Applying AI to Manufacturing: Linear Order Promising and Production Planning

Yury Smirnov
Calico Commerce
333 West San Carlos st.
San Jose, CA 95110
ysmirnov@calicotech.com

Abstract

Many vertical industries within Manufacturing have already entered or are about to enter a new era of mass-customization. Customers expect improved level of service, precise price and date quotes for their personalized orders. Internet communications in general and dedicated e-commerce efforts in particular greatly facilitated the process of taking orders and shipping the requested products to anticipating customers. However, precise, scalable and effective Order Promising and Production Planning still constitute serious challenges for manufacturers.

Specialists in Manufacturing Modeling have already identified the deficiencies of the existing approaches that traditionally split production models into Bills of Materials (BOMs) and Routings (Goldratt 1990). Whereas Artificial Intelligence (AI) understood long ago the benefit of merging states and actions in a combined planning model, an alternative, constructive solution to the BOM/Routing modeling approach has not been explicitly proposed.

Re-configurable products may lead to an exponential explosion of the number of BOMs, if the standard modeling approach of listing all orderable products is followed. Another complication may come from the existence of alternative routings, which are different production processes (actions) that produce the same inventory items (lead to the same states). A selection of a different route may imply substituting already selected group of inventory items by a different group of items, for example, changing a monitor type for PC may require a different video card, which in its turn may need an upgrade of the power supply module. The above feature is called "kitting" in Manufacturing Modeling.

On one hand, a complicated nature of Manufacturing Modeling and a need to capture the AND/OR-logic in presenting inventory items and alternative routings makes it hard to efficiently derive precise price and date quote (Order Promising) and to construct the entire schedule (Production Planning). On the other hand, customers' expectations and a broad spectrum of orderable products state an urgent need for scalable Order Promising and Production Planning functionalities.

In this paper we introduce novel modeling approach that applies some AI modeling techniques to Manufacturing Modeling, allows to avoid the exponential blow-up for re-configurable products and captures the AND/OR-logic without additional modeling efforts. Furthermore, we state a simple, realistic resource sharing assumption. For the introduced type of models, we construct Order Promising and resource allocation (scheduling) procedures that are linear under the stated assumption for any homogeneous objective function.

Introduction

Modern World has emerged a new era of mass customization. This trend is changing the way customers are making purchases, it also has a strong impact on how products are made. Known as on-demand manufacturing, it spans already a diverse set of manufacturing areas from luxury cars, to computers, books and toys just to name few. Actually, the greatest stories of recent success are directly associated with the ability of the vendor to satisfy customers on an individual basis without jeopardizing prices and delivery terms.

Traditional, make-to-stock manufacturing world is undergoing a significant transformation too. In recent years, an increased competition, changing economic environment, and new government regulations have conspired to put pressure on process plant margins. Computer components, for example, decline in value at a rate of about one percent a week. In such an intensive environment, the key test for the system is to ensure that the right products are delivered to the right place at the right time. This test determines the whole suite of requirements for planning, scheduling and execution systems that constitute the back-end of a customer-oriented manufacturer.

E-commerce came as a tool that facilitates accepting orders, captures the logic of customer/vendor communication and tracks shipments. The vast majority of e-commerce efforts are intended for production-free businesses. There are several serious issues of production-based e-commerce that have not been fully addressed. In this paper we attempt to identify those and resolve them effectively under realistic assumptions.

There seems to exist a shift of interest within the order entry process for customizable products. As long as the complexity of the product is below a certain, "Common Sense" level, customers are mostly interested in the availability of the product. However, as soon as the configuration process becomes complicated, non-intuitive, then customers are primarily interested in the correctness of their choices, in the compatibility of the selections. Over time, the "Common Sense" level of the configuration complexity tends to grow along with the educational level of users. E-commerce provides a great hosting environment that promotes such type of education.

This paper introduces a set of order promising and production planning problems arising due to mass-customization and, hence, augmented by the popularity of e-commerce. Besides the problem introduction, we describe another way of modeling manufacturing environemnts in the AI style. We also show that under a simplifying assumption for such models one can build extremely efficient order promising, planning and scheduling algorithms.

Manufacturing Planning Problems

Current state of manufacturing modeling is several dozen years old. Earlier Material Requirement Planning (MRP) systems had to find a compromise between the modeling expressiveness, computational power and slow access to peripheral devices, such as tape-based memory systems. Although MRP itself was viewed as a "Copernican Revolution" (Vollmann, Berry, & Whybark 1992), it has been identified that the current "state-of-the-art" manufacturing modeling used in plain MRP and its successor Manufacturing Resource Planning (MRP II) do not support anymore the growing need in providing efficient grounds for the enterprise level of integration. One of the main question to be answered with respect to the modern state of computer equipment is as follows: "Why is the product structure's file segmented into the BOM and Routing files?" (Goldratt 1990).

The inertia of the field keeps the above splitting principle untouched for old projects, as well for newly launched ones, thus adding even more to the inertia's spinning momentum. However, it is becoming more and more obvious that such a modeling approach has certain limits, in particular when applied to a highly customized product line. For example, a PC with 15 types of hard drives, 10 types of RAM, 5 types of video cards, 5 types of modems, 5 types of I/O Buses, 10 types of monitors, etc. would require at least $15 * 10 * 5 * 5 * 5 * 10 = 187,500$ different Bills of Materials (BOMs). Dealing with dozens and hundreds of thousands of BOMs is a regular situation in on-demand manufacturing.

However, the traditional manufacturing modeling approach requires all BOMs for orderable products to be listed explicitly. This tough requirement leads to an exponential blow-up of the representation of the model, which significantly slows down search and optimization, thus, making an efficient production-based e-commerce an almost impossible problem.

On the other hand, Artificial Intelligence (AI) has evolved as a mature discipline, capable of solving a variety of realistically sized planning and scheduling problems. AI planning systems have gained a definite advantage in computationally intensive problem domains. Whereas previous successful implementations of MRP systems were relying primarily on Operations Research (OR) tools, for example Linear Programming, there are more and more successful examples of AI-based solutions (Ilog, i2, etc.). AI has accumulated enough knowledge to be applied across different engineering areas in a non-traditional manner. Manufacturing modeling, planning and scheduling definitely provide a rich and responsive testing field for innovative AI technologies.

E.Goldradt identified the current manufacturing modeling system as obsolete (Goldratt 1990), even when one considers traditional manufacturing planning driven by forecasts and actual sales orders. The requirement of reactive, on-demand architecture completely rejects existing approaches. Thus, a latest shift from selling a standard suite of products, from a make-to-stock production to mass customization has added new tough requirements on order processing and quoting, on production planning and demand planning.

In this paper we apply a cross-fertilization approach (Smirnov 1997) between AI and manufacturing planning. To be precise, we apply STRIPS-like AI modeling techniques (Filkes & Nilsson 1971) to combine BOMs and Routings in a single model. This novel manufacturing modeling approach allows to avoid the exponential blow-up for re-configurable products and captures the AND/OR inventory item logic without additional modeling efforts. Furthermore, for the introduced type of models, we construct order promising and resource allocation (scheduling) procedures that are linear under a simple, realistic assumption for any homogeneous objective function.

Manufacturing Modeling in the AI Style

AI and manufacturing modeling have a lot of common features. Nonetheless, the latter one still prefers to separate its key modeling items - Bills of Materials (BOMs) and Routings. BOM is a hierarchical collection of inventory items that show all sub-assemblies and raw materials. BOM corresponds to a set of states in AI planning, with manufacturing products being equivalent to goal states in AI, and quantities-on-hand (QOH) - to an initial state. Table 1 presents a short glossary of differences in AI and manufacturing planning terminology.

One of the main drawbacks of BOM/Routings separation is the unnecessary complexity that a modeler has to introduce to maintain inventory items and routing options separately. As we showed in an exam-

Artificial Intelligence	Manufacturing Planning
States	Inventory Items
Goal States (goals)	Products
Actions	Routings
Action Pre-Conditions	Consumable Resources for a Routing
Action Effects	Inventory Items (co-products) of a Routing

Table 1: A Comparison of Key Planning Terms.

ple in Section 1, the number of BOMs that an on-demand manufacturer needs to maintain, grows exponentially on the size of the input. Furthermore, an additional modeling effort is needed to properly represent the logic of alternative Routings and kitting. Some of the inventory items may be built in several alternative way, purchased or transferred from another site. Certain groups of inventory items may form alternative kits that can be substituted only as groups of items. For example, a high-resolution monitor requires a more powerful video card, selection of an IBM's InfoPrint as a faster device for printing the body of a book requires to apply IBM RIPing procedure to convert the Postscript format into an InfoPrint-readable format.

With BOMs on the scale of millions and special grouping (kitting) requirements, manufacturing planning becomes a very inefficient procedure with exponential complexity, when modeled through separate BOMs and Routings. Thus, planning is not the only one to blame, the modeling phase itself transforms input data into BOMs and Routings, number of which may become exponential on the size of the input. Thus, the inefficiency of traditional manufacturing modeling and planning has become an important issue in the Enterprise Management, which has been identified in the literature on manufacturing planning (Goldratt 1990).

On the other hand, AI planning has realized long ago that for certain problem types, it is beneficial to combine states and actions in a single model, thus, preventing the exponential blow-up of the size of the problem domain. We identified that a well-known STRIPS-like modeling approach (Filkes & Nilsson 1971) provides the best fit for manufacturing purposes. Figure 1 shows an example of a Commercial Print Center model that specializes in printing books. States and actions form a directed graph, in which states alternate with actions along any directed path from a book request to a finished product. To proceed with the analysis of benefits of the novel manufacturing modeling approach, we need to state an important resource non-sharing assumption.

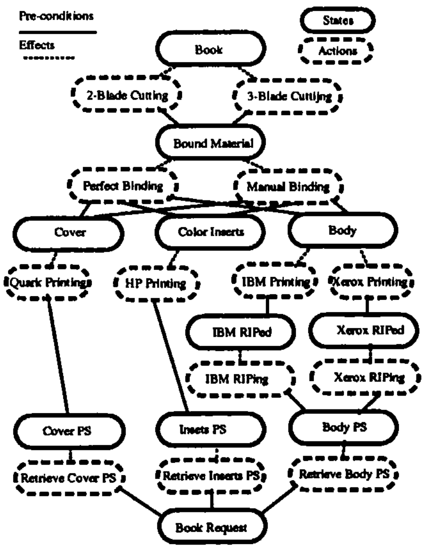

Figure 1: An Example of a Book Publisher Model.

> **Resource Non-Sharing Assumption:** List of all reusable resources that are applicable to the manufacturing process of any particular product does not have repeated entries.

Note that unlike a common assumption about non-repeated entries of inventory items in each BOM adopted by many planning and scheduling software providers, we relax this requirement and allow inventory items to repeat in BOMs as long as they do not induce repeated reusable resources.

Lemma 1 *Under the resource non-sharing assumption, STRIPS-like model captures the AND/OR logic of alternative routings and kitting with the number of states that is linear on the number of different inventory item entries on all BOMs, and the number of actions is less or equal to the number of reusable resources.*

Proof: Each routing contains a list of reusable and consumable resources. Since reusable resources are not shared, the number of actions is equal to the number of routings. If item substitutions are not allowed, then for each routing, consumable resources form the set of required inventory items. Thus, satisfaction of pre-conditions of every action constitute mandatory

("AND") relations. They correspond to di-edges emanating from inventory items (states). Each inventory item (state) is associated with a list of alternative routings (actions) and carries the functionality of an effect of each such action. Thus, the effects of each action (routing) form a set of alternative ("OR" - one out of many) relations.

If item substitution is allowed, then for every allowed substitute one introduces an auxiliary state that denotes one of the substitutes and an auxiliary action that selects a particular substitute. Thus, in the case when substitutes are allowed, the introduced modeling approach keeps track of the AND/OR logic without additional modeling efforts. ■

As Lemma 1 shows, there is no need in listing all BOMs for all possible products. When applied to manufacturing planning, STRIPS-like modeling approach allows to represent the logic of manufacturing processes without explicit exploration of all BOMs and, furthermore, captures the AND/OR logic of alternative routings and kits without any additional modeling efforts.

Order Promising and Production Planning

In Section 3 we introduced a novel manufacturing modeling approach based on AI modeling techniques. This approach resolves the problem of unnecessary exponential growth of the Bills of Materials and captures the logic of alternative routings without additional modeling efforts. In this Section we describe on-line, single-request order promising and production planning algorithms that have linear worst-case complexity. These algorithms utilize the advantages of the newly introduced manufacturing modeling style and may be used as the back-end of an e-commerce system tailored towards re-configurable products.

A modification of a simple, recursive order promising algorithm (OPA) is presented in Table 2, it provides an earliest possible date that a considered order can be satisfied. The difference between State and Action nodes is reflected in the date selection mechanism. Since an action requires all pre-conditions to be satisfied, OPA selects the latest date when all supplies are available. In case of the state node, if there are several alternative actions (alternative routings) that provide the same effect, OPA selects the one that finishes the earliest. *GetEndDate* method takes into consideration all tasks that are currently assigned to a reusable resource that corresponds to the current action node and finds the earliest possible time that a pending production task can be executed.

In the original form, OPA has an exponential worst-case complexity, because the same state and action nodes of the directed graph can be investigated repeatedly. However, this algorithm can be modified to avoid exponential complexity. If one introduces a flag that a node has been investigated, clears it before executing OPA and memorizes the earliest date that an action can be finished or a state can be reached, then repeated recursive requests will not proceed if the flag is already set. If a node has been investigated, the earliest execution date can be obtained from the *earliest* date member of the node object.

```
procedure Order Promising (current node,
                start node, quantity)
  IF (current node.examined)
    return current node.earliest
  IF (current node OF TYPE Action THEN
    Node[ ] predecessors = Predecessors(current node);
    Date earliest date = −∞;
    FOR i:=1 TO #predecessors
      Node predecessor = predecessors[i];
      Date available = Order Promising (predecessor,
                start node, quantity)
      IF available AFTER earliest date THEN
        earliest date = available;
      end
    end
    Time duration = Duration (current node, quantity);
    current node.earliest = GetEndDate(current node,
                earliest date, duration);
    current node.examined = true;
    return current node.earliest;
  end
  IF (current node OF TYPE State THEN
    IF (current node = start node)
      current node.earliest = NOW;
      current node.examined = true;
      return current node.earliest;
    end
  end
  ELSE
    Node[ ] predecessors = Predecessors(current node);
    Date earliest date = ∞;
    FOR i:=1 TO #predecessors
      Node predecessor = predecessors[i];
      Date available = Order Promising (predecessor,
                start node, quantity)
      IF available BEFORE earliest date THEN
        earliest date = available;
      end
    end
    current node.earliest = earliest date;
    current node.examined = true;
    return current node.earliest;
  end
return ∞
```

Table 2: The Single-Request Modified Order Promising Procedure.

Theorem 1 *Under the resource non-sharing assumption, a modified Order Promising algorithm (MOPA) is linear on the number of states and actions.*

Proof: After an edge is traversed for the first time, both of its end nodes have "examined" flags set to true. This memorization "trick" prevents repeated exploration of edges. Hence, a modified Order Promising algorithm is linear on the number of edges. If the number of edges is asymptotically less that the number of nodes, then the entry point is disconnected from the end point and MOPA will figure out that the earliest date is infinite after examining a connected subset of nodes. ■

MOPA outputs only the earliest date, when a current pending request can be satisfied. To perform a complete resource allocation with the earliest start and end dates assigned to tasks according to the best choices of alternative actions, one needs to add another selection flag, memorize the best (chosen) predecessors and perform one more sweep along the model to propagate the selection of best predecessors and assign the dates. This addition converts MOPA into a production planning (scheduling) algorithm - PPA.

Theorem 2 *Under the resource non-sharing assumption, PPA is linear on the number of states and actions.*

In this section we discussed order promising and production planning approaches that find the earliest possible date that a Sales Order (SO) can be furnished. The same approach remains true for many other objective functions besides the earliest availability: Minimal cost, maximal price, maximal profit, etc. We call those objective functions homogeneous. Unfortunately, a mixture of time and price/cost objectives or constraints may lead to an NP-hard class of problem.

Conclusions

We showed that the tradition of separating BOMs and Routings does not always lead to more expressive or compact manufacturing modeling. As the result, order promising and production planning approaches often have to deal with the exponential blow-up of the number of Bills of Materials for re-configurable products. This is an especially hot issue for mass-customization production promoted through e-commerce applications.

If AI modeling techniques are applied to model manufacturing processes, one can often avoid exponential exploration of BOMs and construct extremely efficient on-line order promising and production planning algorithms.

Novel manufacturing modeling approach, order promising and production planning algorithms based on this modeling technique has been filed in 1998 by the author of this paper to the US Beaurau of Patents.

References

Filkes, R., and Nilsson, N. 1971. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2(3-4).

Goldratt, E. 1990. *The Haystack Syndrome.* North River Press.

Smirnov, Y.V. 1997. PhD Thesis. Hybrid Algorithms for On-Line Seach and Combinatorial Optimization. Technical Report CMU-CS-97-171; School of Computer Science, Carnegie Mellon University; 131 pages.

Vollmann, T.; Berry, W.; and Whybark, D. 1992. *Manufacturing Planning and Control Systems.* IRWIN.