CHAPTER

DIMENSIONAL ACCURACY

18

The casting which we make is, of course, never quite perfect in terms of size and shape. To allow for this, tolerances are quoted on engineering drawings. So long as the casting is within tolerance, it will be acceptable.

Some reasons for the casting being out of tolerance include elementary mistakes like the patternmaker planting the boss in the wrong place. This leads to an obvious systematic error in the casting, and is easily recognised and dealt with by correcting the pattern. It is an example of those errors which can be put right after the first sample batch of castings is made and checked. (Even in this simple case, care may be needed if great accuracy is required, as is explained a little later.)

Another common systematic error in castings is the wrong choice of patternmaker's contraction allowance. The contraction of the casting during cooling in the mould is often of the order of 1–2%. However, it depends on several factors, particularly strongly on the strength of the mould and the cores. For instance, in an extreme case, a perfectly rigid mould will fix the casting size; in such a situation, the casting simply would have to stretch during cooling because it would be prevented from taking its natural course of contracting. Such a situation is common for thin wall castings made in steel moulds, such as high pressure die castings (HPDCs). However, in other casting processes in which the mould can yield to some extent, there is often uncertainty. The choice of contraction allowance before the making of the first casting sometimes has to be decided in the absence of previous similar castings that would have provided a guide. Thus, the chosen value for contraction is often not exactly right. This point is taken up at length in a later section on 'net shape', with recommendations on how to live with the problem.

Other errors are less easily dealt with. These are random errors. No two castings are precisely alike. The same is true for any product, including precision-machined parts. The International Standards Organisation (ISO) Standard (1984) for casting tolerances indicates that, although different casting processes have different capabilities for precision, in general the inaccuracies of castings grow with increasing casting size, and the standard therefore specifies increasing linear tolerances as linear dimensions increase. (Nevertheless, it is worth pointing out that the corresponding percentage tolerance actually falls as casting size increases.) Although other work on the tolerance of castings suggests that the ISO standard has considerable potential for further improvement (Reddy et al., 1988), it is the only European standard at this time. This also has to be lived with.

Because of the effects of random errors being superimposed on systematic errors, where great precision is required it is of course risky to move the boss into an apparently correct location simply after the production of the first trial casting. Figure 18.1 illustrates that the random scatter in positions might mean that the boss appeared to be in the correct place first time if the casting happened to be casting 1 in Figure 18.1, or might have been more than twice as far out of place, compared with its average position, if the first casting had been number 2 in Figure 18.1. A sample of at least two or three castings is really needed, and preferably 10 or 100. The mean boss location and its standard deviation from the mean position can then be known accurately and the appropriate actions taken.

At the present time, it will be of little surprise to note that such exemplary action is not common in the industry. This is because companies are not generally equipped with a sufficient number of fast, automated three-coordinate measuring machines. As such standards of measurement become more common, so the attainments in terms of accuracy of castings will increase. However, it is a pleasure to note that advances in metrology are occurring fast and furiously. In particular, the use of lasers and computers has revolutionised measurement of the exterior features of castings. For those fortunate enough to be able to afford it, a similar revolution has happened in X-ray radiographic tomography, so that

Complete Casting Handbook. http://dx.doi.org/10.1016/B978-0-444-63509-9.00018-2

Copyright © 2015 John Campbell. Published by Elsevier Ltd. All rights reserved.

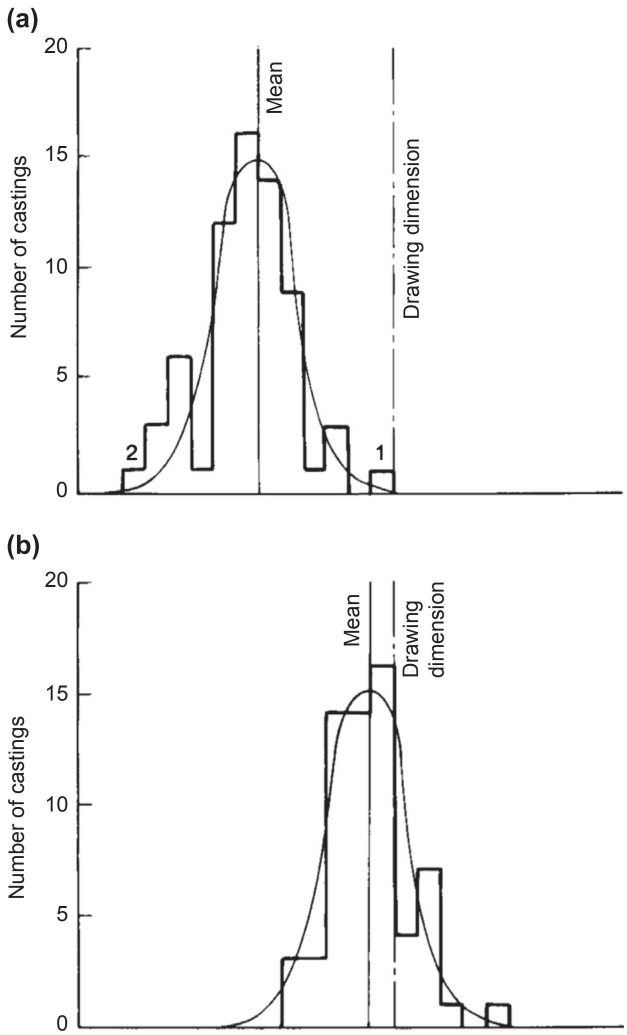

FIGURE 18.1

Statistical distribution of casting dimensions (a) before and (b) after pattern development.

Based on Osborn (1979).

three-dimensional views of the interiors of castings can now be measured with confidence to ensure that hidden wall thicknesses are correct, and that interior cores have not moved or distorted during casting. A further benefit is to ascertain that all core sand has been removed even from inaccessible pockets inside the casting.

Turning now to the accuracy of castings themselves: one of the fundamental requirements for an accurate casting is an accurate rigid pattern and an accurate rigid mould. This may seem self-evident, but it is comparatively rarely achieved.

Figure 10.2 shows several process routes for the manufacture of castings, but in Figure 18.2 they are assessed in terms of their potential for accuracy, and this potential quantified for one particular casting in Figure 18.3. In principle, only the

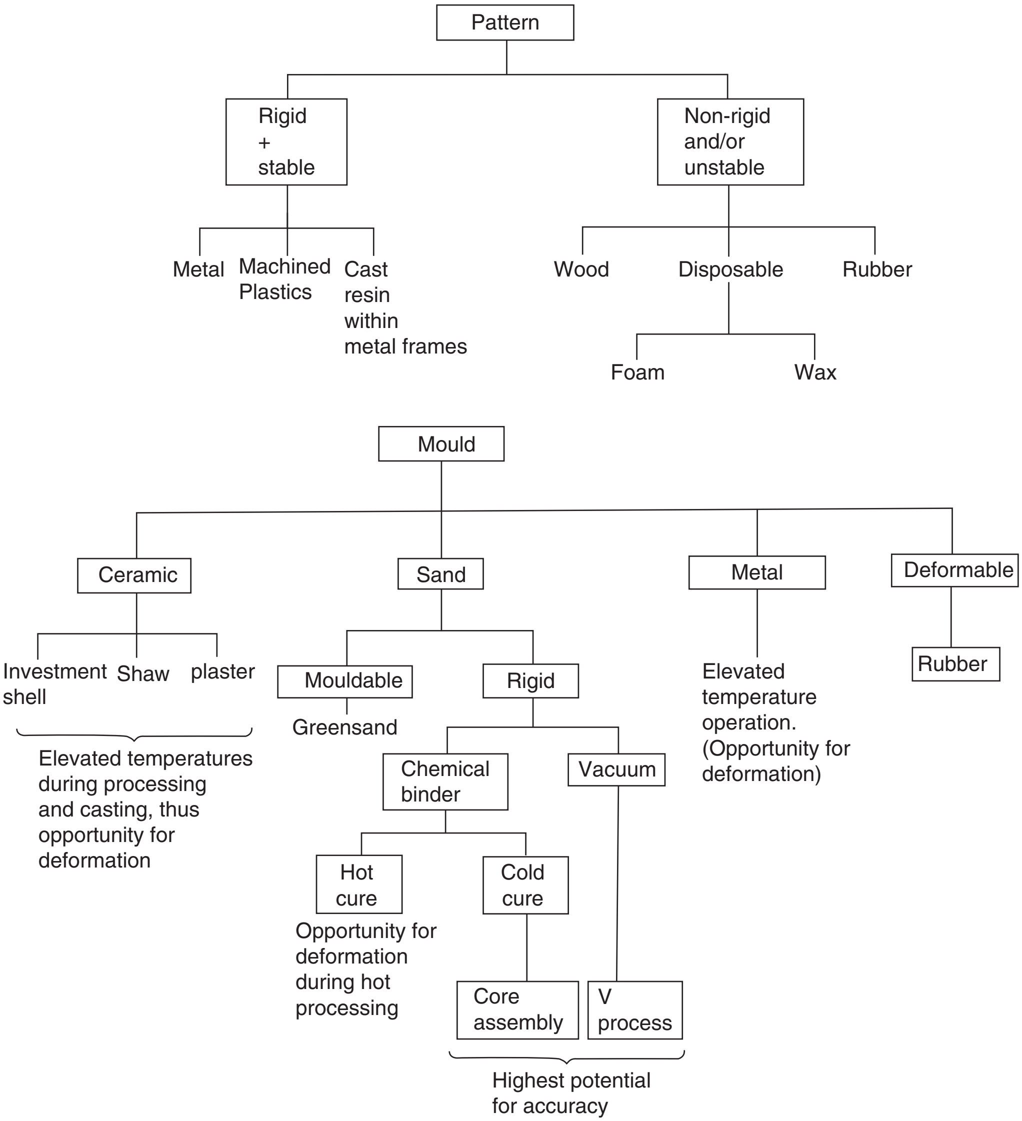

FIGURE 18.2

Patterns and moulds categorised according to their potential for accuracy.

rigid patterns in combination with the hard aggregate mould routes (core assembly and V process) meet all the requirements of accuracy, rigidity, room temperature formation and curing against the patternwork, and room temperature assembly of moulds and cores.

This is not to say that other processes are inappropriate. Some will be adequate in accuracy and have the benefit of low cost. Others will be available locally to the customer, and the foundry may have the capacity to accept the order at that time. There are many factors that influence the choice of an appropriate manufacturing route and choice of manufacturer.

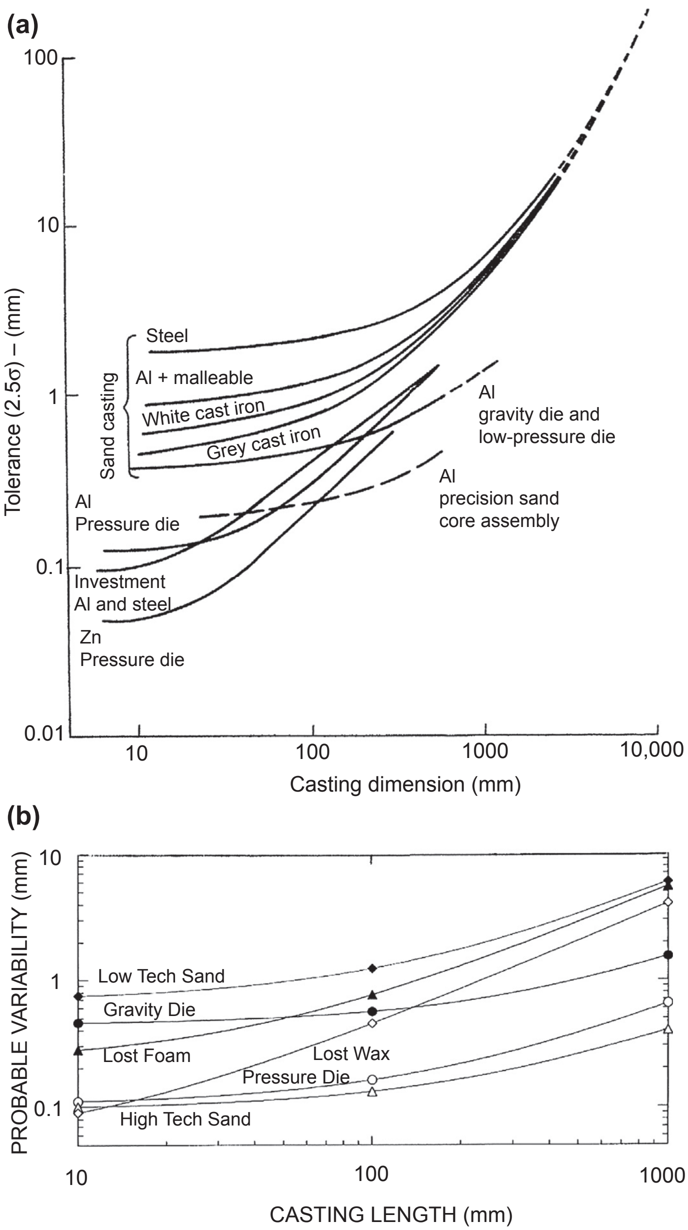

FIGURE 18.3

The average tolerance (taken as 2.5σ) exhibited by various casting processes (sand, and gravity and low-pressure die from IEE TS7I; investment from BICTA 1990; pressure die from ZDA 1990; sand core assembly data are tentative, from early production experience). Relation of casting dimensions with (a) potential variability of a hollow 5 mm wall thickness casting of overall size 250 × 200 × 200 mm, and (b) probable variability of a similar but solid casting (Campbell, 2000b).

The ISO standard for casting accuracy gives the general trend of increase in the size of random errors as casting size increases. However, the casting designer and engineer require much more detailed knowledge of the sources of individual contributions to the final total error. The remainder of this chapter is an examination of these contributions. The reduction of these errors allows the production of castings which can considerably exceed the minimum accuracy requirements specified in the ISO.

The use of computers has of course greatly aided the quest for more accurate and reproducible tooling production and replication. Computer-aided design/computer-aided manufacture has itself benefited from the rationalisation and standardisation of data transfer in the form of Standard Triangulation Language, International Graphical Exchange Standard, and Standard for Exchange of Product Information. The world-wide adoption of such standards has allowed tool sourcing to become a world-wide activity.

18.1 THE CONCEPT OF NET SHAPE

A ‘net shape’ product is one whose shape requires no further modification and thus is ready for use. The concept is one which is usually applied only to parts which are required to assemble and fit closely with other parts.

Most soft plastic toys are made net shape, but the shape is often not critical and the concept has no relevance. However, it is critical in the case of components such as Lego, where each injection moulded plastic part is required to fit accurately into other parts without any further processing; clearly, in this case, no machining can be contemplated on grounds of cost. The component must be net shape.

For many metal components, with perhaps the exception of some powder forming routes, the part can rarely be finished exactly to the required final tolerance in a single forming operation. Thus, in general, a forming operation such as forging or casting is carried out to produce a ‘near-net shape’ product, which is subsequently brought into the required tolerance by a finishing operation of some kind.

If we try to evaluate the various casting processes for their capability in the production of net or near-net shape products, the problem resolves to ‘what accuracy can the various casting processes achieve?’

Answers to this question are tasks not easily limited to a manageable scale. For instance, accuracy is required in terms of straightness, flatness, concentricity, etc. Here we shall consider only the ability to reproduce a length.

Sources of error come in various forms: (1) random errors as a result of variations in processing and (2) systematic errors of many types, including the boss in the wrong place, and more general problems that apply uniformly to parts or to the whole of the product. For instance, these sources include such important features as the mismatch between mould halves. Also included in this category are features such as the thickness of a mould coating when applied uniformly (well or badly) to the whole or part of a mould by spraying or dipping. Systematic errors occur resulting from expansion or contraction during processing, and which are thus a function of the size of the component.

There is also the fundamental difficulty associated with the fact that an intrinsically poor processing route could be carried out with extreme care by one manufacturer, giving a satisfactory product in spite of the technology. This contrasts with the situation in which an intrinsically reproducible process is carried out with ineptitude or even crass irresponsibility by another, therefore giving an unacceptable product. Thus, although in this study the fundamental capability of the process will be emphasised, the ability of the manufacturer to overcome deficiencies in the process by intelligent and diligent effort must not be underestimated. Even so, the view is taken here that processes which are intrinsically reproducible are to be favoured over those which are not. Also, some attempt will be made to place limits on the achievements of both good and bad practice to assess the extent of the tolerance problem.

The only previous comparative study of the accuracy of the various casting processes has been carried out by the author (2000). However, because of the paucity of data on the accuracy of castings produced by different processes (notwithstanding the existence of the ISO, 1984 standard and one notable experimental attempt by the Institute of British Foundryment – IBF, 1979), he found himself driven to making estimates of accuracy of the different processes based merely on his own experience. This, clearly, is hardly satisfactory, but is judged to be better than nothing at this stage, and therefore comes with no apologies. This section considering the problems of net shape is based on the original article, and is offered as a preliminary study, outlining the concepts involved, and providing a framework in which the mechanics of the problem can be understood, and into which better data can be fitted as it becomes available.

18.1.1 EFFECT OF THE CASTING PROCESS

Inaccuracy can be introduced into the product at several stages of the casting process which typically involves the sequential production of a series of shapes, one formed intimately against the other, thus producing a series of positive and negative forms which finally give the desired positive shape. As the number of steps increases, there is an increasing chance to introduce error. The main processes are considered in order of increasing complexity, and by implication, vulnerability to error, later.

Exterior shapes

Die casting is the most direct of the casting processes. Here, a metal mould (a near-net-negative of the required shape) is filled with liquid metal which is allowed to solidify to give the final desired positive form. Pressure die casting involves the injection of the metal at high speed and high pressure into the die cavity, whereas low pressure die casting and gravity die casting use similar cast iron or steel dies which are covered with a protective refractory coating, and which are filled relatively quiescently.

Sand casting involves the additional step of having a positive pattern which is used to produce a negative sand mould to produce the positive cast form. Two main types of sand mould are used: (1) greensand and (2) chemically bonded sand. Although sand has a poor image, and is in fact often used badly and under poor control, it has considerable potential for reproducibility. It often does not require a protective mould coat and thus retains the accuracy of the pattern. Because the pattern never contacts the liquid metal it can be long-lived if it suffers little wear, retaining its original accuracy for up to 10 times longer than die casting processes.

Lost wax (i.e. investment) casting involves a further additional step, because a negative die is required, usually accurately machined from an aluminium alloy. This die to form the wax pattern is long-lived and retains its accuracy well. The die is filled with liquid wax to make a positive pattern. A ceramic shell mould is formed around this, and the wax melted out, to produce a negative into which metal is finally poured to create the final positive shape. Significant scope for variability exists in (1) the expansion and/or distortion of the wax pattern, (2) the stresses of the dewaxing operation, (3) the firing of the shell where sintering and shrinkage of the shell occurs, (4) the expansion and phase changes in the shell which occur on casting (particularly when pouring high temperature alloys and steels) and (5) the variable restraint of the mould on cooling as a result of the prior variations in chemistry and sintering of the ceramic shell.

Lost foam casting is some ways similar to lost wax, but where the disposable positive pattern is formed from polystyrene expanded inside a machined aluminium die. Again, the die is accurate and the least subject to errors of wear compared with all the casting processes. The foam pattern may be subject up to 0.8% shrinkage (depending on the type of foam (Brown, 1992)) and may require being assembled and glued together from individual foam parts. The whole is then coated with a ceramic slurry which is allowed to dry. The assembly is finally placed in a box, loose dry sand is vibrated into place around the pattern to support it during casting and extract the heat during freezing. This action of pouring and vibrating sand into place in this way usually introduces measurable distortion of the flimsy polystyrene pattern. The metal is poured into the foam, displacing it to form the final positive shape. The mould may, of course, distort during the casting and cooling process.

Interior shapes

Hollow parts of castings which can be formed by simple withdrawable shapes can be introduced into die castings. In this case, the core is a permanent feature of the die tooling and is usually made from steel, and is sometimes water cooled. Pressure die cast engine blocks are among the most ambitious aluminium alloy castings made by this process.

More complex interior shapes cannot be withdrawn and thus require to be formed by disposable cores. The most common type is the resin-bonded sand core, in which the resin is designed to break down after casting, allowing the sand to flow out of the cored cavity. Alternatively, in a few instances, cores can be dissolved out (such as salt core by water, or silica core by hydrofluoric acid in the case of Ni-based turbine blades).

The core is supported inside the mould cavity by its prints. These are (positive) core extensions which locate in (negative) print locations in the mould or die. For sand casting, the location of the core can be precise because assembly is all at room temperature, and because the core-to-mould interface is usually kept clear of uncertainties such as a variable thickness of core coating. For gravity die casting, the sand core is often poorly located because the thermal distortion of the die causes the print negatives to be out of location. In addition, the die has a variable thickness of die coat which has been sprayed on, some of which finds its way onto the print locations, causing further deterioration of their precision. Worse still, especially if core loading of the die is lengthy or is delayed, the heat of the die can sometimes soften, or even cause to break down the resin binder in the sand core, causing the core prints to disintegrate and the core to sag.

Once surrounded by liquid metal, the core will be subject to buoyancy forces which will tend to make it float. This is not so bad in liquid magnesium where the density difference with silica sand is near zero. Similar neutral buoyancy applies for the aluminium/zircon sand system. Buoyancy becomes increasingly problematic for the common systems such as silica sand together with liquid aluminium or iron or other dense metals such as copper and steel, the latter systems more particularly so because of the higher casting temperatures involved. Flotation forces have to be withstood by adequate mechanical support of the core, usually by good print design and location, or, often less desirably, by internal steel reinforcement or metallic supports (chaplets) which are cast into the product as permanent features. To ensure good fusion between the insert and the casting is not always easy. Sometimes such features are subsequently machined out.

18.1.2 EFFECT OF EXPANSION AND CONTRACTION

In greensand moulding systems the sand is bonded automatically by the compaction of the sand. However, at compaction pressures above about 1 MPa (10 bar) the sand mould deforms elastically, and, on withdrawing the pattern, is subject to 'spring back'. This general distortion of the mould leads to numerous difficulties relating to core assembly and mould closure. This, to the author's knowledge, is the only common moulding system exhibiting distortion resulting from stress. Most distortions in casting processes arise because of thermal expansion as is discussed later.

When the metal first enters the mould, the mould, and more particularly, the cores, heat up and expand. Thus the mould cavity enlarges and will probably suffer some distortion. This is reasonably reproducible in sand moulds because the pouring temperature is normally under good control (usually ±10°C or better) and the moulding sand is always close to room temperature. Thus the starting conditions are usually reproducible. This is less true for the various processes using metal moulds, where the die temperature of perhaps 300–400°C is often poorly controlled, varying by as much as ±100°C or more.

As the casting cools, it contracts, with the result that some parts of the casting are subject to tensile extension or compression because of geometrical constraint of the mould. The constraints exerted by the disposable sand mould are perhaps less severe than those of the metal die, but in general the casting stays in the sand mould longer and the constraint therefore operates for longer.

Although an attempt will be made to estimate these expansions and contractions, the final result corresponds to the patternmaker's contraction allowance. This is the value based on hundreds of years of experience by patternmakers and toolmakers, and is the allowance they have to provide, making the pattern a little larger than the final casting. Unfortunately, the factor is sensitive to geometry of the casting, and mistakes are often made in the correct choice of the allowance, which can vary between 0% and 1.3% for aluminium castings and up to 2.4% for steel castings. Some recent work in the author's laboratory have highlighted the uncertainties of this factor for cast aluminium and cast irons (Nyichomba and Campbell, 1998). Estimates of this allowance are made in Section 18.5.

In the pressure die casting process, the casting is normally thin-walled and thus rather weak. As it cools in the rigid steel die, contracting onto projections of the geometry, it is therefore forced to stretch plastically. The size of the casting is thus mainly controlled by the time to ejection; larger castings are produced if ejected later.

The hot casting is cooled to room temperature in a variety of ways. This may occur either as a series of individual castings on a cooling conveyor, or as a heap in a bin, or by an immediate quench into water. These post casting operations are likely to create different patterns of distortion. Likewise, heat treatment, even natural ageing, affects many products, particularly heat treatable aluminium alloys. The latter grow slightly (of the order of 0.05%, i.e. 0.5 mm per metre) as hardening precipitates form.

18.1.3 EFFECT OF CAST ALLOY

Zinc castings are most commonly supplied as pressure die castings. As we have seen earlier, the thermal distortions of both die and casting are the least of any of the casting processes because (1) the casting temperatures are low, creating low expansion problems for the die and minor contraction problems for the casting; (2) the dies are accurately machined from steel; (3) the dies use no protective die coating and (4) are supported in a close fitting and rigid steel bolster. Thus zinc pressure die castings are intrinsically capable of meeting many net shape requirements. This is a clear-cut example of near-automatic achievement of net shape by zinc pressure die castings.

The light alloys based on magnesium and aluminium can be cast in moulds of all types. For HPDC, the dies are steel and no protective mould coat is used, so maximising accuracy. For low pressure and gravity die castings, the die is cast iron or steel but is protected by a ceramic die coat. The thickness of the coating is not easily controlled, thus limiting accuracy, plus in addition the dies themselves are rather rough and inaccurate compared to HPDC dies. The dies are also often subject to distortion partly because the die is free-standing (i.e. not supported by a surrounding steel bolster as in the case of HPDC). For sand castings, the moulds can be used without a protective coating, and thus retain their dimensions.

Cast iron is most commonly poured into greensand moulds with no protective coating. Chemically bonded sand moulds and cores on the other hand usually require a refractory wash coat to obtain an acceptable surface finish. Although cast iron dies are possible for cast iron, their use in Western Europe is limited to specialised casting operations, where, again, a die coat is required. At this time, iron dies for iron casting is more widely used in Eastern Europe.

Steel is cast into aggregate moulds. For modest sizes of products, greensand is widely used without a mould coat. Increasingly often, however, steel is being cast into chemically bonded sand moulds with a protective mould coat to achieve an acceptable finish. The coating ranges in thickness from 0.2 to 0.6 mm, being usually 0.3 ± 0.1 mm. The high temperature of steel casting leads to considerable interaction of the surface of the casting with its environment, leading to sand burn-on and oxidation, although much of these problems would be solved immediately by the application of a good mould filling system. Many large steel castings lose up to 0.8 mm or more of their surface during heat treatment because of the flaking off of oxide from the outside surface of the casting, and approximately 0.4 mm from interior surfaces. Because of these interactions, plus the application of mould coatings to prevent them, steel castings are normally the furthest removed from the net shape concept. On top of this, the rough-and-ready patterns often seen in the cast steel industry clearly do not help the attainment of accuracy. For steel components weighing up to a few kilograms, the favoured route for near net shape is therefore lost wax casting in vacuum.

18.2 MOULD DESIGN

18.2.1 GENERAL ISSUES

The problem for the casting engineer is to achieve a successful design of the mould. This problem is not to be underestimated because it requires the simultaneous solving of a list of issues including.

1. The design of the mould and core assembly can be a problem in itself. It is not uncommon to find that it is impossible to assemble the cores because of some shape feature of neighbouring cores has been overlooked. It is all too easy to stumble into such pitfalls in a complex core assembly. When the first set of cores are made from the new patternwork, with its shining new varnish and paintwork, the discovery of such ‘passing problems’, where one core will not pass another and so fit into the assembly, are greeted with embarrassment and dismay.

The other common problem for the casting engineer and toolmaker is the design of the assembly so that cores fit, in logical order, only into the drag if possible (Figure 18.4). Cores in the cope are not usually an option for horizontally parted greensand moulds, because, if the sand strength is not high, they are in danger of falling out of their prints when the cope is turned over and closed onto the drag before casting. Gluing cores into the cope is possible in the case of strong chemically bonded sand moulds. However, gluing takes time and is therefore costly and introduces the danger that any excess glue that exudes out of the join may cause a blow hole defect in the casting if it contacts the metal. In addition, glue applied to a core print may prevent the core from venting, leading to a blow defect from the core. The use of glues should therefore be avoided if at all possible.

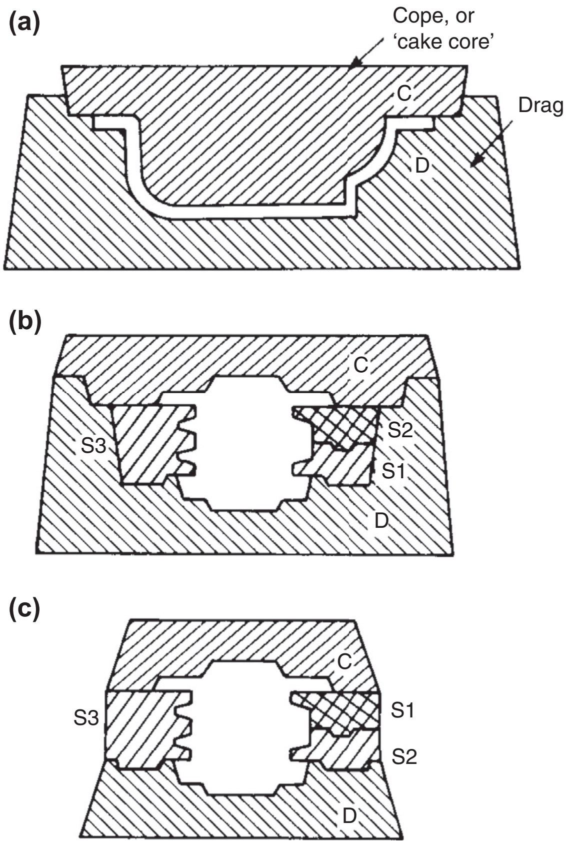

FIGURE 18.4

(a) A simple cake core and drag assembly; (b) a cope and drag with side cores, all located in the drag; (c) an apparently lower-cost alternative to (b), but resulting in loss of dimensional control.

It is common for complex core assemblies to be assembled at a separate station sited off the mould assembly line. Core assembly can then be accurate because the assembly is built up in a jig. The cores are designed to be lifted by the jig, transferring from the assembly station and lowered into the mould as a complete package. Castings that require lengthy core assembly times are not thereby allowed to slow the cycle time of the moulding line.

2. The filling system. The provision of a good filling system and its integration with the rest of the mould and core system is sometimes not easy, and in some cases the additional trouble or expense to provide a good filling system is by-passed. (The minefield of poor castings and high scrap rates is always entered for apparently good reasons.) The filling system design forms a major part of this book. It is mandatory reading. Its rules are strongly recommended to be followed in all cases.
3. The feeding system. Naturally, following the first rule for feeding, it is clearly best if feeders can be completely avoided. I attempt to avoid feeders by the use of chills, fins or pins. However, if feeders are necessary, they are required to be placed high on the casting to feed downwards under gravity. This may force the whole casting to be turned over to orient its heavy sections at the top. Another one of the key issues is to place the feeders so that they are easy to cut off or machine away subsequently.

4. The avoidance of infringement of any of the 10 rules. For instance, convection considerations might force the issue of rotating the mould through 180° after filling. This action usually confers other benefits and makes the integration of the filling and feeding systems powerfully effective and economic. It is a strategy to be recommended, although it is essential to observe the precautions listed in the section on rollover systems.

However, sometimes the solution to all these issues is not straightforward. For instance, much time may be spent attempting to solve the issues with the casting oriented in one direction, only to realise that such an orientation involves insoluble problems. The casting is then turned upside down and the exercise is repeated in the hope of a better outcome. Such experiences are the day-to-day routine of the casting engineer.

Furthermore, the complexity of the issues is not easily solved at this time by computer. There have been many such attempts, but it is fair to say at this stage such efforts have not been developed to such a degree that the professional casting expert is offered any significant help. However, of course, we can look forward to the day when the computer is sufficiently capable to provide useful solutions.

18.2.2 ASSEMBLY METHODS

The design of the mould assembly often simply involves two parts: a cope and a drag, such maximum simplicity makes for maximum accuracy. But for more complex castings the mould assembly can be complicated, requiring many parts, and requiring much discussion between the pattern shop, foundry, and casting designer to find an appropriate solution. Accuracy can now become illusive and troublesome.

The simplest form of construction of moulds, which was popular with many foundrymen, consisted of a drag with a cope in the form of a 'cake core' as shown in Figure 18.4(a). This simple construct did not allow for the placement of any useful bottom-gated running system, and the top pouring which had to be accepted as a consequence was generally acceptable for grey iron castings in greensand moulds. It did not give good results for the majority of casting alloys. The arrangement would, however, have been excellent for a counter-gravity filling technique.

As a general rule, it is useful to ensure that even in the most complicated of mould assemblies, the design of the assembly consists principally of a drag and a cope, and that all parts are interrelated via a single mould part, which will normally be the drag (Figure 18.4(b)). In this case, the assembly of cores will be in the drag, with each core located separately, and the final operation will simply consist of closing with the cope (sometimes called 'a cake core'). The cope also should have features which contact and locate directly on the drag as the key mould part.

Such simple rules are easily forgotten. Figure 18.4(c) shows how it would have been easy to save some sand by abandoning the deep drag construction, and having a core assembly which consists simply of a random heap of assorted shapes: one layer of cores tottering upon another, with finally the cope perched precariously on top. The overall accuracy of the casting now suffers from the accumulation of errors introduced by the intermediate side cores S1, S2 and S3, providing a poor and variable match between the top and bottom features of the casting.

Figure 18.4(b) shows a simple type of cope-drag arrangement with side cores, all located in the drag, apart from side core S2 which rests on S1. It was judged that the small accumulation of errors in the positioning of S2 would be acceptable in this case. If the positioning of S2 had been critical, it could still have been located in the drag by stepping the contour of the drag appropriately.

The dimensional problems which arise in setting cores are examined by Skarbinski (1971). In general, it needs to be said about the accurate printing of cores that the core print should be designed assuming that the core will be produced with errors in its size and shape, and will have to fit into a mould which will also have suffered some distortion during its manufacture. The print also sometimes has to restrict the movement of the core during mould filling because of the build-up of pressure against its front face, and because of buoyancy tending to make it lift. It may also have to allow the relative movement of the core in its print to permit the thermal expansion of the core. All this is a seemingly impossible task to achieve accurately. However, it is usually solvable within practical limits by applying the following simple rules:

1. The print requires *tolerance* where it needs to fit (i.e. must not be made size-for-size; that would have produced an interference fit. The only exception to this is the heights of cores where cores are stacked one on top of the other because in this case the accumulation of errors must be kept to a minimum).

2. The print requires *clearance* where it is not required to fit, and where expansion clearance is required.
3. The number of cores should be reduced to a *minimum*, moulding as much as possible directly in the cope and drag.

Rules often appear pedantic or even pedestrian when they are spelled out! However, the application of the rules involves much work which, unfortunately, is often neglected during the urgency involved in the design and manufacture of patternwork. Although some prints are easily designed, others require much thought, and compromises have to be carefully assessed. Every print requires individual detailed design. It is attention to details such as these that makes the difference between the inadequate and the excellent casting.

Nevertheless, these problems are eliminated if the use of the core can be avoided altogether as suggested by following the useful rule 3. Not only does the application of this principle reduce dimensional errors, but also the addition of each core involves considerable extra tooling cost, and an additional cost in the production of the casting, sometimes approaching the cost of the production of a cope or drag; it is effectively the third piece of sand to be added to the two original mould halves; thus, the costs at this stage may increase by 50% for the addition of the first core. At other times, a small core can save money by avoiding extra complexity of the tooling. Each case needs separate evaluation.

The perhaps unexpectedly high addition to total casting costs resulting from the use of cores arises from the accumulation of several minor operations, most of which are usually overlooked. For instance, the core needs to be scheduled, made, perhaps on a capital-intensive core-making machine, stored, de-flashed, retrieved from storage, transported to the moulding line, and then correctly assembled into the mould. Errors arise as a result of the incorrect core being made or transferred, or sufficient are broken in storage or transit to cause the whole process to be repeated, or its assembly into the mould gets forgotten at the last moment! Cores are therefore almost certainly more expensive than most foundry accounting systems are aware of. (The costs of chills, and of scrapped castings, are similarly illusive and therefore generally underestimated or completely overlooked.)

A further use of cores, in addition to their obvious purpose in providing detail which cannot be moulded directly, is that the running system can often be integrated behind and underneath them, the main runner and gates being located beneath side or end core(s). This is a valuable facility offered by the use of a core and should not be overlooked. In several castings the addition of a core may be for the sole purpose of providing a good running system. Such a core is often money well spent! It could make the difference between a crippling scrap rate and trauma-free production.

Where a complex array of internal cores has to be loaded into a drag, as in the case of the assembly of an automotive crankcase for instance, it is common to arrange for the cores to be preassembled in an assembly jig, and for the complete core package to be lowered into the drag. This saves time on the assembly line.

The problem with the automation of core-assembly systems is finding the core again with sufficient accuracy so that it can be picked up once again after it has been put down on, for instance, a conveyor or a storage rack. This is a difficult job for a robot because extreme accuracy is required, and the cores are often of extreme delicacy. Although this problem is being increasingly well addressed by vision systems and improved robots, clearly, one method of solving this problem is never to put the cores down in the first instance. This simple solution is powerful.

Schilling (1987) has succeeded in developing this concept with a unique system of making and assembling cores in which the cores are not released from one half of the opened core-box until the other half of the core has already been located in the core-assembly package. In this way, the cores are assembled completely automatically and with unbelievable precision. Cores are located to better than 0.03 mm, allowing them to be assembled with clearances which are so small that the cores could not be assembled by hand. In fact, the cores are sprung into place with interference fits. The rigorous application of this technique means that castings need to be designed for the process because the assembly of each core is by vertical placement over the previous core. For instance, any threading of cores in through holes in the sides of other cores, such as often occurs with port cores through the water jacket core of a cylinder head casting, is not possible. This disadvantage will limit the technique to partial application, loading some but not all cores of a cylinder head, for instance. Even this would be an important advance.

The fundamental objection to this super-accuracy technique is the requirement for the plant to run with sufficient reliability. Because the cores are never released, but always under the control of a core-box, there can be no buffer stock of cores between operations. A typical core assembly system involving this principle might involve at least 30–40 machines including core blowers and robots etc. Even if the machines were of excellent reliability, each having an up-time of 95–98%, the multiplication of all those reliabilities results in an overall reliability for the whole operation of usually less than 50%. This is a serious disadvantage that requires an order of magnitude increase in maintenance effort and skill to reduce this drawback. This substantial challenge makes this system difficult to recommend.

A final note in this section relates to cope-to-drag location. This is, of course, of primary importance. Failure to achieve good location results in a mismatch defect. For the case of precision core packages, the sand mould is not contained in a box, and thus is located directly with sand-to-sand locations. Because this is defined from the pattern-work, the location relates perfectly to the casting details, and mismatch is therefore not possible.

In foundries using moulds contained in moulding boxes, however, mismatch is unfortunately all too common and is usually the result of the use of worn pins and bushes which are used to locate the boxes. Southam (1987) analyses the effect of the errors involved in the pin and bush location system. These are numerous and serious. The pin to bush clearance is typically 0.25 mm, and given an apparently acceptable additional wear of 0.35 mm, he finds that the total possible mismatch between cope and drag moulds is as much as 1.5 mm. (Mismatch is a lateral location error, and not to be confused with the vertical precision with which cope and drag meet which is normally of the order of ±0.1 mm.)

He proposes, therefore, a completely different system, in which mould closure is carried out in a special station, where the cope and drag boxes are simply guided by wear blocks fixed to the outside edges of the box. These slide against two guides on the long side of the box, and one guide against the narrow side of the box during moulding and closing operations. The boxes are held against the guides by light pressure from springs or pneumatic cylinders. The system appears deceptively simple, but actually requires a certain amount of good engineering to ensure that it operates correctly on mould closure, as Southam describes. Although Southam calls his method the three-point registration system, it is in reality a classical six-point location system because he uses a further three points to locate the drag in a parallel plane to the cope during closure (the six-point location system is discussed in relation of casting datums as good casting rule 10).

The ability to locate cope to drag with negligible error has several benefits, which Southam lists. The maintenance and replacement of worn pins and bushings is a foundry chore and expense which is eliminated. Instead, only three guides on the closer and three each on the cope and drag pattern need to be checked, and the effect of wear of these parts on mismatch is minimal because the resultant displacement is largely self-compensating. In addition, the foundry will be capable of producing castings with thinner walls, reduced dressing, reduced machining, and improved appearance.

18.3 MOULD ACCURACY

18.3.1 AGGREGATE MOULDS

Buyers of castings have been traditionally prejudiced against sand castings as opposed to the various varieties of metal moulded castings because they have long memories. They recall, or perhaps imagine they recall, the time when most sand castings were made individually by a bench moulder or floor moulder, equipped with a wheelbarrow, watering can and shovel to produce his greensand mix, together with an ancient and battered wood pattern that had seen better days. Working practices encouraged additional variability, for example by rapping or vibrating the pattern to ease its withdrawal from the mould. Lifting off by hand or crane produced similar inaccuracies, as did errors in draft angle from hand-made patterns. All these variations would usually result in an oversize mould. The addition of mould and core washes, and variations in the density of hand ramming of the moulds, would add further variation.

A typical problem in hand-moulded castings was the variability of core location. This arose because a core laid in a print on the main cope/drag joint would usually be a poor fit. If it sat proud of the surface, then on the closing of the mould the core will be forced to settle, by deforming either the cope or the drag. Depending on which happens to have been more densely rammed that day, the core will be accommodated more in the drag or the cope, so that detail on the casting will vary up or down from day to day, and from casting to casting.

Originally, most sand moulds were made using greensand. This is a mixture of sand, clay and water, often in the form in which it was dug out of the local hillside. As such, it would be termed a *natural greensand*. Such natural products were a pleasure to work with, but its variable content of clay, and variability in the type of clay, together with variations in water content, necessarily resulted in rather variable castings.

Later, more controlled greensand mixtures were created by the use of clean, washed silica sands that were mixed with selected clays and other additives to form *synthetic greensands*. These dough-like mixtures were rammed around the pattern to form the mould. The synthetic sands could give rather harder and more reproducible moulds. Additionally, developments away from hand compaction to small-scale automation, such as jolt-squeeze compaction, have slowly given way to massive, fully automated plants with sophisticated high rate compacting mechanisms that produce very high-density moulds of impressive rigidity. Such moulding methods produce the highest accuracy yet achieved by the greensand route.

Studies like those of Bates and Wallace (1966) show how the greensand mould deforms less during casting as the mould density increases, whereas those of Rao and Roshan (1988) confirm that the final castings improve in accuracy as the compaction is improved.

Naturally, the use of a well-controlled synthetic greensand, with a metal pattern, and the mould reproduced by an automated moulding process yielding uniformly hard moulds, was a substantial step forward. Such systems continue to be widely used today for large-volume production with the highest rates of production and highly competitive economics. Despite these benefits, it may be no surprise that the high-pressure moulding process has brought its own crop of new problems in the quest for accuracy. These include distortion of tooling if not adequately constructed, and of spring-back of the mould because of the elastic behaviour of the highly compacted and highly stressed sand mould.

One of the consequences of requiring an accurate pattern and accurate mould is the requirement for *room temperature operation*. Clearly, for all those pattern and moulding processes that have to operate at high temperature raise the danger of distortion of the tooling, and poor temperature control leading to dimensional variation.

A further step in the production of accurate, very hard moulds, has been developed in recent years. This involves the use of chemically bonded sand. The ability to cure the binder at room temperature has had a further important benefit; the tooling remains accurate, suffering neither expansion nor distortion by significant changes in temperature. In fact, providing the tooling is made from metal and/or resin, the tooling never deviates by more than a few degrees from the temperature at which the toolmaker made it. These, together with the fact that the mould or core can cure while in contact with the tooling, are the two vital steps in the construction of accurate moulds. They have revolutionised the concept of the production of accurate castings.

One of the major advantages for accuracy enjoyed by the aggregate mould is not only its reproducible temperature during its formation, but also its reproducible temperature at casting. This is always reasonably close to room temperature. A temperature variation of 20°C will produce a change in length of only approximately 0.02%, or 0.1 mm in 500 mm. This is a useful stable base from which to start, and contrasts with the problems suffered by metal dies that are not easily controlled to this precision. In fact, a variation of die temperature of 100–200°C would not be uncommon, with a consequent change of size of up to 0.2%, or 1 mm in 500 mm. This represents a serious error in a casting that is required to maintain a wall thickness of 2.0 mm. Figure 18.2 is a brief summary of the competing processes.

Nowadays, the hard sand core assembly technique is perhaps the most accurate of all the casting processes, especially for castings exceeding 100 mm in size (Figure 18.3).

Such statements are not easily justified, however, because individual foundries differ widely in their capability, often simply because of attention to detail. Also, individual castings suit some processes and not others. Moreover, as a result of the new standards set by precision sand processes the competitive processes have improved markedly. Thus although it is true to say that sand casting has come a long way in recent years, so have all the others!

Thus, although our castings buyer no longer has any justification for rejecting aggregate moulded castings, he might still be justified to select a permanent mould of some kind. The phrase ‘metal die’ continues to have an *aura* of engineering excellence, and ‘pressure die casting’ has the *impression* of having an advantage of consolidation by pressure. Producers of sand castings clearly need to counter these beguiling marketing ploys, and market their own solid and reliable new achievements.

Core making has gone through a similar revolution. The two main original processes for making cores employed dry sand (i.e. dried greensand) and linseed oil–bonded sands. Both processes had in common the transfer of the core from the core box in its green, or soft, state directly into a specially shaped cradle or carrier, often called, rather confusingly, a ‘drier’. The core at this stage was so weak that it would deform at the touch of a finger. It was then hardened by drying in an oven for an hour or more. Naturally, the opportunities for loss of accuracy were plentiful. Neither process is much used today, to the relief of most casting engineers.

A great improvement in accuracy came with the development of chemical binders that would allow the core to be hardened in the core box. In this way, cores of very precise form could be guaranteed and are nowadays commonplace in the industry.

The original chemical-binder process used sodium silicate that was cured by carbon dioxide gas, often known as 'the CO_2 process' for short. This process was useful for moulds, but never popular for cores for the casting of Al and iron alloys because the heat of casting fused the silicate bond to the sand grains, producing a glassy bond of such strength that the core could not be subsequently removed from the casting. Its strength resembled that of concrete, and for many aluminium alloy castings the heat treated core exceeded the strength of the casting. Attempts to remove the core therefore often led to the destruction of the casting!

The Croning shell process, often known for short simply as the shell process, was another step forward. Here the binder was a phenolic resin, pre-coated onto the sand grains and solidified as a dry coating, thus allowing the sand to be blown or dumped in a dry, free-flowing state onto the pattern. The pattern had to be heated to melt and cure the resin in the thin shell of sand in contact with the pattern. Because of the use of heat for curing there was the danger of distortion of the tooling, and the biscuit-like core was sufficiently thin that it too was subject to distortion if not carefully handled and stored. For aluminium alloy castings, the heat input from modern thin-walled castings gives problems because of the incomplete breakdown of the resin after casting, with the result that de-coring is returning as a problem, although developments continue by suppliers of shell sand in an attempt to overcome this.

The hot box process again involves the curing of the core inside the box, and in general gives much improved de-coring performance. However, the problems of heating and controlling the temperature of core boxes, and the distortion of tooling, and the rate of diffusion of heat into sand, all remain to limit the size of core which it is feasible to produce by this technique. It remains a useful and economical process for cores where such limits are acceptable.

The various cold (or more accurately, room temperature) curing processes are not subject to these limitations. Large products can be produced so that not only internal cores for castings, but also whole moulds, can be manufactured by these processes. This has allowed the opportunity for the creation of the 'core-assembly' technique for the production of a mould. Because the tooling and the mould are all at room temperature, no thermal expansion or distortions arise to impair accuracy. For this reason the core-assembly technique is probably the most accurate moulding process in use today and can produce the most accurate aggregate mould castings. The process rivals the accuracy (if not the surface finish) of other near-net shape processes such as investment casting and HPDC (Figure 18.3).

The control problem faced by users of chemically bonded sands is that of controlling the strength of the cured sand. Continuous mixers are notoriously difficult to control; the quantities of sand and resin that are processed through the mixer change erratically, so that the quality of mixed sand is often not reproducible from time to time during the day. This will, of course, alter the strength of the moulds and cores, resulting in changes to the size and shape of the casting. Some progress is being made in using continuous monitoring and feedback control on the mixers. However, the problem is resolved to a certain extent by the use of batch mixers, where ingredients can be weighed into a mixing bowl. This still leaves the problem that the chemistry of the components of the binder themselves needs to be controlled within close limits. This is an area of continuing concern because organic resins are not easily characterised nor the chemical processing by the resin supplier controlled to the degree a founder might wish. The problem may already be limiting the accuracy of castings produced by this process route.

One interesting hard sand moulding process not requiring control over materials of imprecise chemistry such as binder resins, is the vacuum film moulding process, or V process, in which the mould is sealed on all sides by a plastic film, and the sand is consolidated by the application of a vacuum. An investigation by Grote (1982) into the reproducibility of greensand, cold set and V process, confirmed that cold set (a room temperature cured resin binder) and V process had similar good reproducibility, and both were better than greensand, although his version of greensand was probably less than an optimum system!

A final problem remaining for the sand mould is that of retaining its accuracy when subjected to the shock of being filled with molten metal. This problem includes phenomena such as the thermal distortion of the mould, its deformation under the weight of metal pressing against its walls, and the tendency for cores to float. All these matters are now usually controllable within close limits, particularly for a wide variety of non-silica aggregates and binders. For silica sand, however, we cannot avoid the 1.5% linear expansion at 570°C as it experiences its alpha to beta quartz transition. This is a major factor affecting the accuracy and reproducibility of castings, as illustrated, for instance, in the work by Jennings et al. (2001) in which he found that castings made in zircon sand which suffers no such phase change enjoyed a standard deviation 10 times less.

18.3.2 CERAMIC MOULDS

Ceramic moulds as used in investment casting have enjoyed the reputation of being the most accurate of all the various mould types available. This is a curious perception which may be the result of the process being for many years limited to the production of rather small castings where dimensional problems were naturally too small to be of concern. It probably also relates to the wonderful appearance of investment castings, with such excellent surface finish and sharp, fine detail, giving the comforting impression of accuracy.

However, when ceramic moulds are produced in sizes similar to those of other casting processes, such as precision sand core assembly processes, the dimensional problems of investment castings are seen to be non-trivial. This is hardly surprising in view of the multiple operations involved in the production of a ceramic mould, each of which can introduce errors. These include the injection of the wax into the aluminium die; the temperature of both the wax and the die affect the size and distortion of the wax pattern. The temperatures of the slurry tank and the drying room will similarly contribute an effect during the building up of the shell. The firing of the ceramic involves a certain amount of contraction of the ceramic as the bond is created. The temperatures of the mould and the metal at the time of casting are other key influences. During the time that the molten metal is in the mould, the mould suffers creep as it sinters and softens, often leading to swelling or other distortions of the casting.

Beeley (1972) shows how the expansion characteristics of different ceramic shell materials are strongly affected by how much silica is contained. He also quotes Carter, who showed that a Co-Cr alloy in an investment mould shrinks by 1.6% when the mould is at room temperature, but less than 0.3% when heated to 1000°C. Clegg and Das (1987) evaluate similar problems with Shaw process moulds, finding that the linear dimensions of the mould are greatly affected by the time and temperature used for firing, by the composition of the ceramic, and by the size of casting which is produced.

To summarise the dimensional problems of ceramic moulds, it seems a tribute to the dedication and perseverance of the industry, rather than to the intrinsic merits of the process, that all of these problems are held within tolerance sufficiently well that an acceptable product can usually be made.

18.3.3 METAL MOULDS (DIES)

Castings produced in zinc-based alloys by the pressure die casting process represent a level of accuracy which is hard to beat (Figure 18.3), representing one of the casting industry's most accurate products. This is because of the low temperatures involved. The metal die is accurate, and distorts very little, and the casting itself contracts little during its limited cooling to room temperature. (By extension of this reasoning, plastic injection mouldings can be similarly highly precise.)

For higher melting-point materials, however, the problems mount rapidly. Table 18.1 shows how metal dies suffer increasing problems with thermal shock and fatigue as the temperature of the casting alloy increases. The difficulties are such that although Mo dies have been tried for a stainless steel, the results were not satisfactory and so far as is known no one now uses the process. Even for lower temperature materials such as Al alloys, the degradation of the surface of the die by the development first of fine cracks, leading to general crazing, and finally resulting in major disintegration of parts of the surface, leads to inaccurate and roughened products long before the die becomes so bad that the decision to abandon or rework it becomes unavoidable. These problems are common in long-running aluminium alloy castings made by HPDC and gravity die (permanent mould) casting. Total failure of the die by catastrophic cracking is also sometimes suffered.

Magnesium presents less of a problem when cast into metal moulds because the thermal capacity of magnesium is significantly lower than that of aluminium, resulting in reduced thermal shock to the die. In addition magnesium

Table 18.1 Life of Metal Dies

	Die Life					
Alloy Base	**Rubber**	**Graphite**	**Gravity (Permanent Mould)**	**Low Pressure**	**High Pressure**	**Squeeze**
Zn	Good	Good	Excellent	Excellent	Excellent	Excellent
Mg	–	–	Excellent	Excellent	Excellent	Excellent
Al	Limited	Limited	Satisfactory	Satisfactory	Tolerable	Good
Cu	–	–	Tolerable	Tolerable	Rare	Tolerable
Cast iron	–	–	Tolerable	Rare	–	–
Stainless steel	–	–	–	–	Mo-based die	–
Steel	–	Griffin only	–	Griffin only	–	–

dissolves iron less readily than aluminium, so that 'soldering' to the die, a localised welding effect, with the consequent gradual removal of die surface, is reduced. The attempt to solve the soldering problem in aluminium alloys by increasing the iron content of the alloys to near-saturation levels is only partially successful because some sticking of the casting together with erosion of the die remains common.

The standard of construction of dies for HPDC is necessarily extremely high. This is a consequence of the requirement for the die to fit into, and be operated within, the massive (and expensive) machine tool that is the HPDC machine. The accuracy of fit between the various parts of the die also needs to be excellent to prevent the penetration of liquid metal at the high pressures used for injection. Between each shot, the die is sprayed with a coolant and lubricant and/or parting agent. The parting agent gives minimal build-up problems on the die, contributing significantly to the ability of HPDC to maintain the accuracy of its products (unlike, for instance, gravity die casting, where a rather thick die coat is used).

In fact, it seems that for HPDC, one of the most significant variables affecting the size of castings is the temperature at which the castings are ejected from the die. This is because if ejected early, the casting is hotter and thus cools and contracts over a greater temperature range without the constraint of the die, thus finishing smaller. A casting ejected late cools more in the die, so its contraction is restrained by the rigidity of the die; the casting is stretched plastically, as by medieval torture on a rack, thus finishing larger.

In low-pressure and gravity die casting, the standard of die construction is considerably lower. This is because the die does not need to be fitted into a precise machine, but is usually serviced by simple bolt-on actions such as rack and pinion, or simple hydraulic cylinders which actuate opening and closing of the die. The gravity die casting industry is therefore less capital-intensive, and grows more easily in a piecemeal manner. It is more labour-intensive than the HPDC industry. The two processes are not often mixed in the same foundry.

The lower rates of filling of the die under gravity require the die to be coated with an insulating ceramic layer known as a die coat. This is usually sprayed on at the start of a shift and cleaned off by grit blasting before renewal. The coating may become damaged during the production of a run of castings and might require running repairs by the operator. The die coat is therefore a major variable in the size of gravity and low-pressure die castings. Its thickness of application will vary from day to day because the spray techniques for its application are not easily controlled. Also, parts of the coating on faces with minimal relief or draw will wear during the production run, giving slowly increasing sizes of castings.

The die itself is subject to wear in many ways. In particular, the working face will be slowly eroded by repeated cleaning and recoating, in addition to any thermal damage it may be suffering. The moving parts of the die include slides and other mating components which are subject to general sliding and abrasive wear, particularly if sand cores are used in the mould.

Table 18.1 indicates some of the different materials used for dies. For HPDC dies, the construction material is usually a special hot-work tool steel containing chromium for greater strength and hardness, oxidation resistance and good response to nitriding. It is usually heat treated to achieve its optimum strength and toughness, and finally nitrided to give an extremely hard-wearing surface. Grey iron would not be adequate for pressure dies because it lacks strength and fatigue resistance. Furthermore, its graphite flakes would tend to open up by oxidation, degrading the integrity of surface which would be needed against the penetrating action of liquid metals at high pressure.

Grey iron is, however, widely used for gravity and low-pressure dies. This is because the filling of moulds under gravity or low pressure introduces only gentle pressurisation of the surface by the melt, and the surface is protected by the ceramic die coat. Also, the iron is highly castable, allowing the rough die parts to be produced quickly and usually without problems of porosity or other unsoundness; the castings are easily machined to final size and shape; the material has good thermal conductivity and fair resistance to thermal shock. Weight is easily reduced by sculpting the back of the die to give a mould of even thickness.

Graphite is occasionally chosen as a die material. It is easily machined and is not wetted or attacked by most metals (magnesium would be a spectacular exception!). For zinc alloys, therefore, it has been used with success, although care is required in longer runs to avoid damaging the die by mishandling because its strength is low. The material is not useful for aluminium alloys, because the graphite oxidises increasingly rapidly at temperatures above about 350°C. After two or three casts of aluminium alloy, the die is unusable. Treatment of the surface by a protective silicate wash extends the die life a little.

This experience makes the use of graphite as a die material in the Griffin process all the more impressive. Here, railway wheels are cast in steel displaced into the die by low pressure. The wheels are produced closely to net shape, so that they require little machining (Kotzin, 1981).

18.4 TOOLING ACCURACY

Tooling is taken to include the pattern and its core boxes, or the die, and any measuring or checking jigs and gauges.

The wear of patterns used in sand casting causes the casting to become undersize, whereas the wear of dies used in die casting causes the casting to become gradually oversize. An effect opposite to wear happens as a result of build-up problems on the tooling. Patterns for sand casting are subject to the deposition of small amounts of sand and binder, and the gradual accumulation of release agents. Dies may accumulate layers of die coat.

Distortion is another problem. Wood is a useful and pleasant material for pattern construction. It is easily worked, light to handle and easily and quickly repaired or modified as necessary. Even so, it is not a contender for accurate work because of its tendency to warp and shrink. A good patternmaker will attempt to reduce such movement to a minimum by the careful use of ply, taking care to join layers to cross grain where possible, and affixing strengthening battens. The use of various stabilised woods and simulated wood-like plastic materials has also helped considerably (Barrett, 1967). Nevertheless, the ultimate stability in tooling is only achieved with the use of metal or cast resin in metal frames.

Cast-resin patterns which are backed with wood frames are less reliable; the warpage of the wood distorts the internal resin shape, usually within a month or so. After a year, the tooling is seriously inaccurate, so that cores will not assemble properly. In contrast, resin patterns cast into aluminium alloy frames for strength and rigidity, are usually extremely reliable. However, some resin systems such as polyurethanes tend to suffer from the absorption of solvents from the chemical binders in the sand, and so can suffer swelling and degradation (Gouwens, 1967).

The working temperature of tooling affects the casting size directly; a warm pattern will give a slightly larger casting. If we consider an epoxy resin core-box cast into an aluminium alloy frame, the box will largely take its size from the temperature of the metal frame (i.e. not the internal lining of epoxy resin) which has a coefficient of expansion of $20 \times 10^{-6}\ K^{-1}$. If the temperature at the start of the Monday morning shift is 10°C, and if the returning sand creeps up to 30°C by the end of the morning, then for a 500 mm long casting the 20°C temperature rise will cause the castings to grow by $20 \times 20 \times 10^{-6} \times 500 = 0.2$ mm. This is not large in itself, but when it is added to other random variables the uncertainty in the final casting length becomes increasingly out of control.

Anderson (1987) has emphasised the important requirement that for the most accurate work the pattern or die should be used as an adaptive control element in production. Thus it needs to be built in such a way that it can be modified to produce the required size and shape of the casting. The use of patterns split transversely across their major length is common. The prior insertion of a spacer in this split allows the spacer to be removed and replaced by shims thinner or thicker as necessary. Such simple techniques involve only modest extra expense during the construction of the pattern but are a reassurance against the possibility of major expensive rebuilds later.

18.5 CASTING ACCURACY

For a good general review of accuracy, precision and tolerance concepts in casting manufacture the interested reader is recommended to the paper by Anderson (1987). We shall rely on some of his key points in this section.

18.5.1 UNIFORM CONTRACTION

This section is an examination of the various effects that follow from the strains and consequent stresses in the casting/mould system as a result of the linear contraction of the solid casting as it cools. Uniform contraction and distortion are considered separately. A distorted casting can usually be straightened, whereas a casting that is uniformly oversize or undersize is scrap.

Following Figure 7.1, when the contraction in the liquid state and most of the contraction on freezing will have been completed, the casting has cooled sufficiently to develop some coherence as a solid. Further cooling in the solid state will cause the casting to contract as a whole.

The point at which the casting develops its solid-like character is different for long- and short-freezing-range alloys. For short-freezing-range material the point is reached when the casting has developed a solid skin. For the case of long-freezing-range alloys, the point is marked by the development of a coherent skeleton of solid dendrites.

At this stage of freezing because the casting fits the mould rather well, having been poured in as a liquid (and therefore at that earlier stage fitting perfectly!), as it cools further and contracts, something has to give.

It is not the case that either the casting or the mould will yield. Both yield. Following Newton, the action of the casting on the mould causes an equal and opposite action of the mould on the casting. The degree of yielding of the casting and mould depends on the relative strengths of each. Naturally, this varies greatly from one casting/mould system to another.

The contraction of the casting from its freezing temperature to room temperature can cause the patternmaker sleepless nights. This is because the pattern must be made oversize by an amount known as the contraction allowance (or patternmaker's shrinkage allowance), so that the casting will finally finish at the correct size at room temperature. However, the patternmaker often does not know exactly what allowance to use when he starts to construct the pattern.

This was not so important in the 'bad old days' (perhaps 'the good old days'?) when castings were regarded only as 'rough castings' having plenty of machining allowance. However, now that greater accuracy is being sought in the quest for a 'near-net shape' product, the problem has become serious; the patternmaker risks finding that the wrong allowance was chosen only after the first casting is made! This is the emergency scenario when the tooling has to be modified or remade, but the tooling budget has already been spent, and the deadline for delivery is about to be passed. The contents of this small chapter are recommended to the reader as the only procedure known to the author to avoid this disaster.

The Imperial System of measurements gave us a vast legacy of choice of presentation of patternmaker's shrinkage allowance data. For instance, for simple and heavy aluminium alloy castings many are made using an addition of '5/32 inch per foot' to all the linear sizes. This corresponds to the widely used contraction allowance of '1 in 77'. The author has given up these various units and ratios in favour of a simple percentage, in this case 1.30%.

For other aluminium alloy castings with larger internal cores, such as cylinder blocks, the allowance is only 1/8 inch per foot, or 1 in 96, or, as recommended here, 1.04%.

Other aluminium alloy castings such as sumps (oil pans) and thin-walled pipes contract even less. A contraction of 0.60% is common.

Whereas the patternmaker would originally have chosen a special wooden rule whose scale was expanded by the correct amount so that he could read the dimensions directly, without conversion problems, this clearly limited him to specific contraction values. Nowadays, the greater accuracy requires that intermediate values must be chosen, such as

1.15, or 1.20% etc. for different castings. These are now easy to apply with the use of electronic and digital measuring instruments, which can be programmed for any value of contraction. In fact, a virtual solid model developed in the computer should be capable of allowing for three different contractions along each of the three perpendicular axes.

The different contractions are the result of different degrees of constraint by the mould during cooling. For instance, in the case of zero constraint, a casting such as a straight bar will contract freely to its maximum extent. We can therefore calculate this rather easily, assuming an average linear contraction. For instance for Al-Si alloys the coefficient of thermal expansion is close to $20.5 \times 10^{-6}\ C^{-1}$ and the total cooling from 660 to 25°C. From this, we can predict the contraction as $20.5 \times 10^{-6} \times 635 = 0.0130$ or 1.30%, in exact agreement with practice.

Turning now to the case of high mould constraint, it is possible to envisage an ideal case in which a large box casting with thin walls was cast around a large, rigid sand core. If the wall thickness of the casting is imagined to be vanishingly thin, like a sheet of paper, then its strength will be negligible and the core not compressed at all. Thus the casting will not be allowed to contract; its paper-thin walls will be forced to stretch. We can therefore envisage in principle the case of infinite constraint in which the casting contraction is zero.

In practice, of course, the real world is filled with casting/mould combinations that lie intermediate between the case of zero and infinite constraint; i.e. partway between 0 and 1.30% contraction in the case of aluminium alloy castings.

How can we obtain an estimate of the degree of constraint, so as to be able to predict the contraction allowance exactly? This is the patternmaker's problem. It is a difficult question, to which there is no accurate answer at this time. However, we can obtain a useful estimate by the following procedure which, fortunately, is good enough for many purposes.

In the case of the straight parallel-sided bar casting made in a sand mould, the casting suffers no constraint. We can define this as being a fully dense metal casting, in the case of aluminium having a density of about 2700 kgm^{-3}. This casting will contract the maximum amount, which for aluminium contracting from its melting point is 1.30%. In contrast, our thin-walled box casting has maximum constraint, contains maximum sand, and has (when cast and finally emptied of sand) a density of practically 0 kgm^{-3}. This simple theory gives us the two extreme points on our calibration curve given in Figure 18.5.

Intermediate points are found from measurements on actual castings, taking the volume of the casting divided by the overall volume occupied by the envelope of the casting. The envelope is the shape given by a tight-fitting rubber balloon stretched over the casting. This gives a measure of the amount of restraining sand it contains, compared with the amount of metal in the casting.

When this is carried out accurately, it is found that different varieties of casting are found to lie on a family of approximately parallel curves, all starting and finishing at our theoretical points as shown in the figure. Thus the procedure is not absolute, it does not yield a single universal curve. Nevertheless, it is a helpful guide in the absence of any better alternative at the present time.

In the case of steel castings, the famous result shown in Figure 18.6 can be explained for the first time. Following the procedure that was outlined for aluminium: for the straight bar, the average thermal contraction of steel is around $16 \times 10^{-6}\ C^{-1}$ and the cooling range from freezing point to room temperature is close to 1500°C. Thus the contraction is $16 \times 10^{-6} \times 1500$, which is 2.4%, in agreement with the measured value. We can plot this at the full density of steel of approximately 7850 kgm^{-3} to define our theoretical point, coincident with our measured point, to define the zero constraint condition. The other theoretical point is, of course, the origin (zero contraction at zero envelope density) as before. Working out the area of the sand mould envelopes of the dumbbell and H shapes in Figure 18.6, and dividing by the area of the casting, allows us to plot the two remaining points, giving the nearly linear relation in Figure 18.7.

The Al alloy contraction result from Figure 18.5 is also shown on Figure 18.7 for comparison. Until better methods become available, it seems reasonable to suppose that each foundry will have to determine for itself an equivalent of Figures 18.5 and 18.7 for each of its processes. For instance, it is well known that the values of contraction allowance for greensand are dependent on the hardness of ramming. Similarly, the percentage of binder in chemically bonded sands significantly affects the contraction of the casting. A standard trick to reduce the constraint provided by a central core is to reduce its binder level or to make it hollow.

These relations for sand moulds and sand cores are not expected to apply accurately to metal dies. Here the casting is subject to high mould constraint up to the time of ejection. Clearly, the casting contracts freely only after this instant.

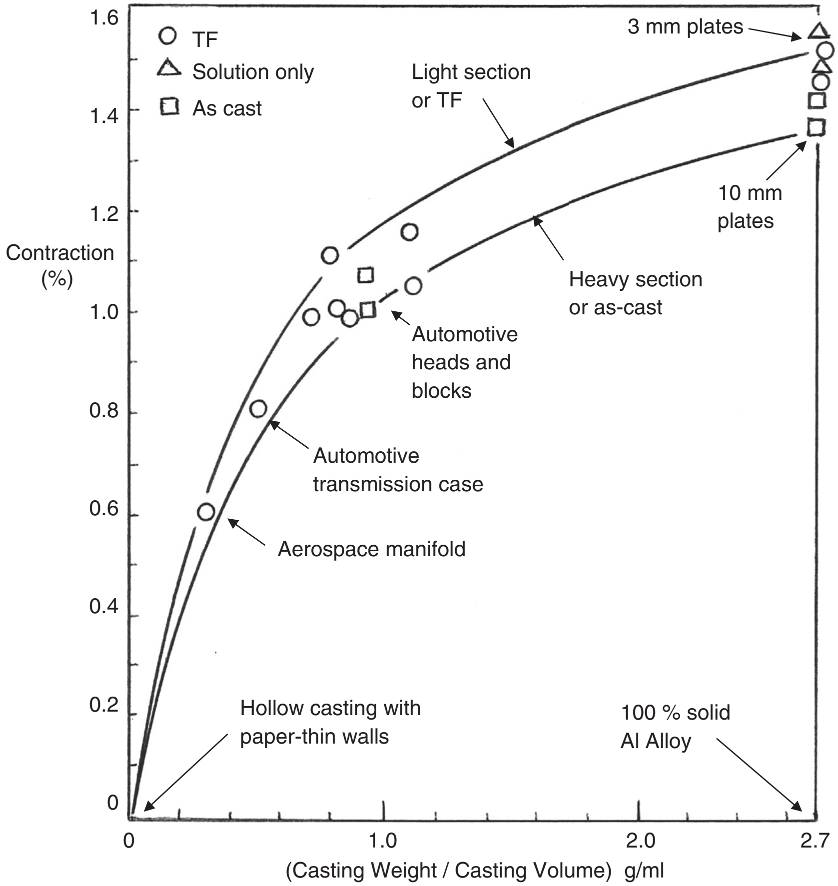

FIGURE 18.5

An experimental result from an automotive and aerospace foundry showing how some castings hardly shrink in size at all when cast, whereas other shrink almost the full theoretical amount of the solid metal. The resistance to shrinkage provided by core and mould geometry accounts for the difference.

For Al-Si alloys cast in gravity and low-pressure dies, the contraction varies between 0.75% for low-silicon alloys, and 0.5% for eutectic (approximately Al-11Si) alloys (Street, 1986), although much of the industry seems to work generally at 0.6%. These low values reflect the high resistance of the die to the contraction of the casting. However, much lower contractions, effectively zero, are occasionally found for thin boxes and window-frame–type castings.

For pressure die casting in magnesium, the contraction allowance is 0.7%, whereas for aluminium alloys it is close to 0.5%. The value is at the lower end of the range for gravity and low-pressure dies, indicating the even greater constraint in high-pressure die design.

These figures for the contraction allowance of die castings are the result of the prior expansion of the die from room temperature to its working temperature, and the subsequent contraction of the casting after its ejection out of the die (we shall assume that its contraction whilst in the die is negligible). We can estimate this quantitatively, taking the die working temperature as roughly 350°C on average (the hot face will be nearer 450°C, but the interior of the die may

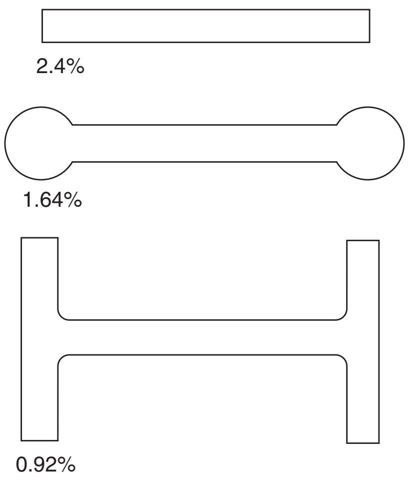

FIGURE 18.6

Contraction of steel greensand castings, showing widely different contractions.

(Steel Castings Handbook, 1970)

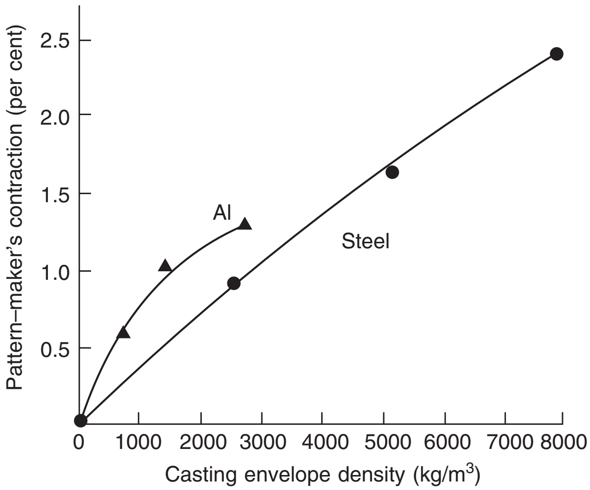

FIGURE 18.7

Contraction allowance for aluminium and steel castings as a function of mould constraint.

be water-cooled), ambient temperature as 25°C, and the temperature of the casting at ejection approximately 500°C, we have:

$$\begin{aligned}\text{Total casting contraction} &= \text{Die Expansion} - \text{Casting Contraction after ejection}\\ &= (350-25)\times 11.7\times 10^{-6} - (500-25)\times 20.5\times 10^{-6}\\ &= 0.60 \text{ per cent}\end{aligned}$$

Using these rough assumptions and simple logic, the answer is seen to be precisely correct. Furthermore, it is clear, as confirmed by experience, that the size of pressure die castings is controlled by the time of ejection of the casting from the die.

For lost-wax castings, the problem is compounded by having to take account of the expansion of the aluminium die into which the hot wax is injected, the contraction of the wax pattern, the expansion of the ceramic shell, and the final contraction of the casting itself. This complicated equation is a major source of uncertainty in the accuracy of what is known as 'precision casting'. Regrettably, it is the reason why many so-called precision castings suffer dimensional out-of-tolerance problems.

It is important to note that the pattern contraction allowance can often be different in different directions, because of different geometrical constraints offered by the casting. Thus each of the x, y and z axes may require a different value.

It is also important to remember that the casting contraction can be greatly affected by the precipitation of gas during solidification. Girshovich et al. (1963) draw attention to this problem in aluminium-, copper- and ferrous-based alloys. The author has sobering and unforgettable experience of a major pickup of hydrogen gas in a 1000 kg holding furnace after the addition of Sr to an Al-7Si-0.4Mg alloy. Over the next 3 days, the castings suffered up to 3 volume% porosity, therefore growing linearly by 1% (the 500 mm long castings growing by 5 mm!). The castings were all outside their machining allowance and were consequently scrapped. Strontium addition was immediately discontinued at that time! The subsequent introduction of better degassing techniques has allowed the question of Sr addition to be revisited. In a related experience with steels, Schurmann (1965) describes how rimming steel ingots that failed to develop any significant rimming action tended to grow a solid crust over the ingot top. The internal pressure in the ingot could not then be relieved by the escape of gas, so the ingot swelled.

Finally, the reader would be forgiven for assuming that the size of the casting was fixed when at last it reached room temperature. However, this is rarely true. For instance, Figure 18.8 shows how the common zinc pressure die casting alloys continue to shrink in size for the first 6 months or so. Alloy A is then fairly stable for the next few decades. Alloy B, on the other hand, starts to reverse its shrinkage after about the first year. At 100°C these changes can be accelerated by about a factor of 250, effectively compressing years into days, as the time axis shows.

The zinc die casting alloy ZA27 (Zn-27%A1) shrinks only about one-tenth of the amount of the lower-aluminium zinc alloys (8 and 12%Al), but its expansion is greater, as seen in Figure 18.9. (Incidentally, the time axis in Figures 18.8 and 18.9 is calculated assuming the useful relation for reaction rates of a factor of two increase in reaction rate for every 10°C rise in temperature.)

Aluminium alloy castings also show size changes. For instance, Al-7Si-0.5Mg alloy contracts by 0.1–0.2% after solution treatment and quenching as alloying elements are taken into solution. The castings grow by 0.05–0.15% during ageing as the alloying elements precipitate once again (Hunsicker, 1980). Gloria et al. (2000) find that Al-8Si-3.3Cu-0.2Mg expands during solution treatment as the CuAl2 phase dissolves.

Figure 18.10 shows the Al-17Si alloy used for wear-resistant applications exhibiting considerable growth at temperatures high enough to allow silicon to precipitate. This growth in service can be reduced by a preage at a minimum of 230°C for 8 h (but a treatment of 260°C for 1 h would be closely equivalent, using our 10°C rule equivalent to doubling the reaction speed).

Even higher percentage growths are exhibited by white cast irons. At high temperatures, greater than approximately 900°C, the breakdown of the metastable cementite to stable graphite takes several days. During this period, the linear dimensions of the casting will grow by up to 1.6% (Johnson and Nohr, 1970). Grey irons will also grow at temperatures down to 350°C, and growth can be catastrophic if the iron is cycled repeatedly through the ferrite/austenite phase change. Walton (1971) observed a growth of 3.5% in only 500 h in a grey iron subjected to cyclic heating to 800°C. Growth can be further enhanced by the internal oxidation of the material. Angus (1976) gives more details of the growth of cast iron at elevated temperatures, and the means by which it can be controlled by control of structure and chemical composition of the iron.

These are just a few examples of the growth and/or shrinkage of castings that can occur in the solid state because of microstructural changes occurring within the alloy. The casting engineer needs to be on guard against such problems.

The changes cause the patternmaker problems when attempting to decide what contraction allowance to use when constructing the pattern. His decision may be right or wrong, depending on whether the foundry check the dimensions of the casting before or (more correctly) after heat treatment, and will depend on the service conditions.

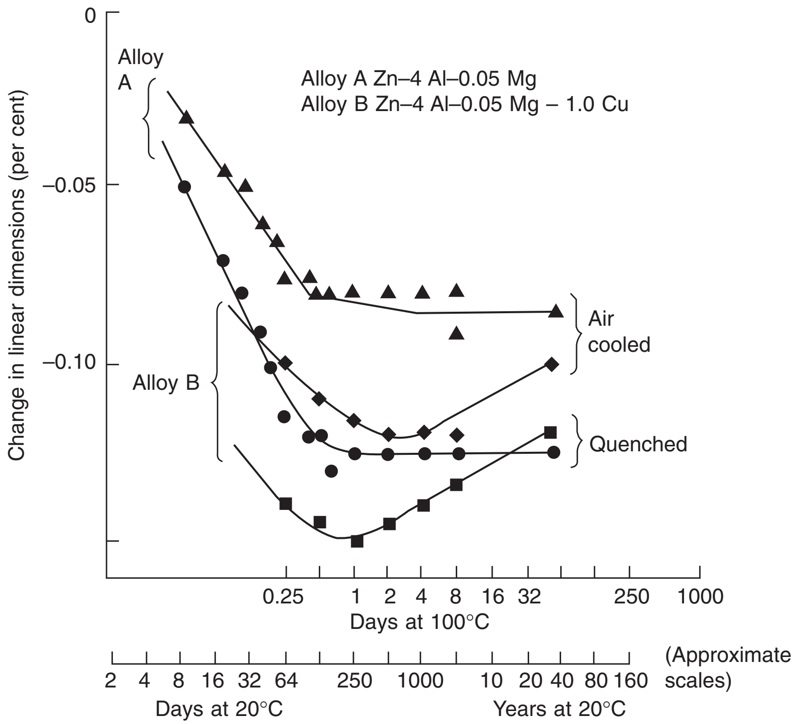

FIGURE 18.8

Zinc pressure die casting alloys (Zn-4Al) showing accelerated ageing at 100°C, or slow shrinkage followed by expansion in alloy B taking the place of decades at 20°C.

Data from Street (1977).

If the allowance was chosen wrongly, giving an undersized casting in grey cast iron, then a heat treatment may save the day by growing the casting by up to 1% or so. However, such good fortune is rare. The choice of contraction allowance before making the casting will remain a difficult and risky decision, and will become more difficult as demands on casting accuracy increase.

18.5.2 NONUNIFORM CONTRACTION (DISTORTION)

If the casting was cooled at a uniform rate and with a uniform constraint acting at all points over its surface, then it would reach room temperature perfectly in proportion, perhaps a little large, or a little small, but not distorted.

In practice, of course, this utopia is never realised. Usually the casting is somewhat large, or somewhat small, and is not as accurate a shape as a discerning customer would prefer. Occasionally, it may be very seriously distorted. We shall examine the reasons for these factors and see to what extent they can be controlled.

Mould constraint

Again, wishing ourselves into utopia, we can envisage that if the constraint by the mould were either zero or infinite, in both cases the casting would be of predictable size and correct shape.

In reality, of course, different parts of the casting experience different degrees of constraint by the mould. One of the most common examples of this problem is a simple five-sided box with its sixth side remaining open, as shown in

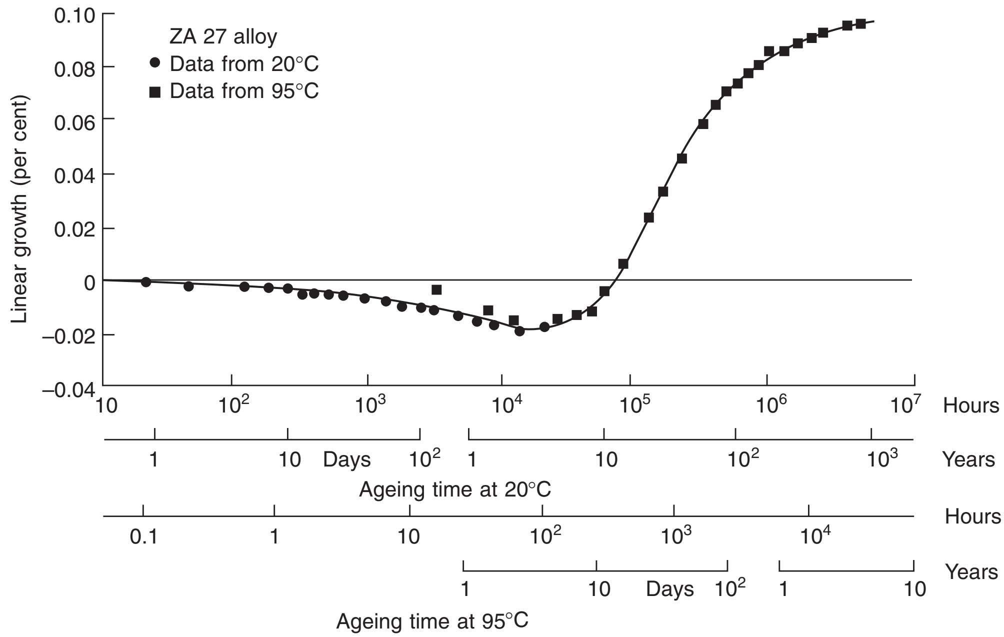

FIGURE 18.9

Zn-27Al alloy dimensional changes with time.

Data from Fakes and Wall (1982).

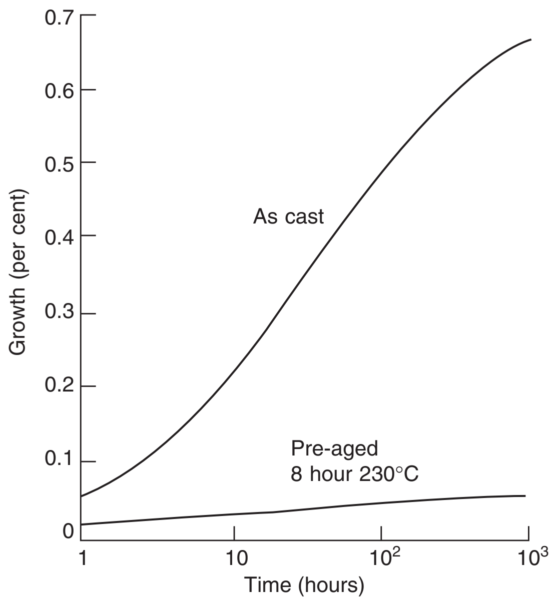

FIGURE 18.10

Permanent growth of A390 (Al-17Si alloy) at 165°C with time (Jorstad, 1971).

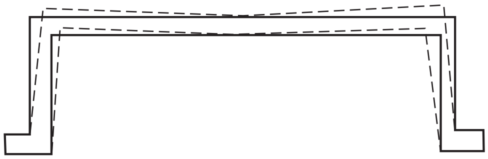

FIGURE 18.11

Distortion of an open-sided box casting during cooling as a result of uneven mould constraint.

Figure 18.11. The closed face wishes to contract as a straight length, whereas the vertical sides have maximum constraint. The result is a compromise, with the straight side shortening with an effective contraction allowance of perhaps 1.2% in the case of an aluminium alloy casting. The vertical sides will be restrained from pulling inwards and so have an effective contraction allowance of perhaps only 0.9%, or perhaps as low as zero if the walls are very thin. During cooling the casting therefore develops a bowed shape with non-parallel walls. For a hard, chemically bonded sand mould the camber will be approximately 1 mm in the centre of the long side of an open box 500 × 100 × 100 mm with walls 4 mm thick.

The box casting may be cast somewhat straighter by several techniques well known to the foundry technologist.

1. The centre core can be made weaker, reducing its constraint on the contraction of the casting. This is achieved either by reducing the percentage addition of binder, or, usually more conveniently, by hollowing out the centre core. The thin shell of sand thereby becomes hotter, giving greater breakdown of the binder in the case of an organic binder, and so allowing the core to collapse earlier.
2. Tie bars can be connected across the open side of the box, thereby holding the walls in place and balancing the effect of the contraction of the closed side. The tie bar need not be a separate device. It can be the running system of the casting, carefully sized so as to carry out its two jobs effectively.

This raises the important issue of the influence of the running and feeding system. Unfortunately, these appendages to the casting cannot be neglected. They can be used positively to resist casting distortion as previously. Alternatively, they can cause distortion and even tensile failure as shown in the simple case in Figure 18.12(a). In addition, if this casting is filled at the flange end, leaving the plate free to contract along its length as shown in Figure 18.12(b), the problem is solved. However, note the important point that whether or not casting ‘a’ has suffered any tensile failure, it will measure somewhat longer than casting ‘b’. Thus different pattern contraction allowances are appropriate for these two different constraint modes.

In more complex castings, the effect of geometry can be hard to predict and harder to rectify if the casting is particularly badly out of shape. Especially for large, thin-walled castings requiring close dimensional tolerance it may be wise to include for a straightening jig in the tooling price. This will be an expensive piece of tooling, usually resisted by the customer.

Casting constraint

Even if the casting were subjected to no constraint at all from the mould, it would certainly suffer internally generated constraints as a result of uneven cooling. The famous example is the mixed-section casting shown in Figure 18.13(a). If a failure occurs it always happens in the thicker section. This may at first sight be surprising. The explanation of this behaviour requires careful reasoning, as follows.

First, the thin section solidifies and cools. Its contraction along its length is easily accommodated by the heavier section, which simply contracts under the compressive load because it is hot, and therefore plastic, if not actually still molten. Later, however, when the thin section has practically finished contracting, the heavier section starts to contract. It is now unable to squash the thin section significantly, which has by now become rigid and strong. The result is the possible bending of the thin section as tension builds up in the thick section. Under its tensile stress, the thick section might therefore stretch plastically, or hot tear, or cold crack.

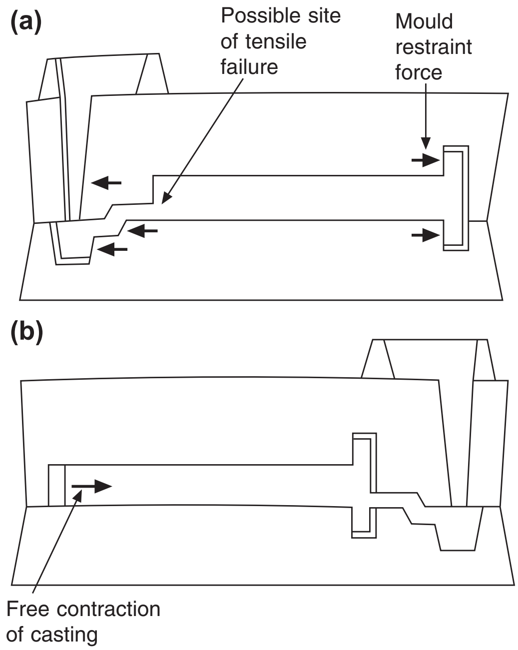

FIGURE 18.12

(a) Effect of the filling and feeding systems imposing constraint on the contraction of a casting. (b) Applying the filling system to the opposite end of the casting eliminates the problem, permitting the casting to contract freely.

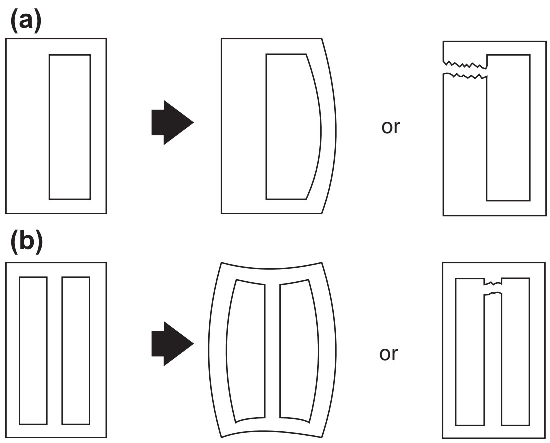

FIGURE 18.13

(a) Thick/thin section casting showing tensile stress in the thick section; (b) even-walled casting showing internal tensile stress.

The example shown in Figure 18.13(b) is another common failure mode. The internal walls of a casting remain hot for longest even though the casting may have been designed with even wall sections. This is, of course, simply the result of the internal sections being surrounded by other hot sections. The reasoning is therefore the same as that for the thick-/thin-section casting previously. The outer walls become cool and rigid, and the internal walls of the casting suffer tension at a later stage of cooling. This tension may be retained as a residual stress in the finished casting, or may be sufficiently high to cause catastrophic failure by tearing or cracking.

The same reasoning applies to the case of a single-component heavy-section casting such as a solid ingot, billet or slab, and especially when these are cast in steel, because of its poor thermal conductivity. The inner parts of the casting solidify and contract last, putting the internal parts of the casting into tension (notice it is always the inside of the casting that suffers the tension; the outside being in compression). Because of the low yield point of the hot metal, extensive plastic yielding occurs at high temperatures. However, as the temperature falls, the stress cannot be relieved by plastic flow, so that increasing amounts of stress are built up and retained.

An example shown in Figure 18.14 shows the kind of distortion to be expected from a box section casting with uneven walls. The late contraction of the thicker walls collapses the box asymmetrically (the casting is at risk from tensile failure in the thicker walls, but we shall assume that neither tearing nor cracking occurs in this case). There is clearly some strong additional effect from mould constraint. If the central core were less rigid, then the casting would contract more evenly, remaining more square.

There is an important kind of distortion seen in plate shaped castings which have heavy ribs adjoining the edges of the plate, or whose faces are reinforced by heavy-section ribs. It is often seen in thin-section boxes that have reinforcing ribs around the edges of the box faces. The general argument is the same as before: the thin, flat faces cool first, and the subsequent contraction of the heavier ribs causes the face to buckle, springing inwards or outwards. This is known as 'oil-can distortion'; an apt name which describes the exasperating nature of this defect, as any attempts to straighten the face cause it to buckle in the opposite direction, taking up its new reversed curvature. It can be flipped backwards and forwards indefinitely, but not straightened permanently. Once a casting exhibits oil-can distortion, it is practically

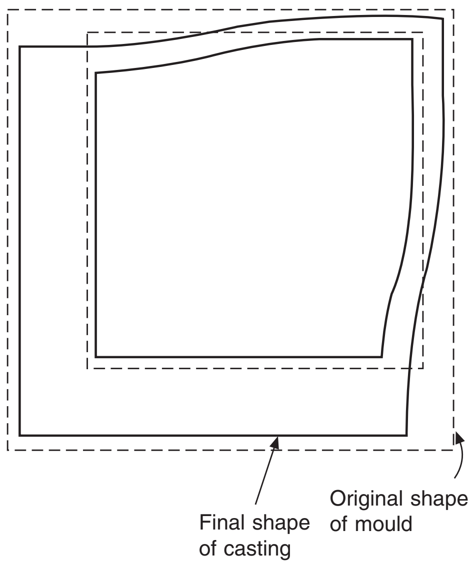

FIGURE 18.14

Distortion in an uneven box section casting because of combined casting and core constraint.

impossible to cure. The effect is more often seen after quenching from heat treatment, where, of course, the rate of cooling is greater than in the mould, and where the casting does not have the benefit of the support of the mould.

Oil-can distortion may be preventable by careful design of ribs to ensure that their geometrical modulus (i.e. their cooling rate) is similar to, or less than, that of the thinner flat face. Alternatively, the cooling rate from the quench needs to be equalised better, possibly by the use of polymer quenchants and/or the masking of the more rapidly cooling areas.

In ductile iron castings that exhibit expansion on solidification resulting from graphite precipitation, the expansion can be used to good effect to reduce or eliminate the necessity for feeders, particularly if the casting cools uniformly. Tafazzoli and Kondic (1977) draw attention to the problem created if the cooling is not uniform. The freezing of sections that freeze first leads to mould dilation in those regions that solidify last. Although these authors attribute this behaviour to the mismatch between the timing of the graphite expansion and the austenite contraction, it seems more likely to be the result of the pressure within the casting being less easily withstood by those portions of the mould that contain the heavier sections of the casting. This follows from the effect of casting modulus on mould dilation; the lighter sections are cooler and stronger, and the thicker sections are hotter and thus more plastic. Any internal pressure will therefore transfer material from those sections able to withstand the pressure to those that cannot. The thinner sections will retain their size while the thicker sections will swell. Tafazzoli and Kondic recommend the use of chills or other devices to encourage uniformity.

In a classic series of articles, Longden (1931–1932, 1939–1940, 1947–1948, 1948) published the results of measurements that he carried out on grey iron lathe beds and other machine beds. The curvature of a casting up to 10 m long could result in a maximum out-of-line deviation (camber) of 50 mm or more. Longden summarised his findings in a nomogram that allowed him to make a prediction of the camber to be expected on any new casting. The reverse camber was then constructed into the mould to give a straight casting. Although Longden's nomogram is probably somewhat specific to his type of machine tool bases, and therefore not useful as a general predictive tool for other castings, it is presented in Figure 18.15 as an example of what can be achieved in the prediction of casting distortion. It is to be expected that castings of other types may exhibit a similar relationship.

Pursuing this nice but very specific example, it is possible to show that Longden's graphical summary can be converted into a quantitative form, where length L and depth D are in metres and wall section thickness w and camber c are in millimetres. Simplifying Longden's nomogram within the limits of accuracy of the original data and making the further approximation that the slight curved lines on the left-hand side are straight lines through the value $L = 1.5$ m, then with perhaps about 10% accuracy for wall thickness w from 10 to 40 mm the camber is given by:

$$c = (L - 1.5)(7.62 - 1.073w) - 660D + 310$$

and for wall thickness w from 40 to 70 mm:

$$c = (L - 1.5)(144 - 2.03w)(1 - D)$$

We now move on to a further type of internal constraint that appears to be universal in castings of all sizes and shapes, and which is rarely recognised, but was investigated by Weiner and Boley (1963) in a theoretical study of a simple slab casting. They assumed elastic–plastic behaviour of the solid, and that the yield point of the solid was zero at the melting point (not quite true, but a reasonable working approximation) and increased as the casting cooled. They found that plastic flow of the solid occurs at the very beginning of solidification. The stress history of a given particle was found to be as follows. On freezing, the particle is subject to tension, and because the yield stress is initially zero, its behaviour is at first plastic. As it cools, the tensile stress on it increases and remains equal to the yield stress corresponding to its temperature until such time as the rate of increase of stress upon it is less than the rate of increase of its yield stress. It then starts to behave elastically. Soon after, unloading begins, the stress on the particle decreases rapidly, becoming compressive, and finally reaches the yield stress in the opposite direction. Its behaviour remains plastic thereafter. Weiner and Boley's analytical predictions have been accurately confirmed in a later numerical study by Thomas and Parkman (1998).

To sum up their findings, in a solidifying material there will be various deformation regimes. These are (1) a plastic zone in tension at the solidification front because the strength of the solid is low; (2) a central region where the stresses are in the elastic range; and (3) a zone at the surface of the casting where there is plastic flow in compression. The overall scheme is illustrated in Figure 18.16.

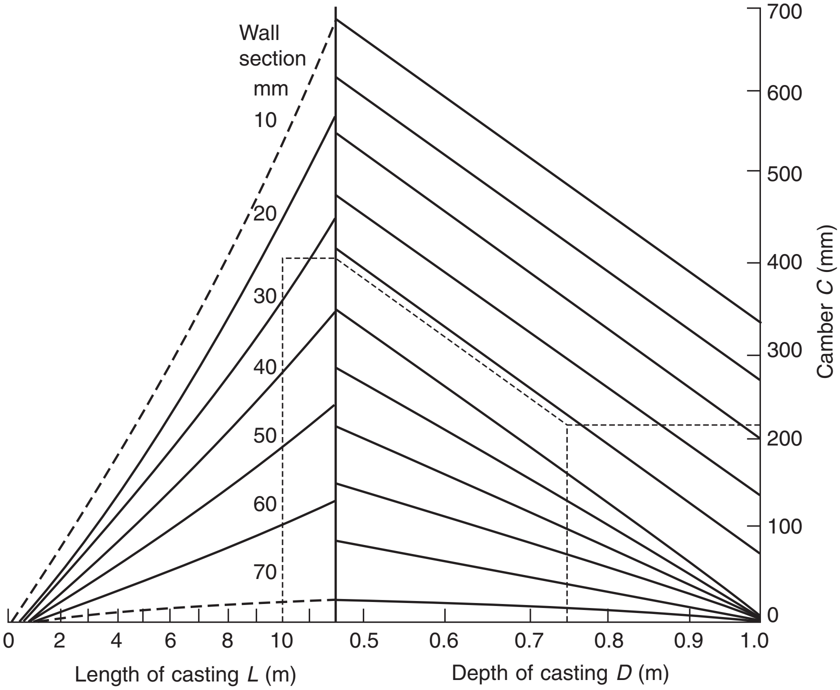

FIGURE 18.15

Camber nomogram based on data from Longden (1984) for machine tool bed castings in grey iron. A 10 m long casting with 25 mm thick side walls 750 mm depth will show approximately 210 mm camber (deviation from straight).

The propagation of the tensile plastic region, the central elastic zone, and the compressive plastic zone are reminiscent of the propagation of the various transformation zones through the sand mould. These waves of strain and stress spread through the newly solidified casting, remaining parallel to the solidification front and remaining at the same relative distances, as illustrated in Figure 18.16.

If the yield stress were not assumed to be zero at the freezing point, but were to be given some small finite value, then the analysis would be expected to be modified only very slightly, with a narrow elastic zone appearing at the solidification front on the right-hand side of Figure 18.16.

The analysis will be fundamentally modified for materials that undergo certain phase changes during cooling. If the crystallographic rearrangement involves a large enough shear strain, or change of volume, as is common in steels cooling through the γ to α transition for instance, then the material will be locally strained well above its yield point, adding an additional plastic front which will propagate through the material. The phenomenon is known as *transformation induced plasticity*. This additional opportunity for the plastic relief of stress will fundamentally alter the distribution of stress as predicted in Figure 18.16. However, the prediction is expected to be reasonably accurate for many other metals such as zinc-, aluminium-, magnesium-, copper- and nickel-based alloys, and for those steels that remain single phase from solidification to room temperature.

The high internal tensions predicted by this analysis will be independent of, and will be superimposed on, stresses that arise as a result of other mould and/or casting constraints as we have discussed previously. It is not surprising, therefore, to note that on occasions castings fail whilst cooling after solidification.

Drezet et al. (2000) showed that the elastic–plastic model would not be fundamentally altered if creep flow behaviour were assumed instead of the elastic–plastic flow behaviour with yield stress a function of temperature. They found that

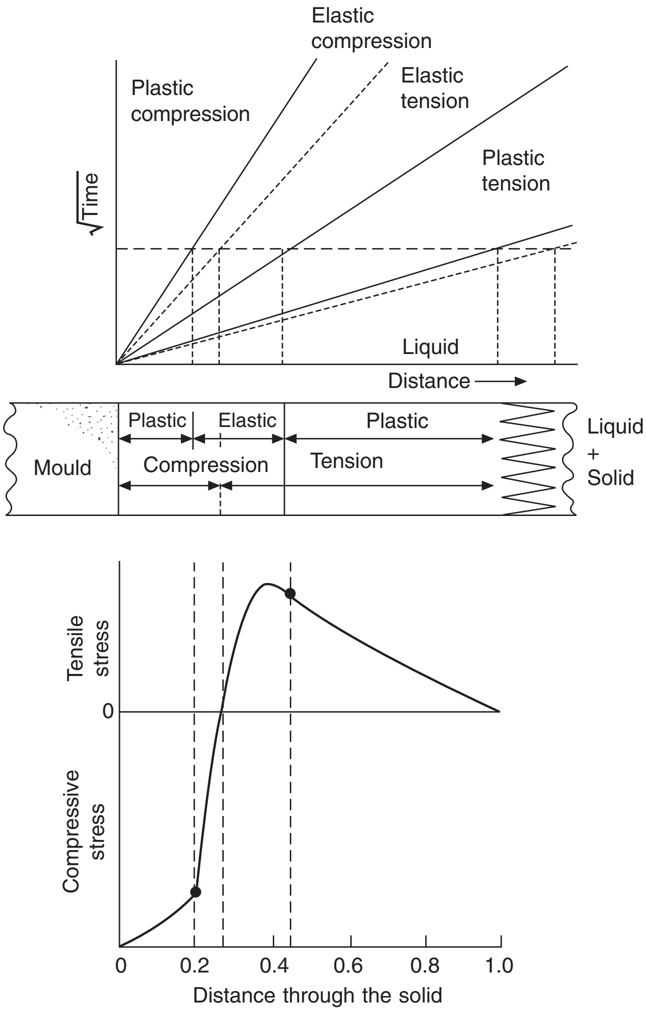

FIGURE 18.16

Elastic/plastic regimes in a simple slab casting.

After Weiner and Boley (1963).

such simulations were insensitive to the rheological model employed, but that the deformation was mainly a simple function of the thermal contraction and the conditions for continuity.

Richmond and Tien (1971) and Tien and Richmond (1982) criticise the model by Weiner and Boley on the grounds that they do not take account of the friction at the casting–mould interface. When Richmond and Tien include this they find that the casting–mould interface is no longer in compression but in tension. This is almost certainly true for large castings in metal moulds such as steel ingots in cast iron ingot moulds, where the pressure between the mould and casting is high, the friction is high, and the mould is rigid. These authors explain the occurrence of surface cracks in steel ingots in this way. However, Weiner and Boley are likely to more nearly correct for smaller castings in sand moulds. Here the interfacial pressure will be less, and the surface more accommodating, and the air gap ensuring that the casting and mould are not in contact in some places. All these factors will reduce the restraint resulting from friction. Thus, their analysis remains probably the most appropriate for medium-sized shaped castings.

18.5.3 PROCESS COMPARISON

It is not easy making a comparison between the various casting processes to ascertain their comparative capabilities in terms of accuracy, or possibly, reproducibility. There are many separate factors that influence the capabilities of different processes, to the extent that it may seem rather surprising that a casting achieves any approach to accuracy at all. However, it is a pleasure to note that the majority of casting processes do rather well, and some excellently. That is not to exclude the certainty that most could probably do better!

There have been several studies attempting to quantify what tolerances are achieved in practice. Useful researches were carried out in Sweden by Villner (1969, 1974) and in the United Kingdom by the Institute of British Foundrymen (IBF Technical Subcommittee T571, 1969, 1971, 1976, 1979). These workers took production castings from several foundries, covering a wide variety of different metals and different processes. Linear measurements taken from a large number of castings were subjected to multiple regression analysis, a statistical technique used to assess quantitatively the effects of a large number of factors simultaneously. The technique identified the following factors as important.

1. The total area of cores projected in the plane of the mould joint.
2. The wall thickness, indicating that thinner walled castings were in general more accurate.
3. The mould joint was generally closed accurately to within approximately 0.1 mm showing that the main mould joint was usually one of the most accurate of the mating interfaces within the mould assembly, contributing only a very small error to any final dimension across this joint. (Note that lateral error, known as mismatch across the joint, was a separate problem. Mismatch could often be more than 10 times the closing error, and could easily scrap castings.)
4. The great value of keeping critical dimensions or reference points all within one half of the mould, or totally within one core was demonstrated. This was a large factor reducing variability.

Apart from the three continuous variables, 'drawing dimension', 'projected area of core' and 'general wall thickness', there were additional factors contributing to variation such a metal type, pattern and mould distortions, temperature variations, coating thickness etc.

Regardless of these numerous effects, the main result was that the standard deviation, sigma (σ), increased with size of the dimension. Although the regression equation derived by these investigators indicated that σ increases linearly with the dimension, this was only, of course, the result of the regression analysis itself being carried out assuming linearity, although the assumption of linearity was probably not too bad as a working approximation at this stage.

Taking therefore an example of a hollow casting 5 mm wall thickness, of overall size approximately $250 \times 200 \times 200$ mm, containing a large body core, Figure 18.3(a) shows the result for 2.5σ, where σ is the expected standard deviation (2.5σ encloses 99% of all expected results). A second study by the author in 2000 considered only probable variation in the length of a simple solid casting using σ as a measure. The results of this rather different casting are shown in Figure 18.3(b). The similarities between the two plots are probably significant because not only length is seen to dominate.

In fact, from Figure 18.3(b), is clear that the variability of processes such as gravity die casting are dominated by factors which are not functions of length and probably include such variable factors as die coating. This contrasts with processes such as lost foam, and more particularly, lost wax, whose variability increases with casting size, reflecting the importance of thermal expansion problems in these processes. Thus lost wax can retain its common name '*precision casting process*' only for very small castings. For large castings, the accuracy of the process becomes no better than low-technology sand casting (although it retains its excellent surface finish and definition of course).

The processes which shine out as having intrinsically repeatable dimensions are pressure die casting and high-technology sand casting. There are good reasons for this. Pressure die casting is a simple, direct process, operated in a rigidly supported steel die and uses no die coat. Sand casting involves more steps, but all the steps are carried out at room temperature so that no significant expansion errors accumulate.

The position of gravity die being systematically below that of the variability of low-technology sand casting was one of the major reasons for the historical choice of 'die casting' as opposed to sand casting for many automotive applications. Now, interestingly, recent advances in sand moulding, both in greensand and chemically bonded sands, have overtaken the accuracy possible in gravity die. Even so, the poor image of sand casting lives on in the engineering profession, whereas 'die casting' has the image of cleanness and precision! It is clear that a re-education of the engineering profession, along with ourselves as founders, will take a generation or more.

From Figure 18.3(a) and (b) the increase in accuracy for a 500 mm long aluminium alloy casting when changing from gravity die to a precision sand process is a factor of between 2 and 4. This is roughly in line with measurements carried out on production runs of cylinder heads by Ford of America who found on average that an accurate sand process yielded castings more than twice as accurate as their standard supplies of gravity die castings. This was a critical finding which led to the adoption of aggregate moulds using the Cosworth Process in North America for automotive cylinder blocks instead of their previous reliance on permanent mould castings.

In conclusion, pressure die casting and high-technology sand casting are the processes with the greatest capability for reproducibility of dimensions. The lost wax process is accurate for small components, but as the casting size increases becomes rapidly poorer, eventually becoming as bad as low-technology sand casting at large sizes. Gravity die casting and lost foam casting show intermediate performance.

A word of warning about Figure 18.3(a) and (b) is probably necessary. Such comparative diagrams are intended to give a general overview of the capability of the various processes, and to this extent they are useful and fair. In particular instances, though, certain foundries may achieve much better results than the norm and, regrettably, some much worse! Also, some castings of regular shape are more easily kept to close tolerance, whereas flimsy or complex shapes may prove very difficult. Thus figures such as Figure 18.3 require to be treated with some caution.

Finally, we need to remind ourselves that cast products are not bought simply for their dimensional reproducibility. Surface finish, internal integrity and many other factors, not the least of which is cost, are important. This section has been aimed at clarifying and quantifying so far as possible only the ability of a process to achieve a level of dimensional control.

18.5.4 GENERAL SUMMARY

The main factors which control the accuracy of the final casting are briefly listed as a summary.

1. Pattern (or tooling) inaccuracy.
2. Mould inaccuracy.
3. Mould expansion and/or contraction because of temperature or pressure.
4. Casting expansion because of precipitation of less dense phases such as graphite or gases.
5. Casting contraction on freezing (solidification shrinkage) causing local sinks.
6. Casting contraction on cooling leading to (a) different overall casting size, depending on the constraint by the mould, and (b) distortion if unevenly constrained or unevenly cooled.
7. Casting overall change of size on heat treatment or on slow ageing at room temperature.
8. Casting distortion if unevenly cooled by an inappropriate quenchant or too rapid quench from heat-treatment temperature.
9. Casting distortion caused by shot blasting. The compressive stresses introduced into the surface by a peening effect can lead to the distortion of the casting as outlined in Chapter 19, *Postcasting Processing*.

18.6 METROLOGY

The technology of metrology has advanced so far over recent years, and continues to advance, such that I seriously thought that this section should be abandoned. However, not all of us have the benefits of modern, computer-controlled, remote-scanning laser techniques that operate as a three-dimensional coordinate measuring machines. Even for those that do have the luxury of remote optical scanning for dimensional inspection, Minetola et al. (2012) explain the concerns of potential pitfalls. Thus, in this transitional period, while many of us struggle with our traditional systems, and possibly struggle with the new technology, there is perhaps room for this section for a year or two more. Those with laser systems or X-ray radiography tomography are permitted to skip this section.

Even if it were possible to produce an absolutely accurate casting, it would not be possible to prove it! This apparently curious statement is the consequence of errors which occur during measurement. Inexact measuring of the casting will cause the random deviations in the measurements, causing dimensions of the casting to appear too large or too small. Svensson and Villner (1974) point out this problem and work out the influence of measuring accuracy on the apparent dimensional accuracy of the casting.

It is clear that even if the casting has dimensions which are perfectly correct; even careful measurement will introduce a certain amount of apparent error, and careless measurement will, of course, introduce even more. The more recent introduction of large-size three-dimensional coordinate measuring machines has significantly reduced these errors, which have been such a traditional problem within the industry.

Even so, problems will remain. For instance, the Swedish workers point out that for small dimensions, and where high accuracy is required, the surface roughness will influence the apparent accuracy of the casting. Thus a change in the surface finish from 75 to 200 μm will give an increase of one tolerance grade in the ISO system.

The surface finish influences the measurement and location processes in other ways. For instance, the modern touch probes, which locate dimensions on the casting with the most delicate of contact pressure, effectively only measure to high spots, thus biasing the measurements in one direction: exterior dimensions on the casting are measured oversize, and cored holes appear undersize.

Results from mark-out equipment using a mark-out table and a mechanically scribed line tend to give more averaged results because minor surface irregularities are cut through.

Similarly, when castings are clamped on to their location points, the small area of the contact points, typically a 5 mm diameter and the high loads which can be exerted by the clamps, ensure that the locating jig point actually indents the surface of the casting by up to 0.25 mm for some aluminium alloy sand castings. Harder materials such as cast irons will, of course, indent less. All surface irregularities are effectively locally smoothed and averaged in this operation. The indentation effect sets an upper limit to the accuracy and repeatability with which castings can be picked up for measurement or machining.

A traditional method of checking the profile of a casting is by the use of template gauges. These are typically sheets of metal which have been cut to the correct contour. On applying them to the casting, the contour on the casting can be seen to be correct or not, depending on the clearance which can be seen between the two. This is an analogue technique which can no longer be recommended in these modern times. The gauges are expensive to make. They are also subject to wear and thus need to be checked regularly and occasionally replaced. However, what is much more serious, they are difficult to use in any effective way. This is because in practice the contours never match exactly. The problem for the user then is how inaccurate can the contour of the casting be allowed to become before remedial action must be taken?

The use of ‘go/no go’ gauges removes the matter of judgement. However, the gauges are again subject to wear, and thus require the cost and complexity of a calibration system. More fundamentally, their use is similarly not helpful in terms of providing useful data to assist process control.

All these difficulties can be removed by the use of a much simpler technique: the use of simple goalpost fixtures which straddle the casting and which are equipped with one or more spring-contact probes, such as dial gauges. The readings from the gauges are read and recorded. The operation becomes even simpler with the use of digital read-out devices (Figure 18.17). Linear transducers are easily fitted and operated, and give an immediate numerical signal of the degree of inaccuracy. Laser techniques give an even faster and contactless benefit involving no wear.

The goalpost would be calibrated and stored on a standard casting and thereby always be seen to be in calibration by being set to zero in this position. (For calibration away from the zero, other readings can be obtained by the insertion of slip gauges under the probe.)

The use of digital electronic read-out in this way allows its incorporation into data-logging and quality-monitoring systems, such as statistical process control. By watching the trends on a daily or weekly basis, the gradual drifts in casting dimensions can be used to predict, for instance, that tooling wear will reach a level which will require the tooling to be replaced in 3 months’ time. Such prior warning allows the appropriate action to be planned well in advance.

Even the use of goalpost systems is outdated by modern digital laser-reading systems, in which the contour is held as a virtual contour inside the computer. Even complex free-form surfaces typical of aerofoils can now be modelled remotely on a screen; the products scanned and checked for conformity within seconds; and the results immediately forwarded electronically to the customer in any part of the world.

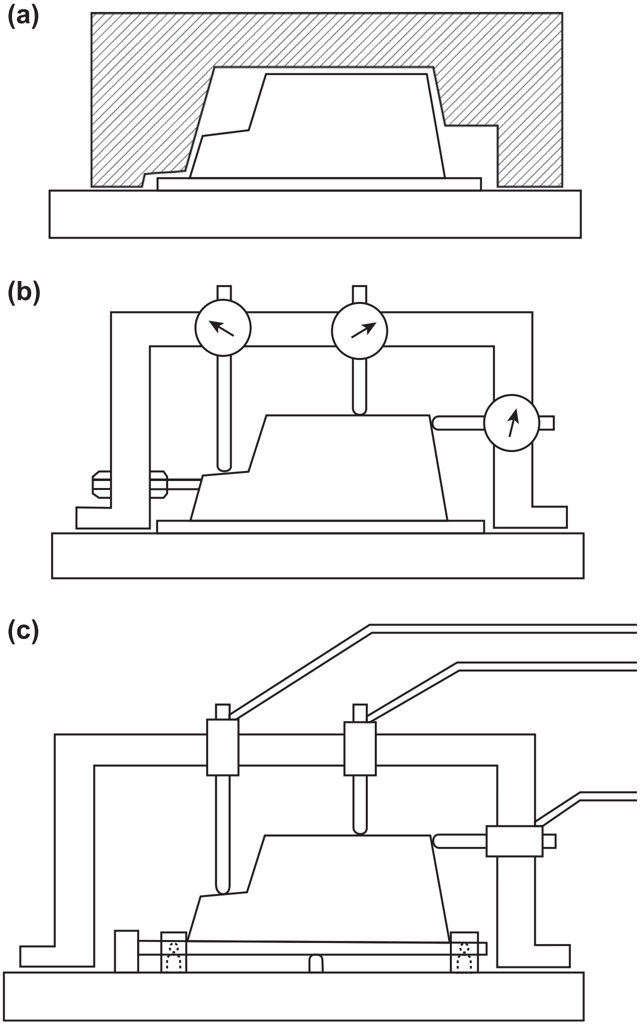

FIGURE 18.17

Checking techniques for size and shape of castings; (a) template; (b) analogue quantification using dial gauges; (c) digital quantification using linear displacement transducers with the casting on a six-point jig. Modern equivalents are laser non-contacting techniques.

Very large castings benefit particularly because the handling of mechanical gauges several meters long, while attempting to measure to fractions of millimeters become difficult if not impossible. The use of a remote laser technique is quick, easy for the operator, and amazingly accurate. Parts of up to 4 m in size can be measured to within 0.03 mm (Tigges, 2010).

Digital techniques have introduced a new level of capability in terms of the accurate measurement of castings. This has raised the stakes for the founder, who now is challenged to improve casting accuracy yet further.