

Theory of constraints at UniCo: analysing *The Goal* as a fictional case study

Ronald S. Tibben-Lembke*

Managerial Sciences Department, University of Nevada, Reno, NV, USA

(Received 13 June 2008; final version received 12 November 2008)

As a fictional case study, Eliyahu Goldratt's novel about manufacturing, *The Goal: A Process of Ongoing Improvement*, presents a constraint-focused approach to production management. As a novel, the book does not emphasise the quantitative details of the plant improvements. However, a great amount of information about the plant is spread throughout the book. By collecting and analysing this data, a concrete picture of the plant's capacity and its improvements may be developed, which can greatly help the book's readers understand and evaluate the cumulative impact from the plant's 'process of ongoing improvement'.

Keywords: production planning; theory of constraints; drum buffer rope

1. Introduction: *The Goal* as a fictional case study

Eliyahu Goldratt's manufacturing novel *The Goal: A Process of Continuous Improvement* has inspired countless professionals in production (and many other fields; Whitford 2004) to embark on their own efforts of continuous improvement. As Rand (1986) writes, 'It's a novel, but it's also a manufacturing text-book, and it's good on both accounts.' Many reviewers have agreed *The Goal* is an easy-to-read way to get an introduction to production realities (Rand 1986, Belis 1994, *The Economist* 1995, Dani 2006). However, no one has taken a detailed look at the numbers presented in the book as a fictional case study.

The concept of drum-buffer-rope (DBR) production control has been discussed considerably in the literature (for just a small, recent sample, see Ye and Han 2008, Jodlbauer and Huber 2008, Watson and Patti 2008). The first formal presentation of DBR was by Goldratt and Fox (1986), but before that detailed presentation, it was first known in *The Goal*, and that is where most practitioners likely first learn about it.

After a casual reading (or even a careful one), the reader may be left with the impression that many improvements were made, but not have a clear idea of exactly what happened when, and how they all fed together to bring about the changes seen in the plant. This paper analyses the data presented during the course of the story. In so doing, there

*Email: rtl@unr.edu

are several objectives:

- (1) To create a complete understanding of the plant's operations from the details scattered throughout the book.
- (2) To assess the impact of the actions taken in the plant.
- (3) To see if all of the details discovered do add up to form a cohesive view of the plant.

The Goal is easy to read and understand which has made it popular with practitioners. Because it puts the reader in the middle of Alex Rogo's chaotic life at UniCo, the reader can see how all of the realities of a plant manager's life affect a person's abilities to make the 'right' decisions. For this reason, it is a very powerful way to help people without production experience to understand these realities. With an appeal like that, it is easy to understand why the third revised edition claims 'Over 3 million copies sold' (Goldratt and Cox 2004).

Because most undergraduate students have not experienced a production environment first-hand, the book has been used in many universities, and the author has used it in an advanced production course for over a decade. Over many readings, the author began to notice more details about the plant's capacity sprinkled throughout the book. A careful look at the details in the book has proved helpful in understanding the changes and their impact in the plant.

The production manager or consultant who takes the time to study the book's problems and solutions can gain a much better appreciation for the benefit gained from each of the small improvements, and how large of a cumulative effect how many small improvements can be made.

In a few cases, the data do not quite add up. Because this book has become perhaps the best-known book in a field dedicated to continuous improvement, it should have consistent technical details. Because the subtitle of the book is *A process of ongoing improvement*, hopefully the few issues presented here can be corrected to further strengthen the book and its message.

After a disclaimer about what this paper will not cover, a brief summary of the book is given, with a discussion of changes made to the book over its various editions. Before beginning a detailed discussion of events in the book, some ambiguities in the timeline of events in the book are clarified. The performance of the plant, and estimates of its capacity, are then considered, before focusing on the capacity improvements made on the two main bottlenecks in the plant, the plant's inventory consumption, and finally lead time improvements. Unless otherwise noted, all page citations are from the third revised edition.

1.1 *What this paper is not*

A lot of time has passed since the book was written. The characters in the book do not use an ERP system, texting, instant messaging, cell phones, pagers, email, voice mail, or even answering machines. This paper is not concerned with such technological dressings. It would be easy to incorporate these technologies into the book, but by the same token, their absence does not in any way diminish the validity of the book's conclusions, nor does it reduce the book's effectiveness in expressing them. The same could be said of other changes in society, like the fact that today no one would consider smoking a cigar in a conference room (p. 34), or that a person cannot get anywhere near the door of an airplane unless they have a boarding pass for the flight (pp. 32–33). If anything, they remind the reader of how long ago the book was written, and how timeless many of its lessons are.

If the book were set in the first decade of the twenty-first century, the characters would likely find management more willing to embark on a process of ongoing improvement. Management at many companies is much more willing to change than it was 20 or more years ago. Many of the other ideas presented in the book are far more widely held today, at least in part, because of the ubiquity of the book.

There are some inconsistencies in the book that could affect the present analysis, such as the fact that the number of heat treat furnaces seems to vary from being given as two ovens (pp. 146, 147, 192) or one (pp. 147, 148).

There are other inconsistencies that are not important, like the fact that the division controller's name changes from Ethan Frost (pp. 25, 55, 249) to Nathan Frost (pp. 261, 262, 268, 319). The hero doesn't even know how old his plant is. One afternoon, he says 'I happen to know the plant is only about fifteen years old,' (p. 37), but something doesn't quite add up when he says that Lou the controller came to the plant 'from corporate about twenty years ago', (p. 44) and Bob Donovan the production supervisor says: 'Hell, I've been at this plant for more than twenty years' (p. 143). These mistakes may be confusing, but do not lead to confusion about operational details.

This paper will focus on the production details of the book. However, this will not be an attempt to determine whether Goldratt's drum-buffer-rope methods are superior to other methods or whether the methods presented are the ideal way to run a plant or not.

1.2 Impact of The Goal

The drum-buffer-rope approach and Goldratt's 'theory of constraints' approach which both appear in *The Goal* have inspired a fair amount of discussion within the operations research literature, starting 20 years ago (Goldratt 1988), and continued ever since, as Watson *et al.* (2007) and Kouvelis *et al.* (2005) note. Watson *et al.* (2007) provide an overview of much of the related literature. Mabin and Balderstone (2000) survey the TOC-related literature, which in their 2003 paper had grown to over 400 articles and papers. The 2003 paper also analyses the results of 81 published case studies.

Spearman (1997) described a number of benefits and shortcomings of the approach. Atwater and Chakravorty (2002) considered a question raised by Spearman (1997), namely the proper level of utilisation for the bottleneck. Silver *et al.* (1998) list some of the other papers that compare the production scheduling approach presented in *The Goal* with other, similar approaches. The ideas in *The Goal* and its theory of constraints (TOC) have been widely discussed in industry so that they are often discussed in other textbooks such as Jacobs *et al.* (2009), Hopp and Spearman (2001), and Vollmann *et al.* (2005).

The academic interest in DBR is somewhat a reflection of the interest in *The Goal* within the business community and its popular press. *The Economist* (1995) wrote: 'The most successful attempt at management-as-fiction is still *The Goal*.' At that point, the book had been translated into 18 languages. *The Economist* also reports that a study of British managers found that although they bought books by many management authors, *The Goal* was the only one they had read cover-to-cover. Writing in *Fortune*, Belis (1994) writes: '*The Goal* offers a far more palatable way to absorb such simple but important home truths than most of the nonfiction management manuals out there.'

Traditional accounting practices are cited as a driver of bad decision making in the book, and a conversation about the relationship between accounting and production decisions has ensued (e.g. Balderstone and Keef 1999, Eden and Ronen 2007).

Goldratt (1997) has also turned his TOC, bottleneck-focused approach to project management which has been studied in that literature (Herroelen and Leus 2007).

2. Plot summary and discussion

For the reader whose knowledge of the book needs refreshing, a brief synopsis follows, from the third revised edition.

2.1 *Brief summary of the book*

Alex Rogo is the manager of the UniCo plant, part of UniWare. His division's Vice President, Bill Peach, tells him that because his plant is losing so much money and has so many late shipments, it will be shut down in three months, unless Alex can dramatically turn things around. Later that same week, Alex's wife Julie leaves him, in large part because he has been spending so many long hours at work.

Alex remembers a chance airport encounter with Jonah, his old college physics professor who has become a manufacturing consultant, and enlists his aid. With Jonah's guidance, Alex and his staff discover the plant's bottlenecks, figure out how to increase the bottlenecks' capacity, and ultimately, how to use the bottlenecks' production schedule to control the release of material into the plant. With a newfound emphasis on shipping late orders first, they clear their backlog, and go in search of new business. Thanks to Jonah's guidance, the plant is such a roaring success, overall, that the plant is saved, and Alex is promoted to run the whole division. He also discovers family life again and saves his marriage (but it's not clear how critical Jonah was to that success).

All of the in-plant action of the book centres on the two bottleneck work stations in the plant. Through Jonah's prodding, Alex and his staff discover that their most expensive numerically controlled (n/c) machine, the NCX-10 (which Alex describes as 'a sexy-looking machine' p. 14) is one bottleneck, and that heat-treat ('It's dirty. It's hot. It's ugly. It's dull', p. 146) is the other. They spend the rest of the book figuring out how to increase the productivity of these bottlenecks, and how to allow the bottlenecks to dictate the pace of the rest of the plant. These two work centres, and their throughput rates, are by far, the most important things in the plant, and most of the current analysis will involve these two work centres.

2.2 *Timeline continuity issues*

The book does not carefully signal the timing of events of the story which, for the reader, can be confusing. For example, at the end of a month, the plant has set a record for shipments, but it's not immediately clear to the reader how much time has passed, to know whether it was the first or second month of the plant's probation. Because a lot of action can happen in a short amount of time in some parts of the book, a lot of pages may be required to describe a short length of time, and the reader can be left confused about the timing of events. To take a detailed look at the book, as will be done here, it is important to be able to know fairly accurately when things happened. After some careful reading, the following timeline can be constructed.

The first 16 chapters take place in what must surely be one of the worst weeks, both professionally and personally, a person could imagine: the plant Alex runs is threatened

with closure if it doesn't improve in three months, and Alex's wife, Julie, leaves him. Most three-month periods of the calendar would last 13 weeks, but for simplicity assuming the months are four-weeks for a 12 week deadline seems reasonable.

Chapter 17 begins the next Monday (week 2), which continues into chapter 18, where the second Tuesday starts, and Alex's staff begins trying to identify the bottlenecks. 'A few days later' (p. 141) database manager Ralph Nakamura admits that he can't find the bottlenecks using the MRP routings. Much later, Ralph says that it took him four days to realise that he couldn't find the bottlenecks (p. 280). If he started Tuesday and tried for four days, that would mean he admitted defeat on Friday, or Monday of the next week. It seems to make more sense for this to be Monday of week 3, as will be shown.

The day Ralph admits defeat, Alex tells his team they've 'only got ten weeks now to make something happen'. Assuming they had 12 weeks to begin with, that is true. But chapter 19 continues 'that evening' when he picks up Jonah at the airport and tells him: 'We've only got two months left' (p. 150). Two weeks apparently evaporated in one afternoon; however, it seems more consistent to assume this is really still Monday of week 3. Chapter 20 starts the morning after Jonah's first visit, apparently Tuesday of week 3. Chapter 21 starts that night and the next morning, which must be Wednesday. If the 'few days' above ended with a Friday, this day would be Sunday, and they would have worked through the weekend, which was not mentioned, and therefore seems unlikely.

In the 'early afternoon', on the Wednesday (of week 3) after Jonah's visit (p. 173), Alex discovers that the NCX-10 is idle. 'A few days pass' (p. 176) while they develop the red and green tag system (red tags indicate parts destined for the bottlenecks, which are to receive higher priority than green tags). 'But at eight o'clock on Friday morning', Alex is in the cafeteria explaining the new system to the employees. Apparently, Thursday lasts 'a few days' at the UniCo plant.

Chapter 22 is Monday, and 'a week has passed' since the tag system went in place, so it is Monday of week 5, and chapter 23 is apparently Wednesday morning of week 6. Chapter 24 fast forwards to the Friday at the end of the month, which must be the second probation month and therefore week 8, and continues to Monday of week 9. Chapter 25 begins with Jonah arriving on Tuesday morning, and chapter 26 continues to Wednesday of week 9. Chapter 27 says: 'May has ended.' 'Two months ago' the threat came, and they have a month to go, so apparently Peach's threat came at the end of March or beginning of April.

The third month passes in chapters 27 to 30. In chapter 28, they halved the batch sizes on non-bottlenecks at the very start of month 3, and went to look for more sales. Chapter 29 takes place on Monday, 'two weeks' from the end of the month, which must be week 11. In this chapter, they halved batch sizes again and accepted a huge order from Bucky Burnside, and the first of four shipments were to ship in 2 weeks, week 13. Chapter 30 jumps two weeks to the end of May, and the three months are up. The first Burnside order has shipped, on time, in week 13. 'The following week', week 14, some creative accounting they used was discovered (p. 248). Two days later, an audit team arrives. 'For a week' Alex waits for fallout from headquarters, which puts the time at roughly the end of week 15. All of Burnside's orders have shipped, even though the last one shouldn't have been able to ship until week 16.

In chapter 31, Alex learns that the plant is safe, and he has two months to prepare for his big promotion, time which passes quickly, while he enlists his team's help in preparing for the next stage of his career.

2.3 Revisions of The Goal

The Goal: Excellence in Manufacturing first appeared in 1984, in a first edition of 262 pages, compared with a 337-page story in the third revised edition. The story is identical to the current version, except that it ends with the point where Alex has been promoted. In the original version, a new division is going to be created, which Alex will head. When Alex asks Jonah for help in figuring out how to run a division, Jonah replies, in the last words of the book: ‘Think about what the goal should be’ (1984, p. 262).

The revised edition from 1986 has a different subtitle: *The Goal: A Process of Ongoing Improvement*. It adds an additional introduction, and has an updated ‘About the author’ section, but the story is completely unchanged for the first 262 pages. A 12-page Epilogue is added, which takes place ‘a couple of years’ later (1986, p. 264). In it, Alex visits with his college buddy Eric on a flight back from Japan, where Alex has been meeting with his sales staff. Eric is impressed that UniCo is selling in Japan, and Alex tells him about what he has learned. Eric can’t wait to get home and start implementing all of the changes Alex made. But Alex tells him not to, that the important thing is to focus on the actual process of ‘ONGOING improvement’ (capitalisation in the original; 1986, p. 269).

The second revised edition from 1992 has been re-typeset, and as a result, p. 261 of this edition is equal to p. 259 of the previous two editions, and it is at this point the second revised edition begins to differ from the previous. The previous ending of the main part of the story is revised, and Jonah still refuses to give Alex all of the answers, but agrees to help Alex with the process of figuring out how he will run the division.

Alex approaches his co-workers about coming to the division office with him, and enlists their aid in figuring out how they exactly achieved their successes, so he can try to have more success in the future, and they agree to start the division ‘on a process of ongoing improvement’. They develop the five-step process of ongoing improvement, known elsewhere as the ‘theory of constraints’.

In that process, they discover inventory manager Stacey Potazenek has been building an inventory of parts going through the bottleneck (BN) processes by releasing fictitious orders. She didn’t want to lose time on the bottleneck work centres, but the capacity of those processes has been increased so much that they are no longer the system’s constraints; customer sales are. The fictitious orders were 20% of total capacity, so they went in search of more orders. Unfortunately, this strained the non-BN processes so much they couldn’t readily replenish the pre-BN inventory buffers, and the BNs starved, and they could not meet the delivery due dates.

In the third revised edition, the story is unchanged, but 46 pages of new material is added after the story, in which Goldratt discusses ways that *The Goal* and TOC have helped organisations in diverse areas, and includes interviews with successful implementations from manufacturing, services industries, and even education.

3. Plant measurements and goals

In order to save the plant, Alex must meet the improvement goals set for him by Bill Peach. Unfortunately, for the reader, there is some confusion about what the target goals are.

3.1 Plant past performance and costs

Bill Peach threatened the plant with closure because it 'is losing money' (p. 5). The plant was 'a loser' when Alex started six months ago, and Alex admitted 'it's gotten worse instead of better' since he came (p. 6). The plant has added new robots fairly recently, apparently within the last year (p. 68), and since that time, the plant has had a harder time shipping orders to customers on time (p. 29) as overdue orders grew rapidly in the last nine months (p. 69), and inventories have increased significantly (pp. 29, 67). At the end of the previous third quarter, six months ago, they discovered the robots were not being used as heavily as they needed to be to justify their purchase, so a lot more material was sent to them (p. 70), which is ultimately discovered to be a big part of the plant's problem.

Aside from the statement about losing money, how the plant is doing financially is not given, but an estimate can be made. The plant controller says the plant has 'total operating expense' of \$1.6 million per month (p. 158). An estimate of 'total operating expense' may or may not include material costs, depending on the accounting practices of the company. As will be shown, it would make more sense if this number does include materials costs.

The plant controller estimates that the selling price of a typical unit is \$1000 (pp. 156, 158). Johnny Jons, the sales manager, told Alex that Bucky Burnside wanted to buy 'a thousand units,' of the popular Model 12, and says 'a thousand units means a little over a million dollars in sales to us' (p. 242), which would put the sales price at just over \$1000. However, later Johnny says that UniCo sells the Model 12 to Burnside for \$827 (p. 312). Despite Johnny's apparent confusion, \$1000 is given by Lou as the typical sales price of a unit (p. 156).

The materials cost of the popular Model 12 is \$344.07 (p. 312). With monthly expenses of \$1.6m, the plant would have to sell 1600 units per month to break even. If the Model 12 material costs are typical, that would mean material costs would be approximately \$550k. Since the plant was in fact losing money, they must have been shipping less than 1600 units, and material costs may have been lower than this. On the other hand, because the plant was building inventory, the plant may have been bringing in more than \$550,000 worth of material each month.

If 'total operating expense' includes material costs, this would mean that total labour, depreciation and utilities would cost \$1.05 million per month. There are 400 people on the day shift at UniCo (p. 36), and 600 people work at the plant (p. 17). If depreciation and utilities costs were zero, this would mean an average annual salary of nearly \$21,000. In 1984 (the year the book was published), the average hourly wage for a US manufacturing worker producing durable goods was \$9.74 (Statistical Abstract of the United States, US Department of Commerce 1986, p. 412.), which would be an annual salary of \$20,259. The median weekly wage for a machinist was about \$395 (Occupational Outlook Handbook, US Department of Labor 1986, p. 424), which computes to \$20,500 per year. Using this calculation UniCo's wages would be slightly above average. If the estimate of 'total operating cost' did not include materials, and the full \$1.6m was for labour, UniCo would be paying its people 50% above the national average.

Assuming total operating expense includes material costs, if allowance is made for the fact that the plant must have utilities and depreciation (for all of those new robots), its average salaries would come out somewhere below the national average, perhaps significantly so. This makes it easy to believe when Bob Donovan says 'we can't seem to attract anybody good with the wages we offer' (p. 146).

Before the improvements in the plant, the old record for shipments in a month was 31 orders worth 'about two million dollars' (p. 195). If the record shipments for one month were \$2 million, typical shipments must have been somewhere below that, and if the plant costs \$1.6 million to run, it is not difficult to see why Alex's plant has not been profitable lately. Also, the plant has already laid off 600 workers (p. 17), some of them just three months before the book starts (p. 4). Operating expenses must have been much higher than \$1.6 million before, and yet the record for shipments was only \$2 million. It would be very hard to see how the plant could have possibly ever been profitable before the layoffs.

3.2 Improvement required to save the plant

When Alex came to the plant it was losing money, and he was put in charge of it because Bill Peach thought Alex could 'change this plant from a loser to ... well, a small winner at least' (p. 6). Peach said in three months he needs to 'turn this plant around' (p. 6). Exactly what had to be accomplished was not specified. In the second month, the plant achieved profitability and was the best-performing plant in the division. The plant's profits were in fact great enough to offset the losses from the other plants (p. 220). But that was not enough to take away the threat of closure. Bill Peach thought that Alex may have only done so well because of the large backlog he was able to ship, and Bill wanted to see better proof of a long-term improvement, in terms of cost reductions. He told Alex: 'It's going to take a ten or fifteen percent reduction in operating expense to make the plant profitable for the long term' (p. 222). When Alex asked for clarification, Bill shifted the criteria from cost reduction to profit improvement and said 'Just give me fifteen percent more on the bottom line than you did this month' (p. 223).

When the plant increased shipments by increasing production without increasing the workforce, labour costs were unchanged, and only materials costs went up, which gives Alex a marginal profit per unit of \$656, assuming material costs of \$344 (p. 312). When he sells an item from inventory he is selling an asset which was valued at full cost of approximately \$700 (p. 312) and profit per unit is only \$300.

In month 2, the plant increased the production of the bottlenecks and shipped \$3 million in orders. Labour costs were unchanged, so labour, depreciation, and utility costs should remain constant at \$1.05 million (assuming, as argued above, that the \$1.6 million figure included material costs). For 3000 units raw material costs would be \$1.03 million so total expenses would be \$2.08 million, and net profits should have been \$0.92 million on sales of \$3 million. Peach told Alex to increase these profits by 15%, which means to keep the plant open, he would need an additional \$0.135 million. If material costs are \$344 per unit, with a \$1000 sales price, an additional 208 units would be sufficient to reach the target.

However, in the book, in month 3, Alex increases production (and sales) by 1000 units, which should have increased profits by \$656,000, or 70%, more than enough to reach the target. But in the story, the increase will be less than 15%, so the controller suggests using a burden calculation based on the last two months instead of the last year (p. 241). After changing the burden, the profit increase comes out to 17% (p. 247). In the end, the original burden is used and they record a 12.8% increase (p. 250). Since production in month 2 was a record, and month 3 was twice the old plant record, compared with the last year the burden should have been close to cut in half, and a larger increase difference between the two profit calculations would be expected.

On the other hand, it is hard to see how the discussion of burden is relevant to the plant's profit improvement. Burden is used to calculate a profit per unit, but actual monthly profits should be determined by actual expenses and revenues. Using burden to compute a unit cost shows a profit per unit of \$300. Suppose the plant shipped 1500 units in the month before the book began. Using the standard cost, Bill Peach would have said Alex had profits of \$450k for the month. But Bill Peach was clearly using actual revenues and expenses to conclude the plant was losing money. Using actual costs and revenues, the plant clearly should have exceeded the 15% target by a huge amount. Finally, on pp. 271–272, Lou refers to a 15% net profit target. It would seem that may be what was intended, all along.

3.3 *Estimates of plant capacity*

The plant's old record month for production was \$2 million, which represents roughly 2000 units. After increasing bottleneck capacity, the plant ships 57 orders in the second month worth 'in round numbers, we'll call it a cool three million' dollars (p. 195), approximately 3000 units. Apparently, the changes made in the plant have allowed them to ship 50% more units than ever before. Given the massive layoffs the plant has experienced, this new record is especially impressive (perhaps unrealistically so). Unfortunately, Alex is confused about the plant's output, because in week 9, he says that volumes have only increased by 20% (p. 211).

When Alex is considering a large potential order from Bucky Burnside in week 11, it has been calculated that the bottlenecks can turn out 100 units of Model 12 per day, which would give a monthly capacity of 3000 units (p. 243). At the end of the book, Alex says that their efforts on the bottlenecks have allowed them to 'squeeze almost twice as much out of them as before' (p. 334). When they were losing money before the book began, they were producing less than 1600 units per month, so this statement is consistent with a capacity of 3000 units per month.

Additionally, in chapter 37, not long after the three months ended, Alex discovered the inventory control manager was putting extra work into the system to avoid unused bottleneck capacity. This unnecessary work represents 'roughly 20%' of the bottlenecks' capacity (p. 308). To use this capacity, he needs to find more than \$10 million more in sales (p. 309), presumably per year. If \$10 million is 20% of current production, the plant's total capacity has been increased to \$50 million per year, or 50,000 units per year. Since the plant is already producing at this level, it is producing around 4167 units per month, and Alex is hoping to find additional sales of roughly 833 units per month. This output level is approaching three times the output at the beginning of the book.

At that point, Alex is prepared to accept an offer of 10% below cost (p. 311), and he agrees to an offer of \$701 per unit. From the comments of Johnny Jons (the director of sales), it appears that this offer is at or below the standard cost, but there is no way to know exactly how it compares to the standard cost. For the sake of argument, assume it is exactly at cost. The material cost is \$344.07 (p. 312). If the standard cost is \$701, the labour and burden must be \$356.93. Alex explains to Julie that for his plant, burden is calculated as 'about three times labor cost' (p. 240), so labour is one fourth of the labour and burden total, or \$89.23. As mentioned, in 1984 manufacturing wages in the US averaged \$9.74 per hour (US Department of Commerce 1986). This means that a typical product would have 9.2 hours of labour content.

An upper bound on the standard cost has to be the \$827 price Bucky Burnside pays (p. 312). Using that as the standard cost and \$344.07 material costs, the labour and burden would be \$482.93, so labour would be \$120.73, which would mean 12.4 hours of labour per unit. The true value must be somewhere between 9.2 and 12.4.

If all 600 workers in the plant were directly involved in production (which is unlikely), there would be $600 * 40 * 4.3 = 102,857$ production hours available per month (assuming 30 days or 4.3 weeks per month). There are 400 workers on the day shift (p. 36), so there are 200 workers total between the second and third shifts. To potentially underestimate the production hours available, assume 200 workers are in managerial or administrative positions. This gives $400 * 40 * 4.3 = 68,800$ production hours available per month.

If an average product has 9.2 hours of labour content, the plant would be able to produce almost 7500 units per month, assuming 100% availability of 400 production workers. At 12.4 hours per unit, the estimate is 5548. As Alex and his plant realise, workers on the bottlenecks need to be activated 100% of the time, but non-bottleneck workers need to be activated less, so actual capacity is below these estimates.

The old plant record was 2000 units in a month. Before the story starts, the plant must have been producing less than that. By some estimates, the plant capacity is 3000 units per month, by others it is over 4000 units per month. Is it possible that total plant capacity could have been raised by 100%? Below, the capacity increases in the bottleneck processes are considered, to answer that question.

4. Plant capacity increases

4.1 *Capacity increases on both bottlenecks*

The plant undertakes a number of different measures to increase the capacities of their two bottleneck (BN) processes. The following changes apply to both bottlenecks.

- First, they moved quality control inspection (QC) to be in front of the bottlenecks, to make sure the BN capacity was not wasted on parts that would later turn out to be defective.
- They made an administrative decision to only work on parts that were needed for waiting customer orders and not waste BN time making parts to be sold 'someday'.
- Worker schedules were changed so that the BNs would not sit idle during workers' lunch breaks; if a break came during a setup, workers would finish the setup before breaking.
- The BN workers were typically working at several different stations, and when a run finished, the BNs would sit idle, sometimes for hours, until personnel came. Dedicated crews were placed at each station, to make sure no time was lost to worker unavailability.
- To make sure that the work was done as quickly and as well as possible, only the best people were put to work on the BNs (p. 191).
- Red tags were placed on each part that would go through a BN, to make sure the BN parts progressed rapidly to the BNs. This ensured a significant queue in front of the BNs, so they would never run out of material to work on. However, before they realised they needed to restrict the flow of material into the plant by having the BNs control the flow, so many red tags flooded the upstream workstations

that the other parts, with green tags, sat in some cases for weeks before a gap in red tags allowed them to be produced.

- Yellow tags were placed on parts that had gone through the BN, to make sure the capacity would not be wasted by preventable errors downstream from the BNs.

These changes had similar, but not identical, benefits at the two bottlenecks and additional steps were taken at the two stations. Also, lot sizes on non-BN machines were cut in half, not once (p. 233) but twice (p. 244). The plant found that this gave an increase in throughput (p. 244) because otherwise a large batch would sometimes take so long to leave a work station that it would cause a non-bottleneck to wait so long that it would lose enough capacity that it would become a temporary bottleneck (p. 238).

4.2 Capacity of the NCX-10 is a fraction of previous capacity

The NCX-10 does jobs formerly done by a series of three machines. Production supervisor Bob Donovan tells Alex that in one 'typical instance' (p. 145) part process times would be 2 minutes on the first machine, 8 minutes on the second machine, and 4 minutes on the third machine, for a total of 14 minutes of processing time. The NCX-10 takes 10 minutes per part, a saving of 4 minutes per part. But the labour savings would appear to be even greater. The old machines each required an operator to run them continually. The NCX-10 only needs operators who load the machine, and then go do something else, only to come back later. Instead of three full-time operators being dedicated to the process, the




Figure 1. Throughput of old machines.




Figure 2. Throughput of NCX-10.




Figure 3. Throughput of old work centres.

plant was using two workers on a very part-time basis, which represents a significant labour savings.

With the old machines, the first machine produced a part every 2 minutes (cycle time = 2 min), so its throughput (TH) is 30 parts per hour (pph). The second machine had CT (cycle time) = 8 minutes, or 7.5 pph. The third machine had CT = 4 min, so TH = 15 pph. The bottleneck of the three is clearly the second machine, which can produce 7.5 pph, so the three machines together could produce only 7.5 pph.

The NCX-10 has CT = 10 min, so TH = 6 pph. It would appear that capacity on these operations went down by 20% when the NCX was brought in, from 7.5 to 6.0. This seems believable, because at the time, because no one realised how important capacity on these processes was. Once Alex and company discover that the NCX-10 is one of the plant's constraints, buying the NCX-10 looks like a bad decision.

But the NCX-10 actually represented a much larger reduction in capacity. Bob said, 'We had two of the first type, five of the second type, and three of the third type' (p. 145). This means the two part-time workers for the NCX-10 represent an even greater labour savings, compared with not three, but 10 full time workers. However, this also means that the first station had capacity of $2 * 30 \text{ pph} = 60 \text{ pph}$. The second station had capacity of $5 * 7.5 \text{ pph} = 37.5 \text{ pph}$, and the third station could produce $3 * 15 \text{ pph} = 45 \text{ pph}$. So the second station is still the bottleneck, but total output is 37.5 pph. The NCX-10 only does 6 pph! It is only capable of 16% as much production as the older machines. The plant needs more capacity than it has, and less than it had with all of the old machines, but it is very hard to believe that any company would have reduced any work centre's capacity by 84%.

Even if the NCX-10 were to run '24/7' (24 hours a day, 7 days a week), and the old machines only ran one shift per week, the NCX-10 can produce a maximum of $7 * 24 * 6 = 1008$ parts per week (ppw), (which is consistent with the estimate when considering the Burnside order) where the old machines produced $37.5 * 40 = 1500$ ppw. In this case, the new machine reduced capacity by only 33%, still quite significantly.

As Alex discovers, the NCX-10 is not running 24/7. A lot of time is lost to breaks and other losses. Its uptime is 585 hours per month (p. 153), out of the possible $24 * 30 = 720$ hours in a typical month. But if it is assumed that the three old machines would have had similar uptime percentages, the fact would remain unchanged that the NCX-10 represents a capacity loss of one third. If those old machines ran 24/7 with a similar uptime ratio, the capacity lost would be the huge 84% seen above. The old machines may have a slightly higher defect rate (p. 184), but it's hard to imagine that even the most obtuse company would reduce capacity of any process as drastically as this plant did.

4.3 *Re-activated old machines have lost capacity*

When Jonah asks if they have the machines that the NCX-10 replaced, Bob says 'Some of them we do' (p. 153). But then he says that they 'got rid of an entire class of machine that we'd need to supplement the NCX-10' (p. 154). The old machines from one of the stations were sold off to make room for more work in process (pp. 154, 184), but Bob gets one of them back, an old 'Zmegma' (pp. 183–184). He also says that combined with a 'Screwmeister' and 'that other machine off in the corner, together they can do all the things the NCX-10 can do' (p. 183). So the plant has its original Screwmeisters and Other Machines (as this paper will call them), and one new-to-them Zmegma.

Because these are the same three machines from before, they have capacities, like before, of 30, 7.5, and 15 pph, per machine, for a combined throughput of 7.5 pph. If they run the three old machines one shift per day, they should realise $8 \text{ hours} * 7.5 \text{ pph} = 60$ additional units per day. The NCX-10 is working 585 hours per month, or an average of 19.5 hours per day (assuming 30 days per month). With a rate of 6 pph, that is 117 per day. The re-activated machines represent an increase of slightly over 50%. Since experienced veteran machinists will be running each machine (p. 190), it would be expected for them to be operating at least at the standard rate, and since they will be dedicated to the machines, they will not be wandering off doing other things and leaving these machines unattended.

Bob estimates running these machines one shift per day will increase capacity 18% (p. 190). An 18% increase over 117 units would be 21 units, a rate of 2.6 pph, roughly 34% of the past capacity. It seems that the machines will not improve capacity nearly as much as would be expected, only improving capacity by one-third of what would be expected.

There are a few possible ways to explain this. First, some small amount of time will be lost to breaks. Second, maybe the machines are old and just don't run as well as they used to. The Zmegma appears to be from 1942 (p. 183). However, the NCX-10 is only 2 years old (p. 145), so the two machines re-introduced into service were being used in this same plant only 2 years ago. It seems reasonable that they would have chosen the most operational machines from those remaining, and any parts that were in poor condition could have been cannibalised from the other machines in the plant, so the machines should be capable of running well.

If the Screwmeister and Other Machine are capable of operating at approximately their old production levels, maybe the Zmegma is the bottleneck machine, and the one they got is just barely operational, capable of only a fraction of the 7.5 pph it should produce. However, since Bob's 'buddy' had 'a couple of these sitting around he'd have no problem parting with' (p. 184), it seems reasonable to expect Bob would have looked them over and chosen the better one, and made sure all of the pieces seemed to be there, maybe cannibalising parts from the other. Since he's been the production manager in the plant for 9 years (p. 8), and worked in the plant for over 20 years (p. 143), he likely knew the machines well in the past, and it's unlikely he'd bring back and suggest the plant start using a machine in terrible condition.

Another way to look at this issue is that if the old machines bring an 18% increase in production, that is 21 units ($117 * 18\%$) per day. If the BN of them produces 7.5 units per hour, when running, that would represent 2.8 hours worth of production per 8 hour shift, which would represent a 35% uptime, which seems highly unlikely, if, as just explained, reasonably soundly-working machines are assumed.

In summary, re-activating the old machines should have represented a 50% increase in capacity for the NCX-10, but somehow, it only increased capacity 18%.

4.4 Capacity increase at the NCX-10

Putting QC in front of the NCX-10 increased the plant's effective capacity. About 5% of the parts coming into the NCX-10 were defective. After the move, 5% fewer parts flow through it. Where they used to produce 100 parts, they now need only produce 95, which means the plant could produce 5.3% more finished goods.

It is not known how many parts that went through the NCX-10 would later become defective, but placing the gold tape on the parts that go through it is assumed to have had some reduction in defects occurring after the NCX-10. If the NCX-10 were in the middle of the process, and if the stages after the NCX-10 had the same defect rate as the steps before it (5%) (p. 180), and no gold parts became defective, the gain from the yellow tags could be as high as 5%. Ralph estimates parts should take 14 days to reach the BN (p. 217), and products are shipping in 2 months (p. 235), so it seems the BNs are earlier than halfway into the production process, and if defects after them are removed, 5% would seem to be an underestimate of the increased output.

In the first chapter, a setup on the NCX-10 took an hour and a half (p. 2). This is not said to be a typical setup time; on the other hand, it is also not said that this was an exceptionally long setup time, so for sake of argument, assume that 1.5 hours is a typical setup time. The machine has typically run 585 hours a month, out of 720 hours per month. In the second month, the plant shipped 57 overdue customer orders (p. 195). It is said that 90% of the overdue orders were waiting for parts from one of the bottlenecks (p. 171). If they all went through the NCX-10, that would be 51 jobs. At 1.5 hours per setup, that would require 76.5 hours, which would leave 58.5 hours for the month for maintenance, or were hours the machine spent idle.

Adding the dedicated crew must certainly have taken away any idle time lost due to a lack of personnel. Before a dedicated crew was put on the NCX-10, it would still frequently sit idle for 20–40 minutes each time a job finished, because the crew needed to tend to it was working at another station (p. 189). If approximately 50 jobs per month pass through the machine, these time losses would represent 25 hours of lost bottleneck time. Because the machine was previously only available 585 hours a month, taking away this idle time represents an effective increase in capacity of just over 4%. If closer to the 58.5 hours of idle time mentioned above were regained, that would be an increase of 10%.

If 4 to 10% more capacity was gained by the dedicated crews, 5% was gained from moving QC, and 5% was gained from the yellow tags, multiplying the benefits (since the impacts are cumulative) the increase should be 15–22% ($1.04 * 1.05 * 1.05$). Adding the old machines added 18% more capacity than what they had before. Together, total NCX-10 capacity has increased 33–40%, which does not appear to be enough to support a doubling of demand.

However, perhaps the plant doesn't need to double NCX-10 output. If it was running 585 hours per month, and if it can produce 6 pph, it should produce 3510 parts per month (ppm), and it appears that each product needs only one NCX-10 part (p. 156). This full capacity was not being realised, because of defective parts being sent through the machine, or its output becoming defective later. If the yellow tags and moving QC drove defects to zero, 3510 good units could be produced in 585 hours. The percentage of jobs waiting for a part from at least one of the bottlenecks is given as 90% (p. 171) or 80% (p. 208). If the long-term average of products waiting for parts from the NCX-10 is 90%, 3600 good parts would be enough to support sales of 4000 units per month (at 80%, production could be 4500 units). By reducing the defective parts after the NCX-10, the plant would almost have enough capacity to satisfy all of its demand.

If 25 to 60 hours of productive time are gained by the dedicated crews, this would represent an additional 150 to 360 units per month. The old machines were stated to have added 18% more capacity, or 632 parts per month (18% of 3,510). Together, these

improvements raised total capacity to 4292 to 4502 units per month, more than enough to support monthly production of 4000 units.

Although not mentioned, the increase in NCX-10 capacity may have been greater. Since the plant appreciates the need to avoid downtime, supervisors would likely make sure setups were performed more quickly. Because only the best workers were put at the BNs, setup time would likely have been reduced, further increasing capacity.

4.5 Capacity increase at heat-treat

Unlike the NCX-10, which requires a constant amount of time for each part, heat-treat (HT) requires a certain amount of time for a batch to stay in the ovens, regardless of the number of pieces in the oven. Because it is not known how large the ovens are, or how many parts will fit in at one time, the number of units they can produce in a period of time cannot be directly computed. However, the percentage increase in capacity can be estimated.

About 7% of parts going to HT were defective (p. 180). Not allowing breaks to interrupt a setup had an impact of unknown size, because they don't know how much capacity was being lost by people taking breaks mid-setup. The red tags increased the size of the buffers at the BNs, so they would not run out of work. The first week that these improvements were in place, the plant shipped 12 orders. Ralph had estimated that 18–20 would be shipped (p. 187), and the reason fewer were shipped was because jobs were being left in the heat-treat ovens much longer than needed, because no personnel were there to take them out (p. 188). If Ralph's estimate of 18–20 orders instead of 12 is correct, then the subsequent decision to put a full-time crew at heat-treat should have increased capacity of HT by 50–67%. A significant impact could have been expected, because Ted Spencer, the HT supervisor, originally had three other centres he was responsible for, in addition to HT (p. 186).

Subsequently, the night shift foreman, Mike Haley, found a way to reduce the setup time. He had his workers stack up parts, ready to be placed onto the tables that roll into the oven as soon as the previous batch came out and cooled down from 1200°. He was also combining multiple parts, to fill up the ovens, instead of only running one part at a time, which was the standard policy (p. 192). These innovations allowed Mike to produce 10% more than the other heat-treat supervisors (p. 191). The conventional wisdom was that 'most of the time these furnaces are running half empty' (p. 146) so unfortunately, these improvements did not lead to that great an improvement as might have been hoped, but they did lead to significant throughput increases, and were disseminated to the other heat treat foremen.

Mike also proposed another way to significantly reduce the changeover time. He suggested building an extra table so that they would not have to wait for the parts to cool down as a part of the setup; the extra table would be stacked up while the other one is in the furnace, and as one came out, the other would be immediately put in. Mike said that it currently 'takes anywhere up to an hour or so to change a furnace load using the crane or doing it by hand. We could cut that down to a couple of minutes' (p. 192) if they built extra tables for each oven. Since parts typically need to be heat treated for 6–8 hours (p. 188), he estimates the tables might save 'a couple of hours a day,' which means 'we can do an extra heat of parts over the course of a week' (p. 193). Using these estimates, a typical heat is 8 hours, with a one hour setup, so in a

168-hour week, each oven can do $168/9 = 18.7$ heats per week. If the setup time were reduced to 10 minutes, total CT would be reduced to 8.167 hours; each oven could do 20.6 heats per week, an increase of 10%.

By adding a full-time crew to HT, they increased throughput by 50 to 67%. Mike Haley's job splitting and setup reductions increased capacity by an additional 10%, each. Since these later increases were cumulative, if it is assumed the full-time crew had an impact of 50%, total capacity has been increased by a factor of $1.5 * 1.1 * 1.1 = 1.815$, or 81.5%. If a 67% increase is assumed from the full-time crew, the increase is 102%, a doubling of capacity. In the short term, to quickly reduce the size of the queue at HT, some work was sent to an outside vendor (pp. 190–191), but this was not going to be a long-term policy.

Finally, Bob discovered that many of the parts that went through HT only did so because the plant was trying to be more efficient at a non-bottleneck machining centre, and was taking larger bites out of the steel in each pass than required by engineering. The larger bites meant the steel had to be heat treated. Going back to smaller bites meant the parts did not need to be heat treated, and reduced the load on HT by 20% (p. 194). This reduction was on top of the 7% reduction gained from moving QC. Together, these mean that the flow of parts coming into HT has been reduced by $1 - (0.93 * 0.8) = 25.6\%$ from the previous levels. For every 100 parts previously coming through HT, the flow has been reduced to 74.4, and the capacity increased to 182–202. The result is that the plant can now take orders for 145–272% more units than before. This is well beyond the 100% increase that would be needed to increase production from 2,000 units per month to 4,000. Thanks to these improvements, they should easily have enough capacity going into the future.

5. Inventory usage

As a result of the increases in throughput and sales, inventories in the plant are drastically reduced, and very quickly; as it turns out, in some cases, a little too quickly.

5.1 Bottleneck buffer stock

As a part of the plant's efforts to increase sales, a much shorter lead time was promised to its customers. Just before the book began, the plant was promising a four month lead time and taking six months to deliver (p. 235). Later, they are willing to reduce the lead time to four weeks (p. 234), and then two weeks (p. 310). In order to reduce lead times this significantly, the plant needed to significantly reduce the amount of work in process (WIP) in the plant.

When Alex discovered that the NCX-10 was a bottleneck (at the start of week 3) it had a huge pile of inventory in front of it, which was estimated to hold 'weeks' of inventory (p. 144). After all of their improvement efforts including the red and green tags, overall work in process was reduced by 12% (p. 195). However, at the bottlenecks, 'the piles of inventory in front of them have grown' (p. 180) and are 'even bigger than before' (at the start of week 5). Alex describes the inventory pile as 'stacked as high as the biggest forklift can reach. It's not just a mountain, but a mountain with many peaks' (p. 204). The buffers are estimated to be 'at least a month or more' (p. 205), and finally estimated to be 5 or 6 weeks worth (p. 216) in week 8.

In the same conversation that the 5–6 week conversation is made, Jonah helps Ralph discover how he could schedule the release of all of the materials so that the bottlenecks are driving the production plans for the plant. Ralph has estimated that it takes parts two weeks to reach the bottleneck queues. He is so confident of this estimate that he believes holding three days' worth of inventory in front of the BNs will be enough to protect them against any variability in production by upstream processes (p. 217).

Only a couple of days later, Alex tells his staff that he is going to promise a lead time of three weeks to the sales manager Johnny Jons, apparently based on the assumption of parts reaching the bottlenecks in two weeks, waiting in the buffer 3 days, and reaching final assembly and being ready to ship in 4 days. He tells Johnny Jons he can deliver in 4 weeks, and Jons tells the sales people to quote a lead time of 6 weeks. Unfortunately there is no way that in two days the BNs could have worked their 5 or 6 week queues down to three days. If they stopped releasing any new orders into the plant, (instead of actively getting more orders), it would have taken at least 6 weeks to produce the first order with a 3- or 4-week lead time. At heat treat, they earlier sent some of the backlog to an outside vendor, and maybe they did that again here (although it's not mentioned), but at the NCX-10, there was no such option.

By the end of the three months, the time it takes parts to reach the BN buffers has been reduced to one week from two (p. 324), and they have reduced total work in process by 50% (p. 272). Because the largest piles of WIP were in front of the bottlenecks, if the buffers in front of them have been slashed from more than a month to a few days, it is easy to see how orders could flow through the plant in two weeks. However, with such drastic reductions on their largest piles of WIP, it would seem that WIP would have been reduced by more than 50%.

5.2 Plant finished goods consumption

It's not clear how much finished goods inventory (FGI) the plant has at the beginning of the story. Early in the book, Alex says the plant has \$20 million in finished goods inventory (p. 40). However, during Jonah's second visit, Alex thinks to himself that there are 1500 units in stock, which would give them \$1.5 million in inventory (p. 209). Of that, they are lucky to sell 10 a month, which would be \$10,000 (p. 209), which means that their finished goods inventory is indeed very slow-moving. All of the competitive products ship as soon as they're done, so the FGI is made up of non-competitive products. It is hard to see how this inventory could quickly be drained down, even with a special effort by marketing, and no such effort is made.

When the inventory control manager realised that BN capacity had increased to the point that it exceeded market demand, she built inventory with the excess 20% of the BN capacity, which represented over 800 units a month. She built a supply that reflected their customer demands, which the plant should not have any difficulty in selling. At the end of month 3 she revealed that she had been doing this. In month 2, the plant shipped the last of its back-ordered items and started producing new incoming orders as they came. This means that her fictitious orders could have been part of the production plan for at most a month, and at most could have represented \$0.8 million worth of easily sold inventory. However, she has 'roughly six weeks' of FGI, produced in 1 month, using 20% of the plant's capacity (p. 306), which seems impossible.

It is hard to imagine how any significant reductions could have been made in the inventories of obsolete or uncompetitive products, although that is what she appears to claim (p. 306). However, the plant was able to significantly reduce its levels of finished goods inventories. At the end of the third month, Alex says 'inventories are about 40% of what they were three months ago,' (p. 247), which would mean total WIP plus FGI has been reduced by 60%. However, a few days later, the controller says that 50% of the WIP has been removed from the plant and 25% of FGI (p. 272). Clearly, both estimates cannot be correct.

A 60% reduction of \$20 million in finished goods would mean \$12 million in shipments out of inventory. In the first month, sales were below \$2 million, in month two they were \$3 million, and in month three they appear to have been \$4 million, for total sales of \$9 million. Clearly \$12 million in shipments from inventory cannot have happened. Assuming the other extremes, a 25% reduction from \$1.5 million in finished goods would be \$375,000, which seems much more likely. However, in light of the \$0.8 million inventory build caused by Stacey, for a net 25% reduction from the original \$1.5 million, they would have to be down to \$1.125 million, and would have to have shipped \$1.175 million out of finished goods inventory, which does not seem possible, given that it appears that each month capacity was sufficient to meet demand.

The conclusion that can be drawn is that because the plant increased production capacity significantly, it does not seem possible that it could have reduced finished goods inventory as much as is claimed.

6. Summary and conclusions

The Goal is written as a novel, and because it's not written like a case study or a textbook, most readers are unlikely to be unaware that hiding beneath the surface is, in fact, a quantitative case study that stands up well to a detailed scrutiny. There are some places where the details in the story do not add up, most notably the increase in NCX-10 capacity from re-activating the old machines, the plant's failure to meet the required profit improvements, the quick consumption of buffer stock, and the reduction of finished goods.

However, as has been shown, the improvement efforts undertaken by Alex and his staff can easily generate sufficient capacity gains to support the dramatic increase in sales that Johnny Jons finds for them.

Although much of the discussion in the literature after the book's publication has focused on the drum-buffer-rope mechanism, the book presents a dramatic example of the cumulative results of continuous process improvement: QC before the BN, reduced BN downtime through faster setups and dedicated crews, splitting lots on the BN, reducing the volume that has to go through the BN, finding additional capacity, and most importantly, keeping the BN stocked with a buffer of inventory before it. Many small actions that each gave a modest increase in effective capacity, when put together, allowed the plant to double its capacity.

The Goal has been a hugely popular book, and at a time when many people did not think that operations were important, it helped non-operations personnel understand the difficulties of production, and appreciate the good that can be done by a continued focus on process improvement.

By collecting and quantifying the impact of the changes, the current paper should provide the book's readers with a much more concrete understanding of the changes in the plant, to help them appreciate what can be accomplished from 'a process of ongoing improvement.'

References

- Atwater, J. and Chakravorty, S., 2002. A study of the utilisation of capacity constrained resources in drum-buffer-rope systems. *Production and Operations Management*, 11 (2), 259–273.
- Balderstone, S. and Keef, S., 1999. Exploding and urban myth. *Management Accounting: Magazine for Chartered Management Accountants*, 77 (9), 26–27.
- Belis, G., 1994. Peeking into your rival's briefcase. *Fortune*, 18 March, 129 (8), 179–180.
- Dani, S., 2006. Book review. The goal: a process of ongoing improvement. *International Journal of Production Research*, 44 (2), 419–420.
- The Economist*, 1995. Tales out of business school. 21 January, 334, 63.
- Eden, Y. and Ronen, B., 2007. *Approximately right, not precisely wrong*. Great Barrington, MA: North River Press.
- Goldratt, E., 1988. Computerised shop floor scheduling. *International Journal of Production Research*, 26 (3), 443–455.
- Goldratt, E., 1997. *Critical Chain: a Business Novel*. Great Barrington, MA: North River Press.
- Goldratt, E. and Cox, J., 1984. *The goal*. Croton-on-Hudson, NY: North River Press.
- Goldratt, E. and Cox, J., 1986. *The goal*, revised ed. Croton-on-Hudson, NY: North River Press.
- Goldratt, E. and Cox, J., 1992. *The goal*, 2nd revised ed. Great Barrington, MA: North River Press.
- Goldratt, E. and Cox, J., 2004. *The goal*, 3rd revised ed. Great Barrington, MA: North River Press.
- Goldratt, E.M. and Fox, R.J., 1986. *The race*. Great Barrington, MA: North River Press.
- Herroelen, W. and Leus, R., 2007. On the merits and pitfalls of critical chain scheduling. *Journal of Operations Management*, 19 (5), 559–577.
- Hopp, W. and Spearman, M., 2001. *Factory physics*, 2nd ed. Burr Ridge, IL: McGraw-Hill Irwin.
- Jacobs, F., Chase, R., and Aquilano, N., 2009. *Operations and supply management*, 12th ed. Burr Ridge, IL: McGraw-Hill Irwin.
- Jodlbauer, H. and Huber, A., 2008. Service-level performance of MRP, kanban, CONWIP, and DBR due to parameter stability and environmental robustness. *International Journal of Production Research*, 8 (15), 2179–2195.
- Kouvelis, P., Chambers, C., and Yu, D., 2005. Manufacturing operations manuscripts published in the first 52 issues of POM: reviews, trends, and opportunities. *Production and Operations Management*, 14 (4), 450–567.
- Mabin, V. and Balderstone, S., 2000. *The world of the theory of constraints: a review of the international literature*. Boca Raton: St. Lucie Press.
- Mabin, V. and Balderstone, S., 2003. The performance of the theory of constraints methodology: analysis and discussion of successful TOC applications. *International Journal of Operations and Production Management*, 23 (6), 568–595.
- Rand, G., 1986. The goal (book). *European Journal of Operational Research*, 26 (3), 412–413.
- Silver, E., Pyke, D., and Peterson, R., 1998. *Inventory management and production planning and scheduling*, 3rd ed. New York: John Wiley & Sons.
- Spearman, M., 1997. On the theory of constraints and the goal system. *Production and Operations Management*, 6 (1), 28–33.
- US Department of Commerce, 1986. *1985 Statistical Abstract of the United States*, 106th ed. Washington, DC: US Bureau of the Census.
- US Department of Labor, 1986. *Occupational Outlook Handbook. 1986–1987 Edition*. Washington, DC: US Department of Labor.

- Vollmann, T., *et al.*, 2001. *Manufacturing planning and control for supply chain management*, 5th ed. Burr Ridge, IL: McGraw-Hill Irwin.
- Watson, K., Blackstone, J., and Gardiner, S., 2007. The evolution of a management philosophy: the theory of constraints. *Journal of Operations Management*, 25 (2), 387–402.
- Watson, K. and Patti, A., 2008. A comparison of JIT and TOC buffering philosophies on system performance with unplanned machine downtime. *International Journal of Production Research*, 46 (7), 1869–1885.
- Whitford, D., 2004. An interview with Eli Goldratt and others, *In*: E. Goldratt and J. Cox., eds. *The goal*, 3rd revised ed. Great Barrington, MA: North River Press.
- Ye, T. and Han, W., 2008. Determination of buffer sizes for drum-buffer-rope (DBR)-controlled production systems. *International Journal of Production Research*, 46 (10), 2827–2844.

Copyright of International Journal of Production Research is the property of Taylor & Francis Ltd and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.