concern is resource scarcity. If some resources (including time) were not scarce, the resource allocation problem would be concerned solely with profit maximization—a relatively easy problem.

In Chapter 5, we evaluated project durations solely in terms of time. A project was either on time or not. Now we must also consider when and for what purposes scarce people, equipment, material, and facilities are used. The PM's performance is judged by the skill with which the trade-offs of time, resources, and performance are managed, so the PM must make constant use of cost/benefit analysis. There are countless questions to be answered. "If we come in late on this project, we face a $1,000 per day penalty. How much project slack do we need and what resources at what costs are required to get it?" "If I hire Cheatem Engineering Associates as design consultants, can I improve project performance by 3 percent without extending the project's due date?" "Adding project slack and hiring a consultant require monetary resources that could be used for other things. Are these the best uses for the dollars?"

At times, the PM is asked to take on a project in which there are the usual time, budget, and performance goals, but which also constrain the trade-offs that the PM may wish to make if required to help the project meet its most important goals. For example, some projects are *time constrained* and must be completed by a fixed time. In such cases, resources (and possibly performance) are variable. Some projects are *resource constrained* and cannot go over budget or use more than a fixed amount of a specific resource. In these cases, time (and possibly performance) is variable. Occasionally, a senior manager suffers from a case of the micromanagement virus and fixes time, cost, and performance, thereby leaving the PM with no flexibility whatsoever. Such projects are certain to fail unless the micromanager has been profligate with the firm's resources, which is highly unlikely for micromanagers. The fault actually lies with the PM who accepts command of such a project. (For those who are thinking that such a PM may find him- or herself without a job following a refusal of an assignment, we would note the senior manager in question is insuring that the PM will fail. Do you want to work for someone who will not allow you to succeed?)

We will start our tour through the wilds of resource allocation by reconsidering the problem of dealing with a pointy-haired boss who insists that a project be completed in much less time than its expected duration.

6.1 EXPEDITING A PROJECT

The unreasonable boss problem in Chapter 5, Section 5.2 could be used as our example here, but a smaller problem will help avoid unnecessary arithmetic. Our problem is set in a deterministic world rather than in a probabilistic one, for the same reason. (Please remember that in reality all projects are carried out under conditions of uncertainty.) Finally, we must also take note of an assumption usually adopted when activities are scheduled, as we did in Chapter 5. That assumption is that all estimates of task duration, whether deterministic or probabilistic, are based on normal or standard resource loadings.

The Critical Path Method

In traditional PERT/CPM, the rules of "standard practice" apply and the *normal* task duration estimate is made with the normal or standard-practice resource usage. Then a second estimate, referred to as the *crash* duration, is made based on the resources required to expedite the task. More resources of the type already used might be added;

more workers and shovels if there is a ditch to be dug. On the other hand, the technology used to dig the ditch might be totally altered, utilizing a backhoe or a Ditch Witch®, for example. When making estimates for crashing, it is important to make sure that the resources required to crash the project are, in fact, available. Using a machine to dig the ditch in three hours instead of the three days required for a worker with a shovel is dependent on the fact that the machine is available and can be on site when needed. (Of course, the warning about resource availability applies equally to normal resource requirements as well as to crash requirements.) There are times when the PM may expedite activities that have little or no impact on the network's critical time, such as when the resources used must be made available to another project. It is important to remember that when we change technology, we may also be changing the level of risk in carrying out the activity. Finally, we must remind ourselves that some tasks *cannot* be crashed. One must not assume that because it takes one woman nine months to carry and bear a child that nine women can accomplish the same result in one month.

Consider the project described in Table 6-1. There is a set of activities, predecessors, normal task duration estimates, crash duration estimates, and estimates for normal cost and crash cost. One crash duration is marked with a single asterisk. For this activity, the task may be carried out in normal time or crashed one day at a time. Another activity is marked with a double asterisk. In this case, the duration must be one or the other; it cannot be broken down to one-day segments. Activities are charged at the "cost per day" (*activity slope*) increments shown in the last column. A given activity may have only two or three technically feasible durations. If an activity cannot be split into one-day segments, the cost is as indicated. The "slope" information for non-or-partially segmented activities is normally given in the slope chart. Activity slope is computed as follows:

$$\text{slope} = \frac{\text{crash cost} - \text{normal cost}}{\text{crash time} - \text{normal time}}$$

When crashing a project, starting with the normal schedule for all project activities, crash selected activities, one at a time, to decrease project duration at the minimum additional cost. To crash a project, follow two simple principles: First, focus on the critical path(s) when trying to shorten the duration of a project. Crashing a noncritical activity will not influence project duration. Second, when shortening a project's duration, select the least expensive way to do it.

Given these guides, consider the network shown in Figure 6-1(a) that was constructed from the data in Table 6-1. It is easier to illustrate the impact of crashing on an activity-on-arrow (AOA) network than on an activity-on-node (AON) network, so we use that approach here. Also, we use dummy activities in this case not to illustrate precedence but to show time durations and slack on the time axis.

Table 6-1 An Example of a Normal/Crash Project

Activity	*Precedence*	*Duration, Days (norm/crash)*	*Cost (norm/crash)*	*Slope ($/day)*
a	—	3, 2	$ 40, 80	40/−1= −40
b	a	2, 1	20, 80	60/−1= −60
c	a	2, 2	20, 20	—
d*	a	4, 1	30, 120	90/−3= −30
e**	b	3, 1	10, 80	−70 (2 days)

*Partial crashing allowed

**Partial crashing not allowed

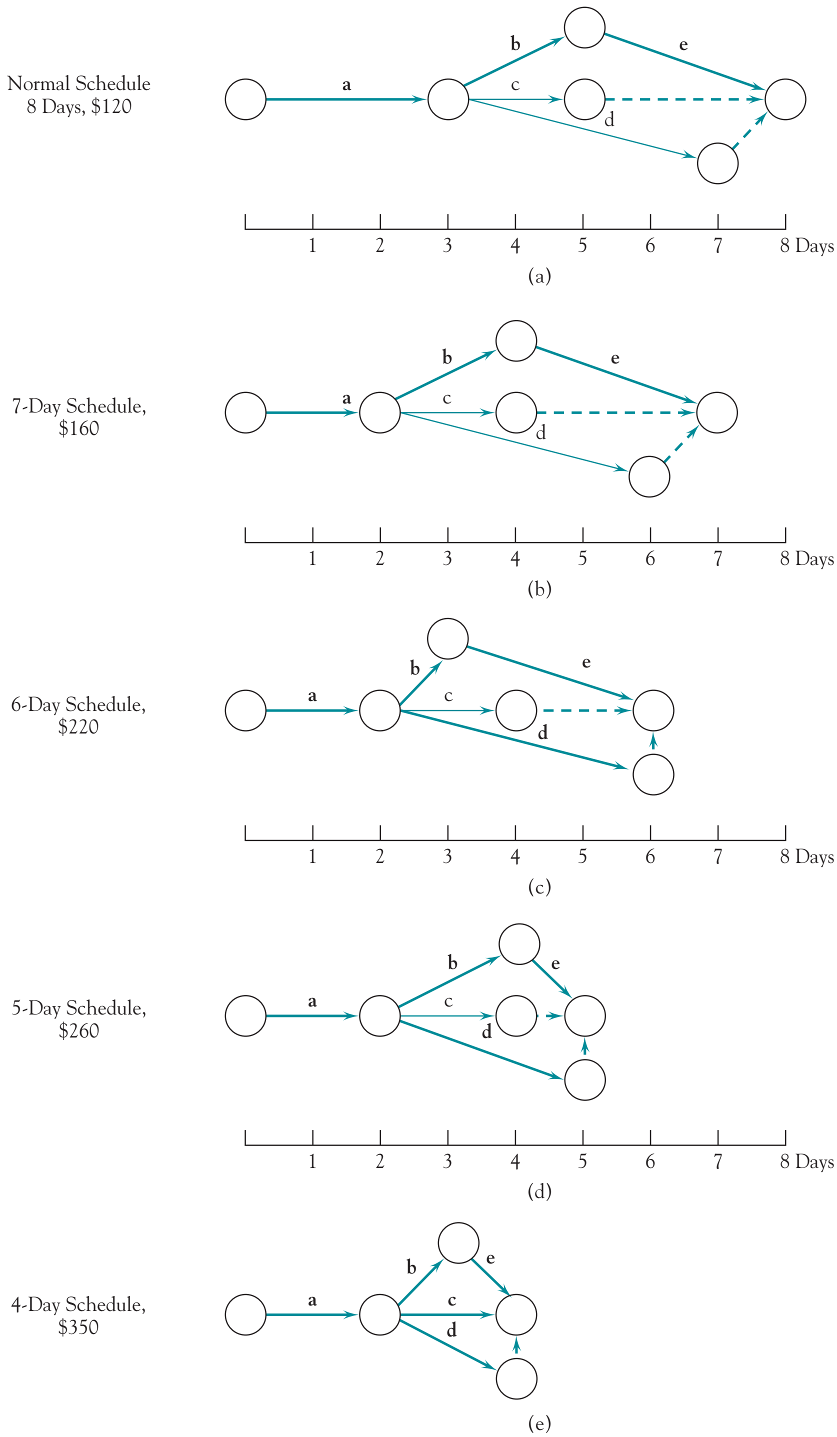

Figure 6-1 A PERT/CPM example of crashing a project, AOA network.

As indicated in Table 6-1, activity **d** can be partially crashed for $30, but it is not on the critical path and will not shorten the project. Activity **e** involves a technological discontinuity and must take either three days to complete at $10 or one day at $80. In general, the impact of having such a technological discontinuity is that the best solution for crashing n days might not be part of the best solution for crashing $n+1$ days. Rather, it may be best to crash the activity with the technological discontinuity at $n + 1$ days and not crash another activity that could be crashed for n days. This situation is illustrated in the discussion that follows.

The network's critical path is **a-b-e**, the project duration is 8 days, and the normal total cost is $120, as illustrated in the network of Figure 6-1(a). The decision about which activities to crash depends on how much we need to reduce the duration of the project. To reduce the total network duration by 1 day, we must reduce the time required by one of the activities along the critical path. Inspecting Table 6-1 to see which critical activity can be reduced at the least cost, we find it is activity **a**, which adds $40 to the project's current cost of $120. Activity **b** could be crashed at an added cost of $60 or we could even crash **e** 2 days for an additional cost of $70. Of course, crashing **e** would only shorten the project duration by one day because when **e** is shortened, the path **a-d-dummy**, seven days long, becomes the critical path and does not allow the project to be shortened to 6 days. Of the three options, crashing **a** is the lowest cost and therefore preferable, see Figure 6-1(b). Notice that crashing **a** also shortens **a-d-dummy** and **a-c-dummy** by 1 day.

Suppose the project must be crashed by 2 days. What are the options? Reconsidering Table 6-1 and Figure 6-1(a), we see that we could crash activity **e** for 2 days ($70), but path **a-d-dummy** (7-days' duration) must also be crashed at least 1 day. We choose **d** ($30/day) because it is cheaper than **a** ($40). The total cost of crashing is $100, and the total project cost is $120 + $100 = $220. Alternatively, we could crash **a** and **b**, also for a cost of $100 ($40 + $60). Arbitrarily, we choose the latter option [Figure 6-1(c)].

Now suppose we wanted to crash the project by 3 days, from the original 8 days down to 5 days. Clearly **e** must be crashed by 2 days, costing $70, and **a** or **b** by a day. We choose **a**, the cheapest, for an additional $40. This leaves **d** to be crashed by 1 day for another $30, resulting in a total crashing cost of $140 and a project cost of $120 + $140 = $260 [Figure 6-1(d)]. Note that we did not crash **b** this time, as we did for 6 days. This is due to the technological discontinuity in activity **e**.

Last, let us consider crashing the project by 4 days down to a project duration of four days. Since we crashed **e**, the technological discontinuity, to reach a 5-day duration, all the remaining activities can be incrementally crashed. Thus, we can simply inspect Figure 9-1(d) to see what else needs incremental crashing to reduce the project by another day. Notice in Figure 6-1(d) that **a-b-e** and **a-d-dummy** are both critical paths. Only **b** and **d** can still be crashed so we crash each by 1 day for an additional cost beyond the 5-day schedule of Figure 6-1(d) of $60 + $30 = $90 for a total project cost of $260 + $90 = $350 [Figure 6-1(e)]. Note that **c** is now critical, therefore *all* paths are critical. Since the critical paths **a-b-e** and **a-c** are at their full extent of crashing, the project duration cannot be further reduced, even though activity **d** could be crashed another day. Thus, Figure 6-1(e) is *not* the all-crash network, although it equals the all-crash time schedule of four days.

Whether or not all this crashing is worthwhile is another matter. On the cost side, Figure 6-2 shows the time/cost relationship of crashing the project. On the benefit side, some projects have penalty clauses that make the parent organization liable for late delivery—and sometimes bonuses for early delivery. Starting at the right (all-normal) side of Figure 6-2, note that it becomes increasingly costly to squeeze additional time out of the project. Charts such as the one shown in Figure 6-2 are useful to the PM in exercising control over project duration and cost. They are particularly helpful in dealing with senior managers who may argue for early project completion dates with little understanding of the costs involved. Similarly, such data are of great benefit when clients plead for early delivery. If the client is willing to pay the cost of crashing, or if

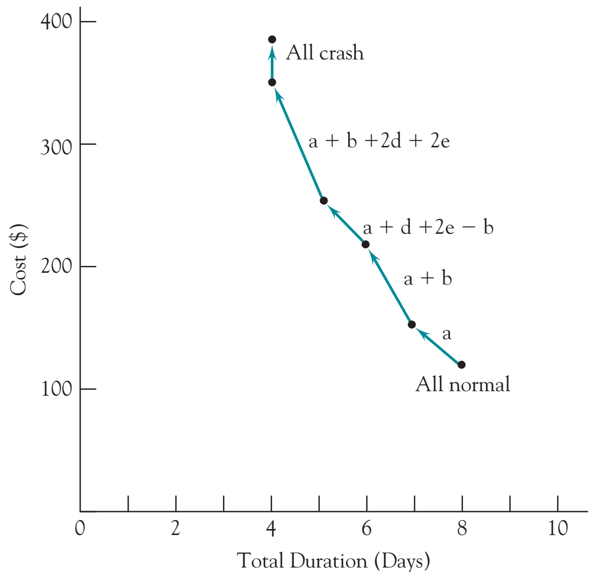

Figure 6-2 CPM crash cost-duration history.

the firm is willing to subsidize the client, the PM can afford to listen with a sympathetic ear. (While we advise the PM to ignore overhead costs over which he or she has no control, it should be noted that indirect costs are often altered when a project is crashed.)

One final note on crashing projects. The same method is used when the task durations are probabilistic, that is, using three time estimates. In this case, optimistic, most likely, and pessimistic activity duration estimates are made for the "normal" resource loading and new optimistic, most likely, and pessimistic duration estimates must be made for crash resource loading. The PM should remember that the variance of both the normal and crash activity times largely depends on the technology used to accomplish the activity in question. Thus the variance of the normal activity time may be quite different from the variance of the crash time. The project budget can be determined in exactly the same way. The solution to project duration and resource cost levels can be reached by using the standard analytical method used in the last chapter, or by simulation, also described in Chapter 5.

Fast-Tracking a Project

In addition to crashing a project in order to expedite it, a project may also be *fast-tracked*. Used primarily in the construction industry, the term refers to an expediting technique in which the design and planning phases of a project are not actually completed before the building phase is started. Usually design and plan are finished before the building is started, so letting them overlap reduces project duration—if the fact that design and planning are incomplete does not result in a significant amount of rework and change orders during the building phase.

For many projects in construction, maintenance, and similar areas, a large proportion of the work is routine. In these cases, fast-tracking rarely causes serious problems. The number of change orders in fast-tracked construction projects is not significantly different from that for similar projects that were not fast-tracked (Kurtulus and Narula, 1985).