A Simulation/Optimization Framework for Locomotive Planning

Artyom Nahapetyan[1], Ravindra Ahuja[1], F. Zeynep Sargut[1],
Andy John[2], and Kamalesh Somani[2]

[1] Innovative Scheduling Inc.
Gainesville Technology Enterprise Center (GTEC)
2153 SE Hawthorne Road, Suite 128
Gainesville, FL 32641, USA
[2] Locomotive Management - CSX Transportation
3019 Warrington Street
Jacksonville, FL 32254, USA

Abstract. In this paper, we give an overview of the Locomotive Simulater/Optimizer (LSO) decision support system developed by us for railroads. This software is designed to imitate locomotive movement across a rail network, and it simulates all four major components of the system; trains, locomotives, terminals, and shops in an integrated framework. It includes about 20 charts that allow evaluating system performance using standard measures. LSO can be used by locomotive management to perform “)-1(what-if” analysis and evaluate system performance for different input data; it provides a safe environment for experimentation. We have tested the software on real data and output showed that the software closely imitates day-to-day operations. We have also performed different scenario analysis, and reports illustrate that the software correctly reflects input data changes.

1 Introduction

All US Class I railroads companies have a centralized group of managers responsible for assigning specific locomotives to specific trains around the clock, 365 days per year. Each manager is responsible for trains originating in a particular geographic region. A director presides over the managers and is responsible for the entire system. Class I railroads typically have thousands of train originations per day, and the managers must assign several thousands of locomotives to those trains. Locomotive assignment consists of assigning sets of locomotives to trains and developing routings for all locomotives while satisfying pulling power requirements of all trains and maintenance and fueling requirements of locomotives.

Many railroads use plan-based locomotive assignment as shown in Figure 1. The locomotive planning problem assigns sets of locomotives to each train in a preplanned weekly train schedule so that each train in the weekly train schedule receives sufficient power to pull its load and the total cost of locomotive

ATMOS 2007 (p.259-276)
7th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems
http://drops.dagstuhl.de/opus/volltexte/2007/1182

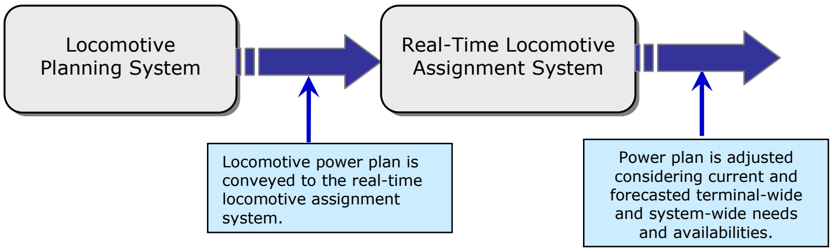

Fig. 1. Role of locomotive planning in real-time locomotive assignment.

usage is minimized (Vaidyanathan et. al. [2007], Ahuja et. al. [2005], and Ziarati et. al. [1997] and [1999]). The resulting plan must honor a variety of business rules, cannot require more locomotives than what is available in the total fleet, and must result in a plan that is relatively simple and repeatable. Another important feature of the locomotive planning problem is that some locomotives may deadhead on trains or light travel. Deadheaded locomotives do not pull the train but are pulled by active locomotives from one place to another. In the case of light travel, a set of locomotives form a group, and one locomotive in the group pulls the others from an origin station to a destination station. Deadheadings and light travels play an important role in locomotive planning, enabling extra locomotives to be moved from surplus locations to locations where locomotives are in short supply. Light travel is not limited by the train schedule, making it much faster than deadheading. However, light travel is costlier, as a crew is required and the move does not generate any revenue, as there are no cars attached.

A power plan specifies which types of locomotives will pull each train and how locomotives will deadhead or light travel to obtain the overall network-wide efficiency. The power plan is a white sheet plan that specifies the locomotive assignment to the trains. It also shows train-to-train connections for locomotives at each terminal. The plan may or may not be fueling or servicing-friendly. Each locomotive must be fueled before it runs out of fuel (typically, around 900 miles) and must be serviced periodically (either after it has traveled a certain number of miles or a certain number of days have elapsed since the last servicing). However, the power plan does not account for locomotive breakdowns, train delays, train cancelation, and adding extra trains. It assumes that all trains run on time and locomotives do not breakdown.

The solution of the locomotive planning problem serves as a blueprint to guide day-to-day real-time locomotive assignment, called tactical locomotive assignment (Chih et. al. [1993]). However, the following disruptions take place in the system and locomotive managers must further refine and adjust the locomotive assignment.

- If a locomotive is due for a regular maintenance, then managers cannot assign it to a train that takes it too far from a shop, as it cannot return before its maintenance-due date.
- Locomotives also break down, and managers must substitute them.
- While generating a locomotive plan, we assume that all trains run on time. However, trains are often delayed and sometimes are canceled altogether. As a result, terminals might not have enough locomotives to depart outbound trains.
- There are usually unanticipated, unscheduled trains that require locomotives not listed in the blueprint.
- Other unplanned events that frequently occur and must be immediately addressed as the data is communicated to the locomotive managers include train derailments, out-of-fuel locomotives, crew no-shows, severe weather, and holding outbound trains to capture priority shipments.

The decision problem faced by locomotive managers is how to change the plan with minimum disruption to the field operations while minimizing the impact on locomotive-related costs. As the operations unfold across the network, the locomotive managers must assess each piece of new data and determine how their current plan should be adjusted and locomotives be assigned to the outbound trains. The managers constantly monitor and adjust daily tactical plans to ensure efficient use of resources while maximizing the on-time operations and protecting the fluidity of the network.

In this paper, we discuss Locomotive Simulater/Optimizer decision support system, which we henceforth refer to as LSO. This decision support system simulates the movement of locomotives across a railroad network. It simulates a real-life environment in which travel times are random variables, locomotives visit shops for quarterly maintenances, and locomotives break down and go to shops for repairs. LSO simulates all of the four major resources involved in locomotive assignment: locomotives, trains, terminals, and shops. It uses the logic similar to that used by locomotive managers and directors to assign locomotives to trains: it uses historical train data to model train delays, historical locomotive data to model locomotive breakdowns, and historical data of shops to model repair and maintenance of locomotives at shops. LSO keeps track of the status, inventory, and detailed plans for individual trains by ID and date, individual locomotives, and individual terminals. As time progresses, LSO collects detailed statistics for locomotives, trains, terminals, and shops. It simulates several months of locomotive assignment in a matter of minutes. After several runs of simulation have been performed, it summarizes the results of these simulation runs and prints various reports and charts.

LSO is an invaluable tool for railroad locomotive management division to make numerous planning and strategic decisions related to locomotive operations. The ultimate goal of locomotive management is to achieve high levels of locomotive productivity and reliable train operations at the lowest possible cost. To achieve this objective, locomotive management must understand (i) the impact of strategy changes on system performance, (ii) where to focus efforts in improving efficiency and effectiveness, (iii) how many resources are required for a given level of system performance, and (iv) how to prepare for and recover from random disruptions. LSO can assist locomotive management in making these decisions. Specifically, it allows testing the efficacy and robustness of the locomotive planning and real-time locomotive assignment systems by simulating a near real-life environment. LSO also enables senior executives, locomotive directors, and locomotive managers to test various management policies, priorities, business rules, and "what-if" strategic questions such as fleet sizing, shop closures, and on-time train performance. The simulation system will show the locomotive director or locomotive managers the downstream implications of changing the system's recommendation in terms of operating cost, train delay, locomotive utilization, consist busting, missed repair commitments, mismatched power, etc. It will also assist locomotive departments in testing service design plans before accepting them and publishing them to the rest of the organization. Indeed, LSO provides a safe environment for experimentation before implementation.

Locomotive operation divisions usually use the following measures to evaluate overall performance of the locomotive assignment procedure, and LSO has about 20 reports and charts that address those measures and allow users to analyze the effect of any strategic changes from different perspectives.

- *Origination performance:* The percentage of trains departing on time from their origins per day.
- *Arrival performance:* The percentage of trains arriving at their destinations on time per day.
- *Dwell time of locomotives:* The amount of time a locomotive spends at a terminal or shop.
- *Out-of-service (OOS) rate:* The percentage of locomotives that cannot be assigned to a train due to breakdowns or maintenance.
- *Setbacks trains:* Percentage of trains held for power (or delayed) in a day.
- *Setbacks hours:* Average delay time of trains due to insufficient power.
- *Consist power plan compliance:* Percentage of trains departing with a set of locomotives specified in the power plan.
- *Locomotive utilization:* Percentage of time a locomotive actively pulls trains, deadheading or light traveling per day.

2 LSO Components and their Relationship

In this section, we provide an overview of LSO components, input and output requirements of the program, and report generating procedure. Figure 2 illustrates the relationship between different components and below we discuss these components in more detail.

LSO requires several types of inputs describing trains, locomotives, terminals and shops. The power plan provides information on trains, origin and destination terminals of the trains, scheduled active and deadheading locomotive requirements, scheduled departure and arrival times and other train related information. However, the power plan does not contain all data required by LSO,

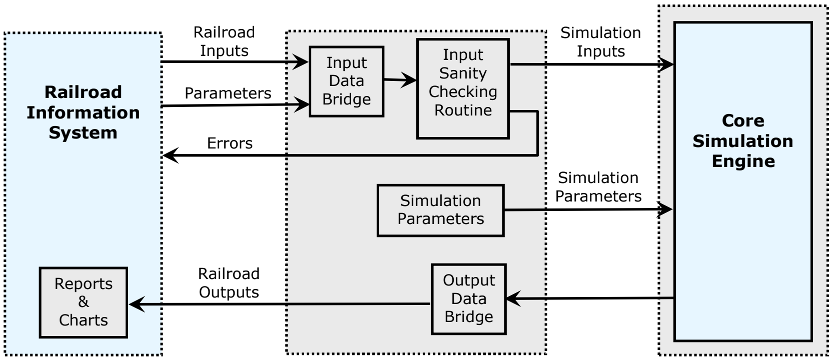

Fig. 2. Overview of LSO components.

and additional data such as properties of locomotive classes and their fleet size, description of the consist types used in the simulation, train consist priorities, probability of sending a locomotive to a shop from a specified terminal, historical travel time of trains, terminal processing distribution, locomotive breakdown rates, etc., is supplied using Excel spreadsheets or Access databases. Using the inputs, LSO performs sanity checking and transforms the data into a format consistent with tables of LSO input database. If during this process the software finds errors in the provided data then it writes corresponding massages into a log file.

After populating tables of the LSO input database, a user can specify simulation parameters and start the simulation. In the beginning, LSO sets up the initial state of the simulation and then executes events from the event list. The events imitate all activities, e.g., train arrivals and departures, locomotive failures, consist busting and terminal processing, consist assignment, locomotive light moves, shop repair procedure, etc., and record statistical data into corresponding tables in the LSO output database. During the simulation process, the module also records all events in a log file for debugging purpose.

Based on the output data, the LSO creates reports describing the overall performance of the system. Specifically, it retrieves data from the LSO output database, performs statistical analysis, and displays reports in Excel spreadsheets in the form of tables and charts. The current version of the engine generates about 20 reports describing train arrival and departure performance, percentage of delayed trains and average delay hours for each terminal, power plan compliance, out-of-service rate, events taking place at a specific terminal at a specific week, details on inventory level of the selected terminal at each simulation day, statistics on shop queue and repair time, details on light moves performed between terminals, etc.

3 Overview of Simulation Engine

Locomotive operations require the interplay of the following major resources: trains, locomotives, shops, and terminals. Figure 3 gives an overview for LSO, and its details are discussed next.

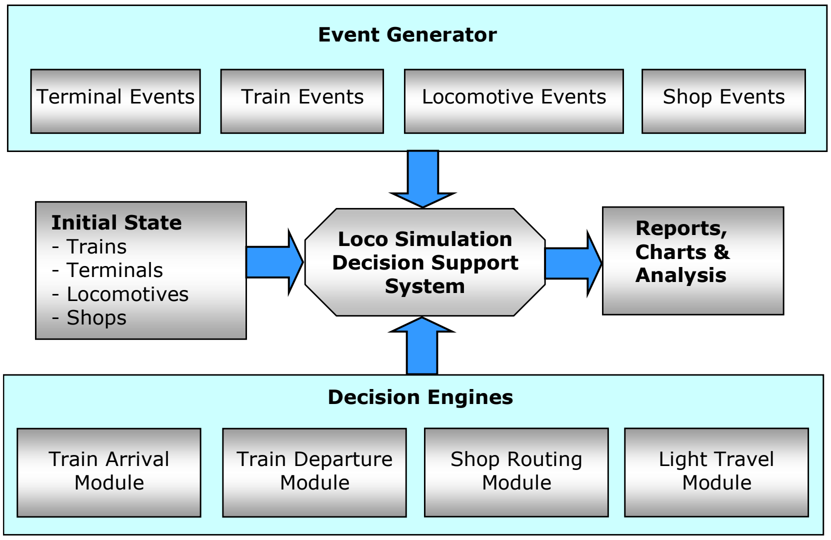

Fig. 3. Overview of LSO algorithmic logic.

We define the state of a system to be the collection of state variables associated with its entities. An event is an instantaneous occurrence that may change some state variables of the system. In the beginning of the simulation, LSO is populated (or seeded) with the current status of the trains, locomotives, terminals, and shops, which constitute the initial state variables. As events take place with respect to the four entities, trains, locomotives, terminals, and shops, the state of the system will change. The simulation engine generates train events according to the train schedule, locomotive events from the historical data of the locomotives, and shop events from the historical data of shops. LSO employs decision engines to assign locomotives to trains, route failed locomotives to the shops, and simulate light travels. It utilizes the locomotive plan as an input that could be generated either manually or using the optimal locomotive plan. As the simulation runs, the engine collects detailed statistics for locomotives, trains, stations, and shops and prints various reports and charts. LSO keeps track of the status, inventory, and detailed plans for individual trains and locomotives by ID, type and date, individual shops and terminals. The system runs on one-minute time increments and simulates trains being ordered, departed, operated over the line of road, and arriving at a destination. Each individual train is modeled deterministically. It is assumed that a particular train occurrence is ready to run at the stipulated time and takes the stipulated time to cross the line of road and arrive at the destination. If locomotives are available and ready by the scheduled departure time of the train, no locomotive delay is attributed to that train, even if it runs later than scheduled. If locomotives are not ready at the time the train is ready, locomotive delay is calculated from the ready time until the train gets locomotives and departs the terminal according to the simulation. The system simulates locomotive breakdowns and the repairs of locomotives at shops. The locomotive simulation assigns locomotives to trains and reposition locomotives via light engine moves.

The length of the simulation period is an input of the system, and the system is designed to simulate pre-specified months of normal operations. Specifically, user can enter the start and end dates and time of the simulation and then run the simulation for the specified time horizon. Users may want to repeatedly simulate the specified time horizon to collect sufficient observations to see system-average results over an extended period of time. The simulation is provided with a fleet of locomotives that can be assigned to the trains. We realize that given the initial state of the system, it requires some warm-up time to reach a steady state before any observations can be taken. We thus need to account for some warm-up period in the simulation, and when determining statistics, we should ignore the data for the warm-up period.

4 Main Simulation Modules and Engines

In this section, we discuss main modules that are necessary to run LSO. We first describe the initial state setup and then engines used in the simulation; subsequent subsections provide a short description of the corresponding components and their functionality.

4.1 Initial State of LSO

Before proceeding to the simulation, LSO creates locomotive, train, terminal, and shop entities and initializes the state of the system and counters. We next discuss each of these procedures in detail.

Entity Construction.

- *Locomotives:* LSO creates a certain number of locomotive entities according to the locomotive class fleet size. Each entity has different attributes describing the locomotive ID, type, class, horsepower, axel count, manufacturer, average time between breakdowns, and other features of the locomotive.
- *Terminals:* LSO considers all origin and destination terminals of the trains and creates corresponding entities. Each entity has attributes describing terminal ID and terminal processing and consist busting time distributions.

- *Trains:* LSO creates a train entity for each train described in the train run table. The attributes of the train describe the train ID, type, priority, tonnage, origin and destination terminals, scheduled departure day and time, list of preferred and accepted consists, planned deadheading locomotives, travel time distribution, and other features of the train.
- *Shops:* LSO creates shop entities according to their location. Each shop has attributes describing the shop ID, type, number of spots, service time distribution, and other features.

Initial State Setup.

- *Locomotive Initial Location:* LSO takes a snapshot of the power plan at a specific time, e.g., Sunday midnight, and distributes the pool of available locomotives among terminals. Specifically, for all trains that are on the way to their destination terminal it creates corresponding consists described in the power plan, assigns them to those trains and triggers train arrival events for the trains at appropriate times. Next it looks at the power plan to count the number of consists at each terminal at the time of the snapshot. These consists constitute the initial inventory at terminals. Finally, it randomly distributes the remaining locomotives, if any, among terminals that have shops.
- *Populate List of Events:* LSO maintains a list of events, which is sorted according to the time they should occur. Some events, e.g., tactical repositioning events, should be triggered at certain points of the planning horizon, and others, e.g., train arrival events, are triggered by other events during the simulation. Before proceeding to the simulation, LSO populates the list by the following known events.
  - *Train departures*
  - *Train arrivals*
  - *Tactical repositionings*
  - *Locomotive Q-maintenances and breakdowns*
  - *Consist assignments*
- *Initialize Simulation Counters:* LSO assigns initial values for all counters used in the simulation.

4.2 Main Modules of LSO

Train Arrival Module: Depending on the condition of active and deadheading locomotives, train arrivals require different actions at the terminal. If no locomotive in a consist fails upon arrival, then the consist can be assigned to an outbound train. However, if at least one of the locomotives fails, the consist must be busted, the failed locomotives are sent to shops, and remaining locomotives and consists can be used to pull other trains. Before a consist is assigned to an outbound train, it also should go through certain terminal activities, which we refer to as terminal processing.

LSO imitates locomotive breakdowns using certain locomotive failure rates. Locomotive Q-maintenance and Breakdown module assigns a “red” status to failed locomotives and locomotives that are due for quarterly maintenance. Train Arrival module checks active and deadheading consists of the train up on arrival. If one of the locomotives in the consist fails, the module creates a consist busting event, which will bust the consist and process failed and good locomotives separately, i.e., route failed locomotives to shops and send good locomotives to terminal processing. The time it takes to bust a consist can either be a random number generated from a pre-specified distribution or a fixed time interval. If arriving locomotives do not have “red” status, then we imitate terminal processing of locomotives, i.e., main track, main line fueling, truck fueling, or servicing. Terminal processing takes a random amount of time generated from a pre-specified distribution. After terminal processing, locomotives are ready for train assignments, and they are stored at the terminal.

Consist Assignment Module: Consist assignment of outbound trains is performed by locomotive managers based on the availability of the preferred consist, availability of accepted consists, consist busting time, and priorities of the outbound trains. Specifically, in the locomotive shortage environment, locomotive managers prefer assigning available locomotives to trains with higher priority. However, if a lower-priority train has been delayed for a certain time, then they try to find a consist to depart the train. Train on-time departure also depends on the availability of the consist given in the power plan, and locomotive managers might delay the train for a certain time if the consist is not available. Mangers continuously monitor consist availability at terminals (i.e., consist inventory, arriving consists, and consist failure) and adjust consist assignment of the departing trains.

Consist Assignment module analyzes the locomotive availability at the terminal. Specifically, it considers all currently available locomotives and locomotives that have already departed on trains and will arrive at the terminal during a certain time horizon. Using collected data, the module tries to find a proper consist for selected trains. During this assignment process, it also takes into account a user-specified amount of time a train can be delayed to assign the preferred consist, i.e., the consist specified by power plan. If a proper consist has not been found for the train, the module considers the consist busting option, i.e., tries to create a consist from available locomotives. If a consist has been assigned to a train before its scheduled departure time, then the train departs on time; otherwise, the train is delayed until a proper consist is assigned to the train by following runs of Consist Assignment module. The module also handles planned locomotive deadheading and light moves. Specifically, if a train has such requirements, then module tries to assign those locomotives to the train. If the number of available locomotives is insufficient, then the module departs the train on time with the available set of locomotives.

Train Departure Module: On-time departure of trains depends on the availability of proper consists, and ideally each departing train should have a proper consist assignment prior to the scheduled departure time. However, if there are not enough locomotives available to power all outbound trains, locomotive managers assign available consists to higher-priority trains and delay lower-priority trains. The managers usually make consist assignment decisions in advance, and at the scheduled departure time trains either have a consist to depart or they should be delayed.

Tactical Repositioning Module: During the real-time locomotive assignment procedure, locomotive imbalances at terminals are created; that is, some terminals may have surplus locomotives while other terminals may face locomotive deficits. These imbalances are created due to various reasons including surplus and deficit locations designed in the power plan, locomotive breakdowns, which create surpluses at shops and deficits at other terminals, train annulments, second section of trains, violation of power plan consist assignments, variance in train travel times. etc. Locomotive managers employ unscheduled deadheading and light travel options to move locomotives from surplus terminals to deficit terminals to restore locomotive balance in the network.

Since LSO imitates the real-time locomotive assignment process, it creates locomotive imbalance at terminals as well. Specifically, if there is an imbalance between the number of inbound and outbound locomotives at a terminal, then the terminal either accumulates certain types of locomotives or encounters a shortage of locomotives. Tactical Repositioning module looks ahead to analyze the inventory level for a user-specified time horizon (from several hours to several days) and determines the surplus and deficit terminals. During this process, it imitates assignment of inbound locomotives to outbound trains using a logic similar to the Consist Assignment module. If a terminal has a shortage of locomotives, LSO computes the demand of the terminal for each locomotive type. After identifying surplus and deficit locations as well as supply/demand of terminals, the module tries to satisfy the demand of deficit terminals by surpluses at surplus locations by solving a multicommodity network flow problem (Ahuja et. al. [1993]). Since speed is of critical issue in simulation, we solve the multicommodity problem heuristically. The solution of this problem yields the tactical repositionings necessary to meet the demand.

Locomotive Q-Maintenance and Breakdown Module: Class I railroads operate thousands of locomotives, and each day some of them break down due to mechanical or weather-related reasons. In the simulation, we assume that locomotives can fail whether they are active, i.e., pulling a train, or inactive, i.e., deadheading or waiting at a terminal. The locomotive failure rate describes the number of times a locomotive class breaks down during a year, and it is an input of the simulation. Although locomotive failures can occur on the way to the destination terminal, locomotive managers can route a locomotive to a shop only when the train arrives at its destination terminal.

According to FRA requirements, each locomotive must undergo preemptive maintenance at some designated shop on or before 92 days have elapsed since its last maintenance. Otherwise, the locomotive must be shut down and moved as a deadhead. This maintenance is also known as a quarterly maintenance or Q-maintenance. When the due date of the Q-maintenance is near (within 4-5 days), locomotive managers try to assign the locomotive to a train that departs to one of the shops. Depending on the manufacturer of the locomotive, it should be sent to an appropriate shop.

Shop Processing Module: Locomotive assignment to a shop is performed by locomotive managers based on (i) the type of repair it requires, (ii) travel time to the shops, and (iii) the number of locomotives at the shops. Different shops have different number of spots to perform repairs; therefore, the capacity and output rates of shops are different. Some shops maintain different spots for broken locomotives and locomotives that are due for Q-maintenance. If the shop is congested, locomotives wait in a queue upon arrival.

In the simulation, we assume two types of repairs: breakdowns and Q–maintenance. After arriving at a shop, a locomotive should wait in the corresponding queue to be processed. If the locomotive is due for Q-maintenance, then the module adds the locomotive at the end of the Q-maintenance queue. Otherwise, the locomotive joins the queue of broken locomotives. Both queues are simulated according to first-in-first-out logic. In the simulation, we allow each shop to maintain three types of spots, i.e., spots for broken locomotives, spots for Q-maintenance, and spots that can perform both repairs. When a spot is ready to seize the next repair request, this module checks the type of the spot and proceeds according to one of the following two cases: (*i*) the spot can perform only one of the repairs, and (*ii*) the spot can perform both repairs. After finishing the repair, the locomotive leaves the shop, goes through terminal processing, and joins the locomotive inventory at this terminal. In addition, the module triggers the next Q-maintenance and breakdown events if necessary.

5 LSO Reports and Charts

During the simulation process, LSO records statistical data into output tables of its database, and based on the collected data constructs various charts and tables describing overall performance of the system. Current version of the software generates about 20 charts and tables using Excel spreadsheets, and in this section we provide an overview of most important reports.

The train on-time performance is one of the most important statistics, and LSO provides several charts that allow analyzing the train on-time performance from different perspectives. The chart in Figure 4 describes the percentage of on-time train departures and arrivals for each day of the simulation. A user can either specify a terminal for which he/she would like to draw the chart or view the chart for all terminals. In the later case, we compute and display the average percentage over all terminals. Figure 5 describes another chart that shows the

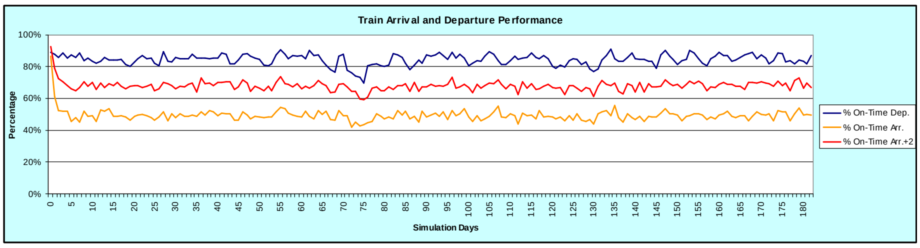

Fig. 4. Train Arrival and Departure Performance.

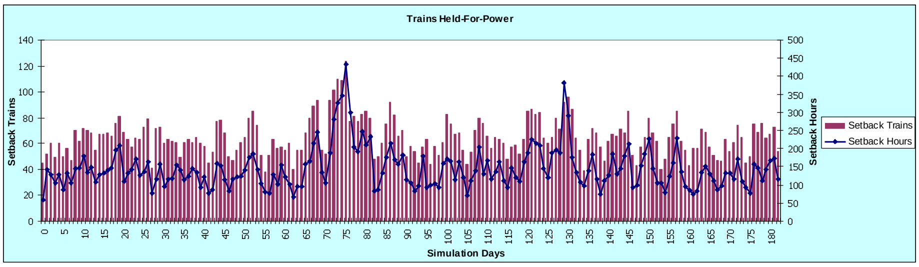

Fig. 5. Trains Held for Power.

total number of delayed trains and the total number of delayed hours for each simulation day. As before, the user can either select a terminal to view the chart or display the data for all terminals. In addition, the software provides two charts that describe average percentage of on-time train departures and arrivals, percentage of delayed trains and average delay hours for each terminal.

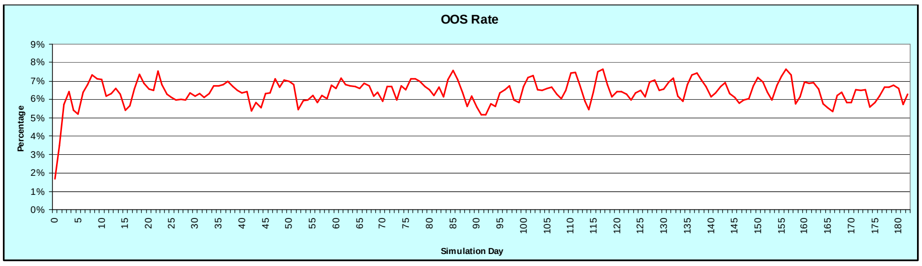

Fig. 6. Out-of-service Rate.

Locomotive managers also employ out-of-service (OOS) rate and percentage of power plan compliance to evaluate the overall performance of the system. Specifically, OOS rate measures the percentage of locomotives that cannot be assigned to trains due to breakdowns and Q-maintenances. Power plan compli-

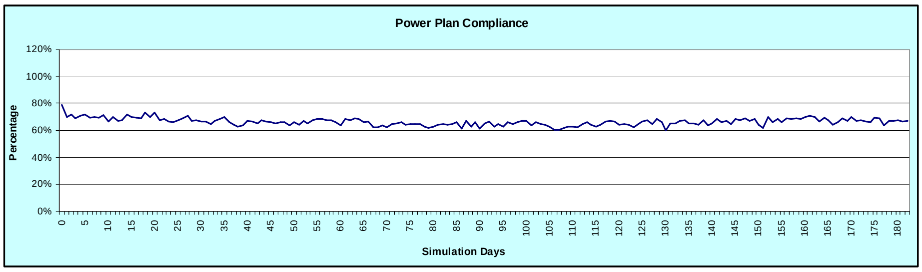

Fig. 7. Power Plan Compliance.

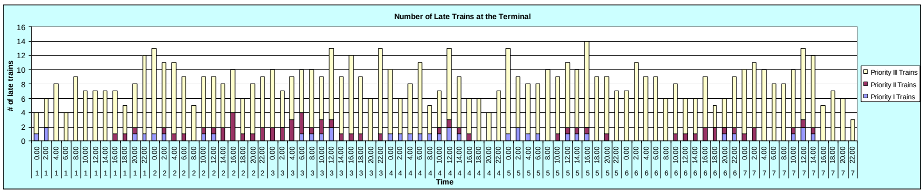

Fig. 8. Number of Late Trains at the Terminal.

ance measures the percentage of trains that have not been assigned the consist specified in the power plan. Charts in Figures 6 and 7 describe the corresponding measures for each simulation day.

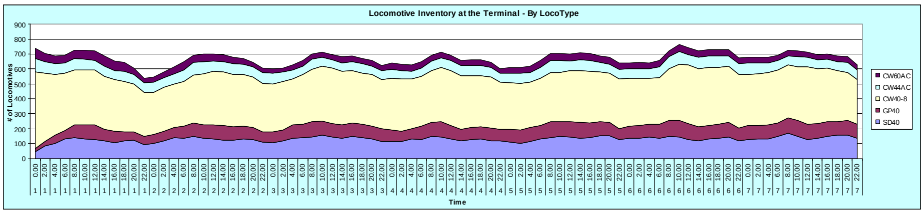

Fig. 9. Locomotive Inventory at the Terminals by Locomotive Type.

In addition to the average numbers, users can choose to view details for each simulation week. The chart in Figure 8 describes the number of late trains in each two-hour bucket for the fourth simulation week. The chart displays the data for each train priority. As before, the user can choose to view the chart for a specific terminal. The software also provides a similar chart for train delayed hours. Users also can look at locomotive inventory of the terminals. Figure 9 describes the locomotive inventory at terminals for each locomotive type for the same fourth simulation week.

In addition to the charts above, LSO generates reports that describe all events taking place at a terminal during a specific week, light moves performed during the simulation, dwell time of locomotives at a terminal and at each simulation day, and statistics on shop repair and queue times.

6 Performing "What-If" Analysis Using LSO

In this section, we describe how the software can be used to perform "what-if" analysis on the system. To illustrate this, we have designed five scenarios that help to understand the influence of different parameters on key measures used by locomotive managers to evaluate overall system performance. In each case, we simulate the process by executing several runs and then present average results in the charts.

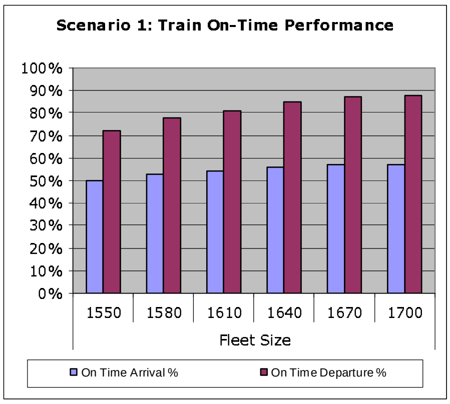

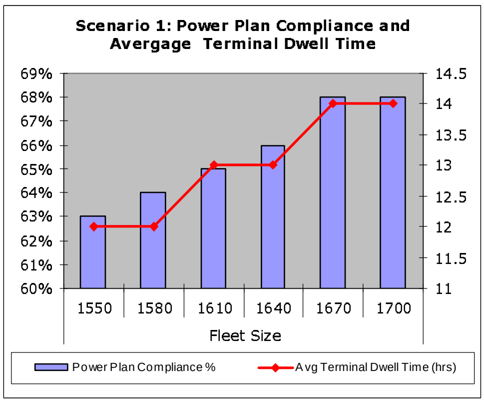

Fig. 10. On-Time Train Performance, Power Plane Compliance and Average Terminal Dwell Time for Different Locomotive Fleet Sizes.

In Scenario 1, we analyze the influence of locomotive fleet size on train on-time departures and arrivals, power plan compliance, and average terminal dwell time. In this experiment, we proportionally change the locomotive fleet size for all five locomotive types used in the simulation. Charts in Figure 10 show that by increasing the locomotive fleet size, we improve train on-time performance as well as the power plan compliance. Since less locomotives are required to move between terminals to restore terminal imbalances, it also increases the terminal dwell time of locomotives.

When locomotive managers assign locomotives to outbound trains, they might delay a train for several hours to assign the consist described in the power plan. Scenario 2 is designed to capture the influence of delay hours on the same three measures used in the previous scenario, i.e., train on-time performance, power plan compliance, and average terminal dwell time. In this experiment,

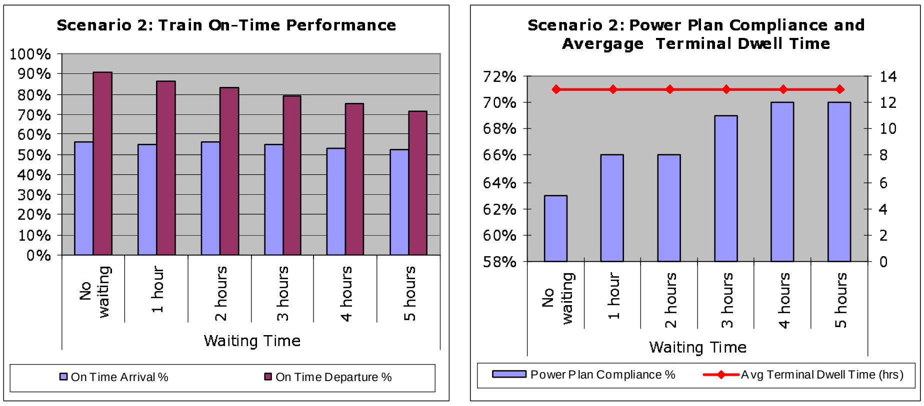

Fig. 11. On-Time Train Performance, Power Plan Compliance and Average Terminal Dwell Time for Different Waiting Hours for Right Consist.

we employ the same delay hours for all three priority trains. In Figure 11, we can see that by increasing the waiting time for the right consist, i.e., consist described in the power plan, we improve power plan compliance but worsen on-time train performance. Note that we do not count these delays towards the terminal deficit; therefore, average number of light moves does not change and the average terminal dwell time of locomotives remains the same.

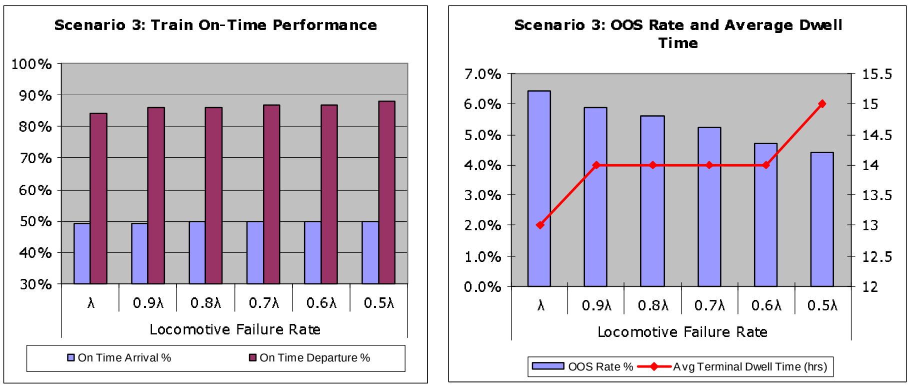

Fig. 12. On-Time Train Performance, Out-Of-Service Rate and Average Terminal Dwell Time for Different Values of Locomotive Failure Rates.

In the next scenario, Scenario 3, we analyze the influence of locomotive failure rates on on-time train performance, out-of-service rate, and locomotive dwell time at terminals. In this experiment, we proportionally change failure rates of

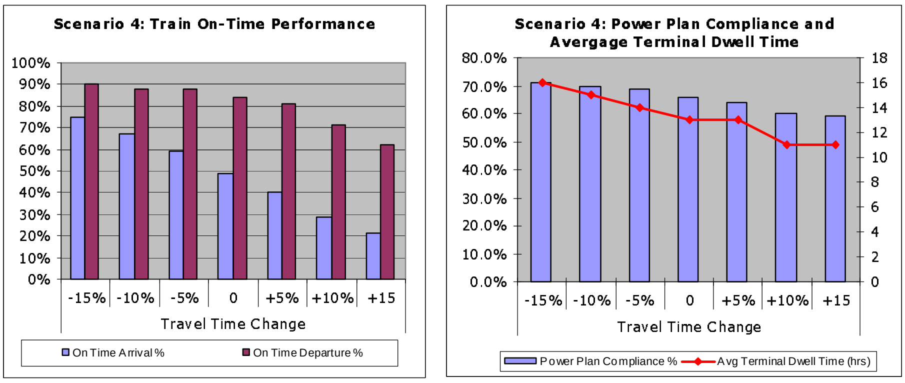

Fig. 13. On-Time Train Performance, Power Plan Compliance and Average Terminal Dwell Time for Different Values of Train Travel Time.

all locomotive classes used in the simulation. Charts in Figure 12 show that by deceasing the locomotive failure rate we reduce the OOS rate of locomotives as it is expected. On the other hand, reducing locomotive failure rate increases the locomotive dwell time at terminals and slightly improves the train on-time performance.

Next, in Scenario 4, we analyze the influence of train velocity on system performance. Specifically, in this experiment, we increase or decrease the train travel time by a certain percentage. Charts in Figure 13 depict that a higher travel time worsens the on-time train performance as well as the power plan compliance. If trains do not arrive on time, outbound trains do not have enough locomotives to depart. As a result, the module considers moving locomotive to those location; therefore, it reduces the dwell time of locomotive at terminals.

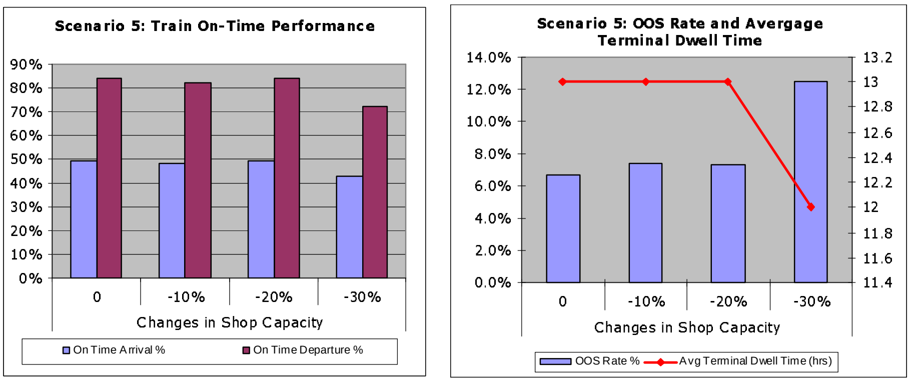

Fig. 14. On-Time Train Performance, OOS Rate and Average Terminal Dwell Time for Different Shop Capacities.

In the last scenario, Scenario 5, we run the simulation for different values of shop capacities. In this experiment, we gradually reduce shop capacities of all 10 shop locations we consider in the simulation. In Figure 14, we can see that a small change in shop capacities slightly changes the OOS rate and does not change on-time train performance and dwell time of locomotives. However, when the capacities are reduced beyond a certain threshold, shops cannot repair all the locomotives which accumulate in queues. As a result, the system shows a huge jump in the OOS rate, reduction in locomotive dwell time at terminals and on-time train performance.

7 Summary and Conclusions

In the paper, we have discussed LSO software, which simulates the movement of locomotives across a railroad network. Specifically, it simulates the locomotive assignment to outbound trains, train arrivals and departures, locomotive breakdowns and maintenances, locomotive repair procedure at shops, terminal processing, tactical repositioning, etc. We have tested the software on real data obtained from CSX Transportation, one of the Class I railroads. The results show that the statistical data of simulation is very close to the figures obtained from day-to-day operations, and the software closely imitates the real-time locomotive assignment and locomotive movement in the network. The software is able to simulate six months of operations in about three minutes. All charts generated in the reports show a very short warm-up period after which the system reaches a steady state.

We have designed several scenarios to test the software and analyze the influence of different input parameters on the system performance. In the paper, we have presented some of these results. In all scenarios, the output data has correctly reflected the changes in the input parameters, and the software shows a stable performance in terms of running time, warm-up period and convergence to a steady state.

References

Ahuja, R.K., Liu, J., Orlin, J.B., Sharma, D., Shughart, L.A.: Solving real-life locomotive scheduling problems. Transportation Science **39** (2005) 503–517.

Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network Flows: Theory, Algorithms, and Applications. Prentice Hall, Englewood Cliffs, NJ (1993).

Chih, K.C., Hornung, M.A., Rothenberg, M.S., Kornhauser, A.L., 1990. Implementation of a real time locomotive distribution system. In Computer Applications in Railway Planning and Management, T.K.S. Murthy, R.E. Rivier, G.F. List, J. Mikolaj (eds.), Computational Mechanics Publications, Southampton, UK, pp. 39-49.

Vaidyanathan, B., Ahuja, R.K., Orlin, J.B., and L.A. Shughart: Real-life locomotive planning: New formulations and computational results. To appear in Transportation Research B (2007).

Ziarati, K., Soumis, F., Desrosiers, J., Gelinas, S., Saintonge, A.: Locomotive assignment with heterogeneous consists at CN North America. European Journal of Operational Research **97** (1997) 281–292.

Ziarati, K., Soumis, F., Desrosiers, J., Solomon, M.M.: A branch-first, cut-second approach for locomotive assignment. Management Science **45** (1999) 1156-1168.