ADVANCES IN COOLING ELECTRONICS WITH CFD

Dr David Tatchell (Flomerics Ltd., UK); Drs. John Parry, CEng & Ian Clark (Mentor Graphics Mechanical Analysis Division, UK)

Dr David Tatchell, Chief Technical Officer (retired)

THEME

Emerging Issues: Current Industrial Applications & Future Industrial Needs.

SUMMARY

The staggering advances in functionality and miniaturisation in electronics equipment over recent decades are too obvious to mention. What is perhaps less clear is the contribution that CFD has made to making this “electronics revolution” possible.

The managing of the heat dissipated in electronics has long been recognised as a major challenge. As functionality has increased, following Moore’s Law, the associated heat dissipation has escalated, to the extent that this has been recognised as a potential limitation on the pace of electronics development. The challenge is to prevent overheating, and failure, of critical components via appropriate cooling strategies. By providing a means of predicting and improving the thermal performance of electronics equipment, CFD-based software has supplied an invaluable design aid in meeting these challenges, and has, over the last two decades, become an essential component in the design process for virtually all types of electronics equipment.

“Electronics cooling CFD” has, however, evolved rather differently from other “mainstream” CFD applications. Instead of utilising established commercial general-purpose software packages, electronics cooling has given rise to the development of a number of specialist application-specific software packages. There are a number of factors leading to this.

Firstly, the software itself has had to be designed and adapted to suit the specific user needs and design environment in the electronics industries. Users are mechanical designers rather than CFD analysts - they work in a fast-moving, multi-disciplinary, design environment – and they need to collaborate with electronic designers using EDA software and with other mechanical designers using MCAD software. The software is expected to contribute at all stages of the design process, from concept, through design exploration and

ADVANCES IN COOLING ELECTRONICS WITH CFD

optimisation, to final verification. These diverse needs have major implications for the software, its user interface, data management, and integration.

Secondly, there are a number of technical challenges that are specific to electronics-cooling applications. The inside of electronics equipment is a complex assembly of many solid objects (printed circuit boards, electronics packages and devices, cabling, fans, heatsinks, etc.). The air flow is confined within narrow regions between these solid objects. As well as convective transport within the air, conduction within the solid objects (which can have extremely complex internal structures) can be critical. Analyses can involve large numbers of such objects (sometimes thousands), and extreme disparities in scale (from meters to micron scale).

All of these provide particular challenges for the CFD, which needs to be adapted accordingly. Special practices for model generation, geometry representation, gridding and grid generation, and turbulence and near-wall modelling have been developed to address these needs. Additionally, electronics systems typically contain objects for which a full, detailed CFD representation is impracticable (such as chip packages, fans, heat pipes etc.), for which special “behavioural models” need to be developed.

These and other specialised industry needs have driven the development of electronics cooling software over the last two decades and more. Special-purpose CFD-based software, optimised for electronics thermal applications and with industry specific user input and control, has been developed to satisfy these needs, and is now widely used in virtually all major electronics companies as a routine part of the product design process.

As electronics becomes ever more widespread, more complex, and more compact, and the industry becomes more competitive, the importance of electronics thermal design continues to become more critical. This is leading to new needs and challenges - to broaden the usage of electronics cooling software within the design teams by addressing a wider range of “user personas” - and to deepen the integration and interoperability with the other software tools used in electronics design processes.

KEYWORDS

Electronics Cooling, Thermal Design, Concurrent CFD, Upfront CFD, History, Technology, Advances.

ADVANCES IN COOLING ELECTRONICS WITH CFD

1: Historical Context

Managing heat dissipation has for some time been recognised as one of the main factors limiting the pace of development in electronics (The Economist, 2003). CFD analysis software, applied to the thermal design of electronics equipment, has, over the last two decades, made a major contribution to addressing this critical issue.

The impact is perhaps best illustrated by a benchmark study (Aberdeen Group, 2007) showing that CFD enables electronics thermal design verification to be completed in one-third of the time that would otherwise be required. By achieving these time and cost savings, CFD is evidently making a major business contribution to addressing a critical problem in a trillion-dollar global industry. Another major success story for CFD!

But electronics cooling has evolved rather differently from other mainstream CFD applications. Rather than utilising the established general-purpose software packages which dominate the commercial CFD market, electronics cooling has given rise to the creation of a number of specialist application-specific CFD packages, which are now routinely used as part of the electronics thermal design process.

This paper considers the challenges of cooling electronics equipment, and of applying CFD to the design processes – and attempts to explain why this application has evolved separately from mainstream CFD applications.

Why Electronics need Cooling

The main heat sources in electronics are the semiconductor (usually silicon) chips within the electronics packages located on circuit boards housed within the equipment.

The challenge in cooling electronics arises from the temperature sensitivities of these heat-generating chips. The mechanisms are complex, and difficult to quantify precisely – but it is well known that overheating of the chips causes premature failure – and of course failure of one chip can often disable the complete equipment. The higher the chip temperature the earlier and more certain the failure. So the thermal design requirement for electronics equipment is often expressed as “maximum allowable junction (i.e. chip) temperature” for each critical component. Generally these critical temperatures are relatively low – typically 85 to 150 deg C. The challenge in the thermal design of electronics is to ensure that all critical components are operating below their specified temperature limit.

The well-known Moore’s Law exponential increase in silicon chip functionality with time has its corollary in a corresponding increase in chip heat fluxes. For example Figure 1 shows the increases in heat flux over five decades for CMOS and Bipolar modules used in mainframe computers. Other classes of modules show similar trends. It is interesting to note that in some cases chip heat fluxes have now reached levels experienced in more-obviously "severe" heat transfer applications such as rocket motors or nuclear blast (Hannemann, 2003). These are not trivial heat transfer problems!

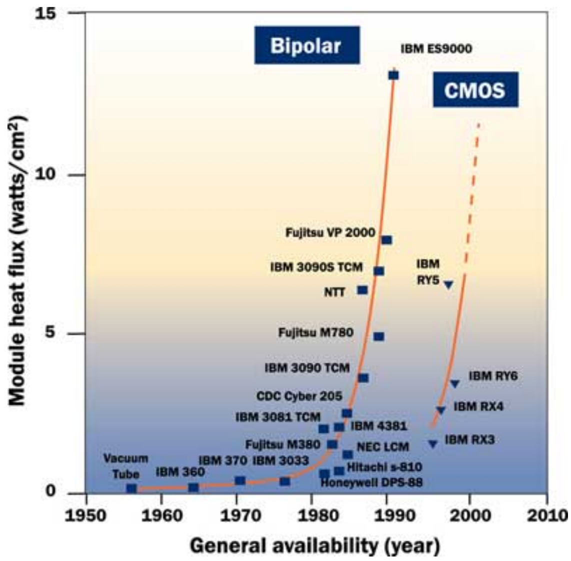

Figure 1: Growth of Bipolar and CMOS Module Heat Flux

While the thermal densities were considerably lower back then, even by the mid-1980s thermal problems were beginning to be a major issue in electronics design. The mechanical engineers (MEs) tasked with cooling the electronics (along with other aspects of mechanical design) were finding that the rules of thumb and hand calculations that they had relied on thus far were insufficient. Equipment was increasingly failing thermal testing at the prototype stage, with financially-damaging consequences in extended product production cycles and delays in product release. Software solutions, to complement experience and build-and-test, were beginning to be sought urgently.

The Need for CFD

The first electronics thermal design solutions focussed on the thermal behaviour of printed circuit boards, rather than the equipment as a whole. While from a thermal standpoint this was a mistake (as explained later), it was perhaps natural in an industry increasingly focussed on the use of EDA (electronics design automation) software for all "electronics" aspects of design. A major component within EDA systems is the PCB design suite, in which the layout of components and the routing of circuitry on the PCB can be created,

ADVANCES IN COOLING ELECTRONICS WITH CFD

displayed and analysed in great detail. Initially thermal solutions focussed on this packaging level.

The PCB thermal solutions that first became available in the mid to late 1980s were either embedded within EDA systems (AutoTherm from Mentor Graphics, Thermax from Cadence) or closely linked (PCBThermal from Pacific Numerix). They offered essentially 2D conduction solutions, for heat transfer within the board itself, and extremely simplified single resistor models for the components located on the board. Convective heat loss from surfaces was approximated via empirically-based heat transfer correlations and presumed global air temperatures. Radiative heat loss from board assembly surfaces was estimated to some presumed global ambient.

As an outcome, PCB thermal solutions provide fields of temperature distribution over the board and discrete component junction temperatures, as shown in Figure 2. From this the designer can assess whether the layout is "thermally acceptable", and, if not, experiment with improvements such as component relocation, or addition of heatsinks to critical components, again represented very simply.

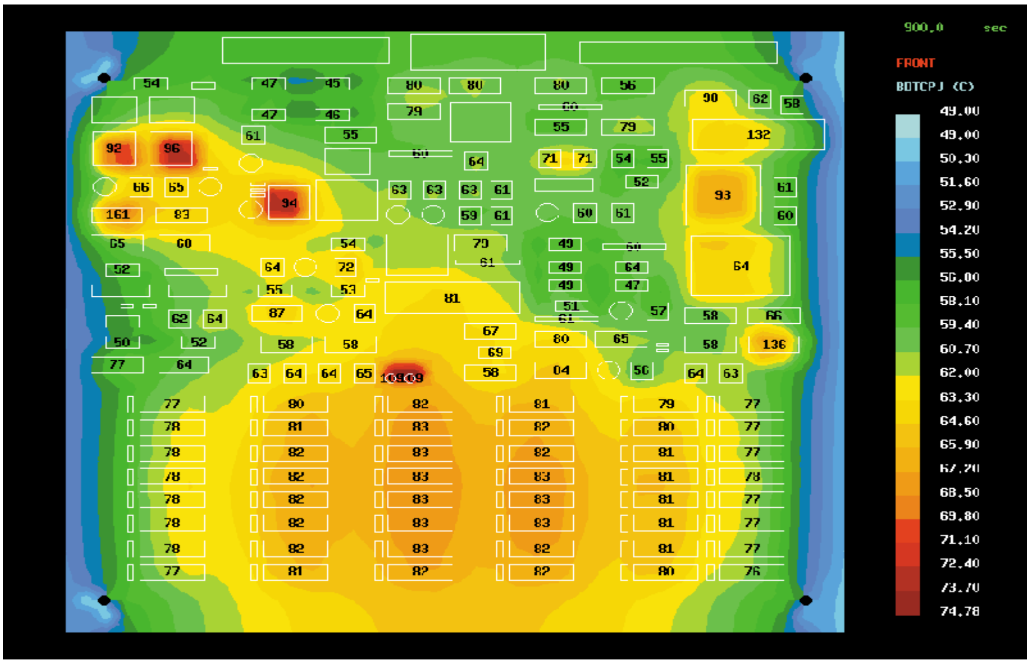

Figure 2: PCB Thermal Solution displaying junction temperatures.

However, in practice such solutions, on their own, have proved to be of limited value. First of all, the link to EDA systems proves a major drawback. The MEs tasked with thermal design do not generally use or have access to the corporate EDA system – and the electrical engineers (EEs) who do, have little knowledge of, or interest in, thermal issues. And, secondly, the usefulness of the results is

ADVANCES IN COOLING ELECTRONICS WITH CFD

limited by the major assumptions and simplifications implicit in the PCB thermal approach – the most critical being the complete neglect of the airflow distribution within the equipment, which (in reality) gives rise to variations in air temperature and surface heat transfer coefficient, the knowledge of which are often crucial in identifying and resolving thermal problems in any air-cooled electronics.

Consequently, by the late 1980s electronics thermal designers were beginning to seek means of predicting and understanding airflow distributions in electronics equipment. The obvious solution was provided by CFD.

The natural first step was to apply the commercial general-purpose CFD codes then available to the electronics cooling problem. IBM Poughkeepsie introduced a number of licences of PHOENICS (from CHAM Ltd), and DEC began using Fluent (from Creare Inc). These had somewhat limited success. Burdick (1991) commented on the experience at IBM Endicott in the late 1980s:- “A small number of engineers attempted to use commercially available general-purpose finite-difference CFD programs at this time but the result of several months of activity was usually fruitless”.

It was only when specialist CFD software, specifically tailored for electronics cooling, became available – most notably FloTHERM from Flomerics, in 1989 – that CFD really began to penetrate the electronics cooling market. And software of this kind has continued to lead the market ever since.

How is Electronics Cooling different from other CFD Applications?

It is interesting to consider why the electronics cooling market has evolved in this way, largely separately from mainstream CFD.

Perhaps the first consideration is the typical user profile, and the design environment of which thermal design is an integral part.

As noted earlier, the users are generally the MEs responsible for all aspects of the physical design of the equipment – that is everything beyond the electronics design, which typically culminates in the PCB layout. The ME is then responsible for the enclosure, appropriate location of the PCBs and other components, and for ensuring structural integrity and the safe, reliable operation of the equipment. In other words, the ME needs to be concerned with a range of issues of which cooling/thermal design is only one (though often a crucial one).

All of this takes place within a fast-moving and highly competitive environment. In electronics the complete design cycle from concept to first customer ship are usually far shorter than in traditional manufacturing industries – in some sectors now as short as nine months. And it is well established that, all else being equal, the first to market with a new-generation product is in a strong position to establish market dominance – so delays in product release of even a few weeks can have a severe financial impact.

Electronics-cooling therefore needs software that can be used speedily and reliably as an integral part of a fast-moving, complex design process. The users are not experts in CFD or fluid dynamics – and therefore need software which requires no user adaptation and minimal user knowledge of CFD concepts, and avoids potentially time-consuming operations such as sophisticated grid generation.

Clearly general-purpose CFD software is far from ideal in satisfying these needs – special-purpose software, optimised for electronics thermal applications, and with industry-specific user input and control is required.

It was largely these considerations that led the electronics cooling CFD market to become established as it did during the early 1990s – and they remain crucial today. And, as the market has evolved since, a number of additional, specific technical requirements have emerged that have reinforced the trends towards specialist electronics-cooling software.

These include the sheer number of discrete objects that make up the analysis problem (sometimes numbering thousands), and the wide range of physical scales that can be encountered in a single analysis, from chip (often a fraction of a millimetre thick) to enclosure (sometimes one metre or so). These lead to rather different requirements in, for example, problem definition, meshing, physical modelling, and solution control as compared with mainstream CFD applications.

Additionally, the complex flow and thermal characteristics of many of the component parts that make up an electronics system provide special challenges. Objects which cannot be represented in full computational detail within the CFD analysis (such as chip packages, fans, heat pipes, thermoelectric coolers, etc.) need to be represented reliably and efficiently, without the user needing knowledge of their detailed design, or of how to model them. Providing parameterised behavioural models of such classes of components has become a major factor in distinguishing the electronics-cooling CFD market from general-purpose CFD.

The remainder of this paper covers in more detail the technical challenges of electronics cooling CFD applications and the solutions adopted, and considers the impact and implications of current and future trends in electronics equipment and electronics design processes.

ADVANCES IN COOLING ELECTRONICS WITH CFD

2: Requirements and Technological Challenges

We have seen from the way Electronics Cooling CFD has evolved that there are a number of high-level requirements for the use of CFD for this application space. As a result there are corresponding technology challenges that, although not unique to electronics cooling, are particularly pronounced for this application.

Over the years we have seen a trend away from centralized centers of thermal excellence in favour of thermal design being undertaken as distributed activity, undertaken by disparate product design teams. Often we observe that EC CFD is used by just one person on a part-time basis, as part of a small multidisciplinary design team, alongside their other responsibilities, for example the mechanical integrity / reliability of the product. Hence we observe the need for software tailored to rapid model building for electronics cooling applications to support the needs of such individuals, as well as thermal experts. We refer to this as <u>User Interface Versatility</u>.

CFD is needed right across the design flow, from concept design to final product validation. Thermal design is an intensive design activity, particularly in the conceptual / architectural design phases, where multiple scenarios need to be investigated to define the best cooling strategy before the product design starts to ‘solidify’ and the range of available cooling options become more limited and expensive to implement. In early design, much of the design information is quite speculative, with daily updates of geometry, material and boundary condition information needing to be input directly into the EC CFD software. The software must therefore provide a rapid design environment, in which geometry can be changed and the simulation results rapidly regenerated. We refer to this as <u>Design Process Centricity</u>.

Once a cooling solution is chosen, the solution needs to be optimized for cost within other design constraints. This is an on-going process, with the cooling solution being refined and re-optimized during the detailed design phase as more information becomes known about the design, and as the design has to be adapted to accommodate other considerations, such as timing, EM, SI, etc.. For example, each heatsink should to be re-optimized every time there is a change to the local air flow environment, or the position of the component on which it is mounted. As a result, for EC CFD it is particularly important that the simulation environment supports automated design of experiment and optimisation techniques that can accommodate geometry changes.

In general, being able to correctly capture trends is of higher importance than simulation accuracy, particularly on coarse meshes, in order to facilitate rapid design optimization. However, design margins are shrinking, and over-design is becoming less tolerated, particularly when products are produced in volume for price-sensitive markets like consumer electronics. Consequently the techniques used cannot compromise the ultimate accuracy achievable in late design.

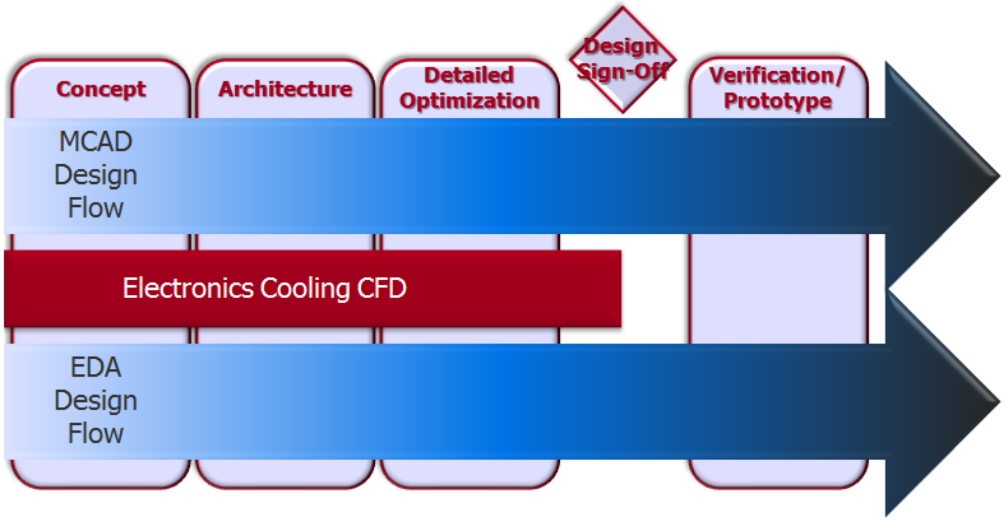

Figure 3: Electronics Cooling CFD positioning with MCAD & EDA design flows

In addition to the above there are a set of technical challenges that are unique to electronics applications, which we will now focus on.

Geometry Capture

During detailed design, the geometry comes from both the EDA and MCAD design flows. One particular challenge is that historically EDA systems deal with 2D representations of the electronics, as both IC and PCB design are done using schematics. PCB design tools require only the component layout, and often do not contain even the most basic geometric information about the components such as component height. Detailed information about the internal geometry of the package is typically unavailable.

Scale Disparity

Miniaturization resulting from Moore's Law has caused an increasing disparity of length scales, between the size of the physical product and the size of the internal components and circuitry. Typically meter to micron scale geometry has to be accommodated within the same model. The presence of small gaps, e.g. in the casing can also have a profound effect on the cooling of the electronics.

As a result, scale disparity continues to increase over time supporting the requirement for behavioural models when the geometry cannot be represented directly within the simulation as is usually the case with PCB traces on multilayer PCBs, and compact thermal models (CTMs) for IC packages to avoid having to model the internal geometry, which is often unknown.

Missing Data

This brings us on to another challenge we list here as unique to electronics cooling applications, that of missing data. Whereas material property data is absent from MCAD systems, so CFD simulations in general suffer this problem, in the case of electronics cooling applications the challenge is more severe, as systems are essentially constructed from many components from many different suppliers, the thermal characteristics of which are typically not well understood. These include IC packages, PCBs heat pipes, fans, Peltier devices, etc.

As noted above, the geometry comes in part from the EDA system, which often do not include any information on the materials being used, and an added complication arises from the assembly of electronics systems, where thermal interface materials (TIMs) and gap pads are used to maximize the thermal contact between different parts of the system in order to implement an effective cooling solution.

A final challenge is the need for operational power information for the active components in order to predict system temperatures under operational conditions, which vary as a function of the product’s usage. Increasingly design for steady-state operation at maximum power, leading to significant over-design is no longer tolerated (Huang 2009, Intel, 2011). Increasingly transient simulations are needed to ensure reliable operation and minimize overdesign.

Flow Regime

In highly-cluttered electronic systems air is forced through channels, containing all sorts of protuberances that induce low Reynolds Number transitional flow. However this wall-induced turbulence is not self-sustaining and the flow would be laminar if the channel were smooth. Turbulence modelling is therefore a particular challenge. Due to the large number of flow channels, objects, etc. combined with a large system residence time, providing a sufficiently fine mesh to perform LES-based modelling is completely impractical within this fast-paced design environment.

Until recently the practicality of using standard two-equation RANS models has been questionable, with the industry favouring zero-equation ‘effective viscosity’ models to impose an estimated turbulent viscosity, since the low mesh densities often used would cause one and two-equations models to predict less realistic turbulent viscosity values than can be estimated based on empirical data and knowledge of the bulk flow velocity.

ADVANCES IN COOLING ELECTRONICS WITH CFD

A key issue with one and two-equation models is the need to refine the mesh near to the surface when used with standard, generalized, and scalable wall function treatments (log law, van Driest, 1/7th Power Law, etc.) to provide a y+ value of ~30 for the near wall cell, with a low mesh size inflation rate out to the core flow. In electronics applications, boundary layers start at the leading edge of components, PCBs, heatsink fins, etc. resulting in a large number of very thin boundary layers to resolve within the system, so the standard advice on y+ simply can't be followed. Consequently LVEL remains the model of choice. However, the recent application of immersed boundary treatments to electronics cooling applications overcomes this drawback.

Mesh Generation

Although generic to CFD, mesh generation for electronics cooling applications presents a challenge due to the sheer number of solid-fluid and solid-solid surfaces that need to be captured, and as a consequence of the need for fully automated optimization including geometry changes, the meshing must also be fully automated with no manual intervention beyond pre-defining the required mesh sizes before meshing is started.

A fortuitous outcome of EDA systems designing components and PCBs in 2-D, with no aesthetic requirements on the unpackaged electronics, is that electronics tend to contain large numbers of Cartesian-aligned Cartesian objects, so Cartesian-based grid systems are the natural choice for this application. Increasingly however, size constraints are forcing electronics designers to angle components on boards, insert DIMMs at an angle, and design heatsinks with non-Cartesian profiles.

Use of simple Cartesian meshes with grid lines that 'bleed' out from an object to the edges of the solution domain have long been recognized as inadequate as they quickly lead to unacceptable mesh counts as increasing geometric detail is added to the model. As a result, the use of locally refined Cartesian overset grids to refine the mesh within and around objects has become prevalent, allowing either porosity or voxelization treatments to approximate non-Cartesian and non-aligned Cartesian objects with acceptable accuracy in many cases.

As the amount of non-Cartesian geometry present within electronics systems has increased, so has the need for more sophisticated meshing strategies. Over recent years we have seen increasing use of Octree meshes with MCAD-embedded CFD in early product design across a range of industries and applications, where the product design process is built upon the company's MCAD system. In electronics, design processes vary considerably from company to company, and embedding CFD within the MCAD system may not facilitate its use, as often much of the early design work will be done outside the MCAD environment, and the design process may be centred around the company's EDA flow. There is therefore a need to provide the simulation approach used in MCAD-embedded CFD within a stand-alone product.

Hardware Environment

Traditionally electronics thermal design has been undertaken alongside the electronic design. The use of HPC infrastructure for electronics cooling CFD has been far less than we observe in other industry sectors, such as automotive, where HPC has facilitated the use of LES to undertake 'high fidelity' CFD to address difficult aspects of the product design, such as aero-acoustics. In electronics cooling applications there are not the same benefits to be gained, as increased simulation precision does not translate into improved product quality. The quality of the simulation model is limited by far greater uncertainty in the input data. To date, good scalar performance with reasonable scaling up to say 8-16 cores has matched the market requirements. Good scaling for a limited number of shared memory nodes is likely to remain the target for hardware performance. The hardware environment may change away from desktop to cloud-based computing, as this will greatly facilitate design space exploration by the use of numerical design of experiment techniques. It will be interesting to see how CFD and other CAE vendors cope with the changes in licensing model needed to support cloud deployment.

Increased Accuracy

As a consequence of design margins shrinking, the need for simulation accuracy is increasing. This however, does not translate directly into a need for higher fidelity CFD. Indeed, since the early 2000s, clock speeds have not increased, capping die-level power density, as shown in Figure 4, with power increases occurring at higher levels of packaging, e.g. the PCB.

What has this to do with accuracy? The answer is that the allowable temperature rise from ambient to junction is not increasing, but as power densities increase within the package, PCB, etc. the proportion of the temperature rise that occurs in the air is diminishing. Put another way, the importance of modelling the conduction within the solid structures is increasing. This explains the emphasis we see being placed on MCAD integration (e.g. for heatsink design) and perhaps more importantly EDA integration, to accurately capture the copper content and distribution on PCBs, effects such as Joule heating in traces, power and ground planes, and on accurately measuring the thermal conductivity of TIM materials, particularly the Type I and Type II materials that are not well suited to being measured in ASTM D5470 based equipment. This is an area where considerable research efforts have been directed over the last decade, because the contribution of thermal interface resistances to the overall temperature rise is increasing (Mentor Graphics, 2012).

ADVANCES IN COOLING ELECTRONICS WITH CFD

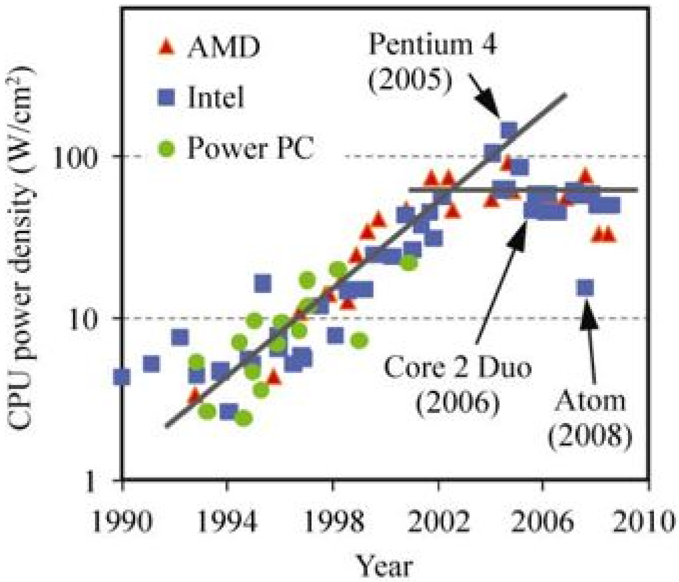

Figure 4: CPU Power Density variation over time (Pop, 2010)

3: Software Advances

User Personas

Traditionally, CFD-based thermal design software has targeted engineering analysts with specialised knowledge of thermal design and the use of CFD techniques. These engineers still form a core group of important users in electronics companies today. The need to perform CFD-based thermal design, however, has now broadened out to include other user personas: electrical engineers, general mechanical design engineers, industrial designers and marketing engineers.

The consequence of this is that the requirements for designing a software solution have become more challenging in terms of: User Interface (UI) design, geometry and attribute pre-processing, interoperability with other MCAD, CAE and EDA software, obfuscation of CFD terminology and functionality, result post-processing and meshing/solver performance.

Software Development

The processes involved in the creation of software have also had to evolve to support the complex requirements of the users being targeted. The most commonly-used approach to software development until relatively recently was documentation heavy, involving the creation of highly-detailed requirements and functional specifications, followed by extremely formal implementation

ADVANCES IN COOLING ELECTRONICS WITH CFD

and quality assurance. The so-called Waterfall methodology originated from the construction and manufacturing industries and was simply used for software development originally as nothing else existed at the time (Bennington, 1983). Its primary benefits are project traceability and a disciplined approach to save on project time and cost, however, its inherent inability as a process to adapt to change makes it a less than optimal method for developing software that is focussed on addressing rapidly evolving and new requirements.

As a result, best-in class companies are increasingly adopting Agile methods and processes, which naturally allow software teams to be more nimble in adapting rapidly to change (Beck, Kent et al., 2001). Agile is essentially a group of software development methods that are iterative and based on teamwork, collaboration and shared responsibilities as a team. One popular method is Scrum, in which development takes place as a series of 2-4 week 'sprints'.

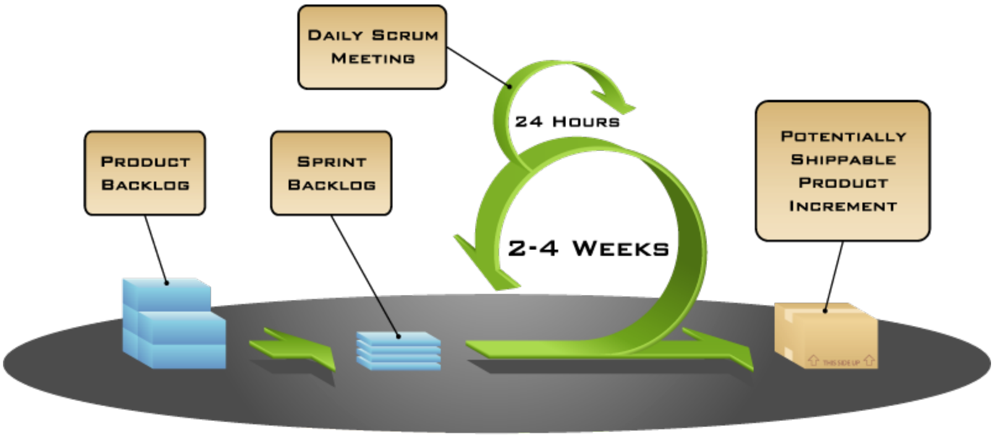

Figure 5: Agile Scrum Process (Mountain Goat Software, 2005)

For the electronics industry, the rate of technology change is accelerating and from a thermal perspective becoming ever more complex. In automotive, there has been an explosion over the last decade in the sophistication and range of electronics systems in vehicles (Fig. 6).

The same applies for gaming, medical, consumer and telecommunications, where the increasing demand for smartphones and tablets, combined with miniaturisation, requires a level of sophistication and innovation in terms of overall design, materials science and cooling techniques unsurpassed in the industry to date.

Only an Agile approach to software development in support of these trends can result in nimble, flexible software that evolves to meet the ever-changing needs of the industry.

ADVANCES IN COOLING ELECTRONICS WITH CFD

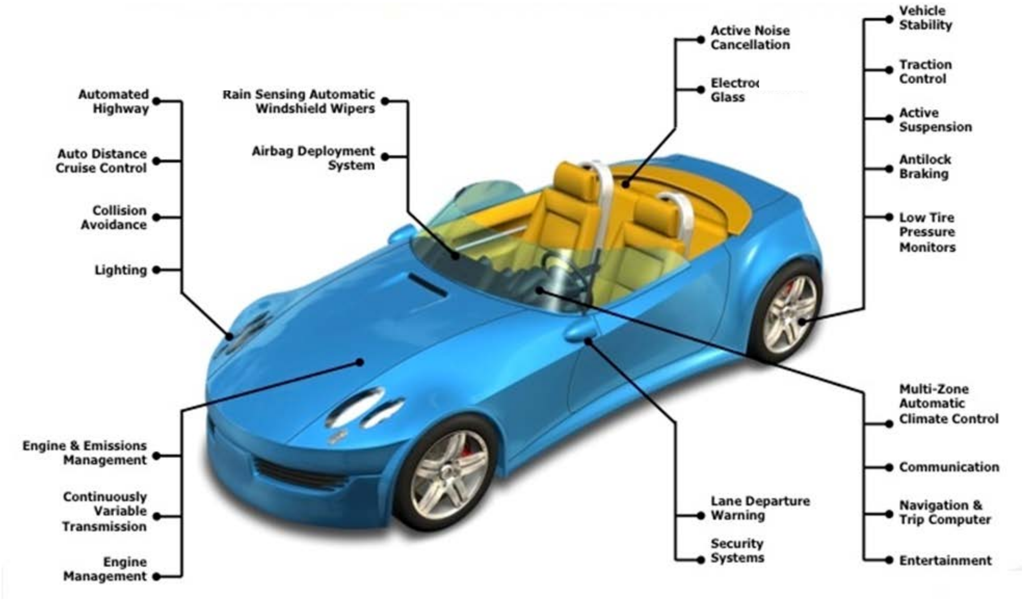

Figure 6: Automobile Electronics Systems Complexity (Chipestimate.com, 2011)

Software Evolution

So what does this actually mean in terms of software design? The primary issue is the level of user interface sophistication and recognition that one-size may not fit all. For many years, software user interfaces assumed an analysts mind-set that limited their usage and spread within electronics companies. As noted earlier, FloTHERM from Flomerics was the first to reach the market with the sole purpose of meeting the needs of the electronics thermal designer and it remains the market leader in the field today (Gary Smith EDA, 2010). In analysing the possible user base today in electronics companies a pyramid of personas can be defined:

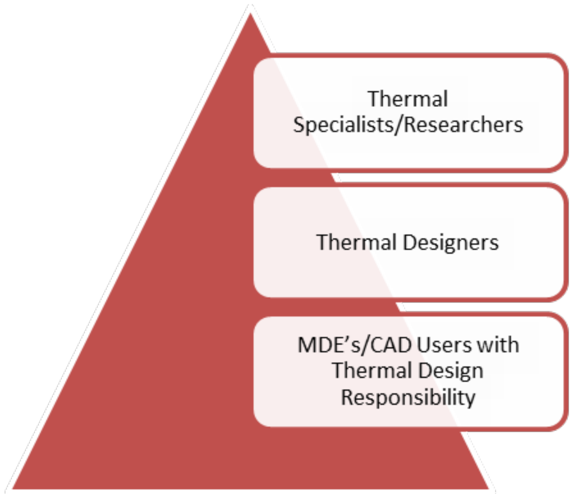

Figure 7: User Persona Pyramid for Electronics Thermal Design

ADVANCES IN COOLING ELECTRONICS WITH CFD

Analysing the differing user personas involved, it is also possible to break the list down further by focussing on their needs on a per industry basis and by region.

Traditionally, the sweet-spot for all CFD software solutions on the market has been the top two tiers of this pyramid. In recent years, however, responsibility for thermal analysis has increasingly been placed on the desk of the general Mechanical Design Engineer (MDE) in a wide range of electronics sectors. Most of these MDE's are thermally-aware, but do expect a software solution to be presented in a familiar form equivalent to today's CAD software in order for them to turn around analyses in a matter of days or weeks at most. The challenge, however, is crafting a solution that addresses these requirements without alienating the thermal designers or specialists who may not have the knowledge and experience to work well with a more CAD-like solution.

The answer is to provide a portfolio of solutions to address this problem and where possible a product UI which can be adjusted easily to suit the needs of anyone on the pyramid of user types. FloTHERM XT, a complementary product to FloTHERM in a portfolio of solutions, is the latest release from Mentor Graphics' Mechanical Analysis Division. It has been built with attention to the user persona requirements, utilising a best-in-class CAD UI toolkit and geometry engine, but with a simple option to either run the software in "full" or "reduced" mode. The latter is a control switch that obfuscates many of the less-common UI features and toolbars, thus reducing feature bloat and encouraging non-CAD trained usage of the software.

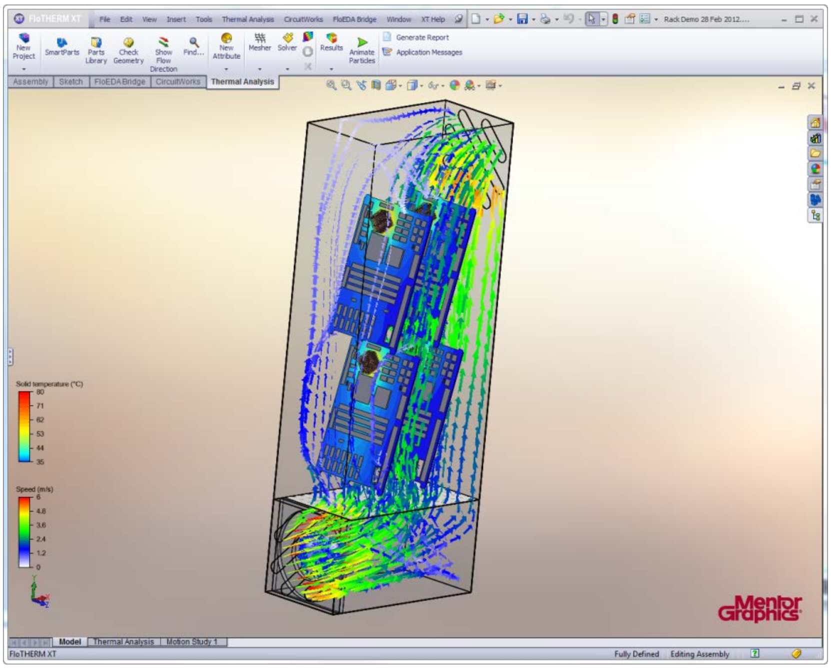

Figure 8: FloTHERM XT User Interface

ADVANCES IN COOLING ELECTRONICS WITH CFD

De-mystifying CFD can be another important element when considering broadening the appeal of CFD based analysis solutions to the non-expert. For CFD software companies, the most common technical support questions usually all involve the CFD aspects: meshing, solution convergence and control, turbulence modelling are the clearest examples of this. In order to address this, it is important to avoid CFD jargon wherever possible. For example, terms like “boundary conditions” are second nature to anyone working with CFD, but may not be so clear for non-specialists and can cause confusion.

As outlined earlier, meshing can take up a significant amount of time and energy in some general-purpose codes and can be a cause of frustration when it goes wrong. Most general mechanical engineers would like to simply have the software do the job for them wherever possible, but with the ability to switch to more manual definition should the need arise, and this has re-enforced the need for more sophisticated meshing strategies. For that reason, there is an increasing use of semi-automatic, object-based algorithms, but with options to adjust the mesh manually where necessary or to allow the freedom and control that is required by the more experienced, and CFD aware, thermal engineers.

Solution control on difficult cases can be a black art in some software codes with trial and error often used to achieve a solution. The more advanced codes, however, will utilise highly stable numerical schemes and solution controls will operate semi-automatically to control the convergence of the solution with only the minimum of user intervention ever being required.

For electronics cooling applications, issues relating to turbulence modelling are rarely, if ever, the largest source of error in the results. It is more likely to be uncertainties in power dissipation, materials, flow rates or interface resistances. That said, it can be a source of concern for some more specialised users. The key in software development for vertically-specialized CFD solutions is to provide the best possible model for the application area of interest and only provide alternatives if there is a clear reason to do so. The best electronics-centric CFD solutions today will provide options for laminar, transitional and turbulent flows, but limit the turbulence models that are available in order to avoid confusion. FloTHERM XT makes use of a general two-equation model combined with a proprietary immersed boundary treatment for near-wall effects that smoothly transitions between the different flow regimes, giving excellent benchmark results appropriate for electronics applications.

Design Collaboration

As described earlier, CFD analysis is needed throughout the design flow, from concept design to final product validation. In electronics, however, there has been an increasing recognition of the need to provide these analysis capabilities in a framework that provides the best possible collaboration between the main

ADVANCES IN COOLING ELECTRONICS WITH CFD

Mechanical and Electrical Design flows. We refer to this as MCAD/EDA Connectivity.

FloTHERM XT, designed with a CAD-centric UI and geometry engine, is designed to fit into the complex mechanical design environments and associated processes that exist in today's electronics companies. Imported CAD and internally-generated geometry function seamlessly together, thus enabling supply-chain integration and offering an opportunity for the product to fit anywhere in the mechanical design flow, from early conceptual design to final product verification.

This is further supported by a unique set of EDA interfacing capabilities, including update functionality, to easily remain concurrent with the latest EDA design changes.

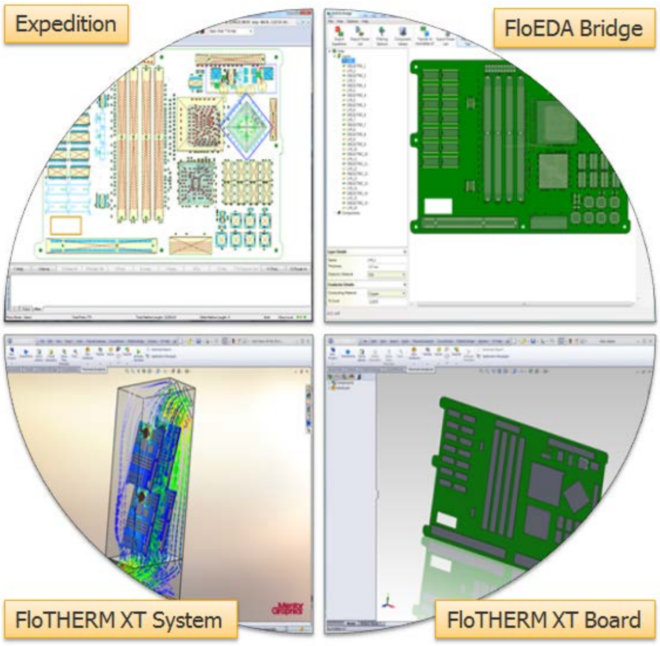

Figure 9: EDA Interfacing in FloTHERM XT

4: Case Study

The scenario for this case study in FloTHERM XT is the concept-to-prototype design of a new wall-mounted internet box, requiring a new arrangement of vents, a stylised shape as dictated by industrial marketing colleagues and a main circuit board dissipating 30% more power than earlier designs.

ADVANCES IN COOLING ELECTRONICS WITH CFD

The first stage concept design involves the creation of a box model with simple representations for the PCB and components, with an initial assessment made of overall cooling strategy and temperatures of the critical components.

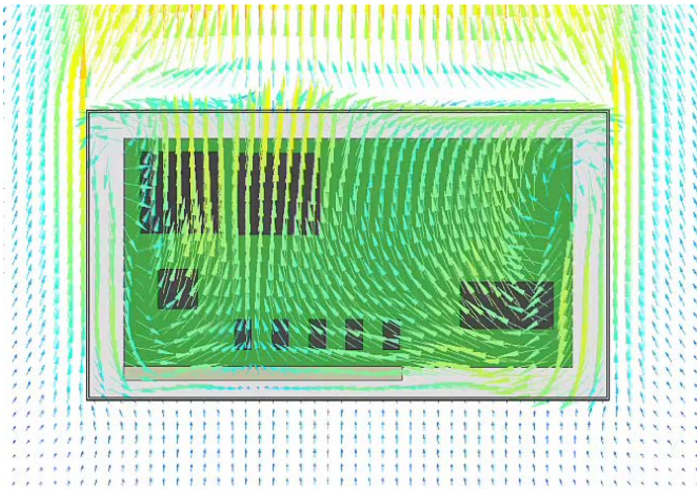

Figure 10: Wall Mounted Internet Box – Concept Design

The design then evolves in stages towards a model of the complete system with a MCAD-based mechanical design for the enclosure, any required heatsinks, interface resistances, and the full board layout imported from EDA via Expedition (Mentor Graphics, 2013).

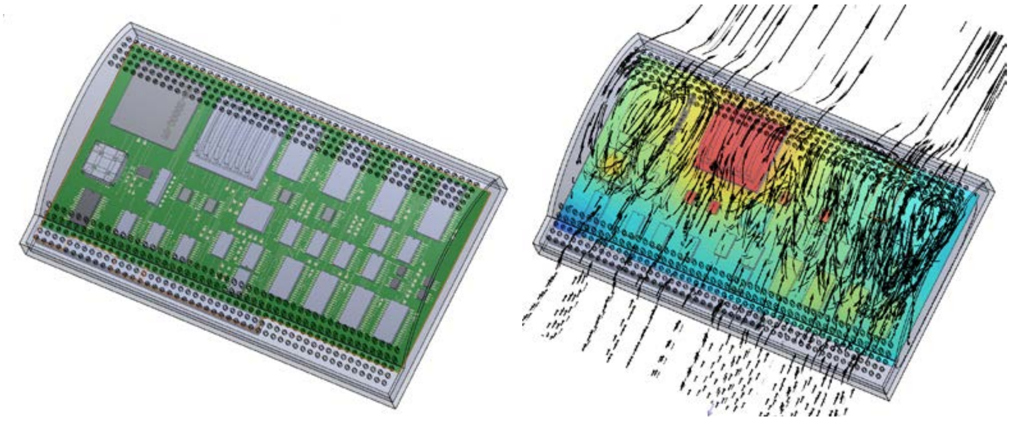

Figure 11: Wall Mounted Internet Box – Detailed Design

The component modelling also evolves as more information is defined. Thermal models may be changed from block-style, simple component models upwards in complexity to 2-Resistor (JEDEC, 2008a), DELPHI (JEDEC, 2008b), or fully detailed models, taking advantage of the advanced library swapping and filtering support within the software.

ADVANCES IN COOLING ELECTRONICS WITH CFD

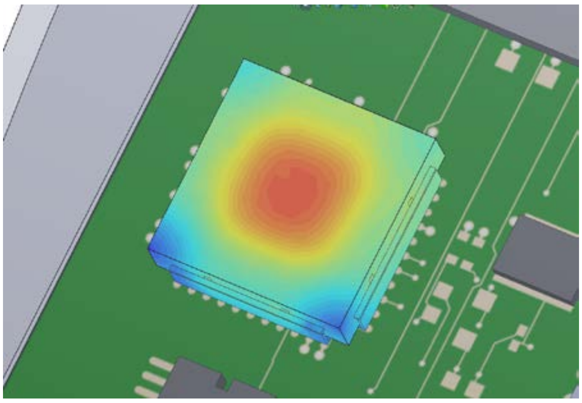

Figure 12: Wall Mounted Internet Box – Detailed Component Modelling

At all times, there is complete history and traceability of both the MCAD and EDA data as the overall system design proceeds towards prototyping and ultimately manufacture.

Conclusion

We have seen how thermal design has evolved since CFD was first used for this application in the late 1980s through to the present day, and the role CFD has played. Electronics cooling CFD has followed a different development trajectory to general-purpose CFD, and we have attributed this to a number of factors:

Electronics products pose a unique set of challenges for thermal simulation: scale disparity; uncertainty over data (component thermal data, power dissipation, material properties, layer thicknesses, interface resistances); and transitional flow regime.

The design times are very short, particularly in market sectors like consumer electronics the design taking place largely within an EDA framework. The pace of miniaturization at the silicon level has an impact at all packaging levels, forcing companies to continually evolve their design processes in order to stay competitive. This has a knock-on effect in the way CFD-based thermal design tools must integrate with and support the overall design flow. One very distinct difference we observe is the use of electronics cooling CFD (particularly Cartesian-based) in conceptual design, starting at the system-level, from an early stage (circa 1990), whilst general-purpose CFD has been slower at making the transition away from verification in late design where the cost and time savings associated with its use are far less.

Another major trend, linked to the rapidly-evolving product creation processes, is the change in user persona. The democratization of electronics cooling CFD away from thermal and simulation experts to the wider community of mechanical design engineers has exceeded that of general-purpose CFD in other industries.

This wider community of users has a broader range of responsibilities than just thermal, and does not place particular value on the CFD technology used to predict the air-side temperature rise, which is becoming a diminishing contribution to the overall junction temperature rise, limiting which remains the principal design criterion.

The aggregate effect of these differences will continue to make electronics cooling focused CFD the best-in-class tool for the majority of electronics cooling applications for the foreseeable future.

REFERENCES

Aberdeen Group (2007), *Electronics – Correct by Design*, Benchmark Report

Benington, Herbert D. (1983). "Production of Large Computer Programs". IEEE Annals of the History of Computing (IEEE Educational Activities Department) 5 (4): 350–361. doi:10.1109/MAHC.1983.10102.

Beck, Kent et al. (2001). "Manifesto for Agile Software Development". Agile Alliance.

Burdick, J. S. (1991), *Electronics Cooling at IBM Endicott*, Proc. of 1st FloTHERM Int. User Conf., Guildford UK, pp. 53-77.

Chipestimate.com (2011), P's Silent Presence in Automotive Market. Available: http://www.chipestimate.com/blogs/IPInsider/?p=92 [17 January 2013]

Gary Smith EDA (2010), PCB Market Trends 2010. Available: http://www.garysmitheda.com/

Hannemann, R. J. (2003), *Thermal Control of Electronics: Perspectives and Prospects*, Rohsenow Symposium on Future Trends in Heat Transfer, Warren M Rohsenhow Heat and Mass Trasfer Laboratory, Massachusetts Institute of Technology, Cambridge, MA. Available: http://web.mit.edu/hmtl/www/papers/HANNEMANN.pdf [17 January 2013]

Huang, J.C., et al. (2009), A Simple Model for Transient Thermal Behavior of Integrated Circuit Package Systems for Use in Determining a Thermal Design Power, Proc. ASME HTSC, San Francisco, USA, 19-13 July 2009)

ADVANCES IN COOLING ELECTRONICS WITH CFD

Intel (2011), Measuring Processor Power: TDP vs. ACP, Whitepaper on Intel Xeon Processor Architecture Analysis, Intel, April 2011. Available: http://www.intel.co.uk/content/dam/doc/white-paper/resources-xeon-measuring-processor-power-paper.pdf [17 January 2013]

JEDEC (2008a), Two-Resistor Compact Thermal Model Guideline, JESD15-3, July 2008. Available: http://www.jedec.org/standards-documents/docs/jesd-15-3 [17 January 2013]

JEDEC (2008b), DELPHI Compact Thermal Model Guideline, JESD15-4, October 2008. Available: http://www.jedec.org/standards-documents/docs/jesd-15-4 [17 January 2013]

Mentor Graphics (2012), *Mentor Graphics Launches New T3Ster DynTIM Tester to Deliver Unrivalled Thermal Interface Material (TIM) Measurement Methodology*. Press release, December 13, 2012. Available: http://www.mentor.com/company/news/mentor-launches-t3ster-dyntim-tester [17 January 2013]

Mentor Graphics (2013), Expedition Enterprise PCB Design Flow. Available: http://www.mentor.com/products/pcb-system-design/design-flows/expedition-enterprise/ [17 January 2013]

Mountain Goat Software (2005), Available: http://www.mountaingoatsoftware.com/system/asset/file/17/ScrumLargeLabelled.png [January 2014]

Pop, Erik (2010), Energy Dissipation and Transport in Nanoscale Devices, Nano Res (3) pp. 147–169. Available: http://www.nanotechnologies.qc.ca/wp-content/uploads/2010/04/fulltext.pdf [17 January 2013]

Sutter, Herb (2005) *The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software*, Dr Dobbs Journal, Vol. 30, No. 3. Available: http://www.gotw.ca/publications/concurrency-ddj.htm [17 January 2013]

The Economist (2003), *Feeling the Heat*, Technology Quarterly, Q1 2003. Available: http://www.economist.com/node/1620884/print [17 January 2013]