Transportation Operations Model (TOM) Technical Manual

Fuel Cycle Research & Development

Prepared for
U.S. Department of Energy
Campaign or Program
Ingrid Busch and Rob Howard
Oak Ridge National Laboratory
October 2012
FCRD-NFST-2012-000425

SUMMARY

The Transportation Operations Model (TOM) is a component of the Transportation Storage Logistics (TSL) model, which assists the analyst in planning and evaluating scenarios for the transportation and storage of spent nuclear fuel (SNF). TOM was originally developed for the DOE's Office of Civilian Radioactive Waste Management (OCRWM) in 2006. TOM provides detailed modeling of the transportation of spent nuclear fuel and the assets needed to accomplish that.

This Transportation Operations Model (TOM) Technical Manual presents the theory and algorithm behind the model. Further information on the Transportation Storage Logistics model is available in the Transportation Storage Logistics Users' Manual, and in the Transportation Storage Logistics Model Data Management Manual.

CONTENTS

FIGURES

TABLES

ACRONYMS

CALVIN	CRWMS Analysis and Logistics Visually Interactive
CRA	Cask Receiving Area
OCRWM	Office of Radioactive Waste Management
SNF	Spent Nuclear Fuel
TOM	Transportation Operations Model
TSL	Transportation Storage Logistics

TRANSPORTATION OPERATIONS MODEL (TOM) TECHNICAL MANUAL

1. Introduction

The Transportation Operations Model (TOM) is a component of the Transportation Storage Logistics (TSL) model, which assists the analyst in planning and evaluating scenarios for the transportation and storage of spent nuclear fuel (SNF). TOM was originally developed for the DOE's Office of Civilian Radioactive Waste Management (OCRWM) in 2006. TOM provides detailed modeling of the transportation of spent nuclear fuel and the assets needed to accomplish that.

OCRWM had also developed another model, called the CRWMS Analysis and Logistics Visually Interactive (CALVIN) model, which provides detailed modeling of operations at the reactor sites and storage facility. In 2012, the decision was made to combine these tools to provide an integrated approach to SNF storage and disposal. The TSL model is the result of this merger.

2. Architecture

The TOM model is written in C# and interacts with a Microsoft SQL Server database. It was designed as a console application that retrieves all of its input data from the database, and stores its results in the database, so that no user interaction is required during the run. A separate TSL user interface has been developed to allow the user to set up runs and analyze the results.

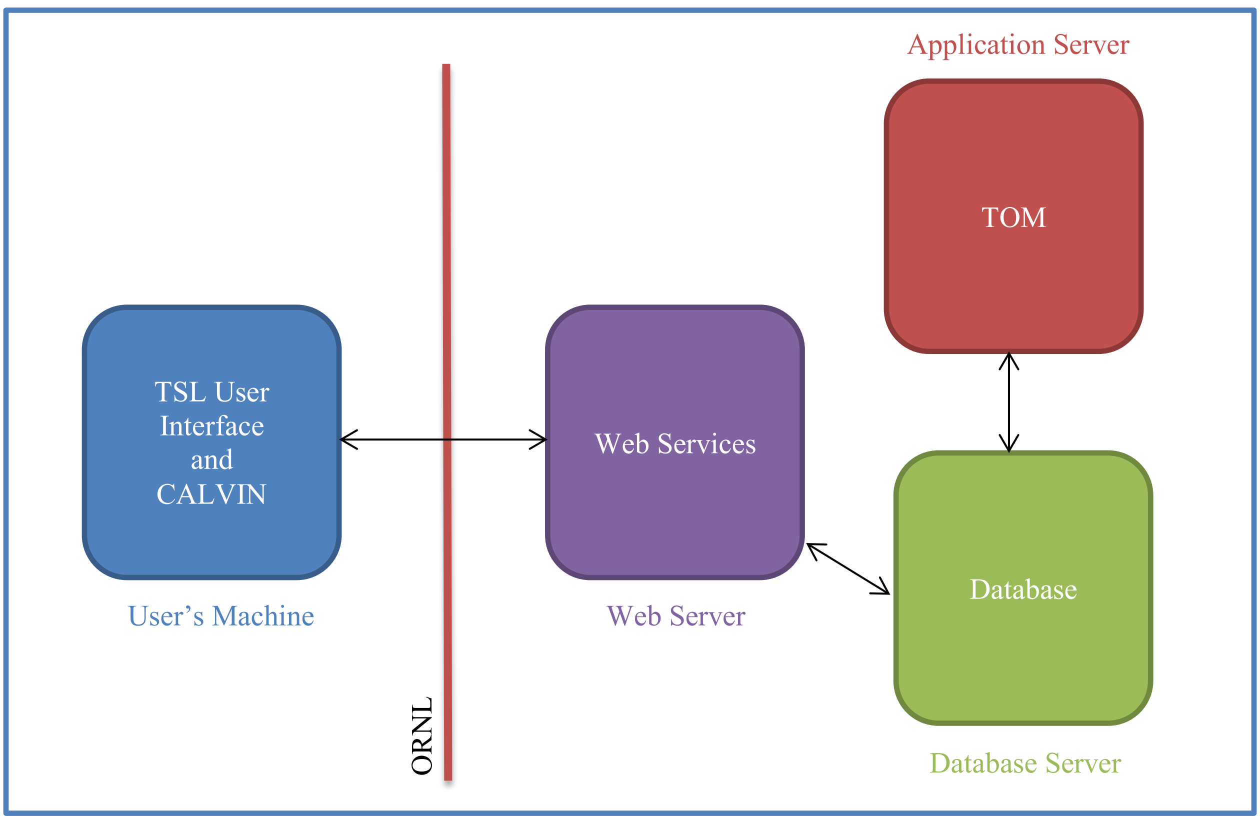

Figure 1. TSL Architecture

The TOM model and the database with which it interacts need to be on the same network, since the connection from TOM to the database is direct.

The sole parameter to the TOM executable is the scenario ID to be run. TOM is controlled by a configuration file, which has the following keys.

Key	Value
server	SQL Server TSL server name
database	SQL Server TSL database name
tempDirectory	directory to be used for temporary files during run
routingServer	SQL Server routing server name
routingDatabase	SQL Server routing database name

Table 1. TOM Configuration Values

3. Inputs

3.1 Scenario

The concept of a scenario is central to the way that TOM operates. The scenario-specific input and output data are all tagged with the scenario's ID. The scenario level data includes the scheduling horizon (the time span for fuel transportation), the sites that are to be used for storage and maintenance, and the default maximum consist size (which may be overridden at individual sites).

Also specified is the year to which costs are to be inflated (typically the current year).

3.2 Casks and Canisters

TOM separates fuel containers into "casks" and "canisters". Canisters require that a transportation overpack cask be brought in to the site so that they can be loaded. Canisters themselves are not tracked nor procured in the TOM model.

TOM is more concerned with casks than with canisters. For the most part, the movement of fuel will require that casks be brought in to the site to be loaded, and the procurement of casks is a major cost driver for the transportation process. Characteristics of casks which are important to TOM include their mode of transportation (rail or truck), their empty and loaded weight, the length of time it takes to load the cask (either with the canister, in the case of an overpack, or with bare fuel assemblies), and how long it takes to accomplish an inspection. Casks are inspected at the Cask Maintenance Facility (CMF) after every use, in addition to a more extensive annual inspection. Costs associated with a cask include its procurement cost (which is charged during the first year of its use), maintenance costs (both routine and annual), and the disposal cost. Casks have a lifetime limit, at the end of which they are disposed.

3.3 Pickup Schedule

A pickup in TOM is a requirement to transport a number of casks of a specified type from one site to another in a given year. There are two versions of a pickup. The first is where a cask has been specified itself. This cask may be a transportation cask for bare fuel, or it may be an overpack cask for a canister.

The pickup data also indicates whether the fuel is already in the cask (for example, if it is a dual-purpose cask that can be used for both storage and transportation). If the fuel is not already in a cask the model will schedule a consist to come with empty casks to be filled on arrival. If the fuel is already in a cask, the turnaround time at the onload site is less than if the cask needs to be loaded after arriving on site.

The second version of a pickup is when the pickup is specified as a canister. In this case, the model assumes that the fuel has already been loaded into the canister, and the consist needs to arrive with transportation overpacks into which the canisters will be loaded. Each canister is associated with a transportation overpack (cask) and the model translates the number of canisters into the number of casks that need to be transported. TOM does not track canisters, and does not bring in empty canisters to be loaded.

A pickup may be transported by several consists that occur months apart. In general, the pickup schedules are generated by CALVIN and input into the database using the TSL user interface.

Scenarios and pickup schedules have a many-to-one relationship. That is, many scenarios can link to the same pickup schedule, but each scenario can be linked to just one scenario. This allows the analyst to investigate the effects of varying size consists and routing strategies on the overall cost, schedule, and acquisitions required.

3.4 Sites

There are several types of sites in TOM. Pickup sites are places where the SNF to be transported is located. For the most part, these are reactor locations. Storage sites include interim storage facilities (ISF) and a repository (MGR). Maintenance facilities are locations to accomplish cask, railcar, and truck trailer maintenance. Any of the sites can be designated to accomplish these tasks, and these roles may be co-located at a site. For example, the maintenance facilities may all be co-located at the repository site.

The model allows for multiple ISFs in a scenario. SNF destined for ISF storage will be move to the nearest ISF.

There is a default maximum consist size for the scenario, but individual sites can have their own maximum consist size, to account for limitations of space for storing casks during the loading procedure.

A site can be characterized by where its rail access is. Some sites have no possibility of rail movement, and these are designated as “truck only” sites. At the other extreme are sites that have rail access to the cask receiving area (CRA) where the railcar can be moved “under the hook” for loading a cask. For the rest of the sites, heavy haul transport (and other equipment) will be needed to accomplish the loading of the SNF onto the railcars.

The type of rail access that a site has dictates the number of locations that need to be specified to modeling the transport. For sites that have rail access that reaches the plant, a rail node location and road node location are specified. For sites that have no rail access, a road location is specified.

For a site that has a nearby (but offsite) rail access, a road network location is specified as well as the intermodal location where the casks will be transferred between railcars and heavy haul transports.

The last group of sites utilize the waterway to move the casks between the rail network and the site. For these sites the locations of the plant’s road network access, water dock near the plant, rail access point, and water dock near the rail site need to be specified.

The locations associated with sites are specified by their latitude and longitude. This reduces the dependence of the model on any one route generation model, and increases its flexibility.

3.5 Transportation Assets

Transportation assets are the vehicles used to move the SNF. The assumption is that for rail transport, DOE will procure cask railcars, buffer railcars, and escort railcars. The cask cars carry the casks, the buffer cars (two per consist) will provide safety in the event of a crash, and the escort car will house the personnel and communications equipment for security. DOE will contract with the railroads to provide the engines to power the trains.

For truck transport, DOE will procure trailers, but will contract with shippers to move the trailers. Two escort trucks will accompany each truck convoy, providing security.

Transportation assets are inspected at their respective fleet maintenance facility at the end of each transportation cycle. In addition, each asset type also has a mileage inspection limit. When this limit is exceeded, it needs to undergo a more thorough (and time-consuming) inspection.

Costs associated with the transportation assets include acquisition costs and costs for routine and intensive inspections. There is no lifetime limit for transportation assets. Once they are procured, they are available from then until the end of the scheduling horizon.

Other ancillary equipment, such as heavy haul transport and cranes will be leased, and their services will be procured on a site-by-site and trip-by-trip basis. Their costs are approximated in the model, but the actual cost is expected to differ by site and circumstance.

4. Routing

4.1 Transportation Cycle

Each trip to transport SNF begins and ends at the fleet maintenance facility. The consist (or convoy, in the case of truck transportation)

1. travels to the pickup site,
2. loads the fuel into casks and onto the transportation asset,
3. travels to the storage facility (ISF or MGR as dictated by the schedule),
4. unloads the cask, unload the fuel, and loads the empty cask onto the transportation asset,
5. travels to the cask maintenance facility,
6. performs cask maintenance,
7. travels to the fleet maintenance facility, and
8. Performs fleet maintenance.

In the case of co-located facilities (such as cask and fleet maintenance facilities), the travel between these operations is eliminated, though the tasks performed at the facilities are done sequentially as indicated above.

4.2 Route Determination

The transportation mode of a specified cask determines whether the fuel will be routed by road or by rail. For road transport, the route consists of one leg, from the pickup site's road network location to the storage site's road network location.

For rail transport, the rail access designation (i.e., how the site will access the rail network) determines the legs that will need to be generated. For a site that has rail access that continues to the owner-controlled area, one leg is generated from the pickup site's rail network location to the storage site's rail network location.

For a pickup site whose rail access location is offsite (but does not need to use the waterway to reach the rail network), a road network leg is generated from the pickup site's road network location to the rail terminal, and then a rail network leg is generated from the rail terminal to the storage site's rail network location.

And finally, for a pickup site whose travel to the rail network involves use of the waterway, a road network leg from the site's road network location to the site's designated dock facility is generated. The second leg is a waterway leg from the site's dock facility to the rail/water terminal, and the final leg is on the railroad network from the rail/water terminal to the storage site's rail network location.

The determination of the network routes (the rail, road, and waterway links) used to transport the fuel is accomplished exogenously to the TOM model. TOM needs the distance of each leg of the trip (for use in trip duration and costing calculations), and that is what is stored in the TSL database. When TOM needs a distance, it first checks the database for the information. If the distance is not stored there, it calls an external system to generate the distance. The travel time is then calculated using the average speed of the relevant asset.

5. Cost Calculation

Costs in TOM are separated into three categories: capital, maintenance, and operations. Each cost parameter is associated with a year (i.e., the year for which the sample cost was given), and operations costs are inflated to the year designated in the scenario.

5.1 Capital

Capital costs include the acquisition and disposal of casks and rolling stock. The costs per cask or rolling stock asset are stored in the database, and are entered in current year dollars. These capital costs are not inflated. The cost of acquisition occurs in the year that the acquisition takes place; there is no advance payment or lead time requirement. There is an implicit assumption that casks can be acquired when they are needed. Whether the cask manufacturing process can accommodate this acquisition schedule is outside the scope of the TOM model.

5.2 Maintenance

Maintenance costs are incurred whenever inspection is done on the casks and rolling stock. These happen at the end of each transportation cycle. For casks, if more than a year has passed since the last annual inspection, it is performed at that time, otherwise, the cask undergoes a standard inspection. For the rolling stock, if the maintenance mileage limit has been exceeded, the more extensive (and expensive and time-consuming) mileage inspection is done, otherwise the rolling stock undergoes the standard inspection. By performing the maintenance when the limits are exceeded, the model may underestimate the amount of maintenance that is done, however, there is sufficient slack built in to the schedule to allow the maintenance to be done to ensure compliance with the required maintenance schedule.

Maintenance does not include maintaining the transportation network infrastructure.

5.3 Operations

Operations costs are those that are directly involved in each transport of SNF. They include transportation asset operations, leased cask handling equipment (such as cranes and heavy haul trucks), 180c charges, and security personnel costs. These costs are the most complicated of the cost calculations. The operations costs can be broken down into categories.

5.3.1 Border Crossing

There are inspections that need to be done whenever a consist or convoy enters a new state, and the costs for these are estimated using a border crossing charge. The cost is given as

$$c_{border\ crossing} = n * s * c_{bc}$$

where n = number of cars in consist (or trucks in convoy)

s = number of state borders crossed

c_{bc} = cost per state border crossing

5.3.2 Crane

Cranes are required to move the rail casks from one asset to another, and to or from a storage location. There is a deployment cost involved in transporting a crane, as well as setting it up once it arrives on location. There is a daily cost for having the crane onsite, as well as an hourly cost for crane operation (which includes operator costs). The crane cost is given by

$$c_{crane} = c_{crt} + c_{cro} * t + c_{crd} * d$$

where c_{crt} = deployment cost for the crane

c_{cro} = hourly operations cost for the crane

c_{crd} = daily onsite cost for the crane

t = length of time crane is in operation (in hours)

d = length of time crane is onsite (in days)

5.3.3 Heavy Haul

Heavy haul transport is required whenever the rail line does not reach the area where the casks are being loaded with the fuel. There is a deployment cost involved in getting the transport to the site, an hourly cost for operation, and a daily cost for having the heavy haul transport on site. The heavy haul cost is given as

$$c_{heavy\ haul} = c_{hht} + c_{hhh} * t + c_{hhd} * d$$

where c_{hht} = deployment cost for the heavy haul transport

c_{hhh} = hourly operations cost for the heavy haul equipment

c_{hhd} = daily onsite cost for the heavy haul equipment

t = length of time heavy haul is in use (in hours)

d = length of time heavy haul is onsite (in days)

5.3.4 Mainline Rail

The mainline rail is the cost for a dedicated train shipment by a mainline carrier. Rail charges for a dedicated train are set by contract between the shipper and the carrier, so the calculated costs in TOM are an approximation of what the actual charges would be. The costs are a function of the weight of the casks, the number of cask cars, and the distance travelled. The mainline cost is given by

$$c_{mainline} = c_{rms} * d^2 + c_{rm} * d + c_r + n * w * (c_{rts} * d^2 + c_{rt} * d + c_{rw})$$

where c_{rms} = cost per mile squared

c_{rm} = cost per mile

c_r = cost per trip

c_{rts} = cost per ton mile squared

c_{rt} = cost per ton mile

c_{rw} = cost per ton

d = distance travelled

n = number of cask cars

w = weight of cask

Note that the cost will differ when hauling a loaded cask than when hauling an empty cask.

5.3.5 Rail Security

Rail security costs are incurred whenever a loaded cask is outside of a site's protected area until it arrives at the storage or disposal location. Empty casks are not guarded. TOM assumes that six guards are required for a security detail, and that each guard is paid 24 hours per day. The rail security charges are given as

$$c_{rail\ security} = c_{rsh} * 6 * h$$

where c_{rsh} = hourly cost per security guard

h = hours security is required

5.3.6 Shortline Rail

Shortline rail companies are local or regional companies that provide transportation from sites to the mainline rail terminals. The shortline cost is a function of the number of cask cars that are being moved. The shortline rail cost is given by

$$c_{shortline} = c_{slt} + n * c_{slc} + 3 * c_{slo}$$

where c_{slt} = per trip cost for shortline rail

c_{slc} = per cask car cost for shortline rail

c_{slo} = per non-cask car cost for shortline rail

n = number of cask cars

The shortline rail cost is only accounted for when the rail path has an indicator for a shortline component.

5.3.7 Switching

The switching fee is the charge to move a railcar from one track to another. In TOM it is assessed when the cars need to be switched from a shortline carrier to a mainline carrier. The cost is a function of the number of cars in the consist. The switching cost is given by

$$c_{switching} = c_{swc} * n + 3 * c_{swo}$$

where c_{swc} = switching fee for a cask car

c_{swo} = switching fee for a non-cask car

n = number of cask cars

5.3.8 Truck

The truck cost is a function of the distance travelled and the weight of the load. It is assumed that there are two drivers, so that the truck can move nonstop. The truck cost is given as

$$c_{truck} = c_{trt} * w + c_{trd} * w * d + \max(c_{trs} * d, m_{trs})$$

where c_{trt} = per ton cost of truck transport

c_{trd} = per ton-mile cost of truck transport

c_{trs} = cost per mile of a second driver

m_{trs} = minimum cost for having a second driver

w = weight of the cask (in tons)

d = distance travelled (in miles)

Note that the cost will differ when hauling a loaded cask than when hauling an empty cask.

5.3.9 Truck Security

Truck security costs are incurred whenever a loaded cask is outside of a site's protected area until it arrives at the storage or disposal location. Empty casks are not guarded. TOM assumes that six guards are required for a security detail, and that each guard is paid for each mile of travel with the cask. The truck security charges are given as

$$c_{truck\ security} = 6 * d * c_{tsm}$$

where c_{tsm} = cost per mile per security guard

d = distance

5.3.10 Waterway

The waterway cost includes the cost of barge and towboat use.

The towboat cost is given by

$$c_{towboat} = c_{wtt} + c_{wtd} * t$$

where c_{wtt} = deployment cost for the towboat

c_{wtd} = daily cost for the towboat

t = waterway travel time (in days)

The barge cost is given by

$$c_{barge} = n * (c_{wbt} + c_{wbd} * t)$$

where n = number of barges required

c_{wbt} = deployment cost for the barge

c_{wbd} = daily cost for the barge

t = waterway travel time (in days)

There is an implicit assumption in TOM that each barge can accommodate three rail casks, so for a consist carrying no more than three casks, only one barge will be required.

6. Assumptions and Caveats

While the present state of the transportation network is known, both rail and road (and to a much lesser extent, waterway) networks are changing entities, and an assumption implicit in the analysis is that the networks as they are now will be usable at such time as SNF is to be moved. Many sites have rails that go up to or into the utilities' control area, but in some cases, these rail lines have not used since construction of the reactors. The lines are still in the network, but it may be that refurbishment will be required if those rail lines are to be utilized in the transportation of the SNF. Of particular concern are the rail crossings and the bridges. As the railroad infrastructure ages, it will no longer be capable of handling the weight of the cask cars. That said, the rail routes and truck routes used by the TOM model were developed using the transportation networks as they currently exist. A caveat to this is that there is an assumed rail line that reaches Yucca Mountain.

Another consideration is that while the cost of trailer transportation over the road network is, for the most part, proportional to the distance travelled, the cost of routing over the railroad network is not as straightforward. Each rail line is owned by a company, and that company will limit which other railroad companies can use the line. The routing of cargo by dedicated rail means that the shipment will be "handed off" from one carrier to another along its path, and DOE will contract with each carrier. Because the details of such contracting are impossible to model, the TOM software includes a proxy rail routing cost function based on distance and weight.

It is assumed that there can only be one consist loading at the reactor at a time. There are economies of scale that can be realized by accomplishing all of the loading needed in a calendar year sequentially. For example, when cranes have to be brought in to the reactor site, there is generally a deployment cost which covers the cost involved in bringing the crane to the site and setting it up. There is also a daily cost for having the crane on-site, as well as an hourly cost for the operation of the crane, which will cover the cost of the operating crew. However, these rates would again be contracted between DOE and the company providing the equipment and crews. Modeling the details of such contracting is complex, and the TOM model assumes that the crane deployment costs are charged for each rail consist.

It is assumed that the unloading capability at the interim storage facilities and repository are unlimited. While this will not actually be the case, it is assumed that the unloading capabilities will be constructed so as to not be the bottleneck in the transportation process. Once the details of the interim storage facilities and repository are established, a constraint should be added to TOM to ensure that the unloading infrastructure is not overtaxed.

Train dynamics and weight restrictions on bridges limit the number of railcars that can be included in a consist. Security concerns limit the number of trucks that can be observed by an escort truck. Also, once casks are loaded on to transportation assets, they must be secured, and not all sites have the facilities for secure storage of railcars or trailers once they are loaded. Limiting the consist size will mitigate these storage issues.

Train consists are assumed to travel with two buffer cars and an escort car containing security. Truck convoys are assumed to travel with two escort trucks.

The operations cost is a function of the distance that the spent fuel is transported. This distance is highly dependent on where the storage facilities and repository are located. Operations costs are expected to have the most variability as the storage locations change.

The capital cost is driven by the number of casks and transportation assets that need to be purchased during the campaign. This is in turn driven by the cycle time needed to transport the spent fuel. While the location of the storage facilities and repository drive the transportation time, transportation is actually only a small contributor to the overall cycle time. Loading, unloading, and maintenance activities make up the majority of the cycle time. Therefore, the acquisitions required are for the most part relatively independent of the location of the storage facility and repository.

Maintenance costs are directly related to the number of casks and transportation assets that are being used, so these costs will vary as the capital cost value does.

A computer generated schedule is an idealized model of how shipments would be accomplished in a perfect world. Machinery does not fail, transportation assets are not delayed, and all needed equipment appears on schedule. In actuality, such a schedule is fragile. Introducing a "contingency factor" to purposefully introduce delays in the schedule would produce buffers that will help guard against such a problematic schedule, but would also tend to overestimate the resources required. It must be kept in mind that TOM is a planning model, and not an operations model.

7. Formulation

The TOM model is based on the premise that it is best to move SNF in the largest size consist allowed. Though this may not always be the overall least costly choice based on equipment required, the coordination and preparations needed before and during each trip dictate that each trip maximize the amount of SNF that can be transported. Therefore, the decision of how large to make each consist, and hence how many trips will be needed is a straightforward calculation. TOM will build as many of the largest-sized consists permitted at the pickup site as possible, and then add another less-than-max-sized consist to move the remaining casks.

The decisions that remain to be made by the model involve the scheduling of the trips to transport the SNF. A mixed integer program can be formulated to state the problem:

Minimize $\Sigma_c k_c z_c + k_w w + k_v v + k_u u + k_r r$

$\Sigma_d x_{id} = 1$ for all trips i

$x_{id} <= y_{id'}$ for all $d <= d' < d + t_i$, for all trips i, days d

$\Sigma_{i \text{ in } Ic}\ s_i y_{id} <= z_{0c} + z_c$ for all days d, cask types c

$\Sigma_{i \text{ in } R}\ s_i y_{id} <= w_0 + w$ for all days d

$\Sigma_{i \text{ in } T}\ s_i y_{id} <= v_0 + v$ for all days d

$\Sigma_{i \text{ in } R}\ y_{id} <= r_0 + r$ for all days d

$\Sigma_{i \text{ in } T}\ y_{id} <= u_0 + u$ for all days d

x_{id} in $\{0,1\}$

Where s_i = number of casks in trip i

t_i = round trip time span of trip i

I_c = trips that move cask type c

$R = \{i \mid \text{trip } i \text{ is by rail}\}$

$T = \{i \mid \text{trip } i \text{ is by truck}\}$

z_{0c} = number of existing type c casks available

w_0 = number of rail cars available

v_0 = number of trucks available

r_0 = number of rail escort groups available

u_0 = number of escort trucks available

The interpretation given to the variables is:

$x_{id} = 1$ if trip i starts on day d

$y_{id} = 1$ if trip i is active on day d

z_c = number of casks of type c to be acquired

w = number of rail cars to be acquired

v = number of trailers to be acquired

r = number of rail escort groups (consisting of two buffer cars and an escort car) to be acquired

u = number of escort trucks to be acquired

While it is straightforward to formulate the mixed-integer programming formulation for the problem that TOM is modeling, solving this problem to optimality is time-prohibitive for all but the smallest of problems.

8. Strip Packing Problem

The problem of scheduling rolling stock and casks is a form of a strip packing problem. In the classical strip packing problem, there is a container of given width and unlimited height, as well as a collection of items of given width and height, and the problem is to pack the individual items into the container to minimize the height of the container required. The items may not overlap and cannot be rotated. Devising an optimization technique for this problem is difficult. (In computer science terms, the problem is NP-complete.) Solution techniques for the strip packing problem have focused on devising heuristics that provide reasonable solutions to the problem.

Once the size of a consist and its routing are determine for a trip, the trip becomes a two dimensional block whose height is the size of the consist and whose width is the length of time required to complete the transportation cycle. Scheduling the trips for a given year is a matter of fitting these trips into a container whose length is one year. The height of the container represents the number of assets needed, so minimizing the height becomes an asset-minimization problem.

There is an existing height to each container, and packing below this height does not lead to an improved solution. This height corresponds to the available assets at the beginning of the year. If the rectangles can be packed into the existing size containers, no new assets will need to be acquired in that year.

If there were only one cask type and one type of rolling stock, this would be the classical problem. However, in general, there are several cask types and several types of rolling stock, including buffer cars and escort vehicles, and this complicates the problem. There are several containers that need to be filled concurrently. There is a container for each type of cask to be scheduled, another for railcars, another for trailers, and two more for both types of escort vehicles. (The railcar escort vehicle container also mirrors the buffer car container, since an escort car is assumed to come with the buffer cars needed to completed the consist; they do not need their own container in the model.)

As an example, a rail trip will result in a rectangle (with height the size of the consist) in the railcar container, the same size rectangle in the appropriate cask container, and rectangle with height of one in the escort railcar container. These rectangles all have the same width, and the starting time of the trip (represented by the distance from the left edge of the container) must be the same for each of these rectangles.

There is an additional complication to the scheduling problem in that the strips representing the assets do not need to be adjacent to be scheduled on a trip together. For example, in the figure below, casks 1, 2, and 3 completed a 3-car trip (in blue) as did casks 4, 5, and 6 (in red). Casks 1 and 2 were scheduled on another trip (in orange) as were casks 4 and 5 (in green). Casks 3 and 6 could be scheduled on the trip indicated in purple, even though their strips are not adjacent in the container.

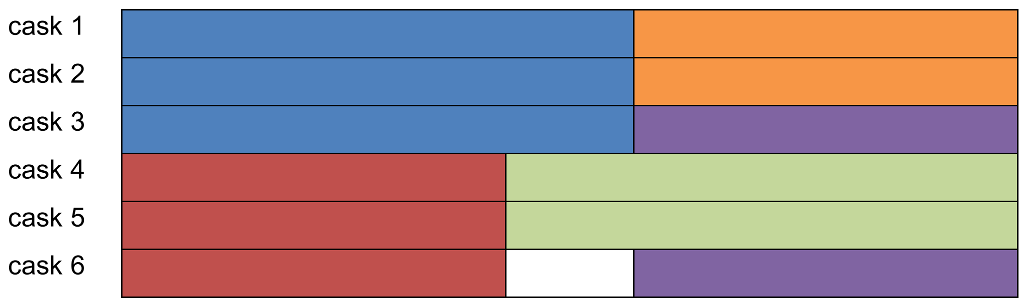

Figure 2. Strip Packing Example

A common technique in developing heuristics to address the strip-packing problem is to order the individual items in order of decreasing difficulty to pack. As a general rule, the larger (measured in height or length) items are packed first, leaving the smaller items to be packed toward the end of the process.

9. Solution Technique

TOM begins by ensuring that the necessary rail, road and waterway routes have been generated for the transportation options that are specified and the sites involved. If any routings are missing, the model calls an external problem to generate the distances required. These distances are then stored in the database.

TOM then determines the number and sizes of consists that will be used for each site in each year. Based on the consist size and near-site transportation scenario, it determines the time required for loading at the pickup site. A larger consist will involve more near-site movement as casks are moved into position, loaded, and moved to a temporary storage/marshaling location while awaiting the loading of the remaining casks.

A list of all trips is generated and is sorted in order of decreasing size of consist, and within each consist size, by decreasing cycle time. In the parlance of the strip-packing problem, this allows for a packing of the container by levels, and tends to minimize the gaps that occur as the container is filled.

The model then traverses the list, and packs the trips into the containers as they are encountered. For each trip, the model generates several lists of available starts times, each based on a different criteria.

First, a list of available loading times for the pickup location is generated. This list contains blocks of contiguous times of length at least as long as the required loading time of the trip being placed. Each block is then shortened by the length of time required for the loading operation. This resulting list is then shifted earlier by the time required to travel to the pickup location, and the resulting list indicates when a consist can leave the fleet maintenance facility (where all cycles start) and have sufficient time to load at the pickup site. If there is no sufficiently suitable block of time to load, the model will not be able to schedule this trip in the year needed, and the pickup will be deferred until the next year.

Next, the model generates a list of cask availability times. It considers the existing casks and the schedule built so far, and builds a list of times where there are enough casks available for the amount of time required by the trip for the complete cycle. Ideally, the model would like to tack this trip to the end of an existing trip of the same consist size, but it may be pulling together a group of casks that previously had been on different trips. If a sufficient group of casks cannot be found, the model will add casks (in effect, raising the height of the container by adding more strips) until a group of sufficient size can be put together.

Each of the blocks of cask availability times in the resulting list Is then shortened by the length of time required for the trip cycle so that the result is a list of blocks of times indicating when the trip can start. This list of start times is then compared with the list of start times based on the loading availabilities. If there are no blocks in the intersection of the lists, that means that there are no casks available when needed to make the loading windows at the pickup site, and new casks will have to be acquired. Sufficient casks are added to the inventory (and strips are added to the cask container) until the intersection of the two lists results in at least one feasible start time.

A similar process is followed to generate a list of escort availability and transit asset availability blocks. These lists are in turn compared to the other previously generated lists until there is at least one feasible start time in the intersection of the lists.

In most cases, there will be several possible start times that could be used, and the model chooses among these start times in order to efficiently use the casks, escort assets, transportation assets, and loading opportunities at the pickup site. This start time is assigned to the trip, and the resource lists are updated to account for the times that each resource is used.

TOM cycles through the list of trips until all trips have been scheduled.

10. Output

All of TOM's output data are stored in the SQL Server database. The output includes information at the scenario level, as well as detailed information about each trip.

The scenario level information includes the date that each asset (casks and rolling stock) was acquired, as well as the maintenance dates for each asset. Summary costs for each year and cost category are also stored at the scenario level.

A trip in TOM is a movement of SNF from the pickup location to the destination location using on consist of railcars or one convoy of trucks. Each trip is associated with a particular set of casks and transportation assets. These assets are given notional asset numbers, so that their use can be tracked through the scenario. The assets assigned to a particular trip are stored in the database. Although an asset may be acquired for a particular trip, it will likely be reused for other trips, and so assigning the acquisition cost to a particular trip would be misleading. Similarly, the maintenance required is the accumulation of usage over several trips, and assigning its cost to a particular trip would also be misleading, so the acquisition and maintenance costs are only reported at the scenario level. The operations costs, however, are attributable to a particular trip, and these are stored at the trip level. Finally, the activity dates (including transportation, loading, and unloading) are stored for each trip.

The volume of output data from TOM can be overwhelming to try to sift through by hand and is best viewed via the TSL user interface. There the user can request graphics and Excel reports that can provide formatted information at the desired level of detail.

Appendix A

Cost Calculations in the System Architecture Study

A-1. Cost Calculations in the System Architecture Study

For the system architecture study, the maximum consist size was set at three cars for a rail transport and three trailers for a truck convoy.

Locations for interim storage facilities and the repository have not yet been determined, so the analysis used a location in the eastern part of the United States as the interim storage facility, and a location in the western United States as the repository. The cask maintenance facility and fleet maintenance facility were assumed to be co-located with the interim storage facility.

While TOM is a data driven model, and the cost parameters are subject to change as the model matures, the values used in the System Architecture study are included here as an example of their range and scope, and to aid in understanding of the model.

Category	Parameter	Year	Value
Border Crossing	c_{bc} (road)	1986	1000
	c_{bc} (rail)	1986	900
Crane	c_{crt}	2006	0
	c_{cro}	2006	500
	c_{crd}	2006	0
Heavy Haul	c_{hht}	1991	10000
	c_{hhh}	1991	0
	c_{hhd}	1991	5000
Mainline Rail	c_{rms}	1991	.0003
	c_{rm}	1991	34.89
	c_r	1991	4562
	c_{rts}	1991	1.61538×10^{-6}
	c_{rt}	1991	.00403846
	c_{rw}	1991	7.95
Rail Security	c_{rsh}	2006	57
Shortline Rail	c_{slt}	1991	51000
	c_{slc}	1991	0
	c_{slo}	1991	0
Switching	c_{swc}	2006	8000
	c_{swo}	2006	5000
Truck	c_{trt} (empty)	1986	.395
	c_{trt} (loaded)	1986	1.161
	c_{trd} (empty)	1986	.00402

	c_{trd} (loaded)	1986	.00476
	c_{trs}	1986	.25
	m_{trs}	1986	120
Truck Security	c_{tsm}	2006	3.7
Waterway	c_{wtt}	2006	0
	c_{wtd}	2006	10000
	c_{wbt}	2006	0
	c_{wbd}	2006	1000