12 Using less metal by design

We use so much metal that we've designed and optimised our production processes to make it with great efficiency. However a feature of that efficiency is that it is much cheaper to make a large volume of material of the same shape than to make each piece of metal a different shape—there are significant economies of scale related to tooling costs, and the speed of continuous as opposed to discrete processes. As a result, it's almost always cheaper to make components with simple geometries, than to use less metal.

The Wright brothers' Flyer

Henry Ford's Model-T

1903 was a good year for engineering: Wilbur and Orville Wright took to the skies in the first powered flight at Kitty Hawk in North Carolina, while a few hundred miles away, Henry Ford set up the Ford Motor Company in Dearborn, Michigan, a suburb of Detroit. The Wright brothers' main contribution to flying was to invent three-axis control, but they also worked on wing geometries and cutting down the weight of their Flyer. And how did they cut down the weight? They removed every possible strut from the airframe, made the rest as thin as possible, used the right materials and built their own engine, with—thanks to the leadership of Napoleon III—a cast aluminium engine block. In order to fly, the Wright brothers had to learn how to design the lightest possible plane, and the aerospace industry has pursued lightweight design ever since. But meanwhile, in Dearborn, Henry Ford was gearing up to produce the Model T, which launched in 1908, and which transformed the world of motoring. For the first time, a car was cheap enough that the people who worked in Ford's factory could afford to buy one. In transforming the car from a luxury product to an affordable one, Henry Ford set in train the whole history of 20th century manufacturing, converting luxuries into commodities, and he did it by ruthlessly pursuing standardisation. Famously Ford's offering was "any colour so long as it's black." In effect Ford discovered and exploited the economies of scale in production: making a high volume of identical parts and goods is significantly cheaper than making a wide variety, because there is no delay moving on from one part to the next, and the people, tools and systems to make parts all improve with experience.

So, in our snapshot of 1903, on the one hand we have the Wright brothers doing anything possible to reduce the weight of their vehicle, and on the other, we have Henry Ford doing everything possible to standardise his. Standardised parts are generally heavier than optimised ones, and that sets up the story of this chapter:

can we use less metal by optimising the design of components? If we optimise the design, what will it cost us?

We'll start with optimising weight, and try to establish guidelines for designing parts with minimum weight. As most components are not perfectly optimised, we'll examine a set of case studies to see how metal saving plays out in practice, and try to understand why we don't minimise weight at present. We'll then revisit our design principles to develop some practical guidelines, and use these to estimate how much of the world's metal could be saved by better design. Finally, we'll look at the business case for saving metal: metal costs money, so why wouldn't you take every opportunity to use less of it?

Basic principles of lightweight design

Since the 1970's the subject of 'structural optimisation' has developed computer aided tools to design components of minimum mass. This is a wonderful subject[1], but mathematically intense, so always demands the best available computers. For any specific problem, it's unlikely that we could beat the computer by hand, but if we rely on the computer, we won't learn how we might try our hand at some other problem. As a result, apart from aerospace applications, optimisation is rare and usually limited to small parts that move rapidly. For example it's worth optimising the heads in an inkjet printer, because reducing inertia allows increased print speed. So in this section we'll try a different approach, and see if we can learn some general principles.

Figure 12.1 shows the simplest example we can imagine, a point load supported by an arm (a cantilever) some distance from a strong stiff wall must deflect less than some limit. This picture might represent a crane, a balcony on a building, or the arm of a robot. We must also ensure that the arm is strong enough but, usually design for stiffness is more demanding than design for strength.

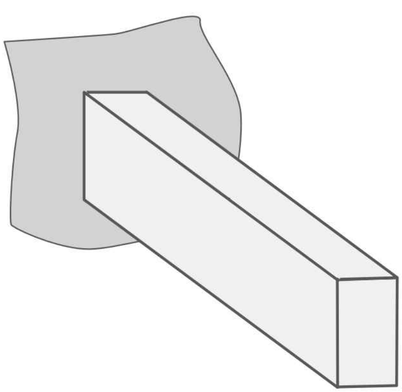

Figure 12.1—**Simple cantilever example**

The arm in Figure 12.1 is uniform, so it is most likely to break at the wall. For a stronger arm, we'd want to make it deeper nearer the wall, and by similar logic, less deep near the load. It turns out that if we want stiffness not strength, the same logic applies. Figure 12.2 shows a more optimised design: the depth of the arm varies so that the arm is as stiff as possible for the load at its tip. This design is 16% lighter than the first one: already good news for the Wright brothers, but bad news for Henry Ford as it's going to be more difficult to manufacture.

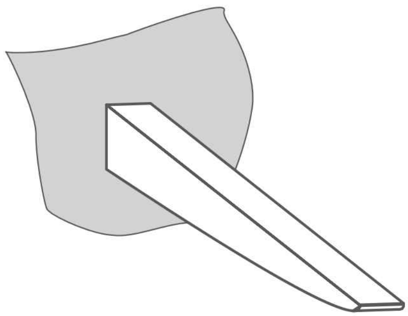

Figure 12.2—**Rectangular beam with depth optimisation**

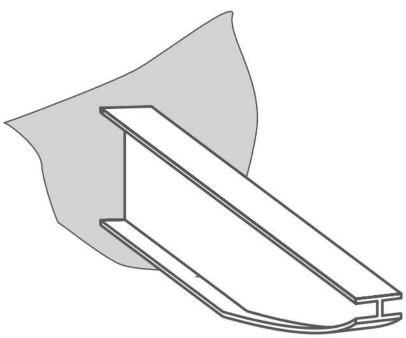

Figure 12.3—**Depth optimised I-beam**

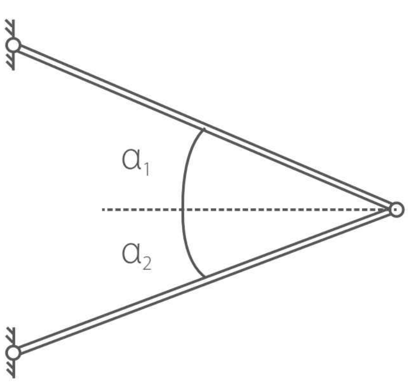

Figure 12.4—**Pin jointed truss**

Although the first beam would fail at the wall, the second one should fail all along its length at the maximum load. However, the whole beam does not fail at once. Failure will start at the upper and lower surfaces. You can easily show this if you have a packet of spaghetti to hand, and don't mind a bit of sweeping up: grasp the bundle of spaghetti firmly at each end, and steadily bend it into a curve of increasing severity. Which is the first strand of spaghetti to snap? It is always the one on the outside of the bundle, the strand at the centre is least likely to snap. Similarly with our beam, failure is most likely at the upper or lower surface, so this is the most useful place to have material, and we can make the middle of the beam thinner. Our third design in Figure 12.3 combines this arrangement of material with the design in Figure 12.2 and now the cross-section of the beam looks like a capital "I". This is the standard form in which structural steel is used in buildings but, because of Henry Ford's concerns, we usually use constant cross-section I-beams, not the variable type we've shown here. I-beams are usually made by rolling with specially shaped rollers, and it's a lot more convenient to make them with the same cross-section along their length. If we'd converted the first design into an equally stiff I-beam but with constant cross-section, we could have saved 54% of the mass, but with the variable depth I-beam in the third picture, we've now saved 85%.

Our variable depth I-beam is beginning to look like two spars working towards a point, and resembles a 'truss' which is familiar from the roof supports of large span buildings such as airports, and from railway bridges. So now let's move to the fourth picture, the simplest possible truss. We have two choices in this design: what's the angle between the two bars in the truss, and what cross-sectional area should they have? For the stiffest design, the best angle between the bars is a right angle, 90° (or equivalently in our picture, angle α_1 and α_2 should be 45°), if we assume that the structure is symmetric. Strictly, the lower strut is in some danger of buckling so might need some extra bracing. At this angle, the distance between the two supports at the wall will be twice the distance of the load from the wall, which may be a problem if we need to conserve space beneath it, but we'll worry about constraints later.

The truss design turns is extremely efficient. If we pick up the bundle of spaghetti again, and grip the two ends firmly, but now just pull, then each strand experiences the same load and is equally likely to fail. This means that we are using the material perfectly efficiently: we will always use less material if we can align loads with members to avoid bending. If the hinges in Figure 12.4 are frictionless, the loads in our truss align perfectly with the spars, so we can use members with constant cross-sections (good for Henry Ford) with just enough area to take the load, and no material wasted (terrific for the Wright brothers).

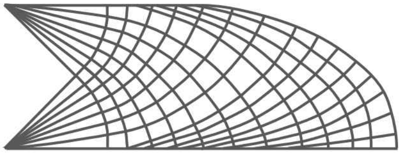

Figure 12.5—**A Michell truss**

We've nearly finished our simple case study, but will end with a remarkable ideal design developed by Anthony Michell—an Australian lubrication engineer, who invented the widely used 'Michell Bearing", set up the Michell Bearing company and as an aside explored the design of 'minimum frames'. In keeping with the theme of the chapter, Michell did this work in 1903 and showed that all minimum frames comprise bars which "form curves of orthogonal systems." So our final solution to the example problem is one of Michell's trusses illustrated in Figure 12.5. The design comprises two 'fans' of lines which always intersect at right angles—that's the meaning of "curves of orthogonal systems." For the Wright brothers, this design looks very interesting, but for Henry Ford rather less so, as the truss requires a complex set of ever shorter bars.

Our simple example has established two key principles for using less metal:

- Avoid bending by using trusses. Truss spars loaded along their length are always more efficient than members loaded in bending.
- If a member must experience bending, it should be designed (like an I-beam) to have the material as far away from its bending axis as possible. If the bending varies along the member, it should have a variable cross-section.

These two rules give us a great starting point for trying to design components with less metal, but before we go on to our case studies, we can add three further principles to guide our search for metal savings.

Firstly, once we've chosen our basic efficient design, we must choose the cross-sectional areas of each member, and as the loading increases, or the required deflection reduces, so the area must increase. Therefore, prior to starting our design, we should aim to reduce the required loads and increase the allowed deflections as much as possible. This seems rather obvious, but in practice, clients or letting agents will often over-specify requirements "to be on the safe side." That's comforting when we're in an aeroplane, but may be quite wasteful if our office is strong enough to have a swimming pool on each floor.

Secondly, our simple example required that we support just one load. What if we also had a second load that should be supported from the same base? Should we support it with a second independent structure, or should we support both loads with the same structure? In most cases we'll save material using one combined structure rather than two separate ones.

Thirdly, we haven't yet discussed the material we're going to use to make the arm, but of course if we use a stronger stiffer material, we'll generally use less material. Material selection is a big topic, because as we saw in chapter 3, there are so many properties we might consider. But fortunately, and as before, we can return to our colleague Professor Mike Ashby, who's book and associated software and databases show us how to choose the best material[2]. We'll illustrate his approach by going back to our simple truss design solution to the example problem. We said that the design must withstand the required load without failing (strength limit) and without exceeding some deflection (stiffness limit). If we set the load at the strength limit of the material, its deflection will decrease as the stiffness increases, or equivalently we can say that the deflection is proportional to the strength limit divided by the stiffness. Our chart in Figure 12.6 shows the properties of a few materials plotted on axes of strength against stiffness, with contours along which the ratio of strength to stiffness is constant.

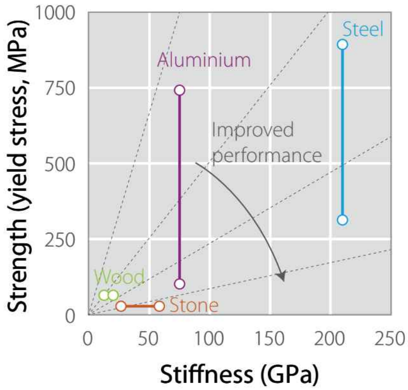

Figure 12.6—**Material selection chart or Ashby chart**

Materials on the same contour, such as steel and aluminium in this case, give similar performance and the arrow shows the direction in which we should move to choose the best material. According to the graph the best material is stone. We could perform a similar search considering other properties: density, for example if the weight of the truss is a significant component of its loading, or cost, or availability and so on.

We've now done enough with our simple example to derive five principles for designing products and components with less metal:

Five principles for using less metal

- Support multiple loads with fewer structures where possible
- Don't over-specify the loads
- Align loads with members to avoid bending if possible
- If bending is unavoidable, optimise the cross-section along the member
- Choose the best material

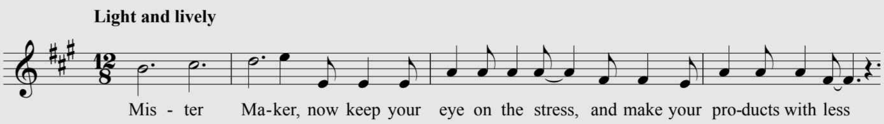

How can these principles be applied in reality? To find out, we'll look at what happens in practice today.

Case studies to explore using less metal in practice

To explore the reality of material saving through efficient design, we've examined five case studies—universal beams, deep-sea oil and gas pipeline, car bodies/crash structures, rebar, and food cans. Globally, annual production of these components uses around 400 Mt of steel and aluminium, nearly 40% of the total production of the two metals. We spent time with the companies making these components to learn about current practice, then we applied our principles from the previous section to propose a new lighter weight design. Then we went back to the companies to see what they made of our suggested change.

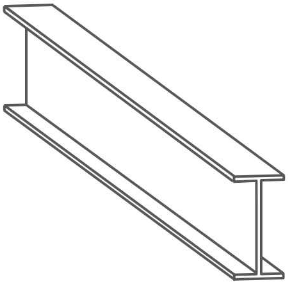

Figure 12.7—**A standard universal beam**

Standard universal beams as illustrated in Figure 12.7 are the key components of steel-framed buildings. They are designed for bending stiffness, and used as horizontal beams to support floors or roofs. They are manufactured in a standardized set of geometries, and listed in catalogues provided by steel producers. However, their geometry, which has a constant cross-section, is chosen for ease of manufacture, so is not perfectly efficient, as we saw with our simple example in the previous section.

In this case study, in order to estimate how much metal we could save through optimised design, we designed a series of beams to cope with a set of standard load cases and then evaluated our findings with experts in the construction industry. Our different beam designs are shown in Figure 12.8 and comprise: standard I-beams; composite floor beams where the concrete floor slab is part of the bending system allowing use of smaller steel sections; open-web joists which are truss structures suitable for lighter loads such as roofs; cellular beams where shaped cells are cut from the web of the beam to save weight; variable cross-section beams where the beam depth or width varies and is optimised for a given loading.

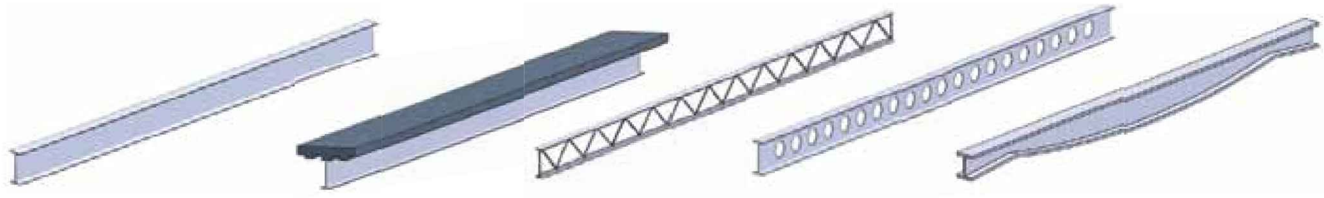

Figure 12.8—**Comparison of different beam designs**

Structural engineers design buildings according to codes of practice written by standards agencies to ensure building safety. In the UK, designs must satisfy European building codes (Eurocode 3 for steel design). So we optimised our beam designs[3] to satisfy these codes in two contexts: for floors and roofs. Our results are summarized in Figure 12.9, showing the weight of steel required for each design.

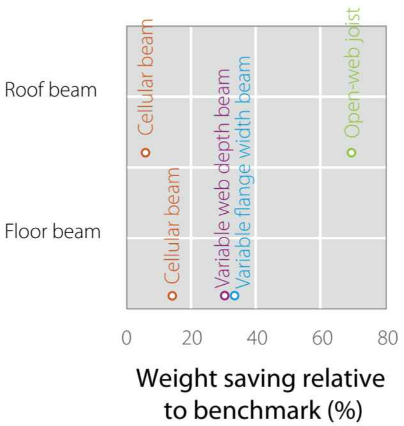

Figure 12.9—**Weight comparison of alternative beam designs**

Composite floor beams are the most common solution for floors in the UK at present, so used these as a reference for the floor load cases. For the roof load cases the current standard design is a standard universal beam. Our results show that weight savings of at least 30% can be achieved by applying our design principles, and higher weight savings are possible in cases where composite floor beams are not currently in use. When we discussed these results with building designers, they said that the improved designs were technically feasible, but would cost more. Clearly this is true, but if Boeing can assemble millions of parts to make a 747, it can be done—and, remembering the brief sketch with which we began chapter 6, the additional cost will be relatively small compared to the total cost of the building.

So, with the Wright brothers, we might be able to reduce the amount of steel in steel framed buildings but what would Henry Ford think? We mentioned earlier that I-beams are currently made by hot rolling with special roll shapes, so when we started thinking about this issue, we also got to work in our lab, and have found a new way to roll optimised I-beams[5]. We hope that Henry would approve.

Although in this case study we've aimed to optimise the cross-section of beams subject to bending, let's not forget our other design principles: we should combine loads, and avoid over specification. The box story on the Velodrome at the 2012 London Olympics on the next page shows that combining loads (supporting the seats and the roof in the same structure) was a key strategy in delivering a materially efficient building.

Our principles also tell us not to over-specify loads. Over-specification occurs in construction because of a process called 'rationalisation'. Typically, a keen young civil engineering graduate might design a building according to the standard codes, and choose the optimum beams. A wizzened old hand then reviews the design, and reduces the number of different beam sections required, because it simplifies life for the contractors who build the building. The cost of steel is low compared to the cost of labour in developed countries, so it's generally cheaper to save labour (by avoiding variety on the construction site) than to save material.

London 2012 Olympic Park

As CO_2 emissions related to the use of buildings are reduced through energy efficiency measures, more attention is focused on the embodied carbon emissions in construction. At the London 2012 Olympic Park more than 90% of embodied carbon is in just three construction materials: concrete, reinforcing steel and structural steel. Each material accounts for approximately 30% of the total. An effective means to reduce embodied carbon in construction projects is to set targets early in the design, preferably in the brief. We found two different stories at the Olympic park.

The architects for the Velodrome had a vision to build a minimum structure building 'shrink-wrapped' around the sport and spectators. As a result the geometry was governed by the track layout and required sightlines; this 'saddle' shape allowed use of a lightweight cable-net roof system where the steel is used in tension to span 130 metres between supports. Despite initial concerns about costs and risks, the contractor could save money and time by using this system and the client approved. The cable-net roof saved 27% of the steel that would have been required in an alternative steel arch option. An advanced dynamic analysis of the seating structure showed that combining the roof, stand and façade support systems, gave performance within accepted limits despite being lighter than code recommendations.

The contract for design the Aquatics Centre was awarded to a signature architect asked to design an iconic building for the London 2012 Games. The roof is a key element—'an undulating roof sweeps up from the ground as a wave'. The shape of the roof could be supported only by a conventional truss system. This was optimised during design but is still over five times as heavy as the roof of the Velodrome's, which has a similar span and area.

The story of these two stadia at the London Olympics demonstrates that specifying lightweight design early in a contract allows significant material savings: finding a favourable form at the start yields greater savings than highly refining a heavier option later on.

Our initial studies, with confidential data, have shown that this two-stage process of rationalisation leads to significant extra use of steel.

So in our exploration of beams in construction, we've see that there are significant opportunities for using less metal, but we don't currently pursue them because of the relative costs of materials and labour.

Sections of pipeline

Deep-sea oil and gas pipeline connects off-shore drilling rigs to shore, and may be installed more than 2 km under the sea. At these depths, a pipe is subjected to a very high water pressure which would tend to crush it if empty, but when in use, oil or gas is pumped through the pipe at a pressure similar to the external pressure. The oil and gas inside the pipe supports the pipe wall, which experiences only a small pressure difference, and could therefore be quite thin. In this case, the pipe wall-thickness would be chosen to avoid dangers from corrosion.

However, it isn't just the use of the pipe that determines the amount of steel required. We also have to solve a different problem: how do you install a 250 km length of pipe two kilometres under the surface of the sea? 2000 metres of water creates a pressure of approximately 200 atmospheres. At that depth there's no sun light, the temperature's around -5°C and you might bump into a Humpback Anglerfish, so it's a difficult environment for lining up and welding steel pipes.

A humpback anglerfish
(Melanocetus johnsonii)

As a result, deep sea pipes are not installed in-situ, but instead are dropped down from the surface. Typically a pipe is made of 30–50 mm thick high grade steel plate. After hot rolling, the plate is cut into lengths, typically 9 m or 12 m, rolled and welded into pipe sections with lengths of around 10 m. Pipes of this type cannot be coiled without damage, so are laid from a ship that slowly steams out to the target oil or gas well by welding each new section onto the existing pipe. This pipe 'string' initially hangs between ship and shore, and as it becomes longer, slowly sinks down to the sea bed. The photo shows the Saipem 7000, one of the worlds' largest pipe-laying ships.

The Saipem 7000 pipe laying vessel[6]

During regular use, the pipe has to withstand only a small pressure difference across its wall. But during laying, the most recent section of pipe to be welded onto the string must support the weight of about 2.5km of the string as it descends to where it reaches the sea bed, and the pipe also experiences significant bending just before it settles. To reduce the weight of the pipe it is laid empty, so the buoyancy of the air inside the pipe reduces its effective weight. Even so, the loading on the pipe during installation, due to self weight and to bending as it is draped onto the sea bed, greatly exceeds the loading it will experience during service.

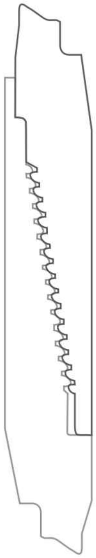

Schematic of a mechanical pipeline connector

Applying our principles for designing with less metal to this case study, we considered whether we could reduce the loads, and whether we had any other material selection options. We could try to reduce the loads during laying by pressurising the pipeline internally during installation, and our calculations suggest that this could reduce the weight of the pipe by around 30%. However, we learnt from industrial partners that the need for corrosion protection may reduce these savings to 10% and generating an internal pressure of 200 atmospheres in the pipe during laying may lead to safety risks.

Alternatively, as the loads during installation, not service, determine metal requirements, is there a different installation system that would allow a reduction in metal use? In shallower water, some pipelines are constructed on the seabed using mechanical connectors like those shown to the left. Could this practice be extended to deeper waters? Alternatively, could we make the pipe from a different material, either to increase its strength to resist crushing pressures, or to improve its corrosion resistance? Potentially "yes", but both solutions need extensive development.

This case study has revealed an important barrier to saving metal by design: line pipe in use today is overweight, not because of over-specification of the design loads, as we saw with beams in construction, but because the pipe experiences much higher loads during installation than in service.

A modern car body

Car body structures originally comprised a stiff chassis on which the body was erected. However, although this is still the basis of truck and heavy vehicle design, most cars today are constructed without a chassis but based on a monocoque: a cage around the passenger area giving better safety than a chassis based design of the same weight. When we look at a car, we can't see this monocoque. Instead we see cosmetic body panels which are only lightly attached so they can be replaced easily if damaged. Their main function is to hold up the paint! The priority design requirements for the monocoque are that it should absorb energy in a crash and provide sufficient strength and stiffness for normal operation. However, the monocoque has many additional functions, such as providing mountings for the engine, drivetrain and wheels, comfort for passengers, and an aesthetically exciting shape that attracts customers, and it must be possible to manufacture the body at acceptable cost. As a result, the design of a car body structure is extraordinarily complex, involving a series of trade-offs between competing intentions. If we skewed this trade-off towards saving weight, how much could we save?

The Lotus Seven lightweight vehicle

One very simple answer to this question comes from the great car designer Colin Chapman, founder of Lotus Cars in the UK. Chapman's design principles included the aim to "simplify and add lightness" and, famously, although we can't endorse this one, "any car which holds together for a whole race is too heavy".

At 500 kg, the Lotus Seven (now produced by Caterham) is one of the lightest cars on the road today. Its lightness is generally used to give very high acceleration, but in a recent competition for fuel efficiency a Lotus Seven was raced for economy, and with a change only to its tyres, and with a different driving style, achieved 160 miles per gallon. So Chapman's commitment to weight saving is inspiring in our quest for a more sustainable future both for reduced fuel consumption, and of course for reduced metal requirements.

We're hoping that our head of department will soon be buying us a Lotus Seven so we can learn more about lightweight design, but meanwhile we'll focus on one part of a more conventional body structure, the car door, and return to our principles for using less metal. The car door must be convenient for passengers, support a window, house various electronic features including loudspeakers, and resist impact in a crash. Applying our five principles in turn:

- **Support all loads with the same structure:** conventionally a car door is designed with the support structure separate from the 'door skin'. If these separate structures are combined, they can be made lighter.
- **Don't over-specify the loads:** the loads on impact are specified by national standards in crash tests, so depend on an average of other vehicles on the road. In future, separating heavy and light vehicles on roads would allow great improvements in safety for lighter weight vehicles.
- **Align loads with members to avoid bending if possible:** doors are supported round their perimeter, so are inevitably loaded in bending. The bending would be minimised if the door were as small as possible but customers prefer larger doors so this strategy has little short term potential.
- **If bending is unavoidable, optimise the cross-section and allow it to vary:** doors would be lighter if thicker in the middle, but this may conflict with passenger comfort or external aesthetics. The schematic to the side shows various alternatives to the conventional door design that would allow weight saving. These designs are currently inhibited by the need to withdraw the window entirely within the door.

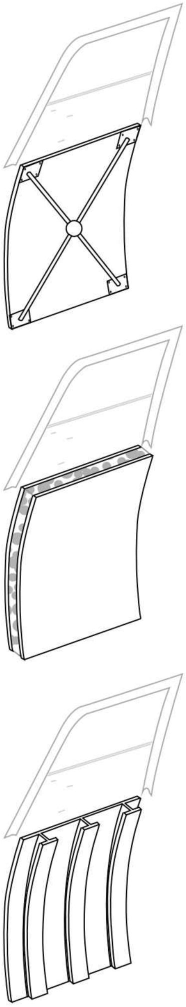

Alternative car door designs

- **Choose the best material:** Using Professor Ashby's material selection tools, we can examine a wide range of alternative materials for the car door. Carbon fibre composites or magnesium sheets could offer weight savings for equivalent energy absorption. However making these materials requires more energy, leads to greater emissions, and composites cannot be recycled. So there is a trade-off between emissions in production and use, and we'll explore this in detail in chapter 16 when we look at life extension. Manufacturing with composites is also more complicated than with metals so costs would increase.

Our principles have revealed many options for saving weight in car body structures. To validate them, we've spent time with a team at Jaguar Land Rover, who are reviewing their door designs, and are now aiming at a 30% reduction in door weight over the next five years. The details of their approach are confidential, but their ambition confirms that a significant weight saving can yet be achieved in these familiar and already highly engineered components.

Steel reinforcing bars

Steel reinforcing bar, commonly known as "rebar" is used extensively to provide structural reinforcement for concrete in buildings and infrastructure. Concrete is a ceramic, so strong in compression but weak in tension, and steel rebars are therefore embedded within it to provide tensile strength. The design of rebar is often constrained by strength, rather than stiffness as we saw in the earlier case study on structural beams, so if stronger steel is used to make the rebar, less mass is required.

Running through our five design principles for using less metal, two apply particularly to the use of rebar: material selection and avoiding over-specification.

In China, where an astonishing 60% of the world's rebar production is used at present, most rebar is made of relatively low strength steel, around two thirds of the strength we generally use in Europe. If we could upgrade all Chinese rebar, from the current mix of strengths to the best in Europe, we would save about 23 Mt or 13% of global rebar production. Why isn't this happening? Improving the strength of Chinese rebar requires a change in composition (in particular an increase in vanadium used in alloying). Vanadium is expensive, but even so this upgrade would reduce costs by around 20%. However, local producers are reluctant to invest in the equipment required for pre-straining, heat treatment and improved control, so Chinese rebar still has low strength.

Are we using the right amount of rebar? We discussed over-specification of loads in construction earlier, but a different issue arises with rebar, where even if the building as a whole has been designed without over-specification, designers, detailers and contractors may make choices leading to excess rebar use. This is because it is easier and quicker to lay out rebar in simple geometries, at a single spacing and with as few different bar diameters as possible. Of course simple layout also reduces the risk of mistakes and makes inspection easier. What would Henry Ford and the Wright brothers make of this?

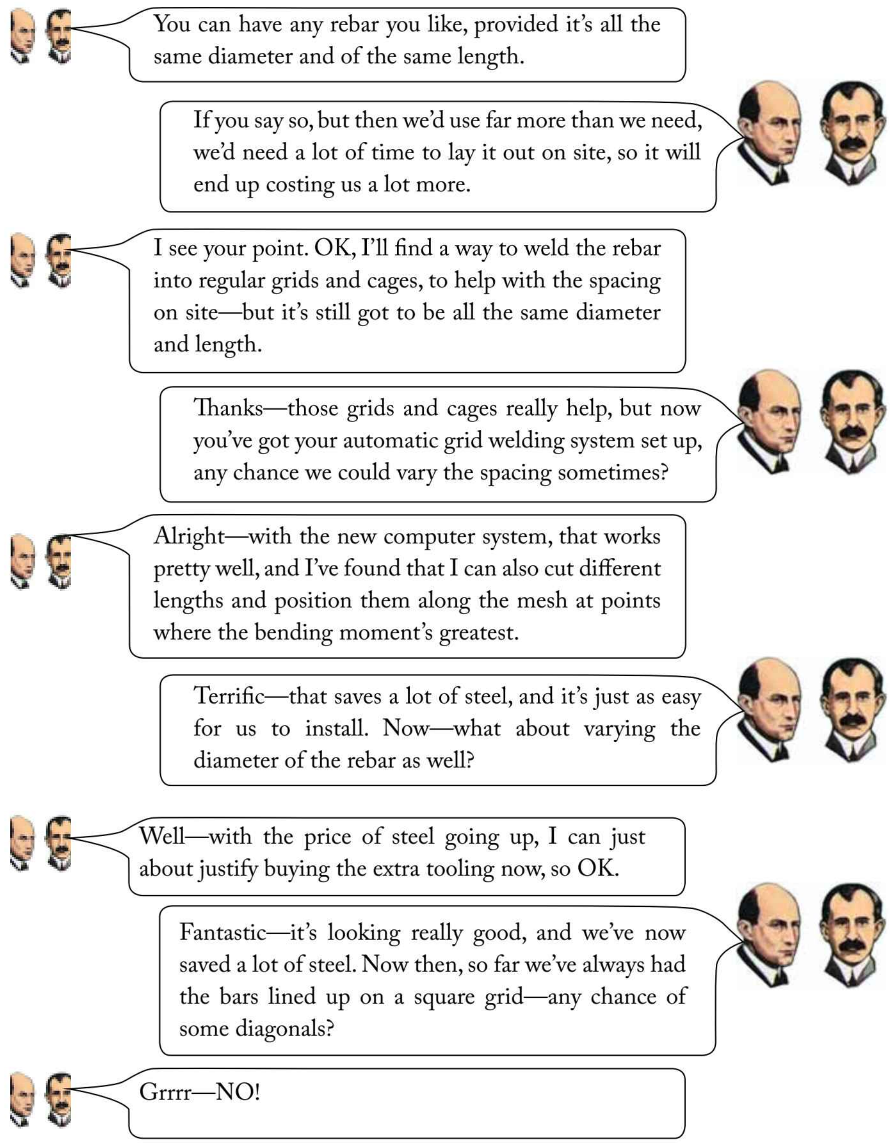

… and that's about where we've got to with the use of rebar today. It is designed carefully, and apart from the simplest applications, it's welded into grids and meshes to ensure the right spacing and separation. Modern computer control systems such as used by Qube in the box story, can design meshes with varying lengths, spacing and diameters, but as yet we still only use meshes aligned with square grids: no diagonals.

However that's the best of what's currently possible, not necessarily what happens in practice. After discussions with industry experts, we estimate that by truly optimising sizes and placement, we could save a further 15% of global rebar production, assuming optimised rebar solutions could be used in 65% of building projects and 50% of infrastructure projects. If we moved to non-orthogonal layouts, yet further savings would be possible, but with increased project complexity and cost.

Food cans experiencing greater loads in the warehouse than at home

Around 100 billion **food cans** are produced each year. In contrast to drinks cans, which have become lighter by around 20% over the past 30 years, food cans have had only modest decreases in weight, and remain around 30% heavier than a drinks can of equivalent volume and aspect ratio. Lighter cans could be produced using existing manufacturing equipment, but this has not been done. Why not?

The performance specifications of food cans are dictated by downstream processing requirements, where the food manufacturer fills the can, caps it, and then sterilises the contents in a cooking process known as "retorting", before stacking the cans to great heights for storage. During the retorting process, the can experiences an implosive pressure of around one atmosphere (the equivalent of being 10 metres below the surface of the sea) due to pressure in the cooking oven, followed by an explosive pressure of nearly three atmospheres as the contents heat and expand. Later, when cans are stacked in a warehouse the can must withstand the weight of all the cans on top of it—potentially as many as 50. Both features of this loading, which occurs before the can is sold to final customers, differ from the treatment of other food packaging such as aluminium pouches, plastic pots and Tetra Pak™. These are sterilized in a balanced retorting process at pressures of around a half an atmosphere, boxed instead of stacked, and handled more carefully. If the same were true for food cans, the can body could be 30% lighter, and in some cases can ends could be replaced by foil closures.

Once safely in customers' homes, cans need be no stronger than drinks cans, so just as with the deep sea oil and gas pipes above, their weight is determined by loads that occur before final use. It is possible to reduce these additional loads:

Reinforcing steel optimisation

Reinforced concrete designs generally include a degree of 'rationalisation' in the selection and layout of reinforcing steel, i.e. bars of the same diameter and same spacing are used across large areas to simplify detailing, identification, laying and checking of the installed reinforcement. This can typically add between 15% and 30% more reinforcing steel than is strictly required to meet performance and code requirements. Qube Design minimise this over-specification by using an advanced finite element approach for designing and detailing reinforcement using the Bamtec prefabricated rolled reinforcement carpet system. Bamtec carpets typically comprise smaller diameter bars (including in addition to normal stock ranges: 14, 18 and 22 mm diameter) which are placed at a reduced spacing to achieve the same reinforcement area required by the design. Bamtec 'rolled carpets' are robotically manufactured with the reinforcement read from the detailed drawings. Complex sequences of bars are used to significantly reduce the degree of rationalisation in the slab, without any loss of stiffness, and with increased crack control. Each bar is spot welded to thin gauge steel straps during manufacture and rolled up, for quick roll out on site. The rolled carpets, together with prefabricated edge curtailment and cages are mainly manufactured offsite. The combination of Qube's approach to design, and the Bamtec carpet system is an attractive example of intelligent innovation leading to real material savings[7].

the balanced retorting process used for foil food pouches could be used for cans, additional support could be provided for light cans in the existing process, and additional (reusable) supports could be used to allow can stacking with reduced total loads. We currently don't do this, because it's cheaper to pay for the additional metal. But if, for example, we had a choice between saving metal or paying for the infrastructure and energy costs of carbon capture and storage, metal saving might be much cheaper.

Practical barriers to saving metal by design and means to overcome them

Armed with our five principles for using less metal by design, we've found significant opportunities to use less metal in each of our case studies, but we've also found that we don't currently take these opportunities because of various practical barriers. Some of these barriers relate to cost: it can be cheaper to use excess metal than to pay the costs of using less. However, our ambition in the book is to look ahead to identify all possible options to cut emissions to 50%

of current levels by 2050, assuming demand for metal services doubles, and we expect that some of them will not be profitable immediately, or they would have been implemented already.

Let's summarise the barriers we identified in the case studies, and look at how we might overcome them:

- **Requirements before final use dominate design:** The service provided by metal components is often multi-faceted. Components may appear to be over-specified for their final use if their design must also satisfy other criteria: the food can must withstand higher pressures during retorting than on the shelf, and the deep-sea pipe experiences higher stress during laying than when pumping gas or oil. However, in both cases we've seen ways to reduce these additional loads to avoid adding metal: the can could be supported during retorting and stacking; the pipe could be joined on the sea bed rather than dropped in a string. So, in response to requirements prior to use: **look for alternative means to reduce loads occurring before final use**.

- **Asymmetric risks of using less metal:** It is generally cheaper to incur extra material costs for an over-designed component than to carry the risk of component failure. As a result, designers are inherently conservative, and in the long chains of companies involved in making final metal products, this conservatism is applied repeatedly. For example, we saw that the beams eventually used in buildings can be significantly over-specified after repeated rationalisation. The solutions to this issue are contractual, and depend on more precision in agreeing risks. For example, current building regulations in the UK specify *minimum* required sections or rebar designs to carry given loads. As a result, everybody involved is motivated to exceed the minimum of their predecessor. However, if the building regulations were changed to specify a target section or rebar requirement instead of a minimum, there would be no motivation to exceed it. So, in response to asymmetric risks: **write standards that specify target not minimum design requirements**.

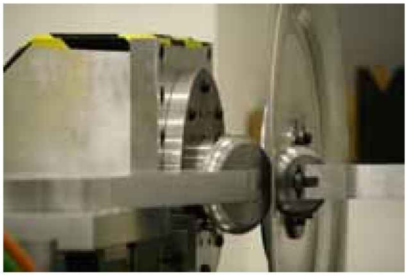

A novel flexible spinning process in our lab

- **Manufacturing minimum weight designs may cost more:** we saw that making variable section structural beams would save weight, but cost more to produce, and in our imagined dialogue between Henry Ford and the Wright brothers, we found that current use of rebar is a compromise between material cost and manufacturing effort. But we also found that there's space for innovation in manufacturing, creating flexibility in forming processes to produce more optimal designs with less metal[4,5]. We won't save metal if optimised parts are

Machining swarf

cut out of big blocks, as occurs in aerospace manufacture: the value of weight saving is so high for aeroplanes that material costs are irrelevant, and typically aircraft manufacturers turn more than 90% of the high quality aluminium they purchase into chips (called swarf). So, in response to concern about manufacturing costs: **develop new flexible metal casting, forming and fabricating processes**.

Two other barriers to using less metal, that came up in discussions with the companies in the case studies, are that customers may perceive lighter weight products as lower quality (this is a concern for luxury car makers for example), and that optimised components may be less robust than those with excess capacity. Both issues could be addressed by good design.

The key to achieving the material savings in these case studies is to foster more detailed collaboration among all the companies involved in converting liquid metal into final components. If product designers, component suppliers, manufacturing bosses, equipment makers, and the producers of intermediate metal stock products—in fact all decision takers between liquid material and final use—were to collaborate in the definition of material service requirements, the assessment of risk, and the build up of manufacturing costs, and evidence is that they could overcome all the barriers that prevent us using less metal.

How much metal could we save and how does this influence emissions?

	Global demand (Mt)	Potential savings (Mt)	
Beams	49	8-21	20–50%
Line pipe	25	3-8	10–30%
Car body	48	10-20	20–40%
Rebar	170	51	30%
Food cans	8	2	30%

Table 12.1—**Estimated weight savings for case study products**

The table summarises our estimates of potential weight savings in our five case studies, showing an average of about 30%. If we assume this estimate also applied to the remaining 60% of steel and aluminium use not covered in our case studies, using less metal by design looks like a dramatic opportunity for saving material: so potentially, we could use 30% less metal than we do at present, with no change in the level of material service provided, simply by optimising product designs and controlling the loads they experience before and during use. In fact, if our estimated reduction is applied across all metal using products, it translates directly to a saving in emissions: optimising designs could lead to a 30% reduction in all emissions associated with steel and aluminium production. In Chapter 11, we found that with one eye open we could save 10–30% of current emissions by efficiency in existing production systems, but now in our first chapter with both eyes open we've found an opportunity greater than this with just one strategy—although it's a strategy that won't be pursued by the metals industries themselves.

In fact the consequence of our 30% saving in metal production would actually be a greater saving in emissions than we've estimated so far, due to three co-benefits we identified in the case studies:

- In any application where a product moves, so particularly in transport, fuel consumption increases with vehicle weight, so lighter vehicles use less fuel as we saw in Figure 2.2. Fuel efficient cars are light weight cars. (Sadly, the recent history of car making is that we've reduced the weight of car body structures, mainly by using higher strength materials, but average car weight has increased, as we continue to want more luxury items in commodity cars. We all now expect air conditioning, electric window winders, great audio systems, and buttons to move our seats around, and no doubt within five years we wouldn't be seen dead in a super-mini that didn't have built in back-massaging in every seat…) If we could only stop compensating for weight saving by translating more of the features of luxury homes into our cars, we'd have lighter cars with lower fuel consumption.

- Lighter weight products may have improved performance: lighter cars accelerate, brake and turn better, lighter robots work faster, and lighter shipping containers can be lifted more rapidly.

- One lighter component may lead to another lighter component, compounding reductions in weight. This is true in office blocks where self-weight exceeds the weight of users, but also applies to oil rigs where the weight of the structure below the surface depends partly on the weight of the topsides, and to trains where lighter weight trains lead to reduced rail wear.

We can't estimate the impact of these co-benefits, but will return to the trade-off between fuel consumption and the use of metal in vehicles in chapter 16.

The business case for using less metal

A sustainable future may not be cheaper than an unsustainable one. But in some industries, for instance in making aluminium drinks cans, using less metal has been a core strategy for many years, driven by profit. So in this section, we'll

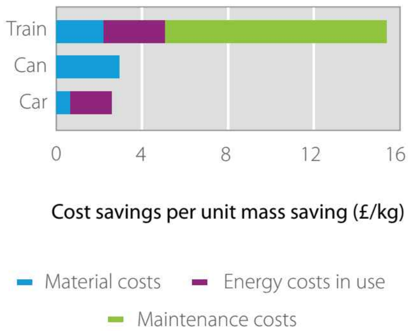

Figure 12.10—**Cost saving distribution for example metal products**

examine the business case for using less metal with three case studies: drinks cans, cars and trains.

To start, we'll estimate the lifetime benefit of using less metal. For each kilogram of metal saved, we save one kilogram of material purchasing. For the drinks can that's the only benefit, but as we saw earlier, the car and train also benefit from reduced fuel consumption over their whole life (say 10 years for the car and a 7 year franchise life for the train). The lighter train gives a further benefit through reduced track wear. Figure 12.10 shows how these costs add up to a predicted benefit to the final consumer.

It looks as if the train owner should have the greatest motivation to save weight, and the car owner and can purchaser should have equal motivation. Is that true?

Aluminium drinks cans today are 35% lighter than they were 30 years ago, driven by the fact that about two thirds of the cost of making a can is the cost of purchasing the aluminium. We use a massive number of these cans (in Europe alone we're using over 50 billion per year) so for the can making industry it's

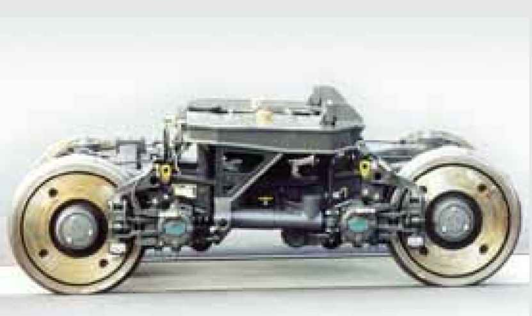

FLEXX Eco-Bogie

Based on early bogie development work by British Rail Research in the early 1990's, Bombardier's FLEXX Eco-Bogie (previously known as the B5000 bogie) is an example of component lightweighting in the rail industry. The integrated design reduces bogie weight by 30% (blue versus grey in plan view), saving approximately two tonnes per bogie. More importantly for track damage, the unsprung mass—that is the mass that is in direct contact with the rail with no suspension—is reduced by 25%, approximately 1 tonne per bogie. The FLEXX Eco-Bogie was developed as part of Bombardier's ECO4 Energy, Efficiency, Economy, Ecology initiative and aims to deliver savings in energy costs, network access charges and maintenance costs. Bombardier estimate that the new bogie results in a 25% lifecycle cost saving. In the UK, where expected track damage influences network track access charges, the lightweighted bogie is expected to save 17% of these charges in the 200 km/hr 16 tonne axle load class compared to a conventional bogie. The lightweighted bogie design is suitable for commuter, regional and high-speed rail applications. Over 1000 units are in operation worldwide. Further units are being manufactured for the Norwegian Railways (NSB) and for the new Bombardier Turbostar.

worth investing in research and development that leads to any possible saving. Surprisingly though, the contracts for can making link the can price to the material price: any reduction in weight achieved by the can maker reduces material purchased, saves money for the can purchaser but gives no benefit to the maker. However, can makers also complete with plastic bottle producers, so are motivated to continue reducing weight to maintain their share of the overall market for drinks packaging.

Cars have grown heavier in the last 40 years, and a typical family car is now around half as heavy again as its 1970's equivalent. The main reasons for this gain in weight are improved comfort, more features, improved performance, larger size and increases in safety. And of course these changes arise because they're what customers want: fuel efficiency is typically about ninth in a list of customer preferences, far behind performance, comfort, style and safety. Customers will pay more for cars with diesel rather than petrol engines—giving better fuel consumption—but only if there is no compromise in other features. If customers won't push the development of lighter weight (and hence fuel efficient) cars by preference, they must then be promoted by legislation—which is precisely what's now happening in Europe, with target emissions for new cars set at 95 gCO_2/km for 2020, compared to a current average of 145 gCO_2/km.

Trains in the UK, having had a constant weight during the 1980's, have become a quarter heavier in the years since to provide increased reliability (for example through having more powered vehicles in the same train), more air conditioning and passenger service systems, better safety and higher performance including tilting bogies. However, this increase in weight is surprising given the dual benefits of reduced power requirements for moving lighter trains, and reduced rail wear which in turn reduces track maintenance and replacement costs. We've found that the low priority given to weight savings for trains in the UK appears to be due to the way the rail industry was privatised: the track is owned by one company, the rolling stock by another, and the trains are operated by a third. So the rolling stock company wants versatile trains with high residual values (which are typically heavier), the track company would like lighter trains causing less wear, and the operating company wants to maximise profits during a relatively short franchise period which is not long enough for it to influence rolling stock development.

In summary, our three case studies have shown quite different motivations for or against reducing metal requirements in these three different industries, not at all linked to our predicted costs: it's not just the size of the benefits of saving weight that motivate change, but their size relative to other costs (only for the can maker was metal purchasing a large fraction of cost) the preference of customers for fuel efficiency against other features and the structure of the industry. So in contrast to our prediction that the train owner should have the highest motivation for weight saving, it's only in drinks cans that this has occurred: both trains and cars have become heavier.

Outlook

We've seen in this chapter that it is possible to define some simple principles for designing goods with less metal, and that if we apply the principles, it looks as if we could reduce global metal production by an amazing 30% without loss of final service. We've identified barriers to adopting this change, and shown they could be overcome, and have also seen that reducing weight has other benefits. However, we've seen that contracts, customers and industry structures may prevent the adoption of weight saving practices, and this suggests that we may need help from policy makers. Given that the massive implementation of carbon capture and storage would also require an input or two from policy makers this isn't too daunting, but we'll leave policy until we've completed our exploration of opportunities with both eyes open.

This chapter would have been much the less without 1903, the Wright brothers and Henry Ford, so let's not forget that in 1903 M.C. Escher, then aged five, moved with his family to Arnhem possibly clutching his first teddy bear invented in 1903 of course, and his box of Crayola crayons first made in 1903, with which he would learn to draw tessellations—which later become a crucial part of his artistic world, and which are central to our exploration of manufacturing yield losses in the next chapter.

Notes

Basic principles of lightweight design

1. Bendsøe and Sigmund (2003) give a thorough introduction to the field of topology optimization. For a more hands on introduction, you can try topology optimization for yourself at http://www.topopt.dtu.dk/.
2. Ashby and Jones (2005) provide a detailed analysis of material selection. Typically the designer will specify the key material parameters, such as strength or stiffness. By comparing these parameters among different materials and classes of material, the most suitable material can be chosen.

Case studies to explore using less metal in practice

3. The design cases we used were a 5 metre long beam taking a floor loading of 50kN/m or a 5 metre long roof beam taking a load of 7.5kN/m.

Practical barriers to saving metal by design and means to overcome them

4. Allwood and Utsunomiya (2006) give a detailed summary of flexible forming processes in Japan, many of which are now being more widely explored.
5. Carruth and Allwood (2011) describe our approach to rolling optimized I-beams.

Images

6. Image author: TeeGeeNo. Used under creative Commons Attribution ShareAlike 3.0 licence (http://creativecommons.org/licneses/by-sa/3.0/
7. Image thanks to Qube design