20/20 HINDSIGHT TO THERMAL DESIGN 2020

ABSTRACT

Computational Fluid Dynamics (CFD) for Electronics Cooling (EC) has evolved differently from general-purpose (GP) CFD, due to the nature of the market it serves and user profile. The benefits are clear – the use of EC CFD in product design has had a profound impact on both time-to-market and cost. Today the EC CFD market is dominated by tools that span the packaging levels, typically from package to system, and represents a consolidation of suites of more package-level specific codes that emerged in the early 2000s. Interfacing with Electronic Design Automation (EDA) and Mechanical Computer Aided Design (MCAD) software has helped their incorporation into existing in-house design practices. MCAD-embedded CFD software has gained in popularity and is being sold effectively by several MCAD vendors for GP applications, and EC CFD. This invited paper considers the evolution of EC CFD in the context of the unique requirements of electronics cooling applications and its future development, with some thoughts on its future.

KEYWORDS

Computational Fluid Dynamics, CFD, thermal design, electronics cooling, IC package PCB, heat sink, reliability.

1. INTRODUCTION

It's worth reflecting on the business impact EC CFD has had on the thermal design of electronic products. Independent research almost 10 years ago found that best-in-class companies that used EC CFD completed thermal design verification more than twice as fast, with the number of board re-spins reduced from typically 2-3 down to 1 or zero [1].

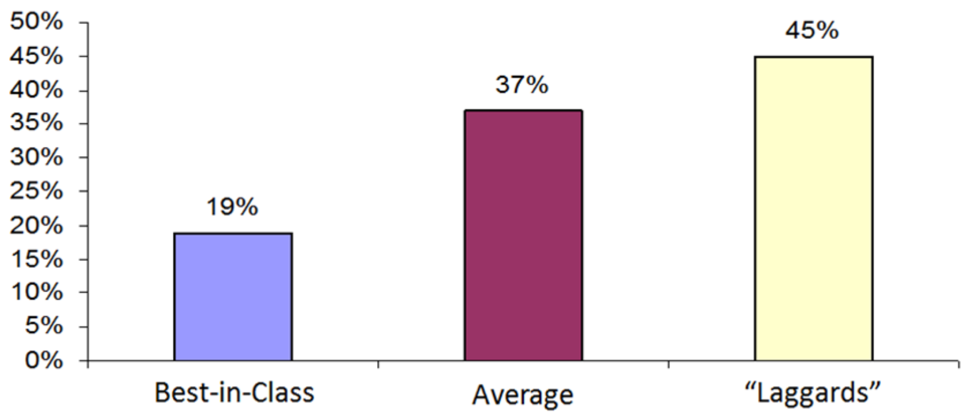

Figure 1: Proportion of Total Design Time to Verify Thermal Design.

2. OBSERVATIONS FROM HISTORY

Prior to the introduction of CFD, mechanical engineers often relied on the use of thermal resistance metrics such as junction-to-ambient ($R\theta_{JA}$) and junction-to-case ($R\theta_{JC}$) in hand calculations to estimate component temperature rises. $R\theta_{JA}$ includes a considerable contribution from the test board and test environment. $R\theta_{JA}$ was never intended to be used as design data, and is not well-suited for that purpose, which is now widely recognized [2].

Early use of EC CFD was focused on system-level design verification when physical prototypes became available. Problems, when found, led to costly late-cycle redesign as depicted in Fig. 2. Early EDA-integrated thermal tools that used correlations for the air-side heat transfer were the main competition for CFD during the early 1990s. Lack of thermal data and an inability to accurately represent the airside heat transfer limited their suitability for system-level analysis, but they still find use in board design [3].

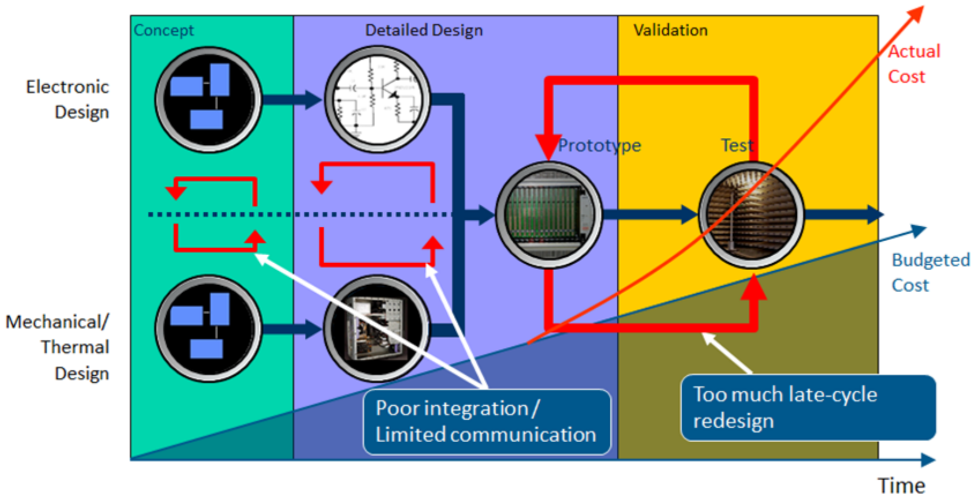

Figure 2: Typical Thermal Design Process (circa 1990).

Major GP CFD codes like PHOENICS, Fluent, and Star-CD had been around for some time when the use of EC CFD started to gather pace. They found relatively little use within this market sector however, held back by the lack of a conjugate[1] heat transfer capability (being the ability to handle heat transfer within solids and the fluid at the same time), time-consuming mesh generation, and limited industry knowledge on the part of the vendors. Burdick [4] commented that "A small number of engineers attempted to use commercially available GP finite-difference CFD programs at this time but the result of several months of activity was usually fruitless".

At introduction, trust in CFD for thermal design was relatively low, perhaps because those evaluating it already had experience with the GP CFD available at the time. As a result, a lot of effort went into evaluating software during a trial period, often lasting more than a month, to confirm that the software could handle the prospect's application. Due to the limited computing power available, simplification was needed to arrive at a tractable model, requiring modeling judgment based on domain expertise and experience, and hand-holding from vendors.

The early adopters were generally experienced thermal engineers that had worked through the bipolar age with a strong background in measurement techniques, gained through build-and-test prototyping. Mainly Mechanical Engineers (MEs), plus the odd physicist, they were converted to using CFD by the insights it provided into system-level air flow to understand the behavior of physical prototypes. This quickly gave way to the desire to predict solid temperatures, as these could easily be measured with thermocouples, the task being to capture as much detail in the model as could be simulated (e.g. Fig 3).

[1] *A term almost unknown outside the general-purpose CFD community*

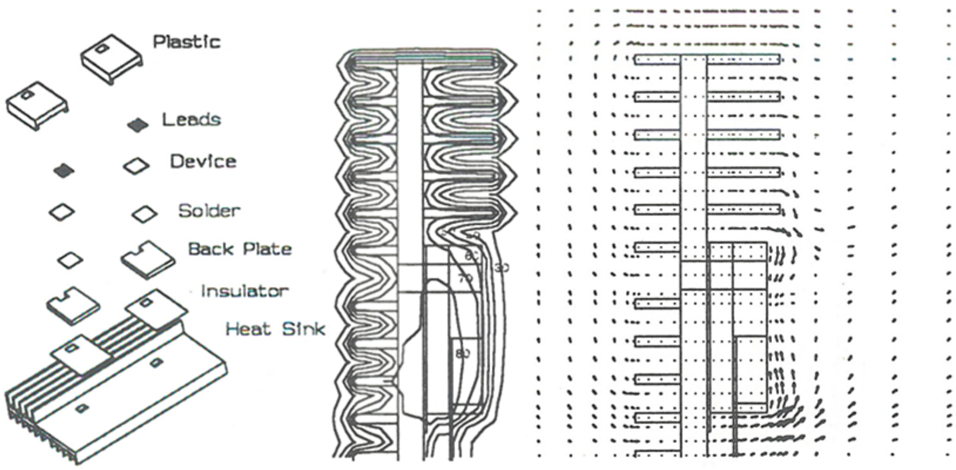

Figure 3: Power Transistor Heatsink Validation Study (IBM Endicott, circa 1990) [4].

To understand how the market evolved beyond this, it's helpful to discuss the challenges presented by electronics cooling applications.

3. UNIQUE FEATURES OF THE ELECTRONICS COOLING MARKET

The first industries to embrace CFD were aerospace and nuclear. Early Eulerian solvers used in aerospace pioneered the use of structured body-fitted meshes for transonic flows led the later finite-volume Navier-Stokes solvers down a body-fitted development trajectory [5]. The need to handle more complex, sometimes moving geometries led to the development of fully unstructured body-fitted codes and mesh generators. Support for user-created coding allowed research scientists to apply CFD to a variety of new industrial problems and academic developments have fed into the GP CFD software of today. Typically used by professional analysts with higher degrees in fluid dynamics and/or numerical methods, these tools are now capable of handling free-surface flows, moving geometry, combustion and chemical reaction, multi-phase phenomena and much more.

While air-cooled electronics requires only basic CFD to conserve mass, momentum and heat in single-phase low-speed flows, the application is not without its challenges. The number of discrete objects, and interfaces between objects, that make up a typical problem is far higher than that encountered in other fields (e.g. thousands, vs. a single aerofoil) with almost every object participating in heat transfer, thus requiring appropriate material and surface properties. The difference in length scales from chip to system/room of 3 to 4 orders of magnitude presents a significant challenge, not least because of the need to capture these within the same computational mesh. The parts that make up an electronics system: chip packages, fans, heat pipes, thermoelectric coolers (TECs), etc. have complex internal structures and/or thermofluidic characteristics.

Flows in electronics systems are often transitional turbulent. Turbulence is created by flow passing over the many small components present in the system, which act as turbulators, and is not self-sustaining, dissipating downstream, between heatsink fins, etc. This class of flow falls well outside the range for which conventional turbulence models were first developed. Another complication is that the thicknesses of the aerodynamic and thermal boundary layers can be quite different due to the localized nature of the heat sources. The heat transfer and fluid flow are strongly coupled in many systems due to buoyancy effects, and the temperature dependence of material properties can be important in silicon, some ceramics, in liquids, and even gas, for example in and around a projector bulb.

Yet a usefully-accurate solution has to be obtained in hours[2] on available, usually desktop, computer hardware. Finally, the target user group demands robustness in terms of problem setup and solution convergence, without having the requisite knowledge of the numerical aspects of CFD that would otherwise be needed to take effective remedial action should the solution prove to be unstable. As a final point, the development time associated with electronics products is much less than that found in the industries that first adopted CFD (aerospace and nuclear), increasing the pressure for fast simulation turnaround to support design decisions.

4. CFD TECHNOLOGIES FOR ELECTRONICS COOLING

Some unique CFD technologies have been developed to meet these needs. Early electronics applications were highly Cartesian in nature, so a Cartesian mesh was the obvious choice. Staggered Cartesian meshes have the additional benefit of being able to tolerate high cell aspect ratios (i.e. above 100:1) without impacting result quality or convergence, making them efficient at handling thin and layered structures like heat sink fins, printed circuit boards (PCBs), etc. They can also be created instantaneously and with 100% reliability.

Over time, originally structured Cartesian meshing approaches have been adapted to become unstructured, allowing everything to be solved as a single matrix. Abutting, overlapping and nested locally-fine regions further help to address the disparity in length scales, and to allow mesh to be associated with objects, including expansion into the fluid region surrounding them. More recently octree-based meshing has found favor.

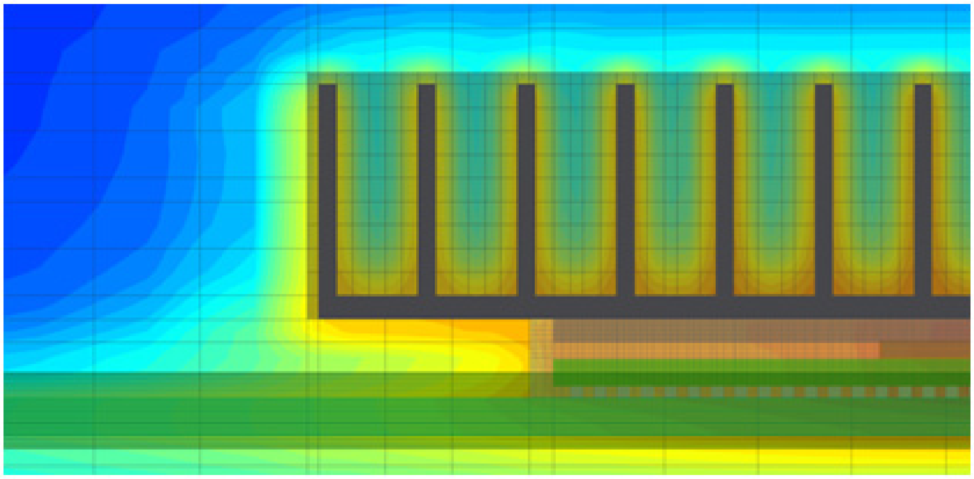

Figure 4: Local Fine Unstructured Cartesian CFD Meshes.

On coarse grids zero-equation turbulence models give better results than higher-order one- or two-equation models that require a fine mesh to adequately calculate the turbulent kinetic energy (k), as its source term depends on the square of the local velocity gradient in each direction. This is particularly true for channel-type flows containing many obstructions that typified early electronics applications, explaining their use from the outset.

The importance of surface-to-surface radiation was initially overlooked due to the early focus on predicting air temperatures. This was quickly addressed with automated calculation of single view factors, typically between user-selected surface pairs. Many electronics systems contain multiple separated fluids, leading tools to support multiple fluids within subdomains separated by solids.

[2] *Delivering far more business value than a better solution obtained in days.*

5. ADDITIONAL NON-CFD TECHNOLOGIES

Lack of knowledge about how to represent the complex flow and thermal characteristics of many of the component parts of an electronics system quickly became a brake on sales, with pressure from industrial customers to improve the situation. The build-up of PCBs and the internal construction of chip packages, plus data on their materials of construction were not available. Often even a single $R\theta_{JC}$ value was not provided, with $R\theta_{JA}$ being more commonly available.

A key benefit of focusing on a specific application area is that it is possible to develop bespoke solutions. Flomerics' in-house Package Level Thermal Initiative (PLTI) in the early 1990s preceded two successful European-funded projects that Flomerics coordinated: DELPHI [6] and SEED [7] with a subsequent project PROFIT, being coordinated by Philips Research [8]. These projects led to the concept of behavioral models for generic parts such as axial fans, chip packages families, and heatsinks. The principle of Boundary Condition Independence (BCI) was established, meaning that the model for the part (to be provided by the part vendor) should contain only information about the part itself, and be free of any information about its environment. The model should predict the thermal performance of the part to an acceptable level of accuracy when an arbitrary thermal environment is specified (by the end user). The Compact Thermal Models (CTMs) of chip packages that resulted from this work has since given rise to substantial additional research that continues to this day [9, 10], and standardization efforts in this area have also borne fruit [11, 12].

At the other end of the accuracy scale, the 'Kordyban model' (an isotropic thermal conductivity of 10Wm-1K-1 applied to a cuboid block to represent a component and/or board), used in the absence of other data, is testament to the tenacity of thermal designers in recognizing that the relative impact of design changes on cooling performance can be judged qualitatively [13].

PCBs have benefitted from considerable study into their thermal representation, covering single isotropic and orthotropic objects to optimize the values to best capture heat spreading in the board. These approaches are still used today, particularly in early design before the board is routed. Later on it became tractable to handle the individual layers as discrete homogeneous objects, and more recently attention has been focused on capturing the local copper content, particularly close to components, where a more detailed representation has the greatest effect on predictive accuracy. Recent developments have also focused on capturing Joule heating effects, either importing these as a matrix of heat sources from an EDA tool, or by representing individual nets, power and ground planes in detail, and performing a DC electrical solution within the CFD tool – a fully-coupled multi-disciplinary simulation.

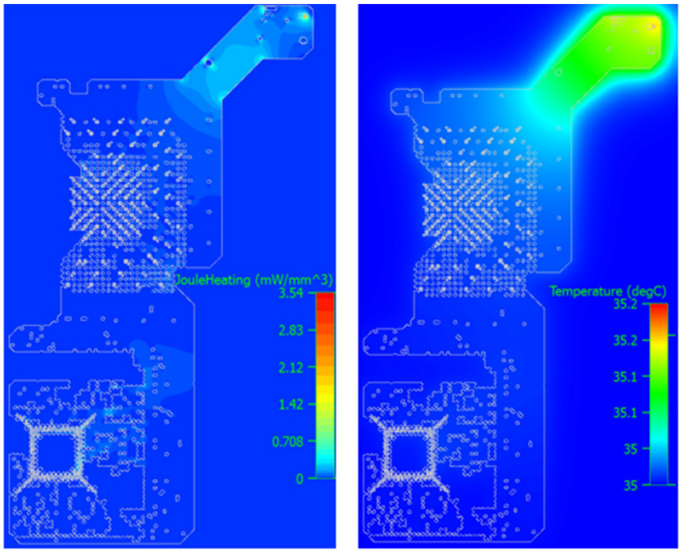

Figure 5: Joule Heating and Temperature in a Section of Split Power Plane.

Today, heat pipes are a common feature in many products, with liquid cooling being employed in various applications. These and other cooling technologies such as synthetic jets [14] and piezoelectric fans [15] continue to underscore the ongoing need for validated behavioral models of the commercial off the shelf (COTS) components that form part of the overall cooling solution.

6. EVOLUTION OF THERMAL DESIGN

Early efforts during physical prototyping typically focused on simple mechanical improvements such fan and vent positioning to improve flow distribution, usually after the electronics design had closed. As space constraints and power densities increased, heat sink optimization during late design became important to minimize weight, system pressure drop and wake effects. The focus on telecoms, computing and later networking continued from the outset until 2000, when the 'dot com' bubble burst and design work in these industries all but stopped as surplus equipment remained unsold in warehouses. Sectors such as defense, aerospace and automotive came to the fore, placing increased emphasis on links to mechanical CAD systems and support for arbitrary geometry.

Early system-level EC CFD software was soon complemented by product offerings at first board- and then package-level, providing suites of software that share models and map onto much of the electronics design flow. FloTHERM PACK [16] is an exception to this general trend, with its early appearance resulting from research started in the DELPHI project to predict junction temperatures accurately in the application environment (Fig. 6).

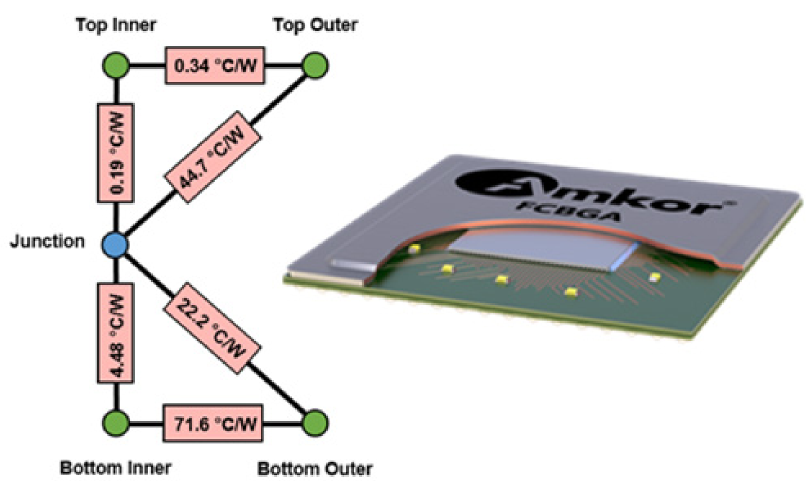

Figure 6: DELPHI network of a 45mm FCLBGA Package (inset), courtesy of Amkor.

The drivers for EC CFD have also shifted, from being performance-related in the computing, telecoms and networking sectors, where interest in predicting junction temperature is due to its impact on switching speed, to being more reliability-related in others like military & aerospace, automotive and industrial electronics. As the complexity of electronic and mechanical products has increased, design processes have become increasingly dependent on MCAD and EDA toolsets for design data and product lifecycle management (PLM). In EC applications, CFD sits at the interface between these main design flows and needs information from both worlds.

By the end of the design process, part details and design powers, PCB layout, details of the board structure etc. are all available within the various EDA toolsets. However, due to the historic nature of electronics design being largely 2D, necessary mechanical information about the board assembly was often lacking, as the board design may carry only component footprint information and a reference designator for the component.

The geometric detail of many other aspects of the product will exist within the MCAD system, and standard part libraries, tied to the PLM system, ensure every nut, screw, washer etc. is accounted for in the bill of materials. But not everything. Geometry of COTS components such as fans, heatsinks, etc. may be absent. Neither system contains sufficient information about the thermal properties of materials, nor do they contain behavioral models of parts needed for the analysis, such as resistance networks for chip packages, fan curves, heat pipe characteristics, TEC performance data, etc. or power information related to the product's operation.

To address these challenges, sophisticated interfaces to both MCAD and EDA systems have been developed. Board designs can be imported in a variety of formats from placement through to final routed designs, with filtering to remove thermally-irrelevant detail. Placement changes made within the EC CFD tool can be back-annotated to the EDA flow, where a board designer can decide whether to accept or reject the change. Imported MCAD geometry could be added to or replaced, for example the CAD geometry for an axial fan, should this exist, could be replaced with the equivalent behavioral model from a library within the EC CFD tool.

Increasing use at board-level using detailed package models to accurately predict junction temperatures in late design fuelled concerns over predictive accuracy, and motivated numerous investigations into the performance of Reynolds-Averaged Navier-Stokes (RANS) turbulence closure models for this class of flow [17, 18]. To better account for the influence of hotspots, surfaces selected to participate in thermal radiation were multiply subdivided, and solar models added, allowing the simulation of outdoor enclosures.

In early design, libraries of behavioral models of common parts, such as fans, heat sinks, chip packages, etc. have an important role to play in efficiently creating a thermal model of an electronics enclosure. Early recognition of this influenced the way EC CFD tools evolved. An early characteristic of EC CFD is an object-centric approach to data storage, with the geometry of a part stored together with its materials of construction; surface information such as roughness, emissivity, and color; design power; and mesh information, allowing the part to be saved and reused in other designs. The support for assemblies and the development of drag-and-drop library functionality, now commonplace within the tools themselves, facilitates model reuse, and sharing across an organization and into the supply chain. This has become as vital to design flow integration as EDA and MCAD interfacing.

7. CURRENT STATUS OF ELECTRONICS COOLING CFD

The unique CFD architecture of EC-dedicated tools has continued with the evolution of Cartesian-based meshes, still generally preferred for their speed of creation and robustness of solution, to hybrid octree-polyhedral meshes, with octree in the fluid region, and arbitrary polyhedral control volumes at solid-solid and solid-fluid boundaries that are not meshed, but constructed directly from the intersection of the CAD geometry with the octree mesh. These SmartCells™ break the normal 1:1 correspondence between mesh cells and solver control volumes typically found in GP CFD.

Octree meshes offer excellent scope for rapidly changing mesh cell size with nested 8:1 volume reduction, allowing mesh to be concentrated only where it is needed to capture angled and curved geometries, with automated refinement to resolve small features, narrow channels etc. and solution-adaptive meshing can be easily applied to resolve gradients in the flow and within solid structures during solution.

Limitations in computing power and short design times still preclude the use of most advanced transient CFD techniques such as Large Eddy Simulation (LES) for electronics product design. Despite some valid accuracy concerns, the use of zero-equation RANS models with first-order differencing schemes on Cartesian-based meshes remains popular [19], alleviated by greater uncertainties relating to the many non-CFD aspects of the model (geometry, layer thicknesses, interfacial resistances, powers, etc.). There is increasing use of 2-equation low Reynolds Number RANS models and second-order differencing schemes, as part of a growing focus on simulation accuracy, as design margins shrink, reliant on increases in model fidelity.

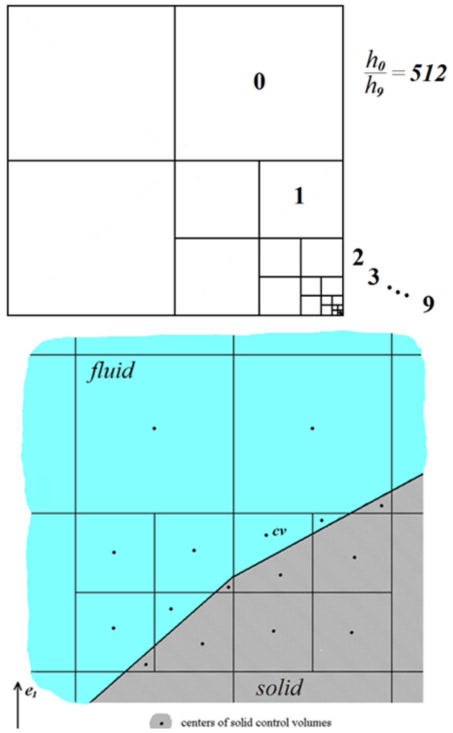

Figure 7: Octree Mesh showing Arbitrary Polyhedral SmartCells™ at Solid-Fluid Interface.

Mesh distortion in body-fitted meshes reduces accuracy and worsens convergence [20]. Achieving high quality body-fitted meshes in electronics applications presents a particular problem due to the complexity and 'clutterdness' of the geometry, which is getting worse over time. This can limit use to late stage design and in some cases may require the model to be partitioned and simulated in parts with assumptions about the boundary conditions. Zoom-in solutions, where a part of a model is analyzed in more detail (e.g. a card slot), by taking the boundary conditions at its periphery from a larger model (e.g. a rack) were automated. They did not gain traction, the lack of feedback between the modeling levels being a concern, e.g. the flow through the card is fixed, and not dependent on the pressure drop, when in practice other card slots provide parallel flow paths.

Tools may now incorporate a solid modeler, e.g. based on ACIS or Parasolid, allowing native parametric CAD geometry to be imported as a complex assembly, complete with mates and the feature history of how it was constructed. Irrelevant features can then be temporarily suppressed to simplify the part for analysis to optimize the trade-off between simulation speed and accuracy, and others changed deliberately to improve the thermal design. The modified CAD part can be exported in its native format for re-import back to the MCAD system. If there is no requirement to change the CAD geometry importing a Standard Tessellation Language (STL) file of a single part might be meet the requirements of the tools targeted workflow, for example in late design.

Other approaches to modelling thermal and solar radiation have been adopted to increase accuracy, first with Discrete Transfer Radiation Model (DTRM), and more recently Discrete Ordinates (DO) that allows for absorption in semi-transparent media, and Monte-Carlo models that can include spectral dependency and optical effects, for example for use cameras, projectors and in automotive headlights, where the use of LEDs is bringing new challenges such as demisting in cold humid conditions.

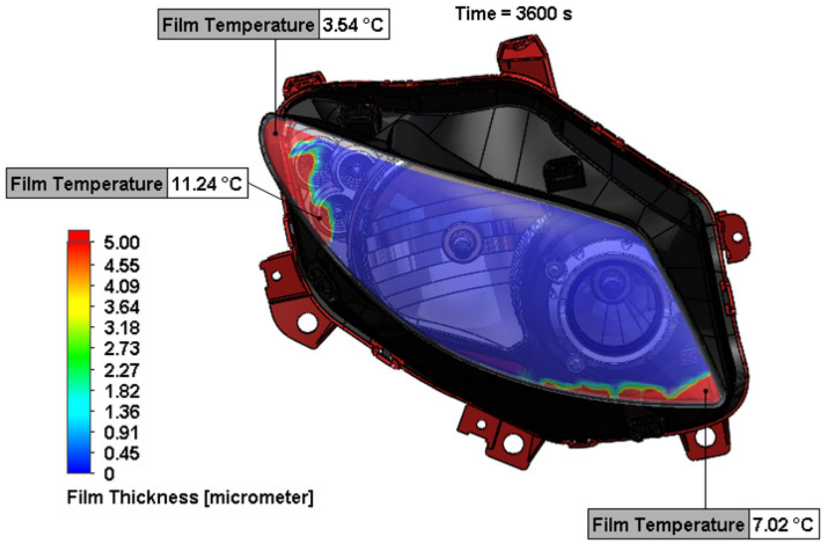

Figure 8: Investigating Water Film Evaporation in Automotive Headlight Thermal Design.

One rather interesting trend is resurgence in thermal testing. Unlike the early days when EC CFD was used to supplement final physical prototyping, thermal testing today is focused on providing input data to underpin virtual prototyping efforts using EC CFD. From an end user perspective this includes confirming and supplementing data in vendor datasheets, measurement of Cauer RC-ladder models, measuring thermal interface material conductivity etc. to achieve the most accurate deterministic simulation model.

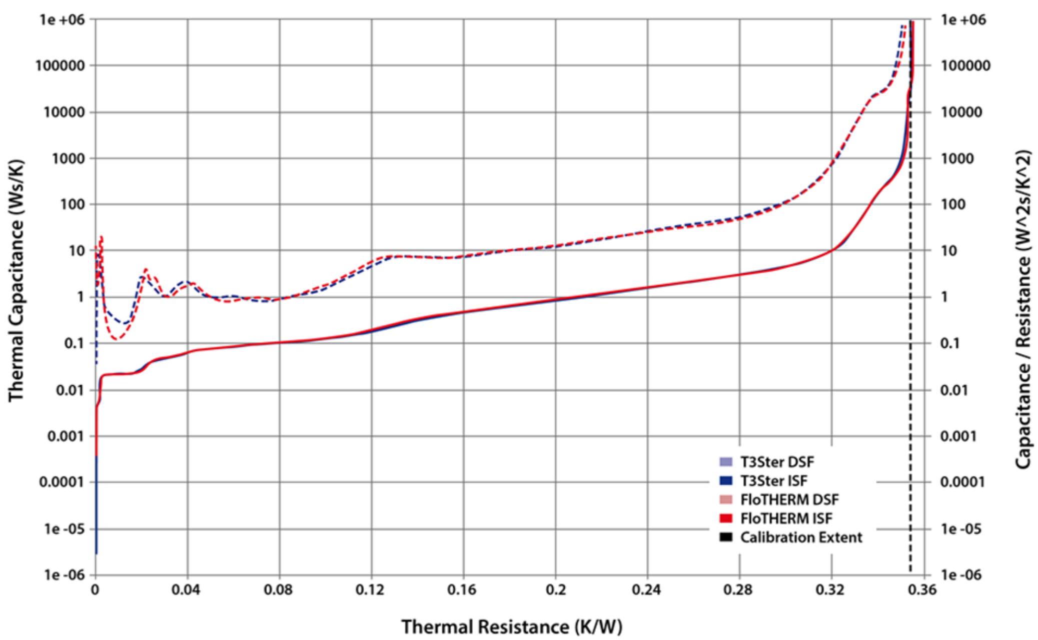

Figure 9: Final (Calibrated) Detailed Model Structure Functions vs. Physical Test.

Vendors can support this effort by calibrating detailed package thermal models they provide to their customers (Fig. 9), or create compact models based on calibrated detailed models.

Cooling adds cost, weight and volume to electronic products without increasing functional performance. The desire to minimize cooling costs against a background of increasing thermal density has led to an emphasis on design optimization throughout the design process. The Cartesian nature of the geometry, use of Cartesian-based meshes, and robust solution techniques has supported fully-automated exploration of the design space. The addition, movement and removal of objects, coupled with space-filling Design-of-Experiment (DoE) approaches with object collision detection and response surface based optimization techniques makes it possible to optimize almost all aspects of the thermal design, including component placement, PCB spacing, heat sink design, etc. from early in the design process where there is the greatest opportunity for design changes and hence cost savings. Such a focus on "front-loading" CFD in the design process is apparent across many industries.

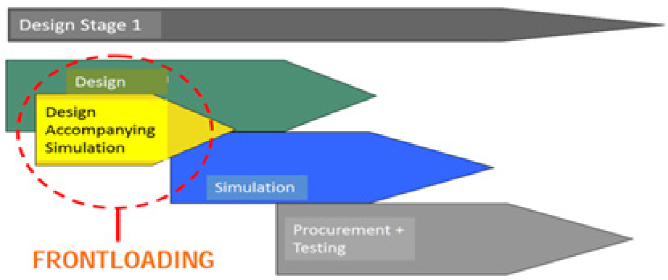

Figure 10: Final (Calibrated) Detailed Model Structure Functions vs. Physical Test.

While some progress has been made in getting suppliers to provide flow/thermal models of the parts they sell, with a few exceptions, openly-available vendor-supplied thermal models has not become commonplace, with models being provided on a request basis. In part, this may be due to a lack of open standards for data exchange between tools necessitating model provision in proprietary vendor formats. An open file interchange format is currently being worked on for resistor-capacitor based CTMs within the JEDEC JC15 committee, with a round robin between CFD vendors underway to harden the file format ahead of it being published.

The increasing need for interference checks that arise with miniaturization [22] is forcing MCAD and EDA vendors more closely together and to communicate changes between their two worlds [23]. Commercially, the distinction between these two worlds is also blurring as EDA companies move into the mechanical space, and mechanical CAD and CAE companies move more into electronics.

Today, stand-alone EC CFD solutions continue to dominate the market. MCAD-embedded CFD exists within at least five MCAD suites and is being used for EC applications, with MCAD-embedded CFD one of the fastest growing niches in CFD in recent years [24]. What is noticeably absent from today's market is an EDA-embedded CFD-based thermal design solution.

Simulation time for EC CFD remains high, as model complexity has grown with increasing computing speeds, the advent of 64 bit operating systems, and lower cost RAM. The limiting factor remains, as it always has, what can be achieved in an overnight run, allowing the thermal impact of design changes to be assessed on a daily basis. What has changed is the level of detail that can be included. Early simulations were limited to ~10,000 cells, whereas today, depending on the technology used, simulations involving 50M-200M cells are quite possible using desktop computing, growing as processing power and memory size broadly track Moore's Law, and through software developments to achieve good performance scaling on multicore machines with shared memory architecture. Moore's law has been accompanied by an increase in power density at all packaging levels, driving the need for more thermal simulation and greater detail at each process node. As models have become more complex, the challenge of interpreting results has increased, requiring both higher fidelity post-processing to produce near photo-realistic results, to help convince management and other stakeholders the results are correct, and so accept recommendations for design improvements.

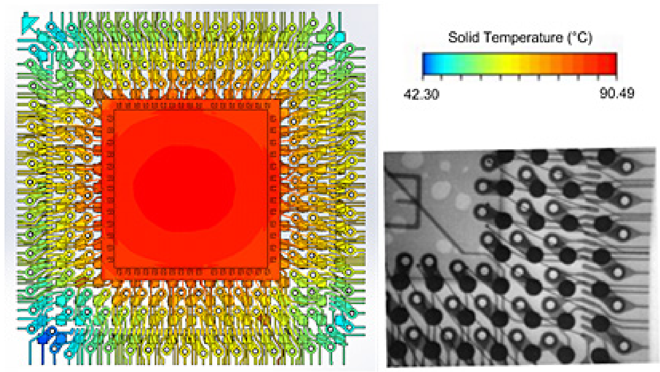

Figure 11: Temperature in FpBGA 208 (inset: X-ray); Courtesy of Thales Corporate Engineering.

As design margins tighten, the focus has moved away from designing for a single, maximum design power, through a time-averaged "thermal design power" to simulating combinations of use cases, with different power profiles, particularly in automotive, and in mobile applications where the focus is on active thermal management with the bottleneck being the perceived outer case temperature. Junction temperature remains the main metric for thermal performance in most applications, despite reliability being largely dependent on the change in temperature, e.g. of the solder joints in the system, and adding heatsinks to reduce temperature can increase solder joint fatigue, as can the use of an inappropriate underfill.

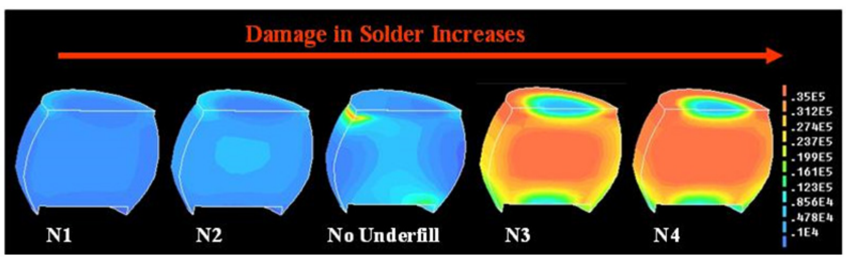

Figure 12: Effect of underfill on creep energy (damage) in critical PBGA solder joint in Eurofighter Avionics system; courtesy of University of Greenwich.

In power electronics, developments in wide band-gap IC technologies such as silicon carbide and gallium nitride have expanded the possibilities for low loss power conversion, bringing electronics into new areas, such as powertrain in ground transportation, albeit strongly encouraged by legislation to cut CO_2 and NO_X emissions, and in more electric aircraft. Accurate prediction of junction temperature excursion during an automotive drive cycle using a validated thermal model of the power module, combined with component reliability measurements from active power cycling is making field lifetime prediction possible [25].

8. BEYOND TOOLSETS

So far the tools themselves have been discussed, largely from a functional standpoint. Today CFD technology is well-trusted, and the available tools quite fully-featured. Attention is now turned to other aspects that are important to the successful commercial deployment of EC CFD from an end user perspective.

Electronic products are often evolutionary in their design, with many parts reused between products, so the advice to "start simple, get that right, and move on from there" often can't be applied in practice, with complex thermal models being created from the outset. This has led to an increasing focus on non-functional aspects related to User Experience (UX) [26]. To support error-free model building, multiple views of the model data, summary information, etc. are provided to assist checking. Largely this focus comes from the user community, following the trend seen in mechanical CAD where very large user communities bring their own challenges for vendors, requiring systems that support a meritocracy by allowing ideas to be discussed and voted on [27].

The need to continuously update and re-run models to check the impact of design changes results in a lot of repetitive work that distracts thermal designers from higher value work on the next generation of products. This has led to a focus on workflow automation. Today, models can be built, solved and the results post-processed without opening the tool, for example, driven by an Excel spreadsheet using VBA scripting.

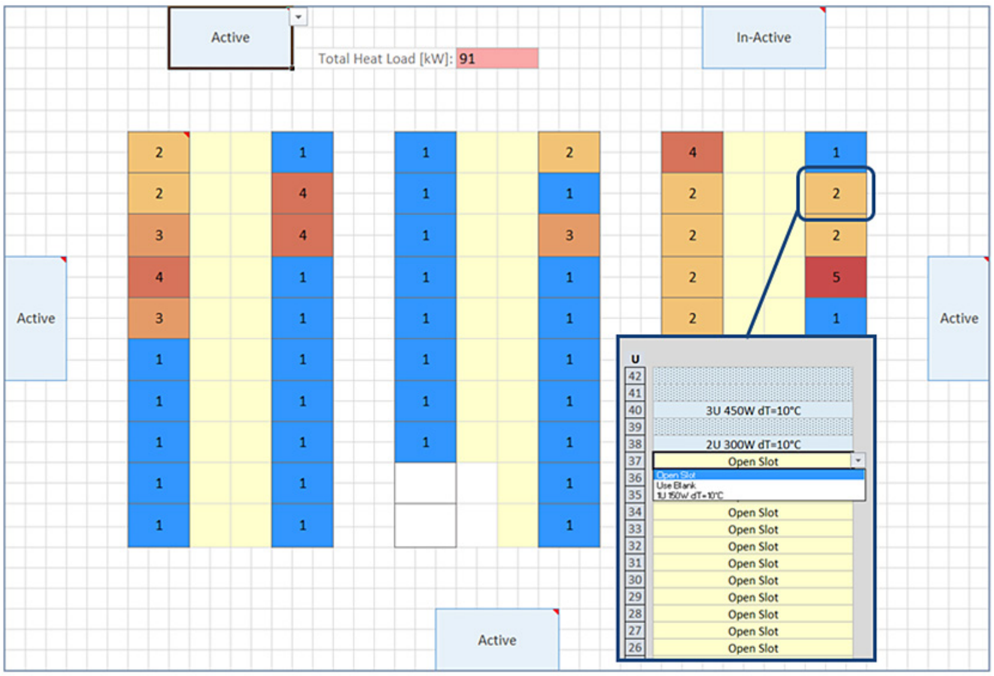

Figure 13: Data Center Model in Excel with Management of Rack Assets Spreadsheet (inset).

Examples include building detailed models of a particular chip package family, and performing wind tunnel and heat spreading 'sizing' calculations in early design previously based on correlations. EC CFD software can also be driven from within other toolsets as diverse as IC design [28] and Data Center Infrastructure Management [29].

It is also appropriate to look beyond the users themselves. Engineering managers are concerned with the ensuring the design process, including simulation, proceeds smoothly. Automation allows an organization to embed best practice into the workflow by capturing the best in-house knowledge and experience as part of the process itself, eliminating dependency on the knowledge and experience of individual engineers, thereby reducing result variability. Actions taken within the software can be recorded, allowing exactly the same actions to be repeated, for example on import of an updated piece of MCAD geometry, all but eliminating manual mistakes from repetitive tasks.

In the wider context of closing the design flow, integration of thermal models into existing circuit simulators used by EEs at chip- and board-level as thermal resistor-capacitor networks is a current trend, enabling reliable electro-thermal simulation. Integration with downstream tools, principally FEA tools for thermo-mechanical stress simulation has become more common.

So what's next?

9. FUTURE DEVELOPMENTS

Over the last 25 years EC CFD codes have followed a development trajectory quite different to that of GP CFD, driven by the needs of application area and a different target user profile. In EC CFD there is not the same driver to use distributed HPC resources to run massive models, as solution accuracy is limited by model data quality to a greater extent than fidelity of the numerics. Cloud computing offers substantial scope for design space exploration by running a large number of design variations concurrently on cloud-provisioned servers. At the other end of the computing scale, tablet-based apps are set to support the 'appification' of CAE tools [30], perhaps making CAE accessible to the large majority of engineers who are not yet users, but will need both customers and vendors to embrace a different business model that respectively has less certainty regarding end user spend.

OpenFOAM [31], a free open source CFD software package has impacted applications requiring large HPC installations for high fidelity CFD, like external aerodynamics in Formula 1, but has made little impact in electronics cooling applications. Vendors of Lattice-Boltzmann method codes have also shown some interest in EC applications, but as yet the technology appears to be less well-suited to heat transfer applications than external aerodynamics.

A single common design environment spanning MCAD and EDA seems impractical, but 3D design in EDA will facilitate bidirectional notification of relevant design changes and data exchange between EDA and MCAD systems. As the EDA and MCAD worlds converge, the potential for EC CFD tools to integrate more tightly with these design environments increases.

As electronics products and the cooling technologies they employ continue to miniaturize, new challenges will continue to appear: micro-channel cooling pushes the limits of applicability of the Navier-Stokes equations requiring as a minimum a slip condition at wall boundaries, and MEMS devices require a multi-domain design approach.

In general, a multi-discipline approach requires 'white box' models, where the geometry is represented in sufficient detail to support each type of analysis, e.g. electrical (finite difference), thermal (CFD, finite volume), thermo-mechanical stress (finite element), electro-magnetic (finite element, transmission line method). This precludes the use of 'black box' behavioral models that can greatly speed the design activity, and so adds both time and complexity to the design process. Further work is needed by vendors, particularly in the development of dynamic compact thermal models for multi-source and multi-die packages to extend the benefits of compact thermal models into the future based on the latest research [32].

Model Order Reduction is a promising approach for system-level modeling, or model-based design, in which complex 3D models are reduced to say the order of 25 equations that can be solved in a matter of seconds. The reduced models retain key characteristics (e.g. key dimensions) from the 3D model, so they respond to changes of these values, enabling fast, multi-domain design space exploration. As yet, toolsets are not in place to take full advantage of this technology, and creating these models can be time-consuming and not particularly straightforward.

Due largely to the pace and complexity of electronics design, EC CFD is commonly used as part of a "design flow" [33] comprised of heterogeneous toolsets, with an emphasis on interfacing to upstream and downstream tools. While EDA-embedded CFD is technically feasible and does offer the potential of allowing thermal issues to be addressed earlier in the design flow, the EEs who use these toolsets expect near-instantaneous results based on their experience with finite-difference based software with imposed airside heat transfer.

Most thermal engineers come from a mechanical rather than electrical background, but are not necessarily designers, and so may not be proficient users of MCAD software. For them, stand-alone EC CFD software, supported by sophisticated MCAD and EDA interfaces may provide the best platform. They are often a scarce resource within their organizations, and this is set to get worse in the future. Adoption of workflow automation for

EC CFD is currently in its infancy, but likely to gain in importance as organizations attempt to leverage their engineering knowledge to seek competitive advantage. EC CFD has matured, and the existing full-featured toolsets and the need for continuous innovation present a high barrier to new market entrants.

10. CONCLUDING REMARKS

At the top of the food chain, the semiconductor industry is currently being reshaped by mergers, concurrent ramp ups at three process nodes: 28, 14/16 and 10nm, and a fundamental change in transistor structure – the FinFET. More sophisticated thermal solutions can be expected as companies strive to achieve cost-effective designs for new products based on these new technologies, with power densities in high-performance applications likely to continue increase, driving innovations in both cooling technology and design practices. Automotive and aerospace is a $2 trillion market, automated to the same level as chip design was in 1970 and already overflowing with electronics [34].

A lack of suitably skilled and experienced labor is expected to increase the demands on EC CFD software with the focus shifting from general simulation platforms to tools focused on solving a specific task, relieving the user of the need to understand the physics involved, or how to simplify the problem for simulation. Rather, the software should work directly with unmodified CAD data from whatever source. This will require a step-change in built-in intelligence and interconnectivity from where we are today.

REFERENCES

1. “Electronics – Correct by Design”, Benchmark Report, Aberdeen Group, February 2007.

2. Integrated Circuits Thermal Test Method Environment Conditions - Natural Convection (Still Air) EIA/JEDEC STANDARD EIA/JESD51-2.

3. Hnatiuc, M & Iov, C.J. “Cooling with Mentor Graphics HyperLynx Thermal”, Proc of 20th International Symposium for Design and Technology in Electronic Packaging, October 2014, Bucharest, Romania

4. J. S. Burdick, “Electronics Cooling at IBM Endicott”, Proc. of 1st FloTHERM Int. User Conf., Guildford UK, Sept. 1991, pp. 53-77.

5. W. N. Dawes, “Turbomachinery computational fluid dynamics: asymptotes and paradigm shifts”, Phil. Trans. of R. Soc. A (2007) Vol. 365, pp. 2553–2585.

6. Harvey Rosten, et al., “Final Report to SEMITHERM XIII on the European-Funded Project DELPHI - the Development of Libraries and Physical Models for an Integrated Design Environment”, Proc. of IEEE SEMITHERM XIII, 20-30 Jan. 1997, pp. 73-91.

7. H. Pape, G. Noebauer, “Generation and verification of boundary independent compact thermal models for active components according to the DELPHI/SEED methods”, Proc. of IEEE SEMITHERM XV, San Diego, Mar. 1999, pp. 201-211.

8. C.J.M. Lasance, “The European project PROFIT: prediction of temperature gradients influencing the quality of electronic products”, Proc. of IEEE SEMITHERM XVII, San Jose CA, March 2001, pp. 120-125.

9. L. Codecasa, V. d’Alessandro, A. Magnani, N. Rinaldi “Matrix Reduction Tool for Creating Boundary Condition Independent Dynamic Compact Thermal Models”, 21st THERMINIC Workshop, Paris, October 2015.

10. J. H. J. Janssen, L. Codecasa “Why Matrix Reduction is Better than Objective Function based Optimization in Compact Thermal Model Creation”, 21st THERMINIC Workshop, Paris, October 2015.

11. Two-Resistor Compact Thermal Model Guideline, EIA/JEDEC STANDARD JESD15-3.

12. DELPHI Compact Thermal Model Guideline, EIA/JEDEC STANDARD JESD15-4. [http://www.jedec.org/standards-documents/results/JESD15-4]

13. Tony Kordyban “Everything You Know Is Wrong - Part II”, CoolingZone Magazine, July 2009.

14. Raghav Mahalingam, Sam Heffington, Lee Jones and Randy Williams, “Synthetic Jets for Forced Air Cooling of Electronics”, ElectronicsCooling Magazine, Vol. 13, No. 2, May 2007, pp. 12-18.

15. Ioan Sauciuc, “Piezo actuators for electronics cooling”, ElectronicsCooling Magazine, Vo. 13, No. 1, Feb. 2007 pp. 12-17.

16. http://www.flopack.com

17. Rodgers, P., Eveloy, V., and Davies, M., “An Experimental Assessment of Numerical Predictive Accuracy for Electronic Component Heat Transfer in Forced Convection: Parts I and II”, Trans. of ASME JEP, Vol. 125, No. 1, 2003, pp. 67-83.

18. K. Dhinsa, C. Bailey, & K. Pericleous, “Investigation into the performance of turbulence models for fluid flow and heat transfer phenomena in electronic applications”, IEEE Trans. on Comp. and Pack. Tech., Vol. 28, No. 4, Dec. 2005, pp. 686-699.

19. Arun Raghupathy, Hoa Do, Brian Philofsky & Gamal Refai-Ahmed “Best Engineering Practice for Thermal Characterization of Stacked Dice FPGA Devices”, International 3D Systems Integration Conference (3DIC), Aug. 2015, Sendai, JP.

20. M. J. Aftosmis, & M. J. Berger, “Multilevel Error Estimation and Adaptive h-Refinement for Cartesian Meshes with Embedded Boundaries” AIAA 2002-0863, 40th AIAA Aerospace Sciences Meeting and Exhibit, 14-17 January 2002, Reno NV.

21. Matthias Benz, “Use of FloEFD in engine development at Porsche Engineering Services GmbH”, Mechanical Analysis U2U Meeting Friedewald Germany, Sept. 2011

22. https://www.mentor.com/pcb/xpedition/3d-layout/

23. “Mentor Graphics and PTC Deliver Industry’s First Bi-Directional ECAD–MCAD Collaboration Capability” Press Release, September 2008.

24. “Global Computational Fluid Dynamics (CFD) Market 2015-2019”, Technavio Report, September 2015.

25. Sarkany, Z., Vass-Varnai, A. & Rencz, M., “Effect of Power Cycling Parameters on Predicted IGBT Lifetime” IEEE Aerospace Conference, 7-14 March 2015, Big Sky, MT, USA, DOI: 10.1109/AERO.2015.7118982.

26. “The Third Wave of CFD”, Ivo Weinhold & John Parry, Proceedings of NAFEMS World Congress, Salzburg, Austria June 2013

27. http://www.brightidea.com

28. “TSMC validates Mentor tools for 3D stacking”, Electronics Weekly, September 2013.

29. “New FloTHERM Product with Automated Calibration for Package Model Accuracy”, Mentor Graphics Press Release, November 2015.

30. “Simulation 20/20: The Next 5 Years – Accessibility” NAFEMS Webinar, October 2015.

31. http://www.openfoam.com

32. Eric Monier-Vinard, Brice Rogié, Valentin Bissuel, Najib Laraqi, Olivier Daniel, Marie-Cécile Kotelon, “State of the Art of Numerical Thermal Characterization of Electronic Component”, EUROSIME 17, Montpellier, France, April 2016

33. https://en.wikipedia.org/wiki/Design_flow_(EDA)

34. Wally Rhines, “EDA Grows by Solving New Problems”, Visionary Talk, 51st DAC, June 2014.

APPENDIX: MILESTONES IN ELECTRONICS COOLING CFD

What follows is an abridged version of events, presenting only what is most pertinent to the topic of this paper, in roughly chronological order.

1970s & 80s:

1974 saw the foundation of CHAM Ltd., the first commercial company to provide a CFD consultancy service, and later software, to industry. PHOENICS debuted in 1981 as the first commercially available software tool in CFD [A1]. Electronic Design Automation also dates back to the beginning of the 1980s, when Daisy Systems, Mentor Graphics Corporation and Valid Logic Systems were all formed. Creare Inc. launched the first version of Fluent in 1983. The use of commercial CFD codes in electronics dates back to the mid to late 1980s when Dereje Agonafer introduced a number of licenses of PHOENICS into IBM Poughkeepsie. At around the same time Fluent from Creare was being used by DEC. The mid 1980s saw rapid growth in chip power in the bipolar-based digital circuitry of the day. The mid to late 1980s saw the emergence of PCB thermal design tools, one of the first being PCBTHERMAL from Pacific Numerix.

EDA companies also started to produce or market their own offerings. Mentor Graphics produced AutoTherm, Cadence acquired and marketed Thermax, and Racal-Redac, now part of Zuken Inc., produced VTAT and later their Thermal Placement tool. The mid 1980s also saw the birth of the first CFD code dedicated to EC when J. P. Bardon (CNET, France) presented THEBES at the ASME IHT conference in San Francisco in 1986. An English version became available in 1987 and was extensively tested by Philips for consumer electronics applications [A2]. 1987 also saw the formation of Innovative Research Inc. During 1988 Fluent Inc. was spun off from Creare, and Flomerics Ltd. was founded. FloTHERM made its debut in late 1989.

1990s:

In 1992, newly-founded Blue Ridge Numerics Inc. released CFDesign, a general-purpose CFD solution tightly integrated with MCAD software, and Daat Research Corp., whose flagship product, Coolit is targeted at EC applications, was founded. In 1994, taking inspiration from FloTHERM, Fluid Dynamics International (FDI) released the first version of IcePak based on its FIDAP FE solver with the user interface written by ICEM-CFD Engineering, now a part of ANSYS Inc.; Mentor Graphics acquired Thebes, which was marketed as AutoFlow; and Harvard Thermal was founded, releasing its Thermal Analysis System (TAS), a conduction and radiation tool for military and defense applications. In August 1995, Fluent Inc. was acquired by Aavid Thermal Technologies, Inc. In May 1996, Fluent acquired FDI, and in 1997 Fluent released the first version of Icepak based on the Fluent UNS solver. In 1998 Flomerics launched FloPACK, a web-based application creating thermal models of chip packages and other electronics parts for use in FloTHERM. 1998 saw Innovative Research Inc. release MacroFlow, using Flow Network Modelling (FNM) for system-level analysis. In 1999 Flomerics released the Command Center to control and co-ordinate distributed processing of multiple jobs simultaneously across a heterogeneous network, and Nika GmbH was founded, producing the first MCAD-embedded CFD product, FloWorks, marketed by SolidWorks Corp. under the COSMOS brand.

2000s:

Aavid was purchased by Willis Stein & Partners, a US private equity investment firm in January 2000. In 2001 Innovative Research launched TileFlow (3D CFD) for Data Center simulation. In 2002 Harvard Thermal began shipping TASPCB aimed at PCB designers and incorporating a CFD capability. In January 2004 Future Facilities formed as a spin-off from Flomerics to market FloVENT to the data center market. Flomerics released FloPCB and Nika's Engineering Fluid Dynamics (EFD) software became available for CATIA V5. Flomerics acquired Hungarian-based MicReD in May to provide model validation and testing services, Nika released EFD.Pro for Pro/ENGINEER in June and Daat released CoolitPCB in July. In August 2004 Harvard Thermal launched Package Thermal Designer (PTD). In October 2005, Ansys acquired Harvard Thermal, marketing TASPCB as Iceboard and PTD as Icechip. 2006 saw the release of FloPCB for Allegro by Flomerics. In May that year Fluent was acquired by ANSYS Inc. and in June Nika GmbH was acquired by Flomerics. Mentor Graphics acquired the BETAsoft product line in May 2007, which previously was developed, sold, and supported by Dynamic Soft Analysis, and is now marketed as HyperLynx Thermal. In December Flomerics released its electronics-specific module for EFD. In June 2008 Flomerics released ThermPaq, adding automated generation of package metrics to FloPACK's compact thermal model generation capability, and an Electronics Cooling module for EFD. Flomerics was acquired by Mentor Graphics' in October 2008, to form Mentor's Mechanical Analysis Division. Future Facilities announced 6SigmaET, to complement its DC offerings at SEMI-THERM 25 in March 2009 In May 2009 Ansys discontinued Iceboard and Icechip.

2010s:

Autodesk acquired Blue Ridge Numerics (CFDesign) in March 2011, and in the Summer, Harley Thermal, funded by ex-Harvard Thermal customers, released Solaria recreating TAS. Mentor Graphics released FloTHERM XT in March 2013, and a LED Lighting Module for FloEFD. Harley Thermal released SolariaPCB in November 2014, recreating TASPCB. In January 2015, Innovative Research Inc. spun out Innovative Research LLC to focus solely on MacroFlow.

The picture is then firstly one of innovation, with both new companies and new products from existing companies emerging onto the market, with dedicated tools created to address different packaging levels (system- board- and package-level). More latterly the picture has been one of consolidation, both in terms of toolsets and of companies through acquisition. In January 2016, Siemens announced it had agreed to buy CD-Adpaco for close to $1bn in cash.

APPENDIX REFERENCES

A1 Dr. Akshai K. Runchal, "Brian Spalding: CFD & Reality", Proc. of CHT-08, May 11-16, 2008, Marrakech, Morocco (CHT-08-012)

A2 Clemens J.M. Lasance, "20 Years of CFD for thermal management at Philips Electronics", Proc. of Theta Workshop, Cairo, January 6, 2007

ACKNOWLEDGMENTS

I'm grateful to Drs. David Tatchell, Clemens Lasance, Ian Clark, Robin Bornoff, Ivo Weinhold, and John Wilson for their comments and suggestions for improving this manuscript. All trademarks used in this paper are recognized as property of their respective owners.

Reprinted with the kind permission of the SEMI-THERM Educational Foundation (STEF), which is dedicated to worldwide educational opportunities and providing resources within the electronics thermal engineering community.