DDMRP

Demand Driven Material Requirements Planning

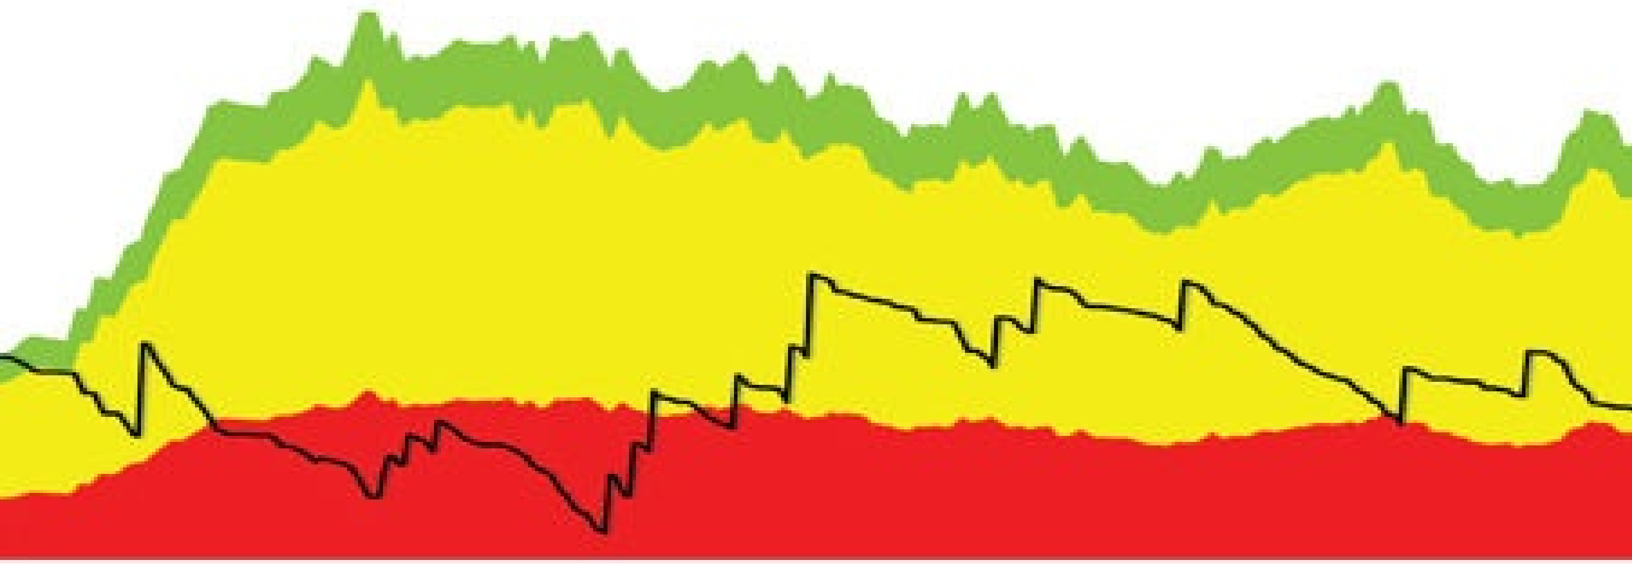

An Intuitive Proven Planning and Execution Method for Today's Complex and Volatile Supply Chains

Ptak and Smith

Demand Driven Material Requirements Planning (DDMRP)

Demand Driven Material Requirements Planning (DDMRP)

Carol Ptak and Chad Smith

INDUSTRIAL PRESS, INC.

Industrial Press, Inc.
32 Haviland Street, Suite 3
South Norwalk, Connecticut 06854
Tel: 203-956-5593, Toll-Free: 888-528-7852
E-mail: info@industrialpress.com

Library of Congress Cataloging-in-Publication Data
Names: Ptak, Carol A., author. | Smith, Chad, 1971 – author.
Title: Demand driven material requirements planning (DDMRP) / Carol Ptak and Chad Smith.
Description: South Norwalk, Connecticut: Industrial Press, Inc., 2016. | Includes bibliographical references and index.
Identifiers: LCCN 2016023395 (print) | LCCN 2016026863 (ebook) | ISBN 9780831135980 (hardcover: alk. paper) | ISBN 9780831193843 (eBook) | ISBN 9780831193850 (Epub) | ISBN 9780831193867 (EMobi)
Subjects: LCSH: Material requirements planning. | Manufacturing resource planning. | Production control—Data processing. | Inventory control—Data processing.
Classification: LCC TS161 .P789 2016 (print) | LCC TS161 (ebook) | DDC 658.5/03—dc23
LC record available at https://lccn.loc.gov/2016023395

ISBN print: 978-0-8311-3598-0
ISBN ePUB: 978-0-8311-9385-0
ISBN eMOBI: 978-0-8311-9386-7
ISBN ePDF: 978-0-8311-9384-3

Sponsoring Editor: Judy Bass
Copy Editor: Judy Duguid
Compositor: Patricia Wallenburg, TypeWriting
Cover Designer: Janet Romano Murray

industrialpress.com
ebooks.industrialpress.com

10 9 8 7 6 5 4 3 2 1

Contents

Foreword
Definitions in This Book
Introduction
About the Authors
Acknowledgments

PART 1 PERSPECTIVE

CHAPTER 1

Planning in the New Normal

The Material Requirements Planning Revolution
Evidence of a Problem
Return on Asset Performance Degradation
Work-Around Proliferation
The Inventory Bimodal Distribution
The New Normal
Summary

CHAPTER 2

The Importance of Flow

Plossl’s First Law
Establishing Flow as the Foundation
Relevant Information and Materials
The Bullwhip Effect

Summary

CHAPTER 3

Material Requirements Planning in the New Normal

What Is MRP?
Distortions to Relevant Information
Demand Signal Input
Nervousness
The Weekly Bucket
Flattening the Bill of Material
Distortions to Relevant Materials
Common Cause Variation
Delay Accumulation
Amplifying the Distortions to Relevant Information and Materials—Batching Policies
Summary

CHAPTER 4

Unlocking a Solution—The Power of Decoupling

Decoupling
Decoupling Point Buffers
Summary

PART 2 BECOMING DEMAND DRIVEN

CHAPTER 5

Supply Order Generation and Execution for the New Normal

MRP Versus Lean—What Can We Learn?
Dependence Versus Independence
Supply Order Generation (Planning Versus Execution)
Lean and Technology
Demand Driven Material Requirements Planning
The History of “Demand Driven”
Position, Protect, and Pull

Summary

CHAPTER 6

Strategic Inventory Positioning

Positioning Factors
Customer Tolerance Time
Market Potential Lead Time
Sales Order Visibility Horizon
External Variability
Inventory Leverage and Flexibility
Critical Operation Protection
Applying the Positioning Criteria
A New Form of Lead Time
Advanced Inventory Positioning Considerations
Distribution Positioning Considerations
Summary

CHAPTER 7

Strategic Buffers

Inventory: Asset or Liability?
Introducing Decoupling Point Buffers
The Green Zone
The Yellow Zone
The Red Zone
Buffer Profiles
Factor 1: Item Type
Factor 2: Lead Time
Factor 3: Variability
Individual Part Attributes
Part Average Daily Usage
Part Lead Time
Part Minimum Order Quantity
Part Location
Calculating Replenished Part Buffer Levels and Zones
The Green Zone
The Yellow Zone

The Red Zone
Continuing with Company ABC
Iterations
Calculating Replenished Override Buffers
Calculating Min-Max Buffers
Summary

CHAPTER 8

Buffer Adjustments

Recalculated Adjustments
Planned Adjustment Factors
Demand Adjustment Factor
Zone Adjustment Factor
Lead Time Adjustment Factor
Summary

CHAPTER 9

Demand Driven Planning

The Shift to Actual Demand
The Net Flow Equation
Qualifying Order Spikes
Supply Order Generation Based on Net Flow Position
Simulating DDMRP Supply Order Generation
Calculating Average On-Hand Inventory
Average On-Hand Range
Average On-Hand Target
Average Open Supply Orders
Decoupled Explosion
Hybrid Model Supply Order Generation
Prioritized Share Supply Order Consideration
Discount Optimization
Freight Optimization
Coverage Optimization
Min-Max Supply Order Generation
Completing the Company ABC Example
Summary

CHAPTER 10

Demand Driven Execution

- Buffer Status Alerts
  - Challenging Priority by Due Date
  - Planning Versus Execution Display
  - Current On-Hand Alert
  - Projected On-Hand Alert
- Synchronization Alerts
  - Material Synchronization Alerts
  - Lead Time Alert
- Summary

CHAPTER 11

DDMRP Impacts on the Operational Environment

- DDMRP Strategic Buffer Criteria
  - The Decoupling Test
  - The Bidirectional Benefit Test
  - The Order Independence Test
  - The Primary Planning Mechanism Test
  - The Relative Priority Test
  - The Dynamic Adjustment Test
- DDMRP Versus Safety Stock and Order Point
  - Safety Stock and the Buffer Criteria
  - Order Point and the Buffer Criteria
- DDMRP Impacts on Scheduling
  - DDMRP and Master Production Scheduling Assumptions
  - DDMRP Shop Floor Scheduling Implications
  - Finite Scheduling with DDMRP?
  - Additional Scheduling Sequence Impacts
- DDMRP and WIP Priority Management
- Summary

CHAPTER 12

DDMRP Metrics and Analytics

- Measuring Relevant Information (Signal Integrity)

Measuring Decoupling Point Integrity
Outlying Event Reports
Measuring Velocity
Driving Improvement in DDMRP
Summary

CHAPTER 13

The Demand Driven Organization

The Demand Driven Adaptive System
Demand Driven Sales and Operations Planning
Conventional Sales and Operations Planning
The Five Steps of Demand Driven Sales and Operations Planning
1. Strategic Business Management Direction and Review
2. Integrated Reconciliation
3. Managing the Portfolio and New Activities
4. Managing Demand
5. Managing Supply
Demand Driven Sales and Operations Planning Projections
Working Capital
Space
Capacity
Projected Order Frequency
Summary
Contribution of Dick Ling

CHAPTER 14

Implications for Technology

Operations and Information Technology—Two Ships Diverging in the Night?
DDMRP Software Compliance Criteria

PART 3 APPENDICES

APPENDIX A

An MRP Example

The Scenario
Simulating the Scenario

APPENDIX B

Simulating DDMRP Buffers

About the Simulation
Widget Simulation Results
Gazoonk Simulation Results
Simulating the Impact of Minimum Order Quantities
Summary
About the Author

APPENDIX C

Applying DDMRP to the Apparel Retail Environment

The Need for a Retail Application of DDMRP
A Retail Apparel DDMRP Example
Special Characteristics and Challenges of the Apparel Retail Environment
The Proposed Model
Retail DDMRP Buffer Zone Considerations
Realized Results
About the Author

APPENDIX D

Demand Driven MRP Dictionary

APPENDIX E

DDS&OP Checklist

Endnotes
References
Index

Foreword

When I was asked to help write the chapter on S&OP for this book, I was complimented and eager to make a contribution. With the development of DDMRP and now the writing of this book, Ptak and Smith have broken through common practice to bring common sense to supply chain management.

If your company is facing variability and uncertainty across your supply chain and the future looks little like the past, then this book holds the answer. DDMRP represents the future of planning in today's complex and volatile supply chains. Inherent flaws in the traditional planning approaches are exposed and resolved for today's complex adaptive supply chains. With the Demand Driven Adaptive Schema and the pivotal position of Demand Driven Sales and Operations Planning, this is not just a better way to plan; it is a better way to run an organization in today's hypercompetitive environment. Operations and strategy can now easily and realistically be connected bi-directionally, allowing both to adapt to critical changes for the best return on shareholder equity.

This book is the ultimate reference for this new way of life across a dynamic adaptive supply chain.

Dick Ling
S&OP Consultant
Author of *Orchestrating Success*

Definitions in This Book

This book will use two sources of definitions. All known and accepted terms that are not new with the advent of Demand Driven Material Requirements Planning (DDMRP) will be defined using definitions from the fourteenth edition of the *APICS Dictionary*. The authors thank APICS for its support of this project. Since 1957, APICS has been the premier professional association for supply chain and operations management and the leading provider of research, education, and certification programs that elevate supply chain excellence, innovation, and resilience.

For terms that are new with the advent of DDMRP, the authors have created a dictionary specific to DDMRP. This dictionary can be found in Appendix D of this book. Translated versions of this dictionary in multiple languages can be found in the download section at www.demanddriveninstitute.com.

Introduction

This is not a book about the intricacies of traditional Materials Requirements Planning (MRP). In 2011, at the request of McGraw-Hill, the authors wrote *Orlicky's Material Requirements Planning*, third edition. That book was 542 pages and provided an expansive view of conventional planning tactics that were born in the 1950s, codified in the 1960s, and commercialized in the 1970s. That book also devoted nearly 100 pages to an emerging alternative method of formal planning and execution—Demand Driven Material Requirements Planning (DDMRP). This book is entirely about that alternative method.

This book will provide an extensive blueprint for DDMRP. It is the authors' intention that this text will serve the same purpose as the first MRP book written by Joe Orlicky in 1975 by ushering in a new era in planning and execution methodology, rules, and tools. The body of knowledge of this alternative method is advancing rapidly and gaining acceptance worldwide. The authors hope that this book will open the door for many other books on specific aspects, applications, and extensions of this method.

About the Authors

Contact the authors at info@demanddriveninstitute.com.

Chad Smith

Chad Smith is the coauthor (with Carol Ptak) of the third edition of *Orlicky's Material Requirements Planning* and the coauthor (with Debra Smith) of *Demand Driven Performance: Using Smart Metrics*. He is a cofounder of and partner in the Demand Driven Institute, an organization dedicated to proliferating demand driven methods throughout the world. Mr. Smith serves as the Program Director of the International Supply Chain Education Alliance's Certified Demand Driven Planner Program.

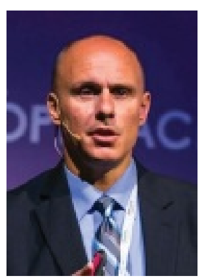

In 1997, Mr. Smith cofounded Constraints Management Group (CMG), a services and technology company specializing in demand driven manufacturing, materials, and project management systems for midrange and large manufacturers. He served as Managing Partner of CMG from 1998 to 2015. Clients, past and present, include Unilever, LeTourneau Technologies, Boeing, Intel, Erickson Air-Crane, Siemens, IBM, The Charles Machine Works (Ditch Witch), and Oregon Freeze Dry. Mr. Smith is also a certified expert in all disciplines of the Theory of Constraints, studying directly under the tutelage of the late Dr. Eli Goldratt.

Chad Smith makes his home in Wenatchee, Washington, with his wife, Sarah, and two daughters, Sophia and Lily.

Carol Ptak

Carol Ptak is currently a partner with the Demand Driven Institute and was previously a Visiting Professor and Distinguished Executive in Residence at Pacific Lutheran University. Before going to academia she was Vice President and Global Industry Executive for Manufacturing and Distribution Industries at PeopleSoft, where she developed the concept of demand driven manufacturing. Ms. Ptak spent four years at IBM Corporation, culminating in the position of Global SMB Segment Executive.

A leading authority in the use of ERP and supply chain tools to drive improved bottom-line performance, Ms. Ptak has an expertise that is well grounded in four decades of practical experience as a successful practitioner, consultant, and educator in manufacturing operations. Her pragmatic approach to complex issues and her dynamic presentation style make her a person in high demand worldwide on the subject of how to leverage these tools and achieve sustainable success.

Ms. Ptak holds an MBA from Rochester Institute of Technology and completed the EMPO program at Stanford University. She is a frequent educator at the university level and presents at many key technical conferences around the world, including conferences in South Africa, France, Israel, Australia, Ireland, and the Netherlands and 11 APICS International Conferences. She is the author of numerous articles and the books *Orlicky's Material Requirements Planning*, third edition, with Chad Smith; *MRP and Beyond; ERP: Tools, Techniques, and Applications for Integrating the Supply Chain; Theory H.O.W.: How Organizations Could Work* with Harold Cavallaro; and *Necessary but Not Sufficient* with Eli Goldratt and Eli Schragenheim. Together with Dean Gilliam, she updated *Quantum Leap*, originally written by John Constanza. Ms. Ptak has lent her name to the internationally coveted Ptak Prize for Supply Chain Excellence that is awarded annually by ISCEA, the International Supply Chain Education Alliance.

Ms. Ptak is certified through APICS at the fellow level (CFPIM) and was certified in Integrated Resource Management (CIRM) with the first group internationally. Ms. Ptak was the President and CEO of APICS for the year 2000. Prior to her election as APICS President, she served at the Society in a variety of positions.

Acknowledgments

If I have seen further, it is by standing on the shoulders of giants.
SIR ISAAC NEWTON TO ROBERT HOOKE, FEBRUARY 5, 1676

This book is truly built on the shoulders of many giants. From the original work of the practitioners who developed MRP, including Joe Orlicky, George Plossl, Richard (Dick) Ling, and Ollie Wight, to the great thinkers behind Lean, Six Sigma, and Theory of Constraints—Taiichi Ohno, W. Edwards Deming, and Eliyahu Goldratt. The authors have stood on the shoulders of these giants to unite these different theories and methodologies and take a leap forward—into a future of planning with relevant visibility that mitigates the volatile, uncertain, and variable world that seems impossible to plan. We have known many of these giants personally and wish to express our appreciation to them.

Collectively the authors would like to thank the International Supply Chain Education Alliance (ISCEA) and members of the Demand Driven Institute Global Affiliate Network for a great partnership in bringing demand driven concepts to the mainstream throughout the world. Additionally, the authors would like to thank various members of the APICS community for their amazing input and support in trying to restore the promise and effectiveness of formal planning. Those people include Keith Launchbury, Ken Titmuss, Bob Reary, and Abe Eshkenazi.

The authors would like to point out particular individuals who have made a lasting contribution to the demand driven body of knowledge. These people include Greg Cass, Debra Smith, Erik Bush, David Poveda, Dick Ling, Paddy Ramaiyengar, Kirk Black, Caroline Mondon, and Laurent Vigouroux.

The authors would like to highlight a few individuals and organizations that have been instrumental in spreading the DDMRP message. Caroline Mondon and the Fapics organization led an amazing charge in France. Philippe Bornert, Bernard Milian, and the whole Agilea team have been instrumental in supporting that charge in France. Ken Titmuss and the SAPICS organization brought these concepts to Africa. David Poveda was instrumental in starting a massive proliferation in South America.

Chad Smith would like thank his wife, Sarah, and two daughters, Sophia and Lily, for putting up with the prolonged absences and locked office door. The support and love of these three people has kept him going. Additionally, Chad would like to thank Carmine Mainiero and Nick Mantenuto for their dedication in bringing demand driven concepts to a global giant in the world of fast-moving consumer goods; what a learning experience! Chad would also like to acknowledge the team at Demand Driven Technologies and its CEO, Erik Bush, in believing in bringing real and sustainable results to customers. Chad would additionally like to thank the team at Constraints Management Group, LLC, for an amazing journey for nearly 20 years. Finally, Chad would like to thank his partner and coauthor Carol Ptak for a very rewarding and fulfilling partnership.

Carol Ptak would like to thank her husband, Jim, for the understanding and the support to keep going and to thank her parents, Dorothy and Bud, who taught her from the youngest age that she was limited only by her imagination. Carol would especially like to thank Chad Smith for an incredible experience and partnership—far beyond any imagination. Chad opened all our eyes to the deeper truth of a new world of planning. It has been an honor and a once-in-a-lifetime experience.

Demand Driven Material Requirements Planning (DDMRP)

PART 1

Perspective

CHAPTER 1

Planning in the New Normal

To truly understand where industry is today, it is necessary to discuss the history behind conventional planning. Where did it come from? What did it replace? What circumstances was it developed for? Is it still relevant and appropriate for the environment of today?

The Material Requirements Planning Revolution

Today most midrange and large manufacturing enterprises throughout the world use a planning method and tool called Material Requirements Planning (MRP). This method and tool was conceived in the 1950s with the increasing availability, promise, and power of computers. Computers allowed for rapid and complex calculations about what and how much was needed to be bought and made given a demand input.

The more complex the products, the more powerful the promise of MRP. The *APICS Dictionary*[1] defines MRP as:

> A set of techniques that uses bill of material data, inventory data, and the master production schedule to calculate requirements for materials. It makes recommendations to release replenishment orders for material. Further, because it is time-phased, it makes recommendations to reschedule open orders when due dates and need dates are not in phase. Time-phased MRP begins with the items listed on the MPS and determines (1) the quantity of all components and materials required to fabricate those items and (2) the date that the components and material are required. Time-phased MRP is accomplished by exploding the bill of material, adjusting for inventory quantities on hand or on order, and offsetting the net requirements by the appropriate lead times. (p. 103)

By 1965 the modern acronym “MRP” was in existence. Then in 1972 capacity reconciliation was incorporated into MRP. This was called closed-loop MRP. The year 1980 saw the significant incorporation of cost accounting into MRP, transforming it into a system known as Manufacturing Resources Planning (MRP II). Finally, by 1990, as client-server architecture became available, MRP II had evolved into Enterprise Resources Planning (ERP). Throughout this progression the definition of the MRP portion of the information system has remained unchanged.

While this is not a book about MRP, a basic level of understanding of MRP will be helpful to the reader. This basic explanation, and even a demonstration of MRP, is included in Chapter 3 and Appendix A, respectively.

Perhaps the most recognized leader of the MRP charge was Joe Orlicky. His 1975 seminal work *Material Requirements Planning: The New Way of Life in Production and Inventory Management* provided the blueprint and codification of MRP that is still the standard today. Consider that when this book was written, only 700 companies or plants in the world had implemented MRP, almost all located in the United States:

> As this book goes into print, there are some 700 manufacturing companies or plants that have implemented, or are committed to implementing, MRP systems. Material requirements planning has become a new way of life in production and inventory management, displacing older methods in general and statistical inventory control in particular. I, for one, have no doubt whatever that it will be the way of life in the future. (p. ix)

MRP did become the way of life in manufacturing. The codification and subsequent commercialization of MRP fundamentally changed the industrial world, and it did so relatively quickly. Orlicky, along with others at the time, recognized the opportunity presented by changing manufacturing circumstances and the invention of the computer that enabled a planning approach never before possible:

> Traditional inventory management approaches, in pre-computer days, could obviously not go beyond the limits imposed by the information processing tools available at the time. Because of this almost all of those approaches and techniques suffered from imperfection. They simply represented the best that could be done under the circumstances. *They acted as a crutch and incorporated summary, shortcut and approximation methods, often based on tenuous or quite unrealistic assumptions, sometimes force-fitting concepts to reality so as to permit the use of a technique.*
>
> The breakthrough, in this area, lies in the simple fact that once a computer

> becomes available, the use of such methods and systems is no longer obligatory. It becomes feasible to sort out, revise, or discard previously used techniques and to institute new ones that heretofore it would have been impractical or impossible to implement. It is now a matter of record that among manufacturing companies that pioneered inventory management computer applications in the 1960s, the most significant results were achieved not by those who chose to improve, refine, and speed up existing procedures, but by those who undertook a fundamental overhaul of their systems. (p. 4)

In his book, Orlicky made the case for a fundamental reexamination of how companies planned and managed inventory and resources. This case was so compelling that the concepts that he brought to the table proliferated throughout the industrial world within two decades. That proliferation remains largely unchanged in the present. Today we know that nearly 80 percent of manufacturing companies that buy an ERP system also buy and implement the MRP module associated with that system.

Perhaps the most interesting and compelling part of the passage from the original Orlicky book is the sentence that is italicized. This was simply common sense that was easily demonstrable with the results of precomputer inventory management systems. Yet could this same description be applied to the widespread use of MRP today? Could it be that conventional planning approaches and tools are:

- Acting as a crutch?
- Incorporating summary, shortcut, and approximation methods based on tenuous assumptions?
- Force-fitting concepts to reality so as to permit the use of a technique?

In the authors' 60+ years of combined manufacturing experience across a wide array of industries, the answer is a resounding *yes* to all these points. By the end of this book, the reader will also be able to understand *why* the answer is yes to all these points. Indeed if the answer is yes to these points, there should be evidence to support the assertion that MRP systems are not living up to their billing—that they are in fact guilty as charged in the previous three bullet points.

Before we review the evidence, let's start with two basic observations about rules:

- **Observation 1.** Most rules are life limited. Rules are instituted most often based on assumptions about the environment at the time they are made. Rules are often made to accommodate certain limitations. When those assumptions or limitations change, the rules must be reexamined to determine whether they are still appropriate. Souder's law states that "repetition does not establish validity." Simply continuing to do something that has always been done does not define whether it is or ever has been the appropriate thing to do. Worse yet, the longer the repetition, the more invalid or inappropriate the rule may be.

- **Observation 2.** "Optimizing" inappropriate rules is counterproductive. Attempts and investment meant to enable or accelerate compliance to rules that are inappropriate can be devastating to an organization. If the rule is not only inappropriate but also damaging, then the organization is at risk to do the wrong things faster.

Evidence of a Problem

There are three areas that point to major issues with the rules and tools of conventional planning featuring MRP.

Return on Asset Performance Degradation

As described above, the United States led the adoption of manufacturing information systems starting with MRP in the 1960s. These systems are expensive to purchase, to implement, and to maintain. The value of these formal planning systems has always been based on the ability to better leverage the assets of a business. Did the widespread adoption of MRP and subsequent information systems enable the U.S. economy to better manage assets?

In late 2013 Deloitte University Press released a report by John Hagel III, John Seely Brown, Tamara Samoylova, and Michael Lui that is quite eye-opening when considered against the progression and adoption rates of information systems.[2] Figure 1-1 is a chart from the report that depicts the return on asset performance of the United States economy since 1965.

There is a steady decrease in return on assets for the U.S. economy from 1965 to 2012. Furthermore, during this time period the same report shows that labor productivity (as measured by Tornqvist aggregation) more than doubled! What is most interesting about this graphic in relation to information systems is that by 1965 we had the modern acronym MRP, but massive proliferation of information systems did not occur until after 1975 and, in particular, after 1980 with MRPII.

Obviously there are many factors at play with this decrease in return on assets, but this report would certainly lead one to realize that the impact of the widespread adoption of MRP, MRP II, and ERP systems (at least in the United States) has not significantly helped companies manage themselves to better returns on asset performance. Indeed, when this decline is taken in combination with the increase in labor productivity, it actually suggests that companies are accelerating their mistakes.

But this is just one point of data, a high-level view with many unrelated factors contributing to these effects. What additional evidence indicts the efficacy of the conventional planning approach?

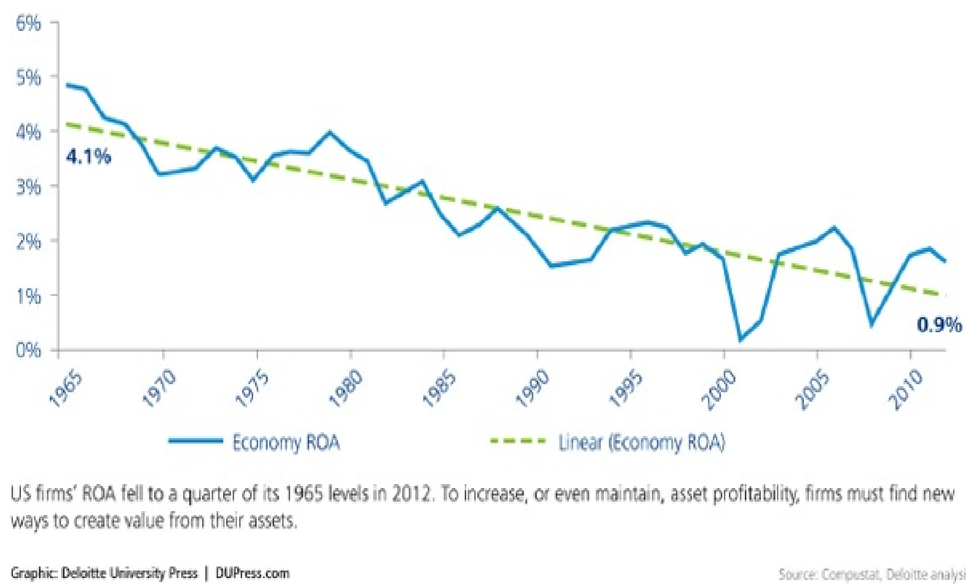

FIGURE 1-1 Return on asset peformance for the U.S. economy

Work-Around Proliferation

In addition to examining the performance of an entire economy over a period of time, next examine the day-to-day actions of the people charged with making decisions about how to utilize assets. One hallmark of supply chains is the presence of supply orders. Supply orders are the purchase orders, stock transfer orders, and manufacturing orders that dictate the flow and activities of any supply chain.

The very purpose of a planning system is to ultimately determine the timing, quantity, and collective synchronization of the supply orders up, down, and across the levels of the network. Inside most manufacturers there are tiers within the planning system where stock transfer orders could prompt manufacturing orders that in turn would prompt purchase orders. Additionally, within most supply chains there are tiers of different planning systems at each organization linked together by these orders and communicating through these supply order signals. For example, purchase orders from a customer can prompt stock transfers or manufacturing orders at suppliers.

Perhaps the biggest indictment of just how inappropriate modern planning rules and tools are can be observed in how frequently people choose to work around them. The typical workaround involves the use of spreadsheets. Data are extracted out of the planning system and put into a spreadsheet. The data are then organized and manipulated within the spreadsheet until a personal comfort level is established. Recommendations and orders are then put back into the planning system, essentially overriding many of the original recommendations.

Consider polling on this subject by the Demand Driven Institute from 2011 to 2014. With over 500 companies responding, 95 percent claim to be augmenting their planning systems with spreadsheets. Nearly 70 percent claim these spreadsheets are used to a heavy or moderate degree. The results of this polling are consistent with other surveys by analyst firms such as Aberdeen Group. This reliance on spreadsheets has often been referred to as “Excel hell.” Validation for this proliferation can be easily provided by simply asking the members of a planning and purchasing team what would happen to their ability to do their job if their access to spreadsheets were taken away.

But why have planners and buyers become so reliant on spreadsheets? Because they know that if they stayed completely within the rules of the formal planning system, approving all recommendations, it would be *very* career limiting. Tomorrow they would undo or reverse half the things they did today because MRP is constantly and dramatically changing the picture. This phenomenon, known as “nervousness,” is explained in Chapter 3.

So what do they do instead? They work around the system. They each have their own ways of working with tools that they have crafted and honed through their years of experience. These ways of working and tools are highly individualized with extremely limited ability to be utilized by anyone but the originator. This is a different, informal, highly variable, and highly customized set of rules.

Worse yet, there is no oversight or auditing of these side "systems." There is no "vice president of spreadsheets" in any company the authors have ever worked in or visited. Everyone simply assumes that the people who created these spreadsheets built and maintain them properly. Consider an article in the *Wall Street Journal's* "Market Watch" in 2013:

> Close to 90% of spreadsheet documents contain errors, a 2008 analysis of multiple studies suggests. "Spreadsheets, even after careful development, contain errors in 1% or more of all formula cells," writes Ray Panko, a professor of IT management at the University of Hawaii and an authority on bad spreadsheet practices. "In large spreadsheets with thousands of formulas, there will be dozens of undetected errors" (Jeremy Olshan, April 20, 2013)

As an example of how disastrous spreadsheet errors can be, consider the role a spreadsheet error played in a $6 billion disaster for JP Morgan in 2012. The following is an excerpt from the zerohedge.com article "How a Rookie Excel Error Led JPMorgan to Misreport Its VaR for Years"[3]:

> Just under a year ago, when JPMorgan's London Whale trading fiasco was exposed as much more than just the proverbial "tempest in a teapot," Morgan watchers were left scratching their heads over another very curious development: the dramatic surge in the company's reported VaR, which as we showed last June nearly doubled, rising by some 93% year over year, a glaring contrast to what the other banks were reporting to be doing.
>
> Specifically, we said that "in the 10-Q filing, the bank reported a VaR of $170 million for the three months ending March 31, 2012. This compared to a tiny $88 million for the previous year." JPM, which was desperate to cover up this modelling snafu, kept mum and shed as little light on the issue as possible. In its own words from the Q1 2012 10-Q filing: "the increase in average VaR was primarily driven by an increase in CIO VaR and a decrease in diversification benefit across the Firm." And furthermore: "CIO VaR averaged $129 million for the three months ended March 31, 2012, compared with $60 million for the comparable 2011 period. **The increase in CIO average VaR was due to changes in the synthetic credit portfolio held by CIO as part of its management of structural and other risks arising from the Firm's on-going business activities."** Keep the bolded sentence in mind, because as it turns out it is nothing but a euphemism for, drumroll, **epic, amateur Excel error!**
>
> How do we know this? We know it courtesy of JPMorgan itself, which in the very last page of its JPM task force report had this to say on the topic of JPM's VaR:
>
> ". . . a decision was made to stop using the Basel II.5 model and not to rely

on it for purposes of reporting CIO VaR in the Firm's first-quarter Form 10-Q. Following that decision, further errors were discovered in the Basel II.5 model, including, most significantly, an operational error in the calculation of the relative changes in hazard rates and correlation estimates. **Specifically, after subtracting the old rate from the new rate, the spreadsheet divided by their sum <u>instead of their average, as the modeler had intended</u>**. This error likely had the effect of **muting volatility by a factor of two and of lowering the VaR.... <u>it also remains unclear when this error was introduced in the calculation</u> ."**

In other words, the doubling in JPM's VaR was due to nothing but the discovery that for years, someone had been using a grossly incorrect formula in their Excel, and as a result misreporting the entire firm VaR by a factor of nearly 50%! So much for the official JPM explanation in its 10-Q filing that somewhat conveniently missed to mention that, oops, we made a rookie, first year analyst error. (Tyler Durden, February 2, 2013)

Perhaps a more interesting question is why are personnel allowed to use these ad-hoc approaches? From a data integrity and security perspective, this is a nightmare. It also means that the fate of the company's purchasing and planning effectiveness is in the hands of a few essentially irreplaceable personnel. These people can't be promoted or get sick or leave without dire consequences to the company. This also means that due to the error-prone nature of spreadsheets, globally on a daily basis there are a lot of wrong signals being generated across supply chains. Wouldn't it be so much easier to just work in the system? The answer seems so obvious. The fact that reality is just the opposite shows just how big the problem is with conventional systems.

To be fair, many executives are simply not aware of just how much work is occurring outside the system. Once they become aware, they are placed in an instant dilemma. Let it continue, thus endorsing it by default, or force compliance to a system that your subject-matter experts are saying is at best suspect? The choice is only easy the first time an executive encounters it. The authors of this book have seen countless examples of executives attempting to end the ad hoc systems only to quickly retreat when inventories balloon and service levels fall dramatically. They may not understand what's behind the need for the work-arounds, but they now know enough to simply look the other way. So they make the appropriate noises about how the entire company is on the new ERP system and downplay just how much ad hoc work is really occurring.

The Inventory Bimodal Distribution

Another piece of evidence to suggest the shortcomings of conventional MRP systems has to do with the inventory performance of the companies that use these systems. To understand this particular challenge, consider the simple graphical depiction in Figure 1-2. In this figure you see a solid horizontal line running in both directions. This line represents the quantity of inventory. As you move from left to right, the quantity of inventory increases; right to left the quantity decreases.

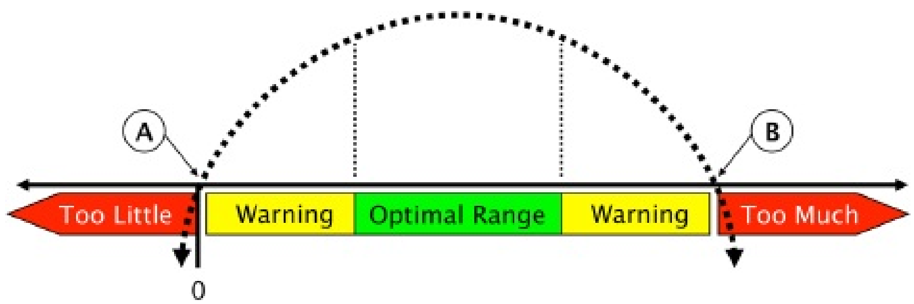

FIGURE 1-2 Taguchi inventory loss function

A curved dotted line bisects the inventory quantity line at two points:

- **Point A, the point where a company has too little inventory.** This point would be a quantity of zero, or "stocked out." Shortages, expedites, and missed sales are experienced at this point. Point A is the point at which the part position and supply chain have become too brittle and are unable to supply required inventory. Planners or buyers that have part numbers past this point to the left typically have sales and operations screaming at them for additional supply.
- **Point B, the point where a company has too much inventory.** There is excessive cash, capacity, and space tied up in working capital. Point B is the point at which inventory is deemed waste. Planners or buyers that have part numbers past this point to the right typically have finance screaming at them for misuse of financial resources.

If we know that these two points exist, then we can also conclude that for each part number, as well as the aggregate inventory level, there is an optimal range somewhere between those two points. This optimal zone is labeled in the middle and colored green. When inventory moves out of the optimal zone in either direction, it is deemed increasingly problematic.

This depiction is consistent with the graphical depiction of a loss function developed by the Japanese business statistician Genichi Taguchi to describe a phenomenon affecting the value of products produced by a company. This made clear the concept that quality does not suddenly plummet when, for instance, a machinist slightly exceeds a rigid blueprint tolerance. Instead “loss” in value progressively increases as variation increases from the intended nominal target.

The same is true for inventory. Chapter 2 will discuss how the value of inventory should be related to the ability of inventory to help promote or protect flow. As the inventory quantity expands out of the optimal zone and moves toward point B, the return on working capital captured in the inventory becomes less and less as the flow of working capital slows down. The converse is also true: as inventory shrinks out of the optimal zone and approaches zero or less, then flow is impeded due to shortages.

When the aggregate inventory position is considered in an environment using traditional MRP, there is frequently a bimodal distribution noted. With regard to inventory, a bimodal distribution can occur on two distinct levels:

1. A bimodal distribution can occur at the single-part level over a period of time, as a part will oscillate back and forth between excess and shortage positions. In each position, flow is threatened or directly inhibited. The bimodal position can be weighted toward one side or the other, but what makes it bimodal is a clear separation between the two groups—the lack of any significant number of occurrences in the “optimal range.”
2. The bimodal distribution also occurs across a group of parts at any point in time. At any one point, many parts will be in excess while other parts are in a shortage position. Shortages of any parts are particularly devastating in environments with assemblies and shared components because the lack of one part can block the delivery of many.

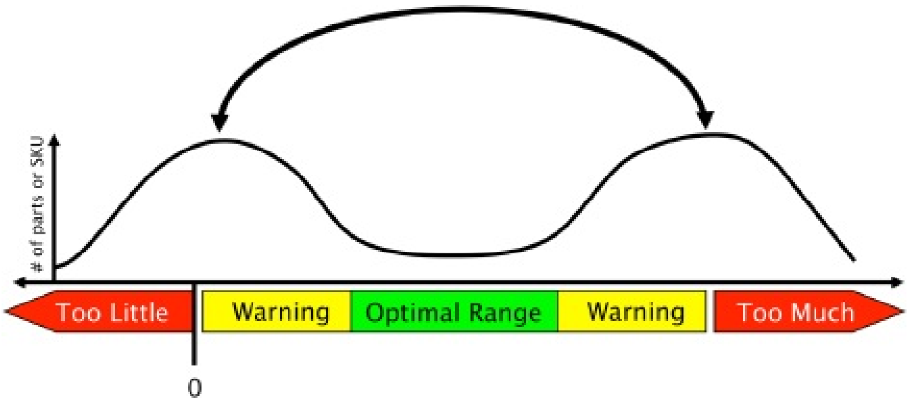

FIGURE 1-3 Bimodal inventory distribution

Figure 1-3 is a conceptual depiction of a bimodal distribution across a group of parts. The bimodal distribution depicts a large number of parts that are in the too-little range while still another large number of parts are in the too-much range. The Y axis represents the number of parts at any particular point on the loss function spectrum.

Not only is the smallest population in the optimal zone, but the time any individual part spends in the optimal zone tends to be short-lived. In fact, most parts tend to oscillate between the two extremes. The oscillation is depicted with the solid curved line connecting the two disparate distributions. That oscillation will occur every time MRP is run. At any one time, any planner or buyer can have many parts in both distributions simultaneously.

This bimodal distribution is rampant throughout industry. It can be very simply described as “too much of the wrong and too little of the right” at any point in time and “too much in total” over time. In the same survey noted earlier, taken between 2011 and 2014 by the Demand Driven Institute, 88 percent of companies reported that they experienced this bimodal inventory pattern. The sample set included over 500 organizations around the world.

Three primary effects of the bimodal distribution are evident in most companies:

1. **High inventories.** The distribution can be disproportionate, as many planners and buyers will tend to err on the side of too much. This results in slow-moving or obsolete inventory, additional space requirements, squandered capacity and materials, and even lower margin performance as discounts are frequently required to clear out

the obsolete and slow-moving items.

2. **Chronic and frequent shortages.** The lack of availability of just a few parts can be devastating to many manufacturing environments, especially those that have assembly operations and common material or components. The lack of any one part will block any assembly. The lack of common material or components will block the manufacture of all parent items calling for that common item. This means an accumulation of delays in manufacturing, late deliveries, and missed sales.
3. **High bimodal-related expenses.** This effect tends to be undermeasured and underappreciated. It is the additional amount of money that an organization must spend in order to compensate for the bimodal distribution. When inventory is too high, third-party storage space may be required. When inventory is too low, premium and fast freight are frequently used to expedite material. Overtime is then used to push late orders through the plant. Partial shipments are made to get the customers some of what they ordered but with significantly increasing freight expenses.

Why the bimodal distribution occurs is explained in Chapter 3. It is a combination of basic MRP traits, the type of demand signal that is typically used in conjunction with MRP, and the complex volatile supply chain environment within which companies now must operate.

The New Normal

Experienced planning and purchasing personnel know that if they simply follow what MRP recommends, they will be in big trouble. Shortages will increase. Excess inventory will increase. Expedites will increase. Intuitively, planners understand that materials and inventory management, under conventional practices, places them in a no-win situation. What happened to the promise of MRP as verbalized by Joe Orlicky in the beginning of this chapter? The answer is exceedingly simple: the world changed and MRP did not.

The circumstances under which Orlicky and his cadre developed the rules behind MRP have dramatically changed. Customer tolerance times have shrunk dramatically, driven by low information and transactional friction largely due to the Internet. Customers can now easily find what they want at a price they are willing to pay and get it in a short period of time.

Ironically, the planning complexity is largely self-induced in the face of these shorter customer tolerance times. Most companies have made strategic decisions that have directly made it much harder to do business. Product variety has risen dramatically. Supply chains have extended around the world driven by low-cost sourcing. Product complexity has risen. Outsourcing is more prevalent. Product life and development cycles have been reduced.

Add on top of this an increased amount of regulatory requirements for consumer safety and environmental protection, and there are simply more complex planning and supply scenarios than ever before. The complexity comes from multiple directions: ownership, the market, engineering and sales, and the supply base. While this complexity has risen, the potential of technology has progressed and accelerated. The lack of significant financial return on technology investments would strongly suggest that this potential, up to this point, has largely been squandered.

Figure 1-4 is taken from *Demand Driven Performance: Using Smart Metrics* by Debra Smith and Chad Smith. The figure shows the tremendous difference in supply chain circumstances between 1965 and 2015.

Circumstance	1965	2015
Supply chain complexity	Low. Supply chains looked like chains—they were more linear. Vertically integrated and domestic supply chains dominated the landscape.	High. Supply chains look more like "supply webs" and are fragmented and extended across the globe.
Product life cycles	Long. Often measured in years or decades (e.g., rotary phones).	Short. Often measured in months (particularly in technology).
Customer tolerance times	Long. Often measured in weeks and months.	Short. Often measured in days, with many situations dictating less than 24-hour turns.
Product complexity	Low.	High. Most products now have relatively complex mechanical and electrical systems and microsystems. Can you even work on a modern car anymore?
Product customization	Low. Few options or custom features available.	High. Lots of configuration and customization to a particular customer or customer type.
Product variety	Low. Example: toothpaste. In 1965 Colgate and Crest each made one type of toothpaste.	High. In 2012 Colgate made 17 types of toothpaste, and Crest made 42!
Long lead time parts	Few. Here the word "long" is in relation to the time the market is willing to wait. By default if customer tolerance times were longer, it stands to reason that there were less long lead time parts. More so, however, is the fact that supply chains looked different. Most parts were domestically sourced and thus often much "closer" in time.	Many. Today's extended and fragmented supply chains have resulted in not only more purchased items but more purchased items coming from more remote locations.
Forecast accuracy	High. With less variety, longer life cycles, and high customer tolerance times, forecast accuracy was almost a nonissue. "If you build it, they will buy it."	Low. The combined complexity of the above items is making the idea of improving forecast accuracy a losing battle.
Pressure for leaner inventories	Low. With less variety and longer cycles, the penalties of building inventory positions was minimized.	High. At the same time that operations is asked to support a much more complex demand and supply scenario (as defined above), it is required to do so with less working capital!
Transactional friction	High. Finding suppliers and customers took exhaustive and expensive efforts. Choices were limited. People's first experience with a manufacturer was often through a salesperson sitting in front of them.	Low. Information is readily available at the click of the mouse. Choices are almost overwhelming. People's first experience with a manufacturer is often through a screen sitting in front of them.

FIGURE 1-4 Changing supply chain circumstances

From Debra Smith and Chad Smith, *Demand Driven Performance: Using Smart Metrics,* McGraw-Hill, 2013, p. 9.

Summary

We appear to have come full circle as MRP, according to observable, prevailing, and widespread effects across the world, now appears to be guilty of the same deficiencies as the techniques that preceded it. Software is simply a tool that translates and reinforces rules into a routine. If the rules behind the software are inappropriate and outdated, then the rules must change before the tools can change. In recent years, however, industry and software providers have attempted to combat increasing complexity with more sophisticated software applications, applications with the old rules still embedded at their core. The net effect is that we have improved the efficiency of doing the wrong or inappropriate things. Money and energy spent to optimize antiquated rules with increasingly sophisticated tools are wasteful, distractive, and counterproductive. Given the current world of increased variability and volatility, conventional planning logic now requires a fundamental overhaul. The authors think Joe Orlicky would agree.

The authors' self-imposed mission was to stand on the shoulders of Joe Orlicky's incredible vision in order to see further. This book proposes elegant and intuitive alternative planning rule sets to address the volatile twenty-first-century landscape. Complexity cannot be combated with more complexity. Effective and simplified rules and subsequent tools are necessary for a company's resources to work more closely in alignment with the market, enabling a demand driven world. There can be no more lip service to small incremental changes that may or may not improve a company's performance; concrete and proven tactics are required that drive sustainable bottom-line results. Where to start?

CHAPTER 2

The Importance of Flow

To understand why precise high-powered tools like MRP are not living up to their potential as well as the direction for a potential solution, let's start at a fundamental level. All for-profit entities have the same objective: to drive shareholder equity. Thus the rules and tools within a for-profit entity must be aligned to that objective. There is in fact a fundamental principle that aligns business rules and tools to that objective.

Plossl's First Law

Manufacturing comprises a bewildering and distracting variety of products, materials, technology, machines, and people skills that obscure the underlying elegance and simplicity of it as a process. The essence of manufacturing (and supply chain in general) is the flow of materials from suppliers, through plants, through distribution channels, to customers; the flow of information to all parties about what is planned and required, what is happening, what has happened, and what should happen next; and the flow of cash.

An appreciation of this elegance and simplicity brings us to what George Plossl (a founding father of MRP and author of the second edition of Orlicky's *Material Requirements Planning)* articulated as the first law of manufacturing:

> All benefits will be directly related to the speed of flow of information and materials.

"All benefits" is quite an encompassing phrase. It can be broken down into components that most companies measure and emphasize. All benefits encompass:

- **Service.** A system that has good informational and material flow produces consistent and reliable results. This has implications for meeting customer expectations, not only for delivery performance but also for quality. This is especially true for industries that have shelf-life issues.

- **Revenue.** When service is consistently high, market share tends to grow—or, at a minimum, doesn’t erode.
- **Quality.** When things are flowing well, fewer mistakes are made that are caused by confusion and expediting.
- **Inventories.** Purchased, work-in-process (WIP), and finished goods inventories will be minimized and directly proportional to the amount of time it takes to flow between stages and through the total system. The less time it takes products to move through the system, the less the total inventory investment. The simple equation is Throughput * lead time = WIP. Throughput is the rate at which material is exiting the system. Lead time is the time it takes to move through the system, and WIP is the amount of inventory contained between entry and exit. A key assumption is that the material entering the system is proportionate to the amount exiting the system. The basis for this equation is the queuing theory known as Little’s law.
- **Expenses.** When flow is poor, additional activities and expenses are incurred to close the gaps in flow. Examples would be expedited freight, overtime, rework, cross-shipping, and unplanned partial ships. Most of these activities are indicative of an inefficient overall system and directly cause cash to leave the organization. These types of expenses were described in Chapter 1 in relation to the bimodal distribution.
- **Cash.** When flow is maximized, the material that a company paid for is converted to cash at a relatively quick and consistent rate. This makes cash flow much easier to manage and predict. Additionally, the expedite-related expenses previously mentioned are minimized, limiting cash leaving the organization.

What happens when revenue is growing, inventory is minimized, and additional and unnecessary ancillary expenses are eliminated? Return on investment (ROI) moves in a favorable direction. Thus the fundamental principle is established that the rules and tools of a business should be built around the *protection and promotion of flow*.

Establishing Flow as the Foundation

It is difficult to foster the flow of relevant information and materials through a system when the components of the system cannot relate their actions to that flow. It's become cliché to say that our organizations have "silos." Those silos typically result in friction, conflict, and communication difficulties between functions. This is because we tend to control segments of our organizations through different metrics. Figure 2-1 lists an organization's primary functions and the respective primary objectives and example metrics to accomplish those objectives.

The actions that each of these functions might take to meet their primary objectives and metrics often come into conflict. As an example, sales typically has a different primary metric than operations. It can frequently be the case that when operations looks to maximize its primary metric, it may compromise or jeopardize the primary metric of sales and vice versa. When quality maximizes its departmental metric, then operations might be adversely affected.

Yet we have already established that when a system flows well, service, revenue, quality, inventories, expenses, and cash are all better controlled. All these elements directly protect the primary objectives of the functions in Figure 2-1. But if flow is not made visible and incorporated into the routine and metrics, then how can it possibly be protected? Flow, if encouraged, measured, and made properly visible, can align all these objectives with the system goal of maximizing return on shareholder equity.

Function	Primary Objective	Example Metric
Planning	Synchronize supply and demand	Shortages and excess inventory
Finance	Drive shareholder equity	Return on average capital
Sales	Capture demand	Order book
Marketing	Create brand awareness and demand	Market share
Operations	Utilize Assets	Overall equipment effectiveness
Quality	Meet specifications	Warranty claims

FIGURE 2-1 Organizational functions, objectives, and metrics

Thus, aligning the functions to the promotion and protection of the flow can be the bridge between local actions and the global benefits. Furthermore, this alignment should significantly raise the quality and timing of relevant information and corresponding relevant materials in a system.

Additionally, the protection and promotion of flow is a unifying concept within major process improvement disciplines and their respective primary objectives. Dr. Eliyahu Goldratt, the inventor of the Theory of Constraints had a primary objective of driving system throughput. This was accomplished by a focus on total system flow. Late in Dr. Goldratt's life, his writings became very specific about the interdependence between the Theory of Constraints and Taiichi Ohno's work with the Toyota Production System (TPS) and flow. Most in the West might say that the goal of TPS and any Lean system is to eliminate waste. When things flow well, there is indeed less waste. But TPS is not just about waste elimination. When Ohno's writings are examined closely, it becomes evident that the primary goal was in fact flow as described in his River Production System for Flow. Additionally, the quality movement driven by Dr. W. Edwards Deming and his 14 points for quality heavily relied on flow. The need for flow is obvious in this framework since improved flow results from less variability.

Any discussion or time spent on ideological battles between these disciplines is a complete waste of time and quite frankly, boring. Focusing on flow is about achieving a common objective through a common strategy based on common sense (also leveraging physics, biology, economics, and management accounting).

Goldratt, Ohno, and Deming did not invent the concept of flow. Their disciplines simply built off the works and concepts of industrial giants that changed manufacturing forever and gave birth to the corporate management structure in use today. To these industrial pioneers the concept of flow was simply common sense. These industrial pioneers include:

- **Frederick Taylor, a founding father of operations management.** Taylor developed the processes for time standards, product routings, tools, methods, and instructions as well as variable costing system and standard variance analysis. He developed the concept of planning as an actual business function.
- **Henry Ford, a founding father of mass production.** The processes used in early Ford production were based on the fact that the slowest task governs flow and that when there is synchronization of activity

to, through, and from those tasks, total system speed and velocity are protected. Ford was well known for his focus on the value of “no wait time.”

- **F.Donaldson Brown, a founding father of management accounting.** During his time at DuPont, Brown developed the DuPont ROI in addition to cost, volume, profit analysis, and flex budgeting. Brown defined relevant information for decision making and pioneered market segmentation at scale—all of which was based on a foundation of the promotion of flow.

Relevant Information and Materials

Yet there is an important caveat to Plossl’s first law that becomes crucial and central to determining whether flow translates to better ROI performance. It has already been hinted at several times in the preceding text. The great basketball coach John Wooden said, “Never mistake activity for achievement.” A company cannot just indiscriminately move data and materials quickly through a system and expect to be successful. Today organizations are frequently drowning in oceans of data with little relevant information and large stocks of irrelevant materials (too much of the wrong stuff) and not enough relevant materials (too little of the right stuff). When this occurs, there is a direct and adverse effect on return on investment.

Thus the flow of information and materials must be *relevant* to the *required* output or market expectation of the system. To be relevant, both the information and materials must synchronize the assets of a business to what the market really wants; no more, no less. Having the right information is a prerequisite to having the right materials. With this is mind, Plossl’s first law can be amended to:

> All benefits will be directly related to the speed of flow of *relevant* information and materials.

But in the highly complex and volatile New Normal, is it even possible to promote and protect the flow of relevant information? What stands in the way?

The Bullwhip Effect

A massive amount of research and literature has been devoted to the phenomenon known as the bullwhip effect. However, very little, if any, of that body of knowledge has been related specifically to the objective of protecting and promoting the flow of relevant information and material. *APICS Dictionary* defines the bullwhip effect as:

> An extreme change in the supply position upstream in a supply chain generated by a small change in demand downstream in the supply chain. Inventory can quickly move from being backordered to being excess. This is caused by the serial nature of communicating orders up the chain with the inherent transportation delays of moving product down the chain. The bullwhip can be eliminated by synchronizing the supply chain. (p. 19)

This definition clearly deals with important points discussed earlier in this book. “Inventory can quickly move from being backordered to being excess” is descriptive of the oscillation effect that occurs with the bimodal distribution. Additionally, this definition deals with both information and materials. “Communicating orders up the chain” is the information component, while “moving product down the chain” is the materials component.

In this respect, the bullwhip is really the systematic breakdown of relevant information and materials in a supply chain. Figure 2-2 is a graphical depiction of the bullwhip effect. The wavy arrow moving from right to left is the distortion to relevant information in the supply chain. The arrow wave grows in amplitude in order to depict that the farther up the chain you go, the more disconnected the information becomes from the origin of the signal, as signal distortion is transferred and amplified at each connection point. An MRP characteristic known as nervousness combined with batching practices creates this transfer and amplification, respectively. Both are explained in Chapter 3.

In the figure, the wavy arrow moving from left to right is the distortion in relevant materials in the supply chain. The wave grows in amplitude from low-level suppliers to the end item producer (OEM) to show the accumulation of delays that occur due to chronic shortages and late shipments. This transference and amplification occurs due to batching practices and the inherent synchronization problems associated with the probability of simultaneous availability; both are explained in Chapter 3.

It could and should be argued that the prevalence of the bullwhip effect is a fourth indicator of conventional planning logic deficiency. Chapter 3 describes how this logic, driven by one key attribute, directly leads to the bullwhip effect.

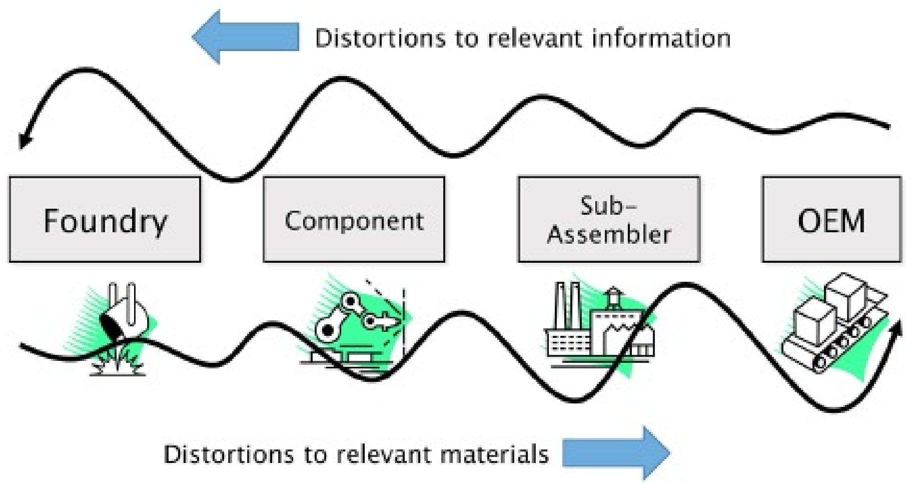

FIGURE 2-2 The bullwhip effect

Summary

The flow of relevant information and materials is the fundamental principle to achieve sustainable success across the supply chain. Is the concept of promoting flow difficult for people to grasp? Titans of early industry like Henry Ford, F. Donaldson Brown, and Frederick Taylor all understood this importance and built their models around it, models that provided the backbone of modern corporate structure. Later thought leaders such as Plossl, Ohno, Deming, and Goldratt built entire methodologies around the concept. The concept is a basic tenet of management accounting.

The concept is also intuitive. In general, most people within an organization seem to intuitively grasp why flow is so important. Yet there is a significant amount of evidence to suggest that most companies are incapable of really managing their assets with this fundamental principle. Next, the planning systems in use throughout the world to plan and manage the use of these assets are examined to discover why this is the case.

CHAPTER 3

Material Requirements Planning in the New Normal

As discussed in Chapter 2, the primary enemy of the protection and promotion of the flow of relevant information and materials is the bullwhip effect. The bullwhip effect exists largely due to the characteristics and configurations of conventional planning systems utilizing MRP. This chapter will describe these characteristics and configurations and highlight one key attribute as a core problem.

What Is MRP?

The *APICS Dictionary* defines Material Requirements Planning (MRP) as a:

> A set of techniques that uses bill of material data, inventory data, and the master production schedule to calculate requirements for materials. (p. 103)

MRP is essentially a calculation hub. The master production schedule feeds demand signals to MRP, which in turn creates a synchronized list of supply orders based on current inventory records (on hand and on order) and product structure (bill of material). The supply orders have date and quantity requirements that define the elements of that synchronization plan. These date and quantity requirements are then fed to a manufacturing execution system. They are turned into transfer orders to distribution sites, manufacturing orders to be scheduled on the shop floor, and purchase orders to be relayed to suppliers. Figure 3-1 shows this conventional planning approach.

The requirements to run MRP are simple and straightforward:

- The master schedule must be stated in terms of the bill of material.
- Unique item numbers exist for every item.

- The bill of material exists at the time of planning (product structure file).
- Inventory records are available for all items (inventory record file).

When these requirements are implemented in the computer system, then the MRP batch program can be run. However, to be considered a Class A user or to expect some kind of reasonable result from the computer system, the following assumptions are made:

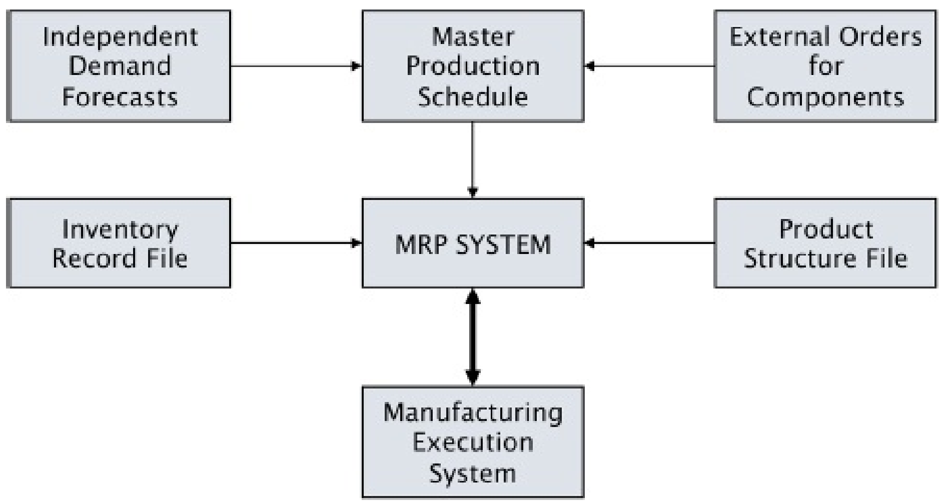

FIGURE 3-1 The conventional planning schema

- File data are 100 percent accurate and complete.
- Lead times are fixed and known.
- Every inventory item goes into and out of stock.
- There is full allocation; no order is started unless all the components are available.
- Components are discrete—things can be counted and measured (no "use as required").
- There is order independence, which means that every order can be started and completed on its own.

MRP was a huge leap forward because for the first time what was required could be calculated based on what was already there compared with what was needed, with the net result time phased. The objective of MRP was to precisely time-phase the requirements and replenishments to dramatically reduce inventory from the previous order point approach where some of everything was kept around all the time. This ability to calculate dependent demand through a bill of material was a significant development. It was no longer necessary to forecast dependent demand—it could be calculated based on the expected demand for the parent part. APICS defines dependent demand as:

> Demand that is directly related to or derived from the bill of material structure for other items or end products. Such demands are therefore calculated and need not and should not be forecast. (p. 46)

MRP evolved because of the advent of the computer, and the age of marketing in the 1950s introduced more product variety and complexity than was managed previously. Order point (the previous method of materials management) clearly could not affordably handle these new requirements. To understand how planners deal with MRP on a daily basis, refer to Appendix A, where a simulated environment demonstrates the day-to-day difficulties associated with MRP.

Yet even if a company has 100 percent of the requirements and 100 percent of the assumptions validated, the conventional planning approach will still be ineffective. The remainder of this chapter will explain why.

Distortions to Relevant Information

The conventional planning approach actually creates the bullwhip effect and its inherent distortions to the flow of relevant information and materials. Some of the ways in which conventional planning creates the bullwhip is related to the manner in which convention chooses to use MRP. Other contributions to the bullwhip are related to hard-coded traits in MRP systems. All of these issues, however, are related to one key and fundamental attribute of MRP.

Demand Signal Input

MRP is essentially a calculator. It needs three basic inputs to perform its calculation. One of those inputs is "demand." Different demand inputs will produce different outputs. The *APICS Dictionary* defines demand as:

> A need for a particular product or component. The demand could come from any number of sources (e.g., a customer order or forecast, an interplant requirement, a branch warehouse request for a service part or the manufacturing of another product. (p. 44)

By this definition, demand can be broken down into two different types: forecasted and actual. Both of the following definitions are from the *APICS Dictionary:*

- **Forecast.** An estimate of future demand. A forecast can be constructed using quantitative methods, qualitative methods, or a combination of methods, and it can be based on extrinsic (external) or intrinsic (internal) factors. Various forecasting techniques attempt to predict one or more of the four components of demand: cyclical, random, seasonal, and trend. (p. 68)
- **Actual demand.** Actual demand is composed of customer orders (and often allocations of items, ingredients, or raw materials to production or distribution). Actual demand nets against or "consumes" the forecast, depending upon the rules chosen over a time horizon. For example, actual demand will totally replace forecast inside the sold-out customer order backlog horizon (often called the demand time fence) but will net against the forecast outside this horizon based on the chosen forecast consumption rule. (p. 4)

The type of demand that is chosen to drive the MRP calculation is a primary determinant of how much relevant information can be produced from MRP. Remember, the flow of information and materials must be relevant to the required output or market expectation of the system. To be relevant, both the information and materials must synchronize the assets of a business with what the market really wants; no more, no less.

A hard-coded trait of MRP is that with a given demand signal, MRP is designed to net perfectly to zero. You make exactly what you need without any excess. In this regard it could be argued that MRP is the perfect JIT system. If the demand signal is perfectly accurate, then the MRP calculation will be perfectly accurate. Given that the math allows no tolerance for error, it seems obvious that MRP should only be given as accurate a signal as possible.

With that in mind, should the demand input to MRP be what a company thinks the market wants to buy or what the customers actually want to buy? Which will produce a more relevant result? As described in the definition of actual demand as well as Figure 3-1, the conventional approach combines both types of demand. Forecast is used to create planned orders, and then demand is adjusted as the picture becomes clearer with actual orders. Why is this problematic?

There are three truths about forecasts:

1. All forecasts start out with some inherent level of inaccuracy. Any prediction about the future carries with it some margin of error. This is especially true in the more complex and volatile New Normal.
2. The more detailed or discrete the forecast is, the less accurate it is. There is definitely a disparity in the accuracy between an aggregate-level forecast (all products or parts), a category-level forecast (a subgroup of products or parts), and a SKU-level forecast (single product or part).
3. The more remote in time or farther out forecasts go, the less accurate they get. Predicting the weather tomorrow is much more accurate than predicting the weather 52 days from today. Yes, history can be used as a basis for a prediction, but the margin of potential error is much higher. It is not uncommon that in many industries the accuracy of a forecast can drop below 10 percent beyond 90 days at the SKU level.

Today many forecasting experts admit that 70 to 75 percent accuracy is the benchmark for the SKU level. Figure 3-2 is the results of a 2012 survey conducted by forecastingblog.com showing the reported forecast error rates across various industries at the SKU level.

Unfortunately, when you start a serial, complex, and interdependent process with an error-prone input, the resulting output integrity must be suspect. Planned orders are derived from these forecasts, and very real commitments of cash, capacity, and materials are directly derived from a prediction that is subject to varying degrees of inaccuracy, sometimes with extremely significant degrees of inaccuracy.

As time progresses, the demand picture changes with the incorporation of actual demand, MRP is rerun, and subsequent changes occur. The result is that we end up with things that we do not need and desperately expedite things we have just discovered that we do need. These are the three effects of the bimodal distribution. Thus the bimodal distribution starts with the use of planned orders based on that forward-looking forecast.

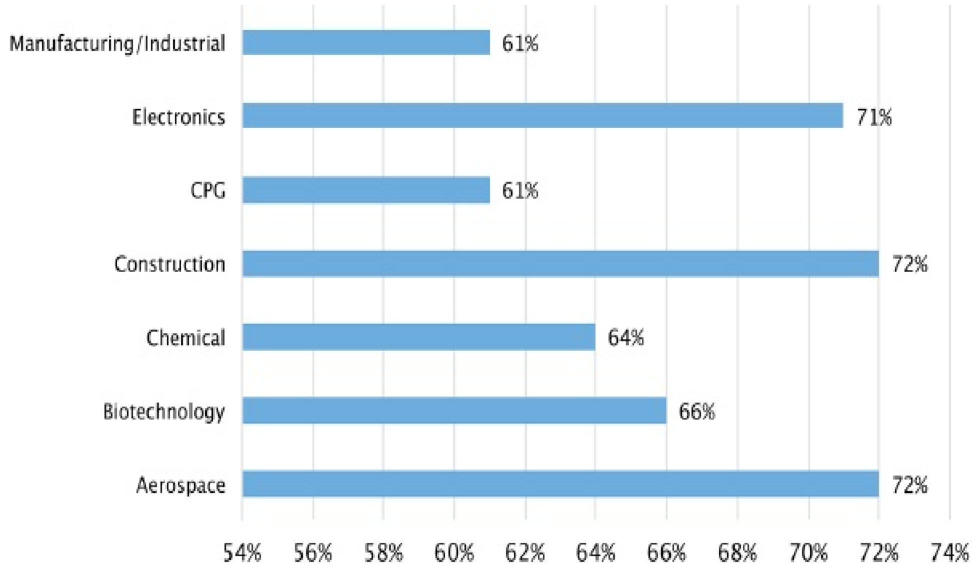

FIGURE 3-2 Average forecast accuracy across industries

Rohan Asardohkar, August 22, 2012, http://www.forecastingblog.com/?p=423

This is a known and accepted routine in most industries despite the waste and performance erosion associated with it. Why would industry intentionally sabotage performance by using an input with known inaccuracy to drive activity and commitments when there is an obvious alternative? Why not just use only sales orders?

The most accurate form of demand input is a sales order. A sales order is a stated intention and commitment to buy from an actual customer in terms of both quantity and time. It is essentially an uncashed check. In this way it is a highly accurate and relevant piece of information. There should be no debate that sales orders are an order of magnitude more accurate than planned orders. So why don't companies simply load only sales orders into MRP?

Using MRP with only sales orders, however, assumes something that does not exist in today's New Normal—enough time. A basic attribute of MRP is to net to zero across the entire network of dependencies. This means that MRP by definition makes all activities dependent on each other. Thus, in order for MRP to be that perfect JIT system, there must be sufficient time to procure and make everything to the stated demand—called "cumulative lead time" (the longest stated chain of time in the bill of material including purchasing lead time).

This means that customer tolerance time would have to be equal to or greater than the cumulative lead time. Today's supply chains, however, are characterized by shorter and shorter customer tolerance times and extended, elongated, and increasingly complex supply chains. There simply is not sufficient visibility to sales orders soon enough to properly plan for them using conventional MRP. Figure 3-3 conceptually shows the disparity between when companies gain visibility to sales orders (actual demand) versus the time that it takes to procure and produce the product (the time frame in which MRP makes it calculations).

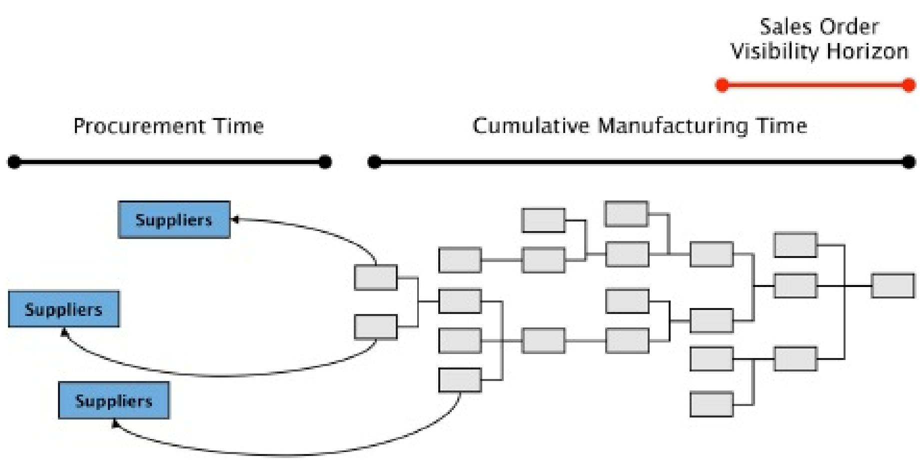

FIGURE 3-3 Manufacturing and procurement times versus sales order visibility

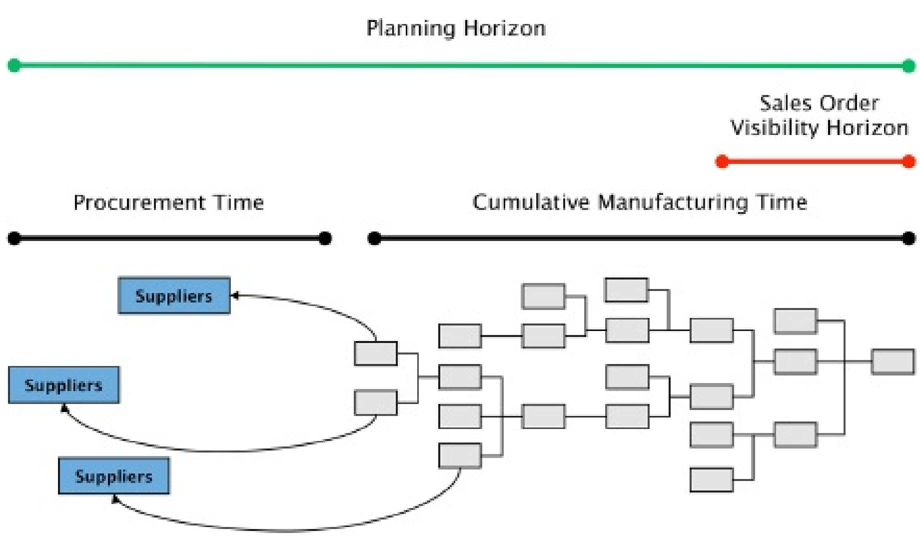

FIGURE 3-4 Planning horizon depiction

With MRP's characteristic of making everything dependent, the only way to find enough time is to attempt to predict what actual demand will look like so that an organization can attempt to ensure that the necessary materials are available in quantity and time as the market places its sales orders. A "planning horizon" extends into the future far enough to cover the cumulative procurement and manufacturing cycle. Figure 3-4 shows the planning horizon covering the cumulative procurement and manufacturing cycle in the example.

This explains the need to load MRP with demand that is largely derived from a forecast and then to make adjustments close in as sales orders become visible. Planned orders for end items are launched at the beginning of the planning horizon. The longer the procurement and manufacturing cycle, the longer the planning horizon must be. The longer the planning horizon, the less accurate the planned orders will be. The less accurate the planned orders, the more course corrections are required. This constant set of corrections brings us to another inherent trait of MRP called "nervousness."

Nervousness

MRP's nature of making everything dependent creates nervousness. Nervousness is the characteristic in an MRP system related to changes in parent demand transferring down and across bills of material. The *APICS Dictionary* defines nervousness as:

> The characteristic in an MRP system when minor changes in higher level (e.g. level 0 or 1) records or the master production schedule cause significant timing or quantity changes in lower level (e.g. 5 or 6) schedules or orders. (p. 86)

Figure 3-5 illustrates the concept of nervousness. The figure illustrates the product structure for an end item called FPA. A timing or quantity change in FPA ripples down through the entire product structure, causing timing and quantity changes at every component position as the system constantly strives to net to zero. The dotted curved arrows depict that change. This creates a constant series of action messages for planners and buyers to review and interpret.

The challenge of system nervousness has been known since the earliest days of MRP. However, the system nervousness was manageable since plans were done once per month. Concepts like firm planned orders, the demand time fence, and the master production schedule were developed to manage the nervousness. But the complex and volatile environment characterized by the New Normal makes the issue a bigger challenge. Given the nature of MRP to make everything dependent, the only way to stop nervousness is to make no changes. Yet that would mean significant service challenges, as the forecasted orders will vary (many times dramatically) from what the market will really desire. What can be done to limit the impact of nervousness? MRP users are forced into compromises in order to slow down the rate of changes.

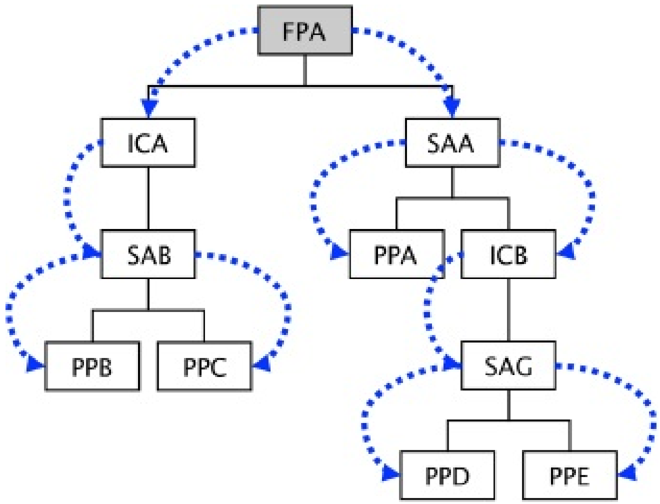

FIGURE 3-5 Nervousness illustrated

The Weekly Bucket

In most conventional environments, planning occurs in weekly buckets. This is a direct effect of the nervousness discussed above—nervousness that is directly related to the use of planned orders with MRP. Planning organizations know that if they ran MRP daily, or worse yet in real time, the resulting nervousness would create chaos. The amount of action flags and messages on the planning screens would be overwhelming.

Instead, a weekly interval is typically used to calm the waters on a daily level. This, however, comes at a price. First, it forces the planning horizon to extend even further (one week). This has a direct correlation to the level of signal inaccuracy at the end of the horizon. Second, it creates a latency that almost guarantees that the level of change between MRP runs will be dramatically larger. Instead of lots of little changes on a daily basis, there are massive changes (and signal distortions) on a weekly basis.

Figure 3-6 depicts the differences in net change impact between daily and weekly MRP runs. The upper left hand bar chart depicts MRP run each day. The level of each change is relatively small but each change ripples through all lower dependencies. The bar chart in the upper right portion of the graphic depicts a weekly MRP run. Days 1–7 are stable (no change) yet Day 8 introduces a significant change (40) that will ripple through the environment. The relative difference in changes is depicted in the chart in the lower left corner of the graphic.

Planning organizations are stuck between these two hard places because of MRP's hard-coded trait of making everything dependent.

Flattening the Bill of Material

Another way to combat nervousness is to reduce the number of connections that MRP sees and calculates against. One way to accomplish this is to "flatten" the bill of material of a product. MRP from a planning and synchronization perspective then becomes blind to intermediate components. Figure 3-7 illustrates the difference between a full product structure (on the left) and a flattened one (on the right). The flattened structure has eliminated all intermediate positions.

While this reduces the amount of changes to intermediates (since there are none) and this reduces the total number of action flags, does it produce more relevant information or actually distort the picture further? The key to more relevant information is not to simply ignore dependencies. When we ignore critical dependencies, we risk oversimplification. Oversimplification means to simplify to the point of error, distortion, or misrepresentation.

The bill of material files used in a planning system should reflect how the product is actually made. Dramatically flattening bills like the example in Figure 3-7 effectively ends any capability to provide visibility to and plan for synchronization between the finished and purchased part levels. The price for this is paid by the manufacturing floor as scheduling and schedule execution become an order of magnitude more difficult.

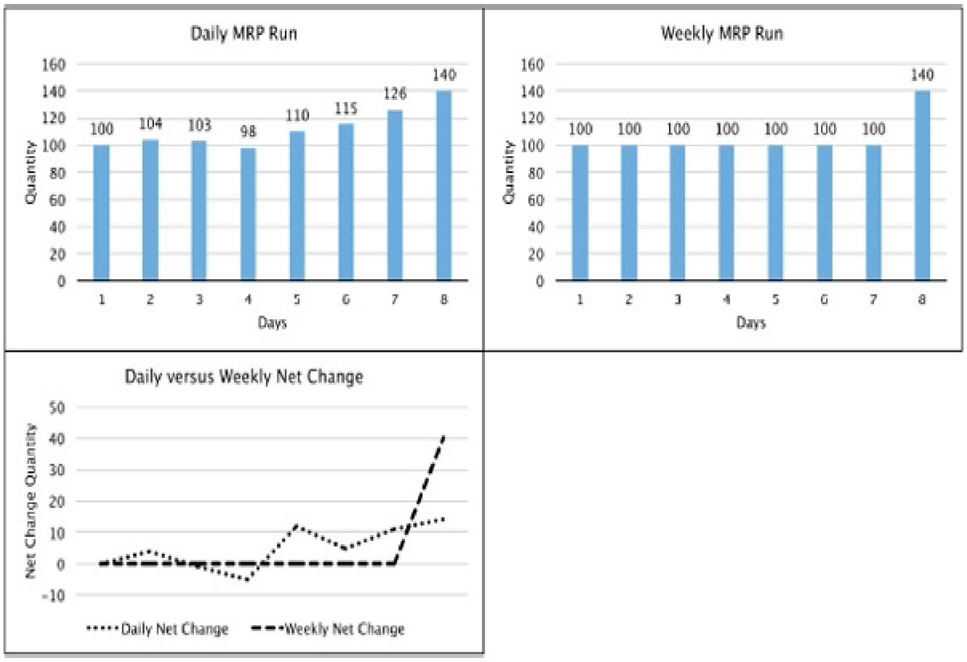

FIGURE 3-6 Daily versus weekly MRP runs

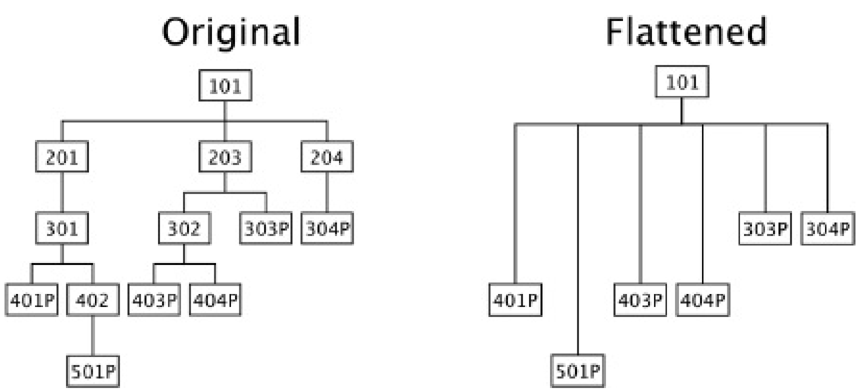

FIGURE 3-7 The flattening of a bill of material

All of these factors combine to mean that MRP is producing plans:

- With high degrees of known error (forecast input)
- In a constant state of change (nervousness)
- With a degree of latency (weekly bucket)
- That may misrepresent the environment (flattened bills of material)

This means that the very nature of MRP combined with the way that it is typically used inevitably leads to distortions to relevant information. Furthermore, all of these distortions to relevant information have been related to one single attribute of MRP. Have you spotted it yet?

Distortions to Relevant Materials

The next consideration is the supply portion of the bullwhip—the distortion of relevant materials. As mentioned previously, MRP creates a synchronized and precise plan at all levels of the bill of material based on its required inputs and assumptions. This plan will happen only if *everything* in the entire dependent network goes precisely according to plan. In almost every modern environment, this is an impossibility for two reasons.

Common Cause Variation

First, there is a basic and inherent level of variability in any environment, even one deemed to be in control. Deming called the normal or random variation that occurs in processes "common cause variation." Normal or random operational variability results in a process that may be statistically within calculated control limits but still varying between those limits. Reducing the gap between the limits is a worthy goal. The elimination of the gap is an impossibility—it would require every process to be perfect.

Delay Accumulation

We know that *any* process cannot be perfect. The collective effect of this imperfection must be examined. Figure 3-8 appeared in the first and third editions of Orlicky's *Material Requirements Planning.* The figure has three columns. The first column is the number of components required to make a parent item. The second two columns are different levels of average component availability. The left column assumes all components have 90 percent availability, whereas the right column assumes 95 percent availability. For example, a parent item with 4 components that average 90 percent availability has a 65.6 percent (.9 × .9 × .9 × .9) chance that all components will be available simultaneously when required. A parent item that has 10 components that have an average of 95 percent availability will have a 59.9 percent chance that all components will be simultaneously available when needed.

Figure 3-9 shows how less than perfect material availability results in an erosion of the probability that all materials will be present when needed.

Remember, MRP assumes full allocation—no order should be started unless all the components are available. In fact, even if many components have extremely high variability or arrive early, the parent order release is still at the mercy of any one missing component.

Figure 3-9 illustrates an environment in which four of the materials have high availability while one component has low availability on average. Components 1, 3, 4, and 5 have extremely high average availability (95 percent, 98 percent, 97 percent, and 99 percent, respectively. Component 2, however, has a relatively low average availability level (72 percent). The impact that component 2 has on the overall probability that all components will be available when required is significant; that probability drops to 64.4 percent. This translates to delays in the planned release.

Number of Components	Availability Level	
	90%	95%
1	90.0%	95.0%
2	81.0%	90.3%
3	72.9%	85.7%
4	65.6%	81.5%
5	59.0%	77.4%
6	53.1%	73.5%
7	47.8%	69.8%
8	43.0%	66.3%
9	38.7%	63.0%
10	34.9%	59.9%
11	31.4%	56.9%
12	28.2%	54.0%
13	25.4%	51.3%
14	22.9%	48.8%
15	20.6%	46.3%
16	18.5%	44.0%
17	16.7%	41.8%
18	15.0%	39.7%
19	13.5%	37.7%
20	12.2%	35.8%
21	10.9%	34.1%
22	9.8%	32.4%
23	8.9%	30.7%
24	8.0%	29.2%
25	7.2%	27.7%

FIGURE 3-8 Probabilities of simultaneous availability

Component Number	Average Availability Level	Probability of Simultaneous Availability
1	95.0%	95.0%
2	72.0%	68.4%
3	98.0%	67.0%
4	97.0%	65.0%
5	99.0%	64.4%

FIGURE 3-9 One problematic material

Thus a simple rule emerges with regard to dependent structures that contain integration points requiring simultaneous inputs to advance to the next stage of the structure or plan. This is a valid description of the plans that MRP generates. This simple rule is “delays accumulate, while gains do not.”

Figure 3-10 conceptually illustrates this effect. A dependent structure is visible at the bottom of the graphic. In this case that dependent structure is a synchronized plan based upon product structure. There are concurrent paths and integration points culminating in a finished item (FPA). Above the structure there is a graphical depiction of delay accumulation. The arrow steadily rises as activity progresses through the planned build. The arrow’s position at any one place depicts both how far along the planned activity path the build is (X axis corresponding to the structure) and the amount of accumulated delay (Y axis).

This effect is only partially impacted by signal accuracy. In other words, the demand signal could be perfect, but delay accumulation will still affect the environment if normal and random variation exist in the resources required to execute those signals. This delay accumulation results in an effect that is frequently referred to as “supply continuity variability.” This forces two profound realizations:

1. From an execution perspective MRP will *never* create a realistic plan in environments of even moderate complexity.
2. Any true solution to the bullwhip effect must address *both* demand signal distortion *and* the material supply distortion (supply continuity variability).

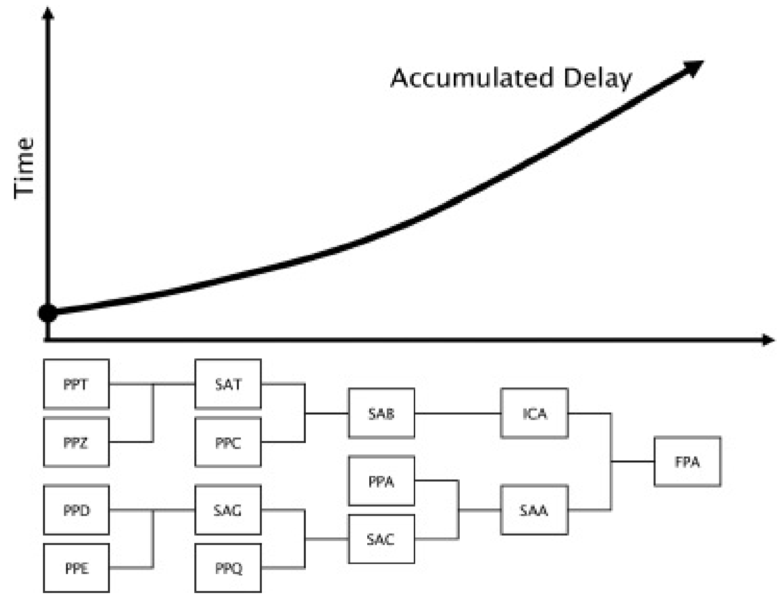

FIGURE 3-10 Illustrating delay accumulation

Amplifying the Distortions to Relevant Information and Materials—Batching Policies

The distortion to relevant information and material inherent in the bullwhip is amplified due to batching policies. Batching policies are determined outside of MRP and are typically formulated to produce better-unitized cost performance or are due to process restrictions or limitations. Batching policies dictate the way that MRP must perform its calculation (demand signal distortion) as well as influence the way in which materials progress through a supply chain (supply continuity variability).

The batching policies that dictate the MRP equation include order minimum—the amount that must always be ordered; order maximum—the largest quantity that can be assigned to an order; and order multiple—a rule that governs ordering between the minimum and the maximum. The order minimum and maximum should be evenly divisible by the order multiple.

An example: an intermediate component can have an order minimum of 100, a multiple of 50, and a maximum of 500. This means that if the intermediate component has a parent demand of 102 pieces, a minimum of 150 (the minimum plus the next multiple) of the component must be ordered to cover that demand. At some point later if the parent requirement changes to 99, the intermediate component requirement drops to 100. The parent changed by 3; the component changed by 50. The effect of this complication is devastating in any environment where ordering policies, particularly minimums and multiples, are dramatically different at each level of the bill of material.

Figure 3-11 is an example of the demand amplification in a more complex environment. An end item (FPA) has three components. All three components have minimums and multiples assigned. A demand of 115 for FPA will yield demands of 150 for Intermediate Component A (ICA), 250 for Intermediate Component B and 200 for Intermediate Component C.

Batching practices can dramatically affect the way that material moves in a supply chain, contributing to or amplifying the accumulation of delays. For example, delay accumulation could occur while an order waits on a truck for other orders to fill up the truck. The transportation batching policy dictates that only full trucks are allowed.

The logic and policies behind batching policies can be very problematic. Most batches are heavily influenced by an emphasis on protecting unit cost and have no consideration for flow. That emphasis on unit cost actually further distorts the flow of relevant information throughout most companies. The assumption that driving to unit cost performance equates to the best return on investment performance is unequivocally and mathematically proven false. Yet industry ignores this fact every day. This subject, however, is technically outside the scope of this text. For an in-depth look at this issue, refer to *Demand Driven Performance: Using Smart Metrics* by Debra Smith and Chad Smith.

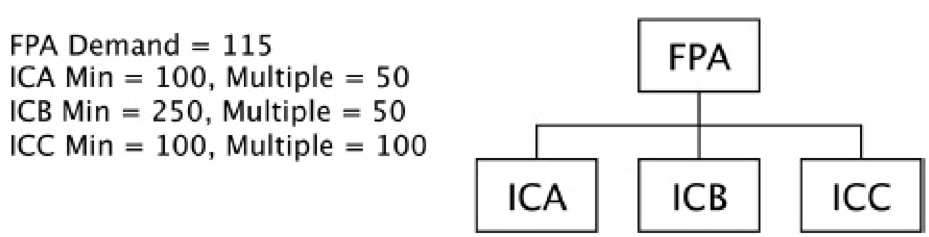

FIGURE 3-11 Batching complications to MRP supply order calculations.

Summary

Are the challenges described in this chapter and Appendix A unknown to seasoned planning professionals? Absolutely not. These challenges are well known and common. They explain the existence of the poor asset performance, the work-arounds, the bimodal inventory distribution, and the bullwhip effect. Additionally, they leave most planning organizations in a huge dilemma: utilize

MRP or ignore it. The answer to this dilemma is almost always the same; do both. Most organizations are simultaneously ignoring and utilizing MRP. Just how much ignoring and utilizing is something that tends to be specific to organizational functions and the individual users. There has to be a better way.

MRP enabled organizations to quickly calculate and synchronize total requirements given a set of demand inputs. This was of particular importance when the company had a deep bill of material or many shared components. The whole purpose of MRP was to synchronize connections and dependencies. In the New Normal there are undoubtedly more connections and dependencies than ever. Thus MRP should be more relevant today than ever. Yet MRP is failing in the New Normal.

MRP's role in the modern supply chain is significant. Even in the New Normal, the heart of every supply chain is manufacturing, and at the heart of manufacturing is MRP—it is the conductor of the supply order symphony in every supply chain. Each node in the supply chain has an MRP system supporting a different manufacturing operation. Therefore, a primary limitation of any supply chain will be how well MRP systems perform not just individually at each node but also collectively throughout the web.

If industry wants more agile manufacturing and supply chains that protect and promote the flow of relevant information and materials, then industry will need a more agile form of MRP. As evidenced in this chapter, companies cannot simply expect to implement conventional MRP better to get the necessary protection and promotion of flow. The first building block of a more agile form of MRP will be explained in the next chapter. This building block will mitigate if not largely eliminate the bullwhip effect by simultaneously addressing *both* demand signal distortion *and* material supply distortion by dealing with the core problem driving the bullwhip effect. This building block is called "decoupling."

CHAPTER 4

Unlocking a Solution—The Power of Decoupling

This chapter establishes a primary solution element to eliminate the bullwhip effect and create a framework for a proven and practical method of planning and execution for the conditions of the New Normal.

Chapter 3 described how the conventional planning approach featuring Material Requirements Planning (MRP) directly leads to the distortions of relevant information and materials that comprise the bull-whip effect. Figure 4-1 summarizes the connection between MRP's core trait of making everything dependent (our previously alluded to core problem) and the distortions to relevant materials and information. The boxes at the tips of the arrows are effects of the boxes at the tail of the arrow.

At the bottom of Figure 4-1 there is a rounded box with the words "MRP treats everything as dependent." There are two primary paths that lead from this box. The first path has to do with distortions to relevant information. That path is noted with dashed rounded boxes with no fill. This path shows that since MRP treats everything as dependent then manufacturing and procurement cycles are simply too long to respond to actual demand. This forces the use of forecasted demand which means the initial signal is in error by definition and that the demand signals will change as the incorporation of actual demand or changes to forecast occur. This triggers nervousness which creates constantly changing signals or leads to distortive behaviors to compensate for the nervousness (weekly buckets and/or BOM flattening).

Figure 4-1 culminates with an effect of distortions to relevant materials. Of course, it will be very difficult to have the "the right material at the needed time" if relevant information is distorted. But even if relevant information was not distorted, if demand was known and accurate and did not change, the effect that “the right material is not ready at needed time” would still exist. This is the second path depicted by the shaded boxes to the right. Since MRP treats everything as dependent, then all of the timing and quantity requirements in its plans are subject to those dependencies. Chapter 3 shows how dependent networks suffer performance erosion. An MRP plan, even with perfect demand information, will only be realistic if everything goes exactly according to plan with no variation.

This core problem of MRP has remained in place in large part because calculation dependency was developed as the real power of the MRP tool. If this dependency calculation was removed, then the true value of the MRP tool has also been removed. Yet as described in Chapter 3 and in Figure 4-1, this trait is the primary culprit in creating the transference and amplifications of variability to the flow of relevant information and materials. Failing to deal with this trait and its effects will guarantee that system flow and return on investment performance will be subpar.

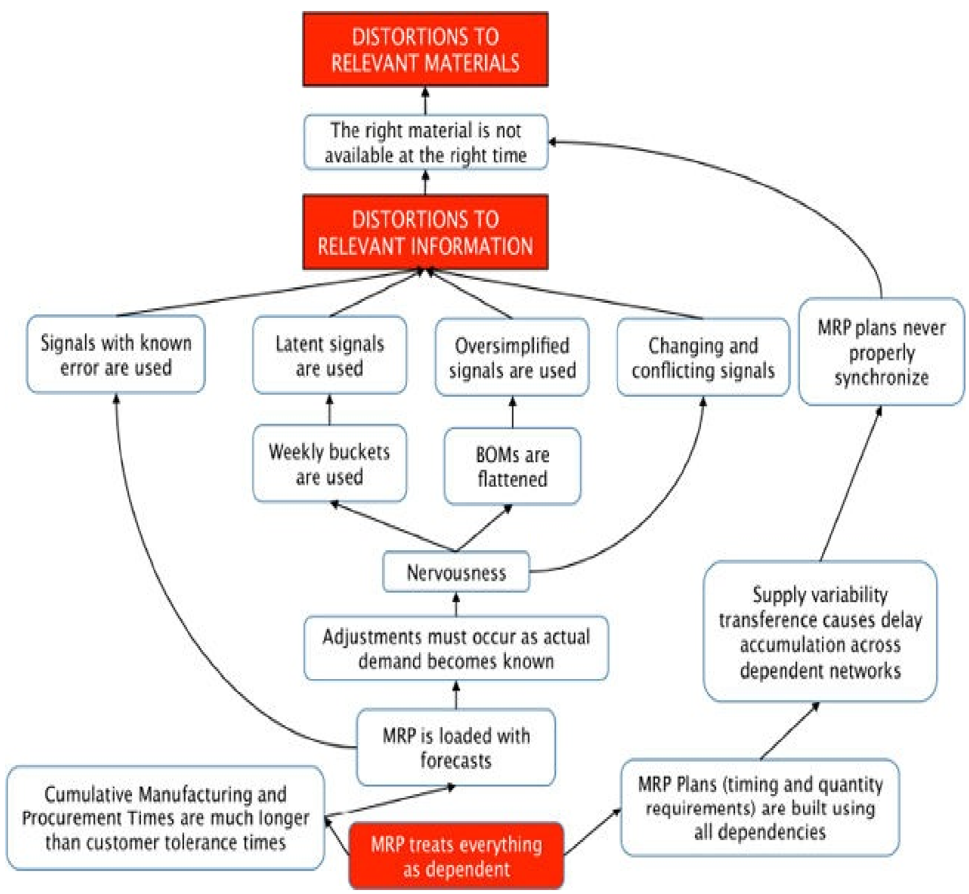

FIGURE 4-1 The core problem of the bullwhip

If the transfer and amplification of variability in the form of demand signal distortion and supply continuity is the biggest enemy to system flow, then we have to design supply chain capability that stops or mitigates the transfer and amplification of variability through the system. But how to do that? The answer cannot be “guess better” or “eliminate all variability.” Industry has tried that for decades, spending fortunes on reengineering efforts and expensive software only to see the problem persist.

Decoupling

The accumulation and impact of supply and demand variability is the enemy of flow. Variability can be systematically minimized and managed, but variability will never be eliminated. The only way to stop nervousness and the bullwhip effect is to stop variation from being passed between the linked parts of the supply chain in both directions.

This is accomplished through a concept called “decoupling.” APICS defines decoupling as:

> Creating independence between supply and use of material. Commonly denotes providing inventory between operations so that fluctuations in the production rate of the supplying operation do not constrain production or use rates of the next operation. (p. 43)

Decoupling disconnects one entity from another. This isolates events that happen in one entity or portion of a system and keeps them from impacting other entities or portions of the system. Think of decoupling as if it were a firewall in a building, automobile, or information system or a break wall around boats in a marina.

The concept of decoupling provides the fundamental break from convention that is needed to mitigate variability. Decoupling breaks the direct connection between dependencies. The places at which the system is decoupled are called “decoupling points.” APICS defines decoupling points as:

> The locations in the product structure or distribution network where strategic inventory is placed to create independence between processes or entities. Selection of decoupling points is a strategic decision that determines customer lead times and inventory investment. (p. 43)

Decoupling is not a new idea. The concept has been around for many years but with no practical way to implement it in MRP. MRP was designed with the explicit intention of tightly coupling everything so that precise equations could be performed in order to synchronize the environment. Limited amounts of decoupling can occur in MRP, but only with complications where costs likely outweigh their benefits (this is discussed further in Chapter 9).

Figure 4-2 is based on Figure 3-3 and depicts the dependent view of an MRP system and the accumulated demand signal distortion (the upper arrow moving right to left) and the supply continuity variability (the lower arrow moving left to right). There is no decoupling; thus the distortion to both relevant information and materials accumulates through the system.

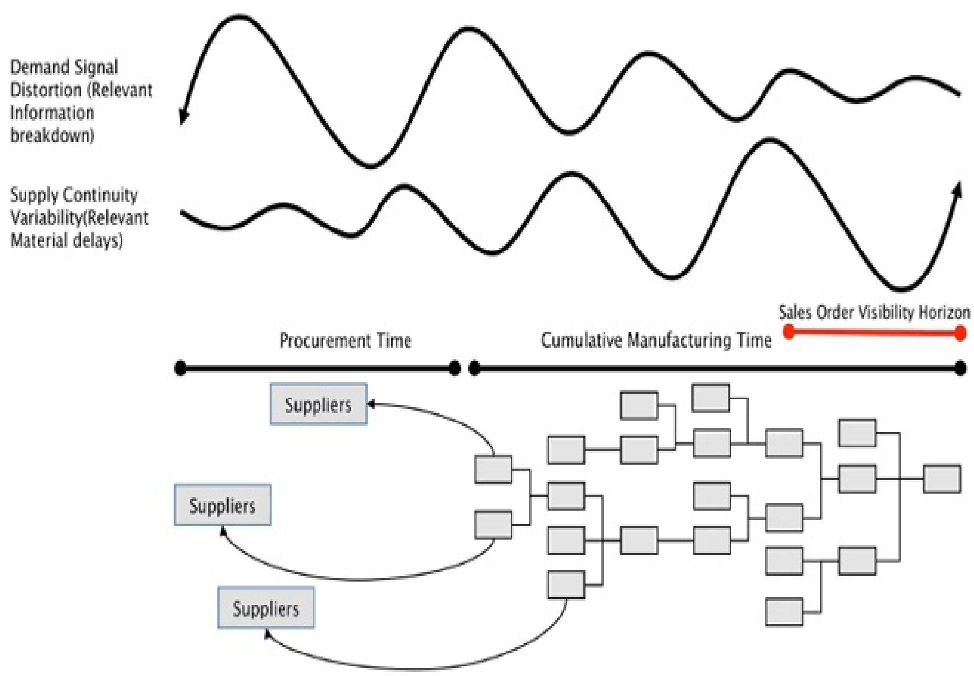

FIGURE 4-2 The MRP approach

Decoupling points represent a place to disconnect the events happening on one side from the events happening on the other side. They delineate the boundaries of independently planned and managed horizons. The determination of their placement is not to be taken lightly—it is a strategic decision that will dramatically affect how the system operates and how effective the overall system will be.

Decoupling Point Buffers

For the decoupling points to maintain their decoupling effect, there must be a level of protection that absorbs demand and supply variability at the same time. This level of protection is a concept called “decoupling inventory.” APICS defines decoupling inventory as:

> An amount of inventory kept between entities in a manufacturing or distribution network to create independence between processes or entities. The objective of decoupling inventory is to disconnect the rate of use from the rate of supply of the item. (p. 43)

Decoupling point inventory in this book will be referred to as a “decoupling point buffer” or simply “buffer.” Decoupling point buffers are quantities of inventory or stock that are designed to decouple demand from supply. Buffers are commonly amounts of inventory that will provide reliable availability to the consumers of the stock while at the same time allowing for the aggregation of demand orders, creating a more stable, realistic and efficient supply signal to suppliers of that stock.

Figure 4-3 depicts the same system as Figure 4-2 but with decoupling point buffers in place. The placement of decoupling point buffers (represented as the tiered bucket icons in the dependent structure) creates independent planning and execution horizons. These horizons are indicated by the dotted lines with rounded terminal points on each side. Demand and supply variability is stopped from further accumulation at those terminal points. This is represented by the wall-like icons labeled “break wall.” This means that the use of decoupling point buffers addresses both components of the bullwhip at the same time and from the same place; it is a bidirectional solution.

Decoupling buffer placement has huge implications for lead times. By decoupling supplying lead times from the consumption side of the buffer, lead times are instantly compressed between buffers and to the customer. This lead time compression has immediate service and inventory implications. Market opportunities can be exploited, while working capital required in buffers placed at higher levels in the product structure can be minimized.

Furthermore, Figure 4-3 reveals an additional lead time compression benefit due to decoupling: its impact on relevant information. As discussed in Chapter 3, MRP users are forced to make commitments to a demand signal that is subject to varying degrees of error (forecasted orders). While there is an alternative much more accurate demand signal (sales orders), MRP’s inability to decouple prevents the exclusive use of that signal.

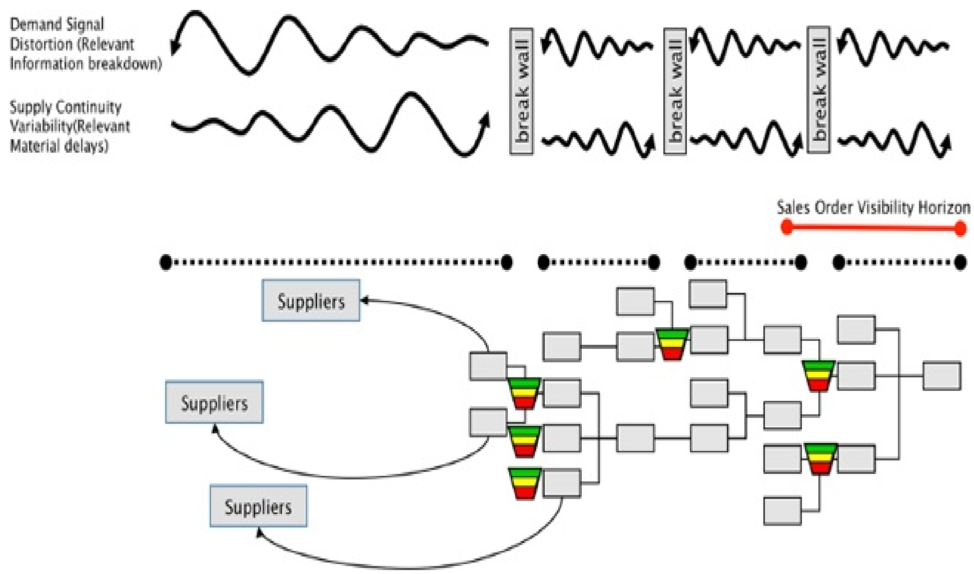

FIGURE 4-3 A system with decoupling point buffers

Yet what Figure 4-3 shows is that when a decoupling point buffer is placed inside the sales order visibility horizon, it will allow for the system to exclusively use that accurate demand input. We have effectively found the time that we believed we lacked that forced the use of forecasted orders in the first place. When decoupling point buffers match the sales order visibility horizon, the demand variability is reduced.

Summary

Decoupling simply makes sense given the basic circumstances that we face today. We have extended supply chains globally and made them more complex and fragile. These longer and more complex supply chains are subject to much higher levels of variability and are much harder to plan. Breaking dependencies in key places will dramatically simplify the planning equation and provides shorter horizons with much more relevant information.

The concept of decoupling poses an ironic situation. In order to promote and protect the flow of relevant information in a system, you must strategically and purposefully slow or interrupt flow at certain critical points. The size of the decoupling point buffers defines the length of the slowdown or interruption at these caching points.

Unfortunately, conventional planning systems are designed to position and then manage decoupling points. The very basic foundation of Material

Requirements Planning was to make everything dependent and only order what was needed, when it was needed, in a mathematically precise way. Decoupling creates a position of independence. The inability to decouple is the primary culprit behind the bullwhip effect and is a major impediment to system flow.

Decoupling is the key that unlocks a decades-old struggle with conventional planning approaches utilizing MRP, a struggle that is becoming more acute in the New Normal. It allows a door to open to a place where daily planning can become obvious, intuitive, and beneficial for supply order generation and management. What is needed is a systematic approach for utilizing decoupling that fundamentally answers these key questions:

- Where to place these decoupling points? The answer is neither "everywhere" nor "nowhere." The answer is simply stated as "somewhere." But how to find that best somewhere? Placing decoupling points will be the subject of Chapter 6.
- How to size the protection at the decoupling point? In order to maintain the integrity of the decoupling point, the buffers must be sized in relation to the specific attributes of the parts, planning, and execution horizons they are protecting. This will be the subject of Chapter 7 and 8.
- How to maintain that protection? Supply orders must be generated and managed in a way that keeps the points properly supplied and intact. These techniques will be explored in Chapters 9 and 10.

PART 2

Becoming Demand Driven

CHAPTER 5

Supply Order Generation and Execution for the New Normal

At this point we are at a crossroads. We can continue to struggle with our conventional planning systems, or we can seek a break from convention, an alternative designed for the New Normal. That alternative design must promote and protect the flow of relevant information and materials. It must systematically break down the distortions to demand signals and material supply that characterize the bullwhip effect through the effective use of decoupling points. But where to go from here?

MRP Versus Lean—What Can We Learn?

The basic elements of this alternative design can be better understood by exploring a chronic conflict between two camps—the believers in Lean methodology and the believers in MRP.

Figure 5-1 illustrates a side-by-side comparison of conventional MRP and Lean approaches. On the MRP side, forecasted demand feeds a master production schedule (MPS). The MPS creates a statement of what will be built. This is then fed to MRP. MRP then explodes through the bill of material, creating synchronized supply orders (date and quantity) as dictated by the product structure. Safety stock is often used at the finished and purchased parts levels in order to absorb variability.

The Lean approach establishes kanban positions, which are independent inventory positions typically placed at each resource position. The kanbans are sized according to a required takt time rate. This rate can be established through a forecast or past consumption. The kanbans are connected with “loops” that provide easy-to-interpret signals for each position to produce or not produce. A "supermarket" can be placed at the intermediate or purchased component level that produces the same easy-to-interpret signal as the kanban. The difference between the supermarket and the kanban is simply that the supermarket is at the part (product structure) level and the kanban is placed at the resource level.

Many Lean implementations attempt to abandon the formal planning approach of MRP because it is seen as inappropriate, transaction intensive, non-value added, even antithetical, to what Lean is trying to accomplish. This causes tremendous friction between planning personnel and those pushing to eliminate these systems. Lean facilitators and advocates often see MRP as an overcomplicated and wasteful dinosaur that simply doesn't work in a customer-centric world.

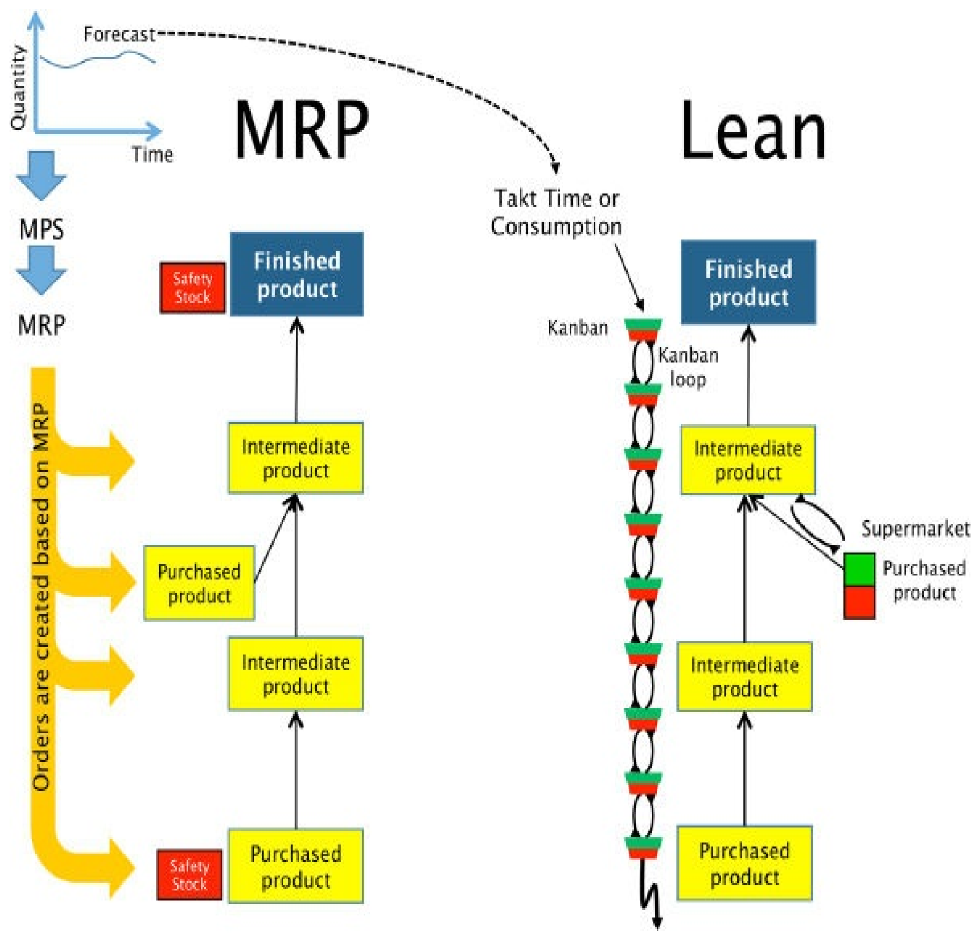

FIGURE 5-1 MRP and Lean comparison

Planning personnel, however, see it in a completely different way. They believe that without the ability to see and synchronize complex and dynamic environments, critical blind spots will exist in the planning process which will lead to shortages, expedites, and even excessive inventory positions to compensate. They see the simple pull approach for managing materials and inventory as a gross oversimplification for the complex planning and supply scenarios that are the norm in today's more volatile environment.

What if both camps are right? What if in many environments today the traditional MRP approach is too complex *and* the Lean approach is too simple? Where would that leave supply chain management? It would create a situation where companies oscillate between the two options, depending on the political wind employing a constantly changing and unsatisfactory number of work-arounds and compromises. Executives get frustrated, in-fighting escalates, efforts are sabotaged, more money and time are spent, and improvements deteriorate to lip service.

Einstein once said, "Any intelligent fool can make things bigger and more complex. It takes a touch of genius—and a lot of courage—to move in the opposite direction." He also said, "Everything should be made as simple as possible, but not simpler." How prophetic.

Can traditional MRP be overly complex? *Without a doubt.* Most people in manufacturing companies don't even fully understand what the planning system is or how it does what it does. Every day, planners are drowning in oceans of data and action messages. The hard-coded rules are rooted firmly in the old "push and promote" methodology that makes MRP ill-suited to today's more volatile and service-oriented world. Furthermore, "fixing" or "cleaning up" the system seems to be a never-ending, transaction-intensive, and expensive journey; the end of which is to always be precisely wrong.

Can Lean be an oversimplification? When it comes to materials and inventory planning, the answer in many environments is *yes.* Oversimplification is defined as "To simplify to the point of causing misrepresentation, misconception, or error."[1] By failing to provide visibility to critical dependencies and relationships with regard to supply, demand, on-hand inventory, and product structure, the Lean tool set can attempt to oversimplify many environments. The larger, more complex, and variable these environments are, the more likely that Lean's simple kanban controls and lack of material planning create an oversimplified approach.

Yet there is one thing they actually agree on: a common objective. Both camps can agree that flow is paramount. Now more than ever, a decisive competitive advantage can be achieved by companies with a high degree of flow through and to their customers. Chapter 2 highlights the need for and benefits of flow. The better the flow of relevant information and materials, the better the service levels and use of working capital. The better the service levels and working capital, the better the bottom line.

Do MRP advocates disagree with this? Certainly not! Materials and processes that flow reliably are the easiest to plan and manage. Having the right things at the right time is the key to flow. MRP's entire reason for existence was to attempt to synchronize environments to have the right things at the right time. But does conventional MRP have deficiencies that hurt flow? Undeniably *yes!*

Do Lean advocates agree with the need for flow? Absolutely. Information and materials that flow reliably generate considerably less waste. But does Lean have a complete tool set for fully protecting and improving flow at the plant, enterprise, and supply chain level in a more complex and volatile world? There seems to be something missing.

So if there is a common objective, why are these camps locked in a chronic conflict? As discussed previously, MRP is a perfect just-in-time system that nets to zero inventory. This sounds incredibly compatible with Lean's approach. Yet the conflict persists because they represent diametrically opposed approaches with regard to two critical factors. These critical factors are essentially two sides of the same coin but are worth discussing specifically. Any solution must involve addressing this inherent opposition in these two areas.

Dependence Versus Independence

Figure 5-2 illustrates the area of this particular conflict. At the top of the structure is the common objective of protecting and promoting flow. MRP and Lean have different critical needs in order to accomplish this objective, and each method has a specific attribute designed to secure its respective critical need.

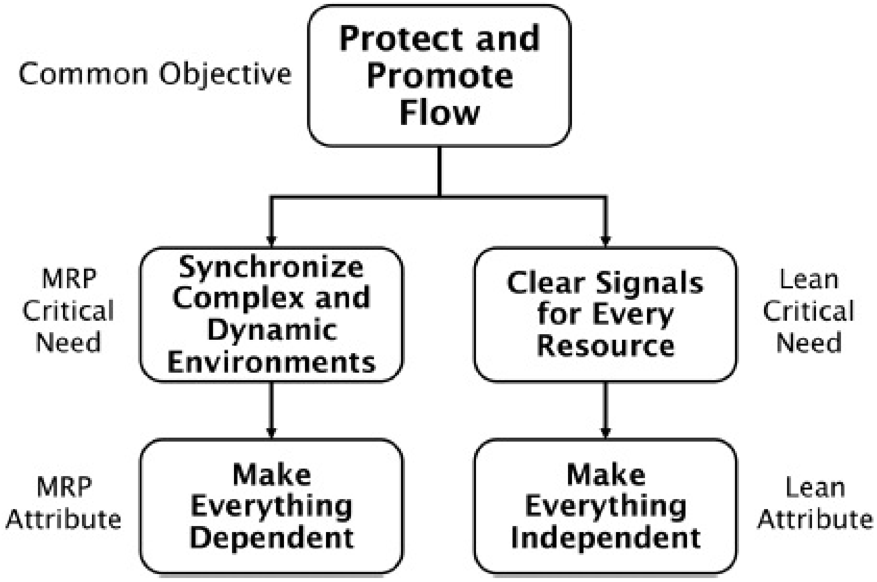

FIGURE 5-2 The dependence versus independence conflict between MRP and Lean

We will explore the MRP side first. MRP advocates understand that the protection and promotion of the flow of relevant information cannot occur without the ability to synchronize complex and dynamic environments. In order to accomplish this synchronization, MRP, as discussed in Chapters 3 and 4, is hard coded to obey the dependencies defined by product structure. A change anywhere creates change everywhere. MRP was designed in this way in order to make sure that the operating environment could understand the impact of changes as they occur.

On the Lean side we see that the protection and promotion of flow requires that resources have clear signals to operate by. When signals become confusing or conflicting, a resource's ability to determine what is correct breaks down. Too many points of data or constantly changing signals create that confusion or conflict. Thus Lean makes everything independent. Resources only need to look to one place (the kanban that they feed) to determine if they should produce or not produce. It is literally that simple.

Making everything dependent versus making everything independent is certainly mutually exclusive. In fact, both sides' attributes break down the other side's critical need. By making everything dependent, MRP creates an incredibly confusing set of constantly changing and conflicting signals. By making everything independent, an environment loses the ability to synchronize to changes that can and will occur. This is particularly true in environments characterized by heavy demand fluctuations, long lead time parts, shared resource bases, and product innovation.

Is there an alternative that can have the benefits of dependence (synchronization) and the benefits of independence (clear signals) at the same time without conflict?

Supply Order Generation (Planning Versus Execution)

Figure 5-3 illustrates the area of another conflict. Once again, at the top of the structure is the common objective of protecting and promoting flow. MRP's need remains unchanged from the previous conflict: synchronize complex and dynamic environments. Lean has a different critical need.

Both sides have a different attribute designed to secure their respective critical need with regard to supply order generation. One could argue that this is simply an extension of the dependence versus independence conflict. That is a valid argument, but there is an additional level of insight that could be brought to bear by discussing them separately.

As described in depth in Chapter 2, MRP is typically loaded with forecasted demand in order to attempt to synchronize the long manufacturing and procurement cycles with anticipated demand. This happens well in advance of consumption. The Lean side, however, seeks to protect and promote flow by pacing to actual demand because the inherent inaccuracy with forecasts directly impedes flow. Resources are squandered on things that are overproduced and overordered in advance, while expedites must accommodate the things that were underproduced or underordered as the picture becomes clearer. The only way to truly know if demand is real is after it has occurred. Consumption is definitive proof of demand.

As with the previous conflict, the attributes seem to be mutually exclusive. When MRP generates supply orders well in advance of anticipated consumption, it loses the capability to pace to actual demand at least by the amount of the forecast error. The longer the planning horizon, the greater the forecast error. When we generate supply orders at the execution level, there is a delay in responding to significant changes, as the supply orders must make their way through the entire connective structure one level at a time; there is rudimentary but extremely slow synchronization at best.

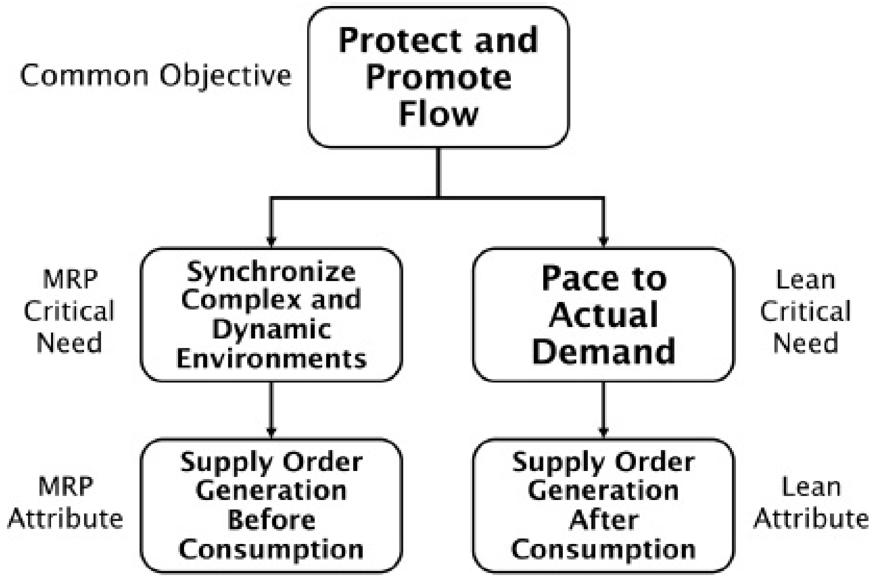

FIGURE 5-3 The supply order generation conflict between MRP and Lean

Is there an alternative that can have the benefits of supply order generation at the planning level (synchronization) and the benefits of supply order generation at the execution level (pace to actual demand)?

It is vital to understand that all the critical needs are required to protect and promote flow. Planning and execution systems must pace to actual demand, provide clear signals for resources, and synchronize complex and dynamic environments.

Focusing on only one critical need and discounting the others almost guarantees challenges to flow. Indeed that seems descriptive of the impasse between the MRP and Lean worlds. When it comes to the protection and promotion of the flow of relevant information and materials, both Lean and MRP have weaknesses in today's more volatile and complex environments. Lean's reliance on independent replenishment kanbans with little to no visibility or connectivity at the plant, enterprise, and product structure level is a problem for the protection and promotion of flow. But the antiquated and complex rules of conventional MRP that govern demand and supply order generation create unrealistic, constantly changing, and generally confusing plans and schedules.

To protect and improve flow, a blend of simple visible pull signals *and* the computational and connective power of technology is necessary. This isn't a compromise for the two sides to live in peace; it must be a harmonious integration where both sides' critical needs are incorporated to create a stronger solution for the protection and promotion of flow. And it must be practical, consistent, and scalable!

What if there is a way to define a solution (rules and tools) that is not overly complex or overly simple? What if there is a way to take key and relevant aspects of both points of view and create an elegant blueprint that will work for and enhance both sides' objectives? This solution must include a level of sophistication that can provide more visibility and synchronization from a planning and execution perspective while at the same time pace to actual demand and promote simple, clear, and highly visible signals across the enterprise. This solution is called Demand Driven Material Requirements Planning (DDMRP).

Lean and Technology

Lean advocates often get accused of being anti-technology, but do Lean advocates really want manufacturing companies to entirely abandon the promise of technology? The answer should be yes when that technology is wasteful, confusing, and not reflective of reality, when it force-fits concepts so as to simply permit their use. Unfortunately, this has been the situation for quite some time with regard to traditional MRP and DRP (distribution requirements planning) systems.

Point 8 of the Toyota Production System states, "Use only reliable, thoroughly tested technology that serves your people and processes."[2] Until now the prevailing materials and inventory planning and execution technology, while thoroughly tested, have been largely inappropriate to serve the people and processes in companies transforming to a leaner approach. Chapter 3 clearly demonstrated that point.

Yet the proliferation and sustainability of Lean implementations has been negatively impacted by the lack of appropriate supply chain materials planning and execution technology. Many well-respected manufacturing analysts have concluded that there is tremendous potential for the incorporation of better planning and visibility software into Lean implementations. Manufacturing needs Lean methods to stay competitive in the more complex environment of the twenty-first century. Lean needs an effective customer-centric planning approach to bring that vision to reality.

What if there were an appropriate technology? What if a reliable, thoroughly tested method for a customer-centric pull-based planning and execution of supply chain materials with high degrees of visibility could be introduced to the

MRP world? Under that condition Lean and MRP would both find an effective solution. This method is Demand Driven Material Requirements Planning.

Demand Driven Material Requirements Planning

This section will serve to introduce DDMRP—its basic foundation and its major components. But first it may help to understand what the term "demand driven" really means and the history behind it. The term was originally defined as the ability to "sense changing customer demand and adapt planning and production while pulling from suppliers all in real time."

The History of "Demand Driven"

The term was pioneered by PeopleSoft in 2002 while Carol Ptak was the vice president of manufacturing and distribution industries. When Oracle acquired PeopleSoft in 2003, the term was largely abandoned. It was then resurrected in 2007 by AMR. In 2010 AMR was acquired by Gartner, and Gartner used the term as part of what it called its "Demand Driven Value Network" approach.

In 2011 the third edition of *Orlicky's Material Requirements Planning* (Ptak and Smith) introduced the initial blueprint for Demand Driven Material Requirements Planning as an alternative formal planning and control logic. The year 2011 also marked the foundation of the Demand Driven Institute by Carol Ptak and Chad Smith. The Demand Driven Institute has published several white papers and case studies on the DDMRP topic. A repository of case studies and white papers on DDMRP is available free at http://www.demanddrivenworld.com.

In 2012 the Demand Driven Institute partnered with the International Supply Chain Education Alliance (ISCEA) to offer the Certified Demand Driven Planner (CDDP) program. The CDDP program was designed to provide consistent global standards for the DDMRP approach and to teach and certify practitioners in those standards. From 2012 to 2015 over 1,000 people took the CDDP program on six continents.

In 2013 *Demand Driven Performance: Using Smart Metrics* was written by Debra Smith and Chad Smith. This book extended the term across the enterprise into finance, scheduling, shop floor control, and strategy, effectively defining the Demand Driven Operating Model. This will be defined and discussed later in this chapter.

In 2016, with *Demand Driven Performance: Using Smart Metrics* as a guide, the Demand Driven Institute and ISCEA released the Certified Demand

Driven Leader (CDDL) program. The CDDL program was designed to provide consistent global standards for the Demand Driven Operating Model (including DDMRP) and to teach and certify practitioners in those standards.

Position, Protect, and Pull

The original definition of "demand driven" still works in today's more mature and larger demand driven body of knowledge. Additionally, this maturation and expansion has provided clarity on precisely what demand driven does not mean. It does not mean "make to order everything." It does not mean "put inventory everywhere." It does not mean "forecast better." Becoming demand driven requires a fundamental shift from the centrality of supply- and cost-based operational methods (commonly referred to as "push and promote") to a centrality of actual demand- and flow-based methods (commonly referred to as "position, protect, and pull"). The term "actual demand" is extremely important in distinguishing it from a rebranded and somehow superior forecasting approach.

Demand Driven Material Requirements Planning is a formal multi-echelon planning and execution method to protect and promote the flow of relevant information and materials through the establishment and management of strategically placed decoupling point stock buffers. DDMRP has roots in many conventional methods. Figure 5-4 shows the methodological foundation for DDMRP.

DDMRP combines some of the still relevant aspects of MRP and DRP with the pull and visibility emphases found in Lean and the Theory of Constraints and the variability reduction emphasis of Six Sigma. Do these methods all just naturally fuse together? No. At a minimum, as noted earlier in this chapter, there are conflicts between Lean and MRP. A similar conflict occurs with MRP and the Theory of Constraints. The final component of this fusion requires a few key innovations that are unique to DDMRP.

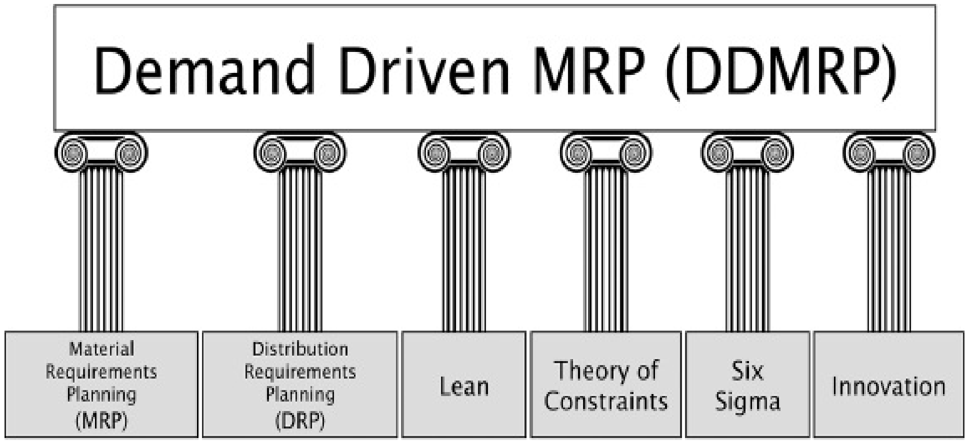

FIGURE 5-4 The methodological foundation of DDMRP

Demand Driven Material Requirements Planning has five sequential components. Figure 5-5 illustrates these components, their sequence, and how they relate to the mantra "position, protect, and pull." These five components are respectively featured in sequential chapters starting with Chapter 6.

The first three components essentially define the initial and evolving configuration of a Demand Driven Material Requirements Planning Model. Strategic inventory positioning will determine where the decoupling points are placed. Buffer profiles and levels will determine the amount of protection at those decoupling points. Dynamic adjustments define how that level of protection flexes up or down based on operating parameters, market changes, and planned or known future events.

The fourth and fifth elements define the actual operational aspects of a DDMRP system: planning and execution. In DDMRP, demand driven planning is the process by which supply orders (purchase orders, manufacturing orders, and stock transfer orders) are generated. Visible and collaborative execution is the process by which a DDMRP system manages open supply orders.

DDMRP is at the heart of the Demand Driven Operating Model, which is defined this way:

> Demand Driven Operating Model—a supply order generation, operational scheduling, and execution model utilizing actual demand in combination with strategic decoupling and control points and stock, time, and capacity buffers in order to create a predictable and agile system that promotes and protects the flow of relevant information and materials within the tactical relevant operational range (hourly, daily, and weekly). A Demand Driven Operating Model's key parameters are set through the Demand Driven Sales and

Operations Planning process to meet the stated business and market objectives while minimizing working capital and expedite-related expenses.

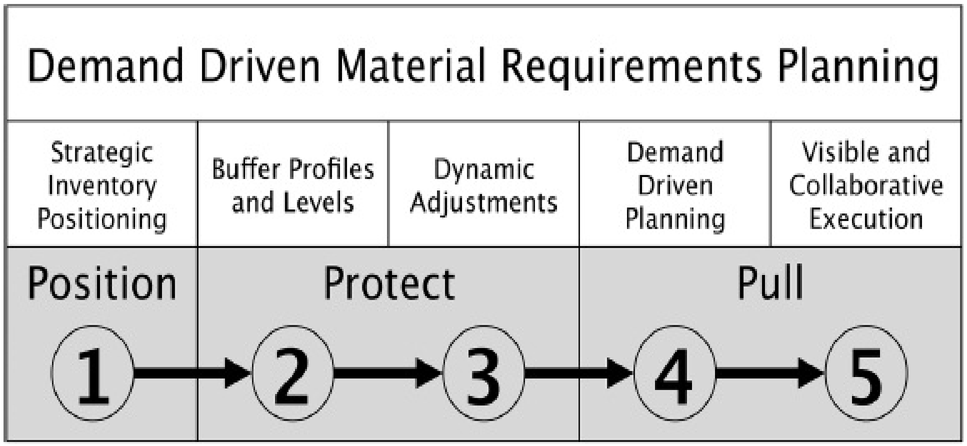

FIGURE 5-5 The five components of DDMRP

Figure 5-6 depicts the Demand Driven Operating Model schema. Model and part parameters, commonly referred to as "master settings," are supplied by the Demand Driven Sales and Operations Planning (DDS&OP) process to Demand Driven MRP. These settings will be thoroughly described in Chapters 6, 7, and 8. DDS&OP is the subject of Chapter 13. These master settings configure the DDMRP system and are combined with inventory and product structure records and actual demand to generate supply orders. If these supply orders are manufacturing orders, they are sent to scheduling. If the supply orders are purchase or stock transfer orders, they go directly into execution.

This book provides limited content on Demand Driven Capacity Scheduling or its respective master setting inputs. What content there is on this topic is found in Chapter 11. Additionally, the execution aspects in this book are confined to the execution elements of DDMRP—the management of open supply orders. For in-depth content on Demand Driven Capacity Scheduling and related shop floor execution tactics, refer to Smith and Smith's *Demand Driven Performance: Using Smart Metrics*. Figure 5-7 depicts the aspects of the Demand Driven Operating Model covered in the subsequent chapters of this book.

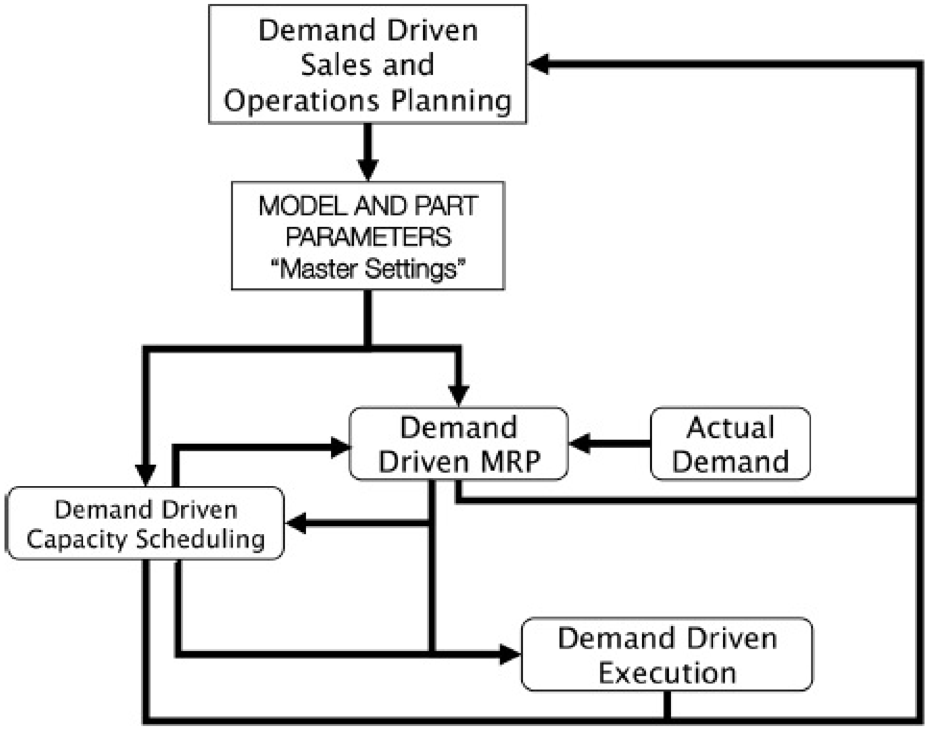

FIGURE 5-6 Demand Driven Operating Model schema

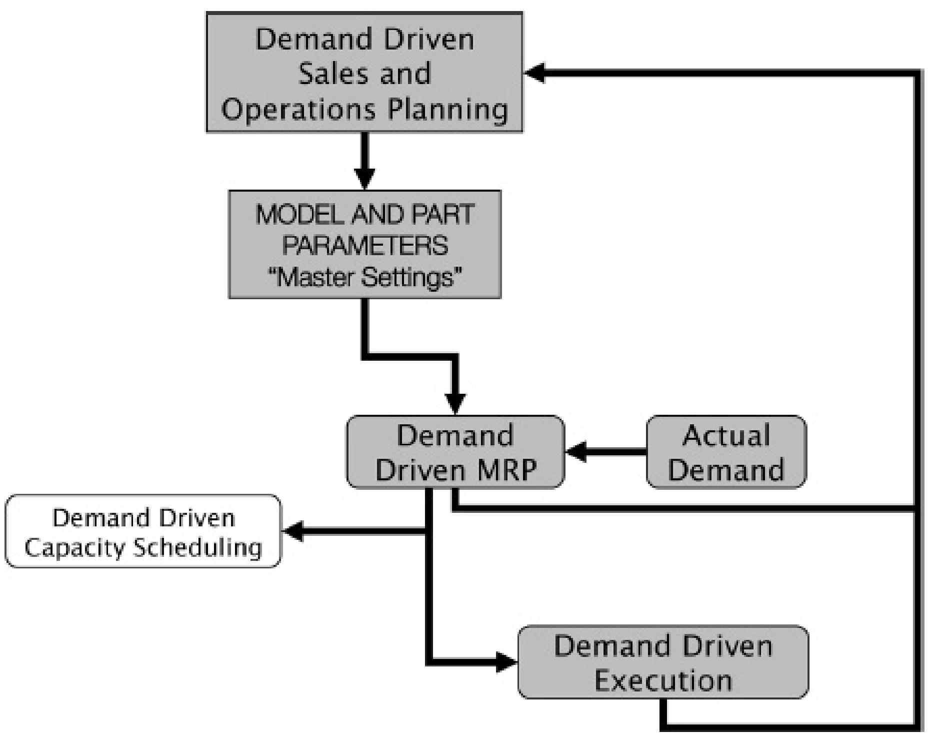

FIGURE 5-7 The emphasis of this book

Summary

This chapter combined with Chapters 3 and 4 has laid the basic foundation for the critical elements of a new supply order generation method for the New Normal called Demand Driven Material Requirements Planning.

What do we know so far about the requirements for DDMRP?

1. It should be based on the protection and promotion of the flow of relevant information and materials. This connects it to driving better returns on investment. This was explained in Chapter 3.
2. It must allow for decoupling in order to mitigate demand signal and supply continuity variability as well as to compress lead times. This was explained in Chapter 4.
3. It should use the most relevant demand information available—actual demand. The problems associated with forecast error were described in Chapter 3 and previously in this chapter
4. It must provide easy-to-interpret signals for all resources. This was described previously in this chapter.
5. It must provide for a way to synchronize complex and dynamic environments.

CHAPTER 6

Strategic Inventory Positioning

At the risk of oversimplifying the everyday tasks of buyers and planners, we should understand that they are constantly dealing with two questions of supply order management. The two questions are these: How much and when? Hundreds if not thousands of books have been written about a wide variety of techniques and tricks to attempt to answer these questions.

The question of how much is a question concerning quantity. Planners and buyers are continually validating, verifying, and supplementing how much they really need versus what MRP is telling them. The question of when is simply a question of timing. Planners and buyers are continually validating, verifying, and supplementing when they really need things versus what MRP is telling them. This is a constantly changing series of wrong answers as system nervousness and the bullwhip impact the environment.

Thus their daily objective degenerates to simply being less wrong. They are constantly challenged about how they historically answered these questions and why things are not available in the time or quantity that they are needed. A common practice is for the planners or buyers to save screenshots of the MRP system in order to create a defense for why they did what they did and when they did it. A frustrating situation indeed.

Perhaps all this activity and series of constantly dissatisfactory answers is not related to the questions, how much and when? Perhaps it is first and foremost related to our failure to ask a more fundamental question.

As discussed earlier, the key to protecting and promoting the flow of relevant information requires the use of decoupling points. Decoupling enables a bidirectional benefit—it mitigates the demand signal distortion (relevant information) and supply continuity variability (relevant materials) inherent in the bullwhip effect. But this raises a question—where should these decoupling points be placed within a supply chain or organization to maximize effectiveness?

Positioning Factors

Most organizations are completely unprepared to deal with this question. First, they lack the knowledge, comprehension, or even capability to even ask the simple question, “Where?” Even if they do ask the right question, they lack the ability to effectively answer that question.

Thus the first component of Demand Driven Material Requirements Planning is determining where the decoupling points and their respective buffers should be placed. This component becomes the cornerstone of the Demand Driven Operating Model discussed in Chapter 5. The selection of these points is a strategic decision that impacts the performance of the supply-demand network in many regards: service, working capital, expedite-related expenses, cash flow, and ultimately return on investment.

Chapters 1 to 4 created the in-depth case about why the question of “where” must be asked. This chapter focuses on how to properly answer this question through the consideration of six key factors.

Customer Tolerance Time

This is the time the typical customer is willing to wait before seeking an alternative source. Customer tolerance time also can be referred to as demand lead time. According to APICS, demand lead time is:

> The amount of time potential customers are willing to wait for the delivery of a good or a service. Syn: customer tolerance time. (p. 45)

Determining this lead time often takes the active involvement of sales and customer service.

Market Potential Lead Time

This lead time will allow an increase of price or the capture of additional business through either existing or new customer channels. Determining this lead time takes the active involvement of sales and customer service. Be aware that there could be different stratifications of market potential lead time. For example, a one-week reduction in lead time may only result in an increase in orders, whereas a two-week reduction in lead time could result in both an increase in orders and a potential price increase on some of those orders. Properly segmenting the market will maximize the possible revenue potential for the company and provide excellent revenue growth control. This is a consideration in Demand Driven Sales and Operations Planning, covered in Chapter 13.

Sales Order Visibility Horizon

The sales order visibility horizon is the time frame in which we typically become aware of sales orders or actual dependent demand. In retail situations, customers do not issue a sales order to a shop in advance of going to the shop. Thus the sales order visibility horizon in this situation is zero. In most manufacturing scenarios, however, there are sales orders conveyed in advance of expecting receipt of the item. Often the sales order visibility either matches or exceeds customer tolerance time. The longer the visibility to sales orders, the better the capability of the environment to see potential spikes and derive relevant demand signal information. In many cases relevant requirements are obscured from planners because all demand (including planned orders based on forecast and safety stock requirements) is aggregated together for aggregate planning purposes.

External Variability

External variability considers both demand and supply variability.

Variable Rate of Demand

This refers to the potential for swings and spikes in demand that could overwhelm resources (capacity, stock, cash, etc.). This variability can be calculated by a variety of equations or determined heuristically by experienced planning personnel. As noted in the *APICS Dictionary,* "Mathematically, demand variability or uncertainty can be calculated through standard deviation, mean absolute deviation (MAD) or variance of forecast errors." If the data required for mathematical calculation do not exist, companies can also use the following criteria:

- **High-demand variability.** Products and parts that are subject to frequent spikes within the customer tolerance time.
- **Medium-demand variability.** Products and parts that are subject to occasional spikes within the customer tolerance time.

- **Low-demand variability.** Products and parts that have little to no spike activity. The demand is stable within the customer tolerance time.

Variable Rate of Supply

This is the potential for and severity of disruptions in sources of supply or specific suppliers. This can also be referred to as supply continuity variability. It can be calculated by examining the variance of promise dates versus actual receipt dates. When first considering the variable rate of supply, the initial variances can be caused by critical inherent flaws in the MRP system. Additionally, those dates often shift due to other shortcomings associated with the way MRP is employed rather than because of the supplier capability. Any critical supplier of a major manufacturer will know exactly which day its customer regenerates its MRP. These suppliers will see a flurry of additional orders, canceled orders, and changes to orders (quantity, specification, and request date).

If the data required for mathematical calculation do not exist, the following heuristics can be used:

- **High supply variability.** Frequent supply disruptions
- **Medium supply variability.** Occasional supply disruptions
- **Low supply variability.** Reliable supply

Inventory Leverage and Flexibility

There are places in the integrated bill of material (BOM) structure (matrix bill of material) or the distribution network that provide a company with the most available options as well as the best lead time compression to meet the business needs. Within manufacturing, these places are typically represented by key purchased materials, subassemblies, and intermediate components. This becomes more critical in environments with BOMs that are deeper and more complex (broader) and have more shared components and materials. This concept will be explored in detail later in this chapter.

Critical Operation Protection

Similar to how variability can impact a bill of material, the longer and more complex the routing structure and dependent chain of events (including interplant transfers), the more important it can be to protect identified key areas. These types of operations include areas where there is limited capacity, or where quality can be compromised by disruptions, or where variability tends to be accumulated or amplified. In Lean, these areas might be referred to as pacesetters. In the Theory of Constraints, they can be referred to as drums. Whatever manufacturing or operational methodology a company ascribes to, these resources typically represent control points that have a huge impact on the total flow or velocity that a particular plant, resource, or area can maintain or achieve.

The preceding six factors must be applied systematically across the entire BOM, routing structure, manufacturing facilities, and supply-demand network to determine the best decoupling positions for purchased, manufactured, and finished items (including service parts) in order to protect and promote the flow of relevant information and drive return on investment performance.

Applying the Positioning Criteria

As an example, let us apply these six factors to a relatively simple environment. In our example, only two finished products are made. Figure 6-1 shows the bill of material for the two products: FPE and FPF.

The numbers in the circles represent the manufacturing or purchasing lead time in days for each discrete part number. For instance, FPE takes 2 days to make when all components are available, and 204P has a purchasing lead time of 20 days.

For each part number in this example, there are three relevant lead times. These are described in the *APICS Dictionary* as:

> **Manufacturing lead time (MLT):** The total time required to manufacture an item, exclusive of lower level purchasing lead time. For make-to-order products, it is the length of time between the release of an order to the production process and shipment to the final customer. For make-to-stock products, it is the length of time between the release of an order to the production process and receipt into inventory. Included here are order preparation time, queue time, setup time, run time, move time, inspection time, and put-away time. (p. 98)

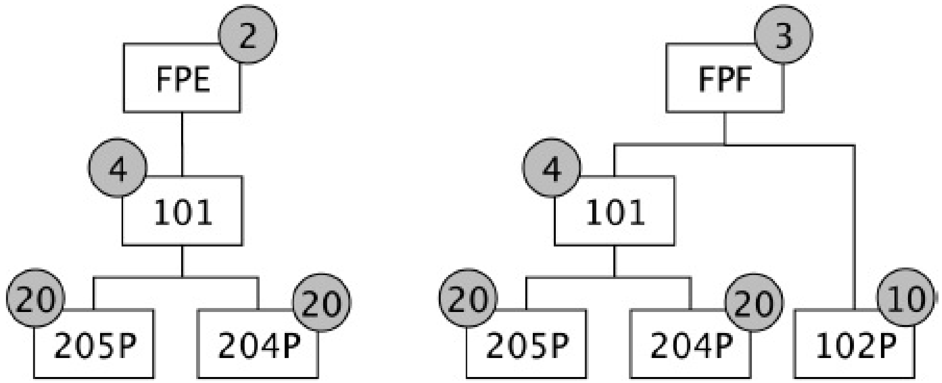

FIGURE 6-1 Product structures for FPE and FPF

> **Cumulative lead time (CLT):** The longest planned length of time to accomplish the activity in question. It is found by reviewing the lead time for each bill of material path below the item; whichever path adds up to the greatest number defines cumulative lead time. (p. 38)

> **Purchasing lead time (PLT):** The total lead time required to obtain a purchased item. Included here are order preparation and release time; supplier lead time; transportation time; and receiving, inspection, and put-away time. (p. 142)

Considering these definitions, for FPE the manufacturing lead time is 2 days, while the cumulative lead time is 26 days (20-day purchasing lead time + 4 days manufacturing lead time for 101 + 2 days manufacturing lead time for FPE). In the case of FPF, the manufacturing lead time is 3 days, while the cumulative lead time is 27 (20-day purchasing lead time + 4 days manufacturing lead time for 101 + 3 days manufacturing lead time for FPF).

To properly apply the six factors, we will need additional information about the environment. Figure 6-2 shows the product and routing structure of both FPE and FPF together. A "routing," as defined by APICS, is "information detailing the method of manufacture of a particular item. It includes the operations to be performed, their sequence, the various work centers involved, and the standards for setup and run." Together, the BOM and the routing paint a relatively complete picture of the view needed to consider positioning for this scenario. Note that no run rates and setup times have been defined, as these will not be relevant for this simple example.

Once a part 205P is introduced to the manufacturing process, it is run through a series of resources (A > B > C > D) and combined with a converted

204P at resource Z. Part 204P is run through a series of resources (B > C > E > F). Resource Z is an assembly operation and the final step in producing intermediate part 101. This conversion process (from 204P and 205P to 101), assuming concurrent activity across paths, takes four days on average. Thus 101's manufacturing lead time is four days.

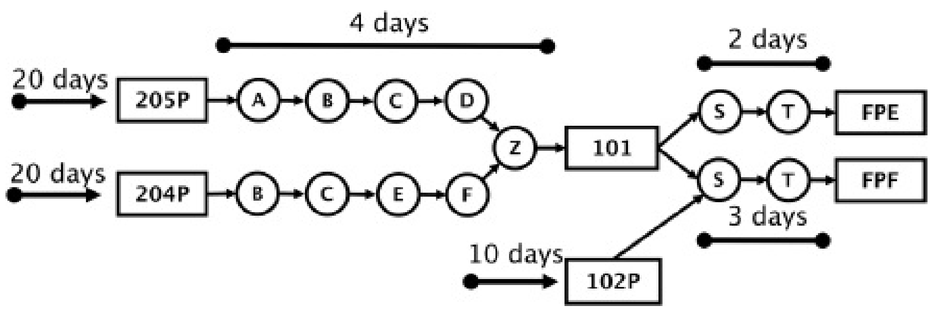

FIGURE 6-2 Product and routing structure for FPE and FPF

Resource Z is a "convergent point." A convergent point is any place where routing legs come together. As discussed in Chapter 3, these points of integration occur most often where significant delays accumulate because all parts must be present for the resource to perform its operation. Resource Z requires a converted 204P from resource F and a converted 205P from resource D at the same time and quantity. This make resource Z a candidate for a resource that we would like to protect as much as possible—a critical operation.

Part 101 is a "point of divergence." A divergent point means that part 101 can be directed into different manufacturing paths culminating in various end items. A divergent point represents a commitment that cannot be practically or cost-effectively reversed. An example would be the introduction of a sheet of steel into a fabrication process. Once the sheet is cut, the options available to use it are narrowed significantly. Thus the decision to cut it precludes it from being used in many other ways.

For this example, part 101 is directed to resources S and T to either begin the process to convert it to FPE or be combined with the purchased part 102P to be finished into an FPF. The conversion into FPE takes two days, and the conversion to FPF, a more complicated build, takes three days. Thus the manufacturing lead time is two days for FPE and three days for FPF.

When checking with sales and customer service, we find that the customer tolerance time for both products is at three days. FPF has lower volumes, as it is a higher-end product, but the market expects it within the same time frame as the lower-end product FPE. Additionally, sales has indicated that there are frequent opportunities in the market for FPE to win quick-turn business. Customers are not inclined to pay more for the items, but the volume would definitely increase with the capability to offer same-day fulfillment. Finally, with the exception of quickturn requests, this company typically receives sales orders at least three days in advance for both products. Occasionally there can be large orders, but those larger orders tend to have at least two weeks of sales order visibility.

When checking with purchasing, we discover that the suppliers for 204P and 205P have decent reliability. Occasional disruptions do happen, but overall both have performed well over the last year. The supplier for 102P, however, is a different story. This supplier is notorious for late deliveries and even routinely produces suspect quality. Figure 6-3 summarizes the positioning criteria information for this example.

Decoupling Point Placement Criteria	Example Answers
Customer tolerance time	3 days for both FPE and FPF
Market potential lead time	FPE has quick-turn (1-day) market available
Sales order visibility horizon	3+ days for most orders
External variability	**Demand:** Large orders are typically known well in advance **Supply:** 204P and 205P have decent reliability. 102P supplier is notorious for poor delivery and quality performance
Inventory leverage and flexibility	101 is a common component for both FPE and FPF
Critical operation protection	Resource Z is an assembly operation that requires both routing paths to be complete before it begin its operation

FIGURE 6-3 Example decoupling point positioning answers

Based on these answers, how should decoupling point positioning be approached in this environment? The impact of each of the criteria on the model is considered:

- **Customer tolerance time.** Three days makes it a requirement to consider decoupling at the end item or 101 and 102P levels. To do anything less will require making product to some sort of anticipated signal or forecast and incur the negatives associated with that.
- **Market potential lead time.** The opportunity for FPE suggests a

benefit for decoupling and stocking at FPE. The additional volume or customers could provide profitable revenue growth.

- **Sale order visibility horizon.** Decoupling at the finished goods or 101 and 102P levels would allow the environment to pace to actual sales orders. This is the most relevant demand signal assuring the alignment of our resources to actual requirements.
- **External variability.** Demand variability does not seem to be a huge issue—large orders are typically known in advance. Supply variability is an issue for 102P. Stocking at 102P would seem prudent.
- **Inventory leverage and flexibility.** Decoupling and stocking at 101 would allow the common component to flow to the end items as required.
- **Critical operation protection.** While the suppliers for 204P and 205P are reliable, decoupling those positions would provide as much protection to resource Z as possible from a product structure perspective.

In consideration of these answers to the positioning criteria, Figure 6-4 shows a model for this environment.

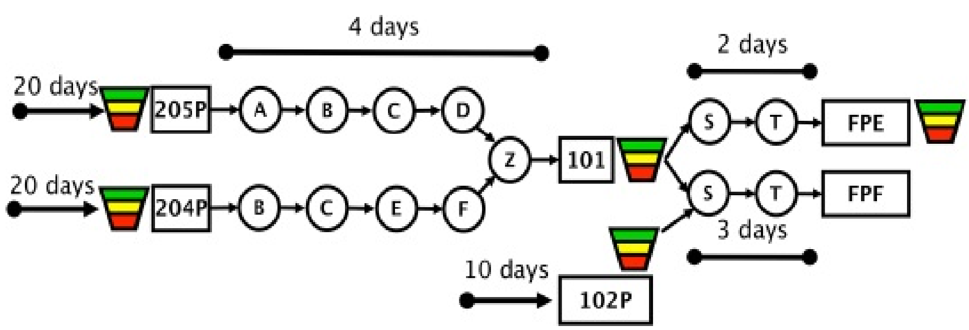

FIGURE 6-4 Decoupling positions based on positioning factor answers

The key elements and benefits of this model include:

- The FPE stock position allows for quick-turn business to be satisfied. This allows for an increase in sales revenue.
- The FPE stock position is minimized due to the short lead time from the decoupling point at 101.

- FPF can move to an assemble-to-order strategy as the lead time (three days) and customer tolerance times (three days) are compatible. Achieving this lead time reliably should be possible for three reasons. First, 101 and 102P are available as needed, decoupling lead time from the front part of the manufacturing process and supplier, respectively. Second, demand variability is not an issue with this product, as large orders are typically known in advance. And third, the buffer at FPE minimizes short-range capacity contention in resources S and T that could affect the ability to consistently achieve the three-day lead time for FPF.
- The decoupling points at 204P and 205P allow supplier variability to be isolated from the concurrent manufacturing processes in front of resource Z, thus minimizing as much as possible from a product structure perspective the variability experienced at resource Z as an assembly operation. More can be done to protect resource Z, but those options are outside the scope of decoupling point considerations. For example, a time buffer can be used in advance of resource Z in order to allow for components to be synchronized effectively. This, however, is at the scheduling and execution level of the Demand Driven Operating Model.

Additionally, we can use this example to illustrate the disadvantages of overflattening bills of material. Figure 6-5 assumes that the intermediate component is removed from the product structure of both FPE and FPF.

Figure 6-5 shows the impact on decoupling point positioning. By removing 101 from the product structure, we expand the manufacturing lead times for both FPE and FPF from two and three days to six and seven days, respectively. Since those lead times are well beyond customer tolerance time, it forces both finished positions to be stocked. Those stock requirements will be relative to that longer lead time as well. Additionally, we lose the inventory leverage of the common item, resulting in higher finished goods inventory levels. Finally, losing the capability to decouple at 101 means that there is a longer sequence of activity to convert raw materials into finished items. The variability is illustrated by the wavy line growing in amplitude in Figure 6-6. This increasing variability may create the need to hold additional stock at the end item level.

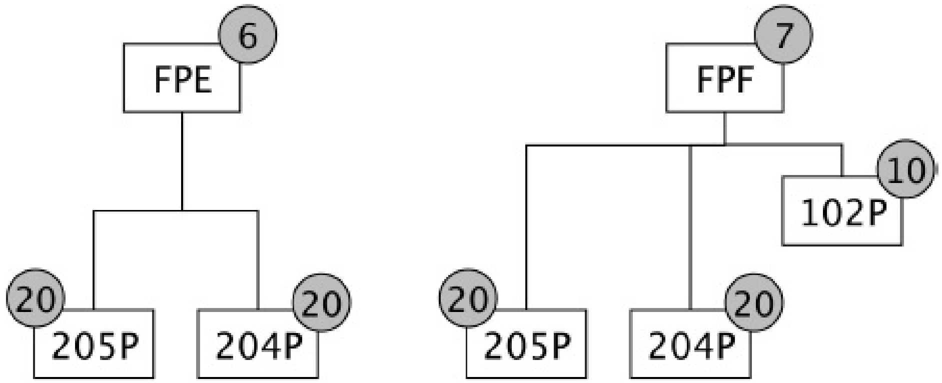

FIGURE 6-5 Flattened FPE and FPF bills of material

This example also serves to highlight the systemic effect of using multiple decoupling points together. When there are multiple tiers of decoupling points, there are bidirectional benefits to each of them. Figure 6-7 depicts the relationship between tiers of decoupling points. The decoupling points essentially protect each other. For example, critical subcomponents are protected from demand variability by the end item decoupling point, while subcomponent decoupling points protect the end item decoupling point from long lead times and large accumulations of supply variability.

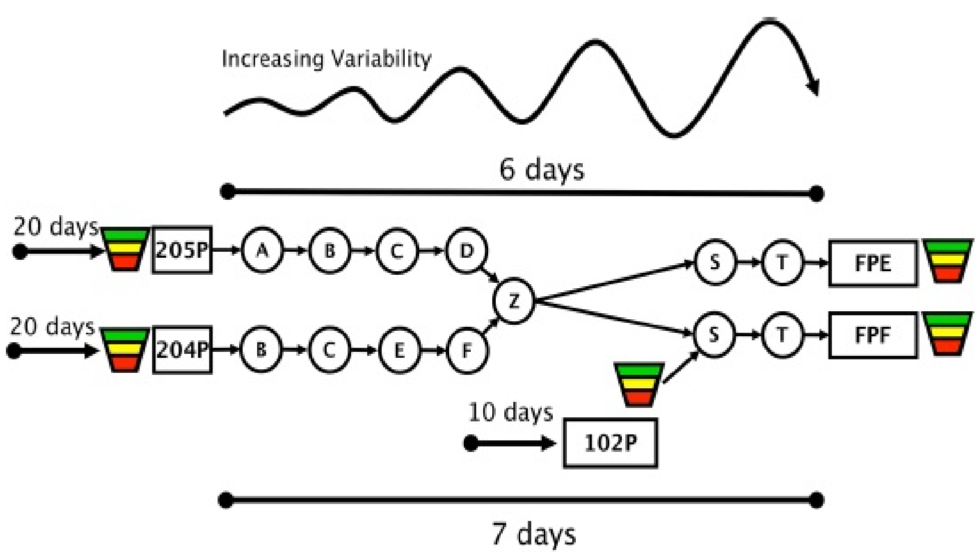

FIGURE 6-6 Example with flattened bill of material

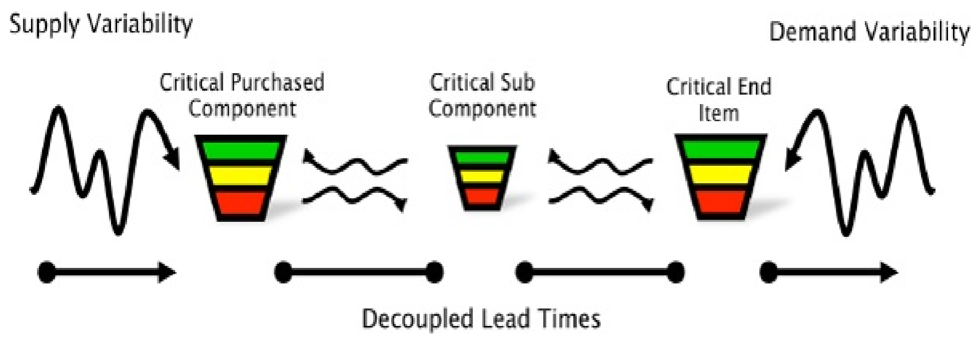

FIGURE 6-7 Illustrating the benefits of tiers of decoupling points

The previous positioning example brings us to a critical realization point that will have significant implications on how Demand Driven MRP really works. By highlighting the importance of tiers of decoupling points and leveraging inventory, this example has led us to another important impact of using decoupling points—an entirely new form of lead time.

A New Form of Lead Time

The concept and necessity of decoupling was introduced and explained in Chapter 4. Decoupling allows for demand signal distortion and supply continuity variability to be simultaneously combated. Decoupling points are the places at which we wage that battle. The use of decoupling points also results in a new type of lead time that must be understood and calculated in order to be able to:

- Compress lead times to required ranges
- Determine realistic due dates when needed
- Set decoupling point buffer levels properly
- Find high-value inventory leverage points for decoupling

The previous positioning example began by using manufacturing and cumulative lead times as factors in determining the right position for decoupling. Within the example, however, a new lead time can be conceptualized, one that can create tremendous opportunity for more complicated manufacturing entities and is a requirement for proper decoupling point buffer-level calculations. These more complex manufacturing entities have depth, breadth, and overlap with regard to their product structures that allow for this opportunity. The word "depth" implies more levels to a product structure. "Breadth" implies more legs to the product structure. And “overlap” implies shared components or materials across product structures.

In order to demonstrate this new lead time, we will take a single product structure for a product called Finished Product D (FPD). That product structure is depicted in Figure 6-8. The numbers in the circles represent the manufacturing lead times from make items and purchasing lead times for purchased items. For example, the manufacturing lead time for part 208 is 5 days. The purchasing lead time for 412P is 45 days. Note that this product structure has four levels (depth) and two major legs (breadth) beginning with components 208 and 210. There is no overlap, as we have not introduced any additional bills of material; that circumstance will be covered in depth later in this chapter.

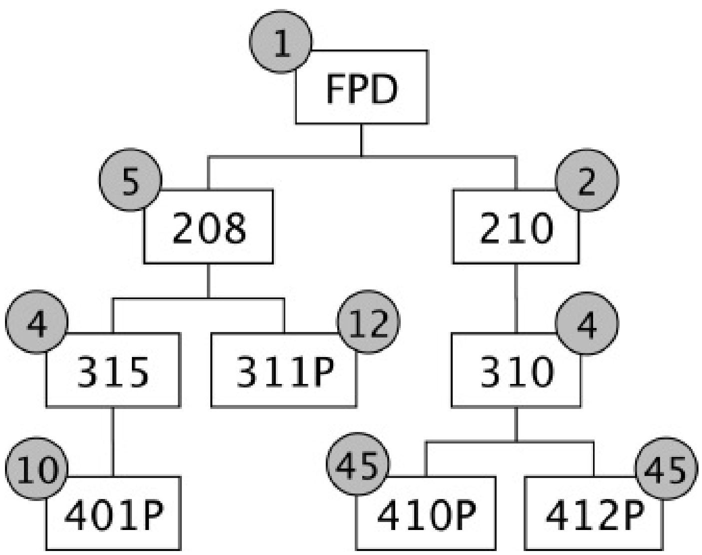

FIGURE 6-8 FPD product structure with lead times

When performing calculations for manufactured items in product structures like the one seen in Figure 6-8, conventional MRP systems only recognize two forms of lead time: manufacturing lead time and cumulative lead time. Once the concept of decoupling is embraced, these lead times become realistic in extreme situations only. In more complex manufacturing operations, these two extremes rarely exist in reality.

In Figure 6-9 we see the FPD product structure with some decoupling points inserted. The decoupling points are represented by the striped bucket icons. In this example FPD, 208, 311P, 410P, and 412P are selected as decoupled and stocked items. Why those points were chosen for decoupling is immaterial to the discussion of problems associated with manufacturing and cumulative lead times.

Let’s examine the problem associated with manufacturing lead time first given the circumstances depicted in Figure 6-9. Manufacturing lead time is the time it takes to manufacture the part exclusive of lower-level lead times. This lead time assumes that all components will be available on the parent order release date. The pathway from the decoupling points at 410P and 412P through 310 and 210 to the completion of FPD is a sequence of dependent events subject to the accumulation of variability. This variability is represented by the increasingly large wavy line moving up through the product structure in Figure 6-10. That pathway is an unprotected sequence of events. Thus the assumption that 210 will be ready precisely when needed for the release of the FPD order is suspect at best. This makes the use of manufacturing lead time an underestimation of the time required to accomplish the production of FPD.

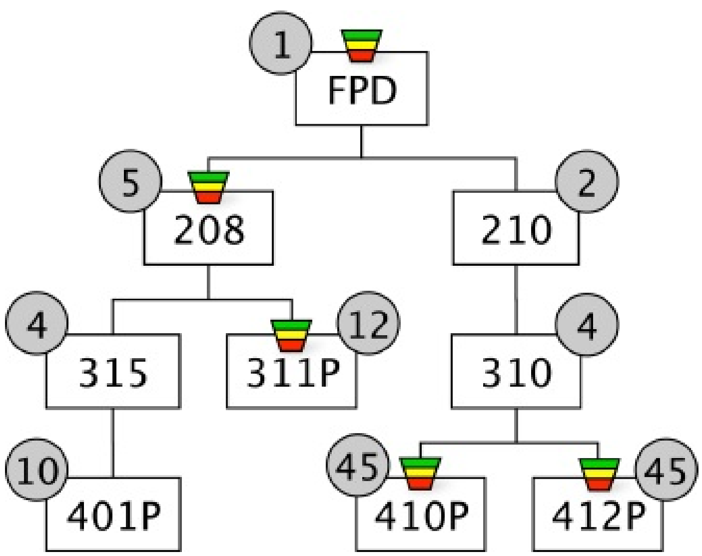

FIGURE 6-9 FPD product structure, lead times, and decoupling point positions

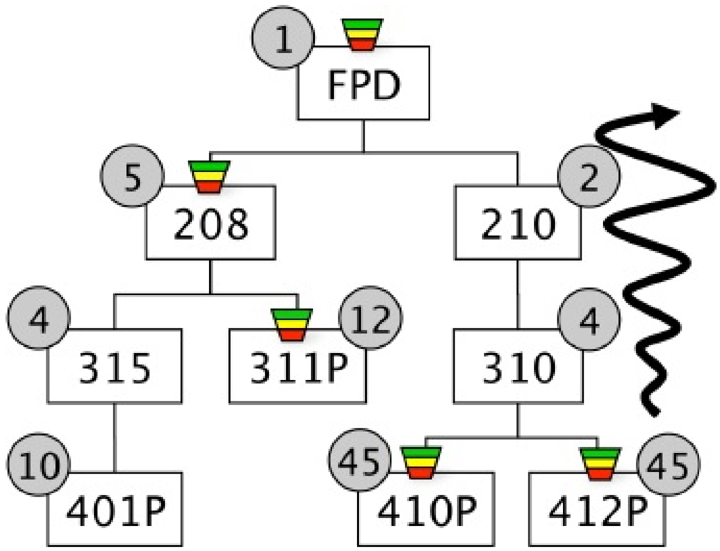

FIGURE 6-10 Variability passed in the FPD product structure

An alternative way to plan FPD in MRP would involve the use of cumulative lead time. Cumulative lead time is the longest sequence in the product structure defined in time. This assumes that no components are available upon order release. In Figure 6-11 FPD's cumulative lead time is depicted as the bolded path terminating in either 410P or 412P. The length of that path is 52 days.

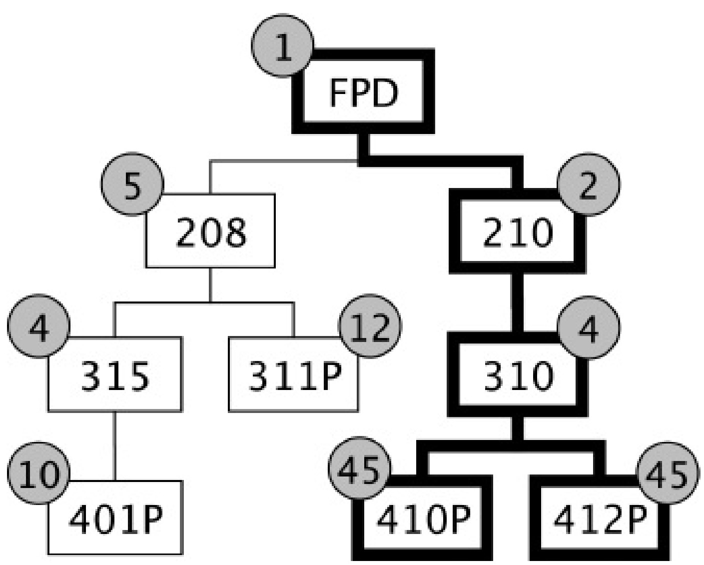

FIGURE 6-11 FPA's cumulative lead time chain

When the decoupling points are considered with FPA's product structure, the immediate problem with using cumulative lead time can be observed. In short, cumulative lead time assumes no decoupling; yet 410P and 412P (as well as 208 and 311P in the other path) are decoupled. Since these components are decoupled, it can be reasonably assumed that they are available upon parent order release. This fact makes the use of cumulative order lead time a gross overestimation of lead time when decoupling points are present.

A new form of lead time is emerging with the use of decoupling points for manufactured parts. This new form of lead time only assumes availability of the component on parent order release at decoupling points. This new form of lead time is called decoupled lead time (DLT). It can be defined as:

> The longest cumulative coupled lead time chain in a manufactured item's product structure. It is a form of cumulative lead time but is limited and defined by the placement of decoupling points within a product structure.

DLT is calculated by summing all the manufacturing and purchasing lead times in that chain. The decoupled lead time always includes the manufacturing lead time of the parent. Any parent item with at least one coupled component will always have a longer decoupled lead time than its manufacturing lead time.

The decoupled lead time path for FPD is depicted in Figure 6-12. It is the bolded large-dashed path connecting FPD and 310. The length of the decoupled lead time chain is seven days. It is calculated by adding the manufacturing lead time for 310 (four days) to the manufacturing lead time for 210 (two days) to the manufacturing lead time for FPD (one day). Figure 6-12 also depicts another important element when using decoupling points. The intermediate component 208 (which is a decoupled position) now has its own decoupled lead time chain. That path is depicted by the bolded small-dashed path connecting 208 to 401P. That decoupled lead time is 19 days.

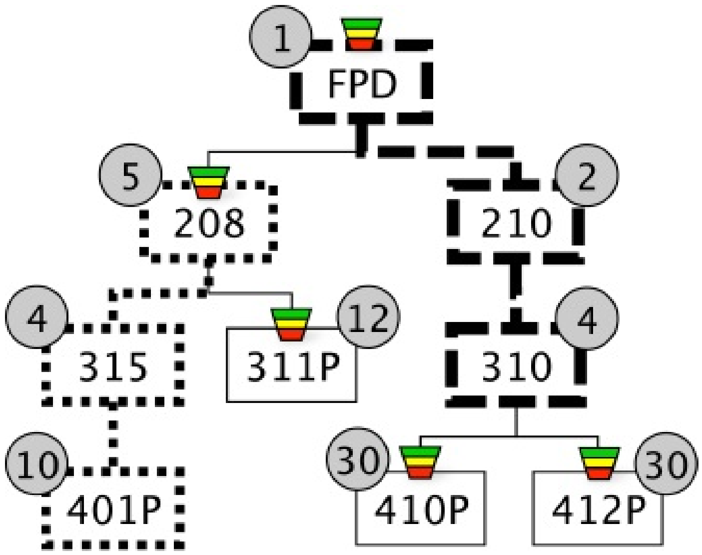

FIGURE 6-12 FPD's compressed lead time chain and 208's decoupled lead time chain

Figure 6-13 shows the decoupled lead time for intermediate component 208 being compressed from 19 days to 9 days by decoupling the purchased part 401P.

The decoupled lead time is simply a qualified cumulative lead time concept. As conventional material planning systems were not designed to use decoupling points, this form of lead time has remained hidden or at best obscured for decades. By using the DLT information, planners can now determine more realistic dates for the replenishment of a part and the inventory levels required for the decoupling point buffers. Of course, using this approach requires each discrete manufactured part number to have a manufacturing lead time or production lead time defined, and those lead times should be as accurate as possible.

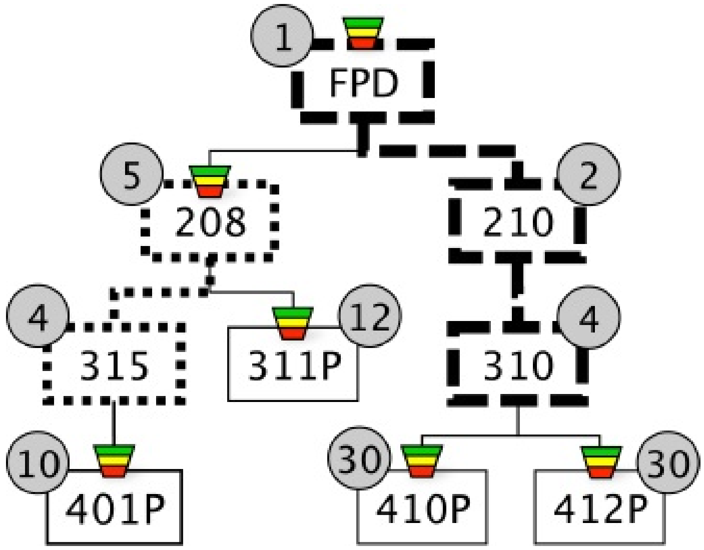

FIGURE 6-13 Decoupling 401P compresses 208's decoupled lead time

So far the value of DLT is shown in determining more realistic lead times and buffer levels with the use of decoupling points. Next let's turn our attention to the use of decoupled lead time to help find new opportunities for decoupling points as well as the elimination of non-value-added stock positions.

Advanced Inventory Positioning Considerations

Understanding decoupled lead time opens a door for more advanced inventory positioning analyses for environments in which there are deeper and broader material structures and where shared components exist across those structures. Many companies have many different products, each with its own unique bill of material. Many companies have a significant amount of shared parts or components—places where the bills of material essentially overlap. This fact combined with the concept of decoupled lead time exposes opportunities that have remained elusive for most of these companies.

Such a company is illustrated in Figure 6-14. Company ABC makes three separate end items: FPA, FPB, and FPC. These items each have a unique bill of material, and there are also shared components. Note that all parents are currently stocked as well as some of the purchased parts.

Finding the right opportunities for additional decoupling point placement, however, requires an old tool that has been around nearly as long as MRP. This concept is called a matrix bill of material. According to APICS, a matrix bill of material is:

> A chart made up from the bills of material for a number of products in the same or similar families. It is arranged in a matrix with components in columns and parents in rows (or vice versa) so that requirements for common components can be summarized conveniently. (pp. 103–104)

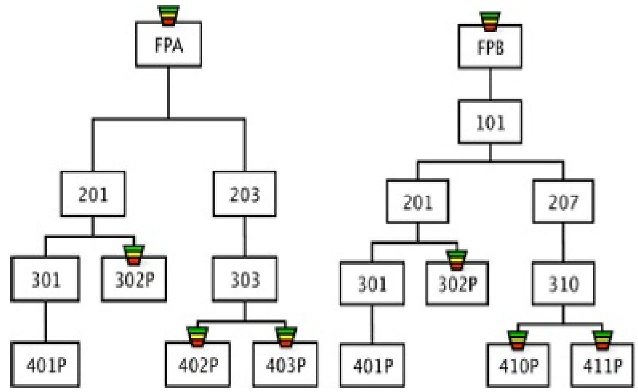

FIGURE 6-14 Company ABC environment with end item parents that share components

Figure 6-15 is the matrix bill of material for the three products illustrated in Figure 6-14. Parent items are displayed along the top (column headers), while components are displayed along the side (row headers). Note that a component will often appear as both a child and a parent. Component 201 is an example of this. It is a child to FPA and intermediate components 101 and 102 but is a parent to intermediate component 301 and purchased part 302P. The shaded parents (FPA, FPB, and FPC) and components (302P, 305P, 402P, 403P, 404P, 410P, 411P,) represent parts that are currently decoupled or stocked. Furthermore the numbers within the grid are not the quantity per parent. Those numbers only represent the number of times that a specific connection appears across all bills of material. For example, the connection between the 201 parent and 301 component occurs in all three bills of material in this environment, so there is a number 3 in that box.

Clearly, this is a very simple example. For companies that have hundreds of end items with deep BOMs and many shared components, the matrix BOM can get quite complex and very large. This fact alone meant the tool was never really in widespread use, particularly in earlier days with limited computing power.

A matrix BOM is a much broader picture than a where-used report. The where-used report is oriented to a particular component to see which parents it goes into and its usage per parent. The *APICS Dictionary* defines a where-used list as:

	Parent items												
Component items	FPA	FPB	FPC	101	102	201	203	207	202	301	303	310	306
101		1											
102			1										
201	1			1	1								
203	1												
207				1									
202					1								
301						3							
302P						3							
303							1						
310								1					
306									1				
305P									1				
401P										3			
402P											1		
403P											1		
410P												1	
411P												1	
404P													1

FIGURE 6-15 Initial matrix BOM for Company ABC

A listing of every parent item that calls for a given component, and the respective quantity required, from a bill of material file. (p. 190)

A where-used list is only a very specific slice of a matrix BOM. The matrix bill of material shows all connections between all parents and all components in an environment. What the matrix bill of material does not show is which connections really matter the most. To gain visibility of those places the matrix bill of material and the concept of decoupled lead time must be used in combination. Figure 6-16 shows the three products with the decoupled lead time chains bolded. The decoupled lead time chain for FPA is 20 days terminating at 401P. The decoupled lead time chain for FPB is 23 days terminating in 401P. The decoupled lead time chain for FPC is also 23 days terminating at 401P.

Immediately we can begin to conceptualize the value provided by this perspective. Since the end items' parents (FPA, FPB and FPC) are stocked, considering additional decoupling positions on the DLT chain will immediately compress the parent's lead times. This means that the level of each parent's buffer will be positively affected. The more the compression of lead time, the more the compression of buffer stock. This relationship will be explored in depth in Chapter 7.

If the parent items were not stocked, the value of decoupling the lead time chains can still be conceptualized. Without stocking the parent parts, FPA, FPB, and FPC would have lead times of 20, 23, and 23 days, respectively. Decoupling on their DLT chains will cause a direct compression of the lead time that can be offered to the market. Any compression to a nonstocked part's lead time must be evaluated against the customer tolerance time and market potential lead times.

The lack of value provided by decoupling component and part positions that do not lie on the decoupled lead time chain also becomes clear. For example, parts 203, 207, and 202 do not lie on the decoupled lead time chains of FPA, FPB, and FPC, respectively. Thus decoupling and buffering those positions will provide no benefit to the parent and would require an infusion of working capital to fund the buffer stock and deliver no tangible benefit.

Figure 6-17 shows the matrix bill of material for this environment with the decoupled lead time chains highlighted. Within the grid, a shaded box denotes that the connections lie on a decoupled lead time chain. For example, component 201 lies on all three decoupled lead time chains.

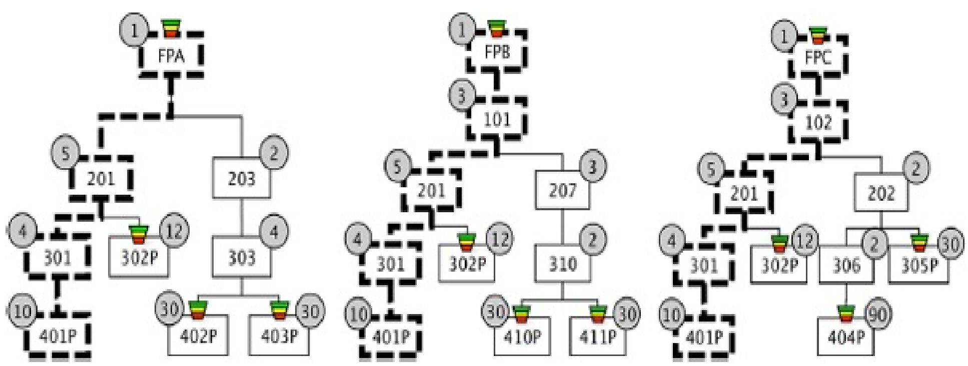

FIGURE 6-16 Decoupled lead time chains defined

Component items	Parent items												
	FPA	FPB	FPC	101	102	201	203	207	202	301	303	310	306
101		1											
102			1										
201	1			1	1								
203	1												
207				1									
202					1								
301						3							
302P						3							
303							1						
310								1					
306									1				
305P									1				
401P										3			
402P											1		
403P											1		
410P												1	
411P												1	
404P													1

Note: Three places to attack: 201, 301 or 401P

FIGURE 6-17 Initial matrix BOM with DLT chains highlighted at Company ABC

Two questions arise when looking at this view. First, where should we consider decoupling on the DLT chains? Second, how can we effectively evaluate the potential financial impact of decoupling components on the DLT chains?

In answering the first question, it is possible to narrow it down to three components. The matrix bill of material shows that components 201, 301, and 401P are on the decoupled lead time chains of all parents. Thus decoupling at any of these points will have at least some value to all parents. But which will provide the most value?

Another relevant factor requires an understanding of one of the critical positioning factors defined previously in this chapter—customer tolerance time. What we do not know yet about this example is the customer tolerance times for each end item. When sales or customer service is asked what the market expects in turnaround time, the answer is three days for all items. That means that when an item is ordered, it must be shipped to the customer within three days. This brings a new perspective to the analysis. We must devise a positioning strategy that allows us to meet the customer tolerance times for each product while effectively leveraging and minimizing inventory investment.

When examining the matrix bill of material, an immediate question must be asked: Is it possible to compress any parent's lead time to the point that it does not have to be stocked? If the answer is yes, then we would be able to potentially move those parents to an assemble-to-order situation, allowing the company to stock common components and let them flow to required parents as needed. This could potentially provide the best leverage for those common components while eliminating the need for finished stock for some end items.

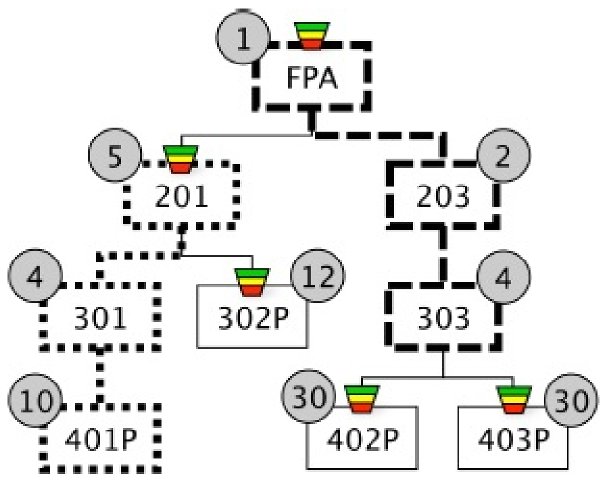

FIGURE 6-18 FPA's shifting decoupled lead time chain

One end item does provide the potential to build a positioning strategy in which it its finished goods stock can be eliminated. When examining FPA's product structure, we can see that buffering 201 will decouple a significant amount of lead time (19 days). When 201 is decoupled, however, the decoupled lead time chain shifts to a path terminating in 303. This is illustrated in Figure 6-18.

This is one important lesson about decoupled lead time chains. These lead time chains can be very dynamic based on decoupling position decisions. For environments with broader and deeper product structures, the shift will often cause an entirely different leg of the product structure to lie on the decoupled lead time chain. This is the scenario with FPA when decoupling 201. This shift is now visible in the new matrix bill of material seen in Figure 6-19.

Decoupling 201 also has implications for FPB and FPC. There is an immediate compression of lead time and shift to their respective decoupled lead time chains. Figure 6-20 depicts the shifts across all parent items accomplished by buffering 201. FPB's decoupled lead time is reduced to 9 days (from 23 days), and FPC's decoupled lead time is reduced to 8 days (from 23 days).

Decoupling 201 does not yet achieve the objective of reducing FPA's decoupled lead time within the customer tolerance time. Another step will be required. Now that the FPA's decoupled lead time chain has shifted, we will need to consider another decoupling point at 203. Decoupling 203 would allow the lead time of FPA to drop to one day—well within customer tolerance time. Figure 6-21 depicts the new decoupled lead time chain for FPA with both 201 and 203 decoupled. Note that 203 now has its own decoupled lead time chain, as it has at least one coupled component.

Figure 6-22 is an updated matrix bill of material showing the impact of decoupling 203. The shading of FPA has been removed, indicating it is now a non-decoupled part. There are only two remaining decoupling points available that impact all three product structures: components 301 and 401P. In reality the impact of the two components is limited to only one product structure's decoupled lead time—the product structure for 201. But since 201 is involved on all three end item product structures, the matrix bill shows 301 and 401P as impacting all three on a decoupled lead time chain.

		Parent items												
		FPA	FPB	FPC	101	102	201	203	207	202	301	303	310	306
Component items	101		1											
	102			1										
	201	1			1	1								
	203	1												
	207				1									
	202					1								
	301						3							
	302P						3							
	303							1						
	310								1					
	306									1				
	305P									1				
	401P										3			
	402P											1		
	403P											1		
	410P												1	
	411P												1	
	404P													1

FIGURE 6-19 The updated matrix bill of material with 201 decoupled

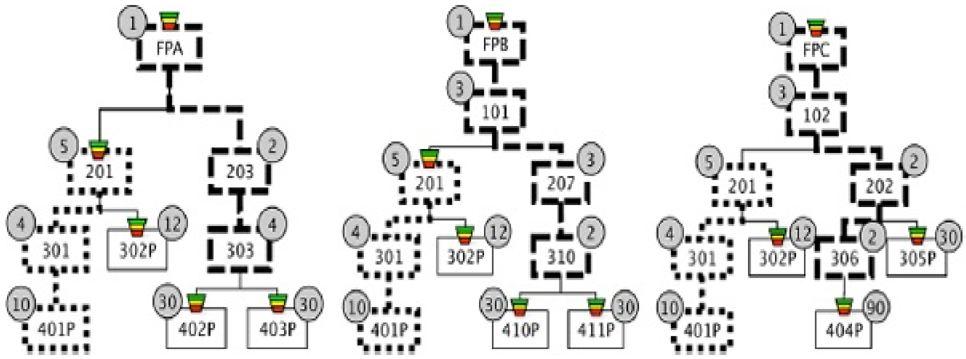

FIGURE 6-20 FPA, FPB, and FPC new DLT with 201 buffer in place

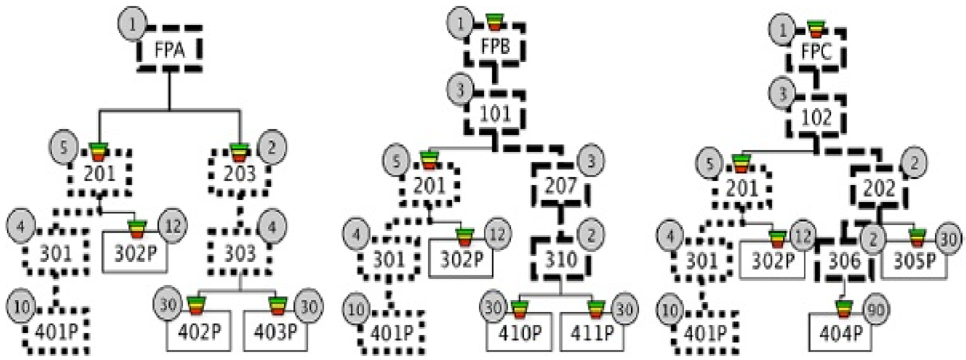

FIGURE 6-21 FPA with a one-day lead time

		Parent items												
		FPA	FPB	FPC	101	102	201	203	207	202	301	303	310	306
Component items	101		1											
	102			1										
	201	1			1	1								
	203	1												
	207				1									
	202					1								
	301						3							
	302P						3							
	303							1						
	310								1					
	306									1				
	305P									1				
	401P										3			
	402P											1		
	403P											1		
	410P												1	
	411P												1	
	404P													1

FIGURE 6-22 Updated matrix bill of material with 203 decoupled

Component 401P would seem to provide the best candidate since it will decouple external variability from the environment. Figure 6-23 shows the impact to 201 of buffering 401P. Its decoupled lead time drops from 19 days to 9 days.

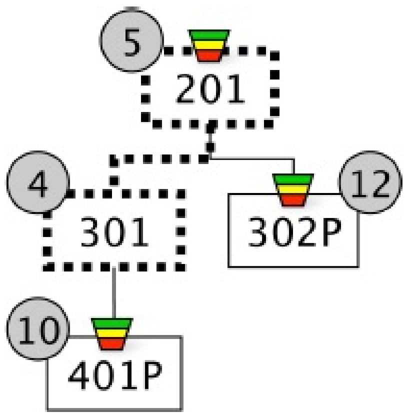

FIGURE 6-23 201's new decoupled lead time

	Parent items												
Component items	FPA	FPB	FPC	101	102	201	203	207	202	301	303	310	306
101		1											
102			1										
201	1			1	1								
203	1												
207				1									
202					1								
301						3							
302P						3							
303							1						
310								1					
306									1				
305P									1				
401P										3			
402P											1		
403P											1		
410P												1	
411P												1	
404P													1

FIGURE 6-24 Matrix bill of material with 401P now buffered

Figure 6-24 shows the matrix bill of material after 401P has been buffered. It is evident that 301 might still provide an opportunity for additional compression to 201. That compression is minimal and would still require a stocking commitment. Furthermore, we can see that the remaining stocked parents (FPB and FPC) have nothing in common to compress their respective decoupled lead times.

Figure 6-25 depicts all the buffered positions after this positioning exercise. Now there is the opportunity to eliminate a finished stock position (FPA) and dramatically compress the decoupled lead times for FPB and FPC as well as the common component 201. What is not answered yet is the financial impact to the business of making these moves. This will require an understanding of how the decoupling point buffers are sized and the amount of average working capital that will be contained in them. This will be covered in Chapter 7, where this example will continue.

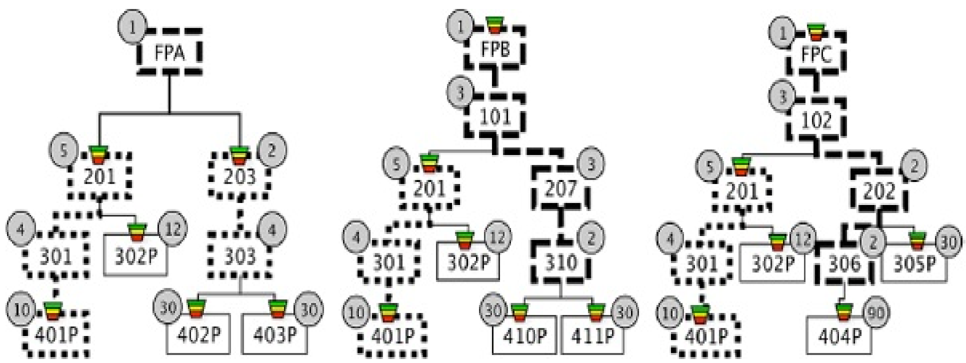

FIGURE 6-25 The final buffer positions for Company ABC

Distribution Positioning Considerations

One final aspect of positioning must be explored. So far we have focused on decoupling point positions within a manufacturing environment where raw materials are converted into sellable products. In distribution there is no conversion. Distribution is about aligning finished product to best meet consumption. In most large companies with both manufacturing and distribution aspects, there is constant tension among planning, manufacturing, distribution, sales, and logistics.

Most distribution networks have regional or local warehouses holding stock. These locations are constantly attempting to balance between the critical requirement to have what the market requires within the time frame it requires it (usually instantly) and the need to turn or convert inventory into cash or profit (not have too much inventory). To understand decoupling point positioning in the context of distribution, an appreciation for what must be decoupled is necessary.

Let's first consider demand variability. Making the right decisions (whether predictions or actions) is inherently easier when the factors being considered are more stable or known. Usually the biggest form of instability for distribution networks is demand variability. Forward locations in a distribution network can often be seen as the front line in the battle against demand variability. If actions are not taken to mitigate this form of variability, shortages, expedites, and inventory imbalance will occur (distortions to the flow of relevant materials), and the visibility to relevant information will be obscured.

Figure 6-26 depicts a simple but typical distribution environment. A sourcing unit or manufacturing plant is feeding a network of four regional warehouses. Each regional warehouse experiences a specific level of demand variability for any specific distributed item. The demand variability is much higher at each of these discrete distribution locations than at the sourcing unit for the same time period. The law of total variance means that aggregating the demand variability from the remote locations creates a natural smoothing effect at the sourcing unit.

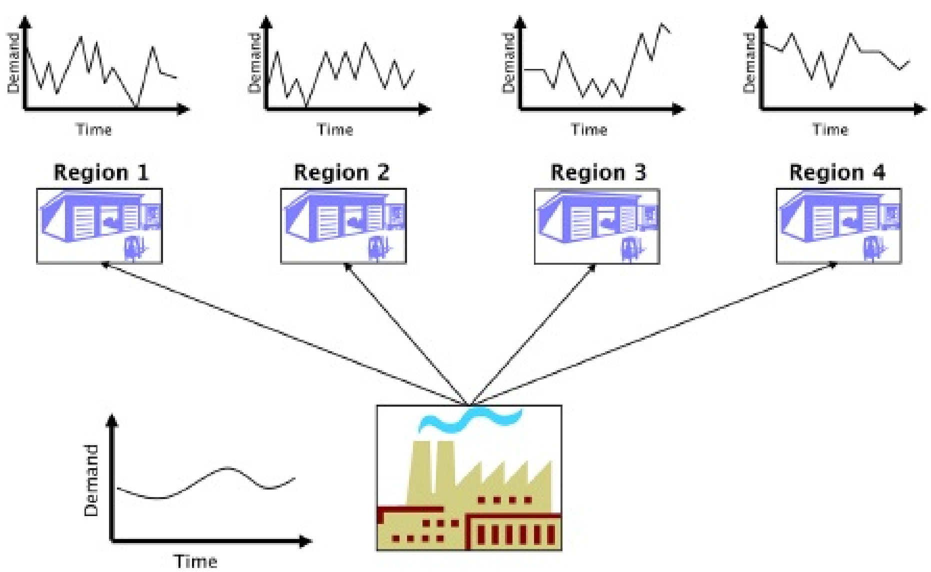

FIGURE 6-26 A typical distribution network

A simple experiment will prove this key point about demand variability smoothing as more events are aggregated. For each warehouse location a fair dice is rolled to simulate daily demand. Two die were used for warehouses 1 and

3, while a single dice was used for warehouses 2 and 4. The different number of die being used is meant to bring some additional reality to the scenario by simulating warehouses that might do higher levels of volume for a specific product. Figure 6-27 illustrates the results of the experiment. Each day the sourcing unit's demand is the summation of all warehouse demand. The column labeled "ADU" is the average daily demand over the 14-day time frame at each location.

Now when the results for each location are charted, a picture similar to Figure 6-27 results. Figure 6-28 charts the daily demand for all locations. The smoothing effect can be easily observed.

This smoothing effect is not just observable through simple line graphs. It can be calculated mathematically by using the coefficient of variance formula, which calculates the normalized measure of dispersion of a distribution. The equation for the coefficient of variation is the standard deviation divided by the mean. It is also known as relative standard deviation. Figure 6-29 depicts the coefficient of variation for each location in the experiment. The column labeled "ADU" serves as the mean for each location. The column labeled "SD" is the standard deviation at each location. Finally, the column labeled "CV" is the coefficient of variation.

This makes a compelling case that the best place in a distribution network to mitigate and manage demand variability is at a point of aggregation where there is less inherent relative volatility. Yet this mathematical fact seems to be lost on the people and organizations running the vast majority of distribution networks. Many distribution networks are designed and managed in a way that prohibits them from taking advantage of this concept.

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9	Day 10	Day 11	Day 12	Day 13	Day 14	ADU
Warehouse 1	12	8	10	8	9	8	9	7	7	5	3	10	5	2	7.4
Warehouse 2	3	6	2	1	2	1	2	3	4	3	6	5	1	1	2.9
Warehouse 3	4	7	5	5	7	10	8	8	5	4	3	9	7	8	6.4
Warehouse 4	4	3	2	1	3	2	1	5	3	4	5	5	2	3	3.1
Sourcing unit	23	24	19	15	21	21	20	23	19	16	17	29	15	14	19.7

FIGURE 6-27 Results of the distribution dice experiment

FIGURE 6-28 Daily demand charted for all locations

	ADU	SD	CV
Warehouse 1	7.3	2.68	0.365
Warehouse 2	2.9	1.72	0.604
Warehouse 3	6.4	2.02	0.315
Warehouse 4	3.1	1.33	0.434
Sourcing unit	19.7	4.04	0.205

FIGURE 6-29 The coefficient of variation at each location

Figure 6-30 depicts the structure of most distribution networks in which the vast majority of inventory is pushed out closest to the point of consumption. The regional warehouses have different levels of inventory of the five products that are distributed. Why does this push effect occur? There are several predominant assumptions or conditions behind the push including:

1. **An attempt to optimize freight costs.** Many distribution networks' primary metric is transportation spend efficiency. Attempting to minimize the transportation cost per unit often leads to large shipment quantities (full trucks) within the network containing product that is not really required at the forward location.
2. **An attempt to optimize sourcing unit costs.** The sourcing unit is typically measured on some form of unitized cost. Smaller runs and additional setups directly harm performance on those metrics. Bigger batches are the rule, and that inventory has to go somewhere since rarely is there space at the sourcing unit to store it.

3. **Location assumption.** Many in an organization, especially in sales, believe that locating the majority of inventory closest to the point of consumption has the greatest potential to meet all demand.
4. **A scarcity mindset.** When capacity or supply is perceived to be scarce, overordering will often occur by regional warehouses attempting to protect their access to product.
5. **Space limitations at the sourcing unit.** If the sourcing unit is not capable of storing inventory, it must be sent out to the network.

One obvious conclusion that Figure 6-30 reveals is that without any stock at the sourcing unit, there is no way to decouple the sourcing unit's lead time from the transportation lead time to the warehouse. That means that the regional warehouses, the places that experience the highest level of demand variability, must account for a much longer lead time as well as plant production variability when ordering. To simplify, there is a much longer and more variable lead time going to a much more variable point of demand.

This is a recipe for poor performance under any set of circumstances. What to be ordered and when to order it becomes a guessing game, as this longer lead time forces a longer planning horizon to be used. Additionally, the plant's capacity is subject to the variability of demand from the distribution centers.

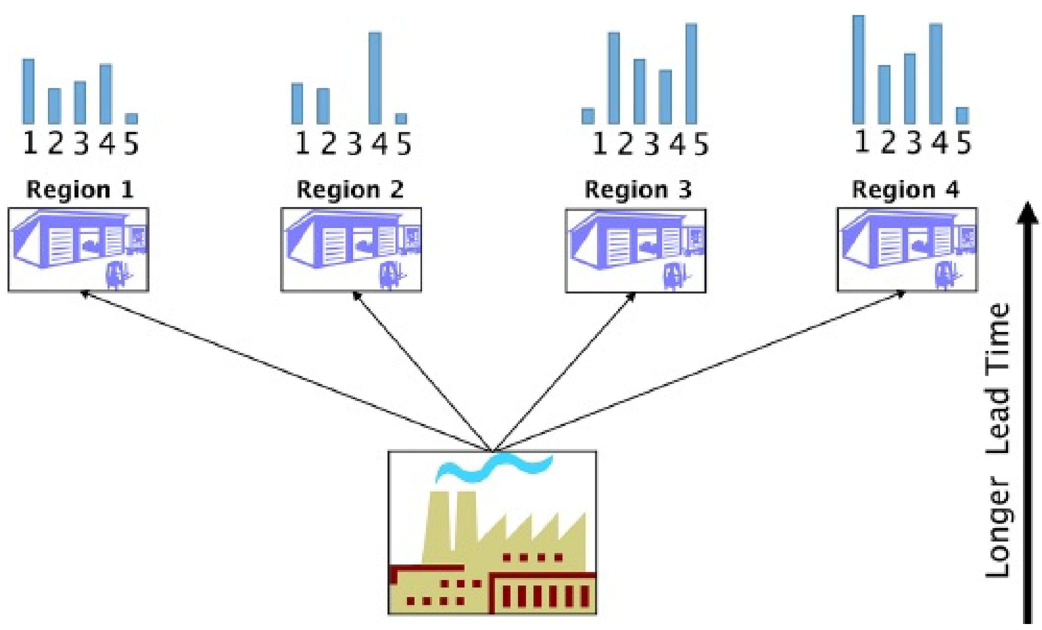

FIGURE 6-30 Inventory placement in a typical distribution situation

This situation often creates circumstances in which one region inevitably does not have enough, whereas others have too much. The result is cross-shipments between the distribution centers, missed potential sales, and expedites placed back into the plant's manufacturing schedule. As these effects occur, they compromise many of the assumptions behind the original network design:

1. Freight costs are far from optimized as cross-shipping or "rebalancing" causes additional and often expensive transportation spend.
2. Frequent break-ins to the schedule caused by emergency orders wreak havoc on the sourcing unit's schedule and operating metrics and fuel the bullwhip effect.
3. Shortages mean missed sales opportunities.
4. Shortages, long lead times, and the sourcing unit's lack of agility can reinforce the scarcity mindset, leading to further distortion to accurate demand signals.

In the aggregate, the system often has enough inventory; it is just located at the wrong place. If the inventory were better aligned in the first place, it would allow for:

- Better demand coverage for all points of consumption
- The minimization or elimination of cross-shipping
- The removal of the sourcing unit's lead time and variability from regional stock-level considerations
- Minimization of the sourcing unit schedule disruptions that reduce available capacity and complicate planning and metrics

But how to achieve this better alignment? The answer is mathematically obvious. This better alignment is created by decoupling at the point of aggregation that we saw in the previous experiment in Figure 6-29. The decoupling "hub" should be created as close to the sourcing unit as possible. Holding inventory closer to the source actually protects the largest portion of potential consumption for the least amount of inventory. A distribution network is shaped like a V. In Figure 6-30, once inventory is pushed out to Region 1, it is unavailable to the other regions in the immediate term. In the longer term, it can be made available, but only through costly realignment activity.

Realigning the inventory to locate the majority of it at this decoupling hub will accomplish the desired benefits. Figure 6-31 is a realigned distribution network showing an inventory hub located at the sourcing unit and small points of inventory at the regional warehouses—the "spokes" of the network.

The insertion of this decoupling hub or central buffer has many benefits:

1. It protects regional locations by ensuring a reliable pipeline of supply defined by transportation time only as opposed to plant lead time plus transportation time. This maximizes availability with the least amount of total inventory requirement.
2. The hub essentially eliminates cross-shipments. Why cross-ship when you can simply ship from the hub? Additionally, a strategy can be employed to ship between hub and spoke that can ensure efficient use of freight resources. This strategy, called prioritized share, will be covered in depth in Chapter 9.

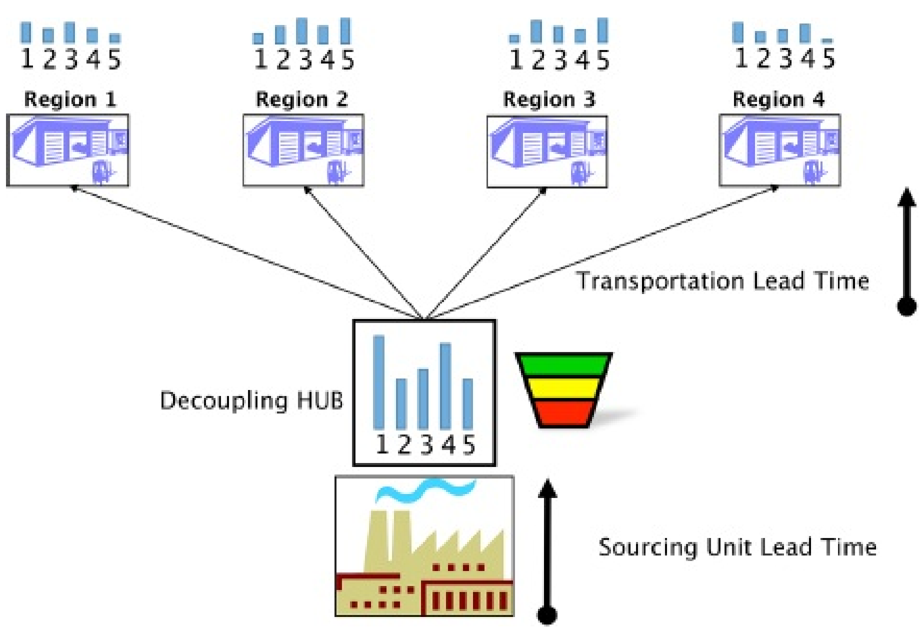

FIGURE 6-31 Decoupled distribution network

3. The planning horizon is dramatically compressed at the most variable demand point (the region), resulting in more relevant demand signals and a minimization of the bullwhip effect.
4. The hub also allows consumption and corresponding resupply signals to the plant to be naturally consolidated into batches that are still

sensible from both a plant capacity and a cost perspective.

5. The decoupling hub also allows the manufacturing facility to schedule close in time to the actual central buffer requirements rather than to arbitrary frozen schedule horizons that further limit flexibility and create frustration for sales

Yet this ideal configuration seen in Figure 6-31 may not be achievable for many distribution networks. For example, a wholesale distributor will probably not be able to force all its suppliers to hold central stocks available on demand to their regional locations. Additionally, space restrictions may exist at sourcing units that will prohibit implementation of a full decoupling hub. Different options will need to be explored to accomplish the same decoupling hub-and-spoke model given different circumstances.

First, let's explore the case of a wholesale distributor. Wholesale distributors buy and receive shipments from many different sourcing units, many of which might be remote to the distributor's facilities or each other. Figure 6-32 illustrates just such a situation. This network has four distribution points. Each point receives separate shipments from suppliers. Obviously this can present a huge challenge with regard to supplier minimum order quantities, freight cost, and space utilization.

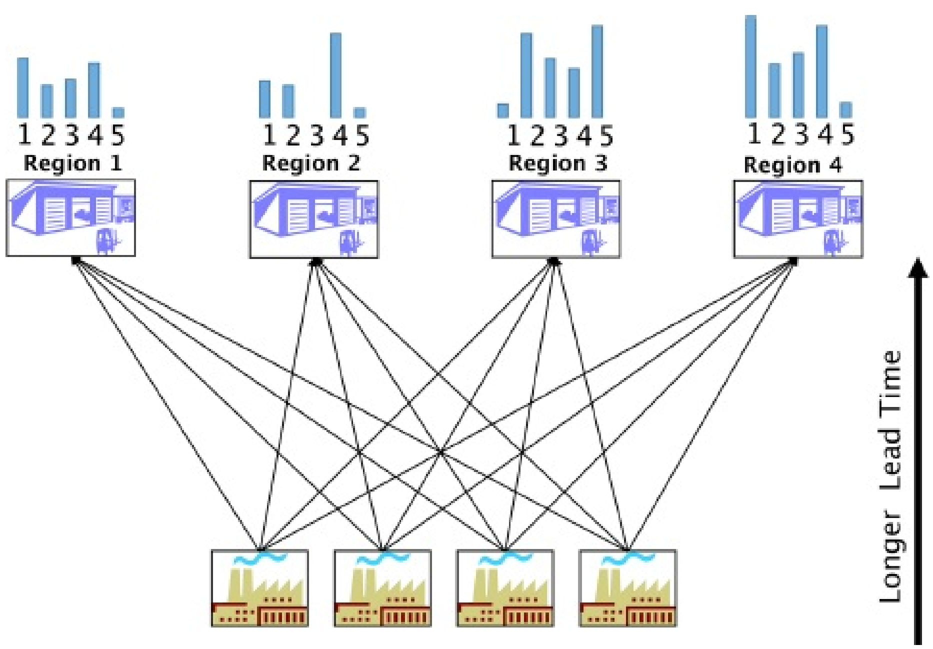

FIGURE 6-32 A typical wholesale distributor's network

As seen before, longer lead times and supply variability are going to inherently more variable points. Furthermore, the minimum order quantities imposed by suppliers can be much larger than the individual distribution point's immediate requirements. This can cause delays in ordering that lead directly to shortages, or it can cause extreme excess inventory positions as the warehouses have to take in months of supply to meet the minimum order size. Space becomes limited and further restricts ordering capabilities.

One option is that every distribution network with more than one distribution point can convert to the hub-and-spoke model in order to take advantage of its benefits. Moving to a hub-and-spoke configuration is not a trivial change in most cases. In the short term, it may require the creation of a hub through either an additional facility or the conversion of an existing warehouse. Figure 6-33 shows the wholesale distribution network seen in Figure 6-32 converted to a hub-and-spoke configuration. In this case one of the warehouses (Region 4) has been converted into the hub.

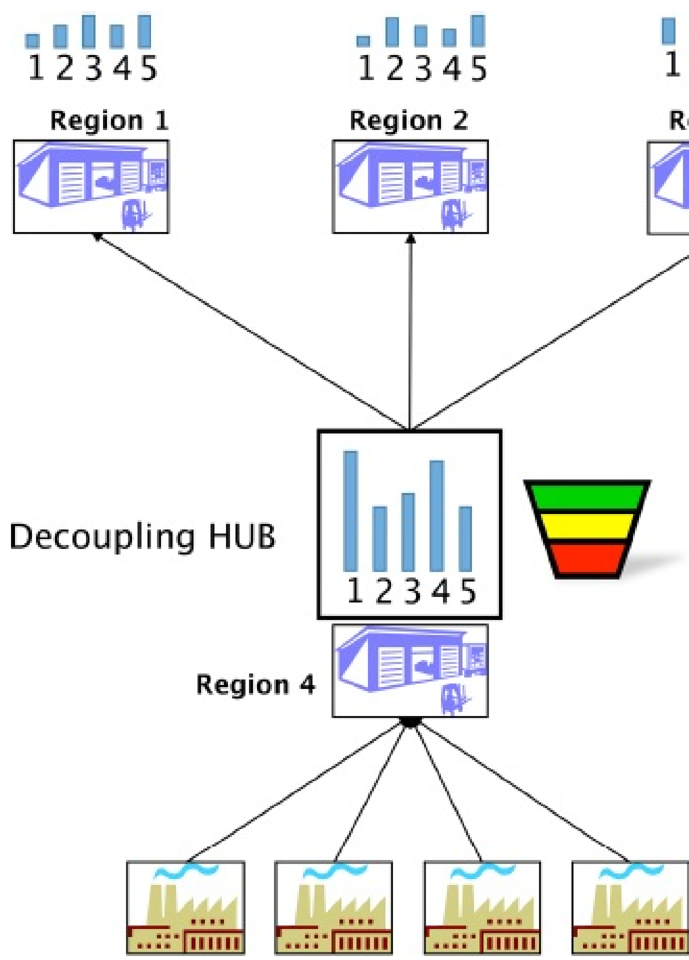

FIGURE 6-33 A wholesale distributor's network with a hub-and-spoke configuration

Some considerations in converting one region to a hub are:

1. **Demand volume.** If this position has the largest customer and volume base, then the stock in the central hub will only be moved once for this customer base.
2. **Proximity to suppliers.** If this location is in close proximity to its suppliers, the central stores might be minimized due to shorter lead times and less variability.
3. **Proximity to quality transportation lanes.** If this location lies close to major transportation lanes, then favorable freight rates (for hub-to-spoke runs) might be negotiated.
4. **Available space.** A hub will tend to hold the majority of inventory. If a location has a larger physical footprint relative to the other locations, then it might make a better hub candidate.

In addition to the previously described benefits of the hub and spoke, there are some additional benefits with regard to this example:

1. **External variability is mitigated.** The central hub decouples supplier lead times, variability, and minimum order quantity requirements from interfering with the ability of the spoke to service its unique customer base and demand variability.
2. **Space is better utilized.** As mentioned in the first benefit, the hub at Region 4 protects Regions 1, 2, and 3 from supplier lead times, variability, and minimum order quantity requirements. Now the lead time and variability are isolated to the transportation time and order frequency between the hub and spoke. This will dramatically compress the total amount of space needed at the spokes.
3. **Purchasing and in-bound freight can be better leveraged.** With one position ordering against a single large stocking point, ordering is consolidated for potential volume and freight discounts.
4. **A hub is better able to fulfill large-quantity orders.** Large orders can be fulfilled and shipped directly from the larger hub position. This position, by being larger, can better absorb these spikes without expedites, back orders, and crossships.

Most detractors of this configuration will claim that this requires moving inventory twice for sales in Regions 1, 2, and 3. But in the typical configuration in Figure 6-32, the prevailing cross-ships also represent “double handling.” Yet the point is still valid, as the model essentially calls to move this inventory twice. This must be considered against better availability and discount potential and the decreased working capital, space, and cross-ship requirements. Additionally, as noted earlier, a concept called prioritized share is explored in Chapter 9 that will help make the most of the freight spend between the hub and spokes.

Let’s expand this concept to a geographically larger distribution network. Each regional warehouse has several suppliers that are located relatively closer to it than other regional ware-houses. Figure 6-34 shows this distribution network. In this case, each regional warehouse is receiving shipments from each supplier. The suppliers are represented by the factory with a letter designation (A–L). There are 12 primary suppliers for the network, and the products produced by each are available in all regional locations. For example, the suppliers of products A, B, and C are in relative proximity to the Region 2 warehouse. The circles with “C” in them are simply meant to show their local customer base.

To apply the hub-and-spoke configuration to this network requires the “multi-hub” concept. As Figure 6-35 shows, the multi-hub configuration allows each warehouse to be both a hub and a spoke. Each warehouse will serve as a hub for the products produced in its relative geographic proximity and a spoke for the products produced in relative proximity to other regional warehouses. For example, Region 2 will be the hub for the products produced by suppliers A, B, and C, and it will serve as a spoke for the products produced by the other nine suppliers.

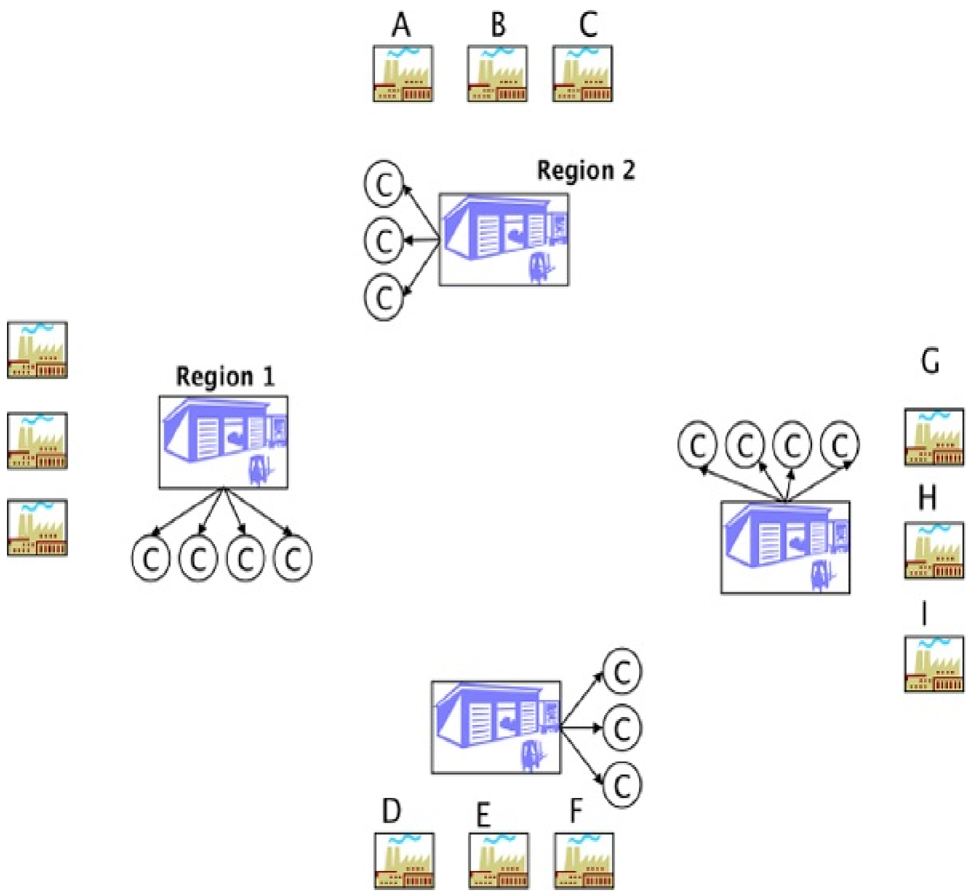

FIGURE 6-34 Supplier proximity to different regional warehouses

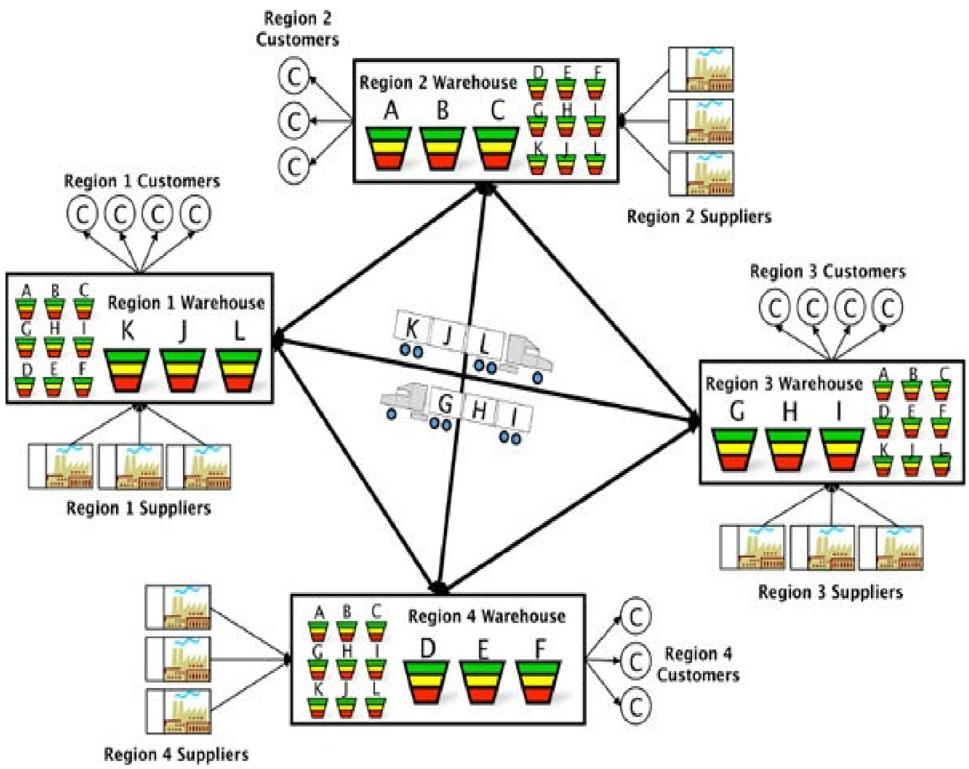

FIGURE 6-35 Multi-hub configuration

This network design brings with it all the benefits of the hub-and-spoke configuration plus one additional transportation-related benefit. The transportation lanes between each warehouse are now bidirectional. For example, trucks leave Region 3 full of G, H, and I products and return from Region 1 with K, J, and L products. This is a transportation system with a built-in backhaul situation that allows a company to negotiate more favorable freight rates or more efficiently use its own trucks.

One additional configuration should be considered. It is essentially a partial hub-and-spoke and is known as the “hybrid.” The hybrid model can be used when there is limited space at a sourcing unit, disallowing the deployment of a full hub at that location. The hybrid model focuses on decoupling the variability between the sourcing unit and the distribution network associated with slow-moving items. Slow-moving items are of particular concern since their minimum quantity requirements in relation to their usage rates often create significant imbalance in the network and scheduling difficulties for the plant.

Figure 6-36 illustrates this hybrid approach. By establishing a hub for slower-moving items, the distribution network has a steady supply of slow-moving items without the space penalties of storing those items for long periods of supply at the regional level. Additionally, the hub allows the sourcing unit to run a more consistent order cycle of faster-moving items while occasionally scheduling reasonable batches of slow-moving items when the hub actually needs resupply. The faster-moving items (products 1, 2, and 3) are sent out to the network as they are produced at the sourcing unit. Loads are supplemented with required quantities of slower-moving items (products 4 and 5) for each region. In Figure 6-36 a manufacturing run of product 1 has just been produced by the sourcing unit and is being sent out to each of the regional warehouses. Each truck additionally has quantities of products 4 and 5 for each warehouse.

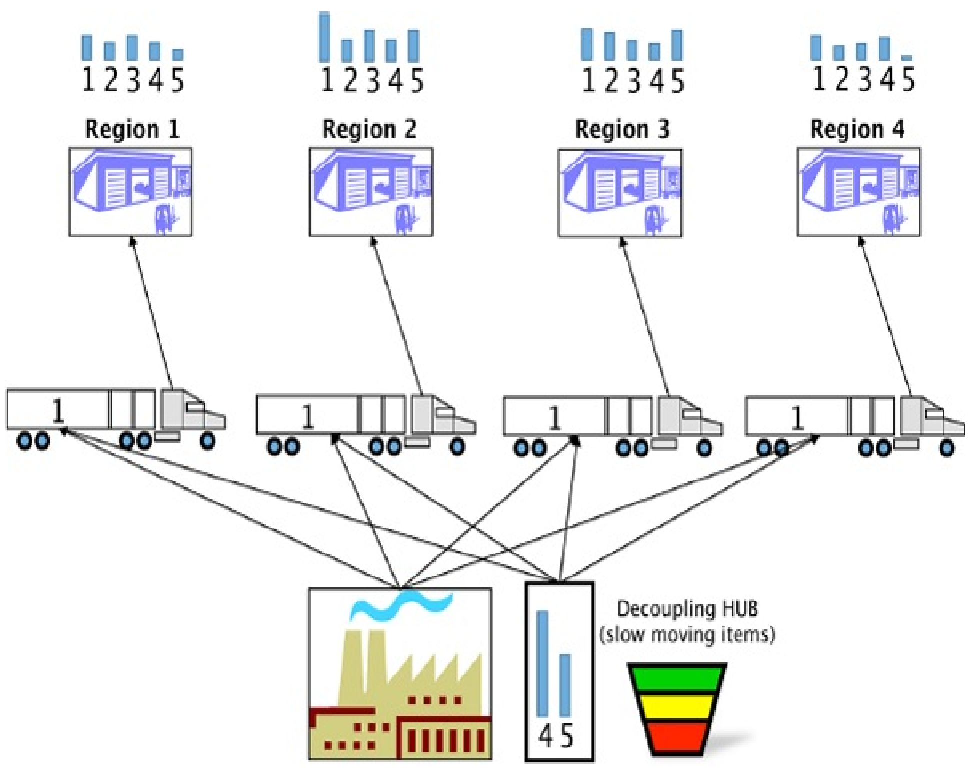

FIGURE 6-36 The hybrid configuration

The hybrid is a compromise and brings with it some additional considerations.

One question that arises is, How should slower-moving items be determined? Slow movers will be identified through the use of a "flow index." The concept of the flow index is covered in Chapter 12.

Other questions include: How should slow movers be fit into the manufacturing schedule? What if a particular fast-moving production run is insufficient to deal with all immediate network demand? What if a particular fast-moving production run is above the immediate network demand?” These questions will be answered after the actual supply order generation mechanism of DDMRP is explained in Chapter 9.

Summary

This chapter has provided an in-depth exploration of the first component of Demand Driven Material Requirements Planning—inventory positioning. Inventory positioning is about answering the most primary question related to promoting and protecting the flow of relevant information and materials. That question is, Where to decouple? As seen in this chapter, answering this primary question is a strategic process with considerations that impact the breadth of a supply chain. Thus planners and buyers cannot answer the question alone. This is a cross-functional and strategic effort that should be discussed with representatives from most aspects of the organization (operations, logistics, sales, purchasing, and finance). This process is further described in Chapter 13 in the section “Demand Driven Sales and Operations Planning.”

CHAPTER 7

Strategic Buffers

The second step of DDMRP is the mechanism that allows a decoupling point to stay decoupled—a buffer. In this case the buffer will be a level of stock that is carefully sized and maintained. This chapter focuses on the sizing considerations of strategic inventory buffers. First, a critical question about inventory must be considered.

Inventory: Asset or Liability?

In order to better understand how to determine the protection levels of decoupling positions, first a question: Is inventory an asset or a liability? There seem to be two prevailing and confusing answers.

According to the company balance sheet, inventory is an asset. For decades many large companies have played paper games with regard to inventory. Despite having no demand, many companies continued to build inventory, realize the accounting value-add from that inventory, and declare accounting profits against it. In the process, companies are drained of cash and may go deeply into debt, but according to generally accepted accounting principles the company was profitable.

Today, with the proliferation of methodologies such as Lean and the Theory of Constraints, in addition to the global economic meltdown of 2008–2010, fewer companies can afford to play these games. Cash is an important focus, and Wall Street also has become aware of the ruse and the penalties associated with too much inventory. But has the pendulum swung too far the other direction?

Edicts to slash inventories can be extremely harmful. For example, let's take the case of a director of purchasing who is given an edict to dramatically and immediately reduce inventory on purchased items. When her staff does the full analysis on the inventory, she discovers that the vast majority of the materials and packaging inventories are on slow-moving or obsolete items. With regard to purchased materials, large-quantity buys were made at significant discounts to drive a positive purchase price variance but have resulted in bloated positions. Furthermore, to exacerbate the situation, engineering made a significant material revision to many of the higher-moving products, further eroding the usage of several materials with high stock positions. At their current rate, these materials will not drain out for over one year. The situation with packaging inventories is even worse. Marketing, without much notification, made significant packaging changes in order to support a rebranding effort. Now the old packaging is essentially useless, but finance is balking at writing it off due to the impact on the profit and loss statement.

What can this director do when most of the inventory is not moving or unusable yet is under direct orders to reduce inventory immediately? There is really only one option to reduce inventory in a very short period of time. That is to cancel or defer open supply orders on the materials and packaging that are actually moving. In just a few short weeks this manager will make a very big dent in inventory value.

Will this director be praised? Stocks of critical, fast-moving items will deplete rapidly. Shortages in materials and packaging will block or delay the manufacturing schedule. Service levels will be adversely impacted. Material and packaging expedites will rise dramatically. Flow will be reduced to a trickle, and operations and sales will scream loudly.

Furthermore, metrics that focus on inventory turnover can be extremely distortive. One of the most important performance measures of the overall health of a manufacturing business is thought to be inventory turnover ratios. In the United States before 1980, these ranged from below one to as high as six. Management thought it was doing well to increase the figure by 50 percent in one year. Many who did so simply fell back to previous ratios the next year, indicating successful crisis management and actions but no permanent or sustainable improvement in performance. Many found that lowering inventories harmed customer service, which impacted revenue and caused higher costs.

With the advent of a leaner emphasis, inventory turns commonly increase—sometimes even dramatically. But is inventory turnover always an indicator of excellent or optimal inventory performance? High turns and high shortages can and often do coincide. This happens frequently when companies lean out their inventory too much and essentially make their supply chains too brittle. They are not improving, they are putting themselves at a competitive disadvantage.

Thus we can conclude that true agility is not synonymous with zero inventories. The key to effectively leveraging the working capital and capacity commitment inherent in inventory is to find the places where inventory can make the biggest positive impact and therefore provide the greatest return. Inventory can decouple otherwise dependent events so that the cumulative effects of variation are not passed or amplified between the dependencies. Thus inventory can be a break wall against the variability experienced from either supply (externally and internally) or demand variability. However, as with any break wall, it is effective only if it is placed and sized properly.

There are two prerequisites for inventory to be a true asset to a business: placement and sizing. Fundamentally, the word "asset" has a specific meaning. Businesses have expectations from the assets on their balance sheet; in particular, they expect a return on those assets. Inventory should be treated no differently. Yet how can we calculate a rate of return on something when simply having more of it distorts the financial picture or conversely slashing it and driving it too low compromises our ability to decouple variation and respond to customer demands?

Perhaps the value of these critical inventory placements can be calculated based on something that we know directly connects to return on investment—the flow of relevant information and materials. As noted in Chapters 1–3, conventional planning systems typically result in a bimodal distribution in relation to inventory levels. Many parts are overstocked, while at the same time some are understocked. Whether overstocked or understocked, there is a breakdown in the flow of relevant information and materials. This means that with regard to inventory, there is a range in which inventory (assuming it is in the right position) is truly an asset, and when outside of that range, inventory truly becomes a liability.

When a company has too much inventory (overages), we know that excess cash, capacity, materials, and space are required. Obsolescence risks are higher. Discounts to liquidate stock cause losses and potentially cannibalize other higher-margin sales. Additionally, work-in-process levels might be higher, expanding lead times beyond the customer tolerance time and hurting sales. From a flow perspective this is certainly less than optimal.

When a company has too little inventory, we know that chronic and frequent shortages prevail, resulting in scheduling delays, missed sales opportunities, costly expedites, additional freight as partial shipments are made, and overtime employed. Once again, from a flow perspective this is certainly less than optimal.

This means that inventory is an asset somewhere between these two points—a nominal point. Furthermore, it means that there is a loss function that occurs in either direction from the nominal. Figure 7-1 illustrates this loss of value as we move toward the extremes of too little or too much and outside of an optimal range.

If we can find a simple way to calculate this optimal range, then we can judge inventory performance against it. If inventory positions are frequently outside this range in either direction, then we know that there are potential improvements to be made. Additionally, over the course of time we can change the parameters of the range through improvement activities (either making the range smaller and/or shifting it to the left). In order to begin this journey, we must explore how to calculate the size of the necessary protection at a decoupling point.

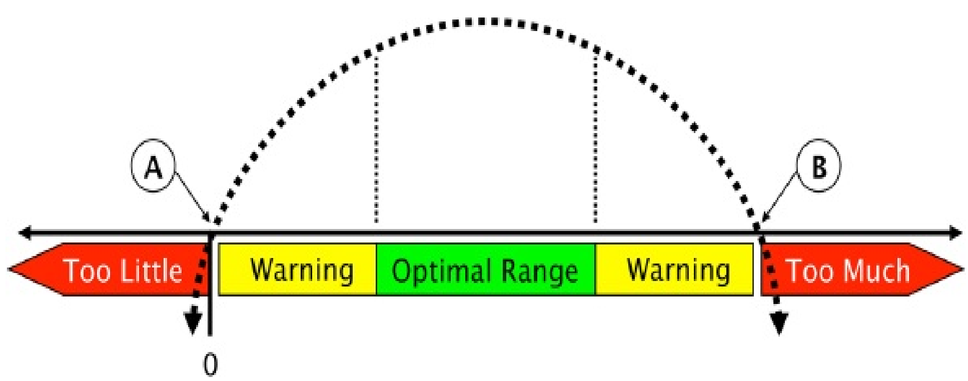

FIGURE 7-1 The inventory value loss (Taguchi) function illustrated

Introducing Decoupling Point Buffers

The protection at the decoupling point is called a buffer. Buffers are the heart of a DDMRP system and serve three primary purposes:

- **Shock absorption.** Dampening both supply and demand variability significantly reduces or eliminates the transfer of variability that creates nervousness and the bullwhip effect. This was covered in Chapters 4 and 6.
- **Lead time compression.** By decoupling supplier lead times from the consumption side of the buffer, lead times are instantly compressed. This was covered in Chapter 6.

- **Supply order generation.** All relevant demand information, supply information, and on-hand information are combined at the buffer to produce a “net flow” equation that determines supply order generation. The buffers are the heart of the planning system in DDMRP. This is covered in Chapter 9.

The word “buffer” implies something substantive enough to be able to accomplish these purposes yet not too large so as to impede flow. This means there needs to be a practical way to calculate what the level of protection should be with the nominal range and specification limits in mind.

DDMRP employs three types of stock buffering methods at decoupling points. The type of method used is based on whether a part is classified “replenished,” “replenished override,” or “min-max.” The bulk of this chapter will focus on the replenished part classification, as it is the predominant method used in a DDMRP system.

1. **Replenished parts.** Replenished parts use strategic and dynamic decoupling point buffers. These parts are managed by a dynamic three-zone color-coded buffer system for planning and execution. The buffer levels are calculated by a combination of globally managed traits relative to the buffer profile into which the part falls and a few critical individual part attributes. These factors are adjusted within defined intervals.
2. **Replenished override parts.** Replenished override parts are strategic and static decoupling point buffers. These parts are managed by a static three-zone color-coded buffer system for planning and execution (as opposed to calculated and dynamic for the replenished parts). Parts are assigned to this category when there are defined limitations (space, process related, and/or cash) or dictated levels of inventory (customer agreements, policy restrictions, etc.) within the planning environment. Without the dynamic nature of the buffer, the color-coding system becomes that much more important for planners to prioritize planning- and execution-related activity.
3. **Min-max parts (MM).** The min-max designation is for nonstrategic and readily available stocked parts and stock-keeping units (SKUs). There is still a role for traditionally defined MM tactics in DDMRP. APICS defines min-max as:

> A type of order point replenishment system where the “min” (minimum) is the order point and the “max” (maximum) is the “order up to” inventory level. The order quantity is variable and is the result of the max minus available and on-order inventory. An order is recommended when the sum of the available and on-order inventory is at or below the min. (p. 105)

Min-max buffers in DDMRP are managed by a simpler two-zone color-coded system that can be dynamically altered or adjusted in the same way as replenished parts.

All these parts’ buffer levels are determined by summing the zones that comprise them. Replenished and replenished override utilize three zones, while min-max utilizes only two zones. Zones are stratifications or layers in the buffer that serve specific purposes and have unique calculations.

The Green Zone

The green zone is the heart of the supply order generation process embedded in the buffer. It determines average order frequency and typical order size. The green zone is determined by one of three factors. Whichever factor yields the greatest number determines the size of the green zone. In this way it represents a conservative view with regard to recommended and average order frequency.

The Yellow Zone

The yellow zone is the heart of the inventory coverage in the buffer. The yellow zone is always calculated as 100 percent average daily usage (ADU) × decoupled lead time (DLT).

The Red Zone

The red zone is the embedded safety in the buffer. The higher the variability associated with the part or SKU, the larger the red zone will be. Calculating the red zone is accomplished with three sequential equations.

Figure 7-2 represents a summary of the purposes and calculations for a DDMRP buffer. Performing the calculation for each of these zones is accomplished through a combination of a grouping assignment (called buffer profiles) and individual part attributes.

Buffer Profiles

A buffer profile is a grouping of parts that have similar characteristics. Buffer profiles allow for the practical and effective global management of massive quantities of strategically decoupled parts. Obviously, many different materials, parts, and end items behave differently. Conversely, many behave very much in the same manner. Buffer profiles are families or groups of parts for which it makes sense to devise a set of rules, guidelines, and procedures that can be applied the same way to all members of a given buffer profile. Devising and revising rules, guidelines, and procedures for hundreds or thousands of parts individually would be overwhelming.

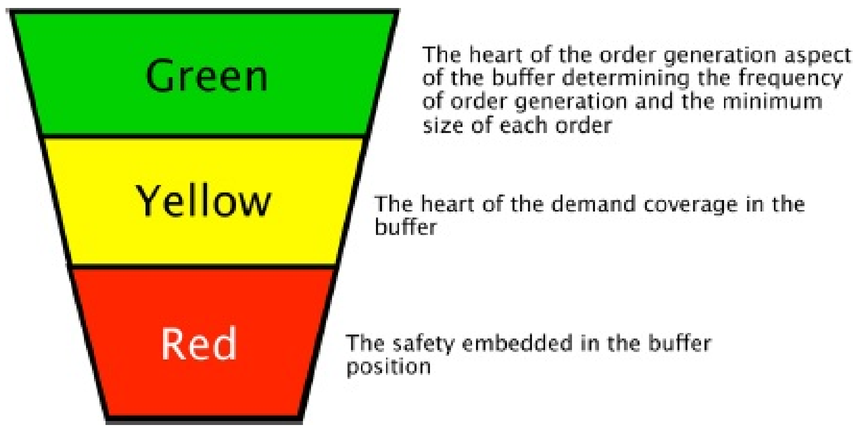

FIGURE 7-2 Buffer zones and purposes

These families should not be confused with the traditional notion of product or marketing families, which tend to be components or end items grouped by like characteristics in terms of physical configuration or markets. With buffer profiles, the familial connection is made based on three specific factors.

Factor 1: Item Type

Item type becomes the primary designator for globally managing families of parts. The groupings will be made by determining whether an item is manufactured (M), purchased (P), or distributed (D). The reasons to group by these designations are:

- **Responsibility.** Companies often designate the control of these different item types to different people or groups.
- **Intuition.** Knowledge about a specific part is frequently limited to the specific group that controls those parts.

- **Organizational control.** There is often a varying degree of direct organizational control over these different item types. Companies will tend to have more direct control over something contained within their facilities. The amount of control that extends to purchased and distributed items often depends on the vertical integration of the enterprise.
- **Categorical differences.** Relative lead time horizons can be very different among these item types. Short lead times for purchased items could be up to a week. Short lead times for manufactured items could be one to two days.

The three item types—manufactured, purchased, and distributed—are typically the minimum number of item-type designations a large supply chain entity should have. There can be others if applying the above criteria leads us to the creation of more. For example, in some environments a distinction can be made between end item manufactured items and intermediate manufactured items with regard to the above criteria. This may call for a different classification called intermediate (I). There is an example later in the chapter that will have an intermediate category.

Factor 2: Lead Time

Lead time is segmented into at least three categories: short, medium, and long. These designations are relative to the company's specific environment and part type. Typically, there is a large distribution spread in the size of lead times associated with purchased parts. This spread could be anywhere from almost zero lead time for on-site supplier-managed inventory to lead times measured in months or years.

Purchased parts that are reliably received with very short lead times are typically not candidates for strategic replenishment designation. Little benefit can be gained from the additional management of these parts. Figure 7-3 details the distribution of lead times for purchased parts identified for strategic replenishment in a sample environment (Company XYZ). Out of 100 parts, 37 are in the short lead time group, 30 in the medium lead time group, and 33 in the long lead time group.

There are differing circumstances that dictate what the parameters defining short, medium, and long within any particular environment will be. Typically, the division point will come down to a comfort level for the buyers in that environment. Later in this chapter the lead time designation influence on buffer levels and zones is discussed.

Manufactured parts have three types of calculated lead times that can be evaluated in order to determine what is short, medium, and long. As discussed in Chapter 6, two of these lead times, manufacturing lead time (MLT) and cumulative lead time (CLT), are problematic when decoupling points are in use; they are either an underestimation or an overestimation, respectively. To this extent, decoupled lead time should be used to determine what is short, medium, and long. Figure 7-4 represents an example of short, medium, and long designations against the distribution of manufacturing parts chosen for replenishment. Note that the lead time definitions of short, medium, and long differ significantly from the purchased part definitions in Figure 7-3.

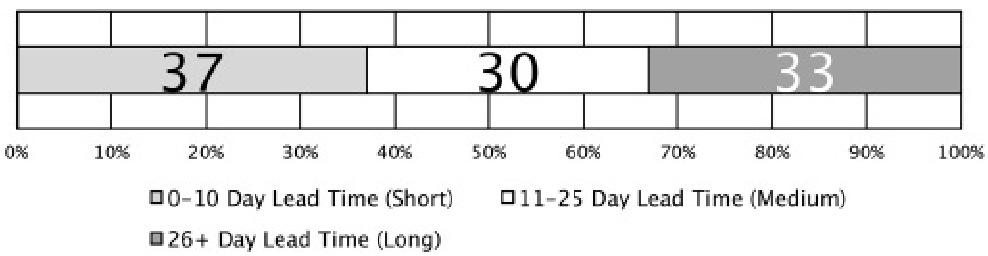

FIGURE 7-3 The distribution of lead time category assignment to purchased parts for Company XYZ

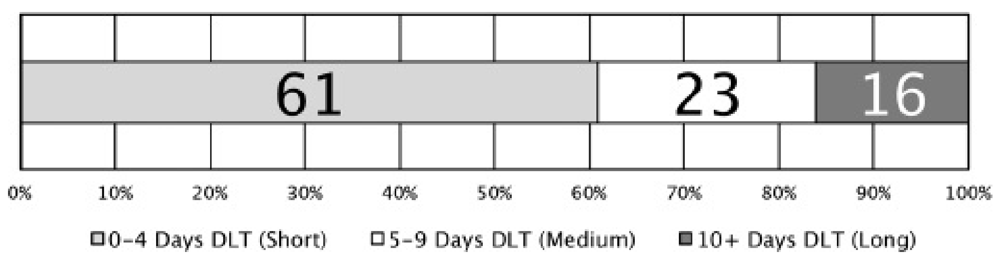

FIGURE 7-4 The distribution of lead time category assignment to manufactured parts for Company XYZ

[The lead time category will then be used to supply a "lead time factor" to parts within a profile. Figure 7-5 is a table of recommended lead time factor ranges assigned to the different lead time categories. The lead time factor is a percentage of ADU within the decoupled lead time of the part. This lead time factor will impact green and red zone calculations for every strategic part within a certain profile.

Notice that *the longer the lead time of the part, the smaller the lead time factor* should be. A smaller lead time factor produces a smaller green zone calculation. Since the green zone determines average order size and frequency, a smaller lead time factor will lead to smaller and more frequent orders. This may seem counterintuitive for many planners and buyers, but the DDMRP approach forces as frequent ordering as possible for long lead time parts (until the minimum order quantity or an imposed order cycle becomes a constraining factor).

This is in direct opposition to the way that many purchased long lead time parts are often handled. Typically, long lead time parts also represent persistent problems and shortages. It is not uncommon that buyers will buy double and triple the order minimum representing months or more of supply just to not have to deal with these parts as frequently. This is especially true if there has been a shortage in the recent past. There is an old saying in purchasing, “Buy double and stay out of trouble.” Of course, much of this behavior is directly associated with the deficiencies of conventional MRP previously described. Once these deficiencies have been addressed, then the behaviors to compensate for them can be reexamined.

DDMRP is about creating and protecting the flow of information *and* materials. For long lead time parts, DDMRP is attempting to create a frequent demand signal relating to actual need and a corresponding supplying “pipeline” delivering a steady stream of supply orders. Figure 7-6 depicts the difference between large infrequent orders and a steadier stream of smaller more frequent orders.

Long Lead Time	20 to 40% Average Daily Usage (ADU) x Decoupled Lead Time (DLT)
Medium Lead Time	41 to 60% ADU x DLT
Short Lead Time	61 to 100% ADU x DLT

FIGURE 7-5 Recommended lead time factor ranges

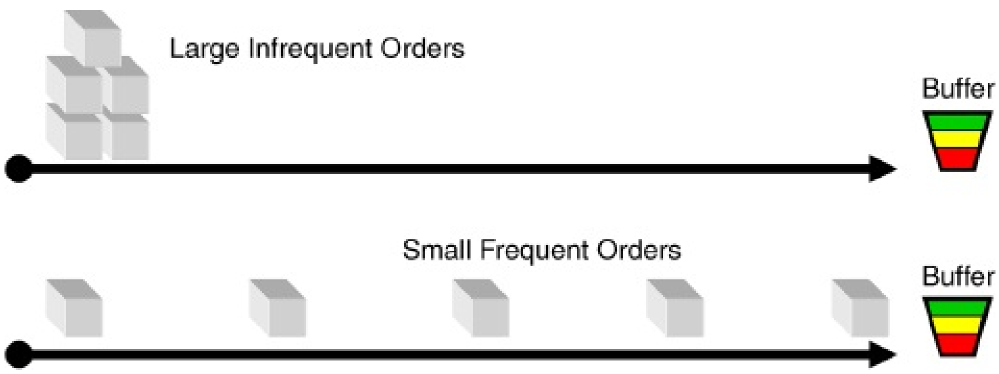

FIGURE 7-6 Large infrequent versus smaller frequent orders

With large infrequent orders, a disruption of a boat, port, truck, etc., could disrupt the entire inbound supply. With smaller more frequent orders, the risk to all the inbound supply is significantly less. Additionally, there can be a cash flow advantage to paying smaller, more frequent invoices as opposed to large infrequent invoices.

Factor 3: Variability

Variability assignment is the next level of assignment. At a minimum, variability can be divided into three segments—high, medium, and low—with the two dimensions of demand and supply variability.

Demand variability is the potential (frequency and size) for spikes in demand with regard to a particular part or SKU number. The variability designation can be calculated by a variety of equations or determined by rules of thumb with intuitive planning personnel.

Heuristically, companies can use the following segmentation for demand variability:

- **High demand variability.** This part is subject to frequent spikes within the lead time.
- **Medium demand variability.** This part is subject to occasional spikes within lead time.
- **Low demand variability.** This part has little to no spike activity—its demand is relatively stable.

Supply variability is the potential for and severity of disruptions in sources of supply for this part or SKU number. This can be calculated by examining the variance of promise dates from actual receipt dates. The caution here is that many of these dates are often determined and managed initially through critical flaws in traditional material requirements planning (MRP). Finally, the number of alternative sources for a part or material can factor into the supply variability equation because the net effect of more sources might be more reliable supply.

Heuristically, companies can use the following segmentation for the three simplest categories of supply variability:

- **High supply variability.** This part or material has frequent supply disruptions.
- **Medium supply variability.** This part or material has occasional supply disruptions.
- **Low supply variability.** This part or material has reliable supply (either a highly reliable single source or multiple alternative sources that can react within the purchasing lead time).

Both forms of variability can also be mathematically calculated and expressed through standard deviation for each particular part. Part variability can be compared with the coefficient of variation within a group of parts. Coefficient of variation is also used in Chapter 6 for the distribution positioning example. This analysis can be useful in understanding the relative distribution of variability within the part population, but it still requires a team to set the specific boundaries of high, medium, and low within the distribution.

As shown previously in Chapter 6, Figure 7-7 illustrates how buffers at different stages within a manufacturing process can experience different levels and types of variability depending on their relationships with each other. Arrowed lines that move from left to right represent supply variability. Arrowed lines that move from right to left represent demand variability. Coming out of the buffer in either direction, they are smoother than when they enter, implying less variability that is passed along.

Purchased parts tend to be influenced almost exclusively by supply variability. One exception is in pure make-to-order (MTO) or engineer-to-order (ETO) environments, where there are no buffers at the subcomponent, intermediate component, or end item level. A pure make-to-order environment would indicate that the inventory positioning factors dictated buffering only for some purchased items. This is an example of why companies cannot skip the inventory positioning step even in MTO and ETO companies. It can dramatically alter which items will end up in which buffer profiles.

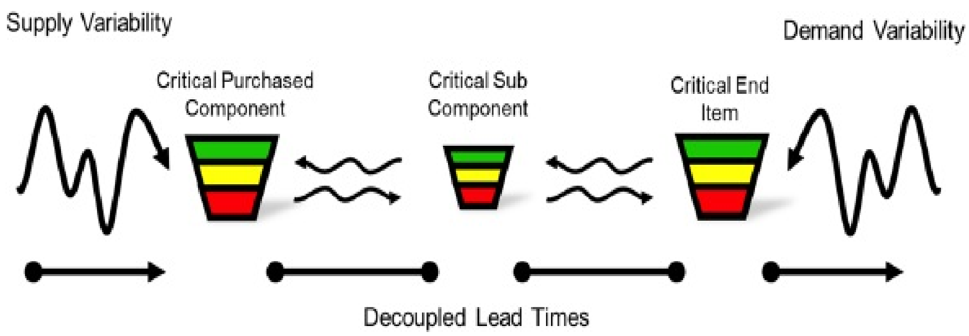

FIGURE 7-7 Multiple buffers and different forms and levels of variability

	Low	Medium	High	Total
Short	12	15	10	37
Medium	6	16	8	30
Long	11	14	8	33
Total	29	45	26	100

FIGURE 7-8 The combination of lead time and variability categories for purchased parts

Figure 7-8 shows the combination of lead time and variability categories for purchased parts from the previous example. This matrix results from the assignment of parts to both lead time and variability categories.

Manufactured parts can be subject to both supply and demand variability depending on how the positioning model is formulated. Manufactured parts are less subject to demand variability if the manufactured item feeds another level of buffered component or end item. These parts are less subject to supply variability if they consume critical parts that are replenished strategically. This is due to the dampening nature of the buffer break walls on the end supply and demand variability.

However, in many cases there can be a blend of demand types experienced by a buffered position. An example of this type of manufactured part is one that is used in subassemblies or the end items (some of which might be buffered) but is also a service part (which might go directly to the customer). This type of manufactured part probably would be subject to more demand variability than a part that fed only some buffered subassemblies or end items. Thus, it is imperative that companies carefully apply the positioning factors described in Chapter 6.

Figure 7-9 shows the combination of lead time and variability categories for manufactured parts from the previous example.

Distributed parts or SKUs will tend to be affected by one variability type depending on their respective locations in the internal supply chain. Distributed parts or SKUs at central buffers can be largely immune from large demand variability if the downstream positions that they feed are sized and managed properly. Part or SKU buffers at downstream locations will be affected almost exclusively by demand variability because they are protected by the central buffer on the supply side. See Chapter 6 for more detail on inventory positioning in distribution networks.

	Low	Medium	High	Total
Short	51	21	5	77
Medium		6	1	7
Long		10	6	16
Total	51	37	12	100

FIGURE 7-9 Lead time and variability categories for manufactured parts

High Variability	61 to 100%+ of Safety Base
Medium Variability	41 to 60% of Safety Base
Low Variability	0 to 40% of Safety Base

FIGURE 7-10 Variability category ranges

The variability category will then be used to supply a "variability factor" to parts within a profile. Figure 7-10 shows the variability factor ranges within each category. The variability factor is applied to another calculation that establishes a base level of safety corresponding to lead time. These calculations will be explained later in this chapter.

Part type, lead time, and variability category assignment are the three basic parameters of buffer profiles. The lead time and variability factor ranges are meant as a conservative guide for planners and buyers to follow. It should be noted that precision in the factor percentage determination is rarely worth the time devoted to it. Roughly right is better than precisely wrong. The difference between 37.6 percent and 38.1 percent will have little effect on the buffer levels. More importantly, the very nature of a buffer allows for quite a bit of imprecision and approximation since the buffers should be adaptive over time. But buffer profiles are only half the necessary requirements to calculate buffers. Next we will turn our attention to the individual part traits that will provide the other necessary condition.

Based on these three factors (part type, lead time category, and variability category), there are 36 basic buffer profiles. Depending on the manufacturing environment, there could be even more derivations and permutations than this. If there is a certain global attribute that makes sense by which parts should be grouped that is not related to variability, lead time, or part type, then another type of buffer profile should be explored and possibly added.

		Part Type				
		Purchased	Manufactured	Distributed	Intermediate	
Lead Time Category	Short	PSL	MSL	DSL	ISL	Low
		PSM	MSM	DSM	ISM	Medium
		PSH	MSH	DSH	ISH	High
	Medium	PML	MML	DML	IML	Low
		PMM	MMM	DMM	IMM	Medium
		PMH	MMH	DMH	IMH	High
	Long	PLL	MLL	DLL	ILL	Low
		PLM	MLM	DLM	ILM	Medium
		PLH	MLH	DLH	ILH	High
						Variability Category

FIGURE 7-11 Basic buffer profile combinations

Figure 7-11 summarizes the 36 different basic buffer profile combinations. Each buffer profile has been designated with a code based on its combination of attributes. Within that naming code, the first letter signifies the part type: “P” for purchased, “M” for manufactured, “D” for distributed, and “I” for intermediate components. Next is the lead time category: “S” for short, “M” for medium, and “L” for long. The third letter represents the variability category: “L” for low, “M” for medium, and “H” for high. For example, a distributed part with medium lead time and low variability is coded as “DML.” A purchased part with a long lead time and high variability is in the buffer profile known as “PLH.” Later examples will use this naming schema with respect to parts’ buffer profile assignments.

Individual Part Attributes

The individual part attributes are properties or numerical values that are specific to the part itself. Many of these properties or values will be found or calculated from the current part master information. In DDMRP there are three specific part attributes that will determine buffer levels for purchased, intermediate, and manufactured buffered items and four specific part attributes that will factor in for distributed buffered items.

Part Average Daily Usage

Average daily usage is a calculated rate of use for each specific part. It is a cornerstone of the buffer equations. Significant changes to the part's ADU will often yield significant impacts to the calculated buffer zones. There are four important considerations in the ADU calculation for each part. These considerations may require planners, buyers, and distribution personnel to consult with other areas of the organization for validation.

Length-of-Period Consideration

Any average is only as relevant as the period over which the equation was applied. If the average is calculated with a shorter horizon, it will be more responsive than a longer horizon. But if the horizon is too short, the ADU will be overreactive and may reproduce the bullwhip effect as the calculated buffers jump between extremes. This will be especially true with products that experience relatively large demand changes within shorter windows of time. Examples would be products that are frequently promoted. Figure 7-12 illustrates this with a product that goes through four major promotional periods within a year. The bars represent the amount of weekly demand experienced by the product. The four major promotional periods coincide with the changing of the seasons and represent significant uplifts in demand.

Figure 7-13 illustrates the differences in ADU values generated by different length-of-period calculations. Three lengths are compared: 52 weeks, 12 weeks, and 1 week. All periods are past looking. Later in the chapter, alternatives to past-looking periods are explored. The ADU values are established by taking the total weekly bucket and dividing by 7. This is assuming a 365-day calendar. Of course, the type of calendar used to calculate the ADU value will make a difference. The calendar should reflect the working calendar of the location of the decoupling point.

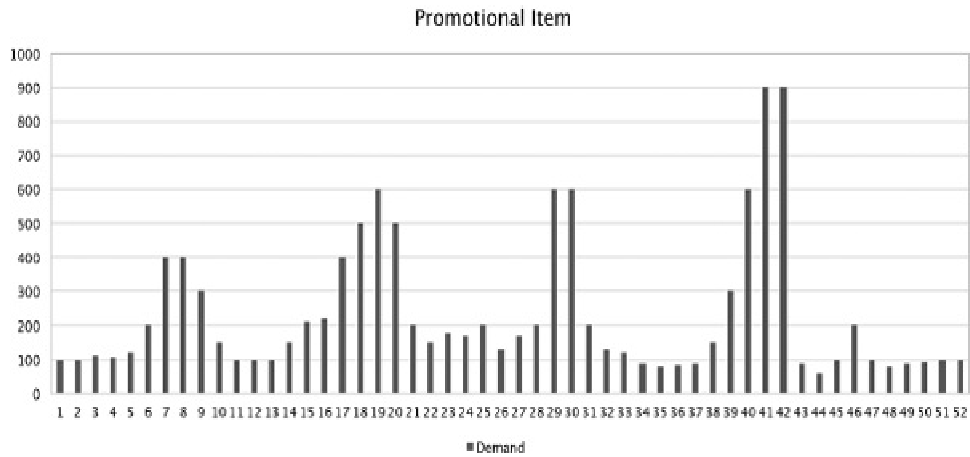

FIGURE 7-12 A frequently promoted item

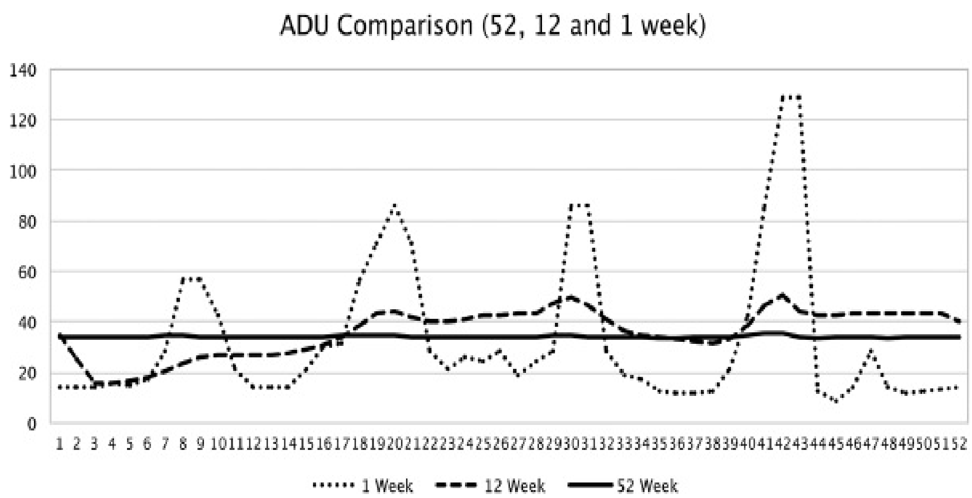

FIGURE 7-13 Differences in ADU value from different period considerations

Figure 7-13 demonstrates a significant difference between a period length rolling 52 weeks in the past versus one that rolls from 1 week in the past. The 52-week roll has very little fluctuation in the calculated ADU, while the 1-week roll produces dramatic swings. Many environments would simply not be able to react to the massive fluctuations in the buffers that would result from a 1-week-period calculation. In this regard, the length-of-period consideration further connects the concept of a stock buffer to time. A stock buffer is meant to buy time, especially when there are shared resources responsible for supporting many different buffers in that environment. In this example ADU moves from a low of under 16 to a high of over 44 when using a 12-week period. In this example, this is sufficient movement to allow the buffer to reasonably flex up and down while not creating a bullwhip effect on the environment.

Frequency-of-Update Consideration

Performing a calculation and updating a part master with the result of that calculation are two different tasks. It is important that there is not too much time between part ADU updates. The longer the time frame, the choppier the ADU will become. This can create buffers that jump or fall in a way that transfers variability to other parts of the organization. Steps should be taken to update the part's ADU on a frequent basis. In most cases this would mean daily or no longer than weekly updates. Ideally, the part's record would be automatically updated as the ADU calculation is performed. Unlike traditional MRP where the increased frequency of MRP updates increases system nervousness, under DDMRP the more frequent the updates of relevant information such as this, the more stable the environment. Infrequent updates (monthly or quarterly) can cause a major bullwhip effect in a DDMRP environment.

Past, Forward, or Blended Consideration

ADU at any point in time will also be heavily influenced by the length-of-period horizon that the consideration is applied to—the past, the future, or a blend of both.

Figure 7-14 shows the difference for the promotional item from the previous example with a 12-week period consideration using an ADU based on 12 weeks of past sales versus an ADU using 12 weeks of forecasted sales. The left Y axis is the demand for the product, while the right Y axis is the calculated ADU for each type of period consideration for any particular week.

An ADU value calculated from a rolling past period can be problematic when there is an expected upsurge in demand that is not being properly reflected in the current or impending ADU. With regard to the promotional example from Figure 7-14, this problem can be seen looking 12 weeks into the past. There is a known upsurge coming in periods 6–9; yet in period 1 (the current period) the ADU is reducing dramatically over the next two periods so that the ADU will be under 16 while the forward-looking value is at 27 and has been trending up. Additionally, the forward-looking ADU appears to better anticipate the major drop-off in demand in week 43.

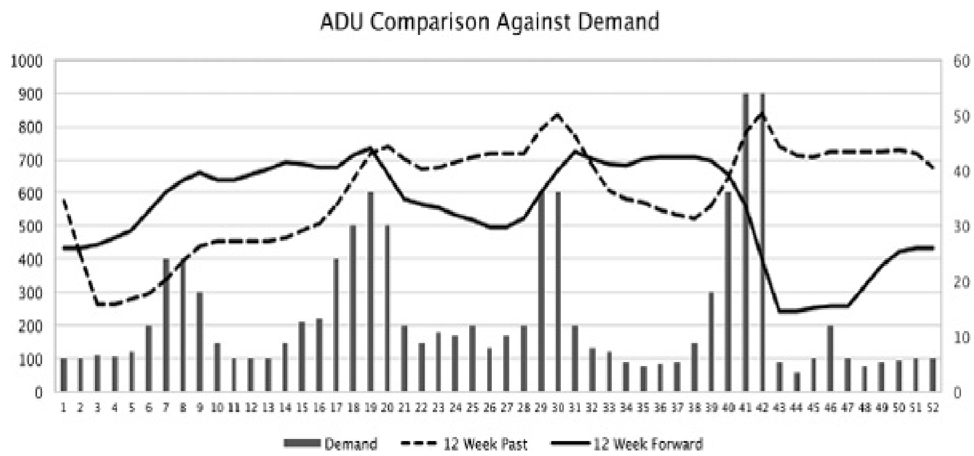

FIGURE 7-14 Past and forward ADU values against the demand for the promotional product

Is forward looking always a better way to calculate ADU? Remember that this scenario assumes that the forecasted rate of sales will actually occur. This is where some degree of intuition must be brought to bear within the environment. How confident is the planning team in the numbers being forecasted? Are these numbers composed of firm orders, or are they more wishful thinking? If these numbers are relatively firm, then it could be appropriate that this part be placed on a forward-looking ADU. "Relatively firm" means that the majority of this demand is already spoken for through a major planned promotion.

Note that a forward-looking ADU is the incorporation of forecast into the DDMRP buffers but not into the DDMRP ordering mechanism. This is extremely important to keep in mind and will be further discussed in Chapters 8 and 9.

If calculating ADU using a forecast seems too aggressive or risky given the environment, there is a final option. Given the fact that the farther one forecasts into the future, the less accurate the demand signal is, then maybe the solution is to limit forward-looking ADU calculation to a shorter future time range that is more accurate. But that amount of time will create an ADU calculation that might be too short and overresponsive. In this case, blending the past sales with the future expected sales would make sense. Figure 7-15 shows the ADU value with a blended approach contrasted with past- and forward-looking versions.

Also another option will be considered in Chapter 8 to dynamically adjust the ADU value using a demand adjustment factor (based on historical patterns or known events) rather than a forward-looking ADU.

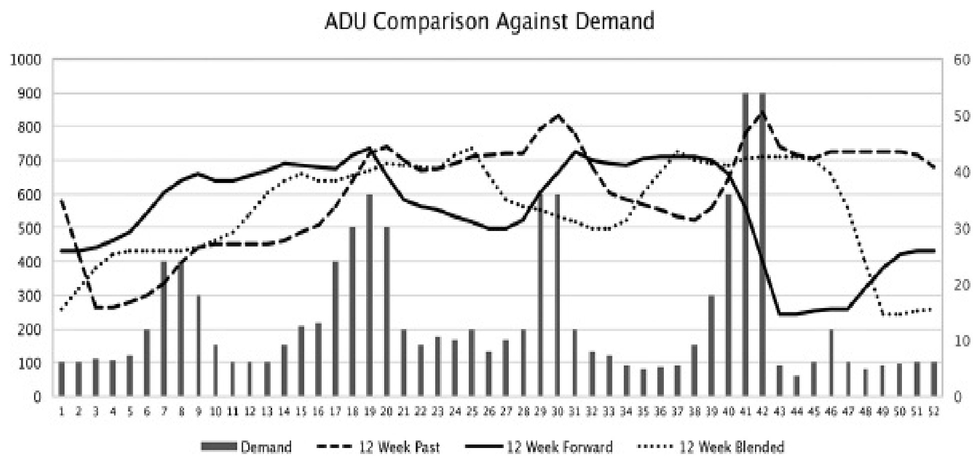

FIGURE 7-15 Past, forward, and blended versions of ADU

ADU Exceptions

If a significant event occurs that has altered demand profiles dramatically within a relevant range, then planners should be cautious about adjusting the ADU. This rapid change in demand should generate what is called an ADU alert. The parameters of an ADU alert must be defined in terms of quantity and time against the calculated ADU. For example, if a part's ADU changes more than *x* percent within a *y* time frame, then the alert is generated.

This requires a high and low threshold and a horizon to be established. For example, if a company sets a high threshold of 300 percent over three days, it would be asking for ADU alerts to be generated when daily demand was triple the daily norm (ADU) for three days in a row. Conversely, a low threshold could be established of 10 percent for five days, meaning an ADU alert would occur when demand was 10 percent of the ADU for five days in a row. These are simply examples; the length and the threshold will be unique to the environment.

If the significant events are anomalous and will revert to normal behavior, then the abnormal usage (or lack of usage) should be excluded in the ADU update equation. However, if the dramatic shift is indicative of what the future might look like for this part, then that information should be included or at the least factored in to some extent through a demand adjustment factor. Demand adjustment factors are discussed in Chapter 8.

Establishing ADU for Items with No History

A company may have no history by which to generate an ADU for strategic items. This could occur in a company that has a legacy system or no system of record at all. If a company has a modern planning system and these data are not available, then the lack of availability would often be due to poor implementation.

The collection of actual usage data is a necessary condition for an adaptive DDMRP implementation. Data need to be collected and phased in with estimations. Planners may need to consult with other relevant personnel to calculate an estimation of ADU that seems realistic for each item. Then over the length-of-period consideration, the actual ADU can be blended with the estimated ADU. At the end of the length-of-period consideration, an ADU that represents actual history over the entire length of period is calculated. Figure 7-16 illustrates this concept with a product that has a 12-week past-looking ADU calculation. An estimated ADU of 3,000 was used to build the buffer, and in the next 12 weeks an ADU of 3,224.9 emerges. Each week the actual usage replaces a week of estimated usage.

Part Lead Time

Another critical individual part input into the buffer equation is the part's unique lead time as measured in discrete units of time (most often in days). For any manufactured or intermediate item, this lead time should be the decoupled lead time of the part. The value of decoupled lead time is discussed in Chapter 6. For purchased parts, the purchasing lead time from the part master should be used. For distributed parts, the transportation lead time from the sourcing unit or hub should be used. In some cases, additional time may be added for necessary staging or receipt if applicable. If the quality assurance or incoming inspection time is significant, then at a minimum this time should be included.

Week	Estimated ADU	Included In 12-Week Estimate	Actual ADU	Included Actual	Sum (Estimate and Actual)	Actualized ADU
1	3,000.0	36,000.0	2,678	2,678	35,678.0	2,973.2
2	3,000.0	33,000.0	3,174	5,852.0	35,852.0	2,987.7
3	3,000.0	30,000.0	2,987	8,839.0	35,839.0	2,986.6
4	3,000.0	27,000.0	3,412	12,251.0	36,251.0	3,020.9
5	3,000.0	24,000.0	3,541	15,792.0	36,792.0	3,066.0
6	3,000.0	21,000.0	3,210	19,002.0	37,002.0	3,083.5
7	3,000.0	18,000.0	3,810	22,812.0	37,812.0	3,151.0
8	3,000.0	15,000.0	2,978	25,790.0	37,790.0	3,149.2
9	3,000.0	12,000.0	3,214	29,004.0	38,004.0	3,167.0
10	3,000.0	9,000.0	3,611	32,615.0	38,615.0	3,217.9
11	3,000.0	6,000.0	2,874	35,489.0	38,489.0	3,207.4
12	3,000.0	3,000.0	3,210	38,699.0	38,699.0	3,224.9

FIGURE 7-16 Actualizing ADU over the length-of-period consideration

Part Minimum Order Quantity

Ordering policies (minimums, maximums, and multiples) complicate planning and supply scenarios but are a fact of life for planners. Many of these ordering policies are based on valid data and sound assumptions; many are not. It is a given that there will be parts and SKUs that do require minimum order quantities. Minimum order quantities (MOQs) can affect buffer levels, especially when they are large in relation to the rate of use. These are called “significant” MOQs and will have a direct impact on the sizing of the buffer through the green zone. The qualifying characteristics that make an MOQ significant will be examined later in this chapter.

Part Location

Location is an attribute that is unique to distributed part types chosen for strategic replenishment assuming more than one forward distribution point. Refer to the distribution example from Chapter 6 for a description of the considerations for distribution. For each part that is distributed through the network, a separate ADU and lead time will exist for each distribution point. Figure 7-17 shows the ADU values from Chapter 6 for distributed part 123.

Part #123		
Location	**ADU**	**Lead Time**
Warehouse 1	7.4	2 Days
Warehouse 2	2.9	2 Days
Warehouse 3	6.4	3 Days
Warehouse 4	3.1	4 Days
Hub	19.7	1 Day

FIGURE 7-17 Distributed part 123 ADU and lead time values by location

Each physical distribution point including the hub has a distinct average daily usage and lead time. While warehouses 1 and 2 have the same value for lead time, it must be understood that those numbers have no relationship to each other; they are separate and distinct transportation times that can change independent of each other.

Figure 7-18 summarizes the combination of the factors discussed so far in this chapter that will combine to create each strategic decoupling point buffer. As previously discussed, individual part attributes will combine with the lead time and variability factors from a buffer profile assignment to create these unique buffer levels.

Calculating Replenished Part Buffer Levels and Zones

Replenished part buffers are composed of three color-coded zones: green, yellow, and red. Each zone has a specific purpose and will vary in size and proportion depending on the combination of the buffer profile and the individual part traits discussed previously in this chapter. It is important to note that the buffer is not simply divided into equal thirds. Understanding the purpose and calculation of each zone is crucial to understanding how DDMRP buffers produce their results as well as how they compare with other stock management techniques.

Figure 7-19 displays the individual and buffer profile attributes of an example part that will be used to learn how to calculate buffers. This part has the buffer profile of MML, meaning it is a manufactured part with a medium lead time and a low variability setting. This is shown in the row in Figure 7-19 called “Buffer Profile.” Additionally, in the “Buffer Profile” row, the lead time and variability category has been assigned a factor. This factor is noted in the parentheses after the lead time and variability designator, respectively. For this example, the lead time factor is 0.5, and the variability factor is 0.33. These factors fall within the ranges for each subcategory discussed earlier in the chapter.

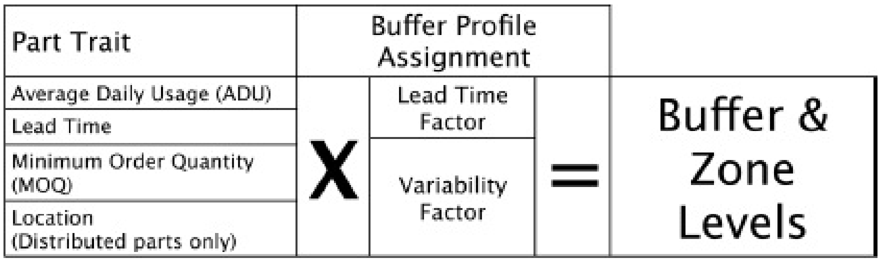

FIGURE 7-18 Buffer profile factors

Example	
Average Daily Usage	10
Buffer Profile	M, M (.5), L (.33)
MOQ	50
Imposed or Desired Order Cycle (DOC)	7 Days
Decoupled Lead Time (DLT)	12 Days

FIGURE 7-19 Example part buffer profile and individual attributes

The Green Zone

As discussed earlier in this chapter, the green zone is the heart of the supply order generation process embedded in the buffer and will determine the average order frequency and typical order size for the part chosen for this example. The green zone is determined by one of three factors. Whichever factor yields the greatest number determines the size of the green zone. Each factor needs to be calculated for our example in order to determine which will qualify as the green zone value.

Option 1: An Imposed or Desired Minimum Order Cycle

An order cycle is simply the number of expected days between orders. It can be an imposed factor through the use of a product scheduling wheel or be a desired average number of days between orders. Either way the equation is the same. The equation for calculating the green zone based on order cycle is simply ADU × desired or imposed order cycle days. For the part in Figure 7-19, this would yield 70.

Option 2: Using a Lead Time Factor

A green zone can be calculated using the lead time factor discussed previously in this chapter. This lead time factor is expressed as a percentage of usage within a full lead time of the part. That percentage will fall within a range corresponding to the lead time category of the part illustrated in Figure 7-19. The formula for producing the green zone value using this technique is decoupled lead time × ADU × lead time factor. For the part in Figure 7-19, this option would yield 60.

Option 3: Minimum Order Quantity (if Applicable)

If the green zone is about supply order generation and frequency and the part has a minimum order quantity, then the minimum order quantity can be relevant in determining the green zone. In short, the green zone should never be less than the minimum order quantity. If the minimum order quantity yields the largest value as a green zone, then that minimum order quantity is deemed significant. The minimum order quantity must be compared against the desired order cycle quantity value and the value created by using the lead time factor. For the part in Figure 7-19, this option would yield 50.

To summarize the options:

- Calculating the green zone using the order cycle yields a green zone of 70 [ADU (10 per day) × order cycle (7 days)].
- Calculating the green zone using the lead time factor yields a green zone of 60 [ADU (10 per day) × DLT (12 days) × lead time factor (0.5)].
- The minimum order quantity is 50.

The largest value is actually the desired order cycle value of 70.

The Yellow Zone

The yellow zone is the heart of the inventory coverage in the buffer. It is the easiest and most straightforward zone to calculate in a buffer. The yellow zone is always calculated as ADU multiplied by the decoupled lead time. In our example in Figure 7-19, the yellow zone of this part would be sized at 120 pieces [ADU (10 per day) × DLT (12 days)].

The Red Zone

The red zone is the embedded safety in the buffer. The higher the variability associated with the part or SKU, the larger the red zone will be. Calculating the red zone requires three sequential equations:

1. **Establish the "red base."** The red base is established by multiplying the lead time factor by the average daily usage by the lead time. This lead time factor corresponds to the same ranges used for the green zone calculation but can have a different numerical value. In our example in Figure 7-19, the part falls in the medium lead time category. For simplicity in the example, we will use the same percentage lead time factor (50 percent, or 0.5) that was used in the green zone calculation. Thus for this example the red base value is 60 units [ADU (10 per day) × DLT (12 days) × lead time factor (0.5)].
2. **Establish the "red safety."** The red safety is calculated as a percentage of the red base. The percentage used is determined by the variability factor. Like the lead time factor, there are ranges of variability factors depending on whether a part experiences high, medium or low variability. Our example part falls in the low variability category with a variability factor of 33 percent (0.33). Our red safety value is calculated as 20 [red base of 60 × 0.33 (variability factor)].
3. **Calculate the total red zone by adding the red base to the red safety.** Our example part will have a red zone of 80 [red base (60) + red safety (20)].

This may seem like a complicated way to calculate a safety level. However, the safety level is related to both lead time and variability through visible and independent factors. The red base is the safety with respect to the ability to recover due to time. The red safety zone factors that number based on the variability factor within that lead time period. Thus parts that have the same variability factor applied but are in different lead time categories will have total red zones that are proportionately different.

Once all zones are calculated, we can add them together to get a total buffer. Figure 7-20 summarizes the buffer calculations and subsequent zone calculations. The top of the buffer is called "top of green" and is the summation of all zones. In this case it is 270 [red (80) + yellow (120) + green (70)].

A few observations can be inferred at this point about the zones and the relationships they may have with each other:

- The average order frequency can be calculated as the green zone divided by the ADU. In this example it is 7 days. This means, on average, with no variability, this part will be ordered every 7 days.
- The red zone divided by ADU tells us just how much embedded safety is within the buffer. The red zone of this buffer represents 8 days of safety [red zone (80) divided by 10 (ADU)].

Example Part Buffer Calculation

Average Daily Usage	10	Green Zone	70
Buffer Profile	M, M (.33), L (.33)		LT Factor: 60 (DLT (12)x ADU (10) x Lead Time Factor (.5))
MOQ	50		Minimum Order Quantity: 50
Imposed or Desired Order Cycle (DOC)	7 days		Order Cycle: 70 (7(OC) x 10(ADU))
Decoupled Lead Time (DLT)	12 days	Yellow Zone	120 (12(DLT) x 10(ADU))
		Red Zone	80 (Red Base (60) + Red Safety (20))
			Red Base: 60 (DLT (12)x ADU (10) x Lead Time Factor (.5))
			Red Safety: 20 (Red Base (60) x Variability Factor (.5))

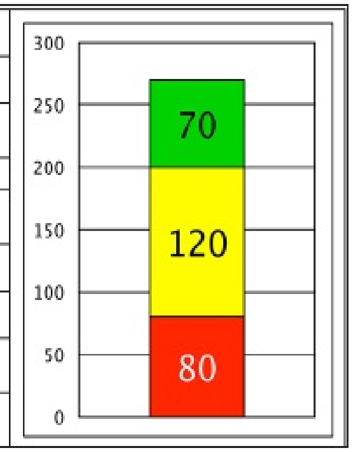

FIGURE 7-20 Example part buffer calculation summary

- The red base would equal the green zone if the order cycle had been lower than the green zone calculated using the same lead time factor.
- The relationship between the green zone and yellow zone can also tell us how many open supply orders we can expect to see at any one time. In our example, dividing the yellow zone (120) by the green zone (70) means that there is an average of around 2 open supply orders at any one time. For longer lead time parts, this estimates how many "orders in the pipeline" can be expected. This can provide a quick way to analyze whether the part fits the buffer profile provided for it and whether the buyer or planner is launching supply orders within a relevant time frame.

Figure 7-21 shows a summary of how buffers are calculated. Note that when calculating the buffers by hand, it is easiest to calculate the yellow zone first since green and red base calculations are a percentage of ADU × DLT (yellow zone). By establishing the yellow zone first, the green and red base equations are quicker to perform.

To further practice calculating buffers as well as continuing to build the case for the value of finding the right inventory positions, the example with Company ABC continues.

Continuing with Company ABC

In Chapter 6 an example called Company ABC was introduced to demonstrate the decoupling point positioning considerations. The various decoupling point positioning considerations were used to create a model that better protected the environment from variability, leveraged inventory, and compressed total inventory requirements in the face of certain market expectations. That model had several iterations. Now the buffers for those different iterations are calculated in the following examples.

The starting situation with Company ABC had three finished items already chosen for strategic buffering due to their relatively short customer tolerance time in relation to their decoupled lead times. Figure 7-22 is the starting situation in Chapter 6 with the three end items (FPA, FPB, and FPC) circled.

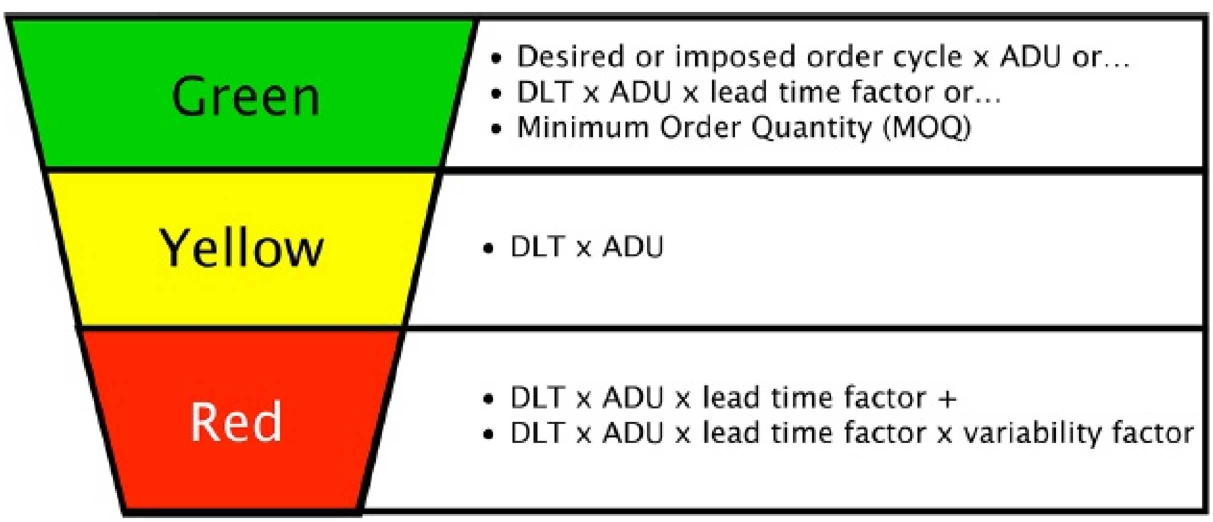

FIGURE 7-21 The buffer equation summary

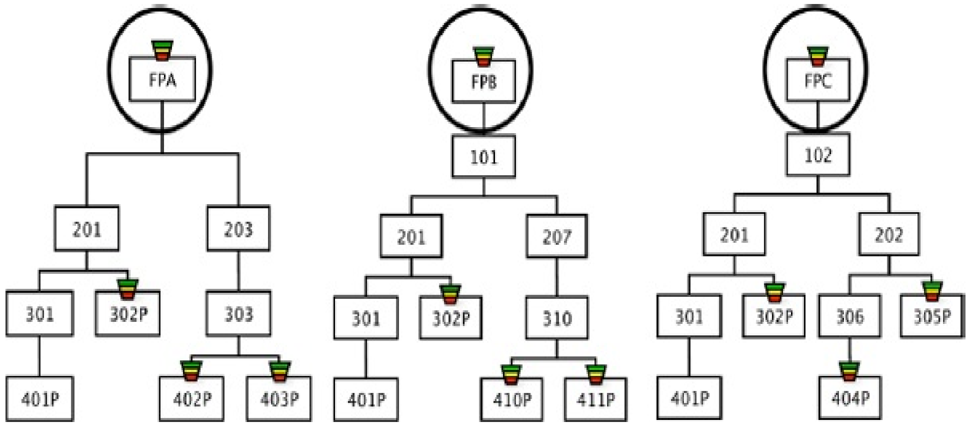

FIGURE 7-22 Initial decoupling points for Company ABC

The buffer profile settings and the individual part attributes are required for each of these end items in order to calculate their buffers. Figure 7-23 contains the buffer profile parameters that Company ABC is using. Note that Company ABC created a separate buffer profile category for intermediate components that impacts their respective variability factors. This is done with the recognition that intermediate components tend to experience less variability since they are protected by parent and/or component positions. These profile settings will be applied to the different iterations of positioning in Chapter 6.

Parent Make Items (M)		Lead Time Factor	Variability Category	Variability Factor
Long Lead Time	8+ Days	.25	High	.75
Medium Lead Time	3–8 Days	.4	Medium	.5
Short Lead Time	1–2 Days	.7	Low	.25
Intermediate Make Items (I)				
Long Lead Time	8+ Days	.25	High	.7
Medium Lead Time	3–8 Days	.4	Medium	.5
Short Lead Time	1–2 Days	.7	Low	.2
Purchased Items (P)				
Long Lead Time	20+ Days	.3	High	.75
Medium Lead Time	11–19 Days	.5	Medium	.5
Short Lead Time	1–10 Days	.7	Low	.25

FIGURE 7-23 Buffer profile settings for Company ABC

Iterations

In the first iteration, the buffer sizes of FPA, FPB, and FPC will start the example. In order to calculate these buffers, the buffer profile settings in Figure 7-23 are combined with each individual part's characteristics. Figure 7-24 has the relevant part data to perform the buffer setting calculation. All parent items have been set to medium variability in this first iteration because they are subject to both demand variability and supply variability from a long, tightly coupled lead time chain.

The buffer profile assignments are the sequence of part type, lead time category (with the lead time factor in parentheses), and variability category (with the variability factor in parentheses). For example, FPA has a profile MLM. This means it is a manufactured part, has a long lead time, and is assigned to the medium variability category.

With the buffer profile settings and the part attributes, the buffers and their respective zones can be calculated for each end item. Figure 7-25 contains the calculations for each buffer. The larger bolded number represents each zone size for each part.

For each part the yellow zone is calculated by multiplying the part's decoupled lead time by the part's average daily usage. For example, FPA has a DLT of 20 and an ADU of 250, yielding a yellow zone of 5,000.

FPA	
Average Daily Usage	250
Buffer Profile	M, L (.25), M (.5)
MOQ	250
Desired Order Cycle (DOC)	3 Days
Decoupled Lead Time (DLT)	20 Days
FPB	
Average Daily Usage	100
Buffer Profile	M, L (.25), M (.5)
MOQ	250
Desired Order Cycle (DOC)	3 Days
Decoupled Lead Time (DLT)	23 Days
FPC	
Average Daily Usage	300
Buffer Profile	M, L (.25), M (.5)
MOQ	250
Desired Order Cycle (DOC)	3 Days
Decoupled Lead Time (DLT)	23 Days

FIGURE 7-24 FPA, FPB, and FPC part attributes at Company ABC

FPA		FPC	
Green Zone	**1250**	Green Zone	**1725**
	LT Factor: 1250 (5000 x .25)		LT Factor: 1725 (6900 x .25)
	Minimum Order Quantity: 250		Minimum Order Quantity: 250
	Order Cycle: 750 (3(DOC) x 250(ADU))		Order Cycle: 900 (3(DOC) x 300(ADU))
Yellow Zone	**5000** (20(DLT) x 250(ADU))	Yellow Zone	**6900** (23(DLT) x 300(ADU)
Red Zone	**1875** (1250 + 625)	Red Zone	**2588** (1725 + 863)
	Base: 1250 (5000 x .25)		Base: 1725 (6900 x .25)
	Safety: 625 (1250 x .5)		Safety: 863 (1725 x .5)
FPB			
Green Zone	**575**		
	LT Factor: 575 (2300 x .25)		
	Minimum Order Quantity: 250		
	Order Cycle: 300 (3(DOC) x 100(ADU))		
Yellow Zone	**2300** (23(DLT) x 100(ADU))		
Red Zone	**863** (575 + 287.5)		
	Base: 575 (2300 x .25)		
	Safety: 287.5 (575 x .5)		

FIGURE 7-25 Buffer calculations for FPA, FPB, and FPC

Each part's green zone is determined by comparing three numbers and taking the largest. In all three parts in this first iteration, the green zone is sized as the lead time factor multiplied by the ADU multiplied by the DLT.

Each part's red zone is arrived at through the red zone sequential equation. First, the base is established by applying a lead time factor. The product of that equation is then multiplied by the variability factor. The red zone base is added to the red zone safety to get the total red zone setting. As Figure 7-25 shows, FPA's base is 1,250 after a lead time factor of 0.25 is applied to full usage within the DLT. That 1,250 is then multiplied by the variability factor of 0.5. This makes the safety portion 625. Next 1,250 and 625 are added together to yield a total red zone of 1,875. *Note:* The same lead time factor for the green and red zones is being used for this example. This is not required but is done so here for simplicity.

Figure 7-26 is the graphical depiction of the end item buffers in the first iteration of the Company ABC example. The green zone for each part is at the top of the stack, the yellow zone is in the middle, and the red is at the bottom. The next iteration involves decoupling 201 in order to reduce the decoupled lead time for each parent particularly FPA. Compressing each of the decoupled lead times for each end item should definitely lower the respective buffer levels, but it will also require building an additional buffer for 201.

Figure 7-27 shows the impact on FPA buffers by decoupling at 201. Decoupling 201 has a major impact on some of the critical inputs to the buffer equation. These changes are highlighted in the shaded boxes (“Buffer Profile” and “DLT After Decoupling”). Decoupling 201 has compressed the decoupled lead time for FPA to seven days. That moves FPA into a different buffer profile for two reasons. First, the shorter lead time moves the part from the long lead time profile (with a lead time factor of 0.25) to the medium lead time profile (with a lead time factor of 0.4). Second, by decoupling at 201 the end items are subject to much less supply variability from that leg of the product structure. This has resulted in FPA moving from the medium variability category (with a variability factor of 0.5) to the low variability category (with a variability factor of 0.25).

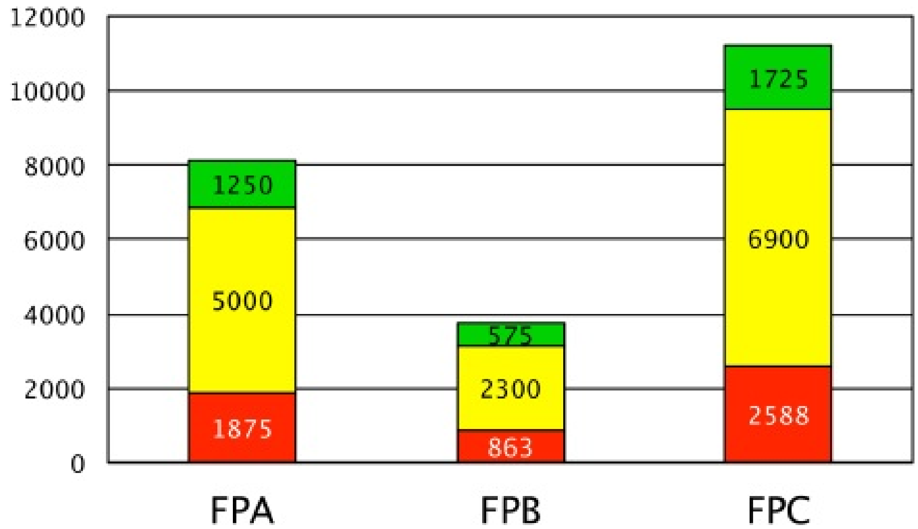

FIGURE 7-26 End item buffer sizes for Company ABC

Buffer Comparison (FPA)

Before 201 Decoupling			
Average Daily Usage	250	Green Zone	1250
Buffer Profile	M, L (.25), M (.5)		LT Factor: 1250 (5000 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle (DOC)	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
Decoupled Lead Time (DLT)	20 days	Yellow Zone	5000 (20(DLT) x 250(ADU))
		Red Zone	1875 (1250 + 625)
			Base: 1250 (5000 x .25)
			Safety: 625 (1250 x .5)
After 201 Decoupling			
Average Daily Usage	250	Green Zone	750
Buffer Profile	M, M (.4), L (.25)		LT Factor: 700 (1750 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
DLT After Decoupling	7 days	Yellow Zone	1750 (7(DLT) x 250(ADU))
		Red Zone	875 (700 + 175)
			Base: 700 (1750 x .4)
			Safety: 175 (700 x .25)

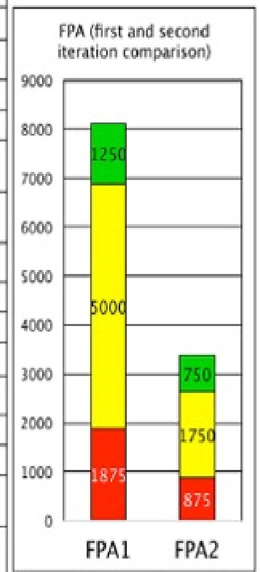

FIGURE 7-27 FPA buffer comparison (first to second iteration)

The impact on the buffer zones of FPA is significant. Figure 7-27 also illustrates a side-byside comparison of buffer sizing, with the first iteration labeled “FPA1” and the second iteration labeled “FPA2.” The top of green level for the first iteration is 8,125, while the top of green level for FPA2 is 3,375.

Similar to the impact on FPA, decoupling 201 also has major implications for FPB. The FPB decoupled lead time has been compressed from 23 days to 9 days. For the same reasons FPB also moves to a medium lead time profile (with a lead time factor of 0.4) and a low variability category (with a variability factor of 0.25). Figure 7-28 depicts the buffer levels between the first and second iteration for FPB.

FPC also experiences a large compression of decoupled lead time (from 23 to 8 days). The same shift in buffer profile occurs with regard to FPC. Figure 7-29 shows the comparison of the first and second iterations.

Obviously, decoupling with 201 reduces the level of average working capital contained in the end item buffers. But by how much? To answer this question there are two requirements. First, the additional required buffer levels in components to support these compressed lead times must be calculated. In this case the buffer required for 201 must be calculated. Second, the average on-hand levels in all buffers involved in the compression (component and end item) must be calculated and compared against the noncompressed scenario. Simply comparing the top of green levels provides a distorted inventory impact picture.

Buffer Comparison (FPB)			
Before 201 Decoupling			
Average Daily Usage	100	Green Zone	**575**
Buffer Profile	M, L (.25), M (.5)		LT Factor: 575 (2300 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 300 (3(DOC) x 100(ADU))
Decoupled Lead Time (DLT)	23 days	Yellow Zone	**2300** (23(DLT) x 100(ADU))
		Red Zone	**863** (575 + 287.5)
			Base: 575 (2300 x .25)
			Safety: 287.5 (575 x .5)
After 201 Decoupling			
Average Daily Usage	100	Green Zone	**360**
Buffer Profile	M, M (.4), M (.5)		LT Factor: 360 (900 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 300 (3(DOC) x 100(ADU))
DLT After Compression	9 days	Yellow Zone	**900** (9(DLT) x 100(ADU))
		Red Zone	**540** (360 + 180)
			Base: 360 (900 x .4)
			Safety: 180 (360 x .5)

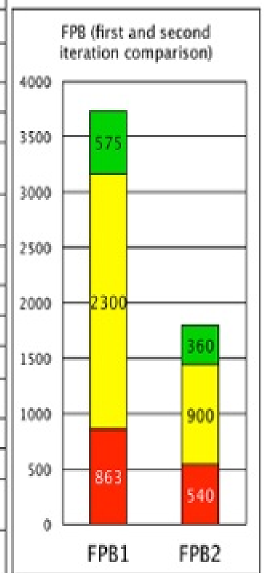

FIGURE 7-28 FPB buffer comparison (first to second iteration)

Buffer Comparison (FPC)			
Before 201 Decoupling			
Average Daily Usage	300	Green Zone	1725
Buffer Profile	M, L (.25), M (.5)		LT Factor: 1725 (6900 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 900 (3(DOC) x 300(ADU))
Decoupled Lead Time (DLT)	23 days	Yellow Zone	6900 (23(DLT) x 300(ADU)
		Red Zone	2588 (1725 + 863)
			Base: 1725 (6900 x .25)
			Safety: 863 (1725 x .5)
After 201 Decoupling			
Average Daily Usage	300	Green Zone	960
Buffer Profile	M, M (.4), L (.25)		LT Factor: 960 (2400 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
DLT After Compression	8 days	Yellow Zone	2400 (8(DLT) x 300(ADU)
		Red Zone	1200 (960 + 240)
			Base: 960 (2400 x .4)
			Safety: 240 (960 x .25)

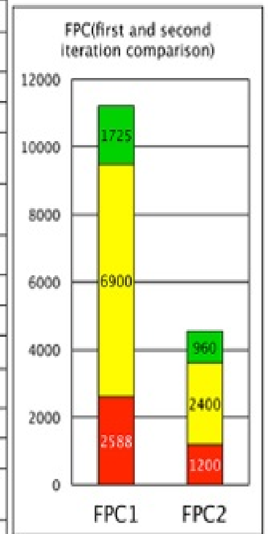

FIGURE 7-29 FPC buffer comparison (first to second iteration)

To meet the first requirement, the buffer levels from the Company ABC example can be calculated. The second requirement necessitates an understanding of how DDMRP buffers plan and receive supply orders. Thus the second requirement will be deferred until Chapter 9 when Company ABC is revisited, including the calculation of the full working capital impact in the example.

Next the decoupling point buffer for 201 is calculated. Figure 7-30 shows the profile assignment and part attributes that will determine 201’s buffer and zone levels. For this example, average daily usage is the sum of the end items (FPA, FPB, and FPC). If 201 was used in other parent items or was sold on its own (e.g., a service part), that demand would have to be incorporated into the ADU. Part 201 is assigned to the manufactured intermediate component, long lead time (lead time factor of 0.25) and medium variability (variability factor of 0.5) profile.

The 201 buffer appears substantial due to a relatively large amount of demand and its long decoupled lead time. The top of green is at 20,070. At face value the vast majority of the inventory savings in the parents appears to have simply shifted to this component. That is true from a quantity perspective only.

Keep in mind that end items have more direct material dollars contained in them, at a minimum. Furthermore, as discussed previously, it will be distortive to focus on the top of green numbers. What is needed to make a fair financial comparison between iterations is what the average working capital levels will be under the different scenarios. This is discussed in Chapter 9.

Buffer Worksheet (201)			
Average Daily Usage	650	Green Zone	3088
Buffer Profile	I, L (.25), M (.5)		LT Factor: 3088 (12350 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	19 days	Yellow Zone	12350 (19(DLT) x 650(ADU)
		Red Zone	4632 (3088 + 1544)
			Base: 3088 (12350 x .25)
			Safety: 1544 (3088 x .5)

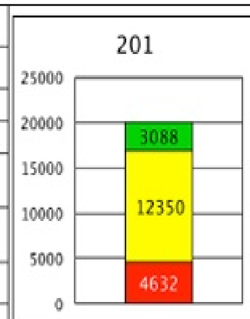

FIGURE 7-30 The 201 buffer profile, part attributes, and calculated buffer levels

The next iteration of the Company ABC example involves buffering the 203 component in order to potentially eliminate the need to hold FPA stock. FPA's customer tolerance time is three days, and buffering both 201 and 203 would allow FPA to be made within one day, thus moving FPA to assemble to order status. Part 203 is an intermediate component with its purchased items decoupled, allowing it to be in the low variability category (variability factor of 0.2). Its decoupled lead time of six days still places it in the medium lead time category. Figure 7-31 shows the profile assignment and part attributes that will determine 203's buffer and zone levels. Figure 7-31 also depicts the completed buffer calculations for 203.

The final iteration of the Company ABC example from Chapter 6 involved compressing the decoupled lead time for 201. The purchased part 401P was selected to be decoupled. By stocking 401P, the decoupled lead time for component 201 is reduced from 19 days to 9 days. However, 9 days still qualifies as a long lead time manufactured part for Company ABC. The lead time category and lead time factor (0.25) remain the same. The variability category has been reduced from medium (variability factor 0.5) to low (variability factor 0.2) because external supplier variability has been mitigated by the 401P buffer position. It is worth noting that after 401P is buffered, 201 becomes the first buffer in the Company ABC example in which the green zone qualifies as the desired order cycle.

Buffer Worksheet (203)			
Average Daily Usage	250	Green Zone	**750**
Buffer Profile	I, M (.5), L (.2)		LT Factor: 750 (1500 x .5)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
Decoupled Lead Time (DLT)	6 days	Yellow Zone	**1500** (6(DLT) x 250(ADU))
		Red Zone	**900** (750 + 150)
			Base: 750 (1500 x .5)
			Safety: 150 (750 x .2)

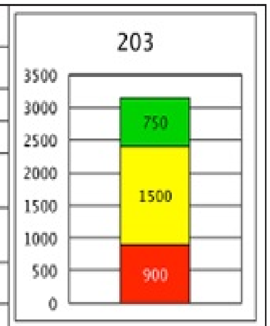

FIGURE 7-31 Component 203's buffer profile, part attributes, and calculated buffer levels

One final buffer calculation is required to complete the Company ABC example from Chapter 6. Part 401P was decoupled in order to compress 201's lead time. Next is the calculation of the 401P buffer. Figure 7-33 shows the profile assignment and part attributes that will determine 401P's buffer and zone levels. Figure 7-33 also depicts the completed buffer calculations for 401P.

Figure 7-34 summarizes the buffer levels for all buffers for each example iteration at Company ABC that were impacted by the positioning example in Chapter 6. The buffers at 302P, 402P, 403P, 410P, 411P, and 404P are not displayed, as there are no significant implications for those buffers. Something that can be noted is that the positioning example tends to create a shift of inventory to the lower levels of the product structure.

Buffer Comparison (201)			
Before 401P Decoupling			
Average Daily Usage	650	Green Zone	**3088**
Buffer Profile	I, L (.25), M (.5)		LT Factor: 3088 (12350 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	19 days	Yellow Zone	**12350** (19(DLT) x 650(ADU))
		Red Zone	**4632** (3088 + 1544)
			Base: 3088 (12350 x .25)
			Safety: 1544 (3088 x .5)
After 401P Decoupling			
Average Daily Usage	650	Green Zone	**1950**
Buffer Profile	I, L (.25), L (.2)		LT Factor: 1463 (5850 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
DLT After Compression	9 days	Yellow Zone	**5850** (9(DLT) x 650(ADU))
		Red Zone	**1756** (1463 + 293)
			Base: 1463 (5850 x .25)
			Safety: 293 (1463 x .2)

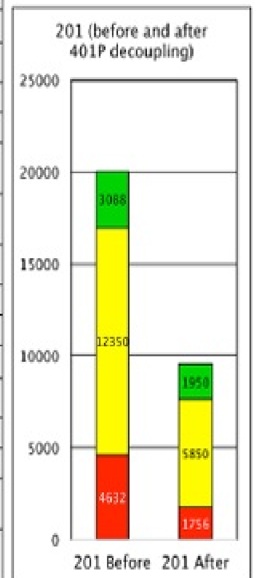

FIGURE 7-32 FPC buffer comparison (before and after buffering 401P)

Buffer Worksheet (401P)			
Average Daily Usage	650	Green Zone	**1950**
Buffer Profile	P, L (.3), L (.25)		LT Factor: 1950 (6500 x .3)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	10 days	Yellow Zone	**6500** (10(DLT) x 650(ADU))
		Red Zone	**2438** (1950 + 488)
			Base: 1950 (6500 x .3)
			Safety: 488 (1950 x .25)

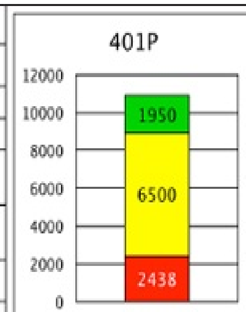

FIGURE 7-33 The buffer at 401P

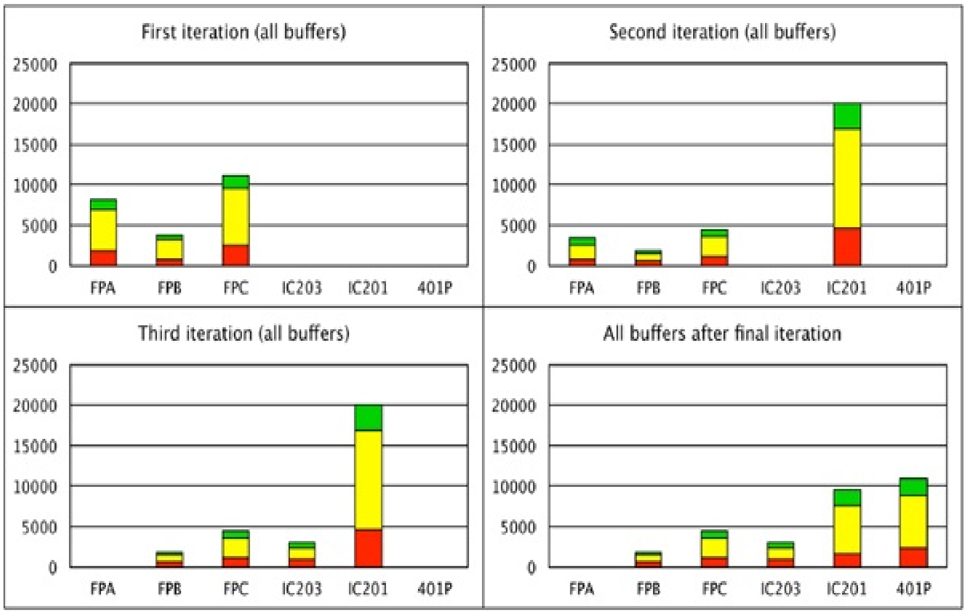

FIGURE 7-34 All buffers from the Company ABC example

In Chapter 9 the Company ABC example is revisited in order to cover the working capital impact with regard to average on-hand levels for the different iterations.

Calculating Replenished Override Buffers

Replenished override parts still utilize the three-zone system of green, yellow, and red. The calculations of these zones, however, are overridden with a modified equation or are essentially user defined depending on the limitations. These limitations would be beyond the limitations already discussed such as minimum order quantity or order cycle. An example might be spaces in a vending machine; there is a finite number of slots in each machine to be utilized. Another example might be cash or space limitations where a company is contractually bound. The zones still serve the same purpose but due to the imposed limitations may have limited effectiveness in their respective capacities. Appendix C, which describes how DDMRP can be applied to the retail environment, provides an example of buffer override logic in order to deal with certain restrictions or limitations.

Calculating Min-Max Buffers

Min-max parts utilize only two zones—green and red. These zones are calculated and adjusted in the same way as replenished parts; there is simply no yellow zone. Figure 7-35 shows the previous example part with a min-max designation. The same green and red zone calculations occur, and the yellow zone is disregarded.

Example Part Buffer Calculation			
Average Daily Usage	10	Green Zone	70
Buffer Profile	M, M (.5), L (.33)		LT Factor: 60 (DLT (12)x ADU (10) x Lead Time Factor (.5))
MOQ	50		Minimum Order Quantity: 50
Imposed or Desired Order Cycle (DOC)	7 days		Order Cycle: 70 (7(OC) x 10(ADU))
Decoupled Lead Time (DLT)	12 days	Yellow Zone	Not Applicable
		Red Zone	80 (Red Base (60) + Red Safety (20))
			Red Base: 60 (DLT (12)x ADU (10) x Lead Time Factor (.5))
			Red Safety: 20 (Red Base (60) x Variability Factor (.33))

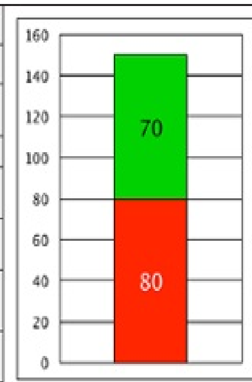

FIGURE 7-35 Example part with min-max designation

Summary

The calculations involved in setting buffer levels are based on quite simple equations. Most people above the age of 12 should be able to perform them with relative ease. The secret of sizing buffers has less to do with formal equations and more to do with the considerations and organization of applying those equations. This requires a process of parameter setting and maintenance at both a global (buffer profile) and individual part level. This process will most likely impact an organization’s Sales and Operations Planning (S&OP) processes and the tools and databases used in that environment. The impact to S&OP will be discussed in depth in Chapter 13 through the introduction and description of Demand Driven S&OP.

CHAPTER 8

Buffer Adjustments

Chapter 7 discussed the considerations in calculating the initial levels of a buffer. Since today's supply chains are incredibly dynamic, these buffers must adjust and adapt to changing conditions. By understanding the equations to set the buffer zones, then the factors that can change a part's buffer over the course of time are also understood. These changes can come from part attribute changes or buffer profile changes.

Recalculated Adjustments

Recalculated adjustments are automated adjustments to buffer levels based on changes to individual part attributes or buffer profile adjustments.

As discussed in Chapter 7, there are three critical factors for all buffered parts that directly impact the buffer equations: ADU, lead time, and minimum order quantity. ADU and lead time tend to have the most dramatic impact because they are involved in all three zone determinations. The minimum order quantity is only involved in green zone determination. The most dynamic part attribute is the ADU, as it is consistently being recalculated and updated.

As an example of each input change, Figure 8-1 shows the buffer inputs over a six-month time frame. Note that every input is static except ADU. The part's decoupled lead time (DLT) and buffer profile inputs [lead time factors (LTF) and variability factors (VF)] do not change. This part seems to experience significant growth over this six-month time frame as the ADU moves from 10 on January 1 to 53 by June 15. At this point, part 1234 does not have a minimum order quantity or order cycle assigned to it. The green zone is calculated using the lead time factor of 0.5 (a medium lead time part).

Figure 8-2 shows the buffer zones adjusting over the course of the six-month period. As the ADU quintuples over that period, the buffer adjusts up accordingly. The rate of this adjustment directly corresponds to the rise in ADU; all other variables remain static. The ADU is represented by the solid line. Its value corresponds to the right hand Y-axis while the buffer zone values relate to the left hand y-axis.

Date	Red	Yellow	Green	ADU	Red Base	Red Safety	DLT	LTF	VF
1-Jan	70	100	50	10	50	20	10	0.5	0.4
15-Jan	84	150	75	15	75	30	10	0.5	0.4
1-Feb	128.8	230	115	23	115	46	10	0.5	0.4
15-Feb	212.8	380	190	38	190	76	10	0.5	0.4
1-Mar	252	450	225	45	225	90	10	0.5	0.4
15-Mar	291.2	520	260	52	260	104	10	0.5	0.4
1-Apr	308	550	275	55	275	110	10	0.5	0.4
15-Apr	324.8	580	290	58	290	116	10	0.5	0.4
1-May	302.4	540	270	54	270	108	10	0.5	0.4
15-May	313.6	560	280	56	280	112	10	0.5	0.4
1-Jun	324.8	580	290	58	290	116	10	0.5	0.4
15-Jun	296.8	530	265	53	265	106	10	0.5	0.4

FIGURE 8-1 Part 1234 data

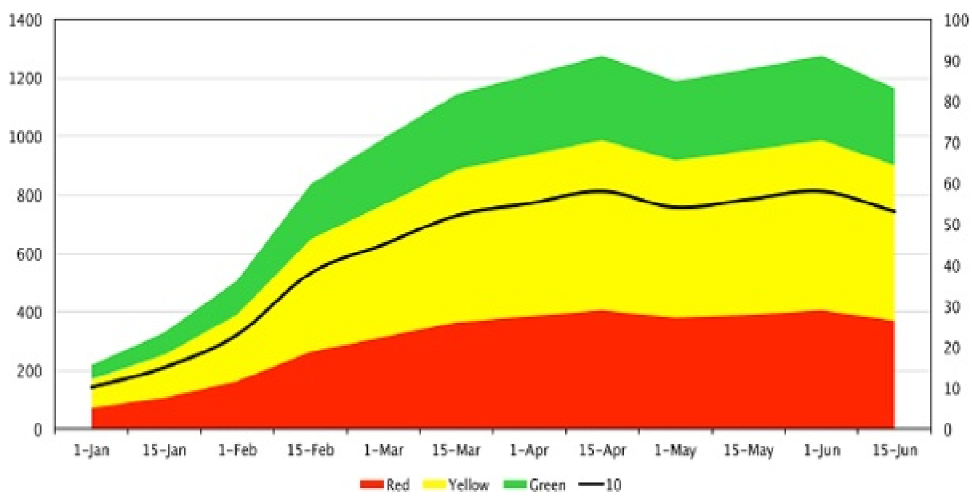

FIGURE 8-2 Part 1234 buffer adjustment over six months

In a DDMRP system, ADU will always be changing since it is frequently recalculated. Within some time periods the level of change could be relatively small, but it is changing nonetheless. What happens when more static but not less significant part attributes change? For example, the part's DLT directly impacts all zones if there is no minimum order quantity present. This is the case so far with part 1234. Figure 8-3 shows the change in the part's DLT on March 15.

The shaded boxes represent the notable changes. First, the DLT has been compressed from 10 to 5 days. If the part is purchased, this could occur by sourcing from a different vendor or a current vendor agreeing to a better lead time now that the company is no longer constantly rescheduling. If the part is manufactured, this compression could be due to additional capacity, process improvement, reduction of work in progress, or additional decoupling (as illustrated with the Company ABC example in Chapters 6 and 7). This shorter lead time has moved the part from a medium lead time factor of 0.5 to a short lead time factor of 0.7. This has a direct impact on the buffer. While the green and yellow zones shrink due to this lead time compression, the red zone actually inflates due to the higher LTF. Figure 8-4 illustrates the impact on the buffer zones of compressing the lead time from 10 to 5 days on March 15 but with the rest of the inputs staying the same.

Date	Red	Yellow	Green	ADU	Red Base	Red Safety	DLT	LTF	VF	MOQ	Green (LTF)
1-Jan	70	100	50	10	50	20	10	0.5	0.4	0	50
15-Jan	84	150	75	15	75	30	10	0.5	0.4	0	75
1-Feb	128.8	230	115	23	115	46	10	0.5	0.4	0	115
15-Feb	212.8	380	190	38	190	76	10	0.5	0.4	0	190
1-Mar	252	450	225	45	225	90	10	0.5	0.4	0	225
15-Mar	291.2	260	182	52	182	72.8	5	0.7	0.4	0	182
1-Apr	308	275	192.5	55	192.5	77	5	0.7	0.4	0	192.5
15-Apr	324.8	290	203	58	203	81.2	5	0.7	0.4	0	203
1-May	302.4	270	189	54	189	75.6	5	0.7	0.4	0	189
15-May	313.6	280	196	56	196	78.4	5	0.7	0.4	0	196
1-Jun	324.8	290	203	58	203	81.2	5	0.7	0.4	0	203
15-Jun	296.8	265	185.5	53	185.5	74.2	5	0.7	0.4	0	185.5

FIGURE 8-3 Lead time change for part 1234

Finally, the part 1234 example is used to illustrate the impact of a significant MOQ on the buffer. In this case a minimum order quantity was imposed on part 1234 on April 15. If this is a purchased part, this could occur due to a new agreement with the supplier (maybe in exchange for the shorter lead time). If this is a manufactured part, this could occur in an attempt to save setups on a capacity constrained resource. The lead time compression from 10 to 5 days on March 15 has been maintained. Figure 8-5 illustrates the imposition of the MOQ. The shaded boxes highlight the relevant factors in this change. An MOQ of 400 has become the green zone value. The column labeled “Green (LTF)” is the green zone value calculated using the lead time factor.

Figure 8-6 illustrates the impact of imposing the MOQ on part 1234 on April 15. There is an immediate jump in the green zone. This will have significant impacts on the amount of working capital contained in the buffer. Working capital implications will be further explored in Chapter 9.

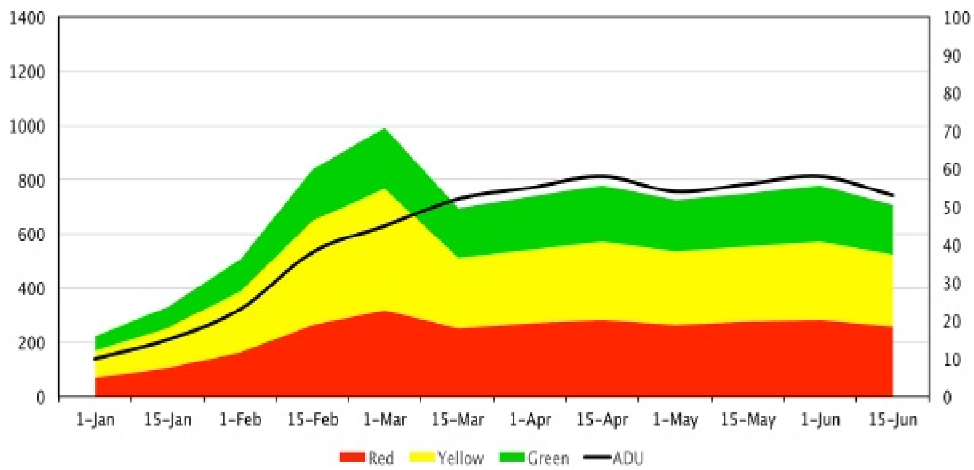

FIGURE 8-4 Part 1234 lead time compression

Date	Red	Yellow	Green	ADU	Red Base	Red Safety	DLT	LTF	VF	MOQ	Green (LTF)
1-Jan	70	100	50	10	50	20	10	0.5	0.4	0	50
15-Jan	84	150	75	15	75	30	10	0.5	0.4	0	75
1-Feb	128.8	230	115	23	115	46	10	0.5	0.4	0	115
15-Feb	212.8	380	190	38	190	76	10	0.5	0.4	0	190
1-Mar	252	450	225	45	225	90	10	0.5	0.4	0	225
15-Mar	353.6	260	182	52	130	91	5	0.5	0.7	0	182
1-Apr	374	275	192.5	55	137.5	96.25	5	0.5	0.7	0	192.5
15-Apr	394.4	290	400	58	145	101.5	5	0.5	0.7	400	203
1-May	367.2	270	400	54	135	94.5	5	0.5	0.7	400	189
15-May	380.8	280	400	56	140	98	5	0.5	0.7	400	196
1-Jun	394.4	290	400	58	145	101.5	5	0.5	0.7	400	203
15-Jun	360.4	265	400	53	132.5	92.75	5	0.5	0.7	400	185.5

FIGURE 8-5 Part 1234 with a minimum order quantity of 400

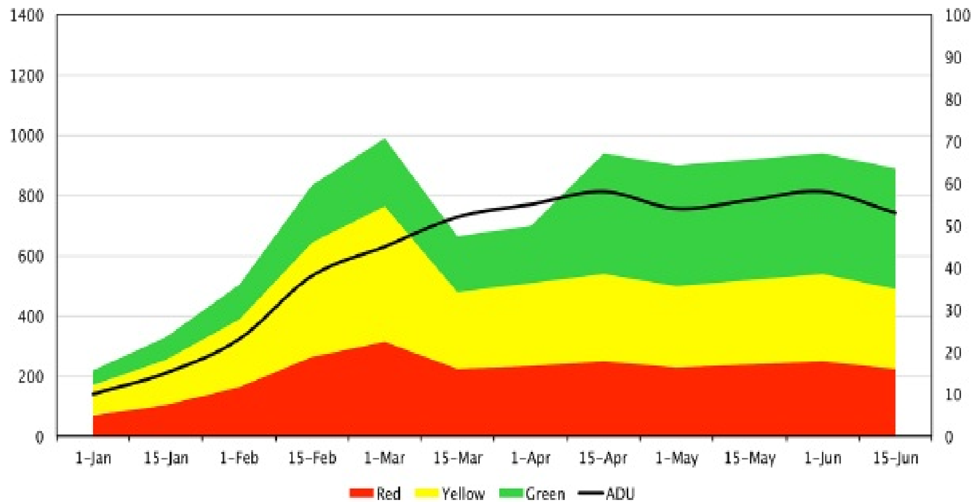

FIGURE 8-6 Part 1234 with a minimum order quantity

A buffer profile change causes a recalculation of the buffers to all parts assigned to that profile. The part 1234 example has already demonstrated the impact of moving a part to a different profile (from medium lead time to short lead time) and applying the different factors associated with that profile. From a global perspective, if there are changes to the factors within a buffer profile, then obviously all parts contained in that profile will be affected simultaneously. For example, if the planning team decides to use 0.6 instead of 0.5 for the lead time factor for medium lead time parts, then all parts with a medium lead time will be affected unless subcategories are created for medium lead time parts.

Planned Adjustment Factors

Buffers also can be manipulated through planned adjustments. Planned adjustments are based on certain strategic, historical, and business intelligence factors. These planned adjustments are manipulations of the buffer equation that affect inventory positions by raising or lowering buffer levels and their corresponding zones at certain points in time. These manipulations tend to be confined to demand input manipulations, zonal manipulations, or lead time manipulations.

Demand Adjustment Factor

The demand adjustment factor (DAF) is a manipulation of the ADU input within a specific time period. This manipulation occurs by adjusting the ADU to a historically proven or planned position based on an approved business case or as a reaction to rapid changes in demand within short periods of time.

Demand adjustment factors should not be indiscriminately used. Do not underestimate the power and flexibility of a properly managed DDMRP system. The buffers are robust. They are designed to absorb variability. The more responsive the plant or supplying resource is, the more robust the buffers' performance will be for higher variability. The higher the variability factor built into the buffers, the more robust the buffers' performance will be for higher variability, albeit with the penalty of additional inventory. The longer the horizon to see spikes, the more robust the buffers' performance will be for higher variability.

With that said, the buffers are only designed to absorb variability to a certain extent. It can often be the case that variability up or down can threaten the effectiveness of the buffer to protect that decoupling point. Demand adjustment factors should be employed when variability threatens to overwhelm the buffers. In this case the adjustment should be up. When variability will cause large amounts of prolonged excess inventory, the adjustment should be down.

How do demand adjustment factors work? Figure 8-7 depicts a sample part called ABC. The buffer-level zones appear to be relatively stable over a significant period of time. These levels, however, can easily be manipulated by changing the ADU value feeding the buffer-level equation at certain points.

Figure 8-8 illustrates part ABC with a DAF implemented from week 13 to week 25. The table below the graph shows the original ADU, the DAF applied to each period, and the adjusted ADU value. For example, in week 18 an adjustment factor of 1.8 has been applied to the original ADU of 34.68 to produce an adjusted ADU of 62.43. Perhaps this is done with the knowledge of a major promotion or seasonality. Any demand adjustment above 1 produces an inflationary effect on the buffer.

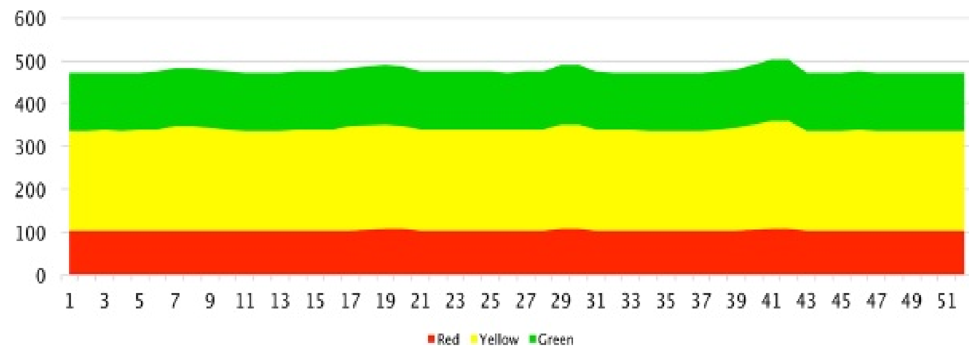

FIGURE 8-7 Part ABC with no demand adjustment

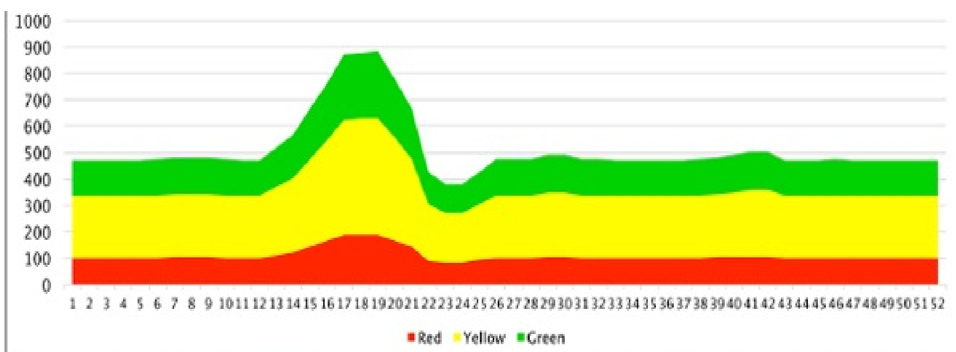

Week	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26
Original ADU	33.59	33.59	33.72	33.89	33.91	34.41	34.68	34.96	34.68	33.86	33.72	33.80	33.78	33.86	33.67
DAF	1.00	1.10	1.20	1.40	1.60	1.80	1.80	1.80	1.60	1.40	0.90	0.80	0.80	0.90	1.00
Adjusted ADU	33.59	36.94	40.47	47.44	54.26	61.94	62.43	62.93	55.49	47.40	30.35	27.04	27.02	30.47	33.67

FIGURE 8-8 Part ABC with demand adjustment factor

In weeks 22–25 a factor less than 1 is applied. This deflates the ADU value feeding the buffer-level equations, causing a trough to appear for a small period of time. Perhaps this was done with the thinking that the market will be saturated and demand will drop off temporarily. This is a “de-promotion” effect. This could also be a seasonal effect. The application of a demand adjustment factor does not change the selection criteria for green zone sizing. The green zone is still determined by the larger value of the MOQ, order cycle, or application of the lead time after the adjusted ADU has been applied.

Demand adjustments are used for common situations such as rapid changes in demand, product transition, and seasonality. These planned adjustments separate what is known, the pattern of the demand, from what is unknown, the level of actual demand that will be experienced.

Rapid Buffer Adjustment

A demand adjustment factor can be employed to immediately raise or lower buffer levels if there is a rapid change in demand that is indicative of a coming trend. This rapid change would be indicated by an ADU alert (discussed in Chapter 7).

An ADU alert does not directly generate a demand adjustment factor. The alert should raise questions that can be considered at the Sales and Operations Planning meeting. Why is the item experiencing significantly heavier demand within the recent past? Did we add a significant new customer or open a new territory? Did we run a promotion that marketing forgot to tell operations about it? Is there an impending perception in the market fueling the demand that we should capitalize on? Conversely, why is the item experiencing significantly lighter demand within the recent past? Was there a natural disaster that caused a major supply chain disruption? Is there negative feedback on social media affecting the product? The answers to these questions could lead to the application of a demand adjustment factor for a period of time until the ADU calculation normalizes to the new level of demand.

Product Introduction, Deletion, and Transition

Demand adjustment factors will often be used with product introductions, deletions, and various forms of transitions.

Product Introduction. When introducing a new product that will be strategically buffered, a company has to establish a buffer position. The demand adjustment factors can be used with regard to a new-product launch. The product will be offered to the market in week 4. The sales and marketing plan for this new product calls for the product to be selling 2,000 units per day 12 weeks from now. That 2,000 will become the baseline for the application of the DAF. Instead of spending the cash and capacity to bring the buffer to full size immediately, a DAF will be applied starting in week 2 and continue to week 11. Figure 8-9 is the demand adjustment schema for this new item.

This planned adjustment will ramp up the planned ADU over a period of time, thus creating a buffer that also grows over time. Additionally, a buffer will be present on the launch date in week 4 to handle initial demand. The demand adjustment factor can be adjusted based on the real performance against that plan if an ADU alert is triggered. Figure 8-10 depicts the buffer zones ramping up.

Week	Projected ADU	DAF	Adjusted ADU
1	2,000	0	0
2	2,000	0.1	200
3	2,000	0.2	400
4	2,000	0.3	600
5	2,000	0.4	800
6	2,000	0.5	1,000
7	2,000	0.6	1,200
8	2,000	0.7	1,400
9	2,000	0.8	1,600
10	2,000	0.9	1,800
11	2,000	1	2,000
12	2,000	1	2,000

FIGURE 8-9 Ramp-up demand adjustment schema

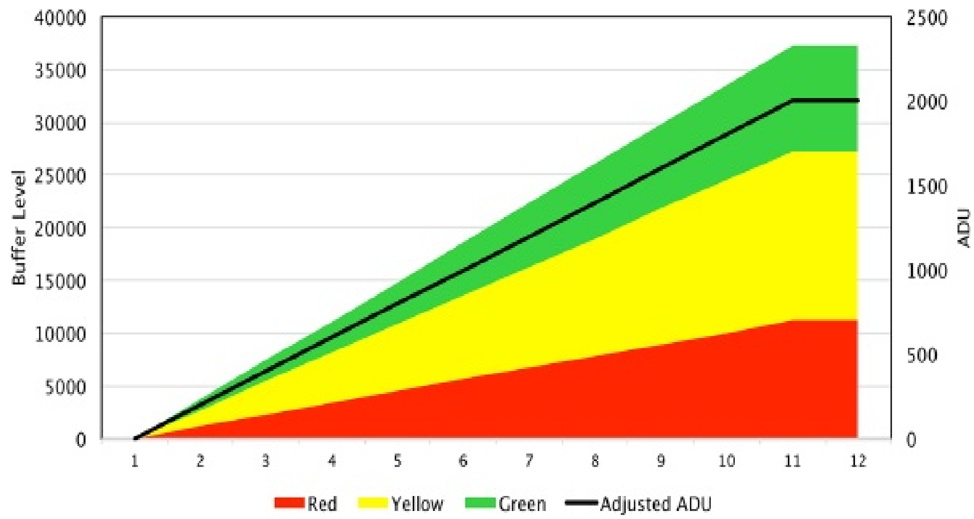

FIGURE 8-10 Buffer ramp-up for the new product

Product Deletion. Next is the application of the demand adjustment factor in bringing a buffer down in anticipation of the product no longer being offered to the market. Figure 8-11 depicts the ramp-down for this example part. The ramp-down uses a DAF below 1 that decreases over a time period until it reaches zero. In this example the application of the DAF begins in week 5 and has a zero value by week 9 when the product is planned to be discontinued.

The ADU is ramped down to create a gradually diminishing buffer level and zone definition. Figure 8-12 illustrates the impact the ramp schema has on the product's buffer.

Week	ADU	DAF	Adjusted ADU
1	300	1	300
2	300	1	300
3	300	1	300
4	300	1	300
5	300	0.8	240
6	300	0.6	180
7	300	0.4	120
8	300	0.2	60
9	300	0	0
10	300	0	0
11	300	0	0
12	300	0	0

FIGURE 8-11 Ramp-down demand adjustment schema

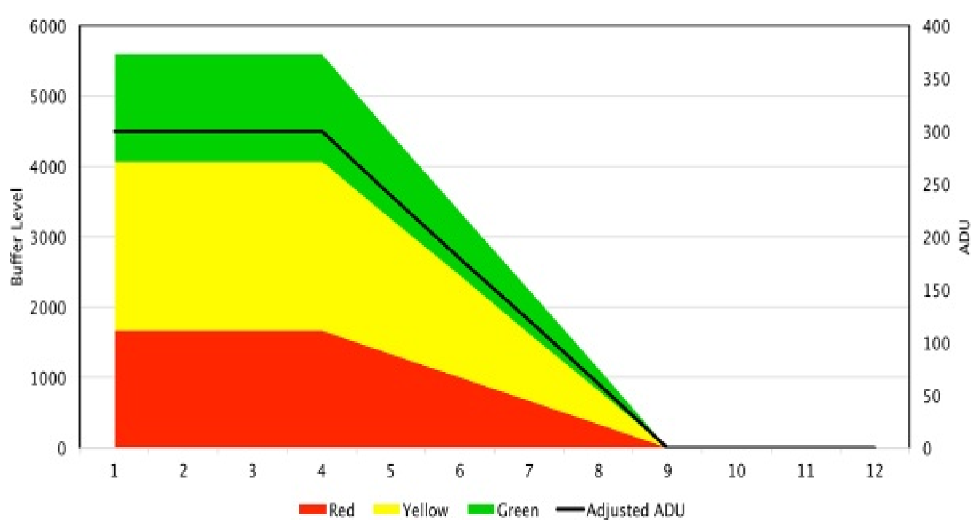

FIGURE 8-12 Buffer ramp-down for the discontinued product

Product Transition. Two types of transition may require the use of the demand adjustment factor. One involves the replacement of an older version of a product with a newer version. In most cases the ADU of the older product can be used as the ADU for the new product. The new product is a replacement and will most likely have the same customer base. Any uptick in demand due to new features, for example, may be able to be actualized through the ADU calculation period once the new product is introduced. While the ADU for both will be the same, be careful to not "double-buffer." Additionally, the risk of obsolete stock should be minimized.

In order to manage this transition, both a ramp-up and a ramp-down are employed simultaneously. The new product will be ramped up, while the old product will be ramped down. In both cases the effectivity date of both items will be used as a culmination point for the ramp-up and ramp-down schemas, which means that the buffers will overlap. Figure 8-13 shows the demand adjustment schema for both products over a 12-week transition period. The new version will be offered for sale on week 9, and the old version will no longer be available on week 12. Thus there is a four-week overlap where both are offered (a soft transition).

Figure 8-14 shows the buffer implications for both versions of the product, the old ramping down while the new is ramping up. By week 11 the new product is at full buffer strength in time for its effectivity date. The adjusted ADU (AADU) is shown for both the new (N) product version and old (O) product version as well as the buffer zone values. The new product is in a stacked-bar format, while the old is in a stacked-line format for contrast.

Week	ADU	DAF (Old)	Adjusted ADU (Old)	DAF (New)	Adjusted ADU (New)
1	30	1	30	0	0
2	30	1	30	0	0
3	30	0.9	27	0	0
4	30	0.8	24	0	0
5	30	0.7	21	0.2	6
6	30	0.6	18	0.4	12
7	30	0.5	15	0.6	18
8	30	0.4	12	0.8	24
9	30	0.3	9	1	30
10	30	0.2	6	1	30
11	30	0.1	3	1	30
12	30	0	0	1	30

FIGURE 8-13 Demand Adjustment Schema over 12 weeks

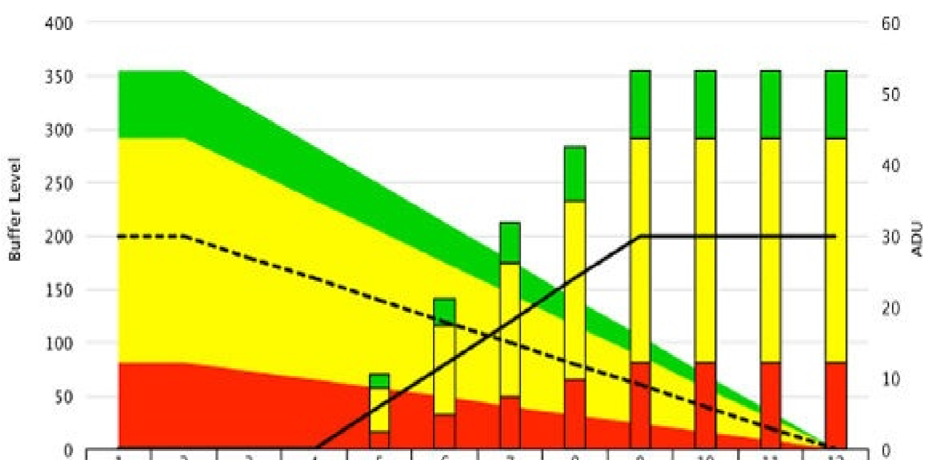

	1	2	3	4	5	6	7	8	9	10	11	12
Green (O)	63	63	56.7	50.4	44.1	37.8	31.5	25.2	18.9	12.6	6.3	0
Yellow (O)	210	210	189	168	147	126	105	84	63	42	21	0
Red (O)	81.9	81.9	73.71	65.52	57.33	49.14	40.95	32.76	24.57	16.38	8.19	0
Green (N)	0	0	0	0	12.6	25.2	37.8	50.4	63	63	63	63
Yellow (N)	0	0	0	0	42	84	126	168	210	210	210	210
Red (N)	0	0	0	0	16.38	32.76	49.14	65.52	81.9	81.9	81.9	81.9
AADU (N)	0	0	0	0	6	12	18	24	30	30	30	30
AADU (O)	30	30	27	24	21	18	15	12	9	6	3	0

FIGURE 8-14 Demand adjustment for the replacement product

A company that utilizes this strategy does not commit resources earlier than necessary. The key is that this example has planned the ramp-up curve to be at

100 percent of ADU by the time the old part is no longer active while at the same time having an amount of new inventory in advance of that cutover date. This strategy will minimize or eliminate obsolete inventory while allowing for a seamless transition to a new part from the market's perspective. This will reduce or eliminate the risks of missed sales due to shortages that tend to occur through poorly managed transitions or variability of demand during that transition.

Another form of product transition involves a known and dramatic impending shift in its demand up or down. This requires ramping up or down as previously described. Figure 8-15 illustrates an example of an expected and significant demand uptick. The row titled "Week" represents the present and future in weekly buckets. Week 1 is the current week. The row titled "Sales" presents the actual sales for week 1 and the projected sales for weeks 2–17. The row titled "ADU (1 Week)" is the ADU value within each weekly bucket, while the row titled "ADU (12 Week)" is the ADU calculated based on a 12-week past period. The row titled "ADU Differential" shows the quotient between the ADU within the week and the ADU based on the last 12 weeks.

In week 6 (five weeks from the current time period) a significant and prolonged upsurge is predicted to occur. Demand surges from a predicted 80 per week to 400 and then to a steady 800 per week for the foreseeable future. Perhaps a new strategic customer has signed a contract calling for large weekly volumes. The ADU differential goes from 1 to 4.8 in the span of only two weeks. Figure 8-16 shows the difference in projected weekly demand ("Demand"), projected ADU within each week ("ADU 1 Week"), and the ADU based on the past 12 weeks ("ADU 12 Week"). It will take until week 17 for the ADU to normalize to the higher level of demand (the ADU differential to 1.0).

Week	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Sales	81	80	80	80	80	400	800	800	800	800	800	800	800	800	800	800	800
ADU (1 Week)	11.6	9.3	10.0	11.4	11.4	57.1	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3
ADU (12 Week)	12.9	12.1	11.5	11.4	11.4	15.2	23.6	32.0	40.5	49.1	57.7	66.4	74.9	83.7	92.4	101.0	109.5
ADU Differential	0.9	0.8	0.9	1.0	1.0	3.8	4.8	3.6	2.8	2.3	2.0	1.7	1.5	1.4	1.2	1.1	1.0

FIGURE 8-15 A large expected upsurge in demand

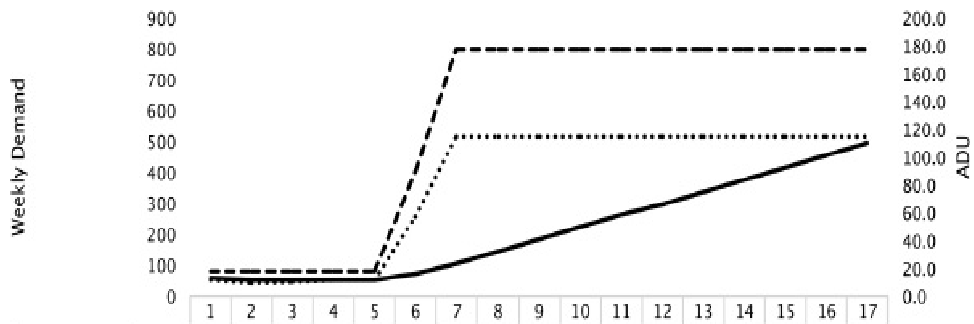

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Demand	81	80	80	80	80	400	800	800	800	800	800	800	800	800	800	800	800
ADU (12 Week)	12.9	12.3	11.8	11.7	11.7	15.5	23.9	32.3	40.8	49.4	58.0	66.7	75.2	83.8	92.4	101.0	109.5
ADU (1 Week)	11.6	9.3	10.0	11.4	11.4	57.1	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3

Demand ADU (12 Week) ADU (1 Week)

FIGURE 8-16 Understated ADU for a prolonged period

This means that the buffers will be dramatically understated for a prolonged period of time. The new agreement will not be able to be supported under these conditions. Figure 8-17 shows the projected lag in the buffer adjusting to the expected upsurge.

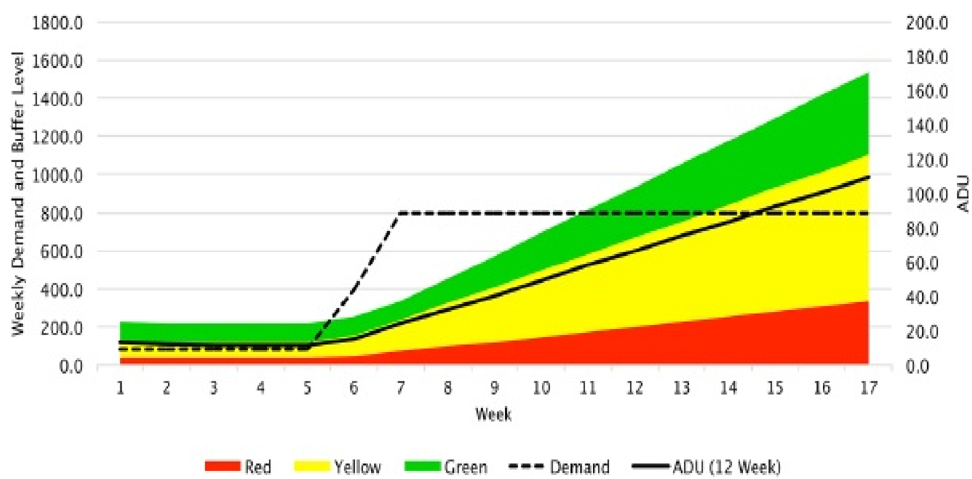

FIGURE 8-17 Lagging buffer adjustment

In this case a demand adjustment factor can be used to flex the buffer up in advance of the impending uptick. The planner will attempt to employ factors starting in week 2 that will ramp the buffers up to the expected required level. By week 7, demand adjustment factors have brought the ADU differential to 1. Once the ramp has been accomplished, the planner will then cut over the buffer calculation to the projected ADU of 114.3 and reduce the DAF back to 1. This cutover actually happens in week 8. Figure 8-18 shows the demand adjustment factors employed from weeks 2 to 7 in order to accomplish the ramp-up.

The new ADU will then be actualized over the horizon. Is this forecasting? Yes, *but* it is tied to a strategic decision and a specific customer agreement. Figure 8-19 shows the buffer reaction to the demand adjustment factor. The assumption is that the necessary resources will be available to accomplish this ramp-up. The buffer pattern provides a visible and realistic picture of what must be accomplished to support this new strategic decision.

Note that these examples are showing the value of being able to simulate the impact on buffers based on changes in part and profile attributes.

Seasonality

Another application of demand adjustment factors occurs with regard to seasonality. Many companies have products with seasonal uplifts and troughs that may pose challenges to DDMRP buffers if not properly addressed. Tackling seasonality will involve both ramp-up and ramp-down adjustments. This was demonstrated in Figure 8-8. But several interactive considerations should be taken into account when considering when and to what extent to apply demand adjustment factors to compensate for seasonality of strategically buffered items:

- **Consideration 1. Severity of the seasonality (length and significance).** The first consideration is the known length and severity of the seasonal swing. The higher the change and the shorter the window, the more severe the seasonality.
- **Consideration 2. Length of the ADU calculation period.** The length of the ADU calculation period must be known and considered in relation to the severity of the season. The longer the ADU calculation period and the more severe the season, the higher the likelihood that the ADU will be dramatically understated during the initial seasonal period.
- **Consideration 3. Past, forward, or blended ADU.** In many cases, a forward-looking ADU negates the need for a planned adjustment factor, as it anticipates the seasonal demand change. A blended ADU approach may be reactive enough depending on the severity of the season and the length of the calculation period. In most cases a past-looking ADU will leave the buffers vulnerable to significant seasonality, and demand adjustment factors will be needed to

compensate.

- **Consideration 4. Lead times of critical components.** Long lead time items (including parent items with long lead time components) become a factor in determining when demand adjustment factors should be applied. Employing a demand adjustment factor too late will mean that despite the buffers being properly sized for the demand change, they will be subject to a lag in supply that is equivalent to the supply lead time. For example, if a demand adjustment factor begins to be applied 6 weeks in advance of the seasonal period, but the lead time of the item or lower-level components is 12 weeks, then the buffer will be essentially undersupplied while waiting for the long lead time item to arrive. This could result in shortages throughout the seasonal period as well as overages coming out of that period. Thus the application of the demand adjustment factor must take lead time into account. In DDMRP this is called a “supply offset.” Even if a forward-looking ADU is employed and the lead time is longer than the ADU period calculation, a supply offset may need to be considered.

Week	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Sales/Week	81	80	80	80	80	400	800	800	800	800	800	800	800	800	800	800	800
ADU (1 Week)	11.6	11.4	11.4	11.4	11.4	57.1	112.4	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3
ADU (12 Week)	12.9	12.3	11.8	11.7	11.7	15.5	23.9	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3	114.3
ADU Differential	0.9	0.9	1.0	1.0	1.0	3.7	4.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
DAF	1.0	2.5	3.5	4.5	6.0	6.0	4.7	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0

FIGURE 8-18 Required demand adjustment factors

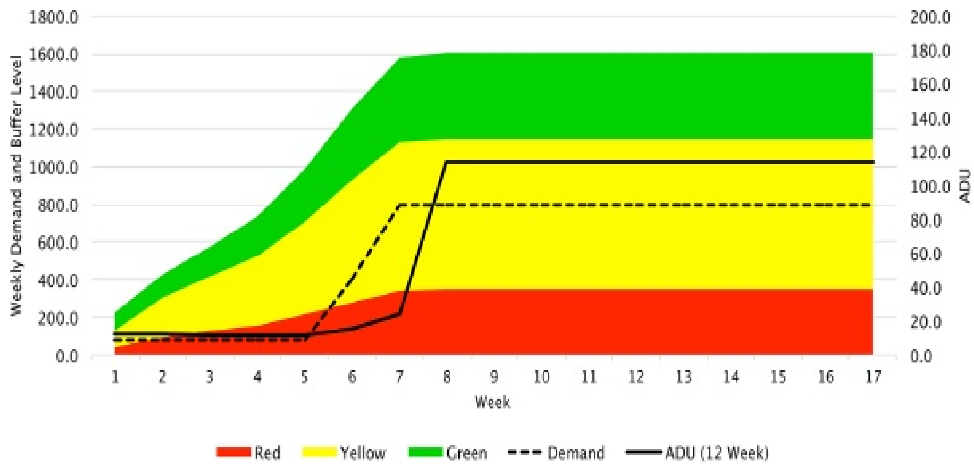

FIGURE 8-19 The buffer with demand adjustment factors applied

If a parent item has a buffered long lead time component, then the cumulative lead time (as opposed to the decoupled lead time) of the parent should be used for the offset—unless the component has a sufficient buffer position (typically a shared component) to absorb the seasonal demand of the parent item without considering the component lead time.

- **Consideration 5. Resource capacity.** Finally, determining when demand adjustment factors should be employed may require the consideration of resource capacity. If the seasonal severity or seasonal ramp-up schema outstrips the supplying capacity, then the application of the demand adjustment factor must be pulled forward in time.

Chicken Truffle Soup Seasonal Example. These considerations are illustrated in a fictitious company called SoupCo. SoupCo makes a type of soup called Chicken Truffle Soup. It is a specialty, low-volume product that has highly seasonal demand. Sales of the soup are about 16,000 cases per year, implying a year-round average of around 307 cases per week. But over 74 percent (11,800 of 15,875 cases) of the yearly Chicken Truffle sales occur within a 12-week window. Figure 8-20 shows the seasonal demand profile of Chicken Truffle Soup.

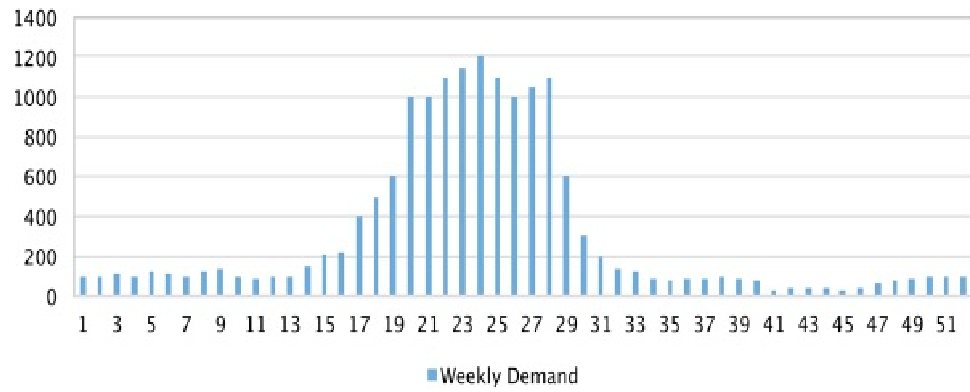

FIGURE 8-20 Chicken Truffle Soup demand profile

SoupCo uses a 12-week past-looking ADU calculation. Figure 8-21 illustrates the ADU calculation using a past-looking 12-week horizon. The past-looking ADU results in a significant lag that leaves the buffer unprepared for the severity of the demand uplift. This means stockouts. Additionally, it will leave a large amount of excess stock past the peak demand period. This potentially could mean stock that would go past the expiration date. The buffer simply will not have the product available in time for the season. Then when it becomes available, the large part of the season will be done.

The lag of a past-looking ADU is further exacerbated by a long lead time associated with the crucial component of the soup—truffles. To make and package Chicken Truffle Soup only takes a week, but the lead time for the truffles is six weeks. Even if the truffles are buffered, that buffer will be dramatically understated based on a lagging ADU. This means that SoupCo must offset the demand adjustment factor by at least six weeks in order to ensure a smooth flow of ingredient supply to support the season. Since the first significant uplift occurs at week 17, this implies that the DAF will need to begin to be applied by at least week 11.

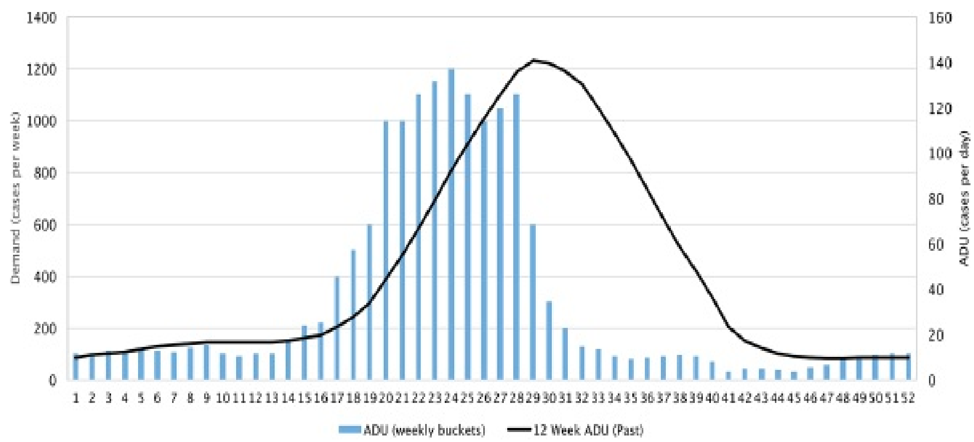

FIGURE 8-21 Chicken Truffle Soup ADU using the past 12-week horizon

Finally, SoupCo has a packaging line capacity constraint to consider. Chicken Truffle Soup is packed on a dedicated specialty line that has limited capacity. This line can only pack 800 cases per week when operating at full capacity. For a period of nine weeks (weeks 20–28) the market will be consuming more than the packaging line can pack. Over that nine-week period the packaging line can only pack 7,200 cases, and yet demand calls for 8,600 cases. That 1,400-case difference represents nearly two weeks of packaging line capacity. Figure 8-22 shows the capacity of the packaging line indicated by the dotted line corresponding to the left Y axis.

This is also a consideration in determining when to start the DAF application. It should be at least week 11 in order to offset the truffle lead time. From week 11 to week 20 (the week in which demand outstrips capacity), there is a total demand of 2,370 cases. Within that same time frame, available capacity is 5,600 cases. The difference is 3,230 cases of available capacity. That means that the packaging line has the capability to absorb the capacity shortage of weeks 20–28 within weeks 8–19 if necessary. An alternative could be to apply smaller demand adjustment factors prior to week 11 in order to create a less steep seasonal ramp-up.

Promotional Campaigns

Still another application for the employment of demand adjustment factors is to compensate for large promotional campaigns. These promotional campaigns are planned to induce demand surges. They can often last one to three months and can be a significant marketing investment. Large promotional campaigns that create large amounts of demand uplift only to result in stockouts are normally viewed as disasters. Thus it is crucial that companies adjust their buffers in advance of these planned campaigns. The same factors are considered in promotional campaigns as for seasonality. Since a company typically has experience with these large campaigns, there is an expectation of the amount of surge expected.

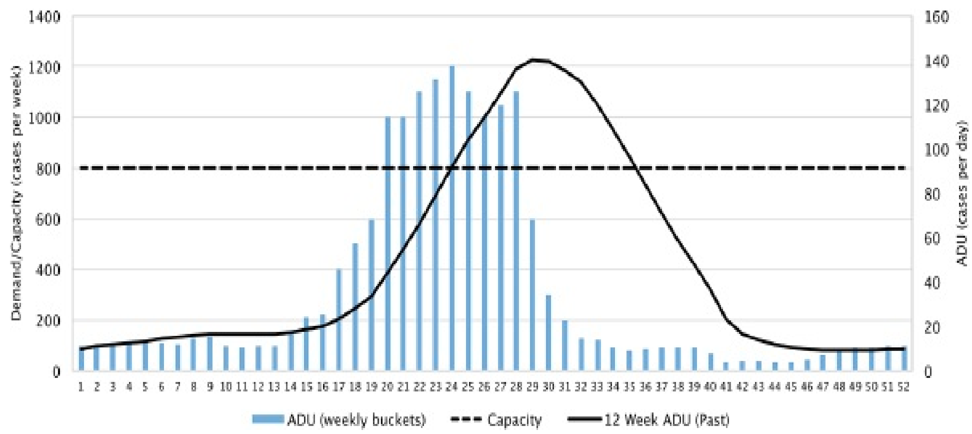

FIGURE 8-22 Considering capacity of the packaging line

Applying DAF to Components

If a demand adjustment is to be applied to a parent item, it may be necessary to apply this factor to the component buffers to ensure a properly supplied parent buffer. If the component is unique to the parent, then the same factor applied to the parent can be applied to the component. If the component is a long lead time item, then its lead time should be considered in determining when to apply the DAF (supply offset).

If the component is a shared item, then the demand adjustments (or lack thereof) of all parents must be considered in determining the amount by which to adjust the component's buffer. Figure 8-23 shows an example with two parent items (FPA and FPB) that have multiple shared components (ICB, PPB, and PPA). The numerical strip in the middle of the graphic corresponds to the level in the product structure in which the part occurs. If these two parent items have different DAFs applied to them, how does it affect the DAF that should be applied to the components?

A relevant factor in determining how the DAF will change at the component level is to account for the parent-to-component ratio of each part across all product structures. In Figure 8-23 that ratio is displayed in the black box to the right of each part name. For example, it takes two ICAs and four PPCs to make one SAA. When a DAF is applied to a parent item's ADU, an adjusted ADU is created. This adjusted ADU is then multiplied by the component ratio to get the component's adjusted ADU. In Figure 8-24 the ratio for each component in the structure is depicted in the columns "FPA Ratio" and "FPB Ratio." PPA is involved in both structures but has a different ratio with its direct parent in each respective product structure.

Figure 8-24 shows the adjusted ADUs for each component item and how that converts to a DAF for each component. In this example FPA has a DAF of 2 (200 percent), whereas FPB has no DAF applied (a DAF of 1). In every case of a common component, there is a blending that occurs based on the DAF of each parent and the ratio relative to each product structure. ICB's ADU has adjusted to 800, which is a DAF of 1.3 of its current ADU.

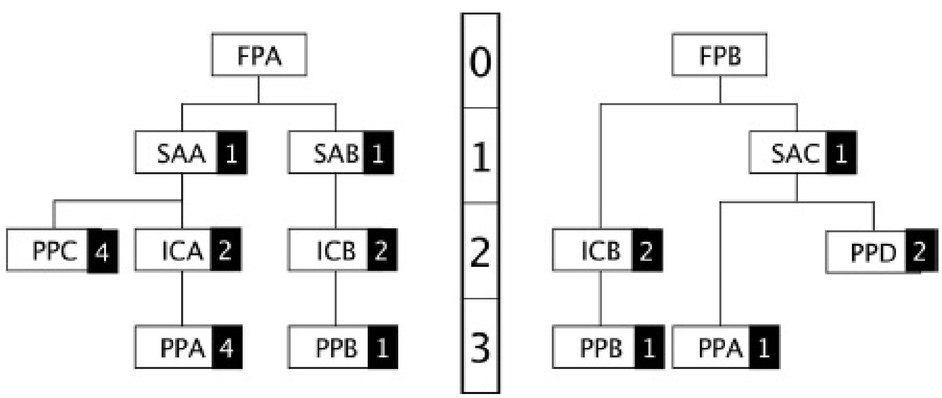

FIGURE 8-23 Product structure for FPB and FPA

Item #	BOM Level	ADU	Adjusted ADU	FPA Ratio	FPB Ratio	DAF
FPA	0	100	200			2
FPB	0	200	200			1
SAA	1	100	200	1		2.0
SAB	1	100	200	1		2.0
SAC	1	200	200		1	1.0
PPC	2	400	800	4		2.0
ICA	2	200	400	2		2.0
ICB	2	600	800	2	2	1.3
PPD	2	400	400		2	1.0
PPA	3	600	1,000	4	1	1.7
PPB	3	300	400	1	1	1.3

FIGURE 8-24 The DAF effect on FBA and FPB components

Zone Adjustment Factor

Another way to adjust part buffers is through zonal manipulations. This could apply to the individual part or a group of parts affected in the same way by the adjustment (not necessarily parts in the same buffer profile). A zone adjustment factor can be applied to any of the three zones of a buffer. The zones of the buffer serve different purposes, so the zone adjustment factor should apply to the appropriate zone based on the rationale for the adjustment.

Green Zone Adjustment

As discussed in Chapter 7, the green zone determines order size and frequency. Thus when order size and frequency need to be adjusted, the green zone can be manipulated up or down. Adjustments up often occur when a supplying resource with significant setup issues encounters a capacity constraint. The green zone of certain items can be raised in order to create less frequent and larger orders to save the capacity lost in additional setups. Note, this is not raising the green zone in order to lower unit cost; it is a deliberate strategy based on protecting the flow of relevant information and materials through a capacity constraint. Any capacity saved at that point directly translates to more total output for the system, and that protects all buffers and customers fed by that resource.

Downward adjustments in the green zone might occur when there is sufficient excess capacity, and market responsiveness is the key. This can be required in some seasonal markets in which a supplying resource, after completing the planned buffer buildup, must carefully watch market demand across products and quickly respond to replenishing items that are unique for this season's high movers. This means that some products may have different MOQs depending on the changing circumstances in relation to supplying capacity.

If the packaging line in SoupCo was not dedicated to only Chicken Truffle Soup, this strategy could be used to augment the use of a DAF. For example, if SoupCo has a whole line of seasonal truffle soups, larger batches of the fast movers could be run in the seasonal buildup in order to prevent erosion of packaging line capacity due to setups. Within the season, smaller batches could then be employed across the different varieties to respond to this particular season's pull.

Yellow Zone Adjustment

The yellow zone is the heart of demand coverage in the buffer. As noted in Chapter 7, the yellow zone is calculated as ADU × decoupled lead time. It is assuming a rate of demand within a response window. In that regard a yellow zone adjustment could be triggered by known or planned events that deal with either of those two components. Typically, a yellow zone adjustment occurs as a response to a planned short-term promotional event or a planned or known supply disruption.

A short-term promotional event is one that does not last long enough to employ a demand adjustment factor schema. It appears within a short window of time where a significant rate of demand change occurs but then demand patterns quickly turn to normal. This window of time is usually within the part's lead time. A factor can be applied to size the yellow zone to the expected demand within that time frame. A part with an ADU of 1,000 and a lead time of 7 days would have a yellow zone of 7,000. If a short-term promotional event lasting a week was expected to triple sales for the week, a factor of 3 could be applied to the yellow zone to size it from 7,000 to 21,000.

Another scenario for a yellow zone adjustment occurs when there is a known or planned interruption in supply. Maybe a critical supplier is located in a region that has an extended holiday, or perhaps there is a planned upgrade to a facility. These events translate to the source being unable to respond for a window of time. For example, if the SoupCo packing line was scheduled to be down for one week, a yellow zone adjustment factor of 2 would double the size of the yellow zone to cover the two weeks of demand. The buffer would then return to normal.

Red Zone Adjustment

The red zone is the embedded safety in the buffer. The red zone should be adjusted when there is a known or planned but temporary change in volatility that does not warrant moving the part(s) to another buffer profile. This is typically associated with a temporary change or transition with regard to purchasing, producing, or distributing an item or group of items. Perhaps a new resource or material is being brought on line that could cause more disruptions in the short term than what the buffer is currently sized for.

Lead Time Adjustment Factor

A lead time adjustment factor could apply to an individual part or a group of parts affected in the same way by what is prompting the adjustment (not necessarily parts in the same buffer profile). The use of a lead time adjustment factor coincides with a planned or known expansion of the lead times of a group of items. For example, if a major transportation route to a warehouse will be temporarily disrupted due to construction, a lead time factor may need to be applied to products at the warehouse that come through that route or its alternate. If the previous lead time was two days but the new lead time is estimated to be three days, on average, then a lead time adjustment factor of 1.5 should be employed for the duration of the construction project. An alternative would be to amend the part master to reflect the lead time change depending on the expected duration for this new lead time.

Summary

The third component of DDMRP is strategic buffer adjustments. Recalculated adjustments allow the buffers to recalculate their respective levels based on key attribute changes of which first and foremost is average daily usage. Other adjustments to buffer levels can be related to known or planned events. These are called planned adjustments. With regard to these events, the buffers will be either significantly overstated or understated if not adjusted. There are many types of planned adjustment strategies to be employed depending on the specific nature of the planned or known event.

CHAPTER 9

Demand Driven Planning

The fourth component of Demand Driven Material Requirements Planning is a proven and intuitive method of supply order generation—demand driven planning. In addition to lead time compression and variability dampening, the buffers placed at the decoupling points are the heart of supply order generation for Demand Driven MRP. They become a focal point for creating, promoting, protecting, and determining relevant information and materials. They also create the opportunity for a more elegant and visible way to generate supply orders. It starts with the consideration of what really is relevant information from a demand perspective.

The Shift to Actual Demand

Protecting the flow of relevant information and materials is the key to protecting and promoting return on investment performance. From a planning perspective, the right materials will not be available without the right information. Yet most conventional planning starts with information that is suspect. The direct connection between forecasted demand and supply order generation means that supply orders are being generated with signals that are known to be wrong. This ties up cash, capacity, space, and time by dealing with the resulting irrelevant materials and forces additional efforts to attempt to get the relevant materials.

The most relevant demand signal is a sales order. It is a known and stated desire from a known customer to buy. Yet using this most accurate piece of demand information has eluded planning professionals for decades. As described in Chapter 4, decoupling opens the door for its effective use. This highly accurate form of demand will provide the demand input into a daily planning equation for each buffered position. This equation is called the net flow equation.

The Net Flow Equation

Previously referred to as the "available stock equation" in DDMRP literature before 2016, the net flow equation provides the supply order generation recommendation signal (timing and quantity) for buffer replenishment. It is a key and unique aspect of DDMRP and should be performed daily on all decoupled positions. The net flow equation is elegant and intuitive while encompassing all the ranges of planning that most experienced planning professionals are concerned about with regard to supply order generation.

The net flow equation is simple:

On-hand + on-order – qualified sales order demand = net flow position

Figure 9-1 illustrates the components of the net flow equation, and the following list explores the components one by one:

- **On-hand.** The quantity of stock physically available. In Figure 9-1 this is represented by the large box labeled "On-Hand" in the middle of the graphic. The smaller boxes represent the actual quantity of units on hand and available for use.
- **On-order.** The quantity of stock that has been ordered but not received. In Figure 9-1 this is represented by the truck on the left-hand side of the graphic that is heading toward the on-hand position. The smaller boxes in the trailer represent the amount of units on order. This could be a single incoming order or several incoming orders. The on-order quantity is the total quantity that has been ordered but not received, irrespective of timing.
- **Qualified sales order demand.** The sum of sales orders past due, sales orders due today, and qualified spikes. In Figure 9-1 this is represented by the orders that are highlighted in "Today" and "Day 3." There is no past due amount represented in this figure. There are two sales orders due today and three sales orders that have combined to create a qualified spike. These two days of qualified demand are added together to get the total amount of qualified demand for today's computation of the net flow equation.

Figure 9-2 gives an example of the net flow equation. There are 6,412 units currently available in stock (on-hand). There are 2,468 that are inbound (on-order). There are 312 units to be shipped out today and 807 units due to ship on day 3. Today's net flow equation for this part is 6,412 (on-hand) + 2,468 (on-order) – 1,119 (qualified sales order demand) = 7,761 (net flow position).

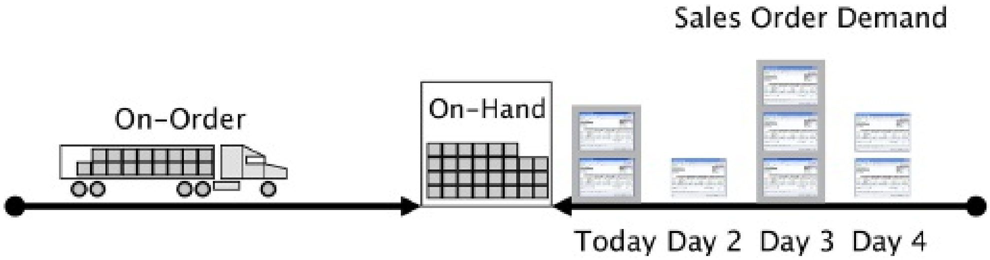

FIGURE 9-1 Illustrating the elements of the net flow equation

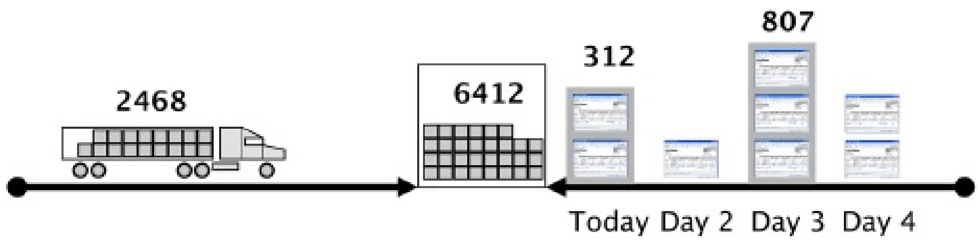

FIGURE 9-2 An example of the net flow equation

The net flow equation answers all the question of any robust planning equation:

- What do I have? The on-hand value.
- What is coming to me? The on-order value.
- What demand do I need to fulfill immediately? Sales orders past due and due today.
- What future demand is relevant? Qualified future spikes.

At this point a significant question still needs to be answered about the demand component of the net flow equation. Sales orders due in the past are the summation of the quantity of stock represented by sales orders we have yet to fulfill but were due before today. Sales orders due today are the total quantity of stock represented by the sales orders scheduled to be fulfilled today. Future qualified demand considers two conditions. In Figure 9-2 the day 3 total order quantity qualifies as a spike amount (807). It is not the number of sales orders that is triggering the spike but the total quantity of stock that those sales orders require within a day.

Qualifying Order Spikes

Two conditions must be met in order to qualify as an order spike. In DDMRP, an order spike is a qualifying quantity of known cumulative daily demand within a qualifying time window that threatens the integrity of the buffer. That means the qualifying level (order spike threshold) and the qualifying time window (order spike horizon) must be defined.

Condition 1: The Order Spike Threshold

The order spike threshold is a level that qualifies a spike in a particular environment. The sales orders for the same part number for each day are totaled and compared against this threshold. If the sum is greater than the threshold, then the entire amount (not just the amount above the threshold) is incorporated into the available flow equation as a qualified spike. An order spike threshold is depicted in Figure 9-3 by the horizontal dotted line. Sales orders for each day are represented by the boxes above the column. For example, day 6 has four sales orders due that day. The cumulative quantity of those sales orders is represented by the column beneath the sales orders. The three sales orders due two days from today (day 3) represent enough cumulative demand to qualify as a spike. This also occurs on day 9, where two sales orders combine to form a qualified spike.

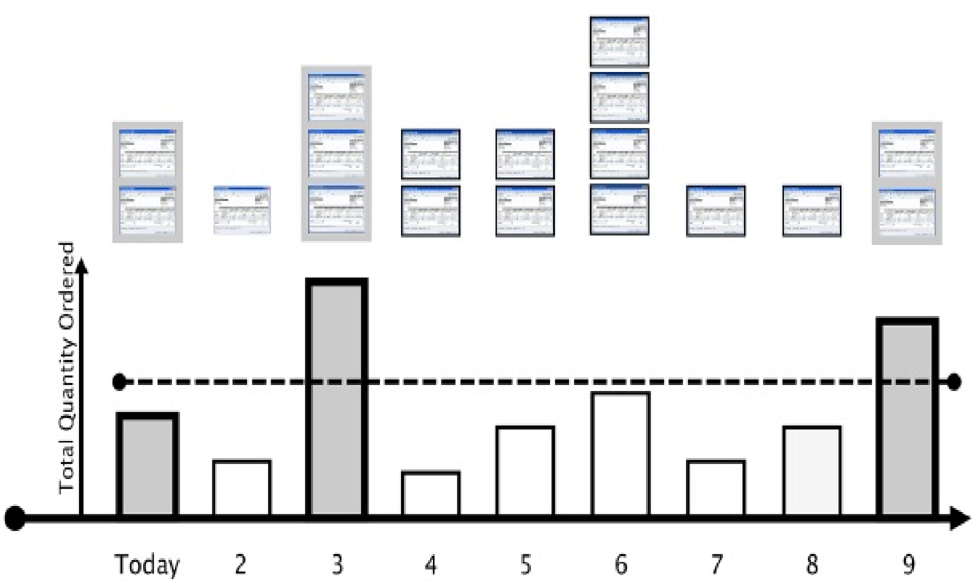

FIGURE 9-3 An order spike threshold

The entire cumulative amount of the spike is qualified and incorporated into the equation. This typically prompts the question of why the whole amount is taken and not just the portion above the threshold. The entire spike amount is taken to guard against successive or clustered spikes creating an overwhelming surge that would compromise buffer integrity if the entire amount were not taken into account. Taking the whole spike represents the most conservative risk-averse approach.

Setting the threshold level involves determining a level of daily demand that jeopardizes the integrity of the buffer. Three alternatives will be explored in determining the threshold level. The first two set the threshold in relation to the red zone of the buffer. Since the red zone is the embedded safety in the buffer, then spikes can be seen relative to their potential consumption of that safety. In early DDMRP implementations a default heuristic spike threshold of 50 percent of the red zone was utilized.

An alternative method relates to the border of the red base and red safety portion of the buffer for finished items. The red safety portion of the red zone directly relates to the variability of the part position. If the finished item is in the high variability category, then that part is subject to frequent spikes. Thus the order spike threshold could be set for the amount equivalent to the red zone base value. Any value above that base amount penetrates into the safety portion of the red zone.

Still another way to determine the order spike threshold has to do with qualifying an order spike in relation to the ADU of a specific part. This is often the most intuitive way for planners and buyers. It can also be historically validated.

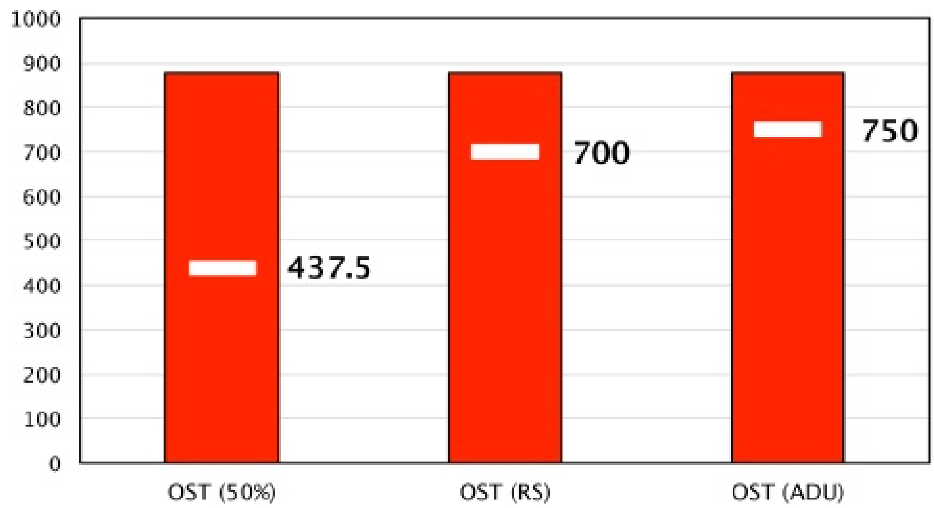

FIGURE 9-4 Different ways to set the order spike threshold for FPA

Figure 9-4 shows the different order spike threshold options for part FPA from the Company ABC example from Chapters 6 and 7. The total red zone value (red base + red safety) after compression was 875. The version labeled "OST (50%)" is the order spike threshold (OST) set at 50 percent of the total red zone (437.5). The version labeled "OST (RS)" is the OST set at the red base value (700). The version labeled "OST (ADU)" is the OST set at three days of average daily usage [3 × ADU (250) = 750].

These alternatives are options for planning teams to evaluate how they would like to determine order spike amounts in their specific environment. Each option will most likely yield a different number of qualified spikes over an extended period of time. In this example, the most qualified spikes will occur with the OST 50 percent method. Thus in most cases the OST 50 percent method is typically the most conservative method of order spike threshold determination. This will tend to yield slightly higher levels of overall inventory but best protect service levels. In reality, the difference in effectiveness between the methods is negligible if the DDMRP model is sensibly constructed, managed, and adapted.

No matter what method is used for OST determination, the OST should always be within the total red zone value. The red zone is the safety embedded in the buffer. Having an OST value greater than that total safety opens the door for insufficient coverage from the buffer.

Condition 2: The Order Spike Horizon

The second condition of qualifying an order spike is to set a qualifying future time window called an order spike horizon (OSH). This horizon is a window of time within which cumulative daily demand can qualify as a spike if that cumulative daily value is above the OST. If the cumulative daily demand is outside that window (farther into the future), then that cumulative daily amount will not be qualified in the net flow equation.

The OSH should be set to at least one decoupled lead time in the future for the buffered part. This allows enough time for the part's buffer to properly compensate for the spike. If the order spike horizon is beyond the sales order visibility horizon, it simply means that qualified spikes will routinely appear within the order spike horizon. In this case the part's buffer profile should be set to a higher variability category in order to compensate for this late qualification.

Figure 9-5 combines the use of an order spike threshold and an order spike horizon. The OSH is represented by the box with hatched lines overlaying the daily sales order quantity positions for a period of seven days (today plus six days in the future). In this case the decoupled lead time of the part is six days. Thus the OSH is today + DLT of the part. When the horizon is imposed, the day 9 cumulative demand is not qualified in the net flow equation. It will take two more days before that demand qualifies as a spike.

Yet this begets an additional question. If we are so concerned with relevant demand signals and sales orders represent the most accurate form of demand signal, then why not take all known demand into account in the net flow equation?

All known sales orders are not included in the net flow equation because that demand is essentially already accounted for in the buffer. If the daily sales order demand is under the threshold, those orders represent normal or average demand. The buffers were built using this average rate of use. Thus what is due today, what was due in the past, and what qualifies as a spike are really all that is relevant from a demand perspective.

Figure 9-6 depicts what an order spike qualification screen could look like in a DDMRP-compliant information system. The part number ("EXAMPLE"), order spike horizon (6 days), order spike threshold (100), and today's date (6/21) are displayed at the top of the screen. The information displayed in the screen corresponds to the example in Figure 9-5. Figure 9-5 has been added beneath the sample screen in order to show how the two are connected. In the sample screen, future days are displayed in an descending manner. The sales orders within the order spike horizon have shaded boxes.

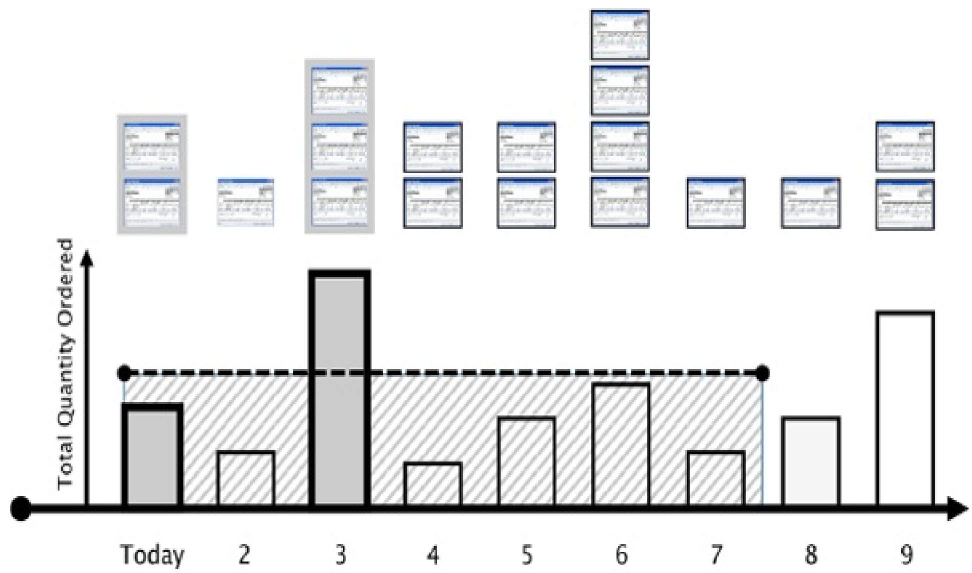

FIGURE 9-5 The order spike threshold and horizon

Part # EXAMPLE Order Spike Horizon = 6 days Order Spike Threshold = 100
Today's Date: 06/21

Demand	Quantity	Due Date
SO# 1234	35	06/21
SO# 1235	40	06/21
SO# 1236	40	06/22
SO# 1237 – SPIKE!	**70**	**06/23**
SO# 1238 – SPIKE!	**50**	**06/23**
SO# 1239 – SPIKE!	**60**	**06/23**
SO# 1240	10	06/24
SO# 1241	25	06/24
SO# 1242	60	06/25
SO# 1243	10	06/25
SO# 1244	10	06/26
SO# 1245	15	06/26
SO# 1246	50	06/26
SO# 1247	20	06/26
SO# 1248	40	06/27
SO# 1249 (OUTSIDE HORIZON)	70	06/28
SO# 1250 – SPIKE! (OUTSIDE HORIZON)	80	06/29
SO# 1251 – SPIKE! (OUTSIDE HORIZON)	60	06/29

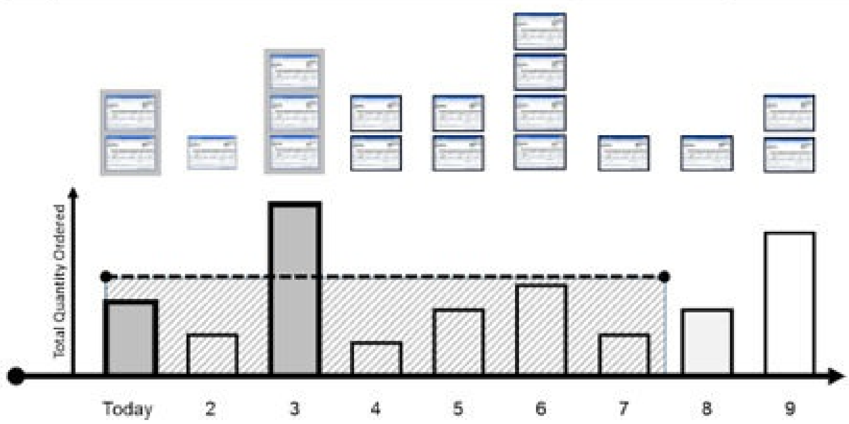

FIGURE 9-6 Sample order spike qualification screen

There are two sales orders (1234 and 1235) totaling 75 units that are due today. The spike two days from today (6/23) is composed of three sales orders (1237, 1238, and 1239) totaling 180 units. That means that today's total qualified demand for the net flow equation is 255 (0 past due + 75 due today + 180 future spike demand on day 3). There is another spike on 6/29 made up of sales orders 1250 and 1251 totaling 140 units but is outside the horizon and labeled as such.

An additional horizon and threshold consideration for finished items that have large and known dependent demand orders against them in the remote future can occur with large planned promotional events in which customers have agreed to take a significant amount of stock within a short window. The event is planned, agreements have been made, and the demand is real. In this case an additional OSH and OST should be formulated. We will call them OSH2 and OST2. They need to be used in combination with each other.

The OSH2 can be increased to the cumulative lead time of the parent part in order to account for the spike against buffered component positions that would be unable to absorb the uplift within the decoupled lead time of the parent. The OST2 will also increase and be applied to the OSH2 period only.

Figure 9-7 depicts the combined use of a close-in order spike horizon and threshold (OSH1 and OST1) and an OSH2 and OST2. The close-in order spike threshold (OST1) is the solid line terminating at day 7 with a solid circle (OSH1). The cumulative lead time of this finished item is 30 days, indicated by the open circle at the end of the dotted OST2 line. The long-range order spike threshold (OST2) is 500 units. Beyond the OSH1 sales order, demand begins to be spotty. Most customers have simply not ordered yet. But the known large orders (1,000 each) to support the planned promotion are visible on days 28 and 29. These orders are significant, and if allowed to only be qualified within the OSH1, they would overwhelm the finished buffer and the component buffers supporting it.

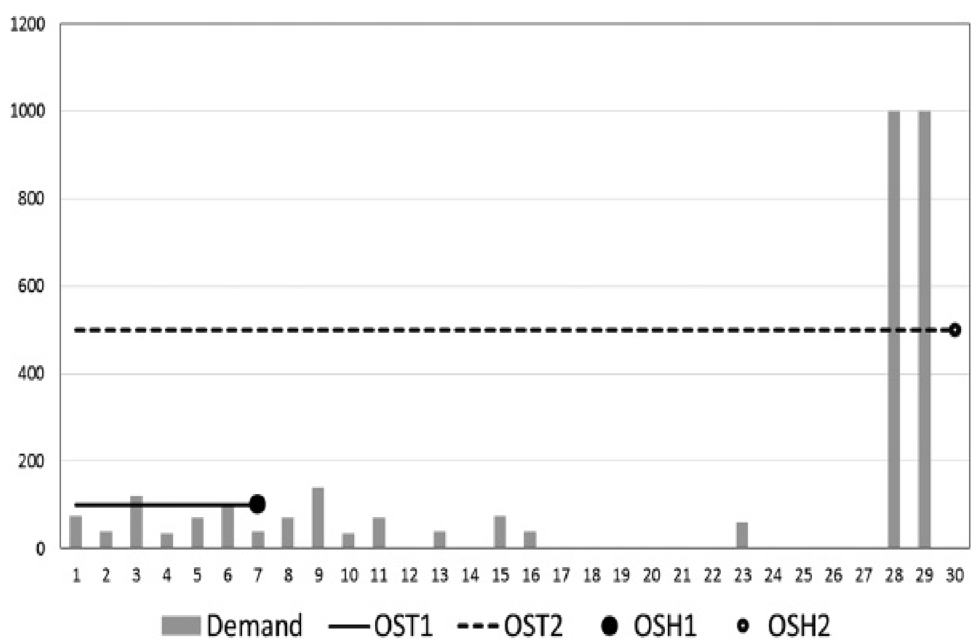

FIGURE 9-7 Short- and long-range order spike thresholds and horizons

In this case both the demands on day 28 and day 29 would be qualified as demand in the net flow equation. This amount of 2,000 would be added to the 180 (sales orders 1237, 1238, and 1239) within the OSH1 and the amount of 75 due today (sales orders 1234 and 1235) to create a total qualified demand of 2,255.

This also means that any promotions or large-order negotiations should be finalized in advance of the OSH2 horizon. Without this commitment, the buffers will always be at risk to significant orders.

One further note on spike qualification involves the continued qualification of spikes. Spikes continue to be qualified each additional time the net flow equation is performed (typically once a day). Figure 9-8 continues the example that we have been using for spike qualification on 6/22 (day 2), one day later than the original date of the example. Today's qualified demand within the OSH1 is 40 (sales order 1236) plus the qualified spike of 180 (sales orders 1237, 1238, and 1239) for a total of 220. The next day (6/23) will result in a new spike qualification of 140 (sales orders 1250 and 1251). That spike will be combined with the 180 due that day (sales orders 1237, 1238, and 1239) for a total qualified demand of 320 within the OSH1.

Cumulative daily spikes continue to be qualified each day from the time in which they were first qualified in order to balance against any on-order quantities that have been generated as a result of the spike inclusion. To disregard the spike after qualifying it on its first day would distort the net flow position and make the buffer appear to be oversupplied. This will be demonstrated later in this chapter with the simulation of the net flow equation over a series of days.

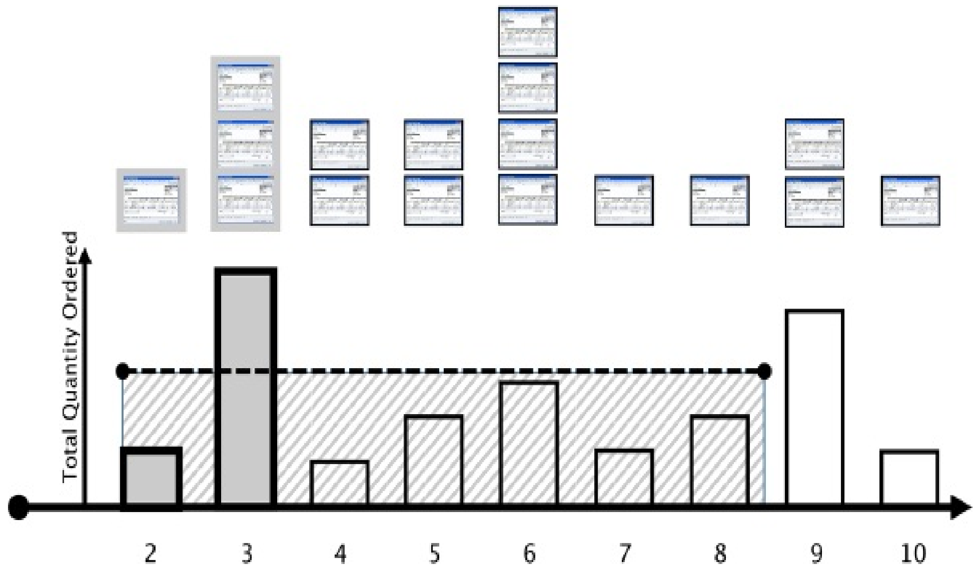

FIGURE 9-8 Qualifying demand on 6/22 within the OSH1

Supply Order Generation Based on Net Flow Position

Each time the net flow equation is computed for each buffer, it yields a number that is called the net flow position within each buffer. The net flow position will dictate whether a supply order will be recommended against the buffered position. If the position is below the top of yellow (TOY), then a supply order is recommended for a quantity that is the difference between the net flow position and the top of green (TOG). Figure 9-9 illustrates the TOG, TOY, and top of red (TOR) positions within the 401P buffer from the Company ABC example from Chapters 7 and 8.

TOG is the summation of all the zones of a buffer. In our 401P example this is 10,888. TOY is the summation of the yellow and red zones, which is 8,938. TOR is the full quantity of the red zone, which is 2,438. When the net flow position is at or below 8,938, a supply order will be generated to restore the net flow position to TOG. This means that the typical order will be at least the full amount of the green zone (1,950) or greater.

Consider the following information for 401P:

On-hand = 2,652
On-order = 6,233
Qualified sales order demand = 712

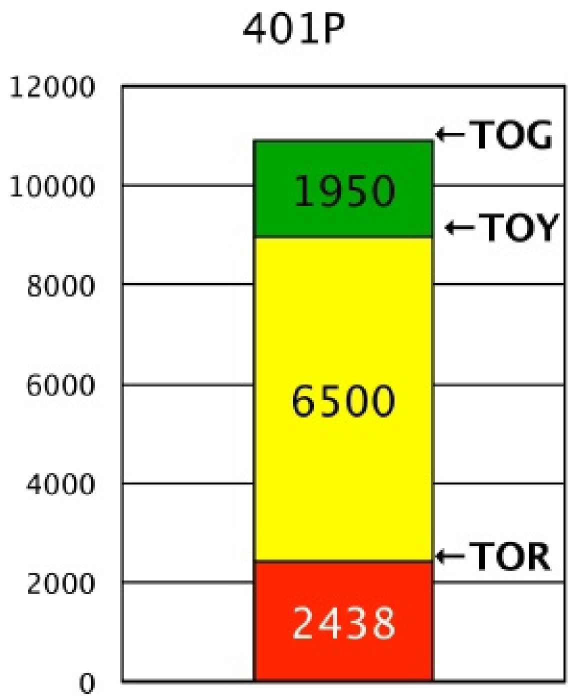

FIGURE 9-9 TOG, TOY, and TOR values for 401P

With this information, then the net flow position for 401P on this particular day is 8,173. This means the net flow position is below TOY. Figure 9-10 shows the 401P net flow position.

The TOG value of the 401P buffer is 10,888. This creates a supply order recommendation of 2,715. Once this supply order is approved, the net flow position changes to 10,888. The order recommendation has now become an on-order quantity. This increases the total on-order quantity to 8,948 (2,715 new on-order quantity + 6,233 old on-order quantity). The net flow equation after order acceptance is 2,652 (on-hand) + 8,948 (on-order) – 712 (qualified sales order demand).

On acceptance, the order is assigned a request date one decoupled lead time into the future. Component 401P is a purchased part. As such, its decoupled lead time is the same as its purchasing lead time. The order would be assigned a due date 10 days from the date of the approval. The type of calendar a business uses must be accounted for in calculating the due date.

Additionally, the net flow position is typically displayed as a percentage of TOG and with the zone color that the position falls within. In the previous 401P example, the net flow position would be displayed as 75.1 percent and be coded yellow. Figure 9-11 illustrates what a basic DDMRP planning screen should display, including the basic elements of the net flow equation and supply order generation logic as well as the lead time and due date for any recommended orders. The figure uses the data for 401P in Figure 9-10.

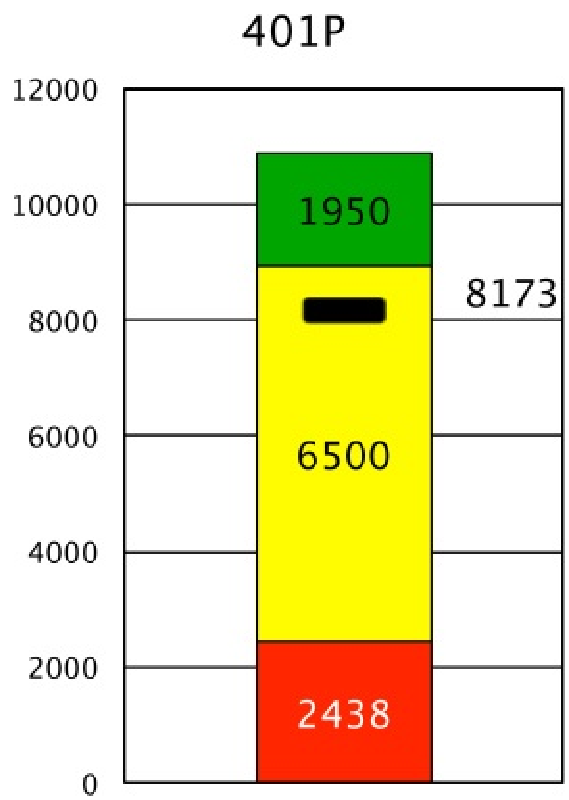

FIGURE 9-10 Net flow position of 401P

Today's Date: 15-July											
Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Recommendation	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
401P	YELLOW 75.1%	2652	6233	712	8173	2715	25-Jul	2438	8938	10888	10

FIGURE 9-11 DDMRP sample planning screen with 401P data

Note that today's date is July 15. An order recommendation for 2,715 has been made with a request date of July 25. This is 10 days from today. The column titled "Planning Priority" is the net flow position expressed as a percentage of the top of green. The box has a yellow shade, denoting that the percentage is inside the yellow zone. There are two expressions of buffer priority. One is discrete and expressed as a percentage. The other is general and is expressed as a color. Thus for each part, a planner or buyer can quickly get a sense of the part's status relative to its planned buffer position. While this is quite powerful on a part-by-part basis, the true power of this view can really only be appreciated when seen with multiple parts.

Figure 9-12 is the same planning screen but now with multiple buffered parts. The sequence is determined by the planning priority column percentage. The lower the percentage, the higher the planning priority. This sequencing now gives a general and discrete sense of relative priority across multiple parts calling for resupply. This is crucial when limitations or constraints are present in an environment. When dollars, time, space, and resource capacity are at a premium, it is extremely advantageous to be able to quickly focus on the highest priority.

With this view we can quickly see that there is a buffered position in immediate trouble. Part 406P has a net flow position colored red and is at 19.8 percent of the top of green. A supply order for 2,606 is recommended and should be immediately approved. In this case we can also see that 401P has the lowest relative priority among all parts calling for supply order generation. We also see a part (404P) that has no recommended supply because its net flow position is in the green zone at 97.6 percent of the top of green. In most cases this part would be filtered out of the planner or buyer view because no planning action is required.

This relative priority distinction is a crucial differentiator between the conventional MRP planning alerts and action messages and the highly visible and focused DDMRP approach. Conventional MRP is a binary system. You are either OK or not OK with regard to each part. There is little sense of how parts compare with each other—you need to either act or not act. Under the DDMRP approach, planners and buyers can quickly judge the relative priority without massive amounts of additional analysis and data queries.

Today's Date: 15-July

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Recommendation	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
406P	RED 19.8%	401	506	263	644	2606	4-Aug	750	2750	3250	20
403P	YELLOW 43.4%	1412	981	412	1981	2579	23-Jul	1200	3600	4560	8
402P	YELLOW 69.0%	601	753	112	1242	558	24-Jul	540	1440	1800	9
405P	YELLOW 74.0%	3400	4251	581	7070	2486	24-Jul	1756	7606	9556	9
401P	YELLOW 75.1%	2652	6233	712	8173	2715	25-Jul	2438	8938	10888	10
404P	GREEN 97.6%	1951	1560	291	3220	0		1050	2550	3300	6

FIGURE 9-12 DDMRP planning screen with multiple buffered items

Sometimes the recommended order quantity could exceed the top of green due to order multiples. In these cases, supply order generation rules can be constructed to incorporate the order multiple but minimize going over the top of green (OTOG). An example would be to limit the inclusion of an additional order multiple to situations in which the OTOG net flow position would be less than the quantity under the top of green without the inclusion.

Once today's orders have been fulfilled and any on-order has been received, on-hand will be adjusted accordingly. This will create an ending on-hand inventory that will be used for tomorrow's net flow equation. Additionally, if on-order is received and converted to on-hand, then the on-order quantity will be adjusted down for tomorrow's net flow equation. If on-order is received but is under a quality or inspection hold, it should still be treated as on-order, as it is truly not available as on-hand.

The following simulation demonstrates how a strategically buffered position behaves over a series of days.

Simulating DDMRP Supply Order Generation

Now that the net flow equation and its role in supply order generation have been discussed, a simulation can be run to demonstrate the daily use of a net flow equation against a buffered position. Figure 9-13 contains the information for a part called "Example." This part has an average daily use of 10 and a decoupled lead time of 7 days and is in the MMM buffer profile category—it is a manufactured item, with medium lead time and medium variability associated with it. The lead time factor is 0.5, and the variability factor is 0.5. The order spike threshold has been set to 50 percent of the red zone (26) over an 8-day order spike horizon. This means that an order spike is defined as 2.6 times the average daily usage or greater within any future daily bucket within the order spike horizon.

Example Part Buffer For Simulation			
Average Daily Usage	10	Green Zone	35
Buffer Profile	M, M (.5), M (.5)		LT Factor: 35 (DLT (7)x ADU (10) x Lead Time Factor (.5))
MOQ	20		Minimum Order Quantity: 20
Imposed or Desired Order Cycle (DOC)	None		Order Cycle: N/A
Decoupled Lead Time (DLT)	7 days	Yellow Zone	70 (7(DLT) x 10(ADU))
Order spike threshold = 26 **Order Spike Horizon = 8 days**		Red Zone	52 (Red Base (35) + Red Safety (17))
			Red Base: 35 (DLT (7)x ADU (10) x Lead Time Factor (.5))
			Red Safety: 17 (Red Base (35) x Variability Factor (.5))

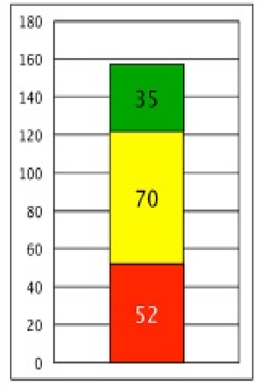

FIGURE 9-13 Simulated replenished part information

A simulation will be run in which the activities against this buffered position for 21 days will be described.

Figure 9-14 demonstrates the starting situation for our example part. The buffer is in the middle of the graphic. Sales order demand is seen over an 8-day horizon to the right of the buffer (labeled "Demand Side"). This horizon corresponds to the order spike horizon (DLT + 1 day). The order spike threshold of 26 is depicted by the dotted line halfway up the red zone of the buffer and extending over the order spike horizon, terminating in a small circle on day 8. Today is day 1, indicated by the number 1 on the arrow to the right of the buffer. The arrow is pointing toward the buffer because it represents demand quantities in daily buckets coming at the buffer. On day 1 we see a demand of 10 due indicated by the light gray bar. All bars are intended to be to the scale of the buffer level. On day 5 we see actual known demand for 5. As time moves forward 1 day, the number 1 will be replaced by 2 on day 2.

On the left-hand side of the buffer, we see another arrow flowing from left to right toward the buffer (labeled "Supply Side"). This represents the inbound supply to the buffer in daily buckets (represented by dark bars). The value of each supply order is specified on the bar. The values on the arrow are negative because the arrow is indicating how many days left to the supply order receipt. There is an order for 35 pieces that will be received by end of day tomorrow. Tomorrow all supply bars will shift right by 1 day. The supply for 35 will then be at –1 on the supply side.

Additionally, below the buffer and the supply and demand sides is the planning screen for our example item. In this case the net flow position has prompted a supply order recommendation. Today’s starting on-hand is 65, indicated by the dotted line inside the buffer in the lower yellow zone. There are two open supply orders totaling 72 units (orders of 37 and 35). One order for 35 is due tomorrow, while the other is due 6 days from now. There is also qualified demand in at least two daily buckets. There are 10 due today and there is a qualified spike on the 8th for 30. There is no past due amount. Total qualified demand is 40 [10 + 30]. The net flow position is 97, indicated by a solid line in the upper half of the yellow zone. This net flow position yields a yellow planning priority of 61.8 percent [97 (net flow position)/157 (top of green)].

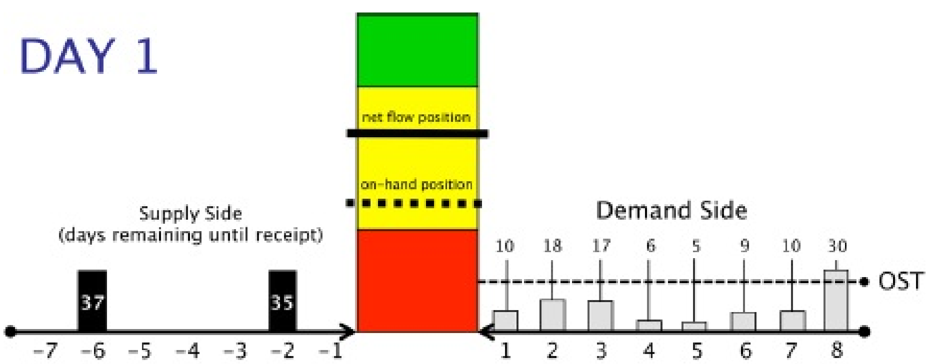

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 61.8%	65	72	40	97	60	Day 8	52	122	157	7

FIGURE 9-14 Simulation day 1

A supply order for 60 is recommended and will be approved. The request date for the order is 7 days from today on the 8th. That newly created order will show up in day 2 (tomorrow) as open supply. Finally, the orders due today will be shipped out. There are no receipts occurring today, so today’s ending on-hand inventory will be 55 units [65 (starting on-hand) – 10 (units due today)].

Day 2 is represented by Figure 9-15. The starting on-hand quantity is 55 (yesterday’s ending on-hand quantity). Here we can see that the demand side has shifted by 1 day to the left. The demand arrow now ranges from the 2nd to the 9th, and the order spike horizon spans to the 9th. On the supply side we see the order for 60 that was created yesterday. There are now three orders of open supply totaling 132 units. A supply of 35 will be delivered today, while a demand of 18 will be fulfilled. That means that ending on-hand inventory for the day will rise by 17 (35 – 18) up to 72. The net flow position is at 139 and is solidly in the green zone (88.5 percent of top of green). Thus there is no additional recommended quantity.

The spike due on the 8th continues to be qualified as part of the net flow equation because it balances against the open supply. If the spike were removed from the net flow equation, the net flow position would rise by 30 to 169 and show over the top of green. This may lead to the impression that the buffer is oversupplied when in fact it is not.

Figure 9-16 shows the continuation of the simulation on day 3. Supply has shifted to the right by 1 day. The order for 35 was received on day 2 and is no longer visible. On-order now totals at 97 (60 + 37). Demand has shifted to the left by 1 day. Qualified demand is at 47 [17 (due today) + 30 (qualified spike due on the 8th)]. Starting on-hand is at 72. Today's net flow position is at 122 (72 + 97 – 47). This position at 122 is a yellow planning priority (77.7 percent of top of green). A new supply order is recommended for 35 [157 (top of green) - 122 (net flow position)]. This new supply order will have a request date of the 10th (7 days from today). Ending on-hand will drop by 17, as no supply receipts will occur and 17 will be sent out the door.

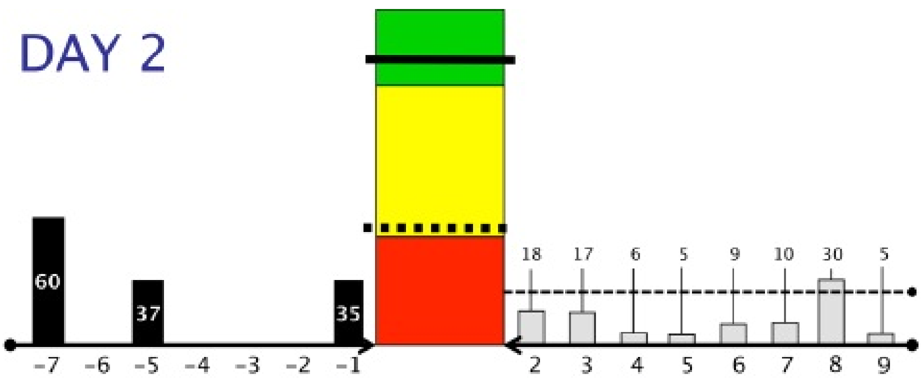

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 88.5%	55	132	48	139	0		52	122	157	7

Ending OH = 72

FIGURE 9-15 Simulation day 2

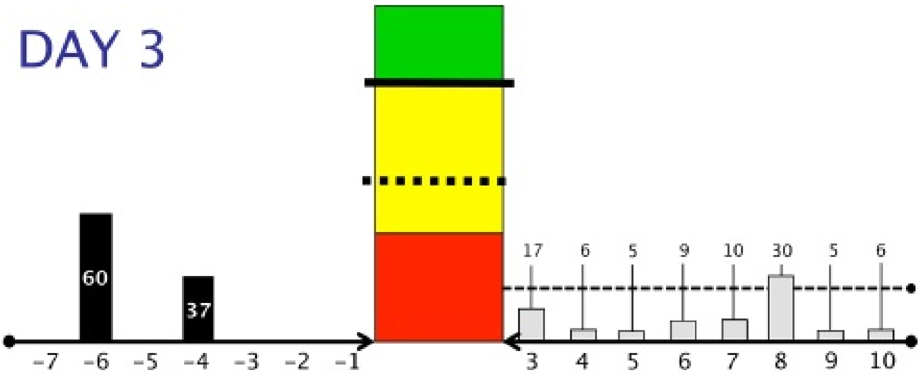

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 77.7%	72	97	47	122	35	Day 10	52	122	157	7

Ending OH = 55

FIGURE 9-16 Simulation day 3

Figure 9-17 shows the continuation of the simulation on day 4. Starting on-hand inventory is at 55. The existing supply orders have shifted to the right, and a new supply order for 35 is now visible. Total on-order is 132. The demand side now ranges from the 4th to the 11th. There are 6 units due to ship today, and the spike continues to be qualified. Total qualified demand is 36. The net flow position is green at 151 (96.2 percent of top of green). There is no supply order recommended. On-hand will drop by 6 to 49, as no incoming supply will be received.

Figure 9-18 shows the continuation of the simulation on day 5. Starting on-hand inventory is now at 49. Supply orders have shifted to the right by 1 day and still total 132. The demand side now ranges from the 5th to the 12th. The qualified demand total is 35 [5 (due today) + 30 (spike)]. Today's net flow position is green at 146 (93 percent of top of green). There is no recommended supply order.

According to the net flow equation, the buffer is properly planned, and yet the on-hand is in the red zone. Is this a problem? This is the first time we can actually see the difference between the planning and execution perspectives of DDMRP. The buffer was built using an ADU value of 10. That means that with a current on-hand of 49 there are nearly 5 days of average coverage still contained in the buffer. The next 3 days are light on demand, and there is open supply scheduled to be received in the short term.

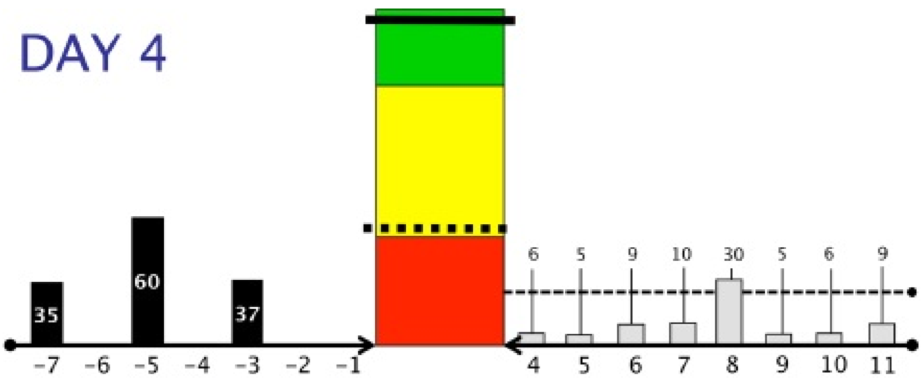

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 96.2%	55	132	36	151	0		52	122	157	7

Ending OH = 49

FIGURE 9-17 Simulation day 4

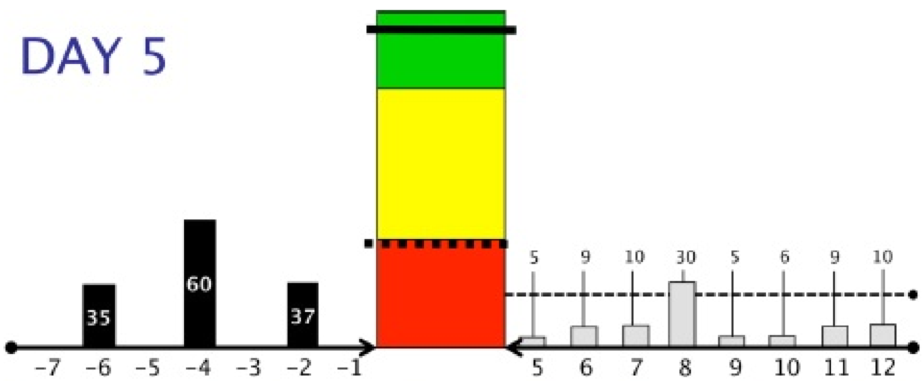

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 93%	49	132	35	146	0		52	122	157	7

Ending OH = 44

FIGURE 9-18 Simulation day 5

On-hand has penetrated into the red zone, but this view of the red zone is built for planning. If the net flow position had dipped into the red zone, then that is a big problem. On-hand dipping into the red is expected. The red zone is part of the buffer and is intended to be used. If on-hand dips too far into the buffer, at some point it becomes a problem, but on day 5 this is not the case. It has only slightly penetrated the red zone, has light upcoming demand, and has a significant amount of open supply due in the near term. Planners at this time should not expedite existing supply or launch a new order, as this is simply not required. Execution urgency is covered in depth in Chapter 10.

The on-hand situation will erode further, as 5 are due today and no receipts will occur. Ending on-hand will drop to 44.

Figure 9-19 shows the continuation of the simulation on day 6. Starting on-hand is at 44. That number still represents over 4 days of average coverage, and a supply order is due to be received today. There is no recommended supply order, as the net flow position is green at 137 (87.3 percent of top of green). Ending on-hand is 72, as 37 are received and only 9 are shipped out.

Figure 9-20 shows the continuation of the simulation on day 7. The on-hand position of the buffer is now 72. On-order has dropped to 95, and qualified demand is at 40. There is no recommended supply order, as the net flow position is green at 127 (80.9 percent of top of green). Ending on-hand will be at 62, as 10 are shipped out and none are received.

Figure 9-21 shows the continuation of the simulation on day 8. Starting on-hand is 62. On-order is 95, and qualified demand is only 30 (the spike has finally come due, and there are no additional spikes within the order spike horizon). There is no additional supply order recommendation, as the net flow position is green at 127 (80.9 percent of top of green). Ending on-hand will actually rise to 92 despite the fulfillment of the spike. This is due to the receipt of the supply order of 60.

Figure 9-22 shows the continuation of the simulation on day 9. Starting on-hand is at 92, and on-order consists of only 35. On the demand side a large drop-in order of 30 has been accepted for delivery tomorrow. This 30 is on top of the 6 that was already ordered. This creates a qualified spike. The demand quantity due on the 10th will now be included in the net flow equation. Thus the total qualified demand is 41. The net flow position on day 9 is 86 [92 (on-hand) + 35 (on-order) – 41 (qualified demand)]. This actually places the net flow position below the on-hand position. A supply order for 71 with a request date of the 16th is recommended. Ending on-hand inventory will be 87, as 5 are shipped out and none are received.

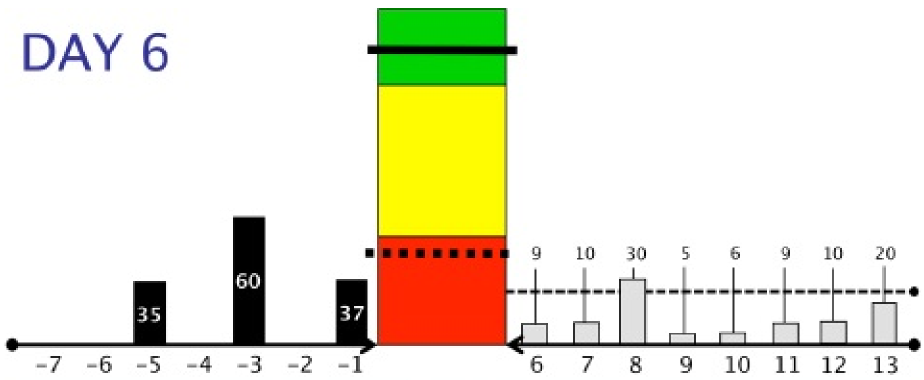

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 87.3%	44	132	39	137	0		52	122	157	7

Ending OH = 72

FIGURE 9-19 Simulation Day 6

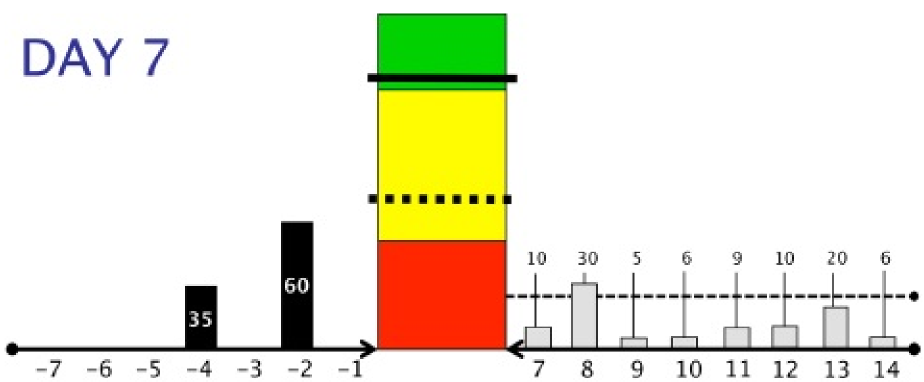

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 80.9%	72	95	40	127	0		52	122	157	7

Ending OH = 62

FIGURE 9-20 Simulation day 7

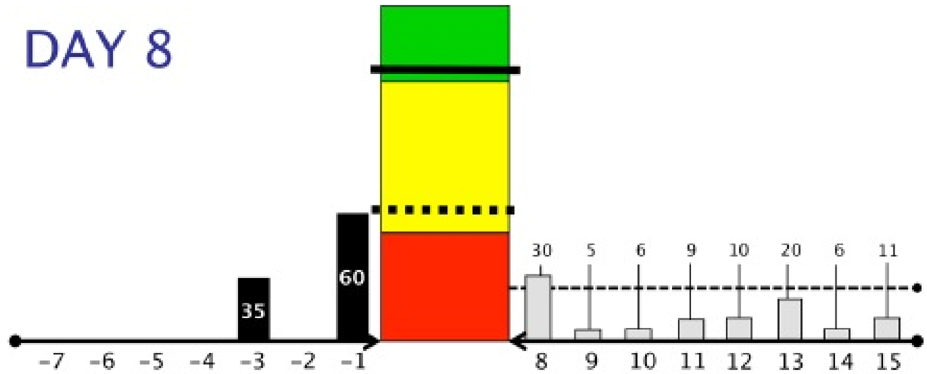

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 80.9%	62	95	30	127	0		52	122	157	7

Ending OH = 92

FIGURE 9-21 Simulation day 8

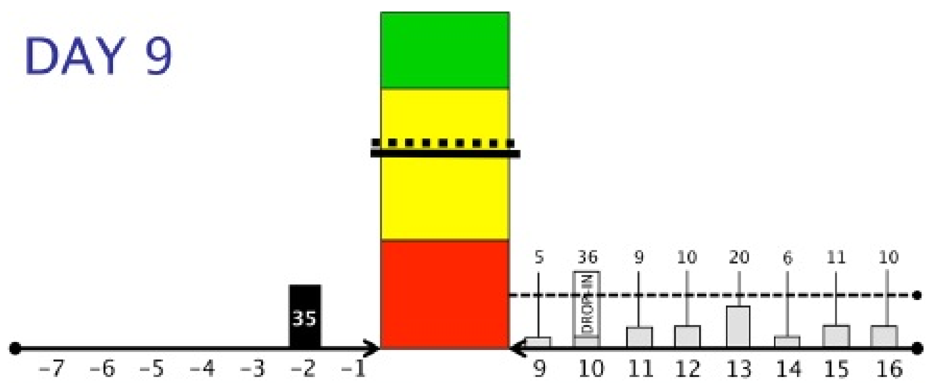

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 54.8%	92	35	41	86	71	Day 16	52	122	157	7

Ending OH = 87

FIGURE 9-22 Simulation day 9

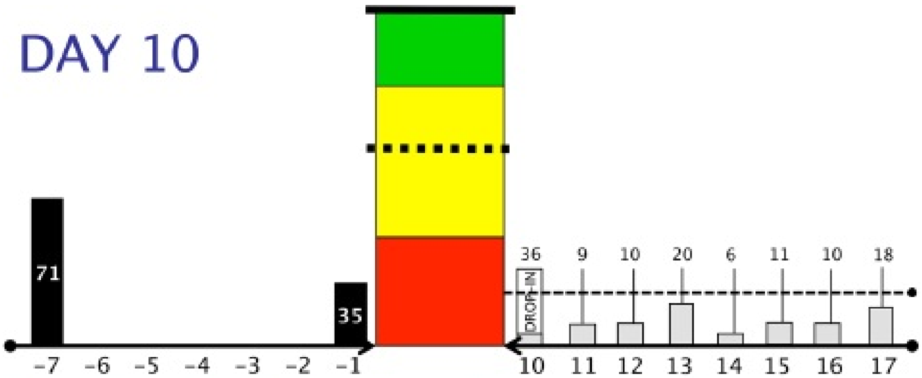

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 100%	87	106	36	157	0		52	122	157	7

Ending OH = 86

FIGURE 9-23 Simulation day 10

Figure 9-23 shows the continuation of the simulation on day 10. Starting on-hand is at 87. On-order now sits at 106 (35 of which will be received today). Qualified demand sits at 36 (including the drop-in order approved yesterday). There is no supply order recommendation, as the net flow position is at the top of green (157). Ending on-hand will be 86, as 35 are received but 36 are shipped out.

Figure 9-24 shows the continuation of the simulation on day 11. The starting on-hand position is at 86. Despite having had a large drop-in order, the buffer seems to be in excellent shape. On-order is 71, and qualified demand is at 9. There is no recommended supply order, as the net flow position is green at 148 (94.3 percent of top of green). Ending inventory is 77, as there are no receipts and 9 are fulfilled.

Figure 9-25 shows the continuation of the simulation on day 12. There is no recommended supply order, as the net flow position is green at 138 [77 (on-hand) + 71 (on-order) – 10 (qualified demand)]. Ending on-hand will be 67.

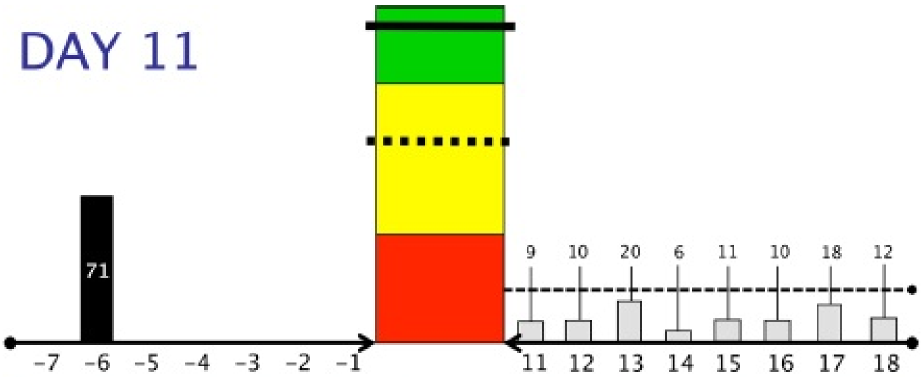

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 94.3%	86	71	9	148	0		52	122	157	7

Ending OH = 77

FIGURE 9-24 Simulation day 11

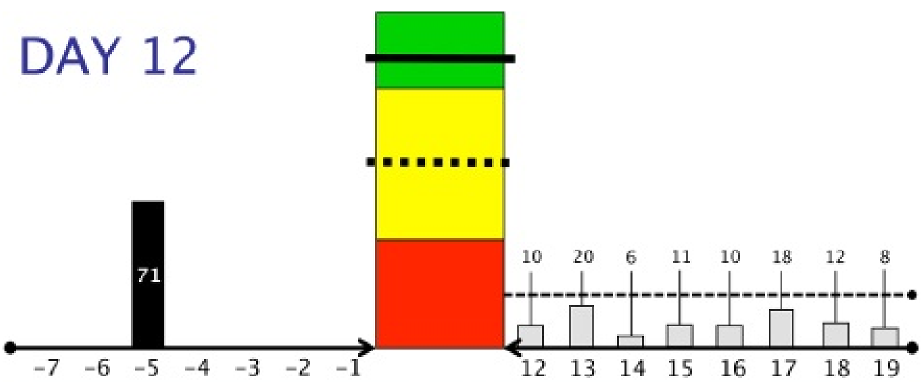

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 87.9%	77	71	10	138	0		52	122	157	7

Ending OH = 67

FIGURE 9-25 Simulation day 12

Figure 9-26 shows the continuation of the simulation on day 13. The net flow position [67 (on-hand) + 71 (on-order) – 20 (qualified demand)] yields a recommended supply order for 39 with a request date of the 20th. Ending on-hand will be 47.

Figure 9-27 shows the continuation of the simulation on day 14. Starting on-hand is 47. The new supply order for 39 is now visible and combines with the existing order of 71 for a total on-order amount of 110. Qualified demand is only at 6. There is no recommended supply order, as the net flow position is green at 151 (96.2 percent of top of green). Ending on-hand will be 41.

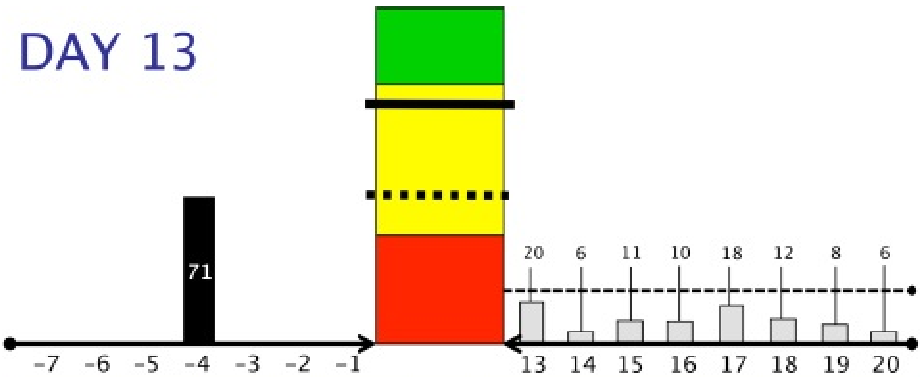

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 75.2%	67	71	20	118	39	Day 20	52	122	157	7

Ending OH = 47

FIGURE 9-26 Simulation day 13

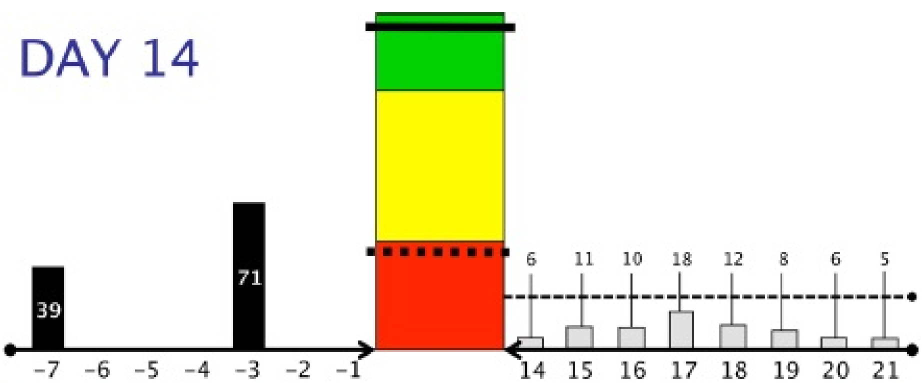

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 96.2%	47	110	6	151	0		52	122	157	7

Ending OH = 41

FIGURE 9-27 Simulation day 14

Figure 9-28 shows the continuation of the simulation on day 15. Starting on-hand is 41; it’s in the red but not deep enough to cause much concern since a large supply order will be received tomorrow. The net flow position is green at 140. There is no supply order recommendation.

Figure 9-29 shows the continuation of the simulation on day 16. On-hand has eroded to nearly half the red zone. If there was no impending supply, this might be cause for major concern. Still the situation dictates that the planner at least check on the incoming order status. Upon being assured that it will be delivered today, no additional action should be required. The severity of the on-hand situation and the level of the response are dictated by the amount of current and near-term future on-hand penetration into the red zone. Once again this is the execution component of DDMRP and is explored in depth in the next chapter.

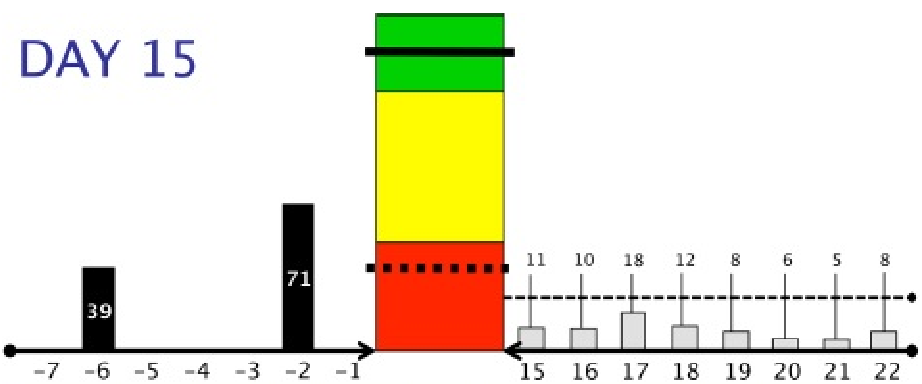

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 89.2%	41	110	11	140	0		52	122	157	7

Ending OH = 30

FIGURE 9-28 Simulation day 15

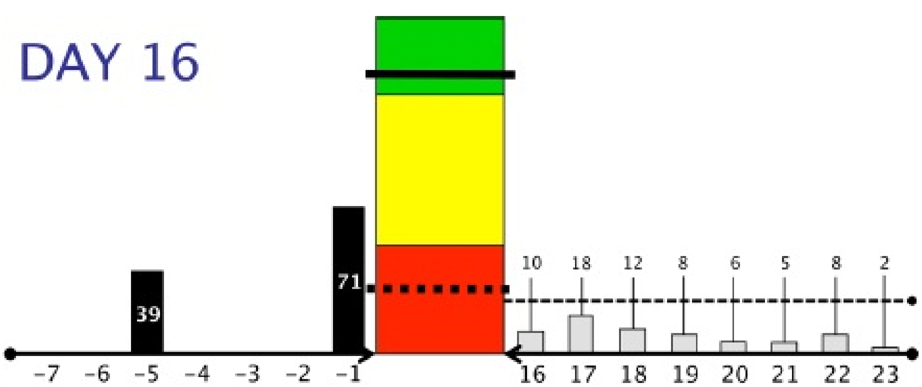

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 82.8%	30	110	10	130	0		52	122	157	7

Ending OH = 91

FIGURE 9-29 Simulation day 16

The net flow position on day 16 is green at 130 (82.8 percent of top of green). There is no supply order recommendation. Ending inventory will be 91, as 71 are received and 10 are shipped out. This is almost too easy—what could possibly go wrong?

Figure 9-30 shows the continuation of the simulation on day 17. Starting on-hand is only 51. What happened? The supply order for 71 was received, but on inspection 40 were found to have quality issues. Those 40 were placed on hold and moved back to on-order status. This brings on-order to 79. The items will need to be reworked and will be unavailable for several days. The “On-Order” column shows a yellow shading to indicate that the on-order quantity contains items on hold. The net flow position is yellow at 112 (71.3 percent of top of green). A supply order for 45 with a request date of day 24 has been recommended. Ending on-hand will be 33, as 18 are shipped out and none are received.

Figure 9-31 shows the continuation of the simulation on day 18. The on-hand situation is cause for concern. The on-hold open supply is still days away from being made available. On-hand still covers the next 3 days of known demand with a supply order due on the 20th. No expedite is ordered or additional supply is launched at this time. There is no supply order recommendation, as the net flow position is green at 145 (92.4 percent of top of green). Ending on-hand inventory will be 21 (12 shipped out and none received).

Figure 9-32 shows the continuation of the simulation on day 19. The starting on-hand situation has now eroded to less than half the red zone. The open supply is still one day away from being received, and the reworked parts have been progressing slower than expected. At this point the planner decides that the on-hand status dictates action. The planner requests an expedite of the closest-in supply order for 39 pieces. The supply order for 39 now has an exclamation mark over it, indicating that it is on expedite status. This order is moved up on the schedule, and overtime is applied to bring the order in at the end of today. Ending on-hand will increase by 31, as 8 are shipped out and 39 are received.

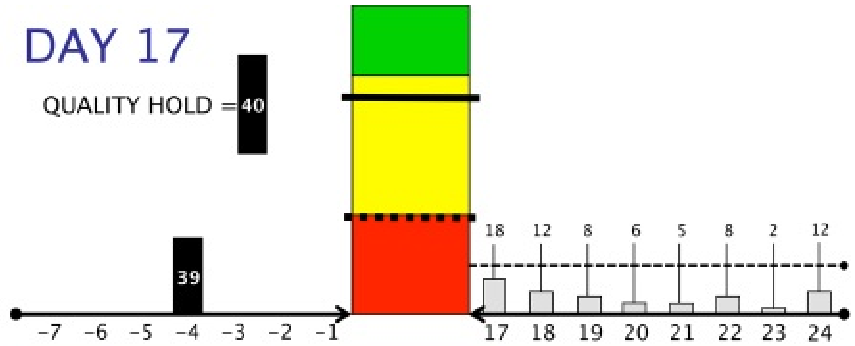

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 71.3%	51	79	18	112	45	Day 24	52	122	157	7

QUALITY HOLD

Ending OH = 33

FIGURE 9-30 Simulation day 17

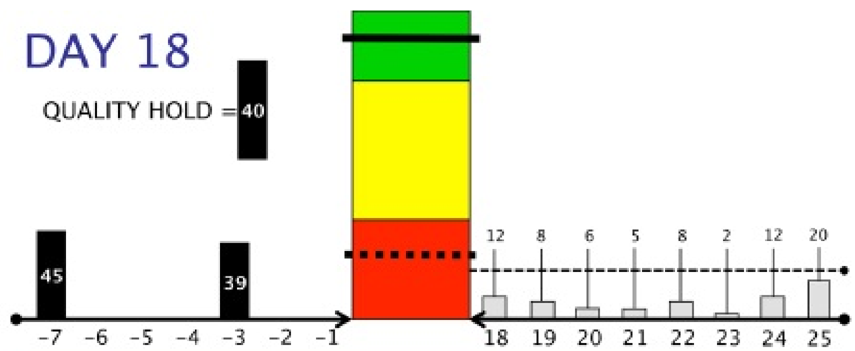

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 92.4%	33	124	12	145	0		52	122	157	7

QUALITY HOLD

Ending OH = 21

FIGURE 9-31 Simulation day 18

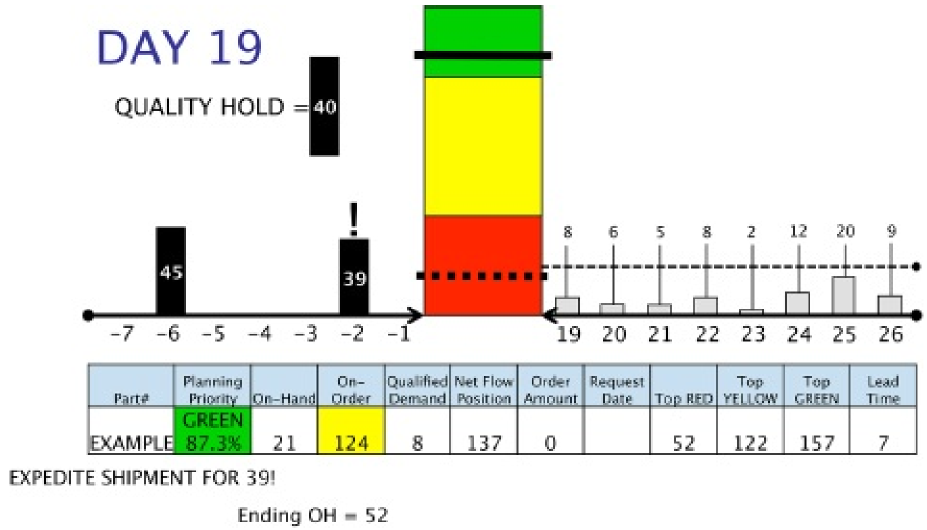

FIGURE 9-32 Simulation day 19

Figure 9-33 shows the continuation of the simulation on day 20. There is no recommended supply quantity, as the net flow position is green at 131 (83.4 percent of top of green). Ending on-hand will be 46.

Figure 9-34 shows the continuation of the simulation on day 21. Starting on-hand is at 86! What happened? The amount on quality hold was moved to on-hand status. On-order is at 45, on-hand is at 86, and qualified demand is only 5. There is no recommended supply order, as the net flow position is green at 126 (80.3 percent of top of green). Ending on-hand will be 81.

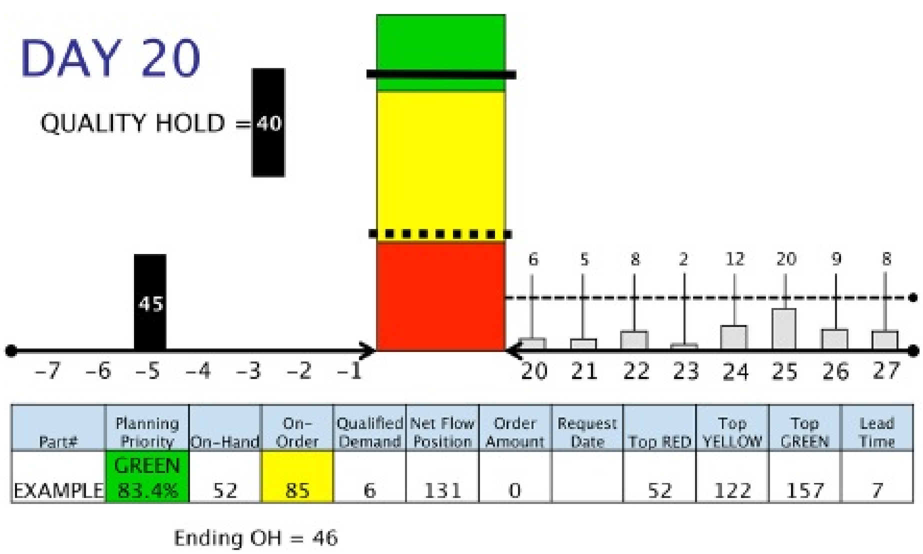

FIGURE 9-33 Simulation day 20

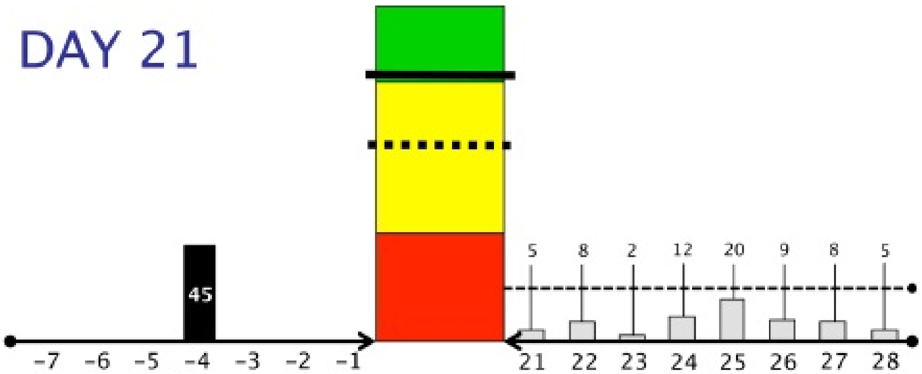

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	GREEN 80.3%	86	45	5	126	0		52	122	157	7

QUALITY HOLD Cleared!

Ending OH = 81

FIGURE 9-34 Simulation day 21

What can be learned from this simulation? First, that the buffers are robust. If properly managed, these buffers can handle many forms of variability such as upticks in demand, supply problems, and drop-in orders, all of which were present in the simulation. Figure 9-35 shows a run chart for net flow and on-hand positions over the simulated 21 days. Each time the net flow position drops into the yellow, it is immediately restored to green by launching a supply order. This is crucial to keeping the buffer properly supplied by maintaining the flow of recommended amounts. Lags and lapses in launching these supply orders will only present additional variability to the buffer and can compromise the ability to maintain the integrity of the decoupling point.

This buffer performed well considering the challenges that were thrown at it. Consider that demand for this period was over 20 percent higher than what the buffer was built for. Over these 21 days, demand averaged 12.4, while the buffer was built for an average daily usage of 10. This buffer should definitely be adjusting up, as this uplift in demand begins to raise the calculated ADU.

This simulation also begins to show us what we can expect from on-hand performance over time with regard to DDMRP buffers. Figure 9-36 is a distribution chart showing the frequency of on-hand value over the simulated 21 days. The × axis is in bin values of 10. Over the course of the simulation, there were two on-hand occurrences between 21 and 30; thus we see a bar with a data label of 2 in the 30 position on the × axis. We see a single uniform distribution (as opposed to a bimodal distribution) ranging in value from 21 to 92. On 12 of the 21 days, the on-hand value was above the top of the red zone. Average on-hand over the 21-day period was at 59.24.

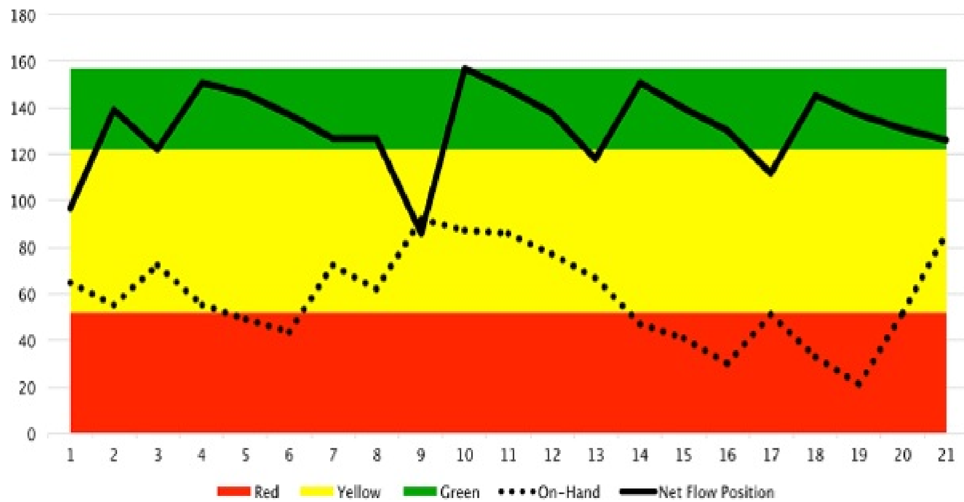

FIGURE 9-35 Reviewing the simulation results

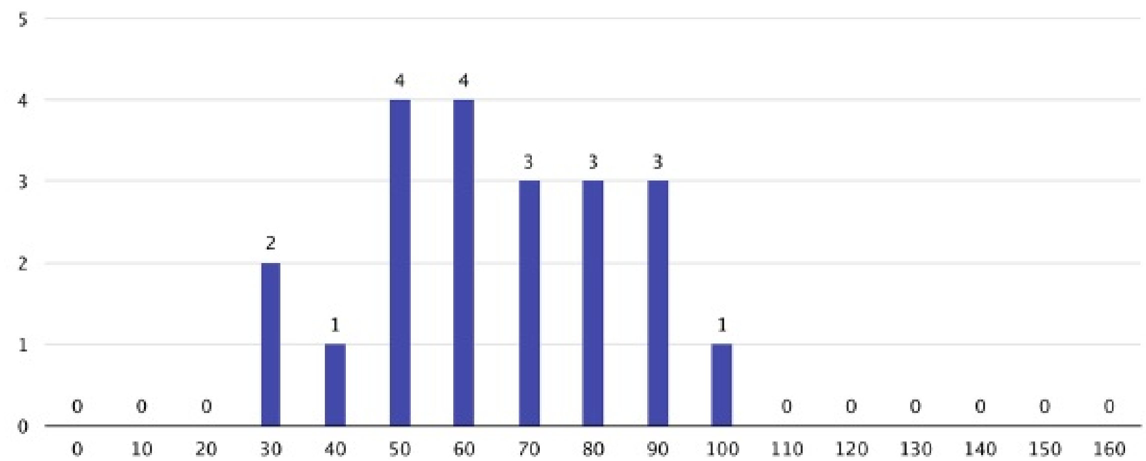

Day	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
On-Hand	65	55	72	55	49	44	72	62	92	87	86	77	67	47	41	30	51	33	21	52	86

AVG = 59.24

FIGURE 9-36 On-hand position frequency over 21 days

Calculating Average On-Hand Inventory

Understanding how the different components of the net flow equation work together is crucial to grasp how to calculate the average on-hand inventory position. When the net flow equation yields a quantity in the yellow zone, a supply order will be recommended for a quantity to bring the net flow position up to the top of the green zone. The green zone defines the average order frequency when the quantity is divided by the average daily usage.

Figure 9-37 depicts the buffer levels for a sample part to demonstrate the average on-hand calculation. With a green zone of 60 and an average daily usage of 10, the average order frequency would be 6 days. The average order quantity would be 60. That means in a perfectly average world, 60 would be ordered every 6 days. Each supply order would be due 6 days apart. If the part has a lead time of 18 days, then we would expect to see three open supply orders at any one time.

Figure 9-38 shows the starting situation for our example. Starting on-hand inventory is 110. The net flow position is at the top of yellow (280). A supply order is being recommended for 60 due in 18 days (day 19 if today is day 1). Open supply shows 180 units, and yet we are launching a new order. That means that an order will be received at the end of today. That order was launched 18 days ago. Tomorrow there will still be three supply orders in the "On-Order" column including the new order.

Figure 9-39 represents the order history of this part, assuming the average order size and frequency of the last 18 days. A supply order is due today (represented in the "–1" column). Two additional supply orders are expected—one for 60 due in 7 days and another for 60 due in 13 days.

Example Part Buffer For Average On-Hand Calculation			
Average Daily Usage	10	Green Zone	**60**
Buffer Profile	M, S (.33), H (.66)		LT Factor: 60 (DLT (18) x ADU (10) x Lead Time Factor (.33))
MOQ	None		Minimum Order Quantity: N/A
Imposed or Desired Order Cycle (DOC)	None		Order Cycle: N/A
Decoupled Lead Time (DLT)	18 days	Yellow Zone	**180** (18(DLT) x 10(ADU))
		Red Zone	**100** (Red Base (60) + Red Safety (40))
			Red Base: 60 (DLT (18)x ADU (10) x Lead Time Factor (.33))
			Red Safety: 40 (Red Base (60) x Variability Factor (.66))

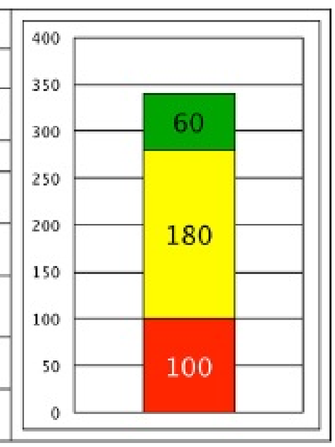

FIGURE 9-37 A sample part

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
ON-HAND EXAMPLE	YELLOW 82.4%	110	180	10	280	60	Day 7	100	280	340	18

FIGURE 9-38 Today's net flow position for the on-hand calculation example

Tomorrow the newly generated supply order will be visible at –18. Figure 9-40 shows the order history on day 2. A supply order was received yesterday. The next supply order is due 5 days from day 2 on day 7.

Returning to our situation today, on-hand will be consumed throughout the day, bringing the ending on-hand to 100. The order for 60 will be received at the end of today and will be reflected in the beginning on-hand inventory for tomorrow (160). Each day until the new supply order is received, on-hand will drain by 10 units. If on-hand inventory starts at 160 on day 2, it will end day 2 at 150. On day 3 on-hand will drop to 140. On day 4 it will drop to 130. On day 5 it will drop to 120. On day 6 it will drop to 110. On day 7 on-hand will drop to 100 and then be restored to 160 on day 8, as the supply order was received for 60 at the end of day 7.

Figure 9-41 depicts this recurring behavior, assuming average demand, order size, and frequency. The on-hand position ranges between a high of 160 and a low 100. Day 1 shows the daily demand drain from 110 to 100. A supply order is received, pushing on-hand up to 160. Demand then drains off for 6 consecutive days back down to 100 before the next supply order is received that day. At the end of day 19, the supply order is received that was created on day 1.

Day	–18	–17	–16	–15	–14	–13	–12	–11	–10	–9	–8	–7	–6	–5	–4	–3	–2	–1
Order						60						60						60

FIGURE 9-39 Supply order history

Day	–18	–17	–16	–15	–14	–13	–12	–11	–10	–9	–8	–7	–6	–5	–4	–3	–2	–1
Order	60						60						60					

FIGURE 9-40 Order history on day 2

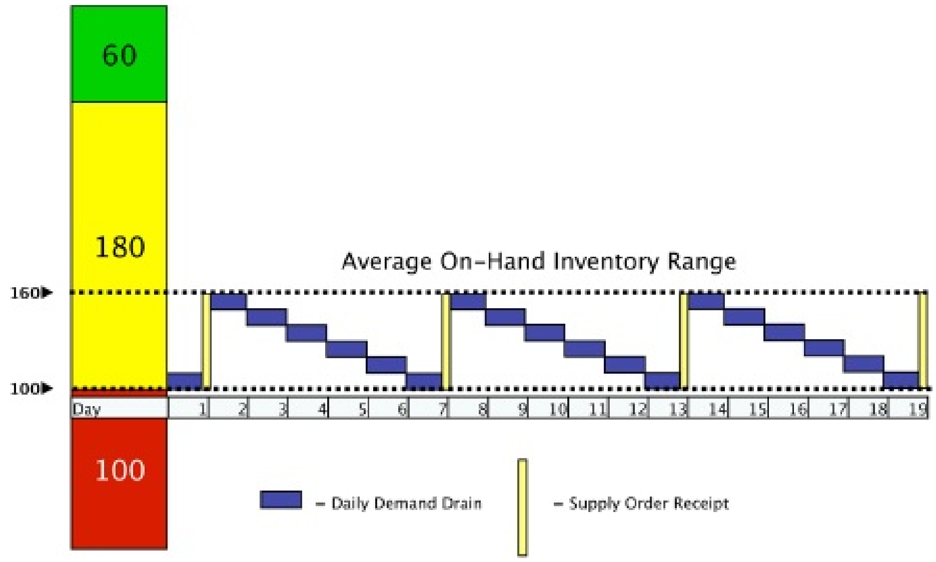

FIGURE 9-41 On-hand position assuming average demand, order size, and order frequency

This exercise gives insight into some interesting characteristics about DDMRP buffers with regard to predicting inventory and open supply ranges.

Average On-Hand Range

With this example a range of expected on-hand is established. This range will become the foundation for making predictions about buffer behavior as well as measuring the performance of a particular replenished part over a past time frame. In DDMRP this range is called the average on-hand range. The average on-hand range corresponds directly to two zones of the buffer: the red and the green. The lower limit of the range is the top of red. The upper end of the limit is the top of red value + the green value. In our on-hand example, that range is 100 (lower) and 160 (upper).

In the beginning of Chapter 7, we discussed the value of inventory and the need to be able to find a range that keeps us balanced between the negatives of having too much and having too little. Figure 9-42 is a duplicate of Figure 7-1, showing the Taguchi loss function of inventory value and the conceptual optimal range. There is a loss of value as we move toward the extremes of too little or too much and outside of an optimal range. We now have a way to calculate what the optimal on-hand range is.

To show how this works, the part will be used that was simulated earlier in this chapter for 21 days. Its top of red zone value was 52. The yellow zone was

70, and the green zone was 35. The optimal on-hand range is the value of the green zone. On either side of the green are the warning ranges. To the left is the low-level warning. That value is the value of the red zone (52). To the right is the high-level warning and is the remaining amount of the yellow zone (35) when the green zone is subtracted. The loss function reaches too little (point A) at zero on-hand and too much (point B) at 122. The 122 is the top of the yellow zone. This is graphically represented in Figure 9-43.

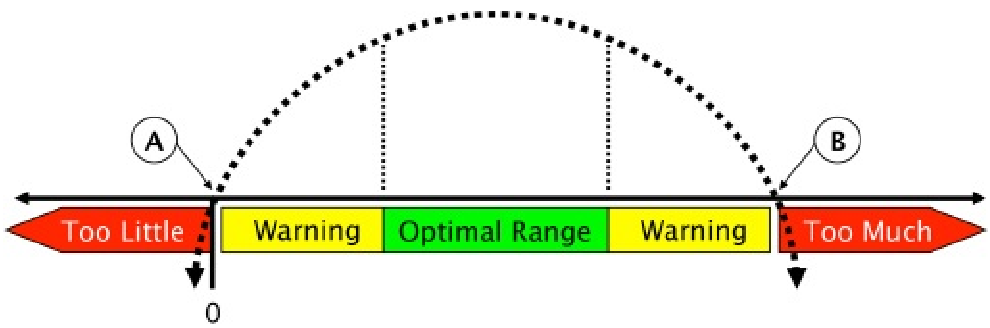

FIGURE 9-42 Former Figure 7-1 showing the optimal zone and Taguchi inventory loss function

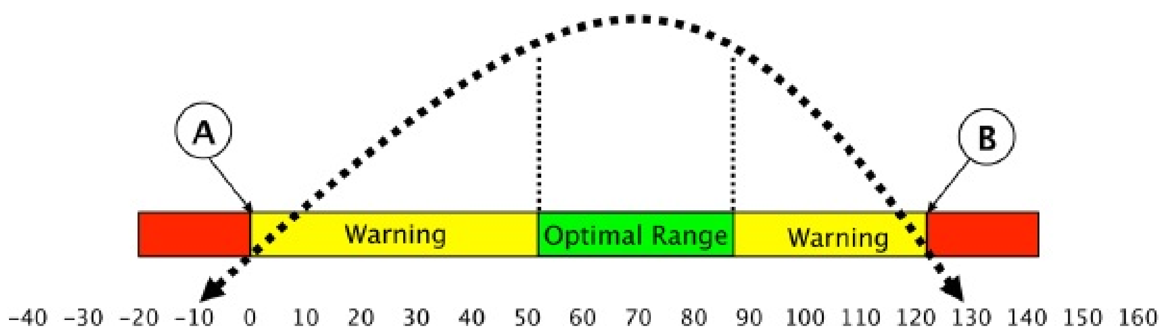

FIGURE 9-43 Inventory loss function for the simulated part's values

When considering past performance, this green range can tell us how well the buffer has performed. Too many severe penetrations below the range can tell us that the buffer's safety was frequently required. Too many on-hand positions above the green range will signal that the buffer might be oversupplied or overstated relative to the actual requirement. Figure 9-44 shows a run chart of the 21-day simulation with regard to the on-hand position on the on-hand range. The run chart displays on-hand values and is converted to a Taguchi loss function view color scheme. The green is still the green zone of 35. It has been stacked on top of the red zone. On the upper side of the green zone is the remaining amount of the yellow zone of 35 (green zone -yellow zone). If on-hand is at zero or less than zero, that is in the "Too Little" range. If on-hand is over the top of yellow value (122), that is in the "Too Much" range.

As we can see, the buffer's on-hand position spent almost the entire first half of the simulation within the optimal on-hand range. Drop-in orders and quality holds created deep penetrations into the low warning range. The on-hand position never went above the optimal range.

Judging past performance is discussed in more depth in Chapter 12.

FIGURE 9-44 Run chart against optimal and warning ranges

Average On-Hand Target

Now that the average range is established, then an average amount can be calculated. In DDMRP this is called the average on-hand target. It is in middle of the range. The equation to calculate the average on-hand target is the top of red value + green zone value/2. In our average on-hand example part, the top of red is 100, and half the green zone value is 30. Thus the average on-hand target is 130. Figure 9-45 now has the average on-hand target illustrated as the solid line through the middle of the average inventory range.

This equation is mathematically similar to an inventory calculation made under the conventional approach, where on-hand is supposed to average out to safety stock + one-half the minimum order quantity. This conventional equation, however, rarely produces a reasonable approximation over longer periods due to MRP's susceptibility to nervousness, variability of demand, inventory inflation typically caused by safety stock ordering, and shortages created by supply continuity variability.

Understanding how to calculate the average on-hand target will now allow the Company ABC example from Chapters 7 and 8 to be finished. With this equation the working capital implications of the buffers can be judged. This is covered later in this chapter.

Average Open Supply Orders

The yellow zone divided by the green zone is the average number of open supply orders. In this example there should be three open supply orders on each day. If the green zone was 90 instead of 60, then there should be an average of two open supply orders at any one time.

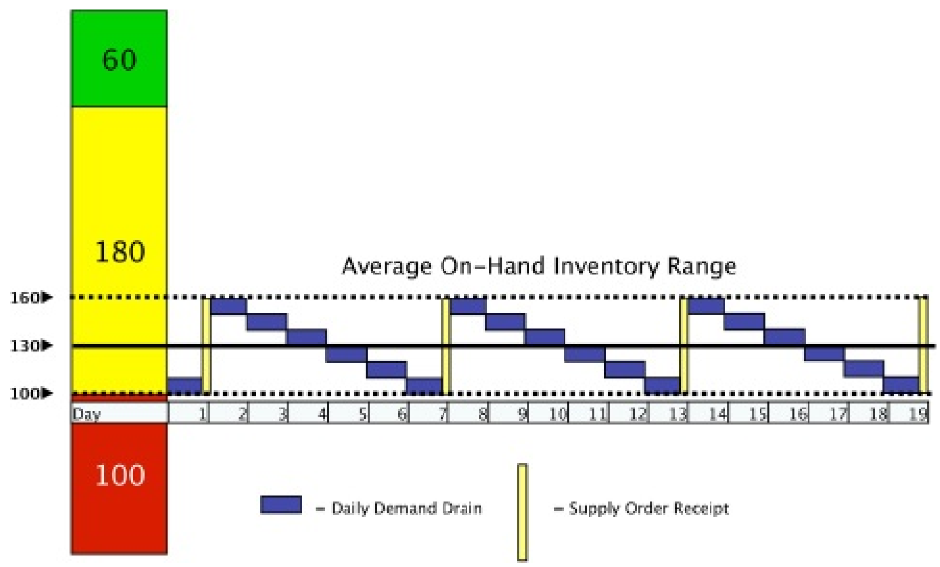

FIGURE 9-45 Average on-hand target of 130

This makes mathematical and intuitive sense when you consider that the purpose of the yellow zone is to provide the heart of the demand coverage within one full lead time. The yellow zone is the pipeline or conveyor belt portion of the buffer. It extends the buffer back into the supply chain. As long as orders are approved when the net flow equation calls for resupply, that pipeline will remain intact.

Of course, these equations are built to averages. Within shorter windows of time, they will not be accurate due to variability. With the 21-day planning simulation, there was plenty of variability within the simulation window. For example, the average on-hand target would be calculated as 69.5 [52 (top of red value) + 35 (green zone value)/2]. Yet at the end of 21 days, the average on-hand amount was 59.24. This was, of course, due to heavier demand than expected and supply delays. Averages are, however, valid to use over longer periods of time—the strategic relevant range. Assuming the model is set and maintained properly, there is a reasonable prediction about on-hand levels to be realized by that model over that longer range.

Decoupled Explosion

When considering decoupling and the DDMRP supply order generation process, an obvious impact emerges. When a supply order is generated at a higher level, decoupling stops the explosion of the bill of material at decoupling points placed at lower levels. The explosion can and will be stopped because that decoupling point is buffered. The net flow equation is then independently calculated at that point. Supply order generation only continues if that position's net flow equation calls for resupply. Under that condition, the explosion then begins again relative to that part's respective components.

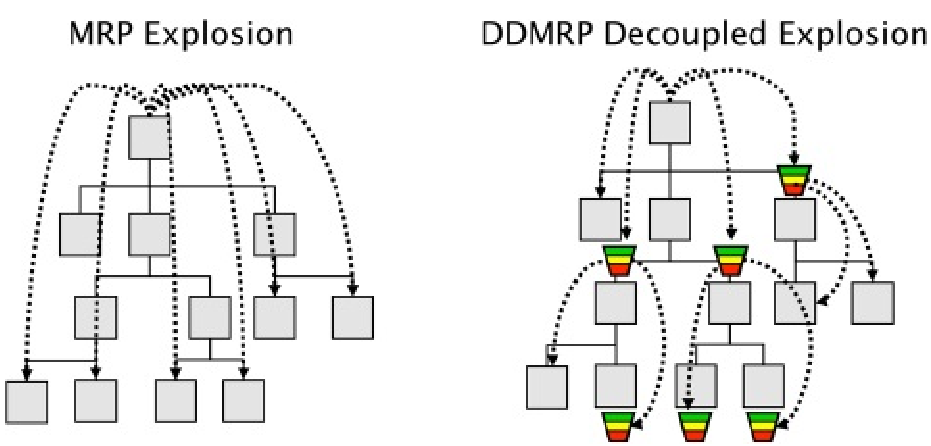

FIGURE 9-46 Explosion differences between MRP and DDMRP

Figure 9-46 illustrates the explosion differences between conventional MRP and DDMRP. MRP uses a full "requirements explosion." APICS defines a requirements explosion as:

> The process of calculating demand for components of a parent item

> requirements by the component usage quantity specified in the bill of material. (p. 149)

This calculation continues through the bill of material over the entire planning horizon. Thus the word "explosion" relates to dependent relationships between a parent item and its component for a projected period of time.

The explosion on the left of Figure 9-46 represents conventional MRP, where any demand at the higher level is typically driven all the way through to the purchased level and projected into the future. There are some exceptions to this rule in MRP, but they are simply that—exceptions. For example, MRP will stop the explosion within a particular leg if there is sufficient on-hand stock to cover the demand requirement in that time period. This occurrence, however, is more happenstance rather than any real plan. It tends to happen only if there is residual on-hand inventory remaining due to changes in schedule or order multiple differences. Remember, MRP fundamentally wants to net all positions to zero projected available balance (for positions with safety stock, the safety stock level is the new zero level).

DDMRP utilizes a "decoupled explosion" depicted on the right-hand side of Figure 9-46. A decoupled explosion is a critical distinguishing characteristic of a DDMRP system. The term itself is an oxymoron. It means "independent dependence." Yet that is exactly what is occurring. This concept is crucial in preventing nervousness because most changes at high-level parents will not be big enough to pass through the buffers at the decoupling points. This means that flow is largely protected against the system nervousness that is transferred and amplified in conventional MRP. This is especially true for decoupling points placed at common components (a common strategy) as the smoothing benefit of aggregation against a bigger calculated buffer level is realized.

Can conventional MRP systems decouple explosions? There are at least three tactics that should be explored with regard to this objective, but all come with a price.

1. Safety stock could be added at desired decoupling points in a sufficient quantity to guarantee on-hand quantities. However, this will create expedite requirements as on-hand balances drop below safety stock. This essentially means there is no real decoupling. The signal is transferred and amplified through the bill of material.
2. A stop explosion order can be used. Many conventional MRP systems allow the deployment of a "stop explosion flag" or an "externally

planned" setting for specifically designated parts. In this case the explosion can definitely be stopped at that position, but there are two subsequent challenges. First, the buffer levels must be constructed at that position, and second, the net flow equation must be employed to restart the explosion. Conventional MRP systems can't automatically restart the explosion, and they most certainly won't restart the explosion according to a net flow position based on a DDMRP buffer construct. This means that it will fall to planners to manually restart the explosion for each decoupled intermediate and purchased position. This has been the scenario for many early DDMRP implementations.

3. A multilevel master production schedule process can be used. In this case an explosion can be stopped and then restarted, but at what cost? Setting up and effectively maintaining multiple master production schedules and their connections will prove extremely difficult for most planning teams. Furthermore, the net flow equation must still be performed against the DDMRP buffer levels between these different master production schedules, and there is no mechanism in conventional MRP systems to effectively do that.

Yet MRP and DDMRP are not always different. Within certain parameters MRP and DDMRP are identical. Figure 9-47 depicts an area in the example in which MRP and DDMRP behave exactly the same way. This area is highlighted within the box on both explosions. Thus the concept of a decoupled explosion illustrates how DDMRP successfully combines both the dependence of MRP and the independence of strategic decoupling required for an effective solution for today's volatile and uncertain environments. There is independence at the decoupling points, but between decoupling points there is dependence. That dependence between decoupling points is no different from conventional MRP. That is why there is still MRP in DDMRP. A decoupled explosion is a cornerstone of the planning mechanism in a DDMRP system and allows the flow of relevant information and materials to be promoted and protected.

So does this mean that with DDMRP, everything will go according to plan? Will the manufacturing lead times always be a perfect reflection of how long it will take to move material through a decoupled lead time chain? Of course not, but that is why there is a buffer, in particular the red zone portion within the buffer. This brings to light another very important distinction between MRP and DDMRP. In MRP, because there is no decoupling, everything must go *exactly* according to plan in order to get the predicted result. In DDMRP the assumption is that almost nothing will go exactly according to plan. DDMRP is roughly right, where MRP is precisely wrong.

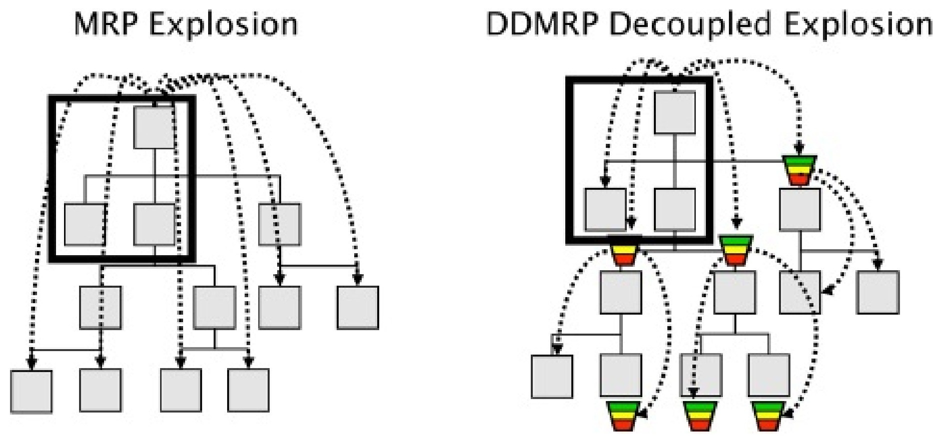

FIGURE 9-47 MRP and DDMRP explosion similarities

Figure 9-48 depicts a product structure for a product called FPD. If today's FPD net flow equation calls for a resupply (below top of yellow) of 25 pieces, a supply order for 25 will be generated for delivery in seven days (the length of the decoupled lead time). This creates a demand for an FPD parent order release six days from today (the decoupled lead time minus the manufacturing lead time of FPD). The component requirements involved in this parent order demand are for components 208 and 210. Component 208 is buffered, and so it is assumed to be available, and no further explosion happens down that leg. The shaded rows in the figure represent the buffered components (208, 410P, 412P) and define where the demand explosion will stop.

Figure 9-48 also depicts the decoupled explosion that drives through the nonbuffered items; however, as shown in the table below the product structure, 210 is unbuffered. Its manufactured lead time is two days. An order release will be created four days from today (FPD decoupled lead time minus 310 manufacturing lead time). This will then generate an immediate release requirement for component 310.

In summary FPD requires 208 and 210 on day 6. Part 210 requires 310 on day 4. Part 310 requires 410P and 412P on day 1. An order source is also available showing what specific component(s) drove the requirement. This is called single-level pegging in MRP and is still very much applicable in DDMRP between decoupling points.

There will always be an immediate release at the terminus of the parent's decoupled lead time chain unless there is residual on-hand inventory that is sufficient to stop the explosion. That immediate release requirement is based on the assumption that the stock is available at the decoupling point and the decoupled lead time for the high-level parent responsible for the original demand will be honored.

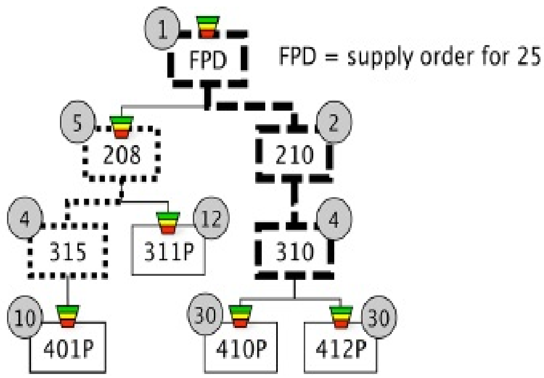

Part #	Demand Allocation	Residual On-Hand	Net Item Requirement	Order Source	Requirement Date	Order Release Date
208	25			FPD	Day 6	Day 6
210	25	0	25	FPD	Day 6	Day 4
310	25	0	25	210	Day 4	Now
410P	25			310	Now	Now
412P	25			310	Now	Now

FIGURE 9-48 A decoupled explosion example

To illustrate how an explosion would stop within the decoupled lead time chain, we will allow for residual on-hand inventory within the example. Figure 9-49 illustrates a situation in which FPD generates the same demand requirement for 25 units. The demand is passed through 210 (evident in the "Net Item Requirement" column), but the explosion stops at 310 because there is residual on-hand inventory. Thus no demand allocation is passed through to 410P and 412P.

What about components being protected against large parent demand spikes? If parent item spikes are not qualified in advance, this could be a big issue for lower levels and further reinforces the need for order spike qualification at the end item level. In the FPD example, the order spike horizon would be at least seven days. If the spike generates a resupply requirement, then 310 would get an immediate requirement for release.

Part #	Demand Allocation	Residual On-Hand	Net Item Requirement	Order Source	Requirement Date	Order Release Date
208	25			FPD	Day 6	Day 6
210	25	0	25	FPD	Day 6	Day 4
310	25	50	0	210	Day 4	Now
410P	0			310		
412P	0			310		

FIGURE 9-49 A decoupled explosion example with residual on-hand inventory

Will 410P and 412P have enough stock on hand to cover the immediate release? Of course, that depends on the size of the spike. At some point no stocking solution will be sufficient to cover enormous parent spikes. But if 410P and 412P are common components, then their respective buffers are built upon a much larger amount of consolidated demand producing a relatively larger buffer. This is evident in the Company ABC example with regard to component 201. This often means that a large spike from a single parent perspective is not a large spike from a component perspective. Additionally, FPD does have an amount of safety if the spike outstrips the on-hand levels of 410P or 412P. The propensity and size of the spikes on FPD determine how that red zone safety level was calculated in the first place..

Figure 9-50 illustrates an FPD demand of 125, qualifying as a spike against the buffered 208 and 412P positions. In the case of 208, the demand is not due for six days but has qualified as a spike against the 208 order spike threshold. That means it will be included in the net flow equation for 208 until it is satisfied. Will it generate a requirement for resupply against the 208 buffer? Perhaps, but that net flow equation is determined separately. That is the entire point of the decoupled explosion.

Figure 9-50 also shows a net requirement of 75 being passed through the 310 position to 410P and 412P. The 50 on-hand quantity for 310 is netted from the demand of 125. For 412P an immediate release of 75 is required, and the demand represents a spike. Will it generate a supply? As with 208 it could generate a supply order. It is definitely a qualified demand in 412P's net flow equation for today.

Let's add an additional layer of complexity to the example to illustrate the similarities of MRP and DDMRP within the decoupled lead time chain. Figure 9-51 shows another parent component generating demand for the common components 210, 310, 410P, and 412P. In this example FPD is calling for a resupply of 25, and FPE is calling for a resupply of 60. This means the gross requirement for 210 is 85. This is seen in the table under the “Order Source” column showing demand from both sources. Component 310 has 50 in residual on-hand and so only passes a netted requirement for 35 to 410P and 412P each.

All of this serves to illustrate that conventional MRP and DDMRP given the same inputs essentially behave the same with regard to nonstocked items. The decoupling points define the boundaries of that common behavior.

Part #	Demand Allocation	Residual On-Hand	Net Item Requirement	Order Source	Requirement Date	Order Release Date
208	125 SPIKE!			FPD	Day 6	Day 6
210	125	0	125	FPD (125)	Day 6	Day 4
310	125	50	75	210	Day 4	Now
410P	75			310	Now	Now
412P	75 SPIKE!			310	Now	Now

FIGURE 9-50 An order spike for 208

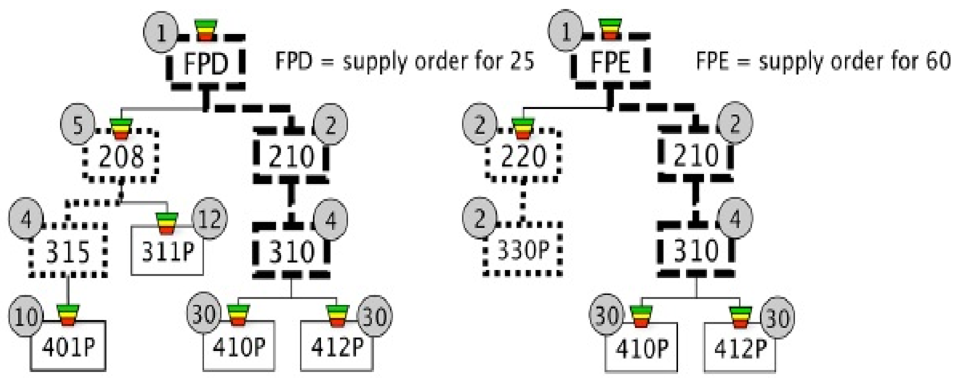

Part #	Demand Allocation	Residual On-Hand	Net Item Requirement	Order Source	Requirement Date	Order Release Date
208	25			FPD	Day 6	Day 6
220	60			FPE	Day 6	Day 6
210	85	0	85	FPD (25), FPE (60)	Day 6	Day 4
310	85	50	35	210	Day 4	Now
410P	35			310	Now	Now
412P	35			310	Now	Now

FIGURE 9-51 Gross requirement for 210

Hybrid Model Supply Order Generation

The hybrid distribution model was introduced in Chapter 6. It involved a mixed model of hub-and-spoke configuration for slower-moving items and make-and-ship on fast-moving items. Figure 9-52 is the hybrid model from Chapter 6.

Products 1, 2, and 3 have buffers at the spoke only, while products 4 and 5 are in the hub-and-spoke configuration.

Generating supply orders for the slow-mover parts in the hub-and-spoke configuration is a straightforward application of the net flow equation at both the spokes and the hub. But how to plan and generate supply orders for the fast movers when there is no clear signal provided by a central buffer? Remember that these fast movers are shipped to the forward distribution locations from the manufacturing plant upon completion. In most cases there will be signals for resupply on fast movers from the spokes at various times and for various quantities. The average order size for a specific forward location is by far smaller than the minimum run size at the sourcing unit. How should the sourcing unit determine when is the right time to run each fast mover?

Figure 9-53 shows the hybrid configuration for product 1 across all locations. There is no central buffer, only spoke buffers. The zone values of the product 1 buffer at each location are given in Figure 9-53. Notice the far right column under "Total." In this case the green zones across all locations have been sized to total the sourcing unit minimum order quantity of 10,000. Each location has a different lead time that is calculated by the plant lead time (seven days) + transportation time. Locations 1 and 4 are three days away from the plant, while locations 2 and 3 are one and two days away respectively.

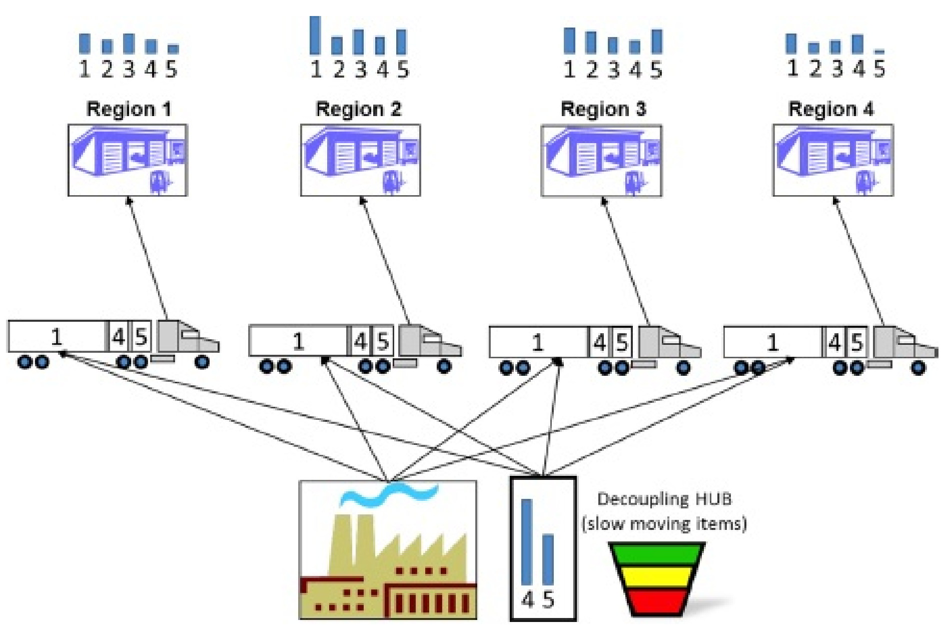

FIGURE 9-52 The hybrid model from Chapter 6

Zone Values	Location 1	Location 2	Location 3	Location 4	Total
Green	4,000	2,000	2,000	2,000	10,000
Yellow	6,500	3,400	2,925	3,500	16,325
Red	4,875	2,550	2,194	2,625	12,244
Lead time	10	8	9	10	
ADU	650	425	325	350	1,750
Sourcing unit MOQ = 10,000					

FIGURE 9-53 Buffer levels across all locations for product 1

The average daily usage for each location is displayed. The total network average daily usage is the summation of all locations' usage. Each location has a different number of days' safety contained in its red zone (location red zone/location ADU). Location 1 has 7.5 days. Location 2 has 6 days. Location 3 has just under 7 days. Location 4 has 7.5 days. In the aggregate the total safety across the network is 7 days [12,244 (total of red zones)/1,750 (total ADU)].

Supply order generation for a fast-moving product must consider the aggregate net flow position across a network. Figure 9-54 displays the situation for product 1 on an example day. The column labeled "NFP" is the net flow position at each location. The "Priority" column is the percentage of the net flow position to the top of green. It is color-coded based on its position relative to the top of yellow. Below the TOY value yields a yellow coding; above the TOY value yields a green coding. The "Order" column is the amount needed to restore each location's buffer to the top of its green zone. The total amount to restore all buffers to the top of green is 10,851. This is greater than the sourcing unit MOQ of 10,000. A supply order is generated despite two locations' net flow position being green.

Location 1					
Product	Priority	NFP	TOG	TOY	Order
1	YELLOW 66.5%	10218	15375	11375	5157
Location 2					
Product	Priority	NFP	TOG	TOY	Order
1	GREEN 78.0%	6200	7950	5950	1750
Location 3					
Product	Priority	NFP	TOG	TOY	Order
1	YELLOW 58.7%	4176	7119	5119	2943
Location 4					
Product	Priority	NFP	TOG	TOY	Order
1	GREEN 87.7%	7124	8125	6125	1001
Total Demanded					10851

FIGURE 9-54 Product 1 priority (example 1)

Location 1					
Product	Priority	NFP	TOG	TOY	Order
1	GREEN 87.5%	13451	15375	11375	1924
Location 2					
Product	Priority	NFP	TOG	TOY	Order
1	YELLOW 68.1%	5415	7950	5950	2535
Location 3					
Product	Priority	NFP	TOG	TOY	Order
1	GREEN 81.7%	5817	7118.75	5118.75	1301.75
Location 4					
Product	Priority	NFP	TOG	TOY	Order
1	GREEN 87.7%	7124	8125	6125	1001
Total Demanded					6762

FIGURE 9-55 Product 1 priority (example 2)

Figure 9-55 shows another example in which a supply order would not be generated. Despite location 2 having a yellow planning priority, the total aggregate demand against the sourcing unit is less than the MOQ.

[This example demonstrates the use of net flow position across a specific group in order to deal with a specific limitation. In this case the group was the distribution locations for a specific product, and the limitation was the sourcing unit MOQ. This example opens the door to a crucial concept with regard to DDMRP supply order generation in relation to limitations associated with those groups. This concept is called prioritized share.

Prioritized Share Supply Order Consideration

Prioritized share is a supply order generation schema applied to a group based on the relative net flow positions within that group and the imposition of a constraint or threshold. The types of group can include within a location, across locations, by supplier, and by product structure or shared component. There are three main uses of prioritized share: discount optimization, freight optimization, and coverage optimization. Each use will be explored.

Discount Optimization

Suppliers frequently offer discounts for meeting certain conditions. These discounts often represent a significant economic advantage to the buying entity if properly managed. Examples might include free freight or a percentage discount for orders that meet a given threshold. This threshold might be that the total dollar amount of the order must reach a specified minimum or that the total order must fill a truck..

One example: Let's say a company buys electronic components from a supplier called SuperTech. SuperTech has a policy that it will pay the freight costs if an order fills a truck. The order does not have to be for just one item. A mixed load can be ordered, and as long as the entire truck is filled, SuperTech will pay the freight. This is a real economic benefit to the customer, but the customer needs to take advantage of it in a way that penalizes it the least with regard to inventory.

There are 30 total pallet positions in a large trailer. The minimum order quantity and multiple for each part is one pallet. This customer buys 12 different components from SuperTech. Many items have green zones sized to their respective minimum order quantities (one pallet). Figure 9-56 shows the buffers of all SuperTech parts. Each part's net flow position is indicated by the black diamond symbol. There are five parts with net flow positions in yellow. Are these five parts enough to fill a truck?

Figure 9-57 is the planning screen for SuperTech parts. The screen is first sorted by priority. Priority is determined first by color and then by percentage of net flow penetration in relation to the top of green. The "Order" column is simply how much quantity is required to move the net flow position to the top of green. The five parts calling for resupply are evident at the top of the planning screen (parts 10, 12, 11, 6, and 9).

Two columns have been added to the planning screen to help determine how to take advantage of the SuperTech free freight offer. "Qty per Pallet" is the number of each part per pallet. "Pallets Req" is the number of pallets needed to restore the net flow position to the top of green ("Order"/"Qty per Pallet"). The parts calling for resupply only total 28 pallets, but 30 pallets are required for free freight. How can the customer best prioritize the share of the additional two spaces in the trailer?

One option would be to order an extra two pallets of the part numbers calling for resupply. This would put at least one or two of those parts in an OTOG net flow position. Another option would be to take the highest-priority green parts. Parts 1 and 2 have relatively significant green zone penetrations. These are parts that are not calling for resupply but have the deepest net flow penetration into their buffers relative to all other green parts (8, 3, 4, 5, and 7). If the green zone is determined by MOQ, then this will cause an OTOG net flow position as well. That is the case for both parts 1 and 2. If the customer must go over the top of green, what is the best decision to minimize excess inventory liability?

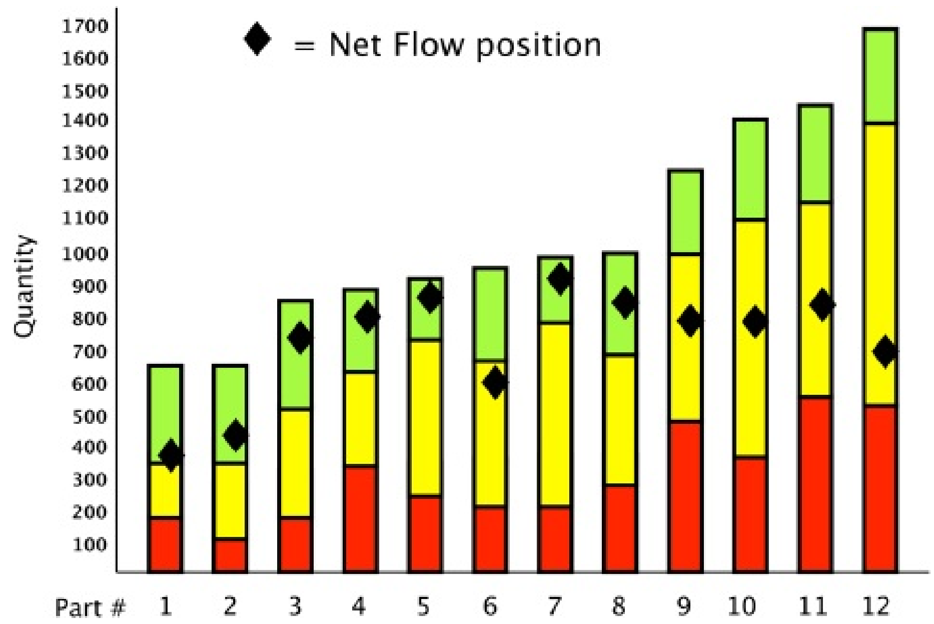

FIGURE 9-56 Purchased SuperTech part numbers and their buffer positions

Part #	Priority	NFP	TOG	Order	Qty/Pallet	Pallets Req
12	YELLOW 41.2%	700	1700	1000	100	10
10	YELLOW 57.1%	800	1400	600	150	4
11	YELLOW 58.6%	850	1450	600	150	4
6	YELLOW 63.2%	600	950	350	50	7
9	YELLOW 64.0%	800	1250	450	150	3
1	GREEN 57.7%	375	650	0	300	0
2	GREEN 69.2%	450	650	0	300	0
8	GREEN 85.0%	850	1000	0	300	0
3	GREEN 88.2%	750	850	0	300	0
4	GREEN 88.9%	800	900	0	250	0
5	GREEN 94.4%	850	900	0	150	0
7	GREEN 95.0%	950	1000	0	200	0

FIGURE 9-57 Planning screen for SuperTech parts

An easy way to determine the minimal over top of green liability of adding two additional pallets to the order is to compare the net flow percentage with one additional pallet for all part numbers. Figure 9-58 shows the impact of ordering an additional pallet than is required for each specific part. The column “Pallets +1” shows an additional pallet for all items. The column “NFP (+ 1 Pallet)” is the net flow position with the additional pallet quantity.

The column “Priority (+ 1 Pallet)” is the net flow position after the additional pallet quantity has been figured in divided by the top of the green zone for each part. When over 100 percent, any priority column should display OTOG and be shaded a light blue. The column “OTOG%” is the percentage amount over the top of green. This can then be easily sorted to show the two items with the smallest OTOG% impact. Figure 9-59 shows the parts sorted by the least amount of OTOG impact on each of their respective buffers.

This analysis shows that ordering an additional pallet of part 6 and ordering one pallet of part 1 will result in the least amount of over top of green liability. Thus the prioritized share approach provides a simple and quick way to take advantage of the free freight offer while minimizing the impact on excess inventory.

The same type of prioritized share analysis can be performed for minimum spend thresholds. Let’s use the same example but with an alternative SuperTech policy. In this case SuperTech offers free freight for orders above $19,000. Figure 9-60 shows the dollar value per pallet of each SuperTech part and the total spend for the items calling for resupply. One pallet is the minimum order quantity and multiple for each item. The five items calling for resupply only total $18,025 of the $19,000 spending threshold. Prioritized share will be used to help determine the best way to spend the remaining amount to meet the threshold while minimizing inventory liability.

Part #	Pallets +1	Qty per Pallet	NFP (+ 1 Pallet)	Priority (+ 1 Pallet)	OTOG%
12	11	100	1,800	OTOG 105.9%	5.9%
10	5	150	1,550	OTOG 110.7%	10.7%
11	5	150	1,600	OTOG 110.3%	10.3%
6	8	50	1,000	OTOG 105.3%	5.3%
9	4	150	1,400	OTOG 112.0%	12.0%
1	1	300	675	OTOG 103.8%	3.8%
2	1	300	750	OTOG 115.4%	15.4%
8	1	300	1,150	OTOG 115.0%	15.0%
3	1	300	1,050	OTOG 123.5%	23.5%
4	1	250	1,050	OTOG 116.7%	16.7%
5	1	150	1,000	OTOG 111.1%	11.1%
7	1	200	1,150	OTOG 115.0%	15.0%

FIGURE 9-58 OTOG percentage with additional pallet for each part

Part #	Pallets +1	Qty per Pallet	NFP (+ 1 Pallet)	Priority (+ 1 Pallet)	OTOG%
1	1	300	675	OTOG 103.8%	3.8%
6	8	50	1,000	OTOG 105.3%	5.3%
12	11	100	1,800	OTOG 105.9%	5.9%
11	5	150	1,600	OTOG 110.3%	10.3%
10	5	150	1,550	OTOG 110.7%	10.7%
5	1	150	1,000	OTOG 111.1%	11.1%
9	4	150	1,400	OTOG 112.0%	12.0%
8	1	300	1,150	OTOG 115.0%	15.0%
7	1	200	1,150	OTOG 115.0%	15.0%
2	1	300	750	OTOG 115.4%	15.4%
4	1	250	1,050	OTOG 116.7%	16.7%
3	1	300	1,050	OTOG 123.5%	23.5%

FIGURE 9-59 Parts sorted by OTOG%

Part #	Priority	NFP	TOG	Order	Qty per Pallet	$ per Pallet	Pallets Req	$ Ordered	NFP After
10	57.1%	800	1,400	600	150	$1,000	4	$4,000	100.0%
12	41.2%	700	1,700	1,000	100	$600	10	$6,000	100.0%
11	58.6%	850	1,450	600	150	$825	4	$3,300	100.0%
6	63.2%	600	950	350	50	$450	7	$3,150	100.0%
9	64.0%	800	1,250	450	150	$525	3	$1,575	100.0%
								$18,025	

FIGURE 9-60 Total dollar amount of parts calling for resupply

The planner will need to find the best combination of additional pallet orders to meet the minimum spend threshold with the least amount of additional spend and the least amount of OTOG% liability. Figure 9-61 is the same as Figure 9-59 but with the additional column of price per pallet (“$ per Pallet”). This column will be used in combination with the “OTOG%” column to select the best combination of parts for additional pallet ordering to satisfy the above conditions.

In this case, the best combination of minimizing additional pallets and OTOG% and meeting the spend threshold with the least amount of additional spend is to order additional pallets of parts 1 and 12. Ordering one additional pallet of each totals $1,000 of additional spend. This will bring the total order to $19,025, only $25 over the required minimum spend. Additionally, this configuration allows for the least amount of OTOG% to capture the free freight.

Part #	Pallets +1	Qty per Pallet	NFP (+ 1 Pallet)	Priority (+ 1 Pallet)	OTOG%	$ per Pallet
1	1	300	675	OTOG 103.8%	3.80%	$400
6	8	50	1,000	OTOG 105.3%	5.30%	$450
12	11	100	1,800	OTOG 105.9%	5.90%	$600
11	5	150	1,600	OTOG 110.3%	10.30%	$825
10	5	150	1,550	OTOG 110.7%	10.70%	$1,000
5	1	150	1,000	OTOG 111.1%	11.10%	$480
9	4	150	1,400	OTOG 112.0%	12.00%	$525
8	1	300	1,150	OTOG 115.0%	15.00%	$1,200
7	1	200	1,150	OTOG 115.0%	15.00%	$1,200
2	1	300	750	OTOG 115.4%	15.40%	$600
4	1	250	1,050	OTOG 116.7%	16.70%	$825
3	1	300	1,050	OTOG 123.5%	23.50%	$300

FIGURE 9-61 Parts sorted by OTOG% with dollar per pallet per part

Freight Optimization

Prioritized share can be used to optimize freight spend from hub to spoke or hub to hub (multihub configuration), assuming that full truckloads will result in better overall freight rates. Just as in the case of the first SuperTech example, the receiving location's (spoke) net flow positions on all replenished items coming from the supplying location (hub) will be analyzed from a prioritized share perspective, and a full truckload will be built to minimize OTOG% liability.

In this case mixed pallets might be a possibility, allowing for multiple green items to be simultaneously brought to top of green status. Depending on the number of green items and the severity of their penetration, this could eliminate any over the top of green ending situations.

Coverage Optimization

The prioritized share schema can be used to allocate scarce quantities of inventory to spoke buffers in order to produce the best aggregate network coverage. This was the first known application of the prioritized share schema.

When there is not enough inventory to cover all of a part's demand across all locations calling for resupply or there is a desire for a matched run out of a discontinued product, prioritized share will step across the locations allocating inventory to the tops of the lowest zones first. First top of red, then top of yellow and then top of green.

This avoids one location that has a deep penetration to take all the inventory, restoring its net flow position to healthy while others remain relatively in trouble. In this application the prioritized share schema is looking to balance all locations with about the same level of protection from a net flow perspective.

Location 1						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	RED 26.8%	4125	15375	11375	4875	11250
Location 2						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	YELLOW 52.1%	4142	7950	5950	2550	3808
Location 3						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	YELLOW 45.1%	3214	7119	5119	2194	3905
Location 4						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	YELLOW 70.7%	5741	8125	6125	2625	2384
Total Demanded						21347
Total Available						15000

FIGURE 9-62 Product prioritized share example

Figure 9-62 shows how scarce supply at a hub is allocated against four locations calling for resupply using the prioritized share. Product 5 has demands for resupply across the network including a red net flow position in location 1. A total of 21,347 units are required to restore all locations to a full net flow position, but there are only 15,000 units available at the hub. If location 1 is allowed to restore its net flow position to 100 percent, it will leave only 3,750 remaining for the rest of the locations. This will result in an unbalanced network from a net flow perspective. Supply orders to all locations will be modified using prioritized share to allocate the 15,000 available to achieve a relatively equal net flow position in all locations.

First, inventory is allocated to bring location 1's net flow position to the top of red. This requires 750 pieces [4,875 (TOR) – 4,125 (NFP)]. This leaves location 1 with the deepest yellow penetration. It has priority for additional inventory share. Now 14,250 pieces remain available. Location 1 has a yellow zone of 6,500. Location 3 (the next deepest penetration in yellow) requires 1,905 to restore it to the top of yellow. Location 2 requires 1,808 to restore it to the top of yellow. Location 4 requires 384 to restore it to the top of yellow. The total requirement to restore all positions to the top of yellow is 10,597. If this demand were less than the total available, prioritized share would look to balance the net flow positions in yellow across the locations.

Since there is enough available inventory to restore all locations to the top of yellow but not enough to restore all to the top of green, prioritized share will need to balance all locations in the green zone. There are 3,653 pieces remaining in available inventory against a total remaining demand of 10,000. The prioritized share schema will balance out each position's priority percentage. Figure 9-63 shows a step-by-step application of prioritized share to this example and the total allocated to each location.

At the end of this sequence, all product 5's net flow positions across all locations are balanced to a priority position of 83.5 percent. Figure 9-64 shows the planning screen with the updated net flow positions.

Total product 5 inventory available		15,000	
Location 1 red allocation		750	
Available remaining		14,250	
Location 1 yellow allocation		6,500	
Location 3 yellow allocation		1,905	
Location 2 yellow allocation		1,808	
Location 4 yellow allocation		384	
Total allocation to TOY (all locations)		11,347	
Available remaining		3,653	
Location 1 green allocation		1,470	
Location 2 green allocation		692	
Location 3 green allocation		828	
Location 4 green allocation		663	
Total green allocation		3,653	
Total inventory allocation (all locations)		15,000	
Location 1	Location 2	Location 3	Location 4
8,720	2,500	2,733	1,047

FIGURE 9-63 Prioritized share sequence

Location 1						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	GREEN 83.5%	12845	15375	11375	4875	0
Location 2						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	GREEN 83.5%	6642	7950	5950	2550	0
Location 3						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	GREEN 83.5%	5947	7119	5119	2194	0
Location 4						
Product	Priority	NFP	TOG	TOY	TOR	Order
5	GREEN 83.5%	6788	8125	6125	2625	0
Total Demanded						0
Total Available						0

FIGURE 9-64 Updated planning screen

Min-Max Supply Order Generation

There are three types of part classifications in DDMRP. To this point this chapter has been exclusively dedicated to replenished parts. Replenished parts are the majority of buffered parts in most DDMRP implementations. The application of the net flow equation is exactly the same for replenished override parts. That leaves only the supply order generation considerations for min-max parts. Min-max parts are nonstrategic buffered positions. The parts assigned to min-max status tend to be low variability and readily available, but they still represent decoupling points.

The daily application of the net flow equation applies to min-max positions. Using the min-max example from Chapter 7, the use of the net flow equation will be demonstrated with a min-max position. In Figure 9-65 the starting on-hand quantity is 40. There are 70 on-order units and 32 in qualified demand. The net flow position is 78 (40 + 70 – 32). A supply order will be generated for 72 units.

Completing the Company ABC Example

The example that was started in Chapter 7 can now be completed. You now have the capability to calculate the average on-hand quantity for each buffer. This will allow us to judge the merits of different decoupling position options. This comparison is made from two perspectives: quantity and cash. The average on-hand target equation yields a quantity. When that quantity is multiplied by the direct material cost of the item, the result is a cash value assigned to the target.

Note that only the direct material cost is considered. This is the best way to make a true working capital comparison without worrying about overhead and labor allocations. These types of allocations can be distortive. Direct material cost cannot be distorted. If those allocations are considered, the picture becomes even more skewed in favor of the decoupling points placed lower in the product structures; those parts don't get that "value added" that higher-level components get and consequently look "cheaper" to hold. However, there is no real difference from a true working capital perspective.

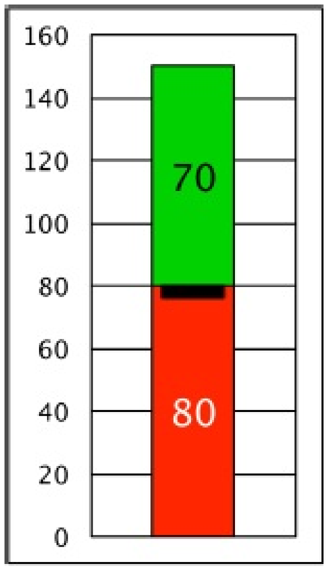

FIGURE 9-65 Net flow position for min-max example

Establishing the direct material cost of all items starts with knowing the direct material costs of all purchased items. Parent item direct material cost is the sum of the component direct material costs. Figure 9-66 shows the three product structures with the direct material costs for all purchased items in dollars below each item.

With the material costs of the purchased items identified, it is then possible to calculate the direct material cost for all parent items. Figure 9-67 shows a summary of the direct material costs for all items in all product structures for this example. Each parent's direct material cost is the sum of its immediate component direct material cost. For example, the direct material cost of 202 ($235) is the sum of the direct material cost of 306 ($125) and 305P ($110). These direct material costs are combined with the calculated average on-hand target equation in order to show the working capital implications for each decoupling iteration.

The Company ABC example began with all end items (FPA, FPB, and FPC) being buffered. Figure 9-68 shows the buffer values in the starting situation before the decoupling point selection example. Each end item's estimated average on-hand dollar value is calculated in the lower right-hand corner of the table. The direct material cost is the total value of the components in each item. The average on-hand target for each item is established by adding the total red zone value to half the green zone value (the average on-hand target equation). The average inventory dollars for each position is established by multiplying the average on-hand target by the direct material cost of the item. The total for all estimated parents at the beginning of the example is $3,260,723.

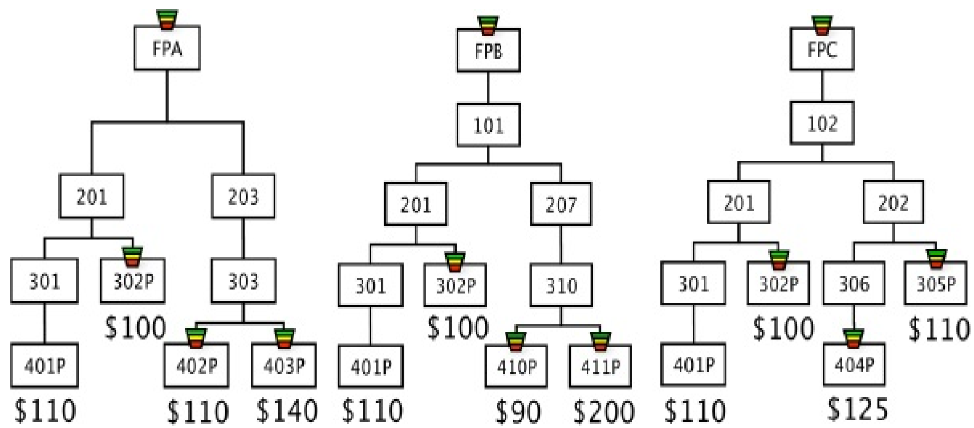

FIGURE 9-66 Product structures with purchased part material costs

Part	Direct Material Cost
FPA	$460
FPB	$500
FPC	$445
101	$500
102	$445
201	$210
203	$250
207	$290
202	$235
301	$110
303	$250
310	$290
306	$125
302P	$100
305P	$110
401P	$110
402P	$110
403P	$140
410P	$90
411P	$200
404P	$125

FIGURE 9-67 All direct material costs

The intermediate component 201 is then selected for decoupling. This compresses the end item buffer levels dramatically but does require an investment in a buffer for 201. Figures 9-69, 9-70, and 9-71 show the reduction in average on-hand target value and average inventory dollars for each end item position when 201 is buffered (shaded boxes). Figure 9-72 shows the average on-hand target and average dollar investment for the 201 buffer.

The investment required for the 201 buffer must be netted against the total reduction in parent item inventory. Figure 9-73 summarizes the inventory picture to this point. By buffering 201, a total of $281,163 average on-hand dollars is expected to be released from the system.

In the next positioning iteration, the component 203 was buffered, allowing for the FPA end item buffer to be completely eliminated. Figure 9-74 shows the required investment for the establishment of the 203 buffer.

FPA	
Green Zone	**1250**
	LT Factor: 1250 (5000 x .25)
	Minimum Order Quantity: 250
	Order Cycle: 750 (3(DOC) x 250(ADU))
Yellow Zone	**5000** (20(DLT) x 250(ADU))
Red Zone	**1875** (1250 + 625)
	Base: 1250 (5000 x .25)
	Safety: 625 (1250 x .5)

FPC	
Green Zone	**1725**
	LT Factor: 1725 (6900 x .25)
	Minimum Order Quantity: 250
	Order Cycle: 900 (3(DOC) x 300(ADU))
Yellow Zone	**6900** (23(DLT) x 300(ADU)
Red Zone	**2588** (1725 + 863)
	Base: 1725 (6900 x .25)
	Safety: 863 (1725 x .5)

FPB	
Green Zone	**575**
	LT Factor: 575 (2300 x .25)
	Minimum Order Quantity: 250
	Order Cycle: 300 (3(DOC) x 100(ADU))
Yellow Zone	**2300** (23(DLT) x 100(ADU))
Red Zone	**863** (575 + 287.5)
	Base: 575 (2300 x .25)
	Safety: 287.5 (575 x .5)

	FPA	FPB	FPC
Direct Material Cost	$460	$500	$445
Average On-Hand Target	2500	1151	3451
Average Inventory $	$1,150,000	$575,250	$1,535,473
Total End Item Inventory $	**$3,260,723**		

FIGURE 9-68 Starting average on-hand inventory position

Buffer Comparison (FPA)

Before 201 Decoupling			
Average Daily Usage	250	Green Zone	**1250**
Buffer Profile	M, L (.25), M (.5)		LT Factor: 1250 (5000 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle (DOC)	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
Decoupled Lead Time (DLT)	20 days	Yellow Zone	**5000** (20(DLT) x 250(ADU))
Direct Material Cost	$460	Red Zone	**1875** (1250 + 625)
Average On-Hand Target	2500		Base: 1250 (5000 x .25)
Average Inventory $	$1,150,000		Safety: 625 (1250 x .5)
After 201 Decoupling			
Average Daily Usage	250	Green Zone	**750**
Buffer Profile	M, M (.4), L (.25)		LT Factor: 700 (1750 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
DLT After Decoupling	7 days	Yellow Zone	**1750** (7(DLT) x 250(ADU)
Direct Material Cost	$460	Red Zone	**875** (700 + 175)
Average On-Hand Target	1250		Base: 700 (1750 x .4)
Average Inventory $	$575,000		Safety: 175 (700 x .25)

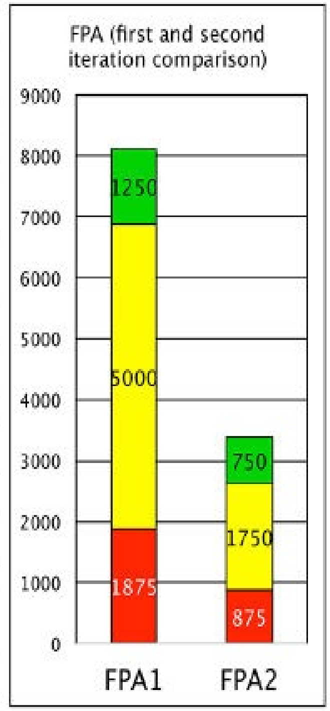

FIGURE 9-69 FPA on-hand and dollar compression

Buffer Comparison (FPB)

Before 201 Decoupling			
Average Daily Usage	100	Green Zone	**575**
Buffer Profile	M, L (.25), M (.5)		LT Factor: 575 (2300 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 300 (3(DOC) x 100(ADU))
Decoupled Lead Time (DLT)	23 days	Yellow Zone	**2300** (23(DLT) x 100(ADU))
Direct Material Cost	$500	Red Zone	**863** (575 + 287.5)
Average On-Hand Target	1151		Base: 575 (2300 x .25)
Average Inventory $	$575,250		Safety: 287.5 (575 x .5)
After 201 Decoupling			
Average Daily Usage	100	Green Zone	**360**
Buffer Profile	M, M (.4), M (.5)		LT Factor: 360 (900 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 300 (3(DOC) x 100(ADU))
DLT After Compression	9 days	Yellow Zone	**900** (9(DLT) x 100(ADU))
Direct Material Cost	$500	Red Zone	**540** (360 + 180)
Average On-Hand Target	720		Base: 360 (900 x .4)
Average Inventory $	$360,000		Safety: 180 (520 x .5)

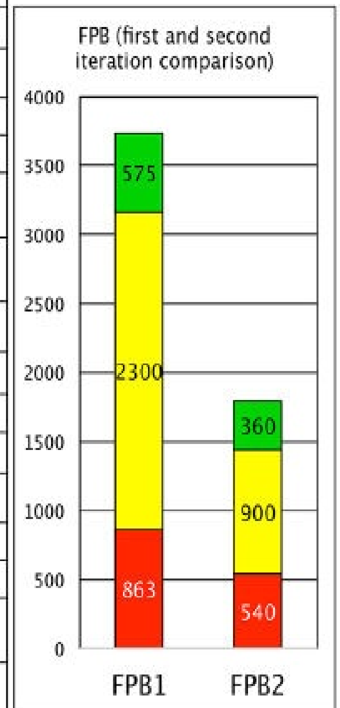

FIGURE 9-70 FPB on-hand and dollar compression

Buffer Comparison (FPC)

Before 201 Decoupling			
Average Daily Usage	300	Green Zone	**1725**
Buffer Profile	M, L (.25), M (.5)		LT Factor: 1725 (6900 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 900 (3(DOC) x 300(ADU))
Decoupled Lead Time (DLT)	23 days	Yellow Zone	**6900** (23(DLT) x 300(ADU)
Direct Material Cost	$445	Red Zone	**2588** (1725 + 863)
Average On-Hand Target	3451		Base: 1725 (6900 x .25)
Average Inventory $	$1,535,473		Safety: 863 (1725 x .5)
After 201 Decoupling			
Average Daily Usage	300	Green Zone	**960**
Buffer Profile	M, M (.4), L (.25)		LT Factor: 960 (2400 x .4)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
DLT After Compression	8 days	Yellow Zone	**2400** (8(DLT) x 300(ADU)
Direct Material Cost	$445	Red Zone	**1200** (960 + 240)
Average On-Hand Target	1680		Base: 960 (2400 x .4)
Average Inventory $	$747,600		Safety: 240 (960 x .25)

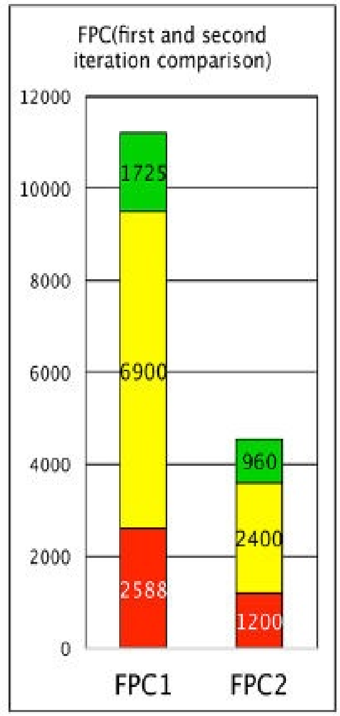

FIGURE 9-71 FPC on-hand and dollar compression

Buffer Worksheet (201)

Average Daily Usage	650	Green Zone	**3088**
Buffer Profile	I, L (.25), M (.5)		LT Factor: 3088 (12350 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	19 days	Yellow Zone	**12350** (19(DLT) x 650(ADU)
Direct Material Cost	$210	Red Zone	**4632** (3088 + 1544)
Average On-Hand Target	6176		Base: 3088 (12350 x .25)
Average Inventory $	$1,296,960		Safety: 1544 (3088 x .5)

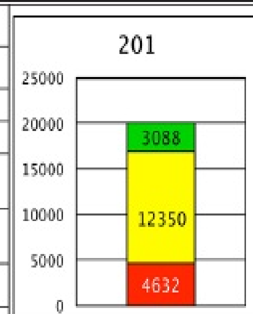

FIGURE 9-72 Component 201 average on-hand target and inventory dollars

	FPA	FPB	FPC	201	Total Inventory
Start	1,150,000	$575,250	$1,535,473	$0	$3,260,723
Iteration 1	$575,000	$360,000	$747,600	$1,296,960	$2,979,560
Reduction					$281,163

FIGURE 9-73 Inventory summary after the first decoupling iteration

Buffer Worksheet (203)			
Average Daily Usage	250	Green Zone	750
Buffer Profile	I, M (.5), L (.2)		LT Factor: 750 (1500 x .5)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 750 (3(DOC) x 250(ADU))
Decoupled Lead Time (DLT)	6 days	Yellow Zone	1500 (6(DLT) x 250(ADU)
Direct Material Cost	$250	Red Zone	900 (750 + 300)
Average On-Hand Target	1275		Base: 750 (1500 x .5)
Average Inventory $	$318,750		Safety: 150 (750 x .2)

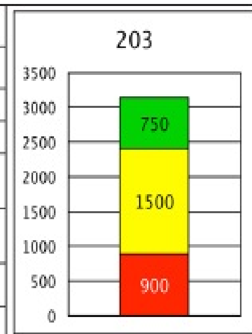

FIGURE 9-74 Component 203 average on-hand target and inventory dollars

	FPA	FPB	FPC	201	203	Total Inventory
Start	$1,150,000	$575,250	$1,535,473	$0	$0	$3,260,723
Iteration 1	$575,000	$360,000	$747,600	$1,296,960	$0	$2,979,560
Iteration 2	$0	$360,000	$747,600	$1,296,960	$318,750	$2,723,310
Total Reduction						$537,413

FIGURE 9-75 Inventory summary after the second decoupling iteration

After the second decoupling iteration, two component buffers (201 and 203) have been added, and one end item buffer (FPA) has been eliminated. Figure 9-75 is an inventory summary after the second iteration, showing a total estimated inventory reduction of $537,413.

The final decoupling iteration for Company ABC establishes a buffered position for the purchased component 401P. This allows for a dramatic compression of the decoupled lead time of 201. Figure 9-76 shows the on-hand target and inventory dollar compression for 201.

However, in order to accomplish the 201 compression, an investment in a buffer at 401P is required. Figure 9-77 shows the on-hand target and inventory dollar investment for 401P.

Finally, this example started in Chapter 7 is complete. Figure 9-78 summarizes the inventory dollar compression through all iterations, showing a total estimated reduction in on-hand inventory investment of $1,325,828.

Buffer Comparison (201)			
Before 401P Decoupling			
Average Daily Usage	650	Green Zone	**3088**
Buffer Profile	I, L (.25), M (.5)		LT Factor: 3088 (12350 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	19 days	Yellow Zone	**12350** (19(DLT) x 650(ADU)
Direct Material Cost	$210	Red Zone	**4632** (3088 + 1544)
Average On-Hand Target	6176		Base: 3088 (12350 x .25)
Average Inventory $	$1,296,960		Safety: 1544 (3088 x .5)
After 401P Decoupling			
Average Daily Usage	650	Green Zone	**1950**
Buffer Profile	I, L (.25), L (.2)		LT Factor: 1463 (5850 x .25)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
DLT After Compression	9 days	Yellow Zone	**5850** (9(DLT) x 650(ADU)
Direct Material Cost	$210	Red Zone	**1756** (1463 + 293)
Average On-Hand Target	2731		Base: 1463 (5850 x .25)
Average Inventory $	$573,510		Safety: 293 (1463 x .2)

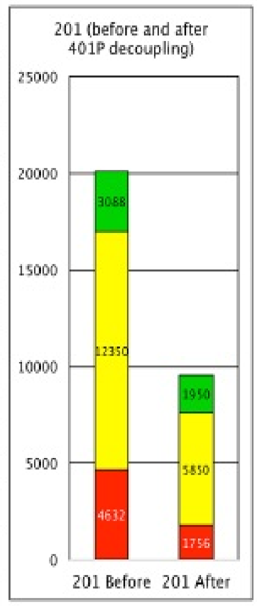

FIGURE 9-76 Component 201 on-hand and dollar compression

Buffer Worksheet (401P)			
Average Daily Usage	650	Green Zone	**1950**
Buffer Profile	P, L (.3), L (.25)		LT Factor: 1950 (6500 x .3)
MOQ	250		Minimum Order Quantity: 250
Desired Order Cycle	3 days		Order Cycle: 1950 (3(DOC) x 650(ADU))
Decoupled Lead Time (DLT)	10 days	Yellow Zone	**6500** (10(DLT) x 650(ADU))
Direct Material Cost	$110	Red Zone	**2438** (1950 + 488)
Average On-Hand Target	3413		Base: 1950 (6500 x .3)
Average Inventory $	$375,430		Safety: 488 (1950 x .25)

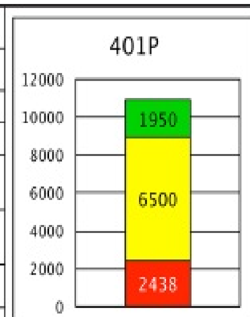

FIGURE 9-77 Component 401P average on-hand target and inventory dollars

	FPA	FPB	FPC	201	203	401P	Total Inventory
Start	$1,175,000	$575,250	$1,575,438	$0	$0		$3,325,688
Iteration 1	$575,000	$360,000	$747,600	$1,296,960	$0		$2,979,560
Iteration 2	$0	$360,000	$747,600	$1,296,960	$318,750		$2,723,310
Iteration 3	$0	$360,000	$747,600	$573,510	$318,750	$375,430	$1,999,860
Total Reduction							$1,325,828

FIGURE 9-78 Inventory summary after the third decoupling iteration

Summary

The planning component of Demand Driven Material Requirements Planning is a simple, intuitive, and highly visible way to generate supply orders. Its use of qualified sales orders means that all components of the supply generation equation (on-order, on-hand, and qualified sales orders) are known and contain relatively little variability. This combined with the decoupling point positioning and buffers means that nervousness, supply continuity variability, and the bullwhip effect are mitigated. The net flow equation and net flow position allow for quick, intuitive, and informative views across groups of items, giving a real sense of relative priority and how to best handle that relative priority.

CHAPTER 10

Demand Driven Execution

In DDMRP, a careful distinction is made between planning and execution. "Planning" is the process of generating supply order requirements using the net flow equation and the elements of decoupled explosion explained in the previous chapter. Planning ends once the recommendation has been approved and becomes an open supply order (purchase order, manufacturing order, or transfer order).

In DDMRP, "execution" is the management of open supply orders against relevant criteria. These criteria are defined in two basic categories necessary to protect and promote flow: buffer status and synchronization. Figure 10-1 depicts the four basic DDMRP execution alerts in these two categories.

Buffer status alerts are designed to show the current and projected status of the decoupling point positions (independent points) across the Demand Driven Operating Model. These alerts use the current and projected on-hand position rather than the net flow position. If there is no on-hand, then the decoupling point is not decoupled and will most likely pass on variability. This can occur despite a green net flow position. The net flow equation is to plan that position only. On-hand tells us if the position is capable of performing its purpose—to maintain decoupling. The two buffer status alerts will be discussed in depth later in this chapter.

Synchronization alerts are designed to highlight problems with regard to dependencies. Dependencies still exist in DDMRP (see the section on decoupled explosion in Chapter 9). These dependencies are about known demand requirements versus projected supply availability. While the buffers mitigate the transference of variability up and down the chain, synchronization is still important in DDMRP between decoupling points and particularly between a decoupling point and the customer. The better the visibility to synchronization problems, the less variability is transferred to and between buffers and to the customer. Two forms of synchronization alerts will be discussed in depth later in this chapter.

Buffer Status Alerts

In order to really understand how and why buffer status alerts work, two key perception changes must occur. The first change requires challenging how conventional priority management typically works. The second change is a shift in the buffer color display between planning (using the net flow equation) and execution (focusing on current and projected on-hand).

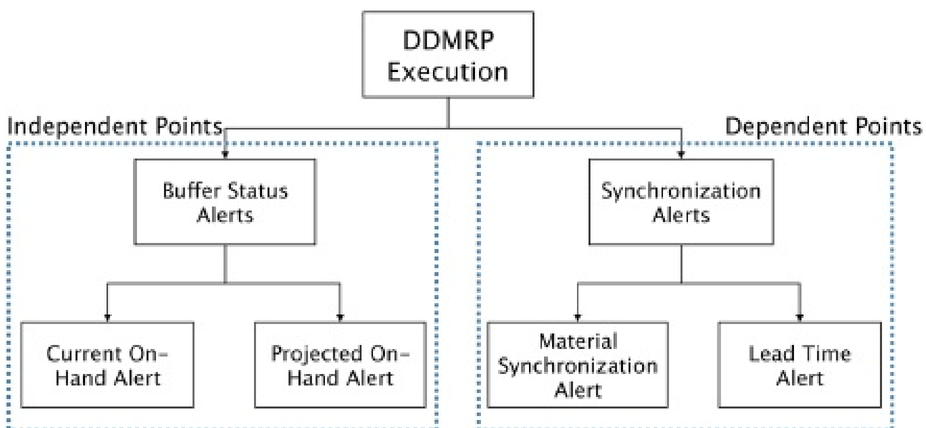

FIGURE 10-1 DDMRP basic execution alerts

Challenging Priority by Due Date

To understand the power behind DDMRP execution, first the problem with conventional planning systems when it comes to execution needs to be understood. The "P" in MRP stands for "planning." Material Requirements Planning inherently is a *planning* system and not an execution system. Conventional MRP systems lack real-time visibility to relative priorities associated with purchase orders (POs), transfer orders (TOs), and manufacturing orders (MOs) throughout the internal manufacturing operation and across the supply chain.

Without this effective priority approach, conventional tools force supply chains (i.e., suppliers, manufacturing, fulfillment, and customers) to employ a rudimentary and arbitrary priority system called priority by due date. Due dates drive everything when it comes to execution. Common practice is that if suppliers are late, it counts against them in their performance report. When it comes to MRP, all plans are developed assuming that dates are maintained. Everything must go according to plan, or the outcome (final due date) is in jeopardy because there is no slack in the planned schedule. If a manufacturer is consistently not able to meet customer due dates, then there are negative business implications. These include lost opportunities and increased expedite-related expenses. Companies are acutely aware of the importance of hitting due dates, especially in this current hypercompetitive market. This ripples throughout an organization, reinforcing the need to measure and act according to priority by due date.

Thus priority by due date becomes the default method to maintain expected customer service levels. When things are closer to being due, they become more important. If they are past due, they are very important. If they are past past due, then they are even more important. Teams of expediters are employed to determine how important things really are. Schedules are constantly changed (causing even more ripples), and overtime and expensive expedites are employed in an attempt to solve the problem. Due date is all that most supply chain personnel know, and due date is all that most supply chain personnel have at their disposal.

What if this mode of operation is more a part of the problem rather than a solution? This book has emphasized from the beginning the importance of protecting and promoting the flow of relevant information and materials. What if the prevailing priority by due date mode of operation is actually more distortive than relevant?

Ask any buyer a very simple question: “Would you rather have your suppliers be on time or never stock you out?” This question usually elicits a very interesting reaction. Most responders blurt out, “On time, of course!” but then pause with a quizzical look on their face as they think about it harder and consider other possibilities. Intuitively, they begin to consider that there may be a misalignment between what companies use as a metric and what effect that metric might have. Suppliers can be consistently 100 percent on time, and yet the company still has shortages. This is especially true when due dates generated out of conventional MRP systems are based purely on the assumptions and the plan at the time of their actual creation. Yet we know that those assumptions are extremely short-lived, as conventional MRP is highly subject to nervousness (demand signal distortion and change) and supply continuity variability (delay accumulation). The “real” requirements are changing constantly, but the due dates stay the same as they were on release. Even if a supplier hits the due dates, there is still a significant chance that there could be shortages or excess inventory based on the changes that have occurred between the time the order was created and the time in which it was received.

Thus priority by due date rarely conveys the real day-to-day inventory and materials priorities. Priorities are not static. They change as variability and volatility occur within the active life span of POs and MOs—the time from when they are opened until they are closed. This life span is called the "execution horizon." Customers change their orders. Quality challenges occur. There can be weather- or customs-related obstacles. Engineering changes happen. Suppliers' capacity and reliability can fluctuate temporarily. The longer the execution horizon is, the more volatile are priorities. This means that the company is more susceptible to adverse material synchronization issues and shortages. This increased variability and volatility guarantees that despite our best attempts at planning, reality will deviate from the plan. Conventional MRP requires that everything go according to plan for the due dates to be the relevant driving force behind priority. More sophisticated scheduling systems like advanced planning and schedule or optimization (APS or APO) only make this variability and volatility worse with its more frequent rescheduling.

Additionally, other challenges are associated with driving priority by due date. Request and promise dates change frequently due to nervousness and supply continuity variability. These changes often create confusion and disagreement between suppliers and customers about the "real" dates. Suppliers could view their on-time performance as high because they delivered to their promise date, whereas customers see it much differently from their view of the request date. It is often the case that MRP systems will even request things with due dates in the past!

Aligning a supplier's schedule with a customer's real priorities under conventional MRP approaches is a huge challenge. A manufacturing company can have several open POs to a supplier all with the same due date. Figure 10-2 is an example of such a case. Note that there are three orders all due on the same day (POs 821158, 831145, and 831162).

Order #	Due Date	Supplier
PO 821158	05/12	PNW Fabrication
PO 831145	05/12	PNW Fabrication
PO 831162	05/12	PNW Fabrication

FIGURE 10-2 Determining PO priority by due date

These orders could have nothing to do with each other from the customer's perspective. From a supplier's perspective, however, they might all converge at a bottleneck resource. If the supplier realizes that it simply cannot fulfill all these orders on the required due dates, how can it decide which is the most important? Options can include calling the buyer or choosing the order based on what the supplier perceives to be its best use of its time or perhaps by the lowest order number.

If the decision is to call the buyer, can the buyer quickly convey the correct priority? In most cases with conventional MRP tools, the answer is simply, "I'll have to get back to you." Determining the correct priority will require an additional amount of data analysis and potential communication with planning. Additionally, whatever answer is derived from this conversation will most likely change as time moves forward and the MRP deck of cards gets shuffled once again. This situation is exacerbated when an APO or APS system is in place.

If the supplier instead does what it perceives to be the best use of its time and capacity, the fact that it might pick the right priority for the customer would be completely coincidental. Picking priority based on order number seems arbitrary at best. Both these options risk supplier capacity being out of alignment with actual customer need. Is there a simple and intuitive way to fix this?

Figure 10-3 is the same graphic as Figure 10-2 but with one additional field added that makes immediate priority determination relatively easy. This view, if provided to the buyer, would allow for a definitive and quick response. Furthermore, if this view were provided to the supplier on a recurring basis, it would negate the need for the phone call in the first place. Additionally, it would reduce the risk that the supplier would make a decision counter to the interest of the customer or one of an arbitrary nature.

Order #	Due Date	Buffer Status	Supplier
PO 831145	05/12	RED - 12.3%	PNW Fabrication
PO 821158	05/12	YELLOW - 52.3%	PNW Fabrication
PO 831162	05/12	YELLOW - 56.1%	PNW Fabrication

FIGURE 10-3 Buffer status included

Just as in in planning, there is a general reference that is color-based with a deterministic percentage. Out of this comes two critical forms of relevant priority not available in conventional MRP. First is a sense of how a single part's status is relative to its own established buffer level necessary to maintain decoupling point integrity. Second, it allows a sense of how a part's status is relative to other parts' statuses. This is crucial for the example in Figure 10-3, where the supplier needs to quickly understand which order is the most important. Figure 10-3 clearly shows that despite order number sequence and due date, the highest priority order is PO 831145.

The priority by due date problem does not just affect the traditional customer-supplier relationship; it has significant implications for a manufacturer as well. The manufacturing floor and manufacturing order priority determination is used to further extend the priority by due date versus priority by buffer status comparison in this next example in Figure 10-4. These are make-to-stock (MTS) manufacturing orders with different due dates to choose from.

If today's date is 5/11, this manufacturing resource has to determine which is the most important to run next. In a typical environment there are three orders that would contend for the highest priority: MOs 831145, 821158, and 831162. The other MOs are typically disregarded. Should the manufacturer simply work the sequence given to it? That is the schedule; yet do we have enough relevant information to make a good decision to best promote and protect flow?

Figure 10-5 now gives the manufacturing floor a much clearer picture on priority. In this case operating according to priority by due date would lead to the wrong sequence from a flow perspective. The order that has the latest due date (MO 645181 with a date of 5/14) has the buffer that is in real jeopardy. Ultimately, priority by buffer status is about aligning efforts to best protect the DDMRP model. The DDMRP model was built under careful and strategic consideration (see Chapters 6, 12, and 13). Aligning efforts to the strategic operating model of a business in real time is the pinnacle of effective execution. In fact, if buffer statuses are displayed, that raises two questions: Why display due dates at all? What additional value do they bring versus the risk of confusion and distortion?

Seeing color and buffer percentages and relating that to priority is intuitive for most people; red is danger, yellow is caution, and green typically means OK. How these colors and percentages are determined are key to effective and collaborative execution.

Order #	Due Date	Order Type	Customer
MO 831145	05/12	MTS	Internal
MO 821158	05/12	MTS	Internal
MO 831162	05/12	MTS	Internal
MO 845172	05/13	MTS	Internal
MO 645181	05/14	MTS	Internal

FIGURE 10-4 Dispatch list for 5/11

Order #	Due Date	Buffer Status	Order Type	Customer
MO 645181	05/14	RED – 13.2%	MTS	Internal
MO 845172	05/13	RED – 36.2%	MTS	Internal
MO 831162	05/12	YELLOW – 62.1%	MTS	Internal
MO 831145	05/12	YELLOW – 72.1%	MTS	Internal
MO 821158	05/12	YELLOW – 87.2%	MTS	Internal

FIGURE 10-5 Manufacturing schedule with buffer status

Planning Versus Execution Display

The buffer status can provide a quick and intuitive way to align efforts to best protect the DDMRP model. Buffer status alerts do not use the net flow equation; they utilize on-hand values only. This separates the activities dedicated to supply order generation from the activities dedicated to open supply order management. Thus buffer status alerts represent a different perspective than DDMRP planning.

As discussed in Chapter 9, the on-hand target range is calculated as top of red (TOR) to TOR + the green zone value. The on-hand position should fluctuate between those values. If the on-hand is consistently above TOR, then we don't have any real concerns. The buffer position is functioning as planned; it has adequate on-hand to ensure the integrity of the decoupling point. It is OK from an execution perspective, and “OK” means green. Thus when considering the buffer status alerts for the management of open supply orders against buffered positions when on-hand is above TOR, the buffer status alert will display green.

Yellow and red are determined by the severity of the on-hand situation in relation to the total red zone value. This should make conceptual sense when you consider the nature of the red zone. The red zone is the embedded safety in the buffer—the cushion against variability. When on-hand dips into that safety zone, a cautionary flag should go up; it should be colored yellow. If on-hand continues to erode into the safety zone, that cautionary flag becomes more urgent; it should change from yellow to red.

This rationale gives what is needed to determine color and percentage. First, there is a point at which yellow should turn to red. This is called the on-hand alert level. The on-hand alert level is the point that escalates priority from yellow to red. The most common way to determine the on-hand alert level is to set it to 50 percent of the red zone. Second, the discrete percentage value is calculated based upon the on-hand level as a percentage of the embedded safety (the red zone). The lower the percentage, the less safety remains.

As an example, take a part with a TOR of 50 and on-hand alert level of 25 (50 percent of the red value). On day 1 an on-hand value of 72 would produce an execution buffer status displaying a green value of 144 percent. The color green is assigned due to the on-hand quantity being above the TOR value. On day 2 an on-hand alert value of 26 would produce an execution buffer status displaying a yellow value of 52 percent. Yellow is displayed because the on-hand value is below the TOR value of the buffer but above the on-hand alert value. On day 3 an on-hand alert value of 20 would produce an execution buffer status displaying a red value of 40 percent. Red is displayed because the on-hand quantity is below the on-hand alert value. Figure 10-6 depicts the buffer status alert for this example part on each successive day.

Part # Example				
Day	On-Hand	TOR	On-Hand Alert	Status
1	72	50	25	GREEN – 144.0%
2	26	50	25	YELLOW – 52.0%
3	20	50	25	RED – 40.0%

FIGURE 10-6 Buffer status alert example

This shift in color assignment can be difficult to immediately grasp because it means that an on-hand value that would represent a yellow designation in planning is actually green from an execution perspective. The key to comprehending this change is simply to understand that the equation and display for planning are separate and distinct from execution. Planning utilizes the net flow equation (of which on-hand is only one component) in order to generate orders. The execution view with buffer status alerts is designed to provide relevant information on what is critical with regard to managing orders that have already been created. Which do we need to expedite? Which can afford to be late?

Figure 10-7 represents the planning simulation from Chapter 9 but displayed from an execution view. This chart only shows the on-hand value against the buffer status alert zones over the course of the 21 days. This gives a clear picture about what a planner would be thinking with regard to on-hand status (as opposed to supply order generation) on any particular day over the 21 days.

The TOR value of this part was 52, and we are setting the on-hand alert level to 50 percent of the red, which is 26. Figure 10-8 represents a day-by-day buffer status for the simulated part. As indicated in both Figures 10-7 and 10-8, there is only one day in which the on-hand level drops below the on-hand alert level. During the simulation, that drop actually prompted the planner to expedite an existing order. This is the purpose of buffer status alerts—to quickly point out where safety has been significantly eroded.

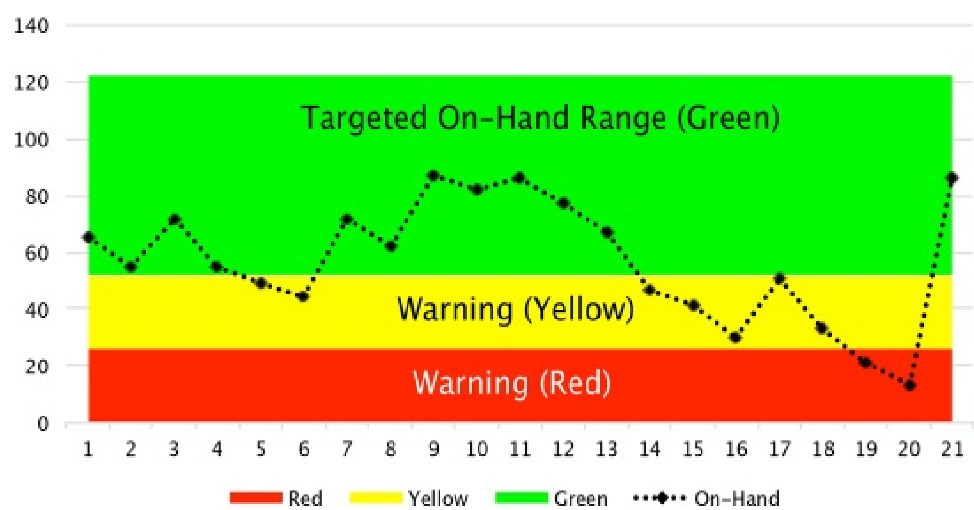

FIGURE 10-7 Planning simulation with execution color scheme

Part # Simulated Part				
Day	On-Hand	TOR	On-Hand Alert	Status
1	65	52	26	GREEN 125.0%
2	55	52	26	GREEN 105.8%
3	72	52	26	GREEN 138.5%
4	55	52	26	GREEN 105.8%
5	49	52	26	YELLOW 94.2%
6	44	52	26	YELLOW 84.6%
7	72	52	26	GREEN 138.5%
8	62	52	26	GREEN 119.2%
9	92	52	26	GREEN 176.9%
10	87	52	26	GREEN 167.3%
11	86	52	26	GREEN 165.4%
12	77	52	26	GREEN 148.1%
13	67	52	26	GREEN 128.8%
14	47	52	26	YELLOW 90.4%
15	41	52	26	YELLOW 78.8%
16	30	52	26	YELLOW 57.7%
17	51	52	26	YELLOW 98.1%
18	33	52	26	YELLOW 63.5%
19	21	52	26	RED 40.4%
20	52	52	26	YELLOW 100%
21	86	52	26	GREEN 165.4%

FIGURE 10-8 Day-to-day buffer status for a simulated part

Buffer status alerts focus on the on-hand position. The current on-hand alert focuses on today's on-hand buffer position against the defined execution buffer status definition. The projected on-hand alert focuses on a projected on-hand level in the near future typically up to one lead time in the future.

Current On-Hand Alert

The current on-hand alert is designed to show personnel what replenished positions are currently in trouble from an on-hand perspective only. For planning and purchasing personnel, these alerts are meant to identify parts where open supply may need to be immediately expedited. For manufacturing personnel, the current on-hand alert provides relevant information about which manufacturing orders should take precedence.

On-Hand Alert Manufactured Parts	
Part #	Status
FPB	RED - 20.7%
IC203	YELLOW - 53.5%
FPC	YELLOW - 99.4%

FIGURE 10-9 Sample manufactured parts' current on-hand alert screen for Company ABC

Figure 10-9 is what a current on-hand alert screen might look like for our previous example for Company ABC. All on-hand alert levels are set to 50 percent of the red zone value. In its simplest form, the alert screen only needs to provide a list of parts with priority sorted from highest to lowest. The on-hand alert screen is displaying for manufactured items only. This screen would be utilized by a planner. A separate on-hand alert screen would most likely be utilized by purchasing personnel for purchased items. Figure 10-10 is the purchasing on-hand alert screen for Company ABC.

Parts in green can be excluded from the list, as they are not encroaching into the safety zone. Of course, after being alerted to a part in trouble, most personnel would need quick access to additional information that may be involved in actions related to attempting to protect the buffer position—most notably open supply orders against the buffered position that are candidates for potential expedite.

Figure 10-11 depicts a drill-down view from the manufactured parts current on-hand alert screen. It displays open supply orders against the FPB position. In each case a request date (the due date generated using decoupled lead time) and promise date (the due date that the scheduler has confirmed) are displayed. Additionally we see that today's date is May 8 and the top of red and on-hand values are displayed that generate an on-hand alert status of 20.7 percent. There are three open manufacturing orders supplying FPB (MOs 15781, 15852, and 15999). Two of these orders are past due. One is severely past due relative to FPB's ADU (100). That would most likely explain the heavily eroded on-hand position.

On-Hand Alert Purchased Parts	
Part #	Status
401P	RED – 26.7%
402P	YELLOW – 58.3%
305P	YELLOW – 59.7%

FIGURE 10-10 Sample current on-hand alert screen for purchased items at Company ABC

Part # FPB On-Hand Buffer Status: 20.7%				Today's Date: 8-May
TOR: 540	Current On-Hand: 112			
Order #	Quantity	Request	Promise	Status
MO15781	363	1-May	2-May	Final QC
MO15852	390	5-May	7-May	In Progress – Final Config
MO15999	382	10-May	11-May	In Progress – Assembly

FIGURE 10-11 Order activity for FPB

The planner's immediate concern must be for MO 15781. A call is made to the final quality control station for information on the order. The planner is told that the order has passed inspection and the entire quantity will be moved to on-hand within the hour. With an additional 363 units, the on-hand level will move to 475 (112 + 363). An on-hand value of 475 will move the on-hand alert from red at 20.7 percent to yellow at 88 percent. Figure 10-12 is the current on-hand alert display after MO 15781 has been completed and the FPB on-hand position has been adjusted. At this point the planner feels comfortable with the position and takes no further action.

Figure 10-13 depicts a drill-down view from the purchased parts current on-hand alert screen. It displays open supply orders against the 401P position. Today's date is May 8. There are three open orders against the 401P position (POs 24366, 24413, and 24587). The buyer's most immediate concern should be for PO 24366, which shows that the order is in transit and the supplier has given a tracking number for it. When checking with the logistics company, the buyer is told that the order will arrive late in the day, possibly after hours, due to inclement weather. The buyer immediately makes arrangements for receiving personnel to stay late to meet the truck.

For distribution networks, the current on-hand alert works the same way as described in the previous examples with the exception that the on-hand status for each part number or SKU must be location specific as well. Figure 10-14 displays the current on-hand alert for a distribution network. There are two locations (Riverside and Dallas) that have severely eroded on-hand positions for product FPT. Additionally, the Chicago location for FPT is near the on-hand alert threshold.

On-Hand Alert Manufactured Parts	
Part #	Status
IC203	YELLOW – 53.5%
FPB	YELLOW – 88.0%
FPC	YELLOW – 99.4%

FIGURE 10-12 The updated current on-hand alert screen

Part # 401P On-Hand Buffer Status: 26.7%			Today's Date: 8-May
TOR: 2438	Current On-Hand: 651		Supplier: Chelan Fabrication
Order #	Quantity	Due Date	Status
PO24366	712	3-May	In Transit – Tracking #127823
PO24413	687	8-May	In Progress at Supplier
PO24587	720	12-May	In Progress at Supplier

FIGURE 10-13 Order activity for 401P

An expanded version of the current on-hand alert display would give a clear picture of buffer integrity across the network for all products. Figure 10-15 is an example of such a view. There are four locations: Riverside, Dallas, Atlanta, and Chicago. The network distributes five products (FPD, FPE, FPG, FPK, and FPT) from these locations. In this view, an additional column can be added that summarizes the overall network health for each product. This is represented by the "Network" column. How this column is populated can be configured based on the individual network (all locations may not have equal impact or strategic importance). In this case a very simple schema creates a color coding for network health; one red location equals a yellow designation; two or more red locations equal a red designation.

This network-wide current on-hand display in combination with prioritized share logic can help personnel decide which locations are in the most trouble when there are limited amounts of supply available. This can also point to better options to help positions in trouble. For example, for product FPT there is limited stock at the hub. The hub must service its own customers as well as the spoke locations. One option to help the Dallas location would be to cross-ship inventory from the Atlanta spoke. It is a less than ideal solution, but it may provide the most effective and expedient way to protect sales in all markets.

On-Hand Alert – Distributed Parts		
Part #	Status	Location
FPT	RED – 21.7%	Riverside
FPT	RED – 29.9%	Dallas
FPD	RED – 30.1%	Riverside
FPT	YELLOW – 51.6%	Chicago
FPE	YELLOW – 66.8%	Chicago
FPD	YELLOW – 64.2%	Atlanta
FPG	YELLOW – 75.1%	Dallas

FIGURE 10-14 Current on-hand alert display for a distribution network

SKU	Chicago (HUB)	Riverside	Dallas	Atlanta	Network
FPD	YELLOW – 97.0%	RED – 30.1%	YELLOW – 81.0%	YELLOW – 64.2%	YELLOW
FPE	YELLOW – 66.8%	GREEN – 112.0%	GREEN – 131.0%	GREEN – 101.2%	GREEN
FPG	GREEN – 115.0%	GREEN – 103.1%	YELLOW – 75.1%	GREEN – 113.2%	GREEN
FPK	GREEN – 140.0%	GREEN – 102.0%	GREEN – 115.0%	GREEN – 121.3%	GREEN
FPT	YELLOW – 51.6%	RED – 21.7%	RED – 29.9%	GREEN – 113.4%	RED

FIGURE 10-15 Network on-hand status

Projected On-Hand Alert

The other on-hand focused buffer status alert is the projected on-hand alert. The projected on-hand alert calculates red zone penetration in the near future in order to warn supply chain personnel about impending buffer integrity problems. In order to do this, the alert requires timing and quantity information from both a demand and supply perspective. Like the net flow equation, it utilizes elements of demand, supply, and on-hand. Unlike the net flow equation, it is not designed to recommend new supply orders but to point out positions that may need existing supply orders expedited in order to maintain the buffer's purpose.

The projected on-hand alert takes today's on-hand inventory and projects on-hand status for each future day based on the average daily usage or the quantity and timing of known demand allocations, depending on which is larger, and the quantity and timing of expected supply order receipts. In this case, due dates are relevant for both demand allocations and supply order receipts but only to make a projection of what buffer status may look like. It is worth mentioning that a projection is exactly that—a projection. It is the best guess about the nearterm future using the best information available: known demand allocations, calculated rate of demand, and known open supply and current on-hand levels.

To demonstrate how the projected on-hand alert works, the 21-day simulation of DDMRP supply order generation is revisited from the previous chapter. The format of the simulation provides an excellent way to articulate how this particular alert works. Figure 10-16 is the simulated environment on day 1 of the simulation.

On day 1 the on-hand inventory is 65 units. The projected on-hand alert will start with that value and use known demand allocations and anticipated supply order receipts to project on-hand levels for each day through one lead time (seven days) of the part beginning with day 2; this also corresponds to the order spike horizon used in the simulation. Figure 10-17 is a projection of on-hand levels against the buffer status alert color zones. The on-hand alert level is 50 percent of the red zone.

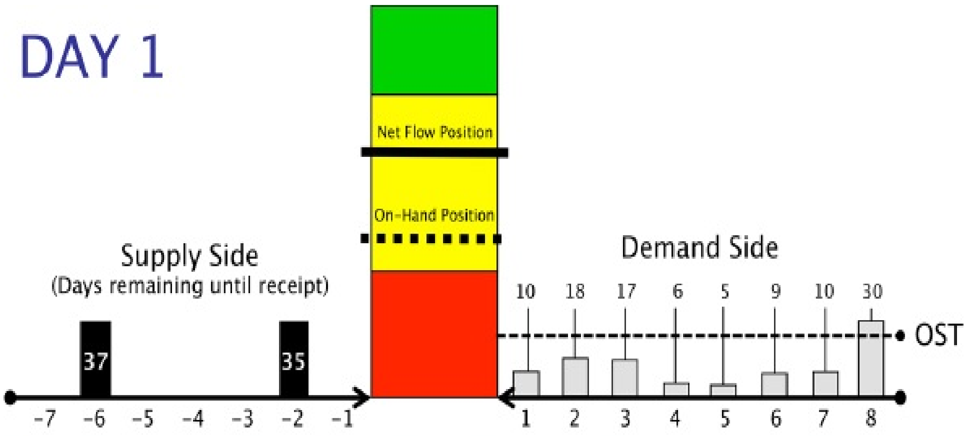

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Amount	Request Date	Top RED	Top YELLOW	Top GREEN	Lead Time
EXAMPLE	YELLOW 61.8%	65	72	40	97	60	Day 8	52	122	157	7

FIGURE 10-16 Day 1

	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Sales Order Demand	18	17	6	5	9	10	30
Expected Supply	35				37		
Projected Starting On-Hand	55	72	55	49	44	72	62
Projected Status	GREEN 105.8%	GREEN 138.5%	GREEN 105.8%	YELLOW 94.2%	YELLOW 84.6%	GREEN 138.5%	GREEN %

FIGURE 10-17 Projected on-hand alert status by day

The projected buffer status is generated by the projected starting on-hand position against a top of red, which is 52. In this case there appears to be little cause for concern if things go even remotely close to plan.

What if the known demand allocations are not available or are only visible in the very nearterm future? In these cases, the projected on-hand alert can project on-hand levels using ADU. Figure 10-18 depicts this simulated part using only ADU as the demand value for each future day. Without accounting for known demand allocations, especially qualified spikes, this form of the alert may overestimate future on-hand positions.

A more conservative approach would be to consider both ADU and known demand allocations and take the greater of the two on any particular day. Figure 10-19 depicts projected on-hand alert status for the simulated part using the highest value each day, either known demand allocations or the average daily usage value. The shaded boxes in the “Average Daily Usage” and “Sales Orders” rows represent the qualifying highest demand input for that day.

	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Average Daily Usage	10	10	10	10	10	10	10
Expected Supply	35				37		
Projected Starting On-Hand	55	80	70	60	50	77	67
Projected Status	GREEN 105.8%	GREEN 153.8%	GREEN 134.6%	GREEN 115.4%	YELLOW 96.2%	GREEN 148.1%	GREEN 128.8%

FIGURE 10-18 Projected buffer status alert with ADU

	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Average Daily Usage	10	10	10	10	10	10	10
Sales Orders	18	17	6	5	9	10	30
Expected Supply	35				37		
Projected Starting On-Hand	55	72	55	45	35	62	52
Projected Status	GREEN 105.8%	GREEN 138.5%	GREEN 105.8%	YELLOW 86.5%	YELLOW 67.3%	GREEN 119.2%	GREEN 100.0%

FIGURE 10-19 Projected on-hand alert using highest daily demand input

When there is less change in an environment, the projected on-hand alert is less impactful. In more stable environments there is less need for advance warning. But what about environments where dates are constantly moving? Suppliers push back delivery dates. Manufacturing is frequently shuffling the schedule, moving jobs forward and backward. In these cases the projected buffer status alert can be extremely valuable.

Figure 10-20 is a sample projected on-hand alert screen for Company ABC purchased parts. The projected on-hand alert should alert to the first on-hand alert-level penetration within one lead time into the future. If a part is projected to be stocked out, a dark red color will be displayed. In DDMRP the dark red color is more severe than red. A negative percentage means that the part is projected to be stocked out with demand. When a part is stocked out with demand, flow is definitely impeded. These potential situations are crucial for visibility.

Projected On-Hand Alert – Purchased Parts	
Today's Date: 3 – May	
Part #	Status
401P	DARK RED −4.3% (5 – MAY)
402P	RED 35.7% (6 – MAY)
404P	RED 49.1% (6 – MAY)

FIGURE 10-20 Sample projected on-hand alert screen for manufactured items at Company ABC

There are three parts on the alert screen, and one requires immediate attention. Part 401P is projecting to be stocked out with demand two days from now on May 5. In order to understand why it is projected to be stocked out will require a drill-down on the 401P order activity. A time-sequenced summary of planned open supply orders is needed with known demand allocations in order to determine potential corrective actions.

Figure 10-21 is an order activity summary for 401P. Demand allocations and open supply orders are listed in the column "Order #." Since 401P is a purchased item, any manufacturing order represents a demand allocation, while purchase orders represent the open supply orders. Each order has a quantity assigned to it. The orders are sequenced based on the "Date" column. For manufacturing orders, the date represents the parent order release date. This is when the order will require the stated quantity of 401P. That quantity is subtracted from the on-hand position on that future date. For purchase orders, the date represents the anticipated delivery of the order based on the best information available. The quantity of those orders is added to the projected on-hand position at those dates. The "Projected On-Hand" column displays the projected on-hand balance after the order quantities have been subtracted or added.

The current on-hand quantity is 2,108. There is one manufacturing order (MO 2574) set to release on May 5 for a quantity of 2,213. This demand allocation against the 401P position exceeds the projected on-hand balance on

May 5 by 105 units. The next supply order is not expected until the day after the manufacturing order's scheduled release. Additionally, another open supply is due the very next day. Apparently this issue was caused by a machine breakdown at the supplier that delayed PO 61325 by three days. An updated supplier promise date of May 6 has immediately generated a projected on-hand alert.

401P (DARK RED -4.3% (5-MAY))					Today's Date: 3-May
Current On-Hand: 2108					
Order #	Quantity	Date	Projected OH	Projected Status	Status
MO2574	2213	5-May	-105	DARK RED -4.3%	Unreleased
PO61325	2112	6-May	2007	YELLOW 82.3%	In Process at Supplier
PO61343	1987	7-May	3994	GREEN 163.8%	In Process at Supplier

FIGURE 10-21 Order activity summary screen for 401P

What can the buyer do? The supplier cannot get the order in sooner than May 6, and the manufacturing orders are not under the buyer's control. The buyer should consult with the manufacturing planner about the best option for the company. Planning will also see this problem but from the perspective of a different type of execution alert—a material synchronization alert. This example will be continued later in the "Material Synchronization Alerts" section.

In most environments, paying attention to projected on-hand alerts and acting appropriately will reduce the number of current on-hand alerts. It will never eliminate the on-hand alerts, as variability in the short term will still play a role with on-hand levels (quality holds, breakdowns, spoilage, etc.).

A mixed-mode operation has both make-to-stock (replenished) orders and make-to-order (nonbuffered and direct to customer) orders moving through the same resource base. In conventional operations most operations defer to the "real" customer order: the make-to-order item and its due date. Yet in DDMRP, this default mode of operation requires reexamination.

It is crucial to protect customer commitments, but it is worth a reminder that a Demand Driven Operating Model is configured to best protect total flow. A make-to-order order represents one customer, while a strategic stock buffer can represent hundreds of customers. Thus the actual priority consideration between stock buffers in trouble and a make-to-order order with a firm customer due date can really only be defined at the time when the conflict occurs. Choices will have to be made, but those choices should be impacted by the relevant information contained in the buffer status alerts.

Synchronization Alerts

As discussed in the previous chapter there will still be many dependencies and synchronization points that are not decoupled in DDMRP. When visibility can be promoted to keep these points from experiencing as little disruption as possible, then there will be less supply continuity variability passed to the decoupling points and customers. When less variability is passed to decoupling points, there is less working capital required to buffer those positions. When there is less variability passed to the customers, sales are typically protected and promoted.

DDMRP uses two types of execution alerts to promote visibility and management to synchronization points: material synchronization alerts and lead time alerts.

Material Synchronization Alerts

Material synchronization alerts display supply shortfalls against known demand allocations. In order to demonstrate a material synchronization alert, the previous projected on-hand alert example (401P) is continued but from a different perspective. That projection has created a potential synchronization problem for an order requiring 401P (MO 2574). The order calls for a quantity of 2,213, and the projected on-hand balance is expected to be short by 105 pieces. Manufacturing orders are under the control of planning, not purchasing. Thus planning personnel need to see when the orders and buffer positions under their control are in potential trouble. Remember, more visibility to relevant information directly impacts the amount of variability experienced, transferred, and amplified in an environment.

Figure 10-22 provides an additional reference for this example. It displays the product structure of FPA. The boxed-in area is the relevant part of the product for this example. The intermediate component 201 is the next decoupling point being fed by 401P and 301. Part 201 is buffered.

Part 301 is a nonbuffered item. Thus a projected stockout at 401P would create a synchronization problem for the demand allocations driven through the 301 position from supply orders generated at 201 (refer to the decoupled explosion section in Chapter 9 for additional information about decoupled explosion behavior). Supply is insufficient to cover the entire demand allocation. It is projected to be short by 105 pieces. When this occurs, a material synchronization alert will be issued for the manufacturing order(s) affected. Figure 10-23 is an example of a material synchronization alert for MO 2574.

The material synchronization alert needs to show planners the relevant factors that trigger that alert as well as the factors involved in a potential reconciliation of the issue. The alert displays the order number (MO 2574) that is impacted, the part number that the order is against (301), the release date of the order (May 5), and the quantity of the order (2,213). The “Shortage” column is displaying the shortage to MO 2574 (–105) and the part number creating the shortage (401P). The “Parent Buffer Status” column is providing a reference to the state of the next-higher buffered position in the product structure (201) that MO 2574 is supplying. In this case it is displaying the current buffer status alert of 201. The buffer appears to be in relatively good shape.

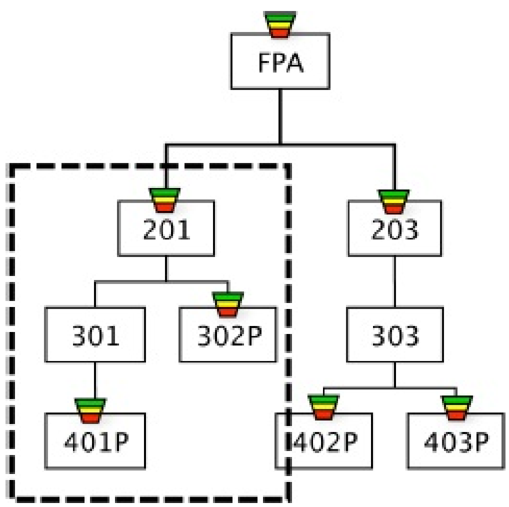

FIGURE 10-22 Product structure for example

Order #	Part #	Release Date	Quantity	Shortage	Parent Buffer Status
MO2574	301	5-May	2213	DARK RED 401P -105	201 - YELLOW 84%

FIGURE 10-23 Material synchronization alert for MO 2574

It should be noted that buffer status alerts are available to the planner for all make item buffers including 201. This material synchronization alert should impact the projected on-hand of 201 (less 105 than planned) on the date that MO 2574 is expected to be received at that position. This amount is insufficient, however, to generate a projected on-hand alert against the 201 position. Figure 10-24 shows the relationships between the execution alerts in this example and who is alerted by them. The buyer sees the current and projected on-hand alerts for 401P, while the planner sees alerts for manufactured items—the material synchronization alerts for 301 and current and projected on-hand alerts for 201.

In the projected on-hand alert example, the buyer's hands are essentially tied with regard to 401P; the supplier simply cannot get the 410P supply orders in sooner. Yet there are options open to the planner to deal with the material synchronization alert for MO 2574.

- **Option 1.** Move the order release date back to coincide with the next planned 401P supply. Since the parent buffer status is in relatively good shape, a short delay of the supply order should have minimal impact on the buffer; there is plenty of safety remaining. This option allows the planner to protect the scarcity in the 401P position if 401P might have additional needs defined in the short term.
- **Option 2.** Reduce the order by 105 pieces. This will get supply to the 201 position as planned from a timing perspective. This option might be selected if the 201 position is currently or projected to be in trouble in the near term.

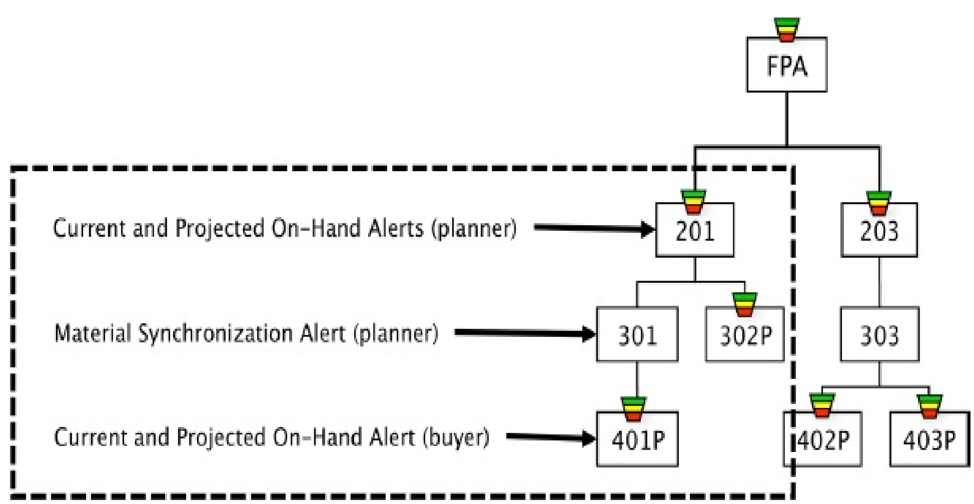

FIGURE 10-24 Execution alert relationships and their personnel assignments

- **Option 3.** Move the order release date back to coincide with the next planned 401P supply but then expedite order MO 2574. This would serve to protect both buffered positions (401P and 201) if necessary but would require additional manufacturing efforts or potential disruption from the expedite.

All these options would remove the material synchronization alert and the projected stockout alert (a projected on-hand alert would still be present but just not as severe).

In DDMRP a material synchronization alert will occur for three primary reasons:

- **Insufficient supply.** Material synchronization alerts are triggered when levels of supply are less than required. This is particularly true when a variance occurs on nonbuffered items. This can happen due to quality or significant yield loss issues. Insufficient supply can also be triggered by periods of significantly heavy demand.
- **Late supply.** When a promise date on a supply order is pushed later in time, it may trigger material synchronization alerts. This can also create a current or projected negative on-hand position in the case of a buffered part. This would drive current or on-hand alerts for the buffered item and material synchronization alerts for demand allocations that are shorted.

  Figure 10-25 is a planned relationship between a manufacturing order (demand allocation) and a supply order. The supply order is required to have full allocation for the manufacturing order on May 5.

  Figure 10-26 shows a supply order that has moved later in time. It is now expected to be received on May 7. This creates a synchronization issue for the manufacturing order on May 5.

  If the promise dates of nonbuffered parts are pushed later in time, then material synchronization alerts typically follow any demand allocations against those nonbuffered positions. The part is nonbuffered and is planned to net to zero. Thus any delay will create a synchronization problem because there is no buffer.

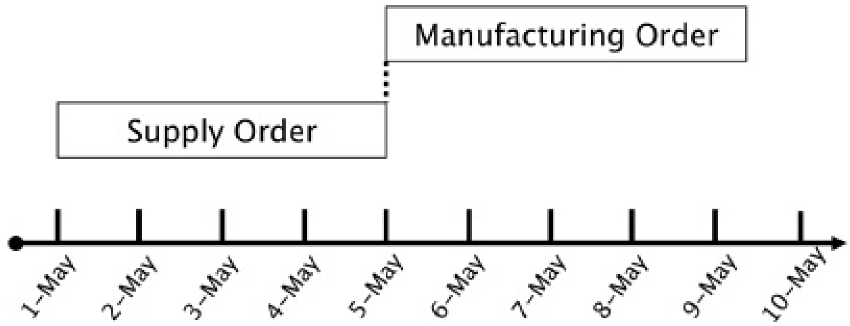

FIGURE 10-25 A planned relationship between a supply order and manufacturing order

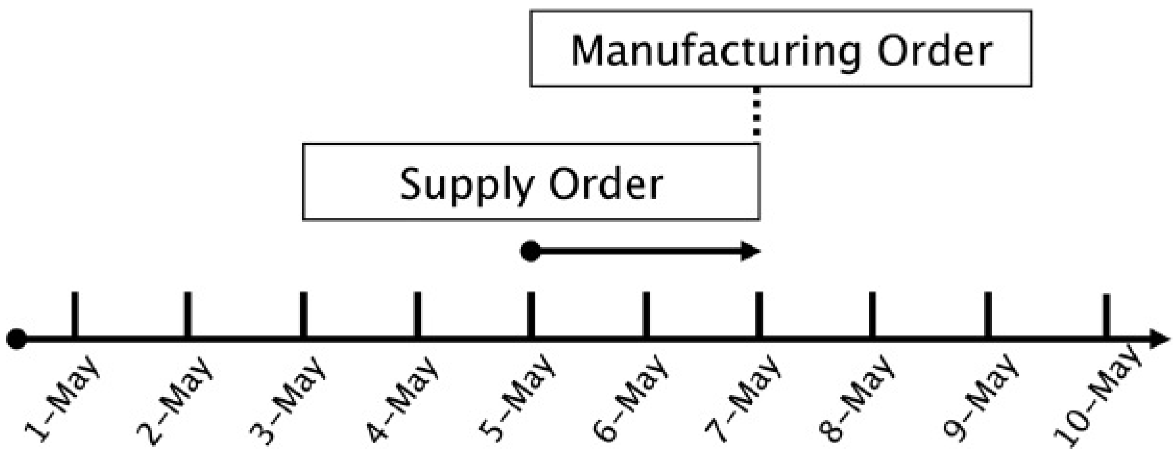

FIGURE 10-26 A supply order expected later than planned

- **Earlier Start Commitment.** If the start date of a demand requirement is changed to an earlier date, material synchronization alerts may follow. This is particularly common when the demand requirement is dependent on nonbuffered components. Since the component is nonbuffered, there is a dependent relationship based on when the demand was originally required and the timing created for the supply order of the component.

  Figure 10-27 illustrates the manufacturing order start date being moved earlier to May 3. This causes a synchronization problem with the supply order that is set to arrive on May 5.

Lead Time Alert

A lead time alert is an execution alert for strategic nonbuffered items. These strategic nonbuffered items may not come in sufficient volume to justify stocking but typically create major synchronization issues when they are required. They could come from a problematic supplier or from a problematic geographic region or be subject to handling or transportation difficulties. These items are planned in the same manner as any nonbuffered item; yet they are paid special attention in the execution horizon.

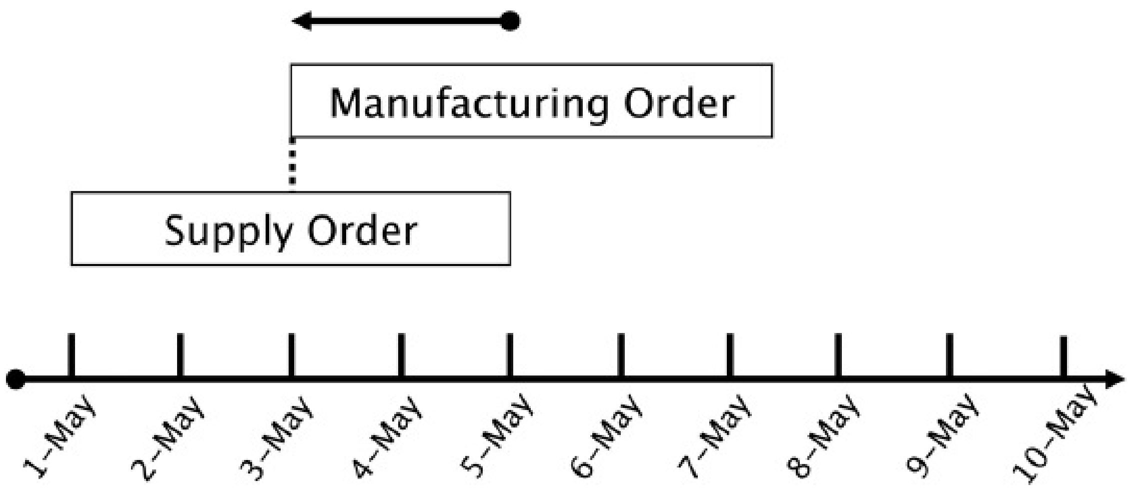

FIGURE 10-27 A manufacturing order moved earlier

The basic ideas behind the lead time alert are as follows:

1. It is beneficial to know about synchronization issues before a synchronization issue occurs. This allows adjustments and contingencies to be planned and executed in order to have better flow performance.
2. Supplier performance can be influenced by how organized a particular customer is with regard to the customer’s orders. If a customer can communicate in a way that helps prioritize a supplier’s actions, there is typically an appreciation and attention paid to that customer’s orders.
3. It can be beneficial to establish a clear “paper trail” showing what led to any synchronization issues or costs to avoid them. This is of particular importance in industries that use financial penalties for late deliveries. Establishing a clear chain of events that led to synchronization issues may provide recourse for cost recovery with problematic suppliers.

Lead time alerts are used to prompt personnel to check up on the status of critical non-stocked parts before those parts become a synchronization issue. Lead time managed (LTM) supply orders are tracked, and at a defined point in the part's lead time, personnel are prompted for follow-up. Typically, this point is two-thirds of the way through the lead time of the part, as measured from the promise date. The final third, called the lead time alert horizon, is then divided into three zones of typically equal proportions (green, yellow, and red). Figure 10-28 illustrates the lead time alert concept.

The lead time managed part has a lead time of nine days. Typically, LTM parts have much longer lead time items (30+ days). A part with a lead time of nine days was chosen as an example in order to make the dates easier to see graphically. A supply order for that part is launched on May 1 and is due nine days later on May 10. The lead time alert horizon is the last three days of the order's lead time. When the order is three days away from being due (on day 7), the order will enter the green zone of the lead time alert horizon. This is the first lead time alert for this part. On day 8 it enters the yellow zone of the lead time alert horizon. On day 9 it is in the red zone of the alert horizon.

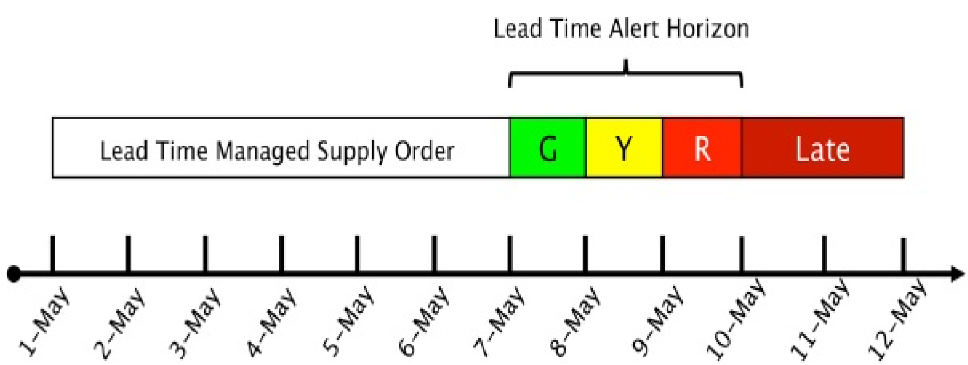

FIGURE 10-28 A lead time managed part and lead time alert zones

The zone colors of the lead time alert horizons are simply a countdown to the expected arrival date of the order. The progression from green to yellow to red is only meant to connote how close the order is to that date. Red simply means that it is due in the near term. If the order does not arrive on the expected date, then it progresses to a dark red color, meaning it is late. If the order goes late, it will most likely create a material synchronization alert for the demand allocation that caused it to be ordered in the first place.

The entry of an order into the lead time alert horizon should prompt a buyer or planner to follow up and document the order status. Any information derived from that follow-up could be valuable in predicting and mitigating the impacts of synchronization issues and, as well, could be valuable in documenting a clear picture about why any synchronization issues would occur. Follow-up and documentation should occur as the order progresses into each zone of the lead time alert horizon.

Note that the lead time alert horizon is not an insertion of additional time into the part's lead time—it overlaps the last portion of the lead time (typically the last third).

Figure 10-29 depicts a simple version of a lead time alert screen. The "Lead Time" column is present in order to understand the size of the lead time alert horizon and each zone. The lead time alert horizon has been set to one-third of the part's lead time. Part PPL has a 90-day lead time, giving it a lead time alert zone of 30 days. Each zone is 10 days in size. The lead time alert horizon for PPF is 20 days. Its green zone value is set to 6, yellow to 7, and red to 7. The lead time alert horizon for FPZ is 15 days. Green, yellow, and red are all set to 5 days.

In the lead time alert screen in Figure 10-29, the date is May 5. There are three orders with lead time alerts. Each order is in a different zone within its own respective lead time alert horizon. PO 112032 is in the red zone of its lead time alert horizon with 7 days remaining until expected receipt. Its total red zone is 10 days. It has been in the red zone for the 3 preceding days.

The "Current?" column in Figure 10-29 tells the planner or buyer whether there has been an updated note on the order for the zone that the order is in. The color green and "YES" indicate that a note has been entered against this zone. Thus there should be a minimum of three notes in each order record as the order progresses through its lead time alert horizon. Both PO 113562 and MO 5741 have yet to have their zone status update recorded. They are colored red and designated with a "NO."

Lead Time Alerts				Today's Date: 5-May	
Order #	Part #	Due Date	Lead Time Alert Status	Current?	Lead Time
PO112032	PPL	12-May	RED - 7 days left	GREEN - YES	90
PO113562	PPF	17-May	YELLOW - 12 days left	RED - NO	60
MO5741	FPZ	19-May	GREEN - 14 days left	RED - NO	45

FIGURE 10-29 A simple version of a lead time alert display for buyers and planners

Lead Time Alerts				Today's Date: 5-May	
Order #	Part #	Due Date	Lead Time Alert Status	Current?	Lead Time
PO112032	PPL	12-May	RED - 7 days left	GREEN - YES	90
PO113562	PPF	22-May	GREEN - 17 days left	GREEN - YES	60
MO5741	FPZ	19-May	GREEN - 14 days left	RED - NO	45

FIGURE 10-30 Updated lead time alert screen

If date revisions occur, this can adjust the lead time alert status. For example, if the zone status follow-up for PO 113562 resulted in the buyer finding out about a significant delay, then the new promise date should be immediately entered into the order record. If the supplier is now saying that due to equipment failure the order will be delayed by five days, the new due date must be adjusted to May 22. Figure 10-30 shows the updated lead time alert screen.

PO 113562 now has a due date of May 22, placing it in the green zone of its lead time alert horizon. Its zone status is current, and it will go through at least two more zone status updates (yellow and red). Furthermore, it has been documented why the date has changed. This date revision has most likely created a material synchronization alert(s) and/or projected on-hand alerts.

The lead time alert horizon does not have to be set as a percentage of a part's individual lead time. A global lead time alert horizon could be applied to all lead time managed parts or groups of parts with similar lead time properties.

Synchronization alerts by themselves can provide excellent visibility for environments in which there is limited opportunity to employ decoupling point buffers such as engineer-to-order and extensive make-to-order.

Summary

This chapter has explained four basic concepts and alerts for supply order management in DDMRP environments. These alerts and the concepts behind them are designed to create highly visible and collaborative execution across the Demand Driven Operating Model. There are nearly limitless permutations of these concepts based on the circumstances of individual environments. All these permutations focus on current and projected decoupling point buffer integrity, order synchronization, and the promotion and protection of the flow of relevant information and materials.

CHAPTER 11

DDMRP Impacts on the Operational Environment

So far this book has described how DDMRP works with regard to supply order generation and management. This chapter is about how DDMRP affects the larger operational environment around it, particularly vertically integrated manufacturers or larger manufacturing environments that have complex scheduling and shop floor execution needs.

DDMRP Strategic Buffer Criteria

Appreciating the unique capabilities of DDMRP buffers and the impact they can have throughout a larger manufacturing environment requires a revisit of the basic elements of the buffers. The previous five chapters have served to describe the positioning, sizing, and operational aspects of these types of buffers. These chapters have defined the necessary conditions for a stock buffer to be a strategic buffer and effectively minimize or mitigate the bullwhip effect. These necessary conditions can be summarized into six critical tests for a stock buffer to be called a strategic DDMRP replenishment buffer. Without meeting these conditions, a buffer is not DDMRP compliant, will not sufficiently dampen the bullwhip effect, and will often force expensive performance-erosive compromises into an environment.

The Decoupling Test

As previously mentioned, decoupling is about creating independence between two linked processes, events, or areas. A stock buffer that passes this test stops the transference and amplification of variability (demand signal distortion and supply continuity variability) up and down the chain as well as breaks the lead time dependency equation for buffer sizing and supply order generation.

The Bidirectional Benefit Test

Well-managed decoupling point buffers provide benefits for both consumers and suppliers of the position. DDMRP buffers clearly pass this test. Processes or areas that consume stock from the position get instant availability and/or shorter times, as the buffered positions are intended to always have stock available. Processes or areas that supply the buffered positions get a consolidated demand signal that corresponds to actual need and order patterns.

The Order Independence Test

Order independence distinguishes inventory that is assigned to a planned stock position from work in process. Work-in-process inventory is already committed to a particular order—it is unavailable to any other orders (without making a diversion between orders). A true planned stock position holds inventory for any potential requirement from a consumer, whether it be different customers or different parent items. This is of particular importance in environments that have shared resources and shared components.

This is essentially the difference between a make-to-stock system and a make-to-order system. Both systems have inventory in them beyond the purchased material level, but the nature or flexibility of the inventory is much different. This means that a DDMRP system is primarily a make-to-stock system with enhanced rule sets about how the strategic stock positions are placed, sized, and managed.

The Primary Planning Mechanism Test

Strategically decoupled stocking points create shorter and independently managed planning and execution horizons. All buffer planning (supply order generation) elements (inputs and outputs) are performed at the buffer itself. Inputs include the buffer parameters and settings and all relevant demand, on-hand, and open supply information, whereas outputs would simply be a supply order with a quantity to restore the position and request date for receipt.

The Relative Priority Test

Both for planning (supply order generation) and execution (supply order management), a buffered position must provide a sense of priority for the position itself as well as an understanding of how that priority relates to other buffered positions. In DDMRP, this is provided through a general reference of color and a discrete reference of percentage. The relative priority test is about a buffered position's ability to convey relevant information that allows operations personnel to quickly determine which items require the most attention or scarce materials or capacity.

The Dynamic Adjustment Test

Any system that is expected to endure must be able to adapt to the complexity and volatility experienced in today's supply chains. In DDMRP, the strategic stocking points are the primary shock absorption system in the promotion and protection of flow in a Demand Driven Operating Model. To that extent the shock absorption capability must be able to adapt (raise or lower its protection level) as change occurs in the environment.

Note that in DDMRP there is one exception to this test: replenished override parts. These parts or positions, due to specific imposed restrictions, have static levels.

DDMRP Versus Safety Stock and Order Point

Now that we have the criteria by which to judge a DDMRP strategic replenishment buffer, some comparisons can be made against more conventional forms of stock buffering and order generation. Demand Driven MRP has frequently been compared with conventional safety stock systems or order point systems. How do these types of conventional mechanisms stack up to the DDMRP buffer criteria?

Safety Stock and the Buffer Criteria

The *APICS Dictionary* defines safety stock as:

> 1) In general, a quantity of stock planned to be in inventory to protect against fluctuations in demand or supply. 2) In the context of master production scheduling, the additional inventory and capacity planned as protection against forecast errors and short-term changes in the backlog. Over-planning can be used to create safety stock. (p. 154)

While there are two definitions, the first definition is incredibly broad and lacks any real substance. The authors acknowledge that limiting the definition of safety stock to the first definition would mean that DDMRP buffers would qualify as safety stock positions with clearly defined rules about placement, sizing, and management.

The second definition, however, defines the conventional approach to safety stock and characterizes what most people refer to when making comparisons between DDMRP and safety stock. While many safety stock approaches have subtle (and often proprietary) differences, they are all built around the same basic principles; defining a stock level to cover against variability (demand and supply) over a defined length of time—typically a planning horizon. The longer the planning horizon, the greater the rate of forecast error and the higher the safety stock. Safety stock supplements supply orders that are generated by a demand allocation explosion using the projected available balance; these are typically forecasted or planned orders that are primarily driven from a forecast. Thus the conventional safety stock approach is a planned supplementary inventory position to guard against variation. If forecast error is high, then the safety stock supplementary position can become quite an extraordinary financial commitment to statistically cover that error.

How does the conventional approach to safety stock relate to the DDMRP strategic buffer criteria?

- **The decoupling test—FAIL.** Safety stock fails this test. The planning horizons and lead time equations are not defined or impacted by the placement of safety stock positions.
- **The bidirectional benefit test—FAIL.** Safety stock fails this test. Safety stock is designed to protect in only one direction. Safety stock is designed to protect supply continuity. It is focused purely on protecting the consumption side of the position. But just how much protection does it really provide?

  It can easily be argued that safety stock actually exacerbates one element of the bullwhip effect—demand signal distortion. The safety stock level becomes the "new zero" that MRP systems attempt to net against. This means that when the available balance is projected to be below the safety stock level, an order is immediately generated and coded as an expedite. Thus any penetration into the safety level creates additional orders with a high degree of urgency. This attribute can create massive amounts of confusion, distortion, and noise for supplying resources to the safety stock position, particularly as the picture changes from MRP run to MRP run. Thus the use of safety stock can be an amplifier to demand signal distortion.

This is contrary to DDMRP thinking, where current and projected on-hand penetration into the safety level (red zone) is expected to be routine. In DDMRP, as the penetration becomes more severe, the focus is on expediting existing supply orders, not creating additional supply orders.

Additionally, conventional safety stock mechanisms are most commonly implemented only at the purchased and end item levels. This limits their dampening effect on supply continuity variability to only those levels, opening the door for nervousness and delay accumulation to be experienced to a larger degree between those levels.

- **The order independence test—PASS.** Safety stock passes this test. Safety stock inventory is not predisposed to any particular order. It is available to any source of consumption.
- **The primary planning mechanism test—FAIL.** Safety stock fails this test. While the conventional safety stock mechanism does generate supply orders, those orders are only supplementary. The vast majority of planning activity (supply order generation) for positions that are safety stocked are forecasted or planned orders and have no relationship to the safety stock buffer itself.
- **The relative priority test—FAIL.** Safety stock fails this test. While the conventional safety stock mechanism will inform that you have positions under safety stock level, there is no ability to prioritize against them. Every penetration is deemed urgent and requires immediate attention. When everything is a priority, there is no priority. It is impossible without significant additional effort and analysis to determine just how urgent each situation is and how much more urgent one situation is over another.
- **The dynamic adjustment test—PASS/FAIL.** In the authors' experiences most safety stock mechanisms are user defined and static. There are glaring exceptions, particularly in the fast-moving consumer goods segment where the safety stock equations are tied to changes in a rolling forecast, forecast error rate, and the days of supply intended for safety stock.

Conventional safety stock mechanisms simply do not score well against the DDMRP buffer compliance criteria, the criteria that have been formulated to better protect and promote flow in the more complex and volatile circumstances of the twenty-first century.

Order Point and the Buffer Criteria

Let's now turn our attention to the comparison of DDMRP buffers to order point. Similar to what was done with safety stock, order point systems will be compared with the DDMRP buffer compliance criteria.

The *APICS Dictionary* defines order point as:

> A set inventory level where, if the total stock on hand plus on order falls to or below that point, action is taken to replenish the stock. The order point is normally calculated as forecasted usage during the replenishment lead time plus safety stock. (p. 117)

At face value this definition seems much more compatible with the DDMRP rule set:

- **The decoupling test—PASS.** Order point passes this test. In conventional order point, the planning horizons and lead time equations are defined by the placement of order point positions. MRP stops the explosion at order point positions because those positions are designed to have on-hand inventory available. This is consistent with a DDMRP approach.
- **The bidirectional benefit test—FAIL.** Order point fails this test. As per the APICS definition, order point incorporates a conventional safety stock mechanism. That incorporation makes order point fail this test for the reasons previously described (exacerbating the demand signal distortion element of the bullwhip effect).
- **The order independence test—PASS.** Order point passes this test. The inventory at the order point position is available to all potential consumers. This is consistent with a DDMRP approach.
- **The primary planning mechanism test—PASS.** Order point passes this test. All supply orders are generated at the order point position including the problematic safety stock orders.
- **The relative priority test—FAIL.** Order point fails this test. Generated

supply orders are simply quantities and date requirements without a way to compare and contrast them. Additionally, the incorporation of a safety stock mechanism further complicates the priority picture.

- **The dynamic adjustment test—FAIL.** Order point fails this test. Most order point systems are user defined and static.

According to the DDMRP buffer compliance criteria, order point systems appear to be closer in nature to DDMRP systems—a distant cousin so to speak. Yet we should be careful in inferring that the closer relationship implies similar results.

Even when order point passes a criterion, there are significant differences worth exploring. These differences will create a major disparity in performance between the two. This is evident when exploring the primary planning mechanism test. While order point passes the planning mechanism test, there are major differences in the nature of the ordering equation (above and beyond the inclusion of the safety stock trigger) that should be understood. Order point systems use a different equation to determine whether a buffer requires resupply.

In order point the equation is limited to on-hand plus on-order. There is no demand element in the equation. No past due sales orders. No sales orders due today. No spike qualification. This means that demand in order point systems is consumption-based and purely historical. It is only recognized after it has occurred. This happens in the form of an on-hand adjustment in the next day's order point equation.

Yet the demand elements in the net flow equation of DDMRP are highly relevant pieces of information. When visibility to relevant information is obscured, we know that there is a direct relationship to the amount of variability experienced, and that is the case with order point systems. They are subject to more demand variability because they disregard highly relevant and accurate sales order information. The susceptibility to variability comes at the price of higher total inventory requirements to cover to protect the buffer position.

DDMRP Impacts on Scheduling

With these six DDMRP buffer characteristics in mind, we will now turn our attention to how DDMRP can affect schedules at both the higher master level and resource scheduling area.

DDMRP and Master Production Scheduling Assumptions

Within the conventional planning schema, the output of the master production schedule (MPS) and Material Requirements Planning is important to understand.

The *APICS Dictionary* defines master production schedule as:

> A set of planning numbers that drives material requirements planning. It represents what the company plans to produce expressed in specific configurations, quantities and dates. (p. 101)

Figure 11-1 displays the conventional relationship between the MPS and MRP as described in Chapter 3. The output of all this activity relies on three crucial assumptions for success. These assumptions must be examined at depth in light of what we now know about conventional planning limitations.

- **Assumption 1.** Demand signals are known and accurate.
- **Assumption 2.** Lead times for supply order release, receipt, and synchronization are realistic.
- **Assumption 3.** Material and capacity are available on the specified dates.

Essentially any schedule at any level in operations makes these three basic assumptions. Why release a schedule if these basic assumptions are not behind it? The relative validity of these three basic assumptions, however, combines to determine just how realistic a schedule is. When the schedule is more realistic, it is more likely to be maintained. Conversely, the less realistic the schedule, the more likely that schedule will be disrupted. Schedule disruptions lead to performance erosion and costly compensation at all levels.

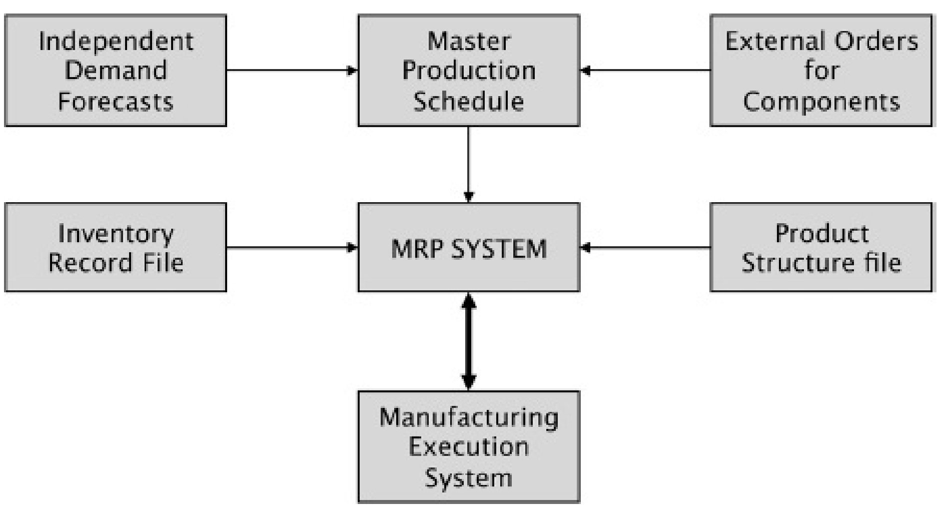

FIGURE 11-1 The conventional MPS-MRP schema

How valid are these assumptions in the conventional MPS and MRP approach? The answer to this question will determine just how realistic the output of the conventional approach is on a routine basis.

- **Assumption 1. Demand signals are known and accurate.** Tying order generation and scheduling directly to forecast means that actual demand will vary from the forecasted demand used to generate the schedule. The longer the planning horizon, the larger the variance between forecasted and actual. As these variances occur, the demand signals change with every MRP run, creating massive numbers of adjustments throughout the environment.
- **Assumption #2. Lead times for supply order release, receipt, and synchronization are realistic.** With no decoupling, delays frequently accumulate, dramatically affecting when orders can be released with full allocation. Safety stock is rarely placed at the intermediate levels in order to provide even partial supply variability dampening. Positions are netted to zero, leaving no margin for error. With everything coupled together and no margin for error, the schedule is much more complex and fragile.

  Furthermore, with no execution capability built into MRP, there is no ability to see how those potential delays will affect the environment in the near term. There is little to no visibility about a problem until the problem has already been encountered, and without decoupling there are guaranteed to be more problems. Under these circumstances synchronization (and flow) quickly breaks down.
- **Assumption 3. Material and capacity are available on the specified dates.** When synchronization breaks down, material and capacity are frequently not available as planned. Material arrives late or gets diverted to cover shortages elsewhere. Capacity is frequently not available due to schedule slides and deviations.

Simply stated, the output of the conventional MPS and MRP process is just not very realistic to start with, and it quickly gets worse in execution, running rampant through operations. As stated in previous chapters this conventional approach will work only if everything goes exactly according to plan. Is this news to planners? Of course not; that is why compromises and workarounds are present in almost every company using this process. Is this news to operational personnel? Of course not; that is why they blame planning for so many of their problems.

How does the DDMRP approach stand up to the same set of basic assumptions? The assumptions don't change; the operational approach does. Will that operational approach produce a realistic output?

- **Assumption 1. Demand signals are known and accurate.** The use of qualified sales orders means that demand signals are much more relevant, accurate, and timely.
- **Assumption 2. Lead times for supply order release, receipt, and synchronization are realistic.** The use of decoupling points creates shorter independently planned and managed horizons. That means that less variability (demand and supply) is passed through the system. This results in synchronization dates that are at the same time more realistic throughout the system yet less important due to the cushion at the buffers. Additionally, the use of decoupled lead time for buffer sizing means that the buffers are realistically sized at the decoupling points to absorb demand and supply variability and maintain the integrity of the decoupling effect.
- **Assumption 3. Material and capacity are available on the specified dates.** DDMRP plans strategic positions to always be available. At decoupling points, DDMRP is designed to never net to zero. DDMRP screams with multiple alerts at planning and operations personnel if those positions are anywhere near zero for net flow and on-hand. The DDMRP buffers represent points of stored capacity and materials. If these positions are maintained with on-hand always available, then the stored capacity and materials at those positions are always available. This means that under DDMRP, this assumption is valid most of the time for the strategic points in the model.

  Additionally, the use of DDMRP's execution components brings degrees of visibility to open supply orders that must be expedited to maintain stock buffer integrity and to meet critical synchronization needs.

DDMRP moves an environment from the conventional statement of "what we can and will build" to a statement of "the capability to build what we can and will sell." Indeed, the whole notion of a conventional master production schedule tied to a supply order generation calculator (MRP) evaporates under a Demand Driven Operating Model. In its place we find a "master settings" component of Demand Driven Sales and Operations Planning (DDS&OP) tied to the supply order generation calculator of DDMRP. Figure 11-2 depicts the master settings feed to DDMRP from DDS&OP.

There are three primary inputs provided by the master settings: buffer profiles, part demand data, and part profile assignment. Buffer profiles are the groupings and settings for replenished parts (part type, variability, and lead time). Part demand data have two elements: planned adjustment factors (if applicable) and average daily usage. Planned adjustment factors are factors to be applied to the ADU of parts or groups of parts. Average daily usage is the average rate of use for each replenished part (past, forward, or blended). Part profile assignment is the assignment of each replenished part to a particular buffer profile.

These inputs are combined with the inventory record file (on-hand and open supply records), the product structure file (bill of material), and external orders for components (sales orders) to perform the net flow equation and generate execution alert reports for each part. Supply orders are conveyed to scheduling (manufacturing orders) and execution (purchase orders and stock transfer orders). Alerts are conveyed to appropriate personnel.

Simply stated, the output of the DDMRP approach produces a much more realistic start to the application of an organization's resources to the fulfillment of market demand. Supply orders are generated based on a model built to promote and protect flow to fulfill known and accurate demand signals. But DDMRP's impact on operations does not end there with valid, prioritized, and synchronized supply orders. In many organizations it will dramatically influence and improve the ability to effectively schedule resources and manage work flow for the best total protection and promotion of flow.

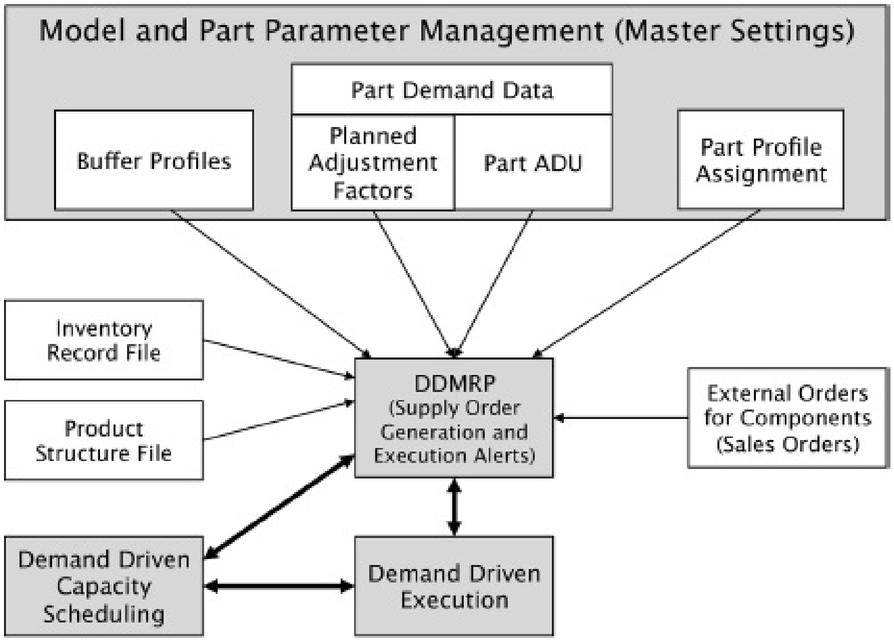

FIGURE 11-2 DDMRP master settings feed

DDMRP Shop Floor Scheduling Implications

With regard to manufacturing orders, the impact of DDMRP on scheduling is not limited to the simple generation of a supply order signal. Conventionally, manufacturing order requirements are generated, and it is up to scheduling to work it out. This presents a huge challenge when resources and materials are shared and there is no relative sense of priority—you simply have to do it all, and MRP assumes that you can, so you do the best you can. Under the DDMRP approach, however, scheduling the application of materials and capacity can be much clearer and even simpler due to two particular impacts.

Impact 1. Decoupled Schedules

So far the DDMRP decoupling effect has been explored at the product structure level in order to generate and manage supply orders. Supply orders exist at the product structure (discrete part number) level. Yet with regard to scheduling manufactured items, there is an additional level that must be understood: the routing level. Routing was defined in Chapter 6 ("Strategic Inventory Positioning").

Simply put, the routing details the way a part number at the lower level of the product structure is converted to a part number at a higher level in the product structure. Figure 11-3 depicts the difference between a product structure view and a product structure view with routing sequences between part numbers. With routings identified, we can see how a part is converted into a higher-level part. The square boxes are the part numbers, and the rounded shaded boxes are identified activities and their relative sequence.

In this instance, converting a purchased item (PPA) into an intermediate component (ICA) requires three activities in a particular sequence: cut, form, and grind. Most routings would include time per piece as well as other pieces of relevant information for the resources performing the action or scheduling the activity. In this case, that information is immaterial to illustrating the DDMRP impact on scheduling. The conversion of the intermediate component (ICA) into the finished item (FPZ) requires three additional sequenced steps: prep, paint, and label.

When decoupling points are placed at the product structure level, they also serve to decouple at the routing level. The routing level is the level at which resource scheduling is created. Having fewer dependencies to schedule dramatically simplifies the resource scheduling process because there are fewer things that have to fit together.

In Figure 11-4, scheduling two groups of three processes independently should be easier than scheduling six processes together. There are simply less dependencies. A simple analogy: Is it easier to schedule two meetings with three busy people or one meeting with six busy people? And that is an oversimplification, because in resource scheduling we need to find time slots in a particular sequence. The chances that open time slots can be found to accommodate three resources in the right sequence are routinely better than routinely finding six open time slots in the right sequence.

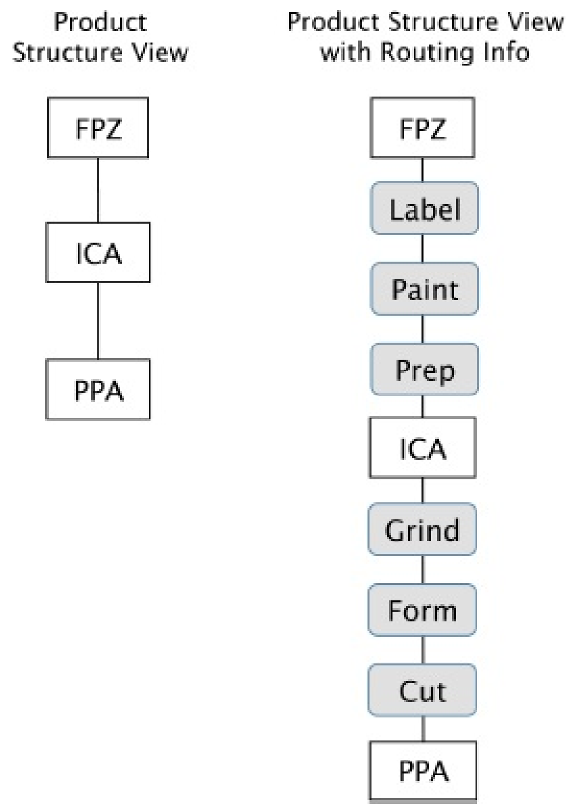

FIGURE 11-3 Routing information in FPZ product structure

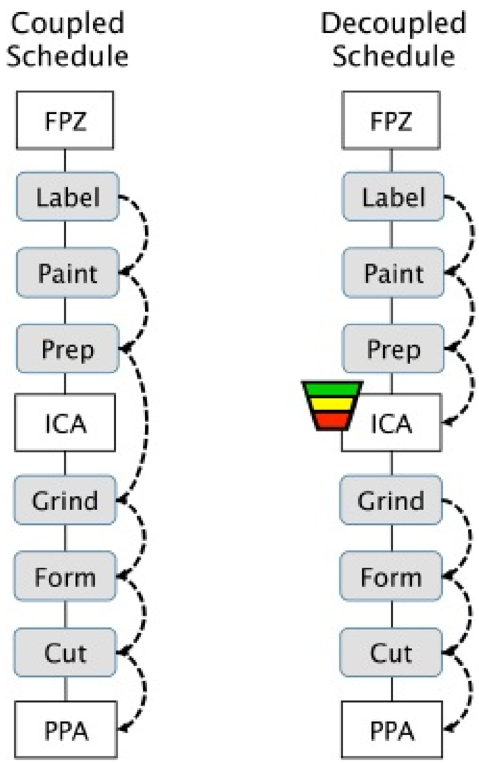

FIGURE 11-4 Coupled versus decoupled scheduling

Figure 11-5 depicts a much larger network of dependencies, one with many different integration points and a significant amount of concurrent activity—a scheduling puzzle across shared resources. The rounded shaded boxes are resource assignments. The rectangular boxes with no shading are part numbers in the bill of material. Shared resources can be identified using the numerical assignment. For example, the same resource (031) is required on the routing for component S12 and component S21.

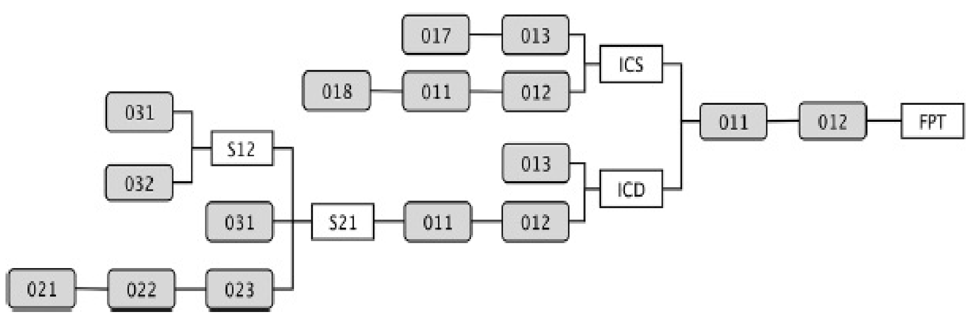

FIGURE 11-5 A scheduling puzzle

[Figure 11-6](#) depicts the same environment but with a decoupled scheduling approach. Decoupling points have been placed at part numbers in the product structure. One large scheduling puzzle has been replaced with four smaller and easier-to-schedule portions defined by the decoupling point placements at components S21, ICD, and ICS.

Perhaps a better depiction of the decoupled schedule approach is to lay them side by side, as they are independent of each other due to decoupling point placement. [Figure 11-7](#) depicts this view.

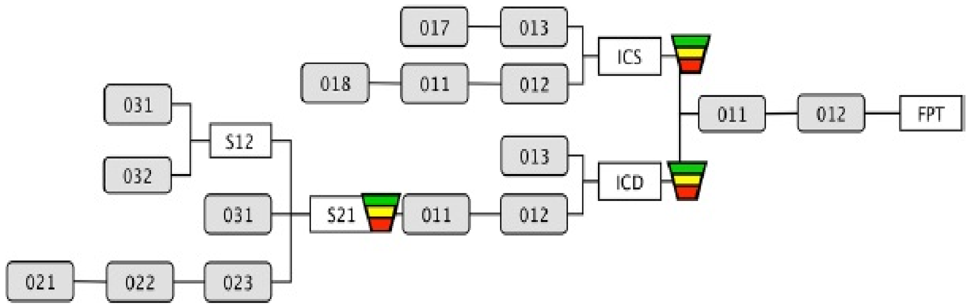

FIGURE 11-6 A decoupled scheduling approach

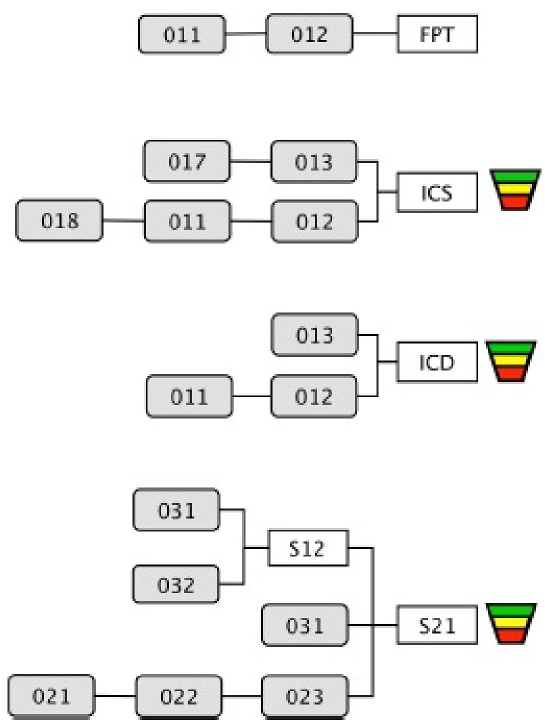

FIGURE 11-7 Decoupled scheduling legs

Yes, there are shared resources between these scheduling legs, but the decoupling point buffers will help build a deconflicted schedule across those resource schedules based on net flow priority. This will be discussed in the next section ("Visible and Priority-Based Scheduling Sequencing").

The stability of the schedule is also impacted by decoupling. The one thing that is constant is change. Run rates can vary, quality issues can occur, and resources can go down from time to time. These events impact schedules. As discussed previously, any change in a highly dependent system creates a ripple that affects everything that is dependent. This was previously referred to as nervousness and was illustrated at the MRP level. Changes in date or quantity in one supply order rippled out to affect dependent supply orders (changes in date and quantity or cancellation), even creating additional supply orders.

Nervousness also exists at the routing level. This is called scheduling nervousness. Changes in scheduled activity or delays at certain operations ripple out to affect the schedule. Advanced planning and scheduling systems perfect this damaging effect. Small changes in schedules are transmitted to the rest of the system, making the schedule a constantly changing work in process. The more frequent the changes, the more chaotic the schedule. If the schedule is constantly changing, then there is no schedule. The more dependencies scheduled together, the more change there will be.

In a decoupled schedule, the changes that occur in one scheduling leg are isolated from the other legs. The exception is when there are shared resources among the legs. Under this circumstance, DDMRP will help mitigate the impact of that variability by providing clear visibility to which schedule has priority for that particular resource. This will be discussed later in this chapter in the "DDMRP and WIP Priority Management" section.

Impact 2. Visible and Priority-Based Scheduling Sequencing

As discussed in Chapter 9, DDMRP's net flow equation determines supply order generation. As supply orders for manufactured items are being generated, the net flow position of each item can be used to establish scheduling sequence across a shared resource base.

Figure 11-8 represents a manufacturing plant with shared resources. This plant makes nine different finished items (FPA, FPB, FPC, FPD, FPE, FPF, FPG, FPH, FPI). All finished items go through the same assembly area. This plant also has 18 different intermediate components. Nine are fabricated (SAA, SAB, SAC, SAD, SAE, SAF, SAG, SAH, SAI), and the others are machined

(ICA, ICB, ICC, ICD, ICE, ICF, ICG, ICH, ICI). All finished items consist of machined and fabricated parts.

Figure 11-9 is the DDMRP planning screen for the finished items in the plant. Supply orders for three items are being recommended (FPB, FPE, and FPA). The planning priority creates the relative priority between the orders and establishes the assembly scheduling sequence. FPB will be scheduled in assembly first.

The same process is applied to the establishment of the machining and fabrication schedules. Figure 11-10 is the DDMRP planning screen for intermediate items. Two machined and two fabricated items are being recommended for replenishment. The planning priority will establish the scheduling sequence for both areas. ICB will be scheduled in machining first, and SAF will be first on the schedule for fabrication.

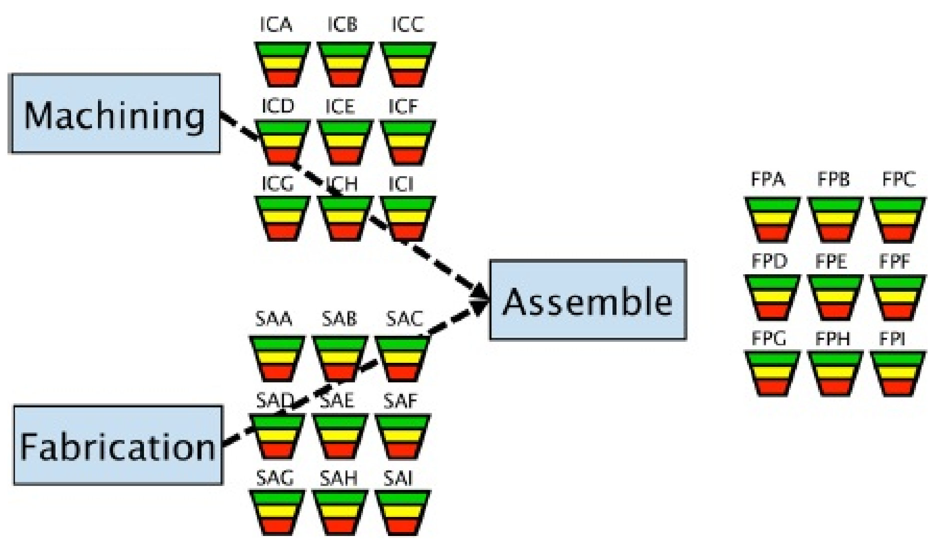

FIGURE 11-8 A sample plant with shared resources

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Recommendation	TOG
FPB	YELLOW (49%)	4012	5453	3200	6265	6500	12765
FPE	YELLOW (68.7%)	4054	3358	540	6872	3128	10000
FPA	YELLOW (70%)	530	4021	213	4338	2162	6200

FIGURE 11-9 Net flow priority for finished items

Part#	Planning Priority	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Recommendation	TOG
ICB	YELLOW (58%)	12000	16359	8000	20359	14743	35102
SAF	YELLOW (72%)	5532	900	960	5472	2128	7600
ICH	YELLOW (74%)	3721	1530	713	4538	1594	6132
ICD	YELLOW (77.3%)	5714	2200	562	7352	2154	9506

FIGURE 11-10 Net flow priority for intermediate items

Finite Scheduling with DDMRP?

Any manufactured item stock buffers should not be thought of simply as units of stock. The only way that you can have stocked positions of manufactured items is to have invested both capacity and materials to create the stock. Thus these buffers should be thought of as storage tanks of capacity and materials. Furthermore, if these buffers are placed at strategic points, then they are strategic storage tanks of capacity and materials. Additionally, these strategic tanks are designed to be maintained at certain prescribed levels based on DDMRP buffer sizing criteria.

What we have learned about DDMRP buffers up to this point can easily be translated into capacity terms. Figure 11-11 relates the buffer of a manufactured item to the capacity of a particular resource.

In this case the buffer of part 123 is being related to assembly capacity. Part 123 takes 30 minutes per piece on average in assembly. That means that a top of green of 455 pieces represents a total of 13,650 minutes of assembly time. The average on-hand of 155 translates to 4,650 minutes of assembly time. This is the average amount of stored assembly capacity that is expected in the buffer. A minimum order size of 120 pieces corresponds to the green zone. This means that each order for resupply will represent at least 3,600 minutes of capacity.

This consideration brings up an interesting extension to or perhaps a redefinition of the term “finite scheduling.” The *APICS Dictionary* defines finite scheduling as:

> Assigning no more work to a work center than the work center can be expected to execute in a given time period. (p. 63)

Part #	Assembly Minutes per Part	Top of Green	Top of Green Assembly Time	Average OH	Average OH Assembly Time	Minimum Order Size	Assembly Minutes per Average Order
123	30	455	13650 minutes	155	4650 minutes	120	3600 minutes

FIGURE 11-11 Translating buffer quantities to capacity terms

This definition is focused solely on the scheduling of a resource. But can this definition be expanded to include the way DDMRP recommends supply? The critical elements of the definition distill down to scheduling capacity to the defined level and not beyond. DDMRP manufactured item buffers are strategic storage tanks of capacity that are replenished up to a level (top of green) and not beyond.

The DDMRP buffers are strategically placed; they are resupplied up to a certain point and not beyond. Supply orders for their replenishment essentially direct activity in the system. This is interesting and similar to a specific type of finite scheduling known as “drum scheduling.” Drum scheduling was most famously introduced to the world in 1984 in the revolutionary book, *The Goal*, by Eliyahu Goldratt and Jeff Cox. The *APICS Dictionary* defines drum scheduling as:

> The detailed production schedule for a resource that sets the pace for the entire system. The drum schedule must reconcile customer requirements with the system’s constraint(s). (p. 52)

Yet in DDMRP, less actual resource scheduling precision is required because the buffers allow for time and flexibility within short ranges of time. They naturally allow for a deconflicted scheduling sequence tied to strategic buffer status (planning and execution priority). Thus DDMRP manufactured item buffers can be defined as “drums.” They are strategically placed. They set the pace of the system. They are supplied up to a defined point and not beyond. For environments that have a fair amount of manufactured item buffers, DDMRP can be characterized as a simple, even elegant, finite scheduling alternative.

Additional Scheduling Sequence Impacts

Even when there are major scheduling complications, DDMRP can provide additional insight into the sequence. For example, in food processing, a major complication is something called allergen sequence.

An allergen sequence defines the order in which food processors must process foods with certain known allergens. For example, all items that do not contain eggs should be run in advanced of all products that do contain eggs. In this case a sequence number can be provided to each item. For example, products without eggs receive a 1, while products with eggs receive a 2.

This sequence assignment can then be combined with the DDMRP planning priority to create a sequenced schedule that is both safe to process and aligns (as best as possible) to actual buffer requirements. Figure 11-12 shows a DDMRP planning queue that is sorted by allergen sequence. All items recommended for resupply are now assigned a sequence based on their allergen assignment.

Despite having the highest planning priority, "NutSurprise" is the last on the sorted allergen sequence list. Running out of sequence would require significant delays for required cleaning activities. This could erode capacity and jeopardize additional stocked positions. In circumstances such as this, it might require the use of higher variability profiles in order to compensate for potential sequencing conflicts.

Part#	Planning Priority	Allergen Sequence	On-Hand	On-Order	Qualified Demand	Net Flow Position	Order Recommendation	TOG
BerryTart(ras)	YELLOW (65.4%)	1	5032	900	960	4972	2628	7600
PieFill(apple)	YELLOW (69.1%)	1	3421	1530	713	4238	1894	6132
OatmealChew	YELLOW (68.9%)	2	5214	1900	562	6552	2954	9506
NutSurprise	YELLOW (58.0%)	4	12000	16359	8000	20359	14743	35102

FIGURE 11-12 DDMRP resupply orders sorted by allergen sequence

Of course, there will be environments and circumstances that dictate additional and more detailed scheduling capability. These include:

1. When more detailed scheduling will result in better lead time and variability control feeding into a stock position. This will allow for the minimization of inventory at the stock position.

2. When series of resources are needed to create a final product and there is no ability to decouple before the finished level in the product structure. In this case, more detailed scheduling efforts might better protect customer commitments. This is particularly necessary when the resources are shared across many items. Config-ure-to-order environments often qualify under this reason.

3. When there are shared resources involved in the production of make-to-stock and make-to-order items (mixed-mode operations). More detailed scheduling will often be needed to better deconflict capacity requirements between make-to-stock and make-to-order items.

When these circumstances are present, a Demand Driven Operating Model uses *control point* scheduling. The *APICS Dictionary* defines control points as:

> Strategic locations in the logical product structure for a product or family that simplify the planning, scheduling, and control functions. Control points include gating operations, convergent points, divergent points, constraints, and shipping points. Detailed scheduling instructions are planned, implemented, and monitored at these locations. (p. 33)

Control points don't decouple. They are places to transfer and amplify control within a given area so the transference and amplification of variability within that area can be minimized. Typically they are placed between decoupling points or between decoupling points and customers. These environments and circumstances in addition to control point determination and scheduling are thoroughly described in the book *Demand Driven Performance: Using Smart Metrics* by Debra Smith and Chad Smith.

DDMRP and WIP Priority Management

After a schedule has been determined, DDMRP can further influence manufacturing resources with regard to released work sequencing. By giving resources or those managing the resources the visibility to the on-hand buffer status corresponding to the orders on their respective dispatch lists, they can easily and quickly determine if the sequence should be altered based on the changes occurring in the environment. Figure 11-13 is a scheduled sequence based on planning priority from the previous example in Figure 11-9.

Part#	Planning Priority
FPB	YELLOW (49%)
FPE	YELLOW (68.7%)
FPA	YELLOW (70%)

FIGURE 11-13 Original scheduled sequence

Should this schedule simply be frozen? Should manufacturing resources always work the sequence in the schedule? Yes, if there are no relevant changes between the time in which the sequence was created and the activity is to be performed. If there are relevant changes in that time range, then visibility must be promoted in order to convey any changes in priority and sequence. Figure 11-14 is an example of the minimum amount of relevant information that resources or those managing the resources would need to see to shift the sequence of the released manufacturing orders. In this case, the released orders have been sequenced not in the order they were scheduled but in the order that corresponds to their buffer status.

Order #	Item #	Buffer Status
MO 819-87	FPA	RED (35.3%)
MO 832-41	FPB	GREEN (133.7%)
WO 211-72	FPE	GREEN (135.1%)

FIGURE 11-14 Released manufacturing orders sequenced by on-hand buffer status

FPA's on-hand situation is becoming critical. A situation such as this may dictate a change in the sequence in which these manufacturing orders are worked. At a minimum it should prompt discussions between appropriate personnel.

Summary

This chapter illustrated the impact DDMRP can have on the operational resources that interact with it. It is the explanation and depiction of how DDMRP fits into a Demand Driven Operating Model—essentially providing the heartbeat of the activity in the model. To summarize, Figure 11-15 is a graphical comparison of conventional MRP, Lean, and DDMRP.

The graphic shows the independent planning horizons between decoupling points. These independent horizons provide demand and supply variability mitigation while allowing for synchronization within the horizon. Supply order generation is illustrated by the dashed arrows pointing to the next lower level decoupling points. The elements of visible and collaborative execution are also illustrated, with the on-hand buffer reports pointing to each leg in the bill of material. This is meant to show that while executing against the initial replenishment signal the statuses of the next buffered positions are visible to allow for course correction if necessary.

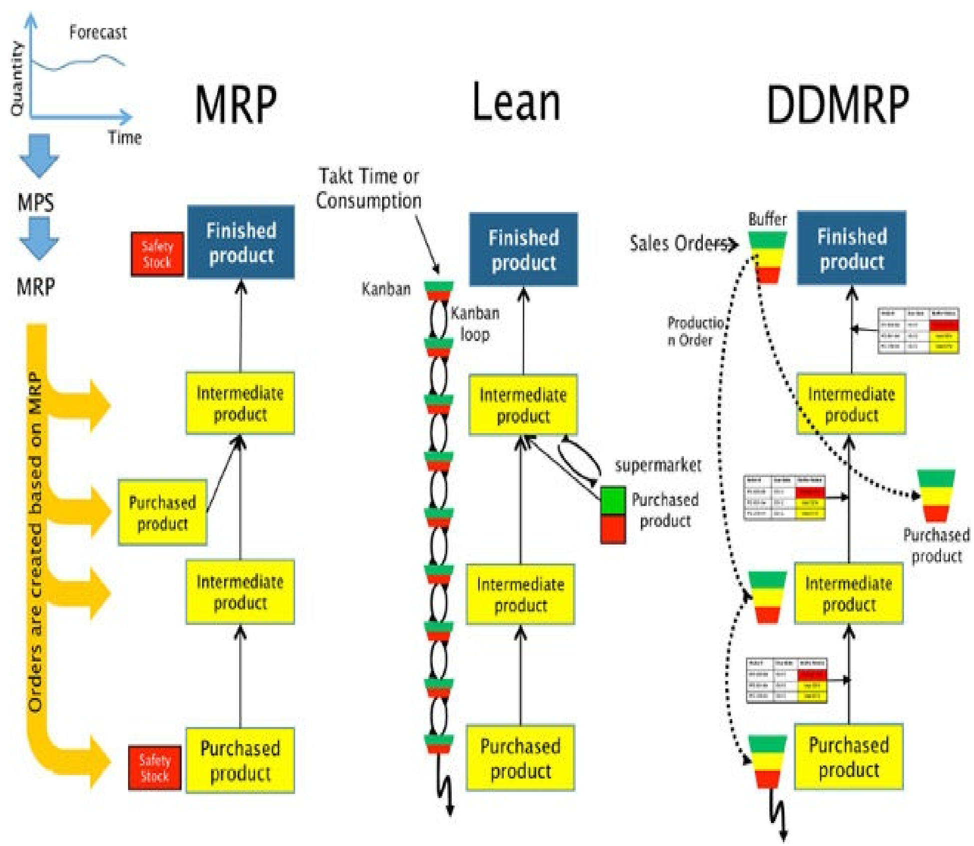

FIGURE 11-15 MRP, Lean, and DDMRP

CHAPTER 12

DDMRP Metrics and Analytics

In the first section of Chapter 7, the two critical prerequisites are established for inventory to act as an asset for the flow of relevant information and materials: placement and sizing. Placement determines the points of independence in the planning, scheduling, and execution process, and the sizing of the inventory position allows that independence to be maintained.

With regard to placement, Chapter 6 discusses key factors to consider when approaching the placement of decoupling points. This is a strategic design or blueprint that is not to be taken lightly. Once that design is completed and implemented, how can visibility be created to make sure that the buffer levels at those decoupling points are protecting and promoting the flow of relevant information and materials for the best return on investment in our complex and dynamic environments? Are the right signals being conveyed without distortion in a timely fashion? Are the right materials available when needed? Is the inventory in excess? Is the blueprint performing as designed? How can it be better? These are the key questions for the metrics and analytics associated with DDMRP. If DDMRP designs are fundamentally different from convention, then it hints that at least some new metrics and analytics may be required to answer these questions.

Measuring Relevant Information (Signal Integrity)

DDMRP planning is the process of generating supply order requirements primarily driven by the decoupling point net flow position. There is an assumption that a system cannot have the relevant materials without first having the relevant information (supply order signal) to act upon. Relevance is determined by timing and accuracy. If there is a delay in conveying the supply order signal, then the DDMRP model is not operating with relevant information and is less responsive. If the supply order signal is inaccurate (too high or too low), the DDMRP model is not operating with relevant information and is either less responsive or wasteful. Thus supply orders should be conveyed as soon as possible and in proper quantity as per the DDMRP model.

Figure 12-1 is an example of a report used to monitor the integrity of supply orders for timing and accuracy. A date range is provided on the left-hand side covering a 15-day period (October 1–15). In this example a company makes four types of soup, which are all buffered: Chicken Truffle, Chicken Noodle, Beef Stew, and Minestrone. Within each product category two columns are provided: “DDMRP Recommendation” and “Actual.”

	Chicken Truffle		Chicken Noodle		Beef Stew		Minestrone	
Date	DDMRP Recommendation	Actual	DDMRP Recommendation	Actual	DDMRP Recommendation	Actual	DDMRP Recommendation	Actual
1-Oct	GREEN		GREEN		OTOG		GREEN	
2-Oct	YELLOW 1200		GREEN		OTOG		GREEN	
3-Oct	YELLOW 1321	1800	GREEN		OTOG	1200	YELLOW 1900	1900
4-Oct	OTOG		GREEN		OTOG		GREEN	
5-Oct	OTOG		GREEN		OTOG		GREEN	
6-Oct	GREEN		YELLOW 2400		OTOG		GREEN	
7-Oct	GREEN		YELLOW 2600		OTOG		GREEN	
8-Oct	GREEN		YELLOW 2950		GREEN		GREEN	
9-Oct	GREEN		YELLOW 3120		GREEN		GREEN	
10-Oct	GREEN		YELLOW 3500		GREEN		YELLOW 2150	2150
11-Oct	GREEN		YELLOW 4200		GREEN		GREEN	
12-Oct	YELLOW 1700		RED 6000		GREEN		GREEN	
13-Oct	YELLOW 1762		RED 6200	6200	GREEN	1200	GREEN	
14-Oct	YELLOW 1900		GREEN		OTOG		GREEN	
15-Oct	YELLOW 1950		GREEN		OTOG		GREEN	

FIGURE 12-1 Tracking net flow signal integrity

The DDMRP recommendation is the quantity recommended to restore net flow to the top of green on that specific date. A quantity will only appear if the net flow position is in the yellow or red zone. “Actual” is the quantity of the approved supply order on that specific date line. For example, on October 3, Chicken Truffle Soup is calling for a supply order of 1,200 in the “DDMRP Recommendation” column. Its net flow position is in the yellow zone of the buffer and requires 1,200 to be restored to the top of green. A supply order was approved on the same day for a quantity of 1,800.

Three of these soups have problems with signal integrity. As described above, a supply order for Chicken Truffle was approved on October 3. Supply order generation was late by one day and well above (1,800) the required quantity (1,321). An initial requirement of 1,200 was known on October 2; yet there was no corresponding supply order. The next day the net flow position eroded further, driving a larger recommended amount. An order was approved, but that order far exceeded the recommended amount. It drives the net flow position over the top of green (OTOG) for the following two days. In this case there are both timing and accuracy issues.

Chicken Noodle exhibits problematic signal integrity from a timing perspective. On October 6 there was a recommended supply order quantity of 2,400 but no approved supply order. This situation persists until the net flow position is driven all the way into red, requiring a resupply of 6,200. Finally, on that day a supply order is approved for the quantity recommended. In this case there are only timing issues.

Finally, Beef Stew appears to be chronically oversupplied. On October 3 a supply order was approved for 1,200 despite the net flow position being in an OTOG position. This drives the net flow position even further above OTOG. Then again on October 13 another supply order is approved for 1,200 despite no order recommendation and a net flow position in the green.

Could there be good explanations behind the planner's behavior for each of the three parts with signal integrity? Absolutely, but questions need to be asked and reasons need to be documented.

Minestrone appears to be the only item that is operating in accordance with the model and method. October 3 produces a recommendation for 1,900, and on the same day an order is approved for 1,900. The net flow position shows green the next day and continues until October 10, when a recommendation of 2,150 is approved with a new supply order for the same amount.

Measuring Decoupling Point Integrity

The first section of Chapter 7 also provided us with a key piece of information that will become an important foundation for the bulk of these new metrics and analytics. Figure 12-2, repeated here from Chapter 7, shows the Taguchi inventory loss function. Inventory will be an asset to flow at the decoupling points when it is maintained between too little (point A) and too much (point B). Between these two points is an optimal range for inventory to be maintained. Within this range it can absorb variability but is not excessive inventory.

This book describes a way to calculate this range as well as the behavior that should occur over time to maintain the range and even the elements required to make improvements. Chapter 7 provides the sizing logic for DDMRP buffers. Chapter 8 provides the logic for adjustments. Chapter 9 provides the planning logic for the buffers. Chapter 10 presents the rationale behind monitoring the current and projected ability to maintain the integrity of the decoupling point.

Chapter 3 describes a typical effect in conventional planning called the bimodal effect. The bimodal distribution is primarily driven by two inherent traits of MRP: its hard-coded trait to net to zero at each part position and its requirement to make everything dependent in its planning equation. Netting to zero means there is no planned buffer and everything is perfect in quantity and timing. Treating everything as dependent means that change anywhere creates change everywhere. These traits create a constant oscillation with regard to a part's position between too little and too much.

To know how problematic any inventory bimodal distribution really is requires the optimal range for each part to be defined. Without this definition any bimodal analysis is limited to whether a position was stocked out or there was a significant amount of excess. The buffer sizing logic of DDMRP provided in Chapter 7 and the planning logic in Chapter 9 give a clear range in which we can judge how severe the bimodal effect really is.

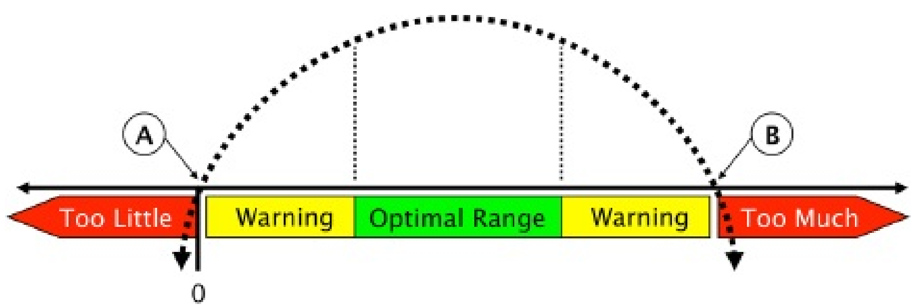

FIGURE 12-2 The inventory value loss (Taguchi) function illustrated

Chapter 7 provides the zone calculations and total buffer size. Chapter 9 illustrates the planning logic to keep the buffer properly supplied and, based on that, the targeted on-hand inventory range. This range is defined as the top of red to the top of red + green. It is this range that becomes the defined "optimal range" in the loss function depiction.

Chapter 9 also shows how to depict a part's past performance against this optimal range. Figure 12-3 is a repeat of Figure 9-43. As described in Chapter 9, the top of the red zone value for the simulated part is 52, its yellow zone is 70, and its green zone is 35. The optimal on-hand range is the value of the green zone. Warning ranges appear on either side of the green zone. The low-level warning range is on the left; that value is the value of the red zone (52). The high-level warning range is on the right; that value is the remaining amount of the yellow zone (35) after the green zone is subtracted. The loss function reaches too little (point A) at zero on-hand and too much (point B) at 122. This value of 122 is the top of the planning yellow zone.

Figure 12-4, which originally appeared in Chapter 9 as Figure 9-44, shows the run chart of the part over the 21-day simulation depicting the on-hand position against the optimal and warning ranges.

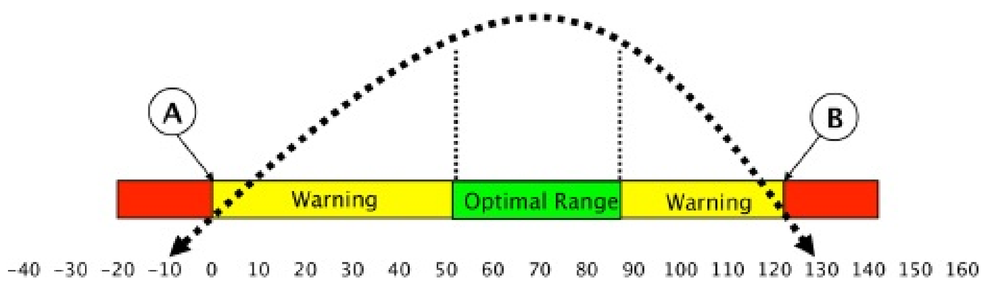

FIGURE 12-3 Optimal range depicted from Chapter 9 planning simulation

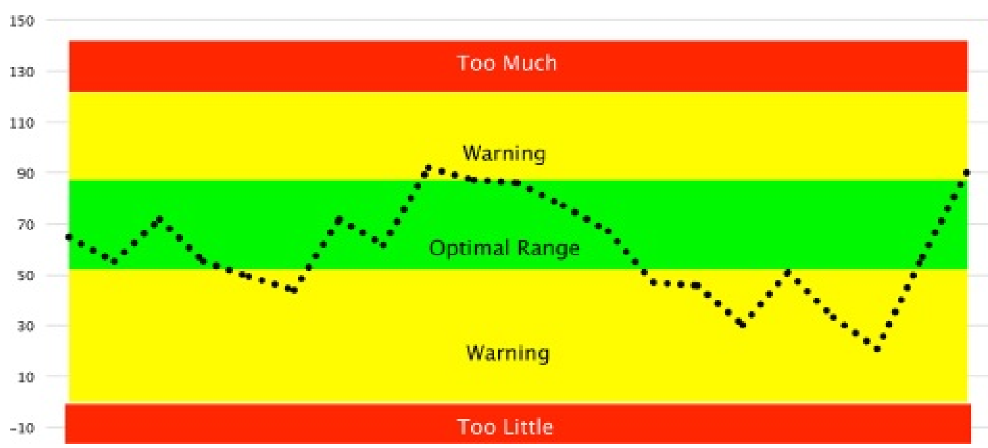

FIGURE 12-4 Simulated part on-hand performance over 21 days

Chapter 10 brings additional clarity to our range analysis with regard to judging the severity of a low on-hand position. The red zone is the embedded safety in the buffer. That safety is expected to be used. The more that on-hand erodes, however, the more severe the situation. The lower warning range is stratified into yellow and red zones. Yellow is the preliminary warning, and red is a more severe warning of a potential stockout. A dark red color is used to depict stockouts. Figure 12-5, originally Figure 10-7, shows the run chart for the simulated part against the execution perspective ranges with the dark red stockout range shown.

Combining the lessons in Chapters 9 and 10, we now have a distinct and specific definition of how to analyze buffer performance over time. Figure 12-6 is the conceptual restatement of the loss function for analysis purposes. This restated version provides tighter ranges to judge how buffers are performing over time.

Each zone serves to provide more specific information about how a buffer is performing against its planned performance—how often it moves out of the nominal range and to what severity. This can provide an indication of how stable and reliable the position is over a period of time. Figure 12-7 depicts the simulated part data from Chapter 9 against the DDMRP analytics color scheme.

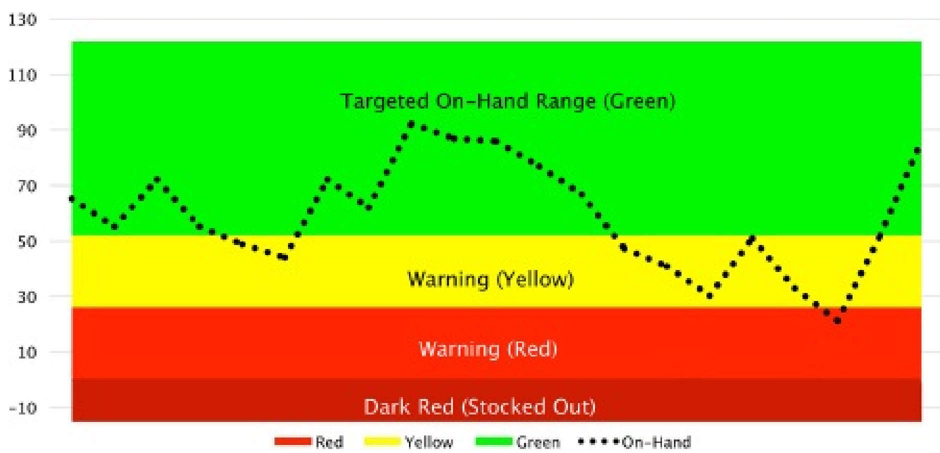

FIGURE 12-5 Planning simulation with execution color scheme

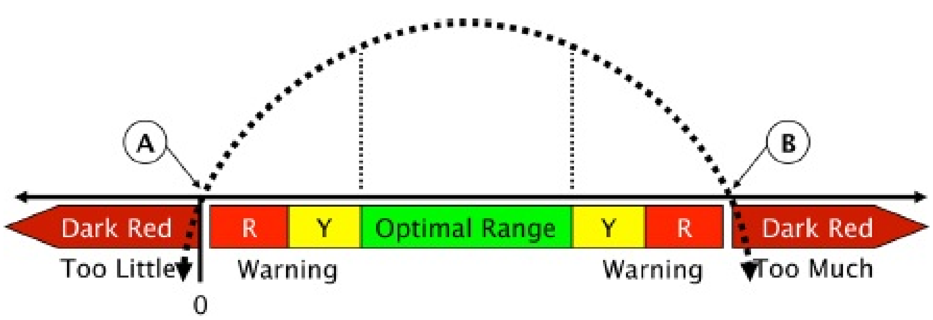

FIGURE 12-6 The restated loss function zones for DDMRP analytics

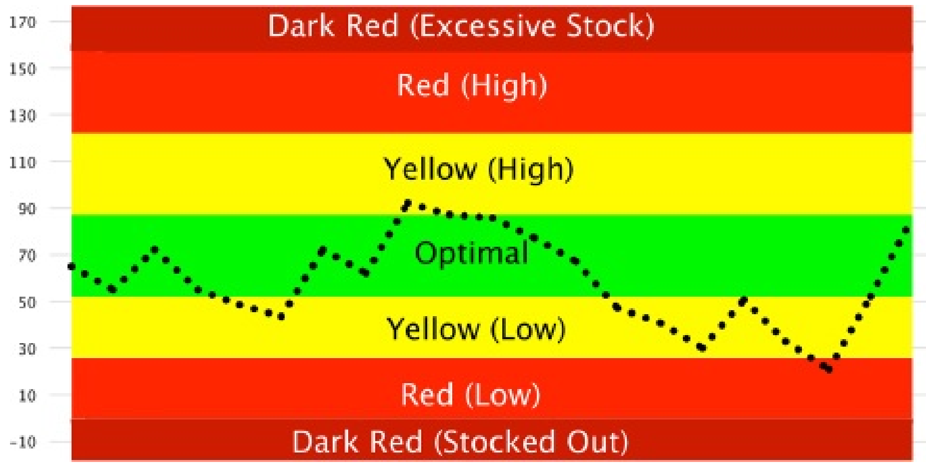

FIGURE 12-7 Simulated part run chart for analytics purposes

The zone labeled "Dark Red (Excessive Stock)" would be triggered when the on-hand level is over the top of the planning green zone of the buffer. The zone labeled "Red (High)" is the value of the planning green zone of the buffer of this part (35). The zone labeled "Yellow (High)" is the remaining amount of the planning yellow zone subtracted from the optimal on-hand range (70 – 35 = 35). The zone labeled "Yellow (Low)" is the upper half of the buffer's planning red zone (on-hand alert threshold set to 50 percent of the red zone). The zone labeled "Red (Low)" is the lower half of the buffer's planning red zone. Finally the zone labeled "Dark Red (Stocked Out)" is an on-hand quantity of zero.

Over the 21-day period, the on-hand level was within the optimal range 12 of the 21 days. Eight times it landed in the low yellow zone. Only once did it land in the low red zone (Day 19). Once it landed in the high yellow zone (Day 9). Run charts like this represent a distribution of occurrences or positions against the loss function scheme. Figure 12-8 depicts the on-hand positions over the 21 days in a distribution curve against the loss function scheme for the simulated part.

With a large enough data set, the distribution of occurrences can be charted, producing a bell curve against that scheme. Figure 12-9 is a conceptual bell curve of occurrences set against the loss function scheme. On both the left and right extremes of the curve, there are dashed boxes starting at the edge of the red zone and moving outward. These boxes become the focus of metrics, analysis, and improvement activities regarding DDMRP buffers.

These outlying positions specifically represent breakdowns and threats to flow. The larger the number of outlying occurrences in either direction, the larger the threat to the flow of relevant materials represented by shortages (to the left) or excess inventory (to the right). Both are targets for improvement and elimination.

Eliminating these outlying positions will improve flow under the current buffer definitions and even allow for the migration of parts to buffer profiles with less variability. This will shrink the distance between the extreme positions, encouraging a tighter distribution. This tighter distribution against smaller buffer definitions yields better flow with less average inventory. This translates directly to better return on average capital. Once this concept is grasped, basic reports emerge to identify these outlying events and the personnel that impact them.

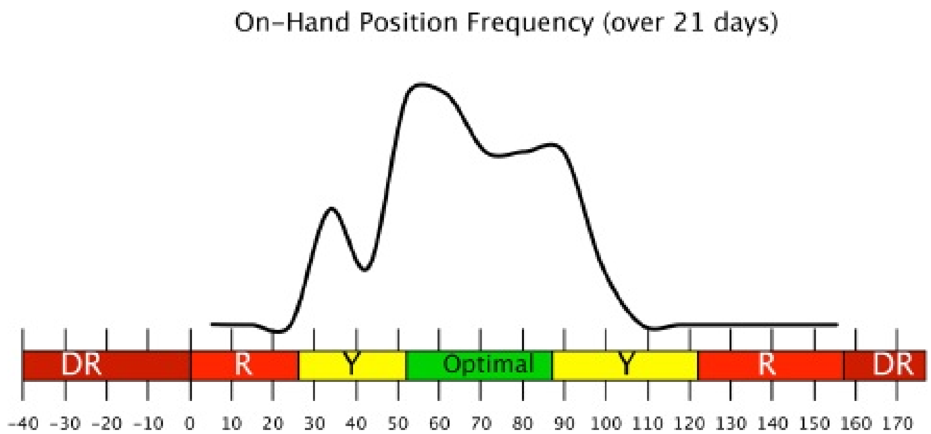

FIGURE 12-8 Simulated part's on-hand distribution curve

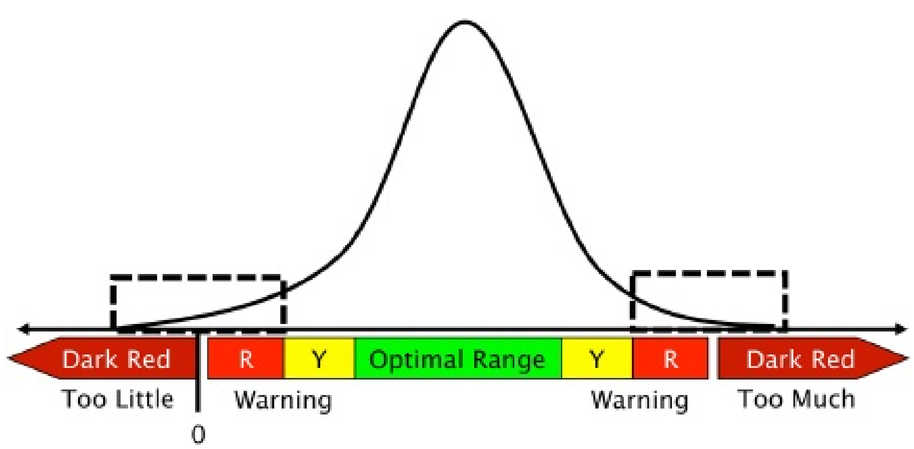

FIGURE 12-9 Outliers of the distribution

Outlying Event Reports

Outlying event reports are designed to promote visibility to the outlying events that break down material flow on both sides of the loss function distribution (the dashed-box areas of Figure 12-9).

Figure 12-10 is a sample of this type of report for purchased items that are disrupting flow on the "too little" side of the distribution curve. The report is for activity over the month of May (the date range May 1–31). The report shows the part number and the supplier of the part. The column "Parent Items" represents the number of parents of which this part is a component. The column "Number of Stockouts" lists the number of times over the date range that the part has been stocked out. The column "Total Stockout Days" shows the total number of days that the part has been stocked out over the date range, as each stockout occurrence could last for multiple days. Finally, the number of times the on-hand position has eroded past the on-hand alert threshold is represented in the column "Number of On-Hand Red." Any stockout would also be counted as an on-hand red penetration since a part would have to pass through the on-hand alert threshold before being stocked out.

Purchased Part Disruption Report			Date Range: May 1–31		
Part #	**Supplier**	**Parent Items**	**Number of Stockouts**	**Total Stockout Days**	**Number of On-Hand Red**
PPA	PNW Heat Treat	4	3	10	5
PPC	PNW Heat Treat	12	3	6	4
PPZ	Columbia Fabrication	2	1	22	1

FIGURE 12-10 Purchased items impacting buffer integrity

In this report PNW Heat Treat has two problematic parts: PPA and PPC. PPA had 3 stockout occurrences representing a total number of 10 stockout days. This endangers 4 separate parent items. The part's on-hand position has eroded severely on 5 occasions over the date range, 3 of which deteriorated to stockout. PPC has 12 parent items, 3 stockout occurrences, and a total of 6 stock days. Only one of the on-hand alert penetrations did not result in a stockout. Columbia Fabrication has one part (PPZ) that was stocked out for 22 days over the date range.

This type of report can be generated for any type of decoupled part position (purchased, intermediate, finished, and distributed). Figure 12-11 depicts an outlier report for disruptions to intermediate part buffers. In this case the supplying resource is identified. ICT is supplied by the Drive Control Cell, has 8 parent items, and has stocked out on 5 separate occasions for a total of 21 stockout days. Additionally, it has penetrated the on-hand alert threshold a total of 9 times, 5 of which resulted in a stockout.

Figure 12-12 is another sample of an outlier report. In this case it is sorted by supplier. The “Number of Parts” column displays how many different parts supplied by this supplier are represented in the report. PNW Heat Treat has 11 items over the date range (May 1–31) that have either stocked out or passed the on-hand alert threshold (shown in the “Number of On-Hand Red” column). PNW has a total of 15 stockout occurrences against the 11 parts. That means that at least one of those parts has stocked out more than once over the date range. Those 15 occurrences have resulted in a total of 29 stockout days. These 11 parts also have generated 19 on-hand alert penetrations (15 of which progressed to stockout.)

Intermediate Part Disruption Report			Date Range: May 1–31		
Intermediate Part	**Supplying Resources**	**Parent Items**	**Number of Stockouts**	**Total Stockout Days**	**Number of On-Hand Red**
ICT	Drive Control Cell	8	5	21	9
ICD	Cylinder Cell	5	2	3	12
ICS	Falcon Lathe	2	0	0	5

FIGURE 12-11 Intermediate buffered items report

Supplier Disruption Report		Date Range: May 1–31		
Supplier	**Number of Parts**	**Number of Stockouts**	**Total Stockout Days**	**Number of On-Hand Red**
PNW Heat Treat	11	15	29	19
Chelan Electronics	5	0	0	10
Columbia Fabrication	1	1	22	1

FIGURE 12-12 Suppliers impacting buffer integrity

These types of outlier reports give visibility to the level of variability experienced by buffers and the sources of that variability. This visibility will be invaluable in driving improvement activities in a DDMRP system.

Turning our attention to the other side of the loss function spectrum, we will focus on samples of outlier reports that make excess inventory visible (the dashed box on the right side of the curve in Figure 12-9).

Figure 12-13 is an excess inventory report by a planner over a 31-day range. It is designed to point out planners (and the parts they manage) that are having difficulty keeping on-hand inventories down. Three planners are represented in this report. The total number of parts under each planner's control is displayed in the column labeled "Number of Parts." Nick has 152 parts. Carmine has 113 parts. Julia has 49 parts. The column labeled "Beginning Excessive On-Hand (Dark Red)" is populated with parts and their cumulative values in the excess on-hand zone. This zone is the zone labeled "Dark Red (Excessive Stock)" in Figure 12-7 and the zone represented on the far right ("Dark Red") in Figure 12-9. The number in parentheses is the total number of parts in the zone in the beginning of the period and their cumulative value just within that zone. Nick had 19 parts in the excessive zone. The value of just the amount in the excessive zone of those 19 parts is $121,633.

The column labeled "End Excessive On-Hand (Dark Red)" is the number of parts and their cumulative zone value at the end of the date range. Nick reduced the number of parts in the excessive zone from 19 to 17 and the total dollar value from $121,633 to $112,361. The reduction of $9,272 is displayed in the column labeled "Excess Inventory Change." The column labeled "Beginning High On-Hand (Red)" represents the number of parts and their value within the zone represented in Figure 12-7 labeled "Red (High)" and the zone labeled "R" on the right side of the distribution in Figure 12-9. Nick began the period with 26 items in the high on-hand zone and ended with 12. High on-hand inventory dropped by $97,939 within the period.

The numbers of parts and their values in the "Excessive" columns are not counted in the "High On-Hand" columns. Adding the numbers together would produce the total amount of on-hand inventory parts and values of concern for each planner. For Nick that would be 29 parts and $179,654 at the end of the period.

Excess Inventory Report by Planner				Date Range: May 1–31			
Planner	Number of Parts	Beginning Excessive On-Hand (Dark Red)	End Excessive On-Hand (Dark Red)	Excess Inventory Change	Beginning High On-Hand (Red)	End High On-Hand (Red), $	High On-Hand Change
Nick M.	152	(19) $121,633.00	(17) $112,361.00	($9,272.00)	(26) $165,223.00	(12) $67,284.00	($97,939.00)
Carmine M.	113	(10) $51,471.21	(12) $67,897.67	$16,426.46	(13) $71,233.00	(25) $172,913.00	$101,680.00
Julia W.	49	(1) $5,123.00	0	($5,123.00)	(7) $9774.00	(1) $674.00	($9,100.00)

FIGURE 12-13 Planners with inflated inventory positions

Excess Inventory by Part				Date Range: May 1–31			
Part #	Planner	Beginning Excessive On-Hand (Dark Red)	End Excessive On-Hand (Dark Red)	Excess Inventory Change	Beginning High On-Hand (Red)	End High On-Hand (Red), $	High On-Hand Change
FPT	Nick M.	$1,256.00	0	($1,256.00)	$2,135.00	$136.00	($1,999.00)
FPH	Carmine M	$5,471.21	$1,781.00	($3,690.21)	$1,124.00	$1,124.00	0
FPI	Carmine M	0	$1,231.00	$1,231.00	$356.00	$1,254.00	$898.00

FIGURE 12-14 Parts with inflated inventory positions

Figure 12-14 is an outlier report displaying parts with inflated inventory positions and their respective change over a defined period (May 1–31). This report includes the specific part number and the planner assigned to it. The part's beginning and ending excessive on-hand position and the change between the two is displayed. The same is true for each part's high on-hand position.

Part FPT is controlled by Nick. Over the date range this part's excessive on-hand value has been eliminated ($1,256). Its high on-hand inventory has been almost eliminated with an ending value of $136. Carmine has managed to reduce FPH's excessive inventory by a significant amount ($3,690.21). Since the excess inventory has yet to be eliminated at the end of the period, the high on-hand remains unchanged. FPI, however, has added inventory over the period. Its high on-hand range was filled (from $356 to $1,254), resulting in an overflow to the excessive inventory range ($1,231).

There are countless derivations and alternatives to organize and display outlier data. These are just a few samples that have been effective in DDMRP implementations to date. Quantity, working capital, and even capacity can be included in these reports in order to highlight visibility to outliers and how they relate critical constraints and concerns to the desired Demand Driven Operating Model.

Measuring Velocity

An important metric in any Demand Driven Operating Model is system velocity. If the decoupling points are strategic in nature, then measuring velocity at those points will provide visibility and insight into total system velocity. If velocity through decoupling points is low, then it will indicate that total system velocity may be breaking down. If velocity through decoupling points is increasing, then it indicates that total system velocity is increasing. But how to measure velocity at decoupling points?

A DDMRP system uses a measurement called “order frequency variance.” Order frequency variance is the difference between planned order frequency and actual order frequency. Planned order frequency is calculated by dividing the green zone of a buffer by the average daily usage. If a buffered item has a green zone of 200 and an average daily usage of 40, then the planned order frequency is five days. Actual order frequency is calculated by dividing the total number of supply orders generated against a position by the number of days in the considered period. If the same part was ordered 7 times over a 22-day period, the actual order frequency is 3.14, greater than the planned order frequency. This metric is more insightful with a longer range.

A “flow index” is another measure of decoupling point velocity. It relates the planned or actual order frequency across a group of buffered items. It is a relative view. Figure 12-15 is an example of a flow index across 100 buffered purchased items. This flow index was created by dividing the parts’ green zone by their respective average daily usage.

The flow index specifically ranges from an order frequency of 1 to 50. Parts with an average order frequency higher than 50 are lumped in the “More” column. The height of the bars corresponds to the number of parts that share an average order frequency within a one-day range. For example, there are four parts that have an average order frequency from 1.01 to 2. The parts to the far right are slow-moving parts. Their green zones are largely out of proportion with ADU. This is typically due to relatively large minimum order quantities.

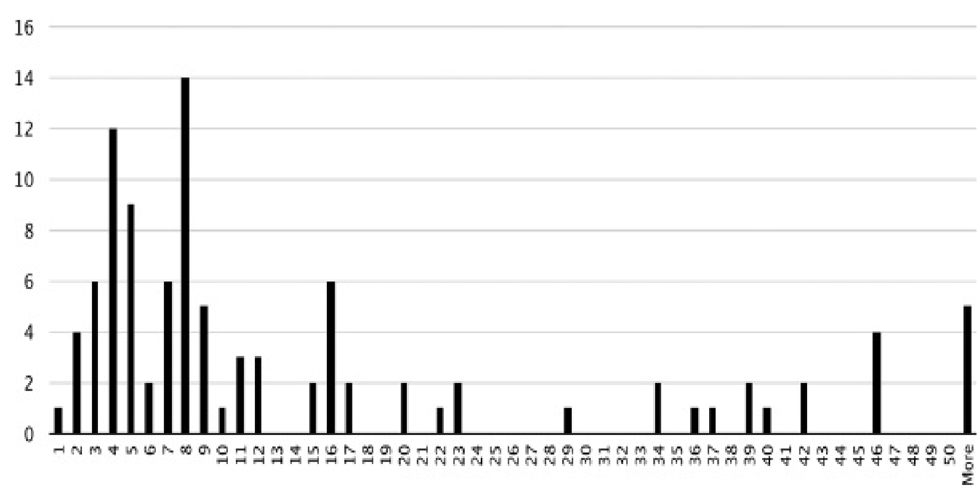

FIGURE 12-15 A flow index for purchased items

There are nine parts with an average order frequency higher than 45 days (four on day 46 and five in the "More" column). Figure 12-16 lists these parts and their respective average order frequency values. Part 182 has the most infrequent order frequency of just under 167 days. Each of these parts has its green zone set as its MOQ.

These types of parts typically pose challenges to planning personnel. They require disproportionately large amounts of inventory. That inventory typically represents shared materials and capacity. Their infrequent order patterns can even complicate scheduling sequences.

Part #	ADU	Green Zone	Order Frequency
148	22	1,000	45.45
181	11	500	45.45
192	22	1,000	45.45
198	11	500	45.45
159	9	500	55.56
162	7	500	71.43
154	11	1,000	90.91
138	3	400	133.33
182	3	500	166.67

FIGURE 12-16 Parts with low order frequency

There is a heightened sensitivity to things that move too slowly. But what about parts that move too fast? Can there be such a thing? Things that move too fast can create unnecessary amounts of transactional activity and additional setups that could result in capacity erosion, especially at a bottleneck resource. Figure 12-17 lists the parts with an order frequency higher than 3.

When severely high-frequency parts are within a group that also has severely low-frequency parts, a "trade" can be made that seeks to smooth flow across the group. Smoothing flow across the group will increase flexibility, thus minimizing total working capital commitments, protecting service levels, and minimizing expedite expenses. This trade prevents capacity erosion from additional setups created by lowering the MOQs of infrequently ordered parts. By raising green zone limits of high-frequency parts, it creates setup space for lowering infrequent part MOQs. Figure 12-18 shows the changes to both high-and low-frequency parts. High-frequency-part green zones have been set to a three-day order cycle (ADU × 3). Low-frequency-part MOQs have been cut in half.

Figure 12-19 is the updated flow index with the changed values for high- and low-frequency parts. The distribution of order frequencies has become tighter.

Driving Improvement in DDMRP

By implementing these simple measures to signal integrity, decoupling point integrity, and operating velocity, DDMRP can transform an environment in a relatively quick time frame and provide a path to continued improvement over time. Figure 12-20 shows a conceptual view of this progression through three stages. Stage 1 is the part in its initial bimodal distribution with larger outlying behavior against the buffer criteria. Stage 2 represents initial DDMRP results stabilizing the part against its initial buffer criteria with fewer outlying events and a single uniform distribution. Stage 3 shows a tighter uniform distribution against a compressed buffer definition with even smaller amounts of outlying occurrences. The arrow in Stage 3 represents the compression of the buffer values.

Part #	ADU	Green Zone	Order Frequency
141	765	400	0.52
187	422	500	1.18
204	65	100	1.54
171	320	500	1.56
220	127	250	1.97
124	321	700	2.18
143	422	1,000	2.37
146	798	2,000	2.51
185	765	2,000	2.61
201	765	2,000	2.61
212	68	200	2.94

FIGURE 12-17 Parts with high order frequency

Changed High-Frequency Parts				Changed Low-Frequency Parts			
Part #	ADU	Green Zone	Order Frequency	Part #	ADU	Green Zone	Order Frequency
141	765	2,295	3.00	148	22	500	22.73
187	422	1,266	3.00	181	11	250	22.73
204	65	195	3.00	192	22	500	22.73
171	320	960	3.00	198	11	250	22.73
220	127	381	3.00	159	9	250	27.78
124	321	963	3.00	162	7	250	35.71
143	422	1,266	3.00	154	11	500	45.45
146	798	2,394	3.00	138	3	200	66.67
185	765	2,295	3.00	182	3	250	83.33
201	765	2,295	3.00				
212	68	204	3.00				

FIGURE 12-18 Changes to high- and low-frequency-part average order frequencies

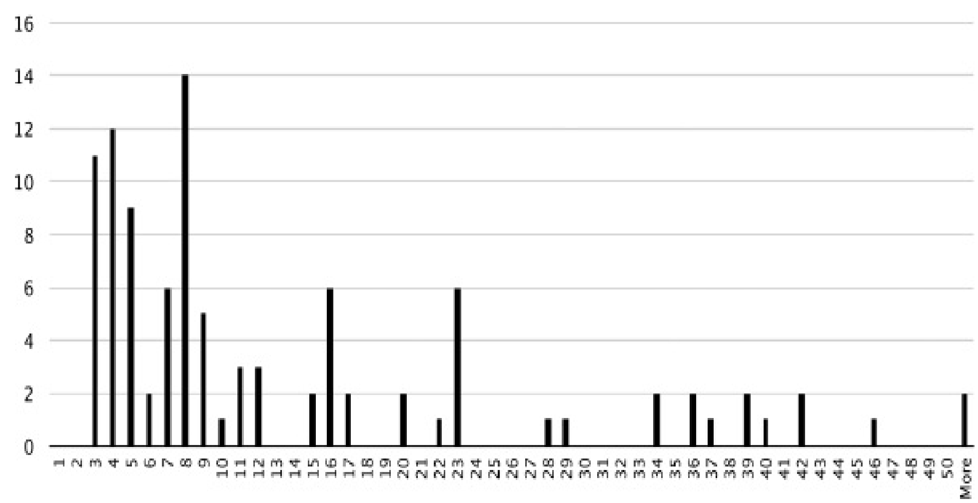

FIGURE 12-19 Updated flow index of purchased parts

Any sustained improvement approach in a DDMRP system has a primary directive: constantly strive to reduce working capital commitments with minimal expedite expenses and without erosion of service levels.. With this definition in mind, compression should not occur unless the majority of outlying occurrences are identified and effectively eliminated.

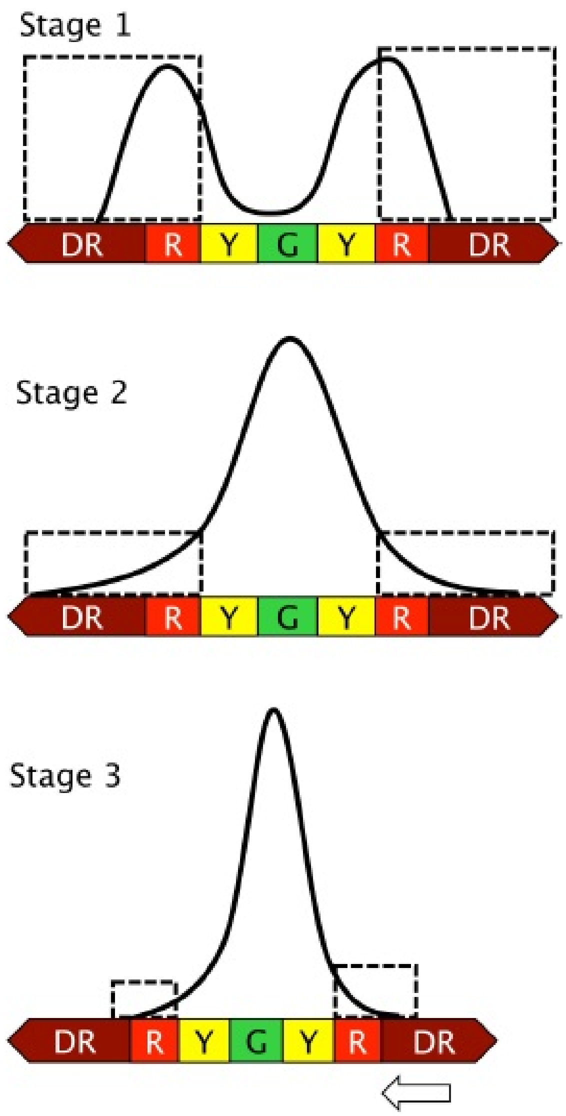

FIGURE 12-20 A company's progression through DDMRP

As described in Chapter 7, there are three primary factors that can combine to create the inventory commitment represented in the levels of a buffer: minimum order quantity, lead time, and the variability factor. Parts with large variability factors are an immediate target for improvement by the elimination of outlying events. The elimination of the outlying events means that the standard deviation of the position is reduced. This results in a much smaller spread in variability. A smaller variability factor will result in a smaller average on-hand commitment. The variability factor directly impacts the planning red zone, and the red zone comprises a significant portion of the average on-hand equation (red plus half the green zone).

Parts with minimum order quantities that qualify as the green zone can be explored for reduction. This will directly impact the planning green zone of the buffer. The larger the MOQ is in relation to average daily usage, the bigger the target for MOQ reduction. A flow index for parts with MOQ-determined green zones is used to identify the parts with the most potential improvement for MOQ reduction.

Parts in long lead time categories can be systematically explored for lead time reduction. Lead time tends to have the single biggest impact on buffer sizing, as it will affect all three zones, if the green zone is a factor of lead time. If the parts are purchased, then alternative methods of supply should be explored with existing or alternative suppliers. In some cases, it may be financially beneficial to pay a higher price per piece if the supplier is one week away as opposed to 90 days away. If the parts are manufactured items, then new decoupling points or faster methods of production can be explored.

Summary

This chapter was about creating visibility to the critical factors that maintain and improve the reliability, stability, and velocity of the flow of relevant information and materials—the very foundation of a DDMRP system. A series of measures, concepts, and reports were featured to ensure this end:

- **Signal integrity.** Properly operating a DDMRP system requires accurate and timely supply order generation. A proven method to make signal integrity visible was introduced. This method measures compliance to DDMRP supply order recommendations in terms of both timing and quantity.
- **Decoupling point integrity.** Properly operating a DDMRP system requires the ability to maintain and improve decoupling point integrity. This requires clear visibility to show how the buffer has performed against the control limits defined by the buffer sizing logic. When the buffer has performed outside those limits, the visibility to the outlying occurrences and their respective causes is vital to systematically working toward their elimination in order to better protect and promote flow.
- **Decoupling point velocity.** Properly operating a DDMRP system also requires visibility to the velocity occurring at the strategic buffers.

Focusing on the above measures will open the door for basic improvement avenues for buffer compression through the reduction of MOQs, lead time, or variability. Buffer compression results in improved ROI.

CHAPTER 13

The Demand Driven Organization

Consider the broader implications that a flow-based strategy has for an organization—the impact on the necessary strategic and tactical components to control, measure, adapt, and improve in the New Normal. The better that organizations are at understanding and implementing these components, then the more successful and sustainable a DDMRP implementation will be and the more sustainable the resulting ROI.

The Demand Driven Adaptive System

This book describes the components and details of Demand Driven MRP. In Chapter 5 Demand Driven MRP was described as a component of the Demand Driven Operating Model. In Chapter 11 the "master settings" of DDMRP were described as a component of a larger process known as Demand Driven Sales and Operations Planning (DDS&OP). Thus DDMRP is a component of a larger framework for the complexity and volatility that organizations must successfully navigate. This larger and emerging framework is a Demand Driven Adaptive System (DDAS). A Demand Driven Adaptive System is defined as:

> A management and operational system designed for complex and volatile manufacturers and supply chains. A Demand Driven Adaptive System uses a constant system of feedback that connects the business strategy to the settings and performance of a Demand Driven Operating Model through a Demand Driven Sales and Operations Planning Process (DDS&OP). A Demand Driven Adaptive System focuses on the protection and promotion of the flow of relevant information and materials in both the strategic (annual, quarterly, and monthly) and tactical (hourly, daily, and weekly) relevant ranges of decision making in order to optimize return on equity performance as change occurs.

Figure 13-1 is a depiction of the DDAS framework. This framework effectively marries two important time horizons: the tactical and strategic. The vertical dashed line bisecting the diagram depicts the border of these two horizons. These two time horizons are incredibly important to understand when it comes to relevant information. In the tactical horizon (hourly, daily, and weekly), the information and decisions that are relevant are distinctly different from the strategic horizon (annually, quarterly, and monthly). For example, Demand Driven MRP as part of the Demand Driven Operating Model exists in the tactical relevant range; it is a supply order generation and management method.

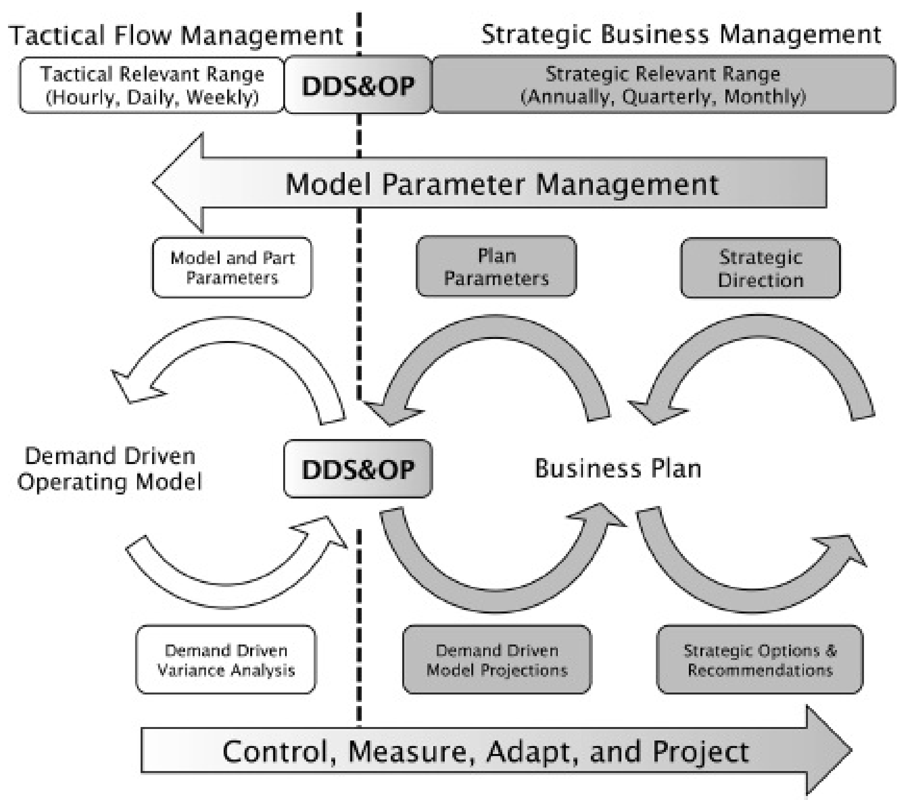

FIGURE 13-1 The Demand Driven Adaptive System schema

Working from right to left, strategic planning takes market intelligence and the desired business objectives to make strategic decisions. This strategic planning process is a closed-loop process and as well is linked to the business plan. The strategic process is affected by the outcome of the feedback loop from DDS&OP, where ROI opportunities and market opportunities are identified. The closed loop for strategic planning has the goal to maximize the system return according to the relevant model factors of volume, rate, and capability. When the loop process is complete, the intended direction is passed to the business plan in terms of financial expectation.

The business plan is also a closed-loop process taking into account the financial expectations from the strategic direction as well as the demand driven model projections. These projections include working capital, lead time, capacity, and other relevant projections. In addition, potential volume opportunities are also identified based on how the current model is performing. Once this process is complete, then the plan parameters are communicated to the DDS&OP process. All these activities to this point are in the strategic relevant range and typically are planned in monthly, quarterly, and annual time buckets.

DDS&OP is the lynchpin between the strategic relevant range and the tactical relevant range. DDS&OP develops the Demand Driven Operating Model based in part on the modeling criteria described in Chapters 6, 7, and 8. This includes decoupling point positions, buffer profiles, and part planning properties. This model is then run in conjunction with actual orders to generate and then manage all supply orders in the model. The DDMRP planning engine and execution alerts are the key to that generation and management, respectively.

Demand Driven Variance Analysis feeds back to DDS&OP to determine the stability, reliability, and velocity of the Demand Driven Operating Model over a past period. This is described in Chapter 12. The variance analysis highlights areas for improvement or changes that need to be made in the operating model. The feedback provides information on how well the current model is performing. The variance analysis also provides input to the DDS&OP process so that DDS&OP can project business financial results based on the current model capability and performance.

Each loop is both a closed loop and a tightly integrated chain to the next loop to provide an aligned adaptive system. Effective deployment allows the company to achieve and sustain improved return on investment. The Demand Driven Adaptive System is how companies will successfully compete in the New Normal. Just as DDMRP provides the engine in the Demand Driven Operating Model, DDS&OP becomes the critical linkage in a Demand Driven Adaptive System.

Demand Driven Sales and Operations Planning

DDS&OP includes managing the model and establishing the parameters for capability. Actual demand runs the system. As opposed to traditional Sales and

Operations Planning (S&OP), where the outcome is a master production schedule—a statement of what can and will be built— DDS&OP defines a capability so the company can profitably build what can and will be sold.

Demand Driven Sales & Operations Planning is required as the linkage between strategy and tactics to do the necessary integrated reconciliation and at the same time take advantage of all the demand driven capabilities that have emerged and can be exploited for new markets. The one thing that hasn't changed is the importance of the people involved to make this process a success. S&OP is about where decisions are made and where decisions should be made. The secret for success is the integrated reconciliation process.

Demand Driven Sales and Operations Planning is defined as:

> A bidirectional integration point in a Demand Driven Adaptive System between the strategic (annual, quarterly, and monthly) and tactical (hourly, daily, and weekly) relevant ranges of decision making. DDS&OP sets key parameters of a Demand Driven Operating Model based on business strategy, market intelligence, and key business objectives aligned with strategic information and requirements. DDS&OP also projects the model performance based on the strategic information and requirements and various model settings. Additionally, DDS&OP uses variance analysis based on past model performance (reliability, stability, and velocity) to adapt the key parameters of a Demand Driven Operating Model and recommend strategic alterations to the model and project their respective impact on the business.

Executed properly, the Demand Driven Sales and Operations Planning process directly links the strategic plans for the business with the review of relevant performance measures for operational execution and continuous improvement. With simulation, DDS&OP develops into a bidirectional integrated business management process. Demand Driven Sales and Operations Planning is the integrated business planning process that provides management with the ability to strategically direct its businesses to achieve and sustain competitive advantage on a continuous basis. It does this by integrating customer-focused business plans for new and existing products with the capability of flexible operations and by recognizing and exploiting the new capabilities that are now possible from the Demand Driven Operating Model.

DDS&OP becomes the bidirectional integration point between strategy and tactics in the Demand Driven Adaptive System schema. Figure 13-2 is the vertical representation of the DDS&OP schema. Plan parameters are communicated to DDS&OP from the business plan that aligns with the financial expectations. The financial expectations come from the strategic planning process that combines market intelligence with business objectives to make the strategic decisions. The plan parameters are then split into three main planning inputs—the demand plan, capability plan, and performance targets. The demand plan is split between the current portfolio and expected new activities.

These three planning inputs collectively create an expectation of where the company is going in the future. They are used to define the model and part parameters that will define the Demand Driven Operating Model as the first three steps of DDMRP. This is further described in the five DDS&OP steps later in this chapter.

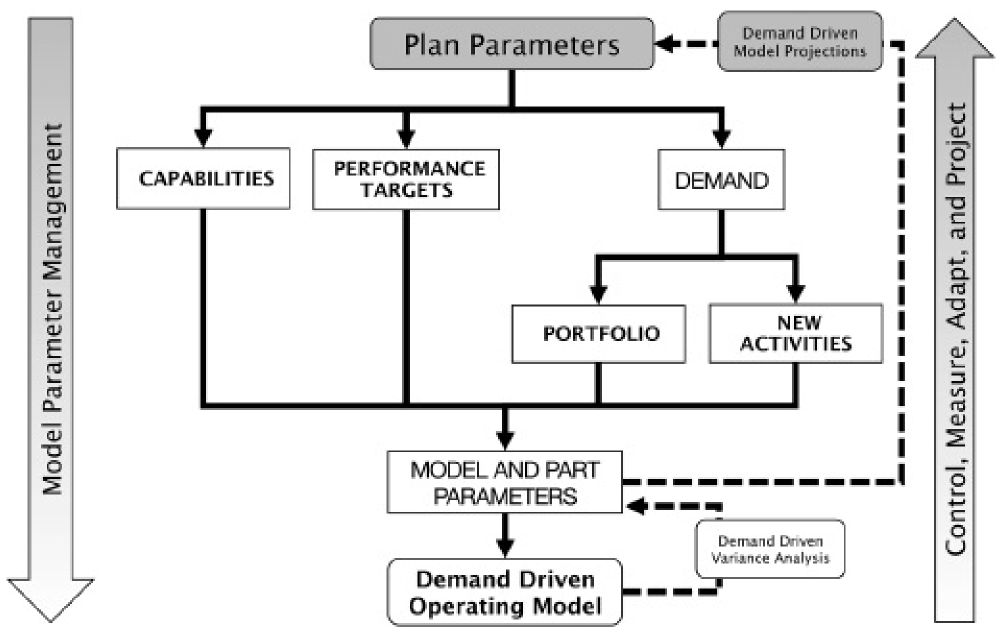

FIGURE 13-2 Demand Driven Sales and Operations Planning schema

The strategy and business plan are typically performed by product lines or marketing product families with a horizon sufficient to secure or change the necessary resources. In DDS&OP this demand plan is translated to the strategic buffer positions as defined in the Demand Driven Operating Model to calculate the requirement for the relevant factors to be considered. For example, this could be space, capacity, working capital, or whatever other limiting relevant factor that exists for the company. This calculation process is further described later in this chapter. These requirements are then considered for feasibility in the DDS&OP process.

The model and part parameters are essentially the capability control settings of the Demand Driven Operating Model. These parameters were referred to in Chapter 11 as the master settings of the model. Note that there is a feedback loop from the model and part parameters to the plan parameters called the Demand Driven Model Projection. This feedback loop provides a forward-looking view of how the Demand Driven Operating Model should perform given any changes in the main planning inputs (demand plan, capability plan, and performance targets). This feedback loop also includes visibility to leadership of current model configuration and performance.

The current model performance is monitored through the other closed loop called Demand Driven Variance Analysis. This feedback loop provides relevant information on how the model has performed historically against targets. This visibility then allows for adjustments to those parameters. Additionally, this information is leveraged through the metrics and analytics described in Chapter 12 to focus on outliers. When the outliers are addressed, then the overall flow dispersion of the process is reduced, flow improves, and the capability of the operational execution is improved. In addition, as the operational execution is improved, new market opportunities can be considered in the business planning and strategic planning processes. Improved operational execution is compared with the competition, and markets can be targeted that value that differentiator. This can include improved customer service delivery, shorter delivery times, flexible configurations, etc. These operational improvements open market opportunities to be leveraged by the business to achieve the strategic plan or provide additional revenue opportunities.

This redefinition from new product to new activities opens the scope outside a business-as-usual situation and increases appeal to a much broader audience in the business. This changing context has a dramatic effect on the view and understanding of the demand and supply process steps. The inherent flexibility of the demand driven MRP (DDMRP) methodology to sense and adapt to real customer demands provides the capability of the overall strategic plan to be roughly right rather than precisely wrong through the traditionally used inflexible master production schedule.

DDS&OP is the lynchpin between the strategic relevant range and the tactical relevant range. Coming from the strategic relevant range, DD&SOP provides model parameter management. Coming from the tactical relevant range, DDS&OP provides the control, measures, adaptation, and projections necessary to run the business.

The overall process begins with the development of a strategic model that aligns with the business plan. This defines the high-level direction for the company—markets, growth, products, etc. This strategic planning is set by the senior-level executives of the company. These strategic plans are then further refined into specific business planning, including the financial and nonfinancial goals and objectives by business unit. The Demand Driven Sales and Operations Planning process is where these high-level plans are connected to the reality of operational capability. Note in Figure 13-1 that there are two very important linkages from the business plan parameters to the model and part parameters. Operational capability and performance targets from the business plan when combined with the demand plan then determine the intended operating model for the company.

DDS&OP does not provide a master production schedule to tactical flow management. In fact, when using DDMRP as the engine to the Demand Driven Operating Model, a master production schedule is not necessary. DDS&OP provides targets in the form of average daily usage (ADU) around which the strategic decoupling buffers are built. Supply order generation is based on the net flow equation, which then replenishes these strategic buffers based on actual consumption. The demand driven model is about capability rather than a definitive monthly or weekly rigid schedule.

DDS&OP addresses the volatile, uncertain, complex, and ambiguous business world with an adaptive capability. The top part of Figure 13-2 is the vision and the development of a common understanding of the goals and objectives of the business. The transition point through DDS&OP is the collaboration between the goals and objectives with the reality of operational capability. This approach distinguishes DDS&OP from conventional S&OP approaches.

Conventional Sales and Operations Planning

MRP was codified in the early 1960s and enabled the concept of time-phased planning and dependent demand. MPS and capacity planning came next and provided the ability to have doable schedules and stability. The result was closed-loop MRP. This level of integrated manufacturing planning led to the introduction of S&OP in the 1980s in order to get management involved to steer the boat.

The Sales and Operations Planning process, according to Richard Ling and Andy Coldrick, is the process that brings together all the plans for the business (customers, sales, marketing, product development, manufacturing, sourcing, and financial) into one integrated set of plans. This process is iterated monthly and reviewed by the senior management team at an aggregate level. The process must reconcile all supply, demand, and new-product requirements at both the detail and aggregate levels and tie to the desired business plan. S&OP provides the definitive statement of how the company plans to compete in its market for the near to intermediate term covering a horizon sufficient to plan for necessary resources in addition to supporting the annual business planning process. Note the use of the word “process.” S&OP does not mean the meetings where decisions are made; S&OP is an ongoing adaptive process to effectively manage the business.

Companies have traditionally focused more on supply and financial considerations in their planning processes. The whole school of thought around supply chain management—and even the name “supply chain management”—reinforces the paradigm that the planning process begins with supply and then must be reconciled to demand. This is a left-to-right process from supply to demand. However, demand and new-product introduction represent the source of most change for a company today. As Charles Darwin said, “It is not the strongest or the most intelligent who will survive but those who can best manage change.” The S&OP process must focus from right to left—from demand—and then enable a supply capability from manufacturing through the supply base that can sense changes in customer demand and adapt planning and production in real time. This was introduced by Ling and Coldrick in the 1990s as “breakthrough S&OP.”

Early attempts at integration often focused on the commercialization and introduction stage of the product development funnel. The aims were to ensure preparation for launch and phase-in/phase-out and to understand possible cannibalization effects, motivated by helping production not to be caught with too little or too much inventory when introducing a new product. Progressive organizations, often those driving very aggressive innovation agendas, realized that connecting only the back end of the process starting at the introduction stage missed significant opportunities to manage the complete innovation funnel in an integrated way.

The scope also was broadened in another direction by those who saw the need to manage new activities beyond the narrower definition of product. Although the list is different in every company, a common theme in opening up this step beyond just product is identification of the activities that:

- Have a significant impact on demand and supply (volume and value) and any other support resources
- Need to be managed across the entire business with decisions driven through a structured review process
- Require visibility and management across a portfolio of activities, leading to better prioritization, resource allocation, and decisions linked to the overall desired business strategy

Planning activity has significantly changed over the last 10 years. As with traditional left-to-right S&OP in the 1980s, the notion of integrating new-product planning with supply and demand planning of the existing portfolio in right-to-left S&OP was something of a breakthrough at the time, despite being common sense. With the increasing focus on innovation and the use of stage and gate decision processes as well as innovation funnel management, the opportunity exists to integrate these emerging approaches and develop them symbiotically with S&OP. The next evolution of S&OP is thus by necessity called Demand Driven Sales and Operations Planning. DDS&OP is accomplished through five integrated steps.

The Five Steps of Demand Driven Sales and Operations Planning

Figure 13-2 shows how the strategic relevant range is translated to the tactical relevant range and vice versa. That bidirectional linkage is DDS&OP and is composed of five steps in sequential order:

1. Strategic business management direction and review
2. Integrated reconciliation
3. Managing the portfolio and new activities
4. Managing demand
5. Managing supply

Underlying this five-step process is a committed management team with communication clarity.

1. Strategic Business Management Direction and Review

The only source of sustainable competitive advantage is to exploit the company's unique operational advantage that provides value to the customer and provides ROI to the company. In the traditional left-to-right S&OP approach, the management review is typically the last step. In DDS&OP, this senior business management discussion must be the first step in establishing the overall intended direction for the company as well as defining the key metrics that will be used to evaluate performance. The senior business management review establishes the level of ownership of the process and specifies how to handle matrix management across multiple divisions, which can span multiple countries and continents. The review is not necessarily a once-a-month meeting but rather is an iterative process for senior management to review performance and set future objectives.

DDS&OP is the linkage from the overall company strategy to the operational capability and vice versa. DDS&OP and DDMRP are symbiotic in nature. Each methodology benefits from the existence of the other. Through DDMRP, operations develops a unique competitive advantage with shorter lead times, lower inventory, and higher customer services through the focus on flow and the leveraging of strategic decoupling buffers. Senior management can then choose to exploit these new capabilities to expand the current market, enter new markets, introduce new products, or engage in any combination of these possible strategies. This kind of scenario planning sets the stage for the integrated reconciliation of the different views that are now possible with this agile capability.

It is important to realize that DDS&OP is not about a monthly meeting but rather a dynamic ongoing process. The senior management direction and review is an ongoing process where senior management has its finger on the pulse of the business with the capability to provide direction between the formal monthly meetings. If not, in a dynamic environment a necessary change could slip by for a full month if this opportunity or change is noted immediately after the formal review meeting.

2. Integrated Reconciliation

Integrated reconciliation highlights the importance of financial involvement and leadership early in the DDS&OP process. This changes the agenda from a discussion about volume by product lines and simply reconciling demand and supply to a true business discussion about capabilities and opportunities for contribution margin. This step is about reconciling different views and scenarios. There is inherent value in discussing these different views, assumptions, and reasons for those scenarios. This increases the understanding of what the numbers really mean, which focuses attention on the assumptions underpinning the numbers, along with opportunities, vulnerabilities, and possible capabilities. The conversation is principally about what has changed both inside and outside the organization since the last review and why. Without assumptions, the conversation is simple: What numbers changed?

When the focus changes from just numbers to assumptions and different scenarios, the need for marrying medium- and long-term forecasting with foresight and market intuition becomes even more apparent.

Some of the questions involved in integrated reconciliation include:

- What is the impact of integrating new activities, demand and supply, on the business (not just the supply network)?
- What are the emerging issues and gaps?
- What are the opportunities and risks?
- What scenarios are relevant to make better decisions in the future?
- What decisions should and could we make, and which ones should be escalated to the senior management review?
- What operational capability can be exploited to provide value to the customer?

Volume and value information and assumption changes are required to answer these questions. Understanding these questions leads to the imperative that finance is an integral part of all five steps, whereas in many traditional S&OP processes, finance is added only at the pre-S&OP meeting and the S&OP meeting rather than being an integral part of the entire process.

The reconciliation step is not a meeting as such, but rather an iterative interactive process run by a senior-level cross-functional team in the business. This team highlights key issues and decisions required for the senior management team’s attention and discussion. In fact, the reconciliation team determines the agenda for senior business management review. Participants in integrated reconciliation are the future executives for the business. This process is truly a key training ground for the next generation of presidents and vice presidents because this team has a holistic understanding of the relevant business drivers. This is fundamentally different from the pre-S&OP meeting in traditional S&OP, where the main focus was on volume, its impact on resources, and attempts to balance capacity and load without considering the financial and strategic impact of those decisions on the business.

Understanding integrated reconciliation broadens the scope of new activities and demand and supply management. Integrated reconciliation as a process leads directly into the senior business management review, which focuses on understanding change including:

- What is our current performance, and how does it compare with what we expected?
- What decisions are still outstanding?
- What decisions have been made already in integrated reconciliation?
- Are we on track with the intended business plan?
- Are we still on track with our strategic intent?

3. Managing the Portfolio and New Activities

Recognizing how different the future will be from the past and present is important in understanding the business issues that connect to the DDS&OP process. DDS&OP is all about managing change and its consequences to the company, its resources, and its people. Figure 13-3 depicts five different portfolio models with their different emphases on DDS&OP. As models 1 through 5 show, a company could have anywhere from a future demand plan devoid of new activity (model 1) to one with a high degree of new-product introduction (model 5). A company operating under model 5 has the new-activity development impetus coming from products that are new to the world. This is by definition a company embracing a strategy of innovation.

The DDS&OP process in portfolio model 1 could be more traditional since there is no new activity in the next few years. Demand and supply balancing would be the emphasis in model parameter management. Because medium- and long-range forecasting for standard products in markets that are not growing is relatively straightforward, the ADU would most likely be calculated from a historical basis. A business embracing this model typically would be following a strategy of cost leadership. An example of this type of business would be commodity chemicals, where the price is set by the market.

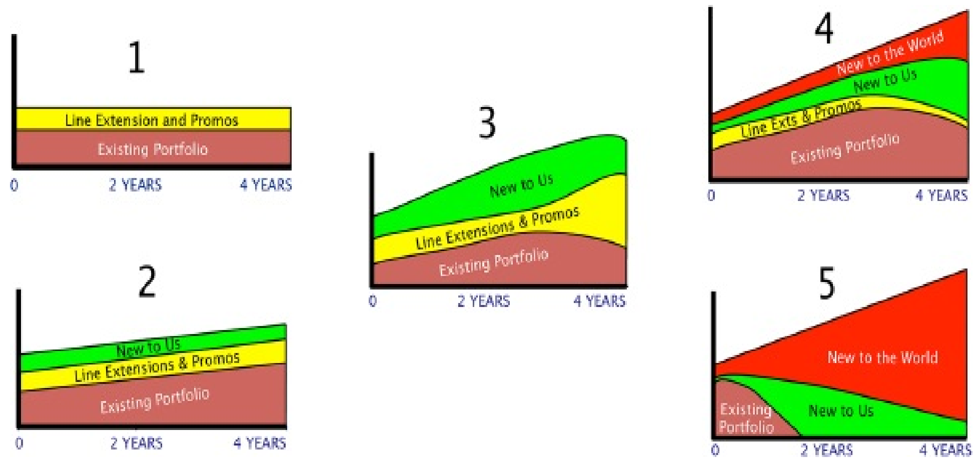

FIGURE 13-3 Different portfolio models

In model 2, there is more new activity, but it is relatively straightforward, and the business appears to have linear growth. New activity would play a part, but it would be a minor role. An example of this would be an industrial chemicals organization whose main business is commodities but provides some specialty chemicals. This company may be acquiring small businesses to augment the new-to-us category and achieving that growth by acquisition. The strategy here is primarily cost leadership, but the response in specialty chemicals could be differentiated service and lead time because of the higher margins on these products.

The most challenging business model for traditional S&OP is portfolio model 5, where today's portfolio will not exist in four years' time. These are businesses with a high degree of technological change and rapid development of new products. There is a high degree of volatility and risk for the business embracing this model. DDS&OP brings attention to those risks so they can be assessed and plans put in place to manage them.

Portfolio management, including new products, is a critical step in the DDS&OP process. This directly links the business plan to the strategic plan. The traditional S&OP model of demand and supply balancing would appear to be of little relevance to executives in these highly volatile environments. The ability to successfully manage uncertainty and a range of numbers in the integrated reconciliation step and the importance of simulation and its impact on profitability have enormous consequences for the business. This is why DDMRP is the logical methodology to effectively manage that variability. DDMRP provides the ability to withstand the uncertainty while leveraging resources, materials, and operational capability to build what can really be sold rather than wasting time and materials on items that are not really required.

Measurements such as time to market and time to profit are immensely important. Manufacturers of electronics, mobile phones, software, and computers are in this portfolio model. The strategy normally followed in these companies is product differentiation coupled with service differentiation. DDMRP provides the company the ability to sense what the actual demand really is and then adapt the supply network to react to that real demand.

Many food and drink companies and fast-moving consumer goods and pharmaceutical companies are examples of portfolio models 3 and 4. Typically, they would follow product or service differentiation or customer relationships. If your business has a portfolio similar to models 3, 4, and 5, spending time only implementing the demand and supply process, such is the case with the S&OP traditional model, provides little insight and value to the management of the company.

Understanding the business strategy is essential to understanding the emphases on the way DDS&OP will work. DDS&OP product portfolio models go hand in hand with understanding the company's strategic models. Strategies are about choices and trade-offs, and each business needs to understand and articulate the principal strategy it is following. It is not unusual to find that an organization might have different business units following different strategies in different areas around the world. DDS&OP must integrate these different strategies to provide a complete picture for the senior business management review. These are typically represented in product lines or families. These product lines or families are then converted to projected ADU in the next section.

4. Managing Demand

Demand management, including the accountability for forecasting, has developed significantly over the past few decades. In the early years of S&OP, a lot of effort went into agreeing to a volume forecast emphasizing a single set of numbers. Demand forecasting was very often part of the supply chain or sales organization, and forecast or master schedule accuracy was seen as the principal measure rather than customer service. Some organizations even went so far as saying, “You did not forecast this; therefore, we cannot make it!” obviously alienating sales and marketing. The thinking that sales and marketing form one homogeneous organization with a single view of the numbers misses the fact that these two functions have different drivers and objectives.

By the mid-1990s, people realized the importance of including sales and marketing in the forecasting process. At the same time, “Customer, Customer, Customer!” was fashionable, and sales became the focal point for forecasting and the one-size-fits-all solution, ignoring the importance of marketing input. In many businesses following a “customer relationship” strategy, sales leadership is appropriate, but in organizations with product and service differentiation strategies, marketing is the principal driver of medium- to long-term demand prediction.

Giving sales single accountability for the forecast led to some organizations spending too much time analyzing detailed history in a futile attempt to get the forecast accurate instead of being with the customers gaining knowledge about future trends. Against this background of trying to get the forecast accurate, there was a growing realization that there is inherent uncertainty that differs by markets, channels, and sectors, as well as with different products and customers. After years of complaining about forecast accuracy and trying to crank the handle faster on the same old detailed forecasting machine, companies began to wake up to forecasting for what it is—predicting the future based on the past! However, given the volatility today, the past bears little resemblance to the future. By no means does this remove the responsibility for forecasting, but it does lead to new and innovative ways of making a more informed prediction. Apart from agreeing with a forecast number, an important piece of knowledge is to understand the likely range (high and low) and forecast confidence factor, together with the supporting assumptions. Providing numbers without documented supporting assumptions is unhelpful. In some companies, the rule is that a forecast number cannot be changed unless an assumption is also changed. Forecasting is very important in the strategic relevant range of the Demand Driven Adaptive System schema.

Today, we understand that a robust demand plan over the necessary time horizon is possible only by reconciling cross-functional views; volume and value must be integrated. This is accomplished in the strategic relevant range. Finance, logistics, and supply chain managers are committed to this output. In general, sales input by major customers (with input from account managers) and channels is important in the short term, typically the first four to six months. Marketing provides information beyond four months based on market share, business goals, and brand or product health and marketing investment. Strategic marketing and research and development in many cases have a role beyond 12 months, particularly in new activities. There must be reconciliation between foresight (i.e., strategic marketing) and forecasting (i.e., nearterm marketing and sales). These are guidelines only to illustrate the collaborative approach and will vary depending on the business. The responsibility of finance, supply chain, and logistics is to ensure that the volume and financial forecast are reconciled and aligned. The demand plan is at an aggregate level in the strategic relevant range, and the aggregate marketing families are chosen, understood, and used by all functions. Simulations at the aggregate level are more helpful than trying to do "what-ifs?" at the stock-keeping unit.

Traditional S&OP tended to use manufacturing families. In demand driven environments and product differentiation businesses today, typically the aggregate product family is the brand or the brand and technology. Why would one choose an aggregate group with little relevance to marketing and sales? The difficulty is then how to translate the forecast from marketing into an input that can be used by operations. This translation is done through the DDS&OP process by utilizing the DDMRP methodology with the strategic decoupling positions. The product family projections are supported by the operating model with those strategic decoupling positions. Those projections require a time-phased view to account for working capital, space, and capacity over time. When an imbalance occurs, then the ADU can be adjusted with the planned adjustment factor to level the load over the planning horizon. This is described in a later section.

5. Managing Supply

Supply management also has broadened in its scope. Traditionally, it applied to just manufacturing, but now it is extended from manufacturing to a wider view of sourcing that encompasses other resources, including external ones. In multinational organizations, it has been extended to supply chain optimization and risk management, making the best sourcing decisions from the scenario planning process. This has challenged the planning capability of many organizations, requiring planners who are capable of moving beyond their traditional role—that of management and execution in detail at single supply points—to the role of optimizing supply networks by testing different scenarios and making informed recommendations and hence the best decisions.

DDMRP allows the evaluation of different supply strategies through the calculation and comparison of the necessary buffers to each scenario, similar to the evaluation of different demand strategies. These different supply strategies are not converted precisely to a master production schedule to drive supply orders as in traditional S&OP. Leveraging the existence of the strategic decoupling buffers, the company now has the capability to sense changing actual customer demand and then adapt supply through the entire supply network in real time to generate the maximum ROI.

Demand Driven Sales and Operations Planning Projections

A critical aspect of Demand Driven Sales and Operations Planning is the ability to project the Demand Driven Operating Model performance given specific capabilities against a projected demand. These performance projections typically include working capital, space, and capacity implications. The DDMRP methods described in this book make these projections relatively easy to derive given the proper inputs.

Chapter 7 described the critical inputs to the buffer equation. Figure 13-4 is a repeat of Figure 7-18 showing the key elements of the buffer equation. Any one of these elements can be changed and the equation rerun in order to judge the impact of that change. For example, the implication of moving a part to a different profile can easily be calculated by changing the lead time or variability factor. The implication of changing any of the part traits can also easily be calculated.

One of these part traits is average daily usage. Chapter 7 described the various considerations for determining the average daily usage in order to calculate current decoupling point buffers. This is one of the master settings of the Demand Driven Operating Model.

DDS&OP can also project more remote periods of time using a projected ADU input. This input will provide a point-in-time picture of what the buffers for a particular part will look like. This picture can then be used to judge the working capital, space, and capacity implications of demand at that level.

The connection from the operational planning using the ADU as described above, through the tactical planning to the DDS&OP process, is accomplished through the use of medium-and long-term forecasts by product families. These medium- and long-term forecasts are then translated to the strategic decoupling buffer ADUs to calculate the strategic buffers that would be necessary to support that level of business. Once these buffers are calculated, then they can be converted to working capital investment, space, or critical resource capacity. These buffers are simulated over the planning horizon to ensure that sufficient resources exist to execute the plan. If there is insufficient capacity in future time periods, then the planned adjustment factors can be used to level the load by building the fast-moving product ahead of the demand surge.

As opposed to the traditional planning by product family that is typical in S&OP, DDS&OP translates that forecast by product family into the projected ADU for strategic buffer items according to the Demand Driven Operating Model, and hence this supports the financial expectation from the business plan to calculate critical resources. Once these calculations are completed, the information can be again displayed in marketing product families to support the DDS&OP reconciliation meeting.

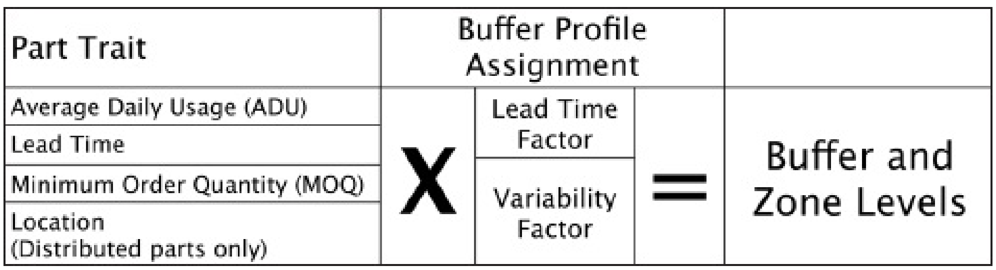

FIGURE 13-4 Buffer equation elements

As an example, consider a company that makes four products (items XYZ, ZYX, ABG, and GJK). Through the demand management process, a projection has been made for six months from now, as described in Figure 13-5. This projected rate of demand is displayed in the "Projected ADU" column. Additionally, Figure 13-5 has the necessary components of the buffer equation (lead time and buffer profile attributes).

XYZ has a current ADU of 100. Six months from now its projected ADU is 150. Its decoupled lead time is five days. The parentheses in the "Decoupled Lead Time" column represent the lead time category and lead time factor. XYZ is in the short lead time category using a 75 percent lead time factor (LTF). The desired order cycle for XYZ is three days. Its minimum order quantity (MOQ) is 500. Finally, XYZ is in the medium variability category using a 50 percent variability factor.

Given these inputs, it is relatively simple to produce the projected buffer levels using the projected ADU. Figure 13-6 shows current versus projected buffer zone values for all four items. Additionally, Figure 13-6 also shows the current versus projected targeted on-hand inventory position (red zone value + one-half green zone value). The row titled "Green" displays the green zone quantity of the buffer as well as the method of calculation.

Item #	Current ADU	Projected ADU	Decoupled Lead Time	Desired Order Cycle	MOQ	Variability
XYZ	100	150	5 (short—75%)	3	500	Medium (50%)
ZYX	50	75	7 (medium—50%)	3	250	Medium (50%)
ABG	25	10	5 (short—75%)	3	250	Low (25%)
GJK	20	200	5 (short—75%)	3	250	High (70%)

FIGURE 13-5 Example company for DDS&OP projections

	XYZ Now	XYZ Projected	ZYX Now	ZYX Projected	ABG Now	ABG Projected	GJK Now	GJK Projected
Green	500 (MOQ)	563 (LTF)	250 (MOQ)	263 (LTF)	250 (MOQ)	250 (MOQ)	250 (MOQ)	750 (LTF)
Yellow	500	750	350	525	125	50	100	1,000
Red	563	845	263	395	118	48	128	1,275
Average on-hand	813	1,127	388	527	243	173	253	1,650

FIGURE 13-6 Current and projected buffer calculations

Working Capital

As described in Chapters 7, 9, 11, and 12, DDRMP buffers are intended to always have stock to maintain their decoupling protection. The average quantity (number of units) of stock is calculated using a simple equation (red zone value + one-half green zone value). Given this equation, an additional equation allows us to convert the average quantity to an average amount of working capital. Figure 13-7 compares the current versus projected average working capital represented by the average on-hand quantity for each of the four items. In each case the average on-hand inventory levels (both current and projected) are multiplied by the working capital per unit. This working capital per unit represents the direct material cost per unit only. This is consistent with the Company ABC example used in Chapters 6, 7, and 9.

To support the predicted future rates of use, an additional \$108,160 in average working capital will be required. This is neither good nor bad. The feasibility is for the business leadership to decide given the circumstances of the business. It simply means that given the current assumptions (same buffer profiles and same part attributes), the buffers will need to contain more capital to support the increased business level.

Those current assumptions can all be challenged over the next six months in order to change the projection. For example, reducing part lead times, reducing direct material costs, or using a lower variability profile would all yield different projections. If there are real capital constraints, then these avenues can be explored to improve the feasibility.

Space

In a similar fashion the current and projected average targeted on-hand levels can be converted to space requirements such as pallet positions. Figure 13-8 compares the current versus projected pallet position requirements represented by the average on-hand quantity for each of the four items. In each case the average on-hand inventory levels (both current and projected) are multiplied by the units per pallet.

In this case an additional column has been inserted called "Projected Maximum Pallet Positions." This represents the number of pallet positions required when all buffers are at the top of their average on-hand range (red zone value + green zone value). This is done to give a sense of range in pallet positions that could be required even when the buffers are deemed to be operating within tolerance. This range might be important if a company has real warehouse or storage limitations. Under the projected rates of demand, an additional 84 pallet positions will be needed on average. As many as an additional 133 positions could be needed if all items are at the top of their average on-hand range.

Item #	Working Capital per Unit	Current Target Inventory Level	Current Average Working Capital	Projected Target Inventory Level	Projected Average Working Capital
XYZ	$100	813	$81,300	1,127	$112,700
ZYX	$90	388	$34,920	527	$47,430
ABG	$80	243	$19,440	173	$13,840
GJK	$50	253	$12,650	1,650	$82,500
		Total	$148,310	Total	$256,470

FIGURE 13-7 Working capital comparison (current versus projected)

Item #	Units per Pallet	Current Target Inventory Level	Current Average Pallet Positions	Projected Target Inventory Level	Projected Average Pallets Positions	Projected Maximum Pallet Positions
XYZ	10	813	81	1,127	113	141
ZYX	20	388	20	527	27	33
ABG	50	243	5	173	4	6
GJK	30	253	9	1,650	55	68
		Total	115	Total	199	248

FIGURE 13-8 Pallet position requirements comparison (current versus projected)

Again, is this feasible? That is for the business leadership to decide given the circumstances of the business. It simply means that given the current assumptions (same buffer profiles and same part attributes), the buffers will need to add more pallet positions to support the increased business level. This could be of strategic importance since additional storage space, if required, could be difficult or costly to obtain in a short period of time. For example, building additional space will require design, permitting, and construction time. Using a third-party warehouse might require a complex logistics plan calling for additional personnel and transportation. No matter what the specific circumstances, it is important that the business gains this visibility well in advance of actual projected need.

Capacity

A significant part of the DDS&OP process is to determine if sufficient capacity exists to support the proposed future. Using DDMRP as the operational planning and execution methodology allows a company to ensure that overall sufficient capacity exists without committing to a master production schedule that forces production to build what is on the schedule, rather than possessing the agility to build what the customer desires to purchase. This is an excellent example of where roughly right will outperform precisely wrong.

Assume that the resource being considered is the lathe department. Currently the lathes are not a constrained resource, but the management team has concerns about the growth in business relative to the overall load on this department. The company has five lathes that are available, each for 1,200 minutes per day (20 hours) of total production capacity; that means there are 6,000 minutes each day of total lathe capacity. For simplicity, all five lathes have the same process capability and rate—they are identical in nature.

Item #	Minutes per Unit on Lathe	Current ADU	Current Average Daily Load	Projected ADU	Projected Average Daily Load
XYZ	30	100	3,000 minutes	150	4,500 minutes
ZYX	25	50	1,250 minutes	75	1,875 minutes
ABG	20	25	500 minutes	10	200 minutes
GJK	10	20	200 minutes	200	2,000 minutes
		Total	4,950 minutes	Total	8,575 minutes

FIGURE 13-9 Lathe department capacity load (current versus projected)

For calculation of rough-cut capacity requirements, the current and projected ADUs are multiplied by the minutes per unit on the lathe. This gives us a rough-cut lathe capacity requirement per day to support the current and projected ADUs. Figure 13-9 shows that there is sufficient current capacity. However, given the current demand projections, it can be seen that the projected load six months from now far exceeds the current lathe capacity. This is detailed in Figure 13-10.

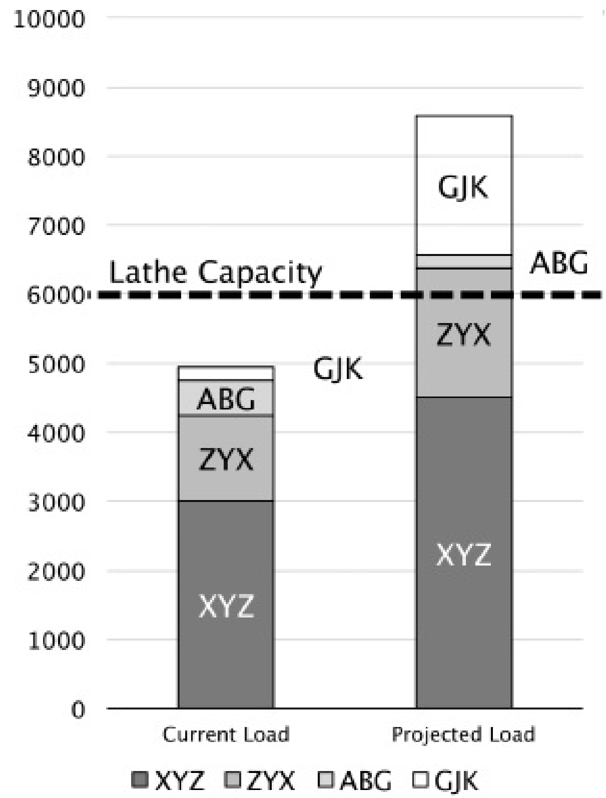

FIGURE 13-10 Composition of lathe load (current and projected)

Now choices must be made in the management meeting. Changing buffer profiles will have no impact since this analysis is based on demand projection, available capacity, and a specific part attribute (minutes per unit on the lathe). One or more of the following things must be considered for manipulation to increase capacity or decrease load.

Lathe Manipulation

Currently the five lathes work 20 hours per day. An additional 1,200 minutes per day could be added to the total lathe capacity if all lathes went to 24 hours per day. That would bring total capacity to 7,200 minutes, still below the required 8,575 minutes required. Additionally, it would require more operators and leave no time for preventive maintenance.

Another consideration would be to add more lathes. If all lathes were available for 1,200 minutes per day, it would take at least three additional lathes to meet the load. This is not a trivial investment. Additionally, it would take more lathe operators. In some parts of the world, the availability and lead time to hire a qualified operator is sometimes more difficult to come by than the acquisition of the machine.

Load Manipulation

One critical factor in calculating this load is the time it takes per unit for each item. The item requiring the heaviest load per unit is XYZ. It takes 30 minutes per unit of lathe capacity. Six months from now it is projected to require 75 percent of available daily lathe capacity. One way to manipulate the load requirement is to attempt to reengineer the product in a way that requires much less lathe time.

Another way to decrease lathe load would be to outsource production of one or more of these items. But which items are the right candidates? In order to answer this question, the company will need to understand the financial return generated by each item that goes across the lathe. A basic tenet of management accounting is that companies' profits maximize when the companies make and sell the products with the highest contribution margin per unit of the scarcest resource. The scarcest resource six months from now is projected to be the lathes. Thus we will need to calculate the contribution margin for each product in relation to its impact on lathing capacity. Figure 13-11 shows the relative cash contribution per lathe minute for each item. Part XYZ has the largest per minute load and the lowest rate of cash return on the lathe. What this means is that when the company is making XYZ, it is getting \$5.83 in cash contribution versus \$7.00 when making GJK.

In this contribution margin calculation, we are considering only the price minus truly variable costs (in this case direct material cost). This is because the only truly variable cost in this environment is the direct material; all other costs are assumed to be fixed within the operational relevant range. The variable cost represents a real cash outlay directly related to each unit of each particular item. We must understand the rate of cash generation at the lathe only—the scarcest resource. Total labor and overhead are irrelevant and will only distort the picture. Obviously, this is not the traditional margin as calculated in the ERP system item master, because standard cost considers nonrelevant costs in the determination of fully loaded costs. Readers that wish to know more about this concept are encouraged to read *Demand Driven Performance: Using Smart Metrics* by Debra Smith and Chad Smith.

Item #	Minutes per Unit on Lathe	Direct Material Cost	Price	Contribution Margin	Contribution Margin per Lathe Minute
XYZ	30	$100	$275	$175	$5.83/min
ZYX	25	$90	$250	$160	$6.40/min
ABG	20	$80	$200	$120	$6.00/min
GJK	10	$50	$120	$70	$7.00/min

FIGURE 13-11 Contribution per margin per minute of lathe time

What does this mean from an outsourcing perspective? If a company has a capacity constraint and it is going to outsource, then it should outsource the item that produces the least return on that capacity constraint. It should keep the items that produce the best return on that resource in-house. In this case, XYZ should be a candidate for outsourcing to bring the required lathe load down to 6,000 minutes per day. That would mean outsourcing at least 86 pieces per day on average.

Demand Manipulation

If the business is incapable of meeting all the demand, then demand could be manipulated down by raising the price of certain items in order to maximize the rate of return. Which items should be chosen for a price increase? The answer to this also lies in examining the cash contribution from the lathe for each item. XYZ is projected to sell 150 per day in six months. It is the lowest cash contributor in terms of the lathe.

By raising the price to $180 per unit on XYZ, the contribution margin per lathe minute becomes identical to ABG. If this does not erode demand enough, then XYZ and ABG should be considered for additional price increases. ABG is a low-volume item, and so its impact in terms of relieving total lathe capacity is limited.

Figure 13-12 shows the price required for XYZ and ABG to provide the same contribution margin as the next-lowest product, ZYX. The XYZ price would have to move to $292 per unit. ABG would have to be priced at $208 per unit. At this point XYZ, ZYX, and ABG could all be considered for additional price increases if the projected ADU erosion was insufficient for the available lathe capacity. One thing is for certain; the company would be making a tremendous amount more return for the same fixed-cost structure.

Projected Order Frequency

An additional factor to consider in managing supply is the current versus projected order frequency. This could be relevant when considering the impact on the number of setups and the impact on inbound logistics.

Item #	Minutes per Unit on Lathe	Direct Material Cost	Price	Contribution Margin	Contribution Margin per Lathe Minute
XYZ	30	$100	$292	$192	$6.40
ZYX	25	$90	$250	$160	$6.40
ABG	20	$80	$208	$128	$6.40
GJK	10	$50	$120	$70	$7.00

FIGURE 13-12 New contribution comparison

Item #	Current Green Zone	Current ADU	Current Order Frequency	Projected Green Zone	Projected ADU	Projected Order Frequency
XYZ	500	100	5	563	150	3.75
ZYX	250	50	5	263	75	3.51
ABG	250	25	10	250	10	25
GJK	250	20	12.5	750	200	3.75

FIGURE 13-13 Current versus projected order frequency

For simplicity we have not dealt with setups related to the lathe example. Yet looking at projected order frequency will give us a good indication of what the impact will be. Figure 13-13 shows the current versus projected order frequency for the four products from our example. As can be seen in the Figure 13-13, the frequency of production for GJK changes dramatically— from every 12.5 days to every 3.75 days. This increase in the frequency of production could dramatically impact capacity if the setup time for GJK is significant on a constrained resource.

The production of ABG will change from every 10 days to every 25 days on average. If the product has a shelf life, this is an issue that may need to be escalated to the management review. This product may be a candidate to discontinue since the contribution margin per unit of constrained capacity is relatively low and there would be a high risk of the product’s shelf life expiring.

In a distribution environment, order frequency will typically relate to the average number of inbound receipts. The more that items are ordered, the more shipments that tend to be received. This can put pressure on the receiving and inspection operations. Is there enough inspection space? Are there enough dock doors? Are there enough personnel?

Summary

The New Normal has radically altered what it takes to sustain and improve a company's competitive advantage. This alteration requires a new form of strategic and tactical management, one that allows a company to see, learn, and adapt its resources to the complexity and volatility of the New Normal. The legacy tactics inherent in the conventional planning and execution systems as characterized by MRP and MPS are simply inappropriate for the circumstances that a company faces today. That inappropriateness translates directly to poor returns on asset performance. These distort and confuse the picture and make strategic analysis and prediction extremely difficult. We are simply starving for relevant information in both the strategic and tactical relevant ranges.

Yet a new way has emerged: the Demand Driven Adaptive System. This approach effectively links through the DDS&OP process the strategic and tactical relevant ranges, providing unprecedented visibility and a mechanism to produce relevant information for adaptation and projection. With more relevant information come more relevant materials and a better return on asset performance.

The results of the demand driven approach speak for themselves. Typical results include:

- Service level above 95 percent
- Inventory reductions of 30 to 50 percent
- Expedite-related expenses down significantly or eliminated

Many case studies are available at www.demanddriveninstitute.com. See Appendix E for a DDS&OP implementation checklist.

Contribution of Dick Ling

We would like to recognize the critical contribution of Richard (Dick) Ling to this chapter. Dick is the originator of S&OP.

Dick Ling has been helping companies large and small with their business planning processes for over 40 years. He has found that most companies can improve their business planning with some help and the right focus. Dick has a well-deserved reputation as an excellent counselor and problem solver.

Dick has experience with IBM, Arista Information Systems, Xerox, and the Oliver Wight Companies. He is an educator, software developer, author, counselor, problem solver, and architect of improved business planning processes. For the past 25 years since he pioneered the development of S&OP, he has been educating companies about S&OP and helping them implement successfully.

Dick views contributing to this chapter in two ways: as a means to help link DDS&OP and DDMRP, a truly exciting alliance, and as a tribute to his longtime friend and collaborator, Andy Coldrick, who passed away in December 2014.

Dick Ling and Andy Coldrick collaborated on S&OP for 25 years. They formed a very strong partnership and specialized in pushing the boundaries of S&OP. They helped businesses all over the world to maximize S&OP's potential to generate more cash and increase profits. Dick created S&OP, and he and Andy have been two of the leading thinkers and consultants on its evolution and advancement.

They led the thinking on aligning the S&OP process with the strategic intent of the business and future portfolio. Before that, they were the first to recognize that new-product activity and financial links to traditional S&OP were being treated as afterthoughts and were not being truly integrated. They pioneered integration of these two pieces and also created the integrated reconciliation step to explode the single-number myth in vogue at the time. The importance of understanding change, assumption management, and scenario planning with a range of views all reinforced management's need for information that built knowledge and know-how rather than data just supplying more and more numbers.

CHAPTER 14

Implications for Technology

Souder's law states, "Repetition does not establish validity" Simply doing something over and over again does not make it the right thing to do; it simply means it is the routine thing to do. The point of software is to enable and reinforce routine at scale and velocity. Yet if that routine is not appropriate, software becomes not an enabler of success but a generator of waste and an inhibitor of the ability to manage assets.

Thus we have come full circle. Joe Orlicky had very descriptive words for precomputer inventory management systems in his groundbreaking book, *Material Requirements Planning*, back in 1975. His description is relevant once again:

> Traditional inventory management approaches, in pre-computer days, could obviously not go beyond the limits imposed by the information processing tools available at the time. Because of this almost all of those approaches and techniques suffered from imperfection. They simply represented the best that could be done under the circumstances. They acted as a crutch and incorporated summary, shortcut and approximation methods, often based on tenuous or quite unrealistic assumptions, sometimes force-fitting concepts to reality so as to permit the use of a technique. (p. 4)

Conventional planning systems are acting as a crutch. They do incorporate summary, shortcut, and approximation methods based on tenuous or quite unrealistic assumptions. Their mandated use force-fits concepts to reality so as to permit the use of the techniques embedded in them. Of course, this means that an alternative way must be proposed, and that will have huge implications for formal planning systems. Are we ready for that change?

Operations and Information Technology— Two Ships Diverging in the

Night?

It appears that we have reached the point of diminishing returns. We can confine ourselves to and keep trying to optimize systems based on tenuous and unrealistic assumptions, or we can seek to break from convention for a true step change in performance. The emergence of complex and volatile supply chains has created a fundamental gap between what the operations function needs in order to stay competitive and innovative and what the current systems allow operations to do. The proof of this gap is found in the widespread proliferation of work-arounds based primarily on individualized, error-prone, and nonscalable spreadsheets. Planning personnel actually believe they are the lesser of the evils. What can close that gap?

There are two major stumbling blocks to effectively closing the gap. The primary obstacle has to do with the divergence between IT and operations that is taking place in most organizations. There is a huge undercurrent of frustration within most supply chain–centric companies with regard to planning and control systems. Operations personnel are frustrated at the lack of solutions made available to them to combat the volatility and complexity they are experiencing in the environment (thus the need for work-arounds). From many perspectives they see information technology closing doors to innovation and improvement rather than opening them.

On the other hand, IT has become extremely frustrated with operations for working around and outside the system. IT is concerned with things like security, data, and transactional integrity. Working outside the system poses risks to all those areas. Even when new techniques deliver promising returns, the techniques can be easily culled by IT because they have not been done in the system. And when informed that the system is incapable of supporting the technique? The response is akin to, “That can’t be; we have one of the best ERP products on the market. It supports all the best practices out there.”

The drive by top management, guided by IT, to force everyone to a uniform use of the existing technology—technology that the company spent a small fortune and tremendous time and effort on—is understandable. It is problematic, however, when considering that the rules behind it are force-fitting techniques based on tenuous and unrealistic assumptions. Operations and distribution are being asked to work in a manner that supports the software rather than having the software work in a manner that best supports the operational flow and the strategic market objectives of the company.

Decades ago IT generally reported to finance, and a logical case could be made based on cash flow and operational needs. The CFO who understood the company's operations and grew up in the company's plants immediately grasped the ideas and the need for change. In the New Normal, information technology is core to the business, increasingly complex, and very political. IT now reports directly to the CEO or COO, and in many instances IT appears to be wholly disconnected from the functions it provides service to and the strategic objectives of the business. In many cases IT has an all too cozy relationship with the ERP provider and implementation partner. Does IT even understand the nature of the company's operations and the global environment it competes in?

This may appear to be a rant against IT. It isn't. The authors are simply relaying what they are seeing inside large multinational corporations. IT projects take forever to negotiate and execute within the company, and both IT and operations people typically walk away shaking their heads and disappointed. Operations and IT seem to be diverging at an alarming rate. We have observed IT organizations that are completely incapable of understanding and orienting around the business needs and the idea of system flow. We have observed that operations people have little knowledge of and regard for the needs and objectives of the IT organization. Both sides simply don't know what they don't know. The solution will require convergence rather than continued divergence. Where to find that point of convergence?

The second obstacle is ontological in nature—what type of reality do major software providers assume supply chains are attempting to control and manage? In the authors' experience working with dozens of different ERP systems, today's ERP, MRP, and DRP systems are simply stuck in the rules from decades past dominated by linear and cost-centric thinking simply because it has always been that way in their history.

Do major software providers even remotely understand the nature of the problem? At the time of this writing, the authors' personal experiences and knowledge of large software providers' inner workings say the answer is "no" or, at best, "not yet." Will the tide turn? Fundamentally, it is very simple. Customers need to start demanding more appropriate planning tools from big software, and that requires the removal of the primary obstacle to improve flow.

DDMRP Software Compliance Criteria

This book has been written in order to convey a blueprint to bring operations and information technology back into convergence with regard to supply chain planning and execution. Additionally, the book is meant to serve as a blueprint to software providers by presenting specifications for coding more effective planning and control systems in the New Normal.

The authors have established the following basic criteria for software to be compliant to the DDMRP method. The compliance criteria are intended to ensure that a piece of software has enough features and functions to implement, sustain, and even improve a DDMRP implementation. The criteria were written in such a way that they ensure compliance to the fundamental principles of the method but allow sufficient open space for competitive differences, creativity, and innovation.

Component 1: Inventory Positioning

- The software must be able to calculate and identify the decoupled lead time for manufactured items.
- If the decoupled lead time calculation cannot be performed, then DDMRP compliance must be limited to purchased and distributed parts only.

Component 2: Buffer Profiles

- The software must be able to group parts into independently managed families with variable settings for zone impact.
- The software must be able to calculate DDMRP buffers and zone values using a combination of buffer profile attributes and the individual part traits of usage, lead time, and order multiple or order cycle.

Component 3: Dynamic Buffer Adjustments

- The software must have a provision for dynamically altering buffers for planned or anticipated events.

Component 4: Demand Driven Planning

- The software must be able to perform the DDMRP net flow equation properly, including qualifying sales order demand (due today, past due, and qualified spikes).
- The software must be able to properly display net flow status (color,

percentage, and quantity) for easy prioritization and supply order generation.

- All elements of the net flow equation should be visible or easily accessible on the planner workbench.

Component 5: Highly Visible and Collaborative Execution

- The software should display alerts based on the on-hand buffer status for decoupled positions.

Of course, larger, more complex entities will need deeper and richer features. This list represents only the basic features that any DDMRP-compliant planning and execution system should have.

PART 3

Appendices

APPENDIX A

An MRP Example

Readers can watch a video version of this example in the video section of http://demanddriven institute.com. The video is called *The Conventional Planning Puzzle—Just How Crazy Does MRP Make Your Life?"* This exercise will assume that all the MRP requirements and assumptions (as described in Chapter 3) are 100 percent true.

The Scenario

This company makes two end items (product A and product L). Figure A-1 is the bill of material for each item. A is made from one B and one C. Each B is made from one D. L is made from one D and one F.

Figure A-2 describes the characteristics of each part. The lot size is a policy decision that controls how many of that part will be ordered at a time. A POS4 is a period of supply for four periods. That means that when A is required, MRP will look forward and build a variable quantity such that A should be built only once every four periods. The other three parts (L, B, D) have minimum order quantities (MOQs) of 250, 100, and 1,000, respectively. These can be set by policy, by economic order quantity, by supplier requirement, or for a variety of other reasons. An MOQ means that the order quantity is fixed but the timing will vary.

The low-level code is the lowest level in any bill of material where that part is structured. Even though D is at level 1 for part L, it is at level 2 for part A, so the lowest level where it is structured would be 2. The lead time is the manufacturing lead time for parts A, L, and B and the purchasing lead time for D. Safety stock has been set on the part levels for A and L and on the purchasing material for D. This is the most common configuration for safety stock since it is intended to protect against variability of supply or demand.

Simulating the Scenario

Figure A-3 is the starting situation for part A. Part A starts with 100 on hand but has 15 allocated already to another order. Also in the first time period, there are 20 scheduled to be received on a previous open order. That means the available inventory would then be 95 = 100 (on-hand) – 15 (allocated) - 10 (requirement) + 20 (scheduled receipt). In the next period the projected on-hand inventory would be 85 = 95 (from period 1) – 10 (requirement). In period 3 the projected available balance falls to 70. Note that in period 5 the projected on-hand quantity is equal to the safety stock. There is no planned order generated until the projected available balance is less than the safety stock, which happens in period 6. The projected on-hand quantity would be 5, which is less than the safety stock, so a planned order would look out the next four periods and include those requirements. In this case the planned order release would be 45 and would occur in period 3 to allow for the three weeks of manufacturing lead time. Period 3 will be the required date for parts B and C, and both will have a requirement of 45.

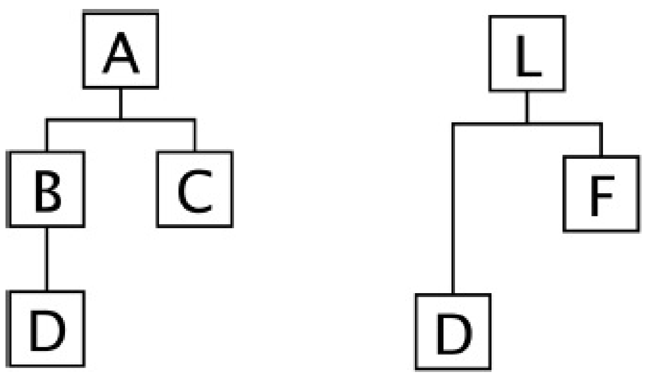

FIGURE A-1 Example bills of material

Part #	Lot Size	Low-Level Code	Lead Time	Safety Stock
A	POS4	0	3	25
L	250 MOQ	0	6	100
B	100 MOQ	1	4	0
D	1,000 MOQ	2	5	100

FIGURE A-2 The static data for each part

The planner would also receive an error message requesting that the scheduled receipt for period 1 be moved to period 5 since there is too much inventory with the arrival in period 1. Unfortunately, the planner ignores that message and allows those parts to arrive.

Figure A-4 is the starting situation for item L. Similar to item A, the available inventory for item L is 280 = 50 (on-hand) – 20 (requirement) + 250 (scheduled receipt). The master schedule notes that there will be two periods of spike demand in periods 4 and 5. There is sufficient inventory for the first order, but the second large order kicks off a planned order receipt in period 5. MRP will plan to receive an order in period 5 to replenish the safety stock. However, the order release should have been done two weeks ago since the lead time for the part is six weeks. The planner will receive a "past due for order release" error message. The requirement of 250 will flow to item D and show past due as a requirement.

Item ID: A

Lot size = POS4 Low level = 0 On-hand = 100 LT = 3 Alloc = 15 SS = 25

	1	2	3	4	5	6	7	8
Requirement	10	10	15	20	25	20	15	10
Scheduled receipts	20							
Projected on-hand	95	85	70	50	25	50	35	25
Planned order receipts						45		
Planned order releases			45					

FIGURE A-3 The initial MRP planning tableau for part A

Item ID: L

Lot size = 250 MOQ Low level = 0 On-hand = 50 LT = 6 Alloc = 0 SS = 100

	1	2	3	4	5	6	7	8
Requirement	20	20	20	120	120	20	20	20
Scheduled receipts	250							
Projected on-hand	280	260	240	120	250	230	210	190
Planned order receipts					250			
Planned order releases	250							

FIGURE A-4 The initial MRP planning tableau for part L

Figure A-5 is the starting situation for part B. B has a starting balance of 50 and no requirements in period 1. This part has a service part demand of 10 units every other week. This is marked as independent demand since the demand is coming directly from the market. The dependent demand is coming from item A and is exactly what is required by A and exactly the timing required by A. However, a problem with item B is that it also has insufficient lead time to be made—two weeks instead of the required four weeks. The planner will receive a past due order release and expedite error message today.

The planner for part D was fine when he left yesterday. In fact, he was probably feeling pretty good since he had twice the safety stock on hand. When the planner arrives this morning, as shown in Figure A-6, he is greeted by an order past due for release and an expedite message. Being a good planner, he picks up the phone and calls the supplier. The first question the supplier asks is how many are *really* needed. The planner responds that 250 units are needed desperately. This is to get the 150 required for the two orders plus replenish the safety stock. This planner went from having too much inventory to too little in the blink of an eye. He tries to solve the problem by offering the supplier an expedite fee, air freight, an increase in price—anything to get those parts shipped in. However, the best the supplier can do is to deliver 200 in period 4 and the balance in normal lead time in period 6. Figure A-7 is the new planning tableau for D given the very best that this planner can do.

Item ID: B

Lot size = 100 MOQ   Low level = 1   On-hand = 50   LT = 4   Alloc = 0   SS = 0

	1	2	3	4	5	6	7	8
Requirement (dependent demand)			45					
Requirement (independent demand)		10		10		10		10
Scheduled receipts								
Projected on-hand	50	40	95	85	85	75	75	65
Planned order receipts			100					
Planned order releases	100							

FIGURE A-5 The first planning tableau for part B

Item ID: D

Lot size = 1,000 MOQ Low level = 2 On-hand = 200 LT = 5 Alloc = 0 SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement	350								
Scheduled receipts									
Projected on-hand	−150	850	850	850	850	850	850	850	850
Planned order receipts		1,000							
Planned order releases	1,000								

FIGURE A-6 The first planning tableau for part D

Item ID: D

Lot size = 1,000 MOQ Low level = 2 On-hand = 200 LT = 5 Alloc = 0 SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement	350								
Scheduled receipts					200		800		
Projected on-hand	−150	−150	−150	−150	50	50	850	850	850
Planned order receipts									
Planned order releases									

FIGURE A-7 Updated tableau for part D

This planner will continue to receive daily error messages requesting that the 200 due in period 4 be expedited, in addition to receiving a message to expedite the 800 from period 6. However, this is the very best that can be done for this part. The typical reaction for a planner is to suppress the expedite messages.

Now what about the other planners? This is where the planners start to break the rules to bring things into balance. The planner for part B is oblivious to the problem with part D since his planner's report only provides information about part B. There is no MRP error message about a component that will not be available. The assumption is made that all parts will be available at the time of the order release. This planner comes to work early this morning and immediately releases an order for B after checking that indeed there is sufficient stock of D for that order. The B part planner does not check the full planning for D since it does not affect his part B. He is anxious to get that order to the shopfloor since it is already a short lead time order. The order is supposed to have a four-week lead time. The planner releases the order and issues it with a big expedite flag that it is due in two weeks. Figure A-8 is the updated view of B.

Now item D is in the situation depicted by Figure A-9. The planner for D has done everything that he can do because the earliest the order can be released is today. No matter how hard the planner for part D works, he cannot release an order earlier than today. He has issued this order with only half the expected lead time, which will likely cause an issue on the shop floor.

The planner for item L stopped for an extra cup of coffee on his way into work this morning and arrives after the planner for item B. Upon his arrival, he is greeted with the realization that not only is the order for L past due for release; he also does not have the supply of D he needs to build L. So he violates the lot sizing rule and releases an order for 100 L since that is all the D that is available. Also he issues the order with a big red expedite sticker and puts a due date of week 5 on it. He is feeling pretty good about this since it gets L back to safety stock in week 5. However, the MRP system will now plan an additional order for week 6 to restore the safety stock. This order is also past due for release, but the planner can't do anything since there will be no D until week 4. Figure A-10 is the new grid for L.

Item ID: B

Lot size = 100 MOQ Low level = 1 On-hand = 50 LT = 4 Alloc = 0 SS = 0

	1	2	3	4	5	6	7	8
Requirement (dependent demand)			45					
Requirement (independent demand)		10		10		10		10
Scheduled receipts			100					
Projected on-hand	50	40	95	85	85	75	75	65
Planned order receipts								
Planned order releases								

FIGURE A-8 The second planning tableau for item B

Item ID: D

Lot size = 1,000 MOQ Low level = 2 On-hand = 100 LT = 5 Alloc = 0 SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement	250								
Scheduled receipts					200		800		
Projected on-hand	−150	−150	−150	−150	50	50	850	850	850
Planned order receipts									
Planned order releases									

FIGURE A-9 The update planning tableau for item D

Item ID: L

Lot size = 250 MOQ Low level = 0 On-hand = 50 LT = 6 Alloc = 0 SS = 100

	1	2	3	4	5	6	7	8
Requirement	20	20	20	120	120	20	20	20
Scheduled receipts	250				100			
Projected on-hand	280	260	240	120	100	330	310	290
Planned order receipts						250		
Planned order releases	250							

FIGURE A-10 New planning tableau for L

Figure A-11 shows the impact on D of the L planner's decisions. Now there is even a higher requirement that is past due, and the buyer tries to call the supplier again to see if the 200 coming in can be increased. The planner vows to increase the safety stock number and the lead time so that this will never happen again!

The net result of this single day is that:

- D is short and cannot support the parent order schedule. D does not recover for five weeks. In one day the part went from having double the safety stock to a severe shortage.
- Expedite fees have been paid to try and expedite D.

- Lead times were violated on three parts—D, B, and L.
- Lot size rules were violated on two parts—D and L.
- Safety stocks would most likely be increased on D and L and further exacerbate the situation

Item ID: D

Lot size = 1,000 MOQ Low level = 2 On-hand = 0 LT = 5 Alloc = 0 SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement	250								
Scheduled receipts					200		800		
Projected on-hand		−250	−250	−250	50	50	850	850	850
Planned order receipts									
Planned order releases									

FIGURE A-11 Impact on part D

As an intellectual exercise, the interesting thing is that if only lot sizing were removed as a rule, the company would have sufficient inventory to fulfill all requirements. See the part planning tableaus in Figure A-12 with only the lot size rules changed for all parts to lot-for-lot (LFL) planning. The safety stocks are still kept as designed. Only one part falls below the safety stock causing an expedite—and that is just to replenish the safety stock, not because it is really needed for production. There is sufficient inventory to cover the immediate requirements. There are still some past due order releases, but they are significantly smaller than before.

However, in the real world, typically there isn't the luxury of being able to order everything only in the quantities that are needed. Usually some kind of batching is employed to save setups or to minimize costs.

This short exercise explores why people use spreadsheets for planning! Meeting commitments means working around the system. The bigger the company and deeper the bills of material, the worse the problem.

Item ID: A

Lot size = LFL   Low level = 0   On-hand = 100   LT = 3   Alloc = 15   SS = 25

	1	2	3	4	5	6	7	8
Requirement	10	10	15	20	25	20	15	10
Scheduled receipts	20							
Projected on-hand	95	85	70	50	25	25	25	25
Planned order receipts						20	15	10
Planned order releases			20	15	10			

Item ID: L

Lot size = LFL   Low level = 0   On-hand = 50   LT = 6   Alloc = 0   SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement		20	20	20	120	120	20	20	20
Scheduled receipts		250							
Projected on-hand		280	260	240	120	100	100	100	100
Planned order receipts						100	20	20	20
Planned order releases	120	20	20						

Item ID: B

Lot size = LFL   Low level = 1   On-hand = 50   LT = 4   Alloc = 0   SS = 0

	Past Due	1	2	3	4	5	6	7	8
Requirement (dependent demand)				20	15	10			
Requirement (independent demand)			10		10		10		10
Scheduled receipts									
Projected on-hand		50	40	20	0	0	0	0	0
Planned order receipts					5	10	10		10
Planned order releases	5	10	10		10				

Item ID: D

Lot size = LFL   Low level = 2   On-hand = 200   LT = 5   Alloc = 0   SS = 100

	Past Due	1	2	3	4	5	6	7	8
Requirement	125	30	30		10				
Scheduled receipts									
Projected on-hand		145	115	115	105	105	105	105	105
Planned order receipts		100							
Planned order releases	100								

FIGURE A-12 Part planning tableaus in an LFL scenario

APPENDIX B

Simulating DDMRP Buffers

This appendix will reveal the results of simulations done using DDMRP tactics in highly volatile and random environments. In order to demonstrate the resilience of DDMRP buffers, the simulations used random input from individuals at the 2014 APICS International Conference and Exposition in New Orleans. The results of the simulation were then displayed in front of a live audience.

About the Simulation

Each participant was handed a card and asked to provide 10 days of demand input, which could be of any value from 0 to 500. There were no other restrictions placed on the input provided. The completed cards were simply sequenced as they were submitted by the visitors. The sequenced demand records were then entered into a DDMRP simulation tool. Figure B-1 is an example of a completed card. In this one card you can see a high degree of demand variability.

The simulation tool contained all relevant features for buffer sizing and supply order generation as described in this book, including dynamically recalculated buffers and order spike qualification as part of the net flow equation. The order spike horizon was set to only three days in advance. This limited the benefit that the order spike qualification would have on the results achieved by the buffer. The order spike threshold was set at 10 percent of the red zone. This will result in more order spikes being qualified than using a 50 percent setting. The simulation obeys lead time, meaning there is no expediting, yet also no supply delays.

Each simulation featured two fictitious parts. The first part is called "Widget." Figure B-2 displays the buffer settings for Widget. Widget is a purchased item. It is in the long lead time category and uses a 25 percent lead time factor. It was placed in a high variability category and given a variability factor of 120 percent.

The second part is called "Gazoonk." Figure B-3 displays the buffer settings for Gazoonk. Gazoonk is a manufactured item. It is in the medium lead time category and uses a 40 percent lead time factor. It was placed in a high variability category and given a variability factor of 120 percent.

To achieve representative simulation results, it is necessary to establish a period of prior history to enable the buffer to settle into an expected behavior. In the simulation tool we create a full year of prior history based on 11 months at an assumed starting ADU. The opening ADU value was set at 30 for both parts for those first 11 months.

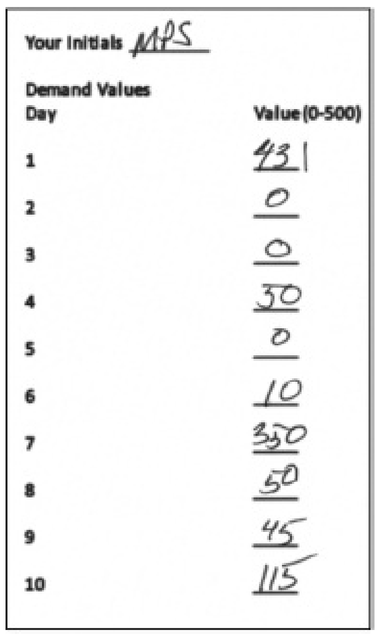
Your Initials MPS

Demand Values

Day	Value (0-500)
1	431
2	0
3	0
4	30
5	0
6	10
7	350
8	50
9	45
10	115

FIGURE B-1 Demand input card example

Buffer profile	P, L (0.25), H (1.2)
MOQ	0
Imposed or desired order cycle (DOC)	0
Decoupled lead time (DLT)	90 days

FIGURE B-2 Buffer data for Widget

Buffer Profile:	P, M (0.4), H (1.2)
MOQ:	0
Imposed or desired order cycle	0
Decoupled lead time:	30 days

FIGURE B-3 Buffer data for Gazoonk

The simulation uses a 90-day past-looking ADU calculation. The final month of prior history is populated with the values provided in the first month of the actual simulation period. This allowed the ADU to normalize to the unknown ADU level provided by the participants. As such, in the month prior to the start of the simulation, the ADU started climbing from the initial setting of 30 to an opening value of 64.32 for the Widget and 81.44 for the Gazoonk. In both simulation trend graphs, you will notice that supply orders arrive early in the simulation due to the prior history. These were based on the increasing buffer requirements in the weeks leading up to the opening day of the simulation window as the normalization was occurring.

Figure B-4 shows the starting buffer levels for both parts given their respective ADU on the opening day of the simulation period. These ADU values, while in the process of normalizing, are still dramatically understated given their respective average demand-level inputs.

Widget Simulation Results

The simulation of the demand for Widget resulted in 100 percent customer service and 6.43 inventory turns for a part with a 90-day lead time. The buffers rapidly increased in size during the first few months of the simulation, as the demand provided by the participants was much higher than the starting ADUs for each part. Figure B-5 is a summary of results from the simulation.

Figure B-6 is a trend graph showing buffer levels and zonal distributions (green, yellow, and red) and on-hand positions (line) over the course of the simulation. The trend graph demonstrates how it quickly adjusted to the greater rate of demand of 151 units per day versus the starting assumption. The actual ADU rate was 2.5 times the opening assumption and pressured the on-hand position, which bottomed out at 1,369 units. However, the incoming supply from orders generated earlier in the simulation put the buffer back into a strong position to maintain service for the balance of the year.

Widget		Gazoonk	
Green zone	**1,158**	Green zone	**513**
	LT factor: 1,158 (5,789 × 0.2)		LT factor: 513 (1,710 × 0.3)
	Minimum order quantity: 0		Minimum order quantity: 0
	Order cycle: 0		Order cycle: 0
Yellow zone	**5,789** [90 (DLT) × 64.32 (ADU)]	Yellow zone	**1,710** [30 (DLT) × 81.44 (ADU)]
Red zone	**3,183** (1,447 + 1,736)	Red zone	**1,505** (684 + 821)
	Base: 1,447 (5,789 × 0.25)		Base: 684 (1,710 × 0.4)
	Safety: 1,736 (1,447 × 1.2)		Safety: 821 (684 × 1.2)
Average on-hand	**3,762**	Average on-hand	**1,762**

FIGURE B-4 Starting buffer levels for Widget and Gazoonk

Simulation Results—Widget	
Average on-hand: 8,551	Minimum on-hand: 1,369
Total demand: 54,967	Maximum on-hand: 14,280
Average daily usage: 151	Service level: 100%
Peak demand: 500	Days stocked out: 0
Supply orders: 20	Annual turns: 6.43
Average order size: 2,911	Average order frequency: 18.25 days

FIGURE B-5 Widget results

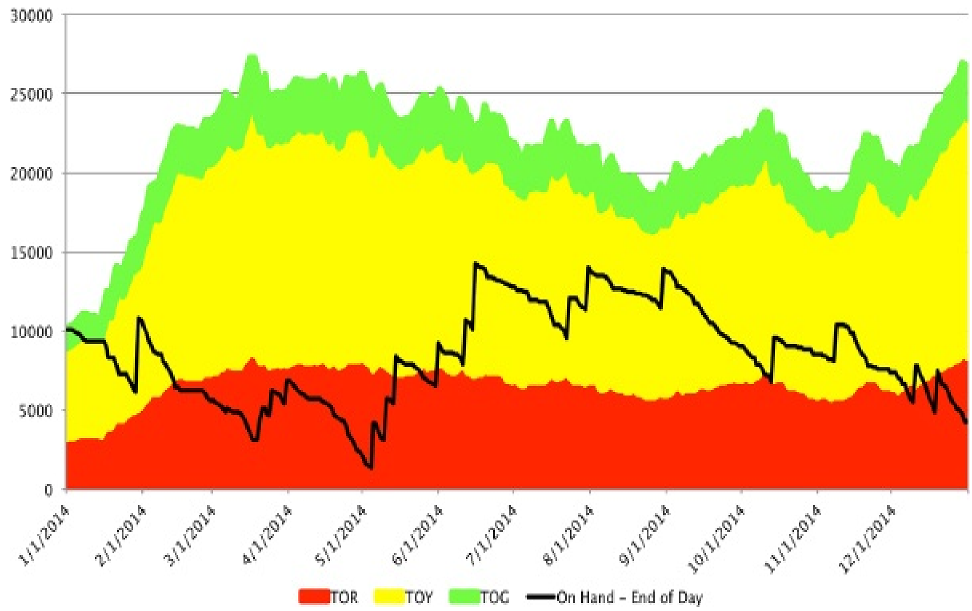

FIGURE B-6 Widget simulation trend graph

The Widget simulation was also run with no order spike qualification. Figure B-7 summarizes the results without order spike qualification.

Without the benefit of three-days forward visibility to sales order demand, the buffer still achieved 100 percent customer service for the year. This makes sense since the order spike horizon is relatively small (3 days) in comparison with the part's lead time (90 days); a small order spike qualification window simply has less value to extremely long lead time parts. Minimum on-hand inventory decreased to 980 units—slightly less than seven days of supply. Figure B-8 shows the trend graph without order spike qualification.

Simulation Results—Widget (no order spike qualification)	
Average on-hand: 8,603	Minimum on-hand: 980
Total demand: 54,967	Maximum on-hand: 13,376
Average daily usage: 151	Service level: 100%
Peak demand: 500	Days stocked out: 0
Supply orders: 20	Annual turns: 6.39
Average order size: 2,936	Average order frequency: 18.25 days

FIGURE B-7 Widget results without order spike qualification

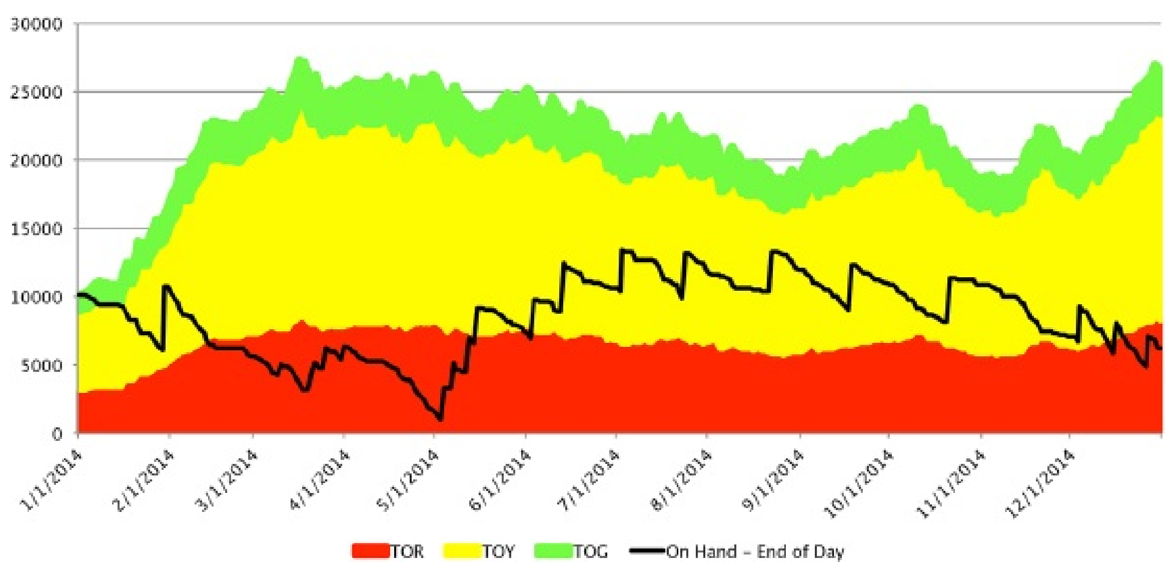

FIGURE B-8 Widget simulation trend graph (no order spike qualification)

In both scenarios 20 supply orders were generated during the year with an average order size of just over 2,900 units.

Gazoonk Simulation Results

Figure B-9 displays the results for the Gazoonk simulation. The simulation resulted in 100 percent customer service and inventory turns of 15.56. The minimum on-hand balance was 380 units and suggests that increased red zone safety coverage would be appropriate to further reduce the risk of stockouts. During the year, 48 supply orders were generated, supporting the rapid turnover rate for the inventory.

Simulation Results—Gazoonk	
Average on-hand: 3,670	Minimum on-hand: 380
Total demand: 57,105	Maximum on-hand: 7,776
Average daily usage: 156	Service level: 100%
Peak demand: 500	Days stocked out: 0
Supply orders: 48	Annual turns: 15.56
Average order size: 1,187	Average order frequency: 7.6 days

FIGURE B-9 Gazoonk simulation results

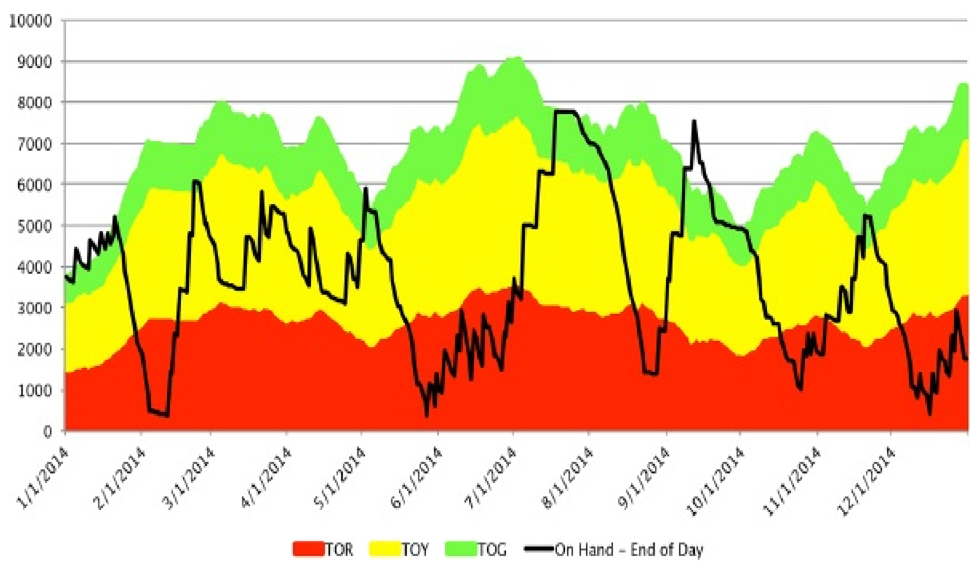

FIGURE B-10 Gazoonk simulation trend graph

Figure B-10 is the Gazoonk simulation trend graph. The buffer size flexed up and down to a greater degree as the average daily usage deviated from the mean in a more “seasonal-type” pattern than Widget. This further illustrates how DDMRP buffers are resilient to changing rates of demand, providing high service while also driving very positive inventory turnover.

As in the Widget example, order spike qualification was turned off and the simulation was run again. Figure B-11 shows the results of the Gazoonk simulation without order spike qualification. Under these conditions the part experienced two days of stockout and a service level of 99.5 percent. The part was also ordered more frequently, as order spikes could not be consolidated into the net flow equation.

Figure B-12 is the trend graph for Gazoonk with no order spike qualification. On several occasions the on-hand positions get dangerously low. In these situations we would expect to see expedite-related activity based on current on-hand or projected on-hand alerts.

Simulation Results—Gazoonk (no order spike qualification)	
Average on-hand: 3,509	Minimum on-hand: (125)
Total demand: 57, 105	Maximum on-hand: 8,328
Average daily usage: 156	Service level: 99.5%
Peak demand: 500	Days stocked out: 2
Supply orders: 51	Annual turns: 16.27
Average order size: 1,137	Average order frequency: 7.1 days

FIGURE B-11 Gazoonk simulation results without order spike qualification

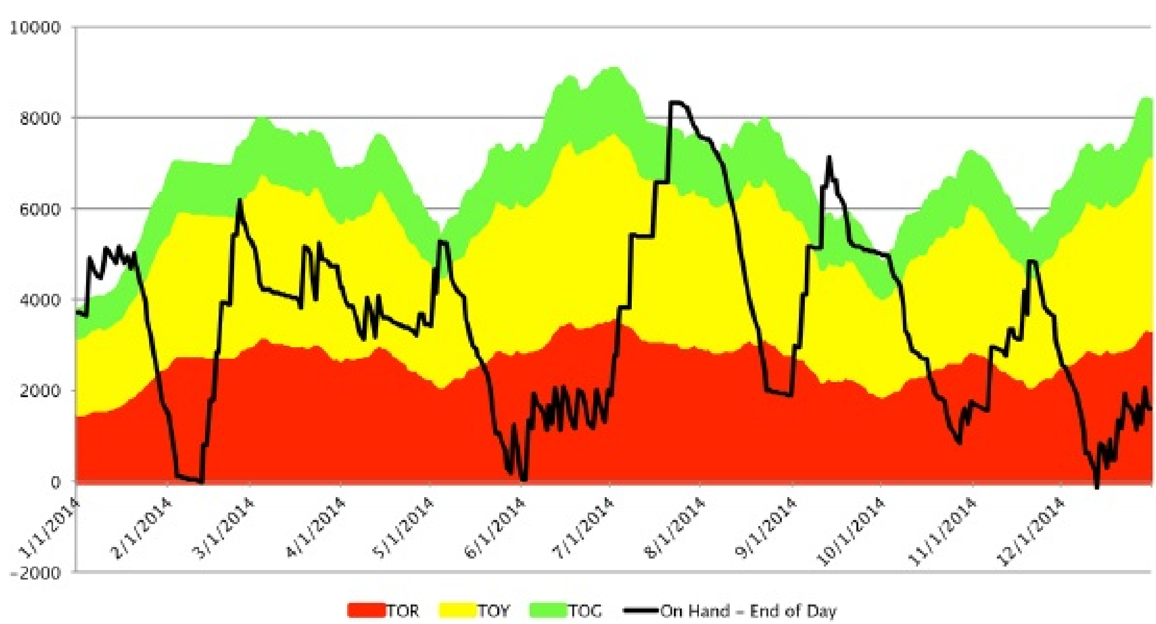

FIGURE B-12 Gazoonk simulation trend graph without order spike qualification

While 99.5 percent customer service would be considered excellent in most companies, an additional simulation was run that adjusted the variability factor to 150 percent. This was an attempt to see what it would take to clear all stockouts in this highly variable environment. Figure B-13 shows the results of the simulation with a larger red zone.

Service improved to 100 percent. Minimum on-hand increased to 235 units, while inventory turnover rate declined slightly to 14.78 annual turns. Again, this result was achieved without any forward visibility to sales order demand. This shows us that even with a limited order spike horizon (3 days) in relation to a part with a 21-day lead time, order spike visibility and qualification make a difference. The closer the order spike horizon, sales visibility horizon, and part lead time are to each other, the more powerful the effect of order spike qualification on the buffer's performance.

Simulation Results—Gazoonk (no order spike qualification with larger red zone)	
Average on-hand: 3,865	Minimum on-hand: 235
Total demand: 57,105	Maximum on-hand: 8,823
Average daily usage: 156	Service level: 100%
Peak demand: 500	Days stocked out: 0
Supply orders: 50	Annual turns: 14.78
Average order size: 1,161	Average order frequency: 7.3 days

FIGURE B-13 Gazoonk simulation results with larger red zone

Simulating the Impact of Minimum Order Quantities

An additional simulation on Gazoonk was performed to demonstrate the impact of a large minimum order quantity (MOQ) on buffer performance. Gazoonk was given an MOQ of 5,000. This represents roughly 33 days of consumption for a part with a 21-day lead time. Obviously this means that the green zone of the Gazoonk buffer will be sized to the MOQ. Figure B-14 summarizes the results of the simulation.

While service was perfect, it came at a cost. Inventory turnover declined roughly 50 percent to 9.07. Supply orders declined from 50 to 10. Average order frequency was over 36 days. Average on-hand became 6,298, which is roughly 42 days of supply. Minimum on-hand increased to 1,141 as a result of the larger and less frequent order size. Figure B-15 shows the Gazoonk trend graph with the large MOQ.

This is somewhat typical of experiences with regard to minimum order quantities representing substantial multiples of usage over the part's lead time. For purchased items, this often represents an ineffective trade-off, as the resulting discount rarely justifies the impact the MOQ has on the flow of materials. The same can be said for minimum batch sizes in production where efficiency metrics cause large "artificial batches" that impede flow and reduce manufacturing responsiveness.

Simulation Results—Gazoonk (Large MOQ)	
Average on-hand: 6,298	Minimum on-hand: 1,141
Total demand: 57,105	Maximum on-hand: 14,740
Average daily usage: 156	Service level: 100%
Peak demand: 500	Days stocked out: 0
Supply orders: 10	Annual turns: 9.07
Average order size: 5,244	Average order frequency: 36.5 days

FIGURE B-14 Gazoonk simulation results with larger MOQ

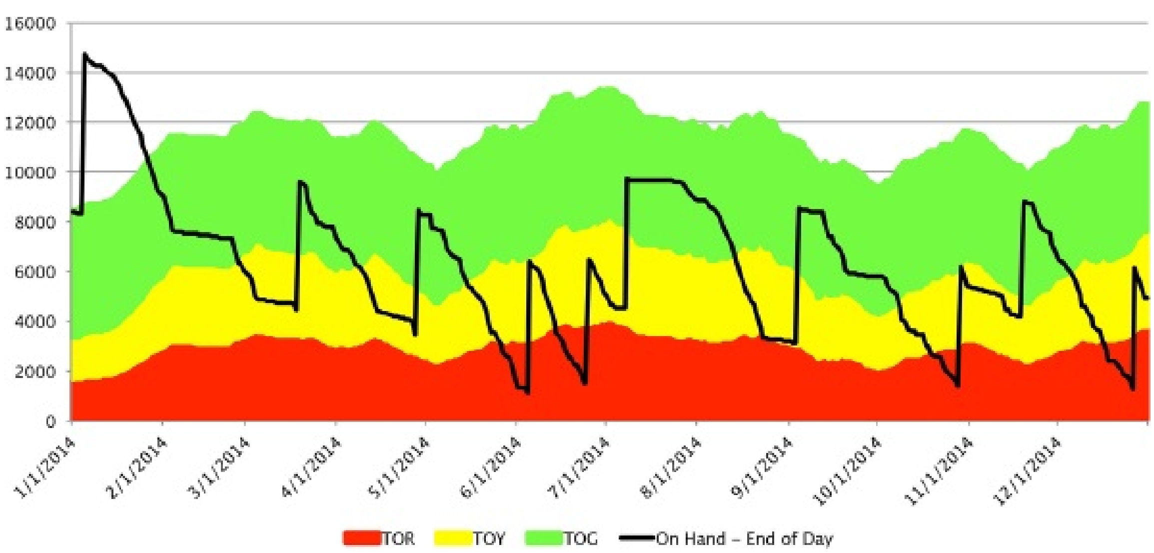

FIGURE B-15 Gazoonk simulation trend graph with large MOQ

Summary

The simulation of buffer performance using random demand values provided by random individuals was a very real and interesting test of the Demand Driven MRP methodology. Other than the upper limit of 500, we had no idea what demand input we'd be getting from the study participants. Both parts that were modeled in the simulation achieved service levels of 100 percent while also driving very solid inventory turnover rates. It's critically important to understand that this performance was achieved with at most three days of forward visibility to demand.

We applied a high-variability safety threshold due to the unknown rate of demand that drove the perfect service levels achieved in the simulation. We also used the simulation to demonstrate how adjusting buffer parameters such as minimum order quantity affects buffer performance.

The core concept of Demand Driven MRP buffers is that they are designed to achieve constant material availability. The resilience of the buffers with regard to that objective was proved in the simulation examples. Supply orders were triggered based on actual sales and the penetration of the buffers. High inventory turn rates were achieved without the prevalent inventory distortions seen in forecast-driven methodologies. DDMRP also provides users with a very easy-to-follow signaling system for planning and supply chain execution.

Additional simulation scenarios and results are available at www.demanddriventech.com.

About the Author

Erik Bush has served as the CEO of Demand Driven Technologies since its formation in the fall of 2011. Under his leadership, DD Tech has quickly moved from a software start-up to a global provider of demand driven supply chain solutions. DD Tech's global client base spans six continents and a vast range of industry segments.

Prior to forming Demand Driven Technologies, Mr. Bush spent 31 years in a variety of roles at IBM. He served as the Vice President of Global Capabilities leading the rapid growth and expansion of IBM’s network of GBS Global Delivery Centers. He led the implementation of sweeping changes that resulted in improved client service, rapid market share growth, and substantially enhanced capabilities. Erik also served as the Vice President of Operations for IBM’s Global Business Services units in Europe and the Americas.

APPENDIX C

Applying DDMRP to the Apparel Retail Environment

The Need for a Retail Application of DDMRP

Retail shops are fundamental links of many supply chains. Worldwide retail sales were estimated at US$22 trillion in 2014. Wal-Mart was the first in the list of the Fortune 500 companies in 2014 and 2015 and will likely maintain that position for several years. Most consumer goods, electronics, apparel, and footwear, all sorts of hardware and home supplies, and many other product categories are delivered to end consumers through retail stores. The actual demand of a considerable amount of goods is driven by retail sales.

But prevailing practices for retail materials planning based on the widely used "forecast-push-and-promote" mode are far from achieving satisfactory results. Current reality proves that the retail industry throughout the world suffers from persistent high stockouts and lost sales coupled with serious inventory excess. In December 2015, the following report was published on the Internet, regarding the performance of retailers in the United States:

> According to a study by IHL Group, "out-of-stocks" accounted for $634.1 billion in lost retail sales for the year ended in the spring—39 percent higher than in 2012. Likewise, overstocks contributed $471.9 billion in lost revenues, up 30 percent from three years prior. When a retailer has too much merchandise, it cuts into its margins.[1]

Developing an effective material requirements planning and execution technique at the retail level that alleviates such severe adverse results is a crucial need.

Demand Driven MRP was designed to be applied mainly in manufacturing companies, including the supply chain portion that stretches from suppliers to distribution centers. The retail level was not considered in its original version. Demand Driven MRP should be able to provide a world-class solution for materials planning and execution at the retail level. However, some of its regular rules do not conform well to materials management in retail due to some specific characteristics of these environments with regard to high uncertainty of new-product sales, sales composition and concentration, buffer sizing, and minimum display quantities

This appendix provides a description of the application of DDMRP to retail, outlining the challenges encountered in these environments and the way they can be overcome, proposing a general framework for what can be called Demand Driven Retail Requirements Planning.

A Retail Apparel DDMRP Example

The first implementation of DDMRP in Latin America, which turned out to be the first DDMRP implementation in the world all the way to the retail level, was led by the author and performed in Maquila Internacional de Confecciones (MIC), a manufacturer of children's clothing based in Medellín, Colombia. This company designs children's garments under licenses from Disney, Mattel, and other similar brands and sells through its own chain of over 90 stores in Colombia, Venezuela, Dominican Republic, and Curaçao.

MIC is also a supplier of small, medium, and large store chains in Colombia. The company designs the products, purchases yarn, and outsources fabrics weaving and finishing. In addition, the company owns the cutting and sewing production facilities needed to manufacture the final products.

What our team found in the initial analysis, when the project started in February 2013, exposed a profound intensification of the usual symptoms of the traditional push-and-promote model. This is the commonly applied model in companies that wrongly focus on reducing unit costs as a way of increasing profits and that show a clear lack of systemic management. These were the main symptoms that were observed:

- Unsatisfactory service levels
- Constant stockouts in finished goods and raw materials
- Long lead times to market
- Purchasing, production, and distribution decisions based on forecasts to the SKU level and intended to minimize unit costs

- Excessive amounts of work-in-process (WIP) inventory throughout the plant to maximize local production efficiencies
- Significant excesses of inventory in the plant warehouse of slow-moving products
- Deliveries to stores every two weeks in order to reduce transportation costs
- A clear silo mentality across the company
- Chronic conflicts between sales and production and between production and purchasing
- Constant “scarcity sensation” in stores despite having excess inventories

It is not surprising that under these practices and conditions, the cash flow of the company was negative.

Supply chains are a combination of inventory and flows that interact in a systemic and dynamic way with feedback loops between both of them. DDMRP deals mainly with inventory. Production scheduling deals with flows, and the manner in which the scheduling is performed determines the production lead times.

Our experience has shown that there is no possible way that a company in the fashion industry (and maybe in many others industries) can achieve good results if there are long lead times across the supply chain. If a company has long lead times in an environment with high uncertainty and variability, it will be almost impossible to become demand driven, properly sensing and quickly adapting to market changes. MIC’s production lead times were as long as 45 days.

We decided that the first task of the project was to substantially reduce production lead times using Little’s law. Little’s law is the fundamental equation of operations of which very few professionals working in manufacturing are aware. In the author’s opinion, working in industrial operations without knowing and using Little’s law is like asking a mechanical engineer to do his or her work without knowing that $F = m * a$.

The scheduling methodology of the Theory of Constraints, referred to as Drum-Buffer-Rope, and the kanban method developed at Toyota both leverage Little’s law to keep WIP at a low and constant level so the lead time is short and constant. Maintaining constant WIP requires that the work that is released into the system be approximately equal to the amount leaving the system.

Our team worked with the company's scheduling team to develop and use a basic spreadsheet that permitted orders to be released in a relatively synchronous manner, defining the sewing operation as the main control point of the system. It was not a world-class solution, but it worked well for this application. The production lead time was reduced to 15 days in a matter of a few weeks—a 67 percent lead time reduction.

At this point, we began with the implementation of DDMRP, leveraging its five components in prerequisite order. However, the lack of formal methodological rules for managing MIC's retail environment meant that we had to develop a working model that could conform to the challenges and particularities of this domain. This construct was developed after spending a significant amount of hours with the planning team of MIC, learning about the existing tools, policies, rules, expectations, and limitations and analyzing how the speed of flow of relevant information and materials could be increased while operating based on actual market demand.

To date (December 2015) this methodology has been fully or partially replicated in four other apparel producers that own their retail stores, obtaining significant and consistent results with regard to sales increases (up to 60 percent), stockouts (less than 1 percent in high movers), and inventory reductions (up to 40 percent), with other qualitative and very valuable effects such as improvements in decision-making procedures, work environment, clarity of purpose, etc.

Special Characteristics and Challenges of the Apparel Retail Environment

There are two significant and specific characteristics of the apparel retail environment that create serious challenges to materials planning:

1. A very high uncertainty of the actual sales of new products launched to the market
2. A very deep and extreme concentration of sales in a relative low number of SKUs

The first one arises from the common practice in the apparel industry of frequently launching new collections and styles following fashion trends. The problem is that there is an extreme uncertainty in the demand for a new product once it arrives in the stores. It is typical to see in these companies that designers and salespeople are constantly surprised by the low sales of a style that they expected to be a market winner, or on the contrary, a not very exciting product becoming a true high mover.

Our experience indicates that only 10 to 20 percent of new styles that are launched to the market will have significant sales figures. These products will be the high movers. Also, around 40 to 60 percent of the new styles will have marginal sales, if any, and will end up in an outlet point sold with large discounts, up to 70 percent. These are the slow-moving items. The remaining styles will have some intermediate sales figures.

The second special characteristic of the retail environment is that the few high-moving styles account for some 40 to 60 percent of the total sales in a retail point. The remaining portion of total sales is spread out in the vast majority of medium- and low-moving styles, each of them having a considerably lower sales volume, producing a very high concentration in the sales distribution as shown in Figure C-1.

This figure depicts real average daily sales and cumulative sales in a retail chain of one of our clients, with 42 stores of a well-known brand. The retail chain carried 761 SKUs. The left-hand vertical axis has the six-month average daily sales for each SKU, sorted in descending order, with average daily sales shown on the solid curve. The right-hand vertical axis corresponds to the cumulative sales of the retail chain represented by the dotted curve.

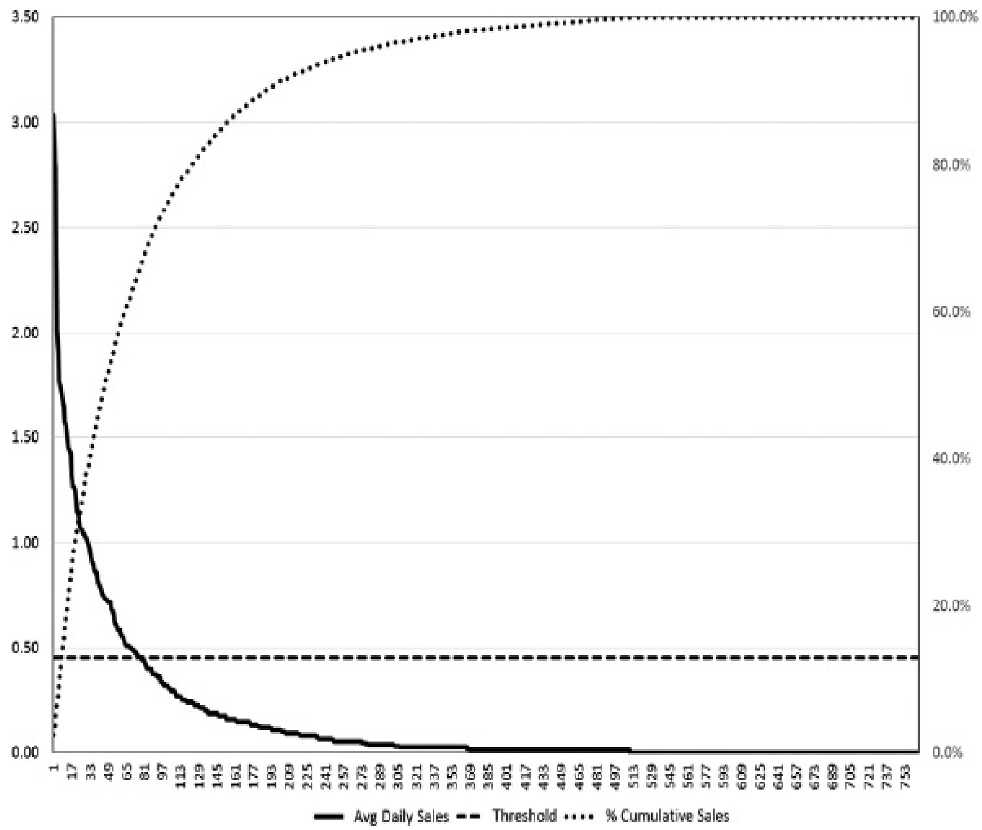

FIGURE C-1 Average daily sales and cumulative sales in a real apparel retail chain

It was a surprise for this client to find that out of the total 761 SKUs, around 440 of them had zero sales in the period. Around 50 SKUs (less than 7 percent) accounted for 60 percent of the total sales of the chain. In this example, the sales concentration is even deeper than what is commonly found.

These figures have strong repercussions for how the entire supply chain of an apparel retail chain should be managed. This also raises sensitive questions about the enormous waste of money, materials, production capacity, labor, transportation costs, etc., caused by slow- and nonmoving items that will hardly be recovered.

They are even of more concern when other industrial sectors like food or medicine are considered, where expiration dates of products require that additional money be devoted to destroy them. There are no legal outlets for food or medicine. The products are simply incinerated due to wrong practices in supply chain planning and execution, while literally millions of people are in desperate need of these very items. Our experience working with pharmaceutical and food companies confirms that this is the case.

As noted earlier, the two general characteristics—extreme uncertainty and deep sales concentration—pose real challenges to material planning in the retail environment, and they call for a model that acknowledges and focuses on these fundamental facts. Under these circumstances, maximizing the speed of flow of relevant information and materials while responding to actual demand is an absolute mandate.

Also as mentioned previously, our experience shows that the standard operating model in this industry is the usual forecast-based push-and-promote, coupled with long delays in receiving the relevant information and with commonly longer-than-30-day production lead times. This model, which we understand is the common denominator in this industry worldwide, works following this general pattern:

1. The design department completes a collection following a previously defined portfolio structure and calendar.
2. The marketing and sales departments forecast expected sales for each style or SKU, for collections that will be released several months ahead.
3. Operations performs a bill of material explosion for all SKUs according to expected sales quantities.
4. Procurement places purchase orders for all the required raw materials, trims, and packaging material.
5. Production manufactures the entire collection against the defined forecast.
6. Products are pushed to the stores, sometimes leaving no buffer at the plant level. No buffers are left in raw materials or trims.
7. The collection is launched and arrives at the stores.
8. The market picks its high-moving items that will be sold out in a matter of weeks.

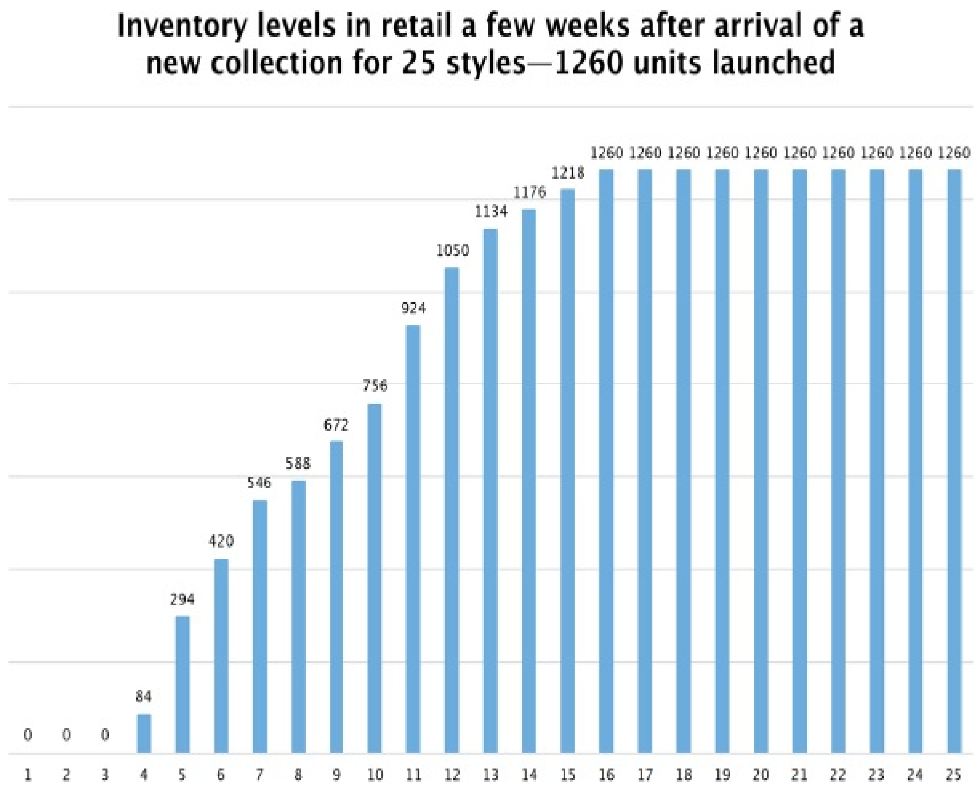

FIGURE C-2 Inventory levels in retail a few weeks after a new collection is launched

9. Stores start having stockouts of the items that account for the majority of the sales, even if a small buffer was left at the central level. This inventory configuration is depicted in Figure C-2. Items 1, 2, and 3 were the high movers. Items 12–25 turned out to be slow movers.

10. There is not an immediate and systematic follow-up of sales levels per SKU. We have encountered companies that measure sales one month after the collection is launched.

11. Even if there were no delays in obtaining and analyzing sales figures, since replenishment lead times in both production and purchasing are usually longer than 30 days, there is no possible way to replenish high-moving items (items that account for 50 to 60 percent of total sales) when the market still wants them. There are significant lost sales.

12. On the other hand, slow movers will not be sold—sometimes not even a single unit. These products consumed materials, working capital, production resources, and capacity and transportations costs. They also occupy very valuable and constrained exhibition space in the stores that could be used by another potential high mover.

13. After some months, these slow-moving items (around 40 to 60 percent of the total collection) are sent to outlet stores and are marked down with deep discounts in order to recover the largest possible amount of their total variable costs and related expenses. This also generates a tremendous opportunity cost, not selling all products at full price in premium stores.
14. The excess inventory is significant, and a substantial amount of working capital was committed for several months, creating long cash cycles.
15. The next collection is designed, and the cycle starts all over again.

The result of this common pattern is devastating: high stock, lost sales, and high inventory excesses with extremely long cash cycles. This is where the figures presented by the IHL group report fit. It is a $1.0+ trillion-dollar problem in the United States alone.

The Proposed Model

At this point, it is absolutely clear that a good model for inventory planning and execution at the retail level must start by acknowledging and addressing the high uncertainty and deep sales concentration issues. In apparel retail (and in any other industry that is constantly launching new short-life-cycle products), it must be recognized that any new product that is delivered to stores is basically a gambling exercise. The rate of success of such a gamble in apparel is rather low —between 10 and 20 percent.

Forecast error is huge. As a matter of fact, some items will not be sold at all at their intended prices.

Since nobody really knows or could know in advance which products will be high movers or slow movers, a radically different model of operation is required.

The proposed model is based on two fundamental propositions of DDMRP, namely:

1. **The first law of supply chain: high-speed flow of relevant information and materials.** Demand sensing must be done on a daily basis, triggering continuous replenishment signals backward in the supply chain. There is an absolute need for

fast identification of high movers and then the need to ensure their full availability, replenishing them in a matter of days.

2. **The first component of DDMRP: strategic inventory positioning.** The retail buffers are replenished by demand from strategic buffers positioned in raw materials, intermediate components, and finished products, carefully and dynamically sized according to the second and third components of DDMRP.

The basic features of the model are conceptually quite simple:

- Produce and send minimum amounts to stores taking the least possible risk.
- Sense the real market demand and identify high movers as soon as possible.
- Replenish the high movers fast.
- Guarantee full availability.

If a new style turns out to be a slow mover, it will not be replenished. Then only a small amount will need to be cleared. The risk taken will be only in the display amounts. The raw materials and trims buffers may eventually be used for other new styles.

This way of operating allows a company to minimize stockouts of high movers, therefore achieving significant sales increases and also minimizing the investment and related expenses in products that turn out to be slow movers and that should not have been produced in the first place (if anyone had the accurate crystal ball to define such behavior in advance).

In order to implement the proposed conceptual model, radical changes in the traditional way of operating must be made. The new pattern should be as follows:

1. The design department should create a buffer of new and unreleased styles according to the company product portfolio.
2. This buffer must be created following the rules indicated by “design for manufacturing” techniques, without losing the brand-specific flavor.
3. Once a set of styles will be launched, the sales department must

produce a forecast. This is absolutely unavoidable if there are long lead times for purchased materials and longer-than-15-day production lead times. If lead times are lower than a week for purchased and produced parts, the system will just adjust rapidly to actual demand without the need for forecasting or large buffers.[2]

4. An initial lot of new products, equal to the visual display of the stores plus a buffer in the plant warehouse, is sized according to the production lead time and expected sales. Obviously, the shorter the production lead time, the lower the buffers will be.
5. Operations performs the material planning explosion for this lot, and launches purchase orders for the required components. If there are purchased parts with long lead times, they should be positioned in strategic buffers.
6. Manufacturing produces the initial lot, and the visual displays are sent to stores.
7. Actual demand is sensed from day 1, and sold items are replenished from the finished products buffer.
8. Within a few days, the planners are able to identify the high runners. Our clients have reported that this trend is easily identified even during a long weekend.
9. An early replenishment production order is issued for high movers. Since either the raw materials and trims are in a buffer or they can be obtained with short lead times, this order will arrive on time to the central warehouse in order to continue replenishing high-moving items, thus avoiding stockouts and maximizing potential sales.
10. Slow-moving items must be cleared in the store or sent to an outlet as soon as possible, making room for new products. This should be a clear company policy.
11. The process is repeated by launching new products taken from the design buffer according to the amount of items that were cleared or sent to outlets and according to the company portfolio structure. Some of the new products will eventually become high movers, which is the key to the retail model: find as many high movers as possible while launching as many products as possible in a coordinated manner and ensure maximum availability of these high movers.

This approach conforms well to how complex adaptive systems (CAS) should be managed. The highly uncertain and emergent demand patterns should be closely sensed, and the entire supply chain execution should quickly adapt to them, in a self-organizing dynamic behavior. Also, in CAS the relevant information that should be monitored is in the tails of the distribution—that is, the high movers that account for a significant portion of sales and brand positioning and the slow movers that should be removed from the stores as soon as possible, allowing for new products to be exhibited and eventually become high movers.

A few months after adopting this Demand Driven Retail Material Planning model, our clients have reported sales increases of up to 60 percent, inventory decreases of up to 45 percent, and therefore significant improvement in cash flow and return on investment.

Retail DDMRP Buffer Zone Considerations

Another major challenge that our team encountered while implementing DDMRP in MIC was the right sizing of the retail buffer zones.

The typical characteristics of the retail environment for clothing are:

- Average daily sales are lower than 0.1 unit in more than 98 percent of the SKUs (or its equivalent: sales frequency greater than 10 days).
- Replenishment lead time to stores is lower than three days. Typically it is one day when DDMRP is implemented.
- Minimum display inventories are required.
- Buffers should take discrete values of 0, 1, 2, 3, or more units.
- There is extreme sensitivity of total inventory to the rounding policies used in the calculation of the buffer zones.
- Higher sales occur on weekends.

These circumstances made it impossible to apply the conventional techniques suggested by DDMRP for buffer sizing because the typical average daily usage and lead time combination would yield levels of zero units for two of the three buffer zones using the regular rounding policies for integer numbers. Additionally, the minimum display quantities would not be respected, and it would not be possible to size buffers of 1 or 2 units, which are the most common in apparel retail environments.

Numerically, if the buffer zones' sizing formulas prescribed by DDMRP are applied to parts with ADU less than 0.1 and lead time equal to 3 days, the yellow zone will be less than an 0.3 unit in more than 98 percent of the SKUs. The red zone will also be less than an 0.5 unit, and the green zone would be equal to 1 unit (the regular MOQ for these environments). But buffer zones must have integer values for obvious reasons.

Rounding down to zero would yield a buffer with zero units in the red and yellow zones and 1 unit in the green zone. The replenishment mechanism of the net flow position would not work since there is not a "below top of green" point. Rounding to 1 unit would unnecessarily and significantly increase the aggregated inventory in the retail chain.

Another requirement of inventory management in retail is related to the minimum display quantities. For visual merchandising reasons, the inventory in a retail store has to have a minimum so the store has a good appearance. This minimum amount should be respected in the buffer sizing. Given these constraints, a new technique and formula for sizing buffers in the retail environment was developed.

The levels of each of the buffer zones should be calculated by considering the average daily usage, lead time, and minimum display; using appropriate rounding rules; and meeting the needs and specific policies of each company regarding product launching and removal. Figure C-3 depicts an example of how the buffer zones are sized according to different ADU values and a given set of the remaining variables.

There are different rules and formulas that could work in different environments and with different company policies, so each specific case should be analyzed. The math is rather simple and should not pose a problem to anyone who would like to apply the rules and formulas to his or her specific environment if the mentioned conditions and requirements are carefully observed.

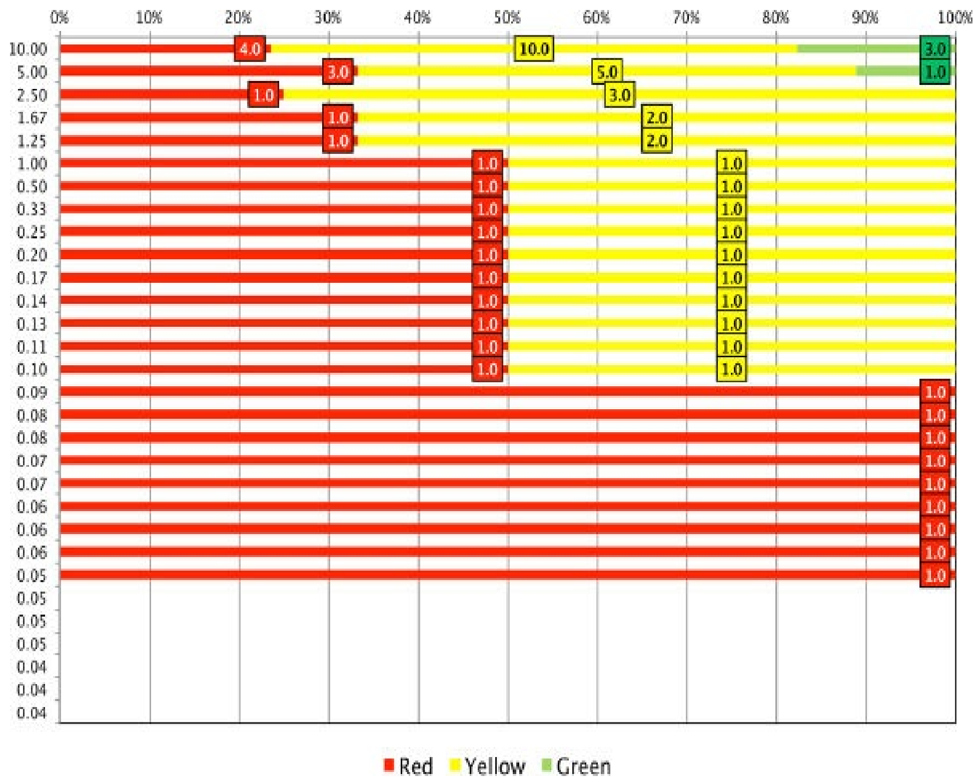

FIGURE C-3 DDMRP buffers for retail—0, 1, 2, or more units

Realized Results

The operations manager of MIC, also a shareholder of the company, has publicly reported the following results achieved after the implementation of DDMRP in the retail store chain[3]:

- 60 percent increase in revenues for the 2013–2015 period
- 40 percent decreased inventory in the retail chain
- Longer product life cycles, requiring less renewal of the product portfolio and reducing the complexity in the supply chain
- Sales of high movers nine times higher during the high season (Christmas) of 2013, compared with those of the same period in 2012
- The elimination of a sense of scarcity in stores despite having a lower inventory
- Decrease in products shipped to outlets and sold with high discounts

- Radical improvement in cash flow of the company
- Synchronization between the different functional areas of the company around the principle of increasing the speed of flow of materials and information

These results are consistent with those reported by other retail implementations of Demand Driven MRP in other countries and other industries. Some of our other current clients report similar results.

In addition, it is worth mentioning that MIC applied the same DDMRP concepts developed for the retail environment to implement a vendor-managed inventory model in 58 stores of El Éxito Group, the largest retailer in Colombia. This implementation significantly contributed to the decision of this group to grant MIC the distinction as "Best Supplier of the Year" in its category for 2014.[4]

El Éxito mentions on its website that this award was given to MIC because "the company developed a supply chain model that has allowed growth levels up to 100%."

In less than 18 months, MIC transformed from being in a deep crisis to achieving this notable recognition. It has continued its demand driven journey, and we are currently implementing a Sales and Operation Planning process leveraging the benefits of the Demand Driven Adaptive System model to ensure the sustainability of the company's results.

About the Author

David Poveda is a civil engineer from the Antioquia School of Engineering, with a master's in project management from the University of British Columbia. He is the founder and General Director of Flowing Consultoría, a Colombian consulting company that specializes in transferring knowledge and supporting implementations of the most advanced methodologies in operations and supply chain management with emphasis in DDMRP and Demand Driven Adaptive Systems. He is an Endorsed Instructor for the Certified Demand Driven Planner Program.

Mr. Poveda served as Managing Director of a midsize steel wire company in the 1990s and since 1997 has dedicated his efforts to consulting with more than 70 companies in South America. He has been an invited speaker to several conferences in supply chain management in Latin America and the United States. He also has served as a member of the board of several industrial companies based in Medellín, Colombia, since 1993.

APPENDIX D

Demand Driven MRP Dictionary

actively synchronized replenishment (ASR) The initial name given to Demand Driven Material Requirements Planning (DDMRP).

ADU *See* average daily usage.

ADU alert An alert indicating a significant change in ADU within a defined set of parameters (quantity and time).

ADU alert horizon A defined shorter rolling range within the broader rolling horizon used to to alert to significant changes to ADU.

ADU alert threshold A defined level of change in ADU that triggers the alert within the ADU alert horizon.

ADU-based recalculation A process of dynamically adjusting strategically replenished buffers incorporating a rolling horizon.

artificial batch Any batch that is not a function of actual demand.

ASR *See* actively synchronized replenishment.

average daily usage (ADU) Average usage of a part, component, or good on a daily basis.

average inventory range This is the expected range in which the on-hand inventory value for a particular buffered part should be on any particular day. The range is defined by the top of the red planning zone of a buffer up to a value of the top of red zone plus the entire green planning zone.

average on-hand position The red planning zone value plus half the green planning zone value of a buffer.

blended ADU ADU calculated based on a combination of historical usage and forecasted usage.

buffer penetration The amount of remaining buffer, typically expressed as a percentage.

buffer profile A globally managed group of parts with similar lead time, variability, control, and order management characteristics.

buffer status alerts Alerts that show the current and projected status of the decoupling point positions across the network of dependencies.

buffer zone A stratification layer within a stock buffer. Typically, buffer zones are color-coded with red, yellow, and green assignments.

CDDL *See* Certified Demand Driven Leader.

CDDP *See* Certified Demand Driven Planner.

Certified Demand Driven Leader (CDDL) A professional certificate from the Demand Driven Institute and International Supply Chain Education Alliance proclaiming that a person has successfully tested for proficiency in the aspects of the Demand Driven Operating Model.

Certified Demand Driven Planner (CDDP) A professional certificate from the Demand Driven Institute and International Supply Chain Education Alliance proclaiming that a person has successfully tested for proficiency in the DDMRP method.

control points Strategic locations in the logical product structure for a product or family that simplifies the planning, scheduling, and control functions (refer to the *APICS Dictionary).*

current on-hand alert An execution alert generated by current on-hand penetration into the red zone of the buffer.

customer tolerance time The amount of time potential customers are willing to wait for the delivery of a good or a service.

DDAS *See* Demand Driven Adaptive System.

DDMRP *See* Demand Driven Material Requirements Planning.

DDOM *See* Demand Driven Operating Model.

DDS&OP *See* Demand Driven Sales and Operations Planning.

decoupled explosion The cessation of a bill of material explosion at any buffered or stocked position.

decoupled lead time (DLT) A qualified cumulative lead time defined as the longest unprotected or unbuffered sequence in a bill of material.

demand adjustment factor (DAF) A planned adjustment that is a manipulation of the ADU input at a specific time period.

Demand Driven Adaptive System (DDAS) A management and operational system designed for complex and volatile manufacturers and supply chains. A Demand Driven Adaptive System uses a constant system of feedback that connects the business strategy to the settings and performance of a Demand Driven Operating Model through a Demand Driven Sales and Operations Planning Process. A Demand Driven Adaptive System focuses on the protection and promotion of the flow of relevant information and materials in both the strategic (annual, quarterly, and monthly) and tactical (hourly, daily, and weekly) relevant ranges of decision making in order to optimize return on equity performance as change occurs.

Demand Driven Material Requirements Planning (DDMRP) A method to model, plan, and manage supply chains to protect and promote the flow of relevant information and materials. DDMRP is the supply order generation and management engine of a Demand Driven Operating Model.

Demand Driven Operating Model (DDOM) A supply order generation, operational scheduling, and execution model utilizing actual demand in combination with strategic decoupling and control points and stock, time, and capacity buffers in order to create a predictable and agile system that promotes and protects the flow of relevant information and materials within the tactical relevant operational range (hourly, daily, and weekly). A Demand Driven Operating Model's key parameters are set through the Demand Driven Sales and Operations Planning process to meet the stated business and market objectives while minimizing working capital and expedite-related expenses.

Demand Driven Sales and Operations Planning (DDS&OP) A bidirectional integration point in a Demand Driven Adaptive System between the strategic (annual, quarterly, and monthly) and tactical (hourly, daily, and weekly) relevant ranges of decision making. DDS&OP sets key parameters of a Demand Driven Operating Model based on business strategy, market intelligence, and key business objectives (strategic information and requirements). DDS&OP also projects the model performance based on the strategic information and requirements and various model settings. Additionally, DDS&OP uses variance analysis based on past model performance (reliability, stability, and velocity) to adapt the key parameters of a Demand Driven Operating Model and recommend strategic alterations to the model and project their respective impact on the business.

DLT *See* decoupled lead time.

dynamic buffers Buffer levels that are adjusted either automatically or manually based on changes to key part traits.

execution horizon The life cycle of an order from the time the order is created or released to the time it is closed.

flow index Average order frequency compared across all parts.

forward ADU ADU calculated based on forecast.

green zone The top layer of a replenished, replenished override, and min-max buffer. If the net flow position is in this zone, then no additional supply is created.

lead time adjustment factor A multiplicative factor applied to a part's lead time.

lead time alert An alert or warning generated by an LTM part. An alert will be triggered whenever the part enters a different zone in the buffer. Green is the first alert to be encountered, followed by yellow and then red.

lead time alert zone The zone associated with the percentage of lead time that provides the definition for lead time alerts. The LTM alert zone has three sections color-coded green, yellow, and red.

lead time managed (LTM) part A critical nonstocked part that will have special attention paid to it over its execution horizon. Typically, LTM parts are critical, long lead time components that do not have sufficient volume to justify stocking. A portion of the lead time of the part (typically 33 percent) will have a three-zoned warning applied to it. That portion is typically divided into three equal sections.

LTM part *See* lead time managed part.

market potential lead time The lead time that will allow an increase in price or the capture of additional business through either existing or new-customer channels.

material synchronization alert An alert against a demand allocation(s) when supply is projected to be insufficient to cover the demand at the time the demand is set to occur.

matrix bill of material A chart made up from the bills of material for a number of products in the same or similar families. It is arranged in a matrix with components in columns and parents in rows (or vice versa) so that requirements for common components can be summarized conveniently (refer to the *APICS Dictionary)*.

net flow equation A planning calculation to determine the planning status of a buffered item. The equation is On-hand + on-order (also referred to as open supply) – unfulfilled qualified actual demand. Also known as the "available stock equation."

net flow position The position yielded by the net flow equation against a part's buffer values. Also known as the "available stock position."

nonbuffered parts All parts that are not stocked.

occurrence-based recalculation A method to adjust buffers based on the number and severity of specific occurrences in a predefined fixed interval.

on-hand alert level The percentage of the red zone used by buffer status alerts in order to determine a yellow or red excecution color designation.

order spike horizon A defined future time frame used to qualify order spikes in combination with an order spike threshold. Typically, the order spike horizon is set to the part's decoupled lead time.

order spike threshold A defined amount used to qualify order spikes in combination with an order spike horizon. Typically, the order spike threshold will be expressed as a percentage of the total red zone (or min value) of a part's buffer.

OTOG *See* over the top of green.

Over the top of green (OTOG) A situation in which either the net flow position or on-hand stock is over the top of the defined green zone, indicating an excessive inventory position.

PAF *See* planned adjustment factor.

past ADU ADU calculated based on historical usage.

planned adjustment factor (PAF) Buffer manipulations based on certain strategic, historical, and business intelligence factors.

planned adjustments Manipulations to the buffer equation that affect inventory positions by raising or lowering buffer levels and their corresponding zones at certain points in time. Planned adjustments are often based on certain strategic, historical, and business intelligence factors.

prioritized share An allocation schema utilizing the net flow positions of a group of parts in order to accommodate a specific limitation or requirement.

projected on-hand alert An alert generated by a low projected on-hand position over a part's DLT based on on-hand, open supply, and either actual demand or ADU.

qualified actual demand The demand portion of the net flow equation composed of qualified order spikes, past due demand, and demand due today.

qualified order spike A quantity of combined daily actual demand within the order spike horizon and over the order spike threshold.

ramp-down adjustment Manipulations to the buffer equation that affect inventory positions, lowering buffer levels and their corresponding zones at certain points in time. Ramp-down adjustments typically are used in part deletion.

ramp-up adjustment Manipulations to the buffer equation that affect inventory positions, raising buffer levels and their corresponding zones at certain points in time. Ramp-up adjustments typically are used for part introduction.

red zone The lowest-level zone in a replenished, replenished override, and min-max part buffer. The zone is color-coded red to connote a serious situation. The red zone is the sum of the red zone safety and red zone base.

red zone base The portion of the red zone sized by the lead time factor.

red zone safety The portion of the red zone sized by the variability factor.

relative priority The priority between orders filtering by zone color (general reference) and buffer penetration (discrete reference).

replenished override part A strategically determined and positioned part using a static (buffer zones are manually defined) three-zoned buffer for planning and execution.

replenished part A strategically determined and managed part using a dynamic three-zoned buffer for planning and execution. Buffer zones are calculated using buffer profiles and specific part attributes such as ADU and DLT.

sales order visibility horizon The time frame in which a company typically becomes aware of sales orders or actual dependent demand.

seasonality adjustment Manipulations to the buffer equation that affect inventory positions by adjusting buffers to follow seasonal patterns.

significant minimum order quantity A minimum order quantity that sets the green zone of a buffer.

spike A comparatively large amount of cumulative daily actual demand that qualifies for inclusion into the net flow equation.

stockout (SO) An item that is not immediately available in stock (refer to the *APICS Dictionary)*.

stockout with demand (SOWD) An item that is not immediately available in stock and has a demand requirement. Also known as negative on-hand.

stockout with demand alert A form of on-hand alert, triggered by a strategically stocked item with a lack of inventory on hand and the presence of a demand requirement.

strategic inventory positioning The process of determining where to put inventory that will best protect the system against various forms of variability to best meet market needs, leverage working capital, and mitigate the bullwhip effect.

supply offset Adjusting the timing of the application of a demand adjustment factor to account for long lead time components.

synchronization alerts Alerts designed to highlight problems with regard to dependencies.

thoughtware The analysis and process employed to define the relevant factors and dependencies in an organization or system in order to construct appropriate business rules and operating strategies that maximize velocity, visibility, and equity. Within the DDRMP framework, thoughtware is commonly referred to with regard to applying the inventory positioning factors.

TOG *See* top of green.

top of green (TOG) The quantity of the top level of the green zone. TOG is calculated by summing the red, yellow, and green zones of a buffer.

top of red (TOR) The quantity of the top level of the red zone.

top of yellow (TOY) The quantity of the top level of the yellow zone. TOY is calculated by summing the red and yellow zones.

TOR *See* top of red.

TOY *See* top of yellow.

yellow zone The middle layer of the buffer level coded with yellow to convey a sense of warning. The yellow zone is the rebuild zone for replenished and replenished override buffers.

zone adjustment factor Adjusting part buffer zones by applying a multiplicative factor to the value of the zone.

APPENDIX E

DDS&OP Checklist

A one-size-fits-all universal checklist for DDS&OP is not helpful because of the uniqueness of each company. DDS&OP provides a dynamic adaptable approach for a company to set its strategy and then recognize and respond to real demand. However, the following checklist is a start to ascertain where your company is at the current time. Rather than a yes or no answer, consider your company's progress along a continuum of these considerations:

- Is the purpose of Demand Driven Sales and Operations Planning understood?
- Does the DDS&OP process have an executive process champion?
- Does the DDS&OP process have process step owners?
- Are values and behaviors recognized as critical for a successful DDS&OP process?
- Is the emphasis on the future, understanding change and its impact on the success of the business?
- Are tactical issues and problems discussed and quickly resolved?
- Is there financial and volume integration?
- Do performance management systems reinforce integrated behavior and the discipline of execution?
- Is the DDS&OP process well documented, with documentation for each step updated on a regular basis?
- Does the business make a distinction between information and data?
- Is there a commitment to data integrity?
- Is there a new-product planning review that provides input to the

managing demand process step and the overall DDS&OP process?

- Is there a monthly review of the future unconstrained demand plan for existing and new products based on inputs from marketing, sales, and finance?
- Is there a DDS&OP process at the supply point level that ensures that there is a valid plan to support orders and shipments?
- Is there a process that involves all the business functions in order to develop an integrated set of plans that reconcile the S&OP projection to the business plan?
- Is information presented to the senior business management team on an exception basis with a focus on understanding and managing changes? Are graphics used extensively to improve scheduling?

As Ralph Waldo Emerson said about life, the same is true about DDS&OP—it is about the journey, not the destination.

Endnotes

Chapter 1

1. All the APICS definitions in the book are from the fourteenth edition of the *APICS Dictionary*, (Blackstone, 2013).
2. You can access the report at http://dupress.com/articles/success-or-struggle-roa-as-a-true-measure-of-business-performance/.
3. You can access the article at www.zerohedge.com/news/2013-02-12/how-rookie-excel-error-led-jpmorgan-misreport-its-var-years.

Chapter 5

1. The definition is from *The Free Dictionary*, www.thefreedictionary.com.
2. Liker, J. (2004). *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer*. New York: McGraw-Hill.

Appendix C

1. www.businesswire.com/news/home/20150506005233/en/Research-Report-Retailers-Lose-1.75-Trillion-Revenue.
2. It can be easily proved that purchasing materials from more expensive suppliers but with short lead times or shipping by plane is a very profitable decision. This fact totally goes against the generalized but very unwise practice of buying from cheaper producers in Asia or elsewhere that have several weeks' or even months' replenishment lead times.
3. Gómez, J. (2015). "Reposición por Demanda," Logismaster Congress, Medellín, Colombia.
4. "Ganadores de Proveedores de Éxito 2014," www.grupoexito.com.co/es/proveedores/concurso-proveedores-de-exito/historia-de-ganadores?id=1232.

References

APICS. (2013). *APICS Dictionary*, 14th ed. Chicago: APICS.

Goldratt, E. M., & Cox, D. E. (1984). *The Goal: A Process of Ongoing Improvement.* Great Barrington, MA: North River Press.

Liker, J. (2004). *The Toyota Way:14 Management Principles from the World's Greatest Manufacturer*. New York: McGraw-Hill.

Noreen, B. G. (2013). *Managerial Accounting for Managers*, 3rd ed. New York: McGraw-Hill Education.

Orlicky, J. (1975). *Material Requirements Planning: The New Way of Life in Production and Inventory Management*. New York: McGraw-Hill.

Plossl, G. (1992). *Orlicky's Material Requirements Planning*, 2nd ed. New York: McGraw-Hill.

Ptak, C., & Smith, C. (2011). *Orlicky's Material Requirements Planning*, 3rd ed. New York: McGraw-Hill.

Smith, D., & Smith, C. (2013). *Demand Driven Performance: Using Smart Metrics*. New York: McGraw-Hill.

Index

actively synchronized replenishment (ASR) lead time. *See* decoupled lead time.
actual demand, 23, 133, 149, 237
ADU. *See* average daily usage.
ADU alert, 107
ADU alert horizon, 107
ADU alert threshold, 107
ADU exceptions, 109
ADU-based recalculation, 105
advanced planning and scheduling (APS/APO), 209
allergen sequence, 246
analytics view, 255
ASR lead time. *See* decoupled lead time.
ASRLT. *See* decoupled lead time.
assets vs. liabilities, 93
average daily usage (ADU), 105
average inventory range, 178
average on-hand position, 176, 181
average open supply orders, 181

batches, 33
batching policies, 33
bidirectional benefit test, 231
bill of material, 28, 45, 60, 182, 239, 250
bimodal distribution, 9, 253, 265
blended ADU, 107
BOMs, 28
breakthrough S&OP, 273
Brown, F. Donaldson, 18
buffer, 39, 93, 96
buffer levels, 111
buffer penetration, 149
buffer profile, 97, 293

buffer simulation, 305
buffer status alerts, 207
buffer zone, 96
bullwhip effect, 18, 32, 35
business plan, 268

calculating average daily usage (ADU), 109
calculating buffers, 111
campaign scheduling, 198
capacity, 142, 238, 246, 283
CDDL, 52
CDDP, 51
certified demand driven leader. *See* CDDL.
certified demand driven planner. *See* CDDP.
challenging traditional S&OP, 272
coefficient of variation, 80, 102
common usage of components, 58
common-cause variation, 30
components, 145
conflict with lean or pull, 45
continuous improvement, 263
control points, 247
conventional planning, 22
convergent point, 62, 247
coverage optimization, 195
critical operation protection, 60
critical positioning factors, 57
cumulative lead time, 61, 68
current on-hand alert, 214
customer tolerance time, 58

DDAS. *See* demand driven adaptive system.
DDMRP, 267
DDMRP and retail, 315
DDOM. *See* demand driven operating model.
DDS&OP checklist, 333
decoupled explosion, 182
decoupled lead time, 69
decoupled schedule, 240
decoupling, 34, 37, 41, 45
decoupling point buffers, 39, 231
decoupling point integrity, 253
decoupling points, 37, 96
delay accumulation, 30
deletion, 135

demand, 149, 237, 278
demand adjustment factor, 131
demand driven adaptive system (DDAS), 267
demand driven material requirements planning (DDMRP), 50, 294
demand driven model projections, 269, 280
demand driven operating model, 53, 268
demand driven performance using smart metrics, 34, 51, 54, 286
demand driven sales and operations planning, 53, 269
demand driven variance analysis, 269
demand manipulation, 286
demand signal, 23
demand signal distortion, 32, 37, 57, 209, 231, 234
demand trend sensing, 133
demand uncertainty, 59
demand variability, 80, 101
demand-driven manufacturing, 51
demand-driven MRP, 50, 294
demand-driven planning, 149
Deming, W. Edwards, 17
dependence, 47
dependent demand, 22
discount optimization, 191
dispatching system. *See* priority control disposition.
distribution inventory, 103
distribution networks, 79, 216
distribution positioning, 79
divergent points, 62
DLT. *See* decoupled lead time.
DRP (distribution requirements planning), 52
drum, 317
drum schedule, 246
due date, 208
dynamic adjustment test, 232
dynamic buffers, 294

Einstein, Albert, 47
enterprise resource planning (ERP), 291
excess inventory, 260
execution, 207, 294
explosion of requirements, 183

finite schedule, 245
five-step process (Ling/Coldrick model), 274
flattening the bill of material, 28, 35
flow, 16

flow index, 49, 262
Ford, Henry, 18
forecast accuracy, 25, 278
forecasts, 23
forward ADU, 107
four sources of variation, 59
freight optimization, 195
frequency up update, 107
goal, 17, 246
Goldratt, Dr. Eliyahu, 17, 246
green zone, 97, 112
green zone adjustment, 146

highly visible priority, 207
hub, 84
hybrid model, 90, 188
hypersensitivity. *See* system nervousness.

independence, 47
integrated reconciliation, 275
introduction adjustment, 133
inventory, 93
inventory, assets vs. liabilities, 93
inventory leverage, 60
inventory positioning, 71
inventory, bimodal, 9
ISCEA, 51
item type, 98

JIT, 24

kanbans, 45
lead time, 99, 109, 237
lead time adjustment factor, 100, 112, 148
lead time alert, 226
lead time alert zone, 226
lead time compression, 96
lead-time-managed (LTM) part, 227
lean, 45, 50, 52
length of period, 105
Ling, Richard (Dick), 288
Little’s Law, 16, 317
load manipulation, 285
location, 110
LTM part, 236

managing the portfolio and new activities, 276
manual adjustments, 133
manufacturing lead time (MLT), 60, 67
market intelligence, 268
market potential lead time, 58
master production schedule (MPS), 236, 272
master settings, 267, 271, 280
material requirements planning. *See* MRP.
material synchronization alert , 222
matrix bill of material, 73
metrics and analytics, 251
minimum order quantity (MOQ), 110, 112, 266
min-max (MM) parts, 96, 125, 198
MLT. *See* manufacturing lead time.
model and part parameters, 270
MRP, 3, 21, 35, 45, 50, 52, 57, 160, 297
MRP nervousness, 27
MRP system requirements, 297
multi-hub, 89
multilevel master schedule, 184
multiple plants, 79

nervousness. *See* system nervousness.
net flow equation, 149
net flow position, 158, 190
new normal, 12, 18, 27, 292

occurrence-based recalculation, 127
Ohno, Taiichi, 17
old S&OP, 272
on-hand alert level, 214
on-hand balance, 150
on-order stock, 158
open orders, 150
order cycle, 112
order frequency variance, 261
order independence test, 232
order spike horizon, 152
order spike threshold, 151
order-point systems, 235
Orlicky, Joseph, 291
Orlicky's MRP 3rd edition, 51
OSH. *See* order spike horizon.
OSH1. *See* order spike horizon.
OSH2. *See* order spike horizon, 156

OST. *See* order spike threshhold.
OTOG. *See* over tip of green.
outlying events report, 257
over top of green (OTOG), 161, 194
overflattening, 45
oversimplification, 47

PAF. *See* planned adjustment factor.
part attributes, 105
past ADU, 107
planned adjustment factor (PAF), 131
planning, 207
planning horizon, 26, 49, 82, 183, 233, 279
planning projections, 280
planning vs. execution, 49, 212
Plossl, George, 15
PLT. *See* purchasing lead time.
portfolio management, 276
position, protect, pull, 52
positioning, 293
positioning factors, 57
primary planning test, 232
prioritized share, 191
priority by due date, 208
priority control, 243
probabilities of simultaneous availability, 71
product deletion, 135
product introduction, 133
product transition, 136
projected buffer status alerts, 218
projected on-hand alert, 218
projected order frequency, 286
purchasing lead time (PLT), 61

qualified actual demand, 150
qualified order spike, 151

ramp-down adjustment, 135
ramp-up adjustment, 133
ramp-up/ramp-down, 136
rapid buffer adjustment, 133
recalculated adjustments, 127
red zone, 97, 113
red zone adjustment, 147
red zone base, 113

red zone safety, 114
relative priority test, 232
relative standard deviation. *See* coefficient of variation.
relevant information, 18, 23, 251
relevant materials, 17, 30, 35, 57, 79, 149, 251, 256, 288
relevant range, 267
reorder-point techniques, 235
replenished override (RO) parts, 96, 125
replenished part, 96
replenishment of stock, 149
requirements explosion, 183
retail, 315
return on assets, 5, 16
RMRP. *See* remanufacturing resource planning (RMRP).
roughly right vs. precisely wrong, 274
routing, 17, 60, 240
run chart, 175

S&OP, 269
safety stock, 184, 233
sales and operations planning. *See* S&OP.
sales order visibility horizon, 58
seasonality adjustment, 140
shock absorber, 96
shop floor control system, 240
signal integrity, 251
significant minimum order quantity, 113
Six Sigma, 52
software compliance, 293
Souder's Law, 291
space planning, 282
spike threshold, 152
stock out (SO), 214
stock out with demand (SOWD), 214
stock out with demand alert, 214
stop explosion, 184
strategic business management, 274
strategic horizon, 268
strategic inventory positioning, 57, 93
strategic relevant range, 268, 274
strategically replenished buffers, 93, 231
supply, 279
supply chains, 293
supply continuity variability, 33, 55, 57, 181, 209, 222, 231, 234
supply generation, 96, 158

Supply order generation, 49
supply variability, 59
supply-chain bullwhip effect, 59, 102
synchronization alerts, 222
system nervousness, 27

tactical horizon, 267
Taguchi function, 10, 95, 255
Taylor, Frederick, 17
technology, 291
Theory of Constraints (TOC), 52
TOG. *See* top of green.
top of green (TOG), 115
top of red (TOR), 115
top of yellow (TOY), 115
TOR. *See* top of red
TOY. *See* top of yellow
Toyota Production System, 17, 50
traditional sales and operations planning, 272
transition, 136

understated, 139

variability, 37, 59, 101
variability category, 104
variability factor, 104
variable rate of demand, 59
variance analysis, 269, 271
velocity, 261

weekly bucket, 28
where used, 72
wholesale network, 85
WIP priority, 247
work-arounds, 7
working capital, 282
work-in-process, 247

yellow zone, 97,113
yellow zone adjustment, 147

zone adjustment factor, 146