LHCb-INT-2015-XYZ
November 24, 2015

VELO Upgrade
Engineering Design Review

Microchannel substrate

Prepared by,
Malcolm John, Matthew Brock
for the VELO upgrade project.

Reviewed by,
Alessandro Mapelli, Eric Thomas

1 Scope

Given the decision to implement microchannel cooling, this review should consider whether it is demonstrated that the main concerns relating to the concept of microchannel cooling has been addressed and thus if the proposed design of the substrate is ready for entering an ordering process. The design feature considered here are:

- The dimensions of the substate
- The channel dimensions and the circuit
- The metallisation pattern
- Assessment of remaining risk factors

2 Introduction

The VELO upgrade TDR [1] describes the decision replace the current silicon-strip based vertex detector with a pixel system based on the Timepix family of pixel chips and to be capable of reading out the detector at 40MHz. A longstanding technical challenge for the adoption on pixel detectors is the ability to cool the active elements without introducing considerable material in the acceptance of the experiment. This is particularly acute at LHCb where the pixel detectors are in a vacuum where convection is impossible and cooling by contact conduction is required. Furthermore the power output of the inner-most chips, which seem most flux of particles, is calculated to be 1.3W/cm^2. Each module has 12 chips (of varying flux and power dissipation) as well as two timing and control chips. The total module power dissipation is expected to be 28.4W.

The VELO project has identified microchannel cooling as the preferred technology choice combined, for the first time, with bi-phase CO_2 as the circulating coolant. This choice has excellent thermal efficiency, the absence of thermal expansion mismatch with silicon ASICs and sensors and a low contribution to a material budget. The obvious challenge facing this concept is the fabrication of small channels inside a silicon wafer and that the mechanical resistance of the bulk material is sufficient to withstand the high pressures typical of a bi-phase CO_2 refrigeration system. The first demonstration of CO_2 in microchannels is reported in Ref. [2] and further demonstrations of pressure tests and fatigue tests are found in Ref. [3].

The benefit of evaporative CO_2 cooling can be appreciated by inspecting the pressure-enthalpy (P-H) diagram reproduced in Fig. 1 were an simple refrigeration cycle is illustrated. With a base temperature in the chiller of $-35°$C it is anticipated that the microchannel

substrate outlet temperature will be −30°C with a pressure close to 14bar. To maintain sufficient mass flow, a differential pressure of 4bar is needed thus the steady-state operational pressure is under 20bar. The P-H diagram also shows that an liquid-phase ambient temperature start-up requires CO_2 at 60-65 bar.

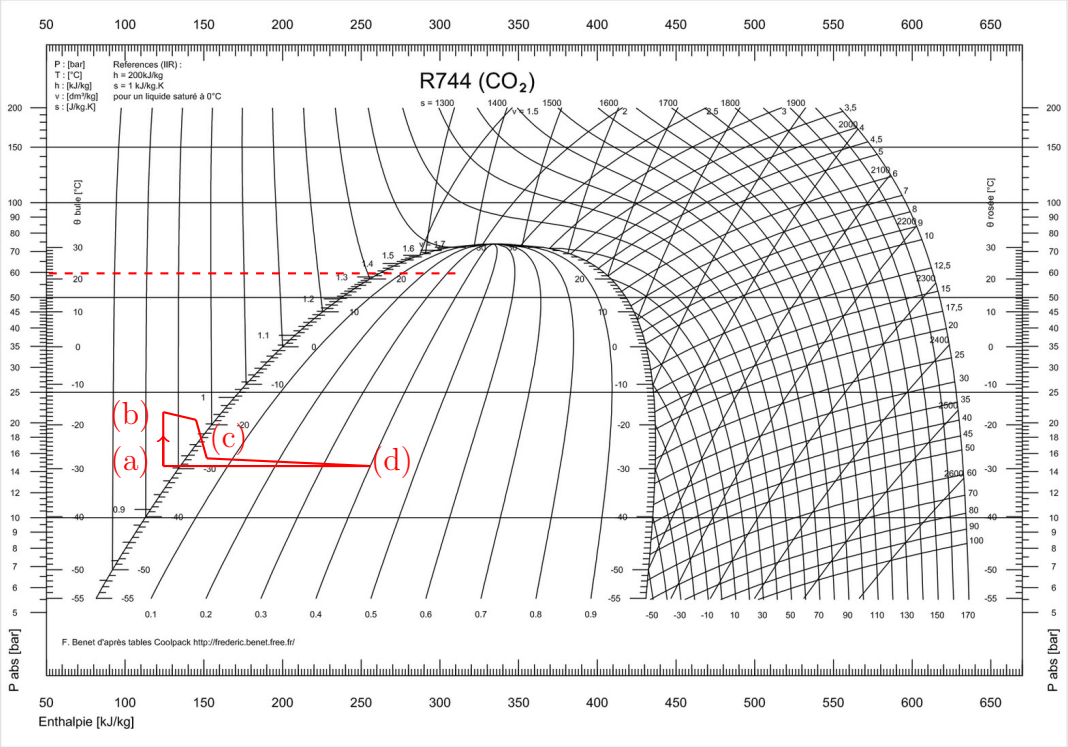

Figure 1: Pressure-enthalpy diagram for CO_2 with an indicative cycle for the LHCb microchannels. The anticipated cooling cycle is driven by the pump which drives the pressure in the system up, about 5bar from (a) to (b). CO_2 is transferred to the detector where a preheater prepares the CO_2 in a state close to boiling. The sudden drop in pressure at the the end of the restrictions (c) triggers boiling. The remaining length microchannels absorb the heat via latent heat of evaporation out to (d) after which the bi-phase fluid flows back to the chiller which returns the CO_2 to a liquid phase at (a). The dashed line highlights boiling at 60 bar and 22°C.

3 Module design

The overall design of the VELO has 52 modules, 26 in each of hemisphere left and right of the beam line. These planar modules are oriented vertically and distributed parallel to the beam axis in such an manner to achieve $> 99\%$ track-finding efficiency over the acceptance of the LHCb spectrometer [4]. All modules should be identical to manufacture, thus the x-y projection of modules placed on the left baseplate in the final assembly is identical to

one the right baseplate, save for a 180° rotation around the beam axis. In compliance with this design axiom, a single of microchannel substrate design is presented here.

Fig. 2 shows the concept of a module in an assembled position and from which one can see the microchannel plate forms the central core of the module on which electrical circuitry, referred to as the *hybrid* and the ASIC chips (not shown) are attached. CO_2 is provided to the microchannel substrate via single connector connected to input and output capillaries. Mechanical support is provided by a carbon fibre *hurdle* consisting of two stiff rods and an adjoining flat section to which the cooling connector is glued.

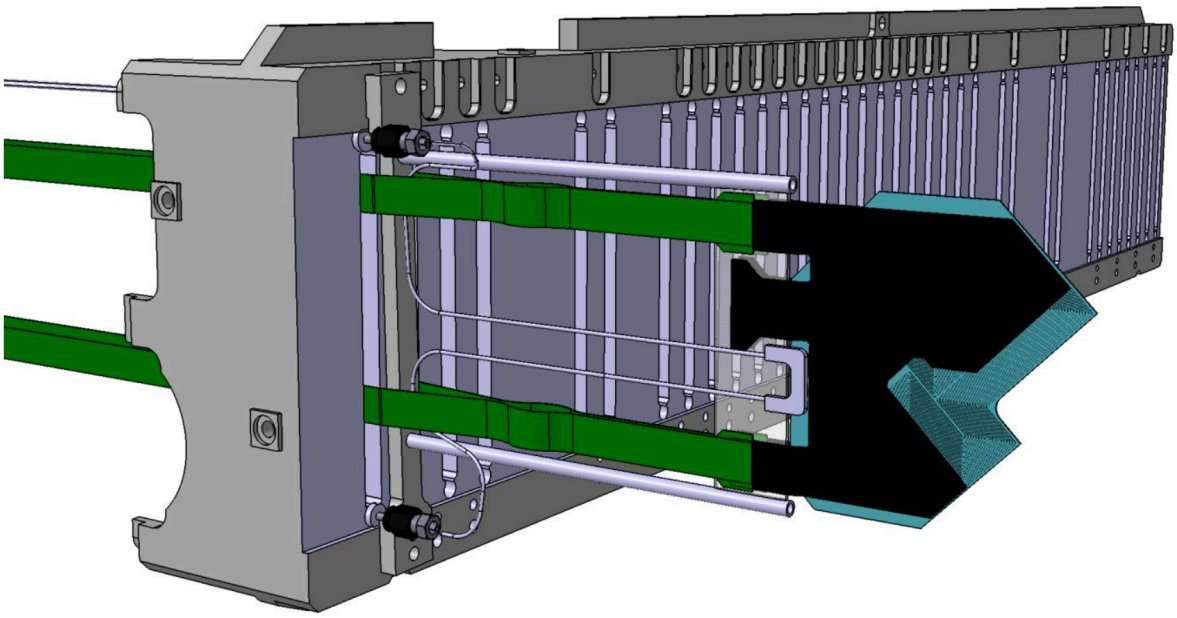

Figure 2: Illustration of a module in its final configuration a the left baseplate. The microchannel substrate (blue) is underneath the electrical circuitry (black) which is connected to long data cables (green). The CO_2 is supplied via a soldered connector (shown here as a U-shape).

Fig. 3 shows the microchannel substrate and hybrid assembly from both sides. The hybrid is designed to that must navigate around the cooling connector to route power, control signals, data and high voltage from the outer edge of the module to the ASIC chips clustered around the beam. The dominant sources of heat per face are:

- six ASIC chips, grouped into two blocks of three, shown in red in the diagram;
- one GBLD on each side, consuming 300mW.

The power dissipation is expected to be 28.4W from the front-end electronics and a negligible output from everything else.

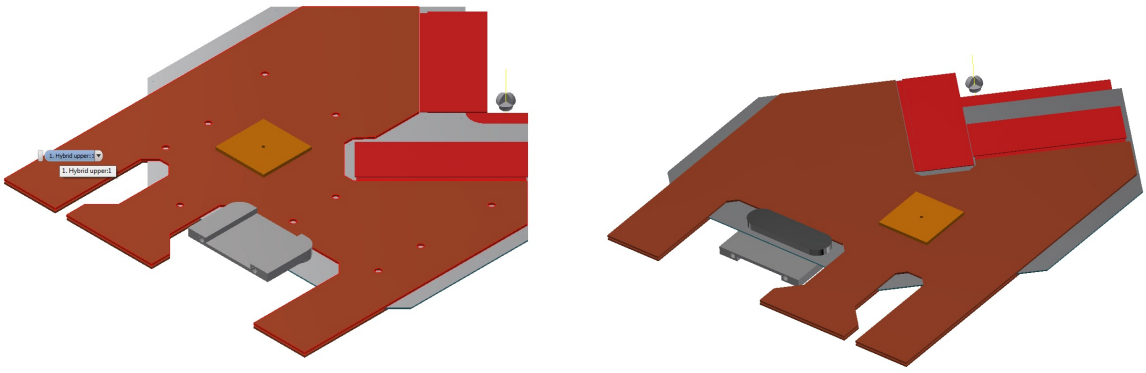

Figure 3: Hybrids and ASICs (red) are glued to both sides of the microchannel substrate as well as a timing and control chip (orange). A glued reinforcement opposite side to the fluidic connector is also shown.

4 Timeline for tendering and production

It is anticipated that call for tender can be launched shortly after the EDR assuming that any iterations to the design are quickly implemented. The invitation to tender will be issued by the grant-holding institute, Manchester, administered by the the LHCb-UK project manager. The tender must the procurement rules of the Official Journal of the European Union in which the notice to tender must appear. At least 35 days must be allowed from dispatch of the notice to receipt of tenders. Requests for additional information must arrive 10 days before the closing date and we are obliged to provide requests for additional information to tenderers at least 6 days before the final date of receipt of tenders. We would require that the quote we valid for at least 90 days from receipt and we envisage placing the order promptly within that window.

It is envisaged that 8-12 preproduction prototypes will be ordered with a six-month delivery time.

5 Suppliers

Conversation has been held with two suppliers, *LETI* in France and *Southampton Nanofabrication Centre* in the UK. Other companies like *CSEM* in Switzerland and *IZM* in Germany are known to be active in this area. Small samples, made from bonding a silicon wafer with embedded microchannels to an second wafer have been acquired from both LETI and Southampton and demonstrated pressure resistance of at least 250 bar and some as high at 700 bar. This should be compared to a burst-disk pressure of in the final system of 100bar.[1]

[1]For operation at $-30°$C the pressure is below 20 bar but values as high as 70 bar are possible during warm start-up at $20°$C.

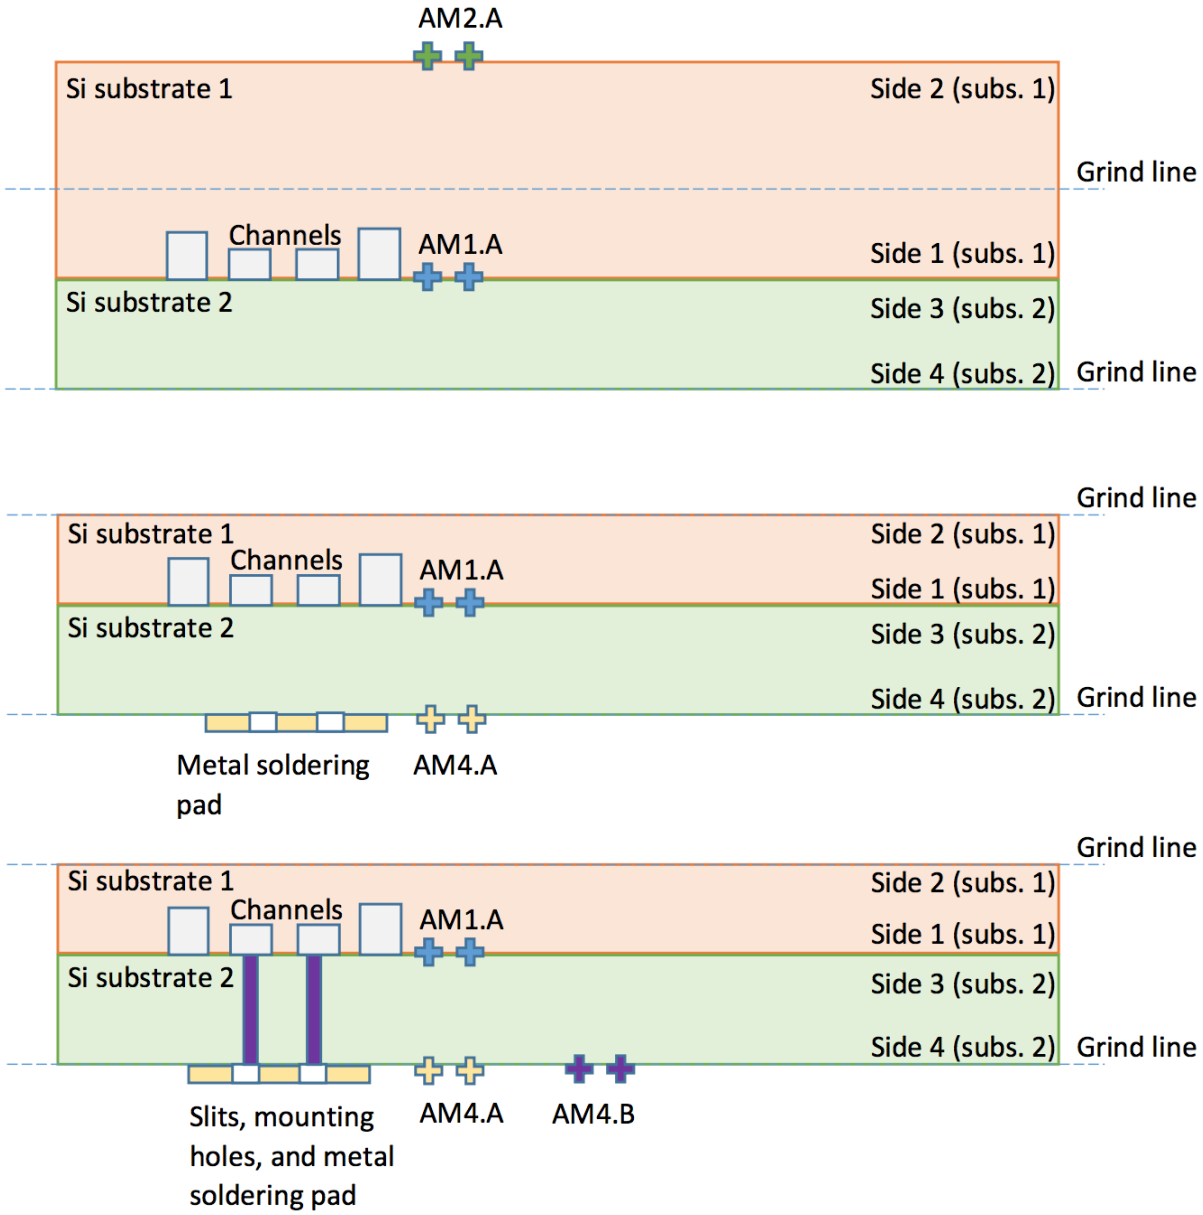

Figure 4: Prospective process flow. (1) Microchannels are etch, in a two stage process, by DRIE before direct bonding to a second wafer. (2) Both faces are ground to a required thickness and metal patterns deposited. (3) Entrance slits are cut by DRIE.

The fabrication is expected to be similar whichever supplier is chosen and it illustrated in Fig. 4.

- The process starts with the *Dry-Reactive-Ion Etching* (DRIE) of microchannels using a technique of alternate application of plasma-etch and deposition of a passivation layer. The process is patented by Bosch and can achieves vertical walls to within

a few degrees. Two depths of channel are desired so this etching needs two stages with a where the shallower channels are masked during the first stage.

- The microchannelled wafer is then direct-bonded to another wafer before being ground down to a suitable thickness. A high-temperature ($\sim 1000^{\circ}$C) annealing step lasting several hours follows which improve the bond strength and can release internal tension in the structure.

- Metal patterns are applied. A large metal pattern is needed for the solder connection of the as well as alignment marks for the subsequent module assembly.

- The entrance holes are DRIE-etched before the device is cut from the wafer.

The accuracy of the fabrication is vital as the entrance holes are on the opposite wafer to the microchannels which are only a few-10s of microns wide. Any alignment marks added during the microchannel etching need to transferred from face-to-face during the grinding step. Nevertheless, prototyping studies (albeit without a grinding step) have shown this to be well controlled, see Fig. 5.

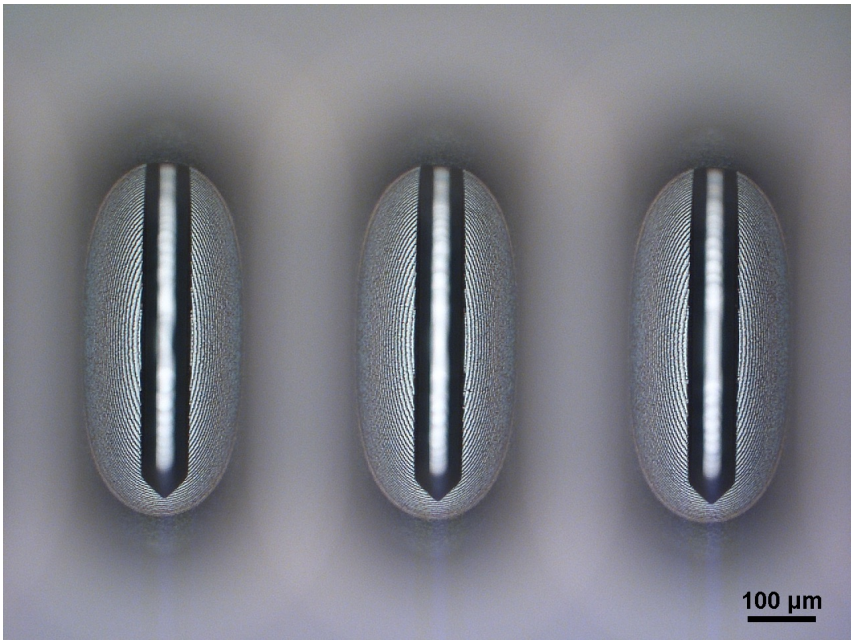

Figure 5: Photograph of a prototype showing the alignment of the DRIE-cut entrance slits (oval, in focus) and the beginning of the microchannel (elongated dark feature, out of focus).

6 The dimensions of the substate

The thickness of the substrate is governed by a compromise of three factors:

- the wish to minimise the thickness and thus the material budget. for comparison, the thickness of the ASIC-sensor assembly is 400μm.
- the necessity to be sufficiently thick to withstand the internal pressure of boiling CO_2. Typical operation pressures when cold will be 14-20 bar but during a warm start-up pressures around 70bar should be anticipated. The cooling circuit will be protected by burst disks, likely 100bar.
- that the mechanical strength of the large substrate is sufficient to withstand shipping, assembly and installation

Studies of the rupture pressure of silicon versus applied fluidic pressure, reported in [3] demonstrate the the inherence pressure resistance of 140μm of silicon, for a 200μm wide channel gives a safety factor of at least four and is thus not considered a critical factor. Thus the baseline design is 400μm as there is little motivation to be thinner than the ASIC-sensor assembly and this thickness gives a 140μm cap on either side of 120μm microchannels. The mechanical integrity is considered further in Sec. 9.

The angular shape of the device is dictated by the ASICs rotation by 45° (x-y plane) with respect to the module support. This design feature is desired as it makes the approach to the RF-foil during installation safer and the closing procedure during physics running more symmetric.

The overall envelope is shown in Fig. 6. The angular shape is dictated by financial need to fit two devices onto one 8-inch wafer leaving a 5mm gutter between the two devices and a 5mm gutter around the edge of the wafer.

A critical design feature is the the internal corner which is of large radius of curvature to minimise it being a concentration point of mechanical stress. A radius of curvature of 5mm matches the retraction of the inner substrate edge with respect to the inner ASIC edge.

The presence of an internal corner presents a difficulty for the dicing the wafer after processing. Standard saw-dicing, which performs complete cuts across a wafer, is physically impossible. Cutting with a pulsed laser has been used repeatedly for mechanical prototypes but it leaves a jagged edge, prone to splinter. The preferred solution is to use DRIE processing to etch away all but one straight edge of the outer envelope of the device. Such a step would require the temporary attachment of a third wafer to sit under the bonded assembly to protect the equipment from the plasma once full penetration is achieved.

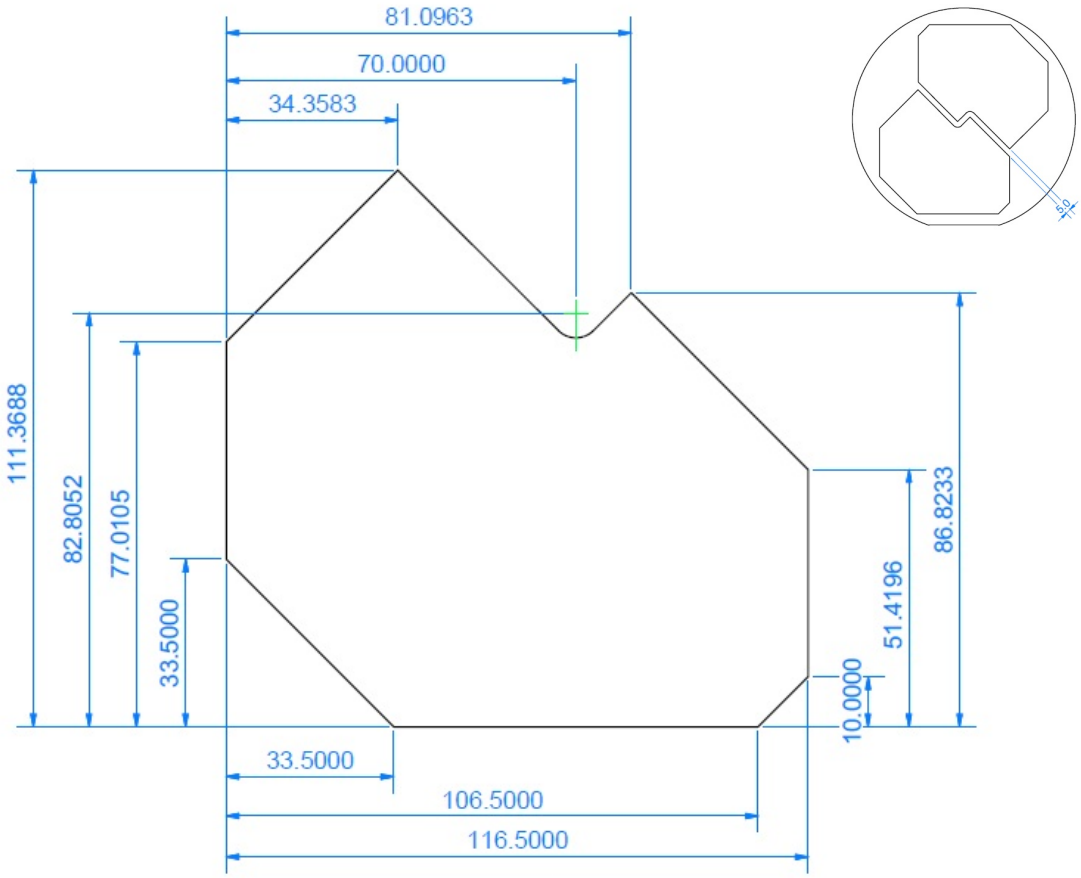

Figure 6: The dimensions of the proposed device. Inset: two devices imaged on an 8-inch wafer.

7 Channel layout

The spacing of microchannels is optimised by thermal simulation using ANSYS. An heat souce producing 1 W/cm^2 is simulated on one side of the substrate. The embedded microchannel dimensions are 200μm×120μm and placed at 700μm pitch (500μm spacing). At this distance the ΔT from heat source to coolant is under 1°C, see Fig. 7.

A race-track circuit is laid out to pass under all ASIC chips and the GBLD chip, see Fig. 8. Nineteen parallel tracks at 700μm pitch provide sufficient surface coverage. The spacing changes from 700μm pitch in two places. (1) Under the GBTx chip, the spacing is increased to 990μm to pass under all of the wide chip. (2) At the entrance and exit the spacing decreases to 450μm to reduce the size of the fluidic connector.

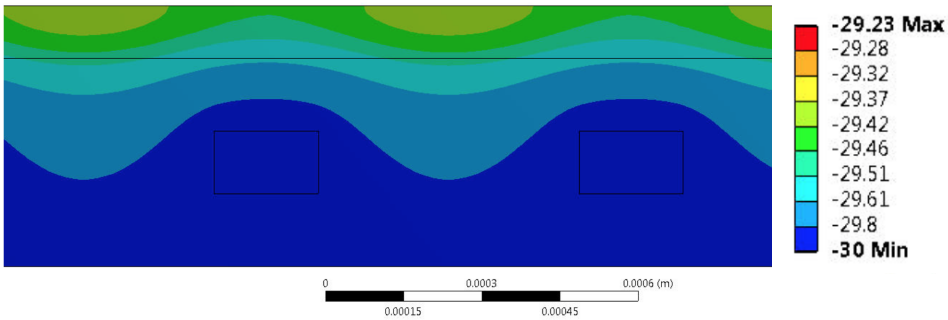

Figure 7: ANSYS simulation of 200μm×120μm microchannels at 700μm pitch with a heat source on one face. Colour scale is in temperature in degrees.

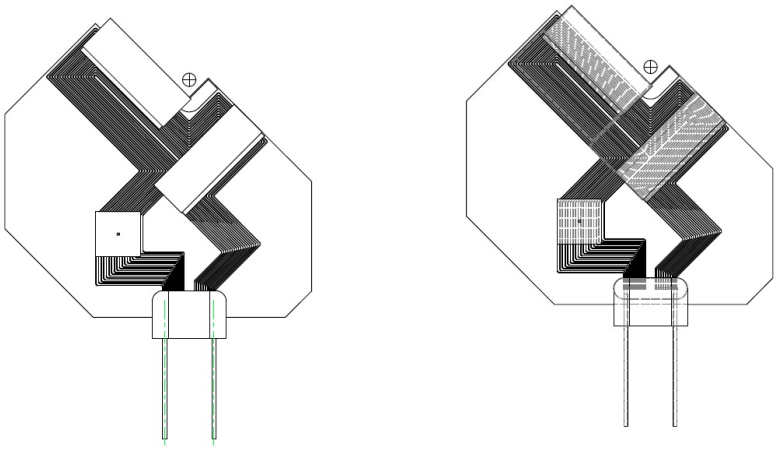

Figure 8: The microchannel route under the various chips.

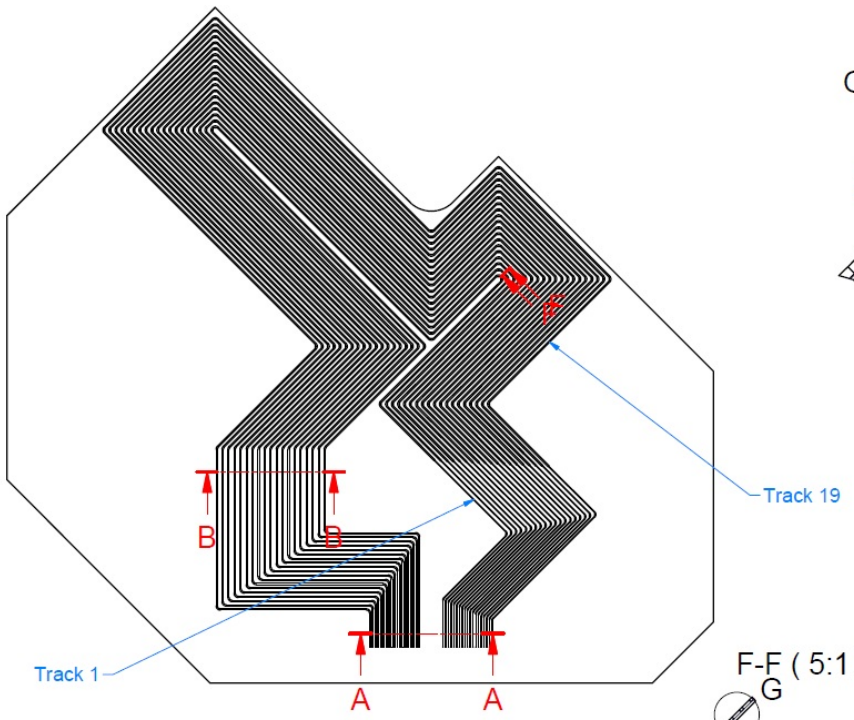

G (20 : 1)

C (20 : 1)

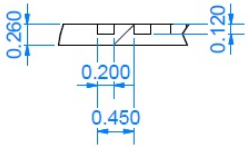

D (20 : 1)

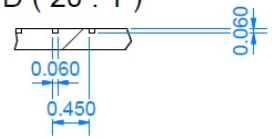

E (20 : 1)

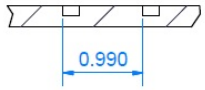

F-F (5:1)

A-A (5:1)

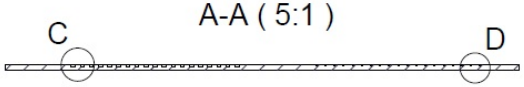

B-B (5 : 1)

UNIVERSITY OF OXFORD DEPARTMENT OF PHYSICS KEBLE ROAD OXFORD, OX1 3RH TEL. +44 (0)1865 283196 FAX. +44 (0)1865 283137	PROJECT:	TITLE: Microchannel layout
CREATOR: brock	GENERAL TOLERANCES: X ± 0.5 X.X ± 0.1 X.XX ± 0.05 X.XXX ± 0.005 ALL DIMENSIONS ARE IN MM UNLESS OTHERWISE STATED DO NOT SCALE	IDENTIFICATION NUMBER:
TECHNICAL REFERENCE:		SCALE: / DATE OF ISSUE: 19/10/15 / MATERIAL:
CHECKED BY:		DOCUMENT TYPE: / NUMBER OFF:
APPROVED BY:	3rd ANGLE PROJECTION / SIZE: A3	DOCUMENT STATUS: / COMPONENT WEIGHT: N/A
		JOB NUMBER: / FINISH:
THIS DRAWING MAY NOT BE USED FOR COMMERCIAL USE		SHEET 1 OF 1

The microchannels run along the inner rim of the substrate with a 1mm distance to the edge. As this is twice the distance from one microchannel to the next, and the inherent strength of the silicon bulk is demonstrated to be strong enough to resist the likely maximum pressure (by an order of magnitude), this distance to the edge seems conservative.

Microchannels of reduced cross-section are used at the beginning of the circuit to even the flow and trigger boiling at the orifice into the main microchannel. Assuming laminar flow and a simple tube model, pressure drop is proportion to length and inversely proportional to the square of the cross-sectional area. Thus for 200μm$\times120\mu$m main channels and 60μm$\times60\mu$m restrictions that are just 10% of total length, the ΔP in the restrictions will be a factor five higher than in the main microchannel. Thus, as long as the restrictions are of equal length, even flow across the nineteen parallel channels can be ensured. Furthermore, as bends in fluidic transfer is a source of pressure drop, the same bend-radius is used for every microchannel at each corner.

The circulating CO_2 enters and leaves the device though a series of small holes etched in the top wafer of the substrate. The holes are oval to avoid internal corners wherever possible. The width of each hole matches that of the microchannel it connects to. This means the entry holes are 60μm wide and the exit holes 200μm wide,[2] see Fig. 9.

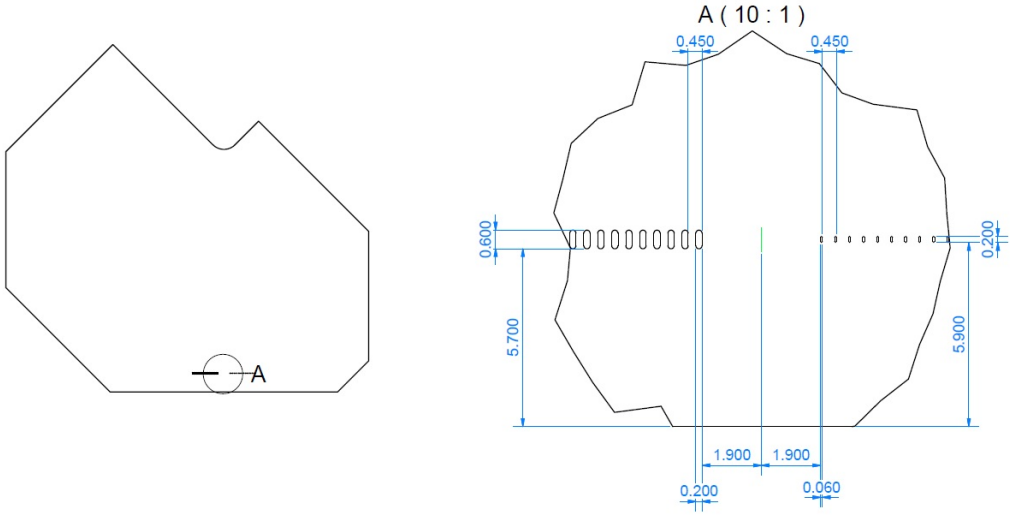

Figure 9: Focus on the entry and exit holes.

[2]The entry hole may need to be enlarged, without serious consequence if the simultaneous DRIE fabrication of these holes requires them to be of more similar lateral dimensions.

Background: demonstration of cooling capacity and realistic ΔP

A full-size prototype of a microchannel network has been fabricated in a silicon-pyrex (as opposed to silicon-silicon) and tested with circulating CO_2 and a set of heaters designed to mimic the head load of a final module. The mass-flow through the circuit and the outlet temperature are measured as a function of applied power. It is seen that over 60W is dissipated before the temperature of the outlet fluid becomes unstable, indicating a risk of dry-out. A measurement of mass flow through the device, and knowledge of latent heat, the maximal theoretical cooling power can be calculated and thus the vapour quality of the coolant is inferred, see Fig. 10. The conclusion is that 60W of heat, which is over twice the required cooling capacity, can be sunk with a mass-flow of roughly 0.3 g/s at $-30°$C. Fig. 11 (taken from Ref. [5]) show a study for this circuit of mass flow versus the applied differential pressure with various heat loads applied. A standard operating point of ΔP=4 bar is anticipated.

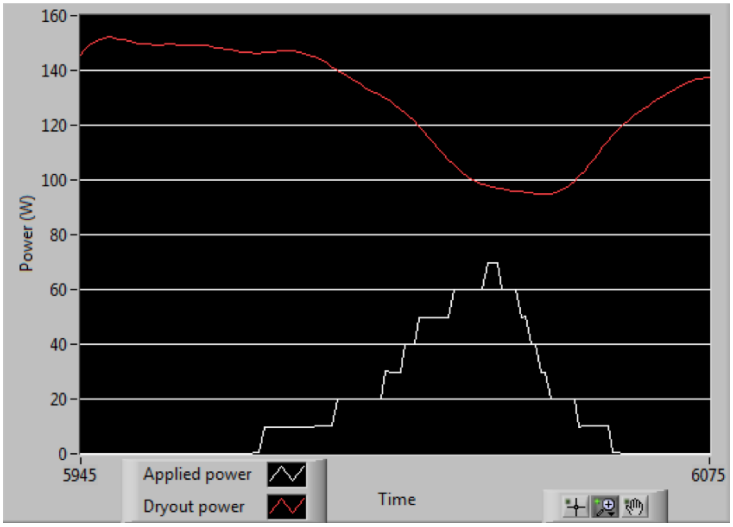

Figure 10: A Labview plot of showing, as a function of time, the power dissipated into a prototype microchannel substrate during stable running of the cooling circuitry (white graph). It also shows the maximum cooling power possible calculated from a real-time measurement of the mass-flow and knowledge of the latent heat of CO_2. A maximum vapour quality of $\sim 80\%$ is inferred.

8 Metallisation

The main need of metallisation is to provide a wettable contact for the soldering the fluidic connector. A 3.5mm wide region is defined around two apertures that contain the 19 entry and exit holes of the microchannel substrate. The two apertures are separated by 3.5mm and there is 100μm clearance from the inner edge of the aperture to the edge of the exit holes. The clearance is larger for the entrance holes because the restrictions are

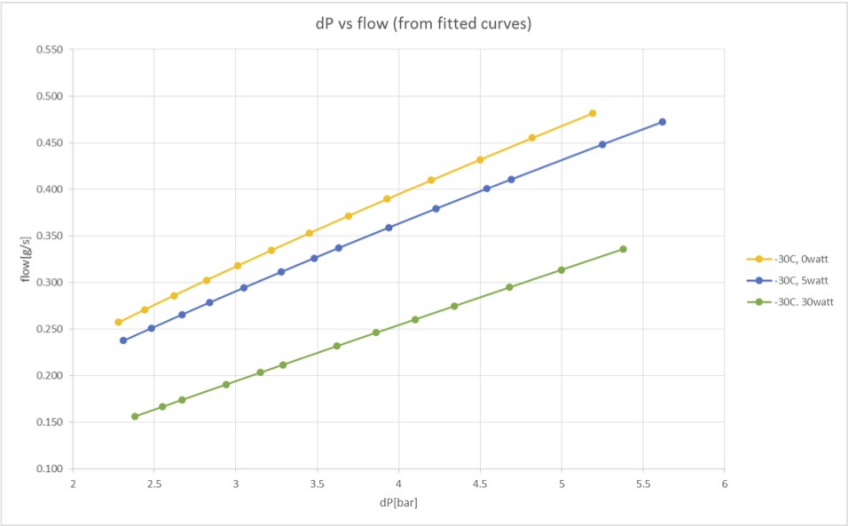

Figure 11: Plots of mass flow versus ΔP for various heat loads.

smaller than the main microchannels see Fig. 12. With 600μm exit holes, the total slit width is 800μm. Each of the two non-metalised areas is 8.6mm long so the total force for 20 bar pressure is 28N.

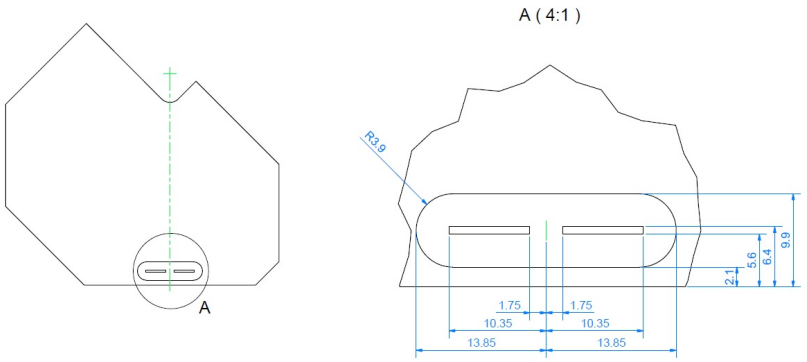

Figure 12: Dimensioned drawing of the connector metallisation.

The metallisation is a stack of three metals:

1. The first layer is 200nm of titanium layer, used for its adhesive reactiveness to the polished silicon surface.

2. The second layer is 500nm of nickel. This is the active material in the later soldering process.

3. The top layer is a protective 500nm gold plating to protect the the nickel underneath from oxidising.

Use of gold for this purpose is common though the thickness suggested here is not. The reasoning is that the soldered joint needs to be as void-less as possible and any weak spots in a thinner gold-plating carries a risk. With *Southampton* a thicker 1000nm nickel layer was attempted but was found not to be a stable deposition and the metal peeled off.[3] However, in all soldering trials to date, there has been no evidence of systematics failure to due to insufficient metallisation.

Alignment marks

It is expected that alignment marks (for the benefit of the module assembly) are deposited during the same metal deposition as above. The location, size and nature of the alignment marks is shown in Fig. 13. The position of the alignment marks is chosen to avoid being covered by either the hybrid or the ASICs during the subsequent module assembly. The style is designed to give an precise point, as defined by the centre of the circular feature, and give an unambiguous orientation, as specified by the 'L'-shape.

The fabrication of such alignment marks with the relatively thick metal deposition described above has been tested with a supplier. The successful result is photographed in Fig. 14. The figure also shows the clear visibility of the same mark, but made with plasma etching which would seem a totally viable alternative should that bring cost advantage.

Solder guides

Extra markers are proposed around the footprint of the fluidic connector to allow visual verification of the precise alignment of the solder connection. See Fig. 15.

[3] Were thicker layers considered desirable, electrolytic nickel deposition would be required.

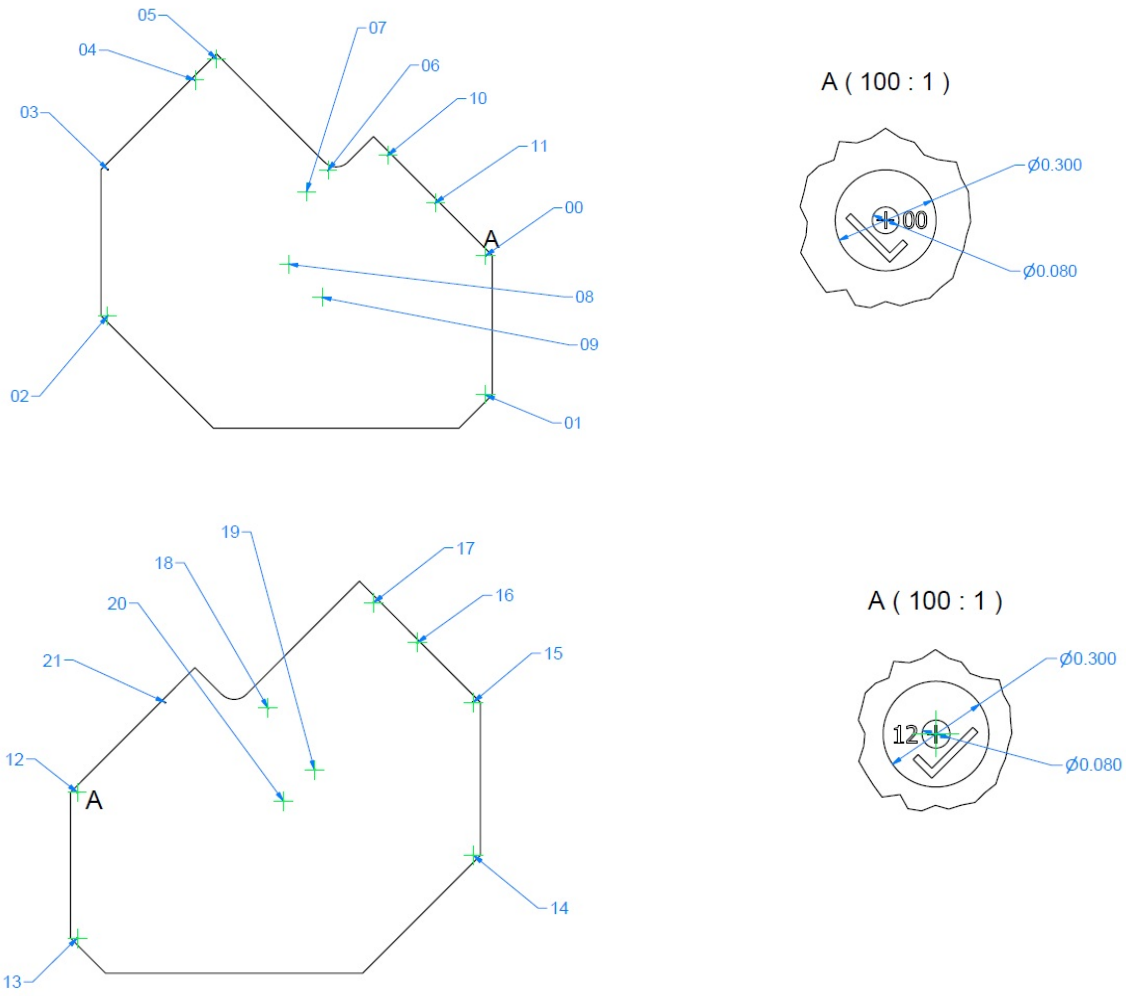

Figure 13: Top: the placement and style of the alignment marks on the connector face. Bottom: the placement and style of the alignment marks on the reverse face.

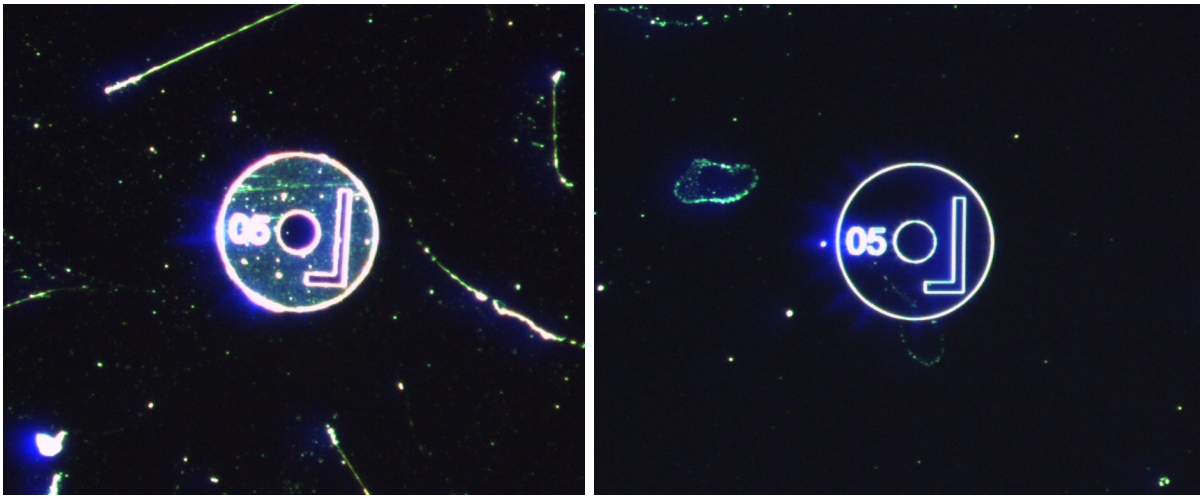

Figure 14: Left: a microscope image of an alignment mark fabricated from deposited metal. Right: the same alignment mark made by plasma etching.

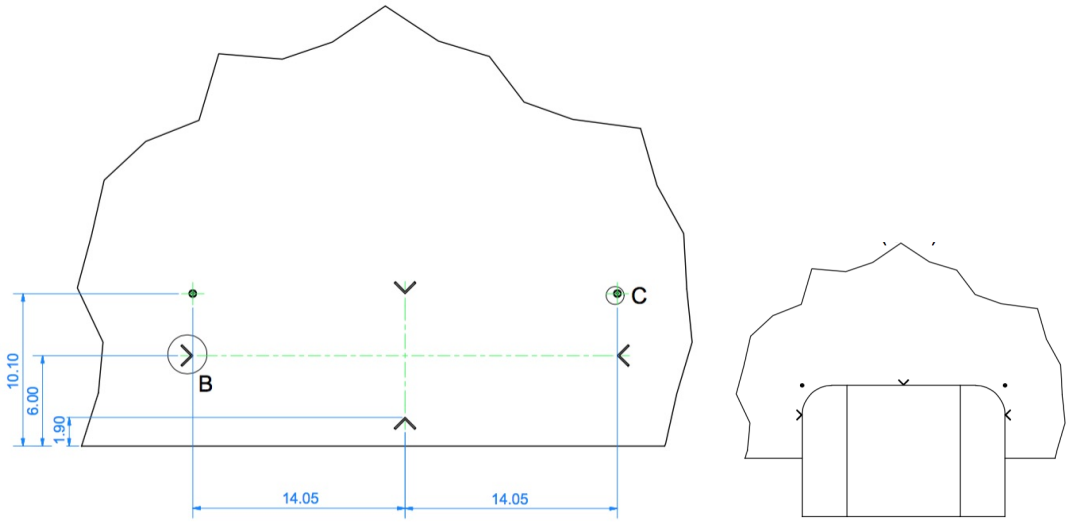

Figure 15: Extra markers to for visual verification of connector alignment.

9 Risk factors

The wider risk of the use of microchannel cooling in an LHC experiment is beyond the scope of this review document. Here, focus is placed on aspects of the design choices that if wrong, will need a redesign, with all the delay and cost that it would require. It is important to note that nearly all aspects of the design are locked in the early stages of fabrication when the masks are made for the etching and metallisation. Thus it is expected that the masks paid for as part of the pre-production series will be reused during the main production. With this in mind, the following risks are identified:

Robustness

The baseline thickness is 400μm, mainly from arguments of material budget.

For the R&D around 30 dummy devices have been cut to the desired angular shape from mechanical-grade silicon wafers using a pulsed laser. Though an accurate log has not been kept, anecdotally around half crack either in the cutting or in transport. Subsequent handling has be very careful and the dummy devices are seen to have very little tolerance to stresses out of the plane. A simple test, see Fig. 16 demonstrates that it is not a inherent problem with a 400μm wafer but perhaps the cutting process cause some warping in the crystal lattice rendering the cut device more fragile.

If this is the case, the production devices may be considerably more robust because (it is assumed) they pass through a long duration, high-temperature annealing step which should release such stress. Nevertheless, the intended thickness of the substrate is a factor that should be considered a risk.

DRIE dicing

Obtaining the final device by DRIE dicing has only been discussed with a supplier but not demonstrated. If during the tendering process it is deemed impossible or prohibitively expensive, laser cutting is the only option. Laser cutting is cleanest for a for a thinner substrate.

Creep

The soldered joint must resist an internal pressure and thus there is, in principle, a risk of long-term plastic deformation. Given the operational pressure is low ($<$ 20bar) and the connection will be held at a cold temperature (-30°C) it is hoped that this risk is low.[4] A simple longevity test is underway, see Fig. 16 where a 108N (11kg) force is applied to a solder connector at room temperature (so 50∘ higher than the operating temperature) for over a year, with not evidence of creep. The internal force

[4]Creep is very temperature dependant. It is thought that it vanishes for temperatures below half of the absolute melting temperature (455/2=227 K= -45°C)

of 20 bar of the non-metallised (thus assumed non-soldered) area inside the connector is 28N.

A possible safeguard is to introduce a screwed reinforcement from the cooling connector to a second plate on the reverse face of the substrate. This requires holes to be made in the substrate, presumably during the proposed DRIE dicing step. It has been observed that holes made in mechanical dummies with laser cutting tend are a failure mechanism (meaning cracks often propagate to/from the holes). The preference is therefore not to include holes in design, but this carries a risk of having no possibility of direct reinforcement.

Time

In targeting installation on the new VeLo by early 2020, the project planning calls for pre-production module testing in the summer of 2016. This module, which is to be equipped with ASICS and irradiated will be tested before launching the full complement of substrates in 2016 Q4. The rolling production of qualified microchannel substrates to the module construction run throughout 2017. All substrates are expected from the supplier by Q2 or 2017 with the assembly and quality control at Oxford complete by the end of that year. Thus there is a risk associated with any delay to the procurement of the pre-production samples, though urgent studies can advance during the tendering process.

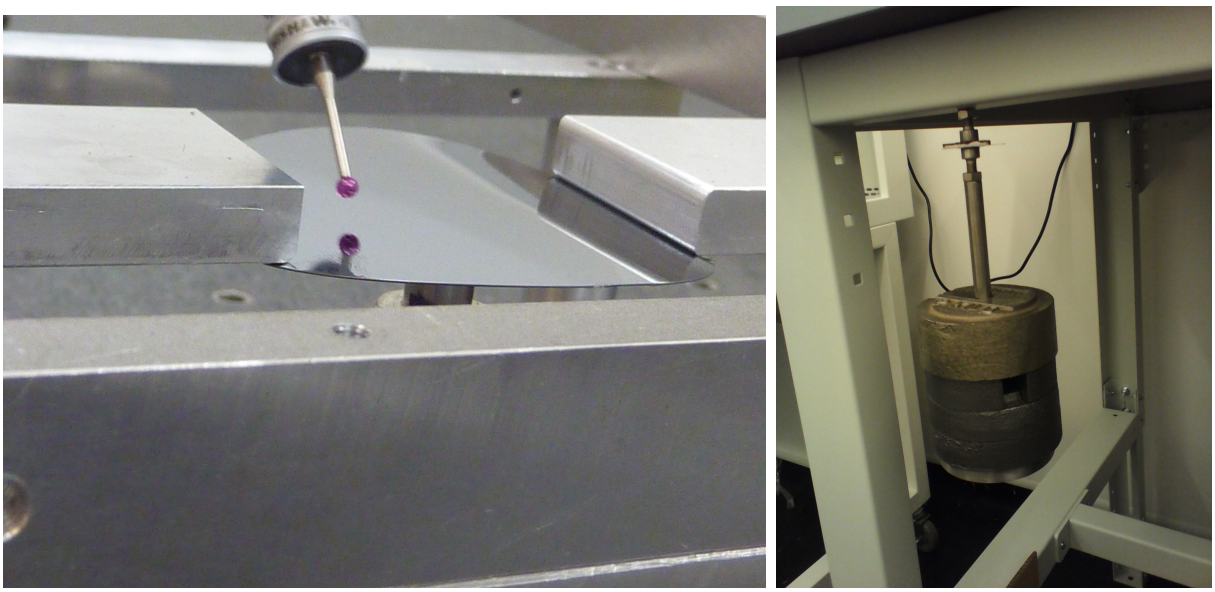

Figure 16: Left: A test of the intrinsic fragility of a 400μm wafer. Right: A longevity test of a solder joint.

References

[1] LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021. LHCb-TDR-013.

[2] A. Nomerotski *et al.*, *Evaporative co_2 cooling using microchannels etched in silicon for the future lhcb vertex detector*, Journal of Instrumentation **8** (2013), no. 04 P04004.

[3] O. A. de Aguiar Francisco *et al.*, *Evaporative co_2 microchannel cooling for the lhcb velo pixel upgrade*, Journal of Instrumentation **10** (2015), no. 05 C05014.

[4] T. Bird, M. Gersabeck, T. Head, and G. Lafferty, *VELO upgrade studies and layout optimisations*, Tech. Rep. LHCb-INT-2014-044. CERN-LHCb-INT-2014-044, CERN, Geneva, Nov, 2014.

[5] J. Buytaert, *Cooling performance of the Snake2 design*, Tech. Rep. Indico presentation.