Using SDBR in Rapid Response Projects

By Eli Schragenheim

Introduction

Simplified DBR (SDBR) has become the standard production planning method for all VV projects. This document describes the **make-to-order** aspects of implementing SDBR. The make-to-stock aspects will be dealt later in time. Many implementations use both make-to-stock and make-to-order. This document makes only brief comments on how to handle both.

One of the necessary conditions for offering successfully RR orders is being able to maintain close to 100% reliability of the regular orders. To achieve such reliability we have to be able to quote safe due-dates that are reasonable to the customers while taking the capacity of the CCR (capacity-constraint resource) into account.

The next stage is to see how can we offer RR orders while NOT tampering with the commitments already given for the regular orders.

The following document first presents the definition of SDBR relative to the traditional DBR, the planned load concept and how it relates to the finite-capacity schedule of the CCR.
Then the technique for providing safe due-dates for regular orders is discussed. This section also makes the distinction between the **production lead-time** and **quoted-lead-time**.
Next the need for reserving capacity on the CCR is presented, and the proposed process is discussed. Determining the release dates for the RR orders and other orders that take advantage of the reserved capacity.
The last part discusses the execution rules for such an implementation.

SDBR, DBR and the planned load

Simplified DBR includes only one buffer: the shipping buffer. This is true for both MTO (make-to-order) orders and MTS (make-to-stock, or rather make-to-availability). This is certainly applicable when there is enough protected capacity even on "the weakest link" in the shop floor, so there is no need to protect the limited capacity of any resource within the shop. We now propose that SDBR could be used effectively also when an internal capacity-constraint-resource (CCR) is active.

Traditional DBR uses a 3-buffer system to protect both the due-dates AND the **detailed finite-capacity schedule** of the CCR. This is far more protection than just keeping the CCR from starvation due to a delay, or lack of priorities, on the non-constraints. The CCR buffer is used to keep the actual sequence intact. Now, when we use only one buffer, protecting the due-date, we should have two concerns:

1. Failing to properly exploit the CCR capacity, especially causing starvation of the CCR or having too much setup time on the CCR.
2. Being unable to smooth the load through the planning time frame. The undesired result could be that we stick to a too long time-buffer or quote an unrealistic due-date.

The first concern could and should be handle by the foremen and the production manager as part of the execution process. In any VV project we recognize the market demand as the only constraint. That means we cannot tolerate true bottlenecks in the shop floor. In other words, we should have protected capacity even on the weakest link in the shop floor. It still does not mean we can ignore any capacity considerations, but it means we are reluctant to give up valid market demand because of lack of capacity.

So, how should we monitor the capacity of a fairly loaded CCR?

We expect the production manager to monitor the smooth CCR execution by avoiding starvation as possible, bearing in mind that when the demand is relatively low some starvation on the weakest link (the resource we treat as a CCR) would naturally happen. **But we do not need to determine the exact sequence at the CCR ahead of time**. The actual sequence should be impacted by the buffer management priorities and possible by setup time considerations when merging some orders to save setup time does not cause serious red penetration into the shipping buffer.

In too many times the actual sequence on the CCR has to be determined by the operators! This is especially true in complex operational environments where the sequence has huge ramifications on quality and on setup and process times. Most such complex environments do not have generic software that can handle the complexity. For instance, I don't know of any current DBR software that is capable to schedule a CCR that consists of several ovens, which are not identical, and in which several orders are grouped together based on their volume and quality considerations. In such a case you should provide a long enough shipping buffer that would naturally provide enough choice for the foreman for real time decisions. Trusting the knowledge of the people on the floor

To handle the **second concern** we need, first of all, a tool to identify situations where too much load threatens the standard time for deliveries. Identifying this situation is the first task of the **planned load**.

Definition: Planned load is the accumulation of the derived load on the CCR (the weakest link) for all the firm orders that have to be delivered within certain horizon of time.

For instance, suppose we have to deliver four orders within the standard time of one week. Order1 needs 10 hours of work on the CCR. Order2 requires 15 hours, Order3 only 4 hours and Order4 10 hours. The planned load is simply the total of 10+15+4+10=39 hours.

Planned load is expressed in hours or days of work. Naturally it should be, at all times, **less than the standard quoted lead-time offered to the market**. The amount of time difference between the standard quoted-lead-time and the planned load represents the current amount of protected capacity on the CCR and thus on the shop floor. Certainly we must have a minimum time difference between the planned load and the standard lead-time to ensure that the last firm order would still have enough time from the CCR until completion. Take the above example. If the regular week is made of five working days, eight hours a day, than the standard lead-time is 40 hours. But, the planned load of 39 hours seems too long to ensure safe delivery of all orders within a week.

While the planned load isn't a "schedule", it does have some key characteristics of a finite capacity schedule. The most important one: assuming the CCR is able to work continuously, **it signals the reasonable time when the CCR is able to work on a new order**, assuming no special priority for that order. This means we can assess a safe completion time taking the planned load into consideration and adding some time buffer to provide enough time for the non-constraints downstream of the CCR. We'll detail the process of quoting due-dates later in this document.

Thus we've overcome the two main concerns from not scheduling the CCR.

But, what do we **gain** by using just one buffer without a detailed schedule of the CCR?

Finite-capacity-scheduling of the CCR is based on the orders and priorities as known at the time of the planning. Any request for a change would either be ignored until the next scheduling session, or necessitates re-scheduling, which results in a shuffle of the older schedule. This means that any local noise would impact the whole system, rather then just the local area. Thus even a small change, like bringing one order backwards in time, creates quite a lot of noise in the shop. The more troubling point is that the "safe due-dates" that were given based on the CCR detailed schedule are not truly "safe".

Scheduling the CCR is especially problematic in make-to-stock, because the priorities could change significantly from one day to another because of a sudden surge of sales.

Hence, the move to SDBR, with the usage of the planned load to determine safe due-dates, and with the capabilities of reserving capacity for rapid-response orders and some other types of orders, seems as a superior method for VV projects when the linkage between Operations and Sales is under full control.

Quoting Safe Due-dates for Regular Orders

The shipping buffer is a liberal estimation of the time required from material release until the safe completion of the order.

We have several key assumptions when we determine the shipping buffer:

1. All known policy constraints have been eliminated. Policies like batching norms, transfer batch made equal to the process batch and early release of orders. The assumption is that

those flawed policies, norms and behaviors that are derived by the efficiency measurements have been eliminated.

2. We control the amount of work-in-process by choking the release of orders according to the load on the CCR. This means that the **shipping buffer** can be much shorter than the **quoted lead-time**, because the quoted lead-time has to consider the wait time for the order to be released to the floor as the CCR has many prior orders to work on. The shipping buffer only considers the time from actual release to completion, taking into account just enough WIP to ensure the smooth work of the CCR.

The current planned load, the accumulation of the load on the CCR, from all the firm orders, looks like the following chart:

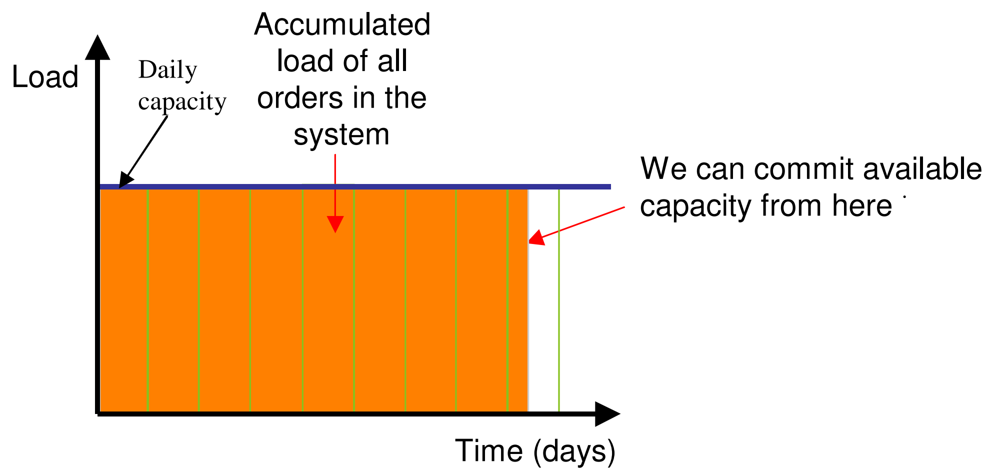

The front of the planned load signals when, on average, the CCR would be able to work on the new order. This is, of course, a very approximate estimation. We do not know the exact sequence the CCR would work on the previous orders. We do not really know how much setup time will be required for the actual sequence, and we are not absolutely certain that the processing time data is precise, even though we like it to be as accurate as possible. So, what we have is a good assumption regarding the timing the CCR could work on the next order to be received.

As we have only one buffer, let's decide to release the materials ½ the time-buffer prior to the time the CCR is supposed to work on it. This is not a very critical assumption, as the timing on the CCR is not all that concrete. We assume that within ½ of the overall shipping buffer enough orders will show up at the CCR to prevent unnecessary starvation. Note that when the CCR is not truly active, meaning it does have significant amount of excess capacity, it should be idle from time to time.

Here it is how it looks like:

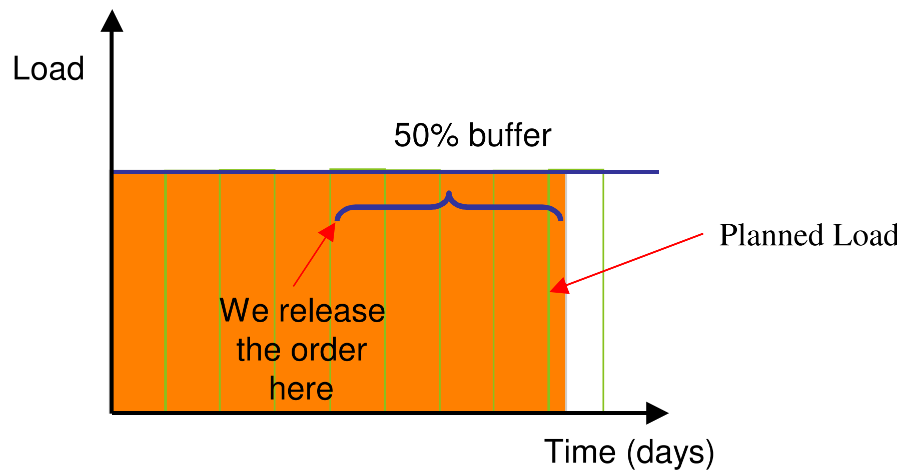

A clarification may be in place here. In the traditional DBR the CCR buffer is designed to protect the detailed schedule of the CCR – not just prevent starvation. The CCR buffer is put specifically for **each instruction** in the schedule, where in SDBR we use **half of the total buffer** (the shipping buffer covers the whole route from raw materials to completion) to determine the material release. All the materials get the same time release for an order. The release is dictated by the front planned load before the new order is added to it.

When should be promise the order? It seems now obvious that we could get a safe delivery time by adding ½ a shipping buffer to the current front of the planned load.

This is how it looks like in the graph.

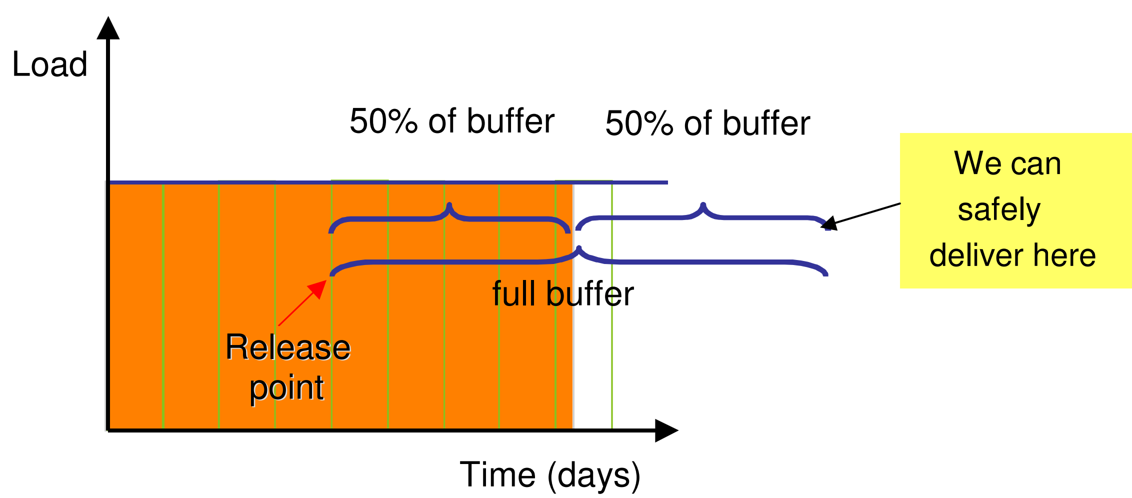

Note that we CAN safely deliver at that time, question is: should we promise to deliver at that date?

The point is that we do not want to offer a very early delivery time, safe as it may be in that point of time, which is earlier than the standard time. When we have rapid response in mind, we better not spoil our future customers with too early delivery.

So, if the safely time is shorter than the standard time then we should extend the buffer until the standard time.

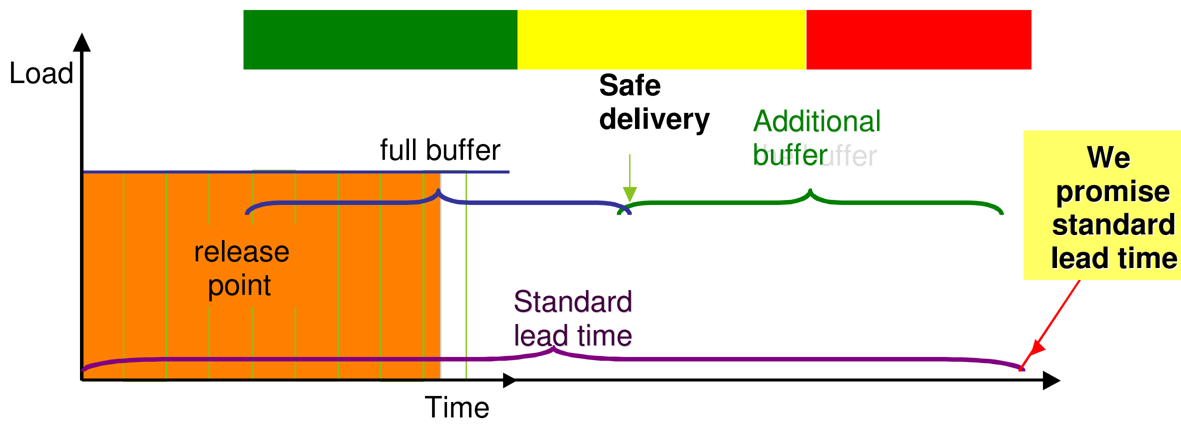

The idea here is to release the order early enough to ensure the CCR is not starved, hoping more demand will show up in the short term. If we'd delay the material release to provide the regular buffer, the possible negative branch is a waste of the CCR capacity. This would happen when a low demand period is followed by a very high demand period. If the capacity of the CCR was properly used in the low period, then it could do more in the high period without threatening the safe delivery within acceptable quoted lead-time.

An example to demonstrate the process:

Suppose that right now we have 30 open regular orders that should take 320 hours on the CCR. The standard lead-time for regular orders is 6 weeks. The factory works 6 days a week with two 8-hours shifts a day. The shipping buffer is 10 days, or 20 shifts.

When can we offer a reliable delivery date for a regular order?

The front of the planned load is at 320 hours from now. We have available capacity of 16 hours (2 shifts) a day. That means 20 days. We work 6 days a week, so the planned load is at 3 weeks and 2 days from now. The reliable due-date is the planned load plus half of the buffer: plus 5 days. So, the safe date is 4 weeks and one day from the current date.
In this case we'll stick to the standard time of 6 weeks from now.

The order will be released to the floor according to the planned load minus half of the shipping buffer: 20 days minus 5 days: 15 days – two weeks and three days from now.

As the delivery date is 6 weeks from now (36 days), the specific order will have a larger than usual buffer of: 36-15=21 days instead of the regular buffer of 10 days. This means we have a substantial amount of excess capacity, which would enable offering the rapid response offer, provided excellent due-date performance is established.

Offering the RR

Ensuring very high reliability of on time delivery is absolutely necessary for launching the RR offering. But, even from just an operational perspective, **it is not sufficient**.

In order to commit to certain clients, whenever the need arises, safe delivery in fast and even in super-fast time frame, the production lead-time must be short enough to provide the ability to do that.

How fast the production lead-time needs to be? The shipping buffer, the liberal estimation of the production lead-time, needs to be **about the length of super-fast time frame**.
One could argue that in many DBR implementations we sometimes accept an order to be delivered in significantly less than the regular shipping buffer. The rational is that the buffer length contains a lot of protection and most orders do not even penetrate into the red zone. Hence, when we are faced with lucrative order to be delivered in, say half of the regular shipping buffer, it is a "mission possible". Of course, we assume that naturally such an order, which penetrates into 50% of the buffer from day one, would get high priority, and other orders would "pave the way" for it to reach completion on time.

While this is quite true for an incidental opportunity, it is quite a different case to commit to such a delivery on a regular basis. It is possible to promise and deliver a very fast delivery once at a time without disrupting the on time delivery of regular orders.
But, how many orders, whose quoted lead-times are about half of the shipping buffer, can you handle at the same time? When you have several "super-priority" orders then there is a real risk that some of the regular orders will be late. And when this happens – your reputation as a very reliable supplier goes away.
Thus we need to have a reliable production lead-time that is approximately the same as the super-fast quoted lead-time. That means we need to reduce **the production lead-time to about a quarter (1/4!) of the standard lead-time**. This would ensure that a new super-fast order would be released immediately, but not automatically get higher priority than regular order that was released at the same day.

Still, even with very short production lead-time we must ensure enough capacity to supply all orders: the fast ones as well as the regular ones upon which the reputation of being very reliable is based on.
When we commit to deliver at a certain date we take into account the current orders in the logbook. If we might get orders whose delivery dates are definitely earlier than the date we just promised, there is a real risk that we won't be able to keep our promise. To overcome this risk we must ensure that capacity commitments will be preserved. Hence, we have to reserve a certain amount of capacity to the RR orders when we consider the safe delivery dates for regular orders.

The idea is to calculate the planned load for regular orders on only 70 or 80% of the available capacity, leaving 20-30% of the capacity for the RR orders (and some other orders as we'll see later).

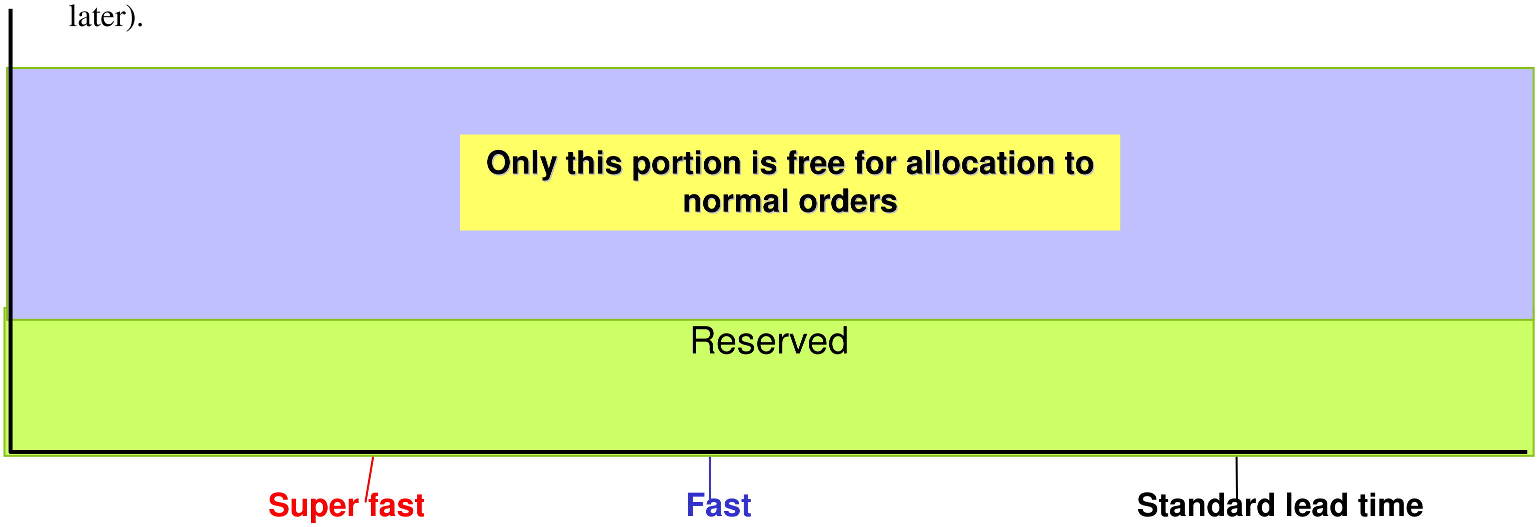

A simple example to illustrate the notion of capacity reservation:

Let's continue the previous example. The standard quoted lead-time is 6 weeks. The shipping buffer is 10 days (20 shifts). The shop works 6 days a week, two shifts a day.
The Fast RR offer is for 3 weeks delivery. The Super-fast offer is for 1.5 weeks, which means 9 working days. Just one day less than the shipping buffer.

Suppose we already launched the RR offer, and we decided to reserve 25% of the total capacity for the RR orders.
Today we have 396 hours of CCR load from firm regular orders in the logbook (yes, the demand went up since the reliability was established and recognized). Then we have also 20 hours of the CCR for Super-Fast orders and another 40 hours for Fast orders.

First let's calculate when can we promise safe delivery of a new **regular order**. As 25% of the capacity is reserved, we have only 12 hours per day (16 hours minus 25%) for regular orders. So, in order for the CCR to process all the existing regular order it would take: 396:12=33 days. To promise safe delivery we need to add half of the shipping buffer: another 5 days. So, we come now to 38 working days from today, or 6 weeks plus 2 days. So, reserving 25% of the capacity means having to quote a somewhat later delivery than the standard. However, that date is fully guarantied. The order will be released 33-5=28 days from now, and the whole shipping buffer will be the regular 10 days buffer.

Now, suppose an additional Super-Fast order is received today. The delivery date is fixed: 9 working days from now.
The order material release is based on the due-date minus the shipping buffer. In this case it is clear we must release the order TODAY. The buffer status for that order is (10-9)*100/10=10% (10% penetration into the buffer due to the requirement to deliver in 9 days rather than 10 days).

So, any new Super-Fast order has some small precedence on other orders that are released today. But, that's not too much to cause any real problems.

A Fast order received today is to be delivered in 3 weeks: 18 working days. The buffer is 10 days, so we'll release the order only 8 working days later than today! See, the reduced shipping buffer allows us to wait until we release an order that is part of the RR offering and we get much better price for it.

Launching the RR Offers

We need to be somewhat careful when we tell the first time the prospects for the RR that we are ready to accept any order for Fast or Super-Fast delivery. **At that time the orders in the logbook were given dates based the full amount of CCR capacity.** Hence, there is a need to reserve the capacity for the RR and the earliest time we can be certain to have the reservation is the current planned load time.

Hence, the RR launch date from when we decide we are absolutely ready is the front of the planned load.
If we take the above example, assume we have 396 hours of the CCR planned-load, based only on regular orders. That means, according to 16 hours a day (no reservation of capacity yet) the calculation come to: 25 working days from now (no need to be more precise).

So, 25 working days of the full CCR capacity are already committed. We can accept Fast and Super Fast order starting from that date! Any new regular order coming today should be based on 25 days from now, plus 75% of the CCR capacity from that day onward.

Managing the RR orders

Once orders are released to the floor, the priorities in the floor are dictated by their buffer statuses. RR orders do not have special priority. In no way the CCR should stick to the sequence that appeared in the planned load. Certainly we **do not** mean to tell the CCR, or any other resource, to work 75% of the day on regular orders and 25% of the RR orders.

Naturally when there are no RR orders the CCR works 100% of its available capacity on regular orders. When this happens then the planned load gets shorter in time, because the calculations consider capacity reservation.

When many RR orders show up they could consume more capacity of the CCR than what was reserved. Then the planned load would increase and new orders would get a later date. The existing regular orders would be squeezed to some degree (as less capacity of the CCR is used for the regular orders), but the shipping buffer and buffer management should be able to support the on time delivery of all orders.

But, we need a mechanism to determine whether the reservation percentage is about right.

The planned load as was defined above contains only the regular orders. The RR orders are supposed to be done by the use of the reserved capacity. Hence, we could fill in the reservation part with the Super-Fast and Super orders and see where it could take us.

In the following figure we show the two load representations: the regular orders planned load, and the reservations part, partially filled in. This is because every order in the reservation part has to be delivered within the Fast time, while the regular planned load includes orders to be delivered up to the standard time.

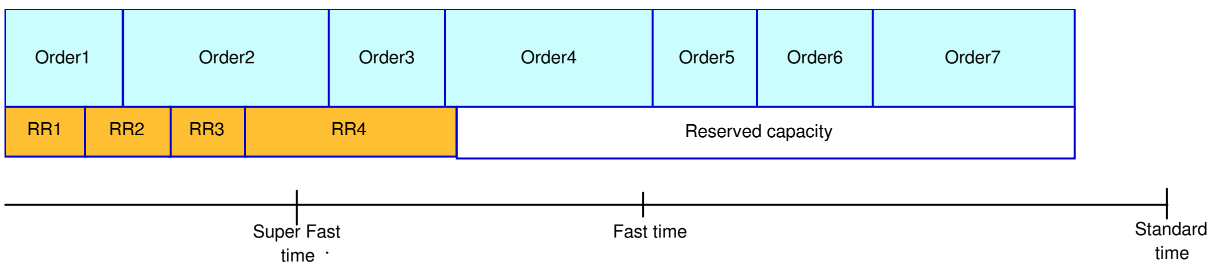

What can we deduce from the above graph? It seems there is no clear threat to the regular orders coming from the amount of Fast and Super-Fast orders. Of course, if all the four RR orders are Super Fast orders then we could have a problem. In such a case all the CCR capacity within the Super-Fast time would be dedicated, first of all, to the orders that are due within very short time. That means certainly the Super-Fast orders, but could be also some regular orders. According to the graph Order1 might be to be delivered before the Super-Fast time also. Such a short term peak could cause a growing pressure on the CCR and a threat to the timely delivery of Order2. Other orders might not be in trouble at all, as currently the reserved capacity is not fully taken.

Certainly we need buffer management on top of the above graphs to manage all the orders, and we need POOGI to have a clearer idea whether the amount of reservations is too high or too low.

For a short-term capacity control a picture of all the orders that are due within the red-zone of the shipping buffer, and a picture of all the orders due in full shipping buffer time is quite relevant. The planned load procedure ensures no peaks of loads on the CCR due to regular orders, but the RR orders might cause temporary peak of load, especially within the short time frame of one shipping buffer.
This kind of graph showing the short-term load should take into account ALL orders, both regular or rapid response, that their respective due-date falls within the specified horizon.

Suppose that in the above graph, all RR orders are Super-Fast, and Order1 is also due in the time frame of Super-Fast from now. Taking all orders and all available capacity we get this relationship of CCR capacity to the time frame:

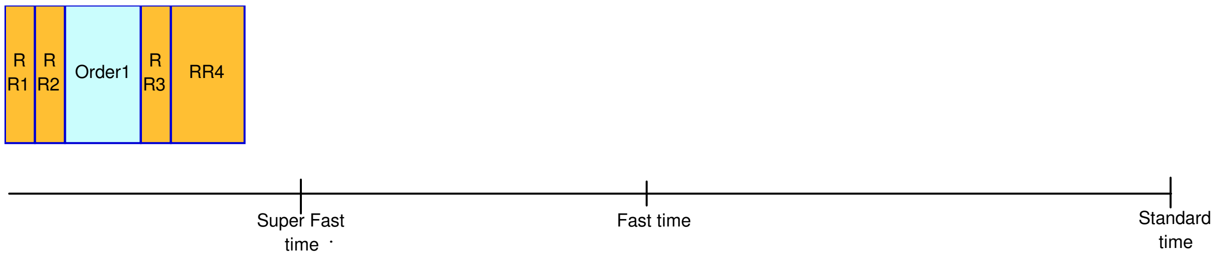

It is sensible to assume the CCR would first work on those five orders, which are due before the other orders and thus we expect their buffer status reflect higher priority.
Are we certain that those five orders will all be delivered on time? This is definitely a problematic situation, and we better not let this state to happen. The problem is clear because there is not much time left for RR4 to provide enough time to reach completion. We need the exact buffer management statuses to really assess the situation.

Another question is whether the short-term peak of load won't have a negative impact on other orders, those that should be delivered sometime later to the Super-Fast time, but not much later. Certainly having the CCR work solely on orders that are due very soon endangers orders that are due later. To have a better idea we should look for the total planned load of all orders that are due in Fast time frame:

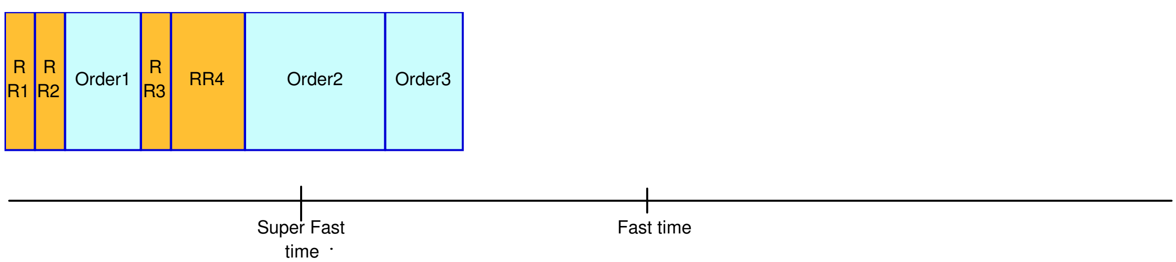

Here the difference between the Fast Time and the CCR load of orders due before or at the Fast time looks reasonable. Actually we'd like to see that there is half a shipping buffer difference to feel about right.

So, right now the Fast time frame looks about right. But, it is quite heavy when you consider the possibility that if more Super-Fast order would show up and take the rest of the capacity. Hence, a good capacity control should deduce that there is a need for higher level of reserved capacity OR elevate the CCR capacity.

All in all we do get signals when capacity pressure starts to build. If there is a quick way to increase capacity when truly needed then the above situation could be easily handled. Otherwise, we need to employ a more cautious capacity planning, reserving more capacity for the RR orders, to ensure we do not threaten the due-date performance of all orders and commitments.

More needs for capacity reservation

The need to reserve capacity on the CCR stems from committing to ship certain future orders earlier than the current planned load of the CCR. In other words, the current planned load shows only part of the actual load, and thus committing to a due-date based on it is not really safe.

The RR orders are typical to such orders that, once they are received, we must ship them earlier than the dates we already quoted to regular orders. There are other types of orders like this:

- ✓ Strategic clients with whom we have a fixed delivery time.
- ✓ Products that have significantly shorter standard lead-times than the main bulk of orders.
- ✓ Products that are made to stock (MTS).

Let's deal first with short-lead-time products. Suppose that while the regular orders have standard lead-time of 6 weeks, a more routine type of products have industry standard lead-time of only 3 weeks. Both types need the CCR, but the routine products need very short processing time, hence the perception in the market is that these products should be delivered pretty fast.

As long as the front of the planned load is such that 3 weeks would still be safe due-date, meaning the front of the planned load is appropriately earlier than 3 weeks then there is no problem. We still would promise safe due-dates to all regular orders within the time frame of the standard lead-time and be very reliable on time delivery.

So, the problem is valid only when the planned load of the CCR goes beyond the standard lead-times of the short-lead-time products. Then when you promise delivery to a long-standard-lead-time item you don't know how many short-lead-time items you might get. Of course you could quote all delivery times based on planned load even when the standard time is short.
You are invited to construct the cloud: quote a shorter lead-time to a routine order versus quote a time based on the planned load.
There are two obvious solutions:

1. Tackle the cause of the problem. If the real constraint is market demand, why do we need to be constrained by a capacity constraint? If we add capacity to the CCR, the planned load would be within the time that would allow us to quote the standard lead-times without any compromises.
2. Use the idea of capacity reservation for those products as well as for the RR orders.

The above discussion is valid for the strategic customers as well.

The make-to-stock products is a different case. It is quite common to have a hybrid environment that has both MTO and MTS. Thus we might have a VV implementation with both RR and VMI.

The concept behind producing to stock is to replenish frequently and quickly. The expectations are to release daily orders for the stocked products based on the sales of yesterday. But, those orders to not have a definite due-date, their respective priority will be determine according to tomorrow's sales.

Of course, the production manager actually releases to the floor MTS orders considering the load on the CCR and the relative priority of the stocked products. Once additional MTS orders are released, some of them would spend more time in the shop and others would be expedited.

So, how should we handle the case of orders without definite due-dates?

I think we need to consider some amount of reserved capacity to ensure the ability to respond to the MTS priorities. But, this amount is lower than the ratio of the load generated by MTS orders to the total load. Thus, once MTS orders are generated their respective load on the CCR should be included in the regular planned load. But, we have to reserve some amount to be able to respond faster when necessary, without sacrificing the reliability of the RR and regular MTO orders.

One important observation: It does not make sense to manage separate capacity reservations. Certainly we can put all of them into one ratio of reserved capacity. Then we should monitor the total amount of reserved capacity by watching the summation of the reserved orders actual load, as was described above.

Simple versus sophisticated solution

Here is a generic cloud for many TOC practitioners, and too many other practitioners.

The relevancy of the cloud to the process described is that we have chosen many simplifying steps and guidelines and some might be tempted to "improve" them. For instance, when we determine the safe due-date we do not consider the net processing time of the order itself. Why? Because it cannot be very significant relative to the buffer size (we still assume it is manufacturing) and many times when we consider an order that was not formally received, we do not know the amount of work on the CCR.
Also, the decision to use ½ of the shipping buffer to dictate the due-date and the release time is also a very simplifying rule, but do we really need to go into details and ask what is the location of the CCR in the routing? I certainly do not think so.

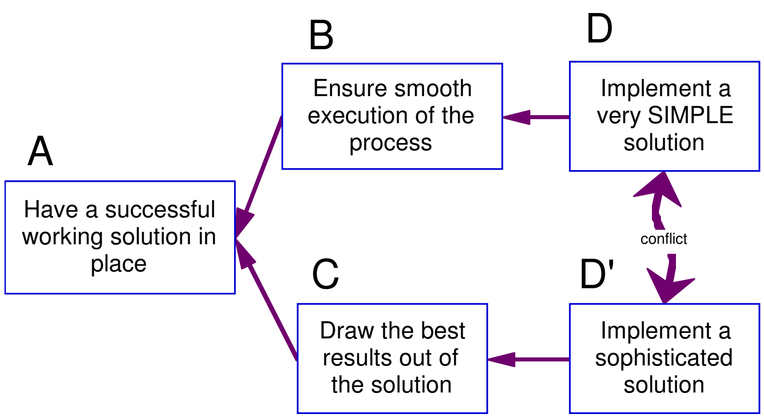

The injection of the cloud is to challenge the assumption that "It is possible to get a significant improvement over the simple solution." The challenge is to recognize the amount of the "noise" in the system and if the "significant improvement" is not larger than the noise than we'll not be able to really materialize the improvement.

Be aware to the amount of noise we have in a typical RR implementation. Usually the assessments of how much CCR time we need are not very good (most of the RR orders are customized products), the assessment of the buffer is very approximate, we have very vague idea how many RR orders will be received etc. Sticking to the simple straightforward solution is so much easier to implement and ensure the people in the floor won't distort it in the future, which is a very serious threat to the implementation.