SYSTEM AND METHOD FOR MANAGING ATP

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to the fields of demand management, supply chain management, and capacity management. More particularly, the present invention relates to a system and method for managing available-to-promise(ATP) and making promises to fulfill customer requests.

THIS SHEET HAS BEEN REPLACED

This Sheet Has Been Replaced

This Sheet Has Been Replaced

BACKGROUND OF THE INVENTION

Manufacturers produce products for sale to customers. In the sales process, customers place demands on manufacturers. A customer demand may consist of a request for a particular quantity of a product by a specific date. This date and quantity information may be collectively referred to as the "customer request" or "request information".

Manufacturing and distribution facilities have limited resources (capacity) and limited inventories (materials). Therefore, every customer request may not be satisfiable in that some may receive no promise, others may receive an inadequate one. Planning and managing which customer requests to promise and fulfill, termed "demand management", is a fundamental and critical activity of most manufacturing and distribution organizations.

Due to material, capacity and other limitations, a manufacturer may not be able to meet a particular customer request. In this situation, the manufacturer typically negotiates with the customer to deliver a quantity of product by one or more dates agreeable to the customer. This date and quantity information may be referred to as the "manufacturer promise" or "promise information". Based on the manufacturer promise, the manufacturer creates operational plans to implement the promise information. Manufacturers may use a combination of diverse software tools in the negotiating and planning processes.

Traditional methods for demand management have several problems. First, such methods and systems are not integrated. Several different tools may be required to implement the entire demand management strategy.

THIS SHEET HAS BEEN REPLACED

Second, such traditional systems and methods are not dynamic. Once a plan is in place, it is difficult for the manufacturer to react to changing circumstances and update the plan. Third, order promising to customers is often done based upon an infeasible plan. Later attempts to find a feasible plan that will satisfy the promises are often futile.

The environment today requires more and more responsiveness. Customers require significant product diversity and want promises to be made to their requests immediately, while on the phone. The traditional way of promising in configure-to-order or make-to-order environments involves submitting the request to the planners and then, a few days or weeks later, after the planners have gone through a planning cycle, receiving a promise or rejection.

Many manufacturing and distribution organizations have several sales offices associated with each manufacturing factory. Each sales office independently promises to supply products from the factory to customers. This is referred to as a "distributed organization". Each sales person in each of the sales organizations needs to be able to make instantaneous promises, simultaneously with other sales people doing the same. In addition, each of those promises need to be fulfillable by a feasible plan.

To better meet customer demand, the manufacturer must build product and/or intermediate items before receiving customer orders. This production is based on projections called "forecast orders". A product produced based on these forecast orders is referred to as "available to promise" or "ATP". ATP consists of quantities of products with associated dates that the

products are scheduled to be available for delivery to the customer.

In distributed organizations a sales office may need approval from the factory before ATP may be promised to meet a customer request. This approval process may take up to a week under current practices. This delay is unacceptable in today's business environment.

This Sheet Has Been Replaced

SUMMARY OF THE INVENTION

In accordance with the present invention, a system and method for managing ATP is provided that substantially eliminate or reduce disadvantages and problems associated with previously developed systems and methods.

More particularly, one embodiment the present invention provides a method for managing ATP in a distributed organization. The distributed organization comprises at least one supplying facility such as a factory. Additionally, the distributed organization comprises a plurality of requesting facilities for each supplying facility. The requesting facilities may comprise, for example, sales offices. The requesting facilities and the supplying facilities may be coupled by a computer network.

In this embodiment, the requesting facilities each store forecast orders in a memory of a computer at the requesting facility. The forecast orders include request information. As defined previously, the request information includes the quantity (or range of quantities) of product requested from the supplying facility and the date (or range of dates) it is needed. A master scheduling software system may be used to selectively plan use of, for example, manufacturing capacity of the supplying facility to meet selected forecast orders based on predetermined criteria. If a feasible and desirable plan can be devised that satisfies the request, then the supplier may make a promise to the customer that the supplier will satisfy the request. The promises to meet the selected forecast orders may be transmitted directly to the customers over a computer network.

THIS SHEET HAS BEEN REPLACED

In environments where customers are not willing to wait for a plan to be developed to get a promise, the supplying facility must create promises in advance that are available for immediate transfer to a customer. In this embodiment, future requests can be forecasted and a plan can be made to satisfy and promise those forecast requests. When an actual customer request is received, one or more (or a portion of) promises made to forecast requests may be instantly reassigned to the customer request.

A technical advantage of the present invention includes the ability in a distributed organization with distributed sales people to allocate some of the promises made to forecast requests to certain sales people, thereby preventing them from simultaneously using the same forecast promise as a promise to a customer, without requiring them to check with each other before making promises. In this embodiment, each sales organization or person can be modeled and each forecast request/promise can be allocated to one such sales entity.

Another technical advantage of the present invention is that the allocation of promises may also be done for business management reasons. For example, a sales organization may be allocated promises based upon how much they are willing to commit to selling. This embodiment allows each sales entity to create its own forecast of what it could sell and establish the level it is willing to commit to selling. Forecast requests are then generated from the committed levels. Promises made to those requests become allocated to that sales entity for it to use to form promises for customer requests.

A further technical advantage of the present invention is that it allows these sales entities to be organized into hierarchies (for example, sales person

within sales office within marketing organization). Promises that are allocated to a sales organization can be used by the sales people within that organization. Coordination is required in such cases to ensure that two sales people do not consume the same promises. But where such coordination is feasible, it is typically desirable to have some allocations that are common among them.

Another technical advantage of the present invention is that customer requests that cannot currently be promised can be queued. As conditions change, the queued requests have the first opportunity to be promised. Without such a queuing mechanism, requests that cannot be promised are forgotten. When new capacity frees, the next customer that happens to make a request gets that newly freed capacity.

An additional technical advantage of the present invention is that an entire distributed organization of suppliers and customers can be modeled along with the requests and promises placed between them. In this way, planners can view, manage, and plan the activity of a whole network where the interfaces between elements must be formal (separate corporations).

Another technical advantage of the present invention is that each sales entity can define the "products" it sells, where a product is an item priced based on the item, the quantity, the order lead time (time from accepting the order to the requested due date), and the customer. For each such product, an independent forecast and commitment can be made, independent forecast requests can be issued, and independent promises can be received. In this way, promises can be allocated for requests with particular characteristics. For example, one product may sell an item for $5 if the order lead time is greater than 6 weeks. Another product may sell the same item for

$10 but with as short as 1 week lead time. Thus, a customer request with 6 week order lead time may be received when all allocations for that product have been consumed. However, if all the allocations for the 1 week order lead time product have not been consumed, then the customer can be given an option: the next available promise for the 6 week order lead time product would be 2 weeks later than your due date, or alternatively you may choose to pay $10 for the 1 week order lead time product and to receive it on time. Such management of products can prevent higher future profits for being sold at lower profits because they are promised first-come-first-served.

A further technical advantage of the present invention is that the forecast requests can specify how they expire. Some may shift out in time if they are not consumed; others may expire and disappear if not consumed. Such auto-maintenance of forecast requests can be very valuable in maintaining accurate forecasts and allocations for hundreds or thousands of products.

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the present invention and the advantages thereof may be obtained by reference to the following description taken in conjunction with the accompanying drawings in which like reference numbers indicate like features, and wherein:

FIGURE 1 is a block diagram of one embodiment of a supply chain model, including site models and seller models, and requests and promises between them;

FIGURE 2 illustrates one embodiment of a forecast entry for one of several forecast periods for one of several products within a seller;

FIGURE 3 illustrates one embodiment of a time horizon with forecast requests and actual requests showing the time horizon moving as time passes and the forecast requests adjusting in response; and

FIGURE 4 illustrates one embodiment of a seller model hierarchy and a product group hierarchy within a seller.

This Sheet Has Been Replaced

DETAILED DESCRIPTION OF THE INVENTION

The Supply Chain, Site, and Seller Models

Figure 1 is a block diagram of one embodiment of a supply chain model, including site models and seller models, and requests and promises between them. FIGURE 1 provides an example supply chain according to the teachings of the present invention. The supply chain model of FIGURE 1 comprises twelve site models, 12, 14, 16, 18, 20, 22, 24, 30, 32, 34, 36, and 38. These site models represent organizational units that may have the capacity and materials to produce or consume items. Each site can place requests for items upon other sites. Requests are in general indicated in FIGURE 1 by triangles 52, 62, 72, and 74. For each request 52, 62, 72, and 74, the site 12, 14, 16, 18, 20, 22, 24, 30, 32, 34, 36, or 38 being requested can make a promise to fulfill (wholly or partially) that request. Promises are in general indicated by inverted triangles 54, 64, and 76.

Other primary members of a supply chain model are seller models. The embodiment of a supply chain of FIGURE 1 consists of a single seller model 50. The seller model 50 is partially depicted in FIGURE 2 and consists of a list of products 110 that seller 50 offers for sale. A product model 110 defines the supplier site, the item at that site, a minimum order lead time, a minimum quantity, and the allowed customer sites. If a customer request fits those criteria of a product, then that request is eligible to be filled by that product, at the pricing specified by that product.

FIGURE 2 illustrates one embodiment of a forecast entry for one of several forecast periods for one of several products within a seller. For each product 110, a forecast horizon 112 is laid out. Forecast horizon 112

THIS SHEET HAS BEEN REPLACED

can be broken up arbitrarily. In this embodiment, three 1-week periods (the first being 114) are followed by three 1-month periods. For each forecast period for each product, a forecast-entry 116 is generated. The 'forecasted' and 'committed' values can be filled in. The value 'forecasted' is the seller's estimate for how much could be sold of that product 110 during that period. The value 'committed' is the quantity the seller is willing to commit to selling.

The committed quantity results in 'forecast' requests being generated in an amount equal to the committed quantity, spread out through the corresponding forecast period according to a forecast policy specified by the product 110. In the embodiment of FIGURE 2, the committed amount results in generation of requests 120 and 124, spaced out in the period 114. The site on which the requests 120 and 124 were placed (specified by the product 110) can then issue promises. Assuming promises 122 and 126 are made for requests 120 and 124, respectively, the value of 'allocated' in the forecast entry 116 for period 114 will be the sum total of the promised quantities.

The allocated amount is the summary amount the seller has available to promise customer requests. When customer request 128 arrives to the seller for product 110 during period 114, the seller can take one or both (or part of one or both) promises that it has already received, break them up or combine them to form a promise for the customer request. The forecast requests are simultaneously adjusted down by the amount of the customer request. So, for example, if the committed value of forecast entry 116 was 500 units, the two forecast requests 120 and 124 were for 250 units each, the two promises 122 and 126 were received for 200 units,

and the customer request 128 was for 300 units, then the two forecast requests 120 and 124 will be adjusted to a total of 200 (i.e., 200 and 0 or 100 and 100 or some other combination, dependent upon the product's forecast policy). The two promises 122 and 126 will be adjusted to a total of 100, and a new promise 130 will be created for 300 units to satisfy request 128. The 'committed' and 'allocated' values of forecast entry 116 do not change as a result, but the 'requested' and 'promised' values do. When 'promised' is equal to 'allocated', then there are no more promises available for promising customer requests.

This process is also depicted in the supply chain model example of FIGURE 1. In FIGURE 1, seller 50 generates forecast request 52 on site 22 for delivery to site 30 (which need not be a physical site). Request 52 results in site 22 generating operation 56 to perform the activity involved in delivering the requested items to site 30. If operation 56 is feasible to perform, then site 22 may choose to create promise 54 to seller 50 that the item can be delivered as requested by request 52.

Site 34 then places request 62 through seller 50 for the same product as request 52. If that customer request 62 is consistent with what seller 50 was forecasting, then seller 50 can reduce request 52, promise 54, and operation 56 by the amount of request 62, and then add promise 64 and operation 66 to fulfill request 62. That simple action did not require replanning through site 22. Effectively, the ability of site 22 to satisfy request 62 had been pre-computed in the form of promise 54. Thus, that promise 54 can be split in order to form promise 64.

A primary caveat is that the load and times of the operation 56 may not be valid when split into operation 66. For example, if operation 56 involved using a truck

to transport the items, then splitting out operation 66 may result in an additional truck being used. If none was available, then operation 66 may have to wait. To compensate for this, each product defines criteria for splitting promises, which can include an amount of time with which to pad the due dates quoted.

Of the site models that make up a supply chain model (as in FIGURE 1), some of the sites can be under the control of that supply chain model, while others can be modeling sites which are planned independently. A field of the site model called 'managed' indicates which sites are managed by this supply chain model and which are not. Two sites that are both managed do not need to make formal promises between each other -- the request will generate an operation and all changes to the requests are immediately passed through the operation to the other site. Requests between a managed Site and an unmanaged site require formal promises. The promises must be made explicitly, and once accepted constitute a rigid agreement between two Sites. Changing that agreement requires both sites' consensus.

Adjustment as Time Passes

Forecasts are often, by their nature, wrong. Thus, as time passes and customer requests arrive faster or slower than expected, it is desirable to modify the forecasts as appropriate. Given a large number of products and numerous forecast periods, automated adjustment is highly desirable.

Thus, the product forecast policy can specify how the forecasted and committed quantities should be adjusted as time passes and actual Requests are received or not.

FIGURE 3 illustrates one embodiment of a time horizon with forecast requests and actual requests showing the time horizon moving as time passes and the forecast requests adjusting in response. The timeline 200 represents the initial state. Forecast requests 202, 204, 206, and 208 have been made in their respective forecast periods. Customer requests are indicated with triangles, as shown. The two customer requests 222 correspond to forecast request 202. The three customer requests 224 correspond to forecast request 204.

Time passes and no more requests are received. The timeline 210 represents that later state. Time has advanced beyond the forecast period of the forecast request 202. The customer requests 222 received during that period were less than that forecast request. One option is to assume the forecast was too high and simply expire the leftover forecast. Another option is to assume the forecast quantity is right, but that the timing is off -- that the total quantity will be requested soon. In the latter case, the forecast request should be moved forward in time and reduced in quantity. This is shown as forecast request 212. There are many other options for how to expire, reduce, or increase forecast requests based on the arrival rate of customer requests that can be encoded in the product's forecast policy.

Allocation to Sellers

FIGURE 4 illustrates one embodiment of a seller model hierarchy and a product group hierarchy within a seller. FIGURE 4 shows two Seller hierarchies. Seller 410 represents an Industrial Products marketing division, and seller 420 represents a Consumer Products marketing division. Within Industrial Products 410, there are

three sales offices that each handle a region: the North is handled by seller 412; the South is handled by seller 414; the West is handled by seller 416. Each sales office is made up of numerous sales people, who are each represented by a seller (for example, Joe is seller 418 and Sally is seller 419).

In many organizations, the sellers may own their own allocations against which they can promise to their customers without consulting the company. However, sellers need not own any allocations. For example, Joe 418 and Sally 419, along with the other sellers in the South sales office 414, may each forecast what they intend to sell. Those forecasts are aggregated up to the sales office seller 414, where they are used as an input. The seller 414 can independently forecast for the whole sales office. That, in turn, is allocated up to the Industrial Products 410 division.

Clearly, forecast requests should not be generated for the forecasts at all three levels -- that would result in triple the requests appropriate.

Instead, each seller can independently commit to selling some or all of the forecast. By committing, forecast requests are created in order to obtain promises which can be used to promise their customers. Those promises are owned by (or controlled by) that seller that committed to selling that amount.

However, it may be that some sellers do not commit at all. For example, none of the salespeople, including Joe 418 and Sally 419 commit to any of the forecast. Instead, the South sales office 414 commits as a whole. That results in allocations to the South Seller 414. Those allocations can be used by any of the sub-sellers, such as Joe 418 and Sally 419. However, such collective usage of the allocations requires coordination. They

must reserve the amount they need before they can actually promise it, since the other sales people may be considering using the same allocations.

A seller is committed to anything its sub-sellers commit to. However, a seller can commit to additional, beyond what the sub-sellers commit to. For instance, each sales person may make a conservative commitment. The sales office will know that some of the sales people will surely sell over their commitment, but it is not clear which sales people. So the sales office can commit to sell additional, and those additional allocations will be available to the first sales people who exceed their personal allocation.

Product Groups

Forecasts tend to be more accurate in aggregate. A monthly forecast will generally be more accurate than a weekly forecast. A forecast for North America will generally be more accurate than a forecast for Texas. Similarly, a forecast for milk will generally be more accurate than for skim milk in pint containers.

Thus, it is important to be able to aggregate up forecasts, modify the aggregated forecasts, and propagate the changes back down to the individual products. The product group model supports this functionality.

Product groups form hierarchies. A product group can have at most one parent product group, and thus can be in at most one product group hierarchy.

Products, on the other hand, can appear in numerous product groups; however, only in one product group of any one hierarchy. A product group defines one consistent hierarchy for aggregation. However, sellers will need to aggregate the products in many different ways. For example, milk products can be aggregated by their

container size (gallon, half gallon, quart, pint), by their fat content (whole, 2%, 1%, skim), by the customer grouping (grocery-store, restaurant, convenience-store), or by brand (ECONO-COW, PUREWHITE).

product group s are depicted in FIGURE 4. Products 450, 452, 454, and 456 are grouped into two product group hierarchies, rooted at product groups 430 and 440. Product group 430 is broken down into product groups 432, 434, and 436.

Advanced Available-To-Promise (ATP)

Each seller has allocation (promises) available for the various products sold. When a customer request comes in to a seller, there may be numerous products that match the request. If the lowest cost product can fully satisfy the request (has sufficient quantity by the requested due date), then the request can simply be promised. Otherwise, a decision may be needed. For example, the customer may be able to choose to have it for a low price but a week later than requested, or by the date requested but 10% higher price. It may be that half the order can be completed on time at the lowest price, but the other half can either be delivered later or for a higher price, and so on. Thus, the ATP can be a list of different products (pricings) with different order lead times, minimum quantities, availability dates, and availability quantities.

Extensible Product Model

The product model type has a forecast policy extension selector that allows additional fields and semantics to be added to a product model. Extension selectors are described in more detail in U.S. Application Serial No. 08/491,153, filed June 16, 1995,

and entitled EXTENSIBLE MODEL NETWORK REPRESENTATION SYSTEM FOR PROCESS PLANNING (Attorney Docket No. 020431.0136), the disclosure of which has been incorporated herein by reference.

In this way, additional forecast information such as forecast error or forecasted variance in either quantity or time or both can be input and used. Additional fields for expected skew during the month can affect how the committed quantity is split out into forecast requests. The expected variance or order arrival rates can affect how forecast requests expire or adjust as time passes, based on the customer requests that have been received.

Although the present invention has been described in detail, it should be understood that various changes, substitutions and alternations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

WHAT IS CLAIMED IS:

1. A software system for managing available to promise and making promises to fulfill customer requests, the software system comprising:

at least one seller model representing a seller that is selling at least one product, the seller model operable to forecast for the at least one product and operable to choose commitment levels creating forecast requests;

the forecast requests receiving promises made by supplier sites; and

the promises available to the seller entity to promise to actual customer requests.

2. The software system of Claim 1, wherein the at least one seller model can comprise a hierarchy such that allocations to a seller can be used by the seller or any sub sellers and such that commitments made by a seller also commit the related sellers.

3. The software system of Claim 1, wherein the software system is located in and executed by a digital computer, the digital computer comprising:

a data storage device;

an execution memory operable to hold the software system; and

a processor coupled to the data storage device and the execution memory, the processor operable to execute the software system.

THIS SHEET HAS BEEN REPLACED

This Sheet Has Been Replaced

4. A software system for managing available to promise and making promises to fulfill customer requests, the software system comprising:

a supply chain model representing a chain of supply, the supply chain model comprising;

site models that represent sites having capacity and that manage material flow; and

seller models that represent sellers and that manage forecasting and purchasing;

wherein commitments between sites is modeled by requests and promises; and

wherein the sellers can post requests on behalf of sites in anticipation of future requests from the sites.

5. The software system of Claim 4, further comprising a queue that allows rejected requests to be queued for consideration whenever capacity frees up.

6. The software system of Claim 4, wherein the software system is located in and executed by a digital computer, the digital computer comprising:

a data storage device;

an execution memory operable to hold the software system; and

a processor coupled to the data storage device and the execution memory, the processor operable to execute the software system.

THIS SHEET HAS BEEN REPLACED

This Sheet Has Been Replaced

7. A software system for managing available to promise and making promises to fulfill customer requests, the software system comprising:

a product model representing a product, the product model specifying a supplier site, an item produced by that site, a minimum quantity, a minimum order lead time, a list of customers allowed to purchase, and pricing for the product;

wherein a customer request having desired characteristics matching the product can be fulfilled by a promise of the product;

such that a list of all matching products and associated available promises can be displayed as available-to-promise for the request.

8. The software system of Claim 7, wherein the product model specifies an extension allowing different forecast information to be recorded and used to compute and distribute forecast requests.

9. The software system of Claim 8, wherein the product model specifies expiration and adjustment semantics for forecasts as time passes and as actual customer requests arrive faster or slower than expected.

10. The software system of Claim 7, wherein the software system is located in and executed by a digital computer, the digital computer comprising:

a data storage device;

an execution memory operable to hold the software system; and

a processor coupled to the data storage device and the execution memory, the processor operable to execute the software system.

THIS SHEET HAS BEEN REPLACED

This Sheet Has Been Replaced