Buffering Value Streams for Flow:

A practical guide for difficult Value Streams

Bill Kerber
Future State Solutions

2004 Copyright Future State Solutions, LLC (888) 541-4500

Introduction

In beginning the process of converting from a current state to a future state value stream, the most common and effective way to plan the transition is to use the seven guidelines that make a value stream lean. These are: Takt time, Finished Goods Strategy, Flow, Pull, Schedule One Point, Interval, and Pitch. However, many plants run into difficulty implementing this approach because the examples used to illustrate these principles are simple and do not address the complexity that they are faced with on a day to day basis. This paper will present solutions to some of the more common difficulties, focusing in particular on how to take customer demand and produce in a Lean manner.

Conventional teaching says that when two operations are not compatible, flow will be broken and a supermarket must be placed in between the two processes to buffer them from each other. This compatibility can be measured in a number of ways, such as differences in cycle times, physical space between the processes, differences in set up times, or simply because the process is shared between more than one value stream. For most plants this would mean that there would be a supermarket between almost every set of processes in the entire facility. Since this would clearly not be feasible, and certainly cannot be considered lean, an alternative to this approach must be found. One possibility is to specify that supermarkets are needed after any operation where we cannot easily define what must be made next by the process. In simple terms, this means that the operation has no visibility of customer demand in that time frame in which a production decision must be made. Placement of supermarkets could then be determined by comparing the lead time through the facility to the demand lead time given by the customer. The contention here is that as long as the production lead time is less than the demand lead time, then production can be made to order, and supermarkets are only needed for leveling demand into the value stream or for raw material. FIFO lanes can move material through the plant in the sequence desired, with no waste of unsold inventory which is present in a supermarket.

We would begin an analysis with a comparison of customer and plant lead times, starting with a picture of how much visibility of customer demand the plant actually has. This can be portrayed in a simple chart depicting the fill rate of customer orders in each interval in the future:

2004 Copyright Future State Solutions, LLC (888) 541-4500

Customer order fill rate over time

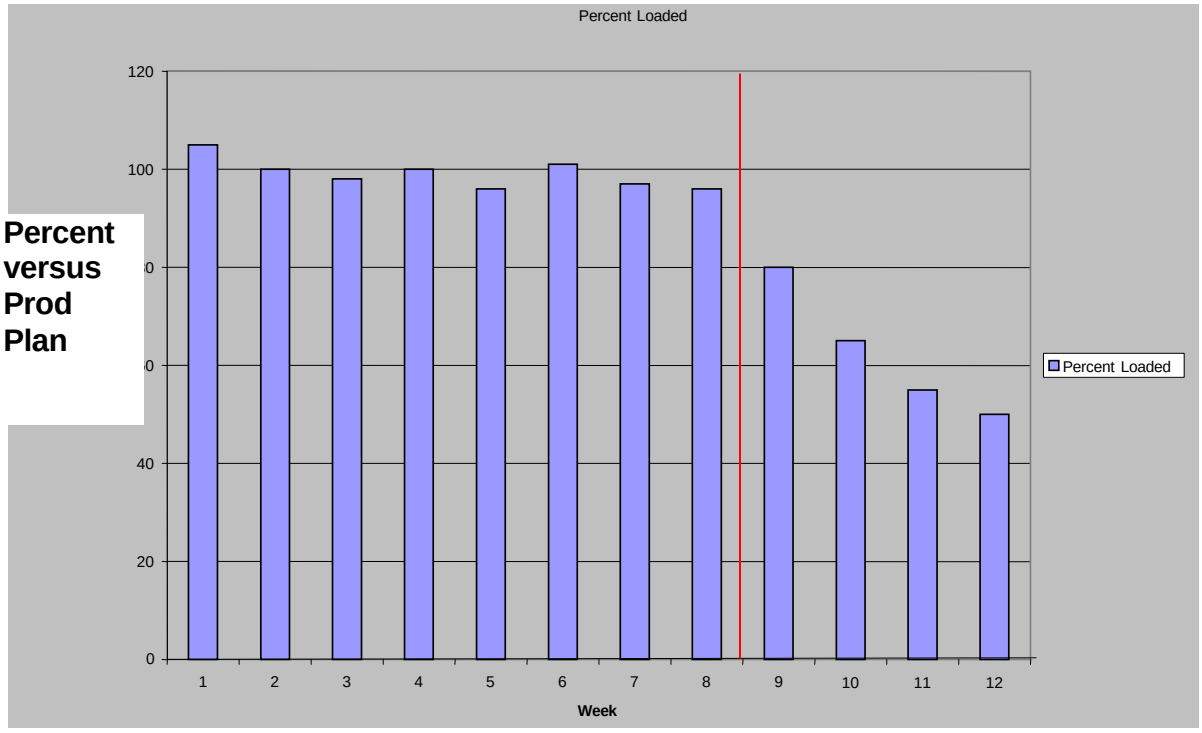

This chart suggests that most products can be made to customer order as long as the plant lead time is less than eight weeks.

Once we have decided that we would like to produce product to customer order, we can begin the process of setting up the factory to do this.

2004 Copyright Future State Solutions, LLC

1. Create a Current State Value Stream map.

The current state map serves several purposes in the process. First it gives everyone involved a clear picture of how the current process is performing. Second, it serves as a basis to give training about what makes a value stream lean. And third, it gives insight into the operating philosophy of the company. Since any long term movement to lean thinking will involve changing the company's approach to the manufacturing process, understanding the current philosophy is a very important step towards becoming lean thinkers. Finally, the current state map provides baseline measurements for processing time and lead time.

2. Map future state, including lead time ladder.

Use the standard seven steps that make the value stream lean as a guideline for this process. These are: takt time, finished goods strategy, flow where you can, pull were you can't, schedule one point, interval, and pitch.

When mapping the Future State it is very important to pay attention to the size of the FIFO lanes which will be created. These FIFO lanes will dictate the lead time through the manufacturing process. As a part of the Future State mapping process, it is vital to get agreement from everyone involved with the value stream that the FIFO lanes will be managed to very specific rules and regulations. Primary among these is that if the FIFO lane is full, then the feeding operations must not make any more parts which will go into the FIFO lane. This process is identical to the discipline which must be present to make supermarkets and pull systems function correctly.

The lead time ladder should use the same calculations as the standard value stream mapping lead time ladder. The variations arise where there are branches in the process flow, in which case each branch should also have a branch of the lead time ladder. Further, many processes may have different cycle times for different items, in which case a range of cycle times from low to high should be recorded in the data box. Don't forget that the lead time ladder records process time, which is not always the same as cycle time, especially where more than one person is used to perform the work, such as in a cell or an assembly line. In all of these cases the process time is usually much longer than the cycle time.

2004 Copyright Future State Solutions, LLC

3. Create Customer order profile as % of demand on hand by week

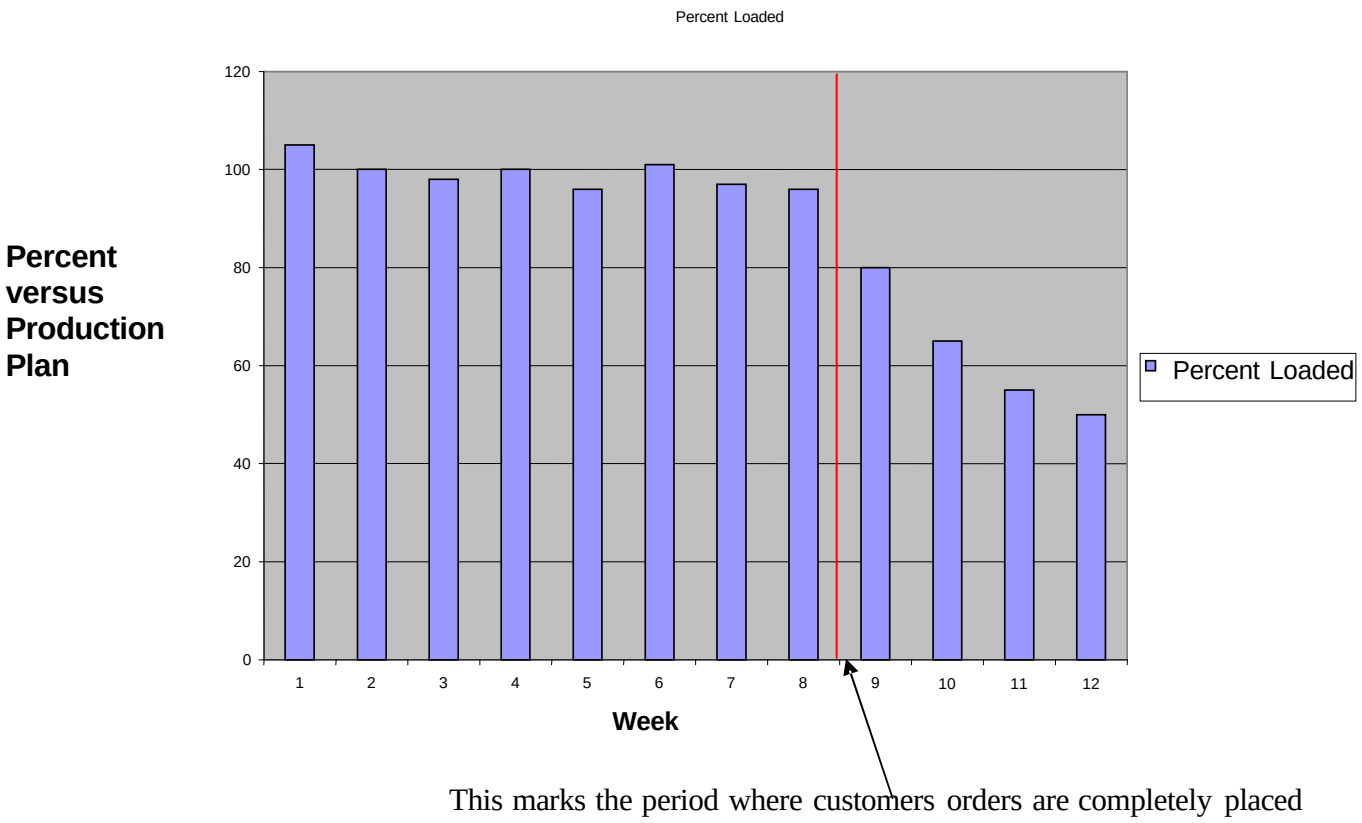

This chart can be created by using the expected production rate as the 100% line in the chart. Alternately you can also use the expected sales rate from the sales and operations plan to mark the 100 percent line for the chart. Either way the chart compares what the expected sales for each interval to the amount of volume which is currently sold for each interval. The information for this chart should be easy to find since most companies track their sales backlog and therefore you can have a look into the future on this chart. In a make to stock company it is not unusual to find that this customer order profile horizon is very short. This typically dictates that all items must be kept in stock unless negotiations can be performed with the customer to produce a few low volume items after a customer order is received.

4. Determine EPEI for the product family

EPEI is every part every interval. This refers to the capability of the process to make the setups required to cover the demand for all of the parts that the machine produces. The

interval is very important in all value streams. It is a measure of process capability and flexibility. Processes that can achieve smaller intervals are much more capable of responding to changes in customer demand, whether from volume changes or mix changes, than processes which have long intervals. In job shops we may have an interval which covers the parts that the customer has actually ordered over a period of time rather than an interval which covers making all of the items in a family. This is because it is very common to find many items that have low demand and infrequent customer orders. In some instances it is therefore not feasible to completely level the mix within a value stream. In these cases we make sure that we level volume, and make the products which the customer has ordered first. We will discuss this process later on.

EMC Supply - Future1-EPEI Calc

Process: Test | Configure and Test

Item Number	Interval Demand	Test C/O	UM	Test C/T	UM	Test s/u r C/O	UM	Test s/u r C/T
18392	32.5000	4	MN	62.72	MN	1.83326	MN	62.72
21000	42.5000	3	MN	82.01	MN	1.83326	MN	82.01
22020	17.5000	5	MN	30.7	MN	1.83326	MN	30.7
28121	35.0000	5	MN	67.54	MN	1.83326	MN	67.54
30000	25.0000	3	MN	46.05	MN	1.83326	MN	46.05
31000	27.5000	5	MN	43.42	MN	1.83326	MN	43.42
38200	10.0000	5	MN	35.09	MN	1.83326	MN	35.09
42005	12.5000	4	MN	21.93	MN	1.83326	MN	21.93
xxxxx	22.5000	5	MN	43.42	MN	1.83326	MN	43.42
.....								
Total C/T				432.88	MN			432.88
Work Time				450	MN			450
Avail for C/O				17.12	MN			17.12
Total C/O Reqd				39	MN			16.5
C/O Avail - Reqd				-21.88	MN			0.62
Target Avg C/O				1.9	MN			1.9

Interval: 450 MN Days per Int: 0.5 Recalc Export OK Mode Chart Help

This is a typical interval calculation, showing the demand for each part in the interval, the cycle time, and the set up time for each item. The total run time required is calculated, then compared to the available time, with the remainder of the available time designated as available for set up. As we can see in the column Test C/T we have seventeen point twelve minutes available for set up during each interval. Currently some setups take 39 minutes

for all of the items in the family. This means that without some improvement we cannot achieve the desired interval. The column labeled test s/u r C/O shows a target set up time of 100 seconds, or 1.83 minutes. When this is achieved we will be able to run parts at the desired interval. This setup target should then be used to drive the Kaizen activity for setup improvements at this process.

Graphically, these two look like this:

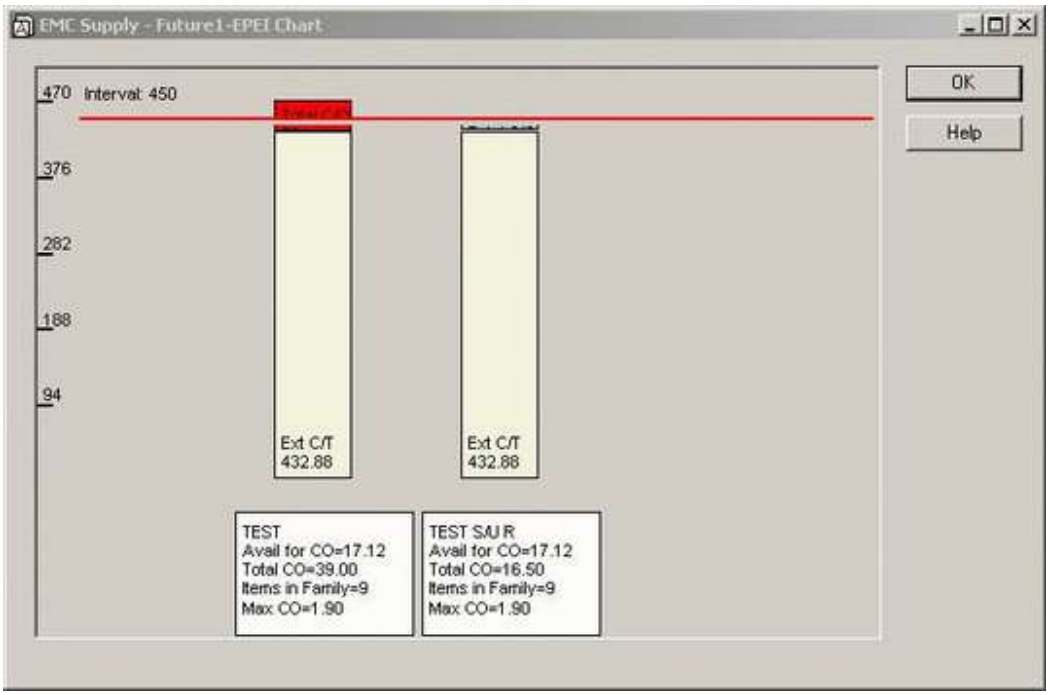

As mentioned before, often times job shops do not get customer orders for all the parts that they sell during a chosen interval. Because of this, it is usually not feasible to make the interval so large that it contains a customer order for every single item. What occurs in this case than is that we select a desired interval, and determine whether we can make each customer order within that interval. In other words, popular items will be grouped together into a single run during the interval, if they have sold multiple times during the interval, while some items that are only sold once will be run once, and many items that are not sold to the customer at all during the interval will not be produced.
In such a case interval is expressed as the volume that will be produced, plus the number of setups [or part numbers] that will be performed over the course of the interval. An example would be: each month we will produce 1,000 units and produce 40 set ups, meaning 40 different part numbers. We could also express this as 50 units and two setups each day in a twenty day month.

Shown below is a typical example of an interval calculation for the above condition. In this case, since we do not know the exact mix which will be sold, we must use the average

cycle time for all of the parts produced by the process, plus the average set up time for all the parts. We then enter the total amount of volume which is required for the interval, plus the number of setups that would be required to make the number of items which we anticipate selling during the interval. The calculation will then determine what the actual setup time needs to be to achieve the interval desired. Again, this should become the target for Kaizen an improvement for the process.

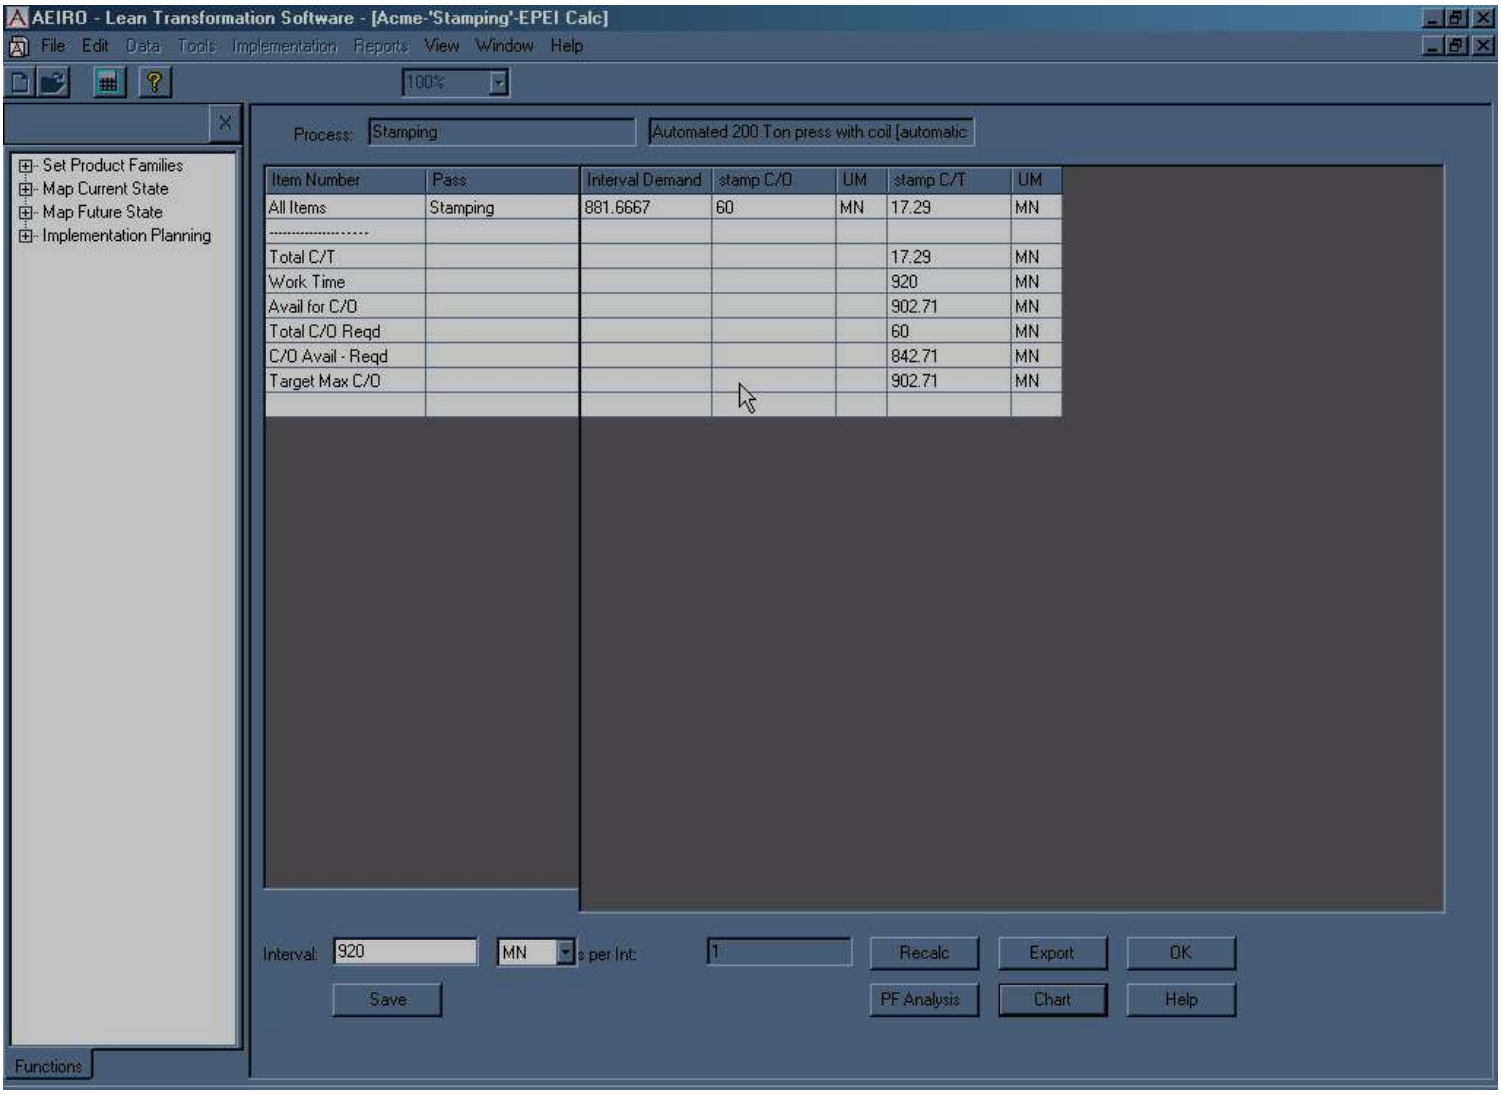

5. Overlay customer lead time and profile over the plant lead time picture, subtracting the EPEI period from the back end (Finished goods).

The EPEI determines the lot size, and in MTO it is simply an accumulation of the orders for the EPEI period. This total amount will need to be delivered at the beginning of the EPEI period to satisfy the first order which triggered the production in the first place. Therefore the longer the interval the less time the factory has to respond to the first demand in an interval, since you must wait until all of the orders are placed in the entire length of the interval.

6. Place the supermarket at the point where the customer order must be started in order to get delivered on time.

This should be shown by a simple over lay of the standard lead time ladder from the value stream map compared to the charts at the top of the paper which shows the customer order fill rates by time. As long as the production lead time is less than the customer order horizon, then we can wait until the customer order is received until we begin the manufacturing process. If, however, the manufacturing lead time is longer than the customer order horizon, then we must begin to manufacture products before the customer order has been received. In this case, we will use the signal from the supermarket to determine what we should begin to manufacture. If there is a natural break in the flow of products, this would be an ideal location for the supermarket. This may or may not to be at the exact point dictated by the lead time comparison, but this should not be of a much concerned, at least in the short run for the value stream managers. As long as the manufacturing lead time from the supermarket to the customer is less than the customers the lead time, then the value stream will be able to provide excellent customer service

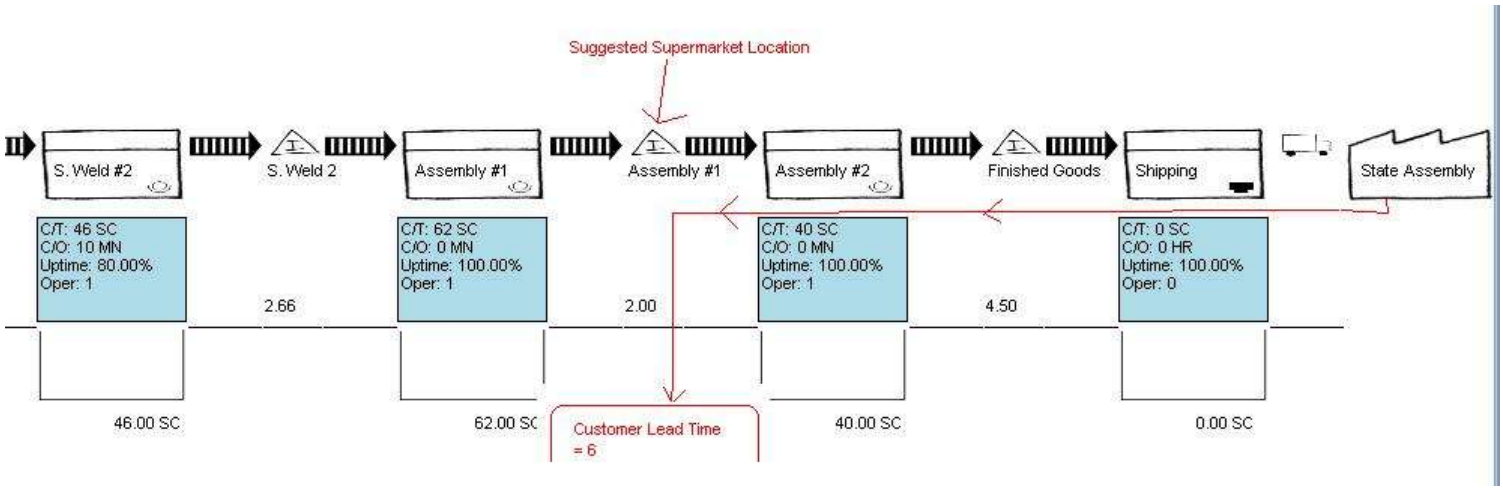

7. Flow these orders via FIFO lanes from the supermarket to the customer.

FIFO lanes allow product to move through the manufacturing facility without being scheduled at each process. The logic simply says that whichever item arrived at the process first should be the one which is worked on first. As long as the size of the FIFO is small in terms of lead time days then this hands off approach will invariably be successful. If, on the other hand the FIFO lanes are long, then it may be necessary to rearrange the priorities of work in front of the process. This is why it is very important to maintain short FIFO lanes and to maintain the discipline at the process which says that the next item in line is the one which should be worked on next. The size of FIFO lanes can be determined based on several factors. One of these factors is whether the two processes being joined by the FIFO lanes have the same cycle time for all of the items. If they don't then the size of the FIFO lane is directly related to the EPEI of these two machines. The main factor is how much tolerance for "indirect down time" management has. In other words, how long can a machine run after the feeding process is stopped for some reason, such as break down or no material.

8. Shared resources need to operate on a train schedule which should be reflected in the lead time for the item.

A train schedule is a fixed sequence and time representation of how product will be run through a process. It should include at least the shared resource being managed, and may also include the processes which feed the shared resource, and perhaps the processes that are being fed by the shared resource. In many cases since we do not know the exact volume that we will need to run of the various products in each interval, gaps are left in between the sequences of parts so that an item that has a heavier demand than usual will still have time to produce all the products required.

Click on the link to see a sample train schedule with two product families in it: www.futurestatesolutions.com .

When dealing with a shared resource, that is to say a process which is used by more than one value stream, it is important to ensure that all of the value streams using the process each get a fair share of the resource. One way to make this happen is by creating a train schedule for the shared resource. By using the interval calculation for the shared resource we should be able to determine the relative amount of time required to produce parts for each of the value streams for an interval. These times can then be blocked off during each interval so that each value stream knows when during the interval that their parts will be produced by the shared resource. This can help immensely in keeping queues in the FIFO lanes short. This process often provides further benefit in that by controlling the sequence of items being produced by the process, setups can often be reduced by creating a sequence that involves only minor adjustments to machines to go from one item to the next. This can then further reduce the interval and improve flexibility of the shared resource to satisfy all the values streams that serves.

9. The Value Stream must be loaded to defined capacity so that processing lead times are relatively fixed.

For the most part it is a good idea to give the factory the amount of work which they were staffed to do. If we give them less work than our plan called for then we risk having productivity be lower than required, and if we give them more work than the plan called for we risk increasing work in process, paying for overtime which is not required, and potentially increasing lead time through the factory.

10. Set up a buffer between demand and supply.

We can buffer a value stream in three different ways: with a supermarket of inventory, with labor, or by varying lead time to the customer. A brief review of each follows.

Inventory is the most common method used to buffer the value stream. However, most of the practices that we use to manage inventory work in opposition to the desire to buffer the value stream from demand fluctuations. This is mainly because the software that we use to manage inventory wants to use a fixed inventory target to determine how much production is needed from the factory. Lean and sales and operations planning both want to treat the capacity of the factory as fixed, while treating both demand and inventory as variable. The implications for this are that the software which is feeding work to the factory will create supply which roughly matches demand. If we're lucky enough to have demand match the plant capacity, this doesn't create any problems. However, this is rarely the case. Because demand varies and plant capacity typically does not, we usually get a workload which does not match the capabilities of the factory. This can be clearly seen by the following charts. The answer to this is relatively simple, although it does represent a significant departure from way production and inventory are currently managed.

When creating a production schedule, the simplest way is to start with what the customer has ordered and determine whether the factory can produce these orders. If the capacity required matches the capacity of the facility then it is a simple matter to introduce the customer orders as the desired production schedule. If the customer orders do not fill the entire production capacity, the schedule should be supplemented with a few high demand high volume parts which can then be placed into a supermarket. If the customer orders require more capacity than the facility currently has, then a few of the items should be pulled from the supermarket of high demand, high volume parts and the remainder of the items should be placed into the production schedule. The implication of this strategy is that as demand fluctuates, the inventory of supermarket parts will fluctuate in the opposite direction, I E. if demand is higher than capacity then inventory will reduce over this interval, while if demand is less than capacity, then the inventory will increase over this interval.

Chart one shows the way a typical ERP system wants to manage supply and demand. Using a fixed inventory target, the system creates supply which exactly matches the demand. This requires the plant to vary capacity to match varying demands.

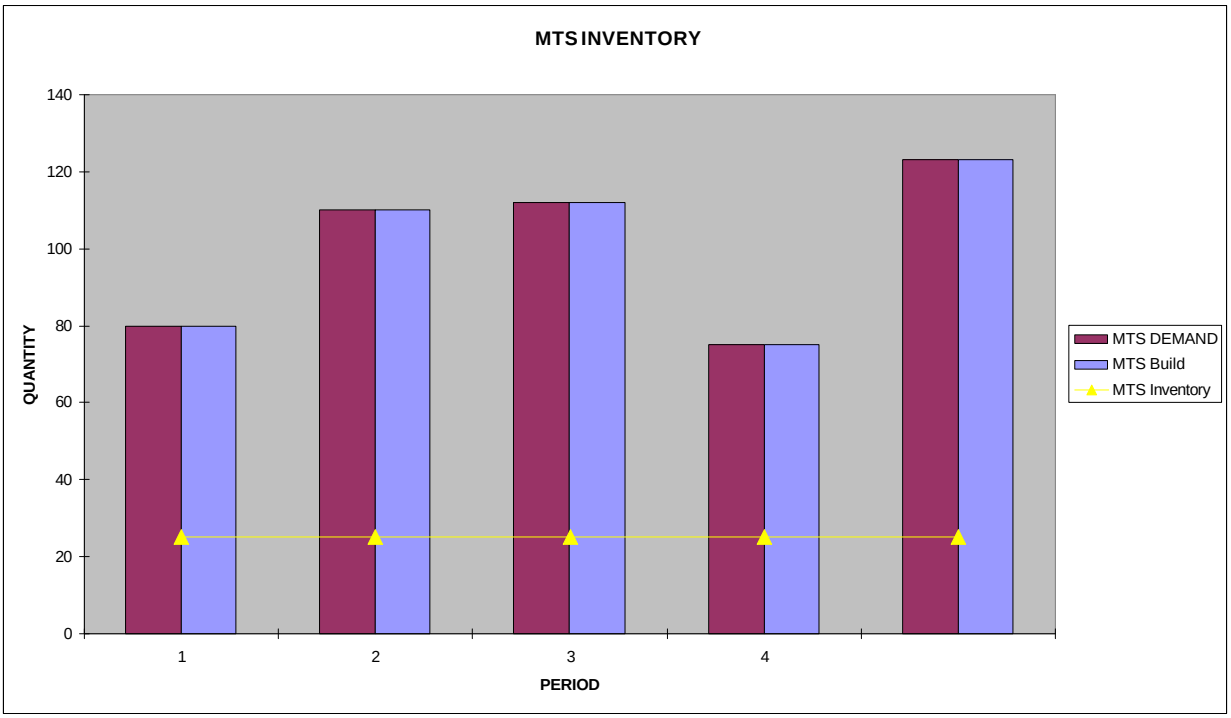

Chart two shows an alternate strategy which basically holds production at a constant rate, and allows inventory levels to vary as demand varies. This is a much easier strategy to execute, as the factory can maintain a constant level of manpower and produce the same amount of material during each period.

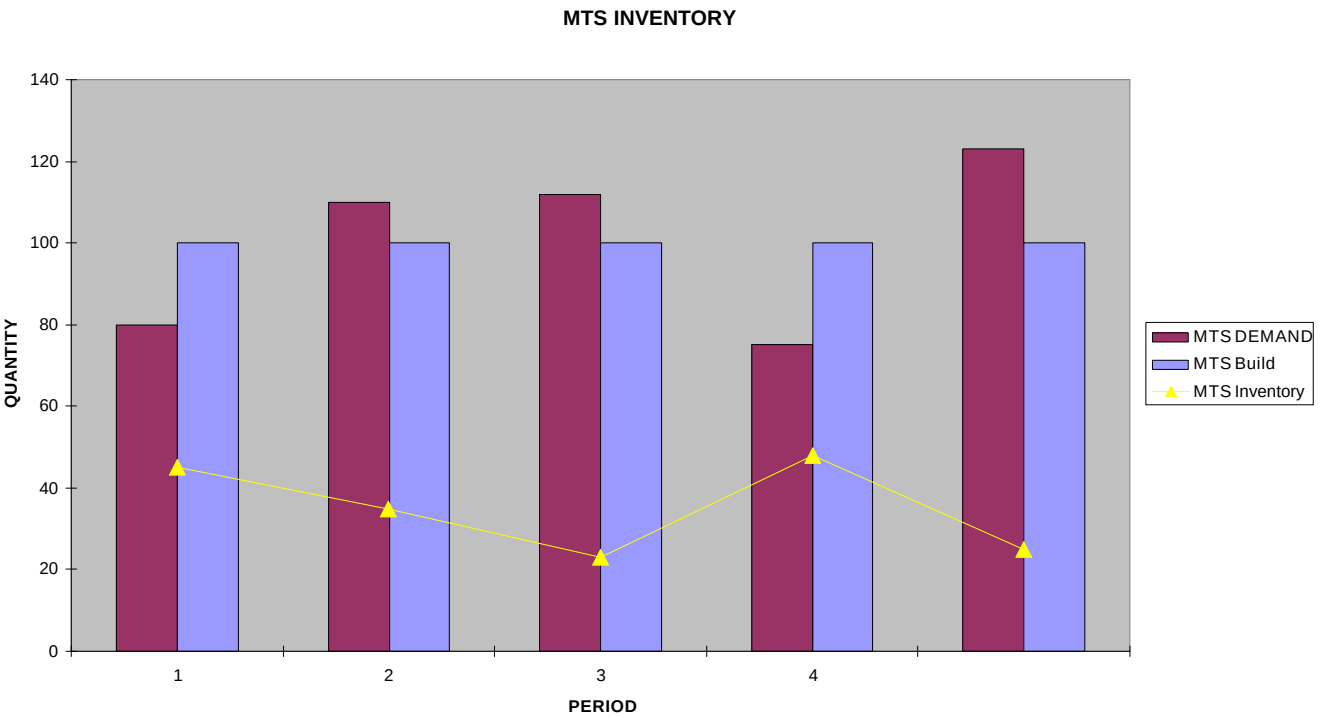

Buffering with lead time is also fairly common in manufacturing facilities. In sales and operations planning terms, this is often referred to as having a variable customer order backlog. What this means is that in order to keep capacity leveled, each customer order gets its own promise date sometime in the future according to when capacity is available to perform the work. If we have sold more than was anticipated, then the customer will experience longer lead times than previously. If we have sold less than was anticipated, then the plant will be able to pull some customer orders forward, and the overall lead time will become shorter. Varying lead times is an acceptable condition in many industries, although there is a danger that in competitive times the company will lose business to a less capable competitor who can simply promised shorter lead times because they're not very busy.

Buffering with labor is much less common than it once was. In this strategy the facility would staff to be able to satisfy the higher level of expected business, and would try to find ways to keep workers busy when business is not at this high level. From a lean perspective, this is a very attractive strategy based on the fact that is very important to engage all levels of the company in the activity of eliminating waste from all processes. These activities could easily be added to the schedule when customer sales are not at a high level. However, from a cost perspective, this strategy does not often make the financial management of the company very happy, because productivity will only be as good as it could possibly be during peak periods, and will of necessity go down during slower periods. Part of the attraction for much offshore processing currently revolves around the fact that the company can afford some extra labor, which is relatively cheap in these areas, and gives the flexibility to meet customer demand, even when it is variable, with relatively low investments in inventory.

11. Finished goods supermarkets for a product family

Ideally, we can level production in the value stream both from a volume perspective, meaning that we make the same quantity of product each interval, plus we can level it for mix, meaning that we make the same products during each interval. Nearly perfect examples of this type of leveling can be found in many automotive final assembly plants. Such plants have a supermarket buffer which consists of some of all items produced. In most job shops, however, there are many items which have small and infrequent demand, making it very difficult to completely level for mix. This is because we either sell a single unit at time and we don't realize a sale each interval, or the items are expensive and we cannot afford to have inventory of such items sitting for long periods of time with no demand. Either of these situations suggests that the best we can hope to achieve is to level the plan for volume, and be prepared to build the actual mix as it comes in from customer orders. If we are buffering for volume, then a small number of high volume parts should be kept in a supermarket of finished goods. The number of items and the quantities can be determined by an analysis of demand patterns, with the items' desirability for being in the supermarket based on the frequency with which the item shows up in the sales order mix. An item that shows up in every interval is much more desirable than an item which has the same volume but is ordered erratically. For an item to be an effective buffer for volume in the intervals that have more demand than planned production, the item must be in the sales

mix for that interval to be of any help. Having the item in inventory ahead of time is only helpful if it can be sold in a period of high demand, reducing the production requirement for that interval.

Demand by MTS and MTO items

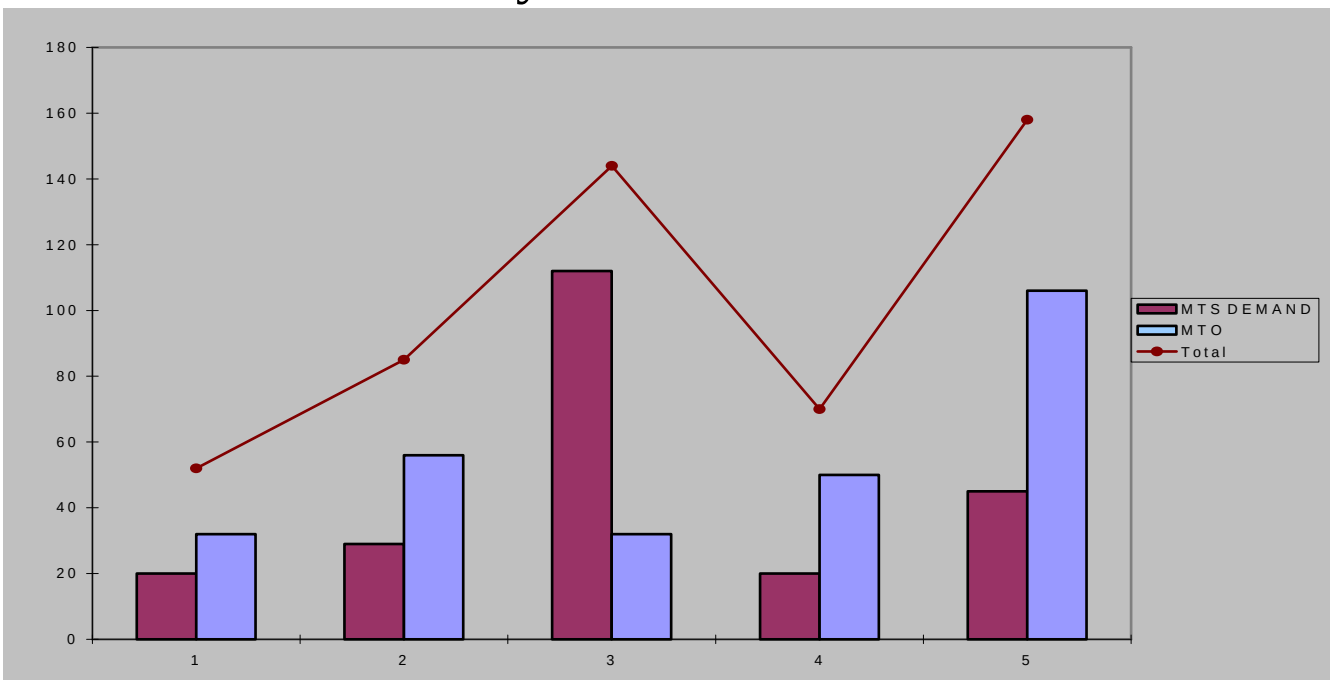

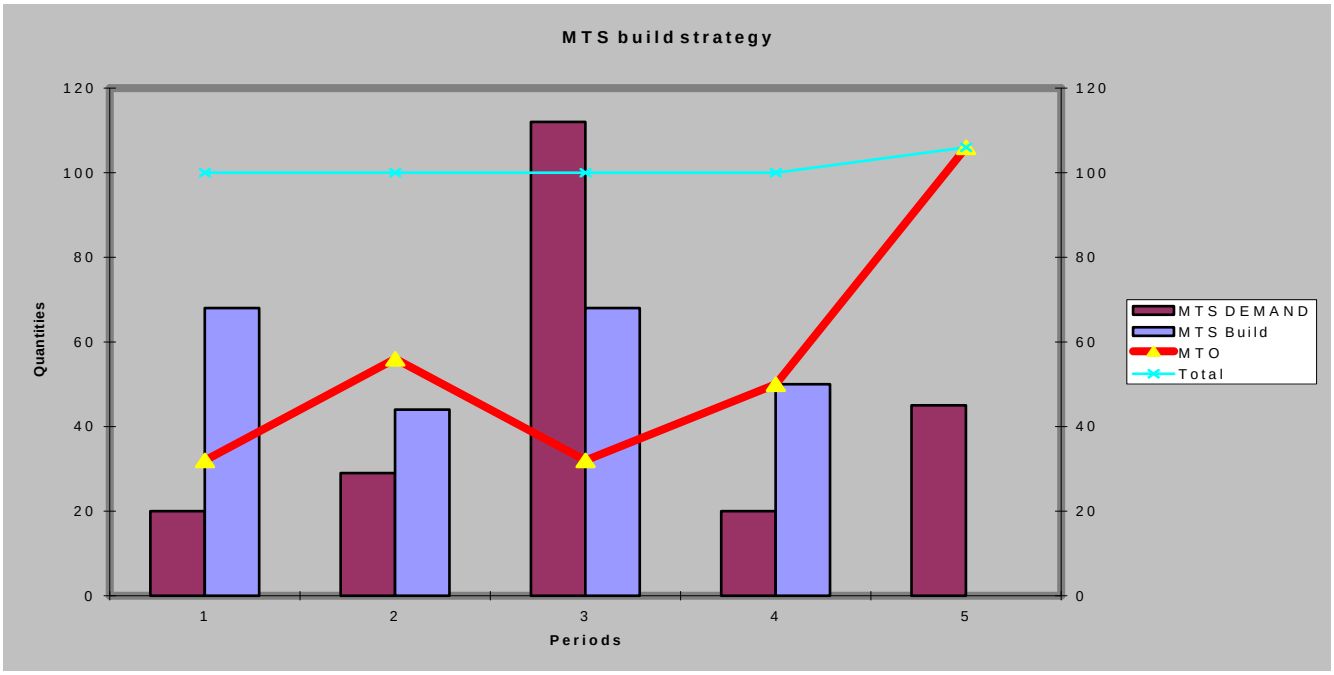

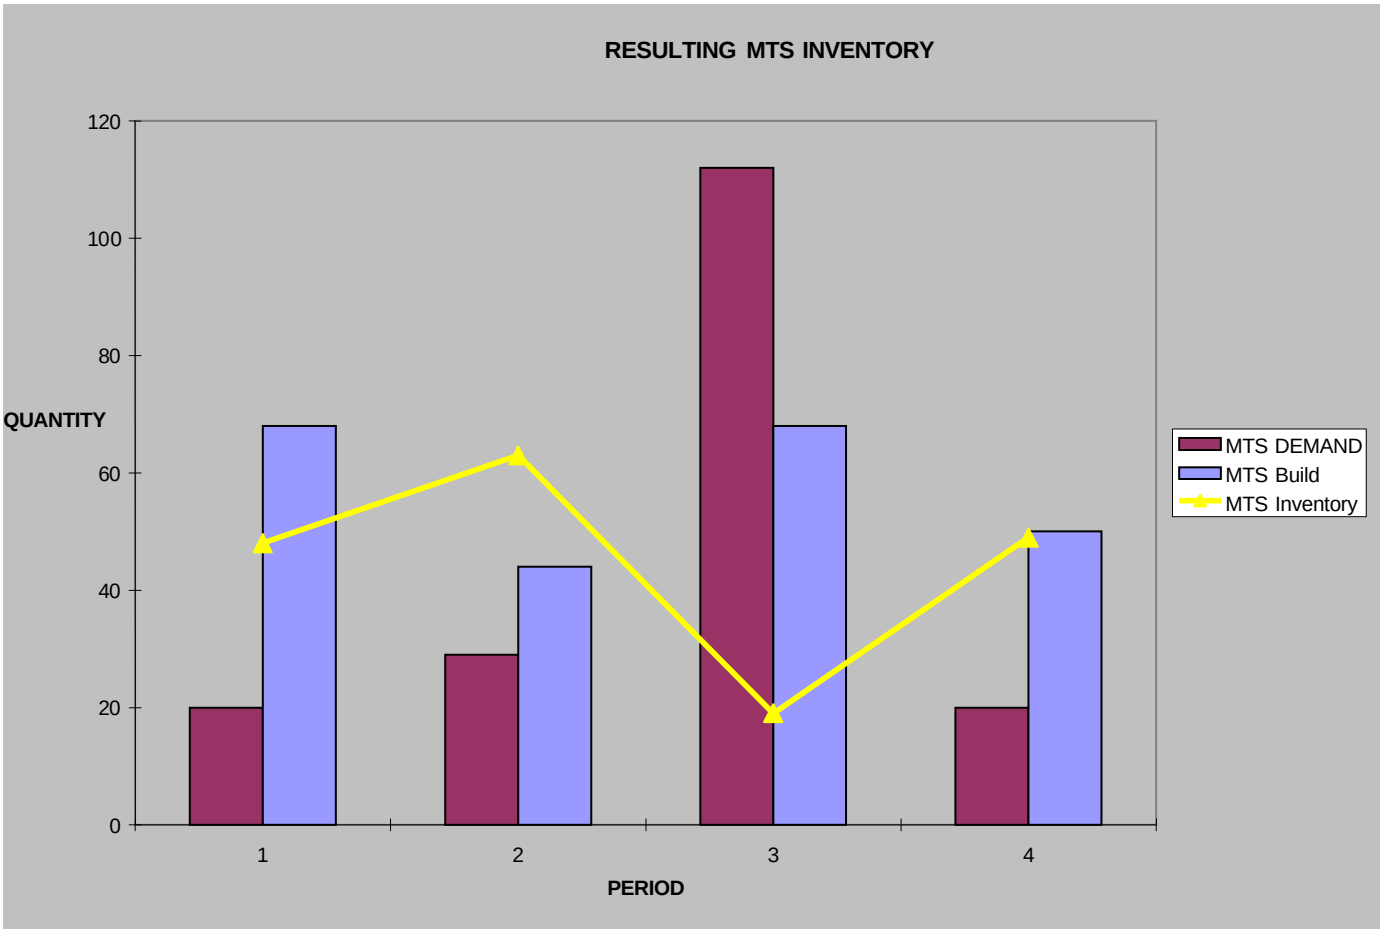

Which items should be in the supermarket?

Obtain a profile of actual customer orders for the last six months or one year. The data should be in weekly buckets or smaller. Group the orders into the desired interval for the pacemaker process. Then total all of the orders for the interval, since the idea of interval is that a part will be set up and run only once during the interval, therefore the quantity that must be produced needs to equal the total of all orders for the interval.

Perform an S&OP to determine desired production rate for an interval. This production rate will be compared to the customer order demand for each interval so that we get a general idea of how far from normal production the actual demand patterns are. Traditionally, this has been captured by using a standard deviation or a mean absolute deviation calculation. This deviation was then translated by a safety factor into an amount of safety stock which needed to be carried beyond the normal lot production stock in order to achieve a certain level of customer service. In the strategy that we are pursuing, we are not necessarily interested in trying to statistically protect an inventory level, but are instead trying to provide 100 percent customer service by either building the item to order or having the item in stock as needed. We will approach the problem by dividing our parts

into two distinct segments. The segment's are items which have erratic and relatively small demands, which will be designated as make to order parts, and items which have high, consistent demand, which will be designated as make to stock parts.

How do we determine which items fall into which category?

In periods where the demand is greater than the production rate, begin identifying the highest volume item, and subtract the demand for this item from the total. The remaining total will be the items which need to be made to customer order, with the idea being that we could have started the interval with the high volume item already in stock. Repeat with the next highest volume item and so on until the total is less than the production rate.

Scan all subsequent periods to see if this net total is less than the production rate. If not, add another item to the list.

These are your supermarket items. All other items will be built to customer order only.

How much should be in the supermarket: minimums and maximums.

There is no target level for this type of supermarket: it is there to smooth demand into production, and therefore will grow and shrink as demand changes. Minimums and maximums should be monitored as a way to check on the overall sales volumes in the value stream. Reaching a boundary should trigger an S&OP review. Reaching the minimum (usually zero on hand) indicates the need to work overtime to satisfy customer demand. Reaching the maximum (supermarket full) indicates that production must be run at less than established capacity.

The first chart below illustrates the concept of managing inventory within a range. Even though the inventory level is changing quite a bit over time within this profile, it has consistently stayed within the expected range as set by the upper and lower control limits. This means that there is no reason to change the production rate for this value stream, unless some significant change in demand is known in the future, such as a seasonal demand change.

The second chart shows an inventory that has gone lower than anticipated, due either to sales larger than expected, production less than expected, or some combination of the two. At this point, it would be prudent to review the expected sales rate, and perhaps make an adjustment to the production rate.

The third chart shows how all you could represent a plan which calls for inventory to be reduced over the course of time. Again, since we still expect variability in demand, we need to manage this reduction by thinking of inventory as fluctuating within a range of expected values, and not react to variability by changing production rates unless inventory goes outside of the expected range.

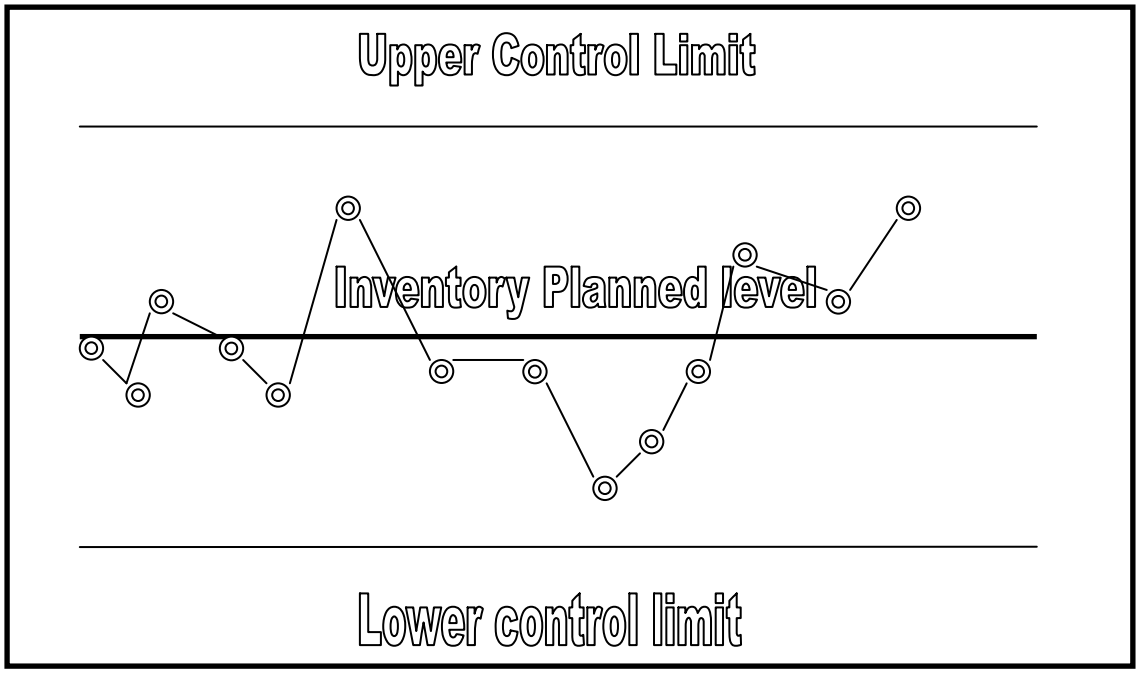

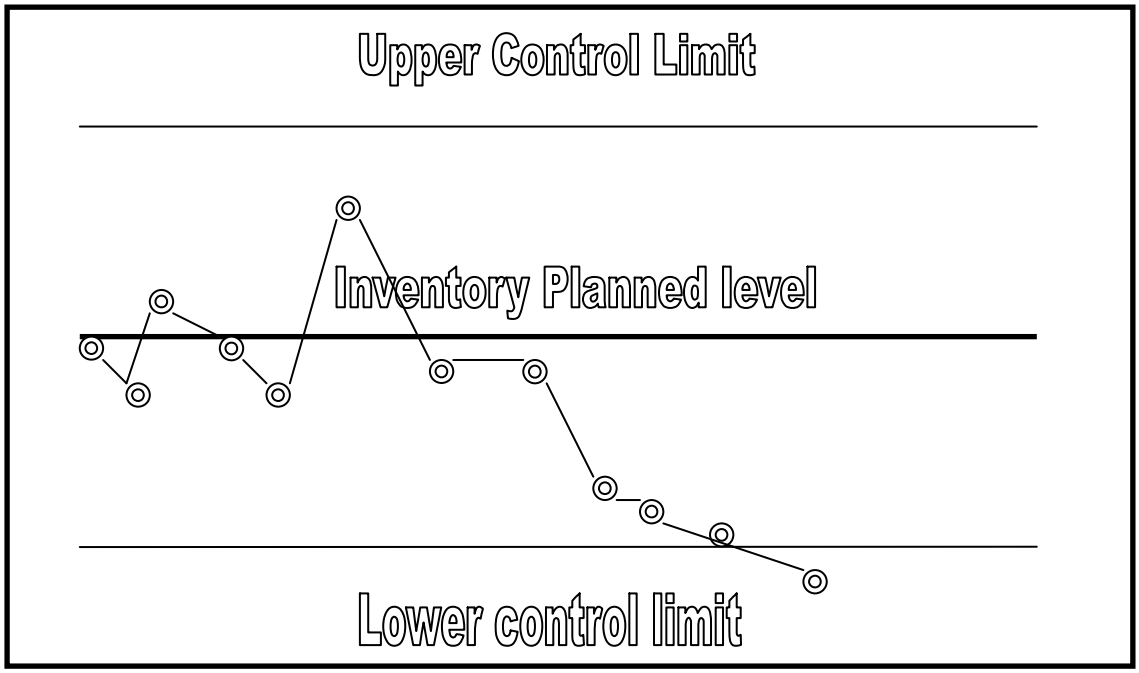

2004 Copyright Future State Solutions, LLC

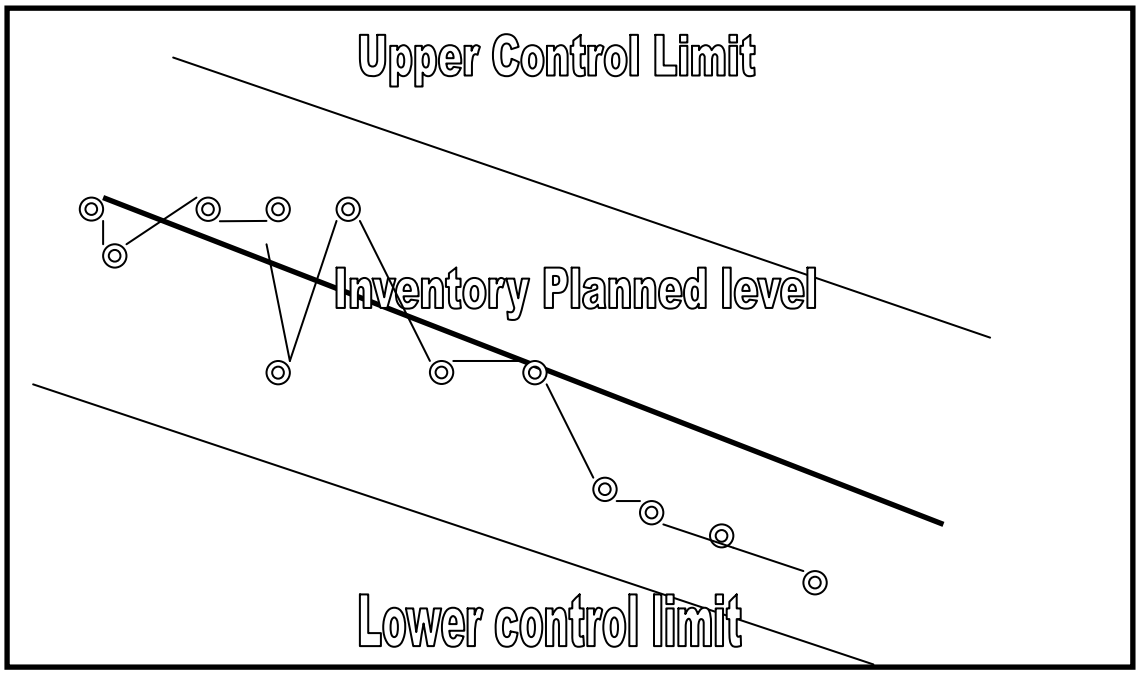

How much should be in the supermarket: starting size and maximum size?

One of the most important things in Lean is to have "organizational discipline." This means that we need to continually recognize what the current state or situations really is, and respond to this in a way that continues to eliminate waste in the future. With regard to sizing supermarkets, what this suggests is that the starting size of the supermarket must be simply the amount of inventory that is currently present. Having an inventory starting point which is less than the amount that is currently on hand suggests that some inventory will have to be stored somewhere other than the supermarket. This immediately fosters a bad habit that says that if we happened to produce more than the supermarket will hold, it is okay because we have another place to store "overflow" inventory. Therefore you should simply start with a Supermarket which holds all of the current inventory. If your calculations show that you have more inventory than you need, simply remove some of the replenishment cards as they come back and don't replace the inventory.

2004 Copyright Future State Solutions, LLC

12. Scheduling the value stream using these strategies:

Gather the customer orders for the interval being scheduled. These should be summarized by item number so that each item is only scheduled once during the interval. Try to schedule the customer orders first, as if no inventory existed (as though the product was shipped already and a pull signal was received). If the customer orders have used up all the available capacity, the schedule is complete and can be sent directly to the pacemaker. If capacity is still available, add a stock item to the schedule, which will be used to replenish the supermarket. Available space in the supermarket should be represented by replenishment cards held by the scheduler. Please note in this partial inventory supermarket that the objective is not to replenish the supermarket, but is instead to build the items which the customer has asked for. The supermarket is only replenished when there is capacity available to do so. If not enough capacity is available, pull from stock and don't replace in this interval. Once the production schedule has been completed, it should be sent to the pacemaker operation and used as the basis for measuring pitch in the value stream. A shipping schedule should also be sent, designating whether the shipments will be made using products which are being produced by the value stream or product which is currently in the supermarket.

Conclusion:

Buffering value streams can be a complicated and difficult task, depending on the complexity of the value stream. In shops which don't have the repetitive characteristics of a value stream such as an automotive assembly plant, the type and placement of the buffer is extremely important to being able to operate a lean value stream. An integrated strategy which ties together a production rate [Takt time], inventory levels [supermarkets], and scheduling of customer orders for 100 percent on time delivery promises to have tremendous power and eliminating waste. Using the process above to answer the questions of Takt time, Finished Goods Strategy, Schedule One Point, and Pull should provide a solid basis for even the most difficult value stream to begin to make major strides in their Lean journey.

2004 Copyright Future State Solutions, LLC