CORPS Capabili

TOC and critical chain methodologies aren't just for manufacturing anymore

ties

BY **MANDYAM SRINIVASAN, DARREN JONES**, AND **ALEX MILLER**

The theory of constraints (TOC) is a management philosophy based on the premise that every organization can be viewed as a system, and every system has a weakest link. Typically, just one aspect of the system, the constraint, limits the organization's ability to achieve its full potential, and so the manner in which the constraint is managed largely determines the organization's throughput. Since its introduction some 25 years ago, TOC has developed rapidly, both in terms of methodology and of where it has found application. One of the more recent developments is the critical chain methodology for managing large projects.

TOC, in particular the critical chain methodology, is finding increasing application within the military, especially in its maintenance, repair, and overhaul (MRO) operations. MRO is a multibillion dollar industry that presents significant opportunities for cost savings resulting from better management practices. For instance, the MRO market for commercial aircraft alone was $37.8 billion in 2001. The Center for Executive Education (CEE) at the University of Tennessee has developed and launched a unique program, entitled Lean MRO, which combines concepts from lean thinking and TOC to train participants in better MRO management practices. As a part of its development efforts, CEE surveyed best practices in the application of lean thinking and TOC across MRO organizations. The Maintenance Center at the Marine Corps Logistics Base, in Albany, Georgia, represents one such successful application.

About four years ago, the Maintenance Center was struggling to complete equipment repairs on time and was faced with an increasing backlog of work. In the center's heavy equipment repair and overhaul lines, asking for "plus-ups," or additional time to complete the work, had become a normal way of doing business. For instance, with the overhaul of the MK-48, a heavy-duty hauler for the Marine Corps, the center was only producing 5 units per month against a demand of 10 per month. Customers were threatening to divert their orders to the private sector.

In an effort to better match scheduling to the realities of the work, the management team contracted with Vector Strategies to implement a critical chain pilot project on the MK-48 vehicle. The pilot project proved very successful and the center began implementation of the critical chain plantwide, generating dramatic performance improvement.

Extension of TOC

The critical chain methodology addresses some of the shortcomings of PERT (program evaluation and review technique), the tool most widely used for project management. PERT is based on identifying the critical path, which is the longest chain of linked events (task dependencies) embedded in the overall project. Focusing only on the longest chain of task dependencies can result in several problems, such as multitasking. The critical chain methodology instead asserts that, in addition to task dependencies, good project management should better address resource constraint and visual management and practice multitasking only when absolutely necessary. The methodology provides a means of determining where time buffers should be placed to prevent unplanned disruptions from delaying the project completion.

Let's look at a simple example that illustrates the critical chain methodology. Figure 1 (next page) represents a project with six activities, each represented by a letter followed by a number. The letter denotes the resource performing the activity, and the number is an estimate of the activity time, appropriately buffered (inflated) to reflect uncertainty in task times. Thus, the box representing the synchronization activity, Z24, indicates that this activity is carried out by resource Z and takes 24 days.

Applying the PERT methodology, the critical path for this project, denoted by the shaded rectangles, takes a total of 66 days. With PERT, each activity in the project is provided with a set of four times, the earliest start time, the earliest finish time, the latest start time, and the latest finish time. These times are broadcast to everyone involved in the project, so they can be closely monitored. The difference between the earliest and latest start time is the slack. Activities on the critical path do not have any slack time and are given the most attention in PERT.

The critical chain methodology questions the manner in which PERT determines and broadcasts the activity times. It suggests that, while one needs to build in some safety buffers to account for uncertainty, buffering *each* activity separately and, furthermore, *broadcasting* these buffered activity times promotes what is termed the "student syndrome," wherein activities are put off until their due date draws near.

The methodology instead suggests that the analyst use historical data to obtain an estimate of an aggressive but

Figure 1 Example project indicating the critical path

Figure 2 Aggressive but possible times and the critical chain

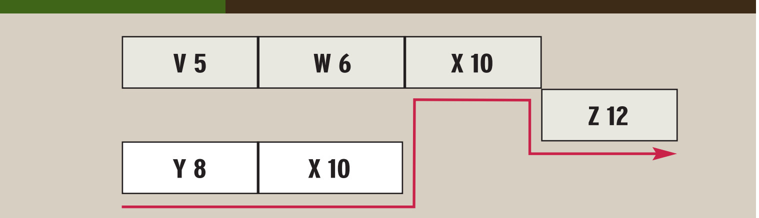

Figure 3 The critical chain with project and feeding buffers inserted

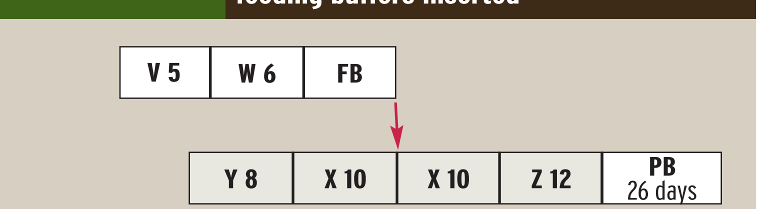

Figure 4 Monitoring the project buffer

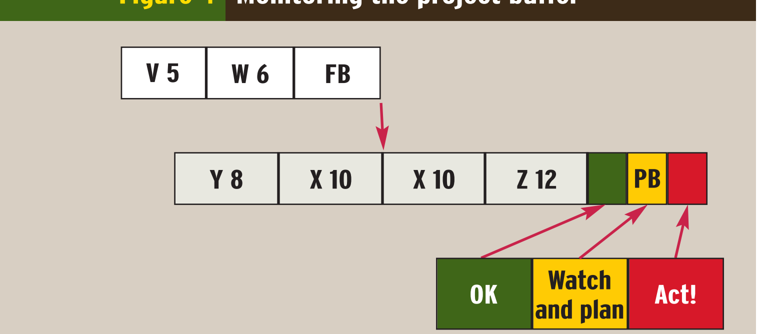

possible (ABP) time for each activity and use this for the activity duration instead. For our example, assume for simplicity that the ABP time is exactly half of the buffered time for each activity. Figure 2 shows the same project except that the task durations are ABP times. This figure also shows the critical chain, indicated by the line with an arrow. Note that the critical chain is quite different from the critical path presented in figure 1. The critical chain explicitly considers the fact that the same resource, X, is used in multiple activities, one of which was not on the critical path earlier. As a result, activity Y8 now becomes a resource that requires close monitoring. Activities V5 and W6, which were formerly on the critical path, are not on the critical chain. The activities on the critical chain are monitored the most closely. Activities not on the critical chain can slip a little without affecting the overall project completion time.

Types of buffers

The series of activities on the critical chain, Y8-X10-X10-Z12, results in an aggressive project completion time of 40 days. The next step is to determine the buffer to cover uncertainty in task times. This buffer is determined by the time released by adopting ABP times, namely, 66 – 40 = 26 days. This is the project buffer, as shown in figure 3, and it buffers against any variation in the completion times of activities times along the critical chain. The estimate for the project duration (the lead time) is the sum of the average activity times for the critical activities, plus the safety time determined by the project buffer.

Figure 3 also shows a feeding buffer that is placed after activity W6 to protect the critical chain from any slippage. Such a feeding buffer is inserted before an activity that feeds into the critical chain but is itself not on the critical chain. The feeding buffers are determined in a manner similar to the way the project buffer is calculated. It must be noted that the project buffer and the feeding buffers are time buffers and not inventory buffers. That is, variation is buffered by capacity, rather than inventory.

Finally, the critical chain is monitored by closely following the rate at which the project buffer is consumed, as denoted by figure 4. Thus, instead of broadcasting due dates for each activity, thereby promoting the student syndrome, the critical chain methodology suggests simply informing the project team on a regular basis whether the rate at which the project buffer is being consumed is under control or not.

Critical chain at the maintenance center

The maintenance center in Albany, Georgia, overhauls and repairs vehicles used by the Marine Corps—fuel tankers, amphibious vehicles, light armored vehicles, earthmoving equipment, trucks, and so on. The overhaul process at the center starts with disassembling each vehicle to determine its work scope—the amount and nature of the work to be done on that product. The work scope also indicates which parts can be repaired and which parts need to be replaced. Parts that require repair are routed through a series of support shops that include cleaning, blasting, painting,

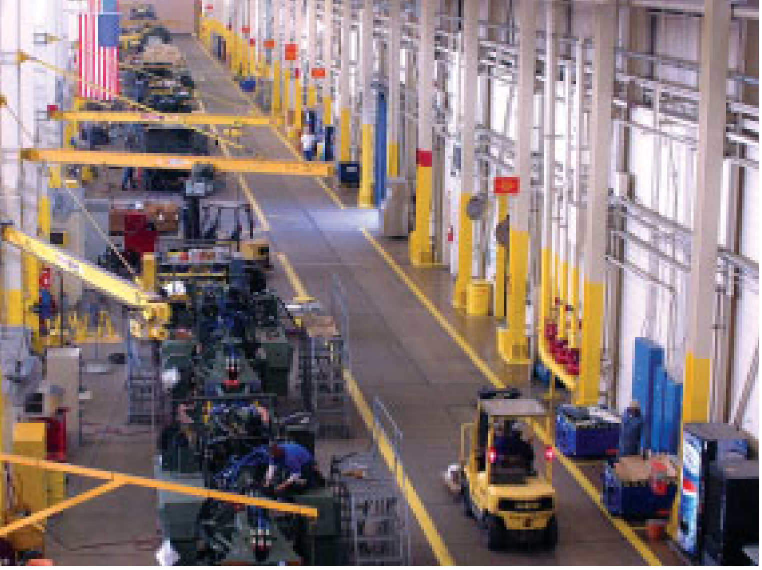

machining, body work, weapons work, and so on. Parts that need replacement are either replaced from existing spare-parts stock or ordered from an external source. To add to the complexity, the original manufacturer no longer may produce the parts that have to be replaced.

Unlike a typical flow shop manufacturing setting, where the enterprise knows the sequence of operations required to complete the finished product, the MRO facility is very much like a pure job shop facility. In the MRO facility, the work scope of a product that arrives at the facility is not known unless the product is disassembled and inspected. There is a tremendous variation in the work scope even for the same type of product, such as the MK-48, and it is difficult to accurately predict the percentage of parts that must be replaced and the percentage of parts that should be repaired. This variation creates problems in using a push system to schedule work.

At the time the pilot project began, scheduling was based on a manufacturing resource planning (MRP II) push system. It was a push system in that products were introduced into the shops without regard to the status of the resources dedicated to the repair activities. Given the high levels of variation in work content across jobs, this led to false starts and delays, increased inventories, and lowered throughput.

Push scheduling created additional noise in the systems by flooding the back shops with jobs that competed for the same resources (equipment and personnel). Disassembled parts were immediately pushed out to the support shops to allow as much time as possible for them to move through the repair cycle. As jobs competed for resources, equipment and personnel often were moved between projects before they completed all the work on a given project. As a result of additional move time and increased set-up times, each of the projects tends to take more time than if they were completed one at a time from start to finish. These were all problems the center's management team hoped to address through their use of TOC and the critical chain methodology.

Finding the real bottleneck

As a first step toward applying TOC, the center's management sought input from throughout the organization on where bottlenecks were believed to be seriously limiting output. Opinions varied as to what were and were not bottleneck activities, but every major activity in the center was believed to be an important bottleneck by at least someone in the facility. In other words, managers believed that every major activity in the center served to limit production.

In applying TOC to address the center's problems, the main shop where the main products were first disassembled and subsequently reassembled, and the support shops (cleaning, repair, and so on), were modeled as the critical chain. The critical chain analysis of the data collected revealed that, contrary to everyone's opinion, the facility had more than enough capacity to carry out the activities required to meet the demand for repair and overhaul of 10 MK-48s per month. In other words, none of the major activities were limiting production.

In fact, the root cause of the consistent shortfalls and high inventory levels was the scheduling system in place that was pushing products out to the shop floor without regard for the status of the resources. The bottleneck was thus not a physical resource constraint. Rather, it was a policy constraint introduced by the scheduling process. This discovery enabled Vector Strategies to use a simplified drum-buffer-rope (SDBR) technique to model and schedule the activities in the shops that processed components removed from the main products.

To add to the COMPLEXITY, the original manufacturer may no longer produce the PARTS that have to be replaced.

TOC employs the drum-buffer-rope (DBR) system to manage production. The traditional DBR model releases orders into the production process such that it synchronizes with the production rate of the least capable resource in the process. This least capable resource is referred to as the capacity constraint resource (CCR). If the CCR works at a rate that is less than the rate of output demanded by the customer, then it is the bottleneck. (Otherwise, the external demand rate, the market, is the bottleneck.)

In the standard DBR model, the production rate of the CCR is the drum, which paces production for the system. The rope in DBR is the mechanism used to release work into the production process. It is a communication device to ensure that raw material is not introduced into the shop at a rate faster than the CCR can handle. If the CCR is not the bottleneck, the rope ensures that raw material is not introduced into the shop floor faster than the customer demand rate. Finally, to prevent the CCR from ever having to wait for work if it becomes free, a time buffer is placed ahead of the CCR to ensure that jobs arrive at the CCR well before they are scheduled for processing at the CCR. Another buffer, called the shipping buffer, protects the situation where the customer's order might be delayed. The standard DBR model is shown in figure 5.

The standard DBR model requires specialized DBR software to implement. For enterprises that already have common material requirements planning systems in place, the alternate technique mentioned earlier, SDBR, can be used when the enterprise is not constrained by any internal resource (the situation at the center as revealed by the initial critical chain analysis.) The drum in SDBR is based on firm orders. As orders come in, a quick check is made on the total load on the CCR. If the CCR is not too heavily loaded, the order is accepted and released into the shop floor for processing.

The only buffer maintained is the shipping buffer. The rope is no longer tied to the CCR schedule. Instead, the material release schedule is directly generated by firm orders received. See figure 6.

The SDBR model has the advantage of not requiring any specialized software. This is a significant benefit for enterprises that might be unwilling or unable to invest in specialized DBR software. Another advantage of the SDBR approach is that it does not have to require two buffers, but needs just one.

Finally, the SDBR approach is more focused on market demand and ties the organization to its customers more directly. The center was able to use an SDBR approach to scheduling in conjunction with the existing MRP II business system. Only the critical chain portion of its solution required additional software, which was Realization Technologies' Concerto. The MRP II system that was used for scheduling now facilitates the SDBR schedules.

Results and rollout

The center's results were impressive. Repair cycle times for the MK-48 were reduced by a factor of 3, from an average of 167 days to an average of 58 days. Work in process levels were reduced from 550 percent of demand to 140 percent of demand. The cost to repair products also went down by 25–30 percent, mainly because the reduction in delays resulted in more throughput without any increase in the cost of repair. The capacity for the MK-48 line is much more flexible and can work with a rate of 10 units per month to as high as 23.

The work carried out to date has made the Albany Maintenance Center a showcase of world-class overhaul and repair. Weekly tours are conducted, hosting officers and executives from government and the private sector.

Based on the success of the MK-48 pilot project, the center expanded the application of TOC and critical chain to additional lines. The second major application focused on a landing assault vehicle, the LAV-25, where cycle times were reduced from 212 days to 119 days. Since then the management team has continued to use TOC to focus its efforts and streamline its repair processes.

Mandyam Srinivasan is The Ball Corporation Distinguished Professor of Business at the University of Tennessee's College of Business Administration. He may be reached at (865) 974-1737, or via e-mail at msrini@utk.edu. Darren Jones is a manager in the production management department at the Marine Corps Maintenance Center in Albany, Georgia. Alex Miller is the William B. Stokely Professor of Management and associate dean for executive education at the University of Tennessee's College of Business Administration.

Figure 5 Traditional drum-buffer-rope model

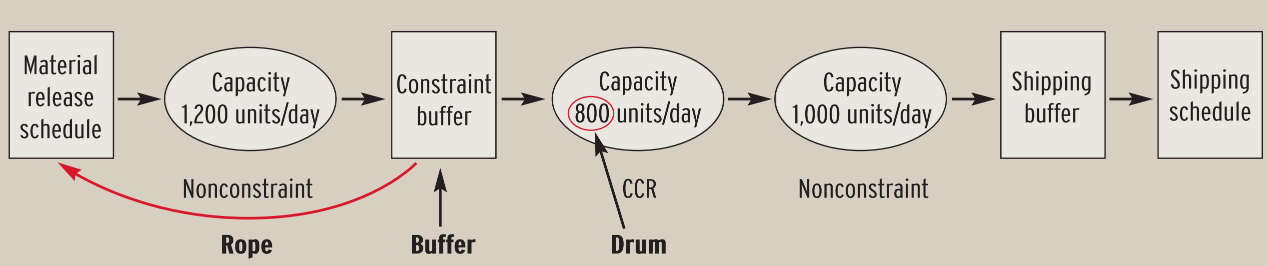

Figure 6 Simplified drum-buffer-rope model