

The Canadian Pacific Railway Transforms Operations by Using Models to Develop Its Operating Plans

Phil Ireland, Rod Case, John Fallis

Canadian Pacific Railway, Gulf Canada Square, Suite 400, 401-9th Avenue SW, Calgary, Alberta, Canada T2P 4Z4
{phil_ireland@cpr.ca, rod_case@cpr.ca}

Carl Van Dyke, Jason Kuehn, Marc Meketon

MultiModal Applied Systems, Inc., 125 Village Boulevard, Suite 270, Princeton, New Jersey 08540
{carl@multimodalinc.com, jason@multimodalinc.com, marc@multimodalinc.com}

North American railways have traditionally practiced tonnage-based dispatching, running trains only when they have enough freight. As a result, their customer service and their use of crews, fixed assets, locomotives, and railcars are poor. Canadian Pacific Railway is using new decision-support tools developed in-house and by MultiModal Applied Systems to create a scheduled railway. These tools use operations research approaches, such as an optimal block-sequencing algorithm, a heuristic algorithm for block design, (very fast) simulation, and time-space network algorithms for planning locomotive use and distributing empty cars. This implementation has saved \$300 million Canadian (US\$170 million) from mid-1999 through autumn 2000. We estimate it has saved at least an additional \$210 million Canadian during 2001 and 2002 in fuel and labor costs alone. Labor productivity, locomotive productivity, fuel consumption, and railcar velocity have improved by 40, 35, 17, and 41 percent, respectively. Furthermore, Canadian Pacific Railway now provides its customers with reliable delivery times and has received many customer and shipping association awards for its improvement in service.

Key words: decision analysis: applications; transportation: freight-materials handling.

Over 3600 freight railroads operate in Canada, Mexico, and the United States. They form a seamless integrated system that provides the world's most efficient, cost-effective freight service. North American railways operate over 170,000 miles of track and produce US\$42 billion in annual revenues. Railways remain the backbone of North America's freight-transportation network. Furthermore, the rail industry is at the center of many critical issues: improving North America's productivity, reducing road congestion, improving transportation safety and border security, and reducing greenhouse gas emissions.

Incorporated in 1881, Canadian Pacific Railway (CPR) is one of Canada's oldest corporations and was North America's first coast-to-coast transcontinental railway. CPR transports rail freight over a 14,000-mile network extending from Montreal to Vancouver and throughout the US Northeast and Midwest. Alliances with other carriers extend CPR's market reach beyond its own network and into the major business centers of Mexico (Figure 1).

In the mid-'90s, CPR was struggling with high costs, low profitability, and rising customer-service requirements. CPR thought its traditional operating strategies would not be adequate for dealing with these issues. CPR needed a new plan.

Although rail is an old technology, today's railways are complex operations. Every day CPR receives approximately 7,000 new shipments from its customers going to destinations across North America and for export. It must route and move these shipments safely and efficiently over its 14,000-mile network of track. It must coordinate the shipments with its operational plans for 1,600 locomotives, 65,000 railcars, and over 5,000 train crew members and take into account the capacity and storage space at 250 yards. Overall, CPR has 6,000 customers shipping via 20,000 distinct origin-destination pairs. In planning, it must also account for track-maintenance windows and connections with other railways. These vital connections account for 40 percent of CPR's business. The railway must manage and integrate a complex set of issues and assets efficiently, seven days a week, 24 hours a day.

To meet rising customer expectations and to make a return on capital investment, CPR decided to make a wholesale change in its operating philosophy.

Like most large North American railways, CPR used a tonnage-based approach in dispatching trains, holding all trains until it had enough tonnage to fill them to capacity. Under the tonnage-based approach, the operating plan may list a train as operating every day,

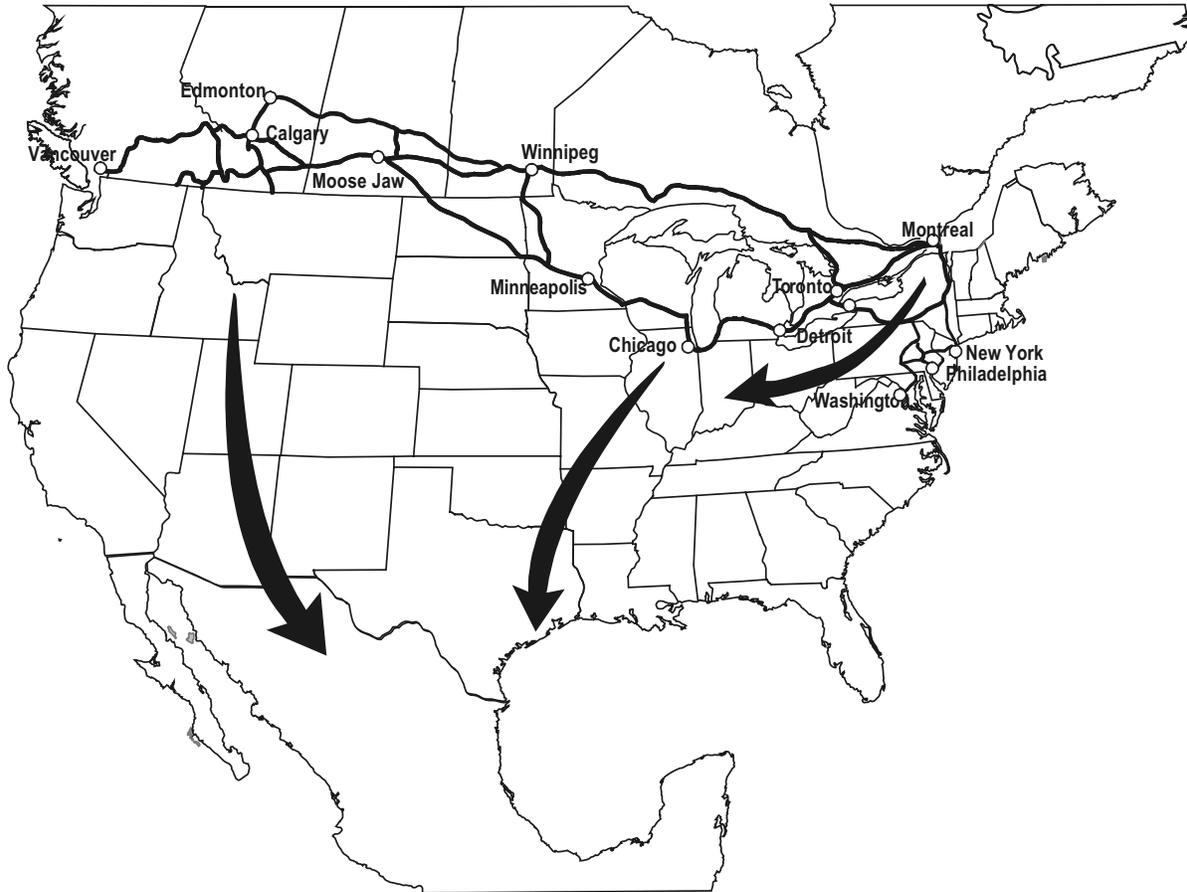


Figure 1: This Canadian Pacific Railway (CPR) system map shows some of the 4,200 locations CPR serves. Because railways interconnect, much of CPR's traffic moves to and from other railways in the US and Mexico.

but if the railway cannot fill enough railcars, it cancels or delays the train. In using this approach, CPR tried to minimize the total number of trains it operated by maximizing their size, which, in theory, minimizes crew costs and maximizes track capacity. However, tonnage-based train planning has serious drawbacks:

(1) The yards cannot fine-tune their operations based on a repetitive schedule, and they require more railcars and greater storage capacity to cope with the traffic variability.

(2) Demands for crew and locomotive resources may increase along with the costs for repositioning crews and equipment.

(3) Most important, customers suffer from unreliable service because the railroad gives train-operation economics priority over customer needs.

The alternative to the tonnage-based approach is a more disciplined, schedule-based approach. Scheduled railway strategies are gaining favor in North America as railways use new management science tools, particularly MultiRail, to craft cost-effective and customer-effective operating plans. CPR, Norfolk Southern, and Canadian National have made the bold-

est moves in this direction. In 1997, CPR began exploring the concept of running a scheduled railway, and it was one of the first railways to adopt a true schedule-based approach that allows it to adjust quickly to changing traffic demands. CPR has become rigorously disciplined in its scheduling.

The schedule-based approach forces trains to run on time, as scheduled, even if they travel with light loads. Until recently, the railway industry shunned scheduled strategies for several reasons:

(1) They require operating trains with low tonnage when customer demand is below expectations.

(2) They depend on railways' systematically forecasting traffic levels by the day of the week, and quickly adjusting the plan.

(3) They require a granular, actionable understanding of each customer's requirements in each corridor.

(4) The needed schedule-based models require sophisticated operations research software to conduct comprehensive and timely analyses of different alternatives.

However, a well-crafted operating plan for a scheduled railway can actually lead to increased train sizes.

Train size becomes a design criterion, and as long as the railway refines its operating plan as traffic patterns change, it will continue to operate large trains.

To address some of these issues, CPR turned to MultiModal Applied Systems and its MultiRail software. MultiRail was first employed by the Saint Lawrence and Hudson division of CPR in 1995 and 1996, which encompassed most of the eastern operations of the railway. This division was able to produce dramatic improvements in its costs and service levels through the careful crafting of a new operating plan using MultiRail, catching the attention of Rob Ritchie, CPR’s current CEO. Under Mr. Ritchie’s leadership, a joint team of CPR and MultiModal employees was formed in 1997 to explore the creation of a new operating strategy for CPR. While many people were involved with this effort, day-to-day technical leadership of the team was provided by John Fallis of CPR and Jason Kuehn of MultiModal. After overcoming a variety of technical and organizational issues, the team implemented a scheduled railway in late 1999. CPR calls the resulting plan the Integrated Operating Plan (IOP).

The Integrated Operating Plan

In 1997, CPR wanted to replace the tonnage business model to improve customer service, operating efficiency and effectiveness, profits and to reduce operating costs. With customers focusing on total supply-chain logistic costs, it had to provide reliable and competitive transit times. CPR found that adding operational capacity did not improve its effectiveness. It launched a number of capital renewal projects to replace the aging locomotive fleet and made selective investments in replacing infrastructure and renewing computer hardware and software. CPR needed to integrate these investments into its operating plan.

Shifting to a schedule-based model from a tonnage approach was a huge challenge for CPR, which had run for 125 years on the old model. It had to change its operations and culture, integrate its capital investments, and improve its financial performance and customer service. This required a massive paradigm shift for the operations team. The objectives included faster railcar velocity, improved locomotive utilization, reduced train starts, and improved customer service (Figure 2).

CPR’s customers want it to transport carloads, but CPR needs to move entire trainloads. For example, on an average day, of the 650 cars customers release to go to Chicago, only 45 of those cars are to move from the entire province of Alberta to Chicago proper. The railway must aggregate these low volumes of traffic in its operating plan.

A railway operating plan describes how railcars should move (the car routings and train plan) and

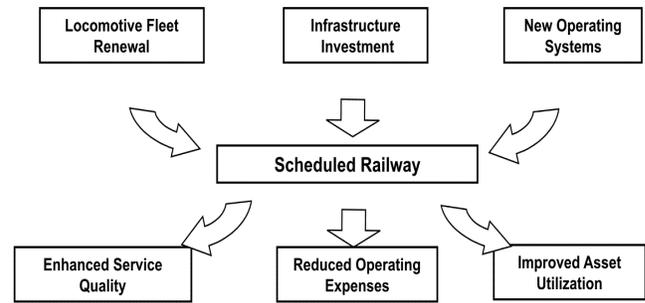


Figure 2: The planning process for the scheduled railway is the focal point for leveraging investments in physical asset, operations, and management science to improve performance in terms of costs, asset utilization, and customer service.

often includes the major assets needed to move the railcars (such as train crews, locomotives, yards, and tracks). Launched in mid-1999, the IOP was designed to improve service and to reduce the number of trains, which are often competing goals.

Fundamental to the railway operating plan is routing cars across the network, through the rail yards, and on the trains. Railways do this with blocking plans, which are made up of elements called blocks.

A *block* is a group of railcars that move together for some portion of their journeys. For example, in a simple blocking plan, a block between A and C can carry traffic destined to all other locations. But a block from C to D can deliver traffic to D only for furtherance to E or F (Figure 3). Often a car is routed on multiple blocks over the network. The blocking plan defines the set of permissible blocks to use for car routing (Table 1).

CPR builds the train plan on top of the blocking plan. The railway aggregates these blocks into trains to move as a single unit. The train designer wants to maximize train size, reduce the complexity of the blocking on the train, eliminate work at intermediate yards, calculate running times between yards, determine block connections, and minimize consumption of fuel.

How train movements are scheduled affects block-connection times between trains at CPR’s yards and, hence, transit times for customers. Spacing the train arrivals and departures at the yards and terminals affects the efficient use of yard resources.

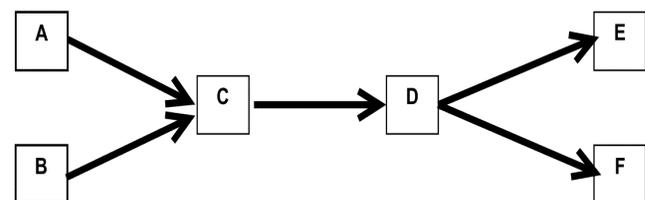


Figure 3: A blocking plan can be represented as a network, in which each link or edge represents a group of cars being moved from one yard to the next.

Blocking Location	Block Destination	Traffic Destinations
A	C	B, C, D, E, F
C	D	D, E, F
D	F	F

Table 1: This sample blocking plan shows where each block is formed, where the block will be broken up to form new blocks, and the composition by destination of traffic assigned to each block. When such plans are expanded to cover the 2,500 to 10,000 locations found on a large railroad and the special rules that apply to specific car types, customers, commodities, and other attributes, they can easily grow to one million entries or more.

Putting this together, CPR develops its IOP based on its traffic and network information and creates a feasible subset of the routing possibilities defined in the blocking plan. It then aggregates these blocks to create the train plans. Once it determines all the train plans, CPR can generate shipment trip plans. A *trip plan* specifies the specific blocks and trains required to move a shipment from origin to destination.

A group of experienced CPR service designers creates the operating plans under the leadership of the authors from CPR, with technical support from MultiModal. Input on the plan design is gathered from a variety of other groups, including both marketing and field operations. Marketing's focus is on the satisfaction of customer-service requirements, while field operations focuses on the ability to execute the plan.

By creating intelligent blocking and train plans, CPR can use its assets efficiently, minimizing crew and locomotive deadheads, routing railcars effectively, and maximizing the use of CPR's track, yards, and terminals.

Our simple blocking example illustrates the blocking concept, but practical problems are much larger. CPR has over 65,000 railcars. In any month, these railcars can take over 10,000 different potential paths, each unique origin-destination combination including a wide variety of traffic types. By refining the blocking plan, CPR gains an opportunity to improve its profitability and operations in the following ways:

(1) It can cut shipment transit times by reducing switching of railcars. Handling and holding railcars in yards often represents over 50 percent of the total transit time. By optimizing the blocking plan, CPR can reduce the number of handlings, thus reducing total transit time.

(2) It can use the time saved by reducing handlings to slow train speeds to reduce fuel consumption, while still maintaining promised transit times. CPR reduced its fuel consumption by 16 percent to 1.25 US gallons per 1,000 gross ton-miles, making it among the best in the industry despite CPR's moving much of its traffic over the Rocky Mountains.

(3) It can balance workloads among yards. By making seasonal adjustments to the blocking plan, CPR

can increase the capacity of the system by moving processing demand from yards near their railcar-processing limit to yards with available capacity.

(4) It can reduce railcar dwell time in yards by rerouting cars to build large enough departing volumes to support more than one departing train per day between processing yards. In addition to the time saved by reducing handlings, increased departure frequencies reduce waiting time in yards, further reducing overall transit times and improving reliability. CPR's railcar velocity at 160 miles per day is among the highest in the industry and has improved by 41.6 percent since 1998.

An intelligent design of the blocking plan is the foundation for producing efficient operating plans.

Routing railcars and moving trains effectively improves operational fluidity, increasing capacity within the nearly fixed plant, and reducing operating and capital costs. Through these improvements, CPR gains opportunities to increase revenue and profitability.

The problem of designing a railway operating plan is to satisfy a set of customer requirements expressed in terms of origin-destination traffic movements, using a blocking plan and a train plan. Thus, the primary variables are the blocks and trains. The constraints are the capacities of the lines and yards, the customer-service requirements, and the availability of various assets, such as crews and locomotives. The objective function in an abstract sense is to maximize profits. However, because of the complex nature of the problem, we focused on various cost metrics, such as car-miles, ton-miles, trains operated, and cars switched between blocks.

The Solution

To develop the operating plan, CPR and MultiModal decomposed the problem into a series of subproblems that are solved sequentially in five steps:

(1) Develop a traffic forecast reflecting each market segment's requirements.

(2) Use these requirements to design the blocking plan.

(3) Design trains based on the blocking plan.

(4) Use simulation to analyze yard and train workloads by the day of week and time of day.

(5) Finally, pass the train schedule on to the planning tools that develop the crew and locomotive cycle plans.

This five-step process is performed in an iterative fashion, both within each step and between steps (Figure 4). Each iteration adjusts the blocks and trains to improve the overall use of yard and train capacity and to improve the routing of the cars. Then customer-service standards are verified for compliance during the simulation step and changes made in the plan when it doesn't meet these standards.

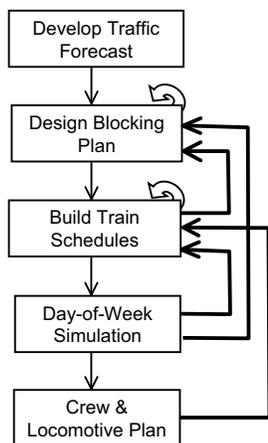


Figure 4: The process of designing the operation plan is decomposed into a series of discrete steps. Feedback loops result in iterative processes within individual steps and between steps.

In over 20 years working on computer applications for railway design, we have found that no single algorithm or model can capture the full complexity of the problem of designing railway operating plans. Our solution works by tackling the entire problem through the use of many separate algorithms within a holistic framework. We know of no other solution in ongoing use that approaches the completeness of our solution.

There are a number of papers that discuss using algorithms to create operating plans, including Assad (1980a, b), Keaton (1989, 1992), Gorman (1998a, b), and Huntley et al. (1995). There have also been survey papers that review railway optimization models (Crainic 2003, Newman et al. 2002), and there is a good Web-site that lists other literature (Kraft 2003). Most prior work focuses on solving subproblems of the overall railway-service-design problem, with few, if any, examples of holistic, integrated solutions. Furthermore, none of these prior efforts have resulted in production solutions employed on an ongoing basis within the railway industry.

Forecasting Traffic

Planning railway operations requires a detailed forecast of car volumes, tonnages, and lengths for each origin-destination pair, and the information must be specific in terms of volume by day of week, type of traffic, load or empty status, and which other railways interchange traffic with CPR.

CPR’s service-design department developed an automated forecasting system that combines last year’s traffic, last month’s traffic, and a high-level revenue forecast produced by CPR’s marketing and sales department, called REVPLAN. The forecasting system provides MultiRail direct access to detailed CPR traffic volumes reflecting both marketing’s projections and

the effects of seasonality. This data drives the entire process of designing operating plans.

Developing the Blocking Plan

The blocking plan is the foundation for the operating plan, determining the car routings, yard workloads, and contributing to customer service.

We design the blocking plan in an iterative, MultiRail-based process (Figure 5). We begin by creating an initial plan. Next, we evaluate the plan and identify potential improvements and test them. The initial plan can be either the one currently used or one algorithmically generated.

Starting with this initial plan and the traffic data, we use an algorithm to generate a block sequence for each traffic movement. We then use these sequences to estimate the expected block volumes and yard workloads and to identify possible improvements. We generally measure a plan’s quality in terms of the number of cars switched and total car-miles, subject to the capacity of the yards. Because there are many trade-offs among the improvement opportunities and many constraints we cannot capture in the computer model, a service-design expert reviews changes to the blocking plan. We repeat this process until we can identify no further major improvements.

Improvements to the blocking plan are primarily found through what we call bypass and circuitry analysis, both of which are supported through MultiRail algorithms.

A bypass is a direct block that eliminates intermediate switching. For example, if cars traveling from A to D are currently switched at C, a bypass block from A to D would eliminate this intermediate handling (Figure 6). We consider various criteria in identifying bypass blocks to ensure that they meet minimum volume requirements and to take into account any interactions with other blocks.

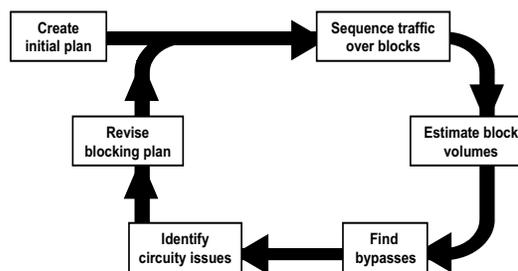


Figure 5: The process of designing the blocking plan is highly iterative once planners create the initial plan. The traffic-sequencing process drives the evaluation process, with the experts in service design acting as gatekeepers determining which changes to include in each iteration. Design criteria include reducing railcar switching, minimizing car miles, and respecting yard capacity.

For example, an A to D bypass block might make the A to C block too small to justify or it might conflict with a C to E bypass. We rely on an expert to assess such complications.

Circuitry is a measure of the difference between the shortest distance a car can travel from its origin to its destination and its actual travel distance as given by its block sequence. MultiRail identifies circuitous movements based on a number of criteria, such as the circuitry percentage and the total number of excess car-miles. In validating a plan, we use circuitry to identify missing blocks and potential improvements to the blocking plan.

The shortest physical path is a function of network factors, such as clearance and traffic type. In the bypass and circuitry analyses, we focus on adding new blocks to the plan. MultiRail also contains algorithms and reports to identify blocks to be eliminated. The full design process takes these removals into account in its iterations.

The ability to rapidly generate block sequences for every traffic movement is central to the process of designing blocking plans. A *block sequence* is a path from the origin to the destination of the traffic over a directed graph composed of the blocks in the blocking plan. Various user-controlled block attributes determine whether we can consider a particular block when finding the shortest path. The cost for each sequence represents the weighted mileage of the block sequence plus mileage-based penalties for each switching activity. We make further cost adjustments based on traffic type and other factors. We consider some constraints only during the solution process, so that we must run the shortest-path algorithm iteratively, restructuring the network between iterations to reflect violated constraints.

MultiModal's block-sequencing algorithm is critical to its effective use and to the overall planning process. To execute the iterative process for designing blocking plans, we must make rapid, large-scale changes to the blocking plan. Current industry practice is to use tables to specify which traffic goes in which block at each yard. Such tables can be huge, containing millions of entries. Making large-scale changes rapidly

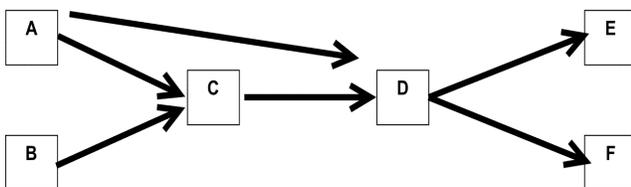


Figure 6: The service designer has a bypass opportunity from A to D. Introducing an A-to-D block would clearly change the volumes on the A-to-C and C-to-D blocks and on any other new blocks under consideration, such as a C-to-E bypass. Making these design trade-offs is the responsibility of the service designer using the software.

is impossible. The algorithmic approach used in MultiRail reduces the number of rules by two orders of magnitude and thus enables the solution strategy we employed.

For example, a yard closure based on a table-based blocking plan would require changes to tens of thousands of entries at each yard that sends blocks to the targeted yard. It would also require changes to tables for a variety of other upstream and downstream locations. In MultiRail, simply raising the cost of the yard to be closed and adding and dropping a few high-level block definitions would be sufficient to complete the yard-closure analysis.

There are other approaches to the optimization of blocking plans, such as large-scale mathematical-programming techniques (Bodin et al. 1980, Barnhart et al. 1998, 2000) and heuristic methods (Ahuja et al. 2003). The concept of a dynamic blocking plan, along with routing algorithms, is described by Kraft (2000). Kraft (2002) gives an excellent overview of the importance of yards and therefore of blocking plans.

Train Plan

The blocking plan lays the foundation for the train plan (Figure 7). Each train's schedule lists departure and arrival times, the blocks of cars it picks up or sets out at each location, crew change points, and locomotive requirements, among other details.

To develop a train plan, we use MultiRail's heuristic algorithms to identify large-volume blocks and to create trains around those blocks. The train size might be smaller than capacity, so we use MultiRail to identify other blocks that can be picked up en route until we estimate the train size is close to capacity. We iterate this process until we have assigned all blocks to at least one train.

Next, we use MultiRail to reestimate the train sizes and refine the day-of-week frequency to further improve capacity utilization. MultiRail's algorithms can accurately calculate the intermediate arrival and departure times of the trains as they travel across the network, but the planner needs to establish the original departure time for each train. Given the departure times, MultiRail employs several algorithms and reports to show the effects of the train plan on connection times and inventory of cars in the yards. The planner uses these calculations to adjust the train times and sometimes the day-of-week frequency to properly balance yard workloads.

Finally, the planner determines crew and locomotive requirements based on the train plan. These requirements are used in subsequent planning steps to develop specific deployment plans for locomotives and crews.

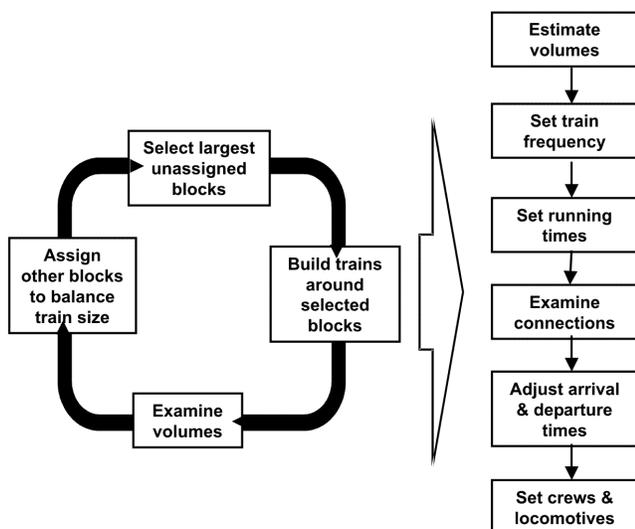


Figure 7: The train designer starts by focusing on assigning all of the blocks to trains. The designer creates trains around the largest anchor blocks first, using the remaining blocks as filler to achieve size goals. Once the designer defines basic trains, he or she must make many refinements to fix their frequencies, their departure and arrival times, their timing at yards, the crew requirements, and the locomotive needs.

What are the characteristics of a good train plan? From a high-level view, a train plan must provide frequent service to meet customers’ needs but contain a minimum of trains to reduce costs. A train should be fast to maximize track capacity and improve service, but slow to save fuel. A good train plan must not overburden yards by sending too many trains through them at once. Yet, bunching trains may reduce the connection times of cars at the yards. The train planners must resolve these somewhat contradictory design criteria. MultiRail provides rapid, interactive feedback on all of these criteria, allowing the planners to focus on perfecting the plans.

Day-of-Week Simulation

To speed the design process, we use average-day analysis in the initial block- and train-plan development work. Ultimately, we must take day-of-week and time-of-day factors into account. To do this, we use MultiRail’s SuperSim tool.

SuperSim calculates the detailed trip plan or itinerary of each origin-destination movement, including the blocks and trains used and the yards where the cars are switched. Because we use the time-of-day and day-of-week car releases, we must typically generate 500,000 to one million trip plans. This simulation can be a bottleneck, inhibiting rapid and thorough analysis.

However, in SuperSim, we use a variety of techniques to speed this process so that we can obtain a solution in a few minutes, rather than in hours or

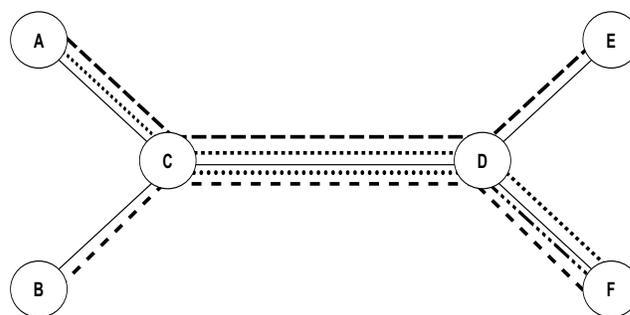


Figure 8: Each nonsolid line in the figure represents a traffic movement. In a blocking plan that assigns all traffic going from C to D to the same block and train, these separate traffic movements will share common trip-plan attributes between C and D and can be processed as a group during the simulation.

days. Outputs from SuperSim focus on yard workload and car inventory, train size, and compliance with customer-service requirements. We use these results to fine-tune the operating plan by

- smoothing workloads at yards,
- making schedule adjustments to improve car connections,
- changing the days trains operate to account for ebbs and flows in car volumes, and
- ensuring that the plan meets customer-service requirements.

How does SuperSim solve the performance problem? Conventional railway trip-planning tools compute trip plans individually. However, the natural aggregation process of building blocks and trains means that we can advance many cars from one location to the next in a single calculation. For example, we may have various flows going from A, B, and C to D, E, and F. If all of these flows travel on the same block from C to D, we can use a single processing step to advance these cars between these two locations, greatly reducing simulation run time (Figure 8).

Other Algorithms

MultiRail includes many additional algorithms and analysis techniques, including

- an interactive trip planner that allows planners to create individual what-if trip plans,
- numerous diagnostics to evaluate and identify plan defects,
- the ability to generate time-distance diagrams to examine line-capacity impacts,
- various reports on workload requirements, and
- the ability to feed the MultiRail data to a variety of CPR real-time and planning systems.

The last major step in the planning process is developing a locomotive cycle plan. MultiRail estimates the tonnage for each train, which an internal CPR

system uses to assign minimum locomotive requirements. These requirements result in an imbalanced, and therefore infeasible, locomotive cycle plan. CPR’s locomotive-planning system devises a feasible plan by deadheading locomotives on existing trains to achieve balance. The algorithm employs a time-space network covering four weeks of train events over the railway’s 250-yard network and uses a depth-first search technique to identify deadhead opportunities. Ahuja et al. (2002) and Luo and Meketon (1997) also did work in developing locomotive plans.

To execute the plan, we use an empty-car distribution model to suggest the routing of empty cars to customers for loading. Several times a day, the model solves a two-week, 250-yard, 30-car type problem to find the least-cost routing for empty cars. The model is based on work initially undertaken by Mark Turnquist (Turnquist and Jordan 1983, Turnquist 1994) of Cornell University, which CSX Transportation subsequently redesigned and reprogrammed.

Results and Conclusions

CPR’s senior managers believe that the company’s adoption of management science tools and operations research techniques has transformed CPR into a more agile, profitable, highly cost-effective, and competitive railway. To quote CPR CEO Rob Ritchie, “CPR’s operations team and its Integrated Operating Plan exceeded its objectives. Today, Canadian Pacific Railway schedules virtually everything it does under its Integrated Operating Plan. It schedules the movement of empty cars to fill customer orders and the movement of the loaded cars to their destinations. It schedules trains in all track corridors and integrates these schedules into those for the yards and terminals. It then schedules track and locomotive maintenance around the operating activities.”

The benefits of successfully implementing scheduled operations have been very significant (all financial figures are in Canadian dollars). One year after the 1999 implementation, CPR performed an audit of the benefits. This audit showed that scheduled operations reduced CPR’s cost base by \$300 million. Since the audit, CPR has analyzed two of its larger expense categories: crew wages and fuel. This analysis showed that

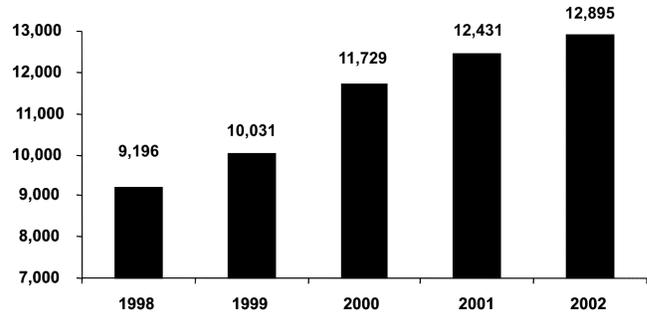


Figure 9: Labor productivity measured as gross-ton-miles per active employee has improved by 40 percent, reflecting an 18.8 percent workforce reduction and a 13.8 percent increase in gross-ton-miles since 1998.

an additional \$210-million savings was attributable to the change in operating practices in 2001 and 2002 (Table 2). Total documented cost savings through the end of 2002 have exceeded half a billion dollars. These savings do not include the benefits from reducing the number of railcars and locomotives owned over the 1999 through 2002 period.

CPR transformed the way it runs its operations and serves its customers by using the algorithms and decision-support tools of MultiRail, as well as traffic forecasting, and locomotive and empty-car planning algorithms. These tools gave CPR an opportunity to leverage new computer systems and capitalize on investments in infrastructure and locomotives.

The new strategies for routing railcars increase train weights and thus decrease train starts, enabling CPR to reduce its workforce by 18.8 percent despite an increase in gross-ton-miles (GTM) of 13.8 percent (Figure 9). These efforts have resulted in an increase in carload train size of over 10 percent. More reliable train schedules facilitate scheduling time for track maintenance and reducing variance in the system and nonproductive time. Aggressive yard bypass blocking reduces railcar processing in yards, which effectively increases yard capacity and reduces yard crew wages and yard fuel consumed.

Reduced horsepower per ton (HP/ton) ratios on trains combined with selective speed reductions enabled by increased car velocity makes the reduction

Category/Year	1998	1999	2000	2001	2002
Fuel (\$)	0	-22,732,441	-62,957,504	-74,823,239	-68,031,806
Road-crew wages (\$)	0	-13,316,997	-31,480,392	-40,564,550	-35,184,042

Table 2: Two sources of the additional Can\$200 million Canadian Pacific Railway saved in 2001 and 2002 were fuel and road-crew wages. We computed the savings using the 1998 fuel and labor-productivity rates to estimate what the railway’s costs would have been in each year had it made no changes in operations. We assumed current costs for labor and fuel and did not include wages for yard crews.

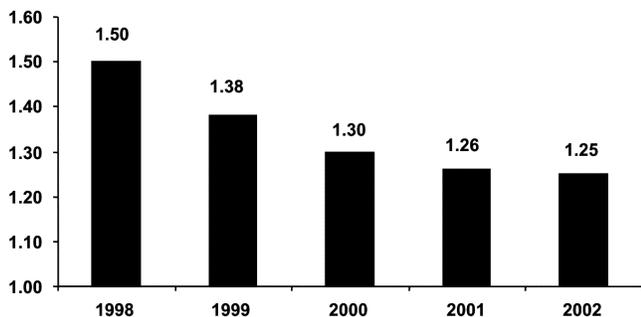


Figure 10: Canadian Pacific Railway has reduced its fuel consumption (US gallons per 1,000 gross-ton-miles) by 17 percent, and it now leads the industry.

in transit times transparent to customers. CPR has also improved fuel consumption by introducing AC powered locomotives (Figure 10).

Aggressive block bypassing and improved connections between trains at yards reduces dwell time in yards, improving railcar velocity. CPR's railcar velocity increased from 113 miles per day in 1998 to 160 miles per day in 2002 (41 percent) (Figure 11). CPR has reduced the fleet it owns or leases from 51,900 in 1998 to 44,300 in 2002 (15 percent) while GTM increased 14 percent. In addition to ownership costs, car fleet size also drives maintenance expense.

Reducing train HP/ton ratios and matching train weight to the pulling capacity of locomotives results in locomotives using their most efficient throttle position and maximum pulling capacity, thereby optimizing their utilization (Figure 12). Locomotive trip plans cycle individual locomotives between scheduled trains and are adjusted to reduce locomotive idle time. CPR plans deadhead moves to balance locomotive supply, adjusting train schedules to improve locomotive productivity.

The IOP succeeds partly because of its flexibility and agility. The plan must be able to accommodate varia-

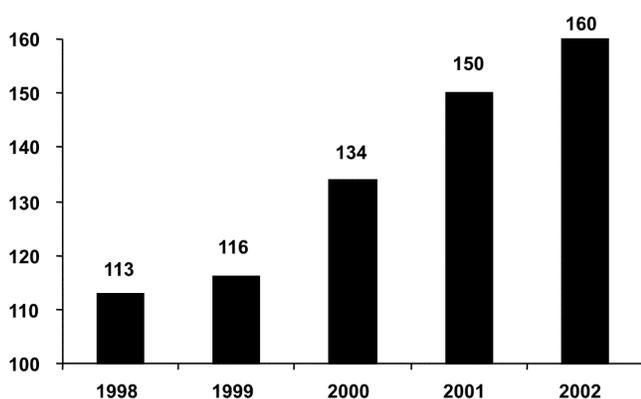


Figure 11: Canadian Pacific Railway's car velocity (car miles per car day) improved by 41 percent between 1998 and 2002.

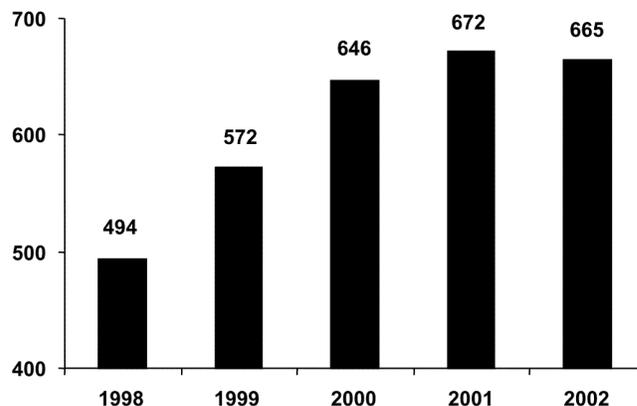


Figure 12: Canadian Pacific Railway's locomotive productivity in terms of thousands of gross-ton-miles per locomotive improved by 35 percent between 1998 and 2002.

tions in traffic levels and resource availability. The network can be affected by a variety of controllable and uncontrollable events, such as extreme weather conditions, derailments, mechanical failures, and fluctuations in freight volumes. Such events harm resource availability because they cause delays, which means that assets tend to sit in queues. This was the case during the last quarter of 2002 when CPR's grain business dropped by about 15 percent and its coal business by about seven percent.

Fortunately, a growth opportunity was developing in the containerized-freight and automotive business sectors. CPR quickly adjusted the IOP to reallocate capacity and resources to these growing markets. As a result, it reported record earnings for 2002 in a challenging North American economy. Looking forward, CPR plans to increase its industrial products—or carload—business faster than the economy. To achieve this growth, it will rely on the IOP to make the railway even more competitive with trucks than it is now.

CPR has improved the reliability of its service and its ability to shift resources quickly to meet customers' needs. It has made these gains while building an outstanding record as the safest major railway in North America for train handling. CPR has been recognized by many customers and shipping organizations for its service excellence and safe product handling, including General Motors, Sears, Shell Oil, Toyota, and Daimler Chrysler.

CPR's adoption of a scheduled strategy went against a long-standing tradition of railway operations based on the tonnage model. This major cultural change within CPR's organization continues to this day. To support this ongoing evolution, CPR is recruiting and training employees with operations research skills and exchanging employees with other railways.

The methods CPR and Multimodal developed are portable to other railways. The success of CPR's

approach to operations planning has captured the attention of railroads in the US, Mexico, Europe, and Brazil. At least two other major North American railways have begun using similar approaches and tool sets to improve their own operating plans.

The tools, techniques, and strategies employed in this effort are a work in progress. The job is never finished in an ever-changing environment. CPR continually searches for and finds areas for improvement. It is using MultiRail and the other tools we described to refine and improve its operating plan as part of an institutionalized, ongoing process. We believe our work on the problem of designing railway operating plans is an example of operations research being applied to a number of key functions in a broad business process. It is this breadth of application that is particularly noteworthy.

References

- Ahuja, Ravinda K., Krishna C. Jha, K. Jain Liu. 2003. Solving real-life railroad blocking problems. Working paper, Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL.
- Ahuja, Ravinda K., Jain Liu, J. B. Orlin, D. Sharma, L. A. Shughart. 2002. Solving real-life locomotive scheduling problems. Working paper.
- Assad, Arjang A. 1980a. Modelling of rail networks: Toward a routing/makeup model. *Transportation Res. Part B* 14(1–2) 101–114.
- Assad, Arjang A. 1980b. Models for rail transportation. *Transportation Res. Part A* 14(3) 205–220.
- Barnhart, Cynthia, H. Jin, Pamela H. Vance. 2000. Railway blocking: A network design application. *Oper. Res.* 48(2) 1–12.
- Barnhart, Cynthia, H. N. Newton, P. H. Vance. 1998. Constructing railway blocking plans to minimize handling costs. *Transportation Sci.* 32(4) 330–345.
- Bodin, Lawrence D., Bruce L. Golden, Allan D. Schuster, William Romig. 1980. A model for the blocking of trains. *Transportation Res. Part B* 14(1–2) 115–120.
- Crainic, Teodor Gabriel. 2003. A survey of optimization models for long-haul freight transportation. R. W. Hall, ed. *Long Haul Freight Transportation Handbook of Transportation Science*. Kluwer, The Netherlands.
- Gorman, Michael. 1998a. An application of genetic and tabu searches to the freight railroad operating plan. *Ann. Oper. Res.* 78 51–69.
- Gorman, Michael. 1998b. An operating plan model improves service design at Santa Fe Railway. *Interfaces* 28(4) 1–12.
- Huntley, Christopher L., Donald E. Brown, David E. Sappington, Bernard P. Markowicz. 1995. Freight routing and scheduling at CSX Transportation. *Interfaces* 25(3) 58–71.
- Keaton, Mark H. 1989. Designing optimal railway operating plans: Lagrangian relaxation and heuristic approaches. *Transportation Res. Part B* 23(6) 415–431.
- Keaton, Mark H. 1992. Designing railway operating plans: A dual adjustment method for implementing Lagrangian relaxation. *Transportation Sci.* 26(4) 263–278.
- Kraft, Edwin R. 2000. Implementation strategies for railroad dynamic freight car scheduling. *J. Transportation Res. Forum* 39(3) 119–137. Published jointly with *Transportation Quart.* (Summer 2000) 54(3).
- Kraft, Edwin R. 2002. The yard: Railwaying's hidden half. Part 1 *Trains Magazine* (June) 46–67; Part 2 *Trains Magazine* (July) 36–47.
- Kraft, Edwin R. 2003. Bibliography. Retrieved September 19, 2003 www.RailReliability.com.
- Luo, Min, Marc Meketon. 1997. A train scheduling model for reducing locomotive requirements. Paper presented at 1997 Spring INFORMS meeting, San Diego, CA.
- Newman, Alexandra M., Linda K. Nozick, Candace Arai Yano. 2002. Optimization in the rail industry. P. Pardalos, M. Resende, eds. *Handbook of Applied Optimization*. Oxford University Press, New York, 704–718.
- Turnquist, Mark A. 1994. Economies of scale and network optimization for empty car distribution. Technical report, Cornell University, Ithaca, NY.
- Turnquist, Mark A., William C. Jordan. 1983. A stochastic, dynamic network model for rail car distribution. *Transportation Sci.* 17(2) 123–145.

Robert Ritchie, President and CEO, Canadian Pacific Railway, stated the following during the presentation of this work for the Edelman Prize: “In the mid-’90s, we were struggling with relatively high costs and low profitability, all in the face of rising customer service requirements. We were not at all sure that our traditional operating strategies were up to the challenge of dealing with these issues. It was apparent to me that we needed a new game plan.

“To meet rising customer expectations and to earn the money to generate a return on the required capital investment, we needed to make a wholesale change in our operating philosophy.

“Canadian Pacific Railway turned to MultiModal and its MultiRail application. Working together, we developed what we believe is the best schedule-based model in the rail industry. We call it our Integrated Operating Plan.

“The wholesale paradigm shift to management science tools and operations research techniques has transformed Canadian Pacific Railway into a more agile, profitable, highly cost-effective and competitive railway.

“The benefits of successfully implementing scheduled operations have been huge. One short year after implementation, we performed an audit of the benefits. This audit showed scheduled operations was responsible for a \$300-million (Canadian) reduction in our cost base. Since the audit, we have analyzed two of the larger expense categories: crew wages and fuel. This analysis showed an additional \$200-million (Canadian) savings was attributable to the change in operating practices. Our total cost savings have exceeded half a billion dollars (Canadian).”