

Finding the Yellow Brick Road: Part 6, Courage!

WARREN B. POWELL

Department of Civil Engineering and Operations
Research
Princeton University
Princeton, New Jersey 08544

DONALD E. MAYORAS

Mau Trucking, Incorporated
90 Jacobs Addition
Ida Grove, Iowa 51445

Sometimes, there is more truth in fiction than in publicly available accounts. This story by Warren B. Powell, a leading researcher of real-time routing and scheduling models and frequent consultant to the motor carrier industry, and Donald E. Mayoras, a senior transportation executive and president of a trucking company, provides a view into the discussions many carriers undertake prior to adopting an optimization model. The account is too long to appear as a single article and instead is being published in serial form. This is the sixth installment.

The cast: Dan Manning, President; Tom Gorman, Chief financial officer; Matt Peterson, Vice-president of operations; Ken Richards, Vice-president of sales; Bill Johnson, Vice-president of management information systems; John Breswick, Director, driver management; Michelle Corwin, Director, customer service; Larry Michaels, Director, driver dispatch; Walter McCormick, Consultant; and Richard Merrick, Lead consultant, precision decision systems.

Allegheny Motor Carriers has been agonizing over the course of action to take to improve their declining profitability. The president, Dan Manning, has favored a more human approach, based on his belief in the principles of total quality management. Others favored the use of advanced satellite communication tech-

nologies to improve driver communication and create a high-tech image of the company in the marketplace. However, after carefully weighing the different options, the company has chosen to implement one of the new optimization models for assigning drivers to loads. This choice was favored strongly by Tom Gorman, the chief financial officer, in part because it was much cheaper than mobile communication, and also because he felt that a lot of the problems could be traced to the inability of the operations staff to handle the complexity of a growing company. Since optimization models tend to expose sloppy business practices, this decision clearly took the most courage on the part of the management team.

The company chosen to implement the optimization model, Precision Decision Systems, is a young startup recommended by Professor

Copyright © 1997, Institute for Operations Research
and the Management Sciences
0092-2102/97/2704/0012/5 00

INDUSTRIES—TRANSPORTATION
DECISION ANALYSIS—APPLICATIONS

INTERFACES 27: 4 July–August 1997 (pp. 12–18)

THE YELLOW BRICK ROAD

Walter McCormick. It features a new technology that seems very promising. The lead consultant on the project, Richard Merrick, has been brought in to describe the process and get Allegheny moving down the road to profitability.

As project leader, Bill Johnson called a meeting of the management team and the project leader from Precision Decision Systems, Richard Merrick.

"As everyone knows, we have contracted with PDS to install their optimization system. Richard Merrick will be the project manager. Rich has been doing this type of work for about five years and has developed what he says is the next generation in these optimization models. We're not the first carrier to use the system, but it's fairly new, and it may give us an opportunity to leapfrog our competition. I have had a chance to spend some time with Rich, and his approach is certainly different from the other vendors. I'd like to use this meeting to give Rich a chance to outline his strategy to all of you. I've also invited Walter back to listen in and keep Rich honest." Walter was sitting against the side wall in between Dan and John Breswick. "Rich, it's your meeting."

Rich stood up and walked to the end of the table where an overhead projector stood. He flipped the switch and put up an overhead transparency.

"First, I'd like to thank everyone for the opportunity to outline our strategy. We're pretty excited about our new systems, and we are looking forward to the opportunity for another successful implementation.

"The biggest challenge we face in implementing optimization models and algorithms for real-time operations is figuring

out what the problem is. There are other software packages out there that can do a decent job of solving a well-defined problem. The difficulty is . . . your problem is not well defined. As a rule, you are going to have data problems, and you are going to have trouble describing parts of your business."

Matt nodded vigorously. Bill put his pencil down and turned toward Rich. "What do you mean?"

"Well, a trucking company is a fairly complicated operation. Aside from matching drivers and loads, you have to manage drivers, the pool of trailers, and the customers. Your drivers are all different, and your customers are all different. Individual loads are different. Your planners know this, but they don't even realize a lot of what they have been doing on a routine basis. You can't possibly list all these issues before the project starts. I doubt that you can even list every piece of data that is important to decision making."

"What do you suggest?" returned Bill.

"Models have to get implemented in cycles. You start with an initial guess of what you think the problem is and try to solve it. Your planners will have an easier job of recognizing decisions they don't like than of articulating how they actually make decisions."

Rich paused while everyone pondered this and then continued.

"The juggling act is how smart to make the model. If you make it too smart at first, you won't be able to discover what is wrong."

Bill leaned forward. "Now I am really confused. The vendors make it seem like the implementation is really easy. You

make it seem like a lot of things will go wrong."

Richard was starting to warm up. "The biggest difference between human beings and computers is how they collect and process data. Humans have developed highly complex processes for collecting and processing data. In particular, they excel at handling bad data. If they have to pick up a load, and the date reads a month ago, they know to disregard the date. A computer can get really confused by this. The computer may show a driver as out of service, but the dispatcher just talked to her, and knows she can be dispatched from her home. And how many of your people conscientiously maintain driver log hours on the computer?"

"We don't need to," Matt responded weakly. "The dispatchers talk to the drivers, and they resolve any log hour problems over the phone."

"That's just my point," Richard responded. "Now, consider putting in a computer model that can look at 500 drivers and find a global optimum assignment that optimizes all the assignments simultaneously." He flipped to a new page and started drawing a series of drivers and loads (Figure 1). "Here's what a network model can do. Assume driver A is closest to load 1, driver B is closest to load 2, and so on." He drew a series of lines connecting drivers and loads. "Now you get down to driver F who will arrive later in the afternoon close to load 6 but can't cover the load because driver F won't arrive in time to pick up the load. A network model will be smart enough to use driver E on load 6, and then assign driver F to load 3, put driver C on load 1, put




Figure 1: A network model will assign drivers to loads (dotted lines) to make it possible for driver F to pick up a load, when the closest load, 6, must be picked up before driver F is available.

driver A on load 4, and leave driver B on load 2."

"That looks great!" interjected Tom. "What's wrong with that?"

Richard continued patiently. "Imagine what the dispatcher sees, trying to dispatch driver A. It makes sense to put him on load 1, which also gets him home. Now you are trading off getting driver A home versus picking up load 6 on time."

"But that's exactly the trade-off we want to make!" Tom noted, somewhat exasperated.

"True, but the real problem with this assignment is that the pickup window for load 6 is wrong. The customer service rep put in a morning pickup because the shipper's dock tends to be busier in the after-

THE YELLOW BRICK ROAD

noon. But there is nothing wrong with an afternoon pickup. The right assignment was to leave driver A on load 1 and let driver F cover load 6. The extra delay at the shipper is OK, because he doesn't have the hours to move after the pickup, and he's just going to go to bed after he gets the load."

"So we need to fix the pickup window?" pursued Tom.

"There is no way for the dispatcher to know that the decision to put driver 1 on load 3 was because of a bad pickup window for load 6." Richard paused. This was the critical issue with global network models, and he had seen problems with diagnostics at several companies trying to use global optimization.

After thinking a bit, Bill picked up the thread. "So the model is too smart?"

"That's a good way to think of it. Actually, the model is not that smart, because the data is wrong. It would be doing good optimization if all the data were correct, but the data never is. The problem is that it is hard to know when the data is the problem. If the tool doesn't allow you to identify data problems quickly and easily, you are simply optimizing the wrong problem."

Richard continued, "The biggest source of bad data is the loads themselves. Dispatch models optimize over the loads they know about but can't optimize the loads they don't know about."

"But most of the systems forecast shipper demands," interjected Bill.

Now it was Richard's turn to be exasperated. "Forecasting is a word that is used very loosely. For most people, *forecasting* means guessing at what loads a

shipper will tender to the carrier and then optimizing over these loads as if they really existed. This works well if you are really certain about the loads but poorly if the loads may not actually materialize. A way around this is to use a concept called mathematical expectation, which considers loads that might be called in, even if the probability is low. This technique is more sophisticated but requires so many approximations to get a tractable model that it is not clear what you get in the end. In addition, the most popular forecasting systems do not work at the level of an individual shipper. They aggregate freight into regions and then try to forecast the aggregate, including shippers with different equipment needs. The result is a lot of sophistication with very little additional value."

"So what do you suggest?" Dan responded. He was growing impatient with hearing about lots of issues that he didn't understand when he needed specific actions.

"Buying a quick fix is really tempting," Richard responded, "and I can understand why a number of companies have found it attractive. But the best approach is to move in stages. Start with the simplest system that brings value and use this system to pull your data and work processes forward. Initially, you will get poor solutions because you have bad data. As you chase these problems down, you will start to identify places where you don't like the answer because the system isn't smart enough. When this happens often enough, it is time to move to the next stages."

Dan jumped in, "I like this. It sounds like this approach will give my people

time to adjust to the new system incrementally, rather than in one big jump, which they probably can't make."

"That's right," added Matt. "I'm not sure how quickly I am going to be able to retrain my people to maintain DOT hours (Department of Transportation regulations on the number of hours a driver can work), customer profiles, trailer pools, and shipper forecasts."

Bill was curious, "What are the different stages?"

"That depends a little on you, but here's a rough outline" (Table 1). He flipped open a new chart and started writing. "Stage 1 is simply finding the best assignment for each driver, taking into account a lot of issues but ignoring other drivers. Let's call this local optimization.

"The second stage is limited global optimization, where the cost of assigning driver A to load 3 considers the cost of bumping driver 3 off load 3 and onto load 4, which in turn bumps driver 4 off load 4 and onto load 5, and so on. It is possible to limit how far this goes." He stopped and turned to the group and added, "... but not using network models or linear programming."

Turning back to the easel, he continued,

Stage 1	Local assignment, with multiple attributes
Stage 2	Limited to full global optimization, matching only
Stage 3	Full tour formation for short loads
Stage 4	Forecasting for routine pickups
Stage 5	Statistical demand forecasting and capacity management

Table 1: Building a system in stages permits the organization to pull its data and work processes forward as the system becomes more sophisticated.

"The end of stage 2 is full global optimization, but only assigning a driver to one load at a time. Stage 3 would consider optimizing full tours, where you can limit how far the tours go into the future."

"What do you mean by a tour?" Tom asked.

"A sequence of two or more loads. If you have short-haul freight or if you have short segments because of relays or terminals, you can actually assign a driver to a sequence of loads."

"Why wouldn't you do that one at a time? Just assign a driver to one load, block that assignment, and then let the model pick the next load."

"Because it ignores sequencing," Richard responded. "Remember how these models work. If you have five drivers and eight loads, a model picks which loads to move using such factors as revenue. So, you run the risk of picking long loads before short loads, because they have higher revenue. In fact, it might be better to do a short load that might take only an hour and then pick up the long load."

He continued, "Stage 4 starts to consider forecasting. My recommendation is to start with simple forecasting, where the planners enter loads from the major shippers that are pretty predictable. The problem with this type of manual forecasting is that it is time consuming. So, at stage 5, we use statistical forecasting of shipper demands combined with a capacity management system."

"Do the stages have to go in that order?" asked Matt. "It seems as if forecasting is important. Do you really have to wait until the end?"

"The answer to your question is no, but

THE YELLOW BRICK ROAD

consider how long each stage takes to pull off. Depending on your data quality, stage 1 may not take more than two or three weeks. It can take much longer if you discover that you are missing key pieces of data. You may have to change your screens and databases and then change your work processes to get the data into the computer. I have seen companies that have never gotten past this stage. Of course, if you can't get this stage right, then there is simply no point in continuing. Your ability to get through stage 1 is really an indication of the maturity of the company."

Tom smiled, but the others squirmed perceptibly in their chairs. Bill in particular understood what was involved in getting quality data into the computer. Matt thought about the impact it would have on his planners and dispatchers.

Rich continued, "Activating stage 2 involves nothing more than turning a knob on the package, but a small percentage of the recommendations the computer makes will be hard to understand. Of course, most recommendations will still be pretty obvious, even with the optimization. But you don't need optimization to handle those cases. It is the cases in which the optimization is doing its job that the recommendations are difficult to understand. Also, you don't get rid of data problems just because you got past stage 1."

"How long will this stage last?" Matt pursued.

"Again, it depends on how well you cleaned up your data. You can start turning the knob on stage 2, allowing it to look one leg away, then two legs, then three, and so on. Once you get past three or

four, you are finished. This whole process might take only a week, but it could take months. Again, it depends on you."

"How about stage 3?" asked Bill.

"The importance of stage 3 depends on the nature of your freight. If you have a lot of short-haul loads, it's really important. If you are mainly a long-haul carrier, your longest tour is going to be two loads out. That is, the model will assign a driver to two successive loads.

"Stage 3 raises two basic issues, one with data and one with operations. The data problem is that you have to predict when a driver is going to finish one load so he or she can start another. If you have short loads and you're trying to assign a driver through a sequence of three or four loads, then the main difficulty is predicting load and unload times. Some shippers get pretty backed up. If they don't control their dock well, then you can get delays of several hours that are hard to predict. Also, monitoring hours of service for a driver becomes more critical in this type of setting. On the other hand, if your loads are longer and your tours typically consist of only one or two loads, then these data items won't be so critical."

"And the second issue?" prodded Matt.

"That has to do with committing drivers to more than one load at a time. It is one thing for the model to plan for drivers to handle sequences of three loads, but that doesn't mean you have to tell the drivers. Since the problem is always changing, you might want to just give drivers their first dispatches and let them call in after they drop off each load."

"That's extra work for the drivers . . .," offered John Breswick.

"Not if you use satellites." interjected Matt, still disappointed that the company was not moving more quickly with satellite communications.

John shot a look at Matt and continued, "... also, it's nice to tell the drivers what's in front of them, so they can plan their work hours better."

Rich joined the discussion, "Satellites are a major technology for improving data collection. You can use optimization methods effectively without satellites, but a key piece of information is driver ETAs (estimated times of arrival). If you don't have good ETA information and don't want to improve it with satellites, you will have to institute other changes to get this data. The model has to know about incoming drivers, and it has to know when they will be available."

"So what about stage 4?" Dan asked, trying to close the discussion.

"Stage 4 is the first forecasting stage. You move to it when you start noticing problems with recommendations the model makes because of the freight it doesn't know about. Your first step is to put in the loads that you know about. These are not real loads, but they are tendered to the model as if they are real loads, except that you would still identify them as forecasted loads. The main challenge here is changes to your screens and database. If you plan for it now, you can flip it on when it is time, so there is no development delay. Right now, it is not on the critical path. If you get the screens developed, then the main delay in activating the feature is getting your planners and customer service people to use the screens."

Bill suddenly looked up with a question. "Do we have to put in a new model for each of these stages?" The potential for escalating software costs was starting to worry him. He was happy to run the project, but the software would be coming out of his budget.

"Up through stage 4, there is only one model," Rich reassured him. "For stages 1 through 4, you just change an input parameter. For stage 4, you need to make a change on your end—it doesn't affect the model. For stage 5, however, you need a new module."

"Does that mean we have to replace the first model?"

"No, the forecasting module operates independently and communicates with the dispatch system. Moving to stage 5 takes your company into an entirely new set of issues and capabilities. You will be able to anticipate shipper needs and forecast capacity. The outputs help the dispatch system to prioritize loads, but they also drive load acceptance, spot pricing, telemarketing, and load solicitation. This really puts the project into an entirely new phase." Richard stopped and looked around.

Sensing closure, Dan added, "Any more questions?"

Sitting away from the table, Walter had slipped John a slip of paper a few minutes earlier, and now John looked up with a question. "What I want to know is, how does the dispatch model handle dual degeneracy?"

Everyone turned and looked at him in surprise. Walter tried to keep a straight face, but John couldn't hold it in. The entire room broke into laughter. (To be continued.)