Thermal performance of BeO and TPG modules

Manchester SCT group

1 Introduction

We have built two dummy end-tapped forward modules, one is based on an old design with a BeO spine, the other has a TPG spine and a seperate hybrid. Each module was mounted on realistic cooling block(s) in an enclosed box, the expected heat loads were applied and the temperature was mapped.

We have also written a simulation program which, after tuning some parameters, is able to account for the temperature distributions which we measured. Our simulation is then used to predict the thermal run-away behaviour that these modules would have if they were built from detectors which had the expected temperature dependent power density.

We then propose some safety factors which should be included to allow for variation in glue and grease layer thickness from module to module. With these safety factors included, the runaway studies are repeated with modules of varying base-plate thickness to find a module design which will just meet the ATLAS requirements.

2 The modules

The BeO module has the design shown in the TDR Figure 11-55. It has a BeO hybrid with a narrow neck which is bridged by the fan-ins. A BeO spine runs down the centre of the module between the wafers. The spine has arms to support the wafers at the wafer gap and at the far end of the module. For practical reasons the spine had to be made from two seperate pieces.

The TPG module is based on more recent ideas shown in drawing NP27-03-103 from Liverpool. We have used the materials and dimensions shown in that drawing except for the BeO which was not available at short notice. Instead we used aluminium pieces 0.78 mm thick to replace the 0.5 mm BeO hybrid and the two BeO washers which make contact with the cooling block.

Both modules were equipped with resistor packs to supply a heat load at the positions of the chips. The wafer power was simulated with a thin heater wire glued in a zig-zag pattern onto the surface of the silicon. Table 1 shows the materials used in construction of these modules. Thermocouples were used to monitor the temperature at nine points on each module at the positions marked in Figure 1.

Component	Material	Thickness (mm)	Conductivity (W/m/K)	Used in
Wafers	Silicon	0.300	126.	both
Cooling blocks	Aluminium	various	220.	both
Glue wafer to spine	BN glue	0.080	0.7	both
Grease	Grease	various	0.9	both
Hybrid and spine	BeO	0.500	280.	BeO
TPG spine	TPG	0.500	1700. , 25.	TPG
TPG coating	Polyimide	0.020	0.2	TPG
Hybrid and washers	Aluminium	0.780	220.	TPG
Glue washers to hybrid and TPG	BN glue	0.010	0.7	TPG

Table 1: Materials used in the thermal modules and in simulation.

3 The cooling

The main cooling contact at the hybrid end of the module has the seven sided shape with an area of about 200 mm^2 which is shown in recent drawings. The module is clamped tightly onto the block by a screw which passes through a hole in the module and into a tapped hole in the block. A thin layer of heat-sink grease is

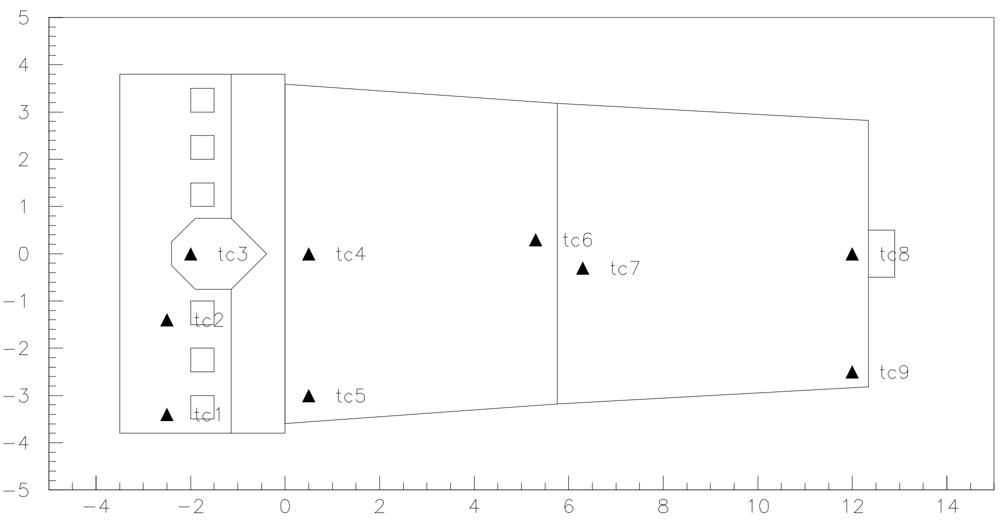

Figure 1: The positions of the thermocouples.

spread on the block to make a good thermal contact. At the far end of the TPG module there was the option of a second cooling contact onto a rectangular block of 50 mm^2. Both blocks were cooled by chilled water running through a channel drilled directly through the block. We found that when a typical power of 5 W is being dissipated through the main cooling block its temperature rises to about 4 degrees above the fluid temperature.

Previous measurements have shown that convection can have a significant effect, especially if the module is poorly cooled by conduction. So in order to minimise the effect of convection we aimed to run in a situation where the temperature of the air in the box is close to laboratory temperature and is also close to the mean surface temperature of the module. The cooling fluid temperature was adjusted to produce this situation.

4 Results

All of our measurements are shown in Table 2 along with the simulation results. The difference between the two sets of BeO results shows that convection is having a large effect on this module. Both modules show an encouragingly low temperature difference of 1.0 to 1.8 degrees between the cooling block and thermocouple 3 which is on the hybrid above the block, when about 5 W is flowing through this junction. In order to analyse the results further it is necessary to look at the simulation.

5 The thermal simulation

We have written a fast thermal simulation program which is specially adapted to the linear shape of the end-tap module. It describes two dimensions in detail. The first axis, X, is a curved axis which runs parallel to the main heat flows in the module. In the hybrid it runs through the 3 chips and curves around to reach the centre of the cooling block, then it runs along the centre-line of the wafers. The other important dimension, Y, is perpendicular to the plane of the module. The module is divided into small cells in the XY plane and heat generation and flow between the cells is followed through short time steps, until equilibrium is reached. The third dimension is described simply as an effective width, which governs the area of contact between cells. It is assumed that temperature is independent of the third dimension.

All the different layers of material are included in detail in the simulation. Non-isotropic thermal conductivity is allowed, so that the TPG can be described properly. A simple model of convection is included, in which all

Module	BeO		BeO		TPG		TPG		TPG		TPG	
Cooling points	1		1		1		1		2		2	
	meas.	sim.	meas.	sim.	meas.	sim.	meas.	sim.	meas.	sim.	meas.	sim.
Hybrid power (W)	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5
Wafer power (W)	1.0	1.0	1.0	1.0	1.0	1.0	2.0	2.0	1.0	1.0	2.0	2.0
Power to air (W)		0.18		0.57		0.10		0.31		-0.29		-0.02
Air Temp. (C)	8.5	8.5	4.7	4.7	3.8	3.8	4.7	4.7	4.1	4.1	3.7	3.7
Hyb. block T. (C)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
End block T. (C)									-2.3	-2.3	-2.3	-2.3
tc 1 (C)	11.8	12.8	11.5	12.4	10.7	10.7	10.6	10.8	10.7	10.7	10.6	10.7
tc 2 (C)	8.0	8.4	7.5	8.1	6.6	6.8	6.7	6.9	6.7	6.8	6.7	6.8
tc 3 (C)	1.2	1.1	1.1	1.0	1.5	1.7	1.6	1.8	1.6	1.7	1.6	1.7
tc 4 (C)	5.0	4.0	3.7	3.0	3.1	2.7	4.2	4.5	1.7	1.4	2.3	2.4
tc 5 (C)					3.2		4.2		1.7		2.3	
tc 6 (C)					3.6	3.6	5.6	6.1	1.6	1.6	2.4	2.9
tc 7 (C)	10.2	10.4	7.2	7.6	3.6	3.7	5.8	6.4	1.5	1.5	2.5	2.8
tc 8 (C)	11.4	11.4	8.1	8.4	4.1	4.0	6.5	6.7	1.0	0.8	1.7	1.9
tc 9 (C)					4.1		6.1		1.4		2.1	

Table 2: Results of measurements and simulation. All temperatures are shown relative to the main cooling block.

the air around the module is at a uniform temperature and it exchanges heat with the module surface according to a convection coefficient of 11 $\mathrm{Wm^{-1}K^{-1}}$.

The only free parameters in the simulation were the thickness of the grease layers and one of the glue layers. We were aiming for a thin glue layer between the washer and the rest of the module, but were not able to measure it accurately after the module was built so we were guided by the temperature measurements to set this layer to the thinnest realistic value of 10 microns. The grease layers were tuned to match the thermal results, which gave 50 microns for the BeO module, and for the TPG module 25 microns at the main block and 40 microns at the far end block.

After tuning these parameters the simulation takes as input the hybrid power, wafer power, cooling block temperatures and air temperature. It gives as output the temperatures shown in Table 2 and the net power which flows from the module to the air. The agreement of the simulation with reality is of order $\pm 10\%$ and $\pm$ 0.5 °C.

6 Extrapolation

The simulation is close enough to reality that we can reasonably use it to calculate what happens if silicon wafer power is a function of temperature. A couple of other changes were made, in order to make the simulation more representative of a module in ATLAS. Firstly we set the temperauture of both cooling block to -10 °C. In the laboratory the far-end cooling block was colder than the main block, but material could be saved in the far-end block which would bring its temperature closer to that of the main block. Secondly we allowed the air temperature to vary so that it follows the mean surface temperature of the module, so convection can not remove heat but can only re-distribute it within the module. Having made these changes, the simulation was run with various power densities and gave the typical thermal runaway plots of Figure 2. It can be seen that with two point cooling the module runs away at a 0 °C power density above 400 μW mm^{-2} and with one point cooling it runs away above 200 μW mm^{-2}. The expected power density after 10 years is 120 μW mm^{-2} so this module with two point cooling is massively over-designed and even with one point cooling it has a safety factor of almost two.

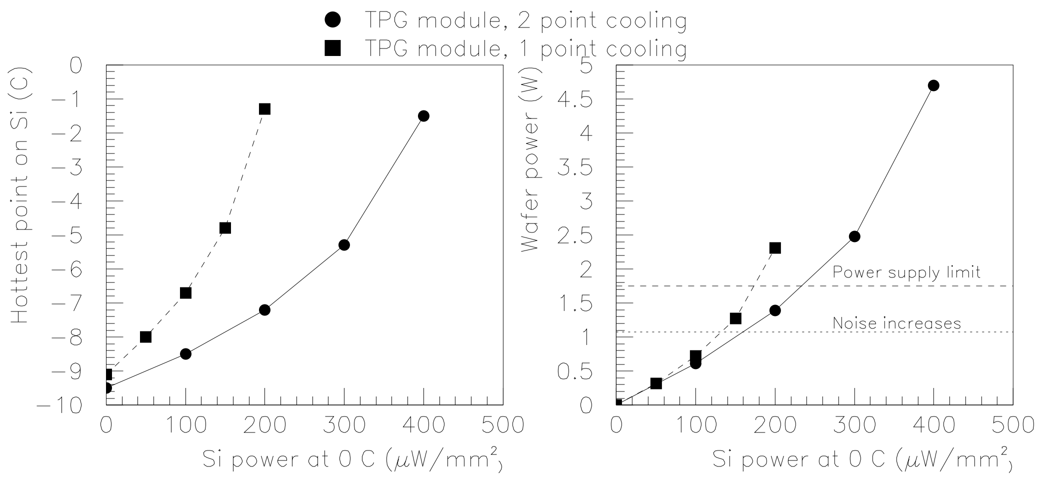

Figure 2: The results of applying temperature-dependent wafer power to the thermal model which fitted our experimental data.

7 Safety factors

Rather than applying a single large safety factor to the power density and then designing a module which will in principle cope with this power, we prefer to derive a safety factor as follows.

The nominal power density of 120 μW mm^{-2} after 10 years operation already includes some safety because it is derived from the worst case among several irradiated detectors. So do we need to add anything to the power density to take account of detector-to-detector variations? *I had thought that a warm-up accident would increase the power density but in fact it does not. It would increase the depletion voltage significantly but the leakage current is simply proportional to radiation dose so it would not change. Thus a warm-up accident would cause the detectors to hit their depletion voltage limit of 350 V at a lower power density than 120 μW mm^{-2}.*

This leads me to argue in the opposite direction: suppose that the SCT gets a bigger dose than expected, either because of a higher flux or because we go on running beyond 10 years. And further suppose that we can avoid some of the annual maintenance warm-ups. This would lead to detectors which could still be depleted but which would have a power density higher than 120 μW mm^{-2}, so we should calculate some sort of "optimism factor" to allow for this scenario. For now we pull a number out of the air; 30%.

Another safety factor should account for variations in the thermal performance of modules. We can expect that the solid components of a module will have highly consistent dimensions and conductivity; the variation comes from the glue and grease layers. We need a careful experimental evaluation of the variability of these layers but we already have evidence that it is difficult to make them very consistent. Section XX contains some warning information about the TPG-glue joint. The different grease thicknesses (25 to 50 microns) which are reported in section 5 are pessimistic because we took no special care in making this grease joint. We believe that setting a specification that the thermal resistance of the glue and grease layers should be within 30% of their nominal values is probably achieveable.

There will be variations in the cooling block temperatures due to imperfections of the cooling system. So the specification of the system should be that it will keep all blocks below some limit (here taken to be −10 C) when the power from each block is 7 W.

8 Optimisation

Based on the numbers above we would say that the optimum module should be dissipating just under 1.75 W (ie the power supply limit) in the wafers when the silicon power at 0 C is 160 μW mm^{-2}, when both the glue and grease joints have 30% worse than nominal impedance and when the cooling blocks are at -10 C. We have simulated this set of conditions and varied the cross sectional area of the spine in order to find the "optimum" module for one-point and two-point cooling. The results of this pocedure are shown in Figure 3.

It turns out that the TPG module with two point cooling needs very little TPG in the spine, in fact according to this simulation it would work with a TPG bar of only 5×0.5 mm^2 but that is not realistic mechanically, so we show results for a module with a central bar of 10×0.5 mm^2.

On the other hand, although the single-point-cooled TPG module that we actually built would meet the power criteria above, once we have increased the impedance of both glue and grease joints by 30% it no longer does so. So we increased the width of the TPG bar to 60 mm, covering the whole module, to get the result shown in Figure 3 which dissipates 1.6 W at 160 μW mm^{-2}.

We also looked into the possibility of a BeO module with two point cooling. We assumed that it could be given a 20×5 mm^2 cooling contact area at the far end. The figure shows that in this case it meets the power criteria.

Both the single-point TPG and the two-point BeO module described here are rather marginal but both would be brought well within the power limit if the cooling block temperature could be reduced from -10 to -12 degrees.

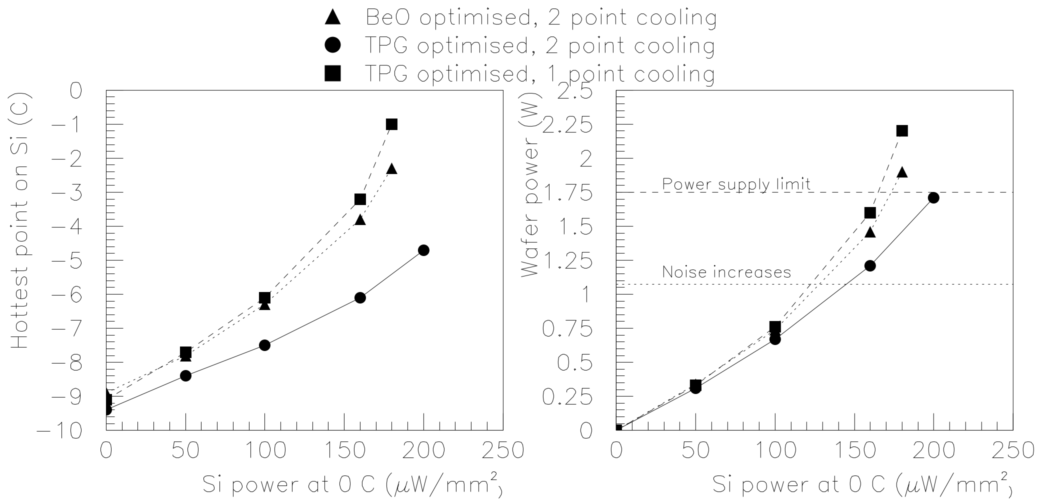

Figure 3: Thermal runaway curves of three alternative module designs.

9 Further tests of the TPG-glue joint

We made a number of the test structures shown in Figure 5 in order to investigate the thermal and mechanical properties of the most critical joint in any TPG module, where the heat comes out of the TPG and into a stronger material which will make contact with the cooling block. Three structures were made with polyimide coated TPG, two of them having a 70 micron glue layer and one having a minimal glue layer which in practice is about 10 microns thick. A further two structures were made with uncoated TPG.

Neither of the uncoated samples survived handling, which gives an indication of how delicate they are. The results of the three coated samples are shown in Figure 4. Although we put 1 W in to the heater, the gradient of the line tc1-tc2-tc3 shows that only 0.85 W was coming out along the aluminium bar. The rest was presumably lost by convection/radiation from the heater itself which was not in very good thermal contact with the TPG. The sample was simulated with the same program as used for the modules. This time there is no approximation involved in the third dimension because the sample really has a constant width. The material properties of Table 1 were used in the simulation, giving the results shown in Figure 4. It can be seen that two of the samples agree rather well with expectations and one of them has a markedly higher temperature drop across the glue joint.

All this really shows is that we failed to make a good glue joint in one case out of three. Somewhere in the bad joint there must be a void covering a significant fraction of the area. Although we have no evidence yet, there is a suspicion that the void could be between the TPG and the polyimide, rather than a bubble in the glue.

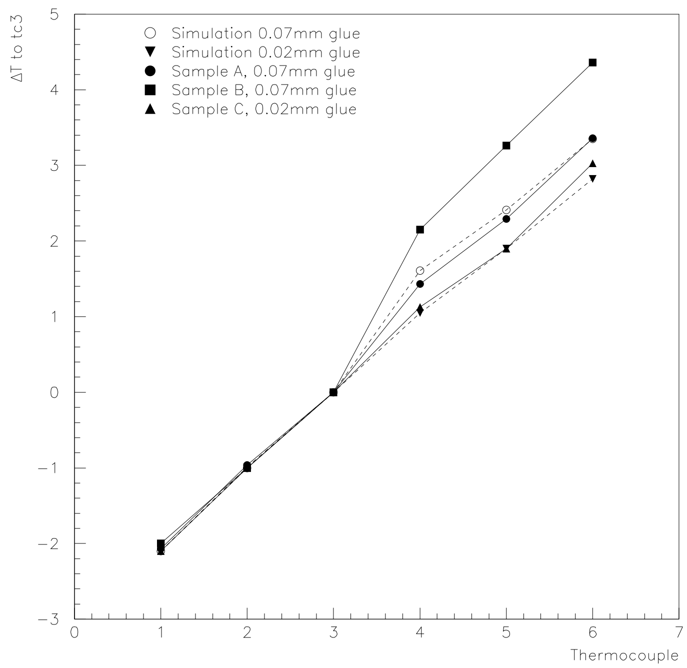

Figure 4: Simulation and measurement of TPG-glue samples. Simulated power was adjusted by 15% to match the measured temperature gradient in tc1-tc2-tc3. The step tc3-tc4 is related to the glue and the transverse TPG conductivity. The gradient tc4-tc5-tc6 is a measure of the longitudinal TPG conductivity.

10 Some Conclusions

Three possible modules designs have been shown to be sufficiently well cooled to avoid runaway in the expected SCT environment;

- a module with single point cooling and a 0.5 mm TPG base covering the whole area between the wafers,
- a module with two point cooling and a 0.5×10 mm^2 TPG spine,
- or a module with two point cooling and a 0.5×20 mm^2 BeO spine.

The decision on which module to use should be based on radiation length and reliability.

After a little experience we consider that uncoated TPG is not acceptable as a module component. Not only because of its fragility but also because even with the most careful handling it continually sheds fragments of graphite.

We are concerned about the reliability of the TPG-glue joint and would like to see a module design in which it is impossible to put a tensile force on the TPG in its weak direction.

More statistics are needed on the glue and grease joints in order to make an informed decision on safety factors.

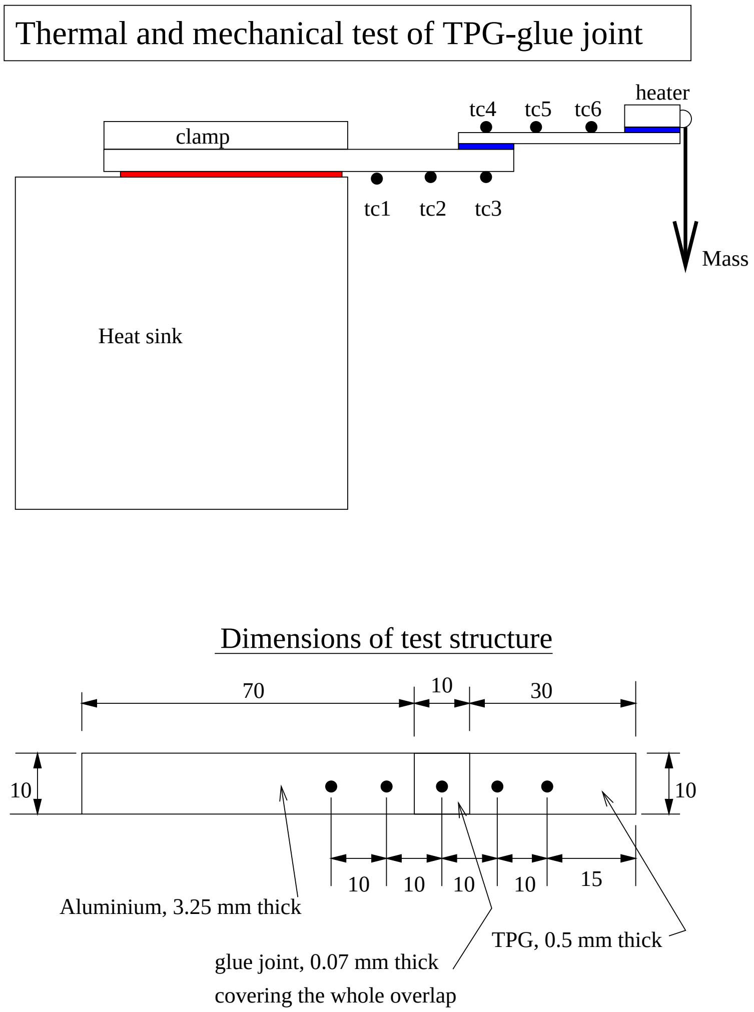

Figure 5: A sketch of the TPG - glue joint test.