Accelerating Business and IT Change: Transforming Project Delivery

ALAN FOWLER AND DENNIS LOCK

GOWER

CHAPTER 4 Planning

Planning is the process of converting the business plan into a practical working document. This chapter considers the process from the point of view that failure is never an option, which means that the planning and subsequent control methods must be designed to ensure that all the intended project benefits are achieved in full and on time.

ATTRIBUTES OF AN IDEAL PLAN

A properly executed plan not only enables the coordination of people working on a project, but also the coordination of that project with other projects. The plan must be able to communicate what people have done, what they are doing now and what they will be expected to do. The plan is at the heart of the project; it is the most informative of all project documents and the common denominator for communication. However, planning is not always easy. Most projects are:

- complex
- subject to risk, uncertainty and change
- wanted as soon as possible.

The planner faces many problems. How can all the possible tasks that might arise be foreseen? How much detail should the plan contain? How should the plan be structured? How can interdependencies be managed effectively whilst avoiding gridlock? Project plans must be capable of meeting these challenges.

In the early years, between 1920 and 1950, planners had to be content with barcharts, now often called Gantt charts after Henry Gantt, the American industrial engineer who introduced them during the First World War. Since the 1950s planners have been able to express the complexity of projects far more effectively using critical path networks. Several variants of the critical path method were developed in America and Europe, but all allow the planner to record task interdependencies in a way that is not easily possible with Gantt charts. However, comprehending critical path networks requires some training; they are not a universally understood language and are therefore not suitable as universal tools of communication. For that purpose, Gantt charts still reign supreme.

Isochron simplifies the problem of multiple interdependencies by arranging project tasks in clusters of autonomous sets. Each set of tasks is bound internally by logical dependencies, but is loosely coupled externally with the other task sets. This minimizes dependencies between different paths in the network. Looking ahead to when the plan will be used for project control, the Isochron method recognizes that not every dependency shown on the plan will prove to be real when the time to start work arrives, and the project manager will often find (and, indeed, must strenuously seek) tasks that can be started with advantage before the network purports to show that they logically can. This approach (which will be explored and explained further in Chapter 8) can introduce an element of risk, but Isochron points out that the possibility of rework is often preferable to the certainty of delay.

Risk management (see Chapter 6) is a subject that has become better understood and more widely applied in recent years, and it is now accepted that planning should not be undertaken without consideration of project risks. A range of techniques has been developed to identify, assess and mitigate the effects of risk. In addition, procedures for controlling subsequent changes to the scope of a project are well known (see Chapter 9), although these are not so well reported in the literature. Whatever the procedures are for dealing with risk, uncertainty and change, the project manager must always be willing to question (and, if necessary, adapt or change) the solutions being followed. The planning method must be able to cope with all these possible sources of change. Thus it must be possible to change a project plan rapidly and easily in line with authorized project changes. The plan must also be updated whenever necessary to reflect progress. Plans drawn on paper or set up on wallcharts are not sufficiently flexible but, fortunately, plans processed by computer can be changed in a matter of minutes.

The best way to ensure that a project is completed on time is to manage the work effectively, using the methods described in Chapter 8. However, that can be done only if the plan provides a reliable 'road map' containing enough milestones (recognition events and value flashpoints) backcast from the end of the project to allow the project manager to question, measure, review and correct progress at relatively short intervals.

All of this raises the question of how much detail to include in the plan. Including a large number of trivial tasks serves only to complicate the plan unnecessarily. Having too few tasks, all with very long uninterrupted durations, is a step in the opposite direction of oversimplification. The experienced planner will know instinctively how much detail to include but a useful rule of thumb is to ensure that every task is interrupted (that is, divided into two or more tasks) at every point where it is expected to pass from the control of one manager, supervisor or technologist to another. Backcasting from the end goal (see the section 'Building backwards from the project benefits' later in this chapter) flows naturally from the use of recognition events and should be a constant guide for pruning out unnecessary detail and digression in the project plan.

As in so many other areas of management, Pareto's rule applies to project planning. Vilfredo Pareto (1848–1923) discovered during population census work that 80 per cent of personal wealth was distributed among only 20 per cent of the population. This has since become known as the 80–20 rule, in which 80 per cent represents the 'insignificant many', leaving 20 per cent as the 'significant few'. This rule is not confined to the distribution of wealth. It is universal in Nature. In work management, if you can focus your attention on the 20 per cent of the work that delivers 80 per cent of the desired results, you will increase your effectiveness fourfold. Taking this a step further, just 20 per cent of that 20 per cent will produce 80 per cent of the 80 per cent, and so on. The important point is that the same is true of the effort and cost invested in a project. Planning using the Isochron approach identifies and prioritizes the key 20 per cent of the project when it works backwards from the recognition events.

The plan must also be communicable. It must be possible to issue work-to lists, visually effective charts and other reports to managers and others after filtering and sorting the data so that managers get information only on tasks for which they are responsible. This is readily achieved by allocating appropriate departmental or resource codes to tasks in a computer system loaded with good project software.

SETTING THE TIME OBJECTIVES

Business plans should always be made before any technical solutions have been assumed, which means that the timescale for the whole project is almost always a given requirement. The outcome required by the business must be delivered, and the timing of the delivery costs and benefits is inevitably important. The first question for planning is not 'when?' but 'what?' and 'how?' and 'can we find a way to do it that makes it possible in the time allowed?'.

Identifying all the value flashpoints and recognition events is a matter of some special skill. In the absence of technical solutions, none of this initial planning can be performed using critical path networks or any other method that depends on detailed knowledge of all the tasks involved. So, even before the project manager has been appointed and the project authorized, the first stage of planning has to take place with very little detailed input. This is top-down rather than bottom-up planning, which means making a plan in coarse details from a project overview rather than attempting to build the plan up from all the detailed (as yet unknown) tasks.

The Isochron techniques involved in backcasting (abductive planning) are a very fast way of getting the planning started and 'right first time'. These techniques kill two birds with one stone by also delivering simultaneously a Pareto-based plan (see, again, the section 'Building backwards from the project benefits' later in this chapter). If the delivery dates of the recognition events and value flashpoints are treated as parameters that are given and key to the value of the project results, it is the question 'how?' (the nature of the approach) that has to be tackled in planning. If the project team were to be shackled to a single solution, where that solution has been given as if it were the objective, then there is a high risk that the deadline will become impossible. If the project end date is non-negotiable, then the project team *must* be given flexibility in choosing the solution. That is true 'agile' project management.

PENALTIES FOR MISSED DEADLINES

It is well known, empirically at least, that projects which run late also exceed their budgeted costs. Conversely, projects that are well planned and managed, and which are completed on or before their target finish dates, tend to run below budget. One important reason for this condition is that every project attracts an element of fixed costs – the costs of accommodating, servicing and managing the project. These fixed costs accumulate relentlessly for as long as people are working on the project, irrespective of progress made, and they will inevitably take the project into overspend if the work extends significantly beyond the date when everything should have been finished.

Engineering and other commercial projects can attract contractual penalties if they are delivered late, often calculated in direct proportion to the number of days or weeks by which the project overruns. Furthermore, when one project runs late, the organization often suffers knock-on effects, because resources and accommodation are tied up for longer than intended, delaying the start of projects and other work that should be following on.

For the kind of business change and IT projects with which this book is primarily concerned, late completion has other penalties. The business cases for such projects are built on the relationship between benefits and the times when those benefits should be realized. The net present value method of project evaluation described in Chapter 3 emphasized this relationship between time and benefits. Late completion of a business change project will delay the value flashpoints that should signal project success and, in a severe case, will result not merely in reduced benefits but in an actual loss to the organization. Such losses are often not simply monetary, but can also adversely affect staff morale and the company's reputation.

The British Library, the Holyrood Scottish parliament building and the air traffic control centre at Swanwick, Hampshire (Duffy, 2002) are just three examples picked at random from high-profile projects that finished well beyond public expectations of their costs and timescales. There are many (far too many) other equally well-known examples where expectations have been mis-set and mismanaged, often for political reasons, and where the projects have been delayed through poor preparation, planning oversights, uncontrolled changes, and contractual and financial difficulties.

BUILDING BACKWARDS FROM THE PROJECT BENEFITS

There are many high-profile projects where late completion would be disastrous and unthinkable, even though the projects are large, complex and fraught with risk. Think, for instance of a city's preparations for the Olympic Games, the Foreign and Commonwealth Office organizing a summit meeting of the G8 prime ministers, a royal occasion such as a coronation, a large public event like a pop music festival or the Chelsea Flower Show, or even an ambitious family holiday expedition. Although these are widely diverse examples, they all share one important common factor. Their finish dates are not negotiable. Failure is simply not an option (even if, sometimes, it comes perilously close). So how is it that these kinds of project do get finished on time, despite their size and complexity?

A large part of the answer lies in the determination of everyone concerned with the project to meet the end objectives. Planners' minds have to be directed not at determining how long these projects should take to complete, but instead at considering how they might, at the end of the project, have met their fixed 'wanted-by' dates; they must be directed not *towards* those dates, but must work everything *backwards* from them.

The obvious (but usually ignored) conclusion to draw from those public projects that traditionally are always ready on time is that the lessons learned should be applicable also to business change projects. If all the benefits set out in a business case are to be achieved in full and on time, then why not set their 'wanted-by' dates in stone and plan backwards from them? Isochron's objective transform technique does just that, using a particular human mental capability of memory and imagination that Tulving (1985) termed *chronesthesia*. In project management this means starting from the end value flashpoint and working backwards, first to identify and then to time all the intermediate value flashpoints and, back beyond those, the recognition events. This concept was introduced in Chapter 3, and is illustrated here in Figure 4.1.

The backcasting approach, together with other techniques described in this chapter, can be used not only at the start of a new project, but also to rescue an existing project that has run out of control. One example of this was a company with several product divisions that had booked a stand in a prestigious trade fair and was producing a number of products for display. These products included large steam and gas sterilizers, a de-Bakey heart-lung machine, electronic patient monitoring equipment for use in operating theatres and samples from a range of prefabricated operating theatre buildings. Six months remained to the exhibition and it was clear that no one knew the current state of progress or whether or not any or all of the exhibits would be ready on time. A 'crisis' project manager, directing a task force, was assigned to this problem. After discovering that the project was in a shambles, with no one having any idea of the current state of progress, let alone a practicable schedule, he decided to start from the only reliable constant, namely the opening day of the exhibition. His first question to a meeting of senior task force staff was:

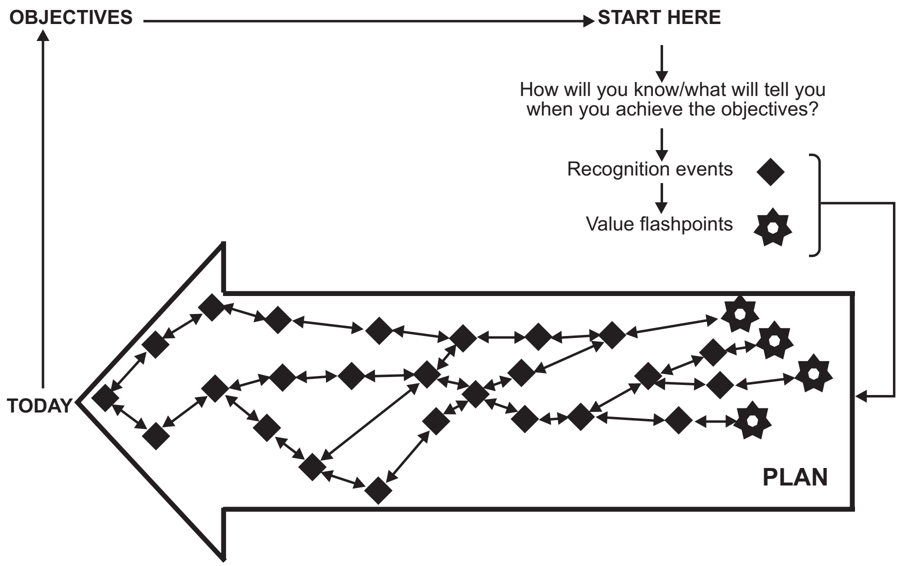

Figure 4.1 Recognition events, value flashpoints and project objectives

'What must we do immediately before the exhibition opens?' Answer (which had to be teased out): 'Final inspection and clean up of the stand.' 'What has to be done immediately before that?' produced more than one constructive answer, and the plan was worked backwards until the existing state of progress was reached to reveal a large number of (previously loose) ends. Using critical path analysis, the final schedule was then derived by working forwards, retracing all the paths and milestones that could only have been revealed by backcasting.

It is apparent from this example that planning needs imagination and, more than that, imagination informed by knowledge and experience. The people at this crisis meeting had to be coaxed and encouraged to use their collective minds to visualize the project at various stages in the future and work backwards from those. The best plans are always those that are reached by consensus, which has two important advantages:

1. Two minds being better than one, and brainstorming being better still, the collective brainpower will probably result in a plan based on the best technical process.
2. People who have helped to formulate the plan and who therefore agree with its implied method and milestones are more likely to commit themselves to its achievement.

Brainstorming with several minds is successful because, quite unlike a computer, the human brain is a *connection* engine. At a conservative estimate, the neuron network that powers an average brain has over *five trillion* connections. Thus the brain is extremely good at making connections – so good in fact that we do not hear all that ‘intuition’ tells us. Put two or more people together in a well-chaired meeting and the results will be vastly better than a plan produced by an individual working alone.

While brainstorming and the collective power of several minds can suggest a viable plan, critical path networks provide the notation that allows the plan to be expressed on paper and, eventually, in the computer.

So, for any new project built on the Isochron transform technique, we can imagine a newly appointed project manager assembling key members of the project team (business and IT) and asking them to think themselves into the future, when the project is complete. Then the manager asks them to cast their minds backwards from the value flashpoints, through the recognition events, until the project start is reached.

The product of the meeting will be a network sketch that contains all the significant known project tasks required to meet the needs and expectations of the sponsor and senior management – but no more! Once the duration of all those tasks has been estimated, the duration and critical path of the project can be calculated. This means making forward and backward passes through the network in a process called time analysis. The forward pass adds all the estimated task durations through each possible path in the network, from which the longest path indicates the earliest predictable start and finish times for every task and thus the minimum duration of the project. The backward pass mirrors this process by subtraction, to establish the latest permissible start and finish times for every task. For any task, the difference between these earliest and latest times is called the float (or slack), which is the amount of time by which one task can be delayed without delaying project completion. Tasks with zero float are critical tasks, and the path or paths that join them are also called critical.

Figure 4.2 summarizes the critical path method, using the example of an insurance company that has decided to move its customer services department to an offshore call centre. The network shown here has been kept simple for clarity, but it is sufficient to illustrate the basic steps of planning a project using the precedence system of critical path analysis. There are other methods of critical path analysis, but the precedence system is the only method supported by modern project management software. The precedence system and the other critical path methods are well described in the literature (Lock, 2003 and Gordon and Lockyer, 2005, being two of the many popular examples).

So far, so good. Working backwards from the value flashpoints should produce a practicable plan from which the project can be managed to achieve all the intended benefits at the times required by the business plan. All very straightforward – but with one enormous snag. When all the duration estimates have been applied, working backwards often finishes up not at the situation *today*, but *yesterday* (or even many days before yesterday). In other words, it seems that not enough time has been allowed to achieve the results because the critical path is longer than the time remaining from today to all the recognition events and value flashpoints.

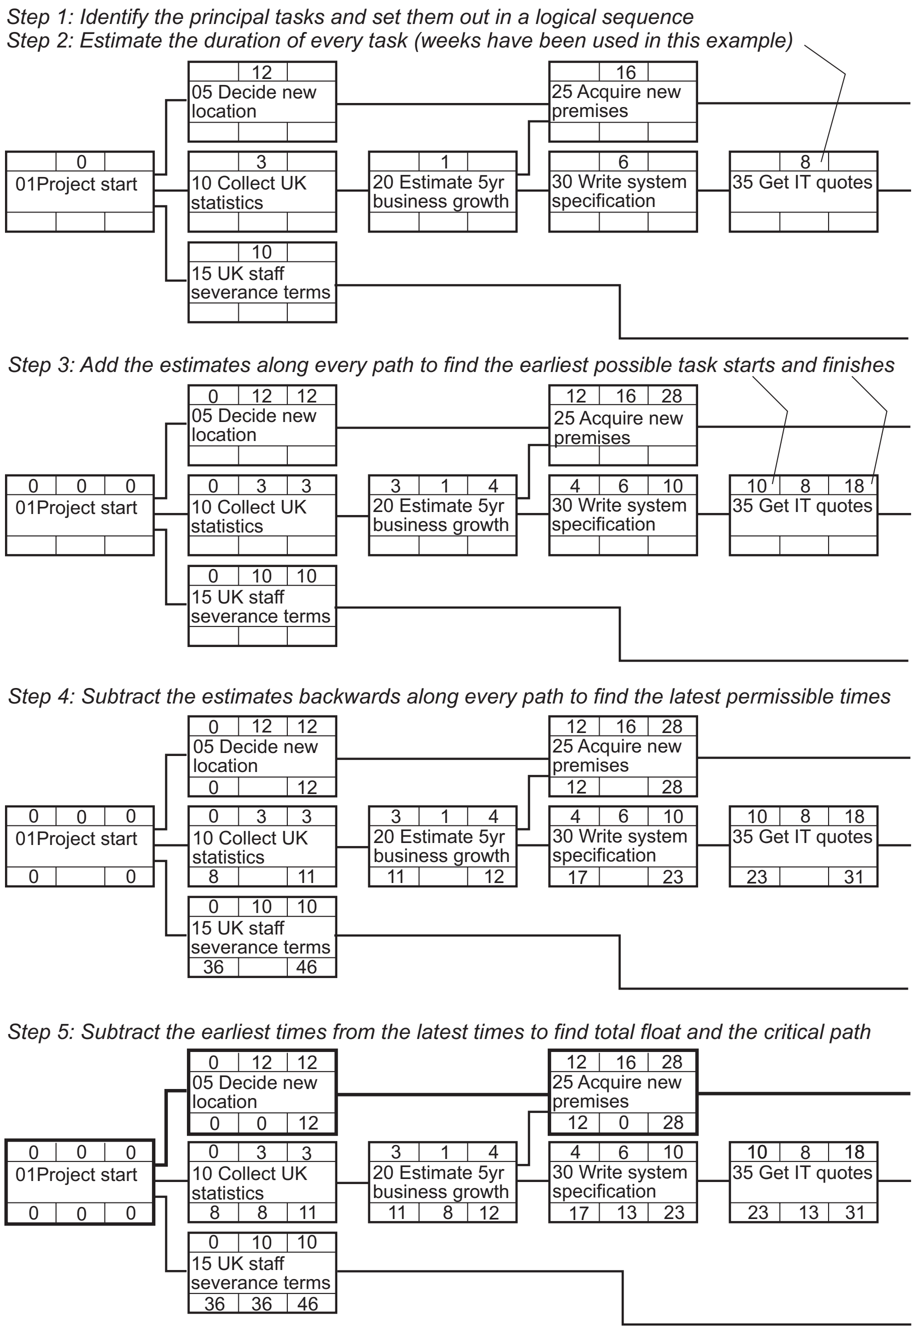

Figure 4.2 Five principal steps in conventional network analysis (precedence system)

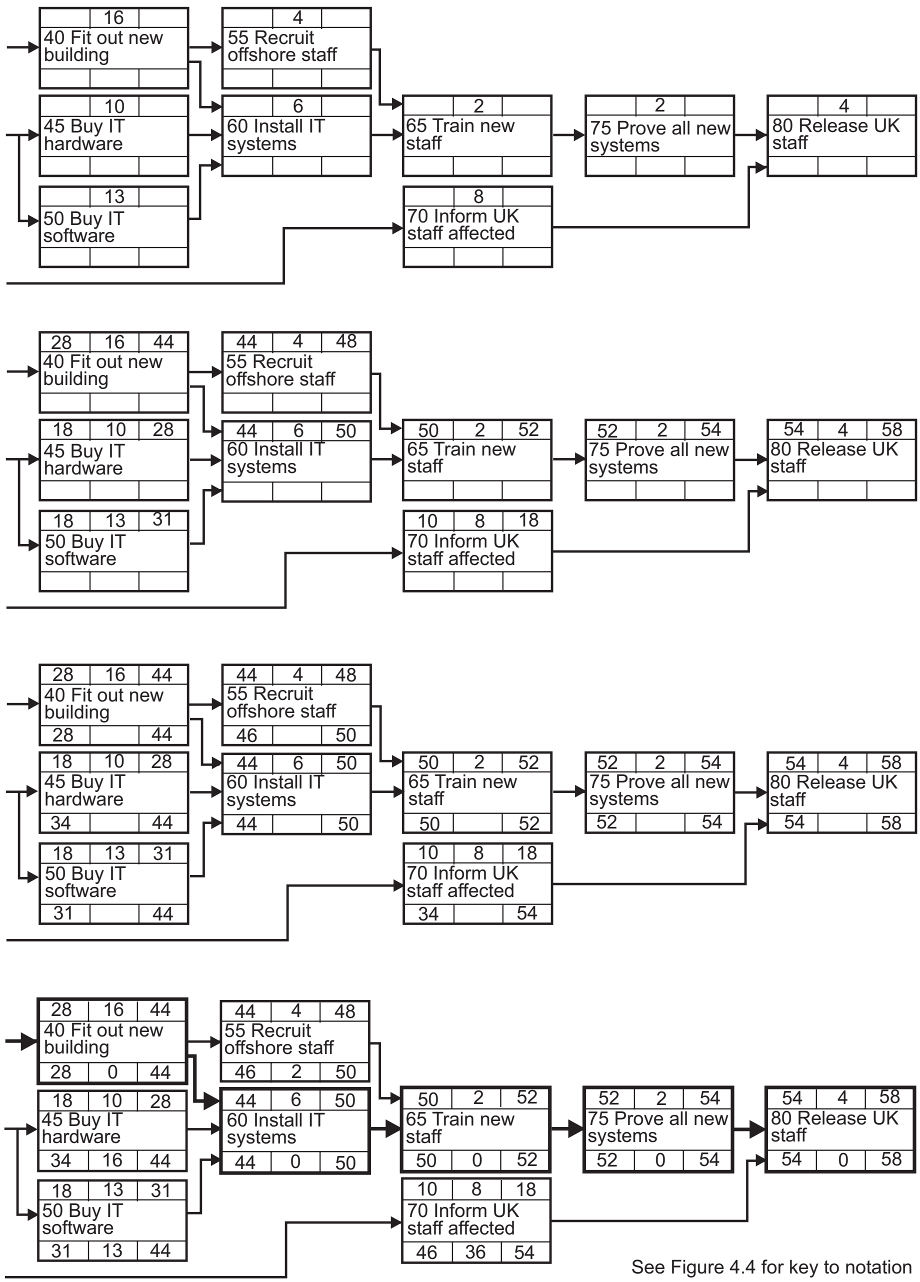

See Figure 4.4 for key to notation

However, this is not the time for despair. The simplest solution is to ask 'Is there another, smarter way to do this?' We regularly follow this practice in our private lives. When obstructed, we simply tackle the job in a different way – and often finish it sooner. There are at least two other well-known methods available to the project manager for accelerating project delivery. One of these is always known as crashing, which simply means throwing more resources at tasks on the critical path or finding alternative ways of accomplishing those tasks. Crashing can be very effective in projects such as construction, where the increased cost of the crashed tasks is more than balanced by the cost saving implicit in finishing the project on time. However, Isochron recommends avoiding that approach because, in a business change or IT project, it simply leads to more heads, more costs, more management overheads and more delay (see Chapter 6). The other well-known method is to take the plan by the scruff of its neck and re-examine it in a process known as fast-tracking.

ACCELERATING PROJECT DELIVERY USING FAST-TRACKING AND CONCURRENCY

Fast-tracking is a process well known to the managers of research and development projects where the time-to-market for a new product has to be as short as possible. If there are techniques that can reduce the total duration of such research and development projects, there is no reason why the same methods should not be used to accelerate business change projects and every reason why they should.

Traditionally, those who plan projects tend to string out the tasks end-to-end. Taken to its perfect limits, fast-tracking plans to do as many tasks as possible concurrently. Concurrency is part of Isochron's methodology. In fact, the company name is based on the term 'isochron', which means a line in space joining places of the same or equal time. The application of this concept to project planning can be illustrated by the very simple analogy in Figure 4.3.

Consider the process of taking shopping through a supermarket checkout. The upper diagram in Figure 4.3 equates to traditional, sequential project planning. Mr Jones joins the queue and, when he can, starts to transfer his goods from the shopping trolley on to the conveyor belt. Mr Jones waits until all the preceding shoppers have paid, loaded their goods into bags and departed. When Mr Jones's turn comes, he watches the checkout operator pass all his goods through the barcode reader. When, and only when he is asked to pay, Mr Jones reaches for his wallet, extracts a credit card, pays, puts the card and till receipts in his wallet, replaces the wallet in his pocket, loads his goods into carrier bags and leaves the checkout.

The next shopper in the queue is Mrs Smith. Her actions on reaching the checkout are depicted in the lower diagram of Figure 4.3. As soon as she can, Mrs Smith loads her goods on to the conveyor belt. She then finds her purse and estimates roughly how much her goods will cost. As soon as Mr Jones departs, Mrs Smith takes her trolley through the checkout to the far side of the till and starts

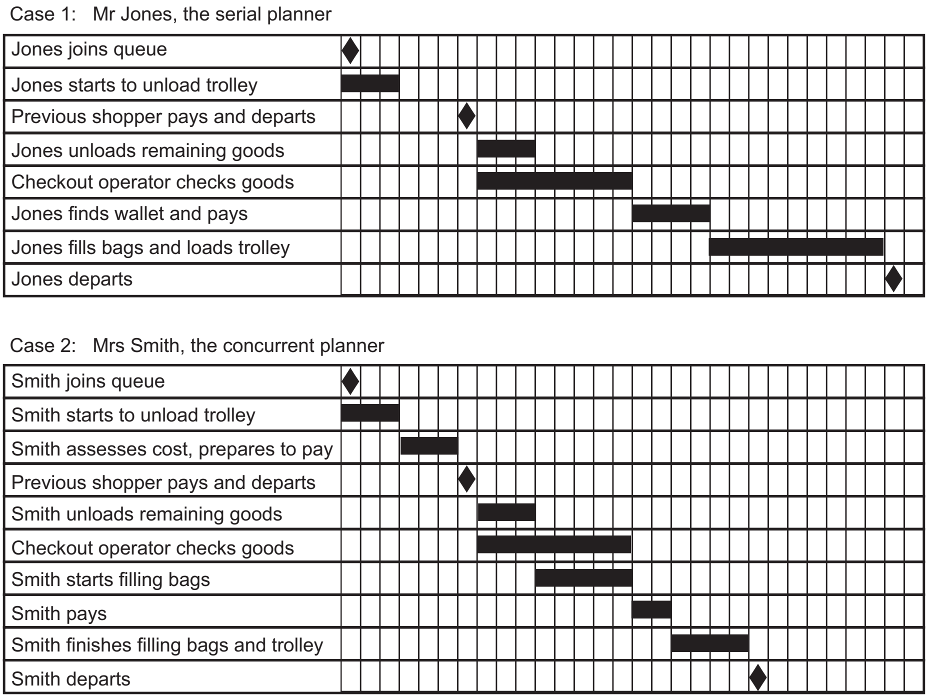

Figure 4.3 Serial and concurrent planning at a supermarket checkout

to prepare bags for packing. As the goods are checked through the barcode reader, she loads them into bags and puts them in her trolley. When the time comes to pay, Mrs Smith already has her purse ready and pays, either with her credit card and ready pen or PIN, or with the right amount of cash. After paying, she replaces the purse in her handbag and leaves. She has taken significantly less time than Mr Jones.

Mrs Smith has adopted the isochronal approach. Fixed in her mind is the outcome – to be walking away with the goods packed and paid for. She has planned her mini 'project' back from the point of leaving the checkout and is already thinking ahead to the chores to be done when she gets home (the outcomes of her next 'project'). She has, perhaps unconsciously, asked herself the following questions:

- What are the key things to be achieved in this checkout process?
- What tasks can be done simultaneously?
- What are the unavoidable dependencies?

Mrs Smith does not, of course, do this explicitly. Her outcome-based approach to planning this tiny project has become routine and instinctive.

Note that the resource levels are the same for both Mr Jones and Mrs Smith, comprising the shopper, the checkout operator, the shopping bags and the checkout equipment. Thus project compression and acceleration do not inevitably need duplication of resources. The values in this simple example may seem trivial, amounting to a few minutes of Mrs Smith's time saved and earlier clearance for subsequent shoppers. Similar planning and efficiencies can seem more important if you are delayed in a check-in or security clearance queue at an airport. In a project or programme of projects they can translate into savings of many months and millions of pounds.

So why do we plan projects any differently? The answer is not that we are, in some way, stupid. Rather it is that, although our *unconscious* thoughts are concurrent in the vast networks of our brain, our *conscious* attention is always sequential and single-channel. Pinch your arm and bite your lip at the same time and feel the sensation of pain flip between the two. Try as hard as you can, it is difficult to feel both pains with equal severity at the same time. 'Sensory gating deficit' prevents us from being overwhelmed with concurrent information from all our senses, memories and thoughts at the same time. It also makes us plan sequentially and causes us to add (Isochron finds) about one-third of redundant sequence and time into our project plans.

We can shed redundant sequence, shorten elapsed time, optimize concurrency of work streams and minimize interdependency by mechanically and systematically reorganizing the tasks in our project. Critical path analysis is a valuable tool for this process. The remainder of this chapter examines how the critical path method can be used to express accelerated plans.

USING CRITICAL PATH ANALYSIS TO EXPRESS ACCELERATED PLANS

With backcasting, skills management and fast-tracking available in the project planner's armoury, dramatic reductions can be made in the planned duration of most projects. Real project networks typically contain more tasks than can be shown with clarity on a book page (or on a computer screen for that matter) but, even in the absence of a large canvas, the principles can be demonstrated using just a tiny fragment of a network diagram, as shown in Figure 4.4.

Finding the minimum path using partial concurrency

The top diagram in Figure 4.4 shows two activities that lie on the critical path network of a first draft network for an imaginary project. An equivalent Gantt chart has been added to give a sense of timescale that networks cannot usually provide. This network plan has been drawn using the very common approach that shows rigid adherence to task interdependencies, with no attempt to accelerate project delivery by any of the methods described above. What this diagram shows is that

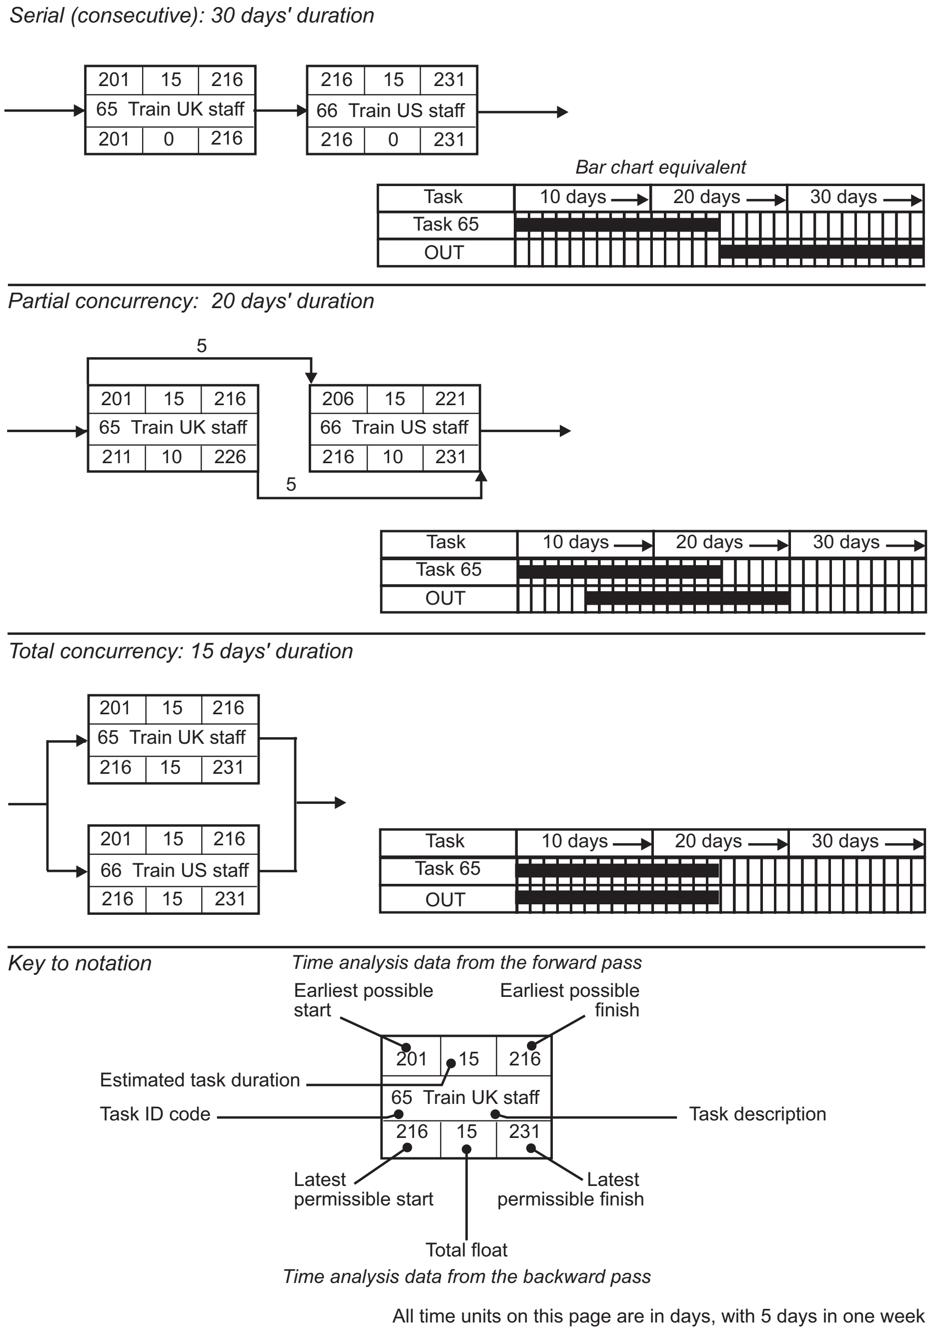

Figure 4.4 Using precedence network notation to express concurrency

Task 65 must be completely finished before Task 66 can even start. The planner, in this case, has assumed that all training of UK staff must be finished before training in America can begin.

A common mistake in planning is to confuse resource constraints with work logic constraints. In this case it might be supposed that the planner has envisaged a training team working for three weeks in the UK, and then flying out to America over a weekend to repeat the process in the American branch of the company for a further three weeks. But why not employ a separate planning team in America and have them carry out the task at the same time? If we were to put that suggestion to the planner, he or she might argue that the second planning team would need a little training themselves, which could be provided by allowing the new team to work alongside the existing team for five days (one calendar week) to absorb the training method. This solution is assumed in the plan headed 'partial concurrency' in Figure 4.4. The simple finish–start constraint between the two training tasks has been replaced by a start-start precedence link which means that the American training can start five days (one week) after the start of training in the UK.

The five-day duration finish–finish link is necessary to complete the path for Task 65 and avoid leaving it with no successor tasks (in which case it would become something that planners call an end dangle). Dangling activities cause broken paths that prevent full time analysis. Devaux (1999) is particularly good at explaining the use of complex links in precedence diagrams.

This partial concurrency solution has reduced the total time needed for these two training tasks from 30 days to 20 days. These tasks were originally critical (with zero float) which means that, for the same project finish date, ten days' total float has been introduced into this path. Or, alternatively, the project could be finished ten days earlier (in which case the training tasks would again become critical).

Finding the minimum level of interdependency to achieve complete concurrency

If the project manager is asked to plan for the absolute minimum staff training time in both the UK and American companies, he or she must consider how to remove the interdependency between Tasks 65 and 66 altogether. This might be achieved by preparing two training teams in advance, so that both can be ready to start work on the same day and finish at the same time with the UK and American operations perfectly concurrent. Figure 4.4 shows that planning for total concurrency could reduce the total training time to just 15 days. So, by planning to eliminate the interdependency of these two critical tasks, the project close (and realization of benefits) has been advanced by three weeks.

A true case example of concurrent planning

The following is based on a true case, but dates and names have been changed to preserve anonymity.

A district council received a directive from the relevant ombudsman to implement a system that would make the allocation of local authority housing to the large number of applicants more impartial and objective. In July 2006 the council's small IT team planned the project and found that, if work were allowed to start in October, the new allocation system should be up and running 12 months later (in September 2007). However, ignoring this prediction, a prominent councillor (let's call him Councillor Sharp) not only announced the project in the local press, but also promised that the new system would be ready to accept applications by November 2006 and actually make the first allocations in January 2007.

The team were thereby presented with a shock requirement to complete their 12-month project in only half that time. An early project start was the first obvious step, but that would still only bring the system on line in July 2007. So, a fast-tracking approach was investigated, in which task concurrency could be maximized.

The team's first step was to list all the known tasks in the project plan and, temporarily ignoring resource constraints, challenge each task in turn to determine whether there was any sound reason why it should not be started immediately (or, at least, tomorrow). After challenging all the tasks in this way, a few were identified that could be indeed be started tomorrow. All the other tasks were dependent in some way on this first small group of tasks.

The team then re-examined all the remaining tasks to ask why each of them should not start immediately after completion of the first group. Again, some could, but the rest would be dependent on this second group. This process was repeated until the state was reached at which the last set of tasks could be started. It was found that this method caused the tasks to fall into a small number of parallel work streams, each largely independent of the others but with the tasks tightly linked by dependencies within their streams.

This revised plan initially suggested project completion by October 2006, but the addition of resource constraints indicated a practicable project completion date of January 2007. In the event the project ran just one month late and finished in February 2007, far in advance of the September completion planned without concurrent working.

CONCLUSION

Two themes run through this chapter. One is treating the time at which achievement of outcomes is required as non-negotiable. This is considered equal in importance with all the other qualities expected or the outcomes. The other theme is using recognition events and backcast planning to produce the most effective and practicable plan, together with alternative and contingency plans. The most effective plan can be made using the following steps:

1. Identify the recognition events from the sponsor and senior management.
2. Ensure that the recognition events and value flashpoints have dates attached.

3. Plan backwards from the value flashpoints to identify high-level key milestones.
4. Plan backwards from the key milestones to identify the technical and administrative processes.
5. Plan for maximum concurrency and minimum interdependency.
6. Express the plan in a critical path network using the precedence system.

REFERENCES AND FURTHER READING

Cross, M. (2002), 'Why Government IT Projects Go Wrong', *Computing*, 11 September, accessible at http://www.computing.co.uk/computing/features/2072199/why-government-projects-wrong.

Devaux, S.A. (1999), *Total Project Control: A Manager's Guide to Integrated Planning, Measuring and Tracking*, New York: Wiley.

Duffy, J. (2002), 'Cleared for a Very Late Take-off', *BBC News Online* at http://news.bbc.co.uk/1/hi/uk/1781780.stm.

Gordon, J. and Lockyer, K. (2005), *Project Management and Project Planning*, 7th edn, London: Financial Times/Prentice Hall.

Lock, D. (2003), *Project Management*, 8th edn, Aldershot: Gower.

Tulving, E. (1985), 'How Many Memory Systems Are There?', *American Psychologist*, 40, pp. 385–98.