

Elsevier Butterworth-Heinemann
Linacre House, Jordan Hill, Oxford OX2 8DP
200 Wheeler Road, Burlington, MA 01803

First published 2004

Copyright © 2004, John Campbell. All rights reserved

The right of John Campbell to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988

No part of this publication may be reproduced in any material form (including photocopying or storing in any medium by electronic means and whether or not transiently or incidentally to some other use of this publication) without the written permission of the copyright holder except in accordance with the provisions of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London, England W1T 4LP. Applications for the copyright holder's written permission to reproduce any part of this publication should be addressed to the publisher

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, e-mail: permissions@elsevier.co.uk. You may also complete your request on-line via the Elsevier homepage (<http://www.elsevier.com>), by selecting 'Customer Support' and then 'Obtaining Permissions'

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication Data

A catalogue record for this book is available from the Library of Congress

ISBN 07506 4791 4

For information on all Elsevier Butterworth-Heinemann publications visit our website at http://books.elsevier.com

Typeset by Newgen Imaging Systems (P) Ltd., Chennai, India
Printed and bound in Great Britain by Biddles Ltd., King's Lynn.

Contents

Dedication	ii		
Preface	vii		
Summary	xi		
Rule 1. <i>Achieve a good quality melt</i>	1		
1.1	Background	1	
1.2	Melting	3	
1.3	Holding	3	
1.4	Pouring	4	
1.5	Melt treatments	5	
1.5.1	Degassing	5	
1.5.2	Additions	6	
1.6	Filtration	7	
1.6.1	Packed beds	7	
1.6.2	Alternative varieties of filters	8	
1.6.3	Practical aspects	8	
Rule 2. <i>Avoid turbulent entrainment (the critical velocity requirement)</i>	9		
2.1	Maximum velocity requirement	10	
2.2	The 'no fall' requirement	13	
2.3	Filling system design	15	
2.3.1	Gravity pouring of open-top moulds	15	
2.3.2	Gravity pouring of closed moulds	16	
2.3.3	Horizontal transfer casting	68	
2.3.4	Counter-gravity	72	
2.3.5	Surface tension controlled filling	75	
2.3.6	Inclusion control: filters and traps	78	
2.3.7	Practical calculation of the filling system	93	
Rule 3. <i>Avoid laminar entrainment of the surface film (the non-stopping, non-reversing condition)</i>	102		
3.1	Continuous expansion of the meniscus	102	
3.2	Arrest of vertical progress	103	
		3.3	Waterfall flow 104
		3.4	Horizontal stream flow 104
		3.5	Hesitation and reversal 106
Rule 4. <i>Avoid bubble damage</i>	108		
4.1	Gravity-filled running systems	111	
4.2	Pumped and low-pressure filling systems	112	
Rule 5. <i>Avoid core blows</i>	114		
5.1	Background	114	
5.2	Prevention	117	
Rule 6. <i>Avoid shrinkage damage</i>	120		
6.1	Feeding systems design background	120	
6.1.1	Gravity feeding	123	
6.1.2	Computer modelling of feeding	124	
6.1.3	Random perturbations to feeding patterns	124	
6.1.4	Dangers of solid feeding	125	
6.1.5	The non-feeding roles of feeders	125	
6.2	The seven feeding rules	126	
Rule 1:	Do not feed	126	
Rule 2:	Heat-transfer requirement	127	
Rule 3:	Mass-transfer requirement	128	
Rule 4:	Junction requirement	132	
Rule 5:	Feed path requirement	133	
Rule 6:	Pressure gradient requirement	138	
Rule 7:	Pressure requirement	140	
6.3	The new feeding logic	142	
6.3.1	Background	142	
6.3.2	The new approach	143	
6.4	Active feeding	145	
6.5	Freezing systems design	146	

6.5.1	External chills	147	10.2.1	Rectilinear systems	177
6.5.2	Internal chills	149	10.2.2	Cylindrical systems	178
6.5.3	Fins	150	10.2.3	Trigonal systems	179
			10.2.4	Thin-walled boxes	179
Rule 7.	<i>Avoid convection damage</i>	157	10.3	Location jigs	180
7.1	Convection: the academic background	157	10.4	Clamping points	180
7.2	Convection: the engineering imperatives	157	10.5	Mould design: the practical issues	181
7.3	Convection damage and casting section thickness	160	10.6	Casting accuracy	182
7.4	Countering convection	162	10.7	Tooling accuracy	183
			10.8	Mould accuracy	183
Rule 8.	<i>Reduce segregation damage</i>	163	10.9	Summary of factors affecting accuracy	186
			10.10	Metrology	186
Rule 9.	<i>Reduce residual stress (the 'no water quench' requirement)</i>	166	Appendix		188
9.1	Introduction	166		The 1.5 factor	188
9.2	Residual stress from casting	166		The Bernoulli equation	189
9.3	Residual stress from quenching	167		Rate of pour of steel castings from a bottom-pour ladle	191
9.4	Distortion	172		Running system calculation record	191
9.5	Heat treatment developments	173		Design methodology for investment castings	194
9.6	Epilogue	174	References		195
Rule 10.	<i>Provide location points</i>	175	Index		199
10.1	Datums	175			
10.2	Location points	176			

Preface

Castings can be difficult to get right. Creating things never is easy. But sense the excitement of this new arrival:

The first moments of creation of the new casting are an explosion of interacting events; the release of quantities of thermal and chemical energy trigger a sequence of cataclysms.

The liquid metal attacks and is attacked by its environment, exchanging alloys, impurities, and gas. The surging and tumbling flow of the melt through the running system can introduce clouds of bubbles and Sargasso seas of oxide film. The mould shocks with the vicious blast of heat, buckling and distending, fizzing with the volcanic release of vapours that flood through the liquid metal by diffusion, or reach pressures to burst the liquid surface as bubbles.

During freezing, liquid surges through the dendrite forest to feed the volume contraction on solidification, washing off branches, cutting flow paths, and polluting regions with excess solute, forming segregates. In those regions cut off from the flow, continuing contraction causes the pressure in the residual liquid to fall, possibly becoming negative (as a tensile stress in the liquid) and sucking in the solid surface of the casting. This will continue until the casting is solid, or unless the increasing stress is suddenly dispelled by an explosive expansion of a gas or vapour giving birth to a shrinkage cavity.

The surface sinks are halted, but the internal defects now start.

The subsequent cooling to room temperature is no less dramatic. The solidified casting strives to contract while being resisted by the mould. The mould suffers, and may crush and crack. The casting also suffers, being stretched as on a rack. Silent, creeping strain and stress change and distort the casting, and may intensify to the

point of catastrophic failure, tearing it apart, or causing insidious thin cracks. Most treacherous of all, the strain may *not quite* crack the casting, leaving it apparently perfect, but loaded to the brink of failure by internal residual stress.

These events are rapidly changing dynamic interactions. It is this rapidity, this dynamism, that characterizes the first seconds and minutes of the casting's life. An understanding of them is crucial to success.

This new work is an attempt to provide a framework of guidelines together with the background knowledge to ensure understanding; to avoid the all too frequent disasters; to cultivate the targeting of success; to encourage a professional approach to the design and manufacture of castings.

The reader who learns to guide the production methods through this minefield will find the rare reward of a truly creative profession. The student who has designed the casting method, and who is present when the mould is opened for the first time will experience the excitement and anxiety, and find himself asking the question asked by all foundry workers on such occasions: 'Is it all there?' The casting design rules in this text are intended to provide, so far as present knowledge will allow, enough predictive capability to know that the casting will be not only all there, but all right!

The clean lines of the finished engineering casting, sound, accurate, and strong, are a pleasure to behold. The knowledge that the casting contains neither defects nor residual stress is an additional powerful reassurance. It represents a miraculous transformation from the original two-dimensional form on paper or the screen to a three-dimensional shape, from a mobile liquid

to a permanently shaped, strong solid. It is an achievement worthy of pride.

The reader will need some background knowledge. The book is intended for final year students in metallurgy or engineering, for those researching in castings, and for casting engineers and all associated with foundries that have to make a living creating castings.

Good luck!

This new book is the second of three books dealing with castings. The three books are (i) *Principles* (the new metallurgy of cast metals; the metallurgist's book) (ii) *Practice* (the practical founder's book) and (iii) *Processes* (an appraisal of the various methods of making castings; perhaps a casting buyer's book). The three are intended as a sequence, dealing with the theory and practice of the casting of metals. At the rate at which new understanding is emerging, an additional text may also be required; (iv) *Properties* (a book for everyone).

The second in the series is devoted to the Ten Rules. These are my own checklist to ensure that no key aspect of the design of the manufacturing route for the casting is forgotten.

The Ten Rules listed here are proposed as necessary, but not, of course, sufficient, for the manufacture of reliable castings. It is proposed that they are used in addition to existing necessary technical specifications such as alloy type, strength, and traceability via international standard quality systems, and other well-known and well-understood foundry controls such as casting temperature etc.

Although not yet tested on all cast materials, there are fundamental reasons for believing that the Rules have general validity. They have been applied to many different alloy systems including aluminium, zinc, magnesium, cast irons, steels, air- and vacuum-cast nickel and cobalt, and even those based on the highly reactive metals titanium and zirconium. Nevertheless, of course, although all materials will probably benefit from the application of the Rules, some will benefit almost out of recognition, whereas others will be less affected.

The Ten Rules are first listed in summary form. They are then addressed in more detail in the following ten chapters with one chapter per Rule.

The Rules originated when emerging from a foundry on a memorable sunny day. The author was discussing with indefatigable Boeing enthusiasts for castings, Fred Feiertag and Dale McLellan, that the casting industry had specifications for alloys, casting properties, and casting quality checking systems, but what did not exist but was most needed was a *process specification*. Dale threw out a challenge: 'Write

one!'. The Rules and this book are the outcome. It was not perhaps the outcome that either Dale or I originally imagined. A Process Specification has proved elusive, proving so difficult that I have concluded that it will need a more accomplished author.

The Rules as they stand therefore constitute a first draft of a Process Specification; more like a checklist of casting guidelines. A buyer of castings would demand that the list were fulfilled if he wished to be assured that he was buying the best possible casting quality. If he were to specify the adherence to these Rules by the casting producer, he would ensure that the quality and reliability of the castings was higher than could be achieved by any amount of expensive checking of the quality of the finished product.

Conversely, of course, the Rules are intended to assist the casting manufacturer. It will speed up the process of producing the casting right first time, and should contribute in a major way to the reduction of scrap when the casting goes into production. In this way the caster will be able to raise standards, without any significant increase in costs. Quality will be raised to the point at which castings of quality equal to that of forgings can be offered with confidence. Only in this way will castings be accepted by the engineering profession as reliable, engineered products, and assure the future prosperity of both the casting industry and its customers.

It is recognized that many users of this book will be students of casting technology. For completeness therefore, the strict description of the Rules as intended as the caster's checklist has been relaxed a little. A small addition has been made to paragraph 10, extending the section describing the requirement for location points. This extension includes related aspects not included elsewhere, such as the accuracy of the whole mould assembly, and the many-sided problems of mould design.

A further feature of the work that emerged as the book was being written was the dominance of Chapter 2, the design of the filling systems of castings. It posed the obvious question 'why not devote the book completely to filling systems?'. I decided against this option on the grounds that both caster and customer require products that are good in every respect. The failure of any one aspect may endanger the casting. Therefore, despite the enormous disparity in length of chapters, none could be eliminated; they were all needed.

Finally, it is worth making some general points about the whole philosophy of making castings.

For a successful casting operation, one of the revered commercial goals is the attainment of

product sales being at least equal to manufacturing costs. There are numerous other requirements for the successful business, like management, plant and equipment, maintenance, accounting, marketing, negotiating etc. All have to be adequate, otherwise the business can suffer, and even fail.

This text deals only with the technical issues of the quest for good castings. Without good castings it is not easy to see what future a casting operation can have. The production of good castings can be highly economical and rewarding. The production of bad castings is usually expensive and damaging.

The 'good casting' in this text is defined as one that meets or exceeds the customer's specification.

It is also worth noting at this early stage, that we hope that meeting the customer's specification will be equivalent to meeting or exceeding service requirements. However, occasionally it is necessary to live with the irony that the aims of the customer and the requirements for service are sometimes not in the harmony one would like to see.

These problems illustrate that there are easier ways of earning a living than in the casting industry. But few are as exciting.

J.C.
West Malvern
3 September 2003

The 10 Rules: Summary

1. Start with a good quality melt

Immediately prior to casting, the melt shall be prepared, checked, and treated, if necessary, to bring it into conformance with an acceptable minimum standard. However, preferably, prepare and use only near-defect-free melt.

2. Avoid turbulent entrainment of the surface film on the liquid

This is the requirement that the liquid metal front (the meniscus) should not go too fast. Maximum meniscus velocity is approximately 0.5 ms^{-1} for most liquid metals. This maximum velocity may be raised in constrained running systems or thin section castings. This requirement also implies that the liquid metal must not be allowed to fall more than the critical height corresponding to the height of a sessile drop of the liquid metal.

3. Avoid laminar entrainment of the surface film on the liquid

This is the requirement that no part of the liquid metal front should come to a stop prior to the complete filling of the mould cavity. The advancing liquid metal meniscus must be kept 'alive' (i.e. moving) and therefore free from thickened surface film that may be incorporated into the casting. This is achieved by the liquid front being designed to expand continuously. In practice this means progress only uphill in a continuous uninterrupted upward advance; i.e. (in the case of gravity poured casting processes,

from the base of the sprue onwards). This implies

- Only bottom gating is permissible.
- No falling or sliding downhill of liquid metal is allowed.
- No horizontal flow of significant extent.
- No stopping of the advance of the front due to arrest of pouring or waterfall effects etc.

4. Avoid bubble entrainment

No bubbles of air entrained by the filling system should pass through the liquid metal in the mould cavity. This may be achieved by:

- Properly designed offset step pouring basin; fast back-fill of properly designed sprue; preferred use of stopper; avoidance of the use of wells or other volume-increasing features of filling systems; small volume runner and/or use of ceramic filter close to sprue/runner junction; possible use of bubble traps.
- No interruptions to pouring.

5. Avoid core blows

- No bubbles from the outgassing of cores or moulds should pass through the liquid metal in the mould cavity. Cores to be demonstrated to be of sufficiently low gas content and/or adequately vented to prevent bubbles from core blows.
- No use of clay-based core or mould repair paste unless demonstrated to be fully dried out.

6. Avoid shrinkage

- No feeding uphill in larger section thickness castings because of (i) unreliable pressure gradient and (ii) complications introduced by convection.
- Demonstrate good feeding design by following all Feeding Rules, by an approved computer solidification model, and by test castings.
- Control (i) the level of flash at mould and core joints; (ii) mould coat thickness (if any); and (iii) temperatures of metal and mould.

7. Avoid convection

Assess the freezing time in relation to the time for convection to cause damage. Thin and thick section castings automatically avoid convection problems. For intermediate sections either (i) reduce the problem by avoiding convective loops in the geometry of the casting and rigging,

(ii) avoid feeding uphill, or (iii) eliminate convection by roll-over after filling.

8. Reduce segregation

Predict segregation to be within limits of the specification, or agree out-of-specification compositional regions with customer. Avoid channel segregation formation if possible.

9. Reduce residual stress

No quenching into water (cold or hot) following solution treatment of light alloys. (Polymer quenchant or forced air quench may be acceptable if casting stress is shown to be negligible.)

10. Provide location points

All castings to be provided with agreed location points for pickup for dimensional checking and machining.

Rule 1

Achieve a good quality melt

1.1 Background

It is a requirement that either the process for the production and treatment of the melt shall have been shown to produce good quality liquid, or the melt should be demonstrated to be of good quality, or, preferably, both. A good quality liquid is one that is defined as

- (i) Substantially free from suspensions of non-metallic inclusions in general, and bifilms in particular.
- (ii) Relative freedom from bifilm-opening agents. These include gas in solution and certain alloy impurities (such as Fe in Al alloys) in solution.

It should be noted that such melts are not to be assumed, and, without proper treatment, are probably rare. (Additional requirements, not part of this specification, may also be placed on the melt. For instance, low values of particular solute impurities that have no effect on bifilms.)

Unfortunately, many melts start life with poor, sometimes grossly poor, quality in terms of content of suspended bifilms. Figure 1.1 gives several examples of different poor qualities of liquid aluminium alloy. The figures show results from reduced pressure test (RPT) samples observed by X-ray radiography. Since the samples are solidified under only one tenth of an atmosphere (76 mm compared to 760 mm of full atmospheric pressure) any gas-containing defects, such as bubbles, or bifilms with air occluded in the centres of their sandwich structures, will be expanded by ten times. Thus rather small defects can become visible for the first time.

We can assume [following the conclusions of *Castings* (2003)] that bifilms always initiate

pores, and that the formation of rounded pores simply occurs as a result of the bifilm being opened by excess precipitation of gas, finally achieving a diameter greater than its original length. Thus the RPT is an admirably simple device for assessing (i) the number of bifilms, but (ii) gas content is assessed by the degree of opening of the bifilms from thin crack-like forms to fairly spherical pores.

If the melt contained no gas-containing defects the radiographs would be clear.

However, as we can see immediately, and without any benefit of complex or expensive equipment, the melts recorded in Figure 1.1 are far from this desirable condition. Figure (a) shows a melt with small rounded pores indicating that the bifilms that initiated these defects were particularly small, of the order of 0.1 mm or less. The density of these defects, however, was high, between 10 and 100 defects per cm^3 . Sample (b) has a similar defect distribution, but with slightly higher hydrogen content. Sample (c) illustrates a melt that displayed a deep shrinkage pipe, normally interpreted to mean good quality, but showing that it contained a scattering of larger pores, probably as a result of fewer bifilms, so that the available gas was concentrated on the fewer available sites. Melt (d) has considerably larger bifilms, of size in the region of 5 mm in length, and in a concentration of approximately 1 per cm^3 . Samples (e) and (f) show similar samples but with increasing gas contents that have inflated these larger bifilms to reasonably equiaxed pores.

Naturally, it would be of little use for the casting engineer to go to great lengths to adopt the best designs of filling and feeding systems if the original melt was so poor that a good casting could not be made from it.




Figure 1.1 Radiographs of RPT samples of Al-7Si-0.4Mg alloy illustrating different bifilm populations (courtesy S. Fox)

Thus this section deals with some of the aspects of obtaining a good quality melt. It should be noted that many of these aspects have been already been touched upon in *Castings* (2003).

In some circumstances it may not be necessary to reduce both bifilms and bifilm-opening agents. An interesting possibility for future specifications for aluminium alloy castings (where residual gas in supersaturated solution does not appear to be harmful) is that a double requirement may be made for the content of dissolved gas in the melt to be high, but the percentage of gas porosity to be low. The meeting of this double requirement will ensure to the customer that bifilms are not present. Thus these damaging but undetectable defects will, if present, be effectively labelled and made visible on X-ray radiographs and polished

sections by the precipitation of dissolved gas. It is appreciated that such a stringent specification might be viewed with dismay by present suppliers. However, at the present time we have mainly only rather poor technology, making such quality levels out of reach.

(For steels, the content of hydrogen is a more serious matter, especially if the section thickness of the casting is large. In some steel castings of section thickness above about 100 mm or so, the hydrogen cannot escape by diffusion during the time available for cooling or during the time of any subsequent heat treatments. Thus the high hydrogen content retained in these heavy sections can lead to hydrogen embrittlement, and catastrophic failure of the section by cracking.)

The possible future production of Al alloys for aerospace, with high hydrogen content but low porosity, is a fascinating challenge. As our technology improves such castings may be found not only to be manufacturable, but offer guaranteed reliability of fatigue life, and therefore command a premium price.

The prospect of producing ultra-clean Al alloys that can be demonstrated in this way to be actually extremely clean, raises the issue of contamination of the liquid alloy from the normal metallurgical additions such as the various master alloys, and grain refiners, modifiers, etc. It may be that for clean material, normal metallurgical additions to achieve refinement of various kinds will be found unnecessary, and possibly even counter-productive.

For ductile iron production the massive amounts of turbulence that accompany the addition of magnesium in some form, such as magnesium ferro-silicon, are almost certainly highly damaging to the liquid metal. It is expected that immediately after such nodularization treatment that the melt will be massively dirty. It will be useful therefore to ensure that the melt can dwell for sufficient time for the entrained magnesium-oxide-rich films to float out. The situation is analogous to the treatment of cast iron with CaSi to effect inoculation (i.e. to achieve a uniform distribution of graphite of desirable form). In this case the volume of calcium-oxide-rich films is well known, so that the CaSi treatment is known as a 'dirty' treatment compared to FeSi inoculation. The author is unsure about in-mould treatments therefore with such oxidizable elements as Ca and Mg. Do they give results as good as external treatments? If so, how is this possible? Some work to clarify this situation would be valuable.

For nickel-based superalloys melted and cast in vacuum, it is with regret that the material is, despite its apparently clean melting environment, found to be sometimes as crammed with

oxides (and/or nitrides) as an aluminium alloy (Rashid and Campbell 2004). This is because the main alloying element in such alloys is aluminium, and the high temperature favours rapid formation and thickening of the surface film on the liquid. This occurs even in vacuum, because the vacuum is only, of course, dilute air. Although thermodynamics indicates that aluminium oxide (and perhaps also the nitride) is unstable in vacuum, there is no doubt that a film forms rapidly, and can become entrained, thus damaging the liquid and any subsequent casting. The melting and pouring processes of these alloys also leaves much to be desired. The alloy is melted by induction, and is poured under gravity via a series of sloping launders, falling several times, and finally falling one or two metres or more into steel tubes that act as moulds. This awfully turbulent primary production process for the alloy impairs all the downstream products.

A recent move towards the production of Ni-base alloy bar by horizontal continuous casting is to be welcomed as the first step towards a more appropriate production technique for these key ingredients of our modern aircraft turbines. Production of improved material is currently limited, but should be the subject of demand from customers. Poor casting technology has been the accepted norm within the aircraft industry for too long. (However, as we all know, the aircraft industry is not alone.)

1.2 Melting

The quality of melts can be significantly affected by the type of melting furnace.

For instance the melting of the charge in a single pot is usual. Furnaces of this type automatically ensure that all the oxides on the surface of the charge material will be incorporated into the melt. Thus crucible furnaces, whether bale-out or tilting, or whether heated by gas or electrical resistance, are all of this type. Also included are induction furnaces and reverberatory (meaning reflected heat from a roof) melting units.

The remelting of aluminium alloy sand castings probably represents one of the worst cases, since the oxide skin on a sand casting is particularly thick, having cooled from high temperature in a reactive environment. When the skin is submerged in the melt it can float about, substantially complete. When remelting scrapped cylinder heads by adding them in to one end of a combined melting and holding furnace, the author has fished out the skin of a complete cylinder head from the opposite end.

The remelting of aluminium alloy gravity die (permanent moulded) castings have an oxide skin that is much thinner, and seems to give less problems. Whether this is a real or imagined advantage in view of the damage that can be caused by any entrained oxide skin, irrespective of its thickness, is not clear at this time.

Induction furnaces enjoy the great advantage of extremely rapid melting. However, they have long been regarded with some reserve by aluminium melters because of the electromagnetic stirring, with the suspicion that oxides may therefore be entrained. With normal inductive coil geometry there is a high-pressure region near the centre of the wall that drives a double torroid (a torroid is a ring shape like a doughnut) in directions away from this point. However, there is no evidence known to the author that the stirring is sufficiently rapid to entrain oxide, although such a problem cannot be ruled out. What is certain is that any oxide will have no opportunity to settle out, but this is also true of most of the above crucible furnaces because of the presence of natural convection. Because of the heat input via the walls of the crucible, and heat loss from the top surface, the convective stirring will be expected to take the form of a simple torroid, the flow direction being upwards at the wall, and downwards in the centre. The only significant difference, if any, between these two stirring modes is the rate of stirring. It is possible that the higher energy in the induction furnace may shred films, whereas the natural convective regimes in other furnaces would be expected to conserve the original film size distribution.

The dry hearth type of furnace is quite different. The charge to be melted is heaped onto a dry, sloping refractory floor, called a hearth. As the charge melts, the liquid alloy flows out of its oxide skin and down the sloping hearth into the main melt. The oxide skins present on the surface of the charge materials remain behind, accumulating on the hearth. The pile of dross is raked off the hearth at intervals via a side door. Such melting units are useful in aluminium and copper alloy production.

1.3 Holding

Holding furnaces can also have a significant effect on melt quality.

Holding furnaces were originally selected for their utility in smoothing the supply of molten metal between batch melters and a fluctuating demand from the casting requirements. An additional advantage was the smoothing of temperature and chemical analysis that was unavoidably variable from batch to batch.

The Cosworth Process was perhaps the first to acknowledge that for liquid aluminium alloys the oxide inclusions in a holding furnace could be encouraged to separate simply by a sink and float principle; the metal for casting being taken by a pump from a point at about midway depth where the best quality metal was to be expected.

In contrast, the holding of melts in closed vessels for the low-pressure die casting of aluminium, or the dosing of the liquid metal, are usually impaired by the initial turbulent pour of the melt to fill the furnace. The total pour height is often of the order of a metre. Not only are new oxides folded into the melt in this pouring action, but those oxides that have settled to the floor of the furnace since the last filling operation are stirred into the melt once again. Finally, these enclosed units suffer from the inaccessibility of the melt. This usually restricts any thorough action to improve the melt by any kind of degassing technique.

The Alotech approach to the design of a holding furnace is patented and not available for publication at the time of writing. It is hoped to rectify this in future editions of this work as details of the process are published. Why even mention it at this stage? The purpose of mentioning it here is to illustrate that even apparently simple equipment such as a holding furnace is capable of considerable sophistication, leading to the production of greatly improved processing and products. It takes the concept of melt cleaning and degassing to an ultimate level that probably represents a limit to what can be achieved. In addition, the technique is simple, low capital cost, low running cost, has no moving parts and is operator-free. At this time the technique is being applied only to aluminium alloys.

1.4 Pouring

Most foundries handle their metal from one point to another by ladle. The metal is, of course, transferred out of the ladle by pouring. In most foundries multiple pours are needed to transfer the liquid metal from the melting unit to the mould.

At every pouring operation, it is likely that large areas of oxide film will be entrained in the melt because pour heights are usually not controlled. It is known that pour heights less than the height of the sessile drop cannot entrain the surface oxide. However, such heights are very low; 16 mm for Mg, 13 mm for Al, and only 8 mm for dense metals such as copper-base, and iron and steel alloys.

However, this theoretical limit, while absolutely safe, may be exceeded for some metals with minimum risk. As long ago as 1928 Beck described how liquid magnesium could be transferred from a ladle into a mould by arranging the pouring lip of the ladle to be as close as possible to the pouring cup of the mould, and relatively fixed in position. In this way the semi-rigid oxide tube that formed automatically around the jet remained unbroken, and so protected the falling stream.

Experiments by Din and Campbell (2002) on Al-7Si-0.4Mg alloy have demonstrated that in practice the damage caused by falls up to 100 mm appears controlled and reproducible. This is in close agreement with early observations by Turner (1965) who noted that air was taken into the melt, reappearing as bubbles on the surface when the pouring height exceeded about 90 mm.

Above 200 mm, Din and Campbell (2003) found that random damage was certain. At these high energies of the plunging jet, bubbles are entrained, with the consequence that bubble trails add to the total damage in terms of area of bifilms.

In general, it seems that the lower the pour height the less damage is suffered by the melt. In addition, of course, less metal is oxidized, thus directly saving the costs of unnecessary melt losses. Ultimately, however, it is, of course, best to avoid pouring altogether. In this way losses are reduced to a minimum and the melt is maintained free from damage.

Until recent years, such concepts have been regarded as pipe dreams. However, the development of the Cosworth Process has demonstrated that it is possible for aluminium alloy castings to be made without the melt suffering any pouring action at any point of the process. Once melted, the liquid metal travels along horizontal heated channels, retaining its constant level through the holding furnace, and finally to the pump, where it is pressurized to fill the mould in a counter-gravity mode. Such technology would also appear to be relatively easily applied to magnesium alloys.

The potential for extension of this technology to other alloy base systems such as copper-based or iron-based alloys is less clear. This is because many of these other alloy systems either do not suffer the same problems from bifilms, or do not have the production requirements of some of the high volume aluminium foundries. Thus in normal circumstances, many irons and steels are relatively free from bifilms because of the large density difference between the inclusions and the parent melt, encouraging rapid flotation. Alternatively, many copper-based and steel foundries are more like jobbing shops, where

the volume requirement does not justify a counter-gravity system, and high technology pumps, if they were available, would have problems surviving the oxidation and thermal shock of a stop-go production requirement.

Despite these reservations, the counter-gravity Griffin Process has been impressively successful for the volume production of steel wheels for rail rolling stock. The process produces wheels that require no machining (apart from the centre hub). The outer cast rim runs directly on the steel rail. The products outperform forged steel wheels in terms of reliability in service, earning the process 80 per cent of the market in the USA. What a demonstration of the soundness of the counter-gravity concept, contrasting dramatically with steel castings produced world-wide by gravity pouring, in which defects and expensive upgrading of the casting are the norm.

1.5 Melt treatments

1.5.1 Degassing

Gases dissolved in melts are disadvantageous because they precipitate in bifilms and cause them to unfurl, and even open further as cracks or voids, thereby progressively reducing the mechanical properties of the cast product.

Treatments to reduce the gas content of melts include vacuum degassing. Such a technique has only been widely adopted by the steel industry. If a melt were simply to be placed under vacuum the rate of degassing would be low because the dissolved gas would have to diffuse to the free surface to escape. The slow convection of the melt will gradually bring most of the volume near to the top surface, given time. The process is greatly speeded by the introduction of millions of small bubbles of inert gas via a porous plug or other technique. In this way the process is accomplished in a fraction of the time.

Traditionally the steel industry has used the technique of a carbon boil, in which the creation and floating out of carbon monoxide bubbles from the melt carries away unwanted gases such as oxygen (actively by chemical reaction with carbon) and hydrogen (passively by simple flushing action). The nitrogen in the melt may go up or down depending on its starting value and the nitrogen content of the environment above the melt, since it will tend to equilibrate with the environment. The more recent oxygen steelmaking processes in which oxygen is injected into the melt certainly reduce hydrogen and nitrogen, but require the oxygen to be reduced either chemically by reaction with C or other

deoxidizers such as Si, Mn or Al etc., or use even more modern techniques such as AOD (argon-oxygen-decarburization). We shall not dwell further on these sophistications, since these specialized techniques, so well understood and well developed for steel, are almost unused elsewhere in the casting industry.

By comparison, the approaches to the degassing of aluminium alloys, containing only hydrogen, has for many years been primitive. Only recently have effective techniques been introduced.

For instance until approximately 1980, aluminium was commonly degassed using immersed tablets of hexachlorethane that thermally degraded in the melt to release large bubbles of chlorine and carbon (the latter as smoke). Alternatively, a primitive tube lance was used to introduce a gas such as nitrogen or argon. Again large bubbles of the gas were formed. These techniques involving the generation of large bubbles were so inefficient that little dissolved gas could be removed, but the creation of large areas of fresh melt to the atmosphere each time a bubble burst at the surface of the melt provided an excellent opportunity for the melt to equilibrate with the environment. Thus on a dry day the degassing effect might be acceptable. On a damp day, or when the flue gases from the gas-fired furnace were suffering poor extraction, the melt could gain hydrogen faster than it could lose it. This poor rate of degassing, combined with the high rate of regassing from interaction with the environment, led to variable and unsatisfactory results.

Rotary degassing came to the rescue. The use of a rapidly rotating rotor to chop bubbles of inert gas into fine clouds raised the rate of degassing and lowered the rate of regassing. Thus effective degassing could be reliably achieved in times that were acceptable in a production environment.

It is to be hoped that this is not the end of the story for the degassing of aluminium alloys. For instance the use of an 'inert' gas is a convenient untruth. Even sources of high purity inert gas contain sufficient oxygen as an impurity to create an oxide film on the inside of the surface of the bubbles. Thus millions of very thin bifilms must be created during this degassing process. For a new rotor, or after a weekend, the rotor and its shaft will have absorbed considerable quantities of water vapour (most refractories can commonly absorb water up to 10 per cent of their weight). Furthermore, the gas lines are usually not clean, or contain long rubber or plastic tubing. These materials absorb and therefore leak large quantities of volatiles into

the degassing line. Thus by the time the gas arrives in the melt it is usually not particularly inert. At this time no-one knows whether the oxides created as by-products of degassing are negligible or whether they seriously reduce properties. What is known is that aluminium alloys are usually greatly improved by rotary degassing. Whether further improvements can be secured is not clear.

Part of the considerable benefit of rotary degassing is not merely the degassing action. It seems that the millions of tiny bubbles attach to oxides in the melt and float them to the top, where they can be skimmed off. After treatment, after the rotor has been raised out of the melt and the surface skimmed, a test of the efficiency of the cleaning action is simply to look at the surface of the melt. If the cleaning action is not complete, during the next few minutes small particles will be seen to arrive at the surface, under the surface oxide film, which is seen to be pushed upwards by the particle. The treatment can be repeated until no further arrival of debris can be seen. The melt surface then retains its pristine mirror smoothness. The melt can then be pronounced as clean as the treatment can achieve.

The action of rotary degassing on oxide films merits further examination. Much has been written on the benefits of small bubbles on the rate of degassing (although the rate of regassing from a damp rotor, or from the environment has been neglected). The action of the bubbles to eliminate films has in general been overlooked. However, the size of bubbles in relation to the efficiency of removal of films is probably critical. For instance, large bubbles will displace large volumes of melt during their rise to the surface, thus displacing films sideways, so that the film and bubble never make contact. On the other hand, small bubbles will displace relatively little liquid, and so be able to impact on relatively large films in their path. Thus such contacted films will be buoyed up to the surface. The mutual contact would be expected to become important when bubbles and films were approximately similar in size. Thus small bubbles will take out correspondingly smaller oxide films. This predicted effect deserves to be demonstrated experimentally at some future date.

An experience by the author illustrates some of the misconceptions surrounding the dual role of hydrogen and oxide bifilms in aluminium melts. An operator used his rotary degasser for 5 minutes to degas 200 kg of Al alloy. The melt was tested with a reduced pressure test (RPT; see below) sample that was found to contain no bubbles. The melt was therefore deemed to be

degassed. The melt was immediately poured into a transfer ladle on a fork lift truck and conveyed to a low-pressure die casting furnace, into which it was poured. The melt in the low-pressure furnace was then tested again by RPT, and the sample found to contain many bubbles. The operator was baffled. He could not understand how so much gas could have re-entered the melt in only the few minutes required for the transfer.

The truth is, of course, that the melt was insufficiently degassed with only 5 minutes of treatment. In fact with a damp rotor the gas level is likely to rise initially (getting worse before it gets better!). The RPT showed no bubbles *not* because the hydrogen was low, but because the short treatment had clearly been sufficiently successful to remove a large proportion of the bifilms that were the nuclei for the bubbles. This high hydrogen metal was then poured twice, each time from a considerable height, re-introducing copious quantities of oxide bifilms that act as excellent nuclei, so that the RPT could now reveal its high hydrogen content.

1.5.2 Additions

Additions to melts are made for a variety of reasons. These can include additives for chemical degassing (as the addition of Al to steel to fix oxygen and nitrogen) or grain refinement (as in the addition of titanium and/or boron or carbon to Al alloys). Sometimes it seems certain that the poor quality of such materials (perhaps melted poorly and cast turbulently, and so containing a high level of oxides) can contaminate the melt directly.

Indirectly, however, the person in charge of making the addition will normally be under instructions to stir the melt to ensure the dissolution of the addition and its distribution throughout the melt. Such stirring actions can disturb the sediment at the bottom of melts, efficiently re-introducing and re-distributing those inclusions that had spent much time in settling out. The author has vivid memories of wrecking the quality of early Cosworth melts in this way: the addition of grain refiners gave wonderfully grain refined castings, and should have improved feeding, and therefore the soundness of the castings. In fact, all the castings were scrapped because of a rash of severe microporosity, initiated almost certainly on the stirred-up oxides that constituted the sediment in the holding furnace.

Additions of Sr to aluminium melts have often been accused of also adding hydrogen because of the porosity that has often been

noted to follow such additions. Here again, it seems unlikely that sufficient hydrogen could be introduced by such a small addition. Perhaps the problem is with stirred sediment, or other more complex reasons as are discussed at some length in *Castings 2003*.

1.6 Filtration

Filtration is perhaps the most obvious way to remove suspended solids from liquid metals. However, it is not without its problems, as we shall see.

The action of a filter in the supply of liquid metal in the launder of a melt distribution system in a foundry or cast house (i.e. a foundry for continuous casting of long products) is rather different from its action in the running system of a shaped casting.

The filters used in castings act for only the few seconds or minutes that the casting is being poured. The velocities through them are high, usually several metres per second, compared to the more usual 0.1 m/s rates in launders. The filtration effect is further not helped by the concentration of flow through an area often a factor of up to 100 smaller than that used in launder systems. The total volume per second rates at which casting filters are used are therefore higher by a factor of around 1000. It is hardly surprising therefore that the filtration action is reduced nearly to zero. However, the effect on the flow is profound. The use of filters in running systems is dealt with in detail later in the book. We concentrate in this section on filters in launder systems.

The treatment of tonnage quantities of metal has been developed only for the aluminium casting industry. The central problem for most workers in this field is to understand how the filters work, since the filters commonly have pore sizes of around 1 mm, whereas, puzzlingly, they seem to be effective in removing a high percentage of inclusions of only 0.1 mm diameter. Most researchers expand at length, listing the mechanisms that might be successful to explain the trapping of such small solid particles. Unfortunately, these conjectures are probably not helpful, and are not repeated here.

The fact is that the important solids being filtered in aluminium alloys are not particles resembling small solid spheres, as has generally been assumed. The important particles are films (actually always double films that we have called bifilms). Once this is appreciated the filtration mechanism becomes much easier to understand. The films are often of size 1 to 10 mm, and so are, in principle, easily trapped

by pores of 1 mm diameter. Such bifilms are not easily seen in their entirety in an optical microscope, the visible portions appearing to be much smaller, explaining why filters appear to arrest particles smaller than their pore diameters.

We need to take care. This explanation, while probably having some truth, may oversimplify the real situation. In *Castings (2003)*, the life of the bifilm was described as starting as the folding in of a planar crack-like defect as a result of surface turbulence. However, internal turbulence wrapped the defect into a compact form, reducing its size by a factor of 10 or so. In this form it could pass through a filter, and finally open once again in the casting as the liquid metal finally came to rest, and bifilm-opening (i.e. unfurling) processes started to come into action.

In more detail now, the trapping of compact forms of films is explained by their irregular and changing form. During the compacted stage of their life, they will be constantly in a state of flux, ravelling and unravelling as they travel along in the severely turbulent flow. Two-dimensional images of such defects seen on polished sections always show loose trailing fragments. Thus in their progress through the filter, one such end could become attached, possibly wrapping over a web or wall of filter material. The rest of the defect would then roll out, unravelling in the flow, and be flattened against the internal surfaces of the filter, where it would remain fixed in the tranquil boundary layer.

1.6.1 Packed beds

In practice, in DC (direct chill) continuous casting plants, filters have been used for many years, sited in the launder system between the melting furnaces and casting units. Commonly, the filtration system is a large and expensive installation, comprising a crucible furnace that contains a divided crucible. One half is filled with refractory material such as alumina balls, or tabular alumina. The flow of the melt down one half of the crucible, through a connecting port, and up through the deep packed bed of such systems has been shown to be effective in greatly reducing the inclusion count (the number of inclusions per unit area). However, it is known that the accidental disturbance of such filters releases large quantities of inclusions into the melt stream. This has also been reported when enthusiastic operators see the filter becoming blocked by the increasing upstream level of the melt, and stir the bed with iron rods to ease the flow of metal. It is not easy to imagine actions that could be more counter-productive.

Some work has been carried out on filtering liquid aluminium through packed beds of tabular or ball alumina (Mutharasan *et al.* 1981) and through bauxite, aluminosilicates, magnesite, chrome-magnesite, limestone, silicon carbide, carbon, and steel wool (Hedjazi *et al.* 1975). This latter piece of work demonstrated that all of the materials were effective in reducing macro-inclusions. This is perhaps to be expected as a simple sieve effect. However, only the aluminosilicates were really effective in removing any micro-inclusions and films, whereas the carbon and chrome-magnesite removed only a small percentage of films and appeared to actually increase the number of micro-inclusions. The authors suggest that the wettability of the inclusions and the filter material is essential for effective filtration.

An interesting application of a chemically active packed bed is that by Geskin *et al.* (1986), in which liquid copper is passed through charcoal to provide oxygen-free copper castings. It is certain, however, that the charcoal will have been difficult to dry thoroughly, so that the final casting may be somewhat high in hydrogen. Because of the low oxygen content it is likely that hydrogen pores will not be nucleated (as discussed in *Castings* (2003)). The hydrogen is expected therefore to stay in solution and remain harmless.

1.6.2 Alternative varieties of filters

Other large and expensive filtration systems include the use of a pack of porous tubes, sealed in a large heated box, through which the aluminium is forced. The pores in this case are of the order of 0.25 mm, with the result that the filter takes a high head of metal to prime it. However, the technique is not subject to failure because of disturbance, and guarantees high quality of liquid metal.

Other smaller and somewhat cheaper systems that have been used include a ceramic foam filter, usually designed to be housed in a box, sited permanently below the surface of the melt. The velocities through the filter are usually low, encouraged by the large area of the filter, usually at least 300×300 mm. The filter is only brought out into the air to be changed when the

metal level either side of the filter box shows a large difference, indicating that the filter is becoming blocked.

The efficiency of many filtration devices can be understood when it is assumed that the important filtration action is the removal of films (not particles). Thus glass cloth is widely used to good effect in many different forms in the Al casting industry.

1.6.3 Practical aspects

It is typical of most filtration systems that the high quality of metal that they produce (often at considerable expense) is destroyed by thoughtless handling of the melt downstream.

Even the filter itself can give difficulties in this way. For instance if the melt exits the filter downwards, or even horizontally, causing fine jets of metal to form in the air, and plunge into the melt, additional oxide defects are necessarily created downstream. The avoidance of this problem is an important aspect of the designing of effective filters into the running systems of shaped castings, as will be discussed later.

Many have wondered whether the filter itself causes oxides because the flow necessarily emerges in a divided state, and therefore must create double films in the hundreds of confluence events. Video observations by the author on a stream of aluminium alloy emerging from a ceramic foam filter with a pore diameter close to 1 mm have helped to clarify the situation. It seems true that the flow emerges divided as separate jets. However, within a few millimetres (apparently depending on the flow rate of the metal) the separate jets merge. Thus the oxide tubes formed around the jets appear to be up to 10 mm long when the melt travelled at around 500 mm s^{-1} but remained attached to the filter. The oxide tubes did not extend further because after the streams merged oxygen was necessarily excluded. The forest of tubes was seen to wave about in the flow like underwater grass. It is possible that more rapid flows might cause the grass to detach as a result of its greater length and the higher speed of the metal. This seems much more likely in the conditions of the running system of a casting.

Rule 2

Avoid turbulent entrainment (the critical velocity requirement)

The avoidance of surface turbulence is probably the most complex and difficult Rule to fulfil when dealing with gravity pouring systems.

The requirement is all the more difficult to appreciate by many in the industry, since everyone working in this field has always emphasized the importance of working with 'turbulence-free' filling systems for castings. Unfortunately, despite all the worthy intentions, all the textbooks, all the systems, and all the talk, so far as the author can discover, it seems that no-one appears to have achieved this target so far. In fact, in travelling around the casting industry, it is quite clear that the majority (at least 80 per cent) of all defects are directly caused by turbulence. Thus the problem is massive; far more serious than suspected by most of us in the industry.

To understand the fundamental root of the problem, it will become clear in Section 2.1 that any fall greater than the height of the sessile drop (of the order of 10 mm) causes the metal to exceed its critical velocity, and so introduces the danger of defects in the casting. As most falls are in fact in the range 10 to 100 times greater than this, and as the damage is likely to be proportional to the energy involved (i.e. proportional to the square of the velocity) the damage so created will usually be expected to be in the range 100 to 10 000 times greater. Thus in the great majority of castings that are poured simply under the influence of gravity, there is a major problem to ensure its integrity. In fact, the situation is so bad that the best outcome of many of the solutions proposed in this book is damage limitation. Effectively, it has to be

admitted that at this time it seems impossible to guarantee the avoidance of some damage when pouring liquid metals.

This somewhat depressing conclusion needs to be tempered by a number of factors.

First; the world has come to accept castings as they are. Thus any improvement will be welcome. This book described techniques that will create very encouraging improvements.

Second; this book is merely a summary of what has been discovered so far in the development of filling system design. Better designs are to be expected now that the design parameters (such as critical velocity, critical fall height, etc.) are defined.

Third; there *are* filling systems that can yield, in principle, perfect results.

Of necessity, such perfection is achieved by fulfilling Rule 2 by avoiding the transfer of the melt into the mould by pouring downhill under gravity. Thus, considering the three directions of filling a mould:

- (i) downhill pouring under gravity;
- (ii) horizontal transfer into the mould (achieved by tilt casting in which the tilt conditions are accurately controlled);
- (iii) uphill (counter-gravity) casting in which the melt is caused to fill the mould in only an uphill mode;

only the last two processes have the potential to deliver castings of near perfect quality. In my experience, I have found that in practice it is often difficult to make a good casting by gravity, whereas by a good counter-gravity process

(i.e. a process observing all the 10 Rules) it has been difficult to make a bad casting. The jury is still out on horizontal transfer by tilting. This approach has great potential, but requires a dedicated effort to achieve the correct conditions.

Thus in summary, filling of moulds can be carried out down, along, or up. Only the 'along' and 'up' modes totally fulfil the non-surface turbulence condition.

However, despite all its problems, it seems more than likely that the downward mode, gravity casting, will continue to be with us for the foreseeable future. Thus we shall devote some considerable time to the damage limitation exercises that can offer considerably improved products, even if, unfortunately, those products cannot be ultimately claimed as perfect. Most will shed no tears over this conclusion. Although potential perfection in the along and up modes is attractive, the casting business is all about making adequate products; products that meet a specification and at a price a buyer can afford.

The question of cost is interesting; perhaps the most interesting. Of course the costs have to be right, and often gravity casting is acceptable and sufficiently economical. However, more often than might be expected, high quality and low cost can go together. An improved gravity system, or even one of the better counter-gravity systems, can be surprisingly economical and effective. Such opportunities are often overlooked. It is useful to watch out for such benefits.

2.1 Maximum velocity requirement

One day, I was seated in the X-ray radiographic room using an illuminated screen to study a series of radiographic films of cylinder head castings made by our recently developed casting system at Cosworth. Each radiograph in turn was beautiful, having a clear, 'wine glass' perfection that every founder dreams of. I was at peace with the world. However, suddenly, a radiograph appeared on the screen that was a total disaster. It had gas bubbles, shrinkage porosity, hot tears, and sand inclusions. I was shocked, but sensed immediately what had happened. I shot out to query Trevor, our man on the casting station. 'What happened to this casting?' He admitted instantly 'Sorry. I put the metal in too fast.'

This was a lesson that remained with me for years. This chance experiment by counter-gravity, using an electromagnetic pump allowing independent control of the ingate velocity, had kept constant all the other casting variables




Figure 2.1 The extremes of velocity entering the mould from the gate; zero, critical and high.

(temperature, metal quality, alloy content, mould geometry, aggregate type, binder type, etc.), showing them to be of negligible importance. Clearly, the ingate velocity was dominant. By only changing the speed of entry of metal we could move from complete success to complete failure.

Thinking further about this, common sense tells us all that there is an optimum velocity at which a liquid metal should enter a mould. The concept is outlined in Figure 2.1. At a velocity of zero the melt is particularly safe (Figure 2.1a), being free from any danger of damage. Regrettably, this condition is not helpful for the filling of moulds. Conversely, at extremely high velocities the melt will enter like a jet of water from a fire-fighter's hosepipe (Figure 2.1c), and is clearly damaging to both metal and mould. At a certain intermediate velocity the melt rises to just that height that can be supported by surface tension around the periphery of the spreading drop (Figure 2.1b). The theoretical background to these concepts is dealt with at length in the first book in this series *Castings (Principles)* (2003). For nearly all liquid metals this critical velocity is close to 0.5 m s^{-1} . This value is of central importance in the casting of liquid metals, and will be referred to repeatedly in this section.

The liquid drop, emerging close to its critical velocity, and spreading slowly from the ingate is closely in equilibrium, its surface tension holding the drop in its compact shape, just balancing the head of pressure tending to spread it because of its density. This slowly expanding drop is closely similar to a sessile (Latin 'sitting') drop. (The word contrasts with glissile drop, meaning a gliding or sliding drop.) A sessile drop of Al sitting on a non-wetted substrate is approximately 12.5 mm high. Corresponding values for other liquids are Fe 10 mm, Cu 8 mm, Zn 7 mm, Pb 4, water about 5 mm.

Recent research has demonstrated that if the liquid velocity exceeds the critical velocity there is a danger that the surface of the liquid metal may be folded over by surface turbulence. If a perturbation to the surface exceeds the height of the sessile drop the liquid can no longer be supported by its surface tension. It will therefore fall back under gravity and may thus entrain its own surface in an enfolding action. At risk of overly repeating this important phenomenon, this entrainment of the surface can occur if there is sufficient energy in the form of velocity in the bulk liquid to perturb the surface against the smoothing action of surface tension. In addition, notice that damage is not necessarily created by the falling back of the metal. The falling is likely to be chaotic, so that any folding action may or may not occur. The significance of the critical velocity is clear therefore: above the critical velocity there is the danger of surface entrainment leading to defect creation. Below the critical velocity the melt is safe from entrainment problems.

(It is perhaps useful to remind the reader that entrainment may still occur as a result of surface contraction. Thus the steady forward advance of the meniscus, Rule 3, remains another central axiom of good filling systems.)

Japanese workers optimizing the filling of their design of vertical stroke pressure die casting machine using both experiment and computer simulation (Itamura 1995) confirm the critical velocity of 0.5 m s^{-1} , finding that air bubbles have a chance to become entrained above this value. (These workers go further to define the amount of liquid that needs to be in the mould cavity to suppress entrainment at higher ingate velocities.)

Normally, the surface is covered with an oxide film, although many other types of films are possible in different circumstances. A common alternative is a graphitic film. There is a chance therefore that if the speed of the liquid exceeds this critical velocity its surface film may be folded into the bulk of the liquid. This folding action is an entrainment event. It leads to a variety of problems in the liquid that we can collectively call entrainment defects. The major entrainment defects are air bubbles and doubled-over oxide films. The author has named these folded-in films 'bifilms' to emphasize their double, folded-over nature. Because the films are necessarily folded dry side to dry side, there is little or no bonding between these dry interfaces, so that the double films act as cracks. The cracks (alias bifilms) become frozen into the casting, lowering the strength and fatigue resistance. Bifilms may also create leak paths, causing leakage failures.

The folding-in of the oxide is a random process, leading to scatter and unreliability in the properties and performance of the product on a casting to casting, day to day, and month to month basis during a production run.

The different qualities of metals arriving in the foundry from batch to batch will also be expected to contain different quantities and different forms of bifilms. Thus the performance of the foundry will suffer further variation. This is the reason for Rule 1. The foundry needs to have procedures in place to smooth variations of its incoming raw material.

Looking a little more closely at the detail of critical velocities for different liquids, it is close to 0.4 m s^{-1} for dense alloys such as irons, steels and bronzes and about 0.5 m s^{-1} for liquid aluminium alloys. The value is 0.55 to 0.6 m s^{-1} for Mg and its alloys. Taking an average of about 0.5 m s^{-1} for all liquid metals is usually good enough for most purposes related to the design of filling systems for castings, and will be generally used in this book.

The maximum velocity condition effectively forbids top gating of castings (i.e. the planting of a gate in the top of the mould cavity, causing the metal to fall freely inside the mould cavity). This is because liquid aluminium reaches its critical velocity of about 0.5 m s^{-1} after falling only 12.5 mm under gravity. The critical velocity of liquid iron or steel is exceeded after a fall of only about 10 mm (these are, of course, the heights of the sessile drops). Naturally, such short fall distances are always exceeded in practice in top gated castings, leading to the danger of the incorporation of the surface films, and consequent leakage and crack defects.

Castings that are made in which velocities everywhere in the mould never exceed the critical velocity are consistently strong, with high fatigue resistance, and are leak-tight (if properly fed, of course, so as to be free from shrinkage porosity).

Experiments on the casting of aluminium have demonstrated that the strength of castings may be reduced by as much as 90 per cent or more if the critical velocity is exceeded. The corresponding defects in the castings are not always detected by conventional non-destructive testing such as X-ray radiography or dye penetrant, since, despite their large area, the folded oxide films are thin, and do not necessarily give rise to any significant surface indications.

The speed requirement automatically excludes conventional pressure die-castings as having significant potential for reliability, since the filling speeds are usually 10 to 100 times greater than the critical velocity.

Over recent years there have been welcome moves, introducing some special developments of high-pressure technology that are capable of meeting this requirement. These include the vertical injection squeeze casting machine, and the shot control techniques. Such techniques can, in principle, be operated to fill the cavity through large gates at low speeds, and without ingress of air into the liquid metal. Such castings require to be sawn, rather than broken, from their filling systems of course. Unfortunately, the castings remain somewhat impaired by the action of pouring into the shot sleeve. Even here, these problems are now being addressed by some manufacturers, with consequent benefits to the integrity of the castings.

Other uphill filling techniques such as low-pressure filling systems are capable of meeting Rule 2. Even so, it is regrettable that the critical velocity is practically always exceeded during the filling of the low-pressure furnace itself because of the severe fall of the metal as it is transferred into the pressure vessel, so that the metal is damaged even prior to casting. In addition, many low-pressure die casting machines are in fact so poorly controlled on flow rate, that the speed of entry into the die greatly exceeds the critical velocity, thus negating one of the most important potential benefits of the low-pressure system. Processes such as the Cosworth Process avoid these problems by never allowing the melt to fall at any stage of processing, and control its upward speed into the mould by electromagnetic pump.

Metals that can also suffer from entrained surface films are suggested to be the ZA (zinc-aluminium) alloys and ductile irons. Carbon and stainless steels are thought to be similar, although in some of these systems the entrained bifilms agglomerate as a result of being partially molten and therefore somewhat sticky. They therefore remain more compact, and float out more easily to form surface imperfections in the form of slag macroinclusions on the surface. For a few materials, particularly alloys based on the Cu-10Al types (aluminium- and manganese-bronzes) the critical velocities were originally thought to be much lower, in the region of only 0.075 m s^{-1} . However, from recent work at Birmingham this low velocity seems to have been a mistake, probably resulting from the confusion caused by bubbles entrained in the early part of the filling system. With well-designed filling systems, the aluminium-bronzes accurately fulfil the theoretically predicted 0.4 m s^{-1} value for a critical ingate velocity (Halvae and Campbell 1997).

Because of the central importance of the concept of critical velocity, the reader will

forgive a re-statement of some aspects in this summary.

- (i) Even if the melt does jump higher than the height of a sessile drop, when it falls back into the surface there is no certainty that it will enfold its surface film. These tumbling motions in the liquid can be chaotic, random events. Sometimes the surface will fold badly, and sometimes not at all. This is the character of surface turbulence: it is not predictable in detail. The key aspect of the critical velocity is that at velocities less than the critical velocity the surface is safe. Above the critical velocity there is the danger of entrainment damage. The criterion is a necessary but not sufficient condition for entrainment damage.
- (ii) If the whole, extensive surface of a liquid were moving upwards at a uniform speed, but exceeding the critical velocity, clearly no entrainment would occur. Thus the surface disturbance that can lead to entrainment is more accurately described not merely as a velocity but in reality a velocity difference. It might therefore be defined more accurately as a critical velocity gradient measured across the liquid surface. For those of a theoretical bent, the critical gradient might be defined as the velocity difference achieving the critical velocity along a distance in the surface of the order of the sessile drop radius (approximately half its height) in the liquid surface. To achieve reasonable accuracy, this approach requires one to allow for the reduction in drop height with velocity. Hirt (2003) solves this problem with a delightful and novel approach, modelling the surface disturbances as arrays of turbulent eddies, and achieves convincing solutions for the simulation of entrainment at hydraulic jumps and plunging jets. Such niceties are neglected here. The problem does not arise when considering the velocity of the melt when emerging from a vertical ingate into a mould cavity. In that situation, the ingate velocity and its relation to the critical velocity is clear.
- (iii) If the melt is travelling at a high speed, but is constrained between narrowly enclosing walls, it does not have the room to fold-over its advancing meniscus. Thus no damage is suffered by the liquid despite its high speed, and despite the high risk involved. This is one of the basic reasons underlying the design of extremely narrow channels for filling systems that are proposed in this book.

2.2 The 'no fall' requirement

It is quickly shown that if liquid aluminium is allowed to fall more than 12.5 mm then it exceeds the critical 0.5 ms^{-1} . The critical fall height can be seen to be a kind of re-statement of the critical velocity condition. Similar critical velocities and critical fall heights can be defined for other liquid metals. The critical fall heights for all liquid metals are in the range 3 to 15 mm.

It follows immediately that top gating of castings almost without exception will lead to a violation of the critical velocity requirement. Many forms of gating that enter the mould cavity at the mould joint, if any significant part of the cavity is below the joint, will also violate this requirement.

In fact, for conventional sand and gravity die casting, it has to be accepted that some fall of the metal is necessary. Thus it has been accepted that the best option is for a single fall, concentrating the loss of height of the liquid at the very beginning of the filling system. The fall takes place down a conduit known as a sprue, or down runner. This conduit brings the melt to the lowest point of the mould. The distribution system from that point, consisting of runners and gates, should progress only uphill.

Considering the mould cavity itself, the requirement effectively rules that all gates into the mould cavity enter at the bottom level, known as bottom gating. The siting of gates into the mould cavity at the top (top gating) or at the joint (gating at the joint line) are not options if safety from surface turbulence is required.

Also excluded are any filling methods that cause waterfall effects in the mould cavity. This requirement dictates the siting of a separate ingate at every isolated low point on the casting.

Even so, the concept of the critical fall distance does require some qualification. If the critical limit is exceeded it does not mean that defects will necessarily occur. It simply means that there is a risk that they may occur. This is because the energy of the liquid is now sufficiently high that the melt is potentially able to enfold in its own surface. Whether a defect occurs or not is now a matter of chance. (This contrasts, of course, with falls of less than the critical height. In this case there is no chance that a defect can occur, the regime being completely safe.)

There is, however, further qualification that needs to be applied to the critical fall distance. This is because the critical value quoted above has been worked out for a liquid, neglecting the presence of any oxide film. In practice, it seems that for some liquid alloys, the surface oxide has

a certain amount of strength and rigidity, so that the falling stream is contained in its oxide tube and so is enabled to better resist the conditions that might enfold its surface. This behaviour has been investigated for aluminium alloys (Din, Kendrick and Campbell 2003). It seems that although the original fall distance limit of 12.5 mm continues to be the safest option, fall heights of up to about 100 mm might be allowable in some instances, possibly depending somewhat on the precise alloy composition. However, falls greater than 200 mm definitely entrain defects; the velocity of the melt in this case is about 2 ms^{-1} so that entrainment seems unavoidable. Also, of course, other alloys may not enjoy the benefits of the support of a tube of oxide around the falling jet. This benefit requires to be investigated in other alloy systems to test what values beyond the theoretical limits may be used in practice.

The initial fall down the sprue in gravity-filled systems does necessarily introduce some oxide damage into the metal. For this reason it seems reasonable to conclude that gravity-poured castings will never attain the degree of reliability that can be provided by counter-gravity and other systems that can avoid surface turbulence.

Of necessity therefore, it has to be accepted that the no-fall requirement applies to the design of the filling system downstream of the base of the sprue. The damage encountered in the fall down the sprue has to be accepted; although with a good sprue and pouring basin design this initial fall damage can be reduced to a minimum as we shall see.

It is a matter of good luck that it seems that for some alloys much of the oxide introduced in this way does not appear to find its way through and into the mould cavity. It seems that much of it remains attached to the walls of the sprue. This surprising effect is clearly seen in many top-gated castings, where most of the oxide damage (and particularly any random leakage problem) is confined to the area of the casting under the point of pouring, where the metal is falling. Extensive damage does not seem to extend into those regions of the casting where the speed of the metal front decreases, and where the front travels uphill, but there does appear to be some carry-over of defects. Thus the provision of a filter immediately after the completion of the fall is valuable. It is to be noted, however, that significant damage will still be expected to pass through the filter.

The requirement that the filling system should cause the melt to progress only uphill after the base of the sprue forces the decision that the runner must be in the drag and the gates

must be in the cope for a horizontal single jointed mould (if the runner is in the cope, then the gates fill prematurely, before the runner itself is filled, thus air bubbles are likely to enter the gates). The 'no fall' requirement may also exclude some of those filling methods in which the metal slides down a face inside the mould cavity, such as some tilt casting type operations. This undesirable effect is discussed in more detail in the section devoted to tilt casting.

It is noteworthy that these precautions to avoid the entrainment of oxide films also apply to casting in inert gas or even in vacuum. This is because the oxides of Al and Mg (as in Al alloys, ductile irons, or high temperature Ni-base alloys for instance) form so readily that they effectively 'getter' the residual oxygen in any conventional industrial vacuum, and form strong films on the surface of the liquid.

Rule 2 applies to 'normal' castings with walls of thickness over 3 or 4 mm.

For channels that are sufficiently narrow, having dimensions of only a few millimetres, the curvature of the meniscus at the liquid front can keep the liquid front from disintegration. Thus narrow filling system geometries are valuable in their action to conserve the liquid as a coherent mass, and so acting to push the air out of the system ahead of the liquid. The filling systems therefore fill in one pass.

A good filling action, pushing the air ahead of the liquid front as a piston in a cylinder, is a critically valuable action. Such systems deserve a special name such as perhaps 'one pass filling (OPF) designs'. Although I do not usually care for such jargon, the special name emphasizes the special action. It contrasts with the turbulent and scattered filling often observed in systems that are over-generously designed, in which the melt can be travelling in two directions at once along a single channel. A fast jet travels under the return wave that rolls over its top, rolling in air and oxides.

For a wide, narrow, horizontal channel, any effect of surface tension is clearly limited to channels that have dimensions smaller than the sessile drop height for that alloy. Thus for Al alloys the maximum channel height would be 12.5 mm, although even this height would exert little influence on the melt, since the roof would just touch the liquid, exerting no pressure on it. Similarly, taking account that the effect of surface tension is doubled if the curvature of the liquid front is doubled by a second component of the curvature at right angles, a channel of square section could be 25 mm square, and be contained just by surface tension. In practice, however, for any useful restraint from the walls of channels, these dimensions require to be at

least halved, effectively compressing the liquid into the channel.

For very thin walled castings, of section thickness less than 2 mm, the effect of surface tension in controlling filling becomes predominant. The walls are so much closer than the natural curvature of a sessile drop that the meniscus is effectively compressed, and requires the application of pressure to force it into such narrow gaps. The liquid surface is now so constrained that it is not easy to break the surface, i.e. once again there is no room for splashing or droplet formation. Thus the critical velocity is higher, and metal speeds can be raised by approximately a factor of 2 without danger.

In very thin walled castings, with walls less than 2 mm thickness, the tight curvature of the meniscus becomes so important that filling can sometimes be without regard to gravity (i.e. can be uphill or downhill) since the effect of gravity is swamped by the effect of surface tension. This makes even the uphill filling of such thin sections problematical, because the effective surface tension exceeds the effect of gravity. Instabilities therefore occur, whereby the moving parts of the meniscus continue to move ahead in spite of gravity because of the reduced thickness of the oxide skin at that point. Conversely, other parts of the meniscus that drag back are further suppressed in their advance by the thickening oxide, so that a run-away instability condition occurs. This dendritic advance of the liquid front is no longer controlled by gravity in very thin castings, making the filling of extensive sections, whether horizontal or vertical, a major filling problem.

The problem of the filling of thin walls occurs because the flow happens, by chance, to avoid filling some areas because of random meandering. Such chance avoidance, if prolonged, leads to the development of strong oxide films, or even freezing of the liquid front. Thus the final advance of the liquid to fill such regions is hindered or prevented altogether.

The dangers of a random filling pattern problem are relieved by the presence of regularly spaced ribs or other geometrical features that assist organizing the distribution of liquid. Random meandering is thereby discouraged and replaced by regular and frequent penetration of the area, so that the liquid front has a better chance to remain 'live', i.e. it keeps moving so that a thick restraining oxide is given less chance to form.

The further complicating effect of the microscopic break-up of the front known as micro-jetting (*Castings 2003*) observed in sections of 2 mm and less in sand and plaster moulds is not yet understood. The effect has not

yet been investigated, and may not occur at all in the dry mould conditions such as are found in gravity die casting.

2.3 Filling system design

Getting the liquid metal out of the crucible or melting furnace and into the mould is a critical step when making a casting: it is likely that most casting scrap arises during this few seconds of pouring of the casting.

Recent work observing the liquid metal as it travels through the filling system indicates that most of the damage is done to castings by poor filling system design. It is also worth reflecting on the fact that every gram of metal in the casting has, of necessity, travelled through the filling system. Leaving its design to chance, or even to the patternmaker (with all due respect to all our invaluable and highly skilful pattern-makers), is a risk not to be recommended.

The early part of this section presents the design background, outlining the general thinking and some of the detailed logic behind the design of filling systems. The detailed calculations that are required to determine the precise dimensions of the various parts of the system are presented later (Section 2.3.7).

2.3.1 Gravity pouring of open-top moulds

Most castings require a mould to be formed in two parts: the bottom part (the drag) forms the base of the casting, and the top half (the cope) forms the top of the casting. However, some castings require no shaping of the top surface. In this case only a drag is required. The absence of a cope means that the mould cavity is open, so that metal can be poured directly in. The foundryman can therefore direct the flow of metal around the mould using his skill during pouring (Figure 2.2). Such open-top moulds represent a successful and economical technique




Figure 2.2 (a) An open and (b) closed mould partially sectioned.

for the production of aluminium or bronze wall plaques and plates in cast iron, which do not require a well-formed back surface. The first great engineering structure, the Iron Bridge built across the River Severn by the great English ironmaster Abraham Darby in 1779, had all its main spars cast in this way. This spectacular feat, with its main structural members over 23 m long cast in open-top sand moulds, heralded the dawn of the modern concept of a structural engineering casting.

Other viscous and poorly fluid materials are cast similarly, such as hydraulic cements, concretes, and organic resins and resin/aggregate mixtures that constitute resin concretes. Molten ceramics such as liquid basalt are poured in the same way, as witnessed by the cast basalt curb stones outside the house where I once lived, that have lined the edge of the road for the last hundred years or so and whose maker's name is still as sharply defined as the day it was cast.

The remainder of this section concentrates on the complex problem of designing filling systems for castings in which all the surfaces are moulded, i.e. the mould is closed. In all such circumstances, a bottom-gated system is adopted (i.e. the melt enters the mould cavity from one or more gates located at the lowest point, or if more than one low point, at each lowest point).

2.3.2 Gravity pouring of closed moulds

The series of funnels, pipes and channels to guide the metal from the ladle into the mould constitutes our liquid metal plumbing, and is known as the filling system, or running system. Its design is crucial; so crucial, that this is without doubt the most important chapter in the book.

However, the reader needs to keep in mind that the elimination of a running system by simply pouring into the top of the mould (down an open feeder, for instance) may be a reasonable solution in certain cases. Although apparently counter to much of the teaching in this book, there is no doubt that a top-poured option has often been demonstrated to be preferable to some poorly designed running systems, especially poorly designed bottom-gated systems. There are fundamental reasons for this that are worth examining right away.

In top gating the plunge of a jet into a liquid is accompanied by relatively low shear forces in the liquid, since the liquid surrounding the jet will move with the jet, reducing the shearing action. Thus although some damage is always done by top pouring, in some circumstances it may not be too bad, and may be preferable to a costly, difficult, or poor bottom-gated system.

In poor filling system designs, velocities in the channels can be significantly higher than the free-fall velocities. What is worse, the walls of the channels are stationary, and so maximize the shearing action, encouraging surface turbulence and the consequential damage from the shredding and entraining of bubbles and bifilms.

Ultimately, however, a bottom-gated system, if designed well, has the greatest potential for success.

Most castings are made by pouring the liquid metal into the opening of the running system, using the action of gravity to effect the filling action of the mould. This is a simple and quick way to make a casting. Thus gravity sand casting and gravity die-casting (permanent mould casting in the USA) are important casting processes at the present time. Gravity castings have, however, gained a poor reputation for reliability and quality, simply because their running systems have in general been badly designed. Surface turbulence has led to porosity and cracks, and unreliability in leak-tightness and mechanical properties.

Nevertheless, there are rules for the design of gravity-running systems that, although admittedly far from perfect, are much better than nothing. Such rules were originally empirical, based on transparent-model work and some confirmatory tests on real castings. We are now a little better informed by access to real-time video radiography of moulds during filling, and sophisticated computer simulation, so that liquid aluminium or liquid steel can be observed as it tumbles through the mould. Despite this, many uncertainties still remain. The rules for the design of filling systems are still not the mature science that we all might wish for. Even so, some rules are now evident, and their intelligent use allows castings of the highest quality to be made. They are therefore described in this section, and constitute essential reading!

It is hoped to answer the questions 'Why is the running system so complicated?' and 'Why are there so many different features?' It is a salutary fact that the apparent complexity has led to much confused thinking.

An invaluable general rule that I recommend to all those studying running and gating systems is 'If in doubt, visualize water'. Most of us have clear perceptions about the mobility and general flow behaviour of water in the gentle pouring of a cup of tea, the splat as it is spilled on the floor, the flow of a river over a weir, or the spray from a high-pressure hose pipe. A general feeling for this behaviour can sometimes allow us to cut through the mystique, and sometimes even the calculations! In addition, the application of this simple criterion can often result in the instant

dismissal of many existing filling systems intended for the production of a reliable quality of casting as being quite clearly useless!

Closed moulds represent the greatest challenge to the casting engineer. There are numerous ways to get the metal into the mould, some disastrously bad, some tolerable, some good. To appreciate the good we shall have to devote some space to the bad (Figure 2.3). If this reads like a sermon, then so be it. *The Good Running System* is a *Good Cause* that deserves the passionate concern of the casting engineer. Too many castings with hastily rigged running systems have appeared to be satisfactory in limited prototype trials, but have proved to have disastrous levels of scrap when put into production. This is normally the result of surface turbulence during filling that produces non-reproducible castings, some apparently good, some definitely bad. This result confirms the nature of turbulence. Turbulence implies chaos; and chaos implies unpredictability. When using a running system that generates surface turbulence a typical scrap rate for a commercial vehicle casting might be 15 per cent, whereas a turbine blade subjected to much more stringent inspection can easily reach 75 per cent rejections.

In general, however, experience shows that foundries that use exclusively turbulent filling methods such as most investment foundries, experience on average about 20–25 per cent scrap, of which 5–10 per cent is the total of miscellaneous minor processing problems such as broken moulds, castings damaged during cut-off, etc. The remaining 15 per cent is composed of random inclusions, random porosity, and misruns—the standard legacy of turbulent running systems: the inclusions are created by the folding of the surface, as are the random pockets of porosity; and the misruns by the unpredictable ebb and flow in different parts of the casting during filling. In sand casting foundries, most of the so-called mould problems leading to sand inclusion are actually the result of the poor filling system designs. With good filling systems sand problems such as mould erosion and sand inclusions usually disappear.

In a foundry making a variety of castings, the 15 per cent running system scrap is made up of difficult castings which might run at 85–95 per cent scrap (almost never 100 per cent!) and easy castings which run at 5 per cent scrap (almost never zero!). The non-repeatable results continuously raise the characteristic false hope that the problems are solved, only to have the hopes dashed again by the next few castings. The variability is baffling, because the foundry engineer will often go to extreme lengths to

ensure that all the variables believed to be under control are held constant.

Only a carefully worked out running system will give filling that is characterized by low surface turbulence, and which is therefore reproducible every time. Interestingly, this can mean 100 per cent scrap. However, this is not such a bad result in practice because the defect will be reproducibly repeated in every casting. It is therefore easy to identify and correct, and when corrected, stays corrected. After the first trials, the good running system should yield reliable, repeatable castings, and be characterized by a scrap rate close to zero.

A good running system, perhaps something like that shown in Figure 2.3b, will also be tolerant of wide variations in foundry practice, in contrast with the normal experience accompanying turbulent filling, in which pouring conditions are critical. Many foundries will know the problem that certain castings can only be poured successfully by certain operators. The good running system will ensure that pouring speed will now be under the control of the running system, not the pourer, and casting temperature will no longer be dictated by the avoidance of misruns, but can be set independently to control grain size without the addition of grain refiners. It is clear, therefore, that a good running system is a good ally in the creation of economical products of high quality.

The elements of a good system are:

1. Economy of size. A lightweight system will increase yield (the ratio of finished casting weight to total cast weight), allowing the foundry to make more castings from the existing melt supply. It may also help to get more castings into a given mould size. This has a big effect on productivity and economy.
2. The filling of the mould at the required speed. In the method proposed in this book, the whole running system is designed so that the velocity of the metal in the gates is below the critical value. This value varies from one alloy to another, but is generally close to 0.5 ms^{-1} . There is now much experimental and theoretical data to support this value (Runyoro 1992). Data on the density of castings produced by gating uphill have shown that air entrapment can occur above approximately 0.5 ms^{-1} (Suzuki 1989). In computer simulations of flow, Lin and Hwang (1988) show that when liquid aluminium enters the mould horizontally at 1.1 ms^{-1} it hits the far wall with such force that the reflected wave breaks, causing




Figure 2.3 (a) Poor top gates and side-fed running system, compared with (b) a more satisfactory bottom-gated and top-fed system (c) poor system gated at joint and (d) recommended economical and effective system.

surface turbulence. These figures confirm the safety of 0.5 m s^{-1} , and the danger of exceeding 1 m s^{-1} .

3. The delivery of only liquid metal into the mould cavity, i.e. not other phases such as slag, oxide, and sand. However, in most cases the overwhelmingly common and unwelcome phase is air (probably contaminated with other mould gases of course). The design of filling systems to achieve the exclusion of air will constitute a major preoccupation in this book.
4. The elimination of surface turbulence, preferably at an early stage in the runner system, but certainly by the time that the metal arrives in the mould cavity. The problem here is that by the time the metal has fallen the length of the sprue to reach the lowest level of the casting, its velocity is well above the critical velocity for surface turbulence. Despite this danger, the running system should, so far as possible, prevent the resulting fragmentation of the stream. Any fragmentation will result in permanent damage to the casting in most alloys. However, if fragmentation occurs, the best that can now happen is that it should be followed by an action to gather the stream together again. In this way the melt enters the mould as a coherent, compact spreading front, preferably at a velocity sufficiently low that the danger of any further break-up of the front is eliminated.
5. Ease of removal. Preferably the system should break off. As a next best option, it should be removable with a single stroke of a clipping press, or a straight cut. Curved cuts take more time and are more difficult to dress to finished size by grinding or finishing. Internal or shielded gates may need to be machined off, in which case the expense of setting up the casting for machining might be

avoidable by carrying out this task later, during the general machining of the casting.

(Note that in general practice it is usually best to assume that there is no requirement for the filling system to act as a feeder, i.e. to compensate for the contraction on solidification. We should ensure that the feeding function if necessary at all, is carried out by a separate feeder placed elsewhere, preferably high up, on the casting (Figure 2.3b). In some cases it is possible to use a running system that can also act as a feeder. These special systems should be used whenever possible. They are considered in Chapter 6. It is worth noting that in investment casting the almost universal confusion between filling and feeding systems is deeply regrettable. In this book the two functions are treated totally separately.)

Because the above list of criteria have been so difficult to meet in practice, there has been a move away from gravity casting as a result of what have been believed to be insoluble barriers to the attainment of high quality and reliability. Uphill filling, against gravity, known as counter-gravity casting (and, more colloquially and less helpfully, as low-pressure casting), has provided a solution to the elimination of surface turbulence. It has seemed to be the ultimate development of bottom gating (Figure 2.4). This development has therefore provided the impetus for the growth of low-pressure die casting, low-pressure sand casting, and various forms of counter-gravity filling of investment castings. A form of high-pressure die-casting has also been developed to take advantage of the quality benefits associated with counter-gravity filling followed by high-pressure consolidation. These different techniques of getting the metal into the mould will all be discussed later.

However, although counter-gravity filling fulfils all the above requirements, our main aim




Figure 2.4 Various direct gating systems applied to a box shaped casting. Possible filter locations are shown as dashed outlines. Note that all of the gravity systems shown here are poor: the sprue base connects directly with the ingate into the casting. All need mechanisms (not shown for clarity) to reduce the velocity of the melt.

in this section is to evaluate gravity filling, to see how far it can meet this difficult set of criteria.

Requirement 3 for good gating is important: only liquid metal should enter the casting. Thus all bubbles entrained by the surface turbulence characterizing the early part of the running system should have been eliminated by this stage. If the running system is poor, and bubbles are still present, their rise and bursting at the liquid surface in the mould violates Rule 4. This violation results in a number of problems, including bubble trails, splash defects, and the retention of the scattering of smaller bubbles that remain trapped under the oxide skin of the rising metal. These cause concentrations of medium-sized pores (0.5–5 mm diameter) at specific locations in the casting, usually at upper surfaces of the casting above the ingates.

The other point in Requirement 3, that dross or slag does not enter the mould cavity is interesting. In the production of iron castings it is normal for the runner to be placed in the cope and the gates in the drag, as is illustrated in Figure 2.3a. The thinking behind this design of system is that slag will float to the top of the runner, and thus will not enter the gates. Such thinking is at fault because it is clear that at least some of the first metal to enter the runner will fall down the first gate that it meets, taking with it not only the first slag but also air. This premature delivery of metal into the mould before the runner is full is clearly unsatisfactory. The metal has had insufficient time to settle down, to organize itself free from dross, oxide and bubbles. The fact that such systems are widely used, and are found in practice to reduce bubble defects in the casting actually reveals how poor the front end of the running system is. Clearly, bubbles are being generated throughout the pour, so the off-take of gates at the base of the runner is valuable in this case.

A more satisfactory system is illustrated in Figure 2.3b. Here the runner is in the drag and the gates in the cope. In this system the runner has to fill first before the gates are reached. Thus the metal has a short but valuable time to rid itself of bubbles and dross, most of which can be trapped in the dross trap or against the upper surface of the runner. Only a limited amount of slag or dross will be unfortunately placed to enter the gate. Provided the velocity of the metal in the gate is not too high, even this slag still has a good chance of being held against the ceiling of the gate, and thus not entering the casting. Figure 2.3d illustrates an optimum system (contrasting with 2.3c), designed to resist the entrainment of air at all stages of the system.

Statement 4 is deceptively simple. However, the requirement of no surface turbulence is so

important, and so central to the quest for good castings, that we have to consider it at length.

Texts elsewhere often refer to turbulence-free filling as laminar filling. The implication here is that turbulence as defined by Reynold's number is involved, and that the desirable criterion is that of laminar flow of the bulk. As discussed in *Castings 2003*, it is not bulk turbulence that is relevant since turbulent flow in the bulk liquid can still be accompanied by the desirable smooth flow of the surface. Our attention requires to be concentrated on the behaviour of the liquid surface. Thus provided we ensure that by 'laminar fill' we mean 'surface laminar fill', then we shall have our concepts correct, and our thinking accurate.

Requirement 4 above is clearly violated by splashing during filling. It can be seen immediately that top gating will probably therefore always introduce some defects (the exception is very thin wall castings where surface tension takes over control of surface turbulence). Figure 2.4 illustrates a poor running system where the metal enters from top or side gates that allow the metal to suffer a free fall into the mould cavity. Bottom-gated systems are always required if surface turbulence is to be eliminated.

However, although bottom gating is necessary, it is not a sufficient criterion. It is easy to design a bad bottom-gating system! In fact, it is possible to state the case more forcefully: a bad bottom-gated system is usually worse than most top-gated systems.

For instance, it is common to see bottom-gated systems proudly displayed with the base of the runner turned so that metal directly enters the mould (Figure 2.5). Such systems are




Figure 2.5 A poor filling system because of direct entry of high velocity metal into the mould cavity.

compact, and appear economical until the percentage scrap figures are inspected. The sequence of events is clear if we consider the fall of the first liquid down the length of the sprue. The high velocity of the metal on its impact at its base is not contained. The resulting splash may be likened to an explosion of high-velocity drops or jets fired like projectiles directly into the mould. The bulk of the metal follows in an untidy fashion, mixed with air and mould gases, and ricochets from the far wall, causing more surface turbulence as the rebounding wave breaks, rolling over and entraining yet more surface and more gas. The elimination of the entrained bubbles by bursting as they rise to the surface of the melt causes additional droplets to be created by splashing. It is important, therefore, to design the down-runner with care so that it will fill quickly, excluding air as quickly as possible, and to design the runner and gate to constrain the metal, avoiding any provision of room for splashing (Figure 2.6a). Further improvements might be allowable as in Figures 2.6b and 2.6c in which the fall heights down the sprue are progressively reduced, reducing velocities in the mould, by simply re-orienting the casting.

The base of the sprue should be the lowest point in the whole system: having reached here, all subsequent flow of the liquid should be uphill, displacing the air ahead in a controlled and progressive advance. So far as possible, the liquid should be slowed as it goes, experiencing as much opportunity as possible to become quiescent before entering the mould. It should finally enter the mould at a velocity less than its critical velocity for the entrainment of defects. In this way a good and reproducible casting is favoured.

2.3.2.1 Pressurized versus unpressurized

In the book *Castings 1991* the author recommended the achievement of velocity reduction by the progressive enlargement of the area of the flow channels at each stage, with the aim of progressively reducing the rate of flow. This is known as an unpressurized running system. The aim was to ensure that the gate was of a sufficient area to make a final reduction to the speed of the melt, so that it entered the mould at a speed no greater than its critical velocity. More recent research, however, has demonstrated that the enlargement of the system, by, for instance, a factor of two as the flow emerges from the exit of the sprue and enters the runner, usually fails to fill the runner. Thus the unpressurized systems unfortunately behaved poorly, entraining bubbles and oxides, because much of the system runs only partly full. The other standard




Figure 2.6 (a) An improved bottom-gated system; (b) and (c) further improved by height reductions.

criticism (but incidentally of much less importance) was that unpressurized systems are heavy, thus reducing metallic yield, and thus costly.

In fact, video radiography reveals that at the abrupt increase in cross-section at the base of the sprue on the entry to the runner, the entrainment of air occurs with dramatic effectiveness. This is because the melt jets along the base of the runner (not filling the additional area provided) and hits the end of the runner.




Figure 2.7 The mode of filling of (a) a pressurized system, showing the jet into the mould cavity; (b) an unpressurized system, showing the fast underjet, and the rolling back wave in the oversized runner (c, d, e) X-ray video frames of an Al alloy filling a mould 100 mm high \times 200 mm wide \times 20 mm deep illustrating the unpressurized system; (f) the final casting showing subsurface bubbles and internal cracks.

The reflected back wave rolls over the underlying fast jet, rolling in oxides and bubbles at the interface between the two (Figure 2.7a). The effect can be long-lived, developing into a stable *hydraulic jump*. The bubbles travel along the interface between the two opposing streams (probably because of the presence of two non-wetting oxide films separating the two flowing streams) and progress to the ingate, usually

collecting in a low pressure zone on one side of the ingate, before proceeding to swim up through the metal in the mould cavity (2.7e). Naturally, these bubbles and oxides bequeath serious permanent damage to the casting (2.7f).

The cast iron foundryman had some justification therefore to champion his own favourite *pressurized* systems. For the benefit of the reader, the so-called *pressurized running system*

is one in which the metal flow is choked (i.e. limited by constriction) at the gate; i.e. its rate of flow into the mould is controlled by the area of the gate, the last point in the running system (Figure 2.7b). This causes the running system to back-fill from this point, and become pressurized with liquid, forcing the system to fill and exclude air. Thus the system entrains fewer bubbles and oxides. However, it also forced the metal into the mould as a jet. Clearly this system violates one of our principal rules, since the metal is now entering the mould above its critical speed. The resulting splashing and other forms of surface turbulence inside the mould introduces its own spectrum of problems, different from those of the unpressurized system, but usually harming both the quality of the mould and the casting.

Thus neither the unpressurized nor the pressurized traditional systems are seen to work satisfactorily. This is a regrettable appraisal of present casting technology.

Because for many years the pressurized systems were mainly used for cast iron, there were special reasons why the systems appeared to be adequate:

1. In the days of pouring grey iron into greensand moulds the problems of surface turbulence were minimized by the tolerance of the metal-mould system. The oxidizing environment in the greensand mould produced a liquid silicate film on the surface of the liquid iron. Thus when this was turbulently entrained it did not lead to a permanent defect (*Castings* (2003)). In fact, many good castings were produced by tipping the metal into the top of the mould, using no running system at all! Nowadays, with the use of certain core binders and mould additives that cause solid graphitic surface films on the metal, and consequently reduce its tolerance to surface turbulence, the pressurized systems are producing defects where once they were working satisfactorily. This problem has become more acute as it has become increasingly common for irons to have alloy additions such as magnesium (to make ductile iron) and chromium (for many alloyed irons).
2. Over recent years the standards required of castings have risen to an extent that the traditional foundryman is shocked and dazed. Whereas the pressurized system was at one time satisfactory, it now needs to be reviewed. The achievement of quality is now being seen to be not by inspection, but by process control. Turbulence during filling introduces a factor that will never be predictable or controllable. This ultimately will

be seen as unacceptable. Reproducibility of the casting process will be guaranteed only by systems that fill the mould cavity with laminar surface flow. At one time this was achievable only with counter-gravity filling systems. Nowadays, as we shall see, we can achieve some success with gravity systems, provided they are designed correctly.

The conclusion given by the author in *Castings* (1991) was 'Unpressurized systems are recommended therefore. Pressurized are not.' This bold statement now requires revision in the light of recent research since we now find that neither system is really satisfactory.

In summary, the unpressurized system had the praiseworthy aim to reduce the gate velocity to below the critical velocity. Unfortunately such systems usually run only part-full, causing damage to the castings because of entrained air bubbles and oxides. The pressurized system probably benefited greatly from its ability to fill quickly and to run full, greatly reducing the damage from bubbles. However, the high velocity of the melt as it jetted into the mould created its own contribution to havoc.

Turning now to another sacred cow of running system design that requires to be addressed. This is the concept of a choke. The choke is a local constriction designed to limit flow. In the non-pressurized system the choke was generally at the base of the sprue, whereas the pressurized system was choked at the ingates into the mould. Unfortunately, a choke is an undesirable feature. Flow rates are usually sufficiently high that the melt will be speeded up through a constriction and emerge as a jet, entraining air once again downstream, with much consequential damage.

All these systems were devised before the benefits of computer simulation and video X-ray radiography. They also pre-dated the development of the concepts of surface turbulence, critical velocity, critical fall height and bifilms. It is not surprising therefore that all these traditional approaches to the design of filling systems gave less than satisfactory results.

In the history of the development of filling systems most of the early work was of limited value because the emphasis was on steady state flow through fully filled pipework, following the principles of hydraulics. This does, of course, sometimes occur late during the filling process. However, the real problems of filling are associated with the priming of the filling system, i.e. its behaviour before the filling system is filled. Thus these early studies give us relatively little useful background on which to base effective designs for real castings.

A completely new approach is described in this book that attempts to address these issues. We shall abandon the concept of a localized choke. The whole of the length of the filling system should experience its walls in permanent contact and gently pressurized by the liquid metal. Thus, effectively, the whole length of the running system should be designed to act like a choke; a kind of continuous choke principle. In all probability, it seems that we really need uniformly pressurized systems. An alternative description might be 'naturally pressurized' systems, because the new design concept is based on designing the flow channels in the mould so as to follow the natural form that the flowing metal wishes to take.

For instance, at the base of the sprue we can define its area as unity. After the right angle bend into the runner, if the stream loses energy so that its velocity falls by 20 per cent, we can expand the channel by this amount. The runner can remain at this area of 1.2 along its length. After turning through a further right angle bend into the gate this gives a series of permissible area ratios of 1:1.2:1.4, although it will be noticed that the ingate velocity has only fallen by approximately 40 per cent from that at the sprue exit.

If the 20 per cent expansion of area after each bend is not entirely allowed (for instance if only 10 per cent expansion were provided) the stream will experience a gentle pressurization. This modest pressure against the walls of the running system will be valuable to counter any effect of bubble formation and will act to support the walls of the running system against collapse (a special problem in large running systems for large castings). Thus to be more sure of maintaining the system completely full, and slightly pressurized, a ratio of 1:1.1:1.2 or even 1:1:1 might be used.

Examples of area ratios are shown in Table 2.1.

Table 2.1 Examples of area ratios (sprue exit area : runner area : gate area)

	<i>Examples of area ratios</i>
Pressurized	1:0.8:0.6 1:1:0.8
Unpressurized	1:2:4 1:4:4
Natural	1:1.2:1.4
Slightly pressurized	1:1:1 1:1.1:1.2
With foam filter in gate	1:1:4
With speed reduction or by-pass designs	1:1:10

From the ratios it is clear that the naturally (or slightly) pressurized system is part-way between the pressurized and unpressurized systems.

However, there is a major problem with the use of these new systems that the reader may already have noticed. The naturally pressurized system has no built-in mechanism for any significant reduction in velocity of the stream. Thus the high velocity at the base of the sprue is maintained (with only minor reduction) into the mould. Thus the benefits of complete priming of the filling system to exclude air are lost once again on entering the mould cavity.

This fundamental problem alerts us to the fact that the naturally pressurized approach requires completely separate mechanisms to reduce the velocity of the melt through the ingates. The options include

- (i) the use of filters;
- (ii) the provision of specially designed runner extension systems such as flow-offs;
- (iii) a surge control system;
- (iv) the use of a vertical fan gate at the end of the runner. Additional mechanisms might be possible in the future, when properly researched, such as
- (v) the use of vortices to absorb energy while avoiding significant surface turbulence.

We shall consider all these options in detail in due course, but the reader needs to be aware that, unfortunately, at this time the use of naturally pressurized systems is in its infancy. In particular, the rules for such designs are not yet known for some features such as joining round or square section sprues to rectangular runners. Filters are not easily incorporated, nor are vortex systems fully understood.

This creates a familiar problem for the foundry person: in the real world, the casting engineer has to take decisions on how to make things, whether or not the information is available at the time to help make the best choice. Thus insofar as the rules are presently understood for the majority of castings, they are set out below, for good or for bad. I hope they assist the caster to achieve a good result. One day I hope we in the industry will all have the better answers that we need.

In the meantime, computers are starting to simulate successfully the flow of metal in filling systems. At the present time such simulations are highly computationally intensive, and therefore slow and/or not particularly accurate. It is necessary to be aware that some simulation packages are still highly inaccurate. However, time will improve this situation, to the great benefit of casting quality.




Figure 2.8 A rogues gallery of non-recommended scrap generating systems. Conical basin and sprue combinations showing (a) perhaps least damaging; (b) basin too large; (c) cup form; (d) basin too small; (e) enlarged sprue to act as a combined basin and sprue.

2.3.2.2 Design of pouring basin

The conical basin

The in-line conical basin (Figure 2.8a), used almost everywhere in the casting industry, appears to be about as bad as could be envisaged for most casting operations. It is probably responsible for the production of more casting scrap than any other single feature of the filling system. It is not recommended.

The problems with the use of the conical basin arise as a result of a number of factors:

1. The metal enters at an unknown velocity, making the estimation of the design of the remainder of the running system problematical.
2. The metal enters at high, unchecked velocity. Since the main problem with running systems is to reduce the velocity, this adds to the difficulty of reducing surface turbulence.
3. Any contaminants such as dross or slag that enter with the melt are necessarily taken directly down the sprue.
4. The device works as an air pump, concentrating air into the flow (the action is analogous to other funnel-shaped pumps in which a fast stream of fluid directed down the centre of the funnel is designed to entrain a second surrounding fluid. Good examples are steam ejectors and the vacuum suction device that can be driven from a compressed air supply). Because air is probably the single most important contaminant in running systems, this is probably the most severe disadvantage, yet is not widely appreciated to be a problem.
5. The small volume of the basin makes it difficult for the pourer to keep full (its *response time* is too short, as explained later), so that air is automatically entrained as the basin becomes partially empty from time to time during pouring. The pourer is usually unaware of this, since the aspiration of air usually takes place under the surface at the basin/sprue junction.
6. The mould cavity fills differently depending on precisely where in the basin the pourer directs the pouring stream, whether at the far side of the cone, the centre, or the near side. Thus the castings are intrinsically not reproducible.
7. This type of basin is most susceptible to the formation of a vortex, because any slight off-axis direction will tend to start a rotation of the pool. There has been much written

steel ingots were produced by a conventional arrangement that consisted of pouring the steel into a central conical cup, affixed to the top of a spider distribution system of ceramic tubes connected to the centre of the base of a group of four or six surrounding ingot moulds. Because the top of the ingot mould remained open during filling, the upwelling cascade of air bubbles in the centre of the rising metal was clear for all to see. (The bottom fill technique was designed to deliver an improved surface condition of the ingot as a result of the gentle rolling action of the liquid meniscus against the wall of the ingot mould as the metal ascended. However, the overall cleanness of the ingot would have been significantly impaired by the passage of so much air. It would have been useful to retain the benefits of the bottom-teemed ladles and yet achieve improved castings by reducing the entrainment of air into the system.)

An example that the author has witnessed many times can be quoted. Bottom-teemed

about the dire dangers of a vortex, and some basins are provided with a flat side to discourage its formation. In fact, however, this so-called disadvantage would only have substance if the vortex continued down the length of the sprue, along the runner and into the mould cavity. This is unlikely. Usually, a vortex will 'bottom out,' giving an air-free flow into the remaining runner system as will be discussed later. This imagined problem is almost certainly the least of the difficulties introduced by the conical basin.

If this long list of faults was not already damning enough, it is made even worse for a variety of reasons. A basin that is too large for the sprue entrance (Figure 2.8b) jets metal horizontally off the exposed ledge formed by the top of the mould, creating much turbulence and preventing the filling of the sprue. The problem is unseen by the caster, who, because he is keeping the basin full, imagines he is doing a good job. The cup shape of the basin (Figure 2.8c) is bad for the same reason. The basin that is too small (Figure 2.8d) has painful memories for the writer: a casting with an otherwise excellent running system was repeatedly wrecked by such a simple oversight! Again, the caster thought he was doing a good job. However, the aspirated air caused a staggering amount of bubble damage in an aluminium sump casting.

The expansion of the sprue entrance to act as a basin (Figure 2.8e) may hold the record for air entrainment (however the author has no plans to expend effort investigating this black claim). Worse still, the top of this awful device is usually not sufficiently wide that the pourer can fill it because it is too small to hit with the stream of metal without the danger of much metal splashed all over the top of the mould and surroundings. Thus this combined 'basin/sprue' necessarily runs partially empty for most of the time. Furthermore, the velocity of the melt is increased as the jet is compressed into the narrow exit from the sprue (this point is discussed in detail later). The elongated tapered basin system has been misguidedly chosen for its ease of moulding. There could hardly be a worse way to introduce metal to the mould.

For very small castings weighing only a few grams, and where the sprue is only a few millimetres diameter, there is a strong element of control of the filling of the sprue by surface tension. For such small castings the conical pouring cup probably works tolerably well. It is simple and economical, and, probably fills well enough. This is as much good as can be said about the conical basin. Probably even this is praising too highly.

Where the conical cup is filled with a hand ladle held just above the cone, the fall distance of about 50 mm above the entrance to the sprue results in a speed of entry into the sprue of approximately 1 m s^{-1} . At such speeds the basin is probably least harmful. On the other hand, where the conical cup is used to funnel metal into the running system when poured directly from a furnace, or from many automatic pouring systems, the distance of fall is usually much greater, often 200 to 500 mm. In such situations the rate of entry of the metal into the system is probably several metres per second. From the bottom-poured ladles in steel foundries the metal head is usually over 1 m giving an entry velocity of 5 m s^{-1} . This situation highlights one of the drawbacks of the conical pouring basin; it contains no mechanism to control the speed of entry of liquid.

The pouring cup needs to be kept full of metal during the whole duration of the pour. If it is allowed to empty at any stage then air and dross will enter the system. Many castings have been spoiled by a slow pour, where the pouring is carried out too slowly, allowing the stream to dribble down the sprue, or simply poured down the centre without touching the sides of the sprue, and without filling the basin at all (which is the trouble with the expanded sprue type). Alternatively, harm can be done by inattention, so that the pour is interrupted, allowing the bush to empty and air to enter the down-runner before pouring is restarted. Even so, because of the small volume of the basin, it is not easily kept full so that these dangers are a constant threat to the quality of the casting.

Unfortunately, even keeping the pouring cup full during the pour is no guarantee of good castings if the cup exit and the sprue entrance are not well matched, as we have seen above. This is the most important reason for moulding the cup and the filling system integral with the mould if possible.

Finally, even if the pour is carried out as well as possible, any witness of the filling of a conical basin will need no convincing that the high velocity of filling, aimed straight into the top of the sprue, will cause oxides and air to be carried directly into the running system, and so into the casting. For castings where quality is at a premium, or where castings are simply required to be adequate but repeatable, the conical basin is definitely not recommended.

Inert gas shroud

A shroud is the cloth draped as a traditional covering over a coffin. This sober meaning does

convey the sense in which the word is used in the foundry.

The inert gas shroud has been adopted in some steel foundries. The device is a protective shield around the metal stream issuing from a bottom-poured ladle, rather like a collar, providing an inert gas environment, usually argon. Its purpose is to reduce re-oxidation of the steel during casting.

It is difficult to believe that a user would think that the short distance between the ladle and the conical basin was influential in any substantial reduction of the oxidation of the melt. Usually, the time involved in this short journey will be probably only a few hundred milliseconds. It is not easy therefore to escape the conclusion that users were in fact tacitly acknowledging the air pump action of the conical basin. The shroud therefore encourages argon to be sucked into the cone instead of air, assuming that the rate of delivery of argon is sufficient (since such pumps usually transfer roughly equal volumes of pumping fluid and entrained fluid).

The beneficial action of an argon shroud is that the reactive gas is simply replaced by an inactive gas. Thus although volumes of bubbles will continue to be entrained with the flow, they at least do not react to produce oxides or nitrides.

In fact of course, the shroud will never be completely protective for various reasons: the gas itself will be contaminated with oxygen, water vapour and other gases and volatiles in the plumbing system that delivers the gas. More important still, the seal of the shroud around the stream cannot be made proof against leakage of air; and finally the outgassing from the mould, especially in the case of an aggregate (sand) mould, will be massive.

Even so, when used appropriately, the shroud is useful. It greatly reduces re-oxidation problems of steels during casting as demonstrated by research carried out by the Steel Founders Society of America (2000). The result emphasizes the damage done by the emulsion of steel and air bubbles that characterizes the average poorly designed casting system.

The shroud has been taken to an extreme form as a long silica tube mounted directly to the underside of a bottom-pour ladle (Harrison Steel, USA, 1999). The tube acts as a re-usable sprue, and is inserted through the top of the mould and lowered carefully, so that its exit reaches the lowest point of the filling system. The stopper is then opened. If the seal between the ladle and tube is good, the filling rate of the mould is high. If leakage of air occurs at the seal the rate of mould filling is significantly reduced,

implying the strong pumping action of the falling stream to create a vacuum in the upper part of the tube, drawing in air if it can, and thus diluting the falling stream with air. Several castings in succession can be poured from one tube. However, after the tube cools the silica fragments, and requires to be replaced. Although this solution to the protection of the metal stream from oxidation is to be admired for its ingenuity, it does appear to the author to be awkward in use. The leakage problem is always an attendant danger.

In general, the author has not opted for the shroud solution, but has preferred to put in place systems that avoid the ingestion of gases into the filling system. These various systems are described below.

Contact pouring

The attempt to exclude air during the pouring of castings is carried to its ultimate logical solution in the concept of contact pouring. In this system the metal delivery system and the mould are brought into contact so that air is effectively sealed out.

The direct contact system is of course necessary, and taken for granted in the case of counter-gravity systems, in which the mould is placed directly over a source of metal. The metal is then displaced upwards by pump or differential pressure.

In the case of gravity pouring, however, the author is only aware of one use in a foundry (VAW, now Hydro Aluminium Limited, Dilligem, Germany) casting aluminium alloy. The melt is brought to the casting station by launder (a horizontal channel). The mould is also brought up to the underside of the launder in the base of which is a nozzle closed by a stopper. When the mould is presented to and pressurized against the nozzle the stopper is opened. After the mould is filled the stopper is closed and the mould can be removed, in this particular case to be rolled immediately through 180 degrees to avoid convection and aid feeding. This system works reliably and well.

The thought of transferring the concept to steel castings, using the stopper in the base of the bottom-poured ladle to deliver directly into the mouth of a sprue is quite another matter. The engineering problems for steels are daunting at this time, but may be solved one day.

The offset basin

Another design of basin (sometimes called a bush) that has been recommended from time to time, is the offset basin (Figure 2.9a).




Figure 2.9 Offset pouring basins (a) without step (definitely not recommended); (b) sharp step (not recommended); (c) radiused step (recommended).

The floor of this basin is usually arranged to be horizontal (but sometimes sloping). The intention is that the falling stream is brought to rest prior to entering the sprue. This, unfortunately, is not true. The vertical component of flow is of course zero, but the horizontal component is practically unchecked. This sideways jet across the entrance to the sprue prevents approximately half of the sprue from filling properly, so that air is entrained once again. The horizontal component of velocity continues beneath the surface of the liquid throughout the pour, even though the basin may be filled.

There has been research using this type of basin over the years, in which the discharge coefficients from sprues have been measured and found to be in the region of 50 per cent or less. These low figures confirm that the sprue is only 50 per cent or less filled, so that the major fluid being discharged is air. The quality of any castings produced from such devices must have been lamentable.

This type of basin is definitely not recommended.

The offset step (weir) basin

The provision of a vertical step, or weir, in the basin (Figure 2.9b and c) brings the horizontal jet across the top of the sprue to a stop. It is an essential feature of a well-designed basin.

Interestingly, this basin has a long history. Sexton and Primrose described a closely similar

design (but without a well-formed step) in their textbook on ironfounding published in 1911. If this basin is really valuable (as is recommended here) the reader will be curious as to why it has been known for so long, but has been extremely unpopular in foundries, whose experience of it has been discouraging. There are several reasons for this bad experience. Sometimes the basin has been made incorrectly, neglecting the important design features listed below. However, more serious than this, it has been usual to place this excellent design of basin on a filling system that completely undoes all the benefits provided by the basin. Thus the benefits of the basin are never realized, and the basin is unjustly blamed.

Despite the revered age of this basin design, the precise function and importance of each feature of the design had not been investigated until recent computer studies by Yang and Campbell (1998). These studies make it clear that

- (i) The *offset* blind end of the basin is important in bringing the vertical downward velocity to a stop. The *offset* also avoids the direct *inline* type of basin, such as the conical basin, where the incoming liquid goes straight down the sprue, its velocity unchecked, and taking with it unwanted components such as air and dross, etc.

In older designs of this device the blind end of the basin was often moulded as a hemispherical cup. This was not helpful,

since metal could easily be returned out of the basin by the sloping sides. The flat floor and near-vertical sides of the basin were therefore significant advantages. In fact the use of sharp corners to the offset side of the basin is positively helpful to avoid metal being ejected by the basin as discussed later.

- (ii) The *step (or weir)* is essential to eliminate the fast horizontal component of flow over the top of the sprue, preventing it from filling properly. Basins without this feature commonly only approximately half fill the sprue, giving an effective so-called discharge coefficient of only approximately 0.5 (how could it be higher if the sprue is only half full?). The provision of the step yields a further bonus since it reverses the downward velocity to make an upward flow, giving some opportunity for lighter phases such as slag and bubbles to separate prior to entering the sprue. Floating debris that has separated in this way is shown schematically in Figure 2.9b, c). Again, early designs were less than ideal because the step was not vertical (Swift 1949) so that its effect was compromised. The step needs a vertical height at least equal to the height of the stream at that point to ensure that it brings the horizontal component of flow to a complete stop. Commonly, this height will be at least a few millimetres for a small casting, and might be 10 to 20 mm for a casting weighing several tonnes.
- (iii) Finally, the provision of a generous *radius* over the top of the step (Figure 2.9c), smoothing the entrance into the sprue, further aids the smooth, laminar flow of metal. Swift and co-workers (1949) illustrated this effect clearly in their water models of various basins. The effect is also confirmed by the computer study by Yang and the author (1998).

The practice of placing a boom, or dam across the top of the basin (Figure 2.10) to hold back floating debris is probably counter-productive. It is seen to interfere with the natural circulation in the basin that will automatically favour the separation of buoyant phases. A dam is not recommended.

In practice, compared to the conical type, the offset step design of basin is so easy to keep full it becomes immediately popular with both caster and quality technologist alike. And, naturally, when teamed up with a well-designed filling system, the basin can demonstrate its full potential for quality improvement of the casting.

An understandable criticism is that the basins are so voluminous that they reduce yield and are




Figure 2.10 Basin with dam (probably not helpful).

thus costly. The usual design is shown in Figure 2.11a. Clearly the yield criticism can be completely met by ensuring that the basin drains as completely as possible by arranging it to be sufficiently higher than the casting. However, of those cases where the basin has to be placed lower and will not drain, the problem is to some extent addressed by the design variant shown in Figure 2.11b. In addition to saving money, this basin works even better because it constrains the melt more effectively. It encourages the funneling of the melt into the sprue with excellent laminar directional guidance.

These offset step basins can be made as separate cores, stored, and planted on moulds, matching up with the sprue entrance when required. However, because they will be required for many different castings, and so will need to mate up with different sprue entrance diameters, there is concern about any mis-match of the basin exit and the sprue entrance. However, the problem is much less acute than mis-match of conical basins, because the speed of the falling stream at this point is considerably lower, in fact only at about its critical velocity. In these circumstances surface tension is able to bridge modest outstanding ledges without significant entrainment of the liquid surface. An overhanging ledge is probably more serious and to be avoided. Thus a selection of stored basins with excess exit diameter is to be preferred. In fact it may be preferable to arrange the bush to have its base completely removed on the sprue side. The bush will then fit practically any mould. Provided the entrance to the sprue on the top surface of the cope is nicely radiused, the metal will probably be adequately funnelled into the sprue (see Figure 2.23).




Figure 2.11 Side and plan views of offset basins (a) conventional rectangular; (b) slimmed shape to streamline flow and improve metal yield.

Ultimately, however, the author prefers to mould the basin integral with the sprue, and so avoiding the link-up and alignment problems. This is easily achieved with a vertical mould joint, but less easy, but still possible, with a horizontally jointed mould.

The basin is easier to use, and works more effectively, if its *response time* is approximately 1 second. To the author's knowledge there is no definition of response time. I therefore adopt a convenient measure as the time for the basin to empty completely if the pourer stops pouring. In practice, of course, the pourer does not usually stop pouring, so that the actual rate of change of level of the basin is usually at least double the response time as defined above. Such times are relatively leisurely, allowing the pourer to maintain a consistent level of melt in the basin. Different pourers or pouring systems may require times shorter or faster than this.

The volume of the basin V_b (m^3) to give a response time t_r (in seconds) at a pouring rate Q ($\text{m}^3 \text{s}^{-1}$) is given simply by

$$V_b = Q/t_r$$

Clearly, when $t_r = 1$ second, $V_b = Q$ when using the recommended SI units.

Offset stepped basin with a bottom-pour ladle

Ladles equipped with a nozzle in the base are common for the production of large steel

castings. The benefits are generally described to be:

- (i) the metal is delivered from beneath the surface of the melt, so avoiding the transfer of slag;
- (ii) for large castings the tipping of a ladle to effect a lip pour becomes impractical;
- (iii) the accuracy of the placing and the direction of the pour is valuable. Even so it is widely known in the trade that foundries using bottom pour ladles suffer dirtier castings than those steel foundries that use lip pour ladles. This follows as a natural consequence of the great difference in pouring speeds into the conical basin, with the consequent great difference in the rate of entrainment of air. (The use of bottom-pour ladles with an offset stepped basin at the entry to the mould has the potential to avoid this central problem. However, it is not without its own set of requirements that need to be studied carefully, as we shall see below.)

The common problem when using an offset stepped basin is that although a pourer using a lip pour ladle can continue to adjust the rate of pour to maintain the level of liquid at the required height in the basin, this is easier said than done if the melt is being supplied from a bottom-pour ladle whose rate of delivery often cannot be controlled, the stopper is either open

or closed. Any attempt to adjust the rate of delivery results in sprays of steel in all directions.

In addition to this problem, as the bottom-teemed ladle gradually empties it reduces its rate of delivery. In the case of pouring a single casting from a ladle, it is fortunate that the filling system for the casting actually requires a falling rate of delivery as the net head (the level in the basin minus the level of metal in the mould) of metal driving the flow around the filling system gradually falls to zero. Even so, it is clear that the two rates are independently changing, and may be poorly matched at times. The match of speeds might be so bad that the basin runs empty, but even well before this moment, filling conditions are expected to be bad. At a filling level beneath the designed fill level in the basin the top of the liquid will appear to be covering the entrance to the sprue, but underneath, the sprue will not be completely filled, and so will be taking down air. It is essential therefore to ensure, somehow, that the level in the basin remains at least up to its designed level. At this time the problem of satisfactorily matching speeds can only be solved in detail by computer. Most software designed to simulate the filling of castings should be able to tackle this problem. However, it is perhaps more easily solved by simply having a basin with greatly increased depth, for instance perhaps up to four times the design depth. The ladle nozzle size is then chosen to deliver at a higher rate, causing the basin to overfill its design level, and so effectively running the casting at an increased speed. This increased speed is far preferable to the danger of underfilling the basin with the consequential ingestion of air into the melt.

In general therefore, a greatly increased depth to the basin is very much to be recommended. The problem of overfilling and increased speed of running may not be as serious as it might first appear. The reason is quickly appreciated. If the rate of delivery from the ladle is 40 per cent higher (a factor of $2^{1/2}$) than the designed rate of filling of the casting, the height of metal in the pouring basin will rise to a level twice as high (provided the basin has been provided with sufficient depth of course). A basin four times the minimum height will accommodate delivery from the ladle at up to twice as fast as the running system was designed for. The increase in pressure that this provides will drive the filling system to meet the higher rate. (Notice that the narrow sprue exit is not acting as a so-called choke, illustrating how wrong this concept is.) Thus the system is, within limits, automatically self-compensating if the basin has been provided with sufficient

freeboard. It is important therefore to make sure that offset stepped basins in collaboration with a bottom poured ladle do have sufficient additional height.

The preferred option to overfill the basin in terms of height is valuable in the other common experience of using a large bottom-pour ladle to fill a succession of castings. Let us take as an example a 20 000 kg ladle that is required to pour nine castings each of 2000 kg. (The final 2000 kg in the ladle will probably be discarded because it will pour too slowly, contain too much slag and be too low in temperature; there are sometimes real problems when pouring successive castings from one ladle.) The first castings will be poured extremely rapidly because the head of metal in the ladle will be high. However, the most serious problem is that the final castings in the sequence will be poured slowly, perhaps too slowly, and so might suffer severe damage from air entrainment.

The important precaution therefore is to ensure that the final casting is still poured sufficiently quickly that the minimum height in the pouring basin is still met. This is a key requirement, and will ensure that the final casting is good. Thus all of the filling design should be based on the filling conditions for the last casting. Clearly, all the preceding castings will all be overpressurized by increased heights of metal in their pouring basins, and so will fill correspondingly faster, with correspondingly higher velocities entering the mould. This should be checked to ensure that the velocities are not so very high as to cause unacceptable damage. Usually, this approach can be made to work out well.

In some cases the first castings may have their pouring basins filled high, but the metal not yet arrived in the feeders to give a signal to the operator to stop pouring. In this case the only option is to monitor the progress of the pour by some other factor, such as precise timing, or better still, a direct read-out load cell on the overhead hoist carrying the ladle.

The matching of the speed of delivery from the ladle with the speed of flow out of the pouring basin is greatly assisted if the rate of delivery from the ladle is known. This is a complex problem dependent on the height of metal in the ladle, its diameter, and the diameter of the nozzle. The interaction of all these factors can be assessed using the nomogram provided in the Appendix.

The sharp-edged or undercut offset weir basin

In addition to the matching of the rate of flow between the ladle and the casting, there are




Figure 2.12 Offset basins for high velocity input. (a) No undercut empties spectacularly upwards (not recommended). (b) The provision of an undercut gives a basin that does not splash a drop. The shaded area can be moulded in a vertical jointed mould to further improve flow and metal yield, and prevents any risk of pouring directly down the sprue.

additional problems with the application of offset weir basins for use with bottom-poured ladles.

As we have discussed above, the velocity of the melt exiting the base of the bottom-poured ladle when the stopper is first opened is soberingly high. This is because the melt at the base of a full ladle is highly pressurized. Effectively it has fallen from the upper surface of the melt in the ladle; often as much as a metre or more. Thus the exit speed is often in the region of 4 or 5 m s^{-1} . This is so high that if this powerful jet is directed into the blind end of a step basin, the liquid metal flashes outwards over the base, hits the radii in the corners of the vertical sides, where it is turned upwards to spray all over the foundry (Figure 2.12a). Such spectacular pyrotechnic displays are not recommended; little metal enters the mould.

The small radii around the four sides of the off-axis well of the basin are extremely effective in redirecting the flow upwards and out of the basin. One solution to this problem is therefore simply the removal of the radii. The provision of sharp corners to all four sides reduces the splashing tendency to a minimum (the top of the weir step leading over to the sprue entrance should still be nicely radiused of course).

The sharp cornered basin is a useful design. However, an ultimate solution to the splashing problem is provided by a simple re-entrant undercut at the base of the basin (Figure 2.12b). (The author demonstrated such a basin in a steel foundry while foundry personnel hid behind pillars and doors. On the opening of the ladle stopper the stream gushed into the basin, but not a drop emerged. The pouring process was quiet; its intense energy tamed for the first time.

The foundry personnel emerged from their hiding places to gaze in wonder.)

The undercut is, of course, a problem for many greensand moulding operations making horizontally parted moulds. This is why the sharp-edged basin is so useful. Even so, where extreme incoming velocities are involved, an undercut edge to all four sides of the filling well of the basin may be the only solution.

The undercut may be difficult to mould, but it can be machined. The upgrading of a sprue cutter to 3-D machining unit equipped with a ball-ended high speed cutter would make short work of the basin, complete with its undercut and sprue entrance, and providing all this within the moulding cycle time. Such a unit would be an expensive sprue cutter, but would be a good investment.

The undercut is not a problem for vertically jointed moulds. Its use on machines such as Disamatics is popular and welcomed by the foundry operators. Its quiet filling is easily controlled, and there is complete absence of splashed metal (commonly seen as pools, sometimes nearly lakes, swimming around on the tops of moulds). The reduction of pouring overspill is a significant contribution to the raising of metal yield in the foundry.

The moulding of the sprue cover (Figure 2.12b) ensures that metal is never poured in error directly down the sprue, and saves a little metal, making a further small contribution to yield. (In some iron foundries, however, the design may be less good at holding back slag since there is now less volume provided for slag to accumulate.)

If the offset stepped basin is successfully maintained full, the head of metal provided by

the height of the down-runner will be steady, and the rate of flow will be controlled by the sprue. The filling rate will be no longer at the mercy of the human operator on that day. The running system will have the best chance to work in accord with the casting engineer's calculations.

Stopper

As a further sophistication of the use of the offset step basin, some foundries place a small sand core in the entrance to the sprue. The core floats only after the bush is full, and therefore ensures that only clean metal is allowed to enter the sprue. Alternatively, a wire attached to the core, or a long stopper rod lifted by hand accomplishes the same task. For a large casting the raising of the stopper will require a more ruggedly engineered solution, involving the benefit of the action of a long lever to add to the mechanical advantage and keep the operator well away from sparks and splashes. However it is achieved, the delayed opening of the down-runner is valuable in many foundry situations.

The early work on the development of filling systems at Birmingham concentrated on the use of the offset step basin. A stopper was not used because it was considered to be too much trouble. However, after about the first 12 months, as a gesture to scientific diligence, it was felt that the action of a stopper should be checked, if only once, by observing the filling of a sprue using the video X-ray radiographic unit, comparing filling conditions with and without a stopper. A stopper was placed in the sprue entrance, sealing the sprue. The metal was poured into the basin. When the basin was filled to the correct level the stopper was raised. The pouring action to keep the basin full was then continued until the mould was filled. The results were unequivocal. The use of a stopper greatly improved the filling of the sprue. It was with some resignation that the author affirmed this result. For all castings after that day, a stopper was always used.

Latimer and Read (1976) demonstrated that the use of a stopper reduced the fill time by 60 per cent. This is further proof that the system runs much fuller.

There seems little doubt therefore that, despite the inconvenience, when the best quality castings are required, a stopper is advisable. Thus the author always recommends its use for aerospace products.

In addition, the use of stoppers is particularly useful for very large castings where different levels of the filling system are activated by the

progressive opening of stoppers as the melt level rises in the mould, so bringing into action new sources of metal to raise the filling speed.

2.3.2.3 Sprue (down-runner)

The sprue has the difficult job of getting the melt down to the lowest level of the mould while introducing a minimum of defects despite the high velocity of the stream.

The fundamental problem with the design of sprues is that the length of fall down the sprue greatly exceeds the critical fall height. The height at which the critical velocity is reached corresponds to the height of the sessile drop for that liquid metal. Thus for aluminium this is about 13 mm, whereas for iron and steel it is only about 8 mm. Since sprues are typically 100 to 1000 mm long, the critical velocity is greatly exceeded. How then is it possible to prevent damage to the liquid? This question is not easily answered and illustrates the central problem to the design of filling systems that work using gravity. (Conversely, of course, counter-gravity systems can solve the problem at a stroke, which is their massive technical advantage.)

For the sprue at least, the problem is soluble. It seems that the secret of designing a good sprue is to make it as narrow as possible, so that the metal has minimal opportunity to break and entrain its surface during the fall. The concept on protecting the liquid from damage is either (i) to prevent it from going over its critical velocity, or (ii) if the critical velocity has to be exceeded, to protect it by constraining its flow in channels as narrow as possible so that it is not able to jump and splash.

Theoretically a design of the sprue can be seen to be achieved by tailoring a funnel in the mould of exactly the right size to fit around a freely falling stream of metal, carrying just the right quantity of metal per second (Figure 2.13). We call the funnel the down-runner, or sprue for short. Many old hands call it the spue, or spew (which, incidentally, does not appear to be a joke).

Most sprues are oversized. This is bad for metallic yield, and thus bad for economy. However, it is much worse for the metal quality, which is damaged in two important ways:

- (i) The sprue takes more time to fill. Air is therefore taken down with the metal, causing severe surface turbulence in the sprue. This, of course, leads to a build-up of oxide in the sprue itself, and much consequential damage downstream from oxide and entrained air. The amount of damage to the metal caused by a poor basin and sprue can




Figure 2.13 The geometry of the stream falling freely from a basin.

be quickly appreciated from the common observation of the blockage of filters. Even with good quality liquid metal, a poor basin and sprue will create so much oxide that a filter is simply overloaded. Such poor front ends to filling systems are so common that filter manufacturers give standard recommendations of how much metal a filter can be expected to take before becoming choked. However, in contrast to what the manufacturers say, with a good basin and sprue (and providing, of course, the quality of the melt is not too bad) a filter seems capable of passing indefinite quantities of liquid metal without problem.

- (ii) The free fall of the melt in an oversized sprue, together with air to oxidize away the binder in the sand, is a potent combined assault that is highly successful in destroying moulds. The hot liquid ricochets and sloshes about, its high speed and agitation punishing the mould surface with a hammering and scouring action. At the same time the pockets of air in this unsteady flow will be displaced through the sand like blasts from a blacksmith's bellows, causing the organic matter in the binder to glow, and, literally, to disappear in a puff of smoke! When the binder is burned away, reclaiming the sand back to clean, unbonded grains, the result is, of course, severe sand erosion. Figure 2.14 shows a typical result for an aluminium alloy casting in a urethane resin-bound mould. An oversized sprue is a liability.




Figure 2.14 An oversized sprue that has suffered severe erosion damage because of air entrainment during the pour. A correctly sized sprue shows a bright surface free from damage.

Conversely, if the sprue is correctly sized the metal fills quickly, excluding air before any substantial oxidation of the binder has a chance to occur. The small amount of oxygen in the surface region of the mould is used up quickly by the burning of a small percentage of binder, but further oxidation has to proceed at the rate at which new supplies of air can arrive by diffusion or convection through the body of the mould. This is, of course, slow, and is therefore not important for those parts of the mould such as the sprue, that are required to survive for only the relatively short duration of the pour. Furthermore, since the liquid metal now fills the volume of the down-runner, the oxide film forming the metal-mould interface is stationary, protecting the mould material in contact with the sprue, and transmitting the gentle pressure of the steady head of metal to keep it intact. The result is a perfectly cast sprue (Figure 2.14), free from sand erosion and oxide laps. A correctly sized sprue for an aluminium alloy casting will shine like a new pin. (But beware, an undersized sprue will too!) Figure 2.3 illustrates some examples of good and bad systems. A test of a good filling system design in any metal is how well the *running system* has cast. It should be perfectly formed.

How then is it possible to be sure that the sprue is exactly the right size? The practical method of calculating the dimensions of the sprue is explained in Section 2.3.7 'Practical calculation of the filling system'. Basically, the sprue is designed to mimic the taper that the falling stream adopts naturally as a result of its acceleration due to gravity (Figure 2.15). The shape is a hyperbola (interestingly, not a parabola as widely stated). Because most sprues




Figure 2.15 *The theoretical hyperbola shape of the falling stream, illustrating the complicating effects of the basin and sprue entrance.*

approximate the shape to a straight taper, the curved sides of the stream encourage the metal to become detached from the walls at about half-way down as shown in this figure. For modest-sized castings this (together with other errors, mainly due to the geometry and friction of the flow in the basin) is simply corrected by making the sprue entrance about 20 per cent larger in area (corresponding of course to about 10 per cent increase in diameter). Thus straight tapered sprues are commonly used, and appear to be satisfactory.

For very tall castings the straight tapered approximation to the sprue shape is definitely not satisfactory. In this case it is necessary to calculate the true diameter of the sprue at close intervals along its length. The correct form of the falling stream can then be followed with sufficient accuracy, and air entrainment during the fall can be avoided.

Using this detailed approach the author has successfully used sand sprues for very large castings (including a steel casting of about 50 000 kg and 7 m high. The sprue was assembled from a stack of tubular sand cores, accurately located by an annular stepped joint. Only

one core box was required, but the central hole required a pile of separately turned loose pieces). The conventional use of ceramic tubes for the building of filling systems for steel castings was thereby avoided, with advantage to the quality of the casting. As an interesting aside, the appearance of this sprue after being broken from the mould was at first sight disappointing. It seemed that considerable sand erosion had occurred, causing the sprue to increase in diameter by over 10 mm (about 10 per cent). On closer examination however, it became clear that no erosion had occurred, but the chromite sand had softened and been compressed, losing its air spaces between the grains to become a solid mass. It had partially softened probably as a result of the use of a silicate binder system; the silicate had probably reacted with the chromite to form a lower melting point phase. Since such a growth in diameter would necessarily have occurred by a kind of creep process, in which pressure, temperature and time would be involved, it follows that much of this expansion would have happened after the casting had filled, since pressure was then highest, the sand fully up to temperature, and more time would be




Figure 2.16 An illustration of various kinds of common junctions or misalignments of the sprue. None are recommended.

available because the time for the solidification of the sprue would be as much as ten times longer than the pouring time. Thus during the pour the sand-moulded sprue would almost certainly have retained a satisfactory shape, as corroborated by the predicted fill time being fulfilled, and the cleanness of the metal rising in the mould cavity was clearly seen to be satisfactory.

Sand-moulded filling systems for steel castings are, of course, prone to erosion if the system design is bad, and particularly if, as is usual, the system is oversized. In this case however, the sand-moulded sprue worked considerably better than the conventional ceramic tube system. However, there would be no doubt that the ceramic tubes would be excellent if they could be specifically designed and produced for the sprues for each individual steel casting. Naturally, at the present time this is not easily arranged. Even so, it may be found to be an economic option in view of the expensive sands and mould coatings required if sand is used alone. In addition, the ceramic tubes are extremely easy and quick to incorporate into a sand mould, often avoiding the problem of creating a new joint line in the mould.

The cross-section of the sprue can be round or square. Some authorities have strongly recommended square in the interests of reducing the tendency of the metal to rotate, forming a vortex, and so aspirating air. This probably was important in castings using conical pouring basins because any out-of-line pouring would induce rotation of the melt. However, the author has never seen any vortex formation with an offset step basin. The problem seems not to exist with good basin design.

In addition, of course, the vortex appears to be unjustifiably maligned. The central cone of

air will only act to introduce air to the casting if the central cone extends into the mould cavity. This is unlikely, and in its use with the vortex sprue and other benign use of vortices, the design is specifically arranged to suppress this possibility. The vortex can be a powerful friend, as we shall see.

The attempt to provide gating or feeding off various parts of the sprue at various heights is almost always a mistake, and is to be avoided. Examples are shown in Figure 2.16. Overflow from such channels can introduce metal into the mould prematurely, where it can fall, splashing, and damaging the casting and mould before the general arrival of the melt via the intended bottom gate. Even if the channels are carefully angled backwards to avoid premature filling, they then act to aspirate air into the metal stream. Thus divided sprues usually either act to let out metal or let in air. They are not easily designed. Extreme caution is recommended. Perhaps one day we shall be able to design such features with complete safety as a result of high quality computer simulation. Those days are awaited patiently.

To summarize: for ease and safety of design at this time, the sprue should be a single, smooth, nearly vertical, tapering channel, containing no connections or interruptions of any kind. The rate of filling of the mould cavity should be under the absolute control of its cross-section area. If, therefore, the casting is found in practice to be filling a little too fast or too slow, then the rate can be modified without difficulty by slight adjustment of the size of the sprue.

Significantly, it is not simply the sprue exit that requires modification in this case. If correctly designed, the whole length of the sprue acts to control the rate of flow. This is what is meant by a naturally pressurized system. We

can get the design absolutely correct for the sprue along its complete length. Although methoding engineers have been carrying out such calculations correctly for many years, somehow only the sprue exit has been considered to act as the choke. We need to take careful note of this widespread error, and perhaps take time to re-think our filling system concepts.

Turning now to a common problem with many automatic moulding units for the manufacture of horizontally parted moulds. It is regrettable that a reverse-taper sprue is usually the only practical option, flagging up a major problem with the design of nearly all of our modern automatic moulding machines. (What is worse, these units also cannot usually provide for a properly moulded basin. Despite such a basin being possible to be machined as mentioned above, production by cutting is usually never actioned.) The sprue pattern needs to be permanently fixed to the pattern plate, and therefore has to be mouldable (i.e. the mould has to be able to be withdrawn off the sprue when stripping the mould off the pattern) as seen in Figure 2.3a. In this case all is not yet lost. The top of the sprue should be designed to maintain its correct size, and the taper (now the wrong sign, remember) down the length of the sprue should be kept to a minimum. (A polished stainless steel sprue pattern can often work perfectly well with zero taper providing the stripping action is accurately square.)

Even though all precautions are taken in this way to reduce the surface turbulence to a minimum, the consequential damage to the melt by a reverse or zero tapered sprue is preferably reduced by the provision of a filter as soon as possible after the base of the sprue. The friction provided by the filter acts to hold back the flow, and thus assist the poorly shaped sprue to back-fill as completely and as quickly as possible, and so reduce the rate of damage. The filter will also act to filter out some of the damage, although it has to be realized that this filtering action is not particularly efficient. The use of filters is dealt with in detail later in Section 2.3.6.3.

We need to dwell a little longer on the importance of the use of the correct taper, so far as possible, for sprues.

The effect of too little, or even negative taper has been seen above to be detrimental to casting quality. Surely, one might expect that the opposite condition of too much taper would not be a problem, since it seems reasonable to assume that the velocity of the metal depends only on the distance of fall. However, this is not true. The head of metal in the pouring basin is the driving force experienced by the melt entering

the sprue. If the sprue tapers to match the natural taper of the falling stream the only acceleration experienced by the melt is the acceleration due to gravity. If, however, the taper of the sprue is greater than this, the melt is correspondingly speeded up as the sprue constricts its area. This extra speed is unwelcome, since the task of the filling system designer is to reduce the speed. The effect of varying taper has been studied by video X-ray techniques. In experiments in which the sprue exit area was maintained constant, a doubling of the sprue entrance area was seen to nearly double the exit speed, with the generation of additional turbulence in the runner. Three times greater entrance area led to such increased velocities in the runner that severe bubble entrainment was created (Sirrell and Campbell 1997). This is one of the reasons why the elongated basin/sprue (Figure 2.8e) is so bad.

This effect is illustrated in Figure 2.17. For the negative tapers (a) and (b) the velocity at the sprue exit is merely that due to the fall of metal. The rate of arrival (kg s^{-1}) is of course controlled by the area of the sprue top. For the correctly sized sprue (c) the velocity and rate of delivery are substantially unchanged, although it will be noticed that the whole of the length of the sprue is now contacting and controlling the stream, to the benefit of the melt quality. Those sprues with too much taper (d) and (e) continue to deliver metal at nearly the same rate (in kg s^{-1} for instance), but at much higher speed (in ms^{-1} for instance) in proportion to the reduction in area of the exit. Far from acting as an effective restraint, the narrow sprue exit merely increases problems.

These effects were studied using real-time X-ray radiography (Sirrell *et al.* 1995) to optimize the taper, measuring the time for the sprue to back-fill, and the speed of the exiting melt (Figure 2.18). This work confirmed that the long-used 20 per cent increase of the area of the sprue entrance was a valuable correction. The consequential 20 per cent increase in velocity into the runner was an acceptable penalty to ensure that the sprue primed faster and more completely despite its straight-wall approximate shape.

Thus to summarize the effect of sprue taper; the taper has to be correct (within the 20 per cent outlined above). Too little or too much taper both lead to damage of the melt.

Multiple sprues

In magnesium alloy casting the widespread use of a parallel pair of rectangular slots to act as the sprue seems to be due to the desire for the




Figure 2.17 A variety of straight tapered sprues. Too little or too much taper is bad. Only the centre taper to match the falling stream is recommended. Even this could be improved by 20 per cent additional entrance area, or better still, shaped to follow the shape of the stream.

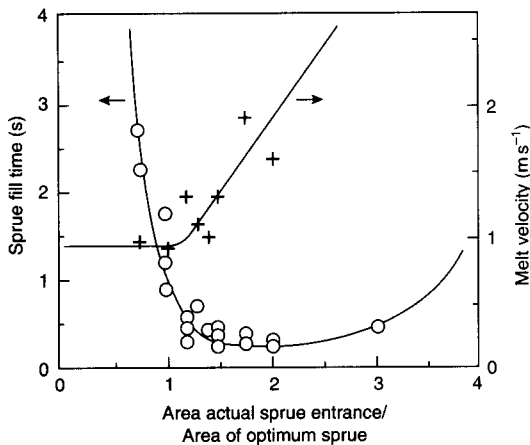


Figure 2.18 Experimental data from video radiographic observations of sprue filling time and velocity of discharge. A taper of 1.2 is shown to be close to an optimum choice (Sirrell 1995).

reduction in vortex formation (especially, as we have noted, if poor designs of pouring basin are employed). Swift *et al.* (1949) have used three parallel slots in their studies of the gating of aluminium alloys. However, the really useful benefit of a slot shape is probably associated with the reduction in stream velocity by the effect of friction on the increased surface area. The slots can be tapered to tailor their shape to that of the falling stream. However, to be strictly accurate, their area should be modified to make allowance for the additional frictional losses.

Such sprues would probably benefit the wider casting industry. A study to confirm the extent of this expected benefit would be valuable.

A really important benefit from the use of a slot sprue appears to have been widely overlooked. This is the accuracy with which it can be attached to a slot runner to give an excellent filling pattern. This benefit is described in detail in the section below concerning the design problems of the sprue/runner junction. It seems that we should perhaps be making much more regular use of slot sprues.

When pouring a large casting whose volume is greater than can be provided from the ladle, it is common to use more than one ladle. The sequential pouring of one ladle after the other into a single basin has to be carried out smoothly because any interruption to the pour is almost certain to create defects in the casting. Simultaneous pouring is often carried out. Occasionally this can be accomplished with a single sprue, but using an enlarged pouring basin, often with a double end, either side of the sprue, allowing the ladles access from either side. Often, however, two or more sprues are used, sited at opposite ends of the mould, so as to give plenty of accessibility for ladles and cranes, and reduce the travel distance for the melt in the filling system. The correspondingly smaller area used when using more than one sprue is an advantage because they fill more easily and quickly, excluding their air more rapidly. Multiple sprues for larger castings are to be recommended and should be considered more often.

For very large castings, an interesting technique can be adopted. Several sprues can connect to runners that are arranged around the mould cavity at different heights. In the pour of a 3 m high iron casting weighing 37 000 kg described by Bromfield (1991), four sprues were arranged to exit from two pouring basins. The sprues were initially closed with graphite stoppers. The trough was first filled. The stoppers to the lowest level runner were then opened. The progress of the filling was signalled by the making of an electrical contact at a critical height of metal in the mould. In other instances witnessed by the author, the progress of filling could be observed by looking down risers or sighting holes placed on the runners. When the next level of runner was reached, announced by the bright glow of metal at the base of the sighting hole, the next level of sprues was brought into action to deliver the metal to this level of runner. In the case of the casting that was witnessed, three levels of runners were provisioned by six sprues. The technique had the great advantage that the rate of pouring did not start too fast, and then slowly decrease to zero during the course of the pour. The rate could be maintained at a more consistent level by the action of bringing in additional sprues as required. In addition, the temperature of the advancing front of the melt could also be maintained by the fresh supplies of hot metal arriving at the different levels, thus reducing the need for excessive casting temperatures to avoid misruns. Again, the significant advantages of multiple sprues are clear.

2.3.2.4 Sprue base

The point at which the falling liquid emerges from the exit of the sprue and executes a right-

angle turn along the runner requires special attention. The design of this part of the liquid metal plumbing system has received much attention by researchers over the years, but with mixed results that the reader should note with caution.

The well

One of the widely used designs for a sprue base is a *well*. This is shown in Figure 2.19a. Its general size and shape has been researched in an effort to provide optimum efficiency in the reduction of air entrainment in the runner. The final optimization was a well of double the diameter of the sprue exit and double the depth of the runner. This optimization was confirmed in an elegant study by Isawa (1993) who found that the elimination of the hundreds and thousands of bubbles that were generated initially reduced exponentially with time. The exponential relationship gave a problem to define a finite time for the elimination of bubbles because the data could not be extrapolated to zero bubbles; clearly the extrapolation predicted an infinite time! He therefore cleverly extrapolated back to the time required to arrive at the last bubble, and used 'the time to the last bubble' to compare different well designs.

However, it should be noticed that both this and all the research into wells had been carried out on water models, and all had used runners of large cross-section that were not easy to fill. The result was a well design that, at best, cleared the liquid of bubbles after about 2 seconds.

For small castings that fill in only a few seconds we have to conclude that such well designs are counter-productive. In these cases it is clear that much of the filling time will be taken up conveying highly damaged metal into the mould




Figure 2.19 A variety of sprue/runner junctions in side and plan views from poorest (a) to best (d). The offset junction at (e) forms a vortex flow along the cylindrical runner.

cavity. Thus the comforting and widely held image of the well as being a 'cushion' to soften the fall of the melt is seen to be an illusion. In reality, the well was an opportunity for the melt to churn, entraining quantities of oxide and bubble defects.

Systematic X-ray radiographic studies started in 1992 have been revealing. They have shown that in a sufficiently narrow filling channel with a good radius at the sprue/runner junction, the high surface tension of the liquid metal assists in retaining the integrity of a compact liquid front, constraining the melt. These investigative studies on dramatically narrow channels in real moulds with real metals quickly confirmed that the sprue/runner junction was best designed as a simple turn (Figure 2.19c and d), provided that the channels were of minimum area.

The studies showed that if a well of any kind was provided, the additional volume created in this way was an opportunity for additional surface turbulence, so damaging the melt. Furthermore, after the well was filled, the rotation of the liquid in the well was seen to act as a kind of ball bearing, reducing the friction on the stream at the turn. In this way the velocity in the runner was increased. These higher speeds observed out of right-angle turns provided by a well were unhelpful. For a narrow turn without a well the velocity of the metal in the runner had the benefit of additional friction from the wall, giving a small (approximately 20 per cent) but useful reduction in metal speed. Thus the conclusion that the filling systems perform better without a well seems conclusive.

On a note of caution, it is perhaps necessary to bear in mind that all this research has been conducted on rather small castings. Even so, there seems no *a priori* reason why the principles should not also apply to large products.

It is unlikely that wells will disappear from the casting scene without strong defence from their supporters. It should be borne in mind that wells may once have been appropriate where large section runners were used.

In summary, despite what was recommended by the author in *Castings 1991*, more recent research confirms that wells are no longer recommended, particularly for narrow section filling systems.

The radius of the turn

It has been shown that for small castings, generally up to a few kilograms in weight, the melt can be turned through the right angle at the base of the sprue simply by putting a right-angle bend into the channel. However, if no radius is




Figure 2.20 The vena contracta problem at a right angle with inadequate radius.

provided, the melt cannot follow the bend, so that a *vena contracta* is created (Figures 2.7a and 2.20). The trailing edge of this cavitated region is unstable, so that its fluttering and flapping action sheds bubbles into the stream.

The *vena contracta* is a widely observed phenomenon in flowing liquids. It occurs wherever a rapid flow is caused to turn through a sharp change of direction. An important example has already been met in the offset pouring basin if no step is provided (Figure 2.9a). This creates a *vena contracta* that showers bubbles down the sprue. However, the base of the down-runner is probably an even more important example if, as is usually the case, speeds are much higher here. The loss of contact of the stream from the top of the runner immediately after the turn has been shown to be the source of much air in the metal. Experiments with water have modelled the low-pressure effect here, demonstrating the sucking of copious volumes of air into the liquid as streams and clouds of bubbles (Webster 1967). This is expected to be particularly severe for sand moulds, where the permeability will allow a good supply of air to the region of reduced pressure.

In fact, when pouring castings late at night, when the foundry is quiet, the sucking of air through into the liquid metal can be clearly heard, like bath water down the plug-hole! Such castings always reveal oxides, sand inclusions and porosity above the gates, which are the tell-tale signs of air bubbles aspirated into the running system.

In contrast, provided that the internal corner of the bend is given a sufficiently large radius, the melt will turn the corner without cavitation or turbulence (Figure 2.19c). In fact, the action of the advancing metal is like a piston in a cylinder: the air is simply pushed ahead of the

advancing front, never becoming mixed. To be effective, the radius needs to be at least equal to the diameter of the sprue exit, and possibly twice this amount. The precise radius requires further research. The action of the internal radius is improved further if the outside of the bend is also provided with a radius (Figure 2.19d).

For larger casting where surface tension becomes progressively less important, the channels are filled only by the available volume of flow. Initially, during the first critical period as the filling system is priming, there is considerable danger of significant damage to the metal.

To limit such damage it is helpful to take all steps to prime the front end of the filling system quickly. This is assisted by the use of a stopper. However, in *Castings* (1991) the author considered the use of various kinds of choke at the entrance to the runner as a possible solution to these problems. Again, recent research has not upheld these recommendations. It seems that any such constriction merely results in the jetting of the flow into the more distant expanded part of the runner.

This finding emphasizes the value of the concept of the *naturally pressurized* system. It is clearly of no use to expand the running system to fulfil some arbitrary formula of ratios, in the hope that the additional area will persuade the flow velocity to reduce. The flow will obey its own rules, and we need to design our system to follow these rules.

The use of a vortex sprue, or even simply a vortex base or vortex runner (Figure 2.19e) to the conventional sprue represent exciting and potentially important new developments in running system design. These concepts are described more fully in Section 2.3.2.12.

2.3.2.5 Runner

The runner is that part of the filling system that acts to distribute the melt horizontally around the mould, reaching distant parts of the mould cavity quickly to reduce heat loss problems.

The runner is usually necessarily horizontal because it simply follows the normal mould joint in conventionally horizontally parted moulds. In other types of moulds, particularly vertically jointed moulds, or investment moulds where there is little geometrical constraint, the runner would often benefit from being inclined uphill.

It is especially useful if the runner can be arranged under the casting, so that the runner is connected to the mould cavity by vertical gates. All the lowest parts of the mould cavity can then be reached easily this way. The technique is

normally achieved only in a three-part mould in which the joint between the cope and the drag contains the mould cavity, and the joint between the lower mould parts (the base and the drag) contains the running channels (Figure 2.21a). The three-part mould is often an expensive option. Sometimes the three-level requirement can be achieved by use of a large core (Figure 2.21b), or the distribution system can be assembled from ceramic or sand sections, and built into the mould as the moulding box is filled with sand (Figure 2.21c). These options are often worth considering, and might prove an economic investment.

More usually, however, a two-part mould requires both casting and running system to be moulded in the same joint between cope and drag. To avoid any falls in the filling system the runner has to be moulded in the drag, and the gates and casting in the cope (Figure 2.3d).

The usual practice, especially in iron and steel foundries, of moulding the casting in the drag




Figure 2.21 Bottom-gated systems achieved by (a) a three-part mould with accurately moulded running system; (b) making use of a core; and (c) a two-part mould with preformed channel sections.

(Figure 2.3a) is understandable from the point of view of minimizing the danger of run-outs. A leak at the joint, or a burst mould is a possible danger and a definite economic loss. This was an important consideration for hand-moulded greensand, where the moulds were rather weak (and was of course the reason for the use of the steel moulding box or flask). However, the placing of the mould cavity below the runner causes an uncontrolled fall into the mould cavity, creating the risk of imperfect castings. It is no longer such a danger for the dense, strong greensand moulds produced from modern automatic moulding machines, nor for the extremely rigid moulds created in chemically bonded sands. For products whose reliability needs to be guaranteed, the arrangement of the runner at the lowest level of the mould cavity, causing the metal to spread through the running system and the mould cavity only in an uphill




Figure 2.22 An external running system arranged around an automotive sump (oil pan).

direction is a challenge that needs to be met (Figure 2.22). Techniques to achieve this include the clever use of a core (Figure 2.21b) or for some hollow castings the use of central gating (Figures 2.23 and 2.24b).




Figure 2.24 Ring casting produced using (a) an external and (b) an internal filling system.




Figure 2.23 Cross-section of an internal running system for the casting of a cylinder.

Webster (1964) carried out some early exploratory experiments to determine optimum runner sizes. We can summarize his results in terms of the comparative areas of the runner/sprue exit. He found that a runner that has only the same area as the sprue exit (ratio 1) will have a metal velocity that is high. A ratio of 2 he claims is close to optimum since the runner fills rapidly and excludes air bubbles reasonably efficiently. A ratio of 3 starts to be difficult to fill; and a ratio of 4 is usually simply wasteful for most castings. Webster's work was a prophesy, foretelling the dangers of large runners that foundries have, despite all this good advice, continued to use.

For the best results, however, recent careful studies have made clear that even the expansion of the area of flow by a factor of 2 is not easy to achieve without a serious amount of surface turbulence. This is now known from video X-ray radiographic studies, and from detailed examination of the scatter of mechanical properties of castings using highly sensitive Weibull analysis.

The best that can easily be achieved without damage is merely the reduction of about 20 per cent in velocity by the friction of the sprue/runner bend, necessitating a 20 per cent increase in area of the runner as has been discussed above. Any greater expansion of the runner will cause the runner to be incompletely filled, and so permit conditions for damage.

Greater speed reductions, and thus greater opportunities for expansion of the runner occur if the number of right-angle bends is increased, since the factor of 0.8 reduction in speed is cumulative from one bend to the next. After three such bends the speed is reduced by half ($0.8 \times 0.8 \times 0.8 = 0.5$). Right-angle bends were anathema in filling system designs when large cross-sections were the norm. However, with very narrow systems, there is less room for surface turbulence. Even so, great care has to be taken. For instance video X-ray studies have confirmed that the bends operate best if their internal and external radii provide a parallel channel. The lack of an external radius can cause a reflected wave in larger channels.

One of the most effective devices to reduce the speed of flow in the runner is the use of a filter. The close spacing of the walls of its capillaries ensures a high degree of viscous drag. Flow rate can often be reduced by a factor of 4 or 5. This is a really valuable feature, and actually explains nearly all of the beneficial action of the filter (i.e. when using good quality metal in a well-designed filling system the filter does very little filtering. Its really important action in improving the quality of castings is its

reduction of velocity). The use of filters is considered later (Section 2.3.6).

There has over the years been a considerable interest in the concept of the separation of second phases in the runner. Jeancolas *et al.* (1969) carried out experiments on ferrous metals to show that at Reynold's numbers below the range 7000–12 000, suspended particles of alumina could be deposited in the runner but at values in excess of 15 000, they could not precipitate. Although these findings underline the importance of working with the minimum flow velocities wherever possible, it is quickly shown that for a steel casting of height 1 m, giving a velocity of flow of 4.5 ms^{-1} , for $\eta = 5.5 \times 10^{-3} \text{ N s m}^{-2}$, and for a runner of 80 mm square, Re is over 100 000. Thus it seems that conditions for the deposition of solid materials such as sand and refractory particles in runners will not be easily met. Even so, every cast iron foundry worker knows that slag will accumulate on the tops of runners, where it is much to be preferred than in the casting. Separation in this case happens because of the great difference in density between the slag and the metal, and because of the large size of the slag droplets. Thus there are some conditions in which a slow runner speed is valuable to assist cleaning the metal.

If there is a choice, the runner should be moulded in the lower half of the mould (the drag). As emphasized previously, this will encourage the runner to fill completely prior to rising through the gates (moulded for preference in the cope) and into the mould cavity.

The basic plan of the filling design starts to become clear: the metal arrives in some chaos at the bottom of the sprue. Here, after this initial trauma, it is gathered together once again by the integrating action of a feature such as a filter to provide some delay and back-pressure, after which it is allowed to rise steadily against gravity, filling section after section of the running system, and finally arriving in the mould in good order at a speed below the critical velocity.

It should be noted that such a logical system and its consequential orderly fill is not to be taken for granted. For instance, a usual mistake is to mould the runner in the cope. This is mainly because the gates, which are in either the drag or the cope, will inevitably start to fill and allow metal into the mould cavity before the runner is full, as is clear from Figure 2.3a. The traditional running of cast iron in this way fails to achieve its potential in its intended separation of metal and slag. This is because the first metal and its load of slag enters the gates immediately, prior to the filling of the runner, and thus prior to the chance that the slag can be trapped

against the upper surface of the runner. In short, the runner in the cope results in the violation of the fundamental 'no fall' criterion. The runner in the cope is not recommended for any type of casting—not even grey iron!

In gravity die castings the placing of the runner in the cope, and taking off gates on the die joint (Figure 2.3a), is especially bad. This is because the impermeable nature of the die prevents the escape of air and mould gases from the top of the runner. Thus the runner never properly fills. The entrapped gas floating on the surface of the metal will occasionally dislodge, as waves race backwards and forwards along the runner, and as the gases heat up and expand. Large bubbles will therefore continue to migrate through the gates from time to time throughout the pour, and possibly even afterwards. Because of their late arrival, it is likely that not only will bubble trails and splash problems occur, but also the advancing solidification front will trap whole bubbles.

This scenario is tempered if a die joint is provided along the top of the runner to allow the escape of air. Alternatively, a sand core sited above the runner can help to allow bubbles to diffuse away.

Even so, the complexity of behaviour of some filling system designs is illustrated by a runner in a gravity die, positioned in the cope, that acted to reduce the bubble damage in the casting (Figure 2.25). This result, apparently in complete contradiction to the behaviour described above, arose because of the exceptionally tall aspect ratio of the runner, which was shaped like a vertical slot. This shape retained bubbles high above the exits to the gates moulded below. In fact it seems that the reduction in bubbles into the casting by placing the gates low in this way only really resulted because of the extremely poor front end of the filling system. This was a bubble-producing design, so that almost any remedy had a chance to produce a better result.




Figure 2.25 Tall slot runner with bottom gates.

However, there is a real benefit to be noted (running systems are perversely complicated) because the gates would prime slowly as a head of metal in the runner was built up, thus avoiding any early jetting through into the mould cavity. This is a benefit not to be underestimated, and highlights the problem of generalizing for complex geometries of castings and their filling systems that can sometimes contain not just liquid metal but sometimes emulsions of slag and/or air.

The tapered runner

It is salutary to consider the case where the runner has two or more gates, and where the stepping or tapering of the runner has been unfortunately overlooked. The situation is shown in Figure 2.26a. Clearly, the momentum of the flowing liquid causes the furthest gate, number 3, to be favoured. The rapid flow past the opening of gate 1 will create a reduced-pressure region in the adjacent gate at this point,




Figure 2.26 (a) An unbalanced delivery of melt into the mould as a result of an incorrect runner design; (b) a tolerably balanced system.

drawing liquid out of the casting! The flow may be either in or out of gate 2, but at such a reduced amount as to probably be negligible. In the case of a non-tapered runner it would have been best to have only gate 3.

Where more than one gate is attached to the runner, the runner needs to be reduced in cross-section as each gate is passed, as illustrated in Figure 2.26b. In the past such reductions have usually been carried out as a series of steps, producing the well-known stepped runner designs. For three ingates the runner would be reduced in section area by a step of one third the height of the runner as each gate was passed. However, real-time X-ray studies have noted how during the priming of such systems, because of the high velocity of the stream, the steps cause the flow to be deflected, leaping into the air, and ricocheting off the roof of the runner. Needless to say, the resulting flow was highly disturbed, and did not achieve its intended even distribution. It has been found that simply reducing the cross-section of the runner gradually, usually linearly, cures the deflection problem. A smooth, straight taper geometry does a reasonable job of distributing the flow evenly (Figure 2.26b).

Kotschi and Kleist (1979) allow a reduction in the runner area of just 10 per cent more than the area of the gate to give a slight pressurization bias to help to balance the filling of the gates. However, they used a highly turbulent non-pressurized system that will not have encouraged results of general applicability. In contrast, computer simulation of the narrow runners recommended in this work has shown that the last gate suffers some starvation as a result of the accumulation of friction along the length of the runner. Thus for slim systems the final gates require some additional area, not less. The author usually provides for this in an *ad hoc* way by simply extending the runner past the final gate, and providing a linear taper to this more distant point (Figure 2.26b). The taper can, of course, be provided horizontally or vertically (an important freedom of choice often forgotten).

Finally, avoid tapering the runner to zero. The thinning section adds no advantage but to provide points on which people keep stabbing themselves in the foundry. It aids safety in the workplace to stop the taper at about 5 mm section thickness.

The expanding runner

In an effort to slow the metal in its early progress in the runner, a number of methods of expanding the area of the runner have been

tried. The simple expansion of the runner at an arbitrary location along the runner is of no use at all (Figure 2.27a). The melt progresses without noticing the expansion. Even expanding the runner directly from the near side of the sprue (shown as having a square section for clarity) is not helpful (Figure 2.27b). However, expanding the runner from the far side of the sprue (Figure 2.27c) does seem to work considerably better. Even here, however, the front tends to progress in two main streams on either side of the central




Figure 2.27 Plan views of a square section sprue connected to a shallow rectangular runner showing attempts to expand the runner (a and b) that fail completely. Attempt (c) is better, but flow ricochets off the walls generates a central starved, low pressure region; (d) a slot sprue and slot runner produce a uniform flow distribution in the runner shown in (e) (recommended) and (f) (probably acceptable).

axis of the runner, leaving the centre empty, or relatively empty, forming a low-pressure region some distance down-stream in the runner. This development of this double jet flow seems to be the result of the attempted radial expansion of the flow as it impacts on the runner, but finds itself constrained by and reflected from the walls of the runner. This situation for high-temperature liquids such as irons and steels leads to the downward collapse of the centre of the runner in sand moulds, since this becomes heated by radiation, and so expands, but is unsupported by the pressure of metal. The closing down of the runner in this way can be avoided by a central moulded support, effectively separating the runner into two separate, parallel runners. In practice I find that a slot runner about 100 mm wide for irons and steels is close to the maximum that can resist collapse.

A further pitfall for the unwary is the possible constricting effect that sometimes occurs as a result of attempting to connect a round or square section sprue on to a thin flat runner (Figure 2.28). For instance, if the runner were paper thin the constriction at the exit of the sprue would be nearly total; only a fraction of the flow would be able to squeeze into the narrow runner. To eliminate a constriction at this point the runner may need to be thickened, or, preferably, the fillet radius at the bend may require to be increased.

Even so, ultimately, it may come as some surprise to the reader to learn that the linking of a round or square section sprue to a slot runner, especially when attempting an expansion of the runner to reduce the velocity of the liquid, is not yet developed. To the author's knowledge, techniques for the satisfactory design of this junction do not yet exist. Some limited expansion in a horizontal plane might be achievable as indicated in Figure 2.28, but should probably be accompanied by at least a partial corresponding reduction in the vertical plane (not shown in Figure 2.28). The reduction in velocity would benefit from the friction provided by the extra surface area, but would probably not be successful to fill an expansion of a factor of 2. Thus the effect is of limited value. More research is required to evaluate what can be achieved by careful runner design.

What seems more certain, is that the distribution of flow would be simpler if a narrow slot sprue were simply to turn to link onto a horizontal slot type of runner (Figure 2.27d and e). The more uniform action of friction might assist better to achieve a modest expansion and corresponding speed reduction. This has yet to be tested. Even so, the use of slot sprues linked to slot runners promises to be a complete solution




Figure 2.28 Potential constriction to flow at the sprue to runner junction.

to the problem of the sprue/runner junction and deserves wider exploration.

2.3.2.6 Gates

Siting

When setting out the requirements for the site of a good gate, it is usual to start with the questions

'Where can we get the gate on?'

and

'Where can we get the gate off?'

Other practical considerations include

'Gate on a straight side if possible'

and

'Locate at the shortest flow distance to the key parts of the casting'.

This is a good start, but, of course, just the start. There are many other aspects to the design of a good gate.

Direct and indirect

In general, it is important that the liquid metal flows through the gates at a speed lower than the critical velocity so as to enter the mould cavity smoothly. If the rate of entry is too high, causing the metal to fountain or splash, then the battle for quality is probably lost. The turbulence inside the mould cavity is the most serious turbulence of all. Turbulence occurring early in the running system may or may not produce defects that find their way into the casting because many bifilms remain attached to the walls of the runners and many bubbles escape. However, any creation of defects in the mould cavity causes unavoidable damage to the casting.

One important rule therefore follows very simply:

Do not place the gate at the base of the down-runner so that the high velocity of the falling stream is redirected straight into the mould, as shown in Figures 2.3c, 2.5 and 2.7b. In effect, this *direct* gating is *too direct*. An improved, somewhat *indirect* system is shown in Figures 2.3b and 2.6, illustrating the provision of a separate runner and gate, and thus incorporating a number of right-angle changes of direction of the stream before it enters the mould. These provisions are all used to good effect in reorganizing the metal from a chaotic mix of liquid and gases into a coherent moving mass of liquid. Thus although we may not be reducing the entrainment of bifilms, we may at least be preventing bubble damage in the mould cavity.

As we have mentioned above, all of the oxides created in the early turbulence of the priming of the running system do not necessarily find their way into the mould cavity. Many appear to 'hang up' in the running system itself. This seems especially true when the oxide is strong as is known to be the case for Al alloys containing Be. In this case the film attached to the wall of

the running system resists being torn away, so that such castings enjoy greater freedom from filling defects. The wisdom of lengthening the running system, increasing friction, especially by the use of right-angle bends, adds back-pressure for improved back-filling and reduces velocity. It also provides more surface to contain and hold the oxides generated during priming.

Total area of gate(s)

A second important rule concerns the sizing of the gates. They should be provided with sufficient area to reduce the velocity of the melt to below the critical velocity of about 0.5 m s^{-1} . The concept is illustrated in Figure 2.1. Occasionally, the author has permitted himself the risk of a velocity up to 1 m s^{-1} and has usually achieved success. However, velocities above 1.2 m s^{-1} for Al alloys always seem to give problems. Velocities of 2 m s^{-1} in film-forming alloys, unless onto a core as explained below, would be expected to have consequences sufficiently serious that they could not be overlooked. With even higher velocities the problems simply increase.

Occasionally, there is a problem obtaining a sufficient size of gate to reduce the melt speed to safe levels before it enters the mould cavity. In such cases it is valuable if the gate opens at right angles onto a thin (thickness a few millimetres) wall. This is because the melt is now forced to spread sideways from the gate, and suffers no splashing problems because the section thickness of the casting is too small. As it spreads away from the gate it increases the area of the advancing front, thereby reducing its velocity. Thus by the time the melt arrives in a thicker section of the casting it is likely to be moving at a speed below critical. In a way, the technique uses the casting as an extension of the filling system.

This is a good reason for gating direct onto a core. This, once again, is contrary to conventional wisdom. In the past, gating onto a core was definitely bad because of the amount of air entrained in the flow. The air-assisted hammer action and oxidation of the binder thus led to sand erosion. With a good design of filling system, however, in which air is largely excluded, the action of the hot metal is safe. Little or no damage is done despite the high velocity of the stream, because the melt merely heats the core while exerting a steady pressure that holds the core material in place. Thus with a good filling system design, gating directly onto a core is recommended.

Returning to the usual gating problem whereby the gate opens into a large-section

casting. If the area of the gate is too small then the metal will be accelerated through, jetting into the cavity as though from a hosepipe. Figure 2.1c shows the effect. In many castings the jet speed can be so high that the metal effectively blasts its way around the mould cavity. Historically, many castings have been gated in this way. At the present time most steels and grey cast iron appear to be cast with this technique. The approach has enjoyed tolerable success while greensand moulding has been employed, but it seems certain that better castings and lower scrap rates would have been achieved with less turbulent filling. In the case of cores and moulds made with resin binders that cause graphitic films on the liquid iron, the pressurized system is usually unacceptable. The same conclusion is true for ductile irons in all types of moulds.

We may define some useful quick rules for determining the total gate area that is needed. For an Al alloy cast at 1 kg s^{-1} assuming a density of approximately 2500 kg m^{-3} and assuming that we wish the metal to enter the gate at its critical speed of approximately 0.5 m s^{-1} it means we need approximately 1000 mm^2 of gate area. The elegant way to describe this interesting ingate parameter is in the form of the units of area per mass per second; thus for instance ' $1000 \text{ mm}^2 \text{ kg}^{-1} \text{ s}^{-1}$ '.

Clearly we may pro-rata this figure in different ways. If we wished to fill the casting at twice this rate (i.e. in half the time) we would require 2000 mm^2 and so on. It can also be seen that the area is quickly adjusted if it is decided that the metal can be allowed to enter at twice the speed, thus the 1 kg s^{-1} would require only 500 mm^2 , or if directly onto a core in a thin section casting, perhaps twice the rate once again, giving only 250 mm^2 .

Allowing for the fact that denser alloys such as irons, steels and copper-based alloys, have a density approximately three times that of aluminium, but the critical velocity is slightly smaller at 0.4 m s^{-1} , the ingate parameter becomes, with sufficient precision, $500 \text{ mm}^2 \text{ kg}^{-1} \text{ s}^{-1}$.

The values of approximately $1000 \text{ mm}^2 \text{ kg}^{-1} \text{ s}^{-1}$ for light alloys and $500 \text{ mm}^2 \text{ kg}^{-1} \text{ s}^{-1}$ for dense alloys are useful parameters to commit to memory.

Gating ratio

In its progress through the running system the metal is at its highest velocity as it exits the sprue. If possible, we aim to reduce this in the runner, and further reduce as it is caused to expand once again into the gates. The aim is to reduce the velocity to below the entrainment

threshold (the 0.4 or 0.5 m s^{-1}) at the point of entry into the mould cavity.

It is worth spending some time below describing an alternative method of defining running systems which is widely used, but erroneous. It is to be noted that it is not recommended!

It has been common to describe running systems in terms of ratios based on the area of the exit of the sprue. For instance, a widely used area ratio of the sprue/runner/gates has been 1:2:4. Note that in this abbreviated notation the ratios given for both the runner and the gates refer back to the sprue, so that for a sprue exit of 1, the runner area is 2 and the total area of the gates is 4. It is clear that such ratios cannot always be appropriate, and that the real parameter that requires control is the velocity of metal entering the mould. Thus on occasions this will result in ratios of 1:5:10 and other unexpected values. The design of running systems based on ratios is therefore a mistake.

Having said this, I do allow myself to use the ratio of the area of sprue exit to the (total) area of the gates. Thus if the sprue is 200 mm tall (measured of course from the top of the metal level in the pouring basin) the velocity at its base will be close to 2 m s^{-1} . Thus a gate of four times this area will be required to get to below 0.5 m s^{-1} . (Note therefore that the old 1:2:4 and 1:4:4 ratios can be seen to be applicable only up to 200 mm sprue height. Beyond this sprue height the ratios are insufficient to reduce the speed below 0.5 m s^{-1} .)

I am often asked what about the problem that occurs when the mould cross-sectional area reduces abruptly at some higher level in the mould cavity. The rate of rise of the metal will also therefore be increased suddenly, perhaps becoming temporarily too fast, causing jetting or fountaining as the flow squeezes through the constriction. Fortunately, and perhaps surprisingly, this is extremely rare in casting geometries. In forty years dealing with thousands of castings I have difficulty recalling whether this has ever happened. The most narrow area is usually the gate, so the casting engineer can devote attention to ensuring that the critical velocity is not exceeded at this critical location, and at the location just inside the mould because of the sideways spreading flow (see below). If the velocity in these two situations is satisfactory it usually follows that the velocity is satisfactory at all other levels in the casting.

Even in a rare situation where a narrowing of the mould is severe, it would still be surprising if the critical velocity were exceeded, because the velocity of filling is at its highest at the ingate,

and usually decreases as the metal level rises, finally becoming zero when the net head is zero, as the metal reaches the top of the mould.

Once again, of course, counter-gravity filling wins outright. In principle, and usually with sufficient accuracy in practice, the velocities can be controlled at every level of filling.

Multiple gates

Premature filling problem via early gates Sutton (2002) applied Bernoulli's theorem to draw attention to the possibility that a melt travelling along a horizontal runner will partly enter vertical gates placed along the length of the runner, despite the fact that the runner may not yet have completely filled and pressurized (Figure 2.29). This arises as a result of the pressure gradient along the flow, and is proportional to the velocity of flow. In real casting conditions, the melt may rise sufficiently high in such gates that cavities attached to the gates might be partially filled with a slow dribble of upwelling metal prior to the filling of the runner, and therefore prior to the main flow up the vertical gates. These dribbles of metal in the cavity are poorly assimilated by the arrival of the main metal supply, and so usually constitute a lap defect resembling a misrun or part-filled casting.

This same effect would be expected to be even more noticeable in horizontal gates moulded in the cope, sited above a runner moulded in the drag (Figure 2.3b). The head pressure required to simply cross the parting line and start an unwanted early filling of part of the mould cavity would be relatively small, and easily exceeded.

Horizontal velocity in the mould When calculating the entry velocity of the metal through the gates, it is easy to overlook what happens to the melt once it starts to spread sideways into the mould cavity. The horizontal sideways velocity away from the gate can sometimes be high. In many castings where the ingate enters a




Figure 2.29 The partial filling of vertical ingates along the length of a runner.

vertical wall the transverse spreading speed inside the mould is higher than the speed through the gate, and causes a damaging splash as the liquid hits the far walls (Figure 2.30). We can make an estimate of this lateral velocity V_L in the following way.

The lateral travel of the melt will normally be at about the height h of a sessile drop. (In a thin wall the height of the flow might reach $2h$, reducing the problem considered below. We shall neglect this complication, and consider only the worst case.) We shall assume the section thickness t , for a symmetrical ingate, area A_i . The melt enters the ingate at the critical velocity V_C , and spreads in both directions away from the gate. Equating the volume flow rates through the gate and along the base of the casting gives

$$V_C \cdot A_i = 2 V_L \cdot h \cdot t$$

If we limit the gate velocity and the transverse velocity to the same critical velocity V_C (for instance 0.5 m s^{-1}) and adopt a gate thickness t the same as that of the casting wall, the relation simplifies to the fairly self-evident geometrical relation in terms of the length of the ingate L_i

$$L_i = 2h$$

The message from this simple formula is that if the length of the gate exceeds twice the height of the sessile drop, even if the gate velocity is below the critical velocity, the transverse velocity may still be too high, and surface turbulence will result from the impact of the transverse flow on the end walls of the mould cavity.

To be sure of meeting this condition, therefore, for aluminium alloys where $h = 13 \text{ mm}$, gates must always be less than 26 mm wide (remembering that this applies only to gates that have the same thickness as the wall of the casting). For irons and steels gates should not exceed 16 mm wide. Since we often require areas




Figure 2.30 Sideways flow inside mould cavity.

considerably greater than can be provided by such a short gate, it follows that multiple gates are required to achieve the total ingate area to bring the transverse velocity below critical.

Clearly, it is a concern that in practice, gate lengths are often longer than these limits and may be causing quality problems from this unsuspected source.

For those conditions where more length (or area) is needed than the above formula will allow, the solution is the provision of more gates. Two equally spaced gates of half the length will halve the problem, and so on. In this way the individual gates lengths can be reduced, reducing the problem correspondingly. Our relation becomes simply for N ingates of total ingate length L_i

$$L_i/N = 2h$$

Or directly giving the number of ingates N that will be required

$$N = L_i/2h$$

These considerations based on velocity through the gate or in the casting take no account of other factors that may be important in some circumstances. For instance, the number of ingates might require to be increased to (i) distribute heat more evenly throughout the mould; (ii) avoid localized hot spots as a result of junction problems (see below); and (iii) provide liquid at all the lowest points in the mould cavity to avoid waterfall effects.

Junction effect When the gates are planted on the casting they create a junction. This self-evident statement requires explanation.

Some geometries of junction create the danger of a hot spot. The result is that a shrinkage defect forms in the pocket of liquid that remains trapped here at a late stage of freezing. Thus when the gate is cut off a shrinkage cavity is revealed underneath. This defect is widely seen in foundries. In fact, it is almost certainly the reason why most traditional moulders cut such narrow gates, causing the metal to jet into the mould cavity with consequent poor results to casting quality.

The magnitude of the problem depends strongly on what kind of junction is created. Figure 2.31 shows the different kinds of junctions. An in-line junction (c) is hardly more than an extension of the wall of the casting. Very little thermal problem is to be expected here. The T-junction (a) is the most serious problem. It is discussed below. The L-junction (b) is an intermediate case and is not further discussed.




Figure 2.31 Maximum allowable gate thickness to avoid a hot spot at the junction with the casting.




Figure 2.32 Geometry of a T-shaped junction.

The reader can make his or her own allowances assuming conditions intermediate between the zero (in-line junction) and T-junction cases.

To help to solve this problem it is instructive to examine the freezing patterns of T-sections. In the 1970s Kotschi and Loper carried out some admirable theoretical studies of T-junctions as shown in Figure 2.32 using only simple

calculations based on modulus. These studies pointed the way for experimental work by Hodjat and Mobley in 1984 that broadly confirmed the predictions. The data are interpreted in Figure 2.33 simply as a set of straight lines of slopes 2/1, 1/1, and 1/2. (A study of the scatter in the data shows that the predictions are not infallibly correct in the transitional areas, so that some caution is required.)

Figure 2.34 presents a simplified summary of these findings. It is clear that a gate (the upright leg of the T) of 1 : 1 geometry, i.e. a section equal to the casting section (the horizontal arm of the T), has a hot spot in the junction, and so is undesirable. In fact, Figure 2.33 makes it clear that any medium-sized gate less than twice as thick, or more than half as thick, will give a

troublesome hot spot. It is only when the gate is reduced to half or less of the casting thickness that the hot spot problem is removed. (Other lessons can be learned from the T-junction results: (1) an appendage of less than one half of the section thickness will act as a cooling fin, locally enhancing the rate of cooling in the manner of a metal chill in a sand mould; and (2) an appendage of section double that of the casting will freeze later without a hot-spot problem. This is the requirement for a feeder when planted on a plate-like casting, as will be discussed in Chapter 6.

In the case of gates forming T-junctions with the casting (Figure 2.34), the requirement to make the gate only half of the casting thickness ensures that under most circumstances no




Figure 2.33 Solidification sequence for T-shaped castings (A = arm, J = junction, L = leg). Experimental data from Hodjat and Mobley (1984).




Figure 2.34 Array of different T-junctions.

localized shrinkage defect will occur, and almost no feeding of the casting will take place through the gate.

For slot gates much less than half the section of the casting that act as cooling fins the effect can be put to good use in setting up a favourable temperature gradient in the casting, encouraging solidification from the gate towards the top feeder. Such cooling fin gates have been used to good effect in the production of aluminium and copper-based alloys because of their high thermal conductivity (Wen *et al.* 1997). (The effect is much less useful in irons and steels.) Also, the common doubt that the cooling effect would be countered by the preheating because of the flow of metal into the casting is easily demonstrated to be, in most cases, a negligible problem. The preheating occurs only for a relatively short time compared to the time of freezing of the casting, and the thin gate itself has little thermal capacity. Thus following the completion of its role as a gate it quickly cools and converts to acting as a cooling fin. Where a gate cannot conveniently be made to act as a cooling fin, the author has planted a cooling fin on the sides of the gate. (This is simple if the slot gate is on a joint line.) By this means the gate is strongly cooled, and in turn, cools the local part of the casting.

The current junction rules have been stated only in terms of thickness of section. For gates and casting sections of more complex geometry it is more convenient to extend the rules, replacing section thickness by equivalent modulus. The more general rule that is inferred is 'the gate modulus should be half or less than the local casting modulus'.

It is worth drawing attention to the fact that not all gates form T-junctions with the casting. For instance, those that are effectively merely extensions of a casting wall may clearly be continued on at the full wall thickness without any hot-spot effect (Figure 2.31c).

Gates which form an L-junction with the wall of the casting are an intermediate case (Sciama 1974), where a gate thickness of 0.75 times the thickness of the wall is the maximum allowable before a hot spot is created at the junction (Figure 2.31b).

It is possible that these simple rules may be modified to some degree if much metal flows through the gate, locally preheating this region. In the absence of quantitative guidelines on this point it is wise to provide a number of gates, well distributed over the casting to reduce such local overheating of the mould. The Cosworth system devised by the author for the ingating of cylinder heads used ten ingates, one for every bolt boss. It contrasted with the two or three

gates that had been used previously, and at least partly accounts for the immediate success of the gating design (although it did not help cut-off costs, of course!).

If the casting contains heavy sections that will require feeding then this feed metal will have to be provided from elsewhere. It is necessary to emphasize the separate roles of (i) the filling system and (ii) the feeding system. The two have quite different functions. In the author's experience attempts to feed the casting through the gate are to be welcomed if really possible; however, there are in practice many reasons why the two systems often work better when completely separate. They can then be separately optimized for their individual roles.

It is necessary to make mention of some approaches to gating that attempt to evaluate the action of running systems with gates that operate only partly filled (Davis 1977). The reader will confirm that such logic only applies if the gates empty downhill into the mould, like water spilling over a weir. This is a violation of one of our most important filling rules. Thus approaches designed for partially filled gates are not relevant to the technique recommended in this book. The placing of gates at the lowest point of the casting, and the runner below that, ensures that the runner fills completely, then the gates completely, and only then can the casting start to fill. The complete prior filling of the running system is essential; it avoids the carrying through of pockets of air as waves slop about in unfilled systems. Complete filling eliminates waves.

As in most foundrywork, curious prejudices creep into even the most logical approaches. In their otherwise praiseworthy attempt to formalize gating theory, Kotschi and Kleist (1979) omit to limit the thickness of their gates to reduce the junction hot spot, but curiously equalize the areas of the gates so as to equalize the flow into the casting. In practice, making the gates the same is rarely desirable because most castings are not uniform. For instance, a double flow rate might be required into part of the casting that is locally twice as heavy.

The design of gates may be summarized in these concluding paragraphs.

The requirement for gates to be limited to a maximum thickness naturally dictates that the gates may have to become elongated into a slot-type shape if the gate area is also required to be large. Limitations to the length of the slots to limit the lateral velocities in the mould may be required of course, dictating more than one gate as explained above. The limitation of lateral velocities in addition to ingate velocities is a vital feature.

The slot form of the gates is sometimes exasperating when designing the gating system because it frequently happens that there is not sufficient length of casting for the required length of slot! In such situations the casting engineer has to settle for the best compromise possible. In practice, the author has found that if the gate area is within a factor of 2 of the area required to give 0.5 m s^{-1} , then an aluminium alloy casting is usually satisfactory. Any further deviation would be cause for concern. Grey irons and carbon steels are somewhat more tolerant of higher ingate velocities.

As a final part of this section on gating, it is worth examining some traditional gating designs.

The touch gate The touch gate, or kiss gate, is shown in Figure 2.35a. As its name suggests, it only just makes contact between the source of metal and the mould cavity. In fact there is no gate as such at all. The casting is simply placed so as to overlap the runner. The overlap is typically 0.8 mm for brass and bronze castings (Schmidt and Jacobson 1970; Ward and Jacobs 1962) although up to 1.2 mm is used. Over 2.5 mm overlap causes the castings to be difficult to break off, negating the most important advantage. The elimination of a gate in the case described by the authors was claimed to allow between 20 and 50 per cent more castings in a mould. Furthermore, the castings are simply broken off the runner, speeding production and avoiding cut-off costs and metal losses from sawing. The broken edge is so small that for most purposes dressing by grinding is not necessary; if anything, only shot blasting is required.

A further benefit of the touch gate is that a certain amount of feeding can be carried out through the gate. This happens because (1) the gate is preheated by the flow of metal through it, and (2) the gate is so close to the runner and casting that it effectively has no separate existence of its own; its modulus is not that of a tiny slot, but some average between that of the runner and that of the local part of the casting to which it connects. Investigations of touch gate geometry have overlooked this point, with confusing results. More work is needed to assess how much feeding can actually be carried out. The result is likely to be highly sensitive to alloy type so that any study would benefit from the inclusion of short and long freezing range alloys, and high and low conductivity metals.

Ward and Jacobs report a reduced incidence of mis-run castings when using touch gating. This observation is almost certainly the result of

the beneficial effect of surface tension control in preventing the penetration of the gates before the runner is fully filled and at least partly pressurized. Only when the critical pressure to force the metal surface into a single curvature of 0.4 mm is reached (in the case of the 0.8 mm overlap) will the metal enter the mould cavity. This pressure corresponds to a head of 30 to 40 mm for copper-based alloys.

With such a thin gate, variations of only 0.1 mm in thickness have been found to change performance drastically. With the runner in one half of the mould and the casting in the other, this is clearly seen to be a problem from small variations in mismatch between the mould halves. The problem can be countered in practice by providing a small gate attached to the casting, i.e. in the same mould half as the casting cavity (Figure 2.35b), so that the gate geometry




Figure 2.35 (a) Touch gate, (b) knife gate, (c) pencil gate, (d) normal and reversed horn gates.

is fixed regardless of mismatch. This is sometimes called a knife gate.

Although it is perhaps self-evident, touch and knife gates are not viable as knock-off gates on the modern designs of accurate, thin-walled, aluminium alloy castings. This is simply because the gate has a thickness similar to the casting, so that on trying to break it off, the casting itself bends! The breaking off technique works only for strong, chunky castings, or for relatively brittle alloys.

The system was said to be unsuitable for aluminium-bronze and manganese-bronze, both of which are strong film-forming alloys (Schmidt and Jacobson 1970), although this discouraging conclusion was probably the result of the runner being usually moulded in the cope and the castings in the drag and a consequence of their poor filling system, generating quantities of oxide films that would threaten to choke gates. The unfortunate fall into the mould cavity would further damage quality, as was confirmed by Ward and Jacobs (1962). They found that uphill filling of the mould was essential to providing a casting quality that would produce a perfect cosmetic polish.

The system has been studied for a number of aluminium alloys (Askeland and Holt 1975), although the poor gating and downhill filling used in this work appears to have clouded the results. Even so, the study implies that a better quality of filling system with runner in the drag and casting impressions in the cope could be important and rewarding.

The fundamental fear that the liquid may jet through the narrow gate may be unfounded. In fact, there may actually be no jetting problem at all. This appears to be a result of the high surface tension of liquid metals. Whereas water might be expected to jet through such a narrow constriction, liquid aluminium is effectively compressed when forced in to any section less than its natural sessile drop height of 12.5 mm. The action of a melt progressing through a thin gate, equipped with an even thinner section formed by a sharp notch was observed for aluminium alloys in the author's laboratory by Cunliffe (1994). The gate was 4 mm thick and the thickness under the various notches was only 1 to 2 mm. The progress of the melt along the section was observed via a glass window from above. The metal was seen to approach, cross the notch constriction, and continue on its way without hindrance, as though the notch constriction did not exist! This can only be explained if the melt immediately re-expands to fill the channel after passing the notch. It seems the liquid meniscus, acting like a compressed, doubled-over leaf spring, immediately expands

back to fill the channel when the point of highest compression is passed.

If the surface turbulence through touch gates is tolerable, or minimal, then they deserve to be much more widely used. It would be so welcome to be able to end the drudgery of sawing castings off running systems, together with the noise and the waste. With good quality metal provided by a good front end to the filling system, and uphill filling of the mould cavity after the gate, it seems likely that this device could work well. It would probably not require much work to establish a proper design code for such a practice.

The pencil gate Many large rolls for a variety of industries are made from grey cast iron in greensand moulds. They often contain a massive proportion of grey iron chills around the roll barrel to develop the white iron wear surface of the roll. It is less common nowadays to cast rolls in *loam moulds* produced by strickling. (*Loam* is a sand mixture containing high percentages of clay and water, like a mud, which allow it to be formed by *sleeking* into place. It needs to be thoroughly dried prior to casting.) Steel rolls are similarly cast.

Where the roll is solid, it is often bottom-gated tangentially into its base. Where the roll or cylinder is hollow, it may be centrifugally cast, or it may be produced by a special kind of top gating technique using pencil gates.

Figure 2.35c represents a cross-section through a mould for a roll casting. Such a casting might weigh over 60 000 kg, and have dimensions up to 5 m diameter by 5 m face length, with a wall thickness 80 mm (Turner and Owen 1964). It is cast by pouring into an open circular runner, and the metal is metered into the mould by a series of pencil gates. The metal falls freely through the complete height of the mould cavity, gradually building up the casting. The metal-mould combination of grey iron in greensand is reasonably tolerant of surface turbulence. In addition, the heavy-section thickness gives a solidification time in excess of 30 minutes, allowing a useful time for the floating out and separation of much of the oxide entrained by splashing. The splashing is limited by the slimness of the falling streams from the narrow pencil gates.

The solidification geometry is akin to continuous casting. The slow, controlled build-up of the casting ensures that the temperature gradient is high, and thus favouring good feeding. The feeder head on top of the casting is therefore only minimal, since much of the casting will have solidified by the time the feeder is

filled. This beneficial temperature gradient is encouraged by the use of pencil gates: the narrow falling streams have limited energy and so do not disturb the pool of liquid to any great depth (a single massive stream would be a disaster for this reason).

Top gating in this fashion using pencil gates is expected to be useful only for the particular conditions of: (i) grey iron; (ii) heavy sections; and (iii) greensand or inert moulds. It is not expected to be appropriate for any metal-mould combinations in which the metal is sensitive to the entrainment of oxide films, especially in thin sections where entrained material has limited opportunity to escape.

Even so, this top pouring, although occurring in the most favourable way possible as discussed above, still results in occasional surface defect in products that are required to be nearly defect-free. The use of bottom gating via an excellent filling system, entering the mould at a tangent to centrifuge defects away from the outer surface of the roll would be expected to yield a superior product. Even vertical-axis centrifugal casting would benefit from better filling design, applying the liquid metal to the rotating mould in a less turbulent fashion. No matter what the casting method, there is no substitute for a good filling system.

The horn gate The horn gate is a device used by a traditional greensand moulder to make a quick and easy connection from the sprue into the base of the mould cavity without the need to make and fit a core or provide an additional joint line (Figure 2.35d). The horn pattern could be withdrawn by carefully easing it out of the mould, following its curved shape. Although the ingenuity of the device can be admired, in practice it cannot be recommended. It breaks one of our fundamental rules for filling system design by allowing the metal to fall downhill. In addition, there are other problems. When used with its narrow end at the mould cavity it causes jetting of the metal into the mould. This effect has been photographed using an open-top mould, revealing liquid iron emerging from the exit of the gate, and executing a graceful arc through the air, before splashing into a messy, turbulent pool at the far side of the cavity (Subcommittee T535 1960). It has occasionally been used in reverse in an attempt to reduce this problem (Figure 2.35b). However the irregular filling of the first half of the gate by the metal running downhill in an uncontrolled fashion and slopping about in the valley of the gate is similarly unsatisfactory. Furthermore, the large end junction with the casting now poses the

additional problem of a large hot spot that requires to be fed to avoid shrinkage porosity.

The horn gate might be tolerable for grey iron in greensand. Otherwise it is definitely to be avoided.



Vertical gate Sometimes it is convenient to place a vertical gate at the end of a runner. Whereas the slowing of the flow by expanding the channel was largely unsuccessful for the horizontal runner, an upward-oriented expanding fan-shaped gate can be extremely beneficial because of the aid of gravity. As always, the application is not completely straightforward. Figure 2.36a shows that if the fan gate is sited directly on top of a rectangular runner, the flow is constrained by the vertical sides of the runner, so that the liquid jets vertically, falling back to fill the fan gate from above. Figure 2.36b shows that if the expansion of the fan is started from the bottom of the runner, the flow expands nicely, filling the expanding volume and so reducing in speed before it enters the mould cavity. This result is valuable because it is one of the very few successful ways in which the speed of the metal entering the mould cavity can be reduced.

The work in the author's laboratory (Rezvani *et al.* 1999) illustrates that this form of gate produces castings of excellent reliability. Compared to conventional slot gates, the Weibull modulus of tensile test bars filled with the nicely diverging fan gate was raised nearly four times, indicating the production of castings of four times greater reliability.


Itamura and co-workers (2002) have shown by computer simulation that the limiting 0.5 m s^{-1} velocity is safe for simple vertical gates, but can be raised to 1.0 m s^{-1} if the gate is expanded as a fan. However, expansion does not continue to work at velocities of 2 m s^{-1} where the flow becomes a fountain. Similar results have been confirmed in the author's laboratory by Lai and Griffiths (2003) who used computer simulation to study the expansion of the vertical gate by the provision of a generous radius at the junction with the casting. All these desirable features involve additional cutting off and dressing costs of course.

Surge control systems The flowing of metal past the gates and into some kind of dump has been widely used to eliminate the first cold metal, diverting it away, together with any initial contamination by sand or oxide. When the dump is filled the gates can start to fill. If there is any raising of back-pressure as, for instance, the accumulation of friction along the length of the runner extension, particularly if


(a)




$t=0.40$




$t=0.50$



$t=0.60$




$t=0.65$




$t=0.75$

Figure 2.36 A vertical fan gate at the end of a runner showing the difference in flow as a result of (a) top connection, and (b) bottom connection to the runner (courtesy X. Yang and Flow-3D).


(b)




$t=0.10$




$t=0.25$




$t=0.40$



$t=0.50$



$t=0.75$



$t=1.13$

Figure 2.36 *Continued*

the runner is narrow, then the gates may start to fill earlier, before the dump is fully filled. This principle is nicely addressed by the use of the Bernoulli equation as used by Sutton (2002). The concept is developed further below.

The *by-pass* principle can be used to generate more important benefits. It can assist gaining control over the initial velocity of the melt through the gates. Usually, the first metal through the gate is a transient jet, the metal spurting through when the runner is suddenly filled. This is not a problem for castings of small height where the jet effect can be negligible, but becomes increasingly severe for those with a high head height. For a tall casting the velocity at the end of the runner is high so the momentum of the melt, shocked to an instant stop, causes the metal to explode through the far gate, and enter the cavity like a javelin. This damaging initial transient can occur despite the correct tapering of the runner, since the taper is

designed to distribute the flow evenly into the mould only after the achievement of steady state conditions.

The problem can be greatly reduced by diverting the initial flow away from the casting. The provision of an additional gate at the end of the runner, beyond the casting, and not connected with the mould cavity, is a valuable technique for the reduction of the shock of the sudden filling of the runner and the impact of metal through the gates. The design of this *flow-off* device is capable of some sophistication, and promises to be a key ingredient, particularly for large, expensive one-off castings. This introduces the concept of surge control systems.

A gate that channels the initial metal into a dump *below* the level of the runner is probably the least valuable form of this technique (Figure 2.37a). The downward facing gate will continue to fill without generating significant back-pressure, the metal merely falling into the trap,




Figure 2.37 By-pass designs showing (a) and (b) cross trap type (better than nothing, but not especially recommended); (c) non-return trap; (d) vertical runner extension for gravity deceleration; (e) and (f) surge control systems using a terminal vortex; (g) surge control system with in-line vortex with axial (central) outlet.

until the instant the trap is completely filled. At that instant, the shock of filling is then likely to create, albeit at a short time later, the spurting action into the mould that it was designed to avoid. However, by this time the gates and perhaps even the mould are likely to contain some liquid, so there is a chance that any deleterious jetting action will to some extent be suppressed. This flow-off dump has the benefit of working as a classical dross trap, of course. A taper prior to the trap can prevent the back-wave reversing debris out of the trap, since there is only room for the inflow of metal (Figure 2.37c).

An improved form of the device is easily envisaged. A gate into a dump moulded *above* from the runner has a more positive action (Figure 2.37d). It provides a gradual reduction in flow rate along the runner because it generates a gradually increasing back-pressure as it fills, building up its head height. When placed at the end of the runner, the gate acts to reduce temporarily the speed into the (real) gates by providing additional gate area, and is valuable for reducing the unwanted final filling shock by some contribution to reducing speed. A simply upturned end to the runner as a runner extension (Figure 2.37b) will help in this way, but its limited area will mean that it generates its back-pressure too quickly, so any benefit of a slow increase in speed through the gates is limited. Additional volume of the dump is an advantage to delay the build-up of pressure to fill the gate.

The economically minded casting engineer might find that some castings could be made as 'free riders' in the mould at the end of such gates. The quality may not be high, especially because of the impregnation of the aggregate mould by the momentum of the metal. Even so, the part may be good enough for some purposes, and may help to boost earnings per mould.

A more sophisticated design incorporates all the desirable features of a fully developed surge control system. It consists of extending the runner into the base of an upright circular cylinder, entering tangentially (Figure 2.37e and f). The height and diameter of the cylinder are calculated to raise the back-pressure into the gates at a steady rate (avoiding the application of the full head from the filling system) for a sufficient time to ensure that the gates and the lower part of the casting are filled. When the cylinder (a kind of vortex dump) is completely filled only then does the full pressure of the filling system come into operation to accelerate the filling of the mould cavity. The final filling of the dump may still occur with a 'bang'—the water hammer effect—announced by the shock

wave of the impact as it flashes back along the runner at the speed of sound. However, this final filling shock will be considerably reduced from that produced by the metal impacting the end of a simple closed runner.

Although the device actually controls the speed of metal through the ingate, it is not called a speed control since its role is over within the first few seconds of the pour. The name 'surge control' emphasizes its temporary nature.

An even more sophisticated variant that can be suggested is the incorporation of the surge control dump in line with the flow from the sprue (Figure 2.37g). The design of the dump as a vortex as before brings additional advantages: on arrival at the base of the sprue and turning into the runner at high speed, the speed creates a centrifugal action. This action is strongly organizing to the melt, retaining the integrity of the front rather than the chaotic splashing that would have occurred in an impact into a rectangular volume, for instance. The rotary action also centrifuges the entrained air, slag (and possibly some oxides) into the centre where they have opportunity to float if the cylinder is given sufficient height. The good quality melt is taken off from the centre of the base. The small fall down the exit of the surge cylinder is not especially harmful in this case because the rotational action assists the flow to progress with maximum friction down the walls of the exit channel. The system acts to take the first blast of high-velocity metal, gradually increasing the height in the surge cylinder. In this way a gradually increased head of metal is applied to the gates. Furthermore, of course, the metal reaching the gates should be free of air and other low density contaminants.

These surge control concepts promise to revolutionize the production of large steel castings, for which other good filling solutions are, in general, either not easy or not practical. The by-pass and surge control devices represent valuable additions to the techniques of controlling not only the initial surge through the gates, but if their action is extended, as seems possible, they can also make a valuable contribution to slowing velocities during the complete vulnerable early phase of filling.

The action of a by-pass to double as a classical dross trap is described further in Section 2.3.6.

2.3.2.7 Direct gating (from gates)

If the casting engineer has successfully designed the running system to provide bottom gating with minimal surface turbulence, then the casting will fill smoothly without the formation of

film defects. However, the battle for a quality casting may not yet be won. Other defects can lie in wait for the unwary!

For the majority of castings the gate connects directly into the mould cavity. I call this simply 'direct gating'. In most cases it is allowable, or tolerable, but it sometimes causes other problems because of the effect it has on the solidification pattern of the casting.

Flow channel structure

Consider the direct-gated vertical plate shown in Figure 2.38a. Imagine this casting being filled




Figure 2.38 (a) Direct bottom-gated vertical plate, and (b) the use of a boss to assist feeding after casting is filled.

slowly to reduce the potential for surface turbulence. If the filling rate was reduced to the point that the metal just reached the top of the mould by the time the metal had just cooled to its freezing point, then it might be expected that the top of the casting would be at its coldest, and freezing would then progress steadily down the plate, from the top to the gate. (At that time the gate would be assumed to be hot because of the preheating effect of the hot metal that would have passed through.) Nothing could be further from the truth.

In reality, the slow filling of the plate causes metal to flow sideways from the gate into the sides of the plate, cooling as it goes, and freezing near the walls. Layers of fresh hot metal would continue to arrive through the gate. The successive positions of the freezing front are shown in Figure 2.38. The final effect is a flow path kept open by the hot metal through a casting that by now has mainly solidified. Rabinovich (1969) describes these patterns of flow in thin vertical plates, calling them *jet streams*. *Flow channel* is suggested as a good name, if somewhat less dramatic. The final freezing of the flow channel is slow because of the preheated mould around the path, and so its structure is coarse and porous. The porosity will be encouraged by the enhanced gas precipitation under the conditions of slow cooling, and shrinkage may contribute if local feeding is poor because either the flow path is long or it happens to be distant from a source of feed metal.

Reducing the subsequent feeding problem in a flow channel by feeding down the channel from above, or by limited feeding uphill from below, is facilitated in thicker sections where the feeding distance is greater (see Chapter 6). Thus bottom gating into bosses can take advantage of the boss as a useful feed path (Figure 2.38b). However, this action increases the problems with slower cooling, leading to enhanced gas porosity and coarse structure.

The flow channel structure is a standard feature of castings that are filled slowly from their base. This serious limitation to structure control seems to have been largely overlooked.

Moreover, the defect is not easily recognizable. It can occasionally be seen as a region of coarse grain and fine porosity in radiographs of large plate-like parts of castings. The structure contrasts with the extensive areas of clear, defect-free regions of the plate on either side. It is possible that many so-called shrinkage problems (for which more or less fruitless attempts are made to provide a solution by extra feeders or other means) are actually residual flow channels that might be cured by changing ingate position or size, or raising fill rate. No research

appears to have been carried out to guide us out of this difficulty.

Nevertheless, in general, the problem is reduced by filling faster (if that is possible without introducing other problems).

However, even fast filling does not cure the other major problem of bottom gating, which is the adverse temperature gradient, with the coldest metal being at the top and the hottest at the bottom of the casting. Where feeders are placed at the top of the casting this thermal regime is clearly unfavourable for effective feeding. In addition, particularly when solidification is slow, the problem of convection may become important. This serious problem is considered later under Rule 7.

2.3.2.8 Indirect gating (from gates)

There is an interesting gating system that solves the major features of the flow channel problem. The problem arises because the hot metal that is required to fill the casting is gated directly into the casting and has to travel through the casting to reach all parts.

The solution is *not* to gate into the casting: a main flow path is created *outside* the casting. It is called a *riser* or *up-runner*. Metal is therefore diverted initially away from the casting, through the riser, only entering subsequently by displacement sideways from the riser as fresh supplies of hot metal arrive. The fresh supplies flood up into the top of the riser, ensuring that the riser remains hot, and that the hottest metal is delivered to the top of the mould cavity. The system is illustrated in Figure 2.39. The system has the special property that the riser and slot gate combination acts not only to fill but also to feed. (The reader will notice that the use of the term 'riser' in this book is limited to this special form of feeder which also acts as an 'up-runner', in which the metal rises up the height of the casting. It is common in the USA to refer to conventional feeders placed on the tops of castings as risers. However, this terminology is avoided here; such reservoirs of metal are called feeders, not risers, following the simple logic of using a name that describes their action perfectly, and does not get confused with other bits of plumbing such as whistlers, up which metal also rises!)

The final parts of the casting to fill in Figure 2.39 will probably also require some feeding. This is easily achieved by planting a feeder on the top of the riser, as a kind of riser extension. This retains all the benefits of the system, since its metal is hot, and hotter metal below in the riser will convect into the feeder. The disadvantages of the riser and slot gate




Figure 2.39 Riser and slot gate to both gate and feed a vertical plate from (a) its side, or (b) its centre.

system are as follows:

1. The considerable cut-off and finishing problem, since the gate often has to be sited on an exterior surface of the casting, and so requires much subsequent dressing to achieve an acceptable cosmetic finish.
2. There appears to be no method of predicting the width and thickness of the gate at the present time. Further research is required here. In the normal gate where it is required to freeze before the casting section to avoid the hot-spot problem in the junction, the

thickness of the gate is held to half of the casting thickness or less. However, in this case the gate is really equivalent to a feeder neck, through which feed metal is required to flow until the casting has solidified. Whereas a thickness of double the casting thickness would then be predicted to avoid the junction hot spot under conditions of uniform starting temperature, the preheating effect of the gate due to the flow of metal through it might mean that a gate as narrow as half of the casting section may be good enough to continue feeding effectively. There are, unfortunately, no confirmatory data on this at the present time.

It is important to caution against the use of a gate which is too narrow for a completely different reason; if filling is reasonably fast then the resistance to flow provided by a narrow slot gate will cause the riser to fill up to a high level before much metal has had a chance to fill the mould cavity. The dynamics of filling and surface tension, compounded by the presence of a strong oxide film, will together conspire to retain the liquid in the riser for as long as possible. The metal will therefore spill through the slot into the casting from an elevated level (the height of fall H_f shown in Figure 2.39c). Again, our no-fall rule is broken. It is desirable, therefore, to fill slowly, and/or to have a gate sufficiently wide to present minimal resistance to metal flow. In this way the system can work properly, with the liquid metal in the riser and casting rising substantially together. Metal will then enter the mould gently.

It is also important for the gate not to be thinner than the casting when the casting wall thickness reduces to approximately 4 mm. A gate 2 mm thick would hold back the metal because of the effect of surface tension and the surface film allowing the metal head to build up in the riser. When the metal eventually breaks through, the liquid will emerge as a jet, and fall and splash into the mould cavity. For casting sections of 4 mm or less the gate should probably be at least as thick as the wall.

In general, it seems reasonable to assume that conditions should be arranged so that the fall distance h in Figure 2.39c should be less than the height of the sessile drop. The fall will then be relatively harmless.

For thinner-section castings (for instance, less than 2 mm thickness) made under normal filling pressure, the feeding of thin-section castings can probably be neglected (as will be discussed in Chapter 6). Thus any hot-spot problems can also be disregarded, with the result that the ingate can be equal to the casting

thickness. Surface tension controls the entry through the gate and the further progress of the metal through the mould cavity, reducing the problems of surface turbulence. Fill speed can therefore be increased.

A further important point of detail in Figure 2.39 should be noted. The runner turns upward on entry to the riser, directing the flow upwards. A substantial upward step is required to ensure this upward direction to the flow. This is a similar feature to that shown in Figure 2.6 and contrasts with the poor system shown in Figure 2.5. If the provision of this step is neglected causing the base of the runner to be level with the gate and the base of the casting, metal rushing along the runner travels unchecked directly into the mould cavity. A flow path would then be set up so that the riser would receive no metal directly, only indirectly after it had circulated through the casting. The base of the casting would receive all the heated metal, and the riser would be cold. Such a flow regime clearly negates the reason for the provision of the system! Many such systems have failed through omitting this small but vital detail.

What rates are necessary to make the system work best? Again we find ourselves without firm data to give any guide. We can obtain some indication from the following considerations.

The first liquid metal to flow through the gate and along the base of the cavity travels as a stream. Being the first metal travelling over the cold surface of the mould, it is most at risk from freezing prematurely. Subsequent flow occurs over the top of this hot layer of metal, and therefore does not lose so much heat from its undersurface. Thus if we can ensure that conditions are right for the first metal to flow successfully, then all subsequent flow should be safe from early freezing. In the limiting condition where the tip of the first stream just solidifies on reaching the end of the plate, it will clearly have established the best possible temperature gradient for subsequent feeding by directional solidification back towards the riser. Subsequent layers overlying this initial metal will, of course, have slightly less beneficial temperature gradients, since they will have cooled less during their journey. Nevertheless, this will be the best that we can do with a simple filling method; further improvements will have to await the application of programmed filling by pumped systems.

Focusing our attention, therefore, on the first metal into the mould, it is clear that the problem is simply a fluidity phenomenon. We shall assume that the height of the stream corresponds to the height h of the meniscus which can be supported by surface tension (Figures 2.30).

If the distance to be run from the gate is L , and the solidification time of the metal is t_f in that section thickness x , then in the limiting condition where the metal just freezes at the limit of flow (thus generating the maximum temperature gradient for subsequent feeding):

$$t_f = L/V \quad (2.1)$$

where V is the velocity of flow (m s^{-1}) of the metal stream. The corresponding rate of flow Q (kg s^{-1}) for metal of density ρ (kg m^{-3}) is easily shown to be:

$$\begin{aligned} Q &= Vhxp \\ &= Lhxp/t_f \end{aligned} \quad (2.2)$$

At constant filling rate the time t to fill a casting of height H is given by:

$$t = t_f(H/h)$$

A 10 mm thick bar (considering the first length of melt to travel along the base of the mould) in Al-7Si alloy would be expected to freeze completely in about 40 seconds giving a flow life for the solidifying alloy of perhaps 20 seconds. The meniscus height h is approximately 12.5 mm, and so for a casting $H = 100$ mm high the pouring time would be $8 \times 20 = 160$ seconds, or nearly 3 minutes. This is a surprisingly long time.

This conclusion is not likely to be particularly accurate, but does emphasize the important point that relatively thin cast sections do not necessarily require fast filling rates to avoid premature solidification. What is important is the steady, continuous advance of the meniscus. Naturally, however, it is important not to press this conclusion too far, and the above first-order approximation to the fill time probably represents a time that might be achievable in ideal circumstances: in fact, if the rate of filling is too slow, then the rate of advance of the liquid front will become unstable for other reasons:

1. Surface film problems may cause instability in the flow of some materials, as is explained in *Castings* (2003). Film-free systems will not suffer this problem, and vacuum casting may also assist.
2. Another instability that has been little researched is the flow of the metal in a pasty mode. Flow channels revealed in the radiograph in Figure 2.40 (Runyoro and Campbell 1990) show the curious behaviour in which channels take a line of least resistance through the casting, abandoning the riser.




Figure 2.40 Radiograph of an Al-7Si-0.4Mg alloy vertical plate filled via a riser and slot gate from the left-hand side, revealing unexpected filling behaviour when cast particularly cool. Remains of thermocouples can be seen (Runyoro 1992).

They adopt the form of magma vents in the earth's crust, and form volcano-like structures at the top surface. (Additionally in these radiographs a metal-mould reaction between A356 alloy and the furan resin binder has produced many minute bubbles that have floated to decorate the upper surfaces of flow channels, revealing the outline of the last regions to remain liquid.)

2.3.2.9 Central versus external systems

Most castings have to be run via an external running system as shown in Figure 2.22. While this is satisfactory for the requirements of the running system, it is costly from the point of view of the space it occupies in the mould. This is especially noticeable in chemically bonded moulds, whose relatively high cost is, of course, directly related to their volume, and whose volume can be modified easily since the moulds are not contained in moulding boxes, i.e. they are boxless. Naturally, in this situation it would be far more desirable if the running system could somehow be incorporated inside the casting, so as to use no more sand than necessary. This ideal might be achieved in some castings by the use of direct gating in conjunction with a filter as discussed later (Section 2.3.6, Direct pour).

For castings that have an open base, however, such as open frames, cylinders or rings, an excellent compact and effective solution is possible. It is illustrated for the case of cylinder and ring castings in Figures 2.23 and 2.24. The runners radiate outwards from the sprue exit, and connect with vertical slot gates arranged as arcs around the base of the casting. Ruddle and Cibula (1957) describe a similar arrangement, but do not show how it can be moulded (with all due respect to our elder statesmen of the foundry world, their suggested arrangement looks unmouldable!), and omit the upward gates. The vertical gates are an important feature for success, introducing useful friction into the system, and making for easy cut-off.

Feeders can be sited on the top of the cylinder if required. Alternatively, if the casting is to be rolled through 180 degrees after pouring, the feeding of the casting can take the form of a ring feeder at the base (later to become the top, of course).

Experience with internal running has found it to be an effective and economical way to produce hollow shapes. It is also effective for the production of other common shapes such as gearboxes and clutch covers, where the sprue can be arranged to pass down through a rather small opening in one half of the casting and then be distributed via a spider of runners and gates on the open side.

However, it has been noted that aluminium alloy castings of 300 mm or more internal diameter exhibit a patternmaker's contraction considerably less than that which would have been expected for an external system. This seems almost certainly to be the result of the expansion of the internal core as a result of the extra heating from the internal running system. For a silica sand core this expansion can be between 1 and 1.5 per cent, effectively negating the patternmaker's shrinkage allowance, which is normally between 1 and 1.3 per cent.

2.3.2.10 Sequential filling

When there are multiple impressions on a horizontal pattern plate, it is usually unwise to attempt to fill all the cavities at the same time. (This is contrary to the situation with a vertically parted mould, in which many filling systems specifically target the filling of all the cavities at once to reduce pouring time. However, such vertically parted moulds have not been subjected to the same degree of study in terms of the defects probably introduced by this system. In the absence of data therefore, they are not described further here. We look forward

to good data becoming available at some future date.)

The reasoning in the case of the horizontal mould is simple. The individual cavities are filling at a comparatively slow rate, and not necessarily in a smooth and progressive way. In fact, despite an otherwise good running system design, it is likely that filling will be severely irregular, with slopping and surging, because of the lack of constraint on the liquid, and because of the additional tendency for the flow to be unstable at low flow rates in film-forming alloys. The result will be the non-filling of a number of the impressions and doubtful quality of the others.

Loper (1981) provided a solution to this problem for multiple impressions on one plate as shown in Figure 2.41. He uses runner dams to retard the metal, allowing it time to build up a head of metal sufficient to fill the first set of impressions before overcoming the dam and proceeding to the next set of impressions, and so on.

The system has only been reported to have been used for grey iron castings in greensand moulds. It may give less satisfactory results for other metal-mould systems that are more susceptible to surface turbulence. However, the design of the overflow (the runner down the far side of each dam) could be designed as a miniature tapered down-runner to control the fall, and so reduce surface turbulence as far as possible. Probably, this has yet to be tested.

Another sequential-filling technique, 'horizontal' stack moulding, has also only so far been used with cast iron. This was invented in the 1970s by one of our great foundry characters from the UK, Fred Hoult, after his retirement at the age of 60. It is known in his honour as the 'H Process'. Figure 2.41 outlines his method. The progress of the metal across the top of those castings already filled keeps the feeders hot, and thus efficient. The length of the stack seems unlimited because the cold metal is repeatedly being taken from the front of the stream and diverted into castings. (The reader will note an interesting analogy with the up-runner and slot gate principle; one is horizontal and the other vertical, but both are designed to divert their metal into the mould progressively. The same effect is also used in the promotion of fluidity as described by Hiratsuka *et al.* 1966.) Stacks of 20 or more moulds can easily be poured at one time. Pouring is continued until all the metal is used up, only the last casting being scrapped because of the short pour, and the remaining unfilled moulds are usable as the first moulds in the next stack to be assembled.




Figure 2.41 (a) Sequential filling for a number of impressions on a pattern plate (after Loper 1981), and (b) sequential filling for horizontal stack moulded castings (H Process).

The size of castings produced by the H Process is limited to parts weighing from a few grams to a few kilograms. Larger parts become unsuitable partly because of handling problems, since the moulds are usually stacked vertically during assembly, then clamped with long threaded steel rods, and finally lowered to the floor to make a horizontal line. Larger parts are also unsuitable because of the fundamental limitation imposed by the increase of defects as a necessary consequence of the increased distance of fall of the liquid metal inside the mould, and possibly greater opportunity to splash in thicker sections.

2.3.2.11 Two-stage filling

There have been a number of attempts over the years to introduce a two-stage filling process. The first stage consists of filling the sprue, after which a second stage of filling is started in which the runner and gates, etc. are allowed to fill.

The stopping of the filling process after the filling of the sprue brings the melt in the sprue to a stop, ensuring the exclusion of air. The melt is then allowed to start flowing once again. This second phase of filling has the full head H of

metal in the sprue and pouring basin to drive it, but the column has to start to move from zero velocity. It reaches its 'equilibrium' velocity $(2gH)^{1/2}$ only after a period of acceleration. Thus the early phase of filling of the runner and gates starts from a zero rate, and has a gradually increasing velocity. The action is similar to our 'surge control' techniques described earlier.

The benefits of the exclusion of air from the sprue, and the reduced velocity during the early part of stage 2, are benefits that have been recorded experimentally for semi-solid (actually partly solid) alloys. These materials are otherwise extremely difficult to cast without defects, almost certainly because their entrainment defects cannot float out but are trapped in suspension because of the high viscosity of the mixture.

Workers from Alcan (Cox *et al.* 1994) developed a system in which the advance of the melt was arrested at the base of the sprue by a layer of ceramic paper supported on a ceramic foam filter (Figure 2.42a). When the sprue was filled the paper was lifted from one corner by a rod, allowing the melt to flow through the filter and into the running system. These authors call their system 'interrupted pouring'. However, the




Figure 2.42 Two-stage filling techniques: (a) paper seal on top of ceramic foam filter, lifted by wire; (b) steel slide gate at entrance to runner; (c) stopper in the base of a snorkel ladle.

name 'two-stage pour' is recommended as being more positive, and less likely to be interpreted as a faulty pour as a result of an accident.

The two-stage pour was convincingly demonstrated as beneficial by Taghiabadi and colleagues (2003) for both partly solid and conventional aluminium casting alloys. These authors used Weibull statistics to confirm the reality of the benefits. They used a steel sheet to form a barrier, as a slide valve, in the runner (Figure 2.42b). After the filling of the sprue the sheet was withdrawn, allowing the mould to fill.

A second completely different incarnation of the two-stage filling concept is the *snorkel ladle*, sometimes known as the *eye-dropper ladle*. It is illustrated in Figure 2.42c. The device is used mainly in the aluminium casting industry, but would with benefit extend to other casting industries. Instead of transferring metal from a furnace via a ladle or spoon of some kind, and pouring into a pouring basin connected to a sprue, the snorkel dips into the melt, and can be filled uphill simply by dipping sufficiently deeply, or by a shallow dip and the melt sucked up by a reduced pressure applied in the body of the snorkel. The ladle is then transferred to the mould where it can deliver its contents into a conventional basin and sprue system, or, in the mode recommended here, lowered down through the mould to reach and engage with the runner. Only then is its stopper raised and the melt delivered to the start of the running system with minimal surface turbulence. The approach is capable of producing excellent products.

Two-stage filling in its various forms seems to offer real promise for many castings.

2.3.2.12 Vortex systems

The vortex has usually been regarded in foundries as a flow feature to be avoided at all costs. If the vortex truly swallowed air, and the air found its way into the casting, the vortex would certainly have to be avoided. However, in general, this seems to be not true.

The great value of the vortex is that it is a powerful organizer of the flow. Designers of water intakes for hydroelectric power stations are well aware of this benefit. Instead of the water being allowed to tumble haphazardly down the water intake from the reservoir, it is caused to spiral down the walls. At the base of the intake duct the loss of rotational energy allows the duct to back-fill to some extent. The central core of air terminates at the level surface of a comparatively tranquil pool, only gently circulating, near the base of the duct. (The spiralling central core of air does not extend indefinitely through the system.)

Several proposals to harness the benefits of vortices to running systems have originated in recent years from Birmingham, following the lead by Isawa (1994). They are potentially exciting departures from conventional approaches. Only initial results can be presented here. The systems merit much further investigation.

Vortex sprue

The benefits of the vortex for the action of a sprue were first explored by Campbell and




Figure 2.43 Vortex sprue (after Isawa and Campbell 1994).

Isawa (1994) as illustrated in Figure 2.43. An aluminium alloy was poured off-axis, being diverted tangentially into a circular pouring basin. The melt spun around the outside of the basin, gradually filling up and so progressing towards the central sprue entrance. As it progressed inwards, gradually reducing its radius, its rotation speeded up, conserving its angular momentum like the spinning ice skater who closes in outstretched arms. Finally, the melt reached the lip of the sprue, where, at maximum spinning speed, it started its downward fall.

The rationale behind this thinking is that the initial fall during the filling of the sprue is controlled by the friction of a spiral descent down the wall of the sprue. Once the sprue starts to fill, the core of the vortex terminates near the base of the sprue; it does not in fact funnel air into the mould cavity. The hydrostatic pressure generated by this system, driving the flow into the runner and gates, arises only from the small depth of liquid at the base of the vortex, so that the net pressure head h driving the filling of the mould is small. Because the level of the pool at the base of the vortex and the level of metal in the mould cavity both rise together, the net

head to drive the filling of the mould remains remarkably constant during the entire mould filling process. Thus the filling of the mould is necessarily gentle at all times.

Despite some early success with this system, it seems that the technique is probably not suitable for sprues of height greater than perhaps 200 or 300 mm, because the benefits of the spiral flow are lost progressively with increasing fall distance. More research may be needed to confirm the benefits and limitations of this design. For instance, the early work has been conducted on parallel cylindrical sprues, since the taper has been thought to be not necessary as a result of the melt adopting its own 'taper' as it accelerates down the walls, becoming a thinner stream as it progresses. However, a taper may in any case be useful to favour the speeding up of the rate of spin, and so assist maintaining the spin despite losses from friction against the walls. Also of course, the provision of a taper will assist the sprue to fill faster, and increase yield. Much work remains to be done to define an optimum system.

Vortex well or gate

The provision of a cylindrical channel at the base of the sprue, entered tangentially by the melt, is a novel idea with considerable potential (Hsu *et al.* 2003). It gives a technique for dealing with the central issue of the high liquid velocity at the base of the sprue, and the problem of turning the right-angle corner and successfully filling the runner. What is even better, it promises to solve all of these problems without significant surface turbulence.

The vortex well can probably be oriented either horizontally or vertically as seen in Figure 2.44. The horizontal orientation may be useful for delivery to a single vertical gate. Alternatively, the vertical orientation is often convenient because the central outlet from the vortex can form the entrance to the runner, allowing the connection to many gates.

Notice that the device works exactly opposite to the supposed action of a spinner designed to centrifuge buoyant inclusions from a melt. In the vortex well the outlet to the rest of the filling system is the outlet that would normally be used to concentrate inclusions. Thus the device certainly does not operate to reduce the inclusion content. However, it should be highly effective in reducing the generation of inclusions by surface turbulence at the sprue base of poorly designed systems.

Once again, these are early days for this invention. Early trials on a steel casting of about 4m height have suggested that the vortex is




Figure 2.44 (a) *Vortex well (with horizontal axis)* and (b) *vortex gate (with vertical axis)*.

extremely effective in absorbing the energy of the flow. In this respect its action resembles that of a ceramic foam filter. To enable the device to be used in routine casting production the energy absorbing behaviour would require to be quantified. There is no shortage of future tasks.

Vortex runner (the offset sprue)

The simple provision of an *offset sprue* causes the runner to fill tangentially, the melt spinning at high speed (Figures 2.19e and 2.45). The technique is especially suited to a vertically parted mould, where a rectangular cross-section sprue, moulded on only one side of the parting line, opens into a cylindrical runner moulded on the joint. The consequential highly organized filling of the runner is a definite improvement over many poor runner designs, as has been demonstrated by the Weibull statistics of strengths of castings produced by conventional in-line and offset (vortex) runners (Yang *et al.* 1998, 2000 and 2003). The technique produces convincingly more reliable castings than conventional in-line sprues and runner systems.

However, at this time it is not certain that the action of the vortex runner is better than other useful solutions such as the slot sprue/slot runner, or the vortex well, but it has the great benefit of simplicity. It promises to be valuable for vertically parted moulds; it deals effectively with the problem of high flow velocities in such moulds because of the great fall heights often encountered.

2.3.3 Horizontal transfer casting

The quest to avoid the gravity pouring of liquid metals has led to systems employing horizontal transfer and counter-gravity transfer. These solutions to avoid pouring are clearly seen to be key developments; both seem capable of giving competitive casting processes that offer products of unexcelled quality. The two major approaches to the first approach, horizontal transfer, are described below.

2.3.3.1 Level pour (side pour)

The 'level pour' technique was invented by Erik Laid (1978). At that time this clever technique delivered castings of unexcelled quality. It seems a pity that the process is not more generally used. This has partly occurred as the result of the process remaining commercially confidential for much of its history, so that relatively little has been published concerning the operational details that might assist a new user to achieve success. Also, the technique is limited to the type of castings, being applied easily only to plate, box or cylinder type castings where a long slot ingate can be provided up the complete height of the casting. In addition, of course, a fairly complex casting station is required.

The arrangement to achieve the so-called level filling of the mould is shown in Figure 2.46. An insulated pouring basin connects to a horizontal insulated trough that surrounds three of the four sides of the mould (a distribution system reminiscent of a Roman aqueduct). The melt enters the mould cavity via slot gates that extend vertically from the drag to the cope. Either side of each slot gate are guide plates that contain the melt between sliding seals as it flows out of the (stationary) trough and into the (descending) mould.

Casting starts with the mould sitting on the fully raised mould platform, so that the trough provides its first metal at the lowest level of the drag. The mould platform is then slowly lowered while pouring continues. The rate of withdrawal of the mould is such that the metal in the slot gate has time to solidify prior to its




Figure 2.45 (a) A conventional runner, and (b) a vortex runner, useful on a vertically parted mould (courtesy of X. Yang and Flow 3-D).

appearance as it slides out below the level of the trough.

In one of the rare descriptions of the use of the process by Bossing (1982) the large area of melt contained in the pouring basin and distribution trough would materially help to smooth the rate of flow from the point of pour to delivery into the mould. Also in this description is an additional complicated distribution system inside the mould, in which tiers of runners are provided to minimize feeding distances and maximize temperature gradients. In general, such sophistication would not be expected to be necessary for most products.

2.3.3.2 Controlled tilt casting

It seems that foundrymen have been fascinated by the intuition that tilt casting might be a solution to the obvious problems of gravity pouring. The result has been that the patent literature is littered with re-inventions of the process decade after decade.

Even so, the deceptive simplicity of the process conceals some fundamental pitfalls for the unwary. The piles of scrap seen from time to time in tilt-pour foundries are silent testimony to these hidden dangers. Generally, however, the dangers can be avoided, as will be discussed in this section.

Tilt casting is a process with the unique feature that, in principle, liquid metal can be




Figure 2.46 Level pour technique.

transferred into a mould by simple mechanical means under the action of gravity, but without surface turbulence. It therefore has the potential to produce very high quality castings. This was understood by Durville in France in the 1800s and applied by him for the casting of aluminium-bronze in an effort to reduce surface defects in French coinage.

The various stages of liquid metal transfer in the Durville Process are schematically illustrated in Figure 2.47a. In the process as




Figure 2.47 Tilt casting process (a) Durville; (b) Semi-Durville; (c) twin-poured tilting die (adapted from Nyamekye et al. 1994); and (d) outline of tilt running system design at the critical moment that metal reaches the far end of the 'sprue'.

originally conceived by Durville, the metal is melted in the same crucible as is used for the tilt machine. No pouring under gravity takes place at all. Also, since he was casting large, open-ended

ingots for subsequent working, he was able to look into the crucible and into the mould, observing the transfer of the melt as the rotation of the mould progressed. In this way he could ensure that the rate of rotation was correct to avoid any disturbance of the surface of the liquid. During the whole process of the transfer, careful control ensured that the melt progressed by 'rolling' in its skin of oxide, like inside a rubber sack, avoiding any folding of its skin by disturbances such as waves. The most sensitive part of the transfer was at the tilt angle close to the horizontal. In this condition the melt front progresses by expanding its skin of oxide, while its top surface at all times remains horizontal and tranquil.

In the USA, Stahl (1961) popularized the concept of 'tilt pouring' for aluminium alloys into shaped permanent mould castings. The gating designs and the advantages of tilt pouring over gravity top pouring have been reviewed and summarized in several papers from this source (Stahl 1963, 1986, 1989).

A useful 'bottom-gated' tilt arrangement is shown in 2.47c, d. Here the sprue is in the drag, and the remainder of the running and gating system, and the mould cavity, is in the cope. Care needs to be taken with a tilt die to ensure that the remaining pockets of air in the die can vent freely to atmosphere. Also, the die side that retains the casting has to contain the ejectors if they are needed. The layout in Figure 2.47c illustrates a unique benefit enjoyed by tilt casting: a single operator can fill both pouring cups from a large ladle prior to starting the tilt. Static gravity casting would require two pourers to fill two pouring basins.

In an effort to understand the process in some depth, Nguyen and Carrig (1986) simulated tilt casting using a water model of liquid metal flow, and Kim and Hong (1995) carried out some of the first computer simulations of the tilt casting process. They found that a combination of gravity, centrifugal and Coriolis forces govern tilt-driven flow. However, for the slow rates of rotation such as are used in most tilt casting operations, centrifugal and Coriolis effects contribute less than 10 per cent of the effects due to gravitational forces, and could therefore normally be neglected. The angular velocity of the rotating mould also made some contribution to the linear velocity of the liquid front, but this again was usually negligible because the axis of rotation was often not far from the centre of the mould.

However, despite these studies, and despite its evident potential, the process has continued to be perfectly capable of producing copious volumes of scrap castings.

The first detailed study of tilt casting using the recently introduced concepts of critical velocity and surface turbulence was carried out in the author's laboratory by Mi (2002). In addition to the benefits of working within the new conceptual framework, he had available powerful experimental techniques. He used a computer controlled, programmable casting wheel onto which sand moulds could be fixed to produce castings in an Al-4.5%Cu alloy. The flow of the metal during the filling of the mould was recorded using video X-ray radiography, and the consequential reliability of the castings was checked by Weibull statistics.

Armed with these techniques, Mi found that at the slow rotation speeds used in his work the mechanical effect of surface tension and/or surface films on the liquid meniscus could not be neglected. For all starting conditions, the flow at low tilt speeds is significantly affected by surface tension (most probably aided by the effect of a strong oxide film). Thus below a speed of rotation of approximately 7 degrees per second the speed of the melt arriving at the end of the runner is held back. Gravity only takes control after tilting through a sufficiently large angle.

As with most casting processes, if carried out too slowly, premature freezing will lead to mis-run castings. One interesting case was found in which the melt was transferred so slowly into the runner that frozen metal in the mouth of the runner acted as an obstructing ski jump to the remaining flow, significantly impairing the casting. At higher speeds, however, although ski jumps could be avoided, the considerable danger of surface turbulence increased.

The radiographic recordings revealed that the molten metal could exhibit tranquil or chaotic flow into the mould during tilt casting, depending on (i) the angle of tilt of the mould at the start of casting, and (ii) the tilting speed. The quality of the castings (assessed by the scatter in mechanical properties) could be linked directly to the quality of the flow into the mould.

We can follow the progress of the melt during the tilt casting process. Initially, the pouring basin at the mouth of the runner is filled. Only then is the tilting of the mould activated. Three starting positions were investigated:

- (i) If the mould starts from some position in which it is already tilted downward, once the metal enters the sprue it is immediately unstable, and runs downhill. The melt accelerates under gravity, hitting the far end of the runner at a speed sufficient to cause splashing. The splash action entrains

the melt surface. Castings of poor reliability are the result.

- (ii) If the mould starts from a horizontal position, the metal in the basin is not usually filled to the brim, and therefore does not start to overflow the brim of the basin and enter the runner until a finite tilt angle has been reached. At this stage the vertical fall distance between the start and the far end of the runner is likely to be greater than the critical fall distance. Thus although slightly better castings can be made, the danger of poor reliability remains. This unsatisfactory mode of transfer typifies many tilt casting arrangements, particularly the so-called Semi-Durville type process shown in Figure 2.47b.
- (iii) If, however, the mould is initially tilted slightly uphill during the filling of the basin, there is a chance that by the time the change of angle becomes sufficient to start the overflow of melt from the basin, the angle of the runner is still somewhat above the horizontal. The nature of the liquid metal transfer is now quite different. At the start of the filling of the runner the meniscus is effectively climbing a slight upward slope. Thus its progress is totally stable, its forward motion being controlled by additional tilt. If the mould is not tilted further the melt will not advance. By extremely careful control of the rate of tilt it is possible in principle to cause the melt to arrive at the base of the runner at zero velocity if required. (Such drastic reductions in speed would, of course, more than likely be counter-productive, involving too great a loss of heat, and are therefore not recommended.) Even at quite high tilting speeds of 30 degrees per second as used by Mi in his experimental mould, the velocity of the melt at the end of the runner did not exceed the critical value 0.5 m s^{-1} , and thus produced sound and repeatable castings.

The unique feature of the transfer when started above the horizontal in this way (mode iii above) is that the surface of the liquid metal is close to *horizontal* at all times during the transfer process. Thus in contrast to all other types of gravity pouring, this condition of tilt casting does not involve pouring (i.e. a free *vertical* fall) at all. It is a *horizontal* transfer process. It will be seen that in the critical region of tilt near to the horizontal, the nature of the transfer is the same as that employed originally by Durville.

Thus the optimum operational mode for tilt casting is the condition of horizontal transfer.




Figure 2.48 Map of variables for tilt pouring, showing the operational window for good castings (Mi et al. 2002).

Horizontal transfer requires the correct choice of starting angle above the horizontal, and the correct tilting speed.

An operational map was constructed (Figure 2.48), revealing for the first time a window for the production of reliable castings. It is recognized that the conditions defined by the window are to some extent dependent on the geometry of the mould that is chosen. However, the mould in Mi's experiments was designed to be close to the size and shape of many industrial castings, particularly those for automotive applications. Thus although the numerical conclusions would require some adaptation for other geometries, the principles are of general significance and are clear: there are conditions, possibly narrowly restricted, but in which horizontal transfer of the melt is possible, and gives excellent castings.

The problem of horizontal transfer is that it is slow, sometimes resulting in the freezing of the 'ski jump' at the entrance to the runner, or even the non-filling of the mould. This can usually be solved by increasing the rate of tilt *after* the runner is primed. This is the reason for the extended threshold, denoted a marginal filling condition, on the left of the window shown on the process map (Figure 2.48). A constant tilt rate (as is common for most tilt machines at this time) cannot achieve this useful extension of the filling conditions to achieve good castings. Programmable tilt rates are required to achieve this solution.

A final danger should be mentioned. At certain critical rates of rise of the melt against an

inclined surface, the development of the transverse travelling waves seems to occur to give lap problems on the cope surface of castings (illustrated later in Figure 2.60). In principle, such problems could be included as an additional threshold to be avoided on the operational window map (Figure 2.48). Fortunately, this does not seem to be a common fault. Thus in the meantime, the laps can probably be avoided by increasing the rate of tilt during this part of the filling of the mould. Once again, the benefits of a programmable tilt rate are clear.

In summary, the conclusions for tilt casting are:

1. If tilt casting is initiated from a tilt orientation at, or (even more especially) below the horizontal, during the priming of the runner the liquid metal runs downhill at a rate out of the control of the operator. The accelerating stream runs as a narrow jet, forming a persistent oxide flow tube. In addition, the velocity of the liquid at the far end of the runner is almost certain to exceed the critical condition for surface turbulence. Once the mould is initially inclined by more than 10 degrees below the horizontal at the initiation of flow, Mi found that it was no longer possible to produce reliable castings by the tilt casting process.
2. Tilt casting operations benefit from using a sufficiently positive starting angle that the melt advances into an upward sloping runner. In this way its advance is stable and controlled. This mode of filling is characterized by horizontal liquid metal transfer, promoting a mould filling condition free from surface turbulence.
3. Tilt filling is preferably slow at the early stages of filling to avoid the high velocities at the far end of the running system. However, after the running system is primed, speeding up the rate of rotation of the mould greatly helps to prevent any consequential non-filling of the castings.

2.3.4 Counter-gravity

There are some advantages to the use of gravity to action the filling of moulds. It is simple, low cost and completely reliable, since gravity has never been known to suffer a power failure. It is with regret, however, that the advantages finish here, and the disadvantages start. Furthermore, the disadvantages are serious.

Nearly all the problems of gravity pouring arise as a result of the velocity of the fall. After a trivial fall distance corresponding to the critical fall height, gravity has accelerated the melt to its critical velocity. Beyond this point there is the

danger of entrainment defects. Because the critical fall distance is so small, being only the height of a sessile drop of the liquid, nearly all actual falls exceed this limit. In other words the energy content of the melt, when allowed to fall even only relatively small distances under gravity, is nearly always sufficiently high to lead to the break-up of the liquid surface. (It is of little comfort at this time to know that foundries on the moon would fare better.)

A second fundamental drawback of gravity filling is the fact that at the start of pouring, at the time the melt is first entering the ingates, the narrowest part of the mould cross-section where volume flow rate should be slowest, the speed of flow by gravity is highest. Conversely, at a late stage of filling, when the melt is at its coldest and approaching the top of the mould cavity, and the melt needs to be fastest, the speed of filling is slowest. Thus filling by gravity gives completely the wrong filling profile.

Thus to some extent, there are always problems to be expected with castings poured by gravity. The long section on filling system design in this book is all about reducing this damage as far as possible. It is a tribute to the dogged determination of the casting fraternity that gravity pouring, despite its severe shortcomings, has achieved the level of success that it currently enjoys.

Even so, over the last 100 years and more, the fundamental problems of gravity filling have prompted casting engineers to dream up and develop counter-gravity systems.

Numerous systems have arisen. The most common is low-pressure casting, in which air or an inert gas is used to pressurize an enclosed furnace, forcing the melt up a riser tube and into the casting (Figure 2.49). Other systems use a partial vacuum to draw up the metal. Yet others use various forms of pumps, including direct displacement by a piston, by gas pressure (pneumatic pumps), and by various types of electromagnetic action.

Clearly, with a good counter-gravity system, one can envisage the filling of the mould at velocities that never exceed the critical velocity, so that the air in the mould is pushed ahead of the metal, and no surface entrainment occurs. The filling can start gently through the ingates, speed up during the filling of the main part of the mould cavity, and finally slow down and stop as the mould is filled. The final deceleration is useful to avoid any final impact at the instant the mould is filled. If not controlled in this way, the transient pressure pulse resulting from the sudden loss of momentum of the melt can cause the liquid to penetrate any sand cores, open mould joints to produce flash, and generally impair surface finish.

When using a good counter-gravity system, good filling conditions are not difficult to achieve. In fact, in comparison with gravity pouring where it is sometimes difficult to achieve a good casting, counter-gravity is such a robust technique that it is often difficult to make a bad casting. This fundamental difference between gravity and counter-gravity filling is not widely




Figure 2.49 (a) Low-pressure casting process, and (b) the usual poor filling technique.

appreciated. In general, only those who have suffered gravity filling and finally accepted counter-gravity appreciate and are amazed by the powerful benefits.

That is not to say that the technique is not sometimes used badly. The usual failure is to keep the metal velocity under control. However, in principle it can be controlled, in contrast to gravity pouring where, in principle, control is often difficult or impossible.

A concern often expressed about counter-gravity is that the adoption of filling speeds below the critical speed of approximately 0.5 m s^{-1} will slow the production rate. Such fears are groundless. For instance if the casting is 0.5 m tall (a tall casting) it can, in principle, be filled in one second. This would be a challenge!

In fact the unfounded fear of the use of low velocities of the melt leading to a sacrifice of production rate follows from the confusion of (i) flow velocity (usually measured in m s^{-1}) and (ii) melt volume flow rate (usually measured in $\text{m}^3 \text{ s}^{-1}$). For instance the filling time can be kept short by retaining a slow filling velocity but increasing the volume flow rate simply by increasing the areas of the flow channels. Worked examples to emphasize and clarify this point further will be given in Section 2.3.7 dealing with the calculation of the filling system.

2.3.4.1 Programmable control

The varying cross-sectional areas of the metal as it rises in the mould pose a problem if the fill rate through the bottom gate is fixed (as is approximately true for many counter-gravity filling systems that lack any sophistication of programmable control). Naturally, the melt may become too slow if the area of the mould increases greatly, leading to a danger of cold laps or oxide laps. Alternatively, if the local velocity is increased above the critical velocity through a narrow part of the mould, the metal may jet, causing entrainment defects.

Counter-gravity filling is unique in having the potential to address this difficulty. In principle, the melt can be speeded up or slowed down as required at each stage of filling. Even so, such programming of the fill rate is not easily achieved. In most moulds there is no way to determine where the melt level is at any time during filling. Thus if the pre-programmed filling sequence (called here the filling profile) gets out of step, its phases occurring either early or late, the filling can become worse than that offered by a constant rate system. The mistiming problems can easily arise from splashes that happen to start timers early, or from

blockage in the melt delivery system causing the time of arrival of the melt to be late.

2.3.4.2 Feedback control

The only sure way to avoid such difficulties is to provide feedback control. This involves a system to monitor the height of metal in the mould, and to feed this signal back to the delivery system, to force the system to adhere to a pre-programmed fill pattern.

One system for the monitoring of height is the sensing of the pressure of the melt in the melt delivery system. This has been attempted by the provision of a pocket of inert gas above the melt contained in the permanent plumbing of the liquid metal delivery system, and connected to a pressure transducer via a capillary.

A non-contact system used by the Cosworth sand casting process senses the change in capacitance between the melt and the mould clamp plate when the two are connected as a parallel plate condenser. Good feedback control solves many of the filling problems associated with casting production.

However, elsewhere, feedback control is little used at this time. The lack of proper control in counter-gravity leads to unsatisfactory modes of filling that explain some of the problems with the technique. The other problems relate to the remainder of the melting and melt handling systems in the foundry, that are often poor, involving multiple pouring operations from melt furnaces to ladles and then into the counter-gravity holding vessel. A widespread re-charging technique for a low-pressure casting unit is illustrated in Figure 2.49b; much of the entrainment damage suffered in such processes usually cannot be blamed on the counter-gravity system itself. The problem arises earlier because of inappropriate metal handling in the foundry before any casting takes place.

The lesson is that only limited success can be expected from a foundry that has added a counter-gravity system on to the end of a badly designed melting and melt-handling system. There is no substitute for an integrated approach to the whole production system. Some of the very few systems to achieve this so far have been the processes that the author has helped to develop: the Cosworth Process (see description later under Rule 7) and Alotech Processes. In these processes, when properly implemented, the liquid metal is never poured, never flows downhill, and is finally transferred uphill into the mould.

Finally, the concept of an integrated approach necessarily involves dealing with convection during the solidification of the casting. This serious problem is usually completely overlooked.

It has been the death of many otherwise good counter-gravity systems, but is specifically addressed in the Cosworth and Alotech systems. The problem is highlighted by the author as Rule 7.

The numerous forms of counter-gravity techniques will be discussed in detail in Volume 3 '*Casting Processes*'.

2.3.5 Surface tension controlled filling

This section starts with the interesting situation that the liquid may not be able to enter the mould at all. This is to be expected if the pressure is too low to force melt into a narrow section. It is an effect due to surface tension. If the liquid surface is forced to take up a sharp curvature to enter a non-wetted mould then it will be subject to a repulsive force that will resist the entry of the metal. Even if the metal enters, it will still be subject to the continuing resistance of surface tension, which will tend to reverse the flow of metal, causing it to empty out of the mould if there is any reduction in the filling pressure. These are important effects in narrow-section moulds (i.e. thin-section castings) and have to be taken into account.

We may usefully quantify our formulation of this problem with the well-known equation

$$P_i - P_e = \gamma(1/r_1 + 1/r_2) \quad (2.3)$$

where P_i is the pressure inside the metal, and P_e the external pressure (i.e. referring to the local environment in the mould). The two radii r_1 and r_2 define the curvature of the meniscus in two planes at right angles. The equation applies to the condition when the pressure difference across the interface is exactly in balance with the effective pressure due to surface tension. To describe the situation for a circular-section tube of radius r (where both radii are now identical), the relation becomes:

$$P_i - P_e = 2\gamma/r \quad (2.4)$$

For the case of filling a narrow plate of thickness $2r$, one radius is, of course, r , but the radius at right angles becomes infinite, so the reciprocal of the infinite radius equates to zero (i.e. if there is no curvature there is no pressure difference). The relation then reduces to the effect of only the one component of the curvature, r :

$$P_i - P_e = \gamma/r \quad (2.5)$$

We have so far assumed that the liquid metal does not wet the mould, leading to the effect of capillary repulsion. If the mould is wetted then the curvature term γ/r becomes negative, so allowing surface tension to assist the metal to

enter the mould. This is, of course, the familiar phenomenon of capillary attraction. The pores in blotting paper attract the ink into them; the capillary channels in the wick of a candle suck up the molten wax; and the water is drawn up the walls of a glass capillary. In general, however, the casting technologist attempts to avoid the wetting of the mould by the liquid metal. Despite all efforts to prevent it, wetting sometimes occurs, leading to the penetration of the melt into sand cores and moulds.

Continuing now in our assumption that the metal-mould combination is non-wetting, we shall estimate what head of metal will be necessary to force it into a mould to make a wall section of thickness $2r$ for a gravity casting made under normal atmospheric pressure. If the head of liquid is h , the hydrostatic pressure at this depth is ρgh , where ρ is the density of the liquid, and g the acceleration due to gravity. The total pressure inside the metal is therefore the sum of the head pressure and the atmospheric pressure, P_a . The external pressure is simply the pressure in the mould due to the atmosphere P_a plus the pressure contributed by mould gases P_m . The equation now is

$$(P_a + \rho gh) - (P_a + P_m) > \gamma/r \quad (2.6)$$

giving immediately

$$\rho gh - P_m > \gamma/r \quad (2.7)$$

The back-pressure due to outgassing in the mould lowers the effective head driving the filling of the mould. It is good practice, therefore, to vent narrow sections, reducing this resistance to practically zero if possible.

It is also clear from the above result that, provided the mould is permeable and/or well vented, atmospheric pressure plays no part in helping or resisting the filling of thin sections in air, since it acts equally on both sides of the liquid front, cancelling any effect. Interestingly, the same equation and reasoning applies to casting in vacuum, which, of course, can be regarded as casting under a reduced atmospheric pressure. Clearly, a vacuum casting is therefore not helpful in overcoming the resistance to filling provided by surface tension (although, to be fair, it may help by reducing P_m by outgassing the mould to some extent prior to casting, and it will help where the permeability of the mould is low, where residual gases may be compressed ahead of the advancing stream. Vacuum casting may also help to fill the mould by reducing—but not eliminating—the effect of the surface film of oxide or nitride).

The case of vacuum-assisted filling (not vacuum casting) is quite different, since the

vacuum is not now applied to both the front and back of the liquid meniscus, thus cancelling any benefit as above, but applied only to the advancing front as illustrated in Figure 2.50. This application of a reduced pressure to one side of the meniscus creates a differential pressure that drives the flow. The differential




Figure 2.50 (a) A plaster mould encased in a steel box using vacuum-assisted filling through the base of the mould. No formal running system is required for such small thin-walled castings. (b) Sand mould to make four cover castings, using narrow slot filling system to maximize benefits from surface tension and wall friction.

pressure acts by atmospheric pressure continuing to apply to the liquid metal via the running system, but the atmospheric pressure in the mould is reduced by applying a (partial) vacuum in the mould cavity. This is achieved by drawing the air out either through the permeable mould, or through fine channels cut through to the section required to be filled (as is commonly applied to the trailing edge of an aerofoil blade section). In this way P_m is guaranteed to be zero or negligible, and P_a remains a powerful pressure to assist in overcoming surface tension as the equation indicates:

$$P_a + \rho gh > \gamma/r \quad (2.8)$$

It is useful to evaluate the terms of this equation to gain a feel for the size of the effects involved. Taking, roughly, g as 10 ms^{-2} , and the liquid aluminium density ρ as 2500 kg m^{-3} and γ as 1.0 N m^{-1} (for steels and high-temperature alloys the corresponding values are approximately 7000 kg m^{-3} and 2.0 N m^{-1}), the resistance term γ/r works out to be 2 kPa for a 1 mm section (0.5 mm radius) and 10 kPa for a 0.1 mm radius trailing edge on a turbine blade.

For a head of metal $h = 100 \text{ mm}$ the head pressure ρgh is 2.5 kPa, showing that the 1 mm section might just fill. However, the 0.1 mm trailing edge has no chance; the head pressure being insufficient to overcome the repulsion of surface tension. However, if vacuum assistance were applied (NB not vacuum casting, remember) then the additional 100 kPa of atmospheric pressure normally ensures filling. In practice it should be noted that the full value of atmospheric pressure is not easily obtained in vacuum-assisted casting; in most cases a value nearer half an atmosphere is more usual. Even so, the effect is still important: one atmosphere pressure corresponds to 4 m head of liquid aluminium, and approximately 1.5 m head of denser metals such as irons, steels and high-temperature alloys. In modest-sized castings of overall height around 100 mm or so, these valuable pressures to assist filling are not easily obtainable by other means. The pressure delivered by a feeder placed on top of the casting may only apply the additional head corresponding to its height of perhaps 0.1 to 0.4 m; only one tenth of the pressures that can be applied by the atmosphere.

For those castings that have sections of only 1 or 2 mm or less, the surface tension wields strong control over the tight radius of the front. Filling is only possible by the operation of additional pressure, such as that provided by the jeweller's centrifuge, or the application of vacuum assistance. Filling can occur upwards or downwards without problems, being always

under the control of the surface tension, which is effectively so strong in such thin sections that it keeps the surface intact. Surface turbulence is thereby suppressed. The liquid has insufficient room to break up into drops, or to jet or splash. The integrity of the front is under the control of surface tension at all times. This special feature of the filling of very thin-walled castings means that they do not require formal running systems. In fact, such thin-walled investment castings are made successfully by simply attaching wax patterns in any orientation directly to a sprue (Figure 2.50). The metal flows similarly with either gravity or counter to gravity, and no 'runner' or 'gate' is necessary.

To gain an idea of the head of metal required to force the liquid metal into small sections, from Equation 2.8 we have:

$$\rho gh = \gamma/r$$

$$h = \gamma/r\rho g \quad (2.8a)$$

Using the values for aluminium and steel given above, we can now quickly show that to penetrate a 1 mm section we require heads of at least 80 and 60 mm respectively for these two liquids.

If the section is halved, the required head for penetration is, of course, doubled. Similarly, if the mould shape is not a flat section that imposes only one curvature on the meniscus, but is a circular hole of diameter 1 mm, the surface then has an additional curvature at right angles to the first curvature. Equation 2.4 shows the head is doubled again.

In general, because of the difficulty of predicting the shape of the liquid surface in complex and delicate castings, the author has found that a safety factor of 2 is not excessive when calculating the head height required to fill thin sections. This safety factor is quickly used up when allowances for errors in the wall thickness, and the likely presence of surface films is taken into account.

The resistance to flow provided by surface tension can be put to good effect in the use of slot-shaped filling systems. In this case the slots are required to be a maximum thickness of only 1 or 2 (perhaps 3 at the most) mm for engineering castings (although, clearly, jewellery and other widget type products might require even thinner filling systems). Figure 2.50b shows a good example of such a system. A similar filling system for a test casting designed by the author, but using a conical basin (not part of the author's original design!), was found to perform tolerably well, filling without the creation of significant defects (Groteke 2002). It is quite evident, however, when filling is complete such

narrow filling channels offer no possibility of significant feeding. This is an important issue that should not be forgotten. In fact, in these trials, this casting never received the proper attention to feeding, and as a consequence suffered surface sinks and internal microporosity (the liquid alloy was clearly full of bifilms that were subsequently opened by the action of solidification shrinkage).

Finally, however, in some circumstances there may be fundamental limitations to the integrity of the liquid front in very thin sections.

- (i) There is a little-researched effect that the author has termed microjetting (*Castings* 2003). This phenomenon has been observed during the filling of liquid Al-7Si-0.4Mg alloy into plaster moulds of sections between 1 and 3 mm thickness (Evans *et al.* 1997). It seems that the oxide on such small liquid areas temporarily restrains the flow, but repeatedly splits open, allowing jets of liquid to be propelled ahead of the front. The result resembles advancing spaghetti. The mechanical properties are impaired by the oxide films around the jets that become entrained in the maelstrom of progress of the front. Whether this unwelcome effect is common in thin-walled castings is unknown, and the conditions for its formation and control are also unknown. Very thin walled castings remain to be researched.
- (ii) In pressure die-castings a high velocity v of the metal through the gate is necessary to fill the mould before too much heat is lost to the die. Speeds of between 25 and 50 m s⁻¹ are common, greatly exceeding the critical velocity of approximately 0.5 m s⁻¹ that represents the watershed between surface tension control and inertial control of the liquid surface. The result is that entrainment of the surface necessarily occurs on a huge scale. The character of the flow is now dictated by inertial pressure, proportional to v^2 , that vastly exceeds the restraining influences of gravity or surface tension. This behaviour is the underlying reason for the use of PQ^2 diagrams as an attempt to understand the filling of pressure die castings. In this approach a diagram is constructed with vertical axis denoting pressure P , and horizontal axis denoting flow rate Q . The parabolic curves are linearized by squaring the scale of the Q values on the horizontal axis. The approach is described in detail in much of the pressure die-casting literature (see, for instance, Wall and Cocks 1980). In practice, it is not certain how valuable this technique is, now that computer

simulation is beginning to be accepted as an accurate tool for the understanding of the process.

2.3.6 Inclusion control: filters and traps

The term 'inclusion' is a shorthand generally used for 'non-metallic inclusion'. However, it is to be noted that such defects as tungsten droplets from a poor welding technique can appear in some recycled metals; these, of course, constitute 'metallic inclusions'. Furthermore, one of the most common defects in many castings is the bubble, entrained during pouring. This constitutes an 'air inclusion' or 'gas inclusion.'

The fact that bubbles are trapped in the casting from the filling stage is remarkable in itself. Why did the bubble not simply rise to the surface, burst and disappear? This is a simple but important question. In most cases the bubble will not have been retained by the growth of solid, because solid will, in general, not have time to form. The answer in practically all cases is that oxide films will also be present. In fact the bubbles themselves are simply sections of the oxide films that have not perfectly folded back together. The bubbles decorate the double films, as inflated islands in the folds. Thus many bubbles, entangled in a jumble of films, never succeed to reach the surface to escape. Even those that are sufficiently buoyant to power their way through the tangle may still not burst at the surface because of the layers of oxide that bar its final escape.

This close association of bubbles and films (since they are both formed by the same turbulent entrainment process; they are both entrainment defects) is called by me bubble damage. We need to keep in mind that the bubble is the visible part of the total defect. The surrounding region of bifilms to which it is connected act as cracks, and can be much more extensive and often invisible. However, the presence of such films is the reason that cracks will often appear to start from porosity, despite the porosity having a nicely rounded shape that would not in itself appear to be a significant stress raiser.

Whereas inclusions are generally assumed to be particles having a compact shape, it is essential to keep in mind that the most damaging inclusions are the films (actually always double, unbonded films, remember, so that they act as cracks), and are common in many of our common casting alloys. Curiously, the majority of workers in this field have largely overlooked this simple fact. It is clear that techniques to remove particles will often not be effective for films, and vice versa. The various methods to clean metals prior to casting have been reviewed in Chapter 1

as a fulfilment of Rule 1. The various methods to clean metals travelling through the filling systems of castings will be reviewed here.

2.3.6.1 Dross trap (or slag trap)

The dross trap is used in light alloy and copper-based alloy casting. In ferrous castings it is called a slag trap. For our purposes we shall consider the devices as being one and the same.

It is good sense to include a dross trap in the running system. In principle, a trap sited at the end of the runner will take the first metal through the runner and keep it away from the gates. This first metal is both cold, having given up much of its heat to the running system en route, and will have suffered damage by oxide or other films during those first moments before the sprue is properly filled.

In the past, designs have been along the lines of Figure 2.37a. This type of trap was sized with a view to accommodating the total volume of metal through the system until the down-runner and horizontal runner were substantially filled. This was a praiseworthy aim. In practice, however, it was a regular joke among foundrymen that the best quality metal was concentrated in the dross trap and all the dross was in the casting! What had happened to lend more than an element of truth to this regrettable piece of folk-law?

It seems that this rather chunky form of trap sets up a circulating eddy during filling. Dross arriving in the trap is therefore efficiently floated out again, only to be swept through the gates and into the casting a few moments later! Ashton and Buhr (1974) have carried out work to show that runner extensions act poorly as traps for dirt. They observed that when the first metal reached the end of the runner extension it rose, and created a reflected wave which then travelled back along the top surface of the metal, carrying the slag or dirt back towards the ingates. Such observations have been repeated on iron and steel casting by Davis and Magny (1977) and on many different alloys in the author's laboratory using real-time radiography of moulds during casting. The effect has also been simulated in computer models. It seems, therefore, to be real and universal in castings of all types. We have to conclude that this design of dross trap cannot be recommended!

Figure 2.26b shows a simple wedge trap. It was thought that metal flowing into the narrowing section was trapped, with no rebound wave from the end wall, and no circulating eddy can form. However, video radiographic studies have shown that such traps can reflect a backward wave if the runner is sufficiently deep.

Also, of course, the volume of melt that they can retain is very limited.

A useful design of dross trap appears to be a volume at the end of the runner that is provided with a narrow entrance (the extension shown in broken outline in Figure 2.37b) to suppress any outflow. It is a kind of wedge trap fitted with a more capacious end. In the case of persistent dross and slag problems, the trap can be extended, running around corners and into spare nooks and crannies of the mould. If the entrance section is less than the height of a sessile drop, it will be filled by the entering liquid, thus being too narrow to allow a reflected wave to exit. It should therefore retain whatever material enters. In addition, depending on the narrowness of the tapered wedge entrance, to some extent the device should be capable of filling and pressurizing the runner in a progressive manner akin to the action of a gate. This is a useful technique to reduce the initial transient momentum problems that cause gates to fill too quickly during the first few seconds. This potentially useful benefit has yet to be researched more thoroughly so as to provide useful guidelines for mould design.

The device can be envisaged to be useful in combination with other forms of by-pass designs such as those shown in 2.37d and e.

Slag pockets

For iron and steel castings the term 'slag pocket' is widely used for a raised portion of the runner that is intended to collect slag. The large size of slag particles and their large density difference with the melt encourages such separation. However, such techniques are not the panacea that the casting engineer might wish for.

For instance the use of traps of wedge-shaped design, Figure 2.51, is expected to be almost completely ineffective because the circulation pattern of flow would take out any material that happened to enter. On the other hand, a rectangular cavity has a secondary flow into which buoyant material can transfer if it has sufficient time, and so remain trapped in the upper

circulating eddy. This consideration again emphasizes the need for relatively slow flow for its effectiveness. Also, of course, none of these traps can become effective until the runner and the trap become filled with metal. Thus many filling systems will have passed much if not all such unwanted material before the separation mechanisms have a chance to come into operation. A further consideration that causes the author to hesitate to recommend such traps is that they locally remove the constraint on the flow of the metal, allowing surface turbulence. Thus the traps might cause more problems than they solve.

Davis and Magny (1977) observed the filling of iron and steel castings by video radiography. They confirmed that most slag retention devices either do not work at all, or work with only partial effectiveness. These authors made castings with different amounts of slag, and tested the ability of slag pockets sited above runners to retain the slag. They found that rectangular pockets were tolerably effective only if the velocity of flow through the runner was below 0.4 m s^{-1} (interesting that this is precisely the critical speed at which surface turbulence will occur, and so cause surface phases to be turbulently stirred back into the bulk liquid). For a casting only 0.1 m high the metal is already travelling four times too fast. For such reasons the experience with slag pockets has been somewhat mixed in practice.

In defence of the historical use of such traps it must be borne in mind that they were traditionally used with pressurized filling systems, heavily choked at the gate, so that the runner was encouraged to fill as quickly as possible, making the trap effective at an early stage of filling. Also, if the choking action was sufficiently severe the speed of flow in the runner may be sufficiently slow to ensure slag entrapment. However, this text does not recommend pressurized filling systems mainly because of the problems that follow from the necessarily high ingate velocities.

Perhaps, therefore, the slag trap has come to the end of its useful life.




Figure 2.51 Various designs of slag pockets: (a) relatively ineffective self-emptying wedge; (b) rectangular trap stores buoyant phases in upper circulating flow.

2.3.6.2 Swirl traps

The centripetal trap is an accurate name for this device, but rather a mouthful. It is also known as a whirl gate, or swirl gate, which is shorter, but inaccurate since the device is not really a gate at all. Choosing to combine the best of both names, we can call it a swirl trap. This is conveniently short, and accurately indicates its main purpose for trapping rubbish.

The idea behind the device is the use of the difference in density between the melt and the various unwanted materials which it may carry, either floating on its surface or in suspension in its interior. The spinning of the liquid creates a centrifugal action, throwing the heavy melt towards the outside where it escapes through the exit, to continue its journey into the casting. Conversely, the lighter materials are thrown towards the centre, where they coagulate and float. The centripetal acceleration a_c is given by:

$$a_c = V^2/r \quad (2.9)$$

where V is the local velocity of the melt, and r is the radius at that point. For a swirl trap of 50 mm radius and sprue heights of 0.1 m and 1 m, corresponding to velocities of 1.5 and 4.5 m s⁻¹ respectively we find that accelerations of 40 and 400 m s⁻² respectively are experienced by the melt. Given that the gravitational acceleration g is 9.81 m s⁻², which we shall approximate to 10 m s⁻², these values illustrate that the separating forces within a swirl trap can be between 4 and 40 times that due to gravity. These are, of course, the so-called 'g' forces experienced in centrifuges.

So much for the theory. What about the reality?

Foundrymen have used swirl traps extensively. This popularity is not easy to understand because, unfortunately, it cannot be the result of their effectiveness. In fact the traps have worked so badly that Ruddle (1956) has recommended not to use them on the grounds that their poor performance does not justify the additional complexity. One has to conclude that their extraordinarily wide use is a reflection of the fascination we all have with whirlpools, and an unshakeable belief, despite all evidence to the contrary, that the device should work.

Regrettably, the swirl trap is expected to be completely useless for film-forming alloys where inclusions in the form of films will be too sluggish to separate. Since some of the worst inclusions are films, the swirl trap is usually worse than useless, creating more films than it can remove. Worse still, in the case of alloys of aluminium and magnesium, their oxides are

denser than the metal, and so will be centrifuged outwards, into the casting! Swirl traps are therefore of no use at all for light alloys (however, notice that the vortex sprue base, although not specifically designed to control inclusions, might have some residual useful effect for light alloys since the outlet is central). Finally, swirl traps seem to be difficult to design to ensure effective action. In the experience of the author, most do much more harm than good.

The inevitable conclusion is that swirl traps should be avoided.

The remainder of this section on swirl traps is for those who refuse to give up, or refuse to believe. It also serves as a mini-illustration of the real complexity of apparently simple foundry solutions. Such illustrations serve to keep us humble.

It is worthwhile to examine why the traditional swirl trap performs so disappointingly. On examination of the literature, the textbooks, and designs in actual use in foundries, three main faults stand out immediately:

1. The inlet and exit ducts from the swirl traps are almost always opposed, as shown in Figure 2.52a. The rotation of the metal as a result of the tangential entry has, of course, to be brought to a stop and reversed in direction to make its exit from the trap. The disorganized flow never develops its intended rotation and cannot help to separate inclusions with any effectiveness.
Where the inlet and exit ducts are arranged in the correct tangential sense, then Trojan *et al.* (1966) have found that efficiency is improved in their model results using wood chips in water. Even so, efficiencies in trapping the chips varied between the wide limits of 50 and 100 per cent.
2. The inlet is nearly always arranged to be higher than the exit. This elementary fault gives two problems. First, any floating slag or dross on the first metal to arrive is immediately carried out of the trap before the trap is properly filled (Figure 2.52c). Second, as was realized many years previously (Johnson and Baker 1948), the premature escape of metal hampers the setting up of a properly developed spinning action. Thus the trap is slow to develop its effect, perhaps never achieving its full speed in the short time available during the pour. This unsatisfactory situation is also seen in the work of Jirsa (1982), who describes a swirl trap for steel casting made from preformed refractory sections. In this design the exit was again lower than the entrance, and filling of the trap was encouraged merely by making the exit smaller than the entrance.




Figure 2.52 Swirl traps showing (a) incorrect opposed inlet and exit ducts; (b) correct tangential arrangements; (c) incorrect low exit; (d) correct high exit.

Much metal and slag almost certainly escaped before the trap could be filled and become fully operational. Since 90 per cent efficiency was claimed for this design it seems probable that all of the remaining 10 per cent which evaded the trap did so before the trap was filled (in other words, the trap was working at zero efficiency during this early stage). Jeancolas *et al.* (1971) report an 80 per cent efficiency for their downhill swirl trap for bronze and steel casting, but admit that the trap does not work at all when only partly full.

3. In many designs of swirl trap there is insufficient attention paid to providing accommodation for the trapped material. For instance, where the swirl trap has a closed top the separated material will collect against the centre of the ceiling of the trap. However, work with transparent models illustrates clearly how perturbations to the flow cause the inclusions, especially if small, to ebb and flow out of these areas back into the main flow into the casting (Jeancolas *et al.* 1971; Trojan *et al.* 1966). Also, of course, traps of such limited volume are in danger of becoming completely overwhelmed, becoming so full of slag or dross that the flow into the casting becomes necessarily contaminated.

Where the trap has an open top the parabolic form of the liquid surface assists the concentration of the floating material in the central 'well' as shown in Figure 2.52d. The extra height for the separated materials to rise into is useful to keep the unwanted material well away from the exit, despite variations from time to time in flow rate. Some workers have opened out the top of the trap, extending it to the top of the

cope, level with the pouring bush. This certainly provides ample opportunity for slag to float well clear, with no danger of the trap becoming overloaded with slag. However, the author does not recommend an open system of this kind, because of the instability which open-channel systems sometimes exhibit, causing surging and slopping between the various components comprising the 'U'-tube effect between the sprue, swirl trap and mould cavity.

It is clear that the optimum design for the swirl trap must include the features:

- (i) the entrance at the base of the trap;
- (ii) the exit to be sited at a substantially higher level;
- (iii) both entrance and exit to have similar tangential direction, and
- (iv) an adequate height above the central axis to provide for the accumulation of separated debris.

In most situations the inlet will be moulded in the drag, and the exit in the cope, which is the most marginal difference in level between the two. At the high speeds at which the metal can be expected to enter the trap the metal will surge over this small ledge with ease, taking inclusions directly into the casting, particularly if the inlet and outlet are in line as shown in Figure 2.52b. This simplest form of cope/drag parting line swirl trap cannot be expected to work.

The trap may be expected to work somewhat more effectively as the angle of the outlet progresses from 90–180 degrees. (The 270 degree option would be more effective still, except that some reflection will show it to be unmouldable

on this single joint line; the exit will overlay the entrance ports! Clearly, for the 270 degree option to be possible, the entrance and exits have to be moulded at different levels, necessitating a second joint line provided by a core or additional mould part.)

When using preformed refractory sections, or pre-formed baked sand cores, as is common for larger steel castings, the exit can with advantage be placed considerably higher than the entrance (Figure 2.52d).

These simple rules are designed to assist the trap to spin the metal up to full speed before the exit is reached, and before any floating or emulsified less-dense material has had a chance to escape.

For the separation of particulate slag inclusions from some irons and steels, *Castings 1991* showed that a trap 100 mm diameter in the running system of moulds 0.1 to 1 m high would be expected to eliminate inclusions of 0.2 to 0.1 mm respectively. The conclusion was that, when correctly designed, the swirl trap could be a useful device to divert unwanted buoyant particles away from ferrous castings.

We have to remember, however, that it is not expected to work for film type inclusions. Compared to particles, films would be expected to take between 10 and 100 times longer to separate under an equivalent field force. Thus most of the important inclusions in a large number of casting alloys will not be effectively trapped. Thus the alloys that need the technique most are least helped.

This damning conclusion applies to other field forces such as electromagnetic techniques that have recently been claimed to remove inclusions from melts. It is true that forces can be applied to non-conducting particles suspended in the liquid. However, whereas compact particles move relatively quickly, and can be separated in the short time available while the melt travels through the field. Films experience the same force, but move too slowly because of their high drag, and so are not removed.

In summary, we can conclude that apart from certain designs of by-pass trap, other varieties of traps are not recommended. In general they almost certainly create more inclusions than they remove.

2.3.6.3 Filters

Filters take many forms: as simple strainers, woven ceramic cloths, and ceramic blocks of various types. Naturally, their effectiveness varies from application to application, as is discussed here.

Strainers

A sand or ceramic core may be moulded to provide a coarse array of holes, of a size and distribution resembling a domestic colander. A typical strainer core might be a cylinder 30 to 50 mm diameter, 10 to 20 mm long, containing 10 or more holes, of diameter approximately 3–5 mm (Figure 2.53a). These devices are mainly used to prevent slag entering iron castings. The domestic colander is usually used to strain aggregates such as peas. These represent solid spherical particles of the order of 5 mm diameter. Thus, when applied to most metal castings, the rather open design of strainers means that they can hardly be expected to perform any significant role as filters.

In fact, Webster (1967) has concluded that the strainer works by reducing the rate of flow of metal, assisting the upstream parts of the filling system to prime, and thereby allowing the slag to float. It can be held against the top surface of the runner, or in special reservoirs placed above the strainers to collect the retained slag. Webster goes on to conclude that if the strainer only acts to reduce the rate of flow, then this can be carried out more simply and cheaply by the proper design of the running system.

This may not be the whole story. The strainer may be additionally useful to laminize the flow (i.e. cause the flow to become more streamlined).

However, whatever benefits the strainer may have, its action to create jets downstream of the strainer is definitely not helpful. The placing of a strainer in a geometry that will quickly fill the region at the back of the strainer would be a




Figure 2.53 Various filters showing (a) a strainer core (hardly a filter at all); (b) woven cloth or mesh, forming a two-dimensional filter; (c) ceramic foam and extruded blocks, constituting three-dimensional filters.

great advantage. A geometry to suppress jetting is provided by the tangential filter print to be discussed later (Figure 2.56). The extruded or pressed ceramic filters with their arrays of parallel pores are, of course, equivalent to strainers with a finer pore size. They also benefit from the tangential placement to the oncoming flow as will be described.

Over the years there has been much work carried out to quantify the benefits of the use of filters. Nearly all of these have shown measurable, and sometimes important, gains in freedom from defects and improvements in mechanical properties. These studies are too numerous to list here, but include metals of all types, including Al alloys, irons and steels. The relatively few negative results can be traced to the use of unfavourable siting or geometry of the filter print. For positive and reliable results, these aspects of the use of filters cannot be overlooked. Special attention is devoted to them in what follows.

Woven cloth or mesh

For light alloys, steel wire mesh or glass cloth (Figure 2.53b) is used to prevent the oxides from entering the casting. Cloth filter material has the great advantage of low cost.

The surprising effectiveness of these rather open meshes is the result of the most important inclusions being in the form of films, which appear to be intercepted by and wrap around the strands of the mesh. Openings in the mesh or weave are typically 1–2 mm; this gives good results, being highly effective in retaining films down to this size range. Significantly, it is also a confirmation of the large size of the majority of films that cause problems in castings, particularly in light alloys.

The use of steel wire mesh is also useful to retain films. The steel does not have time to go into solution during the filling of aluminium alloy castings, so that the material of the casting is in no danger of contamination. However, of course, the steel presents a problem of iron contamination during the recycling of the running system. Even the glass cloth can sometimes cause problems during the break-up of the mould, when fragments of glass fibre can be freed to find their way into the atmosphere of the foundry, and cause breathing problems for operators. Both materials therefore need care in use.

Some glass cloth filters are partially rigidized with a ceramic binder, and some by impregnation with phenolic resin. (The outgassing of the resin can cause the evolution of large bubbles when contacted by the liquid metal. Provided

the bubbles do not find their way into the casting the overall effect of the filters is definitely beneficial in aluminum alloys.) Both types soften at high temperature, permitting the cloth to stretch and deform.

A woven cloth based on a high silica fibre has been developed to avoid softening at these temperatures, and might therefore be very suitable for use with light alloys. In fact at the present time its high-temperature performance usually confines its use to copper-based alloys and cast irons. There are few data to report on the use of this material. However, it is expected that its use will be similar to that of the other meshes, so that the principles discussed here should still apply.

Despite the attraction of low cost, it has to be admitted that, in general, the glass cloth filters are not easy to use successfully.

For instance, as the cloth softens and stretches there is a strong possibility that the cloth will allow the metal to by-pass the filter. It is essential to take this problem into account when deciding on the printing of the filter. Clearly, it is best if it can be firmly trapped on all of its edges. If it can be held on only three of its four edges the vulnerability of the unsupported edge needs careful consideration. For instance, even though a cross-joint filter may be properly held on the edges that are available, the filter is sometimes defeated by the leading edge of the cloth bending out of straight, bowing like a sail, and thus allowing the liquid to jet past. All the filters shown in Figure 2.54 show this problem. It may be better to abandon cloth filtration if there is any danger of the melt jetting around a collapsing filter.

When sited at the point where the flow crosses a joint as in Figure 2.54 a greensand mould will probably hold the cloth successfully, the sand impressing itself into the weave, provided sufficient area of the cloth is trapped in the mould joint. In the case of a hard sand mould or metal die, the cloth requires a shallow print which must be deep enough to allow it room if the joint is not to be held apart. Also, of course, the print must not be too deep, otherwise the cloth will not be held tight, and may be pulled out of position by the force of the liquid metal. Some slight crushing of a hard sand mould is desirable to hold the cloth as firmly as possible.

A rigidized cloth filter can be inserted across the flow by simply fitting it into the pattern in a pre-moulded slot across the runner and so moulding it integral with the mould (Figure 2.54c). However, this is only successful for relatively small castings. Where the runner area becomes large and the time and temperature become too

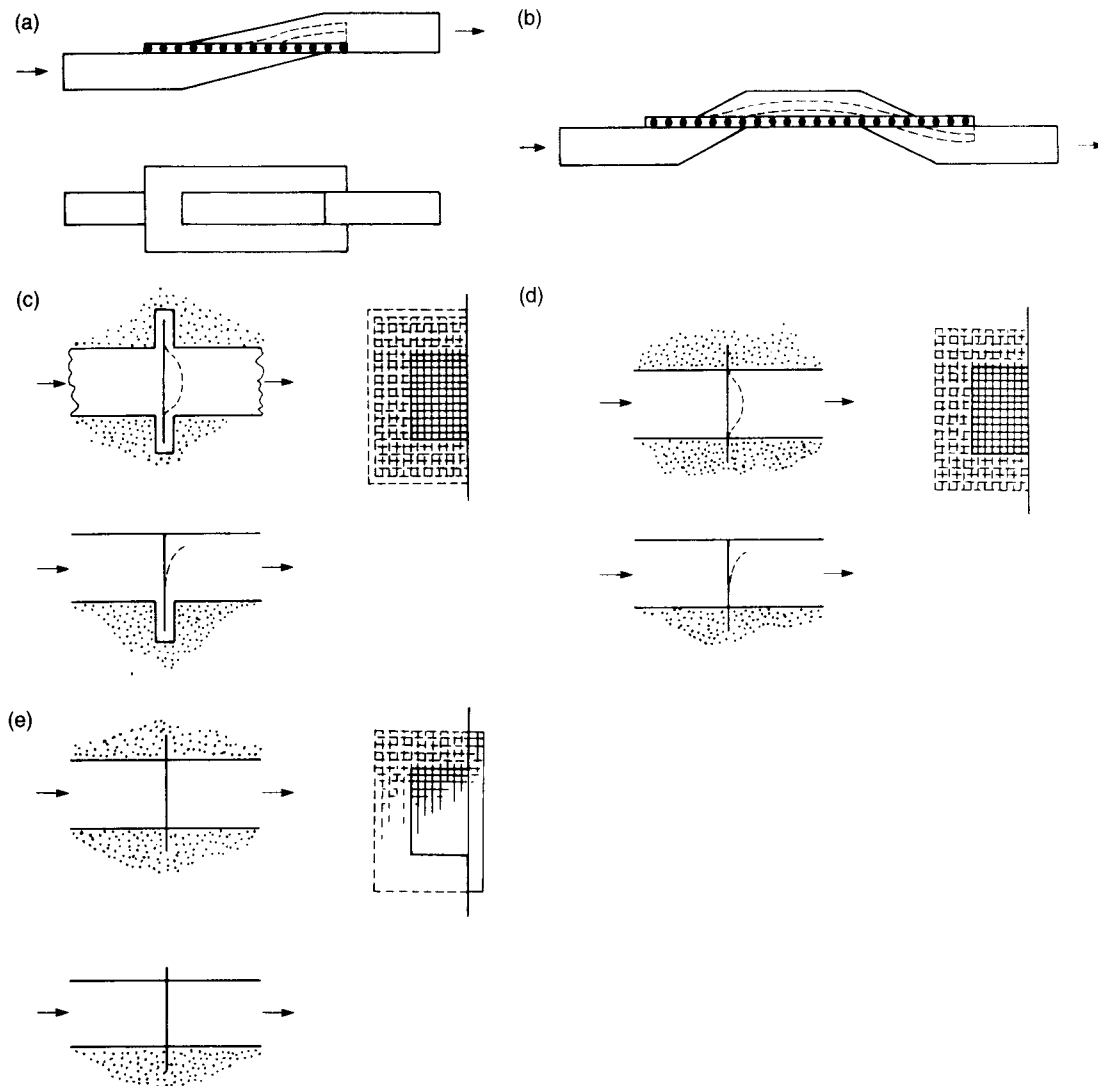


Figure 2.54 Siting of cloth filters (a) in the mould joint; (b) in a double crossing of the joint; (c) in a slot moulded across the runner; (d) in a slot cut in the runner pattern; (e) with an additional upstand across the joint plane to assist sealing.

high, the filter softens and bows in the force of the flow. Even if it is not entirely pulled out of position it may be deformed to sag like a fence in a gale, so that metal is able to flow over the top. This is the reason for the design shown in Figure 2.54e. The edge of the filter crosses the joint line, either to sit in a recess accurately provided on the other half of the mould, or if the upstand is limited to a millimetre or so, to be simply crushed against the other mould half. (The creation of some loose grains of sand is of little consequence in the running system, as has been shown by Davis and Magny (1977); loose

material in the runner is never picked up by the metal and carried into the mould. The author can confirm this observation as particularly true for systems that are not too turbulent. The laminizing action of the filter itself is probably additionally helpful.)

If the filter is introduced at an earlier stage of manufacture during the production of a sand mould, it can be placed in position in a slot cut in the runner pattern. When the sand is introduced the filter is automatically bonded into the mould (Figure 2.54d). Again, an upstand above the level of the joint may be useful

(Figure 2.54e). In any case, when using filters across runners, it also helps to arrange for the selvage (the reinforced edge of the material) of the cloth to be uppermost to give the unsupported edge most strength; the ragged cut edge has little strength, letting the cloth bend easily, and allowing some, or perhaps all, of the flow to avoid the filter. All the cloth filters used as shown in Figure 2.54 are defeatable, since they are held only on three sides. The fourth side is the point of weakness. Failure of the filter by the liquid overshooting this unsupported edge can result in the creation of more oxide dross than the filter was intended to prevent! Increasing the trapped area of filter in the mould joint can significantly reduce the problem.

Geometries that combine bubble traps (or slag or any other low density phase) are shown in Figure 2.55 for in-line arrangements, and for those common occasions when the runner is required to be divided to go in opposite directions. With shallow runners of depth of a few millimetres there is little practical difference in whether the metal goes up or down through the filter. Thus several permutations of these geometries can be envisaged. Much depends on the links to the gates, and how the gates are to be placed on the casting. In general, however, I usually aim to have the runner exit from the filter below the joint.

Cloth filters are entirely satisfactory where they can be held around all four sides. This is the case at the point where gates are taken vertically upwards from the top of the runner. This is a relatively unusual situation, where, instead of a two-part mould, a third mould part forms a base to the mould and allows the runner system to be located under the casting. Alternatively, a special core can be used to create an extra joint beneath the general level of the casting.

Another technique for holding the filter on all sides is the use of a 'window frame' of strong paper or cardboard that is bonded or stapled to the cloth. The frame is quickly dropped into its slot print in the mould, and gives a low cost rigid surround that survives sufficiently long to be effective.

Ceramic block filters

Ceramic block filters of various types introduced in about 1980 have become popular, and have demonstrated impressive effectiveness in many applications in running systems.

Unfortunately, much that has been written about the mechanisms by which they clean up the cast material appears to be irrelevant. This is because most speculation about the filtration mechanisms has considered only particulate

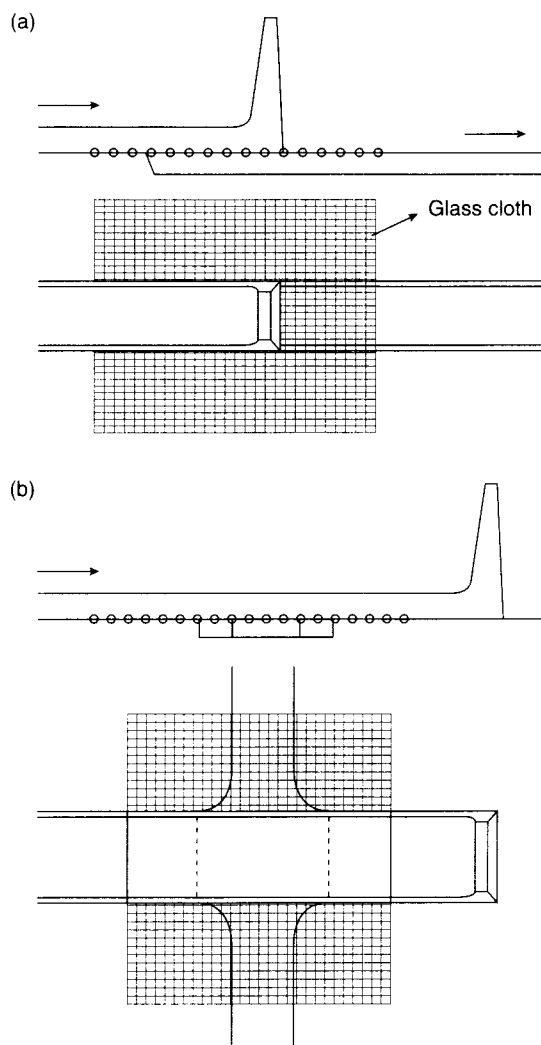


Figure 2.55 Uses of glass cloth filtration (a) for an in line runner; (b) for transverse runners.

inclusions. As has become quite clear over recent years, the most important and widespread inclusions are actually films. Thus the filtration mechanism at work is clearly quite different, and, in fact, easily understood.

In aluminium alloys the action of a ceramic foam filter to stop films has, in general, not been recognized. This is probably because the films are so thin (a new film may be only 20 nm thick, making the doubled-over entrained bifilm still only about 40 nm) that they cannot be detected when wrapped around sections of the ceramic filter. This explains part of the curious experience of finding that a filter has cleaned up a casting, but on sectioning the filter to examine it

under the microscope, not a single inclusion can be found.

The contributing effect, of course, is that the filter acts to improve the filling behaviour of the casting, so reducing the number of inclusions that are created in the mould during the filling process. This behaviour was confirmed by Din *et al.* (2003) who found only about 10 per cent of the action of the filter was the result of filtration, but 90 per cent was the result of improved flow.

A further widespread foundry experience is worth a comment. On occasions the quantity of inclusions has been so great that the filter has become blocked and the mould has not filled completely. Such experiences have caused some users to avoid the use of filters. However, in the experience of the author, such unfortunate events have resulted from the use of poor front ends of filling systems (poor basin and sprue designs) that create huge quantities of oxide films in the pouring process. The filter has therefore been overloaded, leading either to its apparently impressive performance, or its failure by blockage. The general advice given to users by filter manufacturers that filters will only pass limited quantities of metal is seen to be influenced by similar experience. The author has not found any limit to the volume of metal that can be put through a filter without danger of blockage, provided the metal is sufficiently clean and the front end of the filling system is designed to perform well.

Thus the secret of producing good castings using a filter is to team a good front end of the filling system together with the filter. (If the remainder of the filling system design is good, this will, of course, help additionally.) Little oxide is then entrained, so that the filter appears to do little filtration. However, it is then fully enabled to serve a valuable role as a flow rate control device. The beneficial action of the filter in this case is probably the result of several factors:

- (i) The reduction in velocity of the flow (provided an appropriately sized cross-section area channel is provided downstream of course). This is probably the single most important action of the filter. However, there are other important actions listed below.
- (ii) Reduces the time for the back-filling of the sprue, thereby reducing entrainment defects from this source.
- (iii) Smooths fluctuations in flow.
- (iv) Laminizes flow, and thus aids fluidity a little.
- (v) The freezing of part of the melt inside the filter by the chilling action of the filter (as

predicted for Al alloys in ceramic foam filters by computer simulation and found by experiment by Gebelin and Jolly 2003) may be an advantage, because this may act to restrict flow, and so to reduce delivery from the filter in its early moments. The subsequent re-melting of the metal as more hot metal continues to pass through the filter will allow the flow to speed up to its full rate later during filling. (Interestingly, this advantage did not apply to preheated ceramic moulds where the preheat was sufficient to prevent any freezing in the filter).

There are different types of ceramic block filter.

- (i) Foam filters made by impregnation of open-cell plastic foams with a ceramic slurry, squeezing out the excess slurry, and firing to burn out the plastic and develop strength in the ceramic. The foam structure consists of a skeleton of ceramic filaments and struts defining a network of interconnecting passageways.
- (ii) Extruded forms that have long, straight, parallel holes. They are sometimes referred to as cellular filters.
- (iii) Pressed forms, again with long, straight but slightly tapered holes. The filters are made individually from a blank of mouldable clay by a simple pressing operation in a two-part steel die.
- (iv) Sintered forms, in which crushed and graded ceramic particles are mixed with a ceramic binder and fired.

In all types the average pore size can be controlled in the range 2–0.5 mm approximately, although the sintered variety can achieve at least 2–0.05 mm. Insufficient research (other than that funded by the filter manufacturers!) has been carried out so far to be sure whether there are any significant differences in the performance between them. An early result of Khan *et al.* (1987) found that the fatigue strength of ductile iron was improved by extruded cellular filters, but that the foam filters were unpredictable, with results varying from the best to the worst. Their mode of use of the filters was less than optimum, being blasted by metal in the entrance to the runner, and with no back protection for the melt. (We shall deal with these aspects below.) The result underlines the probable unrealized potential of both types, and reminds us that both would almost certainly benefit from the use of recent developments. In general, we have to conclude that the published comparisons made so far are, unfortunately, often not reliable.

For aluminium alloys the results are less controversial, because the filters are highly effective in removing films which have, of course, a powerful effect on mechanical properties. Mollard and Davidson (1978) are typical in their findings that the strength of Al-7Si-Mg alloy is improved by 50 per cent, and elongation to failure is doubled. This kind of result is now common experience in the industry.

For some irons and steels, where a high proportion of the inclusions will be liquid, most filter materials are expected to be wetted by the inclusion so that collection efficiency will be high for those inclusions. Ali *et al.* (1985) found that for alumina inclusions in steel traversing an alumina filter, once an inclusion made contact with the filter it became an integral part of the filter. It effectively sintered into place; despite the fact that both inclusion and filter are solid at the temperature of operation, they behave as though they are 'sticky'. This behaviour is likely to characterize many types of inclusion at the temperature of liquid steel.

In contrast with this, Wieser and Dutta (1986) find that whereas alumina inclusions in steel are retained by an alumina filter, even up to the point at which it will clog, deoxidation of steel with Mn and Si produces silica-containing products that are not retained by an extruded zirconia spinel filter. These authors also tested various locations of the filter, discovering that placing it in the pouring basin was of no use, because it was attacked by the slag and dissolved!

Although these results might have been influenced by the rapid flow rates that appear to have been used in this work, it is a warning that filtration efficiency is likely to be strongly dependent on inclusion and filter types. Ali *et al.* (1985) confirms this strong effect of velocity, finding only at very low velocities measured in mm s^{-1} was a high level (96 per cent) of filtration achieved in steel melts.

Block filters are more expensive than cloth filters. However, they are easier to use and more reliable. They retain sufficient rigidity to minimize any danger of distortion that might result in the by-passing of the filter. It is, however, important to secure a supply of filters that are manufactured within a close size tolerance, so that they will fit immediately into a print in the sand mould or into a location in a die, with minimal danger of leakage around the sides of the filter. Although all filter types have improved in this respect over recent years, the foam filter seems most difficult to control, the extruded is intermediate, whereas the pressed filter exhibits good accuracy and reproducibility as a result of it being made in a steel die; residual

variation seems to be result of poor control of shrinkage on firing.

Leakage control

It is essential to control the leakage past the filter. There are various techniques.

- (i) A seating of a compressible gasket of ceramic paper. This approach is useful when introducing filters into metal dies, where the filter is held by the closing of the two halves of the die. The variations in size of the filters, and the variability of the size and fit of the die parts with time and temperature, which would otherwise cause occasional cracking or crushing of the rather brittle filter, are accommodated safely by the gasket.
- (ii) Moulding the filter directly into an aggregate (sand) mould. This is achieved simply by placing the filter on the pattern, and filling the mould box or core box with aggregate in the normal way. The filter is then perfectly held. In greensand systems or chemically bonded sands the mould material seems not to penetrate a ceramic foam filter more than the first pore depth. This is a smaller loss than would be suffered when using a normal geometrical print. However, the technique often requires other measures such as the moulding of the filter into a separate core, or the provision of a loose piece in the pattern to form the channel on the underside of the filter.

There are other aspects of the siting of filters in running systems that are worth underlining.

- (i) Siting a filter so that some metal can flow by (into a slag trap for instance) prior to priming the filter is suggested to have the additional benefit that the preheat of the filter and the metal reduces the priming problem associated with the chilling of the metal by the filter (Wieser and Dutta 1986). An example is seen in Figure 2.56d.
- (ii) The area of the filter needs to be adequate. There is much evidence to support the fact that the larger the area (thereby giving a lower velocity of flow through the filter) the better the effectiveness of the filter. For instance, if the filter area is too small in relation to the velocity of flow then the filter will be unable to retain foreign matter: the force of the flow will strip away retained films like sheets from the washing line in a hurricane; particles and droplets will follow a similar fate.




Figure 2.56 Ceramic foam filter printing
 (a) conventional; (b) improved early back
 protection of filter by melt; (c) tangential in-line
 filtration in runner; (d) three views of tangential
 transverse runners with reduced fall and well volume.

- (iii) Many filter placements do not distribute the flow evenly over the whole of the filter surface. Thus a concentrated jet is unhelpful, being equivalent to reducing the active area of the filter. The tangential placement of a filter can also be poor in this respect, since the flow naturally concentrates through the farthest portion of the filter. This is countered by tapering the tangential entrance and exit flow channels as illustrated in Figure 2.57b. The provision of a bubble trap reduces the effectiveness of the taper, but the presence of the trap is probably worth this sacrifice. (If the trap is not provided, bubbles arriving from entrainment in the basin or sprue gather on the top surface of the filter. When they have accumulated to occupy almost the whole of the area of the filter the single large bubble is then forced through the filter, and travels on to create severe problems in the mould cavity. The trap is expected to be similarly useful for the diversion of slag from the filter face during the pouring of irons and steels.)

Mutharasan *et al.* (1981) find that the efficiency of removal of TiB_2 inclusions from liquid aluminium increased as the velocity through the filter fell from about 10 mm s^{-1} to 1 mm s^{-1} . Later, the same authors found identical behaviour for the removal of up to 99 per cent of alumina inclusions from liquid steel (Mutharasan *et al.* 1985). However, it is to be noted that these are extremely low velocities, lower than would be found in most casting systems. In the work by Wieser and Dutta (1986) on the filtration of alumina from liquid steel, somewhat higher velocities, in the range $30\text{--}120 \text{ mm s}^{-1}$, are implied despite the use of filter areas up to ten times the runner area in an attempt to obtain sufficient slowing of the rate of flow. Even these

flow velocities will not match most running systems. These facts underline the poverty of the data that currently exists in the understanding of the action of filters.

Wieser and Dutta go on to make the interesting point that working on the basis of providing a filter of sufficient size to deal with the initial high velocity in a bottom-gated casting, the subsequent fall in velocity as the casting fills and the effective head is reduced implies that the filter is oversize during the rest of the pour. However, this effect may be useful in countering the gradual blockage of the filter in steel containing a moderate amount of inclusions.

Use of filters in running systems

In general the correct location for the filter is near the entrance to the runner, immediately following the sprue. The resistance to penetration of the pores of the filter by the action of surface tension is an additional benefit, delaying the entry into the filter until the sprue has at least partially filled. The frictional resistance to flow through the filter once it is operational provides a further contribution to the reduction in speed of the flow. This frictional resistance has been measured by Devaux (1987). He finds the head loss to be large for filters of area only one or two times the area of the runner. He concludes that whereas a filter area of twice the runner area is the minimum size that is acceptable for a thick-section casting, the filter area has to be increased to four times the runner area for thin-section castings. The pressure drop through filters is a key parameter that is not known with the accuracy that would be useful. Midea (2001) has attempted to quantify this resistance to flow but used only low flow velocities useful for only small castings. A slight improvement is available with Lo and Campbell (2000) who study flow up to 2.5 m s^{-1} . Even so, at this time the author regrets that it remains unclear how these measurements can be used in a design of a running system. A clear worked example would be useful for us all.

The filter positioned at the entrance to the runner also serves to arrest the initial splash of the first metal to arrive at the base of the sprue. At the beginning of the runner the filter is ideally positioned to take out the films created before and during the pour. The clean liquid can be maintained relatively free from further contamination so long as no surface turbulence occurs from this point onwards. This condition can be fulfilled if:

- (i) The melt proceeds at a sufficiently low velocity and/or is sufficiently constrained




Figure 2.57 (a) Concentration and reverse flow in a foam filter; (b) tapered inlet and exit ducting to spread flow.

by geometry to prevent entrainment (this particularly includes the provision to eliminate jetting from the rear face of the filter). A low velocity will be achieved if the cross-sectional area of the runner downstream of the filter is increased in proportion to the reduction in speed provided by the filter.

- (ii) Every part of the subsequent journey for the liquid is either horizontal or uphill. The corollary of this condition is that the base of the sprue and the filter should always be at the lowest point of the running system and the casting. This excellent general rule is a key requirement.

Tangential placement

Filters have been seen to be open to criticism because of their action in splitting up the flow, thereby, it was thought, probably introducing additional oxide into the melt. There is some truth in this concern. A preliminary exploration of this problem was carried out by the author (Din and Campbell 1994). Liquid Al alloy was recorded on optical video flowing through a ceramic foam filter in an open runner. The filter did appear to split the flow into separate jets; a tube of oxide forming around each jet. However, close observation indicated that the jets recombined about 10 mm downstream from the filter, so that air was excluded from the stream from that point onwards. The oxide tubes around the jets appeared to wave about in the eddies of the flow, remaining attached to the filter, like weed attached to a grill across a flowing stream. The study was repeated and the observations confirmed by X-ray video radiography. The work was carried out at modest flow velocities in the region of 0.5 to 2.0 ms^{-1} . It is not certain, however, whether the oxides would continue to remain attached if speeds were much higher, or if the flow were to suffer major disruption from, for instance, the passage of bubbles through the filter.

What is certain is the damage that is done to the stream after the filter if the melt issuing from the filter is allowed to jet into the air. Loper and co-workers (1996) call this period during which this occurs *the spraying time*. This is so serious a problem that it is considered in some detail below.

Unfortunately, most filters are placed transverse to the flow, simply straight across a runner (Figure 2.56a) and in locations where the pressure of the liquid is high (i.e. at the base of the sprue or entrance into the runner). In these circumstances, the melt shoots through a straight-through-hole type filter almost as though the filter was not present, indicating the such filters are not particularly effective when used in this way. When a foam filter suffers a similar direct impingement,

penetration occurs by the melt seeking out the easiest flow paths through the various sizes of interconnected channels, and therefore emerges from the back of the filter at various random points. Jets of liquid project from these exit points, and can be seen in video radiography. The jets impinge on the floor of the runner, and on the shallow melt pool as it gradually builds up, causing severe local surface turbulence and so creating dross. If the runner behind the filter is long or has a large volume, the jetting behaviour can continue until the runner is full, creating volumes of seriously damaged metal.

Conversely, if the volume of the filter exit channel is kept small, the volume of damaged melt that can be formed is now reduced correspondingly. Although this factor has been little researched, it is certain to be important in the design of a good placement for the filter. Loper *et al.* (1996) realized this problem, describing the limited volume at the back of the filter as a hydraulic lock, the word lock being used in a similar sense to a lock on an inland waterway canal.

Figure 2.56b shows an improved geometry that enables the back of the filter to be covered with melt quickly. Figure 2.56c shows an improved technique, placing the block filter tangentially to the direction of flow. The tangential mode has the advantage of the limitation of the exit volume from the filter, and providing a geometrical form resembling a sump, or lowest point, so that the exit volume fills quickly. In this way the opportunity for the melt to jet freely into air is greatly reduced so that the remainder of the flow is protected. A further advantage of this geometry is the ability to site a bubble trap over the filter, providing a method whereby the flow of metal and the flow of air bubbles can be divided into separate streams. The air bubbles in the trap are found to diffuse away gradually into sand moulds. For dies, the traps may need to be larger.

An additional benefit is that the straight-through-hole extruded or pressed filters seem to be effective when used tangentially in this way. A study of the effectiveness of tangential placement in the author's laboratory (Prodham *et al.* 1999) has shown that a straight-through-pore filter could achieve comparable reliability of mechanical properties as could be achieved by a relatively well-placed ceramic foam filter (Sirrell and Campbell 1997).

Adams (2001) draws attention to the importance of the flow directed downwards through the filter. In this way buoyant debris such as dross or slag can float clear. In contrast, with upward flow through the filter the buoyant debris collects on the intake face of the filter and progressively blocks the filter.

The tapering of both the tangential approach and the off-take from the filters further reduces the volume of melt, and distributes the flow through the filter more evenly. In the absence of these wedge-like features, only the far side of the filter carries the main flow, whereas the side nearest the upstream end is redundant, experiencing a circulating flow in the reverse direction (Figure 2.57).

Direct pour

Sandford (1988) showed that a variety of top pouring could be used in which a ceramic foam filter was used in conjunction with a ceramic fibre sleeve. The sleeve/filter combination was designed to be sited directly on the top of a mould to act as a pouring basin, eliminating any need for a conventional filling system. In addition, after filling, the system continued to work as a feeder. This simple and attractive system has much appeal.

Although at first sight the technique seems to violate the condition for protection of the melt against jetting from the underside of the filter, jetting does not seem to be a problem in this case. Jetting is avoided almost certainly because the head pressure experienced by the filter is so low, and contrasts with the usual situation where the filter experiences the full blast of flow emerging from the base of the sprue.

Sandford's work illustrated that without the filter in place, direct pour of an aluminium alloy resulted in severe entrainment of oxides in the surface of a cast plate. The oxides were eliminated if a filter was interposed, and the fall after the filter was less than 50 mm. Even after a fall of 75 mm after the filter relatively few oxides were entrained in the surface of the casting. The technique was further investigated in some detail (Din *et al.* 2003) with fascinating results illustrated in Figure 2.58. It seems that under conditions used by the authors in which the melt emerging from the filter fell into a runner bar and series of test bars, some surface turbulence was suffered, and was assessed by measuring the scatter of tensile test results. The effect of the filter acting purely to filter the melt was seen to be present, but slight. The castings were found to be repeatable (although not necessarily free from defects) for fall distances after the filter of up to about 100 mm, in agreement with Sandford. Above a fall of 200 mm reproducibility was lost (Figure 2.58a, b).

This interesting result explains the mix of success and failure experienced with the direct pour system. For modest fall heights of 100 mm or so, the filter acts to smooth the perturbations

to flow, and so confers reproducibility on the casting. However, this may mean 100 per cent good or 100 per cent bad. The difference was seen by video radiography to be merely the chance flow of the metal, and the consequential chance location of defects.

The conclusion to this work was a surprise. It seems that direct pour should not necessarily be expected to work first time. If the technique were found to make a good casting it should be used, since the likelihood would be that all castings would then be good. However, if the first casting was bad, the site of the filter and sleeve should simply be changed to seek a different pattern of filling. This could mean a site only a few centimetres away from the original site. The procedure could be repeated until a site was found that yielded a good casting. The likelihood is that all castings would subsequently be good.

However, the technique will clearly not be applicable to all casting types. For instance, it is difficult to see how the approach could reliably produce extensive relatively thin-walled products in film-forming alloys where surface tension is not quite in control of the spread of metal in the cavity. For such products the advance of the liquid front is required to be steady, reproducible and controlled. Bottom gating in such a case is the obvious solution. Also, the technique works less well in thicker section castings where the melt is less constrained after its fall from the underside of the filter. Figure 2.58c illustrates the fall in reliability of products as the diameter of the test bars increases above 20 mm. Equivalent results would be expected for the increase of plate sections above about 10 mm.

Even though the use and development of the direct pour technique will have to proceed with care, it is already achieving an important place in casting production. A successful application to a permanent moulded cylinder head casting is described by Datta and Sandford (1995). Success here appears to be the result of the limited, and therefore relatively safe fall distance.

Flow rate data through the filter/sleeve combination is necessary to predict the pour times of castings. Such data has been measured by Bird (1989). His results presented here (Figure 2.58d) have been rationalized to apply to 50 mm diameter ceramic foam filters, and relate to Al-Si alloys cast at 720 °C. Clearly, filters of different sizes will pass correspondingly more or less melt per second proportional to their areas, assuming their thickness and pore sizes are sufficiently close. The pores' diameters of approximately 1 and 2 mm in Figure 2.58d refer to the 'pore per inch' categories 20 ppi and 10 ppi respectively.