Superplant:
Creating a Nimble Manufacturing Enterprise with Adaptive Planning Software

Shaun Snapp

Superplant: Creating a Nimble Manufacturing Enterprise with Adaptive Planning Software

For information about this title or to order other books and/or electronic media, contact the publisher:
SCM Focus Press
PO Box 29502 #9059
Las Vegas, NV 89126-9502
http://www.scmfocus.com/scmfocuspress
(408) 657-0249

ISBN: 978-1-939731-21-0

Printed in the United States of America

Cover and Interior design by: 1106 Design

Contents

CHAPTER 1

Introduction

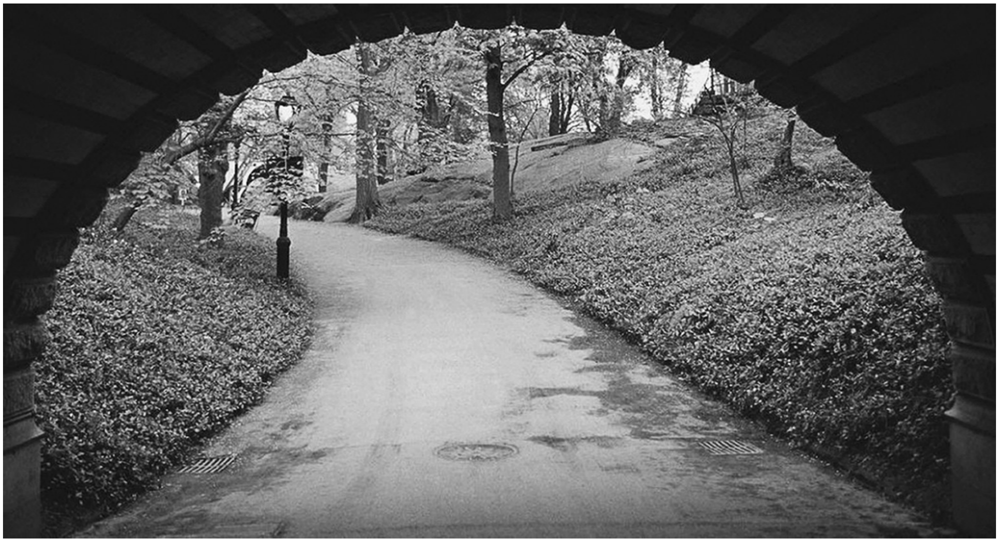

This book addresses several production and supply planning software functionalities that are all related to the location-based adaptability of the supply chain planning application (multi-plant planning and subcontracting, and contract manufacturing planning). This adaptability is a historical weakness of both advanced planning applications as well as

ERP systems. Some of this functionality is rarely found in commercially-available applications, while other functionality is more commonly found but is difficult to implement. For more than a decade and a half, issues of limited location adaptability have plagued many of the projects I have worked on. I was required to deal with these limitations on my projects before I discovered that this functionality had been developed. Therefore, for many years I had a consistent problem but no software solution. However, I can now say confidently that these issues can be addressed with the software from one vendor.

Let's take the three superplant functionalities as an example. In environments where there are dependencies between production activities—such as when a finished good in one factory is fed by semi-finished goods or components (or the components are in a third factory feeding the semi-finished goods plant, which I have in fact seen at several companies)—then the production across the various plants ends up being poorly integrated. For example, the widely-used SAP PP/DS application, which I have implemented, has no way to resolve conflicts between plants that feed each other. Like almost all production planning and scheduling applications, the application cannot even "see" this relationship because its design is such that each plant is seen as ***independent*** during the planning run.[1] While the supply planning system can see the overall supply network, the vast majority of supply planning software can do nothing with respect to multi-plant planning because it cannot create routings that span multiple factories.

Essentially, what has occurred in this scenario is that there is ***one long production line***, which is divided between different plants. These plants could be located close to one another, or as in the example of one of my clients, on different sides of the globe. In an environment of increasingly globalized manufacturing, a very long production line that spans the globe is more common than ever.

1 SAP SNP (Supply Network Planning) can see the overall network and works with SAP PP/DS, but as will be explained in Chapter 4: "Single Versus Multi-pass Planning," sequential supply and production planning runs will not provide multi-plant planning functionality.

There are tricks that can be performed in standard applications that lack multi-plant planning functionality (I chronicle them in my book *Multi Method Supply Planning in SAP APO*) to make the planning systems with a traditional design do the best possible job of integrating these processes. Fundamentally, most supply and production planning software in the market was never designed to meet these requirements; most are designed with a purely sequential approach between supply planning and production planning. The supply plan is created first, as is shown in the graphic below:

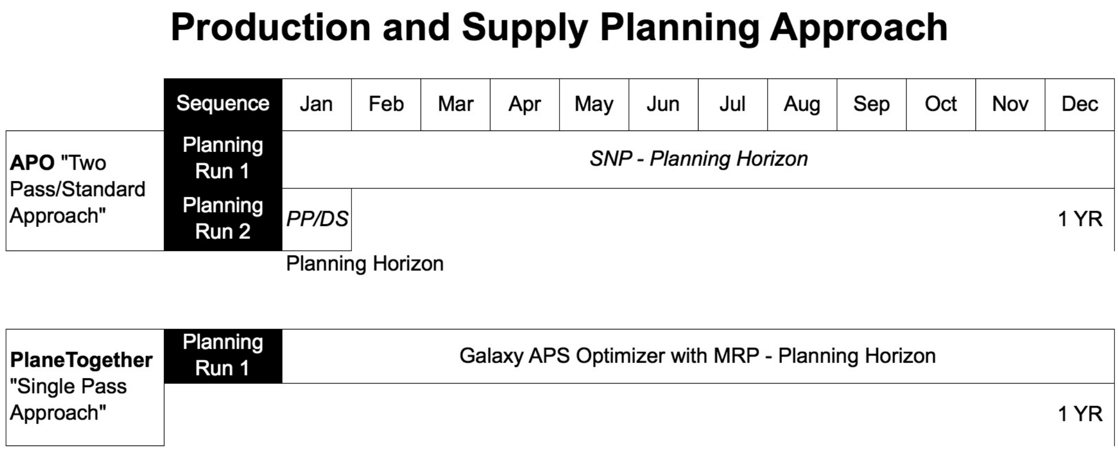

Few vendors have designed their products to meet a requirement like multi-plant planning, or to completely integrate production and supply planning runs.

This book attempts to clearly define specific advanced adaptive functionality: multi-plant planning, subcontracting and multi-source planning. Multi-plant planning is at the intersection of production and supply planning. It is software functionality that has been known as “multi-plant” or less frequently as “multi-site” functionality. However, the term “multi-plant,” while well-known by some of those knowledgeable in production planning software, has never really struck a chord outside of that relatively small group of specialists. It has also not been adopted by many software vendors. In my informal queries on projects, I have found that the vast majority of those who either manage production planning implementations or work with production planning software do not know what the term means. Even for those who have read up on the topic, the term multi-plant has been diluted by a number of software vendors. The most prominent diluter of the term has been SAP, which has adopted the term "multi-plant" to describe the functionality that is interplant supply planning. It's a strange adoption/corruption of the term because all supply planning software can source between supply network locations. Multi-plant planning, on the other hand, means that the software can look across plants, compare alternatives, and choose among alternative routings in order to make the best decision under a particular scenario of where to produce. The stock transfers that are the outcome of this process are a result of the ability to perform a comparison across two or more alternative routings. In contrast, single-plant planning (which is by far the most common software design used by production planning software) cannot do this, and can only choose among alternative routings ***within*** each factory.

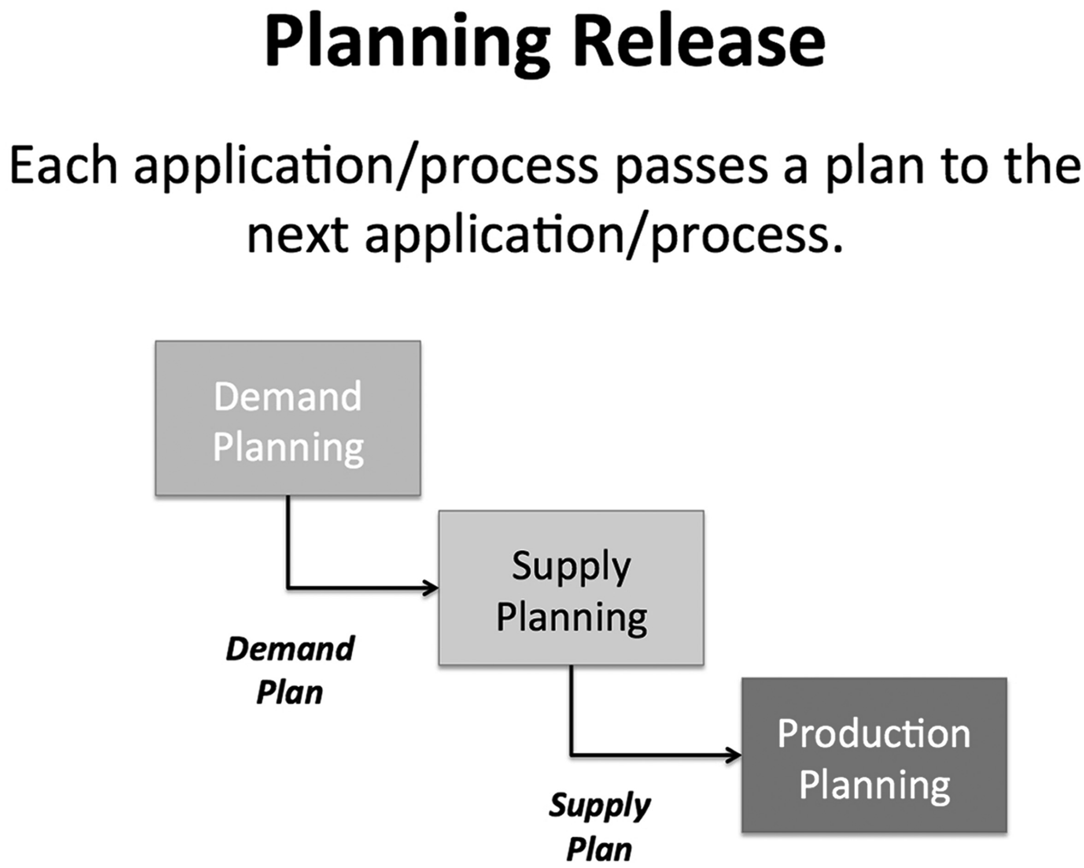

The sequential and "supply planning first" approach means that supply planning is performed prior to production planning and controls the vast majority of the planning horizon. In some applications, supply planning can take into account production resources; however it cannot see the same level of detail as the production planning applications. Therefore the supply planning application creates an oversimplified—or what I call an initial—production plan.

SAP is by far the most popular website returned when one types "*multi-plant planning*" into Google. In fact, when I tried this search, SAP's website(s) were the top six search results as shown in the screenshot on the following page:

Multi-Plant (Site) **Planning**

help.sap.com/saphelp_40b/helpdata/en/f4/.../content.htm

Multi-Plant (Site) Planning. Using **multi-plant planning**, you can carry out material requirements planning for various plants centrally. This planning procedure **...**

Multi-Plant/Site **Planning** (SAP Library - Material Requirements **...**

help.sap.com/saphelp_46c/helpdata/en/f4/.../content.htm

Using **multi-plant planning**, you can carry out material requirements planning for various plants centrally. This planning procedure facilitates the production of a **...**

Shaun Snapp shared this in Gmail – Limited

Multi-Plant Planning with Stock Transfer

help.sap.com/saphelp_40b/helpdata/en/7d/.../content.htm

Multi-Plant Planning with Stock Transfer. Within the stock transfer procedure, goods are produced and delivered within a company. The plant which is to receive **...**

Shaun Snapp shared this in Gmail – Limited

Multi-Plant Planning (SAP Library - Planning Table (PP-REM))

help.sap.com/saphelp_erp2004/helpdata/en/70/.../content.htm

Multi-Plant Planning Locate the document in its SAP Library structure. Purpose. You use this process if you want to plan and then produce a finished product in **...**

Multi-Plant (Site) Total **Planning**

help.sap.com/saphelp_45b/helpdata/en/d7/.../content.htm

Multi-Plant (Site) Total **Planning**. Use. To avoid having to **plan** each plant individually, you can **plan** as many plants as you want in one total **planning** run.

Shaun Snapp shared this in Gmail – Limited

Below are links to SAP Help that describe multi-plant functionality, but which have little to do with actual multi-plant functionality.

http://help.sap.com/saphelp_45b/helpdata/en/d7/5c9366f47811d1a6ba0000e83235d4/content.htm

http://help.sap.com/saphelp_40b/helpdata/en/7d/c276fc454011d182b40000e829fbfe/content.htm

http://help.sap.com/saphelp_46c/helpdata/en/f4/7d2d4d44af11d182b40000e829fbfe/content.htm

Books and Other Publications on Superplant Functionalities

"Superplant" is a term which I created, so one would not expect other books to use this terminology. Superplant is based upon three functionalities: multi-plant planning (primarily covered in academic publications), multi-sourcing (primarily covered in some academic publications and in vendor documentation), and subcontracting (covered—from a software perspective—primarily in vendor documentation).

As with all my books, I performed a comprehensive literature review before I began writing. One of my favorite quotations about research is from the highly-respected RAND Corporation, a "think tank" based in sunny Santa Monica, CA. They are located not far from where I grew up. On my lost surfing weekends during high school, I used to walk right by their offices with my friend—at that time having no idea of the institution's historical significance. RAND's *Standards for High Quality Research and Analysis* publication makes the following statement about how its research references other work.

> *"A high-quality study cannot be done in intellectual isolation: It necessarily builds on and contributes to a body of research and analysis. The relationships between a given study and its predecessors should be rich and explicit. The study team's understanding of past research should be evident in many aspects of its work, from the way in which the problem is formulated and approached to the discussion of the findings and their implications. The team should*

> *take particular care to explain the ways in which its study agrees, disagrees, or otherwise differs importantly from previous studies. Failure to demonstrate an understanding of previous research lowers the perceived quality of a study, despite any other good characteristics it may possess."*

The few books that cover these areas of functionality spend only a few pages on them. And providing more coverage was something I very much looked forward to doing with this book; these topics are complex and require the significant coverage I was able to allocate to them.

I have first-hand experience with some of the vendors that have published the most material on both subcontracting and multi-sourcing—and I disagree with the design approaches most software vendors have taken. As I've seen their products fail quite badly on projects, I will use their poor functionality only as examples of what not to do. It just so happens that the examples of poorly-designed functionality will include one of the best-known software vendors in enterprise software.

The story of multi-plant planning functionality is a bit more straightforward, as I am not aware of more than one vendor that can implement multi-plant functionality—and interestingly, their application was designed to perform multi-plant planning right from the beginning.

The Origin of the Term "Superplant"

The origin of the term "superplant" is easy to explain, because I coined the term. The term actually came to me as I was touring factories outside of the US and discussing the requirements of a previous client who needed their software to manage interdependent factories. During the factory tour I remarked to one of the SAP APO resources that:

> *"It's as if the company's factories are simply individual workstations in a global supply network……like one giant superplant."*

A few other members of the project adopted the term, and I went about developing articles on SCM Focus explaining the concept.

http://www.scmfocus.com/sapplanning/2012/07/26/the-superplant-concept/

http://www.scmfocus.com/productionplanningandscheduling/2013/04/22/multi-plant-superplant-planning-definition/

The Ability to Plan a Single Virtual Plant (for at least some products)

Multi-plant planning is a radical concept because its functionality disrupts the traditional division of responsibilities between the supply planning system and the production planning system. For software to meet the criteria of being multi-plant capable, it must be able to do the following:

1. It must have the ability to create alternate routings that span locations. Multi-plant planning has the ability to treat all resources and all plants as if they are part of one enormous "superplant." This functionality is activated selectively only when it meets the business requirement. Interestingly, with the vendor showcased in this book, resources must actively be assigned to a plant; they are not naturally created within a plant, which is the traditional design approach.
2. It must have intelligence that can allow it to make decisions among alternative routings. Hypothetically it's possible to have some other method, but it is difficult to see this method being anything but an optimizer.

My introduction to multi-plant planning was on a project for a multinational company which produced various components that went into their finished good in different factories separated by more than 4,500 miles. At this time I had never heard of the term multi-plant planning, and was not aware of the previous academic research in this area. My first objective on this project was to simply do what I could in SAP APO (SAP's advanced planning application suite) to meet this multi-plant planning requirement. APO was the tool my client had selected, so that is what I needed to get to work properly; we were going to use a combination of SNP and PP/DS to meet the requirement. Some of the information regarding how this was done was published in the book *Multi Method Supply Planning in SAP APO*. (This book actually addresses several topics. One topic is how to use multiple methods effectively to plan a supply network, but another component was designed to specifically address the needs of a superplant.) We had to do quite a bit of work to make it function, but I was never actually satisfied that the approach we selected—which was based upon adjustments to the supply planning system which customized how product locations were treated by the system—was the best approach for meeting the requirement. (However, I do believe it is the best possible option for those who have already committed to SAP APO.)

After completing this project I contemplated whether any other software vendors had the actual functionality to address the multi-plant planning requirement. Luckily, due to my work either writing articles with best-of-breed vendors or working with them in other capacities, I own a number of their manuals. I decided to read through PlanetTogether's manuals for their application Galaxy APS and checked more on the topic, and I realized that the requirements that I had just made a great deal of effort to test in software with an older software design were covered by ***standard functionality*** in PlanetTogther. Interestingly however, few of the companies that used PlanetTogether actually leveraged this functionality.

The topic of the superplant was supposed to be part of my book *Constrained Supply and Production Planning in SAP APO*. However after forwarding sections of that book that dealt with PlanetTogether to that vendor, the concept seemed to dovetail with their efforts to raise awareness among companies about performing planning in this way. Therefore, I decided to remove the chapter on the superplant from *Constrained Supply and Production Planning in SAP APO* and to turn the topic into its own book. Doing so allowed me to provide much more detail on topics that are of course the chapters of this book. It also allowed for a distinct book in an area where there were none. Furthermore, up until this book, all of my books have been focused primarily on how to set up software to meet business requirements. This is the first book that can be read by executives as well as those interested in the technical aspects of modeling and planning a superplant.

Future of the Term "Superplant"

If the term "superplant" becomes popular, it will no doubt be diluted because software marketers never rest and the competition for software sales is so great. Clear and accurate terminology is to the advantage of those vendors who already have the functionality in question (and the consulting companies aligned with these vendors), but it's disadvantageous to all the other vendors. Those vendors that do not have the functionality (and the functionality is not easy to add to an existing system) will of course look to confuse their potential clients on the topic. Confusing the terminology is one strategy, but not the only strategy that can be employed when functionality is missing. Another strategy is to say that the functionality is not important. This was a major strategy employed by SAP in the late 1990s when companies were becoming interested in advanced planning. At the time SAP had ERP systems to sell, but no advanced planning software. Therefore they told clients that advanced planning was not very important. A third strategy for a vendor to use to counteract a disadvantage in functionality is to propose that, while the functionality may be beneficial, the time is not right for the functionality or for the company to deploy the functionality because there is so much else to do—the company has "*bigger fish to fry*."

I should note that the latter two strategies can both be one hundred percent accurate. Advanced functionality can be enticing, but of less practical importance in the real world than it appears to be conceptually. Secondly, it is also true that a company may find functionality enticing that it is simply not ready to deploy. I have made both of these arguments on numerous occasions myself. However, at the time it was my honest opinion. The difference with the use of these strategies by software vendors who lack functionality is that the opinion ***is developed to minimize a competitive weakness***, and may not be the legitimate opinion of the person proposing it.

At the time of this book's publication, the overwhelming majority of companies that have superplant requirements are meeting them ***without software that can perform superplant planning***. Some of these companies are ready for superplant implementations and some of them are not. Furthermore, companies that do have superplant requirements are on a continuum between those companies that require minor superplant functionality and those companies with high superplant requirements.

The Applications Showcased in This Book

This book will describe the superplant concept, which is an overall adaptive supply and production planning business requirement that is met by very few supply chain applications.

Two different applications are explained in this book. One is SAP APO and the second is PlantTogether. SAP APO is a broad advanced planning suite that includes, depending upon how you count, around ten modules (although the bulk of its implementations are in just four modules). The relevant modules for supply planning and production planning are called Supply Network Planning (SNP) and Production Planning and Detailed Scheduling (PP/DS) respectively. As with all of the APO modules, SNP and PP/DS work in conjunction with each other. While there is some overlap (which will be explained), these are traditional sequential and supply planning first applications. Their design goes back to the mid 1990s, and both products are patterned off of products made by i2 Technologies, one of my former employers. SAP APO is used as an example of the standard non-multi-plant production planning system and for the purposes of explaining how the vast majority of supply and production planning systems work. SAP APO's subcontracting functionality is very difficult to use, as is its multi-sourcing functionality. SAP APO is not alone in this regard, as subcontracting planning and multi-sourcing are areas that are commonly problematic on advanced planning projects. This is due to the fact that few advanced planning applications have been designed from the beginning to treat locations flexibly as if they are ***either*** internal ***or*** external to the enterprise which implements the software. Therefore these planning systems run into problems when the locations are pseudo internal or pseudo external depending upon how one wants to look at it. Unfortunately for many implementing companies, this is increasingly where requirements are going, meaning that many supply and production planning systems are out of date with these requirements. This mirrors the problems that have plagued software vendors with respect to collaboration requirements. As with the treatment of locations, the treatment of collaboration in the design of most supply chain planning applications has been poorly conceived.[2] Thus, the problems SAP APO has performing both these functions is a useful explanation as to the problems that many other software vendors share.

A natural question might be: If SAP APO lacks multi-plant functionality, why is it showcased in this book? The reason is simply that without understanding

[2] This topic is related to collaboration design in software, which is barely discussed. It is covered in the following article: http://www.scmfocus.com/supplychaincollaboration/2013/07/why-must-specialized-supply-chain-collaboration-applications-exist/

the traditional supply planning and production planning design, it is difficult to understand how a superplant-compliant application differs. I anticipate that very few people who read this book currently work with an adaptive application like PlanetTogether's Galaxy APS. Therefore, SAP APO provides me with the ***perfect counterweight*** to Galaxy APS.

I have worked in advanced planning since 1997. I have implemented SAP APO since 2003 and have written five books on APO. I have worked with PlanetTogether since roughly 2010, and have written this book as well as *Process Industry Planning Software: Manufacturing Processes and Software,* which is another book that covers PlanetTogether but focuses on an area greatly underserved by production planning and scheduling software. Therefore, I have a long-term and practical exposure to both applications.

PlantTogether As *the* Superplant Application

The second application to be showcased in this book is called PlanetTogether (PT). I selected PT based upon the fact that it is one of the few production planning applications that can presently perform multi-plant planning. In addition, PT has subcontracting/contract manufacturing functionality that is both usable and straightforward to configure.[3] I became familiar with PT several years ago—both the company and the software—after writing a number of articles on the software. These articles are at the following link:

http://www.scmfocus.com/productionplanningandscheduling/

However, even though PT is capable of supporting the superplant requirement, it should be understood that as of the publication of this book, multi-plant planning is new and usually ***PT is not implemented using the multi-plant functionality. It is sometimes implemented with either its subcontracting or multi-source planning functionality activated***. SCM Focus views this as a big missed opportunity for many companies. One of the major reasons for companies not implementing this functionality is simply the customers' lack of knowledge of these functionalities, something I hope this book is able to address.

[3] Galaxy APS does not, however, perform multi-source planning.

A company may implement Galaxy APS because it is attracted to Galaxy APS's adjustable duration-based optimizer, or how Galaxy APS handles process industry requirements, or because of its flexibility with regards to master data maintenance, or because of its user interface. As I will show further on in this book, many companies that have implemented Galaxy APS do not have superplant requirements and so they do not investigate or have a particular interest in whether or not Galaxy APS contains this functionality. Therefore, I do not want to leave the reader with the impression that superplant planning is the primary function of Galaxy APS, or the main reason it is chosen by companies for implementation, because that is not the case.

Being able to compare and contrast two systems that do supply and production planning and yet work in such different ways is, I think, a main contribution of this book. If I had stopped at showcasing only how multi-plant planning works, I believe the audience for the book would have been much smaller.

Galaxy APS User Interface Introduction

On the next page is the main user interface to Galaxy APS.

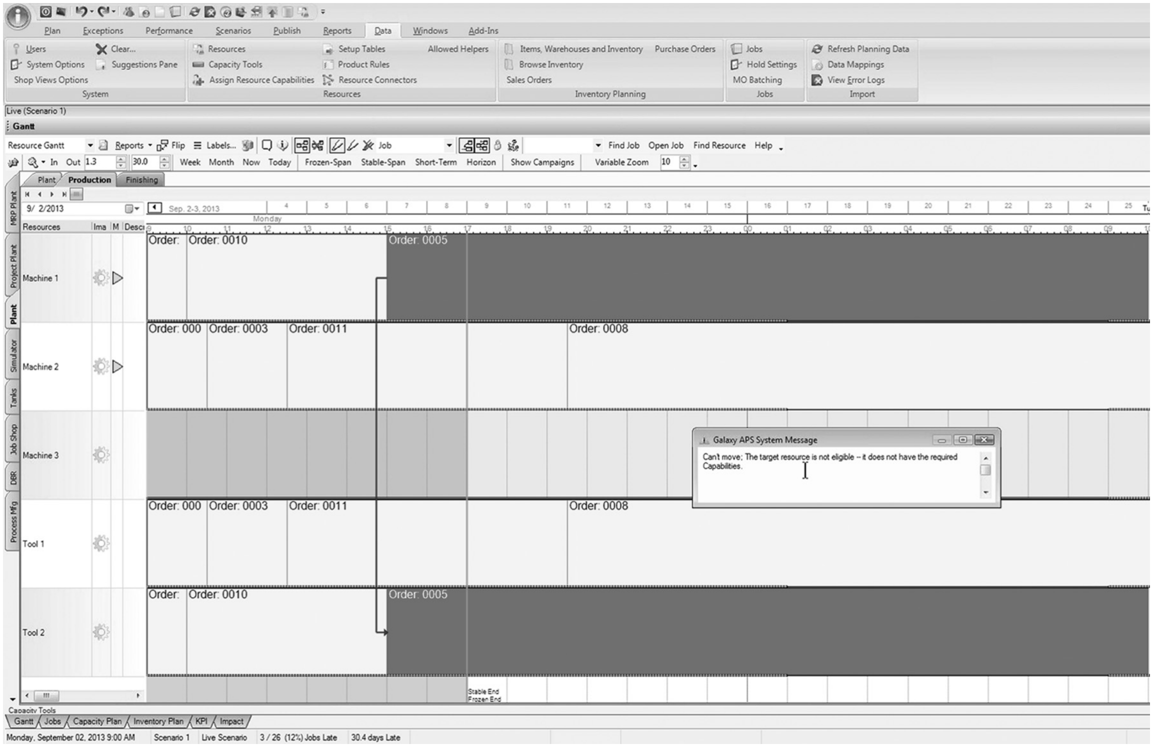

"The system will only schedule work on eligible Resources, and Drag-and-Drop is only allowed to eligible Resources (or a rejection warning is shown),"[4]

Publications on Multi-plant Planning and the History of Method Development

It was quite interesting to review the literature on multi-plant planning necessary to write this book. Clearly, multi-plant planning is an emerging area because most of the substantial documentation on the topic is in the form of academic papers. I have studied the history of a wide variety of supply chain planning methods for other books. From this research I have found that, in most cases, the academic documentation precedes the software development. (Some might think

---

[4] PlanetTogether APS Product Training.

that this would always be the case, but in fact it is not so—as I will explain in just a few sentences.) Furthermore, it is typically the case that the more complex the method, the greater the time lag between the appearance of the method in academic publications and the implementation of the method in software. For instance, inventory optimization and multi-echelon planning—one of the major methods of supply planning—first appeared in an academic paper in 1958, but was not implemented in commercially available software until the late 1990s. There are several reasons for this time lag. Often academics develop mathematics for a method that is so far beyond the hardware capabilities of the time, that the wait for a commercial application will be a long one.[5] Academics generally strive for pushing the envelope. For instance, in the examples of inventory optimization and multi-echelon planning, testing of the method on the hardware of the time could only be performed for a limited supply network and for a single product. Of all the methods I have evaluated over time, only material requirements planning (MRP) and, to a lesser degree, distribution requirements planning (DRP) break this rule of the transition from academic research to commercialization. In both cases the method was developed first in software and in industry, and then was documented in academic or academic-type papers. This history is shown in the graphic on the next page:

[5] Multi-echelon planning research was funded by the Air Force through the RAND Corporation. One of the interesting questions is why the Air Force was interested in funding research that was so far beyond the hardware capabilities of the time that it would be decades before it could be used. This research funding by the Air Force reinforces points made by Noam Chomsky, that a misunderstood feature of US military spending is that it is a clandestine way of directing research money to areas of technology that are in fact part of an industrial policy. Curiously, the Air Force was willing to fund the development of the mathematics for multi-echelon planning, but not for the development of software to actually run a multi-echelon software program—another peculiarity of how it distributed research funding.

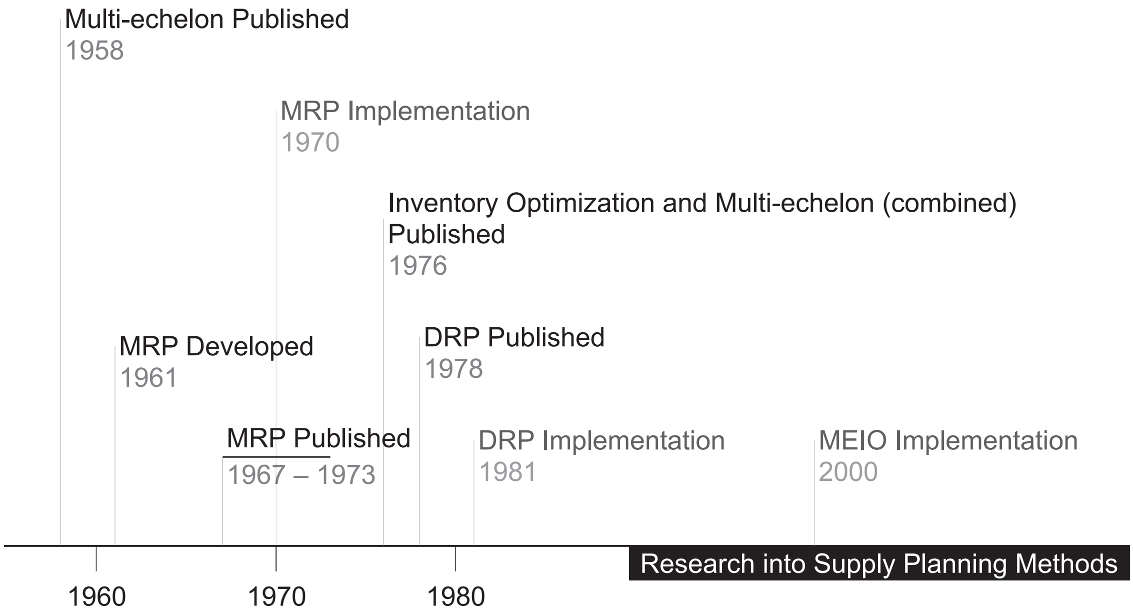

Outside of academic publications and several interesting examples of nonacademic publications that I discuss in this book, the other most prominent publishers on this topic are SAP (which I have already explained describes something that is not multi-plant planning yet uses the term “multi-plant planning”) and PT. PT has begun to write about the topic on its website, and they cover multi-plant planning in their training manuals. (This book was written with input from PT.)

Some of the websites that describe multi-plant planning are not discussing “true” multi-plant planning, but some other use of the term. Therefore, even though they use the terminology, they don’t actually cover the topic according to the officially accepted use of the term. Naturally, I do not count them as contributors in this space.

Generally, what I came away with from my research was that the term “multi-plant” was used lightly, and the concept behind it only weakly explained and socialized. I faced a similar issue—although to a lesser degree—with my book on Inventory Optimization and Multi-echelon Planning (MEIO). This was a method, functionality, and mathematics for which a number of vendors had already developed products, and for which quite a bit of academic literature had been written. However, before my book and sub-site on the topic at SCM Focus, there was no real comprehensive explanation for business people as to how MEIO worked outside of the software manuals of the individual vendors.

At the time of this book's publication, multi-plant planning is at a far earlier stage of development than anything I have written about in the past. For an author like myself, who enjoys bringing the latest innovations in supply chain software from the best vendors, this is exactly the kind of topic—and the stage of development of the topic—that I look for.

Of all the publications I read on the multi-plant topic, it is no coincidence that the one I found more interesting is one of the important publications on multi-plant planning: *An Analytical Framework for Multi-Site Supply Chain Planning Problems*. It has some interesting quotations and frameworks. Several of the ones that I consider the most impactful I have included and commented upon below:

> *"In the past years the multi-site production planning problems have attracted many researchers' attention, but most of the researches put emphasis on the methodology to solve the multi-site planning and scheduling problem. Few of those researches are to analyze the essence and definition of the multi-site production planning problem. The analytical framework of the multi-site production planning problem is proposed in this paper."*

This is an important point because understanding the "analytical framework" essentially means understanding how multi-site or multi-plant planning actually works. The paper also discusses the "structural framework," which I have listed below, along with a description of how each aspect of the structural framework is covered in this book:

- *Multi-site Conceptual Model:* How multi-site planning works from a design perspective is very much a focus of this book.

- *Product Structure (Bill of Materials), Resources and Routings:* How each one of these master data objects must operate and be set up in order to support multi-plant planning. The bill of materials (BOM) is covered to the degree that one overall finished goods BOM must be broken into modular BOMs in order to produce components of the finished good in separate factories. Resources and routings are covered quite extensively.
- *Production Methods:* The planning method used changes how the production and supply planning model must be set up in order to achieve the desired outcome. This book will discuss cost optimization, duration optimization, and duration optimization with a heuristic-based algorithm.
- *Manufacturing Capability and Characteristics:* This book addresses capabilities that are assigned to resources. Understanding how one capability can be assigned to multiple resources that are in different plants is critical to understanding how superplant functionality works.
- *Production Planning Constraints:* This book will assume that constraint-based planning is used in all cases. The most important concept to understand with respect to the use of superplant planning along with constraint-based planning is that a single bottleneck/drum resource may constrain the entire production process regardless of how many plants the routing is strung through.[6] This is a focus of coverage in this book.
- *Key Performance Indicators (KPIs)*:* Optimizing with KPIs is a big part of what makes PlanetTogether work so effectively. Not only will KPIs improve when multi-plant planning functionality is properly deployed, but the Galaxy APS optimizer can be adjusted through the use of KPI adjustment rules. This will be explained in detail in several places in this book.

Different Multi-plant Planning Scenarios

An Analytical Framework for Multi-Site Supply Chain Planning Problems goes on to say the following about how the definitions of multi-plant are used generally.

> *"In the literatures, many researchers had different definitions for the term: 'multi-site' or 'multi-plant.' The multi-site production planning*

[6] To understand constraint-based planning in a way that effectively interoperates with this book, but with a focus on where there is a separate supply and production planning system, please see the book *Constrained Supply and Production Planning in SAP APO.*

> *problem is mainly the production allocation decisions among multiple plants. (Thierry et al.) …manufacturing process of products may require the usage of many resources located in different production units. Furthermore, some alternative manufacturing routings may exist."*

The paper points out that there are multiple multi-plant models, and so it focuses only on the major multi-site models. The paper then goes through the various major models, which are essentially different scenarios. One scenario may have stock transferred between two factories. Another may have stock coming in from one location, or from multiple locations to two or more factories. However, for our purposes, if the requirement were to synchronize production and supply planning across multiple locations, on components that eventually lead to a combined product—most often a finished good, but also an assembly or subassembly—then this would meet the definition of multi-plant planning. When this is the case, the company is in some shape or form planning a "superplant."

The Opportunity of Multi-plant Planning

There is little doubt of the many companies with multi-plant planning requirements. I have experienced the requirements first hand and in depth at one company, but have found these requirements at other companies as well—it's simply that most are unaware what these requirements are actually called and that there is academic work that describes these requirements. They are also unaware that they are losing efficiency by not being able to account for these requirements. Multi-plant planning is an important stage in the evolution of planning software that is related to subcontract and contract manufacturing planning. All of these planning requirements mean the planning system is indifferent to whether the manufacturing location is owned or not owned by the implementing enterprise. Companies have little in the way of information on multi-plant planning, which should not be surprising in the least. In fact, I know from my consulting experience that many companies have multi-plant requirements but are simply not leveraging the software currently available to manage these requirements. In fact, at most companies the internal discussion about doing so has not even begun. Common reasons as to why this is the case are listed below:

1. Many decision makers in companies with multi-plant planning requirements do not know that the functionality to specifically address these requirements exists.
2. Many companies do not include vendors with multi-plant planning functionality in their software selection initiatives.
3. No ERP vendor makes external planning software that performs multi-plant planning. The company would have to be willing to choose a smaller vendor rather than simply purchasing the ERP vendor's external planning system. This is of course a limiting factor, because companies tend to purchase as much software as they can from one vendor, which incidentally is why the enterprise software sector is so monopolistic in nature and why so many companies have such a poor fit between their business requirements and the applications they have purchased.

Furthermore, as companies have moved to more specialization in their factories (co-locating specific manufacturing in global locations), intercompany transfers have become increasingly common. Concentrating similar types of production in factories globally has been occurring for some time. Things like subcontracting, which at first glance would seem to reduce the necessity for multi-location planning, in fact increase the necessity for multi-location planning. Even in instances where third parties are involved—such as with subcontracting—the primary company or OEM ***often wants to plan the activities,*** even if they do not perform the actual execution. In fact, we now have the common scenario where planning factories—or at least partially planning factories that ***are not owned by the company performing the planning—are a common requirement***.

As most enterprise software is not designed to accommodate these requirements, a great deal of effort is needed on the part of companies to both implement and maintain the software. The retort from many vendors might be, "*but we offer supplier collaboration and subcontracting*"—which is true. However, it is also true that these tend to be tricky implementations, and in some cases, such as with supplier collaboration, there are in fact few success stories.

http://www.scmfocus.com/supplychaincollaboration/2010/06/where-are-the-supply-chain-collaboration-success-stories/

Therefore, the business requirements and opportunities for multi-plant planning will only continue to grow. In fact, this is a constant trend in both supply and production planning. Supply chains are simply becoming more global with more interactions of all types between plants. Perhaps more importantly, in the case of locations that are not part of the company's supply network, supply chains are also becoming more ***ambiguous as to who is actually performing the planning***. The business needs within these companies have changed more quickly than the software that supports them.

Subcontracting

Along with multi-plant planning, subcontracting has greatly increased as a planning need within companies. Subcontracting is another form of production where there is ambiguity between the external plant and the internal locations. Another related concept is contract manufacturing, where the product is produced completely by the contract manufacturer but planning responsibility is shared. Some supply chain planning applications can plan subcontracting; however, the functionality I will describe in this book can make comparisons between alternative production that is either internal or external to the company. In some circumstances, production may be outsourced; in other cases it may be planned to be produced internally. Oftentimes inflexible planning systems mean that companies are forced to make these types of decisions "strategically." However in the software described in this book, the alternatives can be set up in the model, and the application can switch between internal and outsourced manufacturing as the situation changes. I cover contract manufacturing along with subcontracting because there are a number of similarities and the functionality that is explained in this book can meet either requirement.

Multi-source Planning

Multi-source planning (or multi-sourcing)—the ability to have the system flexibly choose from external sources of supply—has been a consistent requirement at many companies. However, many companies have also had a problem getting multi-source planning to work the way they want it to work, and so it is not implemented as commonly because of issues with functionality robustness. I have spent years dealing with these issues on projects—projects that have required enormous amounts of configuration and testing because essentially the software I was using was not designed to work the way that the company needed it to work.

I am quite enthusiastic about the business benefits that can be derived from a fully superplant-capable application. But from a purely selfish perspective, I would simply like to be able to implement this software, because the constant conversations about the inability of various applications to meet these needs (and the workarounds that they entail) have become a headache to me as much as to my clients. Therefore, a motivation for writing this book is to explain to the widest group possible that it is not necessary to continue to bang our collective heads against the wall and perform various magic tricks (which always seem to impact some other functionality) on projects in order to meet these requirements. All of the adaptive functionalities mentioned can be implemented both effectively and cost effectively with PT's Galaxy APS product.

The Use of Screen Shots in the Book

I consult in some popular and well-known applications, and I've found that companies have often been given the wrong impression of an application's capabilities. As part of my consulting work, I am required to present the results of testing and research about various applications. The research may show that a well-known application is not able to perform some functionality well enough to be used by a company, and point to a lesser-known application where this functionality is easily performed. Because I am routinely in this situation, I am asked to provide evidence of the testing results within applications, and screen shots provide this necessary evidence.

Furthermore, some time ago it became a habit for me to include extensive screen shots in most of my project documentation. A screen shot does not, of course, guarantee that a particular functionality works, but it is the best that can be done in a document format. Everything in this book exists in one application or another, and nothing described in this book is hypothetical.

Identification of Timing Field Definitions

This book is filled with lists. Some of these lists are field definitions. Lists of field definitions will be all *italics*, while lists that are not field definitions will be only *italics* for the term defined, while the definition that follows is in normal text.

How Writing Bias Is Controlled at SCM Focus and SCM Focus Press

Bias is a serious problem in the enterprise software field. Large vendors receive uncritical coverage of their products, and large consulting companies recommend the large vendors that have the resources to hire and pay consultants rather than the vendors with the best software for the client's needs.

At SCM Focus, we have yet to financially benefit from a company's decision to buy an application showcased in print, either in a book or on the SCM Focus website. This may change in the future as SCM Focus grows—but we have been writing with a strong viewpoint for years without coming into any conflicts of interest. SCM Focus has the most stringent rules related to controlling bias and restricting commercial influence of any information provider. These "writing rules" are provided in the link below:

http://www.scmfocus.com/writing-rules/

If other information providers followed these rules, we would be able to learn about software without being required to perform our own research and testing for every topic.

Information about enterprise supply chain planning software can be found on the Internet, but this information is primarily promotional or written at such a high level that none of the important details or limitations of the application are exposed; this is true of books as well. When only one enterprise software application is covered in a book, one will find that the application works perfectly; the application operates as expected and there are no problems during the implementation to bring the application live. This is all quite amazing and quite different from my experience of implementing enterprise software. However, it is very difficult to make a living by providing objective information about enterprise supply chain software, especially as it means being critical at some point. I once remarked to a friend that SCM Focus had very little competition in providing untarnished information on this software category, and he said, "Of course, there is no money in it."

The Approach to the Book

By writing this book, I wanted to help people get exactly the information they need without having to read a lengthy volume. The approach to the book is essentially the same as to my previous books, and in writing this book I followed the same principles.

1. **Be direct and concise.** There is very little theory in this book and the math that I cover is simple. While the mathematics behind the optimization methods for supply and production planning is involved, there are plenty of books, which cover this topic. This book is focused on software, and for most users and implementers of the software the most important thing is to understand conceptually what the software is doing.
2. **Based on project experience.** Nothing in the book is hypothetical; I have worked with it or tested it on an actual project. My project experience has led to my understanding a number of things that are not covered in typical supply planning books. In this book, I pass on this understanding to you.
3. **Saturate the book with graphics.** Roughly two-thirds of a human's sensory input is visual, and books that do not use graphics—especially educational and training books such as this one—can fall short of their purpose. Graphics have also been used consistently and extensively on the SCM Focus website.

Before writing this book, I spent some time reviewing what has already been published on the subject. This book is different from other books in terms of its intended audience and its scope. It is directed toward people that have either worked with ERP or know what it is; I am assuming that the reader has a basic knowledge level in this area.

The SCM Focus Site

As I am also the author of the SCM Focus site, http://www.scmfocus.com, the site and the book share a number of concepts and graphics. Furthermore, this book contains many links to articles on the site, which provide more detail on specific subjects. This book provides an explanation of how supply and production planning software works and aims to continue to be a reference after its initial reading. However, if your interest in supply planning software continues to grow, the SCM Focus site is a good resource to which articles are continually added.

The SCM site dedicated specifically to supply planning is
http://www.scmfocus.com/ supplyplanning.

The SCM site dedicated specifically to production planning is
http://www.scmfocus.com/productionplanningandscheduling/

The site dedicated to SAP planning is
http://www.scmfocus.com/sapplanning

Intended Audience

This book was written for those with an interest in leveraging leading approaches in the supply network for planning improvement. The specific audience would range from executive decision makers to software implementers to supply and production planners. If you have any questions or comments on the book, please e-mail me at shaunsnapp@scmfocus.com.

Abbreviations

A listing of all abbreviations used throughout the book is provided at the end of the book.

Glossary

The following glossary of terms will be helpful before reading Chapter 3: "Multi-plant Planning" and some have been taken from PT literature. It can be best to breeze through these terms when first reading this chapter, but then return to them when the terms are first used.

1. *Plants*: Correspond to separate physical factories (I use the terms "plant" and "factory" interchangeably).

2. *Resources*: Usually correspond to machines, people, and tools or subcontractors. They are used temporarily to perform an operation. Materials are not Resources since they are "consumed" rather than used temporarily. Each Operation in a Job has one or more Resource Requirements, each of which specifies the need for another additional Resource in the production process. Each Resource belongs to exactly one Department. A Resource is a capital asset that is available and used to perform work. In this broad definition of a Resource, Galaxy APS considers any machine, tool, die, mold, or even skilled laborer as a resource. The general rule is that if it is a production asset that requires scheduling, it will be defined as a Resource inside Galaxy APS.

3. *Routing*: A combination of resources in sequence. (As will be shown, a major component to superplant functionality is how routings can be configured.)

4. *Alternate Path*: An Alternate Path specifies the precedence relationships between Operations thus indicating the path that is followed through the shop to produce a product. Each Manufacturing Order specifies one or more Alternate Paths and exactly one Default Path. When the Manufacturing Order is scheduled, the Default Path is used. One Operation can have multiple Successors and multiple Predecessors.

5. *Capabilities*: This is the work that a resource can do. A Capability is a specification of a skill or function that can be performed by a Resource. For example, a Capability might be "cutting" or "inspection." Each Resource has one or more Capabilities that indicate the types of work that it can perform. In addition, each Job Operation Resource Requirement has one ore more required Capabilities indicating the Capabilities that a Resource must have to be considered eligible to perform the Operation. If an Operation, for example, requires a CNC machine and a CNC operator, then you would create two Capabilities such as: CNC and CNC-Operator. Then the Operation would have two Resource Requirements—one specifying CNC and one specifying CNC-Operator.

6. *Manufacturing Order*: What is created in order to plan and schedule actual production. A Manufacturing Order is a request to manufacture a specific quantity of a specific Product. Each Job contains one or more Manufacturing Orders. Each Manufacturing Order contains one or more Operations and one or more Alternate Paths that describe the sequence in which Operations should be performed.
7. *Operation*: This is the work performed on a resource. An Operation is a definition of a single processing step in a Manufacturing Order. The Operation specifies values such as the Cycle Time, Set-up Time, expected Yield, Resource Requirements, Material Requirements, etc. Each Manufacturing Order contains one or more Operations. Operations can be "connected" to establish Predecessor/Successor relationships using the Alternate Path of the Manufacturing Order.[7]
8. *Drum*: Used in the in the Drum-Buffer-Rope metaphor of the Theory of Constraints where the drum is the constraint, the buffer is the material release duration and the rope is when the material is released.[8]

Corrections

Corrections and updates, as well as reader comments, can be viewed in the comment section of this book's web page. If you have comments or questions, please add them to the following link:

http://www.scmfocus.com/scmfocuspress/production-books/the-superplant-concept/

[7] From the PlanetTogether Glossary.
[8] Galaxy APS can use Drum-Buffer-Rope scheduling if configured to do so.

CHAPTER 2

Understanding a Superplant Conceptually

There are several important differences between a standard production design and a superplant. A good way to begin this chapter is by describing these differences.

1. In a superplant, the bill of material (BOM) is distributed in multiple locations.
2. In a superplant, the production can be placed in multiple locations.
3. Synchronizing the continuous material flow among factories is critical to maintaining production efficiencies. The triggering of stock transports may change depending upon plant proximity. In plants that are within close proximity, the supply planning system may not need to be involved.

    http://www.scmfocus.com/sapplanning/2012/07/24/synchronizing-integrated-factories-with-stock-transfers/

4. In a macro sense, each factory can be thought of as a work center, and the flow through the supply network—for internal locations at least—can be considered a routing.

5. An important consideration with the superplant design can be, depending upon requirements, the prioritization of internal over external demand. This means using both constraint-based planning in addition to being able to automatically prioritize the by-demand type. This is addressed in detail in Appendix C: "Prioritizing Internal Demand for Subcomponents over External Demand."
6. A superplant not only has the ability to evaluate whether to place demand onto internal resources, but also to leverage subcontract and contract manufacturing partners in a similarly flexible manner based upon circumstances. A superplant can perform subcontracting, but may also be able to produce the same item internally, meaning that internal production is compared against subcontracting production.
7. A superplant can constrain a routing, regardless of how many plants it passes through based upon a bottleneck/drum resource that resides in any one of the plants.
8. A superplant can nimbly evaluate the best location to source a product automatically and as part of the standard planning run. A superplant can flexibly alter its source of supply based upon the setup of as many sources of supply as actually exist, and allowing the software to make the determination between alternatives.
9. The best way of thinking about the output of superplant planning is that the output depends upon the particular circumstance. In some circumstances one alternative is selected and in another circumstance a different alternative is selected. These circumstances change all along the planning horizon meaning that different alternatives are the best to select at that point in time.

Now that we have covered how a superplant differs from the standard approach to supply and production planning, we will jump into how each of the functionalities that make up a superplant (multi-plant planning, subcontract planning and multi-source planning) actually work in detail.

CHAPTER 3

Multi-plant Planning

As discussed, multi-plant planning, sometimes called multi-site planning, is the ability to model and make decisions to schedule production between alternate internal production locations that can produce the same product. In some cases these locations may have identical capabilities that would produce a finished good; however, in other cases it may be a comparative process that would produce a component or subcomponent that becomes ***an input to another production process at a downstream factory***.

Who Requires Multi-plant Planning Functionality?

By definition, companies that have components and subcomponents of final finished goods that are moved between factories have a multi-plant planning requirement. This requirement exists in all manufacturing environments (discrete, repetitive, process batch, process continuous). I once consulted for a repetitive manufacturer that produced wires in some factories, and then sent these wires halfway around the world to become input products to finished goods electrical products.[9] Transporting intermediate product between factories is also common in the petroleum industry. Many of these stock transfers of intermediate product are between facilities owned by the same company, as the oil industry is one of the more vertically integrated of all industries.[10]

A company that does not have multi-plant planning requirements when they start out, will have these requirements as soon as they choose to consolidate one stage of manufacturing to a single location in order to benefit from economies of scale and economies of specialization in that manufacturing process. Producing more similar items at fewer factories will tend to increase production efficiency. (This would be for a laundry list of reasons including the ability to group similar machines together; improve maintainability of many similar machines in one plant; more trained and interchangeable machine operators with similar skills; the ability to purchase and install larger, fewer and more efficient machines; etc.)

[9] Generally speaking, not all components that make up a finished good have an equal level of complexity. It can make quite a lot of sense to manufacture lower-complexity components in lower-cost countries and ship the components to higher-cost countries, which can perform the more complex manufacturing operations. Once such a factory is set up and running, it cannot only send its input components to the factory owned by the same parent company, but can have external customers as well.

[10] Vertically integrated companies tend to be a good fit for multi-plant planning software, although what ***is*** in fact vertically integrated is an interesting discussion. For instance, Wikipedia lists Apple as a vertically integrated firm because they *"…control the processor, the hardware and software. Hardware itself is not typically manufactured by Apple, but is outsourced to contract manufacturers such as Foxconn or Pegatron who manufacture Apple's branded products to their specification."* Therefore, according to Wikipedia, the main definition of "vertical integration" is whether the company controls the process, not—as more generally thought to be the case—whether the company actually performs all of the tasks. As Apple does no manufacturing, they could perform multi-plant planning, but the plants would be external plants.

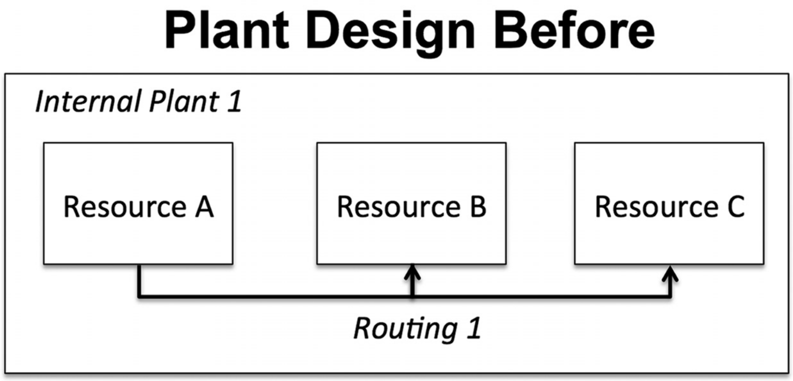

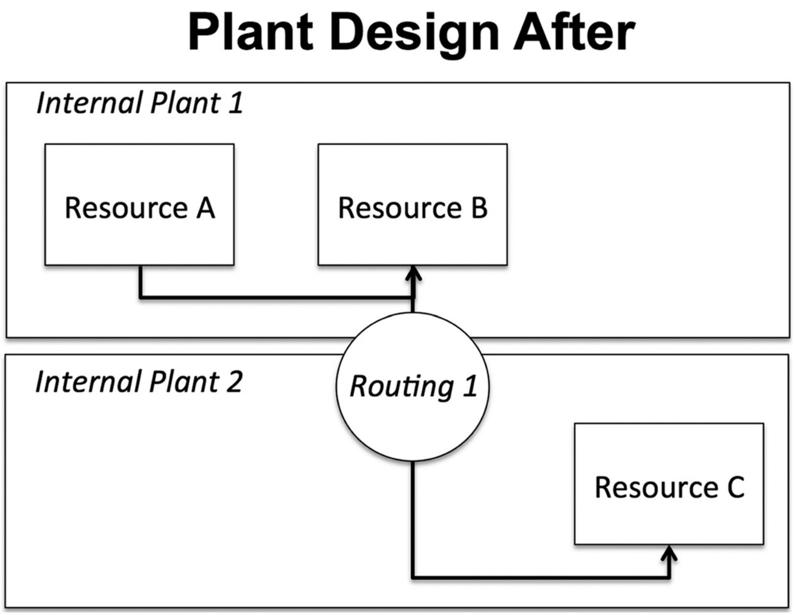

As soon as a routing spans a plant, the company has multi-plant requirements. The more of these relationships a company has, the more it benefits from multi-plant planning functionality.[11]

[11] For the purposes of keeping the explanation easy to follow and not combining too many concepts at once, I have left out the discussion of how the stock transfers that are created by the alternate routings between the factories are managed by the solution (the solution being both the external planning system and the ERP system). This is, however, covered in detail in Chapter 8: "The Superplant and the Integration Between ERP and the External Planning System." I consolidated all stock transfer information here so as to cover stock transfers for all three superplant functionalities in one place.

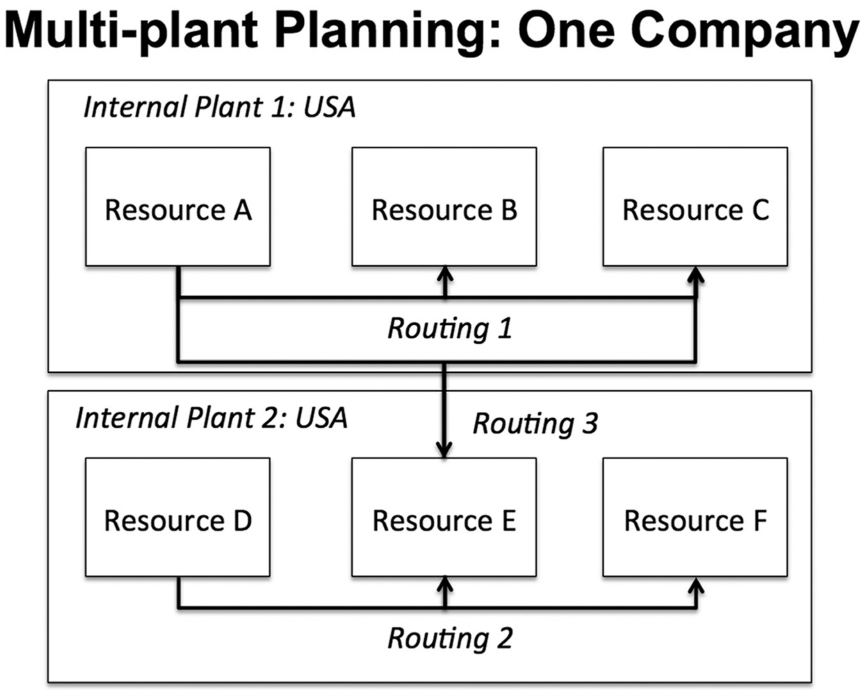

By implementing multi-plant planning software, the company may be able to supply multiple plants. It is able to produce more output with fewer resources because it can receive a higher production utilization from its equipment. Decades of nonmathematical proposals by inventory/lean proponents has blinded companies to the fact that the main emphasis of production is machine efficiency. This can be easily shown using calculations with actual assumptions from factories—these are calculations that lean proponents refuse to perform, relying instead on general statements. Companies that do this receive better production efficiency at the cost of higher planning and coordination costs, as well as higher transportation costs.

There is a tendency to accept current factory configuration as stable. However, throughout the history of manufacturing, factories have gone through many configuration changes. A perfect example of this is during the First Industrial

Revolution, when fossil fuels were first used as a power source over water, wind or even animal power. Also, in the Second Industrial Revolution, steam turbines were moved out of factories and central power plants delivered power to factories over power lines. The power was converted to mechanical movement with electrical motors in the factories. The factories changed enormously before and after both industrial revolutions, not only in their layout but also in their size, specialization, and in almost every other dimension as covered in detail in the article below.

> http://www.scmfocus.com/scmhistory/2013/08/the-electrification-of-production-plants/

Factories are simply the result of the technology, material availability, labor skills and costs, and of the industrial engineering and planning principles of the day. They are subject to change, and indeed will change. Companies that have software capable of multi-plant planning are in a better position to leverage the production efficiencies of manufacturing consolidation. Therefore, software capable of multi-plant planning may be implemented by companies that already have multi-plant configurations, or software capable of multi-plant planning may ***enable a company to move to a multi-plant configuration only because they have no effective way of properly planning and controlling the factory in an alternate configuration***. Loosely translated, when a company has multi-plant planning functionality, not only can it better manage the multi-plant requirements that it currently has, but can actually adjust its factory configuration to better leverage this newfound planning intelligence.

One example of a multi-plant planning scenario is when a routing (which declares the pathway over the individual operations that the manufacturing/production order must pass) must always span factories in order to arrive at a finished product. However, another example that would drive a multi-plant planning requirement is when identical processing resources existing in multiple internal factories can be leveraged under particular circumstances.

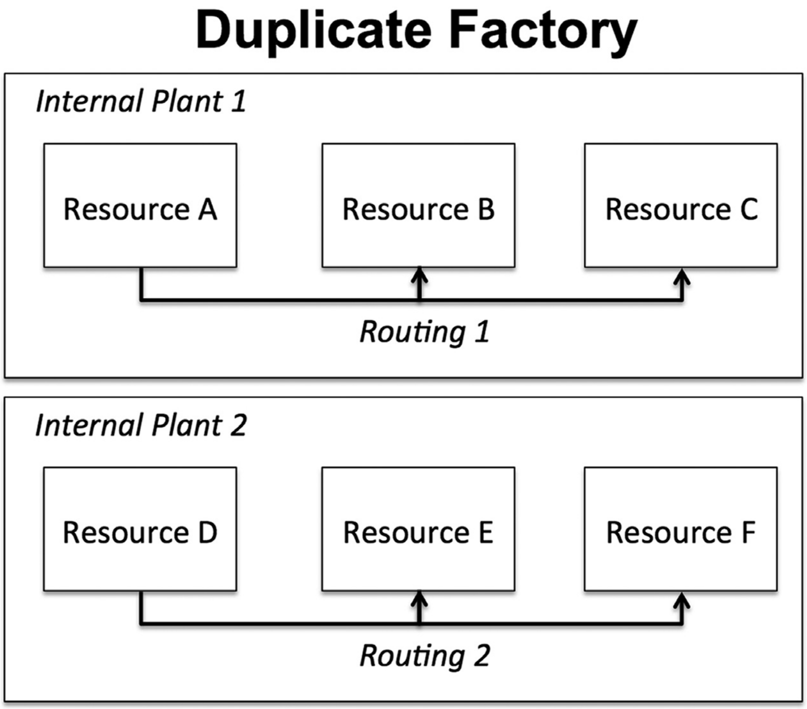

In this scenario there are two duplicate factories that make the same output product. However, when the production rate is particularly high, in one location or the other, the operation that is performed on resource B and E can be "shared." This means making another routing.

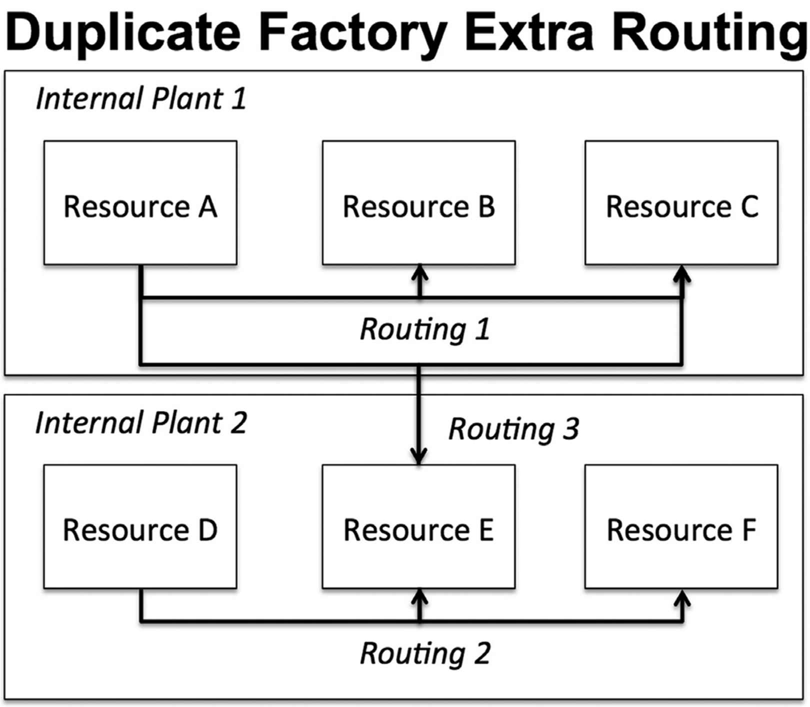

In multi-plant planning, all that is required is to create a new Routing, which allows Plant 1 to call on the capacity of the Resource D in Plant 2. Another Routing could be added which would allow Plant 2 to call on the capacity of Resource E. This may be more efficient than simply satisfying the demand of customers from Plant 2 because considerably more weight is added to the finished product in the final manufacturing step, which is accomplished with Resource C and F.

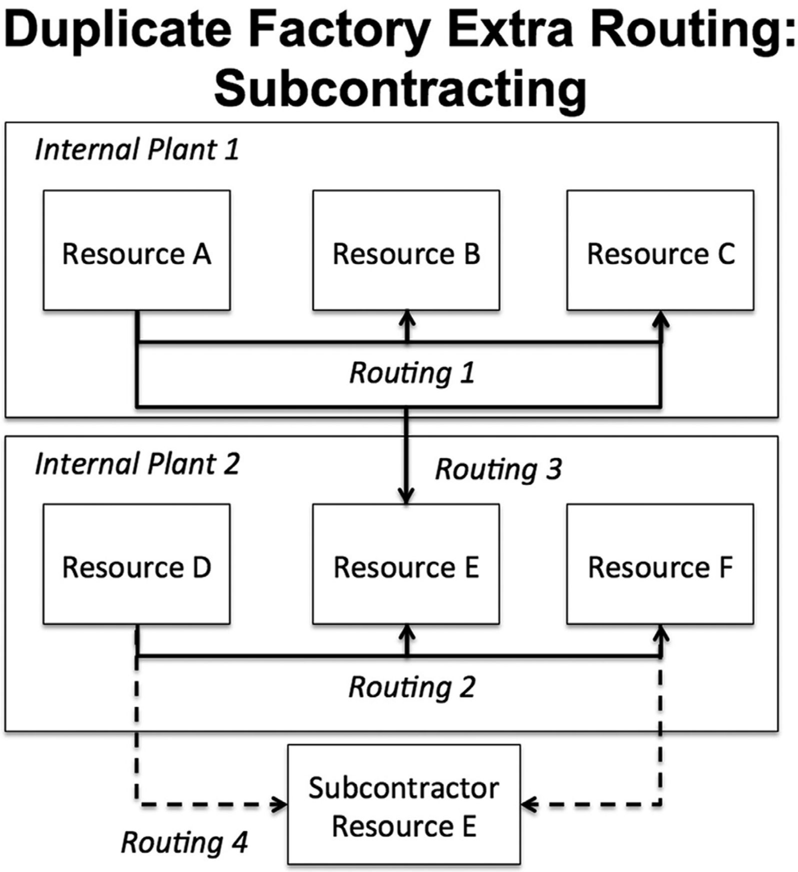

We will get into this topic in Chapter 6: "Subcontracting Planning and Execution." However, with multi-plant planning functionality, we could add another alternate Routing that can reach out to a subcontractor for processing. This would allow Plant 2 to call on the subcontractor for manufacturing capacity. As soon as a company has multi-plant planning functionality, all types of alternative arrangements open up that previously were managed manually.

Understanding the Software Design of Standard Supply and Production Planning Systems

In traditional—and what are by far the most common—supply planning and production planning applications today, if a company were to have one hundred resources spread over ten factories, any resource assigned to one factory is ***unavailable to be combined with other resources in other factories***. Under this approach, the supply planning system controls all logic and flow between the locations in the supply network. The production planning activity and capacity within a plant is limited to just that plant.

This software design was clearly colored by the assumption that production planning occurs in a purely local manner. While we have more advanced methods than MRP at our disposal in production planning (including heuristics, allocation algorithms and optimization), it is interesting that the same location limitations of MRP are used even though these more sophisticated methods have location flexibility. For instance, in SNP, as with many supply planning applications, if location decomposition is used with the optimizer, it will break the overall problem into sub-problems with each sub-problem being the location supply network for just the product which is processed. (For more detail on the topic of decomposition see the following article.)

> http://www.scmfocus.com/sapplanning/2011/10/12/snp-optimizer-sub-problem-division-and-decomposition/[12]

Each one of these locations can be constrained by all of the production resources in all of the plants—however the production, which is planned, is planned independently in each plant. The supply planning system can see the overall network, but production under this design is planned parochially. Production does not *interact* with the supply plan in a mutually adjusting way; instead, the supply plan decides which plant is to produce what, sends the feasible (if constrained) initial production plan to the production planning system, and then simply leaves the details—such as the detailed scheduling—to be worked out by the production planning system. The two systems are kept in synch when, for instance, adjustments to the production schedule are made in the production planning system, but that is the extent of the interaction.

Once this rather inflexible design approach was adopted, all vendors, roughly speaking, adopted the same assumption, even though this assumption is often at odds with reality. An executive who has been shown the same workflow by multiple

---

[12] For a very detailed explanation of decomposition, see the SCM Focus Press book *Supply Planning with MRP, DRP and APS Software.*

software vendors may naturally question what is wrong with it if so many vendors follow the approach. The reason this design is at odds with reality is because factories do not merely accept raw materials and ship finished goods. Instead, many factories receive raw materials and ship out subcomponents. Other factories receive subcomponents and ship out components or subassemblies. Many possible combinations of factories are possible and always have been. Dated designs like this do not allow the full complement of location interactions to be appropriately modeled within the application. This is a complex set of alternatives, which changes depending upon the circumstances at various points in time. There can be very talented production planners, but there is an upward limit to which any human brain can compare and contrast complex sets of alternatives. Loosely translated, it is the perfect problem to hand over to a computer to solve.

In the history[13] of supply chain planning, methods have generally moved from the oversimplified to the more accurate as time has passed and the ability to model more accurately has improved due to advances in both software and hardware. For example, MRP, at one time the leading edge of supply and production planning, uses a number of highly simplified assumptions in order to generate its planning output. Multi-plant planning is yet another example of this continual increase in modeling accuracy. But, it must also be recognized, particularly by implementing companies as they bear the risks, that more advanced approaches have not always led to more implementable software. In fact quite the opposite is more commonly the case, which is why it is not sufficient to simply introduce a more accurate model or a more sophisticated method if doing so causes the likelihood of implementation success to decline appreciably. Unfortunately, too often this has been the history of new introductions of advanced functionality in all of the supply chain planning software categories. It's hard to emphasize this point enough: being sophisticated is not good enough. And just because a software vendor is offering a new technique or just because other companies have jumped on the bandwagon does not mean it's time to implement the software functionality. Many companies can and have purchased the newest, hottest thing, and have

[13] Analysis of history is a major part of our approach at SCM Focus. It relates to everything from our total cost of ownership research to our proposal that it is necessary to assign success probabilities prior to approving implementation budgets.

all failed together. And the really big companies such as Coca-Cola or General Motors seem to fail on their IT initiatives with just as much frequency as any other company—so the fact that some brand name has purchased an application or is implementing a particular functionality means very little. However, through software analysis, this book will show how PT has developed a sophisticated system for providing maximum production location interaction, and one which is also quite implementable.

How Multi-plant Planning Differs from the Standard Design

At SCM Focus we cover multiple areas of supply chain planning software, and through this research we often uncover similarities between the different software categories. For instance, in terms of recent developments in supply chain software, multi-plant planning shares many similarities with the supply planning technique of multi-echelon planning, which I discussed in Chapter 1: "Introduction," in that both methods/technologies expand planning decision-making to be location-agnostic.

With multi-plant planning functionality, each resource is assigned to a particular factory. Routings string together a series of resources through a single factory (the standard approach) or through more than one factory. This requires transportation between factories rather than simple stock movements between workstations within a single factory.

1. In some cases there may be only one routing between more than one plant for a particular product. In that case, the application is performing simple multi-plant planning.
2. In other cases the application may have one or multiple alternate routings—so-called Alternate Paths by PT—and the application must decide upon the best routing/path for a particular circumstance.
3. The application may also compare alternate routings, which are completely contained within one factory, meaning the decision is between routings that string together alternative resources within the plant versus routings which string together resources from multiple external plants. This is not multi-plant planning functionality, and any production planning and scheduling application can do this.

Any number of scenarios is possible. However, the core functionality is first the ability to have routings across factories (this is modeling functionality), and second, a way to make decisions among alternatives (in Galaxy APS this is the optimizer combined with adjustment rules that will be covered in detail further on in this chapter).

Multi-plant Planning Graphically

Below, I use a series of related graphics—beginning at the most simple and becoming more complex—to explain how to gradually build from the simplest scenario (two production lines in a single factory) to the most complicated scenario (multi-plant planning across three factories).

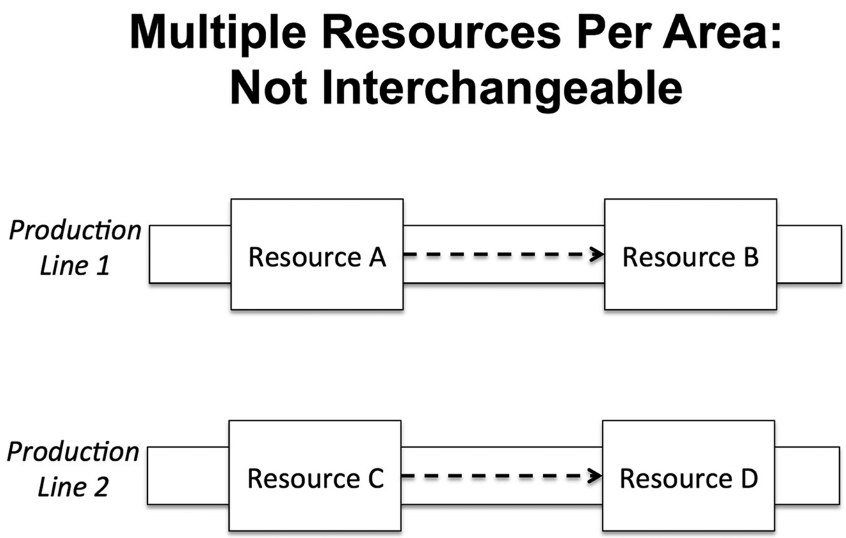

In this first example, we have two production lines in a single factory. Both resources A & C and resources B & D perform the same processing.

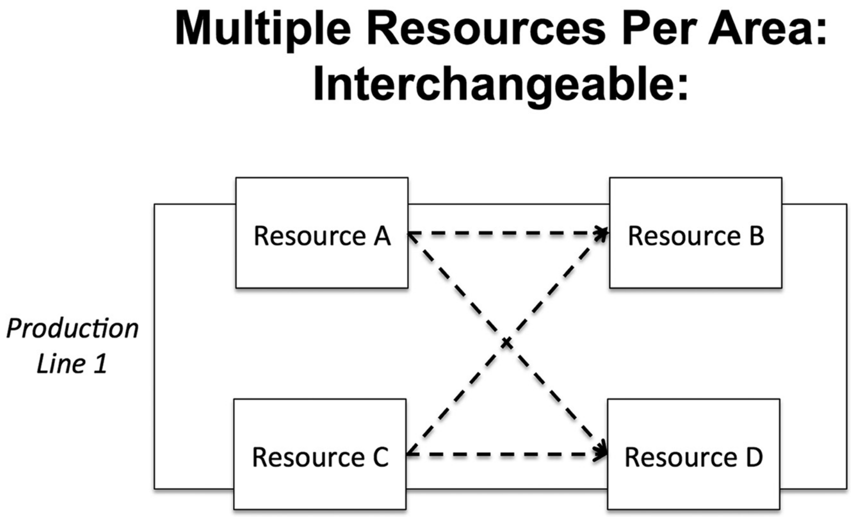

Here, because Resources A & C and Resources B & D perform the same processing, it is possible to aggregate the resources. Resources A & C can feed Resources B &D. This will serve as a baseline for the main point that we want to explain.

Moving to the Next Stage

These first two scenarios shown above have just been a starting point, using the example of resource aggregation within one factory. The following scenarios will not use resource aggregation. Planning level "aggregation" of capacity occurs while at the detailed scheduling level; the work is allocated to specific resources in particular plants. The intention of this clarification is to avoid any misunderstanding on the part of readers that the approach is a bucket-type aggregation of capacity across plants in the scheduling process.

Standard Design

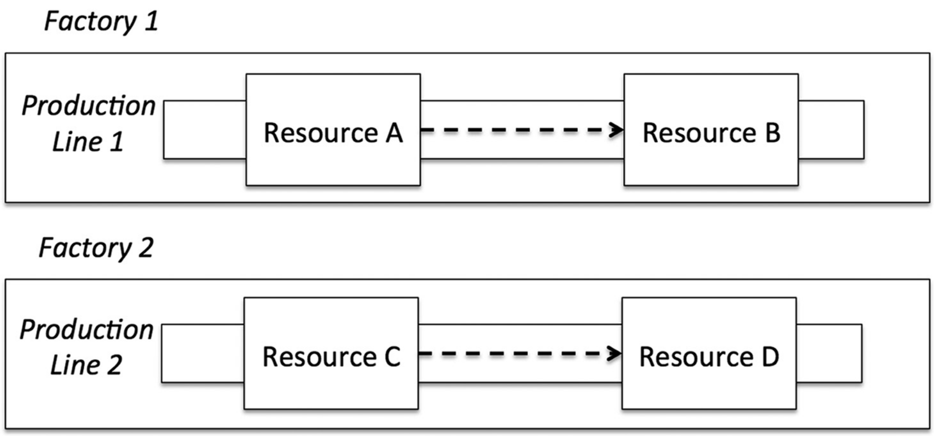

In this scenario, the resources, which are interchangeable, are in different factories. The vast majority of companies plan these locations separately.

Multiple Resources Interchangeable Between Factories

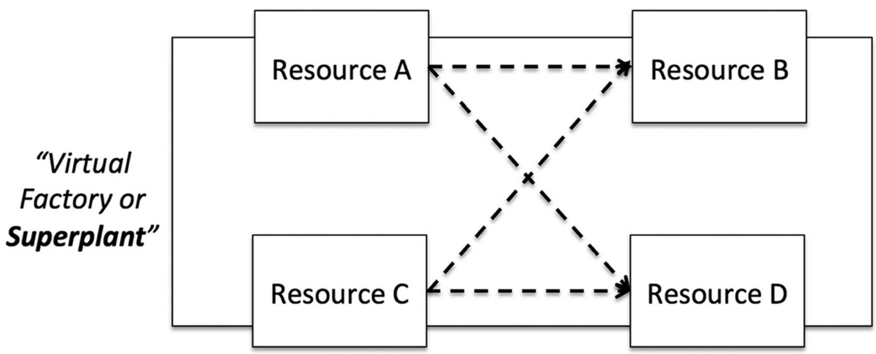

When production planning software has the ability to perform multi-plant planning, the ***locations can be planned as a group****.*

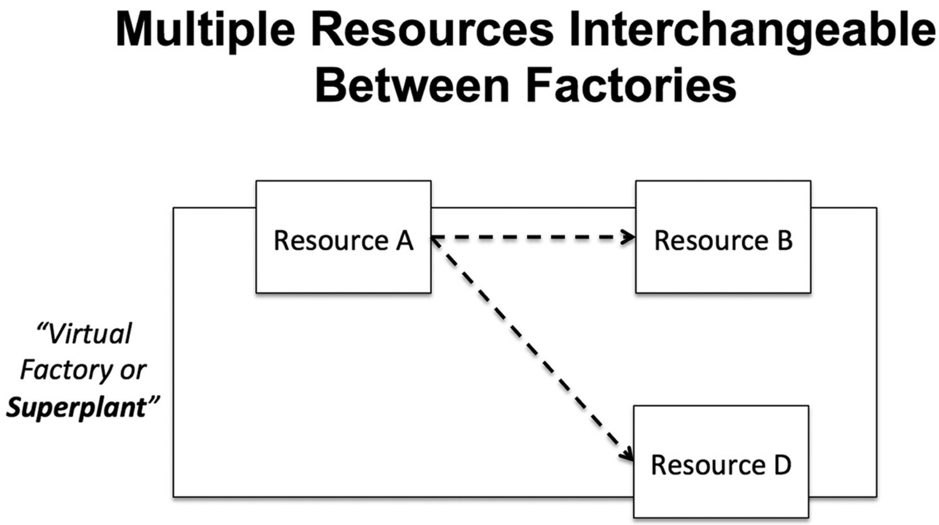

However, with multi-plant functionality, and if there is only one resource in one factory, the resource could service both production lines in both factories.

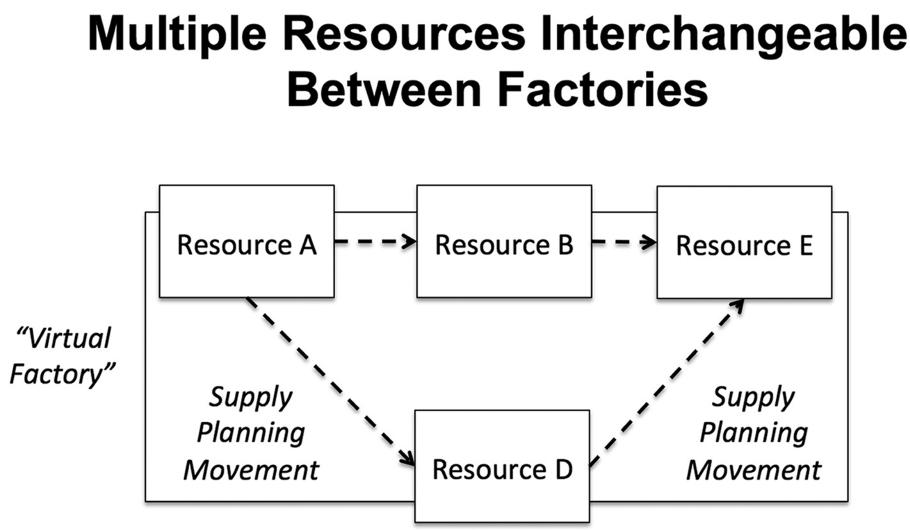

With true multi-plant functionality, the products can move not only to one factory, but also back to the initial factory. These movements are supply planning movements. These supply planning movements are created in response to the requirements that are generated by the selection of the alternate routings within the superplant. The non-multi-plant *planning applications can create stock transfers between locations based upon distribution requirements; however, only multi-plant planning can create stock transfers based upon a selection among alternate routings that are part of an overall production process* ***that runs through multiple factories****.*

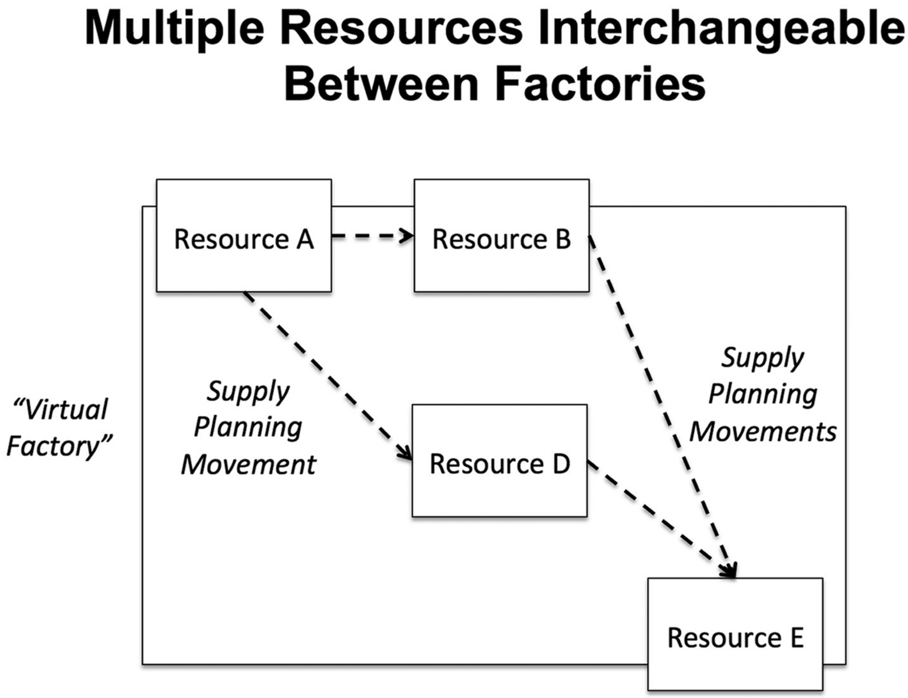

And that is just the beginning. In this example, two factories provide components to a third factory, which contains Resource E. All that is necessary is to model all the possible permutations, so the possible routings are set up in the system. Then, the software will choose the point at which it needs to create supply planning movements between factories when there are alternative resources between the factories. This allows the software to take advantage of the true capabilities of superplant; the software can shift the demand across resources in different factories when the capacity is available. In this example, the software can choose to supply Resource E from either Resource B or Resource D, or both Resource B & D.

Sequential Design (Not Multi Plant)

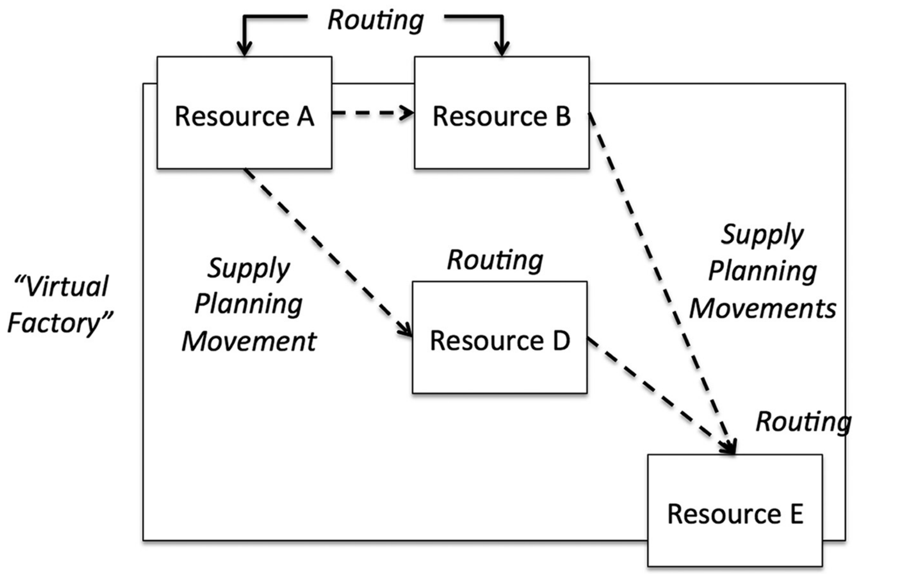

In the non-multi-plant planning design above, notice that each routing is only specific to each location. The supply planning system can pass stock transport requisitions between the locations; however, it cannot plan both supply and production together. Instead the supply plan is created, which then is passed to the production planning system.

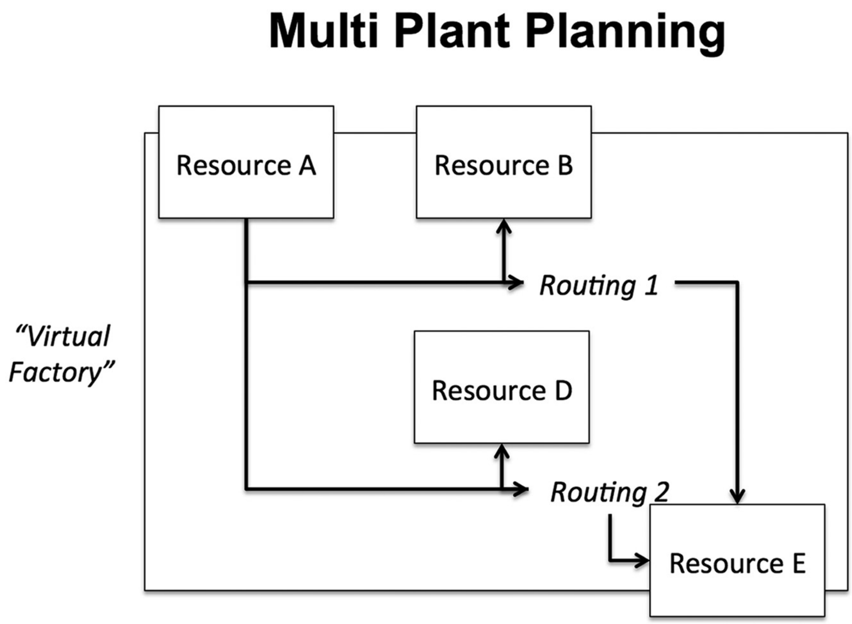

In order for software to be considered multi-plant, routings must be able to span the plants/factories. The multiple routings provide all of the possible permutations that are available within the supply network and the series of plants.

The Criticality of Alternate Paths to the Superplant

The ability to set up and use Alternate Paths within each location is standard production planning functionality, ***while Alternate Paths that span multiple locations is multi-plant planning functionality***. PT's training manual has the following to say about Alternate Paths.

> *"Each Manufacturing Order has one or more Alternate Paths that define the sequence of Operations. If there are multiple ways of making a product then an Alternate Path can be created for each method. For example, you may have a product that goes through a lathe operation followed by a grinding operation. In addition, a new CNC machine might be able to create the same product with a single*

operation. This can be accomplished by setting up an alternate BOM/ Routing that gets imported into Galaxy APS as an Alternate Path. One path would have two operations while the other path would have just one operation. (Note that if the only difference between two processes is the production rate on alternate machines then this can be accomplished using Product Rules or a Custom Add-in to customize the rate of the production on each machine.)

Each Manufacturing Order has one Default Path, which is always used during the Optimize process. From the Gantt View, you can manually drag and drop the job onto a different path by using the Alt Key and dragging the Activity Block [Alt + drag]. The alternate resources will be highlighted and you can drag to the alternate path. After dropping the block, if there is more than one Alternate Path that uses that Resource then you will be prompted with a dialog box to choose which Path to use."

However, Galaxy APS can also auto-select Alternate Paths if allowed. The selection is made during optimizations based on the path preferences and the availability of resources used by the various paths.

The Logic for Decision Making in a Multi-plant Environment

Up to this point, we have focused exclusively on the physical modeling aspects of multi-plant planning. However, how the application makes decisions between the many alternatives that can be presented to a multi-plant planning system depends upon what planning method is used. So let's spend a little time on the planning method to be employed.

The major ways that production planning is performed in advanced planning systems is with either heuristics or optimization. While it is conceivable heuristics could be used somehow, heuristics alone would be a poor method for making choices in a multi-plant environment. This is because heuristics alone tend to work better when there are fewer alternatives to evaluate.[14] Optimization, on the other hand, does have the capability of making these complex comparisons. But that is not the end of the discussion, because the next topic would be which kind of optimizer is to be used. This is because several different types of optimizers are used for production planning.

The oldest and still the most popular is cost optimization. If cost optimization were employed in the case of multi-plant planning, then costs could be applied to different routings, causing the lower cost routing/paths to be selected over the higher cost routings/paths—as long as there was sufficient capacity to meet demand. For instance, PP/DS contains a cost optimizer (PP/DS can use either a wide variety of heuristics—which are described in the article below—a cost optimizer, or in very rare cases, Capable to Match [CTM]—an allocation/prioritization method)

http://www.scmfocus.com/sapplanning/2008/09/21/ppds-and-snp-heuristics/

However, remember that PP/DS does not have multi-plant capabilities. PT's Galaxy APS is not generally deployed with a cost optimizer (although they offer a cost optimizer as an option), but instead tends to use a duration optimizer. Therefore the objective function of the Galaxy APS optimization is to minimize the overall time of the production plan.[15]

---

[14] I say heuristics "alone" because heuristics can be combined with an optimizer to improve the planning output. This is exactly how the Galaxy APS optimizer works in fact. Officially this is referred to as optimization combined with a heuristic-based algorithm. Most supply and production planning systems offer either optimization or heuristics, but have a heuristic-based algorithm interact with the optimizer. However, a detailed analysis of the Galaxy optimization approach is tangential to the main thrust of this book, so it makes sense to not go too far afield. This may be covered in a future book because the PlanetTogether approach to optimization is quite inventive. For those that would be interested in this material, please comment at the book's website page.

http://www.scmfocus.com/scmfocuspress/production-books/the-superplant-concept/

[15] Galaxy APS also has the ability to add optimization adjustment rules, which augment the duration optimizer so it considers or weighs other factors as part of the optimization. However, to begin this analysis without it becoming complex to explain, let us assume that the straight duration optimizer in PT is used without modification.

If there were a differential between the routings in terms of their time, then Galaxy APS would choose the alternate routing/path that could be completed most quickly, as long as there was sufficient capacity on this routing/path. In many cases, moving components between factories takes longer (when there is an actual duplicate resource in the same factory) than simply passing the component to another resource in the same factory. This can be reflected in the routing duration time. To show this, let's take a look at an example.

Multi-plant Planning Example

An example or scenario can help explain how a duration optimizer within an application with multi-plant planning capability would make its decisions.

Let's say that the internal factory durations for resource B and resource D, which are in two different factories, are both 10 hours for 100 units, but the total transit time between the factories is 15 hours. The durations in the routings could be set to reflect this. So let's view the following scenario and then see how the routings durations would be set.

1. It takes 20 hours to process 100 units at resource A.
2. It takes 10 hours to process 100 units at resource B.
3. It takes 10 hours to process 100 units at resource D.
4. It takes 2 hours to move components from resource A to resource B (which are in the same factory).
5. It takes 15 hours to transfer from resource A to resource D (which are in different factories).
6. It takes 12 hours to transfer from resource B to resource E (which are in different factories).
7. It takes 10 hours to transfer from resource D to resource E (which are in different factories).
8. It takes 25 hours to process 100 units at resource E.

Therefore, the following time would be assigned to each routing.

Production and Transit Times

Resource	End Resource (if applicable)	Same Factory or Different Factory	For 100 Units: Processing Time in Hours	For 100 Units: Transit Time in Hours
A			20	
B			10	
D			10	
A	B	Same		2
A	D	Different		15
B	E	Different		12
D	E	Different		10
E			25	

Sum of Hours for Routing 1	69	*Hours*
Sum of Hours for Routing 2	80	*Hours*

**Resources A and E are used in both routings, so these boxes is uncolored.*

If we stick to straight duration optimization, then Routing 1 would be selected as long as there were sufficient capacity at all the resources.

However, imagine that there is a delay of 18 hours on Resource B. This would not mean that the master data would have to be updated because Galaxy APS would see that a production order is already consuming resources when the order for 100 units would be scheduled to be assigned to Resource B. (There is the possibility of changing this production order manually, or having the system change the production order automatically if it were not firmed, but let us assume that the production order that is consuming Resource B when our order for 100 units would like to be scheduled on Resource B is not movable. If too many assumptions change, then we cannot see how any one factor behaves.)

Production and Transit Times

			For Production of 100 Units		
Resource	**End Resource (if applicable)**	**Same Factory or Different Factory**	**Processing Time in Hours**	**Transit Time in Hours**	**Delay**
A			20		
B			10		
B					18
D			10		
A	B	Same		2	
A	D	Different		15	
B	E	Different		12	
D	E	Different		10	
E			25		

Sum of Hours for Routing 1	87	*Hours*
Sum of Hours for Routing 2	80	*Hours*

*Resources A and E are used in both routings, so these boxes is uncolored.

Now we have added a delay of 18 hours to resource B. This can be seen with the new column added to the far right. This reflects the reality that resource B is consumed until 18 hours later than the time when our order for 100 units needs to be processed. Even though the software incurs more transit time, Routing 2 becomes the preferred alternative, allowing the production order to be completed 7 hours earlier.

This example shows the benefit of having multi-plant functionality. However, production and supply planning is situational or based upon a particular circumstance. What may be the best decision on some occasions is not the best decision at different times when conditions change. The less that the planning system can "flex" to different situations, the less powerful and usable it is. And this point is critical—this assessment of power that I just made is ***independent of the method***

deployed. And this gets to a very important point: planning systems tend to be evaluated based upon the sophistication of the methods that they offer—the methods being the decision logic used by the application. However, this is only one measurement of a system's power and ability to meet company requirements. It's an important criterion, but no more important than many other factors such as ease of master data update, modeling power, simulation capability, etc.[16]

To demonstrate this point I will use an example outside of supply planning and production planning, because demand planning software is such an excellent example of this exact issue.

Demand planning software is frequently graded on the sophistication of its forecasting methods (its seasonal models, its exponential smoothing model, etc.). However, the modeling capabilities, ease of adjustment, and ease of use factors are frequently ignored. I am in no way exaggerating when I say that here at SCM Focus we can do a number of things with regards to forecasting using an inexpensive demand planning application that the largest multinationals cannot do, chiefly because we test a demand planning system that grades well in multiple aspects, not simply the sophistication of its forecasting methods. Method sophistication is only one way to measure a demand planning application, or a supply chain planning application for that matter. One could have a very powerful method, but all the other areas

[16] The existence of a sophisticated method within a particular application does not mean that the method is in fact well-designed. It is in fact quite simple and common to take a complex method and design software in a way that makes it difficult to use the method in a beneficial manner. These types of distinctions in the design of software cannot be made by viewing a demo or speaking with software salespeople—but requires having first-hand exposure to the application. This determination can only come from testing the application and reviewing the configuration of the system. This methodology of increasing the exposure to the application is covered in detail in the SCM Focus Press book *Enterprise Software Selection: How to Pinpoint the Perfect Software Solution using Multiple Information Sources*. SCM Focus covers several sophisticated optimizers that are frequently sold as leading-edge but in fact are not and which have problems in implementation that are not simply due to lack of skill in implementing the software, or to the users not understanding the solution.

On the other hand, another optimizer of note runs into consistent problems not because the optimizer is poorly designed or because of users, but because the VPs and directors of supply chain within the implementing company frequently misunderstand that the optimizer calculates lower inventory levels because it is more intelligent than previous methods employed. This causes these companies to override the system output, negating the value of the application. Hopefully this makes clear that each situation must be reviewed individually to get to the root problem.

of the application could be weak, making it difficult to gain much value from it. Alternatively, an application could have a very good technical optimizer, but the optimizer is so difficult to configure, and the output so difficult to troubleshoot and tune, that again, it is difficult to gain much value from the application.[17] It is quite common for companies to buy sophisticated applications which they are never able to properly implement, when they would have been just as well off (and spent considerably less money) implementing far less sophisticated methods.

Multi-plant Modeling in Galaxy APS

Now that we have spent time going over the multi-plant planning environment graphically in order to grasp the this topic conceptually, let's dig into Galaxy APS to see how it is configured.

Galaxy APS has the ability to assign resources flexibly to what are referred to by PT as Capabilities, or Capabilities to Resources. A Capability is the ability to engage in an operation—to perform work. Capabilities in Galaxy APS are powerful because of both how they can be assigned to resources, how easily they can be copied to create new Capabilities (with slight adjustments) along with the fact that *"they can be created for special product attributes, such as item classes, material types of other constraints that define whether a Resource can perform a particular type of task."*[18] [19] The Capabilities Table is shown in the following screen shot:

[17] Interestingly, one particular product has one of the most difficult to configure optimizers as well as one of the most dated optimizer designs. Usually the product is not implemented with its optimizer because the results tend to not make any sense, it takes the most expensive hardware resources to run, and it is extremely difficult to troubleshoot and to tune. Users universally hate it—and yet it is one of the most popular production planning applications in the market! It is amazing how an application can receive the lowest grade on every single one of the criteria used by SCM Focus to evaluate software, and regardless, the software can continue to sell well in the market.

[18] PlanetTogether APS Product Training.

[19] What this means is that the Capabilities can very much depend upon a variety of factors, making them extremely flexible. *"For example, a packing Resource may have the Capabilities: pack, 8 oz, 16 oz, and 32 oz."*

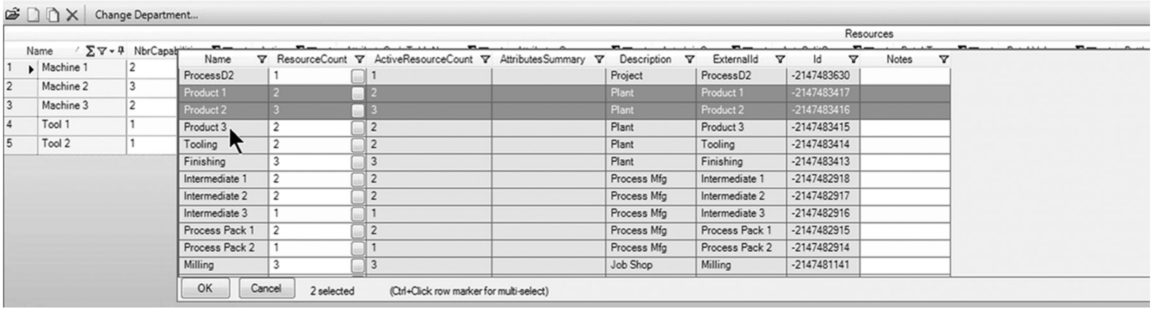

Name	ResourceCount	ActiveResourceCount	AttributesSummary	Description	ExternalId	Id	Notes
ProcessD2	1	1		Project	ProcessD2	-2147483630	
Product 1	2	2		Plant	Product 1	-2147483417	
Product 2	3	3		Plant	Product 2	-2147483416	
Product 3	2	2		Plant	Product 3	-2147483415	
Tooling	2	2		Plant	Tooling	-2147483414	
Finishing	3	3		Plant	Finishing	-2147483413	
Intermediate 1	2	2		Process Mfg	Intermediate 1	-2147482918	
Intermediate 2	2	2		Process Mfg	Intermediate 2	-2147482917	
Intermediate 3	1	1		Process Mfg	Intermediate 3	-2147482916	
Process Pack 1	2	2		Process Mfg	Process Pack 1	-2147482915	
Process Pack 2	1	1		Process Mfg	Process Pack 2	-2147482914	
Milling	3	3		Job Shop	Milling	-2147481141	

In this example, Product 1 and Product 2 Capabilities are already assigned to Machine 1 Resource.

In Galaxy APS the assignment of Capabilities and Resources is bi-directional, meaning that a Resource can also have a Capability assigned to it by opening up the Capabilities table and assigning resources in the same way. This is show in the following screen shot.[20]

[20] I am trying to keep the focus on superplant functionality. However, this ability to manipulate master data so easily in Galaxy APS is of great value not only during the implementation, but also for maintenance after the system is live. This means that users—those that actually know the master data the best—can take a bigger role in controlling the assignments between Capabilities and Resources without the need to rely upon technical resources. In addition to assignment, the creation of new Capacities and Resources is similarly fast and straightforward. These are important factors related to both implementability and maintainability.

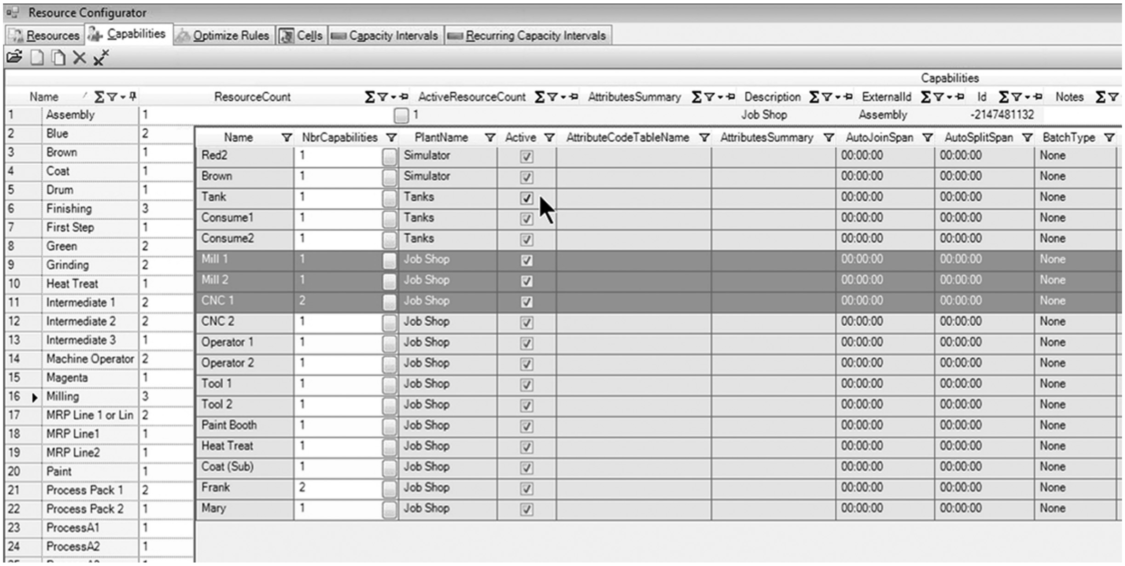

By opening the Capabilities tab in the Resource Configurator, a list of all of the resources that have been placed into the system appears. Currently, the Milling Capability is assigned to three Resources—Mill 1, Mill 2, CNC 1. However, the Resources that are assigned to the Capability are highlighted. To assign a new Resource to the Capability, all that is necessary is to select the line item that one wishes to assign.

Up to this point we can simply assume that all capabilities and all resources that we have discussed are in a single plant. In order to use this same assignment functionality to enable multi-plant planning, all that has to change is that at least one Resource would have to be a different plant. So if in the example above, Mill 1 and Mill 2 were in two different plants, Galaxy would have what it needs to compare two alternate routings/paths.

Moving Beyond Duration Optimization

In the previous paragraphs I described the fact that Galaxy APS uses a duration optimizer. Its objective function is to minimize time. We kept to the paradigm of Galaxy APS as being solely a duration optimizer without getting into its heuristic-based algorithm because it was useful to keep the scenario simple enough in other areas. The focus was on explaining how an optimizer makes decisions with different alternative routings/paths. However, as I alluded to earlier, Galaxy APS's optimizer is not limited to duration minimization. Galaxy APS also has a heuristic-based algorithm that allows other factors to count towards the optimization results. PT calls this an ***Optimization Rule, although it could just as easily have been called an optimization adjustment***. Galaxy APS has a number of available KPIs. These KPIs are controlled with sliders that force the optimizer to weigh a broad range of factors in addition to its objective function. This approach is far more powerful than a simple cost optimizer, such as those within many production planning applications and within PP/DS. This is because PP/DS only allows adjustments to be made along a single dimension: the dimension of costs.[21] After having worked on many optimization projects, the result is clear: using only one dimension for controlling the optimizer greatly limits the control over the planning output. The following article contains a further explanation.[22]

http://www.scmfocus.com/supplyplanning/2011/07/09/what-is-your-supply-planning-optimizer-optimizing/

---

[21] PP/DS "officially" allows its optimizer to be adjusted by noncost factors, but they are minor, and my testing of both the SAP SNP and PP/DS optimizers has demonstrated that any noncost factors that are offered as options do not work reliably enough to be included in any SNP or PP/DS implementation. I have an explanation of drivers that are not related to costs at this article link:

http://www.scmfocus.com/sapplanning/2008/10/11/snp-deployment-and-fair-share/

However, the problem is that many clients thinks that these capabilities work because the SAP marketing documentation and the release notes say they do. This is also a distinction, often lost on implementing companies, that just because functionality can be made to work, does not mean that it is maintainable. It is much easier to simply add functionality so it can be crowed about in the sales process, than actually making the functionality usable over the long term. Implementing companies are not lab environments with large budgets for maintaining marginal functionality. Even the largest companies tend to allocate few financial resources to planning. A quite common feature is to over-invest in the software and implementation stage—and under-invest in the planners and in system maintenance. Unfortunately too many executive decision makers have not spent the time analyzing the history of software implementation, so without this information it is difficult to know which areas to allocate financial resources. Consulting companies will tend to coax their clients into over-investing in the implementation stage because this is where their income is primarily derived. This is a misallocation of IT resources. Furthermore, functionality that is high in maintenance will tend to fall out of use after the implementation. In fact, functionality that was at one time used to sell the software, can fall out of usage for a variety of reasons—but many of them have a common underpinning in the company's under-funding of maintenance, continuous improvement of the solution, and inadequate solution socialization. (For more on solution socialization see the following article http://www.scmfocus.com/inventoryoptimizationmultiechelon/2011/05/socializing-supply-chain-optimization/.)

[22] Inventory optimization may be one of the few exceptions to this rule, because it so effectively matches the requirements of supply planning organizations (minimizing service level for a certain level of inventory, or vice versa).

On the other hand, Galaxy APS's heuristic-based algorithm can work in conjunction with Galaxy APS's duration optimizer to adjust optimization across a wide number of dimensions. How to use an optimization rule in PT is shown in the screen shot below.

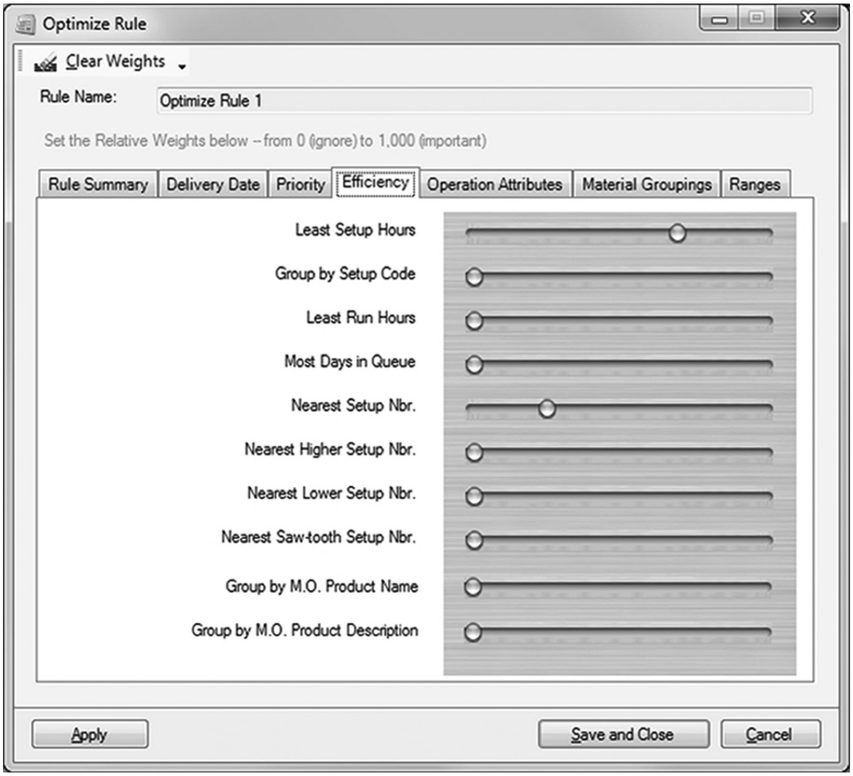

Optimization Rules—or augmentations to the optimization objective function of minimizing duration—can be added. An Optimization Rule is created by moving sliders in any of the Optimization Rule categories (Delivery Date, Priority, Efficiency, Operation Attributes, Material Groupings, and Ranges). A combination of all of the sliders makes up an Optimization Rule, which is then saved. Multiple Optimization Rules can be created and then assigned to the optimization planning run.

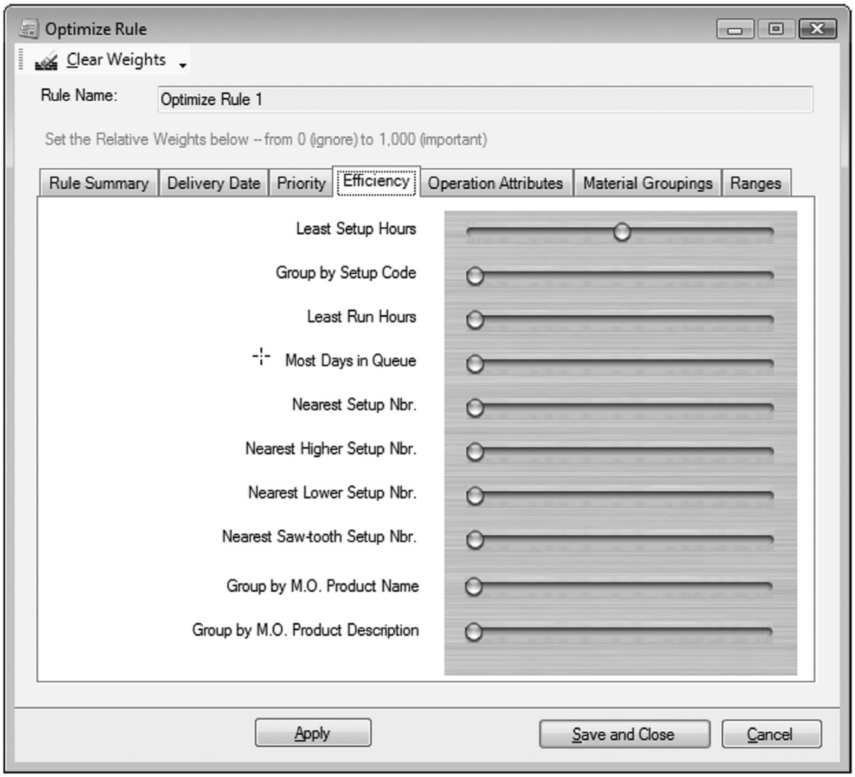

Here is a second Optimization Rule, which has sliders set on a different tab.

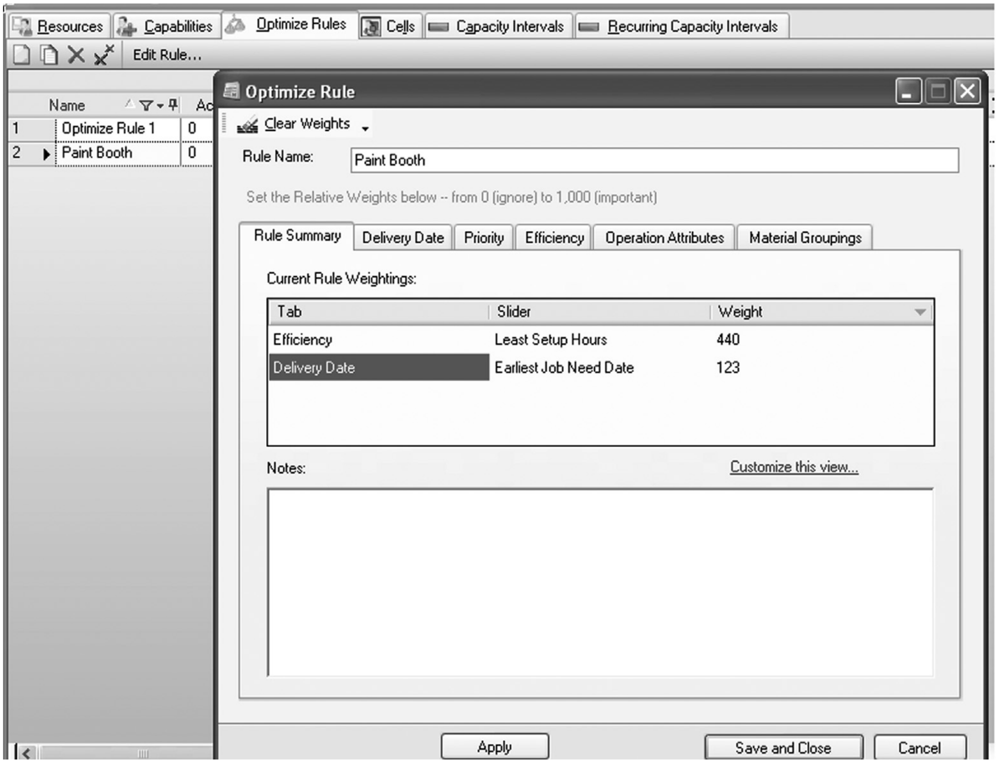

Here you can see the weights that have been applied to each Optimization Rule. You can add any number of Optimization Rules. Beyond that, not only can Optimization Rules be assigned to the overall model, but also to individual Resources as the screen shot on the next page shows.

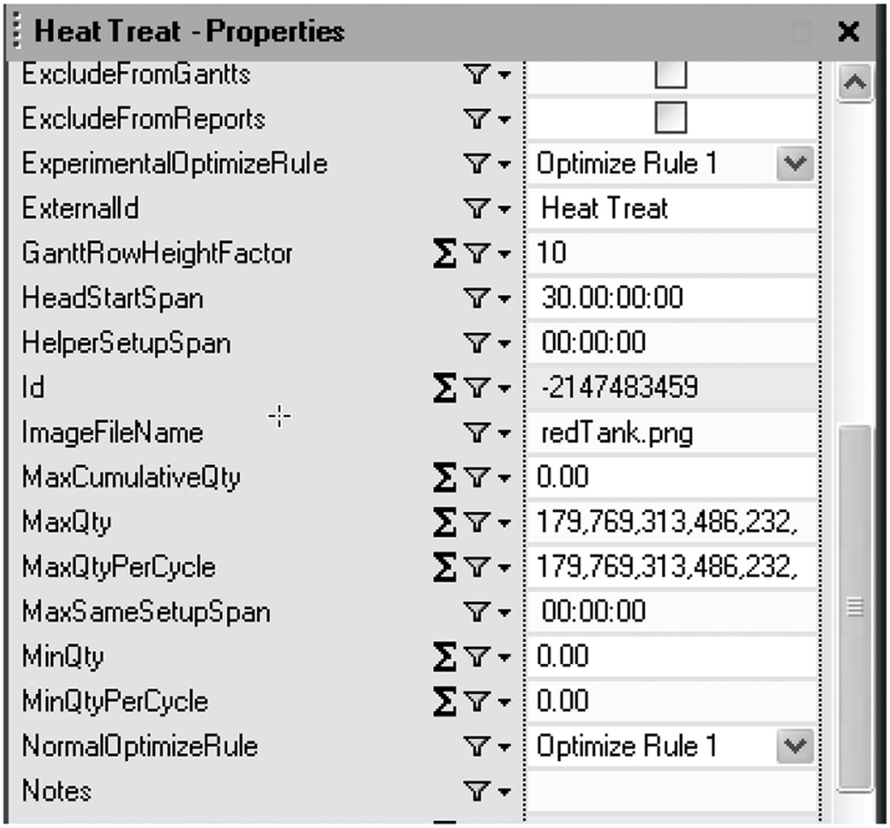

For this Resource the Optimization Rule 1 is applied. Also notice that there is an Experimental Rule. Experimental Rules tend to be used during simulation.

PT's user manual has the following to say on Experimental Rules:

> *"The system has been preconfigured with a default rule that applies to all Resources. During implementation, additional rules can be set up and applied to specific Resources. Each Resource can have a normal Resource rule and an Experimental Rule. During optimization you can specify which rule you want to apply. There are several ways you can experiment with optimization rules."*

Unless a special planning run is completed that breaks out the multi-plant products and locations from the non-multi-plant products and locations (and I am by no means recommending this), the same Optimization Rule that applies to non-multi-plant products and locations applies to multi-plant products and locations, as they are all part of the same optimization planning run.

Multi-plant Planning in PlanetTogether

Now that we have reviewed how multi-plant planning works conceptually, we can dig into how to configure Galaxy APS in order to enable multi-plant functionality. Several quotations from the Galaxy APS training manual will start us off:

> *"If plants are autonomous, meaning there is no interdependency among the plants, they can be optimized separately. There is no fixed relationship between users and plants so planners can schedule plants together or individually as necessary."*

This would be the starting point in terms of configuration. Configuration typically begins with setting up the model for simpler scenarios and then moving to the more complex scenarios. In fact, Galaxy APS can be configured with any number of alternate routings/paths. But in order for the routings to be actually used, a Manufacturing Order must be available to more than one factory.

In the quotation below, PT explains how much flexibility can be provided to manufacturing orders.

> *"Each Manufacturing Order can be locked to a specific Plant by having its Locked Plant field set. If this is not set then the Optimizer will schedule operations based on Capabilities alone (as if the Resources were in the same Plant). There are also 'Can Span Plants' settings in the Manufacturing Order and in the Job to control whether the M.O. or Job can have operations scheduled in multiple Plants."*

Therefore, multi-plant planning is essentially the default setting for Galaxy APS as long as the ***alternate routings exist***. If alternate routings exist—and let's say that they span factories—then the optimizer will attempt to use the alternate routing/path if the circumstances are right for it. However, even if the alternate routing/path would be feasible, Galaxy APS will not use it for planning purposes if Manufacturing Orders ***are locked to a particular factory***.

One might ask, *"Why would one want to lock a Galaxy APS Manufacturing Order to a particular factory?"* There can be several reasons for doing so. For instance, while a product could be made in multiple factories, based upon the final destination to a customer, it may make sense to have the order produced in a particular factory that is in close proximity to this customer, thus reducing lead time and shipping costs. Of course there can be other reasons, which will be explained further on in the book. However, the main point is that locking can be used selectively throughout the model and locked and nonlocked Galaxy APS Manufacturing Orders can be created in the same manufacturing run.

Multi-plant Optimization

The standard design of SAP APO and of most other supply planning and production planning applications is the following:

1. The supply planning application creates the initial supply plan and the initial production plan. If the supply plan has production resources, then either a constraint-based planning method is used or capacity-leveling. The supply plan applies for the supply planning horizon, which is frequently a year or more. The objective of incorporating constraints into the supply planning application is to pass a production plan to production planning that is feasible within each week.[23]

2. The supply plan is then passed to the production planning system, which makes adjustments to the initial production plan based upon much more detailed master data. Also, at this point the planned production jobs are moved around on some type of Gantt chart within the week and sometimes between weeks. This is all for the production planning horizon. An example of how the supply and production horizon interacts for non-multi-plant systems in SAP APO is shown in the screen shot on the next page:

[23] I do not spend time describing the distinctions between capacity-leveling and constraint-based planning in this book. The SCM Focus Press book *Constrained Supply and Production Planning in SAP APO* does cover this topic.

The main focus of most supply chain planning vendors that develop software in this area has been not to integrate the supply planning application and production planning application. Most software vendors simply assume them to be two different things. The weakness of this design is the natural inconsistency between the supply planning application and the production planning application. For example, SNP and PP/DS—and most other supply and production planning applications—work off of a different set of assumptions. In environments where there are dependencies between production, such as when a finished good in one factory is fed by semi-finished goods or components (or the components are in a third factory feeding the semi-finished goods plant—which I have seen at several companies), then the production planning and scheduling across the various plants ends up missing out on a number of planning opportunities that a multi-plant planning system could leverage.

Understanding the Outcomes of Supply and Production Strategies

Executives at most companies are trying to do too many things at once—things that essentially have contradictory outcomes. For example, executives want the reduced manufacturing costs of outsourced manufacturing, but also want to reduce their inventory. In general, it is difficult to meet multiple contradictory objectives but it is even more difficult when the applications selected are not appropriate for the requirement.

The Increased Complexity of the Supply and Production Business Environment

The complexity of the supply and production planning environment has increased significantly in the past three decades. Several reasons for this include the following:

1. *Outsourced Manufacturing*: The outsourcing of manufacturing to low-cost facilities in countries with long lead times from consumer markets.
2. *More Hand-offs in the Supply Network*: Increasingly complex relationships between companies and their suppliers, such as subcontracting, contract manufacturing, VMI, consignment (the transfer of inventory without the transfer of ownership), and intercompany transfer (supply planning must conform to the ownership transfer which is optimized for tax and other objectives unrelated to supply planning objectives).

3. *Growth by Acquisition*: This leads to companies (sometimes) having to combine supply chains that were previously separate.

4. *SKU Proliferation*: Marketing is very much in control of most US companies, and although most new products fail and are most often bad investments for the company, they continue to be introduced at a high rate. According to AcuPoll and Harvard professor Clayton Christensen, roughly ninety-five percent of new products fail.[24] Marketing only includes the costs of these introductions on marketing's direct costs of performing the introduction, never on the cost to supply chain operations and planning to manage all these new products.[25] New products include adjustments to products that make them better, or slight variations on an old product, but also just as importantly changes that cut out costs. Something supply chain planning has had to get used to is that no matter what the cost—changes must be a constant. Many "new" products are not recognizably new to consumers. For instance, in the spirits industry, a small change to the glass in a bottle will mean the creation of a new product number and the need to port the old demand history to the new product number. Product-location databases are in constant flux. What this means for planning is that most of the product database is "dead"—that is, planning systems are filled with product numbers that have no activity. As computerization has improved the ability to apply mathematical methods to planning, marketing has been on the other end, making the job of planning more difficult with not only continual new products, but promotions that erode the quality of demand history. Few companies are intelligent enough to buck this trend.

[24] Amusingly this new product failure rate has been consistent in the modern era—continuing right into the period of "Big Data." It is quite interesting that with all of these improved analytics, the percentage of new products which are successful cannot be improved even one percent by all of these amazing analytic products.

[25] Several examples of this inefficiency include reducing the law of large numbers, a major benefit to high sales items in the supply chain. The law of large numbers means reduced variability on the demand of products—leading to lower forecast error and lower safety stock. Another example is the necessity for shorter production runs, increasing the costs of the production of each produced product, as well as increasing the lead times and the lead time variability (the other factor in safety stock). Shorter production runs scheduled for more products means that each individual product must wait longer to be produced. The common trend has been to grow revenues in small increments, but to grow the product database enormously, so that the company has less average revenue per product. All of this is going on in the background, and companies are still often confused as to why they can't seem to increase their forecast accuracy.

http://www.scmfocus.com/demandplanning/2012/03/how-trader-joes-reduces-lumpy-demand/

As the article above describes, in most cases (the unusual grocery chain Trader Joes being the exception), it is exceedingly rare that supply chain planning departments have any say in SKU proliferation or in the lengthening lead times that they deal with due to the company using highly-integrated outsourced factories. Supply chain departments are expected to respond to the departments that drive strategy in the company, which in the US tends to be marketing, sales and finance.

Multi-plant Planning Versus the Common Manufacturing Trend

While multi-plant planning is required at some companies, other companies are moving in the opposite direction, away from integrated factories with increased outsourcing of subcomponents where the OEMs essentially take on the role of the general contractor. This is of course the contract manufacturing model that was covered earlier. However as both multi-plant planning and subcontracting/contract manufacturer planning are superplant functionalties, the requirement for superplant-capable software continues to rise.

Apple is an example of a company that does no production itself, and instead relies on overseas manufacturers to do all their manufacturing for them. In this model, the OEM makes almost all the profit and contract manufacturers make a margin of just a few percentage points on cost structures, which are based in the lowest-cost countries. It remains to be seen how long these companies can operate this model without becoming displaced by retailers or contract manufacturers that create their own brand, as the electronics company ASUS has done. ASUS began as a contract manufacturer, but they are now a major consumer brand as well as a contract manufacturer for other well-known brands.[26] Other familiar companies are similar. Samsung manufactures an enormous number of LCD

[26] Something that should be pointed out is that through contract manufacturing, an enormous amount of intellectual property (as well as process knowledge) is released outside of the companies that originally controlled this information when they performed their own production. Many companies that have engaged in substantial manufacturing outsourcing, particularly if it was to countries without intellectual property protection, can expect to see low-priced competitors that can do exactly what they do—because these contract manufacturers are currently doing it. These are the types of broader decisions that cannot be made on a simple manufacturing cost comparison basis.

screens for Apple, while competing directly against Apple's iPhone with the far less expensive Galaxy line.

However, for companies that are still vertically integrated—where the plants that provide the components are subcomponents for a finished good—the superplant concept is quite useful for creating a mental model of how the production and supply planning process needs to be designed in the associated planning systems. There are several important differences between a standard production design and the superplant design. I have discussed them throughout this chapter but wanted to list all of them in one location, which I have done below:

1. In a superplant, the BOM is distributed to multiple locations. However, the component and subcomponent BOMs are part of the overall finished goods BOM (either at the company itself, or at a company to which the planning company sells these components and subcomponents). Each BOM for the component and subcomponent are referred to as "modular BOMs." This distributed BOM applies to the multi-plant planning functionality and also to subcontracting superplant functionalities.[27]

2. In a superplant, production takes place in multiple locations, resulting in multiple production orders. The component and subcomponent production orders are triggered by distribution demand (through the supply network), rather than by dependent demand at a single location as part of the finished goods BOM.[28]

3. Synchronizing the continuous material flow between factories is critical to maintaining production efficiencies. However, the triggering of stock transfers may change depending upon plant proximity. In plants that are in close proximity to each other, the supply planning system may not need to be involved, as is described in detail in the

[27] Modular BOMs can be very easily created from master BOMs and version managed—if a company has the right BOM management software. However, most companies do not have this—so modularizing a BOM is much more difficult that it needs to be. The topic of BOM management systems is covered in the SCM Focus Press book *The Bill of Materials in Excel, Planning, ERP and PLM/BMMS Software.*

[28] Distribution demand is the demand that is sent between locations in a supply network. This is explained in detail in the following article: http://www.scmfocus.com/sapplanning/2012/10/15/understanding-the-flow-of-strs-and-prs-through-apo-with-a-custom-deployment-solution/

following article. http://www.scmfocus.com/sapplanning/2012/07/24/synchronizing-integrated-factories-with-stock-transfers/

4. In a larger sense, each factory can be thought of as a work center, and the flow through the supply network—for internal locations at least—can be considered a routing. Of course, this mental model, which is really the mental model of a superplant, only works for applications capable of multi-plant planning.

5. Because production is distributed across factories, when non-multi-plant planning software is used to plan a multi-plant planning problem, ***the overall manufacturing process cannot be capacity-constrained using a single bottleneck resource on a "single global production line."*** In effect, each production line in each factory must be separately constrained, and ***not as a single production process—which it in fact is***. With most production planning software, when all work centers/resources are in a single location, a single bottleneck resource can effectively constrain the entire production process. Generally speaking, this is one of the central concepts to production planning and the reason why there is a strong benefit to managing an entire sequence of work centers/resources ***with a single resource***.[29] This also explains why production resources are by far the most common resource types to be incorporated into supply planning systems, a topic which is covered in depth in the following article: http://www.scmfocus.com/supplyplanning/2011/10/02/commonly-used-and-unused-constraints-for-supply-planning/. However, Galaxy APS can constrain a bottleneck/drum resource along a routing regardless of how many plants it passes through. As would ordinarily be necessary, the lead times between the plants would have to be added into the system—but this is a simple

[29] This is in fact is a problem when a direct sequence of work centers/resources does not apply. A good example of this is with with process manufacturing. Most production planning and scheduling software which has been developed is designed for discrete manufacturing environments—it's the largest production planning software market and the easiest problem to solve—so naturally vendors have tended to target this market first. But, they will then often sell those same applications into other manufacturing environments—not by adjusting the application, but by adjusting their marketing literature. Companies choose inappropriate applications all the time, and then end up with major issues that negatively impact their planning for years—and this is just one example of why. This topic of matching production planning and scheduling functionality to process industry requirements is covered in the SCM Focus Press book *Process Industry Planning Software: Manufacturing Processes and Software.*

matter. Rather than having a transition of several minutes or hours between the feeding resource and the receiving resource, as would be the case if the resources were in the same plant, the lead time might be multiple hours or days when the resources are in different plants.

PlanetTogether and Planning the Superplant

Here are some important points related to Galaxy APS's multi-plant planning functionality:

1. In the Galaxy APS application, there is functionality that allows jobs to be assigned to either one plant or multiple plants. If the job is set to "Can Span Plants," each of its Manufacturing Orders is eligible to be scheduled to different plants. "*Material can flow between plants without having to issue transfer orders. This assumes that both plants are supplying the same warehouse.*"[30] This means Galaxy APS can search through the network of factories and "choose" where to create the Manufacturing Order. Because of this fact, Galaxy APS can treat all of the factory locations as if they are one virtual location. Galaxy APS calculates not only the production lead times, but also the supply planning lead times between the locations when it sets up a schedule that will meet demand.

2. In the example of building chairs, Galaxy APS's Manufacturing Order—which can span plants—may start a cutting operation in one plant while the next operation, a sanding operation, might occur in another facility. The option "Can Span Plants" must be checked as "true" at the Job level in order for the Manufacturing Order setting to take effect. Galaxy APS is designed to mix and match Resources across plants in order to string

[30] What this means is that in order to actually move the material between locations to support the Manufacturing Orders, the stock transfers would have to be created in either APO or the SAP ERP system. If, for instance, APO were used for supply planning, the STOs could be created in APO with an interface, the load built with SNP's Transportation Load Builder (TLB), and then sent over to SAP ERP. If APO or another supply planning system is not in the mix, then the stock transfers could be generated directly in SAP ERP in order to support the combined supply and production plan generated by Galaxy APS.

together the necessary production activities in the correct sequence—regardless of where in the supply network the Resources are located. These are alternate Routings that give Galaxy APS the maximum flexibility to choose the best possible solution—given all of the constraints.

3. Similar to locking an Activity to a Resource, a Manufacturing Order can be locked to a Plant to ensure—for any reason—that the Manufacturing Order is started and finished within that plant. So while Galaxy APS can be given maximal flexibility as described in the previous point, it can also plan like an ordinary production planning system.
4. A common reason for locking a Manufacturing Order to a Plant is proximity to the customer who ordered the product. If a Manufacturing Order is locked to a Plant, every Operation belonging to the Manufacturing Order will also be locked to that Plant.

Manufacturing Order locking can be seen in the screen shot below:

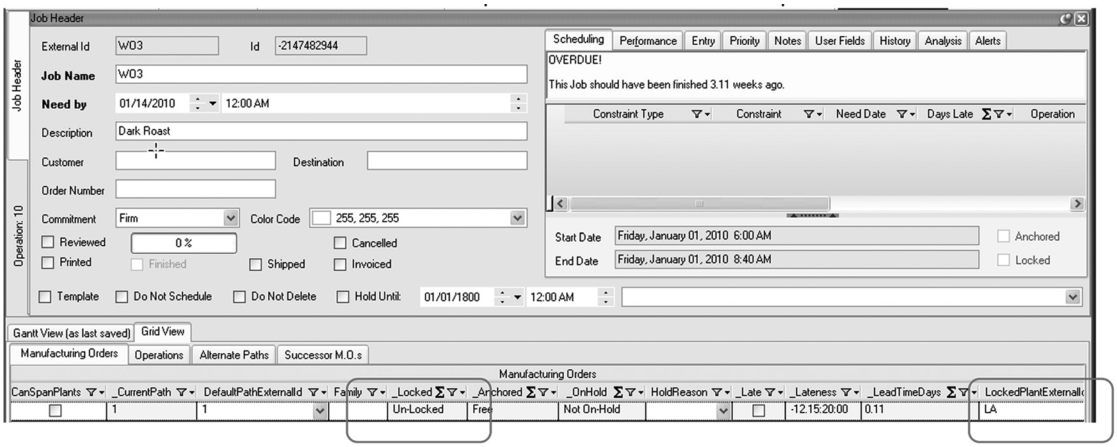

To perform cross-plant or multi-plant optimization, in addition to unlocking the Manufacturing Order you must also actively select to optimize the schedule for “All Plants” as shown in the screen shot on the next page:

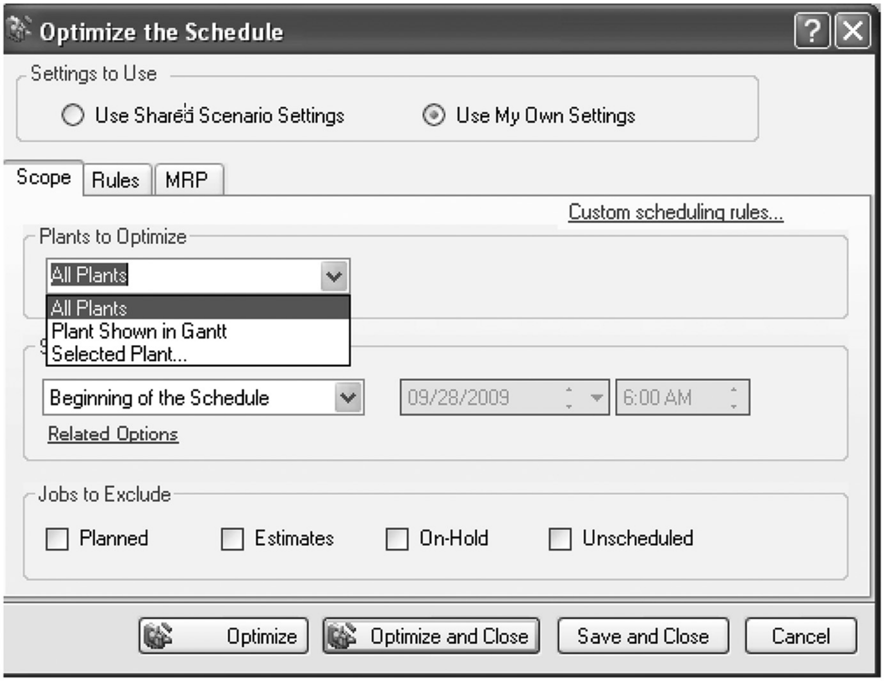

The options available in the drop-down for "Plants to Optimize" are All Plants, Plants Shown in Gantt (which means the plants that are brought up in the user interface) and Selected Plant.

This approach is different from the "supply planning first, followed by production planning" approach that I have described. When Galaxy APS is configured to do so, the Manufacturing Orders are created first, for a long planning horizon, and MRP is run off of the optimized production plan for the total planning horizon in one pass. Here are some of the benefits of this design.

1. Resources can be optimized by factory or across all factories within an Instance.
2. Each factory can be optimized according to its own optimization rules.
3. Users can use their own optimize options or share a common setup.
4. Each factory can have its own Plant Stable Span to enforce schedule stability.
5. Factories can share common resources.

6. Multiple warehouses can be defined for inventory management.
7. Inventory information can be filtered by warehouse.
8. Capable-to-promise (CTP) requests can specify which warehouses the request can draw from.
9. CTP can also be run in each factory where the item can be produced.

PlanetTogether Questions and Answers

The more one works with PT, the more unique both their technology and their implementation approach appears from other production planning and scheduling applications. In speaking with PT representatives, I found some of the following features of PT implementations to be quite interesting. I have presented these features in the form of questions and answers below:

Question #1: How long is the typical Galaxy APS production planning and detailed scheduling horizon?

Answer #1 (from PT): *Most of our clients favor longer scheduling/planning horizons (between one and three months of detailed scheduling and twelve months of planning) for more visibility—especially when they have long lead-time make or buy items.*

At first I found this answer surprising. I questioned why there would be such a large discrepancy between the typical production planning horizons of APO clients and PT clients. I started thinking about the fact that SAP and PT tend to sell to different types of clients. Generally speaking, SAP (and particularly SAP APO) is used at the largest companies, while PT tends to sell to smaller companies. Both vendors have customers in a wide variety of industries, and the production planning horizon and scheduling horizon should normalize around industry rather than company size. However, the answer actually relates to how Galaxy APS is used versus how PP/DS is used.

While Galaxy APS is generally marketed as a production planning and scheduling tool, it can actually perform both supply planning and production planning, depending upon the supply planning requirements.[31] So while SAP customers are implementing both SNP and PP/DS, PT customers are implementing a single product that covers both planning functions. It also means that PT customers are getting much more information about the production schedule than SAP customers because the scheduling horizon of most customers is three months. Details about how the planning horizon and related settings can be configured are covered in Appendix B: "Time Horizons in Galaxy APS."

Question #2: Do most companies that implement Galaxy APS understand that Galaxy APS's unified planning approach is a very different design than the traditional design?

Answer #2 (from PT): ***I don't think they appreciate this difference.*** *I think they just assume that this is how it should work—that the short-term schedule is in sync with the long-term plan. It makes sense if your mind isn't already thinking along the lines of traditional planning systems.*

PT's approach with regards to integrated supply and production planning, as well as the multi-plant functionality, is highly innovative and quite divergent from the standard software design in this area. However, one of the problems with innovative ideas is that they are more difficult to process mentally for the people and companies that could potentially benefit. In fact, when I first read Galaxy APS's documentation, I had to reread it several times, and then e-mail and discuss it with PT. This was all necessary for clarification because it was so different from the other applications I had worked with. In my mental model, the production was assigned to a single location. The only question was how much to make out of that location and when to create the planned production orders.

This led me to conclude that multi-plant functionality really needs to be made quite clear, and to be compared and contrasted with the traditional approach to production planning. Otherwise it will take quite a bit of time for companies to adopt this innovative functionality. This is of course why I decided to explain

[31] At least this has been the Galaxy APS marketing up to the point of the publication of this book.

Galaxy APS along with a contrasting application to show the new versus the traditional design.

Question #3: Are companies able to take advantage of Galaxy APS's ability to choose among many alternate routings and to treat all factories as if they are one large factory? Does this realization sink in later, earlier, etc.?

Answer #3 (from PT): *I think most companies create their work orders to a particular factory in the ERP and don't give Galaxy APS the choice—they schedule as separate/disconnected factories. That said, we're working with a new fifty-plant client now that plans to input alternates to Galaxy APS for the planning so they can capacity-balance across factories as you're suggesting.*

As I stated previously, it took me years to understand that Galaxy APS had this capability. I was aware of PT and had the Galaxy APS application on my laptop, but I simply never thought to ask about the functionality. I found Galaxy APS's capabilities because I happened to have a client with these exact requirements, and also happened to have a PT user manual and access to PT to ask questions. Companies may select PT for a variety of reasons, but they also obtain the capability to perform multi-plant planning and scheduling in addition to the other things that Galaxy APS can do. The second portion of the answer is shown below.

Answer #3 Continued (from PT): *On a related note…given the high volume of data in a fifty-plant twelve-month plan, with this client we're going to take the approach of having a Galaxy APS "Instance" for planning and a separate Instance for detailed scheduling. In this case, the detailed schedules will be imported into the plan. This way the plan will take the detailed schedule into account without impacting its performance. With smaller companies this isn't necessary; they can just use one instance for both.*

I found this answer interesting—primarily because I continually run into situations with SAP APO where clients try to get far too much out of a single server. This condition becomes worse on global APO implementations, which of course increases the processing load while at the same time shrinking the windows available to perform that processing. I have never been on a project where the application is split this way. But it puts Galaxy APS in interesting performance territory. It means that each server can be specialized for each task, and can be made more efficient at performing that task.

Question #4: Do companies communicate to PT that they have benefited from a more stable supply/production plan that is consistent throughout the planning horizon?

Answer #4 (from PT): *I don't recall ever getting this feedback but our consultants might. I do know that in the sales process they highlight this need—for the detailed schedule to drive the early part of the long-term plan.*

I am surprised by this answer. If it comes from a sequential supply planning and production planning process (driven of course by their previous application), then I would be at a loss as to how they would not observe a more stable schedule when using an integrated application like Galaxy APS.

Conclusion

As discussed, multi-plant planning, sometimes called multi-site planning, is the ability to model and make decisions to schedule production between alternate internal production locations that can produce the same product. By definition, companies that have components and subcomponents of final finished goods that are moved between factories have a multi-plant planning requirement. And this requirement applies to all manufacturing environments (discrete, repetitive, process batch, process continuous). A company that does not have multi-plant planning requirements when they start out, will have these requirements as soon as they choose to consolidate one stage of manufacturing to a single location in order to benefit from economies of scale and economies of specialization in that manufacturing process. By implementing multi-plant planning software, the company may be able to supply multiple plants. It is able to manufacture more product with fewer resources because it can receive a higher production utilization from its resources. Companies that have software that is capable of multi-plant planning are in a better position to leverage the production efficiencies of manufacturing consolidation. Therefore, software that is capable of multi-plant planning may be implemented by companies that already have multi-plant configurations, or software that can do multi-plant planning may enable a company to move to a multi-plant configuration only because they have no effective way of properly planning and controlling the factory in an alternate configuration.

Multi-plant planning is a more realistic representation of the real modeling requirements within many companies. This is because factories do not merely accept raw materials and ship finished goods. Instead, many factories receive raw materials and ship out subcomponents. Other factories receive subcomponents and ship out components or subassemblies. Many combinations of factories are possible and always have been at least to some degree. However, better communications, planning and transportation increase the opportunities to combine factories as pieces in an overall virtual production line. In environments where there are dependencies between production, such as when a finished good in one factory is fed by semi-finished goods or components (or the components are in a third factory feeding the semi-finished goods plant—which I have seen at several companies), then the production planning and scheduling across the various plants ends up missing out on a number of planning opportunities that a multi-plant planning system could leverage.

In a superplant, demand-type prioritization can sometimes be a requirement. If a company were to allow an external customer to consume capacity for a component for which there is internal demand, and for which they were constrained, then the in-process portions of the finished good would sit in inventory until the component can be produced. This would be an undesirable outcome.

Companies often require customized designs in their supply planning and production planning systems. Superplant is the distribution of what was previously one production process at one location to multiple production processes at multiple locations. Each production process at each location produces a modular BOM. The overall production process (that makes up the finished good) is then controlled in part by supply planning and in part by production planning—which are distinct and separate applications in most, but not every software vendor's suite of applications. This extra complexity brings up a host of issues ranging from what methods to use for each product location to how to control the stock movements between the related locations.

SAP development is just one example of a software vendor that chose a traditional and sequential design for their supply and production planning. This design is essentially unchanged from when advanced planning software was first implemented broadly in the mid-1990s. SAP APO, particularly the supply and production planning modules, is based upon the solution of i2 Technologies, as SAP maintained a relationship with i2 Technologies in the late 1990s. And i2 Technologies had this same sequential approach between its products Factory Planner and Supply Chain Planner. SAP's history with i2 Technologies is covered in the article below:

> http://www.scmfocus.com/scmhistory/2010/07/the-history-of-apo-and-the-influence-of-i2-technologies/

This is the most common approach to production planning software design, which is currently available in the market, and is a poor fit for companies that want to move to multi-plant planning.

Galaxy APS has multi-plant planning functionality, and can search through a network of factories and "choose" where to create the planned production order. In this way, Galaxy APS can treat all of the factory locations as if they are virtually one location. Galaxy APS is designed to mix and match resources across plants in order to string together the necessary production activities in the correct sequence—regardless of where in the supply network the resources are located.

In this chapter I posed some questions to PT and presented their answers, which indicate that while PT has developed functionality, many companies may purchase Galaxy APS either without knowing the software's full capabilities, or without implementing this multi-plant functionality. This is indicative of how previous experience and familiarity drives software development and software purchases, but also of how software is implemented. There tends to be very little