

A Simple Simulation of Material Management Methods

By Kevin Kohls

Executive Summary

Three material management methods – MRP, Min/Max, and Demand Driven MRP – are evaluated in this paper using a simple simulation. DDMRP was determined to be the best planning method to handle high variability environments more effectively than the other methods. Additionally, in high variability environments Min/Max is a superior solution to forecast driven MRP.

Global manufacturing and supply chain landscape is increasing in complexity and volatility. Higher degrees of variability are the norm.

This paper focuses on planning aspects of the methods only. That means the timing and manner in which supply orders are generated with the assumption that those supply orders will come in when expected. Execution (the real time management of open supply priorities) between the methods was not compared since MRP and Min/Max no execution capability. DDMRP would further distinguish itself against MRP and Min/Max in any environment with its integrated execution and priority management methods.

Basic Methods

In this paper, we'll examine three basic methods and effectiveness using a basic simulation. Those methods are:

Material Requirements Planning (MRP). There are different ways to implement MRP, but the most common element is the reliance on a forecasted demand signal to generate net material requirements. The forecast is reconciled into a master production schedule. In this analysis, we'll be looking at a basic MRP system driven by forecast. The system will be aware of spikes in orders that are coming due to an actual (or forecasted) spike. Orders are placed once every 5 intervals (5 days that make up a standard work week, for simplicity). Safety Stock will be used to cover any higher than expected demand variation. However, the system is not aware of actual demand at the time of order release. Any adjustments to MRP are not automatic; they will have to be done manually.

It should be noted that MRP has greater visibility than Min/Max. MRP systems understand the Bill of Material, process flow, and the location and status of inventory. This information is invaluable to understand how to connect demand and supply signals across the manufacturing system, especially larger systems with lots of shared components.

Min/Max. Min/Max systems are also developed and executed in many different fashions. Again, a basic method will be used. As On Hand stock is drained, the remaining amount is compared to a minimum (Min) Value. Once that level is

reached, a single replenishment order is generated for difference between the maximum (Max) value and the amount of stock left On Hand. The simplicity of Min/Max becomes apparent – there is only one order in the system at any one time. On Hand levels will continue to drain until that order arrives. In scenarios with no variation, the On Hand Buffer will reach zero just as the order arrives. Min values are based on the predicted daily usage and the ordering lead-time. Max values are based on the same variables, but can be expanded if the Minimum Order Quantity is larger than the original Max calculation. A factor can be added to increase the size of both Max and Min, but are not used in this simulation. As with the MRP method, Safety Stock will be used for the same reason – to cover for higher than expected demand. Min/Max is demand driven.

Demand Driven MRP (DDMRP) DDMRP has the advantage of visibility into the Bill of Material, process map, and inventory location and status. However, it does not use forecasted demand in the order generation equation. Instead DDMRP uses something called the “available stock equation.” The available stock equation can be simply stated as (On-Hand + On-Order – Sales Order Demand Allocations due today, due in the past plus qualified spikes). The available stock equation yields a position against a pre-determined zonal distribution (Red, Yellow Green) within the “buffer” or stock position. When the available stock position is in the yellow zone then an order recommendation happens for an amount up to the top of the green zone. In DDMRP, the selection of these buffers is perhaps the most important design element. For simplicity’s sake, our simulation uses only one buffer, and it is assumed to be a position for an end item. It takes seasonality into account, and changes ordering policies based on an actual sales order spikes.

While fairly straight forward, I will not repeat formulas for levels that can be looked up in *Orlicky’s Material Requirements Planning* 3rd Edition (Ptak and Smith, McGraw-Hill, 2011).

For MRP and Min/Max, the following equations are used:

1. **Z:** NORMSINV(Service level): Our desired Service Level is 99%, so $Z = 2.33$.
2. **Safety stock:** $\{Z * \text{SQRT}(\text{Avg. Lead Time} * \text{Standard Deviation of Demand}^2 + \text{Avg. Demand}^2 * \text{Standard Deviation of Lead Time}^2)\}$
3. **Re-order Point (ROP):** $\text{Average Lead Time} * \text{Average Demand} + Z * \text{SQRT}(\text{Avg. Lead Time} * \text{Standard Deviation of Demand}^2 + \text{Avg. Demand}^2 * \text{Standard Deviation of Lead Time}^2)$

Effectiveness

The objective of all these materials management methods is to protect availability. Protecting availability protects flow. Throughput Accounting (or Lean Accounting -

- see The Conflicted Middle: Lean Accounting by the author) emphasizes Throughput (Flow), Inventory, and Operating Expense as the basic metrics that drive Net Profit (NP) and Return on Investment (ROI). To that end, it's critical to maximize the flow of parts to fulfill the customer's demand. However, any parts that are manufactured above the demand of the customer will become inventory. Any Work in Process (WIP) generated to maintain that flow is also considered part of Inventory. The right level of Inventory is financially important. Companies should hold the minimal amount of inventory that ensures flow up to customer demand.

Many companies pursue cost cutting as a means of improving profits. However, if a reduction in Operating Expense disrupts flow to the customer, or causes a major upward shift of Inventory, then a negative impact could be generated on Net Profit or ROI. Thus, the minimal amount of expense that allows the system to maximize flow while maintaining or reducing inventory is desired. Freeing up resources from performing a particular job, such as expediting or planning, is not considered a reduction in cost in this evaluation. The desire here is to reduce the amount of time planners and expeditors spend on these tasks, as well as to reduce the amount of confusion and waste involved with the process. Freed up resources can be applied to other tasks that improve safety, improve quality, increase throughput, etc. Reducing the amount of time used to create manual orders, change inputs or perform expediting is a reflection of the amount of resources used in material management. If we can reduce these factors that generate this waste (overtime, poor decision making, chaotic environment, multi-tasking, etc.) then the associated material management process will be improved.

The effectiveness of the three methods will be determined by a basic set of measures that relate to the desired objectives listed above that drive maximum profitability and ROI. These measures are:

- Inventory – Inventory is inclusive of both Work in Process (WIP) and On Hand buffers. Lower Inventory frees up cash and costs less to track and maintain, and can improve ROI.
- Inventory Coverage – The amount of Inventory needed in the system to theoretically ensure that no shortages are created. It is equal to the Inventory definition above plus the largest shortage of parts to the customer over the timeframe being evaluated. Lower Inventory Coverage is desired, since it demonstrates the correct amount of inventory is being held.
- Service Level – The percentage of times that an order arrives when promised. Also called On Time Performance (OTP). A higher Service Level is desired, since it maximizes flow.
- Lost Parts – the total number of parts overall that are shorted to the customer. It is called Loss Parts because the assumption is that any part that does not make it to the customer will be lost, and that work will go to a competitor. It represents the potential loss in throughput and profits that comes from these shortages. A smaller number of Lost Parts is desired, since it represents a loss to Net Profit that cannot be recovered.

- Warnings – the number of the number of warnings generated over a particular time frame. A planner/expeditor typically prefers a small number of effective warnings.
- Part Impact per Warning – The number of parts impacted by a warning. Warnings that are generated for a small number of parts gives the planner/expeditor a fighting chance to correct the potential shortage. Overall, fewer warnings with a small number of impacted parts is desired by planners and expeditors.

Note that efficiency is not listed in the measures above, since it difficult to demonstrate a correlation between local workstation efficiency and higher profitability. From a materials standpoint driving efficiency at local resources often distorts usage patterns and increases expedites and shortages as parts are directed to places they do not need to go.

The tool used to evaluate these methods will be a simple simulation that will utilize three levels of variation – none, minimal and extremely high. The results will demonstrate the impact that will occur on the measures listed above. By definition a simulation is not meant to be an exact representation of the system being analyzed, but an effective approximation that is useful to those who analyze the results, draw conclusions, and generate recommendations. It has to be approximately correct, not precisely wrong.

It is with this in mind that the simulation described in this paper was created. For ease of use and accessibility, the simulation was created using Excel, utilizing random number function and macros. Experienced simulation engineers will cringe at this simplification, but the level of detail is sufficient for the understanding the comparison.

A real company's situation and information can be used as a basis for this analysis. The purpose of this simulation is to create insight, not a detailed analysis of the future system. A company may decide to pay for the expertise and tools to create such a simulation before a pilot – this effort is not an attempt to replace that type of analysis. Rather, the effort is to increase the speed of an implementation, allowing those impacted to move as quickly as possible from Idea Phase to the Sustain Phase.

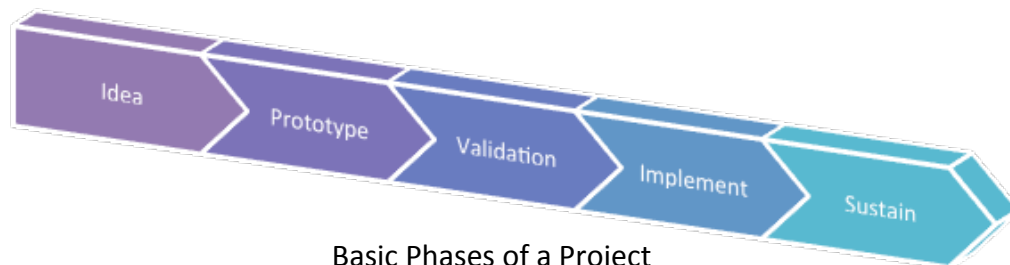


Figure 1 - Implementation Phases of a Project

For one mode of analysis, the simulation will be run using the rules for a particular method with no intervention. Warnings will be ignored, manual intervention is not allowed, and expediting cannot be performed. In another mode (called the Game Mode), the user will play the role of a planner who can intervene to correct a situation that is occurring. The Game Mode will be discussed in another paper.

One of the most important aspects of a simulation is how it handles variability. This is why the Random number function (or RAND(), in this case) is used in as a basis for this variation in our Microsoft Excel™ simulation. A simple spreadsheet analysis that does not consider variation is not the appropriate tool for this type of analysis. The strength in these material management methods will be based on their ability to handle change and variation.

For the set-up of this simulation, the user can define the input data variability. In Figure 2 below, the Average Daily Usage has been characterized by the expected amount of usage over a particular time frame. In this example, an ADU of 8 is the most frequent (Blue Line) but the longer tail on the right results in an overall average of 9.4 units per day, with a Standard Deviation of 3.4. The red histogram bars represent the actual ADU's that occurred during the simulation. This particular run was slightly lower than the designed average at 9.1 units per day, with a Standard Deviation of 3.1.




Figure 2 - Creating Variation for Average Daily Usage

Figure 3 shows the Lead-time variability. The variability shown is one customer's typical experience. For example, there is some chance the ordered material will arrive early, but it typically arrives on the expected date, which is 2 days in this example. The long tail ranging to the right is the variation that causes the most havoc. An order arriving 5 days late can easily cause stock outs in this system.




Figure 3 - Creating Variation for Lead Time

Another type of variability is shown in Figure 4, seasonal variability. Here, month five has the highest demand, 60% higher than the average demand. Months one and twelve are only 70% of the average.




Figure 4 - Creating Seasonality

The typical numbers are entered in the spreadsheet under the Input Data area, shown in Figure 5. Values the user enters are in grey; all the other values (in white boxes) are generated by the spreadsheet based upon DDMRP formulas. The ADU number represents the expected average, but the actual ADU performance is generated by the variation data used to generate Figure 2. The MOQ (Minimum Order Quantity) and Batch (the minimum batch size or minimum order multiple) are also entered here. The other variables are unique to DDMRP. LT Adjust is used to compensate for high Lead Time Variability, while the “Var. Adjust” is used as an overall adjustment for variability in ADU and/or Lead-time. For our purpose, the Adjustment factors are set to .5 unless otherwise noted. Based upon these data, zones for the buffer levels are calculated. The threshold for a “Spike” in orders is also defined.

ADU	ASRLT	LT Adjust	Var. Adjust
16	5	0.5	0.5
MOQ	Batch		
1	1		
		Top Spike Trigger	
Green	40	180	
Yellow	80	140	
Red Safety	20	60	30
Base Red	40		
Iterations	40		

Figure 5 - Input Data into the Excel Simulation

Comparison of Three Methods

For analysis and game purposes, the following three methods will be simulated and compared at their most basic level.

Material Requirements Planning – MRP

The MRP method will relay on the forecast demand and create an order every 5 intervals. The actual demand is not considered, and therefore will have no impact on the number or quantity of orders generated.




Figure 6 – Terms Used in this Simulation for MRP

Min/Max

Here, an order is created whenever the Min level of inventory is reached, and that order is sized to bring the inventory up to the Max level. With this method a Safety Stock (defined earlier) is also used to compensate for an upward variation in demand.




Figure 7 – Terms used in this Simulation for Min/Max

DDMRP

The method used here is as defined in *Orlicky's Material Requirements Planning 3rd Edition* (Ptak and Smith, McGraw-Hill, 2011). Available Stock is monitored, and another order is created when this stock level drops below the Top of Yellow (TOY). The order size is the Top of Green (TOG) less the current Available Stock. Levels here can change based on seasonality. Batch sizes and Minimum Order Quantity are considered when setting the level sizes.

It's important to note that one facet of DDMRP is the ability to prioritize based on warnings, order frequency, etc. This can be a flag sent to a planner for their resolution. This analysis can also result in an automatic level change for future decisions. This ability is not considered in this analysis.

The simulation itself is a local analysis of one work center's performance in matching customer orders with parts sitting in on hand inventory.. Parts incoming to this "Shipping Buffer" have been ordered for a particular amount and arrive in a certain amount of lead-time. If multiple work centers are linked together, stock outs will be generated throughout the chain. For our purposes, we'll assume that the high stock outs for one area will lead to higher stock outs for the system as a whole. This assumption will be challenged in future analysis.




Figure 8 - Terms used in This Simulation for DDMRP

Comparing Methods

Since the methods are inherently different, charts that will be utilized to look at two common facets – the number of parts on-hand and the amount of WIP (or Available Stock in the DDMRP case) being held. When the On Hand level reaches 0 with unmet demand then a stock out condition will be declared. For this comparison, the amount of On Hand inventory can also reflect unmet demand. Thus, On Hand inventory can go below zero. The assumption here is that, despite the stock out, the order will not be cancelled.

For MRP and Min/Max, On Hand and WIP are tracked. For Min/Max On Hand Buffer, the three levels are shown – Safety Stock in Orange, Min in Red, and the Max level is

Green. An order is created when the inventory goes below the top of red, and the order size is the Top of Green (Max) minus the currently On Hand inventory level.

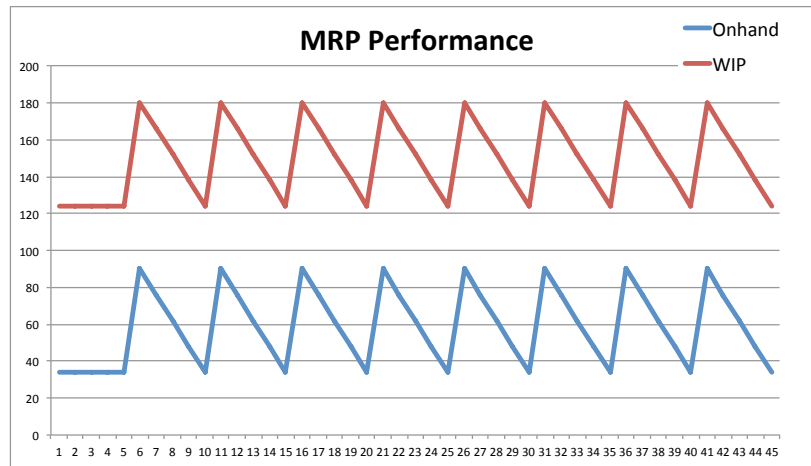
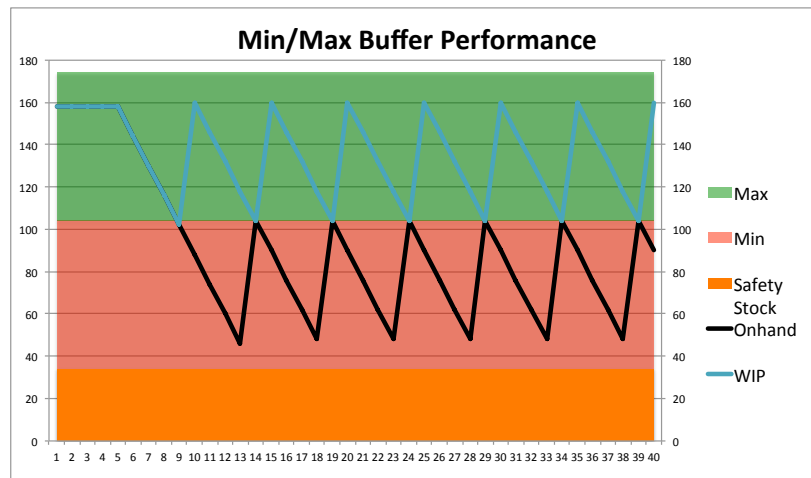


Figure 6 - MRP Inventory Levels for WIP & OnHand



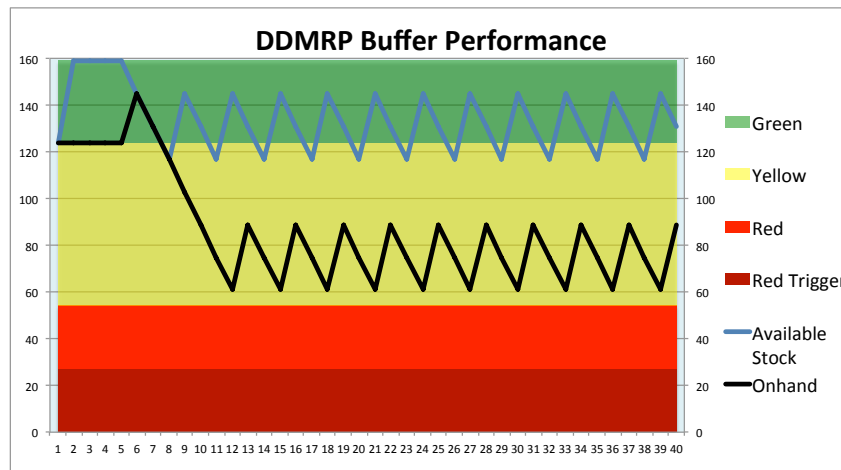


Figure 7 - DDMRP Inventory Levels for Available Stock & OnHand

For DDMRP, the graph looks slightly more complicated. There are two stock levels, Available Stock, and On Hand stock, as described previously. A stock out occurs when the On Hand stock drops below zero. Available Stock is used to determine the need for an order and its size. The buffer levels described previously are also shown in the graph. The darker red trigger zone causes an alarm to be issued with the intent of notifying the planner to expedite. In the simulation, this warning is ignored.

An order is created when the Available Stock reaches Top of Yellow (TOY). The order size is Top of Green (TOG) minus the current Available Stock. The light blue area (not apparent in this chart) demonstrates when the Available Stock is above the TOG.

Basic Simulation Analysis

Minimum Variation For the first comparison, all three methods were run with minimal variation. This is to determine if the method works under a “best case” scenario. A demand schedule that varies by timeframe is used to demonstrate seasonal variability. This is forecasted in both DDMRP and MRP. In the Min/Max method, seasonality variation will be absorbed by both the buffer and its Safety Stock. In the higher demand timeframe, more orders will be generated. Fewer orders will be created in low demand periods.

A large order spike will occur at the peak of the seasonal variability. This will be forecasted in both MRP and detected by DDMRP. Min/Max assumes Safety Stock will cover the spike.

The ADU average and the lead-time average are maintained, and variability occurs around these numbers.

Very High Variation The second comparison will test of effectiveness of the three methods in a highly variability environment:

- The actual demand will be higher than expected – the ADU was 10.2 actual versus the 8 units that were expected.
- Lead-times are also higher than expected, averaging a half a day longer versus what was expected - 4.5 days vs. 4 days.
- A planned order spike will occur on Day 27 for 45 jobs. This is a spike that MRP has in its planning. Min/Max will rely on Safety Stock to compensate for this spike, and DDMRP will include this spike in its Available Stock equation.

This results in a part usage over the projected lead-time to be much greater than expected. The actual units over lead time ends up being 45.7 vs. the expected 32. As a result, usage over lead-time is 43% greater than expectations.

These input variables were obtained by incrementing ADU and Lead-time until all three methods created shortages for at least four out of the ten iterations. No corrections were made despite warning messages being issued. For this analysis, both MRP and Min/Max generate warnings when Safety Stock is penetrated. DDMRP issues an alarm when the On Hand level (On Hand Red) is penetrated.

Ten runs were made, with all three methods facing the same variable in each run.

Minor Variation

With minor variation occurring around the expected ADU and lead-time, the following charts are generated. Seasonality and the spike exist in this scenario.

Note that all three methods work – no shortages occur, and inventories are within expected levels. There are some observations that are worthwhile. The seasonality and the order spike did cause the Min/Max method to penetrate the safety stock. MRP's WIP and On-hand almost matches, indicating that most of the stock is in the On-hand buffer. DDMRP's Available Stock tracks Seasonality fairly well, it's clear that orders based on Available Stock are large as time heads toward the peak season (reducing the number of potential changeovers) and are smaller and more frequent at non-peak times.

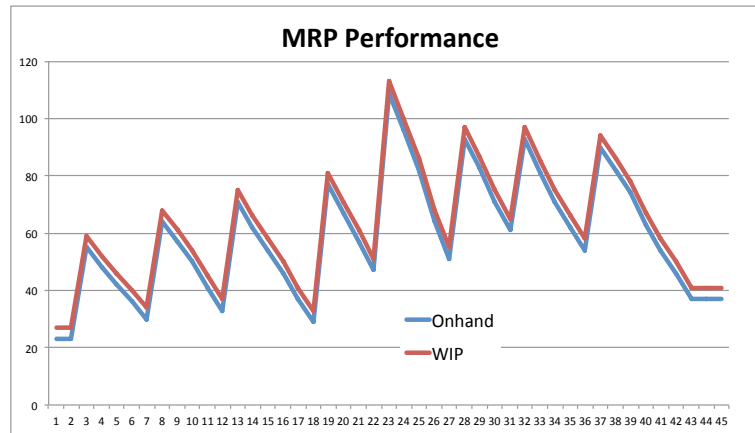


Figure 13- MRP Performance with Minor Variation

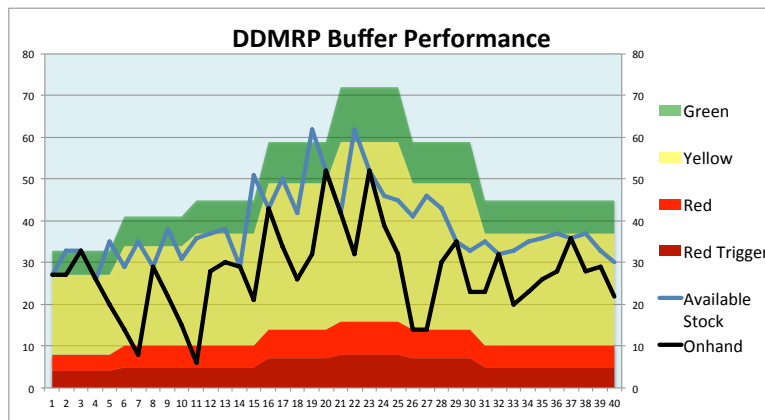


Figure 9 - DDMRP Buffer Performance with Minor Variation

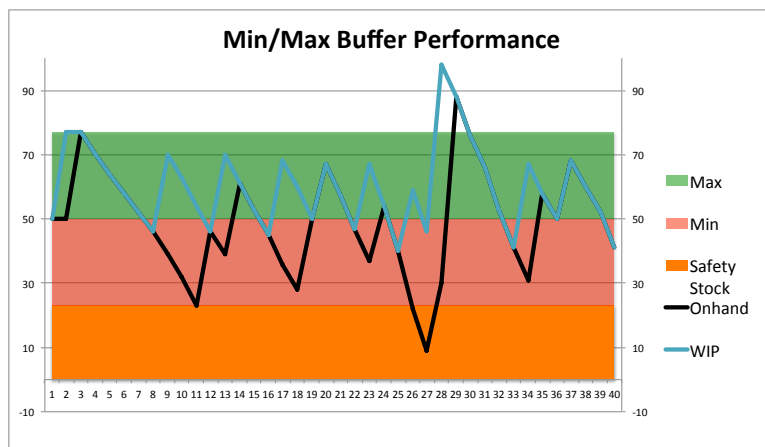


Figure 15 - Min/Max Performance with Minor Variation

High Variation Comparison

The previous comparison gave us an idea of how the three methods work. We also described the high variation scenario where ADU and Lead-time are higher and more variable than expected. As before, we let the systems run without intervention to understand the performance of the each method.

The first area is Inventory, which measures the amount of stock in the system over the timeframe analyzed. It does not take into account shortages. Here, the MRP system appears to be significantly better than the other two methods. It's important to note that the amount of inventory was still going down as the iterations ended. This has implications for cash flow and the cost of inventory – companies that find themselves in financial trouble can use this type of impact to free up cash and reduce costs at the risk of shortages. The reasons can be justified – this is what MRP told us to do.

Min/Max has slightly higher inventory performance, holding about 5% less material in it's On Hand Buffer than DDMRP over this high variation scenario. The cash flow and cost impact will be based on the impact of holding 4 additional parts.

However, for most companies the objective is to provide the customer with all the parts they need to meet their requirements. We can use the next chart (Figure 17) to understand the level of inventory required to ensure no stock outs occur. The Min/Max level jumps to 100 parts, while the DDMRP amount is the same as total inventory held. MRP jumps up as well, but is still about 5% better than DDMRP. It is important to note that in this example, the WIP coverage was trending upward as services levels continued to degrade. At some point in the future, WIP coverage for MRP would degrade past DDMRP without some sort of intervention.

Service Levels are the next metric considered (Figure 18). This is the percentage of time a shortage occurs, without considering the total size of the shortage. Here service levels are significantly higher in the DDMRP method as compared to the other two, at 98% as compared to 91% for Min/Max and 74% for MRP. Of course, this is the analysis for one "workstation." Adding 5 workstations in series multiplies the impact of the shortages (the definition of a workstation and how they are connected can be different for each method). The change in Service Level over 5 workstations is significant and not linear. Service Level with DDMRP drops to 90%, but the impact of Min/Max and MRP is much worse. Despite this, it's clear why Lean advocates would push for Min/Max over MRP – there is a 40% improvement in Service Levels. MRP planners could not allow the system to run to this service level – they would be forced to intervene by manually creating more orders, changing the order frequency and expediting or significantly and intentionally bumping up safety stock levels to prevent the chronic and frequent shortages.

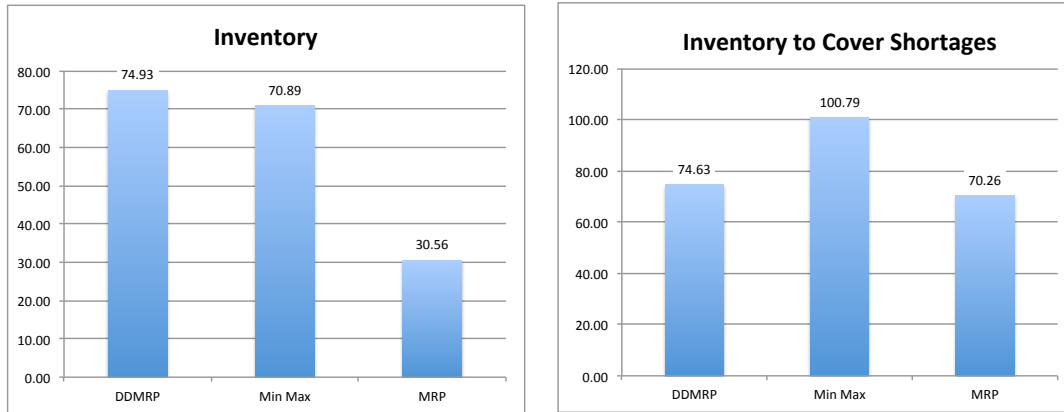


Figure 10 & 17 - Inventory & Inventory Shortage Comparison

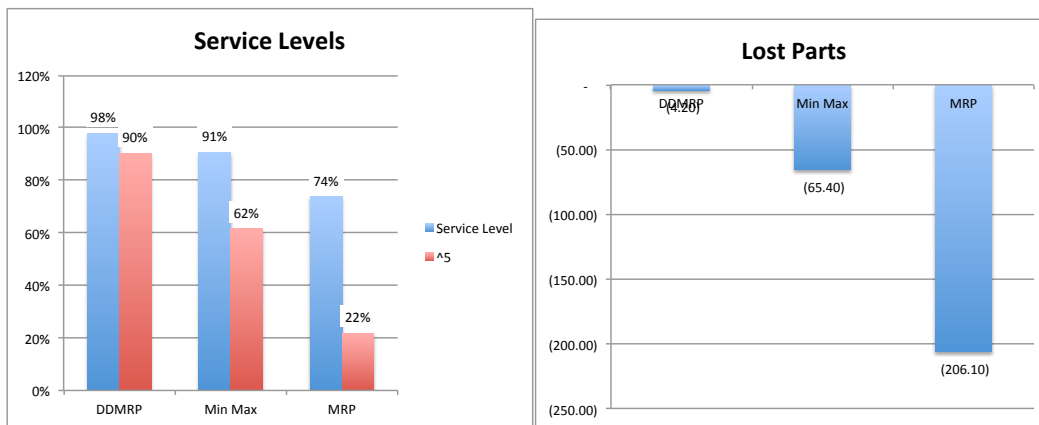


Figure 11 - Service Levels and Lost Parts

Finally, from a customer/supplier perspective, we have to consider the amount of “Lost Parts” the method has generated. We’ll define Lost Parts as lost sales that would occur from a stock out. Our assumption here is that customer would fulfill the shortage from another supplier. If this number is significant, it will have both a negative impact on profitability and the ability to maintain and/or acquire new business from this customer. Again, DDMRP has the best performance, losing only 4 parts. Min/Max would cause 65 lost parts, while MRP would lose 206 parts (Figure 18).

If we take the view of the planner, we would like a method and a system to warn us when the scenario as described in this section is occurring, so that intervention can occur to modify input data, create a manual order and/or expedite an order. From the planner’s perspective, it would be important that the system handle this high variation situation without intervention. Warnings should be infrequent and have a small impact on the customer. If we compare the three methods, we see that the DDMRP system only issued 4.5 warnings on average (Figure 19). The average

predicted impacted was less than 1 part. Min/Max issues 8 warnings over this timeframe, with each warning having the potential impact of 8 parts. MRP is again worse from the planner perspective, issuing 20 warnings, each worth 10 parts. The MRP expeditor would be almost in a constant of chaos, reacting to constant warnings and creating large orders to cover this high variation situation.

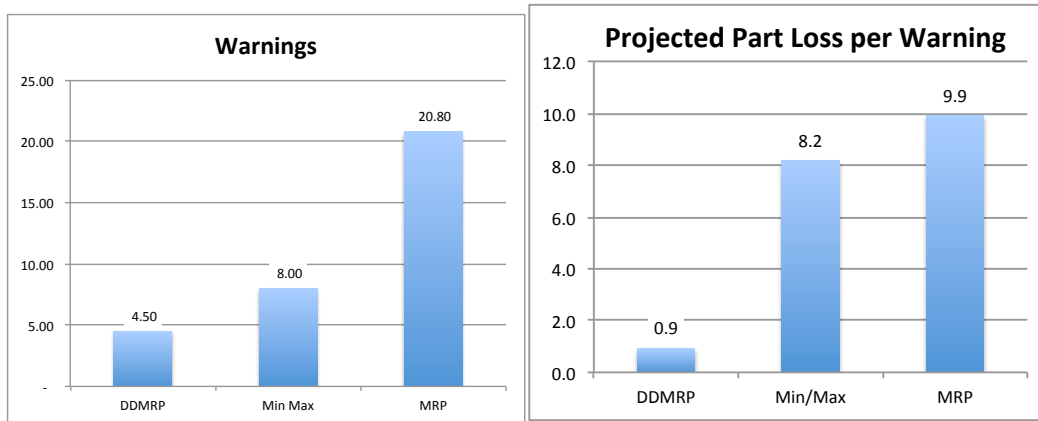


Figure 12 - Warnings and Part Loss per Warning

Conclusions about the Three Methods from this Simulation

1. Each of the method can work with minimal variation in demand and lead-time.
2. MRP is very sensitive to demand and lead time variation. This is a known and pervasive effect called “Nervousness.” Running MRP over time required a significant amount of intervention to try and maintain high Service Levels.
3. Min/Max is a viable method if spikes do not occur frequently, and if minimal expediting has to done because of variation. It is a preferred method over MRP for most metrics assuming the supplier can deliver in time with the frequent orders. In larger environments, however, visibility becomes more limited.
4. DDMRP is a superior solution in a high variability environment, especially those where demand and lead times are higher than expedited. While robust, shortages can still occur if warnings are ignored. The system requires much less attention than other methods, but may more time to set up than a Min/Max method.

Alignment with Experience

While it can be dangerous to start off an experiment with a preconceived viewpoint, it is difficult for anyone with extensive experience in improving operations performance not to see the advantages and disadvantages of each one of these methods. The results from MRP align well with my experience, and I see a large

number of expeditors, planners, and production management involved in touching the system to try and improve performance. Most areas end up using another method that takes data from MRP and uses it to make decisions. The most common are Excel™ spreadsheets and Access™ databases. I have also experienced companies who have gone from a manual Min/Max system to MRP in order to reduce headcount and cut costs. They quickly see a significant drop in performance, ending up with a large amount of inventory for material that is not in demand, and shortages of parts that are in demand. Poor performance with MRP is often blamed upon a “poor forecast.” This results in significant increase in effort and costs focused on improving the forecast. This cost can end up being higher than the projected cost savings. Finally, the errors in forecast have a “whip saw” effect, causing suppliers farther down the supply chain to deal with sudden expediting, and then a large excess of inventory.

The strength of MRP is its visibility through the system and calculation of dependent demand. Orders can usually be found in a particular area, or in a department before or after the build process. Having a system wide visibility and the ability to quickly find orders is beneficial to management, even with all of the other confusion caused by MRP.

Min/Max systems perform significantly better than MRP, and are often used for receiving. The MRP system may generate the orders, but the Min/Max buffers often dictates whether the order will end up in a stock or be sent out to the supplier. Min/Max is often used in Vendor Managed Inventories. The forecast may be used to signal to the supplier to produce parts, but the number held at the plant is determined by Min/Max. Min/Max is often used in many lean operations, but they often require local manual monitoring and ordering.

It should be noted that most companies have the paradigm that high efficiency must be maintained or improved throughout the plant. If this is the case, production of parts that are not in demand will continue regardless of a lack of orders. Worse, the production of unneeded parts with a favorable target allocation will often be run before parts that are in demand because of the positive impact on efficiency.

Middle management is rewarded for keeping their areas of responsibility highly efficient. There are no similar measures for inventory, so creating inventory and pushing it on to the next department allows managers to maintain their high efficiency measures. Tying up cash and increasing material handling costs are not on their scorecards, so they see little reason to run their areas any differently.

All of the systems that tell production to stop making parts will fail unless these measures are changed and the paradigm is shifted. This is biggest paradigm shift in Lean and TOC – make only what is required, and then stop. This is a basic requirement for Min/Max, pull systems, and DDMRP methods.

DDMRP implementations that I have been part of have improved on-time performance by significant amount, often doubling On Time Performance in a pilot

implementation in as little as a month. Inventory levels drop as well. But this method faces the same paradigm shifts that often come with Min/Max due to the directive to cease production because of a lack of orders. Operators and production management often “know better” and keep up production of parts and avoid setups, creating large inventories and filling baskets.

Speed is the next large obstacle, due to a condition called “management churn.” While one executive may strongly support a particular methodology, the next executive may not share the same paradigm, and direct the organization to change to methods that fit their paradigm. Production management on the floor often gives little attention to the change in methods, relying on their Excel spreadsheet to deliver the desired results. When management churns again, it has validated their perception that they “know better.”

Since management churn tends to happen about every two years, implementers of Max/Min or DDMRP will have that amount of time to sell and implement a successful change – to the point that the new method is the accepted methodology at all levels. This includes changing the measures used to evaluate production groups. Time consuming implementation steps must be challenged as well. Writing custom software versus buying existing software looks good from a cost perspective, but has led to both implementation and sustainability problems when management “churns” again.

About the Author

Kevin Kohls has spent over 30 years improving manufacturing systems, from working in both machining and automotive assembly plants and implementing the Theory of Constraints. His team combined automatic data collection, bottleneck analysis, Lean, and Red X (a statistical methodology similar to Six Sigma) into TIP – the Throughput Improvement Process. As an executive, his teams helped validate the design of new automotive plants using Throughput Accounting, Lean, and TOC. His efforts have generated billions of dollars of improved profitability in GM and other companies. This work and those of his peers were recognized by INFORMS, which awarded GM with the prestigious Franz Edelman Award for 2005. He has also been given both the Chairman’s Honor Award and the Boss Kettering Award for Innovation. As a consultant, he has worked with Johnson & Johnson, Magna, Dollar/Thrifty Rental Car, Blue Cross Blue Shield of Michigan, the Detroit Medical Center, Spirit AeroSystems, and many other companies.

Kevin received his undergraduate degree in Electrical Engineering from the General Motors Institute (now Kettering University) and a Master’s degree in Electrical Engineering from Purdue University. He has published his first business/comic book, *The Conflicted Middle: Lean Accounting*, which is available on his website www.kohls-consulting.com, and www.MakeThingsFlow.com. Click on the Comic Book link to purchase a copy.