
Article:
Goodwin, J.C.J, Fletcher, D.I. and Harrison, R.F. (2015) Multi-train trajectory optimisation to maximise rail network energy efficiency under travel-time constraints. Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit. ISSN 0954-4097

https://doi.org/10.1177/0954409715593304

Multi-train trajectory optimisation to maximise rail network energy efficiency under travel-time constraints

Jonathan CJ Goodwin*[1], David I Fletcher[1], Robert F Harrison[2]

[1] Department of Mechanical Engineering, University of Sheffield, UK

[2] Department of Automatic Control & Systems Engineering, University of Sheffield, UK

Corresponding author: Jonathan Goodwin, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield. S1 3JD, UK.

Email: j.goodwin@sheffield.ac.uk

Abstract

Optimising the trajectories of multiple interacting trains to maximise energy efficiency is a difficult but highly desirable problem to solve. A bespoke genetic algorithm (GA) has been developed for the multi-train trajectory optimisation problem and used to seek a near optimal set of control point distances for multiple trains**,** such that a weighted sum of the time and energy objectives is minimised. Genetic operators tailored to the problem are developed including a new mutation operation and the insertion and deletion pairs of control points during the reproduction process. Compared to published results, the new GA was shown to increase the quality of solutions found by an average of 27.6% and increase consistency by a factor of 28. This allows more precise control over the relative priority given to achieving time targets or increasing energy efficiency.

Keywords

Multi-train trajectory optimisation, trajectory planning, train control, energy efficiency, railway network optimisation.

Introduction

Background

The long term increasing cost of energy globally, coupled with concern over CO_2 emissions, means that minimising energy consumption is becoming increasingly desirable for all industries, not least the transport sector which accounted for 27% of Global [1] and 39% of UK [2] energy consumption in 2011. While different parts of the rail industry may have different primary concerns, minimising operational energy consumption is an increasingly pressing problem for all. However, in general, rail is already a relatively efficient transport mode accounting for 8.7% of passenger and 9.0% of freight traffic in the UK, whilst constituting only 1.9% of its transport sector energy consumption in 2011 [3]. This means it is possible to reduce overall energy consumption by modal shift to rail instead of less efficient transport modes such as road and air. Given the projected increases in transport demand, maximising network capacity is also increasing in importance, both economically and environmentally.

Operational methods for minimising traction energy consumption and maximising network capacity, while maintaining competitively short journey times, are often preferable to upgrades in network infrastructure and/or rolling stock. Physical improvements usually require large capital investment and/or only improve performance in a very specific way. In contrast, operational improvements (for example timetabling, rescheduling, train control) can be easier and less expensive to introduce and have the potential to affect several different performance measures.

Work presented here focuses on trajectory optimisation of multiple trains in a network, with the aim of improving overall network punctuality and energy consumption. To do this, operational interaction between trains must be considered.

Multi-train optimisation

McClanachan observed that,

> "If the journey time of one train is extended to save energy, then this could adversely influence the schedules and energy usage of other trains on the same network." [4]

Somewhat surprisingly then, there has been comparatively little work on the problem of multi-train trajectory optimisation, compared to the single-train problem. This is probably due the greatly increased complexity of the combinatorial problem. For a single-train, analytical methods [5] [6] have shown that the optimal trajectory will consist of a combination of only five operational modes: maximum traction, speed holding (using traction or braking), coasting, and maximum braking. There are also a number of works which consider optimisation of train trajectories for more than one train, but without integration of the optimisation between trains. Since the trajectories are not optimised simultaneously this is essentially an extension of single train optimisation, with additional headway constraints placed on the following train, and will be unlikely to optimise the performance of the network overall [7] [8]. To date, most multi-train trajectory optimisation work has focused on applying heuristic optimisation methods, particularly genetic algorithms (GAs,) to find good enough solutions in a reasonable time. GAs are a type of heuristic optimisation pioneered by Holland [9], with Chang and Sim [10] widely cited as the first to apply it to train trajectory optimisation. GAs are based on evolution by natural selection – where a population of different chromosomes (ie. solutions) compete against each other, with the genetic information from the fittest chromosomes more likely to be passed to the next generation. In 2004, Albrecht [11] used a two

level optimisation to minimise the total energy consumption and power peaks of a network. At a low level, the trajectories of individual trains were optimised independently, while at a high level, train movements were synchronised using a GA. In 2007, Miyatake [12] extended the model in [11] to include exchange of energy between the trains and improved optimisation of train trajectories, and later compared the performance of different energy storage devices in [13].

In 2008, Acikbas [14] used a novel approach to the multi-train optimisation of a small network. SimuX software [15], capable of modelling a multi-train system with overhead line voltages, was used to train an artificial neural network of the modelled system. This allowed solutions to be evaluated ~900 times faster than using simulation, with an error of less than 3%, making optimisation by GA feasible. However, trajectories were controlled using only two variables, the coasting and the re-motoring velocity, making the optimisation fairly simplistic.

In 2012, Yang [16] described a GA based optimisation for the multi-train trajectory problem on a branched network. The GA was similar to that proposed by Chang [10], but adapted to work on a network. Each solution defined switching points, between traction and coasting pairs, for all trains on the network. Simulation was then used to estimate the total energy, time and delays (caused by enforcing headway constraints) of the system as a whole.

In 2014, Wang [17] solved the two train problem, where one train is following another, using both mixed integer linear programming (MILP) and pseudospectral methods. Pseudospectral methods were found to give slightly better results than MILP but took two orders of magnitude longer to calculate. The optimisation was also carried out using the greedy (lead train trajectory optimised independently of the second train) and the simultaneous approach. As expected, the simultaneous approach gave slightly better results but took longer to calculate. However, since the number of constraints scaled linearly with the number of train trajectories being optimised it was noted that “the computation time of the bigger [multi-train] problem will be much longer”.

Methods

G1: Model and optimisation

In the current paper the rail network in Figure 1 has been modelled as a finite graph; nodes representing stations, and edges representing bidirectional single-track railway links. Each link has a length, over which a speed limit profile is defined, whereas the nodes have no modelled properties. This network (N1) was previously investigated by Yang et. al. [16] using the modelling methodology and optimisation hereafter referred to as G1. Formulation G1 was implemented in C++, and validated against the published results where these were available. This was chosen as the starting point for this investigation as G1 makes fewer assumptions about the form of train trajectories compared to the older models. It also seemed more readily extendible than Wang's model [17] as it could already consider any number of trains, more complex network structures and operational interactions between trains other than the restriction in headways of following trains.

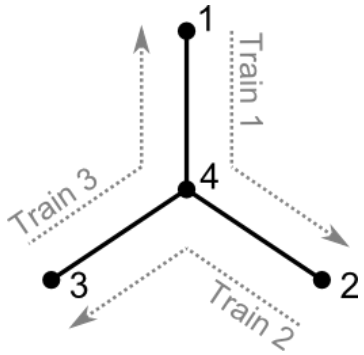

Figure 1. Illustration of network N1, the topology and train journeys of which were previously defined and investigated by Yang et. al. **[16]**. Unless stated otherwise, all optimisation investigated in this paper were applied to this network.

In G1, train motion on each link is defined as alternating sections of maximum traction and coasting, controlled by position vector $\underline{\boldsymbol{x}}$, with application of the maximum braking operation interrupting the final coasting section at distance y to ensure stopping at the end of the link (Figure 2).

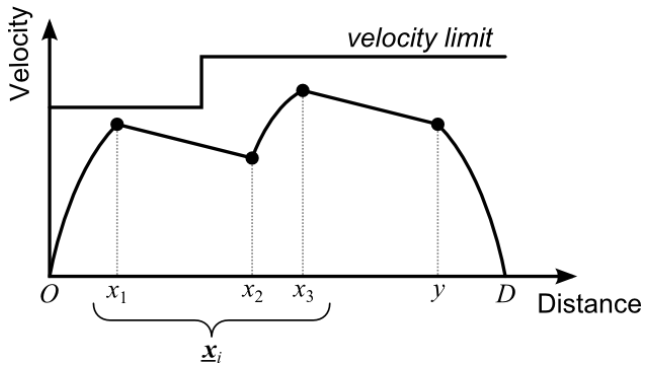

Figure 2. Traction and coasting operation are applied alternately as each link is traversed. The position of the control points determines the trajectory each train follows between origin O and destination D.

Notation:

D	**link length**, the distance between two stations.
x_n	**control point**, given as a distance from the start of the link ($0 \le x < D$).
y	**braking point**, where maximal braking must be applied to come to a stop at distance D.
$\underline{\boldsymbol{x}}_i$	**link control strategy**; made up of a list of control points. $\underline{\boldsymbol{x}}_i = (x_1, x_2 \ldots x_n \ldots x_{n\_max})$ where n_max is an odd positive integer. (Note: $x_0 = 0$, $x_{n\_max+1} = D$)
$\underline{\boldsymbol{X}}$	**network control strategy**; made up of a list of link control strategies. $\underline{\boldsymbol{X}} = (\underline{\boldsymbol{x}}_1, \underline{\boldsymbol{x}}_2 \ldots \underline{\boldsymbol{x}}_i \ldots \underline{\boldsymbol{x}}_{i\_max})$ where i_max is the total number of links traversed by all trains.
$E(\underline{\boldsymbol{X}})$	**total energy consumption** of a particular network control strategy.
$T(\underline{\boldsymbol{X}})$	**total traverse time** of a particular network control strategy.
$G(\underline{\boldsymbol{X}})$	**penalty** for operational interactions between different trains caused by a particular network control strategy.

Train movements defined by each $\underline{\boldsymbol{X}}$ are simulated by implementing Newton's laws of motion using a piece-wise linear approximation (Δt = 1 s). Links are traversed in the order defined by $\underline{\boldsymbol{X}}$ and constraints imposed during simulation to ensure: feasible solutions, safe operation, ride comfort and sufficient time for operations at stations. As well as checking the feasibility of each solution, the simulation allows an objective function to be evaluated for each $\underline{\boldsymbol{X}}$. Target values for the total traverse time and total energy consumption are defined, as $\bar{T}$ and $\bar{E}$ respectively, and the deviation from these targets then formulated into a single equation Eq (1) using a linear weighted sum method.

$$F_{\alpha}(\underline{X}) = \alpha \cdot \max\left\{\frac{E(\underline{X})-\bar{E}}{\bar{E}}, 0\right\} + (1-\alpha) \cdot max\left\{\frac{T(\underline{X})-\bar{T}}{\bar{T}}, 0\right\} \quad \textbf{(1)}$$

where the weighting factor, $\alpha \in [0, 1]$, allows a different relative importance to be placed on energy consumption or traverse time.

Since Eq (1) only considers energy and time spent traversing links, a penalty accounting for operational interactions in stations, $G(\underline{\boldsymbol{X}})$, is added to Eq (1) to give an overall objective score for each network control strategy. $G(\underline{\boldsymbol{X}})$ can be customised for different situations, but here is defined as the sum of departure delays, weighted by the relative priorities of different trains.

$$\text{Objective score} = F_{\alpha}(\underline{\boldsymbol{X}}) + G(\underline{\boldsymbol{X}}) \quad \textbf{(2)}$$

A genetic algorithm (Figure 3) is used to minimise Eq (2), by searching for near optimal $\underline{\boldsymbol{X}}$. Constraint checking is integrated into genetic operators to ensure that any offspring, resulting from the breeding of parent chromosomes, is a feasible solution. The overall process is represented in Figure 3, where the loop will keep iterating new populations of solutions (expected to increase in fitness) until the end condition is reached. For G1 a fixed number of generations is defined, after which the best solution found is accepted.

For consistency, the same parameters as used in Yang's best optimisation are used throughout this paper (Table 1).

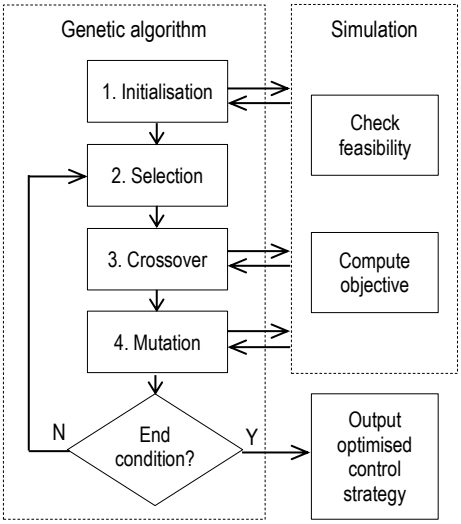

Figure 3. The structure of the optimisation algorithm used in G1.

Table 1. The GA parameters used in this paper (from [16], table 4, Experiment 9)

P_c (probability of crossover)	0.6
P_m (probability of mutation)	0.8
population size	40
number of generations	800
α (weighting between energy and time)	0.3
M (initial mutation distance) *	100 m
Tr (minimum distance between operational transitions) *	500 m
parameter used in roulette wheel selection (also referred to as α) *	0.2

* personal communication [18]

G2: Link-wise mutation operation

The mutation operation proposed by Yang has the advantage that it tends towards the previous solution, which is known to be feasible. However, this places extra constraints on the optimisation process; in this case requiring the same mutation size of all control points on the network. Below, a modified mutation operation is proposed which finds separate feasible mutation sizes for each link independently. This requires the ability to alternately apply a genetic operator to, and then check the feasibility of, the control strategy for each link in the network. A genetic operator that is applied in this way will be called a "link-wise" genetic operator and will be applied using Procedure 1. A mutation operation adapted to work as a link-wise operator is proposed in Procedure 2. Together these procedures allow link-wise mutation to be performed on a population. It is intended that this should place fewer constraints on the optimisation process, thereby allowing better local optimisation.

Procedure 1: Alternating a genetic operation and feasibility checking.

Step 1 For each chromosome (in any order)

Step 2 if P_h < rand [0,1] then go to step 10

Step 3 For each link control strategy (in the order) defined by $\underline{\boldsymbol{X}}$

Step 4 Apply link-wise genetic operator ($\underline{\boldsymbol{x}}'' = h(\underline{\boldsymbol{x}}')$)

Step 5 If $\underline{\boldsymbol{x}}''$ is feasible then go to step 8

Step 6 If $\underline{\boldsymbol{x}}'$ is feasible then $\underline{\boldsymbol{x}}'' = \underline{\boldsymbol{x}}'$ and go to step 8

Step 7 Else, go to step 10

Step 8 Next link

Step 9 $\underline{\boldsymbol{X}}''$ replaces $\underline{\boldsymbol{X}}'$ in the population

Step 10 Next chromosome

where P_h is the probability of applying the link-wise operator $h(\underline{\boldsymbol{x}})$

Procedure 2: Single link mutation

Step 1 Predetermine an initial distance of mutation $M > 0$, let $m = M$

Step 2 Randomly give a mutation vector $\underline{\boldsymbol{d}}$ with the same length as $\underline{\boldsymbol{x}}'$

Step 3 Let $\underline{\boldsymbol{x}}'' = \underline{\boldsymbol{x}}' + m\underline{\boldsymbol{d}}$

Step 4 Correct $\underline{\boldsymbol{x}}''$ to the feasible form (using the procedure in [16])

Step 5 Check validity of $\underline{\boldsymbol{x}}''$ using simulation

Step 6 If $\underline{\boldsymbol{x}}''$ is feasible then end procedure, else let $m \leftarrow m/2$

Step 7 If $m >$ [a small positive distance] then go to Step 3, else end procedure

where $\underline{\boldsymbol{d}}$ is a vector with elements randomly defined as +1 or -1

The mutation operation in G1 was replaced with the link-wise mutation operation (defined in Procedure 1 and Procedure 2) to make optimisation G2.

Unlike G1, mutation in G2 does not guarantee that a feasible network control strategy will be produced. This is the same situation as already existed for the crossover operation. In the case where neither the mutated ($\underline{\boldsymbol{x}}''$) or pre-mutation ($\underline{\boldsymbol{x}}'$) link control sequences are feasible, Procedure 1 will reach step 7 and the current chromosome will not be mutated. However, the improvement in optimisation performance discussed later suggests that, in the system studied, the potential for this event to occur is outweighed by the benefit of having a less constrained genetic operator.

G3: Insertion and deletion operations

As well as having good local optimisation, the other main problem that must be overcome in complex optimisation problems is how to avoid getting stuck in local minima. GAs seek to do this by having

diversity within a population and also the potential to reintroduce lost diversity using mutation. However, as will be discussed in the results section, neither optimisation with the original mutation operation (G1) nor the proposed link-wise mutation operation (G2) appears to be successful in avoiding local minima. In particular, solutions with two distinct patterns of control strategies were observed: those with the second traction operation before the drop in line speed limit, and those with it after. These are illustrated in Figure 4 as A and B respectively. If the population has converged, and only contains one of these control strategy patterns, then the other can only be re-introduced using mutation. However, since the distance of reduced line speed limit (3 km) is large compared to the mutation size (≤100 m), many generations of poorly scoring intermediate strategies make rediscovery of an A-like solution from a population of B-like solutions unlikely (and vice versa). If control points are excluded from a region of the line then, by definition, the mode of train control in this region cannot be changed, which may lead to a suboptimal solution. In this case, solution A fails to exploit the rise in line speed from 20,000 m onwards. Conversely, too many control points in a region may lead to a restricted control strategy, as a minimum distance between operation transitions must be maintained, again leading to suboptimal solutions.

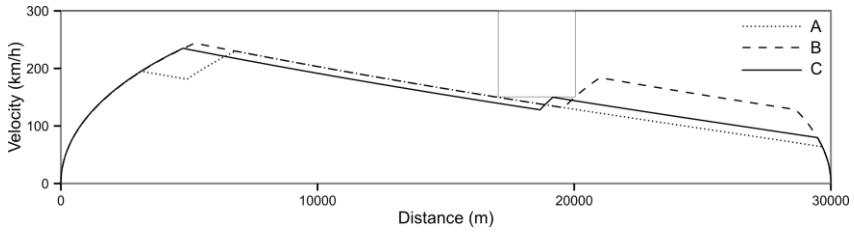

Figure 4. Train trajectories generated by optimisation G2. The lines A and B illustrate proposed local minima observed in the results. Using mutation these must interconvert by passing through some unfavourable intermediate similar to the one illustrated by line C.

It is probable that increasing the population size would cause diversity resulting in a reduced likelihood of getting stuck in local minima, but this would also greatly increase the computational burden from

simulation. In biology there are three classes of single nucleotide mutation: point mutation, insertion, and deletion. Both the original mutation procedure and the link-wise mutation used in G2 are analogous to a DNA point mutation in biology, as one control point is modified, but the total number of control points remains the same. For this reason procedures are proposed for the probabilistic insertion and deletion of pairs of control points (see Figure 5). Chang [10] used similar operations, duplication and deletion, but it is believed the operations proposed here are more effective because:

- The probability of insertions and deletions is biased towards locations where they are most likely to be needed.
- As much as possible, the effect of the operations on the "downstream" trajectory is minimised, decreasing the probability of producing infeasible solutions

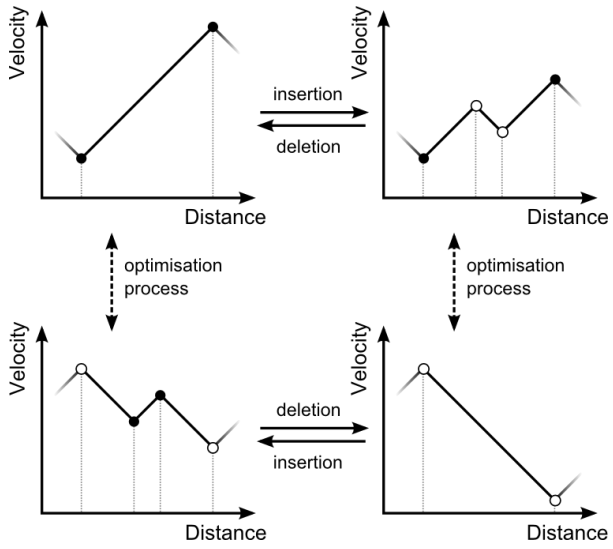

Figure 5. Extracts from train trajectories illustrating how they are affected by the insertion and deletion operations (for simplicity modification of the neighbouring control points has not been shown here). During the optimisation process control points may be moved, by mutation and crossover, extending or contracting the distance for which the traction or coasting operation is applied.

Procedures to enable both of these are detailed below, capturing the following logic. It is proposed that the probability of insertion between two adjacent control points is proportional to the distance between them. This will bias insertion towards areas of the control sequence currently lacking control points. The total probability of insertion or deletion happening on each link was implemented as $P_{ins\_link} = 0.25$ and $P_{del\_link} = 0.25$ respectively (these probabilities were tuned 'by hand' and found to be large enough to give sufficient exploration, but small enough not to impede convergence). Using the notation introduced earlier the probability of inserting a pair of control points between control point n and n+1 is given by:

$$P_{ins\_pair}(x_n) = P_{ins\_link} * \frac{(x_{n+1} - x_n)}{D} \quad \textbf{(3)}$$

where $0 \leq n \leq n\_max$

Similarly, the probability of deleting a pair of control points should be proportional to their 'shortness' to bias for removal of potentially redundant genetic material. The probability of deleting the pair of control points n and n+1 is given by:

$$P_{del\_pair}(x_n) = P_{del\_link} * \frac{\left(1 - \frac{(x_{n+1} - x_n)}{(x_{n\_max} - x_1)}\right)}{(n\_max - 2)} \quad \textbf{(4)}$$

where $1 \leq n \leq (n\_max - 1)$

As can be seen in Figure 5, the insertion or deletion of control point pairs causes downstream changes to the velocity profile of the train. To limit this, and so maximise the chance of insertion or deletion resulting in a feasible solution, two strategies are proposed. The first is to minimise the distance between the inserted pair of control points (ie. $\Delta d = Tr$, the minimum distance between operational transitions).

The second is to move position of neighbouring control points in order to conserve the total distance that each control operation is applied for.

As with link-wise mutation, the insertion and deletion procedures were applied probabilistically to the population using Procedure 1 (step 4), with a probability of P_i = 0.6 and P_d = 0.6 respectively (again, these were tuned 'by hand' in combination with P_{ins_link} and P_{del_link}).

Procedure 3: Link-wise insertion (valid for $n\_max \geq 1$)

Step 1 Let $n = 0$

Step 2 If $P_{ins\_pair}(x_n) <$ rand[0,1] then go to step 14

Step 3 If $(x_{n+1} - x_n) < 2*Tr$ then go to step 11

Step 4 if 0.5 < rand[0,1] then go to step 8

Step 5 if n=0 then go to step 9

Step 6 if $(x_n - x_{n-1}) < 2*Tr$ then go to step 11

Step 7 $x_n \leftarrow x_n - Tr$, go to step 11

Step 8 if $n = n\_max$ then go to step 6

Step 9 if $(x_{n+2} - x_{n+1}) < 2*Tr$ then go to step 11

Step 10 $x_{n+1} \leftarrow x_{n+1} + Tr$

Step 11 if $(x_{n+1} - x_n) < 3*Tr$ then go to 14

Step 12 Let $d = x_n + Tr + ((x_{n+1} - x_n) - 3*Tr)$*rand[0,1]

Step 13 Insert new control points into $\underline{\boldsymbol{x}}$ at position d and $d+Tr$

Step 14 $n \leftarrow n + 1$

Step 15 If $n \leq n\_max$ go to step 2

Procedure 4: Link-wise deletion (valid for $n\_max \geq 3$)

Step 1 Let $n = 1$

Step 2 If $P_{del\_pair}(x_n) <$ rand[0,1] then go to step 10

Step 3 let $d = x_{n+1} - x_n$

Step 4 if $0.5 <$ rand[0,1] then go to step 7

Step 5 if $n = 1$ then go to step 8

Step 6 $x_{n-1} \leftarrow x_{n-1} + d$, go to step 9

Step 7 if $n = (n\_max - 1)$ then go to step 6

Step 8 $x_{n+2} \leftarrow x_{n+2} - d$

Step 9 remove control points x_n and x_{n+1}

Step 10 $n \leftarrow n + 1$

Step 11 If $n \leq (n\_max - 1)$ then go to step 2

G3 was implemented by adding the insertion in deletion operations to G1.

Table 2. Summary of the major innovations of each optimisation procedure defined in this paper.

Optimisation	Innovation
G1	Implementation of the model and GA optimisation described by Yang et. al. in [16]
G2	Introduces a new (link-wise) mutation operation to replace the original mutation operation of G1.
G3	Introduces the new genetic operations of insertion and deletion alongside the original GA optimisation of G1.
G4	Combines the innovations of G2 and G3.

Method of traction energy calculation

On closer inspection of the algorithms in [16] it was found that the traction energy consumption was calculated using the resultant force acting on each train (Δwork = resultant_force*Δdistance) using a piecewise linear approximation. This formulation meant that increased resistance forces at high speed caused a reduction in resultant force and therefore a reduction in the energy use of trains. To enable like-for-like comparison with previously published results, the initial investigation into the performance of optimisations G1 to G4 was performed without changing the method of energy calculation. However, the more realistic formulation of calculating energy using (Δwork = traction_force*Δdistance) was adopted for all subsequent investigations.

Results and discussion

Comparing final optimisation results of G1 to G4

For each of the formulations described above, one hundred independent optimisations were carried out to assess the effectiveness and consistency of the optimisation process. Initialisation of one hundred populations was also performed, without any further optimisation, and the best solution from each population recorded. Comparison of these results is given in Table 3 and Figure 6 followed by a detailed analysis of each individual optimisation in the following sections.

Ideally an optimisation would consistently find the solution that has the lowest score (ie. the global optimal solution); so the smaller the spread in scores, and the lower the scores found, the better the optimisation. However, since the objective score associated with the globally optimal solution is not

known for this system, the performance of each optimisation is quantified relative to the performance of G1 using Eq (5) and Eq (6).

$$\%\ \text{improvement in mean score achieved by GX} = 100 * (S_{GX} - S_{G1}) / (S_{G1} - S_{G0}) \quad (5)$$

$$\sigma\ \text{ratio (fractional improvement in consistency) achieved by GX} = \sigma_{G1} / \sigma_{GX} \quad (6)$$

where GX is any optimisation (G1 to G4), S_{GX} the mean score after optimisation with GX, and σ_{GX} the standard deviation in objective scores after optimisation with GX

Table 3. The result of optimisation using G1 to G4 (each assessed using a sample of 100 independent optimisations).

Optimisation	Objective score		Improvement compared to G1	
	mean	σ	mean /%	σ ratio
None (random initialisation only)	0.1740	0.0349	-100.0*	0.6
G1	0.0550	0.0195	0.0*	1.0*
G2	0.0264	0.0126	24.0	1.6
G3	0.0172	0.0020	31.8	9.7
G4	0.0131	0.0007	35.2	29.5

*by definition.

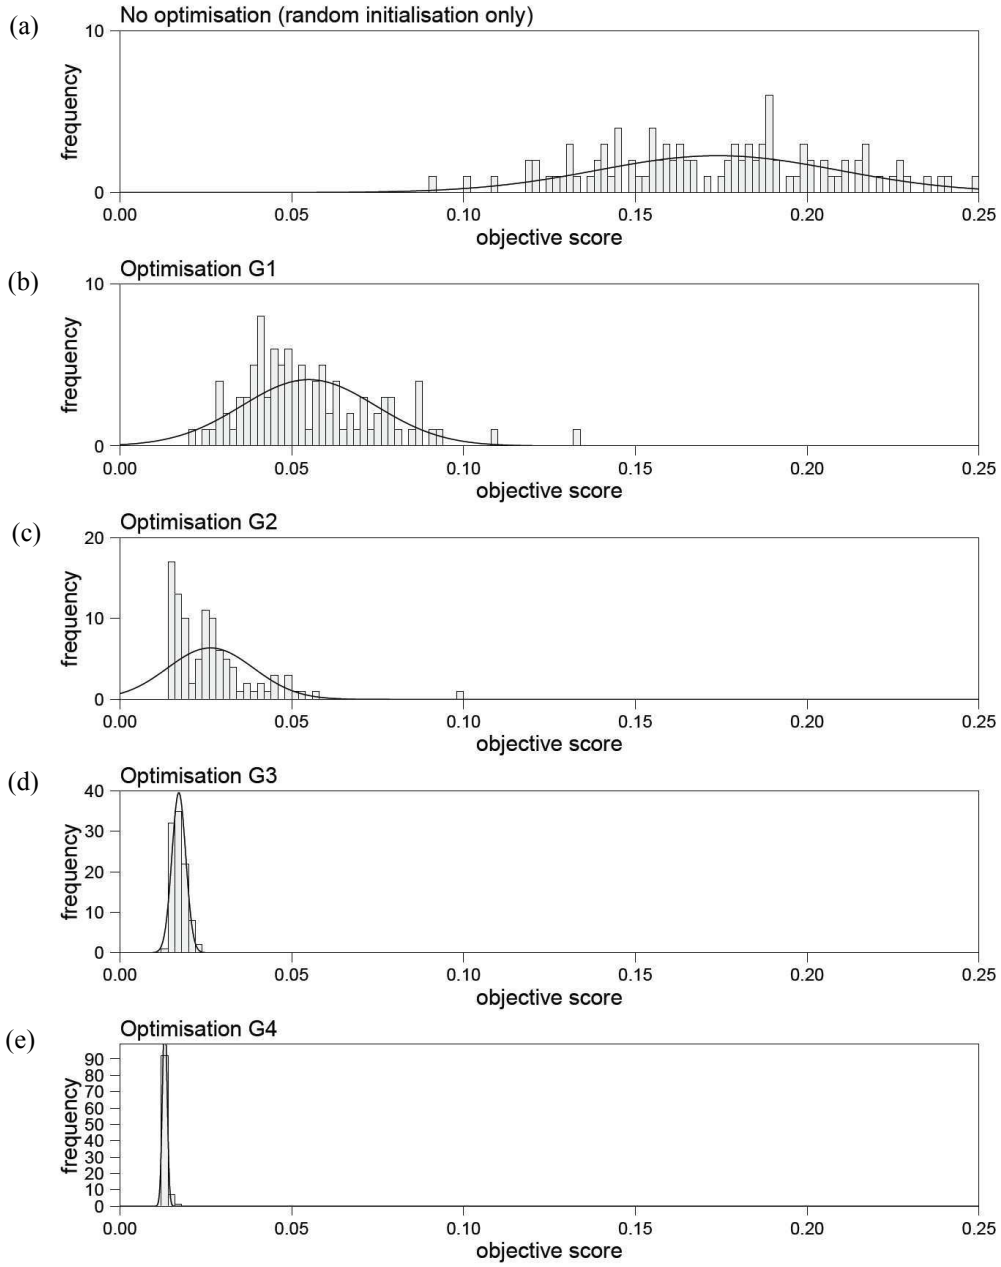

Figure 6. Histograms comparing the distribution of results from different optimisation techniques (lower scores are better). The improvement in optimisation performance from (a) to (e) can be seen by the monotonic decrease in the mean and standard deviation in scores achieved. Normal distribution curves

are shown for clarity, although strictly only the data in (a) is normally distributed having a (Shapiro-Wilk) p-value > 0.05 [19]. The significance of the multimodal distribution observed in (c) is discussed below.

Optimisation using G1

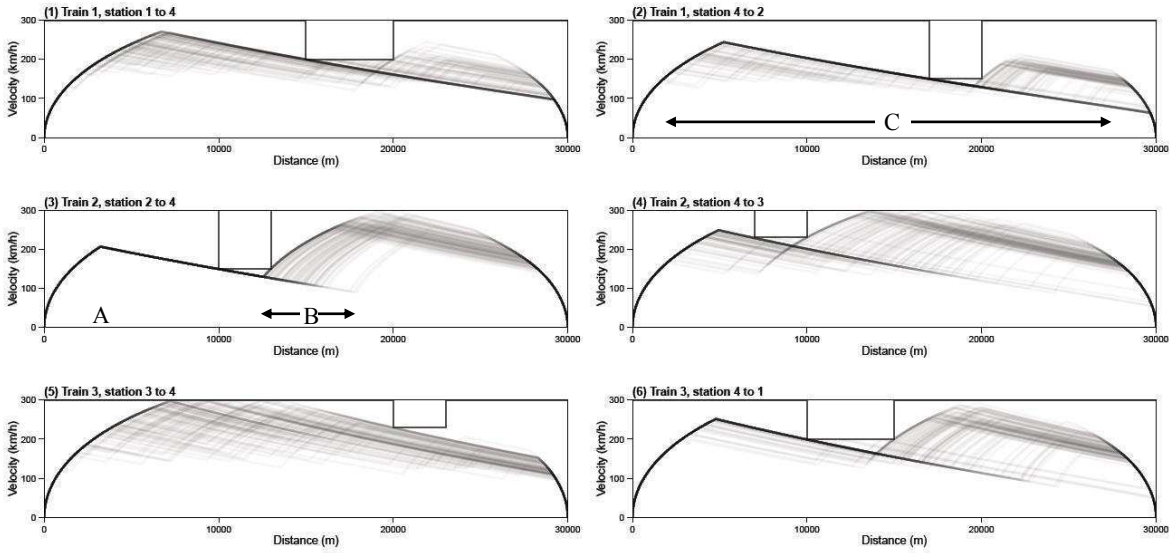

Figure 7. The consistency of train trajectories found using G1 to optimise network N1 (100 independently optimised trajectories overlaid). The positions of control points are labelled to illustrate: A – strong consensus, B – large local variation, C – near global variation.

It can be seen by comparing Figure 6a and Figure 6b that G1 is effective in optimising the system described by Yang [16]. However, after optimisation there is still a large variation in the objective score of results, caused by the trajectories of the optimised results that are illustrated in Figure 7. The trajectories show that in some places there is good consensus in the position of control points found (eg. point A in part 3 of the figure), whereas in other places (eg. points B and C) large variations are clear. Large variation within a single, uninterrupted region of the search space is consistent with either poor

local optimisation or lack of selection pressure where there is no significant change in objective score between different solutions. However, the large variation in objective scores seen in Figure 6b suggests the latter is unlikely. Also, as will be seen for optimisation with G2, if local optimisation is improved then C separates into two local minima. Both of these issues are addressed by the innovations introduced in optimisation G2 and G3 respectively.

Optimisation using G2

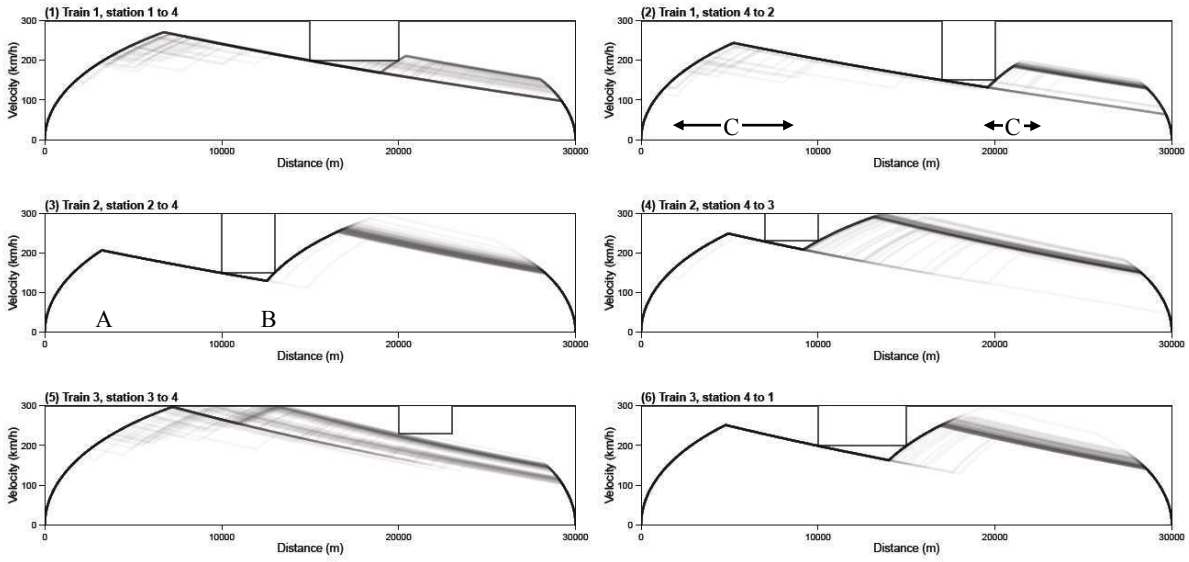

Figure 8. The consistency of train trajectories found using G2 to optimise network N1 (100 independently optimised trajectories overlaid). There is a strong consensus in the positions of control points A and B, but the two distinct locations of C suggest at least two different local minima are present in optimised solutions.

The optimised profiles in Figure 6c have lower objective score values than in Figure 6a or Figure 6b (i.e. better), but no longer appear to be normally distributed and instead a clustering of the results is observed.

This suggests that G2 is finding local minima in the search space and is consistent with improved local optimisation. Both these inferences are supported by analysing the trajectories underlying the distribution of scores, shown in Figure 8.The improvement in local optimisation can be seen for most control points, specifically, the variation in positions found for control point B is much less than in Figure 7. Also, solutions place control point C (the position of the second traction application) in one of two well separated locations. These two types of solution are not easily interconverted using the original mutation alone, so if one is lost from the population the search may become confined to a local minimum (see Figure 4).

Optimisation using G3

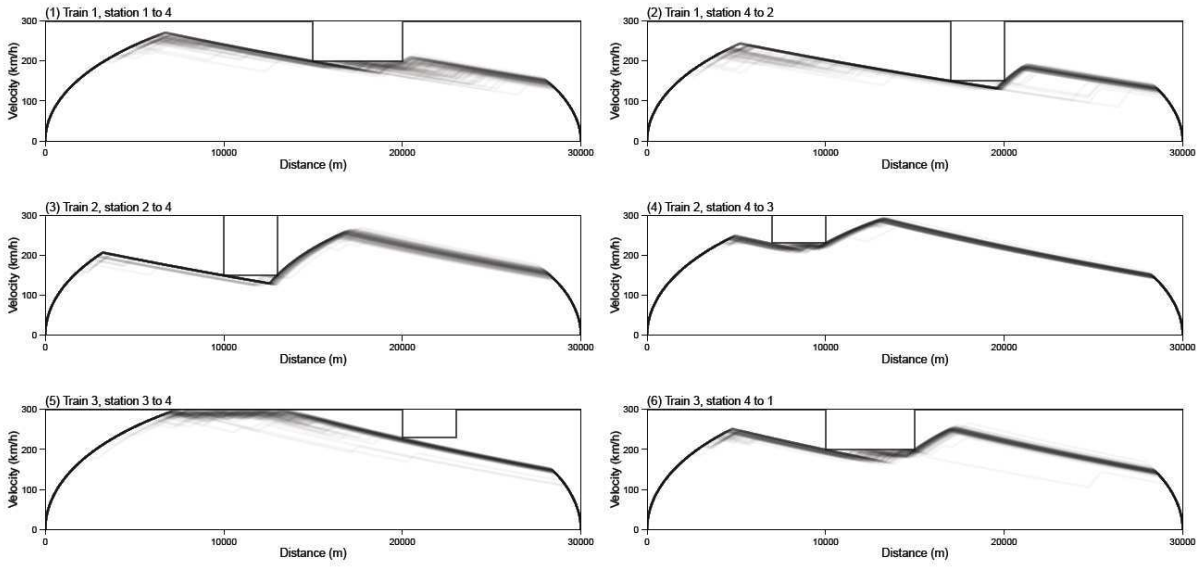

Figure 9. The consistency of train trajectories found using G3 to optimise network N1 (100 independently optimised trajectories overlaid). A clear consensus is seen, though burring of some trajectories suggests there is a small amount of local variation.

Optimisation G3 was specifically developed to address occurrence of local minima in the optimised solutions, highlighted in the results of optimisation G2. It is clear from Figure 9 that this has been successful and that trajectories of solutions found by G3 have a much clearer consensus. Figure 6d also shows that the objective scores resulting from these trajectories have a smaller variance and better average. It is particularly interesting to note that the optimised trajectory of train 3 (station 3 to 4) in Figure 9 now appears to approximate to the optimal profile we expect for a train on flat track: maximum traction, speed holding, coasting, and maximum braking [20]. However, a slight blurring of some trajectories in Figure 9 compared to the equivalent positions in Figure 8 suggests that G2 achieved slightly better local optimisation than G3.

Optimisation using G4

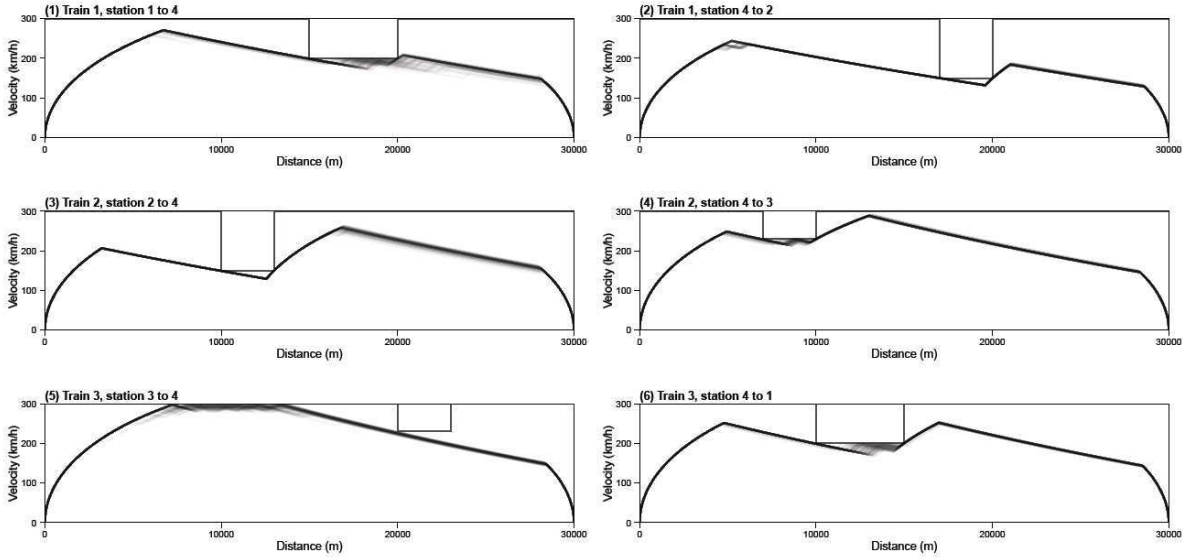

Figure 10. The consistency of train trajectories found using G4 to optimise network N1 (100 independently optimised trajectories overlaid). A clear consensus is observed along with minimal local variation in trajectories.

Optimisation G4 combines the innovations of G2 and G3 allowing it to find solutions with both a clear consensus and very little local variation in trajectories (Figure 10). Figure 6e also shows the improved optimisation performance and consistency. Together these give us much greater confidence that each optimisation using G4 will find a "near optimal" network solution.

Optimisation dynamics

As well as different final results the optimisations, G1 to G4, also displayed different dynamics during the optimisation process. Figure 11 shows that after 800 generations there was still widespread variation amongst the G1 runs, while G4 runs appeared consistently to converge after about 200 generations.

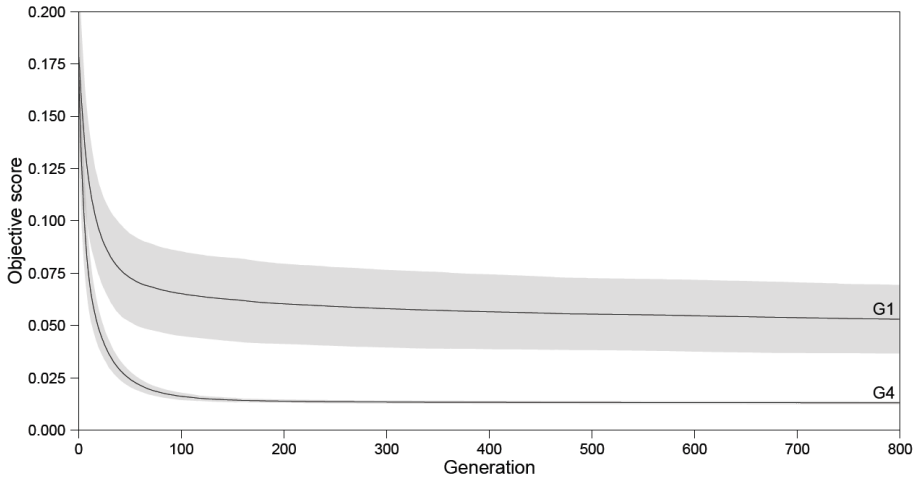

Figure 11. Average genetic algorithm progress from one hundred optimisations using G1 and G4. The grey areas show the one standard deviation about the mean objective score levels.

Trade-off between energy consumption and traverse time

When scoring each network control strategy, $\underline{\boldsymbol{X}}$, both G1 and G4 use Eq (1) to determine the contribution of energy and time. There is a region of the search space, $E(\underline{\boldsymbol{X}}) \geq \bar{E}$ and $T(\underline{\boldsymbol{X}}) \geq \bar{T}$, where $\underline{\boldsymbol{X}}$ does not meet either the energy or the time target. We expect most solutions to be in this region since, in general, going faster uses more energy and there is no improvement in score once the targets have been achieved. In this region Eq (1) reduces to:

$$F_\alpha(\underline{\boldsymbol{X}}) = \alpha \cdot \left(\frac{E(\underline{\boldsymbol{X}}) - \bar{E}}{\bar{E}}\right) + (1-\alpha) \cdot \left(\frac{T(\underline{\boldsymbol{X}}) - \bar{T}}{\bar{T}}\right) \qquad \textbf{(7)}$$

Eq (7) can then be rearranged to give a linear relation between $T(\underline{\boldsymbol{X}})$ and $E(\underline{\boldsymbol{X}})$:

$$T(\underline{\boldsymbol{X}}) = m \cdot E(\underline{\boldsymbol{X}}) + c$$

$$\text{where } m = -\frac{\alpha\ \bar{T}}{(1-\alpha)\ \bar{E}}, \text{and } c = \frac{\bar{T}\ (F_\alpha(X) + 1)}{(1-\alpha)} \qquad \textbf{(8)}$$

This defines a line of constant $F_\alpha(\underline{\boldsymbol{X}})$, a tangent to the Pareto front, along which the combinations of energy and time are equvalent in the cost function. For the above investigations using G1 to G4, α = 0.3, $\bar{E}$ = 4800 kWh, and $\bar{T}$ = 3840 s, so the gradient of this line is, m = –0.3429 (this will vary with the parameters chosen). The intercept c is dependent on the level of optimisation. In solutions from G4 the penalty for delays $D(\underline{\boldsymbol{X}})$ is usually very small (mean = 0.0003, s.d. = 0.0005), so we can assume that the $penalty\ function \sim F_\alpha(\underline{\boldsymbol{X}})$. The lowest G4 score of 0.0131 gives an intercept, $c \sim 5558$. This line of best score is shown on Figure 12, along with the energy and times of solutions obtained using different methods.

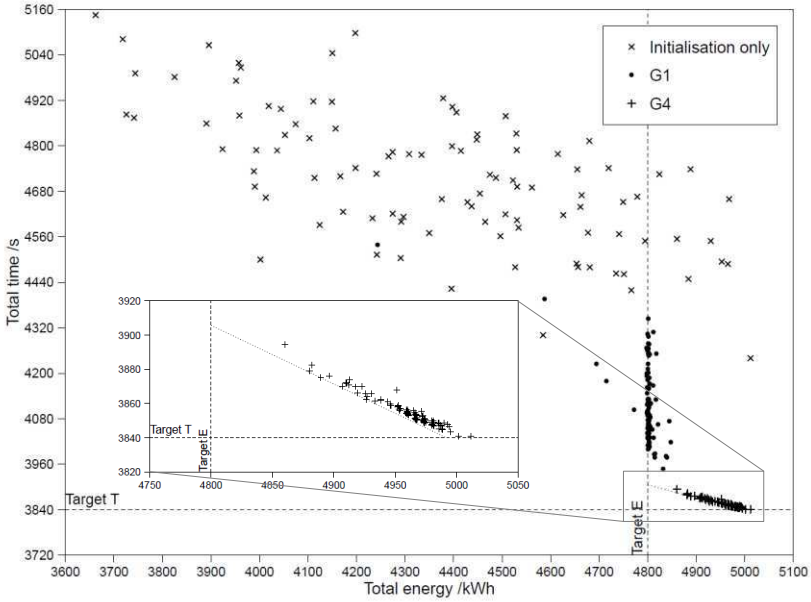

Figure 12. The total energy and traverse times of the network solutions obtained by: random initialisation only, optimisation using G1 and G4 (100 solutions of each). Eq (8) is used to find the line of constant $\boldsymbol{F_\alpha(\underline{X})}$ for the best scoring solution found by G4 (dotted line). It can be seen from the expanded area that G4 solutions vary in energy and time, but all have very similar objective scores.

It is clear from Figure 12 that both optimisations lead to better solutions when compared to the randomly generated initial solutions. However G1 solutions appear to be clustered around the target energy limit but with a large variation in total time, leading to a large variation in score . In contrast, all the G4 solutions are located close to the line of constant $F_\alpha(\underline{\boldsymbol{X}})$, again suggesting that it is a much better and more consistent optimisation. It can also be seen that some solutions found by G4 meet the target time, while others are much closer to meeting the target energy. It seems likely that the trajectories found using lower α, and therefore placing a higher importance on target time, would not be significantly different from the solutions found with α = 0.3 and that increasing α may also have little effect. For this reason, before investigating the effects of varying α, a new method of traction energy calculation is introduced. Not only

is this method based on a more realistic formulation but by increasing the energy consumption at high speeds it also increases the difference between solutions that can achieve the target energy consumption and target traverse time.

Revised method of traction energy calculation

From this point onwards the formulations of G1 to G4 have all been amended to use the more realistic method of traction energy calculation described in the methods section. With this improved formulation optimisation G4 now yields trajectories which appear to exhibit an approximation to speed holding at around 200 km/h. A future version of the model may include this mode of operation directly.

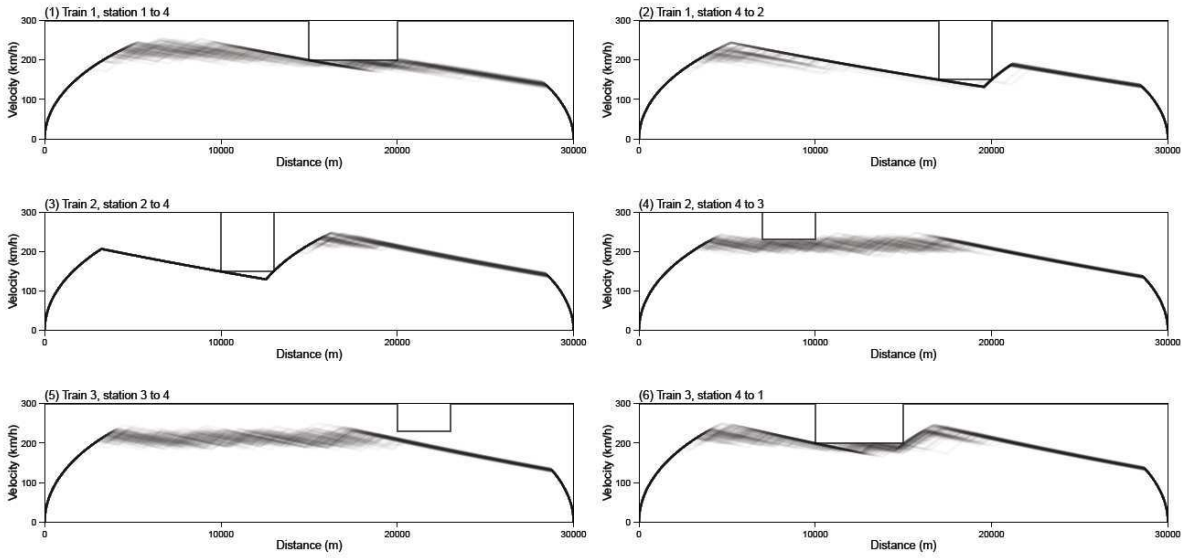

Figure 13. The consistency of train trajectories found using the new formulation of G4 to optimise network N1 (100 independently optimised trajectories overlaid).

Effect of varying α

The weighting parameter $\alpha \in [0, 1]$ in Eq (1) can be varied. A low value of α means the optimisations will prioritise meeting the time target, whereas a high α will prioritise meeting the energy target. By varying α used in the scoring of optimisation G4 (Figure 14) we can see that the optimised objective scores appear to be proportional to α below $\alpha = 0.2$ (low α), and also above $\alpha = 0.4$ (high α). This is consistent with the total time and total energy of solutions being near constant in this region, which Figure 15, showing the output of multiple repeated simulation runs, confirms to be the case.

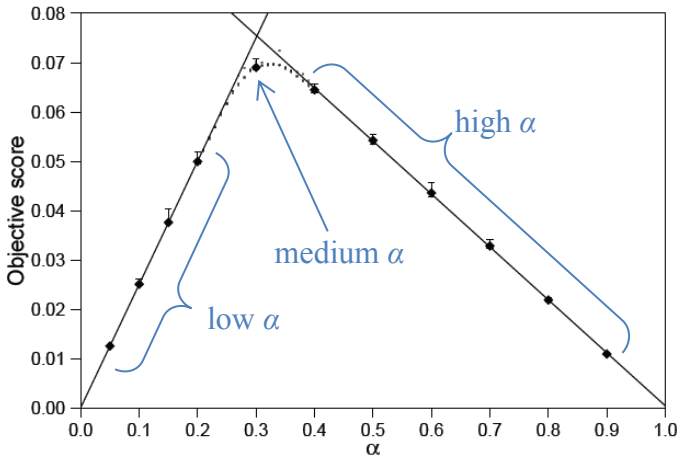

Figure 14. The effect of varying α on the optimal objectives. Dark points are the average of 100 optimisations and have max-min error bars (hardly visible). Two lines are linear regression lines through points at low α (0.05 to 0.2) and high α (0.4 to 0.9). The light grey points (appearing similar to a chain or dotted line) are single optimisations giving a higher resolution at medium α values ($0.2 < \alpha < 0.4$).

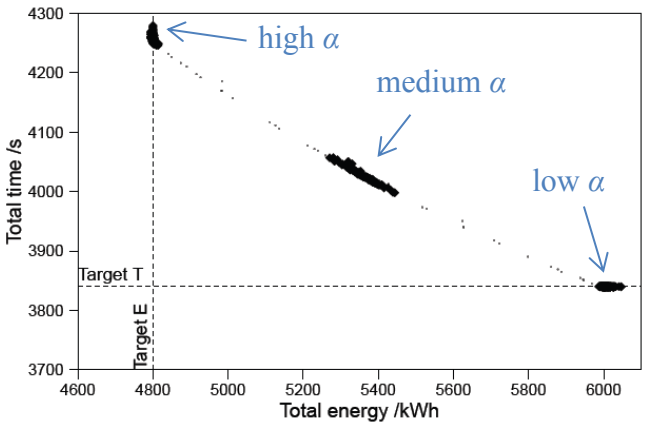

Figure 15. Pareto front of total traverse time against total energy consumption. The dark points are shown for consistency with Figure 14, and all come from sets of 100 repeats.

Figure 15 appears to show a Pareto front similar to those typically found [21] when comparing run times vs energy consumptions of single-train optimisation results. Furthermore, clusters of extreme solutions of min-time and min-energy, as described in [22], are found for low α and high α, respectively. This suggests that the optimisation is effective, though the small number of intermediate solutions means the components of the objective function respond like step functions with regard to variation in α. Plotting the results in an alternative form this can be seen in Figure 16.

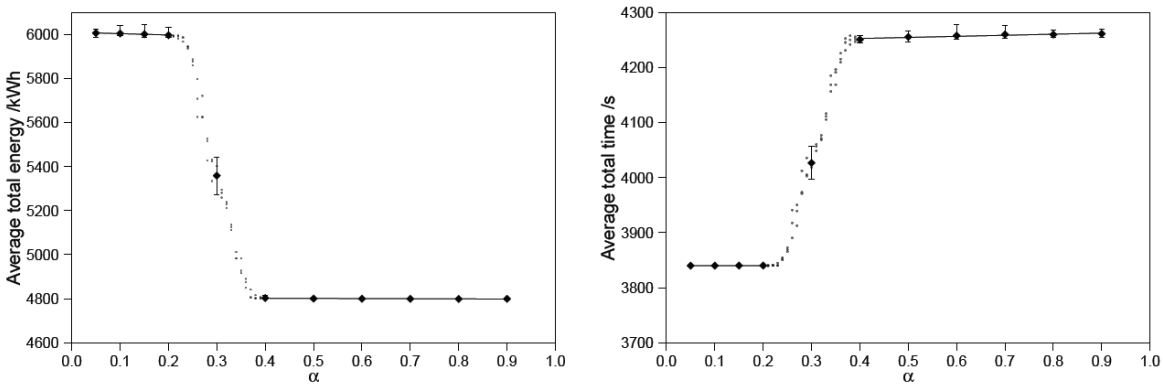

Figure 16. Varying the objective weighting, α, causes a step-like response in the optimised solutions found.

The step behaviour requires further investigation under a broader range of conditions, but could be a very useful property in the context of railway operation. While optimising for either shortest travel time, or least energy usage, it would be difficult to timetable trains subject to a continuous range of travel times on a single route. Much easier to manage would be a distinct division into 'fast' trains, and 'energy saver' trains, with a broad range of optimised driving styles producing one behaviour or the other, i.e. the outcome is resilient to real world application of the optimised strategy. This concept of resilience of optimised strategies is being explored in further research.

Effect of train schedule

A thorough investigation into the scalability of the proposed optimisations is a topic for further research. However, it is important to investigate the characteristics of the optimisations with respect to different train timetables to ensure the results described so far can be generalised and are not just artefacts of the specific timetable defined for network N1. In order to investigate this four new networks were defined – each based on network N1 but with changes affecting the scheduling of trains (Figure 17). The result of applying optimisations G1 and G4 to each of these networks is given in Table 4.

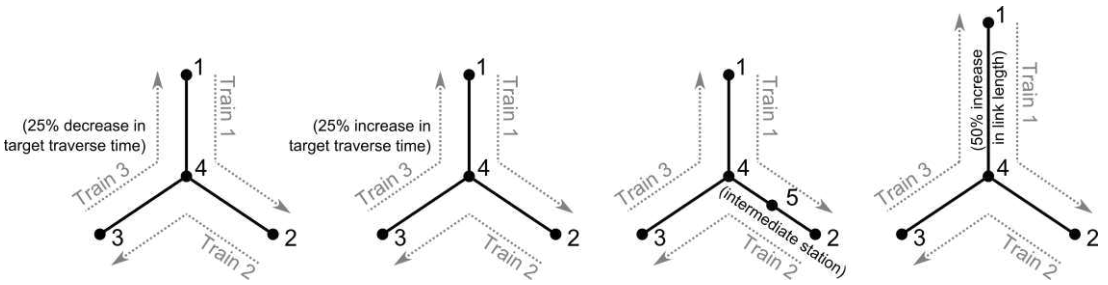

Figure 17. Four networks based on network N1 (see Figure 1). The associated timetables and energy targets are the same as N1 except for the following changes: (N2) 25% decrease in both target traverse

times of train 3, (N3) 25% increase in both target traverse times of train 3, (N4) trains 1 and 2 must dwell at station 5 for at least 30 s and 20 s respectively, (N5) target energy and traverse times are increased by 50% for both the journeys that traverse the longer link between stations 1 and 4.

Table 4. The results from applying optimisation G1 and G4 to networks N1 to N5 (100 independent optimisations were carried out for each combination of optimisation and network).

Network	Objective scores after optimisation						Improvement of G4 compared to G1	
	None (random initialisation)		G1		G4			
	mean	σ	mean	σ	mean	σ	mean /%	σ ratio (σ_{G1}/σ_{G4})
N1	0.185	0.026	0.091	0.011	0.0690	0.0004	23.5	25
N2	0.258	0.032	0.160	0.013	0.1331	0.0005	28.0	28
N3	0.164	0.034	0.065	0.010	0.0457	0.0003	19.4	34
N4	0.244	0.030	0.157	0.010	0.1287	0.0004	32.2	27
N5	0.262	0.025	0.166	0.015	0.1318	0.0006	35.1	26
						Average	**27.6**	**28**

It can be seen from Table 4 that even when the optimisations are applied to networks with different timetables, the overall pattern of improvements (first observed in Table 3) still hold true – G4 finds better scoring solutions than G1 (by an average of 27.6%) and also does so much more consistently (by an average factor of 28). The smallest improvement in mean optimised score, of 19.4%, is observed for network N3. However, rather than suggesting degraded performance of G4 it is thought this may be caused by a chance improvement in the performance of G1 (due to the increased proximity of well optimised solutions to initialisation – see Figure 12). Comparing the relative scores of different networks to N1 we see that the more challenging targets of N2 are consistent with its higher mean score, whereas

N3 has more relaxed targets and resulting in a lower score. The situation for N4 is slightly more complex, with two obvious factors likely contributing to its increased mean score: the additional stop/starts increases energy consumption and the extra dwells have potential to cause knock-on delays at station 4. While it is difficult to pick out either as the dominant cause of increased mean score in N4, the energy and traverse time targets for each journey in N5 are equivalent to those in N1 (when normalised by the distance being travelled). So, considering all train journeys in isolation we would expect similar optimised scores. However, when optimised considering interactions between trains, the mean score of N5 is significantly higher than that of N1. This suggests that the root cause of the increase in score is from interactions between different trains on the network – in this case the delay of train 3 at station 4 as it waits for train 1 to clear the longer link. The significant effect of interactions between trains when evaluating a timetable highlights the fact that multi-train trajectory optimisation is closely linked to the field of schedule optimisation, particularly if energy consumption is considered, as in [23].

Conclusions

Several improvements have been proposed and demonstrated to advance the capability of the multi-train trajectory optimisation originally proposed by Yang *et al.* [16]. Two new genetic operators, tailored to the problem formulation, were developed: a less constraining mutation operation and a procedure to insert and delete pairs of control points. Together, these improvements were shown to optimise an average of 27.6% further than published results when compared to randomly initialised solutions. This was achieved in combination with increased consistency (1/28th of the standard deviation in objective score of solutions), and faster GA convergence (less than 1/4 the number of generations). The resulting optimised trajectories now appear consistent with those expected by optimal control theory.

The improved optimisation consistency allowed a more detailed investigation of the effect of varying α, the weighting between different objectives in the cost function, to be conducted. For the system studied, the components of the objective function respond like step functions with regard to variation in α, causing the optimal objective solutions to switch rapidly between the extreme solutions of minimum time and minimum energy. It is thought this behaviour could be beneficial for application of the optimisation strategies in real-world railway operation.

Acknowledgements

This research was funded by the Engineering and Physical Sciences Research Council, through the E-Futures DTC at the University of Sheffield.

References

[1] IEA, "2013 Key World Energy Statistics," 2013.

[2] Department of Energy & Climate Change, "Statistical press release: Digest of UK energy statistics 2013," 2013.

[3] Department of Energy & Climate Change, "ECUK (Energy consumption in the UK) - Transport data tables," 2013.

[4] M. McClanachan and C. Cole, "Current train control optimization methods with a view for application in heavy haul railways," Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit, vol. 226, no. 1, pp. 36-47, 2012.

[5] E. Khnelnitsky, "On an Optimal Control Problem of Train Operation," IEEE Transactions on Automatic Control, vol. 45, no. 7, pp. 1257-1266, 2000.

[6] R. Liu and I. M. Golovitcher, “Energy-efficient operation of rail vehicles,” Transportation Research Part A: Policy and Practice, vol. 37, no. 10, p. 917–932, 2003.

[7] Q. Gu, X. Lu and T. Tang, “Energy Saving for Automatic Train Control in Moving Block Signaling System,” 14th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 1305-1310, 2011.

[8] N. Zhao, C. Roberts and S. Hillmansen, “The application of an enhanced Brute Force algorithm to minimise energy costs and train delays for differing railway train control systems,” Proc MechE Part F: J Rail and Rapid Transit, vol. 228, no. 2, p. 158–168, 2014.

[9] J. H. Holland, Adaptation in Natural and Artificial Systems, University of Michigan Press, 1975.

[10] C. Chang and S. Sim, “Optimising train movements through coast control using genetic algorithms,” IEE Proceedings - Electric Power Applications, vol. 144, no. 1, pp. 65-73, 1997.

[11] T. Albrecht, “Reducing power peaks and energy consumption in rail transit systems by simultaneous train running time control,” Computers in Railways IX, pp. 885-894, 2004.

[12] M. Miyatake and H. Ko, “Numerical analyses of minimum energy operation of multiple trains under DC power feeding circuit,” European Conference on Power Electronics and Applications, pp. 1-10, 2007.

[13] M. Miyatake and H. Ko, “Optimization of Train Speed Profile for Minimum Energy Consumption,” IEEJ Transactions on Electrical and Electronic Engineering, vol. 5, no. 3, pp. 263-269, 2010.

[14] S. Acikbas and M. Soylemez, “Coasting point optimisation for mass rail transit lines using artificial neural networks and genetic algorithms,” IET Electric Power Applications, vol. 2, no. 3, pp. 172-182, 2008.

[15] M. T. Söylemez and S. Açıkbaş, “Multi-train simulation of DC rail traction power systems with regenerative braking,” Computers in Railways IX, pp. 941-950, 2004.

[16] L. Yang, K. Li, Z. Gao and X. Li, “Optimizing trains movement on a railway network,” Omega, vol. 40, no. 5, pp. 619-633, 2012.

[17] Y. Wang, B. DeSchutter, T. J. J. vandenBoom and B. Ning, “Optimal trajectory planning for trains under fixed and moving signaling systems using mixed integer linear programming,” ControlEngineeringPractice, vol. 22, pp. 44-56, 2014.

[18] L. Yang, Personal communication (22 June), 2013.

[19] R Core Team, “R: A Language and Environment for Statistical,” R Foundation for Statistical Computing (http://www.R-project.org), Vienna, Austria, 2013.

[20] K. Ichikawa, “Application of optimization theory for bounded state variable problems to the operation of a train,” Bulletin of JSME, vol. 11, no. 47, pp. 857-865, 1968.

[21] C. Sicre, P. Cucala, A. Fernández, J. Jiménez, I. Ribera and A. Serrano, “A method to optimise train energy consumption combining manual energy efficient driving and scheduling,” Computers in Railways XII, pp. 549-560, 2010.

[22] Y. Bocharnikov, A. Tobias, C. Roberts, S. Hillmansen and C. Goodman, “Optimal driving strategy for traction energy saving,” IET Electric Power Applications, vol. 1, no. 5, pp. 675 - 682, 2007.

[23] L. Yang, S. Li, Y. Gao and Z. Gao, “A Coordinated Routing Model with Optimized Velocity for Train Scheduling on a Single-Track Railway Line,” Int. J. Intell. Syst., vol. 30, no. 1, p. 3–22, 2015.