6 Process design: heat integration

P.J. FRYER

Introduction

The principles of thermal balances, heat and mass transfer and fluid flow developed in the previous chapters are well suited to the outline design of flowsheets and the design of individual plant items. It is much more difficult to consider the flowsheet as a design problem, and to decide whether the whole process is optimal. The simple economic analysis developed in Chapter 1 will tell whether a plant will make money, but will not produce better designs that make more money. To a large extent, optimization still depends on the skills of the engineer designing the plant. However, in some areas considerable progress has been made in developing design techniques that search for cost-optimal solutions to problems. These techniques are computer based and complex. However, one of them is conceptually very simple, and the principles can be developed without the use of computers. Any process plant will require some heat input and need some cooling. Some of this can be provided from within the plant, i.e. the cooling need on one process stream can be provided by the heating of another; but how much? It is possible, using thermodynamic principles, to find the maximum amount of heat which can be recovered from a system, and thus to design a process which needs the minimum amount of external heating and cooling. This chapter describes the problem, and then develops the solution using a worked example to demonstrate the method.

6.1 Design of process plant

The techniques of material and thermal balances described in Chapter 1 can be used to determine the mass and molar flows and the heat loads within a plant. However, this type of analysis does not indicate whether the plant is optimal: that is, whether a better design could be produced. There are several stages in the design of a food process plant. Each piece of equipment must be designed individually to carry out the job required; however, it is also important to consider the design of the whole system and the way

Chemical Engineering for the Food Industry. Edited by P.J. Fryer, D.L. Pyle and C.D. Rielly.
Published in 1997 by Blackie A & P, an imprint of Chapman & Hall, London.
ISBN 0 412 49500 7

in which the individual pieces of process plant can interact for example, making sure that the process is simple to control. The food industry has, historically, tended to design flowsheets one part at a time, and then assemble the parts into a whole. If plants are designed one part at a time, they can be inefficient.

As an example, consider the simple flowsheet shown in Fig. 6.1, in which:

1. a process stream is heated up to some required temperature,
2. a sterilization reaction is carried out at that temperature; and
3. the process stream is cooled.

In Fig. 6.1(a) the stream is heated and cooled by external steam and water. This requires two process units (heat exchangers) and expense in both heating and cooling. If some of the waste heat from the hot fluid is used to heat the cold fluid, however, as in Fig. 6.1(b), it is possible to reduce the heat

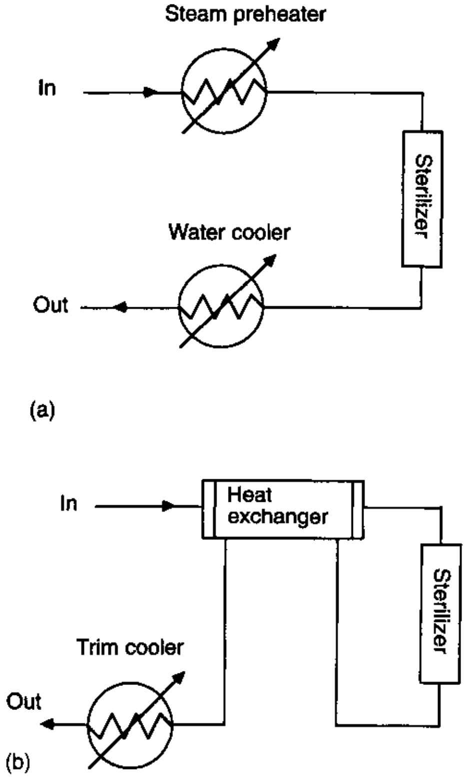

Fig. 6.1 A simple flowsheet for a process requiring feedstock heating, sterilization and cooling of product: (a) inefficient plant; (b) efficient plant.

load and thus save operating costs. It is unlikely that the heating and cooling loads will be the same; the figure includes a trim cooler designed to bring the outlet stream to the right temperature. Although the number of process units is the same, the second plant is more efficient.

The example in Fig. 6.1 is a very simple one in which the solution is obvious; less obvious would be the optimization of a case in which there were (say) 50 process units and 400 streams. Techniques have been developed in the petrochemical industry to consider the design of whole plants; these have been given the general name **process integration**. The use of these techniques to design a very large-scale plant can be very complex and involve the use of very large computer programs. Research into systematic design techniques is under way in both academia and industry. Some of the methods are conceptually very simple, however. The aim of this section is to introduce some of the techniques, using as an example the ways of making heat transfer within process plant as efficient as possible.

There are two basic design problems;

- **process synthesis**, in which the aim is to create a new process, and
- **process analysis**, where the aim is to understand how an existing process works.

The latter is easier than the former; there is no 'right answer' to a process design problem, and millions of possible solutions. The approach of chemical engineers has always been to start with a flowsheet for the process, and to analyse the flowsheet rather than considering the design of each element. It would be ridiculous to suggest that there is any infallible way of solving a synthesis problem using the same sort of easy steps. However, it is important when designing a plant to be able to screen out undesirable options as rapidly as possible, so that time can be spent as usefully as possible, considering only designs that are sensible.

Process synthesis proceeds in five basic steps: development, planning, basic design, detailed design, and improvements. A series of **design estimates** are made during the course of a project. The first stage of the design begins with an order-of-magnitude estimate, with costs accurate to about ±40%. These estimates are then refined through a series of design stages to reach a detailed estimate, which is given to the contractor, which should be accurate to within ±3%. The systematic approach to the preparation of flowsheets and then design estimates has become known as the **hierarchical approach** to design in chemical plant design. Although developed for the chemical industry, it is equally applicable to the food industry. It involves making a series of decisions, as follows.

1. **Should the process be batch or continuous?** The structure of batch process flowsheets is inherently different from that of continuous flowsheets, and it is necessary to decide which to use right at the outset.

Continuous processes are inherently more efficient, but many processes in the food industry, such as beer brewing or whisky distilling, are carried out batchwise because it has proved impossible to produce product continuously with the same texture or quality.

2. **What is the input/output structure?** At this stage an overall material balance on the process should be conducted using the principles described in Chapter 1. A control surface is considered that surrounds the whole process. Into this control surface go raw materials, and out of it come the products and by-products of the process. At this stage the first estimation can be made of the efficiency and thus the economics of the process and any possible environmental problems involved.
3. **What is the recycle structure?** As described in section 1.1.4 on material balances, recycle streams are often necessary: for example, when the conversion in a reaction is low, in which case the unused reactants are recycled back into the process. Requirements for recycles can be determined using the reaction kinetics of the process. Local material balances can be performed at this stage to determine all the flows within the system.
4. **Separation system.** In many cases in the process industries, products must be separated from the outlet stream and concentrated into an economic form. Once the conversion in the reactor is known, the type and size of any separation system needed can be worked out. In general, it is convenient to consider **vapour recovery** (condensation, ab- or adsorption, membranes) and **liquid recovery** (distillation, extraction, membranes) as separate systems. Most food engineering processes are based around liquids rather than gases, of course: only liquid recovery systems are needed.
5. **Heat integration.** Once the structure of the process is known fully, **heat loads** can be considered. Techniques have been developed that minimize the heat needs of a plant; these are outlined below.

The above type of approach is commonly used in computer-aided design processes; however, it is equally applicable to hand design and analysis. The process can be used in reverse to analyse the basic structure of existing flowsheets:

1. remove the heat exchangers;
2. group the distillation columns and separation systems;
3. analyse the resulting recycle structure;
4. consider the overall material balance.

6.2 Second-law analysis: heat integration

Most of the techniques of process design are complex, and beyond the scope of this book. One area, however, is very simple to study, and uses very simple physical principles. The design of structures of heat exchanger networks (HENs) that maximize the amount of energy recovery has been systematized. This is stage 5 of the above scheme. Mathematically it is possible to state the problem of energy recovery quite precisely:

> A set of cold streams initially at supply temperatures T_C^{IN} are to be heated to target temperatures T_C^{OUT} while a set of hot streams at T_H^{IN} are to be cooled to target temperatures T_H^{OUT}. Determine the structure of a network of heat exchangers that will bring all their streams to their target temperatures while minimizing the cost of equipment, steam and cooling water.

The problem is simple to state but much more difficult to solve. Even for very small numbers of process streams the number of possible permutations of networks is enormous, so using a computer program to consider all the possible alternatives is impossible in practice. It is necessary to find solutions that satisfy maximum energy recovery (usually described as MER) together with the minimum number of heat exchangers. This combination will minimize both the operating and capital cost of the flowsheet.

The insights made into this design problem have been developed from both thermodynamic principles and from the mathematics of network theory. The technique identifies pinch points in the system where heat transfer is most difficult; these are the most critical points of any network.

The analysis is based on the second law of thermodynamics. This is always regarded as much less easy to understand than the first law, but is based on a very simple observation: heat can only flow in one direction, from a system with a high temperature to one with a low temperature. The amount of heat contained with a body, which can be found from the first law, is not relevant to the direction of heat flow; a red-hot needle dropped into a swimming pool heats up the pool, despite the fact that the pool contains a much greater amount of heat. The first law of thermodynamics is a statement of conservation of energy; it makes no statement about the direction of heat flow. The second law defines that direction, via the definition of temperature. The core of the second law analysis is the experimental fact that heat can only move in the direction of decreasing temperature. It might be thought from a first-law analysis – the thermal balances studied in Chapter 1 – that heat at a low temperature could be supplied to a fluid at a high temperature; the second law says that this is impossible.

The second law governs the techniques of heat integration. The best way to study this is by an example. A test problem, of four process streams of defined start and end temperatures will thus be solved to demonstrate and explain the method.

6.2.1 Stage 1: identify the process streams

Table 6.1 shows four streams. The table includes the capacity flowrate of the streams, given as $kW K^{-1}$. The capacity flowrate, wc_P, is the product of the

Table 6.1 Process streams

Stream number	Supply temperature (°C)	Target temperature (°C)	Heat capacity flowrate (kWK^{-1})
1	400	60	0.6
2	210	40	1.0
3	20	160	0.8
4	100	300	1.2

mass flowrate and the heat capacity, and is the power required to raise the temperature of the stream by 1 degree.

Determination of the streams is in many ways the most important step in the method. The start and end temperatures are fixed, for example by:

- design requirement within the flowsheet (the need to bring a stream to a sterilization temperature will mean that the upper temperature is set);
- the temperature at which the stream enters the flowsheet;
- the temperature at which the stream leaves the flowsheet (which may be to waste or to some other process) – environmental considerations may determine this.

There is no need for the streams to be materially different: for example, the streams to and from the reactor in Fig. 6.1 are different streams in a heat integration sense because they start and end at different temperatures although they are the same in mass balance terms.

6.2.2 *Stage 2: thermodynamic analysis*

The example above consists of four streams, two to be cooled and two to be heated. The second law states that heat at any temperature can be transferred to a system at a lower temperature. Two **composite curves** can thus be constructed from the data, representing the net amount of heat available at a given temperature. The **hot composite curve** is the thermodynamic sum of the two streams that are to be cooled; there is

0.6 kWK^{-1} available at temperatures between 400 and 210 °C (stream 1 alone) then
1.6 kWK^{-1} between 210 and 60 °C (both 1 and 2) and then
1.0 kWK^{-1} between 60 and 40 °C (stream 2 alone)

The **cold composite curve** is the sum of the two streams that have to be heated; there is

0.8 kWK^{-1} required between 20 and 100 °C (stream 3) then
2.0 kWK^{-1} required between 100 and 160 °C (both 3 and 4) and then
1.2 kWK^{-1} required between 160 and 300 °C (stream 4 alone).

These two curves can be represented graphically on a temperature–enthalpy diagram, as in Fig. 6.2. The horizontal axis is an arbitrary one; the two curves can be moved sideways relative to one another. Any overlap between the two represents the amount of heat recovery that is possible: that is, the amount of heat that can be transferred from the hot streams to the cold. Figure 6.2 thus represents the case where there is no heat recovery: that is, where all heating and cooling duties are satisfied by external heating and cooling analogous to Fig. 6.1(a). In Fig. 6.3, there is some overlap, corresponding to the amount of heat recovery that occurs.

The vertical distance between the curves is the temperature driving force between the streams. This can be used in the basic heat transfer equation:

$$Q = UA\Delta T \qquad (6.1)$$

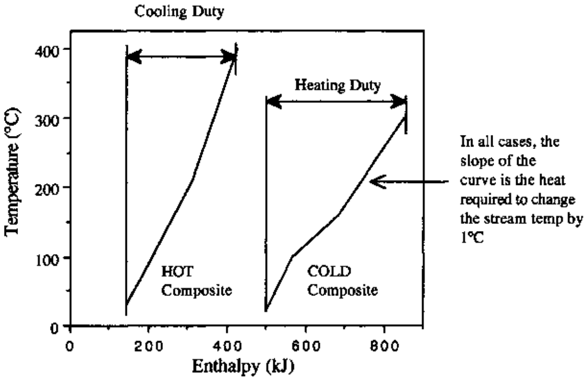

Fig. 6.2 Hot and cold composite curves, showing no heat recovery; all the heating and cooling must be done externally.

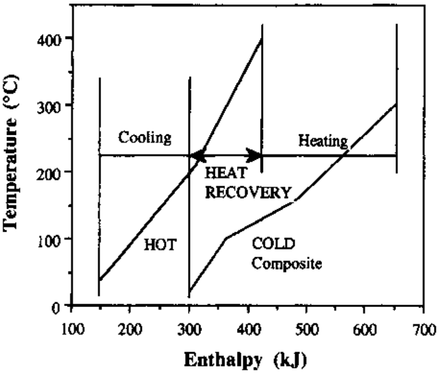

Fig. 6.3 Hot and cold composite curves, showing some heat recovery – overlap of the two curves.

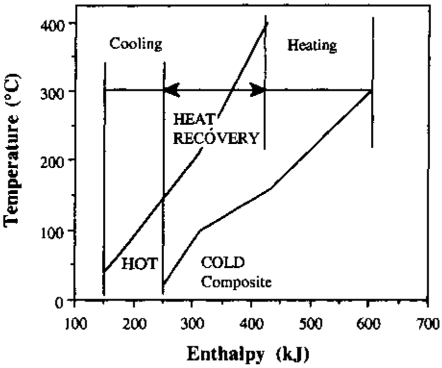

Fig. 6.4 The overlap between the two composite curves again represents the heat recovery; different approach temperatures give different recovery to that of Fig. 6.3.

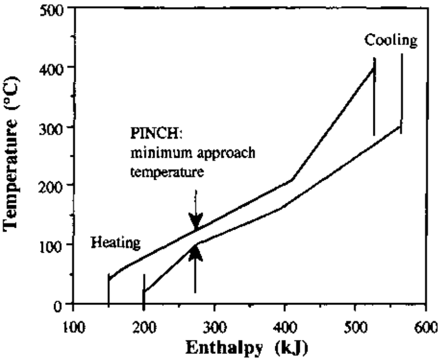

Fig. 6.5 At the chosen minimum approach temperature, heat recovery is maximized.

and is thus a measure of the heat transfer possible. It is only possible to transfer heat from hot to cold streams. Depending on the ΔT selected between the two streams, different amounts of heat can be recovered; compare Figs 6.3 and 6.4, where the overlap of the two composites is different, and the temperature between the two streams is also different.

The smaller the temperature difference between the two streams, the greater is the amount of heat transferred. The selection of a **minimum approach temperature**, or minimum temperature driving force (often written ΔT_{min}) is an important step in the analysis of the system. From equation (6.1) it can be seen that the smaller the driving force the larger is the heat transfer area. In the limit of no temperature difference between the two streams the area required becomes infinite, from equation (4.1). The chosen case corresponds to the composite curves moving until they are separated by the $\Delta T_{\min}$, as in Fig. 6.5. The point where the curves are closest is the **pinch point** for the system. If the two curves are overlapped further, then they will cross, implying that the fluid to be cooled will in fact be heated. This is the key point in the system, where heat transfer is most difficult. The core of the pinch design technique lies in finding this point – once the network is designed to ensure heat transfer in this region can be carried out efficiently, all other heat exchange will be easier. For engineering practicality, a finite $\Delta T_{\min}$ must be selected. Small $\Delta T_{\min}$ leads to a very high heat recovery, but low driving forces imply a large heat transfer area and thus capital cost. If a large $\Delta T_{\min}$ is selected, however, then the heat transfer area is small and the capital cost is low, but the heat recovery is less.

In the example, 20 °C is taken as the $\Delta T_{\min}$. To ensure that this requirement is not violated, and to simplify the arithmetic, it is common to adjust the temperatures of the streams; hot-stream temperatures have $\Delta T_{\min}/2$ subtracted from them and cold-stream temperatures have $\Delta T_{\min}/2$ added to them. This is purely a convenience, but makes the thermodynamics much clearer. It is now possible to plot the hot and cold composite curves on axes of enthalpy versus **adjusted** temperature, as in Fig. 6.6. At the pinch point, where in reality the temperature driving force is $\Delta T_{\min}$, the two adjusted curves touch: this makes the pinch point evident. This leaves an overlap at each end of the curve. Above the pinch point the two curves do not overlap; process heating Q_H is required to bring the cold composite curve to the required temperature. Below the pinch point process cooling Q_C is needed to bring the hot composite curves down to the required temperature. These two values can be seen in the figures.

The second-law analysis thus defines the minimum heat loads for a given approach temperature. For the example of Fig. 6.6, a 20 °C approach results

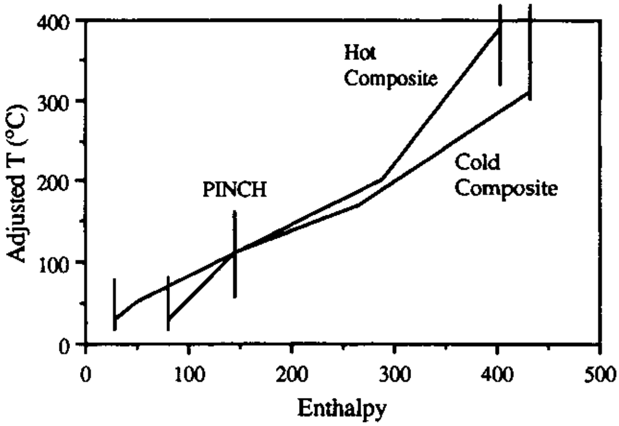

Fig. 6.6 Data of Fig. 6.5 replotted as adjusted temperature: at the pinch point, the two composites touch.

in a heating duty of 30kW and a cooling duty of 52kW; this can be seen from the diagram (but is difficult to read!). If cooling greater than Q_C is given to the system, this will lead to an increase in the process heat Q_H required: this effectively corresponds to the horizontal displacement of the composite curves. This involves a direct transfer of heat from Q_H to Q_C: the thermodynamics suggests that for MER no heat should be transferred across the pinch.

Manipulation of the composite curve in the manner described above is possible for simple systems, and is very useful to demonstrate the concepts of the method. It is less easy to use for larger numbers of streams. In general it is best to construct a **problem table** as in Fig. 6.7. The table includes all the streams at their adjusted temperatures. The left-hand side of the diagram shows the problem as a cascade, with streams shown vertically between

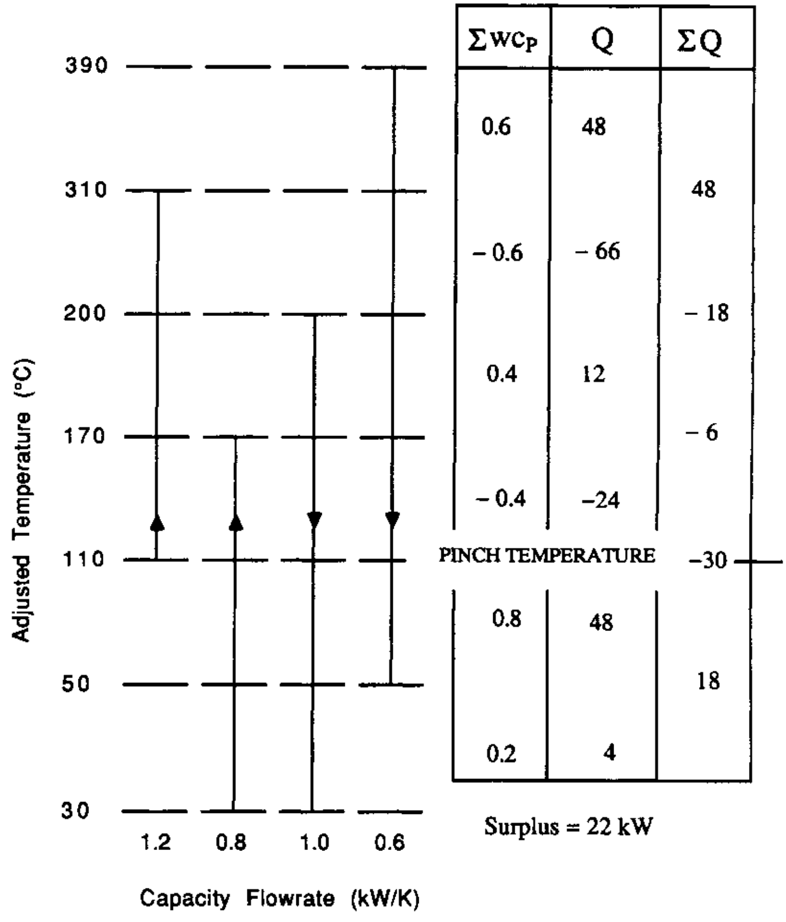

Fig. 6.7 Problem table for the example.

their adjusted temperatures. The diagram defines a series of temperature intervals: in each interval the net capacity flowrate and the net energy requirement are calculated with cooling as positive. For example, between 170 and 200°C, the net capacity flowrate is $(1.0 + 0.6 - 1.2)\,\mathrm{kW\,K^{-1}} = 0.4\,\mathrm{kW\,K^{-1}}$, and so the heat evolved by the system between these two temperatures is 12kW. At any temperature, the second law says that heat can be transferred only to lower temperatures. A heat cascade can be thus be produced with the energy requirements or surpluses from one interval carried down to the next: this is the ΣQ column in the table. The pinch point is that where the net flow is the most negative. As heat cannot be transferred 'uphill' – that is, from a low temperature to a hot – heat must be added to the top of the cascade to make the system operate; 30kW must be added as hot utility making the net flow of heat zero at the pinch, and rejecting $22 + 30\,\mathrm{kW} = 52\,\mathrm{kW}$ as cold utility. This again emphasizes that the two halves of the system are separate, and that no heat should be transferred across the pinch; if extra heating had been added as hot utility it would be transferred down the cascade out to the cold utility.

This method is in general an easier way to find the heating and cooling loads than the composite curve technique, especially for large numbers of streams.

6.2.3 Stage 3: design for MER

The thermodynamics has defined the maximum heat recovery; it is now necessary to design a system of heat exchangers between the streams that produces that recovery. The pinch-adjusted temperature is 110°C; the pinch temperature of the hot stream is thus 120°C and that of the cold stream is 100°C. No heat must flow across the pinch for MER. It is best to design the heat recovery system in two halves:

- **Above the pinch** all hot streams must be cooled from their supply temperatures to the pinch, and all cold streams heated from the pinch to their target temperatures. This must be done without utility cooling.
- **Below the pinch** all hot streams must be cooled from the pinch temperature to their target temperatures, and all cold streams heated from their supply temperatures to the target temperatures. This must be done without utility heating.

Design thus starts at the pinch and works out. To minimize the number of heat exchangers needed, it is important to match heat loads so that, if possible, the whole need of the stream is satisfied in one unit. Care must also be taken to avoid infeasible matches which give a temperature driving force below the minimum:

- exchangers just above the pinch should have $(wc_P)_H < (wc_P)_C$, so that heat transfer increases the temperature difference between the streams;
- those just below the pinch must have $(wc_P)_C < (wc_P)_H$.

Figure 6.8 shows the MER solution obtained following these rules. Each heat exchanger is represented as a vertical line with two circles connecting two horizontal streams; the number in the top circle is the number of the heat exchanger and the number in the bottom circle is the heat load in the exchanger in kW. The structure of the network is obtained using the following sequence of steps:

- **Cold end below the pinch:** All the heat needed to heat stream 3 can be supplied by cooling stream 2 in exchanger 4. No other heat can be used (as stream 4 begins at the pinch!) so the other cooling duties, 36 kW on stream 1 and 16 kW on stream 2, are supplied by cold utility, indicated by C.
- **Hot end above the pinch:** The whole heat duty of 3 can be supplied by stream 1 through exchanger 3, and the whole cooling duty of 2 can be supplied by 4 through exchanger 2. The rest of the duty on 1 can then be used to heat 4 by exchanger 1, leaving a 30 kW heating load to raise 4 to the correct temperature, indicated by H.

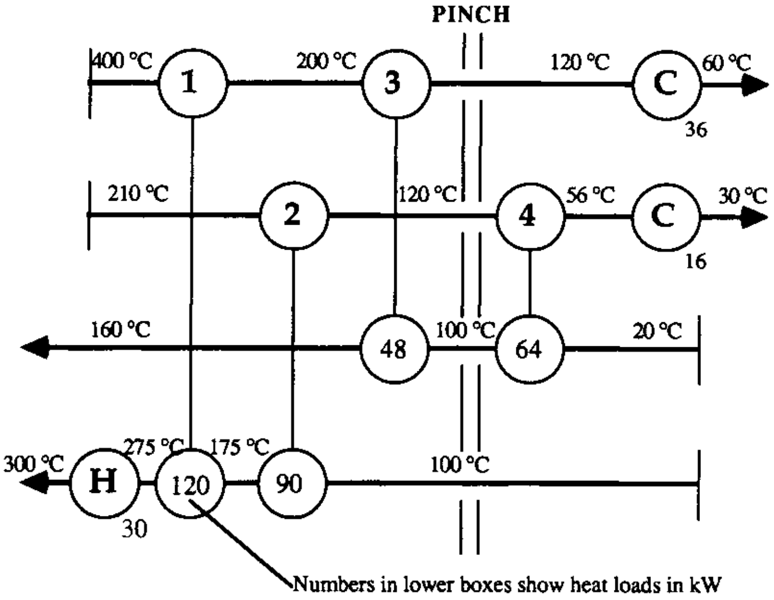

Fig. 6.8 MER solution for the example.

6.2.4 Stage 4: relaxing the solution

The above satisfies MER but contains four heat exchangers and three heaters and coolers. This has the lowest operating costs but may have an excessive capital cost. Network theory – a branch of pure mathematics – can be used to optimize the system, by minimizing the number of process units. The **minimum number of units** possible can be calculated; it is $N - 1$, where N is the number of streams plus the number of utilities, here 6: so the minimum number of units is 5. Network theory also gives a way of identifying which units can be lost, by identifying loops in the system as shown in Fig. 6.9. A loop is a set of connections that can be traced through the network that starts from one exchanger and returns to the same exchanger (for example, loop A in Fig. 6.8), or which passes through a utility, as does loop B, which goes through the two coolers and four heat exchangers. Breaking these loops is possible by removing units.

Design decisions are required here. The MER solution is optimal in terms of the heat loads; other criteria control what, if any, relaxation is made. To make the network controllable, it is probably best to leave the coolers intact. So the loop to be broken is that with the four exchangers. It is best to start by removing the unit with the smallest duty, here exchanger 3. Examination of the loop shows that removal will add 48 kW to loads on 1 and 4, while subtracting 48 kW from 2. This results in smaller ΔT than is allowed; extra heater and cooler duties are needed, as shown in Fig. 6.10. The result is a network with a smaller capital and maintenance cost, but which costs more to operate; the final decision, as in all things, must be made on cost, as discussed in Chapter 1.

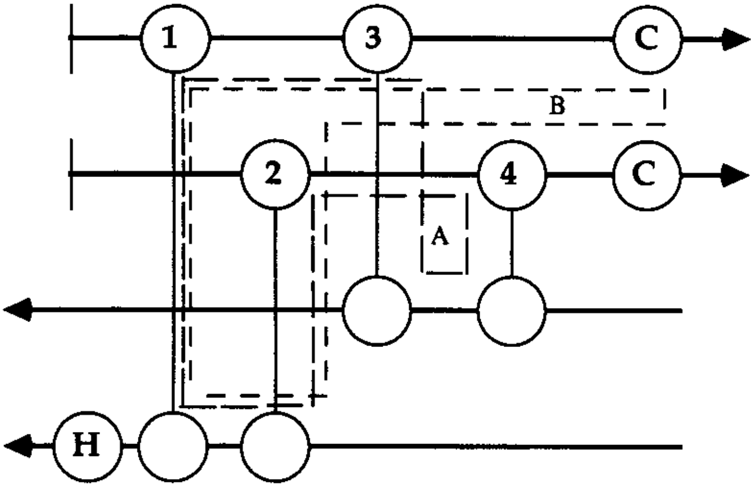

Fig. 6.9 Loops in the MER solution for the example.

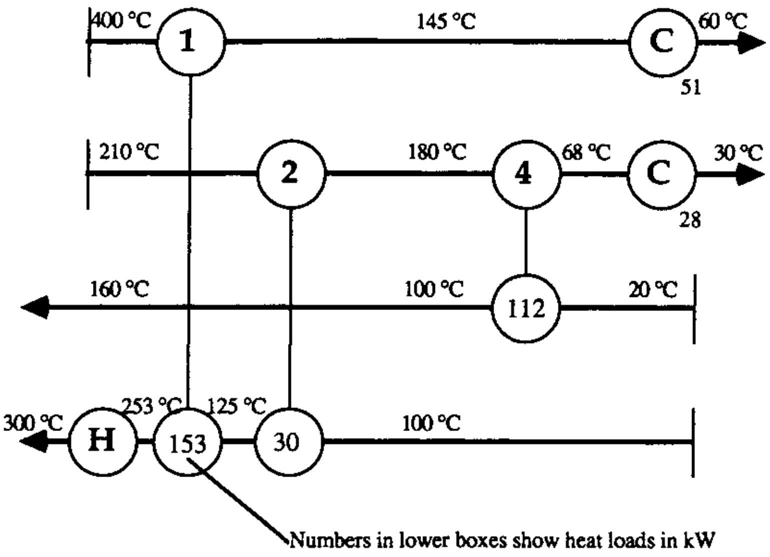

Fig. 6.10 Relaxed solution for the example.

6.3 Heat and process integration in the food industry

The above worked example demonstrates the use of the 'pinch' design technique. The technique has been widely used in the petrochemical industry and has been shown to yield significant energy savings. A petrochemical plant such as an oil refinery contains many heat exchangers and uses a wide range of temperatures, as in the above example. However, a typical food plant will probably contain fewer heat exchangers, and their inter-use may be limited. For example, there may be objections on safety grounds if cannery waste water is used to preheat material for canning; if the exchanger were to leak it is possible that microbes could cross between the two streams. Chemical process plant can involve a number of processes to produce many products; heat generated in one place can be used in another. Food plants tend to be smaller, reducing the opportunities for integration. However, some food plant is already integrated; a plate heat exchanger used as a sterilizer will commonly preheat the feed by exchange with hot product, along the lines of Fig. 6.1.

Another area of interest to the food industry is control of plant, as discussed in Chapter 7. Integrated plant is more difficult to control than unintegrated plant. For example, if the final network of Fig. 6.10 is examined it can be seen that stream 3 is heated only by stream 2. If stream 2 changed its temperature as a result of a problem elsewhere in the process, or if the efficiency of heat exchanger 3 changed as a result of fouling, then it would be difficult to control the 160 °C outlet temperature of stream 3.

However, if the heat were provided by process steam, reaching the required temperature would be simple (but more expensive!).

In practice, the final design may be easy to control and flexible enough to deal with changes in both input conditions and the efficiency of heat transfer; this is termed **process operability**. The pinch design approach is well suited to continuous plants with heat sources and sinks. Many of the operations of the food industry are still batch or semibatch, and when operation is continuous the plant is frequently shut down for cleaning. Unless great care is taken in their design, integrated plants are more difficult to start up and shut down than less thermally efficient systems. Deciding the optimal compromise between operability and operating cost is very difficult, not least because difficulties in operating process plant are not always obvious until the plant has been working for some time.

The integration of batch processes is more complex still: for example, if a biscuit factory has to make 15 products from four lines, what is the most efficient operating strategy, given the shelf-life of the product? It is necessary to design efficient operating schedules to maximize both usage of the plant and profit. This needs an understanding both of the profit resulting from each action and of the limits to the operation of the plant such as the time needed in cleaning shutdowns. This sort of information is common in the fine chemicals industry but less so in the food industry.

Note that, unlike much of the basic work described in this book, such as material and thermal balances, second-law analysis, and ways to develop optimal designs, are an active research area and are thus being continually developed. The example described above is deliberately simple; an attempt to explain the concepts of the method. It has been developed into a complex design tool which can be applied to designing very large-scale plant. More recently, a similar type of analysis has been applied to the study of waste streams from process plant, in an attempt to maximize waste recovery and minimize disposal costs.

Since this area is being actively researched, basic textbooks are scarce. One text which goes into detail of the design of chemical plant is:

Douglas, J.M. (1988) *Conceptual Design of Chemical Processes*, McGraw-Hill, New York.

Conclusions

The ways of calculating energy requirements, that is of setting up energy balances around processing operations, and their basis in the first law of thermodynamics were introduced in the first chapter of this book. Anyone faced with designing or operating a food processing plant will have come across important problems to which those methods don't appear to give an answer. Two such problems are:

- how can we make best use of the various hot streams throughout a plant?
- what is the absolute minimum amount of energy or work that is needed for a particular set of operations?

The second law of thermodynamics is needed to answer these questions, but the first of them can in fact be solved with little new theory beyond that developed on Chapter 1. This chapter has focused on that problem, that is, of energy integration, and has outlined the principles and techniques of 'pinch' technology which can be employed in its solution. The thermodynamic principle underlying the technique is that heat energy cannot flow 'uphill': in other words, heat is only transferred **down** a temperature gradient. Since the rate of heat transfer (and ultimately therefore the size of any heat exchanger) is proportional to the driving force it is not economically or technically feasible to attempt to reduce the temperature difference between a hot and cold stream to zero: in reality a minimum temperature approach (typically 15 to 20 °C) is chosen. The problem is then addressed by preparing composite curves representing the temperature versus heat energy or enthalpy of the hot and cold streams respectively, and adjusting these to maximize the heat recovery from one to the other. The pinch point is the point at which the composite curves come closest to each other, i.e. the minimum temperature approach. The method is illustrated graphically by a simple example, and its extension to more complex and more realistic problems is briefly outlined.

You should now appreciate the difference between the first and second laws of thermodynamics, and be able to analyse plant so as to recognize where opportunities for heat recovery might arise. Whilst the methods of handling complex problems are beyond the scope of this text, we believe that an understanding of the principles behind the technique is of great importance to anyone involved with designing or operating real processes, where energy conservation is economically and environmentally more and more important.