

MODELING THE LOST FOAM PROCESS WITH DEFECT PREDICTION

Flow Science, Inc.
www.flow3d.com
C.W. Hirt and M.R. Barkhudarov
April 1997

Overview

The Lost Foam or Expendable Pattern Casting (EPC) method is an increasingly popular technique for the production of cast parts with fewer inserts and having thinner sections or more detail than is possible with ordinary sand molds. The method uses a polystyrene foam material as an initial pattern for making a sand mold. Metal is then poured into the mold displacing the foam to form the part. Thin sections or relatively small details can be molded in this way because the foam adds support to the sand before it is replaced by metal.

Another advantage offered by a foam core is that it retards the motion of metal during filling. This reduces splashing and air entrapment, which improves the quality of the part. A further advantage is that most of the sand in the mold can be reused because it does not require a binder to stay in place.

In short, the Lost Foam process has many positive features that make it attractive for the production of many types of high volume parts. As with most things in life, however, there is always a price to pay. In this case there are a number of issues that make the Lost Foam process difficult to apply with confidence. Among these are the increased possibility for inclusions by trapped foam residue, increased potential for cold shuts because of the heat loss to the foam, and filling patterns that defy expectations and make gating design more difficult.

In Part I we describe a new Lost Foam modeling capability that has been added to the FLOW-3D program. Although the model is new, it is now at the stage where it can be used to investigate filling patterns and temperature distributions. Thus, the program can be used to design gating, locate possible cold shuts, and identify thin mold sections that may not be able to retain their integrity.

Metal is cooled by the action of removing foam. This means there is a possibility of early solidification or of laps and folds that do not knit completely because of high solid content. These issues are discussed in Part II.

The new model also contains an estimate of the distribution of possible inclusions caused by oxides or foam residue. This feature, described in Part III, not only captures the initial entrapment of impurities by folds or the meeting of different metal fronts, but also predicts the redistribution of those inclusions by subsequent flow processes.

PART I: The Lost Foam Model

Foam is considered to be a special kind of obstacle that can prevent the flow of metal unless it is heated sufficiently to lose its strength. In this sense the movement of metal into foam is controlled more by heat transfer mechanisms than by pressure or inertia of the metal.

Melting and/or vaporization of foam removes heat from the metal. This, in turn, can slow down the rate of heat transfer, although this effect is usually quite small because once the metal reaches its liquidus temperature, it has a large reserve of latent heat that must be removed before the temperature can be lowered further.

When metal enters a control volume containing foam, the heat transferred to the foam from the metal during one increment in time is computed. This amount is then used to compute the amount of foam that is decomposed and raised to the temperature of the metal. Of course, the energy transferred to the foam is removed from the metal.

Removal of the foam leaves room for the metal to advance, which it does through the use of a modeling trick that is equivalent to a local sink in the control volume. The "trick" is better than a sink, however, because it allows the metal to flow smoothly through the control volume into neighboring volume elements.

Should the metal temperature fall below the melt temperature of the foam, it cannot penetrate the foam. In most cases, however, the metal would be solid at temperatures well above the foam melt temperature. Metal is also prevented from moving into foam when its pressure is less than the ambient air pressure (PVOID in the model).

No attempt has been made in our model to directly model the flow of gas produced from vaporized foam. The importance of gas is well known in terms of the effect that different coating permeabilities have on the performance of a casting. With insufficient permeability, gas buildup may actually blow metal back out the filling sprue.

Any model that actually attempts to model the gas will of necessity be complex, requiring more user-supplied data and consuming more computer resources. The gas exists in a narrow gap between the metal and foam. Flow in this region must be relatively unstable with lots of bubbling and jetting of metal across the gap. The rapid flow of gas along the irregular gap would have to be modeled as a separate flow field, including flow losses through the porous coating at side walls. The gas must exhibit some compressibility or at least undergo significant thermal expansion.

Finally, any model of the gas is incomplete without a model of liquified foam and a phase transition model converting the liquid foam to gas. Similarly, some gas passing into the sand may be condensed only to vaporize later as the sand heats up. In view of these complexities, the construction of a model for the dynamic behavior of gas would be very difficult, and it would be expensive to use.

If a direct model of the gas is not feasible, then what can be done? It seems reasonable to want to maintain the smallest possible gas region, otherwise there could be excessive metal splashing and trapping of surface impurities--features that the Lost Foam process is supposed to eliminate. Another reason to minimize the gas region is to prevent the sand from collapsing after the foam has been removed. In other words, coatings should be as porous as possible without sacrificing their strength to keep sand in place. Under this assumption the role of gas is limited.

We propose to include the effects of coating permeability through a change in the metal-to-foam heat-transfer coefficient. The reasoning behind this proposal is that the rate of heat transfer from metal to foam should be proportional to the thickness of the gas region. A low permeability coating would allow gas buildup and reduce heat transfer.

The precise manner and amount of the change in the heat transfer coefficient, and more importantly, the degree to which this idea works, must be determined from laboratory experiments.

Metal-to-Foam Heat Transfer

The foam model involves known or easily estimated physical properties with one exception. That exception is the metal-to-foam heat-transfer rate. This parameter must be determined from experimental data. Simple tests involving the time for metal to penetrate a vertical column of foam should be sufficient for this purpose.

An analysis of the speed at which metal advances into foam in our model indicates a velocity approximately equal to,

$$V_{adv} = \text{HOBS1}(n)/\text{RCOBS}(n),$$

that is, the speed is approximated by the ratio between the heat-transfer coefficient and the product of foam density times specific heat. If a filling time is known for some test case, then HOBS1 can be estimated from this formula.

There is some evidence that the heat-transfer rate may also have a dependence on pressure and temperature at the interface. For this reason the model has a separate subroutine, HFOBCL, for the computation of metal-to-foam heat exchange. If no other information is given, the program uses the constant value of HOBS1(n) supplied by the user. Users can also customize the HFOBCL routine to conform to their own experimental observations.

Whether or not heat exchange can be correlated with pressures, temperatures, and coating permeabilities remains to be seen. Nevertheless, this approach offers an operational procedure to begin modeling. The test examples discussed later will show that the model certainly has considerable promise.

Using the Lost Foam Model

The following properties of foam are required in our model:

IFOB(M) = 1, A flag identifying Obstacle M as foam

HOBS1(M) = Foam/Metal heat-transfer coefficient

RCOBS(M) = Density (macroscopic) times specific heat of the foam

REMFOB(M) = Density times energy to "melt" foam

REVFOB(M) = Density times energy to "vaporize" foam

TMFOB(M) = Foam "melting" temperature

TVFOB(M) = Foam "vaporization" temperature

A foam core can be made up of different densities or types of foam simply by initializing them as different obstacles with their own properties. The shape of the foam pieces is defined in the same way that any other obstacle (geometric object) would be defined using the primitive-based solids modeler in FLOW-3D or by importing one or more CAD files describing the geometry.

Several points of clarification should be made regarding this data. The density of the foam is its macroscopic value not microscopic value because we multiply foam quantities by control volumes to compute such things as total heat capacity.

The energy REMFOB is the heat of transformation approximating the initial degradation of foam into a more or less liquid state, while REVFOB is the heat of transformation for vaporization of the foam material. These energies are usually small in comparison to the thermal energy needed to raise the foam from room temperature to the temperature of the metal, $RCOBS \cdot (T_{metal} - T_o)$. For this reason their precise values are usually not critical for obtaining good model results.

Typical values are REMFOB=0.8E+5 gm/cm/s² and REVFOB=1.6E+5 gm/cm/s² with corresponding temperatures of approximately TMFOB=100 C and TVFOB=747 C.

Foam density is on the order of 0.02 gm/cc but can vary even within a single foam piece by ±10%. Density variations arise because styrofoam beads are blown into a die and then subjected to steam to expand and weld the beads together to form the core. During this process the

beads can pack more or less in corners and behind obstructions in a manner that is difficult to predict much less control. It has recently been shown that foam pattern densities are also higher at surfaces than in central regions [Liu, J., Askeland, R. and Ramsay, C. W., "Effect of Foam Density and Density Gradients on Metal Fill in the Lost Foam Casting Process," AFS Preprint 97-139 (1997)]. This may be caused, for example, by cooling at die surfaces before the foam has fully set.

The specific heat in RCOBS should, strictly speaking, be made a temperature-dependent quantity, but in the present model it is simply a constant. At the suggestion of Prof. S. Shivkumar, we use the specific heat of water as a first approximation, so a first guess value is $RCOBS=8.0E+5 \text{ gm/cm}^3/\text{K}$.

Casting A Horizontal Plate

A simple test example for mold filling has been reported by Yao and Shivkumar (AFS report #95-116). In their tests a 15 by 20 cm plate, 1.3 cm thick, was cast using several different materials. Here we consider Aluminum poured at 1023 K with a 28 cm high sprue. The properties used for the Aluminum in CGS units, with temperatures in degrees K, were:

Density = 2.7
Specific Heat = 8.37E+6
Conductivity = 1.88E+7
Thermal Expansion = 3.0E-5
Liquidus Temperature = 933.0
Solidus Temperature = 933.0
Latent Heat of Fusion = 3.94E+9
Pour (gate) Temperature = 988.2

For the foam properties we used:

Density = 0.02
Metal-to-Foam Heat Transfer Coefficient = 6.0E+6
Density * Specific Heat = 8.0E+5
Density * Melt Energy = 1.6E+5
Density * Vaporization Energy = 0.8E+5
Melt Temperature = 373.0
Vaporization Temperature = 750.0,

Observations showed the plate was filled in about 2.7 s from which we estimated HOBS1 to be about $6.0E+6 \text{ gm/s}^3/\text{K}$. Figure 1 contains the observed filling configuration at a series of times (0.6, 1.0, 1.8 and 2.5 s). Filling is completed at about 2.0 s. When compared to the computed configurations, we see remarkably good correspondence in metal front shapes.

The calculations show some small differences with data in the time to reach each contour. In fact, the calculations exhibit a nearly uniform filling, while experiments seem to indicate an unsteady filling as evidenced by the non-uniform spacing of the interface contours along the horizontal mid-line. Since no reproducibility information was given for the experiment, it is unknown if these variations are real or the result of experimental conditions, e.g., a non-uniform pouring of the metal.

Aside from their agreement with the test data the computed results show some interesting features associated with Lost Foam casting. For one thing, the metal front temperature reaches the liquidus temperature very quickly and then remains there for most of the filling time.

Another observation is the smooth filling pattern, which is very different from what would be observed if the foam was not there. The smooth filling pattern resembles a plastic injection molding process because the front seems to move primarily normal to itself. The high viscosity of plastic makes the material want to advance in the direction of the largest pressure gradient, which is normal to the front. Similarly, in Lost Foam casting, metal wants to advance in the normal direction as well because that is the direction of heat transfer (hence, of foam displacement).

The similarities with plastic flow are not complete, however, because in the Lost Foam case we see a significant flow of metal tangentially along the front, a flow that does not occur in very viscous plastic materials. One consequence of this secondary metal flow is that thermal energy is being well mixed in the metal. In fact, after the plate has completely filled there remains a substantial circulation that reduces the spread between the minimum and maximum temperatures observed in the metal (last frame in Fig.1). Residual momentum of this magnitude will modify the solidification history of the part.

Casting a Sample Part

To illustrate the Lost Foam model for more complicated parts, we have selected a non-proprietary test piece used by the AFS several years ago for demonstrating different software packages. The part was originally intended for gravity pouring and possessed both a sprue and a large conical riser. For our purposes we have reduced the height of the riser (see Fig. 2).

For simplicity we begin with the sprue filled with metal, which in this case is medium carbon steel. Foam occupies the remainder of the mold. We also replaced a specified flow condition at the top of the sprue with a specified pressure corresponding to one atmosphere. Foam properties were assumed to be the same as those used in the previous Horizontal Plate problem, although converted to a different set of units used for this example. The steel properties were taken from the AFS Transactions: 89-117, p.553. In CGS units the steel properties are:

Density = 7.0
Specific Heat = 6.28E+6
Conductivity = 4.19E+6
Thermal Expansion = 1.3E-3

Liquidus Temperature = 1783.3
Solidus Temperature = 1733.2
Latent Heat of Fusion = 2.72E+9
Pour Temperature = 1833.3

Several snapshots of the computed filling progression are shown in Figure 3. Clearly, the Lost Foam process produces a very clean and smooth filling of the entire part. There is no splashing or jetting of metal that could trap air or surface inclusions. The contrast between this case and one without foam is evident from the plot shown in Figure 4.

Including Lighteners

Sometimes it is necessary to modify the sprue and runner system to get metal to different gates at the same time. Oftentimes this is done by removing central portions of foam from the sprue and/or runners. These hollow foam sections are referred to as lighteners.

A simple test of a lightener is a vertical cylinder of foam, 25 cm high, with a cylindrical concentric hole. The inner and outer radii of the foam are 1.25 cm and 2.54 cm, respectively, and the length is 25 cm. For this test the metal enters the cylinder from the bottom with a hydrostatic pressure head of 28 cm of metal, which is sufficient to fill the central hole before the foam at the sides has been completely degraded. All physical parameters for this problem were the same as for the Horizontal Plate example.

Figure 5 shows the progression of the filling, first in the central hole then filling out to the sides as the foam is evaporated. After filling there remains a pronounced circulation in the mold. This is generated by the lightener because the higher velocity central flow is deflected downwards at the sides as the foam is removed. Without the lightener the circulation is much less.

The metal appears not to be penetrating the foam near the ingate at early times. This happens because of low pressures in the recirculating regions on either side of the incoming jet of metal. For a time the pressure in these regions is below atmospheric pressure which is insufficient to advance the front.

As this example shows, the Lost Foam model has the flexibility to model a wide variety of situations. In this case the geometry of the foam was altered to simulate a useful foundry practice. In the next section we look at the consequences of a variation in foam properties.

Variable Foam Density

Since foam is defined in FLOW-3D as a special type of obstacle or collection of obstacles, it can be defined with the same generality as other obstacles in the program. To demonstrate how a change in foam density could affect a casting, consider a simple, right angle junction between two straight channels. In the left side channel, however, let us place foam with a density 10% lower than that in the central vertical channel. Similarly, in the right side channel we place foam having a density 10% greater than that in the central channel.

Aluminum metal enters from the bottom and moves upward. As would be expected, the metal front moves more rapidly into the lower density left channel than it does into the higher density right channel (see Fig. 6). The difference in filling the two horizontal arms is accentuated, of course, with an increase in the length of the arms.

PART II: Mushy Zone Considerations

Molten metal must transfer heat to the foam in order to displace it. This causes the metal at the front to fall fairly quickly to its liquidus temperature. As more heat is removed, the solid fraction at the front begins to increase. Because the heat of transformation of most metals is relatively high compared to the heat required to degrade foam, the solid fraction increases slowly.

An increase in solid fraction has at least two important consequences. One occurs when two fronts come together, for instance, behind an obstruction in the flow. When the solid fraction is high, the surfaces can no longer knit together smoothly.

A second difficulty emphasized by Ken Williams [private communication, Flow Simulation Services, May 1997] is that the viscosity of the metal increases with increasing solid fraction. Experiments reveal that in material such as 319 Aluminum, the viscosity begins a significant increase in magnitude when the solid fraction exceeds about 12% (the point of coherency). Beyond a solid fraction of about 50% (the point of rigidity) the viscosity begins an extremely rapid increase, and the material behaves almost like a solid.

Work is underway to include a realistic solid-fraction dependent viscosity in the flow portion of the Lost Foam model. In the meantime it is still possible to assess the probability that a mold design will have difficulties and to experiment computationally with different remedies.

To illustrate how our model can be used to investigate possible solidification problems at the metal front, consider the situation that arises in a long, vertical sprue connecting several horizontal runners, Figure 7. The sprue is circular with a diameter of 7 cm, and each of the runners is a circular disk, 2 cm thick with a diameter of 24 cm. In one case there are three runners, while in the second there are only two, the middle one having been removed.

Calculational results using aluminum for the metal are shown in Figure 7 at the time that metal has just reached the entrance to the last runner. The shaded region in each plot is the region at the metal front where the solid fraction exceeds 12%. The peak solid fraction along the axis of the sprue is about 0.1 in the three runner case, but 0.2, or twice as much, when only two runners are present.

Clearly the presence or absence of the middle runner is having a significant effect on the metal front. In the two-runner case the solid fraction is already well beyond the point of coherency, which indicates that further filling may be difficult or that laps and folds are not likely to knit well enough to make a sound part.

Because the simulations give us a complete picture of the entire filling process, it is easy to see how the middle runner is affecting the flow. By the time metal reaches the entrance to the middle runner it has already developed a mushy zone. Metal is able to move into the runner as fast as it continues down the sprue because the metalostatic head is more than enough to overshadow what little momentum there is. Consequently, the runner draws off a good portion of the mushy material, which is then replaced by hot metal following along behind the front.

When the middle runner is absent, there is a continuing buildup of solid fraction at the metal front. Given a long enough sprue, it is obvious that solidification difficulties must eventually develop. However, when side channels or pockets are placed along the sprue to remove the solid accumulated at the front, it makes it possible to have long sprues and runners.

Not only is the mushy zone reduced by the presence of side pockets, but any contamination that may have been accumulated at the front will also be reduced by the same mechanism. This topic is discussed in more detail in Part III.

Here is an excellent example of how this model can be used in the design stage of a Lost Foam project to improve the likelihood of a successful casting with a minimum of expensive, time-consuming, trial-and-error adjustment.

PART III: Tracking Defects

Most defects observed in Lost Foam cast parts arise from surface contamination that has been entrained into the body of the part. Surface contamination may arise from several sources including loose sand, oxides and residue from degraded and vaporized polystyrene. Entrainment arises from the meeting of two fronts (folds or laps): material remaining on the surface of a part (wrinkles) and foam material surrounded and trapped by metal due to surface turbulence at the front (bubbles).

Once residue is trapped, it continues to move with the flowing metal and may be diluted and dispersed if the flow exhibits enough shear (possibly turbulence).

For surface contamination to be directly entrained into a body of metal, there must be sufficient kinetic energy in the flow to produce local overturning of the free surface. This concept, termed "surface turbulence" by Professor John Campbell of Birmingham University, has been shown by him to have significant consequences on the quality of cast aluminum parts [Campbell, J., "Invisible Macrodefects in Castings," *J. de Physique* IV, 861 (1993); Flow Science, Inc. report FSI-94-TN41].

A scalar tracking scheme is under development in a collaborative effort between Professor Campbell's group and Flow Science to convert his surface turbulence concept into a quantitative prediction for oxide defects.

With a similar goal in mind, we have coupled a scalar transport model with the Lost Foam casting model. Defects, or at least their likelihood, are represented by a scalar variable that

is initially zero everywhere. As foam is melted or vaporized in a control volume, the scalar quantity in that volume is incremented by an amount proportional to the mass of degraded foam. This scalar is then allowed to advect with the flow using a second order, monotonicity-preserving scheme for increased accuracy. The scalar may be interpreted as proportional to the mass of contamination per unit volume.

At the conclusion of a simulation, the distribution of the scalar indicates locations where residue is likely to be. As we shall see later, most of the scalar tracked residue is found at the last place to fill, as would be expected. In other words, residue accumulates at the front and is pushed along until there is nowhere else for it to go.

More interesting are the situations where two fronts meet, trapping surface residue in the interior of the metal or near corners where flow recirculation regions may develop. Even the circulation remaining in metal after a mold has filled can redistribute the surface-generated contaminants.

Without detailed experimental comparisons to correlate predictions from this type of model with actual defects, it is impossible to assign any significance to the absolute values of the scalar. However, qualitative results from even the simplest of test cases reveal a considerable amount of useful insight into the processes responsible for defect generation and their final distribution in a cast part. In the following sections we shall give a few examples of what can be learned.

Defect Prediction and Plastic Weld Seams

The book Principles of Polymer Processing by Tadmor and Gogos (John Wiley & Sons, New York, 1979) contains a simple, nearly two-dimensional, example of weld seam generation for a case of plastic injection molding. The mold consisted of a long narrow slab with a combination of rectangular and circular obstacles in its interior. Plastic enters the slab at one end through a small gate and flows around the obstacles as it fills the mold. In the wake of each obstacle there is an indication of where two surfaces have met and then have been drawn downstream with the subsequent flow.

A FLOW-3D simulation of this experiment was done as a code validation test. On a whim, marker particles were placed at the leading edge of the plastic entering the mold. Amazingly, at the end of filling the markers were found to be in excellent agreement with the experimentally observed weld seam defects (see Figs. 8a-8b).

A repeat of this computation using the scalar transport concept instead of marker particles shows that the scalar technique is as good or better in some respects than the marker particle method (see Figure 8c). The disadvantage of marker particles is that they can't be subdivided, while the scalar variable has no such restriction. This allows the scalar to give a more continuous estimate of the "probability" of where defects are likely to occur. The scalar method also requires less CPU effort than particles, and it is incorporated more easily into displays of other flow quantities.

Defect Prediction in a Lost Foam Casting

The Tadmor mold was converted (computationally) to a Lost Foam mold, and the plastic was replaced by aluminum, using the same properties for foam and aluminum as for the Yao and Shivkumar Horizontal Plate test described earlier. A pressure inlet condition corresponding to a 28 cm head of metal was defined at the gate. Two snapshots from this simulation are shown in Figure 9, one when the mold is about half full and the other when it is nearly filled.

The scalar variable used to track possible defects was given an initial arbitrary value of 1000.0 in the small notch representing the gate. As the foam is degraded, additional contributions to the scalar at the metal-foam interface are generated (with an arbitrary, multiplicative coefficient of $1.0E+3$).

During filling, the high concentrations are always observed at the interface, as would be expected. At the end of filling, this leads to a high concentration at the location that is last filled. What is of more interest, however, are all the other locations where scalar values are observed.

For example, we see significant values in the two corners nearest the ingate. These values come from material initially pushed into the corners early in the filling period. This material remains there during the remainder of the filling because the flow is almost stagnant in these regions.

The fact that possible defects are trapped and remain in the upstream corners because of stagnant flow is an important observation. The same thing is happening in the wakes of the obstacles. Behind each one there are some residual scalar values. These values are initially created because of the meeting of metal fronts coming from either side of the obstacles. However, it is also the flow stagnation and recirculation in these regions that keeps the scalar values from being mixed into the flow and advected downstream.

A closer look at the computed results also reveals that some amount of the scalar quantity gets left along the side walls. The initial deposition comes from the metal front, and wall shear helps to slow the flow velocity near walls so the scalar can't be diluted as readily from these regions.

The positive correlations observed between computational and experimental data in the case of plastic injection molding offers encouragement for similar defect prediction in the Lost Foam process. More importantly, observations such as these are not peculiar to this simple example. The advantages of a simulation tool like FLOW-3D is that it can quantify the creation and subsequent dilution and advection of possible defects throughout the body of cast parts of nearly any complexity.

Illustration of Defect Prediction in AFS Part

A demonstration of the use of the scalar defect method in connection with the more complicated geometry of the AFS part is given in Figure 10. As noted earlier, this part filled very smoothly, consequently nearly all of the scalar quantity is pushed into the last places to fill.

Defect probability is seen to be greatest in the upper part of the riser, but there are also significant residuals on the upper horizontal surfaces of the part located the furthest from the sprue.

Illustration of Defect Prediction in Lightener Example

The lightener example is somewhat unique because it has a central hole in the foam, which allows rapid filling of this region by metal. One consequence of the hole previously commented on was that it generated a large residual recirculation in the metal after the completion of filling.

When the scalar defect model is coupled with this example, we see that all scalar material is at first pushed to the outside walls of the cylinder because this is the direction in which the foam is being degraded. However, if this material does not leave the mold cavity through the coating (a process not modeled), then the residual recirculation will carry some of it into the central region of the cylinder even though in the beginning this region contained no foam (Figure 11).

Illustration of Defect Prediction in Cross Arm Example

Finally, we look at the distribution of scalar concentration in the case of the cross arm configuration, which had differing foam densities in the horizontal arms. Referring to Figure 12, we see that most of the scalar material is being pushed ahead by the metal fronts, but some residual remains in the junction region.

This residual is trapped in the recirculation zones at the upstream corners of the junction. At the downstream corners the mean flow effectively sweeps these surfaces clean. It is observations such as these that make simple problems worth doing. With simple geometries it is possible to sort out different flow processes and readily see their consequences. A good understanding of elementary situations provides the insight needed to better understand what can happen in real applications.

Summary Comments

A new Lost Foam modeling capability has been incorporated into the FLOW-3D simulation program. Realistic filling simulations can be performed that yield temperature distributions at the end of filling for solidification analysis by the program.

The new model has provisions for spatially non-uniform foam densities (or other properties), includes lighteners, early solidification, and the prediction of incomplete filling when the applied metalostatic head is insufficient to fill the mold.

A companion development has also been described for the prediction of defect distributions using a scalar concentration function. Defects in this model are assumed to arise at the metal front in proportion to the amount of foam consumed. Defect material is then tracked with the flow so that an accurate accounting can be made of where it is likely to end up. Although this model is only preliminary, and will surely undergo further development, its initial predictions have been quite revealing from a qualitative point of view.

The model contains only one unknown parameter whose value must be determined from experimental tests. This parameter, the heat-transfer coefficient between metal and foam, has been placed in a separate subroutine so that users can customize the code for their applications. For instance, depending on the type of foam, material poured, sprue height, or any number of other process variables, users may find it desirable to make the heat-transfer coefficient pressure and/or temperature dependent, or perhaps make it reflect the use of coatings having different permeabilities.

Clearly, the model proposed here is relatively simple and does not address some of the features that researchers have claimed are essential. For instance, the greatest limitation would appear to be the lack of a detailed treatment of gaseous products and their interaction with coatings. This point was discussed above where it was argued that it is better to try simpler models first before attempting what would surely be a difficult task.

Since the present model has some potential for empirically characterizing different coatings, it is premature to conclude that it does not go far enough. Furthermore, there are several extensions of the present model that could be easily realized if comparisons with experimental tests indicate this is warranted.

Possible Future Directions

If the scalar quantity represents residual (or liquid) foam at the metal front, it is not difficult to imagine the incorporation of this information into a more detailed interface model. For example, the heat-transfer coefficient could become dependent on the amount of liquid foam at the interface. Also, sub-models could be constructed to describe the slow vaporization of liquified foam as it continues to be heated and to describe its penetration through coatings into sand. Additionally, the presence of glue layers might be accounted for by initializing glue regions as sources of extra scalar mass that would require more heat to vaporize. In a similar way, glue could be made to change the heat-transfer coefficient.

The surface turbulence model still needs work. In ordinary casting the metal interface can exhibit localized overturning due to turbulence or localized flow structures that are intense enough to overpower gravity and surface tension forces, which tend to keep a surface smooth.

The prediction of how much overturning there might be must be formulated in terms of local variables that measure the competition between the various forces.

Likewise, in Lost Foam casting we need a similar local model of the likelihood of overturning and entrainment of foam material into the metal. In this case, the situation is complicated by the degradation of foam into liquid and gaseous products.

Most publications describing the Lost Foam metal-foam interface usually draw pictures in which they visualize a layer of liquid foam on the metal followed by a gas layer and then the intact foam [Yang, J., Huang, T. and Fu, J., "A Study of Gas Pressure in the Mold of the EPC Casting Process," AFS Preprint 97-136 (1997)]. Sometimes the layering is reversed, with the gas layer adjacent to the metal followed by a liquid layer adjacent to the foam [Lui, J., Askeland, R., Ramsay, "A Study of the Foam-Metal-Coating Interaction in the Lost Foam Casting Process," AFS Preprint 97-137 (1997)]. It's not hard to imagine that neither of these representations is very good, but rather that the interfacial region is highly turbulent, much like the chaotic behavior of liquid water on a hot skillet. Vaporization of liquid in contact with the metal would generate gas bubbles penetrating the liquid layer and inducing more liquid metal to come into contact. This type of picture needs to be quantified and used as a better source description for scalar defect predictions.

There is much yet to do, but what an exciting journey if these efforts lead to a more reliable use of Lost Foam casting with its many economical and environmental benefits.

Acknowledgement

The authors thank Dr. Ken Williams of Flow Simulation Services, Inc. for reviewing this document and making several useful suggestions to improve the presentation.




Fig. 1. Filling of flat plate. Data in first frame from Yao and Shivkumar. Times shown are 0.6, 1.0, 1.8, and 2.5s. Residual flow after filling shown in last frame.

Fig. 2. AFS test casting.






Fig. 3. Filling of AFS part is smooth.





Fig. 4. Comparison of filling patterns with foam (left) and without foam (right).




Fig. 5. Bottom filling of cylindrical column with central lightener channel (hole).




Fig. 6. Bottom filling across junction. Left arm foam density is 10% less and right arm density is 10% larger than foam in central channel.




Fig. 7. Top filling of sprue with side runners. Region at metal front with more than 12% solid fraction is blackened. Side runners reduce solid fraction at front.




Fig. 8. Plastic injection molding test. Marker particles locate weld seams (middle) as does scalar tracking (bottom).




Fig. 9. Lost foam test showing scalar defect tracking.




Fig. 10. Defect prediction for AFS part.




Fig. 12. Defect prediction for junction.




Fig. 11. Defect prediction in lightener example. Residual momentum carries defects from wall to central region.