UNIVERSITY OF
SOUTHAMPTON

A Micro-Fabricated Colloid Microthruster:
High Voltage Electrostatic Fields on a MEMS Device

By Mark Daniel Paine

Doctorate of Philosophy

Department of Aeronautics & Astronautics

December 2002

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

SCHOOL OF ENGINEERING SCIENCES
AERONAUTICS & ASTRONAUTICS

Doctor of Philosophy

A MICRO-FABRICATED COLLOID MICROTHRUSTER: HIGH VOLTAGE ELECTROSTATIC FIELDS ON A MEMS DEVICE

By Mark Daniel Paine

Micro and nanospacecraft are opening up new, more accessible, routes into space. The vastly reduced size of these spacecraft (<100kg) brings mission costs within the reach of smaller organisations and means designers can envision programmes involving large numbers of spacecraft. However, microspacecraft pose significant challenges, perhaps the most pressing of which is the provision of a viable propulsion system. Currently, no system can provide thrusts of sufficiently low levels with high specific impulse and efficiency without suffering from high mass or excessive power consumption.

In recent years, there has been a resurgence of interest in colloid propulsion, sparked by a newly increased understanding of the underlying electrospray process. Electrospray produces sprays of charged droplets with an efficiency and specific charge that makes colloid propulsion very attractive. Traditionally colloid thrusters use hollow metal needles from which the spray is extracted using a nearby high-voltage electrode. However, this is a very problematic approach when mass is restricted, and lacks the dimensional consistency that is required across an array.

The research reported here has produced a colloid microthruster, using micro-fabrication techniques, which aims to solve those problems and enable colloid propulsion as a microthrust technology with a novel approach to thrust throttling. Large arrays of electrospray nozzles were fabricated out of silicon. The amount of fuel lost by evaporation is drastically cut because the nozzles have very small exit diameters. A new packaging system delivers fuel to the nozzles at a prescribed pressure so the liquid remains at the exit, even at high vacuum.

Extraction electrodes are situated on-chip; a silicon dioxide layer insulates them from the substrate. The electrodes are fundamental to the electrospray operation and this work is the first in which they are integrated with microfabricated nozzles. The approach benefits greatly from batch processing and the reduced dimensions mean significantly lower voltages are required. The design must create and withstand extremely high electrical fields produced by kV voltages. To ascertain the suitability of the design I tested the electrode systems and determined the major factors contributing to electrical breakdown. The breakdown strength and destructive mechanism have been discovered for differing geometries and oxide thickness, over a range of experimental conditions that simulate the predicted operational environment.

The results show that the devices can withstand over 1kV even at 10^{-2} mbar pressures of air or fuel vapour. With thicker insulation the devices have achieved breakdown voltages of 3kV. Bulk material breakdown was found to be the destructive mechanism for positively biased aluminium while a negative bias initiates erratic, and eventually damaging, surface arcs. This new knowledge of electrode-nozzle systems, which create strong fields with high voltages, clears the way for future work to fabricate devices capable of producing microthrusts using electrosprays.

Contents

List of Figures

List of Tables

List of Symbols

SYMBOL	MEANING	UNITS
α	Taylor cone semi angle	°
α_s	Mass fraction of satellite drops	
c	Spray final velocity	m/s
C_d	Dissociated concentration of the salts	mol/litre
ε	Relative dielectric constant	---
ε	Relative dielectric constant	---
F	Faraday's constant = 96,500	C/mol
γ	Surface tension coefficient	N/m^2
Isp	Specific Impulse	s
I	Current	Amps
K	Conductivity	Si/m
Λ_d	Ion mobility	Si litre/(m mol)
m	Mass	Kg
m dot	Mass flowrate	Kg/s
η	Efficiency	
q	Charge	C
Q	Volumetric flowrate	m^3/s
ρ	Mass density	Kg/m^3
T	Thrust	N
V	Potential difference across electrodes	Volts

Pressure conversion units

1bar = $1.013*10^5$Pa
1mbar = $1.013*10^2$Pa
1Torr = 133Pa

Chapter 1

Introduction

1.1 The need for micropropulsion

Microspacecraft are becoming increasingly popular, particularly among smaller organisations such as universities, because of the huge reduction in launch costs that results from their size. The cost of launching payload into orbit currently stands at ~£10,000/kg, so the very low mass of microspacecraft (< 100kg) reduces the launch costs by many orders of magnitude compared with traditionally sized spacecraft (1000kg +). Such savings often come with stiff penalties attached though; sharply decreased mass and power budgets await the subsystem designer on a microspacecraft.

These slashed allowances are particularly problematic for the propulsion subsystem as it takes up a large proportion of the spacecraft mass and is critical to the mission. Depending on the mission, the propulsion is used for attitude control and drag compensation as well as achieving the desired orbit. For very small spacecraft, the requirements some of the propulsion tasks impose can be very difficult to meet. The available power is limited to a few watts, and the restricted mass places strict limits on the amount of fuel allowed. Increasingly mission designers are developing scenarios that require propulsion systems capable of delivering microNewtons of thrust. These levels of thrust are desirable for precise positioning and attitude control and would enable groups of spacecraft to fly in formation either for use as large aperture interferometers or to enable data transmission between satellites in distributed systems.

A potential application, which illustrates the need for such a propulsion technology, is the cancellation of small disturbances for scientific missions, for example on a probe to measure gravity waves. This mission demonstrates the kind of stringent demands that the emerging micropropulsion technologies are attempting to fulfil. The requirements for the thrusters desired are low mass, fairly high Isp and efficiency, and microNewton thrusts. The thrust levels required are over a range of 1-20μN with an accuracy of 0.1μN. At present only Field Effect Electric Propulsion thrusters can provide this, however they use high voltages and

suffer from contamination issues because of the use of liquid metals as propellant. Colloid propulsion may provide the answer.

1.2 Colloid thruster development

A colloid thruster is an electrostatic propulsion technology in which an emitter produces charged liquid droplets that accelerate away from the thruster to give the thrusting reaction on the spacecraft. Applying a very strong electrostatic field to the surface of a liquid, which rests at the exit of a small-bore conducting capillary, creates the spray of droplets. The electric field deforms the liquid surface into a cone at the tip of which the field strength is sufficient to overcome surface tension and a jet of liquid emerges. This jet breaks into a spray of charged droplets, which, by falling through the electric potential, accelerate to high velocities.

Experiments on colloid emitters were first performed in the 1960's and the designs progressed quite far into development. Thrusters were designed and successfully tested over significant lifetimes. The performance obtained was in the microNewton range with good efficiencies and high specific impulses. At the time however there was no interest in microspacecraft, or the low thrust levels involved, so the performance was deemed too inadequate to warrant any further work. Furthermore, the mechanism that produces the droplets was too poorly understood for researchers to consciously improve the performance.

Recent years has seen an explosion of research interest in the process that lies at the heart of colloid propulsion. This process is electrospray and in the last ten years it has become an important analytical tool in chemistry, biology and almost all related fields. As a result of this interest a fairly basic, but extremely useful understanding of the underlying physics has been established; it is allowing propulsion engineers to re-examine colloid propulsion in a more revealing light. They now understand the factors that most directly influence the currents and therefore thrust produced by an electrospray and are aware of which operating regimes are the most stable and efficient. At the present time a number of groups worldwide are developing colloid emitters as a micropropulsion technology.

1.3 Objectives of the present research

The very low thrusts produced by a single colloid emitter demand that a thruster consists of large numbers of emitters in an array. The traditional colloid thruster uses metal pipes as emitters, which are welded to a fuel plenum and placed in front of an apertured metal electrode. However, the method results in inconsistent sprays from one emitter to the next, this inconsistency, and more importantly the thruster mass, increases with the number of emitters in an array. A batch production technique is required, in which every emitter is identical.

The electrostatics of colloid propulsion dictate that the smaller the dimensions involved the lower the voltage required to produce a spray. In traditional thrusters upto 10kV is often required – these levels of voltage are dangerous to handle and difficult to produce. Smaller dimensions are obviously desirable in that they reduce the voltage required and can result in a smaller, lighter thruster; but they also lead to a longer thruster life by vastly reducing the amount of fuel lost by evaporation.

The fabrication techniques of Micro-Electro-Mechanical-Systems (MEMS) are taken from the IC industry and are capable of producing micron-sized features in very large batches. The focus of this research has been to try to harness that capability to develop a microthruster that fulfils the potential of colloid propulsion. The work consists of the design and fabrication of a MEMS colloid thruster prototype, followed by testing which quickly identified the critical issue as electrical breakdown. Fabrication of a follow-up batch allowed us to fully investigate the breakdown and identify the mechanisms at work. The result of the work is a MEMS device that can provide the electrostatic and accelerating voltages required to produce electrospray, while operating in conditions representative of an operating thruster environment.

1.4 Thesis layout

After the brief introduction to the subject in this section, chapter 2 provides the reader with a background understanding of the principles involved. For the most part it takes the form of a review of the essential literature with the aim of providing a context and purpose for this work. Possible alternative propulsion techniques are discussed and a development history of colloid propulsion is provided right upto the present day.

Chapter 3 explains the theory behind the electrospray process and leads us to the equations for predicting the propulsive performance of a colloid thruster. The benefits of using MEMS technology to produce a colloid thruster are described. In anticipation of electrical breakdown becoming the crux of the work, this chapter attempts to sketch an outline from the complex picture presented in the literature.

The decisions behind the design of a colloid emitter are described in chapter 4 in terms of the initial design drivers and fabrication limitations. The chapter also explains how the thruster is designed to provide precision thrust control and explains the nanoelectrospray principle, which negates the need for a pump. An electrostatic model is described that predicts the strength of the electrical field near the emitter.

In chapter 5 the fabrication process is described and key aspects explained. The results of that fabrication batch are described in detail with the aim of illuminating potential pitfalls for the reader.

The experimental apparatus developed for testing and evaluation of a colloid microthruster is described in chapter 6. Particular attention is paid to the packaging of such small and fragile devices that nevertheless require complex interaction with the macroscopic world.

Chapter 7 details the first round of experiments; these demonstrated the suitability of the nozzles for delivering the liquid as a stable liquid surface at their exits. The tests also highlighted what was to become the main issue, the problem of electrical breakdown. The main experiments, which produced a new insulation batch and investigated their breakdowns, are also described in chapter 7.

The results of those experiments are presented and discussed in chapter 8. They are interpreted in the context of discovering the breakdown mechanism and values for the mean breakdown voltage are shown. The validity of these experiments as a benchmark for the complete thruster devices is asserted with supportive reasoning.

In chapter 9 the main conclusions are drawn from the work and the suitability of the thruster design is assessed, both in part and totality. Recommendations for future work are presented.

Chapter 2

Background and Literature Survey

2.1 Alternative thruster options

With the growing interest in micro and nanosatellites at space agency, industry and university levels, there is an increasing demand for high performance yet very low thrust propulsion systems. The thrust levels may be required for any or all of station-keeping, attitude control and orbit changing. Each of these is extremely mission dependent although some general comments can be made. Typically station-keeping requires a total velocity change of 5 – 50m/s per year and orbit raising can range all the way upto 7000m/s total [1], high specific impulses are desirable. Attitude and fine-orbit control will typically require 1μN to 1mN [2] or, to look at it another way, impulse bits less than 10^{-6}Ns [3]. The table below [4,11,13] gives a broad overview of the types of thrusters generally considered for use on small spacecraft.

Thruster type	Isp (s)	Typical minimum thrusts (N)
Cold gas	50-75	0.05
Monopropellant	150-225	0.01
Bipropellant	335-450	1
Resistojet	150-500	5E-3
Arcjet	400-2000	2E-3
Ion thruster	2000-6000	1E-2
Plasma	1500	2E-5

TABLE 2.1: Possible thrusters for small spacecraft

All of the above have a long history of design and test, but this is restricted in the main to relatively large devices fabricated with traditional techniques. More and more however, researchers are re-examining each of the thruster types and attempting to miniaturise the devices using new, innovative techniques. These developments are addressed where applicable. Each of these will be dealt with in turn and some new alternatives that are perhaps more promising will be assessed.

2.1.1 Chemical thrusters

This rather broad category includes the first three thruster types in the table above. In cold gas thrusters the gas is stored at high pressure and expands through a nozzle. These are commonly used as they have good accuracy but fairly low impulse bits. The latest design by MOOG [5] produces an I-bit of 0.1mNs with thrust at 4.5mN and represents the most advanced design of this type. The specific impulse of these thrusters is very low however. A notable project is reported [6] in which a cold gas thruster complete with gas handling (sensors, valves, filters etc.) and micronozzles has been made entirely of micro-machined silicon wafers fusion bonded together. No performance results have yet been seen.

Monopropellant chemical thrusters most often use hydrazine, which passes over a heated catalyst. Bipropellant rockets use the combustion caused when a fuel is oxidised, the resulting high-pressure gas expands through a nozzle. Because of the high energies involved the thrusts can be very high, but this is a disadvantage here. To reduce the thrusts much smaller chambers and nozzles are required; the materials used in micro-fabrication can suffer at the extreme temperatures caused by either the heated catalyst of monopropellant or the energetic reactions of bi-propellant systems. Several research groups are attempting to fabricate micro-nozzles and most notable are the nozzles micro-machined in silicon [7]. So far however the very small throat dimensions used seem to result in large viscous losses [8]. Specific impulses are also quite low with the highest values being achieved only by advanced propellants such as CNH_5, N_2O_4 or ClF_3, but the higher temperatures they produce may only exacerbate the material's high temperature problem. It is expected that micro-fabricated nozzles will suffer performance losses because of the inefficiency of their usually flat or square profiles. In addition, limits exist to the shrinking because of the possibility of flame quenching by boundary layers [9].

However, workers [10] have successfully manufactured a very ambitious design in which six silicon wafers are deep ion etched and then bonded together with an alignment accuracy of 3μm. The design is a liquid cooled bi-propellant rocket weighing 1.2g, which delivers 1N of thrust. Tests are continuing.

2.1.2 Electric thrusters

This is an even broader range of thruster types and covers any propulsion option that does not derive the majority of its propulsive energy from the chemical energy of an on-board propellant. The advantage of such systems is that instead, they harness the solar energy

(almost) freely available in space and convert it into electrical power. This can drastically reduce the amount of propellant the spacecraft needs to carry and it need not be so chemically active. Generally, the specific impulse of electric propulsion is much higher; while the thrust is lower than most of the chemical alternatives. Although somewhat at a disadvantage in the traditional area of large spacecraft with high thrust needs, electric thrusters really come to the fore in micropropulsion.

RESISTOJETS

The most basic application of electrical power to aid performance is that of the resistojet. In this thruster, a gaseous propellant (often H_2 or N_2) passes over an electrical heater and rises in temperature. The specific impulses are quite low but this is mostly because most materials cannot withstand the high temperatures (>1600°C) required to achieve Isp > 600s with hydrogen [11]. Power is usually quite high (mainly due to the unfortunate link between high Isp and low efficiency) and lifetime is limited by the high temperatures used. Despite this, researchers have produced micro-fabricated resistojets with which they hope to achieve thrusts as low as 1μN at acceptably low power levels [2]. This illustrates a trade-off that is common among research into micro-thruster technology; since few thrusters promise the low thrusts required at sufficiently low power levels while maintaining the high Isp needed, it is often wise to compromise on one aspect – in this case Isp. At present, no performance results have been published.

ARCJETS

These work by initiating an arc between a concentric upstream cathode and a downstream anode that is also a nozzle. The performance of arcjets depends radically on the propellant used but some general comments can be made. Keeping the hottest gas away from the chamber walls improves the lifetime over that of resistojets. Specific impulse is also higher (partly for the same reason). Fairly low thrust can be produced, but power is often prohibitively high (ranging from 0.5 to 100's of kW) and efficiency rarely passes 40% [11]. As with resistojets, they do have a proven service history that is quite considerable (compared to other electric propulsion methods). However, a low power arcjet has been reported which has an insulating boron nitride insert. Thrusts of 1mN at power levels of 4W were achieved. Isp was only ~100s but this increased to ~350s for power levels of ~25W.

PULSED PLASMA THRUSTERS

In this method, a spark triggers a high current pulse discharge across the surface of a bar of solid propellant, usually Teflon. The spark evaporates a thin layer of the Teflon creating a plasma that is electromagnetically accelerated to high velocities, causing the thrust reaction. These thrusters can produce very small impulse bits with good Isp (upto 1000s) so are well suited to precise positioning missions. They have a simple operation, robust construction and are relatively low cost, with a flight history. However, they are high power (upto 1kW, although this is peak power, the pulsed operation means average power is much lower) with very low efficiencies in the 0.05 – 0.1 range [13, 14]. The most suitable thruster seems to be the coaxial PPT tested at NASA Lewis [15] which has shown itself capable of 180μN thrusts with Isp = 600s, in a compact device.

More recent work from Japan has reported performances of 16.5μNs impulse bits with Isp > 900s at powers as low as 5W for a total mass of 1.2kg [16]

A similar approach is the compact vacuum-arc micro-thruster [17], which uses inductive storage rather than the capacitors that make the PPUs of PPTs so heavy. The CVAM uses a vacuum arc between two concentric metal electrodes that produces a plasma from the ablated material of the cathode. The performance reported gives 1000-3000s of Isp but at powers of 1-100W, the thrust to power ratio is ~ 2×10^{-5}N/W.

ION ENGINES

Ion engines are the most developed of the electric propulsion methods and use a biased grid to electrostatically accelerate an ionised gas. The gas can be ionised in a variety of ways, the most usual being electron bombardment, but other designs use contact ionisation or radio frequency ionisation. The thrusters can be very high efficiency (upto 75%) with high Isp (> 4000s). However, they are large, bulky structures, which are difficult to miniaturise, and have complex, heavy PPUs.

A good deal of research focuses on adapting ion thrusters to the low thrust levels considered here. Staying with the conventional method of ionisation, several groups have analysed how performance changes with decreasing dimensions and powers. An 8cm beam, 100W thruster has been designed [18] with the aim of achieving 4mN thrust at ~50% efficiency and 2000s of Isp. These thrust levels are still too high though and further scaling of the chamber dimensions results in several problems reported in the work of Yashko [19]. The main problem in miniaturising ion thrusters is their need for gas ionisation; at smaller scales

proportionally more electrons and ions are lost to the chamber walls. With reduced chamber sizes the magnetic fields, which confine the electrons so that a substantial proportion of the propellant is ionised, become impractical. It was shown that using the strongest rare-earth permanent magnets the magnetic field is insufficient to continue confining the electrons efficiently when the chamber size is brought much below 1cm. The resulting ion production cost is greatly increased. They are pursuing an alternative geometry, which uses an array of linear chambers instead of the traditional cylinder. Their analysis showed this to be an advantage as the ion cost reduced with the number of chambers used. The design also bypasses another problem of miniaturised ion thrusters, the possibility of electrical breakdown across the acceleration grids [1], as it uses open areas with electrodes on the chamber walls.

An alternative approach is the use of a hemispherical discharge chamber as a hollow cathode turned into an ion thruster [20]. Conventional ion thrusters use hollow cathodes as a source of the electrons, which then enter the chamber and ionise the propellant. In this thruster, the whole of the 1cm diameter hemispherical chamber acts as an anode opposite a 1mm hole in an aluminium hollow cathode. The discharge that takes place between them ionises the argon propellant. The hope is that the discharge will produce levels of ionisation that are higher than the discussed scaled-down conventional chambers, without the need for magnetic containment. The method promises ionisation rates of 2% operating on xenon and theoretically should be scalable to sizes that can achieve the thrust levels required.

HOLLOW CATHODE THRUSTERS

Several research groups have been investigating the use of hollow cathodes on their own as a source of thrust. This direction is being pursued since it has been observed that hollow cathodes produce significant emissions of high-energy ions, which may produce thrust [21]. The work is in the early stages but one piece of research reports calculated thrusts of ~ 400μN with > 500s of Isp [22]. This uses 50W of power and has an overall efficiency of ~ 2%. However, the authors attribute the low efficiency to the use of improperly optimised ion optics on the extraction grids.

Another approach is the use of micro-fabrication techniques. Workers in the University of Southampton's Astronautics research group is developing a miniaturised ion thruster that uses a miniaturised hollow cathode design (previous attempts to use field emitters were unsuccessful), but the work is in the early stages.

HALL THRUSTERS

These electromagnetic thrusters use the Hall effect to confine electrons, which ionise the xenon propellant gas. The ions are then accelerated electrostatically and the reaction is transmitted to the thruster via the electrons. The efficiency is typically in the 45-55% range and Isp is good (upto 2000s [23]) but power is quite high (200-6000W) [13, 24]. Thrusts can be as low as 6mN [24]. The production of hall-type microthrusters suffers from problems associated with the need for gas ionisation (discussed above). Although a low power thruster has been built and tested [25], which achieves low power (50W) and lower thrusts (3mN), efficiency was reduced to around 20% and lifetime was considered an issue.

FIELD EFFECT ELECTRIC PROPULSION (FEEP)

FEEP uses an extremely strong electrostatic field to draw ions directly from the surface of a liquid metal. Developed extensively in Europe [26] the design uses liquid caesium held at the tip of a narrow slit (a few μm high). The ions are then dragged out by the field created by applying upto 10kV on an electrode situated a mm or so away. FEEP fulfils the requirements of thrust perfectly and easily reaches down to thrusts of 1μN. The efficiency is very high as is the specific impulse, reaching upto 10,000s, but this is possibly its severest drawback. The high Isp is derived from the exceedingly high specific charge (q/m) of the metal ions; a side effect of this is that it produces high currents but with little thrust. Consequently the thrust to power ratio is quite low, typically reaching $2x10^{-5}$ N/W. This limits the thrusters to low ΔV and the very low thrust applications of micropropulsion. Other concerns include the very high voltages used, problems with initial wetting of the emitter slit by the propellant and the possibility of contamination of spacecraft surfaces by the liquid metal efflux. Designs reported recently use liquid indium [27], which wets the surface of a sharpened tungsten needle. Tests have revealed performances in the 1-100μN range with very low levels of thrust noise.

2.1.3 Specially designed micro-fabricated thrusters

Several research groups have looked at the problem of micropropulsion from a different angle. Instead of trying to copy designs of conventional thrusters in miniature, they are attempting to use the strengths and advantages of micro-fabrication techniques and materials to design a new kind of thruster. Usually such designs are tailored to a more specific region of the propulsion spectrum. All of the following have only reported early stages of the development so far but the results look promising, at the very least because of the shift in the way the problem of micropropulsion is approached.

VAPORISING LIQUID MICRO-THRUSTER

In this design by the JPL [3], a liquid propellant (ammonia or hydrazine) is pressure fed into a chamber between two silicon wafers. On the surfaces of the channel are heater strips, which use electrical power to vaporise the liquid. The vapour then expands through a converging-diverging nozzle to produce thrust. The first prototypes made had square shaped nozzles with an expansion angle of 70°, which the authors accept is far from ideal but is sufficient at this stage. Since it is designed as an attitude control thruster, the authors feel that the predicted low Isp (~100s) will be adequate. The thrust levels aimed for are 500μN. To be able to achieve the low I-bits required for this application the thruster still requires the development of a fast acting valve.

DIGITAL MICROPROPULSION

The Aerospace Corporation has taken to the use of micro-fabrication with enthusiasm. In this new microthruster concept [28], they use the intrinsic batch nature of micro-fabrication to great advantage. Their design eliminates the need for complex propellant storage and feed mechanisms giving a small, lightweight and fully integrated thruster on a chip. They use a sandwich of silicon and glass layers, which contain an array of small plenums each with a rupturable diaphragm on one side. Each plenum contains either an inert gas or a combustible propellant. The pressure in the plenum is raised, by either heating or ignition of the fluid, until it is high enough to burst the diaphragm. The escaping high-pressure fluid imparts an impulse to the thruster. In this way, each plenum can deliver one impulse bit of a known and reproducible size. By addressing each of the plenums individually, impulses of a specified size can be built up from a number of plenum shots. Although each plenum is only useable once, the batch processes used allow a very large number of them to be manufactured easily and at low cost on one chip. The paper reports I-bits of 100μNs with power levels peaking at 100W; this was using lead styphnate (a solid) as propellant. Much development remains to be done in which efforts will concentrate on improving the amount of propellant consumed (presently ~10%) and decreasing the I-bits.

MEMS MEGA-PIXEL MICROTHRUSTER

Honeywell and Princeton have fabricated a similar thruster [29] to that described above, in which ¼ of a million addressable, single-shot thrusters are fabricated in a 33mm square array. Each shot consists of a charge of lead styphnate, which is heated until it explodes; the energy

from the explosion ignites a nitrocellulose mixture in an overhead cavity. No results can be found but the expected performance was Isp of 100-300s, with each shot producing 5 to $20x10^{-6}$ Ns of impulse at a power of 10mW each.

MICROCHIP LASER THRUSTER

An interesting piece of work [30] produced by MIT Lincoln labs is the pulsed microchip laser thruster. In this work a small diode laser pumps an on-board microchip laser and the pulses ablate propellant material from the surface of a transparent tape, which is advanced by small motors. Using aluminium they estimate measured thrust levels of 3nN - 3μN. Specific impulses have not been measured but power is of the order of watts, and first estimates of mass are < 0.5kg.

2.2 Colloid thruster literature

Development of colloid thrusters was halted 25 years ago when experiments showed that although the technology performed well it could not be flown. This was because achieving sufficiently high Isp (~1000s) required impractically high voltages (10's to 100's of kV) to obtain high enough specific charge. At that time too little was known about the electrospraying process for researchers to be able to lower the voltage to a more manageable level while maintaining high performance. In the past decade however, the electrospraying process has received a great deal of attention. This is because of its importance in the areas of paint spraying and mass spectrometry as a method of creating monodisperse aerosols and producing ions respectively. The electrospraying research work carried out has produced a better understanding of the process and developed several laws, enabling a new interest in colloid propulsion. The amount of literature on electrospraying is very large and a complete survey of it would be a document in itself. The pertinent information is best presented as a summary of the relevant physics and how they apply themselves to the problem of microthrust production. This is done in section 3.4; here the focus is on the spacecraft application of electrospraying.

2.2.1 Previous thruster experiments

In 1967 Perel et al [31] reported the results of their experimental work on a thruster consisting of an array of 73 platinum needles with 0.2mm ID. The thruster was designed to be bipolar to avoid the need for a neutraliser. The propellant used was glycerol as it can dissolve dopants very well. Other liquids were also investigated but for liquid caesium and pure H_2SO_4 the

specific charge values were too high. By using a variety of emitter shapes, they found that the taper of the needle tips played an important part. Those with shallower tapers gave better focusing, higher specific charge and needed a lower voltage for the process to start.

By varying the voltage, V, and mass flow, Q, to the thruster they found trends for the variation of specific charge. Specific charge is an important parameter as it relates directly to the exhaust velocity. They found that q/m increases with both the applied voltage and the conductivity of the liquid but decreased with flow rate. The specific charge distribution of the droplets plays an important part in determining the efficiency of the thruster. Ideal efficiency occurs when all particles have the same q/m. The experiments showed that as the average q/m of the spray increased above a certain level the efficiency decreased, this is because larger values of q/m are caused by a higher proportion of ions (which have q/m much greater than droplets) and therefore the distribution is wider. This experiment also demonstrated good performance with thrusts of 350μN, Isp of 700s and efficiency of 50 to 80%.

In 1973 Mahoney, Daley and Perel published two papers [32,33] relating to the study and development of colloid thrusters using annular slit emitters. The emitters had an inner extractor (-2 to –5kV) at the centre of the annulus (held at +15kV) and were surrounded by a trap electrode (-500V) and an outer extractor. The inner extractor was designed to focus the beam and by moving it, the highest performance position was found to be slightly forward of the emitter. The trap electrode was used to shield the emitter from electrons, which can erode the tip and distort the field.

The dependence of efficiency on the q/m distribution shown by Perel et al was verified in these experiments. They found the emitters produced their optimal performance (thrust=130 μN, Isp=1500s, η=0.8) when the majority of the particles in the spray were heavy, charged droplets.

They tested arrays of three emitters over a 113 hr period. These tests showed the importance of the fluid temperature, as wide variations in the performance were measured and could be attributed largely to temperature effects. The most significant of these was thought to be the temperature effect on fluid viscosity, which altered the fluid flowrate away from the nominal calibration of their feed system.

A separate study was performed by Kidd and Shelton [34] of TRW systems group, in which a 36-needle array was operated for 4350 hrs as part of a long-term qualification test. Again, platinum needles were used with a 0.127mm ID, with a negatively biased shield. The need

for close tolerances was demonstrated as unless these were controlled the performance was severely degraded. Learning from the lessons of previous experiments this thruster incorporated a closed loop temperature and feed control system. As a result, the thruster exhibited good performance (thrust=350μN, Isp=1400s, η=69%) but this was degraded over the duration of the test. Post-test inspections revealed the causes to be tar build up and needle tip erosion, both of which occur because of electron bombardment. By studying the build-up and erosion on needles compared to the relative height of the shield around it, it was found that those with shields slightly above the needle tip fared better. A relation between those needles (thinner) that produced higher q/m particles and higher rates of erosion was also inferred.

2.2.2 More recent work

Due to the growing desire for a low thrust propulsion method, several older, previously unfashionable concepts have seen a resurgence of interest. Primary among them, fuelled by a greater understanding of the underlying physics, has been colloid propulsion.

Prof. Martinez-Sanchez at MIT is largely responsible for this new interest and has published an updated review of the field [35]. In it he describes some brief experimental work as a result of which he finds new propellants capable of producing Isp of ~1000s with voltages less than 5kV. The identification of new propellants is a crucial breakthrough and was only possible due to the recent research into electrospraying. Through that research the importance of several propellant properties was identified, the most important being conductivity and surface tension. He identifies formamide's ability to dissolve large concentrations of ionic salts and therefore achieve the high conductivities needed for high specific charges.

This work has progressed and has caught the eye of a commercial company, Busek of Massachusetts. Several recent papers report the progress [36,37]. A thruster system has been developed that operates in the microNewton range. The thruster consists of 57 steel needles brazed to a plate over a propellant cavity. The needles are 1 cm long and have an ID of 30μm. They are aligned with an extractor grid and above that an accelerator grid, which are separated by ceramic spacers. The extractors are placed ~ 0.5mm away from the needles and the apertures are ~1mm diameter. The voltage required to start spraying with this configuration was ~ 1600V and 2kV of extra acceleration voltage was applied. The fuel used was formamide doped to a conductivity of 0.5S/m. Mass flowrates of 1 to 20×10^{-11} m^3/s were used and thrust levels of 20-200μN with Isps of 400 to 100s respectively were obtained. This shows the drastic decrease in Isp with Q, predicted in the electrospray theory of chapter 3.

The thruster was designed to reach 1-20μN thrusts but as the authors point out this could easily be achieved by using higher conductivity propellants. The performance values were obtained using Time of flight spectrometry. The flowrates were controlled by varying the pressure applied to liquid at the end of a thin pipe. However, some tests were done using potential to control flow rate, this will be discussed more later. This work also tests the use of carbon nanotube field-emitters as a neutralisation source.

A summary paper has demonstrated the Russians continued interest in the concept [38] and they present their state of the art model. Consisting of a monoblock of 20 emitters the published results show that it is still operating at relatively high voltages (15kV upwards) with thrust as low as 25μN per emitter and Isp of 1000s. They do not mention the propellant used.

Work at Stanford is also progressing well with a plan fully underway to perform a flight test of their thruster design [39]. The thruster, comprising 100 hypodermic pipes welded to a common base plate, uses glycerol doped with NaI as propellant. It will fly on the University's Emerald nanospacecraft, scheduled for launch in late 2003. Recent reports of measured performance focus on a single emitter thruster core. The needles used are stainless steel with an ID of 50μm and flow is controlled by a modified insulin pump of the motor controlled plunger type. The performance reported is 1 to 5μN for applied voltages of 2.5 to 4.4 kV but Isp is only 60 –200s for the same V range. These relatively poor values are probably due to the use of a low conductivity propellant. The group are collaborating with researchers at Phrasor scientific who were among the early pioneers of colloid thrusters in the 1970's. They report work [40] on a slit thruster design (similar to FEEP) which has a gap of 75μm filled by a rake substrate which acts as a capillary fuel feed and anchoring for emission sites. The test used glycerol with NaI and obtained thrusts of 3 to 44μN with Isps of 164 to 1110s. The applied voltages were very high with a range of 5.3 to 7.8kV needed to obtain this performance.

More recent work [41] has proposed a novel approach for electrospray acceleration using AC voltages. The charged droplets are linearly accelerated by successive electrodes, which have alternate square waves timed to the frequency of droplet production. By using five electrodes, 1000s of specific impulse should be obtained with only 1kV. However, the authors note that it may be difficult to provide the very high frequencies required and that a spray of uniform specific charge would be needed.

2.2.3 Relevant micro-fabrication work

This research focuses on the effort to miniaturise colloid emitters (the advantages of which are explained in a later section) using micro-fabrication methods. It is therefore wise to see whether anything similar has been fabricated before and to determine the success or otherwise of such attempts.

FIELD EFFECT ELECTRON EMITTERS

The earliest research, which resulted in the fabrication of a similar piece of hardware, is the field-effect electron emitter. Pioneered by Spindt [42] this device emits electrons when a voltage of ~60-100V is applied to a circular aperture gate electrode situated about 1μm away from the sharp pointed tip. The combination of the close proximity of electrodes and the sharpness of the tip produce an extremely strong electric field at the tip, which drags the electrons out of the material. Spindt's tips were made from molybdenum but later efforts used silicon and even polycrystalline diamond [43] because of the lower work functions of these materials, which allow for lower onset voltages. The relevance of these devices lies in their application of strong electrical fields via circular electrodes. They demonstrate the ability of small gaps to hold off strong electrical fields without breakdown, although recent tests done by our group have shown them to fail as pressure is reduced from high vacuum up to ~ 10^{-3} torr [44]. They also demonstrate the use of SiO_2 as an insulator between cathode and gate. The maximum voltage taken by a 1μm thick layer of SiO_2 in the context of field emitters is ~100V.

GRIDS FOR MICRO-ION ENGINES

Some work has been reported [1] in which the authors micro-fabricated structures designed to assess the feasibility of using SiO_2 as an insulator in an ion thruster acceleration grid. The structures were silicon substrates with a polysilicon layer deposited onto them, these acted as the base electrodes. An aluminium counter-electrode was isolated from the polysilicon using a 2.7μm thick layer of low temperature oxide deposited by chemical vapour deposition (CVD). The chips were tested for electrical breakdown in atmospheric pressure, room temperature air and it was found that surface breakdown occurred at 500V for electrode separation gaps of 220μm. When breakdown was not via a surface arc, it was through the bulk of the material and occurred at 767V. This is insufficient for ion engine acceleration. However, the work suffers by the use of low quality oxide, which has half the breakdown strength of thermal oxide, but also, perhaps more significantly, the chips were tested at

atmospheric pressure. This is not representative of the operating conditions and encourages unnecessary surface breakdown.

MICROELECTRONIC FIELD IONISER

A patent filed by Spindt in 1990 [45] showed he realised the potential of the field-emitter geometry to produce ions. The patent describes a method of micro-fabricating arrays of small nozzles each centred in a circular gate aperture. The purpose of this device is however to ionise atoms of a gas which are fed through the nozzle. Though effective, the rather difficult fabrication method results from the need to have a sharp-rimmed conducting nozzle. As will be seen this requirement is lost when the use of conducting liquids is considered. The patent is useful however as it illustrates the feasibility of fabricating arrays of small nozzles and again uses thermal SiO_2 layers as insulation. Spindt states that a vacuum of 10^{-4} Torr is required for operation.

MINIATURISED LIQUID METAL ION SOURCE (MILMIS)

In a survey of liquid metal ion sources [46] the design of a MILMIS is presented. Similar in design to the Spindt patent above it is designed for use in electrospray mass spectrometry but specifically with liquid metals in mind. The limitations of this device are only seen when considering applying the technique to make a colloid emitter. The device is designed so that ionisation occurs at voltages of 100V, but this value is far too small to produce a useable Isp for our thrusters. Increasing the voltage to the levels desired would introduce significant problems to the manufacturing process used. In particular, these include the need for much greater gaps and insulator thickness, which may make the replica method of fabrication far more difficult. A further interesting point is raised however; the use of emitters operating in the opposite polarity mode can be integrated into any colloid microthruster array to provide charge neutralisation. This will be returned to later in this document.

MICRO-FABRICATED ELECTROSPRAY SOURCES.

Several groups have undertaken work to micro-fabricate electrospray sources. One fabricates hollow needles in the plane of a silicon substrate, which extend from the edge by a mm [47]. The emitters have tips that are 5μm x 10μm and are tapered to focus the electrical field. Effective electrospraying was achieved for a number of hours using just a few kV however only 3 emitters were produced along a 1cm edge.

Another group [48] has fabricated monolithic electrospray sources in which a nozzle with 10μm ID and 20μm OD stands to a height of 50μm from the surface of the chip. Made using deep ion etching from both sides this is a very effective nozzle. It is then held in front of a conventional electrospray sampling orifice (which is held at ground potential) and a syringe pump supplies liquid. Good electrospray performance results from a voltage of 1250V.

WORK ON MICRO-FABRICATING COLLOID THRUSTERS

At least three other groups are attempting to use micro-fabrication to produce colloid microthrusters. Very little has been published as yet, indeed this work produced the first publication on the subject. Part of a long iterative design process (with the same starting point as this thesis [49]) has resulted in the fabrication of a micro-fabricated colloid engine [50]. The thruster chip is approximately $30cm^2$ and consists of a fuel manifold that feeds an array of 20 emitters, all fabricated in the plane of a silicon wafer. The emitters are capillaries about 15mm in length with a 12μm square internal profile and a tapered tip. Extraction will be provided by a separate electrode joined and accurately aligned to the emitter chip by a clip structure. Fabrication of the emitters and manifold structure was successful but at the time of publication no sprays had yet been obtained. The design uses high aspect-ratio nozzles that give a large hydraulic resistance so that flow can be controlled by a head pressure provided to the fuel manifold by a pressure tank. However, this pressure is just below that required to start flow – the little extra pull will be provided by turning on the extractor.

X. Jijun of Tsinghua University has published the only other paper [51] that reports the use of the techniques of PCB manufacturing in conjunction with a PZT micropump. The thruster has an array of 300μm diameter orifices in a copper plate and this is joined to a 1.5mm thick PCB chip on the top of which is the extractor electrode. Thrusts of 1.5μN have been directly measured when voltages of 3.5kV are applied. It is with reservations that this work is considered micro-fabrication though, partly because the techniques are not strictly micro-fabrication methods, but also because the dimensions involved are not small enough.

The status of as yet unpublished work, at Queen Mary, London, has been established from personal communications. In collaboration with Rutherford Appleton laboratories they have fabricated nozzles in silicon wafers and have been performing spray tests in vacuum but it is not thought that electrodes are integrated into their design.

Chapter 3

Colloid propulsion

3.1 Introduction to colloid propulsion

Colloid thrusters use a process called electrospraying to produce very low levels of thrust. They belong to the class of microthrust electrostatic propulsion methods. In the electrospray process an electric field is used to produce a spray of charged droplets, as this mass leaves the thruster it causes a thrusting reaction. For fuel, the process uses an electrolytic liquid, which is highly conducting. Historically glycerol doped with salts has been popular and formamide doped with ionic salts (such as LiCl) may be useful. The spray is created by applying an intense electrical field to the liquid surface at the end of a capillary. The electrical forces pull on the liquid meniscus and are balanced by the surface tension of the fluid. As a result, the surface deforms to an equilibrium shape of a cone, at the tip of which the electrical field is at its maximum. The field intensity at the tip causes a jet to form, which breaks up into a spray as it streams away from the capillary. Figure 3.1 illustrates the basic experimental set-up. As with all forms of electrostatic propulsion the (positively) charged particles that leave the spacecraft must be neutralised with an equal current of electrons or negatively charged drops, to stop the spacecraft from becoming charged.

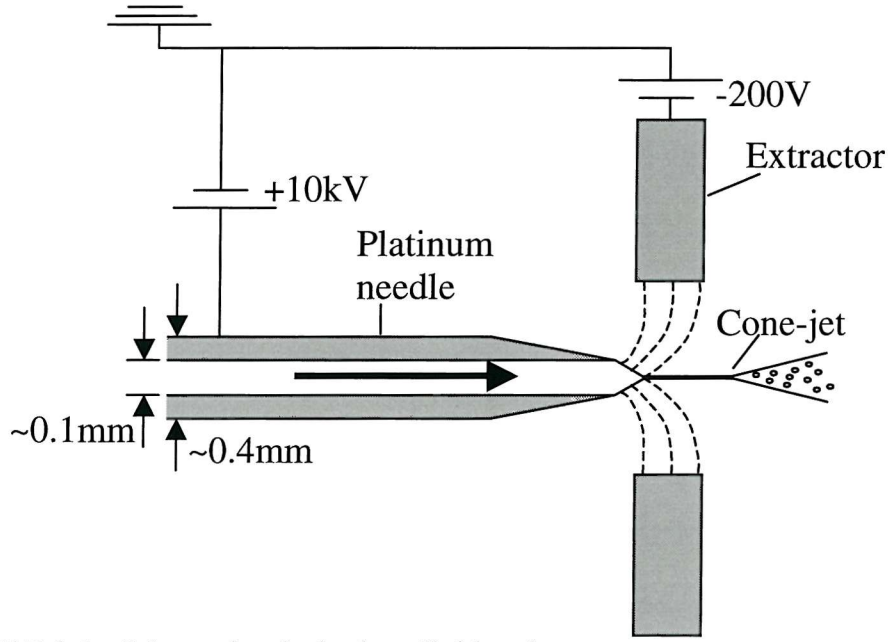

FIGURE 3.1: Schematic of a basic colloid emitter

Colloid propulsion is worthy of our attention because of the combination of suitable performance parameters. The thrust is low enough (< 1μN to a few mN) to be of use in the missions mentioned. Power is low, a few mW, because although the voltages are fairly high (1-5kV) the current is in the microampere range. Efficiencies of up to 70% can be obtained and specific impulses of 1500s have been achieved. Duty cycle is high, especially when compared to PPTs. There is a large degree of flexibility in how the colloid process is implemented. The propellant can be chosen from a wide selection of possible mixtures of almost any solvent and ionic salts. The emitter can be varied in both size and design (for instance annular slits instead of holes) and the flowrates and voltages can be varied over a (relatively) large range.

The proposed missions are increasingly leaning towards the application of miniaturisation. The colloid concept is extremely well suited to size reduction, more so than many other electric propulsion methods. In contrast to the more conventional electric thrusters (arcjets, Hall effect and Ion thrusters), colloid thrusters do not rely on the ionisation of gases; they therefore avoid the life reductions and operating complications that result from the shrinking of a gas discharge chamber. Indeed, besides the fact that the process uses very small diameter capillaries anyway, a reduction in the thruster size creates more intense electrical fields so can actually benefit the process. The miniaturisation of conventional colloid thrusters using MEMS has several major advantages, which are explained in a later section. First, it is useful to understand the physics of the core process – electrospray.

3.2 Electrospray theory

Recent research into the underlying physics of electrospraying has provided us with equations for the performance of an electrospray emitter.

The problem was first tackled by Taylor [68] who attributed the conical shape of the liquid meniscus to the balance of the electrical and surface tension forces. This occurs because the high conductivity of the liquid means the surface is (quasi) equipotential. Since the curvature of a conical section varies as 1/r, the electric field normal to the surface (balancing the surface tension force ~ γ/r) must vary as 1/sqrt(r). The potential is described by a simple spherical harmonic of the form $\phi \sim \mathrm{sqrt}(r) * Q_{1/2}(\theta)$ and since $Q_{1/2}(49.3°) = 0$ this is the angle which yields an equipotential surface. The electrospray process is restricted in this document to a

specific range of applied voltage and mass flow that creates a single cone-jet, as this appears to be the most stable regime and is better understood than regimes in which the spray is intermittent or multiple protrusions occur.

If the jet has a diameter considerably smaller than the capillary the cone rests on then it has been shown that the current through it can be predicted [69]. While independent of the electrostatics, the current has a strong dependence on the mass flow and liquid properties. Experiments [69] with several liquids (including formamide, water, ethylene glycol and triethylene glycol, all doped with LiCl) verified the following equation for the current, I (Amps):

$$I = f(\varepsilon)\left(\frac{\gamma K Q}{\varepsilon}\right)^{1/2} \tag{3.1}$$

$$so,\ as\ \frac{q}{m} = \frac{I}{Q\rho}\ then:\ \frac{q}{m} = \frac{f(\varepsilon)}{\rho}\sqrt{\frac{\gamma K}{Q\varepsilon}} \tag{3.2}$$

Where f(ε) has a value of ~18-25 for $\varepsilon > 40$. However, the function f is derived from only six liquids with conductivities (K) in the range $5*10^{-5}<K<0.83$ S/m so should be treated with caution.

Close to the base of the cone the fluid velocity is negligible and very little surface charge is convected. However, near the cone tip the velocity is high causing charge relaxation effects to set in, and a portion of the current is then carried by convection rather than bulk conduction. This begins to occur when the distance r from the apex is comparable to the electrical relaxation length $L \sim r^* = (Q\varepsilon\varepsilon_0/K)^{1/3}$. This length is important as from experimental photographs it has been seen to be ~2.25 times the diameter of the jet, so $r_{jet} \sim 0.44r^*$ which in turn can determine the size of the drops. The analysis of jet break-up by Rayleigh [70] gives us the relevant equation: $r_{drops} = 1.89r_{jet}$ yielding:

$$r_{drops} = 0.84\sqrt[3]{\frac{\varepsilon\varepsilon_0 Q}{K}} \tag{3.3}$$

The minimum flow rates that produce the cone-jet regime can be calculated by considering the fluid to be used. It is believed [71] that the specific charge must be less than would result if the positive and negative ions contained in the doping salts were to be fully dissociated:

$$\frac{q}{m} = \frac{f(\varepsilon)}{\rho}\sqrt{\frac{\gamma K}{Q\varepsilon}} = \frac{1000FC_d}{\rho} = \frac{1000FK}{\rho\Lambda_d} \tag{3.4}$$

Introducing the non-dimensional group:

$$\varphi_{\min} = \sqrt{\frac{\rho K Q_{\min}}{\gamma\varepsilon\varepsilon_0}} \quad we\ find\ \varphi_{\min} = \sqrt{\frac{\rho}{\varepsilon_o}\left(\frac{f(\varepsilon)\Lambda_d}{1000\varepsilon F}\right)} \tag{3.5}$$

and using formamide (Λ_d=15, ρ=1133 Kg/m^3. f(ε)=25, ε=111) gives φ_{min}=0.396. Rearranging for Q_{min} yields the expression:

$$Q_{\min} = \frac{0.1568\gamma\varepsilon\varepsilon_o}{\rho K} \tag{3.6}$$

Which, used with K=2.2 S/m and γ=0.062 N/m^2, gives Q_{min}= 3.8 10^{-15} m^3/s. It is clear that extremely low flowrates of the fuel are used. This high value of conductivity comes from doping formamide with the ionic salt NaI; the necessary concentration is 3.00mol/l. The conductivity of formamide can be predicted from the concentration of doping salts from the equation, K (S/m) = 1.03N$^{0.694}$. [35]

It is desirable to use flowrates as close to the minimum allowable for a variety of reasons. De La Mora [72] showed that the distribution of droplet sizes depends strongly on the flowrate. Close to Q_{min} the spray is monodisperse, but as Q increases it bifurcates and smaller "satellite" drops are produced. Further increases in Q result in further bifurcations, until at the higher flowrates the size distribution is very broad. Hartman [73] went some way to explaining the production of wide size distributions by examining the method of jet break-up. At low flowrates the jet breaks up in an axial fashion through varicose instabilities of a certain wavelength described by Rayleigh. The result is that the jet breaks into equally sized cylindrical segments which reform into drops. As the flowrate increases the jet has insufficient time to fully flow into the main drop and a portion of it snaps back to form a smaller "satellite" drop. Still higher flowrates (and hence higher currents) mean the surface charge on the jet is increased, as this happens the electric stress starts to dominate surface tension and lateral or kink instabilities begin to appear. The jet starts to whip from side to side and droplets are created with widely varying sizes.

The significance of size distribution is twofold, and both problems are related directly to the droplet specific charge. Since deflection in an electrical field depends on specific charge, a range of values makes it difficult to focus the spray. More importantly though, is the effect that size distribution has on thrust efficiency. Efficiency is the change in the kinetic energy of the exhaust beam (the useful energy) over the electrical power used and can be shown (see appendix) to be:

$$\eta = \frac{T^2}{2\dot{m}IV} = \frac{\left(\alpha_s\sqrt{\left(q/m\right)_s} + (1-\alpha_s)\sqrt{\left(q/m\right)_m}\right)}{\alpha_s\left(q/m\right)_s + (1-\alpha_s)\left(q/m\right)_m} = \frac{\langle c\rangle^2}{\langle c\rangle(c_s + c_m) - c_s c_m} \qquad (3.7)$$

where the subscripts s and m stand for satellite and main droplets respectively, and α_s is the mass fraction of the satellite drops. The graph in figure 3.2 illustrates the effect of bimodal sprays on efficiency. It can clearly be seen that the amount of smaller drops (i.e. those with higher q/m) has a significant effect. The minimum efficiency occurs at a relatively low mass fraction; for ions, which typically have q/m of 10^5 or above there only needs to be a very small proportion for the efficiency to be severely compromised. In the case of droplets, one can see from the q/m=70,000C/Kg curve that the satellite drops need only have 5 times the main drop q/m for efficiency to be decreased by 15%.

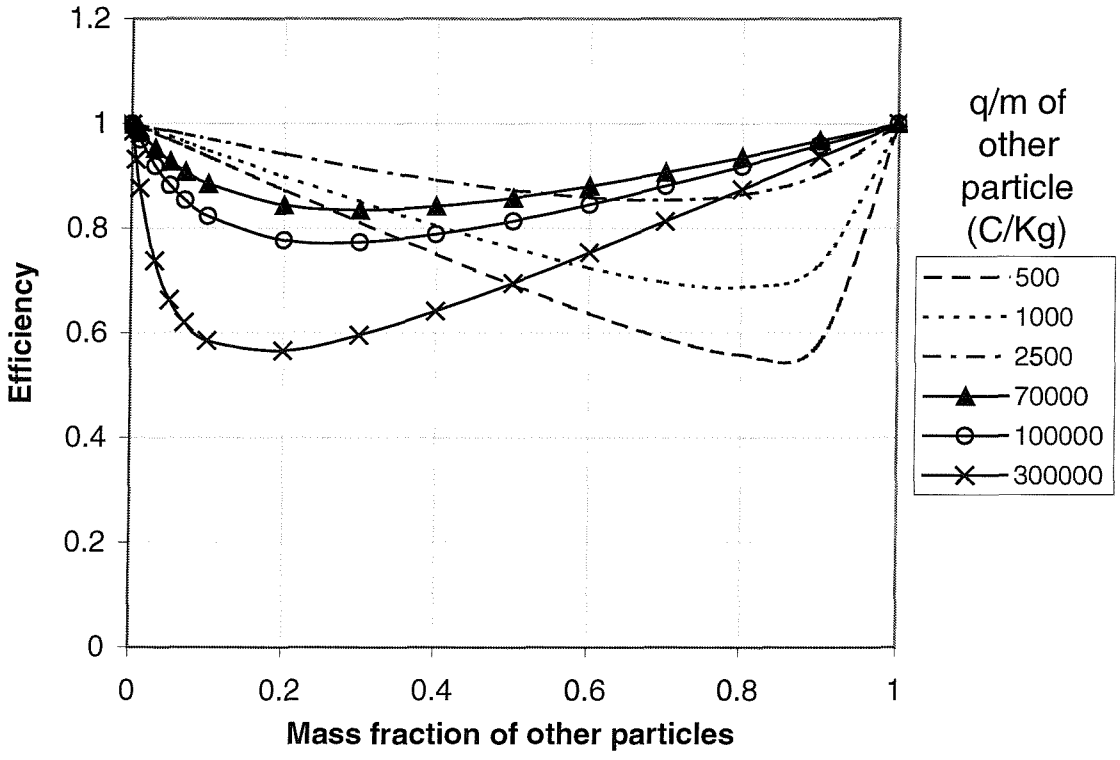

FIGURE 3.2: Effect of other particles on efficiency (main q/m=12,000C/Kg)

Recent work into characterising electrosprays in the regimes useful for colloid thrusters has verified the above theory [74]. This work used formamide solutions of differing conductivities as well as tributyl phosphate solutions. The author kept the voltage difference between the needle and electrode constant but varied the needle voltage and therefore the energy of the droplets. This 'stopping potential' or energy analysis allows determination of the spray composition as well as the voltage loss caused by the formation of the cone-jet. He clearly shows that as flowrate increases small satellite drops are created and with even higher flowrates the kink instability starts to occur. For high conductivity formamide solutions he provides evidence for his prediction that at low flow rates the jet is sufficiently thin to produce electrical fields strong enough to produce molecular ions. He also shows that with these liquids droplets can be produced that have very high specific charges. This and other work [75] shows that the voltage available to accelerate the droplets is decreased by the voltage loss along the cone jet, for formamide this is found to be ~200V, for the lower conductivity tributyl phosphate it is much higher approaching 90% of the available voltage for really low conductivities.

3.3 Benefits of miniaturisation

3.3.1 Starting voltage

The applied potential difference required to start the electrospraying process depends strongly on the diameter of the exit hole and the separation between the liquid tip and the electrode. The equation that determines this 'starting voltage' uses the field solution of a hyperboloid geometry and is [71]:

$$V_{start,hyperb} = \sqrt{\frac{r\gamma}{\varepsilon_o}} \ln\left(\frac{4s}{r}\right) \qquad (3.8)$$

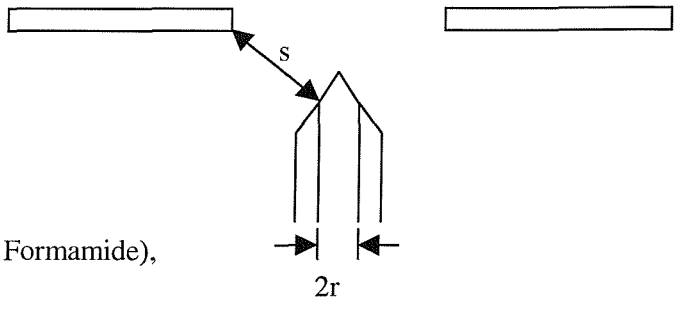

where s and r are illustrated in the sketch and;

γ = surface tension coefficient (0.058N/m^2 for Formamide),

ε_0= dielectric constant of vacuum.

This is also proved in the appendix, as it will be useful later on. So clearly the smaller the geometry of the emitters the lower the starting voltage can be. The thrusters using hollow needles (tested in the 1970's) needed to use high voltages of 5-10kV and this is mostly because the minimum diameters were of the order of 100's of microns.

The MEMS approach however, allows the fabrication of holes with diameters of 1-5μm. This enables starting voltages to be lowered to approximately 500V (depending on an electrode separation of 50μm). It also allows the fabrication of electrodes that are coplanar with the nozzle; such electrodes can be more precisely aligned and placed closer to the nozzle exit. In this case, the geometry means the field strength at the tip could also be estimated by the solution for coaxial conducting cylinders [76], which gives:

$$V_{start,coaxial} = 2\sqrt{\frac{r\gamma}{\varepsilon_o}}\ln\left(\frac{s}{r}\right) \tag{3.9}$$

This gives higher estimates of the starting voltage, typically about 40% higher than with the hyperboloid equation, providing a range of starting voltages to drive the design.

3.3.2 Evaporative Losses

Since the liquid is exposed to the space vacuum (at the liquid tip), a large proportion of the fuel assigned for propulsion can be lost, by evaporation, into space. Evaporative losses can be very expensive (in terms of lifetime vs. mass budget), so it is desirable that they be minimised. The first factor that comes into play here is the volatility of the fuel used, and this prohibits the use of water (which is otherwise attractive due to its high conductivity). Formamide is used instead as it is much less volatile but can be made conductive (by doping with ionic salts).

The second major factor affecting evaporation is the surface area of the exposed liquid; this is the area of the exit hole. The rate of liquid fuel lost by evaporation is directly proportional to the surface area of liquid exposed to the space vacuum, i.e. the square of the nozzle diameter. By using MEMS to reduce the diameter from, say, 100μm to 2μm, the surface area is reduced by a factor of $(100/2)^2 = 2500$, drastically cutting the fuel lost.

3.3.3 Batch Fabrication

In the conventional method of constructing a colloid thruster the emitters were machined using traditional methods in a serial process. To assemble the emitters into arrays capable of useable thrusts the needles were fixed manually. These methods resulted in inconsistency in factors such as tip thickness or diameter and varying needle height or distance from the extracting electrode. These can have significant effects on the electrical field in the tip region causing the performance of each emitter to be different from not only the other emitters but also the specifications. Using the batch fabrication methods of MEMS permits the emitters to be produced in a consistent manner with dimensions within more accurate tolerances. This allows large arrays of the emitters to be fabricated, in a quicker and more reliable way.

3.3.4 Reduced Size/Mass

Miniaturising the emitters and the separation between them and the extracting electrodes allows the overall size of the thruster assembly to be reduced. This is desirable for their use on micro and nanospacecraft. If all other factors are constant then reduced volume means a reduction in the mass used, giving decreased costs and increased mass budgets for payload and other subsystems. This can be seen by comparing the size of a thousand-emitter array fabricated by traditional means with a micro-fabricated version. The most miniaturised emitter array is that by Busek [36, 37] which has an emitter spacing of ~3mm; a 1000 emitter array would therefore be ~95mm by 95mm, giving a surface area of 0.009 m^2. A microfabricated design such as that in chapter 4 can have emitter spacing of ~700μm, giving a surface area of 0.0005 m^2, which is a 18-fold reduction.

3.4 Predicted propulsive performance

To determine the levels of thrust produced use standard electrostatic equations:

$$T = \dot{m}c, \ and \ c = gI_{sp} = \sqrt{\frac{2Vq}{m}}, \ giving \ T = \rho Q\sqrt{\frac{2Vq}{m}} \qquad (3.10)$$

and use equation 3.4 for q/m. For formamide;

$(q/m)_{max}$ = 1000*96500*2.2/(1133*15) = 12,491 C/kg.

So, $T_{(atQmin)} = 1133*3.8*10^{-15}*\text{sqrt}(2*12491*V) = 6.8*10^{-10}*V^{1/2}$.

This illustrates the extremely low thrusts available, for example with V=1kV, T=0.022μN; V=3kV gives T=0.037μN. The graph in figure 3.3 illustrates this. This is only for one emitter and shows the need for large arrays; to obtain 20μN of thrust requires 910 emitters, which is not a problem if a thruster is made using the batch methods of MEMS. Also the flowrate can be increased by a few orders of magnitude if greater thrusts are required, for example $Q=5*10^{-13}$ gives T~1μN per emitter. By using higher flowrates then, a higher thrust colloid array could also be made that would produce a mN of thrust with only 1000 emitters; however this would result in reduced Isp and efficiency.

The chart in figure 3.4 shows the most pressing reason for operation at low flowrates, the rapid drop-off in specific impulse as Q increases. The curve shown is for a voltage of 1000V and it can be seen that Isp falls below 500s when $Q > 6*10^{-15}$ m^3/s. The final chart, figure 3.5, shows the thrust to power ratio of a colloid emitter. The curve shown is again for 1.5μm exit radius using 1kV of applied voltage and NaI doped Formamide as propellant. It is proportional to $Q^{1/4}$ and can be seen to be quite high. When compared to the values for FEEP thrusters, which are typically of the order of $2*10^{-5}$ N/W, the colloid emitters look promising.

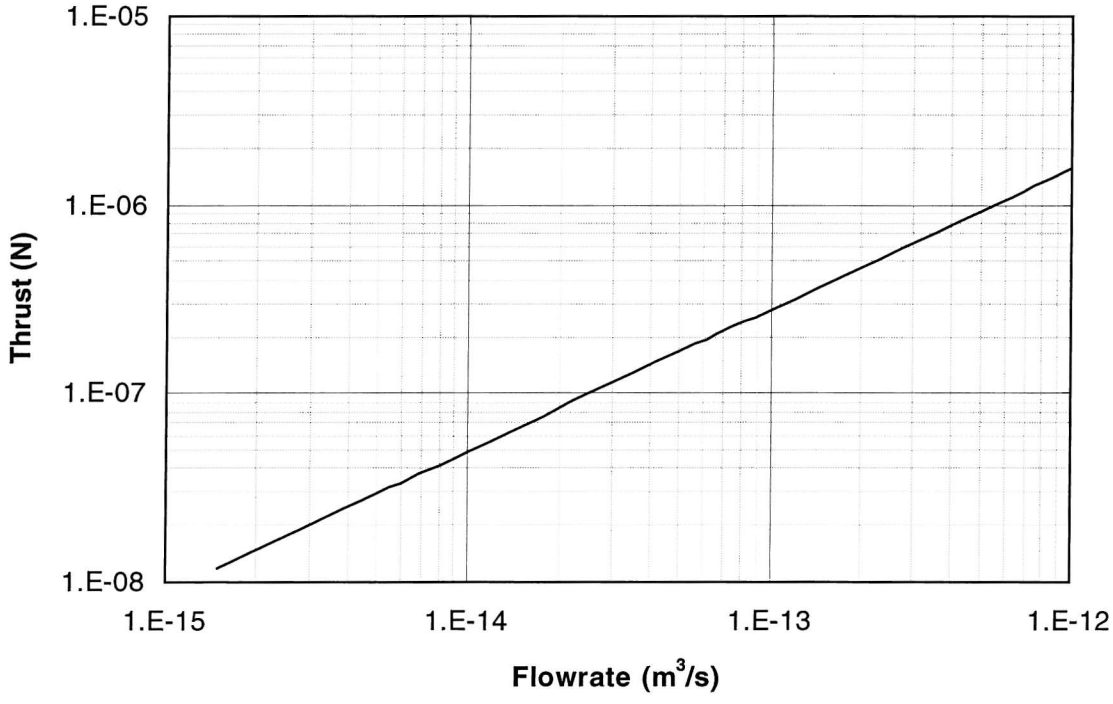

FIGURE 3.3: Thrust variation with flowrate

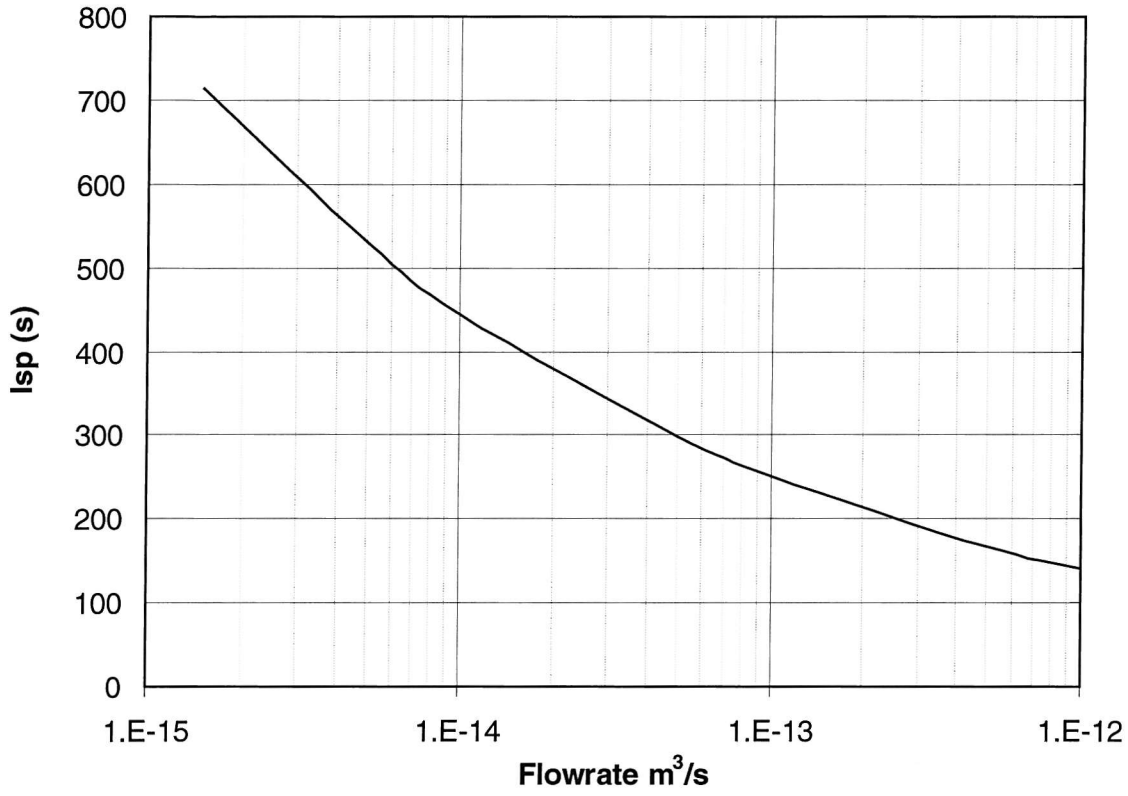

FIGURE 3.4: Specific impulse variation with flowrate

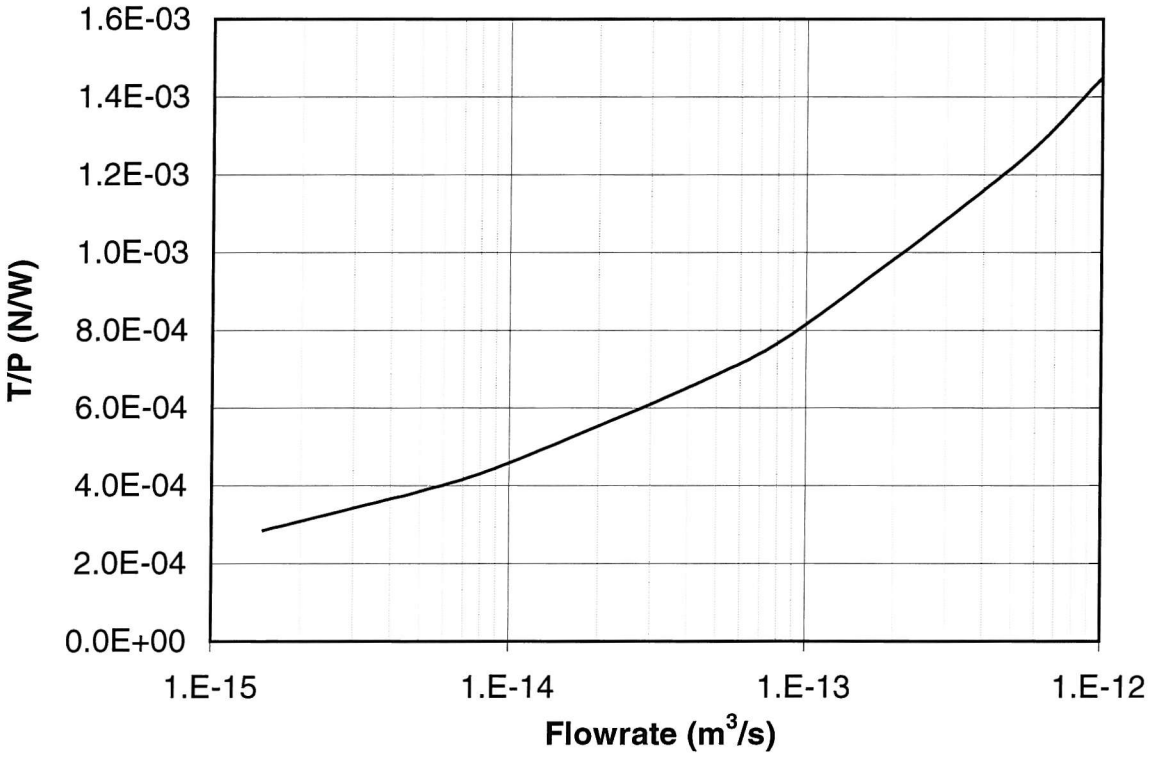

FIGURE 3.5: Thrust to power ratio variation with flowrate

3.5 Electrical breakdown

It is clear from the above discussion that one of the most crucial requirements for a colloid thruster is the electrostatics. To produce an electrospray an extremely strong electric field is required; to obtain high levels of performance high voltages must be utilised. It is crucial then to obtain an understanding of the numerous issues involved in electrical breakdown, since this must be guarded against if the electrostatics are to be sustained.

To maintain an electric field two conductors must be held at different potentials. This can only be achieved if an insulator is placed between them. The insulation prevents the flow of charge between the conductors, which would act to equalise the potential. As fields become stronger however, by either increased potential difference or decreased physical separation, a point is reached when the insulator can no longer perform this function. Depending on the insulator, breaching of this threshold can cause the material's temporary or, more often, permanent failure. Breakdown is an extremely involved subject, and only partial understanding has been realised, even after decades of intensive research. This section will attempt to summarise the most salient points from the literature and try to form a coherent design methodology.

The first useful division of the literature is into the following four areas, which, although broad and often overlapping, can be helpful. Gaseous breakdown is where an insulating gas of varying pressures separates the conductors; if the pressure of the ambient gas is less than about $1x10^{-4}$ mbar then gaseous breakdown is more correctly called vacuum breakdown.

3.5.1 Surface flashover

Whenever an insulating material bridges the space between two conductors at differing potentials, the voltage hold-off capability of the gap is compromised by the insulator surface. The presence of the surface often causes the gap to breakdown at voltages lower than that required for vacuum breakdown due to a process called surface flashover. The generally accepted sequence of events that takes place during flashover is the following [52]:

THE PRE-BREAKDOWN PHASE

1) Primary electrons are emitted from the cathode region, most likely this is at the metal-insulator-vacuum triple junction, and is due to field emission.

2) These electrons bombard the insulator surface and produce secondary electrons which themselves impact the insulator. This regenerative effect causes the insulator surface near the cathode to acquire an increasingly positive charge.

3) The charge increases the field at the triple junction and the field-emitted current therefore grows.

4) Following the electric field, the electrons 'hop' towards the anode and the area of positive charging quickly spreads across the insulator surface.

5) As a result a pre-breakdown current can be measured which is typically less than 1nA

This multiplication of the emitted-electron current is called the cascade effect, whereby each impacting electron produces a number of others, which themselves each produce more. The number of electrons produced depends on the material's secondary electron yield, δ, which is above one in a range of electron energies between ~ 20- 50eV and 1-5keV. The value of δ depends on the insulator material but also on the angle at which electrons hit the surface, with glancing angles give higher values. As the cascade proceeds and the positive charge builds up on the surface the electrons are dragged back to the surface sooner, their energies are then so low that δ tends towards one. This is the steady state condition and the avalanche is said to be saturated.

The insulator gains positive charge through the emission of electrons but the resulting hole in the valence band would have an extremely short lifetime in a perfect insulator. In practice though all insulators have some or all of the following intrinsic defects: point and linear defects, grain boundaries and surface defects. Through electronic excitation these create extrinsic defects, which act as charge localisation points. The relaxation of this charge then depends on factors such as diffusion and the energy level of the trap.

IONISATION PHASE

If the insulator-conductor system is placed in a true vacuum ($\sim< 10^{-7}$mbar) then for a surface flashover to occur an ionisable gas must be produced as part of the flashover process. This gas often consists of molecules desorbed from the insulator either thermally or by electron

stimulation. If the insulator has a poor thermal conductivity then the temperature can rise quickly. Alternatively, in the electron-stimulated model, the positively charged surface desorbs neutrals and positive ions, more neutrals are ionised by electron impact and the ions drift to the cathode. The increased triple junction field causes more electron emission and therefore more desorption etc.

DESTRUCTIVE PHASE

The breakdown can be thought of as being an energy conversion process whereby the electrostatic energy of the electric field is converted into ionisation energy as described above. The final stage is the conversion of the ionisation energy into thermal energy which causes thermalised conducting channels in the insulator surface. Energy is transferred to the surface molecules by interactions with the plasma electrons and the increased heat causes chemical decomposition followed by melting of the material. [53].

FLASHOVER VOLTAGE

The voltage required to cause flashover can be theoretically predicted if a few assumptions are made. By assuming the saturated avalanche model is correct and that ionisation takes place in a critical density of gas desorbed from the surface by electron stimulation then one can obtain the following equation [54]:

$$V_B = \sqrt{\frac{M_{cr} A_1 v_o e l}{2\varepsilon_o \gamma_d v_e \tan\vartheta}} \qquad (3.11)$$

This will be used later, but illustrates the dependence of flashover voltage, V_B on the square root of the insulator surface length, l. The secondary electron emission characteristics of the insulator material play a major part in this equation; the following relations were used in the derivation of the equation [54]:

$$v_e = 5.94 * 10^7 \sqrt{A_1} \qquad 3.12)$$

$$\tan\vartheta = \sqrt{\frac{2A_o}{(A_1 - A_o)}} \qquad (3.13)$$

Where, A_o and A_1 are the creation energy and impact energy respectively, which lead to the critical field angle, θ. Also, v_e is the drift velocity of the gas and v_o is the average velocity of an electron in the avalanche. For the ionisation phase to take place there is a critical amount

of gas, M_{cr}, required for the creation of a gaseous breakdown across the surface; M_{cr} can vary but should be taken in the range of 1 to 4 $x10^{18}$ molecules/cm^2. Since the gas is liberated by electron-stimulated desorption the number of gas atoms which desorb when the surface is struck by electrons is considered by:

$$\gamma_d = \sigma_d Q_{ed} \qquad (3.14)$$

where Q_{ed} can be taken as a small multiple of $10^{16}cm^2$. The amount of gas molecules on the surface therefore has a significant effect - γ_d can vary between 5 and 15 depending on the amount of adsorbed gas. The drift velocity of the gas molecules is assumed to be within the range of 1.5 to 2 $x10^5$ m/s using the gas temperature.

METHODS FOR FLASHOVER SUPPRESSION

The essence of most flashover suppression techniques is to interfere with either the ionisation phase or, more usually, the surface charging in the pre-breakdown phase.

Ensuring that the insulator surface is free from contamination, not just dust, skin grease etc. but also the residues of common cleaning agents, can inhibit the ionisation phase. To limit the amount of gas available for ionisation the specimen can be tested in a high vacuum environment and allowed to outgas for many hours. This allows gas adsorbed on the insulator surface to be desorbed; high vacuum conditions minimise the presence of atmospheric gases that would adsorb back onto the surface.

More effective are the various methods available to limit the charge build-up on the insulator surface. The point of origin of the primary electrons is the cathode triple junction (CTJ) and so great benefit can be derived by insulating the area to decrease the electric field there [52]. Altering the surface properties of the insulator can also help. If a thin conducting or semiconducting coating is applied to the insulator, the flashover voltage can be greatly increased. Coating a quartz specimen with a few nm of tungsten oxide doubled the amplitude of pulsed voltage required to cause flashover [55]. The reason behind this improvement is thought to be twofold. Firstly, the new surface generally has a lower value of δ than the insulator (some have $\delta < 1$) so the cascade is reduced and secondly the reduced surface resistance allows charge build-ups to redistribute more evenly.

Another approach is to use magnetic fields, through the E X B effect, to "lift" the electrons away from the insulator surface. Or, at lower magnetic fields (parallel to the surface), to at

least slow them down so their energies are below the minimum required for $\delta > 1$. It was found [55] that with the magnetic field parallel to the surface but perpendicular to the electrical field that the flashover voltage was doubled with a 0.2T magnetic field. The work found that this improvement was only for pulsed voltages and residual gas pressure (even as low as $1x10^{-6}$ mbar) had a detrimental effect on the improvement.

3.5.2 Gaseous breakdown

When the insulator separating the electrodes is a gas one can study the interactions of electrons with the gas molecules to try to determine when the gas will become conducting. For this to happen there must be a sharp increase in the number of mobile charge carriers i.e. electrons and ions. Under the influence of an electric field a single electron will be accelerated to energies high enough to ionise any gas atom it collides with - producing another electron and a positive ion. These charge carriers move towards their respective electrodes and in the process more collisions occur and more charge carries are produced. The first electron current, produced by cathode irradiation for example, is then multiplied many times. At some magnitude of current multiplication the discharge becomes self-sustaining and this is the point at which breakdown is said to occur. Townsend first described the phenomenon as an avalanche that depends on the gap length and pressure as well as the electrical field.

Clearly one of the most important aspects of this process is the requirement for charge carriers to undergo ionising collisions. If the gap is very small or if the gas pressure is low then this assumption can become invalid. The likelihood of an electron colliding with an atom is determined from the mean free path for an electron. If the gas has a Maxwellian velocity distribution then the equation is [56]:

$$\lambda_{ek} = \frac{1}{n_n Q_{nk} \sqrt{2}} \tag{3.15}$$

where the subscripts e, n and k indicate electron, neutral atom and the collisional process involved respectively. The pressure of the gap is considered through the density of neutral gas molecules, n_n, which can be calculated from the ideal gas law to be:

$$n_n = 9.66 \times 10^{19} \frac{P}{T}, \tag{3.16}$$

where n_n is in molec/cm^3 if P is in Torr and T in K°. Assuming room temperature then the equation for the electron mean free path for ionisation becomes:

$$\lambda_{ei} = \frac{1}{\sqrt{2} \times 3.2 \times 10^{16} P Q_{ei}}, \text{ in cm.} \tag{3.17}$$

Since air is mainly made up of nitrogen and oxygen, one can use these molecules for Q. NIST [57] gives the cross section for ionisation of nitrogen and oxygen by electrons with energies above a few hundred eV as ~10^{16}cm^2. Therefore, one can estimate the electron mean free path for room temperature air for high-energy electrons (such as those accelerated substantially by strong electric fields) to be

$$\lambda_{ei,air} \cong \frac{0.22}{P}, \tag{3.18}$$

which for atmospheric pressure air is about a few microns, but at high vacuum is very large. For example even at a fairly poor vacuum of 10^{-3} Torr, the mean free path is 2.2m, much larger than usual chamber dimensions. However, since ionisation can result from other interactions than the single-electron impact this may be a vast over-estimation. Using the largest cross section that may be suitable, i.e. for a molecule to simply collide with another molecule (if one or other is previously excited this could produce an ion) then:

$$\lambda_{nn,air} \cong \frac{5 \times 10^{-3}}{P} \tag{3.19}$$

even then the mean free path is 5cm for a mTorr vacuum.

An established rule of thumb when considering the likelihood of gaseous breakdown in a gap is Paschen's law [58]. A well-accepted equation for the law is the following:

$$V_B = \frac{APd}{B + \ln(Pd)} \tag{3.20}$$

where for air A=15/cm/Torr and B=1.18 and P is in Torr, d in cm. A plot of which is shown in figure 3.6.

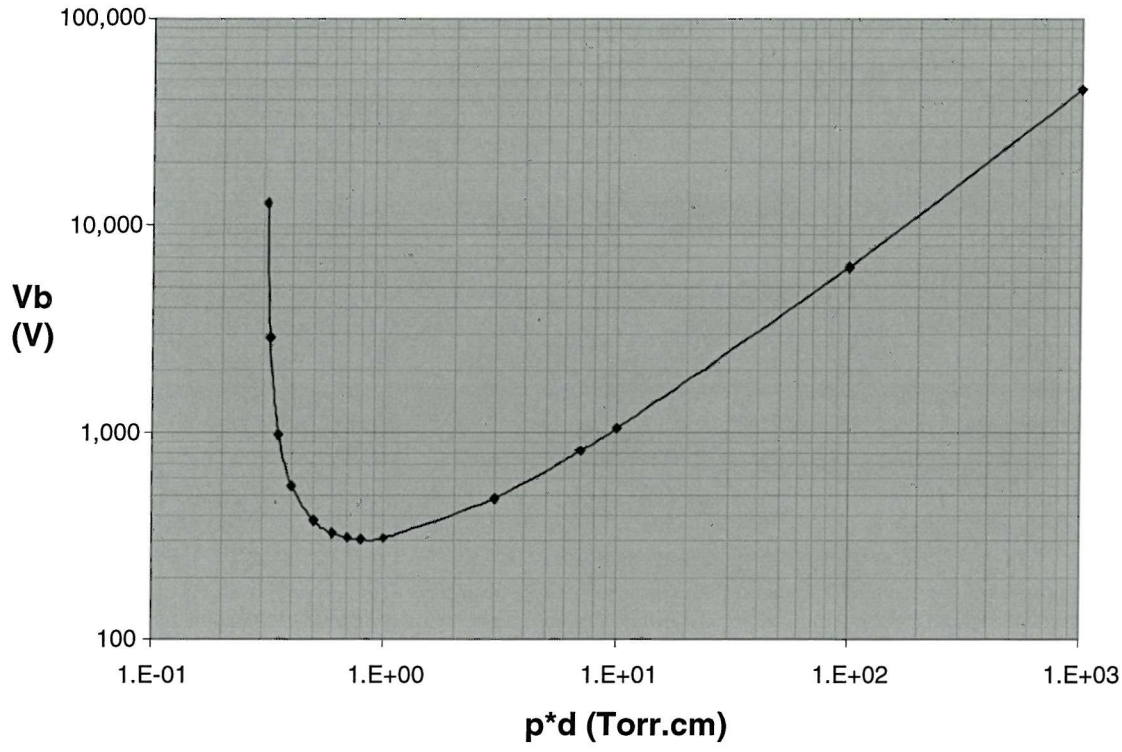

FIGURE 3.6: Paschen curve for room temperature air

The breakdown voltage reaches a minimum of ~ 320V at Pd values of ~ 0.57 Torr.cm. As Pd decreases, it enters the region known as the left of the Paschen minimum and V_B rapidly increases. This is because as the gas density falls the chance of a collision is less – to obtain sufficient ionisation the electrons need to have higher energies (stronger fields) to increase their ionisation probability. However the equation is only valid down to pd values of >0.31, and Paschen's own work only considered values down to 0.1 Torr.cm. [59]. One can see though that a gap of 100μm will give a value of Pd=0.3 for P=30 Torr – well above the pressure the experiments are likely to use.

3.5.3 Vacuum breakdown

Since the field required to cause breakdown of a vacuum between two perfect electrodes is 6.5MV/mm, a vacuum is one of the best insulators. The vacuum of space is far from ideal however, due to the effect of space plasma, radiation and particles from outgassing or thrusters. Consequently, the pressure can vary by several orders of magnitude above that attainable in a laboratory. Provided the spacecraft mission is not in low earth orbit the atmospheric particle density is not likely to be greater than 10^{-7} torr; equivalent to a good

laboratory vacuum. Since outgassing and thruster sourced particles are far harder to predict it would be prudent to use large reserve factors to the limits obtained in terrestrial vacuums.

Vacuum breakdown is a complex issue but has been the subject of considerable research effort. To simplify the problem researchers generally agree on defining the problem as clearly as possible. The required experimental conditions are ultra-high vacuum ($< 10^{-7}$mbar), with broad area edge-free electrodes that have been very finely polished and cleaned of contaminants. These stipulations are needed because they eliminate those elements that are known to alter the breakdown characteristics of a gap in a random way.

BREAKDOWN MECHANISM

Once a breakdown is initiated, it proceeds by way of an arc discharge between the anode and cathode. This is due to an avalanching ionisation process in the relatively high pressure of electrode metal vapour created by the initiation and sustained by the discharge. This is fairly well understood as it is a gaseous discharge. The difficulty in vacuum breakdown lies in the initiating processes.

The first set of processes is based on electron emission. Electrons are emitted from the surface of a metal when the electric field is above ~ $3x10^9$V/m but significant electron emission occurs between flat electrodes when the general field is much lower. This is because isolated sites on the cathode can greatly enhance the electrical field if they have a protruding geometry. Even a highly polished electrode has whisker-like filaments or protrusions when viewed on the microscopic scale. In general the higher the aspect ratio of such sites the greater the field enhancement that causes electron emission to become more likely and, when it happens, more copious.

However, for the breakdown discharge to occur the 'seed' of an ionisable medium is required [52]. Localised melting or vaporisation of the electrode material must produce this. In the case of the cathode, the material is vaporised when the emission site is excessively Joule heated by the emitted current. At the anode the beam of electrons can cause a hot spot that heats the material there to its melting point.

The second possible initiating mechanism is the impact of an accelerated microparticle. Any real world electrode will have a number of very small particles loosely stuck to the surface. Charge will be transferred from the electrode to the particle and when a strong electrical field

detaches it the charged particle will be accelerated across the gap at very high velocities. Upon impact with the counter electrode, the particle's kinetic energy is dissipated into the surface with the potential to vaporise enough microplasma to trigger a breakdown. It is generally agreed however that one or other of the above processes is predominant in a given situation. Breakdown tends to be triggered by electron emission when the gap is small ($d < 0.5$mm) and by microparticle impact for larger gaps ($d > 5$mm) [52].

EXPERIMENTAL OBSERVATIONS

When the voltage applied to the electrodes is slowly increased a very low current is seen (often < 1nA) but as it gets higher large current spikes or transients are seen. These instabilities are short lived but become bigger and more frequent as the voltage increases closer towards the breakdown voltage. They are representative of micro-discharges and are called pre-breakdown currents. It is thought they occur because as voltage increases, sub-critical initiation processes, as described above, occur and then pass. Another plausible explanation is the sudden desorption of pockets of gas from the electrode surface that are then quickly ionised and pass across the gap. By raising the voltage in a way that allows the current spikes to die down before continuing, the electrodes become 'conditioned'. When this has been done, the breakdown voltage is found to be far higher than for unconditioned electrodes. The time constant for the micro-discharges to settle down is approximately 30mins.

Certainly in the case of electron-emission dominated breakdown, one can expect that the breakdown will be influenced by the choice and condition of the electrode chosen to be the cathode. This has been shown to be the case when the electrode with a lower field-enhancing surface and hence a higher breakdown voltage is the cathode. The voltage polarity can have another effect in that if the high voltage electrode is the cathode then the whole of the grounded chamber effectively acts as a potentially gas-desorbing anode, which can aid breakdown.

A limited amount of work has been carried out into breakdown characteristics of very small gaps, i.e. in the micrometric regime. The authors who have investigated this area have found that gaps can withstand very high voltages before breakdown. An empirically based formula [60] was obtained, V_B (kV) = $37 * d^{0.37}$, with d in mm, for gaps between 25 and 200μm giving a breakdown voltage of ~12kV for a 50μm gap. Even smaller gaps in the range of 3-25μm

gave results which fit the formula Vb (kV) = 0.35* $d^{0.8}$, where d is in μm [61]. For instance 5μm held ~ 1.2kV. This was for unconditioned gaps.

3.5.4 Breakdown in solid films

During the sixties and seventies a large amount of research was carried out with the aim of understanding and characterising the breakdown of thin film insulators, in particular SiO_2. While that work has continued to the present day for the case of increasingly thin films interest in thick films (>1μm) has faded somewhat. However, a fair degree of consensus was reached as to the main mechanisms of breakdown through such a film.

MECHANISM

In solid films a breakdown event is defined as one in which the change from the insulating to a conducting state is via a permanent material change, as opposed to switching in which the conduction is only temporary. Typically, such a change takes the form of an I-V curve consisting of a switch from a high voltage-low current state to a low voltage-high current state. When a voltage is applied, an initiation process occurs in which an instability causes current runaway, the voltage then collapses and the low voltage state is established. When the voltage collapses the electrostatic energy stored in the capacitor is being discharged through the breakdown area – sometimes through an arc. Such events are reviewed in detail by Klein [62]. Looking across the range of insulating materials he identifies three main types of process that can cause breakdown: thermal, electronic and electronic modified thermal.

The intrinsic mechanism behind breakdown was the subject of some debate. However, the discovery that breakdown strength increased with temperature and with decreasing thickness, as well as the low power levels involved, argued against thermal runaway being responsible. The experimental data strongly supports the impact-ionisation-recombination breakdown mechanism [63]. In this model electrons enter the oxide conduction band through Fowler-Nordheim tunnelling from the cathode and the field accelerates them. Those that gain energy greater than the band gap impact the lattice atoms and ionisation creates electron-hole pairs. While the electrons are removed swiftly, the holes remain, thereby increasing the cathode field. This multiplies the electron injection in a regenerative feedback manner that is opposed by drift and recombination. This theory is consistent with the seen thickness dependence and

electrode material independence. It is expanded upon in work [64] that derives a theoretical model for the critical field above which recombination (drift is thought to be negligible) is insufficient to balance the hole generation and current runaway occurs. It finds that only a small amount of ionisation is needed for the runaway to occur. The model provides an equation to predict the critical field (E) given information about the material. The equation given is:

$$\frac{e^{H/E}}{(H/E)^2} = M \,, \qquad (3.21)$$

where H is a constant and M is a function of the oxide that is proportional to the thickness. As an example, the authors use experimental data to find a value for M that allows them to predict a critical field E of 7.92MV/cm. The work is particularly thought to be valid for SiO_2 insulators with thickness above 0.1μm. This is supported by the realisation that thicker oxides need higher voltages before breakdown – this means that the electrons can gain sufficient energy (> ~9eV) to cause impact ionisation. This is often not the case with thinner insulators.

BREAKDOWN VOLTAGE DETERMINATION

In previous work, Klein [65] identified the problems associated with breakdown experiments on thin films. When the typically used metal-oxide-silicon capacitor undergoes breakdown it does so at the weakest spot in the film. That is, all films will have flaws or defects and it is undoubtedly at these flaws that breakdown will occur at the lowest voltage. When these are allowed to short the capacitor, it is no longer testable. Furthermore, the breakdown voltage results obtained are not characteristic of the bulk material, only of the weak spot. To avoid this he advocates the use of thin electrodes, <100nm, which allow the occurrence of non-shorting, so called "self-healing", breakdowns. In these events, the instability causes melting of the top electrode around the conducting filament that breaks the conduction path. In this way, he obtains thousands of breakdowns for each test specimen. For a 200nm thick oxide he obtains a curve in which between 30V and 120V about 20 breakdowns occur which correspond to the elimination of weakspots. After this, voltages of about 140V are needed to cause the breakdowns – this is indicative of the true breakdown strength of the specimen. Workers who do not use this approach find a random variation of breakdown field over a very wide range, typically 2 to 10MV/cm. When the specimen is allowed to short with the first breakdown it is important to take a large number of results to give a good statistical picture.

This latter method is suited much more to determining the failure point of an actual device than characterising fundamental properties of the insulator.

EXPERIMENTAL OBSERVATIONS

The final breakdown strength of 7MV/cm from the self-healing studies compares well to that obtained by other workers such as Osburn and Osmond [66] who use the self-healing method to obtain a peak of breakdown strength of ~ 9MV/cm. The distribution of final breakdown voltages was not Gaussian but skewed at the low end with a sharp cut off at the higher voltages. They also found a polarity effect with lower breakdown voltages for the case when the aluminium electrode is negative. No effect was seen for variations in voltage ramp rate between 5V/μm.s and 1MV/μm.s. However, while investigating the effect of aluminium electrode thickness they, somewhat surprisingly, observed a few self-healing events with 1μm thick electrodes.

A further consideration is the amount of time a voltage is applied to the capacitor for. SiO_2 breakdown is not time-independent as some researchers thought – the time to breakdown decreases quasi-exponentially with applied field; the curve is very steep. As an example [62], for a field of 700V/μm the time to breakdown of a sample was found to be ~5 min, for a field of 750V/μm that time was only ~ 1s. While electrode material has no effect on the maximum breakdown strength, it has a strong effect on the time to breakdown (or wear-out time). Researchers found [67] that polysilicon electrodes gave wear-out times many orders of magnitude greater than aluminium.

A detailed description of the damage created by breakdowns was given [65]. The lower voltage breakdowns were seen to be single hole non-shorting events in which a hole was explosively evaporated through the oxide and metal and a pit burnt in the silicon by Joule heating and the action of an arc-like discharge. At higher voltages these breakdowns became propagating – creating other holes in the surrounding area until a large amount of electrode material had been destroyed. These were often permanently destructive. The measured temperature of the aluminium during the arcing phase was ~ 4000°K.

Chapter 4

Design

4.1 Initial design decisions

It was decided [49] that making the thin tapering needle shapes used in previous colloid thrusters would be very difficult using MEMS techniques, and was in fact unnecessary. The main elements of a colloid emitter are a nozzle that delivers the liquid to a suitably sized exit orifice and some means of providing an electrostatic field that is sufficiently strong at the free liquid surface. This field must be able to overcome the surface tension and deform the surface into the cone shape, from the tip of which the jet can develop. Traditionally high aspect ratio tapered needles are used to focus the, otherwise rather low, macroscopic field to the required strengths at the needle tip.

However, this can be achieved without the need for such sharp geometrical enhancements if the dimensions are sufficiently small. Situating an axisymmetric electrode in very close proximity to the liquid surface provides the required electric field at those points that are close enough to the axis of symmetry. It was decided to fabricate a fully integrated device in which the nozzle system and the extraction electrode are co-planar and integrated on the same substrate. An electrode that is on the same surface as a nozzle has the advantage that it can be defined in the same process flow as the nozzle so alignment is easy and accurate. There is no interfacing or assembly of separate components required, whereas a traditionally-fabricated design would require the electrode and the nozzle to be separate and interfacing. This solution is simpler, more accurate and makes full use of the dimensional control and batch production benefits of MEMS technology. It also enables the digital approach described later in this chapter since it permits patterning of the connections needed to allow individual addressing of extraction electrodes.

Another reason high-aspect ratio needles are often used is to provide sufficient hydraulic resistance so that flowrate can be controlled by pumps or by head pressure. This is addressed more fully later in the chapter, but it is not necessary, nor in some cases desirable, to try to

control flow in this way. With this requirement stripped away all that is required of a nozzle is that its exit orifice should be sufficiently small and offset a little from the plane of the surface. It is this fact that allows us to envisage a design in which very large numbers of nozzles can be produced across the surface of a wafer, giving large arrays of the kind needed to make a useable microthruster.

A basic model can be formed to illustrate the design of the emitter, which consists of a structural substrate, an insulating layer and an electrode (see figure 4.1). The substrate has the small-featured nozzles etched into it and an extraction electrode is integrated on the surface.

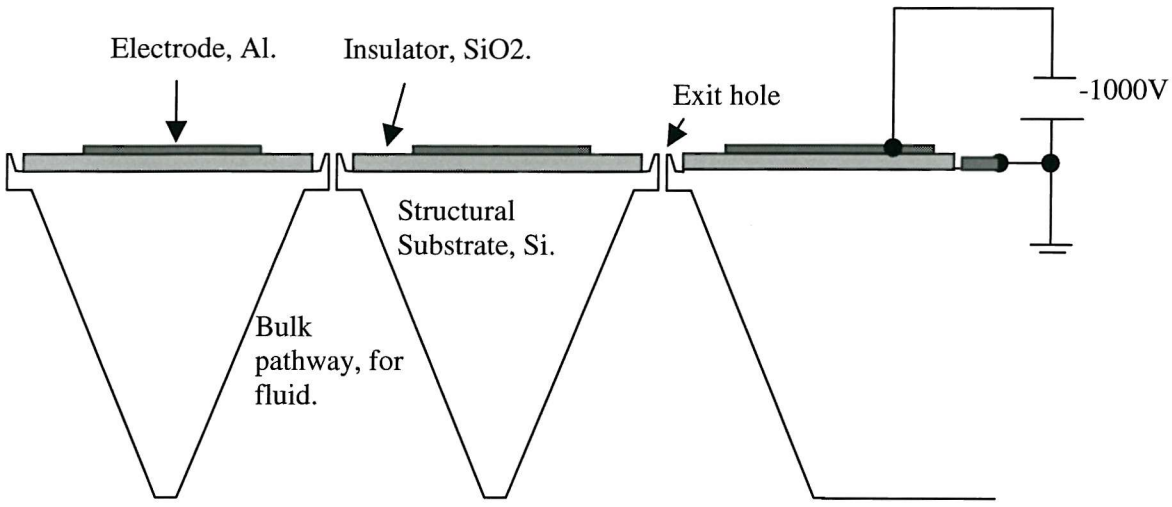

FIGURE 4.1: Cross section of microfabricated emitters

Material selection for each is described below, and they were made subject to the following constraints.

4.1.2 Breakdown

To be able to start the electrospray any thruster design needs to able to create the very strong electrical fields that overcome the liquids surface tension. As mentioned earlier the field strength depends on the geometry so the voltage required can be reduced by shrinking the dimensions. However, since a fairly high value for Isp is needed there is a limit to how low the applied voltage should be – a reasonable minimum is about 1kV.

This presents us with the challenge of ensuring that the design will avoid breakdown at these high voltages. As described in chapter 3, electrical breakdown can occur through a vacuum, across an insulator surface or through the bulk of an insulator material.

4.1.3 Material selection

The materials used in MEMS and the fabrication processes used to shape them have a strong co-dependency. This means that the material and fabrication process must be considered in conjunction and selected appropriately. The factors to be considered during the selection are as follows.

SUBSTRATE

In terms of mechanical and thermal properties it is clear that silicon is the best choice for a structural substrate. All of the features can be fabricated in it, as often silicon is the material of choice in the micro-fabrication field so there are a large number of processes available. The tapered tip can be created easily using either wet or dry etching. Reactive ion etching can produce the narrow holes, and the bulk fluid pathway can be created by using anisotropic wet etching, which although it produces square pyramidal pits is compatible with the design. Additionally, single-crystal silicon is an excellent mechanical material with an intrinsic stability. Because of its high yield strength and low density it has a specific strength higher than that of many engineering materials. It has high thermal conductivity, which is important since the electrospray process depends on the flowrate and therefore on temperature (via viscosity) so heat needs to be conducted away from the liquid. Although no thermal modelling has been performed, possible sources of heat are posited to include ohmic heating and electron/ion bombardment of the emitter surface.

INSULATOR

The materials to be used as insulators need to be able to withstand electrical breakdown. The ability of a material to do this is usually given in terms of the electrical field at which the material will breakdown. If the material is to be used as a deposited film (which typically have thickness less than 10μm) it requires a breakdown strength of 10^8V/m or more. A previous study [49] identified silicon nitride and silicon dioxide to be the best candidates. Other materials such as silicon carbide are not suitable because the fabrication methods available for use with SiC are very limited. The final decision to use SiO_2 is made by

considering the fact that this material has the highest dielectric strength and can be formed into thicker films than SiN.

Field effect emitters (shown in figure 4.2) are in many ways similar in design to the colloid emitter. Both have an electrode with a circular aperture centred on a silicon tip; both a vacuum gap and an insulating layer separate the electrode and substrate. The field emitter can easily use potential differences of upto 100V, which the gap and insulator have been shown to support.

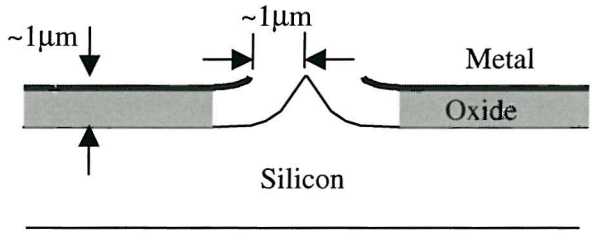

FIGURE 4.2: schematic of a field emitter

A 2 μm layer of SiO_2 will be used as the insulator, deposited either by chemical vapour deposition or grown using thermal oxidation. Thermally grown oxide is of a better quality with a higher breakdown field (1kV/μm against 500V/μm for CVD oxide).

ELECTRODE

The electrode material must be deposited via a process that does not damage the integrity of the surface. This means that a deposition process that causes little ionic or thermal damage would be the most suitable. Compared to sputtering, evaporation causes little damage, it is cheaper than CVD and results in a very pure film with good surface roughness and adhesion, so should be used. Aluminium is the material to be used as it is easily deposited using evaporation and, unlike gold or chromium, does not raise contamination problems that would prohibit its use in CMOS processes.

4.1.4 Geometry selection

SAFE GAP

There must be a sufficiently large gap between the electrode and both the liquid and the substrate (if it is conductive), the purpose of which is to provide vacuum insulation of

sufficient thickness to prevent arcing. Using the formula for micro-gaps in section 3.5.3 it is found that a vacuum gap of ~4μm is sufficient to hold 1kV, if there is an insulator bridging that gap though a design guideline [77] is to apply a safety factor of 10 to this. The resulting value of 40μm gives a minimum, however the gap length will be varied a great deal above this value as this is an important design parameter.

FLUID PATHS

The benefits of using small diameter nozzles have been shown in chapter 3, however this presents a number of restrictions on the fabrication process as a whole. It dictates that an anisotropic etch process must be used, which will give round holes with vertical walls and large aspect ratios. The best process for this is a reactive ion etch, the process available here has a maximum etch depth of 6-8μm. However, for the substrate to be strong enough to withstand handling and other shocks it must be thick and a reasonable minimum thickness for silicon is 380μm. The bulk of the fluid pathway through the remaining 374μm of silicon must then be removed in another way, but this need not produce round profiles or straight (vertical) walls. The KOH etch process is a partially anisotropic process which is highly selective to certain crystallographic planes of silicon. The result is a square pyramidal pit in the surface with walls that are 54.7° to the surface normal.

GROUNDING

To ensure that the silicon substrate is not at a floating voltage it is required that an electrode be placed directly on the silicon surface. This allows us to drive the substrate to ground potential and enforce the desired potential difference.

4.2 Detailed design

In designing the lithography masks advantage is taken of the fact that the micro-fabrication techniques to be used are batch methods. This permits the production of a large number of designs on a single wafer, each with one or more parameters varied. The main geometrical parameters are detailed below.

4.2.1 Dimensions

Several dimensions of the thruster geometry are fixed by the choice of wafers and the feasible limits of the fabrication methods available. They are illustrated in figure 4.3.

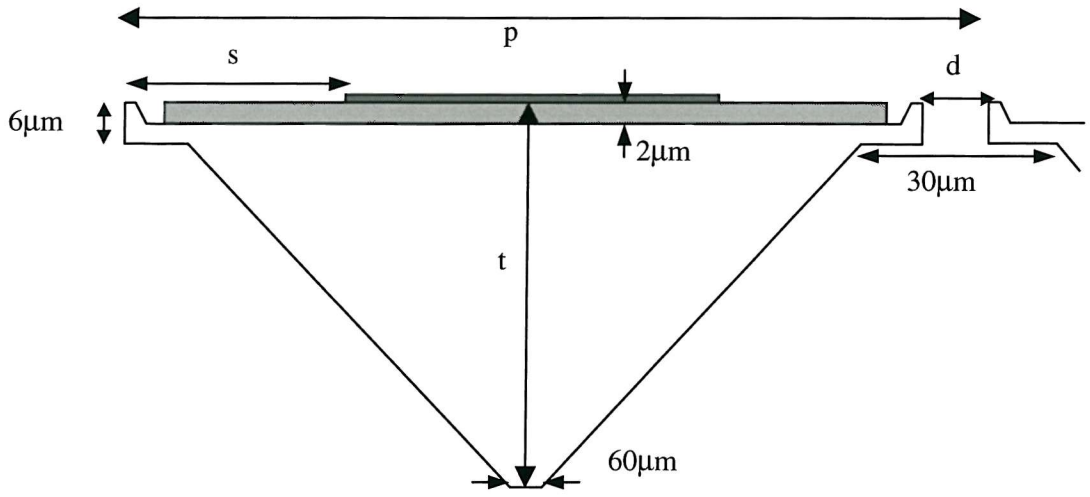

FIGURE 4.3: possible dimensions.

The first of these is the thickness, t, of the substrate. Although wafers with a thickness of 500μm are commonly used I used the thinner, 380μm wafers. This is because the pitch, p, in figure 4.3, which effectively determines the packing and therefore the thrust density of the thruster, depends directly on the wafer thickness (via the 54.7° angle of the KOH etch to be used).

The value of 30μm for the width of the emitter membrane is due to the predicted inaccuracy of back-to-front alignment. The features on the back of the wafer (i.e. the KOH etched bulk pathway) must be aligned to those on the front. This is done by shining an infrared light through the wafer to align special features etched in the front to features on the mask being applied to the rear of the wafer. A tolerance of ± 10μm is applied to this process, and taking into account the diameter of the hole this suggests that ~30μm is a sensible membrane width.

The 2μm thickness of the silicon dioxide insulator layer is the maximum that can be produced by thermal oxidation. This thickness should be capable of withstanding a potential difference of 1kV.

The 60μm at the bottom of the wafer is the separation needed between features for them to be defined during a deep KOH etch. They need to be this wide as the KOH etch is not infinitely selective to the crystallographic planes of the silicon. KOH attacks the 110 and 100 planes at

a rate that is ~ 20 to 100 times slower than the rate at which it attacks the 111 plane. This means that to avoid complete undercutting of the mask used to define features it needs to be of a minimum width. When a hole is etched that is 100μm the 110 plane will also be etched by ~5μm, and the mask will be undercut by 5*cos(54.7) = 2.9μm (as illustrated in the sketch) . Our holes will be ~380+15μm deep so the mask needs to be ~25μm wide. I have applied a safety factor of 2.0.

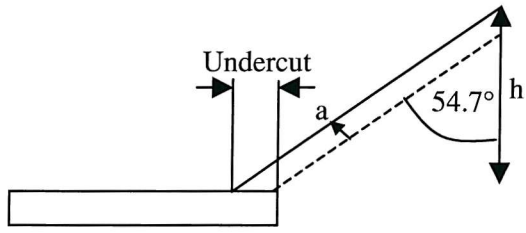

The main parameters to be varied across the wafer are the diameter, d, of the exit hole and the separation, s. The exit diameter will be varied through, 3, 5 and 10μm. This is to allow the effect on the starting voltage to be assessed and to protect against failure through clogging of the nozzles either by particulate contamination or by propellant polymerisation. The separation will be varied through the values of 10, 50, 100, 200 and 300μm. Such a large range is used because of the large degree of uncertainty over the distance required to prevent electrical breakdown. Clearly the smaller values are the most desirable as they produce higher electrical fields and higher thrust densities, but in case of arcing failure it is prudent to produce emitters with larger and therefore safer separations.

Although field emitters and the proposed colloid emitters are very similar, a possible difference between them is the extra oxide that protrudes from underneath the metal electrode. This may or may not be necessary. If the voltage to be applied is much greater than 300V then it may be desirable to have this extra oxide to prevent surface breakdown. As can be seen from figure 4.4, the oxide can be used to increase the distance a surface flashover must cross from just the layers thickness (say 1μm) to a much larger value (10μm or more). Field emitters operate without this extra protection but higher voltages may necessitate it.

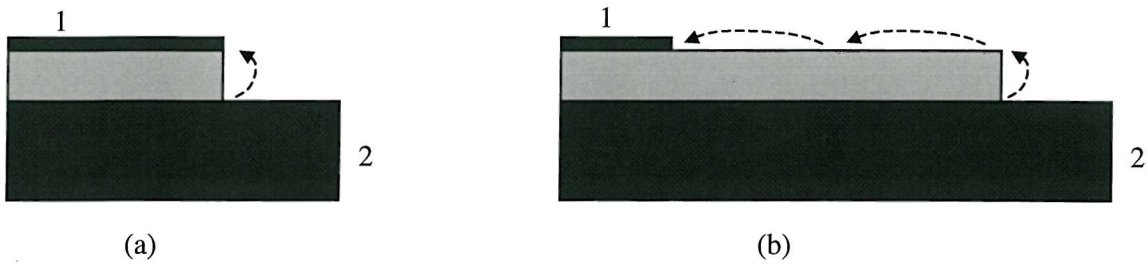

(a) (b)

FIGURE 4.4: Distances for surface flashover between electrodes 1 and 2.
(a) layer thickness, (b) any suitable distance

4.2.2 Thrust control

As the founding intent of this design was to try to fulfil the requirements of NASA's ST5 mission these should be reiterated here. The ST5 spacecraft (a scientific mission to measure gravity waves by cancelling all but gravitational disturbances on a proof mass) required fine thrust control over a micro-Newton range. The exact requirements were for 0.1μN control over the range of 1-20μN. The two most obvious ways of providing this control can be seen by examining the thrust equation:

$$T = \rho Q \sqrt{\frac{2Vq}{m}}$$

i.e. by controlling flowrate or applied voltage. Each of these has disadvantages however.

The most efficient flowrates for operation are those near the minimum, above this value the spray produced has a wide size distribution, because of the varied specific charge and droplet velocities that result, the thruster becomes very inefficient. In addition, controlling flowrates at the extremely low levels required is extremely difficult. No conventional pumps and valves can be used, and at the time of writing no project in the field of microfluidics has solved the problem. Apart from these difficulties, thrust control by altering Q is not an option for our design as nozzles of such small aspect ratios as these do not offer enough hydraulic resistance for the flow to be hydraulically dominated. Since, in the flowrates of interest, the hydraulic resistance of a nozzle is less than the electrostatic pull needed to overcome surface tension, the flow is electrically rather than hydraulically dominated.

Using changes in voltage (with Q held constant) to control the thrust is, at this moment, impractical. The fact that thrust varies with the square root of voltage, means that the 20:1 for ST5 requires a 5:1 range of voltage; upto 5kV would be required. This is too high to insulate against and is at the limit of the spacecraft power bus. Furthermore, when the voltage is increased over ranges such as these the Taylor-Cone mode of operation changes to an increasingly multiple-site emission mode [35]. When these transitions occur the thrust jumps up in difficult to quantify steps. Until the physics of this thrust increase are better understood it is felt that thrust control via voltage is impractical.

DIGITAL THRUST CONTROL

The method of thrust control pursued here then is a digital approach. Since the thrust from an individual emitter is so small and consistent from one emitter to the next, the thrust can be varied across the range by turning on the number of emitters required to produce the desired thrust. This can be done by fabricating chips with a large number of individually addressable electrodes that can be used to turn on small groups of emitters or "clusters". By identifying the number of emitters required to produce the unit of thrust control desired (T_{min}), one can group that number in a cluster and turn on enough clusters to obtain the demanded thrust (see figure 4.5). To meet the ST5 requirements a thruster has been designed that has 20*0.1μN sub-clusters, and 20*1μN clusters. So that, for example, 13.5μN of thrust could be obtained by turning on 5 subclusters and 13 clusters. It should be noted that each emitter operates at the same constant flowrate and applied voltage and therefore performance is unaltered. An example of such a digital approach is shown in figure 4.7 where for reasons related to lack of space the thruster has been designed with 9*0.1μN sub-clusters, 5*0.6μN clusters and 9*0.9μN clusters adding up to a maximum thrust of 12μN.

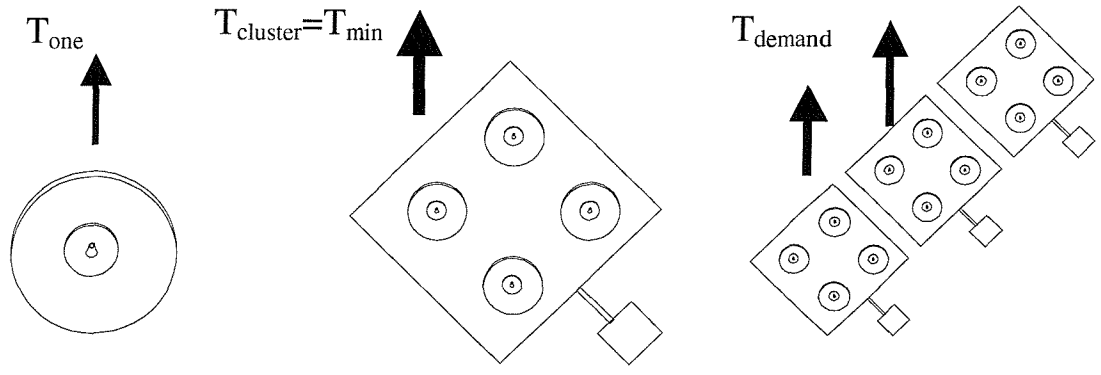

FIGURE 4.5: The digital thrust control concept

It can be safely assumed that the technique of extracting sprays only from particular emitters by applying voltage to the electrodes of those emitters required to be turned on is a sound one. Provided that is, that several conditions are met (refer to the sketch in figure 4.6).

Firstly, to ensure proper wetting of the emitter capillaries it is helpful to apply a pressure head, P_H to the rear of the liquid. Secondly, the pressure head should not be large enough to start liquid flow on its own, so must be less than the pressure exerted on the liquid meniscus by surface tension, i.e. $P_H < P_S$.

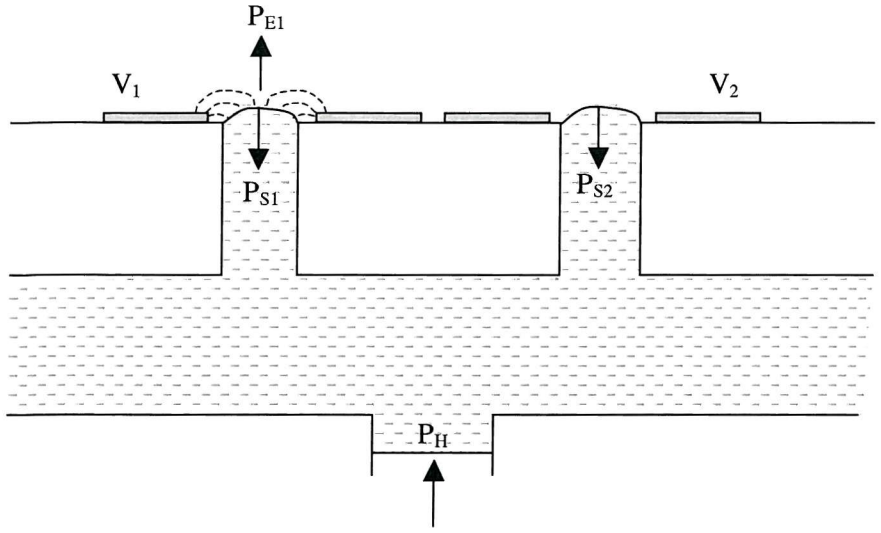

FIGURE 4.6: The pressures affecting two emitters. V_1 is turned on, V_2 is not.

Finally, the pressure exerted by the electric field, P_E, must be able to overcome the surface tension, i.e. $P_E + P_H > P_S$.

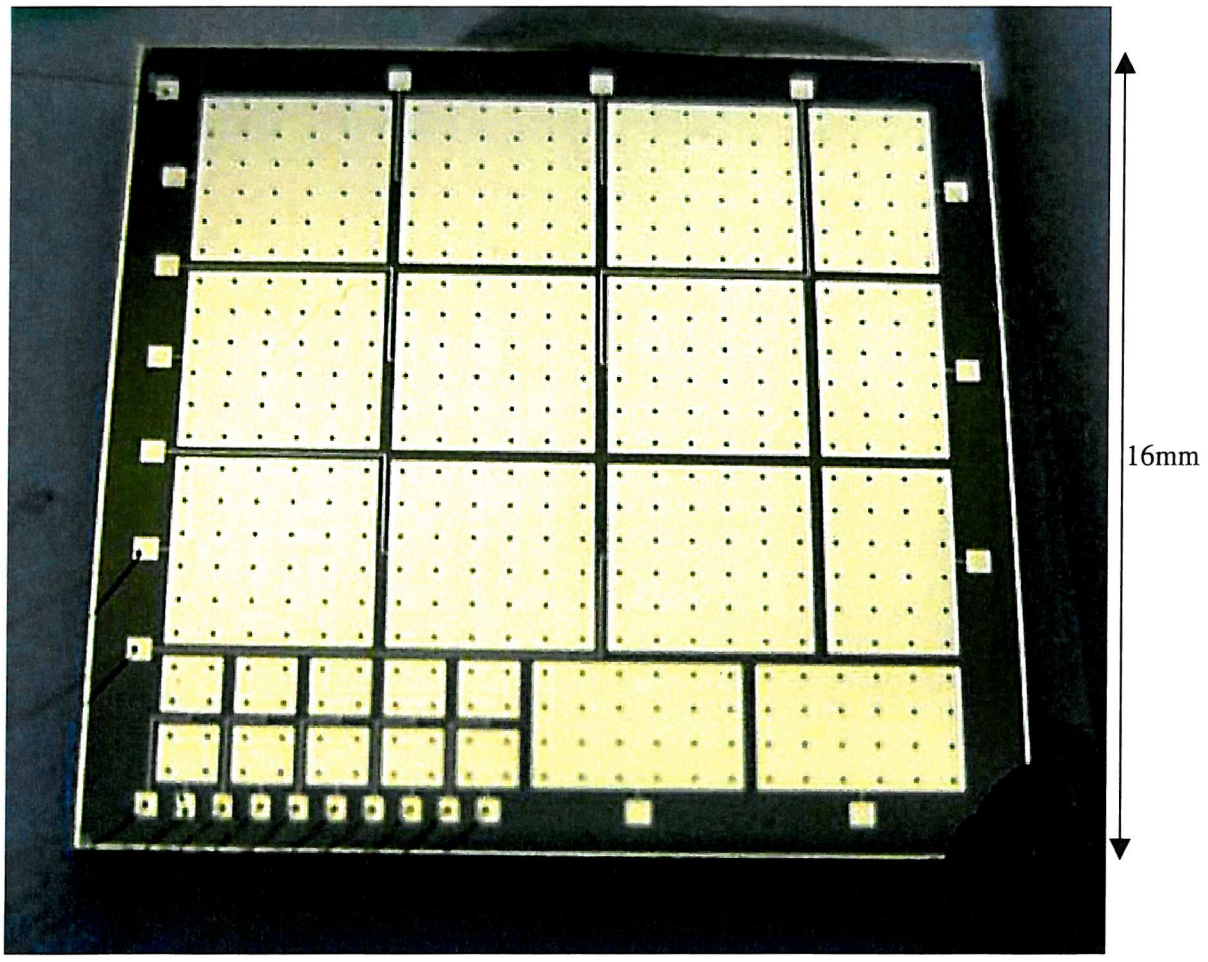

FIGURE 4.7: Photo of a "digital control" chip

The equation (derived in appendix A) describing the flow obtained during electrically dominated flow through a capillary is:

$$Q = \frac{\pi}{8\mu L}\left[\Pr^4 - 2\gamma\cos(\theta)r^3 + \frac{2\varepsilon_0 V^2}{\left(\ln\left(\frac{4s}{r}\right)\right)^2}r^2\right]$$

which exhibits a simple behaviour. For emitters using formamide and an applied voltage of 600V the curves in figure 4.8 are obtained for different emitter hole radii and pressure. As can be seen, in the hole radii range of interest, the flow rate increases as one would expect with both pressure and hole size. However, above a certain hole-size the field is insufficient to start the flow without the addition of a driving pressure, or higher voltage. The curves in 4.9 show how (when pressure head is zero) altering the voltage can alter the flowrate.

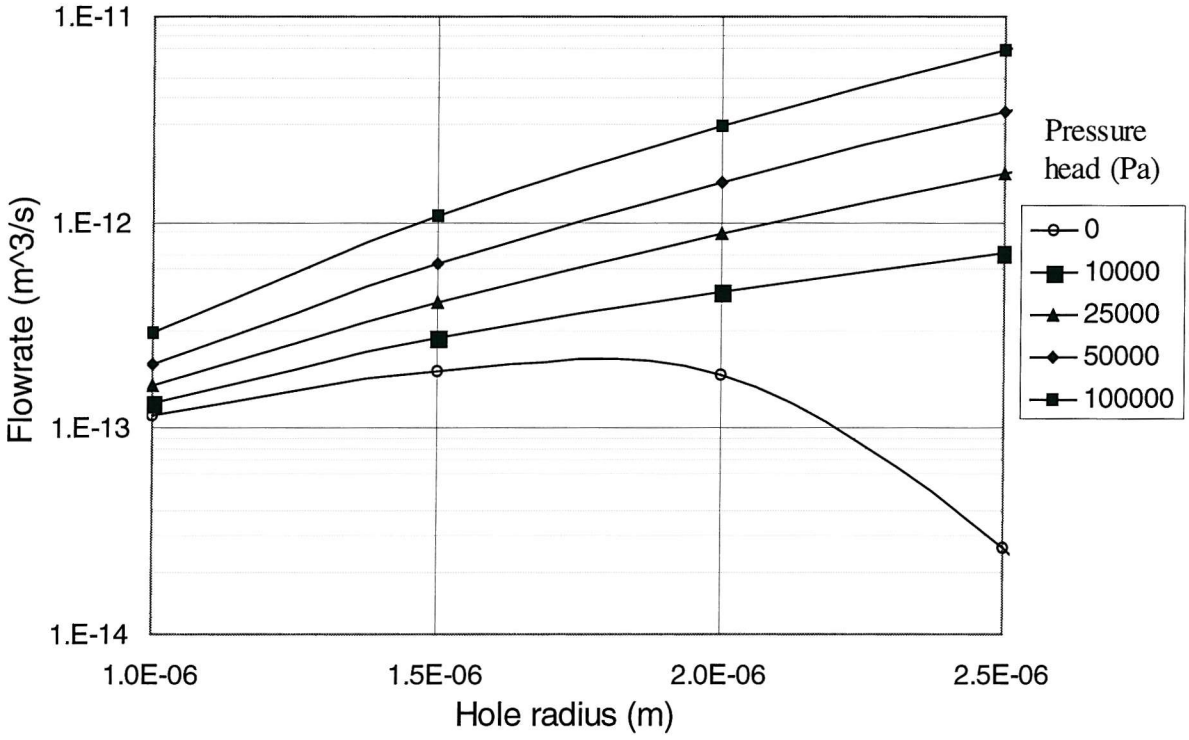

FIGURE 4.8: Flow variation with hole size and pressure (V fixed)

These predictions show a behaviour one would expect. The flow only starts when a certain voltage is reached, dependent on the size of the hole. Then the flowrate increases with the voltage applied, this is reasonable because the electrical stresses on the liquid become stronger and stronger. The flowrates are in the regime of interest i.e. near to the minimum flowrate. There is a very strong dependence of flowrate on the applied voltage however. Taking the curve for a 1.5μm hole as an example, one can see that flow starts to be above Q_{min} (~5×10^{-15} m^3/s) when V = 503V. The flowrate reaches 10^{-13} m^3/s when V = 509V. It is

thought that the range $Q_{min} < Q < 20^* Q_{min}$ is the range in which electrosprays remain in the cone-jet mode. Therefore there is only a limited range of applied voltage in which the spray can be expected to remain stable. There would be no problem applying voltages so precisely, however it does represent a limit to the performance attainable from an emitter; the Isp for the formamide solution operating near Q_{min} with 500V applied is limited to ~330s. This narrow range of allowable voltage is a result of the short length of the emitter nozzles – deeper nozzles with a length of, say 100μm, would allow voltages between 507 and 590V to be used.

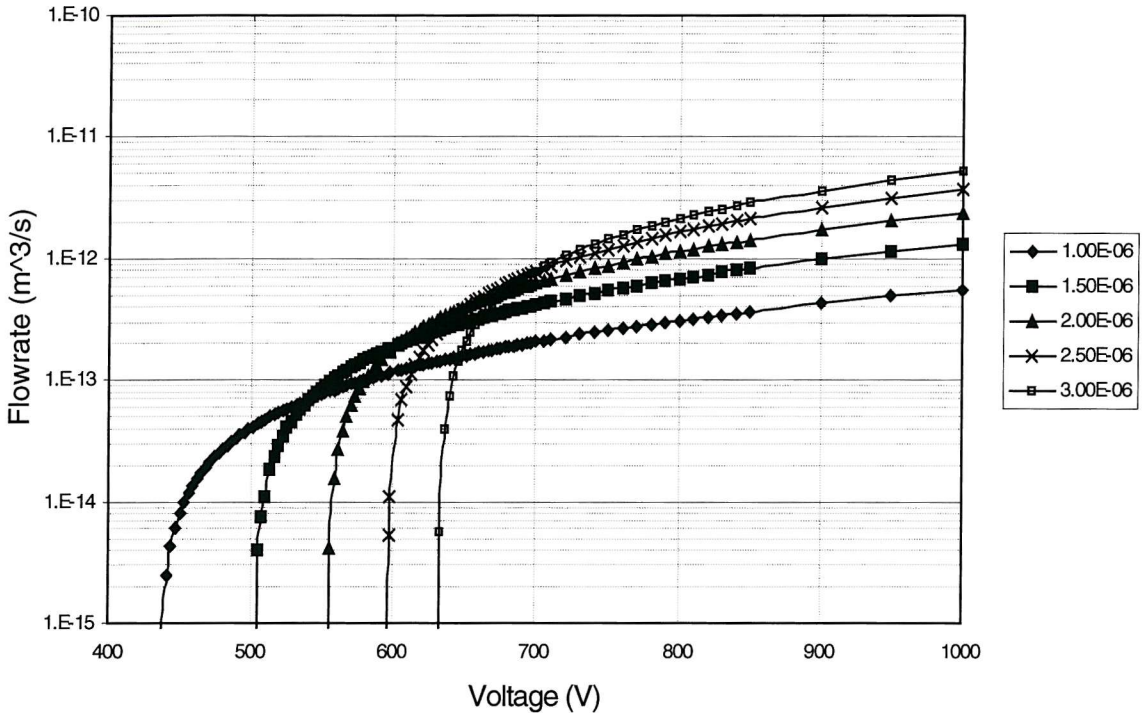

FIGURE 4.9: Flow variation with applied voltage and hole size (P fixed at zero)

A criticism levelled at this work early on was that it was thought unlikely that the theory of electrospraying would be applicable if flowrate is not hydraulically driven. There was doubt whether the $I \alpha Q^{1/2}$ relation would hold, or even whether the flowrates would be sufficient to be above Q_{min} and therefore produce an electrospray. Electrospraying research suggests however that when the voltage controls flow then flowrates above but also close to Qmin result and good cone-jet sprays are produced. Combined with an early version of the equation above it was felt that the criticism was probably not valid, recent work seems to verify that point of view. An experiment done by Busek when characterising their colloid thruster put the electrically dominated flow regime to the test [36]. By altering the onset voltage used to start the spray between 1594V and 1797V and keeping a constant 2kV acceleration voltage they measured the thrust and Isp versus flowrate when the flow driving pressure was zero. The resulting data points lay exactly on the same lines as when the flow was pressure driven

indicating that the I-Q relation is still valid. An attempt was made to model the results of that experiment using the equation given above and using figures from the paper as follows: r = 15μm, L = 1cm, K = 0.5S/m and all other liquid properties for formamide + NaI. The value for electrode separation was not given in the paper but by measuring the photograph in their figure 4 and using equation 3.8 to predict the minimum onset voltage a reasonable estimate for s is ~ 600μm. Figure 4.10 shows the flowrate versus applied onset voltage as taken from that paper and predicted using the above equation. It shows that the equation given is able to predict the flowrate quite well – certainly within one order of magnitude.

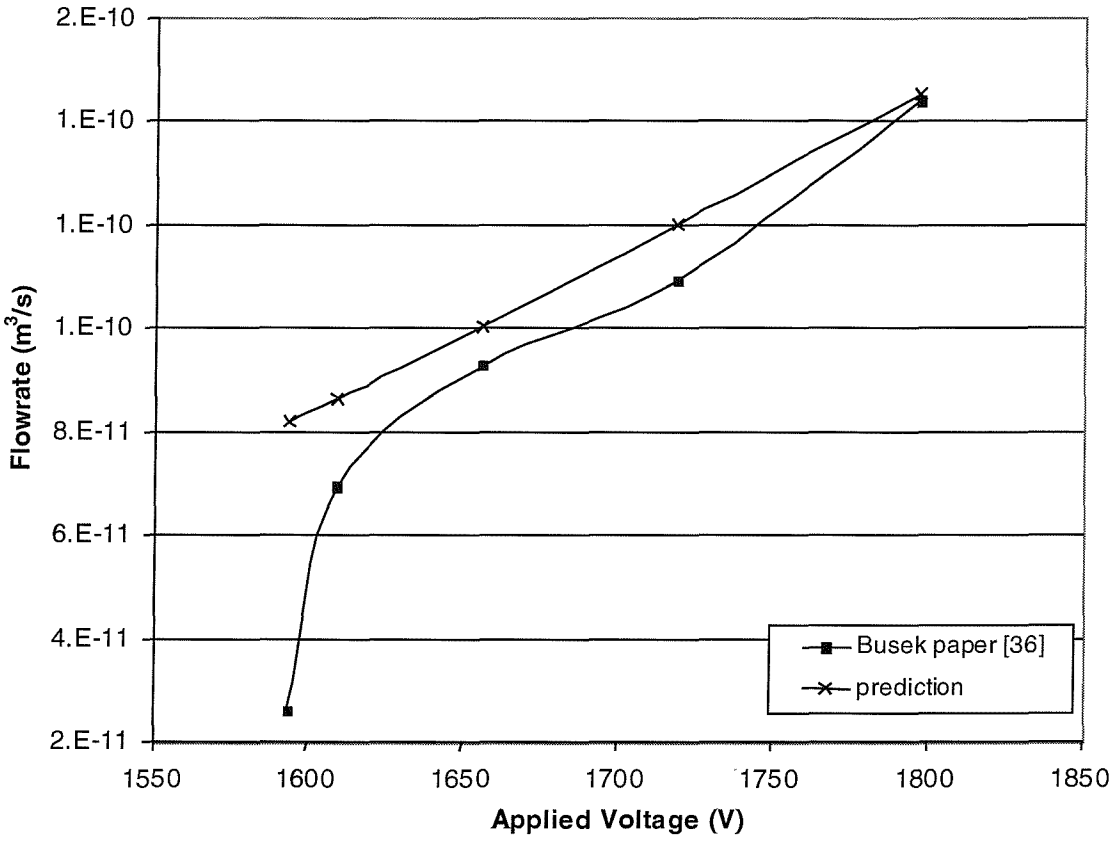

FIGURE 4.10: Predicted and measured flowrate dependence on voltage.

It seems then that there is little need to worry that the thruster will not create a proper ES using the electrically dominated flow regime.

This voltage driven method has already been very successful in the analytical chemistry world. Named nanoElectrospray (nES) it has been used extensively as a source of electrosprayed ions for mass spectrometry. It is performed with no flow-driving pressure head using fused-silica capillaries which have been pulled to a very small diameter – often a few microns or so, they are then coated in gold to make them conductive. Its exponents find

that it gives very good performance [78] and a stable spray is obtained easily (they lack the ruggedness and consistency of dimensions to be used for microthrusters however). I have experience with such set-ups. The use of a nES source with formamide in our lab has allowed us to obtain electrosprays starting at an onset voltage of ~1500V and providing ~100nA of spray current when the applied voltage was 1600V.

4.3 Further design comments

4.3.1 Mask layout of projected chips

The masks for the fabrication run had dies on them with the dimensions in table 4.1, where t is the nozzle rim thickness nozzle rim and o is the separation of the oxide and the nozzle rim.

Label	s (μm)	t (μm)	o (μm)	d (μm)	Size (mm)	No. on wafer	Thrust (μN)
S1	50	1	5	5	16	3	14
S2	100	1	5	5	16	3	14
S3	50	1	5	3	16	3	14
S4	100	1	5	3	16	3	14
T1	280	1	5	5	10	4	10
T2	200	1	5	5	10	3	10
T3	200	1	5	10	10	4	10
F1	300	3	20	5	5	4	1
F2	10	1	5	5	5	4	1
F3	300	1	5	5	5	3	1
F4	300	3	5	5	5	6	1
B1	200	1	5	5	32	1	140

TABLE 4.1: Details of the chips on each wafer

The predicted thrusts are only approximate and take into account the operational flowrate being close to Qmin and 1kV of applied voltage, using NaI doped formamide with a conductivity close to 2 Si/m. The nominal specific impulse will be approximately 500s, which is towards the low end of the desired performance range but if the insulation holds more than 1kV then this can be increased.

The main bulk of the chips on the wafers are of the digital control design i.e. those that are 16mm on a side. However, other, more experimental, chip designs are included. All of these others have only one extraction electrode per chip and are 5mm, 10mm and 32mm in width. On these chips I have varied the electrode separation to wider extremes than on the digital chips.

4.3.2 Ion emission

Because of the small dimensions used, it is quite likely that the electrical field at the liquid tip will be high enough to start field emission of ions. It is well accepted that a field of 1V/nm is required for the onset of ion emission. This is easily achieved in the region of the jet. A simple coaxial cylinder analysis yields the equation (appendix A):

$$E = \sqrt{\frac{2\gamma \cot r}{\varepsilon_o r}}$$

using values for formamide, then at the apex of the Taylor-Cone with the jet radius gives E=1.2V/nm. This mixed regime, where ions and droplets are emitted gives considerably increased Isp (as the ions have very high specific charge so are accelerated to much higher velocities). With an ion to drop mass fraction of 50% the average exhaust velocity is > 2000m/s yielding an Isp of 2000s. If it were attained this would be exceptional performance.

4.3.3 Bipolar operation

As already demonstrated by Perel et al [31] there is a realistic possibility of operating a thruster in a bipolar fashion, i.e. using half of the emitters to produce beams of ions and the other half to produce anions, by biasing the electrodes correctly. In this way, the design would avoid the need for charge neutralisation by an electron source, such as the heated filament or hollow cathodes currently used. The oppositely charged streams of particles will balance each other out and stop any spacecraft charging.

Gamero-Castano performed electrospray experiments using formamide doped with NaI and operated his emitters in both polarities [79]. He found that, over a range of flowrates, a very similar magnitude of current was observed when the electrodes had a positive potential difference (producing I anions) and when a negative bias was applied (producing Na ions).

It is not possible to produce bipolar sprays from emitters that share the same fuel source. However, two identical thrusters could be used in conjunction, one producing negative ions and the other positive ions.

4.3.4 Electrostatic model

An electrostatic model has been produced using Matlab's PDETOOL to solve the electrical field around the emitter. The boundary of an example emitter is drawn using the tool's GUI. Figure 4.10 illustrates the area chosen as the computational domain. It is a vertical slice of the axisymmetric vacuum surrounding the emitter. PDETOOL allows the user to specify the type of PDE to be solved; in this case I used the electrostatic mode, which solves Laplace's equation (space charge need not be considered since it is modelling pre-spraying conditions). The user must define the boundary conditions on each of the domain edges and these are shown in the figure. So that the PDE's take account of the axisymmetrical nature of the problem Laplace's equation must be solved in cylindrical co-ordinate systems. Since the domain is in vacuum, the only effect this has is the dielectric constant of the domain must be multiplied by r.

The domain is then discretised into a triangular grid. To ensure sufficient resolution in the region of the liquid tip the model has a sub-domain (a 10nm wide rectangle) in that area which forces the grid to the required resolution while keeping the rest of the domain at a much coarser level (thus saving considerable computation time).

Edge label	BC type	Value
A	Neumann	0
B	Dirichlet	0
C	Dirichlet	0
D	Dirichlet	0
E	Dirichlet	0
F	Dirichlet	0
G	Dirichlet	0
H	Neumann	0
I	Dirichlet	-1000
J	Dirichlet	-1000
K	Dirichlet	-1000
L	Neumann	0
M	Dirichlet	0

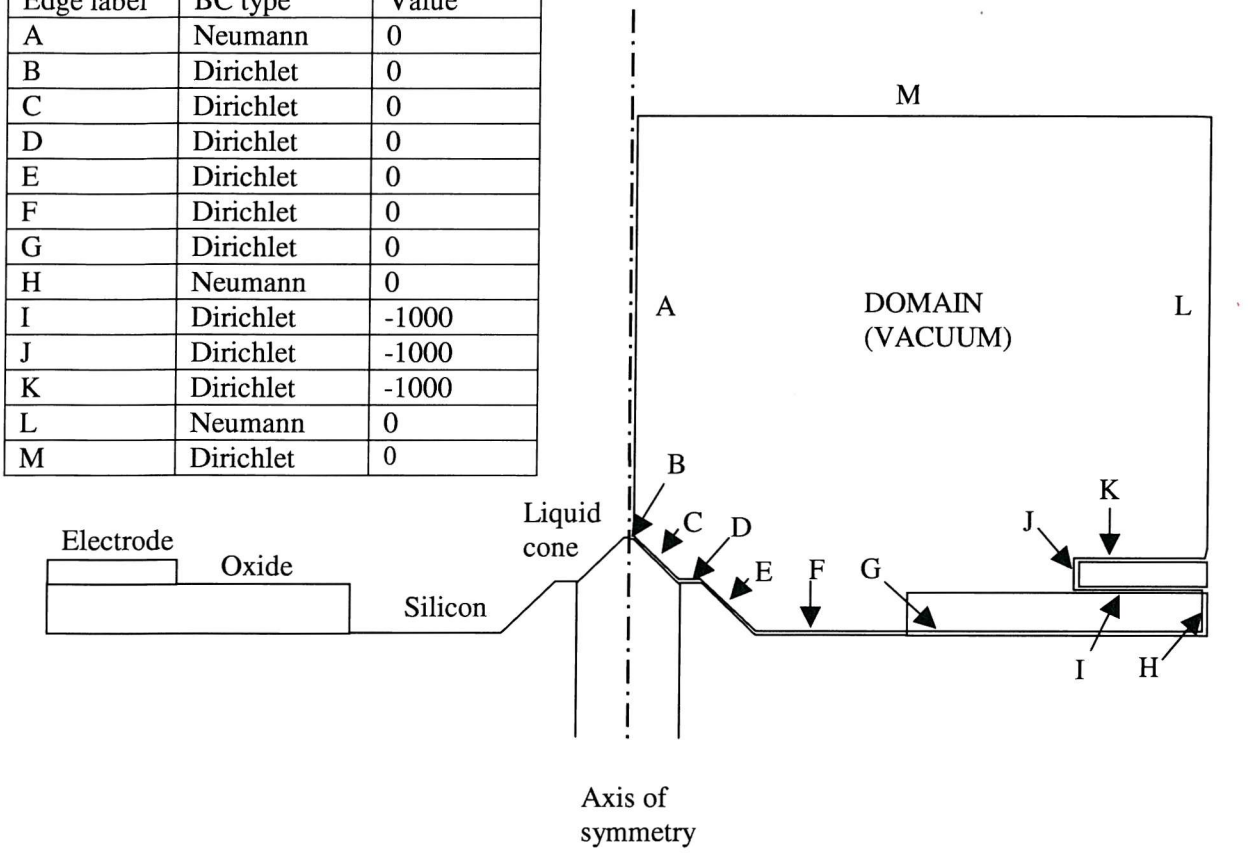

FIGURE 4.10: The domain to be analysed

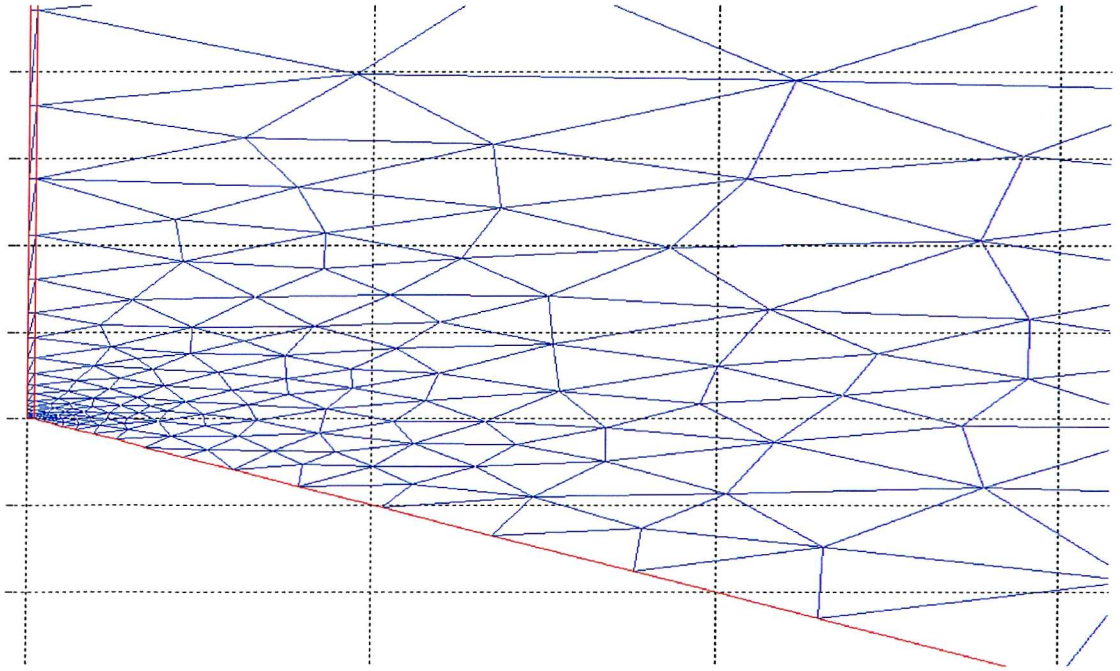

FIGURE 4.11: Increased mesh resolution at the liquid tip.

The PDE's are then solved by PDETOOL and the potential and E field is calculated at each node on the grid. These solutions (along with all the inputs used to create it) are exported to the MATLAB command line and saved as a workspace. The solution is shown in figure 4.12, it shows that E > 1V/nm at the tip of the cone. The field rapidly falls to values of around $2x10^8$V/m away from the tip of the liquid cone.

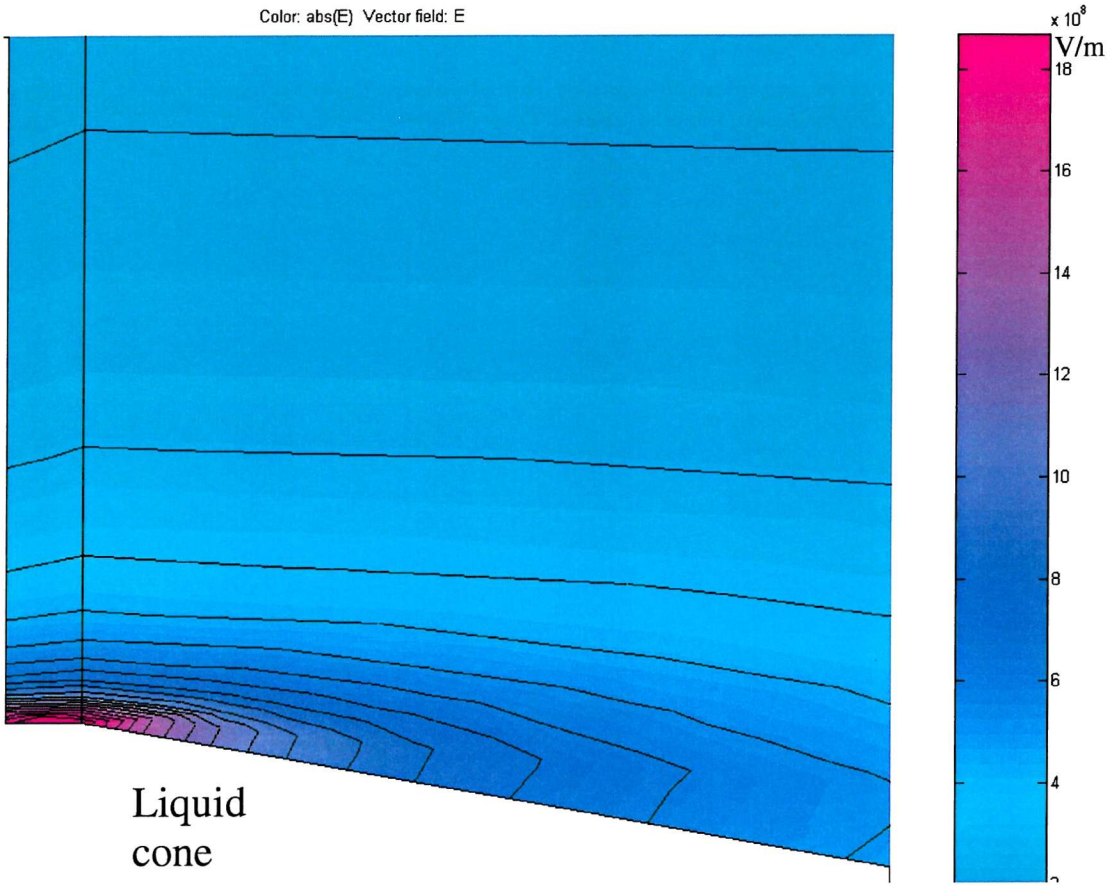

FIGURE 4.12: A close-up of the field solution at the cone tip.

Chapter 5

Fabrication

5.1 Fabrication process

The fabrication process is displayed below in cross section form (the drawings are not to scale for ease of presentation). It uses only standard microfabrication techniques. Steps that may need special explanation are the following:

The use of SOI wafers - Silicon On Insulator wafers have a thin layer of silicon dioxide buried between two much thicker layers of silicon. This is useful in this case as it acts as a precise etch stop for deep etches that will make the bulk pathway.

The use of Boron - By doping the front of the silicon with Boron and allowing it to diffuse inwards to a depth of 6 μm an etch stop can be formed. This is an alternative method of producing the etch stop and the rest of the process is identical to that illustrated below. The process had not been used for quite a while so it had to be tried first. After performing the diffusion, the wafer was cleaved and examined under an SEM; it was clear that the etch depth was 6-7μm deep.

Incomplete KOH etch - The KOH etch will not etch the whole of the bulk fluid pathway, but will instead stop approximately 10μm from the buried oxide layer. Then the ion etch at the end of the process sequence will remove the remaining silicon, stopping on the buried oxide. The reason for this is to maintain structural ruggedness. If the KOH etched the entire pathway, the emitter would have a very thin (2-4μm) and quite wide (~30μm across) square membrane, centred on the tip. During many of the subsequent processes the wafer

has to be held in a vacuum chuck, and it is very likely that the pressure differential across this disk would rupture it. The disk diameter is wide to allow for the large misalignments that will occur between the hole etched from the front and the KOH etch from the back.

Different oxides - The SOI wafers have oxide deposited by the PECVD process whereas the oxide on the boron wafers is grown using thermal oxidation. The reason PECVD was used on the SOI wafers is the incomplete KOH etch described above and the ensuing need for a late-stage ion etch. This approach has the advantage of keeping the membranes intact, as shall be seen later in note F of section 5.2.1, but does mean that the SOI chips have insulation that is only half as resistant to electrical breakdown.

SOI FABRICATION METHOD

Starting with an SOI wafer, 380μm thick, with the oxide layer ~ 8μm from the surface and 0.5μm thick. Both sides polished. Produce the alignment marks on both sides.

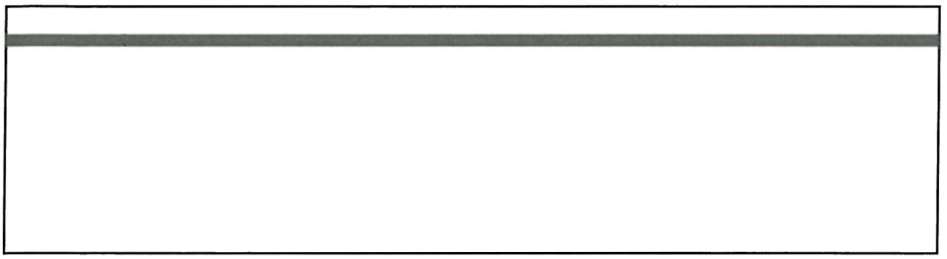

Deposit thin layer of oxide, and a thicker nitride layer (Ox/Ni) on both sides

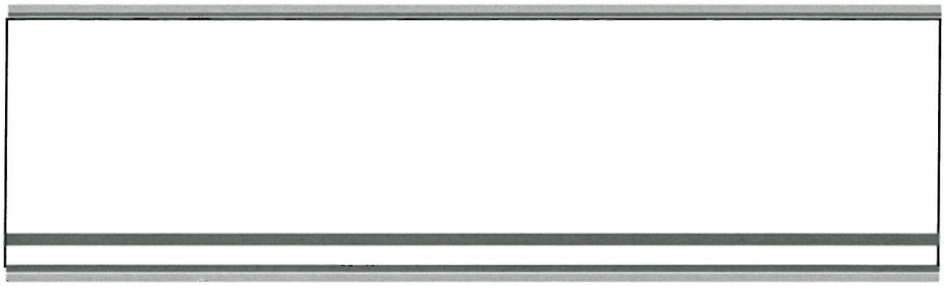

Apply resist and pattern it.

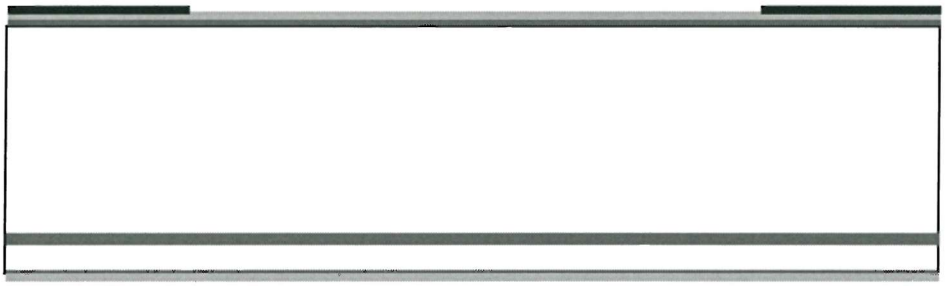

Use resist as mask for patterning the Ox/Ni

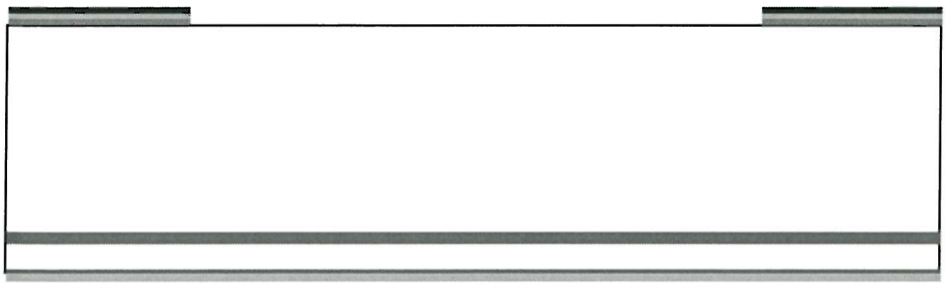

Use patterned Ox/Ni as a mask for the KOH etch, stopping ~10μm from the buried oxide.

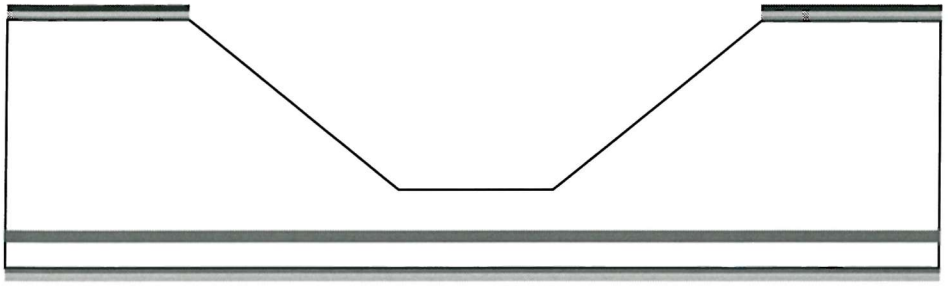

Strip Ox/Ni. Turn wafer over, apply and pattern resist on front.

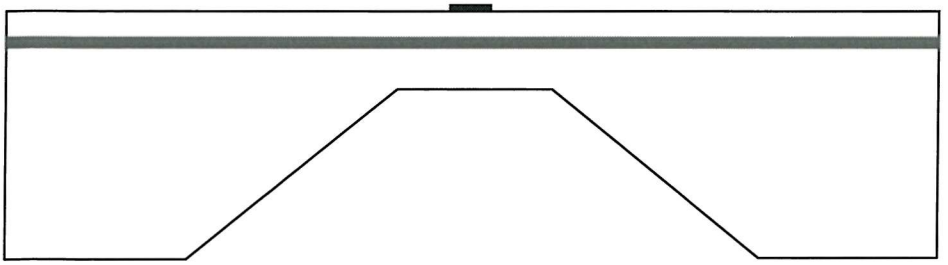

Use resist as mask for SF_6 etching of silicon tips (2μm deep).

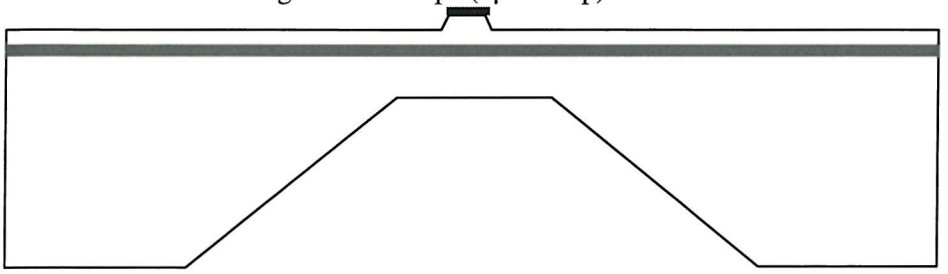

PECVD 2μm layer of oxide

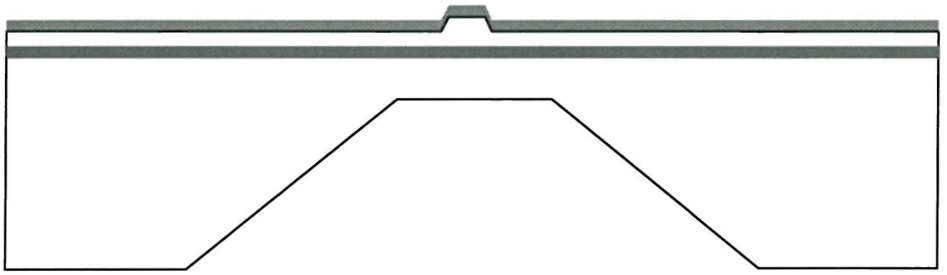

Apply thick resist and pattern

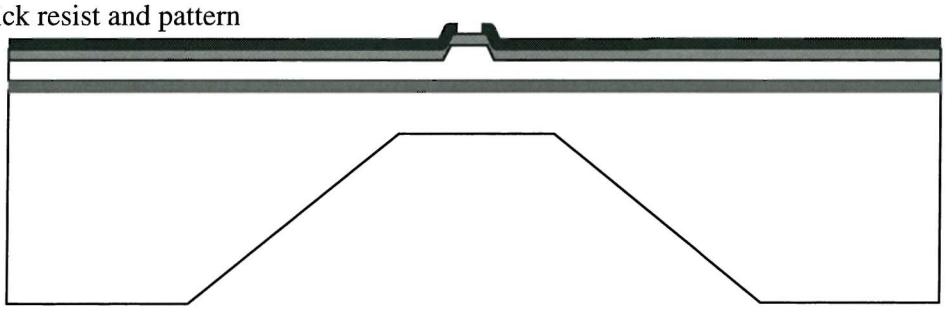

Use resist as mask for the ion etching of the Oxide.

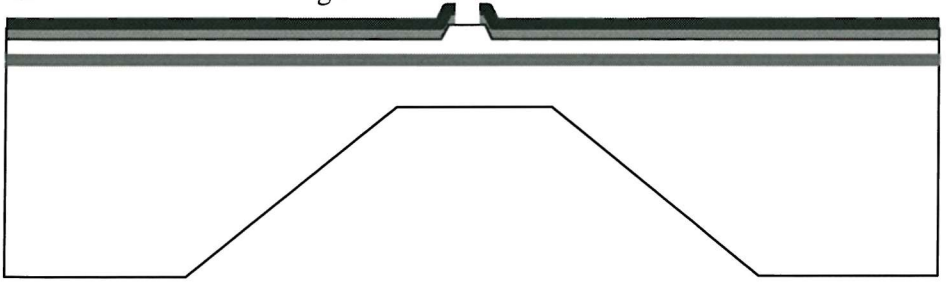

Use Oxide as a mask for the ion etching of the Silicon.

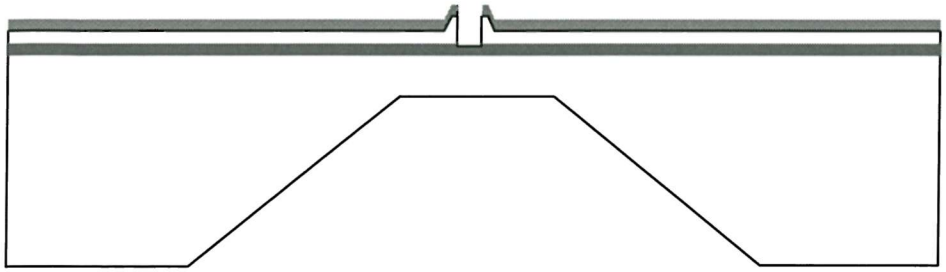

Apply and pattern resist (OE2 layer), use to ion etch oxide

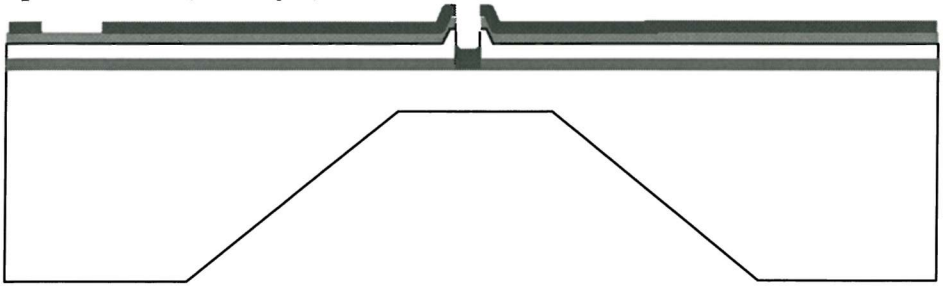

Apply resist, blanket expose and develop.

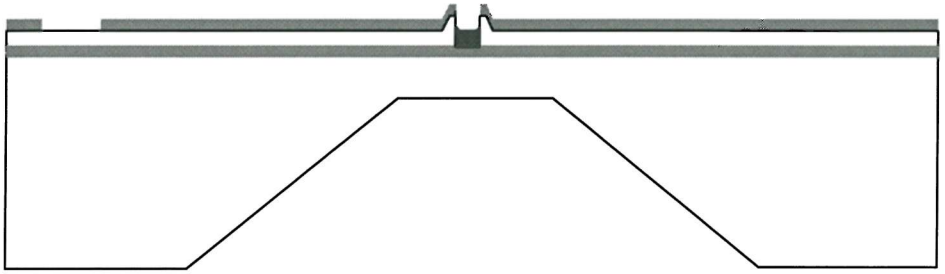

Sputter on Aluminium, apply resist.

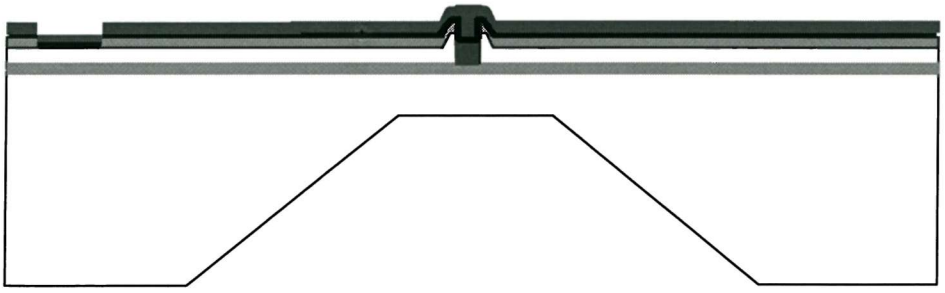

Pattern resist (EL layer) and use as mask to etch metal.

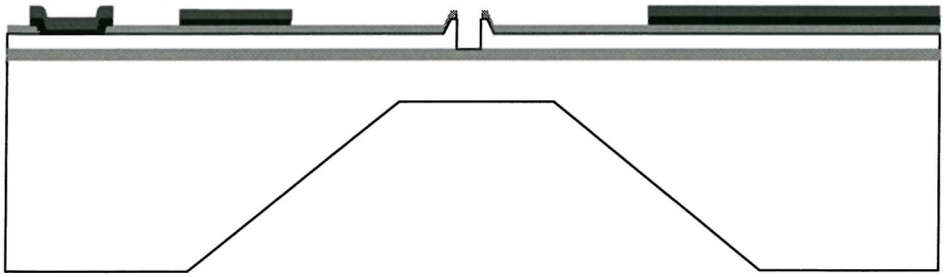

Turn wafer over and ion etch the remaining silicon stopping on the buried oxide layer

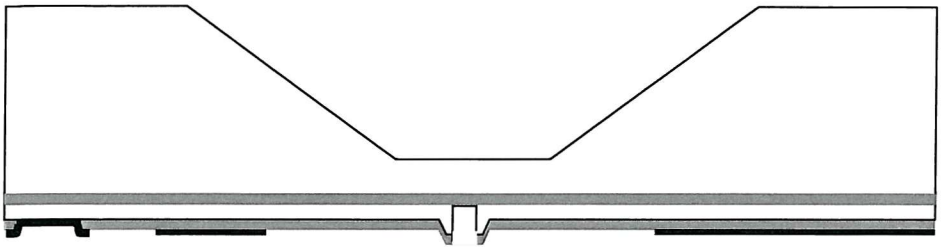

Apply resist and pattern it (OE3 layer).

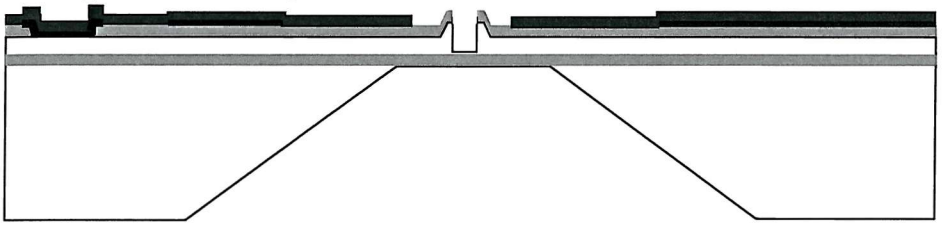

Wet etch away the exposed oxide layers.

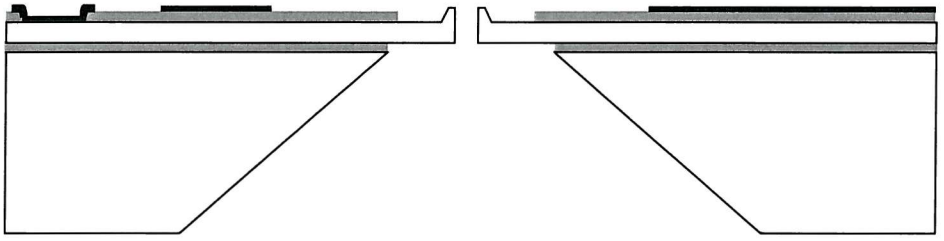

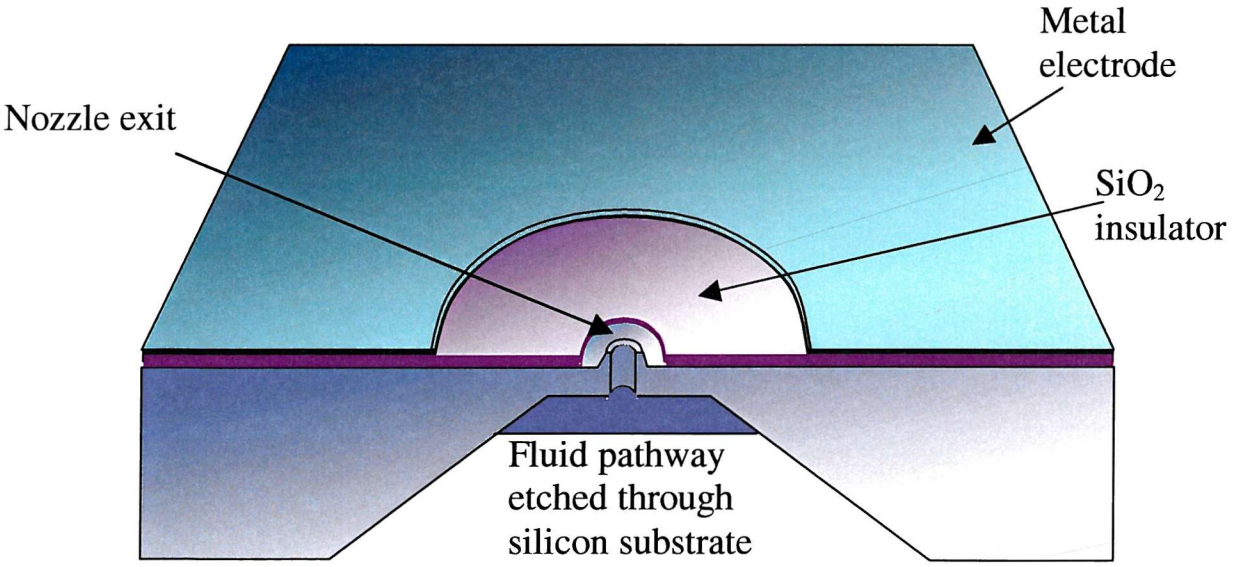

FIGURE 5.1: Schematic of fabrication process steps and final aim

5.2 Results of fabrication

5.2.1 Realisation of design

The fabrication was generally a success as the finished product was at least in some cases faithful to the idealised cross sections of the design. To illustrate the extent of the success a few scanning electron microscope pictures are presented here.

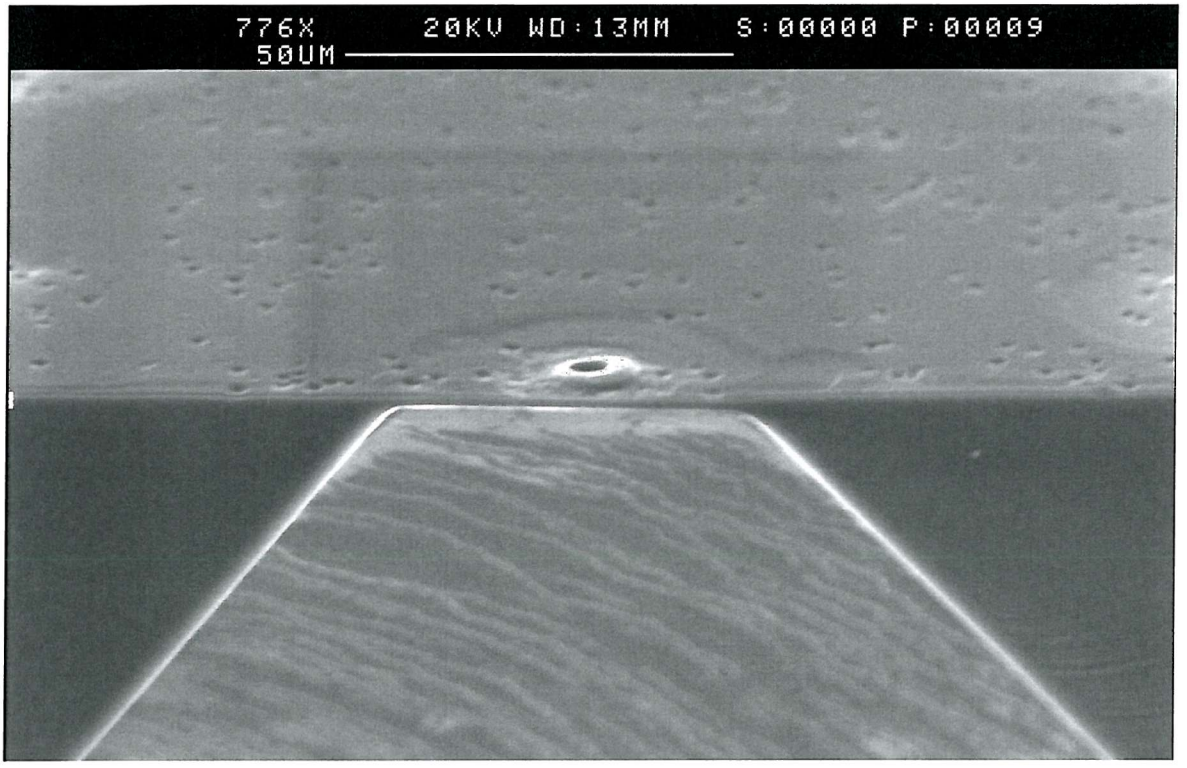

FIGURE 5.2: A tilted cross section SEM picture

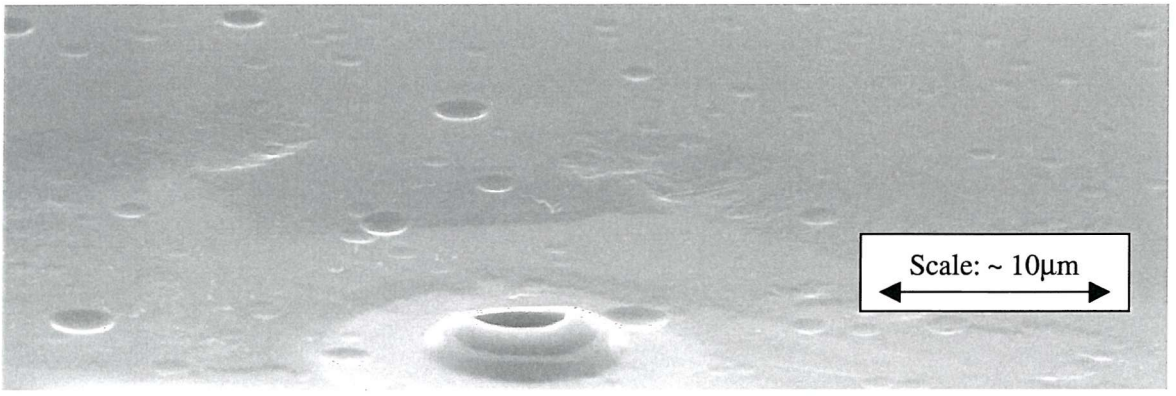

FIGURE 5.3: An SEM close up of a nozzle

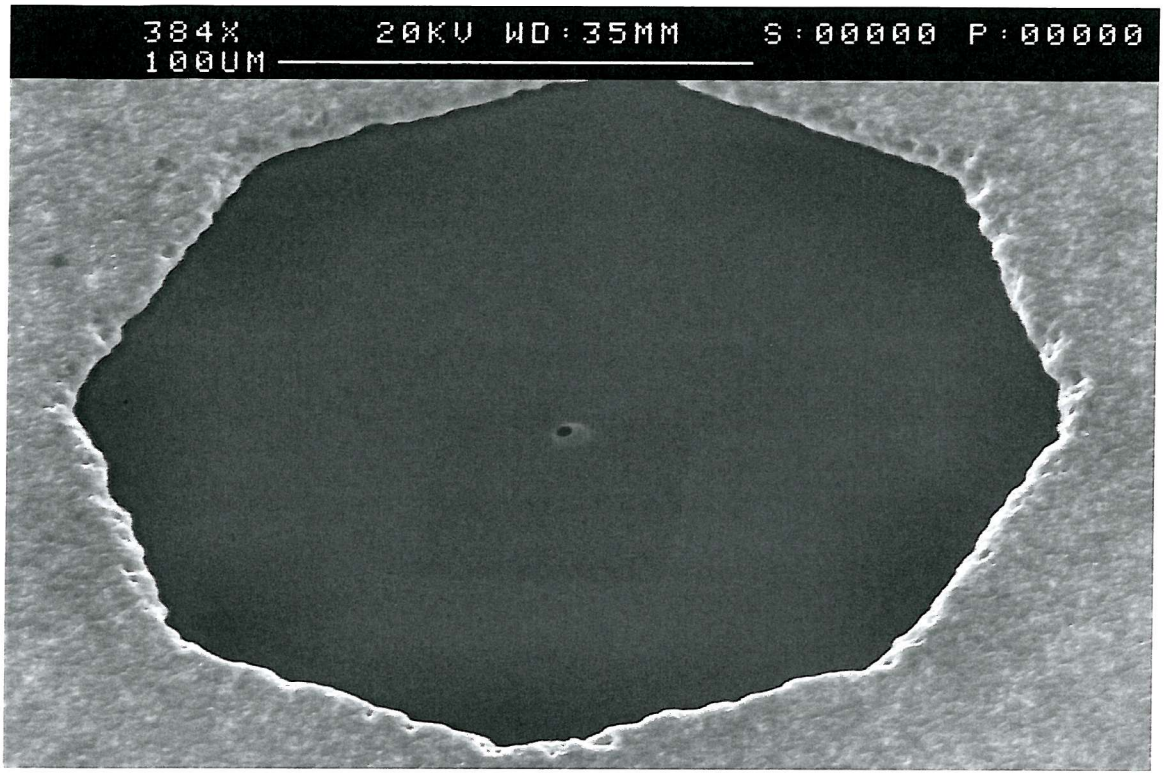

FIGURE 5.4: Tilted plan SEM picture of a 100um emitter.

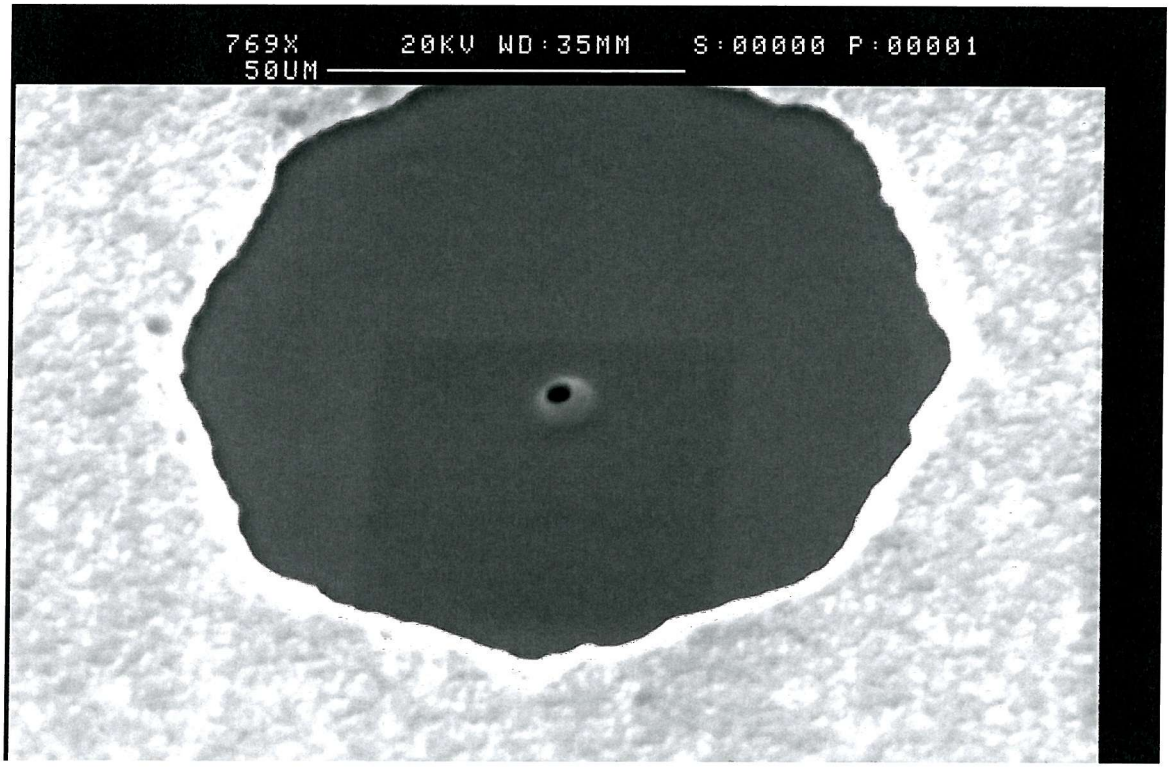

FIGURE 5.5: Tilted plan SEM picture of a 50um emitter.

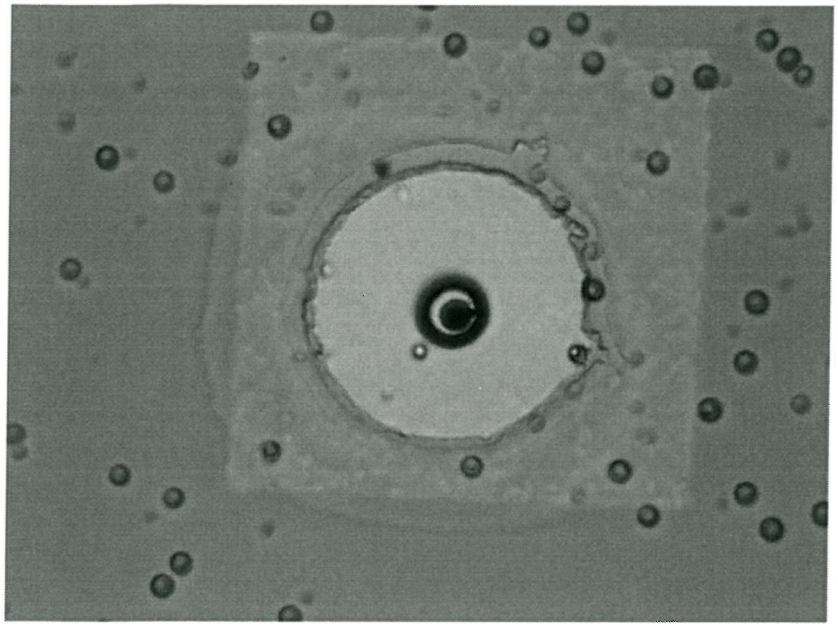

FIGURE 5.6: Optical micrograph of a nozzle.

In Figure 5.2 one of the wafers was cleaved allowing us to take photos of the cross section. This was the closest one to a complete cross section and shows the KOH etched pit and the remaining 2µm thick Si membrane. The nozzle can be seen above the membrane standing approximately 2µm proud of the surface. Figure 5.3 is a close up of a very similar nozzle (from the same wafer). The nozzle has the hole etched through it and a layer of oxide etched back around the side.

Figures 5.4 and 5.5 show two emitters with different sized apertures. On both the Aluminium can be seen around the edges with the aperture etched into it (the aperture is octagonal due to memory limitations of our mask making software). The nozzles are in the centre of the electrode aperture and it can be seen that they rest on top of the square membranes. The nozzles look quite accurately etched and the degree of alignment can be seen on the optical picture in figure 5.6. Also visible in this picture is the thin square membrane and the oxide which has been etched back from the nozzle exterior.

5.2.1 Problems during fabrication

This is not to say however, that there were no problems or that the entire batch of wafers survived and produced 100% yields. The truth is quite different. To understand the problems encountered and the root of deviations from the design it is useful to relate the results of the

fabrication process step by step, as recorded in the cleanroom records. This is done in the following table.

Step No.	Process Step Description	Comments on results
1	Scribe wafers	Wafers labelled on curve with scriber
	Create Boron etch stop	
2	Boron deposition, 1150°C for 6hrs	Did test run to identify operating conditions. Upon cleaving, examination on SEM revealed diffusion depth to be 6μm +/- 1μm. Glass nodules were left. See note A.
	Create alignment structures	
3	Photolith mask S1 on front	Successful; good resolution
4	Etch alignment marks in Si, 1μm deep.	Successful; marks were 1μm deep.
	Etch deep pits in back	
5	Wet oxidation 600nm, 1000°C	Within limits.
6	Deposit Si_3N_4, 160nm	Within limits.
7	Photolith mask WE on back, use double sided alignment.	Quite difficult – the alignment marks should have been reversed to allow easier locating. Alignment was good though – within ~ 10μm. See Note B.
8	Etch oxy/nitride	Successful
9a	KOH etch, stop on boron layer for wafers 4-6,	1st attempt revealed masking flaw. Not enough room between features allowed underetch of the mask and so the edges of the pits merged - destroying features. A new mask with more space solved the problem and etch was successful.
9b	KOH etch, stop 30μm from surface for wafers 1-3	Same initial problem as 9a. New mask worked.
10	Strip oxy/nitride	Successful.
	Make Si tip	
11	Photolith mask TS on front	Successful, good resolution
12	Deep etch Si (partially anisotropic)	Successful. 2μm high tips created with convexly curved slopes.
	Nozzle creation	
13a	Wet oxidation, 2μm wafers 4-6	Successful.

13b	PECVD oxide, 2μm, wafers 1-3	Successful.
14	Photolith mask O1 on fronts	Alignment was critical. Accepted results – holes were within 1μm of the centre – this varied across the wafer. See note C
15	Etch SiO_2 anisotropic	Successful.
16a	Etch Si HBR process to 6μm For wafers 4-6	Successful. Though hard to verify on smaller holes as neither Alphastep or microscope could penetrate into holes. Larger holes in test structure were correct depth.
16b	Etch Si HBR process to 8μm For wafers 1-3	Same as above. Assumed successful but verification difficult.
	Open up contact window	
17	Photolith mask O2 on front	Successful.
18	Etch SiO_2 through to Si	Successful.
	Create electrode	
19	Blanket expose & develop resist into holes	This step was very difficult. Technicians were unable to get resist to stay in holes. Decided to evaporate Aluminium from a 45° angle so it cannot get into the bottom of the holes.
20 alt.	Evaporate 1μm of Al Use oblique angling.	Successful. Inspection of the larger holes in the test structure revealed no Al in bottom.
21	Photolith mask EL on front	Successful.
22	Wet etch Al	Partially successful. Wafers 4-6 were successful. But wafers 1-3 (SOI) had trouble removing the Al from the apertures, required an overetch. See note D
	Open up holes	
23	Ion etch down to buried oxide For wafers 1-3	Partially successful. 1 wafer would not reveal the holes despite overetches. 1 wafer was successful, although in some portions the etch had completely removed the Si membrane.
24	Photolith mask O3 on front	Successful.
25	Wet etch SiO_2	Unsuccessful. All but 2 of the wafers were put into the HF and the photoresist lifted off exposing all the oxide and Al to the HF. The result was that most of the insulation was removed. See Note E. Also on removal of the oxide it became clear that some membranes were broken. See Note F.

TABLE 5.1: Results of important fabrication steps

NOTES

A. The glass nodules can be clearly seen in figure 5.6. They were left on the wafers in the hope they would have no effect on the fabrication and/or operation. In retrospect, this was probably a mistake, albeit a small one. The nodules are ~1-2μm in diameter, a scale that can interfere with the lithography, etching and therefore feature definition. This is probably the cause of the nick in the side of the nozzle in figure 5.6. In addition, the presence of these nodules can introduce non-uniformities into the surface of the insulator causing possible weak points for breakdown or field emission.

B. The alignment of back to front was only just within limits. The nozzle in Figure 5.4 is right on the edge of the membrane and indeed some nozzles were off the membrane – since the thickness of Si increases rapidly away from the membrane these nozzles are not open. However, this was a small problem and applied to only about 1% of emitters.

C. In the initial process definition this step was supposed to be done using E-beam lithography. However, the resist needed for the next step was too thick for this process so optical lithography was used instead, with its lower alignment tolerance of ±1μm. While the vast majority of the holes were placed right on the top of the nozzle, some were right at the edge of the tolerance as shown in figure 5.7.

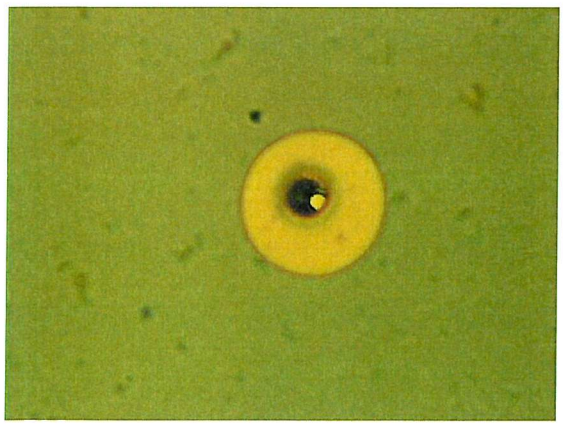

FIGURE 5.7: Misalignment of holes and tips.

On the SOI wafers the aluminium could not be etched away from some of the electrode apertures, even though the aluminium was exposed. An overetch was done and most of

the apertures cleared. Some remained however, one of which is shown in figure 5.8a. Clearly this is a problem as any electrode in which even one aperture has not cleared will be unable to sustain an electric field. Although only about 3-6% of the apertures were afflicted this way the problem has a much wider reach. The 12μN design (shown in figure 5.8b) consists of a number of electrodes with large numbers of apertures.

At a rate of 4% emitter failure, chances of survival seem slim for 24 emitter electrodes and the prospects of 36-emitter ones are very bleak indeed. This did prove to be the case. On the 12μN chips recovered from SOI wafers the survival rates were as follows:

36-emitter electrodes: <10%

24 emitter electrodes: ~20%

4 emitter electrodes: ~60%

As a result the only usable gates from the 12μN chips off SOI wafers were the small sub-clusters with an occasional 24 and (rarely) 36- emitter electrodes.

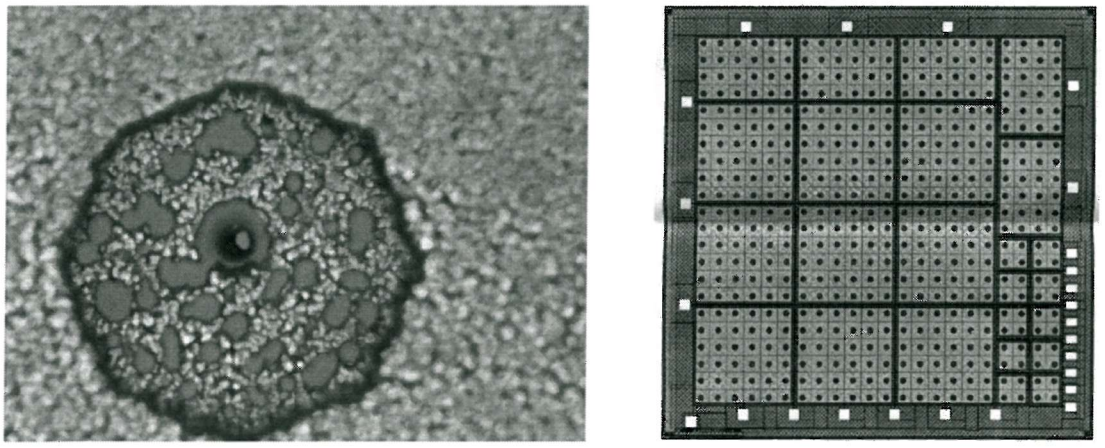

FIGURE 5.8: a) Metal after wet etch. b) emitter number in electrodes

E: This was a very unfortunate occurrence, caused by a lack of previous group experience about the use of resist on Aluminium in HF acid. The photoresist used usually exhibits acceptable adherence to SiO_2 and Si, although an older resist is preferred. However, I learnt that the same resist has very little adherence to aluminium when exposed to HF acid and as a result floats off from the surface of the metal, taking, through viscosity, resist from the surrounding areas as well. Clearly this is very undesirable, as the

aluminium is exposed to corrosive HF (see the very rough surface of the electrode in all photographs), but far worse, the oxide is exposed and completely etched by the acid. Figure 5.9 illustrates this damage very well.

It is clear as well that because the oxide has been removed from near the Aluminium the damage is intensified because the insulation is now removed from the most critical areas. In figure 5.10, the aluminium can be seen resting on top of the oxide. This oxide should extend almost upto the nozzle, but here has been etched back under the aluminium by almost 2μm. Clearly then the benefits of extending the oxide to restrict surface breakdown are lost. The distance for surface breakdown in these wafers is now only 2μm.

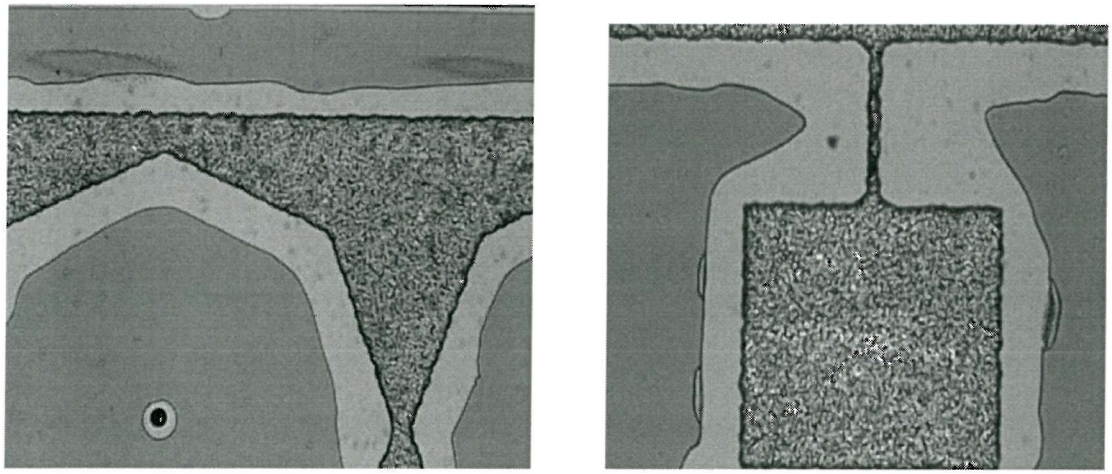

FIGURE 5.9: The extent of oxide removal from the wafers.

An attempt was made to repair one of the wafers by depositing oxide back over the surface in a post-processing step. This was considered unsuccessful however as the oxide coats the aluminium bondpads too, making it very difficult to make electrical connections.

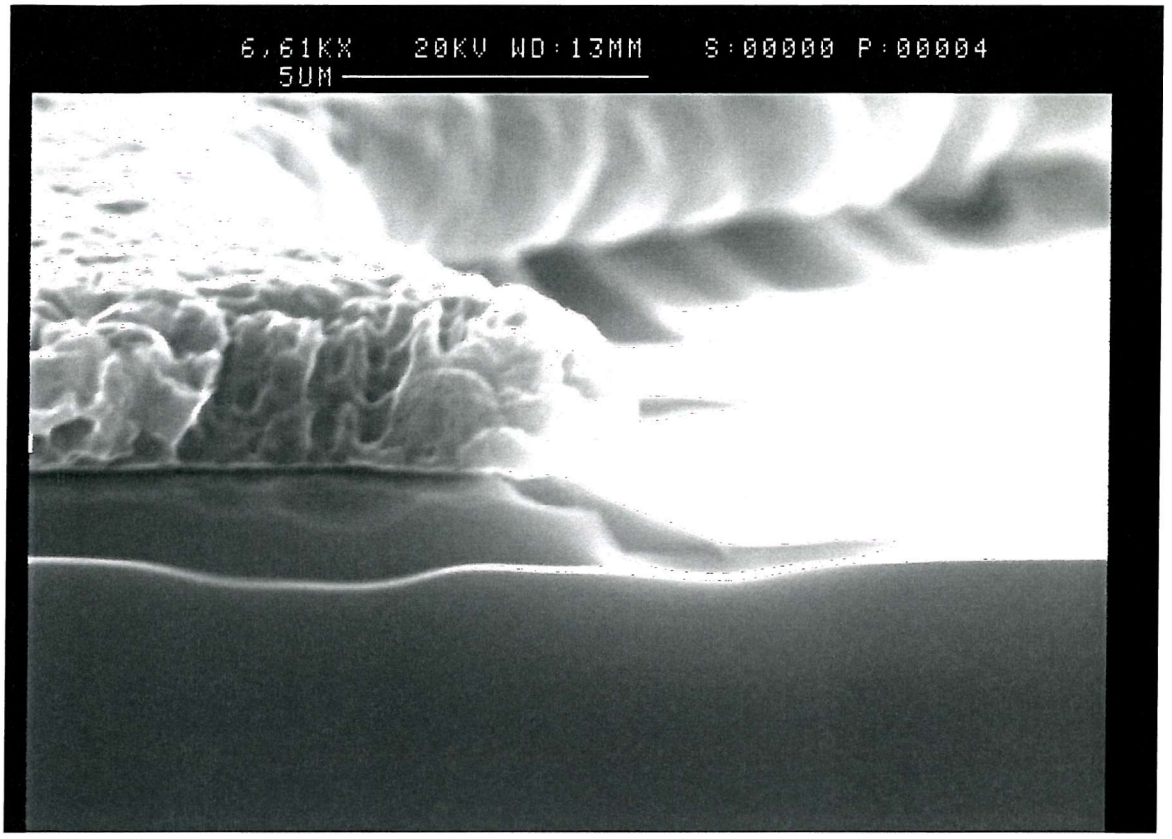

FIGURE 5.10: An SEM close up of the surface layers in cross section.

However, some of the SOI wafers were not processed in the HF at all so they have all of the original oxide.

F: The membranes created by the KOH etch were only 3-4 μm thick but were ~ 30μm wide. The application of an oxide, especially at the 1100°C of thermal oxidation, causes significant stresses in the membrane. In the majority of cases, this causes at least some buckling, and in a number of cases the membranes were destroyed, taking the nozzle with them. The extent of membrane destruction varied from wafer to wafer, with the best wafers exhibiting less than 1% failure rate. The consequence of membrane destruction is severe as it allows a large unrestricted exit path for the fuel. Any liquid coming into contact with the chip passes through the leak site and quickly spreads over the chip surface rendering it useless. However, this was restricted to the chips from the boron wafers since the SOI membranes were much thicker for longer.

5.2.3 Summary of fabrication results

At the end of the fabrication process, there were eight wafers; five of which were boron doped, and three SOI wafers. All of the SOI wafers suffered some damage during processing; two only survived as half wafers and one had ~ a fifth missing. This latter SOI wafer had suffered oxide damage. Of the boron wafers 2 were the initial wafers, which were found to contain defects but were left in the batch anyway. These were much more susceptible to membrane destruction. One of the other three boron wafers was post-processed with the CVD oxide and can be considered at the very least difficult to test.

In summary then I was left with:

- Two non-defective boron wafers all of which suffered from the oxide damage described but otherwise had chips that were of a good quality.

- Two half SOI wafers which had intact oxides and no membrane damage. These were of a very high quality except that a few of the electrodes were incompletely etched out as described in note D.

- One nearly intact SOI wafer, which had damaged oxide.

These wafers were the ones from which it was worthwhile taking chips for test.

Chapter 6

Apparatus design

An experimental facility was designed and assembled and is shown in schematic form in figure 6.1. It consists of a vacuum chamber with the required pumps, access and pressure gauging, a video microscope with the required illumination, and an independently pressurisable fuel chamber. A photograph of the rig is shown in figure 6.2

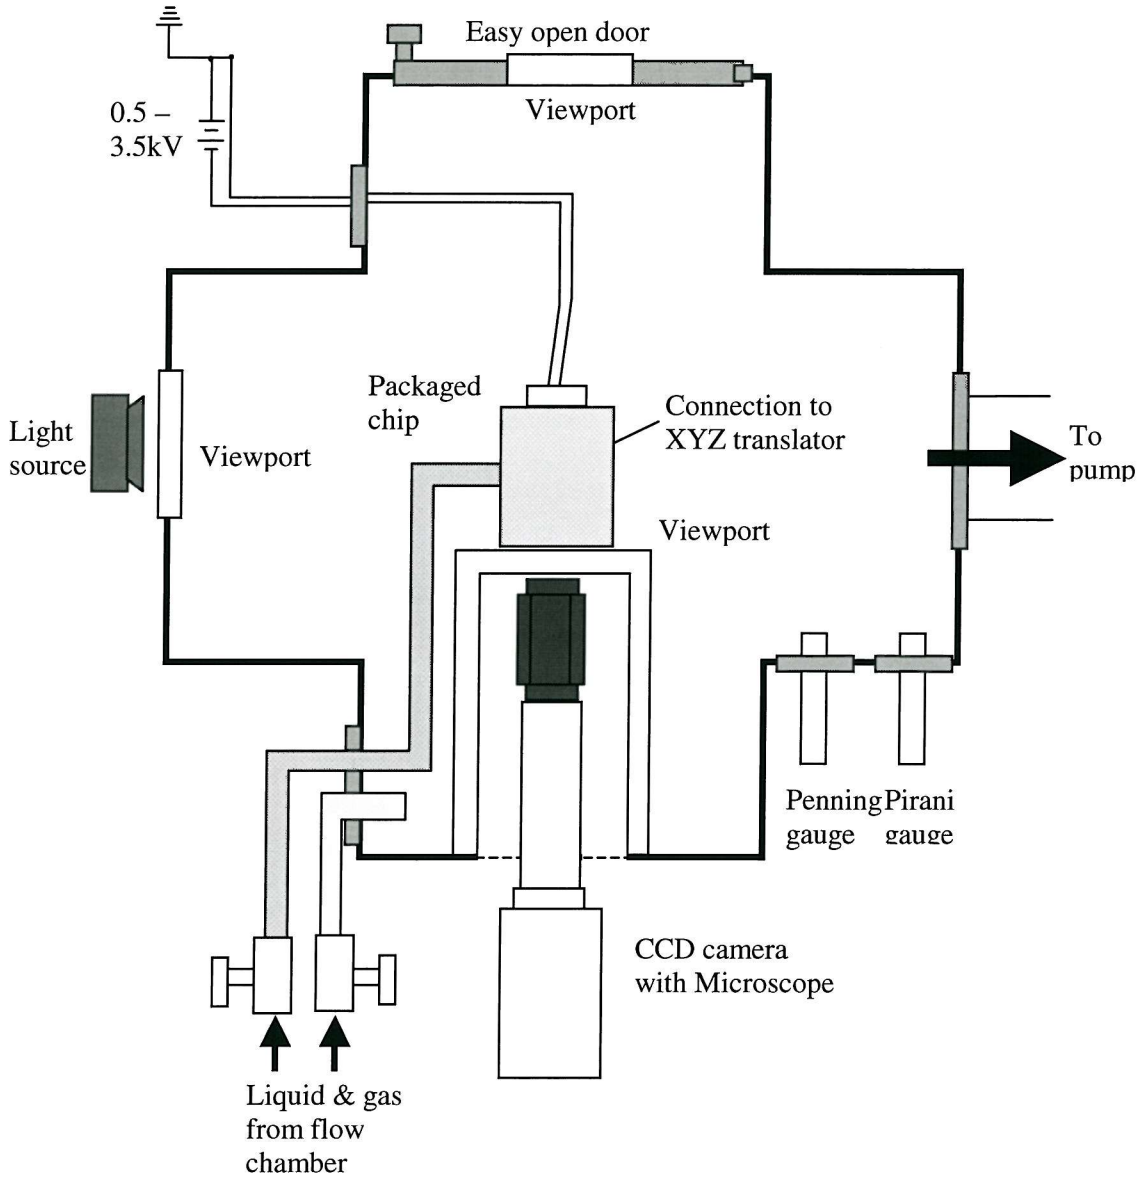

FIGURE 6.1: Schematic of the vacuum chamber set-up

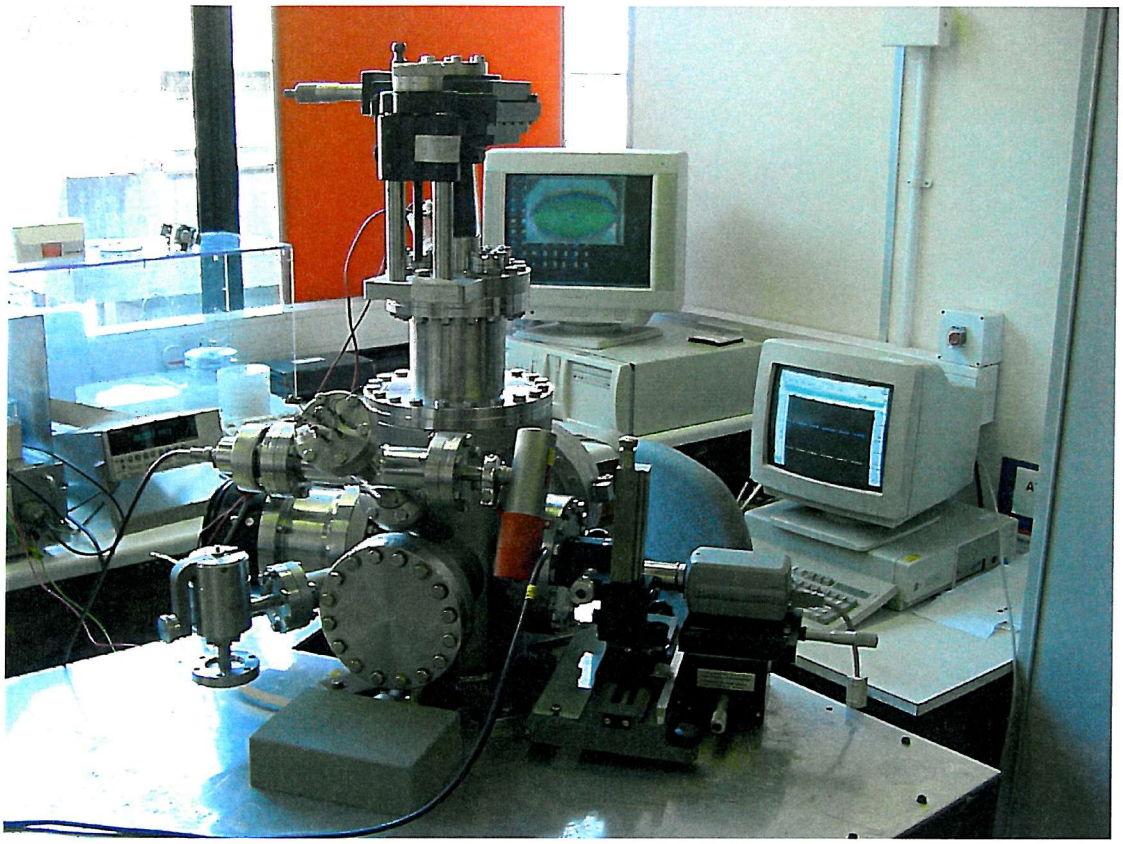

FIGURE 6.2: The experimental apparatus

6.1 Vacuum chamber

The vacuum chamber is a 3D-cross with numerous small diagonal ports. It has a volume of 30 litres and is pumped by a 250 l/s turbo pump and a 3.7m^3/hr rotary vane backing pump. This set-up is capable of reaching vacuum pressures of < 10^{-6} Torr.

The turbo pump is controlled in line with the rotary pump using a flexible coupling, which isolates the turbopump from the vibrations of the rotary. The rotary pump's exhaust is connected to the labs ventilation system and the turbopump is cooled by a tap water supply line at ~ 0.5 l/min. A turbopump controller operates both pumps and ensures that the turbo only accelerates past 50% full speed when the rotary pump has decreased the chamber pressure below 100 mbar. On shutdown the controller decelerates the turbopump and when the speed is at 50% it opens the vent valve. When the chamber pressure is sufficiently high the rotary pump is turned off.

The pressure is measured by a Pirani gauge (down to ~ $5*10^{-3}$ mbar) and a Penning ion gauge (for high vacuum) and displayed on a gauge controller. These pressure gauges are mounted on a small cross connected to the top of the main chamber.

Several electrical feedthrough ports provide electrical connection access for the chip, collector and light source (if needed); while a fluid feedthrough provides the access required to control the fuel flow from the flow chamber (described later). An easy open door (utilising a rubber o-ring seal) allows easy access for loading/adjustment without the need to undo nuts & bolts.

A number of glass viewports allow optical access. One is merely for visual inspection of the interior, useful for checking wires are still connected etc. The other allows the transmission of light into the chamber for the microscopy. The third is a glass tube which projects into the chamber interior allowing microscope access right into the centre of the chamber.

A precision XYZ controller is mounted onto the top of the chamber and penetrates inside. It is used for mounting of the thruster and the translation capability allows us to scan across the surface of the chip and position it into the focal plane of the microscope.

6.2 Microscopy

Due to the extremely small scale of the emitters, it is essential to have a microscopy capability. With a nozzle exit diameter of 5μm or less it is impossible to tell what is happening to the liquid at the end of the nozzle without very high magnification. The chips must be operated in the vacuum chamber but most microscope components are incompatible with vacuum conditions. Therefore, any microscope set-up must have a very long working distance so it can be external to the chamber.

The microscope was designed to allow us to observe and capture images of the liquid surface at the nozzle exit. This is needed so one can determine whether the liquid is being held at the nozzle exit, whether a spray develops and what kind of regime the spray belongs to. It also should allow us to observe arcing and determine the arc sites and therefore danger points on the chip.

An assessment of ready-made video microscopes available on the market revealed them to be prohibitively expensive and not altogether suited to our requirements. It was felt that a large proportion of the secondary components could be developed in the lab at considerable cost saving. Clearly though two major components, the camera and microscope objective, would have to be purchased. The specifications for these was decided on using these requirements:

- Resolution: <1μm
- Field of view: as small as possible, preferably 200μm or less
- Working distance: as large as possible – but at least > 10mm
- Depth of field: at least 3μm

The equations used to design a video microscope system are as follows:

$$Field\ of\ view(FOV) = \frac{CCD\ size(horiz)}{\text{Pr}imary\ magnification} \tag{6.1}$$

$$System\ magnification = PMAG \times \frac{Monitor\ size(diag)}{CCD\ size(diag)} \tag{6.2}$$

$$Camera\ resolution(\mu m) = \frac{1000}{\left(TVL(horiz) \times \frac{1.33}{2 \times CCD\ size(horiz)}\right)} \tag{6.3}$$

$$System\ resolution(\mu m) = \frac{Camera\ resolution(\mu m)}{PMAG} \tag{6.4}$$

Using a 1/3" CCD camera (4.8mm by 3.6mm) with a 20X lens gives:

FOV = 240μm.
System magnification = 1266X (on a 380mm diag monitor)
Camera resolution = 15.3μm
System resolution = 0.76μm

The only 20X lens that is suited to this task is an infinity corrected objective as these have a much longer working distance. The Newport 20X-LD has a working distance of 17mm which is ample for our needs. The Sony DFW-V500 has a 1/3" CCD sensor and provides 30fps of full colour video. This combination is capable of meeting our requirements with the addition of one extra part; the tube lens which allows the objective to be used with a camera.

The rest of the microscope assembly was pieced together from the body of an old microscope in the lab and a few extra interfacing parts found lying around. This system was very inexpensive and allows for coarse positioning of the microscope.

An essential part of any microscopy set-up is the illumination. If done poorly even the best microscope will give bad results. In our situation, the challenge is made more difficult because of the need to illuminate the chip in the cramped confines of the chamber. Furthermore, there are a variety of aspects it is desirable to view, each of which requires a different type of light source. Firstly it is desirable to view the emitter surface i.e. silicon, aluminium and oxide for defects. This is best done with the plane of the chip normal to the axis of the scope. To avoid the image being ruined by glare or shadows a uniform light source is needed, this is best achieved by a ring light, which gives uniform illumination if angled correctly. Again however, such lights are expensive and generally are unsuited to both our required working distance and limited space.

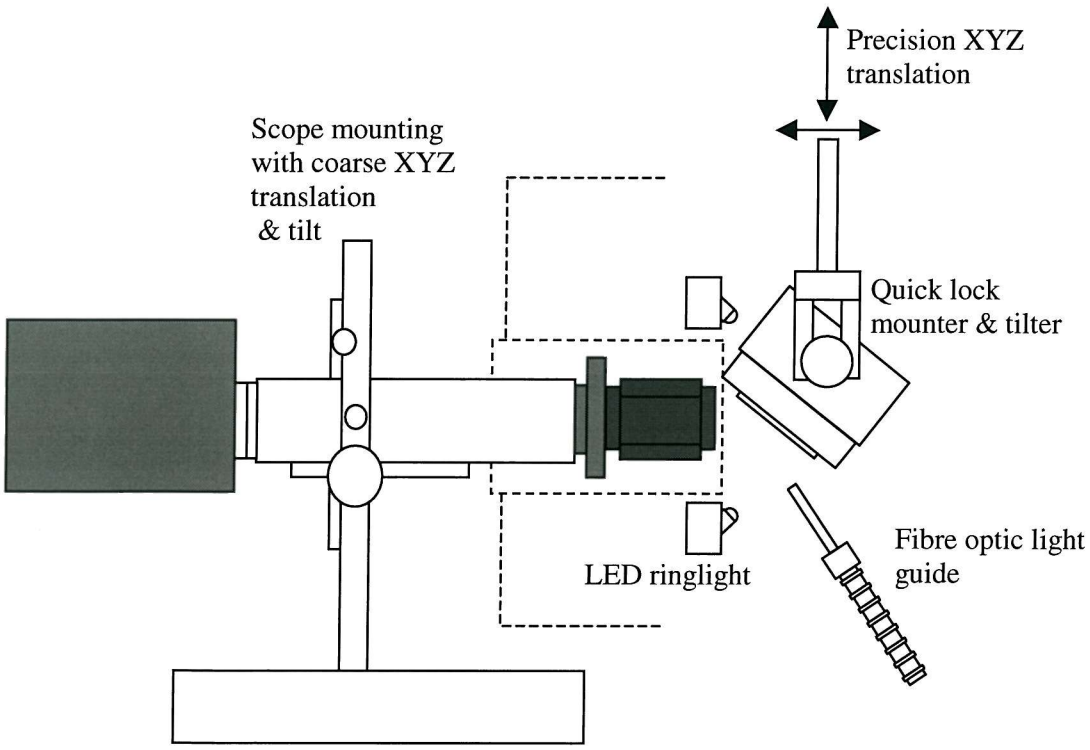

FIGURE 6.3: Microscopy set-up

An attempt was made therefore to provide a basic ringlight illumination for use in the chamber. Superbright LEDs were wired in parallel on a piece of breadboard, and a hole with the diameter of the re-entrant glass port was cut in the board. Each LED is in series with a suitable resistor to limit the current through each LED to 20mA for a 12V and 500mA power

supply. Each LED gives a light intensity of 9.3cd (where cd is a candela, a unit of light intensity), so wiring up 20 gave us a 186cd intensity.

This basic set-up allows us to check the surface of the emitters easily and the nozzle can just about be seen, as shown in figure 6.4. It is most useful for checking the integrity of the electrode and the wire bonds and allows us to locate the specific area of the chip under investigation.

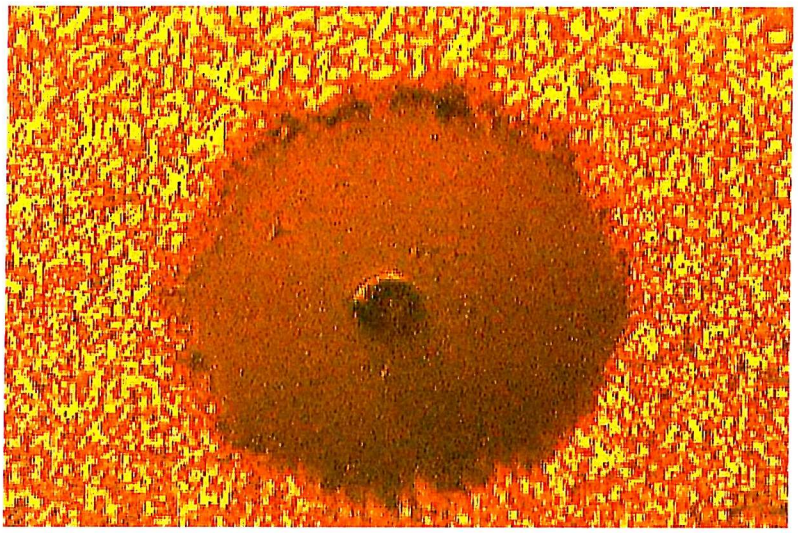

FIGURE 6.4: An emitter in plan using LED ringlight.

However, this type of illumination only gives 1 to 2μm depth of field, which makes it almost impossible to view the emitter in any kind of tilted view. To be able to see any jet, spray or even droplet formation the emitters must be tilted so that more of the liquid is visible. In a tilted view though ordinary illumination means only thin strips of the emitter are in focus at any one time. This is most inconvenient.

A very specific type of illumination called glancing directional illumination completely solves this problem. By directing bright white light at a very acute angle to the surface (~10° above the surface plane) a diffraction effect allows practically the whole of the emitter to come into focus. Light passing through the oxide and bouncing off the rough surface of the aluminium is diffracted, and some portion of the light from each point is sent into the lens at the angle required to bring that point into focus. Not only does this technique give a depth of field of approximately 50-100μm but it also allows the oxide to be seen in detail, as well as flaws on the surface. If the light is set to just the right angle then the nozzle can be seen fairly distinctly, as is illustrated in figure 6.5.

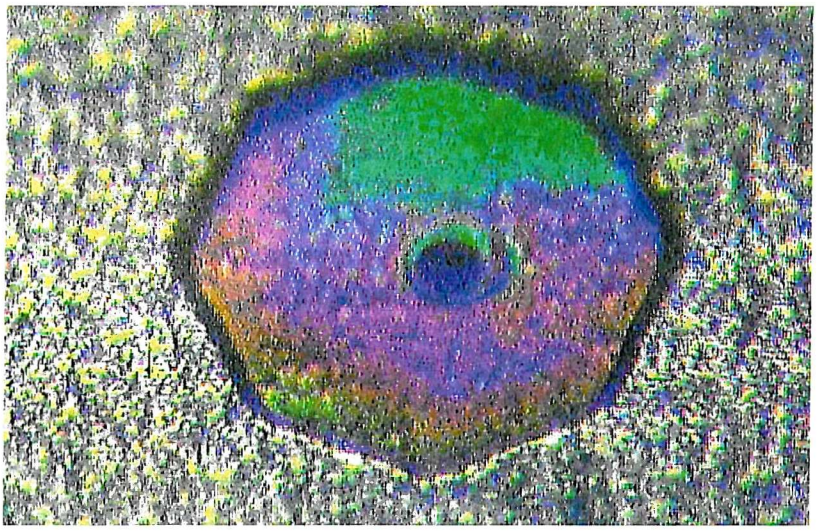

FIGURE 6.5: Tilted plan with glancing illumination

This illumination was achieved using a Halogen cold light source, operated outside the chamber. The light from this source is channelled through a viewport into a flexible optical-fibre light guide inside the chamber. The guide is used to direct the light at the correct angle.

This technique can be used at many different tilt angles; the light source direction just has to be re-adjusted. Figure 6.6 shows a greater tilt in which the profile of the nozzle exit can be seen. This kind of view would clearly show the formation of a liquid Taylor cone on the nozzle – vital for determination of the spray regime.

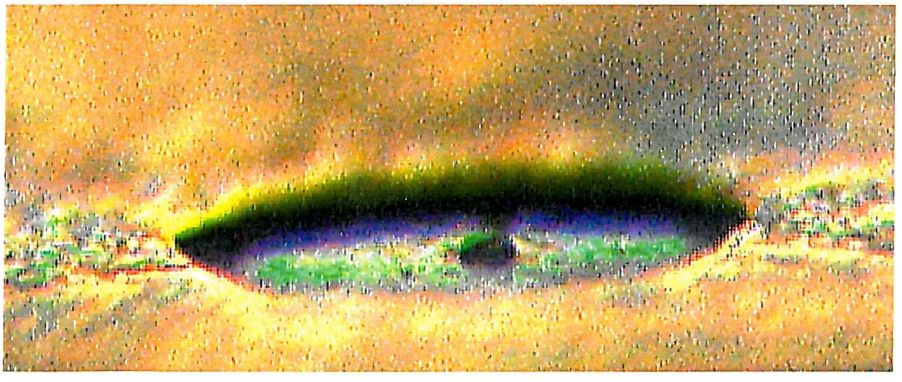

FIGURE 6.6: Extreme tilt using glancing illumination.

One more bonus of this illumination technique is the way that the liquid can easily be seen when it leaves the nozzle. In the next photograph (figure 6.7), a drop of formamide can be seen spreading across the surface of the emitter having left via the nozzle exit. It is very

distinct and can be made even clearer by subtle adjustment of the light direction to make the oxide appear a vivid contrasting green as in figure 6.6.

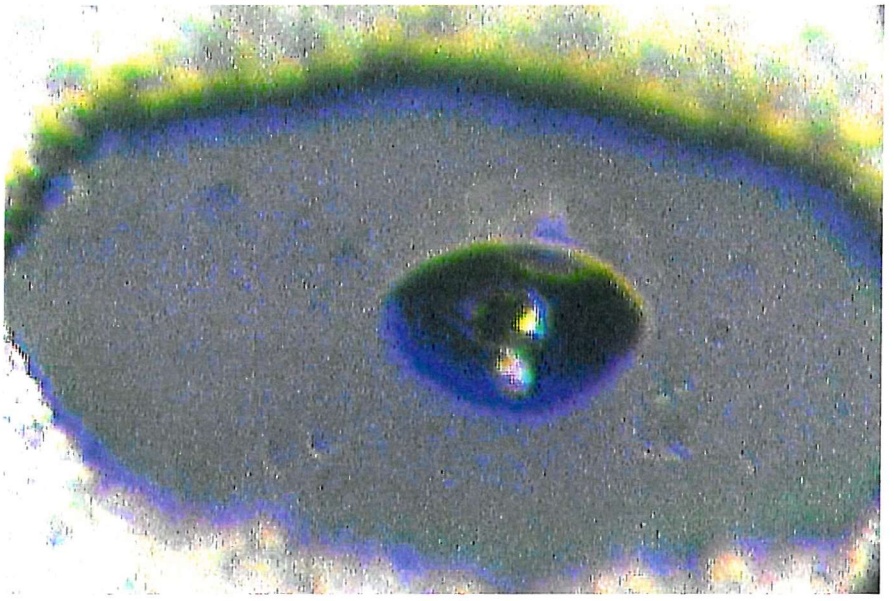

FIGURE 6.7: A drop of liquid formamide on the emitter surface.

6.3 Thruster chip packaging

6.3.1 Packaging considerations

The silicon chips are extremely fragile. At their thickest point they are only a third of a millimetre thick, and by far the majority of the material is less than this due to the deep KOH etches that take up much of the area. Each of the membranes and the immediate surrounding area provides a clear starting point for a fracture especially since some of these are already buckled. This situation is exacerbated by the fact that the emitters were placed in straight lines across the wafer - in the same directions as the crystalline lines of the silicon. As a result, they must be handled with a great deal of care – difficult in a laboratory-testing situation.

Every MEMS design is only half complete if no attention has been paid to the packaging of the final diced chip. Not only does the chip need to be protected from harmful environmental effects but it must also be exposed to selected aspects of that environment. The chip needs to

be interfaced to the outside world, either to be able to do its job, for diagnostic reasons or both. The vital packaging requirements to be considered are:

- Electrical connections
  - Connection to chip surface
  - Interchangeable connection to chambers electrical wiring

- Mechanical
  - Fairly rough and frequent handling during testing
  - Allow open and optical access to front surface
  - Fix chip firmly in position

- Liquid feed
  - Completely vacuum tight interfacing with fuel
  - Must be easily interchangeable from chip to chip
  - Must be washable / easy to purge

- Material choice
  - Compatible with fuel
  - Compatible with cleaning processes
  - Electrically suited.
  - Strong and robust enough to handle testing process

- Mounting system
  - Must be rigid and resistant to disturbing forces
  - Ease of load / unload important (confined chamber)
  - Able to position chip in a range of locations and attitudes

- Whole package system
  - Must be easy to change chip under test.
  - Minimum of chip specific packaging.
  - Low cost – large number of chips to test.

6.3.2 Initial packaging

The initial packaging design was unacceptable and as a result, most of the chips from the first wafer were destroyed during the first round of tests with a very low yield of useable results. It is presented here to illustrate the extent of the problems to be overcome.

The chips were packaged into thruster testing assemblies as shown in figure 6.8. Hinged plastic sample boxes formed the base of the packaging, acting as a structural mount and protection for the chip. The lid could be opened during testing but kept locked at all other times. A square hole was filed in the bottom of the case with sides 2-3mm smaller than the chip and the chip was superglued over the hole to the case. A swagelok fitting screwed into a plastic hose was then bonded to the rear of the case into which a fuel line could be attached.

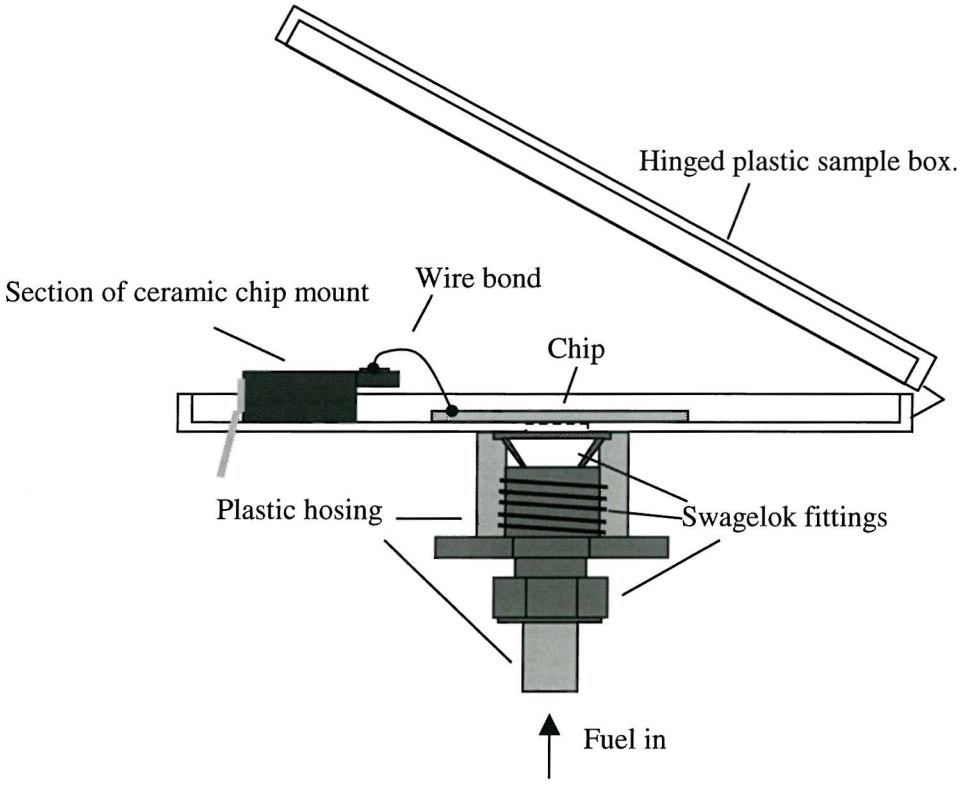

FIGURE 6.8: Schematic of the initial packaging system

Electrical connections to the chip were via ultrasonic wire bonds between the chip and customized sections of ceramic chip mount, also superglued to the casing.

The problems with this packaging design were numerous:

- The plastic cases are incompatible with both IPA and acetone, which are very useful for cleaning contaminants and misplaced fuel from the chip. Upon exposure to these

fluids the plastic becomes opaque, increasingly fragile and eventually breaks under its own weight.

- The plastic is not very strong and is susceptible to fractures – it is unsuitable as a roughly handleable substrate.

- The gold wires and the associated end bonds are very fragile – the lightest touch can break or de-bond them. In this design they are too exposed.

- The use of superglue is inappropriate as the adhesive is corrupted by acetone and to some extent formamide too. This resulted in the chip, the ceramic chip-mount section or often both becoming de-bonded. Many wire bonds and some chips were destroyed as a consequence.

- The liquid interface was wholly unsuitable, partly because of the use of superglue but also because the swagelok fitting was not tight enough in the plastic hose section. As a result, the liquid often leaked severely from the interface as well as from the sides of the chip bondline.

- No provision was made for the mounting of the chip package in the chamber.

Because of the failure of the above design a new packaging system was developed which solves or avoids the problems identified. It is based on MEMS packaging best practice complemented by a few techniques more suited to our particular application.

6.3.3 New packaging

The new packaging system is shown in schematic form in figure 6.12 at the end of this section. The various components are discussed below, with emphasis on how the problems identified above are solved.

LIQUID INTERFACE

The best way to securely bond a silicon chip to a substrate is to select Pyrex as the substrate material and use anodic bonding. In this process, the silicon and Pyrex are heated on a hotplate up to 400°C. At such temperatures, the silicon becomes fully conductive and the ions in the Pyrex can move about much more freely. When a strong electrical field is applied (between the positive silicon and the grounded hotplate) the ions in the glass move towards the hotplate creating a depleted region next to the silicon. This creates a permanent and very strong bond between the two. It is a bond that, if allowed to fully form, can only be undone by destroying one of the materials.

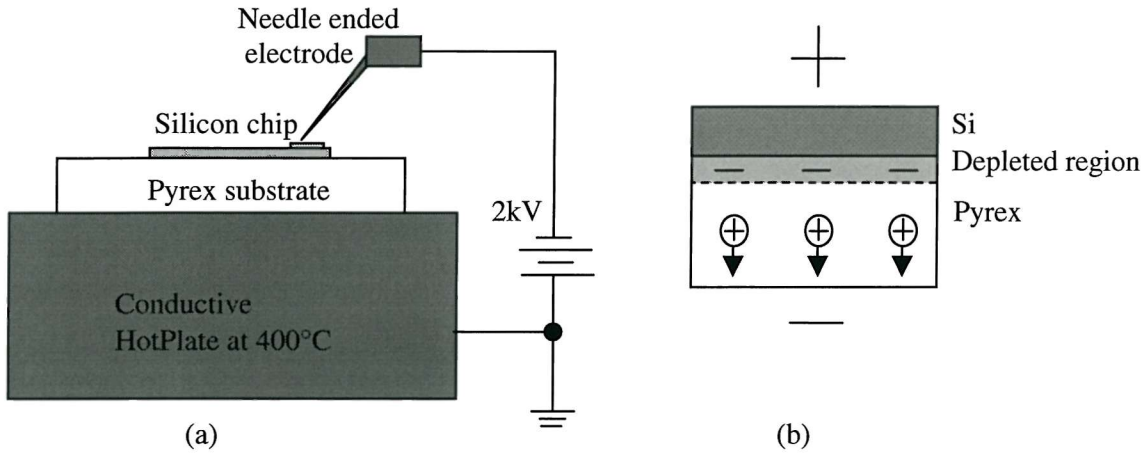

FIGURE 6.9: Anodic bonding. (a) Set-up, (b) Physics of process

The anodic bond is highly suited to our needs as it firmly fixes the chip to a substrate in such a way that no liquid can leak from between them. Furthermore, Pyrex is very attractive because it is strong, robust and fully compatible with the chemicals to be used. Since it is an insulator it also acts as isolation between the chip and the chamber if required. A hole slightly smaller than the size of the chip to be used is machined into the glass using diamond drill bits.

Anodic bonding is not without its problems; the surfaces to be contacted must be clean and without contaminants as well as relatively flat, or the bond may not be fully formed. This has proved to be the case in several instances where because of the presence of a random and so far unidentified layer on the backs of some chips the bond has only formed around 70-80% of the bondline. An incomplete bond is doubly unfortunate because of the irreversibility of the process – the chip cannot be removed without breakage.

However, the problem can be satisfactorily solved by the use of epoxy resin. Although this adhesive is only marginal if used as the only or primary means of affecting a vacuum tight bond, as a backup to anodic bonding it is fine. The anodic bond pulls the surfaces into close enough contact that the application of a line of resin around the periphery of the bondline is sufficient to form a vacuum tight seal.

The rest of the liquid interface is made using mechanical means. The liquid pipe from the fluid cell is fed to the chip package through a plastic tube with swagelok fittings on the end. This fitting connects to a swagelok fitting welded to a stainless steel backing plate. A join between this plate and the Pyrex substrate is effected by a nitrile rubber o-ring which is compressed between them by nuts and bolts. This final element allows chip packages to be simply and quickly interfaced with the feedline.

ELECTRICAL CONNECTIONS

The electrical connections are still made by ultrasonic wire bonding the gold wires between the chip bondpads and the chip-mounts, but now the chip mounts are bonded to the Pyrex using the resin. A strip of plastic is clamped down onto the surface of the chip-mount by the o-ring bolts; this acts as further security against de-bonding but also as protection for the wire bonds on the chip-mount.

MOUNTING

Mounting of the chips is achieved by a steel plate which is secured on the o-ring bolts. The rear of this plate has one half of a quick lock mounting system secured onto it. The mounting system is a steel fork into which a circular holder can be slid, once in place and adjusted to the desired rotation the holder is held securely. The fork is connected to the end of a rod, which represents the end of the precision XYZ translation system.

6.4 Liquid pressure cell

It is important that any liquid fed to the thruster chips is at a prescribable and most importantly determinable pressure. The most pressing reason for this is to stop leakage as described in the results section. Secondly, the liquid flowrate is strongly dependent on the pressure difference between the liquid and the pressure in the chamber. To give us this capability a separate system has been devised, which interfaces with the main chamber.

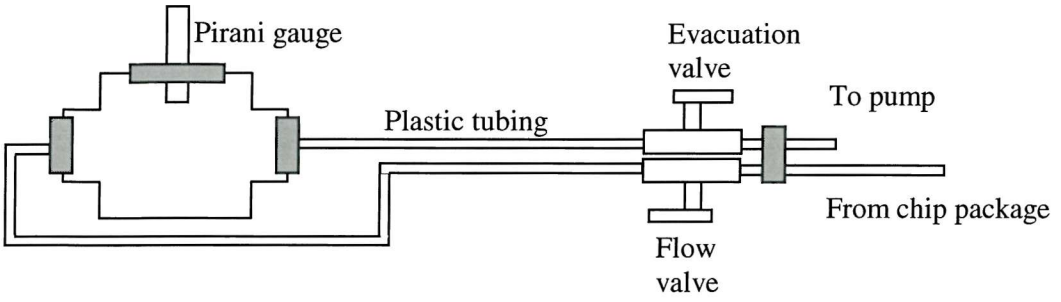

FIGURE 6.10: Schematic of the liquid pressure cell

Shown in figure 6.10, the pressure cell consists of a stainless steel mini-chamber connected to the main chamber by plastic piping and valves. The liquid is stored in the chip package

The Evacuation valve and FLOW valve are used to control the pressure the liquid is subjected to. The pressure under which the liquid is stored is measured using the Pirani gauge connected to the mini-chamber. The procedure for prescribing a liquid pressure is:

- With both valves open the vacuum pumps are started.
- When the pressure in the liquid pressure cell is about 10mbar the EVAC valve is closed.
- Since outgassing in the mini-chamber would cause the pressure to rise, the FLOW valve is then closed.

The liquid in the chip-package is then held at 10mbar and while this can rise due to evaporation over long periods, the pressure cell can be used to check and if necessary correct it.

The liquid is stored in the chip package partially for simplicity reasons but also because doing so allows us to electrically de-couple the liquid from the chamber walls and other external grounding points. This means that, unlike in an earlier design, potentials can be applied to the liquid or, equivalently, to the silicon substrate. A photograph of the chip package is shown in

figure 6.11. An extra Pyrex fuel plenum has been added to allow safe storage of a decent amount of fuel with no contact to the mounting, this too can be de-coupled from the chamber.

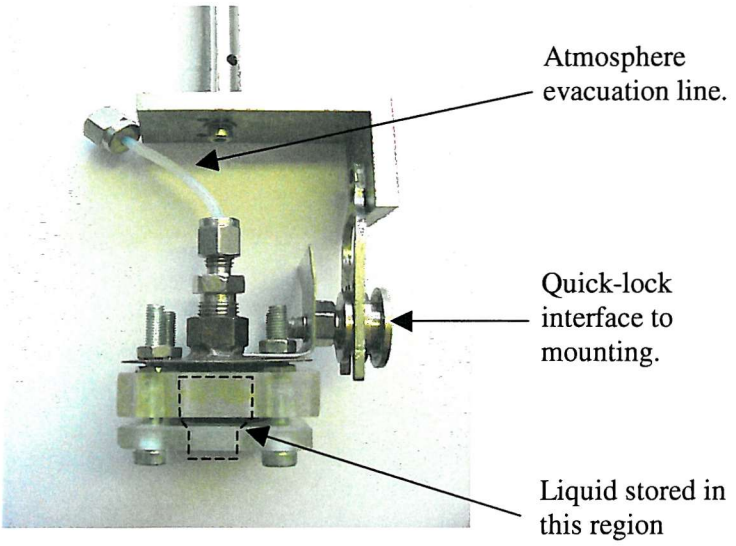

FIGURE 6.11: A chip package with fuel plenum and connections

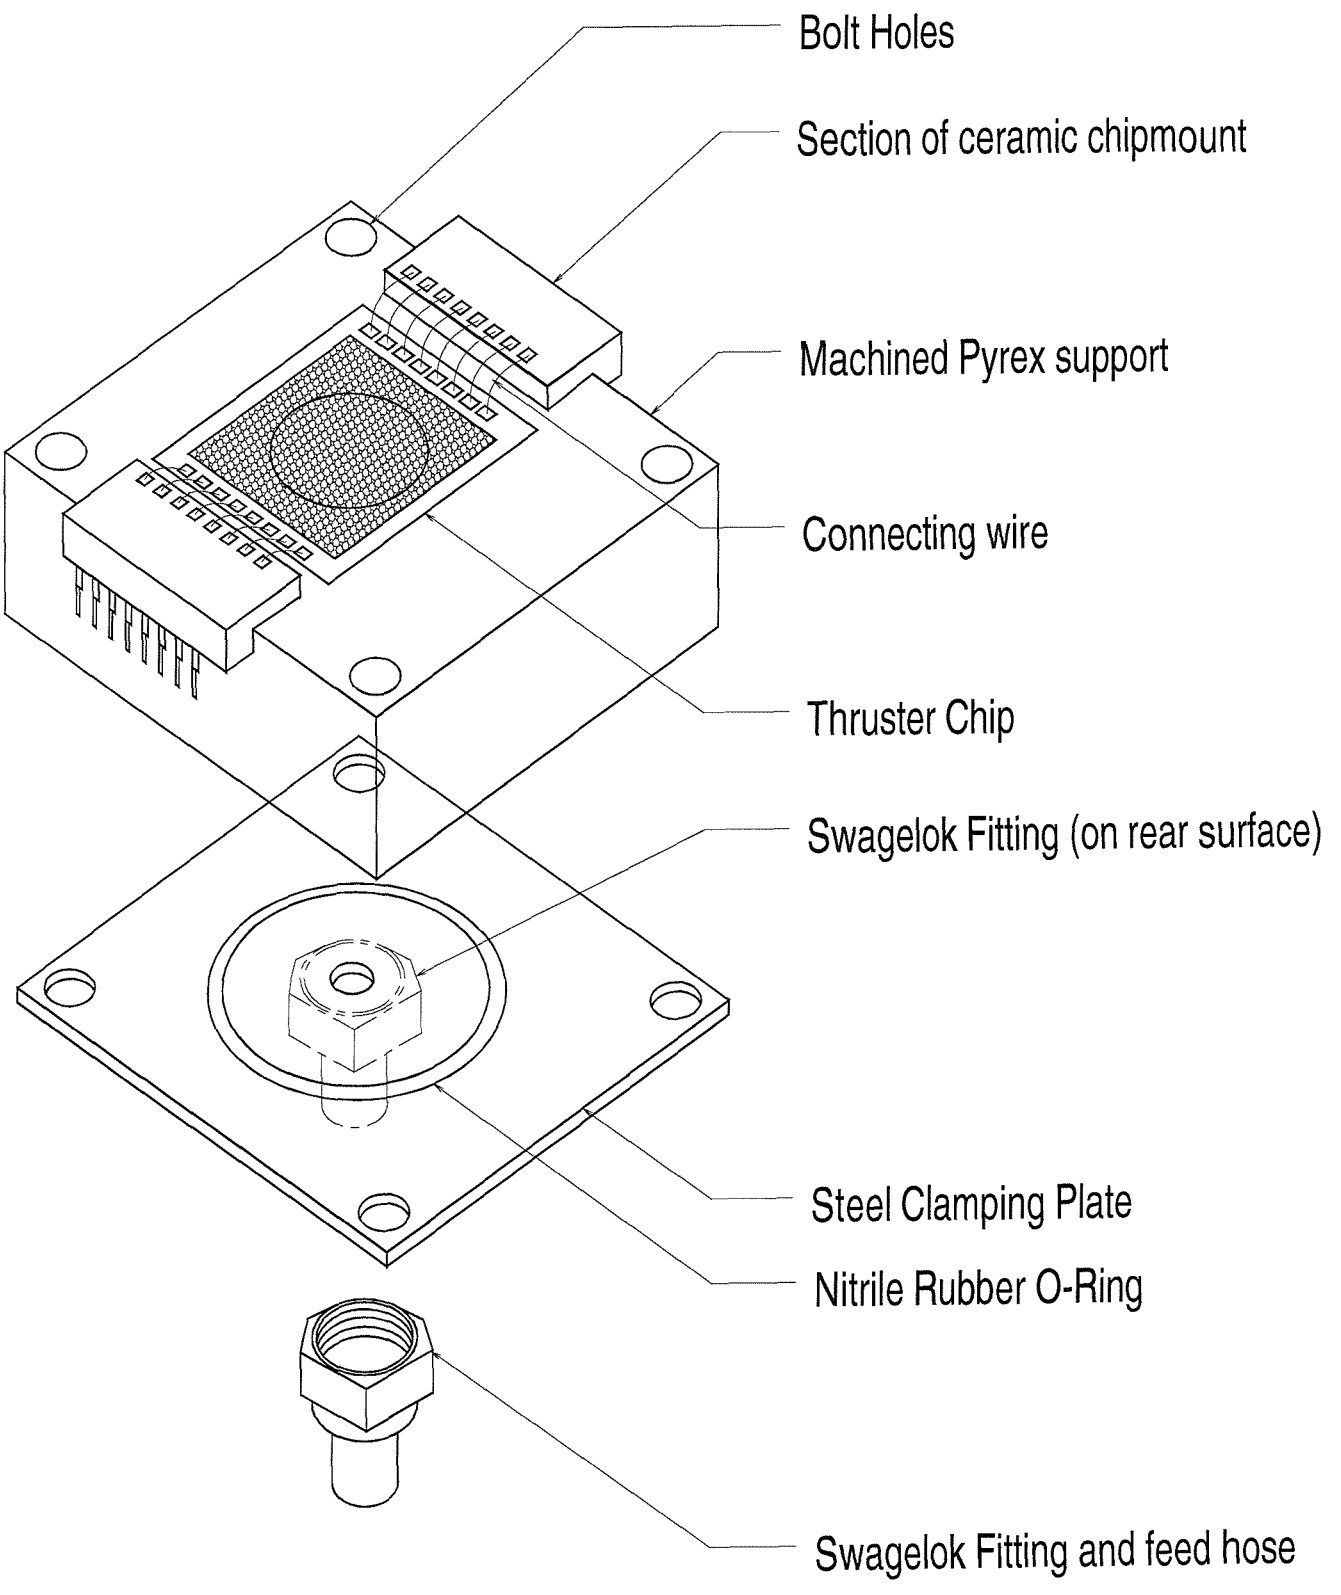

FIGURE 6.12: Main elements of the packaging system.

6.5 Instrumentation

Apart from the microscopy, a few other methods are employed to determine what is happening to the chips during testing. The main method is the use of electrometers to measure currents and potentials. A wiring circuit for doing this is shown in figure 6.13.

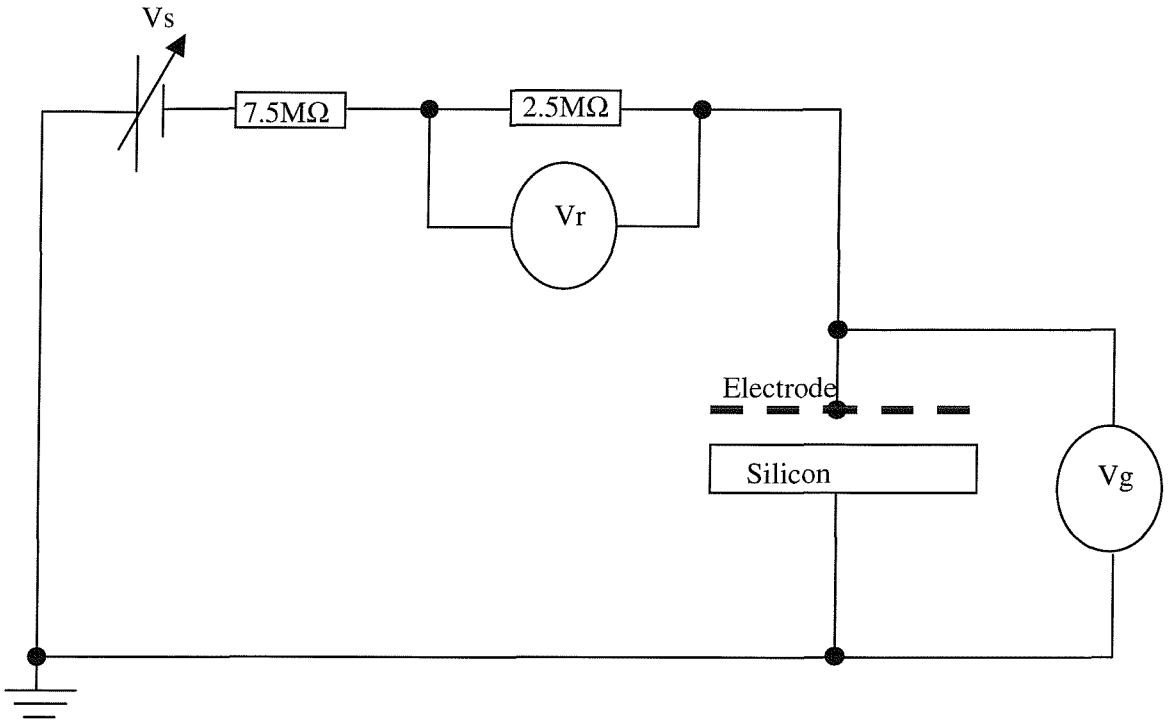

FIGURE 6.13: wiring for the chip testing

A variable power supply is used that can go upto 3.5kV with a maximum of 30mA. Because the power supply has a minimum turn on voltage of 750V the potential divider shown in the figure is used so the minimum voltage one can apply drops to ~400V. The resistors in the circuit act as current limiters so the maximum current that can be drawn if the chip breaks down is < 0.4mA, since the intent is to use little more than 1 kV the current is limited to 100μA. The voltage Vg is measured on a Fluke multimeter and measures the potential being applied directly to the chip electrodes, while Vr is a measure of the current the chips are drawing. Because of the high voltages used safety was paramount so a common grounding point was used which was the stainless steel chamber, itself connected to earth through the mains. Coaxially shielded cable was used where possible and the measuring circuit was contained in a grounded metal shield.

Chapter 7

First results & further experiments

This chapter reports the results of the first round of experiments on the thruster chips and the requirement for more detailed experiments into electrical breakdown. The preparation for the advanced experiments is described.

7.1 Effect of liquid pressure on leakage

The progress of the first experiments was inhibited by fuel leakage. When the liquid leaks at uncontrolled flowrates it first becomes a large drop. At some point, this drop starts to spread across the surface of the emitter. This is aided because the oxide is hydrophilic and in this regard formamide is very similar to water, so the contact angle encourages wetting. Once the liquid has spread to the edge of the aluminium that emitter and others with the same electrode cannot operate since the electric field is destroyed.

As described in chapter 5 a number of the membranes were destroyed during fabrication of the boron doped wafers. The loss of a membrane means that the liquid passes through a large square hole (30 x 30μm) with a length of only 2-3μm, this offers very little flow resistance. As a consequence, liquid quickly and disastrously covers the whole surface of a chip with even just one lost membrane. Many of the first wafers suffered from this setback. A solution that seems to work is to cover over the lost membrane with epoxy resin. This is a suitable sealant at the dimensions needed, provided the liquid is not applied at too high pressures.

Even with lost membranes sealed, or with a chip that has all the membranes intact, leakage was still being seen when the liquid was applied to the chips, though at a slower rate than before. Observation with the video microscope allowed us to diagnose the problem effectively. The initial, and quite serious, concern was that the oxide around the tips was encouraging wetting to the extent that it was overcoming the surface tension of the hole exits.

This concern was alleviated however by real time observation of the emitter holes under the video microscope.

I examined two emitters, one with oxide all the way up the nozzle, and another that had all the oxide missing. When the liquid was allowed to come into contact with the chips the sequence of events was the same for both. In figure 7.1 one of the emitters is shown. Liquid formamide has come out of the nozzle exit and in (a) has crept over the nozzle exterior. About 1 minute later the liquid drop has grown to the size shown in (b) and is just about to make a connection between the nozzle and the electrode. The emitter shown is fully oxidised (as can be seen by the purple colour) but the same pattern occurred on the surface of the non-oxidised chip, so the leakage was probably not because of the surface composition.

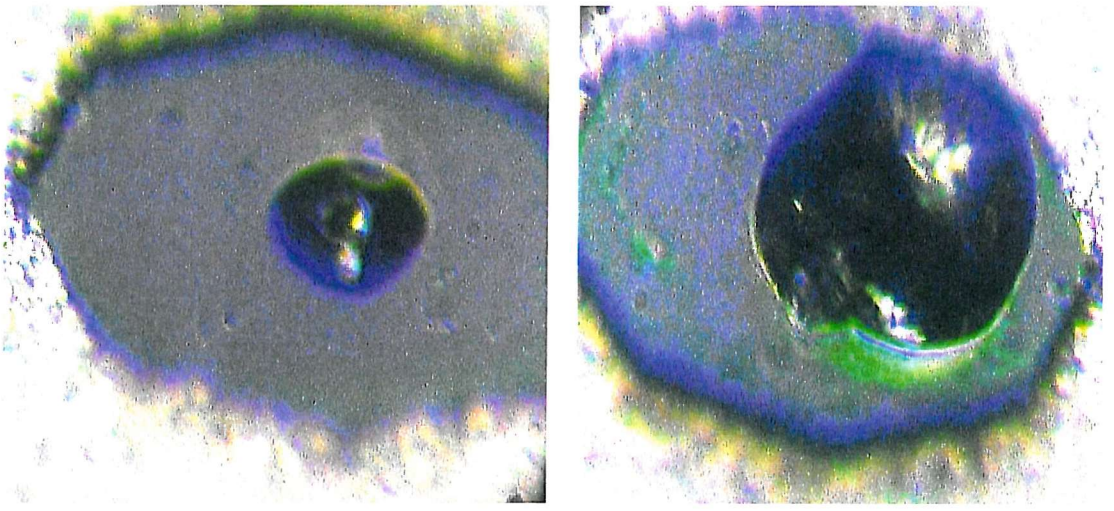

FIGURE 7.1: leakage of fuel; (a) ~ 20s after liquid contact, (b) ~ 1min later.

In chapter four an equation is derived which details the predicted flow from a capillary under a pressure head and a voltage. This equation needs altering however if it is to be of use in predicting leakage. The whole system needs to be considered; one must take account of all the pressure resistances encountered by the fluid flow on the way to the emitters, this includes hydrostatic pressure from height differences, and pipe resistance:

$$\delta P = P_{Hydro} + P_{Pipe} + P_{Surface} + P_{Emitters}$$

$$\delta P = -\rho g \Delta H + \frac{128 \mu L_p Q}{\pi d_p^{\,4}} + \frac{2\gamma \cos\alpha}{N_E r_E} + \frac{128 \mu L_E Q}{\pi N_E d_E^{\,4}} = P_{Liquid}$$

so, rearranging for Q gives:

$$Q = \frac{\left(P_{Liquid} + \rho g \Delta H - \frac{2\gamma \cos\alpha}{r_E} \right)}{\frac{128\mu}{\pi}\left(\frac{L_P}{d_P^4} + \frac{L_E}{N_E d_E^4} \right)}$$

where the subscripts E and P stand for emitters and pipeline respectively. N_E is the number of emitters on the chip. This Q is the total flowrate through the system, to get the flowrate through each emitter divide Q by N_E.

Our set-up uses 1/16" ID pipes that are ~ 60cm long with a ΔH ~ 30cm. Emitters with 1.5μm and 2.5μm radius exit holes were analysed and the results plotted on figure 7.2. The flowrates shown are for each emitter.

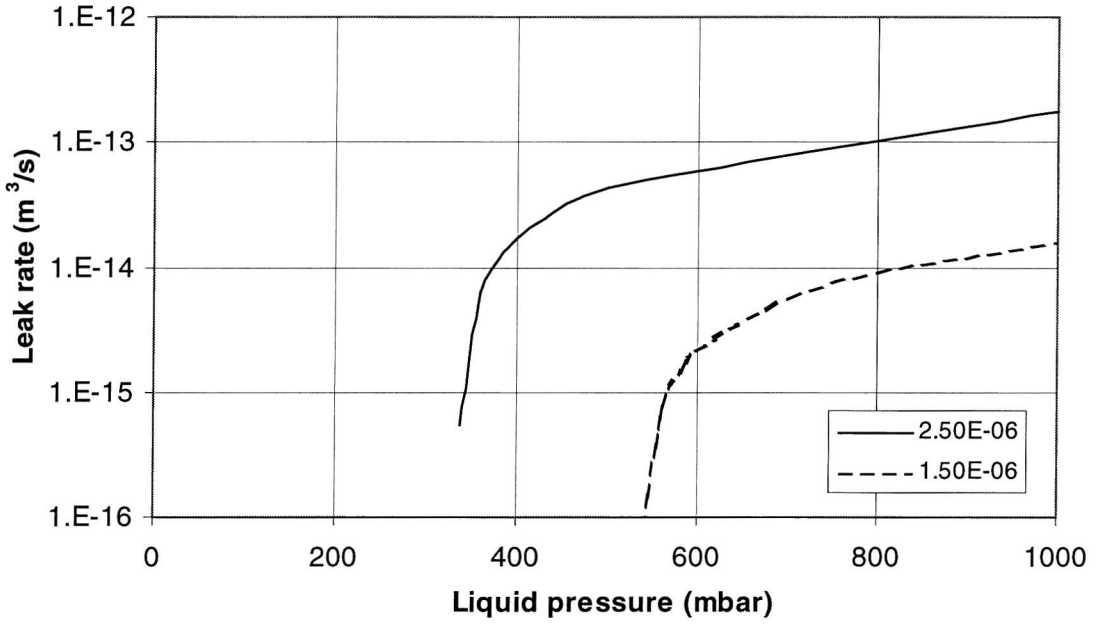

FIGURE 7.2: Plot of leakrate vs. pressure of the liquid

The pressure needed to start flow, or rather leakage, in each case is,

$$P_{Liquid} = \frac{2\gamma \cos\alpha}{r_E} - \rho g \Delta H$$

and equals ~ 474mbar for the 1.5μm hole and ~270mbar for the 2.5μm radius holes.

It can be seen that liquid at atmospheric pressure results in large leak rates. Given that a hemispherical drop of radius 50μm (i.e. wide enough to reach the electrode) only has a volume ~ $2.6*10^{-13}m^3$ the flowrate from the larger holes at atmospheric pressure would short the electrode only seconds after liquid contacts the chip. Every time the first set of chips was exposed to liquid it was at atmospheric pressure so large leaks occurred.

Having realised the above, future experiments were done with the liquid at a controlled, lower than atmospheric, pressure. When the chips were exposed to the liquid at less than the leak-pressure (but high enough to overcome hydrostatic pressure) the liquid held at the end of the nozzle exits. This can just be seen in figure 7.3, the light reflecting off the centre reveals the presence of a small liquid drop that only appeared after the liquid contact, and which remained there without growth.

FIGURE 7.3: A liquid drop held at the surface of a nozzle.

This is an important result as it shows that the nozzles contain the liquid effectively. Groups pursuing colloid microthruster research at MIT and QMUL have expressed concern over wetting of nozzles (even with very high aspect ratio nozzles). There is a consideration of surface composition and wetting angles, but the major difference, and the means by which leakage is avoided here, is our decision to use electrically dominated flow. The removal of a large, constantly driving pressure head, as required by hydraulically dominated flow, removes the main cause of leakage.

To avoid inadvertent exposure of the chips to atmospheric formamide, all parts of the fuel feed on the vacuum side must be cleaned thoroughly with water and acetone before any new chips are connected up. This is because the liquid remains at the controlled pressure until the feed system is disconnected, but rises to atmospheric as soon as the seal is broken to swap chips.

7.2 Electrical breakdown tests

After a number of spurious results caused by improper set-up of the experiment a batch of results were obtained for the breakdown voltages (V_b) of the chips from this first fabrication batch. The chips were selected from wafers 3 and 4; a good boron doped wafer and the fully processed SOI wafer respectively. The boron wafer was one of those on which the insulating oxide was accidentally removed with the results illustrated in figures 5.9 and 5.10. The SOI chips still had all the oxide remaining. All the chips tested were the 12μN precision control types, as these offer a large number of gates to test on each chip, so reducing preparation and experimental uncertainty.

The silicon of each chip was connected to ground and the negatively biased potential Vg connected to each of the gates being tested in turn. The results are tabulated in table 7.1.

Chip Type	Separation (μm)	Min V_B (V)	Max V_B (V)	Average V_B (V)
SOI	50	406	508	473
SOI	100	490	491	491
SOI	50	490	510	500
SOI	100	570	600	580
Boron	2	380	450	420
Boron	2	365	530	451
Boron	2	365	460	413

TABLE 7.1: Breakdown voltage of chips tested

The type of breakdown is difficult to determine, whether it is surface tracking or bulk material breakdown, the result is an irreversible destruction. Before destruction the resistance between the silicon and the electrode is greater than 20MΩ (the limit of our meter), afterwards the resistance drops to a few KΩ and in some cases only tens of Ω. Clearly then the breakdown creates a very conducting path through or along the surface of much of the insulator. Since both surface arcing and bulk breakdown can produce such a situation further testing is needed to determine the nature of the breakdown.

These results illustrated the random behaviour of electrical breakdown and this, together with the low number of results obtained makes these results too statistically weak to draw many

conclusions from. However, one conclusion was clear – the insulation performs very poorly in these test conditions, the breakdown values are well below the theoretical figures from literature. The oxides should be able to withstand 2kV; the factor of four reduction obtained here was puzzling.

One might have expected the SOI chips to withstand larger voltages because they have all the oxide intact and therefore have the 50 or 100μm long insulator surface as extra flashover protection. The fact that the voltages are not appreciably higher does not however mean that that flashover length is inconsequential. The picture is complicated by the fact that the insulation on the SOI wafers, although undamaged, is PECVD SiO_2. Oxides deposited in this way typically have bulk breakdown strengths 50% lower than thermal oxides. This means that a direct comparison cannot be made between the SOI chips and the boron chips.

Clearly, this presents a very serious design problem – if the insulation can only hold 450V before permanent destruction there is no chance of obtaining a spray from the chips. I attached paramount importance to finding the reason for this departure from theory and seeking a resolution.

7.3 Full characterisation of insulation

It is difficult to draw any major, reliable conclusions from the insulation on the first batch of chips, because of the uncertainty caused by the oxide removal. To solve this a new fabrication batch was produced in the Microelectronics group. The process steps are shown in idealised cross section form in figure 7.4. The batch is an incomplete version of the original batch using some of the same masks, in which insulation and the electrodes are fabricated but none of the nozzle steps are performed. Since the nozzle steps are by far the most difficult and time consuming this batch only took a couple of months.

Since the results in section 7.2 suggested that even with extended oxides the insulation might be insufficient, a possible solution was tested. Once the thermal oxide has been grown to the maximum thickness of 2μm the thickness was increased by depositing a further 2μm of oxide by CVD. Only half of the wafers in the batch have this new 4μm thick insulation, the other half have the old 2μm oxide for comparison purposes.

A few modifications were made to the process steps to make successful fabrication much more likely. The final wet etch, which was responsible for the destruction of the first batch oxides, was replaced by a ion etch, as it could not be guaranteed that even the alternative resist will stay firmly on the aluminium. The ion etch is much safer but does require a thicker 2.2μm resist and very long (~3hour) etch times to get through the 4μm. To limit excessive usage of the ion etcher then, the ion etch, which defines the contact window for the silicon bondpad, was replaced by the buffered HF wet etch. This process was successful and so the problem of accidental oxide removal is resolved.

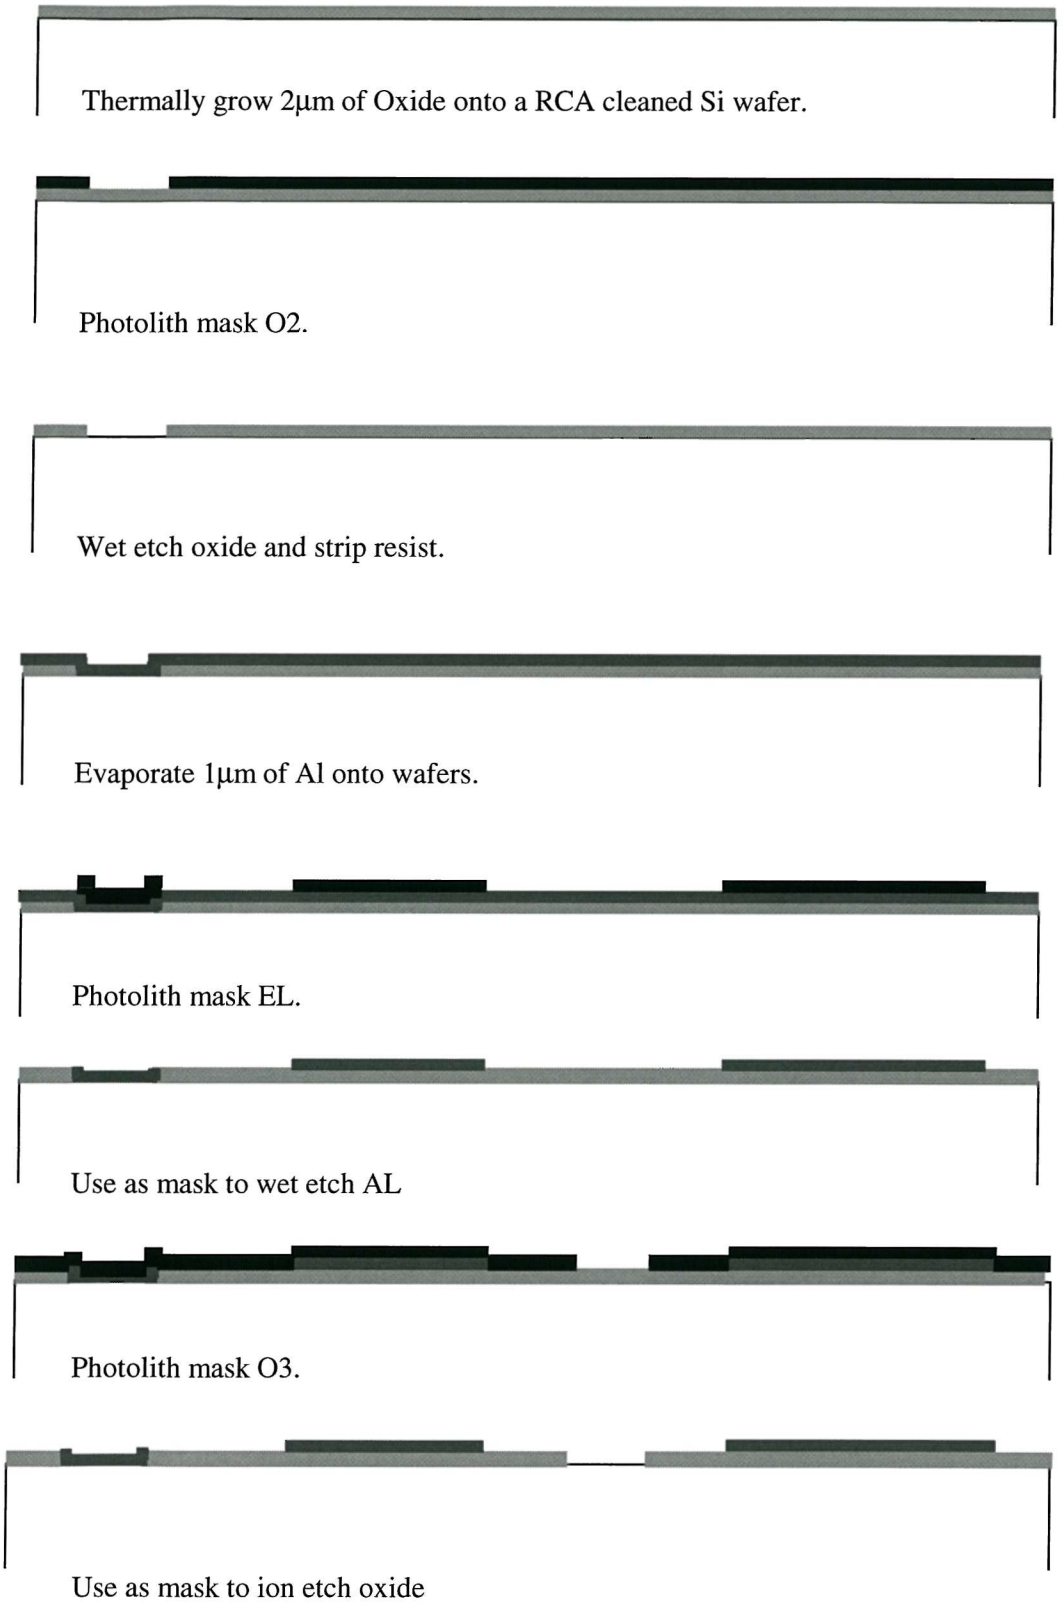

FIGURE 7.4: Fabrication process for insulation - test batch

Once this batch was completed, the chips from all the wafers were packaged. Since no fuel is used in these tests the packaging is much simpler. The emphasis here was on maintaining maximum cleanliness of the chips. The importance of this can be seen by looking at a chip in which no particular care was taken other than to treat the device as fragile. The first photo in figure 7.5 shows a gate on such a chip. Numerous particles can be seen with a wide range of sizes. These could cause disastrous damage during high voltage testing and on some tested gates crater-type damage has been seen.

To combat this the chips were subjected to a stringent cleaning regime. After sawing, the chips spent 10 mins in fuming nitric acid, which is a very harsh cleaning agent. The chips (and the chip mounts) were then bonded to microscope slides using a minimal amount of epoxy resin before undergoing 10mins in an ultrasonic acetone bath. After the chemical cleaning steps, the chips were washed with a water jet before being force dried by a N_2 jet to avoid the residues that form on chip surfaces when solvents are allowed to dry naturally. After wire bonding, they were transported in cleanroom-standard boxes and only handled for the minimum amount of time in the lab, and always with clean gloves.

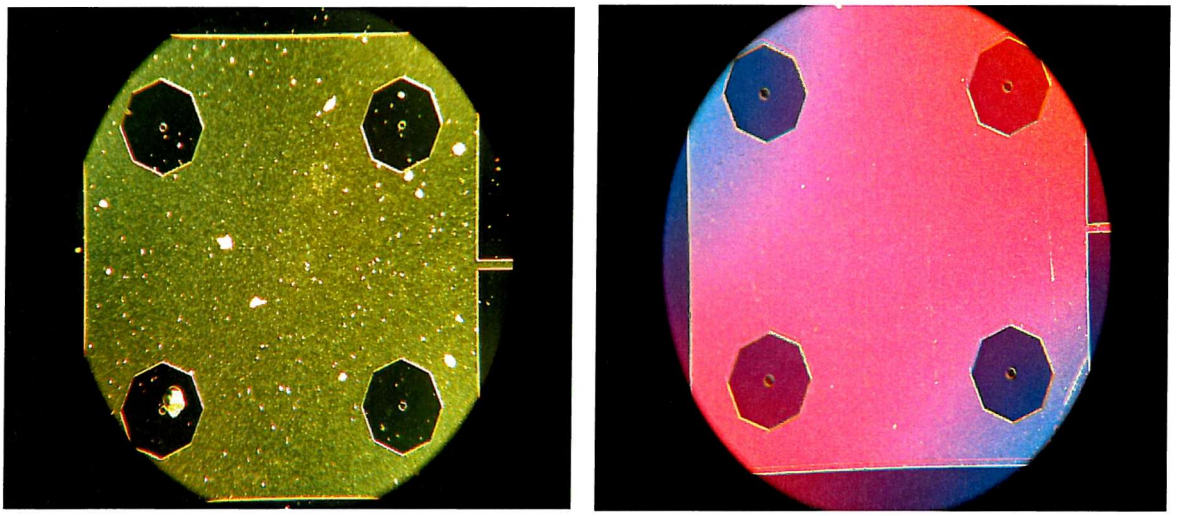

FIGURE 7.5: The contamination on gates before and after cleaning.

The result is a dramatic improvement in the cleanliness of the chip surface and the elimination of staining by residues.

7.4 Approach

With the availability of so many chips for testing a test plan was devised to allow us to elucidate the breakdown mechanisms and influential factors. The various influences were prioritised according to how helpful they might be in determining why the insulation performed so poorly in the preliminary tests. This is shown in table 7.2.

Parameter	Range/ values	Reason / question asked	No. of results	Importance
Electrode separation	50μm and 100μm	Does separation affect V_B – surface tracking?	Very high	Crucial
Oxide thickness	2μm and 4μm	Bulk breakdown?	Very high	Crucial
Ambient pressure	1E-7 to 1E-2 mbar	Should affect surface breakdown & is like real environment	High	High
Voltage polarity	+ve / -ve	Any cathode effect?	Medium	High
Electrode size	Small to large gates on 12μN chips	Is it detrimental to have large grouped arrays?	Medium	Medium

TABLE 7.2: Parameters to be investigated.

The most important question I wished to find an answer to was what is the mechanism of breakdown. Armed with that information one can start to offer solutions/ improvements. As detailed in section 3.5 there are three distinct breakdown mechanisms that are relevant but the discussion can safely be restricted to surface breakdown and bulk breakdown only. This is because surface breakdown will always occur first in a vacuum if the gap is bridged by an insulator surface. A survey of the literature on breakdown uncovers the predominant factors influencing surface breakdown and it is these that the experiments focussed on, i.e. electrode separation, polarity, ambient pressure and surface contamination. Together with the results of similar tests performed on chips with extra thick insulation it was felt this approach would offer up a clear answer to our question.

As with any investigation it is important to separate out the effects of one parameter from the others, this requires that all other factors are kept constant while varying the parameter of interest. To that end a baseline testing condition was established and is the following:

- +ve polarity voltage applied to the aluminium.
- Pressure < 10^{-6} mbar
- Smallest electrode size
- Utmost cleanliness

In the literature numerous mention is made of the influence of the testing circuit. A resistor of 1MΩ or above limits the available current and so the power available for any discharge. Since the benefit of this seems obvious, I implemented this simple step into the testing. Limiting current to relatively low levels must be beneficial – the potential damage of a 1A discharge must be much larger than one limited to 1μA. Transient discharges that are capped in this way are much less likely to cause destruction.

The testing circuit is shown in figure 7.6 and is so designed to allow us to apply voltages of between 300 and 3500V from a variable high voltage supply with an output between 750 and 3500V. Lower voltages can be obtained by altering the potential divider. The high resistances limit the available current to ~5μA/kV and protect the picoammeter from overvoltage damage.

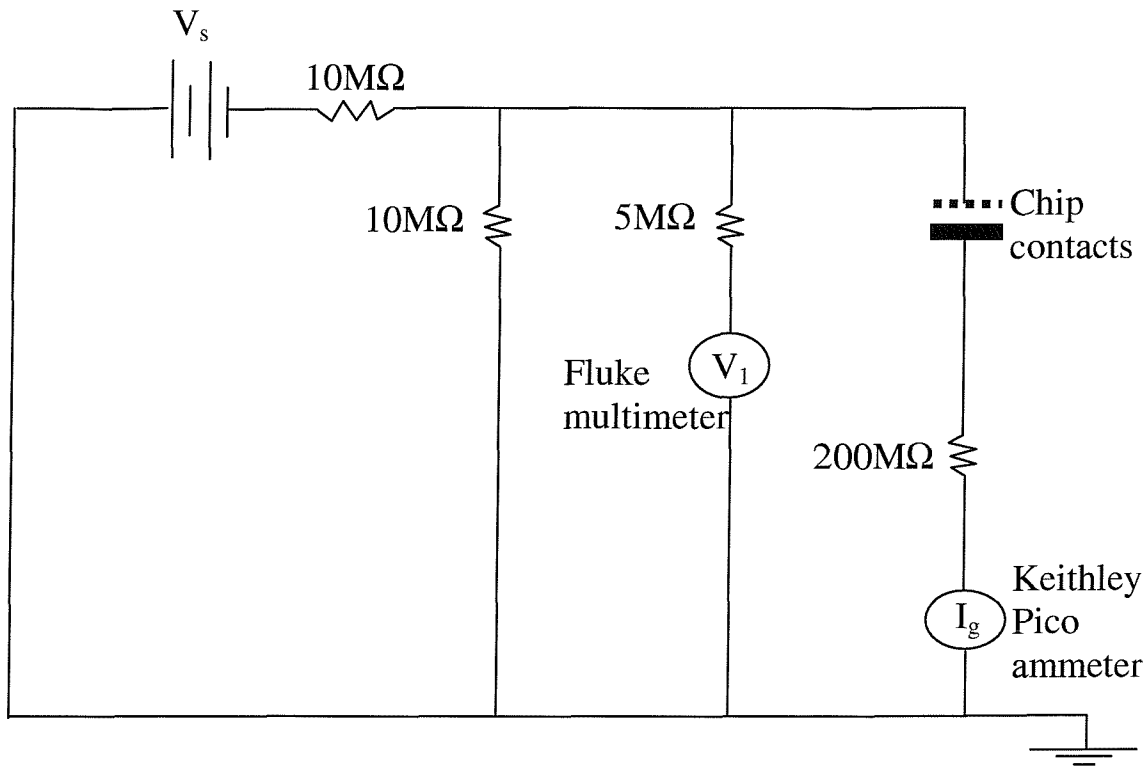

FIGURE 7.6: Main test circuit layout

The method of testing was quite simple and kept consistent throughout. Voltage was applied to the gate(s) of the chip under test when the desired pressure was reached. The voltage was increased in steps at a prescribed rate to give an overall voltage ramprate. The output of the picoammeter was recorded via an A-D converter onto a PC. To aid analysis all voltages, the gate current, chamber pressure and time were recorded in logbooks along with observations. A microscope picture was taken of each gate tested to allow us to compare the condition of the gates before and after test.

Chapter 8

Experimental results and discussion

8.1 Main results

A total of 245 gates were tested, the majority of which were the $1mm^2$ gates on the "12μN" chips. A very large amount of current-voltage data was collected but the most important points can be gleaned from more general overviews which consider only the maximum voltage held before permanent, irreversible breakdown – so called hard breakdown. The following few pages present graphs collating the voltage at hard breakdown (V_B) for the investigated parameters. In this section the main observations will be presented, followed by their interpretation and the conclusions they give rise to.

8.1.1 Electrode separation

To determine the effect of the distance between the aluminium electrode and the exposed silicon (hereafter referred to as the "gap") a large number of gates were tested. The graph shown in figure 8.1 collates the results of 61 gates with 50μm gaps and 55 gates with 100μm gaps.

The graph illustrates the seemingly random nature of breakdown test results and shows the need for large numbers of results. However, with this many results a very clear picture can be seen and acceptable statistics can be obtained - the statistical properties of the results are discussed in section 8.1.6. The mean values for breakdown voltage (V_B) are 1325V and 1330V with standard deviations of 214V and 148V for 50μm and 100μm gaps respectively. Evidently, the gap size has no appreciable effect on the breakdown voltages within this range.

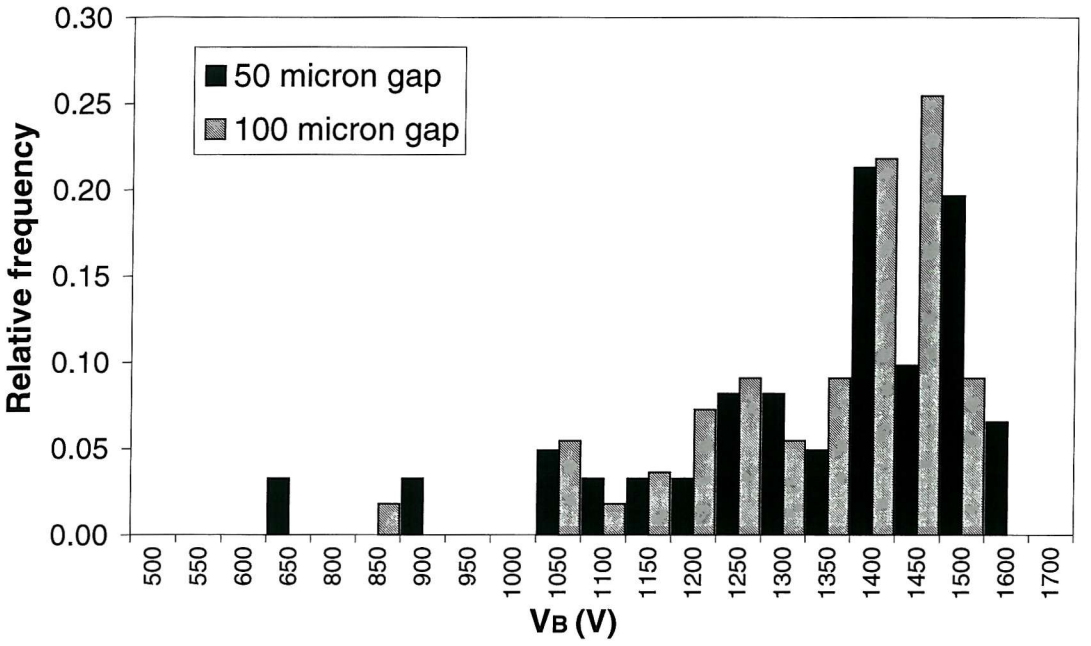

FIGURE 8.1: Breakdown voltage vs. relative frequency for different gaps

8.1.2 Ambient pressure

Tests were done with a wide range of vacuum chamber pressures. This was achieved using a precision leak valve in place of the liquid pressure cell; the valve was adjusted until the required chamber pressure was obtained and was held constant over the course of the test. The graph in figure 8.2 collates the results.

Each data point is the arithmetic mean of all values obtained at that pressure and the standard deviation for each point is between 30 and 60V. No significant effect is seen even with a pressure variation of four orders of magnitude. Tests could not be done at pressures higher than 10^{-2} mbar because glow discharges would then occur with voltages > 1kV for cm sized separations. It was not felt useful to obtain results for such high pressures anyway.

The majority of the data points were obtained using untreated laboratory air as the ambient gas. In an operational environment however the ambient gas would most likely be the vapour of the chosen fuel. This research has been based on that fuel being formamide. To simulate such a situation, I introduced a vial of formamide liquid to the chamber and upon pumping obtained equilibrium between the pump flow and the vapour evaporation rate. As a result, the chamber was filled with formamide vapour at a steady pressure of 10^{-3}mbar. The resulting data point shown is a mean of 1425V with a standard deviation of 227V. Again, no effect is

seen. Although only one pressure was tested, lower pressures are unlikely to encourage breakdown at lower voltages than measured here.

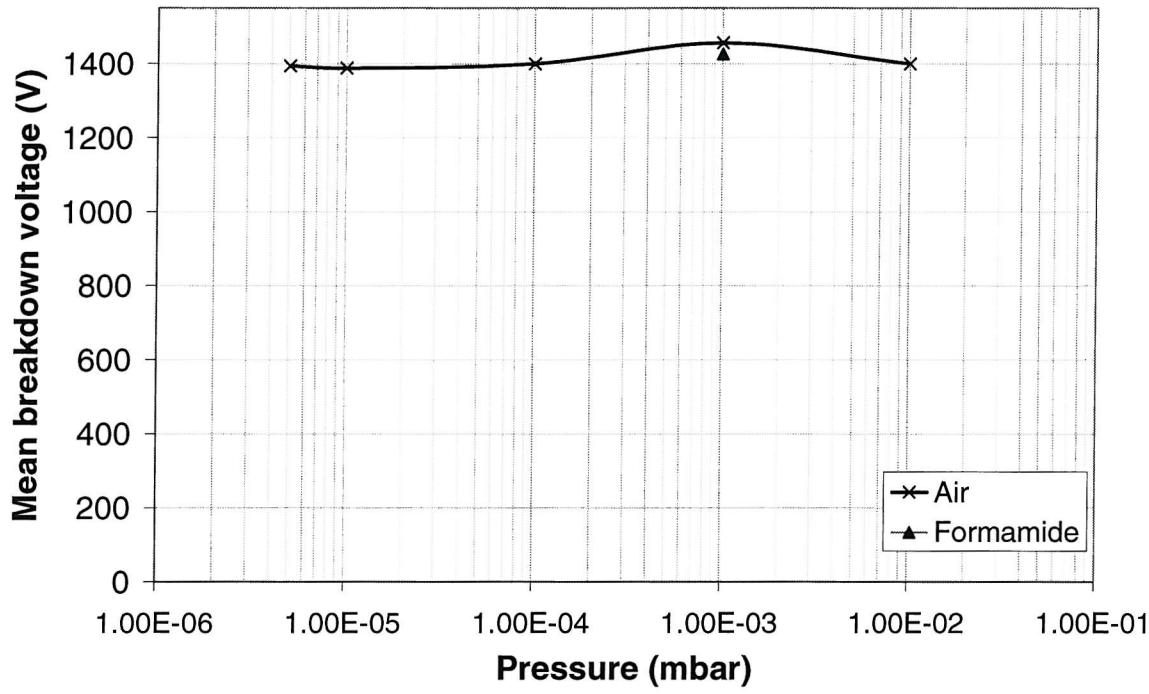

FIGURE 8.2: Effect of ambient pressure on breakdown voltage.

8.1.3 Insulating oxide thickness

In this set of tests chips from the wafers with 4μm of oxide were investigated under the same conditions as the 2μm oxide chips. The results are shown in figure 8.3. As one might expect there is a substantial increase in the breakdown voltage obtained. These results show a wider scatter – maybe this is due to the larger range of possible values. From a total of 79 results, the mean breakdown voltage was 2960V with a standard deviation of 494V. Notice the long tail at the low end and sharp cut-off at the higher voltages similar to that described by Osburn [66] for bulk breakdown.

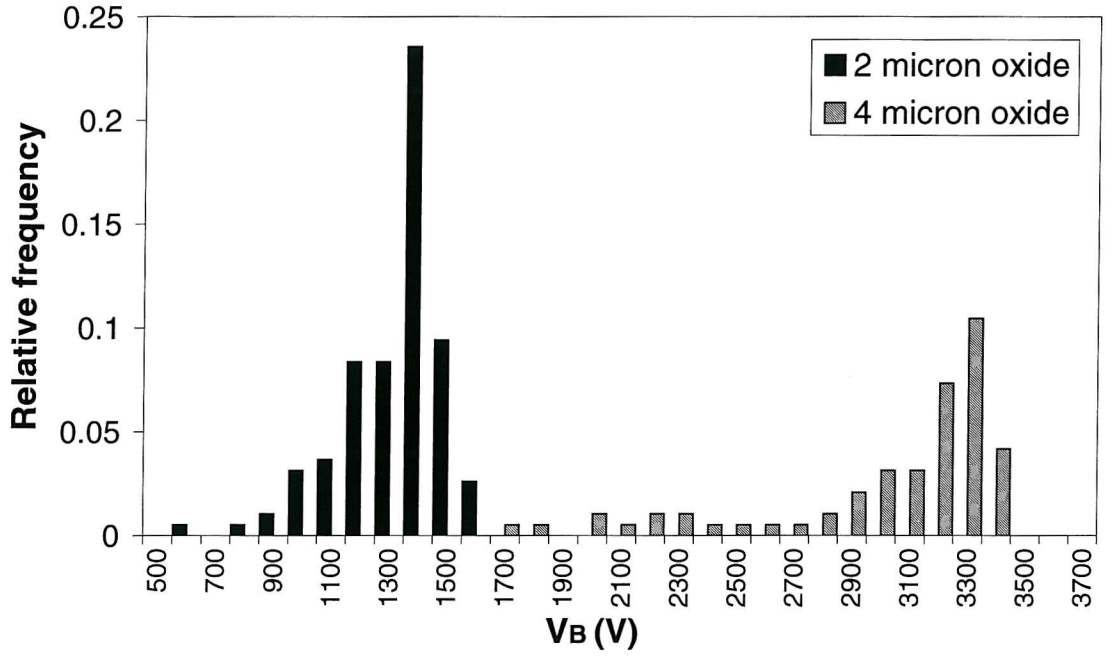

FIGURE 8.3: Effect of insulator thickness on the breakdown voltage.

8.1.4 Voltage polarity

All of the results shown so far have used positive voltage applied to the aluminium. A series of tests was performed however in which negative voltage was used instead. During the tests, very large pre-breakdown currents were observed at relatively low voltages but these are discussed in section 8.3; for the moment hwoever, the focus remains on the voltage at hard breakdown. In figure 8.4 a comparison is made between the results obtained for positively and negatively biased gate voltage, both on 2μm oxide chips.

There is an appreciable decrease in the attainable voltage when the gate is negatively biased. Taken from 42 results, the mean is 1079V with a standard deviation of 144V for negative polarity tests. This should be compared with the mean of 1342V and standard deviation of 180V of the positive polarity tests – a clear difference of almost 300V.

8.1.5 Electrode size

Some investigations were performed to determine the effect of the electrode area on the breakdown voltage. Since the "12μN" chips also have the bigger electrodes (area ~9mm^2) an attempt was made to determine their breakdown voltages. Although the data is limited, it is

still illuminating. The vast majority of these very large electrodes were either already shorted out or broke down immediately on application of the minimum 300V. Those that did permit testing with the usual method gave a mean breakdown voltage of 712V with a 328V standard deviation. However, these values are only illustrative because of the difficulty of deciding whether already shorted or immediately shorting gates should be included in the analysis. What is clear however is that with an increase in area of only one order of magnitude the gates become considerably less reliable and in many cases completely unusable.

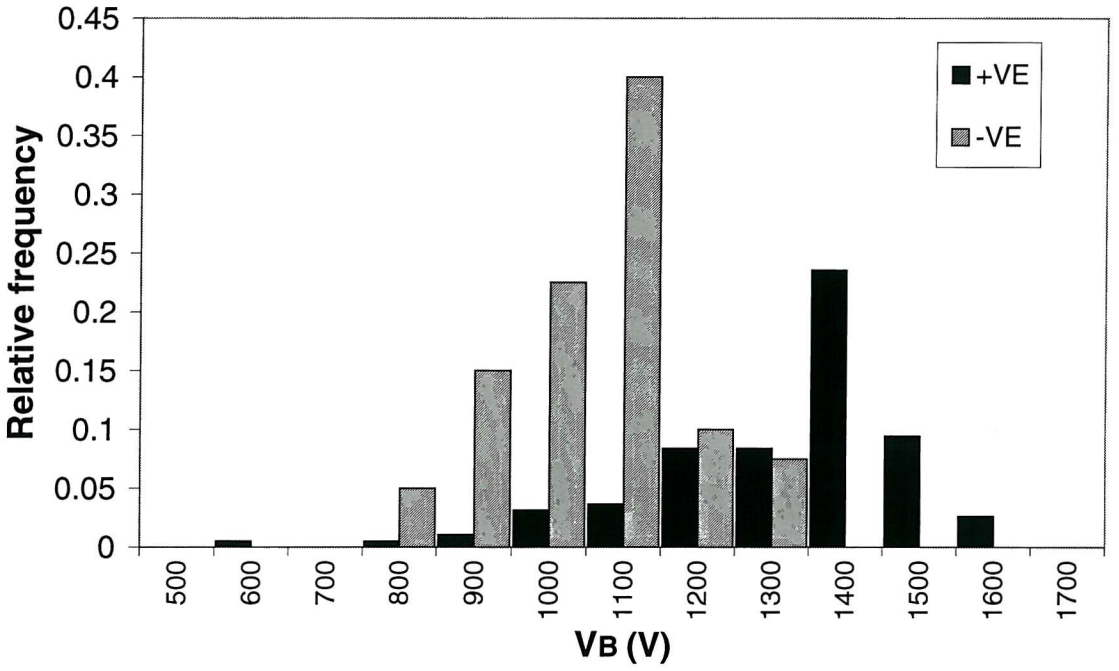

FIGURE 8.4: Effect of the gate voltage polarity.

8.1.6 Statistical properties of results

To gain an understanding of the results an exploratory data analysis was performed . As a first step the descriptive statistics were calculated from the raw data – this provides the location and scale parameters and the shape of the data spread. These are shown in table 8.1.

	50μm	**100μm**	**All 2μm**	**All 4μm**	**Negative**
Mean	1330.2	1325.5	1341.2	2959.7	1073.8
Median	1400	1400	1400	3200	1100
Standard Dev.	214.2	147.8	179.6	493.5	143.7
Skew	-1.3	-1.1	-1.5	-1.5	-0.2
Kurtosis	2.0	1.0	3.0	1.0	0.5

TABLE 8.1: Descriptive statistics of the results

The first thing that stands out is the large amount of skew. The figures given use the formula for Fisher skewness, for which a value of zero represents a normal probability distribution. Negative values indicate the presence of a tail to the left of the mean, i.e. a propensity for small numbers of results much lower than the mean value. All the results are significantly negatively skewed apart from the results for the negative aluminium.

Kurtosis is a measure of the flatness or the size of the tails present in a distribution, the Fisher kurtosis is used here, for which zero indicates a normal distribution. A distribution with long tails will have a positive value. All the distributions have large positive values of kurtosis because of the tails in their data, except that is, for the negative case.

Both of these points can be seen by examining the histograms shown in the previous figures. All of the positive results have large peaks and are heavily skewed to the left of the mean with a long flat tail – the symmetry is very poor. By contrast the negative results have values of kurtosis and skew that are close to zero – the histogram shows a lack of tails and is heavily symmetric. This provides starting points from which to identify the types of distributions the results belong to.

A quantile-quantile plot is commonly used to determine if two populations come from the same distribution. A normal distribution Q-Q plot scatter-plots the quantiles of the experimental data against the quantiles for a standard normal distribution. The fit can be seen from how tightly the data points fit a straight line – curving away from the line indicates skewness while points which are isolated and well away from the line are probably outliers. Figure 8.5 shows a normal Q-Q plot for the 50μm gap and another for the negative results. The plot for the 50μm results has a clear curve away from the line because of the skew and outlying points can be seen in the bottom left hand corner. The negative results are clustered fairly evenly around the normal line with no curvature and very few outliers. This suggests that, while a normal probability distribution is not appropriate for the positive polarity results, it may be for the negative results.

This is verified by a plot of the normal distribution superimposed on data from the negative polarity results as shown in figure 8.6. The plot uses the z scores for which zero is the mean and a value of one is a standard deviation. The fit is good apart from the smaller peak at the mean, this is reflected in the correlation, which is calculated as 92%.

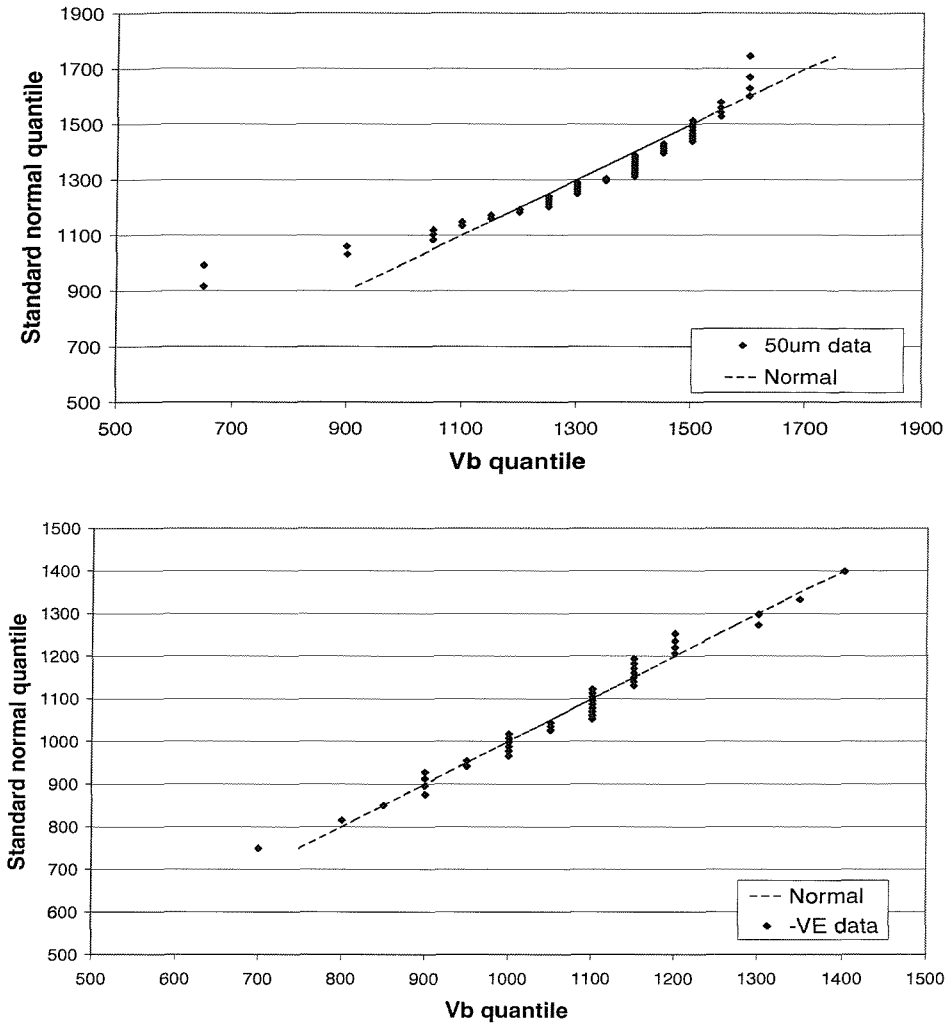

FIGURE 8.5: Normal Q-Q plots for the 50μm (top) and negative (bottom) cases

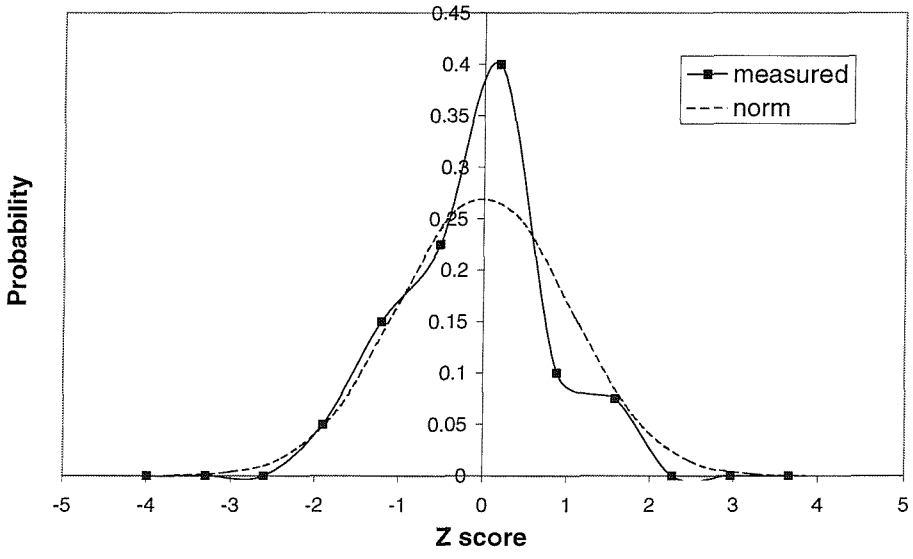

FIGURE 8.6: Comparison of negative results with normal distribution

The positive results clearly do not fit a normal distribution so other types were tested. The most likely distributions for results with such long tails and large peaks are the Weibull and lognormal distributions. Both of these are very flexible due to their use of a shape factor that alters the curve shape obtained. For example, depending on the shape factor used, the Weibull can model exponential or extreme value distributions. The best fits for the data were found using the lognormal distribution. The Weibull function gives a poor fit to the data because there is a very sharp cut off at the high end of the data sets. The two parameter lognormal distribution was calculated using the following equation:

$$f(Y)=\frac{e^{-\left(\frac{(\ln(Y))}{\nu}\Big/2\phi^2\right)^2}}{Y\phi\sqrt{2\pi}} \tag{8.1}$$

where Y is a transformed parameter that changes the right-skewed distribution into a left-skewed one:

$$Y=\frac{\left(V_{B,max}-V_B\right)}{\psi} \tag{8.2}$$

Distributions of this type give very good correlations, and are shown superimposed on the data in figures 8.7 to 8.9, the correlations are 95.7%, 95.3% and 97% for the 50μm gap, 100μm gap and the 4μm oxide cases respectively.

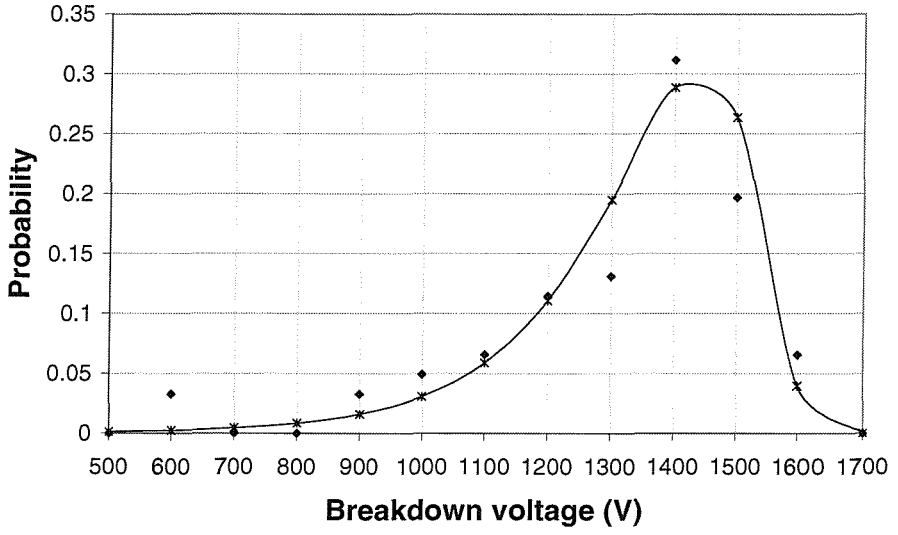

FIGURE 8.7: Lognormal distribution for 50μm data

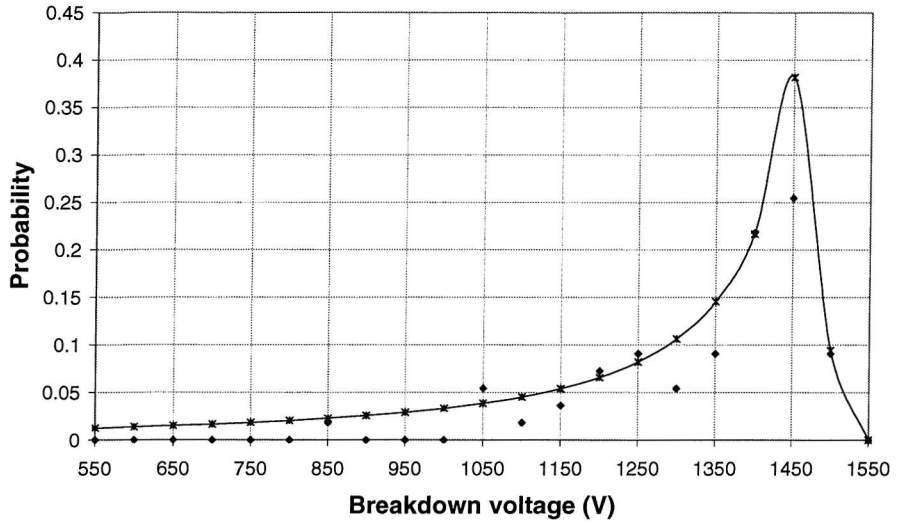

FIGURE 8.8: Lognormal distribution for 100μm data

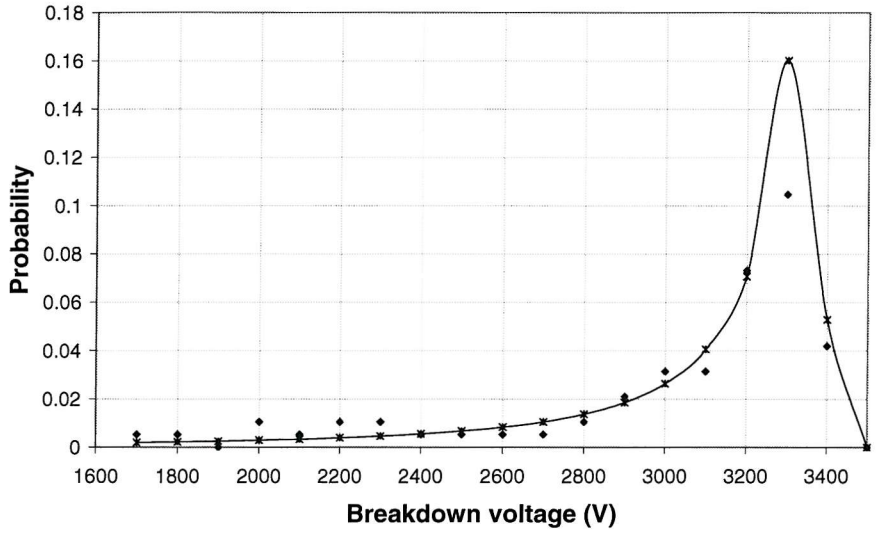

FIGURE 8.9: Lognormal distribution for 4μm oxide data

The lognormal fits the general shape of the data quite well but some points are less well correlated than others, this is possibly because of the limited sample size. This means that it is not possible to state which distribution is the definite characteristic of the experiments; only that the lognormal distribution offers a suitable empirical model.

8.2 Determination of Breakdown Mechanism

It is crucial to determine the mechanism of breakdown as it has serious implications as to whether the design is at all suitable and would operate properly in its intended environment.

The results shown above allow us to say, with a fair degree of certainty, which is the dominant mechanism causing permanent destructive breakdown.

8.2.1 Surface breakdown

In the discussion on surface breakdown in chapter 2, the most influential factors are identified as surface contamination, ambient pressure, insulator material, and surface length. Throughout the tests, attempts have been made to minimise surface contamination as much as possible - using the cleaning procedure described, and, as with insulator material, it is fair to assume that this is a constant. The remaining influences are pressure and insulator length. Results for both of these have been presented.

NO PRESSURE EFFECT

The equation for surface flashover (equation 2.1) includes a critical gas density produced by electron stimulated desorption – this is when the test specimen is in an ultra-high vacuum. It seems reasonable to assume however that high pressures of ambient gas would make breakdown more likely. Equation 2.1 shows the dependence on the desorption probability, which by equation 2.4 is seen to be proportional to the number of gas molecules available for desorption from the surface. At high vacuums the amount of gas adsorbed on a surface can be quite small, as the ambient pressure rises however, this increases considerably. The coverage of a surface with adsorbed gas depends strongly on the pressure of the ambient gas – at high pressures the adsorption is greater than the always occurring thermal desorption and the surface coverage ratio approaches unity [56]. The amount of gas available for desorption is therefore higher when the chamber is held at poor vacuum pressures. Accepting this one might expect that if breakdown was being caused by surface flashover then it would occur at lower voltages when the ambient gas pressure is increased. The complete lack of a pressure effect over such a large range of pressures suggests that the destructive breakdown does not occur either through the "vacuum" or across the insulator surface.

NO GAP DEPENDENCE

Equation 2.1 also predicts a square root dependence of flashover voltage on insulator length. Going from a gap length of 50μm to100μm should cause the flashover voltage to increase by

a factor of √(2), i.e. a 41% increase. If this were the case it would be distinctly observable – if 1400V was measured for 50μm gaps then the 100μm gaps would breakdown at 1980V. There was no statistically significant difference between the mean values obtained for the 50μm and the 100μm long gaps. As with the pressure results this suggests it is unlikely that the gates are being destroyed by surface breakdown. However, the lack of dependence of breakdown voltage on flashover length has only been shown for the range of 50-100μm.

FLASHOVER PREDICTION

One can use equation 3.11 to see whether one would expect flashover to occur across an insulator at the maximum voltages obtained. The curves in figure 8.10 were plotted using the minimum value for the density of gas needed for breakdown, i.e. $M_{cr} = 1x10^{18}cm^{-2}$. and a drift velocity in the middle of the range, $v_o = 1.7x10^5 m/s$.

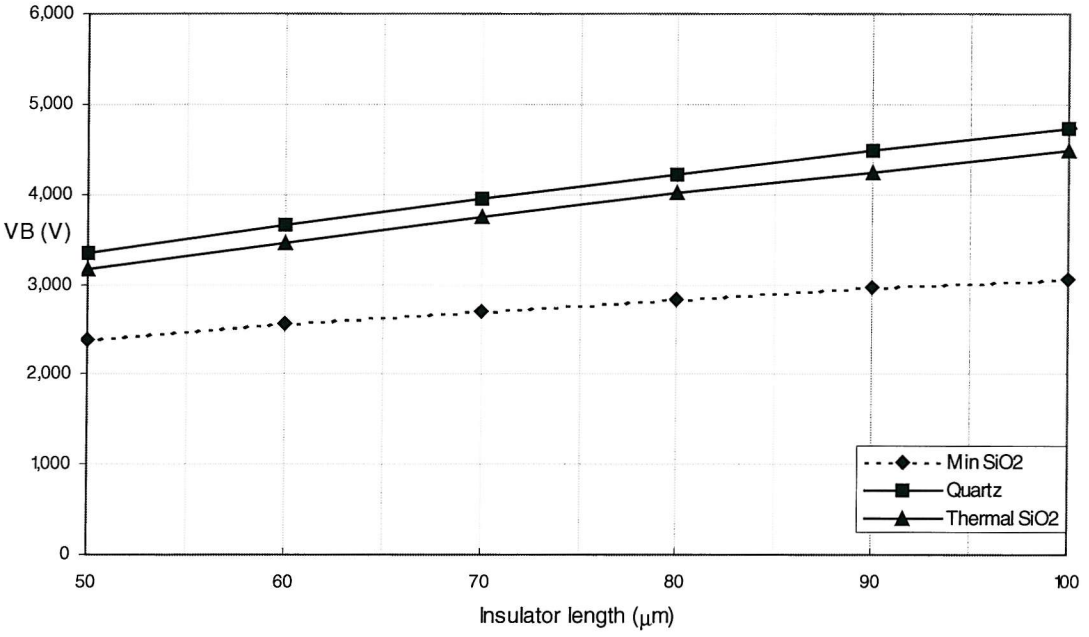

FIGURE 8.10: Voltage required for surface flashover across a gap

Pillai's curve for Quartz has been included (extrapolated to the dimensions of interest) for reference purposes and compares well with our predictions using data for thermal SiO_2, where $\gamma_d = 10$, $A_o = 2eV$ and $A_1 = 20eV$ [80]. It shows that - for a clean sample in a good vacuum - gaps of 50 to 100μm should be able to support 3.2 to 4.5kVon average. The minimum curve shown uses the maximum value of $\gamma_d = 15$ and is included to illustrate the damaging effect that increased desorption probability can have. However, even for that case the minimum voltage required to cause flashover is calculated as 2380V. These are much higher than the

breakdown voltages measured and even the minimum value is outside the 3-sigma range, which for the case of the 50μm gaps is ~ 1973V.

8.2.2 Bulk breakdown

THICKNESS EFFECT

The gates of a chip with 4μm of oxide suffer destructive breakdown at a voltage at least double that of a chip with 2μm oxide but otherwise identical. This is perhaps surprising, given the potential extra problems of an interface between the separate layers of oxide produced by thermal oxidation and CVD. This interface consists of two contacting surfaces with the defects and boundaries that are always present at material surfaces. Such a density of defects might be expected to trap significant amounts of charge; I have discussed how this can encourage breakdown. It is possible this contributes to the longer tail visible on the breakdown distribution for 4μm oxides. Regardless, a simple doubling of oxide thickness would have no effect on breakdown distributions caused by surface breakdown – thickness plays no part in those. Rather, the observation points very strongly to the cause of final destruction being a breakdown through the bulk of the insulating material.

ELECTRODE AREA

The experience from testing gates with much larger electrodes seems to corroborate the proposal that bulk breakdown is to blame. One of the major causes of bulk breakdown is the presence of defects in the insulator. Clearly, when the electrode area goes up by an order of magnitude the area of insulator and so the number of defects subjected to the maximum field increases accordingly.

The breakdown is widely accepted as being destruction that results when the device, which is a capacitor, discharges its stored electrostatic energy through the material. Since that energy depends on the value of capacitance, which is proportional to the electrode area, it is clear that the bigger gates have ~ 10 times the destructive energy available. Note though that the area the energy is discharged through does not increase correspondingly. Breakdown occurs through a defect in the material, only the population of defects increases with the material area, not their size. Both of these points mean that an increase in electrode area can only be

detrimental to the gate reliability, as it was seen to be, if the breakdown mechanism is through the bulk of the material.

FURTHER SUPPORT

Since bulk breakdown occurs internally one would expect factors such as ambient pressure to be irrelevant. Gap length too would have no bearing – reducing this increases the electrical field across the surface only, not through the thickness. In section 3.5.4 mention is made of the theoretical field required to cause intrinsic bulk breakdown by the ionisation – recombination model. If the same experimental constants are used the model (equation 2.11) predicts that field to be equal to ~ 6.55MV/cm for a 2μm thick SiO_2 layer. This prediction of the critical voltage as 1310V is extremely close to the measured mean breakdown voltage of 1330V.

Bulk breakdown would be consistent with the lognormal distributions that were found to be good models for the data. As described in section 2.3.4 this mechanism produces distributions with long tails in the low voltage range – this is due to the presence of defects in the insulator material. However the presence of defects could never produce a symmetrical distribution with tails on both sides, because the breakdown strength of the insulation tends towards the maximum breakdown strength of the defect-free material. This means that there is a definite cut-off at the high end of the voltage range above which no specimen can be expected to survive. The lognormal distribution is very characteristic of such a situation and this is why it has such success when dealing with failure distributions such as crack growth and semiconductor reliability. All the positive aluminium data has this long left-hand tail and sharp high voltage cut off characteristic that one would expect, and has been observed by other workers, for bulk breakdown.

OBSERVED DAMAGE

Upon close examination of the chip under an SEM the type of damage seen is similar to that described as single-hole bulk breakdown of SiO_2 [65]. All of the images in figure 8.11 were from the same chip and corresponded to gates which had the aluminium positively biased. All gates tested in this way had the same damage - characterised by a sudden, explosive discharge that forces a hole through the oxide and damages the silicon. Such damage seems

characteristic of breakdown through the bulk of the oxide since it is hard to imagine how an arc across the surface of the insulator would produce such deeply penetrating and sharply focussed damage.

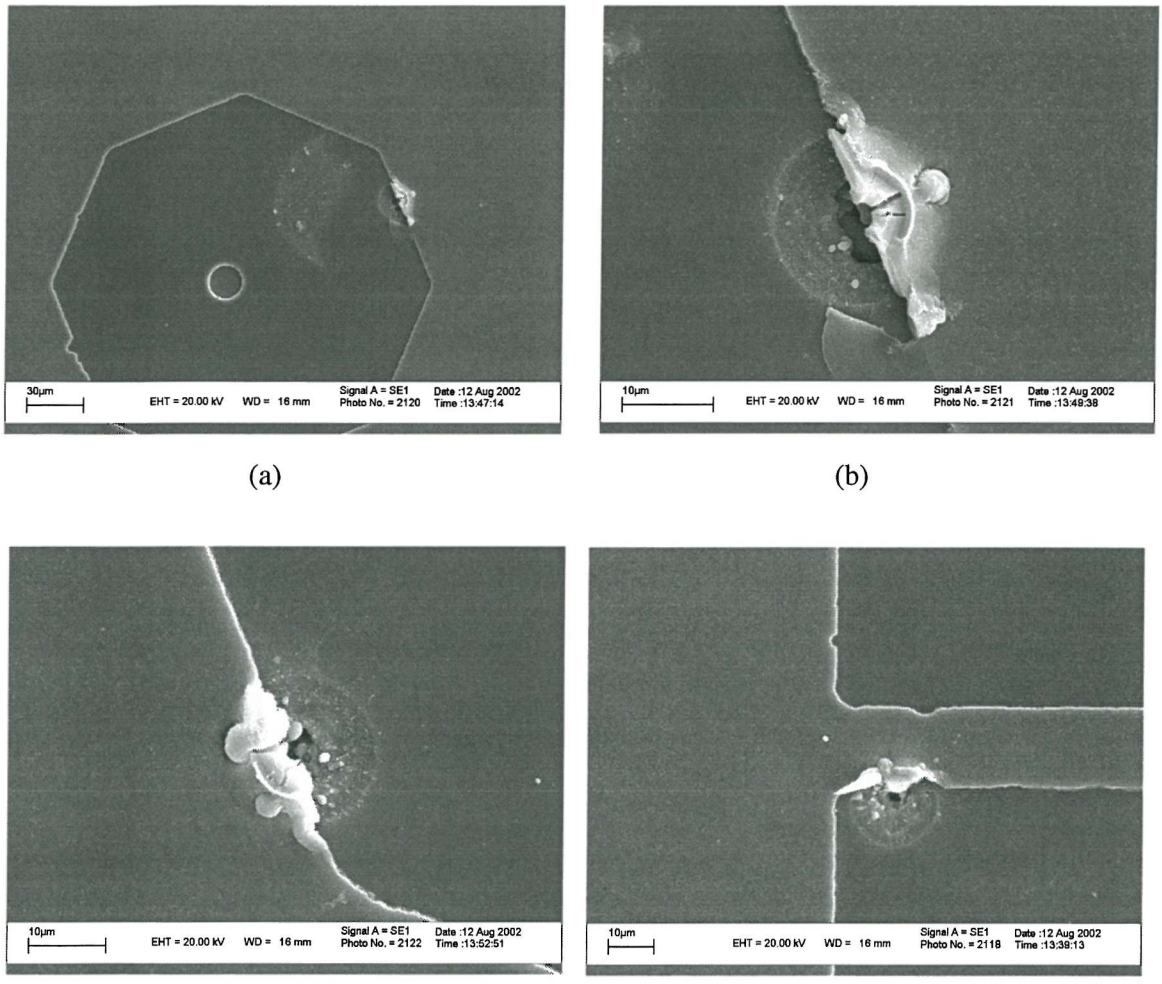

FIGURE 8.11: Breakdown damage for positive aluminium
(a) location of damage site 1, (b) Close-up of site 1
(c) Close-up of site 2, (d) Site 3 (on connection line)

The damage occurred at the edge of the aluminium, which, although it suffered some melting, was not fully evaporated away from the area. The location of the damage is not inconsistent with bulk breakdown since this is the triple junction mentioned in section 2.3.1. Although discussed in the light of surface breakdown this point is significant for bulk breakdown too. The electrical field is greater here than at any other point near the aluminium electrode, because of the presence of edges/corners. This large peak in the field would penetrate into the

adjoining insulator, causing the bulk breakdown to occur preferentially in the insulator beneath the edge of the electrode.

8.3 Negative polarity

8.3.1 I-V characteristics

Not only was the voltage at hard breakdown much lower when the gate voltage was negatively biased, but the I-V characteristics were considerably noisier. When the gate voltage was raised above a typical value of 600-700V there were large transitory peaks of 1μA or more. Figure 8.12 shows a sample plot from the data log for such an instance.

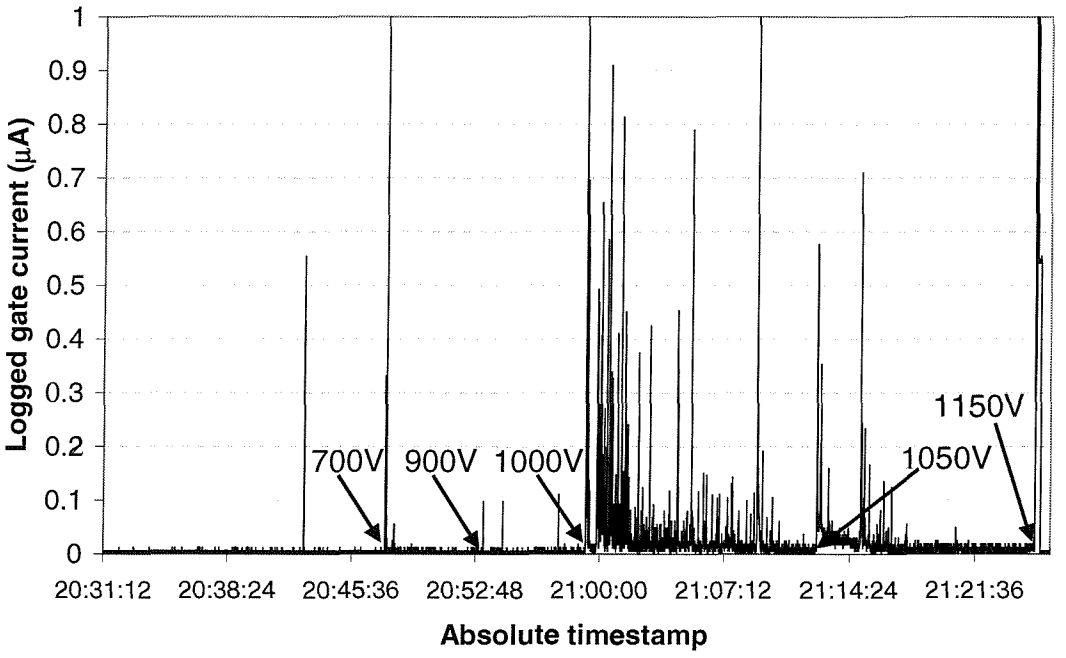

FIGURE 8.12: plot from the data log of a negative polarity test

Notice how the spikes occur with considerably greater frequency and amplitude immediately after a raise in the voltage; they die down somewhat after a period. It is clear that every voltage increase causes a new clutch of activity until eventually the last voltage increase triggers activity severe enough to cause permanent destruction. The frequency of voltage increases and the magnitude of those increases may well have an important effect.

8.3.2 Breakdown mechanism

Observations of the chip while these spikes were occurring revealed them to be indicative of visible arc-like discharges on the surface of the chip, most often in the gap between the electrode and the exposed silicon at the centre of each circular aperture. Occasionally the arcs were between the electrode and the metal or connection line of a neighbouring electrode, which had already been tested to destruction. These arcs make testing of negative polarity quite difficult by confusing the exact voltage being applied to the electrode – every μA of current means 200V is being dropped across the capacitor. Clearly, it would be difficult to operate a device if such an arc is in place. However, the insulation is not necessarily destroyed – the arc quenches if voltage is lowered. This is why the voltages for negative polarity breakdowns are limited to the last voltage before permanent hard-breakdown even though the devices stop operating properly above 700V or so.

On a number of occasions the arc burnt itself out and the current fell to purely resistive levels of a few nA – probably by evaporating away enough of the electrode to interrupt the breakdown path. Post-test observation of the chips in the SEM revealed an interesting picture. All of the images in figure 8.13 are taken from the same chip as figure 8.11 but in this case, the aluminium was negatively biased.

The damage caused seems to be characteristic of a propagating insulator breakdown [65]. This is particularly so in (b) and (c) where there are multiple holes adjacent to each other (shown more clearly in (d)). All the sites differ from those in figure 8.6 in that the aluminium has been evaporated away from the area surrounding the hole. The difference between the extent of the damage can be understood by examining the data logs for the relevant tests. The gate shown in (a) did not produce the large current spikes usually associated with negative tests – instead, when the voltage was 900V the current suddenly rose to 300nA for a short period before final destructive breakdown occurred. By contrast the gate in (c) and (d) suffered extensive current spikes greater than 1μA when the voltage reached 800V, these continued until permanent breakdown occurred when the voltage was 1100V. This link between breakdown damage and the activity of the breakdown current is consistent in the data logs.

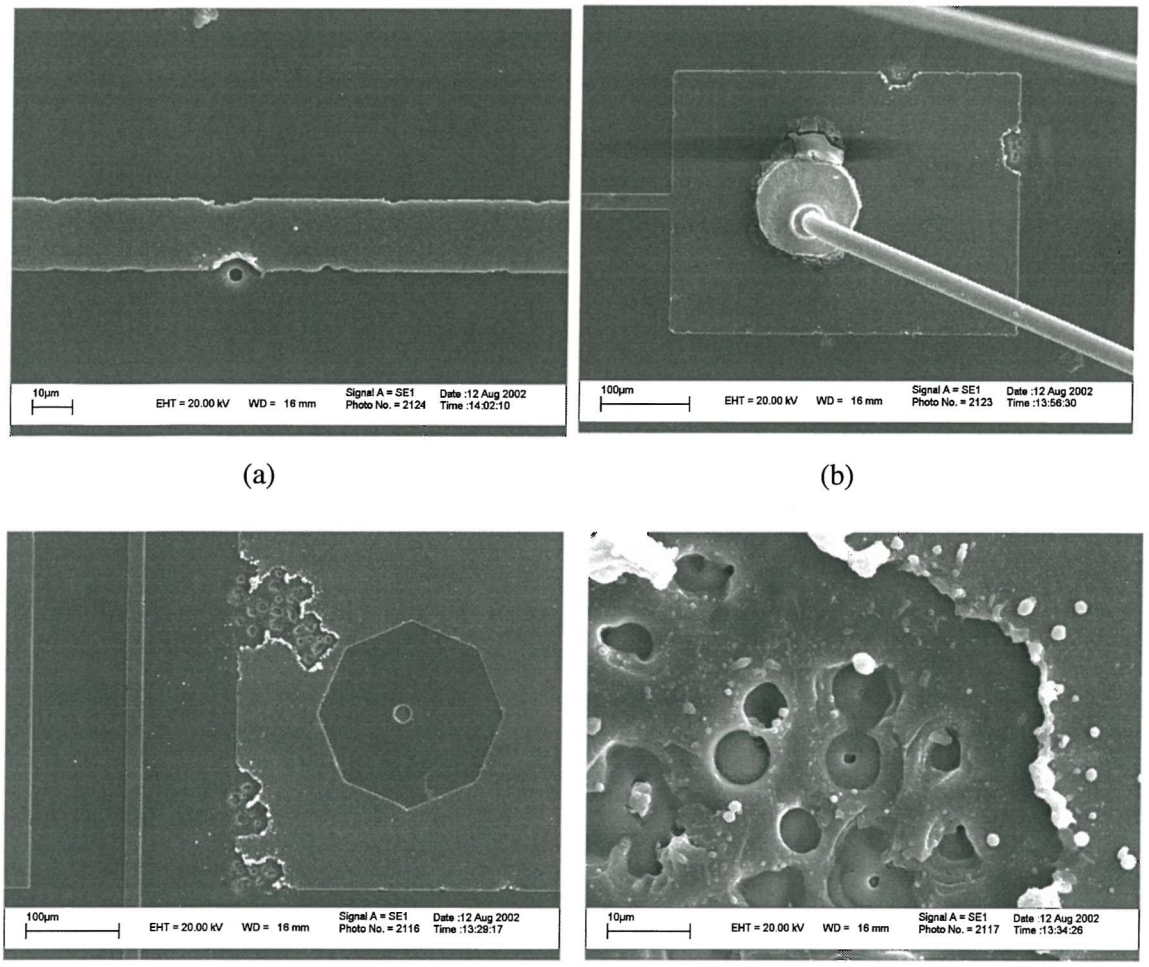

FIGURE 8.13: Breakdown damage for negative aluminium
(a) Site 1 on a connection line, (b) Site 2 on a bondpad
(c) Site 3 on an electrode edge, (d) Close-up of site 3

Clearly then the current spikes are indicative of the arcs that were directly observable; one can assume these arcs to be responsible for the damage shown. The melting back of so much aluminium would require the presence of high-temperature plasma for a relatively sustained period – such as the times required for unaided human observation. The breakdown seems therefore to be propagating from the initial hole to adjacent sites by way of these arcs. But why do these arcs occur only when the aluminium is biased negatively?

The reason for this altered level of damage must be due to the different cathode as it is the only changed parameter. As noted in chapter 2 – the material and surface topography of the electrode chosen to be the cathode can have a significant effect on breakdown. When the aluminium is positive the cathode surface is the largely unaltered, and therefore still smooth, surface of a polished single-crystal silicon wafer.

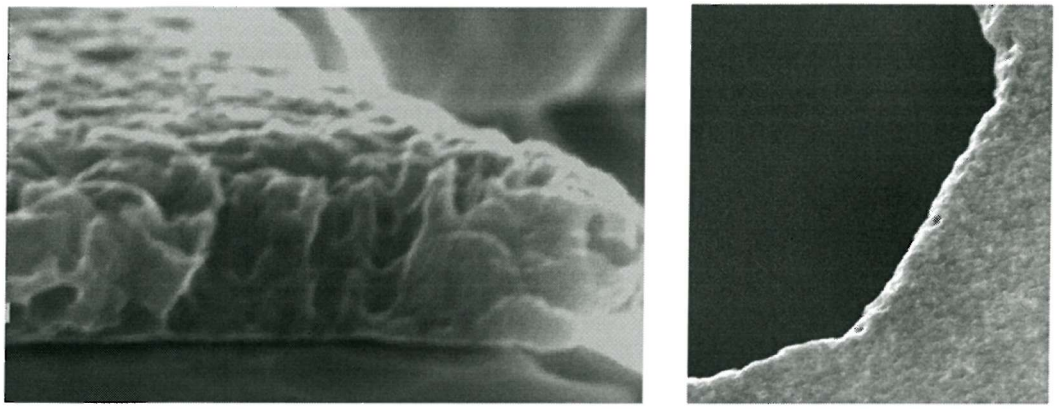

FIGURE 8.14: SEM images showing the roughness of the electrode edge.

In the negative mode the cathode becomes the whole of the aluminium electrode. The edges of the electrode are quite rough and craggy – exactly the type of topography which promotes significant electron emission. Figure 8.14 illustrates this point with SEM close-ups of the electrode edge. The problem is exacerbated by the presence of sharp cornered features originating from mask defects (visible in the bottom left of figure 8.11(a)). As described in chapter 2 – at the edge of an electrode is the metal-insulator-vacuum triple junction, a critical point at which electron emission is most likely to occur. It is reasonable to suggest that a cathode as rough as this would produce a much larger amount of electron emission. The cathode field is close to that needed for electron emission even with no geometrical field enhancement. An enhancement factor of 10 or less is sufficient to provide the required field strength.

The minimum pressure the chips were tested at was 10^{-7}mbar and when this was the case large bursts of pressure were seen to occur concurrently with the current spikes (as logged onto the PC). This suggests that large amounts of vapour were being produced from the chip surface – providing the seed gas for the formation of an arc. Perhaps then the breakdown mechanism changes from a single-hole to a propagating mechanism because large field-emitted electron currents, while insufficient to cause surface flashover, help form and sustain the arcs that spread the oxide-puncturing breakdowns from one site to the next.

It has been shown, with positively biased aluminium, it is much harder to puncture the oxide when the oxide is 4μm thick. Near the final destructive voltage (>3kV) however, a similar pattern of current spikes and visible surface arcs is observed. However, the arcs dominate and spread more vigorously. This can be seen by the widespread aluminium damage shown in figure 8.15, close observation reveals blanket oxide damage as opposed to the sharply defined holes seen previously.

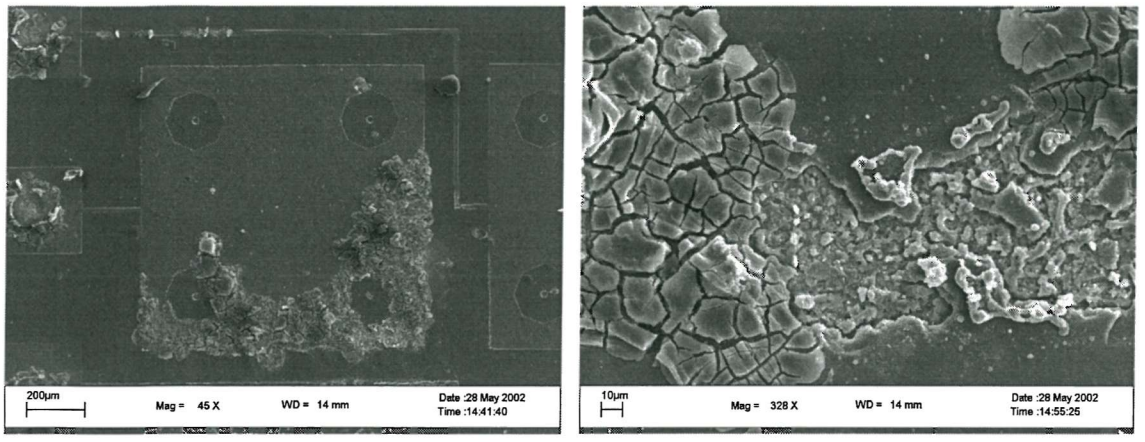

(a) (b)

FIGURE 8.15: Breakdown damage for very high negative voltages
(a) Arc damage on gate, (b) close up of damage

Clearly arcs are occurring when the aluminium is the cathode and the damage differs from that when the aluminium is positive, but it also differs from that seen above. The damage shown in figure 8.13 is so similar to that in figure 8.11, i.e. the very sharply defined holes in the oxide and the pits burnt into the silicon, that it seems reasonable to suggest that a bulk breakdown failure mechanism is at work. However, it also seems likely that the presence of the arcs is directly linked to the earlier onset of the final breakdown seen for negative aluminium. It is most probable that the cause of permanent failure is bulk breakdown, however the picture is complicated by the presence of the arcs, which cause damage and a reduction in the final breakdown strength of the insulator. Arcs were listed as one of the ways in which a propagating bulk breakdown can spread [65]. As the damage seen matches that reported for this failure mechanism, it seems that bulk breakdown propagated via surface arcs was probably responsible for the device failures, when the aluminium was negative. This could well be the reason for the different distribution types seen for negative and positive aluminium tests. If the breakdown mechanism for the negative case was purely bulk breakdown one would expect the data to follow a lognormal distribution as the positive data does. However the negative data is better characterised by a normal distribution – indicating a large amount of randomness, such as that created by the irregular arcs discussed in this section.

8.4 Validity of insulation comparison

The aim of this chapter and chapter 8 is to discover whether the described thruster design is adequate from an electrostatics standpoint. I have found that the SiO_2 insulation and the selected geometries can withstand sufficiently high voltages without breakdown. The present section will address the validity of the above results when examined in the light of actual thruster conditions. The insulation may be able to withstand high voltages in a test environment but how will that same design fare as part of an operational thruster?

8.4.1 Improvement over preliminary breakdown results

The values of breakdown voltage obtained in this experiment are a considerable improvement upon those attained in the preliminary tests of chapter 7. It is important to understand the driving force of that increase as it brings into sharp relief many of the lessons learnt during this work.

The main differences between the preliminary tests and the main insulation experiments are as follows:

- *Surface contamination*: the chips in the preliminary tests were very heavily contaminated. Microscope inspection has showed that chips that were not cleaned had a large amount of particulate contamination ranging up to tens of microns in size. Those chips were handled and stored in a very poor manner so by the time of testing they were no doubt contaminated with organic (grease, dust etc.), inorganic and maybe even metallic material. This contamination is dangerous for several reasons - large quantities of trapped gas are available for sudden release, surfaces can be coated in a conductive film and particles can become detached and cause impact damage. Finally, it has been reported that dielectric particles on the surface of a cathode can be a significant source of electron emission.

- *Testing circuit*: although there was a large series resistor in the circuit used in the preliminary tests the restriction it placed on the current was perhaps insufficient. For every kV there was 0.1mA available compared with less than 5μA in the main insulation experiments.

- *Chamber pressure*: all of the preliminary tests were performed with the chamber pressure above 10^{-4}mbar; this was achieved in 5 minutes of pumping. Although pressure was shown to have no effect in the main insulation experiments this was because breakdown occurred through the insulation bulk – I do not know that this was the case during the preliminary tests. Such a relatively high pressure combined with the lack of time for contaminant outgassing could be significant, when one considers the role of adsorbed gas in surface flashover.

- *Voltage application time*: depending on the breakdown mechanism being investigated the voltage ramprate can be significant. The voltages were raised very quickly (0 – 500V in ~2mins), which gave no time for pre-breakdown current spikes to settle down.

- *Voltage polarity*: all of the results in the preliminary tests were for the case where the silicon was grounded and a negative voltage was applied to the aluminium electrode. It was seen how this causes breakdown to occur in a very unstable way with destructive arcs across the chip surface.

The last point is key – if the chips in the preliminary tests were failing because of surface/vacuum breakdowns then all of the factors above would, to greater and lesser extents, be detrimental. The tests constitute a situation in which a heavily contaminated surface is subjected to bombardment from field-emitted electrons in a background of relatively high pressures. The arcs discussed in section 8.3 are certain to occur with greater ease at much lower fields. Because of the faster voltage rises they would occur with greater frequency and the smaller series resistor makes available much more energy with which to cause damage. It is little wonder that the results were so poor given the combination of potentially detrimental factors described here.

It's very likely though that the chips which accidentally had the oxide etched away would not hold much higher voltages even if they were tested in the manner detailed in chapter 8. Without the extended oxide the gap length is only 2μm – the equation in section 2.3.1 allows us to predict the breakdown voltage for a gap bridged by an insulator. Even for a contaminant free specimen in an ultra high vacuum the predicted voltage is only ~ 350 to 800V.

The factors discussed above may not be the cause of the premature breakdown of the chips from the SOI wafers, although they are probably poor conditions for such tests. Instead, those

chips could have been failing at 500V because of the lower dielectric strength of the PECVD oxide. If this were the case then, as with the damaged oxide chips, improving the experimental conditions would not show much improvement.

The results of the insulation batch indicate that the prescribed insulation and geometries are able to withstand voltages of 1kV and above in ~ 80% of cases. However, a possible concern is that these chips differ from the full thruster design in that they lack nozzles and the resulting liquid surfaces. Unfortunately, there were no more of the original chips available, except for those chips that have damaged oxides. So one needs to ask the question, what possible effect could the presence of a nozzle have upon the breakdown strength of the gate on a complete thruster chip; would any effects be detrimental?

Clearly, if the positive mode breakdown mechanism of such a chip remains through the bulk of insulating material there would be no change in the breakdown strength. Only the material properties of the insulator and electrodes influence the breakdown strength of that mechanism and these would be the same in each case. There would be no change to the external circuit either.

Therefore, the crux of the question is whether the presence of a nozzle and liquid could feasibly change the breakdown mechanism to a gap-related process.

8.4.2 A vapour heavy atmosphere

Before the onset of spraying the liquid in the nozzle represents a free surface, from which formamide vapour will evaporate copiously into the vacuum. The rate of formamide evaporation can be predicted using the Hertz-Knudsen equation:

$$\frac{dN_e}{A_e dt} = \frac{\alpha_v}{\sqrt{2\pi m k_B T}}\left(P^* - P\right).$$

To be conservative α_v can be taken as one, the maximum. If, since it is in a vacuum, P is taken as being zero then one can write:

$$\dot{m} = m\frac{dN_e}{dt} = A_e P^* \sqrt{\frac{m}{2\pi k_B T}}$$

The vapour pressure of formamide is ~10.7N/m^2, molecular mass = 7.97x10^{-26}kg and the temperature can be taken as ~ 300°K. These values give an evaporation mass flux of $\dot{m} = 1.9 * 10^{-2} A_e$ kg/s, where k_B is the Boltzmann constant. This shows the area dependence of fuel loss by evaporation, which was mentioned in chapter 3. A typical emitter has a nozzle radius of 2.5μm giving a n evaporation area, A_e~ 2x10^{-11}m^2; a probable, but very conservative, value for mass flux then is ~ 4x10^{-13}kg/s.

A computer simulation of the gas flow from an emitter nozzle, using mass fluxes of 1x10^{-13} and 1x10^{-12}kg/s, has been performed [81]. The analysis used a Direct Simulation Monte Carlo computational model to predict the motion of gas molecules in the domain outside an emitter. The geometry analysed is a 10 by 10 micron structured mesh with the mass flow of formamide vapour issuing from a boundary representing the liquid surface. The chamber pressure is represented by setting the far boundary to a pressure of 1.33x10^{-6} mbar. However, the low mass fluxes mean that no collisions were seen so the situation is a collisionless expansion into vacuum. For the lowest mass flux of 10^{-13}kg/s the ratio of computer particles to real particles was unity so this was actually a direct physical simulation.

A typical result is shown in figure 8.16 for a mass flux of 10^{-13}kg/s and shows the peak pressure occurring immediately outside of the nozzle with a value of 8.5x10^{-2} mbar. This falls rapidly to values of ~10^{-4} mbar within about 5μm out into the chamber. This illustrates that the formamide vapour does not remain near the gap but rather is pumped away; even at its peak, the pressure is not very high. With a mass flux of 10^{-12}kg/s the peak pressure is only 0.85 mbar and falls to 10^{-3}mbar within 10μm of the nozzle.

Section 8.1 presented results that show the breakdown strength of the insulation chips is unaffected by the presence of formamide vapour even at pressures as high as 10^{-3}mbar. The results of the computer simulation show that even with extremely high evaporation rates the formamide vapour quickly expands to pressures that I have shown to have no effect. Therefore I am confident that the evaporation of formamide from the liquid in the nozzle will not affect the breakdown strength of the devices. The voltage required to cause breakdown should be the same as that presented in section 8.1, whether fuel vapour is evaporating from the nozzles or not.

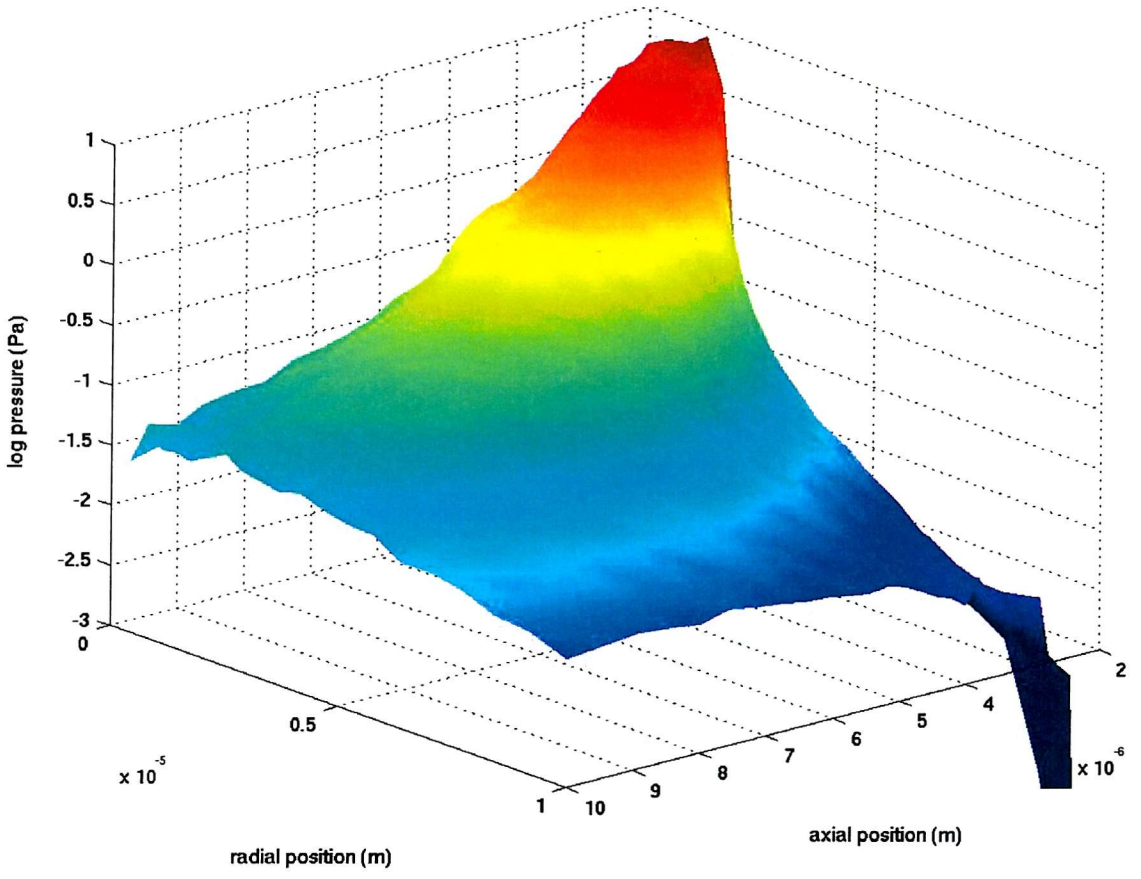

FIGURE 8.16: pressure distribution of formamide vapour around an emitter

8.4.3 Electron emission from the nozzle

For the insulation chips the cathode (during positive polarity testing) is a flat, featureless silicon surface that presents little or no opportunity for field emission of electrons. On the thruster chips though that surface is interrupted by a protruding silicon nozzle which could be a possible source of electrons. This is important since a source of electrons can lead to the occurrence of surface/vacuum breakdown as was seen in section 8.3.

However, the strongest electrical fields are very close to the axis of symmetry i.e. the centre of the nozzle, where there is no silicon, as shown by the electrostatic model and coaxial equation. A micron or so away from the centre and the field drops quickly down to an unchanging value of ~ 10^8V/m for every kV of gate voltage. To attain the electrical field strength of ~3x10^9V/m needed to start field emission a feature would need to have a geometric field enhancement factor, β, of 30 or so. Taking the worst case of representing the nozzle surface as a semi-ellipse of revolution then this value of β can only be obtained by a

feature which has λ=8 or above [52]. λ is the ratio of the semi-major to semi-minor axis of the ellipse and represents a measure of the features aspect ratio. With a height of only 2μm and a base width of > 6μm, the aspect ratio is not even equal to one, which would give a value of three for β.

The SEM image in figure 5.3 shows a close-up of a nozzle – it is very rounded and smooth – even the edges of the hole are smooth curves. This smoothness is caused by the thermal oxidation step during fabrication – during thermal oxidation almost half of the oxide thickness is obtained by consuming the silicon surface. This step is performed after the creation of the nozzle. When the oxide is etched away from the nozzle surface one can see the smooth edges that result from the silicon consumption.

The extremely low aspect ratio and very smooth, edge free surface of the nozzle mean it is very unlikely that the introduction of the nozzle will result in the emission of significant amounts of electrons.

8.4.4 Sprayed droplets on the surface

A potential pitfall of the design from an electrostatics point of view is what happens to the surface between the nozzle and the extraction electrode during spraying. There is the possibility that through collisions or coulombic fissions within the spray that a quantity of fuel may make it back to the emitter surface. If they do then a conductive film could establish itself on that surface – effectively providing a path for the field to short through. While the literature [69] talks about electrospray having a semi-cone angle of upto 30° no research has been carried out into the back-scattering of droplets. It would be very difficult to try to estimate the amount of liquid that makes its way back to the surface. This is not a typical monolayer formation problem not just because of interactions in the spray but also since the electrostatics would force the drops to return preferentially to the electrode. The answer to this question can only really be discovered by operation of a device – an option unavailable to us since there are no suitable chips left.

It seems most probable that this is not a significant cause for concern for the following reasons. Since their surface charge forces them apart, drops could only really collide if a smaller faster drop were to catch upto a larger, slower drop. However, there is a size segregation effect in electrosprays in which smaller drops (of a higher q/m) experience

stronger electrostatic repulsion forces and are pushed to the outside of the spray. Another possible method by which liquid could come back to the surface is through Coulombic fissions; this is where as solvent evaporates, the droplet becomes too small to support the surface charge and the drop fissions into 'daughter droplets'. However, these events are also considered unlikely because the extremely high conductivities and low flowrates to be used mean that charge is shed from the droplets by ion evaporation instead [82].

In light of the above discussion there seems little reason to suspect that the introduction of a nozzle and fuel would have a detrimental effect on the devices ability to withstand electrostatic fields and potentials. Therefore it is asserted that the full thruster chips would break down due to the same mechanism and at the same mean voltage as presented earlier in this chapter for the insulation chips.

Chapter 9

Conclusions

9.1 Main points of this work

The initial aim of this work was to produce a prototype for a colloid microthruster based on MEMS technology. Traditionally-implemented colloid thrusters suffer from high mass, inconsistency in emitter dimensions, large evaporative fuel loss, high starting voltages and difficulties in implementing thrust throttling. Chapters 3 and 4 describe how a new approach using the full potential of MEMS techniques can provide solutions to all of these issues.

A design has been successfully fabricated which has very small diameter nozzles, with a consistent definition; fuel evaporation from such small nozzles would be very low and a spray should be obtainable at voltages less than a kV. This works represents the first to fabricate arrays of nozzles for use in colloid propulsion (although nozzles have been microfabricated for use in Biochemistry – section 2.2.3). A feasible design for a microfabricated colloid thruster has been drawn up which promises to produce thrusts at the very low levels required for microspacecraft with good overall performance. The pitfalls of the microfabrication process were identified and, in some cases, the process was modified to assist in their avoidance in the future. Arrays of 484 emitters were fabricated on a chip that is only 1.6cm wide.

A robust packaging system has been developed which has quickly changeable interconnections. It allows fuel to be fed to the nozzle at a prescribable pressure in a vacuum environment. Tests have shown that the silicon nozzles do not suffer from liquid leakage or external wetting. If pressure is sufficiently low, the liquid is held firmly at the exit of the nozzle. This solves the problems of interfacing a fuel feed to a microfabricated electrospray chip in vacuum for the first time.

Although it remains untested, the concept of throttling thrust using the digital method detailed in section 4.2.2 is an original idea for colloid thrusters. I have shown how the required electrodes can be fabricated – if successful, this approach would be of significant advantage to propulsion systems that wish to use colloid thrusters. It would mean very accurately controlled thrusts over large ranges without degrading performance, without any of the complication or large masses associated with pumps, valves and pipework.

The above paragraphs describe part of the narrowly focused contribution of this work, i.e. the progress made in overcoming the obstacles associated with micro-fabricating a colloid microthruster. A further, and significant, part of that progress is the demonstration of novel on-chip electrodes for providing an electrical field of sufficient strength to start an electrospray. The geometries used and the choices of insulation type and thickness have been proven suitable. Fields greater than $5x10^8$V/m are created in the nozzle vicinity using 1kV applied voltage and the device survives, these values are more than adequate to create an electrospray and with good specific impulse. This holds the door open for researchers to produce a colloid microthruster using MEMS techniques and materials, with the confidence that they can create and safely maintain the electrostatic fields required.

The results of the insulation studies also make a more general contribution. This is the first study of micro-fabricated electrodes on the surface of a device in which very high voltages - upto 3kV, are used. It contributes knowledge to the, otherwise extensively researched, field of silicon dioxide breakdown. Whilst making no claims that the understanding of oxide breakdown has been enriched, the work can claim to have demonstrated the voltage hold-off capability of extremely thick oxides (>2μm) for the first time. A number of observations strongly suggest that electrical breakdown occurs in these materials via a bulk breakdown mechanism (when the aluminium is positive) and the ambient environment has little influence (over the ranges tested). Specifically, the results show that a high yield of chips (~80%) 1kV can safely be taken by 2μm of thermal SiO_2, and upto 3kV can be maintained if a further 2μm of CVD oxide is deposited. Some evidence was also presented which suggests that the breakdown mechanism for negatively biased aluminium is an arc-assisted bulk breakdown. This is significant because voltages of a kV or above are useful for electrostatic applications involving the acceleration of charged particles

As an example of how this may be useful, and staying within the spacecraft propulsion field, I have previously mentioned researchers assessing the suitability of SiO_2 insulation for ion

engine acceleration grids. These grids are required to safely hold potential differences of 1kV. The devices tested were found to be lacking (section 2.2.3) and there have been no further publications in that direction. However, that need not have been the case – this work clearly demonstrates that with a higher quality oxide, more precise experimental conditions would have obtained breakdown voltages more than adequate for their needs. Testing should not be done at atmospheric pressure (as surface discharges uncharacteristic of the intended environment will occur) but I have shown that the insulation works even at pressures typical of an ion thruster operation.

9.2 Assessment of design success

It is necessary to evaluate whether or not the presented colloid microthruster design is suitable for its purpose. I was hampered in this by the lack of chips with the full array of features – the insulation chips have no nozzles and the remaining thruster chips only have damaged or inferior insulation. This means that I am unable to see whether the design works by simply applying sufficient voltage to a chip that is filled with fuel. Therefore I can only assess the totality of the design in theoretical terms. However, the preceding sections allow us to take a critical look at the constituent elements of a colloid microthruster and assess the extent to which our design meets their demands.

9.2.1 Nozzle suitability

Any design for a colloid emitter must consistently define a liquid surface and expose it to the space vacuum; the surface should have a small area for the reasons described in section 3. The liquid needs to be in direct contact with a conductive material and surface tension must hold it in place.

The nozzles designed for this study have accurately etched exit holes with diameters of 3 or 5μm, which give a vast reduction in evaporative fuel loss. The nozzles allow the liquid to reach the exits with a minimal pressure head but the small dimensions mean surface tension keeps the liquid safely at the nozzle exit as demonstrated in chapter 7. The very low hydraulic resistance of the nozzle geometry means that liquid flow is to be controlled by the voltage applied. This has the advantage that liquid only flows out of the nozzles when the

voltage is high enough to start electrosprays. In that case, the liquid will flow in the form of a spray, not into an enlarging drop that spreads over the emitter as is sometimes seen when the flow is pressure driven. As the nozzles are fabricated in silicon, the potential of the liquid is controlled by applying voltage directly to the silicon substrate. This use of voltage-controlled flow is not new to the electrospray field (section 4.2.2) but this is the first design to consider its use in a microfabricated colloid thrusters.

9.2.2 Electrostatics

Aside from the successful delivery of fuel, the other crucial element of a colloid thruster is the provision of a strong enough electric field. The origin of this field is not important – commercial nanoElectrospray instruments simply place the capillary exit to one side of a skimming nozzle that is held at a high voltage relative to the needle. This provides an electrical field that is asymmetrical but is strong enough at the capillary tip to overcome the surface tension and produce the spray. This is all that is required to produce the spray – other electrostatic factors are important due to their influence on beam profiles, focussing etc.

The electrostatic model in chapter 4 and the simple equation in section 3.3 verify that a coplanar electrode set-up with the designed geometries will produce the required electric field if the potential difference is above 800V or so. With the theoretical assessment of the electric field strength in agreement the critical question then is whether the design can withstand voltages of this magnitude – a question that I have provided a firmly positive answer to. In all configurations, provided the experimental conditions of chapter 8 are stuck to, a high yield (~80%) of the devices can successfully hold 1000V or more.

From a propulsion standpoint 1kV of acceleration is enough, providing as it does ~ 500s of predicted specific impulse. There is no reason though why any subsequent design should not use the thicker oxides, which have been shown capable of holding almost 3kV – this would increase specific impulse to 800s.

A downfall of the current design is the disparity between the breakdown strength in the positive mode and that obtained for the negative mode. It is unfortunate that the lower voltages result from having the electrode negative, as it is this mode that would produce positively charged droplets. Electric propulsion techniques utilise positive ions because of spacecraft charging. As positive charges stream away from a spacecraft in orbit an equal current of negative charges must also be emitted otherwise the spacecraft quickly obtains

excess negative charge. This not only then acts to pull back the positive ions but can also lead to arcing around the spacecraft. It is accepted practice to neutralise the beam of positive ions leaving a thruster by using an electron emitter – as the currents are easily obtained and at negligible expense to the mass budget.

However, this is not too severe a problem. The discussion in section 8.3 talks about the occurrence of arcs across the device surface when voltage is taken above 600-700V but it also mentions how these arcs are seen to die down if given time. This is an example of the conditioning effect, which once applied to an insulation gap, remains in place. If a device is conditioned by gently raising it to 900V without hard-breakdown occurring and letting current spikes die down, that device can subsequently hold voltages close to 900V without experiencing arcs. A thruster is then available which would produce positive ions with quite decent specific impulse. Alternatively, a thruster could be operated in a bipolar mode as described in section 4.3.3, so the currents of positive and negative ions neutralise each other.

9.3 Suggestions for future work

This work has clearly shown that micro-fabrication techniques and materials can produce an electrode system capable of creating the electrostatic fields required by colloid microthrusters. From a theoretical point of view the design seems suitable for producing electrosprays, however the lack of undamaged thruster chips has prevented this from being confirmed experimentally. To do so would require the production of a new batch of similar chips but with the fabrication process modified in the following ways:

- To maximise the breakdown strength of the insulation the superior thermal oxide should be used whenever possible. Once the maximum thickness obtainable by thermal oxidation has been reached, the addition of CVD oxide to increase the thickness has been shown to have an impressive effect.
- Thin membranes can buckle or break with damaging effect. Steps should be taken to prevent this, for example using DRIE to produce deeper holes (about 15-20μm) or, saving the membrane thinning step till the end - as adopted for the SOI wafers in this work.
- Since bulk breakdown depends on the area of the electrode (results showed that those electrodes with large electrodes broke down more easily), they should be kept as small as possible. One electrode per emitter would be the ideal.

- Electrode connections should be wide enough to allow for significant overetches – these may be required to clear the aluminium from the apertures.
- Consideration should be given to making the aluminium electrodes thin enough to allow the self-healing breakdown described by Klein – this would allow the weak spots of the insulation to be eliminated, so raising the breakdown voltage of the chips.
- The metal-insulator-vacuum triple junction is thought to be the root cause behind the lower breakdown voltages when the aluminium is negative. There are various ways to alleviate this as described in section 2.3.1.
- To allow a closer examination of individual emitter performance it may be advisable to fabricate chips that have only one emitter – although this would mean that the failure of one nozzle would incapacitate the entire chip.

The first objective of work that fabricates such thruster chips must be to observe the spray produced and characterise it as fully as possible. Questions that will require answering are:

- Do the emitters produce a spray? Is it in the desired cone-jet mode (the most stable and understood regime)? Is there any leakage of propellant over the emitter surface when spraying; does this facilitate breakdown? These questions can be best addressed by video observation of the nozzles.
- How does the mass flowrate of the emitters vary with the applied voltage. Do the current-voltage characteristics follow the $I \alpha Q^{1/2}$ relation and are they independent of electrostatics? Flowrates need to be accurately measured, one method available is via the accurate determination of fluid pressures either side of a flow-tube. Currents should be collected via a secondary electron suppressed collector and all sources of leakage currents need to be eliminated.
- What do the sprays consist of? Are there a large number of satellite droplets and/or ions produced? Time of flight or particle sizing analysers can be useful to determine the q/m and diameters of the produced droplets.
- What happens when individual sub-clusters are turned on with neighbouring sub-clusters turned off? Do only those emitters turned on emit a spray? If so are the answers to the above questions affected?

Answering these questions will give a thorough understanding of the emitters operational characteristics and provide the answers to most questions of thruster performance, such as Isp, efficiency and power. Thrust can be calculated using the same data but researchers may prefer to perform a direct thrust measurement.

Bibliography

1. J. Mueller, W. Tang, W. Li, A. Wallace. "Micro-fabricated accelerator grid system feasibility assessment for micro-ion engines", Paper no. IEPC 97- 071, 25th International Electric Propulsion Conference, Ohio. (1997)

2. S. W. Janson. "Batch-fabricated resistojets: initial results", Paper no. IEPC 97- 070, 25th International Electric Propulsion Conference, Ohio. (1997)

3. J. Mueller, W. Tang, W. et al. "Design, analysis and fabrication of a vaporizing liquid micro-thruster", Paper no. AIAA 97-3058, Proc. AIAA/ASME/SAE/ASEE Joint Prop. Conf.., Seattle, WA, (1997)

4. P. Fortescue, J. Stark. "Spacecraft systems engineering", 2nd edition, John Wiley & sons, UK (1997)

5. http://www.moog.com/Media/1/SPDcat_sec4.pdf Model 58E125A (as of 06/08/02)

6. J. Kohler, U. Simu et al. "Feasibility Demonstration of the Micropropulsion Cold Gas Thruster system", Procs. Of the 3rd round table on micro/nano technologies for space, Noordwijk, (2000)

7. R. Bayt, A. Ayon, K. Breuer. "A performance evaluation of MEMS-based micronozzles", Paper no. AIAA 97-3169, 33rd AIAA/ASME/SAE/ASEE Joint propulsion conference & exhibit, Seattle, WA. (1997)

8. Ketsdever, J. Mueller. "Systems considerations and design options for microspacecraft propulsion systems", AIAA paper 99-2723 35th AIAA Joint propulsion conference, 20-24 June, LA, (1999)

9. C. Bruno. "Chemical microthrusters: effects of scaling on combustion". AIAA paper 2001-3711, 37th AIAA joint propulsion conference, 8-11 July, Utah, (2001)

10. A. London, A. A. Ayon, A. Epstein et al, "Microfabrication of a high pressure bipropellant rocket engine", Sensors and Actuators A 92, pp 351-357, (2001)

11. G. Sutton, D Ross. " Rocket propulsion elements", fig 13.4, 4th edition, John Wiley & sons, USA

12. H. Horisawa, I. Kimura. "Very low-power DC plasma jet microthrusters", AIAA paper 2001-3791, 37th AIAA joint propulsion conference, 8-11 July, Utah, (2001)

13. M. Martinez-Sanchez, J. Pollard. "Spacecraft electric propulsion – an overview", Journal of propulsion and power, V14, no. 5, pp 688-699, (1998)

14. T. Haag. "Thrust stand for pulsed plasma thrusters", Rev. Sci. Instrum. V68, (5), pp 2060-2067, May (1997)

15. M. Wilson, S. Bushman, R. Burton. "A compact thrust stand for pulsed plasma thrusters", Paper no. IEPC 97-122, 25th International Electric Propulsion Conference, Ohio. (1997

16. H. Takegahara, M. Igarashi, et al., "Evaluation of Pulsed Plasma Thruster system for μ-Lab Sat II", IEPC paper no. 01-153, 27th IEPC, October 15-19, Pasadena, (2001)

17. N. Qi, J. Schein et al., "Compact vacuum arc micro-thruster for small satellite systems", AIAA paper no. 2001-3793, 37th AIAA joint propulsion conference, 8-11 July, Utah, (2001)

18. M. Patterson, S. Grisnik, G. Soulas. "Scaling of ion thrusters to low power", Paper no. IEPC 97-098, 25th International Electric Propulsion Conference, Ohio. (1997)

19. G. Yashko, G. Giffin, D. Hastings. "Design considerations for ion microthrusters", Paper no. IEPC 97-072, 25th International Electric Propulsion Conference, Ohio. (1997)

20. M. Young, E. Muntz, A. Ketsdever, D. Wadsworth. "Unique hollow cathode as a code validation experiment and candidate non-magnetic ion micro-thruster", Paper no. AIAA 99-2854

21. P. Gessini, S. Gabriel. "The Hollow Cathode as a Micro-Ion Thruster", IEPC paper no. 01-233, 27th IEPC, October 15-19, Pasadena, (2001)

22. M. Patterson, J. Foster. "Hollow Cathode Micro-Thruster Performance", IEPC paper no. 01-226, 27th IEPC, October 15-19, Pasadena, (2001)

23. Y. Raitses, J. Ashkenazy, G. Appelbaum, M. Guelman. "Experimental investigation of the effect of channel material on hall thruster characteristics", Paper no. IEPC 97-056, 25th International Electric Propulsion Conference, Ohio. (1997)

24. G. Guerrini, C. Michaut, M. Bacal. "An intense hall-type ion source for satellite propulsion", Rev. Sci. Instrum. V69, No. 2, p804, (1998)

25. V. Khayms, M. Martinez-Sanchez. "Preliminary experimental investigation of a miniaturised hall thruster", Paper no. IEPC 97-077, 25th International Electric Propulsion Conference, Ohio. (1997)

26. S. Marcuccio, S. Gianelli, M. Andrenucci. "Attitude and control of small satellites and constellations with FEEP thrusters", 25th International Electric Propulsion Conference, Ohio. (1997)

27. A. Genovese, W. Steiger, M. Tajmar. "Indium FEEP Microthruster Characterization in the 1-100μN range", AIAA paper no. 2001-3788, 37th AIAA joint propulsion conference, 8-11 July, Utah, (2001)

28. D. Lewis, S. Janson, R. Cohen, E. Antonsson, "Digital Micropropulsion", IEEE, MEMS '99, the 12th annual international workshop on MEMS, (1999)

29. D. Youngner, S Lu, E. Choueiri et al. "MEMS Mega-pixel Micro-thruster Arrays for Small Satellite Stationkeeping", SSC00-x-2, 14th Annual/USU Conference on Small Satellites (2000)

30. D. Gonzales, R. Baker. "Microchip Laser Propulsion for Small Satellites", AIAA paper 2001-3789, 37th AIAA joint propulsion conference, 8-11 July, Utah, (2001)

31. J. Perel, T. bates, J. Mahoney, R. Moore, A. Yahiku. "Research on a charged particle bipolar thrustor", Paper no. AIAA 67-728, AIAA Electric propulsion and plasmadynamics conference, Colorado Springs, USA, (1967)

32. J. Mahoney, H. Daley, J. Perel. "Performance of colloid annular emitters", Paper no. AIAA 73-1076, AIAA 10th Electric propulsion conference, Lake Tahoe, USA, (1973)

33. H. Daley, J. Mahoney, J. Perel. "Colloid annular array thruster development", Paper no. AIAA 73-1077, AIAA 10th Electric propulsion conference, Lake Tahoe, USA, (1973)

34. P. Kidd, H. Shelton. "Life test (4350 hours) of an advanced colloid thruster module", Paper no. AIAA 73-1078, AIAA 10th Electric propulsion conference, Lake Tahoe, USA, (1973)

35. M. Martinez-Sanchez, J. Fernandez de la Mora, V. Hruby, M. Gamero-Castano, V. Khayms. "Research on colloid thrusters", Paper no. IEPC 99-014, 26th International Electric propulsion conference, Japan, Oct (1999)

36. M. Gamero-Castano, V. Hruby, "Characterization of a Colloid Thruster Performing in the micro-Newton Thrust Range", IEPC paper no. 01-282, 27th IEPC, October 15-19, Pasadena, (2001)

37. V. Hruby, M. Gamero-Castano, P. Falkos, S. Shenoy "Micro Newton Colloid Thruster System Development", IEPC paper no. 01-281, 27th IEPC, October 15-19, Pasadena, (2001)

38. A Shtyrlin. "State of the art and future prospects of colloid electric thrusters", Paper no. IEPC 95-103, 24th International electric propulsion conference, Moscow (1995)

39. F. Pranajaya, M. Cappelli, "Performance studies of A Colloid Thruster System", IEPC paper no. 01-284, 27th IEPC, October 15-19, Pasadena, (2001)

40. J. Perel, J. Mahoney, "Operational Features of a Linear Slit Colloid Microthruster", IEPC paper no. 01-116, 27th IEPC, October 15-19, Pasadena, (2001)

41. D. Kirtley, J. Fife. "A colloid engine accelerator concept", AIAA paper 2002-3811, 38th AIAA joint propulsion conference, 7-10 July, Indianapolis, Indiana (2002)

42. A. Spindt. "A thin film field emission cathode." J. Appl. Phys. V39, pp 3504-3505 (1968)

43. J. She, S. Huq, J. Chen, S Deng, N, Xu. "A comparative study of electron emission characteristics of Silicon tip arrays with and without Amorphous diamond coating", IEEE (1997)

44. K. Hutchison. "Ion microthrusters for attitude and orbital control of small satellites", PhD 9 month report, University of Southampton (2000)

45. C. Spindt. "Microelectronic field ionizer and method of fabricating the same", US Patent no. 4,926,056. May (1990)

46. J. Mitterauer. "Miniaturized liquid metal ion sources (MILMIS)", IEEE transactions on plasma science, V 19, No. 5, Oct. (1991)

47. L. Licklider, X. Wang et al. "A micromachined Chip-based Electrospray Source for Mass Spectrometry", Anal Chem 72, pp 367-375, (2000)

48. G. Schultz, T. Corso, S. Prosser, S. Zhang. "A Fully Integrated Monolithic Microchip Electrospray Device for Mass Spectrometry", Anal Chem 72, pp 4058-4063, (2000)

49. M. Paine. "Design study for micro fabricated colloidal thrusters", S.M. thesis, Massachusetts Institute of Technology, (2000)

50. L.F. Velasquez, M. Martinez-Sanchez. "A micro-fabricated colloid thruster array", AIAA paper 2002-3810, 38th AIAA joint propulsion conference, 7-10 July, Indianapolis, Indiana (2002)

51. J. Xiong, Z. Zhou et al. "The initial research on colloid micro-thruster system", MNE 2001, paper no. MS-16P, Sept 16-19, Paris, (2001)

52. R. V. Latham. "High voltage vacuum insulation: The physical basis", Academic press, England. (1981)

53. C. Le Gressus, Ph. Maire, J. P. Duraud. "The influence of surface phenomena on the initiation of discharges in vacuum", IEEE transactions on electrical insulation, vol. 24, no. 6, December. pp. 969-977. (1989)

54. A. S. Pillai, R. Hackam. "Surface flashover of solid dielectric in vacuum", J. Appl. Phys. vol. 53, no. 4, April, pp.2983-2987. (1982)

55. L. L. Hatfield, E. R. Boerwinkle et al. "Methods of increasing the surface flashover potential in vacuum", IEEE transactions on electrical insulation, vol. 24, no. 6, December. pp. 985-993. (1989)

56. A. Roth. "Vacuum technology", chapter 2, North Holland publishing company, Amsterdam. (1976)

57. http://www.physics.nist.gov/cgi-bin/Ionization/parsesv.pl?N2 (as of 06/08/02)

58. F. Paschen. Wied. Ann. vol. 37, no. 69. (1889)

59. F. Llewellyn-Jones. "Ionization and breakdown in gases", Methuen & co. Ltd., London. (1957)

60. X. Ma, J. Kim, T. Sudarshan. "High field breakdown characteristics of micrometric gaps in vacuum", Technical Digest of International Vacuum Microelectronics Conference, Korea, (1997)

61. P. G. Muzykov, X. Ma, et al. "High field breakdown of narrow quasi uniform fields in vacuum", J. Appl. Phys. vol. 85, no. 12, June, pp. 8400-8404, (1999)

62. N. Klein. "Switching and breakdown in films", Thin Solid films, vol. 7, pp 149-177. (1971)

63. P. Solomon. "Breakdown in silicon oxide – A review", J. Vac. Sci. Technol., vol. 14, no. 5, Sept./Oct. (1977)

64. N. Klein, P. Solomon, J. Appl. Phys. vol. 47, p4364. (1976)

65. N. Klein. "The mechanism of self-healing electrical breakdown in MOS structures", IEEE transactions on electron devices, vol. 13, no. 11, November. (1966)

66. C. M. Osburn, D. W. Ormond. "Dielectric breakdown in silicon dioxide films on silicon", J. Electrochem. Soc. Vol. 119, pp591-597. (1972)

67. C. M. Osburn, E. Bassous. "Improved dielectric reliability of SiO_2 films with polycrystalline silicon electrodes", J. Electrochem. Soc, vol. 122, pp. 89-91. (1975)

68. G. I. Taylor. "Disintegration of water drops in an electric field", Proc. R. Soc. Lond. A V280, pp383-397. (1964)

69. J. Fernandez de la Mora, I. G. Loscertales. "The current emitted by highly conducting Taylor cones", J. Appl. Phys, vol. 83, no. 6, March, pp. 1-7, (1998)

70. A. G. Bailey. "Electrostatic spraying of liquids", pg. 67, Research studies press Ltd., England. (1988)

71. M. Martinez-Sanchez. Lecture notes from class 16.522, Massachusetts Institute of Technology, (1999)

72. J. Rosell-Llompart, J. Fernandez-de la Mora. "Generation of monodisperse droplets 0.3 to 0.4μm in diameter from electrified cone-jets of highly conducting and viscous fluids", J. Aerosol Sci., vol. 25, no. 6, pp. 1093-1119, (1994)

73. R. Hartman, D. Brunner et al, "Jet break-up in electrohydrodynamic atomization in the cone-jet mode", J. Aerosol Sci., vol. 31, no. 1, pp. 65-95, (2000)

74. M. Gamero-Castano, V. Hruby, "Electrospray as a Source of Nanoparticles for efficient Colloid Thrusters", AIAA paper 2000-3265

75. M. Gamero-Castano, V. Hruby, "Characterization of Electrosprays and Its Relevance to Colloid Thrusters", IEPC paper no. 01-283, 27th IEPC, October 15-19, Pasadena, (2001)

76. J. Crawley. "Fundamentals of applied electrostatics", pg 23, Wiley, New York, (1985)

77. M. Gollor, K. Rogalla. "High voltage design of vacuum insulated power supplies for space applications", XVth International Symposium on discharges and electrical insulation in vacuum, Darmstadt (1992)

78. J. C. Hannis, D. C. Muddiman. "Nanoelectrospray mass spectrometry using non-metalized, tapered (50-10μm) fused-silica capillaries", Rapid Commun. Mass Spectrom. vol. 12, pp. 443-448. (1998)

79. M. Gamero-Castano, J. Fernandez de la Mora. "Direct measurement of ion evaporation kinetics from electrified liquid surfaces", J. of Chemical physics, vol. 113, no. 2, (2000)

80. Y. C. Yong, J. T. L. Thong, J. C. H. Phang. "Determination of secondary electron yield from insulators due to a low-kV electron beam", J. Appl. Phys. vol. 84, no. 8, October, pp. 4543-4548. (1998)

81. F. T. A. Crawford, S. B. Gabriel. "Modelling Small Hollow Cathode Discharges for Ion Microthrusters", AIAA 2002-2101, 33rd AIAA Plasmadynamics and lasers, May 20-23, Maui, HI. (2002)

82. P. Kebarle. "A brief overview of the present status of the mechanisms involved in electrospray mass spectrometry" J. Mass Spectrom. vol. 35, pp. 804-817. (2000)

APPENDIX A: Derivations

Derivation of polydisperity efficiency

$$\eta = \frac{KE\ to\ beam}{Power\ in} = \frac{\frac{1}{2}\dot{m}c^2}{IV}\ ,\ where:\ T = \dot{m}c,\ so:\ \eta = \frac{T^2}{2\dot{m}IV}$$

$$but,\ T = \dot{m}\sqrt{\frac{2Vq}{m}},\ = \dot{m}\sqrt{2V}\left\langle\sqrt{\frac{q}{m}}\right\rangle,\ and,\ I = \dot{m}\left(\frac{q}{m}\right) = \dot{m}\left\langle\frac{q}{m}\right\rangle\ which\ gives:$$

$$\eta = \frac{\dot{m}^2 * 2V\left\langle\sqrt{\frac{q}{m}}\right\rangle^2}{2V\dot{m}\dot{m}\left\langle\frac{q}{m}\right\rangle} = \frac{\left\langle\sqrt{\frac{q}{m}}\right\rangle^2}{\left\langle\frac{q}{m}\right\rangle}$$

$$Now,\ \left\langle\frac{q}{m}\right\rangle = \alpha_s\left(\frac{q}{m}\right)_s + (1-\alpha_s)\left(\frac{q}{m}\right)_m\ \ and,\ \left\langle\sqrt{\frac{q}{m}}\right\rangle = \alpha_s\left(\sqrt{\frac{q}{m}}\right)_s + (1-\alpha_s)\left(\sqrt{\frac{q}{m}}\right)_m$$

where α_i is the mass fraction of ions. Substituting gives:

$$\eta = \frac{\left(\alpha_s\left(\sqrt{\frac{q}{m}}\right)_s + (1-\alpha_s)\left(\sqrt{\frac{q}{m}}\right)_m\right)^2}{\alpha_s\left(\frac{q}{m}\right)_s + (1-\alpha_s)\left(\frac{q}{m}\right)_m}\ \ or\ \ \eta = \frac{\langle c\rangle^2}{\langle c\rangle(c_s + c_m) - c_s c_m}$$

$$where,\ c_s = \sqrt{2V}\sqrt{\left(\frac{q}{m}\right)_s}\ ,\ c_m = \sqrt{2V}\sqrt{\left(\frac{q}{m}\right)_m}\ \ and,\ \langle c\rangle = \alpha_s c_s + (1-\alpha_s)c_m$$

Derivation of electrically dominated flowrate

$$\delta P = P_H - P_S + P_E = \frac{128\mu LQ}{\pi d^4} \quad by\ Pouseille.$$

where P_S is surface tension pressure, $P_S = \frac{2\gamma\cos\alpha}{r}$

and P_E is the electrical field pressure, $P_E = \frac{\varepsilon_0 E^2}{2}$, $and\ E = \frac{2V}{\ln\left(4\,s/r\right)r}$

The equation for E is for an electrode at a distance above the needle electrode, and uses a hyperboloid geometry in the analysis [71].

(A simpler analysis uses two conducting coaxial cylinders [76] which is perhaps the most suitable for our emitter arrangement, both are used in the document to give us a range. The E field for a coaxial analysis is: $E = \frac{V}{\ln\left(s/o\right)r}$ where r can vary and o is the radius of the inner conductor and will be used at the end of this section.)

Using the Hyperboloid equation,

$$so,\quad \frac{16*128\mu LQ}{\pi r^4} = P_H - \frac{2\gamma\cos\alpha}{r} + \frac{\varepsilon_o}{2}\left(\frac{2V}{\ln\left(4s/r\right)r}\right)^2$$

$$\text{which gives, } Q = \frac{\pi}{16*128\mu L}\left[\Pr^4 - 2\gamma\cos(\alpha)r^3 + \frac{2\varepsilon_0 V^2}{\left(\ln\left(4s/r\right)\right)^2}r^2\right]$$

where the subscript of P_H has been dropped for clarity.

The equations for P_S , P_E and E also allow us to prove the starting voltage equation, since for flow to begin $P_E > P_S$ so,

$$\frac{\varepsilon_0 * 4 * V^2}{2\left(r * \ln\left(4s/r\right)\right)^2} > \frac{2\gamma\cos\alpha}{r}$$

$V > \sqrt{\frac{r * \gamma\cos\alpha}{\varepsilon_0}} \ln\left(\frac{4d}{r}\right)$ using the hyperboloid equation.

Using the coaxial equation instead gives, $V > 2\sqrt{\frac{r * \gamma\cos\alpha}{\varepsilon_0}} \ln\left(\frac{d}{o}\right)$, where o is the radius of the silicon that is not covered by oxide. This equation is only valid as a rough guide and is included to provide a wider and more pessimistic range of starting voltage.