

CFD – Bridging The Communication Gap

James J. Quirk¹

Center for Simulation of Advanced Rockets
Digital Computer Laboratory MC-258
University of Illinois at Urbana-Champaign
Urbana IL 61801, USA
jjq@amrita-cfd.org

Abstract

Modern computational fluid dynamics (CFD) is a complex amalgam of numerical analysis, flow physics, computer science, not to mention common sense. As such, it is becoming increasingly difficult for individual workers to remain wholly self-sufficient. This talk will explore some of the inter-disciplinary, communication problems that must be overcome, before this relatively new research field can reach full scientific maturity. To avoid undue subjectivity, the talk will present a series of numerical case studies which serve to illustrate the fundamental changes of perspective in moving from algorithm development, through software implementation, to end application. These case studies have been automated using a purpose-built, computational operating-system and can be reproduced – in their entirety – by any interested party, for no cost other than raw CPU-cycles. The talk will explore how such automation opens up the possibility of improving simulation capability through massed scrutiny: thus doing away with the present cottage-industry mentality, and its concomitant vagaries, where “codes” are crafted on a one-off basis. The present extended-abstract provides some pre-requisites for understanding this computational philosophy so that the talk itself can concentrate on specific applications. Full details of the general approach can be found at <http://www.amrita-cfd.org/FAQ>.

Introduction

There can be little argument that CFD has matured to the point where it earns its keep across a broad range of applications (witness <http://www.cfd-online.com>). Consequently, the provisional title for this talk – *What CFD Can Do For You* – is clearly outdated. Interestingly, this title was chosen by a colleague (an applied mathematician) who graciously stepped into the breach, following my unavailability, when the conference programme was being drawn up. Upon reflection, possibly a more appropriate title for the times would have been – *What You Can Do For CFD*. The very same words, but a talk with a completely different message to convey.

Taken together, the two prospective titles illustrate the gulf in thinking between a CFD software-developer and a CFD end-user. Whereas the typical user begs-borrows-or-steals “a code” for a one-off simulation, focusing firmly on the flow physics of the application, the author of “a code” looks to spread development costs across a wide range of projects with different sets of users, and by necessity is often forced to concentrate more on software issues than on physical issues. Thus to a user: “a code” is purely a means to an end and is often viewed as a disposable entity whose intellectual content may not even be acknowledged. But to a developer: “a code” represents a considerable investment of effort, much of it emotional, and each new user (although welcomed) adds to the programming burden. In this extended abstract it seems appropriate to identify concrete mechanisms for bridging such communication gaps: because, ultimately, miscommunications between related disciplines act to limit CFD’s potential. Hence the actual title of the talk.

The first step in finding a communication-solution is to recognize that there is indeed a communication-problem to solve. For instance, in response to:

*considerable concern with the quality of published numerical results*²

the American Institute of Aeronautics and Astronautics (AIAA) have issued the following editorial policy:

The AIAA journals will not accept for publication any paper reporting (1) numerical solutions of an engineering problem that fails adequately to address accuracy of computed results or (2) experimental results unless the accuracy of the data is adequately presented.

¹Formerly of Graduate Aeronautical Laboratories, California Institute of Technology, Pasadena CA 91125, USA. Now in transit to Los Alamos National Laboratory, New Mexico, USA.

²This quote, along with the next three quotes are taken, without permission, from AIAA information sheet: “Numerical Accuracy and Experimental Uncertainty *AIAA Journals*.”

The effectiveness of the AIAA policy hinges on their chosen definition for numerical accuracy:

The accuracy of the computed results is concerned with how well the specified governing equations in the paper have been solved numerically. The appropriateness of the governing equations for modelling the physical phenomena and comparison with experimental data is not part of this evaluation. Accuracy of the numerical results can be judged from grid refinement studies, variation of numerical parameters that influence the results, comparison with exact solutions, and any other technique the author selects.

and the extent to which they enforce the diktat:

For computational papers, the author must provide an adequate description of the numerical solution procedure, if not documented elsewhere. In addition, the complete governing equations must be specified with sufficient detail along with the input parameters to the code so that a reader could reproduce the results of the paper.

If CFD were a completely mathematically rigorous discipline then there would be no need for the above policy statement, as any investigation worthy of publication would automatically be reproducible via an unbroken sequence of logical steps. Unfortunately, the typical CFD investigation contains numerous software mundanities which act to limit its integrity. This is especially true of modern algorithms, such as adaptive mesh refinement schemes, which do not have a concise mathematical recipe but achieve their ends through convoluted programming. With the best will in the world, the more a CFD investigation relies on custom-written software, the harder it is to satisfy the diktat that the work should be reproducible by the reader. Even supposing the necessary information were available, the software-engineering needed to redo the work would prove prohibitively expensive for all but the most persistent. On this score, matters are compounded by the apparent unwillingness of the scientific community to embrace the need to train its members to program efficiently: a criticism which deserves some explanation.

Casual programmers (i.e. individuals with no formal training in software construction) who switch from an old-fashioned language, such as *Fortran*, to a modern one, such as *C++*, often badger their peers to follow suit: flushed by the successes of an improved programming style. But before doing so they should ponder this assertion, made by the architect of *C++*[4] (my underlining):

... as programs get larger, the problems associated with their development and maintenance shift from being language problems to more global problems of tools and management.

In other words, there is much more to constructing software than selecting a programming language and learning to use its syntax. This perhaps explains, but not pardons, why many “a research code” is unavailable for inspection: the author is embarrassed by its construction.

A similar assertion can be made about CFD simulations:

... as simulations get larger, the problems associated with their development and maintenance shift from being algorithmic problems to more global problems of tools and management.

The upshot being that if the right *tools and management* were put in place then many of the concerns which prompted the AIAA computational policy could be tackled head on. For instance: grid refinement studies, numerical sensitivity studies, and comparisons with exact solutions, are all time consuming activities, and so it should come as no surprise that they are not performed as often as they should be. But if such activities could be fully automated then comprehensive validation tests could be run routinely, and even standardized. Of course, recognizing the need for the types of *tools and management* to do this is a far cry from actually producing them, and an even farther cry from successfully convincing a conservative community of the need to change its work practices.

Back in 1995, I started converting an adaptive mesh refinement algorithm into an interactive teaching aid so as to allow students to explore the practical aspects of compressible, computational fluid dynamics. The aim was not to turn experimentalists into heavy-duty programmers, but to produce a hierarchical, open-ended system in which individuals, reluctant programmers and CFD experts alike, could engage in unambiguous, scientific dialogue. Over time, *Amrita* has grown into a full-blown operating-system for automating numerical investigations[3]. This system, although far from perfect, demonstrates how the AIAA diktat could be realized in a methodical, watertight fashion.

The remainder of this extended-abstract will attempt to add substance to this claim by presenting a concrete numerical example. But before doing so, it is worth noting that the current *Amrita* installation-kit runs to over 520,000 lines of source-code, and so it provides a *tool* to enable any interested party to navigate the source in a painless fashion: see <http://amrita-cfd.org/source-code>. An examination of the source reveals that *Amrita* could legitimately be called any of following: an educational aid; a research tool; a communication protocol; a document preparation system; a quality assurance system; an archive system; a device-independent compilation system; a literate-programming system; a run-time shell for “legacy code”; a computational operating system; a labour saving device; a computational philosophy. In fact, almost anything *except* “a CFD code.”

AmritaMailit::run_T5

Experience shows that reluctant programmers are easily put-off *Amrita* by the sheer size of the system. But the system is large precisely because it is striving to take care of all the tasks a reluctant programmer should not be expected to cope with. It also comes as a shock to many that the system has its own programming language (also called *Amrita*, but note the typographic distinction). The system needs its own language to be able to articulate the wide range of tasks needed to automate a numerical investigation. For instance, this six line *Amrita* fragment:

```
fold::file $abstract -> listing
fold::latex { typeset listing
  \begin{verbatim}
    *listing
  \end{verbatim} }
}
```

produces this listing of the abstract you are currently reading:

```
... ICDERS top matter
Latex2eHead {
  dir = 4ICDERS
  file = jjq.tex
}
... title
... abstract
... introduction
... AmritaMailit::run_T5
... closing comments
LatexTail
Latex
```

Several observations are needed to preempt the inevitable question – *Why bother?*

First, re-read the AIAA diktat; it effectively stipulates that precise and complete documentation is an integral part of a numerical investigation. Here, for preciseness the *Amrita* fragment is typeset from the actual working script, and in turn the script-fragment typesets the abstract from the original³. In so doing, the abstract is *folded* so that its structure is immediately apparent to the reader⁴. A computer scientist would call *Amrita* “a literate programming language,” such languages are designed to produce self-documenting programs. The rationale being that manually prepared documentation is both onerous to produce and rarely a true reflection of a program’s innards, but if the process is automated then it becomes straightforward to produce error free documentation which evolves as the program evolves. Through its use of *program folds*, *Amrita* could also be called “a cooperative programming language.” If it needs to utilise a specialist language to accomplish some task, the new language can be folded into *Amrita* as $\LaTeX$ is folded in here. In essence an *Amrita* program can be viewed as a tree in which each branch can, if needs be, employ a different parser. The main implications of this are: *Amrita* cannot become outdated and superseded by another language, because with the right plumbing it can always fold in the new language (e.g. `fold::python`); reluctant programmers need only unfold an *Amrita* script (fed to them by an expert) to the level of complexity they feel comfortable with; all the activities needed to *validate* a numerical investigation: setting up a problem, running a simulation, plotting results, documenting the method etc., can all be freely intermixed within a single program.

Amrita is not a random collection of tools and much of its utility stems from the seamless manner in which it operates. For instance, consider the *AmritaMailit* shown in Figure 1 (a), it contains all “the user-supplied code” (albeit in an encoded form) needed to cajole the system into producing Figure 2. The packing and unpacking of an *AmritaMailit* takes place automatically in such a fashion that anyone with access to an *Amrita* installation⁵ can now springboard directly off the efforts of the original author. In this case: Prof. Hans Hornung, Director of Graduate Aeronautical Laboratories California Institute of Technology.

³The following is given for completeness, but need not be fully appreciated to follow the main argument. `fold::file` is an *Amrita* keyword which takes the contents of the file identified by `$abstract` and *folds* it following a specific set of rules, depending on the filetype, and stores the resultant text in a *string token* called `listing`. `fold::latex` is a keyword which embeds a block of $\LaTeX$ into the *Amrita* script. The `*` is an operator which expands the contents of `listing` much as `*` dereferences a pointer in C. Full details are given in [3].

⁴`Latex2eHead` is an *Amrita* standard-library procedure which constructs the head of a $\LaTeX$ document – here `jjq.tex`. `LatexTail` writes the tail of the document, and `Latex` runs *latex*, *dvips* and *bibtex* – having first requested *bibitems* from a series of local and remote servers – to produce a *PostScript* file `jjq.ps`.

⁵An *Amrita* kit can be downloaded by anyone with a *genuine commitment* to see the system succeed. Interested parties can obtain a 90 day licence from `licence-mgr@amrita-cfd.org` which automatically rolls over to a life-time licence upon receipt of a suitably polished *AmritaMailit*. Before requesting a licence, please make sure the system will run on your machine: for a list of platforms, see <http://www.amrita-cfd.org/FAQ>.

(e) The *Amrita* fold within *MakeT5Image.amr* which pastes-up the experimental and numerical images into a single figure: *paste* is happy processing both local files and *urls*; ditto for most other *Amrita* commands (e.g. *solver*). Where necessary, digital-signatures are used to safeguard against “Trojan-horse code.”

Figure 1:




Figure 2: Output of *AmritaMailit::run_T5*. The images show experimental shadowgraphs of a one-inch diameter nylon sphere fired into propane, with results from numerical simulations superimposed (see [2] for details). The AIAA definition on numerical accuracy is careful to distinguish the *evaluation* of a simulation by comparison against experiment (as done here) from a true *validation* of a simulation which necessitates comparison against an exact solution for the model employed.

The present *AmritaMailit* unpacks itself to the source tree shown in Figure 1 (b), and the main driver routine is shown in Figure 1 (c). All told, there are just over 300 lines of *Amrita*, which for the competent script writer represents about one day's effort. *Amrita* scripts are short because they merely schedule work to be done by a plugin computational engine, which is then trusted to carry out the work as it sees fit. But this does not mean that an engine is necessarily a black-box. In the case of *Amr_sol*[3], which is built around a block-structured mesh refinement algorithm, the instruction:

```
solver code/body_roe
```

causes a *patch-integrator* to be linked dynamically with the engine (*unix-prompt*> *man dlopen*). Dynamic-linking is a powerful programming technique for constructing programs which can bolster their capabilities on-the-fly. Here *body_roe* is a specialist component for integrating the two-dimensional Euler equations on a body-fitted grid. It amounts to 388 lines of *Fortran* and is the one component within *Amrita* which is close to the classical notion of “a CFD code,” but with one crucial distinction. All the work needed to schedule the *T5* simulation and process its results live external to *body_roe*. Thus it can be thrown away and replaced by another *patch-integrator* of the user's choice without affecting the leg-work needed to produce Figure 2.

Amr_sol could also be termed “a code,” and is also a replaceable item within the software hierarchy, but at over 40,000 lines, much of it dense logic, the label starts to become derogatory. On the other hand, to call *Amrita* “a code” would be as nonsensical as calling *UNIX* “a code.” But some observations are in order lest you think this is just playing with semantics. *Amrita* is a general-purpose, user-programmable system for orchestrating simulations and it is not tied to a single computational methodology (AMR or otherwise). Nor is it tied to solving a single set of equations. Here *EulerEquations* is a standard-library procedure which teaches the plugin a set of symbolic mappings for writing to, and reading from, the computational domain using notation: *RHO*, *U*, *V*, *P* etc. The procedure could have been written by the user, but in the interests of standardization the system provides a definitive version. Thus to satisfy the part of the AIAA diktat which deals with *governing equations*, the clued-in scriptwriter could get away with:

```
grab::info BCG:Latex::[Document] from file EulerEquations.amr
parse token Document
```

which grabs a block of folded-information from the file containing the library procedure and then executes it to typeset up the equations. Inside the script, the named information was introduced using:

```
fold::info'Document {
... body of fold deleted
```

where the identifier *Document* is a moniker chosen by the programmer. Thus the *Document* fold lives inside a *Latex* fold which in turn lives inside a *BCG* fold.

The primary customer of the BCG information is a standard-library routine `BasicCodeGenerator` which constructs CFD components upon demand ⁶. Here the routine `InitSomeResources` runs:

```
EulerEquations
plugin amr_sol
BasicCodeGenerator {
    solver      = body_roe
    scheme      = flux-limited'operator-split
    grid        = body-fitted
    document    = yes
}
```

to produce – on my office machine – the shared-object `code/AMRSO/serial/Solaris/sparc/body_roe.so`. In addition, a 30 page document is constructed which explains the `solver`'s construction. This document can be used either by a student to learn the basic construction of “a CFD code,” or it can be used by an expert to see how to plumb in “third-party code.” In principle, any explicit integration scheme, developed for a topologically rectangular patch, is a candidate for linking with `Amr_sol`. The motivation being that “a serial, single patch code” once linked with the engine becomes – on a parallel machine – “a message-passing, adaptive mesh refinement code.” This is possible because all the AMR machinations and inter-node data transfers are orchestrated external to a patch; a strategy which proves efficient for the target applications for which `Amr_sol` was *designed*[3]. Of course, many details take place behind the scenes to allow this to happen, and therein lies one reason for the seemingly bizarre emphasis of this abstract for a colloquium on the dynamics of explosions and reactive systems.

Even with my extra page allotment, over the standard length abstract, there is not enough room to fulfill the AIAA diktat on a meaningful problem. This statement hinges on the interpretation of *reproduce*; the one point where I part company with the AIAA and most other individuals in the CFD community. The thrust of this paper has been to demonstrate that with the right *tools and management* a CFD investigation – as distinct from an isolated computation – can approach the reproducibility of a theoretical investigation. There is no need to be resigned, as the AIAA appears to be, and bracket numerical investigations with experiments in the reproducibility stakes. Simulations are inferior to experiments in many ways; their superior strengths, however, are portability and repeatability. The fact that CFD developers and end-users are not exploiting these strengths to the full will not result in an apocalyptic spate of planes dropping out of the sky. The problem is more insidious than this and essentially pertains to the division of labour needed to advance CFD to scientific maturity.

Judged by the teraflop yardstick – beloved of computational grand-challenges – Hornung's computations are small beer, because they are so cheap to run. But his intellectual contribution lies not in the CPU power needed to run the simulations, nor with the construction of 300 lines of `Amrita` script, but in the physical observations drawn from an extensive numerical parameter study[2], of which just two results are shown here. Although the experiments shown in Figure 2 are essentially inert, they were inspired by an earlier set in which the projectile initiated a detonation wave[1]. Therefore it is pertinent to ask what script changes would be needed to be able to simulate these other experiments: the answer depends on the choice of reaction model.

`EulerEquations` could be changed to `ReactiveEulerEquations` and given the reaction parameter `model=1step-Arrhenius`, for which `BCG` could be coaxed into generating one of 100+ different *patch-integrators*. Then, after the addition of a heat release and an activation energy the reactive simulation could be run as readily as a homework problem given in an introductory course on CFD – *How?* `Amrita` is designed to allow simulations to be scheduled independently of the mathematical complexity of the physical model, and independently of the cost of the simulation ⁷. The clued-in scriptwriter could incorporate the necessary changes in a literal 5 minutes. But if this reaction model proved inadequate then life becomes more difficult. A change could be made to `model=3step-ChainBranching` which – for the current installation kit – would reduce `BCG` down to just one `solver`, which may or may not be up to the job. Then if further improvements to the reaction model were needed to simulate the experiment, the 5 minute job has turned into a four year PhD thesis: unless outside expertise can be brought in to bridge the gap. Hence the need for “a cooperative computing system.”

In many instances, CFD progress is now limited not by a lack of ideas but by the sheer software effort needed to implement new ideas. To combat this malaise, `Amrita` seeks to provide incentives for third-parties to want to work in situ, by offering the *tools and management* that are needed to push CFD forward, but which have now outgrown a researcher's capacity to craft them singlehandedly; at least in a reasonable time frame. This is done in the hope that specialist components then filter back to improve the system. So that over time, without compromising numerical accuracy, more and more applications become “5 minute jobs.” Thus `Amrita` acts as a conduit between developers and end-users, neither of whom can survive without the other.

⁶“*BCG* codes” are constructed routine by routine, much as this document is constructed paragraph by paragraph. In fact the same “literate-programming techniques” are used in both cases. Interestingly, a reluctant programmer will often utilize the convenience of *BCG* while at the same time condemning the low-level *Amrita* constructs employed as system-overkill.

⁷Serial or parallel, 2D or 3D, `LinearAdvectionEquation` or `RelativisticEulerEquations`, makes no difference to the way the work is scheduled in an *Amrita* script and so has no bearing on either the length or complexity of “user-code.” On the other hand, “the system-code” needed to support this utopia can be extremely onerous to produce.

Closing comments

This abstract is too short to cover all the subtleties of the communication gaps within CFD. For example, a computer scientist recognizing⁸: folds from *Occam*; literate programming from *Web*, or *Icon*; dynamic linking from *Perl*, or *Python*; binding generation from *Swig*; source archiving and navigation from *Cvs*; and so on, might conclude that *Amrita* has nothing, intrinsically, new to offer. But [3] carefully articulates a similar conclusion, together with reasons why the system is not proffered as a panacea. However, a case can be made for an *Amrita*-style system, not because there is a lack of existing *tools*, but because the existing *tools* are too arcane for casual programmers whose intellectual interests lie outside computer science. For instance, when a user issues a `solver` command, *Amr_sol* does not limit itself to linking in a shared-object as would a typical *Perl*, or *Python*, script. It goes on to perform a comprehensive series of checks to ensure that the *patch-integrator* is consistent with the current state of the `plugin` engine. The need for such *management*, over and above raw language features, lies at the root of Stroustrup's assertion[4]. Moreover, the very individuals who benefit most from such safety features, are the reluctant programmers who would have most difficulty articulating, by themselves, the reams of *Perl*, or *Python*, needed for said checks. Thus, at one level, *Amrita* is striving to be a *Mathematica*-cum-*Matlab* for CFD, albeit with an altruistic mission along the lines of *Linux*⁹.

At another level, *Amrita* is striving to provide researchers with a software communication-protocol for propagating CFD-expertise out to community, whether this involves the introduction of new algorithms, or devising standardized tests which reveal weaknesses of current approaches. On this score, it is important to dispel two popular misconceptions: (i) standardization limits intellectual freedom; (ii) simulations which takes days to run, if not weeks, do not need automating. One, CFD is already subject to standardization, whether it be the AIAA diktat on numerical accuracy, or the IEEE-754 rules for floating-point arithmetic; a society cannot survive without some elements of standardization. Two, the longer an end-application takes to run, the more important it is that "the code" is subjected to automated *validation* tests between runs so that "code rot" is detected before it wastes large amounts of CPU time and results in false scientific conclusions. Moreover, if computing power continues to improve at present rates: today's cutting-edge, research problems will eventually become tomorrow's homework problems and so the sooner they are automated, the better.

To come full circle, the answer to – *What You Can do For CFD* – depends on your disposition. A first-step, perhaps, would be to recognize that CFD has become a victim of its own success: there are far too many algorithmic variations on a theme to know precisely when to use one method in preference to another – witness the 100+ schemes that *BasicCodeGenerator* can produce for the *ReactiveEulerEquations*; the all important "code" – which is getting more complex by the year – is rarely open to scrutiny, and so too many claims have to be taken on trust for a discipline with pretensions of mathematical respectability. A second-step, at least for reluctant programmers, would be to recognize that CFD's problems can only be solved through software *tools and management* that transcend specific physical applications, but since the *tools* are essentially labour saving devices, it pays to learn to use them productively: there can be no CFD gain, without at least some software pain, whatever route is chosen. In the case of *Amrita*, the talk will present a collection of numerical investigations: pulsating detonations, detonation diffraction and reflection, which can be picked up by interested individuals to springboard off, adding intellectual content on top of what the present author can achieve singlehandedly.

If you are unswayed by any of the preceding lines of argument in this abstract, and you are unwilling to substantiate counter arguments using *AmritaMailits*¹⁰, then it is probably fitting to end on a light note with an old engineering joke – *It takes a crank to start a revolution!*

Acknowledgement

I am indebted to Hans Hornung whose enthusiasm for *Amrita* has been a constant source of inspiration. I would also like to thank John Buckmaster and Mark Short for giving me the opportunity to get on my CFD soap-box, although I accept full responsibility for any criticisms which come as a result.

References

- [1] J. Bélanger, M Kaneshige and J. E. Shepherd. Detonation initiation by hypervelocity projectiles. In *Shock Waves*, (eds. B.Sturtevant, J.E.Shepherd & H.G.Hornung), Proc. 20th ISSW, World Scientific Press, 1995.
- [2] H. G. Hornung. Shock layer instability near the Newtonian limit of hypervelocity flows. Proc. 13th Australasian Fluid Mech. Conf., eds. M. C. Thompson & K. Hourigan, Monash University, Dec. 1998, 111-118.
- [3] J. J. Quirk. (i) *Amrita* : A Computational Facility (for CFD Modelling). 72 pages. (ii) *Amr_sol*: Design Principles and Practice. 81 pages. In "29th Computational Fluid Dynamics," von Karman Institute Lecture Series, edited by H. Deconinck. ISSN0377-8312 1998. Available at: <http://www.amrita-cfd.org/doc>.
- [4] B. Stroustrup. The C++ Programming Language (2nd ed. p.7). Addison-Wesley, 1991.

⁸If any of these *tools* are unfamiliar to you, a web-search will throw up as much information as you can handle.

⁹See: <http://www.mathematica.com>, <http://www.mathworks.com> and <http://www.linux.org>.

¹⁰This paper is itself generated by an *AmritaMailit*, with the specific aim of allowing flaws – found by critical third-parties – to be corrected.