Traxim

Milwaukee Road
Example Use Case

Modelling the Milwaukee Road

Introduction

This use case provides a basic guide to Traxim inputs and outputs.

All of the input and output files (excluding the output kml file) are downloadable from the Resources page of the TraximRail website.

The alignment of the largely abandoned Western Region of the Chicago, Milwaukee and St Paul Railroad (commonly called the Milwaukee Road, and abbreviated to MR for this document) has been used as the core of the geography for this use case. This is purely a fantasy rail network, though it has some overlap with rail lines still in use, particularly along the western US seaboard. The following screenshot illustrates the example network in Google Earth. Note that all of the track in the fantasy network is assumed to be owned by MR (notwithstanding that some was historically only trackage rights).

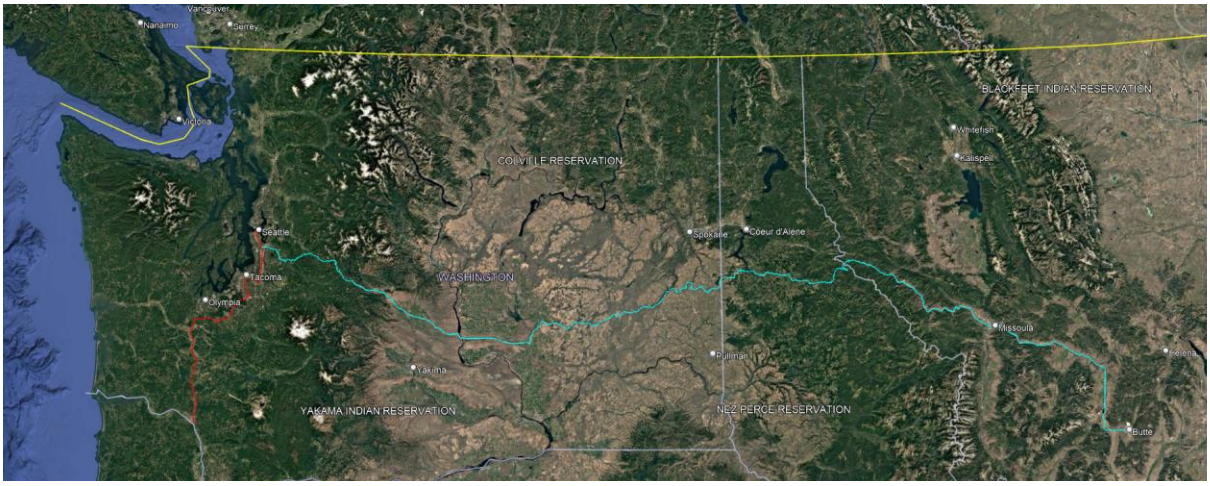

This use case is simply a single scenario to illustrate the use of Traxim to generate a population of timetables for a reasonably congested rail network. The fantasy scenario includes a mixture of train types, including both long distance and commuter passenger, and high speed and bulk freight.

The outputs could be used to quantify section utilisation, and to identify areas of congestion, as the starting point for train plan redesign, or infrastructure investment.

Overview of the Timetable Resolution Process

Traxim is offered as a service.

The Simulation Manager form provides the interface for defining and submitting a scenario. Validation and timetable resolution are undertaken using cloud based virtual machines. Results are sent by email.

In overview, the process is as follows:

1. Create, or update as appropriate, all of the input files to describe the scenario you wish to resolve. This User Manual has detailed instructions for each file type.

2. Once you have your scenario files, go to the Scenario Manager.

   You will first need to select between validation, single perturbation resolution, and batch resolution, and provide a contact email. In the first instance, submit your scenario for validation.

   You will then need to input four fields:
   - Year for simulation.
   - Quarter for simulation.
   - Simulation days.
   - Window days, to define the time period for generating the output data.

   You will then drag and drop, or upload, all of the required scenario input files.

   When you have finished setting up the scenario, choose submit. You will receive an email confirming the scenario, and providing a link to your submitted scenario. Once set up, the scenario can be edited and resubmitted an unlimited number of times.

3. The validation process will check for missing or logically inconsistent data. If there are problems, you will receive an email with a description of the problem. You will need to fix problems as they are identified and resubmit the validation.

4. Once you have a successful validation, you will receive an email with an unresolved timetable pdf. Please review that pdf to ensure that the scenario appears to be as you expected. Note that some scenario problems, such as a train with an inadequate power-to-weight ratio, will not be identified as errors in the validation process.

5. If you need to make changes, rerun the validation to again confirm that the result is as expected and you haven't introduced any errors.

6. When you are confident that your scenario is validating, and it is generating the expected results, submit the scenario for resolution of a single perturbation.

   Once resolution completes, you will receive an email with a pdf of the resolved timetable. Again, review the resolved timetable to confirm that it is what you were expecting. For instance, you may want to check that your train priority settings are giving reasonable crossing solutions. If appropriate, make changes to your scenario and regenerate the sample resolved timetable.

7. Once you are satisfied with your resolved timetable, submit the scenario for batch resolution. To do this, you will need to purchase scenario resolution credits. More detail on pricing will be released when the Traxim resolver becomes publicly available.

8. Once the batch resolution completes, you will be sent an email with your results.

Generating a population of timetables is a computing resource intensive process. For complex scenarios with a long simulation period, a single scenario with 30 perturbations may take a number of hours to resolve.

Input Files

The input files are as follows:

Infrastructure

The definition of the infrastructure is the most complex of the Traxim input files. The infrastructure input to the simulation is a CSV file but due to the need to be able to visualise the network to efficiently edit it, we provide a Network Editor tool (available shortly) that allows the network to be displayed graphically..

The Traxim infrastructure configuration is primarily defined in terms of turnout nodes. There is also a secondary node form to represent intermediate signals, and one for end node points.

A turnout has three branches, “facing”, “through” and “diverge”. An intermediate signal implicitly has two branches, both “through” in nature. An end point has only one connection, to the facing branch.

Branches are connected to each other by links. In network systems theory, these are also called edges.

Each turnout is located in space by its kilometre posting and allocation to a track geometry section.

Two examples of the infrastructure as it appears in the Network Editor are shown below:

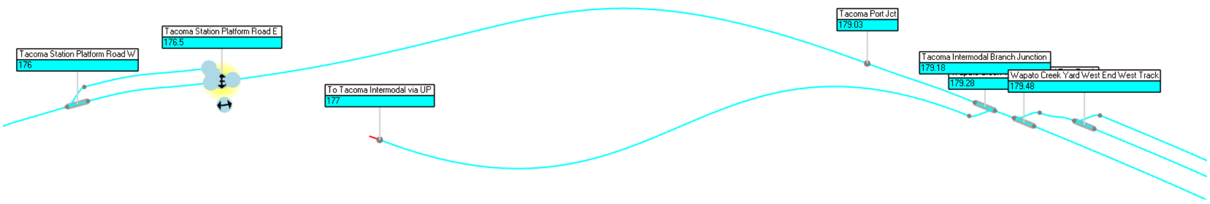

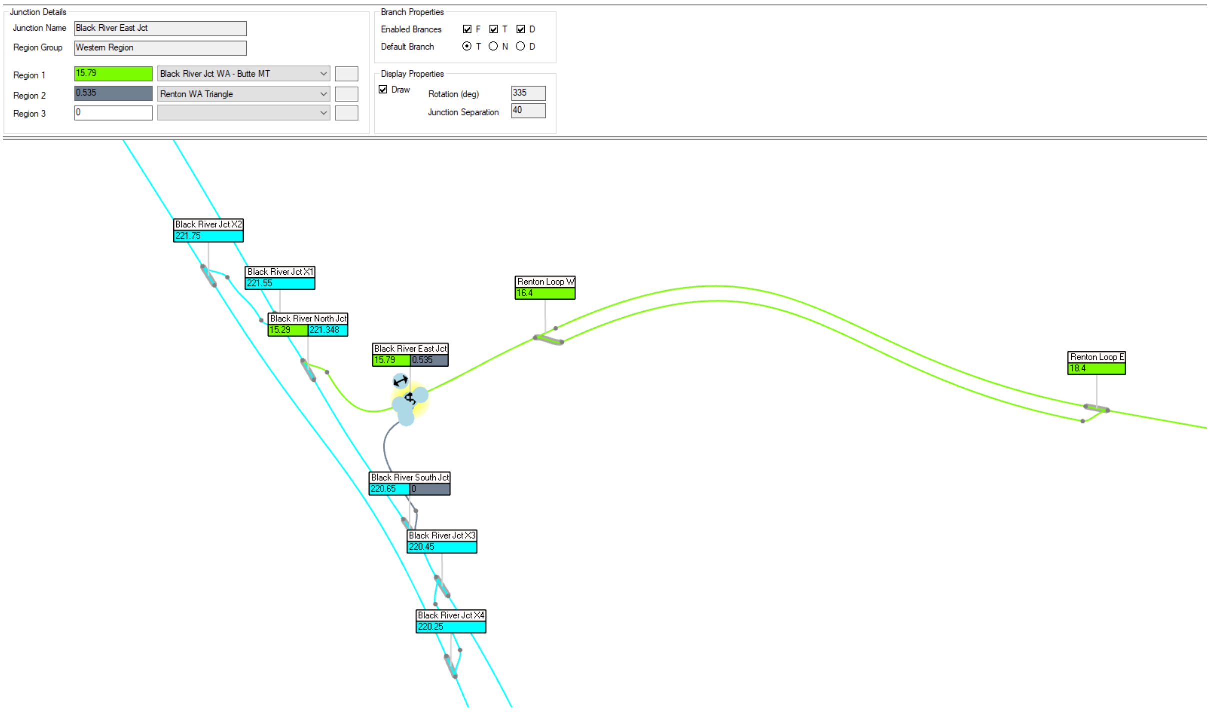

Regions

The primary purpose of the regions file is to integrate the regions naming used in the infrastructure, speedboards and the track geometry files. It also provides a filtering option for analysing results.

The Regions file:

- Designates a colour for each region, which is used in the infrastructure editor and the train graph.
- Specifies which order on the train graph the regions will appear.
- Specifies which direction of travel will be used on the train graph.
- Provides a mechanism to specify region specific transaction delays to be applied when there is a following or opposing conflict.

Speedboards

The speedboard file is a list of speedboards by kilometrage and track segment. Speedboards need to be defined by direction. An unlimited number of speedboard lists can be defined for each direction, allowing different classes of trains to operate to different maximum speed limits.

Train Plan

The purpose of the Train Plan input file is to define the characteristics of the trains to be simulated, including their entry time and entry and exit locations on the network.

The Train Plan file specifies a range of parameters including train name, length, weight, locomotive types and train priority.

Rules

It is frequently desirable to specify a series of events to occur along the journey of a train. The primary ones are that it must pass a specific location (eg a platform), that it must not pass a specific location (eg due to outline gauge limitations) or that it must dwell for a defined period / not depart until a specified time.

The Rules file works in conjunction with the Train Plan file to define these events for each train or group of trains. It is separate to the Train Plan file for ease of data management.

Geometry

Curves and gradients are defined by a track centreline file specifying latitude, longitude, elevation and nominal kilometrage.

The track geometry file is used primarily for dynamically simulating trains. The latitude and longitude, rather than the kilometre points, define the distance to be travelled by the train. Elevation data is used to generate the gradient. The nominal kilometrage cross-references the infrastructure file so that nodes can be positioned relative to the track geometry.

Tractive Effort

Each locomotive type used in the train plan input file needs a number of fundamental performance characteristics to be defined. This information is held in the tractive effort file for that locomotive type.

Output Files

Unresolved Train Graph

The validation process generates a pdf file of the train graph for the undeconflicted train plan. As already noted, you should review the unresolved train graph after the validation process,

The unresolved train plan pdf will look similar to this:

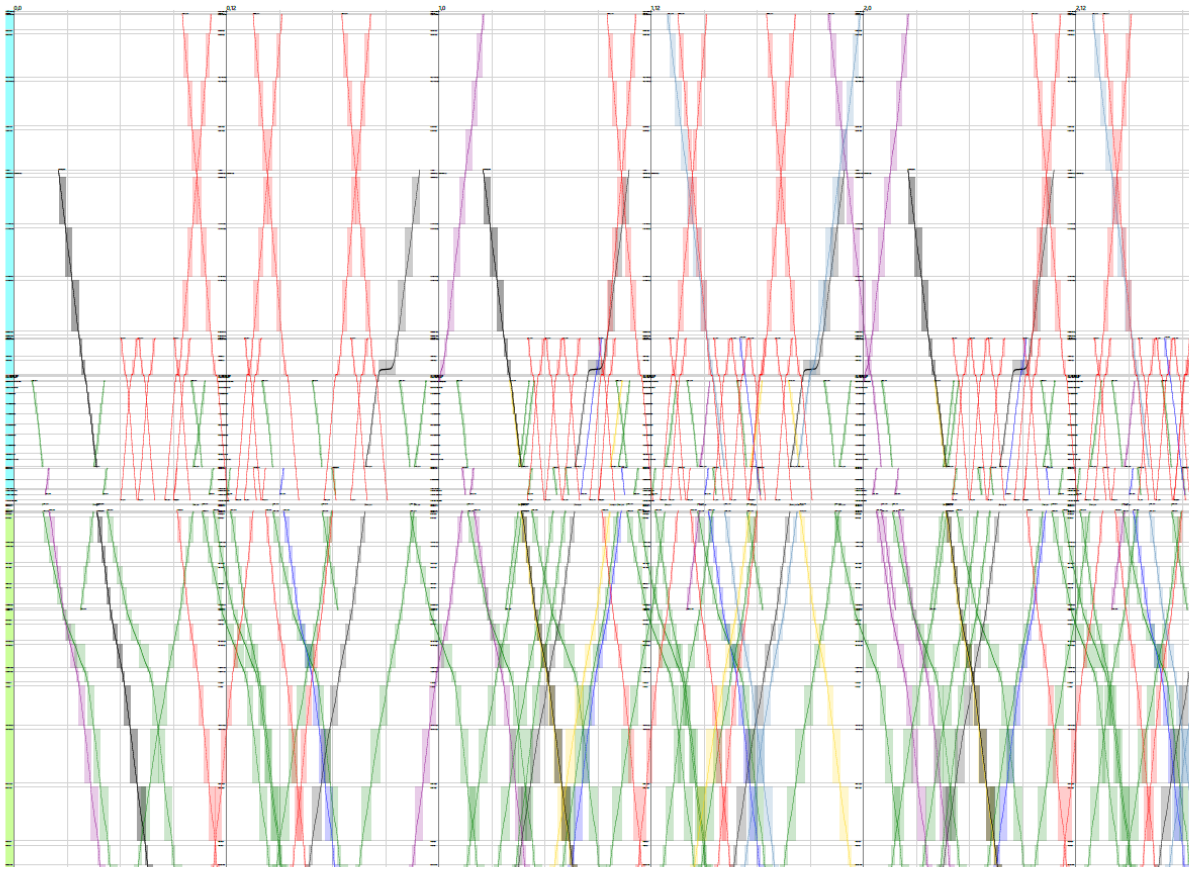

Key points about the train graph are as follows:

- A grey horizontal line represents a node. Note that whether a node is shown or not is user defined in the Infrastructure file. Hence not all nodes will necessarily appear on the train graph.
- Grey vertical lines are time, with the darker lines at 12-hour intervals, and the lighter lines at 3-hour intervals.
- The solid line represents the front of the train and the more lightly coloured box is the section occupation.
- Nodes are spaced by distance. Hence the train path will be more vertical when the train is simulated to be travelling faster, and less vertical when it is travelling slower.
- Note that Traxim sets a minimum train speed of 3 kph. Hence any train will eventually climb a grade even if it has insufficient power to do so in reality. The train path can become almost horizontal where the train is at the minimum speed. Use the train graph to identify trains that may require an adjustment to the consist.
- There is no train linking in the unresolved state. Hence all linked trains will appear as departing simultaneously at zero hours.
- If there are multiple routes for a train between an origin and destination, as opposed to multiple tracks that it can switch between on an individual route, check that the train is following the expected route. This will be apparent from the regions that it appears in on the train graph. In this example scenario there is only one possible route for each train.
- The height and width of the train graph is dictated by the number of simulation days and the length of the network. Depending on the interplay of these two factors, the train paths will be more or less vertical than shown in the sample above.

- Train names appear next to the train path where a train enters or leaves a section. These appear in very small type. You may need to zoom to around 1000% to be able to read them.
- The colour density of an occupation reflects the sum of the trains occupying the section. Be alert for darker occupations along the full length of a train path. This is likely to indicate two identical trains scheduled to depart at the same time.
- You may wish to apply additional intelligent design to the timetable at this point to reduce the bunching of trains. Note though that once crosses are resolved, most trains will shift progressively to the right and that what may appear as bunching in the unresolved timetable will be less intense in the resolved state. That said, having multiple following trains leaving in a cluster is likely to generate congestion, especially if they have similar transit times or if slower trains are scheduled to leave ahead of faster trains.

Resolved Train Graph

The resolved train graph pdf file returned after a successful single scenario submission will be fundamentally similar to the unresolved file, but with all conflicts now resolved. It will look as follows:

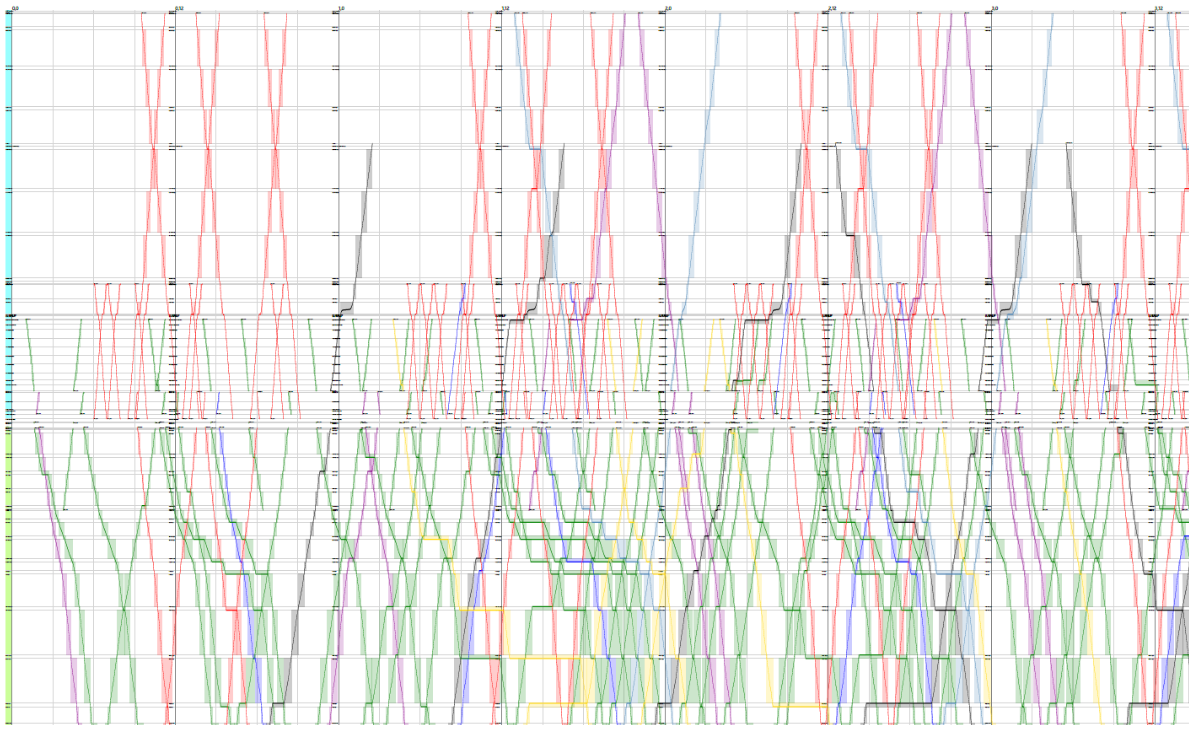

Note that this is provided for a single simulation only. Individual train graphs for each of the 30 perturbations in a batch resolution are not provided.

Train Statistics

The metrics and context information reported by the Train Statistics file are set out below. All of the performance data is as per the analysis window as specified in the Scenario inputs. Accordingly, the performance metrics will vary where otherwise identical trains have different proportions of their journey within the analysis window. Care should be taken to ensure that all metrics are normalized to account for this. Converting metrics to a "per 100 km" basis is usually the most effective method.

TrainName	The unique identifier for the train, as set in the Train Plan, appended with the departure day of the week, and the week number.
ColourGroup	The colour group for the train as set in the Train Plan.
EntryNode	The originating node, as set in the Train Plan.
ExitNode	The route end point node, as set in the Train Plan.
OriginalStartTime	The nominal originating time as set in the Train Plan, in minutes from midnight on day 1.
StarttimeTolerance	The window within which the start time can be randomized, in minutes, as set in the Train Plan.
Starttime	The actual time the train was scheduled to start, averaged across perturbation. This should trend toward OriginalStartTime plus 50% of StarttimeTolerance.
RawTransitMinutes	The dynamically simulated transit time, averaged across perturbations, and windowed. Note that as this is drawn from the resolved timetables, it includes acceleration and deceleration due to deconflicting dwells. Accordingly, trains that are identical as defined in the Train Plan will still have different raw transit times even if they operated fully within the window period.
ScheduledDwellMinutes	Dwell that has been specified for the train in the Rules input as "DwellAt".
FollowingDeconflictingDwellMinutes	Dwell incurred for a cross that was triggered by a following conflict. This includes overtaking.
OpposingDeconflictingDwellMinutes	Dwell incurred for a cross that that was triggered by an opposing conflict.
DeconflictingDwellMinutes	The sum of the following and opposing dwell minutes.
TotalTransitMinutes	The sum of scheduled dwell, deconflicting dwell, and raw transit time, also equal to the time between initial train departure and final train arrival.
TotalOpposingConflictsLost	The number of opposing conflicts where the train was required to deviate from its unconflicted path. A deviation includes both deviations in time (dwells), and routes (diverging to a different track), or both.
TotalOpposingConflictsWon	The number of opposing conflicts where the train forced another train to deviate from its unconflicted path.
TotalFollowingConflictsLost	The number of following conflicts where the train was required to deviate from its unconflicted path.
TotalFollowingConflictsWon	The number of following conflicts where the train

	forced another train to deviate from its unconflicted path.

Track Statistics

The Track Statistics file reports three metrics, each in one hour buckets, as follows:

- Seconds of link occupation, including transaction time.
- Number of trains entering the link.
- Three-hour rolling percentage utilization, including transaction time.

Each of the metrics is the average of the 30 perturbations.

These metrics and the supporting information are defined as follows:

Link Name	The name of the link with the following syntax: Node A name : Node A branch Node B branch : Node B name.
Link type	The connected branch of the node at each end of the link. On a single track line, the Link Type will typically be FF for the single track mainline, TT for the mainline where it is within the bounds of a loop, and DD for a loop track. Intermediate signals and double track can generate any combination of the T, F and D branches. The purpose of providing the Link Type is to simplify filtering and the use of Excel lookup functions such as sumifs, maxifs, and countifs.
Link distance	The distance, in metres, for the link, calculated from the latitude, longitude and altitude specified in the geometry files.
Region	The region as specified by the infrastructure file.
Transaction time - following	The generic transaction time for following dwells, as specified in the regions file.
Transaction time - opposing	The generic transaction time for opposing dwells, as specified in the regions file.
Minutes of link occupation	The sum of seconds the link was occupied by a train, plus the sum of transaction time. The sum of transaction time is calculated as the number of trains that entered the section during the hour, times the average of the following and opposing transaction time (in seconds). This assumes that on average 50% of trains will enter the section travelling in the same direction as the preceding train and 50% will enter from the opposing direction.
Number of trains entering the link	The sum of trains that entered the link during the relevant hour. This can be used to calculate average time occupation per train, and to differentiate between track occupation and transaction time.
Three-hour rolling percentage utilization	The sum of seconds of link occupation for the current and two preceding hours, divided by 3600.

Section utilization is the key metric for capacity analysis. However, there is no objective or correct limit for utilization, or a correct time period over which it should be measured.

The three-hour rolling utilization is provided as a quantification of clustering of trains.

If your traffic has high levels of flexibility over train departure times, such as on a largely bulk network, it is likely to be more appropriate to use a 24-hour utilization rate. If your traffics have relatively inflexible departure windows, such as for some intermodal traffics, a shorter window for utilization may be useful. Either way, peaks and troughs in train numbers may create artificial congestion. This might manifest as high levels of delay, but not as high utilization when calculated over a 24-hour period. In this case, relaxing the constraints around departure windows may be as effective as infrastructure enhancements.

Acceptable levels of section utilization are also influenced by the utilization of nearby sections. If the surrounding sections have relatively low utilization rates, a single section with a high utilization rate, like 80% over 24 hours, may be tolerable, whereas a series of sections all with utilization rates of around 60% may create unsustainable congestion.

The train statistics file is returned as raw data. For better visualisation, you can apply conditional formatting as shown on the example below. This example has also had a filter applied to column B so that only single track sections between loops are shown.

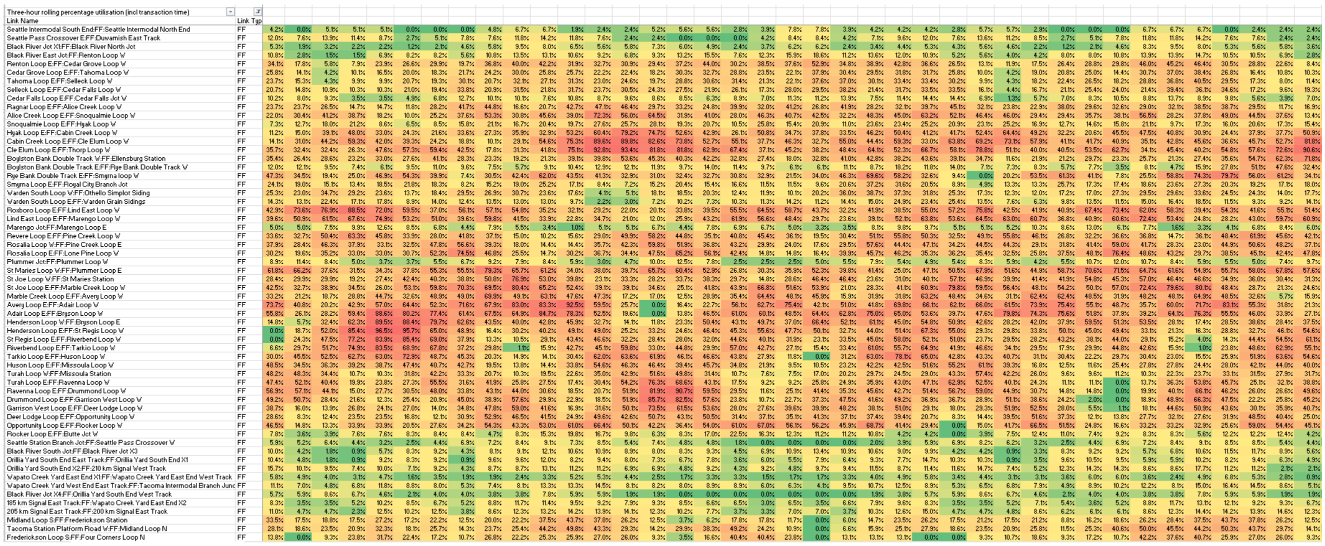

Delay by Location in Google Earth

Traxim generates a Google Earth kml file that shows delay by location. This view allows you to instantly visualise where on your network you are incurring the greatest delay, which is likely to be where you want to focus any capacity enhancement efforts. Google Earth allows you to open multiple kml files, so you can compare delay by location for different infrastructure.

The following image shows delay for the Milwaukee Road Example Scenario. Unfortunately, limitations preclude a kml file being available on the Resources website page. Please contact us if you would like the file.