CRS4-TEC-REP-2010VM03

Free surface simulation with starccmV4

Framework and preliminary simulations resume

V. Moreau,

CRS4, Centre for Advanced Studies, Research and Development in Sardinia

January 28th 2010
Version 1.0

Abstract

The aim of this note is to gather some information on how to run free surface simulations with starccmV4.

Contents

1 Introduction

In the framework of the EUROTRANS IP, a Myrrha-like Target has been foreseen for the XT-ADS. This target presents several difficult features from the CFD point of view. Mainly, it is a two-fluid flow composed of liquid Lead Bismuth Eutectic (LBE) and an extremely rarefied gas which can be considered as an almost perfect vacuum.

In the three feeder design, because of material issues, local velocities should remain below 2.5 m/s, the nominal flow rate being 13 l/s.

2 Technical framework

The free surface simulations are to be performed with starccm+ versions 4.02 and successive. The geometry is elaborate using stardesign. It can be exported to starccm+ in parasolid format. Unfortunately, only relatively simple (block) geometries can be effectively exported. Otherwise,

the mesh can be directly assembled in star-design and, after initialization, the case can be saved in starccm+ format.

3 Theoretical issues

The theoretical framework is based on the multi-fluid representation. That is, all the phases share the same velocity field. As we work with turbulent fields, we work with averaged quantities. The scalar transport equation makes use of the Reynolds averaged velocity. This velocity has the good property of having a clear constraint, with is free-divergence when each fluid is of constant density. However, there is some lack of description leading to suppose that the usual momentum equation is used. But, this equation considers a Favre averaged velocity field, in which the phase velocities are weighted by their mass fraction, and not by their volumetric fraction as in the Reynolds averaging. What are the consequences of this inconsistency is not yet clear. As both averaging are (quite) different only at the interface, it can only affect the behaviour of this interface.

What is observed is that the interface has a slight tendency to diffuse, whether it is due to the former inconsistency or to the diffusive-like behaviour of the turbulence modelling. The tendency to diffuse an interface has been contrasted by using a specific interface capturing convection scheme. This scheme performs quite well in a number of circumstances, but tend to fail in presence of strong flow gradients, that is, in presence of shear, curvature and stagnation point.

From a mathematical point of view, we can see that the multi-fluid approach, which considers a priori sharp interfaces, has no mechanism integrated to re-contract an interface that has been previously smeared. This is independent of the reason why the interface has been smeared, whether it is   correct or not, or if the motive is physical (turbulence) or numerical. In fact, the only mechanism incorporated is the numerical scheme, out of physical understanding.

4 Surface sharpening strategy

As explained before, the surface sharpening strategy actually implemented is numerical in nature and is the surface capturing convection scheme. It is able to recognise a relatively sharp interface and re-contract it. But if the interface has been too much smeared (how much is an issue), then it looses efficiency.

In the VOF+ algorithm, developed by FZK, a complementary strategy has been implemented. First, we describe shortly the VOF+ algorithm. It is based on the combined used of the interface capturing scheme, the cavitation model and an in principle contracting body force localized on the potentially smeared or smearing interface. First, in collaboration with the CD-Adapco developers, the interface capturing scheme has been raised accessible for the cavitation region interface with the cavitating fluid. Second, a body force has been applied in the interface region pushing the fluid in the direction of large volume fraction for small volume fraction and reciprocally. Typically, the force may be put under the form of the gradient of the product of the volume fractions eventually weighted. This strategy has been implemented and has shown to work quite well.

At first glance, it seems that it is the force which is directly responsible for the interface contraction, but it may not be really like this. In fact, the force is likely to be exactly compensated by a variation of the pressure field at the interface without any direct influence on the velocity field. So, what is the real mechanism? When the force is the gradient of some function, the pressure induced variation is equal to the function. As the function must be positive, and null when one of the volume fraction is null, so is the pressure variation. This means that the pressure variation is an overpressure localised on the interface. As by definition, the pressure at the cavitation side of the

interface is exactly the vapour pressure of the fluid, the pressure on the interface region tends to be higher than this value, and the cavitation vapour reduction mechanism is activated, unless in presence of a strong local pressure gradient still imposing a negative pressure at the interface and therefore vapour production. Equilibrium is likely to appear when a slight vapour production in the vapour bulk is compensated by a stronger vapour sink at the interface. The effect is to create an artificial vapour flow towards the interface. The velocity field is no more divergence free and pulls back the fluid possibly smeared.

Understanding this mechanism, we should be able to reproduce the same effect with a more direct control and by passing the cavitation algorithm when it is not specifically needed, for instance in presence of open interface.

Basically, in open configuration, it is enough to set a light gas sink localized on the interface. The function from which the former force is the gradient performs perfectly well (and its rather confusing). The lacking light fluid is then compensated through some available boundary. The stagnation inlet boundary seems for now the better suited.

In closed configuration, the situation is more delicate because we have a global conservation problem. And in fact, the VOF+ algorithm is seemingly better and well suited here. If one has only to contract the interface, then a light fluid sink is enough if there is some mechanism able to produce light fluid elsewhere, such as the stagnation inlet. Otherwise, the fundamental incompressibility of the fluids will create problems.

What we have in fact to do is to mimic the totally compressible behaviour of vacuum using an incompressible light fluid and adequate source/sink terms in such a way that the flow field of the heavy phase is the best possible. In some sense, as the light fluid representing vacuum is almost only an error. We can trick freely on it as long as is reduces the error induced on the heavy fluid.

5 Test cases

Several test cases have been built to test the algorithm variants. They are designed strongly in relation with a possible future Myrrha-like target. On another hand, it may happen that we do not manage to simulate such a target in a satisfactory way. Therefore, we have also in mind to derive a back-up solution which is less demanding from the CFD point of view. If the TRASCO3 Italian MIUR project finally starts (after 8 years or so…) we will have to dimension and simulate a channel target in the continuation of the XT-ADS channel windowless target. The experience gained here could be useful there…

The first test case is based on one of the axial-symmetrical designs of the XT-ADS water experiment. Performed on a 2D basis with triangular cells in StarcdV4, it has not be possible to stabilize the simulation, see Figure 1. The flow behaviour near the symmetry axis is heavily constrained by the imposed exact flow symmetry, leading to a degenerating configuration on the axis. Furthermore, the re-attachment level at which the falling fluid should connect with a more stagnant zone could not be stabilized and showed a very chaotic behaviour. Nevertheless, the contraction algorithm showed to perform quite well.
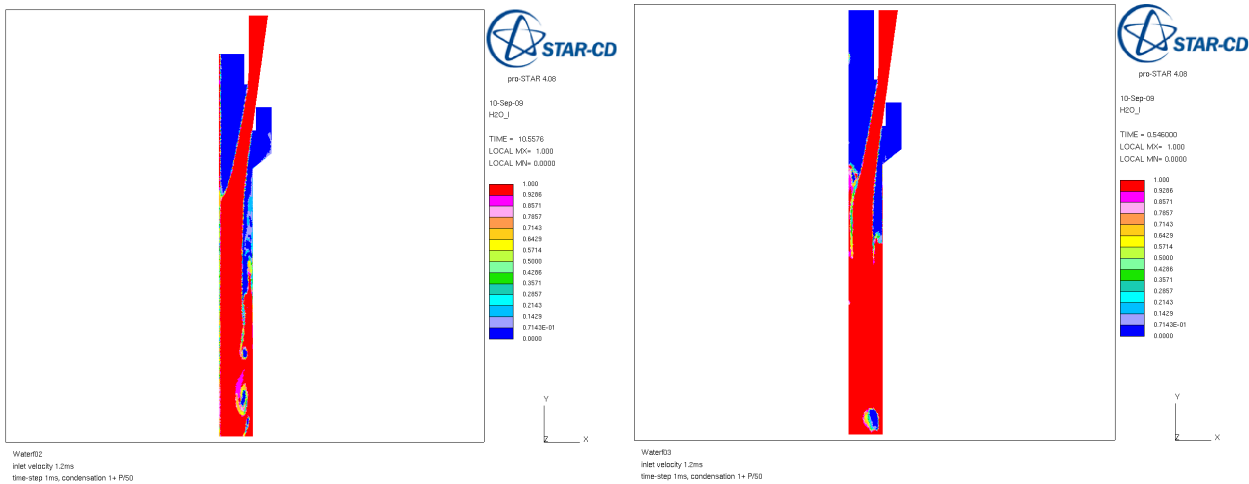
Figure 1: 2D axial-symmetrical slice of a XT-ADS spallation target. The simulation gave problems both at the symmetry axis and at the lower re-attachment level that could not be stabilized.

The second test case is a relatively simple 2D+ free fall simulation. It is illustrated in Figure 2. Water is falling from a downward converging nozzle at about 2m/s and is let fall down freely for about 1m. At first, we have tried to let the falling fluid impact the surface of a relatively stagnant fluid region. Control was performed by adjusting the outlet pressure and also by adjusting some porous resistance parameter in the near outlet region. It has been found out that there seems to be no intermediary stable free surface level. The free surface either rising up to over the nozzle exit, or being transported back out of the computational domain. Both behaviour where obtained by just a slight variation in setting the outlet pressure.

The first objective of this simulation was to stabilize a reasonable stationary solution with a reasonably fixed height of the bottom free surface level. Therefore a relatively large time step was used, together with a strong condensation. The condensation effect performed seemingly well, keeping the interface locally sharp while the flow was perturbed by strong splashes. With no stagnation level, that is, when the falling jet freely exits through the bottom pressure boundary, we could test the usefulness of the condensation. It resulted that when the time step was reduced so that the CFL become order unity or less, the condensation was useless (but not disturbing). By the way, the co-flow was organized in such a way that no strong shear flow would appear close to the falling free surface. The Eulerian sharpening convection algorithm performs therefore perfectly in absence of strong shear, for a controlled CFL below unity.
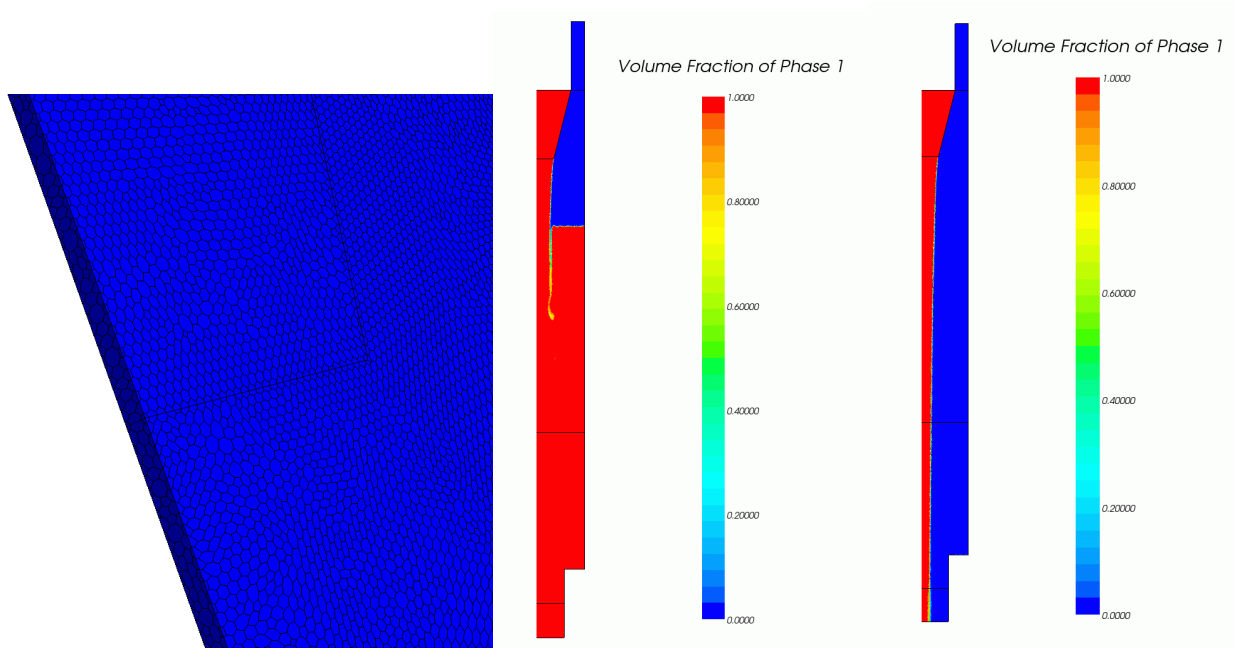
Figure 2: Free-fall test case. Left, a mesh detail. Centre and right: water volume fraction. Centre: free fall connects to a stagnation region. Right: free fall exits through the bottom pressure boundary.

The next series of test cases has been thought with the objective of having a relatively stable free surface subject to a relatively strong shear, associated to the free-surface numerical smearing. Dimensions and operating conditions should be reasonably in line with the foreseen XT-ADS, EFIT and FASTEF targets.

The geometry is therefore enclosed in the external warping of three vertical hexagonal tubes about 15% larger than for XT-ADS to obtain more easily some result at some significant flow rate. The driving idea was to test a channel-like target, even if the geometry does not seems a priory very well suited. In fact, a channel-like target means 3D calculations from scratch (half-domain) with small time steps. That is a foreseeable failure after a long and painful "bath of blood". By the way, hope was to get at least some clue for the main objective (an operable free-surface target). It should also be stressed that related (but without resolved free surface) 2D and 3D simulations had been already performed in the PDS-XADS framework.

Here is the basis idea. The flow rises from one lobe and is distributed to the central spallation region by a differentiated flow filtering grid. The grid is differentiated so has to be less resistive towards the top, therefore promoting a faster (close to) horizontal flow near the free surface than in the bulk, still delivering there a consistent flow. After flowing horizontally in the central region, he flow separated to fall down in the two other lobes.
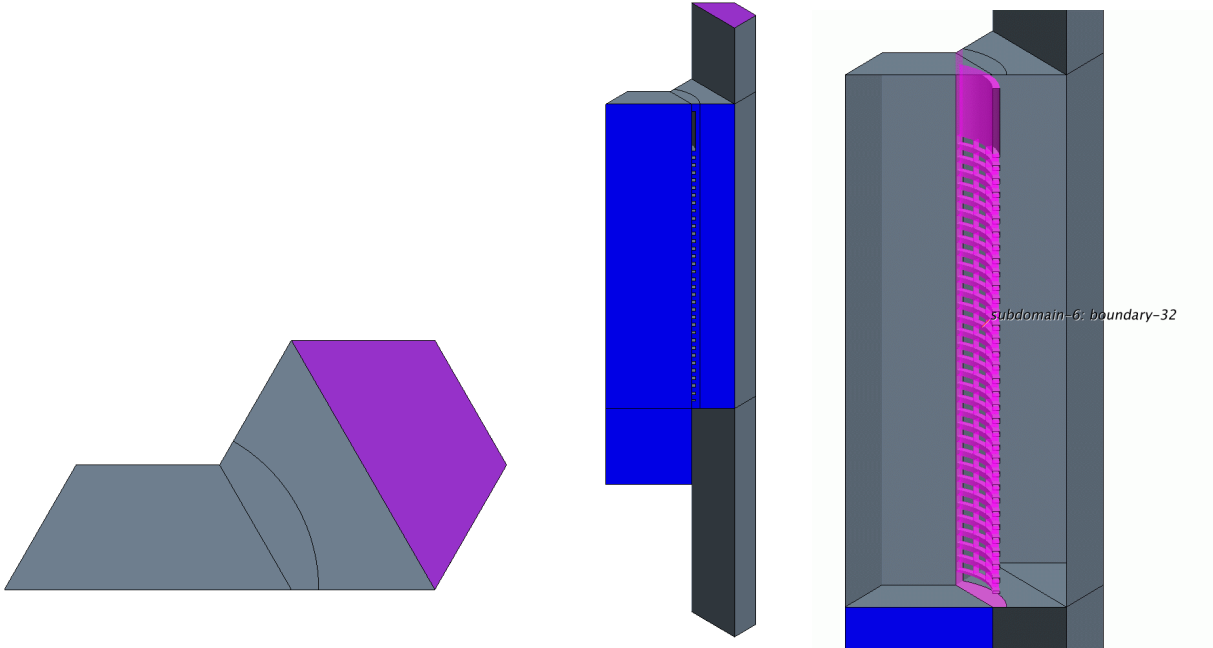
5

Figure 3: CRS4 target. Left and centre: simulation domain. Right: structure of the first internal flow diffuser.
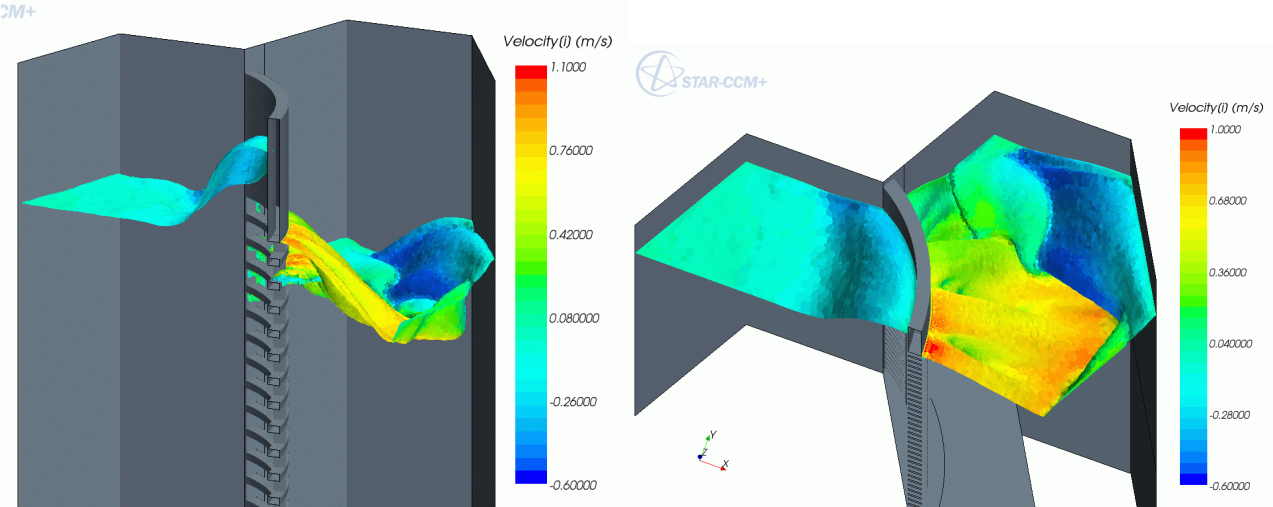

Figure 4: CRS4 target with rounded diffuser. Iso-surface (c=0.5) of volume concentration coloured by velocity magnitude.

The first trials allowed to adjust the grid differentiation to approach the desired flow configuration. It has also been seen that it is unlikely to distribute the flow horizontally in an arc of circle from a bended diffuser because the forced flow divergence was simply resolved by unwanted light phase inclusions and an heavy phase flow separation in the expected free surface proximity. This is illustrated in Figure 4. The flow diffuser would have therefore to operate only in intensity and not on the planar direction. For constructive simplicity and in absence of opposing motivation, we have restrained to a diffuser made of vertical series of small horizontal barrels with variable rectangular section and pitch. Due to erosion concern, cylindrical or at least smoothed barrels would be preferable, but require numerically a too high definition at this stage of the study.

6

A test case taking into account these considerations has been run and gave good preliminary results. As shown in Figure 5, a relatively stable and sharp free surface has been obtained, both sides of the diffuser. This was enough to test several variants of the sharpening algorithm, and also to test successfully the coupling with an ultra-simplified energy release. Warned by the former free-fall simulation, we have checked the usefulness of the sharpening algorithm. The flow after 1.5s of the algorithm switched off is shown on Figure 6, showing a large smearing of the interface. On Figure 7, we show the result of a few seconds of simulation with a tentative alternative sharpening algorithm. However, turning back to the original algorithm, an oscillation of the surface and of the outlet flow rate has been observed, with a typical frequency about 1.1 Hz. The maximum temperature was particularly sensitive to this oscillation even if the flow rate oscillation was order only 4%.
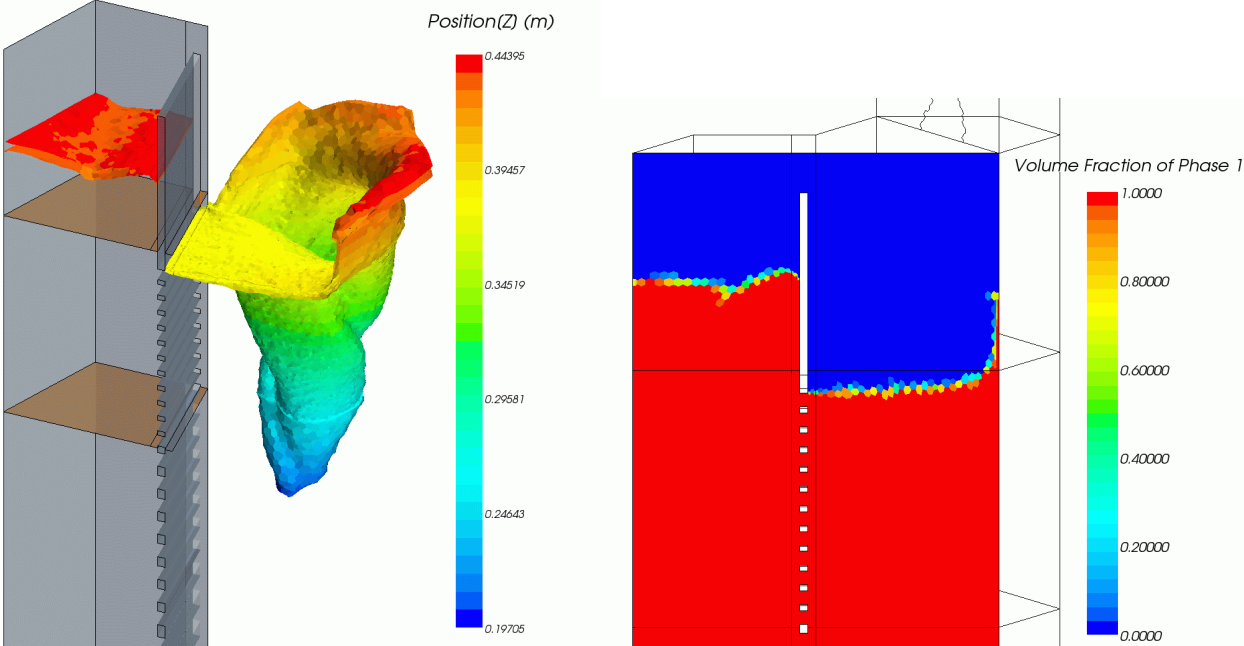

Figure 5: CRS4 target with straight diffuser. Left: two iso-surfaces of the volumetric fraction (0.1 and 0.9) coloured by height. Right: water volume fraction on the symmetry plane. Reference sharpening.

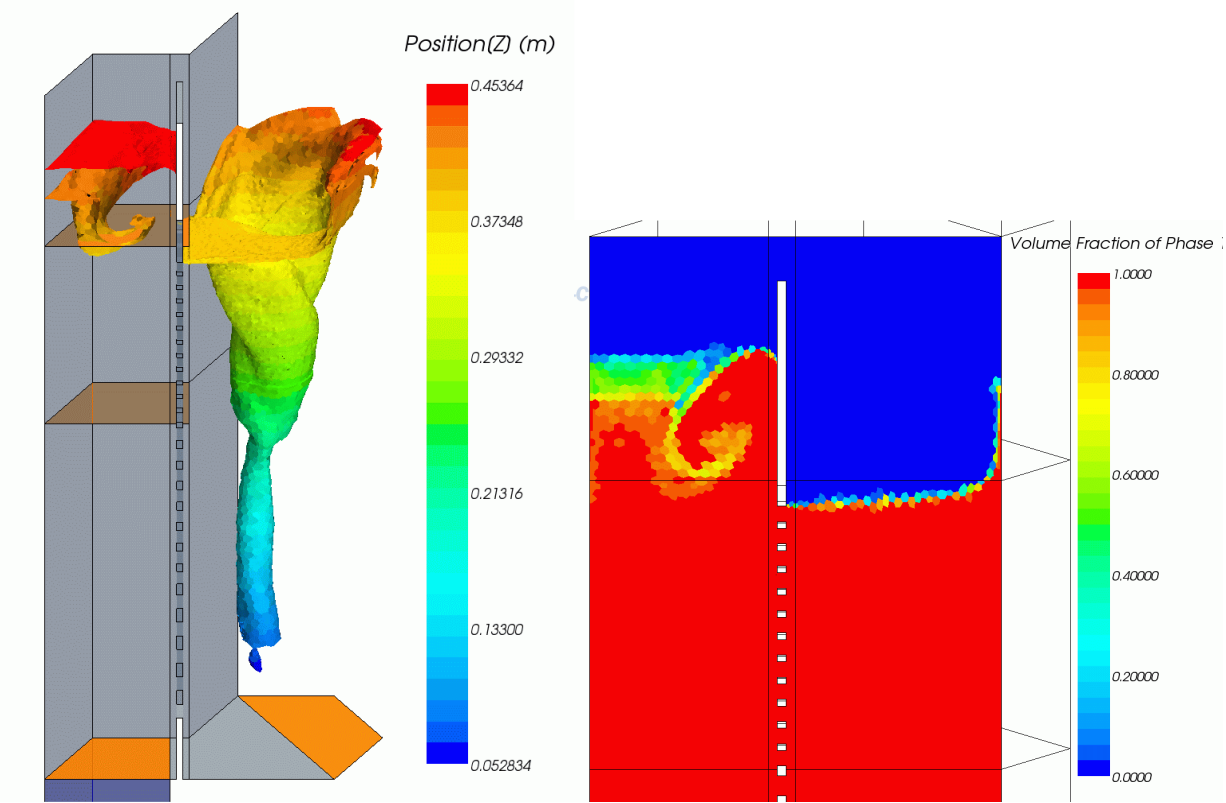Figure 6: CRS4 target. Flow after 1.5s of simulation without sharpening algorithm. Left: volume fraction iso-surfaces 0.1 and 0.9 coloured by height. Right: related water volume fraction on the symmetry plane.
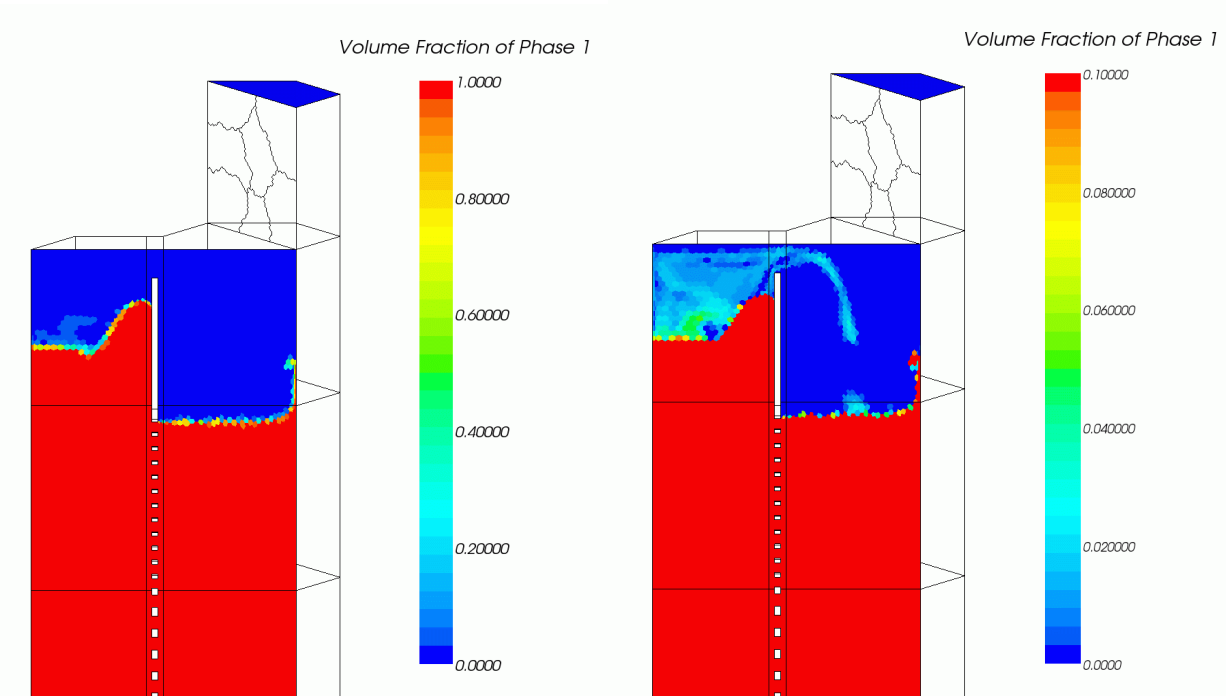

Figure 7: CRS4 target. Few seconds trial of a tentative alternative sharpening algorithm. Water volume fraction on the symmetry axis. Right: the colour scale is concentrated on the low values (between 0 and 0.1).

There was the feeling that this oscillation was due to excessively forced boundary conditions. These are a fixed flow inlet and a pressure outlet. A second top pressure outlet allowed for variations of the light phase volume and for the stabilization of a free surface.

Closing the loop at the bottom seemed more realistic. That is what we have done. The outlet has been connected to the inlet, a bottom region has been dedicated to the resetting of the temperature and another one has been completed by a distributed momentum source simulating a generic pumping device. It should be noted that former trials to operate such closed Eulerian multi-phase loops systematically ended in global failure due to the slight but always increasing mixing of the two phases with an always increasing fraction of light phase entrained into the loop. Thanks to the sharpening convection scheme aided by the condensation feature, we can attempt a new trial with some chance of success.

A closed loop means that we have to wisely choose an initial free-surface level because the overall heavy phase volume is expected to be conserved. We also have to evaluate a priori the global hydraulic resistance of the loop to dimension the pumping power intensity, so that the required flow rate is obtained. However, this can be done rather dynamically from an initial conservative guess.

At the geometric modelling stage, one tends to avoid useless volumes that will have to be meshed and will require additional computing power but will not any useful or better information. Unfortunately, usefulness of some volumes may appear too late… In our case, the plain separation in the upper part allowing the rising flow to buffer has been made too short. The available height for a stagnant flow over the riser is too short and a secondary flow path appears before we reach the desired flow rate (in this case 6.5 l/s for the half-domain). While the flow is limited in this case slightly below 6 l/s, some useful information has been obtained. First, the free surface can and does easily recover from the occasional additional jet flow from above the separator. Second, the flow is slightly pulsed in presence of this overflow (with a period about 0.9s),and further investigation demonstrate that the pulsation continue undisturbed when the overflow has been removed (by lowering the flow rate and the heavy fluid total volume). So, the pulsation seems to be inherent of the geometry under these operating condition. However, while the surface level over the buffer region varies quite a lot, the flow rate in the downward sections varies only by about 3%. Furthermore, the flow is quite regular in absence of overflow close to the exit of the diffuser where the spallation beam is foreseen to impact. Only the surface slopes changes slightly. Third, the condensation is able to keep the heavy fluid integrity in the heavily stressed part of the flow where it swirls down into the two down-coming lobes. No light phase is entrained in the loop in this strongly sheared flow. Forth, the separation of the horizontal flow creates two heavy fluid rolls, one at each side, which do not rebound in the central spallation region but reverse towards the centre of the respective hexagonal lobes. The flow is cut by the re-entrance of the geometry in a similar way one can see at the front of most ship vessels. This feature could be eventually enhanced in upgraded versions of the design. Fifth, we have established an "existence theorem" for the feasibility of such a loop. Coupling with the beam line and operational transients are likely to be studied with meaningful results.

6 Control feature

For marginally stressed flows, the sharpening convection algorithm performs well its duty as long as the control on the CFL (until about 1 on the interface, larger elsewhere) is maintained. So, a strict control on the time step, depending on the local velocity and mesh size must be performed.

This fundamentally precludes the possibility to run steady-state simulations (in fact, they are basically not allowed).

For open problems, a light fluid sink of relatively arbitrary intensity as long as it does not create excessively strong winds, is generally able to maintain a sharp interface and moreover, it is able to re-contract a formerly diffused (may be it by numeric issues or correctly due to a localized strong turbulence) interface in time.

In general, it would be preferable to dissociate the sharpening issue from the cavitation issue (volume source or sink). Therefore, sharpening would preferably be performed independently and without large scale volume generation. However, we still have to swap light volumes from one side of the interface to the other side. Using source and sink terms for an incompressible fluid, this can be obtained with a source which is the divergence of some flux turning to zero out of the smeared interface. The flux should preferably be oriented in direction of the concentration gradient. The simplest such flux is in fact the concentration gradient, leading to the usual diffusion-like term with the sign changed. Such anti-diffusive term has a strong drawback because its overall strength does not scale nicely with the smeared interface width. It becomes arbitrarily strong for sharp interface where there is no need for it and weakens fast for larger interfaces where it is strongly needed. Therefore, we have to derive an anti-diffusive flux with a better scaling property.

Such term has already been derived in the literature. A most known one serves to define the famous Cahn-Hilliard [3] equation. It consists in a Laplacian like anti-diffusive term compensated by a diffusive bi-Laplacian term. The result is an attractive interface size, which in our case could be set in relation to the mesh size. The defect of this term is in the use of the forth order bi-Laplacian, which requires in practical to enlarge a lot the stencil necessary to define the operator and is not very consistent with the necessity of a really sharp interface, no more than 2 cells wide.

To cope with the necessity to avoid to high order differential operators, the Allen-Cahn [4] equation has been derived. The controlling term is now a sharpening scalar function controlled by a diffusive Laplacian term. It should be noted than the Laplacian term is used for opposite effects compared with the Canh-Hilliard equation. The first defect of the Allen-Cahn equation is that the scalar term is not conservative. It is a somewhat small defect when applied to the light phase. Moreover, there is no mass production in the 1D case when the shape is the one of the attractive solution. The second defect is more subtle. It lies in the scaling with the width of the scalar term integral effect, which is directly proportional. This means that the swap of volume can be performed independently of the distance involved. On the contrary, and this is a good feature, the integral effect of the scalar term goes to zero for sharp interfaces. In the framework of the multi-fluid approach, the Laplacian term corresponding to the usual diffusion term can still be set to zero.

The Cahn-Hilliard and Allen-Cahn equations have been abundantly tested ad verified. However, this has been done in multi-phase flows in which there is no large density ratio. In fact, the equations seems to have been derived only for constant density. Care should be taken to generalize to high density ratio cases.

7   Light phase sink term

For the simplest case of a light phase sink term, the sink S has the following form:

$S = -\alpha \rho_l cd$

where ρ_l, is the light phase density, c is the light phase volume fraction and d the heavy phase volume fraction. The issue is to give a reasonable value to α. This is done introducing the

minimum interface width L_0 which is about the mesh size and the actual interface width L which can be approximated by the formula: $L = \dfrac{cd}{|\nabla c|}$. In this way, we can rewrite the source term as a

transport one: $S = -V_s \dfrac{L}{L_0} |\nabla c|$, where $V_s = \alpha L_0$ is the induced external light fluid velocity when the interface width is minimum. This velocity must be low enough to not disturb the heavy phase flow and large enough to have an effect. Taking this velocity to be one order less than the typical heavy fluid velocity, that is V_s=0.1m/s and a typical mesh size about 2 mm, we get α =50.

8 Allen-Cahn type source term

The Allen-Cahn usual source term is in the form S=0.5α f^2(c)' where the prime stands for the derivative in c and typically f=cd. We shall show that this term is in fact a first order approximation of a divergence like term correct in the 1D case (no curvature) for the standard shape. The general term is: $S_G = -\nabla.\alpha L^2 \nabla c$ and for a given L acts exactly has an anti-diffusive Laplacian like term with a local strength independent of the width. If the shape is given as before, L does not depend on space (but may depend on time) and we can take one of its power out of the divergence:
$S_G = -L\nabla.\alpha L \nabla c = -\alpha L\nabla.cdn = -\alpha L(cd)'\nabla c.n - \alpha Lcd\nabla.n = -\alpha L(cd)'|\nabla c| = S$
Because in the 1D case the normal n has zero divergence.
This means that the Allen-Cahn source term is conservative of the global volume (care must be taken in case of variable density) up to curvature effects and deviations from the prescribed profile. However, as we act only on the pseudo-vacuum phase, it can be considered as a better suited term as the previous sink term because is remains much more local in nature and in the 1D case, changes the velocity profile only inside the interface. Here again, an estimate of the velocity pair width is given [dV]=2α L_0 L/L_0.
Also, there is no much meaning to consider the complete divergence term, quite more tedious to compute (and referring to a larger stencil), until you have to consider, in a consistent way, the surface tension curvature effects, still that the vacuum phase is connected.
On another hand, if we want to capture an interface within one or two cells, the discretized form of the source may be quite distant from its mathematical meaning because its integration on the interface width may vary consistently from its theoretical null value. Divergence like terms should not suffer from this drawback, because even if the background field is loosely approximated, its divergence should normally (and automatically) be integrated correctly.

The introduction of the width L is based on an hypothesis on the interface shape. Before, we have implicitly considered a shape with an exponential decay to the asymptotic values. The decay is slower as for the classical diffusion behaviour. In fact, the shape is obtained by equilibrating the source term with diffusion in a 1D framework. For example, if we take f=(cd)$^{0.5}$, then the interface will be shaped as an Arcsinus-like function, and will have a finite extension. This choice is not convenient because f' is not zero for c=0 and c=1. However, intermediary powers between 0.5 and 1 would also give a finite extension shape but with a derivative of f^2 going to zero for c=0 and c=1. In our framework, there is no Laplacian-like diffusion, but only eventually numerical diffusion. So, we should not feel constraint by an a priori interface shape.

9 Interpolated sources

The sink term given before is derived from a tentative equilibrium between a convection flux and the sink term. It implies a large scale flux which condensate on the interface. The Allen-Cahn

source consider an equilibrium with a diffusion term. However, its form does not allow to re-contract a small dispersive tail. In effect, in absence of a nearby "close to unit" region, a low c-value region with only be dispersed by the source term.

It is may be possible, by interpolation of both source terms to strongly mitigate both effects, with a more pleasant result than for only one of the terms. That is, we can hope a global similar effect with a quite lower amplitude of the combined sources.

Blending of the two terms can be done on the following consideration. Let us reinforce the sink term at most up to the limit of an everywhere negative source. Up to a multiplicative constant a, this leads to a constraint on the interpolated source $S_3=a(S_1+S_2)=a(-f+f'f)=af(f'-1) < 0$ on c in $]0,1[$. With f=-cd, this leads to $S_3=-2ac^2d$.

10 Derived Source terms

The volume fraction source induces a mass source. Because all the other equations are written in conservative form, the mass source/sink should arise together with the corresponding momentum enthalpy and turbulence terms. In particular, the non-withdrawn enthalpy that should be associated with a vacuum sink term creates an undue overshoot of the vacuum temperature.

In starccm, the volume fraction equations do not transport the density. The source term has therefore the dimension of a frequency. There is only one energy equation. The energy source associated with the volume fraction source can be obtained by multiplying the later by the product of density, specific heat and temperature (there seems to be no reference temperature to evaluate the Enthalpy, i.e. the reference temperature is 0 K). The same procedure can be applied for momentum and the turbulence variables.

11 Cavitation

When vacuum is entrapped in the heavy fluid, considering it as an incompressible fluid leads to large errors. Basically, vacuum should not oppose resistance to the heavy fluid while the incompressibility implies at least a pressure reaction.

The cavitation model consider that the heavy fluid is filled with a high number low volume concentration of seeds for volume expansion of the heavy fluid vapour. In case of negative pressure (that is, below the saturation pressure, which for Lead and LBE is almost zero, i.e. a fraction of Pascal), the seeds expand to create bubbles and their expansion is limited by the overpressure created by the displacement of the fluid nearby. In absence of better theory, collapse of the bubbles is perform in a symmetrical way. To do this, one has to define a scalar quantity related to mean diameter or volume of the bubbles and derive an evolution equation for this scalar, increasing when pressure is negative and decreasing when it is positive. The hypothesis implied in this approach is that there is a large number of seeds in each computational cells, or considering ensemble averaged, that the number and the side of the seeds vary smoothly in space. This is not very consistent with the objective of having a sharp interface. However, the cavitation model allows the VOF+ model to perform quite well. So, understanding how it works independently of the underlying physics should help derive a functional "vacuum" model.

The problem of vacuum or near vacuum is its extreme reactivity to pressure changes. As said before, we are not, at least in a first time, interested in the "vacuum flow". We just need to avoid that the treatment of the vacuum induces an error on the heavy phase behaviour. So, let us see what has been practically done in the cavitation model. In fact, the bubble size related scalar represent

some kind of reactivity to pressure changes. When the bubbles are extremely small, this reactivity is rather low, the bigger the higher. What is interesting is that this reactivity evolves in time. It increases if the under-pressure is maintained in time so that if the initial reactivity is not enough to get back a neutral pressure, it increases so that this effect will be eventually obtained. In other words, the Cavitation model introduces an adaptive reactivity to pressure changes. Like all adaptive mechanism, it is sensitive to the application and must be calibrated to the specific case. In the cavitation model, calibration is performed indirectly through both the density number of seeds and their initial side. In the Heavy Liquid Metal (HLM) / vacuum context, some more direct criterion should be derived.

Note that while the size of the bubbles in the cavitation model can become arbitrarily large, its effect on pressure reactivity is bounded. So the model exhibits a lower finite positive non zero reactivity bound for seed size and a finite positive higher reactivity bound for infinite bubble size. The model in this aspect is riser wise because it assures a range of reactivity in which both extremes of the range are controlled by the seeds number and initial size. We would like to follow the same strategy, based on more direct control parameters.

12 Surface Tension

Surface tension has in general an apparent small impact on the applications foreseen. This appearance is somewhat misleading. Consider two fluids in a box under gravity. Without surface tension, the solution is a perfectly mixed still mixture, eventually presenting a small mixture fraction gradient if gravitational diffusion is taken into account. In presence of an even very small surface tension between the two fluids, the solution is a perfectly stratified still flow, with the lighter fluid over the heavier one. Surface tension is the residual effect of immiscibility for immiscible fluids already separated on the macroscopic scale. Forcing mixing to an always smaller scale would cause an always increasing reaction force to operate.

For a resolved laminar flow, immiscibility is a priori taken into account to maintain the interface sharpness. This is mainly preserved by the numerical convection scheme but is in a certain sense similar to the divergence-free correction (or projection) for incompressible fluids. For an unresolved turbulent flow, it is however acceptable to have a much more diffuse interface of the mean volume fraction variables. In effect, the real interface is supposed to be deformed by turbulence on unresolved time scales and length scales which could be a priori order the turbulent time and width. Particularly, the turbulent width may be much larger than the mesh size. However, for highly transient configurations simulated with a very small time step, the probabilistic interpretation is more difficult to defend on more than a one or two cells width basis.

13 Conclusion

We have analysed several methods allowing to improve when necessary the sharpness of the interface. We have given some guidelines on what method to use in what case. Effects are then tested and documented in relevant CFD simulations.

14 References

[1] STAR-CD, Version 4.00 Methodology, CD-Adapco, 2006.

[2] Drag Limitation in the in the Three-Feeder Design, Katrien Van Tichelen SCK-CEN, 01-06-2007

[3] www.ctcms.nist.gov/~wcraig/variational/node9.html (Cahn-Hilliard equation)

[4] www.ctcms.nist.gov/~wcraig/variational/node10.html (Allen-Cahn equation)