JUNE 2011 | VOLUME V | ISSUE VI

Qpedia

IN THIS ISSUE

Qpedia

SUBSCRIBE ON LINE AT:
www.qats.com/qpedia

Advanced Thermal Solutions is a leading engineering and manufacturing company supplying complete thermal and mechanical packaging solutions from analysis and testing to final production. ATS provides a wide range of air and liquid cooling solutions, laboratory-quality thermal instrumentation, along with thermal design consulting services and training. Each article within Qpedia is meticulously researched and written by ATS' engineering staff and contributing partners. For more information about Advanced Thermal Solutions, Inc., please visit www.qats.com or call 781-769-2800.

EDITOR
KAVEH AZAR, Ph.D.
President & CEO,
Advanced Thermal Solutions, Inc.

MANAGING EDITOR
BAHMAN TAVASSOLI, Ph.D.
Chief Technology Officer,
Advanced Thermal Solutions, Inc.

NORTH AMERICA
ADVANCED THERMAL SOLUTIONS, INC.
89-27 Access Road Norwood, MA 02062 USA
T: 781.769.2800 | F: 781.769.9979 | www.qats.com

EUROPE
ADVANCED THERMAL SOLUTIONS, B.V.
De Nieuwe Vaart 50 | 1401 GS Bussum
The Netherlands
T: +31 (0) 3569 84715 | F: +31 (0) 3569 21294
www.qats-europe.com

ADVERTISING
TO PLACE AN AD IN QPEDIA:

Contact John O' Day at 781.949.2521

Qpedia was launched in 2007 as a technology eMagazine focused on the thermal management of electronics. It is designed as a resource to help the engineering community solve the most challenging thermal problems.

The eMagazine is published monthly and distributed at no charge to over 17,000 engineers worldwide. Qpedia is also available online or for download at www.qats.com/qpedia.

Qpedia's editorial team includes ATS' President & CEO, Kaveh Azar, Ph.D., and Bahman Tavassoli, Ph.D., the company's chief technologist. Both Azar and Tavassoli are internationally recognized experts in the thermal management of electronics.

For more information on how you can generate awareness about your company or products, contact John O' Day at 781.949.2521.

Qpedia

Features

TWO THERMAL MANAGEMENT POWERHOUSES JOIN FORCES

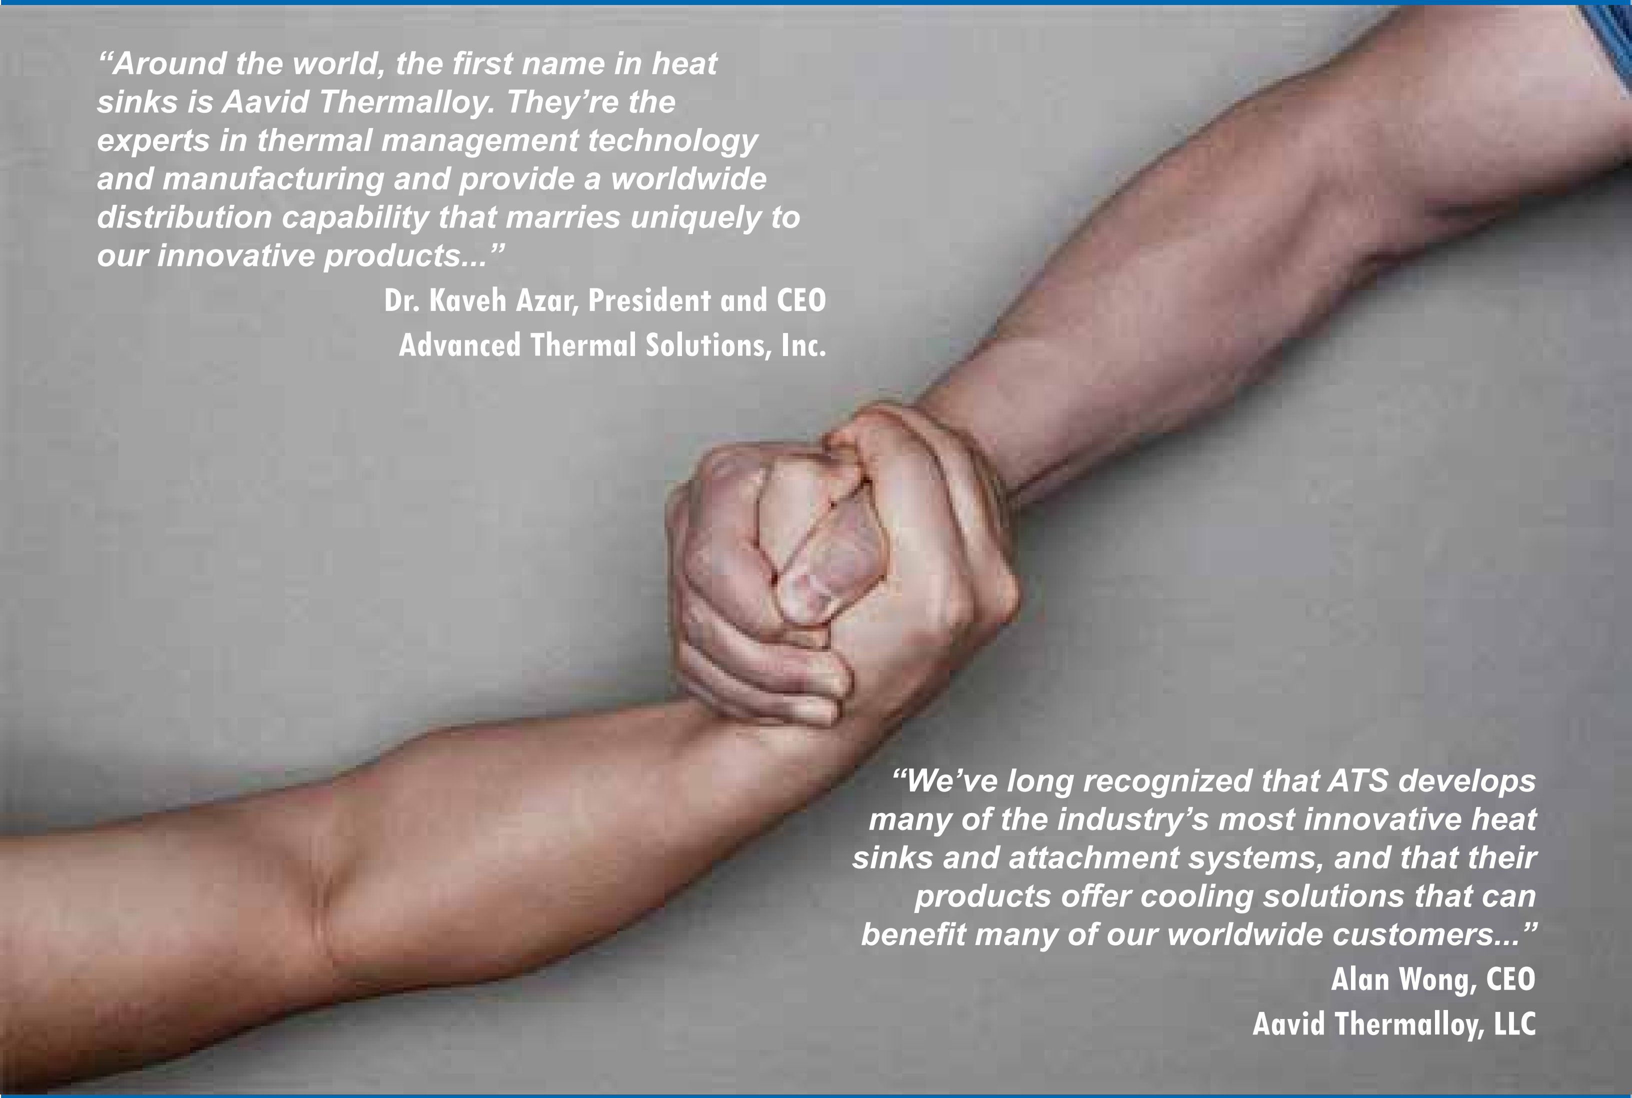

AAVID THERMALLOY TO GLOBALLY DISTRIBUTE ATS HEAT SINK AND ATTACHMENT PRODUCTS

Aavid Thermalloy, LLC, the world's oldest and one of the largest suppliers of thermal management products to the worldwide electronics marketplace, has signed a distribution agreement with Advanced Thermal Solutions, Inc. (ATS), a leader in innovative component cooling solutions. Aavid Thermalloy, with its global customer base and broad network of sales offices, representatives and distributors, will now have access to ATS' high performance heat sink and attachment technologies to address the cooling of today's and tomorrow's electronic devices.

As part of the agreement, ATS will supply Aavid Thermalloy, with its complete line of maxiFLOW™ heat sinks whose patented design and thermal performance is unmatched in the electronics cooling market. In addition, Aavid Thermalloy will also distribute ATS' patented maxiGRIP™ heat sink attachment systems. Its compact design securely attaches heat sinks to hot components on densely populated PCBs, and detaches quickly and cleanly when needed.

Please contact Aavid or ATS for more information regarding product availability.

ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®

Advanced Thermal Solutions, Inc.
89-27 Access Road | Norwood, MA | USA
T: 781.769.2800 | F: 769.769.2800
www.qats.com

AAVID THERMALLOY

Aavid Thermalloy
70 Commercial St. | Concord, NH | 03301 | USA
T: 603.224.9988 | F: 603.223.1790
www.aavidthermalloy.com


Need Direction?

Phase Change Materials

and Transient Loads

In the Thermal Management field, phase change materials (PCMs) have been studied for many years, and they have been put to many uses. They have been used as thermal interface materials (TIMs), where their phase change characteristics allow them to "wet" a surface when they warm up, conforming to imperfect mating surfaces[1]. This gives them excellent performance as a TIM, even though the material itself has relatively poor thermal conductivity [2].

Phase change materials have also been long used for thermal energy storage. This application relies on a completely different property of a phase change material: its latent heat of fusion. As energy is absorbed by a solid material, its temperature rises linearly until it reaches its melting point. To change phases to a liquid, extra energy input is needed, and while the material is absorbing this energy, its temperature remains the same. From a thermal management perspective, this is great, because energy can be dissipated into a phase change material while the cooled object stays at a constant temperature. In many regards, phase change material is the thermal management equivalent of a capacitor in electronics.

Of course, this only lasts as long as there is solid material to melt. And, because many commonly used phase change materials conduct heat poorly, that solid material needs to be close to the thermal interface in order to be effective. Fleischer, Weinstein and Kopec performed a study that illustrates that point [3]. In the study, a volume of paraffin wax PCM roughly 70mm per side was used. The sides were insulated, a heat source was applied to the bottom, and the top was covered by a cold plate maintained at 5ºC. In this setup, the heat flux is transferred from bottom to top through the PCM.

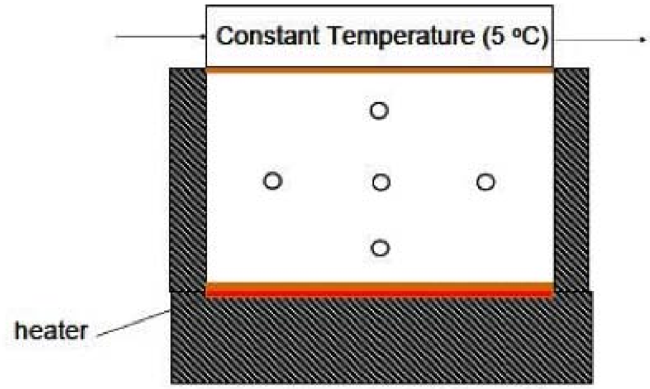

Figure 1- Arrangement of Thermocouples in PCM [3]

Five thermocouples were embedded in the PCM to monitor temperatures along the direction of heat flux, as shown in Figure 1. Thermocouple #3 is at the bottom, #1 is at the top, and #2, 4, and 5 are across the middle.

With an applied heat flux of 7 W, a large temperature gradient develops across the PCM, peaking at 38ºC over a distance of about 30 mm between the top and bottom thermocouples. This gradient can be seen around the 5000 second point in Figure 2. It can also be seen that it takes more than 3 hours (12000 s) for the melt front to reach all of the thermocouples. At this point, it was found that there were still solid pieces of PCM circulating within the liquid PCM, demonstrating how the low thermal conductivity of the PCM contributed to localized temperature differences.

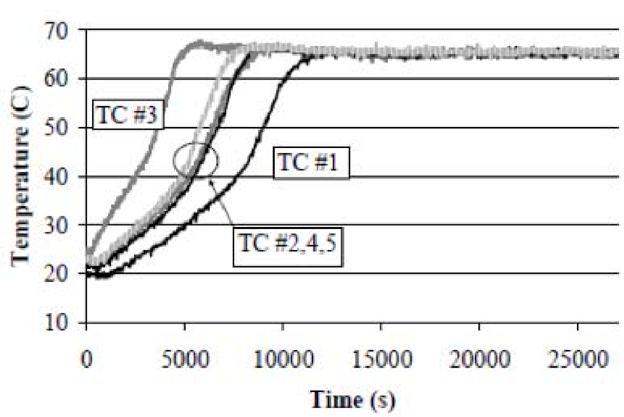

Figure 2- PCM Temperature With 7W Applied Power [3]

One way to work around the low thermal conductivity is to embed a PCM into substrates with higher conductivity, such as aluminum. The use of aluminum foam filled with PCM (see Figure 3) was studied by NASA to provide cooling for LED helmet lights on spacesuits [4]. A particular hydrocarbon family was chosen (which includes paraffin wax) for low toxicity and compatibility with metals such as aluminum. It was found that the aluminum foam spread heat effectively through the PCM, and that desired operating temperatures could be maintained for the target time of 7 hours.

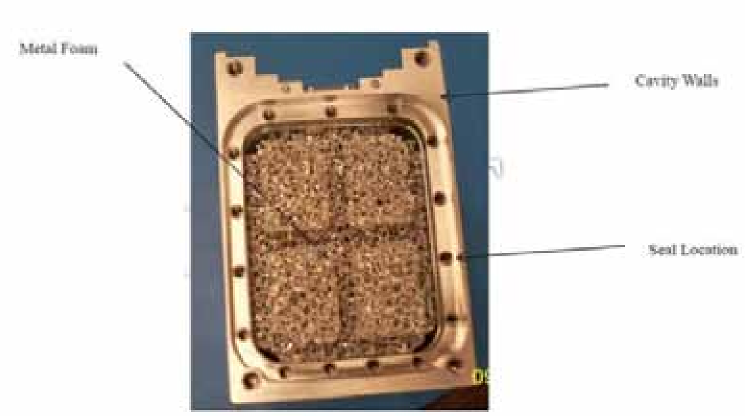

Figure 3- Aluminum Foam With Embedded PCM [4]

One of the drawbacks of using an encapsulated PCM is high price, which means that for now its use is limited to such demanding applications as cooling LEDs in outer space. However, future developments may see more widespread use of encapsulated PCMs for energy storage uses as well as for thermal management [5].

Another strategy to deal with the low thermal conductivity of the PCM is to embed a heat spreader within the PCM. Fleischer, Weinstein and Kopec [3] used a heat sink made of carbon fiber and embedded it in the same volume discussed earlier. Shown in Figure 4 below, the conductivity of the heat sink is directional, and is 500 W/mK along the fin direction. That is approximately 25% better than copper, and other exotic materials such as graphite sheets could be excellent heat spreaders as well, promising thermal conductivities as high as 1500 W/mK [6].

Figure 4- Carbon Fiber Heat Sink in PCM [3]

With 7W power again applied to the bottom of the volume, the effectiveness of the heat spreader is clear. The thermocouple temperatures show that the temperature throughout the volume of PCM is uniform (See Figure 5), and well below the melt point at steady-state.

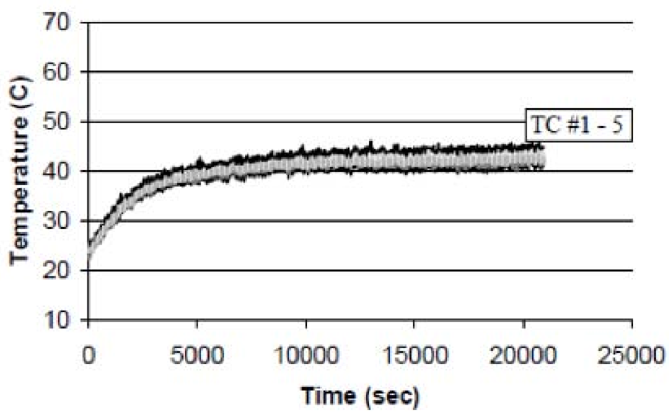

Figure 5- PCM with Embedded Heat Sink, 7W Applied Power [3]

This represents an approximately 20ºC difference in steady-state temperature, which would greatly add to the longevity of electronics in a real-world application. The rate of temperature rise at the lowest thermocouple has also been reduced, indicating the heat spreader has improved the ability of the PCM to regulate temperatures during transient heat loads. Other novel research in Phase Change Materials is ongoing as well, and new applications of PCM in thermal management of electronics are expected to emerge within the next few years.

References:

1. "Tpcm 580 Series Phase Change Material on Laird Technologies," http://lairdtech.thomasnet.com/viewitems/phase-change- tims/tpcm-8482-580-series-phase-change-material?, [June 30, 2011].

2. Wilson, J., (2005, May). "Phase Change Material Thermal Properties," Electronics Cooling. http://www.electronics-cooling.com/2005/05/phase-change-material-thermal- properties/ [Jun. 15, 2011].

3. Fleischer, A., Weinstein, R., Kopec, T., "Analysis of Transient Thermal Management Characteristics of PCM with an Embedded Carbon Fiber Heat Sink," Villanova University, Pennsylvania, 2006.

4. Bue, G., Nguyen, H., Keller, J., "Analysis and Design of Phase Change Thermal Control for Light Emitting Diode (LED) Spacesuit Helmet Lights," National Aeronautics and Space Administration, Texas, 2010.

5. Zalba, B., Marin, J., Cabeza, L., Mehling, H., "Review on thermal energy storage with phase change: materials, heat transfer analysis and applications," Applied Thermal Engineering 23, 2003, pp. 251-283.

6. "KANEKA Electrical and electronic materials," http://www.elecdiv.kaneka.co.jp/english/graphite/index.html , [Jun. 20, 2011].

iQ-200

iQ-200 is a new instrument from Advanced Thermal Solutions, Inc. that measures air flow, air temperature, surface temperature and differential pressure drop simultaneously.
The instrument comes with 16 velocity channels, 12 thermocouples and 4 pressure transducers.
All three functions are bundled in one instrument.

OVERALL DIMENSIONS
(D x W x H)
43 cm x 34 cm x 12.6 cm
(17" x 13.3" x 5")

NUMBER OF VELOCITY CHANNELS - 16

FLOW RANGE
0 to 50 m/s (0 to 10,000 ft/min)

NUMBER OF THERMOCOUPLES
12 (J, K, T, and E types)

NUMBER OF PRESSURE SENSORS
4 (Differential)

PRESSURE RANGE
0 to 0.15 psi

SOFTWARE
IQstage™

VOLTAGE INPUT
110V or 220V

WEIGHT
7.5 kg (15 lbs.)

ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®

89-27 Access Road, Norwood, MA 02620 USA
T: 781.769.2800 | F: 781.769.9979 | www.qats.com

Modern Test Equipment

for TIM Testing

The quality of thermal interface materials (TIM) has become increasingly important as the ever increasing dissipation level of ICs requires more sophisticated solutions to reduce the thermal resistance Rth along the heat-flow path. TIM manufacturers have considerably enhanced TIM's thermal conductivity by using nanoparticles as fillings in new TIM materials. Thermal interface materials (TIMs), as shown in figure 1, are placed between electronic devices and their heat sinks to enhance heat transfer. Thus, for electronics cooling, it is necessary to know the TIM thermal resistance. Most of the TIMs characterization methods are based on the ASTM standard D14570. According to ASTM D14570, a TIM sample is placed between a hot and a cold meter bar with a constant heat flux being applied. The thermal resistance of the sample TIM is calculated from the known heat flux forced through the tester and from the measured temperature drop, including the contact resistance of the TIM to the tester.

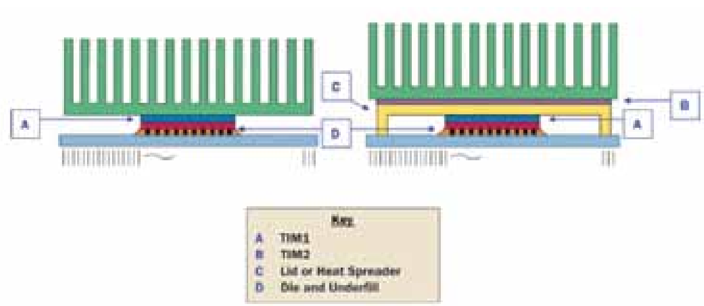

Figure 1 - Schematic of TIM Application [1]

However, there are difficulties in the characterization of new TIMs, especially for TIMs that are made of advanced interface materials: because their thermal resistances approach zero. The ASTM D5470-01 standard uses a steady state heat transfer method that is not accurate for modern TIMs because it was designed for interface materials with substantial thermal resistances. The main source of the ASTM standard uncertainty is its use of higher pressure than used in real applications. The ASTM standard prescribes that the measurements to be done with a clamping force of 3 MPa. For greases and phase change materials, the high pressure will result in a lower gap and a lower measured thermal resistance than in an actual application.

High Accuracy Thermal Interface Material Tester

Trinity College Dublin researchers, in collaboration with Alcatel-Lucent, designed a high accuracy TIM tester. A schematic of the TIM tester is shown in Figure 2. The high accuracy TIM testing setup consists of two OFHC copper meter bars, which have opposing faces with a flatness of a few microns and RMS surface finishes. The lower meter bar (5) is connected to the water jacket (4) that is cooled by a Julabo model F33 HE constant-temperature circulator, with a temperature control of ±0.01 ºC. The upper meter bar (7) is connected to the copper heater (8), which contains two 250-W cartridge heaters. The meter bars are enclosed in 15 mm of insulation to minimize heat leaks. A linear actuator (12) is mounted to the top plate with a load cell. A controller capable of micro-stepping the linear actuator allows for sub-micron changes in displacement with an accuracy of ±0.2% of rated load, or 10 N. Over the meter bar area, this leads to a pressure uncertainty of ±6.25 kPa or ±0.9 PSI. The distance between the opposing faces of the meter bars is measured using an optical micrometer (6). Steel gauge pins were mounted, normally, to the meter bar surfaces, near the contacting surfaces, to serve as optical trips for the micrometer.

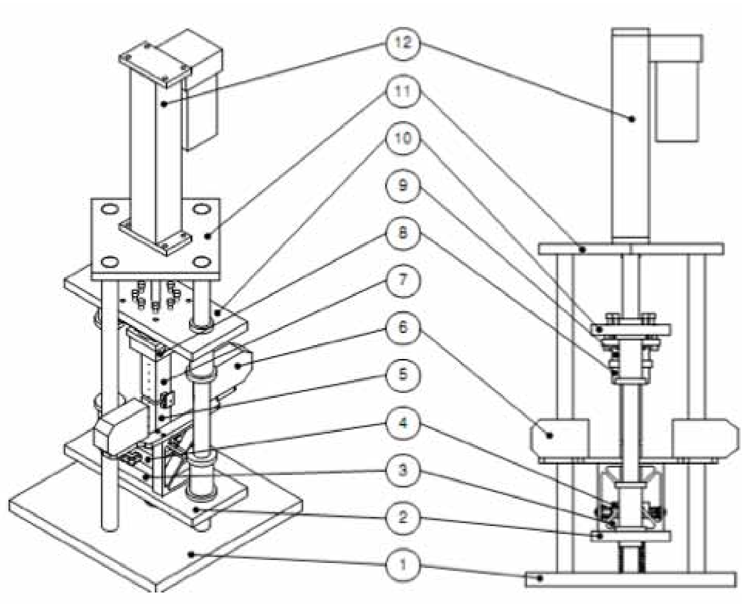

Figure 2 - Schematic of Trinity College TIM Testing Device [2]

Each meter bar was instrumented with four thermistors 0.38 mm in diameter. Thermistors were chosen due to their high temperature sensitivity, which helps to minimize the uncertainty of the temperature measurements. The thermistors were inserted into the holes with a small amount of oil to improve thermal contact and held in place using RTV silicone adhesive. Thermistor resistances were measured using a LakeShore model 370 AC resistance bridge scanner, which has resolution and noise levels well below 0.001K.

The thermistors on both bars were calibrated simultaneously against a Hart Scientific 5611T secondary reference probe whose absolute calibrated uncertainty is ±0.002 K. The entire assembly was placed inside a chamber with a temperature control stability of ±0.01 K, reducing the temperature fluctuations to the level of 0.001K at the thermistor locations. The estimate for the temperature uncertainties for the thermistors was less than ±0.002 K. The micrometer is calibrated by bringing the meter bars into contact and measuring the distance between the optical trips. The optical micrometer allows measurements of samples TIM thickness up to approximately 20 mm with an uncertainty of ±0.15 μm.

A sample temperature distribution along the meter bars in contact with each other, at an interface pressure of 0.4 MPa, was taken to compute the difference in energy balance between the two meter bars, as shown in Figure 3. In every case, the power conducted through the two meter bars agreed with each other. For powers above 2 W, the absolute uncertainties are approximately 1 %. Figure 4 shows the self-contact resistance as a function of heater power for an applied pressure of 0.4 MPa, remaining virtually constant over a range of power levels. Figure 5 shows the effective thermal conductivity of a sample thermal gap pad as a function of pad thickness during compression. As expected, the initial increase in thermal conductivity is due to the reduction in the initial contact resistance between the pad and meter bars. As the pad gets compressed further, the bulk resistance of the sample becomes the main thermal resistance and the effective conductivity asymptotes to approximately 5.4 W/m•K, with a calculated uncertainty of approximately 1.5%. The manufacturer's specified conductivity for this pad is 5 W/m•K.

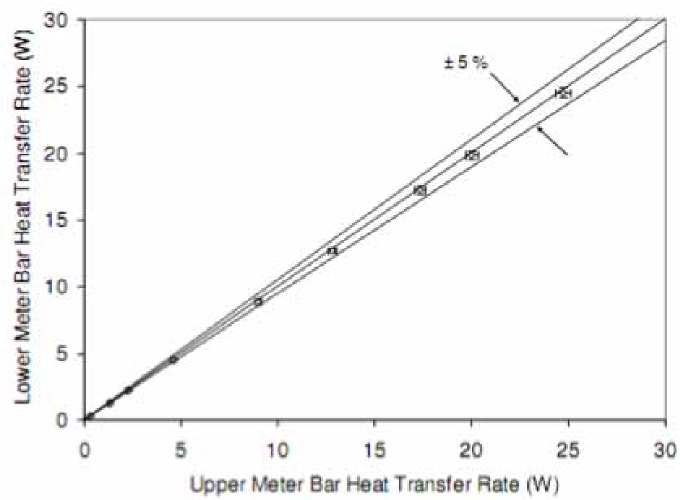

Figure 3 - Difference in Energy Balance Between Upper and Lower Meter Bars [2]

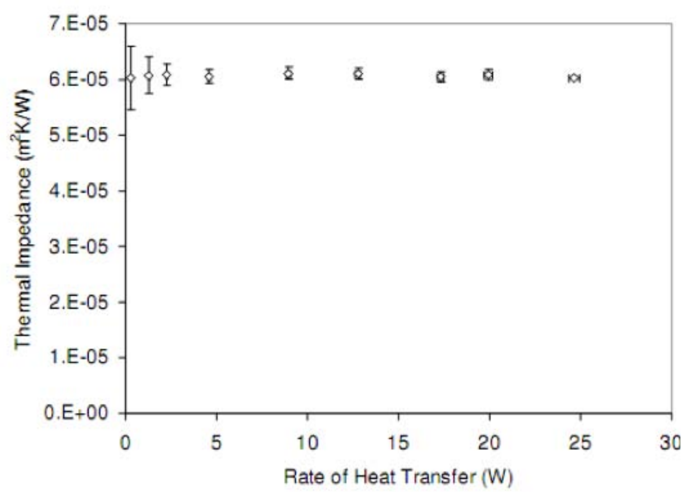

Figure 4 - Self Contact Thermal Resistance of Meter Bars [2]

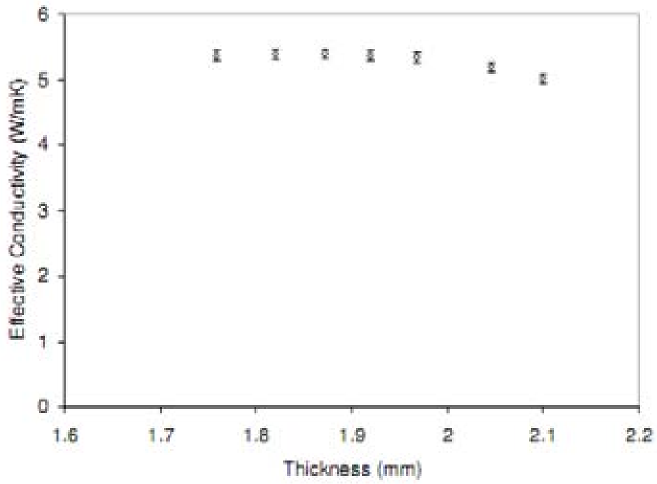

Figure 5 - Thermal Resistance of Sample Gap Pad TIM [2]

A Static TIM Tester

Another high accuracy TIM tester design has been developed by researchers at Budapest University of Technology and Economics. The main advantage of this design is to use microelectronics for temperature and heat flux sensors, making it possible to place these sensors in the closest proximity of the measured sample. The second advantage is its symmetrical structure with reversible direction of the heat flow, eliminating certain inaccuracies of the measurement, such as offset errors of the sensors.

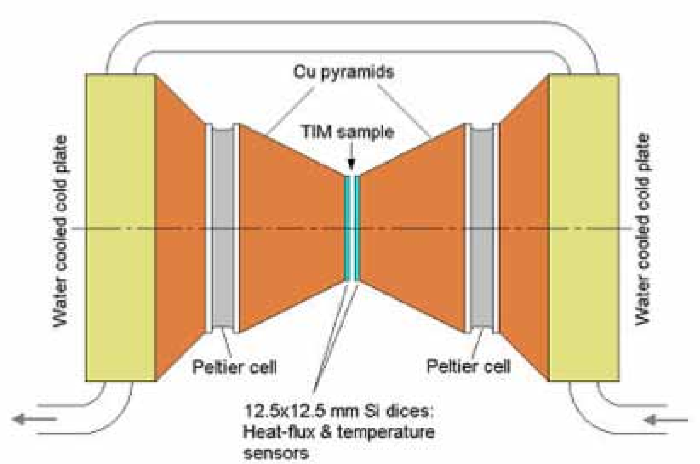

Figure 6 - Schematic of Budapest University TIM Tester [3]

Figure 6 shows a schematic of the heat flux/temperature measurement arrangement. The heat source is provided by two, symmetrically positioned Peltier cells. Both the heat flux and the temperatures are measured by the two silicon dices that contact the sample TIM. Figure 7 shows the mechanical design of the TIM tester. Components 1, 2 and 3 represent water- cooled cold plates, Peltier cells and copper pyramids, respectively. The lower components are fixed to the base plate of the assembly, while the upper components move along the vertical poles. The parallelism of the two grips is adjusted by 3 screws.

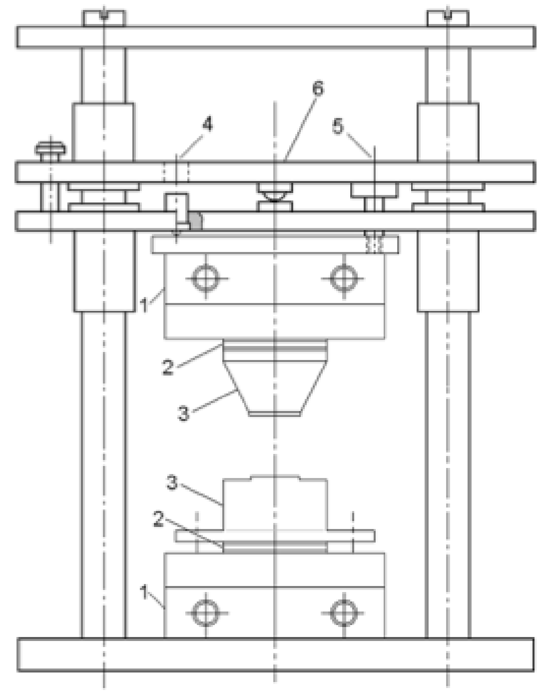

Figure 7 - Mechanical Drawing of TIM Tester [3]

The pressure is set by appropriate weights posed onto the plate 6. Figure 8 shows the actual testing equipment. A digital microscope monitors the region of the TIM sample and the upper/lower sensor chips as well as measures the thickness of the sample. Figure 9 shows a sample image from the microscope.

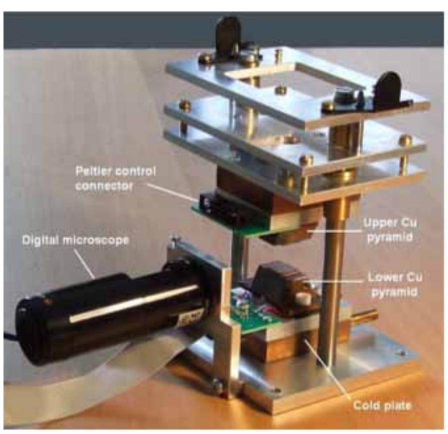

Figure 8 - Actual TIM Testing Equipment [3]

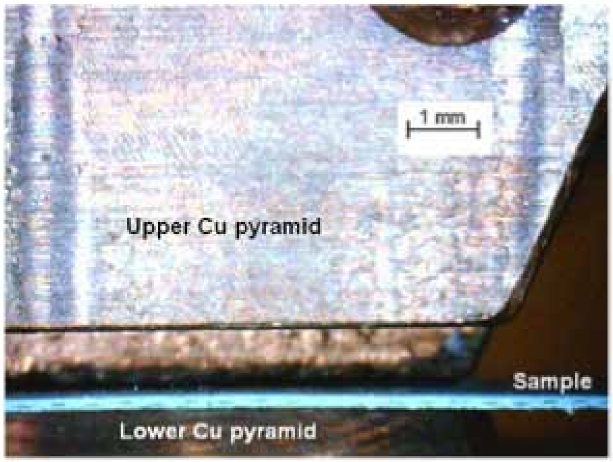

Figure 9 - Sample TIM in the Tester [3]

The two Peltier cells are controlled by the two identical Peltier control units (PCUs) with the thermal state of each Peltier cell measured by two temperature sensors. Figure 10 shows the block diagram of the electronics. The PCUs control the Peltier currents, which set different heat flux and temperature values on the TIM sample. The silicon sensor chip, in contact with the sample, provides two voltages: one proportional to the heat flux and another proportional to the temperature. The output voltage is proportional with the temperature difference between the two sides of the silicon chip. The voltages from the silicon chips are amplified and digitized by the two preamplifier electronics, shown in Figure 11.

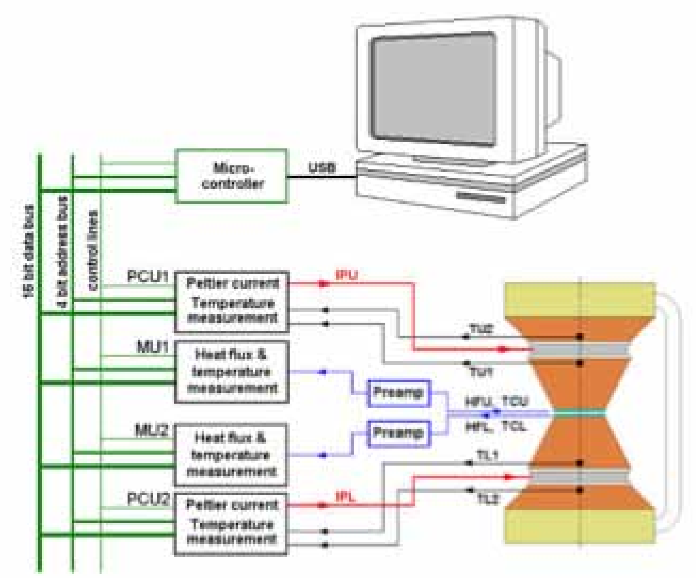

Figure 10 - Purdue Researchers Testing Micro-Channel Heat Sink [3]

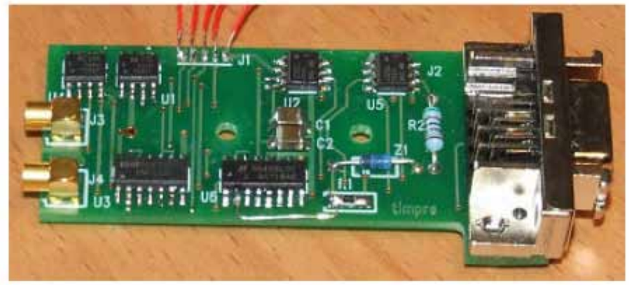

Figure 11 – Preamplifier [3]

Figure 12 shows the layout of silicon chips used to measure the heat flux and surface temperature of the sample TIM. The two sensing areas cover 1 cm^2. Dividing the area into two halves was a need required both by the evaluation electronics and the control of parallelism. The sensitivity of the chip is about 40 $V*cm^2/W$, in heat flux sensor mode. In the temperature sensor mode, we exploit the above mentioned temperature dependence of the sensor resistance. The chip is bounded to a very small printed wire plate, which in turn is connected to the preamplifier. Figure 13 shows an actual silicon chip.

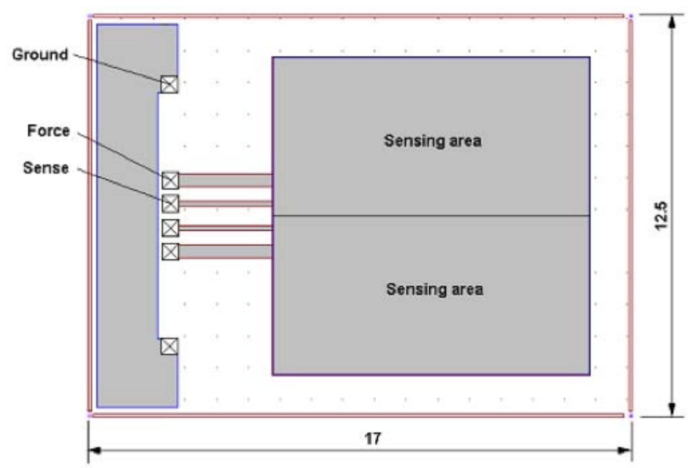

Figure 12: Layout of the Measurement Silicon Chip [3]

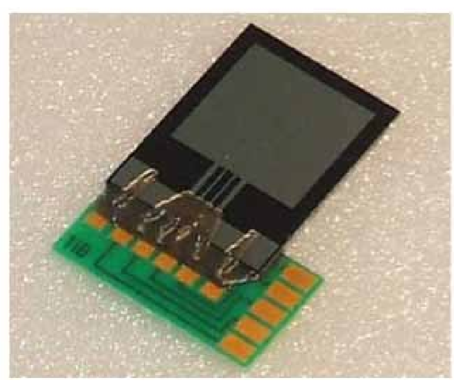

Figure 13 - Actual Measurement Silicon Chip [3]

The resolution of heat flux measurement is the least significant bit (LSB) value determined by the amplifiers, A/D converter and electronic noise. The measurements show this resolution to be 8 mW. The resolution of temperature measurement is ≈ 0.05°C. A dedicated circuit is used to measure the temperature difference between the upper and lower chips. Since the maximum heat flux of the chips is about 40 W, the expected thermal resistance resolution is ~ 0.25 % for a TIM sample with Rth = 0.05 K/W and ~ 1.2 % for 0.01 K/W.

There has been a lot of research in TIM characterization devices to get more accurate prediction of the TIM performance in application. Some prototypes use more accurate thermistors and better calibration to derive more accurate, repeatable thermal resistance measurements. Others are using new electronics, such as the heat flux/temperature detecting silicon chips, to get better measurements. As the challenge of characterizing these modern TIM with near zero thermal resistance increases, these new measurement systems might provide promising solutions.

References:

1. “Developments with metallic thermal interface materials”, Electronics Cooling website http://www.electronics- cooling.com/2007/05/developments-with-metallic-thermal-interface- materials/

2. Kempers, R., et al., “Development of a High-Accuracy Thermal Interface Material Tester”, Proceedings SEMITHERM XXII, March 2006, San Clara, pp.221-226

3. Szekely, V., et al., “Design of a static TIM tester”, Thermal Investigation of ICs and Systems, Sept. 2008, Rome, Italy, pp.132-136.

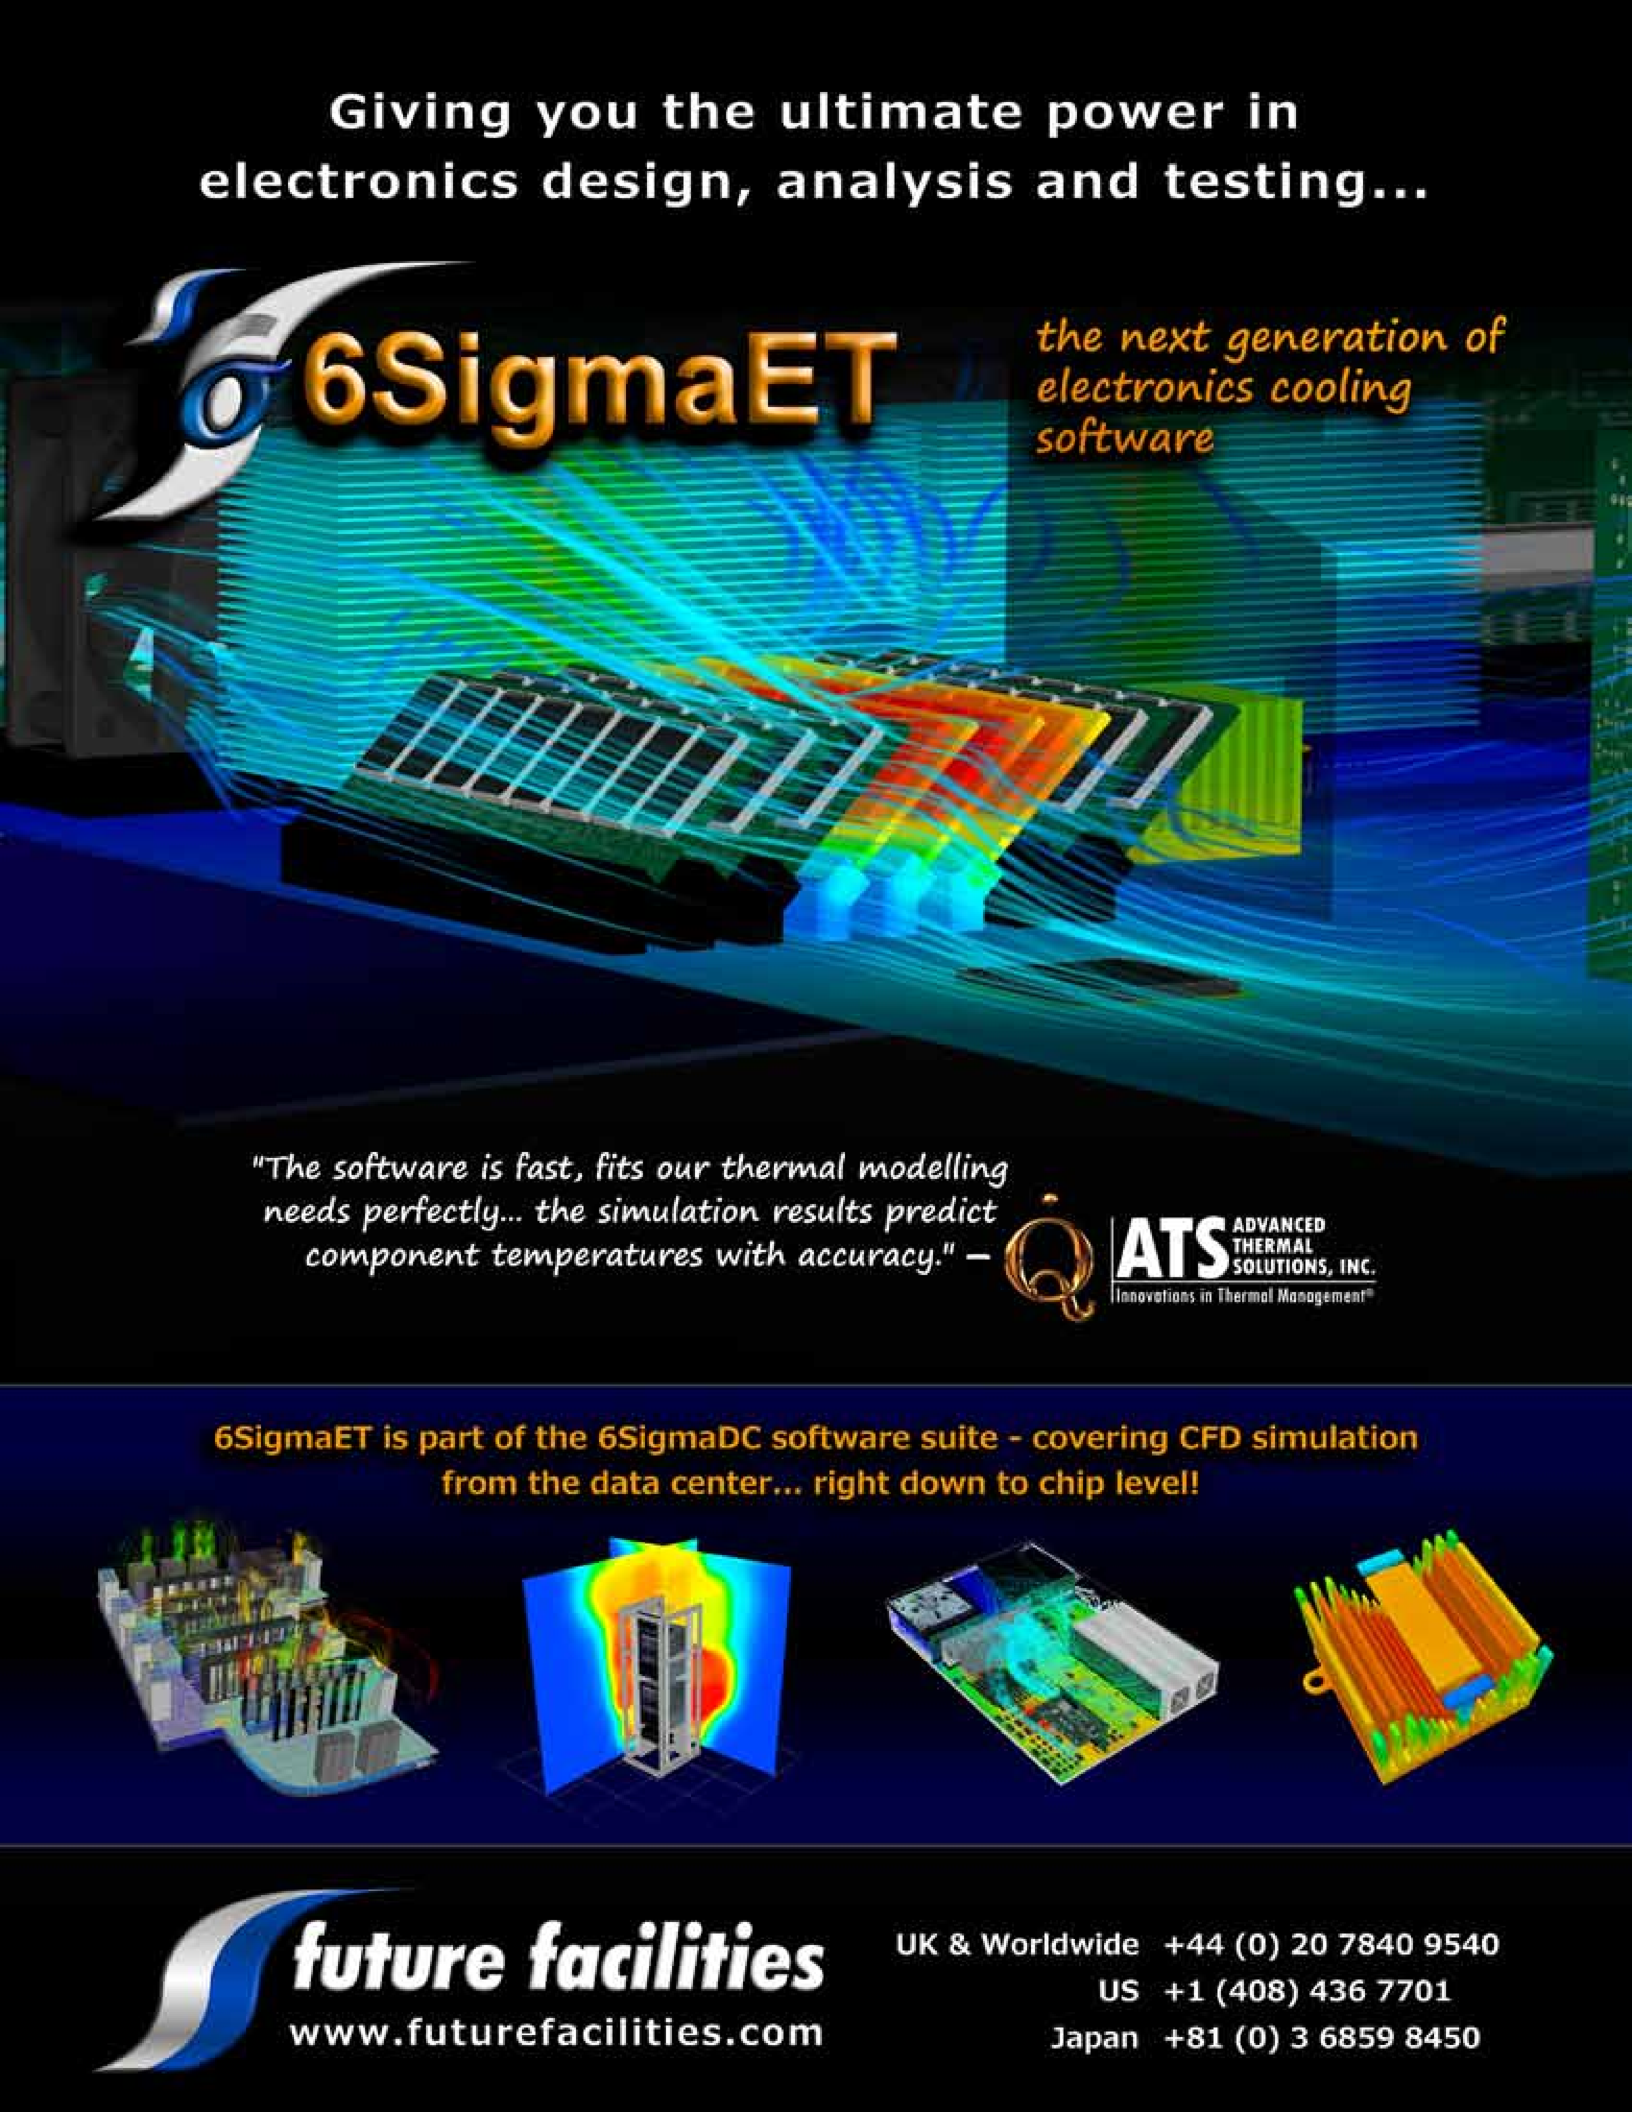
Giving you the ultimate power in electronics design, analysis and testing...
6SigmaET
the next generation of electronics cooling software
"The software is fast, fits our thermal modelling needs perfectly... the simulation results predict component temperatures with accuracy." –
ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®
6SigmaET is part of the 6SigmaDC software suite - covering CFD simulation from the data center... right down to chip level!
future facilities
www.futurefacilities.com
UK & Worldwide +44 (0) 20 7840 9540
US +1 (408) 436 7701
Japan +81 (0) 3 6859 8450

The Benefits of

Supply Air Temperature Control in the Data Centre

David King, Consultant Engineer, Future Facilities Ltd.

Controlling the temperature in a data centre is critical to achieving maximum uptime and efficiency. But, is it being controlled in the correct place? Whilst data centre layouts have moved swiftly towards the segregation of hot and cold air, with hot aisle/cold aisle designs being industry standard, typical air handler temperature control strategies are still rooted in the mainframe and comms rooms principles that have been used for decades.

Return air temperature control aims to achieve a constant ambient temperature within the room by varying the temperature of the air being supplied by the individual cooling units. This often results in a significant variation of the temperature of the air being supplied to the IT equipment, a variation which is out of the hands of the facility or IT manager. Controlling the temperature at the supply side of the air handler, rather than the return, gives control back to the facility manager and makes the temperatures being supplied to servers much more predictable. This predictability leads to the potential for raising water temperatures and reducing running costs as well as savings in capital costs during the construction of new facilities.

This paper uses a Virtual Facility to demonstrate the differences between the two control strategies and quantifies the benefits of switching to supply side control.

Introduction

Over the past 10 – 15 years, the importance of airflow management and close temperature control in the data centre environment has become increasingly evident. Old design strategies that simply called for throwing the necessary kW cooling into the data centre space and regulating the ambient temperature became unable to cope with modern server designs and power densities. Towards the middle of the last decade, data centre layouts that explicitly arranged the space into hot and cold aisles became the de-facto standard for new builds. As the end of the decade approached, more elaborate methods for actually segregating the hot and cold air appeared: containing the cold aisle, containing the hot aisle and adding chimneys to cabinets are three good examples. However, the commonly adopted strategy for controlling the temperature in the data centre is still largely fixed on using the return air temperature, assuming it reflects the ambient air in the space.

It is often forgotten that the only function of the cooling system in the data centre is to provide air of an acceptable temperature and humidity to the IT equipment (at their inlets), rather than occupant comfort or general room air temperatures. Considering this fact, it is obvious that the return air temperature is not an appropriate measure of the IT equipment inlet temperature, and it would seem that explicitly controlling the temperature of the air being supplied by the system would be desirable. However, adoption of this strategy has not been forthcoming. The idea is not new, for example, in 2006 [1] research was conducted that showed CRAC supply temperature is well suited to controlling equipment inlet temperatures, but this work was part of a wider, more complex, control system involving variable fans and grille dampers. Control systems resulting from this research appeared on the market, but their complexity meant that they were not widely adopted and the benefits of supply temperature control were lost in the noise. Recent research [2] has investigated the potential for controlling data centre temperature by using the sensors built into the servers themselves, with encouraging results. However, this work is still in the early stages and this type of control strategy (combined with current equipment features) means that it is not yet mature enough to be available for general adoption. With this in mind, at present, controlling the supply air temperature with a sensor in the supply air stream is the only viable alternative to return air control.

The Virtual Facility

The Virtual Facility is a full 3D computer representation of a data centre. Using models of the objects within the data centre (ACUs, IT equipment, PDUs etc.) and CFD simulation, the Virtual Facility gives an accurate account of the thermal performance of the space. It provides a buffer zone in which planned changes can be tested before they are committed to in the real facility.

The Virtual Facility used for investigating the topics presented in this paper is a hypothetical medium sized data centre. It has a floor area of 625m^2, over 200 cabinets and a maximum capacity of 800kW with N + 2 redundancy. The load is unevenly spread between server cabinets, blade cabinets, storage and networking, to mirror the situation in most real data centres. The design is not perfect, but best practices have been followed where possible and there are no significant airflow problems.

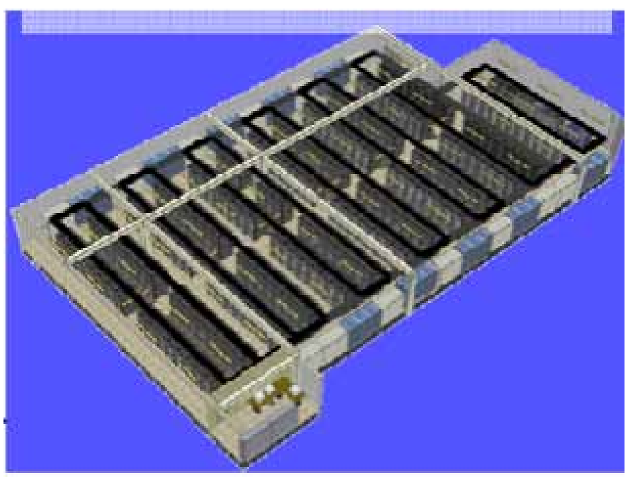

Although units come set up for return control because that is what the market expects (rather than any technical limitations) many newer cooling units come ready for either control strategy out of the box, or at least only require minor modifications. Older units may require slightly more work but the costs involved should be minimal.

Using a Virtual Facility as a testing ground, this paper explores the benefits of making the switch and details the downsides of the old return air strategy.

What is wrong with Return Air Control?

The effectiveness of a return air control strategy is dependent on two important caveats:

(1) The variation in temperature across the entire area of the unit inlet (return) is minimal; and

(2) The air supplied from each unit returns to that unit after being heated, forming a closed control loop.

However, in all but the smallest and simplest of data centres, neither of the above is likely to be true. At best, this will lead to significant inefficiency in the cooling system. At worst, it can cause the cooling system to contribute to hot spots and downtime.

Sensitivity to Sensor Location

Consider the first caveat, that temperature variation across the inlet must be minimal. Nearly all cooling units are controlled using a single, point sensor, but the size of the air inlets can be up to 2m^2. Any significant variation in temperature across this area can result in the sensor reporting a temperature that is higher or, even worse, lower than, the average of all the air that is entering the unit.

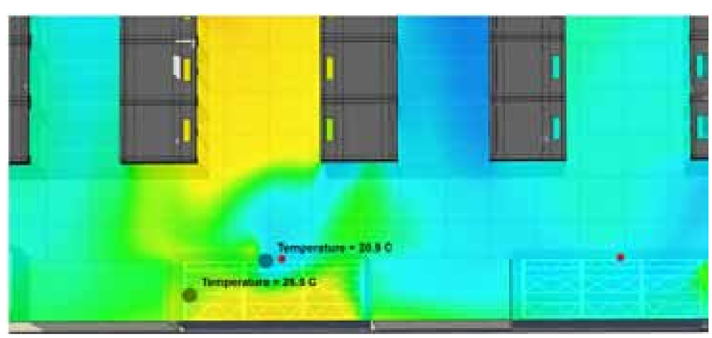

Figure 1 - Temperature Variation Over The Top Of A Cooling Unit

In the Virtual Facility, one of the units is supplying air 6 – 7°C higher than the others, simply because of the location of the sensor. Figure 1 shows a plane of air temperature just above the inlets of the cooling unit. The sensor sits in a small pocket of cool air at the front of the inlet returning a reading of 21°C to the controller, lower than the set- point of 22°C. With an input lower than the set-point the control algorithm tells the unit that it needs to do little or no cooling and it simply passes the air through. However, the average temperature across the inlet is actually 26.5°C and the air that the unit is passing through into the raised floor is 25°C! If the sensor were placed in a different location, on the far left say, then the controller response and unit behaviour would be completely different.

Closed Loop, Open Loop

Cooling unit control algorithms all assume that there is a direct relationship between the variable they adjust (the water valve or the compressor) and the variable they measure (air temperature). For the temperature at the return side to keep this relationship, most, if not all, of the air returning to the unit needs to have been supplied by that same unit. If this is true then as the supply temperature is changed, by adjusting the water valve or compressor, a similar change will occur in the returning air. However, this closed loop seldom exists in the data centre environment. As the loop is broken by air from different cooling units mixing before returning, the ability of the control algorithms to correctly react to conditions becomes diluted. In the Virtual Facility example, the unit with the sensor location issues is also suffering from this dilution problem.

The Virtual Facility allows the tracing of air streams (Figure 2), showing where the unit is pulling its' return air from. The majority is coming from the hot aisle in front of it, but a small portion is coming from the adjacent unit.

In fact the cold air that is recycling from the left hand unit is responsible for the spot of cool air that is fooling the controller of the unit on the right! The net result of this is that, as the left hand unit reduces its supply temperature, the unit on the right will actually increase the temperature at which it is supplying. It is worth noting that this effect can cause supply temperatures to go up and down over time, as the control algorithms fight each other. As the temperature goes up and down, it stresses the joints between components in the IT equipment, reducing the lifetime of the servers.

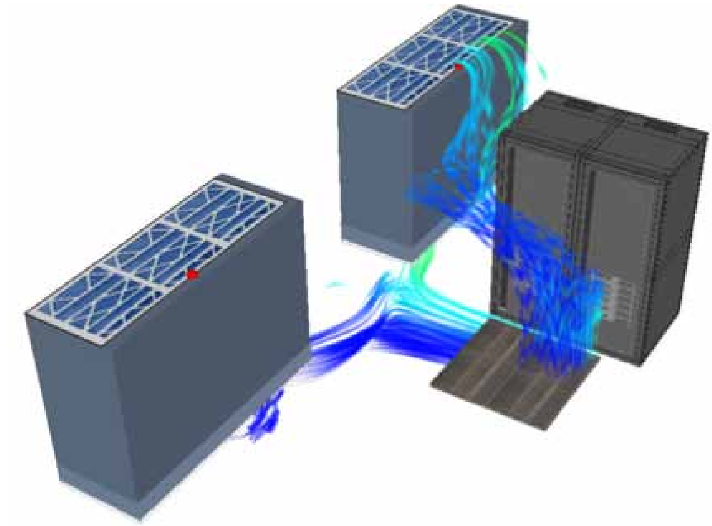

Figure 2 - Crosstalk Between Cooling Units

Curing The Ills

So, how does controlling the supply side overcome these issues? Return control systems are sensitive to sensor placement because it is likely that there will be a significant temperature variation across the return inlet. However, the design of most cooling units results in the off coil temperature being relatively uniform and, hence, the variation in temperature across the supply will be minimal. This almost completely eliminates the dependence of the controller response on the location of the sensor in the supply air stream.

Supply side control overcomes the issue of crosstalk between units, by significantly shortening the closed control loop. With a sensor in the return air path, the air from the unit has to go out into the data centre space and through some equipment before making its way back to the sensor, which leaves a lot of room for the loop to be broken. Moving the sensor in the supply air path guarantees that the control loop remains closed, as it will be only maybe 20 or 30cm away from the cooling coil. Changes made by the controller to the water valve or compressor will now be directly reflected in the temperature measured by the sensor.

The time variation problems with control, resulting from sensor location, are not simply limited to a cooling unit getting misleading messages, when air from another cooling unit returns to it. Figure 3 shows the way in which a single cabinet of equipment and a single cooling unit and sensor interact with different sensor locations. The thermal inertia and mixing that are present result in the system being somewhat under-damped when the sensor is on the return vent. The oscillation reduces as the sensor is moved towards the supply air stream.

This greater stability in the supply air temperature to the equipment should result in improved equipment life, since one of the major causes of failure in electronics is high rates of temperature change.

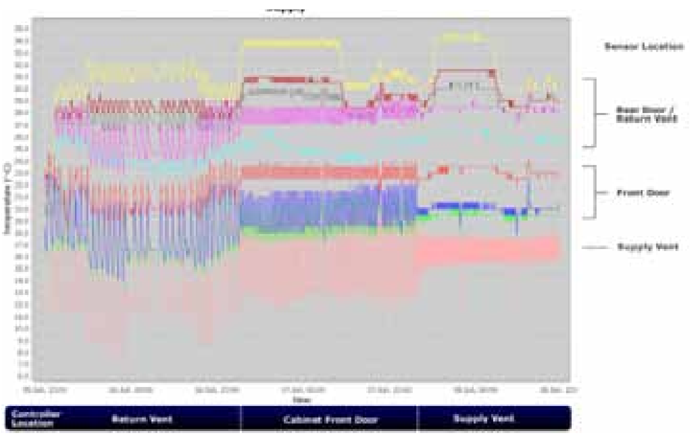

Figure 3 - Supply Temperature Oscillations Varying With Sensor Location

How Does This Affect My Data Centre?

What the previous section has hopefully shown is that, by using a return air control strategy, the facility manager sacrifices nearly all control over the temperature of air being supplied to the IT equipment. Switching to supply air control will hand this back to the facility manager. But, how will this impact the running of the data centre? Many facility managers would argue that, by having the units regulating the ambient temperature by adjusting their supply, temperatures will allow them to react better to shifting loads within the data centre. But, is this the case?

Using the Virtual Facility as a test bed, the two control strategies can be compared in a growing data centre at various stages in its life cycle. For the return control strategy a standard set-point of 22°C (typical of many facilities) was used. The 2008 ASHRAE Guidelines [3] recommend a temperature range for IT equipment of 18 - 27°C, so a set point of 18°C was chosen for the supply control strategy. The most important part of the cooling system is the temperatures at the inlets of the IT equipment, so these are a good metric for comparing the two strategies.

The first example is at the start of the life of the data centre, at around 25% occupancy (Figure 4). The load is unevenly distributed throughout the space, as the strategy is to load server cabinets from either end, working towards the centre.

Comparing the temperatures in the floor void (Figure 5) shows a variation of 9°C in the return control strategy as units are unevenly loaded. The supply control strategy results in an even temperature of 18°C regardless of the load on the cooling units.

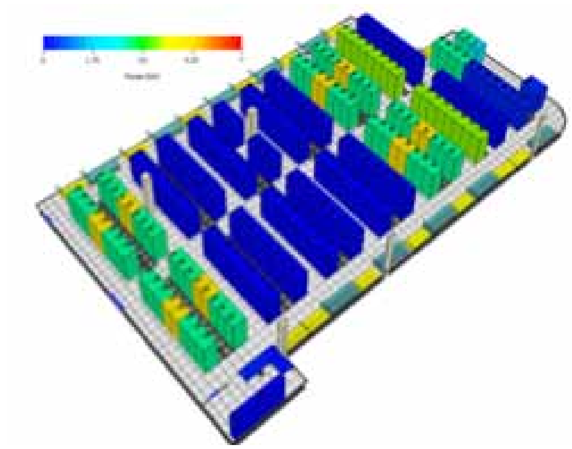

Figure 4 - Power Distribution, 25% Load

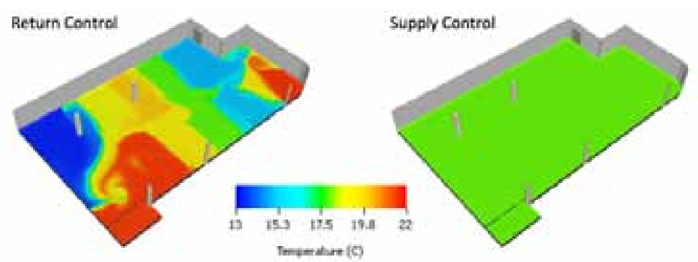

Figure 5 - Floor Void Temperature Variation, 25% Load

What effect does this have on the servers? One method of evaluating performance is to compare the server inlet temperatures with ASHRAE 2008 Class 1 temperature Scale (Figure 6).

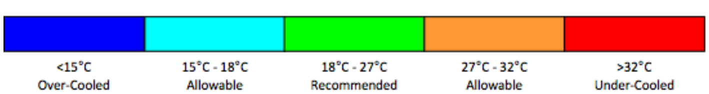

Figure 6 – ASHRAE 2008 Class 1 Equipment Compliance temperature Scale

Figure 7 uses this approach to compare the two control strategies. In the return control case, a significant number of cabinets are being over-cooled to achieve the return condition. In the supply case, no cabinets are being overcooled and the majority are within the ASHRAE accepted range. Interestingly, in both cases, the same cabinets are above the ASHRAE accepted range.

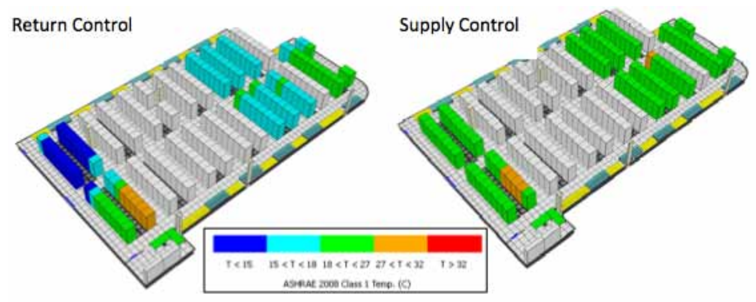

Figure 7 - IT Equipment ASHRAE 2008 Temperature Compliance, 25% Load

The second example is at around 50% occupancy (Figure 8). Here the load is becoming more evenly distributed as the server racks start to fill up. Again, temperature in the floor void varies by up to 9°C in the return case, but is at a constant 18°C in the supply case (Figure 9). As the load has increased, the cooling units in the return case have dropped their supply temperatures. This has resulted in an increasing number of servers being overcooled (Figure 10). In the supply case, the situation remains relatively unchanged, with the majority of servers inside the recommended range.

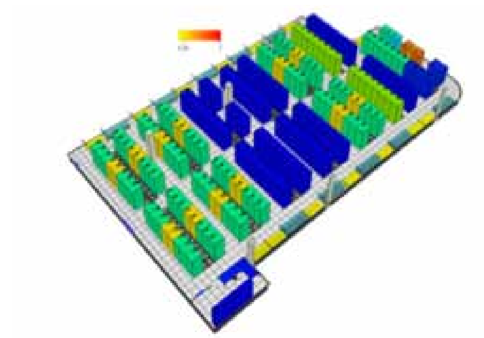

Figure 8 – Power Distribution, 50% Load

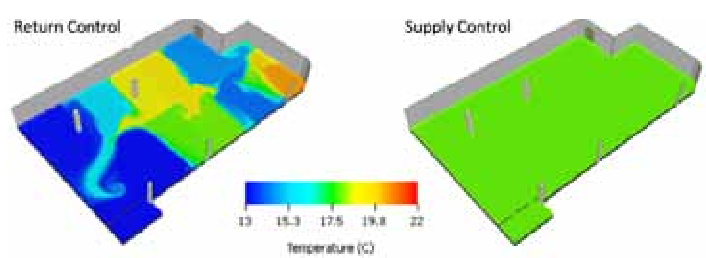

Figure 9 – Floor Void Temperature Variation, 50% Load

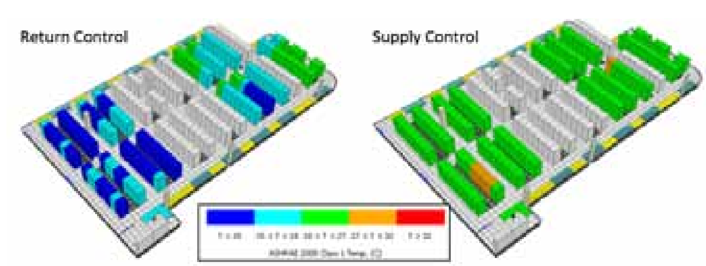

Figure 10 - IT Equipment ASHRAE 2008 Temperature Compliance9, 50% Load

In the last example, the data centre is at the design capacity, with all the cabinets fully populated (Figure 11).

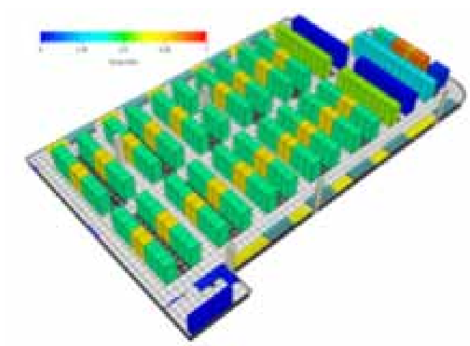

Figure 11 - Power Distribution, 100% Load

Even at full load, the return control strategy still results in a significant variation in floor void temperatures, due to the sensitivity of the control system to sensor location. Once again, the temperature in the supply control case is a constant 18°C (Figure 12). In the return case, the majority of the facility is being overcooled in order to achieve the control criteria, whereas in the supply case the majority of the facility is in the ASHRAE acceptable range (Figure 13).

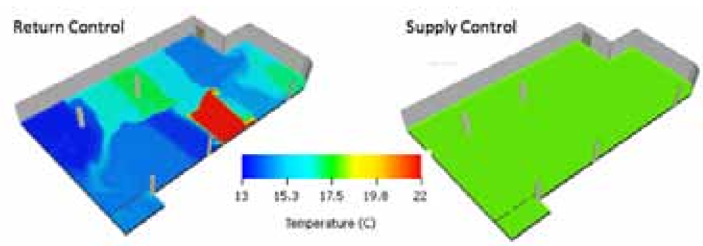

Figure 12 – Floor Void Temperature Variation, 100% Load

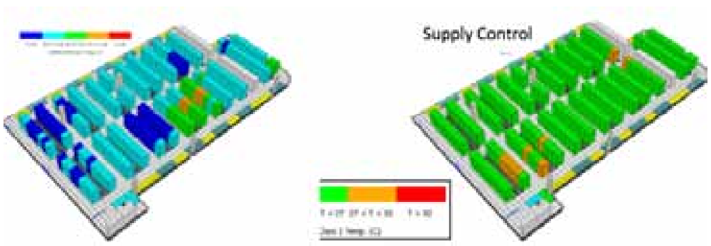

Figure 13 - IT Equipment ASHRAE 2008 Temperature Compliance, 100% load

The Virtual Facility allows for easy export of data to a spread sheet, thus allowing for creation of frequency plots of maximum server inlet temperatures for each case. These plots show how supply control significantly reduces the spread of inlet temperatures throughout the facility. All three graphs show that in the supply control case, there is much smaller grouping around the control temperature, with almost all servers sitting within the ASHRAE recommended range. For all three load conditions (Figure 14 - Figure 16) in the return case, the majority of server inlets are spread between 14°C and 21°C; however, in the supply case, they are more closely grouped between 18°C and 21°C. This over cooling could, to some extent, be offset by using an unconventionally high control set point. However, as the load increases, and especially at full load (Figure 16), the spread of temperature means that increasing the return air control set-point would result in a significant proportion of the equipment receiving air above the recommended range. For the supply control configuration at full load, 99% of servers are within the ASHRAE recommended range, with 90% at or below 21°C. In the return case for the same loading, only 38% of server inlets are within the ASHRAE guidelines, the remainder being overcooled with 92% at or below 21°C.

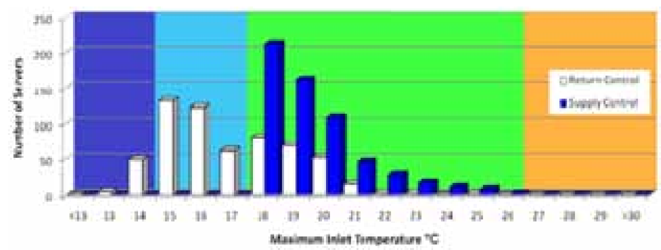

Figure 14 – Frequency Plot Of Maximum Server Inlet Temperatures, 25% load

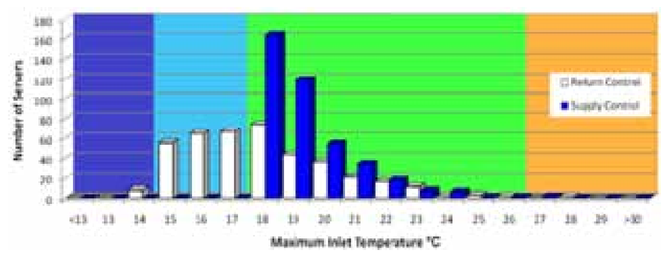

Figure 15 – Frequency plot of maximum server in let temperatures, 50% load

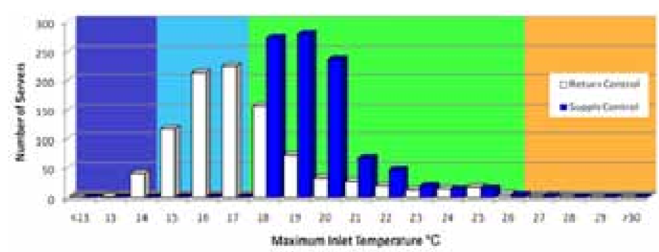

Figure 16 – Frequency Plot Of Maximum Server In Let Temperatures, 100% Load

Money Saving Potential

The chilled water in the Virtual Facility is being supplied at 8°C and returning at 12°C. With the lowest supply temperature in the return case being around 13°C, there is no room for changing the chiller set point without changing the behaviour of the facility. With the supply control strategy, and the resulting constant, predictable temperature in the floor void, there is significant headroom between the chiller set point and the cooling unit set point. This opens the door for potential energy savings by raising the chilled water temperature. At an average UK business rate of 7 penny per kWh, the chiller costs roughly £222,000 per year to run at a set point of 8°C. Chiller efficiency increases as the set point rises, reducing running costs (about 2% per degree [4]). In the Virtual Facility there is the potential to increase the chilled water by up to 7°C and still be able to supply air at 18°C, which would reduce the chiller running cost to £191,000, a saving of £31,000.

Note here, that to achieve the temperature rises and cost savings, the only action taken has been to change the control system. There is still a significant amount of IT equipment that is above the set point of 18°C. By taking action to reduce this by making other design changes, for example some containment of the cold aisle, it is possible that further increases in supply temperature could be achieved. Indeed, in a well designed data centre operating on supply control all server inlets should fall within a degree or two of the control temperature and supply temperature set points of up 24°C or higher can be achieved, with the accompanying savings. Another advantage to supply side control is that it provides a controlled minimum temperature in the data centre. Being able to control the minimum temperature means it can be moved away from the dew point, meaning de-humidification becomes less of an issue. Not running the de-humidifier on the cooling units will reduce running costs further. In temperate climates, humidification is rarely required, and so in these areas new builds could be specified without any humidification/de-humidification equipment, reducing capital costs.

Are There Any Other Knock On Effects?

The Virtual Facility has shown that switching to supply side control has no detrimental effect on the servers. But, are there other knock ons that would prohibit the change? It is often said that raising temperatures in the data centre will affect its resilience. However, a data centre is only as resilient as its least resilient equipment, which in the return control strategy is 6.5°C below thermal shutdown. Controlling the supply temperature to 18°C has reduced this by half a degree to 6°C, however, this reduction is not an artefact of the control system, but of the temperature chosen. It has been shown that rack inlet temperatures have a direct relationship with supply air temperature [1], so if a control temperature of 17°C had been chosen, an increase in resilience by half a degree would have been observed and the same cost savings could have been achieved. Increasing temperatures can increase the amount of power servers draw as fan speeds increase and/or components become less efficient [5]. Fan speeds tend to increase above 22°C and component inefficiencies become significant above 25°C. Switching control strategies has not significantly increased the number of servers above either of these temperatures, so this effect is likely to be negligible.

Finally, conditions for those working in the data centre will have changed, but not necessarily for the worse. Hot aisle temperatures have raised a couple of degrees, from 22-23°C to 25-26°C, which is still acceptable; cold aisle temperatures are at a much more comfortable 18°C, up from 13-15°C.

Conclusion

Using the Virtual Facility as a test environment, it has been shown that there are significant advantages to switching to a supply temperature control strategy. Cooling units are less sensitive to sensor placement and crosstalk between units, giving the facility manager more control over the temperature being supplied to the IT equipment. As the facility grows, a supply control strategy is equally able to cope with the expanding load and keeps server temperatures within a tighter, more controllable band. Lastly, controlling the supply temperature allows for a predictable headroom between the chilled water and supply air temperature. Closing this gap by raising chilled water temperatures will lead to energy savings and reduced running costs. The performance and potential savings will vary from facility to facility and so simulation should be undertaken to understand the benefits on a case by case basis.

References:

[1] Boucher, T., Auslander, D., Bash, C., Federspiel, C. And Patel, C., "Viability of Dynamic Cooling in a Data Centre Environment", Proceedings of ITHERM '04 The Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Las Vegas, NV, USA, June 2004.

[2] Bell, G. C., Storey, B. and Patterson M. K., "Control of Computer Room Air Conditioning Using IT Equipment Sensors" California Energy Commission, 2009.

[3] ASHRAE, "Environmental Guidelines for Datacom Equipment - Expanding the Recommended Environmental Envelope", ASHRAE, Atlanta, 2008.

[4] California Energy Commission, http://www.consumerenergycenter.org/tips/business_summer.html, Aug 2010.

[5] Patterson, M. K., "The Effect of Data Centre Temperature on Energy Efficiency", Proceedings of ITHERM '08 The Eleventh Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, Orlando, FL, USA, May 2008.

[6] Future Facilities Limited "The Effect of Sensor Location on Supply Air Temperature Stability". To be published.

Effect of Different Compounds

and Cooling Enhancements on the Thermal Resistance of a Stacked-Die

As the demands for faster, smaller and multifunctional mobile devices, such as cell phones, PDAs, and netbooks keep increasing, engineers are trying to integrate more functions into a single chip. Such chips are referred as "system on a chip". They may contain digital, analog, mixed-signal and, often, radio-frequency functions – all on a single chip substrate. Stacked multichip design is an effective way to construct such a system on a chip module. It utilizes the third dimension of a chip. By stacking multiple dies on top of each other, it saves the precious board real state. Because the different function die models are in close vicinity, the interconnection performance, signal integrity and assembly are improved, due to short transmission lines. The stacked die package also reduces routing complexity at the next level.

Figure 1 shows a photo of a 3D stacked multichip module. This multichip module has four dies stacked on top of each other. The first die sits on the substrate and it is the largest. The second die sits on top of first die and the two dies are separated by a silicon spacer. Its size is smaller than the first die. The third die sits on top of second die and bonds together with it. The forth die is smallest and is on top of third die. The electrical wires of all dies come down to the interconnectors on the substrate. In most cases, the interconnectors are extended outside through the conventional ball grid array (BGA) on the other side of substrate.

The 3D stacked multichip module packaging also brings many design challenges to packaging and mechanical engineers, especially on thermal management. In a stacked multichip configuration, multiple dies are bonded together by die attaching materials, which normally have low thermal conductivity. The heat flux generated by dies sandwiched by other dies has to overcome extra thermal resistance to be dissipated to ambient or the board.

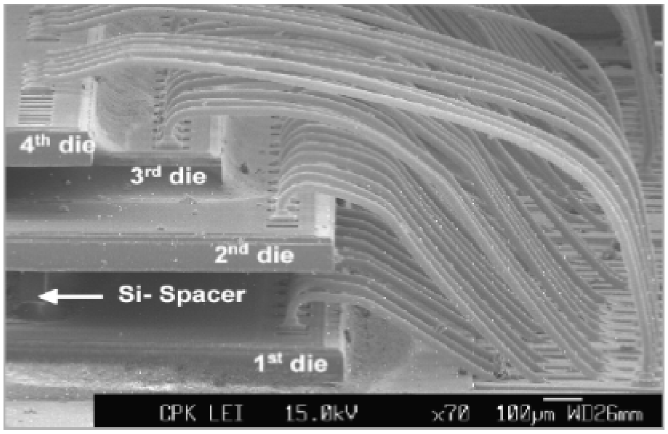

Figure 1. Photo of 3D Stacked Multichip Module [1]

Figure 2 illustrates a typical 3D stacked multichip module. Die attach material is used to bond the chips together. Mold compound is used to encapsulate the dies, wire bonds and substrate.

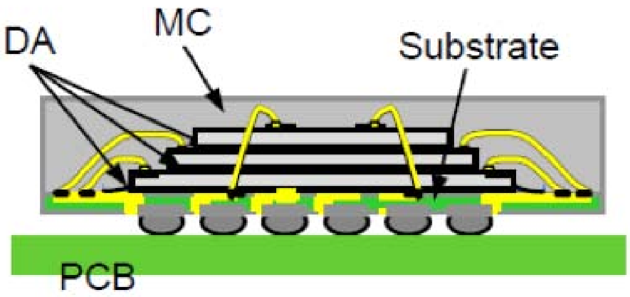

Figure 2. Schematic of Typical 3D Stacked Multichip Module [2]

To investigate the effects of thermal conductivity of die attach (DA), mold compound (MC) and underfill material (UF) on junction temperature of a stacked multichip module, Moon et al. constructed a CFD model to do the parameter sensitivity study. A detailed thermal model for a stacked chip scale package is illustrated in Figure 3. The package has a BGA ball array attached to its bottom.

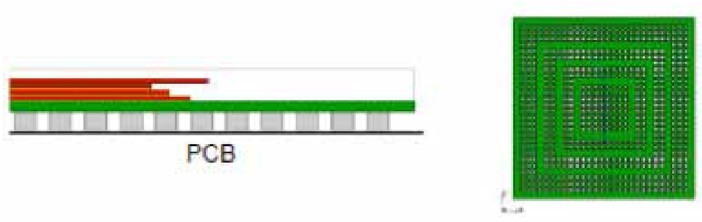

Figure 3. Detailed Thermal Model (Quarter Symmetry) with BGA [2]

To characterize the resistance difference, Moon et al. defined the thermal design power (TDP) ratio as,

$$TDP_{\text{ratio}} = \frac{\theta_{jx,o}}{\theta_{jx}} \quad (1)$$

Where

$$\theta_{jx} = \frac{T_j - T_x}{q} \quad (2)$$

In above equations,

T_j : Maximum junction temperature among stacked dies in the package

q : Total power

x : a for ambient b for board c for package top case

o : Simulation case with the baseline packaging material set

In the baseline CFD study, the model dimensions are shown in Table 1 and the material properties of different parts of the package are shown in Table 2. In most case, typical packaging polymers such as die attaches, mold compound, substrate core, etc., have very low thermal conductivities, as shown in Table 1. In their sensitivity study, the thermal conductivity of die attach (DA), mold compound (MC) or underfill material (UF) varied within from 0.22 to 5 W/m-k.

Parameter	Value
Package body size	8 mm x 10 mm
Substrate thickness	0.210 mm
Die thickness	0.076 mm
Die attach thickness	0.025 mm
Ball height	0.190 mm
Solder ball diameter	0.300 mm
Solder ball pitch	0.800 mm

Table 1. Stacked Multichip Dimensions [2]

Component	Thermal Conductivity (W/m-K)	
Die	120.0	
Die attach adhesive	0.22	
Mold compound	0.88	
Sn-Pb Solder Ball	50.2	
Substrate bulk property	$k_z = 0.4$	$k_{in\text{-}plane} = 46.7$
PCB bulk property	$k_z = 0.6$	$k_{in\text{-}plane} = 46.5$

Table 2. Material Property Assumption for Baseline Model [2]

Figure 4 shows the results of TDP ratio with increase of thermal conductivity of die attach for θ_{ja}, θ_{jb}, and θ_{jc} ,respectively. As the thermal conductivity of die attach increases from 0.22 to 2.0 W/m-k, the TDP ratios increase rapidly. Beyond 2.0 W/m-k, the TDP ratios increase moderately. For θ_{ja} and θ_{jb} , the TDP ratios are around 1.25 at 2.0 W/m-k, which means θ_{ja} and θ_{jb} are 25% smaller than the baseline case.

Figure 5 shows the results of TDP ratios with an increase of thermal conductivity of mold compound for θ_{ja}, θ_{jb}, and θ_{jc} respectively. As is shown in Figure 5, the TDP ratio for junction to case thermal resistance θ_{jc} increases most significantly with the increase of mold compound thermal conductivity. This is understandable, because the θ_{jc} is defined as thermal resistance between junction and case and it is directly related to the mold compound thermal conductivity. However, both θ_{ja} and θ_{jb} increase moderately as the mold compound thermal conductivity increases from 0.88 to 5.0 W/m-k.

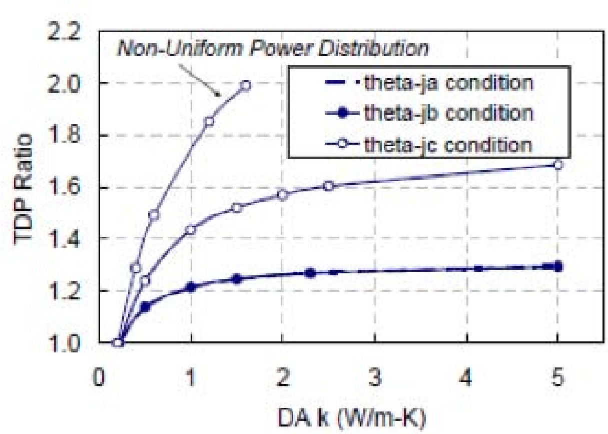

Figure 4. TDP ratios vs. DA thermal conductivity for θ_{ja}, θ_{jb}, and θ_{jc} [2]

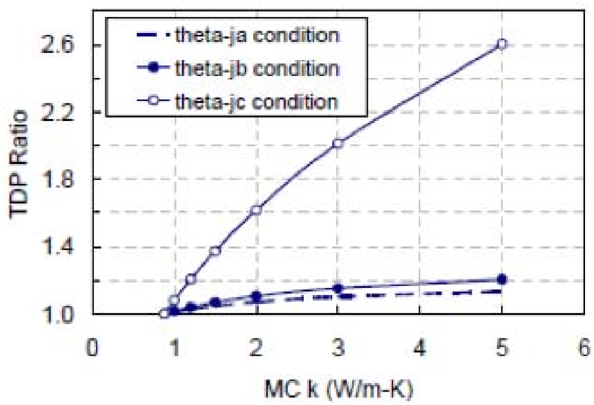

Figure 5. TDP ratios vs. MC thermal conductivity for $\theta_{ja}, \theta_{jb}, and\ \theta_{jc}$ [2]

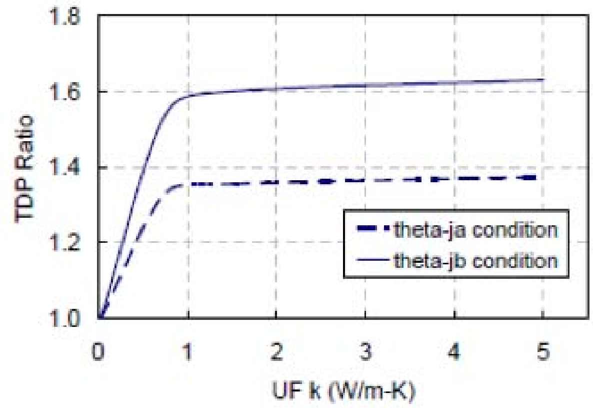

Figure 7. TDP ratios vs. UF Thermal Conductivity for $\theta_{ja}, \theta_{jb}, and\ \theta_{jc}$ [2]

In some mobile and handhold devices, there is no heat sink attached on chip case and no air flow passes through the board. Therefore, most of the heat flux generated by the chip is dissipated outside through board. Under such circumstances, using underfill material to fill voids between the chip substrate and the board helps reduce the thermal resistance between junction and board θ_{jb} . Figure 6 illustrates the usage of underfill for such thermal enhancement.

Figure 7 shows the TDP ratios of θ_{ja} and θ_{jb} with an increase of thermal conductivity of the underfill material. It is not surprising that the junction to board thermal resistance θ_{jb} decreases significantly. When the underfill material thermal conductivity increases from very low to 1.0 W/m-K, the TDP ratio increases by 60% and the θ_{jb} decreases by almost 40%.

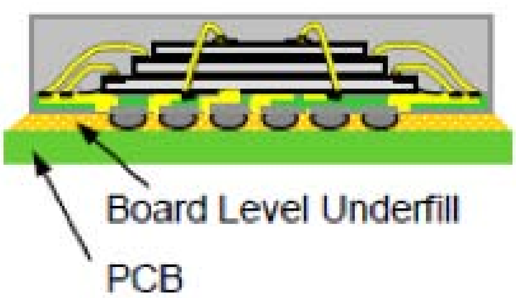

Figure 6. Use of Underfill for Thermal Enhancement [2]

We have discussed the thermal benefits of using better material for die attachment, mode compound and underfill. However, the thermal benefit of using thermally conductive packaging materials varies depending on package architectures and form factors. It is best to combine them with other thermal enhancement techniques. For a stacked multichip package, there are many ways to improve its 3D structure to bring some thermal improvement.

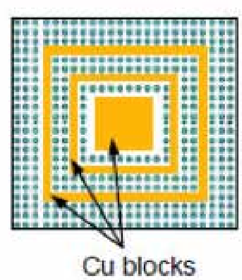

Figure 8. Stacked Chip Package with Copper Strips [2]

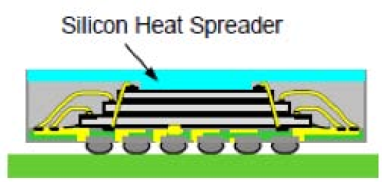

Figure 9. Stacked Chip Package with Heat Spreader [2]

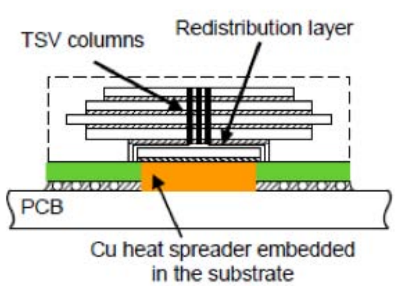

Figure 10. Stacked Chip Package with Thermal Vias and Embedded Heat spreader [2]

Figure 8-10 shows some techniques for such improvement. Heat spreaders can be embedded on top of the chip or under the chip (inside substrate). For the heat spreader under the chip, it can be soldered to the board to help conduct the heat. The thermal vias in vertical direction can smooth temperature between different dies and reduce the stress caused by temperature variation.

The research of Moon et al. [2] has shown that by using better packaging material, the junction temperature of stacked multichip package can be reduced dramatically. However, their study also shows that thermal gain of high conductive material saturates at certain value, normally 2-5 W/m-k. To further reducing die temperature, other creative methods have to be used on the multichip model. Such methods include top and bottom heat spreader, conduction thermal vias, copper thermal strip etc. By working with electrical engineers, die layout can be optimized too, such as placing the high heat flux dies close to the case or the board, integrating a planar silicon heat spreader in the individual die. In most mobile and handhold devices, the chips rely on passive method to dissipate heat away. This requires the engineers to find intuitive ways to solve the thermal management problem.

References:

1. McGrath, J., Sorting Stacked-die, The Challenges of Building SiPs, Advanced Packaging, August 2003.

2. Moon, S., Dizon, M., Chiu, C., and Garcia, E., Optimization of Packaging Materials and Design for Thermal Management in Stacked-Die Packages, ITherm 2006, San Diego, CA, USA, 2006.

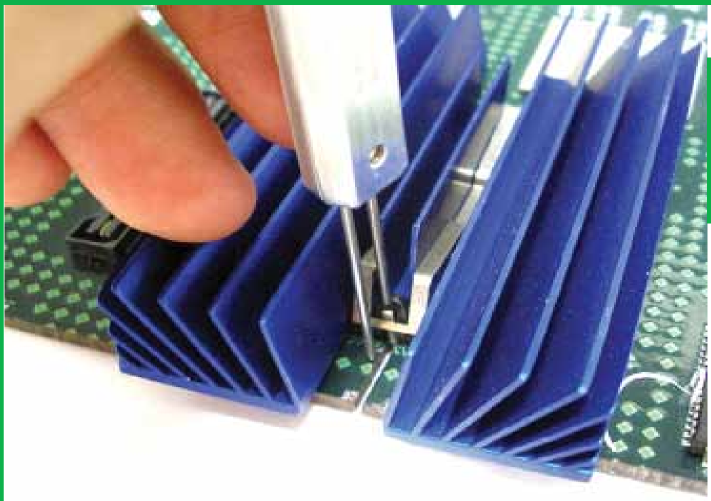

Cooling News

New Products, Services and Events from around the Industry

Highly Thermally Conductive Epoxy System

Master Bond's new EP21ANHT epoxy is designed to help mitigate issues with tightly packed components and miniaturized electronic circuits. The material exhibits thermal conductivity over 22 BTU/in/ft^2/hr/ºF, to provide high performance in demanding uses. The new two-component adhesive, sealant and coating has a convenient 1 to 1 mix ratio by weight or volume and offers room temperature and faster elevated temperature cures. Its coefficient of thermal expansion is 18-20 in/in x 10-6/ºC, with a dielectric strength of >400 volts/mil, and a tensile shear strength greater than 1,000 psi.

Thermal Transfer Stick-on Material

Fujipoly's Sarcon® XR-v-AL is a uniquely formulated silicone thermal interface material that offers the installation flexibility of a putty-like sheet and the performance of a high-end gap filler pad. The material provides thermal conductivity of 6.0 W/m°K and thermal resistance of just 0.04°C-in^2/W (0.29°C·cm^2/W). The non-flammable, peel-and-stick interface material is only 0.11mm (0.0043 in) thick and is quickly applied to components or heat sinks as easy as an adhesive label. The material is available in sheet form on an easy-release carrier in sizes up to a 50 x 50 mm (2 x 2 in). Once applied, this thermal interface material exhibits an extremely low contact resistance and efficiently transfers heat from its source to a heat sink or spreader.

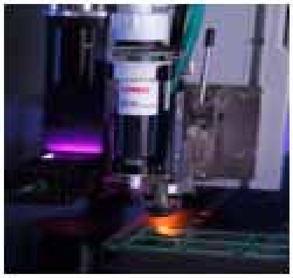

Underfill Has Strong Flow Characteristics, Thermal Cycling Reliability, Low CTE

ME-555 underfill encapsulant from LORD Corporation is a high purity, semiconductor grade epoxy underfill material for the encapsulation of flip chip devices. It is formulated to reduce warping and with the necessary structural strength to handle over-molding processes. Designed with optimum surface tension and viscosity to achieve full coverage without producing bulky fillet or creeping on top of the device, LORD ME-555 will flow consistently without voids, while maintaining a fast flow rate. It can be used under CSP/BGAs and small dies with stand-off heights as small as 25 microns.

Thermally Conductive Epoxy Helps Remove Heat

New 109-12 low-stress thermal epoxy from Creative Materials, Inc., works with high-heat electronics, such as CPUs, power diodes and power transistors, as an interface between heat-producing components and their heat sinks. Other uses include die attachment, PCB fabrication and LED attachment. The epoxy helps today's miniature devices dissipate heat quickly to both extend life and prevent premature component failure. Low viscosity reduces air entrapment while also reducing the quantity of material that must be applied. Designed to keep thermal resistance as low as possible, one-component 109-12, exhibits minimal shrinkage during curing while offering high thermal cycling resistance.

Easy-to-Apply Thermal Grease

Silicone-based thermal grease from Fujipoly provides a thermal conductivity of 2.6 W/m°K, while exhibiting minimal bleed and evaporation characteristics. The unique non-flowing consistency of Sarcon® SG-26SL makes the grease ideal for power converter and high-performance CPU applications that have bond lines as small as 1 mil. The material's low thermal resistance provides for efficient transfer of unwanted heat from board-level components to nearby heat sinks. The compound is easy to apply as well as remove during assembly and rework processes and delivers consistent performance over temperatures that range from -67°F to 401°F. Packaging options include pre-filled syringes and bulk jars for easy stenciling, manual or automated application.

Thin Sensor Measures Temperature and Air Velocity

A flexible, robust candlestick-design sensor from Advanced Thermal Solutions can be used to measure both temperature and air velocity in order to characterize thermal conditions in electronic systems. The candlestick MS 1000-CS-WC sensor is narrow and low profile to minimize disturbance of heat flow in the test domain. Strong base-and-stem design allows continuous repositioning and reading during the testing process. The temperature range is -30°C to 150°C (±1°C) and the air velocity range is 0 to 51 m/s (10,000 ft/min) (± 2%). The sensor's all-plastic design prevents shorting of the electronics.

Upcoming Webinars

Register

Thermal Management of Consumer Electronics

Consumer electronics products are featuring increasingly more sophisticated processors, DSPs and FPGAs. They are even being deployed in military and other hostile environments. These sophisticated technologies allow a wide range of features, but they have higher thermal management needs. Attendees will learn the key thermal issues to consider, including the importance of spreading resistance in components, and system thermal management. They will learn thermal stress factors affecting enclosed cabinets and how to manage direct solar load.

July 28, 2011 at 2:00 p.m. (EST)

How to Decide When You Need to Deploy Liquid for Your Thermal Management

Thermal engineering professionals are becoming more confident employing liquid cooling solutions at device, system, rack, and data center levels. But how do you decide when you need to cool with liquid? How do you find the right liquid cooling system for your application? This webinar provides the best practices for implementing a liquid cooling system at the device level.

August 25, 2011 at 2:00PM (EST)