Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

INTEGRATED MANUFACTURING LOGISTICS: BYPRODUCTS CAN BE CRITICAL

Charles H. White
Bing W. Tsai

DuPont Engineering Technology
DuPont Company
Wilmington DE, 19898, U.S.A.

ABSTRACT

Production Operations can usefully be partitioned into Discrete Manufacturing and Processing Operations. Discrete Manufacturing plants produce products such as automobiles, airplanes, refrigerators, toasters, computers, and such ‘discrete’ products often have quite large work-forces relative to plants such as refineries, distillers, and chemical plants. Processing Operations systems produce ‘stuff’ such as gasoline, paint, beer, ice cream, and chemicals. Processing Operations can then be classified as ‘continuous’, ‘batch’, and ‘hybrid’ (batch and continuous). Many of Process Operations systems are both capital intensive and use relatively low manpower. Process systems often have a lot of expensive equipment; refineries and chemical plants have a lot of large expensive equipment and control systems. Also relative to most manufacturing systems Process plants have few operators (there are some ‘field operators’ and automated control rooms run by a relatively small number of operators and engineers). ‘Continuity of Operations’ is very important in such plants both for economic reasons and technical reasons; specifically, shutdowns interrupt the continuous flow of production and can cause expensive and time-consuming shut-down and start-up situations.

Another factor that can be an important difference is that of co-products, by-products, and waste streams. In many petro-chemical plants there are natural ratios (think of these as ‘recipes’) of input streams and output streams that can be controlled only over a narrow range. For instance in an oil refinery the crude oil is ‘cracked’ into gasoline, kerosene, and several other products and this can be controlled only over a narrow range. The point is that some kerosene will always be produced when producing gasoline. In many chemical systems producing the main product results also in making such co-products and by-products (these really differ only in the relative value associated with these products) as well as waste. Waste is essentially a by-product with a negative value; we need to spend money to dispose of it. Yields are associated with ‘how much good stuff’ we get relative to the input feed streams. In manufacturing operations this is the pounds of good product per pound of input materials. In typical metalworking operations this might be in the 70-95% range. In chemicals production the yields (recall this is the ratio of good product to the input feed materials) can be very low. In fact in this paper we will be discussing a situation where the production process is a solvent based operation where the yields can be (based on the above definition of yield) in the 5% range. The reason for this is that the ratio of solvent to product ranges from 1:1 up to 20:1. This can result in some unusual logistical problems.

1 INTRODUCTION

In this paper we will be dealing with a situation where the current system has one production line (L1) which ‘campaigns’ between a set of products. For each product there is a solvent type. The solvent is used in the production unit (think of it as a sort of carrying medium) and at the tail-end of the production line the spent solvent is separated out as a separate flow-stream. The product stream flows on to other production units while the spent solvent flows first into a tank-farm (SST) and then on into either solvent recovery (SR) or to waste disposal. The plant has a new solvent recovery system that should be able to handle ‘most of’ the spent solvents produced by this production line. Before the solvent recovery unit was constructed the spent solvents were all sent off-site for either recovery or disposal. Spent solvent travels in special tanks loaded onto tank-trucks; determining the number of these special tanks was a major task for this model.

Now just as this new solvent recovery unit has been coming on-line the situation is becoming more complex. A new second production line (L2) is going to be coming on-line soon. And as a further complication additional new products are being introduced that have new solvents. These new additions really complicated the task of

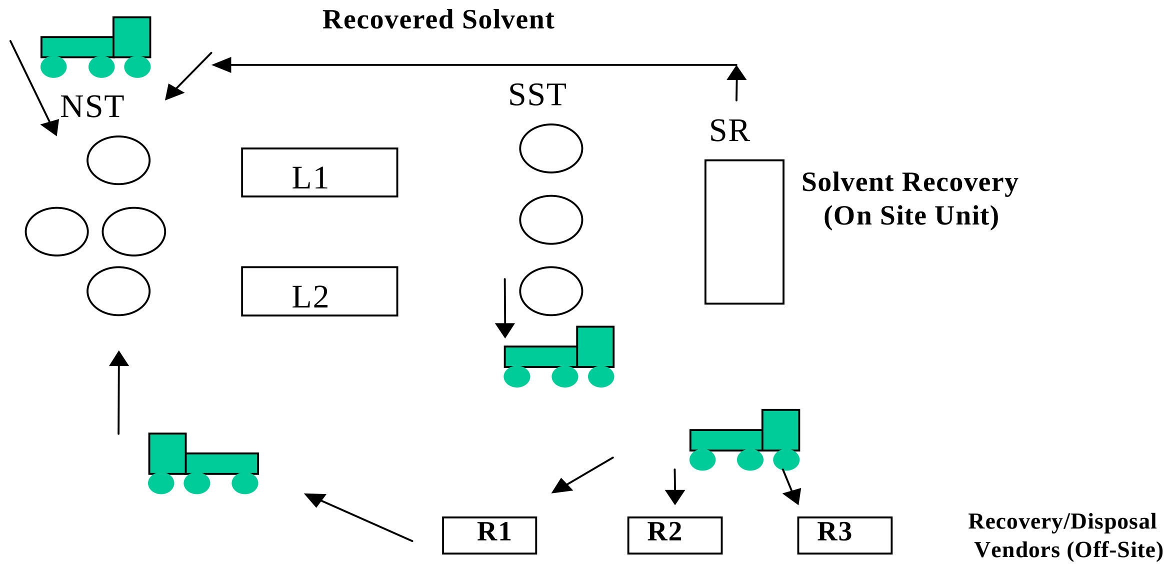

<u>Key Systems Questions:</u>

1. When Line 2 is added, what is the capability of Solvent Recovery to keep up ?
2. When new products and new solvents are added, what is the impact on Solvent Recovery ?
3. As Line 2 starts up and new products and solvents are added, what is the number of Spent Solvent Tanks and Transportation Tanks needed to keep up ?
4. If we can change Solvent Recovery rates and/or transition times, what is the impact ?
5. How large are the dynamic effects due to variability and interruptions ?

Figure 1: Overview of Total System

evaluating whether the solvent recovery unit will be able to handle all of the spent solvent. Will increasing the solvent recovery rate be able to handle any additional needs? Are more spent solvent tanks needed to buffer the solvent recovery unit from the production units?

In this model we needed to deal with the planned production schedules for the two lines. Specifically there is a Production Planning and Scheduling system that considers demand forecasts, planned inventory changes, distribution time-lags, and generates a detailed production schedule for each of the two lines. This schedule can for our purposes be considered to be in the form of two independent lists, each of which contains a sequence of product campaigns with the amount needed and the ‘earliest start’ date for each campaign. Incidentally the line transition times are explicitly considered and appear in the sequence. The production policy is that if we get to a campaign late we immediately do the transition and start production; whereas if we get to the ‘next campaign’ early we must wait (this is a factor that could be manipulated in the dynamic simulation of the production operations).

Figure 1 shows an overview of the total system. It should be noted that the model is at a more detailed level and the model does deal with the specific individual tanks on both the New Solvent side (NST is for New Solvent Tanks) and the Spent Solvent side (SST is for Spent Solvent Tanks). As shown in the overview diagram New Solvent can be (1) new virgin solvent bought from outside vendors, (2) recovered solvent from the on-site solvent recovery unit, or (3) recovered solvent from one of the vendor solvent recovery sites. One of the complications on estimating the number of the special spent solvent tanks needed is the fact that the spent solvent can be shipped to any of three vendor recovery sites (some are DuPont sites with special recovery and/or disposal equipment). Travel times to/from these different sites can differ and all can be variable; hence estimating the ‘average’ travel times for going to and returning from these recovery sites is not a trivial task. The power of simulation is that the model can explicitly deal with each tank-load independently. If we know both what is in the tank and what loads have been recently sent to the three vendors, we have rules for determining where the next load should go; then we have only to sample from the appropriate travel time distribution to get the travel times.

In brief the model needs to be set up so that (1) each of the production units cycles along through its production schedule producing campaigns of known length and thus producing spent solvents at known rates, (2) the spent solvents flow into the appropriate spent solvent tanks, (3) the solvent recovery unit cycles along recovering the spent solvents that are accumulating in the tanks, (4) any excess [or unrecoverable] spent solvent is sent off-site to one of the three vendor sites, (5) new solvent is provided so that the production lines are not forced down to lack of feed solvent, (6) truck tanks are kept at the lowest level possible while still providing for enough tanks that the spent solvent tanks never get too full. On top of that the model needs to explicitly deal with the facts that (a) different products use different solvents, (b) different products run at different rates and have different ratios of solvent to product, (c) the production lines and the solvent recovery unit must have explicit transition times between different campaigns, (d) there are a small number of truck load/unload sites so the number of simultaneous load/unload operations is restricted, (e) there are only a finite number of operators to run the units and make transfers, (f) some units might have random or planned outages, (g) and there might even be 'shift effects' so that some activities cannot take place at certain times of day [or over weekends].

2 SIMULATION TOOL

A wide array of simulation tools is used within DuPont, and they range from internally developed applications to those that are commercially available. Within our Process Engineering Group, we use tools such as Aspen Plus to model steady state processes, and SPEEDUP for modeling transient (dynamic) behaviors during process start-ups. Within Operations Research, which is a sub-group of Process Engineering, we have used several "operations modeling" tools for building dynamic models of plant operations.

ProModel is a discrete-event simulator that is commercially available and is widely used not only within our group, but also across many DuPont plants. Key advantages of using ProModel are in its ability and flexibility to model dynamic and complex processes/operations via the modeling language. In addition, its OLE Automation capability and its built-in features (a standard set of logic codes that can easily be incorporated into a model) allow users to develop robust models in a short time frame.

Another important feature that ProModel provides is the ability to animate processes/operations, real-time. This animation capability not only provides modelers with a tool to verify/validate models, but when used effectively, it is an extremely useful communication tool. It can: (1) convince your clients on the validity of the model and its the impact on the bottom-line results, (2) allow end-users to gain a "world-view" understanding of the whole operation, and (3) assists teams to efficiently evaluate alternatives before any significant capital is spent.

Other tools are available at our disposal. Furthermore, we are frequently evaluating new applications so that when the technology complements our offerings, we will bring the new technology into our group. For the problem at hand, because we are looking specifically at the logistics around solvent recovery (i.e. flow of solvents in- and out-of storage tanks, recovery unit, and loading/unloading stations), and because of our past experience and successes with similar problems, we chose ProModel as the tool for this "operations model".

3 MODEL DETAILS

The simulation is driven by a schedule generated from the Production Planning and Scheduling system. The schedule is produced monthly using revised market and capacity data, and it lists the production campaigns for a two-year horizon for each of the two production lines. Essentially, the schedule lists the products, campaign start-dates for those products, and the production quantity. In addition, the schedule also shows product-related information such as the amount of active ingredient to be used in a production campaign.

Since the schedule provides more information than is needed for the model, only the production-related data (e.g. start-date, production quantity, etc.) is extracted and imported into an Excel spreadsheet. Once in the spreadsheet, we defined this information as "arrivals" in ProModel. As result, when the simulation clock reaches the start-date of a given campaign, an order to start that campaign enters the system at the designated production line that will be "campaigning" the product. If the line is occupied by another campaign, the production order will queue-up and wait for that campaign to complete before making a product transition. Otherwise, if the production line is idle, it will immediately make a transition to start the next production campaign.

3.1 Production Lines

There are two production lines operating independently of each other, with one of the lines being new and is scheduled to go into production in the near future as of this writing. Each line will start production on a packout batch (a packout is used to describe the desired portion of the output), when (1) the line is idle, and (2) the total production quantity has not been met for that campaign.

Although each line consists of series of batch processing steps, and that at any given time, there can be more than one batch going through the line, we have chose not to model this aspect in detail. There are couple reasons: (1) the objective of this model was to provide end-users with a tool to gain a better understanding around the downstream logistics (i.e. solvent recovery and tank-truck requirements), and not the production lines, and (2) the time-frame that we were given was simply too short for us to build a useful model that also included the details around the two production lines. (It might be noted that a detailed operations model of the one existing line was done and the learnings helped guide this modeling effort.) As result, we generalized all individual processing steps into a single process by looking simply at the frequency at which a packout batch exits the production line.

As each batch enters a production line, depending on the product, some ratio of fresh solvents is fed into the process. This fresh solvent ratio ranges anywhere from 1:1 to 10:1 of the packout quantity. The fresh solvents remain in production until the batch processing time for that particular packout has elapsed. Once the packout is processed, some ratio of spent solvents is then produced and transferred into the spent solvent tanks. The ratio for generated spent solvents can range anywhere from 1:1 to 20:1 of the packout quantity.

For both fresh solvent and spent solvent volume ratios, we again created an Excel spreadsheet that gives end-users the flexibility to easily change the ratios as we improve our production processes, and thus lowering the volume of spent solvents generated.

3.2 Tank Farms

The logistics within the tank farm (i.e. flow of solvents from tanks to production lines, to tank trucks, and to other tanks) can get quite complex, especially when environmental concerns are taken into account. On the fresh solvent side, there are a total of eleven (11) tanks dedicated to seven (7) solvents. For each solvent, there is another possible classification between whether a particular solvent is “virgin” or “recovered”. For example, if a particular solvent is recovered, it might be necessary to have two isolated tanks, one for the virgin solvents and the other for recovered, in order to avoid cross contamination between the two (since any cross contamination will lead to shutdown of one or both production lines).

On the spent solvent side, there are a total of eight (8) tanks dedicated to various spent solvents (e.g. solvents for recovery, for incineration, or for off-site reclamation). In addition, if decision is made to start recovering one of the solvents that is currently reclaimed off-site by a contractor, an additional tank will be added to the tank farm, bringing the total number spent solvent tanks to nine (9).

Similar to the logistics on the fresh side, the spent solvent side is just as complex. For example, government regulations specify that if a solvent is recovered on-site, piping connections can not be established between the storage tank(s) and the loading/unloading stations for that particular solvent. Any excess solvents, in this case, will need to be diverted and combined with solvents that are to be incinerated. As one can see, tank capacity issues may arise if (1) both production lines simultaneously are running campaigns that generate large amounts of recoverable spent solvents, and/or (2) the recovery unit is out of service.

Furthermore, in order to address other contamination issues, each storage tank that can load/unload solvents is assigned a unique loading/unloading station. As result, if Solvent A can only be loaded on a particular station (i.e. Station 1) and Solvent C unloaded from that same station, there exists a possibility where both solvents must be processed concurrently in order to avoid shutting down either one of the two production lines. Looking at this subsystem (tank farm operation) as a whole, the logistics of loading/unloading solvents and connecting the loading/unloading stations to tanks can get quite large and complex; these issues must also be addressed with the simulation model.

3.3 Solvent Recovery Unit

The solvent recovery unit is a batch-process unit consisting of two processing steps – distillation and post-treatment. In the first step, a batch-size quantity is sent through the distillation unit at some user-specified rate. When distillation is complete, the unit begins to process the next batch. The distillate (which is the desired fraction from the distillation – this value is data for end-user manipulation) from the column is sent downstream to the post-treatment unit; the heel (the unwanted portion) is sent directly to a tank-truck, rather then a spent solvent tank. (This brings up another interesting logistics problem that will be explained in Section 3.4).

Downstream of the distillation column is the post-treatment unit. The unit will begin processing only when two batches of distillates have been accumulated. Once started, each post-treatment batch moves through the unit at some user-specified rate. At end of the process, the treated (or recovered) solvents are then transferred into fresh tanks based on the solvent type, and whether the tanks have been designated for storing “recovered” solvents only.

3.4 Loading/Unloading Stations

A total of three (3) stations are available for transferring fresh/spent solvents into and from the tank farm. To unload solvents from tanks to tank-trucks, and vice versa, an operator simply needs to connect the appropriate piping between the two. In addition, to prevent cross-contamination, each station has piping established only to a subset of tanks in the tank farm. As result, Station A can load/unload, for example, only organic solvents, while Station B only aqueous solutions. In addition, each station can be docked by only one tank-truck; therefore, each station can only process one type of solvent at any given time.

Having said the above, we can see that the logistics can get quite complex when all tanks in the tank farm are dynamically filling/emptying at the same time (solvents are either being loaded into or unloaded from the system). Furthermore, one must also be concerned with not shutting down the solvent recovery unit because heels from the unit bypasses any storage tanks and are loaded directly onto a tank-truck. Again, instance where two or more tank-trucks are trying to occupy the same loading/unloading station is an issue.

3.5 Shipping Logistics

For spent solvents that can not be recovered on-site, they are either incinerated or reclaimed at contracted facilities. Depending on where the contractors are located, the spent solvents can be sent to the facilities either (1) directly in tank-trucks, or (2) first in trucks using portable containers, then by rail, then lastly by trucks once again (White 1996).

Once a tank-load of spent materials leaves the production plant, the tanks become unavailable for some time typically ranging from three (3) days to over three (3) weeks. During this time period, the tanks are first delivered to and emptied at the contracted facilities, then cleaned at either the contractor sites or transported to other contracted cleaning facilities, before they are returned to the storage area at the production plant.

4 CASE RUNS

A question frequently asked by the end-users has always been whether the existing fleet of tank-trucks will suffice once the new line goes into production? And if not, how many additional tank-trucks would be required in order to keep the processes running?

To help users gain a better understanding of the tank-truck requirements, variables were implemented in the model to monitor the tank-truck usage (i.e. how many are available and how many are currently being used). In addition, model parameters related to the tank-truck logistics were left as inputs that the users can easily change for different case runs. These parameters include (1) the total number of tank-trucks available in the system, (2) the fraction of time a truck-load of spent solvent is sent to a particular contractor for processing, (3) the length of time a tank-truck is unavailable after a contracted facility has been selected and the tank-truck has left the production site (this can be a distribution), (4) solvent tank capacities, and (5) solvent recovery rates and yields. Again, all model parameters are presented to the users in a spreadsheet format for ease of data manipulation.

Various cases were simulated and outputs studied. In cases where the user is interested in the effect of increased production capacity, processing rates (for the products) were varied, and time-series plots for tank levels and tank-truck requirements were plotted. In other cases where the user is interested in the effect of increased solvent recovery capacity (e.g. increasing solvent recovery rate), the recovery rates were varied and the same time-series plot generated. Many other parameters, as well as combination of some parameters, were changed to understand how tank-truck requirements vary. In addition, since production schedules are updated monthly, revisions are made regularly to the production data in order to test how different schedules (i.e. different product sequencing) can affect the tank-truck requirements.

5 OBSERVATIONS AND CONCLUSION

Several interesting observations from the time-series plots. First, we observed that whenever one specific type of solvent is being used in a production campaign, regardless of whether that particular solvent is recovered in-house or not, the need for empty trucks doubles. This is shown with large spikes in the plots of the number of tank-trucks in-use, during periods when campaigns using that particular solvent are in progress.

One major factor contributes to this increase in tank-truck requirement. As stated previously, the ratio of “solvent required” to “product produced” could range from 1:1 up to 10:1, with the ratio of “spent solvent generated” to “product produced” ranging from 1:1 up to 20:1. In this case, the product being manufactured does require a spent solvent to product ratio of 20:1. In addition, simply due to the excess volume of spent solvent produced, the storage tank downstream is limited in how much waste in can hold before having to shut down the production line(s) (holding capacity for this particular solvent is approximately two production batches). Therefore, on one hand, if spent solvent is generated at a rate faster than the time it takes to load a tank-truck, sending that tank-truck to the contractor, and eventually getting it back, then it is likely that the number of tank-trucks needed will be higher than necessary. On the other hand, there is still the question of whether recovering the solvent in-house will alleviate the need for additional tanks. It turns out that, in this case, without making any improvements to the solvent recovery rate, if a decision is made to recover this solvent in-house, the number of tank-trucks needed would remain the same as if the solvent is not recovered.

A second interesting observation is that even though the plant recovery unit is capable of making changeovers to recover several different solvents, there are times when recoverable spent solvents must be sent to contractors for reclamation. The reasons are due to (1) both production lines running large-quantity campaigns that generate two different recoverable solvents, and as result, one of the two storage tanks downstream becomes full while the other solvent is being recovered, and (2) equipment outages force the recovery unit to go down, which causes the upstream storage tanks to fill up as well. The consequences of sending recoverable solvents off-site, of course, is the need for additional tank-trucks, which results in higher operating cost for the plant.

In conclusion, simulation modeling provided the users with a tool to better understand the complex interactions between the operations and the logistics involved in keeping the operations running (refer to additional references and case studies in White, 1996). The model also surfaced hidden operating issues that would not otherwise have been realized, and allowed the users to address these issues by making updates or modifications to the model without any added cost. In the above example, production schedules with different product wheels (e.g. running two large campaigns sequentially rather than concurrently) could be tested to see if the number of tank-trucks required can be lowered. The model could also be used to test whether adding tank capacities, shortening the recovery transitions, or improving recovery rates would further minimize the tank-truck requirements, all without the added cost of testing them in the real system.

At the time this paper is being written, we are beginning another study using an expanded version of this model to determine the feasibility of bringing new products onto this plant. This follow-on study not only brings credibility to operations modeling using discrete-event simulation, but it also shows how good modeling opens new opportunities to improve other areas/aspects of an existing operation; thus benefiting the overall business.

REFERENCES

White, C. H. 1996, *Distribution Logistics in the Process Industries: Railcar Requirements*, 1996 Proceedings of the Winter Simulation Conference.

AUTHOR BIOGRAPHIES

CHARLES H. WHITE is a Principal Consultant in Operations Research with the Process Engineering Section of the DuPont Company. He has 30+ years experience in Consulting and Modeling with a strong focus on Simulation and Mathematical Programming as applied to Business Planning, Plant Design, Research Guidance, Productivity/Asset Improvement, and Distribution Planning. Dr. White has degrees in Physics, Engineering, Mathematics, and Operations Research all from the University of Michigan. During his 30 years with DuPont he has also taught courses at the University of Michigan, the University of Delaware, DuPont, General Electric, and several short courses through the AIChE. He is a member of IIE, Informs, and AIChE. For several years was the DuPont representative to the Materials Handling Research Center at Georgia Tech and was Business Chair for WSC in 1993.

BING W. TSAI is an Operations Research consultant with the Process Engineering Section of the DuPont Company. He holds a B.Sci. in Chemical Engineering from the University of California at Berkeley, and a M.Eng. in Operations Research from Cornell University. Prior to joining DuPont, he worked as a Systems Engineer at BOC Gases. He is a member of IIE, Informs, and APICS.