Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Crew Decision Assist: System for Optimizing Crew Assignments at BNSF Railway

Brian Roth, Anantaram Balakrishnan, Pooja Dewan, April Kuo, Dasaradh Mallampati, Juan Morales

To cite this article:
Brian Roth, Anantaram Balakrishnan, Pooja Dewan, April Kuo, Dasaradh Mallampati, Juan Morales (2018) Crew Decision Assist: System for Optimizing Crew Assignments at BNSF Railway. Interfaces 48(5):436-448. https://doi.org/10.1287/inte.2018.0963

informs
http://pubsonline.informs.org/journal/inte/

INTERFACES
Vol. 48, No. 5, September–October 2018, pp. 436–448
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Crew Decision Assist: System for Optimizing Crew Assignments at BNSF Railway

Brian Roth,[a] Anantaram Balakrishnan,[b] Pooja Dewan,[a] April Kuo,[a] Dasaradh Mallampati,[a] Juan Morales[a]

[a] BNSF Railway, Fort Worth, Texas 76131; [b] McCombs School of Business, University of Texas at Austin, Austin, Texas 78705

Contact: brian.roth@bnsf.com, http://orcid.org/0000-0002-8533-4864 (BR); anantb@utexas.edu (AB); pooja.dewan@bnsf.com (PD); april.kuo@bnsf.com (AK); dasaradh.mallampati@bnsf.com (DM); juan.morales@bnsf.com (JM)

https://doi.org/10.1287/inte.2018.0963

Abstract. Rail is the preferred mode of transport for many categories of freight because of its low cost and energy efficiency. Rail accounts for approximately 40%, measured in ton-miles, of all freight movements in the United States. To maintain their competitive advantage and effectively utilize their large investments in rail infrastructure, freight railroad companies place considerable emphasis on improving the cost efficiency of their operations. Crew costs, including payments to crew members and expenses for crew repositioning and lodging at stations away from the home base, constitute a significant portion of railroad operating expenses. This paper describes the development of an optimization model and solution method and the implementation of a system called "crew decision assist" to support crew scheduling at BNSF Railway. The work was motivated by the company's desire to replace its current manual crew-planning process with a systematic and effective approach. Preexisting crew-scheduling models did not adequately capture all the options and constraints that arise in practice, such as the option to use extra crew members or policies to jointly reposition engineers and conductors. We, therefore, developed a tailored model and solution approach that incorporates various practical features and requirements for crew assignment at BNSF and accounts for uncertainty in train schedules. Our decision support system, based on this method, interfaces with existing information systems to retrieve the necessary data and quickly generate effective crew-deployment plans when train schedules change. The system was recently introduced for use by crew planners at BNSF and has already reduced crew costs, yielding estimated annual savings of several million dollars.

Keywords: railroad operations • crew scheduling • decision support • integer optimization • neighborhood search

Introduction

The rail freight industry is an important component of the transportation sector in many countries and is vital for their national economies. Rail transport offers the advantages of lower cost (for shippers) and higher energy efficiency compared with other modes of transport (American Association of Railroads 2017). In recent years, the railroad industry has also taken significant steps to make train movements safer. In the United States, railroads handle approximately 40% of all freight movements, measured in ton-miles (Federal Railroad Administration 2017), and carry a wide range of materials such as coal, crude oil, chemicals, and many agricultural, industrial and consumer products. BNSF Railway, a leading U.S. Class 1 freight railroad, operates an average of 1,400 trains per day over its vast rail network, which contains 32,500 miles of track and spans 28 states in the continental United States and three Canadian provinces. In 2016, BNSF Railway had operating revenues of over $19 billion, transported 9.7 million carloads, and invested $3.9 billion to expand and upgrade its infrastructure, including tracks, signaling systems, terminals, and rolling stock. To be competitive, freight railroad companies focus on ensuring good utilization of their capital-intensive resources and operating in a cost-efficient manner. BNSF employs approximately 20,000 crew members to operate trains. Because crew costs account for a significant portion of train operating expenses (i.e., crew cost is the highest among all components of BNSF's train operating costs), the effective deployment of train crews is an important priority for railroad companies.

Crew-deployment decisions include assigning crew members to operate the scheduled trains, as well as decisions regarding whether and when to reposition crews between stations, also called deadheading, to ensure adequate crew availability at each station. These decisions determine the payments made to crew members, their layover costs (e.g., lodging and transportation expenses at stations away from their home bases), and expenses for deadheading. For purposes of crew deployment, the rail network is divided into crew districts. Real-time crew-deployment decisions are made by crew planners, each responsible for assigning crews to trains in multiple (5–10) districts. BNSF Railway's previous method for planning crew-to-train assignments and deadheading was manual. The manual crew-planning process was time consuming and was based on individual planners' experience and intuition rather than on a uniform and structured approach that carefully accounts for the cost implications of crew-planning decisions.

The need for decision support for crew scheduling was driven by considerations of financial impact, business continuity, crew safety, and quality of life among the three main stakeholders—senior management, crew planners, and crew members. First, because crew costs constitute a significant portion of overall operating expenses, even a small percentage improvement in these costs can translate to considerable annual cost savings. Therefore, senior managers want to ensure that the crew-planning decisions are close to optimal in terms of minimizing total crew-deployment costs. Second, efficient crew planning is key to uninterrupted and resilient rail network operations. Third, railroad companies take the safety of their operations very seriously, and periodically report safety performance measures with their financial metrics. Further, the Federal Railroad Administration monitors and regulates railroad safety. Finally, and importantly, crew work cycles, including rest periods and time away from home, affect crew members' quality of life. These requirements are complementary rather than conflicting. For example, ensuring adequate crew rest and balanced workloads can lead to safer and more cost-effective crew assignments. Based on these considerations, senior executives created technical and operational teams to develop a strategic solution that would provide cost savings and also meet organizational safety metrics and crew workload policies. Crew planners, facing ever-changing train and crew lineups in the many districts they manage, wanted a system that could be quickly and seamlessly integrated with the current databases and planning platform. BNSF corporate leadership, therefore, asked its operations research group to explore the development of a computerized tool that optimizes crew-assignment and deadheading decisions to minimize total crew-deployment costs. This initiative led to our modeling, solving, and implementing an optimization-based approach, called the "crew decision assist" (CDA) system, to support real-time crew-assignment and deadheading decisions. As train schedules are updated, this system dynamically retrieves the necessary data from existing databases to quickly generate near-optimal crew deployment plans for the next few days. The project entailed collaboration between the operations research group, crew planners, and information systems personnel to frame the problem, gather data, develop and test optimization algorithms, validate the outputs, and implement the approach for use in practice. BNSF crew planners began using the system for real-time decisions in January 2015; initial results indicate that this system has reduced crew costs by several million dollars annually compared with costs based on previous manual decisions.

Crew Scheduling for Freight Trains

Freight railroad companies partition their networks into crew districts, each demarcated by two crew-change stations, one at each end of the track segment covered by the district, and they require that trains passing through a district be operated by crew members assigned to that district. Hence, crew-planning decisions decompose by district. We focus on single-ended districts in which all crew members who are assigned to the district have as their base one of the two stations, called the home station. We refer to the other station as the away station. Each train requires a crew team consisting of one crew member from each occupation (i.e., engineers and conductors for our application context). We are given the set of *trains* that are scheduled to travel through the crew district over the planning horizon (typically, 48 hours). For each train, the schedule specifies its movement direction and the times at which it enters and exits the district. We also know the initial location and rest status (at the start of the planning horizon) of each crew member assigned to the district.

Crew Movements and Transfers

When a train enters a district at either end, assigned crew members from that district operate the train to the other station in the district, disembark at this station, and rest before traveling back. Depending on the intensity of train traffic in the two directions in the district, a station may have, at certain times, too few or too many available crew members relative to the number needed to operate the trains scheduled to traverse the district from that station. In anticipation of such situations, crew members may need to be deadheaded (repositioned) from one station to the other. A crew member can be deadheaded using one of three modes: on a scheduled freight train, on a scheduled public-transportation service (e.g., passenger train), or using a dedicated taxi (i.e., van). These modes vary in timing and costs. Scheduled freight and passenger trains have fixed departure and arrival times, whereas taxis can be dispatched at any time except during blackout windows (i.e., periods of the day in which taxis cannot be dispatched because of safety reasons). To accommodate surges of traffic, the railroad maintains a roster of extra crew members at the home station. Union agreements specify that extra crew members can only be used when no regular crew member is available. We refer to the group of extra crew members as the extra pool, and the group of regular crew members as the regular pool.

We define a trip as any movement of crew members between the two stations, either to operate a train or to deadhead. A connection refers to the transfer of a crew member from an inbound trip to an outbound trip at a station. Before this transfer, the crew member must rest for a minimum period, which depends on the durations of the inbound and outbound trips. There are three types of possible connections at the away station:

• *Full-rest* connections require the rest time between the inbound trip's arrival and the outbound trip's departure to be at least 10 hours plus a lead time (typically, 1.5 hours). To address quality-of-life concerns, we have ensured that the rest time must not exceed a specified maximum value (e.g., 24 hours).

• *Short-rest* connections are possible when the sum of travel times for the inbound and outbound trips does not exceed 12 hours. In this case, the intermediate rest time must be at least four hours plus a lead time.

• *Flip* connections have no rest-time requirement and apply when the crew member's total round-trip elapsed time, from and back to the home station, does not exceed 12 hours.

The types of connections allowed are district specific; for example, some districts permit full-rest only, some also allow flips, and some allow only flips. At the home station, only full-rest connections are allowed, with no upper limit on the rest time. For full- and short-rest connections at the away station, crew members rest at a hotel.

Crew-Deployment Restrictions

Crew-assignment decisions must meet several requirements. The plan must ensure that each scheduled train is assigned the required complement of crew members, one from each occupation, to operate the train. Deadhead decisions must satisfy capacity constraints (i.e., trains and taxis can accommodate only a limited number of deadheading crew members). The connections at each station must meet the appropriate rest requirements, as we discuss above, and must also follow crew-dispatching priorities, which we call crew-rotation rules. These rules specify the chronological sequence in which crew members depart the station. In some districts, a simple "first in, first out" (FIFO) rule applies: for each occupation, crew members must depart the station in the same order in which they arrived. In other districts, a first out, first out (FOFO) rule applies: crew members must depart the station in the order in which they departed the other station on their previous inbound trip. Most districts impose the additional requirement, specified by union agreements, that deadheads consist of an equal number of crew members from both occupations; we refer to this requirement as "occupation pairing."

Crew-Deployment Costs

The costs for crew deployment fall into four broad categories: (1) train operator costs, (2) deadheading costs, (3) layover costs, and (4) costs for using extra crews. Train operator costs refer to the payments, also called trip rates, to members of the crew team for operating a train. Deadheading costs include a trip rate for each deadheaded crew member and any additional cost for transporting the crew member. Deadheading a crew member on a scheduled train entails no additional transportation cost; however, deadheading via other modes does incur such costs (e.g., taxi fixed cost or fare for public transport). If a crew member must wait for a long period at the away station, that crew member receives a payment, which we call a heldaway cost. Specifically, crew members receive a per-hour heldaway payment for every hour they must wait beyond a specified limit, typically 16 hours. Depending on the connection type, costs are also incurred for lodging and meals for crew members who need to rest at the away station. Connections at the home station do not incur such costs because crew members can go home to rest.

Prior Work

Although operations researchers have long studied crew-scheduling problems in public-transportation contexts such as airlines (Barnhart et al. 2003, Gopalakrishnan and Johnson 2005) and passenger trains (Caprara et al. 1998, 2001, 2007; Abbink et al. 2005), the models and methods for these problems do not directly apply to U.S. freight railroads because of the following differences in rail crew-deployment requirements and options. Unlike periodic passenger-transport services that follow fixed, repeating schedules, the itineraries and schedules of freight trains vary from week to week depending on the required cargo movements, network congestion, and resource availability. So crew planning for freight railroads requires dynamic decision-making. Freight railroad companies also differ in their crew operating policies. The duty cycles for train crew members (from home to away and back) are structurally simpler than crew duty cycles for other transport services; however, crew-to-train assignments are subject to dispatching priorities and other rules that do not apply, for example, to airline crew scheduling, and can be difficult to model. For train crews, planners also have a wider choice of deadheading options and can use extra crews when regular crew members are not available to operate a train. Because of these distinctive features, crew scheduling for U.S. freight railroads requires tailored models. The literature on freight train crew scheduling is relatively sparse (e.g., Gorman and Sarrafzadeh 2000, Vaidyanathan et al. 2007, Sahin and Yüceoglu 2011, Jütte et al. 2011). The models proposed in these papers consider only simple dispatching priorities (e.g., FIFO) and do not incorporate some features and options (e.g., fixed cost for taxi deadheads or the option to use extra crews) that must be considered in practice.

First Approach: Integer Programming

Figure 1 provides a visual representation, over a time-space network, of the problem components for rail crew scheduling. This representation facilitates our discussion of the model formulation for the crew-planning problem. Figure 1, left, shows a representative network over which the problem is defined, and Figure 1, right, shows a sample solution. The time-space network has two spatially separated layers (vertical lines), one corresponding to each station; the time axis runs from top to bottom at each station, with the top (bottom) representing the start (end) of the planning horizon. Points on each line correspond to arrival or departure events at the station.

The network has two main types of directed arcs: trip arcs and connection arcs. Trip arcs represent travel on trips (as train operator or deadhead) from the home station to the away station or vice versa. Each trip arc starts at the location and time at which the trip begins and ends at the location and time at which the trip terminates. The solid lines correspond to scheduled trains, and the dotted lines to deadheads (using taxis or public transport). Connection arcs represent the transfer of crew members from an inbound trip to an outbound trip at a station. The time between the arrival of the inbound trip and departure of the outbound trip, reflected in the length of the arc, is the rest time of the crew member making the connection. We also include a source node to model the use of extra crews and a sink node to account for crew members who remain at the away station at the end of the planning horizon. Finally, although taxis can be dispatched at any time, Figure 1 shows only a discrete set of arcs for taxi deadhead trips. As we will explain later, based on the train departure and arrival times, we can prune the set of candidate-taxi start times to such a discrete set without losing optimality.

Formulation

The rail crew-scheduling problem requires assigning the crew team to operate each train, deciding when to use extra crew members, specifying the transfers or connections from inbound to outbound trips at

Figure 1. (Color online) The Crew-Planning Problem Can Be Represented as a Time-Space Network

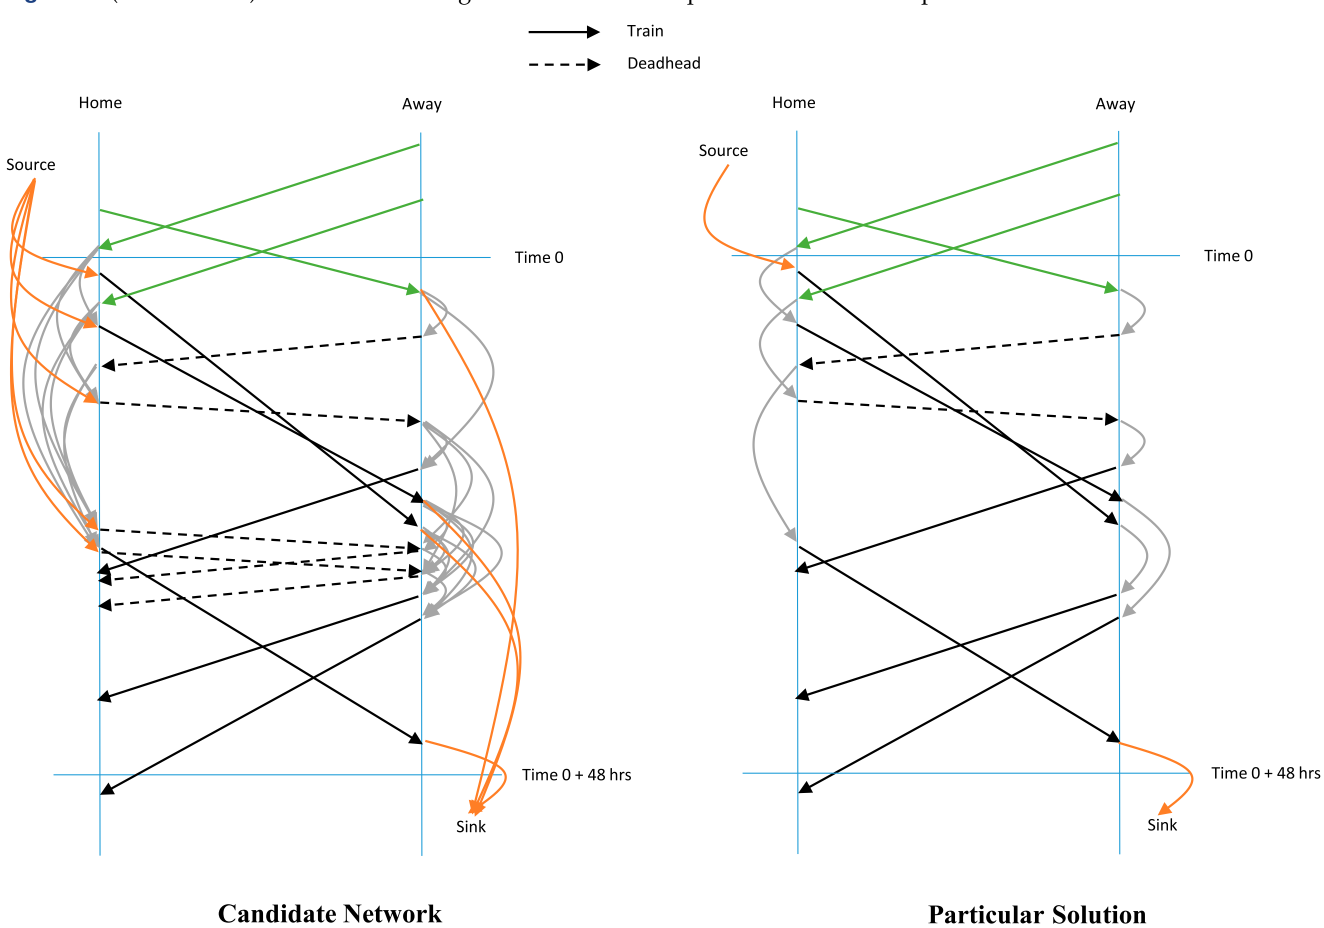

a station, and deciding how many crew members of each occupation to deadhead between stations at what time and using what mode. We can interpret these decisions as flows of crew members on the arcs of the time-space network shown in Figure 1. However, the crew-planning problem is much more complex than a standard minimum-cost network-flow problem because it involves routing multiple "commodities" (one for each combination of occupation and crew pool) and requires many constraints, in addition to flow conservation, to capture the restrictions on crew assignments and connections. Moreover, because using taxis incurs fixed costs, additional binary decision variables are needed to represent whether to use a taxi deadhead trip.

Our integer programming formulation for the rail crew-scheduling problem uses the following three main sets of decision variables:

1. Integer variables representing the flows of crew members on the trip arcs.

2. Binary variables indicating which connection arcs are selected.

3. Binary variables indicating which taxi trip arcs are selected.

The objective function consists of minimizing the total crew-deployment cost during the planning horizon. The coefficients of the decision variables in this function provide wide latitude in modeling different expenses, including crew payments, fixed and variable deadheading costs, connection-dependent layover costs, and costs for using extra crews. We can also incorporate various penalties to capture planner and operational preferences. The model's constraints include equations to link the connection and flow variables, and constraints to ensure that each train is assigned the required complement of crew members, deadheading capacities are not violated, and the connections satisfy rotation rules. In addition, the model contains constraints to appropriately capture taxi fixed costs and meet operational requirements, such as crew pairing and restrictions on using extra crews. The appendix contains our detailed integer programming formulation of the rail crew-scheduling problem.

Our model incorporates several features that previous models (with the partial exception of Balakrishnan et al. 2016) do not consider. First, we allow for the possibility of using extra crews, and capture their costs as well as the policy of using these crews only when no regular crew member is available. Second, we permit using taxis as a deadheading mode, and consider their timing, fixed cost, and capacity. Third, we model the occupation-pairing requirement to concurrently deadhead crew members from both occupations. Fourth and finally, we model the crew-rotation rules that govern the dispatching priorities at each station as constraints on judiciously chosen subsets of connection variables. The model requires many such restrictions to ensure that the solutions generated by the model are implementable in practice. Except for Balakrishnan et al. (2016), other papers on crew scheduling (for trains or other transportation settings) do not consider the general version of crew-rotation rules that our model incorporates.

Data Preparation

Instantiating the decision model for a given district and planning horizon requires first gathering information on train schedules, crew availability, public-transport schedules, various cost parameters, and crew-deployment rules. In addition, we must also determine (1) the potential times at which taxis may be used, which we call candidate-taxi trip generation and (2) the set of feasible inbound-to-outbound trip connections at the away station that satisfy rest requirements, which we call candidate-connection set generation.

Taxis are available as needed to deadhead crew members; the optimization model needs to choose their timing. Instead of treating the taxi dispatch times as continuous decision variables, we judiciously select a discrete set of times for potential (i.e., candidate) taxi trips. This set covers all possible taxi options that any optimal solution may use, and stems from noting that a taxi is required for one of two reasons: (1) to deadhead a crew member away to address a shortage at the away station, or (2) to deadhead a crew member home to avoid a surplus at the away station or a shortage at the home station. Accordingly, for each connection type at the away station, we create a just-in-time home-to-away candidate taxi that is timed to connect, with minimal rest, to each train trip departing the away station. Likewise, for each train trip arriving at the away station and for each connection type from this trip, we create a just-in-time candidate taxi from the away station to the home station that can deadhead the crew members who operated the trip after minimal rest. Taxis departing during blackout periods are omitted; instead, we create candidate taxis departing immediately before and after each blackout period.

Improving Performance

The rail crew-scheduling model is difficult to solve optimally using commercial solvers because it has tens of thousands of variables and millions of constraints. We incorporated several enhancements to reduce the variables and constraints. First, because all crew members in a given occupation and pool have the same capabilities (and costs), our model treats crew assignments as flows rather than considering the trip assignments for each individual crew member. After solving the model, we can easily map the solution to individual trip assignments. Second, a conventional approach to modeling the problem is to include decision variables for crew connections at both the home and the away stations (as we show in the appendix). However, we can reduce the problem size by omitting the connection variables at the home station because no connection costs are incurred at that station. The assignments only need to satisfy the minimum full-rest time requirements, which we accomplish using constraints on the cumulative inbound and outbound flows at the home station. Third, instead of defining a connection variable for every pair of inbound and outbound trips at the away station, we apply a preprocessing method to prune the list of candidate connections based on the connection rules. Specifically, for each inbound trip, we determine the subsequent outbound trips or successors to which a crew member can feasibly transfer, given the minimum and maximum rest-time requirements. Similarly, for every outbound trip, we determine its feasible predecessors. We can omit connections from an inbound deadhead trip to an outbound deadhead trip because cost-minimizing solutions do not use such deadhead-to-deadhead connections. By applying these rules, we not only implicitly incorporate the rest-time requirements but also significantly reduce the number of decision variables. Finally, we developed strategies to strengthen the model and improve its linear programming lower bound by formulating tighter versions of the rotation-rule constraints. This enhancement also reduces the number of constraints in the model. Even after implementing these model improvements, practical instances of the problem formulation are quite large. As an example, for a district with 131 trains and 153 candidate deadheads in the planning horizon, the reduced formulation requires 12,976 decision variables and contains 3.5 million constraints.

A Faster Algorithm: Neighborhood Search

Although we were able to solve the integer program to optimality for most problem instances, the solution times were quite inconsistent, with many instances requiring much more time than is practical for real-time use. In part, the large number of constraints needed to model the rotation rules contributed to the long computational times. Moreover, as a deterministic model, our integer program does not account for uncertainty in train schedules; incorporating uncertainty (e.g., in a stochastic program) greatly increases problem difficulty and solution times. Therefore, we shifted our focus to developing an effective heuristic method.

Key Principle

The crew-scheduling problem is well suited for a heuristic approach because the primary decisions are the deadheading choices. Given a set of deadheads to use, we can readily complete the solution (i.e., determine the corresponding crew-to-trip assignments, extra crew usage, and connections) by applying the crew-rotation rules and rest requirements. We know the initial state of the system (i.e., initial location and rest status of each crew member); therefore, we can track the inventory of rested crew members at each station as the train and deadhead trips arrive at or depart from the station. Then, applying the crew-rotation rule yields the inbound-to-outbound connections, as well as the needed use of extra crew members (when there are crew deficits at the home station). That is, given a deadhead plan, we can fully determine the values of all the decision variables in the rail crew-scheduling model. Therefore, we can use a two-stage approach of first selecting deadheads and then completing the solution. By separating the deadhead-selection decisions from the crew-assignment decisions, this approach overcomes the computational challenges posed by the rotation-rule constraints in the integer program.

Neighborhood-Search Procedure

We use a variable neighborhood-search procedure to explore the space of candidate-deadhead plans. In the following discussions, we refer to any given choice of deadheads as a deadhead plan (d-plan). Given a starting d-plan, the procedure first selects one neighborhood type from among different types, which we define below, and then generates a neighborhood around the incumbent d-plan, completes the solution corresponding to each neighbor and evaluates its total cost, finds the lowest-cost d-plan within this neighborhood, updates the incumbent with this d-plan, and repeats these steps until no improved solution is found for that type of neighborhood. We perform this procedure for all the neighborhood types, considered in a specified sequence.

Each neighborhood type is specified by a combination of the following four attributes:

1. Type of operation: Add, drop, or swap deadheads
2. Number of crew members involved in the operation (one, two, or three in our approach)
3. Whether these deadheads are on the same trip or can be on different trips when multiple deadheads are added, dropped, or swapped
4. Deadhead direction(s) to consider (home-to-away, away-to-home, or both)

Because the number of combinations of these attribute types, and, hence, possible neighborhood types, is large, we selected a subset of neighborhood types that are effective in quickly finding a near-optimal solution. To identify this subset, we conducted a series of computational tests and chose the neighborhood types (and their sequence) based on the results. Starting with the simplest neighborhood types (i.e., add, drop, or swap one deadhead), we examined the heuristic d-plans obtained using these neighborhoods for many real-world test cases, compared them with the d-plans obtained by the optimization model, and added neighborhood types that improved the heuristic d-plans. Using this strategy, we gradually added new types of neighborhoods and dropped neighborhood types that were no longer effective when new ones were added. We also experimented with various sequences for searching the chosen neighborhood types to identify an effective sequence that quickly yields good solutions.

Train-Schedule Uncertainty

Our heuristic approach accounts for random variations in train departure times at the away station. (Typically, trains are ready to depart later than scheduled, not earlier.) We focused on modeling uncertainty in departure times from the away station because the inbound trips that connect to these trains are dispatched early in the planning horizon, and hence have accurate schedules. Moreover, at the away station, train departure delays increase the heldaway expenses for the assigned crew members, whereas heldaway costs are not incurred for crew members departing from home. For each scheduled outbound train at the away station, we consider a train-specific probability distribution for the difference between this train's actual and scheduled departure times. When this train is delayed, the connecting crew members have higher heldaway times. But, we may also encounter situations in which no rested crew member is available when the train is ready to depart. To accommodate this possibility, we permit violating the minimum rest requirement, but with a penalty per hour of shortfall in rest time (relative to the minimum requirement) for the connecting crew member. This penalty reflects the cost of postponing the outbound trip until crew members who arrived on the connecting inbound trips are adequately rested and ready to operate the train. This approach of penalizing violations of rest requirements essentially corresponds to treating the minimum rest requirement as a soft constraint. Based on the probability distribution of train departure times from the away station, our solution approach computes and adds the expected rest-time penalty to the total cost of a d-plan within the local improvement iterations.

Effectiveness of the Heuristic

We applied deterministic versions of both the exact and heuristic algorithms to a sample of 1,728 real-world data sets and compared the run times and solution costs for each problem instance. We found that the heuristic runs nearly seven times faster than the exact algorithm and has more consistent run times. Further, the heuristic found an optimal solution in all but 11 cases (99.2%) and found a solution within 2% of optimality in all but two cases (99.9%). We, therefore, decided that the heuristic was an adequate substitute for the exact algorithm, and focused subsequent development, including the extensions to address train schedule uncertainty, on the heuristic.

From Prototype to Production

Moving our prototype implementation to a production environment for use by planners required a lengthy period in which we validated the model and then integrated it into the users' existing production planning tools. We validated the model in stages, each with iterative improvements, based on user feedback. The first stage entailed examining the solutions to ensure that they did not violate any rules and were implementable in practice, and then reviewing sample solutions with crew-planning experts. After addressing any suggestions and obtaining the approval of the experts, we repeated the process of review and approval with the crew operations managers, and then with the end users of the system (i.e., crew planners) in several districts. This process of reviews and rework culminated in a system that crew planners judged to be suitable for operational use.

In the new CDA system, our decision support approach was integrated with the existing (manual) tool used by BNSF to build deadhead plans. The system modifications ensured that the user's existing business process largely remains unchanged. The CDA system contains a *Parameters* button on the main screen that displays a list of settings (e.g., cost parameters, rotation rules, connection types) applicable to the district. Although the system provides significant flexibility by defining and providing control over many settings, most of the settings do not need to change from run to run; they can be fixed by a manager once and reviewed periodically. So we classified the users into groups based on their permissions to modify various settings. The system permits end users to change only a limited number of settings, whereas the settings their managers can change are less restrictive. To initiate a run of the optimization procedure, users click a *Decision Assist* button on the main screen. The system then gathers all necessary input data, including the upcoming train schedule and current status of all crew members, and passes this information, together with an initial plan (e.g., user-generated or preexisting plan), which we call the presolve plan, to an external server. The server executes the heuristic, generates a list of recommendations to transform the presolve plan into the optimized plan, and passes the recommendations back to the main tool. These recommendations fall into three categories:

- Add: Add a deadhead to the presolve plan
- Remove: Remove a deadhead from the presolve plan

• Modify: Change the departure time and (or) mode of a deadhead in the presolve plan

The user reviews the recommendations by clicking a *Results* button and selects the changes to apply. The system then modifies the presolve plan to incorporate these choices. After reviewing the revised plan, the crew planner commits it (i.e., communicates it to other planning and operations personnel). Because the CDA system required only minor changes to the existing planning tool and each run typically requires fewer than five seconds, the user business process was minimally disrupted.

CDA was rolled out to the user community in stages. We first introduced CDA to managers and crew planners in a few representative districts. Over a period of several months, planners in these districts were instructed to use the system and provide feedback to the development team through an internal bulletin board. The development team incorporated this feedback to improve the system, and then rolled it out to all the crew districts. Initially, planners were encouraged to only experiment with the system. After including further system refinements based on their suggestions, BNSF issued a mandate requiring all crew planners to use the system for decision-making.

Gaining Acceptance

To gain acceptance of CDA by the planner community, we first measured the extent to which planners accepted and used the system, and then employed strategies to increase its acceptance. CDA usage is continuously monitored to identify situations in which the system is not being used (possibly for legitimate reasons). After each CDA run and subsequent commit, the system determines whether the CDA recommendations were followed and classifies the outcomes as follows:

• Full commit: All recommendations were committed
• Partial commit: A subset of recommendations was committed
• Zero commit: No recommendation was committed
• Good standing: No change was required to the presolve plan (i.e., no recommendation was made)

Also, if a deadhead plan is committed without running CDA, the system determines whether CDA was run in the previous 15 minutes. If not, the decision is classified as an unassisted commit. The system tracks each planner's actions by measuring the percentages of commits that fall in each of these categories. The goal is to have a large percentage of commits be classified as either good standing or full commit. Partial commits are considered a partial success and zero commits are failures. Partial and zero commits indicate possible shortcomings of the model. Unassisted commits indicate lack of compliance with the usage mandate, prompting reminders and (or) additional training. These percentages are tracked at various geographic levels (i.e., region, division, district) and by user, facilitating diagnosis of problem areas and users. Finally, changes to CDA settings in each district are recorded to monitor and control the changes because these settings can significantly affect the performance of CDA's optimization module.

CDA usage was initially very low after the system rollout; however, it improved as the user community gradually accepted the system. Figure 2 displays the trends in the proportion of weekly commits in three categories during the ramp-up of CDA and after full

Figure 2. (Color online) Usage of the CDA System Reached 80% Within 1 Year After Deployment

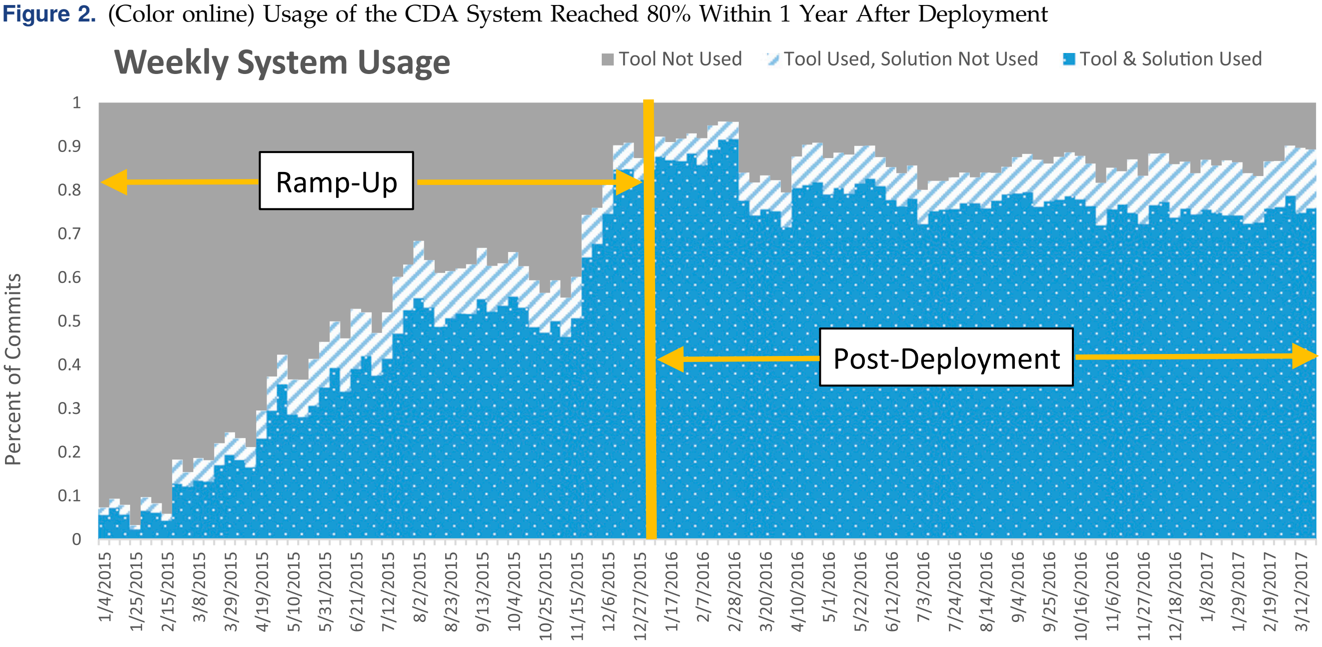

deployment: "Tool Not Used" (unassisted commits), "Tool Used, Solution Not Used" (zero commits), and "Tool & Solution Used" (good standing + full commit + partial commit). As the figure shows, CDA usage (dotted and hatched areas) has stabilized at around 90%, and the percentage of commits in which at least some of the recommendations were used has stabilized at around 75%. Increasing these numbers is a continual process of using feedback to further improve he system and working with crew planners to increase usage.

For many situations in which the CDA solution was not used fully, we found that the settings defined for that district were not appropriate. Maintaining the integrity of these settings is important; so the system generates periodic reports on the changes to the settings. Occasionally, the model needs to be enhanced to address a new or revised requirement; thus, maintenance of the model is also a continual process.

Low-usage patterns by individual users are not easy to address. By comparing the proportion of times that a user does not fully commit the recommended solution with the proportion for other users in that user's peer group, we can identify users who are outliers. Such users tend to be either relatively novice users or, at the other extreme, very experienced users who feel that they do not need model assistance. Usage among novice users can usually be increased with more training. For expert users, providing more detailed explanations of the model's logic and demonstrating its strengths can increase usage. We also identified champions—users who understand CDA very well and support its use—and made them responsible for training their peers. Our objective was to increase usage. This strategy was effective because these champions were better able to relate to their peers and customize their training, rather than having the development team conduct the training.

Impact

To assess the effectiveness and benefits of the CDA-optimized solutions compared with manual solutions, we estimate the aggregate impact of using these two methods over a given period. The two main factors that influence crew costs are the number of deadhead movements and the total heldaway hours (for all crew members) in each crew district. These two values depend on the traffic intensity (i.e., the number of trains traversing the district in each travel direction), which can vary from week to week. Therefore, to assess the benefits of using optimized versus manual plans in a district, we use two nonfinancial metrics, deadhead rate and heldaway rate, defined, respectively, as the total number of deadheads and total number of heldaway hours per period (e.g., week) in that district divided by the number of train starts in that period. The lower the values of these metrics, the lower the cost to BNSF. Since deploying CDA, we have been tracking these two metrics for every district on an ongoing basis.

We cannot directly compare the deadhead and heldaway rates after CDA deployment with the rates observed before deployment because they depend on the characteristics of the train schedules, particularly the number of trains and the directional imbalance (absolute difference in train movements in the two directions) of trains traversing a district. For example, more trains lead to a greater number of opportunities for crew connections, thus reducing crew shortages and heldaways. Larger directional imbalances (i.e., more trains traversing the district in one direction than the other) can necessitate deadheads to equalize the imbalance. At the time CDA was deployed, BNSF experienced structural changes in its traffic patterns due to a changing mix of traffic and various economic factors, causing a notable drop in train flows across many districts.

In an initial comparison of actual pre-CDA and post-CDA deadhead and heldaway rates, we found these rates to be roughly comparable, although post-CDA traffic was lower (implying that the post-CDA rates should have been higher). This observation suggested that the CDA-supported decisions improved on previous manual decisions. To obtain a better assessment of the benefits of CDA, we decided to estimate the deadhead and heldaway rates that each district would have observed whether manual crew plans were applied to the actual traffic flows after CDA was deployed. For this purpose, we developed simple linear regression models for each district, using pre-CDA data, to estimate the two metrics as functions of the number of trains and the magnitude of the directional imbalance in the district. As we noted above, we expect deadhead and heldaway rates to correlate negatively with the number of trains but correlate positively with the magnitude of the directional imbalance.

Figure 3 shows the data points corresponding to the weekly deadhead and heldaway rates versus the number of trains, immediately prior to CDA deployment, for a sample district. Figure 4 shows the data points corresponding to the two rates versus the directional imbalance for the same district. As anticipated, these figures show that deadhead and heldaway rates decrease with number of trains and increase with directional imbalance.

We built such regression models for each district and used them to predict the deadhead and heldaway rates that BNSF would have observed in each district had it used manual planning in the year following CDA deployment. We then compared these predicted values with the values we observed using CDA over the same period. Using the appropriate costs per deadhead and per heldaway hour, we estimated that using CDA yields savings to BNSF of several million dollars per

Figure 3. (Color online) Deadhead and Heldaway Rates Depend on the Number of Trains Traveling Through a District

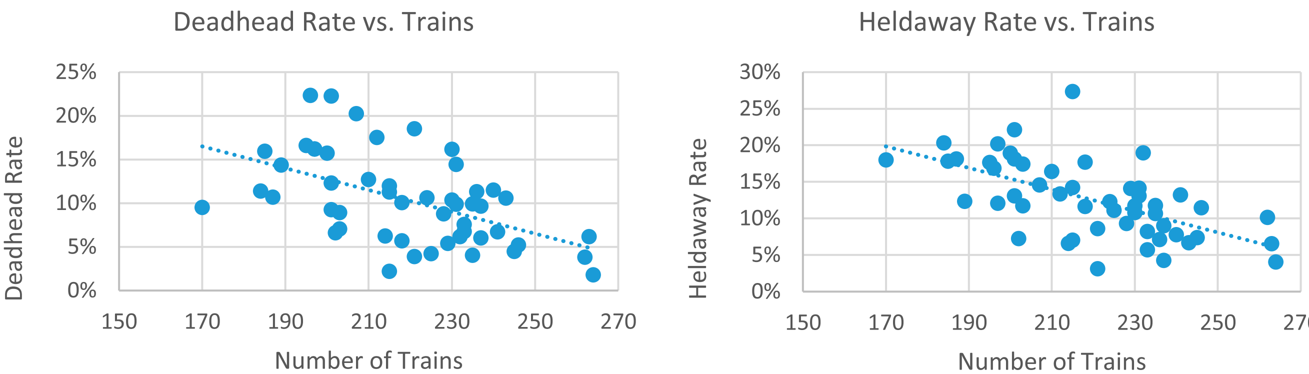

year in crew costs. Senior leadership at BNSF has recognized the impact of CDA:

- The Executive Vice President of Operations stated:

> Given the size and complex nature of our network, an application such as Crew Decision Assist enhances our planners' abilities to deal with the scale and dynamic nature of our operations. In the past two-and-half years, CDA has been adopted by the majority of our planning team and it has been instrumental in lowering costs by several million dollars each year.

- A General Director of Operations stated:

> Crew Decision Assist has provided a technology solution that has changed the way we schedule crews who operate trains. CDA has positively impacted our cost efficiency and the life quality of employees who operate trains by reducing their time away from home. As an example, we recently made the decision to suspend the use of CDA after growth in our volumes resulted in operational challenges. When we did, cost and employee morale deteriorated. When CDA was again initiated, we immediately realized improvements in both.

Conclusions

In this paper, we described the development, implementation, and usage of an optimization model and method for crew planning at BNSF Railway. Our model incorporates features such as different crew-rotation rules (other than FIFO) and usage of extra crews that previous crew-planning models have not addressed. Our solution approach is also novel because it accounts for uncertainty in train schedules. The crew-scheduling algorithm, used as a decision support engine within the existing crew-planning system, has been well received by crew planners, and steady adoption of the system has significantly reduced crew costs. Encouraged by this experience, we next plan to develop optimization-based methods to support crew scheduling in districts that are not single ended.

Our experience with the development, implementation, and adoption of CDA also highlights principles for the successful application of optimization-based decision support in practice. Below, we summarize some pertinent takeaways from this project.

- Formulating a decision problem as an optimization model is valuable to improve and validate our understanding of the problem, and optimally solving it yields benchmarks against which we can compare the quality of solutions using heuristics that may be more viable for practical use.
- Modeling and algorithmic development is an iterative process, which consists of applying the approach

Figure 4. (Color online) Deadhead and Heldaway Rates Depend on the Imbalance in Traffic Between the Two Travel Directions in a District

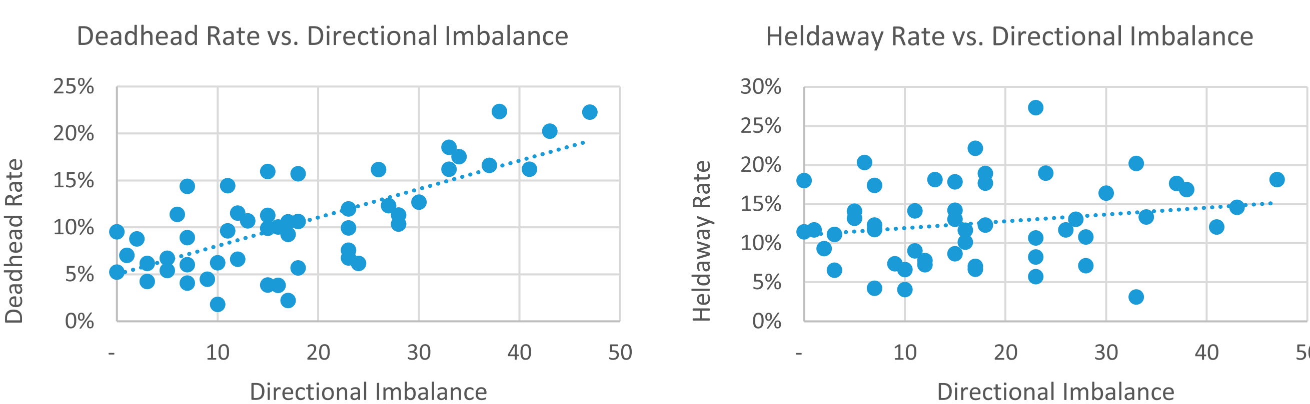

to real data, engaging the users to validate the results, and improving and refining the approach based on feedback.

- In the crew-scheduling context, crew districts vary in their operational policies, requirements and restrictions, and decision options. Moreover, these features change over time, for example, because of new union agreements that change work and rest rules. So developing systems that are flexible and can readily accommodate these variations is important.
- For successful implementation, development teams must anticipate and institute strategies to overcome resistance to change among users by (1) introducing the system in phases, starting with carefully chosen lead users, (2) demonstrating the economic and operational benefits of employing the system, (3) explaining broader applications of the system (e.g., for scenario analysis), and (4) making system improvements in response to user feedback.

These observations reinforce the experience of successful practice-driven operations researchers.

Appendix. Integer Programming Formulation

Sets and Indices

K: Set of occupations, $k \in K$

A: Set of all trips, $i, j \in A$, with j typically denoting the index of an outbound trip and i denoting any trip

A^I: Set of *initial* trips that departed before the planning horizon, each carrying crew members who arrived before or are in transit at the beginning of the planning horizon

$A^{IA} \subseteq A^I$: Subset of *initial* trips arriving at the *away* station

A^F: Set of *future* trips that depart during the planning horizon, including an artificial "sink" trip from each station at the end of the planning horizon, $A^F = A \backslash A^I$

A^T: Set of future *train* trips, $A^T \subset A^F$

A^D: Set of future *deadhead* trips (taxis and public transport), $A^D = A^F \backslash A^T$

A^H: Set of future trips departing the *home* station, $A^H \subset A^F$

P_{jk}: Set of valid *predecessor* trips to trip $j \in A^F$ for occupation $k \in K$

S_{ik}: Set of valid *successor* trips from trip $i \in A$ for occupation $k \in K$

B_i: If the rotation rule at the arrival station of trip $i \in A$ is FIFO, then B_i is the set of trips that arrive at the same station as trip i and arrive before trip i, or if the rotation rule at the arrival station of trip $i \in A$ is FOFO, then B_i is the set of trips that arrive at the same station as trip i and depart before trip i

Parameters

u_{ik}: Number of *regular* crew members from occupation $k \in K$ available on initial trip $i \in A^I$

v_{ik}: Number of *extra* crew members from occupation $k \in K$ available on initial trip $i \in A^{IA}$

b_i: *Capacity* of trip $i \in A^F$ (unlimited for deadhead trips on public transport)

$\underline{m}_{ik}$: *Minimum* number of crew members of occupation $k \in K$ needed on trip $i \in A^T$

$\overline{m}_{ik}$: *Maximum* number of crew members from occupation $k \in K$ who can be assigned to trip $i \in A^F, \overline{m}_{ik} = b_i - \sum_{k' \in K: k' \neq k} \underline{m}_{ik'}$

r_k: *Total* number of regular crew members from occupation k available in the district

t_i: *Departure time* of trip $i \in A$

c_{ijk}: *Connection* cost for a crew member of occupation $k \in K$ to connect from trip $i \in A$ to trip $j \in S_{ik}$ for occupation $k \in K$

d_{ik}: *Crew payment* made to each crew member of occupation $k \in K$ assigned to trip $i \in A^F$

e: Cost for each *extra* crew member used from the home station

f_i: *Fixed* cost of deadhead trip $i \in A^D$ (zero for deadhead trips on public transport)

Decision Variables

X_{ijk} = 1 if a *regular* crew member of occupation $k \in K$ connects from trip $i \in A$ to trip $j \in S_{ik}$ and 0 otherwise

Y_{ijk} = 1 if an *extra* crew member of occupation $k \in K$ connects from trip $i \in A^{IA}$ to trip $j \in S_{ik}$ and 0 otherwise

U_{ik} = Number of *regular* crew members of occupation $k \in K$ assigned to trip $i \in A^F$

V_{ik} = Number of *extra* crew members of occupation $k \in K$ assigned to trip $i \in A^F$

Z_i = 1 if at least one crew member is assigned to trip $i \in A^D$ and 0 otherwise

W_{ik} = 1 if at least one extra crew member of occupation $k \in K$ is assigned to trip $i \in A_h$ and 0 otherwise

Objective Function

Minimize

$$\sum_{\substack{i \in A \\ j \in S_{ik} \\ k \in K}} c_{ijk}(X_{ijk} + Y_{ijk}) + \sum_{\substack{i \in A^F \\ k \in K}} d_{ik} U_{ik} + \sum_{\substack{i \in A^H \\ k \in K}} (d_{ik} + e) V_{ik} + \sum_{i \in A^D} f_i Z_i \tag{A.1}$$

Constraints

Trip-Transfer Linking.

$$\sum_{j \in S_{ik}} X_{ijk} = u_{ik} \qquad i \in A^I,\, k \in K \tag{A.2}$$

$$\sum_{j \in S_{ik}} Y_{ijk} = v_{ik} \qquad i \in A^{IA},\, k \in K \tag{A.3}$$

$$\sum_{j \in S_{ik}} X_{ijk} = U_{ik} \qquad i \in A^F,\, k \in K \tag{A.4}$$

$$\sum_{j \in S_{ik}} Y_{ijk} = V_{ik} \qquad i \in A^F,\, k \in K \tag{A.5}$$

$$\sum_{i \in P_{jk}} X_{ijk} = U_{jk} \qquad j \in A^F,\, k \in K \tag{A.6}$$

$$\sum_{i \in P_{jk}} Y_{ijk} = V_{jk} \qquad j \in A^F,\, k \in K \tag{A.7}$$

Train-Crew Requirement.

$$U_{ik} + V_{ik} \geq \underline{m}_{ik} \qquad i \in A^T,\, k \in K \tag{A.8}$$

Train Capacity.

$$\sum_{k\in K}(U_{ik}+V_{ik})\le b_i \qquad i\in A^T \qquad \text{(A.9)}$$

Deadhead-Selection Forcing and Capacity.

$$\sum_{k\in K}(U_{ik}+V_{ik})\le b_i Z_i \qquad i\in A^D \qquad \text{(A.10)}$$

Occupation Pairing.

$$U_{ik}+V_{ik}-\underline{m}_{ik}=U_{ik'}+V_{ik'}-\underline{m}_{ik'} \qquad i\in A^F, k\in K, k'\in K, k\ne k' \qquad \text{(A.11)}$$

Contingent Extra-Crew Usage.

$$V_{jk}\le \overline{m}_{jk}W_{jk} \qquad j\in A^H,\ k\in K \qquad \text{(A.12)}$$

$$\sum_{i\in A^I\cap P_{jk}} u_{ik}+\sum_{i\in A^F\cap P_{jk}} U_{ik}-\sum_{j'\in A^H:t_{j'}\le t_j} U_{j'k}\le r_k(1-W_{jk}) \qquad j\in A^H, k\in K \qquad \text{(A.13)}$$

Rotation Rule.

$$X_{ijk}+Y_{ijk}+X_{i'j'k}+Y_{i'j'k}\le 1 \qquad k\in K, i\in A, i'\in B_i, j\in S_{ik}\cap S_{i'k}, j'\in S_{i'k}, t_{j'}>t_j \qquad \text{(A.14)}$$

Nonnegativity and Integrality.

$$X_{ijk}, Y_{ijk}\in\{0,1\} \qquad i\in A, j\in S_{ik}, k\in K$$
$$U_{ik}, V_{ik}\ge 0 \text{ and integer} \qquad i\in A^F, k\in K$$
$$Z_i\in\{0,1\} \qquad i\in A^D$$
$$W_{ik}\in\{0,1\} \qquad i\in A^H, k\in K$$

This model has the following structure:

Objective Function

The objective function is to minimize the total crew deployment cost during the planning horizon, including crew payments, and costs for deadheading, layovers, and usage of extra crews.

Constraints

- *Trip-transfer linking*: Constraints (A.2)–(A.7) relate the crew-to-trip assignment and connection variables. Constraints (A.2) and (A.3) ensure that each crew member arriving at the system on an initial trip is connected to a successor trip; Constraints (A.4) and (A.5) ensure the same for crew members arriving on future trips; and Constraints (A.6) and (A.7) ensure that each crew member departing on a future trip has a connecting predecessor trip. These constraints serve to conserve flow: the flow on a trip must equal the flow connecting to the trip and equal the flow connecting from the trip.
- *Train crew requirement*: Constraint (A.8) ensures that each train has the required complement of (regular or extra) crew members to operate the train.
- *Train capacity*: Constraint (A.9) ensures the number of deadheading crew members assigned to each train does not exceed the train's capacity.
- *Deadhead-selection forcing and capacity*: Constraint (A.10) ensures that the binary taxi-selection variable assumes a value of one if one or more crew members are assigned to a taxi trip, and to impose capacity limits on the number of crew members traveling on the taxi.
- *Occupation pairing*: Constraint (A.11) ensures that deadheads contain crew members of both occupations. The number of crew members assigned minus the number required is the number deadheaded, which must be the same for each occupation.
- *Contingent extra crew usage*: Constraints (A.12) and (A.13) ensure that extra crew members are used only when enough rested *regular* crew members are not available to travel on an outbound trip from the home station. Specifically, if one or more *extra* crew members of occupation k are assigned to outbound trip j, Constraint (A.12) forces $W_{jk}=1$; then Constraint (A.13) forces the cumulative number of *regular* crew members who are rested at the home station before trip j departs to be less than or equal to the cumulative number of *regular* crew members dispatched from the home station on or before trip j (which implies that enough regular crew members are not available for trip j). Note that the trip-transfer linking Constraints (A.2), (A.4), and (A.6) ensure that the left side of Constraint (A.13) is nonnegative. Therefore, when $W_{jk}=1$, Constraint (A.13) effectively requires the left side to be equal to zero (i.e., the cumulative number of regular crew members of occupation k who are rested and available before trip j departs must equal the cumulative number dispatched on or before trip j). When $W_{jk}=0$, the constraint is "deactivated."
- *Rotation rule*: Constraint (A.14) ensures that the chosen connections meet the crew-dispatching priorities applicable to the district and occupation. For example, if the rotation rule is FIFO, then connections <i, j> and <i', j'> are "incompatible" if trip i' arrives before trip i and trip j' departs after trip j (because crew members on trips arriving earlier should depart earlier than crew members on trips arriving later), provided trip j is a valid successor to trip i'. Constraint (A.14) prevents incompatible connection pairs from being selected in a solution.

The long computational time to exactly solve the integer program stems primarily from the large number of rotation-rule Constraint (A.14); a typical problem instance has several million of these constraints, but only a few thousand of the remaining constraints. The heuristic easily handles these constraints when it "simulates" the connections made for a given deadhead plan, which is a major contributor to its success.

References

Abbink E, Fischetti M, Kroon L, Timmer G, Vromans M (2005) Reinventing crew scheduling at Netherlands Railways. *Interfaces* 35(5):393–401.

American Association of Railroads (2017) The economic impact of America's freight railroads. Accessed August 13, 2018, https://www.aar.org/issue/freight-rail-economic-impact/.

Balakrishnan A, Kuo A, Si X (2016) Real-time decision support for crew assignment in double-ended districts for U.S. freight railways. *Transportation Sci.* 50(4):1337–1359.

Barnhart C, Cohn AM, Johnson EL, Klabjan D, Nemhauser GL, Vance PH (2003) Airline crew scheduling. Hall RW, ed. *Handbook of Transportation Science*, 2nd ed. (Kluwer Academic Publishers, Norwell, MA), 517–560.

Caprara A, Monaci M, Toth P (2001) A global method for crew planning in railway applications. Voß S, Daduna JR, eds. *Computer-Aided Scheduling of Public Transport*, Lecture Notes Econ. Math. Systems, Vol. 505 (Springer-Verlag, Berlin), 17–36.

Caprara A, Toth P, Vigo D, Fischetti M (1998) Modeling and solving the crew rostering problem. *Oper. Res.* 46(6):820–830.

Caprara A, Kroon L, Monaci M, Peeters M, Toth P (2007) Passenger railway optimization. Barnhart C, Laporte G, eds. *Handbooks in Operation Research and Management Science*, Vol. 14 (North-Holland, Amsterdam), 129–187.

Gopalakrishnan B, Johnson EL (2005) Airline crew scheduling: State-of-the-art. *Ann. Oper. Res.* 140(1):305–337.

Federal Railroad Administration (2017) Freight rail overview. Accessed February 1, 2018, https://www.fra.dot.gov/Page/P0362.

Gorman MF, Sarrafzadeh M (2000) An application of dynamic programming to crew balancing at Burlington Northern Santa Fe Railway. *Internat. J. Services Tech. Management* 1(2–3):174–187.

Jütte S, Albers M, Thonemann UW, Haase K (2011) Optimizing railway crew scheduling at DB Schenker. *Interfaces* 41(2):109–122.

Sahin G, Yüceoglu B (2011) Tactical crew planning in railways. *Transportation Res. Part E* 47(6):1221–1243.

Vaidyanathan B, Jha KC, Ahuja RK (2007) Multicommodity network flow approach to the railroad crew-scheduling problem. *IBM J. Res. Dev.* 51(3):325–344.

Brian Roth is a manager in the Operations Research and Advanced Analytics Group at BNSF. He has over 30 years of experience in the application of mathematical techniques to solve problems in the railway, telecommunications, and defense industries. He received a BS in Applied Mathematics from Carnegie Mellon University, received an MS in Operations Research from the University of Alabama in Huntsville, and completed work on a DSc in Operations Research at George Washington University (ABD).

Anant Balakrishnan holds the Jastrow Chair in Business at the McCombs School of Business, University of Texas at Austin. He has previously held faculty positions at Purdue, MIT, and Penn State. His main research interests lie in developing optimization models and algorithms to support decision-making in supply chain management, transportation, telecommunications, and service operations. He has published articles in the leading operations research and management science journals.

Pooja Dewan is the chief data scientist at BNSF Railway and has led the Operations Research and Advanced Analytics group for the past 14 years. She has also worked as an internal strategy consultant and led Locomotive Distribution and Planning. She received her MS and PhD from Pennsylvania State University and has authored several scientific research publications. She has been a member of INFORMS since 1993 and has served the INFORMS organization in various leadership positions.

April Kuo is a manager of Operations Research at BNSF Railway, where she is supporting service design, travel time prediction, and detector analytics. She received her PhD from University of Maryland, College Park. She has authored several papers for journals such as *Transportation Science*, *Transportation Research Part B*, and *IEEE*. She serves on TRB Transportation Network Modeling Committee and has been actively engaged in the INFORMS Railway Application Section.

Dasaradh Mallampati is a manager in Operations Research and Advanced Analytics at BNSF Railway. He currently supports crew planning and scheduling initiatives. He received his MBA from Northwestern University and his MS in Computer Sciences from University of Tulsa, and did doctoral work in Operations Research at SMU. He has published scientific papers on self-adaptive fuzzy logic control and distributed computing. He worked in airline, residential mortgage, and supply chain areas prior to joining BNSF.

Juan Morales is a director in the Operations Research & Advanced Analytics group at BNSF Railway. He has over 12 years of experience in the railroad industry. He received his PhD from the Georgia Institute of Technology and has published papers in journals such as *Operations Research* and *Transportation Science*. He was awarded the 2007 first prize in the Dissertation Award competition held by the Transportation Science and Logistics Society of INFORMS. He has been an INFORMS member since 2001.