Priority-Based Classification for Improving Connection Reliability in Railroad Yards

Part II of II: Dynamic Block to Track Assignment

The ability to guarantee connections of particular cars to specific outbound trains is a prerequisite to the effective implementation of freight railroad revenue management in the future. This ability would allow railroads to offer reliable "time definite" service with a guaranteed delivery time for each car. This paper describes the block to track assignment component of the Terminal Priority Movement Planner (TPMP), a proof-of-concept decision support system field tested in 1994 at Union Pacific Railroad's Hinkle, Oregon, classification yard. A new method of sorting freight cars by both outbound train and destination yard block, rather than by destination yard block only, is proposed to allow railways to gain precision control over the makeup of each train. Cars can be scheduled through yards based on their delivery commitments, rather than on a first-in-first-out basis, so cars having no remaining schedule slack have first access to available train capacity.

by Edwin R. Kraft

This paper proposes a new method of sorting railcars in classification yards, where the goal is to ensure connections of particular cars to specific outbound trains. By doing this, service reliability can be improved. Cars can be scheduled through yards based on their delivery commitments rather than on a first-in-first-out basis, so cars having no remaining schedule slack time have first access to available train capacity. Although revenue-management based approaches can suggest which cars should be taken on each train (Kraft, 1998; 2000a; 2001; and Kwon, Martland and Sussman, 1998) railroad terminals have never been capable of executing such detailed connection plans. In rail yards today, priority service is provided by exception, outside the normal process workflow, using very expensive methods, such as special switch engine moves for just one or two cars, or extracting specific cars for trains at the "trim" end of the yard, known as "cherry picking." Although some have suggested that yards ought to be physically redesigned to facilitate more "hot car" and "cherry picking" moves, certainly the railroad industry's goal should not be to increase those very costly operations, but rather to reduce the need for them, as the methods are proposed here.

Instead, this paper will show how (with some preplanning) all classification can be accomplished at the hump, which is designed for sorting individual cars, rather than inefficiently selecting cars at the "trim" end of the yard. The process described here allows classification of cars not only by destination, but also by specific train. Cars are thus clas-

Journal of the Transportation Research Forum. Published in *Transportation Quarterly*,
Vol. 56, No. 1, Winter 2002 (107–119). © 2002 Eno Transportation Foundation, Inc., Washington, D.C.

sified more precisely than commonly done today. Classification track space is intensively used to minimize the number of rehump cars generated. This paper focuses on management of classification tracks and rehump activities for classifying cars both by destination block and by outbound train. A specific operating method is proposed to allow yards to gain precision control over the makeup of each train.

A prototype decision support system, the Terminal Priority Movement Planner (TPMP), was developed by the author and field tested in 1994 at Union Pacific Railroad's Hinkle, Oregon, yard. The hump sequencing functionality of TPMP was described in a previous paper (Kraft, 2000b) and further required car scheduling adjustments were discussed in Part I of this paper.

PROBLEM DEFINITION AND SOLUTION APPROACH

Part I of this paper shows how a "feasible connection plan" respecting train capacity limits can be developed by rescheduling low priority cars in excess of capacity to another train. As shown in Figure 7 of Part II, additional tracks may be temporarily needed to separate cars by outbound trains. However, there is no guarantee such track assignments can always be found. If the tracks are not available, low priority cars in excess of train capacity can always be sent to a rehump track. However, the goal of the TPMP pilot at Hinkle was to demonstrate how priority-based classification could be implemented without needing to rehump so many cars.

To accommodate large-scale sorting by train and block, a systematic method of planning block to track assignments is needed. Cars must fit within available track space at all times. To minimize the need for rehumping cars, tracks must be available when needed to start new blocks. Switching efficiency must also be maintained at the trim end of the yard. The problem was considered too complex for manual solution. Automating the search for a pattern of block to track assignments was considered essential to the success of the 1994 Hinkle pilot.

Kraft and Spielberg (1993) proposed a unified mathematical framework for sorting by train block, solving simultaneously for hump sequencing and block to track assignment, but the model was computationally impractical and only capable of solving "toy" problems. For this work for Union Pacific the following year, a solution had to be obtained to practical problems, so the hump sequencing and block to track assignment problems were decoupled and solved sequentially using a heuristic approach. With this approach, after car scheduling trip plans are adjusted for connections and train capacity constraints, a Block to Track Assignment process develops a set of track assignments to implement the plan. To reduce the number of rehump cars generated, the system tries to find track assignments to start new train blocks as soon as possible. Assignments are developed using a heuristic approach, although an optimization approach may do even better.

It is convenient to represent time in discrete "processing intervals," where each processing interval corresponds to the length of time required to process one train or to rehump one classification track. Block to track assignments are effective for the entire processing interval, not depending on the sequence of individual cars within each train. Some requirements for block to track assignment include:

(1) Each train block must be assigned a track when cars first appear at the hump, or else those cars must be sent to a rehump track for reprocessing later.

(2) Capacity of any track must never be exceeded.

(3) Cars require continuous track occupancy from the time they are switched into a

track until they are pulled to form an outbound train, or if cars are sent into a rehump track, until those cars are pulled out of the track for reprocessing at the hump.

(4) Except for rehump tracks, each track may have only one train block assigned at a time; however, after the last car has been processed for a train block, another block may be started in any remaining track space behind it. (In the current implementation, new blocks are not started until the beginning of the next processing interval. This avoids dependency of the solution on the sequence of cars within each train.) Each train block becomes "active" when a track is dedicated to it, and remains active as long as additional cars remain to be processed for the train block and the track still has capacity to receive those cars.

(5) Each train block must receive a dedicated track assignment for at least one processing interval prior to that train's close out time.

(6) If a new train block is started behind a closed out block, it should have a later close out time scheduled than the block ahead of it. This allows each block to be pulled from the trim end of the yard in proper sequence without extra switching. The number of cars scheduled to arrive in the new block may not overflow track space before other blocks ahead on the same track are projected to be pulled out of the yard.

(7) Cars may not miss connections as a result of being assigned to a rehump track. If no rehump processing events are scheduled before train departure, the train block must receive an immediate track assignment. Any train block assigned to a rehump track must be scheduled for reprocessing prior to the close out time of its outbound train.

(8) Special assignment rules are needed to force or "coerce" continued use of tracks already in use for active train blocks to prevent those train blocks from splitting across multiple tracks.

Writing all these constraints in a traditional mathematical programming framework leads to an excessively complex and intractable mixed integer formulation. That will not be attempted here; see Kraft and Spielberg (1993) or Wang (1998) for an example of what such a formulation looks like. In practice, block to track assignments can be found through iterative application of a set of heuristic rules: the approach to block to track assignment proposed here is rule based rather than optimization based.

Management of Rehump Activities

Periodically, cars assigned to rehump tracks must be pulled back and reprocessed over the hump. As described in the previous paper on hump sequencing (Kraft, 2000b), time slots for rehump activities are built into the hump sequence at regular intervals (normally, once every eight hours). Practically, this implies that three separate rehump classifications may be under construction at the same time. Since rehump times are scheduled in advance, rehump cars are assigned to the latest time slot which can still make the planned outbound train connection. If no rehump events are scheduled before the planned trim-out time of a train, the train block must receive an immediate track assignment, and is excluded from consideration as a rehump candidate.

Based on the feasible connection plan developed in Part I of this paper, each train block has a precise number of cars, length, required duration of track occupancy, and since the hump sequence has already been established, a profile describing accumulation of railcars as each inbound train is processed. Track space requirements for each train block can be graphically represented using a well-known analysis tool, the "connection matrix." A connection matrix shows inbound trains along the left or vertical axis; while departing trains are shown across the top. The number of cars connecting from each inbound to each outbound train is shown within each cell of the matrix.

Table 1 shows the number of cars connecting from three inbound trains to three outbound trains. No cars are connecting from trains A, B or C to train Y, or from trains C or D to train Z. A more detailed version of this connection matrix shown in Table 2 includes block as well as train information. In Table 2, blocks are numbered 1 thru 4. The combination of train identifier plus yard block comprises a train block. While only a simple train-to-train matrix shown in Table 1 is needed for hump sequencing (Kraft, 2000b) the detailed matrix of Table 2 is needed to develop block to track assignments. In Table 2, train block "X-2" requires continuous track occupancy, even though train C has no cars for that block. The track requirement associated with each train block can be graphically determined by drawing boxes around and shading groups of cells starting from the first train, and continuing through the last train having cars for each train block, as in Figure 1.

Table 1: Connection Matrix by Outbound Train

	Outbound Trains	X	Y	Z
Inbound Trains	Hump Time	12:00 PM Trim Time	3:30 PM	10:30 AM
A	7:00 AM	2 cars		3 cars
B	8:30 AM	7 cars		6 cars
C	10:00 AM	3 cars		
D	11:00 AM	3 cars	6 cars	
REHUMP	11:45 AM			
E	1:00 PM		4 cars	
F	3:00 PM			

Table 2: Connection Matrix by Train Block

	Outbound Train-Block			
Inbound Trains	X-1	X-2	Y-3	Z-4
A	2 cars			3 cars
B	2 cars	5 cars		6 cars
C	3 cars			
D		3 cars	6 cars	
REHUMP				
E			4 cars	
F				

Cars might not be immediately removed or "trimmed out" from tracks when the last car is classified. Those cars might remain in the classification yard for awhile. For example: even though one train block is complete, other train blocks may still be awaiting additional cars. Even if all the train blocks are complete, crews or locomotives might not yet be ready. Any expected delay to the start of the "trim" operation can be depicted by extending a non-shaded box further downward, as in Figure 2.

During this waiting time, another train block can be started behind a closed-out train block, but cars continue to occupy track space until

Figure 1: Connection Matrix by Train Block

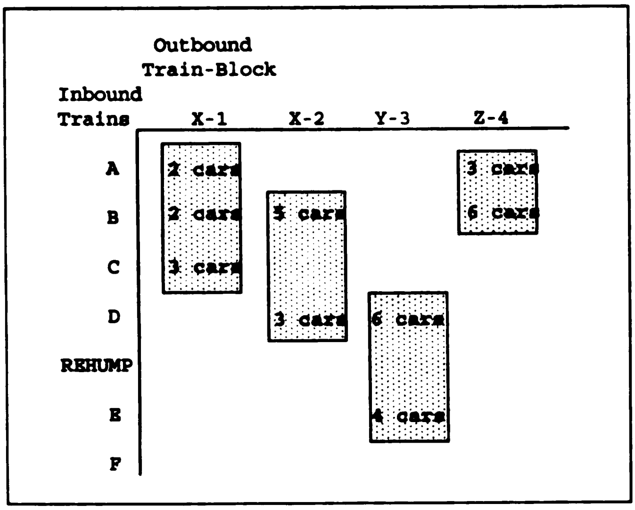

Figure 2: Connection Matrix by Train with Trim Extension

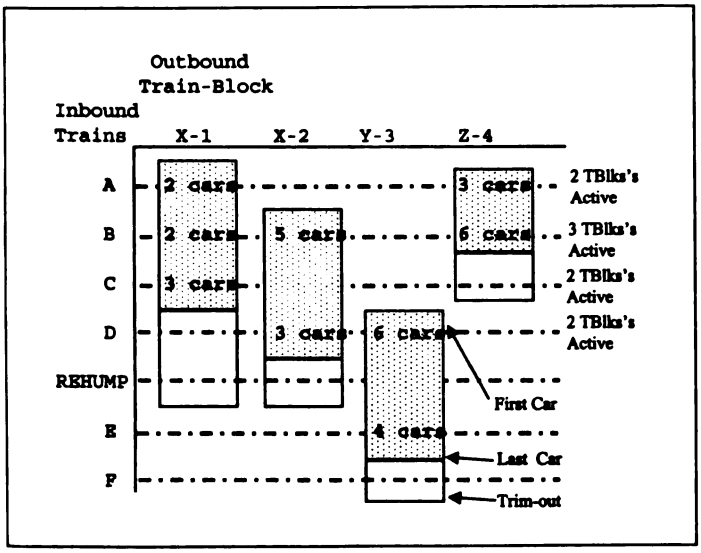

they are trimmed out. The end of the non-shaded box shows when cars are expected to be "trimmed out" or removed from classification tracks. Since we don't want the new train block to overflow into an additional track before the first group of cars can be trimmed out of the yard, this limits the size of the new train block that can be started.

Figure 2 shows a set of train blocks assigned to a hypothetical, two-track yard. By counting the number of shaded bars horizontally across the figure, it can be seen that three train blocks must remain active during processing of train B. But since only two tracks are available, this pattern of block to track assignment would be infeasible.

To reduce the number of tracks needed, some cars can be sent to a rehump track. In Figure 2, train block Z-4 is not an allowable rehumping candidate, since no rehump activities are scheduled between inbound train B and the scheduled trim out time for outbound train Z. (Sending those cars to a rehump track would cause them to miss their connection.) Block Y-3 will not be started until train D is processed. So, only the X-1 or X-2 blocks may be considered as rehumping candidates. The number of cars in each of these two train blocks is shown in Table 3. By sending cars from trains A and B for X-1 into the rehump track, one train block can be eliminated during processing of train B at a minimum cost of rehumping four cars.

After sending the first two connections for X-1 to the rehump track, the result is shown in Figure 3. Now the X-1 train block no longer requires track assignment until train C, but its "close out" is delayed until the rehump track is processed (at which time those cars diverted from trains A and B are put back into the block.) As a result of sending those cars to the rehump track, only two tracks (plus the rehump track, which is assumed to be always available) are required to process train B. However, anytime a block "close out" time is extended in this fashion, another iteration is likely to be triggered. As shown in Figure 3, train D now requires three tracks: yet another iteration is required. The best choice is to rehump cars from train blocks where the close out time doesn't have to be extended (although no such blocks were available in this example).

Fitting Blocks into Available Track Space

Once the number of train blocks has been reduced so it does not exceed the number of tracks in any time period, a trial assignment of those blocks into the tracks can be attempted. This compares the length requirement of each train block to the remaining track space available. If the assignment succeeds, a feasible solution is found. If some train blocks are too long and spill over into another track, the assignment may fail. Then the constraint on the maximum number of active tracks must be further tightened (during the processing interval where the trial assignment failed), requiring additional cars to be rehumped. Through this iterative process, a feasible solution will eventually be found under all but the most congested yard conditions.

Table 3: Car Distribution in Rehump Candidate Blocks

	X-1	X-2
A	2 cars	-
B	2 cars	5 cars
Total*	4 cars	5 cars

* Prior to and including train

Figure 3: Connection Matrix after X-1 Sent to Rehump

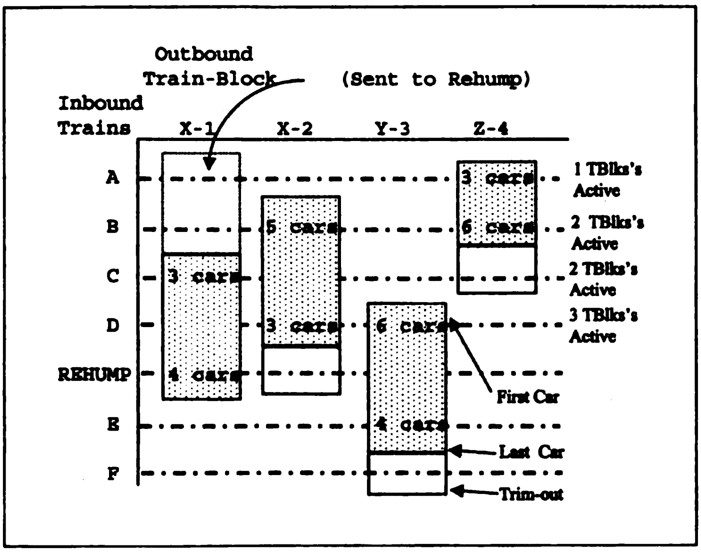

Since cars accumulate over time, it is not necessary for every track to have sufficient room to hold all the cars at first. During each processing interval, it is only necessary for tracks to have sufficient room to hold cars expected to have accumulated by that time. Even a large block can often be started in a small amount of space remaining behind a closed out block, provided the block ahead is scheduled for trimming out before the track overflows.

The block to track assignment process conservatively assumes removal of cars from the classification yard will not commence until the planned "Start Trim" time for each outbound train, even if all the cars would be available to trim earlier. In the Hinkle test, tracks were not assumed to be cleared until two hours after the planned trim-out time for the outbound train.

Train blocks are assigned to tracks in the following priority: first, train blocks already started in the yard must continue using those same tracks. Second, since large blocks too long to fit any track may have very few options, those assignments are processed next. Small blocks with many possible choices are assigned last.

Coerced Assignment

After a block has been assigned to a track and some cars have accumulated, the same track should continue to be used until either the block closes out or the track fills up. This is necessary to prevent cars for the same block being scattered across the yard, leading to difficulties in collecting all the cars from the trim end of the yard. Without special purpose subroutines to lock initial block to track assignments, the rule based heuristics (even with penalty costs built in) kept persistently trying to "improve" the solution.

Unfortunately, data on initial block to track assignments (stored in the yard process control computer) was not directly accessible to the TPMP program. Besides, that computer was only aware of yard block information—the process control computer had no knowledge of the train to which those blocks might be associated. Only data from Union Pacific's Transportation Control System could be made available. Programming logic had to be developed to infer the initial set of block-to-track assignments, based on an examination of the condition of the initial yard inventory.

For coercing initial track assignments, only "active" train blocks—those blocks both having cars in the classification yard and other cars still remaining to be processed—need to be assigned. If no cars remain to be processed, the train block is closed out and do not need to be assigned to any track. If no cars are in the yard yet, then any previous decisions are not binding, and any track could still be assigned.

Split Assignment

Very large train blocks that cannot fit any available track must be "split" into two or more pieces. Since a single block occupies more than one track, this directly reduces the number of blocks that can be built in the yard. This problem can be partially mitigated by reassigning the block to a new track before the first track is completely filled up. Some remaining room on a track may allow starting another (usually small) train block in space behind the first cut of cars.

To decide how to split a block, one of three processing rules might be applied depending on circumstances. If the first connection is too large to fit any available track, the block must be split in the first or current time period: as shown in Figure 4, this is called "splitting from the top." The first track is completely filled to minimize the number of tracks needed. A problem arises when it becomes necessary to split a train block "from the top." Since the length of different railcars varies, the exact sequence of cars is needed to develop the exact length of cars for each track. Instead of relying on the accuracy of this car sequencing data, a safety margin can simply be added to ensure cars can fit the allocated track space, regardless of their standing order.

If all cars fit a track in the first processing interval, a "standard split" can be implemented. As shown in Figure 5, rather than completely filling the track, the block immediately swings to a new track at the beginning of the next processing interval. A different classification may then be started in any remaining space behind the first assignment. In Figure 5, an additional three cars could fit behind the 17 cars already on the track. This strategy also may leave a "gap" of time during which no assignment is needed for the train block at all.

To determine which track should be used for a "standard split" the program determines which possible track assignments

Figure 4: Split "From the Top"
(Note: Longest available track is currently 20 cars)

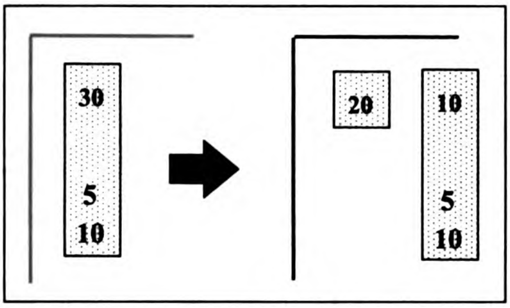

Figure 5: "Standard" Split
(Note: Longest available track is currently 20 cars)

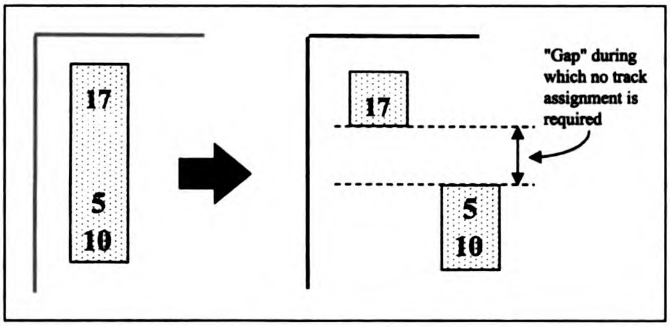

Figure 6: "Cleaner" Split
(Note: Longest available track is currently 20 cars—Fill the track completely before starting a new track)

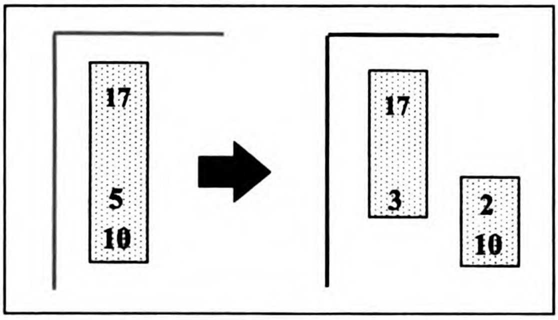

postpone the need to split the block for the longest amount of time. By examining projected track capacities versus expected accumulation of railcars for several time periods in advance, a prediction can be made in which time period each track would overflow. If two or more tracks could hold the block for the same length of time, the shortest track is chosen.

Cars needing mechanical servicing and "no bill" cars without shipping instructions must be handled differently, as shown in Figure 6. Because these cars have no outbound train scheduled, the time they will be removed from the yard is uncertain. Accordingly, no other blocks can be started behind such cars. Since free track space behind such cars provides no advantage, the first track should be completely filled before splitting over to a new track.

The more cars in a yard, the more difficult it tends to be to find track assignments of sufficient capacity to start new train blocks when needed. Under congested conditions, blocks tend to be split more often than when the yard is fluid. Because a single block then occupies more than one track, reducing the number of tracks available for classification purposes, it increases the number of rehump cars. Rather than continuing to add even more cars to an already congested yard, yardmasters often delay or slow down hump processing until cars have been trimmed out to free more track space. Such delays in hump processing cause more missed connections, but improve efficiency of switching operations in the yard.

Normal Block to Track Assignment

Once large train blocks, which require splitting, have all been assigned, smaller blocks can be processed. A decision rule is needed to assign blocks to available tracks. A "greedy" algorithm would first assign whichever block most closely matches the available track space, continuing until all blocks have been assigned. This simple method works fairly well in practice. Consider the example in Figure 7, showing two

Figure 7: Block to Track Assignment Example

2 Blocks	**Method 1:**
A: 1100'	Assign A to 1 (Waste 400')
B: 1400'	Assign B to 2 (Waste 300')
2 Tracks	**Method 2 ("Greedy"):**
1: 1500'	Assign A to 2 (Waste 600')
2: 1700'	Assign B to 1 (Waste 100')

ways to match two blocks with two tracks. In method 1, block A is assigned to track 1, wasting 400', while block B is assigned to track 2, wasting 300'. A "greedy" algorithm shown in Method 2 would first assign the block that wastes the least amount of track space, so it would match B to 1, wasting 100'. This commits assignment of A to 2 (wasting 600') in the second match, as the only remaining option.

Although both assignment methods waste the same total amount of track space, the distribution of that space is quite different. Method 1 distributes wasted space uniformly across all tracks. The problem with Method 1 is the space it leaves in the classification tracks and it is not very useful. Method 2, the "greedy" approach, fills some tracks completely while leaving substantial room behind other assignments. So, there is a high probability another block can be found to fit behind these otherwise poor assignments. Use of a greedy algorithm may not be optimal, but the decision rule does tend to perform reasonably well on average over an extended period of time.

Certain track assignments can be encouraged or discouraged through penalties and prizes. For example, a penalty is assessed for placing a block behind another one having a later scheduled trim out time. A prize is awarded for placing a block behind another block having the same yard block code. Penalties and prizes are not generally used to influence split block assignments. It is important to assign split blocks to the longest possible tracks to minimize the number of tracks taken up. Penalties and prizes could have an adverse effect if they encourage split blocks to occupy more tracks than necessary. Some assignments, such as placement of a new block behind shop cars, are prohibited outright.

Reducing the Maximum Number of Active Blocks

Track assignments may not be found for all blocks. For example, if one or more blocks needed to be "split" consuming more than one track, the assignment process may run out of tracks before it runs out of blocks. Then, another iteration of the assignment process is necessary. A constraint limiting the maximum number of active blocks, called a "choke point" constraint, must be further tightened during each time interval where the assignment process failed.

After counting the number of active train blocks during each time interval, the "choke point" constraint is reduced to one less than the number of current active train blocks in that time period. So the algorithm does not "overcorrect" and rehump more cars than needed, the constraint is tightened by no more than one block per iteration. As these "choke point" constraints continue to be tightened, additional cars are sent into the rehump tracks, until all blocks are successfully assigned. The whole process is summarized in Figure 8.

LESSONS LEARNED AND CONCLUSIONS

The benefits of classifying by train block, rather than only by yard block can be substantial. The main benefit is the ability to guarantee car connections to specific trains, and therefore to apply revenue management principles to determine which cars should move on any given train. This, in turn, would allow a rail carrier to offer a highly reliable service with guaranteed origin-destination delivery times. There also appears to

Figure 8: Process for Dynamic Block to Track Assignment

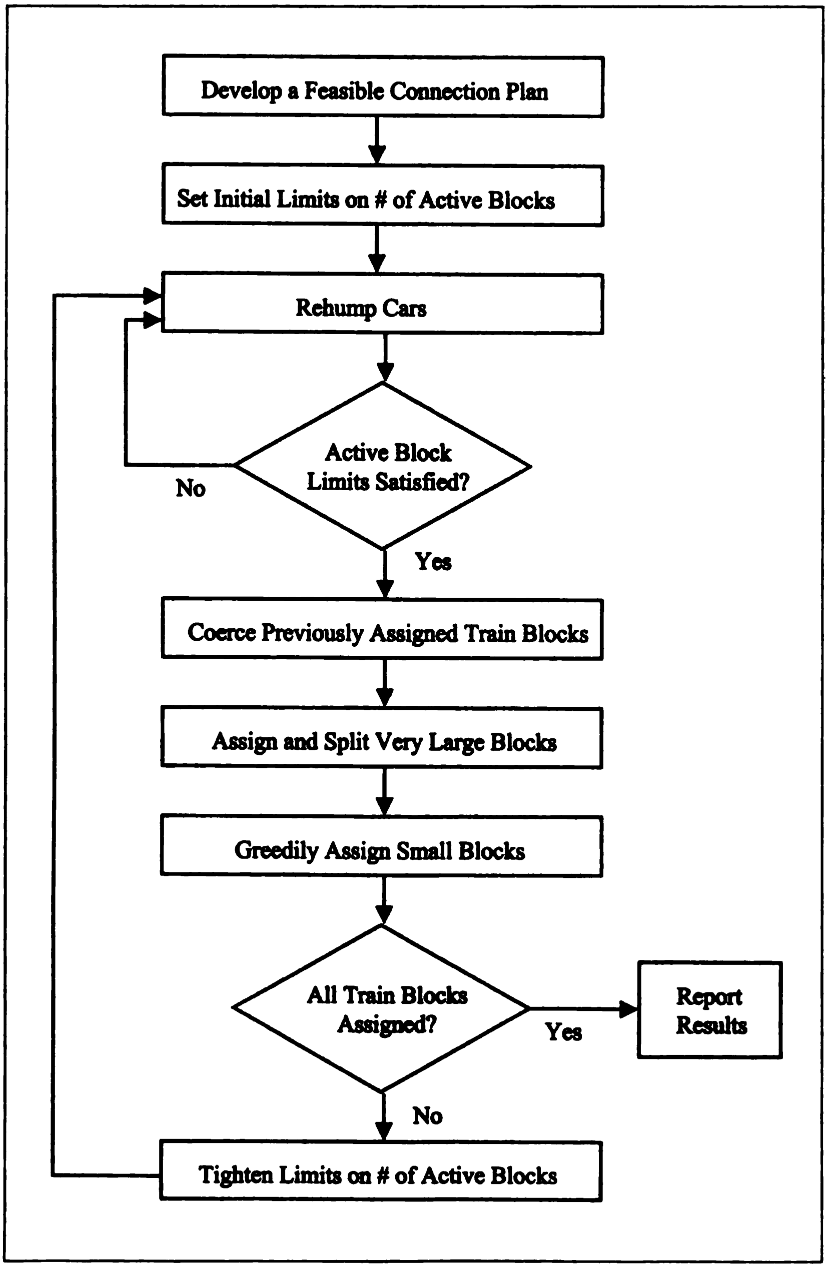

be a substantial increase in the number of blocks that can be built in any given yard. Often as many as three or four train blocks are stacked in departure order on the same track. The ability to flexibly allocate track assignments may also facilitate "opportunistic" blocking, or operation of extra trains pre-blocked to bypass intermediate yards, when traffic volumes permit. However, any plan to operate extra trains must be implemented with sufficient lead time to allow the yard to build the needed train blocks.

However, this operating approach has several disadvantages as well. Nervous behavior (where small changes in input data cause disproportionate impact on model recommendations) was seen in both the hump sequencing and block to track assignment algorithms. Track assignments for shop cars, cleaner, and no-bill cars were changed when model recommendations were periodically refreshed, resulting in these cars being scattered across the yard. Although these technical problems were later addressed, they could have been prevented altogether had time and budget constraints permitted more rigorous simulation model testing. Quantitative metrics on the performance of the algorithms could not be developed for the same reason. Conclusions reported here are based on the subjective assessments of the author and of the implementation team.

The method makes intensive use of every available inch of classification track space, but also tends to widely scatter blocks for the same outbound train across the yard, requiring frequent "crossover" movements for train assembly at the trim end. At Hinkle, this was not considered a problem because of the side-by-side configuration of the classification and departure yards. The implementation team believed that two switching crews could coordinate their activities to assemble outbound trains, so no special constraints on block placement were considered necessary. However in a larger yard with an in-line configuration of classification and departure yards, blocks would have to be clustered to ensure all cars are accessible from the same track at the trim end of the yard.

Although the TPMP test was only a pilot, designation of "priority" cars turned out to be a controversial exercise. It is essential that management trust the process by which priority cars are designated, and that direct and serious consequences actually follow if a priority connection is missed. Unfortunately, this linkage could not be demonstrated in the Hinkle pilot. Along with a fear that the pilot might deteriorate into a "cherry picking" exercise, this may have weakened management commitment to its success.

An undesirable characteristic of the TPMP prototype system was the lack of any user interface to allow yardmasters to interact with the model, perform "what if" analysis, or adjust model recommendations. Model runs were performed by a separate team of managers, operating out of a trailer detached from the main terminal tower, with the results delivered to the hump tower for execution. This had the unfortunate characteristic of taking control away from experienced yardmasters, rather than making them effective participants in the process and contributors to its success. Any future efforts along these lines should take a more inclusive approach.

The method of classifying cars by train-block is a tool for optimization of "scheduled railroad" operations, but "scheduled railroad" conditions did not prevail at the time of the Hinkle trial. The process of classifying by train-block does not tolerate last minute changes. Schlenker (1995) similarly noted a dependency of other complex operating strategies, such as block swapping, local preblocking, and two-stage tandem humping, on schedule adherence. Space is reserved in classification tracks for incoming cars. If cars do not arrive as expected, the plan is disrupted. Similarly, outbound cars must be trimmed before tracks fill up with cars, or humping must stop and the plan falls apart.

The main driver of rehump counts proved to be the level of car inventory in the classification tracks. This occurs because yard congestion increases the number of split blocks and causes the choke point constraints to be tightened. Measures affecting yard inventory, such as the scheduling of trim operations, hump sequencing strategy, and processing rate were found to have much greater impact on rehump counts than any methodology for determining which cars should have priority.

By early 1995, the Burlington Northern-Santa Fe merger had been announced, so Union Pacific's management attention necessarily shifted onto merger related matters. Given the difficulties which occurred in the initial trial of the system on the Hinkle hump, Union Pacific chose not to proceed with further implementation of the TPMP system. Nevertheless, TPMP was considered successful in establishing at least the conceptual feasibility of classifying cars by train block, rather than only by yard block, which has been universal practice previously. Given further technical development of the planning algorithms and a disciplined "scheduled railroad" operating approach, it seemed to be the consensus of the implementation team that the concepts could be made to work.

For the immediate future, the hump sequencing process described in the previous paper (Kraft, 2000b) seems to offer the best prospects for implementation. The most beneficial aspect of the hump sequencing program is its ability to identify jeopardized connections even before cars have arrived in the yard. Scheduling of classification tracks, if attempted at all, should be deferred until a second or third implementation phase. Trying to implement the whole process at one time as was attempted at Hinkle, is too much for a yard to be able to absorb without major operational disruption.

The logical next step would simply be to ask yardmasters to send low priority cars in excess of train capacity into a rehump track. By following that approach, block to track assignments could continue to be managed manually. Indeed, the cost of rehumping a few extra cars is a small price for railroads to pay for improved service reliability and satisfied customers. However, it will probably not take very long before a clever yardmaster devises a better strategy for managing the rehump cars. Once the basic techniques needed to protect priority connections have been worked out in this manner, computerized decision support tools to further refine the process can follow.

Acknowledgments

The financial support of Union Pacific Railroad is gratefully acknowledged, but implies no endorsement of the findings. Any opinions expressed herein are solely those of the author.

References

Kraft, E.R. and Guignard-Spielberg, M. (1993) *A Mixed Integer Optimization Model to Improve Freight Car Classification in Railroad Yards.* Report 93-06-06, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania.

Kraft, E.R. (1998) *A Reservations-Based Railway Network Operations Management System.* Ph. D. Dissertation, Department of Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania.

Kraft, E.R. (2000a) Implementation Strategies for Railroad Dynamic Freight Car Scheduling. *Journal* *of the Transportation Research Forum* 39(3), 119-137, published jointly with *Transportation Quarterly* 54(3).

Kraft, E.R. (2000b) A Hump Sequencing Algorithm for Real Time Management of Train Connection Reliability. *Journal of the Transportation Research Forum* 39(4), 95-115, published jointly with *Transportation Quarterly* 54(4).

Kraft, E.R. (2001) Scheduling Railway Freight Delivery Appointments Using a Bid Price Approach, *Transportation Research* A 36(2), 145-165, forthcoming.

Kwon, O.K., Martland, C.D. and Sussman, J.M. (1998) Routing and Scheduling Temporal and Heterogeneous Freight Car Traffic on Rail Networks. *Transportation Research* E 34(2), 101-115.

Schlenker, M.A. (1995) *Improving Railroad Performance Using Advanced Service Design Techniques: Analyzing the Operating Plan at CSX Transportation,* Master's Thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Wang, X. (1998) *Improving Planning for Railroad Yard, Forestry and Distribution.* Ph.D. Dissertation, Department of Operations and Information Management, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania, 17-59.

***Edwin R. Kraft** is responsible for costing systems development and new computer systems for mail and express operations at Amtrak. He holds a Ph.D. in Systems Engineering from the University of Pennsylvania.*