Blending petroleum products at NZ Refining Company

Geoffrey B. W. Gill
Commercial Department
NZ Refining Company
New Zealand
ggill@nzrc.co.nz

Abstract

There are many petroleum products which New Zealand Refining Company (NZRC) must constantly blend from component tanks. The “operation” at NZRC goes twenty-four hours a days which implies that component tanks are continuously either filling up with a component, or being blended into a final product tank, or neither or both. Products can be made from many different combinations of components, but they must meet specifications (such as 96 octane for premium petrol and 91 octane for regular petrol) and be completed on time for loading onto a ship. Developing a blending schedule is a problem that has features of both a continuous and a batch-type nature. This paper is intended to provide an insight into difficulties that blending possesses and ideas behind developing a good blending schedule.

1 Introduction

Most world-wide Oil Refineries have a core business that can be separated into four distinct processes: Receive, Produce, Blend, and Distribute. The process of Receiving involves getting crude in crude tanks at a Refinery. This includes activities like evaluation and purchasing of crudes, shipping/pipeline logistics and selection of destination crude tanks. Producing is the heart of the Refining business. First the crude gets separated into heavier and lighter fractions when it is processed through a crude distiller. Second the fractions get processed through other process units that refine (improve the chemical and physical properties of) the separated crude fractions. Blending entails combining components (refined crude fractions) in order to obtain petroleum products. Distribution is getting the petroleum products from the Refinery to the market place. Each process on its own is difficult to optimise because of the complexity and number of different options available. Ideally these four processes should be optimised as a whole but the shear enormity of the total combination makes this approach impractical. An overview of the Oil Refining process can be seen in Figure 1 below.

Overview of the Oil Refining Process

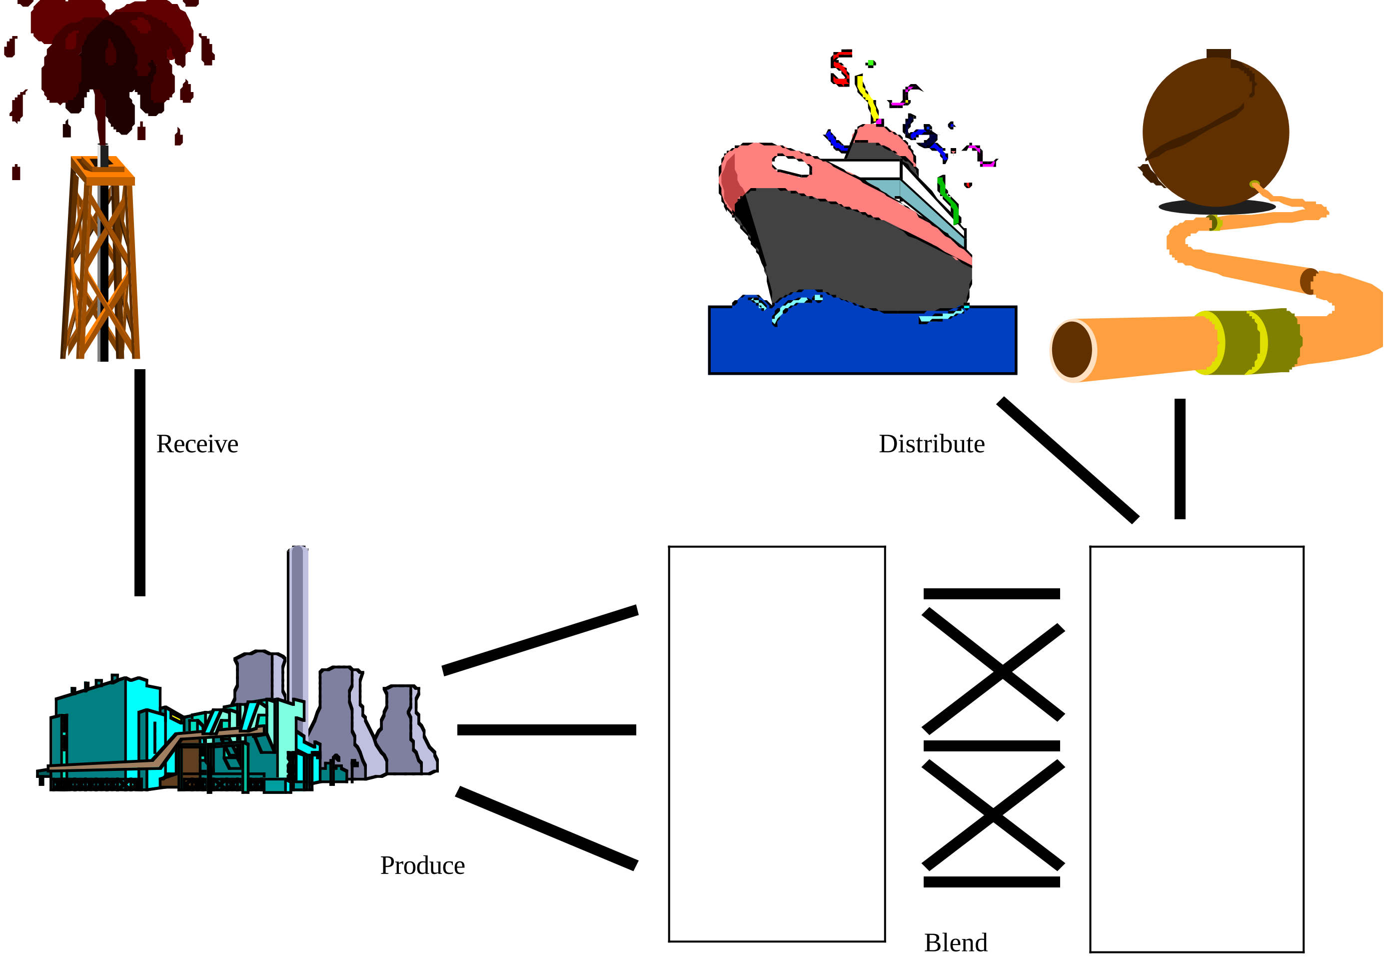

Figure 1

The rest of this paper is going to focus on the blending process. The purpose of the blending process is to obtain petroleum products from refined components that meet certain quality specifications on time for distributing to the market place by ship or pipeline. The objective is to find the combinations of blending recipes that make the “best use” of the components. In this sense the blending problem exhibits features similar to the well known diet problem that can be formulated and solved as an LP. NZRC used to generate blend recipes one at a time by formulating the problem like a diet problem and solving the resulting LP. Although this idea is simple and has worked well in practice it is nevertheless a myopic approach and it runs the risk of unintentionally using a component today that would be far better left for use in a future blend.

It is assumed that the daily quantities and qualities of the component rundowns (the back end of the *Produce* process) are fixed and known. In addition the daily requirements for all products (front end of the *Distribute* process) are also fixed and known. Given this set of “supply and demand” information, the boundaries of the blending problem have been set. A physical description of the problem (from which a mathematical model can be built) will be given along with ideas used in optimising the problem. Finally an assessment of the results given by the model and the model’s limitation will be presented.

2 Description of the blending problem

The quantities and qualities of the component rundowns and product requirements are known on a daily basis. In reality the rundown of all components to component tanks is continuous. Also blending from component tanks to product tanks and lifting from product tanks cannot be performed instantaneously. However the blending model that is used will be formulated in daily granularity.

2.1 Variables

- Each component tank has a set of product tanks where blending can occur into every day. The *movement* (m_{ijk}) is defined as the amount of material transferred from the ith component tank to the jth product tank on the kth day.
- Each product tank has a set of products that can be lifted every day. The *product lifting* (p_{jkl}) is defined as the amount of product l lifted from the jth product tank on the kth day.

2.2 Constraints

i) Each component rundown must go to exactly one component tank every day. (This component rundown schedule is decided upon “a priori.”)

ii) The daily qualities of each component tank are determined by combining the qualities of the component tank on the previous day with any rundown into the component tank on the corresponding day.

iii) The daily qualities of each product tank are determined by combining the qualities of the product tank on the previous day with any movements from any component tanks on the corresponding day.

iv) Each quality is determined by a blend rule (which for some qualities is non-linear).

v) Each component and product tank has a maximum and minimum level which must not be violated on any day.

vi) Each product lifting must be able to be made from one (or more) product tank(s) on the appropriate day and satisfy all quality specifications associated with the product.

vii) Each component rundown has a set of qualities, an economic value in dollars per ton, a quantity and a destination component tank every day.

viii) Each component tank has an initial set of qualities, an initial economic value, and an initial quantity.

ix) Each product tank has an initial set of qualities, an initial economic value, and an initial quantity.

x) Each product has a set of quality specifications, an economic value (selling price), and required quantity every day. Some of these specifications and economic values are in weight while others are in volume.

xi) Due to the daily modelling approximation, a product tank can not be lifted from and blended into on the same day.

2.3 Objective

The objective is a very difficult feature to define accurately in practical terms. There are many different measures of schedule quality which are listed below.

The following are economic drivers and are reliant on accurate predictions of the initial value of component and product tanks as well as the value of the component rundowns.

i) Maximise the value of all components at the end of the time horizon (usually ten days).
ii) Maximise the value of all products at the end of the time horizon.
iii) Maximise the aggregated blend margin between the value of the products blended and the value of the components used.

The other measures of schedule quality are non economic.

i) Minimise the amount of giveaway in products lifted. (Giveaway is defined as the quality of a product being lifted minus the product specification. For example a product lifted of Premium with an octane of 96.3 where the octane specification for Premium is 96.0 has 0.3 giveaway.)
ii) Minimise the total number of movements, or blends. (A blend is a set of movements with the same destination tank and the same day.)
iii) Minimise the deviation from some set of desired ending stock levels.
iv) Minimise the length of time between blends and liftings (ie: a "Just In Time" philosophy)

The dominant objectives that are used at NZRC are maximum value of ending components and maximum value of ending products. In addition there is a smaller drive to minimise the number of blends.

3 Solution process

NZRC have recently implemented a software package called MPMP (Multi Period Multi Product) which is an application built in AIMMS specifically designed for optimising the blend planning problem. First the data is manipulated in such a way so that the problem is ready for optimisation. A snapshot of the exact quantities and qualities of all tanks must be obtained. For the quantities this is not too difficult because all tanks have simple meters which provide the required information. However the qualities have to be predicted in cases where a tank has an upward movement since it was last tested. Once these qualities are obtained, each tank can then have its economic value calculated. The philosophy of the economic calculation is a complex task assuming that there is no market value for components and the only value is a potential one when a component gets blended into a product. When this data is obtained the model can then be generated.

This blend planning problem is a non-linear mixed integer problem, and because these problems are generally difficult to solve "in one hit" there is a three-step solution process which involves solving an LP, then a MIP, and finally a NLMIP.

3.1 LP relaxation

The entire blend planning problem is aggregated into one time period. The consequences of this is that all the component rundowns and the product liftings are aggregated “a priori”. Product tanks are ignored. Any non-linear blend rules are also approximated to linear. The problem is simplified down to one where an orderbook of products has to be blended from all available components. This then becomes very similar in structure to the diet problem.

3.2 MIP relaxation

The solution from the LP is used to help solve the MIP. At this stage all the time dependent information (ie: the daily granulation) and product tanks are modelled. During this phase of the optimisation the scheduling of the blends and the decisions concerning which component tanks to blend into which product tank and which product tanks will be used to satisfy the product demands are optimised. All non-linear blend rules are approximated to linear blending by mass. Typically this phase of the optimisation takes the longest. The MIP is optimised using the commercially available CPLEX solver which is built into the MPMP software package.

3.3 NLP formulation

The tank blending schedule generated in the optimal MIP solution is fixed (all m_{ijk} that are zero in the MIP are not allowed to become greater than zero in the NLP). The blend recipes from the MIP solution are fine tuned to take into account the non-linear blend rules and the complete modelling of the pooling problem which itself is inherently non-linear. Practice has shown that there are seldom infeasibilities encountered during the NLP phase. This is heartening news because it hopefully suggests that the essence of the MIP is fairly close to the NLP. The implication is that there is negligible optimisation opportunity lost in fixing the schedule at the MIP stage. The NLP is optimised using the commercially available OSL solver which is also built into the MPMP software package.

4 Analysis of the model

A model is only an approximation to a real life problem. Due to the approximations a model will have some limitations. This section will identify the approximations made for this blend planning problem and outline the possible consequences.

4.1 Daily granularity

As mentioned earlier, MPMP deals in terms of days. Suppose that a component tank has a rundown assigned to it on a given day and suppose also that the same tank is being used in a blend on that same given day. The question that arises is what the qualities of the material being blended from the component tank should be. One possibility is to assume that the qualities are the same as those that were in the component tank on the previous day, or in another words assume that the blend takes place entirely before the rundown. The second possibility is to assume that the entire rundown takes place before the blend and in this case the resultant qualities obtained from the combination of the rundown qualities and the qualities of the tank on the previous day are used. In reality the true qualities of the material that is going to be blended can lie somewhere between the two extremes above. In a situation where a component tank is being rundown into and blended from simultaneously, the qualities of the material coming out of the component tank would be steadily changing. In addition there is the reality of tank mixers. With all these factors the true qualities of the material in the component tanks is difficult to model let alone optimise with. The assumption made in MPMP is that the rundown always takes place before the blend.

The schedule for the component tank rundowns is decided upon before the optimisation of the blend planning takes place. Unfortunately there are practical situations where component tank ullage is very scarce and a blend has to take place simply so that ullage in a component tank is created and a rundown has some tank to go to. From a modelling point of view, MPMP will sometimes blend from a component tank simply because that component tank lacks the necessary ullage for the rundown that has been assigned to it on the next day. In practice a blend planner has the ability to rundown component into another component tank. In addition it is possible in real life to blend from a component tank for the first half of the day and then rundown into the component tank in the second half of the day, but due to the daily granulation the model does not “see” these types of opportunities.

Similar types of problems occur with modelling the lifting from a product tank and subsequent blending into the product tank on the same day. Fortunately the situation of blending into a product tank and lifting can not happen in reality because the mixture in the product tank needs to be tested and certified as on grade before lifting is allowed.

4.2 Objective functions

This blend planning problem does not finish after 10 days, but continues in real life. In planning a solution to any time dependent scheduling problem that continues after the horizon of the model there are usually two different aspects of the solution process. The firstly aspect is *analysis* where one can look at many different scenarios within the time horizon of the model. The different scenarios can usually be compared with reliable measures. In the blend planning problem such measures can include the amount of giveaway in the products being lifted from the product tanks, or the length of time between a product tank being blended up and lifted from or the number of movements and so on. The important feature of analysis is that the goals are clear and measurable.

The second aspect is *assessment*. The important distinction between analysis and assessment is that the latter involves a static evaluation. At the end of the model horizon it is desirable to plan to end up in a good position. The dilemma that schedulers and planners often face is defining in mathematical terms what a good position to end up in really is. In the blend planning problem we have decided to maximise the value of the contents of all the component and product tanks. In this case the model will do as many blends on the final day as possible whenever there is potential to realise a blending margin. (A blending margin the difference in price between a product and the weighted sum of the value of the components). However the mixture sitting in the product tank on the final day may only have a deceptively high value if further down the track NZRC is not able to sell the mixture at this high value. Earlier in this paper it was said that getting the value of the component tanks correct was very important in order to drive MPMP into obtaining good blend schedules and that will be discussed in the next sub section.

Another attempt at defining a “good” ending position involves quantifying the stock levels. Most people involved with inventory control have a reasonable “feel” for a comfortable level of stock. On one hand, minimum stock can be viewed as optimal because all inventory has a holding cost associated with it. On the other hand, a low level of inventory (or an extremely high level of inventory) carries with it little flexibility or a high level of risk. For example if one of the product tanks that was blended up turned out to be off-grade upon testing, it would need to be fixed up by adding some more of an appropriate component. If an appropriate component was not available then delays to shipping could occur and demurrage (cost incurred due to the delaying of a ship) is a lot more expensive than holding higher than minimum stock levels of components. However, if target stock levels are actually put into the model then there is a risk of the model being forced to blend sub-optimally.

4.3 Pricing of the rundowns and component tanks

There are two factors involved in assessing the value of any commodity: Quality and scarcity. The second factor is very complicated in a blend planning problem where there are continuous rundowns. If there is a shortage of component x initially then it may be thought to be valuable. If this is the case, it may be carefully avoided in blends until a situation develops when there is a surplus of it. After a few days of the same component x suddenly becomes less valuable. At this stage every single blend will try to use component x. Alternatively the value of component x can be fixed depending only on the initial scarcity, but the further into the future we analyse, the less relevant the initial scarcity and corresponding value become. The range of possibilities is large and in most situations placing an economic drive on scarcity can lead the optimisation process along the wrong path. The second factor has proved to be more useful in deriving an economic value.

The building blocks of economic values based on quality are *quality premia*. Every quality has a number associated with it that measures the change in value with a unit change in the quality. For example Octane may have a premia of 2.0. This means that a component with identical qualities as Regular petrol except for an octane of 90 (one less than the octane of Regular petrol) should be valued at $2 less than the value of Regular petrol. These quality premia that are used are based on current world market prices for petroleum products. The tank and rundown values that are generated appear to be reasonable because the blends obtained usually tend to be similar to those that the blend planner would expect.

5 Conclusion

MPMP has recently been in full use at NZRC as well as 5 other Refineries world-wide. It is unusual in the way it makes use of two commercially available solvers. The model gives blend recipes that experienced blend planners agree with. The human can feel comfortable that today’s blend has been optimised taking into account the next 10 days. Due to the constantly changing nature of a refinery and ships that can be delayed at sea, it is not realistic to make a definite plan for the next 10 days. This is somewhat a mixed blessing because the model can provide blend schedules that suffer from the horizon problem. MPMP does not take into account any stochastic modelling and as such has limited use in the evaluation of risk taking. There is still debate concerning what objective drives should be used in MPMP and many practical tests are probably necessary. Perhaps arguably the biggest benefit of MPMP is the ability to look ahead and signal potential problems and opportunities before it is too late in order to take any pro-active action.