Use of Advanced Train Control Systems in Scheduling and Operating Railroads: Models, Algorithms, and Applications

PATRICK T. HARKER

Presented in this paper is an overview of a series of models and algorithms that have been developed for use with advanced train control systems technology on railroads to improve the reliability and costs of operations. After the conceptual framework of a hierarchy of control models is described, examples are used to illustrate the use of the various models at each level.

The railroad industry in the United States is currently undergoing major restructuring of its technology and management practices. Before the deregulation of the industry in 1980 through the Staggers and Motor Carrier acts, railroads were dominated by their operating departments; that is, they were focused on cost reductions at the expense of good marketing techniques [see Keeler (*1*) for a comprehensive review of the state of the rail industry before deregulation]. Such a situation of low cost-low quality (as measured by reliability of arrivals, loss and damage of freight, and so on) was very profitable when the U.S. economy was dominated by bulk commodity production. However, the movement toward the production of high-valued goods and the implementation of more efficient (e.g., just-in-time) inventory policies created a demand for highly reliable and flexible freight transportation services. As a result, railroads today are reinvesting in technology and restructuring their management practices to respond to the market's demand for better transport service.

Recent technological developments in advanced train control systems (ATCS) and high-speed computers have provided railroads with a unique opportunity to automate many functions in rail operations and thus to restructure their management systems. The Burlington Northern (BN) Railroad is precisely in this situation. The BN is one of the largest railroads in the United States, with approximately 25,000 mi of track covering the northwestern and central portions of the country. The BN is considered to be a very "progressive" railroad by most in the industry because of its development of many innovative technologies and management practices. For example, the BN has the highest revenue per employee at corporate headquarters (*2*).

The BN, however, has the same data problem that faces all major railroads. Of the 25,000 mi of track, one-third is "dark territory," in the sense that whenever a train enters this portion of the rail network, the dispatcher knows its position only through voice communication with the train crew. In addition, signal blocks on a railroad like the BN can be long (30 mi), and when a train enters such a block all other trains are prohibited from using that portion of track. Obviously, such a system does not make maximum use of the available track capacity. Furthermore, congestion at yards (terminals) that is caused by too many trains arriving within a short time period is a direct result of poor planning of traffic throughout the rail network and leads to sometimes dramatic underuse of yard capacity.

To overcome the difficulties mentioned above, the BN, in conjunction with Rockwell International, is in the process of developing the Advanced Railroad Electronics System (ARES). As described by Welty (*3*), ARES uses the NAVSTAR Global Positioning System, which is being developed by the U.S. Air Force to provide locational information (plus or minus 50 ft) for each train or maintenance of way vehicle on the system at any point in time (750 to 2,500 trains). In addition to this location information, ARES includes the EMS locomotive system, which provides automated procedures for train handling and energy conservation, and the ROCS dispatching system, which uses the location information from each train to help the dispatchers do a better job of operating the rail lines. Of course, any fully-implemented ATCS system will provide a similar wealth of information.

Thus, an ATCS like ARES provides a wealth of data heretofore not available to railroad management. However, this "wealth" can be more like a "flood" if the proper models and associated algorithms are not available to use this information effectively. The purpose of this paper is to provide an overview of an ongoing research project at the University of Pennsylvania that is attempting to develop such models and algorithms. An overview of the series of problems being studied is given in the first section, details on two of these models are given in the next two sections, and a summary of the progress to date and an overview of future research are given in the last section.

THE CHASE FOR MODELS

In order to use the information generated by an ATCS effectively, a series of models and computational procedures are necessary:

Decision Sciences Department, The Wharton School, University of Pennsylvania, Philadelphia, Pa. 19104-6366.

Schedule policy evaluation
↕
Tactical scheduling of trains
↕
Real-time scheduling
-trains
-locomotives
-crews
-cars
↕
Computer-aided dispatching
↕
Optimal train control

Each level of this model hierarchy is briefly discussed in the following paragraphs.

The first question to ask when implementing an ATCS is whether or not a railroad should run scheduled operations. At first glance, this seems to be a rather odd question, particularly for those accustomed to European or Japanese railroads. However, substantial cost savings can be achieved if a "tonnage" operation is run; that is, trains depart from a yard when sufficient traffic has accumulated. Of course, reliability as measured by the variance of travel time will suffer under such a system compared with a scheduled operation. In either case, the question of which policy to follow in the scheduling of trains should be made at the long-term planning level by incorporating the tradeoffs of crew and equipment costs, service quality, and the ability to affectively route empty cars and locomotives. The ability to address this long-term question requires the development of detailed simulation and analytical models that incorporate a total view of rail operations, not simply a model that focuses on the movements of loaded trains between two points.

Once an overall schedule policy has been decided, this policy must be implemented on a weekly or monthly basis. This tactical scheduling of trains differs from the previously mentioned strategic question in that all trains at the tactical level will have schedules. Thus, for those trains that must be scheduled (passenger, intermodal, etc.), the tactical scheduling procedure will create a set of feasible schedules; that is, a set of schedules that are logically consistent in the sense that an operating plan exists that can achieve the times stated in the schedules with high probability, given the delays encountered by each train as a result of random occurrences (wind, breakdowns, etc.) and interference with other trains. For trains that run on a tonnage basis, scheduled slots would exist. That is, trains would not be permitted to depart at random but instead must depart within a stated time window if they are to be operated on a given day. Thus, a tactical scheduling system must also have the capability to create such slots and check that they are feasible when considered alone and when combined with the other scheduled traffic.

Given the tactical schedules, the purpose of the real-time models is to develop operating plans that will achieve the stated schedules as well as possible, given that events have occurred (breakdowns, crew shortages, etc.) that disrupt the plan of operations on which the tactical schedules are based. For trains, the aim is to develop a plan of arrival and departure times at each major yard or, more generally, at each point at which the planning of the train operations changes (that is, a boundary of the dispatchers' territories). For crews, locomotives, and cars, their movements are planned to guarantee that sufficient resources are available at each yard to achieve the tactical schedule plan.

After defining the arrival and departure times of the trains at the boundaries of the dispatchers' territories (i.e., a planning line), the computer-aided dispatching system attempts to schedule the meets and passes along a rail line along with planned arrival and departure times at intermediate points (sidings, beginnings, and ends of double track, etc.) to assure compliance with the times passed from the train-scheduling model. Several approaches have been proposed for this function (*4*), but all tend to ignore the fact that significant fuel savings can be achieved by pacing trains; that is, to have the trains travel at less than maximum velocity to save fuel. In addition, the planning of meets and passes along with a planned pacing of trains will tend to increase the probability of arriving at the destination on time because it is possible to speed up if disturbances do occur. Planning at maximum velocity does not provide this flexibility.

Finally, the dispatching system provides each train with a specific goal for the time and velocity at which it should reach each point on its path. The engineer and the on-board computer system must then calculate a velocity profile (a combination of throttle and dynamic-air brake settings) that will achieve this goal in a safe and fuel-efficient manner. Again, a pacing problem must be solved for the train, a problem that is now much more complex because of the nature of train forces and handling techniques.

This discussion has described the flow of information down the model hierarchy. Of course, the reverse flow is also very important. The train must constantly inform the dispatching model of its location and performance, the dispatching system must inform the network control model of the status of planning lines, and the performance of the network control system (the interline planner) must be monitored to assess the long-term viability of various schedule policies.

At present, the research program underway at the University of Pennsylvania is attempting to address all of these issues. In the following paragraphs, two topics will be discussed: (a) the computer-aided dispatching system and interline planning model, and (b) a new decision-support system for tactical scheduling. Because of length requirements, all of the details of these models cannot be discussed in this paper. However, reference is made to the relevant technical reports that are available from the author.

TACTICAL SCHEDULE VALIDATION AND CREATION

Given the overall policy concerning the frequency of train departures, the tactical scheduling problem is to create schedules for all trains that are logically consistent; that is, that there are operating plans that can achieve these schedules with high reliability. As described by Assad (*5*), many simulation and optimization models exist for the analysis of rail operations. However, no model exists that can answer the simple question: Is a given set of schedules feasible under the best operating conditions in the sense that there exists a plan of operation that can achieve the scheduled times? If not, what minimal changes can be made to the schedules to make them feasible? If they are feasible under the best circumstances, what is the reliability of achieving these scheduled times when adverse conditions exist? Note that a large-scale optimization model could be developed that would attempt to find optimal schedules, given well-defined cost or profit criteria (see, for example, Crainic et al. (*6*). However, the definition of such an objective function is extremely difficult, given the tradeoffs of marketing concerns, costs, crew, and equipment use. Thus, the approach taken in the Schedule Analysis (SCAN) system (*7*) is to provide a decision-support tool that answers the logical questions of whether or not schedules are feasible, and leaves the marketing-cost tradeoffs to the analyst. As designed, SCAN is meant to support weekly or bimonthly updates to the stated schedules.

SCAN is an interactive decision-support system that contains three modules: a data base system for the updating of track and train data as well as train schedules, an algorithm for checking whether or not a given set of schedules is feasible, and a Monte Carlo simulation technique for the calculation of the reliability of a given set of schedules. The feasibility algorithm takes as input the train schedules, track topology, and the free (unobstructed) meetpoint-to-meetpoint running times for each train, which are calculated by one of many train performance simulators (TEM, TPS, etc.). Given this data, the feasibility algorithm searches for a meet-pass plan that can achieve this given set of schedules. If no plan can be found, the schedules are labeled infeasible and the algorithm presents the plan that would require the minimal change to the schedules to become feasible. The details of this integer-programming-based algorithm can be found in Jovanović and Harker (*7*). If the analyst wants help in changing the schedules to achieve feasibility, SCAN contains a set of heuristics to attain this goal. However, the analyst is encouraged to make these changes manually because of the complex tradeoffs mentioned previously.

Once the schedules have been modified so that they are feasible in the best case, the analyst may wish to know how often feasibility would be maintained under more adverse conditions (adverse weather conditions, breakdowns, etc.). SCAN answers this question through a simulation technique in which probability distributions of the free-running times for the trains are used as input to a Monte Carlo model. The result of this simulation is the percentage of time adherence to the schedules under variable operating conditions can be expected.

To illustrate the working of the SCAN system, consider the example given in Figure 1; this shows the track topology on the vertical axis, the time of day on the horizontal axis, and the schedules for each train as straight lines connecting the departure and arrival times. Looking quickly at this set of schedules, it is tempting to conclude that they are feasible, given the spacing of the schedule lines. However, the analysis of these schedules with SCAN first uncovers the problem that some trains are scheduled to operate faster than is physically possible (i.e., in time lower than the free-running time). Once these problems are resolved, SCAN begins to uncover more subtle problems. For example, in Figure 2, no plan exists that could have Train 3 and Train 34 both arrive on schedule; in the best case, Train 34 would be late by 10 min. Thus, the schedule of Train 3, Train 34 or both must be changed to become feasible. After many such changes, a feasible schedule is achieved, as indicated by the feasible meet-pass plan shown in Figure 3. Once these feasible schedules are found, a simulation analysis finds that the schedules are not very reliable; that is, the schedules were feasible in only 8 percent of the cases in which random delays to the trains were introduced. Thus, more time must be added to certain train schedules to increase this reliability. The details of several other examples that illustate the various features of SCAN can be found in Jovanović and Harker (*7*).

SCAN is currently being used to reschedule a major U.S. railroad as well as to analyze various capital improvements and maintenance policies. The ability to achieve a given set of schedules is obviously influenced by the track topology. The impact of changes in track layout on the performance of the train movements should be carefully considered; with SCAN, this relationship can be made explicit and seems to be a major use of such a system. For example, consider the situation shown in Figure 4, which is a portion of double-tracked railroad with two small pieces of single track. In analyzing this situation with SCAN, the problem that is uncovered is not necessarily that single track exists but, rather, that the speed limits on the portion of single track between MTPNT-2 and MTPNT-3 continually create infeasibilities in the schedules (note the shallow slope of the lines in Figure 4 on this portion of the track). Thus, one way to resolve this problem is to upgrade the single track to allow higher speed limits and not to go to the expense of adding an additional track at this point.

REAL-TIME CONTROL OF TRAIN MOVEMENTS

Once the tactical schedules have been set for the day, the purpose of the real-time scheduling system is to attempt to achieve the times stated in the schedules with a high degree of certainty. In practice, events (breakdowns, accidents, etc.) will occur that may inhibit the system from attaining the scheduled goals. Thus, the real-time models attempt to minimize the deviations from these goals, and, at the same time, operate the trains in a safe and fuel-efficient manner. In this section, two such models will be described, along with the results of preliminary empirical studies.

Network Control of Train Movements: Interline Planning

The interline planning model attempts to minimize the deviations of arrival-departure times at various points on the rail network for each train from the times stated in the tactical schedules. As described by Harker and Kraay (*8*), this problem can be formulated as a large-scale mathematical program. This model takes the following general form:

Minimize disruptions to schedule + block switching delays
+ costs for work rule violations

Subject to:

Crew change constraints

Physical constraints of the trains

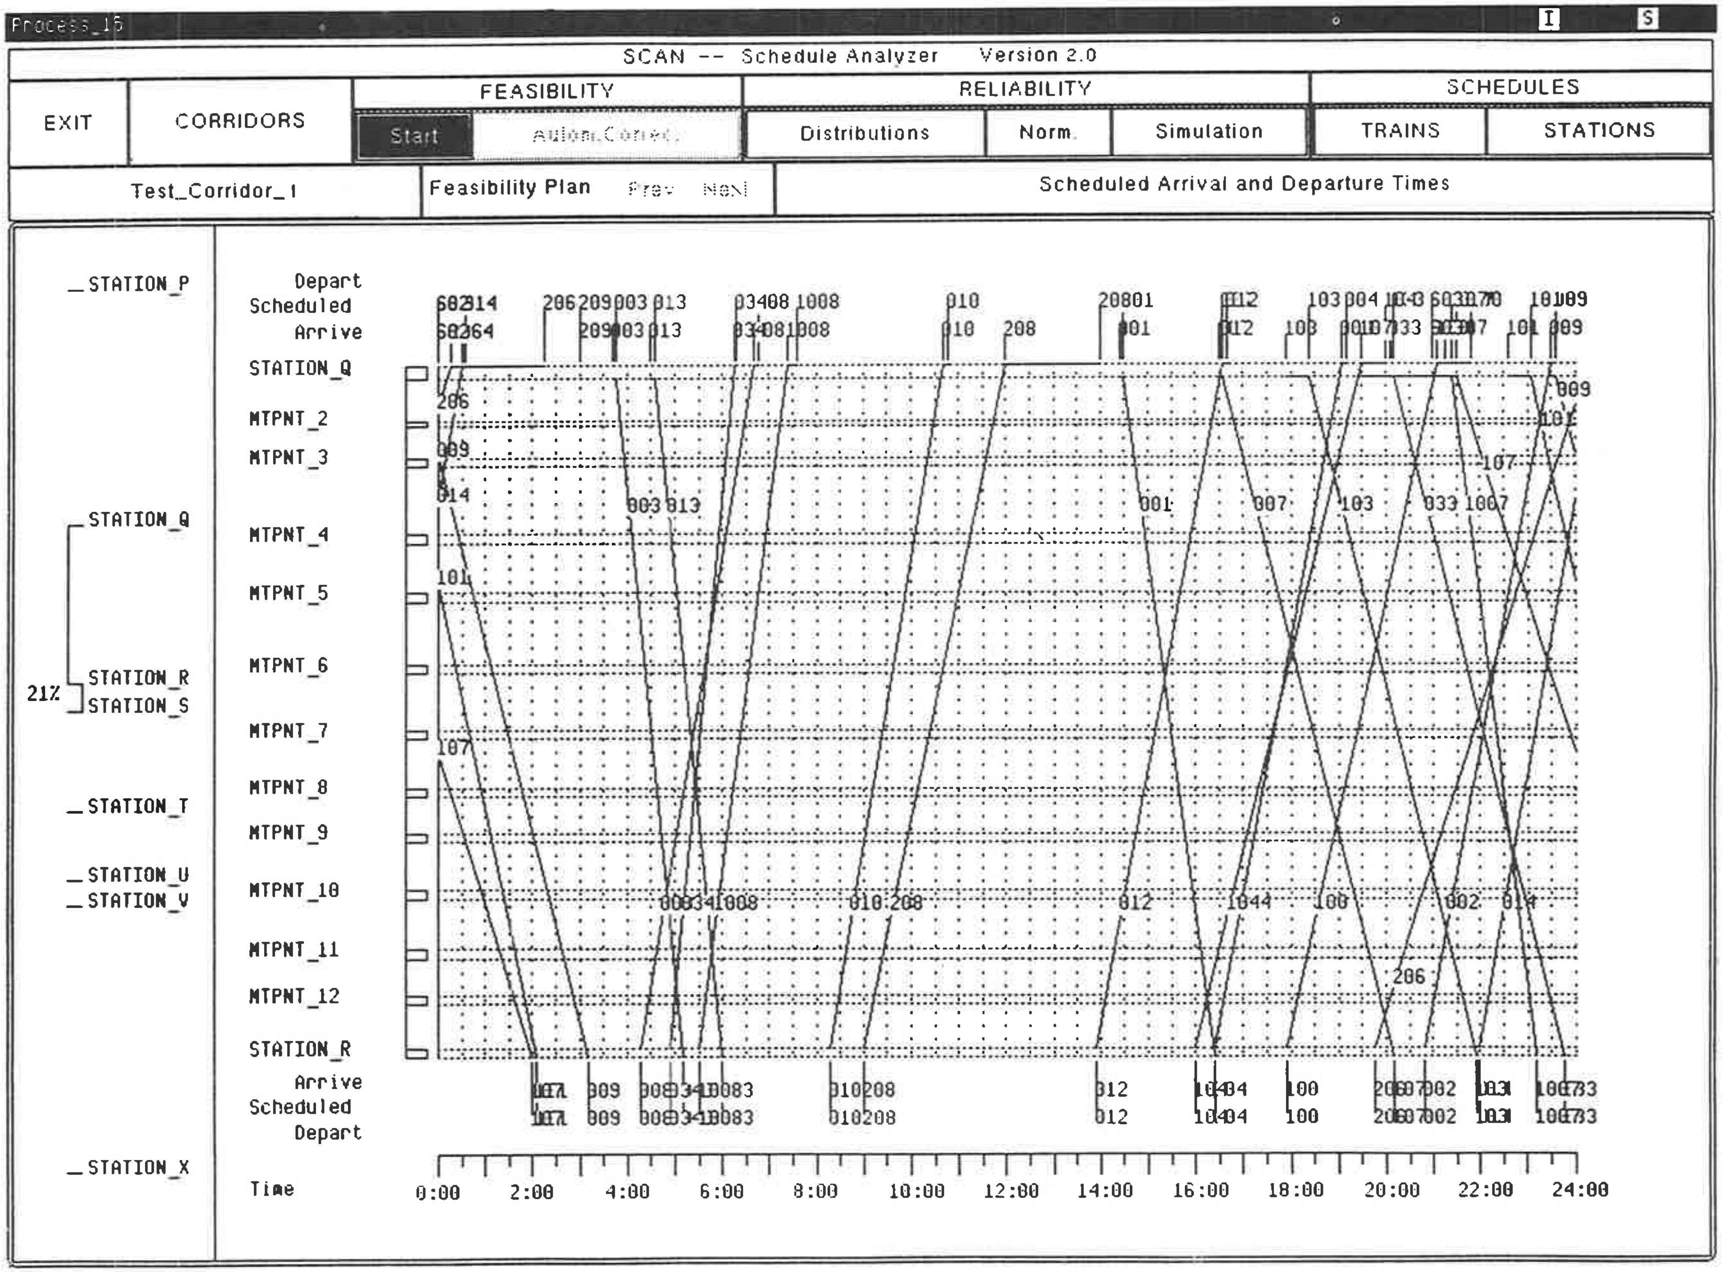

FIGURE 1 Schedules for SCAN example.

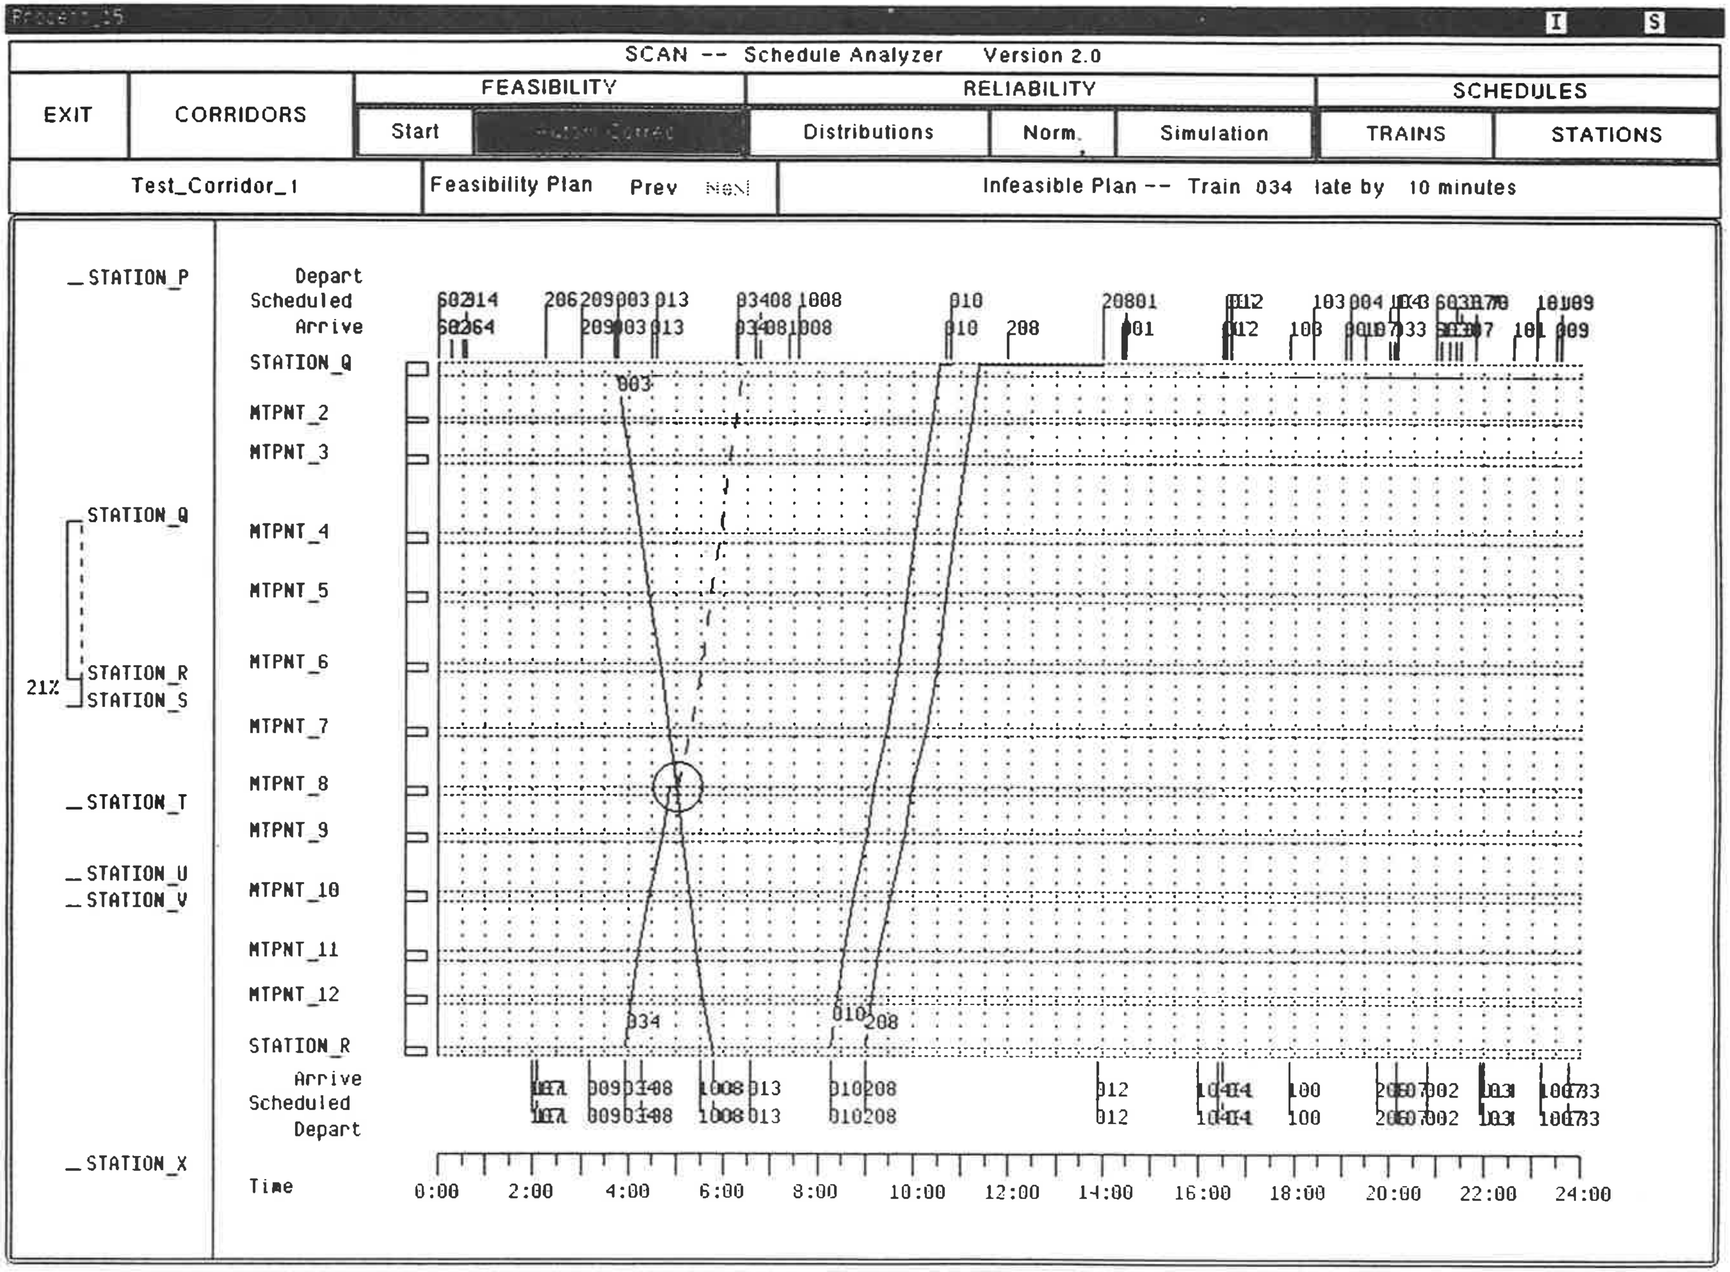

Test_Corridor_1

Infeasible Plan -- Train 034 late by 10 minutes

FIGURE 2 Infeasibilities uncovered by SCAN.

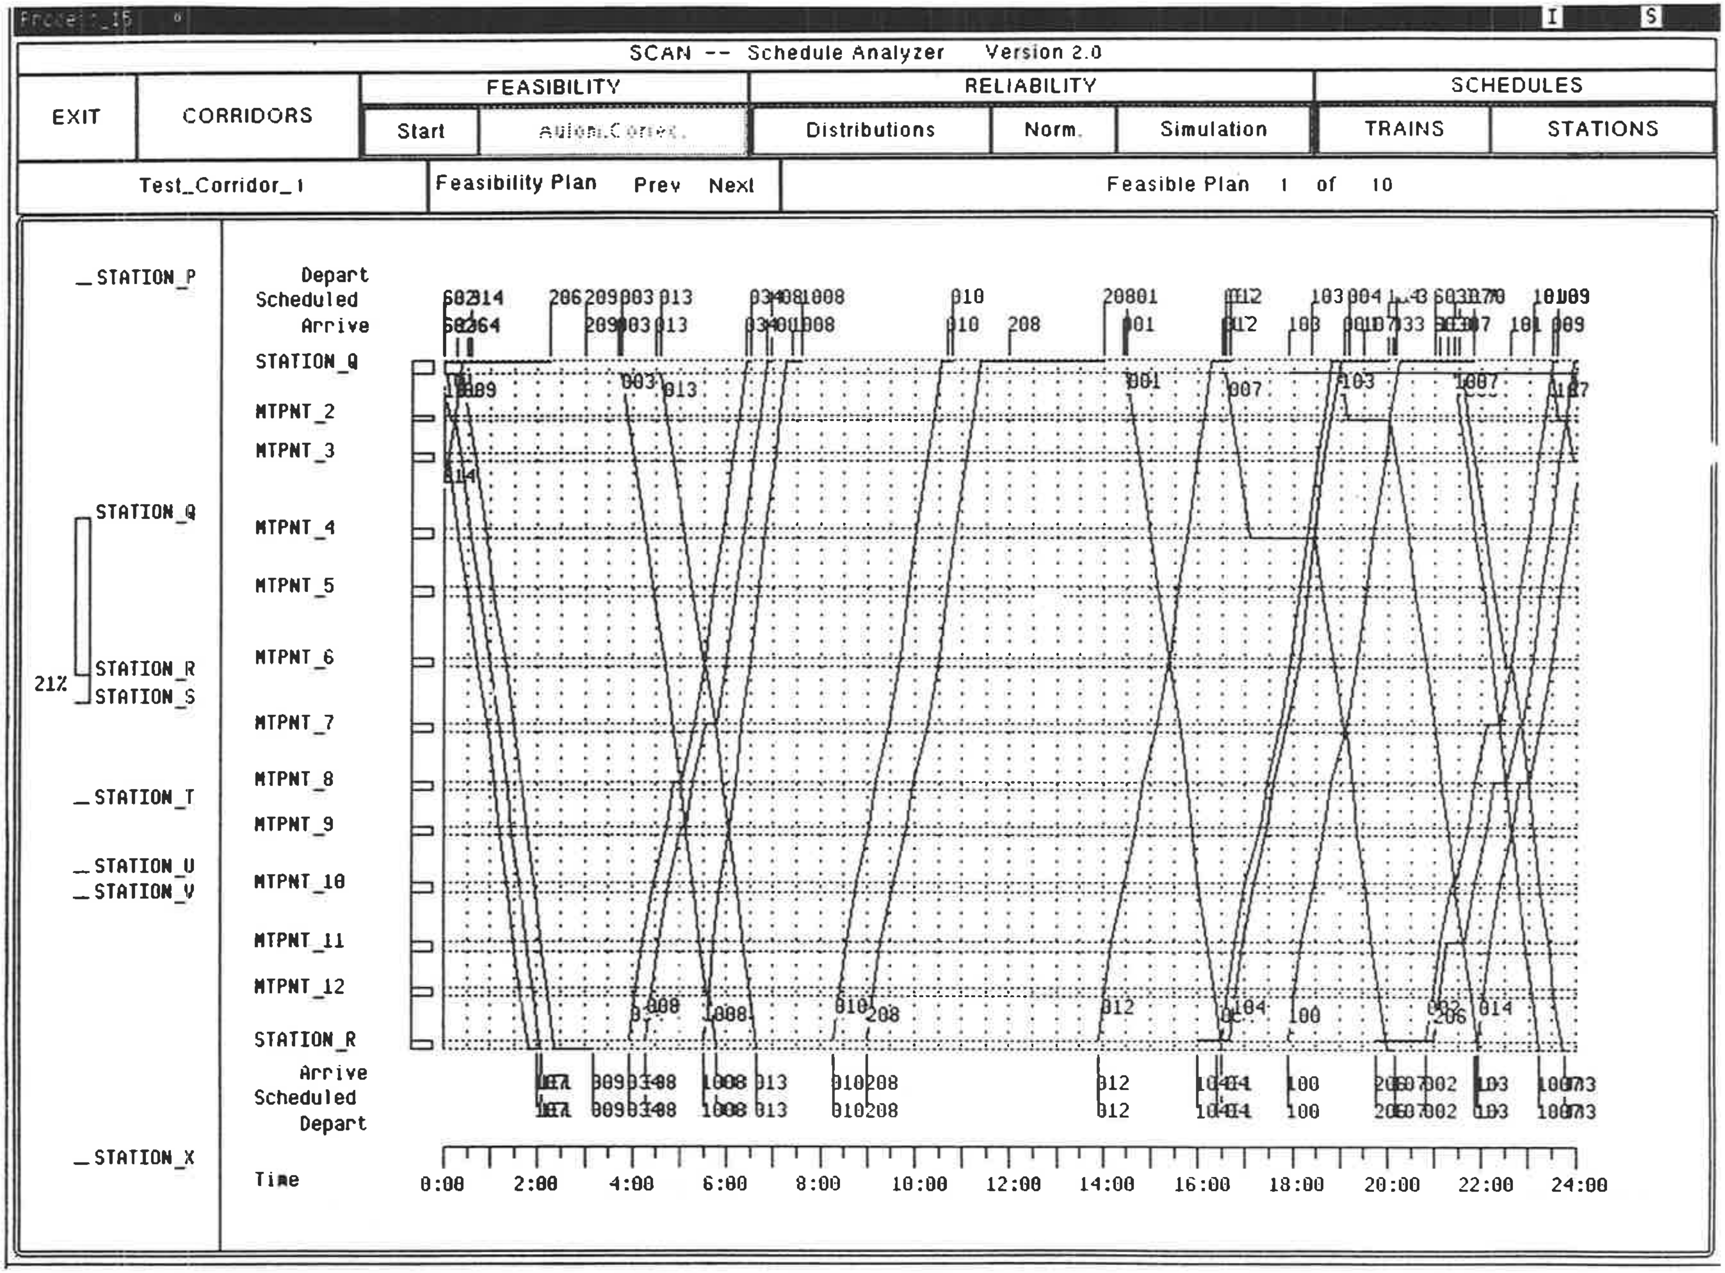

FIGURE 3 One of several feasible operating plans.

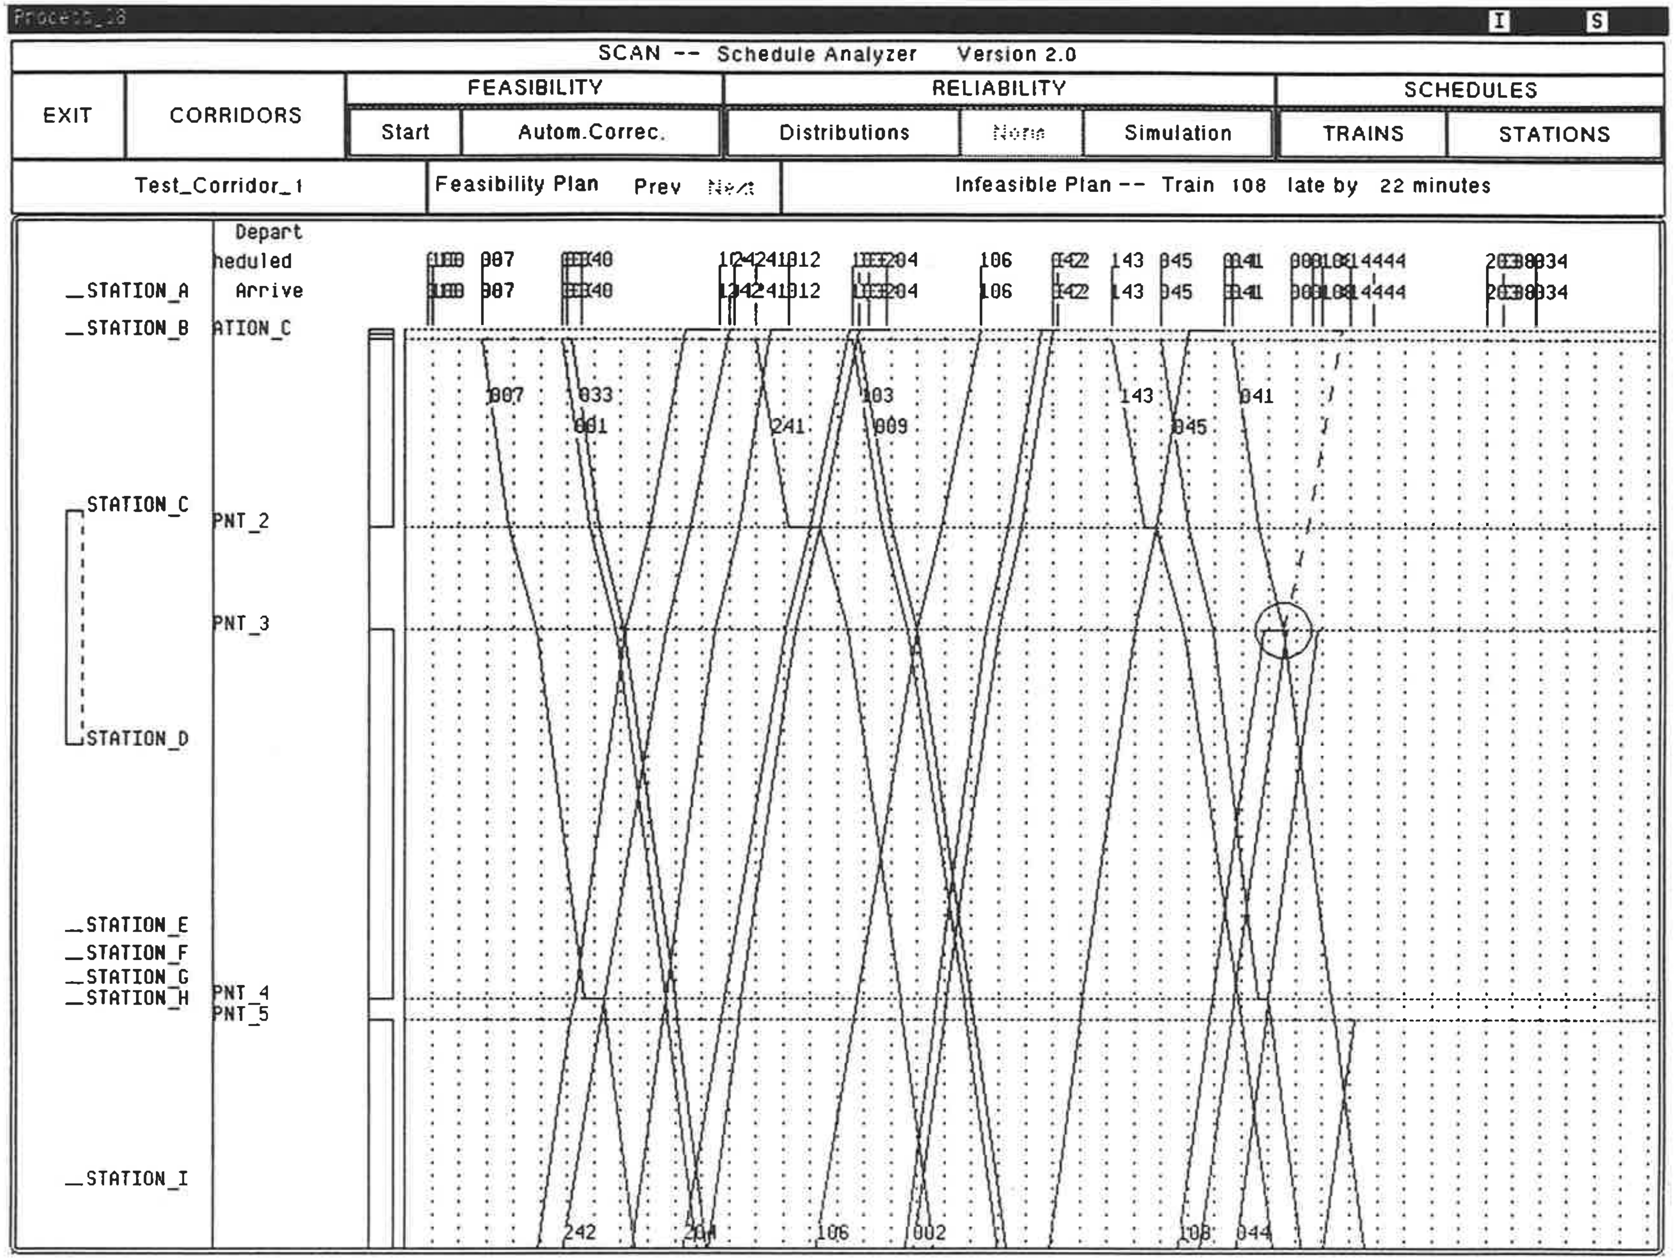

Infeasible Plan -- Train 108 late by 22 minutes

FIGURE 4 Double-track bottleneck example.

Arrival time $\geq$ departure time + free-running time + delays

Logical constraints

The disruptions to schedule can be any metric of the time of arrival-departure at a point (the variables) and of the stated times in the schedule (the data from a SCAN-like system). In particular, these metrics may be weighted because for a given point, it may not be crucial that a particular train arrive on time, but for another train its on-time arrival may be vital. The cost of block-switching delays refers to the fact that cars will most likely have to switch trains at least once in their journey from origin to final destination. Blocks of cars are often scheduled to travel on one train and then switch to another train at a predefined yard. Thus, a precedence relationship is defined for the arrivals of trains at a particular yard by these block-swapping conditions. Of course, if a block of cars misses a particular outbound train, it can travel on another departing train, but with a possible increase in the total travel time for the cars. The cost of the block swapping reflects this increased cost resulting from cars missing their planned connection at a yard. Finally, train crews are required by law to work no more than a prespecified number of hours. If the crews reach this limit, various penalties are assessed; these penalties define the last term of the objective function.

The first set of constraints simply states that crews must be changed at prespecified points on the network. The physical constraints of the train assure that each train departs after it arrives from a particular point, that sufficient time is given to the train if it must perform work at a given point (picking up and setting out cars, maintenance, etc.), and other such conditions. The third set of constraints states that the total running time of a train (arrival at point $i + 1$ minus the departure from point i) must be greater than or equal to the free running time of the train plus any interference delays caused by the meeting and passing of other trains on the system. Finally, the logical constraints ensure that if two trains are scheduled to meet or overtake on a specified portion of the network, then this activity will occur at the stated point.

The interference delays used in the third set of constraints merit discussion. There exists a large amount of literature dealing with the delays encountered by trains operating on single- or double-track railways. However, these models all assume that trains depart randomly according to a uniform or Poisson distribution. In reality, the trains that are considered within the planning horizon of the interline planning model will depart at or near the planned departure time. That is, the departures are not purely random but rather occur with some error around the stated departure time. To correct for this inaccuracy in the literature, Chen and Harker (*9*) have developed a model of delay for scheduled traffic that is formulated as a system of nonlinear equations. Using the successive approximation algorithm, Chen and Harker show how the mean and variance of travel times and hence the reliability of on-time arrival can be efficiently calculated.

The model just described is formulated in Harker and Kraay (*8*) as a mathematical program with a nonlinear objective function, nonlinear constraints because of the delay functions, and integer variables arising from the logical constraints. Research is currently under way to develop algorithms for this problem that are suitable for parallel-computing environments. A preliminary discussion of this research can be found in Herker and Kraay (*8*).

Computer-Aided Dispatching: the Pacing Problem

Once the interline planning model computes the time windows (targets) for the arrival and departures of each train in the network, the goal of a computer-aided dispatching system is to derive a meet-pass plan for the operation of a given planning line (the portion of the rail network between two specified points that makes up a dispatcher's region of authority). There have been many attempts at developing such a system (*4*, *9*). All of these methods try to minimize some measure of cost while assuring that the line is operated safely. Typically, this cost consists of fuel consumption and the cost of arriving early or late at the ends of the planning line. The algorithms are typically simple branch-and-bound methods that implicitly enumerate all feasible plans.

Two problems exist with the current state-of-the-art in computer-aided dispatching. First, by treating the arrival times as a cost rather than as a hard constraint, the models provide the dispatchers with a great deal of freedom to operate their line efficiently. Such freedom typically evolves into a system in which trains are given absolute priorities and some trains are made very late at the expense of others. Furthermore, the dispatchers are often too busy to consider the impacts of late or early arrivals on the performance of the rail network outside their regions of authority. However, it may often be the case that a high-priority train may be delayed to expedite the arrival of a late train even if the latter train has a low priority; priorities are therefore endogenous rather than specified a priori. Also, the minimization of cost along a single planning line may lead to a suboptimal operating plan for the entire network unless the impacts outside the planning region are taken into consideration.

The second problem with the current state-of-the-art involves the hurry up and wait philosophy on which most rail systems operate. Consider, for example, Train 007 in Figure 3. At MTPNT-3, this train arrives $1\frac{1}{2}$ hr earlier than necessary in order to meet the two northbound trains. Because fuel consumption rises as the square of velocity according to the David formulae (*10*), it is far better to pace this train to MTPNT-3 so that it will travel at a lower speed from STATION-Q to this point. Thus, one can simply slow down a train to arrive on time at a planned meet. Can it be done even better? Consider Trains 103 and 100 on the right-hand side of Figure 3. Note that Train 103 arrives approximately 1 hr early at MTPNT-2 for its meet with Train 100. Train 100, on the other hand, arrives $1\frac{1}{2}$ hr early at its destination, STATION-Q. Why not simply slow down both trains? If this were done, Train 100 would not make its meet with Train 007 at MTPNT-7, Train 103 would be late for its meets at MTPNT-10 and MTPNT-11, and so forth. The problem with changing the times of Trains 100 and 103 is that the locations of the meets have been decided, a priori, rather than making this decision simultaneously with the times of arrivals at each meetpoint (and hence, the planned velocity of each train).

The pacing model, as defined by Kraay, Harker, and Chen (*11*), is a mathematical program that attempts to simultaneously find the meet-pass plan (where trains meet or pass) and velocity profiles for each train (their arrival times at each meet-pass point), which minimizes the cost of operating a rail line subject to the scheduled time windows and at the same time conforms to the various operating policies of the railroad. In addition to conserving fuel, this notion of pacing may increase the reliability of train operations. If plans are made in such a way that all trains travel at maximum velocity, then any disruptions can propagate throughout the line, delaying many other trains. By pacing, late trains may have excess power, which will permit them to travel faster than planned to achieve the stated arrival times if disruptions do occur.

The pacing model selects the locations for each meet and overtake, as well as the time of arrival of each train at each intermediate point in the planning line so as to

Minimize cost of fuel + operating penalties

Subject to

Meeting the scheduled time windows at the ends of the planning line
Physical constraints of the trains
Speed restrictions
Logical constraints

The objective function of this model is nonlinear because of the fuel consumption term and the various forms that the operating penalties can exhibit. The time windows simply state that each train should not be permitted to leave the origin yard before the time defined by the interline planner, and should not arrive early or late to the destination yard. The physical constraints portray the physical capabilities of the train vis-à-vis acceleration and deceleration, and the speed restrictions ensure the safe operation of each train. The logical constraints are used to ensure that siding capacities are not exceeded; headways between following trains are maintained; various priority rules are observed; and that any other "reasonable" conditions, such as following trains being permitted to pass one another once at the most (i.e., no leap frogging) are observed. Thus, the pacing model is a large-scale, mixed integer, nonlinear program that must be solved in real time and with a range of solutions—not just one. This latter condition is essential if the model is to be used effectively, because dispatchers may often reject the optimal solution in favor of some other, less optimal solution because of circumstances not considered by the pacing model.

In Kraay et al. (*11*), several alternative algorithms were considered. The best solution procedure is a rounding heuristic in which a velocity profile for each train is computed for each train by not considering the interaction with any other trains. This problem becomes a much smaller nonlinear program that has a special structure. Once these "unconstrained" velocity profiles (and hence, arrival times for each train at each point) have been computed, any conflicts that occur at infeasible points (e.g., a meet in the middle of single track) can be moved to the nearest siding and all of the necessary logical constraints can be obeyed at the same time. This rounding procedure can be accomplished through a modification of the SCAN feasibility algorithm described in the previous section. Once a feasible meet-pass plan has been found via this rounding procedure (the places where trains are scheduled to interact), a nonlinear program with additional constraints is solved in order to compute the times of arrival. This last step is necessary because of the interactions between all trains previously described in the case of Trains 100 and 103; that is, the algorithm must attempt to adjust all the times simultaneously to avoid infeasibilities. In certain cases, this simple rounding procedure can be proven to produce the optimal solution. In other cases, the experimental work reported by Kraay et al. (*11*) shows that this heuristic is quite good.

Preliminary empirical evidence suggests that significant fuel and delay costs can be achieved through the use of this model. In the analysis of current practice, dispatchers tend to become overburdened when many trains are placed under their control. In such cases, they tend to follow the simple practice of dealing first with the highest priority trains, and then progressively moving toward those trains with low priority. The pacing model, by treating all of these decisions simultaneously, often yields significant cost savings. The details of this empirical work will be reported in a subsequent paper. Finally, this notion of pacing extends to many other areas of transportation. For example, the scheduling of barge and ship traffic in a canal (*12*) fits well into this paradigm; these topics will also be explored in the future.

Optimal Control of Train Movements

The pacing model provides the train with the time at which it must reach the next point on its path as well as the velocity at which it should pass this point. The goal of the onboard computer system is to help the engineer achieve this time and velocity constraint in a safe and fuel-efficient manner. This problem has been formulated by Harker and Chen (*13*) as a nonlinear optimal control problem. In fact, both a deterministic model and a stochastic model that take into account the random nature of train performance caused by engine problems, wind, other weather conditions, and so on, have been formulated and analyzed. Research is now underway to develop fast and effective solution procedures for these models.

Summary and Future Research

The hierarchy of models presented in this paper has one goal in mind: to smooth the flow of traffic in rail networks by effectively using the wealth of information available from an ARES-like positioning system. In order to achieve this goal, a simple principle applies: keep it simple! Major policy tradeoffs are made at the top, the SCAN system attempts to implement these policies through the development of tactical schedules, and the real-time control systems develop operating plans that achieve these goals while optimizing performance. Note that this flow of authority is quite different from that typically seen in railroad control systems in the United States; in such systems, cost is typically the driving force. In the schema presented in this paper, the marketing-customer concerns drive the schedules and thus the entire operating philosophy. Simplicity is achieved by clearly stated goals: dispatchers are to obey time windows, engineers the arrival times given by the dispatcher, and so on.

The research that is currently under way at the University of Pennsylvania involves the fleshing out of this hierarchy through the development of the necessary models and algorithms. In addition, various cost-benefit studies are being pursued to ascertain the ability of such a system to improve the reliability and costs associated with freight railroading. In addition, extensions of these concepts to other modes of transportation and, in general, manufacturing processes are currently being explored.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation Presidential Young Investigator Award and by a grant from the Burlington Northern Railroad. The comments of David Kraay, Dejan Jovanović and the staff at the Burlington Northern Railroad are warmly acknowledged.

REFERENCES

1. T. E. Keeler. *Railroads, Freight and Public Policy*. Brookings Institution, Washington, D.C., 1983.
2. T. Moore. Goodbye, Corporate Staff. *Fortune*, Dec. 21, 1987, pp. 65–76.
3. G. Welty. BN and ARES: Control in a New Dimension. *Railway Age*, May 1988, pp. 24–26.
4. E. R. Petersen, A. J. Taylor, and C. D. Martland. An Introduction to Computer-Assisted Train Dispatch. *Journal of Advanced Transportation*, 20, 1986, pp. 63–72.
5. A. A. Assad. Models for Rail Transportation. *Transportation Research*, 14A, 1980, pp. 205–220.
6. T. Crainic, J. Ferland, and J. Rousseau. A Tactical Planning Model for Rail Freight Transportation. *Transportation Science*, 18, 1984, pp. 165–184.
7. D. Jovanović and P. T. Harker. Railroad Schedule Validation and Creation: the SCAN I System, *Transportation Science*, 1988.
8. P. T. Harker and D. Kraay. *Real-Time Scheduling for Rail Networks: Model Description and Proposed Solution Procedures*. Working Paper, Decision Sciences Department, The Wharton School, University of Pennsylvania, Philadelphia, 1988.
9. B. Chen and P. T. Harker. Two Moment Estimation of the Delay on Single-Track Rail Lines with Scheduled Traffic, *Transportation Science*, 1988.
10. E. K. Morlok. Introduction to Transportation Engineering and Planning. McGraw-Hill, New York, 1978.
11. D. Kraay, P. T. Harker, and B. Chen. Optimal Pacing of Trains in Freight Railroads: Model Formulation and Solution. *Operations Research*, 1988.
12. E. R. Petersen and A. J. Taylor. An Optimal Scheduling System for the Welland Canal. *Transportation Science*, 22, 1988, pp. 173–185.
13. P. T. Harker and B. Chen. *Optimal Control of Train Movements*. Working Paper, Decision Sciences Department, Wharton School, University of Pennsylvania, Philadelphia, 1989.

Publication of this paper sponsored by Committee on Electrification and Train Control Systems for Guided Ground Transportation Systems.