icons of CFD

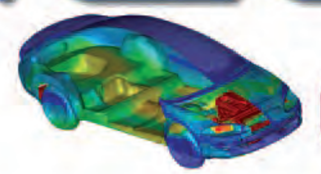

Harvey Rosten

By David Tatchell, PhD

In science and technology, as in other fields of endeavour, it is seldom the case that progress is made steadily over a period of many years. Rather, one often observes periods of intense productive activity during which the foundations are established for a subsequent period of consolidation - progress is made via a series of step changes, rather than a steady monatonic progression.

In CFD we can observe one such period of step change in the decade from the mid 1960s to mid 1970s - and it is noticeable that the "Icons of CFD" covered in this series so far made their major contributions during this period. Prior to the mid 1960s CFD even in its most rudimentary form barely existed – yet by the mid 1970s the main foundations of practical CFD had already been established. The finite-volume approach, the SIMPLE solution method, and the two-equation k-epsilon turbulence model, were devised and proven during this period - and by the mid 1970s had been successfully applied to a range of practical three-dimensional flows, both steady and unsteady. The foundations of present day CFD had been established.

As a result of this progress, the potential for CFD to provide invaluable inputs to engineering design across a wide range of industrial needs was becoming increasingly recognised. However, the CFD software available in the mid 1970s was still far from providing engineering design tools of the kind required to satisfy this need. Most software was being developed in academic and research establishments, for in-house use only. CHAM Limited, the commercial offshoot from Professor Spalding's research activities at Imperial College, had begun providing commercial CFD services in the early 1970s. However, any provision of software to customers was confined to bespoke software developed for specific needs, provided to the customer in open source form at the conclusion of a project.

Clearly (with the benefit of hindsight) this approach was not the best way to put good quality CFD software into the hands of users in industry. Because each development was effectively a one-off, there were no "economies of scale". Rigorous testing was uneconomic, and, because the user could modify the source code, would have been a waste of time - and any systematic after sales support, maintenance or updating was impracticable, and was not attempted.

These deficiencies were recognised, and addressed, during the 1980s - the second "great leap forward" in the history of CFD. This decade saw the blossoming of "CFD as a software business", with the creation of a number of commercial CFD software products (starting with CHAM's PHOENICS in 1981), and the emergence of a number of CFD businesses, focused on developing, supplying, and supporting commercial CFD software packages.

It is during this phase of the evolution of CFD that Harvey Rosten, the subject of this article, made his contributions. Uniquely, Harvey played a leading role in the creation of two of the major CFD products released the 1980s - first PHOENICS, and then Flomerics' FLOTHERM/FLOVENT product.

Harvey studied Theoretical Physics at Queen Mary College, London (1967-71), and then completed an MPhil at the Rutherford High Energy Laboratory (1972-4), on the analysis of the magnetic fields of superconducting magnets for particle generators. On the completion of his MPhil in 1974, Harvey joined CHAM Limited - and entered the world of CFD. This was when I first met Harvey.

As was explained in the earlier "Icons of CFD" article on Professor Brian Spalding, Concentration Heat and Momentum (CHAM) Limited was set up by Spalding as a means of making the outcomes of his CFD research activities at Imperial College available to industry - a direct reflection of Spalding's insistence on the early practical application of his research activities. CHAM started as a consulting operation within Imperial College in about 1970, and grew, and separated from Imperial College, during the 1970s.

Harvey undertook and led a number of projects at CHAM, including (as leader of the Environmental and Process Group) fires and smoke in buildings, glass smelting, cooling towers, gas/liquid flows in undersea pipelines, and the Hall Cell process for aluminium smelting. This range of applications (undertaken by only one of the four groups into which CHAM's work was organised) illustrates both the breadth of demand for CFD services in the 1970s, and the challenge of achieving consistent good-quality outcomes using the "bespoke software development" business model.

By 1979 Spalding recognised the need to create a single CFD code to replace the multiplicity of codes being worked on at CHAM at that time, and which could be marketed to users as a "general-purpose CFD code". This revolutionary decision resulted in the release of PHOENICS, the world's first commercial CFD code, in October 1981.

As CHAM's Software Development Manager, Harvey led the development of PHOENICS, working in tandem with Spalding. The challenge was to incorporate the best techniques that CHAM had developed over the years into a single piece of software, that would cover all of CHAM's then current needs. This meant that the technical requirements for the first release were, to say the least, challenging. The name tells part of the story - PHOENICS stands for Parabolic, Hyperbolic Or Elliptic Numerical Integration Code Series - meaning that it had to work well for one-way boundary-layer flows, for supersonic flows, and for elliptic recirculating flows. It had to handle 3D flows, steady and transient, laminar and turbulent flows, heat transfer, combustion and chemical reaction, and dispersed two-phase flows. Quite a list by any standards!

Another challenge was designing the user interaction with the software so as to achieve the conflicting requirements for a) a single code which could be developed, tested and supported centrally, and which would not be available for the user to modify, and b) the need for users (at CHAM or in customer organisations) to adapt the code to their specific problems and needs. This led to a novel arrangement, in which PHOENICS comprised three parts:

- The "Earth" program - containing all of the general CFD and physical modelling capabilities, and which was developed and maintained by CHAM. All users used identical Earth programs - interacting with it via the other two parts.
- "Satellite" programs - which were where the user specified input data, which was then transferred to Earth as a data file.
- "Ground Stations" - which were user accessible attachments to Earth where users could insert their own coding to interact with or modify the functionality of Earth in any required way. Clearly the key to this was providing access via the Ground Station to what was happening in Earth at the appropriate stages in the solution sequence, and to enable the user to make modifications in as controlled a way as possible.

In view of the ambitious specifications for PHOENICS, it is surprising that Harvey, with Spalding and the handful of developers working with them, completed the first release so quickly. The first installation was at Century Research Centre in Tokyo, on their newly-installed Cray 1 computer. As the only CFD software available at the time, PHOENICS attracted considerable interest, and use grew rapidly during the early 1980s. And, following CHAM's lead, other commercial CFD packages followed - FLUENT from Creare, FIDAP from Fluid Dynamics International, FLOW3D (later renamed CFX) from AEA Harwell, and STAR CD from Computational Dynamics, among others.

By the late 1980s CFD was becoming established alongside other CAE methodologies as an engineering tool with enormous potential, and was attracting interest from an increasingly diverse range of industries and applications. Naturally much of the demand was from the traditional CFD "core markets" of aerospace, defence, power generation, chemical process, automotive, and so on. But in addition, interest was being generated in newer CFD markets, such as building ventilation, electronics cooling, food processing, and consumer goods. However, for such "new applications", while the potential of CFD was clear, what the software of the day would deliver was disappointing.

The difficulty was not with the capabilities of the CFD technology itself - rather, with the accessibility of the software then available. Unlike the CFD core markets, these newer industries were unwilling to invest in CFD specialists able to master the complexities of working in their own coding in, for example, the PHOENICS Ground Station. Rather, what was being sought was CFD-based solutions that could be used by design engineers as a routine part of their day-to-day activities.

This recognition led, in 1988, to the formation of a new CFD business, Flomerics Limited, of which Harvey was co-founder and Technical Director.

Flomerics set out to address the needs of these "new industries" by focussing on selected applications, and developing software addressing only the specific problem in question - specifically, electronics cooling in FLOTHERM and building ventilation in FLOVENT. It was intended that, by focussing in this way, it would be possible to create software that could

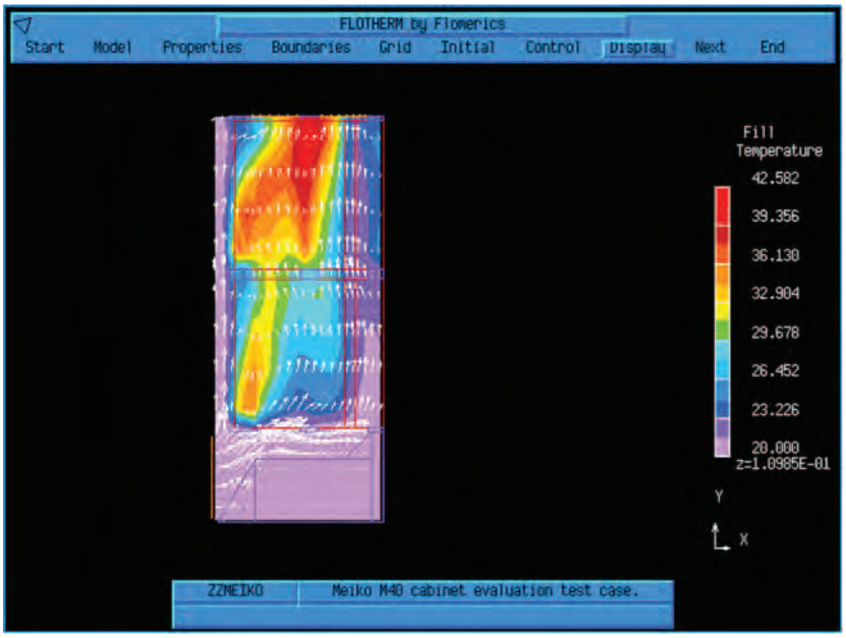

be used successfully by design engineers with virtually no knowledge of CFD.

This led to a software design using a single self-contained graphical user interface through which the user performed all input, control, and post-processing operations, with no other user access whatsoever to data or to coding. Crucially, only what was required was provided, and so in contrast to the approach of other CFD developers – who at the time were delivering general-purpose CFD software - the user's choice of CFD options was kept to the absolute minimum.

Harvey led the Flomerics Development Team from the company's initiation in mid 1988 - indeed the development team (him and two others) operated initially from the top floor of the Rosten family home in New Malden. Harvey played a major part in the design of the software, he led the development of the software, and he himself developed the core solver.

FLOTHERM was released late 1989 and FLOVENT early 1990. The GUIs and the underlying functionality were tailored to the each application, so that they appeared as two separate codes, but the bulk of the code was common to both, so that development and testing could be shared.

FLOTHERM was by far the more successful. Interest was driven by trends in the electronics industries beginning in the 1980s, and continuing to the present day, for accelerating increases in system power and functionality (Moore's Law), and for more compact equipment. These lead inexorably to escalating power densities (more watts/cubic metre), and hence challenges in cooling system design, which were beginning to be recognised in the late 1980s. Typically the mechanical engineers in computer, telecoms and avionics companies who were tasked with (among other things) devising cooling system designs, were finding that past experience and hand calculations were not enough. Prototype equipment was failing thermal testing - and causing expensive delays in product release.

FLOTHERM was designed to satisfy this need - and was increasingly adopted by electronics companies both as a solution to critical problems revealed in testing - and, longer term, as a means of "designing in" the thermal solution, by using thermal analysis at an early stage in the design process, in a manner consistent with the idea of "concurrent engineering" which was emerging at that time.

Harvey continued to lead FLOTHERM/FLOVENT development until 1992, when his work took an unexpected, but related, turn.

It turns out that, in analysing the thermal behaviour of electronics equipment, the modelling of the electronics packages themselves can be crucial. These deceptively simple looking objects have complex internal structures - driven by electrical, mechanical and thermal needs. The main heat sources in electronics equipment are the silicon chips at the heart of these packages - and the main thermal problem is generally the overheating of these chips. The thermal design requirement is thus often expressed as a "maximum junction (i.e. chip) temperature".

Consequently, by the early 1990s the need was being recognised within the electronics industries for a reliable, efficient, standardised way of representing the thermal behaviour of the multitude of package types in use at that time. Harvey took up this challenge.

This resulted in a series of EU-funded collaborative research projects, involving a number of major European electronics companies. Harvey coordinated these, and led Flomerics' contributions. The outcome was an agreed standard methodology for creating and verifying "behavioural models" of electronics packages(Ref 1), which has now been widely adopted as the "industry standard" (Ref 2).

A long way from CFD? To a purist possibly. But in practice, just as a new combustion model (for example) might open up new CFD applications in furnace or fire modelling, so these advances in package-level modelling have proved to be an essential part of realising the full potential of CFD in electronics thermal design.

Tragically Harvey died in 1997 after a short illness. He was posthumously awarded the 1998 IEEE SEMI-THERM THERMI Award in recognition of his contributions to electronics thermal modelling. The annual Harvey Rosten Award for Excellence in the Physical Design of Electronics was instituted in 1998 in his memory (Ref 3).

Harvey once expressed his contribution as "making good science available to industry" - which seems to encapsulate his career nicely. Uniquely, Harvey played a leading part in the creation of two major, innovative, world-leading CFD software products - which have made CFD techniques ("good science") available to many thousands of users in industry and academia throughout the world.

References

[1] 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.

[2] 'DELPHI Compact Thermal Model Guideline' JEDEC Document JESD15-4, October 2008. www.jedec.org/sites/default/files/docs/JESD15-4.pdf

[3] www.rostenaward.org.

AUTHOR INFORMATION

David Tatchell
david.tatchell@tiscali.co.uk

with thanks for the contribution from John Parry, Mentor Graphics.