

**Simplified Drum-Buffer-Rope**  
**A Whole System Approach to High Velocity Manufacturing**  
by Eli Schragenheim and H. William Dettmer



## Introduction

Drum-Buffer-Rope (DBR) is the Theory of Constraints (TOC) production planning methodology originated by Eliyahu M. Goldratt in the 1980s. In fact, the concepts of DBR actually preceded the Five-Focusing-Steps and the notion of the "throughput world" in the development of the TOC paradigm. While the DBR method is much simpler than the older Optimized Production Technology (OPT) algorithm and the recent Advanced Planning and Scheduling (APS) systems, for many production environments, especially those not currently— or consistently— dominated by an active internal bottleneck, an even simpler method can be adopted. We call this method S-DBR, to distinguish it from the traditional model, which we'll refer to as traditional DBR.

S-DBR is based on the same concepts as traditional DBR and is certainly in harmony with TOC and the Five Focusing Steps. What distinguishes it from traditional DBR is its assumption of market demand as the major system constraint, even when an internal capacity constraint temporarily emerges. S-DBR can be easily supported by traditional ERP/MRP systems and it is specifically intended to deal with fluctuating market demand.

In this paper we'll briefly review the basics of traditional DBR, then introduce evolutionary S-DBR model. Then we'll describe the similarities and differences between the two approaches, and we'll point out the circumstances under which the traditional DBR model would still be preferred.

In the presentation at the Constraints Management Special Interest Group (CM-SIG) technical conference, we'll demonstrate the effectiveness of S-DBR using the Management Interactive Case Study Simulator (MICSS), a software package designed to simulate an integrated "whole business" system. For those readers unable to attend the conference, the same business scenario, along with a demonstration version of the MICSS software, will be included on a compact disk that will accompany our forthcoming book on the same subject, due for publication in September 2000.<sup>‡</sup>

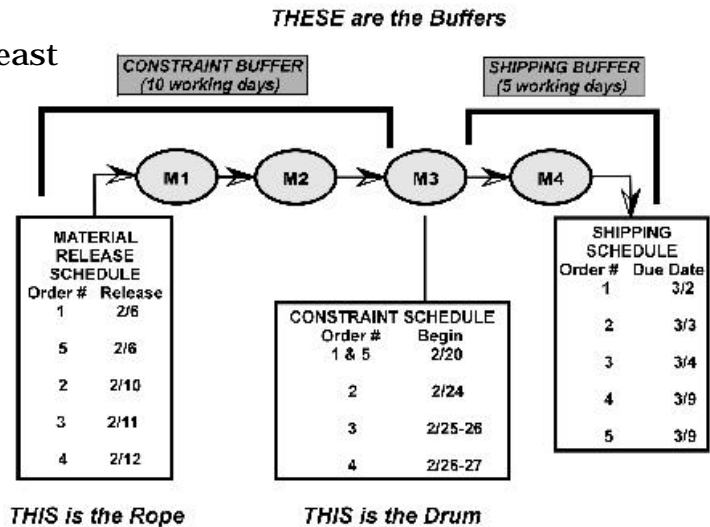
## Traditional Drum-Buffer-Rope (DBR)

The traditional DBR model is designed to regulate the flow of work-in-process (WIP) through a production line at or near the full capacity of the most restricted resource in the manufacturing chain. To achieve this optimum flow, the entry of work

---

<sup>‡</sup> Schragenheim, Eli, and H. William Dettmer. *Manufacturing at Warp Speed: Optimizing Supply Chain Business Performance*. Boca Raton, FL: St. Lucie Press, 2000.

orders into production is synchronized with the current production rate of the least capable part of the process, referred to as the capacity-constrained resource (CCR). The production rate of this CCR is typically likened to the rhythm of a *drum*, and it provides the pace for the rest of the system. The *rope* is essentially a communication device that connects the CCR to the material release point and ensures that raw material is not inserted into the production process at a rate faster than the CCR can accommodate it. The purpose of the rope is to protect the CCR from being swamped with work-in-process. To protect the CCR from being "starved" for productive work to do, a time *buffer* is created to ensure that work-in-process arrives at the CCR well before it is scheduled to be processed. Figure 1 shows a graphic representation of this model. Other sources describe the traditional DBR model in more detail.<sup>1,2,3,4,5</sup>



**Figure 1. DBR - The "Big Picture"**

An underlying principle of TOC is that manufacturing to firm orders with defined due dates is the most desirable situation possible, and preferable to manufacturing to stock. To that end, applying traditional DBR starts with some desired master production schedule (MPS) that includes firm customer orders with delivery due dates. Next, the existence of an internal physical resource constraint is verified. The identification of such a constraint (CCR) can be supported by computerized capacity analysis, but should be validated by production management. Immediately, there are two distinct possibilities:

1. There *is* no capacity constraint currently active, or
2. A definite capacity constraint is identified.

In the relatively few cases in which DBR-specific software is used, there is a discrete procedure to determine which of the two situations apply. In situations where production managers aren't certain whether a specific resource is really an active constraint, Dr. Goldratt has recommended assuming the first possibility (no capacity constraint is active) and allowing the control mechanism (buffer management) to make the case that this isn't true and point instead to the capacity-constrained resource (CCR). Buffer management helps to identify an emerging capacity constraint in two ways. First, it indicates a significant increase in the number of instances in which expediting is needed ("holes" in Zone-1); second, it points to the specific resource where most of the "late" parts are likely to be found. The majority of non-computerized DBR implementations usually assume that a certain resource is the CCR, and managers behave accordingly.

More than one CCR might be active. Dr. Goldratt addressed this possibility in *The Haystack Syndrome*<sup>6</sup> and in the DBR software package, known as "Disaster," developed by the Goldratt Institute in the late 1980s. Management of situations such as this is only possible with the support of computerized DBR software, and even this is not really recommended. The preferred approach is not to allow the emergence of more than one CCR in the first place, either by adding capacity or by imposing restrictions on market demand.

The Drum-Buffer-Rope method strives to achieve the following:

1. Very reliable due-date performance
2. Effective exploitation of the constraint
3. As short response time as possible, within the limitations imposed by the constraint(s).

Conceptually, the three main elements of DBR are: 1) the plan for exploiting the capacity constraint (the "drum"); 2) protection against "Murphy" (the "buffer," expressed in time rather than in things that are stocked somewhere; and 3) a material release schedule (the "rope") that protects the shop floor from excess WIP and priority confusion. DBR assumes that true material constraints are very rare and proper inventory management should ensure material availability as required.

### **DBR When No CCR is Active**

When no CCR is active, there's no reason why all the firm orders shouldn't be delivered on time. The list of those orders constitutes the "drum", which is really the master production schedule (MPS).

The "buffer" is a liberal estimate of the lead-time to move raw materials from their release point through the entire production process to the shipping dock. This estimate should consider the impact of Murphy and allow enough extra time to be reasonably certain that firm orders will be delivered on time. This buffer is called the shipping buffer, because it protects the timely shipping to customers. No CCR buffer is needed, because there is no CCR.

Once a shipping buffer is determined, a material release schedule (MRS) is generated. This constitutes the aforementioned "rope." It precludes any release of materials prior to the date specified in the MRS, ensuring that any WIP is released no sooner than buffer-time from its delivery-due date.

TOC requires full subordination of all other parts of the system to the system's constraint. Since in this situation market demand (lack of enough customer orders to fill available capacity) is the sole constraint, traditional DBR assumes that the flow of WIP to become finished goods, ready for shipping, will be as smooth and fast flow as possible.

### **DBR When A CCR is Active**

When a CCR is confirmed to exist, a finite capacity schedule is generated for it, based on the preliminary master production schedule (MPS). The MPS is subsequently revised, based on the limitations imposed by the CCR. The new MPS and the detailed schedule for the CCR constitute the Drum.

In this situation, three buffers are established as a protection mechanism against variability ("Murphy"):

1. *A shipping buffer.* This is a liberal estimation of the lead time from the CCR to the completion of the order *OR* the lead time from raw materials to completion.
2. *A CCR buffer.* The CCR buffer is a liberal estimation of the lead time from raw material release to the site of the CCR.
3. *An assembly buffer.* This is a liberal estimation of the lead time from the release of raw materials to a process step where parts that don't use the CCR are assembled with parts that do.

The rope is the schedule for release of materials as dictated by the three buffers. Figure illustrates the DBR concept without the assembly buffer. The three schedules are the typical output of DBR planning. The rest of the resources are not specifically scheduled. They are directed to process any order arriving to their site as fast as possible. The rope ensures that no order is released to the manufacturing floor until the CCR or shipping buffer times.

Successful application of DBR depends on actively managing buffers as a means of control. Buffer management allows managers to monitor the state of the buffers and warns when the buffers are in danger of exhaustion. An effective, robust schedule results when DBR's fairly simple planning, which considers the impact of uncertainty, is combined with focused control on the state of the buffers.

### **A Simplified Drum-Buffer-Rope (S-DBR)**

To apply S-DBR, we begin with the presumption that the company is not currently constrained by any internal resource. In other words, the market is most often the overarching constraint for most companies. Is this a realistic assumption? We believe that in most cases it is. The reasons for this thinking will become clearer momentarily.

When the market is clearly the constraint, the combination of the simplicity of DBR planning and the highly focused control afforded by buffer management results in full subordination of operations to sales (the constraint). However, when a CCR begins to emerge the following significant changes are observed:

1. The decreasing capacity of the internal resource constraint may limit the company's ability to respond to the market. Some orders may not be delivered on the required dates. To keep this condition from deteriorating even further, either some of the market demand must be reduced, or capacity must somehow be increased.

2. The actual lead time from raw material release to order completion and shipping increases significantly.
3. Every unit of product needs to pass through two buffers covering various non-constraint operations (assembly and shipping) rather than just one.
4. Buffer management now includes three buffers, each of which must be monitored and managed. This can create conflicts when a single resource must expedite different orders for different buffers. This situation raises the question of the relative importance of the buffers— which is the most critical? And does the relative importance change as the situation changes?

The basic assumption underlying S-DBR is that the market demand is a major constraint even when internal capacity is constrained. The rationale behind this assumption is that if we don't fully satisfy the requirements of the market, our future market demand will decrease. In other words, how well we satisfy our customers today has a direct impact on whether we'll see them in the future. This is the first assumption behind DBR.


*Basic Assumption No. 1:* The market dictates certain requirements that a company must meet. Otherwise, demand for the company's product or service will diminish, and perhaps vanish completely in the future. These requirements imposed by the market sometimes conflict with full exploitation of an internal constraint (CCR).

There are several ramifications of this first assumption:

1. When we decide how to exploit of the CCR, we must also carefully consider the long term impact of this decision on the market. We don't want to refuse a client whose current order consumes a large amount of the CCR's capacity, when that order might be important to our long-term relationship with this customer.
2. Once we've committed to the market, the damage from not fully meeting that commitment might be considerably more severe than that incurred by sacrificing some of the CCR's capacity.
3. The preceding two ramifications imply that whenever a CCR is active, we are really in an interactive constraint

**ASSUMPTIONS:**

1. Differing degrees of CCR exploitation have significant impact on the bottom line
2. Making all the money we can NOW is important
3. The CCR must be exploited as much as possible to maximize Throughput now
4. Exploiting the CCR sometimes requires rejecting orders, delaying deliveries, or compromising some other customer requirement
5. There is no long-term effect to Throughput in rejecting orders, delaying deliveries, etc.
6. Any long term impact of customer dissatisfaction is negligible



**ASSUMPTIONS:**

7. Customers don't care about our need to exploit our CCR
8. Full exploitation of the CCR causes problems in satisfying customer demands
9. Full exploitation of the CCR leaves it vulnerable to "Murphy"
10. Fulfilling market demand often overrides the CCR as a critical limit
11. Our commitments to customers are crucial to retaining our market share

**RETAIN P2. REPLACE P1 WITH THE FOLLOWING INJECTIONS:**

- A. Maintain limited protective capacity at the CCR
- B. Exploit the CCR only with the limits imposed by the requirement to maintain some minimal protective capacity
- C. Select and develop market strategies (segments) carefully to support CCR exploitation (without compromising protective capacity)

**Fig. 2 THE CCR EXPLOITATION DILEMMA**

situation. We have both the CCR and market demand as constraints. The CCR might constrain our Throughput now, but whether (or how well) we satisfy the demands of our market surely affects our Throughput in the future.

4. Internal constraints may come and go, but the market constraint always remains in the background. In order to reasonably subordinate to the demands of the market, we need to maintain some minimal protective capacity on the CCR. The conflict resolution diagram in Figure 2 expresses this condition.

This is why S-DBR assumes the commitments to the market are always the major constraint, though subordinating to a market constraint does not preclude the possibility of having an internal resource that limits the company's market expansion. Such a resource should also be considered a constraint, because deciding what market segments the organization should pursue is dictated by any real or potential capacity limitations. Moreover, as market demand fluctuates, there is a risk of generating too much load on a CCR. Consequently, the load on that resource must be carefully monitored, and we may need to limit our commitments to some market segments.

To briefly summarize, when a capacity-constrained resource is subordinated to commitments made to the market, some protective capacity should be maintained at the CCR. For this reason, it might be argued that in most such cases there might not be a need to adhere to a detailed schedule for the CCR. This conclusion leads us to the second basic assumption of S-DBR:


*Basic Assumption No. 2: A small change to the actual processing sequence at an internal constraint does NOT have much impact on overall system performance.*

When this second assumption is valid, there is no need to schedule the CCR. This leads us to the basic rules of S-DBR.

1. The drum is based on firm orders. As orders come in, we make a quick check of the total load on the CCR. As long as the CCR is not too heavily loaded, the order is accepted as is and released immediately for processing. When the CCR seems to be too heavily loaded to ensure on-time due date performance, short-term measures to relieve the capacity loading are required.

2. The only buffer maintained is the shipping buffer. There is no difference between this tactic and the way traditional DBR handles an environment with no CCR. Although the CCR is not fully protected from starvation, in most cases enough work will accumulate at the CCR's site to prevent this from happening. Even on the few occasions when starvation technically does occur for a short while, all it does is consume a part of the protective capacity. Maintaining this protective capacity helps guarantee due-date performance, and the minimal sacrifice of current Throughput is likely to be more than offset by the future Throughput resulting from full subordination to the market.

3. The rope is no longer tied to the CCR schedule. The material release schedule is directly generated by firm orders received. We mentioned above that we handle transitory high loads on the CCR by using short-term measures to relieve capacity loading, such as overtime or additional shifts. But over the longer term, we need a way to regulate the incoming demand imposed on our CCR. In S-DBR, we tie the rope to the market demand, rather than to the CCR schedule. We do this by instituting longer-term measures to control short-to-medium term future demand. One way to do this is to adjust prices upward, or to provide incentives for customers to wait a little longer for delivery (e.g., charge a premium for shorter quoted lead times).



**Figure 3. S-DBR - A Graphic Depiction**

Figure 3 illustrates S-DBR.

### The Control Part of S-DBR

Buffer management is applied at the shipping buffer. We suggest a small variation to the traditional buffer management methodology. A “red-line” time (equivalent to the notion of Zone-1) should be defined. When the time remaining before the order should be shipped is less than the red-line time, a warning is given to production management. The red-line time is not necessarily a fixed portion of the shipping buffer. As in traditional buffer management, penetration of the red-line initiates expediting of a particular. When the number of "red-line" orders—those that have crossed the red-line— goes up sharply, a true bottleneck is emerging.

S-DBR offers an additional control mechanism that complements buffer management/ red-line control: *planned load*.

**DEFINITION:** Planned load is the total time (in hours) required of the CCR to complete all work that has been formally released into the system, but which has not yet been processed at the CCR.

The planned load consists of all known work that will eventually pass through the CCR. It's an average estimation of the time it takes the CCR to process everything that the company has committed to ship. To determine whether we have an emerging bottleneck, we compare the planned load at the CCR with the shipping buffer. The

result is an immediate indication whether the CCR will become a bottleneck or whether sufficient protective capacity still remains. The planned load supplies less information than a truly finite-capacity schedule, which can be more accurate and can warn of specific orders might be late. But if we assume that we'll refrain from loading the CCR to its limits, the information we get from the comparison of planned load with the shipping buffer is quite sufficient to warn us of a potential problem, long before a "red line" (Zone-1) buffer penetration occurs, and in plenty of time to take remedial action without making it an emergency.

## **Differences Between Traditional DBR and S-DBR**

Here are some key differences between traditional DBR and S-DBR:

*Level of Throughput.* Traditional DBR is capable of squeezing more Throughput out of the CCR in certain peak demand periods, due to the detailed CCR schedule.

*Customer Satisfaction.* The role of the shipping buffer in S-DBR is more dominant than in traditional DBR. When the shipping buffer is added to the CCR buffer, the protection of promised due-dates is less effective. This is essentially the same phenomenon that causes a critical chain project completion buffer to be more effective than a number of buffered intermediate points.

*Focus.* Traditional DBR is usually focused on the internal resource while S-DBR is focused on the market demand.

*Lead time.* Having one buffer, rather than three, enables S-DBR to achieve shorter lead-times. The accumulation of protection is always more effective than spreading it.

*Dealing with peak and off-peak demand.* In most situations market demand fluctuates. Assuming we can't fully level the load on a CCR throughout the year, we can conclude that the CCR is active only around the peak period, and the market demand is the sole active constraint in any off-peak periods. Shifting from three buffers to one buffer and then back to three buffers again represents a huge policy change for traditional DBR, with significant ramifications for management. For this reason, in most cases organizations using traditional DBR would regard the CCR as a constraint even in a period of low demand. This produces sub-optimal results (longer delivery times) in off-peak periods. S-DBR is able to shift smoothly between peak and off-peak periods, as the main focus of planning and control have not changed (satisfaction of the market).

*Support from common information technology (IT) packages.* S-DBR is much easier to plan and control with common MRP systems. In fact, specialized DBR software packages aren't really needed, since MRP systems can be adjusted to support S-DBR. This can be a real benefit to companies that already have MRP systems but might be unable or unwilling to invest in specialized DBR software.



*Ease of implementation.* S-DBR is substantially easier to implement, and the majority of the benefits in doing so accrue more quickly.

## **Integrating S-DBR with the TOC Approach to Purchasing, Finance and Marketing**

Success in managing a manufacturing company requires far more than just planning and control of the shop floor. The real benefits materialize when we use TOC to manage the entire supply chain, end-to-end. Treating market demand as an ever-present constraint and a potential or real CCR as an interactive constraint facilitates a broader view of our system and its environment. Doing so keeps our focus on the strategic implications of our daily management actions. Specifically, it opens our eyes to the importance of identifying the future constraint(s)—or at least the likely candidates to become constraints— before we arbitrarily choose a course of action today.

For instance, consider the notion of “free product.” In the parlance of constraint theory, the term “free product” refers to one that does not pass through a CCR. Capitalizing on a free product can be like a “license to print money.” But if we succeed in exploiting the market for a free product and increase the sales of that particular product, at some point another constrained resource will eventually emerge. This could create interactive constraints, which might cause even more difficult problems. Wouldn’t it be useful to know which resource, currently a non-constraint, might become constraint in the future? Knowing the answer to this question is crucial for keeping our system under control and for deciding where we’d prefer the constraint to lie. Or, in other words, where should the constraint be to give us the best Throughput, Inventory/Investment and Operating Expense performance? There aren't too many strategic questions (answers) more important than that! In order to answer this question, we need to look beyond production management. Our marketing, purchasing and financial policies must be considered along with production policies as an integrated whole. However, detailed discussion of these “integrating” issues is beyond the scope of this paper.

## **Conclusion**

The S-DBR global approach, including the financial, marketing and purchasing approach has been exhaustively tested using the Management Interactive Case Study Simulator (MICSS). The purpose of the MICSS is to demonstrate the feasibility of S-DBR methodology in a semi-complex environment with uncertain market demand, a high level of “Murphy” in operations, and unreliable vendors. The simulation is demonstrated in the presentation and constitutes part of a comprehensive education package.

## **ENDNOTES**

1. Goldratt, Elyahu M. and Robert E. Fox. *The Race*. Croton-on-Hudson, NY: The North River Press, 1986.

2. Umble, M. Michael and M.L. Srikanth. *Synchronous Manufacturing*. Southwestern Publishing Co., 1990.
3. Stein, Robert E. *Re-engineering the Manufacturing System: Applying the Theory of Constraints*. Marcel-Dekker, Inc. 1996
4. Goldratt, Eliyahu M. *Production the TOC Way: A Self-Learning Kit*. North River Press, 1996.
5. Cox, James F. and Michael S. Spencer. *The Constraints Management Handbook*. Boca Raton, FL: St. Lucie Press, 1998.
6. Goldratt, E.M. *The Haystack Syndrome*. Croton-on-Hudson, NY: The North River Press, 1990.

Drum-Buffer-Rope is a manufacturing execution methodology, named for its three components. The drum is the physical constraint of the plant: the work center or machine or operation that limits the ability of the entire system to produce more. The rest of the plant follows the beat of the drum. They make sure the drum has work and that anything the drum has processed does not get wasted. S-DBR is discussed in a paper by Eli Schragenheim and Bill Dettmer, Simplified Drum-Buffer-Rope: A Whole System Approach to High Velocity Manufacturing. Plant types Edit. There are four primary types of plants in the TOC lexicon. Draw the flow of material from the bottom of a page to the top, and you get the four types. Science of Business. Velocity Scheduling System - TOC Drum Buffer Rope. Job Shop Scheduling and Machine Shop Scheduling. Toc Schedule Drums Theory Clever Shops Timeline Tents Percussion. A custom polypropylene sheet and product manufacturer participated in the 14 week Velocity Scheduling System Coaching Program that is based on the Theory of Constraints Drum Buffer Rope -- but modified for custom job shops and machine shops. The dr. Science of Business. Drum Buffer Rope for Manufacturing Systems. In manufacturing, the drum is still the bottleneck. The buffer is the material upstream of the bottleneck and has to make sure that the drum is never starved. The rope is a signal or information from the buffer to the beginning of the line. If the drum processes parts, the buffer moves forward. Simplified Drum-Buffer-Rope is very similar to Drum-Buffer-Rope. The key to simplifying the approach is the assumption that the market or the customer is the largest bottleneck. I.e. in average your system always has enough capacity to satisfy demand. The rope then spans the entire length of the system.