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

SIZING INDUSTRIAL RAIL CAR FLEETS USING DISCRETE-EVENT SIMULATION

William R. Lesyna

E. I. DuPont & Co., Inc.
DuPont Engineering Technology
1007 Market Street
Wilmington, DE 19898, U.S.A.

ABSTRACT

DuPont has many products that use rail cars in various portions of their supply chains. Often these cars are used to deliver final products to a variety of customers at different geographical locations. In many cases it is difficult to optimally size these fleets, since the underlying system is complex, dynamic, and involves random variables.

This paper describes how DuPont has used discrete-event simulation ("DES") to optimally size an industrial rail car fleet used to deliver final products to customers. It explains why it is important to DuPont to optimize the size of our rail car fleets; how such fleets are sized without DES; the value of DES in modeling one particular rail car system; and some of the lessons from building such DES models.

1 RAIL CARS IN DUPONT

DuPont owns or leases about 7,000 rail cars, which would cost about $700 million if bought new. Corporate direction is to own rather than lease these cars. The financial metric called "RONA" (Return on Net Assets) is important in DuPont and causes businesses to want to minimize any investment that does not directly produce revenues and therefore earnings. Hence, the businesses want to minimize capital expenditures such as those required to purchase rail cars.

The Nature of the Problem: The problem being addressed here is: How do we optimize the size of the fleet, given various input parameters such as plant production rates, customer demand rates, hold times at customer sites, and rail car transit times between the production and consumption sites? This is a planning problem rather than a scheduling problem. That is, the question is "How big should our fleet be next year?"; not "How do we schedule our fleet next week?" or "What are the optimal routes for our rail cars?". We have other tools available for the scheduling and route optimization problems.

As part of optimizing the size of the fleet, we also want to specify the best "policies" for managing the fleet. Policy decisions concern issues such as when to use trucks to supplement the rail fleet, how and when to schedule planned production facility shutdowns, willingness to allow customers to hold cars at their sites as part of their raw material inventory, and sharing of fleets between compatible products.

The Nature of the Solution: The models we have built in DuPont to address these problems allow many "what if?" questions to be answered. They are descriptive models that are not inherently optimizing. Optimizing the size of the fleet requires running a series of "what if?" cases, involving the judgment of a modeler and of a "user" who is well familiar with the various issues involving the fleet.

2 SIZING OF RAIL FLEETS WITHOUT SIMULATION

The typical planner who must size a rail car fleet without DES does so using a spreadsheet or possibly an optimization tool such as Linear Programming (LP). While these tools are extremely useful for many problems, they are almost always unsuitable for sizing a rail car fleet since they cannot adequately address several essential features of the system, namely:

Complexity: Typically there will be various "rules" or "logic" that must be followed. An example could be: "Start using trucks to supplement rail cars for delivery of product to customer A when the days supply for that customer drops below some specified level". Such rules can be easily included in a DES model, but are impractical to include in a spreadsheet or LP model.

Dynamics: DES by its nature includes dynamic features of the system, which are very difficult to include in a spreadsheet or LP model. For sizing a fleet, the dynamics of the system are often the whole story. Knowing how many rail cars are full "on average" over the course of a year is of little value, and in fact is likely to be misleading. If you are trying to size a fleet, you are probably very interested in predicting minimum or maximum values (e.g., "Did we ever 'starve' a particular customer?", or "Did we ever run out of empty cars to fill at our production site, thereby threatening the ability of the plant to continue to run?"). A specific example illustrating the importance of understanding the system dynamics will be given below.

Randomness: Real-world rail systems include many random variables. Some examples are: transit times, both on-site and between sites; rail car maintenance times; production and consumption rate changes due to random failures; and production quality problems that may require rail cars to be used for intermediate storage while problems are corrected or product is reworked.

Pressures on the Planner: While the planner responsible for sizing a fleet faces pressure to minimize the size of the fleet to avoid capital and related expenditures, there is often a greater driving force: namely, don't let a lack of empty rail cars ever cause the production facility to have to curtail production, and don't let a lack of full rail cars ever cause a customer to have to curtail consumption. The penalties associated with curtailing production and consumption are often large, and when such a curtailment is due to a shortage of rail cars the cause is quite obvious. As a result, a planner who does not have DES available to size a fleet is very likely to size it conservatively, based on "worst case" assumptions. In practice, we have found such conservatism in a number of fleets in DuPont.

3 VALUE OF DES FOR SIZING A PARTICULAR "SIMPLE" RAIL CAR SYSTEM

Many DES modeling tools are available and many would be applicable for the type of model described in this paper. In DuPont we have found ProModel to be a cost-effective tool, and we have used it for many such applications, including the one described below.

We have sometimes been surprised at how valuable DES can be even for what appears to be a simple rail car system. In one case, a single production facility with only two customers was found to have a fleet that was considerably over-sized. The combination of complex rules, dynamics, and randomness made the analysis of this fleet impossible without DES.

Value of Time Series Plots: This model provides a good example of the value of time series plots for such applications, and of how misleading the use of average values can be. Figure 1 shows a time series plot from the model for empty cars at the production site, over a period of 730 days (2 years). While the average number of cars over this 2-year period is about 80, that provides very little useful information to a planner who is trying to size the fleet. Instead, what is important is what happens at the peaks and valleys. When we hit a maximum for empty

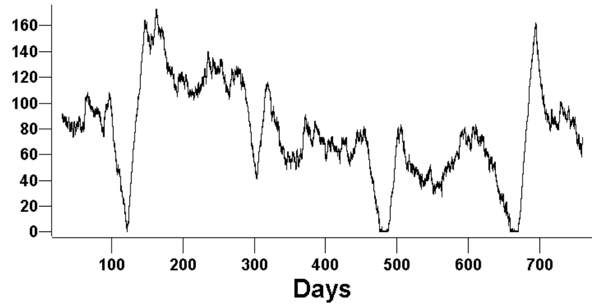

Figure 1: Number of Empty Cars at Production Site

cars, the number of full cars is at a minimum, which may cause problems at the consuming sites. Hitting a minimum number of empty cars (especially at or near zero) at the producing site was especially important and will be covered in more detail below. One of the advantages of DES is that we not only know how the level of empty cars varies over time, but we can focus on the behavior of the system at the critical points in time to greatly increase our understanding of how the system really works. The animation feature included in most DES modeling tools is very useful for this purpose.

Supply of Empty Cars: In this case, the production site was very sensitive to running low on empty cars needed to deliver the finished product, since that carries with it the potential for having to offload the product into SuperSacks that must later be loaded back into rail cars. This of course costs more and takes longer than loading directly into rail cars. Furthermore, running low on empties at the producing site was seen as a harbinger of later starving the consuming sites. As a result, when the number of empty cars at the consuming site was reduced to a specified level, the site would begin using trucks to deliver product to one of the consuming sites. The reasoning was that this scheme would reduce the need for rail cars in that loop and thus, in time, help improve the supply of empty cars at the producing site.

A Surprise About Trucking: This policy of using trucks to supplement the rail cars in certain circumstances had been in place for some time. Once the DES model was available, we were able to run cases to see how well this specific policy worked with different parameters (e.g., "At what level of empty cars should we begin trucking?"). We were surprised when the summary numbers from our case runs showed that system operation worsened when we used the trucking policy, and our first thought was that there must be an error in the model logic.

Using the animation feature built into ProModel, we were able to observe the system in operation and learn why our trucking policy did not have the desired effect. While taking cars out of a delivery loop seemed like a good idea, an unintended effect of that policy is that fewer empty cars were then being returned from that customer site, which was now getting part of its feed via trucks. Our analysis led to the discovery that the implementation of this particular trucking policy was in fact counterproductive.

Importance of Vantage Point: Another insight that came from watching the animation is that running low on empties at the producing site is not a problem as long as there is a steady supply of empties on the way back to the producing site. If your vantage point is simply looking at a staging area for empty cars and you see that there are very few cars at a given point in time, you may take action to solve what you think is a looming problem. However, if your vantage point is expanded so you can see more of the system – namely, the complete loops of empty cars being returned to the producing site, you may find that the future supply of empties is in fact assured, since they will arrive before they are actually needed. To put it another way, even a staging area that is completely empty is not necessarily a problem – if empty cars are being returned just as fast as they are needed.

Sizing the Fleet: After gaining insights into various aspects of the system behavior through observing animations and analyzing time series plots and model output reports, we came to the exciting work of running cases to size the fleet. This involved gradually lowering the number of rail cars in the fleet to see how small we could make the fleet while still obtaining acceptable results. This number of rail cars turned out to be considerably lower than what was originally expected.

Significant changes were made to the actual fleet based on the results of this model, including the removal of $2.3 million of rail car purchases from a previously authorized production expansion project. Subsequent experience proved that the model predictions were accurate.

4 THE HARD PART OF SIMULATION MODELING

While the value of DES is apparent for this type of application, there are many difficulties in building a model for a real-world fleet.

What to model and at what level: Two extremes must be avoided. At one extreme, it is possible to include too much detail in a model. The results can be that the model takes too long and costs too much to develop, and that case runs may require inordinate amounts of computer time. At the other extreme, if the model does not include sufficient detail, the results will not be valid. As a generalization, models are not right or wrong; rather, they are more or less useful in answering real-world questions. The level of detail included in the model is an important factor in how useful it will be.

To help decide whether or not to include some particular level of detail in a model, the most important question to ask is: "Will it significantly affect the answer we are looking for?". In other words, always keep your objective in mind. We have also found that previous experience with DES modeling can lead to good intuition about the appropriate level of modeling to be done. There is some element of "art" here, along with the science.

Clearly we must model more than the rail cars themselves to be able to size the fleet, and here the question of what to model and at what level becomes crucial. We must model some portions of the various interfaces of the fleet. Especially important are decisions about how to model the production facilities, the consumption facilities, and various auxiliary operations, such as trucks, barges, etc., which may interface with the rail car fleet. The decisions about how far to take this modeling must always come back to the fact that our goal is to size the rail car fleet. If, for example, our goal was to determine how to optimize the production capability of the plant that feeds the fleet, the production facilities themselves would likely be modeled very differently. In some cases, the scope may include optimizing the plant production capability along with the size of the fleet. This can be done, but it significantly increases the modeling degree of difficulty. Decomposition of this problem may be helpful, but will not always be appropriate.

Testing: Testing the validity of the model is vital. The first step is to run a "base case" that can duplicate the results from some recent period of operation of the actual system. Then, when "what if?" cases are run, it is important to examine the summary results, the animation, and the time series plots to ensure that the model is operating as intended. It is very easy to assume, once a base case has been validated, that future case runs will all be valid. But changes in input parameters can cause different model logic to be executed and the modeler must ensure that the model is still valid. This work can be tedious, but it is also essential.

"Mass Balance": This type of DES model must have some "driving force", and the modeler must be careful to understand what the force is and how it will operate. The consequence of failing to do this is that the model may not have a "mass balance", meaning that material can build up without bound somewhere in the system, or that some part of the system may starve for material. While this is only one of many kinds of errors that must be rooted out during model validation, it is so fundamental and common that it is singled out here.

As an example, if the input to a model includes both production rates and consumption rates, there must be a mechanism to balance these as the simulation progresses. One way to do this is to adjust production rates based on the dynamic inventory level of finished product, throttling back when it is high, and increasing rates when it is low. In other cases the production facility may be such that its rate is not readily or economically changed, and there must be some outlet available for production that exceeds the demands of the normal consumption sites.

The data: A great deal of input data is required for this type of DES model and the model results can be no better than the quality of the data used to drive it Once again, however, judgment is needed as to the level of detail that should be included. It would be easy to spend an inordinate amount of timing refining certain data, when such refinement may have little noticeable impact on the quality of the ultimate recommendations resulting from the model runs.

Usually, collaboration with others will be needed to collect and analyze this data. For example, in DuPont we have a corporate logistics organization that has access to historical data on the transit times between sites for all of our rail fleets. Help from this organization to sort through the data to find the information needed for the model is essential. Without this data source, it would be very difficult to adequately model the rail transit times, which are often a source of substantial variation. In fact, it is not uncommon that this actual data shows much more variation than what the rail fleet planners thought it did.

5 THE PAYOFF

The payoff from the study of the rail car system described above was considerable:

Avoided capital: Before the model was built, the fleet was believed to be fully utilized, and even required the occasional use of trucks to supplement the rail car fleet. When a project was authorized to increase plant production, $2.3 million of capital funds were included to expand the rail car fleet to handle the higher production. However, as a direct result of the DES model, this $2.3 million for expansion of the fleet was removed from the project.

Reduced operating costs: As explained above, the DES model gave us new insights into the effect of trucking on how the system really operates. As a result, trucking is used less frequently, which results in reduced operating costs. Furthermore, the operating costs (e.g., maintenance) associated with increasing the number of cars in the fleet were avoided.

Improved customer service: Because we now have a much better understanding of how the rail car system actually operates, we are better able to ensure that the supply of full cars to our customers is not interrupted, in spite of various disturbances to the system and in spite of having fewer cars than originally planned.

6 SUMMARY

DuPont has a very large investment in rail cars. Without discrete-event simulation, we have found that rail car fleets are often sized conservatively, since the underlying system is complex, dynamic, and involves random variables. With discrete-event simulation, we can address all of these issues and have profitably applied such models to optimally size rail car fleets. Some of the factors that must considered to make such models successful have been described.

AUTHOR BIOGRAPHY

WILLIAM R. LESYNA is a Principal Consultant in DuPont's Engineering Technology organization. He received his M.S.O.R. degree from the University of Arkansas, and his B.S.E.E. degree from the University of Notre Dame. He has worked for Western Electric in St. Louis, and has spent his last 26 years working for DuPont in Wilmington, Delaware and Nashville, Tennessee. After joining DuPont, he spent 8 years as a consultant in Engineering's OR Group, working with many different businesses and plants. That was followed by 15 years in a wide variety of management jobs, both in businesses and in the corporate Engineering organization. At the end of 1995, he returned to DuPont Engineering's OR consulting group. During his career he has developed many successful OR applications, primarily simulation and optimization models, mostly applied to manufacturing and logistics. He is a member of INFORMS, including the College on Simulation.