Coal-Tipple Operations

John Petty
Ray Eason
Jon Rufenacht
U.S. Military Academy
West Point, NY 10996

Advisor: Charles G. Clark, Jr.

Solution Approach

We formulate a computer simulation in order to calculate operation cost projections. The simulation produces random weekly train schedules and evaluates these schedules with a Pascal program. The simulation has 10,000 weeks of iterations, and is supported by five submodels that aid in arriving at the information required by the Aspen-Boulder management.

Results

1.	Annual cost projection	$\$89,817,000$
2.	Monthly demurrage projection	$\$3,053,000$

Recommendations to Management

- We have formulated an ideal train schedule that minimizes demurrage costs. We recommend that this schedule be implemented in order to make the system more cost effective.
- In order to decrease demurrage costs, use two crews if a train must wait for the tipple to be filled.
- A third loading crew will reduce annual operating costs (cf. Submodel 2).
- This single tipple system can handle a fourth standard train every day. However, we caution that the system will "lag" on Thursdays.
- We recommend that our simulation be verified by comparison with data.
- We recommend that our simulation be used to explore the expected value of hypothetical scenarios.
- We recommend that our simulation be revised to model the operation more realistically.

The Model

Model Assumptions

- Upon arrival to the tipple system, all trains are completely empty.
- All train loading rates remain constant.
- The rate of tipple filling per crew remains constant, so adding a second crew doubles the overall rate.
- Assume that all events are discrete (i.e., trains can arrive only on the hour).

Model Formulation

In order to model the coal operation, we first performed a few hand calculations to get exactly how the operation behaved. After only a few trials, we discovered that many different arrival schedules are possible. It would not be feasible to enumerate every possible train arrival schedule by hand.

Another source of confusion is that the problem does not specify the exact arrival rate of the trains. Without an arrival rate, we were not able to construct easily an algebraic formula or linear program to model the behavior of the operation. Therefore, we decided upon the simulation approach. With the aid of an algorithm and a randomly generated train schedule, we could evaluate every possible scenario. The simulation assists in answering the questions of finding the expected annual costs of the tipple operation and the expected monthly demurrage costs. To assist with answering the other questions, we used five other submodels.

- Simulation (Main Model)
  - Task 1: Random Weekly Schedule Array
  - Task 2: Operation of System and Cost Evaluation
- Hand-calculated Submodels
  - Submodel 1: The Ideal Scenario
    - non-overlapping service times
    - no high-capacity train arrives
    - one crew fills tipple
  - Submodel 2: Minimizing Tipple Loading Costs
    - this is an algebraic optimization formula

- Submodel 3: Minimizing Demurrage Costs
  - look at multiple train arrivals
- Submodel 4: Four Trains?
  - consider four standard trains arriving per day
- Submodel 5: Worst-Case Scenario
  - look at maximized demurrage costs on a Thursday when all trains arrive at the same time

Simulation

The algorithm of the simulation will

- produce a random weekly train schedule;
- evaluate costs for this schedule;
- repeat step (1) for n iterations; and
- tabulate the expected weekly, monthly, and annual costs.

Tasks of the Simulation Model

- generate a random weekly train schedule, and
- evaluate weekly train schedule to model operation behavior and calculate costs.

Simulation Task 1: Random Weekly Schedule Array

Concept: Produce a random variable Y to slot train arrivals within the 168 hours of a week. Produce a random variable X to slot train arrivals of the high-capacity train.

Model formulation: In considering different modeling alternatives, we quickly decided upon a simulation model, because costs are based solely upon the probabilities of trains showing up at crucial times. To a great extent, once a train shows up, the system deals with that train in an extremely deterministic fashion. Since we have no underlying distribution, we found that a simulation model would probably give us the best estimate for an expected value of costs incurred.

Our first step was to determine the probabilities of trains showing up. We decided to look at each day in a global or outside-of-time viewpoint. If we tried to keep track of probabilities as we progressed through the hours of the day, we found that probabilities of train arrival from hour to hour changed, due to laws of conditional probability. However, when observed from the global viewpoint, in each 15-hr window of opportunity for a train to arrive each day, there was a 1-in-15 chance that a given train would show up within a given 1-hr window of time.

Since we decided to consider arrival probabilities on a train-by-train basis, we quickly decided that in evaluating the costs of one day's business, it would simplify matters greatly if we could extract all matters of probability to the beginning of our daily cost evaluations. Another consideration we had to take into account was that days of operation overlap, so that sometimes trains show up one day and don't leave until the next. Therefore, for purposes of observing costs over time, we needed to consider a longer period of time than just one day. For this reason, we extended our schedule length out to one week, or 168 hrs. In each hour, we needed to check to see if any trains arrive and then perform the necessary actions in reaction to the current state of the system. But since we wanted to extract the probabilities to the beginning of the model, we decided to generate ahead of time, independent of cost and time evaluations, a list of times that the trains arrive during a one-week time period.

For each train, this was a simple matter of generating a random number between 1 and 15 and "slotting" it into the appropriate window. For instance, for the first of three trains to arrive on day one of the week, if a 7 was generated, then a train arrived during the seventh hour after 0500, or between 1100 and 1200 hrs. For the purposes of our simulation, we simplified our model to say that any train that arrived during that time period would arrive at 1100 on the hour. This gave us a discrete set, from which we could run a discrete simulation. We continued this slotting until the entire arrival schedule for a week was completed.

Random Variables:

$$\begin{aligned} Y &= \text{a discrete uniform random variable on } 0, \ldots, 14 \\ X &= \text{a continuous uniform random variable on } [0,1] \end{aligned}$$

Assumptions:

- The distribution of arrival times is unknown.
- The random number generator of Turbo Pascal is sufficiently random.

Array Production Algorithm:

- Step 1: Generate a random number Y between 0 and 14 (15 hrs is the window of time between 5 A.M. and 8 P.M.).
- Step 2: Add 6 to Y. (We add 6 to Y because, relative to our timeline, a train cannot arrive before the 6th hour.)
- Step 3: Add a train to the number of trains generated for that hour. Designate a standard train with the value of 1 and a high-capacity train

with the value of 10 (an arbitrary choice). We have "tagged" the hour block by manipulating (6+Y). This "tagged" slot represents the hour that a train arrives. We must now add a value to this slot (1 for a standard train) to signify a train arrival.

- Step 4: Repeat the last three steps three times to slot three trains for that day.
- Steps 5–10: These steps simulate the slotting of the weekly train schedules for the rest of the 168-hr week (we save the high-capacity arrival for Step 11).
- Step 11: Generate a random number X between 0 and 1, add 84, and slot the arrival of the high-capacity train. (On our timeline, the high-capacity train can either arrive during the 84th or 85th hour: evaluate any arrival during the hour blocks at the beginning of the hour up to and including the end of the hour.)

Simulation Task 2: Operation of System and Cost Evaluation

Concept: Evaluate the weekly random schedule in a simulation that calculates labor and demurrage costs.

Model formulation: Once we determine our random arrival schedule, we have to evaluate the week of arrivals as it was generated. The algorithm proceeds hour by hour.

Simplifying Assumptions:

- If trains are waiting to be serviced and a high-capacity train arrives, then the high-capacity train will be loaded first because it has the highest demurrage cost.
- If no trains are waiting in line and no train is being serviced, then it is advantageous to fill the tipple with two crews.
- If you make a train wait because of an empty tipple, then it is more cost-effective to use two crews, in order to minimize demurrage costs (refer to Model 3 for a proof of this simplifying assumption).

Algorithm: (see **Figure 1**)

- Step 1: Check to see if any trains arrived at this hour. If so, how many trains are in the system, and is one of them the high-capacity train? If the high-capacity train arrives: Go to Step 7.
- Step 2: If three standard trains are in the system, at least two are going to have to wait on hold and incur demurrage costs. Add $30,000 to cost for this hour's waiting. Go to Step 5.

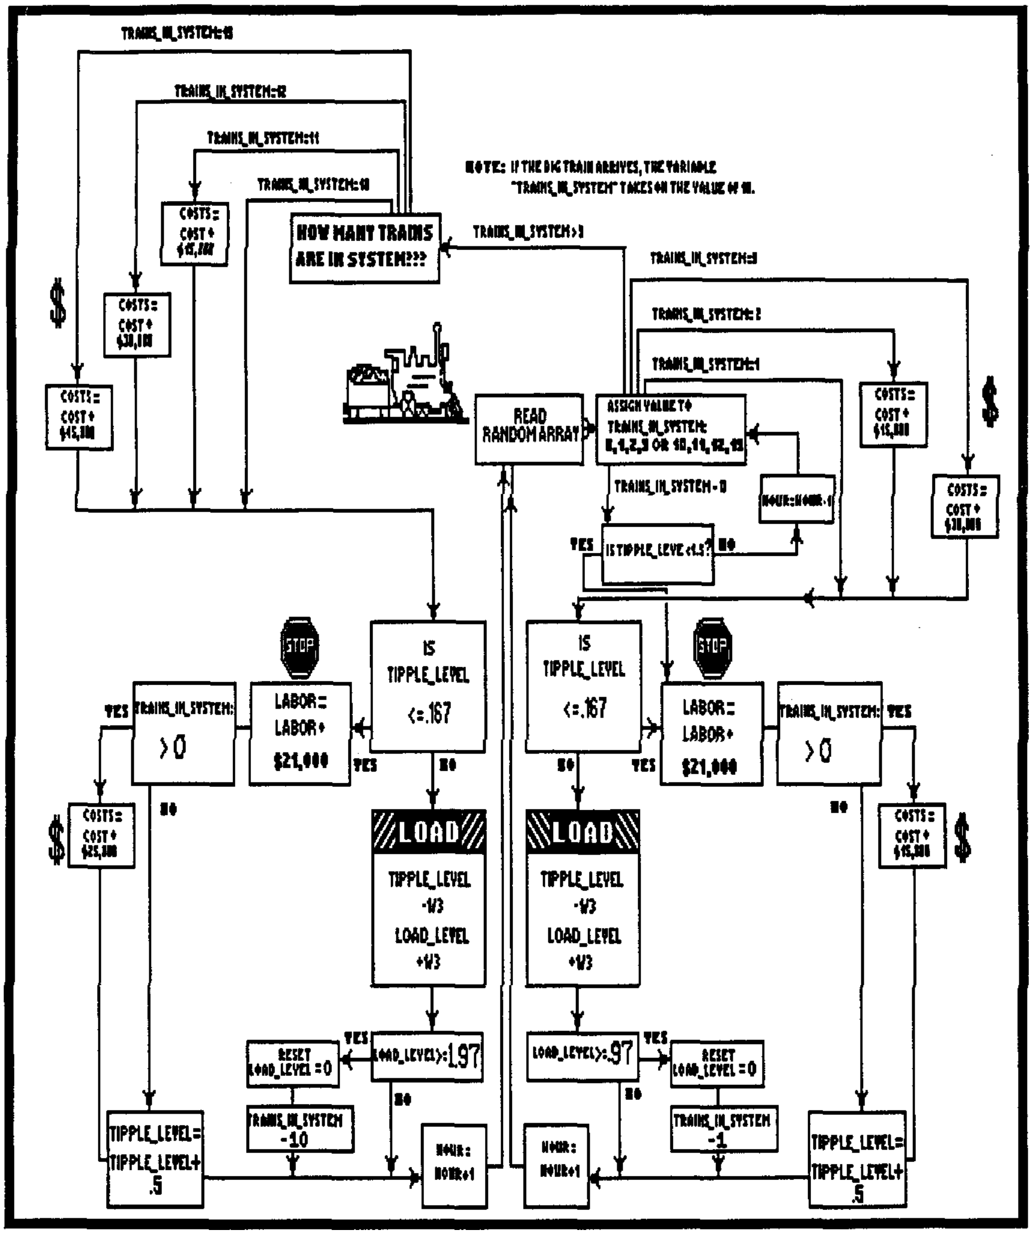

Figure 1. Algorithm for the model.

- Step 3: If two trains are in the system, at least one will have to wait and incur demurrage costs. Add $15,000 to costs for this hour's waiting. Go to Step 5.
- Step 4: If no trains are currently in the system, then let the system be idle, except check to see if the tipple is not at full capacity. If it is not at full capacity, then begin filling it. In 1 hr, two crews can fill one-half trainload into the tipple. Add $11,000 to cost for working crews if they fill the tipple during this hour. Go back and check for next hour to see if trains arrive. Proceed to next hour. Go to Step 1.
- Step 5: By now, we must see if we can actually fill the train that is on-line to be filled. If so, then take one-third of a trainload away from the tipple and put it into the train. If the train is full, send it on its way and reset to 0 the load value of the current train being loaded. Otherwise, keep track of the size of the load in the current train being loaded. Proceed to the next hour. Go to Step 1.
- Step 6: If the tipple needs filling, we should fill it with two crews, add a half trainload to the tipple, add $1,000 demurrage cost for this train to sit on the tracks waiting for the tipple to load, and add $21,000 to the labor costs for the crews at work. Proceed to next hour. Go to Step 1.
- Step 7: Now that a high-capacity train with five engines has arrived, to prevent exorbitant demurrage costs for making this train wait in lieu of another, we will automatically switch to service the high-capacity train. If another train is currently being serviced, then we will put it on hold for now.
- Step 8: If three standard trains are now waiting, all three are going to have to wait for the high-capacity train and incur demurrage costs. Add $45,000 to costs for this hour's waiting.
- Step 9: If two standard trains are now waiting for the high-capacity train to be serviced, both will have to wait and incur demurrage costs. Add $30,000 to costs for this hour's waiting.
- Step 10: If one standard train is now waiting for the high-capacity train to be serviced, then the standard train will have to incur demurrage costs for this hour. Add $15,000 to costs for this hour of waiting.
- Step 11: By now, we must see if we can actually fill the high-capacity train that is on-line to be filled. If so, then no demurrage costs will be incurred for this one train for this hour. If the tipple needs filling, we should fill it with two crews, add a half trainload to the tipple, add to costs $25,000 for the high-capacity train to sit on the tracks waiting and $21,000 for the crews at work. Proceed to the next hour and go back to Step 1 now if the tipple was filled and the high-capacity train had to wait.

- Step 12: If we did not need to fill the tipple, then take one-third of a trainload away from the tipple and put it into the train. If the high-capacity train is full, send it on its way and let the load value of the current train being worked on go to 0 since it has now left. Otherwise, keep track of the size of the load in the current train being serviced. Proceed to next hour and go back to Step 1.

Figure 1 gives a graphic representation of the algorithm. Note that when we check to see if the tipple needs refilling, we check to see if it is less than one-sixth full. This is simply because of the nature of the fractions at which the tipple empties itself and is refilled hour by hour. By setting our refill value at one-sixth or less, we prevent our simulation from accidentally dropping the tipple load value below zero.

Results: Results from the simulation are given in **Table 1**. Two trials may seem insufficient, but realize that each trial models 10,000 weeks of data, for a total of more than 384 years.

Table 1.
Results of simulation. Each trial is for 10,000 weeks.

	First trial	Second trial
Demurrage costs/yr	39,752,000	39,605,000
(Labor + demurrage)/yr	89,904,000	89,730,000

Error Analysis:

- Sources of Possible Error
  - oversimplification of coal operation
  - assumption of constant tipple fill rates
  - assumption of constant train fill rates
- Computer roundoff error. Only addition was used to tabulate cumulative costs. Since multiplication tends to magnify round-off error, we avoided this source of error. In the final calculation of costs per week, however, we used the formula

$$\frac{\text{total costs of } N \text{ iterations of weeks}}{N \text{ weeks}}.$$

If N is on the order of 10,000 and the total cost on the order of magnitude of 10^8, we may lose a few digits of accuracy.

Weaknesses and Strengths

Weaknesses of the Model

- There is no comparison to a real-world data set. The final step of the modeling process should be to verify the model against real-world data. For our problem, real-world data are not available for comparison. This is a major weakness.
- The assumption of discrete hours and discrete events limits the flexibility of the model. For example, if we could calculate demurrage costs for trains arriving at 1:15 and waiting until 1:35, the simulation cost projections would be more realistic.
- The assumption that trains arrive with empty payloads fails to consider that a train may arrive with a fraction of a payload already loaded.
- The assumption that workers are on call at all times fails to account for real-life situations. Can we depend on having two crews available regardless of the time of day, weather conditions, or holidays? Realistically speaking, we cannot. Our model, however, assumes that we can.
- We fail to take into account "hidden" costs. What happens to the operation if the tipple breaks down? What effects does depreciation have on the value of the equipment used in the operation?

Strengths of the Model

- With an 80386 personal computer, a 10,000-week iteration model can be run in less than 10 minutes.
- Our simulation takes a rather complex scenario and simplifies the coal operation into something that is manageable and that we can use to make predictions about the behavior of the system. Simplicity can be a powerful tool for understanding a complex world.

The Submodels

Submodel 1: The Ideal Scenario

If the standard trains could be scheduled to arrive at precise times, what daily schedule would minimize loading costs?

Our simplifying assumptions are:

- non-overlapping service times,
- no high-capacity train arrives,
- one crew fills tipple, and
- the day of the week is not Thursday.

We decided to start with the simplest of cases. This train schedule attempts to fit all of the trains within the window of 5 A.M. to 8 P.M. without multiple trains arriving at the same time. In addition, the one train that is being serviced at the tipple must remain until it is full without having another train arrive. After the train is full, the tipple must be refilled to the level of one trainload before another train arrives. When the next train comes, it will also remain until it's full without another train arrival. Again the tipple will be filled to the level of one trainload before the last train arrives. After all three trains have come and gone, refill the tipple to 1.5 trainloads. We have formulated this scenario as the ideal case. It is ideal because no train ever has to wait and therefore, no demurrage costs are incurred. The labor costs for the tipple loading crew (one crew at $9,000/hr, not two crews at $21,000/hr) are also minimized. If the Aspen Company has control over when trains can arrive and it is feasible to devise the most ideal (minimized costs) scenario, then this simple model satisfies these needs. One ideal schedule that works is listed in **Table 2**.

Table 2.
An ideal schedule of train arrivals.

Time		Tipple level	Fill tipple?
0500	Train A arrives	1.5	—
0800		0.5	Fill 0.5
1000	Train B arrives	1.0	—
1300		0.0	Fill 1.0
1700	Train C arrives	1.0	—
2000		0.0	—
0200		1.5	—

In the ideal schedule, there is one 5-hr gap between trains (3 hrs to fill the train, 2 hrs to fill the tipple with half a trainload) and one 7-hr gap between trains (3 hrs to fill the train, 4 hrs to fill the tipple with one trainload). Therefore, there are two ways for this schedule to work:

- The first time gap between trains A and B is 5 hrs (a 7-hr gap must follow between trains B and C to fill the tipple to one trainload).
- The first gap between trains A and B is 7 hrs (a 5-hr gap must then follow between trains B and C to fill tipple to one trainload).

Table 3.
Two possible ideal train schedules.

	A	B	C
Case a	0500–0800	1000–1300	1700–2000
Case b	0500–0800	1200–1500	1700–2000

Observe these two possible cases in **Table 3**.

In this manner, we can follow the coal operation in an ideal day (no overlapping service times and no demurrage costs). See **Table 4**. For such a day, the daily loading costs are $108,000, with no demurrage costs.

Table 4.
Minimum loading costs on a non-Thursday. Total costs: $108,000.

Time	Tipple (trainloads)	Train A	Train B	Train C	Loading crew
500	1.50				
600	1.50	arrives			
700	1.16				
800	0.83				
900	0.50	leaves			
1000	0.75				$9,000
1100	1.00		arrives		$9,000
1200	0.67				
1300	0.33				
1400	0		leaves		
1500	0.25				$9,000
1600	0.50				$9,000
1700	0.75				$9,000
1800	1.00			arrives	$9,000
1900	0.67				
2000	0.33				
2100	0			leaves	
2200	0.25				$9,000
2300	0.50				$9,000
2400	0.75				$9,000
100	1.00				$9,000
200	1.25				$9,000
300	1.50				$9,000
400	1.50				

We must not forget to include an analysis of the ideal scenario on Thursday. On Thursday, however, there is no way to avoid demurrage costs. The daily loading costs for an ideal Thursday are $210,000, with daily demurrage cost of $135,000, for a total daily cost of $345,000. See **Table 5**.

Table 5.
Minimum loading costs on a Thursday. Total costs: $345,000

Time	Tipple (trainloads)	Train A	Train B	Train C	Big Train	Loading crew
500	1.50	arrives				
600	1.16					
700	0.83					
800	0.50	leaves				$21,000
900	1.00		arrives			
1000	0.67					
1100	0.33					
1200	0		leaves			$21,000
1300	0.50				arrives	
1400	0.16					
1430	0					
1500	0.25				$25,000	$21,000
1600	0.75				$25,000	$21,000
1700	1.25				$25,000	$21,000
1730	1.50					
1800	1.33					
1900	1.00					
2000	0.67			arrives		
2100	0.33			$15,000		
2200	0			$15,000	leaves	$21,000
2300	0.50			$15,000		$21,000
2400	1.00			$15,000		
100	0.67					
200	0.33					
300	0			leaves		$21,000
400	0.50					$21,000
500	1.00					$21,000
600	1.50					

Submodel 2: Minimizing Tipple Loading Costs

Would a third tipple-loading crew at $12,000/hr reduce annual operations costs?

We derive an algebraic formula to model the cost for multiple crews. For notation we will use

Cost	is the total cost of labor,
P	is the percentage of the tipple to be filled,
c	is the number of crews, and
a	is the number of trains in the system.

Our approach is to define cost as a function of the number of crews, the number of trains in the system, and the percentage of the tipple to be filled:

$$\text{Cost} = f(P, c, a).$$

If it takes any one crew 6 hrs to fill the tipple, and all crews work at this standard rate, we know that it takes $6/c$ hrs to fill the tipple. The cost of the tipple loading crew(s) is given by $12{,}000c - 3{,}000$. For 1, 2, and 3 crews, the costs are \$9,000, \$21,000, and \$33,000.

$$\begin{aligned} \text{Cost} &= P \times \text{hours} \times (\text{labor} + \text{demurrage}) \\ &= P \times \frac{6}{c} \times ([12{,}000c - 3{,}000] + 15{,}000a) \\ &= P\left(72{,}000 + \frac{90{,}000a - 18{,}000}{c}\right). \end{aligned}$$

For any integer $a > 0$, cost will be minimized by increasing c. Even though the second crew costs \$12,000/hr, or \$3,000 more than the first crew, the two get the job done in half the time. Demurrage savings of around \$15,000 per train per hour, minus an increase of wages of \$3,000, is well worth the extra crew.

Now let us concentrate on the question: "Would a *third* tipple-loading crew at \$12,000/hr reduce annual operations costs?" Yes! Substitute $c = 3$ into the above formula. Provided the third crew is called in whenever $a > 0$ (a is an integer), the third crew will always reduce that cost. Therefore the annual cost will be reduced also.

Submodel 3: Minimizing Demurrage Costs

How often should the second crew be called out?

Our simplifying assumption is that there are multiple train arrivals for any given hour.

Our reply to the question is that if there is at least one train waiting to be filled, it is cheaper to have two crews working to get the job done in half the time.

What would happen if the trains had to wait for half of a trainload to be filled into the tipple? For one crew, the demurrage cost would be 2 hrs times \$15,000, or \$30,000. For two crews, there would be a demurrage cost of 1 hr times \$15,000. Even for only half a trainload, the increase in demurrage cost is \$15,000 vs. an increase of \$12,000 labor cost. Using two crews would save \$3,000.

Submodel 4: A Fourth Train?

Can this tipple support a fourth standard train every day?

Adding in a fourth standard train on a non-Thursday is possible, and it will not overload the system.

On a Thursday, however, the system will begin to get backed up. Here's why: Looking at the best possible scenario, two trains arrive at 8 P.M., and the tipple is empty until it finishes with the big train at 10 P.M. Not until 5 A.M. the next morning does the last train even begin to get filled. Assuming that we can schedule when the trains arrive, by Saturday evening the system can be back in line again. Therefore, since the system does not overload, it can handle a fourth standard train. Nevertheless, having four standard trains on Thursday will greatly increase the demurrage cost. [EDITOR'S NOTE: For space reasons, we omit the authors' schedules for four standard trains on Thursday, Friday, Saturday, and Sunday.]

Submodel 5: Worst-Case Scenario

Given that our simulation model produces expected costs per week, how can we "double-check" the results to see if the simulation output is reasonable?

We put tremendous effort into developing a flowchart and computer code of the coal operation, so human error is great concern. Even with subroutine checks and debugging, we felt that we needed a method of gauging a "ballpark" figure of expected weekly costs. Therefore, we investigated maximized costs scenarios. From this vantage point, we get some idea of the simulation model verification. [EDITOR'S NOTE: For space reasons, we omit the authors' schedules that justify the following worst-case costs.]

The two scenarios are:

- Three standard trains arrive together:

Labor Costs	\$117,000
Demurrage Costs	\$195,000
Total Costs	\$312,000

- Three standard trains and a high-capacity train arrive together on Thursday:

Labor Costs	\$210,000
Demurrage Costs	\$655,000
Total Costs	\$865,000

Conclusions and Recommendations

Using our simulation and submodels, we were able to answer successfully the questions listed below. We recommend to the management that our simulation model be employed to answer "what if" questions about expected value. Furthermore, we recommend that a third crew be added full time to reduce costs of demurrage. Below are the bottom-line answers to the questions posed by management.

- *What is expected annual cost of the tipple's loading operation?*

Results from first 10,000 week run	\$89,904,000
Results from second 10,000 week run	\$89,730,000
Average	\$89,817,000

- *How often should the second crew be called out?*

  If there is at least one train waiting to be filled, it is cheaper to have two crews working, according to Submodel 3.

- *What are the expected monthly demurrage costs?*

Results from first 10,000 week run	\$3,058,000
Results from second 10,000 week run	\$3,047,000
Average	\$3,053,000

- *If the standard trains could be scheduled to arrive at precise times, what daily schedule would minimize loading costs?*

  See **Tables 4** and **5**.

- *Would a third tipple-loading crew at \$12,000/hr reduce annual operations costs?*

  Yes, according to Submodel 2.

- *Can this tipple support a fourth standard train every day?*

  Yes, according to Submodel 4.

References

Giordano, Frank R., and Maurice D. Weir. 1985. *A First Course in Mathematical Modeling*. Monterey, CA: Brooks/Cole. After we had decided on a simulation approach, we referred to the Harbor System Model (pp. 280–289) to see if we were on the right track; that model is similar to the coal-tipple problem.

Practitioner's Commentary: The Outstanding Coal-Tipple Operations Papers

Ruth Maurer
Dept. of Mathematical and Computer Sciences
Colorado School of Mines
Golden, CO 80401

Introduction

For the system of two crews, three regular trains, and one special train, all three of the final papers generated the same total cost (in the $87–90 million range, which is good for this stochastic situation). However, one team considered only a five-day week, while the others used a seven-day week.

Similar schedules for the three regular trains were achieved in the case where those trains could be scheduled. Costs ranged from $52 million to $59 million, again similar results given the stochastic nature of the problem.

All agreed that use a of third crew would reduce total annual costs but disagreed significantly on the total amount of savings.

All agreed that the system could handle a fourth regular train, but Thursdays would be problematic and costs may soar.

Detailed Analysis

The teams will be discussed in order of performance, first to last.

The team from Cornell University had the simplest solution and the one most amenable to sensitivity analysis. This team used an existing simulation package to build its model, and the parameters of the model can easily be changed to ask "What if —?" types of questions. This team's presentation of approach and results is probably the clearest, except for the the statistical analysis, which is not clear.

The team from the U.S. Military Academy wrote the clearest summary of the problem, approach, and recommendations. They wrote the most thorough statement of its algorithm, having done their own programming in Pascal; they also used spreadsheets to advantage in summarizing results. I disagree with their assumption of train arrivals on the hour—this is just unrealistic. They state as an assumption that "the distribution of arrival times is unknown" when they are actually assuming a discrete uniform distribution; they could just as easily have used a continuous uniform distribution.

The team from the University of Alaska Fairbanks is clearly more oriented in the direction of mathematical statistics, and less in the direction of simulation, than the other teams. This team's solution considers only five scenarios or "rule sets" and chooses the one that minimizes cost. Their solution to the scheduled-regular-trains part of the problem is reasonable, but they don't tell us how they arrived at that particular schedule. Their answer to the four-regular-trains question is sketchy at best. The team is to be commended, however, for testing their primary results against those of a Pascal simulation program (which they wrote). Since the simulation was based on the same logic as the theoretical solution, however, one would expect the results to agree.

About the Author

Dr. Ruth Maurer is presently Associate Professor of Mathematics (Operations Research/Applied Statistics) at the Colorado School of Mines. In addition to considerable professional work as a consultant, she is former Mayor of the city of Golden, Colorado. She also was the Consulting Energy Economist for the First Interstate Bank of Denver and was visiting professor of engineering at the U.S. Military Academy at West Point. For her pro bono consulting work for the Department of the Army, she was awarded the Outstanding Civilian Service Medal and the Commander's Medal. She is the co-author (with R.E.D. Woolsey) of the five books in the Useful Management Series.

Judge's Commentary: The Outstanding Coal-Tipple Operations Papers

Jonathan P. Caulkins
Heinz School of Public Policy and Management
Carnegie Mellon University
Pittsburgh, PA 15213–3890

This problem is deceivingly difficult because it has characteristics that are familiar (queueing, inventory, scheduling) but does not fit neatly into any standard class of problems. As a result, even though teams tried methods as varied as linear programming and simulated annealing, essentially every team also used Monte Carlo simulation.

Unfortunately, many teams plunged too quickly into simulation and neglected to perform supporting analyses. Indeed, several teams used Monte Carlo methods to estimate quantities, such as the expected value of the minimum of independent uniform random variables, that can easily be found exactly. These papers, even when the simulations were well constructed, typically generated more numbers than insight.

The better papers, including all three Outstanding papers, augmented simulations with other analysis. The U.S. Military Academy team produced upper and lower bounds on costs to verify that the results of its simulation were reasonable. When Thursday's operations spill over into Friday, demurrage costs on Friday may be affected. However, exact analysis of the costs of such disruptions is exceedingly difficult, and most teams simply ignored them. The University of Alaska Fairbanks team, in contrast, attempted to construct lower and upper bounds for those costs.

Better papers were distinguished also by a more mature treatment of the assumptions, including performing sensitivity analysis with respect to those assumptions. The Cornell University team stood out because they did not jump to the conclusion that two crews can fill the coal tipple twice as quickly as one. Using a simple graph, the team showed how demurrage costs vary with how much a second crew speeds up the loading process.

The three outstanding papers did not monopolize insightful analysis. For example, several teams gave careful discussions of whether or not to give priority to a high-capacity train. One even considered preemptive priorities and concluded that a standard train should preempt a high capacity train when the high-capacity train is less than about 20% full. There were also several excellent analyses of when to use one crew or two to refill the tipple, as a function of the amount of coal in the tipple, the time of day, and the number of trains still to come that day.

The questions of priority and number of crews were approached at different levels of sophistication. Some teams recommended a reasonable course of action with little explanation. Such recommendations are of little value, because they must be taken on faith and do not bring out general properties. Better papers gave either insightful intuitive arguments that produced more understanding or else mathematical proofs that were more persuasive. The best papers gave both.

More generally, good papers had a sense of perspective. Weaker papers worried excessively about minutiae, such as computer roundoff error and the fact that random number generators do not produce truly random numbers. Better papers recognized the general structure of the problem and designed their approach around it (Cornell); neatly presented and evaluated alternative decision rules (University of Alaska Fairbanks); and distilled their results into a concise, well-written summary (U.S. Military Academy).

When addressing a real-world modeling problem, there is no such thing as having "finished" the problem: There are always more ways of interpreting, structuring, and approaching the problem. In fact, that is the value of publishing the Outstanding papers—reading them helps provoke contestants to think in new ways about the problem that they worked on. Several ideas that were not fully tried in any paper but occurred to the judges as potentially fruitful include explicitly defining a state-space description of the system, and using dynamic programming or Markov decision processes, either as a solution procedure or as a way of structuring the problem.

The judges rewarded teams who used multiple methods, conducted sensitivity analyses, derived and justified insightful properties of good solutions, and had the perspective to understand and acknowledge the limitations of their models. It is very difficult to do well in all of these dimensions; good modeling is an art that takes considerable skill and practice for proficiency. The judges are delighted that so many students have accepted the challenge to become good modelers.

About the Author

Jonathan P. Caulkins participated in the MCM at Washington University, where his teams' entries were judged Outstanding in the first two competitions. Jon earned a master's degree in electrical engineering and computer science and a doctorate in operations research from MIT. Now he is an assistant professor of operations research and public policy at Carnegie Mellon University's Heinz School of Public Policy and Management, where his research focuses on mathematical models of illicit drug markets. Jon was an associate judge for the Coal-Tipple Operations Problem.