FIT Supply Chain Planning

Jurriaan van der Lingen

Fygir logistic information systems
Polakweg 8d
2288 GG Rijswijk
The Netherlands

mailto: jurriaan@fygir.nl
internet: www.fygir.com

Abstract

FIT is an application that enables business engineers to benefit from advanced optimization techniques without the need to become an expert in that field. At the same time, FIT supplies optimization specialists with a problem modeling environment and user interface without the need of becoming an expert in the field of application development. Thirdly, it is an information system with a highly interactive environment that is used for tactical planning of the supply chain at many companies.

At the heart of FIT is an innovative concept of supply chain modelization that bridges the gap between four traditional modeling genres: enterprise models, mathematical models, information object models and classic data models. This paper begins with an introduction to FIT and describes some aspects of this modelization. This presentation ends with a brief survey of mapping a model to ILOG Planner and ILOG Solver.

Introduction

As one of the leading suppliers of production scheduling tools for the process industry, Fygir has long been aware of the problems that can disrupt the smooth running of the whole supply chain. Clients often ask for assistance on issues other than scheduling, especially now that supermarket chains and other influential customers are demanding more flexibility, lower costs and increased logistical performance from manufacturers of consumer packaged goods and the batch process oriented industry.

In respons to this challenge, we formed a new development team to design what we regard as the ultimate planning tool. Combining the optimization techniques ILOG Solver and ILOG Planner with a an advanced application environment, we have come up with one tool that copes with a whole range of supply-chain optimization problems. Simply dragging a different, predefined template into the FIT Modeler enables you to solve a different type of problem, be it plant level planning, multi-site sourcing or Distribution Requirements Planning (DRP) with multiple distribution centers. Advanced users can make these templates themselves.

FIT's modular approach allows for phased implementation. You can begin with the graphical interactive FIT Plan Sheet and later add advanced optimization algorithms, or you can go the other way round. A unique feature is that one model can be used for different optimization techniques. Fygir offers a sophisticated solver based on ILOG's suite of optimization products as standard, but you can replace this with your own algorithms whether they are developed with ILOG or not.

The FIT Integration Kit gives you direct access to virtually any SQL database system and links up effortlessly to ERP systems such as SAP/R3, MFG/Pro, BPCS, Adage and Prism.

"A genuinely innovative planning system that wildly surpasses expectations"

That high praise from the independent US-based Advanced Manufacturing Research marked the recent introduction of FIT. AMR found that FIT's object oriented approach to supply chain optimization set a new standard in the advanced planning software market. And confirmation of their opinion has come with the selection by a number of multinationals of FIT as the preferred package for supply chain optimization. Implementation projects are now demonstrating FIT's power in practice.

A supply chain example

During the presentation at the conference, we illustrate our claims by showing how FIT can be applied to an imaginary multi-site olive oil supply chain. In this paper, we elaborate on this example, which is representative for supply chains in the semi-process industry.

Production consists of three stages:

1. Production of intermediate products (pulp) from raw materials (olives) in a process type environment.
2. Production of end items (types and flavors of olive oil) from these intermediates.
3. Packing of end items in stock keeping units, where each end item may be used for several SKUs (different labels, contents) in a discrete environment.

The different production stages can take place at different locations and intermediate products and end items may be moved from one site to another for further processing. On the supply side, we can distinguish two stages:

A. The procurement of raw materials (crop buying)
B. The distribution of SKUs to customers (supermarket chains, large-scale consumers)

This leads to the following supply chain diagram. Each triangle or box stands for a class, each of which can have many instances (called resources). On the level of instances the supply chain is a network and each arrow in the chain below in reality is a whole bunch of arrows from which one can choose (allocate, source, route). For example, at a certain moment of time a certain quantity of Virgin Oil (an instance of E) may be produced on a site in Spain (an instance of F2) that will be packed somewhere in France (an instance of F3).

[Image]

In FIT Modeler one can define classes at will and indicate which attributes identify the instances. Classes can also have attributes (called properties) that are functions of time. All properties are constrained variables that can bound each other and can be constrained by mathematical expressions. You are completely free to model the classes to fit a particular supply chain. The following is just an example.

RM: Raw Materials. Instances of RM identify the various kinds of olive supplied. These instances may hold information like place of origin, quality and species. Related time-dependent properties include price, availability and stock.

I: Intermediates (pulp). Instances of I contain information like type, specification, recipe, ... that identify the kind of intermediate. Properties that are time dependent are quantity-produced and amount-in-stock are typical time-dependent properties.

E: End items (green oil, black oil, virgin oil). Instances of E contain information that identify end items: type, specification and so forth. Related time-dependent properties include quantity-produced, amount-in-stock, ordered-demand.

P: Packed end items (stock keeping units).

Fi: Various sites acting as a production facility. One physical location can have several instances in the supply chain model: one for each role it performs. A location may, for example, be both a production facility of intermediates (F1) and a production facility of end items (F2).

DC: Distribution center. Instances of DC are the facilities in which SKU's are stored while awaiting distribution to customers. There may be a number of DCs spread over a large geographic area, some of which have an associated packaging facility.

Planning opportunities

Given the above model, here is a list of some of the planning benefits FIT can provide for our imaginary olive oil company.

1. Given a pattern of customer demand / forecast, supply the DC's efficiently, taking transportation and stock-keeping cost as well safety stock levels into account.
2. Given a planned distribution, allocate production (in all three stages) to the various production facilities, taking account of production rate, capacity, costs and labor capacity.
3. Given production allocation, FIT can minimize the transportation cost of intermediates and end items between the production locations.
4. Given price, quality, and availability of intermediates, FIT can choose an optimal recipe for the production of end products.
5. Given the planned production and use of intermediates and the price, availability and quality of olives over the year, establish an optimal purchase plan.

Of course, FIT can search for the solution with the lowest combined cost for all of the five sub-problems at once. However, it can also search for sub-optimal solutions of one or more sub-problems putting the others on hold. For example, in tactical planning the supply of raw materials (5) is probably given during the planning horizon (or at least the immediate future). On the other hand, in a what-if analysis one might start with a forecasted demand for intermediate production in order to arrive at an optimal raw material purchase plan.

As we said, this is just an example. In other supply chain models you will be confronted with other problems and with the benefits associated with their solution. What is important about FIT is that it gives you the flexibility to specify at a high level the type of constraints you want to solve.

Input for FIT

The class model of FIT is an extension of the supply chain model and is maintained and kept in FIT Modeler. Below we will explain what the class model will look like.

The instance model of FIT including routes, recipes, bills of materials, formulas, products, production rates can be read in on a regular basis from an ERP system (if available) with FIT Integration Kit. Alternatively, it can be maintained directly in FIT.

FIT's time-dependent informatio is driven by the demand for end products P at distribution centers DC per time period. We assume that the customer and order quantity determine the DC from which it is delivered. The demand is aggregated from the information in an ERP and/or forecasting system in the interface with these systems. This information is related to the actual (counted) stock of all kinds of products and materials at the various locations and to the future supply of raw materials (and other ingredients).

You can view and edit the time dependent information in FIT Plan sheet, a three-dimensional spreadsheet interface. The three dimensions are resources, properties and time. The cells in the spreadsheet represent values R.p(t) = pR(t) of the property p at resource R at the time t.

Output of FIT

FIT generates output for each product (raw material, intermediate, end item and SKU), for each location (F or DC), and each period t in the planning horizon: the planned quantity produced and/or projected stock level of said product at said location in said time period.

FIT can go on to provide the projected material requirements (flavor additions, packing materials) that follow from this plan. This information can be used by an MRP system to get the orders out for these materials (given lead times, economic order quantity, ..). FIT will not use the availability of these materials as a planning constraint.

Within the production stages, there may be more detailed production steps that can be fed to a scheduling system such as GRIP, Fygir's scheduling solution.

Modeling the supply chain

We will explain the modeling technique step by step. First, consider a basic allocation problem: there is demand for products. These products are kept on stock and can be produced at facilities that have finite capacity. Here is a diagram representing such an allocation problem from the olive oil supply chain.

[Image]

The box and triangle represent enterprise classes from the supply chain. Alongside the classes are the properties attributed to these classes: the demand and stock of the products and the production capacity of the facility.

The circle represents an automatic FIT class. The instances of F3P represent the production at a particular site of a particular product. This class is related to the supply chain classes and inherits (abuse of language - it is not a subclass) their properties.

The diamond represents a scalar attribute of the class relation. Each class has a unit of measurement. For the facility this may be hours of production (hr) and for the product a quantity (pcs). The rate scalar (pcs/hr) converts these to and fro. The lines relating the classes make up linear constraints

1) For all t and all facilities F,
   productionF (t) = Sigma products P = productionP at F (t) * rateF/P (t)

2) For all t and all products P,
   productionP (t) = Sigma facilities F = productionP at F (t).

A generic stock formula relates and constrains the properties of the the two enterprise classes.

3) For all t and all products P,
   stockP (t) = productionP (t) - demandP (t) + stockP (t-1)

This is just an example. You can attribute properties to classes without limit and define constraints on them using mathematical expressions. The automatic classes of FIT allow you to constrain properties across different classes (production on F3 vs demand on P).

The > and < symbol indicate that the property in question is constrained: the amount of intermediate product in stock is not allowed to go below a certain minimum and the production capacity of the facility is finite. FIT Modeler gives you the freedom to model the nature of the constraints: the lower bound of stock may be 0, or a safety stock, or a cover of future demand, or any mathematical expression. The bounds themselves are properties that can vary per time period. There is also a degree of freedom regarding the level at which the constraints are imposed: the facility F may have a general production capacity constraint, or one that varies per product and even per bucket.

4) For all t and all products P,
stockP (t) >= min_stockP (t)

5) For all t and all facilities F,
productionF (t) >= max_capacityF (t)

The $ symbol indicates that the property has a cost associated to it (in $ per unit of measurement). In the example, using the facility has a cost (in $/hr), and storing a product has a cost (in $/ kg). The modeler allows a degree of freedom to indicate which properties have costs and at which level they occur.

This basic example demonstrates what we mean by bridging the gap between enterprise, mathematical, information and data model. From the supply chain (an enterprise model) you only need to name the two classes, name the three properties that go with them and type in a single mathematical formula relating the constraints. This will generate a dynamic set of linear and non-linear constraints (a mathematical model). It automatically creates a user interface for viewing and editing instance data in spreadsheets, table views, inspectors and entry forms (an information model). This is not code generation. You can change and extend your model at the same time you are planning live data and you will see the end user views take shape while you are typing in the modelling formulas. By dragging and dropping icons representing the elements of your model to icons representing a data source, you can hook up immediately to existing ERP systems (with a classic data model).

Dynamic resource model

Those familiar with ILOG Planner will have noticed that the "mathematical" part of basic allocation model example is represented in Planner by the IlcPlan class. You may wonder why three classes are used when one can do the job.
Although the basic allocation model is instructive, it is not very practical. However, you will see that it is a basic building block in the supply chain model. The next model we look at comes closer to a dynamic supply chain. Instead of just producing on demand, we consider what happens when demand comes from different locations (that may or may not be the same as the production facilities) and that the transportation costs and stock levels per storage facility are an issue. The following diagram is another excerpt from the olive oil supply chain.

[Image]

Again, the circles are FIT's automatic classes, the lines are linear relations (with the appropriate constraints) and there is a generic stock formula relating the properties and buckets, as well as a transport formula which states that everything is shipped to a DC as soon as it is produced. The class PDC represents (for each period t, product p and distribution center dc) the amount of product p stored at distribution center dc (waiting to be shipped to the customer). The class F3PDC represents the amount of product p transported from facility f3 to distribution center dc .

Apart from modelling a practical DRP problem without any effort, this example clarifies how you can gradually expand a FIT model to include more and more complexities without losing previous modelling efforts. From here on, we visualize a dynamic resource model diagram as follows.

[Image]

Basic process model.

Another building block. Consider products P to be made using end items E.

[Image]

We assume there is a store of end products that can be drawn from. For the production of an end item e a number of recipes are available. The class EP represents (for each period t) the amount of end product e applied for the production of p. This model represents the problem of finding an optimal mix of formulas, given a production plan for E. The supply has a cost factor associated to it (virtual purchase price, varying over time). The stock is constrained below (and maybe also above). This could involve a safety stock, cover, or some such property. All properties are related by a generic (albeit somewhat different) stock formula. The lines relating the classes have, as always, an associated linear constraint. The inflow factor is a conversion of unit of measurement (kg/pc, ...) combined with an efficiency factor. The constraints and automatic class relations cross-relate the properties of the two enterprise classes.

In short, the Basic Process Model has all the degrees of freedom that the Basic Allocation Model has. In fact, they are identical in terms of FIT components.

Dynamic process model

The basic process model does not include finite capacity. We include this by extending the basic process model in a way that is very similar to dynamic resource model.

[Image]

Again, the circles are FIT's automatic classes. The lower part of the diagram is the basic process model restricted to a facility. The upper part are the enterprise classes with the aggregate information. The F3 production facility class holds the finite capacity constraint. As usual, you can move the constraints and costs to another level. In the next paragraph, we visualize a DRM diagram as follows.

[Image]

Fitting the supply chain together

The prudent observer will have noticed that the DRP and DRM modeling blocks FIT together like a jig saw puzzle. The point to note is that, viewed from the process model, the supply of end products to facilities of type 3 (EF3 as viewed from the DPM) is the same as the demand for end products from facilities of type 3 (EF3 on the customer side in the DRM in the preceding stage). Instead of generating a purchase plan for the end products, we can use the DPM to provide a demand pattern for the preceding DRM (distribution of end products). The supply costs (purchase price from a supplier outside of the supply chain) are then replaced by the actual costs (production and transportation) of that DRM. In this manner one can go on and cover the entire supply chain, as the following diagram illustrates.

[Image]

Using ILOG Planner

Most of you will probably know there is a difference between making a mathematical model and actually solving it. So far we haven't mentioned any non-linear constraints. In fact the LP part of the problem can be solved in negligible time for a single DPM or DRM block in all of the practical cases we have tried.

We have a tried a considerable number of cases where the number of end products ranges from 100 - 1000, the number of allocatable resources ranges from 10 -100 and the number of buckets 13 - 52. In these cases, the total number of constrained variables in the last DRM block can be anywhere between 1000 - 20000. Since the semi-process has a largely divergent production process (many end products, few raw materials), the other blocks become increasingly smaller. Knitting the blocks together does not increase the complexity very much. As a rule of thumb, two similar blocks put together quadruples the time for solving them one by one.

In testing one of the larger single-block problems, we experienced that the LP solver class IlcSimplex that is the heart of ILOG Planner does not perform much better than freeware solvers from the academic community such as lp_solve. It certainly does not match performance of commercial LP solvers. As I understand, ILOG is developing an interior-point LP method to be included in ILOG Planner with a performance that comes close to the best commercial LP solvers available.

In the many cases from various industry types we have considered, we really have encountered only three types of non-linear constraints.

1. The "semi-continuous" constraint ($x = 0$ or $x > C$). Products are often produced in batches and a minimum batch size is required to cover the overhead of setting the machines up.

2. The integer constraint ($x = n * C$, n integer). Some process can only produce in fixed batch sizes. However, the number of variables of the problem that are integer constrained is significantly smaller than the number of real variables (10% at most). The amount-to-produce variable is usually not an integer because even in the discrete stages of the supply chain, the amount is measured in tens of thousands where a non-integer solution is as good as an integer one. Although labor capacity is often a constrained resource it so far has only been used as an upper bound for production and not as an integer constraint. The process and semi-process industry use expensive dedicated machines and relatively expensive materials and requires few human skills that can not be hired on demand.

3. Piecewise linear cost. For example, the price of materials is non-linear in the amount you buy: if you buy in larger quantities you get a reduction. Here the cost function is Fx(x) * x, where Fx is a piecewise linear function, like

Fx(x) = { 100$, if 0 <= x <= 1000

{ 90$, if x > 1000

The function may be different for each variable x . Let's consider a simple example:

minimize: Fx(x)*x + ... ; subject to constraints in x and ... .

Relax the cost function to:

minimize: A*x + b + ... ;
where A is the constant coefficient of the "linear lower bound" function of Fx ;
0 <= b <= B is a variable indicating the error in the cost of x ;
+ an "if-then" constraints further binding b in terms of x;
+ the constraints already there;

[Image]

The relaxed cost is a lower bound for the real cost, and in the branch and bound you can deal with those in a standard manner.

ILOG Planner can collaborate with ILOG Solver to prune the search tree. However, in the situations we encounter where the number of non-linear constraints is significantly smaller than the total number of constraints, the consistency checking takes up so much time that, even though it does prune the search tree, the net effect is a much slower algorithm.

When posing real problems in a single DRM/DPM block with realistic data and non-linear constraints we are on the brink of bad performance, measuring progress in hours rather than seconds, whereas the relaxed problem takes at most a couple of minutes and often just a few seconds. Of course, ILOG Planner is a new product and we only have three months experience working with it. We are confident that we can tweak the algorithm to a good performance on the largest single block-problems.

Unlike LP that increase quadratically with supply chain, the non-linear search tree will reach outer space when chaining the blocks together. It is not realistic to hope for a globally optimal solution on the entire supply chain. We are working on a set of good heuristics for combining sub-optimal single-block solutions into a acceptable global solution.