Interfaces
Vol. 37, No. 1, January–February 2007, pp. 7–21
ISSN 0092-2102 | EISSN 1526-551X | 07 | 3701 | 0007

informs®
DOI 10.1287/inte.1060.0260

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Warner Robins Air Logistics Center Streamlines Aircraft Repair and Overhaul

Mandyam M. Srinivasan
Stokely Management Center, College of Business Administration, University of Tennessee, Knoxville, Tennessee 37996-0562, msrini@utk.edu

William D. Best
402D Aircraft Maintenance Group, Warner Robins Air Logistics Center, Warner Robins, Georgia 31098, william.best@robins.af.mil

Sridhar Chandrasekaran
Strategic Services, Realization Technologies, San Jose, California 95113, sridhar@realization.com

At Warner Robins Air Logistics Center, long lead times for repairing and overhauling aircraft were a serious issue because some models were in short supply. In 2005, the center implemented Critical Chain, an operations research method for project management, to reduce lead time for repairing its C-5 transport aircraft. The implementation took eight months, without using any additional resources. The center returned five additional C-5 aircraft to the US Air Force's operational inventory and generated additional revenue to the Transportation Working Capital Fund estimated at $49.8 million annually. The replacement value for these aircraft is estimated at $2.37 billion. The center is implementing the method for the C-130 and the C-17 cargo aircraft to free up 11 dock spaces. With the additional dock space, the depot can accommodate additional work worth $65 million in 2006. The center expects the additional workload to total $248 million by 2009 (the current C-5 annual operating budget is $295 million). Nonquantifiable benefits include increased responsiveness and casualty avoidance during wartime.

Key words: facilities–equipment planning: maintenance, replacement; project management.

Annual spending on maintenance, repair, and overhaul (MRO) for aircraft is now about $96 billion worldwide, exceeding the $75 billion estimated to be spent on producing new aircraft annually (Stewart 2005). Spending in the military aircraft MRO market alone is estimated to be about $50 billion, with the US accounting for about 40 percent of that amount. To manage the MRO activity effectively, Warner Robins Air Logistics Center (WR-ALC) applies operations research tools for project management and scheduling.

WR-ALC provides MRO services for various weapon systems for its customer, the US Air Force. These systems include transport aircraft, fighter jets, and ground systems that provide maintenance support to these aircraft. WR-ALC provides MRO services to four types of aircraft: the C-5 Galaxy, the C-17 Globemaster, the C-130 Hercules transport aircraft, and the F-15 Eagle fighter jet.

The MRO activity consists of (1) program depot maintenance (PDM), a heavy repair and overhaul of aircraft with long lead times, and (2) unscheduled depot level maintenance (UDLM), a short-cycle maintenance that takes from one to 30 days or more. The US Air Force sends approximately equal numbers of C-5 aircraft to WR-ALC for PDM and UDLM each year. We focus on the C-5 aircraft that require PDM.

The C-5 line has 24 frontline supervisors and about 460 mechanics, organized into skill groups, for example, sheet-metal mechanics, hydraulic mechanics, electricians, and nondestructive inspection (NDI) mechanics. These mechanics are shared across all PDM and UDLM aircraft. Each aircraft typically requires these skill sets multiple times. The C-5 line operates two shifts.

Challenges

When Kelly Air Force Base closed down following the recommendation made by the 1995 Base Realignment and Closure Commission, WR-ALC had to compete in a public/private solicitation for the repair and overhaul of the C-5 aircraft. WR-ALC won the bid in 1997. However, the Air Force did not transfer the level of personnel and equipment from Kelly Air Force Base that WR-ALC expected. It had to retrain its existing mechanics to repair and overhaul C-5s, aging aircraft that Lockheed Martin last produced in 1989. WR-ALC faced numerous challenges.

The first big challenge stems from the inherent complexity in aircraft repair and overhaul. There are three levels of complexity: (1) known work that must be performed on any aircraft that comes in for repair and overhaul, but even this known work can vary, (2) anticipated work that involves work to be performed on some of the aircraft, although it is hard to predict which aircraft need this work, (3) unanticipated work that is due to damage and is quite unpredictable, which may require parts built from scratch or new repair techniques. A typical PDM requires 40,000 to 50,000 man-hours of work and may require an additional 10,000 hours of unanticipated work (Figure 1).

The next challenge stems from the limited resources, such as manpower, facilities, parts, and tools. At any time, several aircraft may compete for these resources. Problems in one aircraft can have a cascading effect on other aircraft when they share resources. For example, repairing a major fuel leak in one aircraft may hold up repair of another aircraft because fuel pits or the skilled mechanics needed for leak checks are unavailable. Such competition for resources seriously affects scheduling and execution.

Several elements are needed for any task: materials, mechanics, work-control documents, tools, and support services. Synchronizing these elements is the third big challenge. To ensure that all the needed elements are ready when the mechanic starts a task, several departments have to work together. In an uncertain environment, with constantly changing priorities, the departments may find it very difficult to agree on the priorities. In addition, every integration point, where multiple tasks or activities must be complete before the next task or activity begins, poses a challenge. For example, in the buildup phase, during which the aircraft is reassembled after individual repairs are completed, over 100 different activities must be completed, each a potential delay.

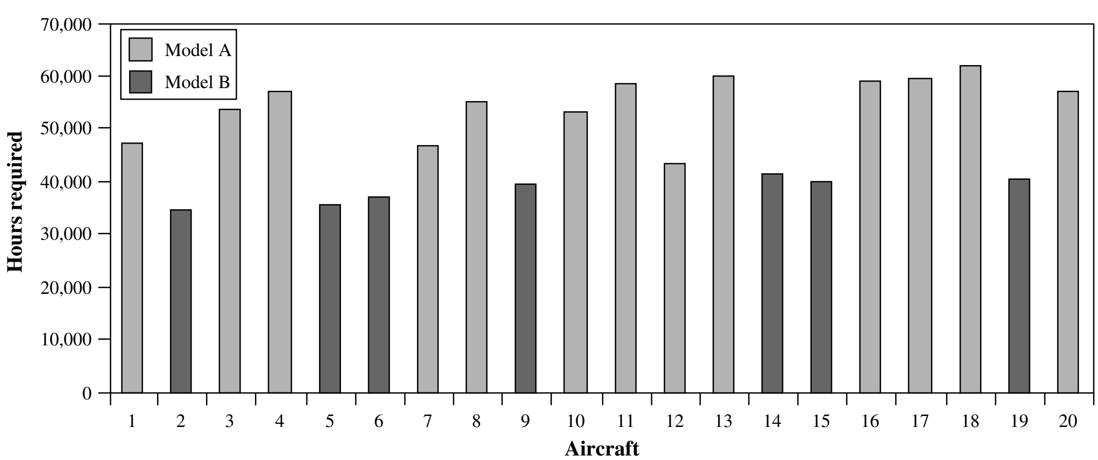

Figure 1: The hours needed to repair and overhaul 20 C-5 aircraft in fiscal year 2005 show large variation.

Shared support organizations, such as the back shops, the NDI department, and engineering, can also affect the schedule. The support organizations are typically shared across weapon systems, such as the C-5, the C-130, the C-17, and the F-15. Competing for and coordinating with these support organizations is the fourth challenge.

WR-ALC had many reasons for taking a lot of time to repair and overhaul the C-5s; during fiscal year (FY) 2000, the C-5 turnaround time was 360 days. Almost all C-5s were delivered late at costs greatly exceeding estimates.

Lean Initiatives

Between FY 2000 and FY 2005, WR-ALC undertook initiatives to reduce turnaround time. It established cell teams, developed standard work for all its cells, deployed production-control boards, applied six-sigma techniques and method sheets (clear job instructions), and instituted parts kits and point-of-use material. It developed a mechanism to pull components from the back shops and worked on reducing travel times within the depot. These efforts greatly improved performance. While part of the improvement came from the learning curve as WR-ALC became proficient at repairing and overhauling the C-5, most of the performance gains came from implementing lean principles to manage repair and overhaul. In FY 2000, it repaired and overhauled 17 aircraft with a turnaround time of 360 days, with about 16 aircraft in the depot at any point in time. In FY 2005, it repaired and overhauled 20 aircraft with an average turnaround time of 240 days and only 12 or 13 aircraft in the depot at any time.

The lean initiatives also improved due-date performance. The center bases depot schedules for aircraft in the fleet on a predetermined maintenance cycle, determining the due date by adding the contract flow days to the input date (the day the aircraft arrives at the depot). Once the aircraft arrives at the depot, the center can adjust the due date if it needs additional work, although that creates customer dissatisfaction. Whereas the center was delivering only 25 percent of the aircraft on time in FY 2002, it delivered all aircraft on time in FY 2004 and FY 2005. The C-5 line won the Shingo gold award for its improvements in September 2005, the first time it was offered to the public sector. WR-ALC was the only public sector organization to win the gold award in 2005.

Continuing Performance Improvements

Despite these achievements, WR-ALC was still under pressure to reduce turnaround time. According to Little's law (Little 1961), a reduction in turnaround time produces a corresponding reduction in the number of aircraft in the depot. Reducing the number of aircraft undergoing repair and overhaul was urgent, especially because only 112 C-5s exist, and those undergoing repair are not in operation. During FY 2004 and FY 2005, WR-ALC always had 11 to 13 aircraft in the depot in various stages of repair. Thus, 10 percent of the fleet were not being used for their intended purpose, carrying cargo and passengers for the military.

WR-ALC was also under pressure to further improve its performance for a number of other reasons. First, the US Air Force was urging it to take on additional maintenance work (UDLMs) that it could not do because of an apparent lack of capacity. Second, the US Air Force wanted to increase the availability of all weapon systems by 20 percent by 2011. This initiative posed a challenge particularly for aging aircraft like the C-5. The pressure to further enhance performance came from the users of the aircraft and from the US Air Force at the corporate level:

The users of the aircraft wanted WR-ALC

—To reduce the number of aircraft in the depot,

—To deliver all aircraft on time with no schedule extensions, and

—To take on more UDLM work without adding capacity.

The US Air Force at the corporate level wanted WR-ALC

—To improve its support for wars,

—To reduce the costs of PDM, and

—To create the capacity for additional work.

WR-ALC realized that reducing the number of aircraft undergoing repair and overhaul in the depot would provide additional benefits:

—It would reduce the competition for facilities, mechanics, and other resources.

—It would allow supervisors and team leads to focus on fewer jets at a time.

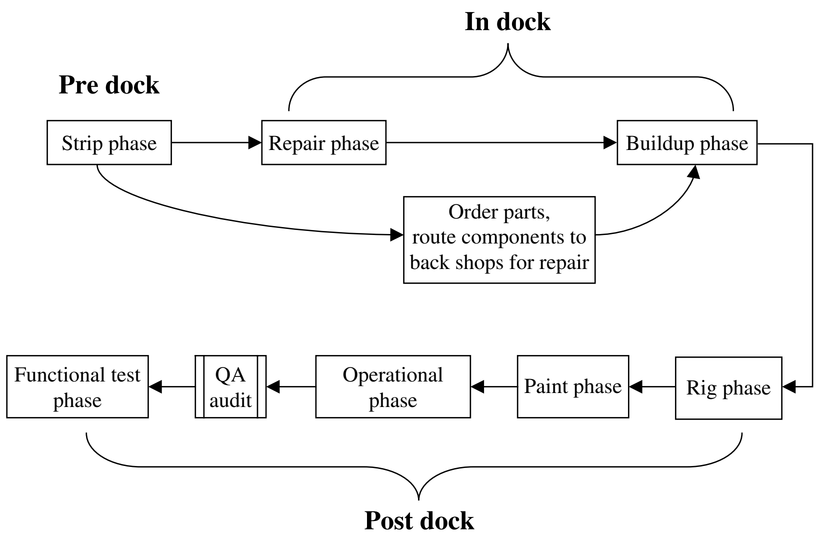

Figure 2: Workers remove arms and fuel from the aircraft and remove and inspect major components in the strip phase. They order parts and route major components to the back shops for repair. In the dock, they perform repair activities on the aircraft to the extent possible while awaiting these parts and components during the repair phase. As parts become available, they continue to reassemble the aircraft (buildup phase). They reconnect systems and manually operate and check them in the rig phase. They scuff and paint aircraft and check for weight and balance (paint phase). They power the aircraft and perform operational tests (operational phase). The aircraft is subject to a quality assurance (QA) audit. Pilots perform flight tests, and mechanics prepare the aircraft and deliver it to the customer (the functional test phase).

—It would give supply, engineering, and other support functions fewer aircraft to support at one time.

—It would allow back shops and material support to focus on fewer parts at any time.

—It would improve management attention, resulting in faster decisions.

—It would reduce paperwork and other administrative tasks.

In sum, reducing the number of aircraft in the depot would speed up maintenance, increasing throughput and promoting a virtuous cycle of ongoing improvement. The question was, how could WR-ALC further reduce turnaround time? After all, it had just completed a very successful lean initiative that reduced turnaround time from 360 days to 240 days. The C-5 was already Warner Robins' best performing line in terms of operating costs and results, due date performance, flow days, and product quality.

Uncovering Problems

To further address air force needs, WR-ALC revisited the repair and overhaul operations. At a macro level, the depot performs the following activities on the aircraft (Figures 2 and 3).

At WR-ALC, managers used Gantt charts to plan and track the many tasks making up these activities; the charts showed when tasks should begin and end. In July 2004, University of Tennessee faculty (Srinivasan et al. 2004) introduced Critical Chain project management (CCPM) (Goldratt 1997, Elton and Roe 1998) to WR-ALC managers. CCPM seemed appropriate for managing the repair and overhaul activity at WR-ALC because each aircraft could be viewed as a project, with a series of tasks and precedence requirements to be completed.

In December 2004, WR-ALC decided to adopt CCPM to schedule repair and overhaul of C-5s even though the C-5 line was already the best performing line at the center. Furthermore, the US Air Force mandated that all MRO activities on aircraft were to be managed using lean manufacturing principles, and some CCPM principles seemed to be in direct conflict with them. For instance, a fundamental step in lean manufacturing, especially for high-volume, low-variety operations, is to pace all processes to customer

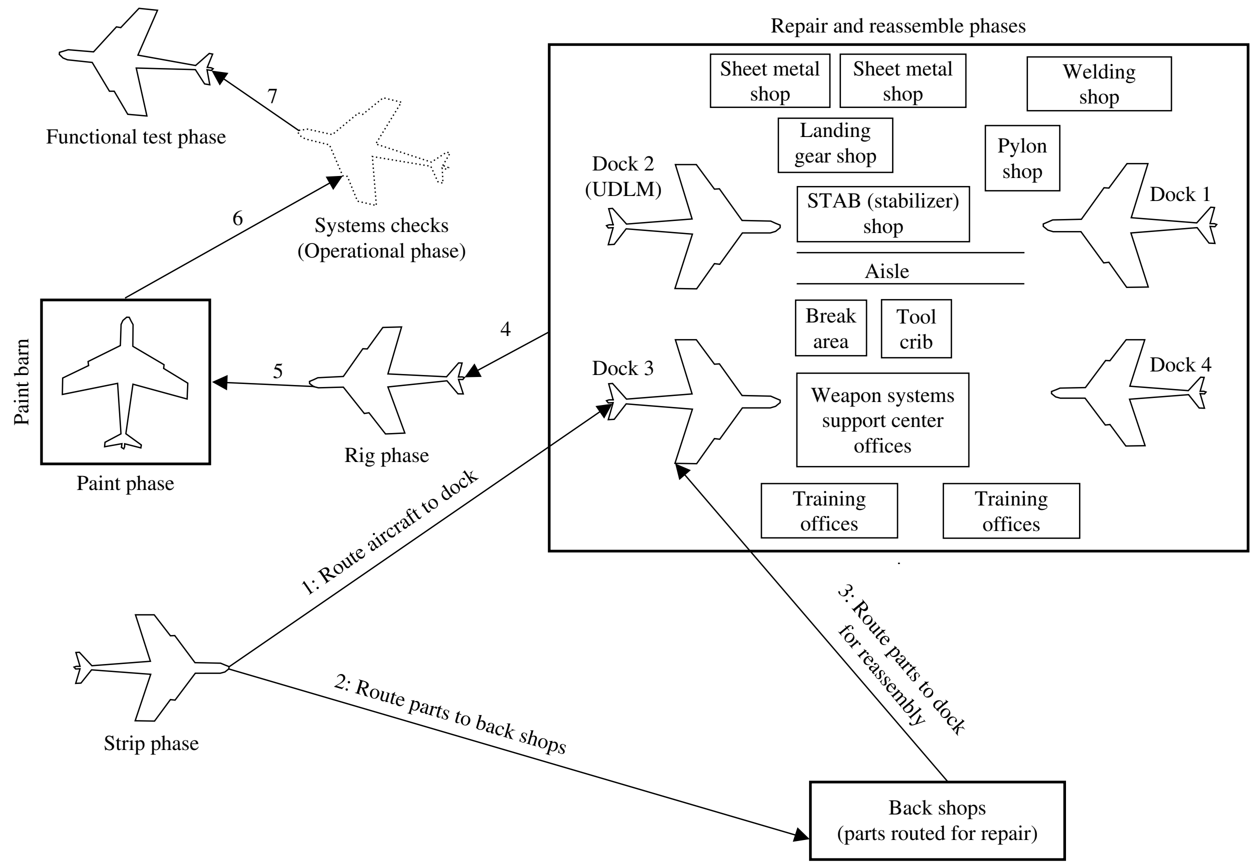

Figure 3: The aircraft moves from one location to another during the repair and overhaul process. Most of the repair and overhaul activity is performed in one of the four docks shown.

demand, which is measured by *takt time*, the time allotted to complete each process step. With repair and overhaul process that vary in complexity and duration, pacing operations to takt time is problematic. When some process steps encounter unexpected or unplanned work to be completed within takt time, quality often suffers. The notion of takt time and the notion of level-loading resources were not readily accepted. We had to overcome huge cultural obstacles to implement CCPM.

WR-ALC decided to adopt CCPM because it realized that the lean initiative had not satisfactorily addressed a number of problems, in particular, the problem of variation. WR-ALC contracted with Realization Technologies to implement CCPM at the center using Realization Technologies' Concerto software. A University of Tennessee faculty member provided guidance and oversight. The three organizations conducted a systematic study of the existing mode of operation at WR-ALC, which revealed numerous problems:

—WR-ALC had buffered task times to allow for the uncertain task durations. Because people saw these safety buffers as waste, they hid them. Everyone knew that task times were inflated but would never admit it. Managers evaluated project durations by simply adding up these buffered task times along the longest sequence of dependent activities. Unfortunately, the projects would consume the hidden safeties in accordance with Parkinson's law (work expands to fill the time available for its completion). Gutierrez and Kouvelis (1991) discuss how Parkinson's law affects project management.

—Managers monitored projects based on milestones that indicated completion of sets of tasks.

Workers tended to procrastinate on tasks until their due dates drew near, exhibiting the student syndrome (Goldratt 1997): they would complete work by the milestone date—or later depending on the delays—but never earlier. Instead of addressing the student syndrome head-on, managers typically deployed more resources on tasks that were falling behind planned completion times.

—Whenever projects conflicted, managers would give priority to the aircraft closest to its finish date (the lead aircraft). (This practice is akin to the shortest-processing-time-first (SPTF) policy.) While this policy may appear reasonable, it had negative side effects. For one, the customer for the aircraft closest to its finish date may not require delivery until much later. In addition, working on that aircraft would invariably delay work on the remaining aircraft, causing aircraft to miss their due dates. Finally this practice often led to mechanics and other resources multitasking across various aircraft, starting new tasks without finishing existing tasks and thus increasing the number of unfinished tasks.

—To limit delays, managers released work as soon as possible (ASAP). For example, immediately after the initial stripping and inspection tasks, they would schedule the repair tasks and then the assembly tasks ASAP, even if all parts were not available. Such a push mode (of starting tasks ASAP) only increased competition for limited resources across aircraft, with supervisors forcing their mechanics to multitask. As a result, work-in-progress (WIP) levels were unnecessarily high. The high WIP level increased competition for parts and put unnecessary pressure on the back shops and the supply system. Other support departments strained to provide the documents, parts, and support services to execute all the open tasks. A combination of these factors actually increased delays, which triggered a vicious cycle of starting more tasks ASAP, leading to further delays (Figure 4).

Not admitting or starting work ASAP might seem to encourage the student syndrome; however, the idea is to resist starting tasks until all the resources and parts needed to complete them are available. Once the manager decides to start a task, the intent is to avoid the student syndrome by completing it as quickly as possible.

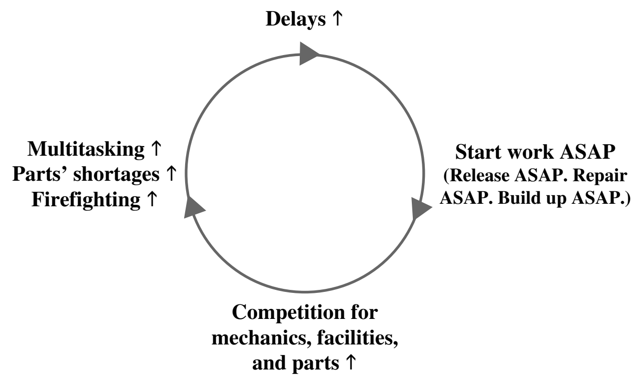

Figure 4: Releasing work into the system as soon as possible results in a vicious cycle of increasing delays.

Our systematic study also revealed that instead of dealing with the uncertainties in the MRO environment during schedule execution, managers were factoring in uncertainties in the project plan. Consequently, they were drawing plans that created waste during execution. Based on the study, we concluded that WR-ALC could not solve the problem with a project-tracking system; it needed a robust execution system that could cope with uncertainties. Critical Chain, an operations research technique for scheduling and managing projects, explicitly recognizes and addresses this problem.

Critical Chain Project Management

CCPM has been much discussed since its introduction in 1997. Some view it as the most important breakthrough in project management since the introduction of PERT/CPM (project evaluation and review technique/critical path method) (Newbold 1998), and others question its innovativeness, arguing that it consists of known concepts presented in a new way (McKay and Morton 1998, Raz et al. 2003).

As Lechler et al. (2005) note, the philosophy behind CCPM is very different from the philosophy behind PERT/CPM, with managers holding a different mindset and following different management practices. CCPM is based on a number of key principles, some of which we outline below:

—Reduce the amount of work in execution largely by releasing new work based on the status of the most loaded resources because they limit the amount of work that can be completed.

—Remove safety buffers from individual tasks and aggregate them into an overall project buffer.

—Do not create precise schedules for resources at planning time; rather, set them during execution based on how much buffer remains. In a multiproject environment, tasks with the lowest buffer ahead of them get the highest priority.

—Avoid multitasking among workers and resources.

Some CCPM principles are counterintuitive. For instance, the principle that work should be released based on the availability of constraining resources may seem counter to the objective of completing work as soon as possible. However, releasing work prematurely into the system overloads already constrained resources and distracts from completing projects in a timely manner.

A distinguishing feature of CCPM is that it explicitly accounts for human behavior. Instead of adding a safety buffer to the time for each task, CCPM prescribes removing the safety buffers from individual tasks and replacing them with a single buffer to protect the overall completion of the project.

In addition, a key tenet of CCPM is that displaying due dates or milestones using, say, Gantt charts, promotes the student syndrome. Instead, CCPM uses visual signals to track the progress of the project, identifying tasks that need attention without providing specific due dates. Managers update the project network regularly. The software highlights tasks that can affect project completion as red tasks so that the supervisor or "team lead" can address problems promptly and set task priorities accordingly.

Project Buffers and Buffer Sizing

Instead of adding safety buffers to individual tasks, CCPM establishes a project buffer (Figure 5). While the concept is straightforward, uncovering hidden safety buffers within the times allotted for each task can be difficult. Goldratt (1997), Leus (2003), and other authors assert that, in the presence of uncertainty, planners typically pad task durations to ensure at least a 90 percent probability of finishing them on time. Goldratt (1997) proposes working with an aggressive time estimate that guarantees finishing tasks on time, 50 percent of the time.

A rough-cut approach to sizing project buffers is to put half the safety time removed from each task

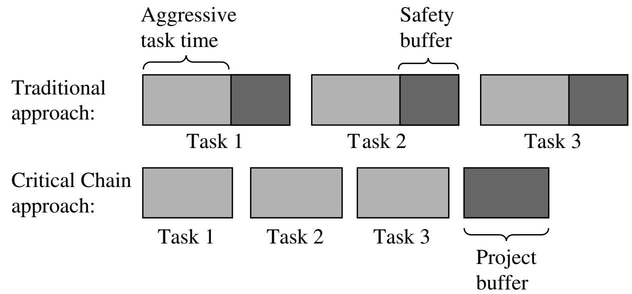

Figure 5: Critical Chain project management removes safety buffers from individual task times to provide a project buffer that protects the project from late completion. The resulting (aggregated) project buffer is smaller than the sum of the individual safety buffers.

into the project buffer. For instance, when all tasks on the critical chain have variance σ^2, padding each task by, say, σ, produces a cumulative safety time of σn with n tasks. The standard deviation of the combined safety times, assuming task independence, is $\sigma\sqrt{n}$. Thus, in this situation the rough-cut approach to sizing the project buffer will provide adequate protection even with $n = 4$ tasks.

A study conducted during lean implementation on the C-5 line indicated the possibility of repairing and overhauling a C-5 in 210 days. Based on this estimate, when we found that the aggressive times for critical chain tasks added up to 105 days, we cut the safety time freed ($210 - 105 = 105$ days) in half and added it back as the project buffer (55 days). We considered several other options (Leus 2003, pp. 145–148) before choosing this simple approach.

Managing Buffers and Setting Priorities

Unlike a traditional project-management technique in which a manager sets milestones for each task and monitors tasks through earliest start and finish times, and latest start and finish times, with CCPM, managers monitor projects and prioritize tasks or projects based on the rate at which buffers are consumed (Figure 6).

Critical Chain Project Management at WR-ALC

We started implementing CCPM by forming a dedicated cross-functional core team consisting of a "team

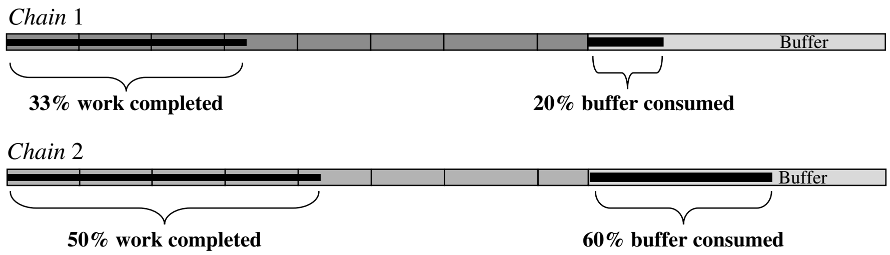

Figure 6: Project 1 (Chain 1) has 33 percent of its work complete and only 20 percent of its buffer consumed. Because the buffer burn rate for this project, defined as the ratio of the percentage of buffer consumed to the percentage of work completed, is less than one, this project is color-coded green by the CCPM software. The second project (Chain 2) is consuming the buffer at a faster rate than the rate of project completion. This project is thus color-coded red by the CCPM software to denote that it should receive a higher priority during execution. CCPM calculates the buffer burn rate on an ongoing basis and reviews it periodically as needed.

lead," a consultant, a representative from Realization Technologies, and a supervisor. The core team participated in an intensive three-day workshop on the concepts underlying CCPM. During the workshop, team members deliberated on how to apply these concepts to the C-5 line. They drew up a project network to model the C-5 repair and overhaul process. The network included the tasks, dependencies between tasks, and the resources assigned to each task. The network included about 450 to 475 tasks. The tasks were identified at a level of detail that would allow supervisors to assign resources to them quite easily, with task durations typically spanning an entire shift.

Subject matter experts, the supervisors, validated the project network to ensure that it reflected the way work was done on the shop floor. The team defined the scope of tasks and handoffs between them clearly, using checklists. The core team listed the anticipated tasks after discussing them with the supervisors, who added these tasks to the network at appropriate places. It was understood that if these anticipated tasks did not materialize during execution, they would be removed from the schedule following initial inspection of the aircraft. It was also understood that unanticipated work would be accommodated as it arose.

With the help of the supervisors, the team challenged and compressed existing time standards for the identified tasks to arrive at aggressive time estimates that eliminated safety buffers. The team identified the critical chain, the longest sequence of tasks after accounting for competition for resources. It identified opportunities for reducing the length of the critical chain by using interchangeable spare parts where possible, so that an aircraft did not have to wait unnecessarily while its own parts were repaired.

The duration of the critical chain, based on these aggressive time estimates, was 105 days. To provide protection against uncertainties, the team added a project buffer to protect the tasks on the critical chain. The network included about 120 feeding chains, each consisting of a set of tasks that were not on the critical chain but fed into it. Tasks on feeding chains were protected by feeding buffers (Patrick 1999). The feeding buffers ensured that the feeding chains did not delay the project.

We deliberated on various options for sizing the project buffer. Ultimately, we adopted a simple buffer-sizing mechanism (Goldratt 1997) to set the project buffer to approximately one-half the duration of the critical chain. With this 50 percent rule, we obtained a 55-day project buffer. We calculated feeding buffers in the same way, using one-half of the duration of the feeding chain.

The team thus arrived at a 160 day target turnaround time, a 33 percent reduction from the existing 240 day turnaround time. We explained to the team that, according to Little's law, the drop in turnaround time to 160-days would reduce the number of aircraft in the depot from 12 or 13 down to seven. The team created project templates to represent the new 160-day flow and loaded them into the Concerto software system. Essentially Concerto automates the processes, providing functionality for planning, execution, and continuous improvement:

—Concerto software helps managers identify the critical chain, calculates and inserts buffers, and provides several what-if capabilities to make trade-offs between overhaul scope and time. The software helps managers to identify constraints and to create a master schedule for starting work on an aircraft.

—Concerto produces reports identifying day-to-day priorities for the entire organization and provides signals warning managers of potential delays. It also highlights areas that need attention to get the schedules back on track.

—Concerto monitors buffer consumption over time and highlights areas whose improvement will most affect performance.

We created a new set of business processes (the production department's daily processes, the support department's daily processes, the C-5 line's weekly buffer review process, and the parts process) for the C-5 line and documented them in detail. We conducted workshops for managers to create consensus regarding CCPM concepts and the new management process. We made presentations to the center's leaders and to support organizations, notably the back shops, and the engineering and functional testing crews, to get their buy-in on the new processes.

We conducted a series of training workshops on CCPM concepts, the business processes, and software use for the front-line supervisors. The C-5 line went live with the new process overnight exactly eight weeks from the start date. During implementation, WR-ALC did not stop accepting new work. Aircraft continued to arrive at the base as originally scheduled. With CCPM, WR-ALC sped up the repair process, mainly by dramatically reducing multitasking, so that existing aircraft were completed at a faster rate than originally planned. The overtime rate remained stable while the line was reducing work in process from 12 or 13 aircraft to seven. After going live, we made several course corrections to incorporate the lessons learned.

Executing Critical Chain Project Management at WR-ALC

The CCPM implementation changed the way the center executed and managed the C-5 repair and overhaul. The center changed the way it released work to production, the way it managed daily work, and the way top management reviewed projects. First, the center increased capacity in the night shift to make it close to that of the day shift by moving some mechanics and support resources to the night shift, increasing throughput at night.

The center reduced multitasking by adopting a novel scheduling approach, *pipelining*. With pipelining, after the mechanics removed components from an aircraft and either sent them for repair to the component shops or placed orders for new components, they ceased working on that aircraft, which was put on hold for approximately 20 days (Figure 7).

Managers initially opposed the introduction of a hold phase because the notion of having an aircraft sit idle, without being worked on, was an anathema. They feared that lack of any activity on an aircraft would increase its turnaround time. However, activity on an aircraft does not necessarily represent progress. Working on an aircraft without all the component parts, a maintenance crew would begin work on an aircraft only to find work interrupted before completion due to a lack of parts. In such a situation, they would idle or the supervisor would try to find another task to keep them busy. The mechanics would start the next task only to find another part was unavailable. The number of unfinished tasks multiplied, and managers complained about the lack of mechanics and shortages of parts.

Thus, instead of keeping the maintenance crews assigned to an aircraft idle for lack of parts while the parts were being repaired or acquired, the supervisors pulled everyone off the aircraft and assigned them to work on another aircraft for which parts were available. By doing that, they foiled Parkinson's law and greatly increased the crew's productivity. They made the mechanics' entire capacity available for fewer aircraft and could allocate enough mechanics to each task rather than spreading them thin. After an initial period during which the center flushed out the WIP, complaints about shortages of mechanics declined.

In general, to inhibit supervisors from releasing work into the system ASAP, management identified several points of control, called *fixer release control* (FRC) points. At these points, only the *fixer* (the highest official responsible for repairing and overhauling

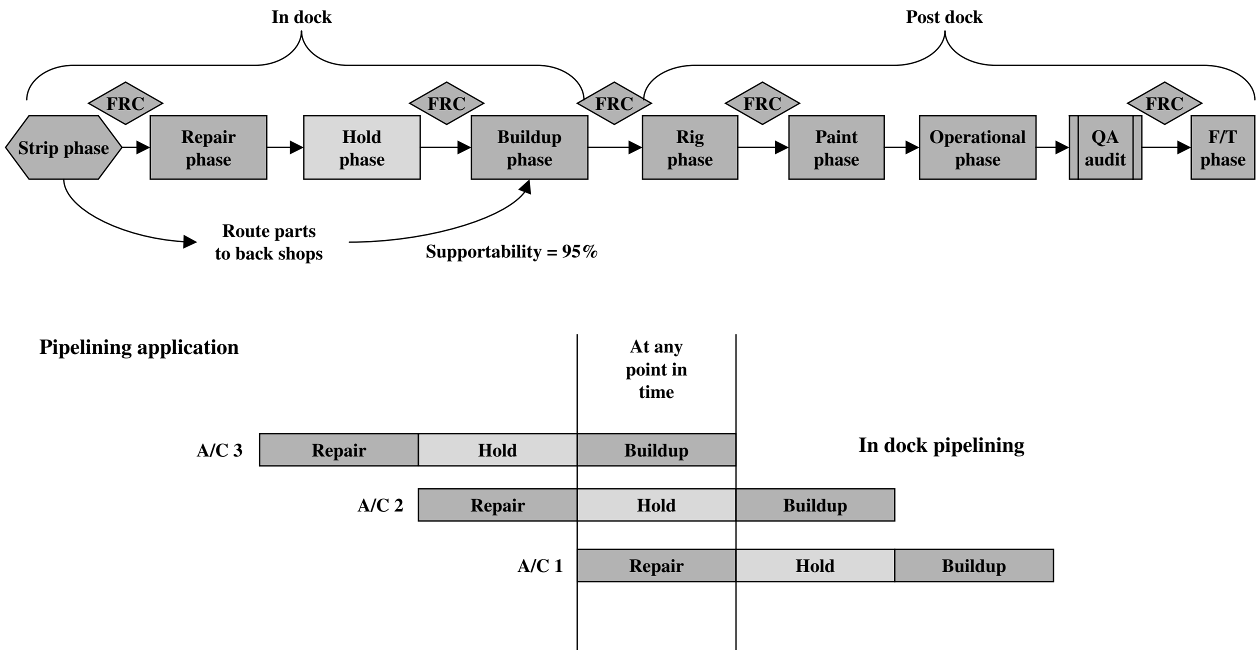

Figure 7: WR-ALC introduced a hold phase, during which period no work is performed on the aircraft. That allows back shops to manufacture or repair parts needed for assembling the aircraft. The aircraft remains in the hold phase until it has at least 95 percent of all required parts. At that time it goes through an aggressive buildup phase. At the fixer release control (FRC) points, only the fixer (the highest official responsible for repairing the aircraft) can permit the aircraft to move forward. The final phase is the functional test (F/T).

the C-5) could permit the aircraft to move forward. At the first FRC point, the fixer decided when to move aircraft into the repair phase. If an aircraft arrived early, the fixer would delay moving it into repair until he or she could release it without slowing the aircraft already in the system.

At the second FRC point, a management rule ensured that an aircraft would not enter the buildup phase until at least 95 percent of the parts were ready. This rule prevented mechanics from working on multiple aircraft at a time and reduced pressure on the back shops and the supply system to attend to demands from multiple projects, allowing them to supply a complete set of parts for one aircraft at a time. Once all the parts were available, the supervisor sent the aircraft on for the buildup phase. The aircraft would then be assembled rapidly with no interruptions. The new process accelerated buildup by several weeks.

At the third FRC point, a rule governed when aircraft could move out of the dock for postdock work—only when 100 percent of the dock repair work was completed. Even though some aircraft now wait a few extra days for mechanics to finish the repair work, the overall process greatly reduces the effects of Parkinson's law at the dock and promotes the smooth flow of postdock activities. The team placed the fourth FRC point just before aircraft entered the paint phase, establishing a rule that mechanics had to finish all rigging tasks before they sent the aircraft on to the paint facility. They identified the paint facility as the constraint, the facility that controlled when arriving aircraft would enter the repair and overhaul process.

By controlling the release of work into the system, called *pipelining* at WR-ALC, we aim to have one aircraft in the repair phase, one in the hold phase, and one in the buildup phase at all times. Pipelining eliminates the tendency of the maintenance crews to multitask and facilitates a smooth flow of work.

Buffer Management

Once the fixer releases an aircraft for repair and overhaul, managers focus on the buffers, which are consumed when tasks are delayed beyond their (aggressive) times. Concerto prioritizes tasks and color codes them as red, yellow, or green, based on their impact on buffers. By managing buffers, the software ensures that all departments follow the same set of task priorities and base their decisions on the same priorities, which facilitates a smooth flow of work.

Every day support departments (planners, schedulers, parts specialists) look ahead at the tasks scheduled to start in the next five days in order of priority and make sure they can support these tasks in all possible aspects, such as parts, paperwork, special tools, and support services. The five-day look-ahead process ensures that mechanics assigned to tasks by production supervisors encounter minimal interruptions due to supportability issues. In addition, to ensure that at least 95 percent of routed and purchased parts are available before aircrafts enter the buildup phase, the software provides buildup dates well in advance, and support departments meet periodically to track and expedite parts.

Supervisors conduct daily production meetings, at which they look at the priority tasks across all aircraft and allocate resources, such as mechanics and tools, to tasks in the order of priority. They assign an adequate number of mechanics to each task and do not move them from the task until they finish the task, unless there is a nontrivial interruption to the task.

The supervisors actively manage tasks and resolve interruptions expeditiously. During every shift, they walk around the shop at least twice and talk to all the mechanics to resolve any problems. They also estimate how long it will take to finish the task and enter this number in Concerto at the end of the shift. The software calculates the impact on the buffers and produces a new priority list for the next day's production meeting. The goal of completing tasks by a predetermined due date has been replaced by a goal of finishing in-process tasks as fast as possible.

Supervisors meet twice a week to discuss the buffer status of all aircraft. They create a buffer-recovery plan for aircraft with a red buffer status (aircraft that have very little buffer capacity to allow them to finish the project on time). They make decisions on expediting tasks during this meeting. Buffer management is now the single mechanism for aligning management decisions during execution.

Challenges During Execution

To ensure that the first aircraft to be entirely scheduled under CCPM would be delivered in 160 days, we had to reduce the number of aircraft in the depot from 13 in April 2005 down to seven in October 2005, flushing the aircraft being worked on out of the depot even as aircraft continued to arrive for repair and overhaul at the old pace.

Second, we had to move from an aircraft-level priority mindset (where the leading aircraft is always the highest priority) to a task-level priority mindset (where a task required on a newly arriving aircraft could have a higher priority than a task remaining on an aircraft very close to its delivery date). Supervisors were initially mystified by the new priority system, but they quickly adopted it when they saw that all the aircraft were flowing through the depot much faster than before.

Third, we wanted to always move resources to high-priority tasks. To improve allocation of mechanics to tasks across all aircraft, we reorganized the production sections. Instead of assigning crews of mechanics comprising all skills to aircraft, we assigned mechanics with certain skills to tasks in any aircraft based on the task's priority. That greatly improved utilization of the skilled mechanics. Initially, supervisors feared that crews would lose the sense of ownership for aircraft. That fear proved to be without basis as the supervisors and mechanics realized that they "owned" all aircraft and would work on a high-priority task in any aircraft. The new structure broke down all the artificial barriers between aircraft and departments that had often led to suboptimal decisions.

Management is seeking to eliminate all milestone dates in meetings, reports, and tracking sheets. We have found that unless supervisors rely solely on managing buffers, people think more about meeting due dates, possibly succumbing to the student syndrome and Parkinson's law and exhibiting the behavior that CCPM strives to eliminate.

Reducing the WIP

WR-ALC implemented CCPM quickly. It went live with the new process, eight weeks from the start date. When implementation started in April 2005,

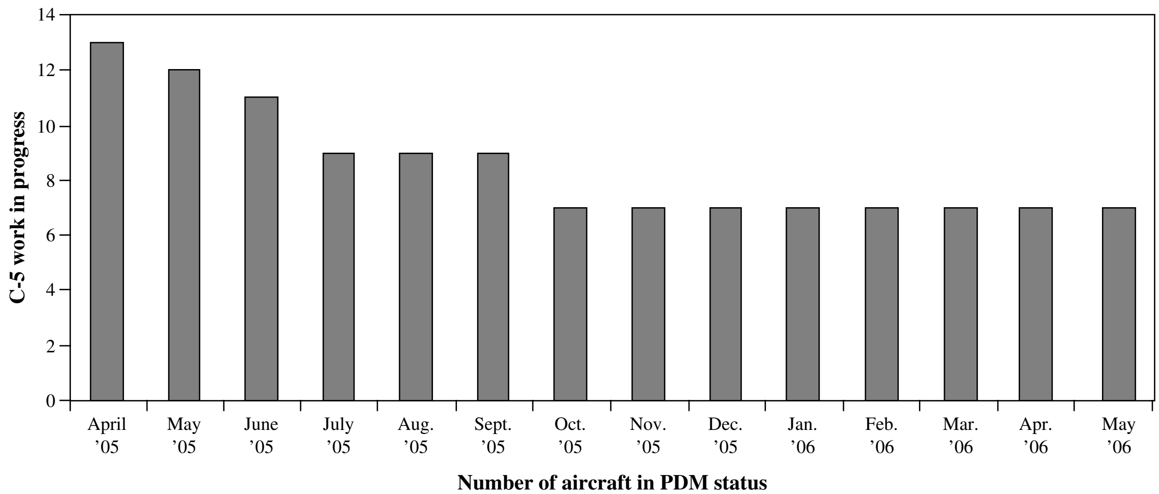

Figure 8: The number of aircraft dropped from 13 at the end of April 2005 down to seven at the end of October 2005 and has held steady at that number since then.

13 aircraft were in process. To achieve the target of 160 flow days without violating the first-come-first-served (FCFS) priority, the center had to get WIP down to seven. To achieve the target, we shortened delivery dates for WIP drastically, making most of the projects red and increasing red tasks. The entire organization rose to the challenge; everyone systematically focused on the high-priority tasks. The depot met the October 2005 target for flushing out WIP without adding overtime or labor. Furthermore, it did so without violating FCFS priority for aircraft already in the depot. Generally, as soon as the critical chain process goes live, the speed of execution picks up. Contributing to the dramatic increase in speed were the following factors:

—The software helps supervisors to allocate resources to the highest priority tasks instead of spreading them thin across multiple tasks. They thus reduce multitasking dramatically and speed up completion of tasks.

—The support departments (parts, paperwork, inspectors, planners, schedulers, engineers) now view the same priorities five days before the mechanics are scheduled to start tasks, which gives them time to provide full support. Mechanics can complete the tasks with few interruptions.

—Supervisors and other managers walk around the C-5 line daily to further cut down interruptions and the effects of Parkinson's law.

—Early warning buffer signals help managers react to potential problems promptly before they cause serious delays.

The reduction in turnaround times reduced the number of aircraft in the depot (Figure 8).

Improved Operations

WR-ALC established a new process for managing operations. It went from cell-level cycle times to one cycle time for the entire aircraft. This was a significant change for the better. It changed management of operations and allowed all the departments to focus on the same set of priorities despite the uncertainties they face every day. As a C5-planner put it, "Everyone is concentrated on the same priorities—engineering, material, production, management, and process improvement."

The CCPM implementation provided valuable information on what activities consistently consumed the buffer. It identified several high-leverage lean events. For example, Concerto identified floorboard replacement as a chronic problem area that consistently consumed the project buffer. A resulting lean event reduced floorboard replacement time by about 45 percent. In the CCPM implementation, we created management processes for both the production departments and the support functions, ensuring that they work together to deliver aircraft in 160 days.

Lean initiatives and CCPM work hand-in-hand to reduce flow days. We have found that organizations realize the power of lean initiatives after they implement CCPM. They can use CCPM performance diagnostics to reveal points in the process where work hours exceed those planned, affecting the critical chain and the project flow time. They can then use lean tools to remove the waste from these bottlenecks instead of conducting lean events in areas that have little impact on the overall flow time.

Communication between departments has increased, particularly concerning task priorities and buffer recovery. People now realize that red tasks do not imply poor performance; rather, they indicate a need for priority, focus, and management attention. The mechanics are generally very pleased with the new management system because they are not being asked to multitask. They appreciate the supervisors limiting interruptions to their working on a task. The supervisors periodically update the mechanics on their progress, giving them prompt feedback on how well they are achieving WR-ALC's overall goals. Most important, they do not feel threatened if they take a little longer to finish tasks than planned because a buffer protects each project from unexpected delays.

Supervisors and "team leads" can now focus on their priorities because (1) they have fewer aircraft to monitor and can identify what stage each aircraft is in at any time, and (2) CCPM's visual guidelines indicate which tasks (across all aircraft) require the most attention on each day.

Applying CCPM to Other Aircraft

WR-ALC's top management is trying to deploy CCPM to the software and development functions, to other support operations, and to aircraft at other US Air Force bases, such as the Tinker and the Hill Air Force Bases. WR-ALC has extended CCPM to the C-130s and the C-17s. The C-130 implementation went live on November 14, 2005 with planned reductions for flow days and WIP that exceed those achieved for the C-5s. Work is underway to implement CCPM for the C-17 cargo aircraft, which should result in huge benefits as the C-17 workload is increasing. Indeed, CCPM is increasing throughput per aircraft dock by 100 percent.

Measurable Benefits

Every day a C-5 spends undergoing repair represents a potential loss in revenue for the US Air Force. (The different areas of responsibility within a military service act as business units, passing dollars between them as revenue and expenses.) Using revenue figures projected by the US Air Force, we made a very conservative estimate of the revenue generated by a C-5 available for service as \$42,000 per day. Assuming 237 operating days a year, a very conservative estimate of the annual revenue generated by the five additional aircraft in operation is $\$42{,}000 \times 5 \times 237 = \$49.8$ million per year. This additional revenue represents an increase in the Transportation Working Capital Fund workload capacity and potential tangible future savings because any additional work performed would be accomplished at reduced cost.

The US Air Force faces a critical shortage of cargo-carrying capacity, especially in light of the war on terrorism. Because of the shortage of C-5s in service, it had been using C-17s to carry cargo that C-5s would otherwise have carried. It is difficult to estimate how much it would cost to build a C-5 aircraft because it is no longer in production. However, we can estimate the cost in another way: The C-5's cargo-carrying capacity is equivalent to that of two C-17s. WR-ALC released between five and six C-5s, which is equivalent in capacity to between 10 and 12 C-17s. The cost of a C-17 (manufactured by Boeing at Long Beach, California) is \$237 million. So, working with 10 C-17s, the cost for replacing the capacity of five C-5s, should that have been necessary, would have been about \$237 million $\times\, 10 = \$2.37$ billion. This is an immediate realization, not a discounted cash flow over a projected time horizon.

As a result of improved efficiencies, the C-5 organization recently reduced the cost of a major UDLM project by \$744,909 over the seven aircraft remaining to be worked on in 2006, and WR-ALC passed the savings on to the customer. This customer also has 97 military personnel stationed at Warner Robins working with the C-5 organization. While a C-5 is undergoing PDM, these military personnel perform other maintenance operations in the form of inspections and modifications that further enhance the aircraft. These maintenance operations must be completed before the aircraft can be put back into service, but they are not part of the PDM contract. The US Air Force recently announced that it is reducing its military staff by over 40,000. The customer asked the center to estimate the cost of taking on the work of the military personnel. The value of this work is approximately $400,000 per aircraft for 16 aircraft each year, or $\$400,000 \times 16 = \$6.4$ million per year. Because of efficiencies gained through CCPM, WR-ALC agreed to take on the customer's workload at no additional cost, saving it $7.2 million per year in military personnel costs and giving back 97 military positions to the customer.

Finally, when WR-ALC implements CCPM across all four types of aircraft, it will free up 11 dock spaces that it can use for additional work. With the freed dock space, the depot can readily accommodate additional customer orders of $65 million in 2006, projected as $248 million in increased orders by 2009. (The current operating budget for the C-5 line is $295 million, which includes revenue from both PDM and UDLM work.)

Aside from the quantifiable dollar savings, the CCPM implementation has provided numerous nonquantifiable benefits: (1) The five C-5s returned to the Air Force provide an additional 180 ton-miles of airlift capability. (2) Unlike C-17s, the C-5s are able to fly into the area of operation. The US Air Force estimates that the five C-5s have reduced the number of dangerous convoys in the war on terrorism, saving the lives of many American troops. (3) With the increase in C-5 availability, the US Air Force has improved its speed in shipping cargo that it would otherwise have transported by surface, for example, by ship.

The CCPM implementation has reenergized lean events, which produced many benefits. For example, CCPM identified the paint barn as a constraint, following which, a lean event reduced the time an aircraft spent in the paint barn by 25 percent. With the freed capacity the center will paint five C-17s in the C-5 paint barn during 2006. In 2007, it plans to use the C-5 paint barn to paint all the C-17s it has contracted to paint for Boeing. (A new paint barn earmarked for the C-17s will not be ready on time.)

Conclusions

Despite several changes in management, CCPM's stability and continuity helps new managers and employees to perform effectively very quickly. CCPM and lean initiatives work very well together: CCPM focuses the organization in conducting successful lean events. It gave WR-ALC the confidence to take on additional workload and to execute it effectively. For example, WR-ALC now completes UDLM work on the C-5 torque deck, which it took on a few years ago, in less than 30 days, fast even by commercial-sector standards. It plans to take on even more PDM work in the future.

References

Elton, J., J. Roe. 1998. Bringing discipline to project management. *Harvard Bus. Rev.* (March–April) 153–160.

Goldratt, E. 1997. *Critical Chain*. North River Press, Great Barrington, MA.

Gutierrez, G. J., P. Kouvelis. 1991. Parkinson's law and its implications for project management. *Management Sci.* **37** 990–1001.

Lechler, T., B. Ronen, E. A. Stohr. 2005. Critical Chain: A new project management paradigm or old wine in new bottles? *Engrg. Management J.* **17**(4) 45–48.

Leus, R. 2003. The generation of stable project plans: Complexity and exact algorithms. PhD dissertation, Department of Applied Economics, Katholieke Universiteit Leuven, Leuven-la-Neuve, Belgium.

Little, J. D. C. 1961. A proof of the queueing formula $L = \lambda W$. *Oper. Res.* **9** 383–387.

McKay, K. N., T. E. Morton. 1998. Critical Chain. *IIE Trans.* **30**(8) 759–763.

Newbold, R. C. 1998. *Project Management in the Fast Lane: Applying the Theory of Constraints*. Saint Lucie Press, Washington, D.C.

Patrick, F. S. 1999. Critical Chain scheduling and buffer management—Getting out between Parkinson's rock and Murphy's hard place. *PM Network* **13** 57–62.

Raz, T., R. Barnes, D. Dvir. 2003. A critical look at Critical Chain project management. *Project Management J.* **34**(4) 24–32.

Srinivasan, M. M., D. Jones, A. Miller. 2004. Applying theory of constraints principles and lean thinking at the Marine Corps maintenance center. *Defense Acquisition Rev. Quart.* (August–November) 134–145.

Stewart, D. 2005. Global and European MRO market. Netherlands Aerospace Cluster Symposium, The Hague, The Netherlands.

At the Franz Edelman Award Competition, May 1, 2006, Ken Percell, chief operating officer, Warner-Robins Air Logistics Center stated, "The increase in C-5 availability has generated an additional 180 million ton-miles of airlift capability. For our Air Mobility Command operators, that will result in revenue generated of 49.8 million dollars per year. While our line required 12 aircraft, global mobility depended on realignment of C-17 aircraft to perform some critical C-5 mission. The additional five C-5s had a replacement cost based on C-17 equivalents of 2.37 billion dollars. This is an immediate realization, which has made it easier for the air force to discontinue C-17 production early as the C-17s return to their original missions.

"Customer operating costs due to maintenance costs have also been reduced. A major operation replacing floorboards will cost 745 thousand dollars less for the seven remaining aircraft to be worked on in 2006 and will be passed back to our customers this year.

"But there is another key consequence that we measure not in dollars but in human lives. The five C-5 ... will immediately reduce the dangerous convoy operations ..., saving uncounted lives that might have been lost in these dangerous operations. This has been a key motivation for our maintainers, and many of these convoy personnel are their neighbors and friends currently deployed to Operation Iraqi Freedom."