Optimal Routing of Locomotives to Shops for Quarterly Maintenances

Andy John
Assistant Vice President – Locomotive Planning & Utilization
CSX Transportation, Jacksonville, FL
Andy_John@csx.com

Ravindra K. Ahuja
University of Florida &
Innovative Scheduling
Gainesville, FL
ravi@InnovativeScheduling.com

InnovativeScheduling

Collaborators

CSX Transportation

- Bernie McCall
- Bill McKee
- Dave Peddicord
- Bill Dawson

Innovative Scheduling

- Artyom Nahapetyan
- Arvind Kumar
- Sapon Tanachaiwiwat
- Zeynep Sargut

Locomotive Shop Router (LSR): Business Problem

- Each locomotive must go to a shop for Quarterly (Q) maintenance within 92 days of its previous maintenance. We want locomotives to reach shops just in time – there is a cost associated with being early or late.

- Shops have different profiles:
  - Maintenance capabilities: Short maintenance (every other 92 day Q), Long maintenance (every 184 days), Annual maintenance (once a year)
  - Shop capabilities may also be constrained by locomotive characteristics, i.e. manufacturer, class of power or special equipment.
  - Each shop has limited capacity for doing these maintenances.

- Business problem: Which shop should each locomotive be routed to, when should they reach the shop and what trains should they take to maximize locomotive productivity while maintaining a balanced shop workload within the current shop capabilities and capacity constraints.

- As time progresses and actual train and locomotive performance is updated, what changes need to be made to re-optimize the solution.

- It is a dynamic and very complex operations research problem.

Before Locomotive Shop Router (LSR)

- 20 locomotive managers assign locomotives to trains. Q maintenance decisions are a subset of the overall locomotive assignment problem.
  - Five territories, each with a separate locomotive manager.
  - Two shifts per day.
  - Each manager had limited visibility to the collective impact of their own and other manager's decisions on locomotive productivity and shop arrival patterns.

- Locomotives do not always reach shops on time. If they are late, they become FRA "dead" and cannot be used as pulling power for trains. These are "past due Q" locomotives.

- The number of locomotives arriving at a given shop could vary widely from one day to another and might be inconsistent with the capacity.
  - Fluctuations in the number of locomotives received impacts shop and locomotive productivity.

Locomotive Shop Routing: Decision Variables

- Determine the shop each locomotive due for Q maintenance should do.

- Determine the trip details to go to that shop.
  - Take Train 1 from Location 1 to Location 2.
  - Take Train 2 from Location 2 to Location 3
  - Take Train 3 from Location 3 to Location 4 (a shop location)

- Generate recommendations which locomotive managers may or may not follow.

- Should work like a GPS which generates routes and we can choose to ignore it. If we do that, LSR should recalculate based on the previous decision.

One Source of Problem

- If we count the number of locomotives with the same maintenance date in the entire network, then they vary widely from one day to another.

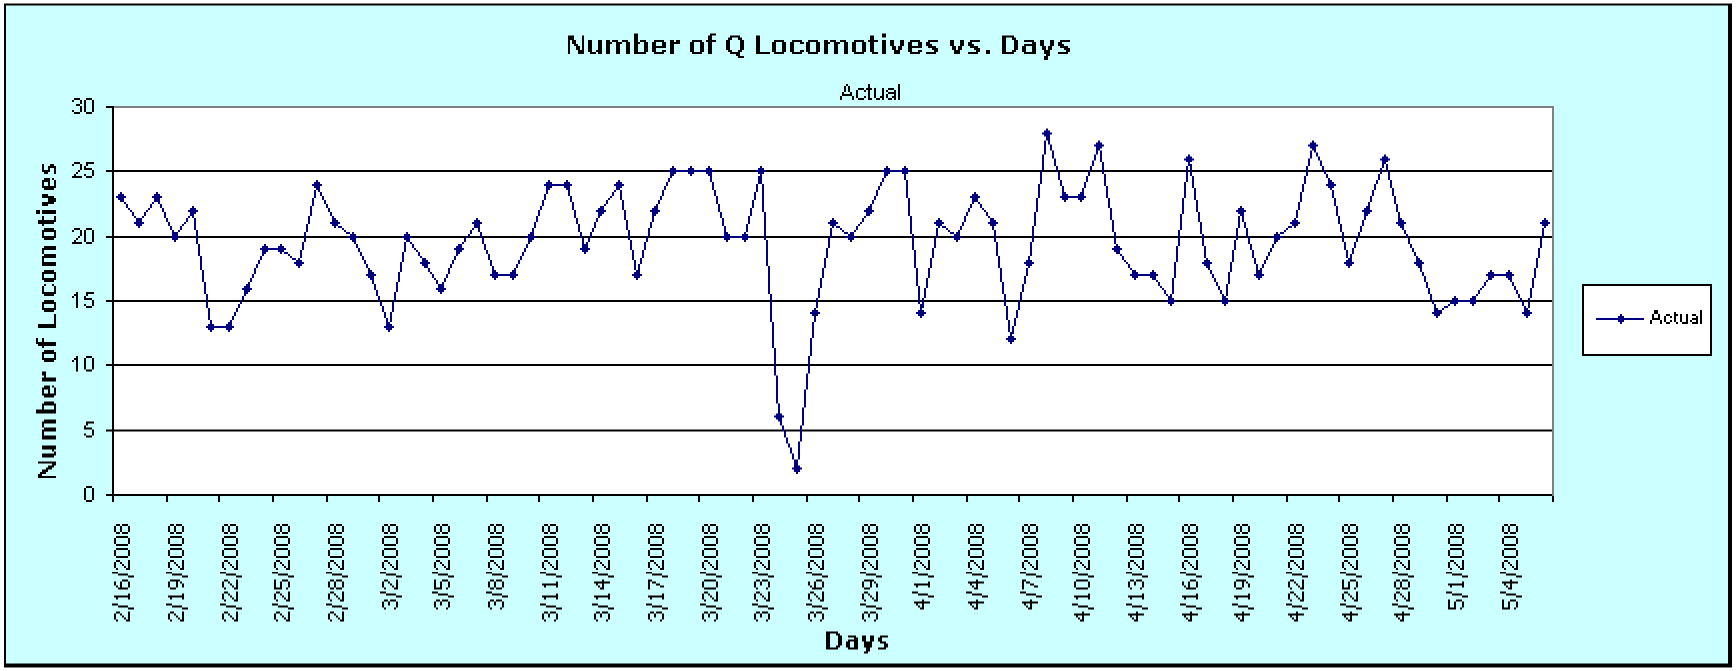

Schedule-Date Balancing

- We need an algorithm that pulls the maintenance dates of some locomotives so that the same number of locomotives are due for Q maintenance each day.

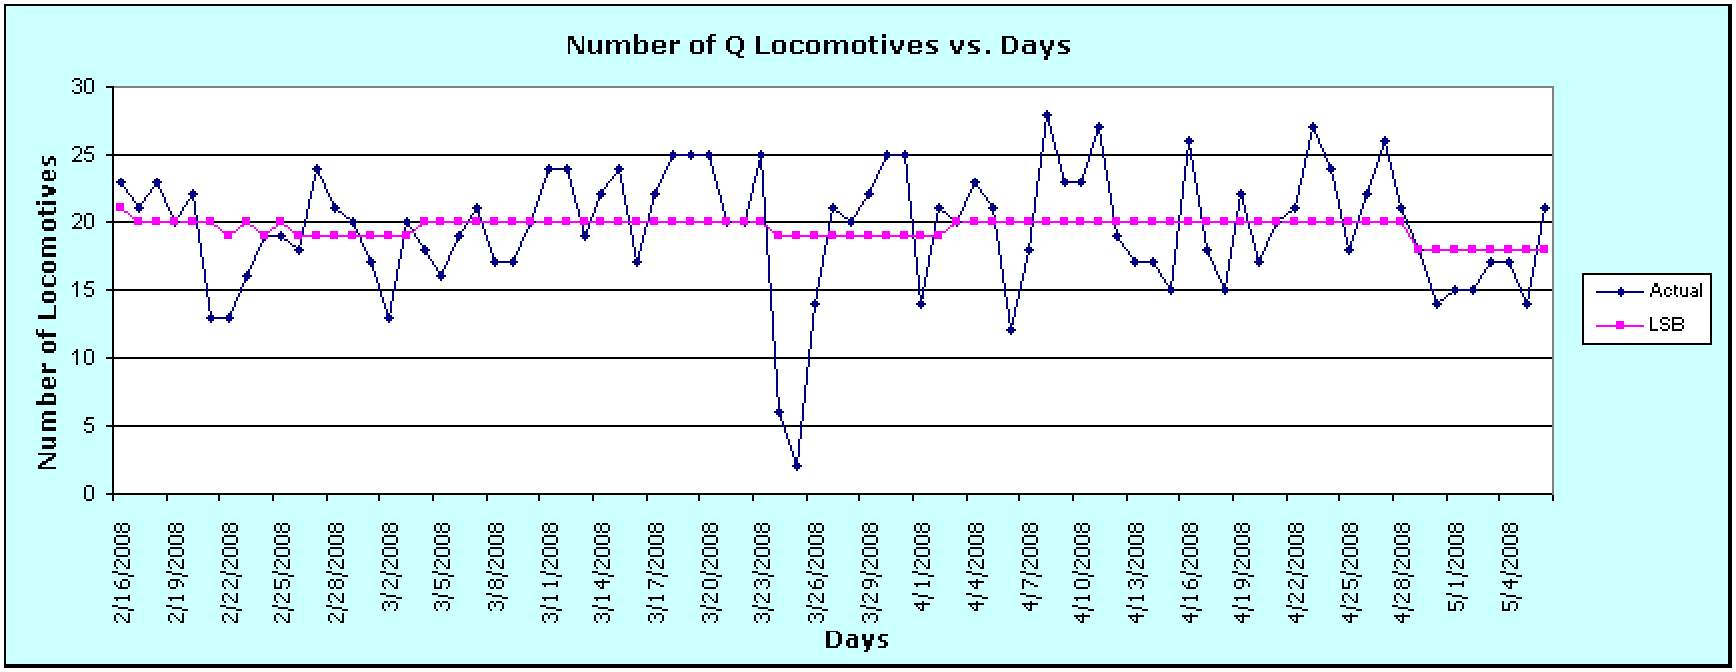

Two Decision Problems

- Locomotive Shop Balancing (Long Term)
  - Adjust the schedule date of locomotives for the next three months so that the same number of locomotives are due for quarterly maintenance each day.
  - Do it every day

- Locomotive Shop Routing (Real-Time)
  - Generate the route of each locomotive due for Q maintenance.
  - Update the route every hour.

Locomotive Schedule Date Balancing

- We have a daily target T for the number of locomotives due for Q maintenance.
- We need to pull the schedule dates of some locomotives so that the number of locomotives due for maintenance each day is equal to T.
- Do not pull a date by more than K days.
- Minimize the penalty:
  - Penalty cost of pulling back locomotives' Q dates
  - Penalty cost of deviating from the target number of locomotives per day
- Some additional business constraints.

Locomotive Schedule-date Balancing Model

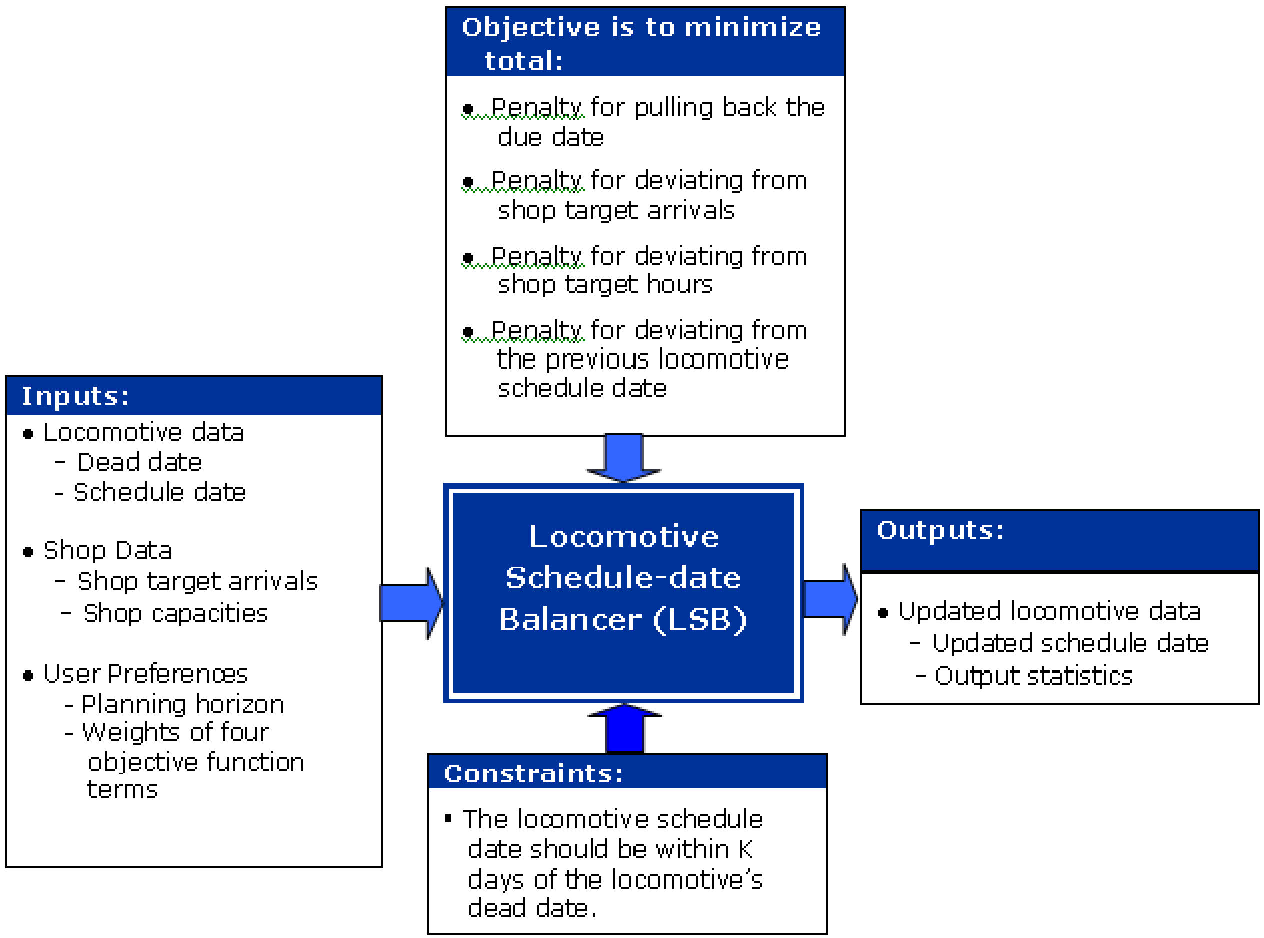

Computational Results

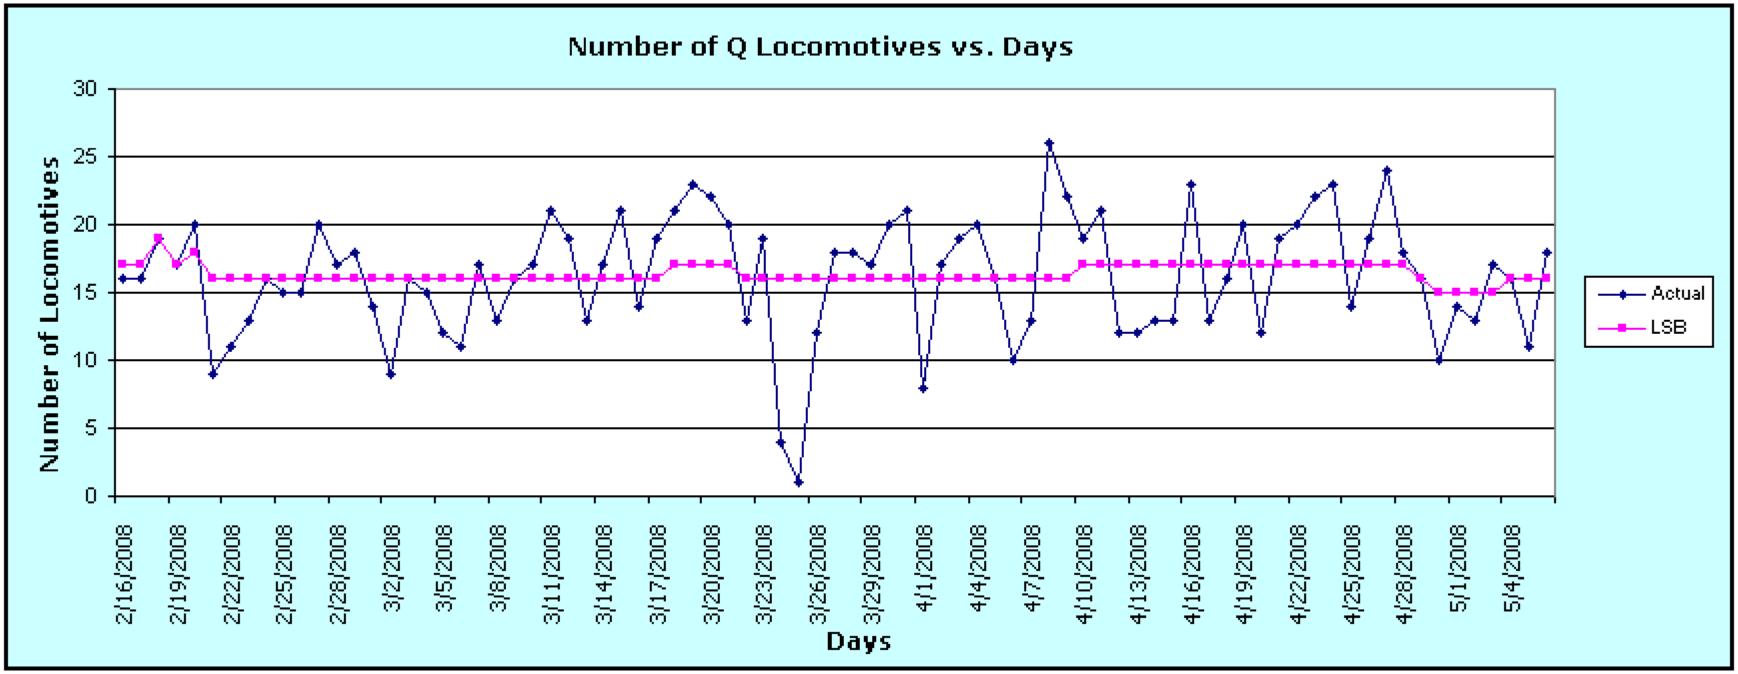

- On the average, a locomotive is pulled back by about one day.

Locomotive Shop Routing Problem

- Locomotive Shop Balancing (Long Term)
  - Adjust the schedule date of locomotives for the next three months so that the same number of locomotives are due for quarterly maintenance each day.
  - Do it every day once a week

- Locomotive Shop Routing (Real-Time)
  - Generate the route of each locomotive due for Q maintenance.
  - Update the route every hour.

Three Types of Locomotives

- Locomotives in the scheduled networks.
  - Shop router generates recommendations for all locomotives in the network.

- Locomotives in the local network.

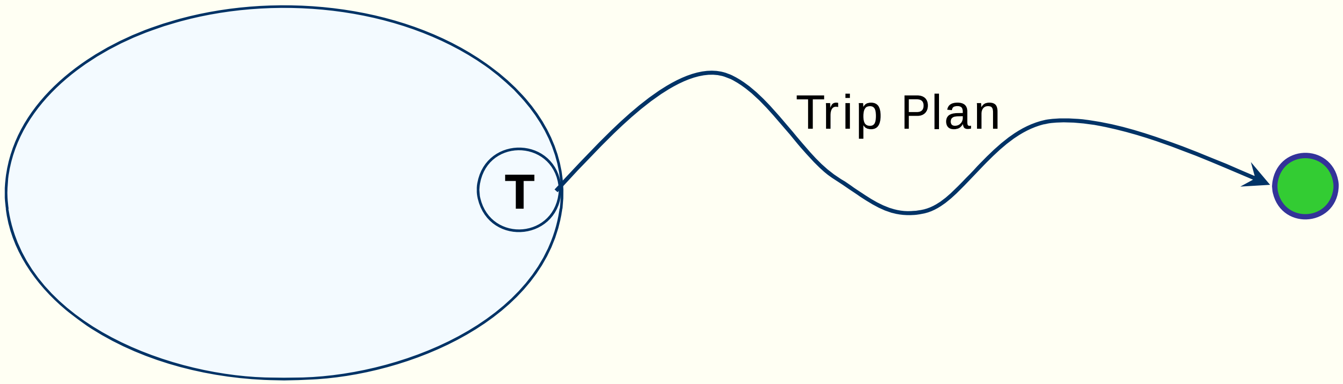

- Locomotives in the unscheduled network.
  - Shop router does not generate recommendation for these locomotives.

The Locomotive Shop Router Problem

Inputs:

- Locomotive data
  - Dead date
  - Schedule date
  - Status
- Shop Data
  - Shop target arrivals
  - Shop capacities
- Train Data
  - Train profiles
  - Current authorized trains
  - Power plan
- Locomotive-Shop preferences
- User Preferences
  - Planning horizon
  - Weights of different objectives

Objective is to minimize the total:

- Penalty for deviating from shop target arrivals
- Penalty for exceeding the shop capacity
- Penalty for not distributing Annuals (LongQs) evenly among the shops
- Penalty for deadheading more healthy locomotives than the number in the power plan
- Penalty for deviating from the previous train and shop assignment

Locomotive Shop Router (LSR)

Constraints:

- Each locomotive should be assigned to a route going to a shop
- Each locomotive has a set of shops where it can be maintained

Outputs:

- Updated locomotive data
  - Updated schedule date
  - Shop assignment
  - Locomotive-train assignment
- Reports
  - Routing details for each locomotive
  - Locomotive-shop assignment details
  - Projected number of arrivals at each shop

Assignment Network

- The model assigns locomotives to shops. An arc denotes the potential assignment of a locomotive to a shop-day.

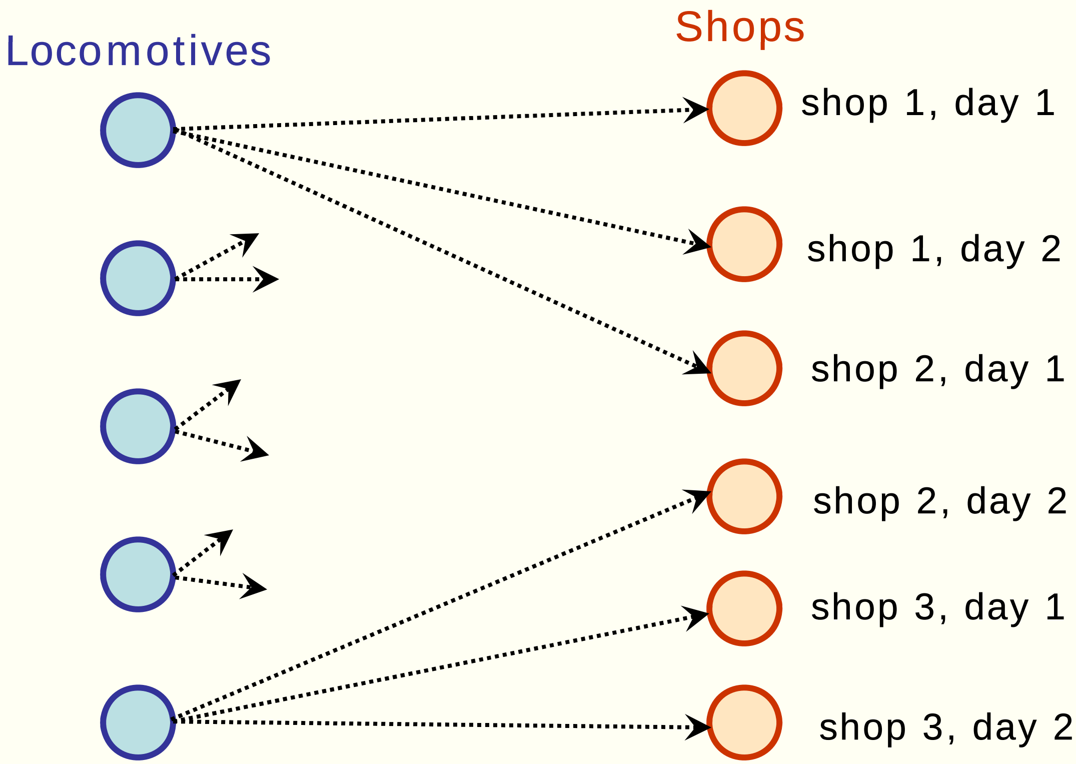

- The resulting problem is a minimum cost flow problem with some side constraints.

Resulting Approach

- Construct the Train (Space-Time) Network for potential routes of locomotives to shops.

- For each potential arc in the Assignment Network, determine a shortest path in the Train Network.
  - This path gives the trip plan of the locomotive to the shop.
  - The cost of this path becomes the cost of the corresponding arc in the assignment network.

- Solve a MIP problem in the Assignment Network to determine the assignment of locomotives to shops.
- The running time of this algorithm is about five minutes.

Shop Router Architecture

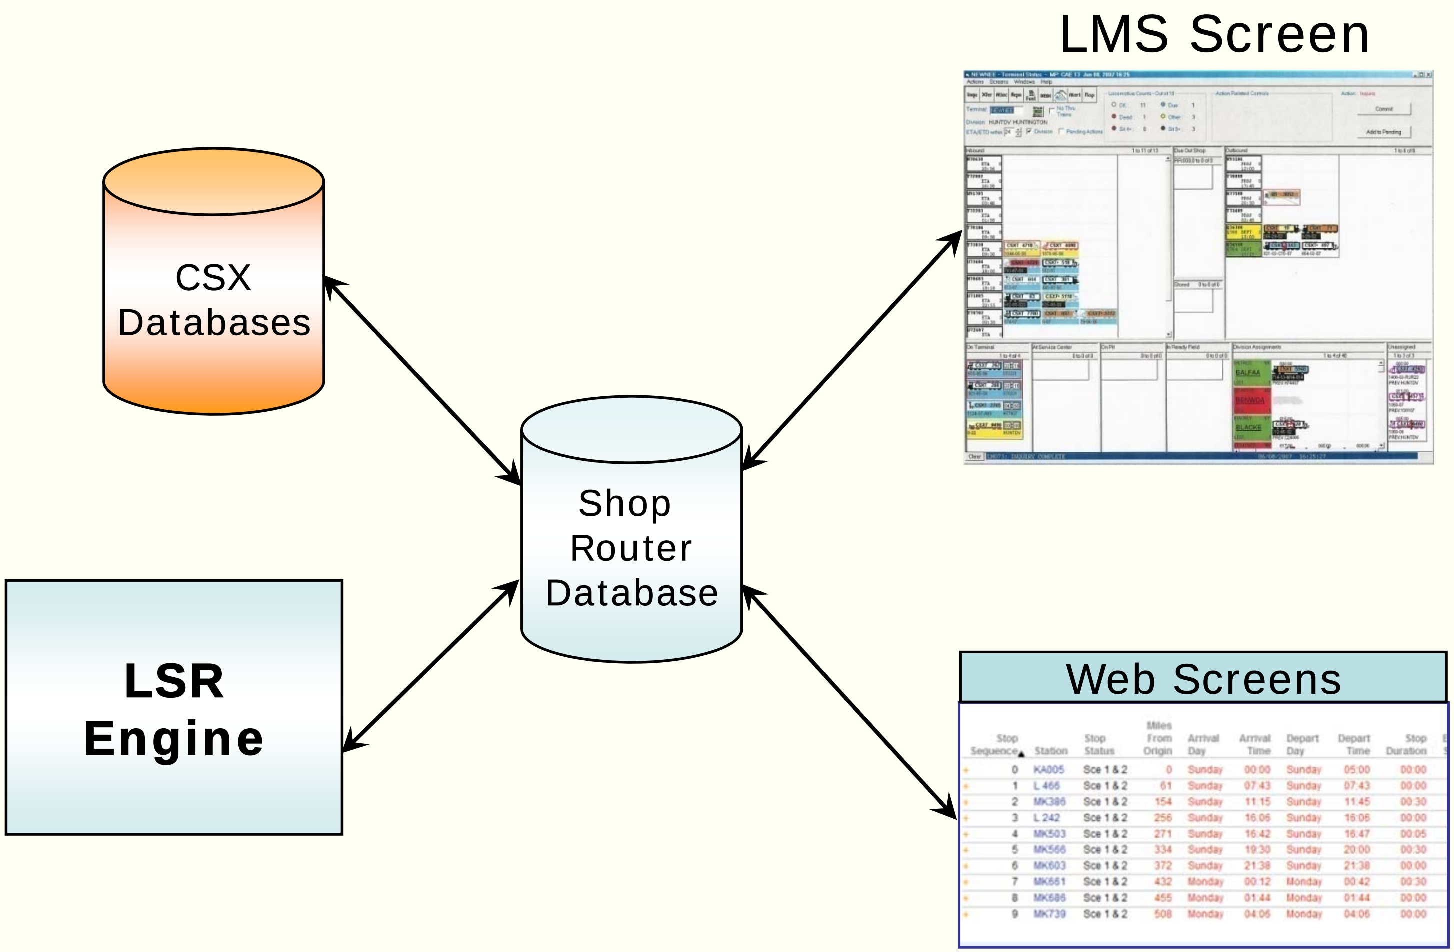

- LSR runs every hour and generates recommendations for locomotive managers. It works like a GPS.

Locomotive Shop Router: Current Status

- The LSR went into production at CSX in September 2008 and results are encouraging.

- We have observed noticeable improvements in past due Q locomotives.

- We have also observed improvements in the daily flow of locomotives to shows. We are exceeding daily shop capacities only infrequently.

Lessons Learnt

- Data analysis is critical. Data guides us in the right direction about the algorithm design.
- Data quality and sanity checking is very important.
- Implementability is critical. Satisfying all or most business constraints is necessary.
- Incremental solutions are needed. We cannot change solutions dramatically from one run to another run.
- User control important.

InnovativeScheduling

OPTIMIZATION IN ACTION