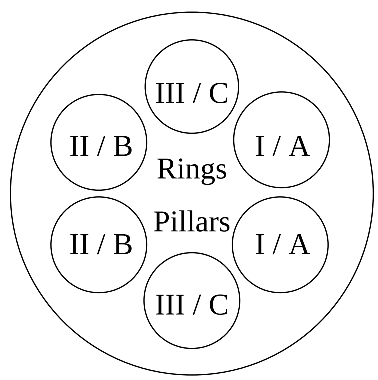

Figure 4-13: Layout of the different device types on a wafer

There are 6 devices per wafer stack, which each provide 3 independent measurements (x, y and θ), for a total of 18 measurements per stack. We have therefore tried to determine the best isometries for each wafer stack. This approach has been par-

[1]An isometry is a transformation which preserves distance, i.e., a rigid-body motion. In this case we are only considering rotations and translations in the plane.

tially successful as it allowed us to reduce the RMS deviations, as shown in Table 4.2. We were pleased to find that the difference between the x and y RMS deviation is substantially reduced with this modeling, as there is no apparent reason why the x axis should have two times more misalignment than the y axis. Looking more closely at table 4.2, there is a visible correlation between electrode type/post geometry and x deviation in the raw data. Pillars correspond to lower values of x than rings, and type C electrodes have larger deviations. Looking at the layout of the devices on the wafer (Figure 4-13), we find that this observation is largely explained by a wafer-level rotation in one of the alignment steps. This is one of the main effects that are eliminated in our isometry modeling.

To see if the remaining error is just due to randomness in the assembly process (due to friction, for example), we assembled and disassembled devices 3ApII and 1CpIII ten times each. The resulting measurements are shown in Figure 4-14, and the standard deviations are in Table 4.2. The assembly/disassembly is repeatable to within a standard deviation of a micrometer. This value is significantly lower that the 5 μm RMS we obtained after isometry modeling. Hence, the 5 μm RMS deviation cannot be explained by randomness in the assembly process.

Up to now, we have considered that when a finger flexes, the fingertip translates in a direction perpendicular to the axis of the finger. In practice this is not exactly true: there is a small rotation of the fingertip, and a small deflection in the axial direction, which could cause a systematic angular offset in the assembled device. Since pillar and ring devices (see Table 4.1 in Section 4.2.1) twist in opposite directions during assembly, we would expect this systematic offset to be in opposite directions for pillar and ring devices. And this offset cannot be explained by isometries. Looking at Table 4.2 some more, we notice that indeed the angles for pillar devices are systematically lower than the angles for ring devices, and this deviation is greater for larger electrode sizes. Averaging over the devices, and after taking into account the isometries we have already determined, we find deflections of 0.5, 1.5 and 1.6 arc minutes for type A, B and C electrodes, respectively. This is the right order of magnitude for the systematic angular deflection we would expect from fingertip motion

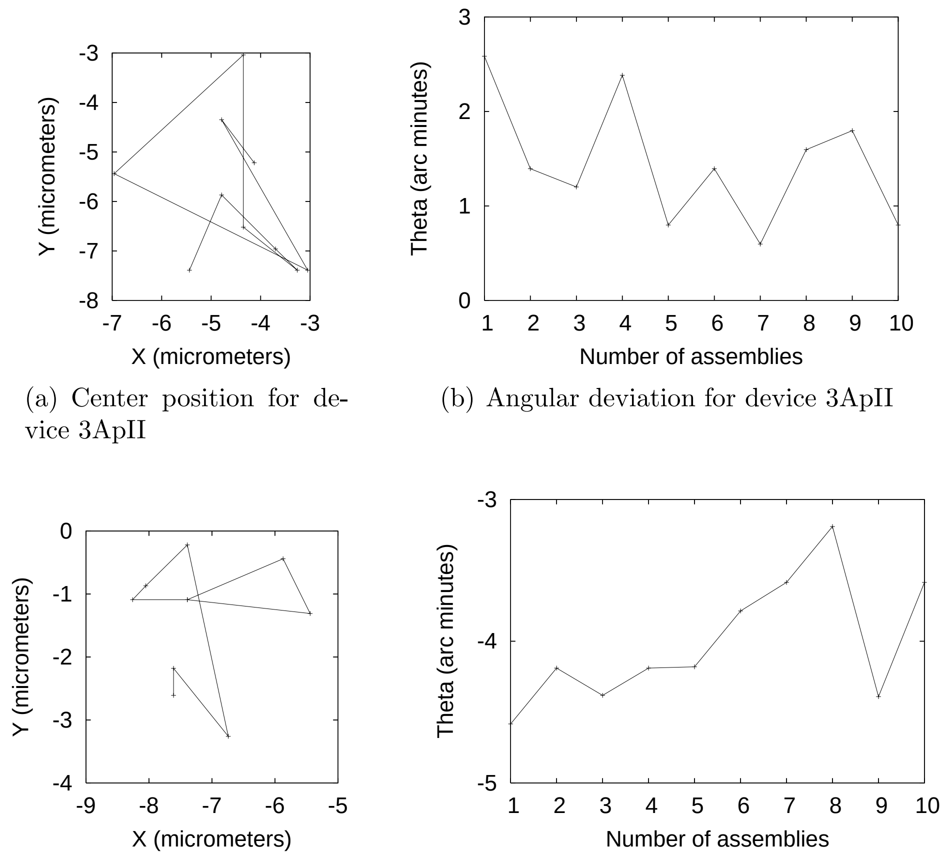

(a) Center position for device 3ApII

(b) Angular deviation for device 3ApII

(c) Center position for device 1CpIII

(d) Angular deviation for device 1CpIII

Figure 4-14: Alignment data for devices 3ApII and 1CpIII, which were assembled and disassembled 10 times each, showing a repeatability of a couple of micrometers

nonidealities.

Another source of error that cannot be explained by isometries occurs if the two sides of a fingertip or notch are not over/under-etched by the same amount. This type of variation would cause a random deviation in our misalignment measurements, but would be consistent when a device is disassembled and reassembled. If this effect contributes to the remaining variation after isometry correction then we hypothesize that making the fingertip and notch smaller should improve alignment accuracy.

4.5 Alternate Assembly Methods

The spring based assembly method we have just discussed is accurate and robust. However, it does not directly solve the problem of insulating the extractor electrode from the emitter electrode. Also, it requires wafer-bonding to assemble the spring wafer to the base wafer, which is always a delicate operation. Therefore, each time a new run of emitters was prepared, we experimented with alternate assembly methods.

The methods we present here use ceramic spacers to achieve both alignment and insulation. In each case, a single wafer yields the emitter and extractor dies from a common two-mask process. Shadow masks were used to separate the emitter die into a wicking black-silicon area and a non-wettable Teflon-coated area. Shadowing of the insulators was implemented to prevent fouling by stray emissions.

4.5.1 Ceramic Ball-Based Assembly

This approach is based on the idea of a kinematic coupling [104], as shown in Figure 4-15(a). Each die is equipped with three holes. If these holes are lined up, the ridge is in front the slot. To line the holes up, we place a 1/32" alumina ball on each hole on the emitter die, we then set the extractor die on top of the balls, so that the balls fit into the holes on the extractor die. Then we use two alumina screws to keep the dies together. As long as the dies are kept together, the balls will keep them lined up, separated, and insulated from each other. The spacing between the dies is determined by the diameter of the balls and the holes they are placed in. As seen in Figures 4-15(d) or 4-15(c), each die actually contains two symmetric pairs of holes. Each pair of holes has a different size. This way, by rotating one of the dies 180°, a different die spacing can be achieved. Depending on which sets of holes are used, four different spacings can be obtained. To avoid over-constraining the system, the holes in the extractor are round, while the holes in the emitter die are rectangular.

During thruster operation, kilovolt potential differences are applied between the dies. The alumina balls and screws ensure die alignment while maintaining electrical insulation. With time, there is a risk of contamination of the alumina surfaces by

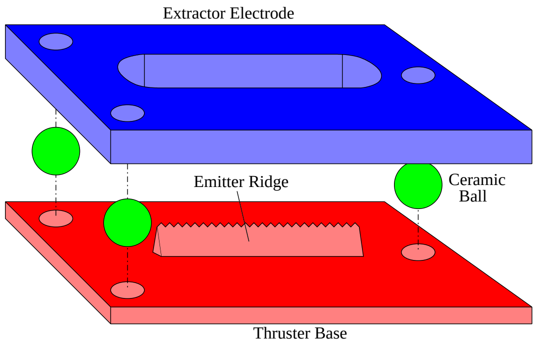

(a) The ball-assembly thruster uses balls to align the extractor slot(s) with the underlying ridge(s)

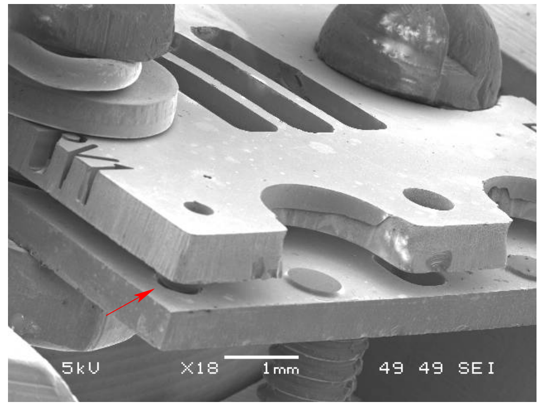

(b) Scanning electron micrograph of an assembled ball-assembly thruster, the red arrow indicates a partly visible ball

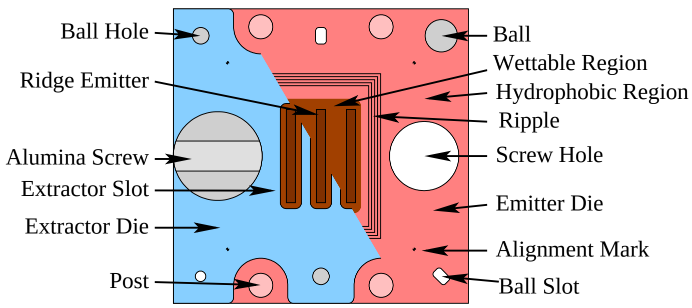

(c) Details of the ball-assembly thruster design

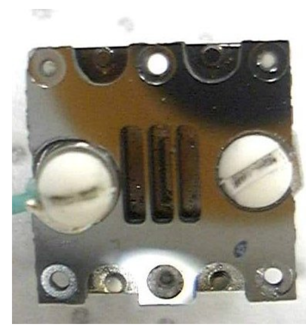

(d) The assembled 1 cm by 1 cm ball-assembly thruster

Figure 4-15: Overview of the ceramic-ball-assembly thruster

material being emitted by the thruster. Figure 4-16 shows how a recess is placed around the hole for the ball on the extractor die, to shadow the ball from material coming from the emitters. A similar recess protects the alumina screws. Moreover, if the extractor die is placed close enough to the emitter die, the emitters end up being completely surrounded by extractor die material. This offers a second level of defense for the insulators.

Figure 4-17 outlines the fabrication of the ball thruster. The basic emitter process in enhanced with a DRIE etch through the wafer from the backside. This etch allows the wafer to be diced before making the emitters, allowing each emitter die to undergo the optimal amount of etching. Indeed, when processing a whole wafer

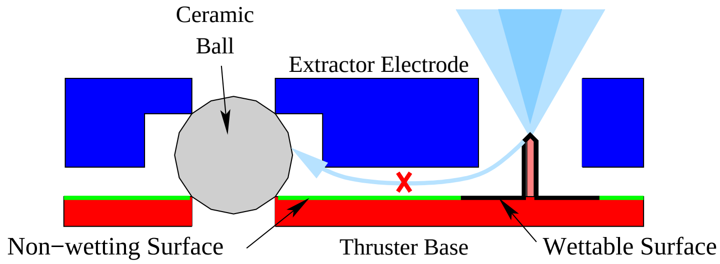

Figure 4-16: Side view showing shielding of the insulators from the emission site

at a time, non-uniformities in the SF_6 etch cause emitters at the periphery of the wafer to be destroyed before the ones at the center have been fully sharpened. The backside through etch is also used to define holes for the ball assembly system, screw holes, and the extractor electrode slots. The process allows us to make both emitter dies and extractor dies from a single wafer. Once the geometry of the devices has been defined using plasma etching, we do some extra plasma treatments to define hydrophobic and hydrophilic regions. The hydrophobic regions are coated in a Teflon-like fluoropolymer. The hydrophilic regions get a black silicon treatment as described in Chapter 3. Finally, finished extractor and emitter dies are assembled using alumina balls for alignment, and alumina screws hold the assembly together.

The reader can refer to [38] for complete details on this thruster concept. The key points are summarized here:

Alignment precision: In the horizontal direction, the alignment accuracy was on the order of 20 μm. The vertical spacing, was highly inaccurate, as it depends on the size of the holes and slots that the balls rest in, and therefore on the exact amount of SF_6 etching that was performed. The SF_6 etching appears to have etched different crystal planes at different rates. Therefore, ball slits with different orientations ended up having different widths. This caused the extractor to be tilted relative to the emitter die, causing 100 μm of difference in spacing from one side of the thruster to the other. Finally, since balls can only be easily found in standard sizes, the only way to adjust emitter to extractor spacing was to vary the amount of DRIE and SF_6 etching done on the grid. In

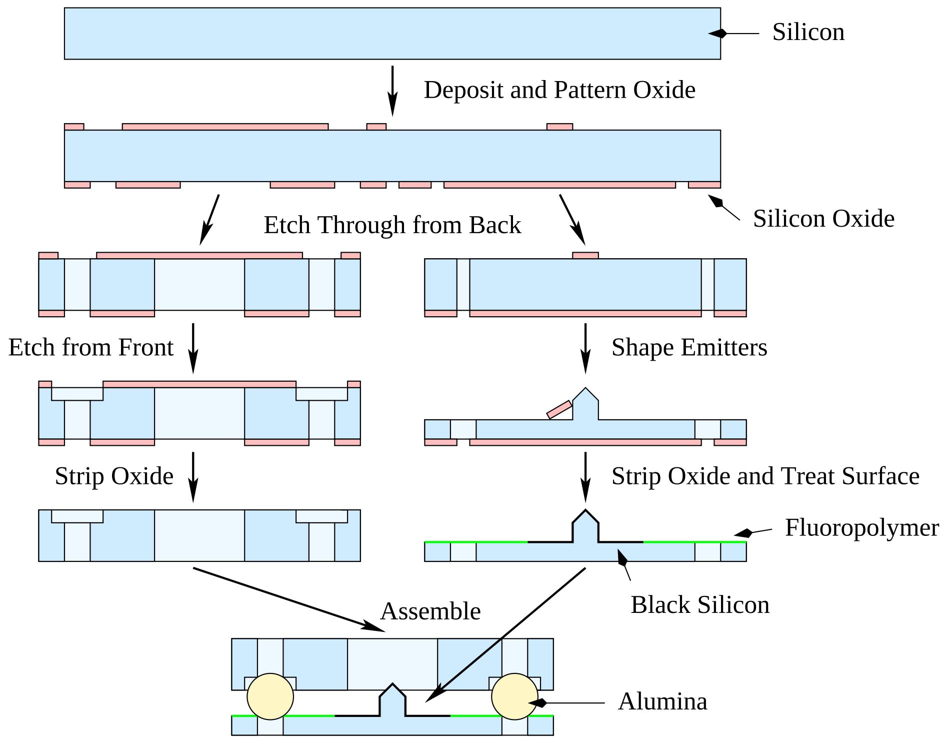

Figure 4-17: Process flow for the ball-assembled thruster

practice, this made it very difficult and time consuming to set the spacing to a desired value.

Ease of assembly: Assembly of these devices was difficult and tedious, despite the use of a jig to hold the dies and screws in place. Assembly would typically take two hours, with a high risk of breaking parts in the process.

Insulation: These devices were successfully fired at over 2 kV, so the insulation scheme was successful. We did get a few short circuits, as it was easy to accidentally contaminate balls with liquid during assembly.

Firing: These devices were able to fire, with about 10% of intercepted current. They never fired very stably. We believe that the poor performance was due to the lack of a good surface treatment on the emitters, at the time.

Process: Doing the DRIE etching on dies turned out to be very unreliable. It was difficult to bond the dies to a handler in a way that provided sufficient heat transfer for the anisotropic etch, and which avoided undesired masking due to upwelling of photoresist through the assembly holes. Therefore yields in this process remained low. Moreover, processing 100 pieces is much more tedious than processing a single wafer, because of the increased amount of handling and the need to mount pieces to a handler wafer before they can be processed in standard fabrication tools. The author strongly recommends against working with pieces whenever possible.

4.5.2 Hybrid Assembly

The ball-based assembly solved the insulation problem but was difficult to work with because of the screws and small balls. This led to a thruster design which combined springs to clamp the thruster together and ceramic spacers to keep the electrodes separated. Figure 4-18 shows the layout of this thruster.

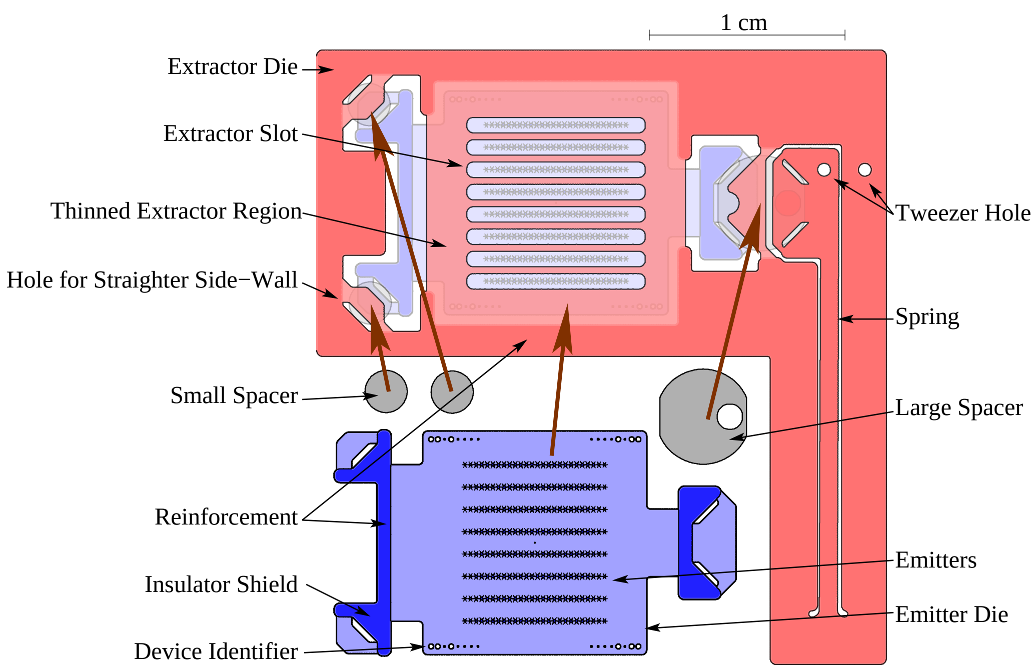

Figure 4-18: Layout of the hybrid thruster

Figure 4-19 shows a side-view of the concept. The extractor and emitter dies are separated by alumina spacers. The silicon parts are shaped in such a way that the conductive path between the two electrodes, along the ceramic surface, is as long as possible. The spacing between the extractor and emitter dies is directly set by the thickness of the spacers. The relative position of the dies is dictated by contact between the spacer side-walls and vertical features on the silicon dies. The insulating surfaces are systematically shadowed from the emitters by side-wall features.

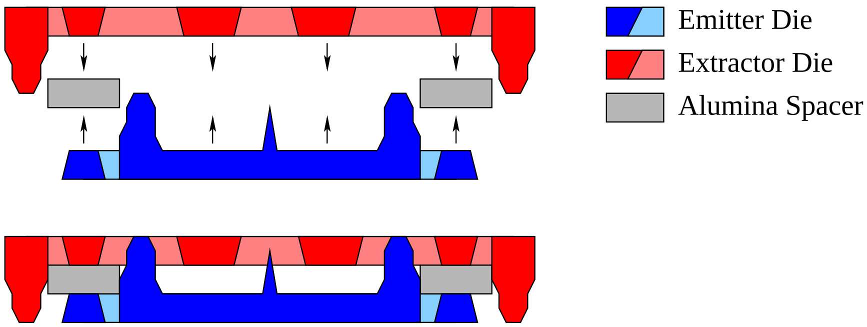

Figure 4-19: Side view of the hybrid approach, unassembled (top) and assembled (bottom)

As shown in Figure 4-20, three spacers are used to position the emitter die relative to the extractor die. First, a pivot is formed using two 2 mm spacers, allowing the emitter die to pivot around its center with the spacers serving as rollers. At the time of this design, we thought that friction might be the limiting factor for the finger-based alignment scheme, and were hoping that the rolling contact would eliminate misalignments due to friction.[2] Kinematically, there are four components, and so $3\times3 = 9$ degrees of freedom (assuming that one component is fixed in space to account for a rigid motion of the whole assembly). Each spacer has two point contacts, which eliminate four degrees of freedom, leaving five. Two of these degrees are contained in the arbitrary orientation of the spacers. Two more are contained in translation of the spacers along their constraining side-walls. The single remaining degree of freedom corresponds to pivoting around the center of the emitter die. The position of the axis of the pivot is set by the position of the spacers and the orientation of the constraining side-walls. By placing the axis in the center of the die, we minimize the amount of misalignment which can be caused by the remaining degree of freedom.

[2] The excellent repeatability of the second batch of finger-based devices disproved this hypothesis.

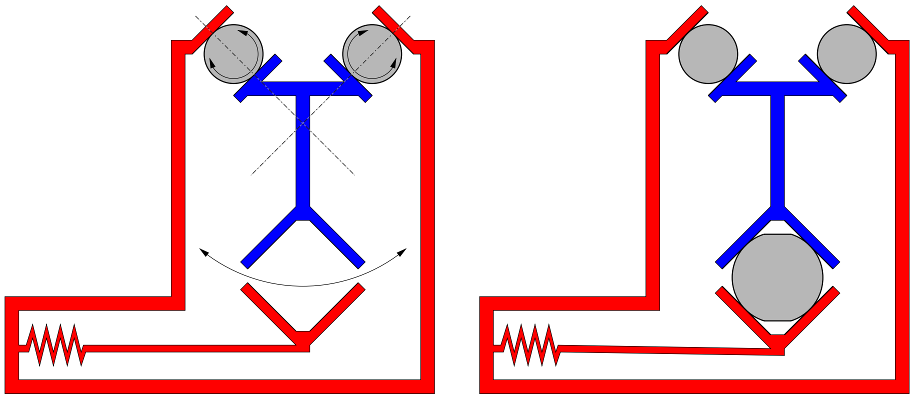

Figure 4-20: Kinematics of the hybrid scheme. When only the two small spacers are present (left), the emitter die can pivot around its center. Adding the third spacer (right) removes this degree of freedom

To constrain the last degree of freedom, we add a third spacer. A spring is used to clamp this spacer in place. The spring adds one degree of freedom, and the extra spacer adds three. Four degrees of freedom are removed by the point contacts with the third spacer. The last remaining degree of freedom allows arbitrary rotation of the spacer around its axis, which leaves the relative position of the emitter and extractor dies unchanged. Thus, the relative position of the emitters and the extractor has been exactly constrained in this scheme.

There is some fabrication uncertainty involved in the exact size of the spacers and position of the constraining side-walls. If the spacers are consistently over or under-cut, or if the side-walls are consistently etched too far in or out, then the position of the emitter die will shift parallel to the direction of the ridges. Thus, when fabrication uncertainty is taken into account, the position of the row of emitters along the slot is not precisely set. Likewise, with the ball assembly concept, the vertical position would depend on process variation. It is unimportant for the emitters to be exactly centered along the direction of the slots, however vertical misalignment can make big changes to the electric field. Thus, we argue that the hybrid scheme has better tolerance for process variation than the ball-based scheme. The original finger-based scheme does not have any processing dependent misalignment, instead it requires more area-intensive springs.

These thrusters are fabricated by using a nested mask: an oxide mask protects regions that shouldn't be etched at all, and a photoresist mask protects regions that need to be etched in 400 μm. Etching begins with a 300 μm etch of the regions where the wafer is to be etched through. Then the photoresist mask is stripped with an oxygen plasma and the emitter formation process forms the emitters and the spacer constraining side-walls, while completing the through etching which had previously started. Then all that remains to do is treat the surface so that it is wettable in the emitter region and hydrophobic elsewhere. The spacers are made using laser-cutting. Because the emitter forming process does not make a sharp corner at the base of the side-walls, a through wafer hole is placed in front of each constraining side-wall. These holes are visible in Figure 4-18, labeled as "Hole for Straighter Side-Wall", and the resulting sidewall profiles have been sketched in Figure 4-19.

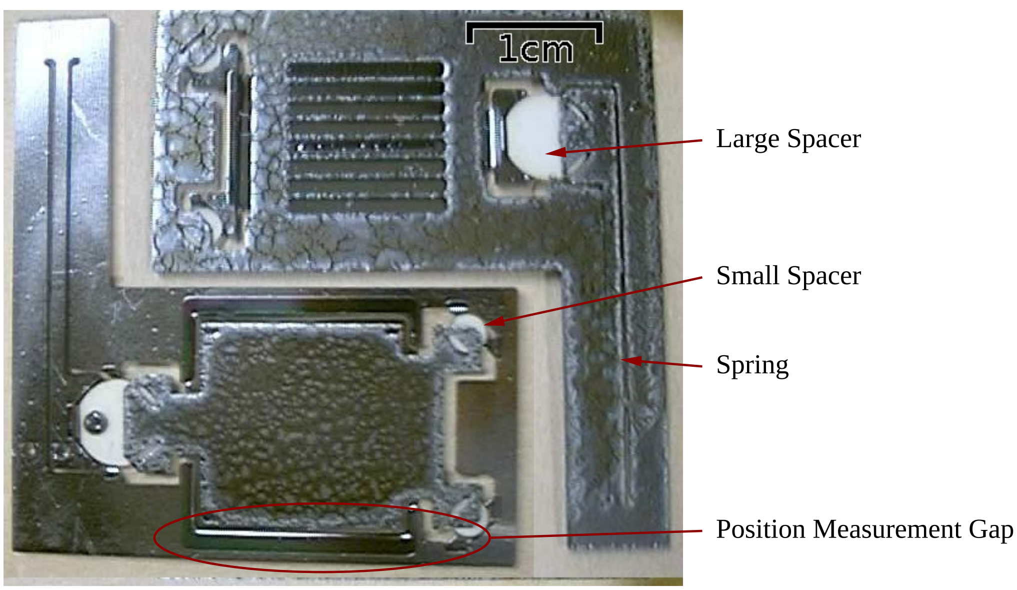

Figure 4-21: An assembled hybrid thruster, emitter side (bottom-left) and extractor side (top-right). (Photoresist which has not yet been stripped gives texture to the silicon.)

Figure 4-21 shows an assembled thruster. The accuracy of assembly can be measured by comparing gap widths on either side of the thruster. The gaps matched to within 50 μm. However, measurements taken along the length of the gap are inconsistent with each other to within about 50 μm, suggesting that the feature edges are not straight. Thus 50 μm in a worst case alignment accuracy.

Assembly takes a few minutes, and is intermediate in difficulty between the two previous approaches. The thruster is placed face-down, and the two small spacers are put into place. The large spacer is placed near its final position, such that it is resting on top of the elevated spring area. The emitter die is put into place. Finally, tweezers are used to deflect the spring and push the large spacer down.

The assembly is not very stable, as the emitter die tends to pop out. This problem can be avoided with a package that continuously presses the emitter die in. The ceramic spacers were delicate to work with, and took many hours each to make, despite their small size. These thrusters were somewhat fragile because much of their area is only 200 μm thick. No further testing was carried out on these devices. Their emitters inherently have vertical side-walls, and we never managed to treat emitters for wettability with vertical side-walls.

4.6 Conclusion

In this chapter, we have investigated methods to assemble two centimeter scale components, such as the emitter and extractor electrodes of an electrospray thruster. Each one of these methods allows the components to be assembled with high precision by hand. By using mechanical clamping means, we are able to decouple the fabrication processes for the device components, and allow greater freedom in material selection.

Table 4.3 compares the three approaches. Of the three methods, we have selected the finger-based approach as the most successful. It is slightly more complex in the processing phase, but its robustness and ease of use overweighs this extra fabrication difficulty in the context of laboratory experimentation, where the thruster will be assembled and disassembled repeatedly. For a production device, that will be assembled only once, and that tries to optimize active area, this tradeoff may shift in favor of one of the other methods. An insulation method for the finger-based method will be presented in Section 5.1.3, which will allow it to be used to make a thruster.

In the finger-based assembly method, the electrode is assembled by hand by rotat-

Method	Advantages	Drawbacks
Finger-based	- 10 μm accuracy, 1.3 μm repeatability. - Fast, easy, reliable assembly. - Insensitive to uniform process variations. - Significant area devoted to springs.	- Difficult to insulate. - Extra spring wafer needed. - Requires wafer bonding. - Insulation is separate.
Ball-based	- Simple process. - Provides insulation. - Maximizes area available for emitters.	- Difficult to assemble. - Poor vertical accuracy.
Hybrid	- Good vertical accuracy. - Simple, single- side, process. - Provides insulation.	- Poor vertical clamping. - Insensitive to uniform process variation along one axis only. - Requires high precision alumina components.

Table 4.3: Comparison of the three assembly methods

ing it until cantilever springs called “fingers” engage in small notches. Hand assembly is possible because the electrode initially only needs to be placed with millimeter accuracy. The assembly motion then forces the electrode into position. In our experiments, once the electrode is assembled, it is positioned within 13 μm RMS of its intended location, which is sufficient for our electrospray thruster application. Analysis of the misalignment shows that 10 μm RMS of error can be attributed to our assembly method. The repeatability of assembly has a standard deviation of 1.3 μm. Our method has proven its robustness: 4% of fingers broke during assembly or manufacturing, and missing fingers did not adversely affect the positional accuracy. We also found that we could disassemble the electrodes by prying them up without causing any damage. Tests with laser-cut polyimide electrodes show that we can easily mix and match processing techniques and materials as long as the mechanical interface of the electrode remains the same.

Chapter 5

The Thruster

We are now ready to build a complete electrospray thruster. Figure 5-1 shows the progression of ideas leading to the complete thruster.[1] We begin (a) with the emitters presented in Chapter 3, and a slotted extractor wafer. To allow tighter spacing, the extractor wafer is thinned (b). The finger-based assembly from Chapter 4 is used to assemble the extractor to the emitters (b). A Pyrex wafer is used to insulate the extractor from the emitters (c). The Pyrex is partially undercut to improve insulation robustness (d). Features are added which allow the emitters to be moved closer to the extractor and even partially recessed (d). To improve control over the vertical positioning of the extractor die, spacers are introduced (e). Liquid traps are added to capture any drops that get close to the insulator region (e).

In Section 5.1, we will now look in detail at each one of these features. The thruster will be made from four wafers: a silicon wafer to make the emitter dies, a Pyrex wafer for insulation, a silicon wafer for the assembly springs, and a middle wafer between the springs and the Pyrex. The details of the fabrication process will be given in Section 5.2. Finally, Section 5.3 will discuss some improvements which could be made to the thruster. Throughout the chapter, the reader may refer back to Figures 5-1, 5-2 and 5-3 to better visualize the parts that are being described.

[1]These steps do not correspond to the actually fabrication steps, which will be reviewed in Section 5.2.

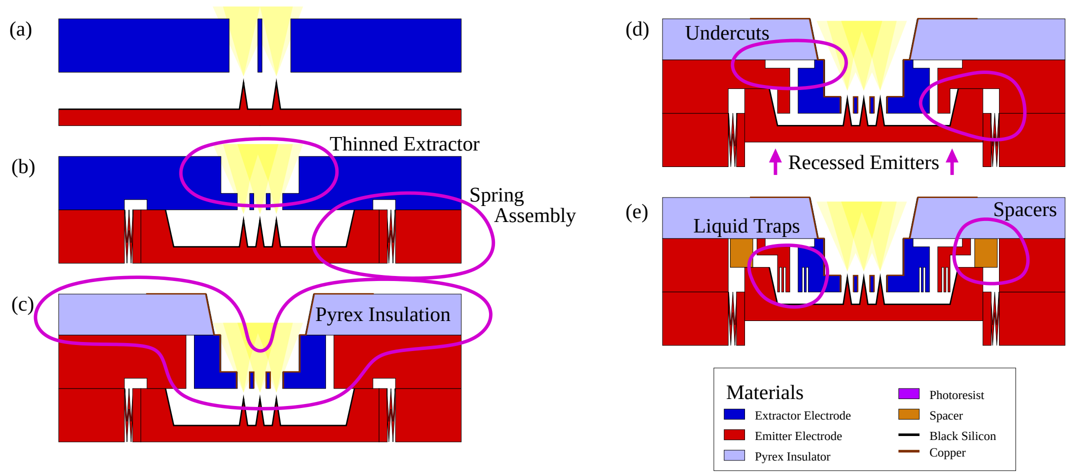

Figure 5-1: Progression of ideas from two electrodes to an fully integrated thruster. (This is not the process flow)

5.1 Description of the Thruster

5.1.1 Extractor Thinning

Beam divergence is one of the limiting factor on emitter packing, a thinner electrode can allow the emitters to be packed closer together without incurring any beam divergence [36]. We opted for a thickness between 200 and 250 μm for the emitter region. The thinned region is wider than the area covered with emitters to be consistent with the expected 30° emission half-angle. See Section 2.2.4 for details on selecting an extractor thickness.

5.1.2 Assembly

The assembly method we developed in Chapter 4 was used to assemble the extractor to the emitters. Originally, we had planned on assembling the extractor component into a spring-equipped emitter component [105]. In the end, the opposite configuration, where simple emitter dies twist into a spring-equipped extractor component, was preferred for the following reasons:

- Emitter formation and the black silicon treatment are very delicate operations. It is preferable to move as much complexity as possible away from the emitter

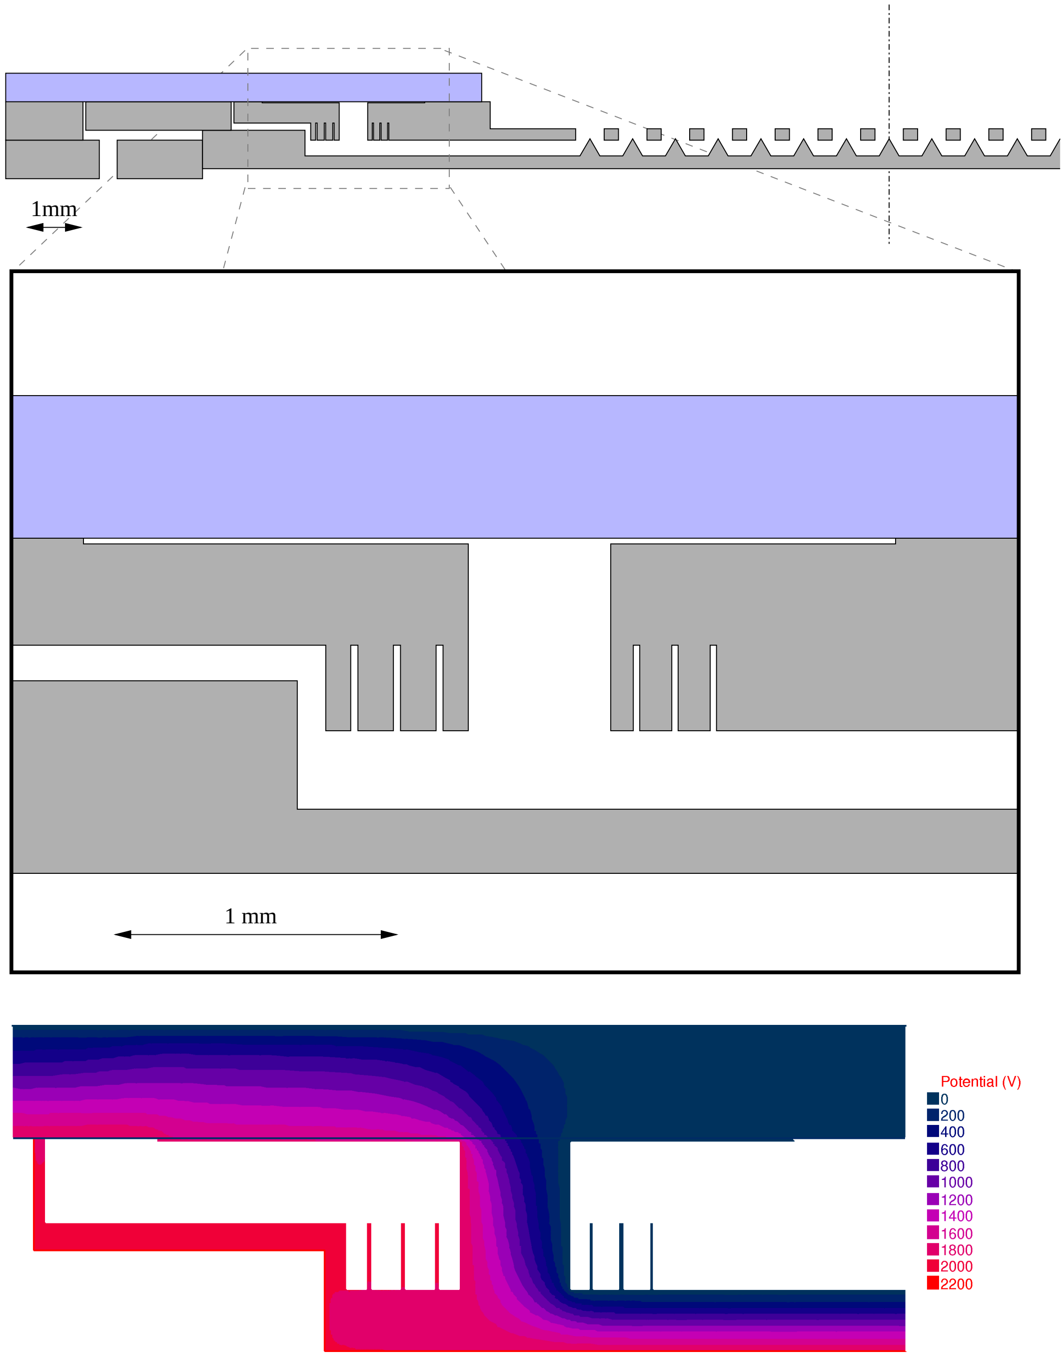

Figure 5-2: Cross-section of the thruster in which all the scales have been preserved (top). Closeup on the insulation region (middle) and the electrical potential in that region (bottom)

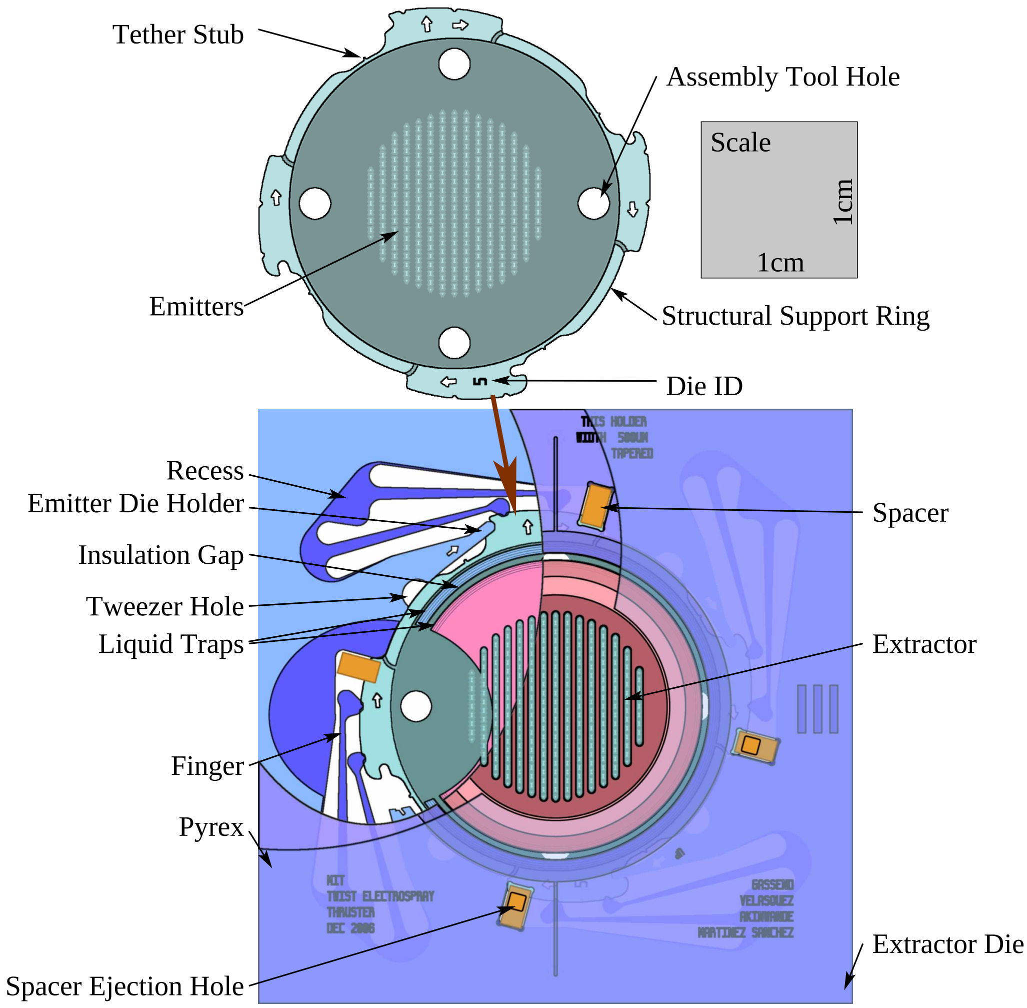

Figure 5-3: The thruster seen from the extractor side

component to reduce the loss when one of these processes fails. Also, it is better to have as few constraints as possible on the process for the emitter component. For example, to avoid sodium contamination of CMOS processes, Pyrex wafers are not allowed in the black silicon treatment machine. If the Pyrex insulation (See Section 5.1.3) had been part of the emitter component, the complexity of the process would have been significantly increased.

- Emitters are generally replaced more frequently than extractors. Indeed, the emitters can be damaged during operation due to electrochemical effects and need replacement. Also, there is a multitude of emitter geometries to explore.

Thus, the emitter dies should be as small and simple to make as possible. This allows a multitude of disposable emitter dies to be fabricated and tested with a small number of extractor components which contain most of the complexity of the thruster.

5.1.3 Insulation of Finger-Based Assembly

So far the device we have described has no electrical insulation between the base and the electrode. In our application, voltages around 2 kV are needed to run the electrospray. Achieving these voltages reliably has been one of the most challenging tasks encountered in this thesis. Figure 5-4 shows different insulation strategies we have tried.

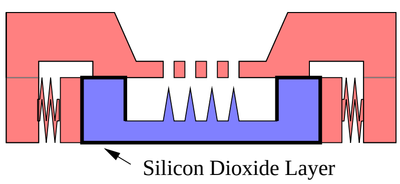

(a) Insulating electrode surface

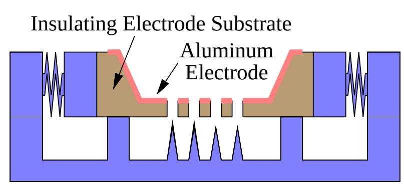

(b) Insulating electrode

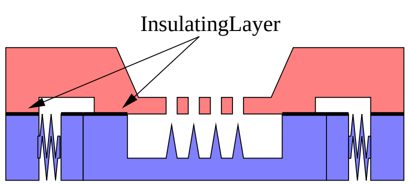

(c) Insulating film at bond plane

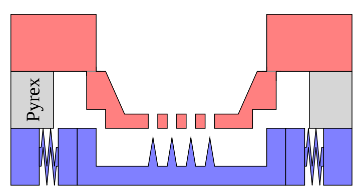

(d) Insulating intermediate wafer

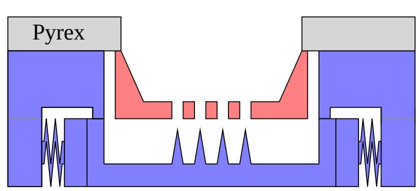

(e) Insulating substrate

Figure 5-4: Different methods of electrically insulating the electrode from the base. Disconnected conductive regions are indicated by different tone and hatches

Insulating Electrode Film

As shown in Figure 5-4(a), we covered the silicon electrodes with a 10 μm layer of PECVD (plasma enhanced chemical vapor deposition) silicon oxide, to insulate the electrode from the rest of the device. Far from the edges, the oxide was able to withstand 5 kV. However, near the edges, the oxide broke down around 700 V. Field enhancement due to the sharp corners or reduced oxide thickness on the sidewalls may be responsible for this poor performance. We have also tried this insulation strategy using an 8 μm parylene film. These devices broke down between 1.5 kV and 2 kV, with leakages before breakdown in the hundreds of nanoamps range.

Insulating Electrode

As shown in Figure 5-4(b), we made the electrode out of 600-μm-thick laser-cut polyimide (Cirlex) film. A thin aluminum film was sputtered onto the polyimide and patterned using the laser. Figure 5-5 shows the assembled polyimide electrodes. These devices were very interesting, as they demonstrate the ease with which the assembly method adapts to a different electrode material. However, the electrical tests were inconsistent: one device was able to withstand voltages up to 5 kV with less than 200 nA leakage current, while others broke down as low as 1.5 kV or had microamps of leakage at 1 kV. Finally, when breakdown occurred, the aluminum film was often damaged and divided into non-connected regions.

Thin Film at the Bond Plane

As shown in Figure 5-4(c), we could bond the top and bottom wafers with an insulator present. Silicon oxide is a likely candidate. We could also use a parylene film and use the bonding method described in [106]. We conducted some tests with parylene in which two parylene samples were bonded at about 170°C in air. These samples had unacceptably high leakage (tens of microamps at 1000 V), but the parylene showed significant discoloration, probably because the bonding was performed in the presence of oxygen. More experimentation is needed to fully evaluate this method.

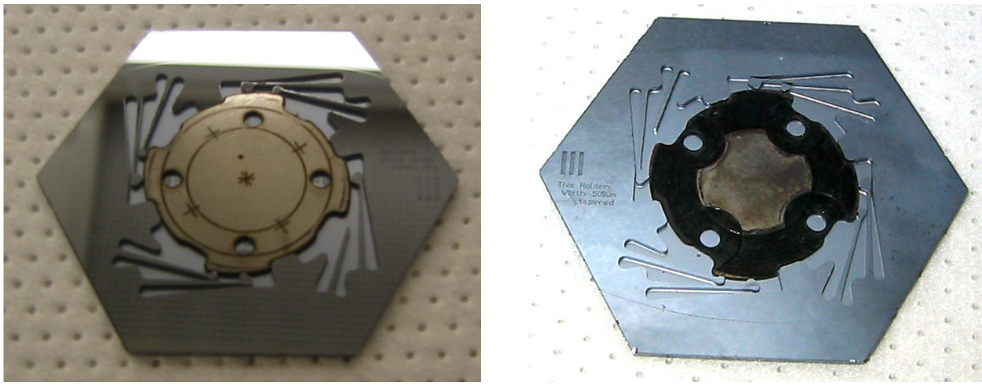

(a) Insufficient gaps in the aluminum film (b) Large gaps in the aluminum film

Figure 5-5: We tried using an insulating polyimide electrode with a thin film of aluminum on it. This experiment demonstrates the flexibility of our assembly method

Insulating Spacer Wafer

As shown in Figure 5-4(d), we tried to use a Pyrex wafer between the top and bottom silicon wafers for insulation. The holes in the Pyrex were cut using an excimer laser at 248 nm. This insulation strategy worked well when we simply anodically bonded unpatterned Pyrex and silicon wafers together. However, with the laser-cut Pyrex wafer, this scheme failed to insulate the top and bottom wafers from each other (in fact, we even observed some shorting during the anodic bond). We conjecture that the laser-cutting has modified the electrical properties of the Pyrex sidewalls. It is likely that other cutting techniques such as ultrasonic grinding, sand blasting [46], or wet etching [55] would preserve these properties. However, the use of a spacer wafer has the added disadvantage that a many hundred micrometer gap is added between the emitters. The insulating substrate method we have developed does not have this drawback.

Insulating Substrate

Figure 5-4(e) shows the strategy that we intend to adopt for our electrospray thruster. In this strategy, the bottom wafer is bonded to a Pyrex wafer, and a trench is cut through the bottom wafer so the central region is insulated from the outer region.

The top wafer is then bonded using thermocompression bonding or fusion bonding (without a high temperature anneal). By varying the trench width, we can increase the surface path across which breakdown must occur. Indeed, with flashover phenomena, breakdown occurs along the insulator surface at much lower voltages than would be expected from the bulk electrical properties.

To evaluate this method, we anodically bonded a Pyrex wafer to a silicon wafer, and used DRIE to etch circular trenches through the wafer, as shown in Figure 5-6. We varied the trench thicknesses between 1 mm and 4 mm to observe the effect of gap length. We then applied voltages across the trenches to determine the electrical characteristics of the insulation.

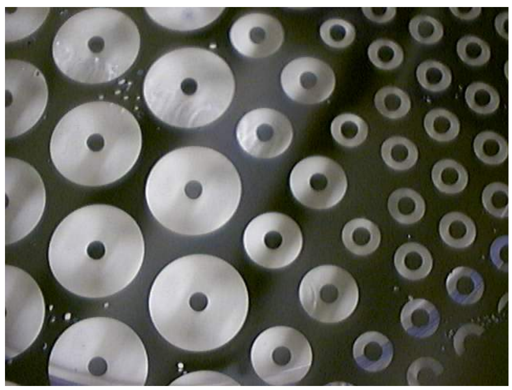

Figure 5-6: The test device that was used to evaluate the electrical properties of the Insulating Substrate insulation method. Trenches 1 mm, 2 mm and 4 mm thick were cut around 2 mm diameter contacts

When these devices, were tested, we observed two different behaviors, which are shown in Figure 5-7(a). Initially, the devices exhibited very low leakage (below 20 pA) with a roughly linear dependence on the applied voltage between 0 and 5 kV. Then, after a sufficiently long time, or at a sufficiently high voltage, the devices irreversibly transitioned to a higher leakage regime (we will call this breakdown). In the high leakage regime, the current varies exponentially with applied voltage, increasing by a factor of ten for every 1000 V. At 2 kV, the leakage is on the order of 1 nA, which is satisfactory for our application. Comparing devices with different gap lengths, larger gaps were able to resist breakdown longer than smaller gaps. For a given gap size, there was a lot of variation on when the breakdown took place, in particular, the amount of charge leaked through the device before breakdown varied by orders of magnitude between devices. Leaving some devices at 5 kV for more than 24 h, we found that the leakage decreases a bit from its value just after the transition (see Figure 5-7(b)).

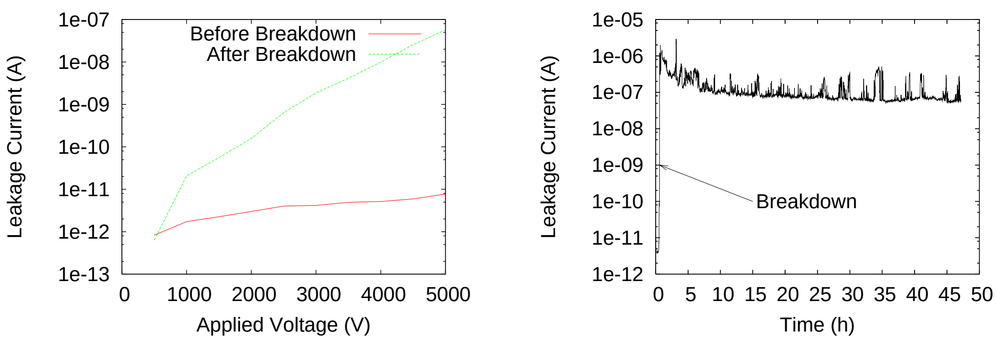

(a) Leakage before and after breakdown (b) Constant voltage test at 5 kV

Figure 5-7: Measurements on a 4 mm-gap substrate-insulation test device

So far we do not fully understand what is happening when these devices break down. However, since their post-breakdown characteristics are suitable for our needs, and since they do not appear to suffer further damage, even when tested at more than twice their operating voltage, we will select this insulation method for the thruster. Section 6.2.8 presents some insulation data on the finished thruster that sheds some light on the leakage we are seeing.

5.1.4 Undercutting the Pyrex

Due to space constraints, we left only a 500 μm gap between the extractor and the surrounding middle wafer in the final thruster. A simple modification can greatly increase the length along the Pyrex to help prevent flashover. Indeed, the top surface of the middle wafer is largely unused. Therefore, the Pyrex can be undercut as in Figure 5-8, leaving a much greater length of Pyrex between opposite polarity electrodes. A millimeter or more can be added on either side without disturbing the features on the bottom of the middle wafer (liquid traps, recesses, etc.).

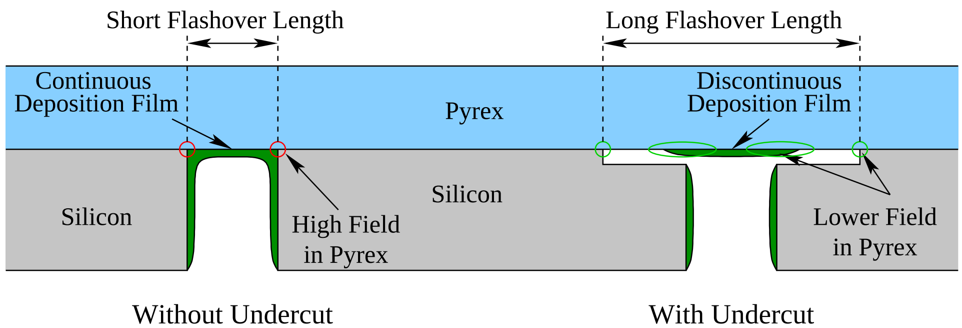

Figure 5-8: Undercutting the insulation trench

This undercutting of the Pyrex also has the advantage of shadowing most of the Pyrex surface, ensuring that it will stay clean even during thruster operation.

Originally, we intended for the undercut to reduce the electric fields in the Pyrex near the sharp silicon corner that it is bonded to. This field could cause charge injection into the Pyrex, and its eventual degradation. Indeed, with the undercut, the sharp silicon corner which contacts the Pyrex is now far from the gap between the extractor and the rest of the middle wafer where high fields are present. However, a conductive layer has been placed above the Pyrex to avoid charging during thruster operation. Thus, there is also a strong field across the Pyrex wafer and large field concentrations even under the outer undercut corner. This high field could be eliminated completely by replacing the conductive layer above the Pyrex by a macro-scale, external, conical ground electrode. The fields in the current configuration and the proposed corrected configuration are shown in Figure 5-9.

One final advantage of the undercuts is that most of the Pyrex surface is not exposed to plasma once the silicon has been etched away. This way, even if plasma exposure degrades the flashover properties of the Pyrex most of the Pyrex surface will be unaffected.

5.1.5 Elevating the Emitters

The spring-based assembly method controls the extractor position in two dimensions. The positioning in the vertical dimension is set by a hard contact between the smooth

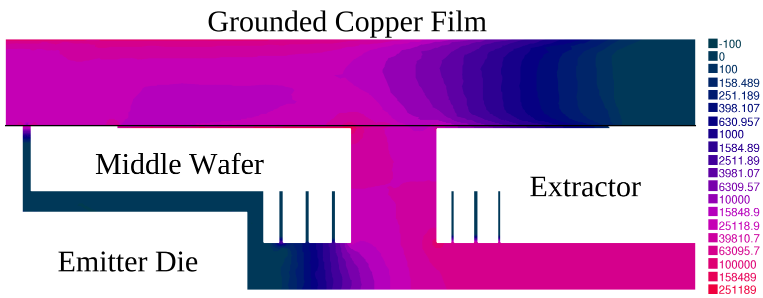

(a) Actual configuration

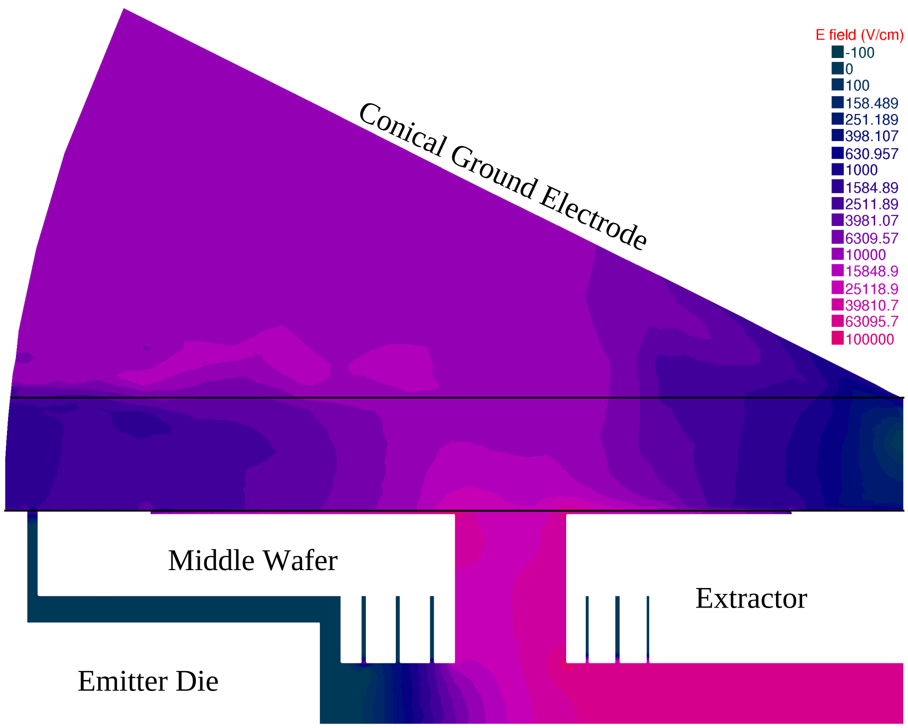

(b) Modified configuration

Figure 5-9: The electric field near the insulator gap

silicon surfaces on top of the emitter die and on the bottom of the middle wafer. The bottom of the extractor is at the same level as that surface, while the tips of the emitter are about 200 μm below that level.

If the tips could be raised relative to the extractor, it should be possible to increase the electric field and decrease the startup voltage. Also, if the emitters can actually be elevated above the base of the extractor then an extra level of shadowing is introduced to protect the insulator from deposition.

This extra elevation can be achieved by providing a recess in the middle wafer so that the emitted die's vertical position is referenced to the recess instead of being referenced to the smooth surface of the middle die. Varying the depth of the recess varies the height of the emitters. However, if the emitters are too high, they will contact the extractor during assembly causing damage to the tips or to the extractor.

The solution to this problem is to assemble the emitters in the horizontal plane at a safe height, and then to change their height. This can easily be achieved by shaping the recess so that the emitter die cannot fit into it until it is within a few tens of micrometers of its final position (Figure 5-10). In this way, the emitters will only be raised once they are safely aligned with the extractor. The shape of the recess has been designed so that the emitter can only be pushed up for one of the four orientations of the extractor (Figure 5-11). This will prevent damage if the emitters are assembled at right angles to the extractor slot, and will allow the recessed assembly to be bypassed if the emitters are assembled at 180° from the correct orientation.

5.1.6 Using Spacers for Vertical Positioning

We now have a method to method to push the emitters up so that they are closer to the extractor. The method has some limitations, however. Firstly, changing the vertical position of the emitters requires a new extractor component to be made. Secondly, the parallelism between the emitter die and the extractor is set by a surface defined by the not-so-uniform DRIE process. Thirdly, this DRIE etch participates in the fabrication of the extractor slots and insulation gap so it cannot be less than about 250 μm.

We can eliminate all these limitations by introducing spacer components to set the height of the emitters relative to the extractor. The spacers are placed in holes that extend through the middle wafer to the Pyrex. When the emitter die is assembled, it contacts the spacers before contacting the bottom of the recess, so the thickness of the spacer sets the final height of the emitters. This height can be adjusted just by changing the set of spacers to a set with a different thickness. Because the spacers are resting on a pristine Pyrex surface, instead of the bottom of a DRIE etched recess, all

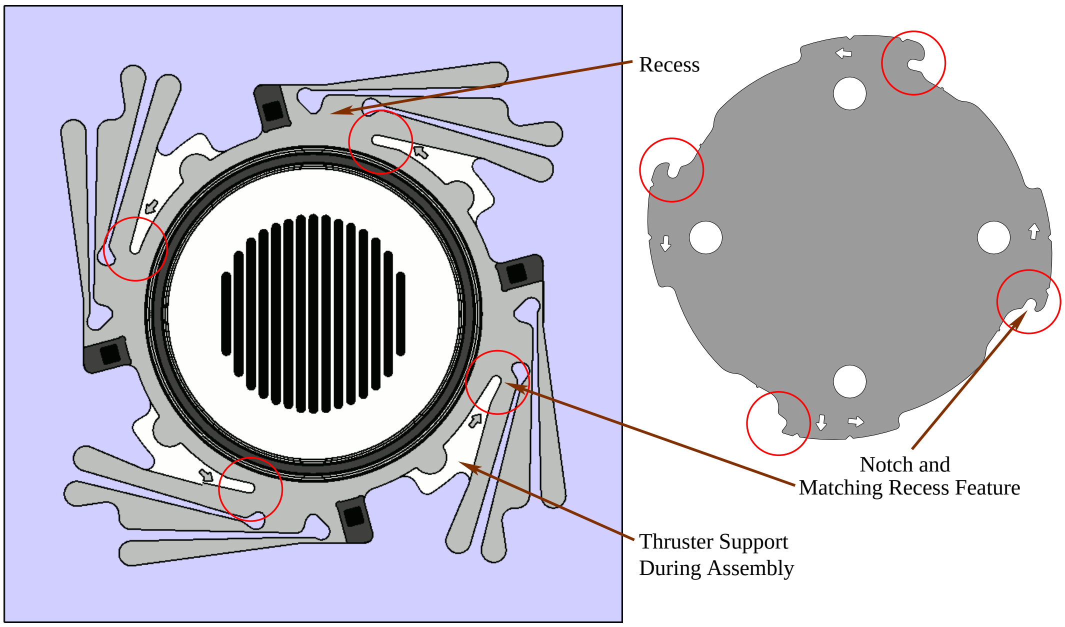

(a) Bottom view of the thruster showing the recess and the features that allow the emitter die to be pushed in only once it is properly aligned

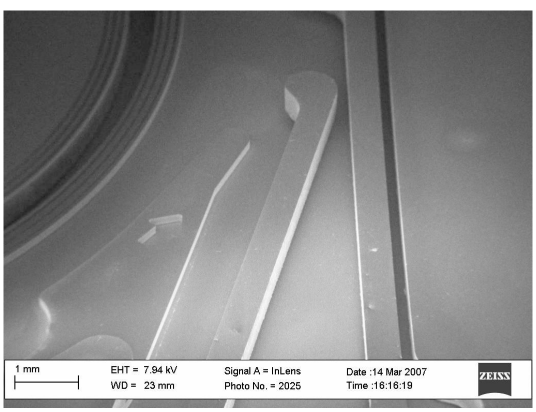

(b) SEM showing a spring and one of the unrecessed areas which holds the emitter die until it is correctly aligned

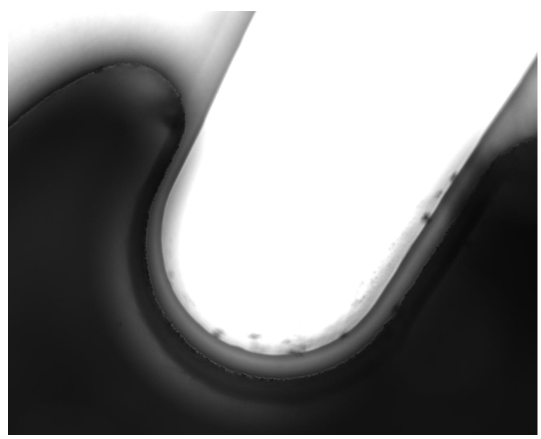

(c) The notch in the emitter die (dark) is aligned with the edge of the recess (light) so the emitters die can be pushed in

Figure 5-10: A carefully shaped recess allows the emitters to be raised once they are properly aligned

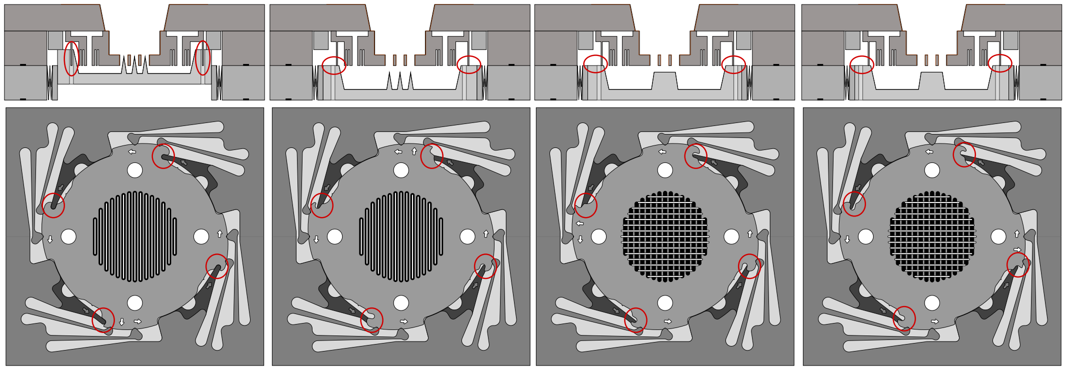

Figure 5-11: The recess is shaped in such a way that the emitters can only be elevated with one orientation of assembly

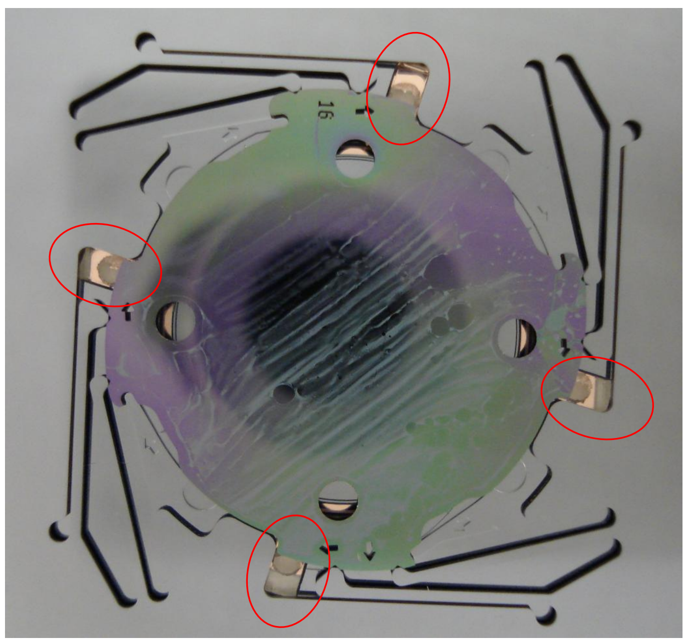

Figure 5-12: Four spacers set the height of the emitters relative to the extractor

the spacers should be dictating precisely the same height. The spacers are rectangular so that hundreds of them can be made in a few minutes using a die saw. They are roughly 1.8×2.6 mm, which is easy to handle with tweezers. The spacers we used were Pyrex because some scrap was available with the desired thickness. Any other material should work just as well.

To check the accuracy of the vertical positioning, we took some cursory measurements. The emitter die was assembled using spacers, and the difference in height between the bottom of the emitter die and the bottom of the spring wafer was measured using an interferometric microscope, in four corners of the die. We repeated this set of measurements twice with the same device after an assembly/disassembly cycle. All the measurements were identical to within $\pm 5\ \mu$m, which is much less than the variation in height between emitters.

5.1.7 Liquid Traps

This engine has excellent properties for avoiding liquid spillage. Only the emitter die has been black silicon treated, so liquid that gets onto the extractor die forms drops rather than spreading across the silicon surface towards the insulating regions. Moreover, the contact areas between the emitter and the extractor die are in a low field region which precludes droplet motion under the influence of electrical forces.

As an added measure of security, we have placed liquid traps around the insulator regions (Figure 5-13). To get from the emitter die or the extractor (if it gets electrosprayed with droplets) to the insulators, liquid has to cross three trenches which are at least 200 μm deep. The electric field is shielded in these trenches and is therefore very low, so there is nothing driving the droplets to cross them. In addition, the trench sidewall away from the insulator has been equipped with triangular groves so that capillarity encourages the liquid fill the grooves and stay on that side of the trench. At their widest point the trenches are 25 μm wide; the triangular groves have an angle of 45° and a spacing of about 15 μm. A spacing of 150 μm was left between trenches to guarantee that they would not merge even if there was a significant taper in the DRIE etching. It should be possible to reduce this spacing and increase the active area of the thruster.

There was not enough liquid migration during thruster operation to evaluate the usefulness of these trenches. In fact, the electric field at the edge of the emitter die was low enough that liquid was not ever drawn up the sidewalls of the emitter die (see Figure 5-14). The liquid traps may become useful for long duration testing.

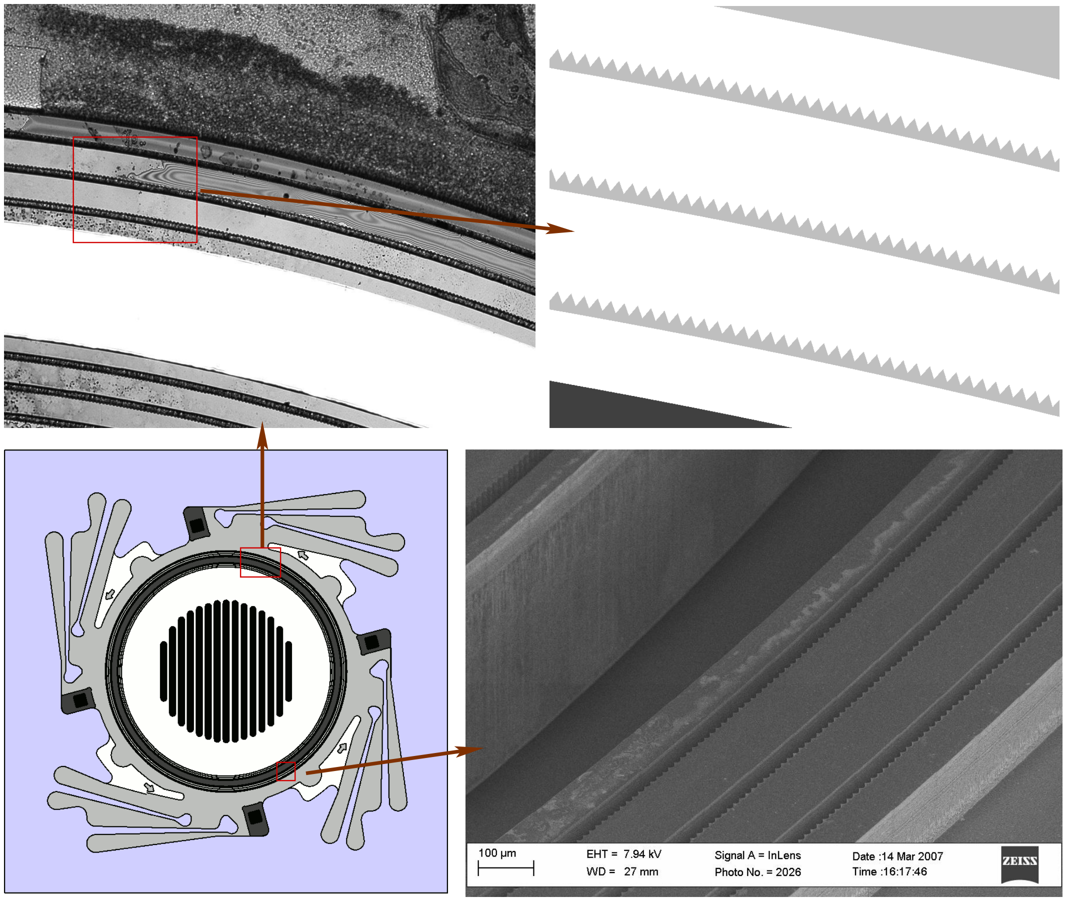

Figure 5-13: Traps to prevent liquid migration to the insulator

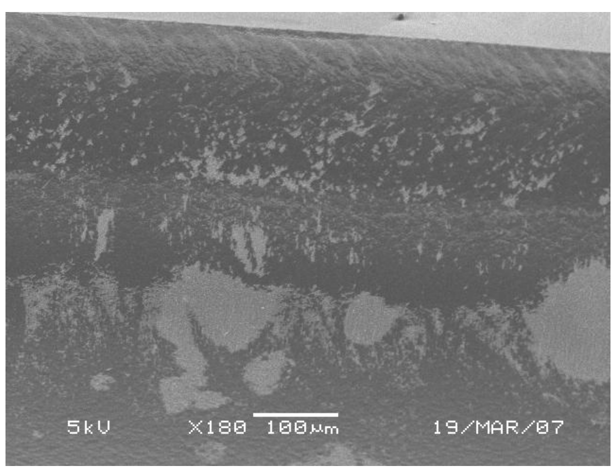

Figure 5-14: The Sidewalls remain dry even after the thruster has been fired. This shows that, as expected, the electric field at the edge of the emitter die is insufficient to draw the liquid up the sidewalls

5.1.8 Active Area

Keeping the spring configuration from Chapter 4, and conservatively integrating all the features we have just reviewed into the thruster design, the area that remains for the thruster is a 12 mm diameter circle. Table 5.1 summarizes the space occupied by each feature, and indicates how this space could be reduced with a less conservative design. We take a safe 750 μm pitch for the extractor slots with a 550 μm opening for each slot. The available area is filled with one of three configurations of the emitters described in Section 3.4: a single one-tip emitter, 216 one-tip emitters (with an extra millimeter of margin around the array), 502 two-tip emitters.

Feature	Conservative (μm)	Aggressive (μm)	Inner Radius (mm)
Emitter Die Diameter			10.60
Structural Support Ring	250	100	10.25
Gap	100	50	10.15
Outer Liquid Traps	500	200	9.65
Insulation Gap	500	200	9.15
Inner Undercut	1000	100	8.15
Pyrex Bonding and Tolerances	1150	500	7.00
Beam Divergence Gap	1000	500	6.00
Outer Insulation Gap	1350		
Inner Liquid Traps	500		

Table 5.1: Width of concentric thruster feature (radius). Features that are not on the critical path are listed at the bottom

5.2 Fabrication

The process flow for the thruster is illustrated in Figure 5-15. We will refer to steps in this figure in the rest of this section. The thruster is a stack made from three six-inch wafers. The Pyrex insulator on top, a silicon wafer in the middle, and the wafer with the springs on the bottom. The silicon wafers are all 675 μm thick, and the Pyrex wafer is 500 μm thick (to speed up laser cutting).

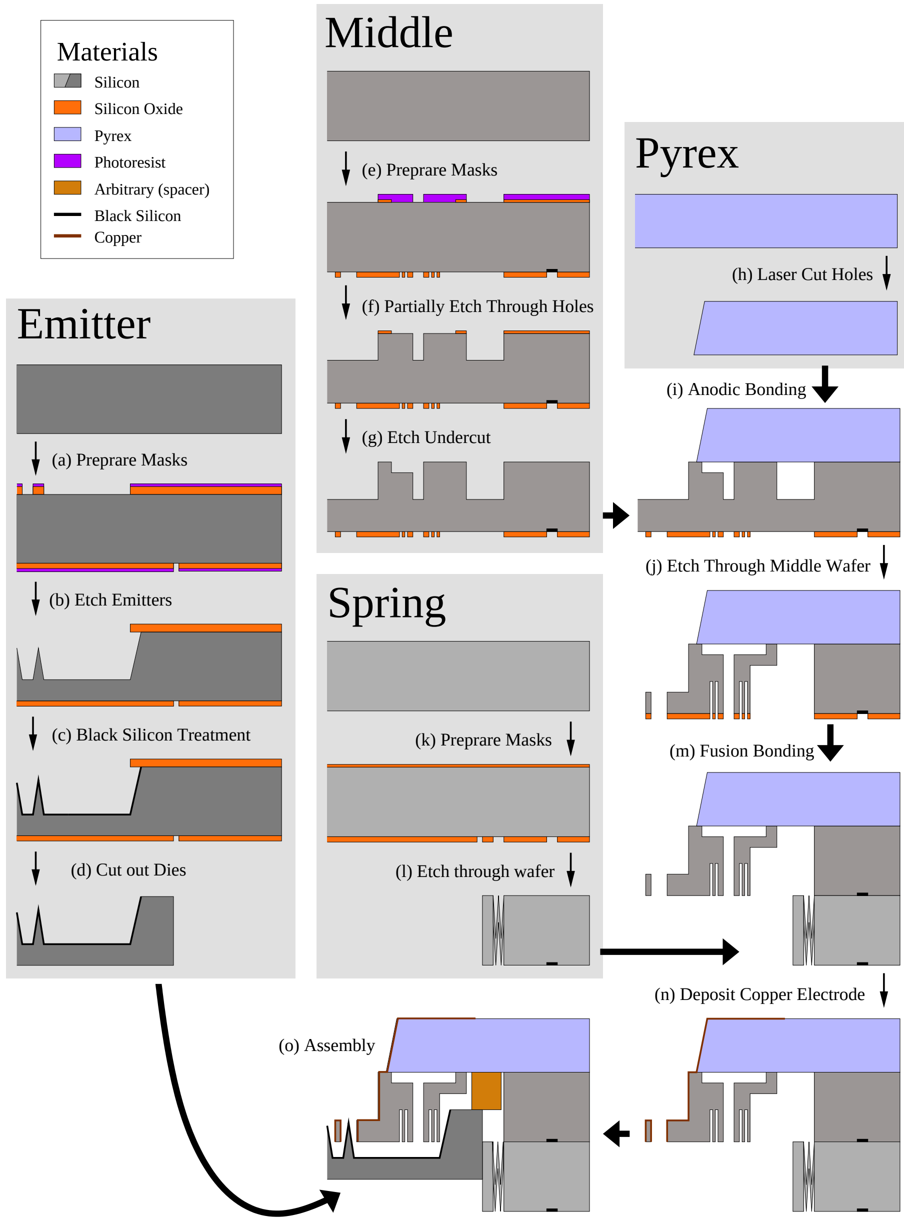

Figure 5-15: Main fabrication steps for the thruster. Only half a die is shown

5.2.1 Masks

Six masks are used to fabricate the thruster. Four masks are used for the extractor component, and two for the emitter dies.

Each emitter wafer contains nineteen dies, nine with one orientation and ten with the other, as shown in Figure 5-16(a). There are six single emitter dies, seven 216 emitter dies, and six 502 emitter dies, equally distributed among the two orientations.

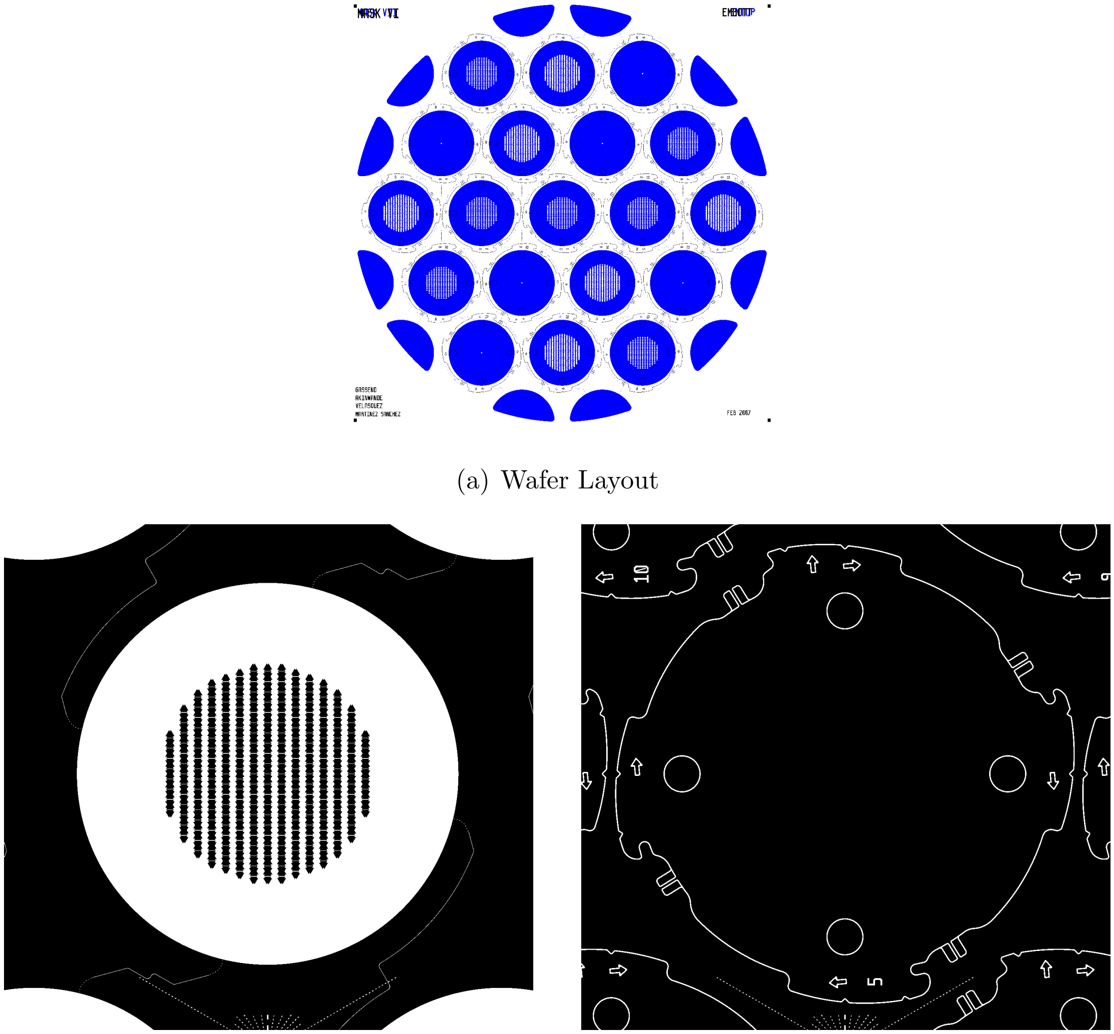

(a) Wafer Layout

(b) Top Mask

(c) Bottom Mask

Figure 5-16: Masks for the emitter dies

Top Mask: This mask defines the etch in Step (b). It describes the geometry of the emitters. It also contains a network of channels to help evacuate gases from the vicinity of the emitters during photoresist baking, when the wafer is target mounted upside-down.

Bottom Mask: This mask defines the etch in Step (d). This mask describes the areas that need to be etched through the wafer: the assembly tool holes, and the outer shape of the die. On this mask, the dies are connected to the bulk of the wafer by four thin tethers. The tethers ensure that the dies remain attached to the wafer even once their shape has been etched out, for ease of handling, and for improved heat transfer. To separate the die from the wafer, the tethers are broken using tweezers.

As with the finger-based assembly test device, there are six dies per extractor wafer, laid out according to Figure 4-13.[2] The masks for the extractor component are shown in Figure 5-17.

Through Mask: This mask defines the etch in Step f for the areas that are to be etched through the wafer, as well as the thinned extractor area. In most cases, through holes are created by etching a thin 70 micrometer trench around the hole. This approach improves side-wall straightness.

Undercut Mask: This mask defines the undercuts in Step g, and was also used to keep track of which device is which.

Recess Mask: This mask defines the etch in Step j, which defines the recesses, the areas that need to be etched through (including the extractor slots and insulation gap), and the liquid traps. It also contains an array of holes to absorb gases that are released after the plasma assisted bond with the spring wafer. For now we do not know if these holes are useful or harmful.

Spring Mask: The spring mask is is the same one used for the test device. It is used to cut out the outline of the fingers in Step l. One electrode component for the test device is generated as a byproduct, but is not used in this process.

5.2.2 Middle Wafer

The middle wafer is prepared using a straightforward DRIE process. A nested mask on the top side defines the insulator undercuts and through holes, and a single mask on the bottom defines the recess and extractor slots.

[2] Referring back to Table 4.1, all the emitter dies now use type B deflection. The post/ring distinction is obsolete.

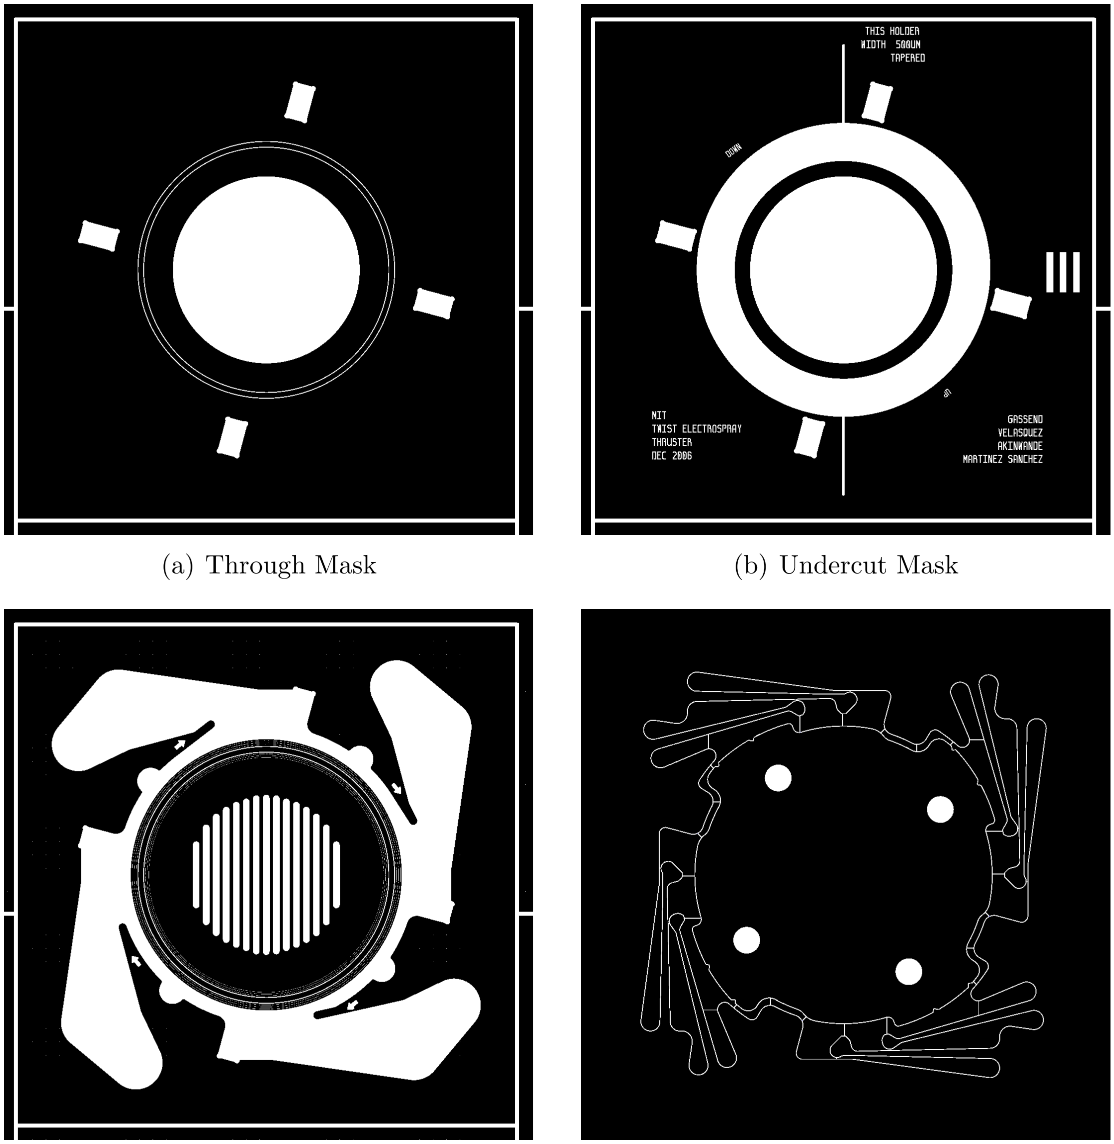

(a) Through Mask

(b) Undercut Mask

(c) Recess Mask

(d) Spring Mask

Figure 5-17: Masks for the extractor components

The virgin 675 μm thick six-inch silicon wafer is RCA cleaned and has a protective thermal oxide grown on it. In our case, two microns of thermal oxide were used, but a few hundred nanometers would have been sufficient. Then, the wafers are cleaned in Piranha, and silicon oxide is deposited on each side using a Novellus Concept One PECVD tool. The goal is to have at least 4 μm of oxide on the bottom side of the wafer, and at least 1 μm on the top. The PECVD oxide is then annealed in nitrogen at 950°C for an hour. Both sides of the wafer are then coated with 10 μm of AZ P4620 photoresist, and prebaked. The top is exposed in ultra-violet with the undercut pattern using contact photolithography. There is then sufficient contrast in the exposed photoresist to use front to back alignment to expose the backside with the recess pattern. The wafers are then developed. The mask for the middle wafer includes trenches separating the dies, intended to reduce die-sawing stresses by cutting the wafer into dies after it has been bonded to the Pyrex. These trenches have caused wafers to break during etching and so they should be painted over using photoresist after the development step. Then the wafer is post baked, and the underlying oxide is etched away in an AME P5000 plasma etcher using an Ar, CHF_3, CF_4 selective oxide etch. A brief SF_6 etch ensured that the alignment marks will remain visible in the silicon even once the oxide is stripped. The photoresist is then stripped in an oxygen plasma, and second photolithography step is used to make a photoresist mask on the top side of the wafer for the features that go through the wafer. The photoresist for this step is dispensed in two steps, separated by a 20 minute prebake, in order to reach a thickness of 20 μm which is guaranteed to mask a 450 μm etch. This completes Step e of the process.

Now that all the masks have been laid out, the DRIE etching begins. The wafer is etched down 450 μm from the front side, defining the through holes and thinning the extractor area (Step f). Then the photoresist is stripped using Piranha, and a short 20 μm DRIE etch is performed to etch the insulator undercuts (Step g).

5.2.3 Cutting Pyrex

The Pyrex wafer is cut using a Resonetics excimer laser at 248nm (Step h). Before cutting, both sides of the wafer are coated with thick photoresist and hard-baked. To maximize cutting speed, a large 1 mm beam is used, at a repetition rate of 100 Hz, with a pulse energy of 450 mJ[3], and a feed rate of 100 μm per pulse. In these conditions, about 50 passes are necessary to get through the Pyrex. With fewer passes at slower feed rates, there is a risk of cracking the Pyrex. More passes at

[3]Not all this energy is actually delivered to the substrate.

higher feed rates lead to more time lost moving from hole to hole between cuts.

The quality of the Pyrex cut is poor and seems to proceed through fracture rather than ablation. Indeed, when the laser is fired at a smooth Pyrex surface, there is no visible effect until some missing chunks are suddenly observed. Moreover, a lot of dust is observed on the Pyrex surface after the cuts are made. Given the large size of the feature being cut, and the low design tolerances on the order of ±1 mm, this poor cutting quality is not a concern. The largest possible beam size is used to speed the cutting process by maximizing the linear distance which is covered by each pulse. A taper on the order of 75° is observed for the laser-cut sidewall. For beam sizes smaller than the wafer thickness, a roughly V-shaped grove is cut in the Pyrex and etching effectively stops. If a large enough beam size is not available, it can be simulated by cutting multiple trenches offset by a few hundred micrometers. These trenches end up merging and allow the cut to extend through the wafer. Figure 5-18 illustrates this method. Once the Pyrex is cut, it is rinsed to remove as much dust as possible. The resist coating is then stripped in Piranha.

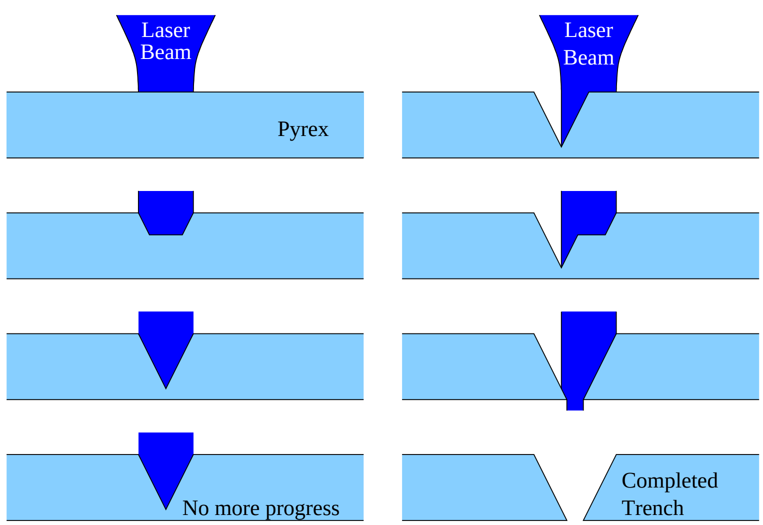

Figure 5-18: If the laser beam is too thin, it cannot cut through the Pyrex. In that case a slightly offset cut can finish the job

5.2.4 Extractor Separation

Once the Pyrex wafer and middle wafer have been prepared, the next step is to bond them and to separate the extractor region from the rest of the wafer (Step i).

First the bonding surfaces are prepared and cleaned. The oxide needs to be removed from the front of the middle wafer. This can be done using the AME P5000 selective oxide etch, or using 49% HF. Because the oxide is thinner on the front of the wafer than on the back, the wafer can be dipped in HF until the front is clear and then promptly rinsed. The progress of the oxide etch is easily observed visually from the colorful fringes in the oxide, which suddenly disappear when the oxide clears. Once the oxide is cleared, a Piranha clean is performed on both the silicon and Pyrex wafers. They are then anodically bonded. Alignment of the two wafers is done by eye, since millimeter precision is sufficient. This avoids patterning fine alignment marks on the Pyrex wafer using the laser cutter, saving a lot of laser cutting time. Indeed, finer cuts take more time because each pulse covers a smaller linear distance along the cut, unless an elongated beam is used. However, the elongated beam requires proper alignment between the beam and the translation stages, that was not available for us.

The stack is then mounted with photoresist onto a handler wafer, Pyrex side down. The mounting is in preparation for a DRIE etch of the Recess Mask (Step j). This etch needs to get through the wafer to separate the extractor region from the rest of the silicon wafer. It has to take place after the anodic bond because otherwise the alignment between the extractor and the rest of the middle wafer would be lost. This step also etches the recess below the springs into which the emitter die can be pushed, clears the holes for the spacers, and defines the extractor slots. This etch is carried out on the oldest and least uniform DRIE machine that is available, "STS1", as it is the only one that will accept Pyrex. We found that once the extractor slots were etched through, the beams of the extractor started to be under-etched, possibly because of over-heating. Rotating the wafer during etching was critical to maximize uniformity. Despite this, some dies etched faster than others, and about half the extractors were destroyed on each of the three runs we performed (see Figure 5-19).

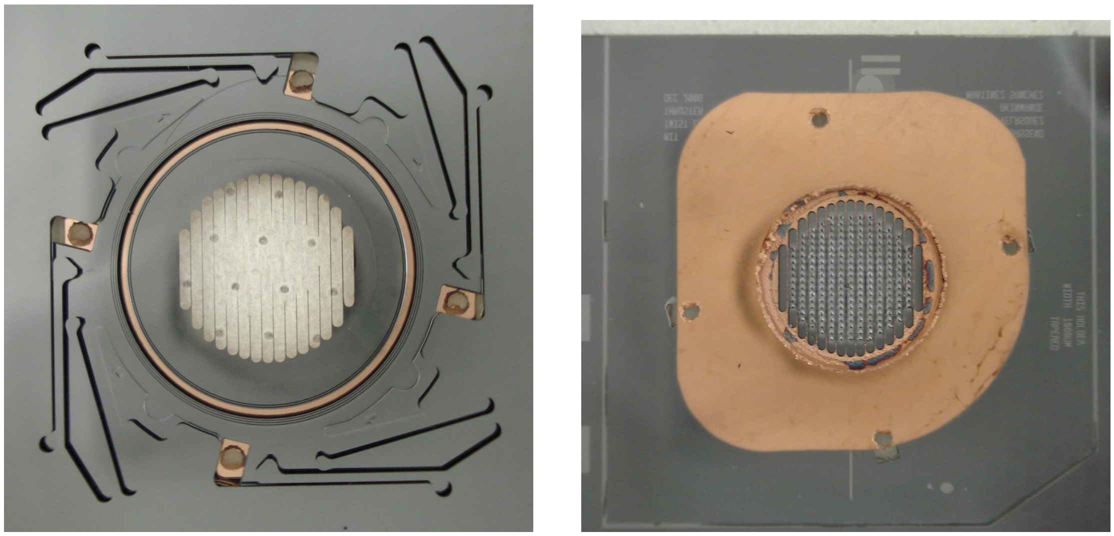

(a) Back side (b) Front with assembled emitters visible

Figure 5-19: Dies with damaged extractor

Once the etch is completed, the stack is cleaned in an oxygen plasma, and dismounted from the handler wafer in Piranha or acetone. It is now ready to be bonded to the spring wafer.

5.2.5 Bonding to the Springs

The spring wafer is prepared exactly as in Section 4.3, except that the alignment mark steps may be omitted (steps k and l).[4] The procedure to bond it to the rest of the thruster differs, because the silicon-Pyrex stack cannot undergo the 1000°C which is usually associated with silicon direct bonding. Instead, we used plasma assisted bonding [107–111].[5] The bond quality requirements in this step application are significantly less than for applications which require sealing at the bond interface. The presence of voids will not change the device performance as long as the spring wafer remains rigidly attached to the middle wafer.

[4]If the alignment mark steps from the finger assembly test device are omitted, the alignment marks from the Spring mask should be etched into the silicon to allow alignment for bonding.

[5]We also did one attempt with thermocompression bonding [112] for which the bond failed during die sawing.

In preparation for bonding, the wafers are cleaned in Piranha, and then stripped of oxide in 49% HF. Once again visual cues are used to determine the end of the oxide etch. The wafers are then placed in an oxygen plasma. In our experiments the plasma exposure lasted 15 to 30 seconds in a 1 kW plasma. The wafers are then dipped in deionized water and dried. Finally, they are aligned and bonded (step m).

Different bonding procedures are possible. Some tests were carried out in which the bonding sequence was carried out on unprocessed wafers. The silicon wafer and the Pyrex-silicon stack were aligned with an EV620 aligner, then transferred to an EV501 bonder to be heated to 400°C in a nitrogen ambient, contacted, and cooled after 2 hours. This procedure led to one good bond and one partial bond on test wafers. The quality of the bond was evaluated by cutting the wafer into 1 cm strips, and further cutting strips into 1 mm slices. None of the slices debonded.

With the device wafers, the wafers were aligned, contacted and briefly pressed on the EV620 aligner using a *silicon direct bonding* recipe. Infrared observation showed that the wafers had bonded over their whole area (Figure 5-20). To avoid the risk of cracking the wafers in the EV620, the stack was then annealed without applied pressure at 400°C for 2 hours in a nitrogen ambient. After the anneal a crack was visible in the Pyrex around one die, and the wafers had partly debonded in that area. The stack was then die-sawed. Half the dies debonded during the die-sawing step. An a posteriori reexamination of Figure 5-20 reveals some slight fringes to the left of the top die. Perhaps the apparently well bonded wafer was in fact barely bonded before the anneal. In this case pressing the wafers in the EV501 bonder may have improved the bond quality and avoided problems at the die-saw step.

5.2.6 Electrical Connection

Finally, a 1-μm-thick copper electrode is sputtered onto the Pyrex side of the stack (step n). An aluminum shadow mask, roughly positioned above the wafer, is used to avoid depositing copper near the edge of the die, which could increase the risk of flashover. The role of the copper is to make a contact between the easily accessible top surface of the thruster and the recessed extractor. Its role is also to make the

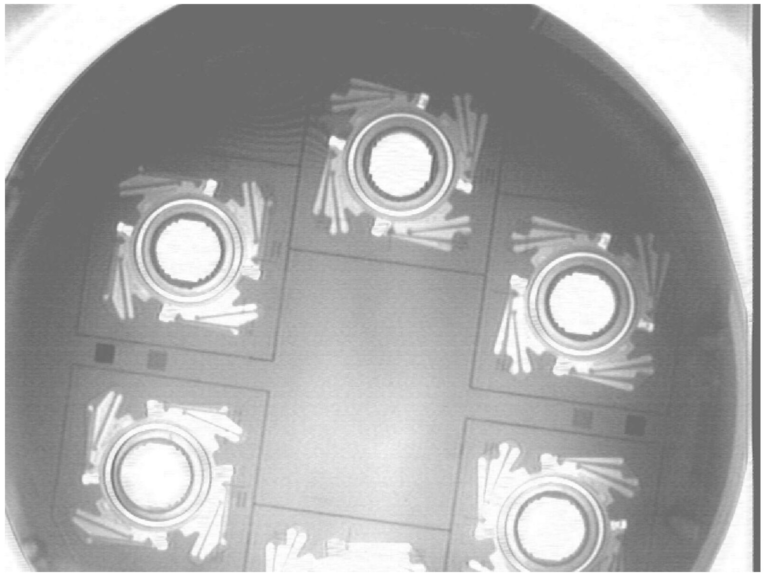

Figure 5-20: Picture of the bonded wafers, before anneal, lit from behind with infrared light

top surface of the thruster conductive to avoid spurious charging during operation. Unfortunately, the deposited copper was not able to make a reliable electrical connection between the top of the Pyrex and the extractor electrode, despite the taper in the Pyrex, so we had to use a drop of silver paint to make the connection (see Figure 5-21).

When the dies were tested up to 1.6 kV in air, we observed arcing between the copper electrode and the middle wafer, through the tweezer holes that had been cut through the Pyrex to facilitate spacer removal. Luckily, the copper had been deposited without an adhesion layer, and we were able to remove the copper around the spacer area by scratching it and lifting it off using tape. We then covered the hole in the Pyrex using Kapton tape, and used copper tape, aluminum foil, and/or silver paint to make the top surface of the thruster conductive. The result is illustrated in Figure 5-21.

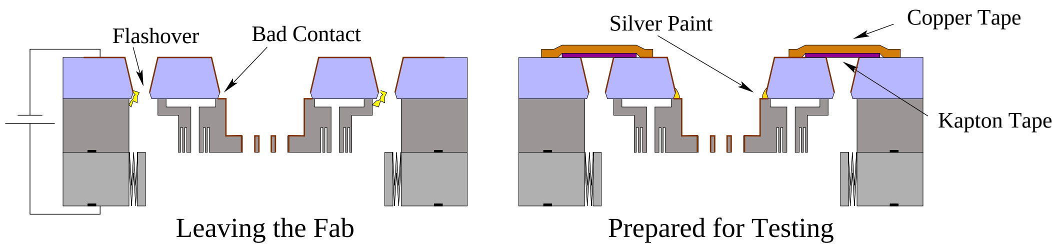

(a)

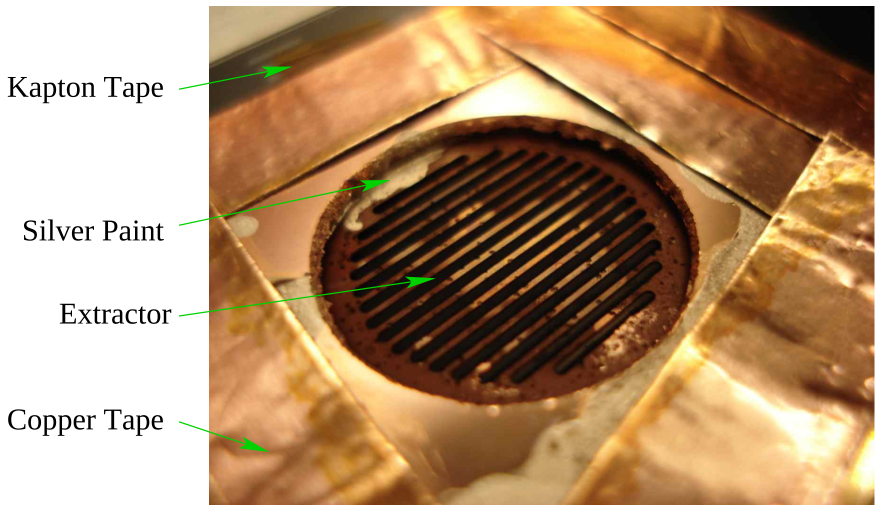

(b)

Figure 5-21: As it leaves the fab, flashover can occur through a tweezer hole, but the contact to the extractor is poor. These problems were fixed by removing some copper, and applying Kapton tape, copper tape and silver paint

5.2.7 Emitters

The emitter die preparation starts by setting up oxide masks on both sides of a silicon wafer. Then emitters are formed on the front side, and the emitters are treated for wettability. Finally, wafer is separated into dies with the correct shape for assembly by DRIE etching from the back side.

Preparation of the oxide mask is identical to the first steps of the middle wafer preparation: oxide is deposited on both sides of the wafer and patterned (Step a). The front side oxide is 6 μm thick, and the back side oxide 4 μm thick.

Then the wafer is mounted with photoresist, front side up, onto a handler wafer. This mounting step was added to avoid wafers breaking during the subsequent emitter-forming etch which goes most of the way through the wafer. The emitters are formed using an SF_6 etch, followed by a tapered DRIE etch (see Section 3.4.1), and finished by another SF_6 etch (Step b). During the final SF_6 etch, the emitters do not all complete at the same time.[6] Therefore, toward the end of the etch, the wafer is removed every ten minutes and inspected. Completed dies are clearly visible by the oxide masks scattered around the array of emitters. Completed dies are coated in OCG825 photoresist using a Q-tip, and the wafer is baked for 30 minutes to dry the resist before continuing the etch.

Once all the dies are ready, the wafer is dismounted in acetone or piranha and cleaned of resist in piranha or in an oxygen plasma. Then the wafer is ready for the black silicon surface treatment. It is critical that the wafer be wetted with water between the SF_6 etching and the black silicon treatment or the black silicon will not form. Thus, if acetone and oxygen plasma were used for dismounting and cleaning the wafer, the wafer should be purposely wetted and dried, for example in a spin-rinser and dryer. The black silicon treatment is made using recipe R_0 from Table 3.1 (Step c).

All that remains now is to separate the dies. First the dies are target mounted onto a handler wafer. This target mounting is somewhat critical because once the dies are separated, each die must be kept in thermal contact with the substrate to avoid overheating. Indeed if a die overheats, the DRIE etch undercuts the oxide mask, spoiling the notches of the assembly system. First, the emitters are covered in OCG825 photoresist using a Q-tip to protect the surface treatment from damage once the assembly tool holes have been etched through.[7] Then the handler wafer is coated in AZ P4620 photoresist, except a 1.5 cm band around the edge of the wafer. The goal of this resist-free band is to allow gases which are released during baking

[6] As usual, emitters at the periphery of the wafer complete first. Single emitter dies etched fastest, then single-tip-emitter arrays, and finally two-tip-emitter arrays. Finally, emitters at the edges of arrays completed before ones in the middle of the array. These observations suggested that the masks from neighboring emitters prevent some low incidence angle ions from reaching the sidewalls and etching them. Moreover, microloading (depletion of reactants near large open areas) seems not to have been a major source of non-uniformity, as it would have had the opposite effect on which emitters completed first. Any effect of azimuthal non-uniformity in the etch was masked by the difference between different array types.

[7] This step is precautionary. In principle, a polymer coating should form on the emitter surface during DRIE etching, and protect it.

to escape around the edge of the wafer. Indeed, channels were left between emitter dies allowing gases to flow from any emitter region to the outside of the wafer. Then the emitter wafer is mounted, emitter side down, onto the handler wafer, and baked at 95°C for 1h. After baking, it was clear that gases were not circulating between the emitter regions and the outside, because the thinned emitter regions were visibly bowed downward. However, when the wafers were placed in a vacuum, gases trapped around the emitters were able to escape, and no further bowing was observed.[8]

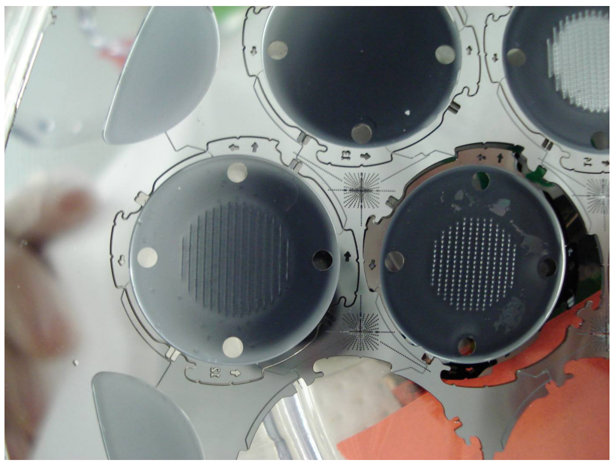

Figure 5-22: Emitter wafer with some dies already removed, some dies still attached, and one die that has just been detached

The wafer is then separated into dies using DRIE (Step d), and unmounted using acetone (optional), followed by a piranha clean. Then the oxide is stripped in 49% HF. After the oxide is stripped, it is important to dip the wafer in piranha or to place it in an oxygen plasma. Otherwise the silicon surface will keep its hydrophobic H-termination which will prevent wetting by the propellant. Finally, the tethers that hold each die to the wafer are broken using tweezers (see Figure 5-22).

5.2.8 Assembly and Disassembly

Assembly and disassembly proceed as in Chapter 4. The emitter die is assembled with a plastic tool that mates with the assembly holes in the die. When the emitter die

[8]Once again, we have not fully explored which part of this procedure is necessary, and which part is overly cautious.

is being recessed to raise the emitter height, the spacers are inserted using tweezers. Then the die is twisted in place, and pressure is applied to the die above each spacer. The downward motion of the die is nearly imperceptible for the person pushing. (One die cracked after repeated vain attempts to push it down when it was already down, this die is the only casualty we have had during assembly/disassembly.) The fact that the die is recessed can be observed by moving the tweezers across the step between the spring wafer and the die. Broken dies have been assembled with no particular difficulty.

For disassembly, tweezers are used to pry up the emitter die. Tweezer holes have been included in the recess so that it is easier to access the edge of the emitter die when it is recessed (see Tweezer Hole in Figure 5-3).

5.3 Process Improvements

This process has only been successfully completed once. Numerous improvements can be made for later generations of the device:

Harmful Features

Some features seem like a good idea during the mask design stage, but end up causing major trouble during processing. We highly recommend updating the masks to remove such features.

- The features to separate the wafer into dies systematically cause dies to break and should be removed. Alternatively, rotating them 45° so that they are no longer parallel to the crystal plane may be sufficient. Initially these features were introduced to reduce the stresses during die-sawing. With the current masks, these features need to be painted over using photoresist to avoid wafer breakage.

- The features to help evacuate gases from the vicinity of the emitters may have contributed to emitter wafers breaking near the end of the emitter etch. It would be useful to test if they are actually needed for gas removal, or if a more conventional target mounting scheme would suffice. If possible, they should be eliminated.

- The array of holes in the recess mask to absorb gases produced after bonding may reduce the quality of the bond. Further tests are needed to see if these holes have a positive or negative effect.

- The spacer removal holes in the Pyrex caused flashover between the front-side copper electrode and the middle wafer. The spacers seem to fall out by gravity when the thruster is upside-down, so the holes in the Pyrex can probably be eliminated without ill-effects. This would not be necessary if the copper electrode was replaced by an external conical electrode as discussed in the next section.

Better Insulation

We shall see in Section 6.2.8 that the Pyrex insulator begins to leak significantly as temperature is increased beyond room temperature. It would be desirable to replace the Pyrex by a less conductive substrate. For example, fused silica could be low-temperature fusion-bonded or thermocompression bonded [112] to the middle wafer, and the exposed bonding metal could be wet-etched away. Non-porous ceramics that are available in wafer form, and that can be polished for wafer bonding are also good candidates. : In the case of Pyrex, a large fraction of the leakage current flows through the wafer to the front-side electrode, as opposed to flowing across the wafer to the extractor. Thus the leakage current could be reduced by a factor of 10 to 100 by using a conical front-side electrode as suggested in Section 5.1.4.

More Emitters

To maximize thrust, the goal is to pack as many emitters as possible into a device. How can this goal be achieved? We briefly address this question from the fabrication point of view, and revisit it in Section 7.7, in which we consider scaling of the thruster.

The first way to increase the number of emitters is to increase the emitter density. Less conservative spacing between extractor slots can be attempted to incorporate more slots. Also, ridges without vertical side-walls can be tried, as they allow tighter tip spacing. Reducing the emitter height would also allow tips to be brought closer together.

The other method is to increase the area available for emitters. This could be done by switching to a full-wafer version of the thruster. Alternatively, it may be possible to save some space in the six-device-per-wafer version of the wafer.

First, space may be saved by reducing the margins that were taken for this design. Currently, the active area has a diameter of 12 mm out 21 mm that is available. If 5 extra active millimeters could be gleaned, the number of emitters could be doubled. From Table 5.1, prime candidates to achieve such a reduction include: insulation gap, inner undercut width, Pyrex size/position uncertainty, and liquid trap spacing.

More space can be saved by redesigning the springs, for example by reducing their number or their clamping force. Also, fancier spring geometries can be attempted such as multiple beams connected in parallel. This increases stiffness by multiplying the number of beams, but reduces the beam length that is needed to achieve a specific deflection.

Thicker emitter wafers

There is nothing preventing thicker emitter wafers from being used. This would increase the robustness of the emitter dies, ease handling, and allow taller emitters to be tried.

Emitter uniformity

The current practice of protecting completed emitters with photoresist, while other emitters continue to be processed gives unsatisfactory results. Figure 5-23 shows an emitter for which the protective resist did not coat the ridges well, causing the double ridges that are visible in the micrograph. To avoid this problem, other masking methods could be used. For example, a piece of Kapton tape could be used to cover the emitter region on dies that are ready.

For a production environment, the best solution would be to vary the emitter mask size across the wafer to compensate for process variation. This approach should allow all the emitters to be ready simultaneously, as long as the etch is sufficiently repeatable.

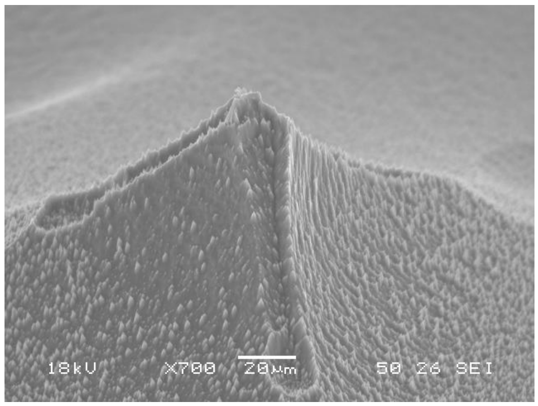

Figure 5-23: Once it had been etched to a point, this emitter was covered in photoresist to prevent it from being destroyed by further etching. However, the photoresist did not coat the ridge lines, allowing the ridge lines to continue etching

Wet etching of Pyrex

Currently the Pyrex is cut using a laser-cutter. Wet etching could be an interesting alternative which is more amenable to batch fabrication, given the lax precision that is needed on the Pyrex cut. For example, a Cr:Au etch mask could be used [113].

5.4 Conclusion

In this chapter we have seen how to use the assembly concept from Chapter 4 to make a complete thruster (see Figure 5-24). The most difficult challenge was insulating the extractor electrode from the emitter electrode. Some enhancements were made to the assembly scheme to allow control over the vertical position of the emitters relative to the extractor.

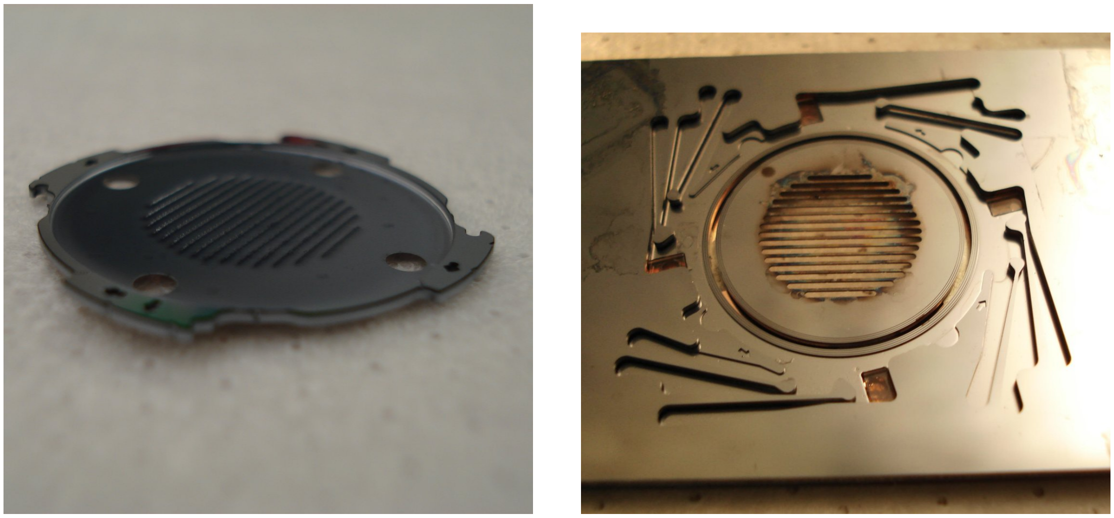

(a) Emitter Die (b) Extractor Die (has been fired)

Figure 5-24: Completed emitter and extractor dies

The fabrication process for the thruster is straightforward DRIE etching with oxide and photoresist masks. Currently yield is limited by two steps. The first one is the backside etch which defines the extractor slots and recesses. For now we have not managed to finish this etch without destroying at least two out of six extractors. The second difficult step is the fusion bonding step. The single attempt which was made on device wafers appeared to bond well but half the devices debonded during die sawing. More experimentation is needed to find the best parameters for this bond. Only two devices were completed, but four more simply need to have springs bonded to them, perhaps using thermocompression bonding.

The completed thrusters were very successful. Assembly, insulation, firing and disassembly all proceeded smoothly. The assembly method proved very useful during testing, as we were able to test many different emitter dies with the two completed extractor components. After use, emitter dies could be disassembled, observed, cleaned, refueled and reused. Once again, the assembly method proved its robustness. Two emitter dies broke losing two of the finger notches. They could still be assembled and fired with no complications. One thruster fell on the floor from table height. It lost two fingers but remained serviceable and continued to provide data for many runs. (Eventually, a crack formed across the insulator region allowing liquid infiltration and short-circuiting.) The next chapter goes into the details of the test firing that was done with the thrusters.

Chapter 6

Experimental Results

We now turn to the experimentation that was performed to characterize the fabricated thrusters. These data were taken over an intense three week period, and we use them to try to give as broad a view as possible of the thrusters. In total, the vacuum chamber was pumped down 16 times, with 7 different emitter dies, to generate these data, and as many tests as possible were carried out during each run. Consequently, there are no runs where the thruster was simply left to fire for a long time, and since some of there experiments were run opportunistically, there are gaps and inconsistencies in the data that would require more experimentation to resolve. Each one of the measurements we have taken could be the object of an extensive systematic study. In writing the present chapter, we have attempted to point areas where more thorough experimentation would be particularly enlightening.

We begin this chapter with some general remarks on the experiments in Section 6.1. Then, in Section 6.2 we look at how much current is emitted by the thruster, as a function of voltage, temperature, and direction, as well as how much current is collected on the extractor electrode. Time-of-flight measurements, which measure the velocity of the emitted particles are presented for EMI-BF_4 and EMI-Im, in Section 6.3. Some measurements in which the thruster was fired against a plate to evaluate emission uniformity across the tips are presented in Section 6.4. The remaining sections discuss emitter damage during operation, and the effect of dust particles contaminating the thruster during assembly. Overall, this chapter focuses on presenting the data. A synthetic interpretation of the data will be given in Chapter 7.

6.1 General Remarks

6.1.1 Electronic Equipment

To take the data which are reported in the rest of this section, the following equipment was used:

Current Measurement: A Keithley 6514 and a Keithley 6517A electrometer were used for current measurements.

Oscilloscope: An Agilent 54835A Infiniium oscilloscope was used to gather waveform data, and to take voltage measurements.

Pulse Generation: Pulses for the time-of-flight gate were provided by a BNC 555 pulse delay generator, and amplified with a DEI PVM-4210 output pulser module which outputs pulses at up to $\pm$950 V.

Power Supplies: The emitter voltage was supplied by a $\pm$5 kV Matsusada AMS-5B5, which was controlled by an Agilent 33220A function generator. The voltage being output was monitored on the 54835A oscilloscope. This setup was accurate to within 20 V. A -3.5 kV Acopian P03.5HN17 power supply was used to set the secondary electron suppression grid voltage, and a +3.5 kV Acopian P03.5HP17 was used for the Einzel lens in the time-of-flight experiments.

High Speed Current Amplifier: A home-made high speed current amplifier built by Lozano [23] was used to measure the high speed time-of-flight signal. It is based around a TI OPA602 or a TI OPA655 operational amplifier.

This equipment was programmed via IEEE-488 General Purpose Interface Bus (GPIB) programmed in C++, and interfacing to the GPIB bus via National Instruments remote VISA.

6.1.2 Thruster Preparation

Propellant was supplied to the emitter die using a syringe. An estimated 0.1 to 1 mm^3 of liquid was deposited on the black silicon treated surface. The liquid could be deposited either on the flat area around the emitters, or directly onto the emitters. There seemed to be less current intercepted on the extractor electrode when the liquid was supplied directly to the emitter area, though more experimentation would be needed to confirm this. Liquid would spread across the black silicon surface within a few tens of minutes, but the drop which was initially deposited would remain visible for many hours. Sometimes the liquid was given only a few minutes to spread before the emitter die was assembled and placed in the vacuum chamber, on other occasions many days elapsed between these two events. No difference was observed between these two cases.

Once wetted, the emitter die was assembled to the extractor component. In some cases emitter dies which had broken and lost one of the pairs of notches were used; they assembled without any difficulty. These broken dies seem to have suffered from more intercepted current than unbroken dies, but more testing would be required to test this hypothesis. In all but the first runs, the emitters were recessed using 500-μm-thick Pyrex spacers. This seems to have reduced intercepted current, but once again more testing is needed to confirm the hypothesis. After use, an emitter die could be disassembled and rewetted. Generally, performance decreased the second time the die was used.

After assembly, the thruster was mounted to a high density polyethylene holder (see Figure 6-1). A flexible clamp on the front side held the thruster in place and provided electrical contact to the extractor electrode. In some cases, some Kapton tape was used in addition to the flexible clamp, particularly when the holder was mounted upside-down. Electrical contact to the emitters was made via the spring wafer, and a strip of copper tape attached to the polyethylene holder. Because of the high voltages involved, little care was taken to make good electrical contacts.

Figure 6-2 shows the electrical connections to the thruster. The extractor elec-

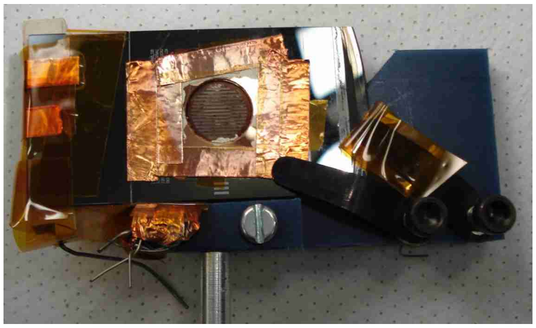

(a) Front

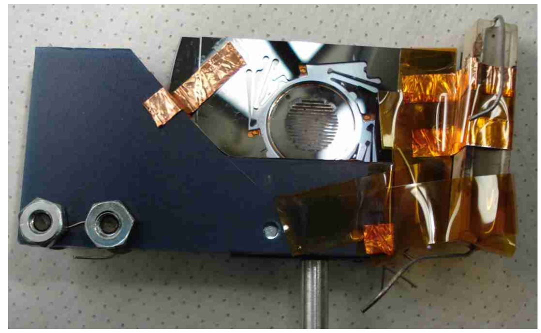

(b) Back

Figure 6-1: The thruster (without the emitters) mounted in its polyethylene holder. A heating resistor is attached to the thruster in these pictures

trode was grounded through a 100 kΩ protection resistor and a Keithley electrometer to measure the current intercepted by the extractor electrode or leaked through the insulation. The emitters were supplied with high voltage through a 100 kΩ resistor. The voltage across this resistor was monitored using a hand-held multimeter, floating at high voltage, to check that all the current from the power supply was being collected by the extractor or the collector electrodes. This was indeed the case, except during the angular measurements presented in Section 6.2.9, and the time-of-flight measurements presented in Section 6.3, in which the collector plate was placed farther away from the thruster. Figure 6-3 shows the electrical characteristics of a thruster for which the extractor was shorted to the emitters. At the operating voltage of the thruster, the electrical resistance was less than 300 kΩ, of which 200 kΩ corresponds to resistances which were intentionally placed in the circuit.

Once mounted to its holder, the thruster was placed in the vacuum chamber, which was pumped down to a base pressure below $1 \cdot 10^{-5}$ torr before starting the experiments. The vacuum chamber, shown in Figure 6-4, was pumped down by two Varian TV-70 turbo-pumps, and the base pressure was measured by a Varian ionization gauge or IMG-100 inverted magnetron gauge.

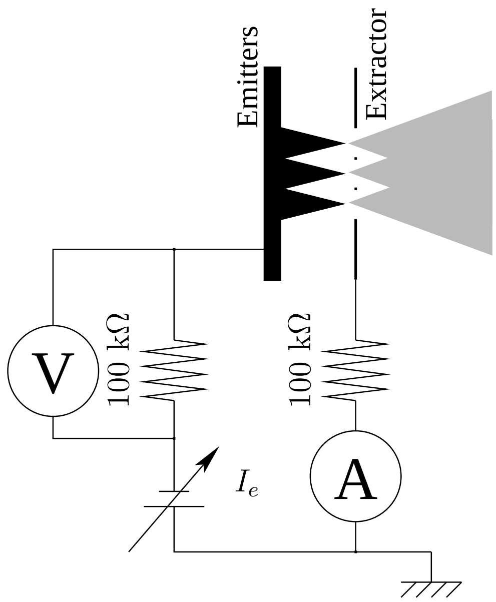

Figure 6-2: The standard electrical connections for firing the thruster

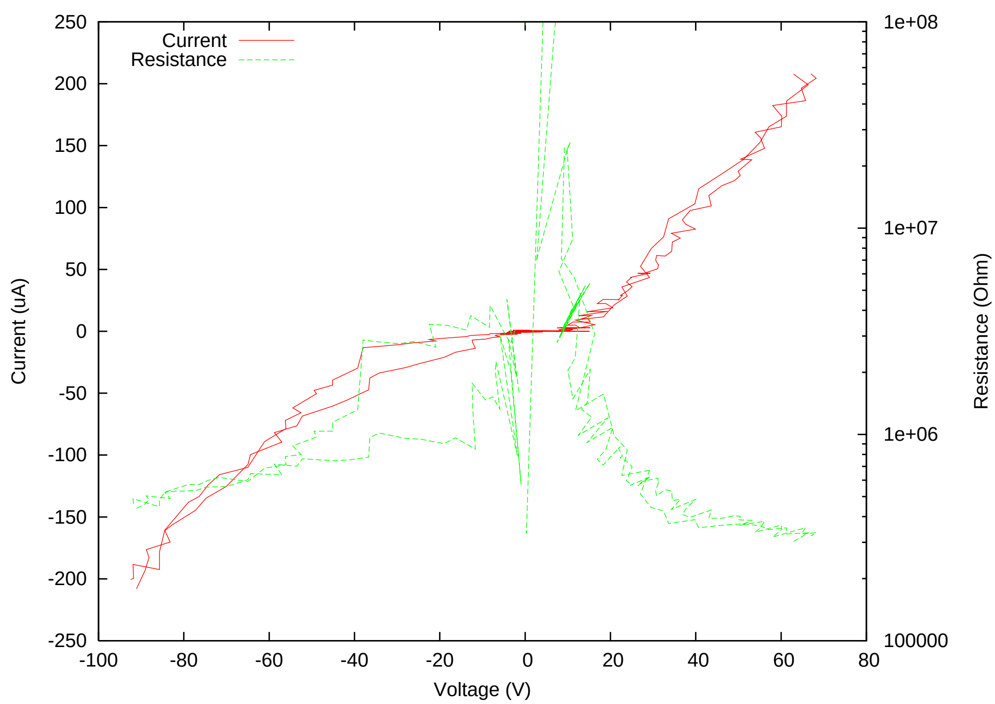

Figure 6-3: The IV characteristic of an intentionally shorted thruster

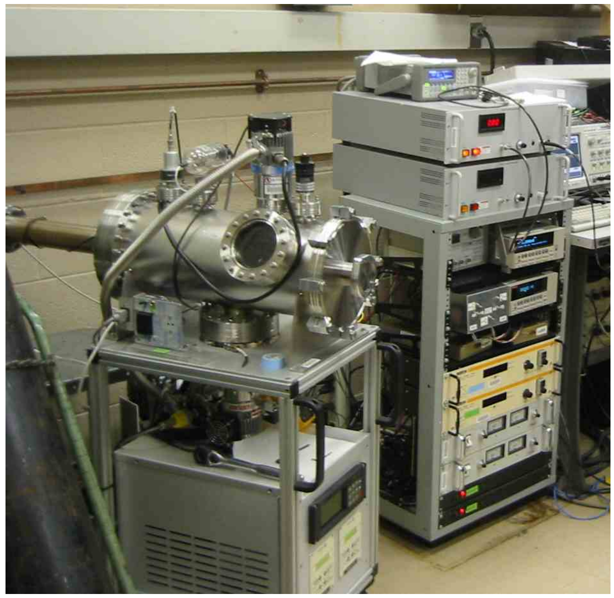

Figure 6-4: The vacuum chamber (left) and power supplies (right) for the experiment

6.1.3 Visual Observation of Thruster Operation

When the thruster was firing at currents greater than 10 μA and the room had been darkened, it was generally possible to see the thruster was firing (See Figure 6-5). A bluish glow was visible, mainly in the area between the secondary electron suppression grid and the collector plate. The emitter region was usually dark, except when the thruster was running low on propellant (see Section 6.2.4).

At high emitted currents (100 μA or more), small sparks were sometimes visible. In one particular case, the whole thruster was be covered in sparks at currents above 400 μA, but only when it was operating in the negative polarity. These sparks were present in the emitter region, but also around the edges of the thruster, and even on some parts of the holder. Perhaps gas being released by electrochemical reactions in the negative polarity [65] sufficiently increased the pressure around the thruster to allow breakdown in all the highly electrically stressed areas. In general, the pressure in the chamber increased during high-current thruster thruster operation, from about $5 \cdot 10^{-6}$ torr to $5 \cdot 10^{-5}$ torr.

Figure 6-5: The thruster firing, as seen through the chamber window. The collector plate is visible on the left, and the thruster can just be made out on the right

6.2 Current-Voltage Characteristics

The simplest characterization of an electrospray array is its Current-Voltage (IV) characteristic. The experimental setup is shown in Figure 6-6. In this experiment, the array was fired against a grounded collector plate. Currents to the collector and extractor were measured using a Keithley electrometer. A secondary electron suppression grid was placed in front of the collector plate, and biased to -50 V. The purpose of this grid was to repel the secondary electrons generated when ions hit the collector back to the collector. Without the grid, these electrons would have accounted for a spurious current to the collector plate.

The IV measurements were entirely computer controlled. Collector current, extractor current, emitter voltage and time were logged. The voltage was usually ramped across the desired range three times to check consistency. Since there is no emission at low voltage, the ramp rate was usually increased at low voltages. About 3 samples per second were taken, with increments between 5 and 250 volts between samples, depending on the experiment. For experiments that reached high currents, larger increments were used to avoid depleting the propellant supply. The ramp had a predefined maximum voltage and current, whichever limit was reached

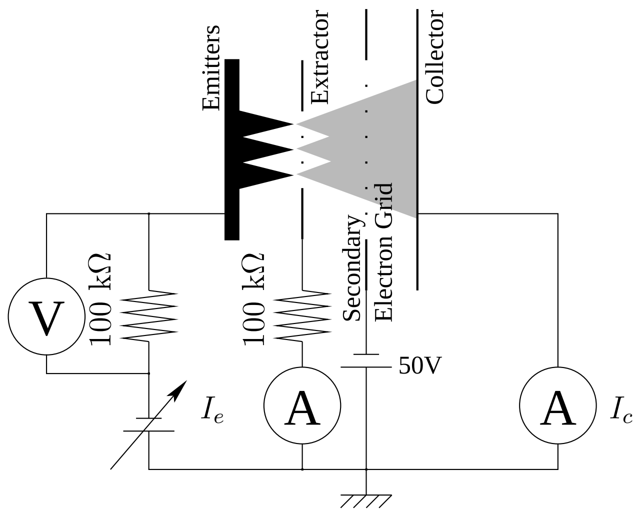

Figure 6-6: Experimental setup to take IV characteristics

first would cause the voltage to ramp back toward zero. Measurements taken with different ramp rates are consistent with each other in all the cases we investigated, except when depletion occurs at high currents.

6.2.1 Stability of Operation

When a thruster was first started, a higher voltage was needed to get it started than during subsequent operation. This extra voltage may have been necessary to complete the wetting of the emitters (see Section 3.4.3). Assuming that this was the correct explanation, we always briefly increased the voltage applied to a freshly started thruster until a collected current in the ten microamp range was reached, to ensure that all the emitters had the opportunity to be wetted. After this conditioning process, the thruster would go through three distinctive phases of operation:

Over-wet Phase: This phase was characterized by high interception on the collector electrode and unsteady emission.

Steady Phase: In this phase the thruster exhibits a repeatable IV characteristic with very low interception current.

Depletion Phase: In this phase, interception remains low, but the current for a given voltage slowly decreases.

6.2.2 Over-wet Phase

Figure 6-7 shows the first few minutes of operation for a thruster that has just been refueled and placed in the vacuum chamber. The IV characteristic is highly inconsistent during this phase. After a few minutes of unsteady operation and large interception, the intercepted current suddenly drops to a few percent only, marking the end of the over-wet phase.

6.2.3 Steady Phase

In the steady phase, the IV characteristic is steady and consistent over repeated voltage ramps. Figure 6-8(a) shows the current per emitter for three different thrusters. One of the thrusters was tested on two different occasions, with the propellant being replenished between the firings. In this plot, the current has been normalized to the number of emitters. The current appears to be linear above 100 nA per emitter, with a 1.2 nA/V slope. The slope is consistent across all the runs. However, the voltage at which the linear portion intersects the x-axis varies between 900 V and 1400 V between samples. The highest current we observed was just over 1 μA per emitter, but at these high current levels, the thruster quickly enters the depletion phase and we were unable to get repeatable IV traces like these.

The low current behavior in Figure 6-8(b) is exponential, but does not show the same consistency between runs that the high current behavior does. The overall shape of the characteristic is similar between different runs, but is shifted up/down or left/right. Looking at the low current behavior, it is difficult to identify a particular starting voltage. The thruster starts firing at current levels below the noise floor, and the current tunes continuously with voltage to six orders of magnitude above the noise floor. The emitter that was refilled emitted more current at lower voltages, the first time it was tested.

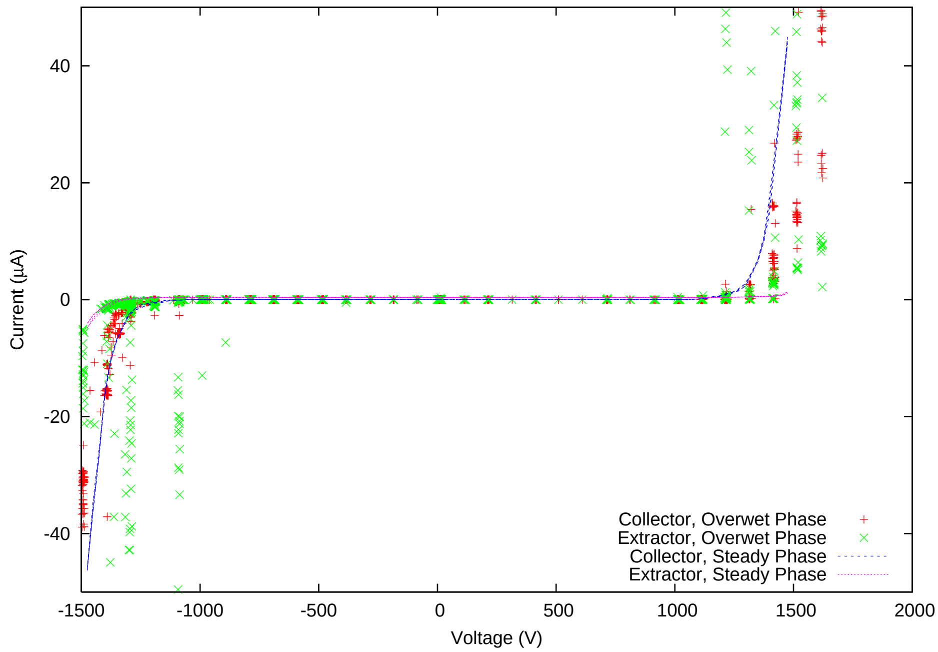

(a) During the first few minutes, the IV characteristic is not very consistent

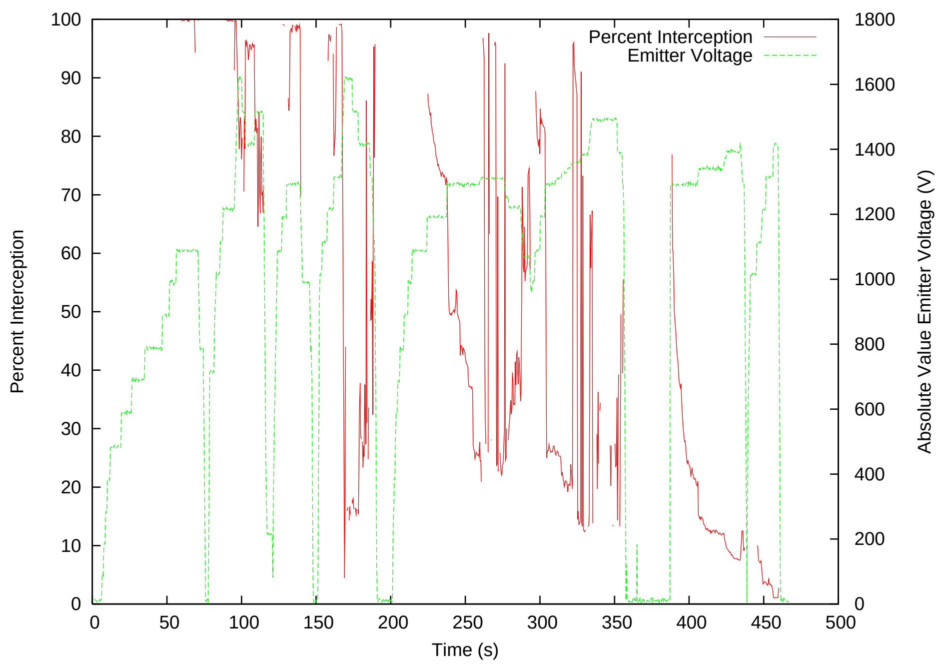

(b) High interception suddenly drops after a few minutes of operation. The voltage was varied during this experiment to monitor changes in the IV characteristic. (Interception is only shown when emitted current is greater than 10 nA to reduce noise.)

Figure 6-7: The first few minutes of thruster operation

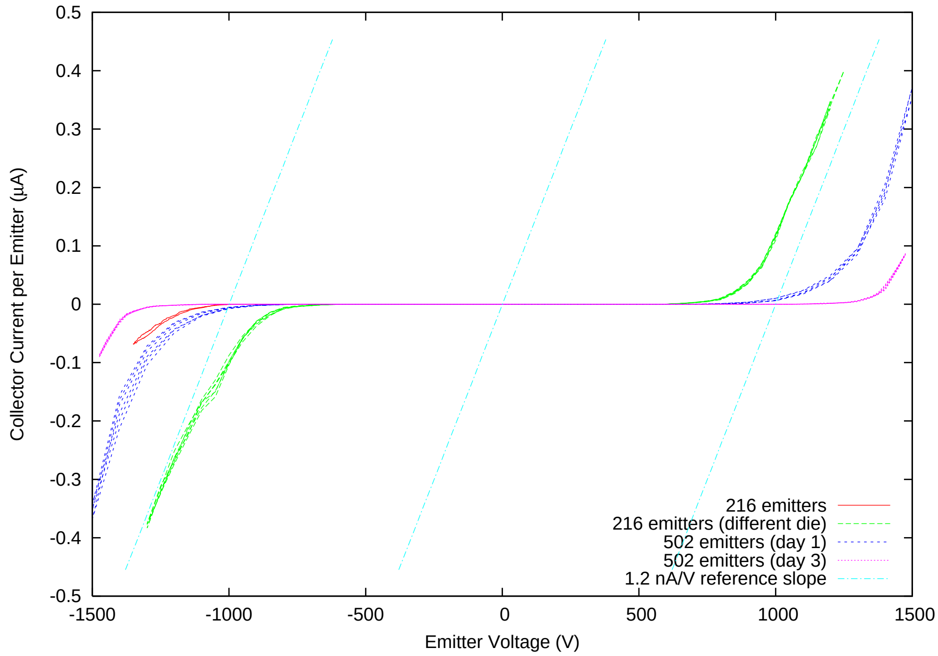

(a) High current behavior

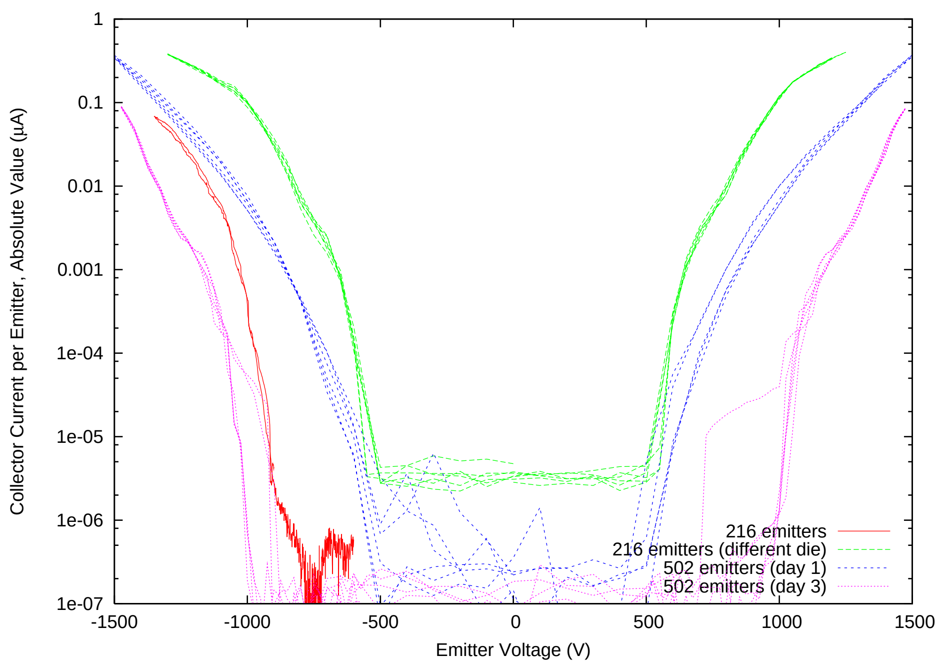

(b) Low current behavior

Figure 6-8: Emitted current for different emitter dies, on different occasions

In Figure 6-9, the current appears to be symmetric to within a factor of 2 with respect to changes in polarity. The amount of asymmetry is not consistent between the different examples shown here.

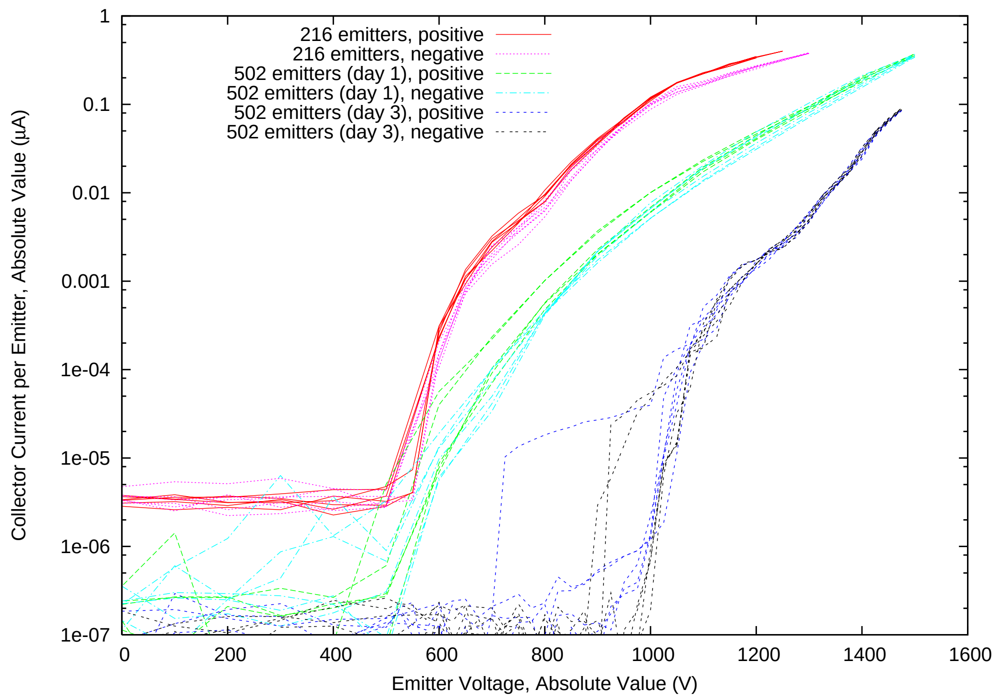

Figure 6-9: Symmetry of the IV characteristic

Figure 6-10 shows the current intercepted by the extractor for the same set of experiments. There is a lot of variation from experiment to experiment. At very low currents, the interception signal is dominated by the noise floor and the capacitive coupling from ramping the emitter voltage. As the emitted current increases beyond the intercepted current noise floor, the proportion of interception drops. It is only when the intercepted current begins to rise above the noise floor that we start to see the real proportion of interception. We observe that the proportion of interception tends to increase with current. It can be well below 1%, even when operation has entered the linear regime, and rarely rises above 10%. For now there is very little consistency of the intercepted current between runs. Effects such as particulate contamination (see Section 6.6), broken emitter dies, different amounts of recessing of the emitters, variations in emitter asymmetry, and the procedure used to wet the emitter die may partly explain these inconsistencies.

The plot labeled "216 emitters" was taken for an emitter die which was broken and not recessed; it exhibits the greatest intercepted current. We do not have sufficient data to determine which of these if either was responsible for the increased emission. This sample also exhibited some hysteresis, with low interception initially, which suddenly jumped up at higher currents. When the voltage was ramped down, the collected current remained high. It seems possible that at higher voltages emission started from an undesirable location on the die, and that once started this source was able to continue to operate even at lower voltages. Once this source was turned off, that sample operated in the low interception regime once again, and the cycle could be repeated.

6.2.4 Depletion Phase

After a long period during which the IV characteristic was very repeatable, the thruster would transition to the depletion phase, in which the IV characteristic slowly shifted towards higher voltages, and the high-current slope would slowly decrease. Figure 6-11 shows successive IV characteristics taken on the 502 emitter (day 1) thruster in the depletion phase. Because the current was constantly decreasing, the different IV traces do not superimpose. Each time it was ramped off, the thruster appeared to partially recover. The sharp decline in current at high voltages in the negative polarity was not consistently observed with other thrusters. Current intercepted by the extractor did not increase during the depletion phase.

During the depletion phase, the light emission from the thruster appeared different than it did during the steady phase. Indeed, significant light could be seen coming from a few different emitters on the thruster, as reported in Section 6.4.4. During subsequent visual inspection of the emitter die, the emitters that glowed more brightly appeared dirty, as described in Section 6.5.

We do not know if the depletion phase is caused by lack of propellant, by fouling of the emitter due to electrochemical reactions, or due to some other cause.

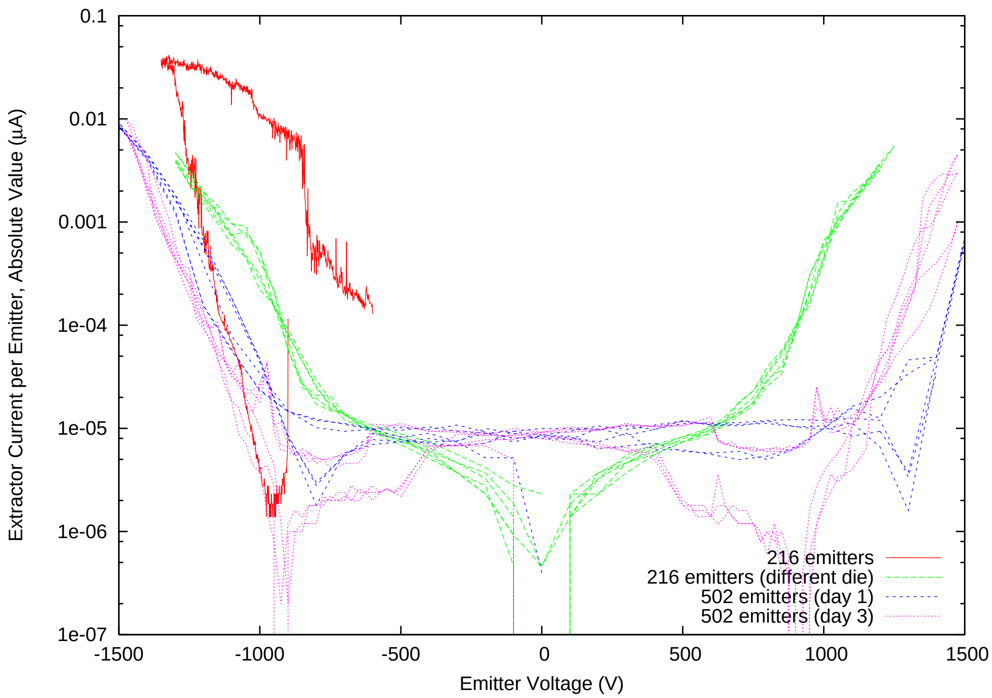

(a)

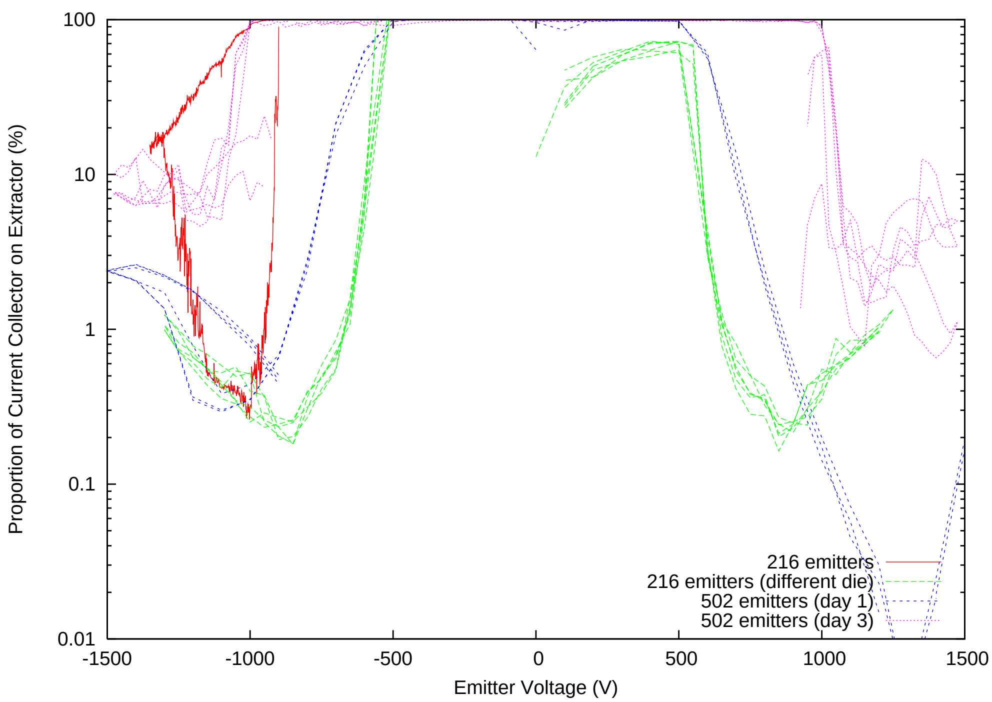

(b)

Figure 6-10: Current intercepted on the extractor. Interception much lower than 1% is possible, even at high currents

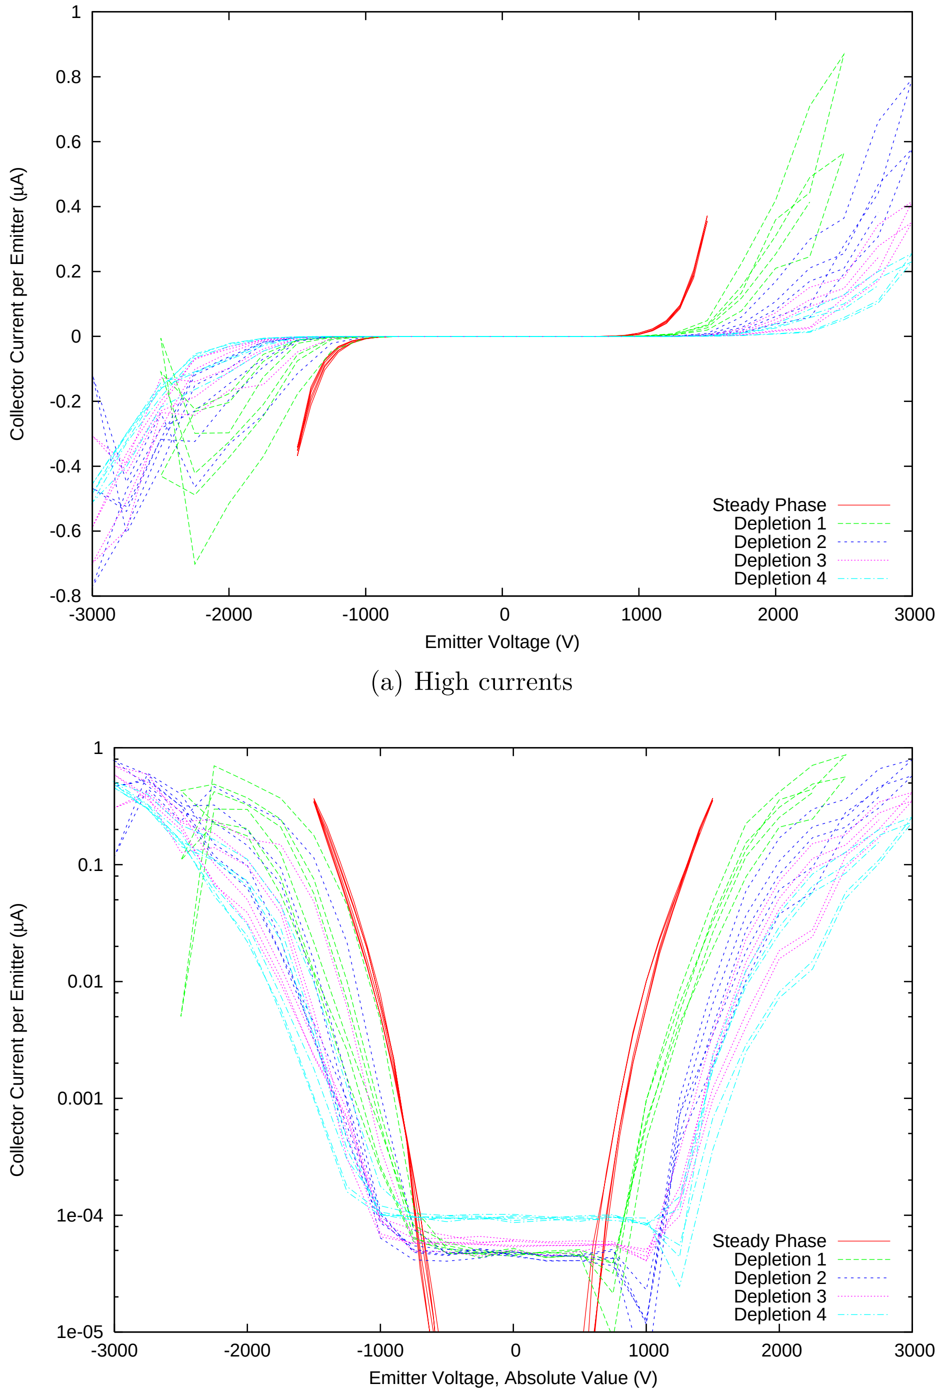

(a) High currents

(b) Low currents

Figure 6-11: In the depletion regime, higher and higher voltages are needed to start the thruster, and the linear regime slope slowly decreases

6.2.5 Single Emitter

In order to understand the IV characteristics for multi-emitter devices, it would be useful to know the characteristic for a single-emitter device. This information would allowed us to distinguish, in the multi-emitter IV characteristics, between the effect of individual emitters increasing their emission, and new emitters beginning to emit.

Unfortunately, we were unable to get a clean IV characteristic from the single emitter dies that were fired. Many reasons can account for this experimental difficulty:

1. Any emissions caused by particles (see Section 6.6) is independent of the number of emitters. Thus emissions that would be lost in the noise with 500 emitters can become dominant when there is a single emitter.

2. The single emitter needles etched faster than all the other needles during fabrication. Therefore, they were the most susceptible to the damage discussed in Section 5.3.

3. With a single emitter, the over-wet phase of operation can be expected to last hundreds of times longer than for dies with hundreds of emitters. We may never have have gotten out of the over-wet phase.

4. Individual emitters may always behave inconsistently. The stability we see with multi-emitter dies may result from the statistical combination of many inconsistent single-emitters. Indeed, there should be more than an order of magnitude noise reduction when moving from a single emitter to 216 or 502 emitters.

6.2.6 Electron Emission

In one case, we observed behavior that was suggestive of electron emission. The IV characteristic in this case (Figure 6-12) was highly asymmetric, with up to 1 mA of current (not captured in this plot[1]) in the negative polarity, at just over 2 kV, versus

[1]For this particular measurement, the chamber was configured for angular dependence measurement, with the collector plate far from the emitters. We estimate that only a third of the emitted current is shown here, based on measurements of the current supplied to the emitters.

only 1 μA of current at 3 kV in the positive polarity. The current intercepted by the extractor in the negative polarity was under 10%. Looking into the vacuum chamber, the usual space glow was not visible, and there was no light at all emanating from the thruster. However, whatever the thruster was pointing at glowed white, particularly if it was insulating.

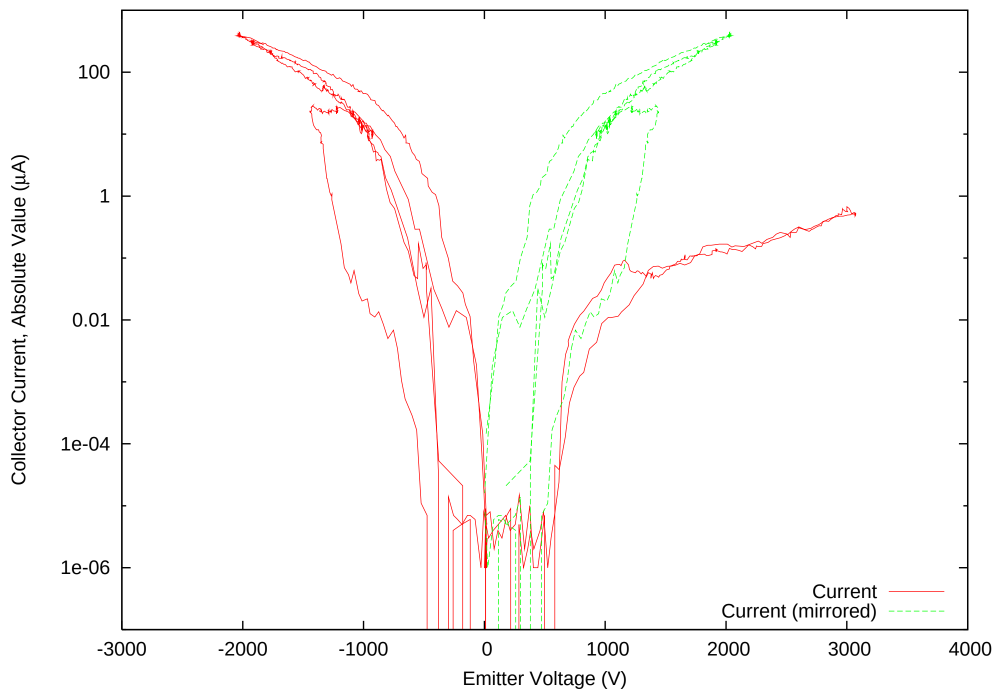

Figure 6-12: Much larger emission in the negative polarity than in the positive polarity may be due to electron emission

We think that this thruster may have been insufficiently fueled, as it behaved as if it was in the depletion phase from the very start of testing. When the thruster was refueled, the electron emission disappeared, and the usual over-wet, steady, depletion sequence was observed.

Figure 6-13 shows the negative part of the IV characteristic plotted in Fowler-Nordheim (FN) coordinates. In these coordinates, the characteristic should be linear if we are witnessing FN emission of electrons, which is characterized by

$$I = aV^2e^{-b/V}, \tag{6.1}$$

where a and b are functions of the effective emission area α (in cm^2), the field enhancement factor β in (cm^{-1}), and the work function ϕ (in eV) [114].

$$a = 1.42 \cdot 10^{-6} \frac{\alpha \beta^2}{\phi} e^{10.4/\sqrt{\phi}} \tag{6.2}$$

$$b = 6.44 \cdot 10^{7} \frac{\phi^{3/2}}{\beta} \tag{6.3}$$

$$\tag{6.4}$$

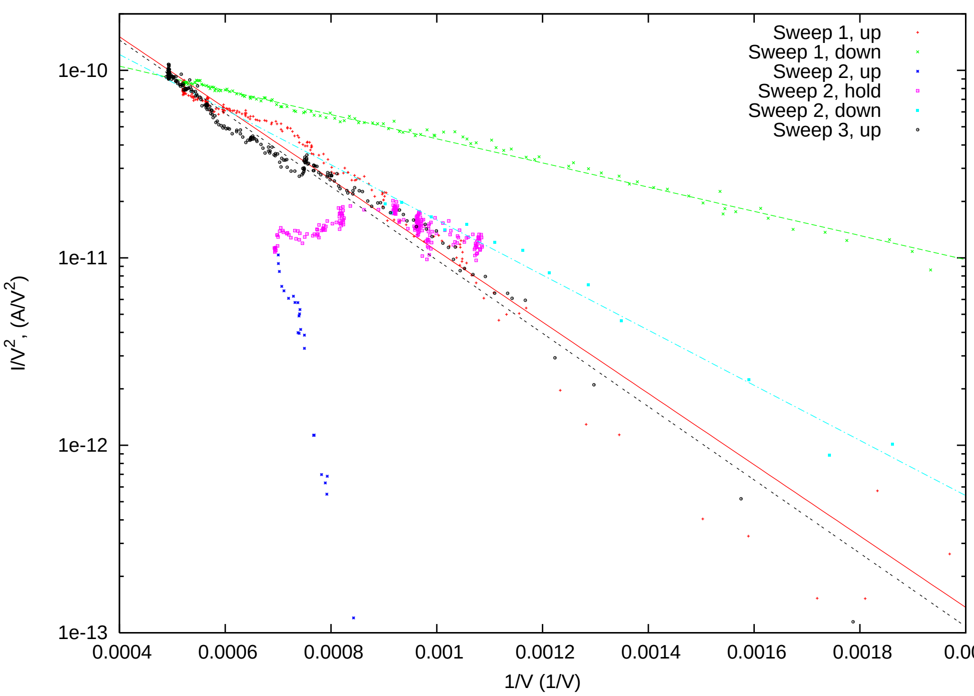

Figure 6-13: Electron emission in Fowler-Nordheim coordinates. For rapid ramps, the characteristic is linear, but it drifts during operation. Fits for the linear portions are provided.

In the FN coordinate plot, different linear characteristics are present for different ramps of the current, suggesting that the emission conditions are not very stable. Table 6.1 shows fits for the different linear portions in Figure 6-13, assuming emission from silicon (ϕ=4.5 eV). The value $1/\beta$ can be used as crude approximation of the radius of curvature of the emission site, and $\sqrt{\alpha/\pi}$ is the radius of a circle with area α. These radius of curvature in the 10 to 100 nm range seem reasonable, and the small emission area suggest that a single emission sites exists, The change in emission parameters with time could be due to motion of liquid on the emitter and/or to erosion of the tip taking place (for example due to background neutrals being ionized and backstreaming to the emitters). The fact that all the fits are nearly concurrent around the maximum operating current, is however, somewhat troubling. Why do the physically relevant parameters α and β change in such a way as to create this concurrence? Is this just a coincidence?

Trace	a	b	β (cm^{-1})	α (cm^2)	$1/\beta$ (nm)	$\sqrt{\alpha/\pi}$ (nm)
1, up	$8.71 \cdot 10^-10$	$4.38 \cdot 10^3$	$1.40 \cdot 10^5$	$1.06 \cdot 10^{-15}$	71	0.18
1, down	$1.91 \cdot 10^-10$	$1.48 \cdot 10^3$	$4.14 \cdot 10^5$	$2.66 \cdot 10^{-17}$	24	0.03
2, down	$4.70 \cdot 10^-10$	$3.38 \cdot 10^3$	$1.82 \cdot 10^5$	$3.40 \cdot 10^{-16}$	55	0.10
3, up	$8.79 \cdot 10^-10$	$4.50 \cdot 10^3$	$1.37 \cdot 10^5$	$1.13 \cdot 10^{-15}$	73	0.19

Table 6.1: Fitted electron emission parameters

6.2.7 Temperature Dependence

In this experiment, the temperature dependence of emission was measured by heating one side of the spring wafer with a 10 Ω resistor, and monitoring the temperature on the other side of the spring wafer using a thermocouple. The experiment was run without a proper feedthrough for the thermocouple, so we expect some inaccuracy in the temperature reading. For some experiments, an Omega CN8200-R temperature controller was used to stabilize the temperature. Experiments were run between 30 and 70°C. At 70°C the resistor was at its power limit, and creep of the adhesives that had been used to attach the resistor to the thruster caused the resistor to partially detach from the thruster.

Two different sets of experiments were run, with the same 216 emitter die. In the *static* experiments, the temperature controller was set to a temperature, and left to settle. Then an IV was taken. This operation was repeated at various temperatures, giving high resolution in voltage and low resolution in temperature. Current measurements in this case were repeatable allowing for about 2°C temperature variation. In the *dynamic* experiments, the heater was turned on, and quick 15 second IV sweeps were made, giving high temperature resolution, but poor voltage resolution. Moreover, since measurements were taken as the temperature was rising or falling, there was a clear lag between the temperature indicated by the thermocouple and the temperature affecting the current, which reached up to about 5°C in these experiments.

Figure 6-14 shows the dependence of the emitted current on temperature for the dynamic experiment. At low currents, the current increased by a factor of about 20 over the temperature range, which, if this were an Arrhenius process, corresponds to an activation temperature of 7700°K (0.67 eV). The increase at high currents was by a much smaller factor of 2 to 3. Between 30 and 70°C, the viscosity of EMI-BF_4 decreases by a factor of 3.1 [74], so our high current data are compatible with current being limited by viscosity.

The IV characteristics from the static experiment are shown in Figure 6-15(a). At 70°C, we observe a sharp decline in current. This decline was observed in all the experiments that were run, though it was not always this sharp (even Figure 6-14 shows this decline). In one case, after observing the decline at 70°C, we repeated the measurements that had been taken on the way up to 70°C on the way back to room temperature. Figure 6-15(b) shows one of these plots, which shows that the downward pass is consistent with the upward pass. The inconsistency of observed results at 70°C may be due a strong sensitivity to temperature at that temperature. For now we have no explanation of what is causing the 70°C current drop, more data would be needed.

6.2.8 Leakage Current

Figure 6-16 shows the temperature dependence of the current between the emitter and extractor electrodes as a function of temperature for the dynamic temperature dependence measurement (see Section 6.2.7). Two exponential trends are clearly visible. Assuming Arrhenius processes, the activation temperature is 10000°K (0.86 eV) at low temperature and 21500°K (1.85 eV) at high temperature.

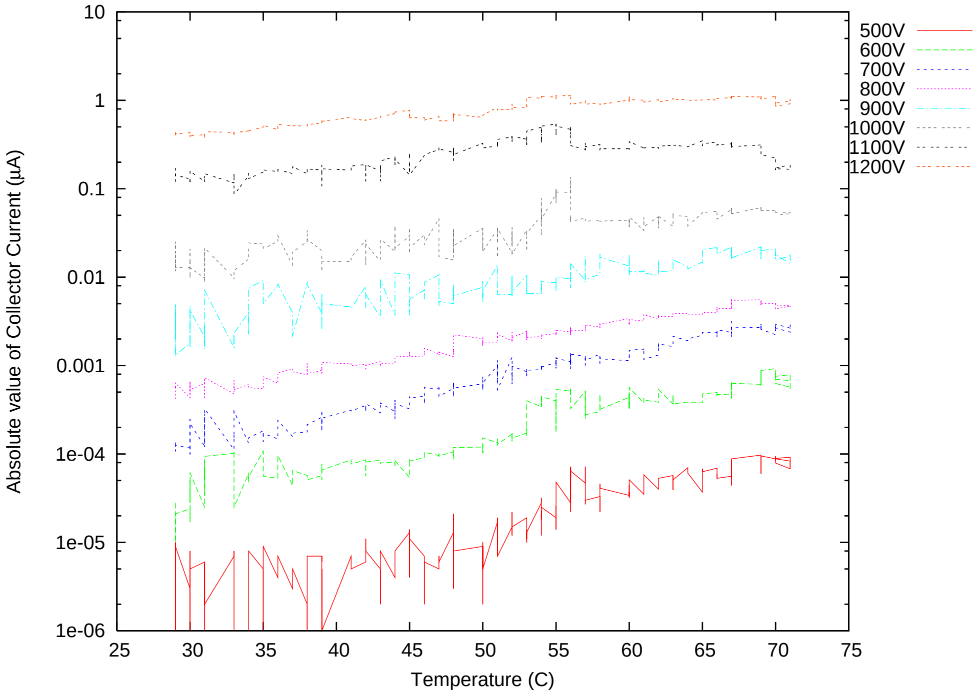

(a) Temperature dependence

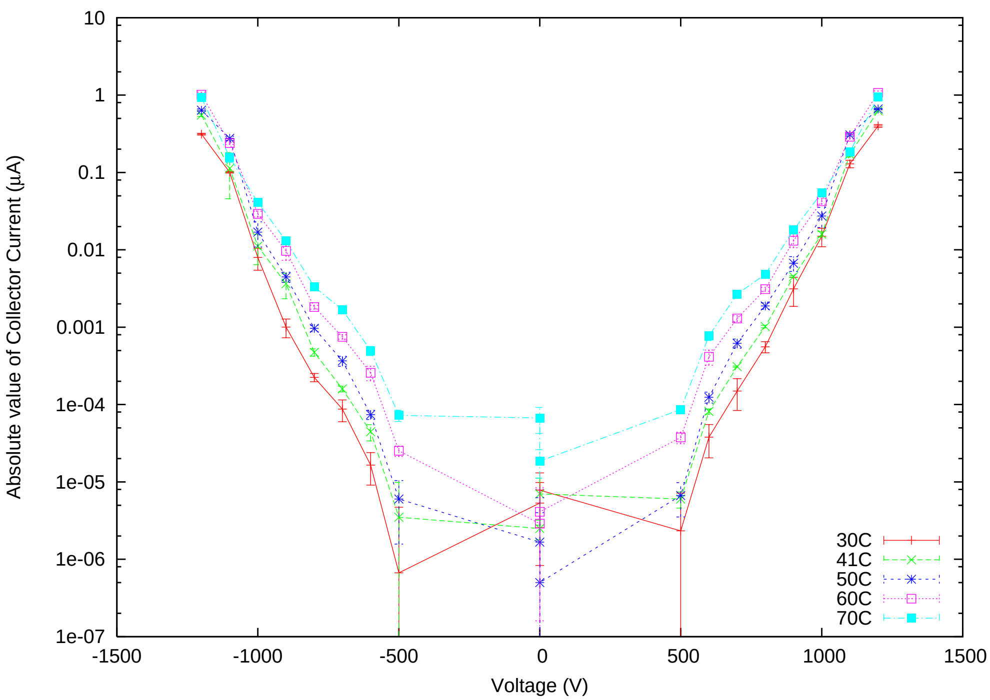

(b) IV characteristics

Figure 6-14: Current dependence on temperature for the dynamic experiment

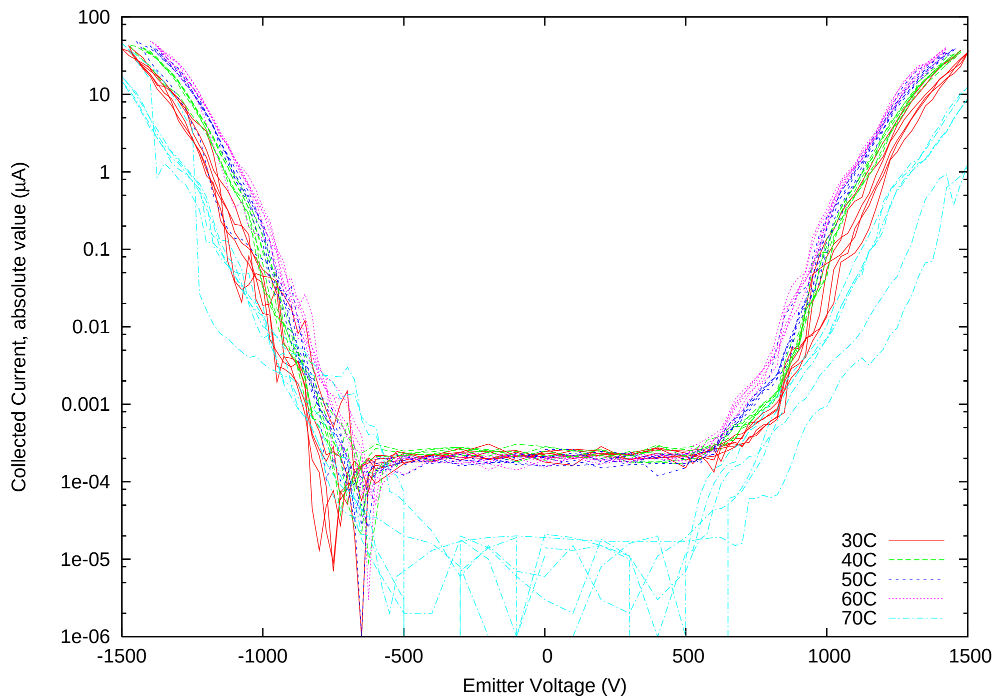

(a) IV characteristics

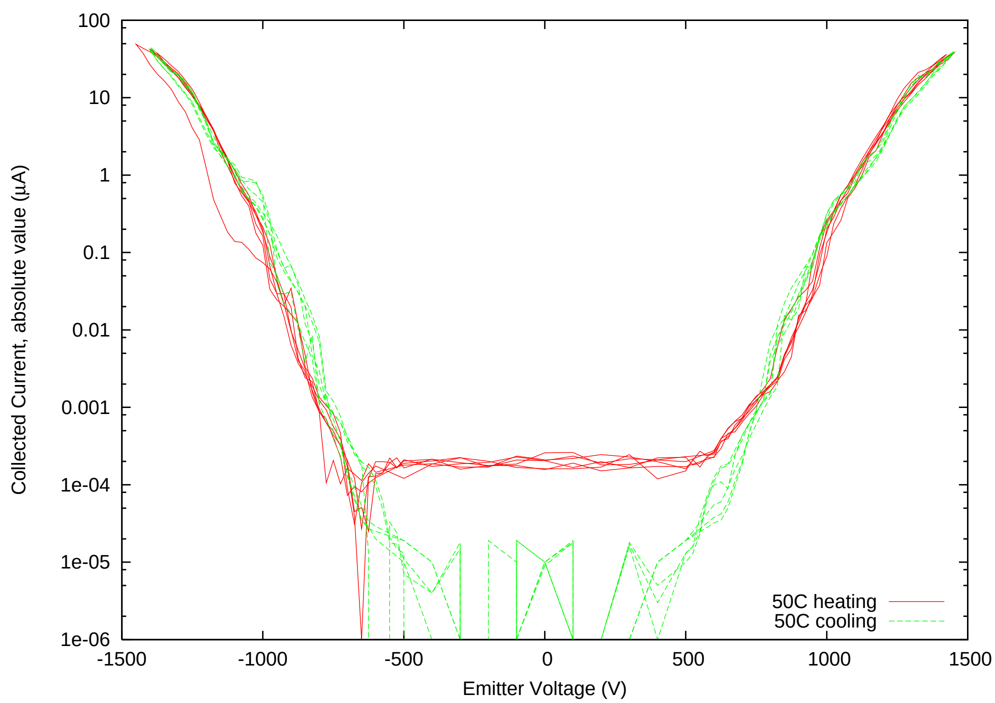

(b) Comparing IV characteristics before and after firing at 70°C

Figure 6-15: Current dependence on temperature for the static experiment

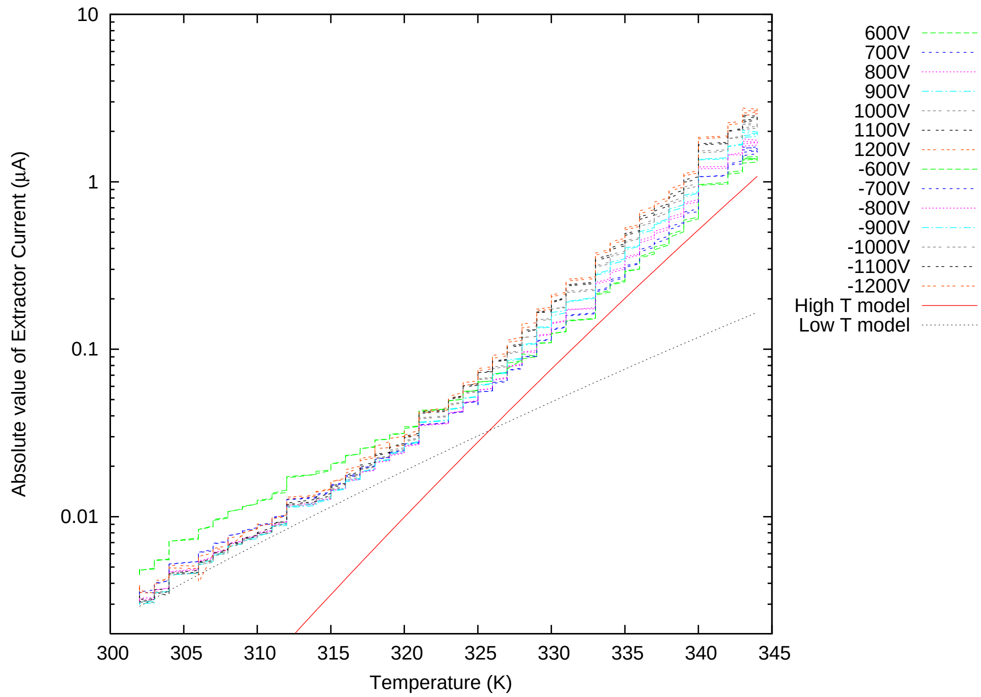

Figure 6-16: Exponential dependence of extractor current on temperature

In Figure 6-17(a) the resistance has been plotted instead of the extractor current. At high temperature, the resistance is independent of the voltage, so an ohmic process is dominating the current to the extractor, while the dominant process at low temperature has a current that increases more than linearly with temperature. The high temperature behavior is consistent with with ionic conduction in the Pyrex insulator. The origin of the low temperature extractor current is currently unknown, and could be beam interception on the extractor.

6.2.9 Angular Dependence

The angular dependence of the emission was measured in the experiment depicted in Figure 6-18. In this experiment, the collector plate was placed farther from the thruster to allow the thruster to be rotated around its axis. Note that because the extractor is slotted, the thruster does not have symmetry of revolution. Thus this experiment should be repeated with many different thruster orientations to fully characterize the angular dependence of the beam. Figure 6-19 shows the angular

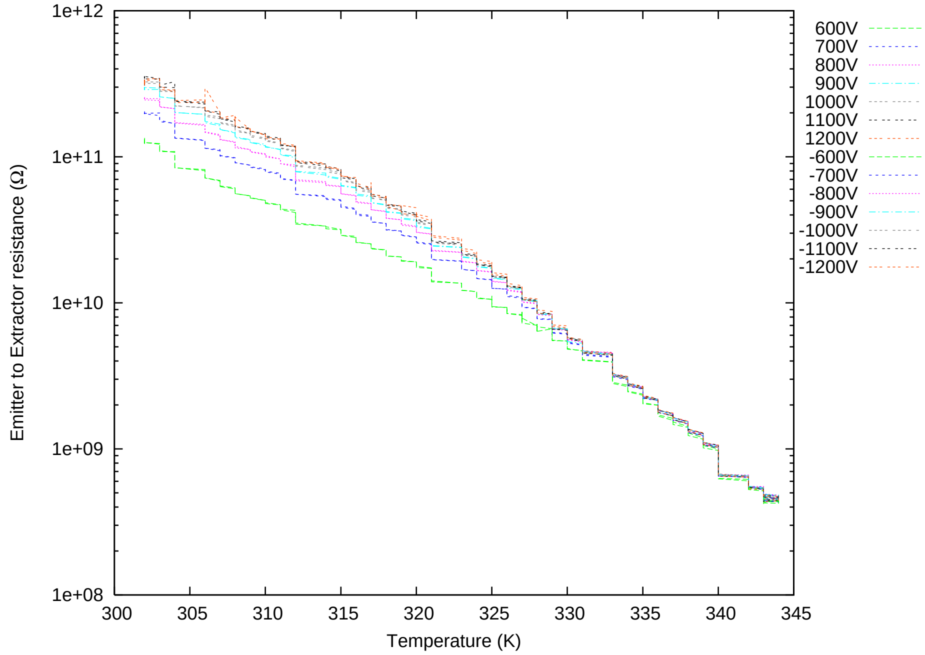

(a) Resistance between extractor and emitters as a function of temperature

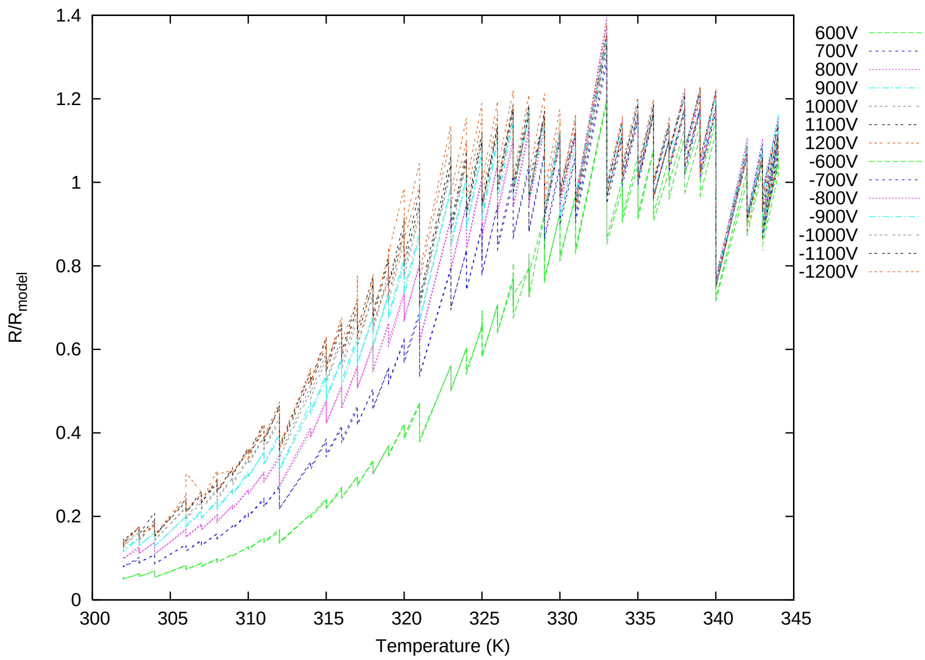

(b) At high temperature, resistance is given by $R_{\text{model}}(T) = 3 \cdot 10^{-19}\, e^{-\frac{21500}{T}}$

Figure 6-17: Leakage is ohmic at high temperature dependence of the collected current.

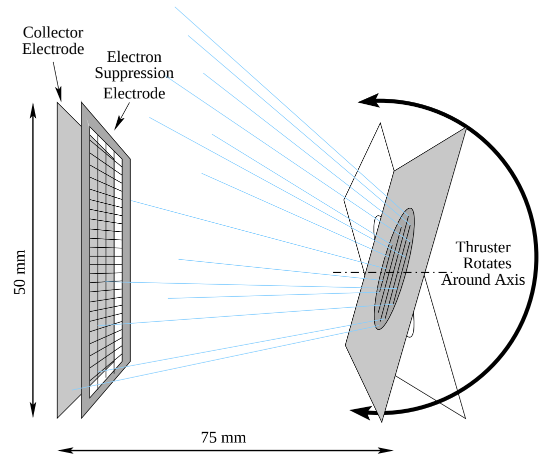

Figure 6-18: The setup to measure the beam angular dependence

Because the collector plate has a non-negligible angular dimension $\Delta\theta$, further processing is needed to reconstruct the angular divergence of the emitted beam. Indeed, the collected current $I_{\mathrm{c}}(\theta)$ is a convolution between the angular beam current density $i_{\mathrm{b}}(\theta)$ and the shape of the collector:

$$I_{\mathrm{c}}(\theta) = \int_0^{\Delta\theta} i_{\mathrm{b}}(\theta)\mathrm{d}\theta \tag{6.5}$$

The current $I_{\mathrm{c}}^{\infty}(\theta)$ for a semi-infinite collector can be inferred from the current $I_{\mathrm{c}}(\theta)$ for a finite collector by:

$$I_{\mathrm{c}}^{\infty}(\theta) = \int_0^{\infty} i_{\mathrm{b}}(\theta)\mathrm{d}\theta = \sum_{k=0}^{\infty} I_{\mathrm{c}}(\theta - k\Delta\theta) \tag{6.6}$$

This expression allows $I_{\mathrm{c}}^{\infty}(\theta)$ to be directly computed from the experimental data for $I_{\mathrm{c}}(\theta)$. The corresponding plots are shown if Figure 6-20. The angular dimension of the beam can now be deduced by looking at the range of angles over which the collected current rises between two arbitrary proportions of the beam current. For example, the half width at half maximum (HWHM) corresponds to the angle over which the current rises from 25% to 50% of its final value.

Doubling the rotation it takes for $I_{\mathrm{c}}^{\infty}(\theta)$ to go from 25% to 50%, we find an HWHM

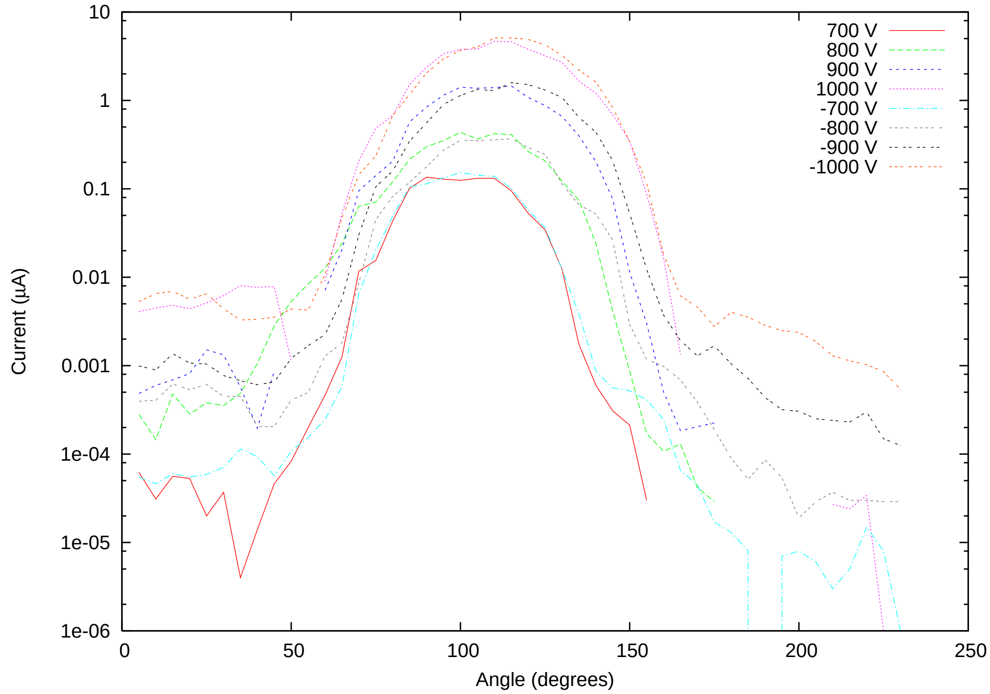

Figure 6-19: Current as a function of thruster orientation

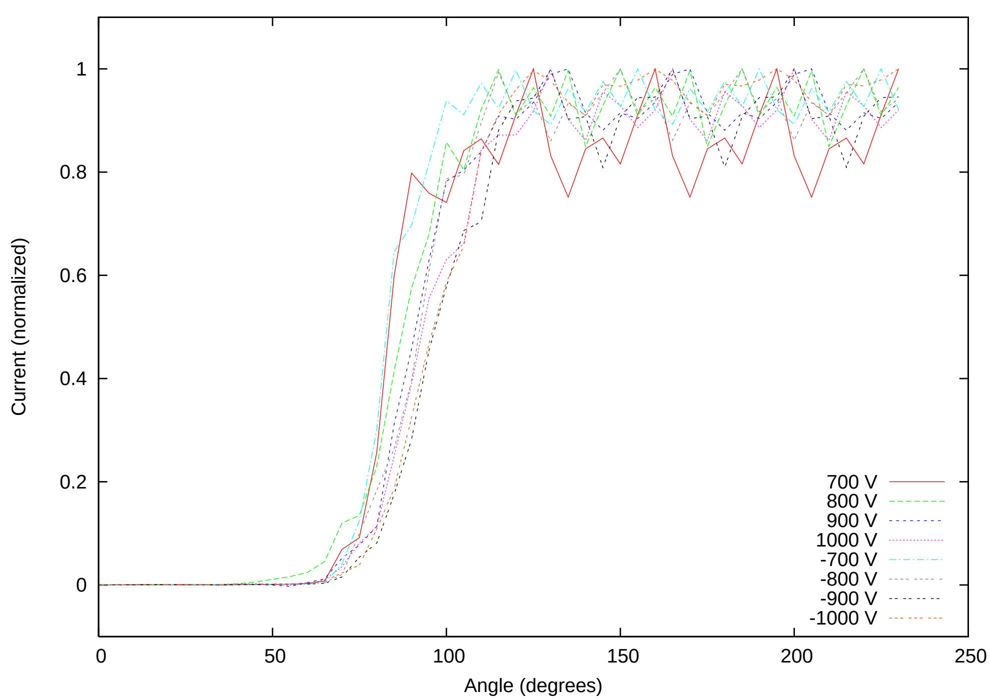

Figure 6-20: Current as a function of thruster orientation, extrapolated to a semi-infinite collector plate

of 10° to 15°. Similarly, 90% of the current is contained within a half-angle of 20° to 25° (the range over which the current goes from 5% to 50%).

6.2.10 Startup Transient

All the data we have taken so far ignore any startup transients. However, there can be rich information present in the startup transient of an electrospray source [115]. In our setup, we observed the startup transient by measuring the voltage across the collector resistor with the oscilloscope. In a typical startup transient, the collected current rises over a sub-millisecond time-scale. Then the current exponentially decays towards a steady state value, with a time scale on the order of one second. At low currents, the amount of overshoot is nearly non-existent. At higher currents, the initial current peak can be over twice the steady-state current. As the current level increases, the duration of the transient slowly decreases. The state of the system does not immediately reinitialize when the array is turned off. For example, if the array is turned off for 100 ms, there is a no overshoot at all when it is turned back on.

We managed to collect a set of startup transients with a thruster that was in the depletion phase. One batch of startup transients is shown in Figure 6-21 superimposed with their exponential fits of the form

$$i(t) = i_0 + i_1 e^{-\frac{t}{\tau}}, \tag{6.7}$$

where i_0 is the steady-state current, i_1 is the overshoot current, and τ is the characteristic decay time. To collect this data, we turned the array on for 5 seconds to measure current, then turned it off for 30 seconds to let the thruster return to its initial state.

Figure 6-22 shows the parameters that have been extracted from the collected data. Each voltage was measured five times; average and standard deviation appear in the plots. The steady-state current seems to vary linearly with the voltage, while the overshoot increases more than linearly from near 0 at the lowest plotted setting. If these data had been taken in the steady phase of operation, it would have been

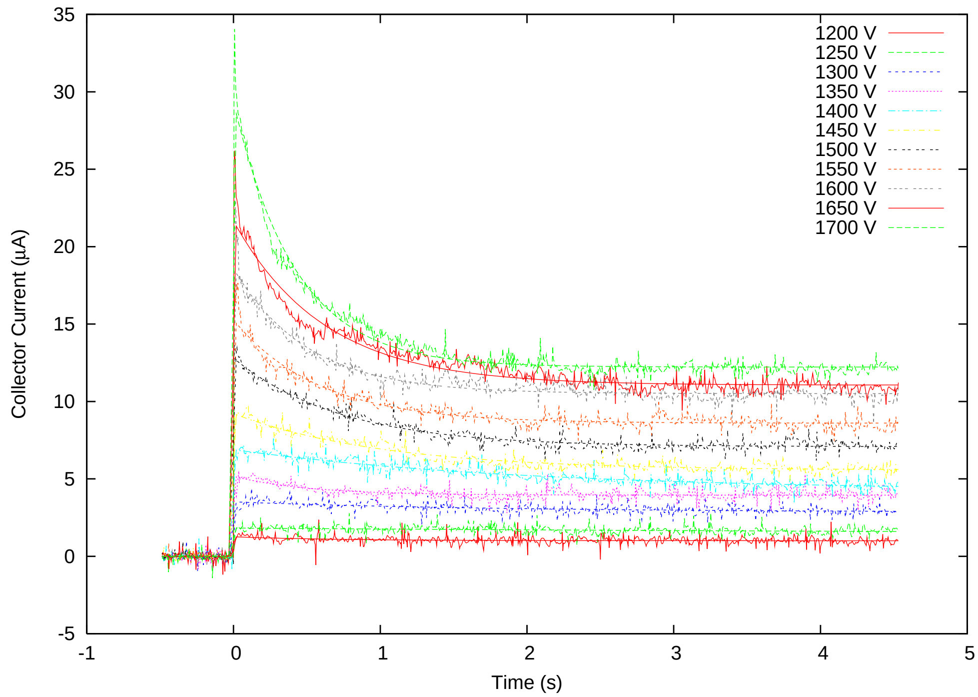

Figure 6-21: Startup transients for different voltages

interesting to see if $i_0 + i_1$ continues the exponential progression of the low current IV characteristic.

An interesting experiment that is left for future work is to take startup transients after the thruster has been briefly turned off for a variable amount of time. This experiment would allow the thruster recovery phase to be probed. For now all we know is that it takes many seconds for the thruster to recover.

6.3 Time of Flight Measurements

Specific impulse is determined by how fast the particles emitted from the thruster are moving. We now present time-of-flight measurements which directly measure this velocity.

To measure time of flight, we use the apparatus depicted in Figure 6-23, which was built by Lozano for his PhD work [23]. The thruster is fired as in previous experiments, but now the collector plate has been placed much farther from the

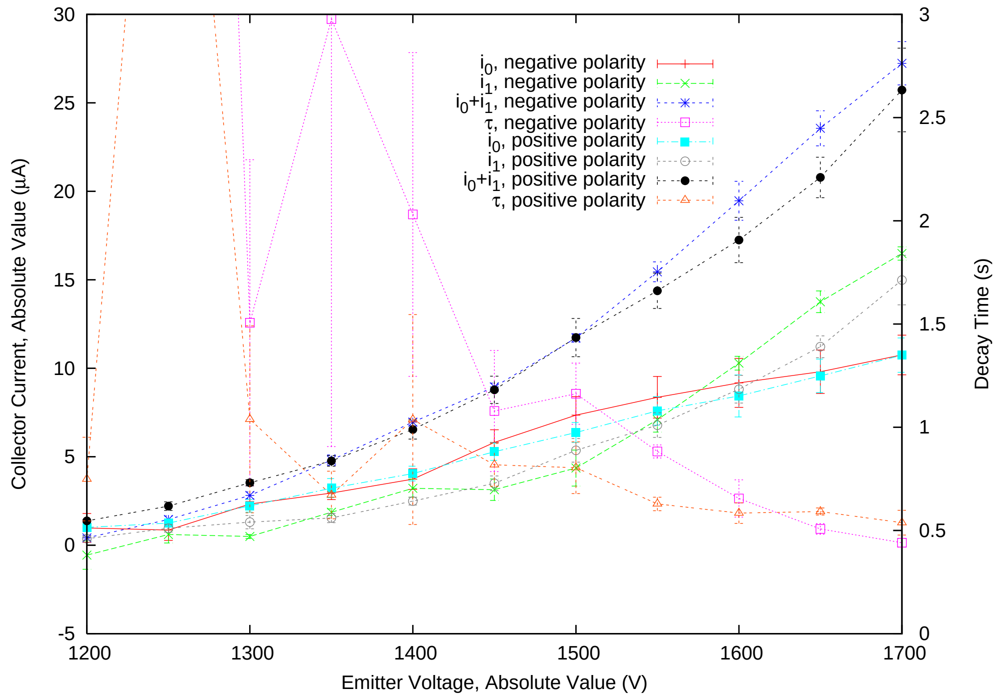

Figure 6-22: Fitted parameters in Equation (6.7) for the startup transients

thruster. The beam from the thruster is passed through a 1/2" aperture, and is collimated by an Einzel electrostatic lens. It then passes through a gate which can be used to suddenly interrupt the beam. The gate consists of two interleaved sets of parallel wires. When the gate is open, the wires are all grounded, and, except for direct collisions, the beam particles are unaffected by the wires and proceed about 80 cm to the collector plate. When the gate is closed, the sets of wires are at opposite voltages up to 950 V. A few centimeters away from the gate, the field created by the wires is imperceptible, but particles close to the wires undergo a strong deflection which prevents them from arriving at the collector plate.

Time-of-flight measurements are taken by switching the gate, and looking at the current transient on the collector plate. For example, if the gate is suddenly opened, particles arriving at the gate will suddenly be able to pass. Depending on their velocity, they will reach the collector electrode after varying amounts of time, and the current detected on the collector electrode will increase as slower and slower particles begin to arrive. When the gate is suddenly closed, the opposite phenomenon occurs:

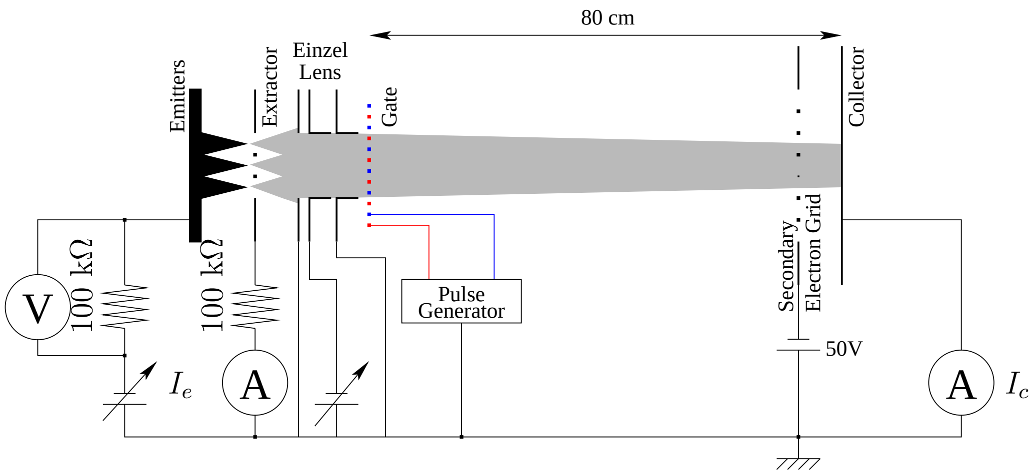

Figure 6-23: Diagram of the time of flight apparatus

first the fast particles are lost from the collected current, then the slower ones. The resulting transient is inverted compared with the transient when the gate is opened.

Signal quality can be greatly improved by averaging the transient over multiple gate openings or closings to reduce noise. The data we present here are the average of 1024 events.

The data have also been preprocessed to translate the time of flight into the mass of the corresponding particle assuming that the energy of the particle is eV, where e is the unit charge, and V is the voltage applied to the emitters (neglecting any resistive voltage drop). In this case, an energy balance yields

$$\frac{1}{2}m\left(\frac{l}{t}\right)^2 = eV, \tag{6.8}$$

where l is the gate to collector distance, and t is the time of flight. This equation can be rewritten as

$$m = \frac{2eVt^2}{l^2}. \tag{6.9}$$

6.3.1 EMI-BF_4

Figure 6-24 shows a typical time-of-flight plot for EMI-BF_4, which proves that the thruster is operating in the pure ion regime. The long exposure plot shows that there are no heavy particles present in the beam, while the short exposure plot shows the details of the particles that are present. Particle masses from Table 6.2 are also indicated in the plots. The monomers EMI and BF_4 are clearly visible, as are the dimers (BF_4-EMI)BF_4^- and (BF_4-EMI)EMI^+, as well as the trimer (BF_4-EMI)$_2BF_4^-$. The proportions of the monomer, dimer and trimer were consistent across the various measurements we made to within 5 to 10%.

Ion	Mass (AMU)
EMI^+	111.2
Im^+	280.2
BF_4^-	86.8

Table 6.2: Masses for EMI^+, BF_4^- and Im^+

A slightly modified pulsed time-of-flight experiment [23] can be used to more easily distinguish the different components in the beam. In this experiment, the gate is opened for just 1 μs. Thus, each particle species now produces a pulse in the time-of-flight plot instead of producing an edge. Figure 6-25 shows the resulting plot[2], which is a smoothed average of the derivative of Figure 6-24. It is now easier to read the peak value of the pulses, which are correct to within a few AMU.

6.3.2 EMI-Im

EMI-Im is a promising ionic liquid for electrospray propulsion. It has lower surface tension than EMI-BF_4, which should lead to lower starting voltages. More importantly, EMI-Im does not contain fluorine, so there is no chance of electrochemical decomposition products damaging the silicon emitters. However, the pure ionic regime is more difficult to reach with EMI-Im than with EMI-BF_4.

Figure 6-26 shows a typical time-of-flight plot for EMI-Im. Once again, pure ion emission is observed, though this time trimers and even tetramers were visible. There was one exception just after we started taking the time-of-flight data. One

[2] In producing this plot, the data were shifted 0.5 μs earlier to account for the length of time the gate was opened.

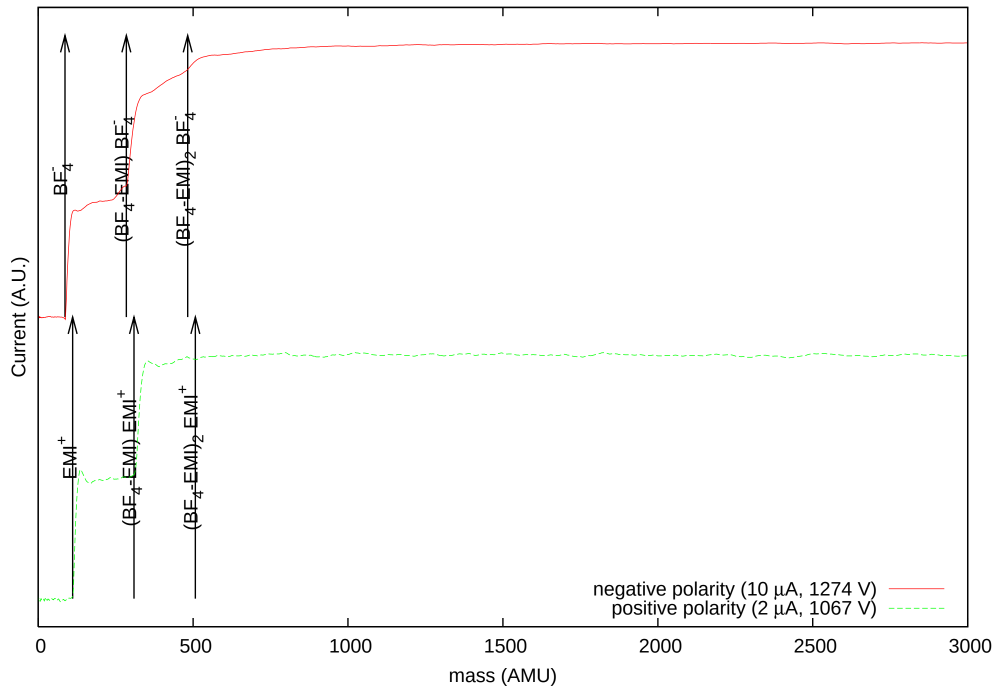

(a) Long exposure

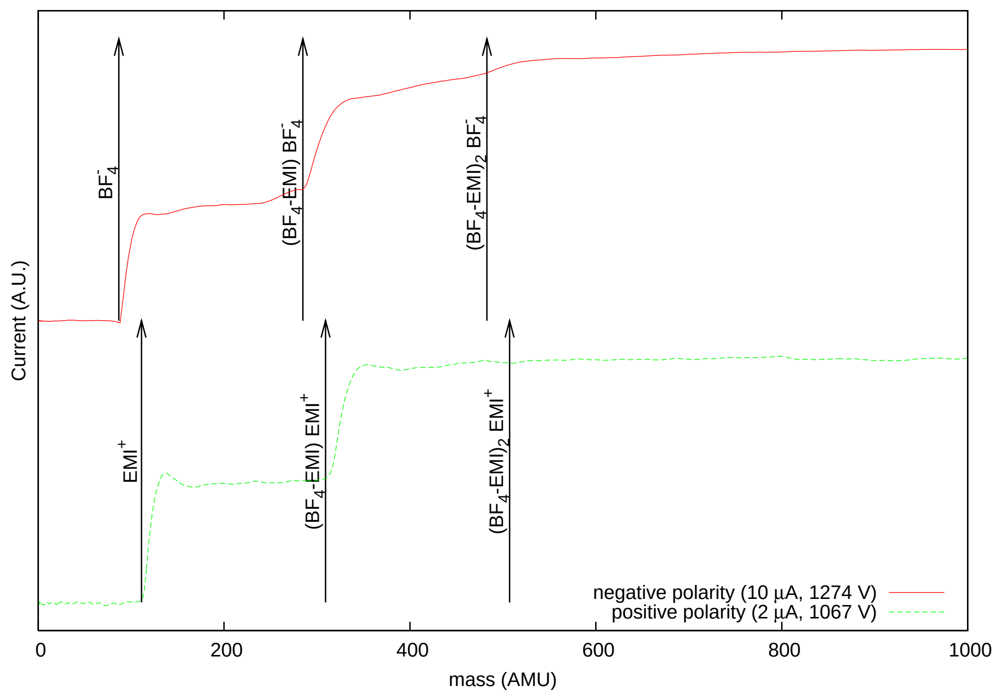

(b) Short exposure

Figure 6-24: Typical time-of-flight measurement with EMI-BF_4

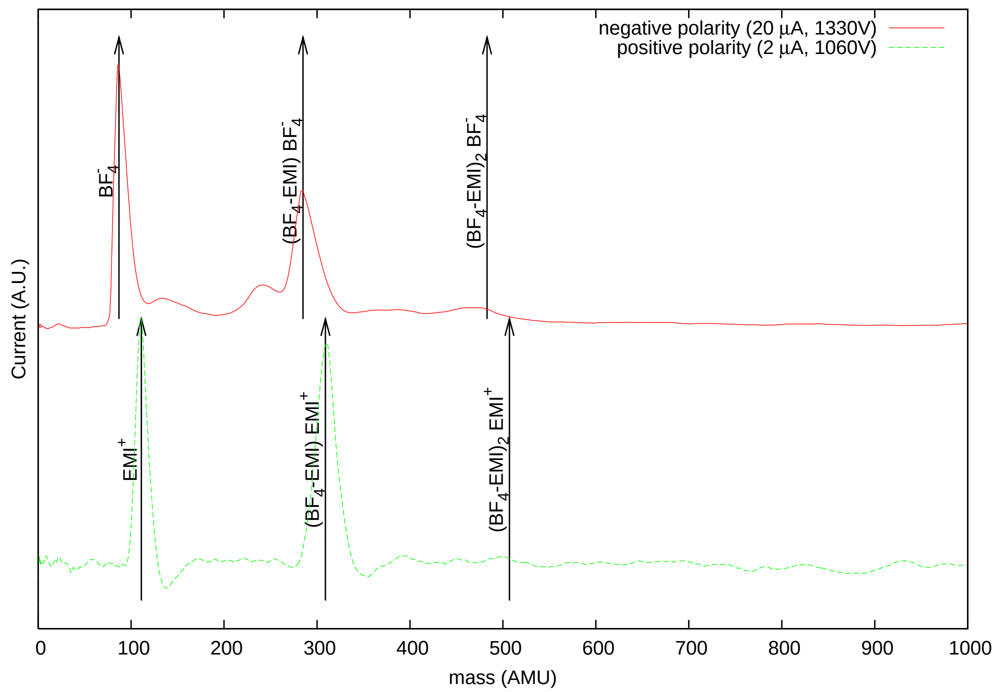

Figure 6-25: Typical pulsed time-of-flight measurement with EMI-BF_4

of the plots in Figure 6-26(a) shows a much longer tail that could indicate a small proportion of droplets. Except for this initial experiment, none of the subsequent data show any sign of droplets, despite covering a wide range of emission currents. Perhaps there is significant droplet emission in the over-wet phase, and when this particular time-of-flight plot was taken, the thruster hadn't completely transitioned to the steady phase. In any case, these results show that it is possible to use EMI-IM in the pure ion emission regime, confirming that EMI-IM is an excellent candidate as an electrospray thruster propellant.

6.4 Emission Imprints

In order to determine if all the emitters are firing, we conducted a series of experiments in which an emitter die was placed in front of a conductive plate, and fired. The experimental setup is shown in Figure 6-27.

Two types of conductive plates were used: polished silicon and Indium-Tin Oxide

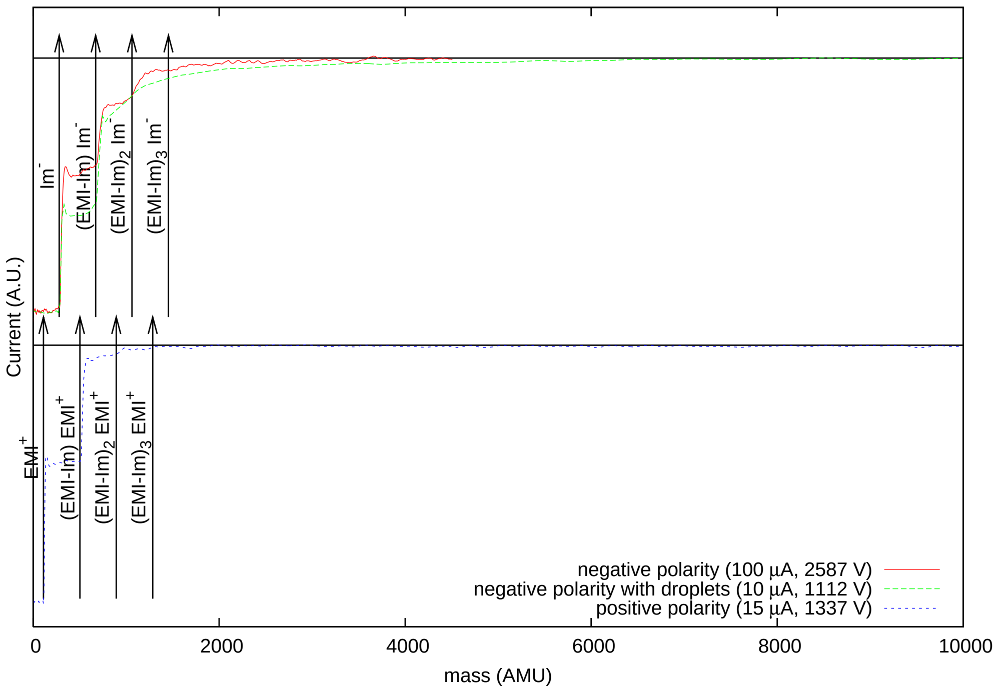

(a) Long exposure

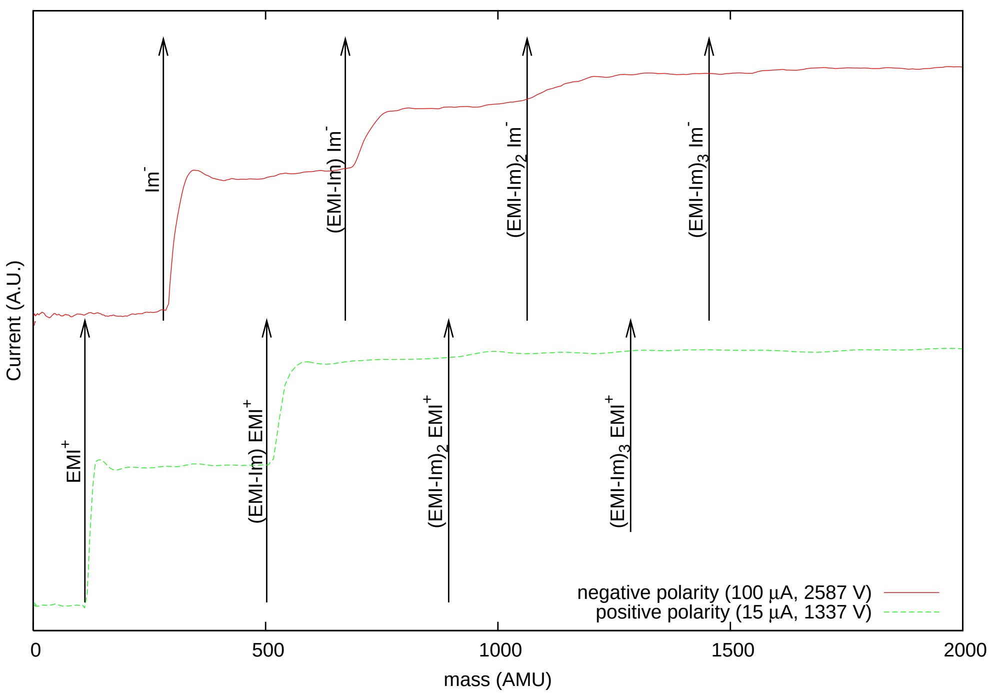

(b) Short exposure

Figure 6-26: Typical time-of-flight measurement with EMI-Im

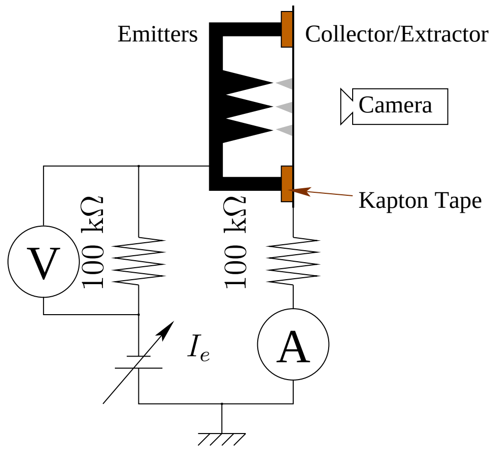

Figure 6-27: Firing against a plate to determine the cumulative emission imprint

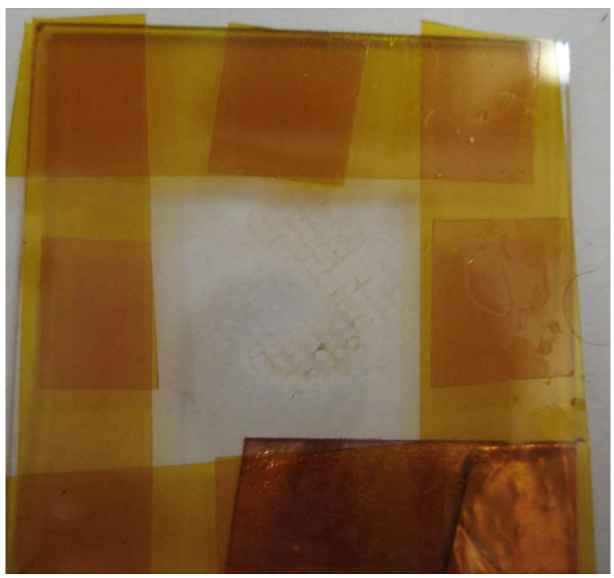

Figure 6-28: Kapton square around the target area to serve as insulation between the target and the emitter die. The copper tape which served to contact the ITO is also visible at the bottom

(ITO)-coated glass. The polished silicon was used because of its availability and smoothness, and the ITO-coated glass was used because it is transparent, but has a conductive layer. Thus, the progress of the experiment could be monitored through the target.

To set up the experiment, 300 μm thick Kapton tape was first applied to the target plate, in such a way as to form an insulating square around the target area (see Figure 6-28). Because the corners of the square were covered by two layers of tape, an extra layer of tape was then sometimes added to the edges of the square so that the thickness would be uniform. We always used the same 216-emitter die for these experiments. The wetted emitter die was placed against the plate, such that the outer wall of the emitter recess was resting on the Kapton. The assembly was then clipped to the usual polyethylene holders (see Figure 6-1). For the experiments with ITO, the clip was placed near the edge of the emitter die to leave a clear line of sight for the camera. Unfortunately, this configuration led to poor clamping, and the emitter die was visibly not well clamped to the Kapton tape on the side opposite the clip.

In this setup, there was nothing to preventing liquid from being drawn up the sidewalls of the emitter die, across the Kapton and up to the target. Nevertheless, we only experienced one short circuit with this configuration, suggesting that it is good enough to produce interesting data.

6.4.1 Silicon Target, High Currents

In this test, the thruster behaved very stably in the positive polarity, but had inconsistent behavior in the negative polarity. Voltages up to 2500 V were applied, and currents up to 70 μA were observed.

Figure 6-29 shows the target upon inspection. Thirteen rows are clearly visible, corresponding to the thirteen rows of emitters. In one corner of the target imprint a black/orange substance is visible. The rows cut through this substance. Microscopic examination reveals that each row is made up of erratic patterns spaced 500 μm apart, and which presumably correspond to each one of the emitters. These erratic patterns have a roughly threefold symmetry, suggestive of the threefold symmetry of the ridge-lines on the emitters.

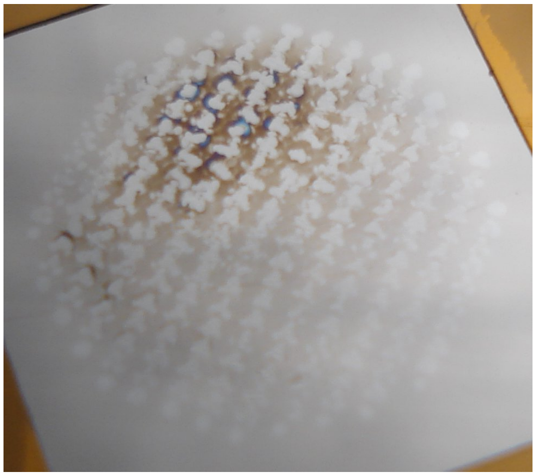

(a) All thirteen rows of emitters fired along their whole length

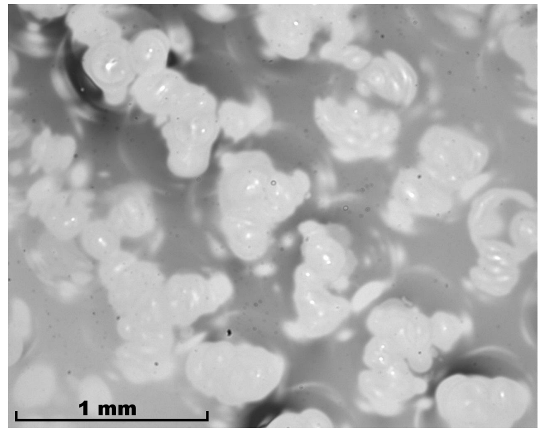

(b) Microscope image of the target, showing an array of imprints with vaguely threefold symmetry

Figure 6-29: The first silicon target after firing

6.4.2 Silicon Target, Low Currents

In this run, we wanted to explore whether the erratic patterns that were previously observed were also present during low current operation. Therefore, we started the thruster, increasing current up to a few hundred nanoamps, and let it fire for about an hour. Then, we removed the target for observation.

This time only nine incomplete rows are visible, suggesting that not all the emitters fired (see Figure 6-30). The emitters that did not fire may have been less sharp than those that did fire, or the emitter die may have been slightly tilted relative to the silicon plate, causing a disparity in electric field between emitters. In any case, in the regions that did fire, all the emitters were firing, as attested to by neat rows of circles, 250 to 350 μm in diameter (see Figure 6-30). Many of the circles contain darker dots up to 200 μm in diameter, and each circle is surrounded by a roughly circular halo 500 μm in diameter.

In this experiment, a single 300 μm layer of tape was used, and the emitters were an estimated 150 μm below the top surface of the emitter die, so the total emitter to plate distance was 450 μm. Table 6.3 summarizes the emission half-angles that would produce the different circles that were observed (i.e., the half angle of the cone with

Circle	Radius (μm)	Emission half-angle (°)
Inner Orange Circle (max)	200	12
Clean Circle (small)	250	15
Clean Circle (large)	350	21
Outer Halo	500	29
Pencils in [37]		9.5
Volcanoes in [37]		14.1

Table 6.3: Approximate feature sizes on the second silicon target

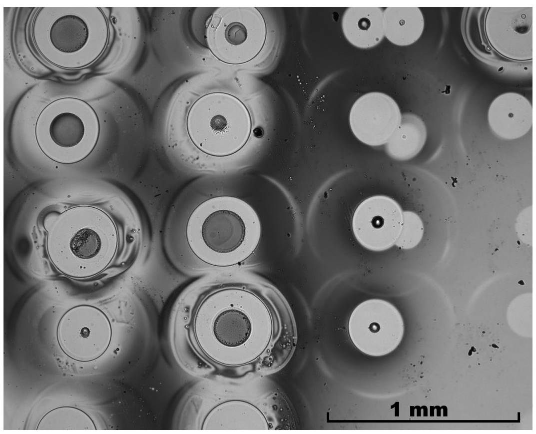

(a) Microscope image showing clean circular marks on the target

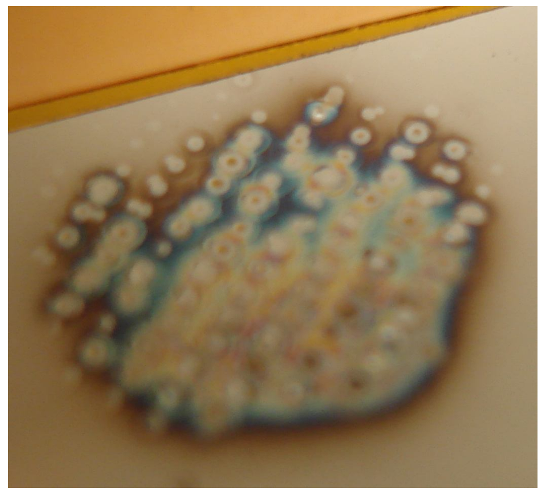

(b) Only nine rows are visible. (The Kapton tape has been moved since the thruster was fired.)

Figure 6-30: The second silicon target after firing

its tip at the emitter tip and which passes through the circle).

To try to elucidate the nature of the various circular features we were observing, we used SEM observation and Energy Dispersive X-Ray (EDX) spectroscopy of the target. Figure 6-31 shows some of the observations. First, we note that the central spot appears to have some topography on the micron scale. The area outside the clean circle did not have this topography (not shown). Next, in EDX spectroscopy, an X-ray detector monitors the X-rays being emitted by a sample under observation in the SEM [116]. These X-rays are produced by the atomic nuclei, excited by the high energy electron beam of the SEM. The energy of the X-rays is characteristic of the atom that emitted them. By correlating X-ray arrival with the SEM scanning, an image of the atoms within about 1 μm of the surface of the sample can be obtained. EDX analysis reveals Carbon and Fluorine in the central spot, but no silicon, confirming that the spot has micron-scale thickness. Around the central spot is an area with pure silicon, and further out a weak fluorine and carbon signal reappears.

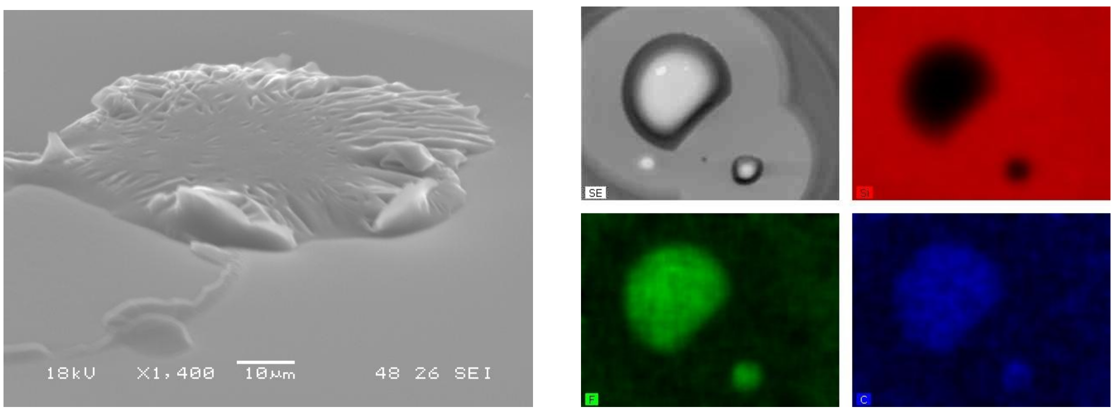

(a) The central dark spot is not flat

(b) EDX of one of the spots

Figure 6-31: SEM and EDX observation of the second silicon target

These imprints could be explained by the hypothetical sequence of events depicted in Figure 6-32. First, in the over-wet phase, droplets are emitted and form a deposit on the target with maximum thickness directly in front of the emitter. Then in steady phase, pure ion emission occurs. The ions sputter away the deposited material over a circular region, getting through the thinner parts before the thicker parts. At higher current levels, the ion beam is broader and the deposited material is completely removed. We have not observed any etching of the underlying silicon by the ion beam. If silicon is being etched, the etch rate appears to be slower than for the deposited material, as no roughness of the silicon is observed. It is likely that we would be unable to detect small amounts of etching from the SEM observations we have made. Cleaning the substrate before observing it would make any etching of the silicon easier to detect.

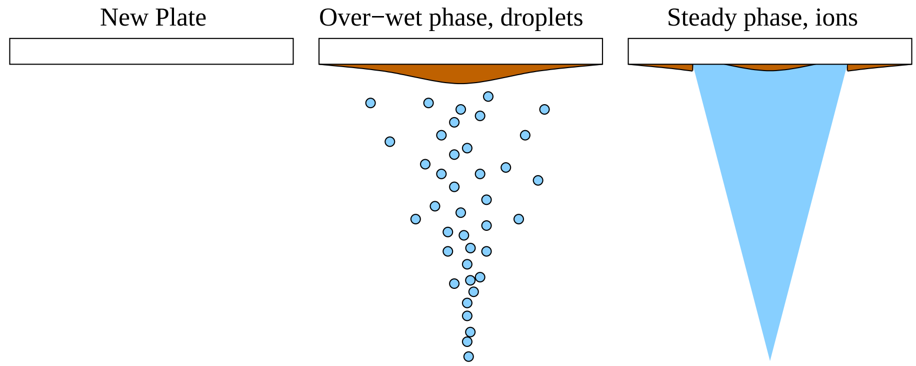

Figure 6-32: A hypothetical set of events which could produce the observed circular patterns on the target. Starting from a clean plate, first a thick layer is deposited in the droplet regime. Then ion emission occurs, clearing off part of the deposited material. A region of deposit remains in the center of the cleaned out area where the deposit was thickest, or where the ion flux was lower

6.4.3 ITO Target

When firing with the ITO target, we were able to film part of the array through the vacuum chamber window. Figure 6-33 shows frames from the movie, matched up with the voltage and current traces of the experiment. Figure 6-34 shows the difference between those same frames and the first image of the plate that was taken, to highlight the changes.[3] The frames that are shown correspond to periods of visible activity of the thruster, and are usually close to times when the emitter voltage was increased.

Frames (a), (b) and (c) are before, during and after the first observed current spike. Before this spike only slight changes from the initial frame were visible. For a few seconds, a blue glow is visible, then distinctive spots are visible on the plate. In frames (d), (e) and (f), the area where the blue glow had been visible becomes grey, except for roughly circular areas above the emitters that remain clear. The circles on the lower-right-hand row slowly increase in diameter. In frames (g), (h) and (i), the grey area spreads farther to the lower right. The circles that slowly grew in size become irregular, and a new row of circles begins to appear. The blue area in

[3]The difference frames cover a slightly larger area. The extra area was cropped out of the unprocessed frames because it underwent no visible change between frames.

Frame (i), unlike the one in Frame (b) does not go away, even when the electrospray is turned off.

6.4.4 Full Thruster Fired Against Target

We did one final ITO-plate test which differed from the others in two respects. First, this time it was a fully assembled thruster that was placed in contact with the ITO plate. Second, the emitter die had already been fired and was in the depletion phase at the start of the experiment. It would be interesting to do some extra experiments in which these two situations have been tried independently. The goal of this experiment was to determine how clean the emissions of the thruster were after the over-wet phase of operation.

All that was observed on the target die was a slight nebulous pattern in front of the emitter area. Microscope observation reveals no distinct pattern other than a few spots (Figure 6-35), and nebulous patches of brightness and darkness. This indicate that the dark orange patches that were visible on the previous targets may be related to the early operation of the thruster, and that once the over-wet (or steady) phase is complete, the thruster operates more cleanly.

During operation we also filmed the thruster through the ITO plate. The series of pictures in Figure 6-36 shows the light emission from the thruster progressing with increasing current from a not-too-uniformly glowing disk, to to a glowing disk with a few very bright spots. The first frame, taken with the light on, is to serve as a reference for the position of the thruster.

Unfortunately, these data were taken in the negative polarity, so there is no way of telling if we are observing electron or ion emission. It would be highly instructive to repeat this experiment in the positive polarity to exclude electron emission. This technique is also very promising for better understanding the over-wet and steady phases of operation.

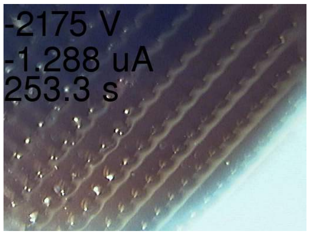

(a)

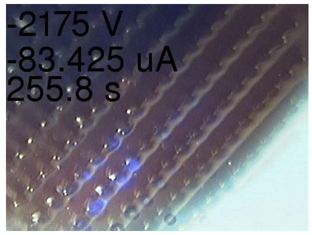

(b)

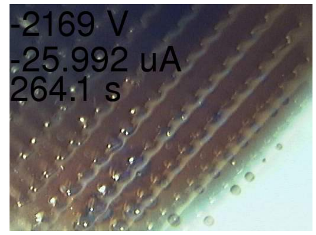

(c)

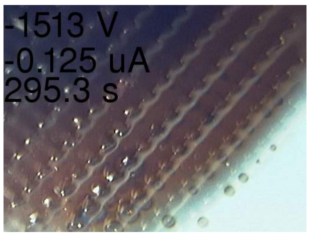

(d)

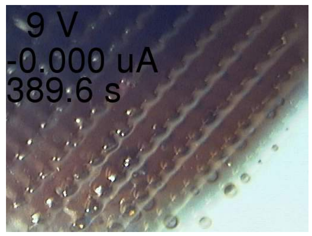

(e)

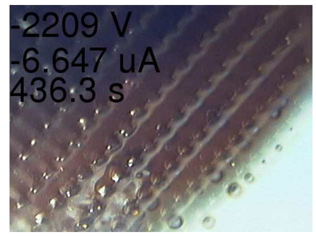

(f)

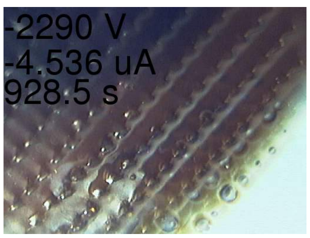

(g)

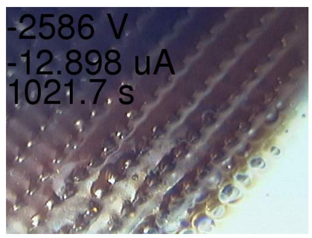

(h)

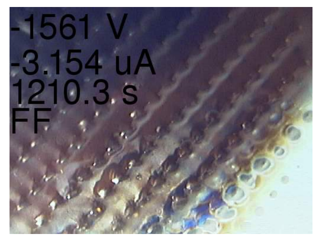

(i)

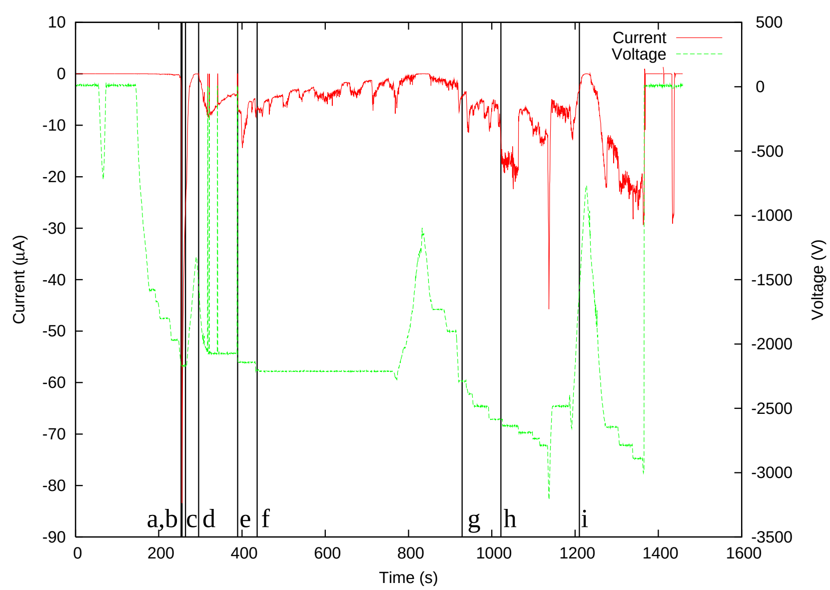

(j) Current and voltage during the experiment. Vertical bars indicate picture times

Figure 6-33: Shots taken while firing through ITO plate. The rows of emitters, located behind the target, are clearly visible

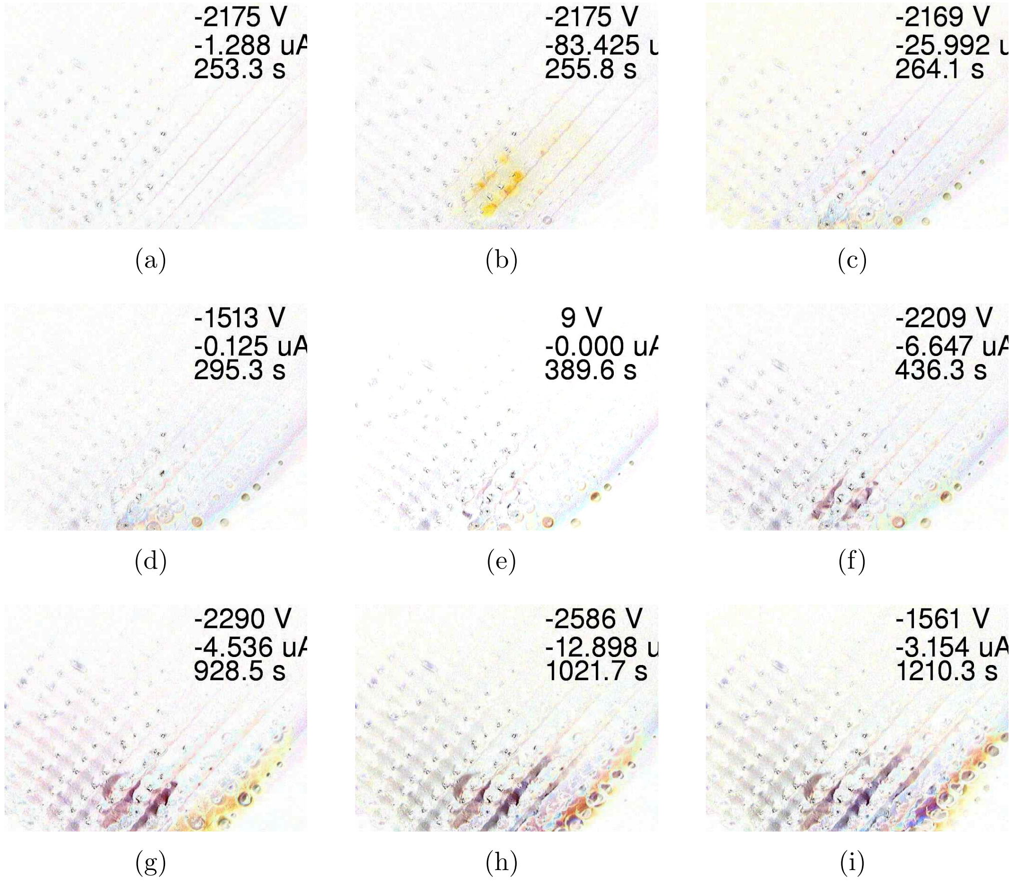

Figure 6-34: Difference between each shot and the initial shot of the plate

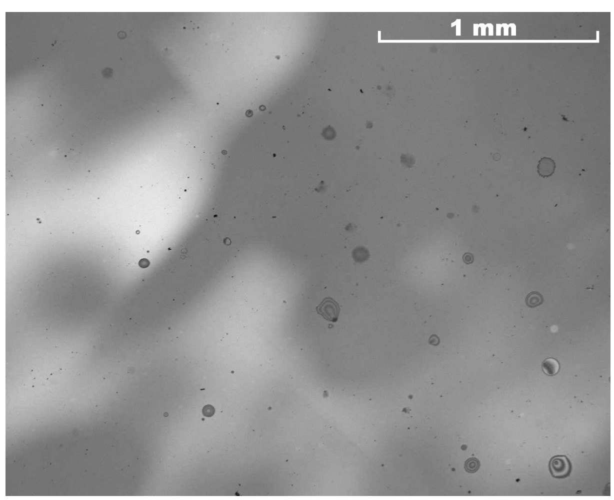

Figure 6-35: The most feature-rich area of a target that was fired upon by a previously depleted thruster. Only nebulous variations in luminosity and a few spots are visible

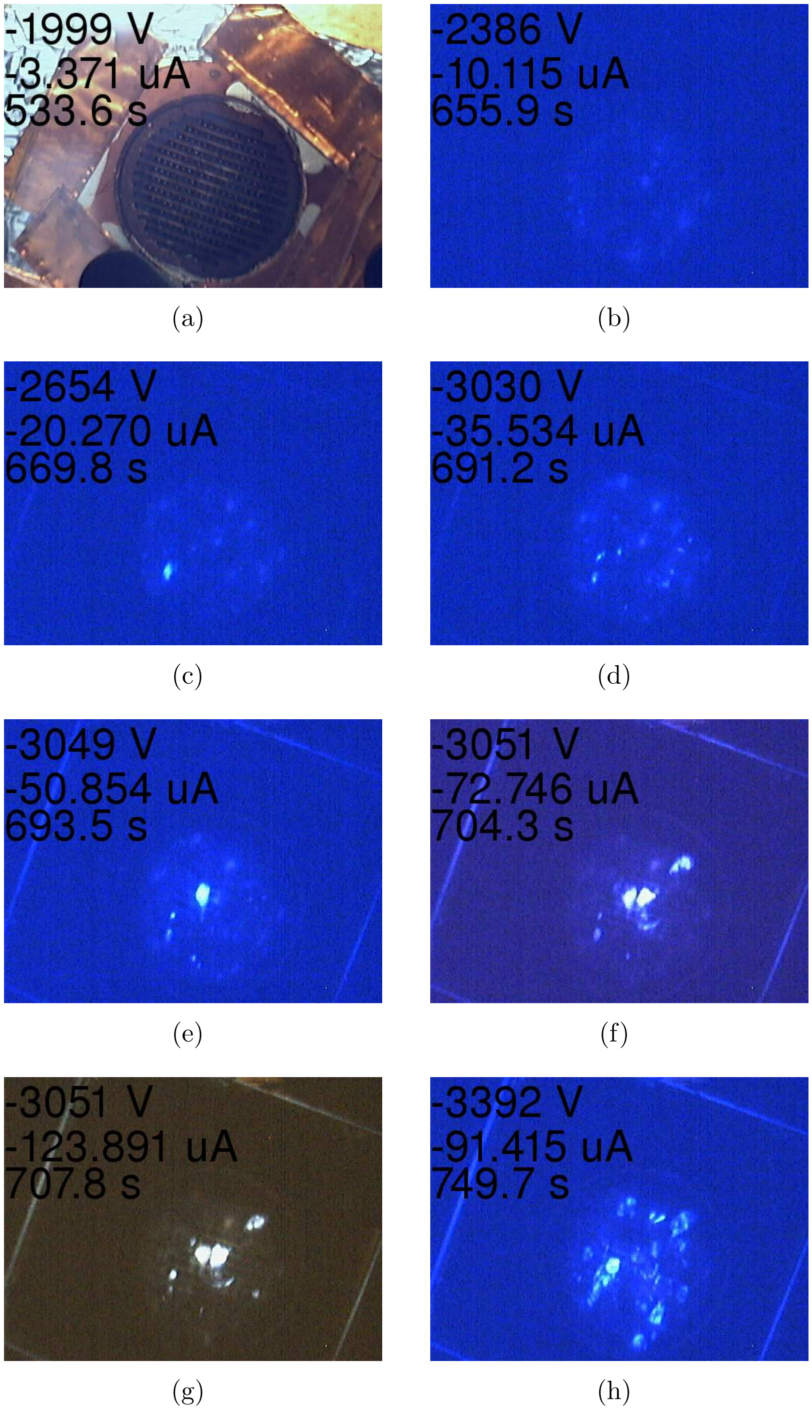

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 6-36: A thruster in the depleted phase, operating at high currents, against an ITO plate

6.5 Damage to Thruster

6.5.1 Emitters

During the first tests, the thruster was operated in the negative polarity, with a few brief bursts in the positive polarity, out of fear of damaging the emitters [65]. In later tests, we were much bolder about firing in the positive polarity, and the thrusters continued to function.

The only degradation we observed was a deposit which formed on some of the emitters. Indeed, during the depletion phase, when the thruster was operating at high current levels, there was light emission from localized areas of the emission region as in Section 6.4.4. When the thrusters were disassembled and inspected, these areas appeared lighter than the rest of the thruster. This difference in color could be reduced with water, acetone, isopropanol and sonication, but not completely eliminated. Figure 6-37 shows one of these soiled areas.

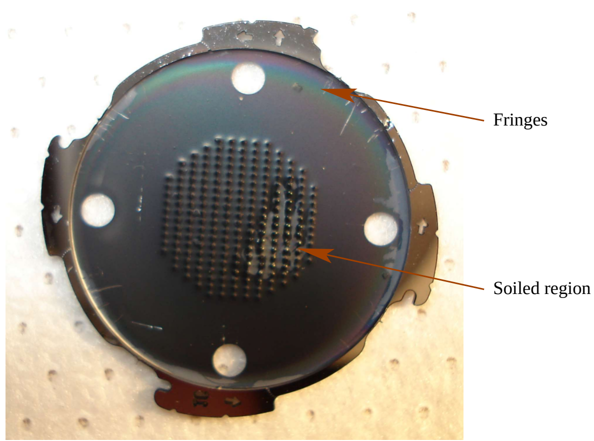

Figure 6-37: An emitter die on which a region has been altered by previous emission. One can also see fringes from the thin film of liquid on the thruster surface

Scanning electron microscope observation reveals that the altered areas of the thruster have material present on the silicon surface. The material takes on a variety of shapes such as ripples or nodules, as seen in Figure 6-38. EDX of the material shows stronger carbon and fluorine signals, and a weaker silicon signal, in the soiled areas than the clean areas (Figure 6-39), suggesting that the observed material could

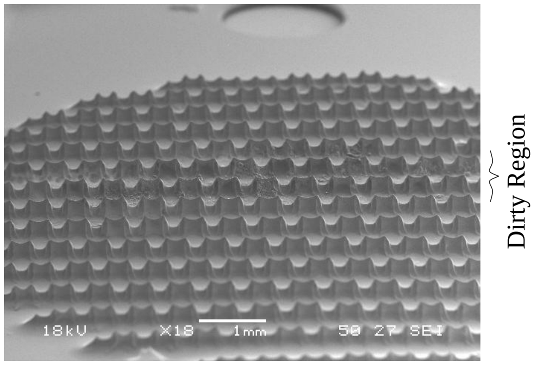

(a) A couple of rows of emitters have globs of material on them

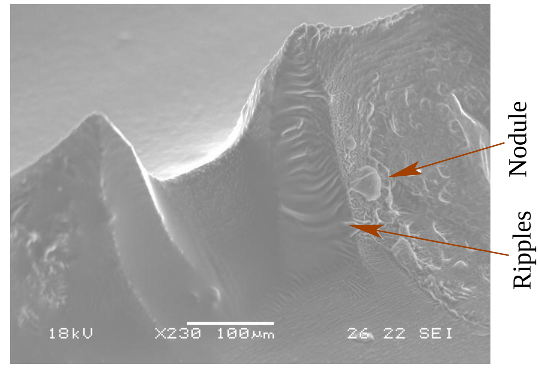

(b) Nodules and ripples are visible on this emitter

Figure 6-38: Altered emitter surfaces (before cleaning) exhibit varied globs of material

be generated from the EMI-BF_4.

While performing EDX analysis with high electron beam currents, we discovered that the deposit can be removed by an electron beam, as can be seen in Figure 6-40. It is possible that heating the sample would have the same effect.

Despite the fact that the emitters did not appear damaged, dies that were cleaned and refueled performed less well the second time they were fired. We have yet to determine what has changed on the die to degrade performance, and what brought about that change. So far, the performance does not seem to further degrade after the second refill, but no systematic testing has been performed to check this assertion.

6.5.2 Extractor

Observation of the extractor after use reveals deposition of material on the extractor beams (Figure 5-24(b)). It seems likely that this material is depositing during the over-wet phase of operation, but this hypothesis needs to be verified to ensure that there will be no long-term wear of the extractor electrode due to beam impingement.

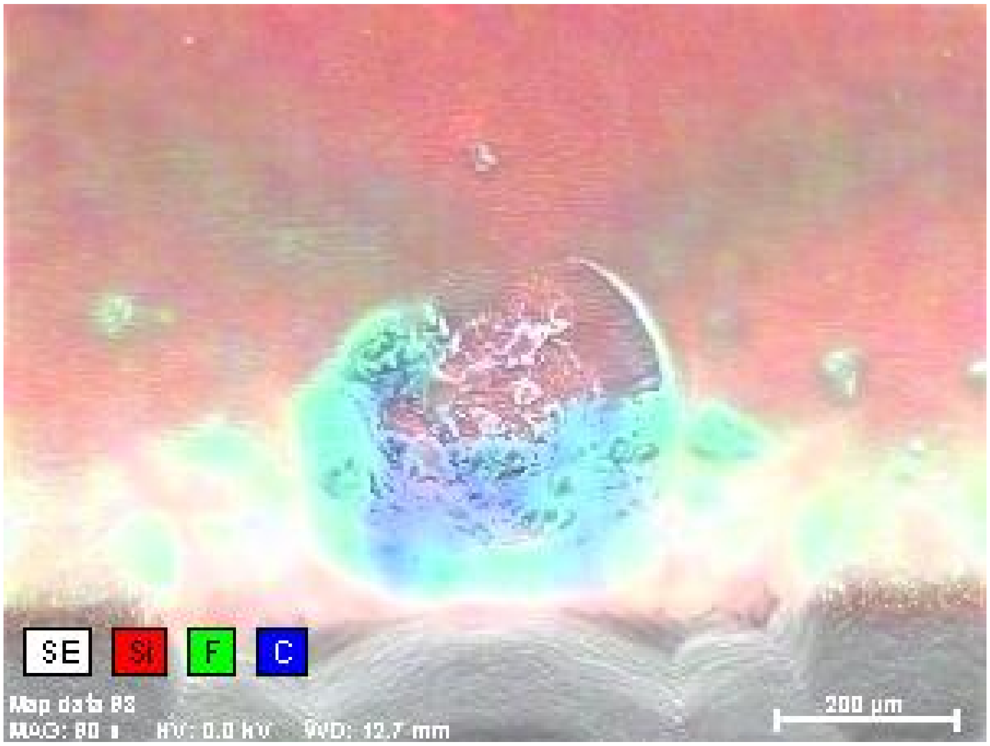

Figure 6-39: EDX analysis of a soiled emitter, seen from above

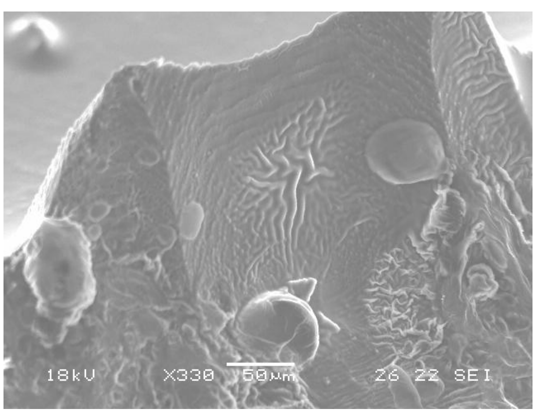

(a) Before e-beam

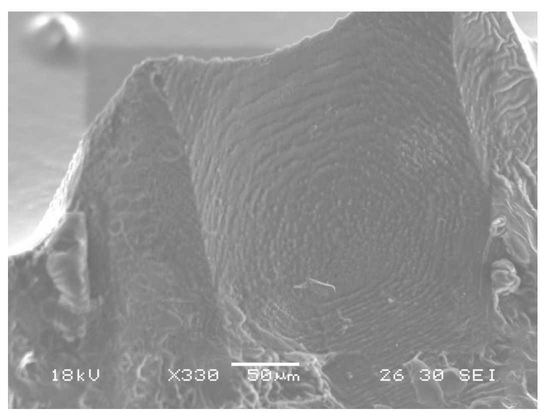

(b) After e-beam

Figure 6-40: A strong electron beam completely removes the deposit

6.6 Effect of Particle Contamination

Once fabrication was finished, none of the assembly and testing was done in a clean-room environment. In one test with a single-emitter die, a white spot about 1 mm in diameter was visible on the wettable region around the emitter. SEM inspection showed that a small fiber, which was invisible to the naked eye, was present on the die at that location (See Figure 6-41). This fiber may have contributed a lot of the extractor current that was observed during this particular test (and probably also some of the collector current). This type of perturbation may have been present in other experiments where it was negligible compared to the signal from the hundreds of emitters that were present. Clearly it would be desirable to do the fueling and assembly steps in a particle free environment to increase repeatability of results.

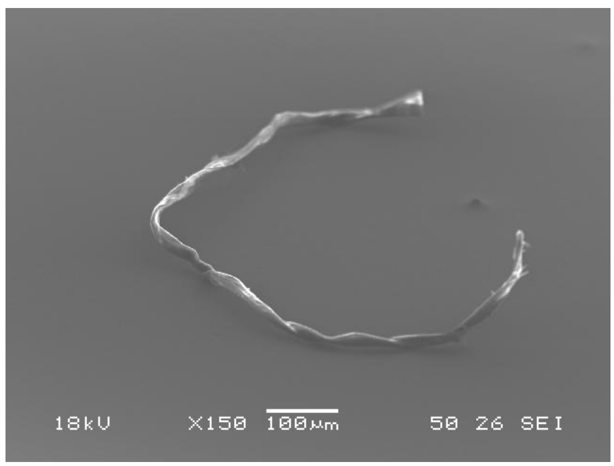

Figure 6-41: A single fiber, invisible to the naked eye, can cause spurious beams which hit the extractor and increase the interception fraction

6.7 Conclusion and Future Work

In this chapter we have reviewed the diverse experiments that were carried out. The general trend is very positive: the thrusters fire, they do not short, they emit more current per emitter than expected, the current intercepted on the extractor is low, and the thrusters emit in the ion regime.

During our review of the data, a number of interesting follow-up experiments were suggested. We review these suggestions here:

Low current interception data: For these experiments, the measurement ranges were selected to include the highest expected current levels (auto-ranging often led to missing data). As a result, the intercepted current was lost in the noise at low current levels. IV characteristic experiments could easily be repeated with a lower range on the extractor current measurement. This would allow the proportion of interception to be observed at lower current levels.

Interception causes: We identified broken dies, particles, amount of emitter recessing, emitter asymmetry, and the method of wetting the emitter die as possible sources of increased interception. These different possible causes should be tested independently to see if any of them actually are related to interception.

Angular dependence: For now the angular dependence of the beam has only been probed in the direction parallel to the slots. It would be interesting to see if the angular dependence in the other direction is similar. Also, this experiment could easily be repeated with a slit in front of the collector to make the interpretation of the measurement more direct.

Electron emission: The existence of electron emission should be explored further, and, if possible, confirmed by time-of-flight. Firing a dry emitter die seems like a likely method to get electron emission.

Orange/black substance: By firing a die that has been prepared to be in the steady or depletion phase at plate, we could see whether the orange/black deposition observed on targets takes place during all phases of operation, or only during the over-wet phase. A better clamping scheme should be used to ensure that the emitter die is parallel to the target. In that case, does the orange/black substance continue to be localized on one side of the target pattern? Does the way the die was wetted change where the orange/black substance is?

Thruster glow: By filming a thruster firing firing at an ITO plate in the over-wet or steady phases of operation, we could check the uniformity of emission in these phases. This should be done in the positive polarity to preclude electron emission from the emitters.

Erratic patterns: At what current levels do the patterns we observe on targets that were fired upon switch from circular to erratic? Does this transition correspond to the transition from exponential to linear IV characteristic?

Temperature dependence: This experiment should be repeated with better temperature control, and with the ability to reach higher temperatures. What is happening at 70°C? Moreover, these experiments have not probed the current ranges at which the IV characteristic becomes linear, to avoid depleting the thruster during the experiment. It would be highly instructive to know if the slope of the linear portion of the IV characteristic changes with temperature.

Transients: Are the transients in the steady phase the same as in the depletion phase? How long does the thruster have to be turned off to return to maximum overshoot after being fired?

Phase duration: By firing a thruster at a fixed voltage or current level, we could determine how long the different phases of operation last, and see how long the propellant supply lasts.

Time of flight in over-wet phase: The time-of-flight measurements were all taken during the steady phase or the depleted phase of operation. One of the EMI-IM time-of-flight traces suggests that droplets were present in the beam at the start of the experiment and material is deposited on target in the first stages of operation, suggesting that droplets may be emitted during the over-wet phase. Time-of-flight data should be taken during the over-wet phase to confirm this hypothesis.

Die wetting methodology: A consistent methodology should be developed for wetting the emitter dies.

Soiled emitters: More ways of removing the material that sometimes forms on emitters should be tried. For example, it would be easy to check whether heating is the mechanism by which the electron beam removes the material. Current observations suggest that this material forms at high current levels in the depletion regime, this hypothesis should be checked by systematic experimentation.

Observation of tips before and after firing: So far, tips have been observed before and after firing, but never the same tip before and after. There is so much variation between tips that the observations we have made so far do not allow us to say if tips have been damaged while firing. To understand why a brand new emitter die produces current at a significantly lower voltage than a refilled die, it would be instructive to observe the *same* emitter before and after firing to see if visible degradation has occurred.

Degradation of performance with refueling: So far emitter dies perform better the first time they are fueled than subsequent times. The degradation in performance seems mainly to occur at the first refueling. It would be useful to study this performance degradation by repeatedly firing some emitter dies in well defined conditions. Determining the causes of the degradation is also an important task. One candidate degradation mechanism is electrochemical reactions of the ionic liquid creating deposits or etching the silicon.

Startup Conditioning: It would be useful to check that the higher voltage needed to first start the thruster is indeed related to the need to draw liquid to the needle tip. For example, a freshly wetted emitter die could be submitted to a voltage that first gets it started, and then immediately observed in an SEM to see if all the emitters are wetted, or only a few.

For now we have mainly presented data with minimal interpretation. Now that all the data have been presented we are ready for Chapter 7 in which we will see what lessons can be learned from our observations.

Chapter 7

Discussion

In Chapter 6, we presented the data that were taken from the fabricated thrusters with minimal interpretation. This chapter is more speculative. Based on the presented data, Section 7.1 establishes the effectiveness of the electrospray array as a thruster, and discusses how best to operate it. Then, Section 7.2 studies the question of how uniform the emissions from this thruster are. Section 7.3 presents some hypotheses that could explain the phases of operation that the thruster undergoes. Section 7.4 comments on the startup voltage of the thruster, or rather the lack of a startup voltage in the measurements. Section 7.5 uses the temperature dependence data which has been gathered from the thruster to see what can be inferred about how the electrospray process is taking place. Section 7.6 considers the startup transients that have been observed. Finally, Section 7.7 considers the benefits and difficulties of scaling the thruster to higher densities and/or greater numbers of emitters.

7.1 Effectiveness as a Thruster

7.1.1 Performance Estimation

We can estimate the performance of the thruster from the IV characteristic, the time-of-flight data, and the angular distribution data. We will assume that the distribution of masses of the emitted particles is given by the time-of-flight data, this will allow us to compute an efficiency η_{poly} due to the polydispersity (many types of particles going at different velocities) of the beam. Then, we will use the angular distribution to estimate the efficiency η_{ang} due to beam divergence. Combined, these two inefficiencies are the kinetic component of the efficiency η_{k} from Equation (1.2.5). We will also assume that there is a non-unity non-kinetic efficiency η_{nk} given by

$$\eta_{\text{nk}} = \frac{|V| - V_{\text{extract}} - |RI|}{|V|}, \tag{7.1}$$

where V_{extract} corresponds to the 7 to 8 eV loss during emission [24], and R is the resistance of the electrical circuit, estimated at 300 $k\Omega$ from Figure 6-3. Assuming a worst case of $I = 500$ μA and $V = 1$ kV (we never saw currents this high at such low voltages), we find that η_{nk} >84%, mainly due to the resistive losses. For for more common current levels in the 50 μA range, at 1 kV, the resistive loss still dominates, but now $\eta_{\text{nk}} = 98\%$. With only extraction losses, we would have $\eta_{\text{nk}} = 99\%$ at 1 kV.

The proportion of monomer and dimer present in the emitted beam in Figure 6-25 is not identical to the proportion in other runs we have taken. Over all the runs we have taken, and ignoring the trimers, we find that in the positive polarity there is between 55% and 60% of monomer, and in the negative polarity there is between 60% and 65% of monomer (these percentages are by current, and hence also by number of particles). Let β be the fraction of monomer (the fraction of dimer is then $1-\beta$). We use Equation (5.82) from [117] i (also Equation (2.1.3) from [23]), which is tedious but straightforward to compute, to evaluate η_{poly}

$$\eta_{\text{poly}} = \frac{\left(1 - (1 - \sqrt{\zeta})\beta\right)^2}{1 - (1 - \zeta)\beta}, \tag{7.2}$$

where ζ is the ratio of dimer mass to monomer mass. Table 7.1 summarizes the results of this calculation, and Figure 7-1(a) shows what η_{poly} would be with different values of β.

We estimate the angular efficiency assuming that the emitted particles have their direction uniformly distributed within a cone that has half-angle $\theta_{\max}$. The angular

Polarity	Negative		Positive	
Species	BF_4^- (monomer)	$(BF_4\text{-EMI})\ BF_4^-$ (dimer)	EMI^+ (monomer)	$(BF_4\text{-EMI})\ EMI^+$ (dimer)
β (%)	60 to 65		55 to 60	
Mass (AMU)	86.8	284.8	111.2	309.2
$\langle Mass \rangle$ (AMU)	156 to 166		190 to 200	
ζ	3.281		2.781	
η_{poly} (%)	93.3 to 94.0		94.4 to 94.8	

Table 7.1: Data used to compute polydispersity efficiency

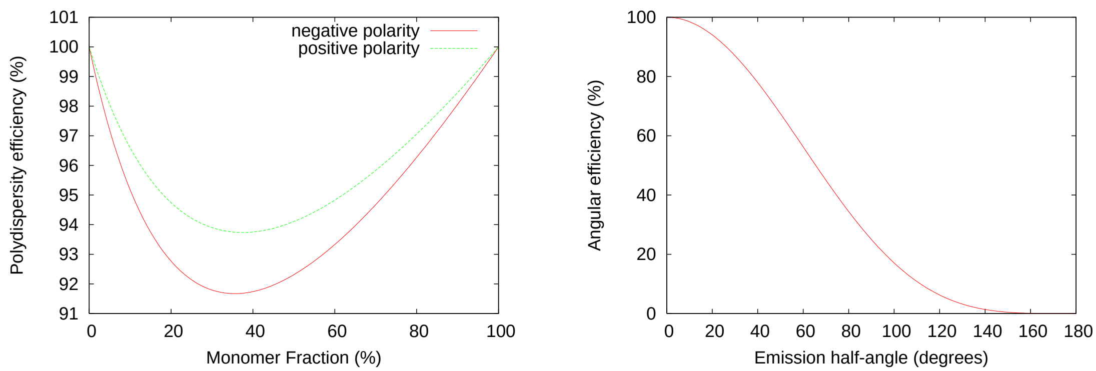

(a) Effect of Polydispersity

(b) Effect of beam divergence

Figure 7-1: Effect of polydispersity and beam divergence on efficiencies for parameters different from ours

efficiency is given by

$$\begin{aligned} \eta_{\text{ang}} &= \left(\frac{\text{Thrust with angular dependence}}{\text{Thrust without angular dependence}}\right)^2 \\ &= \left(\frac{\int_0^{\theta_{\max}} 2\pi \sin(\theta)\cos(\theta)\mathrm{d}\theta}{\int_0^{\theta_{\max}} 2\pi \sin(\theta)\mathrm{d}\theta}\right)^2 \\ &= \cos^4\left(\frac{\theta_{\max}}{2}\right). \end{aligned} \tag{7.3}$$

The resulting efficiency is η_{ang} =94% for $\theta_{\max}$=20° (from Section 6.2.9), Figure 7-1(b) shows the dependence of η_{ang} on $\theta_{\max}$.

To summarize, the predicted inefficiencies are mainly due to beam divergence and the multiplicity of species in the beam (94% efficiency each). The ohmic losses depend on the applied current, only becoming significant at currents above 10 μA. Extraction losses are very low with typically better than 99% efficiency, except at very low voltages. At intermediate current levels, the thrust efficiency $\eta_0 = \eta_{\text{nk}}\eta_{\text{poly}}\eta_{\text{ang}}$ is thus around 87% overall. Figure 7-2(a) shows the values calculated from the IV-characteristic in Figure 6-14(b), under the assumption that β=60% in both polarities. The low-voltage efficiency drop is due to the increasing importance of the extraction energy in η_{nk}, while the high-voltage drop is due to the ohmic loss, and is worse when the current rises at lower voltages.

We now compute thrust using Equation (1.2.5)

$$T = \sqrt{2\dot{m}P_{\text{in}}\eta_0}, \tag{7.4}$$

where P_{in} is just the product of current and applied voltage, and $\dot{m}$ is deduced from the current by

$$\dot{m} = \frac{I\langle M\rangle}{\mathcal{N}e} = \frac{I\left(\beta M_{\text{monomer}} + (1-\beta)M_{\text{dimer}}\right)}{\mathcal{N}e}, \tag{7.5}$$

and

$$\eta_0 = \left(\frac{|V| - V_{\text{extract}} - |RI|}{|V|}\right)\left(\frac{\left(1-(1-\sqrt{\zeta})\beta\right)^2}{1-(1-\zeta)\beta}\right)\cos^4\left(\frac{\theta_{\max}}{2}\right). \tag{7.6}$$

The resulting expression is plotted in Figure 7-2(b), still assuming β=60%.

The specific impulse is obtained by dividing the thrust by $\dot{m}g$. At 1000 V it is about 3050 s in the positive polarity, and 3240 s in the negative polarity, and increases with the square root of the voltage. With an accelerator electrode, this specific impulse should be adjustable over a wide range, limited from above by the available power and maximum operating voltage, and from below by the energy spread of the ions. For the lowest specific impulses, the energy spread of the ions will lead to an reduction in thrust efficiency which has not been considered in this analysis.

7.1.2 Maximum Thrust

The highest calculated thrust for our runs is 13 μN for the thruster is for the 502 emitter die in the positive polarity. The corresponding thrust density is 0.15 N/m^2,

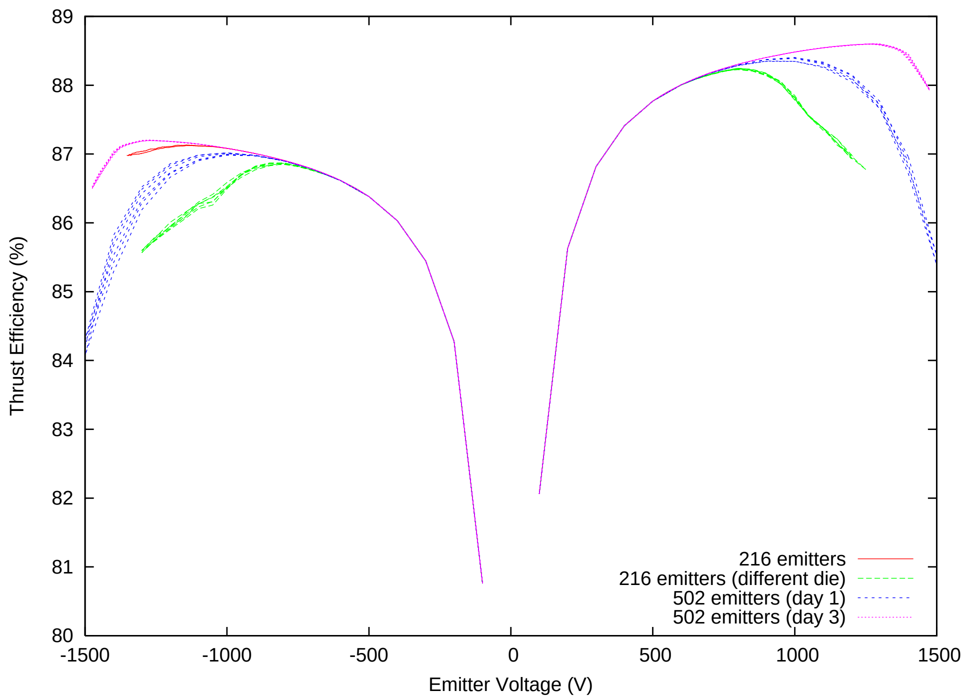

(a) Thrust efficiency

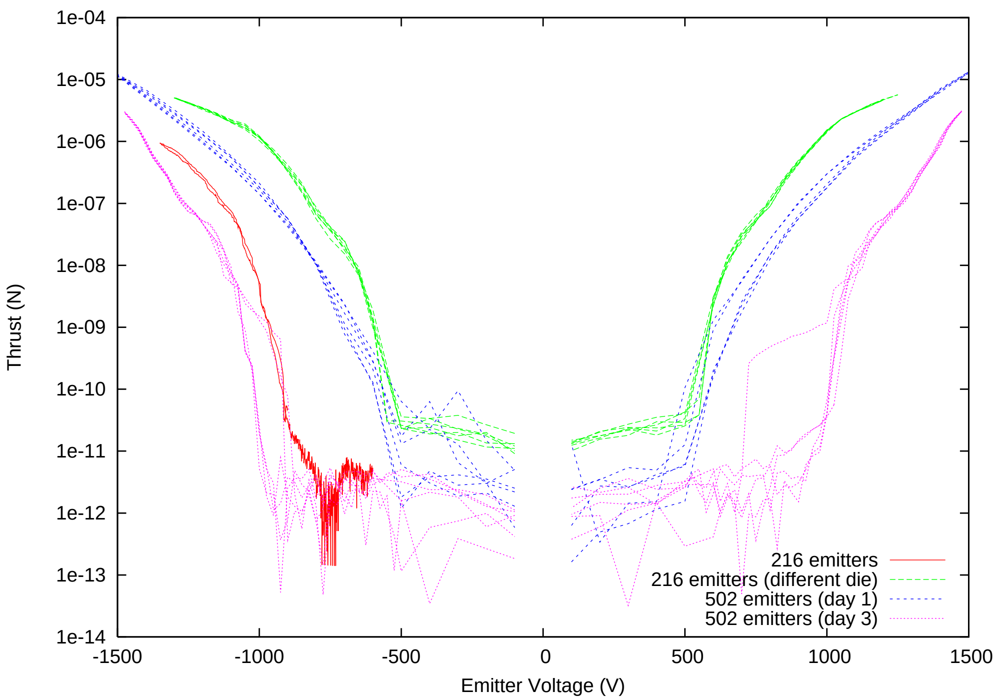

(b) Thrust

Figure 7-2: Thrust efficiency and thrust as a function of applied voltage for some typical runs and the thrust per emitter is 26 nN. The power consumption at maximum thrust is 275 mW. Estimating the mass of the thruster at 5 g, the thrust-to-weight ratio is about $2 \cdot 10^{-4}$. This ratio does not take into account the mass of the power supply. To compare the thruster with other technologies, the mass of the power supply should be included. Some remarkable miniaturized high voltage power supplies are available, for example the Q-series power supplies from EMCO High Voltage Co. (Sutter Creek, CA) come in a 5 g package and can deliver ± 2000 V at 250 μA with better than 50% efficiency. It is therefore conceivable to build a complete thruster with power conversion electronics in under 25 g (accounting for high voltage switches and other components), for a thrust-to-weight ratio of $4 \cdot 10^{-5}$ for the complete system.[1]

This thrust is not an absolute limit, and in fact, this same thruster was run at nearly twice that current level in one test for which the data acquisition failed. What seems to limit operation currently is the rapid depletion of liquid. Perhaps higher thrusts per tip would be achievable with better hydraulics, and in particular with porous feed emitters.

One limit that may be of concern is the erosion of the thruster when the intercepted current is too high. It may therefore be desirable to limit the thrust to levels where there is low interception. Figure 7-3 shows the intercepted current as a function of thrust for the usual set of runs. There are clearly large variations here between thrusters. For the best operating emitter die, (216 emitters, different die), 1% interception is reached at 4 μN. For the second best operating emitter die (502 emitters, day 1), the interception never exceeds 3%. The two remaining examples have high interception. Finding the causes of that interception so that all dies can behave reliably is an important area of future work.

Thrust measurements are currently underway at Busek Co. (Natick, MA) using the torsional thrust balance described in [118]. Hopefully these data will confirm the theoretical predictions.

[1]This figure still does not include the power source itself. We are assuming that during thruster operation a power source that is normally used for other spacecraft operations can be diverted to powering the thruster.

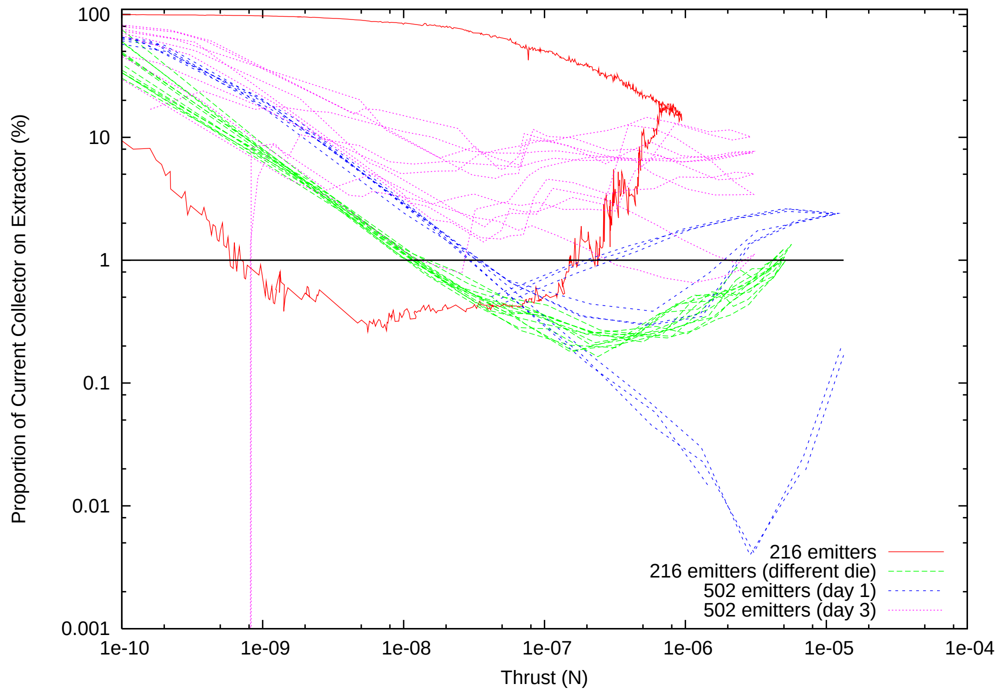

Figure 7-3: Fraction of current intercepted by extractor as a function of thrust

7.1.3 Control

Figure 6-8(b) shows the IV characteristic of the thruster taken on different occasions. By varying the applied voltage, the thrust per emitter can be varied over 6 orders of magnitude. Adjusting the voltage by hand, it is easy to set the emitted current to within 10 to 20% of a desired value. These characteristics imply that this thruster is very easy to control. Thus, it seems that thrust can be tuned over 6 orders of magnitude, with no more difficulty than providing a slow current control loop to set the emitter voltage. This high range of tunability suggests that this thruster would not benefit from schemes which have been proposed to turn on emitters individually or in small groups [41].

Temperature control of the thruster, while possible, only provides a small tuning range. Combined with the fact that temperature significantly increases current, probably by leakage through the Pyrex insulation, temperature does not seem like a useful control parameter. One case in which temperature adjustment may be useful is to increase the maximum current emitted by the thruster. However, until a propellant resupply system is added, higher currents can only be achieved for a few tens of seconds at a time. Temperature control may also be useful to independently set thrust and specific impulse. However, the addition of an accelerator electrode would have the same effect with less complexity.

So far, it appears that operation is just as satisfactory at high or low current levels. The only notable differences are increased interception by the extractor at the highest current levels, and higher sensitivity of current to applied voltage at low current levels.

7.1.4 Propellant Supply

The most critical system that is missing in our thruster is the propellant resupply system. Without it, the thruster lifetime is limited to a few minutes at the highest current levels. There are two main tasks for a propellant supply system: detecting when propellant needs to be added to the thruster, and adding that propellant.

Both of these tasks may be carried out passively, for example by using a porous propellant tank placed in capillary contact with the surface of the emitter die. However, using a porous propellant tank as the only propellant tank seems like an inefficient course of action. In a porous tank, a large fraction of the mass is occupied by the porous material of the tank. If a significant amount of propellant is to be transported, a non-porous tank will be lighter. This inefficiency is compounded by the fact that the thruster may not be able to extract all the propellant from a porous tank. Most of the liquid may be extracted from the large pores, but the smaller corners will no doubt continue to contain liquid. Worse, the mass of liquid in the porous tank may become disconnected, in which case there will be nothing to draw that part of the liquid to the thruster.

The alternative to using a porous tank is to have liquid from a non-porous tank delivered to the surface of the emitter die. In our thruster, the simplest way to get liquid to the surface of the emitter die is through the back of the emitter die. A hole can be made through the die allowing liquid to be fed through. We propose that this hole be placed at the outer periphery of the emitter die where the electric field is low, in order to avoid possible complications arising from the bulk liquid surface deforming in the electric field (See Figure 7-4). A capillary tube can be attached to the backside of the emitter die using a variety of techniques, for example using an Upchurch Scientific Nanoport™ (Oak Harbor, WA) assembly. Propellant can be fed through the tube as needed to resupply the emitters. A micro-valve appropriate for the low flow rates in an electrospray thruster has been designed by Busek Co. (Natick, MA).

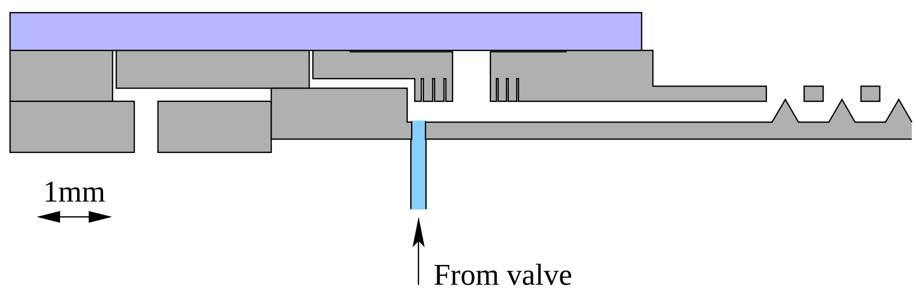

Figure 7-4: Proposed location for propellant feed

This solves the problem of adding propellant to the thruster. The problem of knowing how much propellant remains on the thruster remains to be solved. With our current understanding, it appears that the IV characteristic is a good way to detect when the thruster is running low on propellant. If this hypothesis is confirmed, then the signal to resupply the emitters with propellant can simply be derived from observations of the voltage needed to reach a desired current level.

7.2 Uniformity of Emission

The first step in understanding the IV characteristic, is knowing whether the emitters are firing uniformly. If the emitters are firing uniformly then the IV characteristics we have can be converted into the characteristic for a single emitter just by dividing the current by the number of emitters. This scenario seems unlikely, particularly in the low current regime where a 10% change in applied voltage can increase the current by an order of magnitude. Indeed, a 10% difference in field enhancement between emitters due to non-uniformity of the emitter formation process seems perfectly plausible. In the high current regime, where current varies linearly with voltage, all the emitters probably are emitting similar currents because the sensitivity of current to voltage is much lower.

Observing the silicon targets, we can say for sure from the high current test (Figure 6-29) that all the emitters turn on and start making irregular patterns on the target within about a factor of two of the initial starting voltage. Combined with the fact that slope of the current per emitter is the same for all runs (Figure 6-8(a)), we can be confident that at high current levels, most of the emitters are firing, and have a linear current dependence on voltage. Whatever process is limiting the current at high current levels is providing ballasting to help uniformize the emission across the emitters.

The low current test (Figure 6-30) suggests that the emitters aren't all starting at once, but that emitters that are close to each other seem to start operating at around the same time. This observation suggests that extractor inclination relative to the emitters may be more at fault than differences in geometry between emitters, or that there is no short-range non-uniformity between emitters. This conclusion is supported by Figure 6-33 in which some rows of emitters appear to start at higher voltages than others.

7.3 Phases of Emission

We have identified three phases of thruster operation, of which only the steady phase is characterized by low interception on the extractor, a repeatable IV characteristic, and a low starting voltage. Understanding what causes these different phases is of utmost importance so that ultimately, a thruster can be made to operate only in the steady phase. How can we explain these phases of operation? We present two hypotheses.

7.3.1 The Hydraulics Hypothesis

SEM observation of a depleted emitter shows striking differences compared with an emitter with an emitter that has not been fired yet (see Figure 7-5 or figures 3-23 and 3-22(b)). When the surface has just been wetted, the liquid film appears to be nearly continuous across the emitter surface, with the surface treatment roughness slightly visible near the ridge lines. After firing, the tips of surface-roughness features are visible across the whole surface with a spacing on the order of 10 μm.

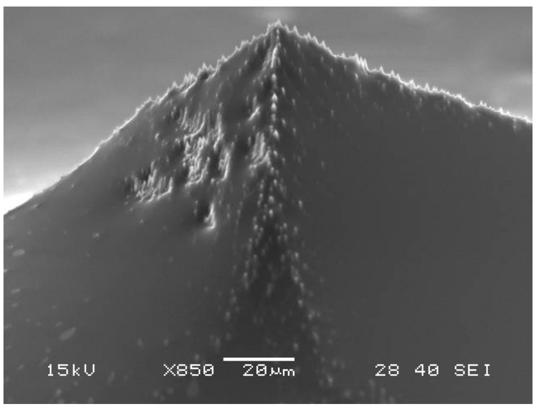

(a) An emitter that has just been filled. Because this emitter is particularly low, it nearly completely wetted as soon as liquid was applied

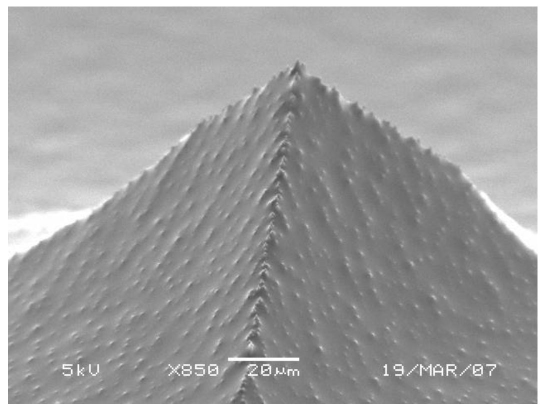

(b) An emitter that has been fired into the depletion phase

Figure 7-5: The liquid surface is markedly different before and after firing

In Section 2.2.3, we had operated under the assumption that liquid would fill the valleys of the surface roughness, and based our hydraulic impedance calculations on that assumption. It seems that depending on the phase of operation, we may be filling surface-roughness valleys at different scales, each scale being characterized by a different hydraulic impedance and pressure of the liquid in the liquid film. Figure 7-6 shows a cut through one side of the emitter. When the thruster is first wetted, so much liquid is present that the surface film floods the surface roughness, and is only supported by the ridge lines of the emitter. A liquid film many microns in thickness is present with near zero curvature (and hence pressure). The resulting hydraulic impedance is very low. As liquid is consumed, the film thins out, impedance increases, and the radius of curvature of the liquid becomes comparable to the radius of curvature of the emitter surface. At this point, the liquid surface starts to touch the highest roughness features. As liquid continues to be removed, the liquid surface eventually gets pinned on even the smallest roughness features, its radius of curvature getting smaller and smaller as these steps occur. We hypothesize that the phases of emission can be explained by studying the effect of the state of the liquid film on emission.

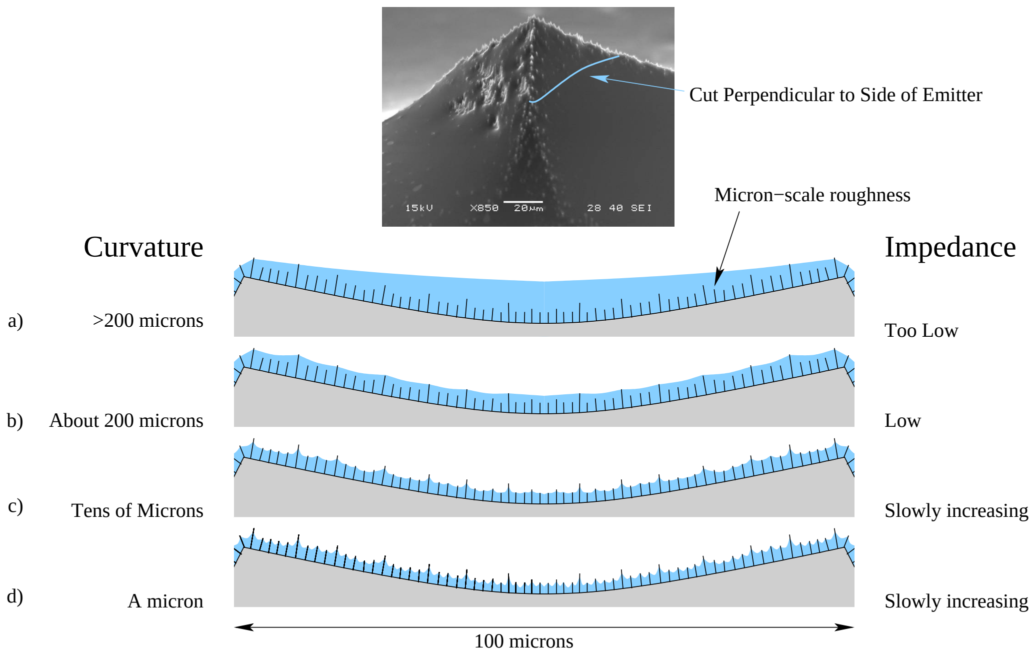

Figure 7-6: A cut perpendicular to the side of an emitter showing the effect of liquid volume on hydraulic impedances and curvatures

We present an outline of how the different phases of operation are caused by the changes in liquid level on the emitter:

Over-wet phase: Initially, there is a low hydraulic impedance, which could allow high flow rates and droplet emission. Droplet emission can explaining the material deposited on the targets in Section 6.4, and could be verified by running a time of flight experiment on a thruster in the over-wet phase. If the beam divergence is greater this regime, then increased interception can be expected. The unsteady nature of the flow could be due to hydraulic interaction between the emitters.

Steady phase: Then, the liquid gets pinned on the tallest roughness features. The increased impedance allows stable operation of the thruster, and the steady phase begins. It is unclear what causes the remarkable stability of the steady phase. Perhaps the electric field is able to draw liquid from the base of the emitter to keep a constant thickness film on the emitter.

Indeed, there is an interplay between radius of curvature R of the liquid and electric field E, since they each adjust the pressure in the liquid. Assuming that the liquid is curved only in one direction, pressure is zero when

$$\frac{1}{2}\epsilon_0 E^2 = \frac{\gamma}{R}. \tag{7.7}$$

Figure 7-7 plots this relation, it can be used to determine whether electric pressure or surface tension dominates in a given situation. Figure 7-8 plots the field around a conical emitter with an extractor that has an axis of revolution, for an applied voltage of 2000 V.[2] The field on the sides of the emitter is on the order of $1 \cdot 10^5$ V/cm which is equivalent to a radius of curvature of 100 μm, this value is close to the radius of curvature which would be expected from the SF_6 etching which produced the emitter, and which the liquid surface should have when it gets pinned on the highest roughness features.

Thus, the electric field might be able to draw liquid from the base of the emitter, without being able to unpin the liquid from the highest roughness features. This would provide regulation of the film thickness on the emitter, and allow nearly steady-state performance.

Depletion phase: Once the level on the base of the emitters starts to reach the level of the lower roughness features, the curvature of the liquid surface rapidly drops into ten of micrometers range or less, and the electric field is no longer strong enough to drain liquid from the base, the emitters start to deplete and

[2] The field was computed using FreeFem++ v1.470002

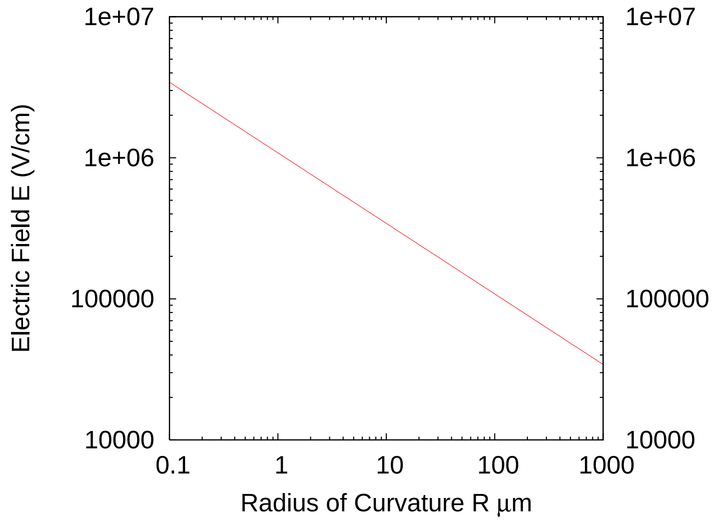

Figure 7-7: Electric fields and radii of curvature which create the same pressure

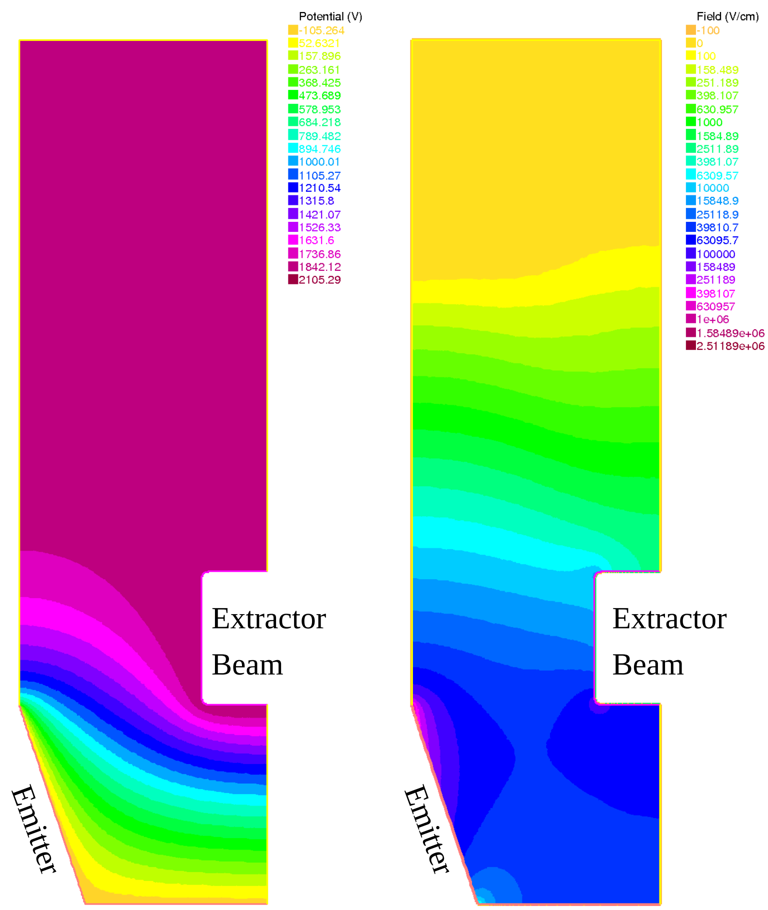

Figure 7-8: The potential and field around a conical needle for a 2000 V applied voltage

the thrust level drops.

This hypothesis suggests a few tests:

- If the over-wet phase is caused by the large amounts of liquid that can fit in the curvature of the emitter surface, then making an emitter with a surface that is positively curved everywhere should eliminate the over-wet phase.

- If the emitter is observed multiple times as it transitions through the phases of operation, it should be possible to gain more insight on the state of the liquid at each phase. (For example, does the over-wet phase indeed end when the liquid level on the emitters reaches the level of the highest roughness?)

7.3.2 The Electrochemistry Hypothesis

According to Lozano [65], electrochemical reactions at the surface of the emitter cause the current to drop if the electrospray process is run for too long in the ionic regime in a single polarity. The depletion regime may simply be the same current drop that Lozano observed, caused by electrochemical reactions. This hypothesis seems particularly tempting given that depletion seems to be accompanied by deposits on the emitter surface. Moreover, when they are refilled, emitters emit less current at a given voltage that when they were originally filled. This hypothesis would be easy to check simply by using Lozano's proposed solution to the problem, which is to alternate the emission polarity at a frequency on the order of 1 Hz. Some work with EMI-Im suggests that in that case the electrochemical effects are reversible, so even slower switching voltages are possible [119].

7.4 Startup Voltage

None of the IV characteristics in Chapter 6 show any evidence of a startup voltage. Emission starts below the noise floor and increases continuously from there. A commonly used expression for the startup voltage is [57]

$$V_{\text{start}} = \sqrt{\frac{\gamma r}{\epsilon_0}} \ln \left(2D/r \right), \tag{7.8}$$

where r is the radius of curvature of the emitter tip, and D is the distance to the extractor electrode. (An approximate form of this expression can be derived by taking $E = V/r$ as the field at the tip, and considering that emission starts when the electric pressure $\frac{1}{2}\epsilon_0 E^2$ exceeds the surface tension $\frac{2\gamma}{r}$.) For a 1 μm tip radius r (large for our emitters) and 250 μm distance D to the extractor, the starting voltage for EMI-BF_4 is 476 V. In all our measurements, the current rose above the noise floor at higher voltages than this, so it is not surprising that we were unable to witness the electrospray startup. In the following section we will see another reason why we may not be able to see the startup voltage.

7.5 Ion Evaporation

Electrospray is a complex process involving the solution of a multi-scale electrostatic free boundary problem, coupled with fluid flow, and charge relaxation. Some success has been achieved in understanding the cone-jet regime, but so far, the pure ion emission regime remains poorly understood [120].

The pure ion emission regime is generally assumed to to follow the Iribarne-Thompson ion evaporation model [121], for which a concise derivation can be found in [23]. In this model, the ion evaporation current density is given by

$$j = \frac{\sigma k T}{h} e^{-\frac{q\phi - \sqrt{q^3 E/(4\pi\epsilon_0)}}{kT}}, \tag{7.9}$$

where k is Boltzmann's constant, T is the absolute temperature, h is Planck's constant, ϕ is the free energy of ion evaporation (in eV), E is the electric field, q is the charge of the emitted ions, σ is the surface charge, and ϵ_0 is the permittivity of vacuum. Setting $\sigma = \epsilon_0 E$, and introducing a field enhancement factor β and an effective emission area α, such that $E = \beta V$ and $I = \alpha J$, we get a formula for the emitted current as a function of applied voltage

$$I = \frac{\alpha \epsilon_0 \beta V k T}{h} e^{-\frac{q\phi - \sqrt{q^3 \beta V/(4\pi\epsilon_0)}}{kT}}. \tag{7.10}$$

Rearranging, we get

$$T \ln\left(\frac{I}{VT}\right) = T \ln\left(\frac{\alpha\epsilon_0\beta k}{h}\right) - \frac{q\phi}{k} + \sqrt{\frac{q^3\beta}{4\pi\epsilon_0 k^2}}\sqrt{V} = AT - B + C\sqrt{V}. \tag{7.11}$$

Therefore, for a liquid surface with constant geometry, $T \ln\left(\frac{I}{VT}\right)$ as a function of T and $\sqrt{V}$ defines a plane. The slope in the T direction is A, the slope in the $\sqrt{V}$ direction is C, and the intercept is $-B$. Thus, by plotting slices of this plane in the T and $\sqrt{V}$ directions, we can deduce A, B and C. We can then get the free energy of ion evaporation ϕ, the field enhancement β and the effective area α using

$$\phi = \frac{BkT}{q} \tag{7.12}$$

$$\beta = \frac{4\pi\epsilon_0 k^2 C^2}{q^3} \tag{7.13}$$

$$\alpha = \frac{e^A h}{\epsilon_0 \beta k} \tag{7.14}$$

Directly applying this method to electrospray seems difficult because the geometry of the liquid surface is unknown a priori. Emission could be taking place from the tip of a Taylor cone with a tip radius that is unknown to us. Moreover, protrusions could emanate from the tip of the cone, further enhancing the field in a difficult to predict manner. It is therefore unexpected to see parallel straight lines appear in figures 7-9, in which the data from Figure 6-14(b) have been replotted. Using Equation (7.11), we find $\phi = 1.12$ V, $\beta = 3.4 \cdot 10^3$ cm^{-1} and $\alpha = 6.9 \cdot 10^{-12}$ m^2 (or $1/\beta = 2.9$ μm and $\sqrt{\alpha/\pi}$ =1.5 μm). For a 500 V applied voltage, the field at the emission site is 0.15 V/nm. The standard deviation of $\log_{10}(I_{\text{meas}}/I_{\text{calc}})$ is 0.23, which corresponds to a factor of 1.7 error.

The fitted parameters have reasonable orders of magnitude, and we are therefore

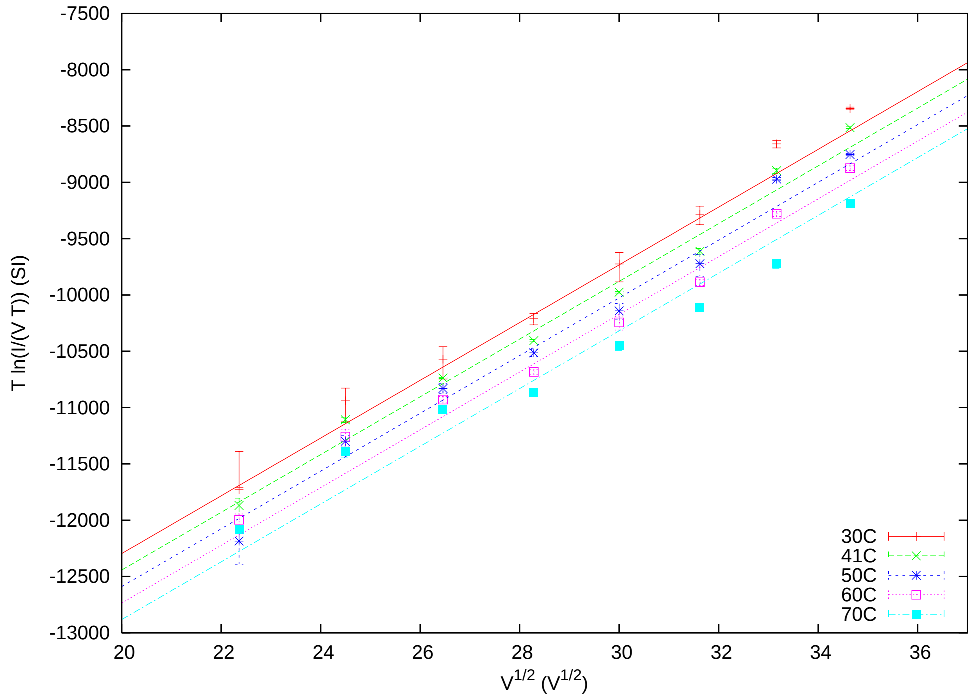

(a) Plot of $T \ln\left(\frac{I}{VT}\right)$ as a function of $\sqrt{V}$. The slope is C and the intercept is $AT - B$

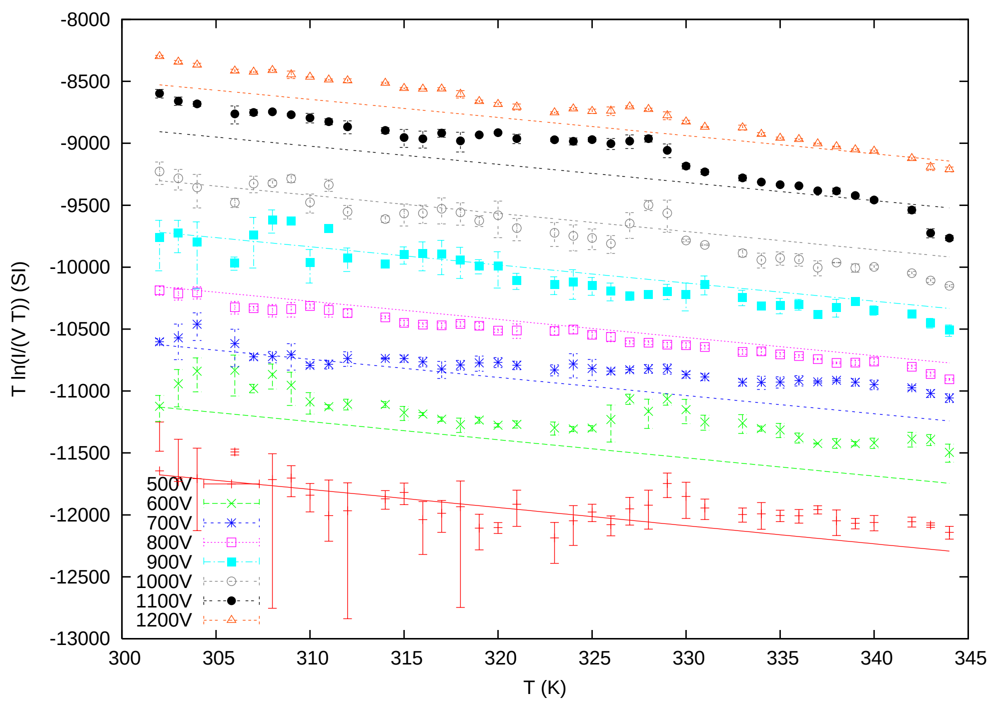

(b) Plot of $T \ln\left(\frac{I}{VT}\right)$ as a function of T. The slope is A and the intercept is $-B + C\sqrt{V}$

Figure 7-9: Modeling ion emission using the Iribarne-Thompson model, taking $\phi =$ 1.12 V, $\beta = 3.4 \cdot 10^3$ cm^{-1} and $\alpha = 6.9 \cdot 10^{-12}$ m^2. Only low current data is available here (up to 1 nA/emitter) led to speculate on what could be allowing a fixed field enhancement factor to be observed. The simplest explanation seems to be that emission is taking place without the formation of a Taylor cone. Thus the liquid surface would have roughly the same geometry whether or not the electrospray is operating. The presence of surface roughness with sharp tips may contribute to the operation without a Taylor cone, since it can pin the liquid surface in a conical configuration, even without an applied voltage (see Figure 7-10).

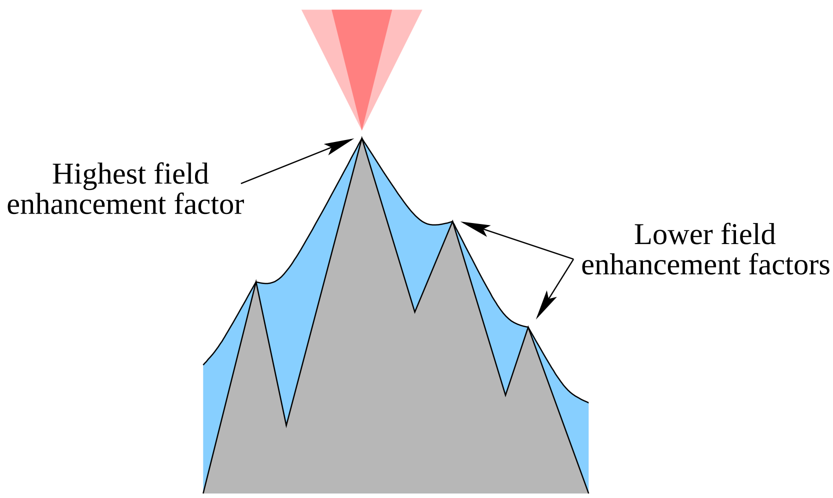

Figure 7-10: Emission site created by sharp roughness features pinning the liquid

Because of the strong dependence of I on β it is likely that only a few emitters are emitting. This situation could change at higher currents than are covered by the dynamic temperature-dependent data. Indeed, at higher currents, we expect that hydraulic impedance will begin to have an effect by limiting the availability of liquid at the emission site. This could have a ballasting effect, causing emission to be more uniformly distributed between the emitters. At sufficiently high applied voltages, emission sites farther from the tip will achieve sufficient fields to begin emitting. Sharp features along the emitter ridges should have the highest field enhancement factor, which could lead to the irregular emission patterns with threefold symmetry in Figure 6-29(b). As the liquid supply depletes, emission sites near the emitter tip will run out of liquid first, as they are farther from the liquid supply. The remaining sites with lower field enhancement factors will require higher voltages to start emitting, leading to the observed behavior of the IV characteristic in the depletion regime. The degraded performance of refilled emitter dies could be explained by damage to the emission sites with the highest field enhancement factor, forcing less enhanced sites to be used.

This hypothesis seems like it can explain many emission properties that we have observed, and it quantitatively agrees with the Iribarne-Thompson model. If it turns out to be correct, it would be analogous to a behavior observed with liquid metal ion sources, for which there is sometimes a small amount of emission precedes the formation of a Taylor cone [57]. In our case, the presence of sharp surface roughness features (as seen in Figure 3-15) may allow a conical liquid surface supported by the surface roughness to form, allowing high field enhancement and ion emission before the surface is destabilized into a Taylor cone.

7.6 Startup Transient

The startup transients in Section 6.2.10 can be easily be interpreted in terms of liquid depletion occurring on the surface of the emitter. At rest, liquid is distributed along the emitter so that the pressure in the liquid is uniform. In operation at steady state, a pressure gradient exists in the liquid, to drive flow towards the tip. To transition from the uniform pressure at rest to the lower pressure during operation, liquid needs to be removed from the surface, resulting in the transient we observe.

As in [73], the liquid level on the surface can be expected to follow a diffusion equation. Thus, the transient we are observing is analogous to the change in heat flux when a heated object is placed in contact with a colder one. Initially, heat transfer is limited by the thermal contact resistance between the two objects, but rapidly a thermal gradient builds up in the contacted objects, and the heat transfer in the bulk becomes dominant. In the electrospray case, the initial current level is presumably emission limited, becoming liquid transport limited as the pressure drops at the emission site and a pressure gradient is set up.

We will not attempt to quantitatively match the observed transient because the discussion in Section 7.3 suggests that we do not quantitatively understand the relation between amount of liquid present on the surface and pressure in the liquid. Nevertheless, if we assume that the change in liquid state on the surface is small, then the pressure gradient should be proportional to the steady-state current current i_0, and the amount of liquid removed from the surface to reach steady state should also be proportional to i_0. Using the notation from Section 6.2.10, the amount of liquid removed from the surface during the transient is proportional to $i_1\tau$. Thus, we expect $i_1\tau/i_0$ to be a constant. This does indeed appear to be the case in Figure 7-11.

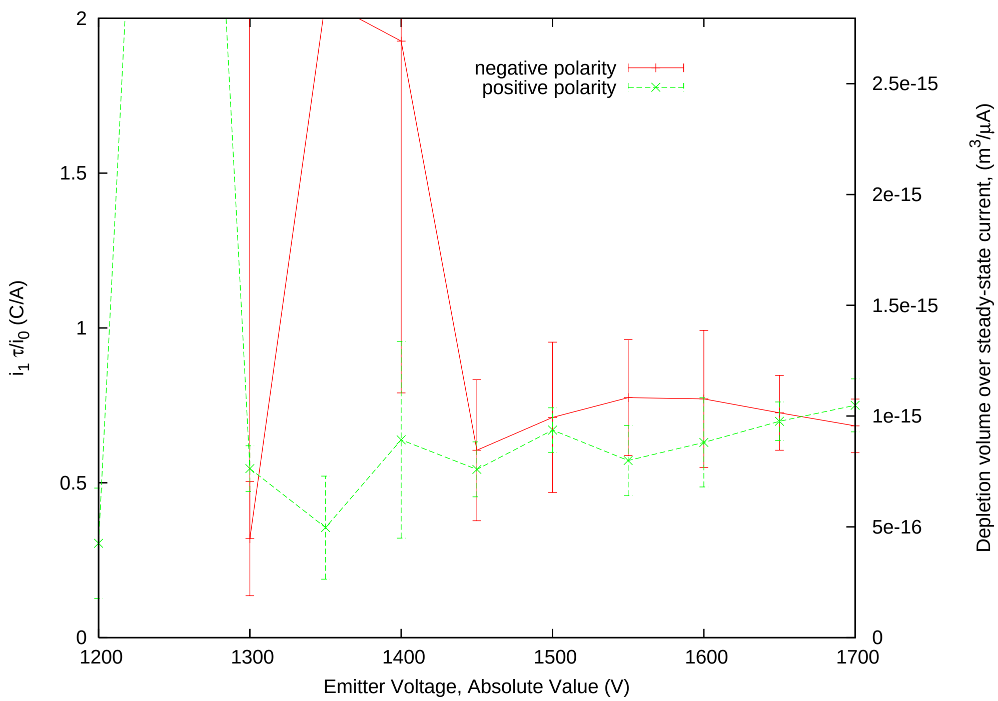

Figure 7-11: The amount of liquid depleted during the startup transient is proportional to the steady-state current. (The conversion from flow-rate to current assumes that emitted particles have an average mass of 175 AMU)

For a current of 100 nA, near the high end of the operating range for a single emitter, the amount of liquid depleted from the surface in steady state is about $1 \cdot 10^{-16}$ m^3, i.e., 100 μm ^{3}. To put this number into perspective, this is the amount of liquid we expect to find in the top 10 μm of the emitter.

This analysis is still very preliminary. A more detailed analysis of the decay times τ should allow the dependence of emission on pressure near the tip to be explored.

7.7 Scaling the Thruster

This thesis has focused on getting together all the ingredients needed to make a functional, fully integrated electrospray thruster that works. Hence, we have deliberately been conservative in the design, and the result is a small emitter area (113 mm^2) with a relatively low emitter density (4.4 tips/mm^2). How can these two numbers be increased to further multiply the total emission and thrust? What hurdles and changes will be involved?

7.7.1 Increasing Emitter Density

Emitter Fabrication: So far, the spacing of the emitters that are in the same extractor slot is limited by the mask size needed to make the emitter. With the current emitters, this mask size can be reduced by making emitters shorter, or by increasing their inclination (which implies improvements in surface treatment). Serrated blade emitters (as in Figure 3-9) without a vertical sidewall can also be made with tight spacing between tips, but there will only be a shallow dip between emitters, which decreases field enhancement. Moreover the resulting emitters will not be as sharp as the three-ridge emitters we have carefully designed. With the use of nested masks, an improved blade emitter can be made, with a large elongated rectangle mask to define a sloping base for the whole row of emitters, and individual small masks to define each individual emitter tip.

Extractor Slot Spacing: To reduce the spacing between extractor slots, the emitter tip height needs to be closer to the top of the extractor. This can be achieved by thinning the extractor. However, thinning will eventually make the extractor unable to withstand the applied electric field; we defer this problem to the next section, as it is also a concern when increasing the active area of the array. Another problem with extractor thinning is that it will compound the problem of non-uniform etching we are already experiencing, and which is causing a third of the extractors to be destroyed during the etch that forms them.

Keeping the same extractor thickness, we can avoid beam impingement by increasing the recessing of the emitters, to the point where the tips are in line with the top of the extractor. The precision with which the height of the emitters has to be set is comparable to the spacing between emitters. Thus, as the spacing is reduced, variations in tip depth between emitters will become a problem. The variation in tip depth can be reduced by making shorter emitters (which will limit the amount of recessing that can be done) or by using nested masks.

Finally, the slot size will eventually be limited by the precision of the extractor assembly methods. With the current assembly method, a slot of size of less than 100 micrometers seems unreasonable.

Field: The tip electric field will be modified in a number of ways by the density increases we have just proposed, and changes in startup voltage should result. Tighter slot spacing, should increase the field by reducing the electrode-to-emitter-tip distance. Shorter emitters, more closely packed emitters and excessive recessing should decrease it. The most critical parameter, however, is probably the radius of curvature of the emitter tip, and we do not know what effect scaling the emitters will have on that radius of curvature.

Space Charge: As the spacing between emitters decreases, we will need to worry that the electric field seen by an emitter will be reduced by ions from neighboring beams, which would cause diminishing returns as the scale is reduced. Some modeling of space charge effects as emitter density is increased can be found in [23, 58].

Overall, decreasing emitter spacing along an extractor slot down to 50 μm seems relatively easy, especially with nested masks. We do not see a fundamental limit to going lower than this to the 10 μm range, but new and unexpected fabrication problems will surely arise as we try to reach more than an order of magnitude below the current size.

It is more difficult to say how much reduction in the spacing between extractor slots will be possible, as there are more constraints involved. The fabrication will rapidly become more difficult, as accurate vertical positioning of the emitter tips and of the top of the extractor will become important. Below 100 μm spacing, the accuracy of the assembly method will need to be revisited or it will become a limiting factor.

7.7.2 Increasing Active Area

We have already discussed, in Section 5.3, ways of increasing the area set aside for emitters by up to a factor of two, by reducing margins in the current design. Even then, less than 20% of the total device area would be devoted to emitters. Improvements in utilization could be achieved by reducing the spring stiffness (and hence their area), or by keeping the same springs and increasing the die size or wafer thickness. In any case, significantly increasing the active area will introduce some new challenges.

Electrostatic Forces: In Section 2.2.4, we estimated the deflection due to the electric field in the current design. As we try to increase the size of the extractor, this deflection will quickly become a problem (as will the deflection of the emitter die) because of the fourth power dependence of deflection on extractor beam length. This problem will be compounded if we try to reduce the extractor thickness while increasing emitter density. The first solution to this problem will be to form clusters of emitters with thicker support ribs between them in a cellular plate configuration. As the active area is further increased support pillars may need to be introduced to reduce the span between the support spacers which are currently located at the periphery of the emitter die. Spacers will be costly in area because they will need to be surrounded by fuel traps and an insulation gap with undercuts.

Fuel Supply: Currently we have proposed that fuel be added to the periphery of the emitter die. As the size of the active area increases, the distance from the periphery to the emitters increases linearly, and the number of emitters per unit length of periphery also increases linearly. As a result, the pressure drop needed to transport liquid from the periphery to the emitter bases will increase quadratically, and is likely to become a problem. Therefore, past a certain scale fuel supply to central regions of the die will have to be introduced.

Alignment: We expect that the precision of the assembly system will not change as long as clamping is at the edge of the emitter die, because the misalignment at the level of each finger will not change. Thermal mismatch will no longer be negligible at the 6" wafer scale because a 20°C difference in temperature between the emitters and the extractor will correspond to a 10 μm size difference across the wafer. Thus, even if the location of the center of the emitter die is insensitive to thermal expansion, the edge of the emitter die will start to see misalignment.

Wafer Bow: Wafer bows of tens of micrometers are typical across a 6" wafer, but decreases quadratically with die size, so on the order of 1 μm is expected for the 1" emitter dies we currently have. As the active area is increased, wafer bow could become a concern, particularly if vertical spacing is being tightly controlled to reduce the spacing between extractor slots.

Emitter Uniformity: Emitters at the periphery of a wafer require at least 20% less etching than ones at the center of a wafer. We are currently able to produce functional emitters across most of the wafer area by masking dies on which the emitters are ready. As dies get larger, we may have to mask clusters of emitters on a single die in order for the outer emitters not to be destroyed before the inner ones are ready. Probably the best solution to this problem is to vary the emitter mask sizes across the wafer to compensate for the etch non-uniformity. However, this approach will require detailed characterization of the etch non-uniformity of the specific etcher to be used.

Assuming that we can increase the thruster size to a full 6" wafer with 50% area utilization, over 18,000 mm^2 would be available for emitters. If we can reduce tip spacing to 50 μm within an extractor slot, and 500 μm between slots, we will be able to pack 40 tips/mm^2, for a total of over 700,000 dies. If this scaling can be achieved, over to 10 mN of thrust could be expected from a single 6" diameter thruster. Such a thruster would have thrust and specific impulse comparable to an ion engine, but with higher thrust efficiency.

To reach this scale, the emitter density can first be increased without changing the extractor component, to allow many emitter dies to be fabricated at once. The next step will be to increase the area of the array. Since many new difficulties arise when increasing the array area, it may be preferable to break a large thruster down into modules smaller than a 6" wafer, that are independently assembled and supplied with propellant. This modular approach also adds a measure of redundancy against short-circuits and other failures.

Chapter 8

Conclusion and Recommendations

In this thesis, we have demonstrated a fully integrated electrospray thruster, weighing 5 g, operating in the ion regime, with low current interception, and specific impulse of 3000 s at 1 kV extraction voltage. An estimated thrust of 13 μN has been achieved at an extractor voltage of 1500 V for a power of 275 mW. This thrust was reached with a 502 emitter array, covering an area of 113 mm^2, so the thrust per emitter is 26 nN, and the thrust density is 0.15 N/m^2. The thrust efficiency is estimated around 85%. We have shown that it is possible to achieve low interception of the emitted beam on the extractor electrode, since there is less than 1% interception over a wide range of operating conditions.

A number of new technologies were needed to make this thruster possible, we highlight the main ones in Section 8.1. This thruster is just one step on the road to a space-worthy electrospray array. Section 8.2 will outline some recommended directions for future work.

8.1 Contributions

Fully Integrated Planar Electrospray Array Operating in the Ion Regime: The first fully integrated electrospray array operating in the ion regime has been reported and characterized. By "fully integrated", we mean that the only interface to the micromachined array is two wires to supply the extractor voltage, and a holder to support the array mechanically.

Easy Emitter Interchange and Adjustable Emitter-to-Electrode Spacing: An assembly method has been implemented which allows electrospray emitters to be replaced in about 2 minutes. This method self-aligns the emitters to within 10 micrometers of their intended location, and allows the emitter-to-electrode spacing to be varied simply by replacing a set of spacer elements. It is also possible to use just about any type of electrospray emitter with an existing extractor component, just by changing the emitter die.

Low current interception: We showed that it is possible to achieve low current interception in a microfabricated array, with hundreds of tips operating in the ionic regime.

Emitter Formation Model: A simple model to predict emitter geometry has been formulated. It has proven its usefulness by allowing emitters to be designed for a specific maximum surface inclination, and by providing methods to improve emitter tip quality.

Mechanism of Liquid Transport: SEM observation of wetted emitters and of progressing wetting fronts has confirmed that liquid is transported within the valleys of the roughened surface treatment. This knowledge allows estimation of surface permeabilities. The details of how liquid transitions from this intra-valley flow to the Taylor cone is still unresolved, however.

Time Resolved Target Observation: The use of ITO coated glass plates to observe the evolution of a target electrode as a function of time greatly increases the amount of information that can be gleaned by firing the thruster against a plate.

Short Turnaround Surface Treatment Test: The evaluation of wicking surface treatments using deionized water has reduced the iteration time for surface treatment exploration from days to minutes.

Horizontal Pyrex Gap Insulation: A method which allows electrical insulation between precisely aligned components, with a long surface length between the components, has been implemented. In this method, the components are fabricated from a single silicon wafer, bonded to Pyrex, and finally separated from each-other. This method facilitates the location of the emitter tips as close as possible to the extractor electrode.

Liquid Spillage Avoidance Methods: Methods to control the motion of liquid towards critical insulator regions have been successfully tested. These methods include: smooth surfaces to prevent wicking, fluoropolymer coatings to limit wettability, passage through low field region to discourage liquid progression, and liquid traps to trap liquid which is migrating towards sensitive locations.

Propellant Supply Detection Method: We have conjectured that the level of propellant on the thruster can be detected from the IV characteristic of the thruster. If this conjecture turns out to be correct, closed-loop control of the liquid-level on the emitter die will be possible without implementing a sensor, thus greatly simplifying the design of the thruster.

Emission Without a Taylor Cone: We have conjectured that emission takes place without the formation of a Taylor cone. If this is indeed the case, analyzing the emission characteristics could be greatly simplified.

8.2 Recommendations for Future Work

We have made recommendations for future work throughout this thesis. This section will serve as a reminder of the main directions for future work.

Wicking materials: We were able to make a wicking surface treatment that works, but it is currently not completely satisfactory because it does not offer good uniformity, and does not treat surfaces conformally. In particular, it cannot treat sidewalls. This is an important direction where improvements can be made, and past work [58] suggests that the solution is at hand. A number of alternate methods of making surface treatments have also been proposed in Section 3.3. An very promising replacement for surface treatments is to make bulk porous emitters, which have the potential to greatly increase performance.

Emitters: The most interesting direction to explore for emitter formation is nested masks, which should allow taller emitters, better process uniformity, and better scaling tradeoffs.

Scaling: Section 7.7 discusses the challenges of scaling the existing thruster. Clearly, there is room to improve thruster density. Initially only the emitter dies need to be changed to get higher density along the extractor slots. The size of the thruster will also need to be increased if thrusts in the millinewton range are desired. Increasing the size will introduce some new challenges, such as significant flexing due to electrostatic forces.

Thruster Process: The yield of the extractor component process is currently very low. We have outlined the difficulties to be overcome in Section 5.3.

Propellant Supply: The main missing piece from this thruster is the propellant supply system. The current design can be enhanced by adding a capillary feed through a hole in the periphery of the base of the emitter die. This hole could be laser-drilled in one of the existing dies, to cheaply explore this option. More details are given in Section 7.1.4.

Accelerator Electrode: The versatility of this thruster could be greatly increased with the addition of an accelerator electrode, to increase or decrease specific impulse. This could potentially be a non-microfabricated accelerator electrode.

Causes of Interception: So far we have observed a lot of variation in current intercepted on the extractor electrode, from emitter die to emitter die. We have proposed dust particles, die fractures, poor die wetting methodology, unintended topography on the emitter die, and lack of recessing as possible candidates for this variation. Since erosion due to intercepted ions is likely to be a lifetime limiting process, it is important to get a better understanding of what causes some dies to intercept current while others don't.

Systematic Characterization: There is a lot of room for improved characterization of this thruster, to confirm or refute some of the hypotheses that have been made in chapters 6 and 7. In particular, Section 6.7 gives a list of interesting experiments to conduct.

Expand to other applications: Electrospray is a very versatile technology. Very little in the solutions we have developed in this thesis is specific to the propulsive application of electrospray. Implementing these solutions in other applications will multiply the importance of this work, and have a significant impact on those fields.

Bibliography

[1] G. I. Taylor. Disintegration of Water Drops in an Electric Field. *Proceedings of the Royal Society of London*, 280(1382):383–397, July 1964.

[2] J. Fernández de la Mora. The effect of charge emission from electrified liquid cone. *Journal of Fluid Mechanics*, 243:561–574, 1992.

[3] M. Gamero-Castano and V. Hubry. Electrospray as a source of nanoparticles for efficient colloid thrusters. *Journal of Power and Propulsion*, 17(5):977–987, Sept.-Oct. 2001.

[4] I. Romero-Sanz, R. Bocanegra, J. Fernández de la Mora, and M. Gamero-Castano. Source of Heavy Molecular Ions Based on Taylor Cones of Ionic Liquids Operating in the Pure Ion Evaporation Regime. *Journal of Applied Physics*, 94(5):3599–3605, September 2003.

[5] P. Lozano and M. Martínez-Sánchez. Experimental Measurements of Colloid Thruster Plumes in the Ion-Droplet Mixed Regime. In *Proceedings of the 38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2002.

[6] J. B. Fenn, M. Mann, C. K. Meng, S. F. Wong, and C. M. Whitehouse. Electrospray Ionization for Mass Spectrometry of Large Biomolecules. *Science*, 246(4926):64–712, October 1989.

[7] Sebastian Kaiser, Dimitrios C. Kyritsis, Peter Dobrowloski, Marshall B. Long, and Alessandro Gomez. The Electrospray and Combustion at the Mesoscale. *Journal of the Mass Spectrometry Society of Japan*, 51(1):42–49, 2003.

[8] Weiwei Deng, James F. Klemic, Xiaohui Li, Mark A. Reed, and Alessandro Gomez. Liquid fuel microcombustor using microfabricated multiplexed electrospray sources. *Proceedings of the Combustion Institute*, 31:2239–2246, 2007.

[9] I. Wuled Lenggoro, Kikuo Okuyama, Juan Fernández de la Mora, and Noboru Tohge. Preparation of ZnS Nanoparticles by Electrospray Pyrolysis. *Journal of Aerosol Science*, 31(1):121–136, 2000.

[10] Weiwei Deng, James F. Klemic, Xiaohui Li, Mark A. Reed, and Alessandro Gomez. Increase of Electrospray Throughput using Multiplexed Microfabricated Sources for the Scalable Generation of Monodisperse Droplets. *Journal of Aerosol Science*, 37(6):696–714, June 2006.

[11] Dan Li and Younan Xia. Electrospinning of Nanofibers: Reinventing the Wheel? *Advanced Materials*, 16(14):1151–1170, 2004.

[12] S. Edward Law. Agricultural Electrostatic Spray Application: a Review of Significant Research and Development during the 20th Century. *Journal of Electrostatics*, 51–51:25–42, May 2001.

[13] E. Pugh. Method and Apparatus for Coating Articles. US patent 1,855,869, 1932.

[14] Boy Hoyer, Gunnar Sorensen, Nina Jensen, Dorthe Berg Nielsen, and Bent Larsen. Electrostatic Spraying: A Novel Technique for Preparation of Polymer Coatings on Electrodes. *Analytical Chemistry*, 68(21):3840–3844, 1996.

[15] Luis F. Velásquez-García. Private Communication, 2007.

[16] S. N. Jayasinghe, M. J. Edirisinghe1, and D. Z. Wang. Controlled Deposition of Nanoparticle Clusters by Electrohydrodynamic Atomization. *Nanotechnology*, 15:1519–1523, 2004.

[17] G. L. R. Mair. Theoretical Determination of Current-Voltage Curves for Liquid Metal Ion Sources. *Journal of Physics D: Applied Physics*, 17(11):2323–2330, November 1984.

[18] S. P. Thompson and P. D. Prewett. The Dynamics of Liquid Metal Ion Sources. *Journal of Physics D: Applied Physics*, 17(11):2305–2321, November 1984.

[19] M. Martínez-Sánchez, J. Fernández de la Mora, V. Hubry, M. Gamero-Castano, and V. Khayms. Research on Colloid Thrusters. In *Proceedings of the 26th International Electric Propulsion Conference (IEPC)*, October 1999.

[20] Robert Harold Bateman and John Anthony David Hickson. Multi-inlet mass spectrometer. European patent EP 0966022, 1999.

[21] Robert B. Cody. Chemical Analysis Method for Multiplexed Samples. US patent US 6,902,938, 2005.

[22] J. F. Mahoney, R. D. Moore, J. Perel, and A. Y. Yahiku. Research and development of a charged-particle bipolar thruster. *AIAA Journal*, 7(3):507–511, 1969.

[23] Paulo Lozano. *Studies on the Ion-Droplet Mixed Regime in Colloid Thrusters*. PhD thesis, Massachusetts Institute of Technology, February 2003.

[24] Paulo C. Lozano. Energy properties of an EMI-Im ionic liquid ion source. *Journal of Physics D: Applied Physics*, 39:126–134, 2006.

[25] M. Tajmar, A. Genovese, and W. Steiger. Indium field emission electric propulsion microthruster experimental characterization. *Journal of Propulsion and Power*, 20(2):211–218, 2004.

[26] Salvo Marcuccio, Angelo Genovese, and Mariano Andrenucci. Experimental performance of field emission microthrusters. *Journal of Propulsion and Power*, 14(5):774–781, 1998.

[27] J. Zeleny. The electrical discharge from liquid points an a hydrostatic method to measure electric intensity at their surface. *Physical Reviews*, 3:69–91, 1914.

[28] F. Kaufman. Ion Rocket. US patent 3,156,090, 1964.

[29] P. W. Kidd and H. Shelton. Life test (4350 hours) of an advanced colloid thruster module. In *Proceedings of the 10th AIAA Electric Propulsion Conference*, October 1973.

[30] E. Cohen, C. J. Somol, and D. A. Gordon. A 100 kv, 10-w heavy particle thruster. In *2nd AIAA annual meeting*, July 1965.

[31] R. E. Hunter. Theoretical considerations on non-uniformly charged expellant beams. Technical Report ARL-TN-60-138, Aeronautical Research Lab, August 1960.

[32] P. Lozano and M. Martínez-Sánchez. Ionic liquid ion sources: Characterization of externally wetted emitters. *Journal of Colloid and Interface Science*, 282:415–421, 2005.

[33] D. Garoz, C. Bueno, C. Larriba, S. Castro, and J. Fernández de la Mora. Taylor cones of ionic liquids from capillary tubes as sources of pure ions for electrical propulsion. In *Proceedings of the 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2006.

[34] S. Castro, C. Larriba, and J. Fernández de la Mora. Capillary vs. externally wetted ionic liquid ion sources. In *Proceedings of the 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2006.

[35] V. Hruby, M. Gamero-Castano, D. Spence, C. Gasdaska, N. Demmons, R. McCormick, P. Falkos, J. Young, and W. Connolly. Colloid thrusters for the new millennium, st7 drs mission. In *Proceedings of the 2004 Aerospace Conference*, March 2004.

[36] L. F. Velásquez-García, A. I. Akinwande, and M. Martínez-Sánchez. A Microfabricated Linear Array of Electrospray Emitters for Thruster Applications. *Journal of MicroElectroMechanical Systems*, 15(5):1260–1271, October 2006.

[37] L. F. Velásquez-García, A. I. Akinwande, and M. Martínez-Sánchez. A Planar Array of Micro-Fabricated Electrospray Emitters for Thruster Applications. *Journal of MicroElectroMechanical Systems*, 15(5):1272–1280, October 2006.

[38] B. Gassend, L. F. Velásquez-García, P. Lozano, A. I. Akinwande, and M. Martínez-Sánchez. A Microfabircated Electrospray Thruster using Ridge Emitters and Ceramic-Ball Extractor Location. In *Proceedings of the 42^{nd} AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2006.

[39] Marc D. Paine and S. Gabriel. A Micro-Fabricated Colloidal Thruster Array. In *Proc. 37^{th} AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, Washington, DC, July 2001. AIAA. 2001-3329.

[40] Marc D. Paine, S. Gabriel, C. G. J. Schabmueller, and A. G. R. Evans. Realization of very high voltage electrode-nozzle systems for MEMS. *Sensors and Actuators A*, 114:112–117, 2004.

[41] Mark D. Paine. Test of a Digital Colloid Thruster for Precise Thrust Throttling. In *Proceedings of the 29^{th} International Electric Propulsion Conference*, November 2005.

[42] M. S. Alexander, J. Stark, and K. L. Smith. Electrospray Performance of Microfabricated Colloid Thruster arrays. *Journal of Power and Propulsion*, 22(3):620–627, May–June 2006.

[43] Jijun Xiong, Zhaoying Zhou, Dong Sun, and Xiongying Ye. Development of a MEMS based colloid thruster with sandwich structure. *Sensors and Actuators A*, 117:168–172, 2005.

[44] Jijun Xiong, Dong Sun, Zhaoying Zhou, and Wendong Zhang. Investigation of the onset voltage for the design of a microfabricated colloid thruster. *IEEE/ASME Transactions on Mechatronics*, 11(1):66–74, 2006.

[45] Eric H. Cardiff, Brian G. Jamieson, Peter C. Norgaard, and Ariane B. Chepko. The nasa gsfc mems colloidal thruster. In *Proceedings of the 40th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2004.

[46] Roger X. Lenard, Stanley H. Kravitz, and Martin Tajmar. Recent Progress in Silicon-Based MEMS Field Emission Thrusters. *American Institute of Physics Conference Proceedings*, 746:954–964, February 2005.

[47] Iulia M. Lazar, Jakub Grym, and Frantisek Foret. Microfabricated devices: a new sample introduction approach to mass spectrometry. *Mass Spectrometry Reviews*, 25:573–594, 2006.

[48] Larry Licklider, Xuan-Qi Wang, Amish Desai, Yu-Chong Tai, and Terry D. Lee. A micromachined chip-based electrospray source for mass spectrometry. *Analytical Chemistry*, 72:367–375, 2000.

[49] Steve. Arscott, Severine Le Gac, Christian Druon, Pierre Tabourier, and Christian Rolando. A planar on-chip micro-nib interface for NanoESI-MS microfluidic applications. *Journal of Micromechanics and Microengineering*, 14:310–316, 2004.

[50] Jin-Sung Kim and Daniel R. Knapp. Miniaturized multichannel electrospray ionization emitters on poly(dimethylsiloxane) microfluidic devices. *Electrophoresis*, 22:3993–3999, 2001.

[51] Wang-Chou Sung, Honest Makamba, and Shu-Hui Chen. Chip-based microfluidic devices coupled with electrospray ionization-mass spectrometry. *Electrophoresis*, 26:1783–1791, 2005.

[52] Patrick Griss, Jessica Melin, Johan Sjödahl, Johan Roeraade, and Göran Stemme. Development of micromachined hollow tips for protein analysis based on nanoelectrospray ionization mass spectrometry. *Journal of Micromechanics and Microengineering*, 12:682–687, 2002.

[53] Gary A. Schultz, Thomas N. Corso, Simon J. Prosser, and Sheng Zhang. A Fully Integrated Monolithic Microchip Electrospray Device for Mass Spectrometry. *Analytical Chemistry*, 72:4058–4063, 2000.

[54] Keqi Tang, Yuehe Lin, Dean W. Matson, Taeman Kim, and Richard D. Smith. Generation of Multiple Electrosprays Using Microfabricated Emitter Arrays for Improved Mass Spectrometric Sensitivity. *Analytical Chemistry*, 73:1658–1663, 2001.

[55] C. M. Waits, N. Jankowski, and B. Geil. Scalable Electrospray Components for Portable Power Applications using MEMS Fabrication Techniques. In *Proceedings of the 25^th Army Science Conference*, November 2006.

[56] P. Campbell, K. Jones, R. Huber, K. Horch, and R. Normann. A Silicon-Based, Three-Dimensional Neural Interface: Manufacturing Processes for an Intracortical Electrode Array. *IEEE Transactions on Biomedical Engineering*, 38(8):758–768, August 1991.

[57] P. D. Prewett and G. L. R. Mair. *Focused ion beams from liquid metal ion sources*. Research Studies Press, 1991.

[58] Luis F. Velásquez-García. *The Design, Fabrication and Testing of Microfabricated Linear and Planar Colloid Thruster Arrays*. PhD thesis, Massachusetts Institute of Technology, May 2004.

[59] A. Bailey. Temperature effects and capillarity in an electrostatic thruster. In *Proceedings of the AIAA 11th Electric Propulsion Conference*, 1975. 75A434.

[60] M. S. Wilm and M. Mann. Electrospray and Taylor-Cone Theory, Dole's Beam of Macromolecules at Last? *Journal of Mass Spectrometry and Ion Processes*, 136:167–180, 1994.

[61] M. S. Wilm and M. Mann. Analytical Properties of the Nanoelectrospray Ion Source. *Analytical Chemistry*, 68(1), January 1996.

[62] John B. Fenn. Method and apparatus for electrospray ionization. US patent US 6,297,499 B1, 2001.

[63] John B. Fenn. Wick injection of colloidal fluids for satellite propulsion. US patent application US 2004/0226279 A1, 2004.

[64] R. Legge, P. Lozano, and M. Martínez-Sánchez. Fabrication and Characterization of Porous Metal Emitters for Electrospray Thrusters. In *Proceedings of the 30th International Electric Propulsion Conference*, September 2007.

[65] P. Lozano and M. Martínez-Sánchez. Ionic Liquid Ion Sources: Suppression of Electrochemical Reactions using Voltage Alternation. *Journal of Colloid Interface Science*, 280(1):149–154, December 2004.

[66] Bruce R. Munson, Donald F. Young, and Theodore H. Okiishi. *Fundamentals of Fluid Mechanics*. Wiley, 2005.

[67] V. M. Starov, S. R. Kostvintsev, V. D. Sobolev, M. G. Velarde, and S. A. Zhdanov. Spreading of Liquid Drops over Dry Porous Layers: Complete Wetting Case. *Journal of Colloid and Interface Science*, 252:397–408, 2002.

[68] Jose Bico, Uwe Thiele, and David Quere. Wetting of Textured Surfaces. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 206:41–46, 2002.

[69] T. Garza, P. Lozano, L. F. Velásquez-García, and M. Martínez-Sánchez. The Characterization of Silicon Wettability and Properties of Externally Wetted Microfabricated Electrospray Thruster Arrays. In *Proceedings of the 29th International Electric Propulsion Conference*, November 2005.

[70] T. Garza, P. Lozano, L. Velásquez-García, and M. Martínez-Sánchez. Propellant Spread over Silicon Electrospray Surfaces. In *Proceedings of the 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2006.

[71] Peter A. Thompson and Mark O. Robbins. Simulations of contact-line motion: Slip and the dynamic contact angle. *Physical Review Letters*, 63:766–769, August 1989.

[72] David Jacqmin. Contact-line dynamics of a diffuse fluid interface. *Journal of Fluid Mechanics*, 402:57–88, 2000.

[73] L. A. Romero and F. G. Yost. Flow in an Open Channel Capillary. *Journal of Fluid Mechanics*, 322:109–129, 1996.

[74] Joan Fuller, Richard T. Carlin, and Robert A. Osteryoung. The Room Temperature Ionic Liquid 1-Ethyl-3-methylimidazolium Tetrafluoroborate: Electrochemical Couples and Physical Properties. *Journal of the Electrochemical Society*, 144(11):3881–3886, November 1997.

[75] Paulo Lozano. Private Communication, 2007.

[76] Warren C. Young, Richard G. Budynas, and Raymond J. Roark. *Roark's formulas for stress and strain*. McGraw-Hill Professional, 2001.

[77] S. Senturia. *Microsystem Design*. Springer, New-York, 2000.

[78] Kuo-Shen Chen, Arturo Ayor, and S. Mark Spearing. Controlling and Testing Fracture Strength of Silicon on the Mesoscale. *Journal of the American Ceramics Society*, 83(6):1476–1484, June 2000.

[79] H. C. Miller. Surface flashover of insulators. *IEEE Transactions on Electrical Insulation*, 24(5):765–786, 1989.

[80] Shawn W. Walker and Benjamin Shapiro. Modeling the Fluid Dynamics of Electrowetting on Dielectric (EWOD). *Journal of Microelectromechanical Systems*, 15(4):986–1000, August 2006.

[81] Todd M. Squires and Stephen R. Quake. Microfluidics: Fluid physics at the nanoliter scale. *Reviews of modern physics*, 77:977–1026, 2005.

[82] R. J. Shul and S. J. Pearton, editors. *Handbook of advanced plasma processing techniques*. Springer, 2000.

[83] Daniel L. Flamm. Mechanisms of silicon etching in fluorine- and chlorine-containing plasmas. *Pure and Applied Chemistry*, 62(9):1709–1720, 1990.

[84] F. Lärmer and A. Schilp. Method of Anisotropically Etching Silicon. US patent 5,501,893, German patent DE4241045, 1996.

[85] H. Blum. A transformation for extracting new descriptors of shape. In W. W. Dunn, editor, *Proc. Symp. Models for the Perception of Speech and Visual Form*, pages 362–380. MIT Press, 1967.

[86] P. E. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.

[87] Henri Jansen, Meint de Boer, Rob Legtenberg, and Miko Elwenspoek. The black silicon method: a universal method for determining the parameter setting of a fluorine-based reactive ion etcher in deep silicon trench etching with profile control. *Journal of Micromechanics and Microengineering*, 5:115–120, 1995.

[88] M. Bengtsson, J Drott, and Th. Laurell. Tailoring of porous silicon morphology in chip integrated bioreactors. *Physica Status Solidi (a)*, 182:533–539, 2000.

[89] R. J. Archer. Stain films on silicon. *Journal of Physics and Chemistry of Solids*, 14:104–110, 1960.

[90] David Mills, Mona Nahidi, and Kurt W. Kolasinski. Stain etching of silicon pillars and macropores. *Physica Status Solidi (a)*, 202(8):1422–1426, 2005.

[91] E. A. Starostina, V. V. Starkov, and A. F. Vyatkin. Porous-Silicon Formation in HF-HNO_3-H_2 Etchants. *Russian Microelectronics*, 31(2):88–96, 2002.

[92] Sai S. Prakash, C. Jeffrey Brinker, Alan J. Hurd, and Sudeep M. Rao. Silica aerogel films prepared at ambient pressure by using surface derivatization to induce reversible drying shrinkage. *Nature*, 374:439–443, March 1995.

[93] Rene Winand. Electrocrystallization – theory and applications. *Hydrometallurgy*, 29:567–598, 1992.

[94] D. Landlot. Microstructural aspects of metal and alloy films fabricated electrochemically or by sputter deposition. In M. Paunovic and D. A. Scherson,

editors, *Electrochemically Deposited Thin Films III*, pages 160–171. Electrochemical Society Proceedings, 1996.

[95] B. Ilic, D. Czaplewski, P. Neuzil, T. Stanczyk, J. Blough, and G. J. Maclay. Preparation and characterization of platinum black electrodes. *Journal of Materials Science*, 35:3447–3457, 2000.

[96] John A. Thornton. High rate thick film growth. *Annual Reviews in Material Science*, 7:239–260, 1977.

[97] N. Dechey, J. K. Mills, and W. L. Cheghorn. Mechanical fastener designs for use in the microassembly of 3d microstructures. In *Proc. 2004 ASME International Mechanical Engineering Congress and RD&D Expo*, November 2004.

[98] M. Last, V. Subramaniam, and K. S. J. Pister. Out of Plane Motion of Assembled Microstructures using a Single-Mask SOI Process. In *Proc. 13th International Conference on Solid-State Sensors, Actuators and Microsystems*, New-York, June 2005. IEEE.

[99] R. Prasad, K.-F. Böhringer, and N. C. MacDonald. Design, Fabrication, and Characterization of SCS Latching Snap Fasteners for Micro Assembly. In *Proc. ASME International Mechanical Engineering Congress and Exposition (IMECE)*, New-York, November 1995. ASME.

[100] Hongtao Han, Lee E. Weiss, and Michael L. Reed. Micromechanical Velcro. *Journal of MicroElectroMechanical Systems*, 1(1), March 1992.

[101] H. Hedler and B. Vasquez. Self-adhering chip. US patent 6,756,540 B2, June 2004.

[102] R. M. Bostock, J. D. Collier, R.-J. E. Jansen, R. Jones, D. F. Moore, and J. E. Townsend. Silicon Nitride Microclips for the Kinematic Location of Optic Fibers in Silicon V-Shaped Grooves. *Journal of Micromechanics and Microengineering*, 8:343–360, 1998.

[103] L. F. Velásquez-García, A. I. Akinwande, and M. Martínez-Sánchez. Precision Hand Assembly of MEMS subsystems using meso-scaled DRIE-patterned deflection Spring Structures: An Assembly Example of an Out-of-plane Substrate Assembly. JMEMS 1456, accepted for publication, *Journal of MicroElectroMechanical Systems*.

[104] Alexander H. Slocum. Design of Three-Groove Kinematic Couplings. *Precision Engineering*, 14(2):67–76, April 1992.

[105] B. Gassend, L. F. Velásquez-García, A. I. Akinwande, and M. Martínez-Sánchez. Mechanical Assembly of Electrospray Thruster Grid. In *Proceedings of the 29th International Electric Propulsion Conference*, November 2005.

[106] Hanseup Kim and Khalil Najafi. Characterization of Low-Temperature Wafer Bonding Using Thin-Film Parylene. *Journal of MicroElectroMechanical Systems*, 14(6), December 2005.

[107] P. Amirfeiz, S. Bengtsson, M. Bergh, E. Zanghellini, and L. Brjessond. Formation of Silicon Structures by Plasma-Activated Wafer Bonding. *Journal of the Electrochemical Society*, 147(7):2693–2698, 2000.

[108] Xuanxiong Zhang and Jean-Pierre Raskinz. Low-Temperature Wafer Bonding Optimal O2 Plasma Surface Pretreatment Time. *Electrochemical and Solid-State Letters*, 7(8):G172–G174, 2004.

[109] Brian E. Roberds Sharon N. Farrens. In situ plasma wafer bonding method. US patent 6908832, June 2005.

[110] A. Weinert, P. Amirfeiz, and S. Bergson. Plasma Assisted Room Temperature Bonding for MST. *Sensors and Actuators A*, 92:214–222, 2001.

[111] Xuan Xiong Zhang and Jean-Pierre Raskin. Low-Temperature Wafer Bonding: A Study of Void Formation and Influence on Bonding Strength. *Jounral of MicroElectroMechanical Systems*, 10(2), April 2005.

[112] Christine H. Tsau, S. Mark Spearing, and Martin A. Schmidt. Fabrication of Wafer-Level Thermocompression Bonds. *Journal of Microelectromechanical Systems*, 11(6):641–647, December 2002.

[113] Kishori Deshpande. Private Communication, April 2007.

[114] M. A. Guillorn, A. V. Melechko, V. I. Merkulov, D. K. Hensley, and M. L. Simpson. Self-aligned gated field emission devices using single carbon nanofiber cathodes. *Applied Physics Letters*, 81(19):3660–3662, 2002.

[115] Paulo Lozano and Manuel Martínez-Sánchez. On the dynamic response of externally wetted ionic liquid ion sources. *Journal of Physics D: Applied Physics*, 38:2371–2377, 2005.

[116] Ron Jenkins, Robert William Gould, and Dale Gedcke. *Quantitative X-Ray Spectrometry*. Marcel Dekker, Inc., 1995.

[117] Vadim Khayms. *Advanced Propulsion for Microsatellites*. PhD thesis, Massachusetts Institute of Technology, June 2000.

[118] Manuel Gamero-Castano. A torsional balance for the characterization of micronewton thrusters. *Review of scientific instruments*, 74(10):4509–4514, 2003.

[119] Yu-Hui Chiu, Geraldine Gaeta, Thomas R. Heine, and Rainer A. Dressler. Analysis of the Electrospray Plume from the EMI-Im Propellant Externally Wetted on a Tungsten Needle. In *Proceedings of the 42^{nd} AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*, July 2006.

[120] Jorge Alejandro Carretero-Benignos. *Numerical Simulation of a single emitter colloid thruster in pure droplet cone-jet mode*. PhD thesis, Massachusetts Institute of Technology, February 2005.

[121] J. V. Iribarne and B. A. Thompson. On the evaporation of small ions from charged droplets. *Journal of Chemical Physics*, 64(6):2287–2294, 1976.