Australasian Transport Research Forum 2011 Proceedings
28 - 30 September 2011, Adelaide, Australia
Publication website: http://www.patrec.org/atrf.aspx

Energy-efficient recovery of delays in a rail network

Amie Albrecht[1], Jaap Koelewijn[2], Peter Pudney[1]

[1]Centre for Industrial and Applied Mathematics, University of South Australia, Australia

[2]Department of Applied Mathematics, University of Twente, The Netherlands

Email for correspondence: amie.albrecht@unisa.edu.au

Abstract

Fuel or energy savings of 10% or more can be achieved by providing train drivers with information on energy-efficient driving strategies. The Energymiser system uses optimal control theory to calculate a control strategy which minimises the mechanical work done by the traction system to move a train from its current location and speed to the next scheduled stop. The calculation takes into account the required arrival time at the next stop as well as train performance parameters and track gradients, curves and speed limits. The optimal speed and control profile is displayed to the driver in the cab.

When a train is delayed, it is more efficient, where possible, to recover the delay gradually over the entire remainder of the journey than to recover quickly. However, the delayed train will often interact with other trains, causing further delays to propagate through the system and thereby impacting on the energy-efficiency of many trains. When the delays are sufficiently small, the interaction locations and sequence of interactions do not change.

We demonstrate a numerical optimisation algorithm to find a set of interaction times that allows each affected train to finish on time wherever possible and that minimises total energy consumption for the whole set of trains. The new interaction times can be transmitted to Energymiser units on each of the trains, which will then calculate new optimal driving strategies to meet these times.

1. Introduction

Energymiser is an in-cab system that displays driving advice to help train drivers improve timekeeping and reduce energy consumption. The system is based on research by Howlett et al (1994, 2009) and similar work by Asnis et al (1985) and Khmelnitsky (2000). It uses optimal control theory to calculate a control strategy which minimises the mechanical work done by the traction system to move a train from its current location and speed to the next scheduled stop. The calculation takes into account the required arrival time at the stop as well as train performance parameters and track gradients, curves and speed limits. Driving advice is displayed to the driver on a graphical display. Fuel or energy savings of 10-20% have been demonstrated.

The system requires target arrival times to be specified at key locations along a journey. For a suburban passenger train, these target times include timetabled arrival times at stops and intermediate timing points at junctions. For trains on single-line rail corridors, such as those which make up most of the Australian interstate rail network, the targets are often short sections of track, called crossing loops, that allow opposing trains to pass.

When designing the schedule for an individual train, Energymiser can be used to determine energy-minimising arrival times at locations between the key locations on the route (Pudney et al, 2009). Similarly, when a train is delayed, Energymiser can recalculate efficient arrival times for intermediate locations. However, trains often interact. For example:

- For opposing trains to cross on a single-line corridor, one train must arrive at the crossing location and pull off the main line and stop to allow the other train to pass. It is important that the stopping train arrives a few minutes before the opposing train, so that the opposing train is not delayed. However, the stopping train should not arrive too early, otherwise it will have wasted energy by travelling too fast.

- Passenger trains often connect with other trains at stations, allowing passengers to transfer between trains.
- At busy junctions, trains must pass through the junction at the correct times to ensure a smooth flow of trains through the junction.

In each of these cases, a delay to one train can cause follow-on delays to other trains. In this paper we extend our previous work by showing how we can calculate (or recalculate) train interaction times so the total energy used by all of the interacting trains is minimised. We will assume that the delays are small enough that the interaction locations and the sequence of interactions do not change. Other delay costs will not be considered but can easily be incorporated.

2. Cost curves

Figure 1 shows speed profiles obtained from Energymiser for three different optimal journeys, corresponding to three different journey times. The data in all examples is from a high-speed passenger railway in the UK. In this case, the train travels from station A to station E with stops at intermediate stations B, C and D. The orange curve shows the track speed limit. The colours of each speed profile indicate driving mode: green is power, grey is coast and red is brake.

Figure 1: Speed profiles for three optimal journeys from A to E with different journey times.

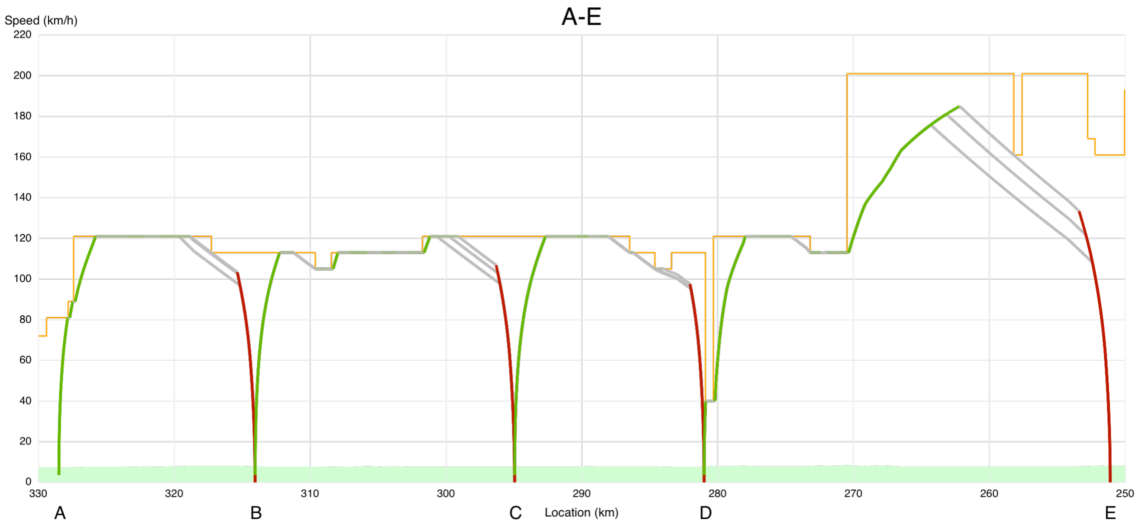

The differences between the speed profiles can be seen most clearly on sections B-C and D-E. The higher speed profiles give shorter trip durations. The duration (in seconds) and energy (in gigajoules) for each of the journey segments on each of the speed profiles are below.

	A-B	B-C	C-D	D-E	Total
high-speed duration (s)	525	681	512	902	2620
high-speed energy (GJ)	0.268	0.299	0.229	0.597	1.393
mid-speed duration (s)	525	684	514	916	2639
mid-speed energy (GJ)	0.267	0.291	0.226	0.561	1.345
low-speed duration (s)	532	689	514	937	2672
low-speed energy (GJ)	0.255	0.280	0.225	0.518	1.278

In general, the cost of an optimal journey increases as the duration decreases because the work done against aerodynamic drag increases with the square of speed.

The left diagram in Figure 2 shows the nominal path of a train as it moves from station A to station E, and then from station E to station H. The horizontal axis is time, and the vertical axis is location. Suppose the departure time from station A is fixed, but we can vary the arrival time at station E. As we increase the scheduled arrival time at E, the A-E duration increases and the cost of the optimal A-E journey decreases, as shown in the lower curve on the right of Figure 2.

Suppose the arrival time at station H is also fixed. As we increase the arrival time at station E, the available duration for the E-H trip decreases and the cost of the optimal E-H journey increases; this is shown in the middle curve on the right of Figure 2. The upper curve shows the sum of the A-E cost and the E-H cost for various arrival times at station E. The overall cost is minimised when the arrival time at station E is about 5800 seconds.

Figure 2: Train graph (left) and section costs (right) for one train.

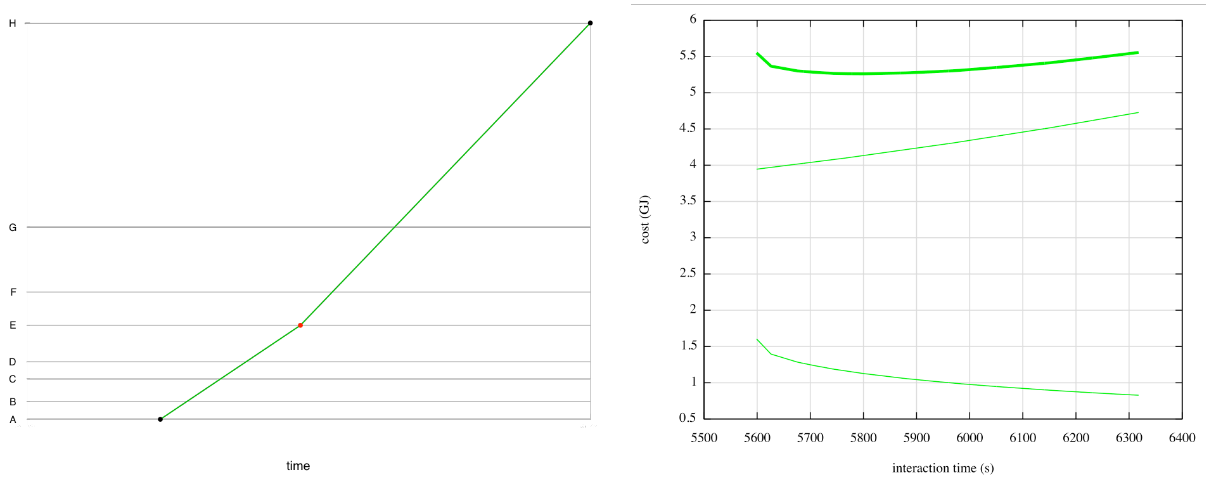

3. Two-train interactions

Now consider what happens when two trains must interact at some location, as shown in Figure 3.

Figure 3: Notation for two interacting trains.

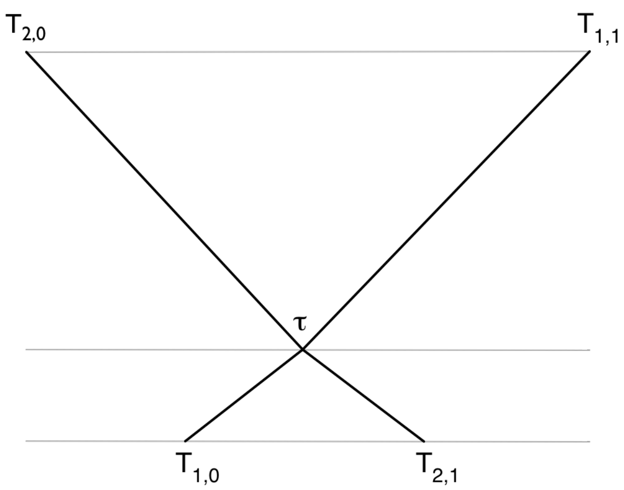

Times $T_{i,0}$ and $T_{i,1}$ for each train i are fixed. We wish to find the time τ that minimises the total cost

$$c = c_{1,1}\left(\tau - T_{1,0}\right) + c_{1,2}\left(T_{1,1} - \tau\right) + c_{2,1}\left(\tau - T_{2,0}\right) + c_{2,2}\left(T_{2,1} - \tau\right)$$

where $c_{i,j}$ is the cost function for segment j of train i.

Figure 4 shows a train graph for two trains. The planned train paths are shown in pale colours. The bolder colours indicate how the interaction time at station E should change when the green train is delayed by 500 seconds. The 'before' and 'after' cost curves are shown below the train graph:

- the thin green lines are the cost curves for the ascending train on each of its two journey sections (before and after the interaction)
- the thick green line is the total cost for the ascending train
- the thin blue lines are the cost curves for the descending train on each of its two sections
- the thick blue line is the total cost for the ascending train
- the red line is the sum of the costs for the two trains.

The delay to the first ascending train reduces the range of possible interaction times, and increases the overall energy cost.

Figure 4: Train graph and section costs ('before' on left and 'after' on right) for a two-train interaction.

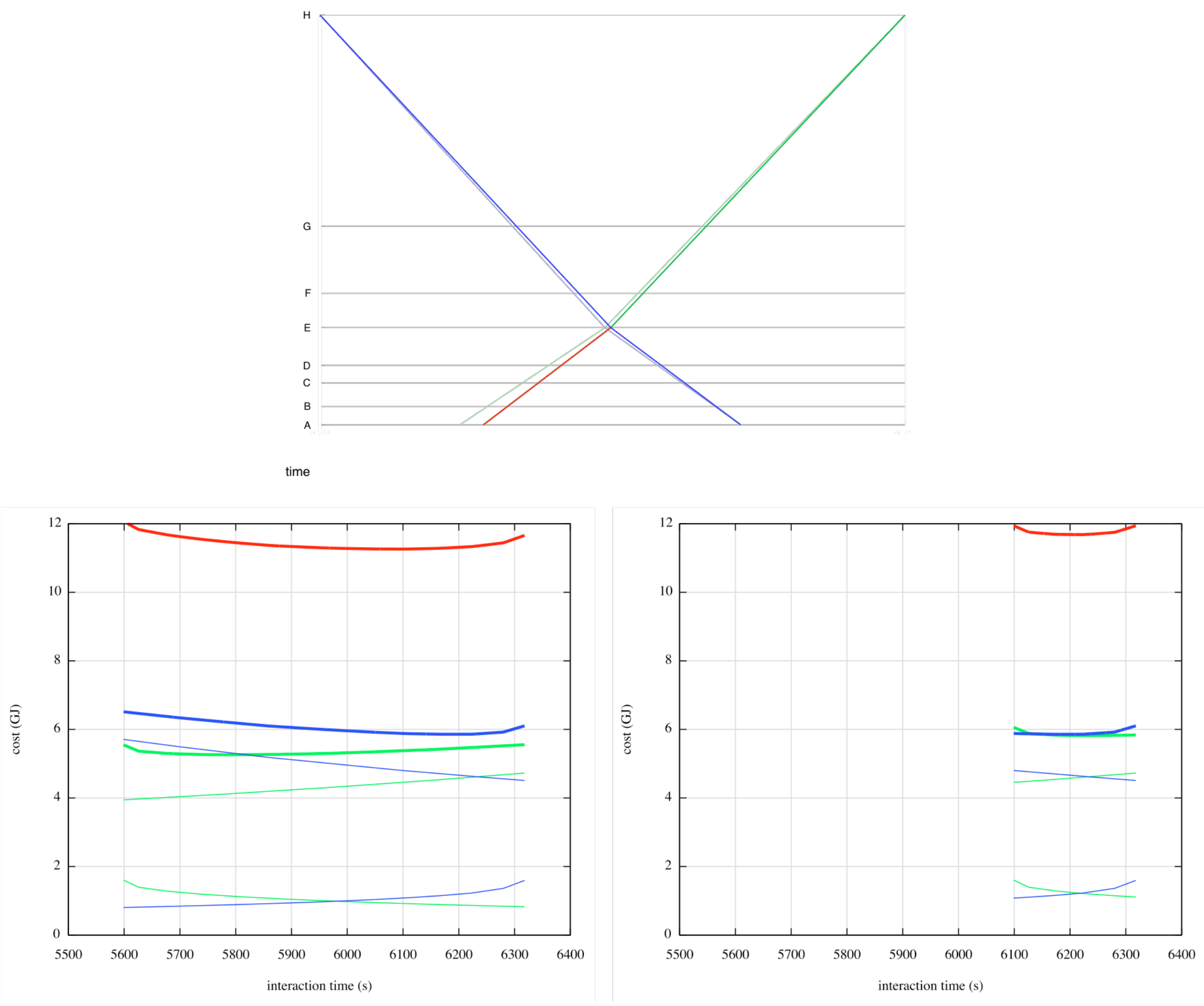

4. Many-train interactions

Figure 5 shows an example of a larger problem, with six trains and seven interactions. The desired start and finish times, indicated by black dots, are fixed. The interaction times at intermediate stops, indicated by red dots, are to be determined. In this particular example, none of these intermediate times are fixed.

Figure 5: Train graph for a six-train interaction.

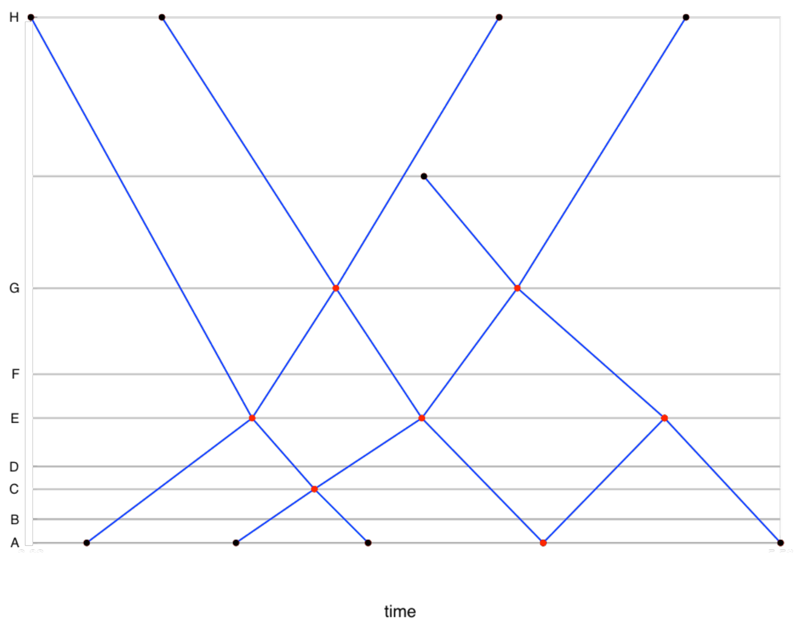

The problem of determining interaction times should be solved when the timetable is initially designed and also when any train is delayed. We can solve such problems as follows:

- Let $\boldsymbol{\tau}$ be the vector of times to be determined.
- Let $c_{i,j}(t)$ be the optimal journey cost on section j of train journey i for section duration t. In our implementation, we represent the cost of each section curve as a pre-calculated piecewise linear function of section duration, from which we interpolate.
- Construct the total cost function. For the example problem, it is

$$c(\boldsymbol{\tau}) = c_{1,1}\left(\tau_0 - T_{1,0}\right) + c_{1,2}\left(\tau_1 - \tau_0\right) + c_{1,3}\left(T_{1,1} - \tau_1\right) + \ldots .$$

We must also include constraints to ensure that the travel duration on each segment is greater than the minimum required travel time.

We can now find $\boldsymbol{\tau}$ which minimises $c(\boldsymbol{\tau})$ by using a numerical optimisation procedure. Because we represent costs as piecewise linear functions of duration, we are restricted to methods that do not require derivatives to be specified. However, each of the cost functions is convex, and so the total cost is also convex.

We used several local search numerical methods to find an optimal solution, including:

- COBYLA (Constrained Optimisation BY Linear Approximation) (Powell, 1994, 1998)
- BOBYQA (Bound Optimisation BY Quadratic Approximation) (Powell, 2009)
- Nelder-Mead (Nelder and Mead, 1965)
- Sbplex, based on Subplex (Rowan, 1990).

Minimum section duration constraints were incorporated into the optimisation using augmented Lagrangian methods (Conn et al, 1991; Birgin and Martínez, 2008). We used the NLopt library (Johnson, nd) to implement the code. COBYLA was the fastest method on all test cases; the time required to solve the six-train problem was less than ten seconds.

In practice, trains may be delayed so that remaining fixed times on their journey are no longer feasible. In this case, we must also adjust the fixed times to make the problem feasible.

In our examples, each interaction requires the two interacting trains to arrive at the same time. Incorporating interactions where one train must arrive a fixed time before the other is straightforward.

5. Conclusions

We have demonstrated methods for formulating and solving the problem of finding interaction times that minimise the sum of the energy costs for a set of interacting trains. Other delay costs could also be incorporated. These methods can be used to calculate, in real time, energy-efficient recovery plans when trains are delayed.

In this problem, the locations of the interactions and the sequence of interactions for each train is fixed. Further work is required to extend these methods to situations where interaction locations and sequences may change.

Acknowledgements

This research was supported under the Australian Research Council's *Linkage Projects* funding scheme (project number LP110100136) and by industry partner TTG Transportation Technology Pty Ltd.

References

Asnis, I A, Dmitruk, A V, and Osmolovskii, N P (1985) Solution of the problem of the energetically optimal control of the motion of a train by the maximum principle, *USSR Computational Mathematics and Mathematical Physics 25*(6), 37-44.

Birgin, E G, and Martínez, J M (2008) Improving ultimate convergence of an augmented Lagrangian method. *Optimization Methods and Software 23*(2), 177-195.

Conn, A R, Gould, N I M, and Toint, P L (1991) A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM Journal of Numerical Analysis 28*(2), 545-572.

Howlett, P G, Milroy, I P and Pudney P J (1994) Energy-efficient train control, *Control Engineering Practice 2*(2),193-200.

Howlett, P, Pudney, P and Vu, X (2009) Local energy minimization in optimal train control, *Automatica 45*(11), 2692-2698.

Steven G. Johnson. The NLopt nonlinear-optimization package. http://ab-initio.mit.edu/nlopt .

Khmelnitsky, E (2000) On an optimal control problem of train operation, *IEEE Transactions on Automatic Control 45*(7), 127-1266.

Nelder, J A, and Mead, R (1965) A simplex method for function minimization. *The Computer Journal 7*, 308-313.

Powell, M J D (1994) A direct search optimization method that models the objective and constraint functions by linear interpolation, pp 51-67 of *Advances in Optimization and Numerical Analysis*, Dordrecht:Kluwer Academic Publishers

Powell, M J D (1998) Direct search algorithms for optimization calculations. *Acta Numerica 7* 287-336.

Powell, M J D (2009) *The BOBYQA algorithm for bound constrained optimization without derivatives*. Technical Report NA2009/06, Department of Applied Mathematics and Theoretical Physics, Cambridge, England.

Pudney, P, Stolz, C and Campbell, P (2009) Using energy-efficient journeys to calculate section running times, AusRailPlus 2009.

Rowan, T (1990) *Functional Stability Analysis of Numerical Algorithms*. PhD thesis, Department of Computer Sciences, University of Texas.