White Paper:

The Transition to a Constraint Based Maintenance Organisation In The Airline Industry

The TOC Centre of Australia Pty Ltd ABN 27 091 919 961 Level 25 Chifley Tower, 2 Chifley Square Sydney 2000
PO Box 1541 Bondi Junction 1355 Phone: +61 2 9293 2844 Fax: +61 2 9315 5891 E-mail: tocca@tocca.com.au
Web: www.tocca.com.au

TABLE OF CONTENTS

1 INTRODUCTION

The airline industry is arguably facing some of the toughest trading conditions in its history. Load factors are reducing, margins are tumbling and many airlines are on the verge of bankruptcy or are already in Chapter 11 and beyond. Assuming that there is little leverage to be gained from the demand side, as more and more seats chase fewer and fewer customers, it is worthwhile to consider what role supply side management can play in improving productivity, and thus the competitiveness of an airline.

This white paper is designed to address the issues surrounding a key area of the supply side – that is aircraft maintenance – and show how a constraint based management approach can significantly increase capacity available from existing assets.

This freed capacity – bought at no cost beside the productivity gain, can be used to:

- Defer new capital outlays
- Service new markets at marginal costing
- Improve flexibility for sector scheduling
- Enhance the airlines competitiveness ahead of any improvement in the trading environment

Before going into the detail of how a constraint based management approach can help in producing the abovementioned benefits, it is worthwhile recalling the words of Deming – founder of the quality movement.

Effective transformation of a system isn't possible without profound knowledge – this profound knowledge comprises of four major areas:

- Appreciation for a system
- Knowledge about variation
- Psychology of individuals, society and change
- Theory of knowledge

In this White Paper we will address all four areas.

2 APPRECIATION FOR A SYSTEM

2.1 The system based view

The system based view of the world starts from the basic premise that if we are to make significant improvement in our performance we must:

- Identify the boundary of the system in question
- Articulate what the intended goal of the system is
- Know the inputs, processes and outputs required of the system
- Be able to measure how we are doing in producing more or less of the system goal

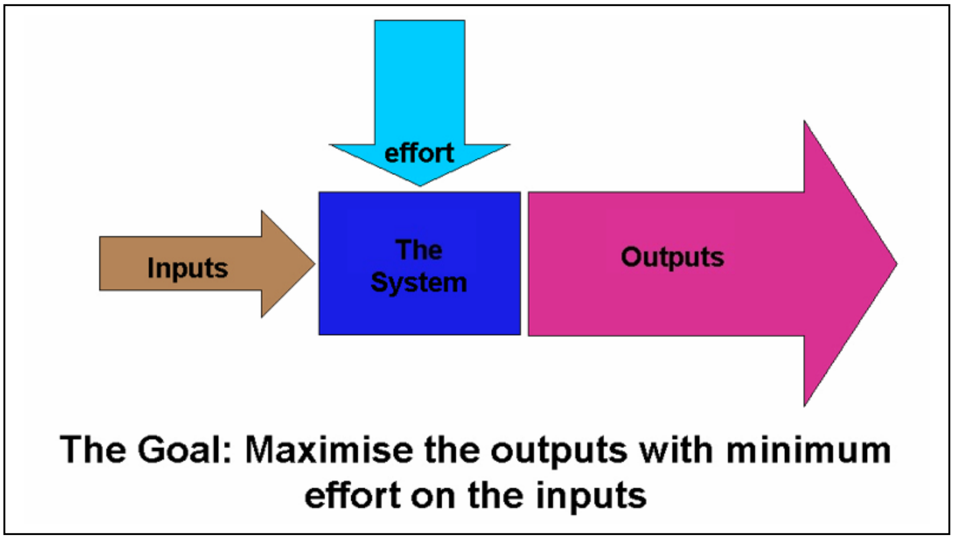

In the case of heavy maintenance in the aircraft industry, we could define the boundary of the system as being the maintenance department of the airline, or indeed that of a maintenance outsource facility.

The goal of the system could be expressed as maximising the availability of a safe aircraft to the airlines operations whilst minimising the cost of so doing.

The inputs could be the plane itself and all the spares and facilities required to keep it in safe condition.

The effort could be translated as the skills and work required in order to perform the maintenance work.

The output is a safe aircraft

2.2 The measures

It is when we enter the realm of system measurements that we find ourselves in a controversial and often contentious arena.

There is a fundamental assumption, which prevails in almost all commerce and industry, that the way to get best performance from a system is to ensure that the individual parts of it are all kept as busy as they are able to be. Implicit in the assumption is that the sum of the components activities (or local optima) is equal to the system outputs (global optima).

The effect of this assumption is that most industry measures are based in some shape or form on activity or utilisation. Rare is the Company that is spared from such a devastating assault on genuine productivity improvement.

How can we so boldly say that a measure such as utilisation is devastating from a productivity point of view?

This is where the insight of the Theory of Constraints (TOC) proves most useful. Simply put, TOC states that any system has within it a constraint -something preventing it from achieving more of its goal. Were this not the case, the output of the system would tend to infinity, or collapse to zero

If the Laws of Thermodynamics are to be believed (and we have found no evidence to the contrary!) there does not exist in the world a perpetual motion machine that will produce an infinite amount of output for a finite amount of input. Put another way, there is always friction, and that which causes the friction we can call the Constraint.

Coming back to our example of the aircraft, we could argue that we would like the turnaround time for maintenance to be zero days, as that will give us its maximum use. As we all know that a zero lead time for the maintenance of an aircraft is impossible, we could define the constraint of aircraft maintenance as being all those things which prevent such a fantasy from happening.

Defining it a little better, we could argue that the constraint of achieving a zero lead-time is a function of:

- The sequence of tasks which need to be undertaken,
- The availability of the facilities required to do that sequence of tasks
- And the availability of the resources and tools needed to do the work

Based on this model, it does not take too much effort to consider that not all tasks will be of equal importance in the completion of the project of maintaining the aircraft.

In fact, we have a useful tool called the Critical Path Method, which defines the longest set of dependant steps through a project network as being the Critical Path. By definition, it is the shortest lead-time in which a project can be completed. Therefore, those tasks, which are not on the Critical Path, are by our definition not a constraint of the project.

Let us assume for the moment that the facility for maintenance is not an issue capable of delaying our project of aircraft maintenance. This may not always be the case, but generally there is enough maintenance hangar facility either in-house, or out-sourced, which can accommodate aircraft maintenance requirements within the timeframe mandated by the aviation safety regulating bodies.

Our next potential area for a constraint lies not only in task dependency, but we must also view the project through the lens of the resource manager, who almost always never has enough of one particular competency or another to satisfy all competing demands of the Critical Path simultaneously. Where there is contention for resources, we must accommodate this by either accepting a longer lead-time, decreasing the scope of the work to be done or adding additional resource of the relevant competency.

A clearer picture should now be coming into view that not all tasks are of equal importance in completing a maintenance check in the shortest possible lead-time.

A direct consequence of this idea then is the flip side of the Theory of Constraints – namely, that that which is not the constraint will not be fully utilised. If the logic of the argument is true, then why do we use local utilisation and activity based methods as a primary measure of system performance?

To quote Deming again:

"The object of any component is to contribute its best to the system, not to maximise its own production…some components may operate at a loss themselves in order to maximise the whole system…"

What measures then are available to us if, as we have demonstrated, utilisation, or local optimisation does not serve the goal of getting the aircraft out in the shortest possible lead-time?

TOC has three fundamental measures, which align the decision-making and behaviours of the shop floor to the corporate goal of increasing ROI for the shareholders:

Throughput: The rate at which the system generates money through sales. In the instance of the airline, we could calculate this as the amount of average contribution margin lost for each day the aircraft is unavailable for service. Another way of stating the definition would be sales revenue less directly variable sales costs – the definition of contribution margin.

Operating Expense: This is defined as all the money the system spends on turning inputs into outputs, or for simplicity's sake, labour and overheads. Note that no distinction is made between direct and indirect labour – the sole criterion for being an overhead is if the expense has to be paid for regardless of the amount of output generated by those costs.

Investment: The assets required to produce Throughput. In our example the biggest asset would be the aircraft itself, probably followed by the inventory, facilities and equipment required to keep the aircraft safe.

If the global measure of corporate success is ultimately a function of return of investment ROI (or even more so in a listed company, the rate at which this improves), then we can link our local measures to the ROI of the company thus:

If net profit is by definition equal to the Throughput (**T**) less the Operating Expense (**OE**), and ROI is net profit divided by investment, then:

$$\mathbf{ROI = \frac{T-OE}{I}}$$

Thus, if one were looking for significant improvement in ROI, the need would be to increase Throughput, whilst holding steady or reducing Operating Expense and Investment.

There is a floor to reduction of operating expense below which one cannot go before it starts to erode service levels, Throughput and competitiveness. Most enterprises will have focused heavily on this area in recent years and would battle to find the kind of quantum gains required to provide the base for long-term competitive advantage.

If however, the focus was on cutting the lead-time of maintenance checks, the benefit of having the craft available for service would have a direct bearing on Throughput and Investment as less aircraft would be required to satisfy the demand requirement for seats.

An increase in Throughput and reduction in Investment would have an impact of an order of magnitude more that trying to shave pennies off operating expense.

In summary, taking a system based point of view, we must abandon the notion that the sum of the activity, or utilisation is what will drive competitive advantage. Focusing on the system as a whole to reduce lead times provides us with the biggest lever available to increasing ROI for the enterprise.

3 KNOWLEDGE ABOUT VARIATION

3.1 Common cause or special cause?

Variation, or the Murphy factor haunts our best-laid plans. Dr. W. Edwards Deming identified two types of variation:

- Common Cause Variation: A cause that is inherent in the system. The responsibility of management.
- Special Cause Variation: A cause that is specific to some group of workers, or to a particular production worker, or to a specific machine, or to a specific local condition.

Dr. Deming notes that managers often make many systems worse by not understanding the fundamental difference between these two types of variation. He also notes, "I should estimate that in my experience most troubles and most possibilities for improvement add up to propositions something like this:

94% belong to the system (responsibility of management)
6 % special."

3.2 The task versus the project

The standard way of approaching the problem of variance in a maintenance environment is to attempt to get better and better estimates of confidence in completing a task in a given amount of time. The focus is placed squarely on the task, which has the negative effect of building in contingency at every possible point to cater for the almost inevitable visit of Murphy. We are certain he will visit, we are just not sure when and how.

The goal we have defined – improved ROI - demands that we complete the project of the maintenance check in as short a lead-time as possible. How do we complete in the shortest possible lead time if we are adding contingency to every task, knowing full well that not every task will require the contingency, but being too unsure of variance to cut the padding out?

3.3 Murphy and dependent events

At the level of the task, we may even be inclined to do research, work-study and the like to get better and better estimates as to the time required to undertake a task. What is often forgotten or not fully understood though is a law of statistics that states that if we have 5 tasks that have to be undertaken sequentially to produce a desired output, each with a confidence level of as high as 90%, the confidence we have for the whole sequence of those tasks is $.9^5$ or 59%. This effect of combining variation with dependent events is often ignored and the consequence is almost always longer lead times than anticipated and hence poor ROI performance. One would have to get processes that yielded six-sigma variation (3.4 defects per million operations) to eliminate this effect of Murphy and dependent events - how likely is this in the maintenance environment?

The maintenance arena, by its very nature, is highly prone to Murphy, as one never knows until a specific maintenance routine has been undertaken whether there will be additional work arising or not. Some inspections yield little or no problem whereas some might require a complete refit.

3.4 The global perspective – aggregating risk

If we attacked the issue from the global perspective though, we may be able to find a way for the law of averages to work in our favour. Suppose for each task we took only the 50% confidence level for task completion. This means that in half of all cases the task at hand would be completed within the specified duration, but equally, fully half would not. Furthermore, let's not get too pedantic about what we mean by 50% confidence level, as it might never be something we could determine by objective empirical means. What we mean by 50% confidence is the amount of time it would take to do a task, assuming the following conditions were met:

- You were ready to work on it immediately
- You were not interrupted from beginning to end of the work package
- All the tools you required were readily available to you
- All the spares you required were readily available to you
- Any assistance you required was readily available to you

In other words, with all else being equal, how much dedicated effort would be required to complete the work package?

Our next step is to add up all of the individual task durations arrived at in this way, and then determine what the longest set of tasks is through the project taking both task precedence and resource availability into account. This combined duration of all these tasks we could only be confident of achieving half the time we undertook the maintenance of our aircraft and no doubt operations would be unwilling to live with such a high degree of uncertainty.

Thus, we need a mechanism whereby we can aggregate the risk of the individual tasks so that we can be reasonably confident (90-95%) that the project as a whole will be completed in the specified lead-time even if the individual tasks making up that chain of events vary in their duration according to the games Murphy is playing on any particular day.

If we were to add an unresourced time buffer equal to 50% of the duration of the longest chain through the project, and call this the Project Buffer, it would provide us with a means of assuring the maintenance check came in on time despite what was happening at the local level.

Furthermore, if we were to protect this “Critical Chain” (Critical Path plus resource dependency) from variation of “Feeding Chains” – those chains in the maintenance process, which have some degree of slack in them – we then have a robust means of dealing with variation.

How so? We first need to have a look at the psychology and behaviour of individuals and the groups they work in.

4 PSYCHOLOGY OF INDIVIDUALS, SOCIETY AND CHANGE

When we are asked how long it will take to complete a certain piece of work, as discussed previously, we are inclined to make a worst case estimate. We know uncertainty exists, we simply don't know where it is going to strike and would thus like to protect our estimates. Furthermore, if the measure of performance is at the level of compliance with task duration, all our behaviours will be geared around making sure that we hit our mark as often as possible, thus perhaps stretching task times even further.

4.1 Multitasking

Besides the effect that Murphy and measures have on task duration estimating, there is also the case to be made that we are never left alone to complete a single piece of work. Always there are interruptions, the demand to keep as many balls in the air as possible because someone, somewhere worked on an assumption that capacity to undertake tasks was to all intents and purpose infinite – or put another way that the resource manager has the people with the competencies he or she requires whenever the demand is made on them. If one has already given a 90% confidence level for task duration estimate, and assuming that not all tasks take their 90% confidence duration the temptation will always be there to get off what you are working on and commence something else which appears more urgent as there will always be time to come back and fix the less urgent thing later. Besides, the measures of utilisation are rewarding me for keeping myself as busy as I can all of the time.

4.2 Student syndrome

Under the existing scheme where task conformance is the measure of performance, I will have given my 90% confidence estimate for all my task durations and I will be inclined to postpone my commencement of a task to the last possible minute. I know that in certain instances I should be able to comfortably do the work required in much less time than estimated in the plan. What I ignore in so doing is that Murphy has a habit of striking when you have the least amount of time to deal with it. This phenomenon of putting off until tomorrow what I could do today is called student syndrome – it being used most often by students in arguing for an extension period for the deadline on an assignment and then doing the work the night before the deadline expires.

4.3 Parkinson's Law

The flipside of student's syndrome is the phenomenon of Parkinson's Law, which states that work will expand to fill the time available. If I have overestimated the duration of my task, I am hardly likely to report that I have finished early. Human behaviour suggests that we might be accused of being poor estimators, or our colleagues might not like the idea of having a tall poppy around. Besides, perhaps this time around the force was with me and next time I won't be so lucky.

Multitasking, student syndrome and Parkinson's Law, when combined with the effect of Murphy and dependent events makes one begin to wonder how anything ever gets done on time.

5 BUFFER MANAGEMENT – BRINGING IT ALL TOGETHER

Let us recap what we have discussed so far and understand how fundamentally our conceptualisation of our project of aircraft maintenance has changed

- **Critical Chain** – we have identified the longest set of dependent steps through the project based on task precedence and resource contention, based on a 50% confidence level.

- **Project and Feeding Buffers** – we have placed the sum of the difference between the 50% estimate and the 90% estimate into the hands of the project manager – that is, the person held accountable for delivering the aircraft to operations on time. By looking at the amount of work left to be done within a particular chain against the amount of aggregate contingency left, this person is able to determine where to focus his or her resources to keep the project from straying off track.

- **Variability: 50% estimating** – the use of a 50% confidence level time estimate, rather than a 90% estimate, results in people working towards the 50% time estimate. This has the effect of reducing the occurrence of Parkinson's Law (work expands to fill the time available) and Student Syndrome (wasting time at the beginning of a project or task). Protection against Murphy's Law (things going wrong unexpectedly) is contained in the strategically placed buffers. "Butt-covering" is reduced as all tasks are only expected to complete on time in only half of all cases. The focus can shift from task conformance to project performance. Emphasis is placed on what needs to be done to keep the flow of the project moving.

- **The measures** – we have replaced the traditional measures of utilisation with Throughput, Investment and operating expense and in any decision we have to make on the shopfloor, we will look at the impact on all three and the resulting change in ROI before responding to a knee jerk based on cost alone.

  Furthermore, we have the ability to measure how the project as a whole is doing by seeing the amount of chain consumed versus its corresponding buffer. This supports the focus on management effort toward those things that are causing undue penetration of buffers and leaving well alone those things that do not require attention.

6 THE THEORY OF KNOWLEDGE

How do we know what we know? What assumptions do we carry with us when asked to improve our systems performance relative to its goal? How much of our habit do we carry with us in determining what would be good for the future?

The best means that mankind has come up with in determining how we know something is through the use of the scientific method. I state my hypothesis and you are under obligation to collect evidence to prove me **wrong.** Only by satisfactorily being able to deal with those instances where my hypothesis does not fit will I be able to produce a better model of reality and thus progress my understanding and deepen my knowledge.

6.1 The Thinking Process

TOC has a robust process of thinking that has been designed using the scientific method to ask the following key questions of our systems:

- What must we change?
- What must we change to?
- How will we bring about this change?

What the Thinking Process is not is brainstorming, or a SWOT analysis or any one of those methods commonly used to develop ideas. It is a logical toolset with firm rules for separating cause from effect and driving down to the systemic root cause of system underperformance.

The Thinking Process examines in a robust way what underlying conflicts perpetuate the systemic root cause – for example: should we hold more inventory and use more working capital with the aim of always having the right materials on hand and thus decrease turnaround time, or should I hold less inventory and risk delaying the completion of the maintenance project?

At $60 000/day of Throughput lost whilst the aircraft waits, these are not trivial conflicts to resolve. What are the assumptions that underlie these core conflicts? How do we surface them so we can examine their validity? What mental model do we carry around unwittingly; not knowing the devastating consequence these assumptions will have on goal accomplishment? The Thinking Process is designed to surface all the assumptions and thus create an environment for creative, breakthrough thinking.

Once developed, the breakthrough thinking is tested, again using a logical construct. What will the future look like if we had our breakthrough idea in place? Would we remove the undesirable effects of our existing process and methods? Can we build virtuous loops into the logic of the cause and effect of my future reality such that it becomes self-perpetuating with success breeding further success? What if we could predict what the potential negative impact of our proposed solution is with our logic and design the elimination or at least minimisation of those effects into my plan for improvement?

At this point, we would have a clear and shared vision of what “good” looks like into the future, but there is still the task of determining what obstacles lie in the path of us getting to that future. Once again we resort to logic and work on the assumption that if we can identify the obstacles to goal achievement (and we are much better at finding obstacles than agreeing on solutions), then by designing the intermediate objectives to our goal around overcoming the identified obstacles, we will be able to reach our goal in bite sized chunks.

These bite-sized chunks can then be broken down to the level of tasks, with specific outcomes, against which competent resources are assigned with durations (50% confidence estimates of course) and sequenced according to the logic of the project.

In a nutshell, the approach starts with identifying the undesirable effects or symptoms for each of the key stakeholders and the core problems or constraints causing these symptoms. The approach ends with a detailed action plan that will be both necessary, sufficient and sequenced correctly to align the organisation to identify, remove and replace the constraining policies, measurements, technology, processes or behaviours.

6.2 Resourcing and sequencing

6.2.1 Phase 1 – the core team

A prerequisite of the change process to constraint-based maintenance is to subject the core team, fully representing the functional areas of the business, to the discipline of the Thinking Process. This has numerous virtues, amongst which are:

- The rigours of a logical process means that a very powerful consensus is built amongst the team on the basis of common sense. It has sometimes been described as the “Pegboard of Knowledge” of the business. Whilst some might be describing a tail, some a trunk and some a tusk, by the end of the process they will all know they are dealing with an elephant!
- Whilst the tools might be those of TOC, the thinking is that of the stakeholders, and as such a strong sense of ownership is gained.
- Because the whole basis of the construction of the tree is logic, anyone can feel free to question the logic with less fear of wounding egos.
- The tool is a powerful means of presenting the thinking of the core team to all other stakeholders – be they subordinate or superordinate in organisational rank.

6.2.2 Phase 2 – second tier senior staff

Once the core team has set the strategic goal and plan, their thinking is presented to the second tier of senior staff. It is not done as a fait accompli, but one would encourage all to use the scientific method described above.

This has the very great advantage of providing a means by which the core team's reasoning can be scrutinised without fear or favour and significant buy in occurs as the logic of the argument is strengthened by the experience and common sense of the growing team.

6.2.3 Roll out

By now, all management and staff will have bought into the plan, understand and be committed to its logic and it will be ready for execution. Now that we have planned the work, it is time to work the plan.

6.3 Frequently asked questions

Some frequently asked questions are listed below with suggestions as to how they may be addressed. It may not be right, but it should prompt the question: How do I know what is right for us?

1. Does the organisational structure we envisage suit the purpose of the task at hand – significantly reducing the lead-time on our maintenance checks?

    Traditionally aircraft maintenance has been organised along the functional areas required to do maintenance. Demarcation has been strong between airframe, sheet metal, avionics and others. The job of coordinating roles was loosely left to the Duty Shift Manager under the advice of a disempowered scheduling team. Many times this leads to sub-optimisation with shop floor staff concerned about being "sold" to another aircraft or trade. Hard jobs are often handballed to an incoming shift as selection of job cards is largely left to leading hands.

    In the new order, it will be essential that maximum use be made of multiskilling. The Critical Chain and buffers will be used to determine who gets to work where with the overall goal being the reduction of lead-time in maintenance checks. It may be necessary for an avionics expert to remove seats as all subordinate to the goal of quicker turnaround times. The master scheduler will require a much more powerful role in determining where resources are allocated and what they are required to do.

2. How will my 3rd party supplier (e.g. a refit workshop) know when to start work on product for my aircraft if I only have a 50% chance of telling him when his product is required?

    If we assume that your supplier can reliably supply you within a specified lead-time, it would be easy to create three tasks for each activity associated with an external supplier, which can be incorporated into the schedule network:

- A notification task that product has been stripped from the aircraft once its predecessors have been completed.
- A task using the third party as the resource, equal in duration to the reliable lead time quoted by the third party
- A task notifying receipt of the product from the third party

If these steps are followed, and a feeding buffer placed before the integration of those tasks back into the project network, by monitoring the buffer, one can proactively manage the refurbishment from notification to tracking degree of completion and then receipt.

3. How will I know when my maintenance lines will be available and when I can offer the aircraft back into service?

    Critical Chain method has demonstrated that it can comfortably reduce aircraft maintenance lead times by 1/3 or more. If we assume in the first instance that both the line and the aircraft are unavailable for the traditional duration, then the worst that can happen if a similar 1/3 reduction in lead-time is achieved is that the capacity remains idle until the next aircraft is scheduled to use the facility.

    The benefit of having the aircraft available for service early at no additional outlay of dollars must be of significant benefit to the airline.

    Once confidence is gained in using the method and maintaining reduced lead-times, the requirement for the facility and the availability of the aircraft can be predicated on the new realities.

4. What will happen if my resources are idle due to improved performance and their utilisation drops?

    A legitimate fear in the change process is that released capacity will lead to retrenchment and the prospect of this will inhibit people from really accomplishing the benefits of the new processes.

    This is fairly easily addressed by in the first instance bringing in-house all out sourced work, and if there is still additional capacity offering the maintenance service to third parties. It would by definition be highly competitive as, with the same resource base, the enterprise would be doing all of its work and be able to look at marginal additional revenue for no additional cost.

5. I have looked at my tasks and they already represent the 50% confidence level. If I add a buffer it will simply overextend the lead-time.

    If the current work practices involve any of the phenomena described above (multi-tasking, student's syndrome, Parkinson's, Murphy) then built into the estimate of the 50% is an unstated acknowledgement that things take longer than they could if they were run on Critical Chain principals.

It is often argued that factory standard times are unrealistic as they represent best achievable in a factory environment – coordinated priorities in an organised structure. Sure begins to sound like 50% confidence!

If worst comes to worst one could cut the duration of all tasks by 1/3 and replace the time saved with buffers. Penetration of the buffers will soon tell you which tasks need closer examination.

6. I will find it difficult to persuade management and staff that it is good for the system that some people are idle some of the time.

   There are those who believe that the Theory of Constraints doesn't apply to them. It's as if they're saying gravity doesn't apply to them. When they have evidence of an unconstrained system, then we may be able to accept their argument. In the meantime, common sense says that if a system has a constraint, then that which is not the constraint, by definition, is not fully utilised. Being not fully utilised is another way of saying idle. Surely it is better to be available to support the critically constrained resource than be measured for utilisation on the mistaken premise that the sum of the local efficiencies is equal to that of the global efficiency?

   Does the Grand Prix driver get out of the car to change the tyres when a second saved in the pits is a second gained in the race?

   Does an orchestra play as loud and fast as it can to ensure all the musicians are earning their keep?

   There are only two choices – you manage the constraint, or it manages you!

7 SOME RESULTS

7.1 Israel Aircraft Industry

The Israeli Aircraft Industry employs about 15,000 people. A major function is to maintain Jumbo Jets used in Passenger service. A particular type of maintenance, called ‘type D,’ normally takes 46 days in the industry. The penalty for non-performance to schedule is very steep...$60,000 per day, because the airlines need the planes back into scheduled service. The company had been paying up to $ 25 million per year in penalties. A letter from the manager to Dr. Goldratt (included on http://www.Goldratt.com) notes, “...we succeeded to drop our average Turn Around Time per Aircraft Visit from three months to two weeks and to increase our backlog from two months to one year.”

7.2 Naval Air Depot at Cherry Point, NC

The air depot is a large and complex repair and overhaul aerospace activity. They perform in-depth scheduled and unscheduled maintenance on a variety of naval aircraft, engines and components. The annual revenue is $650 million and they employ 4,500 people.

The have considerably complex bills of material and routing process. Master Scheduling tasks for repair often considered as projects or job shop environment with substantial variability.

The prototype aircraft selected for proof of concept was the H-46 helicopter. Collaboration with Vector Strategies using our TOC based critical chain scheduling techniques yielded truly impressive results. Over a 2-month period the **amount of aircraft in flow was reduced from 26 to 14**. This first step was crucial; maintaining 14 aircraft in flow matched their available capacity to the customers demand. The system transformed from a push system to a pull system. The clean and strip supervisor Mark Meno remarked, "Pull work through don't push it."

The average turnaround time, due in part to an aging fleet of aircraft, was approaching 225 days. **Within 4 months of implementing the Vector Strategies approach the turnaround time was reduced to 132 days**. This remarkable improvement was accomplished without hiring any additional people. In fact the overtime was now considerably lower than before.

There was a quantifiable improvement of the already high quality product being produced. The amount of rework was reduced and the feedback from customers showed a demonstrable increase in satisfaction. Jerry Frontera, local union steward commented, **"Your approach really works, grievances are down and people doing the work are feeling good."**

7.3 United States Marine Corps Material Command

The mission of the Materiel Command is to provide the highest level of materiel readiness to the United States Marine Corps. They have two maintenance centres that perform depot level repair on their track and wheeled vehicles. The operating forces send vehicles to the centres when depot maintenance and repairs are required and issued a replacement if assets are available.

The maintenance centres are a critical link in the readiness supply chain and have reputations for being very responsive. Similar to private business their customers demand quick turnarounds and competitive prices. Pressure to meet promised schedules is becoming greater. Many of the leading edge best business practices have been implemented and a powerful ERP is in place. But, it is not enough.

A prototype effort cut the cycle time in half within three months. Variation in their processes has been dramatically reduced. The prototype team leader Joe Frisone remarked, "Since we started using TOC, this is the first time we haven't been hit by the end of month bow wave."

They are now committed to implementing TOC throughout both maintenance centres as quickly as possible.

According to Col. Rivers, the commander of the maintenance centre at Barstow, "TOC creates a production planning activity that is detailed and comprehensive. The result is improved use of resources and better scheduling and production flow across cost work centres."

7.4 Pearl Harbour Naval Shipyard

A unique approach to managing submarine maintenance projects is bringing positive results at the Pearl Harbour Naval Shipyard and Intermediate Maintenance Facility (NS&IMF).

The shipyard's Fleet Maintenance Availability Project for Submarines (FMB) is using critical chain project management to improve schedule and performance on submarine availabilities.

For years, critical path project management has been the standard for Navy shipyards and industry as a whole. In the critical path process, a schedule is developed that allots enough time for jobs so that they, and consequently the project, will finish on time with a 90 percent probability.

In critical chain project management, the duration of jobs is shortened so that half of the time they will finish late. The time that has been cut from these jobs is placed in buffers in certain strategic points of the schedule.

“It’s a change in the way we execute and plan work,” said Cmdr. Kent Kettell, Assistant Operations Officer for FMB. “It minimizes chaos and fire fighting.”

Although at face value, the idea of deliberately starting a project late seems counterproductive, critical chain management actually helps the shipyard meet or beat project end dates. In critical path management, even if a job finishes early, the next job in line doesn’t begin immediately because it’s either still scheduled to start at the original date or a resource isn’t available. Critical chain schedules encourage a “relay race” behaviour, where workers finish a job as quickly as possible and, without delay, pass the baton to others on the next job. The intention is to capture gains in schedule and pass them on.

Since critical chain project management was implemented last year, FMB has improved its schedule performance, finishing its last 13 availabilities on time. The average number of jobs done per upkeep has also risen while lowering the number of man-hours required.

Metrics comparing performance in fiscal year 2002 before and after the inception of critical chain project management show 11 percent more jobs done per upkeep while using five percent less man-hours. There was also a 13 percent increase in job completions. All of this was accomplished while the average length of availabilities was reduced 5.6 days.

Shipyard workers see and feel the results. “I’ve seen the difference,” said Roxanne Bataya, a shipfitter supervisor. “We aren’t bouncing around like before.”

Shipfitter Richard Donald agreed. “It seems like we do less jobs at one time. It’s more organized.”